TECHNICAL PAPER

VMware Horizon Desktop Workloads with Hitachi Unified Compute Platform for VMware vSphere on Hitachi Virtual Storage Platform G600 with Active Flash

Lab Validation Report

By Tsuyoshi Inoue

January 2016

Hitachi Data Systems

Feedback

Hitachi Data Systems welcomes your feedback. Please share your thoughts by sending an email message to SolutionLab@hds.com. To assist the routing of this message, use the paper number in the subject and the title of this white paper in the text.

Contents

VMware Horizon Desktop Workloads with Hitachi Unified Compute Platform for VMware vSphere on Hitachi Virtual Storage Platform G600 with Active Flash

Lab Validation Report

Hitachi Data Systems enables organizations to use persistent virtual desktops through active flash in Hitachi Dynamic Tiering on Hitachi Virtual Storage Platform G600 (VSP G600). This lab validation report showcases the following benefits:

- Reduce storage management complexity
- Save 75% of the cost incurred on flash module drives

Hitachi Data Systems has already tested the use of persistent desktop for mixed workloads and validated the test setup. For more information, see Mixed Virtual Desktop Workloads with Hitachi Unified Compute Platform for VMware Horizon View on Hitachi Virtual Storage Platform G1000 Lab Validation Report.

Note — Testing of this configuration was in a lab environment. Many things affect production environments beyond prediction or duplication in a lab environment. Follow the recommended practice of conducting proof-of-concept testing for acceptable results in a non-production, isolated test environment that otherwise matches your production environment before your production implementation of this solution.

Product Features

The following information describes the hardware and software features used in testing.

Hitachi Dynamic Tiering

Hitachi Dynamic Tiering maximizes performance, simplifies data lifecycle management, and optimizes the use of tiered storage.

Instead of manually provisioning space from several storage technologies with different performance and cost characteristics, Dynamic Tiering manages multiple storage tiers as a single entity. Dynamic Tiering presents a virtual volume with embedded smart tiering to monitor access and move data based on demand.

Dynamic Tiering automatically moves infrequently referenced data to lower cost tiers of storage. This automates data placement for higher performance and lower costs. It provides automatic wide-striping performance optimization.

When Tier 1 hardware uses a solid state device or a flash module drive, active flash in Dynamic Tiering provides special care for write endurance. This complements using flash storage in your environment.

Hitachi Virtual Storage Platform Gx00 Models

Hitachi Virtual Storage Platform Gx00 models are based on industry-leading enterprise storage technology. With flash-optimized performance, these systems provide advanced capabilities previously available only in high-end storage arrays. With the Virtual Storage Platform Gx00 models, you can build a high performance, software-defined infrastructure to transform data into valuable information.

Hitachi Storage Virtualization Operating System provides storage virtualization, high availability, superior performance, and advanced data protection for all Virtual Storage Platform Gx00 models. This proven, mature software provides common features to consolidate assets, reclaim space, extend life, and reduce migration effort. New management software improves ease of use to save time and reduce complexity. The infrastructure of Storage Virtualization Operating System creates a management framework for improved IT response to business demands.

Hitachi Storage Virtualization Operating System

Hitachi Storage Virtualization Operating System spans and integrates multiple platforms. It is integrates storage system software to provide system element management and advanced storage system functions. Used across multiple platforms, Storage Virtualization Operating System includes storage virtualization, thin provisioning, storage service level controls, dynamic provisioning, and performance instrumentation.

Storage Virtualization Operating System includes standards-based management software on a Hitachi Command Suite base. This provides storage configuration and control capabilities for you.

Storage Virtualization Operating System uses Hitachi Dynamic Provisioning to provide wide striping and thin provisioning. Dynamic Provisioning provides one or more wide-striping pools across many RAID groups. Each pool has one or more dynamic provisioning virtual volumes (DP-VOLs) without initially allocating any physical space. Deploying Dynamic Provisioning avoids the routine issue of hot spots that occur on logical devices (LDEVs).

Hitachi Compute Blade 500

Hitachi Compute Blade 500 combines the high-end features with the high compute density and adaptable architecture you need to lower costs and protect investment. Safely mix a wide variety of application workloads on a highly reliable, scalable, and flexible platform. Add server management and system monitoring at no cost with Hitachi Compute Systems Manager, which can seamlessly integrate with Hitachi Command Suite in IT environments using Hitachi storage.

Brocade Switches

Brocade and Hitachi Data Systems have collaborated to deliver storage networking and data center solutions. These solutions reduce complexity and cost, as well as enable virtualization and cloud computing to increase business agility.

The solution uses the following Brocade products:

- Brocade 6520 Switch

VMware vSphere

VMware vSphere is a virtualization platform that provides a datacenter infrastructure. It features vSphere Distributed Resource Scheduler (DRS), High Availability, and Fault Tolerance.

VMware vSphere has the following components:

- **ESXi** — A hypervisor that loads directly on a physical server. It partitions one physical machine into many virtual machines that share hardware resources.
- **vCenter Server** — Management of the vSphere environment through a single user interface. With vCenter, there are features available such as vMotion, Storage vMotion, Storage Distributed Resource Scheduler, High Availability, and Fault Tolerance.

VMware Horizon

VMware Horizon transforms static desktops into secure, virtual workspaces that can be delivered on demand. Provision virtual or remote desktops and applications through a single VDI platform to streamline management and easily entitle end users.

Dynamically allocate resources with virtual storage, virtual compute and virtual networking to simplify management and drive down costs. With Horizon, reduce day-to-day operations costs with a single platform that allows you to extend virtualization from the data center to your devices.

Test Environment Configuration

Testing of the persistent desktops solution took place in the Hitachi Data Systems laboratory by using this hardware and software.

Hardware Components

Table 1 describes the details of the hardware components used to test this solution.

Table 1. Hardware Components

Hardware	Description	Version	Quantity
Hitachi Virtual Storage Platform G600	▪ Dual controllers ▪ 8 × 8 Gb/sec Fibre Channel ports ▪ 256 GB cache memory ▪ 20 × 1200 GB 10k RPM SAS HDD, 2.5 inch SFF ▪ 16 × 1600 GB FMD	83-01-20-40/06	1
Hitachi Compute Blade 500 Chassis	▪ 8-blade chassis ▪ 2 Brocade 5460 Fibre Channel switch modules, each with 6 × 8 Gb/sec uplink ports ▪ 2 Hitachi 10GbE LAN Pass-through modules, each with 16 × 10 Gb/sec uplink ports. ▪ 2 management modules ▪ 6 cooling fan modules ▪ 4 power supply modules	SVP: A0231-C-9652 5460: FOS 7.2.1	1
520H B2 Server Blade (Task User Desktops)	▪ Half blade ▪ 2 × 12-core Intel Xeon E5-2697v2 processors, 2.7 GHz ▪ 256 GB RAM ▪ 1 × 2 port 10 Gb/sec Emulex PCIe Ethernet ▪ 1 × 2 port 1 Gb/sec onboard Ethernet	Firmware: 04-40 BMC/EFI: 04-29/10-63	2

Table 1. Hardware Components (Continued)

Hardware	Description	Version	Quantity
520H B2 Server Blade (Extreme Power User Desktops)	▪ Half blade ▪ 2 × 12 core Intel Xeon E5-2697v2 processors, 2.7 GHz ▪ 256 GB RAM ▪ 1 × 2 port 10 Gb/sec Emulex PCIe Ethernet ▪ 1 × 2 port 1 Gb/sec onboard Ethernet	Firmware: 04-40 BMC/EFI: 04-29/10-63	2
520H B2 Server Blade (Infrastructure)	▪ Half Blade ▪ 2 × 12-core Intel Xeon E5-2697v2 processor, 2.7 GHz ▪ 256 GB RAM ▪ 1 × 2 port 10 Gb/sec Emulex PCIe Ethernet ▪ 1 × 2 port 1 Gb/sec onboard Ethernet	Firmware: 04-40 BMC/EFI: 04-29/10-63	2
Brocade 6520 switch	▪ SAN switch with 48 × 8 Gb Fibre Channel ports	FOS 7.2.1a	2
Brocade VDX 6740 switch	▪ Ethernet switch with 48 × 10 GbE ports		2

Software Components

Table 2 describes the details of the software components used to test this solution.

Table 2. Software Components

Software	Version
Hitachi Storage Navigator	Microcode Dependent
Hitachi Dynamic Provisioning	Microcode Dependent
Hitachi Dynamic Tiering, including active flash	Microcode Dependent
VMware vCenter server	5.5.U1a, Build 1750787
VMware ESXi	5.5.U1a, Build 1746018
VMware vSphere Client	5.5, Build 1746248
Microsoft® Windows Server® 2012	Datacenter, R2
Microsoft SQL Server®	2012 SP1
Microsoft Windows® 7	Enterprise Edition, SP1
VMware Horizon	6.0.1, Build 2088845
VMware Horizon Client	3.1.0, Build 2085634
Login VSI	4.0.12.754

Solution Infrastructure

For this testing, the infrastructure servers for VMware Horizon used for this solution are placed on separate infrastructure clusters with dedicated resources.

Table 3 describes the details of the server components required for VMware Horizon.

Table 3. VMware Horizon Components

Server Name	vCPU	Memory	Disk Size	Disk Type	Operating System
View Connection Server	4	16 GB	40 GB	Eager ZeroedThick	Microsoft Windows Server 2012 R2
View Composer	4	12 GB	40 GB	Eager ZeroedThick	Microsoft Windows Server 2012 R2
Domain Controller	2	8 GB	40 GB	Eager ZeroedThick	Microsoft Windows Server 2012 R2
Database Server	4	16 GB	40 GB (operating system) 60 GB (data)	Eager ZeroedThick	Microsoft Windows Server 2012 R2 Microsoft SQL Server 2012 SP1

All these virtual machines were configured with VMware Paravirtual SCSI Controller. The domain controller was deployed to support user authentication and domain services for the VMware Horizon infrastructure.

The compute nodes were configured as follows:

- **Compute Cluster 01:** 2 VMware ESXi hosts, dedicated for task user desktops
- **Compute Cluster 02:** 2 VMware ESXi hosts, dedicated for extreme power user desktops

To separate network traffic for the two different workloads, each compute cluster was configured with the following:

- Its own dvPortGroup
- Separate VLANs for the application virtual machines, which include the following:
  - Extreme power users
  - Task users

To isolate storage traffic, active flash pool virtual volumes were mapped to the VMware ESXi hosts from each compute cluster via separate storage ports, as indicated in Table 4. The storage multipathing policy was set to **round robin**.

