v1.4 | January 2025 | RA-2081

REFERENCE ARCHITECTURE

Parallels RAS on Nutanix AHV

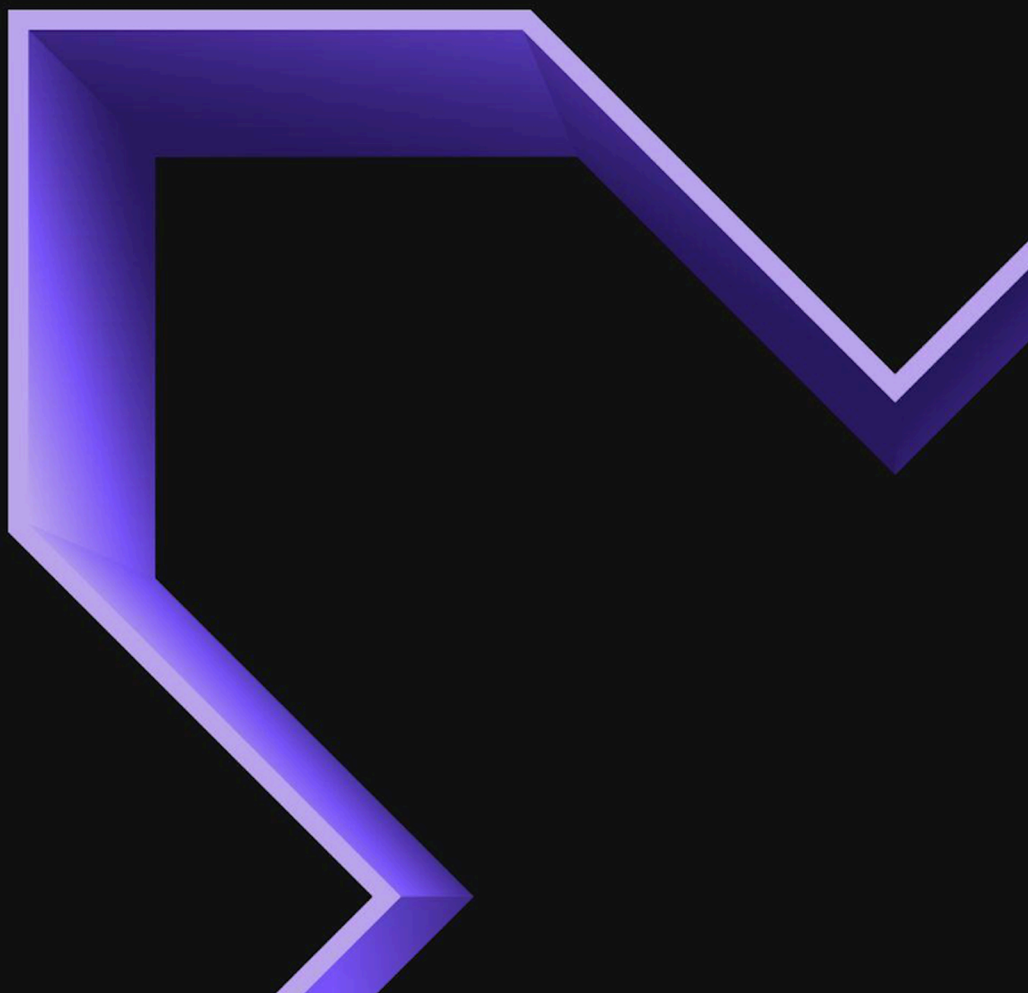

NUTANIX

Legal

Nutanix, Inc.
1740 Technology Drive, Suite 150
San Jose, CA 95110

Contents

1. Executive Summary

Nutanix designed its software to give customers running workloads in a hybrid cloud environment the same experience that they expect from on-premises Nutanix clusters. Because Nutanix in a hybrid multicloud runs Nutanix AOS and Nutanix AHV with the same CLI, UI, and APIs, existing IT processes and third-party integrations continue to work regardless of where they run.

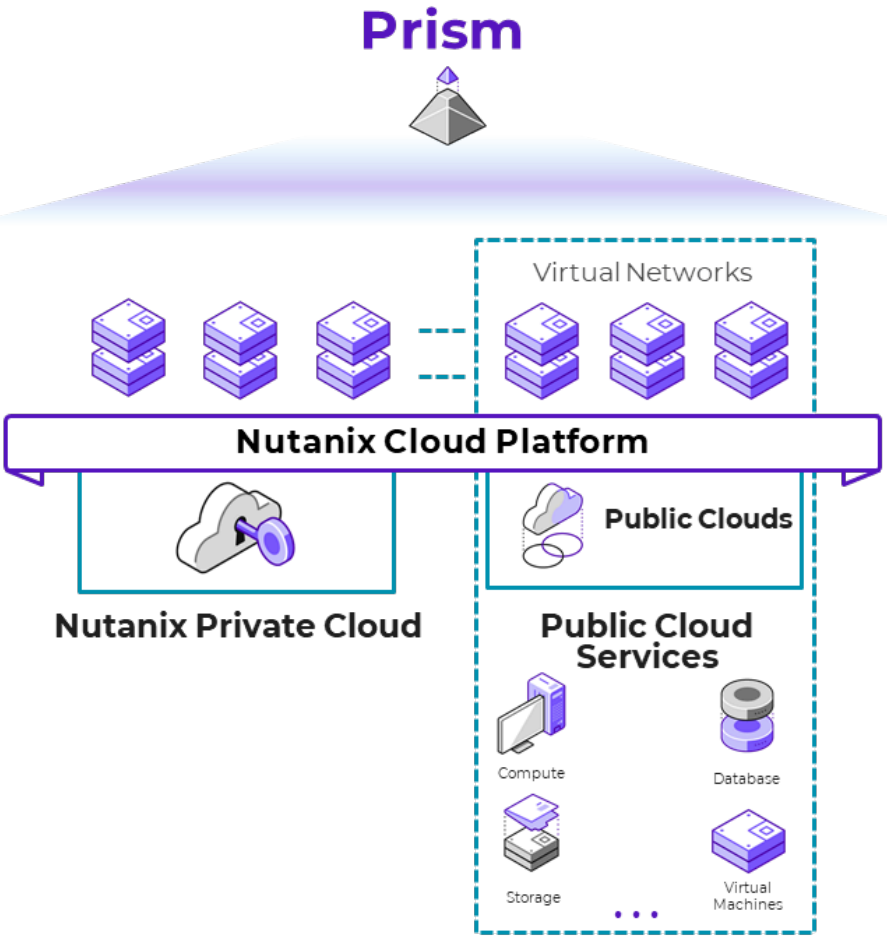

Figure 1: Overview of the Nutanix Hybrid Multicloud Software

AOS can withstand hardware failures and software glitches and ensures that application availability and performance aren't compromised. By combining features like native rack awareness with public cloud partition placement groups, Nutanix operates freely in a dynamic hybrid multicloud environment.

Parallels Remote Application Server (RAS) provides vendor-independent virtual desktop and application delivery from a single platform. Parallels RAS extends Windows Remote Desktop Services by using a customized shell and virtual channel extensions over the Microsoft Remote Desktop Protocol.

In this reference architecture, we make recommendations for designing, optimizing, and scaling Parallels RAS. We used Login Enterprise (Login VSI) and an automated scripting framework on Nutanix to simulate real-world workloads in a Parallels RAS environment.

In addition to desktop and application performance reliability, deploying Parallels RAS on Nutanix provides unlimited scalability, data locality, and a single datastore. Nutanix takes the Parallels commitment to simplicity to another level with streamlined management, reduced rollout time, and decreased operating expenses.

This document covers the following topics:

- Overview of the Nutanix solution
- Overview of Parallels RAS and its use cases
- Benefits of Parallels RAS on the Nutanix Cloud Platform
- Design and configuration considerations for building a Parallels RAS solution on Nutanix
- Process for benchmarking Parallels RAS performance on Nutanix AHV running on Intel Xeon Gold 6342 processors using Nutanix NX hardware

Table: Document Version History

Version Number	Published	Notes
1.0	November 2023	Original publication.
1.1	December 2023	Updated the Parallels RAS on Nutanix AHV with Nutanix Cloud Platform and Parallels RAS Benchmarking and Test Environment sections.

Version Number	Published	Notes
1.2	November 2024	Updated AOS release cycle descriptions.
1.3	December 2024	Refreshed content.
1.4	January 2025	Updated LoginVSI workload links.

2. Parallels RAS Overview

You can access Parallels RAS from anywhere with platform-specific clients and web-enabled solutions, like the built-in Parallels Web Client. Parallels RAS publishes remote desktops, applications, and documents from a single platform, improving desktop manageability, security, and performance. Parallels RAS supports Nutanix AHV, integrates with Prism Element, and enables the publishing of virtual desktops and applications to Parallels Client. Nutanix has validated Parallels RAS up to version 6.5 LTS.

The following components make up the core of Parallels RAS. For more information, see the Parallels RAS documentation.

RAS farm

A Parallels RAS farm is a logical grouping of objects for centralized management. A single database that contains information about all the objects in the farm holds the farm configuration. A Parallels RAS farm consists of at least one site.

RAS site

A site is the next level in the farm hierarchy and groups servers and other objects that provide connection and remote application services. You can use sites to separate management and location functions.

RAS licensing site

A single site in a RAS farm must be configured as the licensing site.

RAS console

The Parallels RAS console administers the RAS environment. By default, the Parallels RAS console installs on the same server where you install other Parallels RAS components, but you can install the console on any computer on your network.

RAS Connection Brokers

RAS Connection Brokers provide load balancing for published applications and desktops. A RAS Connection Broker installs automatically on a server when you

install Parallels RAS. Multiple Connection Brokers can exist in a site to provide resilience and distribute load.

RAS Secure Gateways

A 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. You must install and configure at least one RAS Secure Gateway in every site. You can add gateways to a RAS site to support more users, load-balance connections, and provide redundancy.

