



Parallels Remote Application Server

Administrator's Guide

21.2

Parallels International GmbH
Vordergasse 59
8200 Schaffhausen
Switzerland

Tel: + 41 52 672 20 30
www.parallels.com

© 2026 Parallels International GmbH. All rights reserved. Parallels and the Parallels logo are trademarks or registered trademarks of Parallels International GmbH in Canada, the U.S., and/or elsewhere.

Apple, Safari, iPad, iPhone, Mac, macOS, iPadOS are trademarks of Apple Inc. Google, Chrome, Chrome OS, and Chromebook are trademarks of Google LLC.

All other company, product and service names, logos, brands and any registered or unregistered trademarks mentioned are used for identification purposes only and remain the exclusive property of their respective owners. Use of any brands, names, logos or any other information, imagery or materials pertaining to a third party does not imply endorsement. We disclaim any proprietary interest in such third-party information, imagery, materials, marks and names of others. For all notices and information about patents please visit <https://www.parallels.com/about/legal/>

README

- [Introduction](#)
- [Parallels RAS 21 release history](#)
- [About Parallels RAS](#)
 - [How does it work?](#)
- [About this guide](#)
- [What's new](#)
 - [Parallels RAS 21](#)
- [Terms and abbreviations used in this guide](#)
- [Installing Parallels RAS](#)
- [System requirements](#)
- [Hardware requirements](#)
- [Software requirements](#)
 - [RAS Connection Broker and RAS Secure Gateway \(64-bit versions only\)](#)
 - [RAS Web Administration Service](#)
 - [RAS RD Session Host Agent](#)
 - [RAS Provider Agent](#)
 - [RAS Guest Agent](#)
 - [Remote PC Agent](#)
 - [Parallels RAS PowerShell](#)
 - [Parallels RAS Console](#)
 - [RAS Enrollment Server](#)
 - [RAS Performance Monitor](#)
 - [RAS Reporting](#)
 - [Parallels Client](#)
- [Microsoft license requirements](#)
- [Install Parallels RAS](#)
- [Log in and activate Parallels RAS](#)
 - [Start the Parallels RAS Console](#)
 - [Sign in to Parallels My Account](#)
 - [Activate Parallels RAS](#)
- [Getting Started with Parallels RAS](#)
- [The Parallels RAS Console](#)
- [Set up a basic Parallels RAS Farm](#)
- [Add an RD Session Host](#)

- [Publish applications](#)
- [Invite users](#)
- [Azure Virtual Desktop](#)
- [Conclusion](#)
- [Farm and Sites](#)
- [Connecting to a Parallels RAS Farm](#)
 - [Connecting to a Parallels RAS Farm for the first time](#)
 - [Connecting to a different Parallels RAS Farm](#)
 - [Switching between Parallels RAS Farms](#)
 - [Editing Parallels RAS Farm connections](#)
 - [Connecting to a Parallels RAS Farm via Command Prompt](#)
 - [Arguments](#)
 - [Examples](#)
- [About Sites](#)
- [Sites in the RAS Console](#)
 - [Current Site](#)
 - [Switching between sites](#)
 - [Renaming a Site](#)
 - [Site configuration and health view](#)
 - [Performing tasks on a component](#)
 - [Using the Site Designer](#)
- [Adding a Site to the Farm](#)
- [Replicating Site settings](#)
 - [Overriding Site Replicated Settings](#)
- [Managing Licensing Site](#)
- [Managing administrator accounts](#)
- [Adding an administrator account](#)
 - [Modifying an administrator account](#)
- [Administrator account permissions](#)
 - [Power administrator permissions](#)
 - [Custom administrator permissions](#)
 - [Clone permissions](#)
 - [Delegate permissions](#)
- [Managing administrator accounts](#)
 - [Modifying an account](#)
 - [Moving an account to a different Active Directory group](#)
 - [Handling locked objects](#)
- [Configure RAS Console idle sessions](#)
- [Using instant messaging](#)
- [Joining Customer Experience Program](#)

- [RAS Connection Broker](#)
- [Configuring RAS Connection Brokers](#)
 - [RAS Connection Brokers overview](#)
 - [Automatically connect to an alternative Connection Broker](#)
- [Secondary Connection Brokers](#)
 - [Planning for secondary Connection Brokers](#)
 - [Adding a secondary RAS Connection Broker to a Site](#)
- [Managing Secondary Connection Brokers](#)
 - [Enabling or disabling a secondary Connection Broker](#)
 - [Changing the secondary Connection Broker priority](#)
 - [Promoting a secondary Connection Broker to primary](#)
 - [Configuring auto-promotion](#)
 - [Deleting a secondary Connection Broker](#)
- [Using computer management tools](#)
- [Enabling Connection Broker API](#)
 - [SSL/TLS](#)
 - [API Token](#)
 - [Encryption Certificate](#)
- [RAS Secure Gateway](#)
- [Overview](#)
 - [Installing a RAS Secure Gateway on a dedicated server](#)
 - [How a RAS Secure Gateway works](#)
 - [RAS Secure Gateway operation modes](#)
 - [Planning for high availability](#)
 - [Performance improvements in version 21.0](#)
 - [Baseline specifications](#)
- [Adding a RAS Secure Gateway](#)
- [Manually adding a RAS Secure Gateway](#)
- [Checking the RAS Secure Gateway status](#)
- [Configuring a RAS Secure Gateway](#)
- [Enable or disable a Secure Gateway](#)
- [Set public address](#)
- [Set IP addresses for client connections](#)
- [Site defaults \(Secure Gateways\)](#)
- [Gateway mode and forwarding settings](#)
 - [Using Site defaults](#)
 - [Setting the normal mode](#)
 - [Setting the forwarding mode](#)
- [Gateway network options](#)
 - [Using Site defaults](#)

- [Configuring network](#)
- [SSL/TLS encryption](#)
 - [Using Site defaults](#)
 - [Enforcing HSTS](#)
 - [Configuring SSL](#)
 - [Encrypting Parallels Client connection](#)
 - [Trusted Certificate Authorities](#)
 - [Securing RDP-UDP connections](#)
- [SSL server configuration](#)
- [Configure User Portal](#)
- [Using Site defaults](#)
- [Enable or disable User Portal](#)
- [Client settings](#)
- [Network load balancers access](#)
- [Wyse ThinOS support](#)
- [Secure Gateway security](#)
 - [Using Site defaults](#)
 - [Configuring security](#)
- [Web request load balancing](#)
 - [Redirection URL](#)
 - [Default URL format](#)
 - [Opening a specific User PortalTheme](#)
 - [Web cookie](#)
- [Secure Gateway tunneling policies](#)
 - [Adding a new Tunneling Policy](#)
 - [Managing a Tunneling Policy](#)
 - [Restricting RDP access](#)
- [Configure logging](#)
- [Viewing Secure Gateway summary and metrics](#)
- [Using computer management tools](#)
- [RD Session Hosts](#)
- [RD Session Host types](#)
- [Add an RD Session Host](#)
 - [RD Session Host requirements](#)
 - [Add an RD Session Host](#)
- [Installing the agent manually](#)
 - [Installing RAS RD Session Host Agent manually](#)
 - [Uninstalling RAS RD Session Host Agent](#)
- [Add a template-based RD Session Host](#)
- [Manage RD Session Hosts](#)

- [Manage host pools \(RD Session Hosts\)](#)
 - [Moving an RD Session Host to another host pool](#)
 - [Autoscale](#)
 - [Using host pool defaults](#)
- [Add host pools \(RD Session Hosts\)](#)
 - [Creating a host pool](#)
- [Upgrading Agents \(RD Session Hosts\)](#)
 - [Schedule Agent auto-upgrade](#)
 - [Cancel Agent auto-update](#)
- [Manage templates \(RD Session Hosts\)](#)
- [Creating an RD Session Host template](#)
 - [Requirements](#)
 - [Manual agent installation](#)
 - [Create a template](#)
- [Assigning a template to a host pool \(RD Session Host\)](#)
- [Managing RD Session Hosts based on a template](#)
 - [Viewing RD Session Hosts based on a specific template](#)
 - [Site defaults](#)
 - [Checking the RAS Guest Agent status](#)
 - [Deleting an RD Session Hosts based on a template](#)
 - [Managing RD Session Hosts based on a template that failed preparation](#)
 - [Recreating an RD Session Hosts based on a template](#)
- [Manage hosts \(RD Session Hosts\)](#)
- [Viewing RD Session Hosts](#)
 - [Viewing RD Session Host summary](#)
 - [Available menu options](#)
- [Check an RD Session Host Agent status](#)
- [Change RD Session Host Site assignment](#)
- [View and modify RD Session Host properties](#)
- [Using default settings](#)
- [General](#)
- [Agent settings](#)
 - [Application session lingering](#)
 - [Other settings](#)
 - [Allow Client URL/Mail redirection](#)
 - [Enable Drag and drop](#)
 - [Allow 2XRemoteExec to send command to the client](#)
 - [Use RemoteApp if available](#)
 - [Manage RDP transport protocol](#)

- [Enable application monitoring](#)
- [Allow file transfer command \(Web and Chrome clients\)](#)
- [Enable drive redirection cache](#)
- [Enable Z-order \(Experimental\)](#)
- [Using 2XRemoteExec](#)
- [2XRemoteExec examples:](#)
- [User profile](#)
- [Application Packages](#)
- [Optimization](#)
- [Desktop access](#)
- [RDP printer](#)
- [User profile](#)
- [User Profile Disks](#)
- [FSLogix](#)
 - [Supported FSLogix releases](#)
 - [Configure installation method](#)
 - [Upgrade FSLogix](#)
 - [Configure Site defaults and hosts for FSLogix](#)
 - [Rebooting a host](#)
- [Configure managing existing profiles by Parallels RAS](#)
 - [Before you begin](#)
 - [Preliminary steps](#)
 - [Replicate GPO and FSLogix configuration](#)
 - [Recommendations and testing](#)
- [FSLogix antivirus exclusions](#)
- [FSLogix support requirements](#)
 - [Deprecated Features and Retired Platforms](#)
- [Optimization](#)
 - [Configure optimization](#)
 - [Using custom script](#)
 - [Applying optimization](#)
 - [Upgrade](#)
 - [Inheritance](#)
 - [Additional information](#)
- [Drive redirection cache](#)
 - [Notes](#)
 - [Limitations](#)
- [Configure logging](#)
- [Manage sessions \(RD Session Host\)](#)

- [Using scheduler \(RD Session Hosts\)](#)
 - [Disabling hosts and hosts in pools](#)
 - [Rebooting hosts and hosts in pools](#)
 - [Starting up hosts and hosts in pools](#)
 - [Shutting down hosts and hosts in pools](#)
 - [Recreating hosts and host pools](#)
 - [Sending multiple messages to users](#)
- [Planning for high availability](#)
- [Managing logons](#)
- [Using computer management tools](#)
- [Publishing from an RD Session Host](#)
- [Viewing published resources](#)
- [Virtual Desktop Infrastructure \(VDI\)](#)
- [Supported providers](#)
 - [Hypervisors](#)
 - [Cloud Providers](#)
- [Add a provider](#)
- [RAS Provider Agent information](#)
- [RAS Provider Agent installation options](#)
- [Add a hypervisor provider](#)
- [Add a cloud Provider](#)
- [Microsoft Azure](#)
- [Introduction and prerequisites](#)
 - [Introduction](#)
 - [Prerequisites](#)
- [Create Microsoft Entra ID application](#)
 - [Create an Microsoft Entra ID application](#)
 - [Create a client secret](#)
 - [Give the application read and write access to resources](#)
 - [Finding your Microsoft Azure subscription ID](#)
 - [Summary](#)
- [Add Microsoft Azure as a Provider](#)
 - [Modifying the Provider configuration](#)
- [Microsoft Azure and templates](#)
- [Amazon Web Services](#)
- [Introduction and prerequisites](#)
 - [Introduction](#)
 - [Prerequisites](#)
- [Design considerations](#)
 - [DHCP options set](#)

- [Service quotas](#)
- [Encrypting storage](#)
- [Step 1. Creating an IAM user for programmatic access](#)
- [Step 2. Adding AWS as a Provider](#)
- [Add a Custom Provider](#)
- [Introduction and Support](#)
 - [Introduction](#)
 - [Support](#)
- [Creating and Modifying a Custom Provider Configuration](#)
 - [Modifying the Provider Configuration](#)
- [Custom Provider Scripts](#)
- [Manage VDI](#)
- [Manage providers \(VDI\)](#)
 - [Configure a Provider](#)
 - [Enable or disable a Provider in Site](#)
 - [Properties: configure Provider connection settings](#)
 - [Hypervisor Provider:](#)
 - [Cloud-based Provider:](#)
 - [Credentials: configure username and password](#)
 - [Hypervisor Provider:](#)
 - [Cloud-based Provider:](#)
 - [Advanced](#)
 - [MAC addresses](#)
- [Installing RAS Provider Agent using the installer](#)
- [Checking the RAS Provider Agent status](#)
- [Using a Provider in multiple farms](#)
- [Manage host pools \(VDI\)](#)
- [Add host pools \(VDI\)](#)
- [Delete host pools \(VDI\)](#)
- [Add and delete host pool members](#)
 - [Adding a member to a host pool](#)
 - [Removing a member from a host pool](#)
- [Using a wildcard to filter VMs](#)
- [Managing hosts in pools](#)
- [Upgrading Agents \(VDI\)](#)
 - [Schedule Agent auto-upgrade](#)
 - [Cancel Agent auto-update](#)
- [Manage templates \(VDI\)](#)
- [Virtual desktop templates](#)

- [Multi-provider template distribution](#)
 - [Prerequisites](#)
- [Creating a VM template](#)
 - [Requirements](#)
 - [Manual agent installation](#)
 - [Create a template](#)
- [Step 1: Check and install the Agent](#)
 - [Virtual desktop:](#)
 - [RD Session Host:](#)
- [Step 2: Configure the template](#)
- [Properties](#)
- [Distribution](#)
 - [Microsoft Hyper-V](#)
 - [VMware vCenter](#)
- [Advanced](#)
 - [Hypervisor-based Providers](#)
 - [Nutanix Prism Central Provider](#)
 - [Microsoft Azure Provider](#)
- [Preparation](#)
- [Optimization](#)
- [License keys](#)
- [Summary](#)
- [Host naming](#)
 - [Numeric pattern](#)
 - [Examples:](#)
 - [Reusing VM numbers in a name](#)
- [Parallels Test Template Wizard](#)
- [Modifying template properties](#)
- [How hosts are created from a template](#)
 - [Auto-deletion of hosts](#)
- [Manually adding a host](#)
- [Assigning a template to a host pool \(VDI\)](#)
- [Template maintenance](#)
 - [Full clones](#)
 - [Linked clones](#)
 - [Updating RAS Guest Agent inside a template](#)
 - [Maintaining RD Session Hosts based on a template](#)
- [Template status](#)
 - [Template creation](#)
 - [Template in production](#)

- [Template in maintenance](#)
- [Template removal](#)
- [Managing multi-provider template distribution](#)
 - [Adding or removing Providers from the distribution list](#)
 - [Template distribution status](#)
 - [Distribution details](#)
 - [Leaving maintenance mode](#)
- [Managing template-based hosts](#)
- [Manage hosts \(VDI\)](#)
 - [Viewing host list](#)
 - [Site defaults](#)
 - [Checking the RAS Guest Agent status](#)
 - [Deleting a host](#)
 - [Managing hosts that failed preparation](#)
 - [Recreating a host](#)
 - [Assigning a host](#)
- [Persistent hosts](#)
- [Configuring hosts to interact with RAS Provider Agent in a different subnet](#)
- [Manage sessions \(VDI\)](#)
- [Using scheduler \(VDI\)](#)
 - [Disabling hosts and hosts in pools](#)
 - [Rebooting hosts and hosts in pools](#)
 - [Starting up hosts and hosts in pools](#)
 - [Shutting down hosts and hosts in pools](#)
 - [Recreating hosts and host pools](#)
 - [Sending multiple messages to users](#)
- [Configure logging](#)
- [Enabling high availability for VDI](#)
 - [At least three Connection Brokers are required](#)
 - [An odd number of agents is recommended](#)
 - [Configuring a Provider for high availability](#)
- [Site defaults \(VDI\)](#)
 - [General](#)
 - [User profile](#)
 - [Application Packages](#)
 - [Optimization](#)
 - [Actions](#)
 - [Security](#)
 - [Settings](#)
 - [RDP printer](#)

- [Using computer management tools](#)
- [Viewing Provider summary](#)
- [Remote PC pools in VDI](#)
 - **[Remote PC pools vs. Remote PCs](#)**
- [Adding a Provider](#)
- [Adding Remote PCs to a Provider](#)
- [Adding Remote PCs to a pool](#)
- [Managing Remote PCs in a pool](#)
 - **[Performing power operations](#)**
- [Persistent Remote PCs](#)
- [RAS Guest Agent installation options](#)
- [Azure Virtual Desktop](#)
- [Introduction](#)
 - **[Extended values and capabilities](#)**
- [Prerequisites](#)
 - **[Microsoft Azure subscription](#)**
 - **[Azure Virtual Desktop user license entitlement](#)**
 - **[Permissions and Azure resource providers](#)**
 - **[Microsoft Entra ID application](#)**
 - **[Active Directory](#)**
 - **[Other](#)**
 - **[Additional notes](#)**
- [Enable Azure Virtual Desktop and add a provider](#)
- [Add host pools \(Azure Virtual Desktop\)](#)
- [Manage Azure Virtual Desktop](#)
- [Manage providers \(Azure Virtual Desktop\)](#)
 - **[Choosing the Active Directory Domain Services type](#)**
 - **[Disk storage cost optimization](#)**
- [Manage workspaces \(Azure Virtual Desktop\)](#)
- [Manage host pools \(Azure Virtual Desktop\)](#)
 - [Host pool properties](#)
- [Upgrading Agents \(Azure Virtual Desktop\)](#)
 - [Schedule Agent auto-upgrade](#)
 - [Cancel Agent auto-update](#)
- [Manage templates \(Azure Virtual Desktop\)](#)
- [Create a template](#)
- [Manage existing templates](#)
- [Assigning a template to a host pool \(Azure Virtual Desktop\)](#)
- [Manage hosts \(Azure Virtual Desktop\)](#)
 - **[Add a host](#)**

- [Manage sessions \(Azure Virtual Desktop\)](#)
- [Using scheduler \(Azure Virtual Desktop\)](#)
 - [Disabling hosts and hosts in pools](#)
 - [Rebooting hosts and hosts in pools](#)
 - [Starting up hosts and hosts in pools](#)
 - [Shutting down hosts and hosts in pools](#)
 - [Recreating hosts and host pools](#)
 - [Sending multiple messages to users](#)
- [Site defaults \(Azure Virtual Desktop\)](#)
- [Site defaults for single-session hosts](#)
 - [Host pool settings](#)
 - [User profile](#)
 - [Application packages](#)
 - [Optimization](#)
 - [RDP printer](#)
- [Site defaults for multi-session hosts](#)
 - [Host pool settings](#)
 - [User profile](#)
 - [Application packages](#)
 - [Optimization](#)
 - [RDP printer](#)
- [Using Parallels Client with Azure Virtual Desktop](#)
 - [Requirements](#)
 - [Accessing Azure Virtual Desktop resources in Parallels Client for Windows](#)
- [Verify the deployment](#)
- [Remote PCs](#)
- [Overview](#)
- [Manage host pools \(Remote PC\)](#)
- [Manage hosts \(Remote PC\)](#)
- [Adding a Remote PC to a Farm](#)
- [Admin-initiated Remote PC enrollment](#)
 - [Requirements](#)
 - [Adding a Remote PC to a Farm](#)
 - [Installing Remote PC Agent manually](#)
 - [Uninstalling Remote PC Agent](#)
- [Self-service Remote PC enrollment](#)
 - [Requirements](#)
 - [Configure self-service enrollment](#)
 - [Enrolling a PC](#)

- [Configuring a Remote PC](#)
 - [Properties](#)
 - [Agent Settings](#)
 - [RDP Printer](#)
 - [Configure logging](#)
- [Viewing Remote PC summary](#)
- [Using computer management tools](#)
- [Publishing](#)
- [Overview](#)
- [Publishing a desktop](#)
- [Publishing an application](#)
 - [Publishing an App-V application](#)
- [Publishing local applications](#)
- [Publishing an application with MSIX app attach](#)
- [Publishing a web application](#)
- [Publishing a network folder](#)
- [Publishing a document](#)
- [General management tasks](#)
- [Manage published applications](#)
 - [Configuring a published application](#)
 - [Information — review application properties](#)
 - [Publish from — configure from which servers the application is published](#)
 - [Application — configure application and hosting server settings](#)
 - [Filtering](#)
 - [Routing](#)
 - [Shortcuts — configuring shortcut options for a published application](#)
 - [File extension — configuring file extension associations](#)
 - [License — configuring licensing options for published applications](#)
 - [Display — configuring display settings for a published application](#)
 - [Quick keypad](#)
- [Manage published local applications](#)
 - [Configuring a published application](#)
 - [Information — review application properties](#)
 - [Publish from — configure from which servers the application is published](#)
 - [Application — configure application and hosting server settings](#)
 - [Filtering](#)

- [Routing](#)
- [Shortcuts — configure shortcut options for a published application](#)
- [File extension — configuring file extension associations](#)
- [License — configuring licensing options for published applications](#)
- [Display — configuring display settings for a published application](#)
- [Quick keypad](#)
- [Manage published desktops](#)
 - [Configuring a published desktop](#)
 - [Sites — configuring from which sites a published desktop is available](#)
 - [Publish from — configuring from which RD Session Hosts a desktop is published](#)
 - [Desktop — configuring desktop name, size and other properties](#)
 - [Filtering](#)
 - [Routing](#)
 - [Shortcuts — configuring shortcut options for a published desktop](#)
- [Manage published documents](#)
 - [Configuring a published document](#)
 - [Sites — configuring from which sites a published document is available](#)
 - [Publish from — configuring from which servers a document is published](#)
 - [Application — configuring server-specific document settings](#)
 - [Filtering](#)
 - [Routing](#)
 - [Shortcuts — configuring shortcut options for a published document](#)
 - [File extension — configuring file extension associations](#)
 - [Licensing — configuring licensing options for published documents](#)
 - [Display — configuring display settings for a published document](#)
- [Manage folders](#)
 - [Creating a folder](#)
 - [Managing folders](#)
 - [Adding published resources to a folder](#)
 - [Delegating permissions to custom administrators](#)
- [Site defaults \(Publishing\)](#)
 - [Shortcuts](#)

- [License](#)
- [Display](#)
- [Maintenance](#)
- [Using filtering rules](#)
- [Configuring preferred routing](#)
 - [Overview](#)
 - [Configure preferred routing](#)
 - [Specifying public address when inviting users](#)
 - [Deleting or disabling Gateway or HALBs](#)
- [Understanding session prelaunch](#)
- [Checking effective access](#)
- [Specifying client settings](#)
 - [Select icon resolution](#)
 - [Enable, disable or change the overlay icon](#)
 - [Show password expiration reminder](#)
 - [Session reset](#)
- [Quick keypad](#)
 - [Exporting and importing a Quick Keypad template](#)
- [Session Management](#)
- [Overview](#)
 - [Sessions category](#)
 - [Sessions tabs](#)
- [Session information](#)
 - [Session Setup](#)
 - [Logon details](#)
 - [User Experience](#)
 - [Session details](#)
 - [Client details](#)
 - [Export session information](#)
- [Monitoring settings](#)
- [Managing sessions](#)
 - [User session remote control](#)
 - [Managing processes](#)
- [The Resources tab](#)
- [SSL Certificate Management](#)
- [Generating a self-signed certificate](#)
- [Generating a certificate signing request \(CSR\)](#)
- [Let's Encrypt certificates](#)
- [Requesting a Let's Encrypt Certificate](#)
 - [Issuing a Let's Encrypt certificate](#)

- [Renewing a Let's Encrypt certificate manually](#)
 - [Revoking a Let's Encrypt certificate](#)
- [How Parallels RAS requests certificates from Let's Encrypt](#)
- [Importing a certificate](#)
- [Exporting a certificate](#)
- [Assigning a certificate](#)
 - [Certificate usage](#)
 - [Gateways](#)
 - [HALB](#)
 - [Connection Broker API Server](#)
 - [Connection Broker API data Encryption](#)
- [Auditing certificates](#)
- [Permissions to manage certificates](#)
- [Upgrading from an older RAS version](#)
- [Connection and Authentication Settings](#)
- [RAS Connection Broker connection settings](#)
 - [Choosing authentication type](#)
 - [Smart card certificate](#)
 - [Authentication domain](#)
 - [Authenticating against non domain users](#)
 - [Changing domain password](#)
- [Remote session settings](#)
 - [Declare user session as idle after](#)
 - [FIPS 140-2 encryption](#)
 - [Automatically log out client idle connection after](#)
 - [Log out idle session mode](#)
 - [Cached authentication token timeout](#)
 - [Clear cached authentication tokens \(a button\)](#)
- [Logon hours settings](#)
 - [Prerequisites](#)
 - [Adding a logon hours rule](#)
 - [To specify these settings:](#)
- [Restricting access by Parallels Client type and build number](#)
- [Multi-factor authentication](#)
- [Adding an MFA provider](#)
- [Using RADIUS](#)
- [Connection](#)
- [Automation](#)
 - [Autosend](#)
 - [Parallels Client](#)

- [Automation Conditions](#)
- [Advanced](#)
- [Configuring Azure MFA](#)
 - [Configure Azure MFA](#)
- [Configuring Duo](#)
- [Using TOTP](#)
- [Configuring TOTP](#)
 - [Using Google Authenticator in Parallels Client](#)
- [Configuring Microsoft Authenticator](#)
- [Configuring email OTP](#)
- [Using Deepnet DualShield](#)
- [Supported tokens](#)
- [Configuring DualShield 5.6+ Authentication Platform](#)
- [Configuring Parallels RAS to use the DualShield Authentication Platform](#)
- [Connect to a RAS Farm](#)
 - [Parallels Client](#)
- [Using SafeNet](#)
- [Configuring SafeNet](#)
 - [Parallels Client](#)
- [Configuring MFA rules](#)
- [Allowing users to change domain password](#)
 - [Passing the domain name to Parallels Client](#)
 - [Using a custom link to change domain passwords](#)
- [Allowing users to discover RAS connections via email address](#)
- [Load Balancing and HALB](#)
- [Resource based & round robin load balancing](#)
 - [Selecting load balancing method](#)
 - [Configuring resource counters](#)
 - [Session options](#)
 - [Agent timeout and refresh time](#)
- [Configure CPU optimization](#)
 - [Start](#)
 - [CPU conditions](#)
 - [Exclusions](#)
- [High availability load balancing \(HALB\)](#)
- [Prerequisites](#)
 - [Sizing recommendations](#)
- [Deploying a Parallels HALB appliance](#)
 - [VMware](#)
 - [Microsoft Hyper-V](#)

- [Nutanix AHV](#)
- [Scale HyperCore](#)
- [Deploying a Parallels HALB Appliance](#)
- [HALB Device First Boot Configuration](#)
- [LB Gateway payload](#)
- [LB SSL payload](#)
- [Device Manager](#)
- [Devices](#)
- [Modifying Virtual Server and configuring advanced options](#)
- [HALB Device status and version number](#)
 - [Site tab](#)
 - [Devices tab](#)
- [HALB maintenance](#)
- [HALB connection and session information](#)
- [RAS Multi-Tenant Architecture](#)
- [Overview](#)
- [Architecture description](#)
- [Implementation overview](#)
 - [Licensing](#)
 - [RAS version compatibility](#)
- [User connection flow](#)
- [Deploying Tenant Broker and Tenants](#)
- [Deploying Tenant Broker](#)
 - [The Tenant Broker console](#)
 - [Install RAS Secure Gateways](#)
- [Deploying a Tenant](#)
- [Join a Tenant to Tenant Broker](#)
 - [Overriding Tenant Broker IP address](#)
- [Joining with a secret key](#)
 - [Overriding the Tenant Broker IP address](#)
- [Verify join status](#)
- [Configure network](#)
- [Assign a public domain address](#)
- [Configure an SSL certificate](#)
- [Set up routing for incoming traffic](#)
- [User authentication](#)
- [Unjoining from Tenant Broker](#)
- [Managing Tenants](#)
- [Tenant configuration](#)
- [Deleting a Tenant object](#)

- [Opening a Tenant console](#)
- [Shared Gateways](#)
 - [Tunneling policies](#)
 - [WYSE](#)
 - [Session counters](#)
 - [Client connection routing](#)
 - [Shared gateway maintenance](#)
- [Third-party network load balancers](#)
- [Web Client and Themes](#)
 - [Web Client configuration](#)
 - [Viewing Tenant Themes in Tenant Broker](#)
- [Monitoring Tenants](#)
- [Tenant Broker compatibility and updates](#)
- [Upgrading from an older RAS version](#)
- [Configuring notifications](#)
 - [Tenant event notifications](#)
 - [Tenant Broker event notifications](#)
 - [Common event notifications](#)
- [Communication ports](#)
- [SAML SSO Authentication](#)
- [Introduction](#)
- [System requirements](#)
 - [RAS Enrollment Server](#)
 - [RD Session Hosts](#)
 - [Desktop operating systems \(guest VMs and Remote PCs\)](#)
 - [Parallels Client](#)
- [SAML basics](#)
- [SAML configuration](#)
- [Prerequisites](#)
- [IdP side configuration](#)
- [SP side configuration \(RAS side\)](#)
 - [Enable Web \(SAML\) authentication](#)
 - [Adding an IdP to the RAS Farm](#)
 - [Configuring user account attributes](#)
 - [Attributes configuration tips](#)
 - [Adding an account picture](#)
- [Active Directory user account configuration](#)
 - [Enrollment agent user account](#)
 - [NLA user account](#)
- [Configure certificate authority templates](#)

- [Create an Enrollment Agent template](#)
 - [Issue the certificate template](#)
- [Create a smartcard logon certificate template](#)
 - [Issue the certificate template](#)
- [RAS Enrollment Server configuration](#)
 - [Set up and configure RAS Enrollment Server](#)
 - [Obtain and copy the registration key](#)
 - [Configure AD integration](#)
 - [Using computer management tools](#)
- [RAS Enrollment Server high availability](#)
- [SAML integration examples and tips](#)
- [User account attributes](#)
- [Security tip](#)
- [Parallels Client configuration](#)
 - [Parallel Web Client](#)
 - [Parallels Client for Windows](#)
 - [Parallels Client for Mac](#)
 - [Parallels Client for Linux](#)
- [Parallels client policy configuration](#)
- [Test the SAML SSO deployment](#)
 - [Service provider initiated authentication](#)
 - [Identity provider initiated authentication](#)
- [Error messages](#)
 - [Error messages](#)
 - [Pre HTML5 loading](#)
 - [Post HTML5 loading](#)
 - [Once an application or desktop is launched](#)
- [Parallels Web Client and User Portal](#)
 - [System requirements](#)
- [Configure Web Client](#)
 - [Session persistence based on a cookie](#)
 - [Host header attack protection](#)
- [Configure Themes](#)
- [General settings](#)
 - [General](#)
- [Access settings](#)
- [Message settings](#)
- [Client Public IP Detection](#)
- [Web Client settings](#)
- [URLs](#)

- [Branding](#)
- [Colors](#)
- [Language bar](#)
- [Messages](#)
- [Input prompt](#)
- [Gateway](#)
- [Legal policies](#)
 - [Cookie consent](#)
 - [End User License Agreement](#)
- [Parallels Client for Windows settings](#)
 - [Branding](#)
 - [Messages](#)
 - [Custom menu](#)
 - [Create Windows client package for mass distribution](#)
- [General Theme tasks](#)
- [Delegating session management permissions](#)
 - [Create a Theme and delegate session management permissions](#)
 - [Manage sessions](#)
 - [Settings audit](#)
 - [Using Themes in Parallels Client for Windows](#)
- [Open Parallels Web Client](#)
 - [Launch apps in Parallels Client and fallback to Browser](#)
 - [Parallels Client only](#)
 - [Browser only](#)
- [Main menu options](#)
 - [Settings](#)
 - [Change Password](#)
 - [Detect Client](#)
 - [Download Client](#)
 - [Logout](#)
- [Running remote applications and desktops](#)
- [Using drag and drop functionality](#)
- [Native clipboard experience](#)
- [Other useful features](#)
- [Auto login](#)
 - [Configuration](#)
 - [Using Auto login](#)
- [Direct App access](#)
- [Using the toolbar](#)

- [Using the toolbar on desktop computers](#)
 - [Remote desktop toolbar](#)
 - [Remote application toolbar](#)
- [Using the Toolbar on Mobile Devices](#)
 - [Remote Desktop Toolbar](#)
 - [Remote application toolbar](#)
- [Using the remote clipboard](#)
- [Hiding toolbar items](#)
- [Universal Printing](#)
- [Managing Universal Printing Settings](#)
 - [Enabling and disabling Universal Printing support](#)
 - [Configuring a printer renaming pattern](#)
 - [Printer retention](#)
- [Universal Printing drivers](#)
- [Font management](#)
 - [Excluding fonts from embedding](#)
 - [Automatically install fonts on servers and clients](#)
 - [Resetting excluded fonts to default](#)
- [Universal Scanning](#)
- [Managing Universal Scanning](#)
 - [Configuring a Scanning Rename Pattern](#)
 - [Enabling and Disabling Universal Scanning Support](#)
- [Adding scanning applications](#)
 - [Adding a scanning application](#)
 - [Deleting a scanning application](#)
- [User Device Management and Client Policies](#)
- [Inviting users to connect to Parallels RAS](#)
- [Mass configuring user devices](#)
 - [Exporting and importing Parallels Client settings](#)
 - [Using Parallels Client URL scheme](#)
- [Enabling Help Desk support for users](#)
- [Enabling Help Desk support for custom administrators](#)
- [Monitoring devices](#)
 - [Device states](#)
 - [Switching off device monitoring](#)
- [Getting additional device information](#)
- [Windows device groups](#)
 - [Creating a Windows device group](#)
 - [Adding a Windows device to a group](#)

- [Managing Windows devices](#)
 - [Installing Parallels Client on a Windows computer](#)
 - [Windows device enrollment](#)
 - [Lock a Windows device](#)
 - [Shadow a Windows device](#)
 - [Desktop replacement](#)
 - [Switching to the Admin mode](#)
 - [Configuring local applications when using Parallels Client desktop replacement](#)
- [Windows desktop replacement](#)
 - [Before](#)
 - [After](#)
- [Scheduling Windows devices & groups power cycles](#)
 - [Adding a new scheduler task](#)
 - [Managing scheduled tasks](#)
- [Client policies](#)
 - [Supported Parallels Client versions](#)
- [Add a new client policy](#)
 - [Configure rules for the client policy](#)
- [Configure session settings](#)
- [Appearance](#)
- [Connection](#)
 - [Primary Connection](#)
 - [Secondary Connection](#)
 - [Reconnection](#)
 - [Computer name](#)
 - [Advanced settings](#)
 - [Web authentication](#)
 - [Session prelaunch](#)
 - [Local proxy address](#)
- [Display](#)
 - [Settings](#)
 - [Multi-monitor](#)
 - [Published applications](#)
 - [Desktop options](#)
 - [Browser](#)
- [Printing](#)
 - [RAS Universal Printing](#)
 - [Default printer settings](#)
 - [A Windows 10 and 11 note](#)

- [Scanning](#)
- [Audio](#)
- [Keyboard](#)
- [Local devices and resources](#)
 - **[Clipboard](#)**
 - [Disk drives and folders](#)
 - [Devices](#)
 - [Video capture devices](#)
 - [Ports](#)
 - [Smart cards](#)
 - [WebAuthn](#)
 - [Location](#)
 - [Pen and touch input](#)
 - [Multimedia redirection for AVD](#)
 - [File transfer](#)
- [Experience](#)
 - [Performance](#)
 - [Compression](#)
- [Network](#)
- [Server authentication](#)
- [Advanced settings](#)
- [Configure client policy options](#)
 - [Connection](#)
 - [Logging](#)
 - [Update](#)
 - [PC keyboard](#)
 - **[macOS Keyboard Mapping](#)**
 - [Single sign-on](#)
 - [Advanced](#)
 - [Global](#)
 - [Language](#)
 - [Printing](#)
 - [Windows client](#)
 - [RemoteFX USB redirection](#)
 - [Customer Experience Program](#)
- [Configure control settings](#)
 - [Connections](#)
 - [Password](#)
 - [Import and export](#)

- [Configure Gateway redirection](#)
 - [Gateway criteria](#)
- [Client policy backward compatibility](#)
- [Policy information in Parallels Client](#)
- [Configuring remote file transfer](#)
- [Configure file transfer to a server](#)
- [Configure file transfer in User Portal](#)
- [Configure file transfer for a client policy](#)
- [Reporting](#)
- [System requirements](#)
 - [Operating system requirements](#)
 - [Microsoft SQL Server requirements](#)
 - [Using Microsoft SQL Server 2017 and 2019](#)
 - [Installation locations](#)
- [Install Microsoft SQL Server](#)
- [Install Microsoft SQL Server 2016 or earlier](#)
 - [Install SQL Server Management Tools](#)
 - [Configure Microsoft SQL Server 2016 and earlier](#)
 - [Configure Microsoft SQL Server Reporting Services](#)
- [Install Microsoft SQL Server 2017 or 2019](#)
- [Install Microsoft SQL Server 2022](#)
- [Install Parallels RAS Reporting](#)
 - [Configure RAS Reporting in the RAS Console](#)
 - [Configure advanced settings](#)
- [Running Parallels RAS Reports](#)
 - [Predefined reports](#)
- [GDPR compliance](#)
- [Performance monitor](#)
- [Overview](#)
 - [Components](#)
 - [How it works](#)
- [Install RAS Performance Monitor](#)
 - [Requirements](#)
 - [Prerequisites](#)
- [Using Parallels RAS Performance Monitor](#)
 - [Configure access to Parallels RAS Performance Monitor](#)
 - [Viewing performance metrics](#)
- [Cost Insights](#)
 - [Cost Insights](#)
 - [System requirements](#)

- [Install Parallels RAS Cost Insights](#)
- [Configure Parallels RAS Cost Insights](#)
- [Managing Parallels RAS Cost Insights](#)
- [Definitions and calculations](#)
- [Recovery – add a root administrator](#)
- [Host name resolution](#)
 - [Adding a RAS Connection Broker](#)
 - [Adding a RAS Secure Gateway](#)
 - [Adding an RD Session Host](#)
 - [Add a provider](#)
- [Computer management tools](#)
 - [Requirements for using computer management tools](#)
 - [Available tools](#)
- [Site information](#)
- [Site settings](#)
 - [Auditing](#)
 - [Global logging](#)
 - [URL redirection](#)
 - [Notifications](#)
 - [Client settings](#)
 - [Features](#)
 - [Application packages](#)
- [Using MSIX application packages](#)
 - [Prerequisites](#)
 - [Enabling the Application Packages feature](#)
 - [Creating an MSIX image](#)
 - [Adding an MSIX application package to Parallels RAS](#)
 - [Adding a package to a host](#)
 - [Adding a package to VDI pool](#)
 - [Adding a package to an AVD pool](#)
 - [Adding a package to Group Defaults](#)
 - [Adding a package to Site Defaults](#)
 - [Working with version tags](#)
 - [Working with certificates](#)
 - [Managing MSIX application packages](#)
 - [Package properties](#)
 - [Package statuses](#)
- [Using template versions](#)
 - [Supported providers](#)
 - [Creating a new version](#)

- [Working with version tags](#)
- [Deleting a version](#)
- [Settings audit](#)
- [Settings audit](#)
 - [View the main settings audit list](#)
 - [Turning off audit databases synchronization](#)
 - [Information about modifications](#)
 - [Common tasks](#)
 - [Reverting a modification](#)
 - [View a local settings audit list](#)
 - [Change audit settings](#)
- [Upgrading RAS agents](#)
- [Licensing](#)
- [Configure HTTP proxy settings](#)
- [System event notifications](#)
- [Configuring notification handlers](#)
 - [System events](#)
- [Configuring notification scripts](#)
 - [Command line variables](#)
 - [Notification types](#)
- [Configuring SMTP server connection for event notifications](#)
- [RAS session variables](#)
 - [Examine the registry](#)
 - [Using GetRASVariable.exe utility](#)
- [Resolving z-order issues](#)
- [Maintenance and backup](#)
 - [Keeping Parallels RAS up to date](#)
 - [Backing up the Parallels RAS Farm configuration](#)
- [Importing and exporting Farm settings from the command line](#)
 - [Installing Parallels RAS PowerShell](#)
 - [Using Parallels RAS PowerShell](#)
- [Problem reporting and troubleshooting](#)
 - [Search for a solution](#)
 - [Request support](#)
 - [Send a report](#)
 - [Contacting the local support](#)
- [Logging](#)
 - [Log rotation](#)
- [Suggest a feature](#)
- [Parallels RAS Management Portal](#)

- [Overview](#)
- [Prerequisites](#)
- [Installation](#)
- [Log in to RAS Management Portal](#)
- [Configure RAS Web Administration Service](#)
- [RAS Management Portal user interface](#)
 - [Categories](#)
 - [Admin permissions](#)
 - [Subcategories](#)
 - [Navigation bar](#)
 - [Breadcrumbs](#)
 - [Page header items](#)
 - [Editing](#)
 - [Edit toolbar](#)
 - [Wizards](#)
 - [Modal dialogs](#)
 - [Object properties views](#)
- [Parallels RAS APIs](#)
- [RAS PowerShell API](#)
 - [Parallels RAS requirements](#)
 - [Microsoft Windows component requirements](#)
 - [Installation](#)
 - [RAS PowerShell API versions](#)
 - [RAS PowerShell API concepts](#)
 - [Parallels RAS PowerShell Guide](#)
- [RAS REST API](#)
- [Installation](#)
 - [Configure RAS Web Administration Service](#)
- [Permissions](#)
- [Getting started](#)
- [Logging in and sending requests](#)
 - [Log in to Parallels RAS and obtain a session token](#)
 - [Retrieve information about RD Session Hosts](#)
 - [Retrieve information about a specific RD Session Host](#)
 - [Modify RD Session Host properties](#)
- [More information](#)
- [RAS Web Client API and Parallels Client URL scheme](#)
- [Appendix](#)
- [Microsoft license requirements in Parallels RAS](#)
 - [General](#)

- [RD Session Hosts](#)
- [VDI](#)
- [Remote PC](#)
- [Azure Virtual Desktop](#)
- [FSLogix](#)
- [Microsoft SQL Server](#)
- [App-V](#)
- [Other References](#)
- [Port reference](#)
- [Parallels Client](#)
- [Web browsers](#)
- [HALB](#)
- [RAS Secure Gateway](#)
- [RAS Connection Broker](#)
- [RAS Console](#)
- [SSRS](#)
- [RAS Reporting](#)
- [RAS Web Administration Service \(REST/Management Portal\)](#)
- [RAS PowerShell](#)
- [RAS Provider Agent](#)
- [RAS Enrollment Server](#)
- [RAS RD Session Host Agent](#)
- [RAS Guest Agent](#)
- [RAS Remote PC Agent](#)
- [Tenant Broker](#)
- [Active Directory and Domain Services ports](#)
- [Azure Virtual Desktop](#)
- [RAS performance counters](#)

Introduction

Welcome to Parallels® Remote Application Server (Parallels RAS), an integrated solution to virtualize your applications, desktops and data. Parallels RAS publishes applications and delivers remote and virtual desktops to any device on your network, anywhere.

Parallels RAS 21 release history

The following table lists the Parallels RAS 21 release history. Parallels RAS documentation is updated for every release. This guide refers to the latest Parallels RAS 21 release from the table below. If you are using a newer Parallels RAS release or version, please download the current version of the guide from <https://www.parallels.com/products/ras/resources/>.

```
include "../../../21_Release/10_RAS-21_Admin_Guide/en_US/.gitbook/
includes/ras-21-release-history-table.md" %}
```

About Parallels RAS

Parallels RAS provides vendor independent virtual desktop and application delivery from a single platform. Accessible from anywhere with platform-specific clients and web enabled solutions, like the built-in Parallels Web Client, Parallels RAS allows you to publish remote desktops, applications and documents, improving desktop manageability, security and performance.

Parallels RAS extends Windows Remote Desktop Services by using a customized shell and virtual channel extensions over the Microsoft RDP protocol. Parallels RAS supports all major hypervisors from Microsoft, VMware, and other vendors including Hyperconverged solutions such as Nutanix AHV (AOS) and Scale Computing and Cloud platforms and services such as Microsoft Azure and Azure Virtual Desktop (formerly known as Windows Virtual Desktop), enabling the publishing of virtual desktops and applications to Parallels Client.

The product includes powerful universal printing and scanning functionality, as well as resource-based load balancing and management features.

With Parallels Device Manager Module for Parallels RAS you can also centrally manage user connections and PCs converted into thin clients using the free Parallels Client.

How does it work?

When a user requests an application or a desktop, Parallels RAS finds a least loaded RD Session Host or a guest VM on one of the least loaded Providers and establishes an RDP connection with it. Using Microsoft RDP protocol, the requested application or desktop is presented to the user. Note that in addition to

RD Sessions Hosts and VDI, Parallels RAS can also be used to configure, manage and publish Azure Virtual Desktop resources.

Users can connect to Parallels RAS using Parallels Client (available at no charge), which can run on Windows, Linux, macOS, Android, Chrome, iOS and iPadOS. Users can also connect via an HTML5 browser or Chromebook.

As newer versions of Windows keep on being developed as time goes by, you need to defend the migration cost to your business. Parallels RAS can help. Desktop replacement allows you to extend the lifespan of your hardware and delay migration to the latest OSs to a time that suits you best. The Parallels RAS solution allows you to be very flexible: you can lock machine configurations on the user side, placing your corporate data in an extremely secure position; or you can opt to allow users to run some local and remote applications. Parallels Client Desktop Replacement is able to reduce the operability of the local machine by disabling the most common local configuration options, while guaranteeing the same level of service and security afforded by thin clients, directly from your existing PCs.

About this guide

This guide is intended for system administrators responsible for installing, configuring, and administering Parallels RAS. This guide assumes that the reader is familiar with Microsoft Remote Desktop Services and has an intermediate networking knowledge.

What's new

Parallels RAS 21

Parallels RAS 21.0 introduces a wide range of new features and enhancements focused on delivering greater security, scalability, and a more seamless user experience. This release expands automation and administrative controls, improves accessibility compliance, and provides deeper insights into cost consumption and potential savings that can be achieved through Parallels RAS. Together, these updates help organizations optimize their environments while preparing for future needs.

Below are some new features that were added in Parallels RAS 21:

- Support for Windows Server 2025 Hyper-V Standalone and Cluster.

- FSLogix 25.06 support.
- Support for new client operating systems, including macOS 26, iOS 26, iPadOS 26, Android 16, and more.
- Updated HALB implementation with upgraded third-party packages with improved security, stability, and long-term maintainability.
- The ability to capture the client's public (WAN) IP address and pass it as a RADIUS attribute during MFA authentication, allowing administrators to enforce IP-based security policies, enable contextual access control, and improve auditing across all native clients, the user portal, and web connections.
- New Cost Insights dashboard for Azure Virtual Desktop, which enables administrators to monitor and visualize estimated cloud service costs as well as cloud cost savings directly in RAS Console.
- Fine-grained clipboard redirection and updated clipboard data type restrictions (plain text, RTF, images).
- Improved Organizational Unit (OU) traversal for cross-domain provisioning.
- The option to reassign orphaned hosts to host pools.
- The ability to prevent creating new templates or template versions if the RAS agent is outdated.
- Validation of domain credentials during template creation to prevent cloning failures.
- Seamless, policy-driven URL redirect for User Portal (Web Client) client users to automatically migrate from one farm to another.
- UI/UX improvements for the process of reconnecting to published desktops, with clear visual feedback.
- Support for Microsoft Azure Virtual Desktop (AVD) Sessions in the Windows Client app.
- Clipboard History and Clipboard Sync Integration for Windows devices.
- Conditional RADIUS Automation to provide more granular control for MFA based on user access context.
- The ability to use Windows Credential Manager with the Windows Client app.
- Web Content Accessibility Guidelines (WCAG) 2.2 Level AA Compliance for the User Portal/Web Client.

To learn more the new features of RAS 21, read [this KB article](#).

Terms and abbreviations used in this guide

Note: Starting with Parallels RAS 19, all products and documentation, including this section, use updated terminology. To see what terms were changed, go to <https://kb.parallels.com/en/128943>.

Term/ Abbreviation	Description
RAS Console	
Category	
Farm	A Parallels RAS Farm is a logical grouping of objects for the purpose of centralized management. A Farm configuration is stored in a single database which contains information about all objects comprising the Farm. A Farm consists of at least one Site but may have as many sites as necessary (see Site below).
Site	A Site consists of at least one RAS Connection Broker, RAS Secure Gateway (or multiple gateways), and RAS agents installed on RD Session Hosts, Providers, and Windows PCs. Note that a given RD Session Host, Provider, or PC can be a member of only one Site at any given time.
Licensing Site	
RAS Secure Gateway	RAS Secure Gateway tunnels all traffic needed by applications on a single port and provides secure connections.
Web Client	Web Client allows users to view and launch remote applications and desktops in a web browser. The Web Client functionality is a part of RAS Secure Gateway.
Publishing	The act of making items installed on a Remote Desktop Server, Provider or Remote PC available to the users via Parallels RAS.
RAS Connection Broker	RAS Connection Broker provides load balancing of published applications and desktops.
RAS RD Session Host Agent	RAS RD Session Host Agent collects information from Microsoft RDS hosts required by the Connection Broker and transmits to it when required.
Remote PC Agent	

Term/ Abbreviation	Description
	Remote PC Agent collects information from Remote PC hosts required by the Connection Broker and transmits to it when required.
RAS Guest Agent	RAS Guest Agent collects information from the VDI desktop required by RAS Connection Broker and transmits to it when required.
RAS Provider Agent	RAS Provider Agent collects information from the Parallels RAS Infrastructure and is responsible for controlling VDI through its native API. RAS Provider Agent is built into the RAS Connection Broker and is available by default. It can be used to control multiple Providers in a Parallels RAS Farm.
RAS Provider Agent dedicated	RAS Provider Agent dedicated is similar to the RAS Provider Agent described above with one important difference — it is a separate component that must be installed from the Parallels RAS installer and can only control a single Provider.
RDSH or RD Session Host	RDSH makes applications and a full desktop accessible to a remote client that supports Remote Desktop Protocol (RDP). RDSH replaced Terminal Servicer beginning with Windows 2008 R2.
HALB	
Tenant Broker	Tenant Broker is a special RAS installation that hosts shared RAS Secure Gateways. It is an essential part of the RAS multi-tenant architecture.
Tenant	Tenants are RAS farms that join Tenant Broker (see above) and use shared RAS Secure Gateways and HALB thus eliminating the need to have their own Gateways and HALB deployed.
RAS Enrollment Server	RAS Enrollment Server is an essential component of the SAML SSO Authentication functionality. It communicates with Microsoft Certificate Authority (CA) to request, enroll, and manage digital certificates on behalf of the user for SSO authentication in the Parallels RAS environment.
RAS PowerShell	Parallels RAS PowerShell allows you to perform Parallels RAS administrative tasks using PowerShell cmdlets. You can execute cmdlets in the Windows PowerShell console or you can write scripts to perform common Parallels RAS administrative tasks. A complete guide to Parallels RAS PowerShell is available on the

Term/ Abbreviation	Description
	Parallels website together with other Parallels RAS documentation.
RAS REST API	Parallels RAS comes with various APIs to help you develop custom applications that integrate with it. The RAS REST API is one of them.
RAS Management Portal	Parallels RAS Management Portal is an HTML5 browser-based application that lets you manage Parallels RAS.
RAS Web Administration Service	A Web service that provides the user interface for RAS Management Portal and implements RESTful Web services for the RAS REST API (see above).
Azure Virtual Desktop	Azure Virtual Desktop is a desktop and app virtualization service running on Microsoft Azure, providing access to RD Session Hosts and VDI. Parallels RAS 18 provides the ability to integrate, configure, maintain, support and access Azure Virtual Desktop workloads on top of the existing technical capabilities of Parallels RAS.
FSLogix	FSLogix Profile Container is a remote profile solution for non-persistent environments. Parallels RAS supports FSLogix on RD Session Hosts, VDI, and Azure Virtual Desktop.

Installing Parallels RAS

This chapter describes how to install and activate Parallels RAS.

System requirements

Before installing Parallels RAS, please verify that your hardware and software meet or exceed the hardware and software requirements described below. Please note that although Parallels RAS can be used in Workgroup environment, Parallels recommends using Active Directory to manage users, groups, and machine accounts via group policies.

Hardware requirements

Parallels RAS is extensively tested on both physical and virtual platforms. The minimum hardware requirements approved to run Parallels RAS are outlined below.

- Physical Machines – Dual Core Processor and a minimum of 4GB RAM.
- Virtual Machines – Two Virtual Processors and a minimum of 4GB of RAM.

The server hardware requirements to install and configure Parallels RAS can vary according to end-user requirements.

Typically for an installation of 30 users or under, Parallels RAS can be installed on one high specification server and the resources published directly from it. For more than 30 users, multiple servers may be required.

The below should be considered during the planning stage of a Parallels RAS deployment:

- High specification servers should be used, consisting of multiple CPU cores, a high specification disk transfer rate and plenty of RAM.
- A hypervisor-based virtual machine can be used as long as the resources needed to serve end-users are calculated accordingly.
- It is recommended that RAS Secure Gateway does not exceed 1200 users per server for incoming connections using the Gateway SSL mode.
- When planning VDI Hypervisor resource requirements, extra requirements such as RAM usage per virtual machine and disk space should be taken into account.

When configuring RD Session Hosts, VDI, or Azure Virtual Desktop, please keep in mind that different types of workloads require different session host configurations. For the best possible experience, scale your deployment depending on your users' needs. The following table gives you an idea of how different workload types affect session host configurations.

Workload	Example users	Example apps	Max users per vCPU	Minimum
Light	Basic data entry tasks	Database entry applications, command-line interfaces	6	

Workload	Example users	Example apps	Max users per vCPU	Minimum
Medium	Consultants and market researchers	Database entry applications, command-line interfaces, Microsoft Word, static web pages	4	
Heavy	Software engineers, content creators	Database entry applications, command-line interfaces, Microsoft Word, static web pages, Microsoft Outlook, Microsoft PowerPoint, dynamic web pages	2	
Power	Graphic designers, 3D model makers, machine learning researches	Database entry applications, command-line interfaces, Microsoft Word, static web pages, Microsoft Outlook, Microsoft PowerPoint, dynamic web pages, Adobe Photoshop, Adobe Illustrator, CAD, CAM	1	

Note: Sizing guidelines are based on Microsoft recommendations on RDS or Azure Virtual Desktop multi-session hosts.

For port requirements, please see the **Port Reference** section.

Software requirements

RAS Connection Broker and RAS Secure Gateway (64-bit versions only)

RAS Connection Broker and RAS Secure Gateway are supported on the following operating systems:

- Windows Server 2016 up to Windows Server 2025
- On Windows Server 2016, 2019, 2022, and 2025, both Server Core and Desktop Experience installations are supported.

Note: RAS Connection Broker and RAS Secure Gateway should not be installed on a domain controller or any other machine where a DHCP server is running. This applies generally to all RAS components.

RAS Web Administration Service

The same OS requirements as for RAS Connection Broker (see above). Note that for larger environments (2000 or more concurrent connections), it is recommended to install the component on a dedicated server. For details, please see <https://kb.parallels.com/en/124988>.

RAS RD Session Host Agent

RAS RD Session Host Agent is supported on the following operating systems:

- Windows Server 2012 R2 up to Windows Server 2025
- Windows Server 2016 and newer must be installed using the “Desktop Experience” installation option.
- Windows Server 2012 R2 — Server Core installation option is not supported.

RAS Provider Agent

- Windows Server 2016 up to Windows Server 2025

For the list of supported Providers, see [RAS Provider Agent Installation Options](#).

RAS Guest Agent

- Windows Server 2012 R2 up to Windows Server 2025
- Windows 10 up to Windows 11

Remote PC Agent

- Windows Server 2012 R2 up to Windows Server 2025
- Windows 10 up to Windows 11

Parallels RAS PowerShell

- Windows Server 2016 up to Windows Server 2025
- Windows 10 up to Windows 11

- Windows Management Framework 3.0 and .NET Framework 4.5.2 must be installed

Parallels RAS Console

- Windows Server 2016 up to Windows Server 2025
- Windows 10 up to Windows 11

RAS Enrollment Server

- Windows Server 2016 up to Windows Server 2025

RAS Performance Monitor

- Windows Server 2016 up to Windows Server 2025

RAS Reporting

See [System requirements](#)

Parallels Client

Parallels Client is approved for the following operating systems (both 32-bit and 64-bit systems are supported, where applicable):

- Windows 10, 11
- Windows Server 2012 R2 up to Windows Server 2025
- macOS Ventura 13 to macOS Tahoe 26. Parallels Client runs natively on Intel and Apple silicon processors.
- iOS/iPadOS 17 up to 26
- Android 10 up to 16
- Chrome OS

Note: Parallels Client for Chrome is deprecated. We recommend using Parallels Web Client instead.

Parallels Client for Linux supports the following Linux distributions (x64 versions only):

- Ubuntu 22.04 LTS, 24.04 LTS (Gnome)
- Debian 11 (Bullseye), Debian 12 (Bookworm), Debian 13 (Trixie)
- Fedora 40, 41, 42, 43, 44
- Linux Mint 21, 22

- IGEL OS 11, 12

Parallels Client supports all default window managers of the distributions listed above. If you use a different window manager, your experience may vary from the intended. endhint %}

For a list of supported thin clients and supported hardware from Technology Partners such as Igel, HP, 10Zig, and more, please see the following KB article: <https://kb.parallels.com/124606>.

Attention: With the future release of Parallels RAS 22.0 the following operating systems will no longer be supported for the Parallels Client:

- Windows: 10 x86, Server 2012R2
- macOS: 13
- Android: 10, 11
- iOS/iPadOS: 17
- Fedora: 40, 41
- Debian: 11
- IGEL: OS 11

Microsoft license requirements

For information about Microsoft license requirements, such as Remote Desktop Services Client Access Licenses (RDS CALs) and Virtual Desktop Access (VDA) licenses, please see [Appendix: Microsoft license requirements in Parallels RAS](#).

Install Parallels RAS

To install Parallels RAS:

1. Make sure you have administrative privileges on the computer where you are installing Parallels RAS.
2. Double click the `RASInstaller.msi` file to launch the Parallels RAS installation wizard. If you see a message that begins with “This version of Parallels RAS is only for testing purposes.”, it means that it’s not an official build and should not be used in a production environment.
3. Follow the onscreen instructions.

Note: Please ensure that the presented terms in the license agreement are read and accepted to complete installation and/or upgrade. For programmatic deployment, it is understood that the terms in the license agreement have been read and accepted.

Note: If you are upgrading from one of the major versions (for example, from Parallels RAS 18 to Parallels RAS 19), you will see a message that lists system requirements for every component of the new version. Please read them carefully to make sure that all components can be upgraded in your environment. Note that if you install a component on a system that does not meet its system requirements, the component will not work.

Help us improve our products!

When you install Parallels RAS, you can choose to join Parallels Customer Experience Program. For more information about Parallels Customer Experience Program, see <https://www.parallels.com/about/legal/pcep/>.

1. Proceed to the **Select Installation Type** page and select from the following:
 - **Parallels Remote Application Server.** The default installation that will install RAS Console, RAS Management Portal, RAS Connection Broker, RAS Secure Gateway, RAS RD Session Host Agent, RAS PowerShell, and RAS Web Administration Service on the same machine. This is ideal for testing or small production environments.
 - **Parallels RAS Tenant Broker.** This option installs Tenant Broker. Please note that Tenant Broker must be installed on a server separate from the existing RAS farms. For more information about Tenant Broker, please see the [RAS Multi-Tenant Architecture](#) chapter.
 - **Custom.** Select and install only the components that you require. You can select individual components after you click **Next**. Note that if a component cannot be installed on the current server, it will not be available for installation. See **Software Requirements**.
2. Click **Next**.
3. Review the notice on the **Important Notice** wizard page. If there's a port conflict on your computer, the information will be displayed here. You can resolve the conflict later.
4. Click **Next**.
5. On the **Firewall Settings** page, select **Automatically add firewall rules** to configure the firewall on this computer for Parallels RAS to work properly. See **Port Reference** for details.

6. Click **Next** and then click **Install**. Wait for the installation to finish and click **Finish**.
7. If you are upgrading your RAS installation, it is recommend to reboot all servers where components are upgraded.

When you need to install a particular Parallels RAS component on a different server, run the installation wizard again, select **Custom** and choose the component(s) you wish to install.

Log in and activate Parallels RAS

After you've installed Parallels RAS, run the RAS Console and activate your new Parallels RAS Farm.

Start the Parallels RAS Console

By default, the Parallels RAS Console is launched automatically after you click **Finish** on the last page of the installation wizard. To launch the console manually, navigate to **Start > Apps > Parallels** and click on **Parallels Remote Application Server Console**.

When the Parallels RAS Console is launched for the first time, you are presented with the login dialog. In the dialog, specify the following:

- **Farm:** A Parallels RAS Farm to connect to. Enter the FQDN or IP address of the server where you have RAS Connection Broker installed.
- If you've installed the Parallels Single Sign-On component when installing the RAS Console, you will see the **Authentication type** field from which you can select whether to log on using your credentials or SSO. If you reboot after the installation and select SSO, select **Single Sign-On** and then click **Connect**. Your Windows credentials will be used to log in to the RAS Farm. If you select **Credentials**, enter your credentials as described below.
- **Username:** A user account with administrative privileges on the server where Parallels RAS is installed (usually a domain or local administrator). The account name must be specified using the UPN format (e.g. `administrator@domain.local`). The specified user will be automatically configured as the Parallels RAS administrator with full access rights.
- **Password:** The specified user account password.

- If you select the **Remember credentials** option, this dialog will not be shown the next time you launch the Parallels RAS Console.

The **Edit Connections** button opens a dialog where you can manage your RAS connection. This dialog becomes useful if this is not the first time you are connecting to one or more of your RAS Farms. The left pane of the dialog displays RAS Farms to which previously connected (you can remove a Farm from the list by clicking the [-] icon if you no longer need it). The right pane displays at least the primary Connection Broker for the selected Farm. If you've added a secondary Connection Brokers to a Farm, you can add it to this list by clicking the [+] icon and typing its hostname or IP address (click the "recycle" icon to verify the agent status). This way the RAS Console will try to connect to the primary Connection Broker first and if it fails (e.g. the agent is offline or cannot be reached), it will try to connect to the secondary Connection Broker. For more information about secondary Connection Brokers, please see [Parallels RAS Connection Brokers](#) chapter.

When you are done entering the connection information, click the **Connect** button to connect to the Parallels RAS Farm.

Sign in to Parallels My Account

To activate Parallels RAS, you must register for a Parallels business account. After you logged in to Parallels RAS, you'll see the **Sign In to Parallels My Account** dialog. If you already have an account, type the email address and password you used to register the account and click **Sign In**.

Note: If you use an HTTP proxy server on your network, you will see a dialog asking you to configure the proxy server connection settings. Click the **Configure Proxy** button. In the dialog that opens, select one of the following: **Use system proxy settings** (the default proxy settings from the Internet Explorer will be used) or **Manual HTTP proxy configuration** (specify the settings manually). If your proxy configuration changes, you can re-configure it later by navigating to **Administration > Settings** and clicking the **Configure Proxy** button.

If you don't have a Parallels business account, you can register for one as follows:

1. In the **Sign In to Parallels My Account** dialog, click **Register**. The **Register Parallels My Account** dialog opens.
2. Enter your name and email address, choose and type a password, and enter your company info (all fields are required).

3. Follow the links to Parallels Privacy Policy and Terms of Use. After reading them (and if you agree) select the **I have read and agree to the Parallels Privacy Policy and Terms of Use** checkbox.
4. Click **Register** to register an account. This will create a personal account for yourself and a business account for your organization to which you will be assigned as administrator.

Activate Parallels RAS

After you sign in to Parallels My Account, the **Activate Product** dialog opens asking you to activate the Parallels RAS Farm.

If you already have a Parallels RAS license key, select the **Activate using license key** option and enter the key in the field provided. You can click the button next to the field to see the list of subscriptions and/or permanent license keys you have registered in Parallels My Account. If the list is empty, it means that you don't have any subscriptions or license keys and need to purchase one first.

Note: You can manage your Parallels RAS license using the **Licensing** category in the Parallels RAS console. The management tasks include viewing the license information, switching to a different Parallels My Account, and activating the Parallels RAS Farm using a different license key. For more information, please see the [Licensing](#) section.

If you don't have a Parallels RAS license key, you have the following options:

- Purchase a subscription online by clicking the **Purchase a license** link.
- Activate Parallels RAS as a trial by selecting the **Activate trial version** option.

After entering a license key (or selecting to activate a trial version), click **Activate**. You should see a message that the Parallels RAS Farm was activated successfully. Click **OK** to close the message box.

The first dialog that you see informs you that you have no servers configured that can be used to host published resources. This means that to begin using Parallels RAS, you need at least one RD Session Host, Provider, or a Remote PC configured. We'll talk about configuring a Parallels RAS Farm in the next chapter. For now, click **OK** to close the message box. You will then see the **Applying Settings** dialog. Wait for the initial configuration of Parallels RAS to complete and

click **OK**. You will now see the main Parallels RAS Console window where you can begin configuring the Parallels RAS Farm.

Read on to learn how to quickly add an RD Session Host, publish resources, and invite your users to Parallels RAS.

Getting Started with Parallels RAS

This chapter will help you get started with Parallels RAS. Read it to learn how to use the Parallels RAS Console and how to set up a simple RAS environment.


The Parallels RAS Console

The Parallels RAS Console is a Windows application used to configure and administer a Parallels RAS Farm.

To open the Parallels RAS Console, navigate to **Apps > Parallels** and click **Parallels Remote Application Server Console**. Note that you can open multiple instances of the Parallels RAS Console on the same computer if you want to manage more than one Farm or Site simultaneously without switching between them inside the console. This works with a locally installed Parallels RAS Console and when you run it as a remote application from Parallels Client.

Information: In addition to Parallels RAS Console, Parallels RAS 18 introduced Parallels RAS Management Portal, an HTML5 browser-based console that lets you manage Parallels RAS. Note that at the time of this writing, Parallels RAS Management Portal does not completely replace the desktop RAS Console as some management features are still in development. More features will be added in the upcoming releases. For more information, please refer to **Parallels RAS Management Portal Guide**, which is available on the Parallels website: <https://www.parallels.com/products/ras/resources/>.

The following screenshot and the description below it give you an overview of Parallels RAS Console.



Parallels RAS Console

The Parallels RAS Console consists of the following sections:


1	This section lists categories. Selecting a category will populate the right pane with elements relevant to that category.
2	This section (the middle pane) is available only for the Farm and the Publishing categories. The navigation tree allows you to browse through objects related to that category.
3	This section displays the selected object or category properties, such as servers in a Farm or published application properties, etc.
4	
5	The information bar at the bottom of the screen is used to display the most recent console notification (if one is available).

Set up a basic Parallels RAS Farm

In this section, we'll set up a basic Parallels RAS Farm where all required components run on a single server.

To set up a Parallels RAS Farm:

1. Log in to the Parallels RAS Console.
2. In the console, select the **Start** category. This category gives you access to three wizards that you can use to easily perform essential tasks, such as adding RD Session Hosts, publishing applications, and inviting users to Parallels RAS.




Add an RD Session Host

First, you need to add an RD Session Host to the Farm. In this tutorial, we'll add the local server on which Parallels RAS is installed.

To add an RD Session Host to the Farm:

1. Click **Add RD Session Hosts**. The **Add RD Session Hosts** wizard opens.



2. Click the Tasks menu (or click the **[+]** icon) and select one of the following:

- **Add from Active Directory:** Adds an RD Session Host from Active directory.
- **Add Manually:** Adds RD Session Host by entering its FQDN or IP address.


Note that if you enter the server FQDN, it will be used as the primary method of connecting to this server from other Parallels RAS components and clients. If you enter the IP address, it will be automatically resolved to FQDN, but only if the global option to resolve to FQDN is enabled. To see the current setting of this global option, click **Tools > Options** on the main menu. In the **Options** dialog, examine the **Always attempt to resolve to fully qualified domain name (FQDN) when adding hosts** option. When the option is selected, the IP address of every server/component in the RAS Farm is always resolved to FQDN. When the option is cleared, whatever is specified for a server (IP address or name) is used to communicate with a server. This makes a difference in deployments where an IP address cannot be used to access a server, such as when a server is hosted in the cloud. For more information, see [Host Name Resolution](#).

3. Click **Next**.

4. The page with general settings opens:

Specify the following settings:


- **Add firewall rules.** Add firewall rules required by Parallels RAS in Windows running on the server. See **Port Reference** for details.
- **Install RDS role.** Install the RDS role on the server if it's not installed. You should always select this option.
- **Enable Desktop Experience.** Enable the Desktop Experience feature in Windows running on the server. This option is enabled only if the Install RDS role option (above) is selected. The option applies to Windows Server 2008 R2 and Windows 2012 R1/R2 on which the Desktop Experience feature is not enabled by default.
- **Restart server if required.** Automatically restart the server if necessary. You can restart the server manually if you wish.



5. Click **Next**.
6. Add the server (or servers) to a host pool. Select the desired host pool or create a new host pool. If you are not sure what host pool to choose, select **Default Host pool**. Host pools are described in detail in the [Manage host pools \(RD Session Hosts\)](#) section.
7. Click **Next**.
8. In order for end users to access published resources on the RD Session Host, they must be added to the Remote Desktop Users group in Windows running on the server. This can be done one of the following ways:
 - Adding each user or group directly on the server using standard Windows administrative tools.
 - Adding users or groups through Active Directory.
 - Using the wizard page described below, which is provided for your convenience.

If you already added your users to the Remote Desktop Users group on the given server (or if for any reason you want to use one of the other methods listed above), you can simply click **Next** and skip this page.


To add users to the Remote Desktop Users group using the wizard, select the **Specify users or groups to be added to the Remote Desktop Users group** option and then click the **[+]** icon. In the **Select Users or Groups** dialog, specify a user or group and click **OK**. The selected user/group will be added to the list on the wizard page.



9. Click **Next**.


10. The **User profile** page allows you to select a technology to manage user profiles.

You can select from **User profile disk** or **FSLogix**. User Profile Disks are virtual hard disks that store user application data on a dedicated file share. Microsoft FSLogix Profile Container is the preferred Profile Management solution as the successor of Roaming Profiles and User Profile Disks (UPDs). It is set to maintain user context in non-persistent environments, minimize sign-in times and provide native profile experience eliminating compatibility issues. You can keep the default settings for now. We will talk in detail about [user profiles later in this guide](#).



11. The **Optimization** page allows you to specify settings that will be used to optimize Windows on the RD Session Host for best performance in a Parallels RAS environment.

You can select Windows components, services, and other options that will be disabled, removed, or optimized to ensure a more efficient, streamlined, and improved delivery of virtual apps and desktops. You can keep the default settings or you can modify (or disable if not sure) optimization for now. Optimization is described in [detail later in this guide](#).



12. On the next page, review the settings and click **Next**.

13. The **Install RAS RD Session Host Agent** dialog opens. Follow the instructions and install the agent. When the installation is finished, click **Done** to close the dialog.

14. Back in the wizard, click **Finish** to exit.

If you would like to verify that the RD Session Host has been added to the Farm, click the **Farm** category (below the **Start** category in the left pane of the Parallels RAS Console window) and then click **RD Session Hosts** in the navigation tree (the middle pane). The server should be included in the **RD Session Hosts** list. The **Status** column may display a warning message. If it does, reboot the server. The **Status** column should now say, “OK”, which means that your RD Session Host is functioning properly.


Read on to learn how to [publish an application from an RD Session Host](#).

Publish applications

After you added an RD Session Host, you can publish applications from it.


To publish an application:

1. In the Parallels RAS Console, select the **Start** category and click the **Publish Applications** item in the right pane.
2. The **Publish Applications** wizard opens. On the first page, select one or more servers from which the application should be published. You can select all servers, server host pools, or individual servers.



3. Click **Next**.
4. On the next page, select one or more applications you want to publish.

If you've selected more than one server on the previous screen, the **Show applications not available on all target servers** option becomes enabled. If the option is cleared (default), the folder tree will contain applications that are available on each and every server that you selected. If the option is enabled, the tree will contain applications that may be available on some server(s), but not on the others.



5. Click **Next**. Review the summary information and click **Next** again.

6. Click **Finish** when ready.

To verify that an application has been successfully published, select the **Publishing** category in the RAS Console. The application should be included in the **Published Resources** list (the middle pane).

Invite users

Your Parallels RAS Farm is now fully operational. You have an RD Session Host and published application(s). All you need to do now is invite your users to install the Parallels Client software on their devices and connect to the Parallels RAS Farm.

Note: Consider allowing users to access the published resources by using their email instead of Secure Gateway IP address or hostname. For information on how to do it, see [Allowing users to discover RAS connections via email address](#).

To invite users:

1. In the Parallels RAS Console, select the Start category and click the **Invite Users** item.

2. The **Invite Users** wizard opens. On the first page, specify the mailbox information that should be used to send invitation emails to users.

Specify the following options:

- **Mailbox Server:** Enter the mailbox server name. For example, mail.company.com:500
- **Sender Address:** Enter the email address.
- **TLS / SSL:** Choose whether to use the TLS/SSL protocol. (* **see below**)
- **SMTP server requires authentication:** Select this option if your SMTP server requires authentication. If it does, also type the username and password in the fields provided.

In the **Test Email** section, type one or more email addresses to which a test email should be sent (separate multiple address with a semicolon). Click the **Send Test Email** button to send the email.

Invite Users - Mailbox Configuration

Parallels

Configure the mailbox from where the invitations will be sent from.

Mailbox configuration

Mailbox Server: mail. .com:500
Example: mail.yourcompany.com:500

Sender Address: admin@ .com

TLS / SSL: Do not use

☒ SMTP server requires authentication

Username: admin

Password:

Test email

Separate email addresses by a semicolon to send a Test Email to multiple addresses.

mikef@ .com;andys@ .com

Send Test Email

< Back Next > Cancel Help

3. Click **Next**.
4. On the next page of the wizard, specify target platforms and connection options:
 - In the target devices list, select the types of devices to send an invitation to. Each target device of a particular type will receive an email

with instructions on how to download, install, and configure the Parallels Client software on that device type.

- In the **Public address** field, specify a public FQDN or IP address. This setting is used by the Preferred routing functionality to redirect client connections. Please see [Configuring preferred routing](#).
- In the **Connection Mode** drop-down list, select the RAS Secure Gateway connection mode. Please note that SSL modes require the gateway to have SSL configured. More information can be found in the [Configuring RAS Secure Gateway](#) section.
- In the **Authentication mode** drop-down list, select the authentication mode for your users. For the list of authentication modes, see subsection **Primary connection** in the [RAS Connection Broker connection settings](#) section.
- Optionally, click the **Advanced** button to open the **Advanced Settings** dialog. This dialog allows you to specify a third-party credential provider component. If you use such a component to authenticate your users, specify its GUID in this dialog. For more information, see [Configure Client Policy Options > Single Sign-On](#).

Invite Users - Options

Parallels

Specify target platform:

Name

- ☒ Windows
- ☐ User Portal (Web Client)
- ☐ macOS
- ☐ Linux
- ☐ iOS/iPadOS
- ☐ Android
- ☐ Chrome OS

Specify connection options:

Public address: (IP:Port/SSL Port) ...

Connection Mode: Gateway Mode

Authentication type: Credentials

Advanced

< Back Next > Cancel Help

5. Click **Next**.

6. On the next page, specify the email recipients. Click the [...] button to select users or groups.

Invite Users - Recipients

Parallels

Specify the list of recipients:

rasusers@...com

Review the invitation e-mail:

Dear %RECIPIENT%,

You have been invited by %SENDER% to connect to Parallels Remote Application Server.

%INSTRUCTIONS%

%MANUALINSTRUCTIONS%

Thanks,

System Administrator

Preview Default

< Back Next > Cancel Help

7. Review the invitation email template displayed in the **Review the invitation e-mail** box. You can modify the template text as needed. The template also uses variables, which are explained below.

- %RECIPIENT% — Specifies the name of a recipient to whom the email message is addressed.
- %SENDER% — The sender's email address that you specified in the first step of this wizard when you configured the outgoing email server settings.
- %INSTRUCTIONS% — Includes a custom URL hyperlink for automatic configuration of Parallels Client. The URL uses the Parallels Client URL scheme. For more info, see [RAS Web Client API and Parallels Client URL Scheme](#).
- %MANUALINSTRUCTIONS% — Includes instructions for manual configuration of Parallels Client.

The variables are defined dynamically depending on the type(s) of the target devices and other settings. Normally, you should always include them in the message, so your users will receive all the necessary instructions and links. If you don't include any of the variables, you will see a warning message, but including all of them is not a requirement. To preview the message, click the **Preview** button. This will open the HTML version of the message in a separate window. This is the email message that your users will receive.

8. Click **Next**, review the summary and click **Next** again to send the invitation email to users.

When users receive the invitation email, they will follow the instructions that it contains to install and configure Parallels Client on their devices. Once that's done, the users will be able to connect to Parallels RAS and launch published resources.

* The correct setting depends on the remote mail server configuration and cannot be arbitrary chosen. Please follow this guide table:

Setting	Description
Do not use	
Use TLS / SSL	
Use TLS / SSL if available	
Use TLS 1.2 if available	
Use TLS 1.3 if available	

Azure Virtual Desktop

The **Deploy Azure Virtual Desktop** section in the **Start** category is an optional feature, which allows you to deploy Azure Virtual Desktop in Parallels RAS. The feature is described in detail in the [Azure Virtual Desktop](#) chapter.

Conclusion

In this tutorial, we have configured a simple Parallels RAS Farm with a single RD Session Host and one published application. We then configured a mailbox for

outgoing emails and sent an invitation email to end users with instructions on how to install Parallels Client, connect to the Parallels RAS Farm, and run the published application. Essentially, we've created a fully functional Parallels RAS Farm serving remote applications to end users.

If you wish, you can repeat the tutorial and add more RD Session Hosts, publish more applications, or send an invitation email to users who use different types of devices. The instructions remain essentially the same.

The rest of this guide explains in detail how to configure and use various features of Parallels RAS.

Farm and Sites

Parallels RAS Farm is a logical grouping of objects for the purpose of centralized management. A Farm configuration is stored in a single database which contains information about all objects comprising the Farm. A Site is the next level grouping in the Farm hierarchy which contains servers and other objects providing connection and remote application services.

Connecting to a Parallels RAS Farm

If you have more than one Parallels RAS Farm in your organization, you can use the same Parallels RAS Console instance to manage any of them. By default, the Parallels RAS Console is installed on the same server where you install other Parallels RAS components, but you can install the console on any computer on your network.

Connecting to a Parallels RAS Farm for the first time

When you open the Parallels RAS Console for the first time, it displays the logon dialog on which you need to specify the following:

- **Farm:** A Parallels RAS Farm to connect to. Enter the FQDN or IP address of the server where you have RAS Connection Broker installed.
- If you've installed the Parallels Single Sign-On component when installing the RAS Console, you will see the **Authentication type** field from which you can select whether to log on using your credentials or SSO. If you reboot after the

installation and select SSO, select **Single Sign-On** and then click **Connect**. Your Windows credentials will be used to log in to the RAS Farm. If you select **Credentials**, enter your credentials as described below.

- **Username:** A user account with administrative privileges on the server where Parallels RAS is installed (usually a domain or local administrator). The account name must be specified using the UPN format (e.g. administrator@domain.com). The specified user will be automatically configured as the Parallels RAS administrator with full access rights.
- **Password:** The specified user account password.
- If you select the **Remember credentials** option, this dialog will not be shown the next time you launch the Parallels RAS Console.

After entering the connection properties, click **Connect** to connect to the Farm and open the RAS Console.

Note that the **Edit Connections** button will not display any information on first connect (it is used to edit Farm connections that already exist), so you can ignore it at this point. We will talk about using this button closer to the end of this section.

Connecting to a different Parallels RAS Farm

When you need to connect to a different Parallels RAS Farm, you first need to log off from the Parallels RAS Console in order to see the logon dialog again. To do so:

1. In the Parallels RAS Console, click on the arrow icon next to your user name in the upper right-hand corner and then choose **Log Off** in the context menu.
2. The console will close and the RAS logon dialog will open. The dialog will be populated with the current Farm connection properties.
3. To connect to a different Farm, type the FQDN or IP address of the server where the other Farm is located. Once again, this should be the server where you have the RAS Connection Broker installed.
4. Specify a username and password and click **Connect**. The Parallels RAS Console will connect to the Farm using the connection properties that you specified.

Switching between Parallels RAS Farms

After you connect to more than one Farm from the same Parallels RAS Console instance, you can easily switch between them as follows:

1. In the Parallels RAS Console, click the **Location** drop-down list in the upper left-hand corner (right below the main application menu, where the current Site name is displayed).
2. The lower portion of the drop-down list will contain names of the Farms to which you connected at least once in the past (the upper portion contains one or more Site names for the current Farm). Click a desired Farm name to connect to it.
3. When you click the Farm name, the console will close momentarily and will re-open connected to the Farm that you selected.

Note that you can also switch between Farms by logging off from the console and choosing a desired Farm from the **Farm** drop-down list in the RAS logon dialog. The method described above is more convenient, so this one is just another way to do it.

Editing Parallels RAS Farm connections

As was mentioned in the beginning of this section, the RAS logon dialog has the **Edit Connections** button. When you click it, the **Manage Parallels RAS Farm Connections** dialog opens.

On the left side of the dialog, the **Farm Connections** pane lists Parallels RAS Farms to which you connected at least once in the past. If a connection is no longer relevant, you can remove it by selecting it and clicking the “minus sign” icon at the top. Once a connection is removed, it will no longer appear in the RAS logon dialog and in the Parallels RAS Console (the **Location** drop-down list).

On the right side of the dialog, the **Connection Brokers** pane lists RAS Connection Brokers for the selected Farm connection. By default, the primary Connection Broker is included in the list, but you can add more Connection Brokers if needed. When connecting to a Farm, the Parallels RAS Console will try the primary Connection Broker first. If a connection cannot be established, it will try other Connection Brokers in the order they are listed in the **Connection Brokers** pane. To add a Connection Broker to the list, click the “plus sign” icon and then specify the server FQDN or IP address.

Connecting to a Parallels RAS Farm via Command Prompt

You can connect to a Farm using Windows Command Prompt. This can be useful when you want to automatically connect to specific Farms.

To connect to a Farm using Windows Command Prompt:

1. Navigate to the folder where 2xConsole.exe is located.
2. Open Command Prompt and enter the following command:
`2xconsole.exe <arguments>`. The possible arguments are described below.

Arguments

Argument	Description
-f	The friendly name of the Farm connection.
-s	A Parallels RAS Farm to connect to. Enter the FQDN or IP address of the server where you have RAS Connection Broker installed.
-u	Username.
-p	Password.
-savecreds	When set, the credentials are saved together with the Farm information.

Examples

Create a new Farm without a friendly name and store the credentials:

```
2xconsole.exe -s 10.124.4.113 -u administrator@prls.dev -p password -savecreds
```

Connect to the previously created Farm using the IP:

```
2xconsole.exe -s 10.124.4.113
```

Create a new Farm with a friendly name and store the credentials:

```
2xconsole.exe -f myfarm -s 10.124.5.123 -u administrator@prls.dev -p password -savecreds
```

Connect to the previously created farm using the friendly name:

```
2xconsole.exe -f myfarm
```


About Sites

A Parallels RAS Farm consists of at least one Site, but may have as many sites as necessary.

Sites are often used to separate management and/or location functions. For example, by creating a Site, you can delegate permissions to a Site administrator without granting them full Farm permissions. Or you can have separate sites for different physical locations with the ability to copy the same settings to each Site while using RD Session Hosts, Providers, or PCs that are closer to end users or (depending on your needs) to back-end servers. For instance, it would make sense for a client/server application querying a database to be published from an RD Session Host which is located closer to the database server.

Each Site is completely isolated from other sites within the same Farm. The Farm simply groups sites logically and stores configuration properties of each Site (and the objects that comprise it) in a single database. Sites don't communicate with each other and don't share any objects or data. The only exception to this rule is the RAS Licensing Site which periodically communicates with other sites to obtain statistics.

Individual object settings in a given Site can be replicated to all other sites. This does not mean that settings will be shared between sites. The settings that you choose will simply be applied to other sites. For more information, see the [Replicating Site Settings](#) section.

When you install Parallels RAS, a Farm with a single Site is created automatically. This first Site becomes the RAS Licensing Site and the host for the main Parallels RAS configuration database. When you add more sites to the Farm, the data in this database is automatically synchronized with every Site that you add. When changes are applied to a particular Site, the main configuration database is automatically updated to reflect the changes.

Each Site must have at least the following components installed in order to publish remote applications and desktops for end users:

- Primary RAS Connection Broker
- RAS Secure Gateway. Note that if a Site is joined as Tenant to RAS Tenant Broker, RAS Secure Gateway is not needed. For details, see [RAS Multi-Tenant Architecture](#).

- RD Session Host, VDI, or PC

When you install Parallels RAS using default installation options, the primary RAS Connection Broker and the RAS Secure Gateway are automatically installed on the server on which you perform the installation. You can then add one or more RD Session Hosts to the Site to host published resources. You can also add more sites to the Farm if needed and configure individual components for each Site as you desire.

Sites in the RAS Console

To view existing sites in the Parallels RAS Console, select the **Farm** category in the left pane. Existing sites are listed in the right pane.

Note: The **Farm** node will only be visible to an administrator who has full permissions to manage the Farm. For more information about Farm/Site permissions, please refer to [Managing Administrator Accounts](#).

The **Farm** category displays the configuration of only one Site at a time. If you log in as the Farm administrator, the configuration of the RAS Licensing Site will be displayed. If you log in as an administrator who has access to a specific Site (but not the Farm), the configuration of that Site will be displayed.

Current Site

Click on the **Farm** item in the middle pane to view the list of available sites. The Site which configuration is currently loaded in the console is marked as “Current Site” in the **Type** column. The column also displays other Site attributes. For example, “Licensing Site / Local Site / Current Site”.

Switching between sites

To switch to a particular Site, select **Farm** in the middle pane, then right-click the Site in the right pane and choose **Switch to this Site**. The Site configuration will be loaded into the RAS Console.

The other way of switching between sites is to click the **Location** drop-down list in the upper left-hand side of the RAS Console. The menu lists sites for the current Farm and may also list other Farms if you used this RAS Console to connect to them. For more info, see [Connecting to a Parallels RAS Farm](#).

Renaming a Site

To rename a Site, right-click it and choose **Rename Site**.

Site configuration and health view

When you select the **Site** node in the middle pane, the **Site Info** tab in the right pane displays the list of Parallels RAS components that have been configured for the Site with interactive performance monitoring metrics for each component. Depending on the Site configuration, the list may include RD Sessions Hosts, VDI, Remote PCs, Secure Gateways, Connection Brokers, Azure Virtual Desktop, HALB Virtual Servers and devices, Tenant Broker, Host pools, and Enrollment Server.

To collapse or expand a component group, click an “arrow up” or “arrow down” icon on the right side of the list. Note that if no servers of a particular type have been added to the Site, the group name will not be displayed in the list.

The following information is displayed for each component (the information is updated at an interval of approximately 2 minutes):

- **Address:** Server FQDN or IP address.
- **Status:** Indicates whether the agent software is installed on the server and is functioning properly.
- **CPU:** Current CPU utilization.
- **RAM:** Current RAM utilization.
- **Disk Read Time:** Disk read time.
- **Disk Write Time:** Disk write time.
- **Sessions:** The number of currently active user sessions.
- **Preferred PA:** The name of the RAS Connection Broker designated as preferred for this server.
- **Operating System:** Operating system version installed on the server.
- **Agent Version:** The agent version installed on the server.
- **Hypervisor:** The hypervisor the server is running on.

You can customize this view by clicking **Tasks > Monitoring Settings**. This opens a dialog where you can specify which colors should be used to display different performance counters and their values.

Performing tasks on a component

You can perform a number of tasks on a component displayed in the **Site Info** tab. These tasks are described below.

To configure a component, do one of the following:

- While the **Site** node is selected in the middle pane, right-click a component in the right pane and choose **Show in the editor**.
- Select a component category in the middle pane (e.g. RD Session Hosts, Providers, etc.).

To use server management tools, right-click a component (server), click **Tools** and choose a desired tool. For the complete description of tools, see [Computer Management Tools](#).

Using the Site Designer

Select the **Site** node in the middle and then click the **Designer** tab in the right pane. The tab displays a visual representation of the Site infrastructure. Use the icons at the top to add more components to the diagram as desired. Note that adding a component to the diagram will actually add it to the Site. Double-click a component to view and configure it in a corresponding editor.

Adding a Site to the Farm

To add a Site to the Farm:

1. In the RAS Console, select the **Farm** category in the left pane and then select the Farm in the middle pane.
2. In the **Tasks** drop-down list (the right pane, above the Site list), click **Add** (or click the **+** icon).
3. In the **Add Site** dialog:
 - In the **Site** field, specify a Site name.
 - In the **Server** field, specify the IP address or FQDN of the server where the Primary Connection Broker and Secure Gateway should be installed.
 - Select the **Enable HTML5 Gateway** option to automatically create a self-signed certificate, enable SSL, and enable HTML5 support. For more info, please see [Configure User Portal](#).

4. Click **Next**.
5. The **Site Properties** dialog opens. First, it verifies if RAS Connection Broker is installed on the specified Site server. If it isn't, it will indicate this in the **Status** field.
6. Click the **Install** button to install the agent.
7. In the **Install RAS Connection Broker** dialog, highlight the server name on which the RAS Connection Broker is to be installed.
8. (Optional) Select the option **Override system credentials** to specify and use different credentials to connect to the server and install the agent.
9. Click **Install** to install the Connection Broker and Secure Gateway. Click **Done** once it has been successfully installed.

Once a new Site is created, you can view and manage its configuration by right-clicking the Site in the RAS Console and choosing **Switch to this Site**.

Replicating Site settings

Site-specific settings configured for a given Site can be replicated to all other sites in a Farm. Refer to the table below for the information about which settings can be replicated to other sites.

Category	Section	Options
Farm	VDI > Templates	Auto removal timeout of host pools that fail preparation
Farm	VDI > Desktops	Auto removal timeout
Farm	Settings > Auditing	All settings
Farm	Settings > Global Logging	Logging settings
Farm	Settings > URL Redirection	All settings
Load Balancing	Load Balancing	All settings
Load Balancing	CPU Optimization	All settings
Publishing	Application	

Category	Section	Options
		Site defaults are replicated. Other settings (name, description, icon, etc.) are global and are common to all sites
Publishing	Shortcuts	All settings
Publishing	Extensions	All settings
Publishing	Licensing	All settings
Publishing	Display	All settings
Universal Printing	Universal printing	Printer renaming
Universal Printing	Printer drivers	All settings
Universal Printing	Fonts management	All settings
Universal Scanning	WIA	Scanner renaming
Universal Scanning	TWAIN	Scanner renaming
Universal Scanning	TWAIN > TWAIN applications	Scanning applications
Connection	Authentication	All settings
Connection	Settings	All settings
Connection	Allowed devices	All settings
Reporting	Reporting engine	Reporting engine type
Reporting	Engine specific settings	All settings

To replicate Site settings to all other sites, select **Farm> <site>> Settings** and then select the **Replicate settings** option (at the bottom of the **Auditing** tab). Please note that this option is disabled if you have just one Site in the Farm.

Overriding Site Replicated Settings

If an administrator who has permissions to enable or disable replication settings makes a change to a specific setting, such setting is replicated to all other sites. If an administrator has access to a particular Site only, upon modifying Site settings which have been replicated, the replicated settings are overridden and the option **Replicate Settings** is automatically cleared, therefore such settings will no longer be replicated to other sites.

Managing Licensing Site

The Licensing Site should always be online even if you have other sites in your Farm. If your Licensing Site goes offline, your other sites can still use the maximum number of individual licenses included in your subscription but only for a period of 72 hours. During this time, you need to do one of the following:

- Restore your Licensing Site.
- Promote a different Site to be the Licensing Site in the Farm (see below for instructions).

Please note that if the Licensing Site is offline from 48 to 72 hours and back online three times per month, you will be required to re-activate it using your Parallels RAS licensing key after the third time.

To promote a secondary Site to be the Licensing Site in the Farm:

1. In the RAS Console, navigate to **Farm > Farm**.
2. In the right pane select a Site and then click **Tasks > Set as licensing Site**.
3. You will be asked to activate the new Licensing Site using your Parallels RAS license. Follow the instructions and activate the Site.

Managing administrator accounts

You can have more than one administrator in Parallels RAS. At least one administrator (called the root administrator) must be present at all times. Other administrators can be given the following roles:

- **Root administrator.** Has full permissions to manage a Parallels RAS Farm.
- **Power administrator.** Has most permissions granted by default, but can be configured to have limited permissions to manage certain sites or categories.

- **Custom administrator.** Has no permission by default and can be granted specific permission to view or modify very specific areas or objects in the Parallels RAS Farm.

Read on to learn how to create and manage administrator accounts.

Adding an administrator account

To add an administrator account to the Parallels RAS Farm:

1. In the RAS Console, navigate to **Administration > Accounts**.
2. Click the **Tasks** drop-down list and choose **Add** (or click the **[+]** icon).
3. The **Account Properties** dialog opens.
4. Click the **[...]** button next to the **Name** field. In the **Select User or Group** dialog, select a user or a group.
5. Specify an email address and mobile phone number. Both fields are optional and are disabled if the account specified in the **Name** field is a group.
6. In the **Permissions** drop-down list select a role to assign to the administrator:
 - **Root administrator.** Grants the administrator full permissions to manage the Farm.
 - **Power administrator.** Grants the administrator full permissions by default but allows you to limit them if needed. To grant or deny specific permissions, click the **Change Permissions** button. For additional info, see [Administrator Account Permissions](#).
 - **Custom administrator.** This role doesn't have any permissions by default and allows you grant very specific permissions for a particular category, area, or object in the RAS Console. See [Administrator Account Permissions](#) for details.
7. In the **Receive system notifications via** drop-down list, select **Email** to send all system notifications to the specified email address, or select **None** to disable email system notifications for this account.
8. Click **OK** to add the new administrator account to the Farm.

Modifying an administrator account

To modify an account, select it in the list and click **Tasks > Properties**. This opens the **Account Properties** dialog where you can modify the account information.

To enable or disable an account, select or clear the **Enable account** option at the top of the **Account Properties** dialog.

Administrator account permissions

To set permissions for a RAS administrator, do the following:

1. In the RAS Console, navigate to **Administration > Accounts**.
2. Select an administrator in the list and click **Tasks > Properties**.
3. Click the **Change Permissions** button in the **Administrator Properties** dialog. The following happens depending on what is selected in the **Permissions** field:
 - **Root administrator**. The **Change Permission** button is disabled because the root administrator always has full permissions.
 - **Power administrator**. The **Account Permissions** dialog opens. In the left pane, select one or more sites for which to grant permissions to the administrator. In the right pane, select specific permissions. See the **Power administrator permissions** subsection below for details.
 - **Custom administrator**. A different **Account Permissions** dialog opens where you can set custom permissions. Compared to the **Power administrator** role (see above), this option allows you to grant any permission (view, modify, add, etc.) for entire categories or specific areas or objects in the RAS Console. If a Custom administrator doesn't have permissions to even view a category or tab page, they will not even appear in the RAS Console. Using the **Custom administrator** role, you can limit permissions to one or more very specific tasks. For details, see **Custom administrator permissions** below.

Power administrator permissions

The following permissions can be set for a **Power administrator**:

- **Allow viewing of site information**. Whether the administrator can view the Site information.
- **Allow site changes**. Permissions to modify the following categories: **Site**, **Load Balancing**, **Universal Printing**, **Universal Scanning**. This option is disabled if the **Allow viewing of Site information** option is cleared.
- **Allow session management**. Permission to manage running sessions. This option is disabled if the **Allow viewing of site information** option is cleared.
- **Allow publishing changes**. Permission to modify the **Publishing** category.

- **Allow connection changes.** Permission to modify the **Connection** category.
- **Allow viewing of RAS reporting.** Permission to view reports generated by RAS Reporting.
- **Allow client management changes.** Permission to modify the **Device Manager** category.

In the **Global permission** area, set the following:

- **Allow viewing of policies.** Whether to allow the administrator to view the **Policies** category.
- **Allow policies changes.** Whether to allow the administrator to modify the **Policies** category.

Custom administrator permissions

To set custom administrator permissions, you must be either a root administrator or a power administrator with the “Allow site changes” permission granted.

When you first create an administrator of this type, they will have no permissions. To add permissions, select a Site in the left pane and then click the **Change permissions** button. The **Account Permissions** dialog opens. In the dialog, select a permission type in the left pane.

The permission types are:

- **RD Session hosts groups.** The **Groups** tab in **Farm > RD Session hosts**.

Note: Starting from Parallels RAS 19, per-server RDSH permissions have been deprecated and must be manually replaced with per-group permissions. If you upgrade to Parallels RAS 19 or later from one of the previous versions, during the upgrade you will see a dialog that helps you with the process.

- **Manage Sessions by AD Groups.** Permission for managing user sessions for users that belong to the same AD group as the custom administrator.

Note: Parallels RAS checks all available AD groups to find the ones that include the custom administrator. If you don’t want to check certain AD groups, you can exclude them from the search by clicking the **Exclude AD groups** button in the bottom-left corner of the **Account Permissions** window.

- **Remote PCs.** The **Farm > Remote PCs** view.
- **Secure Gateways.** The **Farm > Secure Gateways** view.
- **Connection Brokers.** The **Farm > Connection Brokers**.

- **HALB.** The **Farm > HALB** view.
- **Themes.** The **Farm > Themes** view.
- **Publishing.** Permissions for individual folders in the **Publishing** category.
- **Connection.** The entire **Connection** category.
- **Device Manager.** The entire **Device manager** category.
- **Certificates.** The **Farm > Certificates** view.
- **Application Packages.** The **Farm > Application Packages** view.

To change global permissions, instead of a specific Site select **Global** in the left pane and then click the **Change permissions** button.

The global permission types are:

- **Monitoring.** The **Monitoring** category.
- **Reporting.** The **Reporting** category.
- **License.** The **License** category.

After you select a permission type, you can set the actual permissions in the right pane. Different permission types may have different sets of permissions. The following list describes all available permissions:

- **View.** View only.
- **Modify.** View and modify.
- **Add.** View, modify, and add new objects (e.g. servers).
- **Delete.** View, modify, and delete an object.
- **Control.** View and control an object. This permission enables the **Tasks > Control** menu (where available), which includes enable and disable logons, cancel pending reboot, install RDS role, reboot, and some other options. Also enables power operations (start, stop, etc., where available).
- **Manage sessions.** View and manage sessions.

The lower portion of the right pane lists individual objects (e.g. servers) if the selected permission type has them. Here, you can set individual permissions for a specific object (not the entire tab, for instance, which otherwise would include all available objects).

The **Global permissions** options at the top of the right pane enables all permissions for all objects for the selected permission type.

Clone permissions

As a root administrator (or a power administrator with sufficient privileges), you can apply (clone) permissions of an existing administrator account to another existing account. This way, you can configure permissions for one account and then quickly apply the same configuration to all other accounts that require them.

To clone permissions, select a source administrator account and click **Tasks > Clone permissions**. In the dialog that opens, select a destination account (or multiple accounts) and click **OK**.

Delegate permissions

There could be a situation when a power administrator needs to grant some permissions to a custom administrator. This cannot be done by modifying permissions because power administrators cannot manage administrator accounts directly. Instead, they can delegate some of their own permissions in a given Site to a custom administrator of their choice.

For example, if a power administrator wants the custom administrator to be able to manage a particular RD Session Host, he/she selects that host in the RAS Console and click **Tasks > Delegate permissions**. This opens a dialog where the administrator can select a custom administrator and specify which permissions (view, modify, etc.) that administrator should have. The **Tasks > Delegate permissions** menu option is available for many objects, such as Providers, host pools (desktops), and some others. If the menu is not available for an object, it means that this functionality is not available for objects of this type.

Managing administrator accounts

To view existing administrator accounts, select the **Administration** category in the RAS Console. The **Accounts** tab lists existing accounts and their properties, including:

- **Group or user name.** Account name, which can be a user or group name.
- **Type.** Account type. Can be one of the following: **User**, **Group**, **Group User**. The User and Group are self-explanatory. The **Group User** is a user who receives Parallels RAS administrative permissions via a group membership. When you initially add a group to the list of Parallels RAS administrators, its members are not displayed on the **Accounts** tab. As soon as a member of

the group logs in to Parallels RAS, the account name is added to the list of administrators as a **Group User** and remains there. Note that you cannot change Parallels RAS permissions for such an account individually outside the group permissions.

- **Permissions.** A security role assigned to an administrator.
- **Email.** Email address.
- **Mobile.** Mobile phone number.
- **Group.** Group name. This column has a value for Group Users only (see the **Type** column description above).
- **Last Modification By.** The name of the user who modified this account in Parallels RAS the last time.
- **Changed On.** The last account modification date.
- **Created By.** The name of the user who created this account in Parallels RAS.
- **Created On.** The date when this account was added to Parallels RAS.
- **ID.** Internal Parallels RAS ID.

Modifying an account

To modify an account:

1. Right-click an account and choose **Properties** in the context menu.
2. Use the **Administrator Properties** dialog to modify the necessary information. For more info, see [Adding an Administrator Account](#).

Moving an account to a different Active Directory group

To move an account to a different Active Directory (AD) group:

1. In the RAS Console, navigate to **Administration > Accounts**.
2. Click the **Tasks** drop-down list and choose **Change account group**. The **Change Account Group** dialog opens.
3. Do one of the following:
 - To assign the account to another group, select the group from the **Choose new group** list. The account will receive the permissions of the new group.
 - To remove the account from the group without assigning it to another one, select **None** from the **Choose new group** list. If you want the

account to keep the permissions of the group, select the **Keep group permissions** option.

4. Click **OK**.

Handling locked objects

When an administrator is working with an object (e.g. a tab in the RD Session Host properties dialog), the object is locked for all other administrators. Therefore, upon trying to access a locked object, an administrator will be alerted with an error that the object is locked and will be denied access to it.

A root administrator (but not power or custom administrator) can release a locked object as follows:

1. On the **Administration > Accounts** tab, click the **Tasks** drop-down list and choose **Show Sessions**.
2. In the **Sessions** dialog, select the administrator who is locking an object and then click the **Send Message** icon (at the top).
3. If the administrator doesn't reply and doesn't release the object, you have an option to click **Log Off**, which will log them off and will unlock the category.

Configure RAS Console idle sessions

If you have a number of administrators using the RAS Console to manage the same Farm, you can configure when an idle RAS Console session should be disconnected. By default, when an administrator opens the console and connects to a Farm but then forgets to log off and goes away, the session will stay active indefinitely possibly locking some of the categories for other administrators. You can change that by specifying the time period after which an idle session will be disconnected (thus unlocking the categories).

To configure idle sessions:

1. In the RAS Console, navigate to **Administration > Settings**.
2. Locate the **Miscellaneous** section (at the bottom) and choose a desired time period in the **Reset idle RAS Console session after** drop-down list.

When a session stays idle for close to the specified time period, the administrator (session owner) will be notified a few minutes in advance that the session is about to be disconnected. If the administrator chooses to stay connected, the time period

is reset. If the administrator does nothing, the session will be disconnected when the time expires.

Using instant messaging

Parallels RAS administrators logged on to the same Farm can communicate with each other using a built-in instant messenger.

To use the instant messenger:

1. In the RAS Console, select the **Administration** category.
2. Expand the drop-down list next to your name (top-right corner of the console screen) and click **Chat**.
3. The **Parallels Remote Application Server Chat** window opens.

To send a message:

1. Type the message text in the lower input panel.
2. In the **Logged on administrators** list box, select a specific administrator or **All** to send the message to an individual or all logged on administrators.
3. Click **Send**.

Your message history is displayed in the **Messages** panel. To clear the history, click **Clear All**.

You can also view the chat history listing all messages between all administrators (not just your own messages). To do so, select the **Administration** node in the console and then select the **Chat History** tab.

Joining Customer Experience Program

Parallels Customer Experience Program helps us to improve the quality and reliability of Parallels RAS. If you accept to join the program, we will collect information about the way you use Parallels RAS. We will not collect any personal data, like your name, address, phone number, or keyboard input.

To join the program:

1. In the RAS Console, select the **Administration** category.

2. In the right pane, click the **Settings** tab.
3. Select the **Participate in the Customer Experience Program** option.

After you join the program, CEP will automatically start to collect information about how you use Parallels RAS. Data collected from you and other participants is combined and thoroughly analyzed to help us improve Parallels RAS.

RAS Connection Broker

RAS Connection Broker provides load balancing of published applications and desktops. A RAS Connection Broker is automatically installed on a server on which you install Parallels RAS and is designated as the primary Connection Broker. Each Site must have a primary RAS Connection Broker but can also have secondary Connection Brokers added to it. The purpose of a secondary Connection Broker is to ensure that users do not experience any interruption of the service due to possible failure of the primary RAS Connection Broker. This chapter describes how to add RAS Connection Brokers to a Site and how to configure them.

Configuring RAS Connection Brokers

To view RAS Connection Brokers installed in a Site, navigate to **Farm > <Site> > Connection Brokers** in the RAS Console. The installed Connection Brokers are listed on the **Connection Brokers** tab in the right pane.

A Site must have at least the primary Connection Broker installed, which is marked so in the **Priority** column. You can also add secondary agents to a Site for redundancy (described in the section that follows this one).

To modify the configuration of a Connection Broker, select it and then click **Tasks > Properties** (or right-click > **Properties**). The **Properties** dialog opens where you can modify the following:

- **Enable Server in site:** Enables or disables the Connection Broker. The option is enabled for secondary Connection Brokers only. It is disabled for the primary Connection Broker.
- **Server:** Specifies the FQDN or IP address of the server that hosts the Connection Broker. To automatically resolve IP address to FQDN, enable the global **Name Resolution** option. For details, see [Host Name Resolution](#).

- **IP:** Specifies the server IP address. Click the **Resolve** button to obtain the IP address automatically using the FQDN specified in the **Server** field. This IP address is used so that multiple Connection Brokers share information in real time.
- **Alternate IPs:** Specifies one or more alternate IP addresses separated by a semicolon. These addresses will be used if RAS Secure Gateways fail to connect to the RAS Connection Broker using its FQDN or the address specified in the **IP** field. This can happen, for example, if Secure Gateways are connecting from a network which is not joined to Active Directory.
- **Description:** A user-defined description.
- **Standby:** If selected, puts a secondary Connection Broker into a standby mode. This means that no agent will connect to this Connection Broker until another Connection Broker goes offline. This option is enabled automatically for any new secondary Connection Broker in excess of the three agents that already exist. It is not recommended to have more than three active Connection Brokers because it may degrade system performance. Using this option you can have more than three agents, but have them in standby mode until they are needed. For more information, see [Secondary Connection Brokers](#).

When done making the changes, click **OK** and then click **Apply** in the main RAS Console window.

The **Tasks** drop-down list on the **Connection Brokers** tab has the following items:

- **Add.** Adds a RAS Connection Broker to the Site. See the section that follows this one for the information on how to add secondary Connection Brokers.
- **Upgrade all Agents.** Upgrades agents to the current version. The item is disabled if all agents are up to date.
- **Tools.** Gives you access to a set of standard server management tools.
- **Troubleshooting.** The **Check agent** menu item verifies that the Connection Broker is functioning properly. It opens a dialog where you can see the verification results and optionally install (or uninstall) the Connection Broker. The **Logging** menu item allows you to configure logging and retrieve or clear log files. For more information, see [Logging](#).
- **Promote to primary.** Promotes a secondary Connection Broker to primary. The current primary becomes a secondary Connection Broker.
- **Refresh.** Refreshes the **Connection Brokers** list.
- **Delete.** Deletes a secondary Connection Broker from the Site. To delete the primary Connection Broker, you first need to promote a secondary Connection Broker to primary.

- **Settings audit.** Opens the **Settings Audit** dialog where you can view the changes that were done to the Connection Broker. For more information, see [Settings Audit](#).
- **Move up** and **Move down.** Changes the priority of a secondary Connection Broker (moves it up or down in the priority list).
- **Properties.** Opens the Connection Broker **Properties** dialog (see above).

RAS Connection Brokers overview

In addition to the Connection Broker editor described above, you can also see the summary about the available RAS Connection Brokers. To do so:

1. In the RAS Console, navigate to the **Farm** > <Site> .
2. The available RAS Connection Brokers are displayed in the **Connection Brokers** group on the **Site Info** tab.
3. To go to the Connection Brokers editor, right-click a RAS Connection Broker and choose **Show in the editor**.

For additional info, see [Sites in the RAS Console](#).

Automatically connect to an alternative Connection Broker

You can configure Parallels RAS to automatically connect to an alternative Connection Broker when one of Connection Brokers is not responding.

To enable automatic connection to an alternative Connection Broker:

1. In the RAS Console, click **Tools > Options** on the main menu (that's the menu at the top of the RAS Console window).
2. In the **Options** dialog, select the **Automatically connect to an alternative Connection Broker when required** option.
3. Click **OK**.

Secondary Connection Brokers

A secondary Connection Broker is added to a Site for redundancy. This way if the primary Connection Broker fails, the secondary Connection Broker is still available to handle the requests. Connection Brokers work in active/active manner to ensure

high availability. In case of a Connection Broker failure, the remaining Connection Brokers are available to handle the load.

When you have one more secondary Connection Broker installed, the runtime data is replicated on each agent, so if any service fails, the downtime is reduced to a minimum. In addition, any active Connection Broker is used for authentication purposes.

The primary Connection Broker performs the same tasks as secondary Connection Brokers but has additional responsibilities. It manages certain processes that must be managed by a single Connection Broker.

The following table lists processes managed by the primary Connection Broker and Secondary Connection Brokers:

Process	Primary Connection Broker	Secondary Connection Broker
Handles administration sessions	Yes (Primary Site Connection Broker)	no
Manages system settings	yes	no
Replicates RAS settings changes	yes	yes
Keeps RAS settings change audit	yes	optional
Monitor Performance counters of Site VMs	yes	yes
Monitor Runtime Sessions	yes	yes
Monitor Runtime Applications / Processes	yes	yes
Manages Application package states	yes	yes
Send Licensing information & heartbeat	yes (Primary site Connection Broker)	no
Monitors other Site states	yes	no
Process and send CEP information	yes	no
Send information to reporting server	yes	no
Manages Schedulers	yes	no
Device Manager Shadowing	yes Primary site Connection Broker)	no

Process	Primary Connection Broker	Secondary Connection Broker
Send System alert email notifications	yes	yes
TOTP - registration	yes	yes
Email OTP - registration	yes	yes
Manages templates versioning and distribution	yes	yes
Clone hosts from templates	yes	yes
Host pool auto scaling	yes Primary site Connection Broker)	no
Manages Web authentication	yes	yes
Manages Credentials authentication	yes	yes
Manages Smart card authentication	yes	yes
Validates Policies	yes	yes
Provides configuration to HALB devices	yes	yes
Distributes configuration changes to other roles & agents	yes	yes

Due to the additional load on the Primary Connection Broker, agent connection load will not be distributed equally. The following system is used:

Number of CBs	Primary CB agent connections	Secondary CB agent connections
1	100%	
2	40%	60%
3	20%	40% each
4	10%	30% each
5 or more	0%	100 / number of secondary PAs

However, if you specify that an agent is managed by a specific Connection Broker, this will be preferred.

Planning for secondary Connection Brokers

RAS Connection Brokers running on the same Site communicate with each other and share the load. The amount of data being transmitted from one agent to another is quite large, so a reliable high-speed communication channel must be ensured (e.g. a subnetwork can be configured for Connection Broker communications).

When adding a secondary Connection Broker to a Site, you specify an IP address for it. Make sure that the IP addresses of all agents belong to the same network segment. The port that Connection Brokers use to communicate with each other is TCP 20030.

In general, the N+1 redundancy approach should be used per Site. Note that for auto-promotion you should have at least three Connection Brokers (auto-promotion is described later in this section). There's no physical limit to how many Connection Brokers you can add to a Site. However, the best results are achieved with only two-three agents present. The three-agent scenario is highly recommended, especially when you have Providers and want to enable [high availability for VDI](#).