Table 4. Compute Clusters and Storage Ports Mapping

Compute Cluster	Number of Hosts	Number of HBAs	Storage Ports (2 in each Virtual Storage Platform cluster)
Task Users	2	2	1A , 1B , 2A , 2B
Extreme Power Users	2	2	3A , 3B , 4A , 4B

For these tests, this was the infrastructure configuration:

- VMware Horizon management and administration components were placed on a separate infrastructure cluster.
- The virtual desktops (task users and extreme power users) were placed on separate compute clusters.

Figure 1 gives a high-level overview of the infrastructure and component placement.

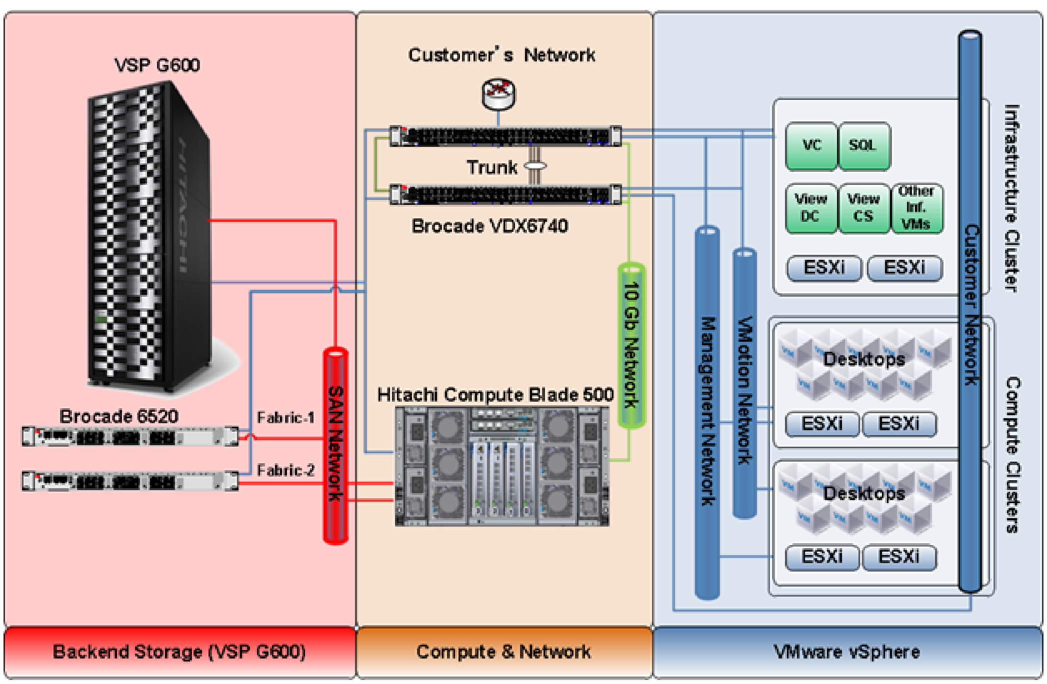

Figure 1

Test Methodology

This describes the test methodology used to test this VMware Horizon environment. The following were the purposes of the tests:

- Show acceptable levels of performance and end user experience when running VMware Horizon workloads concurrently.
- Compare on the same test bed with VMware Horizon a dynamic provisioning pool created with Hitachi Dynamic Provisioning with an active flash pool created with Hitachi Dynamic Tiering with a different disk layout.
- Check if there is any change in the performance and user experience when using active flash in Dynamic Tiering by decreasing the number of flash module drives and replacing them with SAS HDD of the same capacity.

VMware Horizon Configuration

Most of the configuration of the pools for VMware Horizon for the task users and extreme power users is based on previous tests in Mixed Virtual Desktop Workloads with Hitachi Unified Compute Platform for VMware Horizon View on Hitachi Virtual Storage Platform G1000 Lab Validation Report.

Extreme Power Users — Persistent Full Clone DesktopPool

A full clone desktop pool of 180 desktops with dedicated user assignment was configured in VMware Horizon for use in the concurrent workload testing.

The virtual machine template used for the full clones was configured for a heavy power user workload type, as defined in VMware Horizon 6 Storage Considerations Technical White Paper. The Virtual Machines were configured to run a total average load of 150-175 IOPS to simulate software development build and compile tasks.

A 2.5 GB memory reservation was configured on the virtual machines to reduce the datastore space consumed by the vswap files. The virtual machine template was prepared for VDI use by following the guidelines in VMware Horizon ViewOptimization Guide for Windows 7 and Windows 8Optimization Guide for Desktops and Servers in View in VMware Horizon 6 and VMware Horizon Air Desktops and VMware Horizon Air Apps Technical White Paper.

The virtual machine template was configured also for use with Login VSI by following the guidelines established in Login VSI: Documentation.

Table 5 lists the configuration details of the Virtual Machine template used for the extreme power user full clone desktops.

Table 5. Configuration Details of Virtual Machine Template for Full Clones

Operating System	Microsoft Windows 7, 64-bit
vCPU Allocation	2
Memory Allocation	4 GB (2.5 GB reserved)
Desktop Disk and Type	34 GB, thin provisioned
	2 GB, thick provisioned (vdbench test)
SCSI Controller	LSI Logic SAS
Average Steady State IOPS	150-175
High-Density vCPU per Core	7.5

Each VMFS datastore contained 90 desktops. Each ESXi server in the compute cluster used for extreme power users workloads were configured to host exactly 90 desktops to obtain accurate end user experience metrics during testing.

Task Users — Linked Clone Desktop Pool

A linked clone desktop pool with a floating user assignment of 400 desktops was configured in VMware Horizon for use in the concurrent workload testing.

The virtual machine template used for the linked clones was configured for a task user workload type, as defined in VMware Horizon 6 Storage Considerations Technical White Paper. The virtual machine template was prepared for VDI use following the guidelines in Optimization Guide for Desktops and Servers in View in VMware Horizon 6 and VMware Horizon Air Desktops and VMware Horizon Air Apps Technical White Paper.

The virtual machine template was configured also for use with Login VSI by following the guidelines established in Login VSI Documentation.

Table 6 lists the configuration details of the Virtual Machine template used for the linked clone desktops.

Table 6. Configuration Details of Virtual Machine Template for Linked Clones

Operating System	Microsoft Windows 7, 32-bit
vCPU Allocation	1
Memory Allocation	1 GB
Desktop Disk and Type	24 GB, thin provisioned
SCSI Controller	LSI Logic SAS
Average Steady State IOPS	3-7
High-Density vCPU per Core	8.3

Each VMFS datastore contained 50 desktops (plus or minus 2 desktops). Each VMware ESXi server in the compute cluster used for task user workloads was configured to host exactly 200 desktops in order to obtain accurate end user experience metrics during the testing.

Storage Pool Configuration

Table 7 shows the storage pool configuration and disk layout used in when testing this solution on Hitachi Virtual Storage Platform G600. These tests were performed on the same test bed with mixed extreme and task user profile with 580 users, as previous VMware Horizon tests from Mixed Virtual Desktop Workloads with Hitachi Unified Compute Platform for VMware Horizon View on Hitachi Virtual Storage Platform G1000 Lab Validation Report.

Table 7. Storage Pool Configuration

Configuration		Pool Type	RAID Configuration	Drive Quantity
1	Without active flash	Dynamic provisioning (1.6 TB FMD pool)	RAID-10 (2D+2D) × 2 parity groups	8
		Dynamic provisioning (1.2 TB SAS 10k RPM HDD pool)	RAID-10 (2D+2D) × 3 parity groups	12
2	With active flash	Active flash (Tier 1, 1.6 TB FMD pool)	RAID-10 (2D+2D) × 1 parity group	4
		Active flash (Tier 2, 1.2 TB SAS 10k RPM HDD pool)	RAID-10 (2D+2D) × 5 parity groups	20
3	Without active flash	Dynamic provisioning (1.6 TB FMD pool)	RAID-10 (2D+2D) × 4 parity groups	16

Configuration 1 (Without Active Flash): Multiple Dynamic Provisioning Pools with Different Types of Drives

This configuration used multiple dynamic provisioning pools on Hitachi Virtual Storage Platform G600.

Two flash module drive parity groups are required to meet the peak performance for the following:

- During the logon storm
- The capacity demand for extreme power user allocated as full clone and linked clone replica for the task user on the test.

Three SAS parity groups are required to meet the peak performance and capacity demand on linked clone for task user.

Figure 2 illustrates the storage configuration used for Configuration 1.

The desktop pool was configured to use eight datastores for linked clone storage, which were presented to both hosts in the compute cluster used for task user workloads. Each datastore was presented to the cluster as a dynamic provisioning pool virtual volume, with the dynamic provisioning pool containing three RAID-10 (2D+2D) parity groups of 1.2 TB 10k RPM SAS HDD.

The desktop pool was configured to use the following:

- A single datastore for replica storage
- Two datastores for full clone storage

The replica and full clone datastores were presented to the cluster as a dynamic provisioning pool virtual volume with the dynamic provisioning pool containing two RAID-10 (2D+2D) parity groups of 1.6 TB FMD.

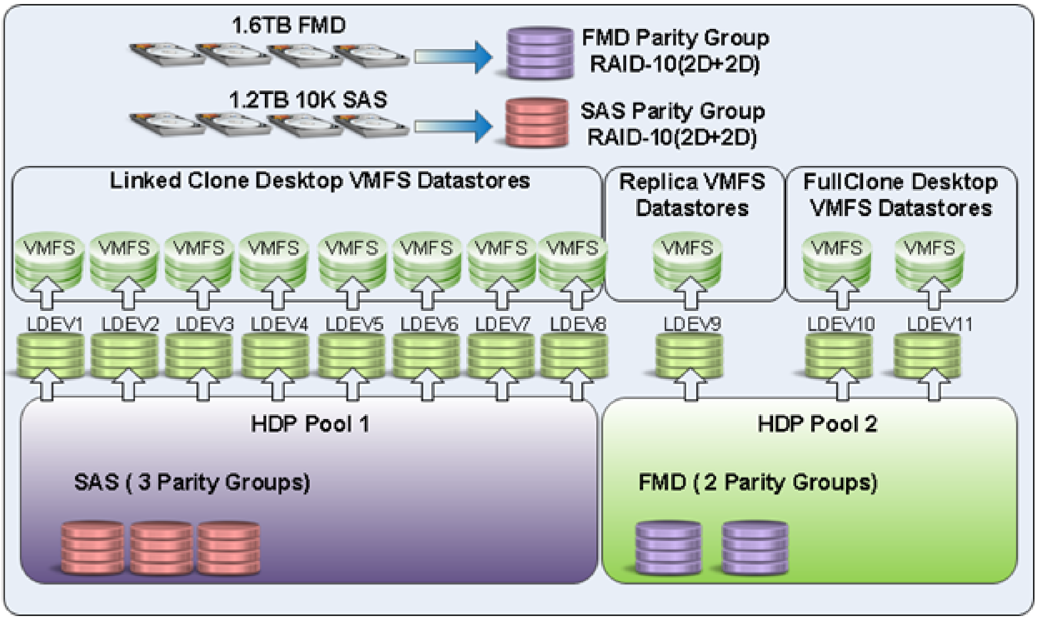

Figure 2

Configuration 2 (With Active Flash): A Single Active Flash Pool (Tiered pool) with Different Types of Drives

This configuration used active flash in Hitachi Dynamic Tiering. It was based on the result of monitoring through the same test bed as Configuration 1.