High availability and load balancing (HALB)

HALB in Parallels RAS load-balances RAS Secure Gateways. A Parallels HALB appliance, a VM with the operating system installed and all relevant settings configured, contains the load balancer.

RAS User Portal

User Portal in RAS Secure Gateway allows users to connect to Parallels RAS and open published resources from a web browser using the Parallels Web Client.

Parallels Client

You can launch a session from an installed client on your endpoint or through an HTML5-based web portal.

Providers

To allow a provider (hypervisor or cloud-based) to function in a RAS farm, install RAS Provider Agent 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.

Remote Desktop Session Host (RDSH)

A Parallels RAS farm uses RDSHs to host published resources (for example, applications, desktops, and documents).

RDSH templates

RDSH templates replicate RDSHs running in VMs. Parallels RAS treats guest VMs from an RDSH template almost like regular RDSHs.

Virtual desktop infrastructure (VDI)

Parallels RAS VDI uses server virtualization to reduce the number of physical servers required to host published resources.

VDI templates

Virtual desktop templates are an essential part of the Parallels RAS VDI and are used to create guest VMs for publishing desktops, applications, and documents. Guest VMs created from a VDI template typically serve a single user.

VDI pool

Pools offer administrators flexibility when managing a large number of guest VMs, especially when they are implemented in large infrastructures.

Groups

With groups, you can combine multiple RDSHs and then publish the resources from the group instead of specifying individual servers.

RASprep

RASprep is the Parallels RAS tool for preparing Windows in a VM after cloning it from a base image. RASprep is similar to Sysprep; however, RASprep works faster because it modifies a smaller number of configurable parameters and requires fewer restarts.

RAS Guest Agent

A guest VM must have the RAS Guest Agent installed and the agent must match the Parallels RAS version.

3. Parallels RAS on Nutanix

The Nutanix modular web-scale architecture lets you start small and expand to meet increasing demand—a node, a block, or multiple blocks at a time—with no effect on performance. This design removes the hurdle of a large initial infrastructure purchase, decreasing the time to value for your Parallels RAS implementation. With Parallels RAS on Nutanix, you can run multiple workloads, all on the same scalable converged infrastructure, while achieving the following benefits:

Modular incremental scale

With the Nutanix solution, you can start small and scale up. A single Nutanix block provides dozens of terabytes of storage and hundreds to thousands of virtual desktops in a compact footprint. With the solution's modularity, you can granularly scale by node, by block, or with multiple blocks, accurately matching supply with demand and minimizing upfront cost.

High performance

By using system memory caching for read I/O and flash storage for read and write I/O, you can deliver high-performance throughput in a compact form factor.

Change management

Maintain environmental control and separation between development, test, staging, and production environments. Snapshots and fast clones can help share production data with nonproduction jobs without requiring full copies and unnecessary data duplication.

Business continuity and data protection

User data and desktops are mission-critical and need enterprise-grade data management features, including backup and disaster recovery.

Data efficiency

Nutanix storage offers both compression and deduplication to help reduce the storage footprint. The compression functionality is truly VM-centric. Unlike traditional solutions that perform compression mainly at the LUN level, the Nutanix solution provides all these capabilities at the VM and file levels, greatly

increasing efficiency and simplicity. These capabilities ensure the highest possible compression and decompression performance, even below the block level.

Enterprise-grade cluster management

Nutanix offers a simplified and intuitive approach to managing large clusters, including a converged user interface that serves as a central point for servers and storage, alert notifications, and the bonjour mechanism that automatically detects new nodes in the cluster.

High-density architecture

Nutanix uses an advanced server architecture coupled with data archiving and compression that can make your desktop hardware footprint five times smaller.

Time-sliced clusters

Like public cloud environments, Nutanix can provide a truly converged cloud infrastructure, allowing you to run your server and desktop virtualization on a single cloud. Gain the efficiency and savings that you require with a converged cloud on a unified architecture.

Parallels RAS on Nutanix AHV with Nutanix Cloud Platform

External devices connect through a customer firewall, proxied through the RAS Secure Gateway to the Connection Brokers through a Parallels RAS HALB appliance. Internal devices also connect directly through the Parallels RAS HALB appliance without traversing the external firewall. Ultimately, the system brokers both VDI and RDSH workloads. For this document, we completed our testing with internal devices.

For more information, see the Parallels RAS documentation for supported deployment architectures.

Parallels RAS Pod Design

The Parallels RAS on Nutanix–hosted virtual desktop control pod has the following design:

- Two Connection Broker roles
- Two Secure Gateway roles

- One provider per Nutanix cluster
- One HALB appliance

> **Note:** We chose to colocate the Connection Broker and Secure Gateway roles for simplicity. For more information on scale considerations for specific components, see Parallels RAS documentation.

The Parallels RAS on Nutanix–hosted virtual desktop services pod has the following design:

- Up to 16 AHV hosts
- One Nutanix cluster
- One datastore
- 100 VDI desktops per host based on load-testing results
- Eight RDSH per host based on load-testing results

> **Note:** Limiting the cluster sizes to 16 nodes reduces the failure domain and the time to patch and update the clusters. This reference architecture uses 16-node building blocks to take advantage of single-rack design, smaller failure domain, and reduced time to patch and update.

We validated Parallels RAS using the Login Enterprise knowledge worker workload with both Windows Server and Windows desktop workloads. We used a 2 vCPU and 4 GB of memory configuration per VM for VDI workloads with 100 VMs per node and an 8 vCPU with 42 GB of memory configuration per VM for RDSH workloads with 8 VMs per node.

For VDI workloads, if you scale a cluster to 16 nodes, you can run 1,500 VDI desktops per cluster. For RDSH workloads, if you scale a cluster to 16 nodes and your workload is similar to the knowledge worker profile, you can host 120 session hosts with 1,800 user sessions.

> **Note:** This calculation reserves one node as a spare (n + 1).

A more resource-intensive workload than the knowledge worker results in a lower density, and a less resource-intensive workload results in a higher density. If you change the vCPU count or memory, the number of Windows VDAs per node and per cluster changes as well.

Nutanix Compute and Storage

Nutanix provides an ideal combination of high-performance compute resources and localized storage to meet any demand. True to this capability, this reference architecture contains no reconfiguration or customization of the Nutanix product to optimize for this use case. The following figure shows a high-level example of the relationship between the Nutanix storage pool and containers, where a single storage pool can host multiple containers with different configurations.

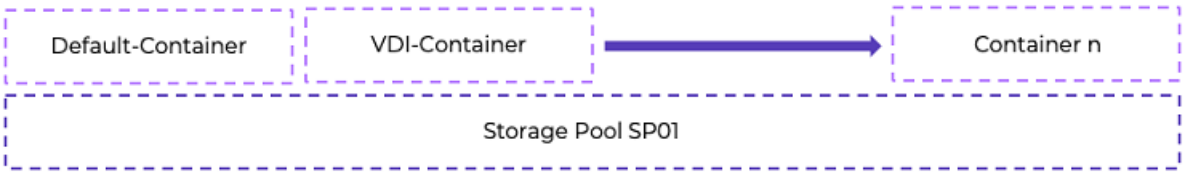

Figure 2: Storage Overview

The following table details the Nutanix storage pool and container configuration.

Table: Nutanix Storage Configuration

Name	Role	Details
SP01	Main storage pool for all data	SSD
VDI	Container for all VMs	AHV datastore
Default-Container	Container for all data (not used here)	AHV datastore

4. Parallels RAS on Nutanix Solution Design

In this section, we describe the design decisions and rationale for Parallels Remote Application Server (RAS) deployments on Nutanix.

Table: Design Decisions

Design Area	Item	Detail	Rationale
General	Software versions	Parallels RAS 19.2	Latest version available
General	Minimum size	3 Nutanix AOS hosts running AHV	Minimum size requirement
General	Scale approach	Incremental modular scale	Allows growth from proof of concept (hundreds of desktops) to massive scale (thousands of desktops)
General	Scale unit	Nodes, blocks, or pods	Granular scale to precisely meet capacity demands; scale in n × node increments
Nutanix AHV	Cluster size	As many as 16 hosts (minimum of 3 hosts)	Isolated fault domains (best practice)
Nutanix AHV	Storage pools	1 storage pool per cluster	Standard practice; intelligent tiering handles data locality
Nutanix AHV	Containers	1 container for VMs	Nutanix handles I/O distribution and localization; n-Controller model

Design Area	Item	Detail	Rationale
Nutanix AHV	Infrastructure services	Small deployments: shared cluster; Large deployments: dedicated cluster	Dedicated infrastructure cluster for larger deployments (best practice)
Nutanix AHV	Features and enhancements	Increase Controller VM (CVM) memory to 24–32+ GB. Turn on deduplication and compression. We set the CVM to 32 GB for the reference architecture.	Best practice
Parallels RAS	Connection Brokers	Minimum: 2 (n + 1); Scale: 1 per additional pod	High availability for the Connection Broker function
Parallels RAS	Secure Gateways	Minimum: 2 (n + 1); Scale: 1 per additional pod	High availability for the Secure Gateway function
Parallels RAS	Load balancing	RAS HALB appliance	Ensures distribution of traffic across the Secure Gateway roles; Parallels recommends a minimum of two
Parallels RAS	Virtual hardware specs	vCPU: 4; Memory: 8+ GB; Disk: 60 GB vDisk	Standard sizing practice
Parallels RAS	Desktop management	Parallels RAS AHV Provider	Allows deployment and management
Parallels RAS	Image preparation	RASprep	A slimmed-down version of Sysprep suitable for our workloads

Design Area	Item	Detail	Rationale
Parallels RAS provider configuration	Provider	A single AHV Provider Agent deployed locally on the RAS Connection Broker	Integrated with a single Nutanix AHV Prism Element (workload cluster) using a local service account for RAS provisioning operations
Active Directory infrastructure	Global catalog and DNS servers	Minimum: 2 (n + 1) per site	High availability for global catalog and DNS; Microsoft best practice
DHCP infrastructure	DHCP servers	Nutanix IPAM	High availability for Nutanix IPAM is built in
DHCP infrastructure	Load balancing	Built-in	Ensures availability of DHCP

Desktop Optimizations

We generated our design with the following high-level desktop optimization guidelines in mind:

- Size desktops appropriately for each use case.
- Disable unnecessary OS services and applications.
- Redirect home directories or use a profile management tool for user profiles and documents.

