

# Virtualized Desktop Infrastructure with Dell EMC PowerMax and VMware Horizon Enterprise 7

October 2019

H17969

## PowerMax Engineering White Paper

### Abstract

This white paper provides details on the performance and space saving values of deploying VMware Horizon virtual desktop infrastructure (VDI) on Dell EMC PowerMax storage systems.

## Copyright

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © [Year or Years] Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Intel, the Intel logo, the Intel Inside logo and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Other trademarks may be trademarks of their respective owners. Published in the USA October 2019 White Paper H17969.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.



# Contents

- Executive summary ..... 4**
- VMware Horizon overview ..... 5**
- PowerMax product overview ..... 6**
- Performance and data reduction assessment ..... 9**

## Executive summary

Virtual Desktop Infrastructure (VDI) provides a way to manage user desktop environments with fewer resources, increased manageability, enhanced security, and greater performance. With VDI, a single image can be used to propagate thousands of desktops. Because this reusable image is built in accordance with company security standards and an approved software stack, it provides IT professionals with a tremendous level of control over their virtualized computing environment, making it far more secure and reliable.

VMware Horizon 7 lets you create, deploy, and update desktop images. It is a centralized desktop virtualization solution that helps you deliver virtualized desktop services and applications to end users from centralized VMware vSphere servers.

The tests described in this paper demonstrate the ability of the Dell EMC PowerMax storage system to support VDI deployment with VMware vSphere and Horizon 7. To perform the tests, we used LoginVSI, the industry standard for load-testing VDI solutions, to generate load and analyze the system's capabilities. In addition, we collected and reviewed storage performance metrics and data reduction efficiencies.

The test environment hardware consisted of a single-brick PowerMax 8000 storage array and twelve Dell PowerEdge R740 servers as compute nodes for desktop deployment. Other servers were used for management and monitoring. We used this environment to generate and support 400 full-clone desktops, or 600 instant-clone desktops running the LoginVSI "Knowledge Worker" workload, which is a well-balanced user type. LoginVSI output showed an exceptionally good base performance ("VSIbase") and a very high estimated maximum performance ("VSImax"). It also showed that the environment could have scaled to support more than twice the number of compute nodes, demonstrating the ability of PowerMax to accommodate intensive workloads.

Before LoginVSI workload phase begins, all the desktops are restarted, creating a 'boot-storm'. Boot-storm is an important aspect of VDI deployment testing, demonstrating what might happen if all the desktops need to start simultaneously after planned or unplanned downtime.

The tests demonstrated high IOPS (over 110,000), high bandwidth (near 6 GB/s), and low latencies (1.4 ms read response time and 1.8 ms write response time) during the boot-storm phase with full-clone desktops. The tests demonstrated even lower latencies (0.4 ms read response time and 0.2 ms write response time) during the LoginVSI workload phase with full-clone desktops. In comparison, instant clones demonstrated sub-millisecond latencies during both the boot-storm and workload phases, because instant clones in general are more efficient.

Instant clones share a virtual disk and therefore provide their own inherent deduplication value; however, PowerMax compression still provides strong data reduction benefits. With full clones, each desktop gets their own virtual disk. In this case, PowerMax compression and deduplication values are both of great benefit for data reduction.

While deploying 400 full-clone desktops, PowerMax provided a data reduction savings of 97% per desktop by reducing the average desktop storage size from 13 GB raw capacity to 0.4 GB actual allocated capacity (12.6 GB saved capacity / 13 GB raw capacity = 97%). These savings were noticed immediately after the desktop initial deployment phase. After

completing the LoginVSI workload phase, the savings were 94% per desktop as the average desktop storage size was reduced from 13 GB raw capacity to 0.84 GB actual allocated capacity (12.16 GB saved capacity / 13 GB raw capacity = 94%). Capacity saving benefits of the workload phase will vary with the amount of data that is changed.

These tests prove that PowerMax is an excellent storage infrastructure choice for deploying the VMware Horizon VDI environment, maintaining high performance, low latencies, and strong data reduction efficiency.

## Audience

This white paper is intended for storage administrators, system administrators, and system architects who are responsible for implementing a virtual desktop infrastructure (VDI) in enterprise environments.

## VMware Horizon overview

VMware Horizon 7 is a centralized desktop virtualization solution that enables organizations to deliver virtualized desktop services and applications to end users from centralized VMware vSphere servers.