The monitor in active flash exposed hot data across the upper tier (Tier 1) on the FMD and the lower tier (Tier 2) on SAS HDD. Active flash also provides the **max performance utilization** performance metric, calculated through the monitor with storage inherent intelligence.

From the monitor result, 0.6 TB was hot data with 15% max performance utilization for active flash during the logon storm against disk layout that has total capacity of approximately 6.4 TB on the Configuration 1 disk layout.

The boot storm duration time is within 5 minutes. This is too short to be monitored by the active flash monitor, due to the minimum duration time of 30 minutes. For the boot storm, only one FMD parity group is necessary to meet the peak IOPS requirement observed during boot storm on this workload.

One FMD parity group on Configuration 2 was removed and three SAS HDD parity groups were added to meet the virtual machine capacity demand for extreme power users allocated as full clone.

Figure 3 on page 12 illustrates the storage configuration used for Configuration 2.

- Full clone, replica, and linked clone desktop pool layout were presented to each host. They were configured as same as Configuration 1, using Hitachi Dynamic Provisioning.
- Each datastore was presented to the cluster as an active flash pool virtual volume that contained five RAID-10 (2D+2D) parity groups of 1.2 TB 10k RPM SAS HDD and a single RAID-10 (2D+2D) parity group with 1.6 TB FMD.
- Active flash in Hitachi Dynamic Tiering relocates data to the proper tier across all the tiers in a single pool configured by multiple media, such as FMD and HDD. This is based on the I/O counter per page (42 MB) in real time or the result of monitor in a certain time period (minimum 30 minutes), according to workload characteristic.
- For the initial tier placement of each user profile on Configuration 2 on Table 7 on page 9 that used the active flash pool, the Tier 1 configuration was with FMD for allocation for the extreme power user. This extreme power use placement generated massive IOPS through all VMware Horizon events. The rest of extreme power user full clone space was allocated from Tier 2 with SAS 10k RPM HDD. Both replica and linked clone for the task user were all allocated from Tier 2 with SAS 10k RPM HDD.

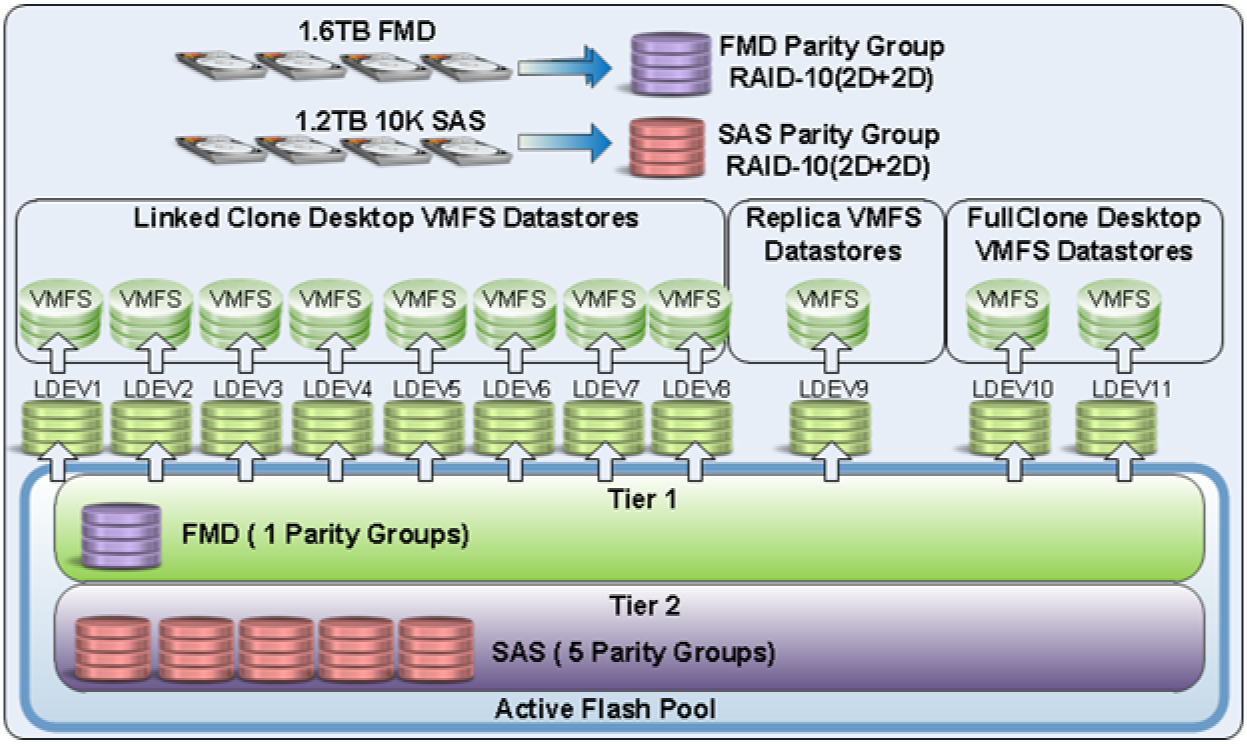

Figure 3

Configuration 3 (Without Active Flash): A Single Dynamic Provisioning Pool with All Flash Module Drives

This configuration used a single dynamic provisioning pool with all FMDs that were all allocated to the following:

- Full cloned extreme users
- Linked cloned task users.

Four FMD parity groups were required to meet the capacity demand for all users.

Figure 4 on page 13 illustrates the storage configuration used for Configuration 3.

- The desktop pool was configured to use eight datastores for linked clone storage, which were presented to both hosts in the compute cluster used for task user workloads. Each datastore was presented to the cluster as a dynamic provisioning pool virtual volume that contained four RAID-10 (2D+2D) parity groups of 1.6 TB FMD.
- The desktop pool was configured to use a single datastore for replica storage. There were two datastores configured for full clone storage. The replica and full clone datastores were presented to the cluster as a dynamic provisioning pool virtual volume that contained four RAID-10 (2D+2D) parity groups of 1.6 TB FMD.

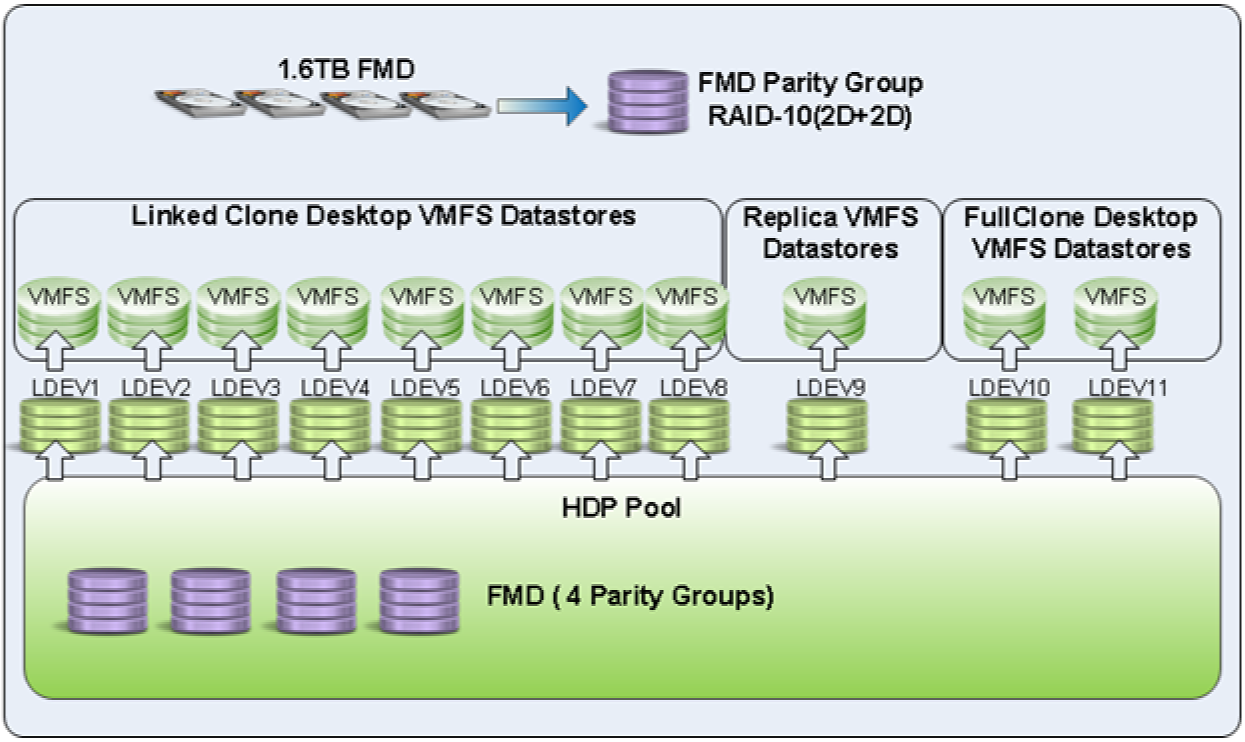

Figure 4

Login VSI Test Harness Configuration

This explains the configuration of the Login VSI test harness.

Extreme Power Users — Persistent and Full Clone Desktop Pool

Login VSI was used to generate a heavy workload on the desktops. Login VSI launchers were configured to initiate up to 15 PCoIP sessions to the VMware Horizon connection server to simulate end-to-end execution of the entire VMware Horizon infrastructure stack.

The standard “heavy” Login VSI workload was modified to include running Vdbench in the background during all test phases in order to generate the additional IOPS necessary to meet the 150-175 IOPS target for the workload.

To ensure the I/O profile used in Vdbench was only applied to the desktop application usage, the individual application I/O profiles were captured from previous testing and used as Vdbench workload definitions.

Table 8 on page 14 lists the applications and associated I/O profiles used to define the Vdbench workload definitions.