For more information on specific Parallels RAS desktop optimizations, see the Optimization section of the Parallels RAS 19 Administrators Guide.

Note: We found that Parallels RAS optimizations were missing significant optimizations found in other tooling. As such, we aligned our image optimizations to match the specifications in Nutanix tech notes on VDI performance.

5. Parallels RAS Benchmarking and Test Environment

The following sections describe the benchmarking method and test environment we used in this reference architecture.

Login Enterprise

Login VSI provides the industry-standard virtual desktop testing platform, Login Enterprise, which helps organizations benchmark and validate the performance and scalability of their virtual desktop solutions. With Login Enterprise, IT teams can reliably measure the effects of changes to their virtual desktop infrastructure on user experience and identify performance issues before they impact the business. Login Enterprise uses synthetic user workloads to simulate real-world user behavior, so IT teams can measure the responsiveness and performance of their virtual desktop environment under different scenarios. Login Enterprise has two built-in workloads: task worker and knowledge worker.

> **Note:** You can't compare the Login Enterprise workloads to the workloads included in the previous edition of Login VSI. The Login Enterprise workloads are much more resource intensive.

The following table includes both workloads available in Login Enterprise.

Table: Login Enterprise Workloads

Task Worker	Knowledge Worker
Light	Medium
2 vCPU	2–4 vCPU
2–3 apps	4–6 apps
No video	720p video

Login Enterprise EUX Score

According to the Login Enterprise documentation, the End User Experience (EUX) score represents the performance of any Windows machine (virtual, physical, cloud, or on-premises). The score ranges from 0 to 10 and measures the experience of one to many virtual users.

> **Note:** Expect your EUX score to drop as you add users to your VDI. More users demand a greater share of a VDI system's shared resources, so performance and user experience decrease.

We interpret EUX scores with the grades in the following table.

Table: EUX Score Grades

EUX Score	Grade
1–5	Bad
5–6	Poor
6–7	Average
7–8	Good
8–10	Excellent

Login Enterprise VSImax

For our test results, we used the 2023 EUX score's version of VSImax. In this version, a number of triggers determine the VSImax (or the maximum number of users). These triggers are CPU- and disk-related operations and can determine whether the user experience is acceptable. EUX scores below 5.5 are one example of a trigger.

We can use this version of the VSImax to run a comparison, but the VSImax on its own doesn't represent the maximum user density accurately. For a more realistic maximum number of users, we suggest using the number of active users when the EUX score is 85 to 90 percent of the initial EUX score.

> **Note:** In the 2023 release of EUX score, the disk-related operations of EUX have an unrealistic impact on storage. In our testing, we discovered that the IOPS are up to 10 times higher when the EUX metrics are enabled, with a read-to-write ratio of 70:30 percent during the steady state. In reality, a knowledge worker has a much lower I/O profile and a read-to-write ratio of 20 to 30 percent reads and 70 to 80 percent writes.

Login Enterprise Metrics

We quantified the evaluation using the following metrics:

- EUX base: The average EUX score of the first five minutes
- EUX score: The average EUX score for the entire test
- Steady state score: The average EUX score starting five minutes after the final logon to the end of the test
- Average logon time: The average user logon time
- VSImax: If reached, the maximum value of sessions launched before the VSI Index Average reaches one of the thresholds
- Maximum CPU usage: The maximum observed CPU usage during the test
- CPU usage during steady state: The average CPU usage during the steady state, or the state when all the sessions are active and using applications; this state simulates user activity during the entire day, rather than just during the logon period

The Baseline and Steady State EUX scores provide additional dimensions to the simulated user experience. The Standard EUX score provides a single score for the entire test duration, including the logon period and the application interaction period. As you add more users to the system that you're testing, it works harder and the user experience diminishes. The Steady State and Baseline EUX scores describe the user experience during specific periods of the test run.

Baseline EUX score

The Baseline EUX score represents the best possible performance of the system and is the average EUX score from the best five minutes of the test. This score indicates how the system performs when it's not under stress. Typically, you capture the Baseline EUX score at the beginning of the test before the system is fully loaded.

Steady State EUX score

The Steady State EUX score represents the period after all users log on and the system starts to normalize. The Steady State EUX score is the average of the EUX scores captured between five minutes after all sessions log on through the end of the test.

Login Enterprise Graph

A Login Enterprise graph shows the values obtained during the launch for each desktop session. The y-axis on the left side measures the EUX score, the y-axis on the right side measures the number of active sessions, and the x-axis represents the test duration in minutes. We configured our benchmark test to sign in all sessions in 48 minutes, followed by a steady state of 10 minutes.

Parallels RAS on Nutanix Test Environment

In this section, we describe the hardware we used for this reference architecture.

Management Infrastructure Cluster

We used one Nutanix NX-3060-G7 cluster with four nodes to host all infrastructure and Parallels RAS services and the Login Enterprise appliance. Active Directory services and Domain Name System (DNS) also ran in this cluster, which we designated the management infrastructure cluster. Four nodes provided resources to host these servers. The following table shows the Parallels RAS configuration.

Table: Parallels RAS Configuration

VM	Quantity	vCPU	Memory	Disks
Connection Broker and Secure Gateway (combined)	2	4	8 GB	1 × 60 GB (OS)
HALB appliances	1	2	4 GB	Default

Login Enterprise Launcher Cluster

To initiate the sessions to the virtual desktops, Login Enterprise uses launcher VMs. Depending on the display protocol used, one launcher VM can host up to 25 sessions. For this reference architecture, we used one Nutanix NX-3060-G7 cluster with four nodes to host 75 launcher VMs. Each launcher VM had 4 vCPU and 6 GB of memory.

Virtual Desktop Cluster

Four Nutanix NX-3155G-G8 nodes formed the cluster to host all virtual desktops. The next tables contain the specifications of this cluster.

Table: Virtual Desktop Cluster Specifications

Parameter	Setting
Block type	Nutanix NX-3155G-G8
Number of blocks	4
Number of nodes	4
CPU type	Intel Xeon Gold 6354
Number of CPUs per node	2
Number of cores per CPU	18
Memory per node	0.98 TiB
Disk config per node	6 × 1.55 TB SSD
Network	2 × 25 GbE

Table: Nutanix Software Specifications

Parameter	Setting
Nutanix AHV version	20220304.423
Nutanix AOS version	6.5.3.6
CVM vCPU	12
CVM memory	32 GB
Redundancy factor	2
Number of containers for worker VMs	1
Datastore specifications	Compression: On; Compression delay: 0; Deduplication: Off

Table: Parallels RAS Software Specifications

Parameter	Setting
Parallels RAS	19.2
Parallels HALB Appliance	19.1 (latest)

Table: Windows 10 Template Image Configuration

Parameter	Setting
Operating system	Windows 10 22H2 (x64)
Windows updates	01/09/23
CPU	2 vCPU
Memory	4 GB
NICs	1
Virtual network adapter	Nutanix VirtIO Adapter
Virtual SCSI controller 0	Nutanix VirtIO SCSI passthrough
Virtual disk	64 GB
Virtual CD/DVD drive 1	Client
Applications	Adobe Acrobat DC, Microsoft Edge, Microsoft Office 2019 (x64)
Parallels RAS Guest Agent	19.2

Table: Windows Server 2022 Template Image Configuration

Parameter	Setting
Operating system	Windows Server 2022
Windows updates	01/09/23
CPU	8 vCPU
Memory	42 GB
NICs	1
Virtual disk	80 GB
Virtual floppy drive	Removed
Virtual CD/DVD drive 1	Client
Applications	Adobe Acrobat DC, Microsoft Edge, Microsoft Office 2019 (x64)
Parallels RAS Guest Agent	19.2
Parallels RDSH Agent	19.2

6. Parallels RAS on Nutanix Test Validation

This section provides the details and results of our Parallels RAS performance tests with both Windows 10 and Windows Server 2022 on Nutanix AHV. We ran multiple test iterations to ensure accuracy.

Test Objectives

Our overall objective was to determine the session capacity we can host on Nutanix while running the Login Enterprise tests with Parallels RAS. We used a Windows 10, version 22H2 image for VDI workloads and Windows Server 2022 for RDSH workloads. We tested with the Login Enterprise knowledge worker profile.

We had the following specific objectives:

- Determine the maximum number of sessions we can host on this system with the Login Enterprise knowledge worker workload while maintaining a good EUX score.
- Show the linear scalability of the Nutanix platform.
- Determine if performance differences exist when using Parallels RAS full clone or linked clone template configurations to provision the workloads on the Nutanix platform.
- Determine the time to provision 400 Windows 10 desktops on four nodes.
- Determine the time to provision 32 Windows Server 2022 RDSHs on four nodes.

General Specifications

Table: Login Enterprise Specifications

Specification	Value
Product version	5.3.3
EUX version	EUX2023
Workload	Knowledge worker

Table: Broker Specifications

Specification	Value
Delivery type	RAS
Desktop Broker version	19.2
Session configuration	RDP

Considerations for Test Results

We used two methods to deploy workloads based on current Parallels RAS capability:

- For VDI workloads, we used the Add VM wizard from the RAS console and selected **Template** and **Quantity**.
- For RDSH workloads, we configured autoscale on the appropriate RAS group. We specified a minimum number of hosts and enabled the group to activate the autoscale logic. By design, a two-minute delay exists between enabling the group and when autoscale begins provisioning.