Horizon 7 operates with VMware vCenter Server to create desktops from virtual machines that are running on VMware ESXi™ compute nodes and deploy these desktops to end users. Horizon 7 can also be installed on RDS hosts to deploy desktops and applications to end users. Horizon 7 uses an existing Active Directory infrastructure for user authentication and management.

After the creation of a desktop or application, authorized end users can use Web-based or locally installed client software to securely connect to centralized virtual machines, back-end physical systems, or RDS hosts.

## Horizon 7 Components

Horizon 7 consists of the following major components:

- Horizon Connection Server - A software service that acts as a broker for client connections by authenticating and then directing incoming user requests to the appropriate virtual machine, physical system, or RDS host.
- Horizon Agent - A software service that is installed on all guest virtual machines, physical systems, or RDS hosts to allow them to be managed by Horizon 7. Horizon Agent provides features such as connection monitoring, virtual printing, USB support, and single sign-on.
- Horizon Client - A software application that communicates with Connection Server to enable users to connect to their desktops.
- Horizon Administrator - A web application that enables Horizon 7 administrators to configure Connection Server, deploy desktop and application pools, manage machines, control user authentication, initiate and examine system events, and perform analytical activities.

- vCenter Server - A server that acts as a central administrator for ESXi hosts that are connected on a network. A vCenter Server instance provides the central point for configuring, provisioning, and managing virtual machines in the datacenter.

### How the components fit together

End users start Horizon Client to log in to Horizon Connection Server. This server, which integrates with Windows Active Directory, provides access to remote desktops hosted on a VMware vSphere server, a physical PC, or a Microsoft RDS host. Horizon Client also provides access to published applications on a Microsoft RDS host.





Figure 1. VMware Horizon architecture<sup>1</sup>

## PowerMax product overview

The Dell EMC PowerMax family consists of two models, as shown in the following figure:

- PowerMax 2000 – Designed to provide customers with efficiency and maximum flexibility in a 20U footprint.
- PowerMax 8000 – Designed for massive scale and performance, all within a two-floor-tile footprint.

<sup>1</sup> Image source: <https://techzone.vmware.com/resource/how-does-horizon-7-work>.




**Figure 2. PowerMax 2000 and PowerMax 8000**

Both PowerMax storage arrays have, at their foundation, the trusted Dynamic Virtual Matrix architecture and a new version of HYPERMAX OS management software—PowerMaxOS 5978—that has been rewritten for the non-volatile memory express (NVMe) platform. PowerMaxOS can run natively on both PowerMax storage arrays and on legacy VMAX All Flash systems as an upgrade. PowerMax storage arrays are true all-flash arrays, specifically targeted to meet the storage capacity and performance requirements of the all-flash enterprise data center.

**PowerMax architecture**

PowerMax configurations consist of modular building blocks called PowerMax bricks, as shown in the following figure. The modular brick architecture reduces complexity and facilitates simpler system configuration and deployment.



**Figure 3. PowerMax 2000 and PowerMax 8000 bricks**

The initial PowerMax 2000 or 8000 brick includes a single engine consisting of two directors, two system power supplies (SPS), and two 24-slot 2.5” NVMe Drive Array Enclosures (DAEs).

PowerMax storage arrays can scale up and scale out. Customers can scale up by adding Flash Capacity Packs. For PowerMax 2000, Flash Capacity Packs are available in increments of 13.2 TBU for RAID5 (7+1) or 11.3 TBU for RAID5 (3+1). For PowerMax 8000, Flash Capacity Packs are available in increments of 13.2 TBU.

PowerMax storage arrays scale out by aggregating up to two bricks for the PowerMax 2000 storage array, and up to eight bricks for the PowerMax 8000 storage array in a single system with fully shared connectivity, processing power, and linear scalability.

The PowerMax update of Q3, 2019 includes several new capabilities:

- Storage Class Memory (SCM) drives powered by dual port Intel® Optane™ drives are used as persistent storage alongside NAND SSD flash drives in the same DAEs.
- FC-NVMe (NVMe over Fiber Channels Fabrics) support delivers end-to-end NVMe connectivity from host to flash drives.
- 32Gb/s FC and FC-NVMe front-end modules provide higher front-end port bandwidth.
- Other improvements include increased IOPS, increased bandwidth, and reduced latencies.