Table 8. Vdbench Application I/O Profiles Used in Workload Definitions

Application	% Read	% Write	% Random	% Sequential
Microsoft PowerPoint®	19	81	75	25
Adobe Acrobat Reader	22	78	71	29
Microsoft Outlook®	17	83	81	19
Microsoft Excel®	18	82	74	26
Mozilla Firefox	36	64	55	45
Microsoft Internet Explorer®	17	83	80	20
Microsoft Web Album	27	73	68	32
Microsoft Media Player	23	77	61	39
Microsoft Word	26	74	62	38

Login VSI was configured to stagger logons of all 180 users with a single logon that occurs every 5.5 seconds. This ensured all 180 users logged into the desktops and began the steady state workloads within a 17 minute period in order to simulate a moderate logon storm.

Task Users — Linked Clone Desktop Pool

Login VSI was used to generate a light workload on the desktops.

Login VSI launchers were configured to initiate up to 15 PCoIP sessions to the VMware Horizon connection server to simulate end-to-end execution of the entire VMware Horizon infrastructure stack.

The standard "light" Login VSI workload was used to simulate a task user workload.

Login VSI was configured to stagger logons of all 400 users with a single logon occurring every 6 seconds. This ensures that all 400 users logged into the desktops and began the steady state workloads within a 40 minute period to simulate a moderate logon storm.

Test Cases

These are the test cases used.

Test Case 1: Concurrent Boot Storm for Virtual Desktops

In this test, an immediate power on of 580 desktops for linked clone and full clone was performed through the VMware Virtual Infrastructure Client. This boot storm test contains the following mix of desktop and server types:

- 400 linked clone desktops configured for task user workloads
- 180 full clone desktops configured for extreme power user workloads

This test case was performed to ensure that the storage assigned to the desktop pools performed adequately under stress, and to determine the amount of time necessary for the desktops to be ready for logon by the end user.

Test Case 2: Concurrent Logon Storm and Steady State for Extreme Power Users and Task Users

Login VSI allows for staggered user logons within a desktop pool and then looping the configured workload for a specified amount of time before logging off the session. For this test case, both desktop pools started their logon storm at the same time and each pool logged on a single user every 5.5 seconds.

All the Login VSI sessions were configured to log off approximately five hours after logging on to ensure the following:

- Each desktop ran a steady state workload loop at least twice
- Both desktop workloads (extreme power users and task users) ran together for at least three hours

Test Results Summary – Test Case 2

This is a summary of the results for Test Case 2.

Application Experience

Login VSI was used to measure the end user application response time and to determine the maximum number of desktops that can be supported on the tested infrastructure running the specified workloads.

Application Response Times

Login VSI reported the time required for various operations to complete within the desktop during the test.

- All operations for all configurations completed in less than 3.1 seconds.
- If the zip high compression (ZHC) metric from the linked clone desktops is removed from the analyzed metrics, all other operations for all configurations completed in less than 2.6 seconds.

These performance metrics prove that the tested infrastructure provides adequate end user experience for 400 task users and 180 extreme power users.

Table 9 on page 15 shows the operation abbreviations used in Login VSI and a description of the action taken during the operation.

Table 9. Login VSI Operation Descriptions

Login VSI Operation	Description
FCDL	File Copy Document Local
FCDS	File Copy Document Share
FCTL	File Copy Text Local
FCTS	File Copy Text Share
NFP	Notepad File Print
NSLD	Notepad Start/Load File
WFO	Windows File Open
WSLD	Word Start/Load File
ZHC	Zip High Compression
ZNC	Zip No Compression

VSI Max Metric

VSI max is a metric that indicates the number of desktop sessions that the tested infrastructure can support running a specified workload. Several metrics are monitored during the concurrent workload testing.

While the VSI max metric from Login VSI gives a good indicator of when the maximum user density has been reached due to infrastructure resource saturation, during this testing the densities were lowered slightly below VSI Max while carefully monitoring the guest CPU ready percentage, ESXi CPU load and utilization, and guest CPU CoStop values. This metric ensured that being able to guarantee end user experience while still having enough overall resource headroom to support bursts of CPU or I/O usage by the underlying guest operating system.

For more detailed metrics on the infrastructure such as storage, see “Results Analysis,” starting on page 19.

Figure 5 on page 16 shows the task user application experience metrics as reported by LoginVSI. Figure 6 on page 17 shows the extreme power user application experience metrics as reported by LoginVSI.

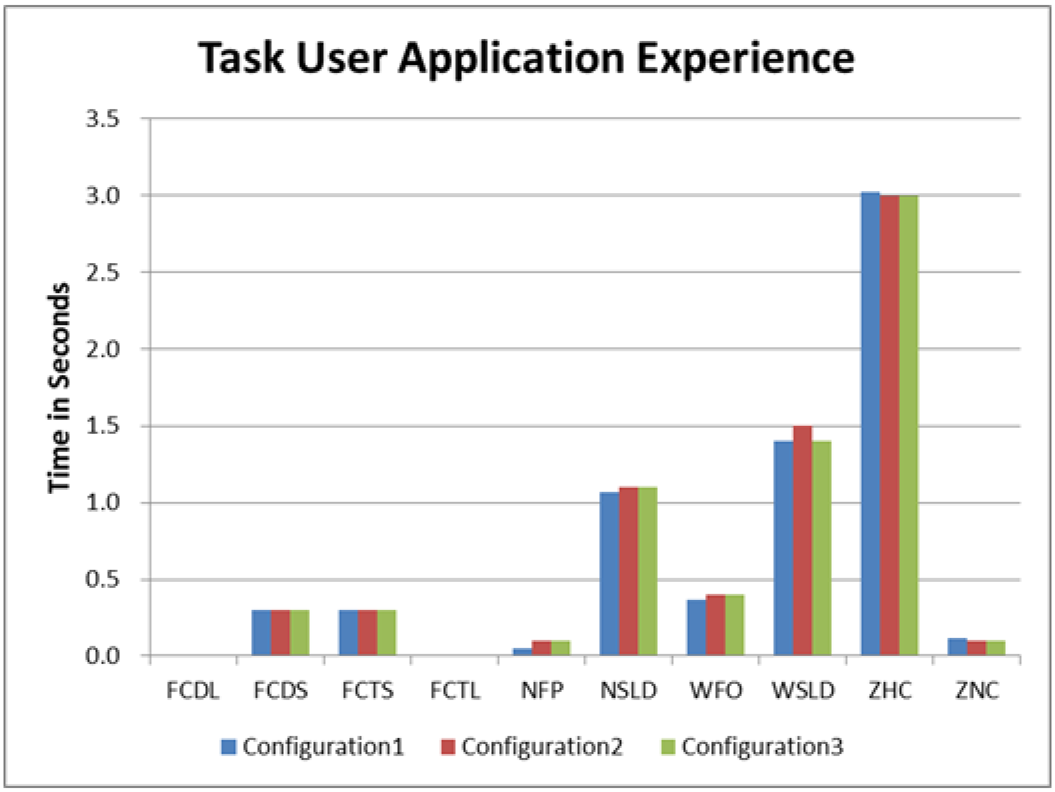

Figure 5

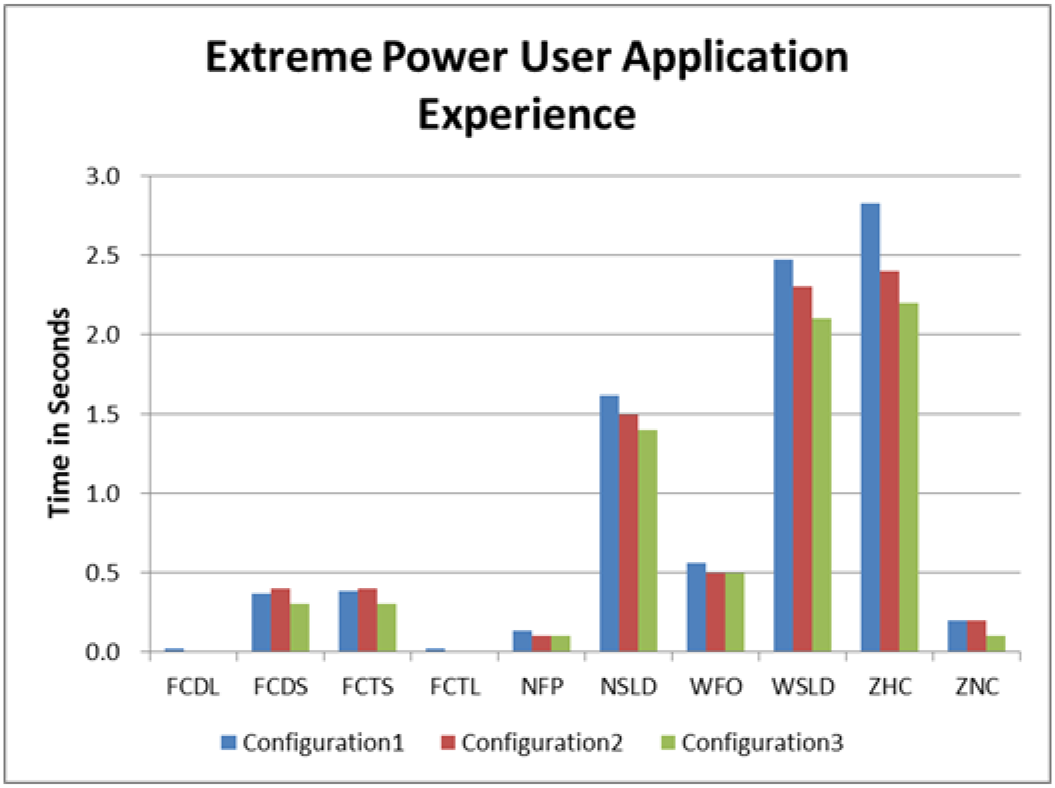
Extreme Power User Application Experience
Time in Seconds
3.0
2.5
2.0
1.5
1.0
0.5
0.0
FCDL
FCDS
FCTS
FCTL
NFP
NSLD
WFO
WSLD
ZHC
ZNC
Configuration1
Configuration2
Configuration3

Figure 6

Conclusion

Hitachi Dynamic Tiering with active flash provides the ability to combine FMD and HDD in a configuration, therefore improving the cost performance efficiency. As demonstrated in the test cases, this was achieved by going beyond the multiple physical media boundary in a single pool, simplifying storage management.

As observed in Figure 5 on page 16, there is no difference across all configurations used in the test. For extreme users, Figure 6 on page 17 shows a 0.5 seconds difference in NSLD, WSLD, and ZHC, which is not a disruption in user experience. The test result also confirms that active flash ensures no disruption in terms of storage performance and user experience against Hitachi Dynamic Provisioning, after reducing the FMDs as listed in Table 10.