Adding more than two secondary Connection Brokers to a Site may have an adverse effect and actually degrade the system performance due to the data replication taking place. Note that this does not apply to secondary Connection Brokers in standby mode, which is explained in [Configuring RAS Connection Brokers](#).

Adding a secondary RAS Connection Broker to a Site

To add a secondary Connection Broker:

1. In the RAS console, navigate to **Farm** > <Site> > **Connection Brokers**.
2. Click the **Tasks** drop-down list and choose **Add** to launch the **Add RAS Connection Broker** wizard.
3. The **Server** field specifies the FQDN or IP address of the server that hosts the RAS Connection Broker. To automatically resolve IP address to FQDN, enable the global **Name Resolution** option. For details, see [Host Name Resolution](#).
4. The **IP** field specifies the server IP address. Click the **Resolve** button to obtain the IP address automatically using the FQDN specified in the **Server** field.

5. The **Alternative IPs** field specifies one or more alternative IP addresses, separated by a semicolon. These addresses will be used if RAS Secure Gateways fail to connect to the RAS Connection Broker using its FQDN or the address specified in the **IP** field. This can happen, for example, if Secure Gateways are connecting from a different network, which is not joined to Active Directory.
6. Select the **Install a Secure Gateway with a Connection Broker** option if you also want to install a RAS Secure Gateway on the specified server. If you select this option, you may also select the **Enable HTML5 Gateway** option (for more info, see [Configure User Portal](#)).
7. Select the **Add Firewall Rules** option to automatically configure the firewall on the server. See **Port Reference** for details.
8. Click **Next**.
9. On the next page, click **Install** to install the RAS Connection Broker on the server. The **Installing RAS Redundancy Service** dialog opens.
10. Select the server on which the RAS Connection Broker is to be installed and click **Install**.
11. Click **Done**.
12. Click **OK** to add the server to the Farm.

Managing Secondary Connection Brokers

Enabling or disabling a secondary Connection Broker

To enable or disable a secondary Connection Broker in a Site, select it in the **Connection Brokers** list and then select or clear the check box at the beginning of the row.

Changing the secondary Connection Broker priority

Each secondary Connection Broker is given a priority. To change the priority, select a secondary Connection Broker and use the “Up arrow” and “Down arrow” icons

(or **Tasks > Move up**, **Move down**) to move it up or down the list. The higher the agent is in the list, the higher the priority.

Promoting a secondary Connection Broker to primary

If the primary Connection Broker cannot be recovered, you can promote a secondary Connection Broker to primary as follows:

1. Open the RAS Console on the Connection Broker server that you would like to promote (all required files are automatically installed when a server is added to a Site as a secondary Connection Broker).
2. Select the **Farm** category and navigate to the **Connection Brokers** node.
3. Select the Connection Broker and then click **Tasks > Promote to primary**.
4. Click **OK** once the process is finished.

Configuring auto-promotion

If the primary Connection Broker goes offline, you will need to promote a secondary Connection Broker to take its place. The auto-promotion feature can do it automatically after a specified time period.

By default, auto-promotion is turned off. To enable it, do the following:

1. In the RAS Console, navigate to **Farm > <Site> > Connection Brokers**.
2. Select the **Auto-promotion** tab in the right pane.
3. Select the **Enable auto-promotion** option and specify the time period after which the next secondary Connection Broker should be promoted to primary. The time period can be set between 15 minutes and 72 hours (the default value is 30 min).
4. Select the **Enable fallback** option if you want the original Connection Broker to become primary again should it go back online. For the Licensing Site, this eliminates license activation if fallback happens within 72 hours. The license activation countdown is always displayed in the RAS Console, so the administrator can check if the original primary Connection Broker recovers within this time period or not. If the original agent goes back online after the 72-hour period (and if the Farm has been already reactivated), it will become a secondary Connection Broker.

Note: To enable auto-promotion, you need at least three active Connection Brokers in a Site. If you have less than three, the auto-promotion is ignored.

Please also note that auto-promotion must be disabled if you have a single Site with Connection Brokers split across different locations with bad WAN links. If there's no link between Connection Broker located remotely, the third Connection Broker acts as a witness to prevent split-brain.

When auto-promotion takes place, the RAS administrator will receive notifications via email about the following events:

- A secondary Connection Broker has been promoted to primary.
- Auto-promotion of a secondary Connection Broker has failed.
- Auto-promotion failback completed.

Deleting a secondary Connection Broker

To delete a secondary Connection Broker, select it in the list and then click **Delete** in the **Tasks** drop-down list.

Using computer management tools

You can perform standard computer management tasks on a server hosting the RAS Connection Broker right from the RAS Console. These include Remote Desktop Connection, remote PowerShell, Computer Management, Service Management, Event Viewer, IPconfig, Reboot, and others. To access the **Tools** menu, select a server, click **Tasks > Tools** and choose a desired tool. For requirements and usage information, see [Computer Management Tools](#).

Enabling Connection Broker API

It is possible to manage sessions in Parallels RAS without using the RAS Connection Brokers to carry out the management. This is useful for organisations that wish to manage sessions via a third-party method. The steps to enable this are below.

In addition to the steps below, configuration via API is needed. This API is documented [here](#).

Navigate to **Farm > Site > Connection Brokers** and follow the instructions below:

1. Choose **Tasks > Site Defaults** (or choose the properties of an individual Connection Broker to assign to a specific Connection Broker)
2. Choose the **API** tab
3. Choose **Enable API** (if configuring on an individual Connection Broker, you will need to deselect **Inherit Default Settings**)

SSL/TLS

1. Select the **API Port** option and specify a port number for secure API communication
2. In the **Accepted SSL Versions** drop-down list, select the SSL/TLS version required.
3. In the **Cipher Strength** field, select a desired cipher strength.
4. If Custom is chosen in the **Cipher Strength** field, in the **Cipher** field, specify the cipher. A stronger cipher allows for stronger encryption, which increases the effort needed to break it.
5. The **Use ciphers according to server preference** option is ON by default. You can use client preferences by disabling this option.
6. In the **Certificates** drop-down list, select a desired certificate. For the information on how to create a new certificate and make it appear in this list, see the [SSL Certificate Management](#) chapter.

API Token

To use an API token, set the HTTP Authorization header value set when making API requests.

When the **Enable API** option is selected, an **API Token String** is generated. Choosing **Regenerate** will issue a new API Token String.

Encryption Certificate

The chosen encryption certificate will be used by the Connection Broker API when creating runtime data between backend components.

To set this, navigate to **Farm > Site > Connection Brokers** and select the **API** tab (if this tab is not visible, it indicates that **Enable API** has not been selected).

If **Automatic** is selected, the default method of using the site certificate will be used. The list shows certificates available with the `Third-Party Encryption` property selected.

For the information on how to create a new certificate and make it appear in this list, see the [SSL Certificate Management](#) chapter.

RAS Secure Gateway

RAS Secure Gateway tunnels all Parallels RAS data on a single port. It also provides secure connections and is the user connection point to Parallels RAS.

At least one RAS Secure Gateway must be installed and configured in every Site. Note that if a Site is joined as Tenant to RAS Tenant Broker, RAS Secure Gateway is not needed. For details, see **RAS Multi-**[Tenant Architecture](#).

Multiple gateways can exist depending on your requirements. Read this chapter to learn how to add, configure, and manage RAS Secure Gateways.

Overview

You need to install at least one RAS Secure Gateway for Parallels RAS to work. You can add additional Gateways to a RAS Site to support more users, load-balance connections, and provide redundancy.

Installing a RAS Secure Gateway on a dedicated server

If you are installing a RAS Secure Gateway on a dedicated server, you can also install the Parallels RAS console on the same server. The console will have limited functionality but will allow you to perform some important management operations on the Gateway, including:

- Setting the Gateway operation mode (normal or forwarding, see below for details).
- Assigning a RAS Connection Broker that will manage the Gateway.
- Setting the Gateway communication port.
- Viewing the Gateway information, such as host OS version, Parallels RAS version, available IP addresses, and other.

The RAS Console in such an installation scenario (when connected to the local computer, not the RAS Farm) will only have two categories that you can select in the left pane: **Gateway** and **Information**. To manage the Gateway settings, select **Gateway** and then click **Change Ownership** in the right pane. To view the information select the **Information** category.

When the RAS console is connected to a Parallels RAS Farm (i.e. the server where RAS Connection Broker is running), you can manage RAS Secure Gateways by navigating to **Farm > <Site> > Secure Gateways**.

How a RAS Secure Gateway works

The following describes how a RAS Secure Gateway handles user connection requests:

1. A RAS Secure Gateway receives a user connection request.
2. It then forwards the request to the RAS Connection Broker with which it's registered (the Preferred Connection Broker setting by default).
3. The RAS Connection Broker performs load balancing checks and the Active Directory security lookup to obtain security permissions.
4. If the user requesting a published resource has sufficient rights, the RAS Connection Broker sends a response to the gateway which includes details about the RD Session Host the user can connect to.
5. Depending on the connection mode, the client either connects through the gateway or disconnects from it and then connects directly to the RD Session Host server.

RAS Secure Gateway operation modes

RAS Secure Gateway can operate in one of the following modes:

- **Normal Mode.** A RAS Secure Gateway in normal mode receives user connection requests and checks with the RAS Connection Broker if the user making the request is allowed access. Gateways operating in this mode can support a larger number of requests and can be used to improve redundancy.
- **Forwarding Mode.** A RAS Secure Gateway in forwarding mode forwards user connection requests to a preconfigured gateway. Gateways in forwarding mode are useful if cascading firewalls are in use, to separate WAN connections from LAN connections and make it possible to disconnect WAN segments in the event of issues without disrupting the LAN.

Note: To configure the forwarding mode, a Parallels RAS Farm must have more than one RAS Secure Gateway.

Planning for high availability

When adding RAS Secure Gateways to a Site, the N+1 redundancy should be configured to ensure uninterrupted service to your users. This is a general rule that also applies to other Parallels RAS components, such as Connection Brokers or RD Sessions Hosts.

Performance improvements in version 21.0

Starting with version 21.0, RAS Secure Gateways have been enhanced with multi-threaded processing. This improvement significantly increases scalability by allowing the Gateway to take advantage of multiple CPU cores. As a result, adding additional CPUs to a Gateway server will directly improve its performance. In many cases, this eliminates the need to deploy additional Gateway servers solely to handle increased user load, simplifying infrastructure planning and reducing management overhead.

Baseline specifications

For planning purposes, a single RAS Secure Gateway configured with 4 vCPUs and 8 GB of RAM can support up to 1,200 concurrent users. These specifications provide a balanced starting point for most environments. Depending on workload characteristics, security requirements, and network conditions, additional resources may be allocated to further improve performance.

Adding a RAS Secure Gateway

To add a RAS Secure Gateway to a Site, follow these steps:

1. In the RAS Console, navigate to **Farm** > <Site> > **Secure Gateways**.
2. With the **Secure Gateways** tab selected in the right pane, click **Tasks** > **Add** to start the **Add RAS Secure Gateway** wizard.
3. Enter the server FQDN or IP address (or click the [...] button to select a server from the list). To automatically resolve IP address to FQDN, enable the global **Name Resolution** option. For details, see [Host Name Resolution](#).
4. Select the gateway mode from the **Mode** drop-down list.

5. If you selected the **Forwarding** mode in the step above, select the destination gateway in the **Forward To** drop-down list. You can also select a specific IP address in the **On IP** drop-down list if the Gateway server has more than one.
6. Select the **Enable HTML5 Gateway** option to automatically create a self-signed certificate, enable SSL, and enable HTML5 support. For more info, please see [Configure User Portal](#).
7. Select the **Add Firewall Rules** to automatically configure the firewall on the server hosting the gateway. See **Port Reference** for details.
8. Click **Next**.
9. On the next page, click **Install** to start the RAS Secure Gateway installation.
10. Click **Done** when the installation is finished.

Manually adding a RAS Secure Gateway

To manually install a RAS Secure Gateway and add it to the Farm, follow these steps:

1. Log into the server where you'll be installing the RAS Secure Gateway using an administrator account.
2. Copy the Parallels RAS installation file (`RASInstaller.msi`) to the server and double-click it to launch the installation wizard.
3. Follow the onscreen instruction and proceed to the installation type page. Select **Custom** and click **Next**.
4. Click on **RAS Secure Gateway** in the feature tree and select **Entire Feature will be installed on local hard drive**.
5. Ensure that all other components in the selection tree are cleared and click **Next**.
6. Click **Install** to start the installation.
7. When the installation is completed, click **Finish** to close the wizard.
8. Open the RAS Console and specify the RAS Connection Broker that will manage the gateway.

Checking the RAS Secure Gateway status

To check the status of a RAS Secure Gateway, right-click it in the list and then click **Check Status** in the context menu. The **RAS Secure Gateway Information** dialog opens.

The dialog displays the gateway information, including:

- **Server:** The name of the server on which the gateway is installed.
- **Gateway:** The gateway verification status (e.g. Verified).
- **Version:** The gateway software version number. The version number must match the Parallels RAS version number.
- **OS Type:** Operating system type and version.
- **Status:** Display the current RAS Secure Gateway status. If the status indicates a problem (e.g. the gateway did not reply or the gateway software version is wrong), click the **Install** button to push install the gateway software on the server. Wait for the installation to complete and check the status again.

Configuring a RAS Secure Gateway

To configure a RAS Secure Gateway:

1. In the RAS console, navigate to **Farm** > <Site> > **Secure Gateways**.
2. In the right pane, right-click a Secure Gateway and click **Properties**.
3. The **RAS Secure Gateway Properties** dialog opens.

Read on to learn how to configure the RAS Secure Gateway properties.

Enable or disable a Secure Gateway

A RAS Secure Gateway is enabled by default. To enable or disable a Secure Gateway, open the **RAS Secure Gateway Properties** dialog and select or clear the **Enable RAS Secure Gateway in site** option on the **General** tab.

Set public address

The **Public address** field on the **General** tab specifies a public FQDN or IP address of the Secure Gateway. This setting is used by the Preferred routing functionality for redirecting a client connection. Please see [Configuring preferred routing](#).

Set IP addresses for client connections

- **Use IP version:** Select the IP version(s) to use.
- **IP(s):** Specify one or more IP addresses separated by a semicolon, or click **Resolve** to resolve the IP address automatically. These are the available addresses on the Secure Gateway server. To specify IP addresses that should be used for client connections, use the **Bind to IP** section (see below).
- **Bind to IP:** Use this section to specify on which IP address (or addresses) the Secure Gateway will listen for client connections. You can select a specific address or **<All available addresses>**, in which case all of the IP addresses specified in the **IP(s)** field will be used.
- **Remove system buffers for:** These fields (one for each IP version) can be used when the connection between the Secure Gateway and the Parallels Client has a high latency (such as the Internet). This option will optimize traffic for better experience on the Parallels Client side. You can select a specific address, all available addresses, or none. What this option will do is delay the internal socket to match the performance of the external socket. If the internal network is fast and the external is slow, RDP detects the fast internal socket and sends a lot of data. The problem is that this data cannot be sent fast enough from the Secure Gateway to the Client, thus ending up with a bad user experience. Enabling this option will optimize the data exchange.

You can specify the following IP options:

IP addresses for incoming client connections for a Secure Gateway are specified on the **General** tab of the **RAS Secure Gateway Properties** dialog. RAS Secure Gateway recognizes both IPv4 and IPv6. By default, IPv4 is used.

Site defaults (Secure Gateways)

RAS Secure Gateway **Properties** dialog consists of tabs, each containing their own specific set of options. All tabs, except **Properties**, have one common option **Inherit default settings**. When you select this option, all fields on a tab are grayed out and the settings are inherited from Site defaults. To view (and modify if necessary) Site default properties for Secure Gateways, click the **Site Defaults** link, which is available on all tabs mentioned above. The link opens the **Site default properties** dialog. You can also open this dialog by clicking **Tasks > Site defaults** while on the **Farm > Site > Secure Gateways** tab.

The subsequent sections describe individual tabs and available options in the Secure Gateway **Properties** dialog.

Gateway mode and forwarding settings

A RAS Secure Gateway can operate in [normal and forwarding modes](#). To set the desired mode and configure related settings click the **Mode** tab in the **RAS Secure Gateway Properties** dialog.

Using Site defaults

To use Site default settings, click the **Inherit default settings** option. To specify your own settings, clear the option. For more info, see [Site Defaults \(Gateways\)](#).

Setting the normal mode

To set the normal mode, in the **Gateway mode** drop-down list, select **Normal**.

The **Forward requests to HTTP Server** option allows you to forward requests that do not belong to RAS Secure Gateways (gateways handle HTML5 traffic, Wyse, and URL scheme). To specify multiple servers, separate them with a semicolon. An HTTP server can be specified using an IPv6 address if necessary. Please note that the HTTP server must support the same IP version as the browser making the request.

The **Preferred Connection Broker** drop-down list allows you to specify a RAS Connection Broker that the Secure Gateway should connect to. This is helpful when Site components are installed in multiple physical locations communicating through WAN. You can decrease network traffic by specifying a more appropriate Connection Broker. For the Secure Gateway to select a Connection Broker automatically, select the **Automatic** option.

Setting the forwarding mode

To configure the forwarding mode, in the **Gateway mode** drop-down list, select **Forwarding**.

Specify (or select) one or more forwarding Secure Gateways in the **Forwarding RAS Secure Gateway(s)** field.

Note: The forwarding mode allows you to forward data to a Secure Gateway listening on IPv6. It is recommended that forwarding Secure Gateways are configured to use the same IP version.

Gateway network options

The **Network** tab is used to configure RAS Secure Gateway network options.

Using Site defaults

To use Site default settings, click the **Inherit default settings** option. To specify your own settings, clear the option. For more info, see [Site Defaults \(Gateways\)](#).

Configuring network

By default RAS Secure Gateway listens on TCP ports 80 and 443 to tunnel all Parallels RAS traffic. To change the port, specify a new port in the **RAS Secure Gateway Port** input field.

RDP port 3389 is used for clients that require basic load balanced desktop sessions. Connections on this port do not support published resources. To change the RDP port on a gateway select the **RDP Port** option and specify a new port. When setting your own port, please make sure that the port number does not conflict with the standard “RD Session Host Port” setting.

Note: If RDP port is changed, the users need to append the port number to their connection string in the remote desktop client (e.g. [ip address]:[port]).

Broadcast RAS Secure Gateway Address. This option can be used to switch on the broadcasting of the Secure Gateway address, so Parallels Clients can automatically find their primary Secure Gateway. The option is enabled by default.

Enable RDP UDP Data Tunneling. To enable UDP tunneling on Windows devices, select this option (default). To disable UDP tunneling, clear the option.

Device Manager Port. Select this option to enable management of Windows devices from the **Device Manager** category. The option is enabled by default.

Enable RDP DOS Attack Filter. When selected, this option denies chains of uncompleted sessions from the same IP address. For example, if a Parallels Client initiates multiple successive sessions with each session waiting for the user to provide credentials, Parallels RAS will deny further attempts. The option is enabled by default.

SSL/TLS encryption

The traffic between Parallels RAS users and a RAS Secure Gateway can be encrypted. The **SSL/TLS** tab allows you to configure data encryption options.

Using Site defaults

To use the Site's default settings, click the **Inherit default settings** option. To specify your own settings, clear the option. For more info, see [Site defaults \(Gateways\)](#).

Enforcing HSTS

The **Configure** button in the HSTS section allows you to enforce HTTP Strict Transport Security (HSTS), which is a mechanism that makes a web browser to communicate with the web server using only secure HTTPS connections. When HSTS is enforced for a RAS Secure Gateway, all web requests to it will be forced to use HTTPS. This specifically affects the [RAS User Portal](#), which typically accepts only HTTPS requests for security reasons.

When you click the **Configure** button, the **HSTS Settings** dialog opens, where you can specify the following:

- **Enforce HTTP strict transport security (HSTS):** Enables or disables HSTS for the Secure Gateway.
- **Max-age:** Specifies the max-age for HSTS, which is the time (in our case, in months) that the web browser should remember that it can only communicate with the Secure Gateway using HTTPS. The default (and recommended) value is 12 months. Acceptable values are 4 to 120 months.
- **Include subdomains:** Specifies whether to include subdomains (if you have them).
- **Preload:** Enables or disables HSTS preloading. This is a mechanism whereby a list of hosts that wish to enforce the use of SSL/TLS on their Site is hardcoded into a web browser. The list is compiled by Google and is used by Chrome, Firefox, Safari, Internet Explorer 11, and Edge browsers. When HSTS preload is used, a web browser will not even try to send a request using HTTP, but will use HTTPS every time. Please also read the important note below.

Note: To use HSTS preload, you have to submit your domain name for inclusion in Chrome's HSTS preload list. Your domain will be hardcoded into all web browser that use the list. **Important:** Inclusion in the preload list cannot easily be undone. You should only request inclusion if you are sure that you can support HTTPS for your entire Site and all its subdomains in the long term (usually 1-2 years).

Please also note the following requirements:

- Your website must have a valid SSL certificate. See [SSL server configuration](#).
- All subdomains (if any) must be covered in your SSL Certificate. Consider ordering a Wildcard Certificate.

Configuring SSL

By default, a self-signed certificate is assigned to a RAS Secure Gateway when the gateway is installed. Each RAS Secure Gateway must have a certificate assigned, and the certificate should be added to Trusted Root Authorities on the client side to avoid security warnings.

SSL certificates are created on the Site level using the **Farm > Site > Certificates** subcategory in the RAS Console. Once a certificate is created, it can be assigned

to a RAS Secure Gateway. For information about creating and managing certificates, refer to the [SSL Certificate Management](#) chapter.

To configure SSL for a Secure Gateway:

1. Select the **Enable SSL on Port** option and specify a port number (default is 443).
2. In the **Accepted SSL Versions** drop-down list, select the SSL version accepted by the RAS Secure Gateway.
3. In the **Cipher Strength** field, select a desired cipher strength.
4. In the **Cipher** field, specify the cipher. A stronger cipher allows for stronger encryption, which increases the effort needed to break it.
5. The **Use ciphers according to server preference** option is ON by default. You can use client preferences by disabling this option.
6. In the **Certificates** drop-down list, select a desired certificate. For the information on how to create a new certificate and make it appear in this list, see the [SSL Certificate Management](#) chapter.

The **<All matching usage>** option will use any certificate configured to be used by Secure Gateways. When you create a certificate, you specify the “Usage” property where you can select “Gateway”, “HALB”, or both. If this property has the “Gateway” option selected, it can be used with a Secure Gateway. Please note that if you select this option, but not a single certificate matching it exists, you will see a warning and will have to create a certificate first.

Encrypting Parallels Client connection

By default, the connection from the client to the Secure Gateway is encrypted and is set to **Gateway SSL Mode**. In this mode, the certificate assigned to the Secure Gateway is used to encrypt client communication, and the Secure Gateway is used to communicate through to the backend resources. To alter this in Parallels Client, open **Connection Properties** and set the connection mode accordingly.

To simplify the Secure Gateway and Parallels Client configuration, it is recommended to use a certificate issued by a well-known third-party Trusted Certificate Authority. Note that some web browsers (Chrome, Edge, etc.) use the operating system certificate store when connecting to the RAS User Portal.

Trusted Certificate Authorities

The Parallels Client for macOS relies on the system's trusted certificate authorities (CAs). In contrast, the Windows version of the Parallels Client maintains its own certificate store, which contains the list of trusted CAs used for validating server certificates.

If you want to add a custom root certificate (a certificate issued by Enterprise CA), Parallels Clients should be configured as follows:

1. Export the certificate in Base-64 encoded X.509 (.CER) format.
2. Open the exported certificate with a text editor, such as Notepad or WordPad, and copy the contents to the clipboard.
3. Follow one of the instructions below depending on your operating system.

Adding a certificate in Parallels Client for Windows

To add the certificate in Parallels Client for Windows:

1. Identify the location of your `trusted.pem` file. On most systems, it should be inside the installation directory. This file contains certificates of common trusted authorities.
2. Open the file named `customtrusted.pem` in the directory where `trusted.pem` is located. If this file does not exist, create it. This file will store certificates of custom trusted authorities.
3. Paste the content of the exported certificate into the file. If the file exists and contains other certificates, add the certificate to the end of the file.

Adding a certificate in Parallels Client for Linux

To add the certificate in Parallels Client for Linux:

- On IGEL OS 12:
 1. Open the file named `customtrusted.pem` in the directory `$HOME/.config/2X/Client/`. If this file does not exist, create it. This file will store certificates of custom trusted authorities.
 2. Paste the content of the exported certificate into the file. If the file exists and contains other certificates, add the certificate to the end of the file.
- On other Linux distributions:
 1. Identify the location of your `trusted.pem` file. On most systems, this should be the `<installation directory>/share/` directory. This file contains certificates of common trusted authorities.

2. Open the file named `customtrusted.pem` in the directory where `trusted.pem` is located. If this file does not exist, create it. This file will store certificates of custom trusted authorities.
3. Paste the content of the exported certificate into the file. If the file exists and contains other certificates, add the certificate to the end of the file.

Once the `customtrusted.pem` file is created, the Parallels Client will verify that the root CA is present in either of the two files.

Note: When the Parallels Client is upgraded, the contents of `trusted.pem` are replaced with those from the new version. The `customtrusted.pem` file, however, remains unchanged.

Securing RDP-UDP connections

A Parallels Client normally communicates with a RAS Secure Gateway over a TCP connection. Recent Windows clients may also utilize a UDP connection to improve WAN performance. To provide the SSL protection for UDP connections, DTLS must be used.

To use DTLS on a RAS Secure Gateway:

1. On the **SSL/TLS** tab, make sure that the **Enable SSL on Port** option is selected.
2. On the [Networktab](#), make sure that the **Enable RDP UDP Data Tunneling** option is selected.

The Parallels Clients must be configured to use the **Gateway SSL Mode**. This option can be set in the **Connections Settings > Connection Mode** drop-down list on the client side.

Once the above options are correctly set, both TCP and UDP connections will be tunneled over SSL.

SSL server configuration

When configuring RAS Secure Gateway to use SSL encryption, you should pay attention to how the SSL server is configured to avoid possible traps and security issues. Specifically, the following SSL components should be rated to determine how good the configuration is:

- The certificate, which should be valid and trusted.

- The protocol, key exchange, and cipher should be supported.

The assessment may not be easy to perform without specific knowledge about SSL. That's why we suggest that you use the SSL Server Test available from Qualys SSL Labs. This is a free online service that performs an analysis of the configuration of an SSL web server on the public Internet. To perform the test on a RAS Secure Gateway, you may need to temporarily move it to the public Internet.

The test is available at the following URL: <https://www.ssllabs.com/ssltest/>.

You can read a paper from Qualys SSL Labs describing the methodology used in the assessment at the following URL: <https://github.com/ssllabs/research/wiki/SSL-Server-Rating-Guide>.

Configure User Portal

User Portal is a functionality built into RAS Secure Gateway that allows users to connect to Parallels RAS and open published resources from a web browser using the Parallels Web Client. The client works similarly to a platform-specific Parallels Client, but does not require any additional software to be installed on users' computers or devices. All that users need is an HTML5-enabled web browser.

This section describes how to configure User Portal in the Parallels RAS Console. For the information about how to use it, please refer to the [Parallels Web Client and User Portal](#) chapter.

Note: To use Web Client and User Portal, SSL must be enabled on a RAS Secure Gateway. When enabling the client, please verify that SSL is enabled on the **SLL/TLS** tab or on your network load balancer. Please also note that the **User Portal** tab is only available if the gateway mode is set to "Normal". For more information, see [Gateway mode and forwarding settings](#).

To configure User Portal, click the **User Portal** tab in the RAS Secure Gateway properties dialog and then set the options described in the subsequent sections.

For the information on how to configure the Web Client URL and how to access the client from a web browser, please [Web request load balancing](#).

Using Site defaults

To use Site default settings on the **User Portal** tab, click the **Inherit default settings** option. To specify your own settings, clear the option. For more info, see [Site defaults \(Gateways\)](#).

Enable or disable User Portal

To enable or disable User Portal, select or clear the **Enable User Portal** option. This disables User Portal, so users will no be able to connect to User Portal using the Web Client.

Client settings

The **Client** section allows you to specify application launch methods and other Web Client settings.

- **Launch sessions using:** When a user tries to open a resource from the User Portal web page, the resource can open right in the web browser or it can be launched in a platform-specific Parallels Client installed on the user's computer (e.g., Parallels Client for Windows). This option specifies which client will be used. Compared to Web Client, platform-specific Parallels Client includes a richer set of features and provides end users with a better overall user experience. Select one of the following:
 1. **Browser Only:** Users can run remote applications and desktops using Parallels Web Client only. Use this option if you don't want your users to install a platform-specific Parallels Client.
 2. **Parallels Client Only:** Users can run remote applications and desktops in Parallels Client only. When a user connects to Parallels RAS using Parallels Web Client, they will be asked to install the platform-specific Parallels Client before they can launch remote applications and desktops. A message will be displayed to the user with a link for downloading the Parallels Client installer. After the user installs Parallels Client, they can still select a remote application or desktop in Parallels Web Client but it will open in Parallels Client instead.
 3. **Parallels Client with fallback to Browser:** Both Parallels Client and a browser (HTML5) can be used to launch remote applications and desktops. Parallels Client will be the primary method; Parallels Web

Client will be used as a backup method if a published resource cannot be launched in Parallels Client for any reason. A user will be informed if a resource couldn't be opened in Parallels Client and will be given a choice to open it in the browser instead.

- **(Parallels Client with fallback to Browser and the Parallels Client only)**

Additionally, you can configure Parallels Client detection by clicking on the **Configure** button:

- **Detect client:** Select when Parallels RAS tries to detect platform-specific Parallels Client.
 - **Automatically on sign in:** Parallels RAS tries to detect platform-specific Parallels Client immediately.
 - **Manually on user prompt:** Parallels RAS shows users a prompt where can they select whether they want to detect platform-specific Parallels Client .
- **Client detection timeout:** Time period during which Parallels RAS tries to detect platform-specific Parallels Client.
- **Allow users to select a launch method:** If selected, users will be able to choose whether to open remote applications in a browser or in Parallels Client. You can enable this option only if the **Launch session using** option (above) is set to **Parallels Client with fallback to Browser** (i.e. both methods are allowed).
- **Allow opening applications in a new tab:** If selected, users will be able to open remote applications in a new tab in a web browser.
- **Use Pre Windows 2000 login format:** Enables legacy (pre-Windows 2000) login format.
- **Allow embedding of User Portal into other web pages:** If selected, the User Portal web page can be embedded in other web pages. Please note that this may be a potential security risk due to a practice known as clickjacking.
- **Allow file transfer command:** Enables file transfer in a remote session. To enable file transfer, select this option and click the **Configure** button. In the dialog that opens, select **Client to server only** (transfer files from client to server only), **Server to client only** (transfer files from server to client only), **Bidirectional** (transfer files in both directions). For more information, see [Configuring Remote File Transfer](#).
- **Allow clipboard command:** Enables clipboard operations (copy/paste) in a remote session. To enable the clipboard, select this option and click the **Configure** button. In the dialog that opens, select **Client to server only** (copy/paste from client to server only), **Server to client only** (copy and paste from server to client only), **Bidirectional** (copy and paste in both directions).

For more information about using the clipboard, see [Using the Remote Clipboard](#).

- **Allow cross-origin resource sharing:** Enables cross-origin resource sharing (CORS). To enable CORS, select this option and click the **Configure** button. In the dialog that opens, specify one or more domains for which access to resources should be allowed. If you don't specify any domains, the option will be automatically disabled. In the **Browser cache time** field, specify for how long the end-user's browser will cache a resource.
- **Use a client IP detection service:** If selected, allows configuring an IP detection service to report IP addresses of connected Parallels Web Client applications. To enable a client IP detection service, select this option and click the **Configure** button. In the dialog that opens, provide the URL to the IP detection service you want to use. You can press the **Test** button to ensure the API works as expected. When you click the **Test** button, the Connection Broker will take the role of the client and call the API. If successful, you will be presented with a window showing the IP address of the Connection Broker.

Network load balancers access

The **Network Load Balancers access** section is intended for deployment scenarios where third-party front-end load balancers such as Amazon Web Services (AWS) Elastic Load Balancers (ELBs) are used. It allows you to configure an alternate hostname and port number to be used by the Network Load Balancer (NLB). This is needed to separate hostnames and ports on which TCP and HTTPS communications are carried out because AWS load balancers don't support both specific protocols over the same port.

The following options are available:

- **Use alternate hostname:** Select this option and specify an alternate hostname. When the alternate hostname is enabled, all platform-specific Parallels Clients will use this hostname to connect to the RAS Farm or Site.
- **Use alternate port:** Select this option and specify an alternate port number. The port must not be used by any other component in the RAS Farm or Site. To reset the port number to the default value, click **Default**. When the alternate port is enabled, all platform-specific Parallels Clients will use this port to connect to the RAS Farm or Site. Note that RDP sessions in Web Client will still be connecting to the standard SSL port (443).

Note: Please note that using an alternate host or port is not suitable in a multi-tenant environment as Tenant Broker RAS Secure Gateways are shared between Tenants, which would require different configurations.

In addition, the AWS Application Load Balancer (ALB), which handles HTTP/s traffic required by the Parallels Web Client, only supports specific cookies that are usually automatically generated. When a load balancer first receives a request from a client, it routes the request to a target and generates a cookie named AWSALB, which encodes information about the selected target. The load balancer then encrypts the cookie and includes it in the response to the client. When sticky sessions are enabled, the load balancer uses the cookie received from the client to route the traffic to the same target, assuming the target is registered successfully and is considered healthy. By default, Parallels RAS uses its own ASP.NET cookie named `_SessionId`, however in this case you must customize the cookie specifying the mentioned AWS cookie for sticky sessions. This can be configured using the **Web cookie** field on the **Web Requests** tab. Please note that this functionality is available in Parallels RAS 17.1 or newer.

Wyse ThinOS support

To publish applications from the Parallels RAS to thin clients using the Wyse ThinOS, select the **Enable Wyse ThinOS support** option on the **Wyse** tab.

Note: The Wyse tab is only available if the gateway mode is set to normal. See [Gateway mode and forwarding settings](#) for more info.

By enabling this option, the RAS Secure Gateway will act as a Wyse broker. You need to make sure that DHCP option 188 on your DHCP server is set to the IP address of this gateway for thin clients that will be booting via this Secure Gateway. Once the DHCP server is configured, click the **Test** button to verify the DHCP server settings.

The **Do not warn if server certificate is not verified** option can be selected (enabled) if a Wyse device shows an SSL warning when connecting to a RAS Secure Gateway because the hostname does not match the certificate. When the option is selected, the Secure Gateway will send Wyse clients the following parameters in the `wnos.ini` file: `SecurityPolicy=low` `TLSCheckCN=no`, which will disable SSL checks. Note that the option is not required if a certificate has the following:

- The CNAME set to the FQDN of the RAS Secure Gateway.

- The SAN set to the RAS Secure Gateway IP address.

Note that if you use a custom wnos.ini in “C:\Files (x86)” folder on Secure Gateway, the Secure Gateway will not send the SSL check parameters.

If you configure DHCP option 188 to set the broker address to a given Secure Gateway, you can verify this by clicking the **Test** button.

Secure Gateway security

You can allow or deny user access to a Secure Gateway based on a MAC address. This can be accomplished using the **Security** tab in the **RAS Secure Gateway Properties** dialog.

Using Site defaults

To use Site default settings, click the **Inherit default settings** option. To specify your own settings, clear the option. For more info, see [Site defaults \(Gateways\)](#).

Configuring security

Access Control

Here, an administrator can restrict the usage of the Secure Gateway to specific modes of operation. For example, to enable the exclusive use of the Connection Broker API for user connections, but not the use of the User Portal directly.

The options available are:

- **Block access to standard RAS sessions.** Blocks new authentication to standard RAS sessions, preventing application listing and remote session launch. Existing sessions are not affected.
- **Block access to User Portal sign in page.** Blocks access to the User Portal user interface, including portal login and related pages such as cookie policy and download. Other APIs may still be used.
- **Block access to sessions via Web API.** Blocks new session launch through the Web API. Existing sessions are not affected.
- **Block access to deprecated Web API.** Blocks deprecated Web APIs to reduce security and compatibility risks. This option is ON by default for new installations as well as upgrades. If the option “Block access to sessions via Web API” is enabled, this option is automatically enforced.

- **Block access to sessions via Connection Broker API.** Blocks new session launch through the [Connection Broker API](#). Existing sessions are not affected.

MAC Address Filtering

To configure a list of allowed or denied MAC addresses, click the **Configure button** and select one of the following options:

- **Allow all except.** All devices on the network will be allowed to connect to the Secure Gateway except those included in this list. Click **Tasks > Add** to select a device or to specify a MAC address.
- **Allow only.** Only the devices with the MAC addresses included in the list are allowed to connect to the Secure Gateway. Click **Tasks > Add** to select a device or to specify a MAC address.

Please note that the Secure Gateway MAC address filtering is based on ARP, so the client and server must be on the same network for the filtering to work. It does not work across network boundaries.

Web request load balancing

Note: The **Web** tab is only available if the gateway mode is set to normal. See more in [Gateway mode and forwarding settings](#).

The **Web** tab allows you to tweak settings necessary for load balancing in certain scenarios. Here you can specify a redirection URL for web requests and a session cookie name to maintain persistence between a client and a server.

Redirection URL

An original web request can reach the gateway one of the following two ways:

- The request is sent directly to the gateway over the local network using its IP address or FQDN. For example, <https://192.168.10.10>.
- The request is sent to a HALB device that load-balances this and other gateways in the Farm. The HALB device often faces the Internet (i.e. located in DMZ) and so its DNS name can be used in the original request URL. For example, <https://ras.msp.com>. The HALB device is then distributes the request to a gateway.

When the gateway receives the web request, it takes the URL specified on the **Web** tab and sends it back to the web browser for redirection.

Technically, you can enter any URL here, and the original web request will be redirected to that URL. The primary purpose of this field, however, is to give end users an easy way to access User Portal from their web browsers. Here's how it works:

1. A user enters the Load Balancer DNS name in a web browser. For example, `https://ras.msp.com`.
2. The Load Balancer receives the request and distributes it to the least-busy RAS Secure Gateway for processing.
3. The gateway receives the original URL and replaces it with the URL specified in the **Default URL** field. See the **Default URL format** subsection below.
4. The replaced URL is then sent back to the web browser, which uses it to open the User Portal login page.

Default URL format

The default URL format is the following:

`https://%hostname%/userportal`

- The `%hostname%` variable is automatically replaced with the name of the server that received the original request, which in our example is the Load Balancer DNS name. If you wish, you can replace the variable with a specific host name or IP address (e.g. this or some other gateway). For example, `https://192.168.5.5/userportal`. If you do this, the web requests will always be forwarded to the specified host and will open the User Portal on it. Hard-coding a host may not be very practical, but you can do this nevertheless.
- `userportal` is a constant and is the path to the User Portal login page.

In our example, the resulting URL that the web browser will use to access the User Portal is the following:

`https://ras.msp.com/userportal`

The fact is, a user could simply use the above URL from the start, but thanks to the redirection feature, users only need to enter the server DNS name (or FQDN/IP-address on the local network) instead of the entire URL.

Opening a specific User PortalTheme

User Portal Themes is a feature that allows you to custom design the User Portal look and feel for different groups of users. Themes are described in detail in [Parallels Web Client and User Portal](#).

The default web request URL opens the default Theme. To make it open a specific Theme, add the Theme name at end of the URL as follows:

```
https://%hostname%/userportal/?theme=<theme-name>
```

where <theme-name> is the name of a Theme without brackets or quotes.

For users to open a specific Theme, the URL that they enter in a web browser must contain the Theme name, but in this case the format is as simple as the following:

```
https://<server-name>/<theme-name>
```

Using our Load Balancer DNS name example from above, the URL may look like the following:

```
https://ras.msp.com/Theme-E1
```

For additional information, please see [Configure Themes > URLs](#).

Web cookie

The Web cookie field is used to specify a session cookie name. RAS Web Client session persistence is normally set by the user IP address (source addressing). If you can't use source addressing in your environment (e.g. your security policy doesn't allow it), you can use the session cookie to maintain persistence between a client and a server. To do so, you need to set up a load balancer that can use a session cookie for persistence. The default cookie name is ASP.NET_SessionId. Note that if you are using Amazon Web Services (AWS) or other third-party load balancers, you may need to specify their own cookie name. See [Network load balancers access](#) for more information.

Secure Gateway tunneling policies

Tunneling policies can be used to load balance connections by assigning a group of RD Session Hosts to a specific RAS Secure Gateway or RAS Secure Gateway IP address.

To configure tunneling policies, navigate to **Farm > <Site> > Secure Gateways** and then click the **Tunneling Policies** tab in the right pane.

The **<Default>** policy is a preconfigured rule and is always the last one to catch all non-configured Secure Gateway IP addresses and load balance the sessions between all servers in the Farm. You can configure the **<Default>** policy by right-clicking it and then clicking **Properties** in the context menu.

Adding a new Tunneling Policy

To add a new policy:

1. Click **Tasks > Add**.
2. Select a Secure Gateway IP address.
3. Specify to which RD Session Host(s) the users connecting to that specific Secure Gateway should be forwarded. If you select **None** (no forwarding), read the **Restricting RDP access** section below.

Managing a Tunneling Policy

To modify an existing Tunneling Policy, right-click it and then click **Properties** in the context menu.

Restricting RDP access

You can use tunneling policies to restrict RDP accesses through the RAS Secure Gateway port. To do so, on the **Tunneling Policies** tab, select the **None** option at the bottom of the tab (this is the default setting in a new Parallels RAS installation). By doing so, you are restricting native MSTSC from accessing the gateway through its port (the default port is 80). As a result, when someone tries to use MSTSC at IP-address:80, the access will be denied. Same will happen for an RDP connection from a Parallels Client.

There are a couple of reasons why you would want to restrict RDP access. The first one is when you want your users to connect to the RAS Farm using the Parallels RAS connection only, but not RDP. The second reason is *to prevent a DDoS attack*.

A common indication of a DDoS attack taking place is when your users cannot login to a RAS Farm for no apparent reason. If that happens, you can look at the Controller.log file (located on the RAS Connection Broker server, path C:) and see that it is full of messages similar to the following:

- [I 06/0000003E] Mon May 22 10:37:00 2018 - Native RDP LB Connection from Public IP x.x.x.x, Private IP xxx.xxx.xx.xx, on Secure Gateway xxx.xxx.xx.xx, Using Default Rule
- [I 06/00000372] Mon May 22 10:37:00 2018 - CLIENT_IDLESERVER_REPLY UserName hello@DOMAIN, ClientName , AppName , PeerIP xxx.xxx.xx.xx, Secure GatewayIP xxx.xx.x.xx, Server , Direct , desktop 0
- [I 05/0000000E] Mon May 22 10:37:00 2018 - Maximum amount of sessions reached.
- [I 06/00000034] Mon May 22 10:37:00 2018 - Resource LB User 'hello' No Servers Available!
- [W 06/00000002] Mon May 22 10:37:00 2018 - Request for "" by User hello, Client , Address xxx.xxx.xx.xx, was not served error code 14.

These messages tell us that a DDoS attack is in progress on the RDP port. By restricting RDP access through Secure Gateway tunneling policies, you can prevent this from happening.

Configure logging

A RAS Secure Gateway is monitored and logs are created containing relevant information. To configure logging and retrieve or clear existing log files, right-click a gateway, choose **Troubleshooting > Logging** in the context menu, and then click **Configure**, **Retrieve**, or **Clear** depending on what you want to do. For the information on how to perform these tasks, see the [Logging](#) section.

Viewing Secure Gateway summary and metrics

You can view the summary information for all available RAS Secure Gateways in one place as follows:

1. In the RAS Console, select the **Farm** category and then select the **Site** node in the middle pane.
2. The available RAS Secure Gateways are displayed in the **Gateways** group in the right pane.
3. To go to the main Gateway view/editor, right-click a server and choose **Show in the Editor**.

You can also view the detailed information about a RAS Secure Gateway by navigating to **Information > Site** in the Parallels RAS Console. The information on this page includes general information, such as OS version, RAS version, Gateway mode, as well as the information about various types of connections, sessions, cached sockets, and threads.

Using computer management tools

You can perform standard computer management tasks on server hosting the RAS Secure Gateway right from the RAS Console. These include Remote Desktop Connection, PowerShell, Computer Management, Service Management, Event Viewer, IPconfig, Reboot, and others. To access the **Tools** menu, select a server, click **Tasks** (or right-click) > **Tools** and choose a desired tool. For requirements and usage information, see [Computer management tools](#).

RD Session Hosts

RD Session Hosts are used to host published resources (applications, desktops, documents, etc.) in a Parallels RAS Farm. Read this chapter to learn how to add, configure, and administer RD Session Hosts.

RD Session Host types

Beginning with Parallels RAS v16.5, you can create and add to a RAS Farm the following types of RD Session Hosts:

- Individual servers. These can be physical boxes or virtual machines treated as physical servers. For information on how to create these types of servers, see [Adding an RD Session Host](#).
- Virtual machines (VMs) created from a template, which is a part of RAS Virtual Desktop Infrastructure (VDI). The main advantage of using VMs is the ability to create as many of them as you require from a single template. For information on how to create these types of servers, see **Add a template-based RD Session Host**.

Considering that template is a part of RAS VDI, some aspects of creating, provisioning, and managing RD Session Hosts based on a template differ from the regular RD Session Hosts (individual servers). When reading these sections, please pay attention to whether or not a particular functionality applies to RD Session Hosts based on a template.

Add an RD Session Host

RD Session Host requirements

An RD Session Host must have the Remote Desktop Services (RDS) role installed. You can install RDS right from the RAS Console, as described later in this section.

To push install the RAS RD Session Host Agent on a server, the following requirements must be met:

- The firewall must be configured on the server to allow push installation. Standard SMB ports (139 and 445) need to be open. See also **Port reference** for the list of ports used by Parallels RAS.
- SMB access. The administrative share (\\$) must be accessible. Simple file sharing must be enabled.
- Your Parallels RAS administrator account must have permissions to perform a remote installation on the server. If it doesn't, you'll be asked to enter credentials of an account that does.

- The RD Session Host should be joined to an AD domain. If it's not, the push installation may not work and you will have to install the Agent on the server manually. See [Installing the Agent manually](#) section.

Note: The rest of this section applies to regular RD Session Hosts only. If you are looking for the information on how to add an RD Session Host based on a template, see [Add a template-based RD Session Host](#).

Add an RD Session Host

To add an RD Session Host to a Site:

1. In the RAS Console, navigate to **Farm > Site > RD Session Hosts**.
2. Click **Tasks > Add**. This opens the **Add RD Session Hosts** wizard. Note that you can also open the wizard from the **Start** category as describe in [Set up a basic Parallels RAS Farm](#).
3. Click the Tasks menu (or click the **[+]** icon) and select one of the following:
 - **Add from Active Directory:** Adds an RD Session Host from Active directory.
 - **Add Manually:** Adds RD Session Host by entering its FQDN or IP address.

Note that if you enter the server name (hostname or FQDN), it will be used as the primary method of connecting to this server from other RAS components and clients. If you enter the IP address, it will be automatically resolved to FQDN, but only if the global option to resolve to FQDN is enabled. To see the current setting of this global option, click **Tools > Options** on the main menu. In the **Options** dialog, examine the **Always attempt to resolve to fully qualified domain name (FQDN) when adding hosts** option. When the option is selected, the IP address of every server/component in the RAS Farm is always resolved to FQDN. When the option is cleared, whatever is specified for a server (IP address or name) is used to communicate with a server. This makes a difference in deployments where an IP address cannot be used to access a server, such as when a server is hosted in the cloud. For more information, see [Host name resolution](#).

4. Click **Next**.

5. On the next page, specify the following options:

- **Add firewall rules.** Add firewall rules required by Parallels RAS in Windows running on the server. See **Port reference** for details.
- **Install RDS role.** Install the RDS role on the server if it's not installed. You should always select this option.
- **Enable Desktop Experience.** Enable the Desktop Experience feature in Windows running on the server. This option is enabled only if the Install RDS role option (above) is selected. The option applies to Windows Server 2008 R2 and Windows 2012 R1/R2 on which the Desktop Experience feature is not enabled by default.
- **Restart server if required.** Automatically restart the server if necessary. You can restart the server manually if you wish.
- **Add server(s) to host pool.** Add the server (or servers) to a host pool. Select the desired host pool in the list box located below this option. If you are not sure what host pool to choose, select **Default Host pool**. Host pools are described in detail in the [Manage host pools \(RD Session Hosts\)](#) section.

6. Click **Next**.

7. Add the server (or servers) to a host pool. Select the desired host pool or create a new host pool. If you are not sure what host pool to choose, select **Default Host pool**. Host pools are described in detail in the [Manage host pools \(RD Session Hosts\)](#) section.

8. Click **Next**.

9. The next page allows you to add users and groups to the Remote Desktop Users groups in Windows running on the server. This is necessary for your Parallels RAS users to be able to access published resources hosted by an RD Session Host. To specify users and/or groups, select the option provided and then click the **[+]** icon. In the **Select Users or Groups** dialog, specify a user or a group and click **OK**. The selected user/group will be added to the list on the wizard page.

Note: If you skip this step and your users are not members of the Remote Desktop Users group on an RD Session Host, they will not be able to access published resources. If you already used (or want to use later) standard Windows tools to add users to the Remote Desktop Users group, you can skip this page.

1. Click **Next**.

2. The **User profile** page allows you to select a technology to manage user profiles. You can select from **User profile disk** or **FSlogix**. User profile disks are virtual hard disks that store user application data on a dedicated file share. Microsoft FSLogix Profile Container is the preferred Profile Management solution as the successor of Roaming Profiles and User Profile Disks (UPDs). It is set to maintain user context in non-persistent environments, minimize sign-in times and provide native profile experience eliminating compatibility issues. For complete instructions, please see [User profile](#).
3. The **Optimization** page allows you to specify settings that will be used to optimize Windows on the RD Session Host for best performance in a Parallels RAS environment. You can select Windows components, services, and other options that will be disabled, removed, or optimized to ensure a more efficient, streamlined, and improved delivery of virtual apps and desktops. For the complete description, please see [Optimization](#).
4. On the next page, review the settings and click **Next**.
5. The **Install RAS RD Session Host Agent** dialog opens. Follow the instructions and install the agent. When the installation is finished, click **Done** to close the dialog.
6. Back in the wizard, click **Finish** to close it.

If you would like to verify that the RD Session Host has been added to the Farm, click the **Farm** category (below the **Start** category in the left pane of the Parallels RAS Console window) and then click **RD Session Hosts** in the navigation tree (the middle pane). The server should be included in the **RD Session Hosts** list. The **Status** column may display a warning message. If it does, reboot the server. The **Status** column should now say, "OK", which means that your RD Session Host is functioning properly.

Read on to learn how to [publish an application from an RD Session Host](#).

Installing the agent manually

You may need to install the RAS RD Session Host Agent manually if the automatic push installation cannot be performed. For instance, an SMB share may not be available or the firewall rules may interfere with the push installation, etc.

Installing RAS RD Session Host Agent manually

1. Log in to the server where the RAS RD Session Host Agent is to be installed using an administrator account and close all other applications.
2. Copy the Parallels RAS installation file (`RASInstaller.msi`) to the server and double-click it to launch the installation.
3. Follow the onscreen instructions and proceed to the installation type page. Select **Custom** and click **Next**.
4. Click on **RAS RD Session Host Agent** and select **Entire Feature will be installed on local hard drive** from the drop-down list.
5. Ensure that all other components are deselected and click **Next**.
6. Click **Install** to start the installation.
7. Click **Finish** once the installation is finished.

The RAS RD Session Host Agent doesn't require any configuration. Once the agent is installed, highlight the server name in the RAS Console and click **Troubleshooting > Check Agent** in the **Tasks** drop-down list to update the server status.

Uninstalling RAS RD Session Host Agent

To uninstall RAS RD Session Host Agent from a server:

1. Navigate to **Start > Control Panel > Programs > Uninstall a Program**.
2. Find **Parallels Remote Application Server** in the list of installed programs.
3. If you don't have any other Parallels RAS components on the server that you want to keep, right-click **Parallels Remote Application Server** and then click **Uninstall**. Follow the instructions to uninstall the program. You may skip the steps below.
4. If you have other RAS components that you want to keep on the server, right-click **Parallels Remote Application Server** and then click **Change**.
5. Click **Next** on the Welcome page.
6. On the **Change, repair, or remove** page, select **Change**.
7. On the next page, select **Custom**.
8. Select **RAS RD Session Host Agent**, then click the drop-down list in front of it, and click **Entire feature will be unavailable**.
9. Click **Next** and complete the wizard.

Add a template-based RD Session Host

A template-based RD Session Host is a clone of a virtual machine running on a hypervisor or a cloud-based provider. When you create a template, you select a preconfigured VM with the operating system and applications already installed. Individual hosts (VMs) are then created as clones of the template. The clones can be created in advance or on as-needed basis (configurable when you create a template). This functionality allows you to create and configure an RD Session Host running in a virtual machine and then create as many copies of it as you require.

To add a template-based RD Session Host to a Site:

1. Create a template as described in [Creating an RD Session Host template](#).
2. Assign the template to a host pool as described in [Assigning a template to a host pool](#).
3. Add individual RD Session Hosts to the host pool. Do one of the following:
 - If you want to add RD Session Hosts manually, go to the host pool properties, select the **Servers** tab and click **Tasks > Add** (or click the **[+]** icon). In the dialog that opens, select the number of RD Session Host you want to create and click **OK**.
 - If you want Parallels RAS to add RD Session Hosts automatically when certain conditions are met, configure autoscaling as described in [Manage host pools \(RD Session Hosts\)](#).

Manage RD Session Hosts

Read this section to learn how manage RD Session Hosts components in Parallels RAS.

Manage host pools (RD Session Hosts)

When you publish resources in Parallels RAS, you need to specify one or more servers that host them. Host pools allow you to combine multiple RD Session

Hosts and then publish the resources from the host pool instead of specifying individual servers.

The main benefits of using RD Session Host host pools are as follows:

- They simplify the management of published resources.
- They allow you to use RD Session Hosts created from a template. More on this later in this section.

Each RD Session Hosts must belong to a host pool. Parallels RAS comes with a built-in host pool named **<Default>** that you can use. Note that an RD Session Host can be a member of one host pool only. You cannot add the same server to multiple host pools.

Moving an RD Session Host to another host pool

To move an RD Session Host from one host pool to another:

1. In the RAS console, navigate to **Farm > <Site> > RD Session Hosts**.
2. Select an RD Session Host.
3. Click **Tasks > Assign to host pool** or right-click the RD Session Host and select **Assign to host pool** in the context menu.
4. In the **Assign to Host pool** dialog, select the host pool you need.

Note: The settings of the new host pool will apply to the RD Session Host.

Autoscale

The settings on the **Autoscale** tab of the host pool properties determine how RD Session Hosts are created from the specified template. The settings are described below.

Template: Specifies the template assigned to the host pool.

Enable autoscale: Enables autoscale.

Configure: Configures the autoscale settings:

- **Min number of hosts to be added to the host pool from the Template:**
Specifies the minimum number of servers that will be added to the host pool automatically when the template is assigned to the host pool. This number of servers will remain in the host pool irrespective of utilization.

- **Max number of hosts to be added to the host pool from the Template:**

This option allows you to set a limit on how many servers in total can be added to the host pool from the template. A template can be shared between host pools. By setting a limit for each host pool, you can ensure that the combined number of servers in each host pool will not exceed the template limit. Consider the following examples:

- If the template is used by a single host pool, then this number can be up to the **Maximum hosts** setting of the template.
- If two or more host pools share the same template, then the combined number from all host pools must be less or equal to the **Maximum hosts** settings of the template.

When you save a host pool, a validation will be performed against other host pools (if any) and you will see an error message if the numbers don't match. Note that when a server cannot be created on request due to an error, a "Template error" event is triggered and the administrator will receive an alert message.

- **Add new or power on existing hosts when workload is above (%):**

Specifies a workload threshold in percent. When the actual workload is above this value, a new server (or servers) will be created and added to the host pool (if not already available). The host pool workload percentage is calculated using the following formula:

$$\text{Host pool Workload} = (\text{Current Sessions} / \text{Max Sessions}) * 100$$

In the formula above:

- Current Sessions is the total number of all sessions on all servers in the host pool. This includes static (standalone) servers and servers created from the template (host pools). Note that servers that are disabled, being drained, or have the agent status of 'Not Verified' are not included in the calculation.
- Max Sessions is a setting that you specify on the **Agent Settings** tab (either inherited from Site defaults or overridden for this host pool) and the maximum number of sessions allowed for the host pool.

Consider the following examples:

RAS Host pool 1 — mixed server types (static and host pools), different agent status:

- RDSH-1, Status: OK, Max Sessions 10, Current Sessions: 2, Type: Static
- RDSH-2, Status: Disabled, Max Sessions 20, Current Sessions: 0, Type: Static
- RDSH-3, Status: OK, Max sessions 10, Current Sessions: 4, Type: Host
- RDSH-4, Status: Drain Mode, Max sessions 10, Current Sessions: 3, Type: Host

For the host pool above, the workload is calculated as (Current Sessions / Max Sessions) * 100 or $((2 + 4) / 20) * 100 = 30\%$

Note that servers RDSH-2 and RDSH-4 are not included in the workload because the former has the agent disabled and the latter is in drain mode.

RAS Host pool 2 — mixed server types (static and host pools), different agent status:

- RDSH-1, Status: OK, Max Sessions 10, Current Sessions: 0, Type: Host
- RDSH-2, Status: OK, Max Sessions 10, Current Sessions: 2, Type: Host
- RDSH-3, Status: Not Verified, Max sessions 10, Current Sessions: 0, Type: Host

Host pool Workload = (Current Sessions / Max Sessions) * 100 or $((0 + 2) / 20) * 100 = 10\%$

Please note that a host pool will always make sure that it has at least one server available, even if the workload is zero percent.

- **Number of hosts to be added to the host pool per request:** Specifies how many servers should be created when the workload goes above the threshold value. This setting works together with the **Add servers from template when workload is above (%)** setting described above. When a host pool sends a request to the template to create additional servers, the value specified here will determine the number of servers that will be created.
- **Drain and power off hosts from host pool when workload is below (%):** Specifies a workload threshold in percent. When the actual workload is below

this value and remains there for a period specified in the **Workload remains below this level** field, excessive hosts will be switched to drain mode or powered off. The period of time can be selected from the drop-down list or you can type your own integer value using “weeks”, “days”, “hours”, “minutes”, or “seconds” as a unit measure. The server(s) with the least number of sessions will be switched to drain mode. As soon as all users are logged off from a server, it is unassigned from the host pool. At that point, the server becomes available to other host pools on demand.

- **Remove hosts from host pool after drain and power off:** Specifies if hosts should be removed from the host pool after being drained and powered off.

Tip: Servers are unassigned from the host pool only when all user sessions on that particular server are logged off. In case user sessions are still present, such as user sessions in idle, active or disconnected state, autoscaling does not log off user sessions and does not unassign the server from a host pool.

Note: Parallels recommends setting viable timeouts for idle time and disconnected sessions either in Windows Host pool Policies or in the **Site Default Properties** dialog to make the drain mode effective. GPOs can be used to forcibly log off a user session, however this should be used carefully as this may result in data loss.

Using host pool defaults

RD Sessions Hosts assigned to a host pool have various settings that they can inherit from the host pool defaults. This makes it simpler to configure a single set of settings for all servers instead of configuring each server individually. A Site also has its own default settings (Site defaults). Moreover, an RD Session Host host pool can inherit these Site defaults. This gives you the following choices when inheriting default settings by an RD Session Host:

- Configure Site defaults and make the host pool inherit these settings. The RD Session Hosts assigned to the host pool will therefore also inherit Site defaults. This is the default scenario for a new host pool. Site defaults can be configured by navigating to **Farm > <Site> > RD Session hosts** and clicking **Tasks > Site defaults**.
- Configure default settings for a given host pool. This way you can have multiple host pools, each having its own host pool defaults (different from Site defaults). Therefore, the servers assigned to a host pool will inherit the host pool's defaults.

To configure default settings for a host pool, open the **Host pool Properties** dialog (**Tasks > Properties**), select a desired tab (except the **General** tab, which doesn't have any defaults) and select or clear the **Inherit default settings** option. If you clear the option, you can specify your own defaults. All servers that are (or will be) assigned to this host pool will inherit these settings. Note that inheritance works independently for each individual tab on the host pool properties dialog.

For information on how default settings are configured for an RD Session Host, see [View and modify RD Session Host properties](#).

Add host pools (RD Session Hosts)

Creating a host pool

To create an RD Session Host host pool:

1. In the RAS console, navigate to **Farm > <Site> > RD Session Hosts****Host pools**.
2. Click **Tasks > Add** (or click the **[+]** icon).
3. Select **Enable Host pool in site** to enable the host pool. Specify the name and the description for the new host pool.
4. Click **Next**.
5. On the **Provisioning** page, select whether this host pool will contain template-based or standalone hosts:
 - **Template:** (Template-based RD Session Hosts only) Hosts will be created dynamically from a template. You will need to create or select an existing template in the next step or later. Choosing **Template** as the provisioning type ensures a homogeneous host pool, which is recommended to provide consistent user experience across the host pool. For more information about creating template-based RD Session Hosts, see section [Add a template-based RD Session Host](#).
 - **Standalone:** (Template-based and standalone RD Session Hosts) Select one or more hosts that already exist. You'll be able to do it in the next step or you can do it later. Prior to adding hosts to host pools, ensure that hosts are domain joined and have network access to the domain environment. Note that the Standalone provisioning is considered "unmanaged" as it lacks some of the functionality, such as Autoscaling.
6. Click **Next**.

7. Depending on the selection made on the **Provisioning** page (above), do one of the following

- **Standalone**: Select one or more hosts from the list to be included in the host pool (you can also add hosts to the pool later).
- **Template**: Select a template from the list or click **Create new** to create a new template and specify the template settings. **Versions**: If you selected an existing template, select one of its versions. **Enable autoscale**: (Multi-session hosts) Enable and configure autoscale.

8. Click **Next**.

9. (Templates only) On the **Provisioning** page, specify the following options:

- **Template name**: Choose and type a template name.
- **Host name**: A pattern to use when naming new hosts.
- **Number of hosts deployed on wizard completion**: The number of hosts to deploy once the template is created. Please keep in mind that this will take some time because the hosts will be created one at a time.
- **Host state after the preparation**: Select the power state that should be applied to a host after it is prepared. Choose from **Powered on**, **Powered off**, or **Suspended**. Note that when the power state is set to **Power off** or **Suspended**, the number of running (fully ready and waiting for incoming connections) hosts is controlled by the **Keep available buffer** setting (see above). For example, let's say the **Maximum hosts** value is set at 200, the number of guest hosts deployed on wizard completion is 100, and the power state after preparation is **Powered off**. The result of such a configuration will be 100 clones deployed and powered off.
- **Maximum hosts**: Specify the maximum number of hosts that can be created from this template.
- (Only for templates created from VMware vCenter instances) Select one of the following:
 - If you want to create all hosts in the datastore that contains the template, select **Create hosts in the primary datastore**.
 - If you want to create hosts in several datastores, select **Create hosts in specific datastores**. If you select this option, you will be prompted to select the datastores in the **Datastores** list on the next page.
- (Only for templates created from AWS instances) **Virtual networks**: Select Availability Zones for the host pool. Choose from **Use the subnet associated with the template** and **Use specific subnets available in the VPC**. If you select **Use specific subnets available in**

the **VPC**, then after clicking **Next** you can manually select one or several Availability Zones.

- Autoscale settings: Configure autoscale settings as described in [Manage host pools \(RD Session Hosts\)](#).

10. Click **Next**.

11. On the **Specifications** page, override settings that are defined at the template level.

- To override the Organizational Unit in Active Directory where hosts are joined, select **Override Organizational Unit specified in template properties**. In the target OU field, provide the distinguished name of the Organizational Unit by using the browse (...) button or entering it manually.
- (Microsoft Azure only) To override the High availability settings, select **Override the high availability settings specified in template properties**. Select an availability option which can be 'no infrastructure redundancy required', 'availability set', or 'availability zone'. When 'availability zone' is selected, specify the zone where hosts should be deployed. Note that High availability settings cannot be changed after the host pool is created.
- (Microsoft Azure only) To override the host size, select the **Overwrite the size specified in template properties** option and choose the desired size from the drop-down list. The selected size will only be used by this host pool. Other host pools using the same template will be unaffected. Note that available sizes may depend on the location, size, and power state of the host pool members and the template. Also note that overwriting the size requires a host reboot.
- (Nutanix only) To override the Cluster settings, select **Override the cluster specified in Template Properties**. In the Custer field, provide the cluster name by using the browse (...) button or entering it manually.
- (Nutanix only) To override the Network settings, select **Override the network specified in Template Properties**. In the Network field, provide the network name by using the browse (...) button or entering it manually.

12. On the **User profile** page, you can select from **Do not manage by RAS** (user profiles will not be managed) or **FSlogix**. Microsoft FSLogix Profile Container allows to maintain user context in non-persistent environments, minimize sign-in times and provides native profile experience eliminating compatibility issues. For complete instructions, please see [User profile](#).

13. Click **Next**.

14. On the **Summary** page, review the template summary information. You can click the **Back** button to correct some of the information if needed.
15. Finally, click **Finish** to create the host pool and close the wizard.

After you create a host pool and later publish resources from it, you can view the list of resources by right-clicking a host pool and choosing **Show published resources** (or click **Tasks > Show Published Resources**). For more information, see [Viewing published resources hosted by RD Session Hosts](#).

Upgrading Agents (RD Session Hosts)

You can enable and configure automatic updates for all RD Session Host Agents in a host pool.

Schedule Agent auto-upgrade

To schedule Agent auto-upgrade:

1. Go to **Farm > Site > RD Session Hosts > Host pools > Properties > Auto-upgrade** tab.
2. Clear the **Inherit default settings** options if you want to modify them for this host pool.
3. Select the **Enable auto-upgrade maintenance window** option. During the maintenance window, all hosts in the host pool will try to download Agent upgrades. The upgrades will be downloaded and installed as soon as all users log out of their hosts. New logons from users are prohibited (drain mode). If the users don't log off during a maintenance window, the upgrades won't be installed until the next window.
4. Specify the start date and time, duration, and recurrence settings for this event. To make this a one-time event, select **Never** in the **Recur** drop-down list.
5. (Optional) If you want to forcefully log off all users and download the upgrades at the end of a maintenance window, select the **Force logoff of current sessions at the end of the maintenance window duration** option.
6. (Optional) Configure a message that will be sent to users before or during a maintenance window. Click the **Configure messages** button and specify the message title, body, and the time period when it should be sent.

Cancel Agent auto-update

To cancel Agent auto-update:

1. Go to **Farm > Site > RD Session Hosts > Host pools**.
2. **Select Tasks > Cancel auto-upgrade maintenance window**.

Manage templates (RD Session Hosts)

RD Session Host templates are designed specifically to give you the ability to replicate RD Session Hosts running in virtual machines. Hosts created from an RD Session Host template are treated by Parallels RAS almost like regular RD Session Hosts. The main difference is, you can create as many hosts from a single template as you require, thus automating RD Session Host provisioning according to your needs.

RD Session Host templates are supported on the following VDI platforms:

- Microsoft Hyper-V
- Microsoft Hyper-V Failover Cluster
- VMware VCenter
- VMware ESXi
- SC//HyperCore
- Nutanix AHV (AOS)
- Microsoft Azure
- Amazon Web Services

RD Session Host templates support Windows Server 2008 R2 up to Windows Server 2022 as a guest OS. Compared to regular [RD Session Hosts](#), servers created from an RD Session Host template do not support earlier versions of Windows Server. The reason is, these servers run in VMs and require the RAS Guest Agent installed in them, so the guest OS requirements are limited by Windows Server versions supported by RAS Guest Agent.

Please note that the following standard RAS VDI features are not available when using RD Session Host templates:

- Pool management
- Persistent hosts
- Session management
- Publishing from a specific Template

- Some other strictly RAS VDI specific features.

For the information on how to provision RD Session Hosts created from a template, see [Manage host pools \(RD Session Hosts\)](#).

Creating an RD Session Host template

Requirements

To complete the tasks described in this section, the following requirements must be met:

- Requirements described in the “Requirements” subsection of [Creating a VM template](#).
- Network Discovery UDP port 137 must be enabled for a domain firewall profile in the guest OS. This can be done via domain group policies or manually in the guest OS.

Manual agent installation

Normally, you will push install the necessary agent software in a source VM right from the Parallels RAS console. However, you can also install the software manually by running the Parallels RAS installer in Windows in the VM. When doing so, use the **Custom** installation option and select the following agent components RAS Guest Agent and RAS RD Session Host Agent to be installed in the source VM.

Create a template

To create an RD Session Host template:

1. Add one of the supported provides, as described in [Add a Provider](#).
2. Go to **Farm > Site > RD Session Hosts > Templates** tab.
3. In the **Tasks** drop-down menu, click **Add** (or click the **[+]** icon).
4. In the dialog that opens, select a host from which you would like to create a template and click **OK**.
5. The **Create Parallels Template Wizard** opens. Each wizard page is described below in the order they appear on the screen.

6. Verify that the Agent is installed and install it manually if needed as described in [Step 1: Check and install the Agent](#). This step only appears if an on-premises Provider is used.
7. Configure the template as described in [Step 2: Configure the template](#).

Assigning a template to a host pool (RD Session Host)

When you create RD Session Host host pools, you can assign a template to a host pool. This can be done when you create or modify a host pool, or it can be done from the **Templates** tab.

To assign a template to a host pool:

1. Go to **Farm > Site > RD Session Hosts > Templates** tab.
2. On the **Templates** tab, select a template.
3. Click **Tasks > Assign to host pool**.
4. Select the template version in the **Version** dialog.
5. A dialog opens listing existing RD Session Host host pools. Host pools that already have a template assigned are not shown in this list by default. To display them, select the **Show host pools with assigned template** option. The template that they are currently using is displayed in the **Template** column.
6. Select one or more host pools and click **OK**.

To remove a template from a host pool or host pool:

1. Select a template and click **Tasks > Remove from host pool**.
2. A dialog opens listing all host pools to which this template is assigned.
3. Select the host pools to remove the template from and click **OK**.

Note that if a host pool has hosts created from the template that you are removing, they will be removed as well. A message is displayed where you need to confirm the removal.

Managing RD Session Hosts based on a template

Viewing RD Session Hosts based on a specific template

To view the list of RD Session Hosts based on a specific template:

1. Go to **Farm > <Site> > RD Session Hosts > Templates**.
2. Select a template and click **Tasks > Show servers**.

Site defaults

RD Session Hosts based on a template inherit the template settings. To view the settings, note on which template an RD Session Host is based and then view properties of that template, specifically the **Settings** and **Security** tabs. For more information, see [Site defaults](#). Note that you a template can inherit Site default settings or you can specify your own custom settings for it.

Checking the RAS Guest Agent status

A guest RD Session Hosts based on a template must have RAS Guest Agent installed and the agent must match the Parallels RAS version. The agent is installed by default when an RD Session Host is created from a template. If the RD Session Host was created using the native hypervisor tools, it may not have the agent installed in it. In such a case, the RD Session Host will be able to serve only the remote desktop. To enable it to server applications or documents, you'll need to install the agent yourself.

To check if the RAS Guest Agent is installed and up to date:

1. Go to **Farm > <Site> > RD Session Hosts > RD Session Hosts**.
2. Continue as described in [Managing hosts](#), subsection "Checking the RAS Guest Agent status".

The RD Session Host template must also have the RAS RD Session Host Agent installed.

To check if the RAS RD Session Host Agent is installed and up to date:

1. Go to **Farm > <Site> > RD Session Hosts > Templates**.
2. Select a template in the list and then click **Tasks > Troubleshooting > Check agent**.

Deleting an RD Session Hosts based on a template

See [Managing hosts](#), subsection “Deleting a host”.

Managing RD Session Hosts based on a template that failed preparation

See [Managing hosts](#), subsection “Managing hosts that failed preparation”. Notice that in case of RD Session Hosts, you have to go to **Farm > <Site> > RD Session Hosts > RD Session Hosts** and click **Tasks > Site defaults** to see Site Defaults.

Recreating an RD Session Hosts based on a template

If something happens to a RD Session Hosts based on a template and it becomes unusable, you don't have to delete it and create a new one. Instead, you can recreate it keeping its name, MAC address, and other properties. This way none of the other Site settings, which may rely on a broken RD Session Host, will be affected. Another reason for recreating an RD Session Host is to apply changes made to the template (when you exit from maintenance without executing the Recreate command).

Please note that recreated RD Session Hosts can keep the the following properties:

- MAC address is kept on ESXi, vCenter, Hyper-v, Hyper-v Failover Cluster, Nutanix AHV (AOS), and SC//HyperCore.
- BIOS UUID is kept on ESXi and vCenter.
- DRS groups are kept on vCenter.

Note: If an RD Session Host based on a template was already assigned to an RD Session Host host pool, it cannot be recreated.

To recreate one or more guest RD Session Host:

1. In the Parallels RAS Console, navigate to **Farm** > <Site> > **RD Session Hosts** > **Templates**.
2. To recreate all deployed RD Session Host, click the **Tasks** drop-down list and choose **Recreate all servers**.
3. To recreate a specific host (or multiple hosts), click **Tasks** > **Show servers**. This will open the dialog which will list RD Session Hosts. Select one or more RD Session Hosts and then click the **Tasks** > **Recreate**.

When you recreate a RD Session Host based on a template:

- The procedure deletes the RD Session Host and creates a new one from the same template.
- The new RD Session Host retains the same computer name as the one it replaces.
- If an RD Session Host is running, all unsaved data in its memory will be lost. For this reason, an important data should be saved to an external storage.

Manage hosts (RD Session Hosts)

This section describes how to configure and manage an existing RD Session Host.

Viewing RD Session Hosts

To view the list of RD Session Hosts for the current Site:

1. In the RAS Console, navigate to **Farm** > <Site-name> > **RD Session Hosts**.
2. The available RD Session Hosts are displayed on the **RD Session Hosts** tab in the right pane.

You can filter the **RD Session Hosts** list as follows:

1. Click the magnifying glass icon, which is located on a toolbar above the list.
2. An extra row is displayed at the top of the list where you can type a string in one or more columns that will be used to filter the list.
3. For example, if you want to search for a server by its name, enter the text in the **Server** column. You can type the entire server name or the first few characters until a match is found. The list will be filtered as you type and only the matching server(s) will be displayed.

4. If you type a filter string in more than one column, they will be combined using the logical AND operator.
5. To remove the filter and display the complete list, click the magnifying glass icon again.
6. If you click the magnifying glass icon one more time, you'll see that the filter that you specified earlier is still there. To remove it completely, simply delete the filter string(s) from the column(s).

Viewing RD Session Host summary

In addition to the RD Session Hosts editor described above, you can also see the summary about the available RD Session Hosts. To do so:

1. In the RAS Console, select the **Farm** category and then select the **Site** node in the middle pane.
2. The available servers are displayed in the **RD Session Hosts** host pool in the right pane.
3. To go to the RD Session Host editor (described above), right-click a server and choose **Show in the Editor**.

For additional info, see [Sites in the RAS Console](#).

Available menu options

You can perform a number of tasks on the an RD Session Host using menus. To do so, click the **Tasks** drop-down list and choose a desired option, or right-click a host and choose an option from the context menu.

Please note that not all menu options are available for RD Session Hosts based on a template. If an option is not available for this host type, it will be either disabled or hidden. These include:

- **Assign to host pool.** Host pool assignment is performed automatically for template-based hosts.
- **Delete.** Deleting a host (which is a VM) can only be done on the template level (the **Host List** dialog).
- **Properties.** RD Session Hosts of this type don't have individual properties. Some essential properties are inherited from **Default Server Properties** (see [View and Modify RD Session Host Properties > Agent Settings](#)).

- **Control** (logon commands). Drain mode is managed automatically by the host pool to which a template-based host belongs.

Check an RD Session Host Agent status

An RD Session Host must have RAS RD Session Host Agent installed in order to publish remote applications and desktop from it. In addition to this, Remote Desktop Services (formerly Terminal Services) must also be installed.

Normally when you add an RD Session Host to a Site, the RD Session Host Agent and Remote Desktop Services are installed by default. However, if you skipped the installation (or uninstalled the agent or RDS from the server), you can check their status and take appropriate actions if needed.

To check the status of RD Session Host Agent and RDS, do the following:

1. First, check the **Status** column in the **RD Session Hosts** list. The column should display “OK”. If so, the Agent is installed and functioning properly. If not, read on.
2. In addition to the description, the **Status** column uses a color code to indicate the agent status as follows:
 - Red — not verified
 - Orange — needs update
 - Green — verified
3. Right-click a server and click **Troubleshooting > Check agent** in the context menu. The **Agent Information** dialog opens.
4. If the agent is not installed on the server, click the **Install** button and follow the instructions on the screen.

After the agent installation is complete, you may need to reboot the RD Session Host. You can do it right from the Parallels RAS Console by selecting the server and clicking **Tasks > Control > Reboot**.

Change RD Session Host Site assignment

You can assign an RD Session Host to a different Site in your Farm if needed. Please note that this functionality is only available if you have more than one Site in your Farm.

To change the Site assignment:

1. Right-click an RD Session Host and then click **Change Site** in the context menu. The **Change Site** dialog opens.
2. Select a Site in the list and click **OK**. The server will be moved to the **RD Session Hosts** list of the target Site (**Farm** > <new-site-name> > **RD Session Hosts**).

View and modify RD Session Host properties

Note: The information in this section does not apply to RD Session Hosts based on a template. Hosts of that type don't have individual properties and are managed on the template level. For more information, see [Manage host pools \(RD Session Hosts\)](#) and [Templates](#).

To configure an RD Session Host:

1. In the RAS Console, navigate to **Farm** > <site> > **RD Session Hosts**.
2. Select a server and click **Tasks** > **Properties**.
3. The server properties dialog opens where you can configure the RD Session Host properties.

The dialog is described in the subsections that follow this one.

Using default settings

The server properties dialog consists of tabs, each containing their own specific set of properties. All tabs, except General, have either **Group Defaults** or **Site defaults** link, which allows you to view and modify default settings. If you want the

properties on a particular tab to inherit default settings, select the **Inherit default settings** option. When you do, the default settings will be inherited from one of the following:

- **Group defaults.** Groups are described in [Grouping and cloning RD Session Host Servers](#).
- **Site defaults.** Note that a group may also inherit Site defaults, but this can be overridden in the group properties dialog where you can specify custom settings for a group.

To view or modify default settings, click the **Group Defaults** or **Site defaults** link. Note that each individual tab can inherit default settings independently from other tabs.

To specify custom settings for an RD Session Host, clear the **Inherit default settings** option and use the controls on a given tab to set the desired options.

General

Select or clear the **Enable Server in site** option to enable or disable the server. A disabled server cannot serve published applications and virtual desktops to clients.

Other elements on this tab are:

- **Server:** Specifies the server FQDN or IP address.
- **Description:** An optional server description.
- **Change Direct Address:** Select this option if you need to change the direct address that Parallels Client uses to establish a direct connection with the RD Session Host.

Agent settings

Each RD Session Host in a RAS Farm has an RAS RD Session Host Agent installed through which it communicates with other Parallels RAS components. Use the **Agent Settings** tab to configure the agent.

To use default settings, select the **Inherit default settings** option. See [Using default settings](#).

To configure the agent, set the options as described below.

Application session lingering

- **Disconnect active session after:** Specifies the amount of time each session remains connected in the background after the user has closed a remote application. This option is used to avoid unnecessary reconnections with the server.
- **Logoff disconnected session after:** This setting allows you to control how long it takes for a session to be logged off after it is marked as “disconnected”.

Other settings

- **Port:** Specifies a different remote desktop connection port number if a non-default port is configured on the server.
- **Max sessions:** Specifies the maximum number of sessions.
- **Preferred Connection Broker:** Select a Connection Broker to which the RD Session Host should connect. This is helpful when Site components are installed in multiple physical locations communicating through WAN. You can decrease network traffic by specifying a more appropriate Connection Broker.

Allow Client URL/Mail redirection

When a user tries to open a URL or an HTML Mailto link in a remote application, the link can be redirected to the client computer and open in a local default application (a web browser or email client) instead of an application on the remote host. To enable this functionality, select the option and click the **Configure** button. In the dialog that opens, select one of the following:

- **Replace Registered Application:** This option uses an alternative method of redirecting a link. It replaces the default web browser and mail client with “dummy” apps on the remote server side. By doing so, it can intercept an attempt to open a link and redirect it to the client computer.
- **Support Windows Shell URL namespace objects:** The Shell URL namespace objects support means that Parallels RAS can intercept actions in published applications that use the Shell namespace API to open links, which is a standard behavior in most applications. The ability to disable support for Shell URL namespace objects is for compatibility with older versions of Parallels RAS. You may disable this option if you want the behavior of an older version of Parallels RAS (RAS v16.2 or earlier).

Please note that you can configure a list of URLs that should never be redirected, even if the redirection is enabled. This can be done on the **Farm > Site > Settings > URL Redirection** tab. See more in [Site settings](#).

Enable Drag and drop

Allows you to set how the drag and drop functionality works in Parallels Clients. To enable drag and drop, select the option, click the **Configure** button and then select from the following:

- **Server to client only:** Drag and drop to a local application, but not in the opposite direction.
- **Client to server only:** Drag and drop to a remote application only.
- **Bidirectional:** Note that this option has changed since Parallels RAS 17.1. In the past, it was a checkbox that would enable or disable drag and drop which worked in the “Client to server only” mode. When upgrading from an older version of Parallels RAS, and if the checkbox was enabled, the “Client to server only” option is selected by default. If the option was disabled, the “Disabled” option will be set. You can change it to any of the new available options if you wish.
- **Note:** At the time of this writing, the drag and drop functionality is only supported on Parallels Client for Windows and Parallels Client for Mac.

Allow 2XRemoteExec to send command to the client

Select this option to allow a process running on the server to instruct the client to deploy an application on the client side. Read more about 2XRemoteExec in the **Using RemoteExec** subsection at the end of this topic.

Use RemoteApp if available

Enable this option to allow use of remote apps for shell-related issues when an app is not displayed correctly. This feature is supported on the Parallels Client for Windows only.

Manage RDP transport protocol

Select the transport protocol that will be used for connections between Parallels Client and a server. To do this, select this option and click the **Configure** button.

Enable application monitoring

Enable or disable monitoring of applications on the server. Disabling application monitoring stops the WMI monitoring to reduce CPU usage on the server and network usage while transferring the information to RAS Connection Broker. If the option is enabled, the collected information will appear in a corresponding RAS report. If the option is disabled, the information from this server will be absent from a report.

Allow file transfer command (Web and Chrome clients)

Enables file transfer in a remote session. To enable file transfer, select this option and click the **Configure** button. For more information, see [Configuring remote file transfer](#).

Enable drive redirection cache

Improves user experience by making file browsing and navigation on redirected drives much faster. For details, see [Drive redirection cache](#).

Enable Z-order (Experimental)

Select this option if you encounter an issue related to z-order in applications published with Parallels RAS.

Using 2XRemoteExec

2XRemoteExec is a feature that facilitates the servers ability to send commands to the client. This is done using the command line utility 2XRemoteExec.exe.

Command line options include:

Command Line Parameter	Parameter Description
-s	Used to run the 2XRemoteExec command in 'silent' mode. Without this parameter, the command will display pop up messages from the application. If you include the parameter, the messages will not be displayed.
-t	Is used to specify the timeout until the application is started. Timeout must be a value between 5000ms and 30000ms. Note that the value inserted is in 'ms'. If the timeout expires the command returns with an error. Please note that the application might still be started on the client.
-?	Shows a help list of the parameters that 2XRemoteExec uses.
"Path for Remote Application"	The Application that will be started on the client as prompted from the server.

2XRemoteExec examples:

The following command displays a message box describing the parameters that can be used.

```
2XRemoteExec -?
```

This command runs Notepad on the client.

```
2XRemoteExec C:\Windows\System32\Notepad.exe
```

In this example, the command opens the C:\readme.txt file in the Notepad on the client. No message is shown and 2XRemoteExec would wait for 6 seconds or until the application is started.

```
2XRemoteExec C:\Windows\System32\Notepad.exe "C:\readme.txt"
```

User profile

Use this tab to configure user profile settings.

To use default settings, select the **Inherit default settings** option. See [Using default settings](#).

For complete instructions about configuring user profiles, see [User profile](#).

Application Packages

The **Application Packages** tab allows you to manage MSIX application packages on RD Session Hosts and groups.

To use default settings, select the **Inherit default settings** option. See [Using default settings](#).

Adding a package to an RD Session Host

See [Using MSIX application packages](#), subsection “Adding a package to a host”.

Adding a package to a VDI pool

See [Using MSIX application packages](#), subsection “Adding a package to a VDI pool”.

Managing applications installed from MSIX packages

The following actions are available from the **Task** drop-down list:

- **Add**: Add a new package to the RD Session Host.
- **Retry Staging**: Manually trigger re-staging of all added packages.
- **Refresh**: Refresh the list of the packages.
- **Delete**: Delete the selected package.

Optimization

The **Optimization** tab allows you to specify settings that will be used to optimize the RD Session Host for best performance in a Parallels RAS environment. You can select Windows components, services, and other options that will be disabled, removed, or optimized to ensure a more efficient, streamlined, and improved delivery of virtual apps and desktops.

To use default settings, select the **Inherit default settings** option. See [Using default settings](#).

For the complete description, please see [Optimization](#).

Desktop access

The **Desktop Access** tab allows you to restrict remote desktop access to certain users.

To use default settings, select the **Inherit default settings** option. See [Using default settings](#).

By default, all users who have access to remote applications on an RD Session Host can also connect to the server via a standard RDP connection. If you want to restrict remote desktop access to certain users, do the following:

1. On the **Desktop Access** tab, select the **Restrict direct desktop access to the following users** option. If you have the **Inherit default settings** option selected, click the **Edit Defaults** link to see (and modify if needed) the default configuration. The rest of the steps apply to both the **Server Properties** and **Default Server Properties** dialogs.
2. Click the **Add** button.
3. Select the desired users. To include multiple users, separate them by a semicolon.
4. Click **OK**.
5. The selected users will appear in the list on the **Desktop Access** tab.

Users in this list will still be able to access remote applications using Parallels Client, but will be denied direct remote desktop access to this server.

Note: **Computer Configuration > Administrative Templates > Windows Components > Remote Desktop Services > Remote Desktop Session Host > Connection > Allow users to connect remotely using remote desktop services** must be set to **Not configured**, otherwise it takes precedence.

Please note that members of the Administrator group will still be able to connect to the remote desktop even if they are included in this list.

RDP printer

The **RDP Printer** tab allows you to configure the renaming format of redirected printers. The format may vary depending on which version and language of the server you are using.

To use default settings, select the **Inherit default settings** option. See [Using default settings](#).

The **RDP Printer Name Format** drop-down list allows you to select a printer name format specifically for the configured server.

Select the **Remove session number from printer name** and the **Remove client name from printer name** options to exclude the corresponding information from the printer name.

User profile

User profile is a collection of settings and application data associated with a specific user. In a non-persistent remote environment, such as Parallels RAS, user profiles must be maintained to provide consistent user experience. This is achieved by storing user profile data in a network location to minimize sign in times and optimize file I/O between host, client, and the profile storage.

Parallels RAS supports the following technologies to manage user profiles:

- [User profile disk](#): [RD Session Hosts only] These are virtual hard disks that store user application data on a dedicated file share. This disk is mounted to the user session as soon as the user signs in to a session host. The disk is unmounted when the user logs out.

Note: The User Profile Disks technology is no longer being actively developed by Microsoft. It's recommended to migrate profiles to [FSLogix](#). Please note that the **User profile disk** option is not available for VDI and Azure Virtual Desktop due to obsolescence.

- [FSLogix](#): A remote profile solution for non-persistent environments. FSLogix Profile Container redirects the entire user profile to a remote location and maintains user context in non-persistent environments, minimizing sign-in times and providing native profile experience eliminating compatibility issues.

FSLogix Profile Container is the preferred profile management solution as the successor of Roaming Profiles and User Profile Disks.

User profiles can be configured for the following:

- RD Session Hosts
- VDI
- Azure Virtual Desktop

User profile settings are configured for the above on the Site level (Site defaults) and can also be configured for individual components if the RAS administrator decides to use custom settings for a given component.

To configure user profile on the Site level, navigate to **Farm > Site**, click the **Tasks > Site defaults** menu and choose one of the following:

- **RD Session Host**
- **VDI**
- **AVD multi-session hosts**
- **AVD single-session hosts**

In a Site defaults dialog that opens, select the **User profile** tab. The user interface for configuring optimization is the same for all of the above.

The subsequent sections describe in detail how to configure the user profile functionality.

User Profile Disks

To configure User Profile Disks, specify the following settings:

1. When in the host “Properties” dialog, clear the **Inherit default settings** if you want to specify different settings for this host.
2. In the **Technology** section, select **User profile disk**.
3. In the drop-down list, select one of the following:
 - **Do not change**: Keep the current server settings (default).
 - **Enabled**: Enable the User Profile Disks functionality.
 - **Disabled**: Disable the functionality.
4. Click the **Configure advanced User Profile Disks settings** button to open the **User Profile Advanced Settings** dialog.

5. On the **Disk** tab, specify the following:
 - **Disk location:** If you selected **Enabled** in the previous step, specify a network location where the User Profile Disks should be created. Use the Microsoft Windows UNC format to specify a location (e.g. \\\RAS\\users\\disks). Please note that the server must have full control permissions on the disk share.
 - **Maximum size:** Enter the maximum allowed disk size (in gigabytes).
6. On the **Folders** tab, specify the following:
 - **Store all user settings and data on the user profile disk:** All folders, except those specified in the exclusion list, will be stored on the user profile disk. To add or remove folders to/from the exclusion list, click the **[+]** or **[-]** buttons.
 - **Store only the following folders on the user profile disk:** Only folders specified in the inclusion lists will be stored on the user profile disk. There are two inclusion lists. The first one contains standard user profile folders (e.g. Desktop, Documents, Downloads, etc.) and allows you to select the folders that you want to include. The second list allows you to specify additional folders. Click the **[+]** or **[-]** buttons to add or remove folders.

Note that when you enable User Profile Disks, you need to restart the server for the changes to take effect.

FSLogix

Note: If you have existing FSLogix Profile Containers and would like their configurations to be managed by Parallels RAS, please read additional instructions in [Configure managing existing profiles by Parallels RAS](#). endhint %}

Supported FSLogix releases

Parallels RAS has been tested with FSLogix releases up to and including release 26.01.

Configure installation method

Before you configure FSLogix for a specific server or a template (described later in this guide), you need to configure the FSLogix installation method on the Site level as follows:

1. Navigate to **Farm > Site > Settings** and select the **Features** tab. Here you need to select a method that Parallels RAS will use to install FSLogix on individual hosts. You can select from one of the following:
 - **Install manually**: Select this option if you want to install FSLogix on every host yourself. If this option is selected, Parallels RAS will not attempt to install FSLogix on a host.
 - **Install online**: This option installs FSLogix on session hosts from the Internet. Select one of the supported FSLogix versions from the drop-down list or select **Custom URL** and specify a download URL. Click the **Detect latest** button to automatically obtain a URL of the latest FSLogix version.
 - **Install from a network share**: Select this option if you have the FSLogix installation files on a network share and specify its location.
 - **Push from RAS Connection Broker**: This option allows you to upload the FSLogix installation archive to the RAS Connection Broker server. When you enable FSLogix on a session host, it will be push installed on the host from the RAS Connection Broker server.
2. When done, click **Apply** in the RAS Console to apply your changes to Parallels RAS.

Upgrade FSLogix

The dialog described above can also be used to upgrade FSLogix to a newer version. To upgrade, do one of the following:

- Select **Install online** and choose from one of the provided FSLogix builds or specify a custom URL. The **Detect latest** button obtains a URL for the latest stable FSLogix build.
- Download a new version from the Microsoft website, place it on a network share or upload it to the RAS Connection Broker server and then select **Install from a network share** or **Push from RAS Connection Broker**, whichever applies.

If FSLogix is already installed on one or more hosts and a new version of FSLogix becomes available when you do one of the above, FSLogix will be upgraded on

hosts that have it installed. Note that if you specify a version that is earlier than the version installed on a host, then FSLogix will be downgraded.

Configure Site defaults and hosts for FSLogix

To configure Site defaults or individual hosts for FSLogix, do one of the following:

- For Site defaults, navigate to **Farm > Site** and click **Tasks > Site defaults > RD Session Hosts** (or **VDI** to configure defaults for VDI, or one of the **AVD** options to configure site defaults for Azure Virtual Desktop).
- To configure individual hosts, navigate to **Farm > Site > RD Session Hosts**. Right-click a host and choose **Properties**.
- When you add an RD Session Host to a Farm, the FSLogix settings are specified on the **User profile** page.

In the **Site defaults** or **Properties** dialog, select the **User profile** tab and specify the following options:

1. If you are in the host **Properties** dialog (or in a wizard where you add a new host or template), clear the **Inherit default settings** option if you want to specify different settings for this host.
2. In the **Technology** section, select **FSLogix**.
3. The **Deployment method** field shows the currently set deployment method as configured on the Site level (see the description above). You can click the **Change...** link and select a different method. Note that this will modify the Site setting, which will change it for all hosts in the Site.
4. If you want to use Profile Containers, select the **Use Profile Containers** options. Click the **Configure** button to configure settings:
 - **Users and Groups** tab: Specify include and exclude user and group lists. By default, Everyone is added to the FSLogix profile include list. If you want some user profiles remain local, you can add those users to the exclude list. Users and group can exist in both lists but exclude takes priority.
 - **Folders** tab: Specify include and exclude lists for folders. You can select from common folders or you can specify your own. Please note that folders must reside in user profile path.
 - **Disks** tab: Specify the settings of the profile disk. **Location type:** Select a location type for profile disks (SMB Location or Cloud Cache) and then specify one or more locations. **Location of profile disks:** Location(s) of profile disks. These are the locations of VHD(X) files (the VHDLocations setting in the registry as specified in the FSLogix

documentation). **Profile disk format:** Select from VHD or VHDX according to your requirements. VHDX is a newer format and has more features. **Allocation type:** Select **Dynamic** or **Full**. This setting is used in conjunction with the **Default size** setting (see below) to manage the size of a profile. Dynamic causes the profile container to use the minimum space on disk, regardless of the allocated Default size. As a user profile is filled with more data, the amount of data on disk will grow up to the size specified in Default size, but will never exceed it. **Default size:** Specifies the size of newly created VHD(X) in megabytes.

- **Advanced** tab: This tab allows you to modify advanced FSLogix registry settings. To modify a setting, select it and click **Tasks > Edit**. By default, the settings are disabled. To enable a setting, select the checkbox in front of its name. A description for each setting is provided in the RAS console. For further information regarding FSLogix Profile Containers configurations, visit <https://docs.microsoft.com/en-us/fslogix/profile-container-configuration-reference>.

5. If you want to use Office Containers, select the **Use Office Containers** options. Click the **Configure** button to configure settings:

- **Users and Groups** tab: Same as above.
- **Disks** tab: Same as above.
- **Advanced** tab: Same as above.

6. Click the **Configure general settings** button to configure FSLogix settings for all types of containers:

- **App Services** tab: This tab allows you to modify advanced FSLogix registry settings. For more information about these settings, see <https://learn.microsoft.com/en-us/fslogix/reference-configuration-settings?tabs=profiles#app-services-settings>.
- **Cloud Cache** tab: This tab allows you to modify Cloud Cache settings. For more information about these settings, see <https://learn.microsoft.com/en-us/fslogix/reference-configuration-settings?tabs=ccd#fslogix-settings-profile-odfc-cloud-cache-logging>.
- **Logging** tab: This tab allows you to modify logging settings for profile containers. For more information about these settings, see <https://learn.microsoft.com/en-us/fslogix/reference-configuration-settings?tabs=logging#fslogix-settings-profile-odfc-cloud-cache-logging>.

Rebooting a host

When you enable FSLogix for a new host while running the wizard, no additional steps are necessary. On wizard completion, the host is rebooted and is added to

the active load balancing. An existing host must be rebooted manually using the **Tasks > Tools > Reboot** menu option.

Configure managing existing profiles by Parallels RAS

This topic describes how to configure existing FSLogix Profile Containers to be managed by Parallels RAS. FSLogix Profile Container configuration defines how and where the profile is redirected. Normally, you configure profiles through registry settings and GPO. Parallels RAS gives you the ability to configure profiles from the Parallels RAS Console or RAS Management Portal without using external tools.

Before you begin

Before you configure FSLogix Profile Containers in Parallels RAS, make note of the following:

- You don't have to change the profiles themselves; existing profiles stay the same.
- You can keep using your existing FSLogix Profile Container locations, such as SMB network shares or Cloud Cache.

Preliminary steps

Perform the following preliminary steps:

1. Back up your existing profiles. It is highly unlikely that profile data can be lost or corrupted, but it is best practice to have a valid backup prior to any change in profile configuration.
2. Turn off the GPO configuration of FSLogix Profile Containers. This step is important because you cannot have both GPO and Parallels RAS management of FSLogix profiles enabled at the same time.
3. Before configuring FSLogix profiles for a host in a RAS Farm, make sure there are no user sessions running on the host. As a suggestion, you can make the transition in a maintenance window out of working hours.

Replicate GPO and FSLogix configuration

To configure existing FSLogix Profile Containers in Parallels RAS, you need to replicate your existing GPO to the FSLogix configuration in Parallels RAS. This can be done in the Parallels RAS Console or the Parallels RAS Management Portal.

To configure profiles in the RAS Console:

1. Follow the instruction from the [FSLogix Profile Containers](#) section and open the **Disks** tab.
2. In the **Location of profile disks** list box, specify existing SMB or Cloud Cache locations where you keep your FSLogix profiles. Also, specify the profile disk format, allocation type, and default size.
3. Configure the rest of FSLogix settings you may have on your servers, such as user exclusions, folder exclusions, and others.

To configure profiles in the RAS Management Portal:

1. Navigate to **Infrastructure > RD Session Hosts**.
2. Click a host in the list and then click **Properties**.
3. In the middle pane, click **User Profile**.
4. Specify the settings as described in steps above for the RAS Console.

Please note that at the time of this writing RAS Management Portal can only be used to configure RD Session Hosts to use FSLogix Profile Containers. For other host types, please use the desktop-based RAS Console.

Recommendations and testing

When performing steps in the previous section, do not configure multiple (or all) servers in a RAS Farm right away. Begin with a single server (e.g. an RD Session Host) and then test it with a single user connection. After that, configure some other servers and test the same user logging in to multiple servers consecutively to confirm the profile is loaded and personalization is retained irrespective of a session host. If all is good, configure other hosts, host pools, or Site defaults.

Your RAS users can now connect to Parallels RAS using pre-existing FSLogix Profile Containers, which are now managed centrally through Parallels RAS.

FSLogix antivirus exclusions

Make sure to configure the following antivirus exclusions for FSLogix Profile Container virtual hard drives. Make sure to check the following information with your security team.

Exclude files:

- %Programfiles%.sys
- %Programfiles%.sys
- %Programfiles%.sys
- %TEMP%*.VHD
- %TEMP%*.VHDX
- %Windir%*.VHD
- %Windir%*.VHDX
- \.file.core.windows.net**.VHD
- \.file.core.windows.net**.VHDX

Exclude processes:

- %Programfiles%.exe
- %Programfiles%.exe
- %Programfiles%.exe

When configuring optimizations, you can specify files and processes to exclude in the Windows Defender ATP category. For more information, please see [Optimization](#).

FSLogix support requirements

FSLogix follows Microsoft's Modern Lifecycle Policy. To remain eligible for support, customers must stay current with the latest FSLogix release. Updates and fixes are only provided for the most recent version. If issues are encountered on an older version, customers are required to upgrade before receiving support.

Support for FSLogix is available only to customers who meet product [eligibility requirements](#) and run the software on supported operating systems. The following operating systems are currently supported:

- Windows 10 (including multi-session SKUs from Azure Marketplace)
- Windows 11 (including multi-session SKUs from Azure Marketplace)

- Windows Server 2016
- Windows Server 2019
- Windows Server 2022
- Windows Server 2025

For a complete list of supported platforms, see:

<https://learn.microsoft.com/en-us/fslogix/troubleshooting-fslogix-product-support>

Deprecated Features and Retired Platforms

- 32-bit architecture: Support retired as of February 11, 2025.
- Windows Server 2012 R2: Support retired as of February 11, 2025.

For a complete list of FSLogix Deprecated Features and Retired Platforms, see:

<https://learn.microsoft.com/en-us/fslogix/overview-feature-deprecation>

Optimization

Beginning with version 18, Parallels RAS includes built-in automated optimization capabilities for RD Session Hosts, VDI, and Azure Virtual Desktop workloads. Different preconfigured optimizations for multi-session (such as RD Session Hosts) or single-session (such as VDI) hosts are available for administrators to choose from manually or automatically to ensure a more efficient, streamlined and improved delivery of virtual apps and desktops.

Preconfigured optimizations were designed to be easily updated to support future releases of Microsoft Windows. Moreover, custom scripts may also be used within the tool to make use of already available optimizations to be deployed on Parallels RAS workload machines.

Over 130 image optimizations are available out-of-the-box and divided into the following main categories:

- UWP application packages (removal; available for VDI only)
- Windows Defender ATP (turn ON or OFF, disable real-time scan, exclude files, folder, processes, and extensions)
- Windows components (removal)
- Windows services (disable)
- Windows scheduled tasks (disable)
- Windows advanced options (Cortana, system restore, telemetry, custom layout)

- Network performance (disable task offload, ipv6, etc.)
- Registry (service startup timeout, disk I/O timeout, custom, etc.)
- Visual effects (best appearance, best performance, custom)
- Disk cleanup (delete user profiles, image cleanup, etc.)
- Custom scripts (.ps1, .exe, .cmd, and other extensions/formats)

For the complete list of optimization categories and components, please see <https://kb.parallels.com/125222>.

Optimizations are applicable to RD Session Hosts, VDI desktops, Azure Virtual Desktop, and Remote PC pools (through VDI) based on:

- Windows Server 2012 R2 and later
- Windows 7 SP1
- Windows 10
- Windows 11

Configure optimization

Optimization can be configured for the following:

- RD Session Hosts
- VDI
- Azure Virtual Desktop

Optimization settings are configured for the above on the Site level (Site defaults) and can also be configured for individual components if the RAS administrator decides to use custom settings for a given component.

To configure optimizations on the Site level, navigate to **Farm > Site**, click the **Tasks > Site defaults** menu and choose one of the following:

- **RD Session Host**
- **VDI**
- **AVD multi-session hosts**
- **AVD single-session hosts**

In a Site defaults dialog that opens, select the **Optimization** tab. The user interface for configuring optimization is the same for all of the above.

Note: Before applying optimization, make sure you have a saved state of session hosts as you will not be able to revert changes after they are applied. endhint %}

To configure optimization:

1. If you are in the host **Properties** dialog or in a wizard, clear the **Inherit default settings** options if you want to modify them for this host.
2. Select the **Enable optimization** option.
3. Choose optimization type from the following:
 - **Automatic**: Predefined and preconfigured optimization will be used automatically.
 - **Manual**: Gives you full control over which optimization options to use and allows you to configure each one. This option also gives you an option to use a custom optimization script that will be executed on the host.
4. If you selected **Manual** in the previous step, configure optimization categories and components according to your requirements. See **Configure optimization** below.
5. **Force optimization on all enabled categories**: This is a special option that should only be used in situations when some parts of optimization failed to apply to a host for some unforeseen reason (e.g. the host went offline unexpectedly). When you select this option, then click **OK** and then **Apply** in the RAS Console, the entire optimization configuration will be applied to the host. This way you can make sure that changes that you made to optimization components last time, and that were not applied to the host, will be applied again. The state of the **Force optimization on all enabled categories** option (selected or cleared) is not saved because this is a one-time action, so the next time you open the dialog, the option will be cleared again. Note that in a standard scenario, when you make changes and then apply them to a host, you don't need to select this option, because normally you want to apply just the changes that you made, not the entire optimization configuration.
6. The **Category** list contains optimization categories that can be configured. To include a category in optimization, select the corresponding checkbox. Some categories contain multiple components, which can be configured individually, some have settings that can be customized. To configure category settings or components, highlight the category and click the gear icon (or click **Tasks > Properties**, or simply double-click a category). Depending on the category selected, you can do the following:
 - Configure category settings (choose from available options, select or clear individual settings, specify values, add or remove entries).
 - Add or remove underlying components to include or exclude them from optimization (use the plus- and minus-sign icons). When adding a

component (where available), you can select from a predefined list or you can specify a custom component.

- In some cases (specifically registry entries) you can double-click an entry and specify multiple values for it.
- If you remove a predefined component, you can always get it back in the list by clicking **Tasks > Reset to default**. You can also use this menu to reset category settings to default values if they were modified.
- The last optimization category in the list is **Custom script**. You can use it to execute an optimization script that you may have available. Read the **Using custom script** subsection below for details.

7. When done, click **OK** to close the dialog.

Using custom script

The **Custom script** optimization category is used to execute an optimization script on a target host. Before configuring this category, make sure that the script exists on target hosts and that the path and file name are the same on each host.

To configure the Custom script optimization:

1. Enable the **Custom script** category in the list (select the checkbox), then highlight it and click **Tasks > Properties**.
2. In the dialog that opens, specify the command to execute, arguments (if required), the initial directory, and credentials that will be used to execute the script.
3. Click **OK**.

When you apply the optimization to a host, the script will be executed as part of applying other optimization parameters.

Applying optimization

After you enable optimization for a host and then click **Apply** in the RAS Console, the following will happen the next time the host communicates with Parallels RAS:

1. The host status changes to **Optimization pending** and the host enters the drain mode. At this stage, you can stop optimization by selecting a host in the list and clicking **Tasks > Stop optimization**.
2. Once all users are logged off, the host status changes to **Optimization in progress**.
3. After all optimization settings are applied, the host will reboot.

4. After the reboot, the host returns to operation and its status changes to **OK**.

Note: By design, the host will be rebooted after optimization completion even if it is failed.

Optimization results are logged on a host at the following location:

%ProgramData%.log. Open the file and search for entries similar to the following:

- [I 78/00000009/T10C4/P0FD4] 11-30-20 10:09:19 - Image Optimization completed with 98 successful and 0 unsuccessful optimizations.

Note: By design, Optimization has less priority than Reboot/Disable schedule. For example, it is expected if a host changes the status from “Optimization pending” to Disabled/Reboot when schedule starts.

Upgrade

When Parallels RAS is upgraded from an older version:

- The optimization feature is disabled.
- The inheritance is off.

To use optimization after the upgrade, the administrator needs to enable it manually either in Site defaults or in the host pool/host pool settings.

Inheritance

Component	Optimization	Inherits from
RDSH Site defaults	Yes	None
RDSH Host pool	No	None
RDSH standalone	Yes	RDSH Site defaults
RDSH template	Yes	RDSH Site defaults
RDSH from template	No	None
VDI Site defaults	Yes	None
VDI Desktop standalone	Yes	VDI Site defaults
VDI Desktop template	Yes	VDI Site defaults
VDI Desktop from template	No	None
Azure Virtual Desktop Site defaults	Yes	None
	No	None

Component	Optimization	Inherits from
Azure Virtual Desktop host pool - hosts from a template		
Azure Virtual Desktop host pool - standalone hosts	Yes	AVD multi-session hosts Site defaults or AVD single-session hosts Site defaults.
Azure Virtual Desktop template	Yes	AVD multi-session hosts Site defaults or AVD single-session hosts Site defaults.
Azure Virtual Desktop hosts from template	No	None

Additional information

Please note the following:

- Some optimizations may fail and generate warnings if they had been already applied.
- Some optimizations may fail and generate warnings depending on OS specifics. For example, removal of UWP apps may fail because apps are already absent.

Drive redirection cache

This topic explains the **Enable drive redirection cache** option, which is available in a dialog where you configure RAS RD Session Host, VDI, Azure Virtual Desktop, or Remote PC agents. When the option is enabled, browsing folders on redirected drives becomes much faster thanks to the caching mechanism explained below.

Native RDP is not efficient for file and folder enumeration when using drive redirection, which results in slow and sluggish user experience. The **Enable drive redirection cache** option forces the session host to run the kernel-based driver (RasRdpFs). This optimizes how the communication is carried out compared to standard RDP and also adds caching of the folder structure on the session host (RDSH, VDI, or Azure Virtual Desktop). The driver starts as soon as the setting is pushed to the session host via **Apply** in the RAS Console. When this happens, all

new sessions will have this functionality enabled. The existing sessions need to be reconnected to use this optimization.

Notes

- A session host must run a 64-bit operating system.
- The cache is per session and is paged into the driver memory.
- On log off or disconnect, the cache is purged.
- If the number of cached folders in the session exceeds the threshold, and the user accesses a new non-cached folder, then the oldest accessed folder is replaced in the cache.
- When the option is switched off, all currently active user sessions will lose the cache (the driver is stopped and the cache is purged). This happens transparently to the user, but file and folder enumeration become slow.
- When the option is switched on, all currently active user sessions will not automatically have the cache enabled. To use this functionality, the existing sessions will need to be reconnected.

Limitations

- The option is applicable only to sessions initiated by the following versions of Parallels Client:
 - Parallels Client for Windows versions 18 and later
 - Parallels Client for macOS versions 19 and later
- Similar to native RDP changes made on the client side (in a remote session), requires manual refresh (F5) in a redirected folder on the server side.

Configure logging

An RD Session Host is monitored and logs are created containing relevant information. To configure logging and retrieve or clear existing log files, right-click a host, choose **Troubleshooting > Logging** in the context menu, and then click **Configure**, **Retrieve**, or **Clear** depending on what you want to do. For the information on how to perform these tasks, see the [Logging](#) section.

Manage sessions (RD Session Host)

Please see [Session Management](#).

Using scheduler (RD Session Hosts)

The **Scheduler** tab allows you to create scheduler tasks that will be performed on individual hosts or host pools at a specified time.

Note: When the scheduled event is triggered, affected hosts are disabled in Parallels RAS and their status is displayed as “Disabled (scheduler)” or “Pending reboot (scheduler)”. You can cancel these states by right-clicking a host on the **Hosts** tab and choosing **Control > Cancel disabled state (scheduler)** or **Control > Cancel pending reboot (scheduler)**.

Disabling hosts and hosts in pools

To disable a host or a host in a pool:

1. Click **Tasks > Add > Disable host** or **Disable host pool**.
2. On the **General** tab, select the **Enable Schedule** option.
3. Specify a name for this schedule and an optional description.
4. Select a host or a pool in the **Available** list and click **Add**. The host or pool will appear in the **Target** list.
5. Select the **Trigger** tab and specify start date and time, duration, and recurrence settings for this event. To make this a one-time event, select **Never** in the **Recur** drop-down list.
6. Select the **Options** tab. It contains the following options:
 - **Message list:** Configure a message that will be sent to users before the host goes offline. Click **Tasks > Add** and specify the message title, body, and the time period when it should be sent.
 - **On disable:** Specify what should happen to current sessions when a scheduled task triggers. Select the desired option from the **On disable** drop-down list.
 - **Enforce schedule for currently inactive hosts:** This option is only enabled when you have an active message in the list. If the option is enabled, hosts that are currently offline are also monitored, and if such a server comes back online during the scheduled task execution, the task is applied to it too.

7. Click **OK** to save the schedule.

Rebooting hosts and hosts in pools

To reboot a host or a host in a pool:

1. Click **Tasks > Add > Reboot host** or **Reboot host pool**.
2. On the **General** tab, select the **Enable Schedule** option.
3. Specify a name for this schedule and an optional description.
4. Select a host or a pool in the **Available** list and click **Add**. The host or pool will appear in the **Target** list.
5. Select the **Trigger** tab and specify start date, time, and recurrence settings for this event. To make this a one-time event, select **Never** in the **Recur** drop-down list. In addition, specify the following options for the “Reboot host pool” task:
 - **Condition:** Trigger only if a session state is Down can be used to clear sessions that are stuck in a Down state. Not needed under normal circumstances.
6. Select the **Options** tab. It contains the following options:
 - **Message list:** Configure a message that will be sent to users before the host is rebooted. Click **Tasks > Add** and specify the message title, body, and the time period when it should be sent.
 - **Enable Drain Mode** and **Force server reboot after:** The two options work together. If you select the **Enable Drain Mode** option, then when the task triggers, new connections to a host are refused but active connections will continue to run and can be reconnected. The server will be rebooted when all active users close their sessions or when **Force server reboot after** time is reached, whichever comes first. For active users not to lose their work, create a message that will advise them to save their work and log off.
 - **Enforce schedule for currently inactive hosts:** This option is only enabled when you have an active message in the list. If the option is enabled, hosts that are currently offline are also monitored, and if such a server comes back online during the scheduled task execution, the task is applied to it too.
7. Click **OK** to save the schedule.