For more information about the PowerMax architecture and features, see these documents:

- [Dell EMC PowerMax Family Overview White Paper](#)
- [Dell EMC PowerMax Family Data Sheet](#)
- [Dell EMC PowerMax Family Specification Sheet](#)

## PowerMax compression and deduplication

The PowerMax storage array uses a strategy that is targeted to provide the best data reduction without compromising performance. The PowerMax Adaptive Compression Engine (ACE) combines these attributes:

- **Hardware acceleration** – Each PowerMax engine is configured with two hardware compression modules (one per director) that handle data compression and decompression. These hardware modules are also capable of generating Hash IDs that enable deduplication and are more powerful than the modules used with VMAX All Flash arrays.
- **Optimized data placement** – Based on compressibility, the data is allocated in different compression pools that provide a compression ratio (CR) from 1:1 (128 KB pool) up to 16:1 (8 KB pool) and are spread across the PowerMax back end for best performance. The pools are dynamically added or deleted based on need.
- **Activity Based Compression (ABC)** – Typically, the most recent data is the most active, creating an “access skew.” ABC relies on that skew to prevent constant compression and decompression of data extents that are frequently accessed. The ABC function marks the busiest 20 percent of all allocated data extents in the

system, and skips the compression workflow for those extents. Data extents that are highly active remain uncompressed, even if their storage group has compression enabled. As the data extents become less active, they are automatically compressed while newly active extents are part of the “hottest” 20 percent (if enough free storage capacity is available).

- **Fine grain data packing** – When PowerMax compresses data, each 128 K track is split into four 32 K buffers. All buffers are compressed in parallel. The total of the four buffers results in the final compressed size and determines in which compression pool the data is allocated. Included in this process is a zero reclaim function that prevents the allocation of buffers with all zeros and no actual data. For a small size read or write, only the necessary buffers participate, not all four.
- **Extended data compression (EDC)** – Data that is already compressed automatically goes through additional, more powerful compression if it remains untouched for over 30 days, increasing storage efficiency even more.

Additionally, note the following:

- Compression is enabled or disabled at a storage group level for ease of management. Generally, most databases can benefit from storage compression. Customers might decide not to enable compression if the data is fully encrypted or if the storage group contains data that is continuously overwritten.
- When compression is enabled, all new writes benefit from inline compression. If the storage group already contains data when compression is enabled, it goes through background compression with low priority (relative to application I/Os).

In addition to more powerful hardware compression modules, the PowerMax storage array also introduces data deduplication (dedupe) capability. PowerMax deduplication is automatically enabled or disabled when compression is enabled or disabled (compression and deduplication cannot be managed separately).

For more information about the PowerMax Adaptive Compression Engine, see [Data Reduction with Dell EMC PowerMax](#).

## VMware APIs for Array Integration (VAAI)

VAAI was first introduced by VMware as an improvement to host-based VM cloning. It permits, for example, the offloading of the VM cloning process to the storage array, making it much more efficient: instead of the host copying all blocks of a VM from the array and back to it, VMware passes the job to the array which uses internal mechanisms to save both host and network resources and reduce storage latency. The XCOPY (extended copy) command offloads the operation to the storage array. PowerMax supports all the VAAI primitives, including Full Copy, Write Same, Atomic Test and Set, and Unmap.

## Performance and data reduction assessment

### Overview

#### Test goals and sizing considerations

The assessment tests targeted the following goals:

- Demonstrate that PowerMax can support large scale VDI deployments

- Demonstrate that PowerMax can provide suitable data reduction ratios critical for full-clone VDI deployments

During the tests, we quickly realized that the number and type of servers used to deploy desktops were the key factor in achieving higher desktop scale. Although the PowerMax 8000 storage system was configured with only a single brick (engine)—the smallest configuration for this model—it performed extremely well and could have supported additional load if we had added more servers.

In the initial test environment, we used eight Dell PowerEdge 740 servers with 128 GB of RAM each for the desktop users. We later added four more PowerEdge 740 servers with 256 GB of RAM each (see hardware configuration in [Table 1](#)).

With our testing infrastructure we were able to test a total of 400 full-clone desktop users, and 600 instant-clone desktop users. The PowerMax storage array was never saturated and could have supported additional desktop scale or alternatively, other workloads.