Table 10. Drive Cost Comparison

Configuration		Pool Type	RAID Configuration	Drive Quantity
1	Without active flash	Dynamic provisioning pool with FMD	1.6 TB FMD, RAID-1 (2D+2D) × 2parity groups	8
2	With active flash	Active flash pool	Tier 1: 1.6 TB FMD,RAID-10 (2D+2D) × 1 parity group	4
3	Without active flash	Dynamic provisioning pool with FMD	1.6 TB FMD, RAID-10 (2D+2D) × 4 parity groups	16

Configuration 2, with only four FMDs, provides approximately the same performance as Configuration 1 with eight FMDs and Configuration 3 with 16 FMDs. By using Hitachi Dynamic Tiering with active flash, you can save about 75% of the FMD cost and use cost effective storage, such as the SAS HDD, to obtain approximately the same performance as that of Hitachi Dynamic Provisioning with 100% FMDs.

Results Analysis

This is an analysis of the test results.

Test Case 1: Boot Storm for Virtual Desktops

The immediate power on of 580 desktops took 3 minutes and 4-20 seconds. This was measured from the time that the desktops were concurrently powered on from VMware vCenter until the time all 580 desktops showed as *available* within VMware Horizon Administrator.

Six minutes of metrics are graphed to illustrate the periods prior to power-on and after the desktops are marked as available and ready to login.

Storage Infrastructure

Multiple performance metrics from Hitachi Virtual Storage Platform G600 were analyzed to ensure that the storage could support the stress of an immediate power on of all of desktops.

Pool IOPS

Figure 7 and Figure 8 illustrate the total combined IOPS during the boot storm for the LDEVs used within the pool used for the virtual desktops.

- Figure 7 on page 20 shows there is no difference from Configuration 1 to Configuration 3 on boot storm.
- Figure 8 on page 20 illustrates that higher IOPS observed in linked clone for both configurations comes from the number of powered-on virtual machines. There is no dependency on type of user profile on boot storm.
- Table 11 on page 21 shows the I/O characteristic observed in boot storm.
  (1) Configuration 1 Total IOPS peak at approximately 41,021
  (2) Configuration 2 Total IOPS peak at approximately 42,692
  (3) Configuration 3 Total IOPS peak at approximately 42,013.

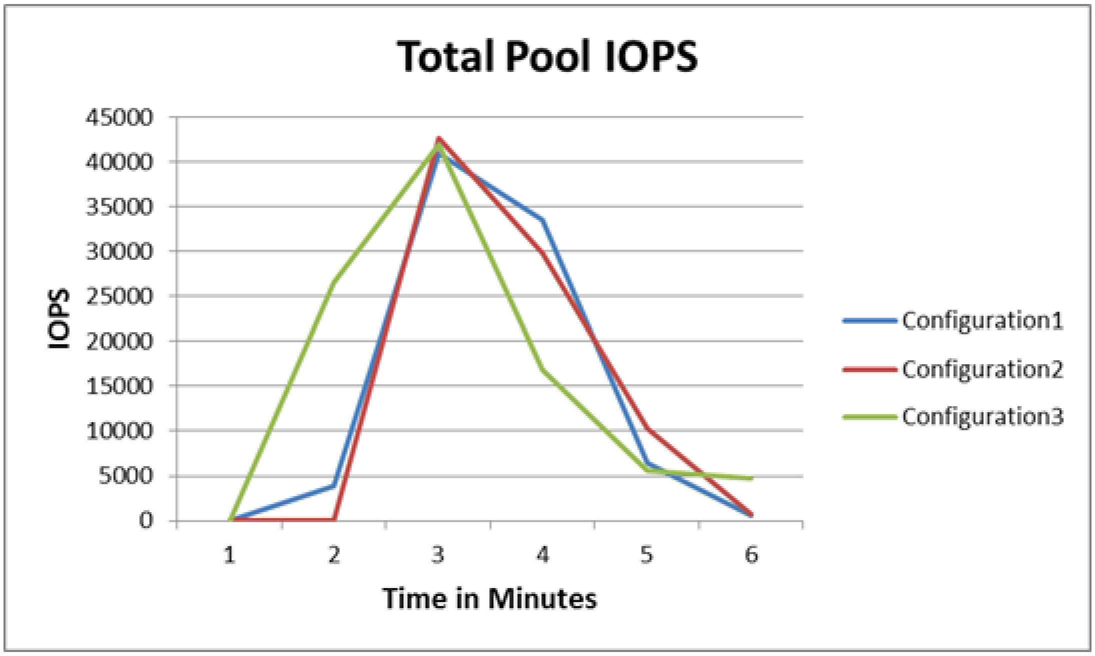

Figure 7

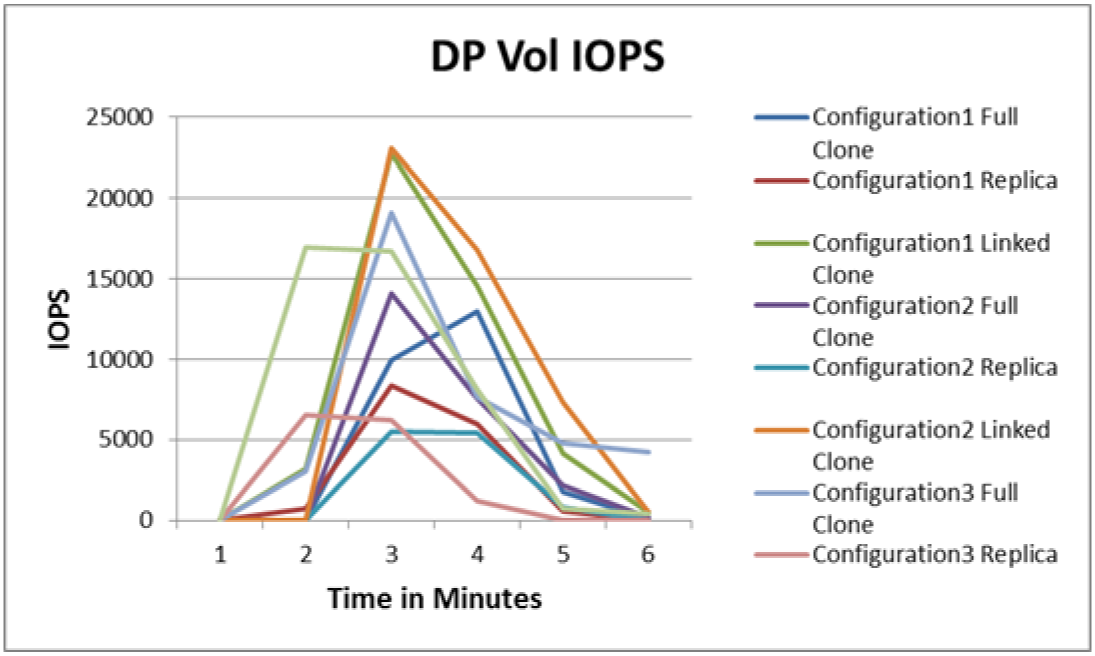

Figure 8

Table 11. I/O Characteristic of VDI Event

VDI Event	% Random Read	% Random Write	% Sequential Read	% Sequential Write
Boot Storm	53	24	18	6

Processor and Cache Write Pending

Figure 9 on page 21 and Figure 10 on page 22 illustrate the management processor utilization and cache write pending rate observed during the boot storm.

- MPU utilization of all configurations did not rise above 27%.
- Cache write pending of all configurations did not rise above 9%.

These metrics show that, even during a boot storm, Hitachi Virtual Storage Platform G600 still had plenty of resources for additional workloads.