Starting up hosts and hosts in pools

Note: This task applies only to hosts and host pools based on a template.

To start up a host or a host in a pool:

1. Click **Tasks > Add > Startup host** or **Startup host pool**.
2. On the **General** tab, select the **Enable Schedule** option.
3. Specify a name for this schedule and an optional description.
4. Select a host or a pool in the **Available** list and click **Add**. The host or pool will appear in the **Target** list.
5. Select the **Trigger** tab and specify start date, time, and recurrence settings for this event. To make this a one-time event, select **Never** in the **Recur** drop-down list.
6. (“Startup host pool” task only) Select the **Options** tab. It contains the following options:
 - **Percentage of members**: Select this option to specify the percentage of hosts that must be started up in each pool.
 - **Specific number of members to be started**: Select this option to specify the number of hosts that must be started up in each pool.
7. Click **OK** to save the schedule.

Shutting down hosts and hosts in pools

To shut down a host or a host in a pool:

1. Click **Tasks > Add > Shutdown host** or **Shutdown host pool**.
2. On the **General** tab, select the **Enable Schedule** option.
3. Specify a name for this schedule and an optional description.
4. Select a host or a pool in the **Available** list and click **Add**. The host or pool will appear in the **Target** list.
5. Select the **Trigger** tab and specify start date, time, and recurrence settings for this event. To make this a one-time event, select **Never** in the **Recur** drop-down list.
6. Select the **Options** tab. It contains the following options:
 - **Message list**: Configure a message that will be sent to users before the host is rebooted. Click **Tasks > Add** and specify the message title, body, and the time period when it should be sent.
 - **Enable Drain Mode** and **Force server shutdown after**: The two options work together. If you select the **Enable Drain Mode** option, then when the task triggers, new connections to a host are refused but active connections will continue to run. The server will be shut down when all active users close their sessions or when **Force server reboot after** time is reached, whichever comes first. For active users

not to lose their work, create a message that will advise them to save their work and log off.

- **Enforce schedule for currently inactive hosts:** This option is only enabled when you have an active message in the list. If the option is enabled, hosts that are currently offline are also monitored, and if such a server comes back online during the scheduled task execution, the task is applied to it too.

Recreating hosts and host pools

Note: This task applies only to hosts and host pools based on a template.

To recreate a specific host or all hosts in a host pool:

1. Click **Tasks > Add > Recreate host from template** or **Recreate host pool from template**.
2. On the **General** tab, select the **Enable Schedule** option.
3. Specify a name for this schedule and an optional description.
4. Select a host or a pool in the **Available** list and click **Add**. The host or pool will appear in the **Target** list.
5. Select the **Trigger** tab and specify start date, time, and recurrence settings for this event. To make this a one-time event, select **Never** in the **Recur** drop-down list.
6. Select the **Options** tab. It contains the following options:
 - **Message list:** Configure a message that will be sent to users before the host is rebooted. Click **Tasks > Add** and specify the message title, body, and the time period when it should be sent.
 - **Enable Drain Mode, Force host recreation after, and Force host pool recreation after:** The options work together. If you select the **Enable Drain Mode** option, then when the task triggers, new connections to a host are refused but active connections will continue to run. The server will be recreated when all active users close their sessions or when the time specified in **Force host recreation after** or **Force host pool recreation after** is reached, whichever comes first. For active users not to lose their work, create a message that will advise them to save their work and log off.
 - **Enforce schedule for currently inactive hosts:** This option is only enabled when you have an active message in the list. If the option is enabled, hosts that are currently offline are also monitored, and if such

a server comes back online during the scheduled task execution, the task is applied to it too.

Sending multiple messages to users

For **Disable Host** and **Disable Host Pool** tasks, you can only send a message before the scheduled task is triggered. Hence, when creating a message, you can only select the “before” option when specifying when the message should be sent. You can create more than one message if needed and send them at different time intervals, so the users are notified more than once before the task executes.

For **Reboot Host** and **Reboot Host pool** tasks, you can send a message before or after the scheduled task is triggered. The “after” option is available for these tasks because you have the ability to enable the drain mode, which will keep the active sessions running for some time. During this time, you can send multiple messages to active users reminding them that they should finish their work and close their sessions. To use the “after” option, the **Enable Drain Mode** option must be selected. Please also note that the “after” time interval and the **Force server reboot after** setting should be coordinated. For example, if the force reboot occurs before the “after” time elapses, active users will not have a chance to see the message.

Planning for high availability

When adding RD Session Hosts to a Site, the N+1 redundancy approach should be used to ensure uninterrupted service to your users. This is a general rule that also applies to other Parallels RAS components, such as Connection Brokers, RAS Secure Gateways, or possibly Providers.

Managing logons

The logon management feature allows you to enable or disable logons from RD Session Hosts. The feature performs the same tasks as the `change_logon` command-line utility.

Attention: When autoscaling is enabled, the following control options are unavailable because manual intervention could cause the autoscaler to start operating unexpectedly.

Note: For RD Session Hosts based on a template, the drain mode (which disables logons) is handled automatically by the group to which a host belongs. For more information, see [Using Scheduler](#).

To manage logons:

1. In the Parallels RAS Console, navigate to **Farm** > <Site> > **RD Session Hosts**.
2. Select an RD Session Host, click **Tasks** > **Control**, and choose one of the following:
 - **Enable logons:** Enables logons from client sessions, but not from the console. This option performs the same action as the `change logon /enable` command.
 - **Disable logons:** Disables subsequent logons from client sessions, but not from the console. Does not affect currently logged-on users. This option performs the same action as `change logon /disable` command.
 - **Drain:** Disables logons from new client sessions, but allows reconnections to existing sessions. Drain is kept even after reboot until the admin enables logons.

Note that while a host is in drain mode, administrators may still log on to the physical console or remotely log on using the `/admin` or `/console` command-line option for MSTSC. This allows administrators to remotely maintain the RDS host via **Tools** > **Remote Desktop**.

- **Drain until reboot:** Disables logons from new client sessions until the computer is restarted, but allows reconnections to existing sessions. Drain is kept until the host is restarted. Same action as the `change logon /drainuntilrestart` command.

To see the current logon control mode for an RD Session Host, click **Tasks** > **Control**. The checked-out option indicates the current logon control mode of the selected RD Session Host. To do this check from the command line, execute the `change logon /QUERY` command on the host.

Please also note the following:

- When applying a logon control mode on a host, ensure that the agent status is updated accordingly.

- You must set the logon control options for the hosts one by one. If you need to do it for a group of hosts, you can use the scheduler (see [Using scheduler](#)).
- `/Drain` disables logons from new client sessions, but allows reconnections to existing sessions. `Drain` differs from `Drainuntilrestart` as far as it requires administrator intervention, whereas the latter enables logons automatically after a restart.
- **Computer Configuration / Administrative Templates / Windows Components / Remote Desktop Services / Remote Desktop Session Host / Connection / Allow users to connect remotely using remote desktop services** must be set to **Not configured**, otherwise, it takes precedence.

Using computer management tools

You can perform standard computer management tasks on an RD Session Host right from the RAS Console. These include Remote Desktop Connection, PowerShell, Computer Management, Service Management, Event Viewer, IPconfig, Reboot, and others. To access the **Tools** menu, select a host, click **Tasks** (or right-click) > **Tools** and choose a desired tool. For requirements and usage information, see [Computer management tools](#).

Publishing from an RD Session Host

See [Publishing](#).

You can also publish resources using a publishing wizard in the **Start** category, as described in the [Set up a basic Parallels RAS Farm](#) section. The **Start** category publishing wizard is a simplified version that gives you convenient options of selecting the resources that you want to publish. You may try both approaches and choose the one that better suits your needs.

Viewing published resources

When you want to remove an RD Session Host or an RD Session Host group from a Site, you might want to see the list of published resources hosted by the host or

host pool. This way you can see which resources will be affected. You can do so as follows:

1. In the Parallels RAS Console, select **Farm RD Session hosts**.
2. To see published resources for a specific RD Session Host, select the **RD Session hosts** tab. To see published resources for a group, select the **Groups** tab.
3. Right-click a host or a host group and choose **Show published resources** (or click **Tasks > Show published resources**).
4. The **Published Resources** window opens displaying the list of published resources for the selected host or host pool. Resource information includes:
 - **Name**. Resource name.
 - **Status**. Enabled or disabled.
 - **Type**. “Application” is used for published applications, URLs, network folders, etc. “Desktop” is used for published desktops.
 - **Path**. For published applications, specifies a path to the execute file, URL, or UNC path.
 - **Parameters**. Published application parameters (if any).
 - **Published from**. Site, host pool(s), or host(s).
5. To refresh the list, press F5 or click the “recycle” icon (top-right).
6. To filter the list, press Ctrl-F or click the magnifying glass icon and then specify the filter criteria for desired column(s).

Virtual Desktop Infrastructure (VDI)

Parallels RAS VDI (Virtual Desktop Infrastructure) enables you to use host virtualization to reduce the number of physical host required to host published resources. Parallels RAS VDI supports numerous virtualization technologies, including hypervisor and cloud-based platforms.

Parallels RAS VDI also includes the Template functionality, which gives you the ability to create a template from a preconfigured host (virtual machine) and then automatically clone hosts and RD Session Host VMs from it.

Supported providers

Parallels RAS supports hypervisor-based Providers and cloud-based Providers.

Hypervisors

The following hypervisors are supported:

- Microsoft Hyper-V (Windows Server 2016 up to Windows Server 2025)
- Microsoft Hyper-V Failover Cluster (Windows Server 2016 up to Windows Server 2025)
- VMware vCenter 7.x*, 8.0
- VMware ESXi 7.x*, 8.0
- SC//HyperCore 9.4, 9.6
- Nutanix AHV (using Prism Element)** AOS 6.10 LTS, AOS 7.3
- Nutanix AHV (using Prism Central) p.c. 7.3, p.c. 7.5
- Remote PC — This is a special type that allows you to create pools of Remote PCs. See [Remote PC pools](#).

* VMware ends support of vSphere 7 on October 2, 2025. While you can still use this version with Parallels RAS 21, it is recommended to upgrade to vSphere 8.0 or newer to ensure long-term support.

** Nutanix connection via Prism Element is being deprecated by Nutanix in October 2026. Therefore the Provider using Prism Central should be used by all customers moving forward.

Cloud Providers

- Microsoft Azure
- Amazon Web Services

Add a provider

RAS Provider Agent information

In order to function in a RAS Farm, a Provider (hypervisor or cloud-based) needs RAS Provider Agent to be installed in the Farm. RAS Provider Agent acts as an interface between other RAS components and a Provider. RAS Provider Agent conducts all communications with a Provider through the provider's native API.

Parallels RAS has two types of RAS Provider Agents that can be installed in a Farm:

- **Built-in:** This RAS Provider Agent is built into the RAS Connection Broker and is installed automatically when you install Parallels RAS. The agent can handle multiple Providers and can also be configured for high availability.
- **Dedicated:** This RAS Provider Agent is installed manually. It can handle only a single Provider. If you want to use this agent type with more than one provider, you need to install a separate instance for each provider.

Both built-in and dedicated RAS Provider Agents are compatible with all types of Providers supported by Parallels RAS. Which agent you choose to install depends only on your requirements. When possible, it is always recommended to use the built-in Provider Agent for high availability and business continuity.

What to read next:

- If you are adding a Provider that will use the built-in RAS Provider Agent, you may skip to [Add a Provider](#).
- If you want to install a dedicated RAS Provider Agent on a host of your choice, read the [RAS Provider Agent installation options](#) section, which follows this one.

RAS Provider Agent installation options

If you are installing a dedicated RAS Provider Agent, you first need to determine where it will be installed. Depending on the Provider type, the following options are available:

- The host on which the hypervisor is running. This option is available for Microsoft Hyper-V only.
- A supported version of Windows Server running on a physical box or in a virtual machine. For supported Windows Server versions, see **Software requirements > RAS Provider Agent**.

The following table lists RAS Provider Agent installation options for each supported Provider:

Provider	Built-in Agent (part of PA)	Agent on a Provider	Agent on a Windows Server (VM or HW)
Microsoft Hyper-V	Yes	Yes	Yes*
Microsoft Hyper-V Failover Cluster	Yes	No	Yes*
VMware VCenter	Yes	No	Yes*
VMware ESXi	Yes	No	Yes*
SC//HyperCore	Yes	No	Yes*
Nutanix AHV (AOS)	Yes	No	Yes*
Nutanix Prism Central	Yes	No	Yes*
Remote PC (see the Note below)	Yes	No	Yes*
Microsoft Azure	Yes	No	Yes*
Amazon Web Services	Yes	No	Yes*

* High Availability is not available with these Provider Agent installation options. For details, see [Enabling high availability for VDI](#).

Note: The **Remote PC** is a special type that can be used to create and manage pools of Remote PCs as part of hosted desktop infrastructure (HDI). When you add a Provider of this type, you can manage it like one of the real Providers with some limitations, such as you cannot create templates and use some other strictly VDI-specific functions. The main feature when using this type is the ability to create pools of HDI-based Remote PCs (e.g. HPE Moonshot System, Atrust Remote PC Array) and making PCs persistent by assigning an individual PC to a specific user. For more info, see [Remote PC pools](#).

In the table above, find the Provider type that you are using and see where the RAS Provider Agent can be installed. Depending on the available choices, do one of the following:

- **Built-in Agent:** The agent is a part of RAS Connection Broker, so it is already installed. When possible, it is always recommended to use the built-in Provider Agent for high availability and business continuity.
- **Agent on a the provider:** This option is only available if you are using Microsoft Hyper-V. You can simply install the agent on the host, as described in [Add a Provider](#).
- **Agent on a Windows Server (VM or HW):** To use this option, make sure you have a physical box or a virtual machine running a supported version of

Windows Server. You will need to specify its FQDN or IP address when adding a Provider to the Farm.

Add a hypervisor provider

This section describes how to add a [hypervisor-based Provider](#). For the information on how to add a cloud-based Provider, see [Add a cloud Provider](#).

To add a Provider:

1. In the RAS Console, navigate to **Farm > Site > Providers**.
2. On the **Providers** tab, click **Tasks > Add** and select the provider you want to add.
3. The **Add Virtualization Provider** wizard opens.
4. In the **Name** field, specify the name for the provider.
5. In the **Description** field, type an optional description.
6. In the **Address** field, specify the host's FQDN or IP address. For SC//HyperCore, you can specify IP addresses for several nodes. For Nutanix Prism Central, specify the Address of Prism Central, not the individual clusters.
7. Specify a user name and password to log in to the host.
8. Click the **Manage Credentials** button to specify the accounts that will be used to deploy RAS agents.
9. Click the **Advanced Settings** link to open the **Advanced Provider Settings** dialog. The dialog allows you choose the following options:
 - **Use dedicated Provider Agent:** Select this option if you will install (or have installed) the RAS Provider Agent yourself. Clear the option if you will use the [built-in RAS Provider Agent](#).
 - **Agent address:** This option becomes enabled if you select the option above it. Specify the FQDN or IP address of the host where the RAS Provider Agent is (or will be) installed. This can be either a physical box or virtual machine.

- **Preferred Connection Broker:** Select a RAS Connection Broker to be the preferred agent for this Provider. For more info, see [Enabling high availability for VDI](#).

10. Click **Next**.

11. The wizard will now try to connect to the RAS Provider Agent. If you specified **Use dedicated Provider Agent** option in the previous (optional) step, but haven't installed the agent yet, click **Install** and follow the instructions to push install the agent on the specified host.

Please note that for the remote installation to work, the following requirements must be met:

- The firewall must be configured on the host to allow push installation. Standard SMB ports (139 and 445) need to be open. See also **Port reference** for the list of ports used by Parallels RAS.
- SMB access. The administrative share (\\$) must be accessible. Simple file sharing must be enabled.
- Your Parallels RAS administrator account must have permissions to perform a remote installation on the host. If it doesn't, you'll be asked to enter credentials of an account that does.
- The target host should be joined to an AD domain.

If push installation cannot be performed for any reason, you can install the agent manually using the installer. See [Installing RAS Provider Agent using the installer](#).

12. If you've selected **Microsoft Hyper-V Failover Cluster** as the Provider type, the page opens where you can disable MAC address management for hosts. Note that you should only do it if you are using Microsoft System Center Virtual Machine Manager (SCVMM) or other solution to manage MAC addresses. See the explanation below.

MAC address management is required when using Microsoft Hyper-V Failover Cluster as a Provider. This is to avoid duplicate MAC addresses, which may occur when a host is migrated to a different node in the cluster and the MAC address is released and reused on the original node. If that happens, such a host can no longer be managed in a Farm. Parallels RAS uses a pool of static MAC addresses at the Provider level to automatically generate and assign MAC addresses to hosts. This way, when a host is migrated to a different node in the cluster, its MAC address will not be reused

for a different VM and no duplicate MAC addresses will occur. The pool has 10,000 reserved MAC addresses in the range displayed in the **Starting MAC address** and **Ending MAC address** fields on the wizard page.

As was said above, if you are already managing MAC addresses using SCVMM or other solution, clear the **Enable MAC address management** option.

13. Click **Next**.

14. If you've selected **VMware vCenter** as the Provider, another page opens (the page will not open for any other host type). On this page, you can specify a vCenter resource pool. This allows you to enumerate VMs by selecting a cluster (root resource pool) or an individual resource pool within a cluster. To choose a resource pool, select the **Use specific resource pool** option and then click the [...] button next to the **Resource Pool** field. In the dialog that opens, select a desired resource pool. Note that if you leave the **Use specific resource pool** option cleared, all VMs from the entire vCenter cluster will be retrieved (max number is 35,000). Click **OK** when done.

15. Click **Finish** to close the wizard.

Add a cloud Provider

This section describes how to add a [cloud-based Provider](#). For the information on how to add a hypervisor provider, see [Add a hypervisor Provider](#).

Microsoft Azure

Introduction and prerequisites

Introduction

Organizations using or interested in using Microsoft Azure can provision, scale, and manage VDI and RD Session Host workloads directly from the Parallels RAS console and deploy onto Microsoft Azure using Azure Resource Manager (ARM). Parallels RAS uses a service principal with required permissions on relevant Azure

resources (subscription and resource groups) to authenticate, provision, and manage the resources.

Prerequisites

To use Microsoft Azure as a Provider, you need the following:

- An existing Microsoft Azure account and subscription.
- The necessary Microsoft Azure providers must be enabled, including Microsoft.ResourceGraph, Microsoft.Resources, Microsoft.Compute, Microsoft.Network.
- An ARM virtual network and subnet in your preferred region with connectivity to AD services. Microsoft Entra ID with Active Directory Domain Services (AADDS), Domain Controller in Azure IAAS or hybrid with connectivity to an on-premises domain can be used.
- Site-to-site VPN or ExpressRoute is required if a hybrid RAS deployment is used.
- A configured VM to be used for VDI or RD Session Host as a template.
- Effective March 31, 2026, Microsoft will have/has retired default outbound access. The Subnet and VNet you use must explicitly allow outbound connectivity. For more information, see [Azure's outbound connectivity methods](#) for the most commonly used approaches.

Adding Microsoft Azure as a Provider is a two-step process:

1. First, you need to create an application in Microsoft Azure to access the resources in your subscription. This step is described in the [Create a Microsoft Entra ID application](#) section.
2. Once the application is created and registered, you can add Microsoft Azure as a Provider in the Parallels RAS Console. This step is described in [Add Microsoft Azure as a Provider](#).

Read on to learn how to perform the steps above.

Create Microsoft Entra ID application

To complete the steps below, you must have a Microsoft Azure subscription and account. If you don't have a subscription, you need to purchase one first.

Create an Microsoft Entra ID application

An Microsoft Entra ID application is used with the role-based access control. You need to create an Microsoft Entra ID application to access resources in your subscription from Parallels RAS.

To create an Microsoft Entra ID application:

1. Log in to the Microsoft Azure portal.
2. Open the portal menu and select **Microsoft Entra ID**.
3. In the left pane, select **App registrations**.
4. Click **New registration** (at the top of the right pane).
5. The **Register an application** blade opens.
6. In the **Name** field, type a name you want to use for the application.
7. In the **Redirect URI (optional)** section, make sure that **Web** is selected in the drop-down list. Leave the URI field empty.
8. Click **Register** (at the bottom left).
9. The new Microsoft Entra ID app is created and its blade is displayed in the portal.

Note the following app properties, which are displayed at the top of the right pane:

- **Display name**
- **Application (client) ID***
- **Directory (tenant) ID***
- **Object ID***

* Copy and save these properties. You will need to specify them later when adding Azure as a Provider in the RAS Console.

Create a client secret

A client secret is a string that the application uses to prove its identity when requesting a token. It essentially acts as an application password. You will need to specify this string in the RAS Console when adding Azure as a Provider.

To create a client secret:

1. If you are not on the application page anymore, navigate to it from the **Home** page by selecting **Microsoft Entra ID > App registration** and then clicking the app in the right pane.
2. In the left pane, click **Certificates & secrets**.

3. In the right pane, click **New client secret**.
4. Type a client name and select a desired expiration option.
5. Click **Add**. The new client secret appears in the **Client secrets** list.
6. **IMPORTANT:** Copy and save the client secret (the **Value** column). If you leave this page without copying the secret, it will be hidden and you will not be able to retrieve it later.

Give the application read and write access to resources

The Microsoft Entra ID app that you created must have read and write access to Azure resources. The following instructions demonstrate how to give the application read and write access to a resource group. You can also give access to a specific resource or to your entire Azure subscription. For more information, please see the Microsoft Azure documentation.

To give the app write access to the resource group where new VMs will reside:

1. In the Azure portal menu, select **Resource groups**.
2. Click a resource group where the new VMs will reside.
3. In the left pane, select **Access control (IAM)**.
4. In the right pane, locate the **Grant access to this resource** box and click **Add role assignment**.
5. On the **Role** tab of the **Add role assignment** page, select **Privileged administrator roles**, then the **Contributor** role.
6. Click **Next**.
7. On the **Members** tab, select the **User, group, or service principal** option.
8. Click on the **Select members** link and enter the name of the previously created application in the **Select** field. Select the application in the drop-down list and click **Select**.
9. Click **Next**.
10. On the **Review + assign** tab, confirm that the configuration is correct and click **Review + assign**.

To give the app read access to the resource group:

1. Repeat steps 1-4 from the list above.
2. On the **Role** tab of the **Add role assignment** page, select **Job function roles**, then the **Reader** role.
3. Repeat steps 6-10 from the list above.

Note: If you would like to give the application read access to your entire subscription (not just a specific resource groups), select **All services** in the Azure portal menu, then navigate to **Categories > All > Subscriptions** and select your subscription. Select **Access control (IAM)** in the middle pane and click **Add** in the **Add a role assignment** box. Repeat steps 2-4 from the list above.

Finding your Microsoft Azure subscription ID

When you'll be adding Microsoft Azure as a Provider in the RAS Console, you will need to specify your Azure subscription ID. If you don't remember it, here's how to find it in the Microsoft Azure portal:

1. In the portal menu, choose **All services**.
2. In the **Categories** list, click **All**.
3. In the right pane, click **Subscriptions**.
4. Click a subscription and then copy and save the value from the **Subscription ID** field.

Summary

When you complete all of the above steps, you should have the following values saved and ready to be used to add Microsoft Azure as a Provider in the RAS Console:

- **App (client) ID:** Application ID.
- **Directory (tenant) ID:** Tenant ID.
- **Client secret:** Client secret (application key).
- **Subscription ID:** Your Microsoft Azure subscription ID.

Read on to learn how to add Microsoft Azure as a Provider in the RAS Console.

Add Microsoft Azure as a Provider

To add Microsoft Azure as a Provider:

1. In the RAS Console, navigate to **Farm > Site > Providers**.
2. On the **Providers** tab, click **Tasks > Add > Microsoft Azure**.
3. The **Add Cloud Computing** wizard opens.
4. In the wizard, specify the following:
 - **Name:** Name of the provider.

- **Description:** Description of the provider.
 - **Manage credentials:** the administrative accounts that will be used to deploy Parallels Agents.
 - **Authentication URL:** Prepopulated with the Microsoft authentication site URL. Unless otherwise required or indicated, keep the default value provided.
 - **Management URL:** Prepopulated with the Microsoft Azure management site URL. Unless otherwise required or indicated, keep the default value provided.
 - **Resource URI:** Prepopulated with the Microsoft Azure resource URI. Unless otherwise required or indicated, keep the default value provided.
 - **Tenant ID:** The “Directory (tenant) ID” value of the Microsoft Entra ID app that you created earlier.
 - **Subscription ID:** Your Microsoft subscription ID.
 - **Application ID:** The “App (client) ID” value of the Microsoft Entra ID app that you [created earlier](#).
 - **Application key:** The “Client secret” value of the Microsoft Entra ID app that you [created earlier](#).
5. Click the **Advanced Settings** link to open a dialog where you can configure the following optional settings:
- **Use dedicated Provider Agent:** When this option is cleared (default), the built-in RAS Provider Agent will be used. If you want to use a dedicated RAS Provider Agent, select this option and specify the host FQDN or IP address.
 - **Agent address:** This option becomes enabled if you select the option above it. Specify the FQDN or IP address of the host where the RAS Provider Agent is (or will be) installed. This can be either a physical box or virtual machine.
 - **Preferred Connection Broker:** Select a RAS Connection Broker to be the preferred agent for this Provider. For more info, see [Enabling high availability for VDI](#).
6. Click **Next**. The wizard will display the new Provider information and will indicate the RAS Provider Agent status. If everything is OK, click **Finish** to exit the wizard. If something is not as expected, click **Back** and correct any mistakes if necessary.

The new Provider will now appear on the **Providers** tab in the RAS Console. Complete the Provider addition as follows:

1. Click **Apply** to apply the changes.

2. Verify the value of the **Status** column. If it's anything other than **OK**, right-click the Provider and choose **Troubleshooting > Check agent**. Verify the agent status and install it if necessary, then click **OK**. The **Status** column on the **Providers** tab should now say **OK**.

Modifying the Provider configuration

To view and modify the Provider configuration, right-click it and choose **Properties**. In the dialog that opens, view and modify the Provider properties.

Microsoft Azure and templates

When creating a template for cloning VMs in Microsoft Azure, you need to select an Azure resource group where VM clones will be created. Note that this must be a group to which you granted permissions to the Microsoft Entra ID application. You also need to select a VM size and disk type to be used for cloned VMs. These settings are specified on the **Advanced** page of the **Create Template Wizard**.

Both Virtual Desktop and RD Session Host templates can be created with Microsoft Azure as a Provider. When VMs are cloned, you will see them appear in the RAS Console. At the same time, you can also see them in the Microsoft Azure portal.

Note: If there are multiple RAS installations using the same subscription, then the workaround is to change the Provider Agent application read access from subscription level to resource group level or a set of resource groups. This is necessary to avoid a situation when a given Provider Agent intersects with the set of resource groups of another Provider Agent application.

For complete information about creating and using templates, including Microsoft Azure specifics, please see the [Templates](#) section.

Amazon Web Services

Introduction and prerequisites

Introduction

Amazon Web Services (AWS) is a leading cloud platform provider offering over 200 fully featured services from data centers globally. Parallels RAS 19 provides the ability to integrate, configure, maintain, support, and access Amazon EC2 workloads on top of the existing capabilities of Parallels RAS.

Support is targeted at multi-session (RDSH), single session (server-based VDI) server operating systems, and other Microsoft operating systems, if your organization holds licenses for them. For more information about using Microsoft operating systems with AWS, see <https://aws.amazon.com/windows/faq/>.

Parallels RAS Console allows you to do the following:

- Manage Amazon EC2 instances
- Create and manage templates
- Create and manage instance pools
- Configure autoscaling
- Enable, reboot, start up and shut down instances via schedules
- Configure image optimization
- Use FSLogix Profile Container and MSIX app attach
- Change instance types and storage types

Prerequisites

- An AWS account. If you do not already have an account, you can create it for free at aws.amazon.com/ec2/.
- A working Microsoft Active Directory environment to join the Amazon EC2 cloned instances to your domain.
- A preconfigured Virtual Private Cloud (VPC) as your virtual network and security groups that act as a virtual firewall for your EC2 instances.
- A preconfigured Amazon EC2 instance, which will be used later as a Parallels RAS template, running on Windows Server 2012 up to Windows Server 2022.

Design considerations

This section contains design advice that you might want to keep in mind when using AWS in Parallels RAS.

DHCP options set

You might need to use an AWS DHCP options set to specify a custom DNS pointing to the domain controller so that the VMs created from templates are able to join the Active Directory domain. If the custom DNS is not set, the default AWS public DNS will be used, and the VMs won't be able to communicate with the domain controller.

For information on how to configure DHCP options sets, see <https://docs.aws.amazon.com/vpc/latest/userguide/DHCPOptionSet.html>.

The Provider Agent and Guest Agents need to be on the same subnet for the Guest Agent to discover the Provider Agent using broadcasts. If this is not possible, then a registry setting with the IP of the Provider Agent needs to be added on the VM as described here: <https://kb.parallels.com/en/124157?language=en>.

Service quotas

Sometimes solutions scale in usage, invocations, number of instances, and so on. Due to this, the standard AWS service quotas can be reached. For more information about AWS service quotas, see https://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html.

Parallels RAS integrations are subject to the EC2 and EBS endpoint limits as specified here:

- <https://docs.aws.amazon.com/general/latest/gr/ec2-service.html>
- <https://docs.aws.amazon.com/general/latest/gr/ebs-service.html>

Encrypting storage

The storage of clones created from RAS templates will be encrypted if the AWS administrator enables encryption of the RAS template VM storage in AWS Management Console.

Encryption can be enabled by default or explicitly when launching a new EC2 VM:



The screenshot shows the 'Storage (volumes)' configuration page in the AWS Management Console. The 'Encrypted' checkbox is checked, and the 'KMS key' is set to '(default) aws/ebs'. The 'Size (GiB)' is set to 30, and the 'Volume type' is set to gp2. The 'Delete on termination' checkbox is also checked.

Storage type	Device name - required	Snapshot
EBS	/dev/sda1	snap-0c3ab6d4de940dccb

Size (GiB)	Volume type	IOPS
30	gp2	100 / 3000

Delete on termination	Encrypted	KMS key
Yes	Encrypted	(default) aws/ebs Key ID: alias/aws/ebs

For more information about encryption, see <https://aws.amazon.com/blogs/compute/must-know-best-practices-for-amazon-ebs-encryption/>.

Step 1. Creating an IAM user for programmatic access

To create the IAM user account, you can use the AWS Management Console, the AWS CLI, Tools for Windows PowerShell, or AWS API operation. In this example, we will be using the AWS Management Console:

1. Sign in to the AWS Management Console and open the **IAM** page at console.aws.amazon.com/iam.
2. In the navigation pane, choose **Users** and then click the **Add users** button.
3. Under **Set user details** section, provide a user name such as "ParallelsConnector".
4. Under **AWS access type**, select **Access key - Programmatic access**, as the Parallels RAS Console will be using APIs to communicate with your AWS account. This will create an access key for the IAM user. You can view or

download the access keys when you get to the **Final** page. Click **Next** to proceed to the permissions page.

5. On the permissions page, you can create a user group for the new IAM user to be a part of. This is recommended as its beneficial for management purposes, although not mandatory.
6. If you are not using groups, choose **Attach existing policies directly**. A list of the AWS managed and customer managed policies in your account will appear.
7. Filter policies and choose **AmazonEC2FullAccess**, which is an AWS managed preconfigured policy, and click **Next** to proceed to the next page.
8. Optionally, on this page, you can use the tags to organize, track, or control access for this user.
9. Once the tags are ready, click **Next** to see all of the choices you made up to this point. When you are ready to proceed, click **Create user**.
10. To view the user's access key ID and secret access keys, click **Show** next to each password and access key that you want to see. To save the access keys, choose **Download CSV** and then save the file to a safe location.

Please note that this is your only opportunity to view or download the secret access keys.

11. Save the user's new access key ID and secret access key in a safe and secure place to be used next in Parallels RAS Console.

Note: For security reasons, it is recommended to regularly change keys of the IAM user as described in <https://aws.amazon.com/blogs/security/how-to-rotate-access-keys-for-iam-users/>.

Proceed to [Step 2. Adding AWS as a Provider](#).

Step 2. Adding AWS as a Provider

To configure Amazon Web Services as a Cloud Computing provider:

1. In the RAS Console, navigate to **Farm > Providers**.
2. Click the **Tasks** drop-down menu and choose **Add** (or click the **[+]** icon).

3. In the menu, select **Amazon EC2**. The **Add Cloud Computing Provider** wizard opens.
4. In the Wizard, specify the following:
 - **Name:** Name of the provider.
 - **Description:** Description of the provider.
 - **Manage credentials:** the administrative accounts that will be used to deploy Parallels Agents on the session hosts (Amazon EC2 instances). The current RAS administrator is already present in this list, but you can other accounts.
 - **Access Key ID:** Your access key ID.
 - **Secret Access Key:** Your secret key.
5. Click **Next**.
6. Wait until Parallels RAS validates the settings and click **Next**.
7. Select the Region that you will use. In most cases, the best Region would be the one closest to you. You can also choose one of opt-in AWS Regions by selecting the **Opted-in Region** option or specify a custom EC2 endpoint URL by selecting the **EC2 Endpoint URL** option.
8. Click **Finish**.
9. Proceed to creating a Template as described in [Creating a VDI template](#). During template creation you can configure the instance type for the clones and the storage including Type, Size, and IOPS. Note that you can also do this from **Farm > RD Session Hosts > right-click the template > Properties**.

Add a Custom Provider

Use a Custom Provider to integrate with platforms that Parallels RAS does not natively support as tier 1 providers. Your script or executable connects to the platform and acts as middleware. Parallels RAS. It can support VM discovery, guest information, power actions, cloning, and template workflows. This section covers support scope, provider configuration, required variables, and script guidance.

Introduction and Support

Introduction

Parallels RAS Custom Provider is a framework that allows organizations to integrate hypervisors not currently supported as built-in Tier 1 providers in RAS.

Using a script-based connector that acts as middleware, customers can onboard a custom provider, enumerate virtual machines, retrieve guest information, perform supported power operations, promote VMs to a RAS Template, clone VMs, and perform Template versioning through Parallels RAS. Custom Provider lets organizations integrate any hypervisor or cloud platform with Parallels RAS by supplying a script or executable that implements the Parallels RAS connector interface

Support

Parallels supports the Custom Provider framework and its integration with RAS. Hypervisor-specific scripts, provider-side logic, and any required third-party dependencies are customer- or partner-owned unless explicitly agreed otherwise.

Creating and Modifying a Custom Provider Configuration

To add a Custom Provider:

1. In the RAS Console, navigate to **Farm > Site > Providers**.
2. On the **Providers** tab, click **Tasks > Add > Custom Provider**
3. Provide a **Name** and **Description** for the Custom Provider. In this example, we use Proxmox as the example platform.
4. Optionally, provide **credentials** for communication with the **Connection Broker**, and provide credentials for deploying the **Guest Agent**. This is similar to how other providers are created.
5. Note that **Check credentials** does not work until the Provider connection details and variables are first configured. You can optionally go back to this section to perform the check credentials.
6. Click **Next**.
7. Provide the **command**, **command arguments**, and **working directory** of your custom provider code

1. Example values:

```
C:\Program Files\PowerShell\7\pwsh.exe
-NoProfile -NonInteractive -ExecutionPolicy Bypass
-File C:\CPF_Scripts\Parallels-RAS-CPF-Basic.ps1
```

```
C:\CPF_Scripts
```

8. Click **Next**.
9. Customize the variables and **values** to be passed to the custom provider code. The information provided here depends on how the custom provider code is implemented.
10. Additional variables can be added as required. For example, if you introduce various logging levels in the custom code, you can customize the log level here to pass to the script.
11. Use the plus sign to add new variables, which can be either regular or secure.
12. Click **Next**.
13. Verify the details and confirm that the status is verified and valid
14. Click **Next**.
15. Confirm the details provided in the summary
16. Click **Finish** and **Apply** the configuration.
17. Confirm that the provider's status is OK.

Modifying the Provider Configuration

To view and modify the Provider configuration, right-click it and choose **Properties**. In the dialog that opens, view and modify the Provider properties.

Custom Provider Scripts

Organizations are responsible for developing their own scripts for use with the Custom Provider. This allows for full flexibility. Consult the dedicated [Custom Provider Framework guide](#) for more information on developing, using, and troubleshooting scripts for use with the Custom Provider.

Manage VDI

Read this section to learn how manage VDI components in Parallels RAS.

Manage providers (VDI)

Read this section to learn how to modify the configuration of a Provider in Parallels RAS.

Configure a Provider

To configure an existing Provider:

1. In the RAS Console, navigate to **Farm > Site > Providers**. .
2. In the **Providers** tab, select a Provider and click **Tasks > Properties**. The **Properties** dialog opens.

Note: Some of the properties described below may be unavailable on some hosts. This depends on the Provider type.

Enable or disable a Provider in Site

By default a Provider is enabled. To enable or disable a Provider, use the **Enable provider in site** option on the **General** tab.

Properties: configure Provider connection settings

The **General** tab has different properties depending on whether it's a hypervisor-based or cloud-based provider.

Hypervisor Provider:

- **Type:** Provider type.
- **Subtype:** Hypervisor version. If the hypervisor version that you are using is not listed, select **Other**.
- **Host:** The Provider host IP address.
- **Port:** Port number on which the Provider listens for incoming connections.
- **Resource pool:** This field is enabled for VMware vCenter only. If you've specified a vCenter resource pool while adding a Provider, the pool will be displayed here. The [...] button allows you to specify a different pool (or select one if the field is empty), but only if no hosts from the current pool have been created or used in Parallels RAS in any way. If Parallels RAS detects any current usage, you will see a warning message and will not be able to change it. If you still want to select a different resource pool, you'll have manually do a full clean up in the RAS Console, so that no usage of any kind exists.
- **Description:** An optional description.

- **Dedicated Provider Agent:** Select this option if you have a dedicated RAS Provider Agent installed on a different host. Enter the host FQDN or IP address in the **Agent address** field.

Cloud-based Provider:

- **Type:** Cloud-based Provider type (e.g. Microsoft Azure).
- **Name:** Provider name.
- **Description:** An optional description.
- **Credentials:** Credentials for the account used for installing RAS Guest Agent.
- **Dedicated Provider Agent:** Select this option if you have a dedicated RAS Provider Agent installed on a different host. Enter the host FQDN or IP address in the **Agent address** field.

For description of the remaining properties of Microsoft Azure, please see [Add Microsoft Azure as a Provider](#).

Credentials: configure username and password

The **Credentials** tab has different properties depending on whether it's a hypervisor-based or cloud-based host.

Hypervisor Provider:

- Specify the username and password to log in to the Provider. Click the **Check Credentials** button to verify the credentials that you've entered.

Cloud-based Provider:

- See [Add Microsoft Azure as a Provider](#) or [Adding AWS as a Provider](#).

Advanced

The **Advanced** tab allows you to configure a provider to automatically change the type of the used managed disk to Standard HDD for VMs that are not currently in use. When a VM is started, the managed disk is automatically changed to the original type. This feature allows you to reduce the cost of maintaining VMs.

To enable disk storage cost optimization:

1. Right-click a provider in the list and choose **Properties**.
2. In the provider properties window, select the **Advanced** tab.
3. Select the **Enable disk storage cost optimization** option.
4. Select the desired option in the **Set timeout before enabling storage cost optimization** drop-down list.

MAC addresses

This tab is only displayed for Microsoft Hyper-V Failover Cluster as the Provider. It is used to enable or disable MAC address management for hosts. For more information, please see [Add a hypervisor host](#) (read the description of the step where the MAC address management is configured).

Please note that MAC address management is available in Parallels RAS since version 18. In new Parallels RAS 18 installations, this functionality is enabled by default when a Provider of type Microsoft Hyper-V Failover Cluster is added to a farm. In older Parallels RAS versions, the functionality is disabled for existing Providers, but is enabled by default when a new Provider is added.

Installing RAS Provider Agent using the installer

By default, Provider Agent is installed together with Connection Broker. However, if you want to install Provider Agent on a separate server or the push installation from the RAS Console cannot be performed for any reason. If that happens, you can install the agent by running the installer directly on the target server.

Note: You can only use these instructions to install Provider Agent in Windows.

To install the dedicated Provider Agent Agent:

1. Log in to the server where you want Provider Agent installed using an administrator account and close all other applications.

Copy the standard Parallels RAS installer (RASInstaller.msi) to the server and run it:

1. When you get to the **Select Installation Type** page, select **Custom** and click **Next**.
2. Click on **RAS Provider Agent dedicated** and select **Entire Feature will be installed on local hard drive** from the drop-down list.
3. Ensure that all other components are cleared (excluded from the installation) and click **Next**.
4. Click **Install** and follow the onscreen instruction to install the agent.

The Provider Agent does not require any configuration. Once it is installed, go back to the RAS Console, highlight the server name and click **Troubleshooting > Check Agent**. If the agent is installed properly, the status should change to **Agent Installed**.

To uninstall the Provider Agent from a server:

1. Navigate to **Start > Control Panel > Programs > Uninstall a Program**.
2. Find **Parallels Remote Application Server** in the list of installed programs.
3. If you don't have any other Parallels RAS components on the server that you want to keep, right-click **Parallels Remote Application Server** and then click **Uninstall**. Follow the instructions to uninstall the program. You may skip the rest of these instructions.
4. If you have other RAS components that you want to keep on the server, right-click **Parallels Remote Application Server** and then click **Change**.
5. Click **Next** on the Welcome page.
6. On the **Change, repair, or remove** page, select **Change**.
7. On the next page, select **Custom**.
8. Select **RAS Provider Agent dedicated**, then click the drop-down list in front of it, and click **Entire feature will be unavailable**.
9. Click **Next** and complete the wizard.

Checking the RAS Provider Agent status

To verify that the RAS Provider Agent is installed and functions properly, do the following:

1. First, you can look at the **Status** column in the **Farm > Site > Providers** list. If there's a problem with the agent, the column will display an appropriate description. Note that in addition to the description, the **Status** column uses a color code to indicate the agent status as follows:
 - Red — Not Verified
 - Orange — Needs Update
 - Green — Verified
2. Right-click a host and then click **Troubleshooting > Check agent** in the context menu.
3. The **Provider Agent Information** dialog opens displaying the information about the Provider Agent, VDI Services, and other related info.
4. If the Provider Agent is not installed, click the **Install** button and follow the onscreen instructions. See [RAS Provider Agent installation options](#) for more info.

Using a Provider in multiple farms

This topic describes how to use the same Provider in multiple RAS Farms simultaneously. To better understand the problem and the solution, consider the following hypothetical example:

- Let's say we have a hypervisor with two available virtual machines.
- We also have two Farms (1 and 2).
- Our intention is to use the first host to host resources in Farm 1 and the other to be a template in Farm 2. Both hosts will run simultaneously on the hypervisor, but each one will be available in its respective Farm only.

The problem is, RAS Guest Agent can normally communicate with one RAS Provider Agent, but since each Farm has its own RAS Provider Agent, this will not work out of the box. The solution is to make the RAS Guest Agent running in a host to be aware of only one specific RAS Provider Agent with the ability to change the assignment as needed.

The assignment is done via Windows registry. All hosts belonging to VDI pools and host clones created from a template need to have a new String value 2XVDIAgent specifying the RAS Provider Agent name or address. To add the value:

1. Log in to Windows running in the virtual machine, open the registry editor (regedit) and locate the following keys:
 - 32-bit systems: `HLKM\Software\Parallels\GuestAgent`
 - 64-bit systems:
`HLKM\Software\WOW6432Node\Parallels\GuestAgent`
2. Add a String value named 2XVDIAgent. The value data should be specified as follows:
 - If a dedicated RAS Provider Agent is used, the value must be set to the FQDN or IP address of the server where the agent is installed.
 - If the built-in RAS Provider Agent is used with manual agent selection, the value must be set to the FQDN or IP address of the RAS Connection Broker.
 - If the built-in RAS Provider Agent is used and the agent is selected automatically (high availability), the string must contain FQDNs or IP addresses of all RAS Connection Brokers separated by a semicolon (i.e. <PA1 address>;<PA2 address>;<PA3 address>).

Note that you can include names or IP addresses of multiple Connection Brokers for the manual agent selection scenario as well (the second bullet in the list above). This way you will not need to change the value every time you switch the preferred Connection Broker for a Provider.

Manage host pools (VDI)

Pools offer administrators more flexibility when managing an extensive number of hosts, especially when they are implemented in large company infrastructures. The RAS Console provides you with the framework and tools needed to create a complete pool management foundation. To manage pools, in the RAS Console, navigate to **Farm** > <Site> > **VDI** and then click the **Pools** tab.

Add host pools (VDI)

To add a host pool:

1. In the RAS console, navigate to **Farm** > <Site> > **VDIHost pools**.
2. Click the **Tasks** drop-down list above the **Host Pools** list and then click **Add** (or click the plus-sign icon). This opens the **Add VDI host pool** wizard.
3. Select **Enable Host pool in site** to enable the host pool. Specify the name and the description for the new host pool.
4. Click **Next**.
5. On the **Provisioning** page, select whether this host pool will contain template-based or standalone hosts:
 - **Template**: Hosts will be created dynamically from a template. You will need to create or select an existing template in the next step or later. Choosing **Template** as the provisioning type ensures a homogeneous host pool, which is recommended to provide consistent user experience across the host pool.
 - **Standalone**: Select one or more hosts that already exist. You'll be able to do it in the next step or you can do it later. Prior to adding hosts to host pools, ensure that hosts are domain joined and have network access to the domain environment. Note that the Standalone provisioning is considered "unmanaged" as it lacks some of the functionality, such as Autoscaling.
6. Depending on the selection made on the **Provisioning** page (above), do one of the following
 - **Standalone**: Select one or more hosts from the list to be included in the host pool (you can also add hosts to the pool later).
 - **Template**: Select a template from the list or click **Create new** to create a new template and specify the template settings. **Versions**: If you selected an existing template, select one of its versions.
7. Click **Next**.
8. (Templates only) On the **Provisioning** page, specify the following options:
 - **Template name**: Choose and type a template name.
 - **Host name**: A pattern to use when naming new hosts. For example, Windows10-RAS-%ID%.
 - **Number of hosts deployed on wizard completion**: The number of hosts to deploy once the template is created. Please keep in mind that this will take some time because the hosts will be created one at a time.

- On the **Specifications** page, override settings that are defined at the template level.
 - To override the Organizational Unit in Active Directory where hosts are joined, select **Override Organizational Unit specified in template properties**. In the target OU field, provide the distinguished name of the Organizational Unit by using the browse (...) button or entering it manually.
 - (Microsoft Azure only) To override the High availability settings, select **Override the high availability settings specified in template properties**. Select an availability option which can be 'no infrastructure redundancy required', 'availability set', or 'availability zone'. When 'availability zone' is selected, specify the zone where hosts should be deployed. Note that High availability settings cannot be changed after the host pool is created.
 - (Microsoft Azure only) To override the host size, select the **Overwrite the size specified in template properties** option and choose the desired size from the drop-down list. The selected size will only be used by this host pool. Other host pools using the same template will be unaffected. Note that available sizes may depend on the location, size, and power state of the host pool members and the template. Also note that overwriting the size requires a host reboot.
 - (Nutanix only) To override the Cluster settings, select **Override the cluster specified in Template Properties**. In the Cluster field, provide the cluster name by using the browse (...) button or entering it manually.
 - (Nutanix only) To override the Network settings, select **Override the network specified in Template Properties**. In the Network field, provide the network name by using the browse (...) button or entering it manually.
- **Host state after the preparation:** Select the power state that should be applied to a host after it is prepared. Choose from **Powered on**, **Powered off**, or **Suspended**. Note that when the power state is set to **Power off** or **Suspended**, the number of running (fully ready and waiting for incoming connections) hosts is controlled by the **Keep available buffer** setting (see above). For example, let's say the **Maximum hosts** value is set at 200, the number of guest hosts deployed on wizard completion is 100, and the power state after preparation is **Powered off**. The result of such a configuration will be 100 clones deployed and powered off.

- **Maximum hosts:** Specify the maximum number of hosts that can be created from this template.
 - (Only for templates created from VMware vCenter instances) Select one of the following:
 - If you want to create all hosts in the datastore that contains the template, select **Create hosts in the primary datastore**.
 - If you want to create hosts in several datastores, select **Create hosts in specific datastores**. If you select this option, you will be prompted to select the datastores in the **Datastores** list on the next page.
 - (Only for templates created from AWS instances) **Virtual networks:** Select Availability Zones for the host pool. Choose from **Use the subnet associated with the template** and **Use specific subnets available in the VPC**. If you select **Use specific subnets available in the VPC**, then after clicking **Next** you can manually select one or several Availability Zones.
 - **Autoscale settings:** Configure autoscale settings as described in [Manage host pools \(RD Session Hosts\)](#).
9. Click **Next**.
 10. On the **Host pool settings** page, specify the following options:
 - **On session:** Select when an action triggers.
 - **Perform action:** Select an action.
 - **After:** Select how much time must pass before action triggers.
 11. Click **Next**.
 12. On the **User profile** page, you can select from **Do not manage by RAS** (user profiles will not be managed) or **FSLogix**. Microsoft FSLogix Profile Container allows to maintain user context in non-persistent environments, minimize sign-in times and provides native profile experience eliminating compatibility issues. For complete instructions, please see [User profile](#).
 13. Click **Next**.
 14. (Standalone only) On the **Optimization** page, configure optimization as described in [Optimization](#).
 15. On the **Summary** page, review the template summary information. You can click the **Back** button to correct some of the information if needed.
 16. Finally, click **Finish** to create the host pool and close the wizard.

Delete host pools (VDI)

To delete a host pool:

Right-click it and then click **Delete** (or click the minus-sign icon, or **Tasks > Delete**).

Add and delete host pool members

A VDI pool can contain different types of members. These could be all available hosts, specific hosts, and hosts created from a template.

Adding a member to a host pool

To add a member to a pool:

1. Double-click a pool in the **Host pools** list.
2. Select the **Members** tab.
3. Click the plus (+) button and choose a member type from the following list:

- **All hosts in provider.** All hosts that are located on a given Provider. After clicking this options, you'll be able to select a Provider.

Note: Parallels does not recommend to use this type because there's a possibility that hosts with unsupported OS will be added (e.g. Linux, HALB etc). If you need to use this type, please do it carefully or use a [wildcard with appropriate hosts names](#).

- **Host.** A specific host located in the Farm. After clicking this options, you'll be able to select a host from the list.
- **Resource pool.** A group of hosts that were natively configured in the hypervisor as a pool. Please note that a hypervisor may use a different term for pools (e.g. "resource pools"). After clicking this option, you'll be able to select a resource pool from the list, if any are available.
- **Template.** Hosts that are automatically created from a template. After selecting this option, you'll be able to select a template. For more information about templates, refer to [Templates](#).

4. After you click one of the above menu items, you will be presented with the list of the available hosts, hosts, or templates from which you can make your selection.

Note: To avoid issues with overlapping members, a given pool can have members of the same type only. For example, if the first member that you add to a pool is a

host, any additional member can be a host, but not a template, Resource pool, or all hosts on a specified host. If you want to use members of different types, you must create a separate pool for each member type (i.e. one pool for hosts, another pool for templates, etc.). This requirement is enforced in the UI by disabling the member type choices once the first member is added to a pool.

Removing a member from a host pool

To remove a member from a pool:

1. Double-click a pool in the **Host pools** list.
2. Select the **Members** tab.
3. Select the pool member that you wish to delete.
4. Click the minus (-) button.

When a member is removed from a host pool, it is deleted.

Using a wildcard to filter VMs

Use the **Wildcard** input field at the bottom of the **Pools** tab to specify a wildcard to indicate which hosts should be available for users. If a VM name matches the wildcard, it will be available. If not, the users will not be able to use it. Use the asterisk operator (*) to specify a wildcard (e.g. ABC*, *ABC*).

Managing hosts in pools

Hosts that belong to a pool (and other hosts and desktops) are managed on the **VDI > Desktops** tab, where you can perform all of the standard desktop management operations from the **Tasks** menu. The operations include Recreate, Delete, Upgrade all Agents, Assign, Unassign, Show sessions, Start, Stop, Restart, Suspend, Reset, and others. The Restart operation (graceful) has a 10 min timeout. If not completed during this time, the Reset operation (forced) will be used.

By default, the **Desktops** tab displays all of the desktop available in the Farm (you may need to scroll the list to see all available desktops). To see just the hosts that belong to a specific pool, select a pool in the **Pools** tab and click **Tasks > Show hosts in Pool**. This will switch you to the **Desktops** tab where the list will be automatically filtered to include only the VMs that belong to the selected pool.

Upgrading Agents (VDI)

You can enable and configure automatic updates for all VDI hosts in a host pool.

Schedule Agent auto-upgrade

To schedule Agent auto-upgrade:

1. Go to **Farm > Site > VDI > Host pools > Properties > Auto-upgrade** tab.
2. Clear the **Inherit default settings** options if you want to modify them for this host pool.
3. Select the **Enable auto-upgrade maintenance window** option. During the maintenance window, all hosts in the host pool will try to download Agent upgrades. The upgrades will be downloaded and installed as soon as all users log out of their hosts. New logons from users are prohibited (drain mode). If the users don't log off during a maintenance window, the upgrades won't be installed until the next window.
4. Specify the start date and time, duration, and recurrence settings for this event. To make this a one-time event, select **Never** in the **Recur** drop-down list.
5. (Optional) If you want to forcefully log off all users and download the upgrades at the end of a maintenance window, select the **Force logoff of current sessions at the end of the maintenance window duration** option.
6. (Optional) Configure a message that will be sent to users before or during a maintenance window. Click the **Configure messages** button and specify the message title, body, and the time period when it should be sent.

Cancel Agent auto-update

To cancel Agent auto-update:

1. Go to **Farm > Site > VDI > Host pools**.
2. **Select Tasks > Cancel auto-upgrade maintenance window**.

Manage templates (VDI)

Templates are used to automate the creation and deployment of hosts in Parallels RAS. A template is based on an existing virtual machine created with one of the hypervisors supported by Parallels RAS. Once a template is ready, it can be used

to create clones (hosts) that will inherit all properties of the template. The resulting hosts can then be used to host published resources.

Read the following topics to learn how to create and use a template.

Virtual desktop templates

Virtual desktop templates are the essential part of the Parallels RAS VDI. They are used to create hosts for publishing of desktops, applications, documents, etc. The guest OS support is the same as for RAS Guest Agent (which must be installed in a VM). See [Software requirements](#).

Hosts created from a Virtual desktop template normally serve a single user. They are managed entirely from within the RAS VDI, which includes such features as creating persistent VMs, managing VDI sessions, publishing resources from a specific Virtual desktop template, and others.

Multi-provider template distribution

When you create a template, it is normally managed by a single Provider, which is the same provider to which the source virtual machine belongs. Clones are deployed from the template and run on the same Provider, which was used to create the template. Cloned hosts are usually stored on a centrally shared storage, such as Storage Area Network (SAN).

Beginning with Parallels RAS 18, admins have the ability to create a template and distribute it to multiple Microsoft Hyper-V hosts. As a result, the template configuration, as seen in the Parallels RAS Console, is shared among multiple Microsoft Hyper-V hosts, while each host has its own copy of the template residing on its local storage. This makes it possible to deploy clones not only to a centrally shared storage, but also to local disks of multiple independent Microsoft Hyper-V hosts. Scaling out is easily carried out by adding as many Microsoft Hyper-V hosts to the template distribution list as necessary.

Template distribution is configured on the **Distribution** page of the **Create Parallels Template Wizard**, which is described in the subsequent sections. If you are planning on using the template distribution functionality, please read the **Prerequisites** subsection below before running the wizard.

For more information about managing multi-provider distribution for a template, see also [Managing multi-provider template distribution](#).

Prerequisites

- Template distribution is supported on a standalone Microsoft Hyper-V Server, 2016, 2019, 2022, and 2025.
- All target Providers must have identical:
 - Provider type and subtype.
 - Folder path where hosts will reside.
 - Virtual switch name to which hosts are connected.
- All target Providers have to be configured with their FQDN under **Farm > Providers > Properties > Address**.
- Hyper-V hosts must be domain-joined. The current implementation uses a full VM copy of the template to distribute the template to other hosts (local storage) via the Hyper-V Live Migration mechanism.

Note: Full clones can also be moved to other hosts via Live Migration, but the process is time-consuming (equal to the first copying of the template). endhint %}

- The Microsoft Hyper-V server hosting the source VM may also be used as a target host.
- Always ensure that enough storage space is available prior to choosing target hosts to which the template will be distributed and on which clones will be created.
- Hyper-V settings must have Kerberos authentication enabled and appropriate delegations configured in AD:
 - Go to Hyper-V settings for the host machines and enable Live Migration using Kerberos.
 - Go to Active Directory Users and Computers and for each Hyper-V host server enable delegation for “cifs” and “Microsoft Virtual System Migration Service” for all servers you want to migrate To and From.

Note: If authentication isn't working, try changing the “Use any authentication protocol” option.

Creating a VM template

Requirements

To complete the tasks described in this section, the following requirements must be met:

- For hypervisor-based hosts, make sure to install the tools and/or drivers according to the hypervisor provider's recommendations, and note any specific requirements.

For Nutanix, install Nutanix Guest Tools (if required) after the cloning process rather than on the template — see the Nutanix reference [here](#).

- Make sure you know account credentials that will allow you to push install the agent software on a VM. If you run the Parallels RAS console using such credentials (e.g. a domain admin), you will not be asked to enter them during the agent installation. If you run the console using a different account, you'll be asked to enter credentials when you install the agent.
- The guest OS (Windows) running in the VM must be configured to obtain an IP address from a DHCP server.
- For users to access published resources in a host, the RDP port must be open locally or via Group Policy in Windows running in the VM. The default RDP port is 3389.
- For RD Session Host templates, Network Discovery UDP port 137 must be enabled for a domain firewall profile in the guest OS. This can be done via domain group policies or manually in the guest OS.

Manual agent installation

Normally, you will push install the necessary agent software in a source VM right from the Parallels RAS console (as described later in this section). However, you can also install the software manually by running the Parallels RAS installer in Windows in the VM. When doing so, use the **Custom** installation option and select RAS Guest Agent to be installed in the source VM.

Create a template

To begin creating a template:

1. In the RAS Console, navigate to **Farm** > <Site> > **VDI**.
2. Select the **Templates** tab in the right pane.
3. In the **Tasks** drop-down list, click **Add** (or click the “+” icon)
4. In the dialog that opens, select a host from which you would like to create a template and click **OK**.
5. The **Create Parallels Template Wizard** opens. Each wizard page is described below in the order they appear on the screen.
6. Verify that the Agent is installed and install it manually if needed as described in [Step 1: Check and install the Agent](#). This step only appears if an on-premises Provider is used.
7. Configure the template as described in [Step 2: Configure the template](#).

Step 1: Check and install the Agent

This step only appears if an on-premises Provider is used. It will not appear for Azure Virtual Desktop and cloud providers.

In this step, the wizard will check if the selected VM has the RAS Guest Agent installed. Wait for it to finish and then examine the **Status** field (closer to the bottom of the page). Depending on the result, do one of the following:

- If the agent is installed, click **Next** to continue. You may stop reading here and jump to [Step 2: Configure the template](#).
- If the agent is not installed, you need to install it as described below.
- If the agent is outdated, you’ll be prompted to update.
- If your domain credentials are invalid, you’ll be prompted to update them.

To install the agent, first click the **Customize Guest Agent deployment settings** link and specify the options in the dialog that opens. None of the options are forced, so you can select or clear them depending on your needs. Note that depending on the template type, the options are different, as described below.

Virtual desktop:

- **Add firewall rules:** Automatically configure firewall rules in the host.

- **Allow remote desktop connections:** Select to automatically configure remote desktop access in the VM.
- **Specify users or groups to be added to the Remote Desktop Users group:** Select this option and then click the **[+]** icon to add specific users to the group.

RD Session Host:

- **Add firewall rules:** Automatically configure firewall rules in the host.
- **Note:** Network Discovery UDP port 137 must be enabled for a domain firewall profile in the guest OS as a separate step. This can be done via domain group policies or manually in the guest OS.
- **Install RDS role:** Install the RDS role in the host.
- **Enable Desktop Experience:** Enable the Desktop Experience feature in Windows.
- **Restart server if required:** Restart the VM if required.
- **Specify users or groups to be added to the Remote Desktop Users group:** Select this option and then click the **[+]** icon to add specific users to the group.

When done specifying the options, click **OK** to close the dialog.

Now click the **Install** button and follow the onscreen instructions to install the agent software.

Hint: If the host cannot be reached by its name specified as hostname, double-click the host name and change it to the correct IP address.

Once done, verify that the agent is installed by looking at the **Status** field on the **Check Agent** wizard page. If so, continue to the next section that describes [Step 2: Configure the template](#).

Step 2: Configure the template

Once the agent is installed, and the **Status** field on the **Check Agent** wizard page confirms this, click **Next**. The VM will now be powered off (wait for the power off operation to finish). Once the VM is powered off, the template configuration Step begins.

The subsequent wizard pages are described in the sections that follow this one. Please note that many of the wizard pages inherit the information from Site default

settings, but you can override it if needed. To specify your own settings, clear the **Inherit default settings** option. To see and edit default settings, click the **Edit Defaults** link. For more information, see [Site defaults](#).

Properties

On the **General** page, specify the following options:

- **Template name:** Choose and type a template name.
- **Clone method:** Whether to create linked or full clones. A full clone is a complete copy of a template. As such, it occupies as much space on the physical hard drive as the source template and takes a significant time to create. A linked clone is a copy of a template made from a snapshot that shares virtual disk with the source template, therefore it takes much less space on the physical hard drive and it takes only a couple of minutes to create.

You should use full clones if your application and OS updates are too slow (full clones take longer to create, but they provide the best possible performance). Otherwise if your updates are fast enough, use linked clones as it takes much less time to create them.

Note: If the **Create a linked clone** option is greyed out, it means that the current version of Parallels RAS does not deliver linked clones with your chosen provider.

- (Microsoft Azure only) **Security type:** Select other Trusted Launch Virtual Machines or Standard. When Trusted Launch Virtual Machines is selected, additional security features can be enabled as needed.

Distribution

Use the **Distribution** page to configure template distribution. Note that this page will only appear if the source VM is created with Microsoft Hyper-V or VMware vCenter.

Microsoft Hyper-V

To configure template distribution:

1. Select the **Enable multi-provider template distribution** option.
2. In the **Available** list, select one or more providers and click **Add** (or **Add all** to add all available providers). Note that only providers of the same type and subtype as the source VM are displayed in this list.
3. In the **Number of providers for concurrent distribution** field, specify the number of concurrent distribution operations. The template is distributed to target hosts using Hyper-V Live Migration, which first exports the virtual machine to a file and then moves it to the destination host. For each host in the **Target** list, a Live Migration operation must be performed. The number specified here dictates how many network copy operations should be started at the same time. The larger the number, the more network resources will be required. Note that virtual machine exports (the first step of Live Migration) are always done one VM at a time, so the number you specify here affect only the copy operations.

Note: The **Enable multi-provider template distribution** setting cannot be modified (selected or cleared) once the template is created. If later you decide to turn it on or off (enable or disable the feature), you will need to delete and re-create the entire template. You can, however, add or remove Providers to/from an existing template.

When done, click **Next** to proceed to the next wizard page.

For the description of this feature and requirements, please see [Multi-provider template distribution](#).

VMware vCenter

Datastores are containers that store VMware virtual machines. By default, all hosts created from a template are stored in the same datastore as the template. You can choose to create hosts in several datastores, which will ensure that some hosts stay operational even if one datastore fails.

To select datastores for your hosts:

1. Select one of the following:
 - If you want to create all hosts in the datastore that contains the template, select **Use the datastore containing the VM being converted to a template**.
 - If you want to create hosts in several datastores, select **Select additional datastores to use for template distribution**.
2. (Optional) If you selected **Select additional datastores to use for template distribution**, select the datastores in the **Datastores** list.
3. Click **Next**.

Advanced

The **Advanced** page has different properties for different types of Providers. The differences are described below.

Hypervisor-based Providers

- **Cluster Shared Volume (CSV), Network share:** These two options appear if you are using Hyper-V Failover Cluster. They allow you to select a type of storage where hosts will be created. Select a desired option and then click the [...] button next to the edit field. Depending on the option selected, specify a Cluster Shared Volume or network folder. Note that a shared folder must be compatible with SMB 3.0. Please also note that the same credentials used to register Microsoft Hyper-V host as a Provider will be used to access the SMB file share for Hosts.

Please also read the important note below.

Note: To use this functionality, you need to set SMB constrained delegation (resource-based) using Windows PowerShell. **Important:** Windows Server 2012 forest functional level is required.

On a server running Windows 2012 R2 and above install the Active Directory PowerShell module using Powershell. Note that you don't need the module on a Hyper-V host or SMB file servers.

Run the following cmdlet:

```
Install-WindowsFeature RSAT-AD-PowerShell
```

Delegate SMB delegation on a file server (cluster) for every node of Hyper-V cluster. For example if you are running a four-node Hyper-V cluster and you use a Scale-Out File Server cluster FS-CL01 as virtual machine storage:

```
Enable-SmbDelegation -SmbServer FS-CL01 -SmbClient Hyperv-01
Enable-SmbDelegation -SmbServer FS-CL01 -SmbClient Hyperv-02
Enable-SmbDelegation -SmbServer FS-CL01 -SmbClient Hyperv-03
Enable-SmbDelegation -SmbServer FS-CL01 -SmbClient Hyperv-04
```

Mandatory: verify applied settings (the actual delegations) as follows:

```
Get-SmbDelegation -SmbServer FS-CL01
```

- **Folder:** This option is available if you are using Hyper-V, VMware vCenter, or Nutanix AHV (AOS). It specifies a folder where hosts will be created.
- **Use a separate network interface for LAN access:** This option is available if you are using any of the Hyper-V or VMware providers. Specifies the network interface that will be used by Connection Broker and Provider Agent. This is useful if a template has several network interfaces and you want to use a specific one for communication with Parallels RAS. If you select this option, you also need to specify the following:
 - **Address:** IP address of the network interface.
 - **Subnet mask:** Subnet mask of the IP address.
- **Resource pool:** Specifies a VMware resource pool.
- **Physical Host:** Available for VMware vCenter. Specifies a physical host where hosts will be created.
- **Enable hardware acceleration graphics licensing support:** This option is available if you are using VMware vCenter or VMware ESXi. Select it to allow vGPU-enabled hosts to unregister their vGPU licenses from the license server on shutdown.

Nutanix Prism Central Provider

- **Cluster:** Select the cluster where clones will be created. The list of clusters is that which the Provider has access to from Prism Central. This setting can be over-ridden in the hostpool but will apply by default to all hostpools.
- **Network:** Select the network that clones are going to join. This list of networks is that which the Provider has access to from Prism Central. This

setting can be over-ridden in the hostpool but will apply by default to all hostpools.

Microsoft Azure Provider

- **Resource group:** Select an Azure resource group where the cloned VM will be created. Note that this must be a group to which you granted permissions to the Microsoft Entra ID app. For details, see [Create a Microsoft Entra ID application](#).
 - **High availability:**
 - **No infrastructure redundancy required:** No redundancy in Azure will be configured.
 - **Availability set:** If selected, hosts will be deployed from the template in an availability set. Note that the maximum number of hosts that can be deployed in an availability set is 200 (this is an Azure limitation).
 - **Availability zone:** Hosts will be deployed in a specific availability zone that is selected.
- Note:** The availability option can be overridden in the settings of the host pool using a given template. This gives you the ability to specify a different availability option at the host pool level.
- **Size:** Select a VM size to be used for cloned VMs.
 - **Virtual Network:** Select a virtual network to use for cloned VMs.
 - **Subnet:** Select a subnet to use for cloned VMs.
 - **OS disk type:** Select a disk type to be used for cloned VMs.

Preparation

Use the **Preparation** page to select and configure an image preparation tool.

Note: When you specify properties on this page, they are remembered in your personal configuration file on the local machine. The next time you decide to create another template, the fields here will be populated automatically using the values you used the last time.

First, select whether you want to use RASprep or Sysprep. The advantages of using RASprep and the differences between the two tools are described below.

RASprep is the Parallels RAS tool for preparing Windows in a VM after cloning it from a base image. RASprep performs the following tasks during the initial startup of each new VM:

- Creates a new computer account in Active Directory for each host.
- Gives the host a new name.
- Joins the host to the Active Directory domain.

Compared to Sysprep, RASprep works much faster because it modifies a lower number of configurable parameters and requires less reboots.

Note: Due to API limitations, RASprep cannot be used on Windows Server 2008 machines.

The following table lists the main differences between RASprep and Sysprep:

Operation	RASprep	Sysprep
Delete local accounts	No	Yes
Generate new SIDs	No	Yes
Unjoin the parent host from the domain	No	Yes
Change computer name	Yes	Yes
Join the new instance to the domain	Yes	Yes
Language, regional settings, date and time customization	No	Yes
Number of reboots	1	2 (seal, mini-setup and domain joining)

After selecting the preparation tool, specify the following options:

- **Owner name:** Owner name (optional).
- **Organization:** Organization name (optional).
- **Administrative password:** Local Windows administrator password.
- **Domain:** A domain name for the VM to join.
- **Administrator:** Domain account.
- **Password:** Domain account password.
- **Target OU:** Specify the full distinguished name (DN) of the organizational unit (OU). Use the [...] button to browse Active Directory and select an OU. When prompted, choose whether to browse all domains or only the current domain (in most cases, the current domain is sufficient).

Note: If the Target OU cannot be validated, you may choose to ignore the warning. This is **not recommended**, but can be necessary in certain network configurations.

Optimization

The **Optimization** page allows you to specify settings that will be used to optimize Windows running in hosts for best performance in a Parallels RAS environment. You can select Windows components, services, and other options that will be disabled, removed, or optimized to ensure a more efficient, streamlined, and improved delivery of virtual apps and desktops. For the complete description, please see [Optimization](#).

After reading the **Optimization** section mentioned above, please also note the following VDI specifics:

- Optimization is disabled by default when you create a new template. If you plan to enable it, you should make a backup (create a full clone) of the source host before doing so. You can also create a template with optimization disabled, then create a snapshot, and only then enable optimization. Making a backup is a good idea because once optimization settings are applied, they cannot be rolled back.
- To enable optimization for an existing template, the template must be in maintenance. A template in the “Ready” state has the **Optimization** tab disabled.
- When optimization is enabled or modified for an existing template and the template exits the maintenance mode, the administrator will be asked to recreate existing hosts, so that optimization settings are applied to them. Note that hosts must be recreated (now or later) to receive optimization settings.
- When optimization is applied to a template, its status changes to **Optimization in progress** (among others). At this stage, you can select the template in the list and click **Tasks > Stop optimization**, which will cancel the operation.

License keys

On the **License Keys** page, specify the license key information that will be used to activate virtual machines created from this template.

First, select the license key management type that you are using in your organization (KMS or MAK). Parallels recommend to use KMS because MAK has limited activations.

Key Management Service (KMS): If you are using KMS, click the **Finish** button to save the template configuration information. Virtual machines that will be created from this template will look for KMS in DNS (at the end of the OS mini-setup and domain joining) and will be activated accordingly.

Note: If you are using KMS activation and RASPrep, the source host must be activated using KMS before you create a template from it. If the host has already been activated using another method (retail key or MAK), you need to convert it to KMS activation. For the information on how to do it, please read the following article from Microsoft: <https://technet.microsoft.com/en-us/library/ff793406.aspx>.

Multiple Activation Keys (MAK): If you are using MAK, do the following:

1. Click the **Add** button and type a valid key in the **License key** field.
2. In the **Max. guests** field, specify the key limit. The limit should be greater than or equal to the max guests in the template (which you set on the first page of the wizard)
3. Click **OK**.

Note: Parallels RAS does not keep the old MAK key in hosts if it was updated in the Parallels template properties.

Summary

On the **Summary** page, review the template summary information. You can click the **Back** button to correct some of the information if needed.

Finally, click **Finish** to create the template and close the wizard.

Host naming

This section describes the host naming pattern that you specify on the [Properties](#) page of the template creation wizard.

Each time a new host is created, a name for it is generated automatically based on the pattern that you specify in the **Host name** field. The complete name format is as follows:

`<prefix>%ID:N:S%<ending>`

where:

- **<prefix>** is an alphanumeric string that must begin with a letter (not a digit).
- **%ID:N:S%** is a numeric pattern used to automatically generate a unique host ID. See the **Numeric pattern** subsection below.
- **<ending>** is a free-form alphanumeric string.

Numeric pattern

The numeric pattern in the VM name has the following format:

`%ID:N:S%`

The elements in the pattern above are:

- **ID** — Must be included as is.
- **N** — The number of digits to use, including leading zeros. Use “0” if you don’t want to insert leading zeros.
- **S** — The starting number. This element is optional. If you don’t include it, the number will start with 1.

Examples:

- `%ID:3%` — This pattern will generate 3 digit numbers with leading zeros, such as “001”, “002”, “003”... “998”, “999”.
- `%ID:3:200%` — This example will generate 3 digit numbers starting from 200, such as “200”, “201”, “202”... “998”, “999”.

- VDI-R1-%ID:3:100% — This is a complete name with an alphanumeric prefix and a numeric pattern. The resulting names will look like the following: “VDI-R1-100”, “VDI-R1-101”, etc.

When creating a name pattern, follow the rules listed below. If any of these rules are not observed, you will see an error message and will have to correct it:

- The name must start with a letter. A digit is not allowed as the first character.
- The alphanumeric part of a name can contain letters, digits, and a hyphen. No other characters are allowed.
- The total length of the name must not exceed 15 characters.
- The name can include just one numerical pattern (%ID:N:S%), which must be placed at the end or in the middle of the name.

The pattern that you specify is also validated against the value of the **Maximum hosts** field. If the pattern doesn’t cover the maximum number of hosts, you will get an error and will have to correct it.

Reusing VM numbers in a name

When you delete a host, the number that was assigned to it becomes unused. The next host that is created will be given this number, so there are no gaps in numbering.

Parallels Test Template Wizard

The **Parallels RAS Test Template Wizard** is used to test the health of a template. The wizard allows you to see that all post-prep activities for a template complete correctly. This includes checking DHCP settings, DNS registration, correct VLAN, joining the AD domain, correct target OU, etc.

To open the wizard, right-click a template in the Parallels RAS console and choose **Test**. The test procedure consists of the following steps:

1. The template is switched temporarily to the “Test” mode designed specifically for this purpose. Please note that while the template is in this mode, all other operations are blocked until the test is finished and the template exits the test mode.
2. A host is cloned from it to be used for testing. The VM is kept on the server for the duration of the test and will be deleted afterwards.

3. A series of tests is then run on the host to test the template from which it was created.
4. Once the test is complete, a report is displayed on the screen showing the test results.

When the wizard starts:

1. The **Welcome** page opens. Read the info that it contains and click **Next** when ready.
2. The next page displays the list of individual tests that will be performed, including:
 - **Check host Agent:** This test tries to communicate with the RAS Guest Agent installed in the VM. If the agent responds, it means that the VM has been created and started successfully.
 - **Check domain membership:** Checks that the computer has joined the AD domain.
 - **Check target OU:** Checks that the RDP connection to the computer is possible with domain credentials.
 - **Launch Parallels Client:** This test launches Parallels Client and establishes a connection with the host.
3. While the test is running, the progress indicator is displayed on the screen. If needed, you can cancel the test at any time by clicking the **Cancel** button.
4. Once all tests are completed, you will see a page displaying the test results:
 - **Success:** If all tests complete successfully, the temporary host will be marked for deletion and the template will be switched back to the normal operation mode.
 - **Failure:** If one or more tests fail, you will see the corresponding info and will be able to download the log file by clicking the **Download log file** link. You will also have an option to switch the template to maintenance mode, which will prevent creating hosts from it until it is fixed.
5. Click **Finish** to close the wizard.

Modifying template properties

If you need to change the configuration of an existing template, select it in the **Templates** list and click **Tasks > Properties**. This opens the **Template Properties** dialog, which consists of tabs containing the same properties as the wizard pages described in [Step 2: Configure the template](#).

How hosts are created from a template

After a template is created, Parallels RAS begins creating hosts from it, one virtual machine at a time. The number of VMs created at this time is determined by the **Number of hosts deployed on the wizard completion** property (all property names here and later refer to the **Create Template Wizard** described earlier).

The number of VMs available at any time will never go below the number specified in the **Keep available buffer** property. To comply with this rule, a new VM is automatically created when needed. At the same time, the total number of VMs will never exceed the number specified in the **Maximum hosts** property.

Please note that creating a new host from a template takes some time, especially when a template is configured to create full clones (linked clones are created much faster). If a host is in the middle of being created, and no other VMs are available, a user (or users) who need it will have to wait until the VM is ready.

If a host encounters a problem during the preparation stage, it will remain on the server in unusable state. You can identify such VMs by the **Failed to create** value in the **Status** column. Unless a VM like this is repaired or recreated, it will be automatically removed after the time period specified in the **Auto remove hosts which failed preparation after** field in Site defaults (Farm > <Site> > **VDI** > **Desktops** > **Tasks** > **Site defaults**). You can view the details of the failure to create a host by clicking on the **Details** link in the **Status** column or double-clicking the column. You can also choose **Recreate** or **Delete** from the same menu. For more information on how to recreate a host, please see the [Template maintenance](#) section.

Auto-deletion of hosts

A host is automatically deleted when it is not used longer than specified in the **Delete unused hosts after** field in template properties.

Manually adding a host

Hosts are created from a template automatically. In a situation when one or more additional hosts are required, you can add (create) them manually.

To add a host:

1. In the RAS Console, navigate to **Farm** > <Site> > **VDI** > **Hosts**.
2. Click the **[+]** icon at the top of the list.
3. In the **Add Hosts** dialog that opens, select a template from which to create a new host,
4. Specify the number of hosts to create. If the number you specify exceeds the **Maximum hosts** value set in the host pool properties (taking into account the number of VMs that already exist), you'll see a warning message and will need use a lower number or change the maximum host number on the **Provisioning** tab of the host pools properties.
5. Click **OK** to close the dialog.
6. After you click **Apply** in the RAS Console, the new hosts will appear in the list on the **Desktop** tab with the **Status** column saying "Cloning". Once the cloning is complete, the new hosts become available to users.

Assigning a template to a host pool (VDI)

When you create VDI host pools, you can assign a template to a pool. This can be done when you create or modify a pool, or it can be done from the **Templates** tab.

To assign a template to a host pool:

1. On the **Templates** tab, select a template.
2. Click **Tasks** > **Assign to pool**. A wizard opens.
3. On the **Versions** page, select the template version that will be assigned to the host pool.
4. (Optional) On the **Host pool** page, select the host pools that you want to recreate on schedule and click the **Next** button. You will see a dialog that allows you to schedule recreation. Configure the schedule according to your needs and click **Next**.
5. Click **Finish**.

To remove a template from a host pool:

1. Select a template and click **Tasks** > **Remove from pool**.
2. A dialog opens listing all host pools to which this template is assigned.
3. Select the host pools to remove the template from and click **OK**.

Note that if a host pool has hosts created from the template that you are removing, they will be removed as well. A message is displayed where you need to confirm the removal.

Template maintenance

A template can be put into to a special mode called “maintenance”, which is primarily used to update or install software in the guest operating system. While in this mode, the template becomes unavailable for all normal tasks, including creating new hosts, and it becomes possible to start it as a regular virtual machine. Once the virtual machine is running, you can install or update software in the guest OS or perform administrative tasks in the operating system.

Depending on whether a template is configured for full or linked clones, the maintenance mode is used slightly differently, as described below.

Full clones

If your template is configured to create full clones, do the following:

1. Select a template and click **Tasks > Maintenance** and select the [template version](#) to be put in maintenance. The template becomes disabled (grayed out) and all operations on it are suspended. The status of template in the **Status** column changes to **Entering maintenance** and when completed, changes to **Maintenance**.
2. Using native tools of the hypervisor, start the template as a normal virtual machine.
3. Install Windows updates or software as necessary.
4. When done, shut down the virtual machine.
5. Back in the RAS Console, select the template and click **Tasks > Maintenance** again to exit the maintenance mode. You will see a dialog where you can choose whether to create a new version or discard the changes. Select **Create a new version**.

Note: One template can have up to five versions. If you want to create another version, you will have to delete an already existing one.

1. Create a new template version as described in section [Using template versions](#), subsection “Creating a new version”. When you update a full clone template, the changes will only affect future clones. For existing clones to have these updates, they must be recreated. You can choose to recreate

existing hosts now or you can postpone it. Please note that recreating a full clone is a time consuming process. Also, a new app may be installed in a full clone VM or a user profile may be changed while the recreation is in progress, all of which will be lost. To minimize impact on users, it makes sense to schedule a maintenance window during which the clones can be recreated.

Linked clones

Since linked clones share the virtual hard disk with a snapshot of a template, you need to take additional steps compared to full clones.

First, you need to notify host users to save their data and log off. This is necessary for existing hosts to include the updates that you will install in the template. Once all users are logged off, do the following:

1. Select a template and click **Tasks > Maintenance** and select the [template version](#) to be put in maintenance. The template becomes disabled (grayed out) and all operations on it are suspended. The status of the operation is displayed at the bottom of the window.
2. Using native tools of the hypervisor, start the template as a normal virtual machine.

Attention: For Nutanix Prism Central, the template versioning process does not power on the original template. Instead, it creates a temporary copy, and this copy (and not the original) is the one you should edit.

In RAS, you can do it by navigating to **Tasks > Tools > Remote Desktop** with the **Template** selected, or by navigating to **Templates > Versions** in Prism Central, whilst the template is in maintenance mode in RAS.

3. Install Windows updates or software as necessary.
4. When done, shut down the virtual machine.
5. Back in the RAS Console, select the template and click **Tasks > Maintenance** again to exit the maintenance mode. You will see a dialog where you can choose whether to create a new version or discard the changes. Select **Create a new version**.

Note: One template can have up to five versions. If you want to create another version, you will have to delete an already existing one.

1. Create a new template version as described in section [Using template versions](#), subsection “Creating a new version”. Please note that if you create a new version without recreating linked clones, you will have to recreate them manually or by using scheduler.

Updating RAS Guest Agent inside a template

A template must have the latest version of RAS Guest Agent installed in it. The agent is installed when you create a template. When a new version of RAS Guest Agent becomes available, it should be updated. To update the agent, the maintenance mode must be used as described above. To simplify agent updates, Parallels RAS monitors all installed agents and notifies the administrator when an update is available.

When the RAS Console starts, all installed agents are checked and a message is displayed if one or more agents need to be updated. This applies to servers in the RAS infrastructure and the templates. The message will ask if you want to update all agents. If you click **Yes**, you are presented with a dialog listing all servers and templates on which an agent needs to be updated. You can select or un-select a server/template to include it in the bulk update procedure or exclude it. Once you’ve made your selection, click **OK** to start the update. Follow the onscreen instructions and update the agents.

Full vs. linked clone templates: When you update RAS Guest Agent in a template, you also need to update Agents in hosts that were created from this template. This update is done differently for full and linked clone templates. Please read the instructions below for the explanation.

When you update the Agent in a linked clone template, you’ll be asked if you want to recreate all hosts that were created from this template. You can click **Yes** and they will be automatically recreated to match the template.

When you update the Agent in a full clone template, full clone hosts are not automatically recreated. You will be asked if you want to recreate them. If you decide to do so, please note that full clone VMs are complete machines, so recreating them is a time-consuming process. Alternatively, you can update the agent in these VMs by push-installing it from the RAS Console. This can be done by clicking **Tasks > Upgrade all Agents** while on the **VDI > Desktops** tab.

To manually check the RAS Guest Agent status in a template, click **Tasks > Check agent**. If the agent is up to date, a message box is displayed confirming this. If a newer version of RAS Guest Agent is available, you'll see a dialog asking you to update it. Please note that the difference in updating full and linked clone templates (as described above) applies to this scenario as well.

Maintaining RD Session Hosts based on a template

If you need to do a scheduled maintenance of RD Session Hosts that were created from a template, please follow these steps:

1. Create a schedule that fits your maintenance window to drain a desired RD Session Host group.
2. During maintenance (or right before it) switch the template into maintenance mode. Then apply the necessary changes.
3. The schedule disables groups provisioned by the template (while the maintenance window lasts) which leads to removing (unassigning) all hosts from them.
4. Release the template from maintenance and click **Yes** when asked whether to recreate all clones.
5. Enable groups which were disabled in step 3 (above). At this point, the groups will begin receiving hosts to comply with **Keep Available Buffer** setting
6. From this point forward, groups are provisioned with VMs on demand.

Template status

To verify that a template is functioning as intended, you can examine its status in the main template list in the RAS Console (the **Status** column). When a template is functioning properly, the **Status** column displays "Ready", which means that hosts can be created from it as needed. When a template is being created or when it's in maintenance, or when it's being removed, the status will change accordingly.

Note that one of the other columns in the table is **Agent status**, which is the status of the RAS Guest Agent installed in a template. Compared to servers in the RAS infrastructure (Connection Broker, Gateways, RD Session Hosts, etc), the agent status is not as important in a template as the template status. This is because a template is not a regular virtual machine and is not normally running, so checking the agent status in a stopped VM doesn't make much sense. This is why the agent status for a template in the RAS Console is usually **Not Available**, which is

perfectly normal. The only situation when the **Agent status** displays a meaningful value is when the template is in Maintenance and running like a regular VM, in which case the agent is also running and its status can be verified.

The tables below describe what the **Status** and **Agent status** columns will show for various template states or transitions.

Template creation

Status color	Template status	Agent status	Description
Gray	Not available	Not available	When a provider is disabled or Provider Agent is disconnected or the template does not exist.
Gray	Not applied	Not applied	Awaiting admin to click Apply in the RAS Console after the wizard completion.
Orange	Creating	Not available	When using Azure Gallery as a source (no pre-created host available).
Orange	Agent installation	Not available then OK	Deploying the agent to a newly created or available VM to be used as a template.
Orange	Deployment in progress	Optimization pending	When the optimization is waiting to be applied. The admin can stop the optimization at this point.
Optimization in progress	When the optimization is in progress. The admin can still stop the optimization at this point.		
OK then Not available	Internal procedure when converting a VM to a template. Once the conversion is completed, the template		

Status color	Template status	Agent status	Description
	status changes to "Ready".		
Red	Creation failed	Not available	A problem has occurred. For example, a quota hit or a resource creation issue in Azure. The admin can retry the action by clicking Tasks > Retry last action.
Red	Agent installation failed	Not available	Possible network issues, file share limitations, or an issue with admin rights. The admin can try Tasks > Retry last action.
Red	Deployment failed	The actual status (OK, Not available, etc.)	A problem has occurred. For example, a quota hit, storage space, or snapshot creation from provider issue. The admin can try Tasks > Retry last action.
FSLogix not available	FSLogix agent not found.		
FSLogix not updated	FSLogix agent needs updating.		

Template in production

Status color	Template status	Agent status	Description
Green	Ready	The actual status (OK, Not available, etc.)	Template is ready.
Green	Cloning	The actual status (OK, Not available, etc.)	A host is being cloned from the template.
Orange	Needs update	Needs update	RAS Guest Agent needs updating.

Template in maintenance

Status color	Template status	Agent status	Description
Orange	Maintenance	The actual status (OK, Not verified, etc.)	Host used as the template is up and running,
Optimization pending	When the optimization is waiting to be applied. The admin can stop the optimization at this point.		
Optimization in progress	When the optimization is in progress. The admin can still stop the optimization at this point.		
Needs update	Host used as the template is up and running but RAS Guest Agent needs updating.		

Template removal

Status color	Template status	Agent status	Description
Gray	Marked for deletion	The actual status while the host used as the template is still running (OK, Not verified, etc.)	The template is in the process of being deleted.

Managing multi-provider template distribution

For the description of the Multi-Provider Template Distribution feature, please see the [Multi-provider template distribution](#) section.

Adding or removing Providers from the distribution list

You can add or remove a Provider to/from a distribution list at any time using the template **Properties** dialog. To open the dialog, right-click a template on the **VDI > Templates** tab and choose **Properties**.

Template distribution status

After you complete the **Create Parallels Template Wizard** and create a template, or when you add/remove a Provider to/from an existing template, you can monitor the template distribution status on the **Templates** tab. The status is displayed in the **Distribution** column and may have the following values:

- **Distributing** — the distribution is in progress (the template is being distributed to target hosts).
- **OK** — the template has been successfully distributed to all specified hosts.
- **Removing / Adding** provider — A Provider is being added or removed.
- **Failed to distribute** — indicates that an error has occurred during the distribution operation.

Distribution details

The **Tasks > Distribution details** menu on the **Templates** tab opens a dialog where you can view the current distribution **State** and **Progress** indicators for the Providers that use this template.

The **Progress** column displays the same values as the **Distribution** column in the main template list (see above).

The **State** column may display one of the following:

- **Ready** — The Provider is ready.
- **Not available** — The Provider is not responding.
- **Needs update** — The template distribution operation may need to be performed again. You can click the **Retry** button to retry the template distribution operation for this host.

Leaving maintenance mode

When a template leaves the maintenance mode, a prompt is usually displayed saying that “All hosts must be recreated because the template has been modified. Do you want to recreate them now?”. If the administrator clicks **Yes** and the template uses multi-provider distribution, Parallels RAS verifies the status of each provider. If a Provider is not responding, a message is displayed, asking the administrator to check the provider status. You can bring the provider back online and try recreating the hosts again. If this cannot be done at this time, you can recreate the hosts later.

Managing template-based hosts

Hosts and other desktops are managed on the **VDI > Desktops** tab, where you can perform all of the standard desktop management operations from the **Tasks** menu. The operations include Recreate, Delete, Upgrade all Agents, Assign, Unassign, Show sessions, Start, Stop, Restart, Suspend, Reset, and others. The Restart operation (graceful) has a 10 min timeout. If not completed during this time, the Reset operation (forced) will be used.

By default, the **Desktops** tab displays all of the desktop available in the Farm (you may need to scroll the list to see all available desktops). To see just the hosts that belong to a specific template, select a template in the **Templates** tab and click **Tasks > Show hosts**. This will switch you to the **Desktops** tab where the list will be automatically filtered to include only the VMs that belong to the selected template.

For more information, see [Managing hosts](#).

Manage hosts (VDI)

There are two basic types of hosts when using Parallels RAS VDI: template-based and non-template based. This topic describes management tasks for both host types, indicating whether a task applies to a particular host type.

Viewing host list

To view the list of non-template based hosts, select **Farm > <Site> > VDI > Desktops**. If you have a filter applied to the list, remove it by click the magnifying

glass icon. Without the filter, the list shows all desktops available in this RAS Farm, including hosts (both template-based and non-template based), hosts from a pool (RAS or native), and pool-based Remote PCs. Therefore, the **Desktops** tab is a location where you can view all of your desktops in one place. Here you can perform all of the standard desktop management tasks accessible from the **Tasks** menu, including Recreate, Delete, Assign, Unassign, Start, Stop, Restart, Suspend, Reset, Show sessions, and others. The Restart operation (graceful) has a 10 min timeout. If not completed during this time, the Reset operation (forced) will be used.

To view the list of hosts created from a template, select **Farm** > <Site> > **VDI** > **Templates**. Select a template and click **Tasks** > **Show hosts**. You will be switched to the **Desktops** tab where the list of desktops will be filtered to include only those that belong to the template. As was mentioned above, you can perform all of the standard desktop management operations on this tab, including power operations, which are described in detail later in this section.

For the list to include only the hosts from a particular pool, select a pool in the **Pools** tab and click **Tasks** > **Show hosts in Pool**.

The filter in the **Desktop** tab can also be applied manually by clicking the magnifying glass icon and entering the filter criteria in the fields that appear at the top of the list.

Site defaults

Hosts created from a template inherit the template settings. To view the settings, note on which template a host is based and then view properties of that template, specifically the **Settings** and **Security** tabs. For more information, see [Site defaults](#). Note that you a template can inherit Site default settings or you can specify your own custom settings for it.

Non-template based hosts have their own settings, some of which (specifically Settings and Security) are inherited from [Site defaults](#). To see settings for a non template-based VM, navigate to **Farm** > <Site> > **VDI** > **Desktops**. A host that doesn't belong to a template is identified by an empty value in the **Template** column. Right-click a template and choose **Properties** (note that template-based hosts do not have this menu option).

Checking the RAS Guest Agent status

A host must have the RAS Guest Agent installed and the agent must match the Parallels RAS version. The agent is installed by default when a host is created from a template. If a host was created using the native hypervisor tools, it may not have the agent installed in it. In such a case, the host will be able to serve only the remote desktop. To enable it to server applications or documents, you'll need to install the agent yourself.

To check if the RAS Guest Agent is installed in a host and is up to date:

1. Select a host in the list and then click **Tasks > Troubleshooting > Check agent**.
2. The **Guest Agent Information** dialog opens displaying the information about the RAS Guest Agent.
3. If the agent is not installed, click the **Install** button and follow the instructions. The agent will be push installed in Windows running inside the host.

Deleting a host

To delete a template-based host, select it and then click the **Tasks > Delete**.

Important: You should delete a host only from the RAS Console. You should not try to delete a host using the hypervisor's native client or web interface. If you do, it may delete not only the VM but its parent template as well (which will also invalidate all other hosts created as linked clones from this template). The reason for this is some native hypervisor clients treat linked clones as standalone VMs. Parallels RAS treats linked clones as clones, not as standalone VMs.

Managing hosts that failed preparation

If a template-based host encounters a problem during the preparation stage, it remains on the server but cannot be used. You can identify such VMs by the "Failed to create" value in the **Status** column. Unless a VM like this is repaired, it will be automatically removed after the time period specified in the [Site defaults](#). To see Site defaults:

1. Select **Farm > <Site> > VDI > Desktop** and then click **Tasks > Site defaults**.
2. In the dialog that opens, on the **General** tab, view or modify (if needed) the **Auto remove hosts which failed preparation after** option. You can set any

of the available time periods by selecting it from the drop-down list or you can type a desired value, such as “8 days” or “12 hours”.

Recreating a host

If something happens to a template-based host and it becomes unusable, you don't have to delete it and create a new one. Instead, you can recreate it keeping its name, MAC address, and other properties. This way none of the other Site settings, which may rely on a broken host, will be affected. Another reason for recreating a host is to apply changes made to the template (when you exit from maintenance without executing the Recreate command).

Please note that recreated VMs can keep the the following properties:

- MAC address is kept on ESXi, vCenter, Hyper-v, Hyper-v Failover Cluster, Nutanix AHV (AOS), Nutanix Prism Central, and SC//HyperCore.
- BIOS UUID is kept on ESXi and vCenter.
- DRS groups are kept on vCenter.

To recreate one or more hosts:

1. In the Parallels RAS Console, navigate to **Farm > <Site> > VDI > Templates**.
2. To recreate all deployed hosts, click the **Tasks** drop-down list and choose **Recreate all hosts**.
3. To recreate a specific host (or multiple hosts), click **Tasks > Show hosts**. This will open the **Desktops** tab, which will list hosts. Select one or more hosts and then click the **Tasks > Recreate**.

When you recreate a host:

- The procedure deletes a VM and creates a new one from the same template.
- The new host retains the same computer name as the one it replaces.
- If a host is running, all unsaved data in its memory will be lost. For this reason, an important data should be saved to an external storage.

Assigning a host

If a template-based host becomes unassigned from its host pool, the Parallels console shows a dialog indicating that a templated host is not assigned to a host pool at a particular site.

To assign the host back to the host pool:

1. Open the hosts VDI hosts overview
2. Select one or more unassigned hosts
3. Right-click select **Assign to host pool**
4. Select the host pool you want to assign the host to and click **OK**.

Persistent hosts

A host is called persistent when it is assigned to a particular user or device. To make a host persistent, do the following:

1. Begin publishing a desktop or a resource from a host.
2. When specifying **Virtual Guest Settings** options, select **Enable static assignment to host**.
3. Complete the publishing wizard.
4. As a result, the VM will be assigned to the first user or device who uses a desktop or a resource. For information on how to switch between user and device assignments see section [Site defaults \(VDI\)](#), subsection “General”.

You can also manually assign a host to a user or device. To do so:

1. Navigate to **Farm > <Site> > VDI > Desktops**.
2. Select a host and click **Tasks > Assign**. You can assign a host to a user or device. For information on how to switch between user and device assignments see section [Site defaults \(VDI\)](#), subsection “General”.
3. To assign to a user, specify the user registered in Active Directory.
4. To assign to a device, select one of the following options:
 - **Add from Active Director** allows you to add a device joined to a domain in Active Directory.
 - **Add from known device** allows you to add a device known to RAS Device Manager.
 - **Add custom entry** allows you to input the device name manually
5. As a result, the host will be assigned to the selected user.

To view persistent hosts, navigate to **Farm > <Site> > VDI > Desktops**. A persistent host is identified by the **Persistent** value in the **Assignment** column.

To remove persistence from a host, do one of the following:

- Select a host on the **Desktops** tab and then click **Tasks > Unassign**.

- Navigate to **Farm** > <Site> > **VDI** > **Desktops** and click **Tasks** > **Site defaults**. In the dialog that opens, use the **Auto remove persistence if host was not used for** option to select the time period after which persistence should be automatically removed. You can also type any desired time period, such as “1 week 3 days”.

Configuring hosts to interact with RAS Provider Agent in a different subnet

RAS Guest Agent installed on a host discovers RAS Provider Agent by sending broadcast packets over UDP port 30006 to all devices in its subnet.

The discovery packet won't be received if the Agents are on different broadcast domains. In this case, you need to specify the Provider Agent's IP through the registry.

To specify the Provider Agent's IP:

1. On the host, open the Windows registry.
2. Add a new string value called `2XVDIAgent` at:
 - 32-bit systems: `HLKM > Software > Parallels>GuestAgent`
 - 64-bit systems: `HLKM > Software > WOW6432Node > Parallels>GuestAgent`

The value should be specified as follows:

- If a dedicated RAS Provider Agent is used, the value must be set to the FQDN or IP address of the server where the Agent is installed.
- If the built-in RAS Provider Agent is used with manual agent selection, the value must be set to the FQDN or IP address of the RAS Publishing Agent.
- If the built-in RAS VDI Agent is used and the Agent is selected automatically (high availability), the string must contain FQDNs or IP addresses of all RAS Publishing Agents separated by a semicolon (i.e., <PA1 address>;<PA2 address>;<PA3 address>).

Note: You can also include names or IP addresses of multiple Publishing Agents for the manual agent selection scenario (the second bullet in the list above). This

way you will not need to change the value every time you switch the preferred Publishing Agent for a VDI provider.

Note: If you add the registry key to a host used as a template, its clones will inherit the key.

1. Restart the RAS Guest Agent service on the host to apply the changes.

Manage sessions (VDI)

Please see [Session Management](#).

Using scheduler (VDI)

The **Scheduler** tab allows you to create scheduler tasks that will be performed on individual hosts or host pools at a specified time.

Note: When the scheduled event is triggered, affected hosts are disabled in Parallels RAS and their status is displayed as “Disabled (scheduler)” or “Pending reboot (scheduler)”. You can cancel these states by right-clicking a host on the **Hosts** tab and choosing **Control > Cancel disabled state (scheduler)** or **Control > Cancel pending reboot (scheduler)**.

Disabling hosts and hosts in pools

To disable a host or a host in a pool:

1. Click **Tasks > Add > Disable host** or **Disable host pool**.
2. On the **General** tab, select the **Enable Schedule** option.
3. Specify a name for this schedule and an optional description.
4. Select a host or a pool in the **Available** list and click **Add**. The host or pool will appear in the **Target** list.
5. Select the **Trigger** tab and specify start date and time, duration, and recurrence settings for this event. To make this a one-time event, select **Never** in the **Recur** drop-down list.
6. Select the **Options** tab. It contains the following options:
 - Message list: Configure a message that will be sent to users before the host goes offline. Click **Tasks > Add** and specify the message title, body, and the time period when it should be sent.

- **On disable:** Specify what should happen to current sessions when a scheduled task triggers. Select the desired option from the **On disable** drop-down list.
 - **Enforce schedule for currently inactive hosts:** This option is only enabled when you have an active message in the list. If the option is enabled, hosts that are currently offline are also monitored, and if such a server comes back online during the scheduled task execution, the task is applied to it too.
7. Click **OK** to save the schedule.

Rebooting hosts and hosts in pools

To reboot a host or a host in a pool:

1. Click **Tasks > Add > Reboot host** or **Reboot host pool**.
2. On the **General** tab, select the **Enable Schedule** option.
3. Specify a name for this schedule and an optional description.
4. Select a host or a pool in the **Available** list and click **Add**. The host or pool will appear in the **Target** list.
5. Select the **Trigger** tab and specify start date, time, and recurrence settings for this event. To make this a one-time event, select **Never** in the **Recur** drop-down list. In addition, specify the following options for the “Reboot host pool” task:
 - **Complete in:** Specify the time to complete the task.
6. Select the **Options** tab. It contains the following options:
 - **Message list:** Configure a message that will be sent to users before the host is rebooted. Click **Tasks > Add** and specify the message title, body, and the time period when it should be sent.
 - **Enable Drain Mode** and **Force server reboot after:** The two options work together. If you select the **Enable Drain Mode** option, then when the task triggers, new connections to a host are refused but active connections will continue to run and can be reconnected. The server will be rebooted when all active users close their sessions or when **Force server reboot after** time is reached, whichever comes first. For active users not to lose their work, create a message that will advise them to save their work and log off.
 - **Enforce schedule for currently inactive hosts:** This option is only enabled when you have an active message in the list. If the option is enabled, hosts that are currently offline are also monitored, and if such

a server comes back online during the scheduled task execution, the task is applied to it too.

7. Click **OK** to save the schedule.

Starting up hosts and hosts in pools

Note: This task applies only to hosts and host pools based on a template.

To start up a host or a host in a pool:

1. Click **Tasks > Add > Startup host** or **Startup host pool**.
2. On the **General** tab, select the **Enable Schedule** option.
3. Specify a name for this schedule and an optional description.
4. Select a host or a pool in the **Available** list and click **Add**. The host or pool will appear in the **Target** list.
5. Select the **Trigger** tab and specify start date, time, and recurrence settings for this event. To make this a one-time event, select **Never** in the **Recur** drop-down list.
6. (“Startup host pool” task only) Select the **Options** tab. It contains the following options:
 - **Power on all members:** Select this option to start up all hosts assigned to specific users.
 - **Percentage of members:** Select this option to specify the percentage of hosts that must be started up in each pool.
 - **Specific number of members to be started:** Select this option to specify the number of hosts that must be started up in each pool.
7. Click **OK** to save the schedule.

Shutting down hosts and hosts in pools

To shut down a host or a host in a pool:

1. Click **Tasks > Add > Shutdown host** or **Shutdown host pool**.
2. On the **General** tab, select the **Enable Schedule** option.
3. Specify a name for this schedule and an optional description.
4. Select a host or a pool in the **Available** list and click **Add**. The host or pool will appear in the **Target** list.
5. Select the **Trigger** tab and specify start date, time, and recurrence settings for this event. To make this a one-time event, select **Never** in the **Recur** drop-down list.

6. Select the **Options** tab. It contains the following options:
 - **Message list**: Configure a message that will be sent to users before the host is rebooted. Click **Tasks > Add** and specify the message title, body, and the time period when it should be sent.
 - **Enable Drain Mode** and **Force server shutdown after**: The two options work together. If you select the **Enable Drain Mode** option, then when the task triggers, new connections to a host are refused but active connections will continue to run. The server will be shut down when all active users close their sessions or when **Force server reboot after** time is reached, whichever comes first. For active users not to lose their work, create a message that will advise them to save their work and log off.
 - **Enforce schedule for currently inactive hosts**: This option is only enabled when you have an active message in the list. If the option is enabled, hosts that are currently offline are also monitored, and if such a server comes back online during the scheduled task execution, the task is applied to it too.

Recreating hosts and host pools

Note: This task applies only to hosts and host pools based on a template.

To recreate a specific host or all hosts in a host pool:

1. Click **Tasks > Add > Recreate host from template** or **Recreate host pool from template**.
2. On the **General** tab, select the **Enable Schedule** option.
3. Specify a name for this schedule and an optional description.
4. Select a host or a pool in the **Available** list and click **Add**. The host or pool will appear in the **Target** list.
5. Select the **Trigger** tab and specify start date, time, and recurrence settings for this event. To make this a one-time event, select **Never** in the **Recur** drop-down list.
6. Select the **Options** tab. It contains the following options:
 - **Message list**: Configure a message that will be sent to users before the host is rebooted. Click **Tasks > Add** and specify the message title, body, and the time period when it should be sent.
 - **Enable Drain Mode**, **Force host recreation after**, and **Force host pool recreation after**: The options work together. If you select the **Enable Drain Mode** option, then when the task triggers, new connections to a host are refused but active connections will continue

to run. The server will be recreated when all active users close their sessions or when the time specified in **Force host recreation after** or **Force host pool recreation after** is reached, whichever comes first. For active users not to lose their work, create a message that will advise them to save their work and log off.

- **Enforce schedule for currently inactive hosts:** This option is only enabled when you have an active message in the list. If the option is enabled, hosts that are currently offline are also monitored, and if such a server comes back online during the scheduled task execution, the task is applied to it too.

Sending multiple messages to users

For **Disable Host** and **Disable Host Pool** tasks, you can only send a message before the scheduled task is triggered. Hence, when creating a message, you can only select the “before” option when specifying when the message should be sent. You can create more than one message if needed and send them at different time intervals, so the users are notified more than once before the task executes.

For **Reboot Host** and **Reboot Host pool** tasks, you can send a message before or after the scheduled task is triggered. The “after” option is available for these tasks because you have the ability to enable the drain mode, which will keep the active sessions running for some time. During this time, you can send multiple messages to active users reminding them that they should finish their work and close their sessions. To use the “after” option, the **Enable Drain Mode** option must be selected. Please also note that the “after” time interval and the **Force server reboot after** setting should be coordinated. For example, if the force reboot occurs before the “after” time elapses, active users will not have a chance to see the message.

Configure logging

To configure logging and retrieve or clear existing log files, right-click a Provider, choose **Troubleshooting > Logging** in the context menu, and then click one of the following, depending on what you would like to do: **Configure**, **Retrieve**, or **Clear**. For the information on how to perform these tasks, see the [Logging](#) section. Please also read the important information below.

Note that logging of Provider operations is performed on the RAS Provider Agent level. When you configure logging for a Provider, you are essentially configuring it

for the RAS Provider Agent that services this Provider. This means that if you are using the built-in RAS Provider Agent, its logging configuration applies to all Providers that it services. Consider the following scenarios:

- When you retrieve log files for a specific Provider serviced by the built-in Provider Agent, the files will contain logs for all Providers serviced by the same agent.
- If you clear log files for a particular Provider, you should be careful because the logs will be cleared for all Providers if they are serviced by the same built-in Provider Agent. The RAS Console will prompt you if you try to delete such a shared log.

If a Provider has a dedicated Provider Agent, which services this host only, none of the above applies.

Enabling high availability for VDI

High availability for VDI means that a Provider must never lose a connection with a Provider Agent. If the connection is lost, the hosts will become unavailable for user connections. High availability for VDI is accomplished by installing at least three RAS Connection Brokers. This way, if one of the Connection Brokers goes offline (and with it the built-in Provider Agent), the Provider will be automatically assigned to the Provider Agent running in the next available Connection Broker.

To configure high availability for VDI, use the information and instructions below.

At least three Connection Brokers are required

Make sure you have at least three RAS Connection Brokers installed and running. When RAS Connection Brokers from your site are online, high availability is enabled automatically. You may also have additional Connection Brokers in standby mode, but you must have at least three agents in the active state for the high availability functionality to work. All Connection Brokers must be able to communicate with each other.

An odd number of agents is recommended

To properly control a possible split-brain situation, strictly more than half of all available Connection Brokers should be able to communicate with each other at any given time. Consider the following examples:

- Let's say there are three Connection Brokers in a Site. All of them can communicate with each other. If one of the agents suddenly loses a connection with the other two, the two agents will know that they are in the majority and will take over the Provider hosts that are currently managed by the first agent.
- Let's now say that there are four Connection Brokers. If one of them loses a connection to the remaining three, the same scenario will occur as in the example above. But if two agents simultaneously lose a connection to the other two, none of the two groups will be in the majority and therefore none will be able to make a decision who should take over the Provider hosts. In a situation like this, steps must be taken to prevent a split-brain scenario, which will happen if the agents continue to operate independently from each other. As a solution to this problem, all agents will simply abandon all Providers at the same time, so no data loss or any other problem can possibly happen.

For the reasons explained above, you should always install an odd number of Connection Brokers. This way, one of the groups of agents will always be in the majority and will continue to handle all Providers. Please note that the general recommendation (regardless of the high availability functionality described here) is to have three RAS Connection Brokers running in a Site. For details, see [Secondary Connection Brokers](#).

Please also note that Connection Brokers in [standby mode](#) don't participate in the high availability operations. These agents stay inactive until one of the active Connection Broker goes completely offline. When that happens, an agent in standby mode is activated and takes place of the lost agent. From this point forward, it is considered a part of the high availability setup. When the lost agent is brought back online, everything goes back to what it was before.

Configuring a Provider for high availability

Parallels RAS can maintain high availability in the following ways:

- Parallels RAS automatically selects a Connection Broker for the Provider. If this Connection Broker goes down, Parallels RAS moves the workload to one of the remaining Agents.
- You select a Connection Broker for the Provider manually. If this Connection Broker goes down, Parallels RAS moves the workload to one of the remaining Agents.

Use one of the following to configure a Provider for high availability:

- For an existing Provider, open the **Properties** dialog, select the **Agent Settings** tab and in the **Preferred Connection Broker** field, select **Automatic** or manually select your preferred Connection Broker.
- When adding a new Provider, on the second wizard page where you specify the host type and address, click the **Advanced Settings** link and in the **Preferred Connection Broker** drop-down list select **Automatically** or manually select your preferred Connection Broker. Note that the **Automatic** option is selected by default when there are three or more Connection Brokers available.

Site defaults (VDI)

Site defaults are settings that are defined on a Site level and can be used by templates and hosts (both template-based and non-template based). By default, templates (described later in this chapter) inherit Site default settings, but you can override them if needed when you configure a template. Non-template based hosts also use Site default settings by default and you can also override them if needed when you configure these VMs.

To view and modify Site defaults, do the following:

1. Navigate to **Farm** > <Site> > **VDI**.
2. Select the **Desktops** tab in the right pane.
3. Click **Tasks** > **Site defaults**. This opens the **Site Default Properties** dialog, which is described below.

Note that any modifications you make to Site defaults are immediately applied to all hosts in the current Site that use them.

General

The **General** tab contains the following properties:

- **Session readiness timeout:** The maximum amount of time it should require to establish a session. If the specified timeout is reached, and the session is still not ready, the user will see an error message and will have to try to log in again.
- **Protocol:** Specifies a protocol that Parallels RAS uses to communicate with a host.
- **Auto remove hosts which failed preparation after:** If a host encounters a problem during the preparation stage (for any reason), it remains on the server but cannot be used. You can identify such VMs by the “Failed to create” value in the **Status** column (**Farm** > <Site> > **VDI** > **Desktops**). Unless a VM like this is repaired, it will be automatically removed after the time period specified in this field. You can set any of the available time periods by selecting it from the drop-down list or you can type a desired value, such as “8 days” or “12 hours”.
- **Desktop assignment type:** Specifies whether the persistent hosts are assigned by the UPN (the **User** option) or device hostname (the **Device** option). Each host will be automatically assigned to the first user or device who uses a resource published from it with persistent assignment enabled. You can also assign hosts manually. For more information, see [Persistent Hosts](#).
- **Auto remove persistence if guest was not used for:** The time period after which persistence should be automatically removed. You can also type any desired time period, such as “1 week 3 days”.

Note: Beginning with RAS 17, the default setting for this option is **Never**. Please keep that in mind.

User profile

Configure this tab as described in [User profile](#).