We tested using a single, full-HD screen as the client and limited the frames per second to 16. Using multiple screens or other screen resolution settings affects the results.

Note: We used several Parallels RAS templates in our testing due to the way Parallels RAS handles provisioning.

Table: Parallels RAS Template Layout

Name	OS	Provisioning Type
W10-22H2-FC	Windows 10 Enterprise	Full clone
W10-22H2-LC	Windows 10 Enterprise	Linked clone
SRV-2022-FC	Windows Server 2022	Full clone
SRV-2022-LC	Windows Server 2022	Linked clone

In Nutanix AHV deployments, Parallels RAS deploys VMs in full clone mode by cloning the template VM directly and linked clone mode from a Nutanix snapshot of the template VM.

Parallels RAS Provisioning Times

We ran multiple provisioning tests to identify how long it took to deploy workloads on Nutanix. We also looked for any difference in Parallels RAS full clone versus linked clone provisioning.

The following table displays the results of the provisioning tests.

Table: Parallels RAS Provisioning Times on AHV

Quantity	Workload	Clone Type	Provisioning Time	Provisioning Method
400	VDI	Full clone	4,742 seconds (79.03 minutes)	Manual
400	VDI	Linked clone	4,699 seconds (78.32 minutes)	Manual
32	RDSH	Full clone	706 seconds (11.77 minutes)	RDSH group autoscale
32	RDSH	Linked clone	687 seconds (11.45 minutes)	RDSH group autoscale

Note: Currently, no autoscale configurations are available for virtual desktop infrastructure workloads in Parallels RAS.

Linear Scalability

This section contains graphs that show the linear scalability of the following test runs:

- Windows 10 tests with one and four nodes with 100 VMs per node
- Windows Server 2022 tests with one and four nodes with 8 VMs per node
- Both tests using linked clone provisioning with Parallels RAS

The results show good EUX scores in all scenarios.

Windows 10 VDI Desktops

Table: Target VM Specifications

Specification	Value
CPU sockets	1
Cores per socket	2
Memory	4 GB
GPU profile	None
Secure boot	True
Virtual TPM	False
Credential guard	False
Operating system	Windows 10
Operating system version	22H2-19045.2364
Desktop Broker Agent version	19.2
Office version	Microsoft Office Professional Plus 2019 x64
Clone type	Linked clone

Table: Test Specifications

Specification	RAS-1n-W10-100VMs	RAS-4n-W10-400VMs
Number of nodes	1	4
Number of VMs	100	400
Number of sessions	100	400

EUX Scores

The EUX base score for RAS-1n-W10-100VMs and RAS-4n-W10-400VMs was 8.6. The following graph shows the EUX scores from the tests. A higher score indicates a better user experience.

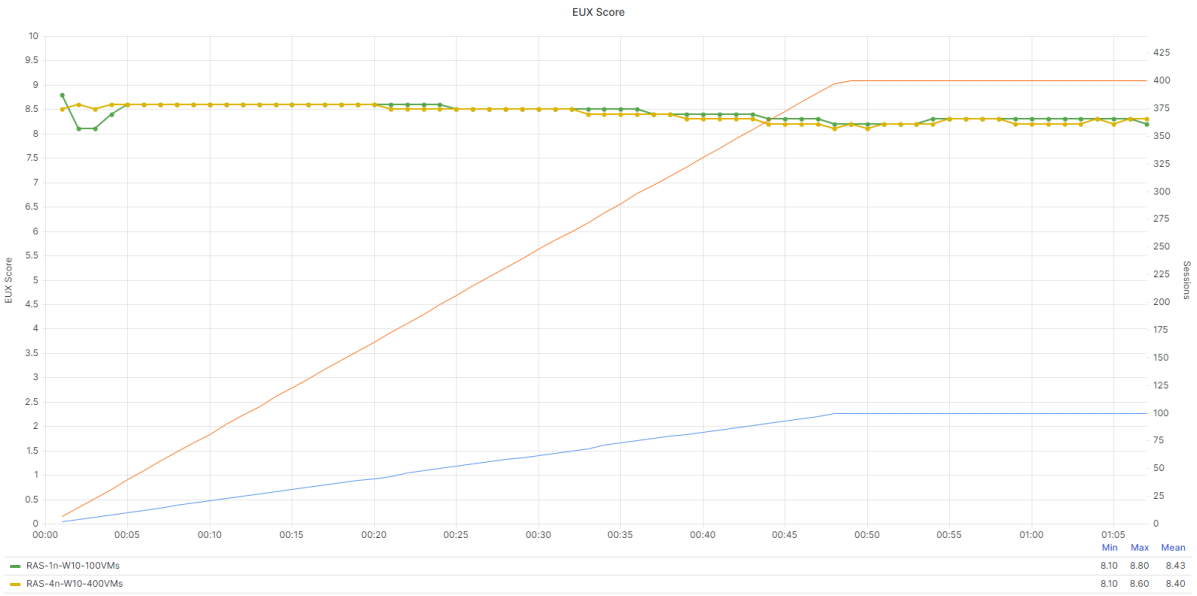

Figure 3: Windows 10 Linear Testing: EUX Score with Sessions

Steady State Scores

The following table details the EUX scores during the steady state.

Table: Windows 10 Linear Testing: EUX Steady State Score

Configuration	EUX Score (Steady State)	Difference
RAS-1n-W10-100VMs	8.3	Highest EUX Score
RAS-4n-W10-400VMs	8.2	–1.2%

Logon Time Performance

The following table captures logon metrics throughout the tests.

Table: Windows 10 Linear Scale: Logon Time Metrics (Averages)

Metric	RAS-1n-W10-100VMs	RAS-4n-W10-400VMs
Average logon time	6.5 seconds	6.7 seconds
Logon rate per minute	2.1	8.33
Total logon time	6.50 seconds	6.71 seconds
Group policies	2.25 seconds	2.26 seconds

Metric	RAS-1n-W10-100VMs	RAS-4n-W10-400VMs
User profile load	0.55 seconds	0.56 seconds
Connection	2.76 seconds	2.93 seconds

The following graph shows the linear scalability of logon times over the test runs. A lower result represents better performance.

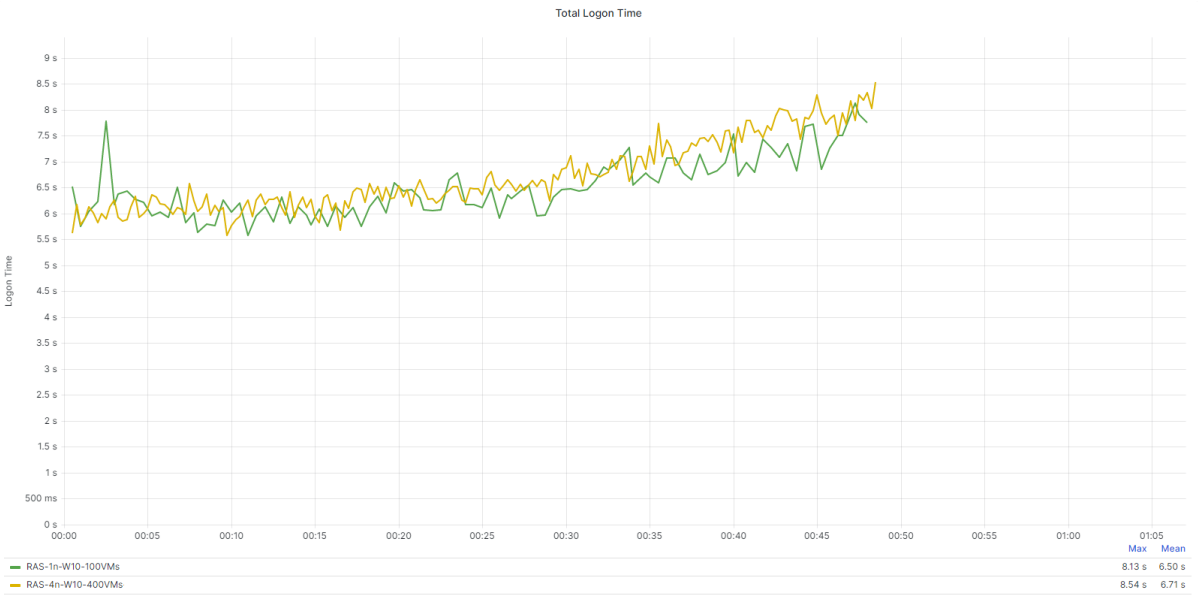

Figure 4: Windows 10 Linear Testing: Total Logon Time

Application Performance

The following table shows the linear scalability detail for application performance.

Table: Windows 10 Linear Testing: Steady State Performance Comparison

Application	RAS-1n-W10-100VMs	RAS-4n-W10-400VMs	Difference
Microsoft Edge: Logon	0.10 seconds	0.10 seconds	Equal
Microsoft Excel: App start time	0.94 seconds	0.97 seconds	3.19% slower
Microsoft Excel: Open document	1.16 seconds	1.19 seconds	2.59% slower

Application	RAS-1n-W10-100VMs	RAS-4n-W10-400VMs	Difference
Microsoft Excel: Open window	0.43 seconds	0.44 seconds	2.33% slower
Microsoft Excel: Save file	0.47 seconds	0.48 seconds	2.13% slower
Microsoft PowerPoint: App start time	0.94 seconds	0.97 seconds	3.19% slower
Microsoft PowerPoint: Open document	1.57 seconds	1.58 seconds	0.64% slower
Microsoft PowerPoint: Open window	0.38 seconds	0.38 seconds	Equal
Microsoft PowerPoint: Save file	0.97 seconds	0.97 seconds	Equal
Microsoft Word: App start time	0.87 seconds	0.88 seconds	1.15% slower
Microsoft Word: Open window	0.42 seconds	0.42 seconds	Equal
Microsoft Word: Open document	1.20 seconds	1.22 seconds	1.67% slower
Microsoft Word: Save file	0.36 seconds	0.37 seconds	2.78% slower

Windows 10 Power Consumption

During the four-node test, we monitored one node's power usage. The following chart shows this host's power usage over the test duration.