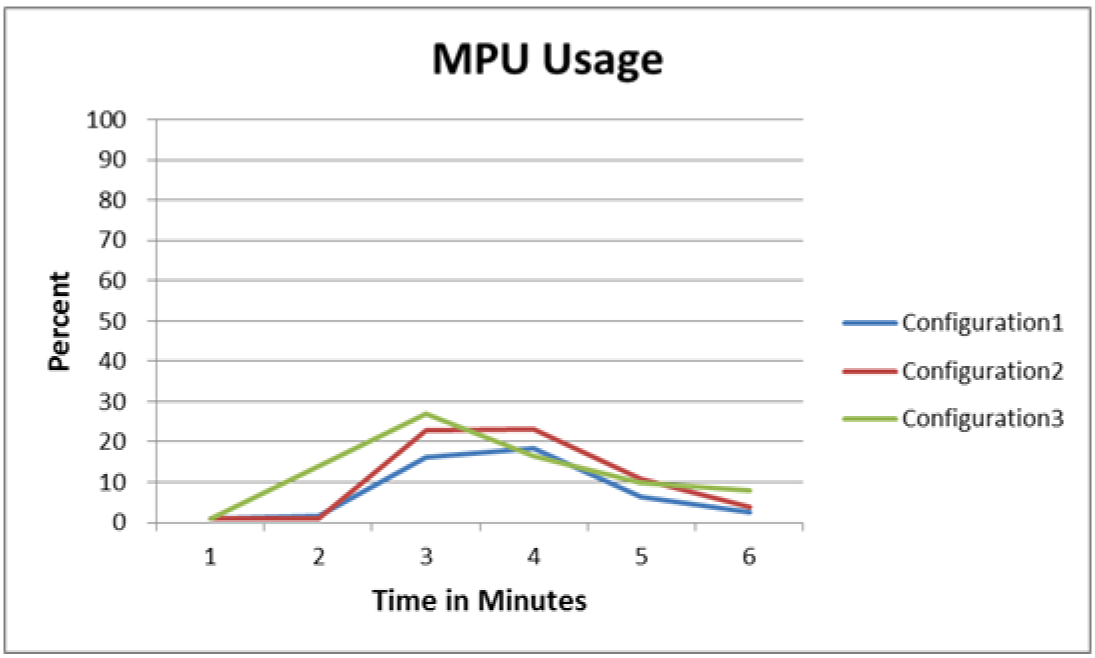

Figure 9

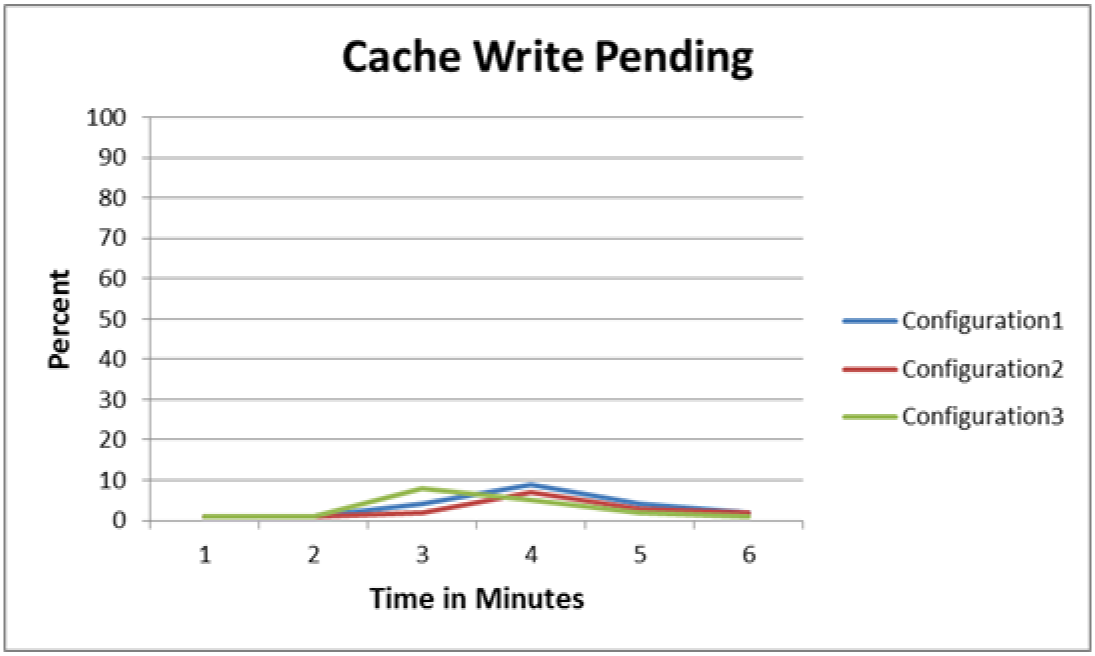

Figure 10

Physical Disk

Figure 11 on page 23 illustrates the average physical disk busy rate for the LDEVs used within the pool used for the virtual desktops.

- The Configuration 1 disk busy rate on FMD pool did not rise above 33% during the boot storm.
- The Configuration 1 disk busy rate on SAS 10k RPM HDD pool peaked at 97% during the boot storm.
- The Configuration 2 disk busy rate on Tier 1 (FMD) did not rise above 76% during the boot storm.
- The Configuration 2 disk busy rate on Tier 2 (SAS 10k RPM HDD) did not rise above 54% during the boot storm.
- The Configuration 3 disk busy rate on the FMD pool did not rise above 50% during the boot storm.

The FMD busy rate of 76% was observed in Configuration 2 with active flash, which is higher than the 33% observed in Configuration 1 with Hitachi Dynamic Provisioning. This comes from the fewer amount of FMD parity groups on Configuration 2, but still safe state from storage performance perspective.

The SAS 10k RPM HDD busy rate of 54% observed in active flash is lower than the 76% observed in Configuration 1 with Hitachi Dynamic Provisioning. This comes from having more SAS 10k RPM HDD on Configuration 2.

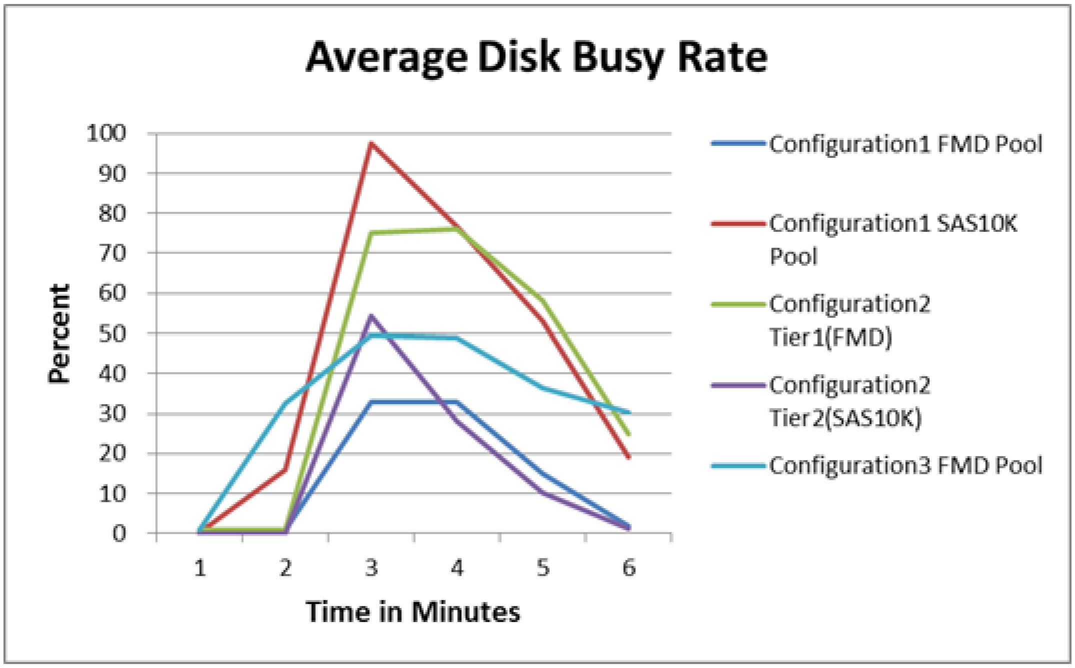

Figure 11

Storage Port Latency

Figure 12 on page 24 illustrates the average storage latency observed on Hitachi Virtual Storage Platform G600 during the boot storm.

- **Extreme power user**
  - Average latency on ports used for the Configuration 1 extreme power user peaked at 1.5 milliseconds during the boot storm.
  - Average latency on ports used for the Configuration 2 extreme power user peaked at 3.5 milliseconds during the boot storm.
  - Average latency on ports used for the Configuration 3 extreme power user peaked at 1.7 milliseconds during the boot storm.
- **Task user**
  - Average latency on ports used for the Configuration 1 task user peaked at 5.6 milliseconds during the boot storm.
  - Average latency on ports used for the Configuration 2 task user peaked at 2.1 milliseconds during the boot storm.
  - Average latency on ports used for the Configuration 3 task user peaked at 1.3 milliseconds during the boot storm.

The storage port latency maximum of 3.5 msec, related to the FMD observed in Configuration 1 with active flash, is higher than maximum of 1.5 msec observed in Configuration 2 with Hitachi Dynamic Provisioning. This comes from the same reason that causes the differences in disk busy rate illustrated in Figure 11. The storage latency of less than 15 msec does not matter from storage performance perspective.

The storage port latency maximum of 2.1 msec, related to the SAS HDD observed in Configuration 1 with active flash, is lower than the maximum 5 msec observed in Configuration 1 with Hitachi Dynamic Provisioning. This comes from the same reason that causes the differences in disk busy rate illustrated in Figure 11 on page 23.

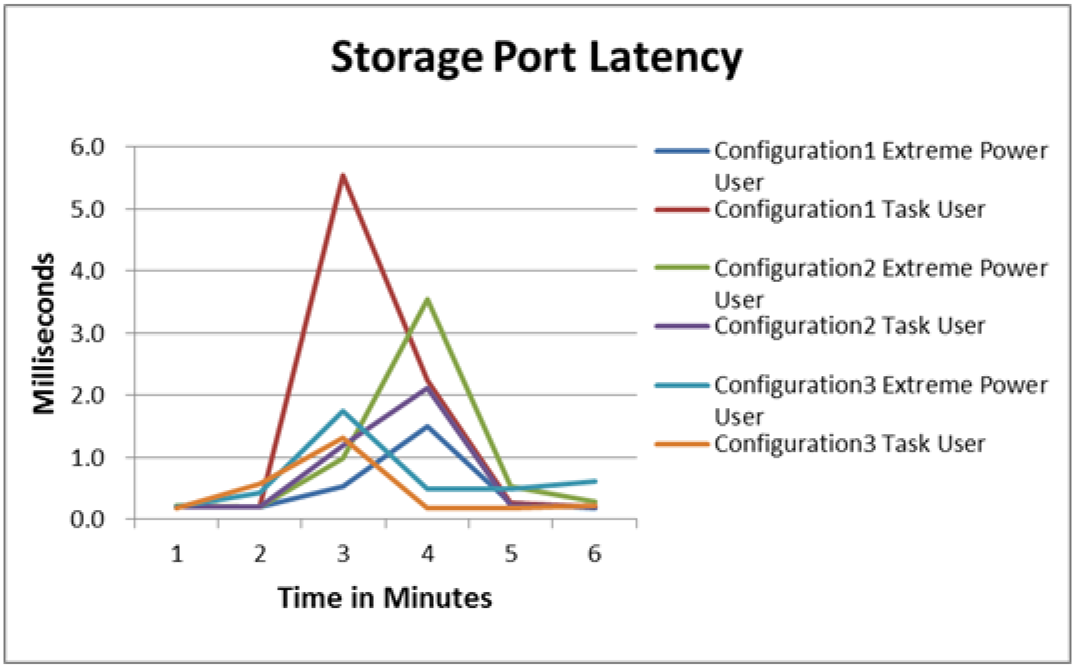

Figure 12

Test Case 2: Logon Storm and Steady State for Extreme Power Users and Task Users

This test generated multiple performance metrics from VMware ESXi hypervisors, Hitachi Virtual Storage Platform G600, and Login VSI test harnesses during the logon storm and steady state operations. This test contained the following mix of desktop and server types:

- 400 linked clone desktops configured for task user workloads
- 180 full clone desktops configured for extreme power user workloads

Compute Infrastructure

Esxtop was used to capture performance metrics on all four hosts in the VMware Horizon clusters during the logon storm or steady state operations for the desktops. The following metrics illustrate the performance of the guest virtual machines.

Guest IOPS Performance — Extreme Power User

Figure 13 on page 25 shows in-guest IOPS performance metrics obtained from esxtop during the logon storm and steady state operations for the extreme power users. The graphs show four representative desktops from the 180 that ran the workloads.