Application Packages

Configure this tab as described in [Using MSIX application packages](#).

Optimization

Configure this tab as described in [Optimization](#).

Actions

Actions: The two drop-down lists here specify an action to perform on session disconnect or logoff.

Note for Nutanix AHV (AOS): Nutanix AHV (AOS) does not support the suspend operation for its VMs. If **Suspend** is selected in the **Perform action** field, no action will be applied to a Nutanix AHV (AOS) VM when a session disconnect occurs (a corresponding error will be recorded in the Provider Agent log).

Security

On the **Security** tab, you can specify whether to automatically grant users Remote Desktop connection permissions on hosts. Here's how it works. Instead of manually adding each user to the Remote Desktop Users (or Administrators) group, you can enable this option to do it automatically. When a user logs on, he/she will be automatically added to the specified group and will therefore have the Remote Desktop connection (or full Administrator) permissions on the server. When the user logs off, they will be removed from the group (i.e. the group membership will only exist for the duration of the session).

The more important benefits of this feature are as follows:

- You don't have to permanently add your users to the Remote Desktop Users groups. This way, a user will never be able to establish a Remote Desktop session with a server outside of Parallels Client.
- By automatically adding a user to the Administrators group, you can give them rights to install applications and perform other administrative tasks. Once again, the user will only be able to do it from Parallels Client but never by connecting to the server using standard Remote Desktop tools.

Settings

The **Settings** tab contains the following:

- **Disconnect active session after:** The amount of time a session remains logged in after the user closes a published application. The default timeout is 25 seconds. Note that this only works for applications, but not published desktops (when a user closes a desktop, the session is logged off). This timeout is used to avoid unnecessary logins when a user closes one application and then opens another.
- **Preferred Connection Broker:** Select a preferred Connection Broker to which this Provider should be assigned. This can be helpful when Site components are installed in multiple physical locations communicating through WAN. You can decrease network traffic by specifying a more appropriate Connection Broker.
- **Allow URL/Mail redirection:** When a user tries to open a URL or an HTML Mailto link in a remote application, the link can be redirected to the client computer and open in a local default application (a web browser or email client) instead of an application on the remote host. This option allows you to enable or disable the redirection. Click the **Configure** button to choose from the following options:
 1. **Replace registered application** — this option uses an alternative method of redirecting a link. It replaces the default web browser and mail client with “dummy” apps on the remote server side. By doing so, it can intercept an attempt to open a link and redirect it to the client computer.
 2. **Support Windows Shell URL namespace objects** — the Shell URL namespace objects support means that Parallels RAS can intercept actions in published applications that use Shell namespace API to open links, which is a standard behavior in most applications. The ability to disable the support for Shell URL namespace objects is for compatibility with older versions of Parallels RAS.
- **Enable drag and drop:** Allows you to set how the drag and drop functionality works in Parallels Clients. Click **Configure** and choose from “Disabled” (no drag and drop functionality), “Sever to client only” (drag and drop to a local application only), “Client to server only” (drag and drop to a remote application only), “Bidirectional” (drag and drop in both directions).

Note: At the time of this writing, the drag and drop functionality is only supported on Parallels Client for Windows and Parallels Client for Mac.

- **Manage RDP transport protocol:** Selects the transport protocol that will be used for connections between Parallels Client and a server. To do this, select this option and click the **Configure** button.
- **Allow file transfer command (Web and Chrome clients):** Enables file transfer in a remote session. To enable file transfer, select this option and click the **Configure** button. For more information, see [Configuring remote file transfer](#).
- **Enable drive redirection cache:** Improves user experience by making file browsing and navigation on redirected drives much faster. For details, see [Drive redirection cache](#).
- **Enable Z-order (Experimental):** Select this option if you encounter an issue related to z-order in applications published with Parallels RAS.

RDP printer

The **RDP Printer** tab allows you to configure the renaming format of redirected printers. The format may vary depending on which version and language of the server you are using. Select the **RDP Printer Name Format** option specifically for the configured server:

- **Prntername** (from Computername) in Session no.
- Session no. (computername from) **Prntername**
- **Prntername (redirected Session no)**

The other RDP Printing option available is **Remove session number** from printer name, which will do what it says.

Using computer management tools

You can perform standard computer management tasks on a server right from the RAS Console. The tasks include Remote Desktop Connection, PowerShell, Computer Management, Service Management, Event Viewer, IPconfig, Reboot, and others. To access the **Tools** menu, select a server, click **Tasks > Tools** and choose a desired tool. For the complete description, please see [Computer management tools](#).

Viewing Provider summary

In addition to the Provider editor described earlier in this chapter, you can also see summary about the available Providers. To do so:

1. In the RAS Console, select the **Farm** category and then select the **Site** node in the middle pane.
2. The available servers are displayed in the **VDI** section in the right pane.
3. To go to the Provider editor, right-click a server and choose **Show in the Editor**.

For additional info, see [Sites in the RAS Console](#).

Remote PC pools in VDI

Remote PC pools is a Parallels RAS feature that allows you to create pools of standalone (preferably domain-joined) PCs and optionally assign them to specific users. The Remote PC pools functionality is integrated into RAS VDI to take advantage of the infrastructure that already handles host pools.

Remote PC pools vs. Remote PCs

Remote PCs are standalone machines (physical or virtual) that can be used to host published resources in Parallels RAS. Remote PCs are managed in the Parallels RAS Console in the **Farm** > <Site> > **Remote PCs** section. The [Remote PCs](#) chapter describes this functionality in detail. Remote PC pools described here are handled separately and differently from standalone Remote PCs. They are managed in the **Farm** > <Site> > **VDI** section of the RAS Console.

Adding a Provider

To set up a Remote PC pool in the RAS Console, you first need to add a Provider of type **Remote PC**. This is a special type that exists specifically for the purpose of creating and managing Remote PC pools. It is not a real Provider, so it doesn't need a hypervisor installed. It simply uses the existing VDI functionality to create and manage computer pools. Note that when you add a Provider of this type, you can manage it like any other real Provider with some limitations, such as you

cannot create templates and use some other strictly VDI (hypervisor)-specific functionality.

To add a Provider of type **Remote PC**:

1. Navigate to **Farm** > <Site> > **Providers**.
2. On the **Providers** tab, click **Tasks** > **Add**.
3. Select one of the following:
 - **Remote PC dynamic**: This approach assigns PCs using the information from Active Directory. All you have to do is specify an organizational unit (OU) containing computer accounts to be assigned to the host.
 - **Remote PC static**: Using this approach, Remote PCs are assigned to the Provider by entering their FQDN or IP address (one by one) or by importing a list from a CSV file.
4. In the wizard that opens, specify the following:
 - **Name**: Name of the provider.
 - **Description**: Description of the provider.
 - **Address**: FQDN or IP address of a server that will manage Remote PC pools. This must be a server with RAS Provider Agent installed. You can use the RAS Connection Broker server, since it has the RAS Provider Agent built in, but it can be any other server running a dedicated RAS Provider Agent.
 - **Username**: Account name in the UPN format (e.g. administrator@domain.local). This must be a domain user account with administrative rights on the server specified above. Using a local Windows account is also possible with some limitations and only when using the static PC assignment (see below). Using a domain account is recommended.
 - **Password**: Account password and an optional description.
5. Click the **Manage Credentials** button to specify the accounts that will be used to deploy RAS agents.

6. Click the **Advanced Settings** link to open the **Advanced Provider Settings** dialog. The dialog allows you choose the following options:

- **Use dedicated Provider Agent:** Select this option if you will install (or have installed) the RAS Provider Agent yourself. Clear the option if you will use the [built-in RAS Provider Agent](#).
- **Agent address:** This option becomes enabled if you select the option above it. Specify the FQDN or IP address of the server where the RAS Provider Agent is (or will be) installed. This can be either a physical box or virtual machine.
- **Preferred Connection Broker:** Select a RAS Connection Broker to be the preferred agent for this Provider. For more info, see [Enabling high availability for VDI](#).

7. Click **Next**.

8. The wizard will display the new Provider information and will indicate the RAS Provider Agent status. If everything is OK, click **Next**.

9. If you selected **Remote PC dynamic** in step 3, specify the following:

- **Target OU:** Organizational unit (OU) containing computer accounts to be assigned to the host. You can click the [...] button to browse Active Directory. Note that a maximum of 1000 Remote PCs can be returned in a single AD/OU search result.

Note: When using dynamic assignment, Remote PCs must be domain-joined. You cannot manage such PCs using a local Windows user account.

When you use the dynamic assignment, you have an option to install RAS Guest Agent on every PC by adding a Group Policy to the organizational unit with a script to deploy RAS Guest Agent. The following is an example of such script:

```
msiexec /i RASInstaller-<version & build>.msi  
ADDLOCAL=F_GuestAgent /qn+ /norestart
```

Other agent installation options are described in [RAS Guest Agent installation options](#).

- **Subnet mask:** Subnet mask used for calculating the directed broadcast address from the IP addresses of the Remote PCs. It is

used for sending a directed broadcast of the Wake on LAN magic packet.

10. If you selected **Remote PC static** in step 3, do one of the following:

- Click **Tasks > Add** and type FQDN or IP address of a PC you want to add. You can click the [...] button to search for it. Next, specify the subnet mask used for calculating the directed broadcast address from the IP address of the Remote PC. It is used for sending a directed broadcast of the Wake on LAN magic packet. After this, enter the MAC address of the computer you are adding. Note that all fields are mandatory.
- Click **Tasks > Import from CSV file** and then select a CSV file containing the list of computers. The CSV file must contain two columns: (1) FQDN or IP address; (2) MAC address. Once again, both columns are mandatory and must contain valid values.

Parallels RAS 18 (and newer) supports a maximum of 1000 Remote PCs per Provider.

Note: To be manageable, Remote PCs should be domain-joined. In case of static assignment described here, it is possible to add non-domain joined PCs, but you will have to create the same local user account on each and everyone of them. Using a domain account and domain-joined PCs is recommended.

- Click **Finish**.

Adding Remote PCs to a Provider

You can add Remote PCs to a Provider in two ways:

- In the wizard while adding the provider as described in [Adding a VDI Host](#).
- After creating the provider, as described in this section.

To add Remote PCs to a Provider after it has been created:

1. Right-click a Provider that you created and choose **Properties**.
2. Select the **Remote PCs** tab.
3. Add Remote PC as described in steps 9 and 10 in [Adding a VDI Host](#).

Adding Remote PCs to a pool

Note: To be managed in a Remote PC pool, a Remote PC must have RAS Guest Agent installed. For more information, see the [RAS Guest Agent installation options](#).

Once you assigned PCs to a Provider, you can add them to a Remote PC pool as follows:

1. In **Farm** > <Site> > **VDI**, select the **Pools** tab.
2. Add a new pool by clicking **Tasks** > **Add** in the **Pools** pane.
3. Select the pool that you've created and then in the **Members** pane, click **Tasks** > **Add** and choose one of the following:
 - **All Hosts in Host:** Adds all Remote PCs assigned to the Provider. When you click this option, a dialog opens allowing you to select a Provider. Select the host and click **OK**.
 - **Host:** Adds an individual Remote PC. In the dialog that opens, select a desired Remote PC and click **OK**. Another dialog may open asking you to upgrade RAS Guest Agent on a Remote PC (the agent is required for a PC to be managed in a pool). Click **OK** to upgrade (or install) the agent. You can also upgrade the RAS Guest Agent on one or more PCs at another time as described in [RAS Guest Agent installation options](#).

Once you add one or more Remote PCs to a pool, they will appear in the **Pool management** tab and in the **Desktops** tab.

Tip: If you need to disable the pool for maintenance, you can do so by clearing the checkbox in front of the pool name.

Managing Remote PCs in a pool

Management of pool-based Remote PCs includes assigning a PC to a specific user, upgrading the RAS Guest Agent, viewing and modifying PC properties, performing some standard administrative tasks, and some others.

To manage Remote PCs in a pool:

1. In **Farm** > <Site> > **VDI**, select the **Desktops** tab.

2. Note that the list on this tab includes all managed desktops, including hosts and pool-based Remote PCs. You can order the list by the **Pool** column to see Remote PCs assigned to a particular pool.
3. Select a Remote PC, click the **Tasks** drop-down list and choose one of the options described below. Note that not all options available in the **Tasks** menu are applicable to Remote PCs. The list below describes only the options that you can use with pool-based Remote PCs.

The **Tasks** menu options that apply to Remote PCs are:

- **Upgrade all Agents.** Upgrade RAS Guest Agent in all Remote PCs (and hosts) in the list.
- **Assign.** Assign a Remote PC to a specific user (make a PC persistent). Click the menu option and specify a user.
- **Unassign.** Remove the user assignment (persistence) from a Remote PC.
- **Show sessions.** Switches the view to the **Sessions** tab and displays the session information.
- **Tools.** Allows to perform a set of standard operations, such as establishing a remote desktop connection, pinging, rebooting/shutting down a Remote PC, and others. For the description of power operations, please see **Performing power operations** below.
- **Troubleshooting.** Check and install/upgrade the RAS Guest Agent in a Remote PC.
- **Reset properties.** Resets Remote PC properties to their default values. See **Properties** below.
- **Properties.** Opens a dialog where you can view and modify Remote PC settings. The **General** tab allows you to temporarily disable the Remote PC in a pool (use the **Do not use this host** option). This is specifically useful when you need to perform maintenance tasks on a PC. You can also view and modify the Remote PC display name, computer name, and the port number on which it communicates with the Provider. For the description of **Settings** and **Security** tabs, see [Site defaults](#).

Performing power operations

For remote power operations to work, WMI must be enabled in Windows running in a VM and TCP ports 30004 and 30005 must be open. At the time of this writing, this is not automated in Parallels RAS but will be in future versions.

To perform power operations on a host (start, stop, restart, suspend, reset), open the **VDI > Desktops** tab, select a host, then click **Tasks** and choose an operation

that you want to perform (for start and stop operations, you can click the corresponding icons at the top). The restart operation (graceful) has a 10 min timeout. If not completed during this time, the reset operation (forced) will be used.

Please note if you are using Nutanix AHV (AOS), the suspend operation is not available (the **Suspend** icon is disabled). The reason for this is Nutanix AHV (AOS) does not support the suspend operation on its virtual machines.

Persistent Remote PCs

A persistent Remote PC is a PC assigned to a particular user. Once a PC is assigned, no other user can connect to it.

There are two ways to make a Remote PC persistent:

- When you publish a resource (application, desktop, etc.) from a pool-based Remote PC using the publishing wizard, you can select the **Persistent** option in the **Virtual Guest Settings** section. This way, a Remote PC in a pool will be assigned to the first user that opens the published resource. For more info, see **Publishing from a pool-based Remote PC**.
- You can also assign a Remote PC to a user manually. To do so, navigate to **Farm > <Site> > VDI**, select the **Desktops** tab, then select a Remote PC in the list and click **Tasks > Assign**. In the dialog that opens, specify the target user.

To remove persistence from a Remote PC, select it in the **Desktops** tab and click **Tasks > Unassign**.

RAS Guest Agent installation options

To be managed in a Remote PC pool, a Remote PC must have RAS Guest Agent installed. This can be done using one of the following options:

- When you add an individual Remote PC to a pool, you'll be asked to upgrade the agent. Follow the onscreen instructions and install or upgrade it.
- When you add all Remote PCs in a host to a pool at once, you can add them first and then use the **Tasks > Upgrade all Agents** menu option in the **Desktops** tab.

- When you assign Remote PCs to a Provider via Active Directory, you can have a Group Policy in the OU with a script to deploy the agent. See [Configuring the Provider > Dynamic \(VDI subtype\)](#).
- To install or upgrade the agent on an individual Remote PC, select it in the **Desktops** tab and click **Tasks > Troubleshooting > Check agent** option. In the dialog that opens, click **Install**.
- Finally, you can install RAS Guest Agent manually by running the Parallels RAS installer on a Remote PC and selecting to install the RAS Guest Agent component.


Azure Virtual Desktop

Azure Virtual Desktop (formerly known as Windows Virtual Desktop) is a desktop and app virtualization service running on Microsoft Azure, providing access to RD Session Hosts and VDI, including the new offering of Windows 10 and Windows 11 Enterprise multi-session hosts. Parallels RAS 18 provides the ability to integrate, configure, maintain, support and access Azure Virtual Desktop workloads on top of the existing technical capabilities of Parallels RAS.

Introduction

The diagram below illustrates a hybrid deployment of Parallels RAS and Azure Virtual Desktop with the following characteristics:

- Workload hosts are available both on-premises through standard Parallels RAS deployment and on Microsoft Azure through the service.
- Azure Virtual Desktop objects such as workspaces, host pools, desktop and RemoteApp groups are created and configured from the Parallels RAS Console.
- Azure Virtual Desktop hosts (multi-session or single-session) contain both Azure Virtual Desktop Agent and RAS Agent for management and configuration purposes.
- Parallels Client for Windows is connecting to both Parallels RAS Secure Gateway and Azure Virtual Desktop service providing resource availability to end-users from a single interface.



Microsoft Azure subscription

You need a Microsoft Azure subscription, including:

- An Azure Tenant ID.
- An Azure subscription with sufficient credit.

Azure Virtual Desktop user license entitlement

Customers with the licenses listed below are entitled to use Azure Virtual Desktop at no additional charge apart from Azure compute, storage, and network usage billing.

To run Windows 10 and Windows 11 with Azure Virtual Desktop, you need to have one of the following per-user licenses:

- Microsoft 365 F3, E3, E5, A3, A5, Student Use Benefits or Business Premium
- Windows 10 Enterprise E3, E5
- Windows 10 Education A3, A5
- Windows 10 VDA per user

To run Windows Server 2012 R2, 2016, 2019, 2022:

- Per user or per device Remote Desktop Services (RDS) Client Access License (CAL) with active Software Assurance (SA).

For further information, please refer to Microsoft licensing requirements at <https://docs.microsoft.com/en-us/azure/virtual-desktop/overview>.

Permissions and Azure resource providers

The following highlights permissions and resource providers to be registered in the subscription:

- Permissions to enable resource providers on your Azure subscription and create virtual machines (VMs).
- The necessary Microsoft Azure resource providers (**Azure Portal > Subscription > Resource Providers**) must be enabled, including **Microsoft.ResourceGraph**, **Microsoft.Resources**, **Microsoft.Compute**, **Microsoft.Network**, **Microsoft.DesktopVirtualization**.

Microsoft Entra ID application

For detailed information about creating a Microsoft Entra ID application, please see [Create a Microsoft Entra ID application](#).

Once a Microsoft Entra ID Application is created, give the application the following API permissions in the Microsoft Azure Portal (**Microsoft Entra ID > App Registrations > API permissions > Add a permissions > Microsoft.Graph > Application permission**):

- **Group > Group.Read.All**
- **User > User.Read.All**

Note: Please make sure that when adding Graph API permissions, User and Group, the permission type is “Application”, not “Delegated”.¹

Give the application read and write access to resources:

- The Microsoft Entra ID application that you created must have read and write access to Azure resources as described in [Create a Microsoft Entra ID application](#). Look for “Give the application read and write access to resources”.

Roles and permissions for the application should include:

- “User Access Administrator” role for the application from **Subscription > Access Control (IAM)**.
- “Contributor” role at the Resource group level from **Resource group > Access Control (IAM)**.

If a resource group creation is required, also assign the contributor role at the subscription level **Subscription > Access Control (IAM)**.

Note: If you would like to also view/read resources outside the resource group, make sure that the application is also given read access at the subscription level.

Active Directory

- A Server Active Directory environment or Azure Active Directory Domain Services (AADDs). See <https://azure.microsoft.com/services/active-directory-ds/>.
- Azure AD Connect — AD must be in sync with your Microsoft Entra ID, so users can be associated between the two.

- The user must be sourced from the same Active Directory that's connected to Microsoft Entra ID. Azure Virtual Desktop does not support B2B or MSA accounts.
- The user configured in the Parallels client with access to Azure Virtual Desktop resources must exist in the Active Directory domain the session host it is joined to.

Other

- Azure Virtual Network providing session hosts connection to the domain.
- Session hosts must be domain-joined to Active Directory.
- (optional) Site-to-site VPN or ExpressRoute is required if a hybrid Parallels RAS deployment is used.
- (optional) Shared network location to be used for FSLogix Profile Containers, which may run on Azure Files or Azure NetApp Files.

Note: At the time of writing, Windows 7 is not supported by Parallels RAS as an Azure Virtual Desktop session host.

Attention: Effective September 30, 2025, Microsoft retired default outbound access. The Subnet and VNet you use must explicitly allow outbound connectivity. For more information, see [Azure's outbound connectivity methods](#) for the most commonly used approaches.

Additional notes

Please also note the following Provider and Azure Application requirements for different RAS Farm and RAS Site scenarios:

- Same RAS Farm, same RAS Site: The same Farm, Site, and Application ID is possible to be used for both VDI and Azure Virtual Desktop. Build the guest VM list with Azure Virtual Desktop tags for Azure Virtual Desktop provider and guest VMs with VDI tags (or no tags) for Azure Provider.
- Same RAS Farm, same RAS Site: It is recommended to use different Azure Applications for multiple providers of the same type. For example, multiple Azure Virtual Desktop or multiple Providers, but not mixed.
- Same RAS Farm, different RAS Sites, or different RAS Farms: The point above applies. Alternatively, different RAS Farms or Sites can (and must in this case) reside in different virtual networks with no communication to a common set of VMs.

Important: It is recommended that Parallels RAS-managed Azure Virtual Desktop objects are managed through the Parallels RAS console. Configuration changes outside the Parallels RAS console may result in a broken state of Azure Virtual Desktop objects. For such cases, Parallels RAS provides the ability to repair objects. For example, auto-created friendly names and associated tags for workspaces and host pools can also be viewed from the Microsoft Azure portal; however, they are not to be edited as these are used to ensure proper functionality.

Azure Virtual Desktop deployment in Parallels RAS is done by completing a series of wizards, including:

1. Enable Feature and Add Azure Virtual Desktop Provider.
2. Add an Azure Virtual Desktop Workspace.
3. Add an Azure Virtual Desktop host pool and then add standalone or template-based hosts to the host pool.
4. Publish Azure Virtual Desktop resources.

You can run all from the **Start** category as part of a single deployment procedure. Read on to learn how to do it.

Enable Azure Virtual Desktop and add a provider

Azure Virtual Desktop integration must first be enabled in the RAS Farm. This can be done from two places in the RAS Console:

- Using the **Deploy Azure Virtual Desktop** wizard in the **Start** category.
- By going to **Farm > Site > Settings** and selecting the **Features** tab.

The instructions below are for enabling and deploying Azure Virtual Desktop from the **Start** category. The **Features** tab in **Site > Settings** has the same elements as the **Enable Feature** page described below.

Note: If you haven't enabled Azure Virtual Desktop in the RAS Farm yet, the wizard pages will open in the order described below. If Azure Virtual Desktop is already enabled (e.g. you ran the wizard before or enabled Azure Virtual Desktop from Site settings), the first two pages will be skipped and the first page you'll see is **Add Azure Virtual Desktop Provider** where you need to enter the provider information.

To begin the deployment:

1. In the Parallels RAS Console, select the **Start** category and launch the **Deploy Azure Virtual Desktop** wizard.
2. **System and user requirements:** On the first page, read system and user requirements. Click a link at the bottom of the page to read a Parallels KB article for more information. Click **Next**.
3. The **Client feature set** selection specifies which client features will be available when you open a published resource in Parallels Client. Select from the following options:
 - **Standard:** Standard feature set. This is identical to opening and running a published resource using the Microsoft Windows Desktop client, also known as Remote Desktop (MSRDC) client, or the Microsoft Windows App, which are the clients used to access apps and desktops from Azure Virtual Desktop.
 - **Advanced:** This option also uses the Windows Desktop client but adds support for drag and drop, and session prelaunch.
 - **Advanced with fallback:** This option first tries to open a published resource using the Advanced feature set. If Advanced doesn't work for any reason, it will try to open the resource using the Standard option.

[Note: The Advanced feature set is deprecated. The Advanced or Advanced with fallback options should no longer be used. We recommend using the Standard mode for all new and existing deployments to ensure compatibility and long-term support.](#)²

4. This completes the task of enabling Azure Virtual Desktop in the RAS Farm. Click **Next** to advance to the next page.
5. **Add Azure Virtual Desktop Provider:** On this page you need to specify your Microsoft Azure Tenant ID, Subscription ID, Application ID, and a secret key. This is similar to setting up Microsoft Azure as a Provider in Parallels RAS. For the explanation of how to specify these properties, please see [Add Microsoft Azure as a Provider](#). Please note that under subscription details, URIs/URLs may be edited during creation of a provider. The **Feed URL** setting, which by default is <https://rdweb.wvd.microsoft.com/api/arm/feeddiscovery>, may also be edited once an Azure Virtual Desktop provider is created.

6. Click **Next**, review the summary and click **Finish**. Note that changing Microsoft Azure app permissions after a provider is created may require a restart of the Parallels RAS redundancy service so new permissions are loaded and used.

This completes the first wizard in the series. On the last page, the **Launch Azure Virtual Desktop Workspace wizard** option is enabled by default. This will automatically open the next wizard where you can add an Azure Virtual Desktop workspace.

A workspace is a logical grouping of application groups in Azure Virtual Desktop. Each Azure Virtual Desktop application group must be associated with a workspace for users to see the remote apps and desktops published to them.

To add a workspace:

1. Select whether you want to create a new workspace or select an existing one:
 - To select an existing workspace, click the [...] button next to the **Name** field.
 - To create a new workspace, type a name and optional description. Select an existing or create a new resource group. Specify a location. Note that the location that you select here will be used for all Azure Virtual Desktop objects, including workspaces, host pools, and application groups.

In the **Friendly name field**, specify the friendly name that will be used for the workspace in Azure Virtual Desktop and Parallels RAS.

2. Click **Next**, review the summary and click **Finish**.

The last page of the wizard has the **Launch Azure Virtual Desktop host pool wizard** option selected by default. This will automatically open the **Add Azure Virtual Desktop Host Pool** wizard when you click **Finish**.

Add host pools (Azure Virtual Desktop)

A host pool is a collection of one or more identical virtual machines (VMs) within an Azure Virtual Desktop environment. Each host pool contains an application group that users can access.

To add a host pool:

1. In the RAS console, navigate to **Farm > <Site> > AVD > Host pools**.
2. Click the **Tasks** drop-down list above the Pools list and then click **Add** (or click the plus-sign icon). This opens the **Add AVD host pool** wizard.
3. Select whether you want to create new or select an existing host pool:
 - To select an existing workspace, click the [...] button next to the **Name** field.
 - To create a new host pool, select **Create new host pool** and select the provider, workspace, name, description, resource group, and location.

In the **Friendly name field**, specify the friendly name that will be used for the host pool in Azure Virtual Desktop and Parallels RAS.

4. Click **Next**.
5. On the **Configuration** page, specify the following:
 - **Host pool type:** Select from **Pooled** (multi-session hosts) or **Personal** (single-session hosts).
 - **Publishing type:** Select from **Application** or **Desktop** depending on what you want to use the pool for.
 - **Load balancer:** Select a load balancer type. Breadth-first load balancing allows you to evenly distribute user sessions across the session hosts in a host pool. Depth-first load balancing allows you to saturate a session host with user sessions in a host pool. Once the first session host reaches its session limit threshold, the load balancer directs any new user connections to the next session host in the host pool until it reaches its limit, and so on.
 - **Limit number of sessions on host:** For a pooled (multi-session) pool type, specify the maximum allowed number of sessions on a host.

- **Power on host on-demand:** Specify whether a powered down host should be powered on when a user tries to connect to it. Note that this applies only if all session hosts in the host pool are powered off.
- **Default license type:** Select the Azure license type.
- **Service updates validation:** Select the **Validation environment** option if you want to make this host pool a validation environment for Microsoft service updates.

6. Click **Next**.

7. On the **Provisioning** page, select whether this host pool will contain template-based or standalone hosts:

- **Template:** Hosts will be created dynamically from a template. You will need to create or select an existing template in the next step or later. Choosing **Template** as the provisioning type ensures a homogeneous host pool, which is recommended to provide consistent user experience across the host pool. For information on how to create a template, see [Create a template](#).
- **Standalone:** Select one or more hosts that already exist. You'll be able to do it in the next step or you can do it later. Prior to adding hosts to host pools, ensure that hosts are domain joined and have network access to the domain environment. Note that the Standalone provisioning is considered "unmanaged" as it lacks some of the functionality, such as Autoscaling.

8. Click **Next**.

9. Depending on the selection made on the **Provisioning** page (above), do one of the following

- **Standalone:** Select one or more hosts from the list to be included in the host pool (you can also add hosts to the pool later).
- **Template:** Select a template from the list or click **Create new** to create a new template and specify the template settings. **Versions:** If you selected an existing template, select one of its versions. **Enable autoscale:** (Multi-session hosts) Enable and configure autoscale.

10. Click **Next**.

11. (Templates only) On the **Properties** page, specify the following options:

- **Template name:** Choose and type a template name.

- **Maximum hosts:** Specify the maximum number of hosts that can be created from this template.
- **Number of hosts deployed on wizard completion:** The number of hosts to deploy once the template is created. Please keep in mind that this will take some time because the hosts will be created one at a time.
- **Host prefix:** A pattern to use when naming new hosts.

12. Click **Next**.

13. (Templates only) On the **Settings** page, specify the following options:

- **Keep available buffer:** The minimum number of hosts to always keep unassigned and session free for the template. As soon as the number of free and unassigned desktops drops below the setting value, it forces the template to create another host. The template uses its own settings for host creation including initial power state.
- **Host state after the preparation:** Select the power state that should be applied to a host after it is prepared. Choose from **Powered on**, **Powered off**, or **Suspended**. Note that when the power state is set to **Power off** or **Suspended**, the number of running (fully ready and waiting for incoming connections) hosts is controlled by the **Keep available buffer** setting (see above). For example, let's say the **Maximum hosts** value is set at 200, the number of guest hosts deployed on wizard completion is 100, and the power state after preparation is **Powered off**. The result of such a configuration will be 100 clones deployed and powered off.
- **Delete unused hosts after:** Select what to do with unused hosts to save resources. Choose whether to never delete them or specify the time period after which they should be deleted.

14. Click **Next**.

15. (Templates only) On the **specifications** page, override settings that are defined at the template level.

- To override the Organizational Unit in Active Directory where hosts are joined, select **Override Organizational Unit specified in template properties**. In the target OU field, provide the distinguished name of the Organizational Unit by using the browse (...) button or entering it manually.
- To override the High availability settings, select **Override the high availability settings specified in template properties**. Select an

availability option which can be no infrastructure redundancy required, availability set, or availability zone. When availability zone is selected, specify the zone where hosts should be deployed. Note that High availability settings cannot be changed after the host pool is created.

- To override the host size, select the **Overwrite the size specified in template properties** option and choose the desired size from the drop-down list. The selected size will only be used by this host pool. Other host pools using the same template will be unaffected. Note that available sizes may depend on the location, size, and power state of the host pool members and the template. Also note that overwriting the size requires a host reboot.

16. On the **Assignment** page, specify users or groups to be assigned to the application group in the host pool. This is necessary for users to have access to published applications or desktops. Click **Tasks > Add** and specify a user or group. An application group of type **Desktop** or **RemoteApp** (whichever is appropriate) will be created and associated with the host pool automatically on wizard completion.
17. On the **User profile** page, you can select from **Do not manage by RAS** (user profiles will not be managed) or **FSLogix**. Microsoft FSLogix Profile Container allows to maintain user context in non-persistent environments, minimize sign-in times and provides native profile experience eliminating compatibility issues.
18. Follow the onscreen instructions and complete the wizard.
19. On the **Summary** page, review the template summary information. You can click the **Back** button to correct some of the information if needed.
20. Finally, click **Finish** to create the host pool and close the wizard.

Note: In the case of using the Advanced Client Feature Set, RemoteApp groups are not required for publishing applications since the Desktop App Group with Parallels seamless technology will be used to provide application publishing from configured desktop app groups.

Next step

[Verify the Azure Virtual Desktop deployment](#)

Manage Azure Virtual Desktop

Manage providers (Azure Virtual Desktop)

An Azure Virtual Desktop provider in Parallels RAS is a collection of IDs and other properties that give you access to Azure resources. Properties include Tenant ID and Subscription ID, among others. Normally, an organization is given one Tenant ID by Microsoft, but there could be multiple subscription IDs owned by the same organization. For each Tenant ID and subscription ID combination, a provider must be configured in Parallels RAS.

To manage Azure Virtual Desktop providers, navigate to **Farm > Site > Providers** and select the **Providers** tab.

To add a new provider, click **Tasks > Add** and select **Azure Virtual Desktop**. For the information on how to complete the wizard, please see [Add Microsoft Azure as a VDI Host](#). To view and modify some of the existing provider properties, right-click a provider in the list and choose **Properties**.

Other provider management tasks can be accessed from the **Tasks** menu, including:

- Checking provider status: **Tasks > Troubleshooting > Check status**.
- Configure and manage logging: **Tasks > Troubleshooting > Logging**.

Choosing the Active Directory Domain Services type

Parallels RAS 18.3 supports virtual machines joined to Windows Server Active Directory Domain Services and Azure Active Directory Domain Services. By default, Parallels RAS is configured to work with Windows Server Active Directory Domain Services, but you can change this if needed.

To choose the type of Active Directory Domain Services for a provider:

1. Right-click a provider in the list and choose **Properties**.
2. In the provider properties window, select the **Credentials** tab.

3. In the **Active Directory Domain Services type** drop-down list, select the type of Active Directory Domain Services:
 - If your users are created with Windows Server Active Directory and your virtual machines are joined to Windows Server Active Directory Domain Services, select **Windows Server AD DS** (selected by default).
 - If your users are created with Windows Server Active Directory and your virtual machines are joined to Azure Active Directory Domain Services, select **Azure AD DS**.
 - If your users are created with Microsoft Entra ID and your virtual machines are joined to Azure Active Directory Domain Services, select **Azure AD DS**.

Disk storage cost optimization

You can configure an Azure Virtual Desktop provider to automatically change the type of the used managed disk to Standard HDD for AVD hosts that are not currently in use. When an AVD hosts is started, the managed disk is automatically changed to the original type. This feature allows you to reduce the cost of maintaining AVD hosts.

To enable disk storage cost optimization:

1. Right-click a provider in the list and choose **Properties**.
2. In the provider properties window, select the **Advanced** tab.
3. Select the **Enable disk storage cost optimization** option.
4. Select the desired option in the **Set timeout before enabling storage cost optimization** drop-down list.

Manage workspaces (Azure Virtual Desktop)

A workspace is a logical grouping of application groups in Azure Virtual Desktop. Each Azure Virtual Desktop application group must be associated with a workspace for users to see published remote apps and desktops.

To manage Azure Virtual Desktop workspaces, navigate to **Farm > Site > Azure Virtual Desktop** and select the **Workspaces** tab.

To add a workspace:

1. Click **Tasks > Add** to open the **Add Azure Virtual Desktop Workspace** wizard.
2. Select a provider at the top of the wizard page (if you have more than one). You can also create a new provider right from this page. If you wish to do so, click the **New provider** button to open another wizard. For details, see [Manage providers](#).
3. After selecting (or creating) a provider, complete the workspace wizard as described in [Add an Azure Virtual Desktop workspace](#).

To view properties of an existing workspace, right-click it and choose **Properties**. You can enable or disable the workspace and modify the workspace description and friendly name. Other properties are read-only. Note that if you disable the workspace, all associated objects, including host pools and published resources will also be disabled.

Manage host pools (Azure Virtual Desktop)

A host pool is a collection of one or more identical virtual machines (VMs) within an Azure Virtual Desktop environment. Each host pool contains an application group that users can access.

Host pools can be configured a number of different ways depending on the intended purpose. The following table describes different options that you can choose when creating a host pool.

Option	Description
Personal vs. pooled	
Application vs. desktop	A host pool can only publish applications or desktops, but not both at the same time. When you create a host pool, you choose a publishing type from Desktop or Application . An application group of the appropriate type (Desktop or RemoteApp) for the host pool is created automatically. Note that you cannot change the publishing type later. If you decide that you want to change it, you'll have to delete the existing host pool and create a new one.

Option	Description
Template vs. standalone	

To manage Azure Virtual Desktop host pools, navigate to **Farm > Site > Azure Virtual Desktop** and select the **Host pools** tab.

Host pool properties

To view and modify host pool properties, right-click it and choose **Properties**. In the dialog that opens, select tabs and view or modify host pool properties as described below.

General

On the **General** tab, you can enable or disable the host pool. Note that if you disable it, all hosts and published resources will also be disabled.

You can also modify the host pool description and view general host pool properties.

In the **Application group** section, you will see the name of the application groups created for the host pool.

The **Friendly name** field shows the friendly name used for the workspace in Azure Virtual Desktop and Parallels RAS.

Configuration

On the **Configuration** tab, examine the host pool configuration properties. You should be familiar with them from when you created a host pool.

You can modify the following properties on this page:

- **Load balancer**
- **Limit number of sessions on host**
- **Power on host on-demand**
- **Validation environment**

For the explanation of configuration properties, see [Add an Azure Virtual Desktop host pool](#).

Autoscale

This tab is shown only for host pools with **Template** as a provisioning type. Here you can select a template if one has not been specified for the host pool yet. You can also create a new template by clicking the **Create new** button, which will open a wizard. If you don't have any templates, the only selection available is **None**, which means that there's no template to create hosts from. If that's the case, you need to create a template first and then select it here. See also [Manage templates](#).

The **Autoscale settings** section contains settings that determine how hosts (virtual machines) are created from the specified template. These settings work the same as Autoscale settings for RD Session Host groups. The only difference is, in Azure Virtual Desktop we deal with hosts and host pools, while in RD Session Host groups we talk about servers and groups, otherwise the settings work in a similar manner. For details, please see the **Autoscale** subsection in [Grouping and cloning RD Session Hosts](#).

Hosts

The **Hosts** tab lists hosts from this host pool. You can examine the status of a host and other properties by looking at the values in the table.

The **Status** column should indicate "OK" if a host is operating normally. To verify the agent status, right-click a host and choose **Check agent**.

Installing agents

If you see a message **Agent did not reply**, click **Install** to try and install the agent. If everything goes well, the agent will be updated and the **Status** column should say **OK**. You can also upgrade all agents by clicking **Tasks > Upgrade all Agents**.

Repairing agents

If you see **Broken (Host unavailable)** in the **Status** column, the Azure Virtual Desktop agent has become unresponsive.

To repair the agent:

- Click **Tasks > Repair**.

This will uninstall the Azure Virtual Desktop agents, remove the host from the host pool, reinstall the Azure Virtual Desktop agents and add them to the host pool. During this process the **Status** column will show the status **Deployment in progress**.

Adding hosts to a pool

To add a new host to the pool:

1. Click **Tasks > Add**.
2. Depending on the host pool provisioning type, do one of the following:
 - If the host pool provisioning is configured as **Standalone**, select one or more hosts from the list. You can also select the **Show hosts in existing host pools not managed by RAS** option to show hosts that exist in other host pools on Azure.
 - If the host pool provisioning is configured as **Template**, you cannot manually add hosts here. Instead, use the [Hosts](#) tab in the main Azure Virtual Desktop view.
3. Click **OK**.

Specifications

The **Specifications** tab allows you to override settings defined at the template level.

To override the Organizational Unit in Active Directory where hosts are joined, select **Override Organizational Unit specified in template properties**. In the target OU field, provide the distinguished name of the Organizational Unit by using the browse (...) button or entering it manually.

To override the host size, select the **Overwrite the size specified in template properties** option and choose the desired size from the drop-down list. The selected size will only be used by this host pool. Other host pools using the same template will be unaffected. Note that available sizes may depend on the location, size, and power state of the host pool members and the template. Also note that overwriting the size requires a host reboot.

Note: High Availability settings cannot be changed after the host pool is created. To override High Availability settings at the host pool level, specify them during host pool creation.

Application packages

See **Using [MSIX application packages](#)**.

Assignment

The **Assignment** tab displays Active Directory users and groups assigned to Microsoft Entra ID objects. In order for users to see published desktops and applications, they must be assigned to the application group available in the host pool.

To create a new assignment:

1. Click **Tasks > Add**.
2. In the **Select User or Group** dialog, specify a user or group and click **OK**.
3. Follow the onscreen instructions and complete the assignment. Note that additional filtering in the **Publishing** category may be used to control Azure Virtual Desktop resource availability in Parallels Client. For details, see **Publish resources**.

User profile

By default, this tab inherits its settings from Site defaults. If you wish to specify custom settings, clear the **Inherit default settings** option. For the information about configuring user profile, please see [Site defaults \(Azure Virtual Desktop\)](#).

Optimization

The **Optimization** tab allows you to specify settings that will be used to optimize session hosts for best performance in a Parallels RAS environment. You can select Windows components, services, and other options that will be disabled, removed, or optimized to ensure a more efficient, streamlined, and improved delivery of virtual apps and desktops. By default, this tab inherits its settings from Site defaults. If you wish to specify custom settings, clear the **Inherit default settings** option. For the information about configuring optimization options, please see [Site defaults \(Azure Virtual Desktop\)](#).

Settings

This tab allows you to configure settings such as sessions timeouts, client URL/Mail redirection, drag and drop and others. By default, this tab inherits its settings from Site defaults. If you wish to specify custom settings, clear the **Inherit default**

settings option. For the information about configuring host pool settings, please see [Site defaults \(Azure Virtual Desktop\)](#).

RDP printer

The **RDP Printer** tab allows you to configure the renaming format of redirected printers. By default, this tab inherits its settings from Site defaults. If you wish to specify custom settings, clear the **Inherit default settings** option. For the information about configuring user profile, please see [Site defaults \(Azure Virtual Desktop\)](#).

Agent auto-upgrade

See [Upgrading Agents \(Azure Virtual Desktop\)](#).

AVD agent

The **AVD agent** tab allows you to schedule updates for Azure Virtual Desktop agents installed on the hosts in the host pool. If you do not schedule the updates, they will be installed automatically when Microsoft releases a new version.

For more information about Azure Virtual Desktop agents, see <https://learn.microsoft.com/en-us/azure/virtual-desktop/whats-new-agent>.

Note: Microsoft recommends always installing the latest stable version of the agent so that hosts do not lose connectivity or be subjected to security vulnerabilities.

To schedule AVD agent updates:

1. Go to the **Farm > Site > Azure Virtual Desktop > Host pools > Properties > AVD agent** tab.
2. Select the **Scheduled Microsoft AVD agent updates** option.
3. Select the time zone for the updates:
 - If you want to use the host's local time zone, select the **Use local session host time zone** option.
 - If you want to use a specific time zone, select it in the **Time zone** drop-down menu.
4. Specify the date and time for the maintenance window.

Note: If, after four attempts, the agent fails to update all hosts in the host pool, the update will happen the next time the hosts are turned on. `endhint %}`

1. (Optional) You might want to create two maintenance windows to ensure that updates are installed during the designated time. To do that, select the **Add an additional window** option and specify the date and time for the second maintenance window.

Upgrading Agents (Azure Virtual Desktop)

You can enable and configure automatic updates for all Azure Virtual Desktop hosts in a host pool.

Schedule Agent auto-upgrade

To schedule Agent auto-upgrade:

1. Go to the **Farm > Site > Azure Virtual Desktop > Host pools > Properties > Auto-upgrade** tab.
2. Clear the **Inherit default settings** options if you want to modify them for this host pool.
3. Select the **Enable auto-upgrade maintenance window** option. During the maintenance window, all hosts in the host pool will try to download Agent upgrades. The upgrades will be downloaded and installed as soon as all users log out of their hosts. New logons from users are prohibited (drain mode). If the users don't log off during a maintenance window, the upgrades won't be installed until the next window.
4. Specify the start date and time, duration, and recurrence settings for this event. To make this a one-time event, select **Never** in the **Recur** drop-down list.
5. (Optional) If you want to forcefully log off all users and download the upgrades at the end of a maintenance window, select the **Force logoff of current sessions at the end of the maintenance window duration** option.
6. (Optional) Configure a message that will be sent to users before or during a maintenance window. Click the **Configure messages** button and specify the message title, body, and the time period when it should be sent.

Cancel Agent auto-update

To cancel Agent auto-update:

1. Go to **Farm > Site > Azure Virtual Desktop > Host pools**.
2. Select **Tasks > Cancel auto-upgrade maintenance window**.

Manage templates (Azure Virtual Desktop)

An Azure Virtual Desktop template is a virtual machine from which other virtual machines are created as clones of the original VM and added to a host pool as session hosts.

To manage Azure Virtual Desktop templates, navigate to **Farm > Site > Azure Virtual Desktop** and select the **Templates** tab.

Create a template

To create a template:

1. On the **Templates** tab, click **Tasks > Add**. This opens the **Create Parallels Template Wizard**.
2. On the first page, select an Azure Virtual Desktop provider (if you have more than one).
3. Select a template type from the following:
 - **Multi-session**: Multiple concurrent user sessions are allowed on a single host running a Windows Server operating system, Windows 10 or Windows 11 Enterprise multi-session.
 - **Single-session**: A single user session is allowed on a single session host.
4. On the **Template Source** page, select a source from the following:
 - **Custom host**: Displays a list of existing virtual machines.
 - **Azure Gallery**: Allows you to select an image and create a new virtual machine from it. Depending on the template type, multi-session or single-session, commonly used marketplace images such as Windows 10 Enterprise multi-session, are predefined to be easily chosen and created as a template. Select a location and specify the local

administrator username and password. The **Browse all images** button opens a dialog where you can choose any other image from the Marketplace or Shared Image Gallery. When choosing an image from the Shared Image Gallery, select from a list of publishers, SKUs, offers, and other options.

- Specify the local administrator username and password

5. On the **General** page, specify the following settings:

- **Template name:** The name of the template.
- **Description:** The description of the template.
- **Security type:** Select other Trusted Launch Virtual Machines or Standard. When Trusted Launch Virtual Machines is selected, additional security features can be enabled as needed.

6. On the **Hosts** page, select virtual machine properties from the predefined Azure values according to your needs:

- Specify an Azure resource group.
- Select an availability option:
 - **No infrastructure redundancy required:** No redundancy in Azure will be configured.
 - **Availability set:** If selected, hosts will be deployed from the template in an availability set. Note that the maximum number of hosts that can be deployed in an availability set is 200 (this is an Azure limitation).
 - **Availability zone:** Hosts will be deployed in a specific availability zone that is selected.
 - **Note:** The availability option can be overridden in the settings of the host pool using a given template. This gives you the ability to specify a different availability option at the host pool level.

- Select a virtual machine size.

Note: The virtual machine size can be overridden in the settings of the host pool using a given template. This gives you the ability to specify a different VM size at the host pool level.

- Select a disk type.
- Select a virtual network and subnet.

Note: In case using Accelerated networking for the Template, make sure you select the appropriate host size for session hosts that support accelerated networking.

7. On the **Preparation** page, select an image preparation and specify the required options. This is similar to how an image is prepared for a RAS VDI template. There are some minor differences, but the configuration procedure is essentially the same. For details, please see [Preparation](#).
8. On the **Optimization** page, configure optimization settings. These settings are inherited from Site defaults but custom settings can be specified if needed. For details, please see [Site defaults \(Azure Virtual Desktop\)](#).
9. On the **New template version** page, specify the name and description and select the tags for the version. You can select several tags.
10. On the **Summary** page, review the settings and click **Finish** to create the template.

Manage existing templates

To modify an existing template, right-click it and choose **Properties**. Some properties cannot be modified, while many can. For the description of individual properties and settings, please refer to instructions in [Create a template](#).

To delete a template, select it in the list and click **Tasks > Delete**. Note that at the time of this writing, there's a known issue that if a template is deleted in the RAS Console, the template and associated hosts may not be completely removed from Microsoft Azure. To make sure that all such objects are removed, it is recommended to do it from the Azure portal.

Assigning a template to a host pool (Azure Virtual Desktop)

When you create a host pool and set its provisioning type as Template, you need to assign an existing template to it. This can be done when you create or modify a host pool, or you can assign a template to a host pool on the **Templates** tab.

To assign a template to a host pool:

1. Select a template and click **Tasks > Assign to host pool**.
2. Click **Tasks > Assign to pool**. A wizard opens.

3. On the **Versions** page, select the template version that will be assigned to the host pool.
4. (Optional) On the **Host pool** page, select the host pools that you want to recreate on schedule and click the **Configure** button. You will see a dialog that allows you to schedule recreation. Configure the schedule according to your needs and click **Next**.
5. Click **Finish**.

To remove a template from a host pool:

1. Select a template and click **Tasks > Remove from host pool**.
2. The **Remove from Host Pool** dialog opens listing all host pools using the selected template.
3. Select one or more host pools to remove the template from and click **OK**.
4. Note that if a host pool has hosts, they will be removed. You will see a message and need to confirm the removal.
5. If one or more host pools are being locked by another administrator, you will also see a message and will have to repeat the operation later when a pool is unlocked.

Manage hosts (Azure Virtual Desktop)

To manage Azure Virtual Desktop hosts, navigate to **Farm > Site > Azure Virtual Desktop** and select the **Hosts** tab.

The list displays hosts from all available host pools. You can apply a filter to the table to see hosts from a particular pool or using other criteria. To apply a filter, click the magnifying glass icon and specify the filter in a column (or columns) of interest.

Tasks that you can perform on a host are accessible from the **Tasks** menu and include the following:

- **Add**: Add a host to one of the available host pools. See the **Add a host** subsection below.
- **Assign**: This option is enabled for hosts from a Personal host pool. It allows you to assign the selected host to a user. If a host is already assigned to another user, you'll be asked if you want to change the assignment. Select an Microsoft Entra ID user when asked. The assignment is done in Azure, so the host status will change to "Assigning" for the duration of the operation.

- **Unassign:** Removes the user assignment from the selected host, see **Assign** above. This menu option is enabled for hosts that are currently assigned to a user. The status of the host changes to “Unassigning” for the duration of the operation.
- **Search:** Allows you to search for a host in the list by applying a filter.
- **Show sessions:** Switches to the **Sessions** tab with a filter applied to show the selected host sessions.
- **Show published resources:** Displays a list of resources published from the selected host.
- **Show application packages:** Displays MSIX application packages added to the selected host.
- **Control:** Control options, including enable or disable logons on the selected host, cancel a pending reboot (originated by scheduler), cancel a disabled state (originated by scheduler). See [Using scheduler](#) for details.
- **Start, Stop, Reset, Restart:** Power operations that can be performed on the selected host. The Restart operation (graceful) has a 10-minute timeout. If not completed during this time, the Reset operation (forced) will be used. If you choose **Stop**, the host will be stopped and deallocated.
- **Upgrade all Agents:** Upgrades agents on every host in the list (if necessary).
- **Stop optimization:** When an optimization is applied to a host, it can be canceled in the beginning stages. For more information, see [Optimization](#).
- **Tools:** Standard RAS tools, including Remote Desktop, computer management, service management, event viewer, Powershell, and others. For the complete description, please see [Computer management tools](#).
- **Troubleshooting:** Allows you to check the agent status and update it if necessary. Also allows you to manage logging.
- **Details:** Shows details when a host was not created due to a failure. The option opens a dialog describing the reason of failure and some additional information.
- **Change license type:** Change the type of Azure license.
- **Recreate:** Recreates a host.
- **Delete:** Deletes a host from the list and from the host pool to which it belongs. The host (virtual machine) itself is kept or deleted depending on the host pool provisioning type. A host created from a template will be completely removed. A standalone host is not deleted, which means that the virtual machine stays intact.
- **Refresh:** Refreshes the list.

Add a host

You can add a host to a host pool from the **Hosts** tab. To do so:

1. Click **Tasks > Add**.
2. In the **Add Hosts** dialog, select a target host pool. Depending on the provisioning type configured for the selected host pool, do the following:
 - **Standalone**: Select one or more hosts from the list. You can also select the option at the bottom to show hosts in other existing host pools on Azure that are not managed by Parallels RAS.
 - **Template**: Specify the number of hosts to add to the pool from the template.
3. Click **OK**.

Manage sessions (Azure Virtual Desktop)

To view and manage Azure Virtual Desktop sessions, navigate to **Farm > Site > Azure Virtual Desktop** and select the **Sessions** tab. Sessions from all hosts in all host pools are displayed in the list.

For a detailed information about managing sessions, please see [Session Management](#).

Using scheduler (Azure Virtual Desktop)

The **Scheduler** tab allows you to create scheduler tasks that will be performed on individual hosts or host pools at a specified time.

Note: When the scheduled event is triggered, affected hosts are disabled in Parallels RAS and their status is displayed as “Disabled (scheduler)” or “Pending reboot (scheduler)”. You can cancel these states by right-clicking a host on the **Hosts** tab and choosing **Control > Cancel disabled state (scheduler)** or **Control > Cancel pending reboot (scheduler)**.

Disabling hosts and hosts in pools

To disable a host or a host in a pool:

1. Click **Tasks > Add > Disable host** or **Disable host pool**.
2. On the **General** tab, select the **Enable Schedule** option.
3. Specify a name for this schedule and an optional description.
4. Select a host or a pool in the **Available** list and click **Add**. The host or pool will appear in the **Target** list.
5. Select the **Trigger** tab and specify start date and time, duration, and recurrence settings for this event. To make this a one-time event, select **Never** in the **Recur** drop-down list.
6. Select the **Options** tab. It contains the following options:
 - **Message list**: Configure a message that will be sent to users before the host goes offline. Click **Tasks > Add** and specify the message title, body, and the time period when it should be sent.
 - **On disable**: Specify what should happen to current sessions when a scheduled task triggers. Select the desired option from the **On disable** drop-down list.
 - **Enforce schedule for currently inactive hosts**: This option is only enabled when you have an active message in the list. If the option is enabled, hosts that are currently offline are also monitored, and if such a server comes back online during the scheduled task execution, the task is applied to it too.
7. Click **OK** to save the schedule.

Rebooting hosts and hosts in pools

To reboot a host or a host in a pool:

1. Click **Tasks > Add > Reboot host** or **Reboot host pool**.
2. On the **General** tab, select the **Enable Schedule** option.
3. Specify a name for this schedule and an optional description.
4. Select a host or a pool in the **Available** list and click **Add**. The host or pool will appear in the **Target** list.
5. Select the **Trigger** tab and specify start date, time, and recurrence settings for this event. To make this a one-time event, select **Never** in the **Recur** drop-down list. In addition, specify the following options for the “Reboot host pool” task:
 - **Complete in**: Specify the time to complete the task.

6. Select the **Options** tab. It contains the following options:
 - **Message list**: Configure a message that will be sent to users before the host is rebooted. Click **Tasks > Add** and specify the message title, body, and the time period when it should be sent.
 - **Enable Drain Mode** and **Force server reboot after**: The two options work together. If you select the **Enable Drain Mode** option, then when the task triggers, new connections to a host are refused but active connections will continue to run and can be reconnected. The server will be rebooted when all active users close their sessions or when **Force server reboot after** time is reached, whichever comes first. For active users not to lose their work, create a message that will advise them to save their work and log off.
 - **Enforce schedule for currently inactive hosts**: This option is only enabled when you have an active message in the list. If the option is enabled, hosts that are currently offline are also monitored, and if such a server comes back online during the scheduled task execution, the task is applied to it too.
7. Click **OK** to save the schedule.

Starting up hosts and hosts in pools

Note: This task applies only to hosts and host pools based on a template.

To start up a host or a host in a pool:

1. Click **Tasks > Add > Startup host** or **Startup host pool**.
2. On the **General** tab, select the **Enable Schedule** option.
3. Specify a name for this schedule and an optional description.
4. Select a host or a pool in the **Available** list and click **Add**. The host or pool will appear in the **Target** list.
5. Select the **Trigger** tab and specify start date, time, and recurrence settings for this event. To make this a one-time event, select **Never** in the **Recur** drop-down list.
6. (“Startup host pool” task only) Select the **Options** tab. It contains the following options:
 - **Percentage of members**: Select this option to specify the percentage of hosts that must be started up in each pool.
 - **Specific number of members to be started**: Select this option to specify the number of hosts that must be started up in each pool.
7. Click **OK** to save the schedule.

Shutting down hosts and hosts in pools

To shut down a host or a host in a pool:

1. Click **Tasks > Add > Shutdown host** or **Shutdown host pool**.
2. On the **General** tab, select the **Enable Schedule** option.
3. Specify a name for this schedule and an optional description.
4. Select a host or a pool in the **Available** list and click **Add**. The host or pool will appear in the **Target** list.
5. Select the **Trigger** tab and specify start date, time, and recurrence settings for this event. To make this a one-time event, select **Never** in the **Recur** drop-down list.
6. Select the **Options** tab. It contains the following options:
 - **Message list**: Configure a message that will be sent to users before the host is rebooted. Click **Tasks > Add** and specify the message title, body, and the time period when it should be sent.
 - **Enable Drain Mode** and **Force server shutdown after**: The two options work together. If you select the **Enable Drain Mode** option, then when the task triggers, new connections to a host are refused but active connections will continue to run. The server will be shut down when all active users close their sessions or when **Force server reboot after** time is reached, whichever comes first. For active users not to lose their work, create a message that will advise them to save their work and log off.
 - **Enforce schedule for currently inactive hosts**: This option is only enabled when you have an active message in the list. If the option is enabled, hosts that are currently offline are also monitored, and if such a server comes back online during the scheduled task execution, the task is applied to it too.

Recreating hosts and host pools

Note: This task applies only to hosts and host pools based on a template.

To recreate a specific host or all hosts in a host pool:

1. Click **Tasks > Add > Recreate host from template** or **Recreate host pool from template**.
2. On the **General** tab, select the **Enable Schedule** option.
3. Specify a name for this schedule and an optional description.

4. Select a host or a pool in the **Available** list and click **Add**. The host or pool will appear in the **Target** list.
5. Select the **Trigger** tab and specify start date, time, and recurrence settings for this event. To make this a one-time event, select **Never** in the **Recur** drop-down list.
6. Select the **Options** tab. It contains the following options:
 - **Message list**: Configure a message that will be sent to users before the host is rebooted. Click **Tasks > Add** and specify the message title, body, and the time period when it should be sent.
 - **Enable Drain Mode**, **Force host recreation after**, and **Force host pool recreation after**: The options work together. If you select the **Enable Drain Mode** option, then when the task triggers, new connections to a host are refused but active connections will continue to run. The server will be recreated when all active users close their sessions or when the time specified in **Force host recreation after** or **Force host pool recreation after** is reached, whichever comes first. For active users not to lose their work, create a message that will advise them to save their work and log off.
 - **Enforce schedule for currently inactive hosts**: This option is only enabled when you have an active message in the list. If the option is enabled, hosts that are currently offline are also monitored, and if such a server comes back online during the scheduled task execution, the task is applied to it too.

Sending multiple messages to users

For **Disable Host** and **Disable Host Pool** tasks, you can only send a message before the scheduled task is triggered. Hence, when creating a message, you can only select the “before” option when specifying when the message should be sent. You can create more than one message if needed and send them at different time intervals, so the users are notified more than once before the task executes.

For **Reboot Host** and **Reboot Host pool** tasks, you can send a message before or after the scheduled task is triggered. The “after” option is available for these tasks because you have the ability to enable the drain mode, which will keep the active sessions running for some time. During this time, you can send multiple messages to active users reminding them that they should finish their work and close their sessions. To use the “after” option, the **Enable Drain Mode** option must be selected. Please also note that the “after” time interval and the **Force server reboot after** setting should be coordinated. For example, if the force reboot occurs

before the “after” time elapses, active users will not have a chance to see the message.

Site defaults (Azure Virtual Desktop)

When you configure Azure Virtual Desktop components and objects in the RAS Console, some of the properties are inherited from Site defaults. When you see the **Inherit default settings** option in a dialog or tab page, it means that settings can be either inherited from Site defaults or custom values can be specified for a given object.

To view and configure Site defaults for Azure Virtual Desktop, navigate to **Farm > Site**, click the **Tasks menu** and choose one of the following:

- **AVD multi-session hosts:** Opens a dialog to configure Azure Virtual Desktop [site defaults for multi-session hosts](#).
- **AVD single-session hosts:** Opens a dialog to configure Azure Virtual Desktop [Site defaults for single-session hosts](#).

Each dialog is described below.

Site defaults for single-session hosts

Host pool settings

Configure the following settings:

- **Disconnect active session after:** Specifies the amount of time each session remains connected in the background after the user has closed the published application. This option is used to avoid unnecessary reconnections with the server.
- **Logoff disconnected session after:** This setting allows you to control how long it takes for a session to be logged off after it is marked as “disconnected”.
- **Session readiness timeout:** The maximum amount of time it should require to establish a session. If the specified timeout is reached, and the session is still not ready, the user will see an error message and will have to try to log in again.

- **Allow URL/Mail redirection:** When a user tries to open a URL or an HTML Mailto link in a remote application, the link can be redirected to the client computer and open in a local default application (a web browser or email client) instead of an application on the remote host. This option allows you to enable or disable the redirection. Click the **Configure** button to choose from the following options:
 1. **Replace registered application** — this option uses an alternative method of redirecting a link. It replaces the default web browser and mail client with “dummy” apps on the remote server side. By doing so, it can intercept an attempt to open a link and redirect it to the client computer.
 2. **Support Windows Shell URL namespace objects** — the Shell URL namespace objects support means that Parallels RAS can intercept actions in published applications that use Shell namespace API to open links, which is a standard behavior in most applications. The ability to disable the support for Shell URL namespace objects is for compatibility with older versions of Parallels RAS.
- **Enable drag and drop:** Allows you to set how the drag and drop functionality works in Parallels Clients. Click Configure and choose from “Disabled” (no drag and drop functionality), “Sever to client only” (drag and drop to a local application only), “Client to server only” (drag and drop to a remote application only), “Bidirectional” (drag and drop in both directions).

Note: At the time of this writing, the drag and drop functionality is only supported on Parallels Client for Windows and Parallels Client for Mac.

- **Allow 2XRemoteExec to send command to the client:** Select this option to allow a process running on the server to instruct the client to deploy an application on the client side.
- **Manage RDP Shortpath:** Configures RDP Shortpath. RDP Shortpath establishes direct UDP-based connection between a remote desktop client and session host. Direct connection reduces the dependency on the Azure Virtual Desktop gateways, improves the connection reliability, and increases the bandwidth available for each user session. This applies to RDP and RAS connections. A session host requires a restart for this setting to become effective.

RDP Shortpath can be used in two ways:

- Managed networks, where direct connectivity is established between the client and the session host when using a private connection, such

as a virtual private network (VPN). To allow access to the RDP Shortpath listener across network security boundaries, Azure Network Security Group must be configured to allow inbound UDP port 3390. VPN or ExpressRoute is required or each session host should have public IP address.

- Public networks, where direct connectivity is established between the client and the session host when using a public connection. To allow access to the RDP Shortpath listener no inbound ports are required because outbound ports are being used.

When both RDP Shortpath for public networks and managed networks are enabled, a first-found algorithm kicks in, and whichever connection gets established first for that session is used. In most scenarios, when you have configured RDP Shortpath for managed networks you will want to have it take the precedence, and that is usually the case because building the session for RDP Shortpath for public networks takes a little longer.

Click the **Configure** button to enable and configure RDP Shortpath:

- **Use RDP Shortpath:** Enables RDP Shortpath.
- **Use a smaller default range of ports:** Limits the range of ports that remote desktop client can use to connect to a session host. The default range is 49152-65535. This only applies to RDP Shortpath for public networks.
- **Enable applications monitoring:** Enable or disable monitoring of applications on the server. Disabling application monitoring stops the WMI monitoring to reduce CPU usage on the server and network usage while transferring the information to RAS Connection Broker. If the option is enabled, the collected information will appear in a corresponding RAS report. If the option is disabled, the information from this server will be absent from a report.
- **Allow file transfer command (Web and Chrome clients):** Enables file transfer in a remote session. To enable file transfer, select this option and click the **Configure** button. For more information, see [Configuring remote file transfer](#).
- **Enable drive redirection cache:** Improves user experience by making file browsing and navigation on redirected drives much faster. For details, see [Drive redirection cache explanation](#).

- **Use RemoteApp if available:** Enable this option to allow use of remote apps for shell-related issues when an app is not displayed correctly. This feature is supported on the Parallels Client for Windows only.
- **Enable Z-order (Experimental):** Select this option if you encounter an issue related to z-order in applications published with Parallels RAS.

User profile

The **User profile** tab allows you to configure the user profile functionality. You can select from **Do not manage by RAS** (user profiles will not be managed) or **FSlogix**. Microsoft FSLogix Profile Container allows to maintain user context in non-persistent environments, minimize sign-in times and provides native profile experience eliminating compatibility issues. For complete instructions, please see [User profile](#).

Application packages

The **Application packages** tab allows you to add, remove, or otherwise manage MSIX application packages on single-session hosts in the Site. For the complete description, please see subsection “Adding a package to Site Defaults” in section [Using MSIX application packages](#).

Optimization

The **Optimization** tab allows you to specify settings that will be used to optimize a session host for best performance in a Parallels RAS environment. You can select Windows components, services, and other options that will be disabled, removed, or optimized to ensure a more efficient, streamlined, and improved delivery of virtual apps and desktops. For the complete description, please see [Optimization](#).

RDP printer

The **RDP Printer** tab allows you to configure the renaming format of redirected printers. The format may vary depending on which version and language of the server you are using.

Set your RDP Printer Name Format specifically for the configured server by choosing any of the below options from the RDP Printer Name Format drop-down list:

- **Prntername** (from Computername) in Session no
- **Session no.** (computername from) **Prntername**
- **Prntername** (redirected Session no)

The **Remove session number from printer name** does just that, it removes the session number from the name, so it's not visible.

Site defaults for multi-session hosts

Host pool settings

Configure the following settings:

- **Disconnect active session after:** Specifies the amount of time each session remains connected in the background after the user has closed the published application. This option is used to avoid unnecessary reconnections with the server.
- **Logoff disconnected session after:** This setting allows you to control how long it takes for a session to be logged off after it is marked as “disconnected”.
- **Session readiness timeout:** The maximum amount of time it should require to establish a session. If the specified timeout is reached, and the session is still not ready, the user will see an error message and will have to try to log in again.
- **Allow URL/Mail redirection:** When a user tries to open a URL or an HTML Mailto link in a remote application, the link can be redirected to the client computer and open in a local default application (a web browser or email client) instead of an application on the remote host. This option allows you to enable or disable the redirection. Click the **Configure** button to choose from the following options:
 1. **Replace registered application** — this option uses an alternative method of redirecting a link. It replaces the default web browser and mail client with “dummy” apps on the remote server side. By doing so, it can intercept an attempt to open a link and redirect it to the client computer.

2. Support Windows Shell URL namespace objects — the Shell URL namespace objects support means that Parallels RAS can intercept actions in published applications that use Shell namespace API to open links, which is a standard behavior in most applications. The ability to disable the support for Shell URL namespace objects is for compatibility with older versions of Parallels RAS.

- **Enable drag and drop:** Allows you to set how the drag and drop functionality works in Parallels Clients. Click Configure and choose from “Disabled” (no drag and drop functionality), “Sever to client only” (drag and drop to a local application only), “Client to server only” (drag and drop to a remote application only), “Bidirectional” (drag and drop in both directions).

Note: At the time of this writing, the drag and drop functionality is only supported on Parallels Client for Windows and Parallels Client for Mac.

- **Allow 2XRemoteExec to send command to the client:** Select this option to allow a process running on the server to instruct the client to deploy an application on the client side.
- **Manage RDP Shortpath:** Configures RDP Shortpath. RDP Shortpath establishes direct UDP-based connection between a remote desktop client and session host. Direct connection reduces the dependency on the Azure Virtual Desktop gateways, improves the connection reliability, and increases the bandwidth available for each user session. This applies to RDP and RAS connections. A session host requires a restart for this setting to become effective. To allow access to the RDP Shortpath listener across network security boundaries, Azure Network Security Group must be configured to allow inbound UDP port 3390. VPN or ExpressRoute is required or each session host should have public IP address.

Click the **Configure** button to enable and configure RDP Shortpath:

- **Use RDP Shortpath:** Enables RDP Shortpath.
- **Use a smaller default range of ports:** Limits the range of ports that remote desktop client can use to connect to a session host. The default range is 49152-65535.
- **Enable applications monitoring:** Enable or disable monitoring of applications on the server. Disabling application monitoring stops the WMI monitoring to reduce CPU usage on the server and network usage while transferring the information to RAS Connection Broker. If the option is enabled, the collected information will appear in a corresponding RAS report.

If the option is disabled, the information from this server will be absent from a report.

- **Allow file transfer command (Web and Chrome clients):** Enables file transfer in a remote session. To enable file transfer, select this option and click the **Configure** button. For more information, see [Configuring remote file transfer](#).
- **Enable drive redirection cache:** Improves user experience by making file browsing and navigation on redirected drives much faster. For details, see [Drive redirection cache explanation](#).
- **Use RemoteApp if available:** Enable this option to allow use of remote apps for shell-related issues when an app is not displayed correctly. This feature is supported on the Parallels Client for Windows only.
- **Enable Z-order (Experimental):** Select this option if you encounter an issue related to z-order in applications published with Parallels RAS.

User profile

The **User profile** tab allows you to configure the user profile functionality. You can select from **Do not manage by RAS** (user profiles will not be managed) or **FSlogix**. Microsoft FSLogix Profile Container allows to maintain user context in non-persistent environments, minimize sign-in times and provides native profile experience eliminating compatibility issues. For complete instructions, please see [User profile](#).

Application packages

The **Application packages** tab allows you to add, remove, or otherwise manage MSIX application packages on multi-session hosts in the Site. For the complete description, please see subsection “Adding a package to Site Defaults” in section [Using MSIX application packages](#).

Optimization

The **Optimization** tab allows you to specify settings that will be used to optimize a session host for best performance in a Parallels RAS environment. You can select Windows components, services, and other options that will be disabled, removed,

or optimized to ensure a more efficient, streamlined, and improved delivery of virtual apps and desktops. For the complete description, please see [Optimization](#).

RDP printer

The **RDP Printer** tab allows you to configure the renaming format of redirected printers. The format may vary depending on which version and language of the server you are using.

Set your RDP Printer Name Format specifically for the configured server by choosing any of the below options from the RDP Printer Name Format drop-down list:

- Printername (from Computername) in Session no
- Session no. (computername from) Printername
- Printername (redirected Session no)

The **Remove session number from printer name** does just that, it removes the session number from the name, so it's not visible.

Using Parallels Client with Azure Virtual Desktop

After you deployed Azure Virtual Desktop in Parallels RAS and published resources from it, you can use Parallels Client to access applications and/or desktops that you published. This topic describes Parallels Client requirements and provides additional information about accessing published resources.

Requirements

Parallels Client requirements for opening Azure Virtual Desktop apps and desktops are as follows:

- Parallels Client for Windows:
 - Parallels Client for Windows version 18 or newer (Basic or Full versions).
 - Supported Windows versions: Windows 10 and Windows 11. Note that Windows Server operating systems are not supported.
 - Windows updates: Update for Universal C Runtime for Windows (KB2999226). Microsoft Windows 10 incorporates this by default.

- Microsoft .NET Framework 4.6.0 or later is required. Microsoft Windows 10 incorporates .NET Framework 4 and has it enabled by default.
- In case the Parallels RAS policy is configured to use the Azure Virtual Desktop Client, the Microsoft Windows Desktop client, also known as Remote Desktop (MSRDC) client, must be installed. The client is downloaded and installed automatically upon launching an Azure Virtual Desktop resource from Parallels Client (if not already installed on a supported Windows client device). You may also download the client using the following link: <https://go.microsoft.com/fwlink/?linkid=2068602>.
- In case the Parallels RAS policy is configured to use Microsoft Windows App, the Windows App is automatically downloaded and extracted upon launching an Azure Virtual Desktop resource from Parallels Client. A manually installed version of Microsoft Windows App will not be used by the Parallels RAS Client for Windows.

Note: If you are using the [Standard](#) client feature set option and Windows 10 and Windows 11 Enterprise Virtual Desktop as desktop OS where the Parallels Client is running (nested), the administrator needs to have the Windows Desktop client preinstalled using the per-device installation as highlighted by Microsoft in the following article: <https://docs.microsoft.com/en-us/windows-server/remote/remote-desktop-services/clients/windowsdesktop-admin>.

Note: Starting March 27, 2026, the Remote Desktop client for Windows (MSI) will no longer be supported by Microsoft. Use Parallels RAS policies to configure using the Microsoft Windows App for launching Azure Virtual Desktop resources.

- Microsoft Teams support: Starting with Parallels RAS 19, you can redirect audio and video from an Azure Virtual Desktop host running Windows 10 or 11 Multi-session, Windows 10, or 11 Enterprise to Parallels Client for Windows. For information on how to install Microsoft Teams on a Azure Virtual Desktop host, see <https://docs.microsoft.com/en-us/azure/virtual-desktop/teams-on-avd>.

Note: Video feed is not available if Microsoft Teams is accessed as a published application using the Advanced client feature set.

Note: At the time of this writing, it is not yet possible to use other platform-specific Parallels Clients to launch Azure Virtual Desktop resources. However, you can use Parallels Web Client to do that on certain operating systems (see below).

- Parallels Web Client:
 - Windows, Mac, or Linux operating system installed on the user's machine.
 - System requirements as listed in [Parallels Web Client and User Portal](#).

Accessing Azure Virtual Desktop resources in Parallels Client for Windows

When you connect to Parallels RAS from Parallels Client for Windows, all published resources, including Azure Virtual Desktop resources, are listed and made available for a user to access. Note that Azure Virtual Desktop resources are only shown in Parallels Client running on supported versions of Windows (see above).

The **Client feature set** options **Advanced** and **Advanced with feedback** are deprecated and should no longer be used. Use **Standard** mode for all new and existing deployments to ensure compatibility and long-term support. Standard mode contains most features that were previously only available in Advanced mode, such as RAS Universal Printing and Scanning and accelerated file redirection. Standard mode currently does not support session prelaunch and drag and drop.

To view and modify this setting, navigate to **Farm > Site > Settings**, select the **Features** tab, and select a desired setting in the **Client feature set** drop-down list.

Parallels RAS policies can be configured to use the **Microsoft Windows App** for launching Azure Virtual Desktop resources. Using the Microsoft Windows App requires the Client feature set to be set to Standard.

Note: Starting March 27, 2026, the Remote Desktop client for Windows (MSI) will no longer be supported by Microsoft. Use Parallels RAS policies to configure using the Microsoft Windows App for launching Azure Virtual Desktop resources.

Verify the deployment

To verify the Azure Virtual Desktop deployment, do the following:

1. Navigate to **Farm > Site > Settings** and select the **Features** tab. Verify that the **Enable Azure Virtual Desktop management** option is selected and the **Status** section says **Available** and displays the version number.
2. Navigate to **Farm > Site > Azure Virtual Desktop**. Select the following tabs and verify that corresponding components are configured and functioning properly:
 - **Providers**
 - **Workspaces**
 - **Host pools**
 - **Templates** (if you created a template, it should be listed on this tab)
 - **Hosts** (should list one or more session hosts)

Remote PCs

Overview

In addition to RD Session Hosts, Virtual desktops, and Azure Virtual Desktop, resources can also be published from a standalone Remote PC running a [supported version of Windows](#). A Remote PC can be a physical box or a virtual machine treated as a physical PC, but typically they are physical computers. If you have virtual machines on your network, it makes sense to use them as part of the VDI infrastructure as described in the [VDI and Virtual Desktops](#) chapter.

However, if you don't need the guest VM cloning functionality or, for example, if your end users require full administrative permissions on a PC for customization, you can use a virtual machine as a Remote PC.

Note: Remote PCs can also be combined into pools and managed as pool members. Remote PC pools use the RAS VDI infrastructure and work differently than individual Remote PCs described in this chapter. For more information see [Remote PC pools](#).

This chapter describes how to add a Remote PC to a Farm and how to publish remote applications and desktop from it.

Manage host pools (Remote PC)

See [Remote PC pools in VDI](#).

Manage hosts (Remote PC)

Adding a Remote PC to a Farm

Remote PCs can be added to a RAS Farm using one of the following methods:

- [Admin-initiated enrollment](#) from the RAS Console by specifying the PC's IP and MAC addresses and installing (remotely or directly) RAS Remote PC Agent on it.
- By allowing users to use [self-service Remote PC enrollment](#).

Read the subsequent sections for the description of each method.

Admin-initiated Remote PC enrollment

Requirements

In order to add a Remote PC to a RAS Farm, RAS Remote PC Agent must be installed on it. The requirements to push install RAS Remote PC Agent are as follows:

- The firewall must be configured on the server to allow push installation. Standard SMB ports (139 and 445) need to be open. See also **Port reference** for the list of ports used by Parallels RAS.
- SMB access. The administrative share (\\$) must be accessible. Simple file sharing must be enabled.
- Your Parallels RAS administrator account must have permissions to perform a remote installation on the PC. If it doesn't, you'll be asked to enter credentials of an account that does.
- The PC should be joined to an AD domain. If it's not, the push installation may not work and you will have to install the Agent on it manually. Please see **Installing Remote PC Agent manually**.

Adding a Remote PC to a Farm

Follow the below procedure to add a Remote PC to a Farm:

1. In the RAS Console, select the **Farm** category and click the **Remote PCs** node in the navigational tree.
2. Click **Add** in the **Tasks** drop-down list to launch the setup wizard.
3. Specify the IP address or FQDN of a Remote PC. Click the **Get MAC** button to obtain the PC's MAC address. To automatically resolve IP address to FQDN, enable the global **Name Resolution** option. For details, see [Host name resolution](#).
4. Click **Next**.
5. In this step, the Parallels RAS checks if the Remote PC Agent is installed on the specified PC. If it's not installed, click **Install** to push install the agent on the PC. If the push installation of Remote PC Agent fails for any reason, you can install it manually. See **Installing Remote PC Agent Manually** below.
6. Click **Add** to add the Remote PC to the Parallels RAS Farm.

Installing Remote PC Agent manually

You may need to install the Remote PC Agent manually if the automatic push installation cannot be performed for any reason. To do so:

1. Log into the PC where the Remote PC Agent is to be installed using an administrator account and close all other applications.
2. Copy the Parallels RAS installation file (`RASInstaller.msi`) to the PC and double click it to launch the installation.
3. Follow the onscreen instructions and proceed to the installation type page. Select **Custom** and click **Next**.
4. Click on the **Remote PC Agent** and select **Entire Feature will be installed on local hard drive** from the drop-down list.
5. Ensure that all other components are deselected and click **Next**.
6. Click **Install** to start the installation. Click **Finish** once the installation is finished.

Remote PC Agent does not require any configuration. Once the agent is installed, select the Remote PC name in the Parallels RAS Console and click **Troubleshooting > Check Agent**. If the agent is installed properly, the status should change to **Agent Installed**.

Uninstalling Remote PC Agent

To uninstall Remote PC Agent from a server:

1. Navigate to **Start > Control Panel > Programs > Uninstall a Program**.
2. Find **Parallels Remote Application Server** in the list of installed programs.
3. If you don't have any other Parallels RAS components on the server that you want to keep, right-click **Parallels Remote Application Server** and then click **Uninstall**. Follow the instructions to uninstall the program. You may skip the rest of these instructions.
4. If you have other RAS components that you want to keep on the server, right-click **Parallels Remote Application Server** and then click **Change**.
5. Click **Next** on the Welcome page.
6. On the **Change, repair, or remove** page, select **Change**.
7. On the next page, select **Custom**.
8. Select **Remote PC Agent**, then click the drop-down list in front of it, and click **Entire feature will be unavailable**.
9. Click **Next** and complete the wizard.

Self-service Remote PC enrollment

As an alternative to the admin-initiated Remote PC enrollment described above, end users can add their preferred PCs to a RAS Farm in a self-service manner. A computer enrolled this way is automatically added to a RAS Farm as a Remote PC and a corresponding published desktop is created with access granted (through automatic filtering) to the user who performed the self-enrollment. The Remote PC can then be accessed by the user from any device, anywhere.

Note: This feature applies to standalone Remote PCs. It is not intended for Remote PCs configured through VDI technology.

Requirements

- Users enrolling their PCs must have local administrator privileges in Windows to install Parallels RAS Remote PC Agent.
- Mailbox must be configured in a RAS Farm to send an invitation email to end users.

Configure self-service enrollment

To configure self-service Remote PC enrollment:

1. In the RAS Console, navigate to **Farm > Site > Remote PCs**.
2. In the right pane, click the **Tasks** menu and choose **Self-service enrollment**.
In the dialog that opens, specify the options described below.
3. To enable self-service enrollment, select **Allow** or **Allow until**. For the latter, specify the date and time. To disable the functionality (e.g. temporarily), select **Do not allow**.
4. In the **Settings** section, specify a publishing folder in which the PCs will appear as published resources. You can select an existing folder or create a new one. Click the [...] button to select or create a folder.
5. The **Remote PC invitation hash** field contains a hash that must be specified when enrolling a PC. The hash can also be copied from here and used separately for scripting purposes. IT administrators can make use of this hash along with the `msiexec` command to silently install and configure Remote PC on users' behalf. For details, see **Enrolling a PC** below.
6. To send an invitation email to users that will contain instructions on how to enroll their PCs, click the **Send via email** button.
7. In the dialog that opens, specify the recipients by typing (or pasting) their email addresses. You can also click the [...] button and select the recipients.
8. In the **Review the invitation email** text box, review or modify (if needed) the email. The variables used in the email are set internally and are substituted with their values in the actual email. To preview the final email text, click **Tasks > Preview**.
9. Click **Send** to send the email. If you don't want to send the email at this time, click **Cancel** to return to the previous dialog where you can click **OK** to save the changes.

Enrolling a PC

When a user receives the invitation email, they will follow the instructions that it contains to enroll their PC. The installation process consists of the steps described below.

Log in to the Remote PC. Download or copy the `RASInstaller.msi` file (the Parallels RAS installer) and run the following command with administrative privileges (the invitation email will contain this command with the hash value already in it):


```
msiexec /qb /i ADDLOCAL=F_PCAgent ADDFWRULES=1 SELFENROLL=[  
[OVERRIDEUSER=user@domain] [OVERRIDEPAIP=ip of PA]  
[OVERRIDEHOST=published name]
```

The following arguments can be used to customize the enrollment. Such arguments are required if a Remote PC is not in the Active Directory domain:

- **OVERRIDEPAIP:** The IP address of one of the Connection Brokers in the Farm Site. Use this if the standard installation fails to connect using the system detected IP address.
- **OVERRIDEUSER:** Use this argument if you don't want to register the Remote PC to the user logged in to the machine.
- **OVERRIDEHOST:** Use this argument if you want to change the published item name from the hostname of the Remote PC.

Once the installation is complete, launch Parallels Client and log in using the local machine credentials or the ones specified in the **OVERRIDEUSER** argument. Look for your Remote PC using the IP address or the name specified via **OVERRIDEHOST** in the list of published resources and launch the desktop.

Configuring a Remote PC

To view the properties of a Remote PC, highlight the computer in the navigation tree and click **Tasks > Properties**. This opens the Remote PC properties dialog.

Properties

By default, a PC is enabled in the Farm. When it is disabled, published applications and virtual desktops cannot be served from it. To enable or disable a PC in the Farm, select or clear the **Enable Remote PC** option.

If the IP or MAC address of a Remote PC has changed, modify them using the **Remote PC** and **MAC Address** input fields.

The **Change Direct Address** option allows you to specify an IP address that Parallels Client can use to connect to the PC directly. This address is only used in the Direct Connection mode and it could be an internal or external IP address.

Note: The Wake On Lan option should be enabled in BIOS so the machine could be automatically turned on. If you are using a virtual machine, the option is usually supported by a hypervisor natively or via a 3rd party software. To test if the Wake

On Lan option is turned on, close the **Remote PC Properties** dialog and then click the **Test Wake on LAN** button, which is located below the **Remote PCs** list.

Agent Settings

Each Remote PC in the Farm has a RAS Remote PC Agent installed to conduct communications between Parallels RAS and the PC. The agent can be configured on the **Agent Settings** tab page.

- **Logoff active session after:** The amount of time a session remains logged in after the user closes a published application. The default timeout is 25 seconds. Note that this only works for applications, but not published desktops (when a user closes a desktop, the session is logged off). This timeout is used to avoid unnecessary logins when a user closes one application and then opens another.
- **Session readiness timeout:** The maximum amount of time it should require to establish a session. If the specified timeout is reached, and the session is still not ready, the user will see an error message and will have to try to log in again.
- **Port.** Specify a different remote desktop connection port number if needed.
- **Preferred Connection Broker:** Select a Connection Broker with which the Remote PC Agent should communicate. This can be helpful when Site components are installed in multiple physical locations communicating through WAN. You can decrease network traffic by specifying a more appropriate Connection Broker.
- **Enable drag and drop:** Allows you to set how the drag and drop functionality works in Parallels Clients. Click **Configure** and choose from “Disabled” (no drag and drop functionality), “Sever to client only” (drag and drop to a local application only), “Client to server only” (drag and drop to a remote application only), “Bidirectional” (drag and drop in both directions).

Note: At the time of this writing, the drag and drop functionality is only supported on Parallels Client for Windows and Parallels Client for Mac.

- **Manage RDP transport protocol:** Selects the transport protocol that will be used for connections between Parallels Client and a server. To do this, select this option and click the **Configure** button.
- **Allow file transfer command (Web and Chrome clients):** Enables file transfer in a remote session. To enable file transfer, select this option and click the **Configure** button. For more information, see [Configuring remote file transfer](#).

- **Enable drive redirection cache:** Improves user experience by making file browsing and navigation on redirected drives much faster. For details, see [Drive redirection cache](#).

RDP Printer

The **RDP Printer** tab allows you to configure the renaming format of redirected printers. The format may vary depending on which version and language of the server you are using.

Set your RDP Printer Name Format specifically for the configured server by choosing any of the below options from the RDP Printer Name Format drop-down list:

- Printername (from Computername) in Session no.
- Session no. (computername from) Printername
- Printername (redirected Session no)

The other RDP Printing options available in the RDP Printer tab are:

- Remove session number from printer name
- Remove client name from printer name

Configure logging

A Remote PC is monitored and logs are created containing relevant information. To configure logging and retrieve or clear existing log files, right-click a Remote PC, choose **Troubleshooting > Logging** in the context menu, and then click **Configure**, **Retrieve**, or **Clear** depending on what you want to do. For the information on how to perform these tasks, see the [Logging](#) section.

Viewing Remote PC summary

In addition to the Remote PCs editor described in this chapter, you can also see the summary about the available Remote PCs. To do so:

1. In the RAS Console, select the **Farm** category and then select the **Site** node in the middle pane.
2. The available servers are displayed in the **Remote PCs** group in the right pane.

3. To go to the Remote PCs editor, right-click a server and choose **Show in the Editor**.

For additional info, see [Sites in the RAS Console](#).

Using computer management tools

You can perform standard computer management tasks on a server right from the RAS Console. These include Remote Desktop Connection, PowerShell, Computer Management, Service Management, Event Viewer, IPconfig, Reboot, and others. To access the **Tools** menu, select a server, click **Tasks > Tools** and choose a desired tool. For requirements and usage information, see [Computer management tools](#).

Publishing

Overview

Resources that can be published in Parallels RAS include:

- Installed applications
- Containerized applications
- Packaged applications
- Desktops
- Documents
- Web applications
- Network folders

This chapter describes how to publish resources hosted on servers managed by Parallels RAS and management tasks that you can perform on resources that have been already published. The publishing functionality described here is accessed from the **Publishing** category in the RAS Console.

Publishing a desktop

To publish a remote desktop:

1. In the RAS Console, select the **Publishing** category and click the **Add** icon below the **Published Resources** tree. This will launch the publishing wizard.

Note: If the wizard has all options disabled, it means that there are no resources (servers) in the Farm from which publishing can be configured.

1. In the first step of the wizard, select **Desktop** and click **Next**.
2. In the **Desktop Type** step, select the server type from which to publish and click **Next**.

Note: If you want to publish from a pool-based Remote PC, select the **Virtual Guest** option. The **Remote PC** option there is for standalone Remote PCs.

1. Select one or more servers which desktops you want to publish. You can select all available hosts, host pool(s), or individual hosts. Please note that if you have just one available server, this page will not be displayed.
2. Click **Next**.
3. Enter a desktop name, an optional description, and change the icon if needed.
4. (RD Session Host only) Configure the following settings:
 - Select the **Connect to administrative session** option if you want users to connect to the administrative session.
 - Select **Exclude from session prelaunch** if needed. For details, see [Understanding session prelaunch](#).
 - Select the **Start automatically when user logs on** option if you want to open a desktop as soon as a user logs on.
5. (VDI only) Select the **Enable static assignment to guest VM** option to mark a guest VM as persistent the first time a user connects to it.
6. (Remote PC only) Click the [...] button in the **Remote PC Settings** section to select a Remote PC from which the application should be published. In the box that opens, double-click a PC to select it.
7. Specify the desired screen resolution using the **Desktop Size** drop-down list. To set a custom width and height of the screen, select **Custom** in the **Size** drop-down list and specify the desired values in the fields provided.
8. In the **Multi-Monitor** drop-down list, select whether the multi-monitor support should be enabled, disabled, or whether the client settings should be used.

9. On the next page, specify the initial status of the resource. Choose from **Enabled** (end users can launch the resource), **Disabled** (the resource will not appear in Parallels Client), **In maintenance** (the resource will appear in Parallels Client but users will not be able to launch it). When a resource is in maintenance and a user tries to launch it, they will see a message. To customize the message, click the **Configure** button. For more info, see [Site defaults \(Publishing\)](#).
10. When done, click **Finish** to publish the desktop.

Publishing an application

To publish an application, follow the below procedure:

1. In the RAS Console, select the **Publishing** category and then click the **Add** icon below the **Published Resources** tree (or right-click inside the **Published Resources** box and click **Add** in the context menu). This will launch the publishing wizard.

Note: If the wizard has all options disabled, it means that there are no resources (servers) in the Farm from which publishing can be configured.

1. On the **Select Item Type** wizard page, select **Application** and click **Next**.
2. On the **Select Server Type** page, select the server type from which to publish and click **Next**.

Note: If you want to publish from a pool-based Remote PC, select the **Virtual Guest** option. The **Remote PC** option there is for standalone Remote PCs.

1. On the **Select Application Type** page, select one of the following available options:
 - **Single Application.** Choose this option to fully configure the application settings yourself such as the executable path etc.
 - **Installed Application.** Choose this option to publish an application that is already installed on the server, therefore all of the application settings are automatically configured.
 - **Predefined Application.** Choose this option to publish a commonly used Windows application such as Windows Explorer.
 - **Application Packages.** Choose this option to publish an application from an MSIX application package. The process of publishing applications from MSIX

application packages is described in [Publishing an application with MSIX app attach](#).

1. Click **Next**.
2. On the **Publish From** page, specify from which host the application should be published. You have the following options:
 - (RDSH) **All Hosts in Site**. If selected, the application will be published from all servers that are available on the Site.
 - (RDSH) **RD Sessions Host Host Pools**. Select this option and then select individual host pools to publish the application from.
 - (RDSH) **Individual Hosts**. Select this option and select individual servers to publish the application from.
 - (VDI) **Host pools**. Select the host pool to publish the application from.
 - (AVD) **Host pools**. Select the host pool to publish the application from.

The page will be skipped if the application type that you are installing is **Predefined Application**.

3. Click **Next**.
4. Depending on the application type that you selected on the **Select Application Type** page, the next wizard page will be one of the following:
 - If you selected **Single Application**, the **Application** page will open where you have to specify the application settings manually (more about this option later in this section).
 - If you selected **Installed Applications**, the **Installed Applications** page will open listing available applications (the applications are grouped by functionality). Select an application you wish to install and click **Next**. Follow the instructions to complete the wizard.
 - If you selected **Predefined Application**, the **Select Predefined Applications** page will open listing available applications. Select an application you wish to publish and click **Finish**.
5. If you selected **Single Application** on the **Select Application Type** wizard page, the **Application** page will open. Specify the application settings as follows:

Note that if you populate the **Target** field first using the “browse” button ([...]), the application **Name**, **Description**, and icon will be chosen automatically. You can override this selection if you wish.

- **Name.** Choose and type a name for the application.
- **Description.** Type an optional description.
- **Run.** Select the application window state (normal window, minimized, maximized).
- **Exclude from session prelaunch.** For details, see see [Understanding session prelaunch](#).
- **Start automatically when user logs on.** Select this option if you want to start an application as soon as a user logs on. This option works on desktop versions of Parallels Client only.
- **Change Icon.** Change the application icon (optional).
- **Server(s).** Allows you to specify the rest of the server parameters individually for each server the application was published from. Select a server from the drop-down list box and specify the parameters. Repeat for other servers in the list.
- **Target.** Specify the application executable path and file name.
- **Start in.** If the **Target** field is valid, this field will be populated automatically. You can specify your own path if needed.
- **Parameters.** If the application accepts startup parameters, you can specify them in this field.

6. (VDI only) Select the **Enable static assignment to guest VM** option to mark a guest VM as persistent the first time a user connects to it.
7. (Remote PC only) Click the [...] button in the **Remote PC Settings** section to select a Remote PC from which the application should be published. In the box that opens, double-click a PC to select it.
8. On the next page, specify the initial status of the resource. Choose from **Enabled** (end users can launch the resource), **Disabled** (the resource will not appear in Parallels Client), **In maintenance** (the resource will appear in Parallels Client but users will not be able to launch it). When a resource is in maintenance and a user tries to launch it, they will see a message. To customize the message, click the **Configure** button. For more info, see [Site defaults \(Publishing\)](#).
9. When done, click **Finish** to publish the application.

Publishing an App-V application

Microsoft Application Virtualization (or App-V) is an application streaming solution from Microsoft. Beginning with Parallels RAS v16.5, a support for App-V application publishing is available in the Parallels RAS console.

At the time of this writing, the App-V support implements scenarios where application provisioning is performed by means of App-V components:

- Applications are sequenced by the administrator according to Microsoft guidelines.
- Applications are stored on a network share created by the administrator (SMB, HTTPs).
- App-V Management and Publishing servers are used to publish applications for a specific AD groups that must be synced manually by the administrator with RAS publishing groups used for App-V application publishing.
- App-V client is installed and configured manually by the administrator.

The process of deploying and publishing an App-V application is as follows:

1. Package an applications using the App-V Sequencer.
2. Deploy the application to an RD Session Host using the App-V Management Console, Microsoft SCCM, etc.
3. Provision the application.
4. Verify that users can launch the application from the RD Session Host.
5. Publish the application from RAS Console (see below for instructions).
6. Launch the application from a Parallels Client.

To publish an App-V application:

1. In the Parallels RAS Console, select the **Publishing** category.
2. Click the **[+] Add** icon at the bottom of the right pane. The publishing wizard opens.
3. On the **Select Item Type** page, select the **App-V application** option.
4. Click **Next**.
5. Select the server type from which to publish an application and click **Next**.
6. Select a server or a group to publish from and click **Next**.
7. On the **Installed Applications** page, select one or more App-V applications and click **Next**.
8. Review the summary and complete the wizard.

Once an App-V application is published, it can be launched from a Parallels Client.

Note: To avoid launch issues, use AutoLoad=2. More details in <https://blogs.technet.microsoft.com/technetstosup/2013/11/12/autoload-setting-in-app-v-5-0/>.

Publishing local applications

To publish a local application, follow the below procedure:

1. In the RAS Console, select the **Publishing** category and then click the **Add** icon below the **Published Resources** tree (or right-click inside the **Published Resources** box and click **Add** in the context menu). This will launch the publishing wizard.

Note: If the wizard has all options disabled, it means that there are no resources (servers) in the Farm from which publishing can be configured.

1. On the **Select Item Type** wizard page, select **Application** and click **Next**.
2. On the **Select Server Type** page, select **Local Device** and click **Next**.
3. On the **Application** page, specify the application settings as follows:
 - **Name.** Choose and type a name for the application.
 - **Description.** Type an optional description.
 - **Start automatically when user logs on.** Select this option if you want to start an application as soon as a user logs on. This option works on desktop versions of Parallels Client only.
 - **Change Icon.** Change the application icon (optional).
 - **Target.** Specify the application executable path and file name.
 - **Start in.** If the **Target** field is valid, this field will be populated automatically. You can specify your own path if needed.
 - **Parameters.** If the application accepts startup parameters, you can specify them in this field.
 - **Run.** Select the application window state (normal window, minimized, maximized).
4. Click **Next**.
5. (Optional) On the **Alternate access by URL** page, specify the URL to the web version of the application. Parallels RAS will launch the web version if it cannot access the locally installed application.
6. Click **Next**.

7. On the next page, specify the initial status of the resource. Choose from **Enabled** (end users can launch the resource), **Disabled** (the resource will not appear in Parallels Client).
8. When done, click **Finish** to publish the application.

Publishing an application with MSIX app attach

To publish an application from an MSIX application package:

1. Add a package to a session host as described in [Using MSIX application packages](#).
2. Complete steps 1-2 as described in [Publishing an application](#).
3. On the **Select Server Type** page, select **RD Session Host, Virtual Guest**, or **Azure Virtual Desktop** and click **Next**.
4. On the **Select Application Type** page, select **Application Packages**.
5. Perform steps 5-7 as described in [Publishing an application](#).
6. The **Installed Applications** page will open listing available applications. Select an application you wish to publish and click **Next**.
7. Perform steps 10-11 as described in [Publishing an application](#).
8. The **Summary** page will open. It contains information about the application you selected for publishing. Click **Next**.
9. Click **Finish**.

Publishing a web application

A web application is like any other application that you can publish using the standard application publishing functionality. However, to simplify publishing of straight URL links to web applications, a separate publishing item type is available that allows you to accomplish this task with minimal number of steps.

To publish a web application:

1. In the RAS Console, select the **Publishing** category and then click the **Add** icon below the **Published Resources** tree (or right-click inside the **Published Resources** box and click **Add** in the context menu). This will launch the publishing wizard.

Note: If the wizard has all options disabled, it means that there are no resources (servers) in the Farm from which publishing can be configured.

1. On the **Select Item Type** wizard page, select **Web Application** and click **Next**.
2. On the **Server Type** page, select the server type from which to publish and click **Next**.

Note: If the wizard has all options disabled, it means that there are no resources (servers) in the Farm from which publishing can be configured.

Note: Resources published from the local device load faster than resources published from other servers. This is possible because Parallels RAS does not have to connect to the servers via RDP.

1. On the **Publish From** page, select the server(s) to publish from. Note that if you have just one server or selected **Local Device** on the previous step, the **Publish From** page will not appear.
2. On the **Web Application** wizard page that opens, specify the web application name, description, window state, and the URL. Select the **Force to use Internet Explorer** option if needed. To browse for a specific application icon, click **Change Icon**.
3. (Local device publishing only) You can check that the application can be accessed by clicking the **Test** button and start an application as soon as a user logs in by selecting the **Start automatically when user logs on** option.
4. (VDI only) Select the **Persistent** option to make a guest VM persistent. For more info, see Persistent guest VMs.
5. (Remote PC only) Click the [...] button in the **Remote PC Settings** section to select a Remote PC from which the application should be published. In the box that opens, double-click a PC to select it.
6. On the next page, specify the initial status of the resource. Choose from **Enabled** (end users can launch the resource), **Disabled** (the resource will not appear in Parallels Client), **In maintenance** (the resource will appear in Parallels Client but users will not be able to launch it). When a resource is in maintenance and a user tries to launch it, they will see a message. To customize the message, click the **Configure** button. For more info, see [Site defaults \(Publishing\)](#).
7. When done, click **Finish** to publish the application.

When published, the web application will appear in the **Publishing > Published Resources list**, just like any other application.

Publishing a network folder

You can publish a filesystem folder via UNC path to open in Windows explorer. To minimize the number of configuration steps, a special publishing item is available that allows you to publish a network folder from an RD Session Host.

To publish a network folder:

1. In the RAS Console, select the **Publishing** category and then click the **Add** icon below the **Published Resources** tree (or right-click inside the **Published Resources** box and click **Add** in the context menu). This will launch the publishing wizard.

Note: If the wizard has all options disabled, it means that there are no resources (servers) in the Farm from which publishing can be configured.

1. On the **Select Item Type** wizard page, select **Folder on the file system** and click **Next**.
2. On the **Select Server Type** page, select the server type from which to publish and click **Next**.

Note: If you want to publish from a pool-based Remote PC, select the **Virtual Guest** option. The **Remote PC** option there is for standalone Remote PCs.

Note: Resources published from the local device load faster than resources published from other servers. This is possible because Parallels RAS does not have to connect to the servers via RDP.

1. On the **Publish From** page, select the server(s) to publish from. Note that if you have just one server or selected **Local Device** on the previous step, the **Publish From** page will not appear.
2. On the **UNC Folder** wizard page, specify the usual application properties.
3. In the **UNC path** field, enter the UNC path of the folder you wish to publish. Click the [...] button to browse for a folder (it may take some time for the **Browse for Folder** dialog to open).
4. (VDI only) Select the **Persistent** option to make a guest VM persistent. For more info, see Persistent guest VMs.
5. (Remote PC only) Click the [...] button in the **Remote PC Settings** section to select a Remote PC from which the application should be published. In the box that opens, double-click a PC to select it.

6. On the next page, specify the initial status of the resource. Choose from **Enabled** (end users can launch the resource), **Disabled** (the resource will not appear in Parallels Client), **In maintenance** (the resource will appear in Parallels Client but users will not be able to launch it). When a resource is in maintenance and a user tries to launch it, they will see a message. To customize the message, click the **Configure** button. For more info, see [Site defaults \(Publishing\)](#).
7. Click **Finish** to publish the folder and close the wizard.

When published, the network folder will appear in the **Publishing > Published Resources list**, just like any other application. If you select it and then click the **Application** tab, the application settings will be as follows:

- The **Target** property will always be set to `PublishedExplorer.exe`. This binary is created automatically (via agents pushing) and is simply a copy of the standard `explorer.exe` executable.
- The **Parameters** property specifies the network folder that we want to publish. The folder path can be in any format that the `explorer.exe` can handle.

Please note that although you have all standard application property tabs enabled for this publishing item, at least the following items should be ignored, as they are completely irrelevant:

- **Publish From**
- **File Extensions**

Publishing a document

To publish a document, follow the below procedure:

1. In the RAS Console, select the **Publishing** category and then click the **Add** icon below the **Published Resources** tree (or right-click inside the **Published Resources** box and click **Add** in the context menu). This will launch the publishing wizard.

Note: If the wizard has all options disabled, it means that there are no resources (servers) in the Farm from which publishing can be configured.

1. On the **Select Item Type** wizard page, select **Document** and click **Next**.
2. Select the server type from which to publish and click **Next**.

Note: If you want to publish from a pool-based Remote PC, select the **Virtual Guest** option. The **Remote PC** option there is for standalone Remote PCs.

Note: Resources published from the local device load faster than resources published from other servers. This is possible because Parallels RAS does not have to connect to the servers via RDP.

1. Specify the content type of the document you want to publish. You can select the content type from the predefined list or specify a custom content type in the **Custom content types** input field.
2. Click **Next** when ready.
3. On the **Publish From** page, select the server(s) to publish from. Note that if you have just one server or selected **Local Device** on the previous step, the **Publish From** page will not appear.
4. On the **Application** page, enter a name, an optional description, a Window state, and an icon if needed.
5. Use the [...] button next to the **Target** input field to browse for the document. All other fields will be automatically populated. To edit any of the auto populated fields, highlight them and enter the required details.
6. (Optional) In the **Parameters** input field, specify the parameters to pass to the application when it starts.

Note: Use the **Server(s)** drop-down list to specify different document settings for a specific server in case the document is configured differently on that particular server. The settings will be saved for each server you select individually.

1. (VDI only) Select the **Persistent** option to make a guest VM persistent. For more info, see Persistent guest VMs.
2. (Remote PC only) Click the [...] button in the **Remote PC Settings** section to select a Remote PC from which the application should be published. In the box that opens, double-click a PC to select it.
3. On the next page, specify the initial status of the resource. Choose from **Enabled** (end users can launch the resource), **Disabled** (the resource will not appear in Parallels Client), **In maintenance** (the resource will appear in Parallels Client but users will not be able to launch it). When a resource is in maintenance and a user tries to launch it, they will see a message. To customize the message, click the **Configure** button. For more info, see [Site defaults \(Publishing\)](#).
4. Click **Finish** to publish the document.

General management tasks

To view published resources, select the **Publishing** category in the Parallels RAS Console. In the middle pane, expand the **Published Resources** node (if it's collapsed) to see the resources.

Right-click a resource to open a context menu. The menu has the following options:

- **Add**: Starts the publishing wizard, which you can use to publish a resource.
- **New Folder**: Allows you to add a folder to the **Published Resources** tree. Folders are described in the [Manage folders](#) section.
- **Find**: Allows you to search the list for a resource by name.
- **Duplicate**: Duplicates a selected resource. You can publish multiple resources of the same type, but configure them differently according to your needs.
- **Delete**: Deletes a published resource from the Parallels RAS Farm. This only removes the published resource item from the Farm. The actual application is not affected. To avoid accidental deletions, a dialog box is displayed asking for your confirmation.
- **Status**: Choose from **Enabled**, **Disabled**, **In maintenance**. When a resource is disabled or in maintenance, it is unavailable to users. A disabled resource will not appear in Parallels Client in the available resources list. A resource that is in maintenance will appear, but will be grayed out (in User Portal it will say so in the resource name). If a user tries to open the resource, a message is displayed. You can customize this message in [Site defaults for published resources](#).

Note that when you set the status of a published folder, all subfolders (if any) and resources that it contains will inherit the parent folder's status.

- **Delegate Permissions**: Opens the **Delegate Permissions** dialog where you can add users and grant them publishing permissions.
- **Settings audit**: Allows you to see recent changes to published resources and revert them. The changes that can be reverted include Create, Delete, and Update.

- **Verify Target(s):** Verifies that the target specified for the selected resource is valid. To see the target, select a resource and then click the **Application** tab.
- **Convert Filters to Secure Identifiers:** If filtering for a resource is specified using WinNT or LDAP, you can use this option to convert it to SID (Secure Identifier). For more information, see [Using filtering rules](#).
- **Running Instances:** Opens the **Running Processes** dialog. For more information about the dialog, please see [Managing sessions > Managing running processes](#). When the dialog is opened, a filter is applied to the process list to include only the processes for the selected published resource (a resource ID is used as a value). You can further filter the list to include only the process for a particular user (the **Username** column).

The action items at the bottom of the screen allow you to perform the following actions:

- **Add:** Same action as the **Add** menu item described above.
- **New Folder:** Same action as the **New Folder** menu item described above.
- **Delete:** Same as the **Delete** menu item described above.
- **Move Up:** Moves a selected published resource item up the list.
- **Move Down:** Moves a selected published resource item down the list.
- **Disable:** Same as the **Disable** menu item described above.
- **Sort:** Sorts resources alphabetically. For this action item to become enabled, you must select the **Published Resources** node (the topmost one) or a folder containing individual items.
- **Find:** Same as the **Find** menu item described above.
- **Running Instances:** Same as the **Running Instances** menu item described above.
- **Effective Access:** Allows you to view which published resources are available for a specific user. For complete details, see [Checking effective access](#).

After making any changes to published resources, please don't forget to click the **Apply** button to commit them to the Parallels RAS Farm.

Manage published applications

Configuring a published application

When publishing an application using a wizard, you specify multiple application parameters, such as name, executable path, etc. You can modify these options after the application is published.

To modify published application settings:

1. In the RAS Console, select the **Publishing** category and then select a desired application in the **Published Resources** tree.
2. Use the tabs in the right pane to change the application settings as described in the following subsections.

Information — review application properties

The **Information** tab displays application properties.

Publish from — configure from which servers the application is published

You can specify RD Session Hosts host pools from which an application is published on the **Publish From** tab.

Application — configure application and hosting server settings

The **Application** tab displays application- and server-specific settings.

You can modify the basic application settings (name, description, window mode) as needed. Other settings include:

- **Run.** Select the application window state (normal window, minimized, maximized).
- **Status:** Choose from **Enabled**, **Disabled**, **In maintenance**. When a resource is disabled or in maintenance, it is unavailable to users. A disabled resource will not appear in Parallels Client in the available resources list. A

resource that is in maintenance will appear, but will be grayed out (in User Portal it will say so in the resource name). If a user tries to open the resource, a message is displayed. You can customize this message in [Site defaults for published resources](#).

Note that when you set the status of a published folder, all subfolders (if any) and resources that it contains will inherit the parent folder's status.

- **Start automatically when user logs on:** Select this option to start an application as soon as a user logs in. This option works on desktop versions of Parallels Client only.
- **Exclude from session prelaunch:** For details about using this option, see [Understanding session prelaunch](#).
- **Change Icon.** Change the application icon.
- **Target.** Specify the application executable path and file name.
- **Start in.** If the **Target** field is valid, this field will be populated automatically. You can specify your own path if needed.
- **Parameters.** If the application accepts startup parameters, you can specify them in this field.

The **Server settings** section contains server-specific options that you can configure. If an application was published from multiple servers, the **Host(s)** drop-down list can be used to select individual servers and set **Target**, **Start in**, and **Parameters** values for a particular server. As an example, you should do this when different servers have the application installed in different folders, so that the **Target** and **Start in** field values would be valid on each server.

If an application is published from an MSIX package, the **Server settings** section is replaced with the **Application information** section. The options here are the same except for the **Change Application** button that allows you to select a different application for publishing.

To save the currently displayed server settings as default, click the **Save as Default Settings** button. To apply the saved default settings to a server, click the **Use Default Settings** button. These two buttons give you the flexibility of using custom settings or defaults in different server configuration scenarios. Please note that when you save settings as default, Parallels RAS will check if this Site has applications with per-server settings and will display a message asking if you

would like those servers to use the new default settings. If you say, “No”, the servers will keep their unique settings. The defaults will still be saved.

To verify that the specified **Target** and **Start In** values are correct for all servers, click the **Verify Target(s)** button. The **Target Verifier** dialog opens listing each server and the verification status in the **Progress** column. If the application is installed at a different path on one of the servers, the **Progress** column will indicate an error. In such a case, close the **Target Verifier** dialog and then select the server in the **Server(s)** drop-down list. Specify new values in the **Target**, **Start In**, and (if necessary) **Parameters** fields specific for that server. Click **Apply** to save your changes.

The **Target Verifier** dialog can also be used to verify the targets for all published applications at once. To do so, right-click **Published Resources** (the root node of the **Published Resources** tree) and then click **Verify Target(s)** in the context menu. This time, the **Target Verifier** dialog will contain all published applications and their verification status.

Filtering

Please see [Using filtering rules](#).

Routing

Please see [Configuring preferred routing](#).

Shortcuts — configuring shortcut options for a published application

Click the **Shortcuts** tab to enable the creation of the application shortcut on the user’s desktop and in the Start and Auto Start folders. When the **Auto Start** option is selected, the application will start automatically on computer startup. To use Site default settings, select the **Inherit default settings** option. You can view or modify Site defaults by clicking the **Site Defaults** link. See [Site defaults \(Publishing\)](#) for more info.

Note: Shortcuts are not available on all operating systems.

File extension — configuring file extension associations

To modify file extension association for a particular published application, click the **File Extensions** tab.

To add, remove, or modify an entry, select the **Associate File Extensions** option. To add a new extension to the list, click **Add** in the **Tasks** drop-down list (or click the + icon) and specify the desired extension.

To modify an existing association, highlight the extension and click **Properties** in the **Tasks** drop-down list (or double-click the **Parameters** column) and type the parameter.

License — configuring licensing options for published applications

Click the **Licensing** tab to configure the following licensing options:

- **Disable session sharing:** If this option is enabled, it allows you to isolate a given published application to one session. If the same application is launched more than once, the instances of the application will share the same sessions. A different application, however, will start in its own session.
- **Allow users to start only one instance of the application:** If this option is enabled, a user can only launch a single instance of the application.
- **Concurrent Licenses.** Use this option to specify the maximum number of concurrent instances the application can run. E.g. if the license of the application allows you to only run 10 instances of the application, set the Concurrent licenses option to 10 so once such limit is reached, other users cannot initiate other instances.
- **If limit is exceeded:** From this drop-down list you can specify what action should the Parallels RAS take in case any of the above licensing configured limits are exceeded.

To use Site default settings, select the **Inherit default settings** option. You can view or modify the default settings by clicking the **Site Defaults** link. See [Site defaults \(Publishing\)](#) for more info.