Each server is limited by its CPU and memory utilization in terms of how many desktop users it can support. When we tested Horizon full-clones, each of the servers with 128 GB RAM could support 30 desktop users per server at 80% CPU and memory utilization, and each of the servers with 256 GB RAM could support 40 desktop users per server at 80% CPU and memory utilization. When Horizon instant-clones were tested, each of the servers with 128 GB RAM could support 45 desktop users at 80% CPU and memory utilization, and each of the servers with 256 GB RAM could support 60 desktop users at 80% CPU and memory utilization. It should be noted that a higher user density can be reached per server with additional RAM and if the CPU is allowed to reach 100% utilization.

Clone type	RAM	Number of Desktop Users	CPU and Memory Utilization
Full Clone	128 G	30	80%
Full Clone	256 G	40	80%
Instant Clone	128 G	45	80%
Instant Clone	256 G	60	80%

We also measured IOPS and bandwidth demand on the storage during the ‘boot-storm’ phase, in which all the desktops were restarted simultaneously. During the boot storm, we noted the maximum IOPS and bandwidth generated, and the low response time that the PowerMax storage maintained. For example, the boot-storm phase of the 400 full-clone desktops generated over 110,000 IOPS and bandwidth near 6 GB/sec. PowerMax maintained an average read latency of just 1.4 ms during this test. This is important because one of the critical aspects of VDI is the ability of the business to restart operations quickly after planned or unplanned downtime.

### Full clone and instant clone considerations

VMware Horizon full clones are complete and independent copies of a virtual machine and operate separately from the original parent VM. Full clones do not share virtual disks, and therefore have higher demand for storage capacity and performance compared to instant clones. In addition, unlike instant clones, where the image is freshly created for

each user logging in and destroyed when the user logs out, a full-clone image is not destroyed and can be reused by the same or other users.

Instant clones share a virtual disk replica of a parent VM and therefore consume much less storage than a full clone. Only the modified data is unique to each instant-clone desktop.

Instant clones are more efficient than linked clones because they utilize vmFork technology. This technology allows memory sharing of unmodified data from the master image, and instant-clones can be created from a running VM instead of from a powered-off one. Using instant clones provides improved provisioning, updates, and memory utilization. Because instant clones are an improvement over linked clones, linked clones are not covered in this white paper.

### VMware VAAI

VAAI integration with PowerMax is important because, during our testing, the VM cloning process for both full and instant clones was offloaded to the storage array. Desktops were provisioned more quickly because of the more available compute resources.

## Test environment

We tested the hardware and software configuration described below to assess the performance and space-saving benefits of running a VDI environment on Dell PowerEdge servers with the PowerMax storage system.

### Hardware

The test environment included a total of twelve Dell PowerEdge R740 servers for hosting the desktops. In addition, a mixture of other Intel-based servers were used for the VDI management infrastructure. All the physical servers were running ESXi 6.7, with Horizon management and LoginVSI software components consisting of virtual machines. [Table 1](#) shows the main hardware components used.

**Table 1. Hardware components**

Device	Quantity	Configuration	Description
PowerMax 8000	1	Single engine 5978.221.221 8 FA ports 4 x 12 TB devices used for 4 x VMFS datastores for desktops	Single engine
ESXi Hosts (LoginVSI launchers)	4	20 CPUs x 2.8 GHz (Intel® Xeon® E5-2680 v2) 96 GB memory	LoginVSI test launchers
ESXi Hosts (LoginVSI and management hosts)	5	24 CPUs x 2.6 GHz (Intel® Xeon® E5-2690 v3) 128 GB memory	Used for LoginVSI and Horizon management tasks

Device	Quantity	Configuration	Description
ESXi Hosts (VDI compute nodes)	8	Dell R740 48 CPUs x 2.6 GHz (Intel® Xeon® gold 6126) 128 GB memory	Used for desktop provisioning
ESXi Hosts (VDI compute nodes)	4	Dell R740 48 CPUs x 2.6 GHz (Intel® Xeon® gold 6126) 256 GB memory	Used for desktop provisioning

To keep the configuration simple while allowing for more I/O parallelism and scale, we used four VMFS datastores, each containing a single 12 TB thin device. Horizon View used these datastores to deploy the full- and instant-clone desktops.

## Software

Table 2 shows the main software components used.