Figure 5: Windows 10 Linear Testing: Single Host Power Usage

Windows Server 2022 RDSH

Table: Target VM Specifications

Specification	Value
CPU sockets	1
Cores per socket	8
Memory	42 GB
GPU profile	None
Secure boot	True
Virtual TPM	False
Credential guard	False
Operating system	Windows Server 2022
Operating system version	21H2-20348.1366
Desktop Broker Agent version	19.2
Office version	Microsoft Office Professional Plus 2019 x64
Clone type	Linked clone

Table: Test Specifications

Specification	RAS-1n-SRV2022-8V-120U	RAS-4n-SRV2022-32V-480U
Number of nodes	1	4
Number of VMs	8	32
Number of sessions	120	480

EUX Scores

The EUX base score for RAS-1n-SRV2022-8V-120U and RAS-4n-SRV2022-32V-480U was 8.7. The following graph shows the EUX scores from the tests. A higher score indicates a better user experience.

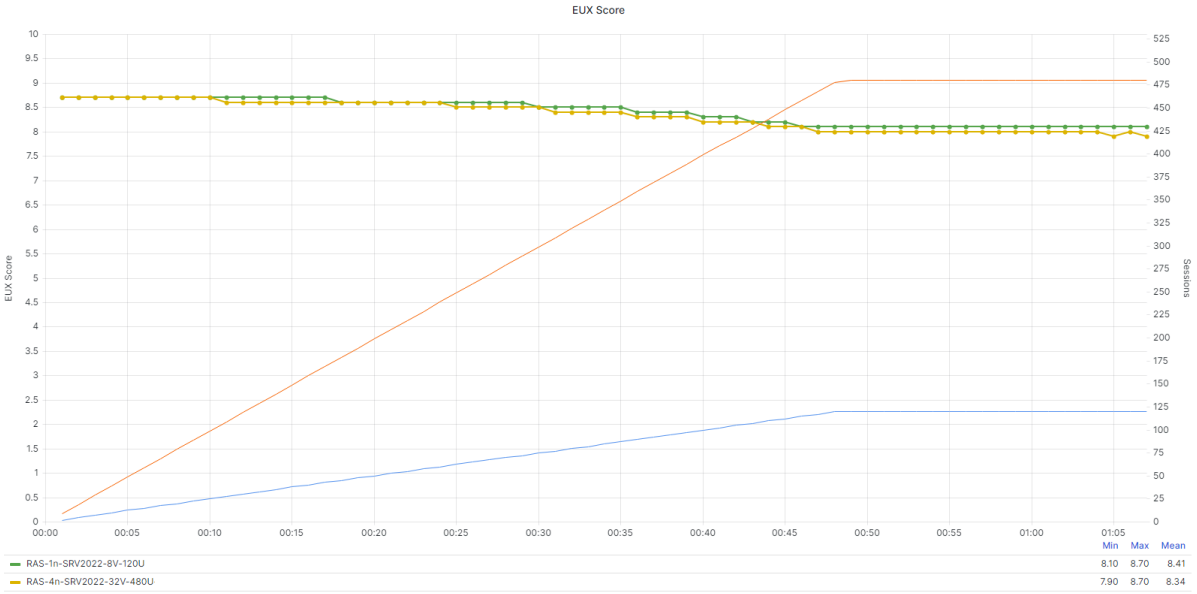

Figure 6: Windows Server 2022 Linear Testing: EUX Score with Sessions

Steady State Scores

The following table and graph detail the EUX scores during the steady state.

Table: Windows Server 2022 Linear Testing: EUX Steady State Score

Configuration	EUX Score (Steady State)	Difference
RAS-1n-SRV2022-8V-120U	8.1	Highest EUX score

Configuration	EUX Score (Steady State)	Difference
RAS-4n-SRV2022-32V-480U	8.0	–1.2%

Logon Time Performance

Table: Windows Server 2022 Linear Testing: Logon Time Metrics (Averages)

Metric	RAS-1n-SRV2022-8V-120U	RAS-4n-SRV2022-32V-480U
Logon time	3.8 seconds	3.9 seconds
Logon rate per minute	2.53	10
Total logon time	3.84 seconds	3.92 seconds
Group policies	1.28 seconds	1.29 seconds
User profile load	0.73 seconds	0.76 seconds
Connection	1.45 seconds	1.48 seconds

The following graphs show the linear scalability of logon times over the test runs. A lower result represents better performance.

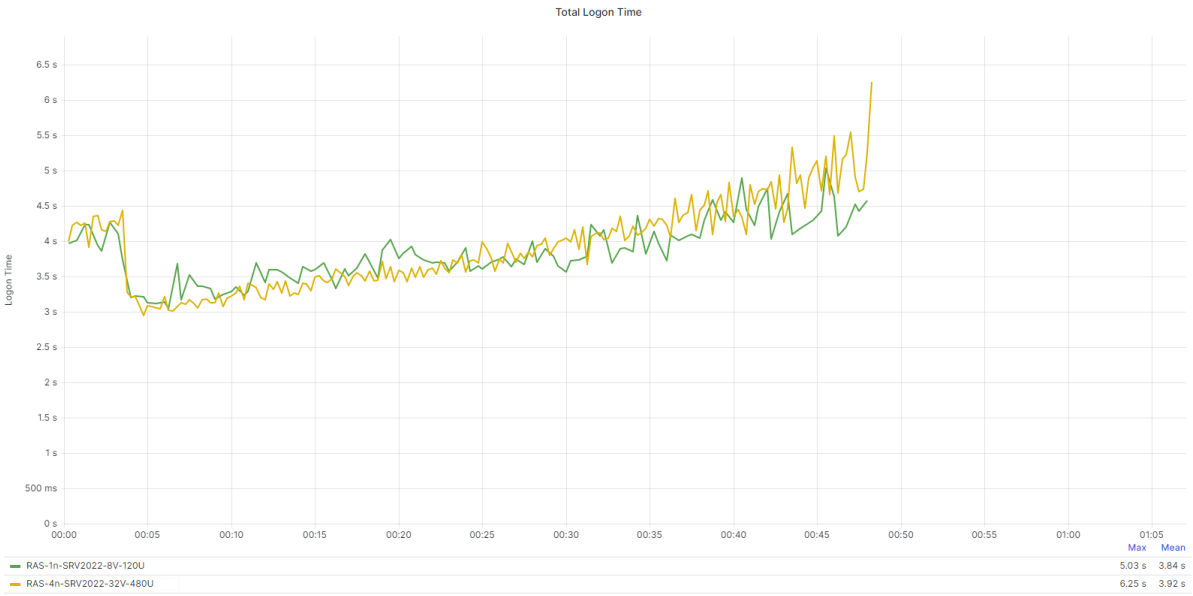

Figure 7: Windows Server 2022 Linear Testing: Total Logon Time

Application Performance

The following table shows the linear scalability detail for application performance.

Table: Windows Server 2022 Linear Testing: Steady State Performance Comparison

Application	RAS-1n-SRV2022-8V-120U	RAS-4n-SRV2022-32V-480U	Difference
Microsoft Excel: App start time	0.97 seconds	1.06 seconds	9.28% slower
Microsoft Excel: Open document	1.07 seconds	1.12 seconds	4.67% slower
Microsoft Excel: Open window	0.53 seconds	0.56 seconds	5.66% slower
Microsoft Excel: Save file	0.50 seconds	0.52 seconds	4.00% slower
Microsoft PowerPoint: App start time	0.89 seconds	0.98 seconds	10.11% slower
Microsoft PowerPoint: Open document	1.74 seconds	1.84 seconds	5.75% slower
Microsoft PowerPoint: Open window	0.49 seconds	0.52 seconds	6.12% slower
Microsoft PowerPoint: Save file	1.20 seconds	1.28 seconds	6.67% slower
Microsoft Word: App start time	1.00 seconds	1.10 seconds	10.00% slower
Microsoft Word: Open window	0.56 seconds	0.59 seconds	5.36% slower
Microsoft Word: Open document	1.04 seconds	1.06 seconds	1.92% slower
Microsoft Word: Save file	0.43 seconds	0.47 seconds	9.30% slower

Windows Server 2022 Power Consumption

During the four-node test, we monitored one node's power usage. The following chart shows this host's power usage over the test duration.

Figure 8: Windows Server 2022 Linear Testing: Single Host Power Usage

Parallels RAS Linked Clone vs. Full Clone Provisioning

In this section, we compare the results of a Login Enterprise test on four nodes, using linked clone and full clone provisioning as the deployment methods.

Windows 10 VDI Desktops

Table: Full Clone and Linked Clone Test Specifics

Specification	Full Clone	Linked Clone
Number of nodes	1	4
Number of VMs	400	400
Number of sessions	400	400

EUX Scores

The EUX base score for the Windows 10 full clone and linked clone provisioning was 8.6. The following graph shows the EUX scores from the tests. A higher score indicates a better user experience.