Display — configuring display settings for a published application

On the **Display** tab, you can configure the following options:

- **Wait until all RAS Universal Printers are redirected before showing the application:** Enable this option to wait for printers to be redirected before the application is loaded. You can also specify the maximum wait time (in seconds) for the Universal Printers to be redirected. Please note that redirecting a printer may take some time. To avoid confusion, a progress bar is shown to the user while the printers are being redirected.
- **Color Depth, Resolution, Width, Height:** Select the desired display settings for the application.
- **Start the application as maximized when using mobile clients:** This option applies only to Parallels Client running on mobile devices. When the option is selected, the application will start on a mobile device in the maximized state. This gives users the best experience while working with a remote application. This option gives the RAS administrator an easy way to always maximize an application without taking any additional steps.
- **Start in fullscreen mode for Wyse ThinOS:** If selected, the application will start in fullscreen mode in Wyse ThinOS. In some cases, the bottom part of an application may be covered by the taskbar, not allowing to see the entire application window. When this option is used, the taskbar will be hidden and the entire application window will be visible.

Note that to specify custom display values, the **Inherit default settings** checkbox must be cleared; otherwise Site defaults settings are used. To view and modify Site defaults, click the **Site Defaults** link. See [Site defaults \(Publishing\)](#) for more info.

Quick keypad

The **Quick Keypad** section allows you to select a Quick Keypad template that should be assigned to this application. The **Quick Keypads** link below the drop-down list takes you to the **Quick Keypad** category in the console where you can configure keypad templates. For more information, please see the [Quick keypad](#) section.

Manage published local applications

Configuring a published application

When publishing an application using a wizard, you specify multiple application parameters, such as name, executable path, etc. You can modify these options after the application is published.

To modify published application settings:

1. In the RAS Console, select the **Publishing** category and then select a desired application in the **Published Resources** tree.
2. Use the tabs in the right pane to change the application settings as described in the following subsections.

Information — review application properties

The **Information** tab displays application properties.

Publish from — configure from which servers the application is published

You can specify RD Session Hosts host pools from which an application is published on the **Publish From** tab.

Application — configure application and hosting server settings

The **Application** tab displays application- and server-specific settings.

You can modify the basic application settings (name, description, window mode) as needed. Other settings include:

- **Status:** Choose from **Enabled** and **Disabled**. When a resource is disabled, it is unavailable to users. A disabled resource will not appear in Parallels Client in the available resources list. If a user tries to open the resource, a message is displayed. You can customize this message in [Site defaults for published resources](#).

Note that when you set the status of a published folder, all subfolders (if any) and resources that it contains will inherit the parent folder's status.

- **Start automatically when user logs on:** Select this option to start an application as soon as a user logs in. This option works on desktop versions of Parallels Client only.
- **Change Icon.** Change the application icon.
- **URL.** The URL to a web version of the local application. If Parallels RAS cannot access the locally installed application, it will launch the web version. You can check that the application can be accessed by clicking the **Test** button.
- **Target.** Specify the application executable path and file name.
- **Start in.** If the **Target** field is valid, this field will be populated automatically. You can specify your own path if needed.
- **Parameters.** If the application accepts startup parameters, you can specify them in this field.
- **Run.** Select the application window state (normal window, minimized, maximized).

Filtering

Please see [Using filtering rules](#).

Routing

Please see [Configuring preferred routing](#).

Shortcuts — configure shortcut options for a published application

Click the **Shortcuts** tab to enable the creation of the application shortcut on the user's desktop and in the Start and Auto Start folders. To use Site default settings, select the **Inherit default settings** option. You can view or modify Site defaults by clicking the **Site Defaults** link. See [Site defaults \(Publishing\)](#) for more info.

Note: Shortcuts are not available on all operating systems.

File extension — configuring file extension associations

To modify file extension association for a particular published application, click the **File Extensions** tab.

To add, remove, or modify an entry, select the **Associate File Extensions** option. To add a new extension to the list, click **Add** in the **Tasks** drop-down list (or click the + icon) and specify the desired extension.

To modify an existing association, highlight the extension and click **Properties** in the **Tasks** drop-down list (or double-click the **Parameters** column) and type the parameter.

License — configuring licensing options for published applications

Click the **Licensing** tab to configure the following licensing options:

- **Disable session sharing:** If this option is enabled, it allows you to isolate a given published application to one session. If the same application is launched more than once, the instances of the application will share the same sessions. A different application, however, will start in its own session.
- **Allow users to start only one instance of the application:** If this option is enabled, a user can only launch a single instance of the application.
- **Concurrent Licenses.** Use this option to specify the maximum number of concurrent instances the application can run. E.g. if the license of the application allows you to only run 10 instances of the application, set the Concurrent licenses option to 10 so once such limit is reached, other users cannot initiate other instances.
- **If limit is exceeded:** From this drop-down list you can specify what action should the Parallels RAS take in case any of the above licensing configured limits are exceeded.

To use Site default settings, select the **Inherit default settings** option. You can view or modify the default settings by clicking the **Site Defaults** link. See [Site defaults \(Publishing\)](#) for more info.

Display — configuring display settings for a published application

On the **Display** tab, you can configure the following options:

- **Wait until all RAS Universal Printers are redirected before showing the application:** Enable this option to wait for printers to be redirected before the application is loaded. You can also specify the maximum wait time (in seconds) for the Universal Printers to be redirected. Please note that redirecting a printer may take some time. To avoid confusion, a progress bar is shown to the user while the printers are being redirected.
- **Color Depth, Resolution, Width, Height:** Select the desired display settings for the application.
- **Start the application as maximized when using mobile clients:** This option applies only to Parallels Client running on mobile devices. When the option is selected, the application will start on a mobile device in the maximized state. This gives users the best experience while working with a remote application. This option gives the RAS administrator an easy way to always maximize an application without taking any additional steps.
- **Start in fullscreen mode for Wyse ThinOS:** If selected, the application will start in fullscreen mode in Wyse ThinOS. In some cases, the bottom part of an application may be covered by the taskbar, not allowing to see the entire application window. When this option is used, the taskbar will be hidden and the entire application window will be visible.

Note that to specify custom display values, the **Inherit default settings** checkbox must be cleared; otherwise Site defaults settings are used. To view and modify Site defaults, click the **Site Defaults** link. See [Site defaults \(Publishing\)](#) for more info.

Quick keypad

The **Quick Keypad** section allows you to select a Quick Keypad template that should be assigned to this application. The **Quick Keypads** link below the drop-down list takes you to the **Quick Keypad** category in the console where you can configure keypad templates. For more information, please see the [Quick keypad](#) section.

Manage published desktops

Configuring a published desktop

When publishing a desktop using a wizard, you have to specify the desktop settings, such as display size, etc. You can modify these options after the desktop has been published.

To modify a published desktop, select it in the **Published Resources** tree in the **Publishing** category.

Sites — configuring from which sites a published desktop is available

By default, a published desktop is available through all of the available sites. To restrict access to a specific Site or a Site group, select a desktop in the **Published Resources** tree and then click the **Sites** tab in the right pane. Select the sites from which the desktop should be available.

Note: For the **Sites** tab to be available, you need more than Site in a farm.

Publish from — configuring from which RD Session Hosts a desktop is published

When configuring an RD Session Host desktop, you can specify from which servers it should be published. To do so, click the **Publish From** tab and select the desired servers.

Desktop — configuring desktop name, size and other properties

Depending on the desktop type, click the **Desktop**, **Remote PC Desktop**, or **Virtual Desktop** tab to configure the desktop name, description, icon, resolution, status, and other settings.

- **Status:** Choose from **Enabled**, **Disabled**, **In maintenance**. When a resource is disabled or in maintenance, it is unavailable to users. A disabled resource will not appear in Parallels Client in the available resources list. A

resource that is in maintenance will appear, but will be grayed out (in User Portal it will say so in the resource name). If a user tries to open the resource, a message is displayed. You can customize this message in [Site defaults for published resources](#).

Note that when you set the status of a published folder, all subfolders (if any) and resources that it contains will inherit the parent folder's status.

- **Connect to administrative session:** Select this option if you want users to connect to the administrative session. Note that a user connecting to a desktop with this option enabled must have administrative privileges; otherwise "Access is denied" error will be shown to the user.
- **Start automatically when user logs on:** Select this option if you want to open a desktop as soon as a user logs in. For the information about **Exclude from session prelaunch** option, see [Understanding session prelaunch](#).
- **Desktop Size:** Select a desired desktop size from the drop-down list.
- **Multi-Monitor:** Select whether the multi-monitor should be enabled, disabled, or whether the client settings should be used.

Filtering

Please see [Using filtering rules](#).

Routing

Please see [Configuring preferred routing](#).

Shortcuts — configuring shortcut options for a published desktop

Click the **Shortcuts** tab to enable the creation of a shortcuts on the user's desktop and in the Start and Auto Start folders. When the Auto Start shortcut is enabled, the application will start automatically on computer startup. To use Site default settings, select the **Inherit default settings** option. See [Site defaults \(Publishing\)](#) for more info.

Note: This option is not available on all operating systems.

Manage published documents

Configuring a published document

When publishing a document using a wizard, you have to specify the document settings. These options can be modified after the document has been published.

To modify a published document, select it in the **Published Resources** tree in the **Publishing** category and then use the tabs in the right pane to configure the published document settings.

Sites — configuring from which sites a published document is available

By default, a published document is available through all available sites. To restrict access to a specific Site or a Site group, click the **Sites** tab in the right pane. Select the sites from which the document should be available.

Note: For the **Sites** tab to be available, you need more than one Site in a Farm.

Publish from — configuring from which servers a document is published

Click the **Publish From** tab and select the servers from which the document should be published. Please note that a server must have the application installed that can open this particular document type.

Application — configuring server-specific document settings

By default, the settings configured in the **Target** (application path), **Start In**, and **Parameters** fields apply to all servers a document is published from. If a document exists in a different folder on one (or more) of the servers, you can specify the above settings for a specific server or servers individually.

To do so:

1. Click the **Application** tab and.

2. Select a server in the **Server(s)** list.
3. Specify the **Target**, **Start In**, and **Parameters** (optional) properties. The values that you specify will apply to the selected server only. Repeat the steps for other servers if needed.
4. Click the **Verify Target(s)** button to verify the document path on all servers from which this application is published. The results are displayed in the **Target Verifier** dialog where you can see whether the target is correct or not for each server.

Filtering

Please see [Using filtering rules](#).

Routing

Please see [Configuring preferred routing](#).

Shortcuts — configuring shortcut options for a published document

Click the **Shortcuts** tab to enable the creation of shortcuts on the user desktops, shortcuts in the **Start** folder and shortcut in the **Auto Start** folder. When the **Auto Start** shortcut is enabled, the application will start when the user's computer is started.

Note: This option is not available on all operating systems.

File extension — configuring file extension associations

To modify file extension association for a particular published document, click the **File Extensions** tab. To add a new extension to the list, click **Tasks > Add** and specify the extension. To modify the extension parameters, highlight the extension and click **Tasks > Properties**.

Licensing — configuring licensing options for published documents

Click the **Licensing** tab to configure any of the below licensing options:

Select the **Inherit default settings** option to use the defaults. To specify your own settings, clear the option and set the following options:

- **Disable session sharing.** If this option is enabled, it allows you to isolate a given published application to one session. If the same application is launched more than once, the instances of the application will share the same sessions. A different application, however, will start in its own session.
- **Allow users to start only one instance of the application.** If this option is enabled, a user can only launch a single instance of the application.
- **Concurrent Licenses.** Use this option to specify the maximum number of concurrent instances the application can run. E.g. if the license of the application allows you to only run 10 instances of the application, set the Concurrent licenses option to 10 so once such limit is reached, other users cannot initiate other instances.
- **If limit is exceeded.** From this drop-down list you can specify what action should the Parallels RAS take in case any of the above licensing configured limits has been exceeded.

To use Site default settings, select the **Inherit default settings** option. See [Site defaults \(Publishing\)](#) for more info.

Display — configuring display settings for a published document

Click the **Display** tab to configure the color depth of the published document, resolution, width and height. If these options are left at their default values, the client-specified options will take over.

You can also enable the option to wait for the Universal Printers to be redirected before the application is loaded. When enabling this option, you can also configure the maximum wait time (in seconds) for the Universal Printers to be redirected. To use Site default settings, select the **Inherit default settings** option. See [Site defaults \(Publishing\)](#) for more info.

Manage folders

Folders are used to organize published resources and to facilitate filtering options.

There are two types of folders that you can create in the **Published Resources** tree in the Parallels RAS Console:

- **Folders for administrative purposes.** Folders of this type are intended for Parallels RAS administrators (users of the Parallels RAS Console). They are used to logically organize published resources in the Parallels RAS Console but they do not appear in the Parallels Client launchpad on user devices. These folders are used to help administrators manage published resources more efficiently.
- **Regular folders.** These folders are similar to administrative folders described above but they do appear in the launchpad on user devices. You normally use these folders to group published resources by type (e.g. office applications, specific business applications, utilities, etc.).

Creating a folder

To create a new folder:

1. In the RAS Console, select the **Publishing** category.
2. Right-click anywhere in the **Published Resources** tree and choose **New Folder** (or click the **[+] New Folder** icon at the bottom).
3. In the **Folder** dialog, specify a folder name and an optional description.
4. To make it a folder for administrative purposes, select the **Use for administrative purposes** option. To publish a regular folder, clear the option. See above for the explanation of the two folder types.
5. When creating a regular folder, you can change its icon by clicking the **Change icon** button. Administrative folders use a built-in icon that cannot be changed. Icons appear in the **Publishing** category in the Parallels RAS Console and in the Parallels Client launchpad (regular folders only).
6. On the next page, specify the initial status of the resource (the folder). Choose from the following options:
 - **Enabled:** End users can see the folder and will be able to launch published resources that it contains.
 - **Disabled:** The folder will not appear in Parallels Client.
 - **In maintenance:** The folder will appear in Parallels Client but users will not be able to launch resources that it contains. If the folder has

subfolders, they inherit the status of the parent folder, which means that none of the resources contains in any of the folders in the hierarchy will be accessible to users. When a folder is in maintenance and a user tries to launch a resource from it, they will see a message. To customize the message, click the **Configure** button. For more info, see [Site defaults \(Publishing\)](#).

7. Click **Finish** to create the folder.

Managing folders

To modify an existing folder:

1. Select a desired folder in the **Published Resources** tree.
2. The **Information tab** in the right pane displays the folder information (read-only).
3. On the **Folder** tab, you can see and modify the folder name and description. You can also select or clear the **Use for administrative purposes** option to change the folder type (see above for an explanation). To change the folder icon, click the **Change icon** button. Note that the button is disabled if the **Use for administrative purposes** option is selected.
4. The **Filtering** tab specifies filtering options. Once set, these options will be inherited by all published resources in that folder. For more information, please see [Using filtering rules](#).
5. For the information about **Routing**, please see [Configuring preferred routing](#).

Adding published resources to a folder

To add a published resource to a folder, first add it to the root location and then drag it to the desired folder.

Delegating permissions to custom administrators

If you have custom administrators in your Farm, you can delegate permissions to them to manage a folder. This is specifically useful when a power administrator needs to grant permissions to a custom admin. To grant folder rights:

1. Right-click anywhere in the **Published Resources** pane.
2. In the context menu, select **Delegate Permissions**.

3. In the dialog that opens, select a user to grant folder permissions to. In the lower right pane of the **Delegate Permission - Publishing** dialog, select permissions (view, modify, add, delete) for a desired folder you want the user to have. For more information about custom administrators, see [Managing administrator accounts](#).
4. You can also grant folder rights to a custom administrator via the Administration category as described in [Configuring Administrator Accounts Permissions](#).

Site defaults (Publishing)

The **Default Settings** dialog allows you to view and modify Site default settings for publishing. Published resources can inherit the following groups of settings from Site defaults:

- Shortcuts
- License
- Display
- Maintenance

To open the **Default Settings** dialog, navigate to **Farm > Site**. Click the **Tasks** menu and choose **Site defaults > Published resources**. The dialog consists of tabs, which are described below.

Shortcuts

In this tab specify whether and how the published resource shortcuts should be created on the user's computer. The following options are available:

- **Create shortcut on Desktop.** If selected, a shortcut will be created on the user's desktop.
- **Create shortcut in Start folder.** If selected, a shortcut will be added to the **Start** folder. You can specify the target subfolder name and path in the field provided. The default (and only) %Groups% variable will add additional subfolders as they appear on the host server where the published resource is hosted. For example, if the resource is located in "Myapps > Games" on the host server, the same folder structure will be added to the path. Note that you cannot use any custom variables.

- **Create shortcut in Auto Start folder.** If selected, the published resource will start automatically on computer startup.

License

The **License** tab contains the following options:

- **Disable session sharing.** If this option is enabled, it allows you to isolate a given published application to one session. If the same application is launched more than once, the instances of the application will share the same sessions. A different application, however, will start in its own session.
- **Allow users to start only one instance of the application.** If this option is enabled, a user can only launch a single instance of the published resource.
- **Concurrent Licenses.** Use this option to specify the maximum number of concurrent instances the published resource can run. For example, if the license of the application allows you to only run 10 instances of the application, set the **Concurrent licenses** option to 10, so once this limit is reached, other users cannot initiate other instances.
- **If limit is exceeded.** Specifies which action should Parallels RAS take in case any of the licensing limits configured above are exceeded.

Display

The **Display** tab contains the following options:

- **Wait until all RAS Universal Printers are redirected before showing the application.** Enable this option to wait for printers to be redirected before the application is loaded. You can also specify the maximum wait time (in seconds) for the Universal Printers to be redirected. Please note that redirecting a printer may take some time. To avoid confusion, a progress bar is shown to the user while the printers are being redirected.
- **Color Depth, Resolution, Width, Height.** These options specify the desired display settings for the application.
- **Start the application as maximized when using mobile clients.** This option applies only to Parallels Client running on mobile devices. When the option is selected, the application will start on a mobile device in the maximized state. This gives users the best experience while working with a remote application. This option gives the RAS administrator an easy way to always maximize an application without taking any additional steps.

Maintenance

The **Maintenance** tab allows you to specify a message that users will see when trying to launch a published resource in maintenance. When a resource is in maintenance, it will still appear in Parallels Client, but will be grayed out (in User Portal, it will say so in the resource name). If a user tries to open the resource, they will see the message that you specify here. If you modified a message, but want to return the default one, select a message in the desired language and click **Tasks > Reset to default**. To reset messages in all languages, click **Tasks > Reset all to default**.

You can replicate the Site settings described above to other sites in your Parallels RAS Farm. To do so, select the **Replicate settings** option in a desired tab. All settings contained in the tab will be replicated.

Using filtering rules

Filtering rules is a feature that allows you to control who can access a particular published resource. Each rule consists of one or several criteria for matching against user connections. In turn, each criteria consists of one or several specific objects that can be matched.

You can match the following objects:

- User, a group the user belongs to, or the computer the user connects from.
- Secure Gateway the user connects to.
- Client device name.
- Client device operating system.
- Theme.
- IP address.
- Hardware ID. The format of a hardware ID depends on the operating system of the client.


Notice the following about the rules:

- Criteria are connected by the AND operator. For example, if a rule has a criteria that matches certain IP addresses and a criteria that matches client device operating systems, the rule will be applied when a user connection matches one of the IP addresses AND one of the client operating systems.

- Objects are connected by the OR operator. For example, if you only create a criteria for matching client device operating systems, the rule will be applied if one of the operating systems matches the client connection.
- The rules are compared to a user connection starting from the top. Because of this, the priority of a rule depends on its place in the rule list. Parallels RAS will apply the first rule that matches the user connection.
- The default rule is used when no other rule is matched. You can set it to either **Allow if no other rule matches** or **Deny if no other rule matches**, but no criteria is available for this rule.

To create a new rule:

1. In the RAS Console, navigate to **Publishing**.
2. Select the **Filtering** tab.
3. Click the **Tasks** drop-down menu and choose **Add** (or click the **[+]** icon)
4. Specify criteria for the rule. You will find the following controls:
 - **Allow if** and **Deny if**: specifies whether the resource must be accessible when a user connection matches all the criteria. Click on the link to switch between the two options.
 - **(+)**: adds a new criteria. If you want to match a Secure Gateway, a client device name, a client device operating system, a Theme, an IP address or a hardware ID, click **(+)**. In the context menu that appears, select the type of an object that you want to match and add the specific objects in the dialog that appears. The new criteria appears on the next line.
 - **(X)**: Deletes a specific object from matching. For example, you want to delete IP address 198.51.100.1 from matching, click **(X)** next to it. This control appears when at least one object is added. If all objects in a criteria are deleted, the criteria is removed.
 - **is** and **is not**: specifies whether the resource must be accessible when a user connection matches the criteria. Click on the link to switch between the two options. This control appears when at least one object is added.
 - **configure**: edits the list of objects to be matched. Click this link to add or delete new objects. Note that for the first criteria (**User or group**) this link is called **everyone**. It will change to **configure** once you specify objects for this criteria.



Configuring preferred routing

Overview

Preferred routing is a useful feature when Parallels RAS users with different geo located deployments are connecting to the same Parallels RAS Farm/Site. A common access layer usage (RAS Secure Gateway, HALB, or a third-party load balancer) is not optimal if a resource is located in a different data center in the same RAS Farm/Site. The solution is to configure a preferred access layer server for a specific published resource, in which case any user would connect to a default Secure Gateway, but would be redirected using proximity rules set by the administrator. Typically, using the Secure Gateway closest to the session host provides improved user experience, reduced internal network traffic and associated costs along with providing better use of resources.

Note: Preferred routing doesn't apply to Azure Virtual Desktop published objects. endhint %}

Here's how preferred routing works:

1. Parallels Client establishes a connection with a Secure Gateway using a standard authentication.
2. Through the RAS Connection Broker, the resource's preferred route (if configured) is identified.
3. Parallels Client receives the preferred public address to launch the resource.
4. Parallels Client then tries to launch the resource through the redirected address and falls back to the original Gateway if it fails.

Configure preferred routing

To configure preferred routing for a published resource:

1. In the RAS Console, select the **Publishing** category.
2. Select an existing published resource and then select the **Routing** tab.
3. Select the **Enable preferred routing** option.
4. Click **Tasks > Add**. The **Add preferred route** dialog opens. Read on.

In the **Add preferred route** dialog:

1. Type a name for this route and an optional description.
2. From the **Type** drop-down list, select one of the following:
 - **HALB Virtual Server**: Select a RAS HALB virtual server from the list below the **Type** field. Note that for a RAS HALB virtual server to appear in the list, the HALB server must have a public address specified, as far as you cannot add HALB server here.
 - **Secure Gateway**: Same as for HALB virtual server (described above), the **Public address** field must have a value for the Gateway to appear in the list. See the **Public address** field when you create or configure a RAS Secure Gateway.
 - **Custom**: A third-party load balancer will be used. Select this option, then click **Tasks > Add** and specify the server properties in the list below the field. The properties include **Name**, **Description**, **Public address**, **Port**, and **SSL port**. You can add as many servers as required and then select one of them to be used for a given published resource.

When configuring preferred routing, please also consider the following:

- If routing fails, an automatic fallback to the originating address is carried out.
- If routing is enabled in the RAS Console, but not configured, the administrator will see an error message and will have to either configure or disable it.
- It is recommended to use Folders (configured for administrative purpose) in case routing needs to be configured on many resources. Routes are not inherited from parent folder if the routing is enabled on the child object, i.e. only one set of routes is used.
- Same user credentials will be used when redirecting the RDP traffic to the same RAS Site. Users will not be asked to re-enter credentials.

- If you had existing session tunneled in a particular Secure Gateway and using session sharing, the same session workflow path will be used (if a published resource is also available on the same session host) irrespective of configured routing.
- Routing is supported in SAML environments.
- Supported clients are Windows, macOS, Linux, Android, iOS, Web.

Specifying public address when inviting users

When you use the **Invite Users** wizard, you can specify a public address on the second page where you specify target platforms and connection options. This way, you can set a preferred routing for a group of users in a specific geo location. For more information, see [Invite users](#).

Deleting or disabling Gateway or HALBs

If an administrator tries to delete a Gateway or HALB when they are being used as preferred routing, the information about objects using it will be shown on the screen, so no accidental deletion takes place.

Understanding session prelaunch

When a user opens a remote application or desktop, a session must first be launched. Launching a session can take time, which will result in the user waiting for the application to start. To improve user experience, a session can be launched ahead of time, before the user actually opens an application. When a session is prelaunched, it will all happen in the background, so the user will not see any windows or message boxes on the screen. When the user starts an application, it will open using the prelaunched session, so it will start very quickly.

For the information about how to configure session prelaunch, see [Client policies](#).

When you configure session prelaunch, the following options are available:

- **Off.** No session prelaunch is used.
- **Basic.** A session is prelaunched as soon as the user gets the application listing. The assumption is, the user will open an application within the next few minutes. The session will stay active for 10 minutes. If the user doesn't open an application during that time, the client will disconnect from the session.

- **Machine Learning.** When the application listing is acquired, a session is prelaunched based on user habits. With this option enabled, Parallels Client will record and analyze the user habits of launching applications on a given day of the week. A session is started a few minutes before the user usually opens an application.

You can configure rules when session prelaunch must not be used. The following options are available:

- Specify dates on which the prelaunch must not be used.
- Exclude a published resource from the session prelaunch decision making. If a resource is excluded from the analysis, it is never considered by Parallels Client when making a decision whether to prelaunch a session. For example, when you have a server on which you never want to prelaunch sessions, you can flag all published resources hosted by that server as to be excluded from session prelaunch. To exclude a published resource from session prelaunch, in the RAS Console, navigate to **Published Resources**, select a resource and then select the **Exclude from session prelaunch** option.

Checking effective access

Filtering rules described in the [previous section](#) allow you to configure who can access a particular published resource. If a Parallels RAS user cannot see one or more published resources in Parallels Client, you would normally have to check filtering settings for each resource to make sure that it is published for a given user. The Effective Access functionality simplifies this task by allowing you to view in one place which published resources are available for a user and which are not.

To open the **Effective Access** dialog, select the **Publishing** category in the Parallels RAS Console and then click the **Effective Access** item in the toolbar at the bottom of the window (if you don't see the item, maximize the console window). You can also open the dialog by right-clicking anywhere in the **Published Resources** pane and choosing **Effective Access** in the context menu.

The **Effective Access** dialog allows you to specify a user (and optionally additional criteria) and then view published resources this user is allowed to access. To choose a user, do one of the following:

- Type the user name in the **User** field, or click the [...] button next to it and use the **Select User or Group** dialog to select a user.

- Select a device owned by this user from the list of known devices. To do so, click the **Select a Device** button then select a device. Note that if a device has never been used to connect to this Parallels RAS Farm, it will not be included in the list. For more information, see the [Monitoring devices](#) section. After selecting a device, click **OK** to return to the **Effective Access** dialog. All of the fields will be automatically populated using properties of the selected device.

Once you specify a user, enter the additional criteria if needed (all fields except **User** are optional):

- **Client.** Client name assigned to a device. This could be a computer name, FQDN, or a custom name that the user could have set in Parallels Client.
- **IP Address.** Client IP address.
- **MAC.** Client MAC address.
- **Gateway.** RAS Secure Gateway name through which the client connects to the Farm.

The **Manage groups** button allows you to preview how user access changes if the user is added to one or more groups. When you click the button:

1. The **Manage Groups** dialog opens listing groups to which the user already belongs.
2. Click the **[+]** button to add the user to one or more additional groups. Note that this will only be a simulation; the user will not be actually added to any additional group.
3. To remove a “simulated” group, select it in the lower pane and click the **[-]** button.
4. Click **Close** to return to the **Effective Access** dialog.

Finally, to view the effective access information for the specified user, click the **View** button. This opens the **Effective Access - Summary** dialog, which displays the following information:

- The left pane contains the complete list of resources published in the current Site. To view only the resources that the specified user can access, select the **Show only allowed published resources** option. If the user is not allowed to access a resource, the resource name is highlighted in red.
- The right pane contains information whether the user is allowed to access a resource selected in the left pane and whether filtering is enabled for the

selected resource. Additional information may include filtering details and extended group membership.

By looking through the resource list, you can see which resources the user can or cannot access and take appropriate actions if necessary. If needed, you can export the effective access information to a CSV file. To do so, click the **Export** button and specify a file name. The CSV file has the following columns:

- **Name.** Application name.
- **ID.** Application ID.
- **Accessible.** Whether the application is accessible to the user (Yes or No).
- **Rule.** Filtering rule. If no rules are configured for the application, the column will have no value.

Specifying client settings

To specify client settings for published resources, navigate to **Farm** > <Site> > **Settings** and select the **Client Settings** tab. On this page, you can specify how published application icons are displayed on the client side and some other options.

Select icon resolution

Published resources are displayed in Parallels Client as icons or as a list. You can specify which resolution should be used when the resources are displayed as icons. Select from the following options:

- **Send standard resolution icons.** Standard resolution icons.
- **Send high resolution icons.** High resolution icons. Please note that this option will use more network bandwidth.

Enable, disable or change the overlay icon

Note: These configuration changes are applicable to desktop clients only (Windows, Mac, Linux). They have no effect on mobile and Web Clients.

The other options on this tab is to enable, disable or update the overlay icon. An overlay icon is placed over a standard application icon to indicate that it's a remote application served by Parallels RAS. When you launch a remote application from Parallels Client, the application icon is displayed on the local desktop (e.g. on the

taskbar in Windows or Dock in macOS). By using an overlay icon, you give the user the ability to tell at a glance which of their running applications are remote Parallels RAS applications and which are local (or any other kind).

Parallels RAS uses the Parallels logo as the overlay icon by default. However, the administrator may also change this to use the standard Microsoft RemoteApp overlay icon. When using Parallels logo as the overlay icon, an application icon on a local computer will look like the following sample icons:





As you can see, these are standard icons used by the Windows Calculator and Paint applications with the Parallels logo icon (red parallel lines in the right corner) displayed on top of them. When a user notices the overlay, they'll know right away that this is a remote application served by Parallels RAS, not a local Windows app.

Show password expiration reminder

You can automatically remind your Parallels RAS users to change their domain password when it nears the expiration date. To enable this functionality, select the **Show password expiration reminder** option. When it is enabled, a Parallels Client user whose password is about to expire will see a notification right after they connect to Parallels RAS. Note that the option is disabled by default.

Session reset

You can force user sessions to reset on user logoff by selecting the **Force session reset on logoff from Parallels Client** option. This is useful for resetting frozen user sessions.

Quick keypad

The **Quick Keypad** category in the Parallels RAS Console allows you to define custom keys to perform common actions in published applications running on mobile devices. Custom keys appear above the standard keyboard in iOS and Android and can be tapped just like any other key on the virtual keyboard.

This feature is designed for users who run published applications on a phone or a tablet. When a particular software requires repeated selection of certain menu or toolbar items, using custom keys can significantly improve user experience. For example, let's say a user has some data entry task which requires them to press **File > New** and **File > Save** menu items over and over again. If you define two custom keys to perform these actions, the user will see them above the standard keyboard in iOS or Android, so instead of tapping the application's native menu items (which can be cumbersome), they can tap these keys, which is much easier and quicker.

To define custom keys, select the **Quick Keypad** category in the Parallels RAS Console. The **Quick Keypads** view in the right pane allows you to create a Quick Keypad template. A template is created for a specific application (or a group of applications with the identical UI design) and contains shortcuts to perform common actions in an application. Once a template is created, you assign it to a published application or a group of applications, so each application (or a group) has its own Quick Keypad.

To create a Quick Keypad template:

1. Click the **Tasks** drop-down list and choose **New Quick Keypad** (or click the **[+]** icon).
2. Specify a Quick Keypad template name (e.g. "Office apps").
3. You can organize a Quick Keypad using a multi-level menu system. If you want to do this, click the **New menu** item and specify the menu item name. You can add sub-menu items too. To move a menu item across the tree, simply drag and drop it to the desired tree node.
4. When you have your basic menu structure defined, you can add shortcuts (or you can do it any order you like).
5. To add a shortcut, click the **New shortcut** item.
6. In the **Label** field, enter the name (e.g. "New").
7. Click the **Shortcut** field and press a shortcut on the keyboard as you would in the target application. For example, the standard shortcut to create a new

document in many applications is Ctrl+N, so to input this shortcut, you would press and hold Ctrl and then press N. The shortcut will appear in the field as “Ctrl+N”. You can input up to three shortcuts in this field.

8. To add another shortcut to the template, click the **New shortcut** item again. Repeat until all desired shortcuts are defined.
9. Click **OK** to close the dialog. The new template will appear in the **Quick Keypads** list.

To modify the template, right-click it and choose **Properties**.

You now need to assign the template that you created to an application (or multiple applications). To do so:

1. Right-click a template and choose **Assign to Application** (you can also use the **Tasks** drop-down list or click the “link” icon).
2. In the **Assign Quick Keypad Template** dialog, select one or more applications to which the template should be assigned.
3. Click **OK** when done.

When a remote user runs an application on their mobile device and opens a virtual keyboard, they will see the extra keys corresponding to shortcuts that you defined for a Quick Keypad template. Tapping a key will perform the corresponding action (e.g. Ctrl-N, which will open a new document).

Exporting and importing a Quick Keypad template

To easily move a Quick Keypad template from one Parallels RAS Farm to another, use the Import and Export functionality. To export a template, right-click a template and choose **Export**. Specify the file name and location and click **Save**. To import a template, right-click on an empty space in the **Quick Keypads** list and choose **Import**. You can also perform these actions using the **Tasks** drop-down list.

Session Management

Overview

When users connect to Parallels RAS and establish a session, the session information is displayed in the Parallels RAS Console in the following locations:

- The **Sessions** category (new since Parallels RAS 18.1).

- The **Sessions** tab in RD Session Hosts, VDI, and Azure Virtual Desktop views (**Farm > Site > RD Sessions Hosts > VDI > Azure Virtual Desktop**).

The **Session** category displays user sessions for all available host types, including RD Session Hosts, VDI, and Azure Virtual Desktop. This is the place where you can view all current sessions irrespective of the type of a server hosting a session. Individual **Sessions** tabs display sessions for their respective host types.

Sessions category

When you select the **Session** category (in the main category list), the following two tabs are displayed in the right pane of the RAS Console:

- **Users**: Lists user sessions for all available host type.
- **Resources**: Lists currently running published resources (apps and desktops) from hosts of all types.

You can filter the lists in the **Sessions** category by clicking **Tasks > Search** (or clicking the magnifying glass icon) and specifying the criteria in one or more column headings. For example, you can filter the list on the **Users** tab by host type using the **Source** column, which can contain one of the following values:

- **RDSH**: RD Session Host
- **VDI**: Virtual desktop
- **RemotePC**: Remote PC through VDI
- **AVD**: Azure Virtual Desktop

Sessions tabs

The **Sessions** tabs display user sessions for their respective host types. To see sessions for a particular host, you can filter the list by host name.

Please note that when you open the **Sessions** category or a **Sessions** tab, some of the columns in a list may not be populated right away. This is because it takes some time to calculate these values. The examples of such columns include **Logon duration**, **UX Evaluator**, **Latency**. Simply wait a few seconds and the values will appear in the list.

Most of the information in this chapter is common to both the **Sessions** category and **Sessions** tabs. Specifics and differences are described where applicable.

Session information

To see the complete information for a specific session, right-click it and choose **Show information**. This opens the **Session Information** dialog where session properties are grouped by functionality.

The following groups are displayed:

- **Session Setup:** Contains general session information.
- **Logon Details:** Displays logon metrics that can be used to evaluate the logon process.
- **Session Details:** Displays the current session state, logon time, in/out data size, and general session information.
- **Connection Details:** Displays connection and authentication details.
- **User Experience:** Displays metrics that can be used to evaluate user experience.
- **Client Details:** Displays information about the user device and Parallels Client type and version.

Parallels RAS 18 introduces over 25 new session detail metrics available. The following tables give an overview of these new and some of the important preexisting metrics.

Session Setup

Metric	Description
Session host*	Session host name
Source*	

* New since Parallels RAS 18.1

Logon details

Metric	Description
Logon duration*	Time taken to logon excluding the time waiting on UI.

Metric	Description
Logon duration breakdown*	
User Profile*	User Profile method in use: FSLogix, User Profile Disk, or Other (also contains additional information, such as error code).

* New since Parallels RAS 18.0

User Experience

Metric	Description
UX Evaluator*	This is the time interval measured at the client between the first step (user action) and the last step (graphical response displayed).
Connection quality*	Connection quality rating (poor – excellent)
Latency*	Network latency
Transport Protocol*	TCP or UDP (over RDP)
Bandwidth availability*	Bandwidth availability as seen from the client
Reconnects*	Number of reconnects the current session suffered from inception (excluding graceful ones)
Last Reconnects*	Number of reconnects suffered from the current device session (excluding graceful ones)
Disconnect reason*	The last session disconnect reason

* New in Parallels RAS 18.0

Session details

Metric	Description
Session State	Active, Idle, Disconnected, etc
Logon time	Time and date when the session was established
Session Length	Time the session has been established
Idle Time	Time the session has been idle
Incoming Data*	Amount of data received from the client

Metric	Description
Outgoing Data*	Amount of data sent to the client
Resolution	Session resolution
Color Depth	Session colors depth
Bandwidth Usage*	Bandwidth used by the client

* New since Parallels RAS 18.0

Client details

Metric	Description
Device name	Name of the device from which the session was established
IP Address	Client private IP address
Client OS*	The operating system on which the client is running
Client OS version*	The operating system version on which the client is running
Client version*	The Parallels Client version is use
Policy	The policy applied to the user session

* New since Parallels RAS 18.0

Export session information

To export the session information to a CSV file, click the **Export** button in the **Session Information** dialog and specify the location and file name.

You can also export session information from the main session list by clicking **Tasks > Export**. Note that depending on what is selected in the list, the following will be exported:

- A single session — the information about that session is exported.
 - Multiple sessions — the information for all selected sessions is exported.
 - No selection — the information about all current sessions is exported.
- Exported CSV includes the exported session details along with export detail in the following format:

Session details (%Server type% such as RD Session Hosts) from Parallels RAS Farm %Farm name% and Site %Site name% exported by %Administrator% on %date% at %time%

Monitoring settings

The **Monitoring Settings** functionality allows you to add colors to thresholds to identify Warning and Critical levels for better aid to Administrators or helpdesk.

To configure monitoring settings, in the **Sessions** category or **Sessions** tabs, click **Tasks > Monitoring settings**. The dialog opens where you can configure settings for various session metrics:

1. Select a metric for which you want to enable color coding.
2. Specify a threshold in the **Warning** and **Critical** columns. The Warning threshold is denoted by the orange color. The Critical threshold is denoted by the red color.
3. If you want to see just the Critical color (red), then set both thresholds to the same value, in which case only the red color will be used when the threshold is reached.

When a metric with color coding enabled is below any of the specified thresholds, it is highlighted with the green color in the session list. When a threshold is reached, the value of a metric is highlighted using the corresponding threshold color (orange or red). Note that critical threshold value can be greater or equal to warning threshold value. In case both warning and critical values are equal then the critical color coding is used which is red.

Monitoring Settings are set globally, which means that other RAS admins will be able to see and change them.

Managing sessions

To manage a session (or multiple sessions at the same time), select one or more sessions and then use the **Tasks** drop-down list to choose from the following actions:

- **Refresh**. Refresh the list.
- **Disconnect**. Disconnect the selected session(s).
- **Log off**. Log off the session(s).
- **Send message**. Opens the **Send Message** dialog where you can type and send a message to the session owner(s).
- **Remote control**. Remotely control the selected user session. To establish a connection, domain or local Windows account credentials (whichever the

user used to log in to this computer) of the current RAS Console administrator will be used. Note that the current user (specifically if it's the local Windows user) may not be permitted to connect to the remote computer. In such a case, use the **Remote control (prompt)** option (described below). See also the **User session remote control** subsection below for important information.

- **Remote control (prompt)**. Same as above but prompts you to enter credentials. Use this option when the current user credentials cannot be used to control a session.
- **Show processes**. Display and manage running processes. See **Managing processes** below for details.

Note: When sessions are in the Down state, recovery is impossible; the host must reboot. Logoff, remote control, and similar functions are disabled for that session. Other users on the same host can continue working.

If this issue reoccurs, check if your hosts have enough CPU and memory, or set a schedule to reboot the host automatically that will Trigger only if a session state is Down.

User session remote control

The **Remote Control** and **Remote control (prompt)** menu options (see above) allow you to shadow a user RDS session. There are limitations as described below:

- Parallels RAS cannot shadow RDS sessions running on Windows 7 and Windows Server 2008 R2. This doesn't work even with native tools.

Managing processes

The **Tasks > Show processes** option opens the **Running Processes** dialog where you can view running processes for one or more hosts.

Note: You can also open the **Running Processes** dialog by right-clicking a server in the main host list and choosing **Show Processes**. This will open the **Running Processes** dialog with a filter applied to it to display only the processes that belong to the selected host.

In the **Running Processes** dialog, use the **Show processes from** drop-down list to filter the list using the following options:

- **Selected Session.** Displays processes for the session selected in the **Sessions** list.
- **Selected Server.** Displays all running processes for the server on which the selected session is running.
- **All Servers.** Displays all running processes for all available servers.

You can also filter the list by specifying a search criteria for one or more columns. To do so, click the magnifying glass icon (top right) and then type a desired text in one or more columns. The list is filtered as you type to match the specified criteria.

The **Tasks** drop-down list in the **Running Processes** dialog includes the following options:

- **Refresh.** Refresh the list.
- **Kill process.** Kill the selected process.
- **Go To Published Item.** Enabled when you select a process that belongs to a running published resource. Brings up the main Parallels RAS Console window and navigates to the corresponding published resource.
- **Disconnect.** Disconnect the session.
- **Log off.** Log off the session.
- **Send message.** Send a message to the session owner.
- **Remote control.** Remotely control the selected user session.

The Resources tab

The **Resources** tab in the **Sessions** category displays currently running published resources (apps and desktops).

Some of the notable columns in the list are:

- **ID:** The published resource ID (as seen in the **Publishing** category).
- **Published name:** Published resource name (as seen in the **Publishing** category).
- **User:** Session owner.
- **Session ID:** Session ID.
- **Session Host:** Session host name.
- **Source:** Session source (RDSH, VDI, RemotePC, AVD).

To perform a task on a resource, click the **Tasks** menu. Some of the tasks include:

- **Search:** Allows you to filter the list using one or more columns (e.g. User, Session ID, Session Host, etc.).
- **Show user session view:** Switches to the **Users** tab and applies a filter to display the session to which the selected resource belongs.
- **Go to published resource:** Takes you to the **Publishing** category and displays the selected resource information.
- **Show information:** Displays the resource summary info and the session information. The session information includes the same metrics as described in [Session information](#).

SSL Certificate Management

The Parallels RAS Console includes a certificate management interface that allows you to manage all of your SSL certificates in one place.

Certificates are managed on a Site level. Once a certificate is added to a Site, it can be used within the Site.

To manage certificates, in the RAS Console, navigate to **Farm > Site > Certificates**. The **Certificates** tab in the right pane displays the existing certificates. When you install Parallels RAS, the <Default> self-signed certificate is created automatically, so you will see at least this certificate in the list. The default certificate is also automatically assigned to all new RAS Secure Gateways and HALB.

The subsequent sections describe certificate management tasks in detail and provide additional certificate information and instructions.

Generating a self-signed certificate

To generate a self-signed certificate, navigate to **Farm > Site > Certificates**. Click **Tasks > Generate self-signed certificate**. In the dialog that opens, specify the following options:

- **Name:** Type a name for this certificate. This field is mandatory.
- **Description:** An optional description.
- **Usage:** Specify whether the certificate should be used for RAS Secure Gateways or HALB, or both. This selection is mandatory.

- **Key size:** The certificate key size, in bits. Here you can select from the predefined values. The default is 2048 bit, which is the minimum required length according to current industry standards.
- **Country code:** Select your country.
- **Expire in:** The certificate expiration date.
- **Full state or province:** Your state or province info.
- **City:** City name.
- **Organization:** The name of your organization.
- **Organization unit:** Organizational unit.
- **E-mail:** Your email address. This field is mandatory.
- **Common name:** The Common Name (CN), also known as the Fully Qualified Domain Name (FQDN). This field is mandatory.
- **Alternative names:** Specify one or more subject alternative names (SANs). Click the [...] icon and then add one or more DNS or IP addresses. Note that because Parallels Client for mobile devices doesn't support the SAN field, it's safest to set your common name to the name that most mobile devices will be using.

Click **Save** to generate the certificate. When done, the certificate will appear in the **Certificates** list in the RAS Console with the **Status** column indicating **Self-signed**.

To view the certificate info, right-click it and choose **Properties**. In the dialog that opens, examine the properties and then click the **View certificate info** button to view the certificate trust information, details, certification path and the certificate status. You can also view the certificate info by right-clicking it and choosing **View certificate info**.

Generating a certificate signing request (CSR)

To generate a CSR, navigate to **Farm > Site > Certificates**. Click **Tasks > Generate a certificate request**. In the dialog that opens, specify the required information. The information is exactly the same as for the [self-signed certificate described above](#). If you need an explanation, please refer to the list of options described in that section.

After entering the information, click **Generate**. Another dialog will open displaying the request. Copy and paste the request into a text editor and save the file for your

records. The dialog also allows you to import a public key at this time. You can submit the request to a certificate authority now, obtain the public key, and import it without closing the dialog, or you can do it later. If you close the dialog, the certificate will appear in the RAS Console with the **Status** column indicating **Requested**.

To submit the request to a certificate authority and import a public key:

1. If the certificate request **Properties** dialog is closed, open it by right-clicking a certificate and choosing **Properties**. In the dialog, select the **Request** tab.
2. Copy the request and paste it into the certificate authority web page (or email it, in which case you will need to come back to this dialog later).
3. Obtain the certificate file from the certificate authority.
4. Click the **Import public key** button and finalize the certificate registration by specifying the key file and the certificate file.

Let's Encrypt certificates

Requesting a Let's Encrypt Certificate

Let's Encrypt is a global Certificate Authority (CA). This organization is a non-profit and does not charge fees for their certificates. Each certificate is valid for 90 days. RAS Console allows you to issue, automatically renew and revoke Let's Encrypt certificates.

Issuing a Let's Encrypt certificate

To issue a new Let's Encrypt certificate:

1. In the RAS Console, navigate to **Farm > Certificates**.
2. Click the **[+]** button to the left of the **Tasks** drop-down menu and select **Issue Let's Encrypt certificate**.
3. Select the **I have read and accept Let's Encrypt EULA** option.
4. In the **Expiration emails** field list specify the email addresses that will receive notifications from Let's Encrypt.
5. Optionally, change the time when certificates are renewed automatically in the **Automatically renew certificates before expiration** field.
6. Click **OK**.

7. In the **Issue Let's Encrypt certificate** dialog, specify the following:
 - **Name:** Name of the certificate.
 - **Description:** Description of the certificate.
 - **Usage:** HALB and/or Secure Gateway.
 - **Key size:** Key size.
 - **Country code:** Code of your country.
 - **Full state or province:** Name of your state or province.
 - **City:** Your city.
 - **Organization:** Name of your organization.
 - **Organization unit:** Name of your organization unit.
 - **E-mail:** Email address of your organization.
 - **Common name:** Valid domain name of a publicly accessible HALB or Secure Gateway.
 - **Alternative names:** Valid domain names of a publicly accessible HALBs or Secure Gateways.
8. Click **Save**.

Renewing a Let's Encrypt certificate manually

To manually renew a Let's Encrypt certificate:

1. In the RAS Console, navigate to **Farm > Certificates**.
2. Right-click the Let's Encrypt certificate that you want to renew.
3. In the context menu, select **Control > Renew**.

Revoking a Let's Encrypt certificate

To revoke a Let's Encrypt certificate:

1. In the RAS Console, navigate to **Farm > Certificates**.
2. Right-click the Let's Encrypt certificate that you want to revoke.
3. In the context menu, select **Control > Revoke**.
4. In the **Revoke Certificate** dialog, select the reason why you want to revoke the certificate.
5. Click **Revoke**.

How Parallels RAS requests certificates from Let's Encrypt

When you create a new Let's Encrypt certificate using Parallels RAS, the following process is carried out:

1. Parallels RAS Primary Connection Broker that hosts the licensing role makes the initial request to the Let's Encrypt server to create an account.
2. Account creation confirmation is received. Parallels RAS creates a CSR and sends it to the Let's Encrypt server.
3. A list of challenges is received, and Connection Broker reads the HTTP token sent by the Let's Encrypt server.
4. Secure Gateway or HALB retrieves the tokens from the Connection Broker.
5. Once ready, Connection Broker notifies the Let's Encrypt Server.
6. Let's Encrypt starts the verification process by going to the Secure Gateway or HALB and confirming the availability of the token.
7. Challenges are completed including confirmation that the Secure Gateways or HALB can reply to the domain mentioned.
8. Assuming that the challenge is completed successfully, Parallels RAS requests a certificate.
9. Valid certificate is downloaded from the Let's Encrypt server to Connection Broker.
10. Connection Broker distributes the certificate to the Secure Gateways or HALB.

Importing a certificate

To import a certificate from a file, on the **Certificates** tab, click **Tasks > Import certificate**. In the dialog that opens, specify the following:

- **Name:** Type a name for the certificate.
- **Description:** An optional description.
- **Private key file:** Specify a file containing the private key. Click the [...] button to browse for the file.
- **Certificate file:** When you specify a private key file (above) and have a matching certificate file, it will be inserted in this field automatically. Otherwise, specify a certificate file.

- **Usage:** Specify how the certificate will be used.

Click **OK** when done. The certificate will appear in the list in the RAS Console with the **Status** column indicating **Imported**.

To view the certificate info, right-click it and choose **Properties**. In the dialog that opens, examine the properties and then click the **View certificate info** button to view the certificate trust information, details, certification path, and the certificate status. You can also view the certificate info by right-clicking it and choosing **View certificate info**.

For imported certificates, the **Properties** dialog has an additional tab **Intermediate**. If the original certificate included an intermediate certificate (in addition to the root certificate), it will be displayed here. You can paste a different intermediate certificate here if you wish.

Exporting a certificate

To export a certificate to a file, on the **Certificates** tab, click **Tasks** > Export certificate, specify a filename and click **Save**. You can later import the certificate in a different Farm or Site by clicking **Tasks** > **Import certificate** and specifying the certificate file in the **Private key file** field.

Assigning a certificate

After you add a certificate to a Site, you can assign it to a RAS Secure Gateway, HALB, Connection Broker API server, Connection Broker API data encryption, or a combination, depending on the usage type that you specified when you created the certificate (described at the [beginning](#) of this section). More on the certificate **Usage** option below.

Certificate usage

Certificate **Usage** is an option that you specify when you create a certificate. When setting this option, you can choose from the following:

- **Secure Gateway:** If selected, makes the certificate available for RAS Secure Gateways.
- **HALB:** If selected, makes the certificate available for HALB.

- **Connection Broker API server:** If selected, makes the certificate available for Connection Broker API integration.
- **Connection Broker API data encryption:** If selected, makes the certificate available for API data encryption.

You can select one of the options above or a number. For details on how to create a certificate and choose these options, please see [Generating a self-signed certificate](#) and [Generating a certificate signing request \(CSR\)](#).

When you configure SSL/TLS for a RAS Secure Gateway, HALB, or Third-Party API later, you need to specify a certificate. For the information on how to do this, please see [SSL/TLS encryption](#) and [Configuring HALB in the RAS Console](#). When you select a certificate, the following options will be available depending on how the **Usage** option is configured for a particular certificate:

- **<All matching usage>:** This is the default option, which is always available. It means that any certificate on which the **Usage** selection matches the object type will be used. For example, if you are configuring a Gateway and have a certificate that has **Usage** set to “Gateway”, it will be used. If a certificate has both Gateway and HALB usage options selected, it can also be used with the given gateway. This works the same way for HALB when you configure the LB SSL Payload. Please note that if you select this option for a Gateway or HALB, but no single matching certificate exists, you will see a warning and will have to create a certificate first.
- Other items in the **Certificates** drop-down list are individual certificates, which will or will not be present depending on the certificate’s **Usage** settings. For example, if you configure LB SSL Payload for HALB and have a certificate with the **Usage** option set to “HALB”, the certificate will appear in the drop-down list. On the other hand, certificates with **Usage** set to “Gateway” will not be listed.

As another example, if you need just one certificate, which you would like to use for all of your Gateways, you need to create a certificate and set the **Usage** option to “Gateways”. You can then configure each Gateway to use this specific certificate, or you can keep the default **<All matching usage>** selection, in which case the certificate will be picked up by a Gateway automatically.

Gateways

To assign a certificate to a RAS Secure Gateway:

1. Navigate to **Farm > Site > Secure Gateways**.

2. Right-click a gateway and choose **Properties**.
3. Select the **SSL/TLS** tab.
4. In the **Certificates** drop-down list, select the certificate that you created.
5. Click **OK**.

Note: You can also select the **<All matching usage>** option, which will use any certificate that either has the usage set to Gateway. endpoint %}

HALB

To assign a certificate to a HALB, navigate to **Farm > Site > HALB**. Assuming that your HALB is enabled and configured, and the **LB SSL Payload** option is selected, follow the instructions below:

1. Click **Configure** next to the **LB SSL Payload** option.
2. A certificate must be used when the **Mode** option is set to **SSL Offloading**.
Once again, assuming it is selected, continue to the next step.
3. Click **Configure**.
4. In the **SSL** dialog, select the certificate in the **Certificates** drop-down list.

As with gateways, you can also select the **<All matching usage>** option, which will use any certificate that has the usage set to HALB or both HALB and Gateway.

Connection Broker API Server

This usage is only needed if access to Parallels RAS is needed via API to a Third-Party Connection Broker (more information [here](#)). To assign this certificate type, navigate to **Farm > Site > Connection Brokers** and follow the instructions below:

1. Choose **Tasks > Site Defaults** (or choose the properties of an individual Connection Broker to assign to a specific Connection Broker)
2. Choose the **API** tab
3. Choose **Enable API**
4. Then select the certificate in the **Certificates** dropdown

Note: Different Connection Brokers within the same site can have a different configuration or use a shared configuration at the site level (i.e., each Connection Broker can inherit the site's default).

Connection Broker API data Encryption

This usage is only needed if access to Parallels RAS is needed via API to a Third-Party Connection Broker (more information [here](#)), and a specific certificate is needed to encrypt traffic (if no certificate is selected for Third-Party Encryption, traffic will still be encrypted using the default site certificate). To assign this certificate type, navigate to **Farm > Site > Connection Brokers** and follow the instructions below:

1. Choose the **API** tab (this tab is only visible if the Enable API option is chosen from the “Third-Party API step above -
2. Then select the certificate in the **Encryption Certificate** dropdown

It is important that Connection Brokers on the same site are set to use the **same** Encryption certificate.

Auditing certificates

All actions that you perform on certificates are audited and can be viewed later. Note that reverting certificate changes is not possible. If you need to revert to a previous state, you'll have to delete a certificate and create a new one.

To audit certificates:

1. In the RAS Console, navigate to **Farm > Site > Certificates**.
2. Click **Tasks > Settings audit**.
3. The dialog opens where you can view the history of certificate actions. Note that the **Revert** button is disabled. As noted at the beginning of this section, reverting a certificate action is not possible.
4. To view details for a particular audit entry, double-click it.

Permissions to manage certificates

Root and Power administrators always have rights to manage certificates. Custom administrators don't have them by default. To grant permissions to manage certificates to Power administrators, the **Certificates** global permission type is used.

If you are a Root or Power administrator, you can set certificate permissions as follows:

1. In the RAS Console, navigate to **Administration > Accounts**.
2. Select a Custom administrator account and click **Tasks > Properties**.
3. In the **Account Properties** dialog, click **Change Permissions**.
4. In the **Account Permissions** dialog, select a Site in the left pane and click **Change permissions** (or click the **Edit** link in the right pane).
5. In the left pane (Permission type), select **Certificates**.
6. In the right pane (Global permissions), select one or more permissions.
7. When done, close all dialogs.

A RAS administrator can also delegate his/her permissions to a custom administrator. To do so, navigate to **Farm > Site > Certificates** and click **Tasks > Delegate permissions**. In the dialog that opens, delegate permissions to a desired Custom administrator.

Upgrading from an older RAS version

When you upgrade Parallels RAS from a version prior to RAS 17.1 to a RAS 17.1 (or newer), every certificate that is used by RAS Secure Gateways and HALB is enumerated and only unique certificates are added to the **Certificates** subcategory. Gateways and HALB are then linked 1-to-1 to the certificates they were using before the upgrade.

Other actions related to an upgrade include the following:

- The **Inherit defaults** option in gateways is turned off after the upgrade.
- If a gateway is disabled during an upgrade, the Connection Broker still has the information about the certificate that the gateway uses, so the gateway is configured properly when it comes back online.
- Site defaults settings are configured to use the default self-signed certificate.
- When a new gateway is added, it is configured to use the default self-signed certificate, provided the Site defaults are not changed afterwards.

Connection and Authentication Settings

A Parallels RAS administrator has the ability to customize how users connect to Parallels RAS. This chapter describes connection and authentication settings that can be configured according to your organization requirements. It then explains how to use two-factor authentication for higher level of security.

RAS Connection Broker connection settings

RAS Connection Broker connection settings can be accessed from the **Connection** category.

Choosing authentication type

Select the **Authentication** tab. In the **Allowed authentication types** section, select one of the following options:

- **Credentials.** The user credentials are validated by the Windows system on which RAS is running. The credentials used for Windows authentication are also used to log in to an RDP session.
- **Smart Card.** Smart card authentication. Similar to Windows authentication, smart card credentials can be shared between both RAS and RDP. Hence, smart card credentials only need to be entered once. Unlike Windows authentication, the user only needs to know the smart card's PIN. The username is obtained automatically from the smart card, so the user doesn't need to provide it.

Note: Smart card authentication is not supported in Parallels Client for Linux.

- **Web (SAML).** SAML SSO authentication. For more information, see [SAML SSO Authentication](#).
- **Web + Credentials.** The same as **Web (SAML)**, but users are prompted to enter credentials when they launch a published application. To enable the **Web + Credentials** method, you must configure your IdP and RAS as described in [IdP side configuration](#) and [SP side configuration](#).

Note: The **Web + Credentials** method works only in Parallels Client for Windows and User Portal.

Note that if smart card authentication is disabled, RAS Connection Broker will not hook the Local Security Authority Subsystem Service (LSASS). Smart card authentication can be used in Parallels Client for Windows, Mac, and Linux. Please also note that smart cards cannot be used for authentication if Parallels Client is running inside an RDP session.

Smart card certificate

A valid certificate must be installed on a user device in order to use smart cards. To do so, you need to import the certificate authority root certificate into the device's keystore.

A certificate must meet the following criteria:

- The “Key Usage” field must contain digital signature.
- The “Subject Alternative Name” (SAN) field must contain a user principal name (UPN).
- The “Enhanced Key Usage” field must contain smart card logon and client authentication.

Authentication domain

To specify an authentication domain, select one of the following:

- **Specific:** Select this option and type a specific domain name.
- **All trusted domains:** If the information about users connecting to Parallels RAS is stored in different domains within a forest, select the **All Trusted Domains** option to authenticate against multiple domains.
- **Use client domain if specified:** Select this option to use the domain specified in the Parallels Client connection properties. If no domain name is specified on the client side, the authentication is performed according to the settings above.
- **Force clients to use NetBIOS credentials:** If this option is selected, the Parallels Client will replace the username with the NetBIOS username.

Note: If a certificate on your smart card does not contain a user principal name (UPN) in the “Subject Alternative Name” (SAN) field (or if it doesn't have the

“Subject Alternative Name” field at all) you have to disable the **Force clients to use NETBIOS credentials** option.

Recommendation: After changing the domain names or some other authentication related changes, click the **Clear cached session IDs** button on the **Settings** tab.

Authenticating against non domain users

In order to authenticate users sessions against users specified on a standalone machine you must enter the [workgroup_name] / [machine_name] instead of the domain name. For example if you would like to authenticate users against a list of local users on a machine called SERVER1 that is a member of the workgroup WORKGROUP, enter the following in the domain field: WORKGROUP/SERVER1.

Changing domain password

You can configure Parallels Client to use a custom URL for changing domain passwords.

To make Parallels Client use a custom URL for changing domain passwords:

1. Select **Use a custom link fro the “Change domain password” option**.
2. Add the link to the text field below.

Remote session settings

The **Settings** tab in the **Connection** category allows you to configure the following remote session options.

Declare user session as idle after

This option affects reporting statistics, whereby a session is declared idle after the amount of time specified without any activity.

FIPS 140-2 encryption

The **FIPS 140-2 encryption** property allows you to specify whether FIPS-encrypted connections are allowed or even enforced on RAS Secure Gateways.

When you allow (or enforce) the encryption, the Gateways will use the FIPS 140-2 encryption module. You can choose from the following options:

- **Disabled.** FIPS 140-2 encryption is disabled on RAS Secure Gateways.
- **Allowed.** RAS Secure Gateways accept both FIPS-encrypted and non-FIPS-encrypted connections.
- **Enforced.** RAS Secure Gateways accept FIPS-encrypted connections and will drop any non-FIPS-encrypted connection.

Note: For FIPS 140-2 encryption to work, a FIPS compliant certificate must be installed on each RAS Secure Gateway.

When you enable FIPS 140-2 encryption, the encryption status is displayed on the **Information > Site** tab in the RAS Console. Look for the **Encryption** property of a RAS Secure Gateway.

Note: If you use FIPS, the minimum allowed version of TLS is automatically set to 1.2. endhint %}

FIPS 140-2 encryption is supported in all versions of Parallels Client except for the following:

- Parallels Client for Windows installed on Windows 8.1 and earlier
- Parallels Client for Linux downloaded from the IGEL App Portal

Note: If you need to use FIPS in IGEL OS, you can download the tarball version of Parallels Client for Linux from the [Parallels site](#). Note that installing applications outside of the IGEL App Portal is not endorsed by IGEL and Parallels will not be able to provide support for this case. The tarball needs to be extracted to a persistent storage area, or it will be erased upon system reboot.

- Parallels Client for Android
- Parallels Client for iOS
- Parallels User Portal

Note: Parallels Client for ARM64 does not support FIPS 140-2.

Please also note that when FIPS 140-2 encryption is enforced, it is enforced all users in a given Farm. If there's a necessity to force FIPS for one user group and not forced for another, a new Farm must be deployed for this purpose.

Automatically log out client idle connection after

Specifies the time period after which an idle client connection should be logged out. Once the connection is logged out, the user is disconnected from Parallels RAS and is presented with the **Connections** dialog in Parallels Client as a way to notify them that they were logged out. They can use the dialog to log back on if desired.

Log out idle session mode

Specifies when the automatic logout of an idle client connection is triggered:

- With no active session, the client connection is logged out only if the user does not have any active published applications or desktop sessions running.
- With or without active session, the client connection is logged out regardless of whether the user still has active published applications or desktop sessions.

Note: In the Parallels RAS Web Client, automatic logout behavior depends on how resources are opened:

- Applications opened in the same browser tab as the User Portal launcher do not trigger a client logout.
- Applications opened in a separate browser tab and published desktops (which always open in a separate tab) allow the User Portal launcher session to log out independently.

Cached authentication token timeout

Specify the amount of time that a session is cached for (higher amount of time reduces AD transactions).

Clear cached authentication tokens (a button)

Clears all cached session information.

Logon hours settings

Note: This feature is not supported on Parallels Clients earlier than version 19 and Parallels Client for Chrome. Creating a logon hours rule restricts the ability to connect to published resources (within a site) using any of these clients.

Logon hours restrictions provide an ability to restrict user access to published resources during specified time frames using flexible expression-based rules.

Prerequisites

Time zone redirection is required to be set on the server in order for the feature to work as intended.

To enable group policy setting Allow time zone redirection:

1. On the Active Directory server, open the Group Policy Management Console.
2. Expand your domain and Group Policy Objects.
3. Right-click the GPO that you created for the group policy settings and select **Edit**.
4. In the **Group Policy Management Editor**, navigate to **Computer Configuration > Policies > Administrative Templates > Windows Components > Remote Desktop Services > Remote Desktop Session Host > Device and Resource Redirection**.
5. Enable the setting **Allow time zone redirection**.


Adding a logon hours rule

To add a new logon hours rule:

1. In the RAS Console, navigate to **Connection** and select the **Logon hours** tab.
2. Click **Tasks > Add** (or click the **[+]** icon).
3. In the **Name** field, specify the name of the rule in the .
4. In the **Description** field, specify the description of the rule

5. In the **Criteria** section, specify criteria for the rule. You will find the following controls:

- **(+)**: adds a new criteria. If you want to match a Secure Gateway, a client device name, a client device operating system, an IP address or a hardware ID, click **(+)**. In the context menu that appears, select the type of an object that you want to match and add the specific objects in the dialog that appears. The new criteria appears on the next line.
- **(X)**: Deletes a specific object from matching. For example, you want to delete IP address 198.51.100.1 from matching, click **(X)** next to it. This control appears when at least one object is added. If all objects in a criteria are deleted, the criteria is removed.
- **is** and **is not**: specifies whether the logon hours rule must be applied when a user connection matches the criteria. Click on the link to switch between the two options. This control appears when at least one object is added.
- **configure**: edits the list of objects to be matched. Click this link to add or delete new objects. Note that for the first criteria (**User or group**) this link is called **everyone**. It will change to **configure** once you specify objects for this criteria.



6. In the **Logon hours** specify the hours when users are permitted to log on. To deny logon during a certain day or period of time, select that day or time and click the **Logon denied** button that is located to the right of the table.

7. Click **OK**.

8. Click **Apply**.

Note: If no logon hours rules are specified, access to published resources is not restricted. If rules are specified, but the user connection does not match any of them, the user is denied access.

You can also specify the following settings for a logon hours rule:

- **Do not allow Parallels Client to connect outside of allowed logon hours:** If selected, a Parallels Client is not allowed to connect to resources published on the site.
- **Disconnect user session if the time has elapsed:** If selected, shows users a notification that their sessions are going to be disconnected. After selecting this option, you can specify the settings below:
 - **Notify user before disconnect:** Time when Parallels RAS notifies the user before the client is disconnected from the Farm.
 - **Allow user to extend session time:** If selected, allows user to extend the session.

To specify these settings:

1. In the RAS Console, navigate to **Connection** and select the **Logon hours** tab.
2. Select the rule that you want to configure.
3. Click the gear icon to the left of the **Task** menu. The **Options** dialog opens. From here, select the options that you want.

Restricting access by Parallels Client type and build number

You can specify a minimum requirement for the Parallels Client type and version number in order for it to connect to the Parallels RAS Farm or to list published resources. In addition, you can set the Parallels Client security patch level (described later in this section).

To specify Parallels Client requirements:

1. In the RAS Console, select the **Connection** category and click the **Allowed Devices** tab.
2. The **Allow only clients with latest security patches** option specifies the Parallels Client security patch level. If the option is selected, only clients with

the latest security patches applied will be allowed to connect to Parallels RAS. This option must normally be selected to protect your environment from vulnerabilities. You should only clear it if you must use an older version of Parallels Client with no security patches installed.

3. In the **Mode** drop-down list, select from the following options:
 - **Allow all clients to connect to the system.** No restrictions. All Parallels Client types and versions are allowed full access.
 - **Allow only the selected clients to connect to the system.** Allows you to specify Parallels Client types and versions that are allowed to connect to the Parallels RAS Farm. Select the desired Parallels Client types in the **Clients** list. To set the **Minimum build** value, right-click the client type and choose **Edit**. Type the version number directly in the **Minimum build** column.
 - **Allow only the selected clients to list the published items.** Allows you to specify Parallels Client types and versions that can list published resources. Compared to the option above, this one does not restrict Parallels Clients connecting to Parallels RAS. Select this option and then select the desired Parallels Client types in the **Clients** list. To set the **Minimum build** value, right-click the client type and then click **Edit** in the context menu. Type the version number directly in the **Minimum build** column.

If a restriction is configured and a Parallels Client is excluded from the list, the user running it will receive a corresponding error message and will be advised to contact the system administrator.

Multi-factor authentication

Parallels RAS allows you to use multi-factor authentication for access control. When multi-factor authentication is used, users will have to authenticate through two successive stages to get the application list. While the first level will always use native authentication (Active Directory / LDAP), the second level can use one of the following solutions:

- [RADIUS](#)
 - Azure MFA (RADIUS)
 - Duo (RADIUS)
 - FortiAuthenticator (RADIUS)
 - TekRADIUS
 - RADIUS

- [TOTP](#)
 - Google Authenticator
 - Microsoft Authenticator
 - TOTP (Time-based one-time password)
- [Email OTP](#)
- [Deepnet](#)
- [SafeNet](#)

Multi-factor authentication is more secure because instead of using a standard user name and password, it uses a static user name and a one-time password generated by a token.

Learn how to add an MFA provider in the [Adding an MFA provider](#) section.

See also [Configuring MFA rules](#).

Adding an MFA provider

To add an MFA provider:


1. In the RAS Console, navigate to **Connection** and select the **Multi-Factor authentication** tab.
2. Click **Tasks > Add** (or click the **[+]** icon).
3. Select your MFA provider. A wizard will open.
4. In the Wizard window, specify the following parameters:
 - **Name:** Name of the provider.
 - **Description:** Description of the provider.
 - In the **Themes** table select the Theme(s) that will use this MFA provider.
5. Click **Next**.
6. Do one of the following:
 - If you use RADIUS, configure the setting as described in [Connection](#) and click Finish.
 - If you are using a TOTP provider other than Google Authenticator, configure the setting as described in [Configuring TOTP](#).
 - If you use email to send OTPs, configure the setting as described in [Configuring email OTP](#).
 - If you use Deepnet DualShield, configure the setting as described in [Configuring Parallels RAS to use the DualShield Authentication Platform](#). For information about configuring DualShield Authentication

Platform, see section [Configuring DualShield 5.6+ Authentication Platform](#).

- If you use SafeNet, configure the setting as described in [Configuring SafeNet](#).
- If you use Google Authenticator, configure the setting as described in [Configuring Google Authenticator](#).

Using RADIUS

The below diagram shows the double hop perimeter network scenario with RAS Connection Broker connected to a RADIUS server (RADIUS is located in Intranet but it can be placed in DMZ).



To configure RADIUS properties:

1. In the Parallels RAS Console, navigate to **Connection > Multi-factor authentication**.
2. Double-click the MFA provider that you want to configure.

Read on to learn how to configure RADIUS provider settings.

Connection

The **Connection** tab lets you specify the following options:

- **Display name:** Specify the name of the OTP connection type that will be displayed on the Logon screen on the client side. This should be the name that your users will clearly understand.

- **Primary server** and **Secondary server**: These two fields allow you to specify one or two RADIUS servers to include in the configuration. Specifying two servers gives you an option to configure high availability for RADIUS hosts (see below). Specify a server by entering its hostname or IP address or click the [...] button to select a server via Active Directory.
- When two RADIUS servers are specified, select one of the following high availability modes from the **HA mode** drop-down list: **Active-active (parallel)** means the command is sent to both servers simultaneously, the first to reply will be used; **Active-passive (failover)** means failover and timeout are doubled, Parallels RAS will wait for both hosts to reply.
- **HA mode**: See **Primary server** and **Secondary server** above. If only the **Primary server** is specified, this field is disabled.
- **Port**: Enter the port number for the RADIUS Server. Click the **Default** button to use the default value.
- **Timeout**: Specify the packet timeout in seconds.
- **Retries**: Specify the number of retries when attempting to establish a connection.
- **Secret key**: Type the secret key.
- **Password encoding**: Choose from PAP (Password Authentication Protocol) or CHAP (Challenge Handshake Authentication Protocol), according to the setting specified in your RADIUS server. Both options are available in UTF-8 format to support passwords with special characters (E.g., umlaut characters or §).

Click the **Check connection** button to validate the connection. If the connection is configured correctly, you will see a confirmation message.

Specify additional properties as required:

- **User Prompt**: Specify the text that the user will see when prompted with an OTP dialog.
- **Forward username only to RADIUS server**: Select this option if needed.
- **Forward the first password to Windows authentication provider**: Select this option to avoid a prompt to enter the password twice (RADIUS and Windows AD).

[Note: For the Azure MFA server, this option is always enabled and cannot be turned off.³](#)

- **Enforce Message-Authenticator validation from RADIUS server**: Select this option if needed.

- Please also read a note at the bottom of the dialog (if available) suggesting certain setting specifics for the selected RADIUS solution.

If your RADIUS solution requires configuring attributes, click the **Attributes** tab and then click **Add**. In the dialog that opens, choose a desired preconfigured vendor and attribute:

- In the **Vendor** drop-down list, select a vendor.
- In the **Attribute** list, select a vendor attribute.
- In the **Value** field, enter a value for the selected attribute type (numeric, string, IP address, date, etc).

In certain scenarios you may need to add vendors and attributes that are not listed in this dialog. For the information about how to add vendors and attributes, please see the following KB article: <https://kb.parallels.com/en/125576>.

Click **OK** and then click **OK** again .to close all dialogs.

Automation

The **Automation** tab in the RADIUS **Properties** dialog allows you customize the OTP experience for Parallels Client users by configuring security verification methods and custom commands to be sent to a RADIUS server during the MFA login process. Different security verification methods can be assigned priority and configured to be automatically used.

With this functionality configured, users can choose their preferred security verification method from a predefined and configurable list including Push notification, Phone Callback, SMS, Email, and Custom. The methods appear as clickable icons on the OTP dialog in Parallels Client. When a user clicks an icon, a command is sent to the RADIUS server and the corresponding verification methods is used.

To configure a verification method (also called “actions” here and in the Parallels RAS Console), on the **Automation** tab, click **Tasks > Add**. In the **Add Action** dialog, specify the following properties:

- **Enable Action:** Enables or disables the action.
- **Title:** The text that will appear on the clickable icon in Parallels Client (e.g. “Push”).
- **Command:** The OTP command to be used when the action icon is clicked in Parallels Client. Consult your MFA provider for command specifications.

- **Description:** A description that will appear on the user's screen as a balloon when the mouse pointer hovers over the action icon.
- **Action message:** A message to show to the user in the connection progress box.
- **Select an image:** Select an image from the provided gallery. The image is used as the action icon in the OTP dialog in Parallels Client.

When done, click **OK** to save the action. Repeat the steps above for other actions.

You can move the actions on the **Automation** tab up or down the list. This dictates in which order the action icons will be displayed in Parallels Client.

Optionally, you can configure automation conditions. This allows you to control which verification method (action) needs to be presented to which user. To start configuring automation conditions, select **Enable Automation Conditions**. A new tab, Automation Conditions, will become available. Open that tab for further configuration.

Autosend

There's one more option that you can configure for an action. It is called **Autosend**. The option can be enabled for one action only, making it a default action, which will be used automatically without user interaction.

To enable the **Autosend** option, select an action on the **Automation** tab and click **Tasks > Autosend**. To disable the option, click the same menu again. If you enable **Autosend** for a different action, it will be automatically disabled for the previous action.

There are two possible ways to make an action execute automatically in Parallels Client:

- Client is receiving the action icon configuration for the first time and one of the actions has **Autosend** enabled.
- Enabling the **Remember last method used** option in **Policies > Session > Connection > Multifactor authentication**. When the option is enabled, and Parallel Client receives the policy, the last method successfully used by the user will become the default automatic method.

Parallels Client

When the user logs in to Parallels RAS via MFA, the OTP dialog is shown in Parallels Client with the actions icons positioned above the OTP field. The user clicks an icon and the authentication is carried out according to the predefined action. For example, if the user clicks the “Push” icon, a push notification is sent to the user mobile device where they can simply tap “Approve”. Or there could be a “Text me” icon, in which case a text is sent to the user mobile phone with a one-time password. If one of the actions has the **Autosend** option enabled, then this action is used automatically.

If a user always uses the same authentication method, they can make it the default one. To do so, the user enables the **Remember last method used** option in the **MFA authentication** section of the connection properties. Depending on the platform, the option can be found at the following locations:

- Parallels Client for Windows / Linux: **Connection Advanced Settings > MFA authentication**
- Parallels Client for Mac: **Advanced > MFA authentication**
- Parallels Client for Chrome: **Advanced Settings**
- Web Client: **Settings**
- Parallels Client for iOS: **Connection Settings > MFA authentication**
- Parallels Client for Android: **Settings > MFA authentication**

As was already mentioned above, the **Remember last method used** can also be configured in Client Policies in the RAS Console. The option is enabled by default.

Automation Conditions

The **Automation Conditions** tab in the RADIUS **Properties** dialog allows you to customize the OTP experience by assigning specific security verification methods (actions) to particular users. With this functionality configured, you have fine-grained control over which user is presented with which action in a single MFA object. If you do not see the Automation Conditions tab, make sure to enable it by selecting **Enable Automation Conditions** on the **Automation** tab.

To configure an Automation Condition, click **Tasks > Add**. In the **Add Condition** dialog, specify the following properties:

- On the left-hand side, click the plus-sign and enter the **username pattern** to apply the action to.

- On the right-hand side, click the plus sign and select the **action(s)** you want to assign to users matching the username pattern.

Repeat these steps for every user you want to apply a specific set of actions to and click **OK** to confirm.

Advanced

The **Advanced** tab lets you specify the error messages sent by the RADIUS server that will not be shown by Parallels Client. This can be useful if an error message is confusing for the user or disrupts user experience.

By default, the “New SMS passcodes sent.” is added to the list of ignored messages for DUO Radius. This is done to make authentication via SMS easier for the user. It’s not recommended to remove this message from the list of ignored messages.

To add a new message to the list of ignored messages:

1. On the **Advanced** tab, **Tasks > Add** (or click the **[+]** icon).
2. Type the exact text of the error message you want to be ignored. Messages are not case sensitive. Please note that you have to specify only the text sent by the RADIUS server. For example, if Parallels Client shows an error that reads “Code [01/00000003] Logon using RADIUS failed. Error: New SMS passcodes sent.”, you need to add “New SMS passcodes sent.” to the list.

Configuring Azure MFA

Before reading this section, please read the following important note.

Note: As of July 1, 2019, Microsoft will no longer offer MFA Server for new deployments. New customers who would like to require multi-factor authentication from their users should use cloud-based Azure Multi-Factor Authentication. Existing customers who have activated MFA Server prior to July 1 will be able to download the latest version, future updates, and generate activation credentials as usual: <https://docs.microsoft.com/en-us/azure/active-directory/authentication/howto-mfaserver-deploy>.

For new deployments, it is recommended to use Azure NPS Extension <https://>

docs.microsoft.com/en-us/azure/active-directory/authentication/howto-mfa-nps-extension or Azure MFA Service along with SAML configuration in RAS. endhint %}

Configure Azure MFA

Depending on the user location, there are four scenarios for the cloud MFA service:


User location	MFA in the cloud	MFA Server
Microsoft Entra ID	Yes	
Microsoft Entra ID and on-premises AD using federation with AD FS (is required for SSO)	Yes	Yes
Microsoft Entra ID and on-premises AD using DirSync, Azure AD Sync, Azure AD Connect - no password sync	Yes	Yes
Microsoft Entra ID and on-premises AD using DirSync, Azure AD Sync, Azure AD Connect - with password sync	Yes	
On-premises Active Directory		Yes

An Azure account with Global Administrator role is required to download and activate MFA Server. Syncing with Microsoft Entra ID (via AD Connect) or a custom DNS domain aren't required to setup an MFA Server which runs exclusively on-premises.

Users need to be imported into MFA Server and be configured for MFA authentication.

Parallels RAS authenticates users with MFA Server using the RADIUS second level authentication provider. MFA Server thus needs to be configured to allow RADIUS client connections from the RAS server.

The authentication process goes through the following stages:



In stage 2 the user can be authenticated using either RADIUS or Windows AD. A prompt to enter the credentials twice (in stage 1 and 6) is avoided by enabling the option to forward the password.

Configuring Duo

For instructions on how to configure Parallels RAS with Duo RADIUS, please read the following Parallels KB article: <https://kb.parallels.com/124429>.

Using TOTP

This section explains how to integrate TOTP MFA providers with Parallels RAS.

Configuring TOTP

To configure TOTP settings:

1. Specify the following:

- **Display Name:** The default name here is TOTP. The name will appear on the registration dialog in Parallels Client in the following sentence: “Install TOTP app on your iOS or Android device”. If you change the name, the sentence will contain the name you specify, such as “Install <new-name> app on your iOS or Android device”.
- **User Prompt:** Specify the text that the user will see when prompted with an OTP dialog.
- **Account name in authenticator:** By default, the Site and Farm names are shown. Alternatively, you can choose to display the Display Name or the Theme Name instead. This helps end users clearly identify which environment they’re working in—especially when they have access to multiple environments.

On the next page of the wizard, specify:

- The **User enrollment** section allows you to limit user enrollment if needed. You can allow all users to enroll without limitations (the **Allow** option), allow enrollment until the specified date and time (**Allow until**), or completely disable enrollment (the **Do not allow** option). If enrollment is disabled due to an expired time frame or because the **Do not allow option** is selected, a user trying to log in will see an error message saying that enrollment is disabled and advising the user to contact the system administrator. When you restrict or disable enrollment, Google Authenticator or other TOTP providers can still be used, but with added security, which would not allow further user enrollment. This is a security measure to mitigate users with compromised credentials from enrolling in MFA.
- **Show information to unenrolled users:** Select whether unenrolled users can see the **The user name or password is incorrect** error when they enter incorrect credentials:
 - **Never (most secure):** Unenrolled users see a TOTP prompt instead of the error.
 - **If enrollment is allowed:** Unenrolled users see the error if user enrollment is allowed. Otherwise, they see a TOTP prompt.
 - **Always:** Unenrolled users always see the error.

- The **Authentication** section allows you to configure TOTP tolerance. When using Time-based One-Time Password (TOTP), it is required to have the time synchronized between the RAS Connection Broker and client devices. The synchronization must be performed against a global NTP server (e.g., `time.google.com`). Using the **TOTP tolerance** drop-down list, you can select a time difference that should be tolerated while performing authentication. Expand the drop-down list and select one of the predefined values (number of seconds).

[Note: Changing time tolerance should be used with caution, as it has security implications since the time validity of a security token can be increased, thus creating a wider time window for potential misuse.](#)⁴

[Note: When using TOTP providers, it is required to have both Connection Brokers and client devices time synchronized with a global NTP server \(e.g., `time.google.com`\). Adding TOTP tolerance increases the one-time password validity, which might have security implications.](#)⁵

- The **Reset User(s)** field in the **User management** section is used to reset the token that a user receives when they log in to Parallels RAS for the first time using the TOTP provider. If you reset a user, they'll have to go through the registration procedure again (for instructions on doing this for Google Authenticator, see [Using Google Authenticator in Parallels Client](#)). You can search for specific users, reset all users, or import the list of users from a CSV file.

2. Click **Finish**.

Please also note that the TOTP available time is calculated as the default 30 seconds + x amount of seconds in the past + x amount of seconds in the future.

To configure Google Authenticator settings:

1. Specify the following:

- **Display Name:** The default name here is Google Authenticator. The name will appear on the registration dialog in Parallels Client in the following sentence, "Install Google Authenticator app on your iOS or Android device". If you change the name, the sentence will contain the name you specify, such as "Install <new-name> app on your iOS or Android device". Technically, you can use any authenticator app (hence

the ability to change the name), but at the time of this writing only the Google Authenticator app is officially supported.

- **User Prompt:** Specify the text that the user will see when prompted with an OTP dialog.
- The **User enrollment** section allows you to limit user enrollment via Google Authenticator if needed. You can allow all users to enroll without limitations (the **Allow** option), allow enrollment until the specified date and time (**Allow until**), or completely disable enrollment (the **Do not allow** option). If enrollment is disabled due to expired time frame or because the **Do not allow option** is selected, a user trying to log in will see an error message saying that enrollment is disabled and advising the user to contact the system administrator. When you restrict or disable enrollment, Google authenticator or other TOTP provider can still be used, but with added security which would not allow further user enrollment. This is a security measure to mitigate users with compromised credentials to enroll in MFA.
- **Show information to unenrolled users:** Select whether unenrolled users can see the **The user name or password is incorrect** error when they enter incorrect credentials:
 - **Never (most secure):** Unenrolled users see a TOTP prompt instead of the error.
 - **If enrollment is allowed:** Unenrolled users see the error if user enrollment is allowed. Otherwise, they see a TOTP prompt.
 - **Always:** Unenrolled users always see the error.
- The **Authentication** section allows you to configure TOTP tolerance. When using Time-based One-Time Password (TOTP), it is required to have the time synchronized between the RAS Connection Broker and client devices. The synchronization must be performed against a global NTP server (e.g. time.google.com). Using the **TOTP tolerance** drop-down list, you can select a time difference that should be tolerated while performing authentication. Expand the drop-down list and select one of the predefined values (number of seconds). Note that changing time tolerance should be used with caution as it has security implications since the time validity of a security token can be increased, thus a wider time window for potential misuse.

Note: When using Time-based One-time Passwords (TOTP) providers, it is required to have both Connection Brokers and client devices time synchronized with a global NTP server (e.g. time.google.com). Adding TOTP tolerance increases the one-time password validity, which might have security implications.

- The **Reset User(s)** field in the **User management** section is used to reset the token that a user receives when they log in to Parallels RAS for the first time using Google Authenticator. If you reset a user, they'll have to go through the registration procedure again (see **Using Google Authenticator in Parallels Client** below). You can search for specific users, reset all users, or import the list of users from a CSV file.

2. Click **Finish**.

Please also note that the TOTP available time is calculated as the default 30 seconds + x amount of seconds in the past + x amount of second in the future.

Using Google Authenticator in Parallels Client

Important: To use Google Authenticator or other TOTP provider, the time on a user device must be in sync with the time set on the RAS Connection Broker server. Otherwise, Google authentication will fail.

Google Authenticator is supported in Parallels Client running on all supported platforms, including mobile, desktop, and Web.

To use Google Authenticator, a user needs to install the Authenticator app on their iOS or Android device. Simply visit Google Play or App Store and install the app. Once the Authenticator app is installed, the user is ready to connect to Parallels RAS using two-factor authentication.

To connect to Parallels RAS:

1. The user opens Parallels Client or Web Client and logs in using his/her credentials.
2. The multi-factor authentication dialog opens displaying a barcode (also known as QR code) and a secret key.

3. The user opens the Google Authenticator app on their mobile device:
 - If this is the first time they use it, they tap **Begin** and then tap **Scan a barcode**.
 - If a user already has another account in Google Authenticator, they tap the plus-sign icon and choose **Scan a barcode**.
4. The user then scans the barcode displayed in the Parallels Client login dialog.

If scanning doesn't work for any reason, the user goes back in the app, chooses **Enter a provided key** and then enters the account name and the key displayed in the Parallels Client login dialog.
5. The user then taps **Add account** in the app, which will create an account and display a one time password.
6. The user goes back to Parallels Client, clicks **Next** and enters the one time password in the **OTP** field.

On every subsequent logon, the user will only have to type their credentials (or nothing at all if the **Save password** options was selected) and enter a one time password obtained from the Google Authenticator app (the app will continually generate a new password). If the RAS administrator resets a user (see the **Reset Users(s)** field description at the beginning of this section), the user will have to repeat the registration procedure described above.

Configuring Microsoft Authenticator

See [Configuring TOTP](#).

Configuring email OTP

To configure sending OTPs via email:

1. Configure an SMTP server as described in [Configuring SMTP server connection for event notifications](#).
2. Specify the following:
 - **Name:** The name that will appear in RAS Console.
 - **(Optional) Description:** The description of MFA.
 - **Themes:** The Themes that use the MFA.

- **Display name:** The name that will appear in Parallels Client.
- **OTP Length:** The length of an OTP. Can be between 4 and 20 numbers.
- **OTP Validity:** The time period when an OTP is valid. Can be between 30 and 240 seconds.
- **User Prompt:** Specify the text the user will see when prompted with an OTP dialog.
- **E-mail subject:** The subject of an email containing an OTP.
- **E-mail content:** The content of an email containing an OTP.
- **Allow users to enroll using external emails:** Select this option if you want users to enroll using external email addresses. You can store external emails in RAS Storage or an AD Attribute. If you want to store emails in an Active Directory Custom attribute, you must specify the name of the attribute in the field **AD Custom Attribute**. You can make sure that you have the permission necessary for storing email addresses in an AD attribute by clicking **Validate**.
- The **User enrollment** section allows you to limit user enrollment if needed. You can allow all users to enroll without limitations (the **Allow** option), allow enrollment until the specified date and time (**Allow until**), or completely disable enrollment (the **Do not allow** option). If enrollment is disabled due to an expired time frame or because the **Do not allow option** is selected, a user trying to log in will see an error message saying that enrollment is disabled and advising the user to contact the system administrator. When you restrict or disable enrollment, Google authenticator or other TOTP provider can still be used, but with added security which would not allow further user enrollment. This is a security measure to mitigate users with compromised credentials to enroll in MFA.
- **Show information to unenrolled users:** Select whether unenrolled users can see the **The user name or password is incorrect** error when they enter incorrect credentials:
 - **Never (most secure):** Unenrolled users see a TOTP prompt instead of the error.
 - **If enrollment is allowed:** Unenrolled users see the error if user enrollment is allowed. Otherwise, they see a TOTP prompt.
 - **Always:** Unenrolled users always see the error.

Using Deepnet DualShield

This section explains how to integrate Deepnet DualShield Authentication Platform 5.6 or higher with Parallels RAS.

You may also read the following documentation on DualShield Authentication Platform:

- DualShield Authentication Platform – Installation Guide
- DualShield Authentication Platform – Quick Start Guide
- DualShield Authentication Platform – Administration Guide

Supported tokens

The following is the list tokens supported by Parallels RAS:


- MobileID (FlashID is not integrated with MobileID)
- QuickID
- GridID
- SafeID
- SecureID (RSA)
- DigiPass (Vasco)

If using hardware tokens such as SafeID the token information must first the XML file provided. Click on 'Import' and browse for the XML file provided. After the XML file has been imported each hardware token must be assigned to a user.

Configuring DualShield 5.6+ Authentication Platform

After following all the specified steps in “DualShield Authentication Platform – installation Guide” a URP is automatically opened in your internet browser (<http://localhost:8073>) which allows you to logon to the Management Console of DualShield.

Login in to the DualShield Management Console with the default credentials (User: sa, Password: sa). You will be prompted to change the default password.




Applications are set to provide a connection to realm, as the realm contains domains of users who will be allowed the access to the application.

Realm is set for multiple domain users to be able to access the same application.


You need to create an Application which Parallels RAS will communicate with. Click on **Authentication > Application Wizard** and enter the information shown below and press **Next**.

Specify the LDAP Server settings as shown below and press **Finish**.


After you have configured the application you need to configure an Email or SMS gateway which are used by DualShield server to communicate with the end user. In this document we will be using an Email gateway. Select Gateways from the Configuration menu.



Configure your email gateway.




Click **Edit** to enter your SMTP server information



Configuring Parallels RAS to use the DualShield Authentication Platform

To configure Deepnet DualShield settings:



1. Specify the following:
 - **Server:** Hostname of the Deepnet server.
 - **Enable SSL:** Whether to use SSL when connecting to the Deepnet server.
 - **Port:** Port used for connection to the Deepnet server.
 - **Agent:** Agent name that will be used during registration.
2. Click the **Check Connection** button to test that the authentication server can be reached and to verify that the RAS Console is registered as a DualShield agent. If you see the “DeepNet server not valid” message, it could be due to the following:
 - The specified server information is incorrect
 - You need to allow auto registration of the Parallels components as a DualShield agent.
3. If you need to allow auto registration of the Parallels components as a DualShield agent, do the following:
 1. Go back to the DualShield Management Console and select **Agents** from the **Authentication** menu as shown below.
 2. Select **Auto Registration**.
 3. Select the **Enabled** option and set the date range.



_DualShield_10.png

Once the Agent Auto Registration is set, go back to the RAS Console and select **Yes**. You should see a message that the Dual Shield agent has been successfully registered.

Please note that all RAS Connection Brokers must be registered with Deepnet DualShield server. If you are using secondary Connection Brokers, you need to close all open windows until you can press **Apply** in the RAS Console. This will inform all the agents to self-register as DualShield agents.



4. Go back to the RAS Console and click **Next**.

5. Specify the following:

- **Application:** Name of the Application created in [Configuring DualShield 5.6+ Authentication Platform](#).
- **Default domain:** Domain that will be used if the domain was not specified by the user, in the Theme properties or in the Connection settings.

6. In the **Mode** drop-down list, and select how you want your users to be authenticated:

- **Mandatory for all users** means that every user using the system must log in using two-factor authentication.
- **Create token for Domain Authenticated Users** will allow Parallels RAS to automatically create software tokens for Domain Authenticated Users. Choose a token type from the drop-down list. Note that this option only works with software tokens, such as QuickID and MobileID

- **Use only for users with a DualShield account** will allow users that do not have a DualShield account to use the system without have to login using two-factor authentication.
7. In the **Allow channels** section, select the channels that will be used to send OTPs to users.
 8. Click **Finish**.

Connect to a RAS Farm

Parallels Client

Once DualShield has been enabled the users will have two-factor authentication. If using software tokens such as QuickID the administrator does not have to create a token for each user. RAS Connection Broker will automatically create the token when the user tries to log in for the first time.

When a user tries to access a RAS Connection from Parallels Client, they are first prompted for the Windows username and password. If the credentials are accepted, RAS Connection Broker will communicate with the DualShield server to create a unique token for that user.

If using MobileID or QuickID, an email about where to download the appropriate software will be sent to the user.

If using QuickID tokens, the application will ask for a One-Time Password which is sent by e-mail or SMS.

When asked for OTP, enter the One-Time Password to log in to the Parallels ApplicationServer XG Gateway.

Using SafeNet

SafeNet Token Management System provides a high-value of protection via secure tokens which makes it a perfect tool for second-level authentication in Parallels RAS.

Configuring SafeNet

To configure SafeNet settings:

1. In the **Connection** section, enter the valid URL into the **OTP Service URL** field. To verify that the connection with the OTP Service can be established, click the **Check connection** button.

Note: RAS Connection Broker communicates with the SafeNet Token Management System Server. It is highly recommended to have this behind a firewall for security reasons.

1. Click the **Authentication** tab.
2. In the **Mode** drop-down list, select how you want your users to be authenticated.

The available modes are:

- **Mandatory for all users:** every user using the system must login using two-factor authentication.
- **Create token for Domain Authenticated Users:** Allows Parallels RAS to automatically create software tokens for Domain Authenticated Users. Choose a token type from the drop-down list. Note that this option only works with software tokens.
- **Use only for users with a SafeNet account:** Allows users that do not have a SafeNet account to use the system without having to login using two-factor authentication.

1. In the **TMS Web API URL** field, enter the location of the SafeNet API URL.
2. In the **User Repository** field, enter the user repository destination.
3. Click **Finish**.

Parallels Client

In **Parallels Client — New Account Info** dialog:

1. Enter any four digits in the **OTP PIN** number field (these digits will be required further on in the process).
2. Enter your email address and then click on **OK**.

3. Log into your email account and retrieve the email containing the information you will need to activate your SafeNet authentication. An example of this email is shown below.

Activation Key: YZQHoczZWw3cBCNo

Token Serial: 4F214C507612A26A

Download MobilePASS client from: <http://localhost:80/TMSService/ClientDownload/MobilePASSWin.exe>

**Login with domain credentials.*

**Place the attached seed file in the same folder as the MobilePASS client.*

Enter the One-Time Password to log into the RD Session Host Connection.

Application PIN: 4089

4. Download the MobilePASS client from the URL provided in the email.
5. Enter the Activation Key found in the SafeNet email.
6. Next, input the application PIN found in the email into the **MobilePASS PIN** field.
7. Click **Generate** to generate the eToken number and then click **Copy**.
8. Combine the OTP PIN and eToken in this order: OTP + eToken.
9. Enter this value into the Parallels Client and click **OK** to log in.

Configuring MFA rules

Multi-factor authentication (MFA) can be enabled or disabled for all user connections, but you can configure more complex rules for specific connections. This functionality allows you to enable or disable MFA for the same user, depending on where the user is connecting from and from which device. Each MFA provider has one rule that consists of one or several criteria for matching against user connections. In turn, each criteria consists of one or several specific objects that can be matched.

You can match the following objects:

- User, a group the user belongs to, or the computer the user connects from.
- Secure Gateway the user connects to.
- Client device name.
- Client device operating system.
- IP address.
- Hardware ID. The format of a hardware ID depends on the operating system of the client.

Notice the following about the rules:


- Criteria and objects are connected by the OR operator. For example, if a rule has a criteria that matches certain IP addresses and a criteria that matches client device operating systems, the rule will be applied when a user connection matches one of the IP addresses OR one of the client operating systems.

To configure a rule:

1. In the RAS Console, navigate to **Connection** and select the **Multi-Factor authentication** tab.
2. Double-click on the provider you want to create the rule for.
3. Select the **Restrictions** tab.
4. Specify criteria for the rule. You will find the following controls:
 - **Enable MFA if** and **Disable MFA if**: specifies whether the MFA provider must be enabled when a user connection matches all the criteria. Click on the link to switch between the two options.
 - **(+)**: adds a new criteria. If you want to match a Secure Gateway, a client device name, a client device operating system, an IP address or a hardware ID, click **(+)**. In the context menu that appears, select the type of an object that you want to match and add the specific objects in the dialog that appears. The new criteria appears on the next line.
 - **(X)**: Deletes a specific object from matching. For example, you want to delete IP address 198.51.100.1 from matching, click **(X)** next to it. This control appears when at least one object is added. If all objects in a criteria are deleted, the criteria is removed.
 - **is** and **is not**: specifies whether the MFA provider must be enabled when a user connection matches the criteria. Click on the link to switch

between the two options. This control appears when at least one object is added.

- **configure**: edits the list of objects to be matched. Click this link to add or delete new objects. Note that for the first criteria (**User or group**) this link is called **everyone**. It will change to **configure** once you specify objects for this criteria.



Allowing users to change domain password

Users can change their domain password directly from Parallels Client. In some cases, a user may be forced to change their domain password (e.g. when the password is about to expire). When changing a password, a username must be supplied in the UPN format (e.g. user@domain.com).

Passing the domain name to Parallels Client

Since users may not know the name of their domain, you can configure Parallels RAS to pass it to the client side automatically, so users don't have to enter it.

The domain name may be specified in the RAS Console in the following locations:

- The **Connection > Authentication** tab. The tab page is [described earlier in this section](#). To force a domain name to the client side, select the **Specific** option and specify a domain name.
- In the **Theme Properties** dialog. Themes are described later in this guide in the [Configure Themes](#) section. Note that when you specify a domain name for a Theme, it overrides the domain name specified on the **Authentication** tab page (see above). To specify a domain name for a Theme, open the

Theme properties dialog, select the **General** category, select the **Override authentication domain** option, and specify a domain name.

When Parallels Client connects to Parallels RAS, the domain name, specified as described above, is passed back to it. When the user opens a dialog in Parallels Client to change their domain password, the domain name is automatically added to the user name and the username field is grayed out. This way the user doesn't have to specify the domain name.

Using a custom link to change domain passwords

If your users connect to Parallels RAS with Azure Active Directory Domain Services or a third-party IdP, you should configure changing domain passwords via a custom link. The link should point to the page that allows to change password in your service.

To specify the custom link for changing domain passwords:

1. Navigate to **Connection > Authentication > Change domain password**.
2. Select the **Use a custom link for the "Change domain password" option**.
3. In the text field below, specify the custom link for changing domain.

Note: Once a user connects, this URL will be cached on their device and the **Change Domain Password** option will remain available even when disconnected. The cached URL is accessible to all users on that machine.

If you update this URL, clients will continue using the old cached URL until their next successful connection.

Allowing users to discover RAS connections via email address

You can allow users to log in to the RAS Farms by using their email addresses. This way, users can gain access to applications and desktops published on a Farm without knowing the server address or hostname. All native Parallels Clients support finding Parallels RAS Farms by entering an email address.

For users to connect to a Farm using their email addresses, first you need to create a new TXT record in the forward lookup zone of users' domain on your DNS server. The specific way to do this depends on the configuration of the DNS server.

The syntax of the TXT record is as follows:

Host: `_prlsclient`

Text: `hostname:port/theme;connmode=X;authmode=X`

The following parameters are available for the text field:

- `hostname`: Hostname of the server where the Secure Gateway resides. This parameter is mandatory.
- `port`: Port on which the Secure Gateway listens for incoming connections. This parameter is optional.
- `theme`: Theme. This parameter is optional.
- `connmode`: connection mode. This parameter is optional. Possible values are 0, 1, 2, 3, where:
 - 0: Gateway mode
 - 1: Direct mode
 - 2: Gateway SSL
 - 3: Direct SSL
- `authmode`: Authentication type. This parameter is optional. Possible values are 0, 1, 2, 3, 4, where
 - 0: Credentials
 - 1: SSO
 - 2: Smartcard
 - 3: Web (SAML)
 - 4: Web + Credentials (Windows client)

Examples of the text value:

`hostname`

`hostname:port`

`hostname:port/theme`

`hostname;connmode=2;authmode=1`

After the DNS record is configured, users will be able to log in using their email addresses. For information on how to do that on specific clients, see *Parallels Client Guides*.

Load Balancing and HALB

This chapter describes load balancing options that you can use in Parallels RAS.

Resource based & round robin load balancing

Load Balancer in Parallels RAS is designed to balance RD Session Host connections from Parallels Clients.

The following types of load balancing are available:

- **Resource based.** Distributes sessions to servers depending on how busy the servers are. A new incoming session is always redirected to the least busy server.
- **Round robin.** Redirects sessions in sequential order. For example, let's say there are two RD Session Hosts in the Farm. The first session is redirected to server 1, the second session is redirected to server 2, and the third session is redirected to server 1 again.

Both methods are explained in this and the following subsections. Load Balancing options can be configured from the **Load Balancing** category in the RAS Console.

Selecting load balancing method

Load balancing is enabled by default when more than one server is available in a Site. The resource based load balancing is the default method. Load balancing method can be selected from the **Method** drop-down list.

Configuring resource counters

Resource-based load balancing uses the following counters to determine if a given server is busier than other servers and vice versa:

- **User sessions:** Redirect users to a server with the least number of sessions.
- **Memory:** Redirect users to the server with the best free/used RAM ratio.
- **CPU:** Redirect users to the server with the best free/used CPU time ratio.

When all of the counters are enabled, the Load Balancer adds the counter ratios together and redirects the session to the server with the most favorable combined ratio.

To remove a counter from the equation, clear the checkbox next to the counter name in the **Counters** section.

Session options

Reconnect to disconnected sessions: Enable this option to redirect incoming user sessions to a previously disconnected session owned by the same user.

Reconnect sessions using client's IP address only: When reconnecting to a disconnected session, the Parallels RAS will match the username requesting reconnection with the username of the disconnected session to match the sessions. With this option enabled, Parallels RAS will determine to which disconnected session to reconnect the session by matching the source IP address.

Limit each user to one session per desktop: Enable this option to ensure that the same user does not open multiple sessions. Please note that for this option to work, your session host must also be configured to restrict each user to a single session. In Windows Server 2012(R2), it's the "Restrict Remote Desktop Services users to a single Remote Desktop Services session" option in Local Group Policy Remote Desktop Services Remote Desktop Session Host Connections.

Disable Microsoft RD Connection Broker: If this option is enabled, the Microsoft RD Connection Broker will not interfere with the RAS brokering done by the RAS Connection Broker if it is installed. Please note that this option will only work with Windows Server 2012 and above.

Agent timeout and refresh time

You can also change the default timeout and refresh time for RAS agents running on the servers. If you believe that it takes too long to wait for an agent to respond or if the timeout is not long enough, you can specify your own values.

To change default timeouts:

1. Click the **Configure** button.
2. In the dialog that opens, specify the time period in seconds in the **Declare agent dead if not responding for** field. If the agent is not responding within this time period, the server is excluded from the load balancer.

3. In the **Agent Refresh Time** field, specify the number of seconds needed to check if the agent is reachable.

Configure CPU optimization

The CPU optimization functionality allows you to optimize CPU load balancing according to your requirements. When configured, the CPU load balancer will lower the priority of a process when its CPU usage exceeds a specified value for a specified number of seconds. The load balancer will revert the priority to its original level when the process has been running below a certain percentage for a certain number of seconds.

To configure CPU optimization, select the **Enable CPU Optimization** option and then specify the values as described below.

Start

Specifies when the CPU optimization should be activated. The **Total CPU usage exceeds** field specifies the system wide CPU usage in percent.

CPU conditions

Specifies thresholds per process when a specific process exceeds or falls below the specified CPU percentage. Here you can specify **Critical** and **Idle** values. The CPU load balancer will adjust other priorities with respect to these values.

Please note that CPU usage values are attenuated and calculated based on the agent refresh time configured on the [Load Balancing](#) tab.

Exclusions

Use the **Exclusions** list to specify processes that should be excluded from CPU optimization. Click **Tasks > Add** to select a process. To remove a process from the list, select it and click **Tasks > Delete**.

Irregular values for critical/idle may cause issues (processes set to idle due to incorrect configuration). If there are issues with getting the CPU usage counter, optimizations cannot be applied.

Log files can be found in %ProgramData%.log. Use the log to confirm thresholds. You can check the CPU usage performance counter on Windows.

Note: Since the critical/idle thresholds are calculated based on the highest process CPU usage (not the absolute CPU usage), this value is not reflected in the logs when changing priorities.

Absolute CPU usage equals to total CPU usage. For example, if there are 2 processes taking 30% each, the total CPU usage is 60%. The usage threshold when CPU load balancer kicks in is 25% (default).

The highest process CPU usage is the CPU usage of the process taking the most CPU. For example, if you have three processes, two taking 10% and the third taking 40%, the highest CPU usage is 40%.

High availability load balancing (HALB)

High availability load balancing (HALB) in Parallels RAS is a functionality that load balances RAS Secure Gateways. The load balancer is built into a Parallels HALB appliance, which is a preconfigured virtual machine with the operating system installed and all relevant settings configured.

Parallels HALB appliance has been tested to work for the following hypervisors:

- Microsoft Hyper-V
- VMware ESXi
- Nutanix AHV
- Scale HyperCore

Please note that other hypervisors may also be used, and support will be provided, but only on a best effort basis.

HALB is deployed in Parallels RAS on a Site level. You can have multiple HALB configurations per Site, which are called Virtual Servers. Each Virtual Server has its own IP address (called Virtual IP or VIP) and is assigned one or more HALB appliances (also called HALB devices in the Virtual Server context) that perform the actual load balancing. An HALB Virtual Server is a virtual representation of HALB devices. It provides traffic distribution to HALB devices when they are properly configured. Since the IP address of a specific Virtual Server is the single

point of contact for the client software, it is recommended to have at least two HALB devices per Virtual Server for redundancy.

Multiple HALB devices assigned to a Virtual Server can run simultaneously, one acting as the primary and others as secondary. The more HALB devices a Virtual Server has, the lower the probability that end users will experience downtime. The Virtual Server is assigned the IP address of the primary HALB device, which is shared with secondary HALB devices. Should the primary HALB device fail, a secondary is promoted to primary and takes its place using the same IP address for client connections.

Note: Please note that when a secondary HALB device is promoted to primary, a user may experience up to two disconnects. The first disconnect will occur when an HALB device goes down. The second disconnect may happen when a device goes back online. The disconnects cannot be avoided because the virtual IP address has to be transferred from one HALB device to the other, which means that the first device has to stop communications over this address, while the other device will have start it. Note that disconnects don't affect user sessions. Users are able to reconnect to their sessions and no user data is lost.

Setting up High Availability Load Balancing consists of the following steps:

1. Deploying one or more Parallels HALB appliances (devices).
2. Configuring one or more Virtual Servers in the RAS Console.

Read on to learn how to download and deploy a Parallels HALB appliance.

Prerequisites

The below highlights the prerequisites required to use HALB:

- Firewall or router in front of a HALB configured to preserve the source IPs of client devices

Sizing recommendations

Traffic Profile	Concurrent TCP Sessions	Recommended vCPUs	Recommended RAM	Notes
	Up to 12,500	2 vCPUs	2 GB	

Traffic Profile	Concurrent TCP Sessions	Recommended vCPUs	Recommended RAM	Notes
Low-intensity payloads				Safe and stable for small payloads and low send frequency
Moderate-intensity payloads	Up to 12,500	4 vCPUs	2–4 GB	Provides headroom and improved CPU efficiency under sustained load
Moderate-intensity payloads	Up to 20,000	8 vCPUs	4 GB	Required to avoid CPU saturation at high concurrency
All profiles (general guidance)	Any	—	4 GB (preferred)	4 GB RAM provides additional safety margin and resilience during prolonged high-load operation

Deploying a Parallels HALB appliance

To download a Parallels HALB appliance, visit <https://www.parallels.com/products/ras/download/links/>.

On the **Download Parallels Remote Application Server** web page, scroll down to the **Download Optional Server Components** table and find the **Parallels Remote Application Server HALB Appliances** row. The row contains the following download links:

- HALB Appliance OVA
- HALB Appliance VHD
- HALB Appliance QCOW2

The appliance type that you need to download depends on the hypervisor that you are using. Please follow the instructions below for your hypervisor type.


VMware

For VMware, the appliance can be imported using the OVA appliance file. The VM is created pre-configured with 2 CPUs and 2048 MB of RAM (which can be changed later).

Microsoft Hyper-V

For Microsoft Hyper-V, the appliance can be imported using the VHD file.

Note: When setting up a Hyper-V machine, choose the **Generation 1** configuration options.



Nutanix AHV

For Nutanix AHV, the appliance can be imported using the QCOW2 file. Once downloaded, follow the recommended steps:

1. From the **Prism Element Console**, navigate to **Settings > Image Configuration > Create Image > Upload a file** and pick the downloaded QCOW2 disk image.

2. Once successful, it should report its status as **ACTIVE**.

HALBv21	DISK	ACTIVE	32 GiB	 

3. Navigate to **VM > Create VM**.

4. Under **Boot Configuration**, choose **Legacy BIOS**.

Boot Configuration

☐ UEFI 

☒ Legacy BIOS

Set Boot Priority

Default Boot Order (CD-ROM, Disk, Network) 

5. Choose **Add New Disk** and set the parameters as follows:

1. Set **Type** to **DISK**.

2. Set **Operation** to **Clone from Image Service**.

3. Under **Image**, select the previously added disk image.

Add Disk?×

Type

DISK

Operation

Clone from Image Service

Bus Type

SCSI

Image ?

HALBv21

Logical Size (GiB) ?

32

Please note that changing the size of an image is not allowed.

Index

Next Available

Cancel

Add


6. Click **Add**.

7. Navigate to **VM** > {VM_name} > **Power On** > **Launch Console**.

Scale HyperCore

For Scale HyperCore, the appliance can be imported using the QCOW2 file. Once downloaded, follow the recommended steps:

1. Navigate to **Control Center** > **Media** > **Upload**.



2. Click **Create VM**.
3. Set **Boot Type** to BIOS.
4. Under **Drive**, select the **Attach Uploaded Drive** option and pick the downloaded QCOW2 image.

Create VM

Name

HALBv21

Description

Tags

OS

Other

▼

Drivers

Performance

▼

Boot Type ?

BIOS

▼

CPU

1

▼

Memory

1

GiB

▼

Drive

100

GB

▼

RAS_HALB_Applian... ▼

34.36 GB

×

⊞ Attach Uploaded Drive

⊞ Add Drive

Boot From

None

▼

Create

5. Remove all other drives except for the one attached in the previous step.

6. Click **Create**.




Deploying a Parallels HALB Appliance

After you download a Parallels HALB appliance, you need to import it to a hypervisor running on a separate machine connected to the same local network as Parallels RAS. For the information on how to import a virtual appliance, please consult your hypervisor documentation.


HALB Device First Boot Configuration

1. Choose between static IP or DHCP.

[Note: When using DHCP \(default\), the admin must make sure the IP of the HALB device does not change when the DHCP lease expires by adding a DHCP reservation.⁶](#)



2. For the static option, specify the network parameters.




3. [OPTIONAL] Configure mailbox notifications.

4. Wait for the HALB device to be ready for configuration from the RAS console.

5. In the **Advanced** menu, the following settings are available:

1. **Networking.** Reconfigure the network parameters of this device.
2. **Rescan.** Rescan the Appliance services state, since they don't refresh automatically.

3. **Reset.** Use to reset the existing Appliance configuration due to a glitch or adding it to a different HALB virtual server.