**Table 2. Software components**

Component	Description/Version
Hypervisor	ESXi 6.7
Horizon broker technology	VMware Horizon 7 version 7.5.1
Horizon broker database	Microsoft SQL Server 2016
Horizon agent	Version 7.5.1
Horizon client	Version 4.10
Management VM OS	Microsoft Windows Server 2016 (Connection Server & DB)
Virtual desktop OS	Microsoft Windows 10 Pro (64-bit)
Office application suite	Microsoft Office Professional 2016
LoginVSI test suite	Version 4.1.32.1

## Management software infrastructure

Table 3 lists the components used for the management of the VDI environment. Each component was a virtual machine and was deployed on servers other than those used for desktops, thus keeping the management infrastructure separate from the desktop deployment infrastructure.

**Table 3. Management server infrastructure**

Component	Quantity	vCPU	RAM (GB)	NIC	VMDK (GB)
VMware vCenter Server and SQL Server	1	4	16	2	250

Component	Quantity	vCPU	RAM (GB)	NIC	VMDK (GB)
Horizon Connection Server	1	4	16	2	250
Login VSI Management Console and File Server	1	2	4	2	40
Login VSI Launcher	36	4	4	1	40
DHCP Server	1	2	8	2	250
Active Directory and DNS	1	2	16	2	250

### LoginVSI Knowledge Worker

The LoginVSI Knowledge Worker profile emulates user actions such as opening a Word document, modifying an Excel spreadsheet, browsing a PDF document, web browsing or streaming a webinar. This simulates typical "advanced" user behavior and demonstrates the performance of Dell EMC PowerMax in such scenarios. The amount of I/O latency detected in the storage array is of the greatest importance. This is a parameter that directly influences the end user experience in addition to CPU and memory stats on the ESXi hosts. We chose Microsoft Windows 10 (64-bit) as the desktop operating system. Additionally, Office 2016 suite, Adobe Reader and the latest Oracle JRE were installed on the test systems.

Table 4 lists the settings used for the parent virtual machine or "golden master". From this virtual machine we deployed both our full and instant-clone desktops, which were exercised by the Knowledge Worker workload.

**Table 4. Virtual machine settings for golden master**

Guest operating system	Windows 10 Pro (64-bit)
VM Version	14
Number of vCPUs	2
Memory	4 GB
Number of vNICs	1
Number of virtual disks	1
VMDK size	40 GB

### Full-clone test results

#### Full-clone desktop deployment phase

In this section, we describe the performance during the desktop deployment phase of the 400 full-clone desktops.


The first phase of the LoginVSI test is desktop deployment. VMware Horizon creates a single replica of the 'golden-master' operating system (OS) image, and then clones this replica for each new desktop. Once cloned, each desktop performs customization steps. This process is repeated for all the desktops generated in the test. When all the desktops are provisioned, LoginVSI restarts all of them, creating a 'boot-storm'.

The eight Dell PowerEdge 740 servers with 128 GB RAM accommodated 30 full-clone desktops at 80% CPU and memory utilization each, and the Dell PowerEdge 740 servers

with 256 GB RAM accommodated 40 full-clone desktops at 80% CPU and memory utilization each.


**Note:** Based on Dell EMC Ready Solutions, 768 GB of RAM is recommended for a balanced Dell EMC PowerEdge 740 server compute node. More information can be found in the VDI Info Hub for Ready Solutions at: <https://infohub.delltechnologies.com/t/solutions/vdi/>.

Full-clone provisioning and OS customization of all 400 desktops were completed in about 45 minutes. As shown in Figure 4, maximum IOPS were 110,651 with write and read response time at 1.8 ms and 1.4 ms respectively.



**Figure 4. Full clone desktop deployment phase – PowerMax IOPs and response times**

As shown in Figure 5, maximum throughput during this time was about 6 GB/s.



**Figure 5. Full clone desktop deployment phase – PowerMax throughput**

Average CPU utilization of the four PowerEdge R740 (250 GB memory) and eight PowerEdge R740 (128 GB memory) servers was about 80%. Average memory utilization was about 72% for the PowerEdge R740 (250 GB memory) and 97% for the PowerEdge R740 (128 GB memory).

Type of PowerEdge Server	Number of Servers	RAM (GB)	CPU Utilization %	Memory Utilization %
R740	4	250	80	72
R740	8	128	80	97


### Full-clone desktop workload phase

In this section we describe the performance during the desktop workload phase of the 400 full-clone desktops.