- IOPS peak at 654 during the logon storm
- IOPS peak at 188 with an average at 153 during steady state

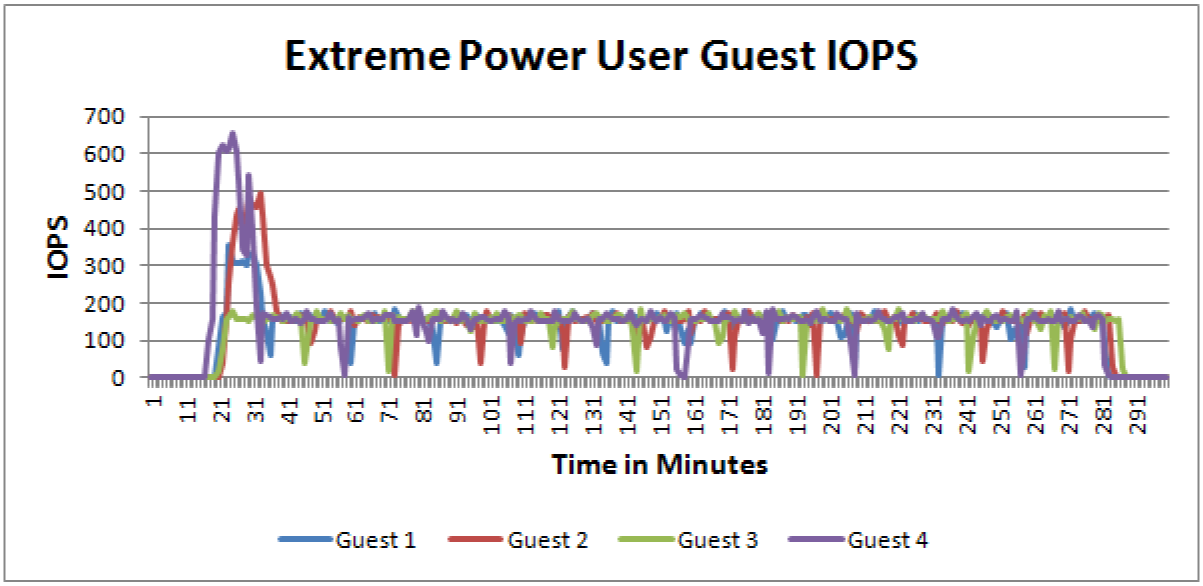

Figure 13

Guest IOPS Performance - Task User

Figure 14 shows in-guest IOPS performance metrics obtained from esxtop during the logon storm and steady state operations for the task users. The graph shows four representative desktops from the 400 desktops that ran the workloads.

- IOPS peak at 23 during the logon storm
- IOPS peak at 19 with averages between 3 and 4 during steady state

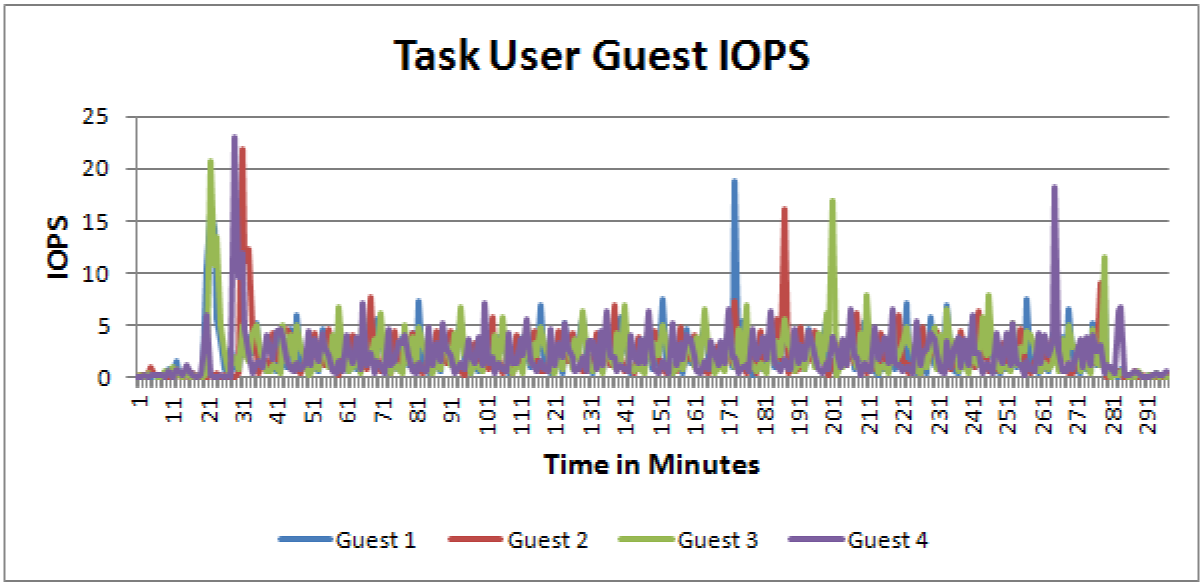

Figure 14

Storage Infrastructure

Multiple performance metrics were captured from Hitachi Virtual Storage Platform G600 during the logon storm and steady state operations for the virtual desktops. The following metrics illustrate the performance of the storage array.

Pool IOPS

Figure 15 and Figure 16 on page 27 illustrate the total combined IOPS during the logon storm and steady state operations for the virtual desktops. Table 12 on page 27 shows the I/O characteristic observed in the logon storm and steady state operations.

- Configuration 1 total IOPS peak at approximately 55,309
- Configuration 2 total IOPS peak at approximately 53,550
- Configuration 3 total IOPS peak at approximately 54,545

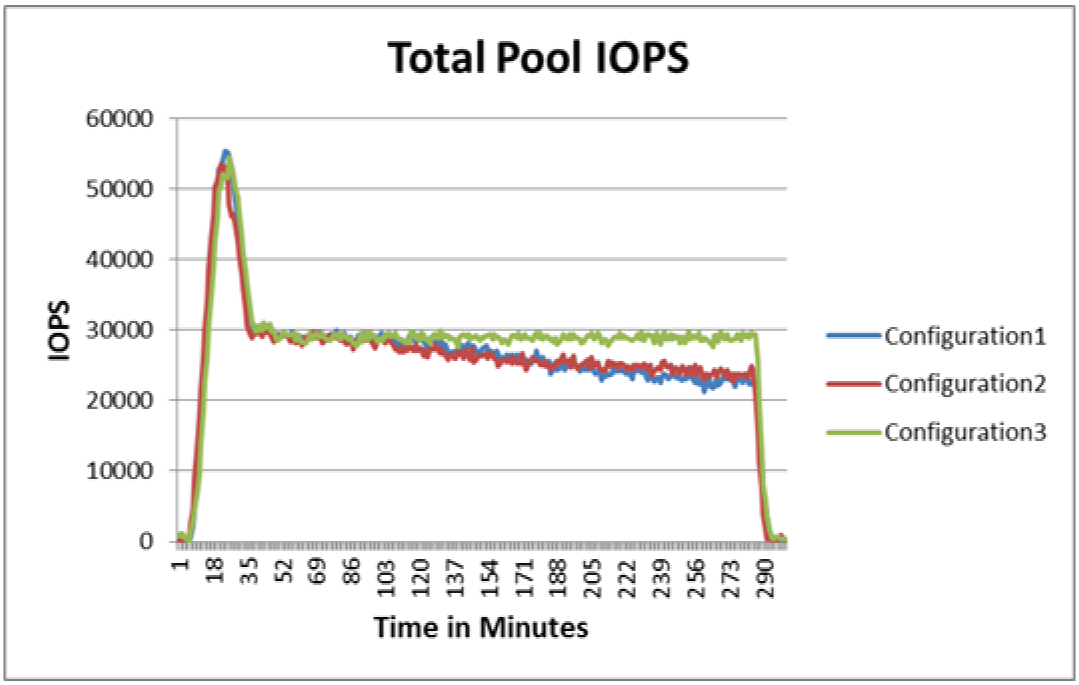

Figure 15

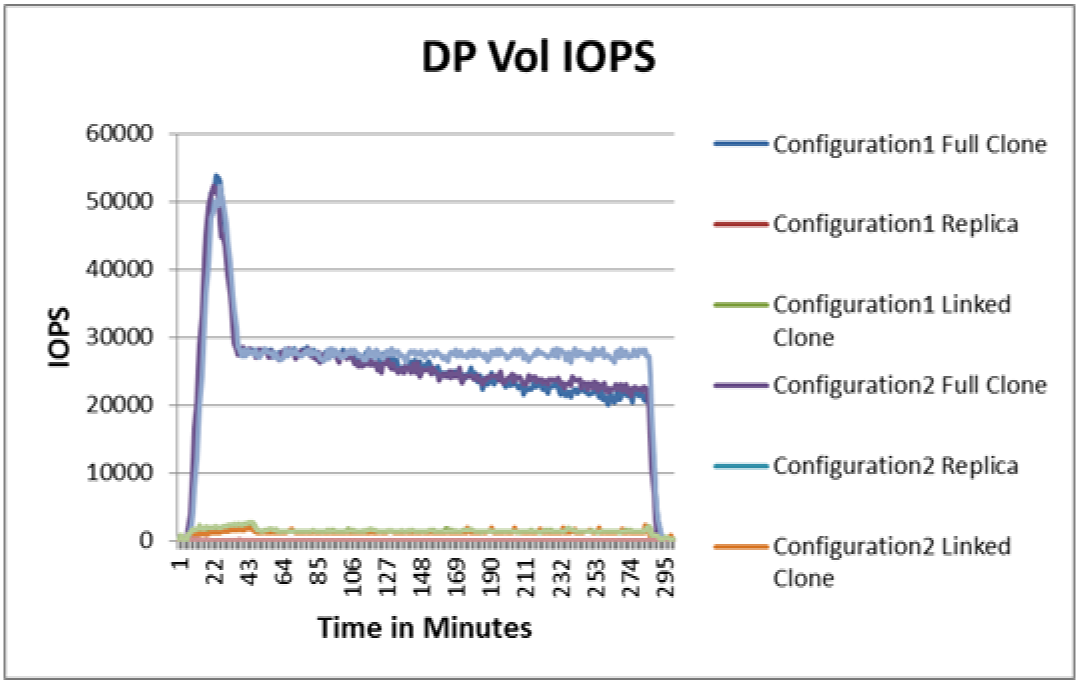

Figure 16

Table 12. I/O Characteristic of VDI Events

VDI Events	% Random Read	% Random Write	% Sequential Read	% Sequential Write
Logon Storm	19	75	1	6
Steady State	18	77	0	4

Processor and Cache Write Pending

Figure 17 on page 28 and Figure 18 on page 29 illustrate the management processor utilization and cache write pending rate observed during the login storm/steady state operations for the virtual desktops.

- MPU utilization of all patterns did not rise above 47% during the logon storm, and stays below 46% during the steady workloads. This indicates that there is more than 53% headroom during steady workloads.
- Cache write pending of both patterns did not rise above 44% during the logon storm, and stays below 49% during the steady workloads. This indicates plenty of headroom for additional workloads.