To add a HALB virtual server:

1. In the RAS console, navigate to **Farm** > <Site> > **HALB**.
2. On the **Virtual Servers** tab in the right pane, click **Tasks** > **Add**. The **HALB Configuration** wizard opens.
3. Make sure the **Enable HALB** option is selected.
4. Type a name for this virtual server and an optional description.
5. In the **Public address** field, type a public FQDN or IP address of this server. This is used by the Preferred routing functionality for redirecting client connections. Please see [Configuring preferred routing](#).
6. In the **Virtual IP** section, specify the virtual IP address properties that will be used for incoming client connections by a HALB device that you will assign to this Virtual Server later.
7. In the **Settings** section, select one or more of the following options. Note that at least one “LB” option must be selected. If you skip an option at this time, you can add it later in the virtual server properties dialog:
 - **LB Gateway Payload:** Enables load balancing of normal (unsecured) gateway connections.
 - **LB SSL Payload:** Enables load balancing of SSL connections.
 - **Client Management:** Enables management of Windows client devices connected through HALB.
8. Click **Next**.

From this point forward, depending on the payloads that you selected in the previous step, a wizard page will open where you can configure the payload properties. These pages are described below.

LB Gateway payload

Configure load balancing for normal connections:

1. Set the port number used by HALB devices to forward traffic to RAS Secure Gateways. The port is configured on a gateway. The default port is 80.
2. In the **Gateways** list, select a RAS Secure Gateway to be load-balanced. Please note that only one IP address per gateway can be used. If you have more than one entry for the same gateway with different IP addresses, you can select just one.

LB SSL payload

Configure load balancing for SSL connections:

1. Set the port number used by HALB devices to forward SSL traffic to RAS Secure Gateways. The port is configured on a gateway. The default port is 443.
2. Select the SSL mode from **Passthrough** or **SSL Offloading**. By default, SSL connections are tunneled directly to gateways (referred to as Passthrough) where the SSL decryption process is performed.

The **SSL Offloading** mode requires an SSL certificate to be assigned to HALB. When you select it, click **Configure** and specify the following:

- **Accepted SSL Version:** Select an SSL version.
- **Cipher Strength:** Select the cipher strength of your choice. To specify a custom cipher, select **Custom** and then specify the cipher in the **Cipher** field.
- The **Use ciphers according to server preference** option is ON by default. You can use client preferences by disabling this option.
- **Certificates:** Select a desired certificate. For the information on how to create a new certificate and make it appear in this list, see the [SSL Certificate Management](#) chapter.

The **<All matching usage>** option will use any certificate configured to be used by HALB. When you create a certificate, you specify the “Usage” property where you can select “Gateway”, “HALB”, or both. If this property has the “HALB” option selected, it can be used with HALB. Please note that if you select this option, but not a single certificate matching it exists, you will see a warning and will have to create a certificate first.

3. Select a gateway to be load-balanced. Note that only one IP address per gateway can be used.

Device Manager

Configure Windows client device management, select a gateway that will manage Windows client devices. Note that only one IP address per gateway can be used.

Devices

To assign HALB devices to the Virtual Server:

1. Click **Tasks > Add** and select or specify a HALB device. If you haven’t deployed any HALB devices (appliances) yet, you can still save the Virtual Server configuration and assign HALB devices to it later. At least two HALB devices are recommended per Virtual Server. For more info, see [High Availability Load Balancing \(HALB\)](#). HALB device priority is set by positioning a device in the list. The device at the top is the primary HALB device. Devices under it are secondary HALB devices. To promote a device to primary, simply move it to the top of the list.
2. Finally, click **Finish** to save the Virtual Server settings and close the wizard.

The new virtual server will appear in the list in the RAS Console.

Modifying Virtual Server and configuring advanced options

To modify the Virtual Server settings, right-click it and choose **Properties**. The tabs in the **Properties** dialog have the same options as the wizard pages described above. The only exception is the **Advanced** tab, which is described below.

To view and configure advanced Virtual Server options, select the **Advanced** tab. The options that you see on this tab are applied to all HALB devices assigned to a Virtual Server. This list gives you simple access to the HALB device options without logging in to the virtual machine directly. Please note that changing any of

these values may potentially lead to undesired results. You should only change them according to specific network requirements.

The following advanced settings are available:

Option	Default value	Description
Enable RDP UDP tunneling	Enable	Enables RDP clients to transfer RDP over UDP traffic through HALB devices.
Maximum TCP connections	2000	Sets the maximum number of concurrent TCP connections.
Client inactivity timeout (s)	150	Maximum inactivity time on the client side in seconds.
Gateway connection timeout (s)	30	Maximum time to wait for a connection attempt to a gateway to succeed in seconds.
Client connection queue timeout (s)	30	When a device's Max TCP connections is reached, connections are left pending in a queue for the period of this timeout (seconds).
Gateway inactivity timeout (s)	150	Set the maximum inactivity time for gateways in seconds.
Amount of TCP connections per second	1000	Set a limit on the number of new connections accepted per second on an HALB device.
Gateway health check intervals (s)	5	Set the interval between two consecutive health checks in seconds.
VRRP virtual router ID	15	Used to differentiate multiple instances of VRRP running on the same network.
VRRP authentication password	-	Enable password authentication for VRRP communication between HALB devices used by for failover synchronization.
VRRP broadcast interval (m)	1	Minimum time interval in minutes for refreshing gratuitous ARPs while device is in active state.
	2	

Option	Default value	Description
VRRP health script check interval (s)		Set the interval between invocations of the script that ensures local HALB services are up and running (seconds).
VRRP health script check timeout (s)	10	Execution timeout for the script that ensures local HALB services are up and running (seconds).
VRRP advertisement interval (s)	1	The time interval between the advertisement packets that are being sent between HALB devices in the same VRRP group (seconds).
Keep existing load balancing settings	Disable	Keep load balancing configuration currently present on the device and do not overwrite with new settings.
Keep existing VRRP/keepalived	Disable	Keep VRRP/keepalived configuration currently present on the device and do not overwrite with new settings.

HALB Device status and version number

HALB device status and version information can be verified in two places in the RAS Console, which are described below.

Site tab

You can view HALB devices and related information on the **Site** tab in the RAS Console. To see it, navigate to **Farm > Site**. Note the **Agent** and **Agent Version** columns. The two columns are described below.

The **Agent column** can have the following values:

- **Not verified** (red) - The agent is not verified and cannot communicate. If you see this, verify the agent.
- **Needs update** (yellow) - The agent is functioning normally but is an older version. If you see this, you should update the agent to the latest version.
- **Agent OK** (green) - The agent is OK. No actions are necessary.

The **Agent Version** column displays the actual agent version, including the Parallels RAS version and build numbers.

You can also ping a HALB device by right-clicking it and choosing **Tools > Ping host**. For additional information about using the **Tools** menu and the **Ping** tool specifically, please see [Computer management tools](#).

Devices tab

The HALB devices agent status and version can also be viewed in the main HALB subcategory. To see it, navigate to **Farm > Site > HALB** and select the **Devices** tab. The agent information displayed here is the same as on the **Site** tab described above.

HALB maintenance

When you need to replace or repair a HALB device (virtual machine), you can simply remove it from the Virtual Server configuration and then add the repaired or new device later. If you need to temporarily remove all HALB devices from a Virtual Server configuration, you can do that too.

You can also disable the Virtual Server during maintenance by clearing the **Enable HALB** option on the **General** tab in the Virtual Server properties dialog.

HALB connection and session information

To see the number of TCP connections per HALB device, navigate to **HALB > Devices** and examine the **TCP Connections** column in the device list. To refresh the list, click **Tasks > Refresh**.

To see the session information per Virtual Server, navigate to **Farm > Site**. The session count is displayed for each Virtual Server in the **Session** column.

RAS Multi-Tenant Architecture

Overview


Beginning with RAS 17.1, Parallels introduces a new multi-tenant architecture, with the addition of Parallels RAS Tenant Broker, enabling organizations to share components from the same RAS infrastructure among different Tenants while keeping client data segregated and reducing costs.

The RAS multi-tenant architecture offers the following advantages to Service Providers and organizations:

- **Cost savings** due to reduction of number of RAS Secure Gateways and High Availability Load Balancers (HALBs) while maximizing resource usage and consolidation.
- **Faster onboarding** of new tenants/customers.
- **Simplified centralized management** of multi-tenant environments.
- **Extended market reach** through reduction of operational costs for organizations of any size by allowing cost scaling through shared infrastructure.

Architecture description

The following diagram illustrates a typical Parallels RAS deployment that uses the RAS multi-tenant architecture.



- Firewalls and HALB are installed in a DMZ and are shared by Tenants.
- Tenant Broker is a special RAS installation that hosts shared RAS Secure Gateways and HALB, and can also use RAS access layer. Tenant Broker is

installed using the **Parallels RAS Tenant Broker** option in the Parallels RAS installer. Tenant Broker can be installed in its own domain or outside of a domain.

- Tenant farms are deployed just like traditional on-premises RAS environments and are joined to the Tenant Broker. Each Tenant Farm has its own RAS Connection Brokers and servers hosting published resources (VDI, RD Session hosts, or Remote PCs). No local RAS Secure Gateways and HALB (or third-party load balancers) are needed.
- Tenants are joined to the Tenant Broker and each Tenant is represented as a Tenant object in the Tenant Broker.
- Parallels Clients (both platform-specific and Web) connect to shared gateways in the Tenant Broker. When a client connects to User Portal, a Theme from the corresponding Tenant is always used depending on which Tenant the client belongs to.

Implementation overview

The following is an implementation overview of the RAS multi-tenant architecture:

- Tenants are deployed as separate individual Farms or Sites. Tenants deployed as separate Farms are completely independent and never communicate with each other. If tenants are deployed as Sites, every Site must join the Tenant Broker separately.
- Shared resources include RAS Secure Gateways (including User Portal) and High Availability Load Balancers (HALB).
- A Tenant Farm doesn't need its own RAS Secure Gateways and HALB. However, deployments with Gateways and HALB are possible if you need them for internal connections. For example, if you have different policies for internal and external connections, you might want to install a Gateway and HALB to serve local users.
- The network configuration of a Tenant requires the Tenant Connection Broker to Tenant Broker Connection Broker connectivity. Additionally, shared RAS Secure Gateways need to communicate with servers hosting published resources and the Tenant's Connection Broker. Depending on the implemented network architecture, it might require a VLAN to VLAN connectivity, VPN, etc. These communications require only a limited number of open ports. For the complete list, see [Communication ports](#).
- Communications with a Tenant domain are always performed from a local Tenant Connection Broker and never from the Tenant Broker infrastructure.

- Every Tenant must have a unique public domain address, which can be assigned a number of different ways. For example, a service provider can register a subdomain (e.g. Tenant1.Service-Provider.com) and assign it to a Tenant. Another approach could be using a private domain address (e.g. RAS.Tenant1.com) and have it routed to RAS Secure Gateways in the Tenant Broker. Note that different public domain addresses can resolve to the same IP address if needed.
- When a Tenant is joined to the Tenant Broker, shared RAS Secure Gateways become aware of the Tenant and its configuration and can connect to the Tenant's RAS Connection Broker(s). A route must be set for the incoming Tenant's traffic from the Internet to RAS Secure Gateways (or HALB) in the Tenant Broker.
- Tenant Broker comes with its own RAS Console allowing you to manage shared resources, Tenant objects and certificates, monitor Tenant performance, and carry out standard RAS administration tasks.
- All Tenant Themes are made available in the Tenant Broker. When user connects via a shared RAS Secure Gateway in the Tenant Broker, the corresponding Tenant Theme is presented to the user.
- Different SSL certificates can be used for different Tenants.

Licensing

Tenant Broker doesn't need a license. Licenses are managed on a Tenant level.


RAS version compatibility

Parallels RAS multi-tenant architecture is available in Parallels RAS 17.1 and newer. The following limitations apply when using older versions of Parallels RAS:

- Parallels Clients older than RAS 17.1 are incompatible with shared gateways and therefore cannot be used to connect to a Tenant Farm via the Tenant Broker.
- Parallels RAS installations older than RAS 17.1 are incompatible with Tenant Broker and cannot be joined as Tenants.

User connection flow

The following diagram illustrates the RAS user connection flow through Tenant Broker:



Shared RAS Secure Gateways installed in Tenant Broker are able to work with multiple concurrent user sessions in multiple Tenant farms. On the diagram above, you can see two users (1 and 2) connecting to different Tenant Farms (Tenant 1 Farm and Tenant 2 Farm). Both connections are tunneled through the same Gateway and then delivered to the correct Tenant Farm.

The connection flow consists of the following steps:

1. (1A), (2A) — A user initiates a RAS connection to a public address registered in the Tenant Broker. The (1A) connection goes to the Tenant 1 public address; the (2A) connection goes to the Tenant 2 public address.
2. (1B), (1C) — The shared Gateway makes a decision where to forward a user connection based on a hostname used in the initial connection (1A, 2A). After that each client establishes a RAS session with a Connection Broker of their respective Tenant Farm. Tenant's Connection Broker authenticates the user against Active Directory of the Tenant. After that, the user receives the list of published applications available to him or her.
3. (1D), (2D) — A user start a Remote User Session to a published application. The shared Gateway requests from Tenant's Connection Broker an address of a server to forward the remote session to and forwards it.

The mapping of public addresses to Tenants is configured on shared Gateways by the Tenant Broker Connection Broker.

Deploying Tenant Broker and Tenants

A typical scenario of deploying the multi-tenant architecture of Parallels RAS consists of the following steps:

1. Deploy Tenant Broker.

2. Deploy a traditional RAS Farm to operate as a Tenant.
3. Configure network between the Tenant Broker and the Tenant to allow the following connections:
 - Shared RAS Secure Gateways to Tenant RAS Connection Brokers.
 - Shared RAS Secure Gateways to resources hosts.
 - Tenant RAS Connection Brokers to Tenant Broker RAS Connection Broker.

For the information about ports numbers, please see [Communication ports](#).

4. Create a Tenant object and a corresponding invitations hash in the Tenant Broker console, or create a secret key (more on this later in this chapter).
5. Join the Tenant to the Tenant Broker using the invitation hash or the secret key.
6. Assign a public domain address to the Tenant. This can be done at this point (after you join a Tenant) or it can be done in advance if you wish. Either way it has to be done or the clients will not be able to connect to the Tenant Farm.
7. Set up routing for incoming Tenant traffic from the Internet to shared RAS Secure Gateways and HALB.
8. Configure a certificate for the Tenant. By default, a self-signed certificate created during the installation will be used.
9. Test the client connectivity.

The subsequent sections describe the steps above in detail.

Deploying Tenant Broker

First you need to install Tenant Broker on a dedicated server. Please note that if you have Parallels RAS already installed on a computer where you are planning to install Tenant Broker, you need to uninstall it first. The two installation versions cannot coexist on the same machine.

To install Tenant Broker:

1. Run the standard Parallels RAS installer.
2. On the **Select Installation Type** page, select **Parallels RAS Tenant Broker**.

3. Click **Next** and follow the onscreen instructions.

Once the installation is finished, run the Parallels RAS Console.

When the console starts, you'll see that it has a different set of categories and managed objects compared to the standard RAS Console. The purpose of the Tenant Broker console is to manage shared resources and Tenants. It is not used to manage RD Sessions Hosts, VDI, or any other standard RAS resources because they are deployed and managed in individual Tenant Farms.

The Tenant Broker console

You can manage the following categories and object in the Tenant Broker console:

- **Farm.** This category allows you to manage Tenants, Gateways, Connection Brokers, HALB, and Certificates. The **Settings** subcategory allows you to manage global logging and the Tenant Broker itself.
- **Administration.** Allows you to perform management tasks similar to the standard RAS Console: Accounts, Settings, Mailbox, Reporting, Settings Audit.
- **Information.** Lists services and components running in the Tenant Broker and their status.

As with the standard RAS Console, every time you modify any of the objects, you need to click the **Apply** button for the changes to be saved in the configuration database.

Install RAS Secure Gateways

By default, Tenant Broker does not have any RAS Secure Gateways installed. To add a Gateway, log in to the Tenant Broker console, navigate to **Farm > Secure Gateways** and click **Tasks > Add**. If you already have one or more RAS Secure Gateways, which are not used in any other RAS Farm, you can also add such a Gateway to the Tenant Broker. Please note that existing RAS Secure Gateway installations must be RAS version 17.1 or newer. Gateways from older RAS versions cannot operate as shared gateways.

To install a new gateway, run the Parallels RAS installer on a desired server, choose **Custom** and select the **RAS Secure Gateway** component. After the installation is finished, go back to the Tenant Broker console and add the gateway to the Tenant Broker.

Deploying a Tenant

A Tenant Farm is deployed just like a traditional Parallels RAS Farm. The only difference is, when installing the Farm, you don't need to install RAS Secure Gateways in it.

Note: If you decide to install a local (private) RAS Secure Gateway in a Tenant Farm (e.g. for local connections), you can do that, but please keep in mind that you cannot mix HALB and Gateways from the Tenant Broker and a Tenant Farm. The HALB appliance installed in the Tenant Broker will not support this scenario.

To set up a Parallels RAS Farm to be used as a Tenant:

1. Run the Parallels RAS installer.
2. On the **Select Installation Type** page, select **Custom**.
3. Click **Next**.
4. Make sure that the following components are selected for installation:
 - RAS Connection Broker
 - Parallels RAS Console (optional; you can have the RAS Console installed on a different machine)

Other components are optional. You can install them now or you can install them later if needed.

5. Click **Next** and follow the onscreen instructions to complete the installation.

Join a Tenant to Tenant Broker

Once the Tenant Farm is operational, you can join one or more sites in it to the Tenant Broker.

Note: A Tenant is a Site in a separately deployed Parallels RAS Farm. When you join a Tenant to Tenant Broker, you join a Site. When you want to join the whole Farm, you do it one Site at a time. Of course, if you have just one Site in a Farm (and have no plans to create more sites), you are essentially joining the whole Farm.

There are two ways you can join a Tenant: (1) Using an invitation hash or (2) Using a shared secret key. The difference between the two is as follows:

- **Invitation hash.** An invitation hash is an automatically generated encrypted string that can be used to join a single Tenant to Tenant Broker. Invitation hash is a property of a Tenant object, which is created in the Tenant Broker console. You email the hash to the Tenant Farm administrator, so they can use it to join the Tenant Broker. Once used, an invitation hash cannot be used again by any other Tenant.
- **Shared secret key.** A shared secret key is similar to an invitation hash, with one important difference. It can be used to join an unlimited number of Tenants. A Tenant object is not pre-created for a secret key in the Tenant Broker. Instead, the object is created when the key is used to join a Tenant. Because of its unlimited usage capability, only the Tenant Broker admins should have access to a shared secret key. This scenario is useful when there are multiple Tenants, all managed by the same Tenant Broker administrator.

The invitation hash scenario is described below. For the secret key scenario see [Joining with a secret key](#).

First, you need to generate an invitation hash and create a Tenant object on the Tenant Broker side:

1. Log in to the Tenant Broker.
2. In the RAS Console, navigate to **Farm > Tenants**.
3. Click **Tasks > Add**.
4. In the **Tenant properties** dialog, specify the following:
 - **Name:** Type a Tenant name (this can be any name that you like).
 - **Public domain address:** If you've already assigned a public domain address to the Tenant, specify it here. If not, you can leave it blank. The address is not required for the Tenant to join the Tenant Broker. However, without the address specified here, end users will not be able to connect to the Tenant, so you will need to come back and fill it in later. For details, see [Assign a public domain address](#).
 - **Clients in gateway mode connect to published tenant resources by server IP:** When selected, clients will use the Tenant IP address instead of the DNS name. You can use this option when a Tenant farm does not share the same DNS provider as the Tenant Broker farm.
 - **Do not show billing information:** When selected, billing information is not shown in the [Licensing](#) category of the Tenant.

- **Description:** Type an optional description.
 - **Connection Brokers:** This field is disabled and will be populated automatically when the Tenant joins the Tenant Broker. See more in [Tenant configuration](#).
 - **Tenant invitation hash:** This is the hash that the admin of the Tenant Farm will need to use to join the Tenant Broker. A hash is generated automatically when you open this dialog. To generate a new hash, click **Create new hash**.
 - **Send via email.** You can give the invitation hash to the Tenant admin directly or you can use this button to send it via email. When you click the button, you'll see a dialog where you can enter the recipients and where you can review and modify the email message. By default, the message contains instructions on how to join the Tenant Broker. Please note that SMTP settings must be configured in the RAS Console before you can use the email option. You can configure SMTP first and then return to this screen to complete this step.
5. Click **OK** to close the **Tenant properties** dialog. The new Tenant will appear in the **Tenants** list in the console. At this time, the Tenant is not joined yet. Read on to learn how to join it.

To join the Tenant to the Tenant Broker:

1. Log in to the Tenant Farm.
2. In the RAS console, navigate to **Farm > Site**. Note that you are joining a Site to the Tenant Broker, not the whole Farm, so if you have more than one Site, you need to join them one by one.
3. Click **Tasks > Join Tenant Broker**.
4. In the **Join Tenant Broker** dialog, enter the invitation hash that you obtained from the Tenant Broker in the previous steps (or, if you are an admin of a Tenant Farm, the one you received in the invitation email).
5. Click **Join**.

On successful join, you will see a message welcoming you to the Tenant Broker. If the primary Connection Broker in your Tenant Farm can't reach the Tenant Broker, you will see a corresponding error message. Make sure that the Tenant Broker computer is reachable from the machine where you have the Tenant's RAS Connection Broker running.

Overriding Tenant Broker IP address

The Tenant Broker IP address is detected automatically when you generate an invitation hash (or a secret key) and is embedded into the hash. If a Tenant can't reach the Tenant Broker using this address, you have the ability to override it as follows:

1. Log in to the Tenant Broker.
2. In the RAS Console, navigate to **Farm > Settings** and click the **Tenant broker** tab.
3. Select the **Override Tenant Broker address in tenant invitations and secret keys** option.
4. Enter the desired IP address in the field provided.

When done, the specified IP address will be used instead of the auto-detected address when generating an invitation hash or secret key. When the hash is used on the Tenant side to join the Tenant Broker, the Tenant will use this address to connect to the Tenant Broker.

Once used on the Tenant side, an invitation hash binds the Tenant Farm to the corresponding Tenant object in the Tenant Broker and the tenancy becomes effective.

Joining with a secret key

In addition to an invitation hash, you can join a Tenant to the Tenant Broker using a secret key. [As described earlier](#), a secret key can be used to join an unlimited number of Tenants to the same Tenant Broker.

To create a secret key:

1. Log in to the Tenant Broker console.
2. In the RAS Console, navigate to **Farm > Settings**.
3. Select the **Tenant broker** tab.
4. Select **Allow RAS Farms to register in Tenant Broker using a secret key**.
5. Optionally, select **Do not show billing information** to hide billing information in the **Licensing** category of Tenants joined with secret keys.
6. The secret key is generated automatically. To generate a different key, click **Generate**.

7. If you want to register Tenants as subdomains, specify the domain part of the hostname in the **Domain** field. For example, to use “subdomain.domain.com” as a Tenant host name, specify “domain.com”.

Once you have the key, you can use it to join one or more Tenants to the Tenant Broker.

Note: Due to its unlimited usage capability, only the Tenant Broker administrator should have access to a shared secret key. Secret keys can be practical when the Tenant Broker administrator manages Tenant Farms, so instead of generating a hash for every Tenant, he/she can use a single secret key to join all of them to the Tenant Broker.

To join a Tenant using a secret key:

1. Log in to the Tenant.
2. In the RAS Console, navigate to **Farm > Site**.
3. Click **Tasks > Join Tenant Broker**.
4. In the **Join Tenant Broker** dialog, specify the following:
 - Enter the secret key in the first field from the top. If the Tenant is able to reach the Tenant Broker, the **Tenant Broker** field will be populated automatically.
 - The **Tenant Name** field is populated automatically based on the name of the current Site, but you can specify a Tenant name of your choosing. The name you enter will be used in the Tenant Broker to name the corresponding Tenant object.
 - In the **Public domain addresses** field, you can specify public domain addresses that will be used to access the Tenant. Configuring this is optional. If the **Domain** field is configured in the Tenant Broker settings (see above), you may enter subdomain only rather than the full domain address.
5. Click **Join**.

On successful join, you will see a message welcoming you to the Tenant Broker. If the primary Connection Broker in your Tenant Farm can't reach the Tenant Broker, you will see a corresponding error message. Make sure that the Tenant Broker computer is reachable from the machine where you have the primary Connection Broker running.

Overriding the Tenant Broker IP address

The Tenant Broker IP address is detected automatically when you generate a secret key and is embedded into it. If a Tenant can't reach the Tenant Broker using this address, you have the ability to override it as follows:

1. Log in to the Tenant Broker.
2. In the RAS Console, navigate to **Farm > Settings** and click the **Tenant broker** tab.
3. Select the **Override Tenant Broker address in tenant invitations and secret keys** option.
4. Enter the desired IP address in the field provided.

Verify join status

After you join a Tenant to the Tenant Broker, you should verify that the procedure was successful.

First, verify the Tenant Broker status in the Tenant console:

1. Log in to the Tenant Farm.
2. In the RAS Console, navigate to **Farm > Site** and select the **Site** tab in the right pane.
3. You should see the **Tenant Broker** section with the **Status** column, which should say **OK**. If the status is **Not verified**, make sure that the Tenant Broker is operational (or contact the Tenant Broker admin if you are not him or her).

You can also see additional Tenant Broker information by right-clicking it and choosing **Properties**. The information includes the following:

- **Name:** The Tenant Broker name.
- **Primary address:** The primary RAS Connection Broker address.
- **Secondary address:** The secondary RAS Connection Broker address (if available).

You should then verify the Tenant status in the Tenant Broker console:

1. Log in to the Tenant Broker.
2. In the RAS Console, navigate to **Farm > Tenants**.

3. In the **Tenants** tab, find the Tenant of interest and examine the **Status** column, which should say **OK** if the Tenant is joined properly. For other possible **Status** column values, see [Tenant configuration](#).

Configure network

After deploying a Tenant, you need to configure networking between Tenant Broker and Tenant in order to allow the following communications:

- Tenant Connection Broker > Tenant Broker Connection Broker: port 20003
- Tenant Broker Gateway > Tenant Broker Connection Broker: port 20002
- Tenant Broker Gateway > Tenant Connection Broker: port 20002
- Tenant Broker Gateway > Servers hosting published resources: port 3389

These are standard RAS ports, which are also described in the **Port reference** section.

Assign a public domain address

Every Tenant must have a unique public domain address for end users to connect to it through Tenant Broker. Although every Tenant must have a unique public domain address, it is not required for every Tenant to have a unique IP address. Different public domain address can be configured to resolve to the same IP address to reach the Tenant Broker shared Gateways. This way the Tenant Broker is still able to forward traffic to the right tenant based on the hostname requested by an end user.

A public domain address can be chosen a number of different ways. For example, a service provider can register a subdomain (e.g. Tenant1.Service-Provider.com) and assign it to a Tenant. Another approach could be using a private domain address (e.g. RAS.Tenant1.com) and have it routed to RAS Secure Gateways in the Tenant Broker. For testing purposes, you can even use an IP address.

The **Public domain address** is also a property of a Tenant object in the Tenant Broker console. After joining a Tenant to the Tenant Broker, you must ensure that this property contains the correct address. Otherwise end users will not be able to connect to the Tenant through the Tenant Broker.

To verify (and set if necessary) the Tenant's public domain address:

1. Log in to the Tenant Broker.
2. In the RAS Console, navigate to **Farm > Tenants**.
3. Right-click a Tenant and choose **Properties**.
4. In the **Properties** dialog, verify that the **Public domain address** field contains the correct address.

Configure an SSL certificate

The public domain address assigned to a Tenant must have a matching certificate. The Tenant Broker admin must create a certificate for every Tenant in the Tenant Broker console. Shared RAS Secure Gateways must then be configured to use these certificates. Tenant certificates are created and managed in Parallels RAS the same way as other certificates using the **Farm > Site > Certificates** subcategory. For the complete information about how to create certificates and how to assign them to RAS Secure Gateways and HALB, please see the [SSL Certificate Management](#) chapter.

When a user connects to the Tenant's public domain address, a certificate with the common name matching the requested public domain address is selected automatically for every connection. The first available certificate is used which might not be the self-signed (say it was deleted)

If no matching certificate is found, the default self-signed certificate will be used, but the user will see a certificate warning in the web browser.

Set up routing for incoming traffic

One other thing that you have to do after you join a Tenant to the Tenant Broker, is set up routing for the incoming traffic from the Internet to shared RAS Secure Gateways or HALB.

User authentication

User authentication in the RAS multi-tenant architecture is performed by the RAS Connection Broker running in the Tenant Farm. The Connection Broker is selected randomly by a shared RAS Secure Gateway. If the Connection Broker is unavailable, then it's marked accordingly and no communication is conducted with

it from the same shared gateway for a period of time. The gateway checks the Connection Broker status periodically and resumes communications as soon as the agent becomes available.

Unjoining from Tenant Broker

To unjoin a Tenant from the Tenant Broker, do the following:

1. Log in to the Tenant Farm.
2. In the RAS Console, navigate to **Farm > Site**.
3. Click **Tasks > Unjoin from Tenant Broker**.

The Tenant will be unjoined from the Tenant Broker. As a result, the Tenant users will no longer be able to connect to the Tenant Farm through the Tenant Broker.

Managing Tenants

Tenant configuration

To see the list of existing Tenants in the Tenant Broker console, select **Farm > Tenants**.

The **Status** column indicates the Tenant status, which can be one of the following:

- **OK** — The Tenant has joined and has been verified.
- **Not Joined** — The Tenant object was created for the Tenant and the invitation hash was generated, but the Tenant has not joined the Tenant Broker yet.
- **Not Verified** — The Tenant has joined, but no connection to the Tenant's RAS Connection Broker has been established yet. This status is usually displayed for a minute or so immediately after the Tenant joins the Tenant Broker. Once the connection is established, the status changes to **OK**.

This status can also appear when the Tenant Broker loses a connect with the Tenant's primary Connection Broker. Shared gateways will be able to process connections only if they are still able to communicate with the Tenant's Connection Broker on their own. They are independent from the Tenant

Broker's Connection Broker, but Tenant's Connection Broker is still required to authenticate users.

- **Disabled** — The Tenant is disabled in the Tenant Broker configuration. You can enable and disable Tenant objects as described below.

To see and modify Tenant properties, click **Tasks > Properties** (or right-click > **Properties**). The **Properties** dialog opens where you can view and modify the following properties:

- **Enable Tenant:** Enable or disable the Tenant object in the Tenant Broker.
- **Name:** The Tenant name (must be unique).
- **Public domain address:** The unique address that end users connect to from the outside (e.g. RAS.tenant.com, tenant1.MSP-FARM.com, etc.). See more in [Assign a public domain address](#).
- **Clients in gateway mode connect to published tenant resources by server IP:** When selected, clients will use the Tenant IP address instead of the DNS name. You can use this option when a Tenant farm does not share the same DNS provider as the Tenant Broker farm.
- **Do not show billing information:** (Only for Tenants joined with an invitation hash) When selected, billing information is not shown in the [Licensing](#) category of the Tenant.
- **Forward tenant sessions tunneled through gateway using server IP:** When a client session is forwarded to a server hosting published resources, either the server name (FQDN, hostname) or IP address can be used. When this option is selected (default) the IP address is used to forward the session internally. When the option is cleared, the configured host name is used.
- **Description:** An optional Tenant description. The Tenant description is a property that exists and can be viewed only in the Tenant Broker console.
- **Connection Brokers:** An IP address of one or more RAS Connection Brokers installed in the Tenant Farm. This is a read-only field.
- **Tenant invitation hash:** The hash that was used to join the Tenant to the Tenant Broker. This is a read-only field.
- **Automatically log out idle client connection after:** The time period after which an idle client connection should be logged out. For information on how to configure this property, see [Remote session settings](#).

Deleting a Tenant object

A Tenant object can be deleted any time. To delete an object, click **Tasks > Delete** (or right-click > **Delete**). This deletes the Tenant configuration from shared RAS Secure Gateways, so no RDP sessions can be established from the gateway to the deleted Tenant anymore. The Tenant's RAS Console will show the Tenant Broker status as "Join Broken" after this. To completely remove any references to the Tenant Broker, the Tenant admin needs to unjoin the [Tenant from the Tenant Broker](#).

Opening a Tenant console

As a Tenant Broker admin, you can open the Tenant console right from the Tenant Broker console. To do so, navigate to **Farm > Tenants**, right-click a Tenant and choose **Open tenant console**. This will open a new instance of the RAS Console and will prompt you to log in to the Tenant Farm. Please note that the Tenant Farm must be configured to allow remote console connections, which means that the corresponding port must be open on the Tenant Connection Broker and you need to know the credentials of the Tenant Farm administrator.

When you log in to a Tenant from the Tenant Broker console, the Tenant Farm is automatically added to the **Location** drop-down list (in the upper left-hand corner of the RAS Console window), so you can connect to the Tenant again by simply selecting it in the **Location** list.

Shared Gateways

All RAS Secure Gateways that exist in the Tenant Broker are shared among Tenants. For the most part, shared gateways operate similarly to standard RAS Secure Gateways but there are differences, which are described below.

Tunneling policies

Tunneling policies are allowed. Tunneled connections are sent to a Tenant Farm mapped to the public address used. The policies however are limited to "None" and "All servers in Site".

WYSE

WYSE is not supported.

Session counters

For each shared gateway, a session counter is displayed in the Tenant Broker console. To see how many sessions a gateway is running, navigate to **Farm > Site** and examine the **Sessions** column in the **Gateways** section.

Client connection routing

Each shared gateway is aware of a configuration of each existing Tenant and is able to route client connections to a correct RAS Connection Broker running in a Tenant Farm. The routing works as follows:

1. A new client connection is established.
2. A shared gateway determines which Tenant the client belongs to based on the Tenant configuration.
3. The correct RAS Connection Broker in the Tenant Farm is selected for this connection.
4. Two-factor authentication and application listing requests are forwarded to the selected RAS Connection Broker. All subsequent client operations are also carried out using that Connection Broker. See also [User authentication](#).

Shared gateway maintenance

When you need to take a shared RAS Secure Gateway offline for maintenance, you can do it the same way it's done in a traditional Parallels RAS Farm. You disable the gateway and wait for active sessions to drain. To see the number of active sessions for a gateway, navigate to **Farm > Site**. The session count is displayed in the **Sessions** column.

You can safely take shared Gateways offline. Parallels Clients will reconnect to the same sessions automatically.

Third-party network load balancers

Third-party network load balancers are possible to use with shared RAS Secure Gateways the same way they are used with traditional (not shared) RAS Secure Gateways.

Web Client and Themes

One of the important features of the RAS multi-tenant architecture is the ability to use a shared User Portal (which is a part of the RAS Secure Gateway) for all browser-based client connections, while at the same time using tenant-specific Web Client Themes defined on the Tenant side. This allows Service Providers to implement white-labeling by creating unique custom Themes for individual Tenants.

An Web Client Theme is created in a Tenant Farm. The user interface and the functionality remain the same as with a traditional Parallels RAS Farm. When Tenants join the Tenant Broker, Themes are pulled from the Tenant's RAS Connection Broker and added to the configuration of every shared RAS Secure Gateway.

When connecting to a Tenant Farm via the Web Client, a user must enter the Tenant public domain address (not the gateway address). The correct Theme is then used by the shared gateway as follows:

- The default Tenant Theme is used when the user enters the default URL: `https://<public-tenant-address>`.
- A specific Theme is used when the user adds the Theme name after the Tenant address: `https://<public-tenant-address>/<Theme-name>`

Web Client configuration

The Web Client is normally configured on the RAS Secure Gateway level (the **User Portal** tab in the gateway **Properties** dialog). When configuring a Theme, you have the ability to override the gateway settings by specifying them for a specific Theme in a Tenant Farm. To do so, in the Tenant RAS Console, select a Theme, open its properties and then select the **Gateway** category where you can specify your own settings. For more information, see [Web Client Theme settings > Secure Gateway](#).

Viewing Tenant Themes in Tenant Broker

If you are a Tenant Broker administrator, you can view Tenant Themes right in the Tenant Broker console:

1. In the Tenant Broker console, select **Farm > Tenants**.
2. Select a Tenant and click **Tasks > View tenant themes**.
3. The dialog opens where you can view Themes that were pulled from the Tenant and added to the configuration of every RAS Secure Gateway in the Tenant Broker.

Use this functionality to ensure that all Tenant Themes are properly synchronized on the Tenant Broker side, so when users connect to a Tenant through Tenant Broker, the appropriate Theme is used.

Monitoring Tenants

Parallels RAS Performance Monitor is a RAS component used to analyze Parallels RAS deployment bottlenecks and resource usage. RAS Performance Monitor can be used to monitor Tenants and view their performance metrics right from the Tenant Broker console.

To configure RAS Performance Monitor to collect information about Tenants:

1. Install RAS Performance Monitor as described in [Parallels RAS Performance Monitor](#) chapter.
2. Log in to the Tenant Broker console.
3. In the console, navigate to **Administration > Reporting**.
4. Select the **Enable RAS Performance Monitor** option (the **RAS Performance Monitor configuration** section).
5. In the **Server** and **Port** fields, specify the name or IP address of the server where you have RAS Performance Monitor installed.
6. Click **Apply**.
7. Now open a Tenant console and repeat steps 3 to 6 above, so both Tenant Broker and the Tenant are configured to use the same RAS Performance Monitor. This way, when Tenant(s) report their performance data to the RAS Performance Monitor, it can be viewed on the Tenant Broker side.

Tenants will report statistics to RAS Performance Monitor and you can view these statistics in the Tenant Broker console. When viewing the data in the RAS

Performance Monitor dashboard, you can switch between Farms and sites, so you can select a specific Tenant and view its performance metrics.

Tenant Broker compatibility and updates

When updating Parallels RAS to a newer version, the following rules apply to the RAS Multi-Tenant architecture:

- Parallels RAS Tenant Broker supports Tenants up to two major RAS versions older. For example, if you upgrade RAS Tenant Broker from RAS 17 to RAS 18 (or the next major version when it becomes available), it will support Tenants running Parallels RAS 17.
- When doing updates, you should first update the Tenant Broker, so your RAS Multi-Tenant installation remains fully operational. You can update Tenants later during their own maintenance window.

Upgrading from an older RAS version

If you have an existing Farm running RAS v16.x and would like to join it as a Tenant to Tenant Broker, follow these steps:

1. Upgrade the Farm to RAS 17.1 (or newer).
2. To join your Farm as a Tenant to Tenant Broker, follow the instructions in the [Deploying Tenant Broker and Tenants](#) section.
3. Once the Farm is joined, you can remove local RAS Secure Gateways if you are not planning on using them for local connections. See [Implementation overview](#) for additional info.

Configuring notifications

System event notifications are used to alert RAS administrators about system events via email. You can configure system event notifications in **Farm > Site > Settings > Notifications**. For the complete description of this functionality, please see [System event notifications](#). The rest of this section describes notifications, which are specifics to Tenant Broker and Tenants.

Tenant event notifications

As a Tenant Broker administrator, you can receive notifications about the following Tenant events:

- **New Tenant enrollment.** Triggers when a new Tenant joins the Tenant Broker.
- **Tenant unjoins the broker.** Triggers when a registered Tenant unjoins the Tenant Broker.
- **Tenant status alert.** Triggers when the RAS Connection Broker in a Tenant Farm goes offline.

When a Tenant event occurs, the Tenant Broker administrator receives an email containing the following information (depending on the event type):

- Tenant name.
- Tenant Broker name.
- Tenant enrollment method (invitation hash or secret key).
- Tenant status.
- Date.

To enable Tenant notifications, do the following:

1. Log in to the Tenant Broker.
2. In the RAS Console, navigate to **Farm > Site > Settings > Notifications**.
3. In the **Notification handlers** section, click **Tasks > New > Tenant events**.
4. In the **Tenant Events Notification Handler Properties** dialog, specify the following:
 - On the **General** tab, select the **Send email to RAS administrators** option and specify one or more email addresses separated by a semicolon.
 - On the **Settings** tab, either select the **Use the default settings** option (to use Site defaults) or clear it and specify your own settings.
5. Click **OK** to save your settings and close the dialog.

Tenant Broker event notifications

A Tenant Farm administrator can receive notifications when the Tenant Broker becomes unavailable. This usually happens when the RAS Connection Broker in the Tenant Broker goes offline. The notification handler is configured the same way

as described above, but this one is configured in the Tenant Farm (not the Tenant Broker).

Common event notifications

In addition to the **Tenant events** handler, you can configure notifications for common events, such CPU utilization, Memory utilization, RAS Agent events, etc. The only limitation here when it comes to Tenant Broker is the Tenant Broker has a limited set of system events for which notification handlers can be configured (see the list of available handlers below). This is due to the fact that the Tenant Broker doesn't have RD Sessions Hosts, Provider, licensing limits, published resources, etc. A Tenant Farm has the complete set of notification handlers, so the Tenant admin can configure any of them.

The following notification handlers are available in the Tenant Broker:

- CPU utilization
- Memory utilization
- Number of Secure Gateway tunneled sessions
- Number of failed Secure Gateway tunneled sessions
- RAS Agent events

For additional information, please see [System event notifications](#).

Communication ports

Tenant Broker and Tenants communicate with each other using the following ports:

- Tenant Connection Broker > Tenant Broker Connection Broker: port 20003
- Tenant Broker Gateway > Tenant Broker Connection Broker: port 20002
- Tenant Broker Gateway > Tenant Connection Broker: port 20002
- Tenant Broker Gateway > Servers hosting published resources: port 3389

These are standard RAS ports, which are also described in the **Port reference** section.

SAML SSO Authentication

Parallels RAS 17.1 and newer support the Security Assertion Markup Language (SAML) authentication mechanism. SAML is an XML-based authentication that

provides single sign-on (SSO) capability between different organizations by allowing user authentication without sharing the local identity database.

As part of the SAML SSO process, the new RAS Enrollment Server communicates with Microsoft Certificate Authority (CA) to request, enroll, and manage digital certificates on behalf of the user to complete authentication without requiring the users to put in their Active Directory credentials. Service providers and enterprises with multiple subsidiaries don't have to maintain their own internal Identity Management solutions or complex domains/forest trusts. Integrating with third-party Identity Providers allows customers and partners to provide end users with a true SSO experience.

Introduction

Security Assertion Markup Language (SAML) is an XML-based authentication that provides single sign-on (SSO) capability between different organizations by allowing user authentication without sharing the local identity database. Parallels RAS 17.1 and newer support the SAML authentication mechanism.

SAML (2.0) SSO was introduced in Parallels RAS 17.1 supporting HTML5 initiated authentication using Web Client or Parallels Client for Windows. Parallels RAS 18 extends the client support for initiating SAML authentication using the default OS browser or the browser embedded in Parallels Client.


As part of the SAML SSO process, the new RAS Enrollment Server communicates with Microsoft Certificate Authority (CA) to request, enroll, and manage digital certificates on behalf of the user to complete authentication without requiring the users to put in their Active Directory credentials. Service providers and enterprises with multiple subsidiaries don't have to maintain their own internal Identity Management solutions or complex domains/forest trusts. Integrating with third-party Identity Providers allows customers and partners to provide end users with a true SSO experience.

Supported delivery options are:

- Web Client
- Web Client initiated SAML for Windows
- Web Client initiated SAML for Mac and Linux
- Web Client initiated SAML for Android, iOS and iPadOS
- Parallels Client for Windows initiated SAML authentication
- Parallels Client for Mac initiated SAML authentication

- Parallels Client for Linux initiated SAML authentication
- Parallels Client for iOS/iPadOS initiated SAML authentication
- Parallels Client for Android initiated SAML authentication

The below high-level logical diagram depicts SAML authentication and login process within a Parallels RAS environment:



The SAML authentication and login steps on the diagram above are:

1. RAS Secure Gateway redirects the Parallels Client login request to the IdP site.
2. The user authenticates with IdP.
3. IdP redirects the user to the RAS Secure Gateway with the SAML Assertion.
4. The user is authenticated using the SAML Assertion and the user is logged in.
5. The list of the available RAS published resources is retrieved.
6. The user chooses a published resource and launches it from Parallels Client.
7. The launch request from the user is sent to the server side and the resource is started on the available server.
8. A Parallels RAS session is established.
9. User certificate is processed:
 - Certificate is requested.
 - Certificate is created.
 - Encryption is preformed using the certificate.
10. Smartcard logon.

System requirements

RAS Enrollment Server

- Windows Server 2016 up to Windows Server 2025

RD Session Hosts

- Windows Server 2012 R2 (x64 bit versions) up to Windows Server 2025

Desktop operating systems (guest VMs and Remote PCs)

- Windows 7 up to Windows 11

Please note that 32-bit operating systems are not supported.

Parallels Client

- Parallels Client version 18 or 19 is required.
- Supported platforms include Windows, Mac, Linux, iOS, and Android.

SAML basics

Security Assertion Markup Language (SAML) is a standard for exchanging authentication information between identity and service providers. SAML authentication is a single sign-on mechanism where a centralized identity provider (IdP) performs user authentication, while the service provider (SP) only makes access control decisions based on the results of authentication.

The main benefits of using SAML authentication are as follows:

- Service providers don't need to maintain their own user databases. User information is stored in a centralized database on the identity provider side. If a user has to be added or removed, it only needs to be done in a single database.

- Service providers don't need to validate users themselves, so there's no need for a secure authentication and authorization implementation on the provider's side.
- Single sign-on means that a user has to log in once. All subsequent sign-ons (when a user launches a different application) are automatic.
- Users don't have to type in credentials when signing in.
- Users don't have to remember and renew passwords.
- No weak passwords.

The single sign-on process

SAML single sign-on can be initiated on the service provider side or on the identity provider side. The two scenarios are outlined below.

The SAML single sign-on process initiated on the service provider side consists of the following steps:

1. A user opens Parallels Client (one of the [supported versions](#)) and connects to the service provider.
2. The service provider sends a message to the identity provider, asking to authenticate the user.
3. The identity provider asks the user for a username and password.
4. If the user credentials are correct, an authentication response (assertion) is sent to the client and then passed to the service provider. The response contains a message that the user has logged in successfully; the identity provider signs the assertion.
5. The user is presented with the published applications list. When the user launches an application, there's no prompt for credentials.

Single sign-on can also be initiated on the identity provider side, in which case the basic steps are the following:

1. A user logs in to identity provider via a web browser and is presented with a list of enterprise applications, including Parallels RAS.
2. Once Parallels RAS is selected, the assertion is sent to the client, then passed to the service provider configured for Parallels RAS.
3. Users are presented with the RAS published applications list.
4. When the user launches an application, there is no prompt for credentials.

SAML configuration

Prerequisites

To configure SAML in Parallels RAS, you need the following:

1. Microsoft Active Directory with the following two user accounts present:
 - **Enrollment agent user**: used to enroll certificates through RAS Enrollment Server (ES) on behalf of the authenticated user.
 - **NLA User**: used to initiate the NLA connection with RD Session Hosts and/or VDI guests.

See [Active Directory user account configuration](#) for required permissions and delegations. Note that Azure Active Directory Domain Services (AADDS) are not supported to be used with SAML SSO.

2. Microsoft Enterprise Certification Authority (CA) including the following templates:
 - Enrollment Agent Certificate Template
 - Smartcard Logon Certificate Template
3. Third-party Identity Provider (IdP) such as Azure, Okta, Ping Identity, Gemalto SafeNet, and others. This is where the user accounts will reside. User accounts in IdP must be synchronized with the Microsoft Active Directory environment. Please consult with the provider on how to properly synchronize users.
4. Domain Controllers must have Domain Controller certificates. The certificates on the Domain Controllers must support smart card authentication. Certificates are created using the Microsoft CA certificate template named Domain Controller Authentication. Manually created Domain Controller certificates might not work. If you get an error “Request Not Supported”, you may need to recreate Domain Controller certificates. Make sure RD Session Hosts and VDIs have the root certificate issued by the CA in the Trusted Root Certification Authorities store.
5. A Parallels RAS Farm with RD Session Host and/or VDI workloads (running on 64-bit OS).

6. For security reasons, the RAS Enrollment Server is recommended to be installed on a dedicated host. The host should be a standalone server that does not have any other components and roles installed.
7. Both SAML and RAS Enrollment Server configurations are Site-specific settings within the RAS environment. RAS administrators must have “Allow viewing of site information” and “Allow site changes” permissions delegated.

Note: Prerequisite knowledge of Microsoft Active Directory and Group Policy configuration is required for some of the above tasks.

Azure Active Directory Domain Services (AADDs) and Azure Virtual Desktop access are not currently supported with Parallels RAS SAML SSO.

IdP side configuration

On the identity provider side, you need to do the following:

1. Log in to your preferred IdP platform and create a generic or RAS specific SAML-based application to be used with the Parallels RAS environment.
2. Configure the application and take note of the following configuration properties to be added in Parallels RAS later:
 - **Entity ID**
 - **Logon URL**
 - **Logout URL**
 - **Certificate (base64)**
3. Alternatively, you can export a metadata file to be imported in Parallels RAS. For additional help, see [SAML integration examples and tips](#).

SP side configuration (RAS side)

On the service provider side (the Parallels RAS side), you need to enable Web (SAML) authentication and add the identity provider to the RAS Farm.

Enable Web (SAML) authentication

1. In the RAS Console, navigate to **Connection > Authentication**.
2. In the **Allowed authentication types** section, select the **Web (SAML)** option.

Adding an IdP to the RAS Farm

To add an IdP:

1. In the RAS Console, navigate to **Connection > SAML**. If the tab page is disabled, make sure you enabled Web (SAML). See above.
2. Click **Tasks > Add**.
3. In the **Add Identity Provider** wizard, specify a provider name.
4. In the **Use with Theme** drop-down list, select a [Theme](#) to which the IdP will be assigned. If you don't have a specific Theme yet, you can use the default Theme or you can select "<not used>" and assign a Theme later. Note that there can be multiple IdPs configured in the same RAS Farm. However, at this time, one IdP can be assigned to one Theme.
5. Select one of the following methods that the wizard will use to obtain the IdP information:
 - **Import published IdP metadata:** Import from an XML document published on the Internet. Specify the document URL taken from the IdP side configuration.
 - **Import IdP metadata from file:** Import from a local XML file downloaded from the IdP application. Specify the file name and path in the field provided.
 - **Manually enter the IdP information:** Select this option and then enter the information manually on the next wizard page.
6. Click **Next**.
7. If the configuration was imported in the previous step, the next page will be populated with data obtained from the XML file. If you've selected to enter the IdP data manually, you'll have to enter the values yourself:
 - **IdP entity ID:** Identity provider entity ID.
 - **IdP certificate:** Identity provider certificate data. To populate this field, you need to download the certificate from the IdP side, then open the downloaded file, copy its contents and paste it into this field.
 - **Logon URL:** Logon URL.
 - **Logout URL:** Logout URL.

Select the **Allow unencrypted assertion** option if needed.

Note: By default, the **Allow unencrypted assertion** option is disabled. Ensure that the IdP configuration is set to encrypt assertion or change the default setting within the RAS configuration.

1. At this point, you can configure service provider (SP) settings to be imported on the IdP side (IdP portal). You can do it now or you can do it later. To do it now, follow the steps below. To do it later, click **Finish** and then, when needed, open the identify provider object properties, select the **SP** tab and do the same steps as described below.
2. To configure SP settings, click the **Service provider information** button.
3. In the dialog that opens, enter the host address. The IdP will redirect to this address, which should be accessible from the end user browser.
4. The other fields including **SP Entity ID**, **Reply URL**, **Logon URL** and **Logout URL** are prepopulated based on the host address. The SP Certificate is autogenerated.
5. Next step is to complete the IdP configuration based on the values above. These values can be manually copied or exported as a metadata file (XML). Click the **Export SP metadata to file** link. Save the metadata as an XML file. Import the XML file into your IdP.
6. Close the dialog and click **Finish**.

Configuring user account attributes

When user authentication is performed by the IdP, user account attributes in Active Directory are compared with the matching attributes in the IdP user database. You can configure which attributes should be used for comparison as described below.

The following table lists available attributes:

RAS name	SAML name *	AD name	Description
UserPrincipalName	NameID	userPrincipalName	User Principal Name (UPN) is the name of a system user in an email address format.
Immutable ID	ImmutableID	objectGUID	A Universally Unique Identifier.
SID	SID	objectSid	An ObjectSID includes a domain prefix identifier that uniquely

RAS name	SAML name *	AD name	Description
			identifies the domain and a Relative Identifier (RID) that uniquely identifies the security principal within the domain.
sAMAccountName	sAMAccountName	sAMAccountName	The sAMAccountName attribute is a logon name used to support clients and servers from previous version of Windows, such as Windows NT 4.0 and others.
Custom	Email	Mail	A custom attribute that can be used to allow any SAML attribute name to match any AD attribute value. By default, it is the email address.

* The attributes in the **SAML name** column are editable and can be customized based on the IdP that you are using.

To configure attributes:

1. In the RAS Console, right-click an IdP that you've added in previous steps.
2. In the IdP **Properties** dialog, select the **Attributes** tab. On this tab, you can select or clear the attributes to be used for comparison or create custom ones:
 - Attributes that are selected will be compared for a match.
 - The names of all of the preconfigured SAML attributes (the IdP side) can be modified to match the AD attributes as required.
 - The custom attribute can be used to allow any SAML attribute name to match any AD attribute value. By default, it is the email address.
3. Configure and enable the desired attributes as needed based on the attributes configured on the IdP side.
4. Click **OK** to close the dialog.

Note 1: Multiple attributes are used in the presented order. If an attribute fails, the next configured attribute is used. Only one attribute is used at a time (in either/or fashion).

Note 2: If multiple AD users are configured with the same AD attribute value, user matching will fail. For example, if the email attribute is chosen and different AD users have the same email address, attribute matching between IdP account and AD User account will not be successful.

endhint %}

Attributes configuration tips

- When possible, use automation for user synchronization (such as Microsoft Azure AD Connect for Azure IdP configuration) between your Active Directory and the IdP to minimize user identity management overhead.
- Choose a user identification attribute that is unique to your environment, such as the User Principal Name (UPN) or Immutable ID (ObjectGuid) when possible. Alternatively, you can use other unique identifiers such as email address. In this case make sure that the **Email address** field in the user object in the AD is configured. If you use Microsoft Exchange Server, use the **Exchange Addresses** tab and Exchange policies.
- If using UPN as an attribute, you can also configure alternative UPN suffixes. This can be done from Active Directory Domains and Trusts (select root > right-click to open the **Properties** dialog). Once a new alternative UPN suffix is created, you can change the UPN on the user object properties from Active Directory Users and Computers.

Adding an account picture

For additional personalization, you can add a custom account picture which will be shown on the Windows logon screen during the user login when using Single Sign-On. This can be done as described in <https://kb.parallels.com/en/129028>.

Active Directory user account configuration

The enrollment agent user and NLA user must be created in Microsoft Active Directory. The following describes how to create these users.

Enrollment agent user account

The enrollment agent user account is required to enroll certificates through RAS Enrollment Server on behalf of the authenticated user. Please note that the enrollment agent user requires logon privileges on the machine where RAS Enrollment Server Agent is installed.

NLA user account

The NLA User is needed to initiate the NLA connection with RD Session Hosts and/or VDI guests. Please note that the NLA user requires log on privileges to the session host.

The NLA User must be a member of the Remote Desktop Users group and be granted the **Allow log on through Remote Desktop Services** permission. At the same time the NLA User must be prohibited to logon via Remote Desktop Services.

To exclude the NLA User account, it must be assigned the **Deny log on through Remote Desktop Services** user right.

To achieve both goals, you can use local or domain GPOs (linked to OU or domain wide).

A restart of the device is not required for this policy setting to be effective. Any change to the user rights assignment for an account becomes effective the next time the owner of the account logs on.

Group Policy settings are applied through GPOs in the following order, which will overwrite settings on the local computer at the next Group Policy update:

1. Local policy settings
2. Site policy settings
3. Domain policy settings
4. OU policy settings

Create a new GPO or use Default Domain Policy GPO as follows:

1. Open the Group Policy Management Console (GPMC).
2. Open or create a GPO linked with the OU where the RDSH or VDI objects reside.

3. Navigate to **Computer Configuration > Windows Settings > Security Settings > Local Policies > User Rights Assignment** and open “Allow log on through Remote Desktop Services” option.
4. Choose to add User or Group..., add the NLA user and click **OK**.

Note: The option will override default settings (on workstation and servers: Administrators, Remote Desktop User; on domain controllers: Administrators) therefore do not forget to add the groups like local administrators group or domain admins group.

1. Navigate to **Computer Configuration > Windows Settings > Security Settings > Local Policies > User Rights Assignment** and open the **Deny log on through Remote Desktop Services** option.
2. Choose to add User or Group..., add the NLA user and click **OK**.

Configure certificate authority templates

Create an Enrollment Agent template

To create the Enrollment Agent template:

1. From the Certificate Authority server, launch the Certificate Authority management console (MMC) from Administrative Tools.
2. Expand the CA, right -click on the “Certificate Templates” folder and select **Manage**.
3. Right-click the Enrollment Agent template and choose **Duplicate Template**. The new template properties window opens. On the **General** tab, configure the following properties:
 - **Template display name:** PrlsEnrollmentAgent
 - **Template name:** PrlsEnrollmentAgent
 - **Validity period:** 2 years
 - **Renewal period:** 6 weeks
 - **Publish certificate in Active Directory:** ON
 - **Do not automatically re-enroll if a duplicate certificate exists in Active Directory:** OFF

Note: The display name can be any name you choose, however the template name must match the template name highlighted above.

Properties of New Template

Subject Name	Server	Issuance Requirements
Superseded Templates	Extensions	Security
Compatibility	General	Request Handling
Cryptography	Key Attestation	

Template display name:
PrisEnrollment Agent

Template name:
PrisEnrollmentAgent

Validity period: 2 years
Renewal period: 6 weeks

☒ Publish certificate in Active Directory
☐ Do not automatically reenroll if a duplicate certificate exists in Active Directory

OK Cancel Apply Help

1. Select the **Cryptography** tab and set the following values:
 - **Provider category:** Legacy Cryptographic Service Provider (read-only).
 - **Algorithm name:** Determined by CSP
 - **Minimum key size:** The desired minimum key size up to 4096 bits

In the section **Choose which cryptographic providers can be used for requests**, choose **Requests must use one of the following providers**. In the following list of providers, select your desired provider.

Properties of New Template [X]

Subject Name	Server	Issuance Requirements
Superseded Templates		Extensions
Security		
Compatibility	General	Request Handling
Cryptography		Key Attestation

Provider Category: Legacy Cryptographic Service Provider

Algorithm name: Determined by CSP

Minimum key size: 2048

Choose which cryptographic providers can be used for requests

☐ Requests can use any provider available on the subject's computer

☒ Requests must use one of the following providers:

Providers:

- ☒ Microsoft Strong Cryptographic Provider
- ☐ Microsoft Base Smart Card Crypto Provider
- ☐ Microsoft Enhanced Cryptographic Provider v1.0
- ☐ Microsoft Base Cryptographic Provider v1.0
- ☐ Microsoft Enhanced RSA and AES Cryptographic Provider


Request hash: Determined by CSP

☐ Use alternate signature format

OK Cancel Apply Help

1. Select the **Security** tab and do the following:

- Click **Add**.
- Add the enrollment agent user account.
- Allow (select) the “Read” and “Enroll” permission. Click **Apply** and **OK**.



Issue the certificate template

To issue the certificate template that you've created:

1. Run Certificate Authority again and right click on **Certificate Templates**, select new and click on **Certificate Template to Issue**.
2. Select the certificate template you've created in the previous steps (i.e. Prls Enrollment Agent) and click **OK**.
3. The certificate template should appear in the **Certificate Templates** list.

Note: After creating the Enrollment Agent template and the Smartcard Logon template (described later), you should restart the **Active Directory Certificate Services** service in Windows.

Create a smartcard logon certificate template

To create a smartcard logon certificate template:

1. From the Certificate Authority server, launch the Certificate Authority management console (MMC) from Administrative Tools.
2. Expand the CA, right-click on the “Certificate Templates” folder, and select **Manage**.
3. Right-click on the “Smartcard Logon” certificate template and then select **Duplicate**.
4. The new template properties open in the **General** tab. Type a template name in the text box. Note that the real name automatically appears in the second text box with no spaces. Remember this name. You will need it later to configure of SAML feature. The options on this tab should be configured as follows:
 - **Template display name:** PrlsSmartcardLogon
 - **Template name:** PrlsSmartcardLogon
 - **Validity period:** 1 year(s)
 - **Renewal period:** 6 weeks
 - **Publish certificate in Active Directory:** OFF
 - **Do not automatically re-enroll if a duplicate certificate exists in Active Directory:** OFF

Note: The display name can be any name you choose; however, the template name must match the template name highlighted above.

Properties of New Template

Subject Name	Server	Issuance Requirements
Superseded Templates		Extensions
Security		
Compatibility	General	Request Handling
	Cryptography	Key Attestation

Template display name:

Template name:

Validity period: years

Renewal period: weeks

☐ Publish certificate in Active Directory

☐ Do not automatically reenroll if a duplicate certificate exists in Active Directory

OK Cancel Apply Help

1. Select the **Cryptography** tab and set the following:

- **Provider category:** Legacy Cryptographic Service Provider (read-only).
- **Algorithm name:** Determined by CSP
- **Minimum key size:** The desired minimum key size, up to 4096 bits

In the section **Choose which cryptographic providers can be used for requests**, choose **Requests must use one of the following providers**. In the following list of providers, select your preferred provider.

When using SAML with EntraID, and receiving in-session Primary Refresh Tokens (PRT) is required, select **Microsoft Enhanced RSA and AES Cryptographic Provider**.

Properties of New Template

Subject Name | Server | Issuance Requirements

Superseded Templates | Extensions | Security

Compatibility | General | Request Handling | **Cryptography** | Key Attestation

Provider Category: Legacy Cryptographic Service Provider

Algorithm name: Determined by CSP

Minimum key size: 2048

Choose which cryptographic providers can be used for requests

☐ Requests can use any provider available on the subject's computer

☒ Requests must use one of the following providers:

Providers:

- ☒ Microsoft Enhanced RSA and AES Cryptographic Provider
- ☐ Microsoft Strong Cryptographic Provider
- ☐ Microsoft DH SChannel Cryptographic Provider
- ☐ Microsoft Enhanced Cryptographic Provider v1.0
- ☐ Microsoft Enhanced DSS and Diffie-Hellman Cryptographic Pr

Request hash: Determined by CSP

☐ Use alternate signature format

OK Cancel Apply Help

1. Select the **Issuance Requirements** tab and set the following:

- **CA certificate manager approval:** OFF
- **This number of authorized signatures:** 1
- **Policy type required in signature:** Application policy
- **Application policy:** Certificate Request Agent
- **Same criteria as for enrollment:** ON

Prls Smartcard Logon Properties ? X

Superseded Templates		Extensions		Security	Server
General	Compatibility	Request Handling	Cryptography	Key Attestation	

Subject Name Issuance Requirements

Require the following for enrollment:

☐ CA certificate manager approval

☒ This number of authorized signatures:

If you require more than one signature, autoenrollment is not allowed.

Policy type required in signature:

Application policy

Application policy:

Certificate Request Agent

Issuance policies:

Add...

Remove

Require the following for reenrollment:

☒ Same criteria as for enrollment

☐ Valid existing certificate

☐ Allow key based renewal (*)


Requires subject information to be provided within the certificate request.

* Control is disabled due to [compatibility settings](#).

OK Cancel Apply Help

1. Select the **Security** tab and do the following:

- Click **Add**.
- Add the enrollment agent user account.
- Allow (select) the “Read” and “Enroll” permissions. Click **Apply** and **OK**.



Issue the certificate template

To issue the certificate template that you've created:

1. Run Certificate Authority again and right click on **Certificate Templates**, select **New**, and click on **Certificate Template to Issue**.
2. Select the certificate template you created in the previous steps (i.e., **Pr1s Smarcard Logon**), then click **OK**.
3. The certificate template should appear in the **Certificate Templates** list.

Note: After creating the Smartcard Logon template and the Enrollment Agent template (described earlier), you should restart the **Active Directory Certificate Services** service in Windows.

RAS Enrollment Server configuration

RAS Enrollment Server communicates with Microsoft Certificate Authority (CA) to request, enroll, and manage digital certificates on behalf of a user for SSO authentication in the Parallels RAS environment.

Note: For security reasons, RAS Enrollment Server should be installed on a secure, dedicated server similar to an Active Directory Domain Controller or Certificate Authority with no other Parallels RAS components installed. `endhint %}`

Set up and configure RAS Enrollment Server

You can remotely install the RAS Enrollment Server Agent on a specified server from the RAS Console. You can also install the Agent by running the standard RAS installer on the desired server.

To remotely install the RAS Enrollment Server:

1. In the RAS Console, navigate to **Farm > Site > Enrollment servers**.
2. Click **Tasks > Add**.
3. Specify the FQDN or IP address of the server where you want the RAS Enrollment Server Agent to be installed.
4. Click **Next**.
5. In the **Enrollment Server Agent Information** dialog, click **Install** and follow the onscreen instructions.

To install the RAS Enrollment Server using the Parallels RAS installer:

1. Run the Parallels RAS installer on the server where you want the RAS Enrollment Server Agent installed.
2. On the **Select Installation Type** page, select **Custom** and click **Next**.
3. Clear all other components and select the Parallels RAS Enrollment Server component.
4. Click **Next** and follow the onscreen instructions.
5. Once the RAS Enrollment Server is installed, open the RAS Console and navigate to **Farm > Site > Enrollment servers**.
6. Click **Tasks > Add**.
7. Enter the Enrollment Server FQDN or IP address and click **Next**.
8. Follow the onscreen instructions to add the server to the Farm.

Obtain and copy the registration key

If you perform a manual installation using the RAS installer, it is necessary to place a registration key file on the Enrollment Server host. This step is not required if the RAS Enrollment Server Agent was remotely deployed from the RAS Console.

First, you need to obtain the registration key file as follows:

1. Open the RAS Console and navigate to **Farm > Site > Enrollment servers**.
2. Click **Tasks > Export registration key**.
3. Save the key to a file named *registration.crt*.

Once you have the registration.crt file, copy it to the following folder on the server where you have the RAS Enrollment Server installed, by default in the following path:

C:\Program Files (x86)\Parallels\ApplicationServer\x64

Note: It is mandatory for the registration key file to be named “registration.crt”.
endhint %}

Configure AD integration

After you added the RAS Enrollment Server in the RAS Console, you need to configure AD integration for it as follows:

1. In the RAS Console, navigate to **Farm > Site > Enrollment Servers**.
2. Select the **AD Integration** tab.
3. In the **Certificate authority (CA)** section, select the **Type** to be AWS Private Certificate Authority or Microsoft Active Directory Certificate Services (AD CS).
4. Specify the configuration string of your Enterprise CA where the new certificate templates, (Prls Enrollment Agent and Prls Smartcard Logon) were created. This should be done in the following format:

CAhostname.domain\issuing CA name

Alternatively, you can click the [...] button to select a CA. (This is only available for Microsoft Active Directory Certificate Services). For configuration details, see [Configure certificate authority templates](#).

5. If **AWS Private Certificate Authority** was selected, configure the following:
 1. Specify the **Region** if it was not automatically extracted from the specified configuration string.
 2. Specify to **use an Access Key ID and Secret Access Key** configured on an existing EC2 provider (if available) or:
 3. Specify the **Access Key ID** and **Secret Access Key** to be used.
6. In the **Enrollment Agent** section, specify the Enrollment Agent username and password. For configuration details, see [Active Directory user account configuration](#).
7. In the **NLA user** section, specify the NLA username and password. For configuration details, see [Active Directory user account configuration](#).
8. Click the **Validate AD integration settings** button to make sure that the information you've entered is valid.

Using computer management tools

You can perform standard computer management tasks on a RAS Enrollment Server host right from the RAS Console. These include Remote Desktop Connection, PowerShell, Computer Management, Service Management, Event Viewer, IPconfig, Reboot, and others. To access the **Tools** menu, click **Tasks > Tools** and choose a desired tool. For requirements and usage information, see [Computer management tools](#).

RAS Enrollment Server high availability

For high availability, multiple Enrollment Servers (ESs) can be added to each Site. All enabled and verified ESs will be used in an active/active fashion. Upon user login, requests from workload VMs such as RD Session Hosts or VDIs are equally distributed among the available ESs. In case of failures on a particular ES, the next available ES is selected and the SAML SSO authentication process continues. Specifically required for manual deployment of multiple ESs, it is important to note that all ESs in the same site share the same registration key which is required to be deployed in the specified path as mentioned in the [RAS Enrollment Server configuration](#) section.

Note: Multiple ESs do not share a common certificate repository store and all certificates are segregated on each ES. This means that in case of multiple ESs, same user might have different certificates available on different ESs.

SAML integration examples and tips

For examples of how to integrate various Identity Providers with Parallels RAS, please read the **SAML SSO Authentication Examples** guide, which is available on the Parallels website at <https://www.parallels.com/products/ras/resources/>.

User account attributes

When user authentication is performed by the IdP, user account attributes in Active Directory and the IdP are compared with each other for a match. Attributes to be compared are configured on the IdP and in the RAS Console. For details, see [SP side configuration \(RAS side\)](#).

Security tip

For security reasons, it is advisable to configure enrollment agent restrictions for a CA to allow only the newly created Enrollment Agent User permissions to enroll certificates on behalf of the users. To do so, follow the steps below.

1. Open the Certification Authority snap-in, right-click the name of the CA, and then click **Properties**.
2. Click the **Enrollment Agents** tab, click **Restrict enrollment agents**, and click **OK** on the message that appears.
3. Under **Enrollment agents**, click **Add**, type the name of the Enrollment agent user created in the previous steps and then click **OK**. Click **Everyone**, and then click **Remove**.
4. Under **Certificate Templates**, click **Add**, select the templates that were created (PrIs Enrollment Agent and PrIs Smartcard Logon) and then click **OK**. When you have finished adding the names of certificate templates, click **<All>**, and then click **Remove**.
5. Under **Permissions**, click **Add**, type the names or groups, which are the users or group expected to login to the RAS environment using SAML, and then click **OK**. Click **Everyone**, and then click **Remove**.

6. If you want to block the enrollment agent from managing certificates for other users, computers, or groups, under **Permissions**, select this user, computer, or group, and then click **Deny**.
7. When you are finished configuring enrollment agent restrictions, click **OK** or **Apply**.

Note: The user or group that you applied enrollment agent restrictions to must have a valid enrollment agent certificate for the CA before they can act as an enrollment agent, whether restricted enrollment agent permissions have or have not been configured.

Parallels Client configuration

Parallel Web Client

No additional configuration is needed.

Parallels Client for Windows

See Parallels Client for Windows User's Guide.

Parallels Client for Mac

To configure Parallels Client for Mac for SAML SSO authentication:

1. Select a connection (or create a new one) and open its properties.
2. On the **Connection** tab, in the **Login** section, select **Web** as authentication type.
3. Select the **Advanced** tab and select or clear the **Use default OS browser** option (**Web authentication** section). If the option is selected, the SAML SSO login dialog will open in the default browser. If the option is cleared, the browser built into the Parallels Client will be used.
4. Close the dialog to save the connection properties.

When connecting to Parallels RAS, a dialog will open in a web browser asking the user to enter credentials to be verified by the identity provider. If the credentials are valid, the list of published applications will appear in the Parallels Client.

Parallels Client for Linux

To configure Parallels Client for Linux for SAML SSO authentication:

1. Select a connection (or create a new one) and open its properties.
2. On the **Connection** tab, in the **Log in** section, select **Web** as authentication type.
3. If you want to configure additional settings, do the following:
 1. Install the QtWebEngine library.
 2. Select the **Advanced Settings** tab and click the **Connection advanced settings** button.
 3. Select or clear the **Use default OS browser** option in the **Web Authentication** section. If the option is selected, the SAML SSO login dialog will open in the default browser. If the option is cleared, the browser built into the Parallels Client will be used.
 4. If you are using the built-in browser, select or clear the **Open browser window to complete log out** option in the **Web Authentication** section. If this option is selected, a URL will open to perform the logout from SAML. By default, this web page will not be displayed, but if you need to interact with the browser, you can enable this option.
4. Close the dialog to save the connection properties.

Parallels client policy configuration

In addition to specifying SAML SSO options directly on the client side, you can also set them via the Parallels Client Policy in the RAS Console. To do so:

1. In the RAS Console, select the **Policies** category.
2. Open the policy properties and navigate to **Session > Connection > Primary Connection**.
3. In the **Authentication type** drop-down list, select **Web**.
4. Navigate to **Session > Connection > Web authentication**.
5. Select or clear the **Use default OS browser** option. If the option is selected, the SAML SSO login dialog on the client side will open in the default browser. If the option is cleared, the browser built into the Parallels Client will be used.

Note: To launch SAML SSO login dialog with built-in browser while using Parallels RAS Console 19.3 or later, use Parallels RAS Client for Windows 19.3 or later

See also [Parallels Client configuration](#).

Test the SAML SSO deployment

When you have the SAML SSO authentication configured, you can test it as described below.

Service provider initiated authentication

1. Use a web browser to open the RAS Web Client (or use a platform-specific Parallels Client), specifying the Theme to which you assigned the identity provider.
2. Start a published application. Check that the application session was started successfully.
3. Create one more Theme, add one more IdP providers and then connect specifying the new Theme. Launch another application.

Identity provider initiated authentication

1. Use the web browser to connect to the identity provider portal.
2. Start a published application. Check that the application session was started successfully.

Error messages

Error messages

Error messages appear in the web browser when something goes wrong with SAML SSO authentication.

Pre HTML5 loading

Error message	Notes
Unable to parse SAML Assertion	There was an error while parsing and validating the SAML Assertion. Further details can be found in HTML5 Logs. Most common causes:

Error message	Notes
	SAML Response is not valid for this audience: The most probable cause for this issue is having wrong configuration on the IDP, especially the Entity ID URL. The entity ID URL in the assertion will not match with the Entity ID provided in the SP SAML settings. Expected 1 Assertion or 1 EncryptedAssertion; found 0: The Assertion / EncryptedAssertion tag was not found in the response. The Web Client will be expecting an encrypted assertion while the IDP is sending a non encrypted one. This can either be fixed by changing the IDP settings to send an encrypted assertion or tick the checkbox found in 'RAS Console > Connection > SAML > IDP Settings > Allow unencrypted assertion' SAML Response is not yet valid: This might happen if the time of the server where RAS Gateway is installed is incorrect, for instance 4 seconds behind. In this case the assert will be created before actually trying to parse it. SAML Response is no longer valid: This might happen if the time of the server where RAS Gateway is installed is incorrect. In case it's manually set in the future, assert might be seen as not valid anymore while trying to validate it.
SAML Assertion body is empty	SAML Assertion was not found in the response. Further details can be found in HTML5 Logs
Unable to create SAML logout request	

Error message	Notes
	There was an error while creating SAML logout request. Further details can be found in HTML5 Logs.
Unable to create SAML logout response	There was an error while creating logout response. Further details can be found in HTML5 Logs.

Post HTML5 loading

Error code	Error message	Notes
0x00000029	SAML IdP settings not found. IdP Id:'xxx'	Check the Identity Provider settings. Check if the IdP metadata are correctly imported.
0x0000002A	SAML IdP info keys loading failed. IdP Id:'xxx'	Check if the IdP certificate is present in the IdP settings.
0x0000002B	SAML Theme mismatch	Check if the theme is correctly set in the IdP settings.
0x0000002C	Logon using SAML failed. Error: 0x00001	See errors below
0x00000029	No Enrollment Sever available	Check Enrollment server(s) status
0x0000002A	Missing NLA User Configuration	Enter NLA User details
0x00000003	Logon using SAML failed. Error: Failed to match AD User. 0x00000006	Check if the Attributes settings are correct in the IdP properties.
0x00000003	Logon using SAML failed. Error: Failed to validate and decrypt the response. 0x00000009	Check if the IdP certificate is present in the IdP settings.
0x00000003	Logon using SAML failed. Error: Assertion not encrypted. 0x0000001C	Check if the IdP settings for the logon request are correct.
0x00000003	Logon using SAML failed. Error: Failed to decrypt the assertion. 0x0000001D	Check the SP certificate is correctly set in the IdP settings.
0x00000003	Logon using SAML failed. Error: Failed to verify assertion. 0x0000001F	Check if the IdP certificate is present in the IdP settings.

Once an application or desktop is launched

Error message	Description and reference
Invalid username or password	The user certificate is valid, but the domain controller did not accept it. Check the Kerberos logs on the domain controller.
The system could not log you on. Your credentials could not be verified.	Check connectivity with the domain controller and check that the appropriate certificates installed.
The request is not supported	The “Domain Controller” and “Domain Controller Authentication” certificates on Domain Controller require enrolling, even if they are already available.
The system could not log you on. The smartcard certificate used for authentication was not trusted.	The intermediate and root certificates are not installed on the machine where the error is shown. The CA root certificate and any intermediate certificates must be added to the “Trusted root certificates” in the local computer account.
You cannot logon because smart card logon is not supported for your account.	The user account has not been fully configured for smart card logon.
No valid smart card certificate could be found.	Check the configuration of the PrlsSmartcardCertificate. The extensions might not be set correctly, or the RSA key is less than 2048 bits.
Bad Request	Check the configuration of the PrlsSmartcardCertificate. The extensions might not be set correctly, or the RSA key is less than 2048 bits.

Parallels Web Client and User Portal

Parallels Web Client is a client application that runs in a web browser. Users can use Parallels Web Client to access User Portal that allows to view and launch remote applications and desktop from a web browser.

Compared to platform-specific Parallels Clients (Parallels Client for Windows, Parallels Client for iOS, etc.), Parallels Web Client does not require end users to install additional software on their computers or mobile devices. Feature-wise,

platform-specific Parallels Clients give users more options than Parallels Web Client. Nonetheless, Parallels Web Client is a fully-featured platform-independent client providing end users with an alternative method of working with remote resources published via Parallels RAS.

System requirements

Secure Gateway (hosts the User Portal and Web Client):

- Windows Server 2012 R2 or higher.

Client side:

- Any HTML5-enabled web browser, except Internet Explorer.

Configure Web Client

The Web Client is a part of RAS Secure Gateway. To be used by end users, the User Portal must be enabled and configured in the RAS Console as described in [Configure User Portal](#).

Session persistence based on a cookie

RAS Web Client session persistence is normally set by user's IP address (source addressing). If you can't use source addressing in your environment (e.g. your security policy doesn't allow it), you can use the Session Cookie to maintain persistence between a user and a server. To do so, you'll need to set up a load balancer that can use a session cookie for persistence. The cookie that you should use is ASP.NET_SessionId. If you are using a load balancer that doesn't use ASP.NET, you can specify a different cookie on the Web Requests tab of the RAS Secure Gateway Properties dialog. For more information, see [Web request load balancing](#).

Host header attack protection

You can enable host header attack protection for the User Portal URL. This security measure will ensure that the Host headers of the users' HTTP requests to User Portal cannot be changed in transit, and users who access User Portal via a

browser are always redirected to one of your Secure Gateways and not any other hosts.

To enable host header attack protection:

1. Navigate to **Farm > Farm > Tasks > Properties**.
2. Select the **Enable HTTP Host header attack protection** option.
3. (Optional) If you use additional hostnames or IPs for your Secure Gateways, you can add them to the list of the allowed addresses by selecting **Tasks > New** (or clicking the plus-sign icon) in the **Access addresses** section.

Note: The default hostnames and IP addresses of Secure Gateways and HALBs are added to the list automatically.

Configure Themes

Themes in Parallels RAS is a functionality that allows you to do the following:

- Allow access to a Theme to specified groups of users while configuring certain Theme properties that will apply to these groups. This functionality is supported by Parallels Client on all available platforms.
- Customize the appearance of User Portal, which enables you to implement custom branding of User Portal for different groups of users. Note that this functionality is only available for RAS Web Client and Parallels Client for Windows.

To manage Themes, in the Parallels RAS Console, navigate to **Farm > <Site> > Themes**. The **Themes** view in the right pane displays the available Themes. The list contains at least one default Theme. This Theme cannot be removed but you can customize it as needed. In addition to the default Theme, you can create your own Themes.

To create a new or modify an existing Theme:

- Click **Tasks > New Theme** (or click the **[+]** icon) to create a new Theme.
- Double-click an existing Theme (or right-click it and choose **Properties**).

The **Theme Properties** dialog opens. Use the dialog to create a new or modify an existing Theme. The instructions in the subsequent sections apply to both scenarios.

General settings

The Theme settings described below apply to Parallels Client for all available platforms.

General

Select **General** in the left pane and specify the following Theme properties:

- **Enable Theme:** Enable or disable the Theme (the default Theme cannot be disabled).
- **Name:** Specify a Theme name.
- **Description:** Specify an optional Theme description.
- **Path:** Specifies a postfix for the Theme login page URL. This field is populated automatically with the Theme name when you save it, but you can specify a name of your choice. The complete URL of the Theme login page is comprised of “https://<host-name>/” followed by the name specified in this field. For the explanation of what the <host-name> should be, please see [Web request load balancing](#).

For example, if you name the Theme “Theme-S1”, the complete URL is https://<host-name>/Theme-S1. When you save the Theme, the URL is displayed on the **Themes** tab in the RAS Console (the **User Portal URL** column).

Please note that the URL described above is the short version, which is easier to remember and use. The full version is:

https://<host-name>/userportal/?theme=<team-name>

Both the short and the long versions are equally valid.

Note: You can find the link for connecting to the Theme via browser in the **User Portal URL** field and via a native Parallels Client in the **Client connection** field.

Note: The following reserved words are not allowed as theme names: userportal, rashtml5gateway, 2xhtml5gateway, getclintip, urlschemaeditor, and rasgeolocatorinfo.

Access settings

The Theme settings described below apply to Parallels Client for all available platforms.

- **Override authentication domain:** Allows you to specify a domain name that will be passed to Parallels Client, so users don't have to enter it manually. This setting overrides the domain name setting in **Connection > Authentication**. For more information, see [Allowing users to change domain password](#).
- **Limit access to this Theme to members of these Active Directory groups:** If this option is cleared, any Parallels RAS user can access the Theme if they know its URL. To limit access to a particular group (or groups), select this option and then click **Tasks > Add** (or click the **[+]** icon) and select the desired group(s).
- **MFA provider:** Select an MFA provider for the Theme.
- **(Default theme only) Allow gateway tunneling connections:** Allow gateway tunneling connections for the default theme.

Message settings

Select **Messages** in the left pane and specify a post-logon message (up to 500 characters). The post-logon message appears as a message box immediately after the user successfully logs in. The message can be overridden for Web Client and Windows client individually (see **Messages** for each client in the subsequent sections).

Client Public IP Detection

The Theme settings described below apply to Parallels Client for all supported platforms.

When configured, a Parallels client will use the selected service to retrieve the endpoint's public IP and send it to the Parallels RAS Connection Broker along with other connection information. The received endpoint's public IP is stored in the Parallels RAS variable %CLIENTPUBLICIP%, which can be configured as attribute values sent to a RADIUS provider.

The following settings can be configured:

- **Do not use client public IP detection:** The client's public IP will not be captured and sent to Parallels RAS Connection Broker.
- **Use Parallels client public IP detection service:** The client's public IP is captured using a Parallels-managed [public IP detection service](#) and sent to Parallels RAS Connection Broker.

Use the test button to test the Parallels-managed public IP detection service.

The public IP of the machine that runs the Parallels RAS console should be returned. Clients need to be able to access the Parallels-managed public IP detection service via HTTPS.

- **Use Custom client public IP detection service:** The client's public IP is captured using a custom public IP detection service that can be specified as the service URL.

Use the test button to test the specified public IP detection service.

The public IP of the machine that runs the Parallels RAS console should be returned. Clients need to be able to access the specified public IP detection service via HTTPS. The custom service must return a JSON response containing the public IP in one of the following fields: "ip", "IP", "ipString", or "ipAddress".

Example JSON response:

```
{  
  "ip": "84.104.46.194"  
}
```

Web Client settings

The **User Portal (Web Client)** category allows you to configure Theme settings for User Portal. These settings affect how the User Portal looks and behaves in a web browser.

Note: To see how your Web Client Theme looks, click the **Preview User Portal** button in the lower left-hand corner of the dialog at any time.

URLs

The **URLs** category is used to specify the Theme login page URL and add additional URLs to the User Portal page:

- **Show Parallels Client download URL.** If selected, users will see the **Download Client** link on the Web Client page, which can be used to download, install, and configure Parallels Client on users' computers.
- **Override download URL for branded Parallels Client (Windows):** Specifies a location from which your Windows users will download Parallels Client for Windows. By default, Parallels Client is downloaded from the Parallels web site. If you use a branded version of Parallels Client, you can specify its location in this field.
- **Footer URLs.** This option allows you to specify custom URLs that will be placed in the Web Client footer. To add a URL, click **Tasks > Add** and specify a URL, a text that will appear on the page footer, and a tooltip text. When entering similar URLs, you can duplicate an existing one by right-clicking it and choosing **Duplicate** (or select an entry and click the “duplicate” icon next to the [-] icon). If you've added multiple URLs, you can reorder them by clicking the up or down arrow icons or selecting **Up** or **Down** items in the **Tasks** menu. The URLs will appear in the footer in the order listed (you can click the **Preview HTML4 Theme** button to see how it looks).

Branding

The **Branding** category allows you to configure the branding of User Portal.

The following properties can be customized:

- **Webpage title:** Specifies the title that appears on the webpage. You can type any title you like.
- **Login to:** Specifies a name that will appear in the User Portal login dialog. For example, if you type “ABC” here, the login page will say, “Log in to ABC”. There are two predefined variables that you can use here: %FARM% (the actual Farm name; this is the default value) and %SITE% (the Licensing Site name).
- **Company logo:** Displays the image which is displayed on the User Portal page header. To change the image, select browse and then specify the

image file. Note that changing the logo image also removes the default **Remote Application Server** part from the page header.

- **Favicon icon:** Displays the currently set favicon icon. To change the icon, click **Browse** and select an icon file.

Colors

The **Colors** category allows you to customize the appearance of the User Portal pages.

The following properties can be customized:

- **Header options:** Configures the appearance of the header.
- **Work area image:** Configures the background image for User Portal. You can choose one of the preset images or upload an image from a URL or a local computer. The supported file types are JPG, PNG, and SVG. The size of the image file must be under 10 MB.

Note: If you want to upload an image from a URL, make sure that cross-origin resource sharing is enabled at **Secure Gateway > User Portal > Properties**.

- **Colors:** Specifies the desired colors for various User Portal elements, such as header, footer, work area, buttons, etc.

Language bar

Select languages that will appear in the language selector on the User Portal page. The selector appears as a language flag icon on the page header to the right of the user name.

Messages

On this pane, you can specify pre-logout and post-logout messages:

- A pre-logout message will appear on the **Log in** page.
- To override the default post-logout message (see **Messages** in the beginning of this topic), select the **Override post-logout message** option and enter a message.

The messages must be 500 characters or less.

Input prompt

Input prompts specified here will appear on the login page to help users enter their username and password correctly in the fields provided. For example, the default **user@domain** login prompt will appear as a light gray text in the login field, hinting the user that they should enter their name in the UPN format. Predefined input prompts are provided for every supported language. You can specify your own prompts if needed.

Gateway

The **Secure Gateway** category can be used to override the default User Portal settings, which are configured for the RAS Secure Gateway. Normally, you shouldn't override the gateway settings if you are running a traditional Parallels RAS Farm and using a single Theme in a Site. Scenarios when overriding the settings may be needed include the following:

- You have multiple Themes for different groups of users and would like different Themes to behave differently in terms of application launching methods and restrictions.
- You are using RAS multi-tenant architecture where RAS Secure Gateways are running in Tenant Broker and are shared by Tenants, which are separate Farms. Themes in this kind of deployment are defined on the Tenant level, so each Tenant can have its own Web Client look and feel. Since gateways are shared by Tenants, it is logical to configure these settings on a Theme level, which is exactly what the **Secure Gateway** category allows you to do. For the complete description of what Tenant Broker and Tenants are, please read the [RAS Multi-Tenant Architecture](#) chapter.

To override the RAS Secure Gateway settings, select the **Override Secure Gateway settings for the Theme** option and then specify your own settings. For the description on how to configure these settings, see [Configure Web Client](#).

Legal policies

Cookie consent

Select the **Enable cookie consent** option to show a notification about User Portal cookie policy to users on first time use. This provides users with information regarding the use of cookies and the option to accept.

End User License Agreement

Select the **Enable EULA** option to show the Parallels End User License Agreement (EULA) to users on first time use. User intervention is required to read and accept the agreement to complete the login process.

Parallels Client for Windows settings

Panes under the **Windows client** heading allow you to configure Theme settings for Parallels Client for Windows. By configuring a Windows client Theme, you can make the client appear to end users as your organization requires.

Branding

On the **Branding** pane, specify the following:

- **Company name:** Used to create the Start menu hierarchy: Start Company Name App Name.
- **Application name:** Displayed in the app caption and the Start menu entry name.
- **Connection banner:** Displayed when a connection is being established.
- **Application icon:** The application icon used for the Start menu and by the main app window.

Messages

To override the default post-logout message, select the **Override post-logout message** option and enter a message.

Custom menu

The **Custom Menu** pane allows you add a menu item to the **Help** menu in white-labeled Parallels Client for Windows. For example, if you enter “&Notepad” in the **Menu item** field and “notepad.exe” in the **Command** field, a new menu item will appear under the **Help** menu in every white-labeled Parallels Client for Windows connecting to this Farm. The item will be named **Notepad** (with the “N” being the shortcut) and it will open the Notepad.exe application when clicked. The **Command** field can contain an executable name, a URL, or any other command that can be properly executed on a Windows machine. For instance, you can add a menu item specifying a URL of your Helpdesk solution, so your users can easily reach it when needed.

Create Windows client package for mass distribution

After configuring a Windows Client Theme, you can create a Windows client package for mass distribution as follows:

1. While still in the Windows client section of the **Theme Properties** dialog, click the **Generate Windows client package** button.
2. In the dialog that opens, specify the following options:
 - Specify the target folder on your local computer where the package will be created.
 - Select or clear the “Open the folder in Windows Explorer ...” option as needed.
3. Click **Generate**. This will create the ClientDownloader.exe file. When you run the file, it will download the latest version of Parallels Client for Windows installer (MSI) and will apply your custom Theme to it.

You can now distribute the installer to end users. When they run the installer, it will install Parallels Client for Windows with all customizations (start menu shortcuts, desktop shortcut, images and icons) specified in the Theme. In the future, if you need to upgrade an installed copy of Parallels Client for Windows to a newer version, you don't need to repeat the instructions described above. Simply upgrade the older version and the branding features will remain intact.

General Theme tasks

When you are done customizing a Theme, click **OK** to save it and return to the Parallels RAS console.

You can also perform the following actions on the **Themes** tab in the Parallels RAS Console:

- **Duplicate a Theme** — right-click a Theme and choose **Duplicate** (or select a Theme and click **Tasks > Duplicate**).
- **Preview User Portal** — right-click a Theme and choose **Preview User Portal** (or **Tasks > Preview...**).
- **Delete a Theme** — right-click a Theme and choose **Delete** (or **Tasks > Delete**).

When done creating or modifying Themes, click **Apply** in the Parallels RAS Console to commit the changes to Parallels RAS. You can now test the Theme by opening its URL in an HTML5-enabled web browser.

Delegating session management permissions

If your organization has multiple user groups, all sharing centralized Parallels RAS resources, you have the ability to delegate session management permissions to an administrator of a particular group. When you do, the administrator can see and manage Parallels RAS sessions only for users who belong to that group.

Here's how this functionality works:

1. A separate Theme is created for each group. Session management permissions for the Theme are delegated to a custom administrator (see [Managing Administrator accounts](#)).
2. When a custom administrator logs in to the Parallels RAS Console, they are presented with a limited user interface displaying sessions that belong to the Theme (or multiple Themes) that the administrator is allowed to manage.

The rest of this section describes how to configure and use this functionality.

Create a Theme and delegate session management permissions

If you don't have a Theme for a user group, you need to create it. Follow the instructions [provided earlier in this chapter](#). To delegate session management permissions, you specifically need to do the following:

1. When specifying settings on the **General** page, select the **Limit access to this Theme to members of these Active Directory groups** option and add one or more groups.
2. After creating or configuring the Theme, close the **Theme Properties** dialog, then right-click anywhere in the list and choose **Delegate Permissions**.
3. If you already have a custom administrator account that you would like to use, it will appear in the list. If you don't have an account, create one as follows:
 1. Click **Tasks > Add**.
 2. In the **Account Properties** dialog, click the [...] button next to **Name** and select an account.
 3. The **Permissions** field is read-only and set to **Custom administrator** (the type that must be used here).
 4. Populate the rest of the fields (email, mobile, etc.) as needed.
 5. Click **OK**.
4. Back in the **Delegate Permission** dialog, select the administrator in the left pane.
5. In the lower portion of the right pane, select permissions (view, modify, manage sessions) for the desired Theme. You can also set permissions in the upper portion of the right pane, but they will apply to all existing Themes, and this is probably not what we are trying to do here.
6. Click **OK**.

Manage sessions

Once the above is complete, the custom administrator can manage sessions that belong to the specified Theme(s). To manage sessions:

1. Run the Parallels RAS Console and log in using the credentials of a custom administrator.
2. The right pane will contain sessions that belong to the members of the group(s) assigned to the Theme.

3. To manage a session, select it, then click the **Tasks** drop-down list and choose a desired option (Disconnect, Log off, Send message, etc.).

Settings audit

Any changes to administrator permissions are recorded in the settings audit. Possible actions are create, update, and delete. You can view the changes by going to **Administration > Settings audit** or **Farm > Themes > Settings Audit**.

Using Themes in Parallels Client for Windows

In order for a user to use a corresponding Theme, the connection properties must be properly set. To do so:

1. In Parallels Client for Windows, right-click a connection and choose **Connection Properties**.
2. On the **Connection** tab, the server name must be followed by the Theme name after a forward slash, as in `Server-name/Theme-name`.

When the administrator views sessions in the RAS Console, a client using a Theme can be identified by the Theme name in the **Theme** column.

Open Parallels Web Client

To open Parallels Web Client in a web browser, enter one of the following in a web browser, depending on your setup:

- The DNS name of an HALB device or HALB Virtual Server (if in use). For example, `https://ras.msp.com`.
- The FQDN or IP address of a specific RAS Secure Gateway. For example, `https://ras-gw1.company.dom`.

For more information about the Web Client URL, please see [Web request load balancing](#).

When you open the Web Client in a web browser, the login page is displayed.

Note: By default, when a user opens Web Client in a web browser for the first time, the cookie consent message is displayed at the top of the page in accordance with the GDPR regulation. To read the Parallels cookie policy, the user clicks the provided link. To agree with the policy, the user clicks **Got it** to close the message

and continue. The RAS administrator can disable the cookie consent message in the [Theme settings dialog](#).

To log in to Parallels RAS, specify your user name in the UPN format (username@domain.com) and password and click **Log in**.

Note: If Parallels RAS is configured to use Google Authenticator as a second-level authentication provider, an additional dialog opens where the user can either scan a QR code or use a secret key to generate a one time password (OTP). For details, please see [Using Google Authenticator](#).

Once the user is logged in, one of the scenarios described below takes place depending on how the Web Client is configured on the server side. For details, please see [Configure Web Client](#).

Launch apps in Parallels Client and fallback to Browser

With this option configured on the server side, you will see a dialog box in the web browser with the following options:

- **Detect Client.** Determines if Parallels Client is installed on the local computer. If it is not installed, opens the Parallels Client installation page. Follow the instructions to install Parallels Client.

Note: If you don't have administrative permissions on this computer, a dialog will open saying so. The dialog has two buttons: **Install Full Client** and **Install Basic Client**. If you know credentials of an administrative account on this computer, click **Install Full Client** and enter the credentials when asked. The installation will continue using these credentials and the full version of Parallels Client will be installed. If you don't know the credentials, click **Install Basic Client**. The basic version of Parallels Client will still work but some of the functionality will be missing.

After the installation, you should see Parallels Web Client displaying published resources that you can use. Please also note a link in the lower left corner of the screen displaying the Parallels Client version and build number.

You can now run remote applications and desktop in Parallels Client or in Parallels Web Client. The default method for running applications and desktops is Parallels

Client. To run a remote application or desktop in Parallels Web Client, right-click it (or tap and hold on a mobile device) and then choose Parallels Web Client.

- **Use Web Browser.** Closes this dialog box and opens the main Parallels Web Client screen. Remote applications or desktops will be launched in the web browser. When you open Parallels Web Client the next time, you will again see the same dialog box with the same options.

Parallels Client only

When this option is configured on the server side, you will see a dialog box prompting you to install Parallels Client. Click the link provided to open the Parallels Client download and installation page and follow the instructions. After you install Parallels Client, the main User Portal screen opens displaying published resources that you can use. If you now double-click or tap a resource, it will be launched in Parallels Client.

Browser only

With this option configured, the main User Portal screen opens with no additional prompts. Remote applications and desktops will be launched in the web browser.

Main menu options

To open the Parallels User Portal main menu, click or tap the “person” icon in the upper-right. You can select from the menu options described below.

Settings

Allows you to configure the following settings:

- **Clipboard redirection:** Enable or disable the clipboard in a remote session. Select from the following options: **Bidirectional** (copying is allowed in both directions), **Server to client only**, **Client to server only**, **Disabled** (copying in either directions is not allowed).
- **Sound:** To play the sound on the local computer, select the **Bring to this computer** option. If sound is not supported by your browser, the menu will be disabled and you’ll see a corresponding text message below it.

- **Remote audio recording:** Enable or disable the sound input redirection from the local computer to the remote application. For example, if you would like to use a microphone in Skype or a similar app for teleconferencing, you need to enable audio recording in User portal. Select **Record from this computer** to enable recording or select **Do not record** to disable it.

Note: Audio input is supported in Chrome, Firefox, Edge and Safari 11. If your browser doesn't support audio input, this setting will be disabled and you will see a text message instead.

- **Redirect links:** Select a desired redirection option from the following: **Do no redirect**, **Redirect URLs**, **Redirect email**, **Redirect all**. When redirection is enabled, a link will be opened on the local computer.
- **Pen and touch input:** Enable or disable pen input redirection with pressure sensitivity support. Please note that the eraser button is not supported.

Note: Pen input redirection is supported with the following software: Chromium-based browsers running on Windows 8.1 or later, Google Chrome running on Chrome OS. This functionality was tested on Chrome OS 97.X and 98.X. endhint %}

- **Redirect printers:** Select a printer redirection option: **RAS Universal Printer** (uses the RAS Universal Printing technology) or **Do not redirect** (printers will not be redirected).
- **Keyboard mode:** Select **Universal Keyboard** or **PC Keyboard**. If you have problems typing certain characters, try selecting **PC Keyboard** and then selecting a proper layout in the **Keyboard Layout** drop-down list (see below).
- **Keyboard layout:** Select a keyboard layout (e.g. English (US), English (UK), Japanese). To enable this drop-down list, the Keyboard Mode option must be set to PC Keyboard.
- (ChromeOS only) **Use a shortcut for the Windows key:** Specify the shortcut that will be used in place of the Windows key.
- **Auto login:** Enable or disable auto login in User Portal. If this option is on, and the user credentials have been saved before, the user will not have to enter them again. This option may not be available if a Client Policy was applied where this option is turned off. Note that the auto login option is supported on the latest Chromium-based browsers, such as Google Chrome and Microsoft Edge. For more information, please see [Auto login](#).
- **Connection timeout, seconds:** Specify the connection timeout.

- **MFA: Remember last method used:** If using multi-factor authentication, enable this option so the last method used is remembered and used by default.
- **Always ask for credentials when starting application:** If this option is enabled, a user will be asked to enter credentials when starting an application even if the session is still active. You can use this option as added security to prevent unauthorized users to access applications. For example, if a user disconnects from a session, no one else will be able to take over the session and run remote applications. As another example, if a user leaves a device with an open User Portal displaying the app listing (with or without running RDP sessions) then any user who tries to open a new application or another instance of a running application will be prompted for credentials. Please note that the **Auto login** option (described above) must be disabled for this functionality to work; otherwise saved credentials will be used automatically.

Change Password

Allows the user to remotely change their domain password. When the password is being changed, the password requirements are displayed on the screen, so the user can follow them for the new password to be accepted. If you use a third-party identity provider, you can configure User Portal to use a custom URL for changing passwords via [RAS Connection Broker connection settings](#). These options can be disabled through Client Policies (**Control settings > Password > Prohibit changing password**).

Detect Client

Determines if Parallels Client is installed on the local computer. If Parallels Client is not installed, gives user an option to install it or skip the automatic Parallels Client detection on subsequent logons.

Download Client

Opens a web page with instruction on how to download and install Parallels Client.

Logout

Ends user session with Parallels RAS and logs the user out.

Running remote applications and desktops

To launch a remote application or desktop in User Portal, do one of the following:

- Double-click (or tap on a mobile device) an application or a desktop icon. The resource will open inside a web browser or in Parallels Client depending on the server-side User Portal configuration (**RAS Secure Gateway Properties > User Portal > Launch sessions using** option).
- Right-click (or tap and hold on a mobile device) an application or a desktop to display a context menu. The menu will appear if the **Allow user to select launch method** or **Allow opening applications in a new tab** (or both) options are selected on the **RAS Secure Gateway Properties > User Portal** tab in the RAS console. The menu allows you to choose whether to open the resource in Parallels Client or Parallels Web Client (depending on the setting mentioned above) and it also allows you to choose whether to open an application in the same or new tab in the web browser.
- If a resource cannot be opened in Parallels Client due to an error, a message will be displayed with an option to open it in the web browser instead.

Please note that to open a resource in Parallels Client from the HTML5 page, a URL with a custom scheme is used. When you double-click on a resource on the HTML5 page, the URL is executed and is then passed to Parallels Client which uses the instructions that it contains to open the resource. For more information see [RAS Web Client API and Parallels Client URL scheme](#).

Using drag and drop functionality

Parallels Web Client supports drag and drop functionality when running remote applications and desktops.

Note: The **Allow file transfer command** option must be enabled on the Gateway for the drag and drop functionality to work. See [Configure Web Client](#).

Here's how to use drag and drop when working with a remote application:

1. Select a file on your local computer.
2. Drag and drop the selected file to an app. The 'Save as' window will pop up.

3. Enter a name for the file and save it. The file will be saved on the server hosting the app.

You can also drag and drop files between two remote apps running on different hosts.

Here's how to use drag and drop with a remote desktop:

1. Select a file on your local computer.
2. Drag and drop the selected file to a remote desktop. The 'save as' window will pop up.
3. Enter a name for the file and save it. The file will be saved on the desktop on the server that hosts it.

Native clipboard experience

In Parallels RAS 18.2 and newer, you can copy and paste plain text between a local device and remote session in both directions. Simply use Ctrl+C/Cmd+C to copy (Ctrl+X/Cmd+X to cut) and Ctrl+V/Cmd+V to paste. This functionality is fully supported in Chromium based browsers (Chrome, Edge Chromium, Opera) and Internet Explorer. On Firefox, only server to client copy/paste is supported. Other web browsers do not support this functionality.

Other useful features

Other useful functionality on the main User Portal screen includes the following:

- **Favorites list.** You can add a remote application or a desktop to the Favorites list, so you can easily find them. To do so, point to or tap an application or a desktop and then click or tap the "star" icon. To view the list, select the **Favorites** tab at the top of the list. To remove a resource from the list, point to it and click the "X" icon (or point to or tap the resource icon and then click or tap the start icon).
- **Search.** To search for a resource, begin typing its name in the **Search** box (upper right). The list will be filtered as you type to contain only the resources with matching names.
- **List view.** You can switch between the grid and the list view by clicking the icon below the search box. The list view allows you to see the descriptions of the published resources.

- **View a description in the grid view.** To view a resource description in the grid view, position the mouse pointer over it. The description will appear as a tooltip. This could be helpful if one or more resources are published using the same name. By reading the description, you can distinguish between them.
- **Taskbar.** When you launch a remote application or a desktop, its icon is added to the taskbar at the bottom of the screen. When the taskbar is full, items of the same type are grouped to save space. You can click or tap on a group to see the list of all running instances and to switch to or close a particular instance.

Auto login

Auto login for RAS User Portal facilitates the frequent use of the portal by providing auto login option for the user without requiring user intervention to input their credentials. With Auto login enabled, as soon as user opens the Parallels User Portal, user will be automatically logged in, able to see the list of resources that were made available by the administrator and to launch resources accordingly. This experience can be configured from the Web Client settings or centrally controlled from [RAS policies](#). It is set to accelerate user login and increase user experience by reducing the number of times users are prompted to login to the Parallels User Portal.

To use this feature, the following requirements must be met:

- The Secure Gateway that hosts the User Portal and Web Client must have a valid and trusted (by end point devices) certificate.
- There is only one set of valid credentials already saved for a given domain (including subdomains). Auto login is not applicable in case of using shared devices due to multiple user accounts accessing the same Parallels RAS environment.
- Not using incognito — when using incognito, and only one set of credentials is available, sign-in pop up still shows, so auto login cannot happen without the user intervention.
- Auto login is supported on Chromium-based browsers only, such as Google Chrome and Microsoft Edge.

Configuration

The following settings control the Auto login functionality:

- When a user logs in to User Portal for the first time, they need to click **Save** when asked to save the password. For this to happen, the **Offer to save password** and **Auto sign-in** options must be enabled in the browser (these are the default settings in a Chromium-based browser).
- The user needs to confirm to reduce the number of times they asked for credentials. This will enable the Auto login option.
- The RAS administrator may also use RAS policies to enforce auto login (enabled/disabled). This can be done from **Policy > Session > Connection > Primary Connection > Auto Login**.
- An expiry time of 60 days is set when the **Auto login** option is enabled due to security.

Using Auto login

The following describes how Auto login works:

1. The user opens the User Portal web page in a browser and logs in. Note that [direct app access is also supported](#).
2. On the first login, User Portal will suggest the user to enable Auto login.
3. When the user opens the User Portal (or uses a direct app link) the prompt to enter credentials will not appear.

To see the Auto login setting in User Portal, click on the user icon in the top right and click **Settings**. Examine the **Auto Login** setting.

Direct App access

Specific published resources in Parallels RAS 18 may be directly accessed through RAS Web Client. This can be achieved with the introduction of a new parameter, *appid*, which allows administrator to use links to access the published resource directly. This allows a more flexible and easier way for users to access Parallels RAS published resources such as using browser shortcuts or bookmarks or third-party portals such as Azure My Apps Portal to access independent SAAS applications and Parallels RAS virtual apps and desktops.

To launch a published resource directly, you need to specify a URL using one of the following formats:

URL format	Description
https://FQDN? appid=	
https://FQDN/ <Theme- name>? appid=<app- ID>	This format is similar to the one above, but specifies a Theme name.
https://FQDN/ userportal? theme=&appid=	This format is the same as the one above, but uses the full URL specification. It is listed here just for reference.

Supported parameters:

Parameter	Description
appid	The published item (application or desktop) to be launched.
overrideparams	[Optional]. URL Encoded Override arguments that needs to be passed to the published application.

Example:

https://FQDN?appid=14&overrideparams=C%3A%2Ftest.txt

When opening a published resource using a direct link, the [Auto login](#) option is also used depending on the settings.

Using the toolbar

User Portal includes a special toolbar that becomes available when you launch a remote application or desktop. The toolbar appears differently for remote desktops and remote applications. The toolbar has also slightly different functions for desktop computers and mobile devices. The differences are explained in the subsequent topics.

Using the toolbar on desktop computers

Remote desktop toolbar

When you launch a remote desktop in a web browser on a desktop or laptop computer, the toolbar appears as follows:



The top area of the toolbar is used to drag the toolbar up or down. Click and hold it and then drag the toolbar to the desired position. The arrow icon is used to show or hide the toolbar items.

The main toolbar items are (from top to bottom):

- **Full screen.** Display the remote desktop in full screen on the local computer.
- **Upload a file.** Upload a file from the local computer to the remote server. After clicking this item, you are presented with two dialogs, one after another. In the first dialog, select a file on the local computer you wish to upload. In the second dialog, select a location on the remote server where you want to save the file.
- **Download a file.** Download a file from the remote server to the local computer. After clicking this item, select a file on the remote server you wish to download. Depending on your web browser configuration, the download will start automatically or you will be asked to select a destination folder on your local computer.
- **Shortcuts.** Display the **Shortcuts** menu (see below for the menu description).
- **Clipboard.** Display the remote clipboard. Please see [Using the remote clipboard](#) for more information.

The **Shortcuts** menu allows you to send keystrokes and key sequences to the remote desktop:

- **Escape.** Sends the “Escape” keystroke to the remote desktop.
- **Tab.** Sends the “Tab” keystroke.
- **Backspace.** Sends the “Backspace” keystroke.
- **Print screen.** Sends the “Print Screen” keystroke. The screen will be printed to the clipboard of the remote desktop from where you can paste it into an application (e.g. Paint) running on the same remote computer.
- **Windows Key.** Sends the “Windows logo key” keystroke.
- **Control+Alt+Delete.** Sends the “Ctrl+Alt+Delete” key sequence.

Remote application toolbar

When you launch a remote application, the toolbar is embedded into the page footer and it's collapsed by default. To expand the toolbar, click the “arrow-up” icon in the lower right-hand corner.



The toolbar items are (from top to bottom):

- **Download.** Download a file from the remote server to the local computer. After clicking this item, select a file on the remote server you wish to download. Depending on your web browser configuration, the download will start automatically or you will be asked to select a destination folder on your local computer.
- **Upload.** Upload a file from the local computer to the remote server. After clicking this item, you are presented with two dialogs, one after another. In the first dialog, select a file on the local computer you wish to upload. In the second dialog, select a location on the remote server where you want to save the file.
- **Clipboard.** Display the remote clipboard. Please see [Using the remote clipboard](#) for more information.

Using the Toolbar on Mobile Devices

Remote Desktop Toolbar

When you launch a remote desktop in a web browser on a mobile device, the toolbar appears as follows:



The small arrow icon at the top is used to show or hide the toolbar items.

The main toolbar items are (from top to bottom):

- **Upload a file.** Upload a file from the local device to the remote server. Note that in iOS, you can upload from the Photos folder only.
- **Download a file.** Download a file from the remote server to the local device (not available in iOS).
- **Shortcuts.** Display the **Shortcuts** menu (see below for the menu description).
- **Clipboard.** Display the remote clipboard. Please see [Using the remote clipboard](#) for more information.
- **Keyboard.** Display the native keyboard. This opens your mobile device native keyboard so you can type in an application on the remote desktop.
- **Arrow.** The arrow icon is used to switch between the two available mouse input modes:

Mode 1: The first mode (the arrow icon is white) follows the movement of your finger on the screen and performs a click on a remote desktop where you tap.

Mode 2: The second mode (the arrow icon is red) displays a virtual mouse pointer on the remote desktop and allows you to move that pointer to a precise position with your finger. When you tap anywhere on the screen, the click on the remote desktop is performed at the precise position of the virtual mouse pointer.

The **Shortcuts** menu allows you to send keystrokes and key sequences to the remote desktop:

- **Escape.** Sends the “Escape” keystroke to the remote desktop.
- **Tab.** Sends the “Tab” keystroke.
- **Backspace.** Sends the “Backspace” keystroke.
- **Print screen.** Sends the “Print Screen” keystroke. The screen will be printed to the clipboard of the remote desktop from where you can paste it into an application (e.g. Paint) running on the same remote computer.
- **Windows Key.** Sends the “Windows logo key” keystroke.
- **Control+Alt+Delete.** Sends the “Ctrl+Alt+Delete” key sequence.

Remote application toolbar

When you launch a remote application, the toolbar is embedded into the page footer and it's collapsed by default. To expand the toolbar, click the “arrow-up” icon in the lower right-hand corner.



The toolbar items are (from top to bottom):

- **Download.** Download a file from the remote server to the local device (not available in iOS).
- **Upload.** Upload a file from the local device to the remote server. Note that in iOS, you can upload from the Photos folder only.
- **Clipboard.** Display the remote clipboard. Please see [Using the remote clipboard](#) for more information.
- **Keyboard.** Display the native keyboard. This opens your mobile device native keyboard so you can type in an application on the remote desktop.


Using the remote clipboard

The Remote Clipboard allows you to copy and paste text between your local client system and a remote application or desktop. The clipboard is accessed from the toolbar.

Note: In Parallels RAS 18.2 and newer, you can copy and paste plain text between a local device and remote session without using the Remote Clipboard. For more information, see [Native clipboard experience](#).

To use the clipboard:

1. Expand the toolbar click the **[A]** icon.
2. This opens the **Remote Clipboard** window. On the screenshot below, a remote desktop toolbar is shown. A remote application toolbar looks differently, but it functions exactly the same.