In this phase, LoginVSI performs activity in each of the desktops, based on a 'knowledge-worker' workload profile (an advanced user). [Figure 6](#) shows the LoginVSI test results. The horizontal axis shows the number of active users and the vertical axis shows the response times in milliseconds. The blue line in the chart follows the progression of the "VSI average" compared with the number of active sessions. This is an aggregated metric, using average application latencies as more desktop sessions are added over time.

VSI<sub>max</sub> is an important metric as it shows the number of concurrent desktop sessions beyond which the user experience deteriorates. The VSI<sub>max</sub> threshold is indicated by a flat red line (not shown in the image below, as it was outside of the graph's scale).

**In our test environment we never reached the VSI<sub>max</sub> threshold of 1659.** Our VSI Index Average was only 770, indicating that the PowerMax storage system could have sustained *more than twice* the tested workload, if more compute nodes (desktop servers) were added to the test environment.



**Figure 6. LoginVSI Knowledge Worker full-clone test results**

[Figure 7](#) shows the LoginVSI test results summary. The baseline performance score was 658 for the knowledge worker test run, which, as shown at the bottom of the image, is considered by LoginVSI to be very good performance.

Test result review

**400** sessions were configured to be launched in **2700** seconds.

In total **0** sessions failed during the test:

- **0** sessions was/were not successfully launched
- **0** launched sessions failed to become active
- **400** sessions were active during the test
- **0** sessions got stuck during the test (before VSImax threshold)


With **400** sessions the maximum capacity VSImax (v4.1) **knowledgeworker** was not reached with a Login VSI baseline performance score of **658**

Login VSI index average score is **896** lower than threshold. It might be possible to launch more sessions in this configuration.

Baseline performance of **658** is: **Very good**

**Figure 7. LoginVSI full-clone test result summary**

Figure 8 shows the max IOPS and response times during the Knowledge Worker test run. We can see a read response time of 0.40 ms and write response time of 0.23 ms, which indicates a superior desktop-user experience. The test run took about 60 minutes.



**Figure 8. Full-clone workload phase – PowerMax IOPS and response times**

### Full-clone data reduction considerations

Full-clones, as the name implies, require the same raw storage capacity as the source VM (the ‘golden master’ in this case). However, PowerMax inline deduplication, compression, and thin-provisioning provide strong data reduction efficiencies. These efficiencies reduce or eliminate many of the drawbacks of the full clones increased storage consumption.

Table 5 shows the PowerMax storage efficiencies achieved during the deployment of the 400 full clones test environment. We collected the metrics by looking at the PowerMax Storage Groups (SGs) and Storage Resource Pool (SRP) information from before and after the deployment phase of the 400 desktops. It is important to note that the PowerMax storage system was dedicated for these tests to make sure no other workloads caused changes to storage allocations.

**The total desktop allocated capacity** was 5,187.9 GB. This information is available by aggregating the *Total Allocated Capacity* reported by Unisphere or SE for each of the SGs containing the 12 TB devices set as datastores for the desktops' storage. This capacity is reported prior to data reduction benefits of compression and deduplication. The 5,187.9 GB shown in the table is the *difference* in the aggregated *Total Allocated Capacity* from before and after the desktops' deployment.

Based on that, we can conclude that the average desktop raw image size was 13 GB ( $5,187.9 / 400 = 12.97$ ). This already indicates storage savings due to thin provisioning (both VMware and PowerMax). In fact, the vmdk disk image attached to the VM 'golden master' was sized at 40 GB, but due to thin provisioning, only consumed as much capacity as was actually written to it.

**The total desktops compressed capacity** was 3,061.8 GB. Compressed capacity information can be *calculated* for each SG by multiplying its *Total Allocated Capacity* by its *Compression Ratio* (both metrics are reported by Unisphere or SE). As before, we performed this calculation for each of the SGs comprising the desktops' storage and reported in the table the *difference* in total compressed capacity from before and after the desktop's deployment phase.

Based on that, we can conclude that the average desktop *compressed* image size was 7.7GB ( $3,061.8 / 400 = 7.7$ ). This demonstrates a compression rate of 1.7:1 per desktop image ( $13 / 7.7 = 1.7$ ). However, the most important data reduction factor for full-clone desktops is storage deduplication, which we will review next.