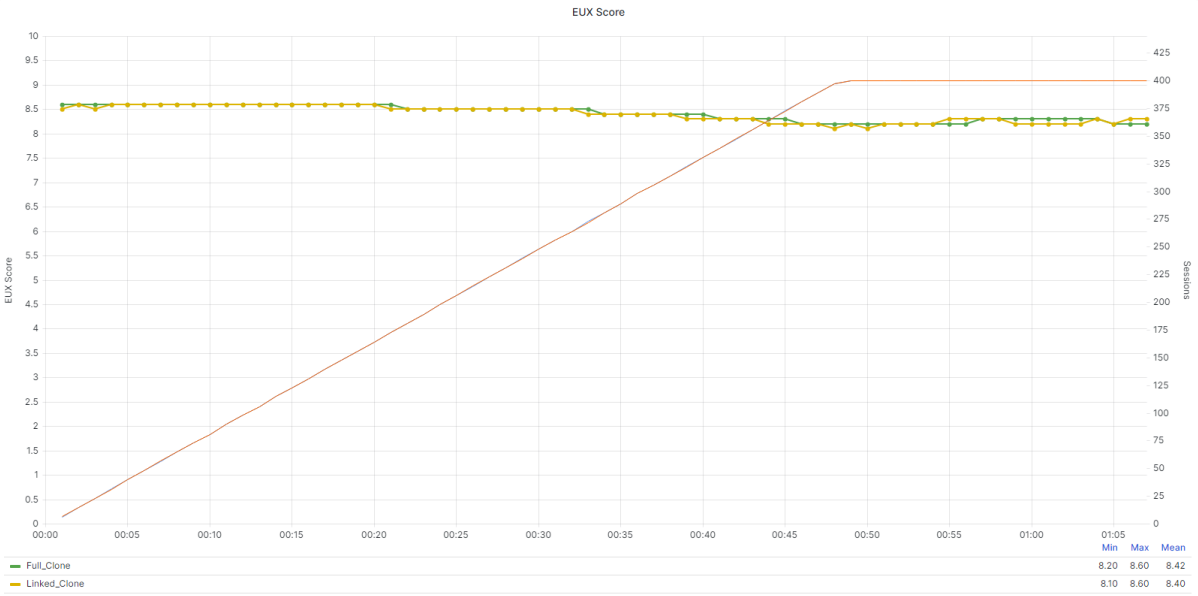

Figure 9: Windows 10 Full vs. Linked Clone Provisioning: EUX Score with Sessions

Steady State Scores

The following table and graph detail the EUX scores during the steady state.

Table: Windows 10 Full vs. Linked Clone Provisioning: EUX Steady State Score

Provisioning Method	EUX Score (Steady State)	Difference
Full clone	8.3	Highest EUX score
Linked clone	8.2	–1.2%

Logon Time Performance

The following table captures logon metrics throughout the tests.

Table: Windows 10 Full vs. Linked Clone Provisioning: Logon Time Metrics (Averages)

Metric	Full Clone	Linked Clone
Logon time	6.7 seconds	6.7 seconds
Logon rate per minute	8.37	8.33
Total logon time	6.72 seconds	6.71 seconds
Group policies	2.27 seconds	2.26 seconds

Metric	Full Clone	Linked Clone
User profile load	0.56 seconds	0.56 seconds
Connection	2.91 seconds	2.93 seconds

The following graph shows logon time performance over the test runs.

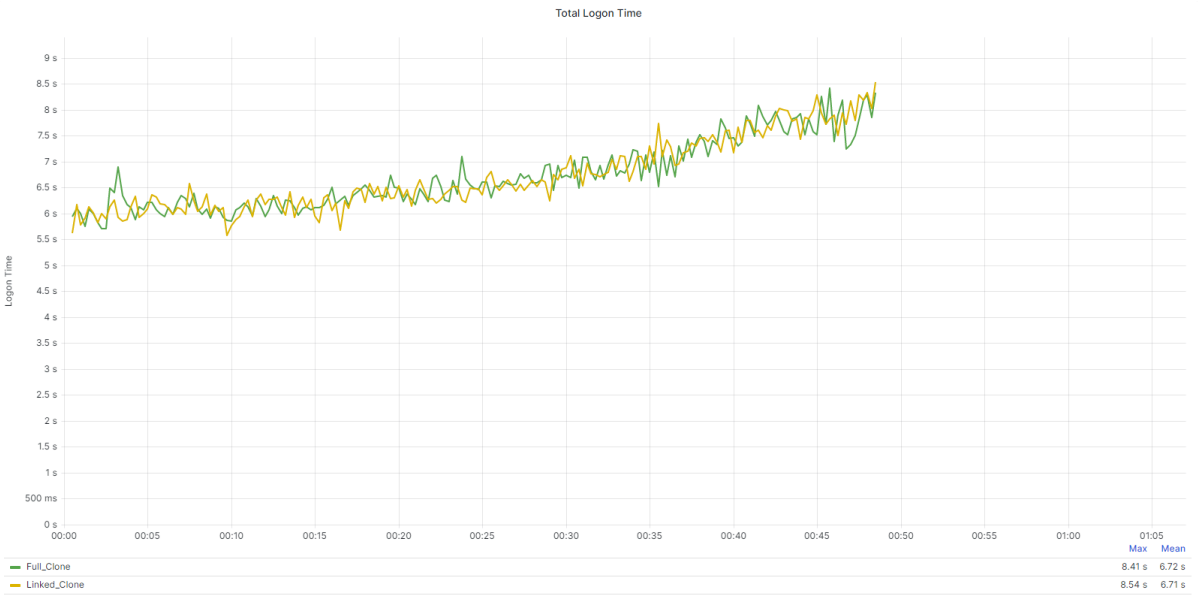

Figure 10: Windows 10 Full vs. Linked Clone Provisioning: Total Logon Time

Application Performance

The following table shows the detail for application performance.

Table: Windows 10 Full vs. Linked Clone Provisioning: Steady State Performance Comparison

Application	Full Clone	Linked Clone	Difference
Microsoft Edge: Logon	0.10 seconds	0.10 seconds	Equal
Microsoft Excel: App start time	0.97 seconds	0.97 seconds	Equal
Microsoft Excel: Open document	1.18 seconds	1.19 seconds	0.85% slower

Application	Full Clone	Linked Clone	Difference
Microsoft Excel: Open window	0.43 seconds	0.44 seconds	2.33% slower
Microsoft Excel: Save file	0.48 seconds	0.48 seconds	Equal
Microsoft PowerPoint: App start time	0.96 seconds	0.97 seconds	1.04% slower
Microsoft PowerPoint: Open document	1.57 seconds	1.58 seconds	0.64% slower
Microsoft PowerPoint: Open window	0.38 seconds	0.38 seconds	Equal
Microsoft PowerPoint: Save file	0.97 seconds	0.97 seconds	Equal
Microsoft Word: App start time	0.88 seconds	0.88 seconds	Equal
Microsoft Word: Open window	0.42 seconds	0.42 seconds	Equal
Microsoft Word: Open document	1.22 seconds	1.22 seconds	Equal
Microsoft Word: Save file	0.36 seconds	0.37 seconds	2.78% slower

Cluster Resources

The following table summarizes the cluster usage during the tests.

Table: Windows 10 Full vs. Linked Clone Provisioning: Cluster Performance Summary During Tests (Averages)

Metric	Full Clone	Linked Clone
CPU usage	61.8%	62.1%
CPU ready time	0.46%	0.51%
Memory usage	33.3%	33.3%
Controller IOPS	12,119	12,231
Controller write IOPS	5,774	5,774

Metric	Full Clone	Linked Clone
Controller read IOPS	6,345	6,407
Controller latency	< 1 ms	< 1 ms
Controller write I/O latency	< 1 ms	< 1 ms
Controller read I/O latency	< 1 ms	< 1 ms

The following table and graph compare the cluster CPU usage during the tests.

Table: Windows 10 Full vs. Linked Clone Provisioning: Cluster CPU Steady State

Provisioning Method	Cluster CPU (Steady State)	Difference
Full clone	79.7%	Lowest CPU value
Linked clone	79.8%	+0.1%

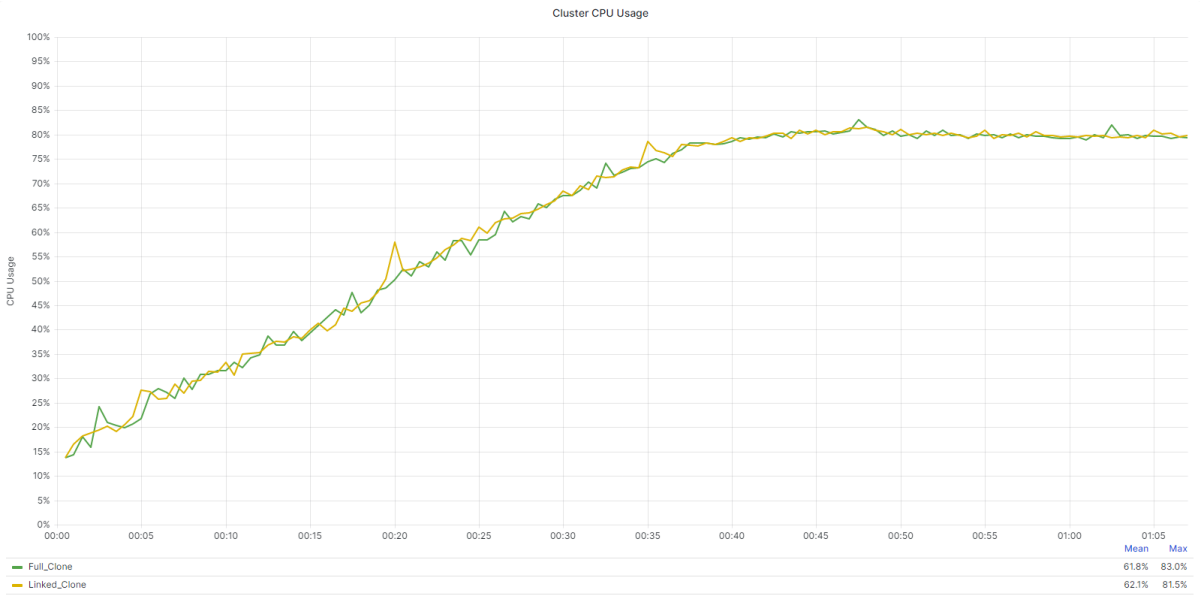

Figure 11: Windows 10 Full vs. Linked Clone Provisioning: Cluster CPU Usage Overall

Note: The previous graph shows the overall test results, whereas the previous table shows the steady state results.