3. To copy text from the local computer to a remote application, type (or paste) it in the **Remote Clipboard** window. The text is automatically saved on the remote computer clipboard, so you can use a standard paste command (e.g. Ctrl+V) to paste it into a remote application.
4. To copy text from a remote application to the **Remote Clipboard** window, highlight it and use the standard copy command (e.g, Ctrl+C). The text will appear in the **Remote Clipboard** window from where you can copy it to a local application.

Hiding toolbar items

You can hide the clipboard and file transfer items on the toolbar if you believe that it's a security risk. The clipboard can be disabled on a RAS Secure Gateway or Client Policy level.

To disable the clipboard for a Secure Gateway:

1. In the Parallels RAS Console, navigate to **Farms** > <Site> > **Secure Gateways**.
2. Right-click a desired RAS Secure Gateway and choose **Properties**.

3. Select the **User Portal** tab and clear the **Allow clipboard command** option in the **Restrictions** section.

You can also disable the clipboard on the Client Policy level, which will disable it for a given user or user group on any Gateway they connect to:

1. In the Parallels RAS Console, select the **Policies** category.
2. Right-click a policy and choose **Properties**.
3. Select the **Connection Properties** item in the left pane and then select the **Local Resources** tab in the right pane.
4. In the **Local devices and resources** section, clear the **Clipboard** option.

Note: Please note that when enabling or disabling the clipboard on a client policy level, this will also affect the clipboard functionality on desktop and mobile versions of Parallels Client. This means that if you disable the clipboard, the desktop and mobile device users will not be able to use their local clipboard when working with a remote application.

You can also disable the file upload and file download items on the toolbar. For instructions, please read the [Configuring remote file transfer](#) section.

Universal Printing

Printer redirection enables users to redirect a print job from a remote application or desktop to their local printer, which can be connected to the user's computer or be a local network printer attached via an IP address. RAS Universal Printing simplifies the printing process and solves most printer driver issues by eliminating the need for a remote server to have a printer driver for a specific local printer on the client side. Therefore, a user can print regardless of which printer they have installed locally, and the RAS administrator doesn't have to install a printer driver for each printer connected to the local network.

Managing Universal Printing Settings

To configure RAS Universal Printing, select the **Universal Printing** category in the RAS Console.

By default, the Universal Printing driver is automatically installed together with an RD Session Host Agent, VDI Guest VM Agent, or a Remote PC Agent. Therefore,

upon adding a server to the Farm, the Universal Printing is already enabled. The Universal Printing driver is available as a 32 bit and 64 bit version.

Enabling and disabling Universal Printing support

To enable or disable the Universal Printing support for a server, right-click the server in the **Servers in Site** list and click **Enable** or **Disable** in the context menu.

Configuring a printer renaming pattern

By default, Parallels RAS renames printers using the following pattern: %PRINTERNAME% for %USERNAME% by Parallels. For example, let's say a user named Alice has a local printer named Printer1. When Alice launches a remote application or desktop, her printer is named Printer1 for Alice by Parallels.

To change the default printer renaming pattern, select the Universal printing category. On the Universal printing tab, specify a pattern in the **Printer rename pattern** field. To see the predefined variables that you can use, click the [...] button next to the input field. The variables are:

- %CLIENTNAME% — the name of the client computer.
- %PRINTERNAME% — the name of a printer on the client side.
- %SESSIONID% — RAS session ID.
- %USERNAME% — the name of the user connected to RAS.
- <2X Universal Printer> — This is a legacy mode where only one printer object will be created in the RDP session.

You can also use certain other characters in a printer renaming pattern. For example, you can define the following commonly used pattern: Client/%CLIENTNAME%#/%PRINTERNAME%. Using this pattern (and the user named Alice from the example above), a local printer will be named Client/Alice's Computer#/Printer1

You can specify a different printer renaming pattern for each server in the **Servers in Site** list.

Note: Redirected printers are only accessible by the administrator and the user who redirected the printer.

Printer retention

When client-defined printers are redirected to a remote session, it takes time and impacts overall session establishing time. To improve user experience, you can reuse previously created user's printers. To do so, on the **Universal printing** tab, set the **Printer retention** option to **On**.

Universal Printing drivers

A system administrator can control the list of client-side printer drivers which should be allowed or denied the Universal Printing redirection privileges.

Using this functionality you can:

- Avoid server resource overloading by non-useful printer redirection. Since the majority of users choose to redirect all local printers (this is default setting), a large number of redirected devices is created on the server which are not really used. It's mostly related to various paperless printers like PDFCreator, Microsoft XPS Writer, or various FAX devices.
- Avoid server instability with certain printers. There are some printers that might create server instability (spooler service component) and as the result deny printing services as a whole for all connected users. It is very important that the administrator has the ability to include such drivers to the "deny" list to continue running printing services.

To specify printer drivers:

1. In the Parallels RAS Console, navigate to **Universal Printing > Printer drivers**.
2. In the **Mode** drop-down list, select which printers should be allowed redirection from the following options:
 - **Allow redirection of printers using any driver** — (default) This option places no limitation on the type of driver a printer is using to use redirection privileges.
 - **Allow redirection of printers using one of the following drivers** — Only the printers using the drivers listed in the box below the **Mode** field are allowed redirection. To add a printer driver to the list, click the **Tasks > Add** (or click the **+** icon) and type the printer driver name in the edit field provided.
 - **Don't allow redirection of printers that use one of the following drivers** — This is probably the most useful option in the context of this

feature. The printers that use the drivers specified in the list will be denied redirection privileges. All other printers will be allowed to use redirection. To add a printer driver to the list, click the **Tasks > Add** (or click the **+** icon) and type the printer driver name in the edit field provided.

3. To delete a printer driver from the list, click **Tasks > Delete** or click the minus-sign icon.
4. When done making changes, click the **Apply** button to save the changes.

Please make a note of the following:

- When adding a printer driver to the list, type the printer *driver* name, not the printer name.
- The driver names comparison is case insensitive and requires full match (no partial names, no wildcards).
- The settings that you specify on this tab affect the entire Site (not an individual server).

Font management

Fonts need to be embedded so that when printing a document using Universal Printing the document is copied to the local spooler of the client machine to be printed. If the fonts are not present on the client machine the print out would not be correct.

To control the embedding of fonts within a print job use the **Fonts Management** tab page and check/uncheck the option **Embed Fonts**.

Excluding fonts from embedding

To exclude a specific font type from being embedded, click **Tasks > Add** in the **Exclude the following Fonts from embedding** section and select a font from the list.

Automatically install fonts on servers and clients

To automatically install a specific font type on servers and clients, click **Tasks > Add** in the **Auto install fonts** section and select the fonts from the list.

Note: By default, fonts added to the auto-install list will be excluded from the embedding list because the fonts would be installed on the Windows clients,

therefore, there is no need for them to be embedded. Clear the option **Automatically exclude font from embedding** in the select font dialog so the font is not excluded from the embedding list.

Resetting excluded fonts to default

To reset the list of excluded fonts to default, click **Tasks > Reset to default**.

You can also specify a universal printing compression policy. For more info see [Client Policies > Experience](#).

Universal Scanning

Scanner redirection enables users who are connected to a remote desktop or accessing a published application to make a scan using the scanner that is connected to the client machine. This chapter describes how to configure and use RAS Universal Scanning services.

Managing Universal Scanning

Universal Scanning uses TWAIN and WIA redirection to let any application using either technology hardware connected to the client device for scanning. With Universal Scanning there is no need to install a specific scanner driver on the server.

Note: The server feature **Desktop Experience** is required in order to enable both WIA and TWAIN scanning on RD Session Hosts.

To configure Universal Scanning, select the **Universal Scanning** category in the RAS Console.

By default, the Universal Scanning driver is automatically installed with RD Session Host, Guest VM, and Remote PC agents. Therefore, upon adding a server to the Farm the Universal Scanning is installed.

Configuring a Scanning Rename Pattern

By default, Parallels RAS renames scanners using the following pattern: %SCANNERNAME% for %USERNAME% by RAS. For example, if a user named

Lois, who has SCANNER1 installed locally, connects to a remote desktop or published application, her scanner is renamed to “SCANNER1 for Lois by RAS”.

To change the pattern used to rename scanners, specify a new pattern in the **Scanner rename pattern** input field. The variables that you can use for renaming are:

- %SCANNERNAME% — client-side scanner name.
- %USERNAME% — username of the user connected to the server.
- %SESSIONID% — ID of the active session.

You can configure a different renaming pattern specifically for each server in the list.

Note: Redirected scanners are only accessible by the administrator and the user who redirected the scanner.

Enabling and Disabling Universal Scanning Support

To enable or disable the WIA or Twain Universal Scanning support for a particular server, click the **WIA** tab or the **TWAIN** tab, then right-click a server and click **Enable** or **Disable** in the context menu.

Adding scanning applications

Adding a scanning application

TWAIN applications that will use the Universal Scanning feature have to be added in the TWAIN tab by selecting the **TWAIN Applications** button so they can use the Twain driver, hence making it easier for the administrator to set them up.

To add an application to the list of scanning applications:

1. With the **Universal Scanning** category selected in the RAS Console, click the **TWAIN** tab.
2. Click the **Twain Applications** button (below the **Servers in Site** list) and then click **Add**.
3. In the **TWAIN Applications** dialog, click **Tasks > Add** and browse for the application executable. Select the executable and click **Open**.

Note: Some applications might use different or multiple executables. Make sure that all required executables are added to the list of scanning applications. endhint %}

Deleting a scanning application

To delete a scanning application from the list, highlight it and click **Tasks > Delete**.

Note: If you delete an application from the list, the installation of the application will not be affected.

You can also specify a universal scanning compression policy. For more info see [Client Policies > Experience](#).

User Device Management and Client Policies

This chapter describes tasks that a Parallels RAS administrator can perform to manage user devices, such as desktop computers, phones, or tablets.

Inviting users to connect to Parallels RAS

Parallels RAS supports multiple platforms ranging from desktop PCs and Mac computers to mobile devices and ChromeApps. The Invitation Email feature is designed to reduce the complexities involved in the installation and client rollout process. This feature allows the administrator to send client installation and automatic configuration instructions to end users right from the Parallels RAS Console.

Before proceeding, please confirm that you've configured the mailbox as described in [Configuring SMTP server connection for email notifications](#). To send an invitation email to users, use the **Start** category in the RAS Console. For more information see [Invite users](#).

Mass configuring user devices

If you need to configure Parallels Client that is already installed on multiple devices in your organization, you can simplify the procedure by using one of the following mass configuration options:

- By exporting Parallels Client settings to a file and then importing them into all other Parallels Client installations.
- Using the Parallels Client URL scheme.

Exporting and importing Parallels Client settings

Parallels Client includes the Export/Import functionality that lets you export RAS or RDP connection settings to a file and then import them into Parallels Client running on another device. This functionality is available on all platforms, including desktop and mobile versions of Parallels Client (except Parallels Client for Chrome App). The Export/Import functionality is accessed in Parallels Client as follows:

- **Windows, Mac, Linux:** On the main menu, click **File > Export Settings** or **File > Import Settings**.
- **iOS/iPadOS:** To export connection settings, tap the [...] icon in the top right corner and choose **Share Connection**. To import, select the file that you exported earlier and choose to open it with Parallels Client.
- **Android:** To export connection settings, tap the menu icon (three vertical dots) in the top right corner and choose **Share connections**. To import, select the file that you exported earlier and choose to open it with Parallels Client.

For more information about exporting and importing connection settings, see the Parallels Client Guide for a desired platform.

Using Parallels Client URL scheme

Parallels RAS uses a URL scheme to perform actions in Parallels Client installed on user devices. Specifically, the URL scheme can be used to configure RAS and RDP connections using predefined settings. For the information about the URL scheme please see [RAS Web Client API and Parallels Client URL scheme](#).

The URL scheme is used in invitation emails when you send an email to your users to install Parallels Client on their devices. An invitation email includes a link, which is a complete URL that uses the Parallels Client URL scheme. When you

mass install Parallels Client on user devices, you simply [send an invitation email to your users](#). If you need to reconfigure existing Parallels Client installations (and don't want to do it by sending an invitation email), you can do the following:

1. Create an invitation email containing configuration profiles for all required platforms and send it to yourself.
2. Open the email and copy Parallels Client configuration URLs to a local intranet portal.
3. Let your users know where the URLs are.
4. To configure Parallels Client, your users will need to simply click a URL for their platform. This will automatically configure Parallels Client on their devices.

Enabling Help Desk support for users

Parallels Client provides users with the ability to send a support request, together with a problem report, to your organization's help desk.

Note: At the time of this writing, this functionality is only available in Parallels Client for iOS and Parallels Client for Android. Support for other clients will be added in future releases.

To enable Help Desk support for users, do the following:

1. In the RAS Console, select the **Features** category.
2. Select the **Enable Helpdesk functionality in Parallels Client** option and specify your help desk email address in the field provided. This email address will be updated in Parallels Client every time a user connects to Parallels RAS from it.

Help desk can be accessed in Parallels Client from the Help section (or menu). When the user selects the **Request support from helpdesk** item, a local email client will open. The following information will be prefilled in the email:

- Help desk email address (the one you set in the RAS Console).
- Application name.
- A screenshot.
- User name.
- Application version.
- Operating system version.

The user can provide their own description of the request.

Enabling Help Desk support for custom administrators

RAS console and Management Portal provide custom administrators with the ability to send a support request to your organization's help desk.

To enable Help Desk support for custom administrators, do the following:

1. In the RAS Console, select the **Features** category.
2. Select the **Overwrite the local support actions with the following URL** option and specify the link to to your local support portal in the field provided. This link will open when a custom administrator clicks on **Help > Request Support** in RAS Console or **Help and Support > Request Support** in Management Portal.

Monitoring devices

Device monitoring allows you to view devices which are connected to the Farm or have established a connection at least once in the past. To monitor devices, select the **Device Manager** category in the Parallels RAS Console and click the **Devices manager** tab in the right pane. The information for a device includes:

- Device name
- IP address
- State (see below for the list of states)
- Last user (who used a device)
- MAC address
- OS version
- Parallels Client version
- Group (if a device is a member of a device group)
- Gateway name (the RAS Secure Gateway a device is connected to)
- Gateway IP address

Device states

Devices that connect to Parallels RAS can have any of the following states:

- **Off:** Device is switched off.
- **Connected:** Device is connected.
- **Logged On:** Devices is logged on to the system.
- **Standalone:** Device has previously connected to the Parallels RAS but is not using Parallels Client, therefore it cannot be managed.
- **Not Support:** Device is not supported by the Parallels RAS.
- **Foreign Managed:** Connecting to the Farm but managed by a different Farm.
- **Not Manageable:** Client not manageable due to incompatible client version or uninstalled component.
- **Locked.** Device has an active session in locked status.
- **Pair Pending.** Connection should be refreshed on the client side; port UDP 20009 is blocked from the client to gateway; client management port is disabled on the gateway.

Switching off device monitoring

If you use third-party endpoint management solutions and do not need Parallels RAS device monitoring, you can switch it off on the **Device Manager > Options** tab. It helps you to save computing resources and may improve performance of the RAS Console.

To switch off device monitoring:

1. In the RAS Console, navigate to **Device Manager > Options**.
2. Clear the **Enable device manager** option.
3. Click **Yes** and **Apply**.

As soon as you switch off device monitoring, the RAS Console stops tracking connected devices and deletes device connection history, which is displayed on the **Device Manager > Device Manager** tab. After that you can still view information about current connections in the **Sessions** category.

Getting additional device information

To see the additional device information, right-click a device and choose **Get Device Information** in the context menu. In the dialog that opens, review the following properties:

- **Name:** Device name.
- **IPs:** Device IP address (or multiple addresses if applicable).
- **MAC Address:** MAC address.
- **State:** State (see below for the list of states).
- **Last User:** The user who logged in from this device the last time.
- **Last Logon Time:** The time of last logon.
- **OS Version:** The operating system version running on the device. Windows portable and U3 clients are marked as “Portable”.
- **Client Version:** Parallels Client version installed on the device.
- **Gateway IP:** The RAS Secure Gateway IP address (the gateway the client is using).
- **Secure Gateway:** The RAS Secure Gateway name.
- **Last Activity:** The date and time when any activity was detected from this device.

Windows device groups

The **Windows device groups** tab (**Device Manager** category) allows you to group managed Windows devices and administer them together.

Creating a Windows device group

To create a Windows Device Group:

1. Navigate to the **Windows device groups** tab in the **Device Manager** category and click **Tasks > Add**.
2. On the **Main** tab page, specify a group name and an optional description.
3. On the **OS Settings** tab, set the following options:
 - **Disable removable drives.** Disable mounting of removable drives on managed Windows device.
 - **Disable Print Screen.** Disable the **Print Screen** key.
 - **Replace desktop.** This feature makes a Windows computer behave like a thin client. It limits users from changing system settings or

installing new applications. The administrator can add local apps (which are already installed on a computer) to the app list in addition to published resources from Parallels RAS. If you select this option, specify an administrator password in the **Admin Mode Password** field (below) to be used to switch a computer between user and admin modes.

- **Kiosk mode.** Enable the kiosk mode. This will disable power cycling functions (reboot, shutdown) on computers in the group. Note that power functions will still be available when the computer is switched to the Admin mode.
 - **Use client as desktop.** If this option is selected, Parallels Client will run in full screen mode. A user will not be able to minimize it. Select this option to overcome an issue with Parallels Client breaking out of the kiosk mode on Windows 8.x. The issue may manifest itself in the tile-based UI or while using the “drag to close” feature.
 - **Admin Mode Password.** Specify a password to switch between user and admin modes when a Windows desktop is replaced (see **Replace desktop** above).
4. On the **Firewall Settings** tab, enable or disable the firewall and add the inbound ports if necessary.
 5. On the **Shadowing** tab, select the **Request Authorization** option to prompt a Windows device user before remotely controlling their desktop. If enabled, the user can choose to decline the connection. For more information, see [Managing Windows devices](#).

Adding a Windows device to a group

To add a Windows device to a group:

1. Navigate to the **Device Manager > Device Manager** tab.
2. Select one or more devices, then click Tasks (or right-click) and choose **Move to Group**.
3. Select a group and click **OK** to save the settings.

The administrator can now perform standard Windows power operations (Power On, Power Off, Reboot, Logoff, Lock) on groups of devices.

Managing Windows devices

The Device Manager feature allows the administrator to convert Windows devices running Windows 7 up to Windows 11 into a thin-client-like OS. In order to be managed, Windows devices must be running the latest version of Parallels Client for Windows.

Read the instructions below to learn how to set up Parallels Client on a Windows computer and how to enroll and manage it in Parallels RAS.

Installing Parallels Client on a Windows computer

To install and configure Parallels Client for Windows, follow the steps below. You can also read the **Parallels Client for Windows User's Guide** for complete instructions on how to install and configure Parallels Client.

1. Download the Parallels Client for Windows from <https://www.parallels.com/products/ras/download/client/>.
2. Double-click the `RASClient.msi` or `RASClient-x64.msi` and follow the on-screen instructions to complete the installation wizard.
3. Create a new Parallels RAS connection by clicking **File > Add New Connection**. <https://www.parallels.com/products/ras/download/client/>.
4. Select **Parallels Remote Application Server** and click **OK**.
5. Next, configure the following connection properties:
 - **Primary Connection** — Specify the Parallels RAS FQDN or IP address.
 - **User Credentials** — Enter username, password, and domain.
6. Click **OK** to create the connection and then double-click it to connect to Parallels RAS.

Upon completion, the Windows device will appear in the Parallels RAS Console in **Device Manager > Devices**.

Windows device enrollment

You can configure Parallels RAS to enroll a Windows device automatically or you can opt to do it manually.

To manually enroll a Windows device in Parallels RAS:

1. In the RAS Console, navigate to **Device Manager > Devices**.

2. Select a device on the **Devices** tab.
3. Click **Tasks > Manage Device**.

The device state will change to **Pair pending** until the device reconnects. Ensure the **Device Manager Port** option is enabled for a gateway. To verify this:

1. Navigate to **Farm > <Site> > Secure Gateways**.
2. Select a gateway and click **Tasks > Properties**.
3. Click the **Network** tab and make sure that the **Device Manager Port** option is selected

Once the device reconnects, the enrollment process is complete and the device state is updated to **Logged On**, which indicates that it's now managed by Parallels RAS. The user running Parallels Client on their Windows PC can also verify that the PC is managed by clicking **Help > About** on the main Parallels Client menu. The information includes the RAS Secure Gateway information that the Parallels Client uses to communicate with Parallels RAS.

You can also set Parallels RAS to automatically manage Windows devices. To do so:

1. In the RAS Console, select the **Device Manager** category.
2. Click the **Options** tab.
3. Enable the **Automatically manage Windows devices** option.

The administrator can now check the state of the device and perform power operations, such as Power On, Power Off, Reboot, and Logoff.

Note: Devices running some older versions of Parallels Client cannot be managed and are marked as **Not Supported**.

Lock a Windows device

To lock a Windows device that has an active session, select it in the list and then click the **Lock** item in the toolbar at the bottom. Note that the **Lock** icon is only enabled when the selected device is in the **Logged On** state.

You can also lock a device (or a device group) using the scheduler, which is described in the [Scheduling Windows devices & group power cycles](#) section.

Shadow a Windows device

By shadowing a Windows device, you gain full access to the Windows desktop on the device and can control local and remote applications.

To shadow a Windows device:

1. In the RAS Console, navigate to **Device Manager > Devices**.
2. Select a device and click the **Shadow** item in the toolbar at the bottom.

The Windows user will be prompted to allow the administrator to take control over the device and can choose to deny access. The **Request Authorization** prompt can be deactivated by the administrator. To do so:

1. In the Parallels RAS Console, select the **Device Manager** category and click the **Windows Device Groups** tab in the right pane.
2. Right-click a group and choose **Properties**.
3. In the **Windows Device Group** dialog, select the **Shadowing** tab and clear the **Request Authorization** option.

Desktop replacement

The **Replace desktop** feature limits users from changing system settings or installing new applications. When this feature is enabled, the Windows desktop is replaced by Parallels Client, which converts it into a thin-client-like OS without actually replacing the operating system. This way the user can only deploy applications from Parallels Client, which gives the administrator a higher level of control over connected devices.

Additionally, the Kiosk mode allows you to limit the user from power cycling a device (power actions are still available in the Admin mode; see below for details.).

To enable the **Replace desktop** feature:

1. In the **Device Manager** category, select the **Windows Device Groups** tab.
2. Right-click a group and choose **Properties**.
3. Click the **OS Settings** tab.
4. Enable the **Replace desktop** option and optionally the **Kiosk mode** option.
5. Click **OK**.

Note: This feature requires an administrative password set to switch between User and Admin mode on the Windows device.

Switching to the Admin mode

In User mode, the user is restricted to use only the applications provided by the administrator. In order to change system settings, switch the device to the Admin mode.

To switch to the Admin mode, right-click on the system tray icon and select **Switch to admin mode**. Type the password when prompted.

The following table outlines features that are available in Admin and User modes.

Parallels Client Global Options		X
Parallels Client Farm Connection Properties		X
Configuration of Local Applications		X
Add a new RAS Connection		X
Add a new RDP Connection		X
Manage Standard RDP Connections and Folders		X
Display Settings	X	X
Mouse Settings	X	X
Printer Settings		X
Task Manager		X
Control Panel		X
Command Prompt		X
Windows Explorer		X
Import / Export Settings		X

Configuring local applications when using Parallels Client desktop replacement

With the **Replace Desktop** option enabled, the administrator's goal should be to deploy remote applications or remote desktops and use the native OS to simply deploy the software needed to connect remotely. However, in some instances, local applications may be required. The administrator still has the ability to configure local applications to be shown within the Parallels Client Desktop Replacement, however it is necessary to switch to the Admin mode prior to it.

Publish a local application according to the following steps:

1. Shadow the user's session or use the user device station directly.
2. Switch the Parallels Client Desktop Replacement to admin mode.

3. Click **File > Add New Application**
4. Fill in the application information
5. Applications added will be visible in the Application Launcher.
6. Switch back to user mode once all the applications needed are configured.

Windows desktop replacement

This section explains what happens when the **Replace Desktop** option is enabled and why it is useful to administrators.

When enabled, the Replace Desktop feature allows the administrator to convert a standard desktop into a limited device similar to a Thin Client without replacing the operating system.

The end user will not have access to Windows Explorer, Taskbar, or any other Windows components that usually allow them to install new applications or change system settings. The user can now only deploy applications configured within the Parallels Client, including remote applications, remote desktops, and locally configured applications. Local applications are allowed, so if a specific application is needed but is not available remotely (e.g. software that communicates with specific peripherals), the user can still deploy it.

When the **Replace Desktop** option is enabled, the following features take effect on the corresponding versions of Windows (7, 8, 8.1, 10, 11):

Feature	7	8	8.1	10	11
Replace Desktop with Parallels Client	X	X	X	X	X
Disable Start Button	X	X	X	X	n/a
Restrict Control Panel Access	X	X	X	X	X
Disable Windows Key	X	X	X	X	X
Disable the Task Manager	X	X	X	X	X
Disable Quick Access Toolbar	n/a	n/a	n/a	n/a	n/a
Disable Security Manager/Action Center Notifications	X	X	X	X	X
Lock the Taskbar	X	X	X	X	X
Remove Pinned Applications	X	X	X	X	X
Disable Metro Screen (user logs directly to desktop)	n/a	X	X	X	X
Disable Hot Corners	n/a	X	X	X	X
Disable Charm Hints	n/a	X	X	X	X
Disable Help Aids	n/a	X	X	X	X


Feature	7	8	8.1	10	11
Disable Windows Sidebar	X	n/a	n/a	n/a	n/a

In this mode, the user also has access to the Mouse and Display Control Panel applets. The user cannot change the Parallels Client Global Options and the Client Farm Connection Options. Advanced management features can be enabled if the device is switched into administration mode.


If the Windows Desktop Replacement feature is switched off, all the restrictions are removed and the standard desktop is made available to the user.

The following are the screenshots of a Windows 10 desktop before and after the **Replace Desktop** option is enabled.

Before



After



Scheduling Windows devices & groups power cycles

The **Scheduler** tab of the **Device Manager** category can be used to schedule automatic power operations on devices.

Adding a new scheduler task

To schedule a task:

1. On the **Scheduler** tab, click **Tasks > Add** to open the **Device Scheduler Properties** dialog.
2. Select the **Enable this scheduled entry** option.
3. Select an action in the **Action** drop-down list:
 - **Device group switch on**
 - **Device group log off**
 - **Device group switch off**
 - **Device group reboot**
 - **Device group lock**
4. Select a device group in the **Target** drop-down list.
5. Specify the task start date and time.

6. Select the **Repeat** option from the following choices:
 - **Never** (a task will run only once, as specified in the **Start** and **Time** fields)
 - **Every day**
 - **Every week**
 - **Every 2 weeks**
 - **Every month**
 - **Every year**
 - **On specific day(s) of the week.** When selecting this option, select the day(s) of the week.
7. Enter a task description in the **Description** field.
8. Click **OK** to create the task.

Managing scheduled tasks

To modify an existing task, right-click it in the **Schedule List** and click **Properties** in the context menu.

To enable or disable an event, right-click it, click **Properties**, and then select or clear the **Enable this scheduled entry** option.

To execute a scheduled task immediately, right-click it and click **Execute Now** in the context menu.

To delete a task, right-click it and then click **Delete**.

Client policies

The **Policies** category allows you to manage Parallels Client policies for users connecting to a Farm. By adding client policies, you can group users and push different Parallels Client settings to user devices forcing them to function as your organization requires.

Settings that can be enforced on user devices include RAS connection properties, display, printing, scanning, audio, keyboard, device, and others. Once you create a policy and push it to a client device, the user of the device cannot modify the settings that the policy enforces. In Parallels Client, this will manifest itself as hidden or disabled connection properties and global preferences.

Supported Parallels Client versions

Parallels Clients for all platforms are supported.

Note: Starting with Parallels RAS v16.5, a new approach is used to manage client policies. In the previous versions, a client policy would apply the full set of parameters and replace the client settings, completely hiding an enforced category. In RAS v16.5 and newer, client policy settings are split into smaller groups with the ability to configure and enforce each group on the client side individually. For information on how this affects existing client policies that were created in an earlier version of Parallels RAS, please read [Client policy backward compatibility](#).

Add a new client policy

To add a new client policy:

1. Select the **Policies** category and then click **Tasks > Add** in the right pane. The **Policy Properties** dialog opens.
2. The left pane contains a navigation tree allowing you to select a group of options to configure. You can search for options using the **Find** field in the upper left corner of the dialog. If multiple options are found, you can navigate between them using arrows.
3. Make sure the **Policy** node is selected, and then specify a policy name and an optional description.
4. In the **Apply policy to** section, click **Tasks > Add** (or click the plus sign icon) and specify rules that define what object the policy applies to (see below).

Configure rules for the client policy

By default, a client policy applies to configured users, computers, and groups in all situations. Optionally, you can specify rules that define when the policy should be applied. This functionality allows you to create different policies for the same user or computer, which will be applied depending on where the user is connecting from and from which device. Each rule consists of one or several criteria for matching against user connections. In turn, each criteria consists of one or several specific objects that can be matched.

You can match the following objects:

- User, a group the user belongs to, or the computer the user connects from.


- Secure Gateway the user connects to.
- Client device operating system.
- IP address.
- Hardware ID. The format of a hardware ID depends on the operating system of the client.

Notice the following about the rules:

- Criteria are connected by the AND operator. For example, if a rule has a criteria that matches certain IP addresses and a criteria that matches client device operating systems, the rule will be applied when a user connection matches one of the IP addresses AND one of the client operating systems.
- Objects are connected by the OR operator. For example, if you only create a criteria for matching client device operating systems, the rule will be applied if one of the operating systems matches the client connection.
- The rules are compared to a user connection starting from the top. Because of this, the priority of a rule depends on its place in the rule list. Parallels RAS will apply the first rule that matches the user connection.
- The default rule is used when no other rule is matched. You can set it to either **Apply policy if no other rule matches** or **Do not apply policy if no other rule matches**, but no criteria is available for this rule.

To create a new rule:

1. Select the **Policy** node.
2. In the **Apply policy to** section, click **Tasks > Add**. The **New rule properties** dialog opens.
3. Specify the name and the description of the rule.
4. In the **Criteria** section, specify criteria for the rule. You will find the following controls:



- **Apply policy if** and **Do not apply policy if**: specifies whether the policy is applied or not applied when a user connection matches all the criteria. Click on the link to switch between the two options.
- **(+)**: adds a new criteria. If you want to match a Secure Gateway, a client device operating system, an IP address, or a hardware ID, click **(+)**. In the context menu that appears, select the type of object that you want to match and add the specific objects in the dialog that appears. The new criteria appears on the next line.
- **(X)**: Deletes a specific object from matching. For example, you want to delete IP address 198.51.100.1 from matching, click **(X)** next to it. This control appears when at least one object is added. If all objects in a criteria are deleted, the criteria is removed.
- **is** and **is not**: specifies whether the policy is applied or not applied when a user connection matches the criteria. Click on the link to switch between the two options. This control appears when at least one object is added.
- **configure**: edits the list of objects to be matched. Click this link to add or delete new objects. Note that for the first criteria (**User, group or computer**) this link is called **everyone**. It will change to **configure** once you specify objects for this criteria.

Configure session settings

Items under the **Session** node in the **Policy Properties** dialog include connection, display, printing, network, and other settings that will be enforced on a client if defined and enabled.

For a particular group of settings to be enforced on a client device, it must be selected (checked). Unselected groups will not be enforced, so end users will be able to configure them themselves. For example, you can check the **Connection**

node, but only check the **Primary connection** and **Secondary connections** groups under it. This will enforce only the two selected groups of settings on client devices.

Appearance

To configure the appearance of Parallels Client, select the **Appearance** node and then configure the groups of settings described below.

- **Parallels Client interface.** Select the style of interface for Parallels Client for Windows.
- **Prompt user to switch to Modern interface.** Select this option if you want the user to see a prompt that allows them to switch to Modern interface.

Connection

To configure connection properties, select the **Connection** node and then go through each child node, configuring their respective properties.

Primary Connection

The primary connection always defaults to the primary RAS Secure Gateway, but you can modify the following connection properties:

1. Specify a friendly name for the connection.
2. **Auto login:** Enable or disable auto login in the RAS User Portal. If the option is disabled, auto login will be disabled in the User Portal, and the user will not be able to change it. For more information, see [Auto Login](#).
3. In the **Authentication type** drop-down list, select the desired method of authentication:
 - **Credentials.** The user will have to enter credentials to log on.
 - **Single Sign-On.** This option will be included in the list only if the Single Sign-On module is installed during Parallels Client installation. The credentials that the user used to log on will be used to connect to the remote server.
 - **Smart Card.** Select this option to authenticate using a smart card. When connecting to the remote server, a user will need to insert a smart card into the card reader and then enter a PIN when prompted.

- **Web.** If selected, the SAML SSO authentication is allowed. For more information, see [SAML SSO Authentication](#).
- **Web + Credentials.** The same as **Web**, but users are prompted to enter credentials when they launch a published application. To enable the **Web + Credentials** method, you must configure your IdP and RAS as described in [IdP side configuration](#) and [SP side configuration](#).

Note: The **Single Sign-On** and **Smart Card** options are not supported in the User Portal.

Note: The **Smart Card** option is not supported in Parallels Client for Linux.

Note: The **Web + Credentials** method works only in Parallels Client for Windows and User Portal.

Note: The allowed authentication type(s) must be specified in RAS Console in [Connection > Authentication](#), otherwise, the **Authentication type** policy will not work.

1. Select or clear **Save password** as needed (if credentials are used for authentication). This means forcing a client to save the password for this connection.
2. Specify the domain name (if credentials are used for authentication).

Secondary Connection

If you have more than one RAS Secure Gateway, you can define a secondary connection, which will be used as a backup connection in case the primary gateway connection fails.

To add a secondary connection:

1. Select the **Secondary connections** item.
2. In the **Secondary connections** pane, click **Tasks > Add** and specify a server name or IP address.
3. Select the connection mode and modify the default port number if necessary.

If you have multiple secondary connections, you can move them up or down in the list. If the primary connection cannot be established, Parallels Client will use secondary connections in the order listed.

Reconnection

In this pane, specify what to do if the connection is dropped:

- **Reconnect if connection is dropped.** If this option is selected, Parallels Client will try to reconnect if the connection is dropped. The **Connection retries** property specifies the number of retries.
- **Show connection banner if reconnection is not established within.** Specifies the number of seconds after which the connection banner will be displayed in Parallels Client. This will inform the user that the connection was dropped and will allow them to take action on their own.

When using the Parallels RAS clients for Windows or Mac, the user is presented with a greyed-out version of the published desktop or Published App(s) during the reconnect stage. This overlay indicates that the session is in the process of reconnecting, helping users understand that the session is not lost, but the connection to the remote session is being restored. This applies to published desktops and published apps.

Computer name

Specify the name that a computer will use during a remote desktop session. If set, this will override the default computer name. Any filtering set by the administrator on the server side will make use of the **Override computer name** setting.

Advanced settings

- **Connection timeout.** The Parallels Client connection timeout value.
- **Show connection banner if connection is not established within.** Specifies the number of seconds after which the connection banner will be displayed. This will inform the user that the connection cannot be established and will allow them to take actions on their own.
- **Show desktop if published application does not start within.** If a published application is not launched within the time period specified in this field, the host server desktop will be shown instead. This is helpful if an error occurs on the server side while launching an application. By showing the server desktop, the user can see the error message.

Web authentication

- Select or clear the **Use default OS browser** option. If the option is selected, the SAML SSO login dialog on the client side will open in the default browser.

If the option is cleared, the browser built into the Parallels Client will be used.

Note: To launch SAML SSO login dialog with built-in browser while using Parallels RAS Console 19.3 or later, use Parallels RAS Client for Windows 19.3 or later.

- The **Open browser window to complete log out** option is used when the built-in browser is used. In this case, there's no control over the SAML log out, so when this option is selected, a URL will open to perform the logout from SAML. By default, this web page will not be displayed, but if you need to interact with the browser, you can enable this option.

For more info, see [SAML SSO Authentication](#).

Session prelaunch

When a user opens a remote application, a session must first be launched. Launching a session can take time, which will result in the user waiting for the application to start. To improve user experience, a session can be launched ahead of time, before the user actually opens an application.

To enable (or disable) session pre-launch, choose one of the following in the **Mode** drop-down list:

- **Off.** No session prelaunch is used.
- **Basic.** A session is pre-launched as soon as the user gets the application listing. The assumption is, the user will open an application within the next few minutes. The session will stay active for 10 minutes. If the user doesn't open an application during that time, the client will disconnect from the session.
- **Machine Learning.** When the application listing is acquired, a session is pre-launched based on user habits. With this option enabled, Parallels Client will record and analyze the user habits of launching applications on a given day of the week. A session is started a few minutes before the user usually opens an application.

When a session is prelaunched, it will all happen in the background, so the user will not see any windows or message boxes on the screen. When the user starts an application, it will open using the prelaunched session, so it will start very quickly.

You can configure rules that govern when session pre-launch must not be used. The following options are available:

- Use the **Exclude sessions prelaunch** list to specify dates on which the prelaunch must not be used. Click on the plus-sign icon and select a date. The list can contain multiple entries.
- You can also exclude a published resource from the session pre-launching scheme altogether. This way, the resource is excluded from the analysis and is never considered by Parallels Client when making a decision whether to prelaunch a session. For example, when you have a server on which you never want to pre-launch sessions, you can flag all published resources hosted by that server to be excluded from session pre-launch. To exclude a published resource from session prelaunch, in the RAS Console, navigate to **Published Resources**, select a resource, and then select the **Exclude from session prelaunch** option.

Local proxy address

The setting in this section specifies to which IP address to bind the local RDP proxy. Select the **Use 127.0.01 IP address when using Gateway mode in VPN scenarios** option. You should have this setting enabled. Disabling it may lead to users not being able to open applications or desktops when using a VPN. This setting applies to Parallels Client for Windows only.

Display

To configure display settings, select the **Display** node and then configure the groups of settings described below.

Settings

Select the desired video acceleration mode and color depth.

Multi-monitor

Specify which monitors should be used for a session if more than one monitor is connected to the user's computer.

The following options are available:

- **All:** All displays.

- **Primary:** User's primary display.
- **Selected:** User can select one or several displays manually. To use this option for a published desktop, you need to select **Full Screen** in **Publishing** category > select the published desktop > **Desktop** tab > **Desktop Size**.

Published applications

Specify the options as follows:

- **Use primary monitor only:** Select this option to start published applications on the primary monitor. Other monitors connected to a user's computer will not be used.
- **Use dynamic desktop resizing:** Select this option if you want published resources to use the display settings of the local desktop.
- **Enable Z-order mode (Experimental):** Select this option if you encounter an issue related to z-order in applications published with Parallels RAS. Most issues can be resolved by selecting **(4) Improved heuristic mode**. If this does not help, try the solution described in [Resolving z-order issues](#).

Desktop options

Specify the desktop options as follows:

- **Smart-sizing:** Choose a smart sizing option. The **Scale (fit to window)** option scales a remote desktop to fit the connection window. The **Resize (update resolution)** option updates the resolution dynamically (without the need to reconnect) based on the window size. To disable smart sizing, select **Disabled**.
- **Embed desktop in launcher.** Enable this option to access a published desktop inside Parallels Client.
- **Span desktop across all monitors.** Enable this option to span published desktops across all connected monitors.
- **Connection bar in full screen.** Specify whether the connection bar should be pinned, unpinned, or hidden when connecting in full-screen mode.

Browser

This section applies to Parallels Web Client only. Specify whether a remote application should open in the same or a new tab in a web browser by default.

Printing

The **Printing** node in the **Policy Properties** dialog allows you to configure printing options.

In the **Technology** section, select the technology to use when redirecting printers to a remote computer:

- **None.** No printer redirection will be used.
- **RAS Universal Printing technology.** Select this option if you want to use RAS Universal Printing technology.
- **Microsoft Basic Printing Redirection technology.** Select this option if you want to use Microsoft Basic printing technology.
- **RAS Universal Printing and Microsoft Basic redirection technologies.** Select this option to use both Parallels RAS and Microsoft technologies.

Note: The following rules apply when using printing in RAS HTML 5 Client. If **None** or **Microsoft Basic Printing** is selected, then no printing redirection will be available in a remote session. If **RAS Universal Printing** or **RAS Universal Printing and Microsoft Basic Printing** is selected, then RAS Universal Printing will be used in a remote session.

RAS Universal Printing

If you selected **RAS Universal Printing technology**, use the **Redirect Printers** drop-down list to specify whether to redirect all printer on the client side, default printer only, or specific printers.

If you select **Specific only** in the step above, click **Tasks > Add**. Type a printer name and then click the **Options** button. In the dialog that opens, specify settings described below.

In the **Choose Format** drop-down list, select a data format for printing:

- **Print Portable Document Format (PDF).** Adobe PDF. This option does not require you to install any local applications capable of printing a PDF document. All the necessary libraries are already installed together with Parallels Client.
- **View PDF with external application.** To use this option you must have a local application installed which is capable of viewing a PDF document. Note that not all applications are supported. For example, the built-in PDF viewer

in Windows is not supported, so you must have Adobe Acrobat Reader (or a similar application) installed.

- **Print PDF with external application.** This option works similar to the View PDF option above. It also requires an application capable of printing a PDF document installed locally.
- **Enhanced Meta File (EMF).** Use vector format and embedded fonts.
- **Bitmap (BMP).** Bitmap images.

In the **Client printer preferences** section, select one of the following:

- **Use server preferences for all printers.** If this option is selected, a generic printer preferences dialog will be shown when a user clicks **Print** in a remote application. The dialog has only a minimal set of options that they can choose.
- **Use client preferences for all printers.** With this option selected, a local printer preferences dialog will open when a user clicks **Print** in an application. The dialog will contain a full set of options for a particular printer that the user has installed on their local computer. If they have more than one printer installed, a native preferences dialog will open for any particular printer that they choose to print to.
- **Use client preferences for the following printers.** This option works similar to the **Use client preferences for all printers option** (above), but allows users to select which printers should use it. Select this option and then select one or more printer in the list below. If a printer is not selected, it will use the generic printer preferences dialog, similar to the first option in this list.

Default printer settings

To configure default printer settings, click the **Change Default Printer settings** button.

The default printer list shows printers that can be redirected by the client to the remote computer:

- To disable the default printer, select **<none>**.
- To redirect the default local printer, select **<defaultlocalprinter>**.
- When **<custom printer>** is selected, you can specify a custom printer. The first local printer that matches the printer name inserted in the **Custom** field will be set as the default printer on the remote computer.

Select **Match exact printer name** to match the name exactly as inserted in the **Custom** field. Please note that the remote printer name may not match the original

printer name. Also note that local printers may not redirect due to server settings or policies.

The **Force Default printer for** option specifies the time period, during which a printer will be forced as default. If the default printer is changed during this time after the connection is established, the printer is reset as default.

Select the **Update the remote default printer if the local default printer is changed** option to change the remote default printer automatically when the local default printer is changed. Please note that the new printer must have been previously redirected.

A Windows 10 and 11 note

Windows 10 and 11 have a feature that automatically sets the default printer to the one used most recently or more often. This can break the default printer control on RD Sessions Hosts, guest VMs, and Remote PCs. To resolve this issue, the default printer management in Windows 10 and 11 should be disabled. To disable this feature using the Group Policy, do the following:

1. Open the group policy editor.
2. Navigate to **User Configuration > Administrative Templates > Control Panel > Printers**.
3. Find the **Turn off Windows default printer management** policy and enable it.
4. Force the group policy to all computers attached to the domain.

You can also disable the default printer management in Windows 10 and 11 locally by using the GUI or the registry editor:

1. On a Windows 10 or 11 computer, click **Start**, then click the “gear” icon which will open the **Settings** page.
2. On the **Printers and Scanners** tab, set the **Let Windows manage my default printer** option to **OFF**.

Using the registry editor:

1. Open the registry editor (regedit).
2. Navigate to HKEY_CURRENT_USER\NT.
3. Create a new DWORD item and name it **LegacyDefaultPrinterMode**.
4. Change the item’s Value data to hexadecimal and set the value data to **1**.

In addition to disabling the default printer management, the **Download over metered connections** option should be enabled in **Settings > Devices > Printers & Scanners**.

Scanning

On the **Scanning** node in the **Policy Properties** dialog, you can specify a scanner that should be used when one is required by a published application:

- **Use.** Allows you to select a scanning technology. RAS Universal Scanning uses TWAIN and WIA redirection allowing an application to use either technology depending on the hardware type connected to the local computer. If you select **None**, scanning will be disabled.
- **Redirect Scanners.** Select scanners attached to your computer for redirection. You can select **All** (all attached scanners will be redirected) or **Specific only** (only the scanners you select in the provided list will be redirected).

Audio

This node in the **Policy Properties** dialog allows you to configure remote audio playback and recording settings.

In the **Remote audio playback** section, Use the **Where** drop-down list to select one of the following remote audio playback options:

- **Bring to this computer.** Audio from the remote computer will play on your local computer.
- **Do not play.** Audio from the remote computer will not play on your local computer and will be muted on the remote computer as well.
- **Leave at remote computer.** Audio will not play on your local computer but will play normally on the remote computer.

Use the **Quality** drop-down list to adjust the audio quality:

- **Dynamically adjust based on available bandwidth.** This option will increase or decrease the audio quality based on your connection speed. The faster the connection, the higher audio quality setting will be used.

- **Always use medium audio quality.** The audio quality is fixed at the medium level. You can use this option when you don't require the best possible audio quality and would rather use the available bandwidth for graphics.
- **Always use uncompressed audio quality.** The audio quality is fixed at the highest level. Select this option if you have a very fast connection and require the best possible audio quality.

The **Enable recording (if applicable)** option allows you to enable audio recording on the remote computer. For example, you can speak into a microphone on the local computer and use a sound recording application on the remote computer to record yourself.

Keyboard

On the **Keyboard** node in the **Policy Properties** dialog, select how you want to apply key combinations (e.g. Alt+Tab) that you press on the keyboard:

- **On the local computer.** Key combinations will be applied to Windows running on the local computer.
- **On the remote computer.** Key combinations will be applied to Windows running on the remote computer.
- **In full screen mode only.** Key combinations will be applied to the remote computer only when in the full-screen mode.

Select or clear the **Send unicode characters** as needed.

On the **Keyboard** node in the **Policy Properties** dialog, you also have the ability to enable **Redirect remote keyboard input** to allow keyboard input and events from remote applications to be sent to the currently active applications on the local device.

Local devices and resources

Use the **Local devices and resources** node in the **Policy Properties** dialog to configure how local resources are used in a remote session.

Clipboard

Enable or disable the clipboard in a remote session. The **Clipboard redirection** drop-down menu allows you to limit the functionality of the clipboard:

- **Client to server only:** Copy and paste from the client to a server app only.
- **Server to client only:** Copy and paste from a server app to the client only.
- **Bidirectional:** Copy and paste in both directions.
- **Disabled:** The clipboard is disabled.

The **Allow clipboard redirection for** section allows you to further customize the selected **Clipboard redirection**. For each option, select the checkboxes to allow that clipboard type from either client to server, server to client, or both.

Note: Customization of **Clipboard redirection** options applies to Parallels Client for Windows only.

The **Advanced** section allows you to specify more advanced options to further customize the clipboard redirection.

- Select **Allow remote clipboard content to be saved in local clipboard history** to allow clipboard items copied from a remote session to be saved in the client's local clipboard history (Ct r l+V). This feature is currently only supported on the Parallels RAS Client for Windows.
- Select **Allow remote clipboard content to sync across devices** to allow clipboard data received from a remote session to be synced via the client's Cloud Clipboard feature. This feature is currently only supported on the Parallels RAS Client for Windows, and a Windows signed-in Microsoft account is required.

Note: The advanced settings apply to Parallels Client for Windows only.

For more information on clipboard sync and clipboard history on Windows, visit [Clipboard History for Productivity | Microsoft Windows](#)

Note: Clearing the **Clipboard** option also disables the Remote Clipboard functionality for affected users in Parallels Web Client. For more information, please see [Using the remote clipboard](#).

Disk drives and folders

Select the **Allow disk drives and folders redirection** option and select the local drives you want to redirect, or select **Use all disk drives available**.

Select the **Redirect as read-only drives** option to redirect all selected disk drives in read-only mode.

Note: When you select the **Redirect as read-only drives** option, the drag-and-drop functionality in Parallels Client for Windows becomes limited. Users will be able to use drag-and-drop only for copying file paths from local to remote computers.

Note: When you select the **Redirect as read-only drives** option, the drag-and-drop functionality in Parallels Client for Mac becomes limited. Users will be able to use drag-and-drop only for copying data from local to remote computers.

If you select the **Use also disk drives that I plug in later** option, disk drives that you connect to a local computer later will be automatically available in a remote session.

Note: This option applies to Parallels Client for Windows only.

In the **Cache** drop-down list, you can select whether to enable drive redirection cache, which makes file browsing and navigation on the redirected drives much faster:

- **Disable:** Drive redirection cache is disabled.
- **Enable:** Drive redirection cache is enabled.
- **Fast mode:** Same as above, but certain decorative features of File Explorer are disabled in favor of faster browsing.

Note: This option applies to Parallels Client for Windows only.

Devices

On this pane, specify whether to redirect local devices in general, use all devices available, and also devices that will be plugged in later.

Local devices that can be redirected include supported Plug and Play devices, media players based on the Media Transfer Protocol (MTP), and digital cameras based on the Picture Transfer Protocol (PTP).

Please note that disk drives and smart cards are redirected using dedicated **Disk drives and folders** and **Smart cards** options.

Video capture devices

Specifies video capture devices to redirect from a user device to the remote session. This is a high-level redirection that allows redirecting a composite USB device, such as a webcam with a microphone.

- **Allow devices redirection:** Allows to choose which video capture devices to redirect.
- **Use all devices available:** Redirect all available devices.
- **Use also devices that I plug in later:** A device that is plugged in after a session is started will also be used. Note that if this option is disabled, you will need to restart a session for a newly plugged-in device to become available.

Ports

Select whether to redirect LPT and COM ports.

Smart cards

Select whether to redirect smart cards or Windows Hello for Business. Note that if a smart card is selected as the authentication type in the **Primary connection** pane, the smart card redirection is automatically enabled, and this option is dimmed and inactive.

WebAuthn

Select whether to redirect Windows Hello and security keys from the local device to the remote session.

Note: This policy only applies to Parallels Client for Windows.

Location

Select whether to redirect location information from the local device to the remote session.

Note: This policy only applies to Parallels Client for Windows.

Pen and touch input

Enables or disables the following functions:

- Pen input redirection with pressure sensitivity support.

Note: Pen input redirection is supported on remote desktops with the following operating systems: Windows Server 2016 up to Windows Server 2025, Windows 10 version 1607 up to Windows 11.

- Windows touch input redirection. Windows touch input redirection allows users to use Windows native touch gestures from touch-enabled devices, including touch, hold, and release actions. The actions are redirected to remote applications and desktops as corresponding mouse clicks. This option allows you to disable touch input redirection in case of app compatibility issues.

Note: This policy only applies to Parallels Client for Windows and Parallels Web Client.

Multimedia redirection for AVD

Allows watching video content played in a browser on a remote Azure Virtual Desktop host. To use this feature, you also need to configure redirection on your AVD hosts as described at <https://learn.microsoft.com/en-us/azure/virtual-desktop/multimedia-redirection?tabs=edge#requirements>.

Note: This policy is applicable to Parallels Client for Windows 10 1909 and later, and Windows 11.

Note: Multimedia redirection on Azure Virtual Desktop is not available when using the Advanced client feature set.

Note: Multimedia redirection on Azure Virtual Desktop is currently in preview. For the list of websites that support multimedia redirection, see <https://learn.microsoft.com/en-us/azure/virtual-desktop/multimedia-redirection-intro>.

File transfer

Enables file transfer in a remote session. To enable file transfer, select this node and then select a desired option in the **Allow file transfer** drop-down list in the right pane. For additional information, see [Configuring remote file transfer](#).

Experience

The **Experience** node in the **Policy Properties** dialog allows you to tweak connection speed and compression.

Performance

Choose your connection speed to optimize performance: Choose a connection type according to your situation and then select experience options you want enabled. If you are connecting to a remote server on a local network that runs at 100 Mbps or higher, it is usually safe to have all of the experience options enabled. If you choose **Detect connection quality automatically**, the experience options will be enabled by default, but some may be dynamically disabled depending on the actual connection speed.

Enhance windows move/size: Enable this option if your users experience graphics artifacts (dark squares) while moving or resizing a remote application window on their desktops. The issue may manifest itself when a remote application is hosted on a Windows Server 2016, 2019 or 2022 and when the **Show contents of window while dragging** option is enabled. The issue does not appear with any other versions of Windows.

Compression

It is recommended to enable compression to have a more efficient connection. The available compression options are described below.

Enable RDP Compression: Enables compression for RDP connections.

Universal printing compression policy: The compression type should be selected based on your environment specifics. You can choose from the following options:

- **Compression disabled.** No compression is used.
- **Best speed (uses less CPU).** Compression is optimized for best speed.
- **Best size (uses less network traffic).** Compression is optimized to save network traffic.
- **Based on connection speed.** The faster the connection speed, the lower compression level and the minimum data size to compress are used.

Universal scanning compression policy: This drop-down list has the same options as the universal printing compression above. Select the compression type based on your environment specifics.

Network

Use the **Network** node in the **Policy Properties** dialog to configure a proxy server for Parallels Client.

Select the **Use proxy server** option and then select the protocol from the following list:

- **SOCKS4.** Enable this option to transparently use the service of a network firewall.
- **SOCKS4A.** Enable this option to allow a client that cannot connect to resolve the destination host's name to specify it.
- **SOCKS5.** Enable this option to be able to connect using authentication.
- **HTTP 1.1.** Enable this option to connect using a standard HTTP 1.1 protocol connection.

Specify the proxy host's domain name or IP address and the port number.

For SOCKS5 and HTTP 1.1 protocols, select the **Proxy requires authentication** option. For authentication, select the **Use user logon credentials** option or specify a user name and password in the fields provided.

Server authentication

Use the **Server authentication** node in the **Policy Properties** dialog to specify what should happen if authentication of an RD Session Host, Remote PC, or Guest VM fails.

In the **If authentication fails** drop-down list, select one of the following options:

- **Connect.** The user can ignore the certificate of the server and still connect.
- **Warn.** The user is alerted about the certificate and still has the ability to choose whether to connect or not.
- **Do not connect.** The user is not allowed to connect.

Advanced settings

The **Advanced Settings** node in the **Policy Properties** dialog allows you to customize the default behavior of Parallels Client.

You can specify the following properties:

- **Use client system colors:** Enable this option to use the client system colors instead of those specified on the remote desktop.
- **Use client system settings:** Enable this option to use the client system settings instead of those specified on the RD Session Host.
- **Create shortcuts configured on server:** For each published application, the administrator can configure shortcuts that can be created on the client's desktop and the Start menu. Select this option to create the shortcuts, or clear the option if you don't want to create them.
- **Register file extensions associated from the server:** For each published application, the administrator can create file extension associations. Use this option to either register the associated file extensions or not.
- **Redirect URLs to the client device:** Enable this option to use the local web browser when opening 'http:' links.
- **Redirect MAILTO to the client device:** Enable this option to use the local mail client when opening 'mailto:' links.
- **Always ask for credentials when starting applications:** If this option is enabled, a user will be asked to enter credentials when starting an application even if the session is still active. You can use this option as added security to prevent unauthorized users to access applications. For example, if a user disconnects from a session, no one else will be able to take over the session and run remote applications. As another example, if a user leaves a device with an open User Portal displaying the app listing (with or without running RDP sessions) then any user who tries to open a new application or another instance of a running application will be prompted for credentials. Please note that the [Auto login](#) option must be disabled for this functionality to work; otherwise saved credentials will be used automatically.
- **Allow Server to send commands to be executed by client:** Enable this option to allow commands being received from the server to be executed by the client.
- **Confirm Server commands before executing them:** If this option is enabled, a message is displayed on the client to confirm any commands before they are executed from the server.

- **Network Level Authentication:** Check this option to enable network level authentication, which will require the client to authenticate before connecting to the server.
- **Redirect POS devices:** Enables the Point of Service (POS) devices such as bar code scanners or magnetic readers that are attached to the local computer to be used in the remote connection.
- **Use Pre Windows 2000 login format:** If this option is selected, it allows you to use legacy (pre-Windows 2000) login format.
- **Disable RDP-UDP for gateway connections:** Disables RDP UDP data tunneling on the client side. You can use this option when some clients experience random disconnects when RDP UDP data tunneling is enabled on the RAS Secure Gateway (the **Network** tab in the gateway **Properties** dialog), while other clients are not.

Configure client policy options

The **Client options** node allows you to configure client policy options. Select the node and then select and configure individual items under it as described below.

Connection

On the **Connection** pane, specify the following options:

- **Connection Banner.** Select a banner to display while establishing a connection.
- **Automatically refresh connected RAS connections every [] minutes.** Select this option and specify the time interval to automatically refresh a connection. This will refresh the published resources list in Parallels Client.
- **When all sessions are closed.** Specifies what happens when all user sessions are closed:
 - **Do nothing.** Nothing happens.
 - **Lock workstation.** The computer is locked.
 - **Sign out from workstation.** The current user is signed out from their account.

Note: The **Lock workstation** option is not supported on the [devices managed in the Kiosk mode](#).

Logging

Specify a log level for Parallels Client. Choose from the following options:

- **Standard**
- **Extended**
- **Verbose**

You should normally use the Standard logging. When you have an issue with Parallels Client, you can temporarily raise the log level by selecting Extended or Verbose and setting the start date/time and a duration. Note that the start date and time correspond to the local client time zone. Parallels Client must be running in order for the logging to take place. If Parallels Client is launched when Extended or Verbose levels are supposedly already in effect, the level will stay on for the remainder of the original duration setting. If a policy changes during this time, the actual log level settings will be reapplied accordingly.

Update

Select **Check for updates on startup** and specify an update URL if you want Parallels Client to check for updates when it starts. The URL can point to the Parallels website, or you can store updates on your local network and use this local URL. For the information on how to configure a local update server, please read <https://kb.parallels.com/123658> .

The section **Azure Virtual Desktop Client** allows you to configure settings related to how resources published in Azure Virtual Desktop are launched using the Parallels Client for Windows.

- **Use Azure Virtual Desktop client.** If this option is selected, the Microsoft Azure Virtual Desktop (Remote Desktop) client will be installed and used to connect to resources published in Azure Virtual Desktop.
- **Use Windows App (verified version).** If this option is selected, the latest Parallels-verified version of the Microsoft Windows App client will be downloaded and used to connect to resources published in Azure Virtual Desktop. The latest Parallels-verified version of the Microsoft Windows App client is specified in the XML file hosted by Parallels at <https://download.parallels.com/ras/WVDCClientSupport.xml>
- **Use Windows App custom (unverified version).** If this option is selected, a specific version of the Microsoft Windows App client will be downloaded and used to connect to resources published in Azure Virtual Desktop. You can use a specific version of the Microsoft Windows App by hosting your own

custom version of the XML file. Note that the URL you specify needs to be accessible from all client devices targeted by this policy.

Note: These options work with Parallels Client for Windows only. Parallels Client for Mac can be updated only from the App Store. Parallels Client for Linux does not support this feature.

PC keyboard

These settings control the on-screen software keyboard used by Parallels Client for iOS, Parallels Client for Android, and the User Portal (HTML5). When the PC keyboard is enabled, Parallels Client displays a custom on-screen keyboard that sends scancodes to the remote session. When it is not enabled, the default operating system keyboard is used, which sends Unicode characters instead.

Specify the following options:

- **Force use PC keyboard.** Select this option to force the use of the PC keyboard on supported clients. End users can also enable this manually in the client (for example, in Parallels Client for Android under **Settings > User input > Use PC keyboard**).
- **Keyboard layout.** Select a keyboard layout from the drop-down list. Note that the selected layout can and will only be used in a Parallels Client version that supports this particular layout.

Note: These settings apply to Parallels Client for iOS, Parallels Client for Android, and the User Portal (HTML5) only. Desktop clients (Parallels Client for Windows, Mac, and Linux) always send scancodes, and there is no policy option to apply a keyboard layout to them.

macOS Keyboard Mapping

To determine whether local macOS shortcut keys are mapped to the remote session and the corresponding Microsoft keyboard shortcuts, set the shortcuts here. If you enable the policy, be sure to set at least one shortcut key, or users may lose functionality.

macOS KeyText	Windows Key Mapping
^ Left Control	Left Control
^ Right Control	Right Control
⌘ Left Command	Left Windows

macOS KeyText	Windows Key Mapping
⌘ Right Command	Right Windows
⌥ Left Option	Left Alt
⌥ Right Option	Right Alt (AltGr)
⇧ Left Shift	Left Shift
⇧ Right Shift	Right Shift

Single sign-on

Parallels Client for Windows comes with its own SSO component that you can install and use to sign in to Parallels RAS. If you already use a third-party credential provider component on your Windows computers, you first need to try whether the single sign-on works right out of the box. If it doesn't, you need to configure Parallels RAS and Parallels Client to use the Parallels RAS SSO component to function as a wrapper for the third-party credential provider component.

To use Parallels RAS SSO as a wrapper, specify a third-party component, select the **Force to wrap third party credential provider component** option and specify the component's GUID in the field provided. You can obtain the GUID in Parallels Client as follows:

1. Install Parallels Client on a computer that has the third-party component installed.
2. In Parallels Client, navigate to **Tools > Options > Single Sign-On** (tab page).
3. Select the "Force to wrap..." option and then select your provider in the drop-down list.
4. Click the **Copy GUID to Clipboard** button to obtain the component's GUID.

You will also need to specify the component's GUID when setting up an invitation email in the RAS Console. If you haven't set up an invitation email yet, you can do it as follows:

1. In the RAS Console, select the **Start** category and then click the **Invite Users** item in the right pane.
2. On the second page of the wizard (target platform and connection options), click the **Advanced** button.
3. In the dialog that opens, select the **Force to wrap third party SSO component** option and specify the GUID of the component.

For more information, see the [Invite users](#) section.

After the policies are applied to Windows computers, Parallels Client will be automatically configured to use the specified third-party credentials provider.

Advanced

Use this pane to specify advanced client options, as described below.

Global

- **Always on Top.** With this feature enabled, other applications will no longer mask the launcher.
- **Show connection tree.** Displays the connection tree.
- **Minimize to tray on close or escape.** Enable this feature to place the Parallels Client into the System Tray when you click on the **Close** button or hit escape.
- **Enable graphic acceleration (Chrome client).**
- **Do not warn if server certificate is not verified.** When connected to a RAS Secure Gateway over SSL, and the certificate is not verified, a warning message will be displayed. You can disable this warning message by enabling this option.
- **Swap mouse buttons.** When enabling this setting, the mouse buttons will be swapped on the remote computer.
- **DPI aware.** This will force a published application to be DPI-aware depending on the client's DPI settings. This feature works on Windows 8.1 or higher.
- **Add RAS Connection automatically when starting web or shortcuts items.** This option will add the connection preferences in the Parallels Client when starting an item contained in a connection that is not yet listed.
- **Do not show prompt message for auto add RAS connection.** Enable this option to disable prompt messages when adding auto connections.
- **Close error messages automatically.** When a session disconnects because of an error, the error is automatically dismissed after 15 seconds.
- **Clear session cookies on exit.** When a user logs on, a Parallels RAS logon cookie is kept on the client side. This will allow the user to connect again with Parallels RAS without re-authenticating. Check this option to delete any cookies when the user closes the Parallels Client.
- **Enable extended logging.** Enables extended logging.

- **Turn off UDP on Client:** Turns off UDP traffic from Parallels Client for Windows.

Language

Specify a language that Parallels Client should use. The **Default** option uses the main language used by the client's operating system.

Printing

- **Install missing fonts automatically.** If automatic fonts are installed on the server, they will be available when a session connects.
- **Raw printing support.** When enabling this setting, printing will still work for applications sending data in RAW format.
- **Convert non distributable fonts data to images.** During RAS Universal Printing, if a document includes non-distributable fonts, each page is converted to an image.
- **Dynamic printer redirection.** This feature allows you to change your redirected printers without signing out of Parallels Client.
- **Cache RAS Universal Printing embedded fonts.** Caching of embedded fonts locally to speed up the RAS universal printing process time.

Windows client

- **Hide Launcher when application is launched.** If this option is enabled, the launcher will be minimized in the system tray after an application is launched.
- **Launch automatically at Windows startup.** This option will place a shortcut in the start menu folder of the client and the Parallels Client will launch automatically on Windows startup.

RemoteFX USB redirection

- **Allow RDP redirection of other supported RemoteFX USB devices to all users.** This setting applies to Parallels Client for Windows only. Outside Parallels RAS, the standard RemoteFX USB redirection feature must be enabled via Group Policy in order to work. When you select the "Allow RDP redirection ..." option on this screen, it will do the same as GPO, which is update the corresponding registry setting in Windows on a client machine. Parallels Client for Windows relies on this feature to be enabled in Windows registry in order to redirect USB devices. When the policy containing this setting is applied on a client machine, the user will see a message that

RemoteFX USB redirection was enabled and that they will need to restart Windows.

Customer Experience Program

- **Participate in Customer Experience Program.** This setting allows you to join Parallels Customer Experience Program. For more information about Parallels Customer Experience Program, see <https://www.parallels.com/about/legal/pcep/>.

Configure control settings

Control settings options allow you to control various actions on the client side. These options affect the following Parallels Clients:

- Windows
- Linux
- Mac
- Android
- iOS

Connections

On the **Connections** pane, select (or clear) the following options:

- **Prohibit adding of RAS connections.** When a user presses the **Add Connection** button, an RDP connection is always created.
- **Prohibit adding standard RDP connections.** When a user presses the **Add Connection** button, a RAS connection is always created

Password

On the **Password** pane, specify the following options:

- **Prohibit saving username.** Parallels Client will not display the username of the last user who logged in. Selecting this option automatically enables the **Prohibit saving password** option.

Note: This setting is only available on Linux and Windows clients.

- **Prohibit saving password.** The option to save the password will not be shown to the user for that particular connection. A password is never saved on a disk, but kept in memory until the user closes the application.
- **Prohibit changing password.** The option to change the password will not be shown in the context menu for that particular connection.

Import and export

On the **Import and Export** pane:

- **Prohibit importing settings.** If this option is selected, the user cannot import connection settings to Parallels Client.
- **Prohibit exporting settings.** If this option is selected, the user cannot export connection settings from Parallels Client.

Configure Gateway redirection

Redirection options allow you to move your existing users from one RAS Secure Gateway to another gateway within the same Farm, or you can even redirect users to a gateway in a different Farm.

Note: When setting gateway redirection, make sure that the gateway criteria (the **Criteria** node) does not conflict with it. Read the **Gateway criteria** subsection at the end of this section for the explanation.

To configure redirection options:

1. Select the **Redirection** node in the left pane of the **Policy Properties** dialog.
2. In the right pane, specify the new connection properties, including:
 - **Gateway address**
 - **Connection mode**
 - **Port number**
 - **Alternative address**

When this policy is applied to user devices, the following will happen:

- Parallels Client connection settings are automatically updated on each device.
- Parallels Client tests the new connection. If succeeded, the current connection policies are removed and new policies are added.

- If Parallels Client cannot connect to Parallels RAS using new settings, the application list will not be shown and an error message will be displayed saying that the redirection policy has failed to apply. The user will be advised to contact the system administrator.

Gateway criteria

If a policy has both **Redirection** and **Criteria** settings enabled and configured, a situation may occur when the policy is applied in an infinite loop on the client side, which will result in an error. Consider the following possible scenarios when this may happen:

- Parallels Client connects to gateway “A” and applies a policy, which redirects it back to gateway “A”. This will continue to loop until Parallels Client gives up and displays an error to the user, which will say, “Failed to apply redirection policy....”.
- Parallels Client connects to gateway “A” and applies policy “P1”, which redirects it to gateway “B”. As expected, Parallels Client connects to gateway “B” and applies policy “P2”, which redirects it back to gateway “A” where it all began. This will also continue to loop until Parallels Client gives up and displays the same error message as described above.

Once again, this may only happen if the **Criteria** node is enabled and specified gateways conflict with each other. To avoid it, make sure that the **Gateway criteria** option on the **Criteria** pane is set to **if Client is connected to one of the following gateways** and that the same policy is not applied again when Parallels Client is redirected to a new gateway.

Client policy backward compatibility

Starting with Parallels RAS v16.5, a new approach is used to manage client policies. In the previous versions, a client policy would apply the full set of parameters and replace the client settings completely hiding an enforced category. In RAS v16.5 (or newer), client policy settings are split into smaller groups with the ability to configure and enforce each group on the client side individually. For example, the administrator wants to re-design the policies to disable clipboard redirection only, leaving the rest of the local devices and resources settings available for the end users to control. In the previous version, this would not be possible. The new design allows an administrator to easily achieve this goal.

This section explains how the backward compatibility is achieved with older clients and how new clients retain compatibility with older server-side installations.

The new client policies implementation handles compatibility issues as follows:

- All settings found in older policies are sent to the client as if being sent from an older Parallels RAS server. When a client receives the policy, the **Connection properties** and **Options/Preferences** settings are set correctly from the old design point of view. If, however, the policy is configured in such a way that the user cannot change anything, the entire tab will be hidden (no need to display the options if all of them are disabled).
- The Parallels RAS Console handles old-style policy settings as if they are new and displays them using the updated graphical user interface.
- In terms of policies, when a Parallels RAS v16.5 client connects to a previous version of Parallels RAS, the client keeps working normally and all of the policy settings are functioning as expected.

Policy information in Parallels Client

When a policy is applied to a user device, the information about it is displayed in Parallels Client. The information can be used to verify that the correct policy was delivered to a user device. The following information is included:

- **ID**: The policy ID as displayed in the **ID** field in the **Policies** list in the RAS Console.
- **Version**: The policy version number as displayed in the **Version** field in the **Policies** list in the RAS Console.
- **RAS Connection**: The name of the connection through which the policy was delivered. Displayed only on mobile devices and in Web Client.

By comparing the information above in Parallels Client running on a user device and the information in the RAS Console, you can see which policy was applied to a user device.

To see the applied policy information for a connection:

- In Parallels Client for Windows / Mac / Linux, open the **Connection Properties** dialog. The information is displayed at the bottom of a tab page to which the policy was applied.
- In Parallels Client for Android, the information is displayed at the bottom of the **Settings** screen.

- In Parallels Client for iOS, open the Edit RAS Connection screen and tap **View Applied Server Policy** (as the bottom).
- In RAS Web Client, the information is displayed in the **Settings** dialog.

Please note that when all of the connection properties in Parallels Client are managed through client policies, the user can still open the **Connection Properties** dialog, but it will contain a single tab displaying the applied policy information. If only some of the connection properties are managed through policies, the user will be able to see those tabs and the applied policy information that they contain.

When a policy includes global policy options, you can view the applied policy information in Parallels Client as follows:

- In Parallels Client for Windows and Linux, open the **Options** dialog (click **Tools > Options**).
- In Parallels Client for Mac, open **Preferences** (click **Parallels Client > Preferences**).

The applied policy information is displayed at the bottom of the dialog, similar to how it is displayed for the connection.

Configuring remote file transfer

Parallels RAS provides end users with the ability to transfer files remotely to and from a remote server.

Note: At the time of this writing, file transfer is supported in Parallels Web Client and Parallels Client for Chrome only. Bidirectional file transfer is supported in Parallels Web Client only. In Parallels Client for Chrome, you can only enable or disable file transfer.

To make the remote file transfer functionality flexible, Parallels RAS allows you to configure it on the following three levels:

- [RD Session Host, Provider, or Remote PC](#)
- [RAS User Portal](#)
- [Gateway settings in Web Client](#)
- [Client policy](#)

File transfer settings that you configure on each level take precedence in the order listed above. For example, if you enable file transfer on a User Portal, but disable it

on an RD Session Host, file transfer will be disabled for all users who connect to the given RD Session Host through the given User Portal. As another example, you can enable file transfer on an RD Session Host and then disable it for a particular Client policy (or an User Portal). This way you can control which clients can use file transfer and which cannot.

Read the subsequent sections to learn how to configure file transfer on each level.

Configure file transfer to a server

To configure remote file transfer for an RD Session Host, Provider, or Remote PC, do the following:

1. In the Parallels RAS Console, select the **Farm** category and then select a desired server type (RD Session Host, Provider, Remote PCs) in the middle pane.
2. Right-click a desired server in the right pane and choose **Properties**.
3. Select the **Agent Settings** tab.
4. Select the **Allow file transfer command** option and click the **Configure** button. A dialog opens where you can specify remote file transfer options as described below.
5. In the **Direction** drop-down list, select one of the following:
 - **Client to server only**: Transfer files from client to server only.
 - **Server to client only**: Transfer files from server to client only.
 - **Bidirectional**: Transfer files in both directions.
6. In the **Location** field, specify a UNC path to a folder to be used as the default upload location. This path will also be used as the default source location when a user tries to download a file from a remote server. You can select from one of the locations predefined in the drop-down list or you can specify your own. Standard Windows environment variables, such %USERNAME%, %USERDOMAIN%, %USERPROFILE%, can be used. If the location is not found during an upload or download operation, the standard (default) download location will be used.
7. The **Do not allow to change location** option prohibits the user to change the UNC path specified in the **Location** field. If the option is enabled, the

user cannot select a different location while trying to upload or download a file. If the option is cleared, the user can specify a different location.

Important: Please note that the **Do not allow to change location** option cannot prevent the user from accessing the specified remote location directly. For example, a user can try to upload a file, note the default location's UNC path (to which he/she has access), then open it in File Explorer and copy it to any folder in his/her profile. To prevent such a scenario from happening, you need to implement additional measures to control locations other than the location that you specify here.

8. **Suppress prompt for upload files:** Select this option to upload files immediately, without displaying a confirmation prompt for each individual file. When the option is cleared (default), the user is asked to confirm every file upload. Enabling this option also makes the **Existing file mode** drop-down list available.

Existing file mode: Specifies how an uploaded file is handled when a file with the same name already exists in the target location. This option is available only when **Suppress prompt for upload files** is enabled. Select one of the following:

- **Create new file:** the uploaded file is saved under a new, unique name and the existing file is kept.
- **Overwrite existing file:** the existing file is replaced with the uploaded file.

Close transfer window when complete: Select this option to automatically close the file transfer window once all transfers have finished. When the option is cleared (default), the transfer window stays open after completion so the user can review the results.

Configure file transfer in User Portal

To configure remote file transfer in User Portal, do the following:

1. In the Parallels RAS Console, navigate to **Farm** > <Site> > **Secure Gateways**.
2. Right-click a desired RAS Secure Gateway in the right pane and choose **Properties**.
3. Select the **User Portal** tab.

4. Select the **Allow file transfer command** option and click the **Configure** button. In the dialog that opens, select one of the following:
 - **Client to server only**: Transfer files from client to server only.
 - **Server to client only**: Transfer files from server to client only.
 - **Bidirectional**: Transfer files in both directions.

For more information about configuring User Portal, see [Configure User Portal](#).

Configure file transfer for a client policy

To configure remote file transfer for a client policy, do the following:

1. In the RAS Console, select the **Policies** category.
2. Right-click a desired policy in the right pane and choose **Properties**.
3. In the left pane, navigate to **Session > Local devices and resources**.
4. Select the **File transfer** node.
5. In the right pane, select one of the following in the **Allow file transfer** drop-down list:
 - **Client to server only**: Transfer files from client to server only.
 - **Server to client only**: Transfer files from server to client only.
 - **Bidirectional**: Transfer files in both directions.

For additional information about client policies, see [Client Policies](#).

Reporting

Parallels RAS Reporting is an optional RAS component that allows Parallels RAS administrator to run and view predefined and custom Parallels RAS reports. Predefined reports include user and group activity, device information, session information, and application usage. You can also create custom reports using your own criteria. Read this chapter to learn how to install and configure Parallels RAS Reporting and how to use it.

System requirements

Operating system requirements

Parallels RAS Reporting can be installed on a server running one of the following Windows Server versions:

- Windows Server 2025
- Windows Server 2022
- Windows Server 2019
- Windows Server 2016
- Windows Server 2012 R2

.NET Framework 3.5 and .NET Framework 4.5 or higher must be installed.

Microsoft SQL Server requirements

Parallels RAS Reporting can be used with the following Microsoft SQL Server versions:

- Microsoft SQL Server 2022
- Microsoft SQL Server 2019
- Microsoft SQL Server 2017
- Microsoft SQL Server 2016

Beginning with RAS 17.1, SQL Server Reporting Services (SSRS) and the SQL Server database engine can be deployed on separate hosts.

Using Microsoft SQL Server 2017 and 2019

Microsoft SQL Server 2017 and 2019 allow you to install the database engine and SQL Server Reporting Services (SSRS) on different hosts. Parallels RAS 17.1 (and newer) supports this deployment scenario and gives you the ability to use SQL Server Reporting Services and the SQL Server database engine installed on separate hosts.

Installation locations

RAS Reporting must be installed on the same server where SQL Server Reporting Services are running. Please note that if you have SSRS and the database engine

installed on different hosts, RAS Reporting must be installed where the SSRS are installed.

The following table contains RAS and SQL Server version compatibility information and locations where components necessary to use RAS Reporting can be installed:

RAS Reporting version	SSRS version	SQL Server version	Installation locations
19.0	2022	2022	SSRS - same host as RAS Reporting SQL Server - can be a different host
17.1, 18.0, 19.0	2019	2019	SSRS - same host as RAS Reporting SQL Server - can be a different host
17.1, 18.0, 19.0	2017	2019	SSRS - same host as RAS Reporting SQL Server - can be a different host
17.1, 18.0, 19.0	2017	2019	SSRS - same host as RAS Reporting SQL Server - can be a different host
17.1, 18.0, 19.0	2017	2017	SSRS - same host as RAS Reporting SQL Server - can be a different host
17.1, 18.0, 19.0	2017	2016	

RAS Reporting version	SSRS version	SQL Server version	Installation locations
			SSRS - same host as RAS Reporting SQL Server - can be a different host

Microsoft SQL Server must be installed as a named instance (not default or unnamed instance), because an instance must have a name for RAS Reporting to work. You can specify an instance name when you install Microsoft SQL Server (or when you create a new SQL Server instance in a multi-instance scenario). For complete details, please read the **Install Microsoft SQL Server** section that follows this one.

Install Microsoft SQL Server

Note: For Parallels RAS installations running on multiple servers, it is recommended that Microsoft SQL Server is installed on a dedicated server.
endhint %}

Install Microsoft SQL Server 2016 or earlier

To install a SQL Server instance (SQL Server 2016 or earlier):

1. Run the Microsoft SQL Server installation program and select the **Custom** installation type. Wait for the necessary files to be downloaded to your computer.
2. Once the files are downloaded, the **SQL Server Installation Center** window opens.
3. On the **Installation** page, select **New SQL Server stand-alone installation or add features to an existing installation**.
4. Click **Next** and follow onscreen instructions until you get to the **Feature Selection** page.

5. On the **Feature Selection** page, make sure that at least the following SQL Server features are selected for installation:
 - **Database Engine Services**
 - **Reporting Services - Native**
6. Click **Next**.
7. On the **Instance Configuration** page, select the **Named instance** option and enter an instance name. When naming the instance, you have the following options:
 - Enter “RASREPORTING”, which is the default instance name used by Parallels RAS Reporting. If you use this name, you don’t have to specify it later when installing RAS Reporting and configure it in the RAS Console. This is the recommended option.
 - If you would like to use a different name, you can do that, but you will have to make sure that you use this name when installing and configuring RAS Reporting. The instructions for installing RAS Reporting (described later in this chapter) indicate where the instance name must be specified, so if you follow the instructions, you won’t miss it. Note that the instance name cannot contain dashes, dots and some other characters.

After entering the instance name, make sure that it is also set in the **Instance ID** field.
8. Click **Next** and proceed to the **Database Engine Configuration** page.
9. On the **Database Engine Configuration** page, select the **Server Configuration** tab and add the following users to the **SQL Server administrators** list:
 - Local administrator (e.g. Administrator)
 - AD administrator (if you are just testing Parallels RAS Reporting on a local server, you can exclude this account).
 - SYSTEM (click **Add**, type “SYSTEM”, click **Check Names** and click **OK**; the account will appear in the list as “NT AUTHORITY”).
10. Complete the wizard using default settings on remaining pages.
11. Wait for the SQL Server installation to finish. On the **Complete** page, make sure that the installation was successful and exist the wizard.

Install SQL Server Management Tools

You should also install SQL Server management tools, specifically the SQL Server Management Studio. It is not required by RAS Reporting, but it is an essential SQL Server management tool that you might find useful. If you've never worked with SQL Server Management Studio before and not sure whether you need it, we suggest you install it. For example, you can use it to view RAS Reporting database tables, constraints, and stored procedures, which may help you better understand the RAS Reporting database design. The installation link is provided on the **SQL Server Installation Center** window.

Configure Microsoft SQL Server 2016 and earlier

When using Microsoft SQL Server 2016 and earlier version, it must be configured for remote connections as follows:

1. Open Microsoft SQL Server Management Studio.
2. Right-click on the server and select **Properties**.
3. Go to **Connections** and select **Allow Remote**.
4. Open SQL Server Configuration Manager and go to **SQL Server Network Configuration > Protocols** for RASREPORTING.
5. Right-click on **TCP/IP** and choose **Properties**.
6. Make sure the **Enabled** property is set to **Yes**.
7. Select the **IP Address** tab and locate the **IPAll** section. Set the **TCP Dynamic Ports** field to be blank and the **TCP Port** field to "1433".
8. Restart SQL Server. To do so, in the SQL Server Configuration Manager, right-click the SQL Server service and choose **Restart**.
9. After the restart, in the SQL Server Configuration Manager, right-click on **SQL Server Browser** and choose **Properties**.
10. Select the **Service** tab and set the **Start Mode** property to **Automatic**.
11. Start the SQL Server Browser.

Configure Microsoft SQL Server Reporting Services

To configure Microsoft SQL Server Reporting Services, follow these steps:

1. Run the Reporting Service Configuration Manager (**Start > Apps > Microsoft SQL Server 2016 > Reporting Services Configuration Manager**).

2. In the **Reporting Services Configuration Connection** dialog that opens, do the following:
 - Make sure the **Server Name** field contains the name of the server hosting the SQL Server instance.
 - Make sure the **Report Server Instance** field contains the name of the SQL Server instance that you've created earlier. If you used the default Parallels RAS name, it will appear as "RASREPORTING". If you used a different instance name, select that name.
3. Click **Connect**. If the connection is successful, the **Reporting Services Configuration Manager** window opens.
4. Select the **Web Service URL** category (not to be confused with Web Portal URL) in the left pane and set the following properties in the right pane:
 - **Virtual Directory**: Make sure that the directory name is "ReportServer_RASREPORTING". If you used a different name for the SQL Server instance, you should see that name instead of the "RASREPORTING" part.
 - **TCP port**: Set the port number to 8085.
5. Click the **Apply** button to apply the settings.
6. Select the **Web Portal URL** category in the left pane and then do the following:
 - Make sure that the **Virtual Directory** field is set to "Reports_<InstanceName>", where "InstanceName" is the name of your SQL Server instance. The default Parallels RAS name would be "Reports_RASREPORTING".
 - Examine the **URLs** field. Make sure that the port number after the server name is 8085. If it's not, click the **Advanced** button and change the port number.
7. Verify that you can access the Reporting Services Web Portal by clicking the URL on the **Web Portal URL** page. This should open the SQL Server Reporting Services home page in a web browser.
8. Click **Exit** to close the Reporting Services Configuration Manager.

Install Microsoft SQL Server 2017 or 2019

Microsoft SQL Server 2017 and 2019 allow you to install the database engine and SQL Server Reporting Services (SSRS) on separate hosts. Parallels RAS 17.1 (and newer) supports this deployment scenario and gives you the ability to use

SQL Server Reporting Services and the SQL Server database engine installed on separate hosts.

For step-by-step instructions on how to install and configure Parallels RAS Reporting Service with SQL Server 2019 and Microsoft SSRS 2019, please read the following Parallels KB articles:

- **Microsoft SQL Server 2017 and 2019 single server installation:** <https://kb.parallels.com/125164>.
- **Microsoft SQL Server 2017 and 2019 multi-server installation:** <https://kb.parallels.com/125156>.

Install Microsoft SQL Server 2022

Follow the instructions provided at <https://kb.parallels.com/125164>.

Install Parallels RAS Reporting

To install Parallels RAS Reporting:

1. Log in to the server where you have Microsoft SQL Server Reporting Services installed. Make sure you use the account with administrative privileges (AD).

Note: As was mentioned earlier, SQL Server 2017 and newer allow you to install SQL Server database engine and SQL Server Reporting Services (SSRS) on different hosts. You need to be logged in to the server where you have SSRS installed.

1. Download the latest version of Parallels RAS Reporting from <https://www.parallels.com/products/ras/download/links/>.
2. Once downloaded, double-click the RASReporting-xxx.msi file to run the installation wizard.
3. Follow the onscreen instructions and proceed to the **Database connection** page. Specify the SQL Server database engine location:
 - **Location:** If the SQL Server database engine is installed on the local server (together with SSRS), select **Localhost**. If the SQL Server is installed on a different server, select **Remote** and then specify the server connection properties (see below).

- **Server:** If you selected **Remote**, specify the FQDN or IP address of the server where you have SQL Server installed.
 - **Username:** Specify the username to log in to SQL Server.
 - **Password:** Specify the password.
4. On the same page, specify the SQL Server instance name. The default instance name is RASREPORTING. If you would like to use a different instance, you can specify it on this page. If the instance doesn't exist, you need to create it first.
 5. Click **Next**.
 6. On the **Viewing Reports User** page, you need to specify an Active Directory user who will be granted permissions to access the RAS reporting database. The default user is "rasreportingview" (note that the user must be created in Active Directory before it can be used here). You can specify a different Active Directory user if you wish, but you will need to change the reporting settings in the RAS Console before you can view reports (this change is described later in this chapter when the RAS reporting configuration is explained).
 7. Click **Next** to install Parallels RAS Reporting.

Configure RAS Reporting in the RAS Console

To configure Parallels RAS Reporting:

1. Log in to the Parallels RAS Console.
2. Select the **Administration** category and click the **Reporting** tab in the right pane.
3. In the **Reporting** tab, select the **Enable RAS Reporting** option.
4. In the **Server** field, specify the FQDN or IP address of the server hosting your SQL Server instance. The value in the **Port** field is used by the service which receives data from the RAS Connection Broker. The default port is 30008.
5. Specify a user login option by selecting one of the following:
 - **Prompt user for login details** — If this option is selected, the Parallels RAS Console user will be prompted to enter credentials before they can run a report.
 - **Use following credentials** — If this option is selected, the specified username and password will be used. The default (built-in) user name is RASREPORTINGVIEW. If you specified a different user when you installed RAS Reporting, specify that user credentials here.
6. To test the database connection, click the **Test connection** button.

Configure advanced settings

These settings are optional, so you can configure them according to your needs.

To access advanced settings:

1. On the **Administration > Reporting** tab page, click the **Tracking Settings** button. The **Advanced Setting** dialog opens.
2. In the **Session Information** section, specify the following options:
 - **Enable Tracking**. Records sessions data (affects all reports except server reports).
 - **Retain information for**. Select for how long the information should be kept in the database.
3. In the **Server Counters Information** section, specify the following:
 - **Enable Tracking**. If selected, server counter data is recorded (affects server reports only).
 - **Retain information for**. Select for how long the information should be kept in the database.
 - **Track CPU / Memory counter when change is more than (%)**. Use these two options to set the minimum CPU and Memory resource usage required to record data.
4. The **Custom reports** section is used to enable custom reports in the Parallels RAS Console. Select the **Enable custom reports** option and specify a folder name where custom reports will be stored (or use the default “Custom reports” name). Note that this is a virtual folder located on the SQL Server Reporting Services side, so you need to specify just a name (not a traditional path). You will see the folder in the Parallels RAS Console in the **Reporting** category together with other (predefined) folders that contain reports.

Running Parallels RAS Reports

To view Parallels RAS reports, select the **Reporting** category in the RAS Console. The report information is displayed as follows:

- The middle pane lists the available reports. See the **Predefined reports** subsection below for the complete list. The blue “folders” icon (at the top of the list) groups reports by type or displays all of them as a flat list. The “refresh” icon refreshes the report list by retrieving it from the database (this

can be useful when you enable/disable the reporting functionality or when you add custom reports, which may not appear in the list automatically).

- When you initially open the **Reporting** category, the right pane contains just the **Information** tab page, which informs you whether Parallels RAS Reporting is active. If it's not, you need to make sure that it is installed and enabled.
- The blue "square" icon in front of the **Tasks** drop-down list (upper right-hand side of the RAS Console) expands the reporting interface into full screen. The **Tasks** drop-down list allows you to perform the following actions: **Duplicate** (duplicates a report tab page), **Full screen** (on/off), various **Close Report** options, **Delegate Permissions** (allows you to grant permissions to view reports to other RAS administrators, such as Power and Custom administrators who don't have these rights).

To run a report, double-click it in the middle pane. The report opens in a tab page in the right pane:

- Most of the predefined reports include controls that you can interact with, such as **From/To** dates, **Sort By**, **Sort Order**, **Chart Type**, **Server Name**, and others depending on the report type. When you change a value in any of these controls, click the **View Report** button to apply the new values/options and re-run the report.
- The main report area (lower portion where the data is represented as a graph, text, or numbers) includes a menu bar with icons that allow you to change the magnification, list through report pages (if more than one is included), search for text, save a report to a file, print a report, and export it to one of the available formats (Word, Excel, PowerPoint, PDF, or a data feed).

Note: The first time the reports are viewed, you may be requested to add `https://<server domain/ IP>` as a trusted website. This will appear depending on the Parallels RAS machine's Internet Explorer Enhanced Security Configuration.

Predefined reports

Parallels RAS Reporting includes a number of predefined reports in the following groups:

1. **Users reports.** This group includes reports about how end users are interacting with Parallels RAS:
 - **Sessions activity for all users** — shows all sessions produced by all users in the system. The report shows information about each session and includes active time, idle time, disconnected and total time. A user is identified by username and IP address. The Secure Gateway information is also included.
 - **Sessions activity for user** — shows all sessions produced by a single user. The report shows information about each session and includes active time, idle time, disconnected and total time.
 - **Application usage for user** — shows applications used by a specified user, including number of times used and total time.
 - **Device usage for user** — shows information about devices used by a user. The report includes information such as device vendor, device model, and total time used.
 - **Operating system usage for user** — shows the operating system being used by a specified user.
 - **Full user information** — shows detailed information about a specified user.
2. **User groups reports.** These reports obtain information about how groups of users are interacting with Parallels RAS:
 - **Sessions activity for all groups** — shows all sessions produced by all groups in the system. The report includes active, idle, and disconnected time.
 - **Sessions activity for group** — shows all sessions produced by a group in the system. The report shows information about each session produced by each user in the group and includes start, end, active, idle, disconnect and total time.
 - **Application usage for group** — shows applications used by a specified group, including number of times used and total time.
 - **Device usage for group** — shows information about devices used by users as members of a specified group. The report includes device vendor, model and total time used.
 - **Client operating system usage for group** — shows the operating system used by members of a particular group.

3. **Devices reports.** This group includes reports about the devices that are connecting to Parallels RAS.
- **Devices used** — shows all devices using the system. The report includes a device manufacturer, model, and the number of sessions opened by the device.
 - **Client operating system used** — shows devices and corresponding operating systems that are using the system.
 - **Parallels client version used** — shows information about a device model, Parallels Client version used, and session information.
4. **Servers activity reports.** This group includes reports about the activity of Parallels RAS server components:
- **Sessions activity for RD session hosts** — shows the session activity of users on a particular RD Session Host. Report includes start, end, active, idle and disconnect time.
 - **Sessions activity for VDI provider** — shows the session activity of users on a particular Provider. Report includes start, end, active, idle and disconnect time (standalone Hyper-V and VMware ESXi only).
 - **Sessions activity for AVD provider** — shows the session activity of users on a particular AVD provider. Report includes start, end, active, idle and disconnect time.
 - **Session activity for RD Session Host Pool** — shows the session activity of users of a particular RDSH host pool. Report includes start, end, active, idle and disconnect time.
 - **Session activity for VDI Host Pool** — shows the session activity of users of a particular VDI host pool. Report includes start, end, active, idle and disconnect time.
 - **Session activity for AVD Host Pool** — shows the session activity of users of a particular AVD host pool. Report includes start, end, active, idle and disconnect time.
 - **Gateway tunnelled sessions** — shows tunnelled session information for a specified Gateway.
5. **Server health reports.** This group includes reports on server CPU and RAM usage for different components of Parallels RAS.
- **RD Session hosts health** — shows server CPU and RAM usage for a specified server in the Farm.
 - **Providers health** — shows server CPU and RAM usage for a specified provider in the Farm.
 - **Remote PCs health** — shows server CPU and RAM usage for a specified Remote PC in the Farm.

- **Gateways health** — shows server CPU and RAM usage for a specified Gateway in the Farm.
- **Connection Brokers health** — shows server CPU and RAM usage for a specified Connection Broker in the Farm.
- **Enrollment servers health** — shows server CPU and RAM usage for a specified Enrollment server in the Farm.

6. Application reports. Reports related to applications.

- **Activity for all applications** — shows information about applications used in the system. Report includes information such as application name, number of times used and the total usage time. When viewing this report, select “All applications” or “RAS published applications” depending on your needs. When the second option is selected, the report will not include non-published applications and duplicates.
- **Activity for application** — shows usage of an application by individual users during a specified time period. The information includes start time, end time, and total time for each session. Other information, such as host server name and session ID is also shown.

7. Logon duration reports. Reports that show detailed information about user logon duration. They also show information about connection duration, authentication duration, RAS policy lookup duration, host preparation duration, group policy load time, and desktop load time.

- **Logon duration for all users** — shows minimum, maximum, and average logon duration for all users on every server.
- **Logon duration for user** — shows minimum, maximum, and average logon duration for a specified user on every server.
- **Logon duration for RD Session Host** — shows minimum, maximum, and average logon duration on a specified RD Session Host for every user.
- **Logon duration for VDI Provider** — shows minimum, maximum, and average logon duration for specified Provider for every user.
- **Logon duration for AVD Provider** — shows minimum, maximum, and average logon duration for specified AVD provider for every user. The report also shows this information for connection duration, authentication duration, RAS policy lookup duration, host preparation duration, group policy load time, and desktop load time.
- **Logon duration for RD Session Host Pool** — shows minimum, maximum, and average logon duration on a specified RD Session Host Pool for every user.
- **Logon duration for VDI Host Pool** — shows minimum, maximum, and average logon duration for specified VDI Host Pool for every user.

- **Logon duration for AVD Host Pool** — shows minimum, maximum, and average logon duration for specified AVD Host Pool for every user.
8. **UX Evaluator reports.** Reports related to UX Evaluator, which is the time interval measured at the client between the first step (user action) and the last step (graphical response displayed).
- **UX Evaluator for all users** — shows UX Evaluator for all users on every server.
 - **UX Evaluator for user** — shows UX Evaluator for a specified user on every server.
 - **UX Evaluator for RD Session Host** — shows UX Evaluator for a specified RD Session Host for every user.
 - **UX Evaluator for VDI Provider** — shows UX Evaluator for a specified Provider for every user.
 - **UX Evaluator for AVD Provider** — shows UX Evaluator for a specified AVD provider for every user.
 - **UX Evaluator for RD Session Host Pool** — shows UX Evaluator for a specified RD Session Host Pool for every user.
 - **UX Evaluator for VDI Host Pool** — shows UX Evaluator for a specified VDI Host Pool for every user.
 - **UX Evaluator for AVD Host Pool** — shows UX Evaluator for a specified AVD Host Pool for every user.
9. **Transport protocol reports.** Reports that show how long each transport protocol is used during sessions.
- **Transport protocol for all users** — shows information about the used transport protocols for all users.
 - **Transport protocol for user** — shows information about the used transport protocols for a specified user.
 - **Transport protocol for RD Session Hosts** — shows information about the used transport protocols for a specified RD Session Host.
 - **Transport protocol for VDI Provider** — shows information about the used transport protocols for a specified Provider.
 - **Transport protocol for AVD Provider** — shows information about the used transport protocols for a specified AVD Provider.
 - **Transport protocol for RD Session Host Pool** — shows information about the used transport protocols for a specified RD Session Host Pool.
 - **Transport protocol for VDI Host Pool** — shows information about the used transport protocols for a specified VDI Host Pool.
 - **Transport protocol for AVD Host Pool** — shows information about the used transport protocols for a specified AVD Host Pool.

10. **Connection quality reports.** Reports that show information about connection quality.
- **Connection quality for all users** — shows information about connection quality for all users.
 - **Connection quality for user** — shows information about connection quality for a specified user.
 - **Connection quality for RD Session Hosts** — shows information about connection quality for a specified RD Session Host.
 - **Connection quality for VDI Provider** — shows information about connection quality for a specified Provider.
 - **Connection quality for AVD Provider** — shows information about connection quality for a specified AVD Provider.
 - **Connection quality for RD Session Host Pool** — shows information about connection quality for a specified RD Session Host Pool.
 - **Connection quality for VDI Host Pool** — shows information about connection quality for a specified VDI Host Pool.
 - **Connection quality for AVD Host Pool** — shows information about connection quality for a specified AVD Host Pool.
11. **Latency reports.** Reports that show information about session latency.
- **Latency for all users** — shows information about session latency for all users.
 - **Latency for user** — shows information about session latency for a specified user.
 - **Latency for RD Session Hosts** — shows information about session latency for a specified RD Session Host.
 - **Latency for VDI Provider** — shows information about session latency for a specified Provider.
 - **Latency for AVD Provider** — shows information about session latency for a specified AVD Provider.
 - **Latency for RD Session Host Pool** — shows information about session latency for a specified RD Session Host Pool.
 - **Latency for VDI Host Pool** — shows information about session latency for a specified VDI Host Pool.
 - **Latency for AVD Host Pool** — shows information about session latency for a specified AVD Host Pool.
12. **Bandwidth availability reports.** Reports that show information about bandwidth availability.
- **Bandwidth availability for all users** — shows information about bandwidth availability for all users.

- **Bandwidth availability for user** — shows information about bandwidth availability for a specified user.
 - **Bandwidth availability for RD Session Hosts** — shows information about bandwidth availability for a specified RD Session Host.
 - **Bandwidth availability for VDI Provider** — shows information about bandwidth availability for a specified Provider.
 - **Bandwidth Availability for AVD Provider** — shows information about bandwidth availability for a specified AVD Provider.
 - **Bandwidth availability for RD Session Hosts Pool** — shows information about bandwidth availability for a specified RD Session Host Pool.
 - **Bandwidth availability for VDI Host Pool** — shows information about bandwidth availability for a specified VDI Host Pool.
 - **Bandwidth Availability for AVD Host Pool** — shows information about bandwidth availability for a specified AVD Host Pool.
13. **Session disconnect reports.** Reports that show the most frequent reasons for disconnecting and the number of reconnections.
- **Session disconnect for all users** — shows top disconnect reasons and the number of reconnections for all users.
 - **Session disconnect for user** — shows top disconnect reasons and the number of reconnections for a specified user.
 - **Session disconnect for RD Session Hosts** — shows top disconnect reasons and the number of reconnections for a specified RD Session Host.
 - **Session disconnect for VDI Provider** — shows top disconnect reasons and the number of reconnections for a specified Provider.
 - **Session disconnect for AVD Provider** — shows top disconnect reasons and the number of reconnections for a specified AVD Provider.
 - **Session disconnect for RD Session Hosts Pool** — shows top disconnect reasons and the number of reconnections for a specified RD Session Host Pool.
 - **Session disconnect for VDI Host Pool** — shows top disconnect reasons and the number of reconnections for a specified Provider.

GDPR compliance

The Parallels RAS reporting database contains information about users, which may possibly include personal user information. To conform to GDPR, Parallels RAS gives you the ability to clear user data from the database at any time.

Parallels RAS Reporting Tools is a simple application that you can use to perform this task. The tool is installed automatically when you install Parallels RAS.

To clear user data:

1. On the computer where you have Parallels RAS installed, navigate to C : \Program Files (x86)\Parallels\RAS Reporting.
2. In the folder specified above, locate and run the **RASReportingTools** application.
3. When the application starts, enter a user name in the **User data** field and click **Find user**. If the user is found, the user information is displayed. If the user is not found, it means that the RAS reporting database doesn't have any information about that user.
4. To see the user information contained in the RAS reporting database, click the **Show full user information** button. This will open the **Full User Information** report in a web browser (note that this report is also available in the **Reporting** category in the RAS Console). Examine the report to determine if any of the user information is subject to GDPR requirements.
5. To clear the user data, go back to the **Parallels RAS Reporting Tool** app and click the **Clear user data** button. When asked, confirm that you want to clear the data.

Performance monitor

Overview

Parallels RAS Performance Monitor is a browser-based dashboard designed to help administrators analyze Parallels RAS deployment bottlenecks and resource usage. The dashboard provides a visual display of performance metrics, which can be viewed in the Parallels RAS Console or in a web browser.

Components

Parallels RAS Performance Monitor consists of the following components:

- **InfluxDB database** — a database for storage of system performance data.
- **Grafana dashboard** — a browser-based dashboard providing a visual display of performance metrics.

- **Telegraf service** — a service that collects performance data on a server where it is installed. The service is installed automatically when you add a server to a Parallels RAS Farm and install a corresponding RAS Agent on it (e.g. RAS Secure Gateway Agent, RD Session Host Agent, Remote PC Agent, etc.).

How it works

The Telegraf service is stopped by default, so it doesn't collect any data. To start the service on each server in the Farm, the performance monitoring functionality must be configured and enabled in the Parallels RAS Console. Once enabled, the Telegraf service begins collecting a predefined set of performance counters at a fixed time interval (10 seconds). It then sends the collected data to the InfluxDB database for storage. To view performance metrics, the Parallels RAS administrator uses the dashboard (Grafana), which displays the visual representation of performance counters in real time.

The performance metrics are grouped in the dashboard by type (Session, CPU, Memory, Disk, etc.), so the administrator can view each group of metrics separately. The administrator can also select whether to view performance metrics for one or more specific servers or for all servers in the Farm or Site. In addition, the administrator can select a specific Site for which the data should be displayed.

Install RAS Performance Monitor

Requirements

Parallels RAS Performance Monitor is a separate component of Parallels RAS with its own installer. It can be installed on a dedicated server or on a server hosting any of the Parallels RAS components. When you run the installer, the InfluxDB database is automatically installed, and Parallels RAS-specific Grafana dashboards are created. For additional info, see the Installation subsection below.

The following firewall rules (open ports) are automatically added on the server where you install Parallels RAS Performance Monitor:

TCP port 8086 (used by the InfluxDB database).

Prerequisites

Grafana

The RAS Performance Monitor requires Grafana to visualize collected Parallels Remote Application Server performance metrics. This is installed and updated separately from Performance Monitor. This allows customers to manage and update Grafana independently of Parallels RAS release cadence and allows customers to stay current with their supported Grafana versions and editions while aligning with Grafana license terms based on their requirements.

Before installing Performance Monitor, Grafana should be installed on the same server where Performance Monitor will be installed. Configuration of a database connection is not required, as this will be configured during the Performance Monitor installation itself.

When installing Grafana, it is recommended to run Grafana as a service.

Both Grafana Enterprise and OSS are supported.

Currently supported versions of Grafana are version 11.x through to 13.x.

Upgrades

Important note for customers carrying out Performance Monitor upgrades from versions prior to 20.2.4.2 and 21.1.1.2.

When upgrading from a previous version (where RAS Performance Monitor installed Grafana directly), the installer will detect the bundled Grafana instance and remove it during the upgrade. As a result, any custom dashboards, data sources, users, or other configurations stored in that bundled Grafana will be removed. To mitigate this, customers are advised to back up any custom Grafana content before upgrading. The Parallels RAS dashboards themselves will be re-created during the Performance Monitor upgrade. Additional dashboards outside the scope of RAS can be added separately, provided they do not interfere with RAS-specific dashboards.

Installation

To install Parallels RAS Performance Monitor:

1. Download the Parallels RAS Performance Monitor installer from [here](#).

2. Run the installation wizard (the `RASPerformanceMonitor.msi` file) and follow the onscreen instructions.
3. Close the wizard when finished.

Using Parallels RAS Performance Monitor

Configure access to Parallels RAS Performance Monitor

To enable data collection and view the dashboard:


1. In the RAS Console, navigate to **Administration > Reporting**.
2. Select the **Enable RAS Performance Monitor** option (the **RAS Performance Monitor configuration** section).
3. Enter the FQDN or IP address of the server where you have the InfluxDB database and Grafana dashboard installed.
4. Click **Apply** to commit the changes.

Once you perform the steps above, the Telegraf service is started on each server in the Site, and the data collection begins.

Viewing performance metrics

Note: You should give Parallels RAS Performance Monitor some time to collect performance data before you can view it (about 1 hour on initial installation).

To view performance metrics, select the **Monitoring** category in the RAS Console. The data is displayed in the right pane. The logon is performed automatically, so no logon credentials are required.



The buttons on the **Monitoring** tab (below the dashboard area) are as follows:

- **Home.** Displays the home page. The button is useful when you click on an external link in the dashboard, which may take you to an external web page.
- **Refresh.** Reloads the current page.
- **Open in browser.** Opens the performance dashboard in a web browser.

To view metrics of a specific type, expand the desired category in the dashboard's main area. The categories include:

- **Session Information.** Displays the information about active sessions and disconnected sessions.


[Note: AVD and VDI sessions are not shown here.](#)⁷

- **CPU usage.** CPU counters.
- **Free memory.** Physical memory counters.
- **Disk usage.** Disk I/O counters.
- **Network usage.** Network interface I/O counters.
- **System information.** System information counters.
- **RAS Gateway information.** RAS Gateway counters.

Performance metrics are displayed in the dashboard as a graph. Different counters are displayed using different colors.

To zoom in on a particular area of a graph, select a rectangular block with a mouse. You can also use the **Zoom** controls at the top of the dashboard for time range zoom out, shift time forward, or shift time backward. To select a specific time range, click the “clock” icon at the top and then specify a time range.

By default, the dashboard opens in kiosk mode. To exit it, press **Esc**. To cycle view mode, click the “monitor” icon in the upper right. When you exit the kiosk mode, the **RAS Infrastructure Health** page is displayed:



The menu at the top has the following items:

- **Hosts.** Allows you to select one or multiple servers for which the performance metrics should be displayed. To display the data for all servers in the Site, select **All**. Please note that if you don't see any servers in the list, you need to wait for Parallels RAS Performance Monitor to collect the initial set of statistics. This only happens on initial installation.
- **Instance.** This item allows you to select a specific counter instance (if there's more than one). For Network counters, it is usually the name of a network interface. For Disk counters, it is a disk name. Other types of counters don't usually have multiple instances.
- **Farm & Site.** Select a Site for which to display the data. Selecting **All** displays the data for all sites in the Farm. If you have another RAS Farm, and the RAS Performance Monitor is configured and enabled in it, you can also select a Site from that Farm.
- **Agent Type.** Select a RAS agent type.
- **Group.** Select an RDS group.

For more information about performance metrics and their meaning, please refer to the following articles from Microsoft:

- <https://technet.microsoft.com/en-us/library/cc976785.aspx>
- <https://technet.microsoft.com/en-us/library/2008.08.pulse.aspx>

See also [RAS Performance Counters](#).

By default, any user can access the Performance Monitor page and view performance metrics. To increase security, you can configure the RAS Performance Monitor to use credentials so that only authorized users can view it.

First, remove anonymous authentication from the Grafana configuration file as follows:

1. Open file C:\FilesPerformance Monitor.ini.

2. In the file, look for the following:

```
##### Anonymous Auth
#####
```

```
[auth.anonymous]
```

```
# enable anonymous access
```

```
enabled = true
```

3. Change "enabled = true" to "enabled = false".

[Note: The user will be prompted to change the admin password automatically after disabling the anonymous access. After that, the password can be changed following the Grafana official documentation⁸ here.](#)

4. Restart the Grafana service.


5. Select the **Monitoring** category and log in to Grafana using the following credentials:

- **User:** admin
- **Password:** admin (if you changed the password before, use the current password).

6. Once logged in, press “Esc” and then click the “shield” icon > **Users**.



7. Click **New user** and create a new user.



8. You now need to add the user to your organization's list. To do so, in the **Users** list, click **Edit** to edit the user and then set the organization and make the user a **Viewer**.

9. Click **Add** to add the user to your organization's list. The user can now view the RAS Performance Monitor statistics.

Newer versions of Grafana may be incompatible with a browser installed in the OS. A modern Edge browser is supported and will be embedded automatically in the RAS Console. But if it's not so, we suggest using a redistributable browser that comes in a form of a plugin. Plugins are available as separate msi packaged and distributed the same way as all other Parallels RAS distributions.

To avoid potential issues, Parallels RAS monitors which versions of Grafana and browser are installed and informs the RAS administrator about it when necessary. In a situation like that, when you select the **Monitoring** category, you will not see a dashboard as usual. Instead, a message is displayed that an updated browser plugin must be installed. To install the plugin, follow the instructions below.

1. In the main menu (at the top), click **Tools** and choose **Plugins**.
2. In the **Plugins** dialog, locate the **Browser Engine** plugin. The **State** column should indicate that it is not installed.
3. Select the plugin and click one of the following buttons:
 - **Install online** — Downloads the plugin from the Internet and installs it. Click this button if you can access the Internet from your environment.
 - **Install offline** — In an offline environment, admin can't get the list of plugins and instead uploads the plugin zip from a local file server. Only installed plugins are listed in the plugins table.
4. After the plugin is installed, you can go back to the Monitoring category and this time you will see the Grafana dashboard.

To update or remove a plugin (if necessary), use the same instructions as above, but select a plugin that is marked "installed" and then click **Update online** (or offline) or **Uninstall**. Please note that custom dashboards are now maintained after upgrades from Parallels RAS 18 to future versions.

Cost Insights

Cost Insights

Warning: The costs shown in **Cost Insights** are estimates and may differ from your actual Azure invoices. Final billing is determined solely by Microsoft at the time of invoicing and may depend on such factors as currency exchange rates, discounts, or adjustments that are not reflected in this information.

This chapter focuses solely on the session host servers in each selected host pool. Costs associated with Parallels RAS infrastructure components, such as the Secure Gateway and Connection Broker, as well as Azure infrastructure components like the Azure Load Balancer, are currently not included in the calculations below.

System requirements

Parallels RAS Cost Insights can be installed on a server running one of the following Windows Server versions:

- Windows Server 2025
- Windows Server 2022
- Windows Server 2019
- Windows Server 2016

Other requirements include:

1. Parallels RAS Cost Insights requires Parallels RAS Reporting to be configured first.
2. Parallels RAS Cost Insights requires Azure Virtual Desktop to be configured first, including an Azure Virtual Desktop Provider.
3. Parallels RAS Cost Insights needs to be installed on the same machine as Parallels RAS Reporting and Microsoft SQL Server Reporting Services.
4. Parallels RAS Cost Insights required an App Registration to be configured in Azure. This App Registration can be the same as configured on the provider, and needs at least reader permission on all Resource Groups where Azure Virtual Desktop resources are located, and Billing Reader permissions on the Azure subscription.
5. Accessing Parallels RAS Cost Insights from within the Parallels RAS console requires the WebView2 runtime. This is installed by default on most operating systems, except for Windows Server 2022. If the WebView2 runtime is

missing, install it manually or follow the steps provided upon initially opening Cost Insights, which guide you through installing the Parallels RAS Browser Plugin.

Install Parallels RAS Cost Insights

To install Parallels RAS Cost Insights:

1. Log in to the server where Parallels RAS Reporting is installed. Make sure to use an account with administrative privileges (AD).
2. Double-click the `RASInstaller.msi` file to launch the Parallels RAS installation wizard. If you see a message that begins with “*This version of Parallels RAS is only for testing purposes*”, it means that it’s not an official build and should not be used in a production environment.
3. Follow the onscreen instructions.
4. Proceed to the **Select Installation Type** page and select **Custom**.
5. Select the component **Cost Insights** only, and unselect all other components.
6. Click **Next**.
7. On the **Firewall Settings** page, select **Automatically add firewall rules** to configure the firewall on this computer for Parallels RAS to work properly. See **Port Reference** for details.
8. Click **Next** and then click **Install**. Wait for the installation to finish and click **Finish**.
9. If you are upgrading your RAS installation, it is recommended to reboot all servers where components are upgraded.