**Table 5. Storage savings after desktop deployment phase (before workload)**

Desktop metric	Value
Total number of desktops	400
Total desktops allocated capacity (GB)	5,187.9
Avg. allocated capacity per desktop (GB)	13.0
Total desktops compressed capacity (GB)	3,061.8
Avg. compressed capacity per desktop (GB)	7.7
Total exclusive allocated capacity (GB)	157.6
Avg. exclusive allocated capacity per desktop (GB)	0.4
Per desktop data reduction (times X)	32.5
Per desktop capacity saved (GB)	12.6
Per desktop capacity saved (%)	97.5
Overall data reduction @ 400 desktops (times X)	32
Overall capacity saved @ 400 desktops (GB)	5,078.6
Overall capacity saved @ 400 desktops (%)	97

**The total desktops exclusive capacity** was 157.6 GB. This number should *not* be confused with the SG *Exclusive Allocated Capacity* metric reported by Unisphere or SE. Instead, this number is simply calculated as the *difference* in the *SRP Allocated Capacity* from before and after the desktops were deployed (*SRP Allocated Capacity* is reported by

Unisphere or SE). Therefore, it represents the actual capacity added to the SRP by the desktops, after taking into account all storage data reduction benefits, including compression and deduplication<sup>2</sup>. Again, remember that there were no other users or workloads in the system to affect the SRP allocations besides our test environment.

Based on that, we can conclude that with both compression and deduplication, the actual storage capacity required by the 400 full-clone desktops was only 157.6 GB, or 0.4 GB on average per desktop ( $157.6 / 400 = 0.4$ ). [Table 5](#) also shows that the per-desktop capacity savings was 12.6 GB ( $13 - 0.4 = 12.6$ ), or 97%. The data reduction rate is not 100% deduplication among the desktops because of the small amount of unique information that each desktop contained.

The overall data reduction benefits of PowerMax compression and deduplication in this environment after the deployment phase can be described as follows:

- We started with an average desktop capacity of 13 GB (after thin provisioning benefits, but before compression and deduplication)
- We ended with 0.4 GB of actual capacity allocated on average per desktop

**This is a 97% data reduction per desktop ( $12.6 / 13 = 97\%$ ) after the deployment phase.**

Remember that these numbers apply to the state of the test environment that occurred right after the desktops were deployed, but not yet used. As users actively use the desktops, additional unique (exclusive) data will start to accumulate. Many aspects of the desktop will remain unchanged (such as the operating system and software binaries). In fact, very often VDI users save their applications' data to shared network drives and not to the virtual desktop storage. However, as a reference point, we can look at the changes in capacity from the time *after the LoginVSI workload phase completed*, as summarized in [Table 6](#).

After the workload phase, the SRP allocations grew by 176.6 GB. The overall added capacity of all 400 desktops during both *deployment and workload* phases, was therefore,  $157.6 + 176.6 = 334.2$  GB, or 0.84 GB per desktop ( $334.2 / 400 = 0.84$ ). Since the average desktop image size was 13 GB, this means that 12.16 GB of capacity was saved ( $13 - 0.84 = 12.16$ ).

The overall data reduction benefits of PowerMax compression and deduplication in this environment after the deployment and workload phases can be described as follows:

- We started with an average desktop capacity of 13 GB (after thin provisioning benefits, but before compression and deduplication)
- We ended with 0.84 GB of actual capacity allocated on average per desktop

**This is a 94% data reduction per desktop ( $12.16 / 13 = 94\%$ ) after the workload phase.**

---

<sup>2</sup> In a customer deployment the SRP allocations will be affected by Activity Based Compression (ABC). In our engineering lab we temporarily disabled ABC to collect the data reduction metrics prior to ABC affect. Remember that ABC keeps 20% of the most active allocated storage capacity uncompressed, and automatically compresses it once it falls below the 20% threshold.

Keep in mind that in a real VDI deployment, the workload phase is different for different users and your results may vary. For this reason, we have included data reduction information from both before and after the LoginVSI workload phase.

**Table 6. Storage savings after desktop workload phase**

Desktop metric	Value
Total number of desktops	400
Total desktops allocated capacity (GB)	5,187.9
Avg. allocated capacity per desktop (GB)	13.0
Total exclusive allocated capacity (GB)	334.2
Avg. exclusive allocated capacity per desktop (GB)	0.84
Per desktop capacity saved (GB)	12.16
Per desktop capacity saved (%)	94

## Instant-clone test results


### Instant-clone desktop deployment phase

In this section, we discuss the performance during the desktop deployment phase of the 600 instant-clone desktops.