The following graph shows the difference in the cluster controller IOPS.

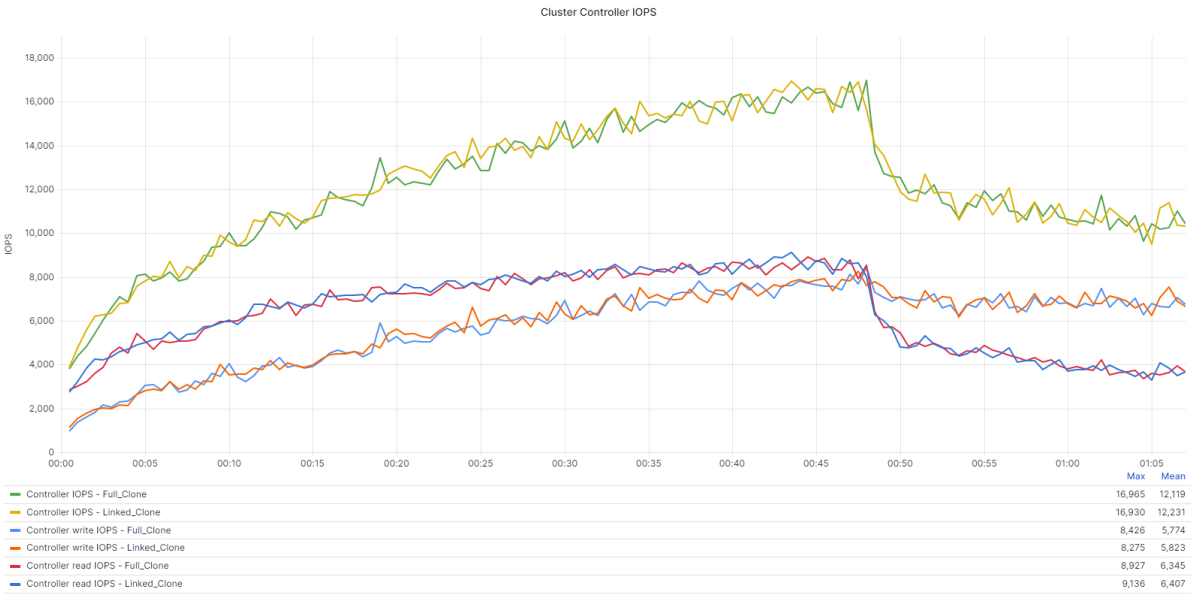

Figure 12: Windows 10 Full vs. Linked Clone Provisioning: Cluster Controller IOPS

Note: As mentioned previously, the I/O load in the current release of Login Enterprise EUX 2023 is much higher than is realistic for a knowledge worker. However, this result shows that Nutanix storage can handle the load.

Windows Server 2022 RDSH

Table: Windows Server 2022 RDSH Test Specifics

Specification	Value
Single node test	False
Number of VMs	32
Number of sessions	480

EUX Scores

The EUX base score for the Windows Server 2022 full clone and linked clone provisioning was 8.7. The following graph shows the EUX scores from the tests. A higher score indicates a better user experience.

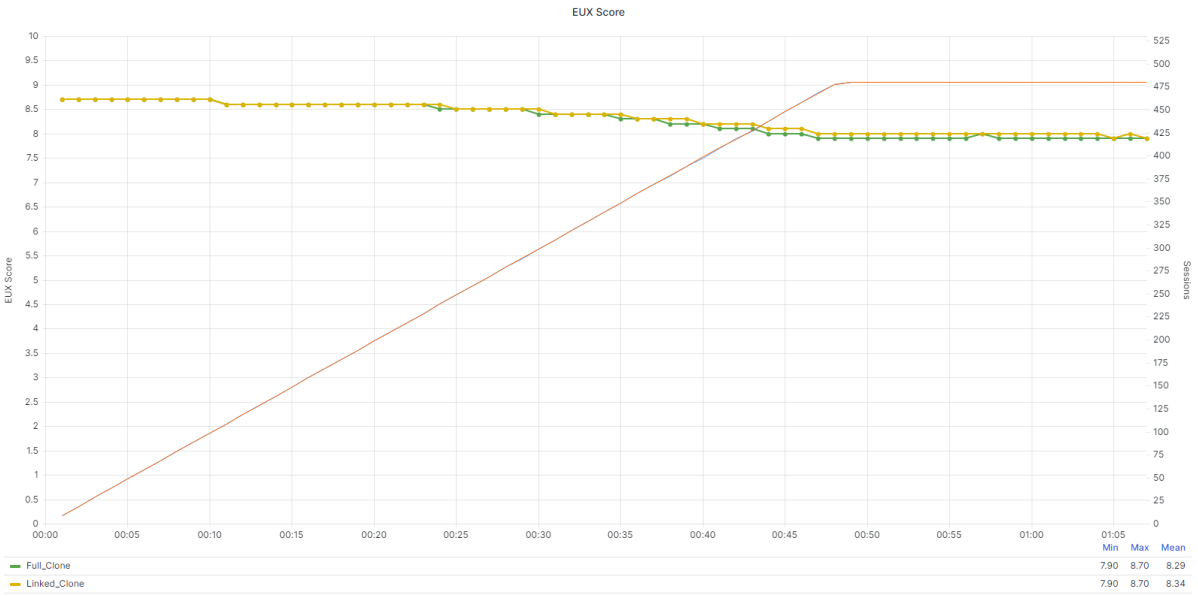

Figure 13: Windows Server 2022 Full vs. Linked Clone Provisioning: EUX Score with Sessions

Steady State Scores

The following table and graph detail the EUX scores during the steady state.

Table: Windows Server 2022 Full vs. Linked Clone Provisioning: EUX Steady State Score

Provisioning Method	EUX Score (Steady State)	Difference
Full clone	7.9	–1.2%
Linked clone	8.0	Highest EUX score

Logon Time Performance

The following table captures logon metrics throughout the tests.

Table: Windows Server 2022 Full vs. Linked Clone Provisioning: Logon Time Metrics (Averages)

Metric	Full Clone	Linked Clone
Logon time	4.0 seconds	3.9 seconds
Logon rate per minute	10	10

Metric	Full Clone	Linked Clone
Total logon time	4.04 seconds	3.92 seconds
Group policies	1.34 seconds	1.29 seconds
User profile load	0.79 seconds	0.76 seconds
Connection	1.49 seconds	1.48 seconds

The following graph shows the logon time performance over the test runs.

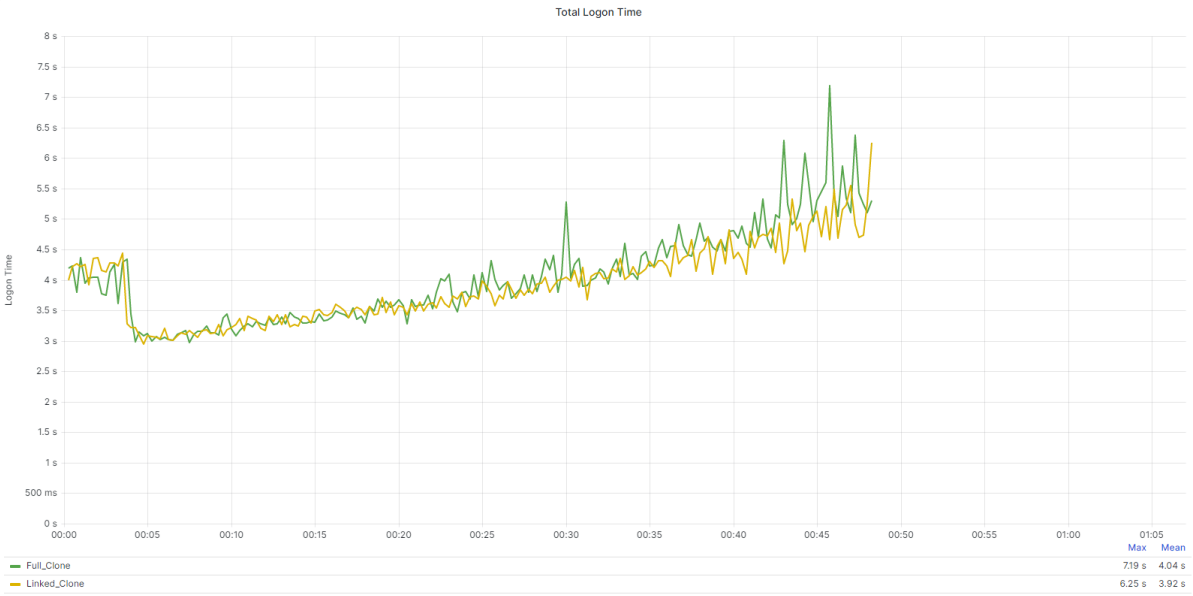

Figure 14: Windows Server 2022 Full vs. Linked Clone Provisioning: Total Logon Time

Application Performance

The following table shows the detail for application performance.

Table: Windows Server 2022 Full vs. Linked Clone Provisioning: Steady State Performance Comparison

Application	Full Clone	Linked Clone	Difference
Microsoft Excel: App start time	1.11 seconds	1.06 seconds	4.50% faster
Microsoft Excel: Open document	1.16 seconds	1.12 seconds	3.45% faster

Application	Full Clone	Linked Clone	Difference
Microsoft Excel: Open window	0.59 seconds	0.56 seconds	5.08% faster
Microsoft Excel: Save file	0.55 seconds	0.52 seconds	5.45% faster
Microsoft PowerPoint: App start time	1.05 seconds	0.98 seconds	6.67% faster
Microsoft PowerPoint: Open document	2.16 seconds	1.84 seconds	14.81% faster
Microsoft PowerPoint: Open window	0.56 seconds	0.52 seconds	7.14% faster
Microsoft PowerPoint: Save file	1.39 seconds	1.28 seconds	7.91% faster
Microsoft Word: App start time	1.16 seconds	1.10 seconds	5.17% faster
Microsoft Word: Open window	0.62 seconds	0.59 seconds	4.84% faster
Microsoft Word: Open document	1.15 seconds	1.06 seconds	7.83% faster
Microsoft Word: Save file	0.50 seconds	0.47 seconds	6.00% faster

Cluster Resources

The following table summarizes the cluster usage during the tests.