Configure Parallels RAS Cost Insights

To configure Parallels RAS Cost Insights:

1. Log in to the Parallels RAS Console.
2. Select the Administration category and click the **Reporting & Cost Insights** tab in the right pane.
3. In the **Reporting** tab, select the **Enable RAS Cost Insights** option and click **Apply**.
4. Select the **Farm category** and click the **Providers** tab in the right pane.
5. Open the provider you want to collect costs for and switch to the **Advanced** tab.
6. Select **Collect costs for this provider**.

7. Optionally, modify the app registration to use, or use the same app registration as configured on the provider. Consult the requirements to confirm the specified app registration has the required permissions.
8. Click **Apply**.
9. Select the **Administration category** and click the **Reporting & Cost Insights** tab in the right pane.
10. To test the RAS Cost Insights connection, click the **Test connection** button.

Managing Parallels RAS Cost Insights

The **Cost Insights** page allows you to review and analyze the estimated Azure infrastructure costs associated with your session host pools. From this page, you can monitor both current and historical trends, evaluate cost efficiency, and identify opportunities for optimization using Parallels RAS features.

1. Log in to the Parallels RAS Console.
2. Select the **Cost Insights** category.

[**Note:** After the initial configuration of Cost Insights, cost data may take up to 24 hours to appear. Once initial processing is complete, costs will be visible.](#)⁹

3. Use the **Time Range** dropdown to select the reporting period, such as This week, Last month, or a custom date range. The selected range determines the time window used for all displayed cost calculations, metrics, and reports.
4. Use the **Selected Host Pool** dropdown to focus the analysis on a specific host pool. If multiple host pools exist, you can switch between them to compare costs and utilization.
5. The **Virtual Machine Type** field displays the VM size(s) used by the selected host pool (for example, *Standard_B2s*). This is for reference only and cannot be used as a filter. It helps you understand the cost breakdown in the context of the VM instance type(s) currently in use.
6. Click on the information icons next to each metric (e.g., *Host Pool Cost*, *Compute Cost*, *Parallels RAS Current Savings*) to see detailed definitions.
7. Compare **Parallels RAS Current Savings** with **Parallels RAS Potential Savings** to understand the actual and possible benefits of enabling optimization features. This helps you determine whether enabling disk storage optimization or power management can further reduce costs.
8. Use **Total Cost per User** to benchmark the cost efficiency of your host pool. Combine this with the **# of Active Users** to see how effectively your host pool resources are being utilized.

9. Scroll to the **Historical Analysis** graphs to view cost and cost-per-user trends over time. These visualizations help you correlate peaks in costs with usage patterns, scaling events, or configuration changes.

Definitions and calculations

General

All costs shown in the **Cost Insights Dashboard** are based on the UTC time zone. This is because both the Azure Cost Management API and the Azure Portal operate exclusively in UTC, and all cost data in Azure is emitted and stored in UTC.

For example, when selecting “Yesterday” as the time range, the dashboard displays costs from 00:00:00 to 23:59:59 UTC for that calendar day.

The costs shown in **Cost Insights** are estimates and may differ from your actual Azure invoices. Final billing is determined solely by Microsoft at the time of invoicing and may depend on such factors as currency exchange rates, discounts, or adjustments that are not reflected in this information.

Host Pool Cost

The total Azure cost for all session host virtual machines that are part of the selected host pool(s) within the specified time range. This includes three major components: **Compute**, **OS Disk Storage**, and **Network**. **Compute** represents the VM size and runtime costs, **OS Disk Storage** represents the managed disks attached to each session host, and **Network** represents outbound and inter-region traffic costs. This metric provides you with a consolidated view of the overall infrastructure cost directly associated with operating the session hosts in the selected host pool(s).

Calculation Details

Host Pool Cost = Σ (Compute cost of each session host) + Σ (OS disk cost of each session host) + Σ (Network cost of each session host)

Compute Cost

The compute cost of all Azure session host virtual machines that belong to the selected host pool(s) within the specified time range. The compute cost is based on the VM size (vCPU, memory, and GPU, if applicable) and the number of hours

each session host runs. Costs are calculated according to the Azure pricing model for the chosen VM sizes. This metric reflects the single most significant portion of host pool cost and is directly influenced by power management configurations, VM uptime, and workload demand.

Calculation Formula:

Compute Cost = Σ (Compute cost of each session host)

Storage Cost

The operating system (OS) disk storage cost for all session host virtual machines in the selected host pool(s) within the specified time range. This is based on the size and type of managed disks (Standard HDD, Standard SSD, or Premium SSD) provisioned for each session host. Storage cost reflects the recurring per-GB/month rate charged by Azure for disk allocation, regardless of how much of the disk space is actively used. Temporary disks, data disks, and ephemeral OS disks (if configured) are excluded from the calculation.

Calculation Formula:

Storage Cost = Σ (OS disk cost of each session host)

Network Cost

The network cost associated with all session host virtual machines in the selected host pool(s) during the specified time range. This cost is derived from Azure bandwidth charges, specifically for data transfer out across continents and inter-region ingress/egress between Azure regions. Intra-region and inbound data transfer are typically not billed, but outbound and cross-region transfers can add up significantly depending on workload patterns. This metric consolidates those Azure network costs attributable to session host activity.

Calculation Formula:

Network Cost = Σ (Network cost of each session host)

Parallels RAS Current Savings

The realized cost savings that have been achieved through the configuration of Parallels RAS optimization features in the selected host pool(s) during the specified time range. This currently includes two mechanisms:

- **Disk Storage Optimization:** savings by using optimized disks when VMs are stopped and deallocated.
- **Power management:** savings from powering VMs on or off based on actual demand.

This metric represents *actual* avoided costs compared to what would have been incurred if all session hosts had run continuously without optimization.

Calculation Formula:

- **Disk Storage Optimization Savings** (if enabled on the RAS provider):

Σ for each session host:

$((\text{Deallocated hours}) \times (\text{unoptimized disk price per hour})) -$

$((\text{Deallocated hours}) \times (\text{optimized disk price per hour}))$

- **Power management Savings** (if enabled):

Σ for each session host:

$(\text{Deallocated hours}) \times (\text{compute price per hour for VM type})$

Parallels RAS Potential Savings

The additional cost savings that could have been achieved if Parallels RAS optimization features (disk storage optimization and power management) had been fully enabled and utilized for the selected host pool(s) and time range. Unlike **Current Savings**, which represents realized reductions, this metric highlights the potential for further savings. It is a projection based on observed usage patterns and assumes full adoption of RAS scaling and storage policies.

Calculation Formula:

Potential Disk Storage Savings (if optimization is disabled):

Toff = time the host was deallocated

C_unopt = unoptimized disk cost per hour

C_{opt} = optimized disk cost per hour

Potential Storage Savings = $\sum$ for each session host: $T_{off} \times (C_{unopt} - C_{opt})$

Potential power management Savings (if power management is disabled):

$\sum$ for each session host:

(Hours the VM was running with no active sessions) $\times$ (compute price per hour for that VM type)

Total Cost per User

The host pool cost, divided by the number of active users during the selected time range. This provides an average cost-per-user metric that helps you measure the efficiency of your deployment. By normalizing total infrastructure cost against peak user load, this metric offers a benchmark for cost-effectiveness and helps you compare host pool performance against other pools or industry baselines.

Calculation Formula:

Total Cost per User = Host Pool Cost $\div$ Maximum # of Active Users

Maximum # of Session Hosts

The highest number of session host virtual machines that were simultaneously running across all selected host pool(s) during the specified time range. This metric reflects the historical peak scale of the environment and helps you analyze capacity requirements, evaluate power management performance, and plan for future demand.

Calculation Formula:

Maximum # of Session Hosts = Max (Count of session hosts in “Running” state over the selected time range)

of Active Users

The number of unique users who were active over the selected time range for all selected host pool(s). Active means the user had at least one active session selected in the time range for all selected host pool(s).

Calculation Formula:

Maximum # of Active Users = Max (Count of users with active sessions over the selected time range)

Historical Analysis: Cost per User

A time-series graph showing how the cost per user has varied over the selected time range for all selected host pool(s). Costs are normalized against active user counts for each time interval (e.g., hourly, daily, or weekly). This allows you to visualize trends in cost efficiency, identify peak inefficiencies, and assess the impact of scaling or optimization changes over time.

Historical Analysis: Host Pool Cost

A time-series graph showing the total host pool cost (compute, OS disk storage, and network) over the selected time range for all selected host pool(s). This allows you to monitor historical cost patterns, correlate costs with user activity or scaling events, and identify opportunities for cost optimization.

This chapter describes common Parallels RAS management tasks, including Farm status monitoring, license management, backup management, and others.

Recovery – add a root administrator

This topic addresses a possible issue when the root administrator is not available or the domain is changed. In such events, the system becomes inaccessible. If you encounter this issue, you can quickly add a root administrator by executing the following command on the server hosting the primary RAS Connection Broker:

```
2XRedundancy -c -AddRootAccount user [domain]
```

Please note that an open Parallels RAS console will not be notified about the new account since this is an emergency recovery. You need to log out and then log in again to see the new account in the **Administration** area.

Host name resolution

When you add a server component (Connection Broker, Gateway, RD Session Host, Provider, etc.) to a RAS Farm you have to specify its FQDN or IP address. It is normally up to you whether to use FQDN or IP address. On the other hand, the server IP address can change in the future. If that happens, you will have to

reconfigure the corresponding component in the RAS Farm. On the other hand, the server FQDN usually stays the same, so if you used it instead of the IP address, no RAS configuration changes will be necessary. For this reason, Parallels RAS gives you an option to always resolve IP addresses to FQDNs for all server components in a Farm.

To always use name resolution, do the following:

1. In the RAS Console, click **Tools > Options** on the main menu (that's the menu at the top of the RAS Console window).
2. In the **Options** dialog, select the **Always attempt to resolve to fully qualified domain name (FQDN) when adding hosts** option.
3. Click **OK**.

When you now try to add a component to a Farm and enter its IP address instead of a name, it will be automatically resolved to FQDN. If the FQDN cannot be determined, you will see an error message and will be asked if you would like to use the IP address instead.

The examples below demonstrate how the automatic name resolution works for different components.

Adding a RAS Connection Broker

1. On the **Connection Brokers** tab, click **Tasks > Add**.
2. In the **Server** field, enter the server IP address.
3. Click **Next**.
4. In the dialog that opens, observe that the IP address has been resolved to FQDN and the **Server** field contains the FQDN.

Adding a RAS Secure Gateway

1. On the **Gateways** tab, click **Tasks > Add**.
2. In the **Server** field, enter the server IP address.
3. Click **Resolve**. This will copy the IP address to the **IP(s)** field and will enable the **Next** button.
4. Click **Next**.
5. In the **Installing RAS Secure Gateway** dialog, observe that the server IP address is replaced with the FQDN.

Adding an RD Session Host

1. On the **RD Session Hosts** tab, click **Tasks > Add**.
2. On the first page of the wizard, enter the server IP address and click the plus-sign icon.
3. Observe that the server is added to the list, but the IP address is substituted with the FQDN that was automatically resolved.

Add a provider

1. In **Farm > Site > Providers** tab, click **Tasks > Add**.
2. Select the provider you want to add.
3. In the **Address** field, enter the IP address of a Provider.
4. Enter the remaining properties and click **Next**.
5. Observe that the Provider address is replaced with the FQDN.

Computer management tools

When you need to perform standard Windows computer management tasks, you can do it without leaving the RAS Console. The tasks include Remote Desktop Connection, Computer Management, Service Management, Event Viewer, PowerShell, Reboot, and others. To perform these tasks, use the **Tools** menu, which is accessible from the **Site** menu and individual Parallels RAS infrastructure servers and session hosts.

Requirements for using computer management tools

Some of the tools require an appropriate target host configuration before you can use them in the RAS Console. Please read the following requirements and make sure they are met.

To use Remote Desktop, remote connections must be enabled on a target host. You can verify that by using the standard Windows Remote Desktop Connection application and see if you can connect to a remote server.

PowerShell related tools require PowerShell remoting enabled on a target server. To enable PowerShell remoting, run the `Enable-PSRemoting` cmdlet on a target

computer in PowerShell window with administrator privileges. Please note the following:

- The cmdlet configures a computer to receive PowerShell remote commands.
- The cmdlet starts the WinRM (Windows Remote Management) service, among other tasks. To see if the WinRM service is running, use the `Test-WSMan` cmdlet.
- When you execute the cmdlet, it will ask you to confirm every task that it wants to perform. To execute the command silently, use the `-Force` option.
- If you receive an error saying that “WinRM firewall exception will not work since one of the network connection types on this machine is set to Public”, you can try to execute the cmdlet with the `-SkipNetworkProfileCheck` option, or you can change the network connection type on this host to Domain or Private.

To use PowerShell to manage a remote host, you also need to add the host to the TrustedHosts list on the computer where you have the RAS Console installed. To view the current TrustedHosts list, execute the following command in PowerShell window:

```
Get-Item WSMan:\localhost\Client\TrustedHosts
```

To add a host to the TrustedHosts list, use one of the options described below. Please note that all examples below, except the last one, always overwrite an existing TrustedHosts list. To add a specific computer to an existing list, use the last example (the one with the `-Concatenate` parameter).

Add all computers to the list:

```
Set-Item WSMan:\localhost\Client\TrustedHosts *
```

Add all domain computers:

```
Set-Item WSMan:\localhost\Client\TrustedHosts *.domain-name.dom
```

Add specific computers:

```
Set-Item WSMan:\localhost\Client\TrustedHosts ,[]
```

Add a computer to an existing list (this is the only example that will not overwrite an existing TrustedHosts list):


```
Set-Item WSMAN:\localhost\Client\TrustedHosts -Concatenate
<ComputerName>
```

Available tools

The table below describes the tools available in the **Tasks > Tools** menu and their execution strings.

Tool	Execution string	Description
Remote Desktop	mstsc.exe / v:<selectedRDShostName>:<port> /admin	Launch a standard RDP connection to the selected RDS host.
Computer Management	compmgmt.msc / computer:<selectedRDShostName>	Launch Computer Management locally with connection to the selected host.
Service Management	services.msc / computer:<selectedRDShostName>	Launch Services Management locally with connection to the selected host.
Event Viewer	eventvwr.msc / computer:<selectedRDShostName>	Launch Event Viewer locally with connection to the selected host.
Shared Folders	smgmt.msc / computer:<selectedRDShostName>	Launch Shared Folders locally with connection to the selected host.
Powershell	Enter-PSSession -ComputerName <selectedRDShostName> [- Credential username]	Launch Powershell locally with connection to the selected host.
IPconfig	• Powershell remote connection to selected host ◦ Get-NetIPConfiguration	Provides network configuration for the selected host.
Ping	• Powershell remote connection to selected host ◦ Test-NetConnection -ComputerName www.microsoft.com	Provides ICMP reply with status and RTT for the selected host.

Tool	Execution string	Description
	Select -ExpandProperty PingReplyDetails FT Address, Status, RoundTripTime	
Netstat	• Powershell remote connection to selected host ◦ Get-NetTCPConnection	Displays network connections for Transmission Control Protocol on the selected host.
Reboot	shutdown /m \<selectedRDSHostName> /f /r /t 0	Reboot the selected host.
Shutdown	shutdown /m \<selectedRDSHostName> /f /s /t 0	Shutdown the selected host.

Note that individual tools availability depends on the server type. For example, HALB has only **Ping** in the **Tools** menu.

Site information

To view the Site information, select the **Information** category in the RAS Console.

The **Site Information** tab displays information about available servers, Connection Brokers, Secure Gateways (see [Viewing Gateway summary and metrics](#)), and sessions on the local computer. To view information about running applications, select the **Show application information** option (at the bottom of the page).

The **Local Information** tab shows the status of RAS components running on the local server.

Site settings

To view and configure common Site settings in the RAS Console, navigate to **Farm** > <Site> > **Settings**.

Auditing

The **Auditing** tab allows you to configure application auditing. When enabled, application auditing monitors processes running in the Site and records this information in the audit file. To view the information, click the **View Audit** button (at the bottom of the page). The information is also displayed on the **Information > Site** page and in RAS Reports.

To enable or disable application auditing, use the **Auditing** drop-down list (at the bottom of the page). The **Clear Audit File** button clears the current audit.

The **Filtering the following processes** list allows you to specify processes that will be excluded from the audit. Use the **Tasks** drop-down list to add or delete a process. You can also use the **Task** menu to import and export a process list from/to a CSV file. The **Task > Properties** menu item allows you to edit a process name. The **Default** menu item resets the list to contain the default set of standard processes.

Global logging

The **Global logging** tab allows you to specify the log level for Parallels RAS components. Logs are used by Parallels RAS support engineers to analyze possible issues with a Parallels RAS installation. To specify the log level, select one or more servers in the list and click the **Configure Logging** item. In the dialog that opens, select one of the following:

- **Standard** — This is the standard log level that records only the most important events. Unless you are asked by Parallels RAS support to use one of the log levels described below, you should always use this one.
- **Extended** — This logging involves more information than the standard logging, but it slows down the system because of the additional information that it needs to collect.
- **Verbose** — Verbose logging involves even more information than the extended logging and can slow down your system significantly.

Please note that to avoid degraded performance, extended and verbose logging should only be enabled for a limited time period (enough to collect the necessary information for analysis). You can set this time period using **Reset to the standard level after** option. The default value is 12 hours. In specific cases, a Parallels support engineer will advise you whether this time period should be set to a

different value. Once this time period is over, the log level will be reset back to standard.

To retrieve a ZIP archive containing the collected log files, click the **Retrieve** item and then specify a location where you want the file to be saved. The **Clear** item clears all logs.

You can also set the log level for an individual server by navigating to the page where servers of that type are listed (e.g. RD Session hosts, Gateways, etc) and clicking **Tasks** (or right-click) > **Troubleshooting** > **Logging**. The context menu that opens has the same **Configure**, **Retrieve**, and **Clear** options as described above. The **Log Level** column in the server list indicates the currently set level.

URL redirection

The **URL redirection** tab allows you to create redirection rules that specify URLs which will be redirected when the **Allow Client URL/Mail redirection** option is enabled for an RD Session Host, virtual machine, or Remote PC (**Agent Settings** tab in the corresponding server properties). You can also deny redirection of specific URLs. A URL either will be opened on the client side if redirection is allowed or on the remote session host if redirection is denied.

Note the following about redirection rules:

- Redirection rules apply from top to bottom.
- Only the first rule that match the URL applies.
- A URL will be redirected if any part of it matches a rule. For example, `https://www.parallels.com`, `www.parallels.com`, `remoteapplicationserver`, and `www.parallels.com/remoteapplication` are all valid rules for redirecting `https://www.parallels.com/remoteapplicationserver`.

To add a redirection rule:

1. Configure Parallels Client on your client devices to allow URL redirection.
2. For certain applications, you need to enable URL redirection on your RD Session Host, , or Remote PC (see the **Agent Settings** tab in the corresponding server properties).
3. (RD Session Hosts only) For certain applications, you need to enable the **Replace registered applications** option on the server where the application is published. To do this, go to **Farm > RD Session hosts > right-click on your RD Session host > Properties > Agent Settings > Configure**.
4. In the RAS Console, navigate to **Farm > Settings**.

5. Select the **URL redirection** tab.
6. Click **Tasks > Add** (or click the **[+]** icon).
7. In the **URL** field, specify the URL that must be redirected.
8. In the **Action** drop-down list, select **Redirect** or **Do not redirect**.
9. Click **OK**.
10. Click **Apply**.

Notifications

See [System event notifications](#).

Client settings

See [Specifying client settings](#).

Features

See [FSLogix Profile Container](#) and [Enable Windows Virtual Desktop and add a provider](#).

Application packages

See [Using MSIX application packages](#).

Using MSIX application packages

Parallels RAS 19 provides a new and modern application delivery method - Application Packages, which is based on MSIX app attach technology. MSIX app attach is a Microsoft's application layering solution that allows you to dynamically attach applications (containerized MSIX packages) to a user session. Separating the application from the operating system makes it easier to get more control by providing the right application for the right user. Additional third party solutions such as appCURE may be used for application moderation with Parallels RAS.

Prerequisites

- [RD Sessions Hosts](#), [VMs](#), or [AVD hosts](#).

- MSIX app attach requires hosts running Windows Server 2025, Windows Server 2022, Windows 11, Windows 10 version 2004 or later.
- A network share where the MSIX images will be stored. Storage requirements and recommendations are highlighted here: <https://docs.microsoft.com/en-us/azure/virtual-desktop/app-attach-file-share>.
- All hosts (computer accounts) must have read permissions on the network share where the MSIX images are stored.

Note: In Parallels RAS 19, MSIX app attach applications can be deployed and managed directly from the Parallels RAS console only when using Windows Server 2022.

Enabling the Application Packages feature

To start working with MSIX application packages, you need to enable the Application Packages feature.

To enable the Application Packages feature:

1. Navigate to **Farm > Site > Settings** and select the **Application packages** tab.
2. Select the **Enable Application Packages feature** option.

Next, you need to add the package to Parallels RAS.

Creating an MSIX image

To create an MSIX package from any desktop installer such as MSI, EXE, ClickOnce, or App-V you can use the MSIX Packaging tool <https://docs.microsoft.com/en-us/windows/msix/packaging-tool/tool-overview>.

To expand MSIX-packaged applications into MSIX images you can use the MSIXMGR tool <https://docs.microsoft.com/en-us/azure/virtual-desktop/app-attach-msixmgr>.

Adding an MSIX application package to Parallels RAS

To add an MSIX application package to Parallels RAS:

1. Navigate to **Farm > Site > Application Packages**.
2. Click **Tasks > Add** (or click the **[+]** icon). The **Add from MSIX Image** wizard opens.

3. In the **MSIX Image path** field, specify the path to your image or click the **Browse** button to select it in File Explorer. The file must be located on a network share. You can add packages from VHD, VHDX, and CIM images. All hosts (computer accounts) must have read permissions on the network share where the MSIX images are stored.
4. In the **Package** drop-down list select the package that you want to add.
5. In the **Display Name** field specify the name that will be used for this package in Parallels RAS. After that, the rest of the fields will be populated automatically.
6. Click **Finish**.

Next, you need to add the package to a host. After being added, a packaged application behaves the same way as a regular application as if it was installed on a host.

Adding a package to a host

To add a package to a host:

1. Make sure that the package is added to Parallels RAS as described above.
2. Go to **Farm > Site > RD Session Hosts > RD Session Hosts**.
3. Double-click the host that you want to install the package to.
4. In the properties dialog, on the **Application Packages** tab, click **Tasks > Add** (or click the **[+]** icon).
5. In the first column to the left, select the packages that you want to install on the host.
6. In the **Version** column, select the version of the package. It is recommended to make use of [version tags](#), which can facilitate application version updates. If the selected packages have dependencies, you will see a warning that lists all of them.
7. Click **OK**.

The selected packages will be added to the host.

Adding a package to VDI pool

To add a package to a VDI pool:

1. Make sure that the package is added to Parallels RAS as described above.
2. Go to **Farm > Site > VDI > Pools**.
3. Double-click the pool that you want to install the package to.

4. In the properties, on the **Application Packages** tab, clear option **Inherit default settings**.
5. Continue from Step 4 as described in the subsection “Adding a package to a host” above.

The selected packages will be added to all VMs in the pool.

Adding a package to an AVD pool

To add a package to an AVD pool:

1. Make sure that the package is added to Parallels RAS as described above.
2. Go to **Farm > Site > Azure Virtual Desktop > Host pools**.
3. Double-click the pool that you want to install the package to.
4. In the properties, on the **Application Packages** tab, clear option **Inherit default settings**.
5. In the properties, continue from Step 4 as described in the subsection “Adding a package to a host” above.

The selected packages will be added to all hosts in the pool.

Adding a package to Group Defaults

To add a package to Group Defaults:

1. Make sure that the package is added to Parallels RAS as described above.
2. Go to **Farm > Site > RD Session Hosts > Groups**.
3. Double-click the group or pool that you want to install the package to.
4. In the properties, continue from Step 4 as described in the subsection “Adding a package to a host” above.

Adding a package to Site Defaults

To add a package to Site Defaults:

1. Make sure that the package is added to Parallels RAS as described above.
2. Open group or pool properties as described above.
3. On the **Application Packages** tab, click **Site Defaults**.
4. In the dialog that opens, continue from Step 4 as described in the subsection “Adding a package to a host” above.

Working with version tags

You can use version tags to simplify package management. For example, you can assign different tags to packages that are ready for publishing and these that are still in the testing stage. By default, Parallels RAS uses three tags: **Production**, **Pre-production**, and **Custom**. You can rename tags, but you cannot add or delete them.

To rename a tag:

1. Navigate to **Farm > Site > Settings** and select **Application packages** tab.
2. Select the tag that you want to rename.
3. Click **Tasks > Edit**.
4. Change the name of the tag and press Enter.

To use tags, you need to assign them.

To assign tags to a package:

1. Navigate to **Farm > Site > Application Packages**.
2. Double-click the package.
3. In the **Version tag** section, select the tags that you need. You can assign several tags to one package.

To remove all tags assigned to a package:

1. Navigate to **Farm > Site > Settings** and select the **Application packages** tab.
2. Select the package.
3. Click **Tasks > Remove all tags**.

Working with certificates

Parallels RAS uses code signing certificates to ensure authenticity and content integrity of MSIX application packages.

The following code signing certificates can be used:

- Self-signed certificates
- CA certificates
- Internal CA certificates

You can provision code signing certificates via GPO or let Parallels RAS install them to hosts automatically. The code signing certificate of a package must be trusted by all hosts that use that package.

Parallels RAS allows you to add certificates to hosts automatically. This option is recommended for self-signed certificates.

To enable automatic certificate provisioning:

1. Navigate to **Farm > Site > Settings** and select the **Application packages** tab.
2. Select option **Provision package certificates automatically**.

Certificate expiration dates are shown in **Farm > Site > Application** packages.

Managing MSIX application packages

You can manage the added packages on the **Farm > Site > Application Packages** tab.

The following actions are available in the **Task** drop-down list:

- **Add:** Adds a new package.
- **Change version tag:** Assigns tags to a package.
- **Remove all tags:** Removes all tags from the package.
- **Show published resources:** Opens the list of all published applications from the package and the hosts they were published from.
- **Show assigned session hosts:** Opens the list of hosts the selected package is assigned to.
- **Search:** Allows you to search for a package in the list by applying a filter.
- **Delete:** Deletes the package from Parallels RAS.
- **Settings audit:** Opens the Settings Audit dialog where you can view the changes that were done to the packages.
- **Refresh:** Refreshes the package list.
- **Properties:** Shows the properties of the package (see below).

Package properties

The following settings are available in the **Application package properties** window:

The **General** tab:

- **Enable application package in site:** Select this option to enable the package.
- **Package:** Name of the package.
- **Display Name:** Name used for the package inside Parallels RAS.
- **Version:** Version of the package.
- **Publisher:** Common name of the publisher.
- **MSIX image path:** Path to the MSIX image.
- **Version tag:** Tags assigned to the package. You can change the assigned tags to the package from here.
- **Applications:** List of applications added from the package.
- **Dependencies:** All dependencies of the package.
- The **Certificate** tab:
 - **Key size:** Size of the certificate.
 - **Expiration date:** Certificate expiration date.
 - **Common name:** Common Name specified in the certificate.
 - **View certificate info:** Shows information about the certificate.

Package statuses

Status color	Package status	Description
Green	Ready	Package is enabled and ready for registration.
Green	In use	Package is being used in a session.
Orange	Disabling	Waiting for deregistration in sessions.
Red	Staging failed	A problem with registration has occurred. You can retry registration as described in Application Packages .
Red	Not found	The image file or network location unavailable. The admin can retry staging.
Red	Certificate missing	Package certificate is missing on the host.
Red	No version found	There is no application package marked with the tag selected in the host configuration.
Not applied	Not applied	Settings were changed but not applied

Using template versions

Template versions allow you to safely test changes on your hosts and perform rollback if necessary.

Supported providers

Template versions are supported for the following providers:

Hypervisor Providers:

- Microsoft Hyper-V
- Microsoft Hyper-V Failover Cluster
- VMware ESXi
- VMware VCenter

Cloud Providers:

- Microsoft Azure

Virtualization services:

- Azure Virtual Desktop

Creating a new version

To create a new version:

1. Do one of the following:
 - To create a new version for an RD Session Host template, navigate to **Farm > Site > RD Session Hosts > Templates**.
 - To create a new version for a VDI host template, navigate to **Farm > Site > VDI > Templates**.
 - To create a new version for an AVD host template, navigate to **Farm > Site > Azure Virtual Desktop > Templates**.
2. Select the template, enter the maintenance mode, make changes and exit. You will see a dialog that prompts you to create a new template version. Select the **Create a new version** option.

Note: One template can have up to five versions. If you want to create another version, you will have to delete an already existing one.

3. Click **Next**.

4. On the **New template version** page, specify the name and description and select the tags for the version. If the tag was previously assigned to another version, it will be removed from this version. You can select several tags.
5. Click **Next**.
6. (Optional) On the **Select host pools** page, select the host pools that you want to recreate on schedule and click the **Configure** button. You will see a dialog that allows you to schedule recreation. You will only see the step if you specified a version tag that is in use by an existing host pool. Configure the schedule according to your needs and click **Next**.
7. Click **Finish**.

Working with version tags

To rename a tag:

1. Navigate to **Farm > Site > Settings** and select **Template versions** tab.
2. Double-click the tag that you want to rename.
3. Change the name of the tag and press Enter.

To reassign tags of a version:

1. Do one of the following:
 - To reassign tags for RD Session Host template, navigate to **Farm > Site > RD Session Hosts > Templates**.
 - To reassign tags for a VDI host template, navigate to **Farm > Site > VDI > Templates**.
 - To reassign tags for an AVD host template, navigate to **Farm > Site > VDI > Azure Virtual Desktop > Templates**.
2. Select a template and click **Tasks > Versions**. The **Versions** dialog opens.
3. Select a version and click **Tasks > Properties**.
4. In the **Version tag** section, select the tags that you need. If the tag was previously assigned to another version, it will be removed from this version. You can assign several tags to a version. If another host pool uses the same version with the same tags, you will see a dialog that allows you to recreate hosts in that host pool.

Deleting a version

To delete a version:

1. Do one of the following:
 - To delete a version of an RD Session Host template, navigate to **Farm > Site > RD Session Hosts > Templates**.
 - To delete a version of a VDI host template, navigate to **Farm > Site > VDI > Templates**.
 - To delete a version of an AVD host template, navigate to **Farm > Site > VDI > Azure Virtual Desktop > Templates**.
2. Select a template and click **Tasks > Versions**. The **Versions** dialog opens.
3. Select a version and click **Tasks > Delete**.

Settings audit

Settings audit allows you to see recent changes to versions.

To see settings audit:

1. Do one of the following:
 - For RD Session Host template version settings audit, navigate to **Farm > Site > RD Session Hosts > Templates**.
 - For a VDI template version settings audit, navigate to **Farm > Site > VDI > Templates**.
 - For an AVD template version settings audit, navigate to **Farm > Site > VDI > Azure Virtual Desktop > Templates**.
2. Select a template and click **Tasks > Versions**. The **Versions** dialog opens.
3. Select a version and click **Tasks > Settings audit**.

Settings audit

Parallels RAS gives you the ability to audit the modifications that were made to a Parallels RAS Farm, including changes to any of the components, objects, resources, and users. This information is stored in a database, so it can be reviewed and possibly reverted, if needed. The information is stored in the primary database and replicated to a local database on the computer where Parallels RAS Console is running.

You can view the list of modifications using one of the following options:

- By navigating to **Administration > Settings audit**. The tab displays the main list of all changes to any components/objects in the Farm. If a modification can be reverted, you can do it here.
- By clicking **Tasks > Settings audit** on any pane in the RAS Console that supports this functionality. Compared to the main list (described above), you will only see modifications to the same types of components or objects that are managed on a given pane. You can also revert a modification here if it can be reverted. If the **Settings audit** menu option is not available in a particular pane, it means the functionality is not available for the component or object types that this pane manages.

The following describes in detail how to view and revert Farm modifications.

View the main settings audit list

To view the main list of all modifications for a Farm, do the following:

1. In the Parallels RAS Console, select the **Administration** category and then click the **Settings audit** tab.
2. The sync process checks that the local audit database is in sync with the primary database and updates it if necessary (you may see a progress indicator while it syncs).
3. Once the syncing is complete, the **Settings audit** tab will be populated with data. Each entry in the list corresponds to a modification made either by a RAS administrator or by a system service.

Turning off audit databases synchronization

By default, Parallels RAS synchronizes audit databases across all Connection Brokers. This will take more time as the databases grow. You can turn off database synchronization, in which case you will only have access to the audit database on the current Licensing Connection Broker. If you change the Licensing Connection Broker, you will not have access to the previous audit database unless you enable synchronization.

To turn off the synchronization of audit databases:

1. On the **Settings audit** tab, click the **Tasks** drop-down list and select **Settings**.

2. Clear the **Replicate administrators' audit data on all Connection Brokers** option.
3. Click **OK**.

Information about modifications

The information for each entry in the list includes the following:

- **Date:** Date and time of the modification.
- **Session:** Session ID.
- **Username:** The name of the administrator or RAS service that was responsible for the modification. RAS services may include System (redundancy service) and Connection Broker (controller service).
- **Action:** The action that was performed, such as Connect, Disconnect, Create, Update, Switch site, and others.
- **ID:** The affected object's ID.
- **Site:** The number and name of the affected Site. "Global" means the change affected all sites.
- **Type:** The modification type. This usually makes sense when viewed together with the **Action** value.
- **Name:** The value in this column is displayed for some entries and can provide additional information, such as the name of the changed object.

Common tasks

You can perform the following actions on the list:

- To refresh the list, click the "recycle" icon (top right).
- To view details for an entry, double-click it (or select an entry and click **Tasks > View entry**).
- To search for a specific entry (or entries), click the magnifying glass icon (top right). An extra row is added at the top of the list, allowing you to enter the search criteria. You can type a string to search for in one or multiple columns. The search is performed as you type, and the list is filtered to include only the matching entries. To cancel filtering and display the complete list, click the magnifying glass icon again.

Reverting a modification

To revert a modification in the main list:

1. Double-click a desired entry on the **Settings Audit** tab.

2. The **Audit Entry** dialog opens. While here, you can click the **Next** and **Previous** buttons to go to the next or previous item as they are displayed in the main list.
3. To revert the change, click the **Revert** button. If the button is disabled, it means that the change cannot be reverted.

Changes that can never be reverted include the following:

- Any changes made by the System or Connection Broker (as displayed in the **Username** column).
- Changes made in previous versions of Parallels RAS that did not include this feature.
- Changes related to administrator accounts.

View a local settings audit list

You can also view and revert configuration changes for a specific type of RAS components or objects. When you are on a particular pane (or tab) in the RAS Console, look for the **Tasks > Settings audit** menu option (or right-click > **Settings audit**). If it's there, then you can view the changes and revert them if needed. Consider the example below.

Let's say you want to see the changes made to RD Session Hosts. To do so:

1. In the RAS Console, navigate to **Farm > <Site> > RD Session Hosts**.
2. Click **Tasks > Settings audit**.
3. The **Settings Audit** dialog opens, listing all known modifications that were made to RD Session Hosts. The modifications may include creating, moving, deleting, or updating an RD Session Host. The type of the modification is displayed in the **Action** column in the list.
4. To revert a modification, select it and click the **Revert** button (in the lower right of the dialog). If the button is disabled when you select a particular entry, it means that the modification cannot be reverted.

The local settings audit functionality is available for most of the major components and objects in the Parallels RAS Console. This includes RD Session Hosts (including Groups and Scheduler), VDI, Remote PCs, Gateways, Connection Brokers, Themes, Publishing, Quick Keypad, and many others. Once again, when you view a particular pane, look for the **Tasks > Settings audit** menu option (or right-click > **Settings audit**). If it's there, then you can view the changes and revert them if needed.

Change audit settings

You can also configure audit log archiving. Archived log entries are moved to a separate database. This is to ensure that the audit database does not grow too large, in which case, performance when accessing the audit logs could degrade. To enable audit log archiving:

1. On the **Settings audit** tab, click the **Tasks** drop-down list and select **Settings**.
2. Select **“Settings”**
3. select **“Enable audit log archiving.”**

Optionally, change the default archiving settings by specifying the following properties:

1. **Archive audit entries after.** The threshold to archive a log entry.
2. **Start:** The time at which the archiving action begins.
3. **Recur:** The recurrence of the archiving action.
4. **Day:** The day to run the archiving action (only applicable if recur is set to “Every week”).

Upgrading RAS agents

When you add Parallels RAS components to a Farm, you install a corresponding RAS Agent on them. This includes RAS Connection Broker, RD Session Host Agent, Provider Agent, Guest Agent, Remote PC Agent. In addition to the functionality that allows you to check agent status, and update it if necessary, you can do a bulk agent update or upgrade.

There are two ways you can find out if agents need to be updated. You can be notified by Parallels RAS or you can check the status and initiate the update procedure manually.

When the Parallels RAS Console starts, you may see a message box saying that Agents need to be installed or updated. You can start the update procedure by clicking **Yes** in this dialog. You will then see a list of all servers on which an Agent needs to be updated where you can decide whether to include a server in the bulk update procedure or exclude it. Once you’ve made your selection, follow the onscreen instructions and update the Agents.

To initiate the procedure manually, click the **Task > Upgrade all Agents** in the RAS Console where this menu is available (most of the views where it makes sense). You can also right-click inside the view and choose **Upgrade all Agents**. Follow the onscreen instructions and select the servers on which an Agent requires an update or upgrade. Please note that if all Agents on all servers displayed on a given pane are up to date, the menu option will be disabled.

For example, to upgrade all primary Connection Brokers in all sites, select **Farm > Farm** and then click **Tasks > Upgrade all Agents** (or right-click inside the pane and choose **Upgrade all Agents**). To upgrade all Agents on all servers in a Site, select **Farm > <Site>** and click **Tasks > Upgrade all Agents**. Similarly, to upgrade Agents on all RAS Secure Gateways, select **Farm > <Site> > Secure Gateways** and use the same **Tasks > Upgrade all Agents** menu item. For other components, do exactly the same. Note that if you use the same credentials on all servers, you will have to enter them only once. The update procedure will remember the last entered credentials and will try to use them on all servers. If the credentials don't work on a server, you'll be asked to enter them again.

Please note that after you click the **Tasks > Upgrade all Agents** menu, the dialog that opens will contain the hosts on which an Agent needs updating or upgrading. The **Status** column in the list will indicate that and the host will be preselected for the upgrade. Unverified Agents will also be included in the list but will not be preselected. You can select them if you believe that an Agent has to be upgraded on them too.

Note: When updating an agent in a template (VDI), full and linked clone templates are updated differently. For important information, please read the [Template maintenance](#) section.

Licensing

The **Licensing** category allows you to manage your Parallels RAS license. When you click on the **Licensing** category, the **License details** tab displays the license information. Note that depending on the license type (prepaid subscription, SPLA, NFR, trial), different information is displayed. Please also note that an NFR (Not for resale) license can be either a prepaid subscription or SPLA, so different information can also be shown for an NFR license.

The information that can be displayed on the **License details** tab includes the following:

- **License Type:** The type of your Parallels RAS license (e.g. prepaid subscription, SPLA, NFR, trial).
- **License Key:** The license key that was used to activate the Farm (only last characters are shown).
- **Support Expiration Date:** The support plan expiration date.
- **Upgrade Insurance:** The upgrade insurance expiration date which, in case of using subscription based licenses, is in line with Expiration date, since such licenses are automatically eligible for upgrades.
- **Expiration date:** The license expiration date (or the number of days remaining for a trial license).
- **First Activation:** The date the Farm was first activated.
- **Peak Users:** Peak concurrent users to date. For a prepaid subscription, you can use this value to evaluate whether you might need to upgrade your subscription to include more concurrent users.
- **Usage Today:** SPLA license only. The number of maximum concurrent users recorded for today. Note that “today” starts at midnight UTC.
- **Current Period Usage:** SPLA license only. The combined usage for all Farms activated with the same license key (SPLA licenses allow the activation of multiple farms with the same key).
- **Billing period started:** SPLA license only. The current billing period start date.
- **Billing period ends:** SPLA license only. The current billing period end date.
- **Current Users:** The number of users currently connected to your Parallels RAS Farm.
- **Maximum allowed concurrent users:** Prepaid subscription and NFR licenses only. The maximum number of concurrent users that your license allows. For example, if you own a prepaid subscription and need more concurrent connections, you need to upgrade your subscription.
- **Parallels Account user email:** Parallels business account email that was used to activate the Farm.
- **Parallels Account user name:** Parallels business account user name.
- **Parallels Account company:** Parallels business account company name.

Please note that you can also see these values (and more) in your Parallels Account. For more information, please read the **Parallels RAS Licensing Guide** and **Parallels RAS SPLA Guide**, which are available on the Parallels website.

The **View Active Users** button opens a dialog where you can view currently active users and license usage. Use the toolbar buttons to refresh the list and to copy the information to the clipboard.

The **Manage license** button allows you to switch to a different Parallels account and to activate Parallels RAS using a different license key. Click the button to open the **Sign in to Parallels My Account** dialog. Use the dialog to sign in using an existing account or click **Register** to create a new account. If you are creating a new account, you'll also have to register a Parallels RAS license key in it and activate your Parallels RAS Farm using that key (see below).

To activate Parallels RAS using a different license key:

1. In the **Sign In to Parallels My Account** dialog, type the email address and password you used to register your account and click **Sign In**. You'll see the **Activate Product** dialog.
2. Select the **Activate using license key** option and enter the key in the field provided. You can click the button next to the field to see the list of subscriptions and/or permanent license keys you have registered with Parallels My Account. If the list is empty, it means that you don't have a subscription yet and need to purchase one first.
3. To purchase a subscription online, click the **Purchase a license** link.
4. After entering a license key, click **Activate**. You should see the confirmation message that your Parallels RAS was activated successfully.

Configure HTTP proxy settings

If you use an HTTP proxy server on your network, you need to configure it in the RAS Console. The proxy server settings will be used during Parallels RAS license updates and by other features that communicate with the Parallels cloud.

To configure a proxy server:

1. In the RAS Console, navigate to **Administration > Settings**.
2. In the **HTTP Proxy settings** section, click the **Configure Proxy** button.
3. In the dialog that opens, select one of the following options:
 - **No Proxy server** — if you don't use a proxy server.
 - **Manual HTTP proxy configuration** — select this option to specify the settings manually. The **Detect Settings** button will attempt to detect the proxy settings automatically.

The **Proxy requires authentication** option allows you to specify or omit credentials for the proxy server. If your proxy server uses an IP address to authenticate clients, you can omit the credentials. Otherwise, select this option and specify a user name and password.

4. Click **OK** to save the settings.

System event notifications

You can configure system event notifications on the **Farm > Site > Settings > Notifications** tab. Notifications are used to alert the administrator about system events via email. When you configure notifications, the settings apply to all servers in the Farm.

To configure notifications, you first need to configure notification handlers where you can specify threshold values (where available) and whether an administrator should be notified via email. You can also configure notification scripts, which will be automatically executed when an event occurs.

Configuring notification handlers

To configure notification handlers:

1. In the RAS Console, navigate to **Farm > Site > Settings**.
2. Select the **Notifications** tab.
3. Click **Tasks > New** (or click the plus-sign icon) and choose an event for which to create a handler. For the list of events and their descriptions, please see the **System Events** subsection below.
4. A dialog opens where you can specify the event handler setting.

On the **General** tab, specify the following options:

- The threshold value (a number or percentage). Not available for some events (such as Licensing, Agent, and some other events).
- The direction (whether the event should trigger when the value rises above or drops below the specified value). Not available for some event (same as above).
- Whether to notify the administrator via email.

- Additional emails (separated by commas or semi-colons) to which to send event messages.
- Whether to execute a script when the event triggers. Here you need to select the **Execute a notification script** option and then choose a script from the drop-down list. Before you can use this option, you need to create one or more scripts as described in [Configuring notification scripts](#).

On the **Criteria** tab, specify the following:

- **Type:** Select the type of objects that trigger notifications.
- **All servers in site:** Select this option to include all available servers.
- **Available:** Select objects that trigger notifications.

On the **Settings** tab, specify the following:

- **Use default settings:** Select this option to use default settings. To edit defaults, click the **Edit Defaults** link. To use custom settings, clear this option and specify the options as described below.
- **Notification handler grace period:** Specify a time period (in minutes) to wait from the event occurrence until the notification is triggered. Some events may trigger but last for a very short period of time. For example, a CPU usage can sharply jump above the specified threshold but quickly return to normal. For such events, it would probably make sense not to trigger the notification right away. This option allows you to specify the delay.
- **Notification interval:** Specify the minimum time interval (in minutes) between the last and the next notification. Allows to prevent multiple notifications to be emailed to administrators in rapid succession (i.e. prevents spamming).
- **Send one notification and suspend further notifications until recovered:** When this setting is enabled, a notification will be raised only once, and after that it will be suspended until the values monitored by the notification have recovered. For example, if the CPU usage is above the threshold for the whole day, instead of executing the notification handler multiple times, RAS would execute it only once.

5. When done, click **OK** to save the notification handler.

Please note that the mailbox should be configured in the RAS Console for the outgoing email functionality to work. This mailbox is usually set up when you run the RAS Console for the first time and then use the **Start** category to set up your

RAS environment. You can also set up a mailbox as described in [Configuring SMTP server connection for event notifications](#).

To enable or disable an event handler, select or clear the checkbox in the first column, or right-click an event and choose **Enable** or **Disable**. To modify a handler, right-click it and choose **Properties**. To delete a handler right-click and choose **Delete**.

System events

You can create notification handlers for the following system events:

- **CPU utilization.** Triggers when CPU utilization rises above or drops below a specified value.
- **Memory utilization.** Triggers when memory utilization rises above or drops below a specified value.
- **Number of RDSH sessions.** Triggers when the number of active sessions rises above or drops below a specified value.
- **Number of disconnected RDSH sessions.** Triggers when the number of disconnected sessions rises above or drops below a specified value.
- **RDSH session utilization.** Triggers when the number of RDSH sessions rises above or drops below a specified percentage of the maximum number of sessions.
- **RDSH disconnected sessions utilization.** Triggers when the number of RDSH disconnected sessions rises above or drops below a specified percentage of the maximum number of sessions.
- **AVD session utilization.** Triggers when the number of AVD sessions rises above or drops below a specified percentage of the maximum number of sessions.
- **AVD disconnected sessions utilization.** Triggers when the number of disconnected AVD sessions rises above or drops below a specified percentage of the maximum number of sessions.
- **Number of Secure Gateway tunneled sessions.** Triggers when the number of gateway tunneled sessions rises above or drops below a specified value.
- **Number of failed Secure Gateway tunneled sessions.** Triggers when the number of failed gateway tunneled sessions rises above a certain value within an hour.
- **RAS Agents events.** Triggers when an agent event occurs (e.g. agent disconnects or reconnects).
- **Licensing events.** Triggers when a licensing event occurs. One notable event here is the license usage reaching a predefined threshold. Specifically,

when the license usage reaches 90% of all available licenses, you will receive an email, so you have time to decide whether you have enough licenses or need to add more. Other events include license activation/deactivation, license expiration, grace period starting/ending, license information changes, problem communicating with the licensing server.

- **Authentication server events.** Triggers when a connection issues occurs with an authentication server.
- **Published items events.** Triggers when a published item event occurs (e.g. the concurrent instance limit for an application is reached).
- **Template events.** Triggers when a VDI event occurs (e.g. a template is not found).
- **Tenant events.** Triggers when a Tenant event occurs. For more info, see [RAS Multi-Tenant Architecture > Configuring notifications](#).

Please also see the **Notification Types** table in the [Configuring notification scripts](#) section.

Configuring notification scripts

To configure notification scripts:

1. On the **Notifications** tab, click **Tasks > New** (or click the plus-sign icon) in the **Notifications scripts** section.
2. In the dialog that open, specify the following options:
 - **Script name:** Enter a friendly name for the script.
 - **Command:** The command to execute.
 - **Arguments:** Command line arguments to pass to the command. An argument can be one of the predefined variables, which Parallels RAS will automatically replace with an actual value. See the **Command Line Variables** table below (the ID column contains the values that can be used here).
 - **Initial directory:** The full path to the current directory for the process. The string can also specify a UNC path.
 - **User name, Password:** These are optional fields that you can specify if you would like to execute the command under a specific user account.
3. When done, click **OK** to save the notification script item.

To modify a notification script, right-click it and choose **Properties**.

To delete a script, right-click and choose **Delete**. Please note that if a script is used by a notification handler, you will see a warning message. If you choose to delete it anyway, the script association will be removed from all notification handlers where it is used and all affected handlers will be automatically configured to send an email alert.

Command line variables

The following table lists command line variables that you can use as arguments when executing a script (see the **Arguments** option description above):

Variable
(FARM – NAME) </td> <td> Name of the RAS Farm which has raised the notification. </td> </tr> <tr> <td> (SITE-NAME)
(SERVER – ADDRESS) </td> <td> IP address or FQDN of the server which has raised the notification. It could be an RDSH server, the server hosting a RAS Connection Broker, RAS Secure Gateway, etc. </td> </tr> <tr> <td> (TRIGGER-ADDRESS)
(THRESHOLD – VALUE) </td> <td> The threshold value that is assigned to the notification handler. If a notification type doesn't support thresholds, the arguments should be replaced with an empty string. </td> </tr> <tr> <td> (NOTIFICATION-TIME)

Variable
(\$NOTIFICATION-TYPE)

Notification types

The following table lists supported notification types (the ID column represents values that are passed to the (\$NOTIFICATION-TYPE) command line variable):

Notification type	ID	Description
CPU utilization	1	This notification is triggered when CPU utilization rises above or decreases below a certain value.
Memory utilization	2	This notification is triggered when memory utilization rises above or reaches below a certain value.
Number of active session	3	This notification is triggered when the number of active sessions rises above or decreases below a certain value.
	4	

Notification type	ID	Description
Number of disconnected sessions		This notification is triggered when the number of disconnected sessions rises above or decreases below a certain value.
RAS Agent reconnect	5	Agent reconnected.
RAS Agent disconnect	6	Agent disconnected.
VDI template is missing	7	This notification is triggered when an a VDI event occurs (e.g. a template is not found).
Published application limit exceeded	8	This notification is triggered when a published item event occurs (e.g. an application's instance limit is exceeded).
Multi CB communication error	9	Multiple CB communication error.
Authentication provider not reachable	10	This notification is triggered when a connection issue occurs with an authentication server.
% of RDSH session out of the maximum specified value	11	This notification is triggered when the number of RDSH sessions rises above or decreases below the specified percentage of the maximum number of sessions.
Gateway is tunneling X number of sessions	12	This notification is triggered when the number of gateway tunnelled sessions rises above or decreases below a certain value.
% of RDSH disconnected session out of the maximum specified value	13	This notification is triggered when the number of RDSH disconnected sessions rises above or decreases below the specified percentage of the maximum number of sessions.
Connection Broker auto promotion	20	Connection Broker auto promotion
Connection Broker auto promotion failed	21	Connection Broker auto promotion failed.
Connection Broker auto promotion failback	22	Connection Broker auto promotion failback.
CA not available	30	This notification is triggered when a connection issue occurs Certificate Authority.
License site switched to failed over mode	50	Connection Broker failover mode.
License site is offline	51	Licensing site is offline.
License site reconnected	52	Licensing site is back online.

Notification type	ID	Description
IP of Licensing CB changed	53	Licensing Connection Broker IP change.
Hostname of Licensing CB changed	54	Licensing Connection Broker Hostname change.
IP of secondary CB changed	55	Non Licensing Connection Broker IP change.
Hostname of secondary CB changed	56	Non Licensing Connection Broker Hostname change.
Templates max guests reached	60	Template max guest limit reached.
Template max servers reached	61	Template maximum server limit reached
Templates cloning failed	62	Template cloning failed.
License activated	100	This notification is triggered when a licensing event occurs (e.g. a farm has been successfully activated).
License deactivated	101	License was deactivated.
License max usage	102	The maximum license usage has reached x%.
License about to expire	103	If license is about to expire, notify every day saying how many days left.
License expired	104	License expired.
License trial expired	105	Trial period expired
License grace period start	106	Grace period started.
License grace period end	107	Grace period ended.
License disabled	108	License was disabled.
License information changed	109	License information changed
License failed to communicate with server	110	Failed to communicate with licensing server.
License no file	111	Failed to load license file.
License invalid version	112	Invalid license file version.
License invalid signature	113	Invalid license signature.
License invalid license	114	System errors.
License invalid MAC address	115	Invalid MAC address (hardware change).
	116	Migration grace period started.

Notification type	ID	Description
Licensing unsigned grace period		
Tenant enrolled	200	This notification is triggered when an event related to any registered tenant occurs (e.g. a new tenant was added to Tenant Broker or tenant becomes unavailable).
Tenant status changed	201	Tenant status changed
Broker status change	202	Tenant Broker status changed
Tenant disenrolled	203	A Tenant has unjoined the broker.
Standard Farm tunnel session failed	220	Standard tunnel session failed.
Broker Farm tunnel session failed	221	Tenant Broker tunnel session failed.

Configuring SMTP server connection for event notifications

The **Mailbox** tab in the **Administration** category allows you to configure an SMTP server for outgoing emails. The SMTP server is required for the administrator to receive system event alerts (as described in the previous sections) and to send invitation emails to users.

To configure an SMTP server:

1. In the RAS Console, select the **Administration** category and then click the **Mailbox** tab.
2. In the **Mail Server** field, type your mail server FQDN or IP address.
3. In the **TLS / SSL** drop-down list, select whether to use it the protocol.*
4. Select the **SMTP server requires authentication** option if required and then type the SMTP server username and password in the fields provided.
5. In the **Sender information** section, type the sender email address (e.g. your email).
6. The **Test mailbox settings** section can be used to test your SMTP server configuration. Enter one or more email addresses separated by a semicolon. Click **Send Test Email** to test the settings.

* The correct setting depends on the remote mail server configuration and cannot be arbitrary chosen. Please follow this guide table:

Setting	Description
Do not use	
Use TLS / SSL	
Use TLS / SSL if available	
Use TLS 1.2 if available	
Use TLS 1.3 if available	

RAS session variables

When a remote user starts a published application or desktop, a set of session variables is created by Parallels RAS on the host server. The variables contain information about the client machine, which you can examine if needed. The variables are always updated, so on connect/reconnect they always contain the latest values.

The following RAS session variables are available:

Variable name	Description
TUX_REMOTECLIENT_PLATFORM	Name and version of the operating system running on the client machine. For example, "Windows 8.1 Enterprise Edition (WOW 64)", "iPhone OS 9.2.1", "Android 6.0", etc.
TUX_REMOTECLIENT_MAC	MAC address of the client machine.
TUX_REMOTECLIENT_IP	

Variable name	Description
	IP address of the client machine as seen by the client.
TUX_REMOTECLIENT_LANG	Language used by the GUI on the client machine: EN, FR, RU, DE, ES, IT, PT, NL, JP, CS (Chinese Simplified), CT (Chinese Traditional), KR (Korean). Note that on macOS, iOS, and Android devices, the language is reported as the one used in the OS but only if it's a supported language. If it's not supported, it will default to EN.
TUX_REMOTECLIENT_MACHINE	Client's computer name. For example, "Bob's iPad mini 1st generation", "BobPC", "Bob's iMac", etc.
TUX_REMOTECLIENT_LOGIN	The username (including domain) that was used to log in to Parallels RAS. For example, myuser@somedomain.
TUX_REMOTECLIENT_VERSION	Parallels Client version.
TUX_REMOTECLIENT_VENDOR	Device vendor name. For example, "Asus", "Apple", "Google", etc.
TUX_REMOTECLIENT_MODEL	Device model name. For example, "Nexus 5", "iPad2.6", etc.

You can view RAS session variables and their values using one of the following two methods:

- By examining the Windows registry on the host server.
- By executing the GetRASVariable.exe utility (provided by Parallels RAS).

Each method is described below.

Examine the registry

To see the variables, run `regedit` and navigate to `HKEY_CURRENT_USER\Software\Parallels\Shell\<Session-ID>`, where `<Session-ID>` is the ID of a session as displayed in the RAS Console (e.g. 2, 3, 4, etc.) The variables for a particular session are listed under the session ID node. On

user connect/reconnect they are updated to reflect the actual client configuration. The variables exist for the duration of a session and are removed from the registry once the session is terminated.

Please note that in addition to the variables listed in the table above you may see other (undocumented) variables under a session ID. Those are for internal Parallels RAS use only and should be ignored.

Using GetRASVariable.exe utility

The GetRASVariable.exe utility is located in the Parallels RAS installation folder (e.g. C:\Files (x86)). To obtain a value of a variable, execute the utility from the command line passing the variable name as parameter (see the table above). The utility will output the value to the screen.

The following example displays the value of the TUX_REMOTECLIENT_MACHINE variable:

```
GetRASVariable.exe TUX_REMOTECLIENT_MACHINE
```

Resolving z-order issues

If you encounter an issue related to z-order in applications published with Parallels RAS, you can try doing the following:

1. Open the Windows Registry on the computer where Parallels Client is installed.
2. Create a value named ZOrderMode with the type DWORD in the HKEY_CURRENT_USER\Software\Parallels\AppServerClient\<guid> key, where <guid> is the ID of the Farm.
3. Specify one of the following values:
 - 1: Simple window manager rendering, no window clipping support. Works best with rectangular window areas.
 - 2: Simple window-to-window copying with a fixed refresh rate, supports window clipping. Supports all window area shapes.
 - 3: Heuristic mode, which switches between modes 1 and 2.
 - 4: Improved heuristic mode, which can require more resources.

Maintenance and backup

Keeping Parallels RAS up to date

By default, Parallels RAS checks for updates each time the RAS Console is started. If you wish to change this behavior:

1. Select the **Administration** category and click the **Settings** tab.
2. Select or clear the **Check for updates when launching Parallels RAS Console** option according to your needs.
3. Select or clear the **Automatically update RAS Session Host Agents** option according to your needs.
4. To check for updates manually, click the **Check Now** button.

Backing up the Parallels RAS Farm configuration

To backup the Parallels RAS Farm configuration:

1. Select the **Administration** category and then click the **Settings** tab.
2. Click the **Export Settings** button.
3. You'll see a message box saying that all sites will be synchronized. Click **Yes** to continue with export or click **No** to abort it.
4. Specify the file name and target folder and click **Save**.

Note: The export procedure only exports the Parallels RAS Farm configuration data. Unrelated objects, such as downloaded OS, etc. are not included in the exported file.

To restore a Parallels RAS Farm configuration from a backup file, click the **Import Settings** button and select a backup file (the default file extension is `.dat2`). When you import a configuration from a file, your existing Farm configuration will be completely replaced with it.

You can also export and import a Parallels RAS Farm configuration from the command line. For complete instructions, please read on.

Importing and exporting Farm settings from the command line

Parallels RAS PowerShell allows you to perform the majority of Parallels RAS administration tasks from the command line.

This section contains information about using PowerShell to export and import Farm settings. To learn more about Parallels RAS PowerShell, please visit <https://www.parallels.com/products/ras/resources/> and download (or view online) the **Parallels RAS PowerShell Guide**.

One of the uses of exporting and importing Farm settings is running automated tests. Specific configurations can be created, exported, and then imported for specific test scenarios. You can also use this functionality with Windows task scheduler for regular backups of Farm settings.

Installing Parallels RAS PowerShell

RAS PowerShell is installed by default when you run the default Parallels RAS installation. If you haven't installed it (or to install it on a different computer), do the following:

1. Run the Parallels RAS installer.
2. Select **Custom** and then select the **Parallels RAS PowerShell** component.
3. Complete the wizard and install Parallels RAS PowerShell.

Using Parallels RAS PowerShell

The complete up-to-date information about Parallels RAS PowerShell can be found in the **Parallels RAS PowerShell Guide**. The guide includes the **Getting Started** chapter to help you quickly get started with Parallels RAS PowerShell, as well as the complete reference and code samples. Please visit <https://www.parallels.com/products/ras/resources/> to view or download the guide.

Use the instructions below to export and import Parallels RAS Farm settings.

To import the Parallels RAS PowerShell module, open the PowerShell console and execute the following command:

```
Import-Module PSAdmin
```

Create a Parallels RAS session (use the name or IP address of the server where you have Parallels RAS installed):

```
New-RASSession -Server "server.company.dom"
```

To export Farm settings, execute the following command (substitute the path and filename of the backup file with your own):

```
Invoke-RASExportSettings "C:\Backup\RAS-backup.dat2"
```

To import Farm settings:

```
Invoke-RASImportSettings "C:\Backup\RAS-backup.dat2"
```

Problem reporting and troubleshooting

If you are experiencing an issue with Parallels RAS, you can search for a solution right from the RAS Console. If you can't find a solution, you can send a support request to Parallels. This section describes how to accomplish these tasks.

Search for a solution

To search for a solution from the RAS Console:

1. In the console, click **Help** on the main menu and choose **Troubleshooting and request support**.
2. The **Troubleshooting** dialog opens.
3. In the **Select Category** drop-down list, select a category you are having a problem with. The area in the middle of the dialog will be populated with a list of existing KB articles related to that category.
4. Click an article of interest to read in a web browser.
5. You can also click **Knowledge Base Index** or **Forums** links to go to the Parallels knowledge base or Parallels forums.

Request support

If you can't find a solution for your problem using the options described above, you can send a support request to Parallels. When you do, the collected logging

information is retrieved and attached to the email, so that Parallels Support can analyze it. See [Logging](#) for more information.

Note: A support request creates a support ticket, which is then sent to Parallels Support. If you already have a request support ticket, you can send just the system report to Parallels without creating an additional (and identical) ticket. See the **Send a report** subsection below. Please note that if you don't have a valid RAS subscription or a support contract, the ticket will not be created. In order to receive support, you will need to purchase a subscription or support contract. `endhint %}`

Before you request support, please make sure that you have a mailbox setup in the RAS Console. If you haven't set up a mailbox yet, do it as follows:

1. In the RAS Console, navigate to **Administration > Mailbox**.
2. Enter your outgoing email server information, your email address, and the security/authentication information if needed.
3. You can send a test email by entering an email address in the field provided and clicking the **Send Test Email** button.

To send a support request to Parallels:

1. In the **Troubleshooting** dialog, click the **Send Support Request** button.
2. The **Contact Support** dialog opens.
3. Enter your full name and your company name.
4. Enter the subject. This will be used as a subject in the email that will be sent to Parallels Support.
5. In the **Enter your query** box, describe the issue the best you can.
6. Use the **Attachment** field to attach a file to the email. Click the [...] button to browse for a file. You can attach a picture or any other file that you think might help the Parallels Support to find a solution. Please note that the log files and the Parallels RAS settings are collected and attached to the email automatically, so you don't have to do it yourself.
7. In the drop-down list at the bottom of the dialog, select whether you want to send the email or save it (including the automatically collected information) as a zip file.
8. Depending on the action selected in the previous step, click **Send** to send the email or **Save** to save it as a zip file on your local drive or a network folder.

Send a report

If you already have a support request ticket, you can send just a system report to Parallels without creating a new ticket.

To send a report:

1. In the console, click **Help** on the main menu and choose **Upload System Report to Parallels**.
2. A dialog opens displaying the progress bar.
3. Once the system report data is collected and sent to Parallels, a message box is displayed containing the report number.
4. Click **OK** to finish.

Contacting the local support

If your root administrator added a URL to the local support portal, you can navigate to it using the **Request Support** option in the **Help** menu.

Logging

Parallels RAS components are monitored and logs are created containing relevant information. Logs are used by Parallels RAS support engineers to analyze possible issues with a Parallels RAS installation. As a Parallels RAS administrator you have the ability to set the log level for a specific component or multiple components. By default, the standard log level is used, which collects and saves only the essential information. A Parallels RAS support engineer can ask you to enable the extended or verbose log level when an additional information is required to analyze an issue.

To set the log level for a specific component/server, navigate to the page in the RAS Console where the components of that type are listed (e.g. RD Session hosts, VDI, Gateways, Connection Brokers), select a component and then click **Tasks** (or right-click) > **Troubleshooting** > **Logging** > **Configure**. This opens the **Set Log Level** dialog where you can choose a log level from the following:

- **Standard** — This is the standard log level that records only the most important events. Unless you are asked by Parallels RAS support to use one of the log levels described below, you should always use this one.

- **Extended** — This logging involves more information than the standard logging, but it slows down the system because of the additional information that it needs to collect.
- **Verbose** — Verbose logging involves even more information than the extended logging and can slow down your system significantly.

Please note that to avoid degraded performance, extended and verbose logging should only be enabled for a limited time period (enough to collect the necessary information for analysis). You can set this time period using **Reset to the standard level after** option. The default value is 12 hours. In specific cases, a Parallels support engineer will advise you whether this time period should be set to a different value. Once this time period is over, the log level will be reset back to standard.

To retrieve a ZIP archive containing the log files, click **Tasks** (or right-click) > **Troubleshooting** > **Logging** > **Retrieve** and then specify a location where you want the file to be saved. The **Clear** item in the same context menu clears all logs.

Note that you can also set the log level on the **Farm** > <Site> > **Settings** > **Global logging** tab, where you can see RAS components of all types in one list. For more information, see [Site settings](#).

Log rotation

Parallels RAS log rotation works as follows:

1. When the total size of all log files reaches a predefined size (200 MB by default), the logs are archived. This is done log by log by appending the current timestamp to the filename and starting a new empty log file.
2. A new ZIP file is created for each old log named %logname%_%DATE%.zip . (e.g. console_10.06.2018.zip, controller_10.06.2018.zip).
3. Renamed old logs are moved to the ZIP file. Parallels RAS keeps five ZIP files by default.
4. When the maximum number of archived files is exceeded, the oldest file is deleted.
5. This log rotation mechanism guarantees that the total log file size never exceeds $X * Y * Z$ MB, where X is the total size of all log files (200 MB by default), Y is the maximum number of ZIP files (5 by default) and Z is the number of RAS components.
6. The X and Y values from the example above are pre-configured in Windows registry on a computer hosting a given RAS component. The default values are the same for every RAS component. To modify the values, navigate to

HKEY_LOCAL_MACHINE > SOFTWARE > Wow6432Node > Parallels and set the LogMaxSize and LogMaxBackups values for a RAS component. Restart the services after modifying the values.

Suggest a feature

If you have an idea of a new feature for Parallels RAS, we would like to hear from you! To suggest a feature, in the RAS Console, click **Help** on the main menu and choose **Suggest a Feature**. This will take you to the **Parallels RAS Feature Suggestion** web page where you can communicate your ideas to us. Please note that you must be signed in using your Parallels account email address and password to post in the feature suggestion forum.

Parallels RAS Management Portal

This chapter gives you an overview of Parallels RAS Management Portal. For the complete information, please refer to **Parallels RAS Management Portal Guide**, which is available on the Parallels website at <https://www.parallels.com/products/ras/resources/>.

Overview

Parallels® RAS Management Portal is a modern web-based configuration and administration console designed for Parallels RAS administrators using a desktop/laptop computer or a mobile device to carry out configurations and day-to-day activities.

Parallels RAS Management Portal provides administrators with ability to:

- Centrally deploy, manage, and configure essential Parallels RAS components such as RD Session Hosts, Connection Brokers and Secure Gateways.
- Publish various resources from RD Session Hosts.
- Configure FSLogix Profile Container settings.
- Configure printing and scanning settings.
- Manage SSL certificates.
- Configure connection settings and MFA (Google Authenticator or other Time-based One-time Password (TOTP) apps such as Microsoft Authenticator).

- Monitor and manage user sessions.
- Manage administrative accounts and sessions
- Configure mailbox.
- Manage your license.
- Contact support and provide necessary system reports.

Note: More features and capabilities that are currently available in the desktop-based Parallels RAS Console will be included in Parallels RAS Management Portal in future releases until it becomes the main management tool for Parallels RAS. Management of Azure Virtual Desktop capabilities included in Parallels RAS Management Portal are experimental and expected to be released in upcoming versions.

Prerequisites

RAS Management Portal can run in any modern web browser supporting HTML5 except for Internet Explorer.

Make sure your Windows Server has the following updates installed (RAS Management Portal depends on them):

- Windows Server 2012 R2: KB2999226

Newer versions of Windows Server do not require any specific updates.

The web service listens to web requests on the following ports by default:

- HTTPS: 20443
- HTTP: 20080

Installation

To enable RAS Management Portal in a RAS Farm, you need to install the RAS Web Administration Service component. The component is installed automatically when you do a clean Parallels RAS install using the “Typical” installation option. You can also install the component using the “Custom” installation option and choosing the “RAS Web Administration Service” as the component to install. For example, if you want to install RAS Management Portal on a dedicated machine, you should use the “Custom” installation option and select “RAS Web Administration Service” as a component to install.

After the RAS Web Administration Service is installed, you need to configure it. Specifically, you need to specify a RAS Farm that the RAS Management Portal will be used to manage, and you also need to configure a number of other parameters. For complete instructions, please see [Configure RAS Web Administration Service](#).

Log in to RAS Management Portal

To open RAS Management Portal on the machine where you've installed the RAS Web Administration Service, navigate to **Apps > Parallels** and click **Parallels RAS Management Portal**.

To log in to RAS Management Portal from a remote computer, enter the following URL in a web browser:

```
https://<server-address>:20443
```

The <server-address> is the FQDN or IP address of the server where the RAS Web Administration Service is installed. By default, port 20443 is used for HTTPS connections. You can change the port number if needed as described in [Configure RAS Web Administration Service](#).

On the **Welcome** page, enter your RAS administrator username and password and click **Sign in**.

Configure RAS Web Administration Service


Before you begin, you may need to configure the RAS Web Administration Service as described below:

1. In RAS Management Portal, click the "User" icon in the upper right-hand corner and choose **Configure Management Portal**.
2. You will be asked to sign in again. Note that the RAS Web Administration Service must be running on the local server for this sign in to work. This is necessary to prevent users from remote servers to enter the RAS Web Administration Service configuration pages.
3. Enter the username and password of a member of local administrators or domain administrators and click **Sign in**.

4. The **RAS Management Portal Configuration** page opens.
5. In the **RAS Farm Address** field, specify the RAS Farm address that this RAS Management Portal will manage. This is the RAS Connection Broker address installed in the Farm.
6. In the **Advanced Settings** section, specify the following:
 - **Certificate:** A certificate to use for this connection. Click **Upload** to select a certificate.
 - **Certificate Password:** The certificate password.
 - **Port:** The port number on which RAS Management Portal listens for connections. The default port is 20443. This port number is chosen not to conflict with RAS Secure Gateway ports. You can change it to 443 (if possible), in which case the port number doesn't need to be included in the connection URL. You can also change it to any custom port. For example, the default "URL": "https://*:20443" can be changed to "URL": "http://*:20080".
 - **Admin Session Timeout:** The timeout after which the admin session will be disconnected.
 - **Polling Interval:** The interval at which RAS Management Portal will update the information displayed in it. You can increase this number up to 30 seconds if you have a large number of admins working at the same time and/or if you have a large number of hosts, sessions, etc.
7. Click **Save** when done.

RAS Management Portal user interface

All navigations in the RAS Management Portal start from the sidebar on the left, which lists management categories. The **Site** category is selected by default.



Categories

The following table lists all available categories that can be managed in the RAS Management Portal. The Root Administrator can see and manage all categories. Administrators of other types (Power, Custom) may need permissions to see a particular category.

Category	Description
Site	Displays the current Site overview.
Infrastructure	RAS infrastructure management, including RD Sessions Hosts, VDI, Gateways, Connection Broker, etc.
Sessions	Session management.
Publishing	Publishing and published resources management.
Monitoring	RAS Performance Monitor.
Site Settings	Connection, authentication, FSLogix, Universal printing and scanning.
Help and Support	Help and support.
Farm Settings	Displayed at the bottom of the sidebar on the left, this category manages global Farm settings, such as Administrators, Mailbox, Licensing.

Each category is described in detail later in this guide.

Admin permissions

Some categories and actions in the RAS Management Portal may not be viewed or allowed depending on the Admin permissions configured in the desktop RAS Console. For the information about how to configure administrator permissions,

please refer to the **Parallels RAS Administrator's Guide**. In the guide, look for the **Administrator Account Permissions** topic. The guide is available on the Parallels website at <https://www.parallels.com/products/ras/resources/>.

Subcategories

Some categories have subcategories (namely **Infrastructure** and **Site Settings**). When you selected a category, the right side of the RAS Management Portal may include one or more additional panes where you can select a subcategory.

Navigation bar

Some components have their settings and information grouped by functionality (e.g. General, Properties, Sessions, etc.). When you view component properties, a navigation bar is displayed in the middle allowing you to browse these settings. When you select an item in the navigation bar, the settings are displayed in the right pane.

Breadcrumbs

As you select categories, subcategories, individual items, a breadcrumb trail is displayed at the top of the page to show where you are. To take one or more steps back, click a link in the trail.

Page header items

The page header includes the following items:

- The Farm and the current Site names. If you have more than one Site, you can select one from the drop-down list. The RAS Management Portal will switch to that site allowing you to manage the Site components.
- The “User” icon is a drop-down list with the following items: Current user name (e.g. **Administrator**); **About** (opens the About dialog); **Give feedback** (takes you to a web page where you can give feedback to Parallels); [Configure Management Portal](#), **Logout** (logs you out).
- **Apply All Changes**: This button applies changes that you’ve made in the RAS Management Portal to Farm components. When you create or modify components and objects, the changes are not applied to Farm components automatically and don’t have any effect on the Site or Farm. When you click the **Apply All Changes** button, the changes are applied across the Farm or Site. Note that you shouldn’t always click this button every time you make a

change. If you are working on a task that requires multiple changes in different areas, complete all of them and then click the **Apply All Changes** button, so all changes are applied together.

Editing

When you open a view where you can modify some settings, the view is normally read-only. To enable editing, click the **Edit** button in the upper right-hand corner. The button name changes to **Save**. When done editing, click **Save**. To discard the changes, click **Cancel**.

Please note that an object that is opened for editing by an admin cannot be edited by another admin at the same time. If you try to enable editing for such an object, you will get an error with the name of the admin who has the object locked.

Edit toolbar

Some views (specifically lists) have a toolbar in the upper right-hand corner from which you can execute actions. To see a toolbar item name, hover over it with the mouse. The standard items (icons) on the toolbar are the following:

- **Show filter**: Specify a filter to show only the entries that match it.
- **Select columns**: Select table columns to display or hide.
- **Add**: Add a new entry. For example, add a new Gateway or RD Session Host, etc.
- **Refresh**: Refresh the view.
- **Ellipsis**: The ellipsis menu may have different items in different types of views. Some items have a corresponding toolbar items (e.g. **Add**, **Refresh**).

Other items may be present depending on the view you are in. For example, **Show running processes** and **Show sessions**.

Wizards

When you add a component to a Farm, a wizard usually opens which takes you through a series of pages where you specify component settings and options. A wizard has the usual **Next** and **Back** navigation buttons, and the **Cancel** button that closes the wizard and cancels the operation.

Modal dialogs

Clicking some menu and navigation bar items brings up a modal dialog. These are usually items that require you to confirm an action or enter additional information.

Object properties views

All objects (components) in the RAS Management Portal have properties. To view these properties, you select a category and a subcategory and click the object name in the list. This opens a view where object properties are displayed with its own navigation bar from which you can configure the object, perform actions. and view additional information.

Parallels RAS APIs

Parallels RAS comes with APIs to help you develop custom applications that integrate with it. This includes RAS PowerShell API and RAS REST API.

In addition, the RAS Web Client API and Parallels Client URL scheme allow you to integrate with Parallels Client for Windows/macOS/Linux/iOS/Android and the Web Client.

RAS PowerShell API

RAS PowerShell API is intended for RAS administrators who would like to automate their RAS administration. The API includes commands to perform most of the RAS management tasks.

Parallels RAS requirements

The Parallels RAS PowerShell API version must match the version of the RAS Connection Broker with which it communicates. Since the two components can be installed separately, you need to make sure that their versions match.

Microsoft Windows component requirements

The following components must be installed on the computer where you'll be executing Parallels RAS PowerShell cmdlets:

- Windows PowerShell 3.0 or higher

- Microsoft .NET Framework 4.5.2 or higher

Installation

To install Parallels RAS PowerShell, run the standard Parallels RAS installer, choose **Custom** installation, and then select to install the **Parallels RAS PowerShell** component. Follow the onscreen instructions to install the component.

RAS PowerShell API versions

The RAS PowerShell API has undergone changes in Parallels RAS 18 as follows:

- The RAS PowerShell module name has changed from PSAdmin to RASAdmin.
- Most of the commands now have the “RAS” prefix, such as RASGW or RASApply.
- API versions: Version 2.0 (latest) and version 1.0 are supported for backward compatibility.

Note that the API version 1.0 is still available in the current RAS PowerShell module. If you have existing scripts that use the older module and command names, you can use them with minimal changes. To do that, you need to load the API version 1.0 when you import the RAS PowerShell module. See below for more information about API versions.

Version 2.0

This version is the one loaded by default by the system or if the `-RequiredVersion` parameter is not specified when importing the module. See **RAS PowerShell API concepts** for examples.

Version 1.0

This version keeps backward compatible with the old PSAdmin module to allow administrators to use existing scripts with minor modification. This version includes:

- Cmdlet aliases
- Aliased parameters
- Returns old and new properties

RAS PowerShell API concepts

To quickly get started with RAS PowerShell, do the following:

1. Open the Windows PowerShell console.
2. Import the Parallels RAS PowerShell module using one of the following commands:
 - `Import-Module RASAdmin` – Loads the current API (version 2.0).
 - `Import-Module RASAdmin -RequiredVersion 1.0` – Loads the API version 1.0.

3. Create a Parallels RAS session by executing the `New-RASSession` cmdlet (see example below). Substitute the server name (in quotes) with the name or IP address of your Parallels RAS Licensing Server. Type your RAS administrator username and password when prompted:

```
New-RASSession -Server "server.company.dom"
```

4. Execute the following cmdlet to see the list of cmdlets included in the Parallels RAS PowerShell module:

```
Get-Command -Module RASAdmin
```

5. Execute other cmdlets. For example, try executing the `Get-RASW` cmdlet to retrieve information about RAS Secure Gateway(s). The example below returns information about all RAS Secure Gateways available in the RAS Licensing Server Site:

```
Get-RASGW
```

6. To see help for a cmdlet, execute `Get-Help` passing a cmdlet name:

```
Get-Help Get-RASGW
```

7. To apply changes you've made to the Farm configuration, use the `Invoke-RASApply` cmdlet (this performs the same action as the **Apply** button in the RAS Console):

```
Invoke-RASApply
```

8. To activate a Parallels RAS license, use the `Invoke-RASLicenseActivate` cmdlet:

Invoke-RASLicenseActivate

When executing the cmdlet above, you'll be prompted to enter your Parallels account email address and password. You can include an optional -Key parameter and specify a Parallels RAS license key. If omitted (as in the example above), Parallels RAS is activated as a trial.

Parallels RAS PowerShell Guide

To view and download the new **Parallels RAS PowerShell Guide** version 2.0, visit Parallels website at <https://www.parallels.com/products/ras/resources>.

RAS REST API

This section gives you an introduction to the RAS REST API. Read it to learn about system requirements, installation, configuration, and basic usage.

Installation

To enable RAS REST API in a RAS Farm, you need to install the RAS Web Administration Service. It can be installed on the RAS Connection Broker server or any other server. If you install the service on a separate server, you will need to change its configuration (after the installation) to point to RAS Connection Broker. By default, the configuration points to "localhost".

Note: If you've already configured and are using Parallels RAS Management Portal, you may skip this step because you should already have the RAS Web Administration Service installed.

To install RAS Web Administration Service:

1. Run the Parallels RAS installer on the RAS Connection Broker or any other server.
2. On the **Select Installation Type** page, select **Custom**.
3. On the next page, select to install the **Parallels RAS Web Administration Service** component.
4. Click **Next** and follow the onscreen instructions.

Configure RAS Web Administration Service

If the RAS Web Administration Service was installed on a separate server, you need to modify its configuration and specify the RAS Connection Broker server address. You may also want to change the port number and certificate information in the same configuration file. The configuration of the RAS Web Administration Service is saved as a JSON file at the following location:

C:\Files (x86).json

When modifying the service configuration, please note the following:

- In the configuration JSON file, the RAS Connection Broker address is specified using the “LicenseServer” parameter.
- The default HTTPS port number is set to 20443. This number is chosen not to conflict with RAS Secure Gateway ports. You can change it to 443 (if this does not conflict), so when opening the portal, you don’t have to include the port in the URL.

Permissions

To access any of the RAS REST resources, the user executing a request must have sufficient rights to access a particular resource. These are basically the same rights a RAS administrator has in the Parallels RAS Console. For example, a root administrator can access any of the RAS REST resources. On the other hand, a power administrator who doesn’t have rights to modify Site settings (as an example) will not be able to access a corresponding REST resource. Similarly, a custom administrator who, for instance, only has rights to view and modify RD Session Hosts will be able to access just that particular REST resource and no other.

Getting started

Applications communicate with Parallels RAS by sending HTTP or HTTPS requests. Parallels RAS answers with a JSON file in a response to every HTTP request.

All HTTP requests that you will use to retrieve and manage Parallels RAS resources have the following base structure:

`https://<API-host>/api/<URI>`

The parameters in the above URL are:

- <API-host> is the IP address or FQDN of the server on which the RAS Web Administration Service is installed.
- <URI> is a path to a REST resource that you would like to work with.

Logging in and sending requests

This section contains an example of RAS REST API usage that can help you quickly get started. The example demonstrates how to:

1. Login to Parallels RAS and obtain a session token.
2. Retrieve the information about all available RD Session Hosts.
3. Retrieve the information about a specific RD Session Host.
4. Modify RD Session Host properties.

Log in to Parallels RAS and obtain a session token

Before you can access any of the resources, you need to log in to Parallels RAS using administrator credentials and obtain a session token. This is accomplished by sending the following request:

`POST https://<API-host>/api/session/logon`

Request headers: The logon request must contain just the Content-Type request header. Subsequent requests must additionally contain the `auth_token` header, as you'll see in the examples that follow this one.

Content-Type: `application/json; api-version=1.0`

Request body: The request body must contain the RAS administrator user name and password.

```
{
  "username": "USER",
  "password": "PASSWORD"
}
```

Response: After sending the logon request, you will receive a reply containing the session token, which you will use in all subsequent requests:

```
{
  "authToken": "Lj+KddoJkANhzvbDRvB=K=DFCroRjXJHeeWGbGIIRKaz-EXplbmhVWvWTiDVqtOq"
}
```

Retrieve information about RD Session Hosts

Now that we have the session token, we can send requests to access various resources. In this example we'll first obtain the information about all available RD Session Hosts. In the example that follows, we'll obtain the information about a specific RD Session Host.

To retrieve the RD Session Host info, send the following request:

```
GET https://<API-host>/api/RDS/Host
```

Request headers: This time the `auth_token` request header must also be included and must contain the session token that we've obtained earlier.

Content-Type: `application/json; api-version=1.0`

auth_token: `Lj+KddoJkANhzvbDRvB=K=DFCroRjXJHeeWGbGIIRKaz-EXplbmhVWvWTiDVqtOq`

Response: The response will look similar to the following (with multiple RD Session Hosts in the Farm each block of the result set will contain the information about an individual server).

```
[
  {
    "directAddress": "IP_ADDR",
    "rasTemplateId": 0,
    "inheritDefaultAgentSettings": true,
    "inheritDefaultPrinterSettings": true,
    "inheritDefaultUPDSettings": true,
    "inheritDefaultDesktopAccessSettings": true,
    "port": 3389,

    "restrictDesktopAccess": false,
    "restrictedUsers": [],
    "server": "IP_ADDR",
    "enabled": true,
    "description": "",
```

```
        "siteId": 1,  
        "id": 2  
    }  
]
```

Retrieve information about a specific RD Session Host

To retrieve the information about a specific server, we'll use the same request as above but will add the server ID at the end:

```
GET https://<API-host>/api/RDS/Host/2/
```

The response will also be similar to the example above and will contain the information just for the specified server.

Modify RD Session Host properties

In this example we'll modify a property of the RD Session Host that we retrieved earlier. For simplicity let's modify the "description" field.

The request to modify properties of an RD Session Host has the following syntax:

```
PUT https://<API-host>/api/RDS/Host/2/
```

Note the 2 at the end of the request, which specifies the ID of the RD Session Host that we want to modify.

Request headers:

- **Content-Type:** application/json; api-version=1.0
- **auth_token:** Lj+KddoJkANhzvbDRvB=K=DFC CroRjXJHeeWGbGIIRKaz-EXplbmhVWvWTiDVqtOq

Request body:

```
{  
    "description": "description was updated!"  
}
```

Response: If the PUT request succeeds, you will get an empty response with code 204: No Content. To verify that the "description" field was in fact modified, let's use the same GET request that we used earlier: `GET https://<API-host>/api/RDS/Host`

As we can see, the result now contains the updated “description” field:

```
[
  {
    "directAddress": "IP_ADDR",
    "rasTemplateId": 0,
    "inheritDefaultAgentSettings": true,

    "server": "IP_ADDR",
    "enabled": true,
    "description": "description was updated!",
    "siteId": 1,
    "id": 2
  }
]
```

More information

Parallels RAS REST API comes with the **Parallels RAS REST API Guide**. The guide contains more examples and the complete resource and schema reference. To view and download the guide, visit <https://www.parallels.com/products/ras/resources/>.

RAS Web Client API and Parallels Client URL scheme

RAS Web Client API and Parallels Client URL scheme allow you to integrate with Parallels clients.

Using the RAS Web Client API or the URL scheme, you can implement an in-house solution, such as an application hub or web portal, for authenticating users and launching remote applications, desktops, and other published resources. Such an implementation is possible by integrating a custom solution with Parallels Clients, including Parallels Clients for supported platforms (Windows, macOS, Linux, iOS, Android) and RAS Web Client.

The following is a quick summary of the API and the URL scheme:

- **RAS Web Client API** — provides connection, user authentication, and resource launching methods called from a web browser via Web Client.
- **Parallels Client URL scheme** — a custom URL scheme that allows you to perform actions in a Parallels Client installed on a user device. Actions include configuring a connection, authenticating a user, and launching published resources.

RAS Web Client API and Parallels Client URL scheme are described in detail in the **Integrating with Parallels Clients** guide, which is available for download on the Parallels website at the following location: <https://www.parallels.com/products/ras/resources/>.

Appendix

Microsoft license requirements in Parallels RAS

This section is intended to provide guidance on typical Microsoft license requirements in a Parallels RAS environment, and should not be used as an exclusive list. For further information, please refer to your Microsoft licensing partner.

Microsoft license requirements include:

General

- Licensing all Windows Server and Desktop Operating Systems (OS's) is required and is not included within Parallels RAS.
- The Windows Server Operating System (OS) to be accessed must be covered by Microsoft Windows Server Client Access Licenses (CALs). A Windows server hosted within Microsoft Azure can include access rights in the per-minute charge for the virtual machine.
- The appropriate Microsoft access licenses are required, no matter how access is gained to the remote operating systems, which applies to both virtual and physical OS.

RD Session Hosts

If Windows Server is accessed remotely (for non-administrative work), then you need a Remote Desktop Service (RDS) Client access license (CAL):

- RDS CALs are required for users or devices that want to access a Windows Server remotely. The following types of RDS CAL are available:
 - RDS Device CAL: Permits one device (used by any user) to use Remote Desktop Services functionality on any of your servers.
 - RDS User CAL: Permits one user (using any device) to use Remote Desktop Services functionality on any of your servers.
 - RDS External Connector: Permits multiple external users to access a single Remote Desktop server. If you have multiple servers, you need multiple external connectors in addition to any required Windows Server External Connectors.

You may choose to combine RDS Device CALs and RDS User CALs simultaneously with the server software. Regular User or Device CALs are required in addition to the RDS User or RDS Device CALs.

- RDS SAL is a service that provides a Microsoft Remote Desktop Service Subscriber Access License (called an “RDS SAL”) on Virtual Machines created in Compute Resource. This makes it possible for three or more users to connect to a remote desktop (RD Session Host) for a specific Virtual Machine in Compute Resource (for SPLA partners).

Read more:

License your RDS deployment with client access licenses (CALs): <https://docs.microsoft.com/en-us/windows-server/remote/remote-desktop-services/rds-client-access-license>.

VDI

In the case of remotely accessing a single-user desktop OS, or Virtual Desktop Infrastructure (VDI), Virtual Desktop Access (VDA) licenses are required. Microsoft licenses Windows by access device:

- Virtual desktop access rights are a benefit of Windows Client Software Assurance (SA). Customers who intend to use PCs covered under SA have access to their remote VDI desktops at no additional charge.

- Customers who want to use devices that do not qualify for Windows Client SA, such as thin clients, will need to license those devices with a Virtual Desktop Access (VDA) to access a Windows VDI desktop. Windows VDA is also applicable to third-party devices, such as contractor or employee-owned PCs.
- Any user with one of the following Microsoft subscriptions: Microsoft 365 F3, E3, or E5 will not need a separate VDA license when accessing Windows client OS (Windows 11) on their own or from outsourced infrastructure. This does not cover access to what Microsoft terms “Listed Providers”, which are specifically excluded.

Read more:

- Windows 11 licensing portal: <https://www.microsoft.com/en-us/Licensing/product-licensing/windows>.
- Configure VDA for Windows Subscription Activation: <https://learn.microsoft.com/en-us/windows/deployment/vda-subscription-activation>.
- Microsoft Listed Providers: <https://www.microsoft.com/licensing/docs/view/Listed-Providers>.

Remote PC

Parallels RAS Remote PC allows access to both physical and standalone virtual Windows PCs. The Microsoft licensing position is similar for both scenarios. In the case of remotely accessing a single-user desktop OS, Virtual Desktop Access (VDA) licenses are required. Microsoft licenses Windows by access device:

- Virtual desktop access rights are a benefit of Windows Client Software Assurance (SA). Customers who intend to use PCs covered under SA have remote access to their physical desktops at no additional charge.
- Customers who want to use devices that do not qualify for Windows Client SA, such as thin clients, will need to license those devices with a Virtual Desktop Access (VDA) to access a Windows desktop remotely. Windows VDA is also applicable to third-party devices, such as contractor or employee-owned PCs.
- Any user with one of the following Microsoft subscriptions: Microsoft 365 F3, E3, or E5 will not need a separate VDA license when accessing Windows client OS (Windows 11) on their own or from outsourced infrastructure. This does not cover access to what Microsoft terms “Listed Providers”, which are specifically excluded.

Read more:

- Windows 11 licensing portal: <https://www.microsoft.com/en-us/Licensing/product-licensing/windows>.
- Configure VDA for Windows Subscription Activation: <https://learn.microsoft.com/en-us/windows/deployment/vda-subscription-activation>.
- Microsoft Listed Providers: <https://www.microsoft.com/licensing/docs/view/Listed-Providers>.

Azure Virtual Desktop

- Access to Windows 11 Enterprise multi-session and Windows 11 Enterprise desktops and apps is provided at no additional cost via AVD (this is licensing only and excludes compute, storage, and networking costs) if you have one of the following per-user licenses:
 - Microsoft 365 E3/E5/A3/A5/F3/Business Premium/Student Use Benefits
 - Windows Enterprise E3/E5
 - Windows Education A3/A5
 - Windows VDA per user
- Access to desktops powered by supported versions of Windows Server Remote Desktop Services is provided at no additional cost (excluding compute, storage, and networking costs) if you have a per-user or per-device RDS CAL license with active Software Assurance (SA).

Read more:

- Azure Virtual Desktop pricing overview: <https://azure.microsoft.com/en-us/pricing/details/virtual-desktop/>.
- Prerequisites for Azure Virtual Desktop: <https://learn.microsoft.com/en-us/azure/virtual-desktop/prerequisites?tabs=portal>.

FSLogix

You are eligible to access FSLogix Profile Container, Office 365 Container, Application Masking, and Java Redirection tools if you have one of the following licenses:

- Microsoft 365 E3/E5
- Microsoft 365 A3/A5/ Student Use Benefits
- Microsoft 365 F1/F3

- Microsoft 365 Business
- Windows Enterprise E3/E5
- Windows Education A3/A5
- Windows VDA per user
- Remote Desktop Services (RDS) Client Access License (CAL)
- Remote Desktop Services (RDS) Subscriber Access License (SAL)

FSLogix solutions may be used in any public or private data center if a user is properly licensed.

Read more:

- FSLogix Overview: <https://docs.microsoft.com/en-us/fslogix/overview>.
- Prerequisites for FSLogix: <https://learn.microsoft.com/en-us/fslogix/overview-prerequisites>.

Microsoft SQL Server

SQL Server is required if using Parallels RAS Reporting. SQL Server installation may be based on:

- SQL Express which is free but has a database size limit of 10 GB.
- SQL Server commercial edition Standard or Enterprise, using Core-based licenses or Server + CAL-based licenses.

Read more:

- SQL Server 2022 licensing guide: https://download.microsoft.com/download/9/3/d/93d32de6-f268-45ed-ba25-2f9a6756b6af/SQL_Server_2022_Licensing_guide.pdf.
- SQL Licensing Summary: <https://www.microsoft.com/en-us/licensing/product-licensing/sql-server>.

App-V

App-V is not separately licensed. It is included in other license agreements such as Microsoft Volume Licensing, Windows Software Assurance, Microsoft Remote Desktop Services (RDS) CAL, and as part of a wider Microsoft licensing agreement. For instance, with an RDS CAL (either per-user or per-device), the App-V client may be used on an RD Session Host to deliver App-V applications.

To license App-V correctly, it is recommended that you engage with a Microsoft Partner (solution provider) knowledgeable on Microsoft Volume Licensing (list of Microsoft Partners: <https://pinpoint.microsoft.com/en-us/search?type=companies&competency=100010>).


Note: Application Virtualization will reach end-of-life in April 2026. It is recommended to consider a transition to an alternative delivery method, such as MSIX app attach, before this date.

Other References

For a detailed list of Microsoft Licensing Product Terms, please see <https://www.microsoft.com/licensing/docs>.

Port reference

The following diagram illustrates communication ports used in Parallels RAS.



The above diagram include SAML SSO components such as RAS Enrollment Server, however it does not include Tenant Broker.

Tip: If you are reading the PDF version of this guide, click the following link to view the full-sized diagram in a web browser: https://download.parallels.com/ras/v19/docs/en_US/Parallels-RAS-19-Administrators-Guide/index.htm#47092.

Parallels Client

Source	Destination	Protocols	Ports	Description
Parallels Client	HALB	TCP, UDP	80, 443	Management and user session connections.
		TCP, UDP	20009	Device Manager shadowing via Firewall (indirect network connection).
	RAS Secure Gateway Forwarding mode	TCP, UDP	80, 443	Management and user session connections.
		TCP, UDP	3389	Optional - Used for user session if RDP load balancing is enabled (Standard RDP).
		UDP	20000	Secure Gateway lookup broadcast.
	RAS Secure Gateway Normal mode	TCP, UDP	80, 443,	Management and user session connections.
		TCP, UDP	3389	

Source	Destination	Protocols	Ports	Description
		TCP, UDP	20009	Optional - Used for user session if RDP load balancing is enabled (Standard RDP).
		UDP	20000	Device Manager shadowing via Firewall (indirect network connection)
				Secure Gateway Lookup Broadcast.
	Session host (VDI, RDS, RemotePC)	TCP, UDP	3389	Used for user session connections in Direct Mode only. RDP connection is always encrypted.
	Azure Virtual Desktop Services	TCP	443	Azure Virtual Desktop Gateway connection
		UDP	3390	Used for user session connections in

Source	Destination	Protocols	Ports	Description
				ShortPath mode only.
	Microsoft site	TCP	443	Download Microsoft Remote Desktop (MSRDC) client or Microsoft Windows App.
	Parallels site	TCP	80, 443	Check for updates and download Parallels Client.

Web browsers

Source	Destination	Protocol	Port	Description
Web browser (HTML5) and Let's Encrypt service	RAS Web Admin Service [RAS Management Portal]	TCP	20443	Admin access to HTML5 based Management Portal of RAS environment
	HALB	TCP	80, 443	End-user access to Parallels RAS Web Client (on Secure Gateway in Normal mode) through the HALB Note: Ports 80 and 443 must be open for

Source	Destination	Protocol	Port	Description
				incoming requests when using Let's Encrypt.
	RAS Secure Gateway	TCP	80, 443	End-user access to Parallels RAS Web Client (on Secure Gateway in Normal mode) Note: Ports 80 and 443 must be open for incoming requests when using Let's Encrypt.

HALB

Source	Destination	Protocol	Port	Description
HALB	HALB	VRRP	112	HALB to HALB communication used for automatic assignment of VIP to active HALB.
	RAS Secure Gateway in Forwarding Mode	TCP, UDP	80, 443	Management and user session connections.

Source	Destination	Protocol	Port	Description
				Management and user session connections.
	RAS Secure Gateway in Normal Mode	TCP, UDP	80, 443	Device Manager shadowing via Firewall (indirect network connection).
		TCP, UDP	20009	

RAS Secure Gateway

Source	Destination	Protocol	Port	Description
				Management and user session connections.
RAS Secure Gateway in Forwarding mode	RAS Secure Gateway in Normal mode	TCP, UDP	80, 443	Optional - Used for user session if RDP Load Balancing is enabled.
		TCP, UDP	3389	
	RAS Performance Monitor	TCP	8086	Agent (Telegraf service) sends collected performance data to InfluxDB.
RAS Secure Gateway in Normal mode	Remote Desktop Services	TCP, UDP	3389	RDP Connections.

Source	Destination	Protocol	Port	Description
	RAS Connection Broker	TCP TCP, UDP	20002 20009	RAS Connection Broker service port - communications with RAS Secure Gateways and the RAS Console (in Normal mode only). Device Manager shadowing via Firewall (indirect network connection) if RAS Console runs on RAS Connection Broker
	RAS Performance Monitor	TCP	8086	Agent (Telegraf service) sends collected performance data to InfluxDB.
	Localhost	TCP	20020	Communication with User Portal web server (NodeJS).

RAS Connection Broker

Source	Destination	Protocols	Ports	Description
RAS Connection Broker	AD DS controllers	TCP TCP TCP,UDP UDP	389, 3268 636, 3269 88 53	LDAP LDAPS Kerberos DNS
	RAS Connection Broker	TCP	20001 20030	Redundancy service. Communication between RAS Connection Brokers running in the same site.
	Parallels Licensing Server	TCP	443	RAS Connection Broker (primary Connection Broker in Licensing Site) communicates with Parallels Licensing Server (https://ras.parallels.com). Note: Not required for Tenant Broker RAS Connection Broker (see the Tenant Broker section).

Source	Destination	Protocols	Ports	Description
	RAS Performance Monitor	TCP	8086	Agent (Telegraf service) sends collected performance data to InfluxDB.
	RAS RD Session Host Agent	TCP, UDP	30004	Server for Connection Broker requests.
	RAS Provider Agent	TCP, UDP	30006	Provider Agent communication port.
	RAS Remote PC Agent	TCP, UDP	30004	Remote PC Agent Communication Port (agent state, counters and session information)
	2FA Server(s)	TCP, UDP	8080, 80 1812, 1813	Deepnet/ Safenet Radius
	RAS Enrollment Server	TCP	30030	RAS Connection Broker Sends RAS Enrollment Server connection Request
	RAS Reporting	TCP	30008	Master RAS Connection Broker communicates with RAS Reporting (installed on the same host as SSRS).
	RAS Remote Installer Service	TCP	30020	Remote agent pushing

Source	Destination	Protocols	Ports	Description
	RAS RD Session Host Agent RAS Guest Agent RAS Remote PC Agent RAS Connection Broker RAS Secure Gateway RAS Enrollment Server	TCP	135, 445, 49179	Remote Install Push/Takeover of Software
	SMTP	TCP	587	Notifdispatcher is the service which sends the emails using port specified in the Mailbox settings (+SSL/TLS)
	Let's Encrypt Service	TCP	80, 443	Communication between the Let's Encrypt client (available in the primary Connection Broker) and a Let's Encrypt server.

RAS Console

Source	Destination	Protocols	Ports	Description
RAS Console	RAS Reporting	TCP	30008	RAS Console is connected to primary RAS Connection Broker which communicates with RAS Reporting (installed on the same host as SSRS). SSRS talks to SQL via TCP 1433 (or dynamic if 1433 is not established in the settings).
	SSRS	TCP	443	Reports retrieval.
	HALB	TCP, UDP	31006	Used for configuration.
	Parallels Client	TCP	50005	Shadowing from the RAS Console in case of direct network connection.
	RAS RD Session Host Agent	UDP, TCP	30004	
	RAS Guest Agent			
	RAS Remote PC Agent	UDP, TCP	30004	
	RAS Provider Agent	UDP, TCP	30006	
	MFA Server(s)	TCP, UDP	8080, 80, 1812, 1813	Deepnet / Safenet / Radius
	Microsoft site	TCP	80, 443	Check for updates and download Parallels Client
	Parallels site	TCP	80	Check for updates and download Parallels Client
	RAS Secure Gateway	TCP	80, 443	
	RAS Performance Monitor	TCP	20002, 20001	RAS browser plugin connection to Grafana.

Source	Destination	Protocols	Ports	Description
	RAS Connection Broker	TCP	20002, 20001	Communication with Connection Broker and redundancy.
	RAS Enrollment Server	TCP, UDP	30030	
	Wyse Broker	UDP		
	SMTP	TCP	587	RAS Console can send test emails using port specified in the Mailbox settings (+SSL/TLS)

SSRS

Source	Destination	Protocol	Port	Description
SSRS	Microsoft SQL Server	TCP	1433	RAS Console is connected to RAS Reporting

RAS Reporting

Source	Destination	Protocol	Port	Description
RAS Reporting Service	MS SQL	TCP	1433	Store RAS activity information
	SSRS	TCP	8085, 443	Enumeration of reports (incl. custom reports)

RAS Web Administration Service (REST/Management Portal)

Source	Destination	Protocols	Ports	Description
RAS Web Administration Service	RAS RD Session Host Agent	TCP	30004	Log retrieval
	RAS Guest Agent	TCP	30010	Log retrieval
	RAS Provider Agent	TCP	30006	Log retrieval

Source	Destination	Protocols	Ports	Description
	RAS Connection Broker	TCP	20002, 20001 30020	
		TCP	135, 445	Remote Install Push/ Takeover of Software (pre-RAS 18).
	RAS Reporting Service	TCP	3000	Integration of RAS Reporting in Management Portal iFrame

RAS PowerShell

Source	Destination	Protocols	Ports	Description
RAS PowerShell	RAS RD Session Host Agent	TCP	30004	Log retrieval
	RAS Guest Agent	TCP	30010	Log retrieval
	RAS Remote PC Agent	TCP	30004	Log retrieval
	RAS Provider Agent	TCP	30006	Log retrieval
	RAS Connection Broker	TCP	20002, 20001	

RAS Provider Agent

Source	Destination	Protocols	Ports	Description
RAS Provider Agent	RAS Connection Broker	TCP	20003	Connection Broker communication port.

Source	Destination	Protocols	Ports	Description
	RAS Guest Agent			
	RAS Performance Monitor	TCP	8086	Agent (Telegraf service) sends collected performance data to InfluxDB - applicable to Hyper-V only.
	Hyper-V	TCP	135, 49152-65535	Used to check if the host is powered on and send export, import, delete, shutdown, restart or suspend commands.
	Nutanix AHV (AOS)	TCP	9440	Used to check if the host is powered on and sends clone, delete, shutdown, restart commands (RestAPI calls, PoSH, remote ncli).
	VMWare	TCP	443	Used to check if the host is powered on and sends clone, delete, shutdown, restart and suspend commands.
	Microsoft Azure	TCP	443	Used to check if the guest is powered on and sends clone, shutdown, restart commands (via REST).
	Azure Virtual Desktop	TCP	443	Used to check if the host is powered on and sends clone, shutdown, restart commands (via REST).
	AWS	TCP	443	Used to check if the host is powered on and sends clone, shutdown, restart commands (via REST).
	Scale	TCP	443	Used to check if the host is powered on and sends clone, shutdown, restart commands (via REST).
	Remote PC over VDI	TCP	135, 49152-65535	Used to check if the host is powered on and sends

Source	Destination	Protocols	Ports	Description
				shutdown, restart or suspend commands.

RAS Enrollment Server

Source	Destination	Protocol	Ports	Description
RAS Enrollment Server	AD DS controllers	TCP	389, 3268	LDAP
		TCP	636, 3269	LDAPS
		TCP,UDP	88	Kerberos
		UDP	53	DNS
	RAS Connection Broker	TCP	20003	Settings synchronization and performance counters.
		UDP	20003	Deny Connection Request
	Certificate Authority (CA)	TCP	135	DCOM/RPC ports
		TCP	dynamic range 49152 - 65535	

RAS RD Session Host Agent

Source	Destination	Protocols	Ports	Description
RAS RD Session Host Agent	RAS Connection Broker	TCP, UDP	20003	Used for communications with RAS Connection Brokers.

Source	Destination	Protocols	Ports	Description
	Localhost	TCP	30005	For internal commands (memshell, printer redirector).
	FSlogix	TCP	443	Download FSlogix installer
	RAS Performance Monitor	TCP	8086	Agent (Telegraf service) sends collected performance data to InfluxDB.
	RAS Enrollment Server	TCP	30030	RAS RD Session Host Agent (PrIsSCDriver) connects to get logon credentials.

RAS Guest Agent

Source	Destination	Protocols	Ports	Description
RAS Guest Agent (used by Azure Virtual Desktop)	Provider Agent	TCP, UDP	30006	
	Localhost	TCP	30005	For internal commands - memshell, printer redirector)
	RAS Performance Monitor	TCP	8086	Agent (Telegraf service) sends collected performance data to InfluxDB
	RAS Enrollment Server	TCP	30030	RAS Guest Agent (PrIsSCDriver) connects to get logon credentials
	FSlogix	TCP	443	Download FSlogix installer

RAS Remote PC Agent

Source	Destination	Protocols	Ports	Description
RAS Remote PC Agent	RAS Connection Broker	TCP, UDP	20003	Used for communications with RAS Connection Brokers
	Localhost	TCP	30005	For internal commands - memshell, printer redirector)

Source	Destination	Protocols	Ports	Description
	RAS Performance Monitor	TCP	8086	Agent (Telegraf service) sends collected performance data to InfluxDB
	RAS Enrollment Server	TCP, UDP	30030	RAS Remote PC (PrIsSCDriver) connects to get logon credentials
	FSlogix	TCP	443	Download FSlogix installer

Tenant Broker

Source	Destination	Protocols	Ports	Description
Tenant - RAS Connection Broker	Tenant Broker - RAS Connection Broker	TCP	20003	Tenant's RAS Connection Broker communicates with Tenant Broker to join Tenant Broker, synchronize configuration and statuses

Active Directory and Domain Services ports

For Active Directory and Active Directory Domain Services port requirements, please see the following article: <https://technet.microsoft.com/en-us/library/dd772723%28v=ws.10%29.aspx>.

Azure Virtual Desktop

The Azure virtual machines you create for Azure Virtual Desktop must have access to the following URLs in the Azure commercial cloud:

Address	Outbound TCP port	Purpose	Service tag
*.wvd.microsoft.com	443	Service traffic	AzureVirtualDesktop
gcs.prod.monitoring.core.windows.net	443	Agent traffic	AzureCloud
production.diagnostics.monitoring.core.windows.net	443	Agent traffic	AzureCloud

Address	Outbound TCP port	Purpose	Service tag
*xt.blob.core.windows.net	443	Agent traffic	AzureCloud
*eh.servicebus.windows.net	443	Agent traffic	AzureCloud
*xt.table.core.windows.net	443	Agent traffic	AzureCloud
*xt.queue.core.windows.net	443	Agent traffic	AzureCloud
catalogartifact.azureedge.net	443	Azure Marketplace	AzureCloud
kms.core.windows.net	1688	Windows activation	Internet
mrsglobalsteus2prod.blob.core.windows.net	443	Agent and SXS stack updates	AzureCloud
wvdportalstorageblob.blob.core.windows.net	443	Azure portal support	AzureCloud
169.254.169.254	80	Azure Instance Metadata service endpoint	N/A
168.63.129.16	80	Host health monitoring	N/A
https://download.parallels.com/ras/Configuration_01-20-2022.zip	443	Joining a host to a host pool	AzureVirtualDesktop

The following table lists optional URLs that your Azure virtual machines can have access to:

Address	Outbound TCP port	Purpose	Azure Gov
*.microsoftonline.com	443	Authentication to Microsoft Online Services	login.microsoftonline.us
*.events.data.microsoft.com	443	Telemetry Service	None
www.msftconnecttest.com	443	Detects if the OS is	None

Address	Outbound TCP port	Purpose	Azure Gov
		connected to the internet	
*.prod.do.dsp.mp.microsoft.com	443	Windows Update	None
login.windows.net	443	Sign in to Microsoft Online Services, Microsoft 365	login.microsoftonline.us
*.sfx.ms	443	Updates for OneDrive client software	oneclient.sfx.ms
*.digicert.com	443	Certificate revocation check	None
*.azure-dns.com	443	Azure DNS resolution	None
*.azure-dns.net	443	Azure DNS resolution	None

For up to date information, please also visit the Microsoft website at <https://docs.microsoft.com/en-us/azure/virtual-desktop/safe-url-list#required-url-check-tool>.

RAS performance counters

The following table lists performance counters available in Parallels RAS per component:

Parallels RAS Gateway (2XProxyGateway.exe)

ID	Name	Description
ras_gw_tot_conn	Total connections	The total number of Connections with the Gateway.
ras_gw_tot_threads	Total threads	The total number of threads running on the Gateway.

ID	Name	Description
ras_gw_rpd_sess	RDP tunneled sessions	The number of tunneled RDP sessions.
ras_gw_rpd_sess_s	RDP SSL tunneled sessions	The number of tunneled RDP sessions over SSL.
ras_gw_html	HTTP connections	The number of tunneled HTTP sockets
ras_gw_html_s	HTTPS connections	The number of tunneled HTTPS sockets
ras_gw_html5	HTML5 connections	The number of tunneled HTTP5 sockets
ras_gw_html5_s	HTML5 SSL connections	The number of tunneled HTTP5 sockets over SSL
ras_gw_cm	Device Manager connections	The number of Parallels Device Manager connections
ras_gw_cm_s	Device Manager SSL connections	The number of Parallels Device Manager connections over SSL
ras_gw_wyse	Wyse connections	The number of Wyse connections
ras_gw_wyse_s	Wyse SSL connections	The number of Wyse connections over SSL
ras_gw_rdpudp	RDP UDP tunneled sessions	The number of RDP UDP connections
ras_gw_rdpudp_s	RDP UDP DTLS tunneled sessions	The number of RDP UDP connections over DTLS
ras_gw_cache_sock	Cached sockets	The number of cached sockets between Gateway and Connection Broker
ras_gw_idle_threads	Idle threads	The number of idle threads on the Gateway
ras_gw_client	Client connections	The number of Parallels Client connections
ras_gw_client_s	Client SSL connections	The number of Parallels Client connections over SSL

Parallels RAS Connection Broker (2XController.exe)

ID	Name	Description
ras_pa_avg_client_connection_time		

ID	Name	Description
	Average time for client connection	The average client connection time.
ras_pa_avg_client_auth_time	Average time for user authentication	The average time taken to authenticate a user.
ras_pa_avg_client_policy_time	Average time to retrieve user policy	The average time taken to retrieve the user's policy.
ras_pa_avg_client_rep_time	Average time to send client telemetry	The average time taken to send client telemetry. Used by CEP.
ras_pa_avg_client_applist_time	Average time to retrieve user's published items	The average time taken to retrieve user's published items list.
ras_pa_avg_client_appicons_time	Average time to retrieve icons	The average time taken to retrieve published items icons.
ras_pa_avg_client_getidle_time	Average time to start up a request	The average time taken for the start-up request.

Parallels RAS RDS Agent (2XAgent.exe)

ID	Name	Description
act_sess	Active RDS sessions	The number of active RDS Sessions.
disc_sess	Disconnected RDS sessions	The number of disconnected RDS Sessions.

1. Common Management Tasks



2. Common Management Tasks



3. **Common Management Tasks**



4. **Common Management Tasks**



5. **Common Management Tasks**



6. **Common Management Tasks**



7. **Common Management Tasks**



8. **Common Management Tasks**



9. **Common Management Tasks**