For instant-clones as for full-clones, in the first phase of the desktop deployment, VMware Horizon creates a single replica of the 'golden-master' operating system (OS) image, and then clones this replica for each new desktop. Once cloned, each desktop performs customization steps. This process is repeated for all the desktops generated in the test. When all the desktops are provisioned, LoginVSI restarts all of them, creating a 'boot-storm'.

Because more resources are shared when using instant clones, each desktop server could accommodate more desktops, as compared to the full-clone test. The eight Dell PowerEdge 740 servers with 128 GB RAM were able to accommodate 45 instant-clone desktops each, and the Dell PowerEdge 740 servers with 256 GB RAM were able to accommodate 60 instant-clone desktops each without affecting CPU or memory utilization.

Instant-clone provisioning and OS customization of all 600 desktops finished in about 45 minutes. The maximum IOPs measured during that phase was 72,193 IOPs with write and read response times at 0.72ms and 0.17ms respectively, as shown in [Figure 9](#).



**Figure 9. Instant clone provisioning PowerMax IOPS and response times**


The average CPU utilization of the four PowerEdge R740 servers (256 GB memory) and of the eight PowerEdge R740 servers (128 GB memory) was about 80%. The average memory utilization was about 70% for the PowerEdge R740 server (256 GB) and 95% for the PowerEdge R740 server (128 GB).

### Instant-clone desktop workload phase

In this section we describe the performance during the desktop workload phase of the 600 instant-clone desktops.

Figure 10 shows the LoginVSI test results of the test environment with 600 instant-clone desktops. (see explanation about the graph components in [Full-clone](#) ).

As with the full-clone tests, instant-clone testing **never reached the VSImax threshold of 1675**. Our VSImax average was only 835, indicating that the PowerMax storage system could have sustained more than twice the tested workload, if more ESXi servers had been added.




**Figure 10. LoginVSI Knowledge Worker instant-clone test results**

More details about the LoginVSI test methodology can be found here: [Analyzing LoginVSI Results](#).

Figure 11 shows the LoginVSI test results summary. The baseline performance score was 675 for the knowledge worker test run, which as shown at the bottom of the image, is considered by LoginVSI to be very good performance.

Successfully completed Login VSI test with 597 **knowledgeworker** sessions. VSI<sub>max</sub> (system saturation) was not reached.



With 597 sessions the maximum capacity VSI<sub>max</sub> (v4.1) **knowledgeworker** was not reached with a Login VSI baseline performance score of 675

Login VSI index average score is 874 lower than threshold. It might be possible to launch more sessions in this configuration.

Baseline performance of 675 is: **Very good**

Figure 11. LoginVSI instant-clone test results summary

Figure 12 shows the maximum IOPS and response times achieved during the Knowledge Worker test run. The results included a read response time of 0.15 ms and write response time of 0.40 ms, which indicates a superior desktop-user experience. The test ran for approximately 60 minutes.




Figure 12. Knowledge Worker instant-clone test run PowerMax IOPS and response times

### Login VSI tests conclusion

Our testing demonstrated the performance and space-saving capabilities of the PowerMax storage system in VDI environments. Thanks to the built-in compression, deduplication, and thin provisioning capabilities of PowerMax, we provided a 97.5% storage savings with full clone desktops. Using the industry-standard tool for VDI testing,

LoginVSI, we were able to show extremely low baseline performance scores and never reached a VSI<sub>max</sub> limit or produced an unstable environment. In addition, we have provided sizing guidelines, based on the hardware used in the tests.

With increased emphasis on security and consolidation, companies can leverage the power of Virtual Desktop Infrastructure to conduct business operations. The improvements in server technology such as faster CPUs, higher core-count, and affordable large DRAM allow for higher desktop user density.

Dell EMC's PowerMax storage system improves the deployment of virtual desktops by offering high reliability and availability, data reduction, high-performance, and many other features. With the Q3 2019 platform update, SCM drives and Machine Learning data placement add even more performance, and 32Gb/s front-end modules add even more bandwidth per port.

The choice of server type affects the user-density per compute node, and the choice of full clone vs. instant clone affects the type of deployment (on-demand, or persistent desktops) based on the company's needs. PowerMax is a very capable storage platform that is flexible to support changing needs and quickly adjust to changes in workloads as necessary.

More information and deployment options for VDI with Dell EMC storage can be found in the VDI Info Hub for Ready Solutions at:  
<https://infohub.delltechnologies.com/t/solutions/vdi/>.