Table: Windows Server 2022 Full vs. Linked Clone Provisioning: Cluster Performance Summary During Tests (Averages)

Metric	Full Clone	Linked Clone
CPU usage	59.3%	57.0%
CPU ready time	0.49%	0.37%
Memory usage	28.2%	28.2%
Controller IOPS	8,557	8,647
Controller write IOPS	5,486	5,312

Metric	Full Clone	Linked Clone
Controller read IOPS	3,071	3,334
Controller latency	< 1 ms	< 1 ms
Controller write I/O latency	1.22 ms	1.23 ms
Controller read I/O latency	< 1 ms	< 1 ms

The following table and graph compare the cluster CPU usage during the test.

Table: Windows Server 2022 Full vs. Linked Clone Provisioning: Cluster CPU Steady State

Provisioning Method	Cluster CPU (Steady State)	Difference
Full clone	78.9%	+1.1%
Linked clone	77.8%	Lowest CPU value

Figure 15: Windows Server 2022 Full vs. Linked Clone Provisioning: Cluster CPU Usage Overall

Note: The previous graph shows the overall test result, whereas the previous table shows the steady state results.

The following graph shows the difference in the cluster controller IOPS.

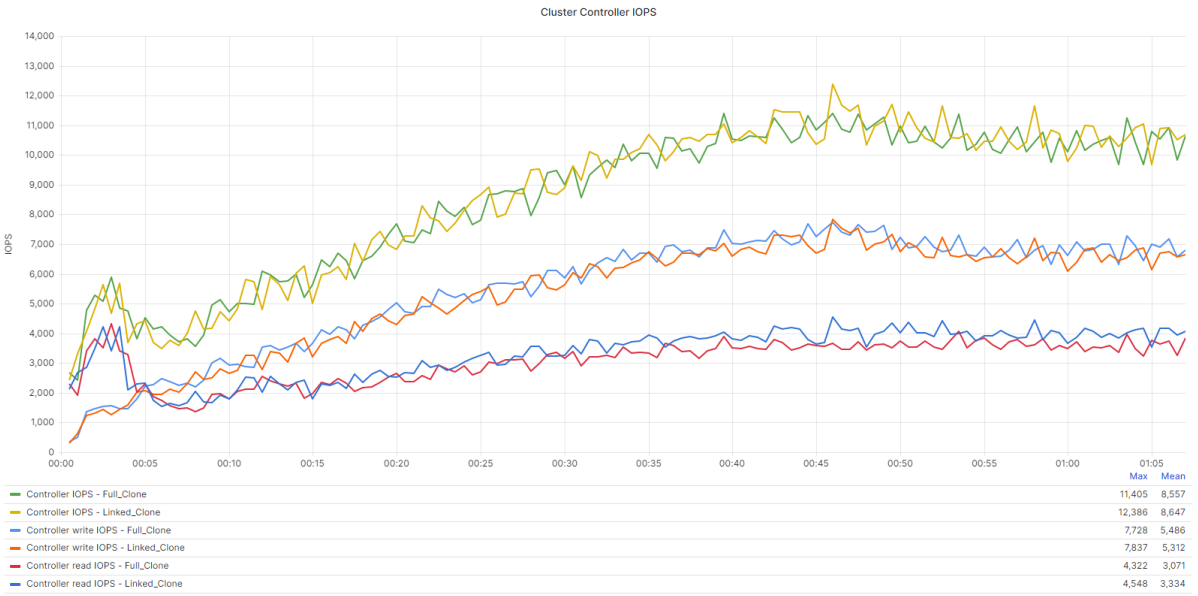

Figure 16: Windows Server 2022 Full vs. Linked Clone Provisioning: Cluster Controller IOPS

Note: As discussed earlier, the I/O load in the current release of Login Enterprise EUX 2023 is much higher than is realistic for a knowledge worker. However, this result shows that Nutanix storage can handle the load.

RAS Linked Clone vs. Full Clone Provisioning Summary

Our testing shows a negligible performance difference between full and linked clone provisioning when deployed on Nutanix AHV, which makes sense because the only difference is that full clones use the base template as the clone source and linked clones use a snapshot of the template.

7. Conclusion

The Parallels Remote Application Server (RAS) and Nutanix solution provides a single high-density platform for virtual desktop delivery. With this modular, linear scaling approach, you can grow Parallels RAS deployments easily. Localized and distributed caching and integrated disaster recovery enable quick deployments and simplify day-to-day operations. Robust self-healing and multistorage controllers deliver high availability in the face of failure or rolling upgrades.

On Nutanix, available host CPU resources drive Parallels RAS user density, rather than any I/O or resource bottlenecks for virtual desktops. Login Enterprise test results showed densities at 100 sessions per Nutanix node for VDI and 120 sessions per Nutanix node running RDSHs while running the Login Enterprise knowledge worker workload.

Our results indicated a negligible difference between Parallels RAS linked and full cloning capability. Additionally, if using Parallels RAS provisioning capabilities to linearly scale the solution, the system maintains a good Login Enterprise EUX Score regardless of how many nodes you have.

8. Appendix

Hardware Configuration

Storage and compute:

- Nutanix NX-3155G-G8
- Per-node specs:
  - CPU: 2 × Intel Xeon Gold 6354
  - Memory: 0.98 TiB
  - Disk: 6 × 1.55 TB SSD

Network:

- Arista 7050Q: L3 spine
- Arista 7050S: L2 leaf

Software Configuration

Nutanix:

- AOS 6.5.3.6
- AHV build 20220304.423
- CVM: 12 vCPU, 32 GB of memory

Parallels RAS:

- Parallels RAS 19.2
- Parallels RAS HALB Appliance 19.1

Virtual Desktop and RDSH Servers:

- Windows 10, version 22H2: build 19045.2364

- Windows Server 2022: build 20348.1366

Login Enterprise Configuration Changes

We made the following changes to the Login Enterprise configuration prior to running the tests to better reflect a VDI workload.

Login Enterprise Actions

```
{
    "diskmydocs": {
      "App": "diskspeed",
      "Arguments": "folder=\"{myDocs}\" blockSize=4k bufferSize=32K
 writeMask=0x5555 cachePct=97 latencyPct=99 threads=1 duration=250",
      "Label": "MyDocuments"
    },
    "diskappdata": {
      "App": "diskspeed",
      "Arguments": "folder=\"{appData}\"  blockSize=50k bufferSize=4K
 writeMask=0x5555 cachePct=97 latencyPct=99 threads=1 duration=250",
      "Label": "LocalAppdata"
    },
    "cpuspeed": {
      "App": "cpuspeed",
      "Arguments": "d=250 t=4",
      "Label": "CPU"
    },
    "highcompression": {
      "App": "compressionspeed",
      "Arguments": "folder=\"{appData}\" cachePct=95 writePct=100
 duration=250 threads=1 -high",
      "Label": "Compression"
    },
    "fastcompression": {
      "App": "compressionspeed",
      "Arguments": "folder=\"{appData}\" cachePct=95 writePct=100
 duration=250 threads=1",
      "Label": "CachedHighCompression"
    },
    "appspeed": {
      "App": "appspeed",
      "Arguments": "folder=\"{appData}\" duration=10000
 launchtimestamp={launchTimestamp}",
      "Label": "App"
    }
}
```

Login Enterprise Tuning

```
 {
  "PerformancePenalty": 3.0,
  "BucketSizeInMinutes": 5,
  "NumSamplesForBaseline": 5,
  "CapacityRollingAverageSize": 3,
```



```
  "MaxBaselineForCapacity": 4000,
  "CapacityTrigger": "<80%",
  "SteadyStateCooldownWindow": 5,
  "BaselineScoreWindowSize": 5,
  "Actions": {
    "DiskMyDocs": {
      "Weight": 0,
      "NominalValue": 8500,
      "CapacityTrigger": "<25%"
    },
    "DiskMyDocsLatency": {
      "Weight": 0,
      "NominalValue": 1200,
      "CapacityTrigger": "<5%"
    },
    "DiskAppData": {
      "Weight": 0,
      "NominalValue": 14000,
      "CapacityTrigger": "<25%"
    },
    "DiskAppDataLatency": {
      "Weight": 0,
      "NominalValue": 1700,
      "CapacityTrigger": "<5%"
    },
    "CpuSpeed": {
      "Weight": 0,
      "NominalValue": 50000,
      "CapacityTrigger": "<55%"
    },
    "HighCompression": {
      "Weight": 1,
      "NominalValue": 2500,
      "CapacityTrigger": "<5%"
    },
    "FastCompression": {
      "Weight": 1,
      "NominalValue": 2500,
      "CapacityTrigger": "<5%"
    },
    "AppSpeed": {
      "Weight": 6,
      "NominalValue": 2500,
      "CapacityTrigger": "<45%"
    },
    "AppSpeedUserInput": {
      "Weight": 1,
      "NominalValue": 500,
      "CapacityTrigger": "<35%"
    }
  }
}
```

References

1. Login Enterprise

2. Login Enterprise EUX Score
3. Login Enterprise Workload Templates
4. Parallels RAS 19 Administrators Guide
5. Parallels RAS Documentation

About Nutanix

Nutanix offers a single platform to run all your apps and data across multiple clouds while simplifying operations and reducing complexity. Trusted by companies worldwide, Nutanix powers hybrid multicloud environments efficiently and cost effectively. This enables companies to focus on successful business outcomes and new innovations. Learn more at Nutanix.com.

List of Figures