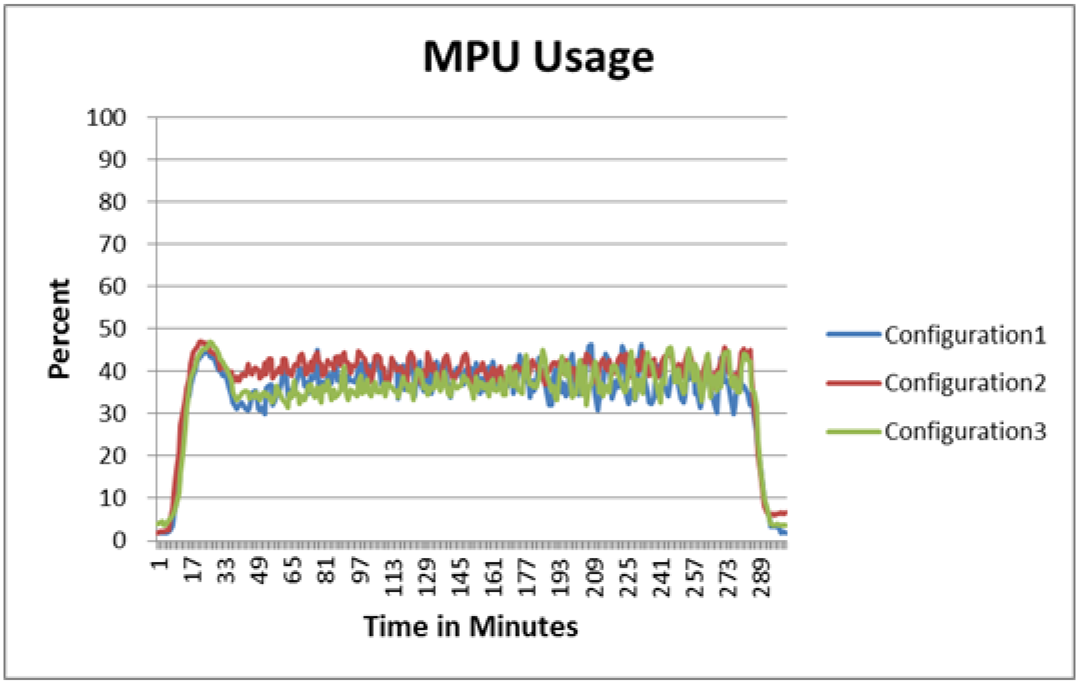

Figure 17

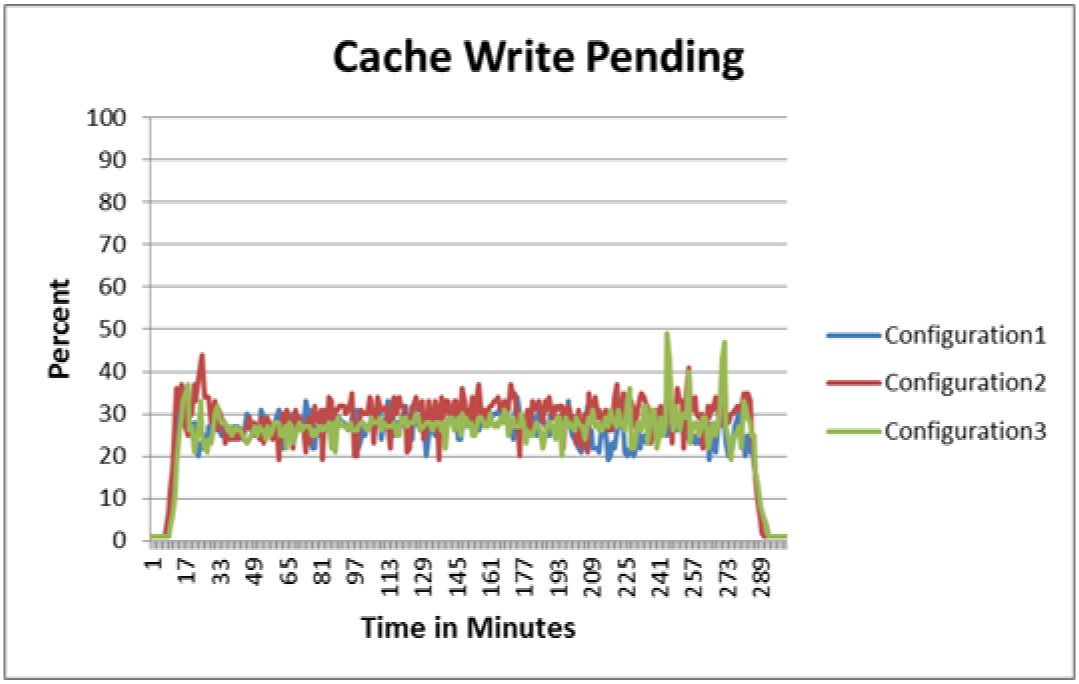

Figure 18

Physical Disk

Figure 19 on page 30 illustrates the average physical disk busy rate for the LDEVs used within the pool during the logon storm and steady state operations.

- **FMD**
  - The Configuration 1 disk busy rate on the FMD pool is approximately 93% throughout the logon storm and steady state workloads.
  - The Configuration 2 disk busy rate on Tier 1 (FMD) is approximately 99% throughout the logon storm and steady state workloads.
  - The Configuration 3 disk busy rate on the FMD pool is approximately 99% throughout the logon storm and steady state workloads.
- **HDD**
  - The Configuration 1 disk busy rate on the SAS 10k RPM HDD pool did not rise above 60%.
  - The Configuration 2 disk busy rate on Tier 2 (SAS 10k RPM HDD) peaked at 100% during the logon storm and remained below 25% during steady state workloads.

The disk busy rate related to the SAS 10k RPM HDD observed in Configuration 2 with active flash stayed at 100% versus the under 50% that was observed in Configuration1 in the 15 minute logon storm. This comes from massive IOPS generated by a half of the extreme power users residing on the SAS 10k RPM HDD in Configuration 2 with active flash due to the decrease of one FMD parity group against Configuration 1. This does not matter from the maximum performance utilization of 51% calculated by the active flash monitor and user experience across the entire logon time and application response described in “Storage Port Latency” on page 23.

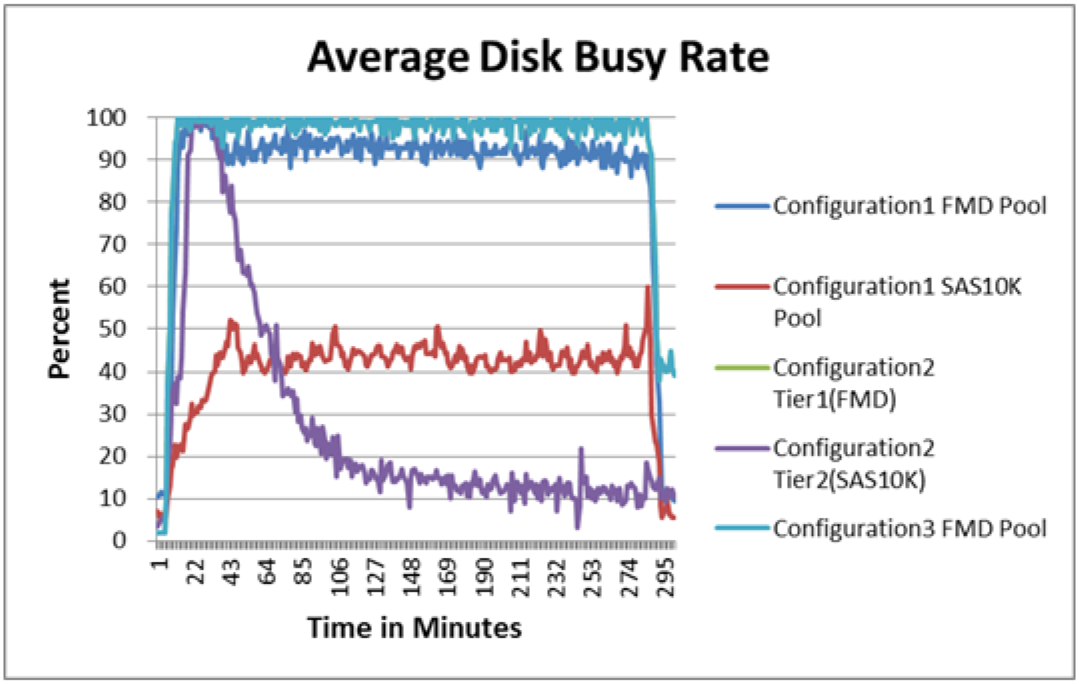

Figure 19

Storage Port Latency

Figure 20 on page 32 illustrates the latency observed on the front end Fibre Channel ports of Hitachi Virtual Storage Platform G600 during the logon storm or steady state operations.

- **Logon storm**
  - **Power users**
    - Average latency on ports used for the Configuration 1 extreme power users peaked at 1.1 milliseconds during the logon storm.
    - Average latency on ports used for the Configuration 2 extreme power users peaked at 2.1 milliseconds during the logon storm.
    - Average latency on ports used for the Configuration 3 extreme power users peaked at 1.6 milliseconds during the logon storm.
  - **Task users**
    - Average latency on ports used for the Configuration 1 task users peaked at 0.7 milliseconds during the logon storm.
    - Average latency on ports used for the Configuration 2 task users peaked at 3.5 milliseconds during the logon storm.
    - Average latency on ports used for the Configuration 3 task users peaked at 0.4 milliseconds during the logon storm.

- **Steady state**
  - **Power users**
    - Average latency on ports used for the Configuration 1 extreme power users peaked at 0.5 milliseconds during steady state.
    - Average latency on ports used for the Configuration 2 extreme power users peaked at 0.7 milliseconds during steady state.
    - Average latency on ports used for the Configuration 3 extreme power users peaked at 0.8 milliseconds during steady state.
  - **Task users**
    - Average latency on ports used for the Configuration 1 task users peaked at 1.9 milliseconds during steady state.
    - Average latency on ports used for the Configuration 2 task users peaked at 0.9 milliseconds during steady state.
    - Average latency on ports used for the Configuration 3 task users peaked at 0.4 milliseconds during steady state.

The storage port latency maximum of 2.1 msec related to the FMD observed in Configuration 1 with active flash on logon storm is higher than the maximum of 1.1 msec observed in Configuration 2 with Hitachi Dynamic Provisioning. This comes from the same reason that caused the differences in the disk busy rate illustrated in Figure 19 on page 30. The storage latency under 15 msec does not matter from the storage performance perspective. The storage port latency maximum of 3.5 msec related to the SAS HDD observed in Configuration 1 with active flash during the logon storm is higher than maximum of 0.7 msec observed in Configuration 1 with Hitachi Dynamic Provisioning.

This comes from the same reason that causes the differences in disk busy rate illustrated in Figure 19. The storage latency under 15 msec does not matter from storage performance perspective.

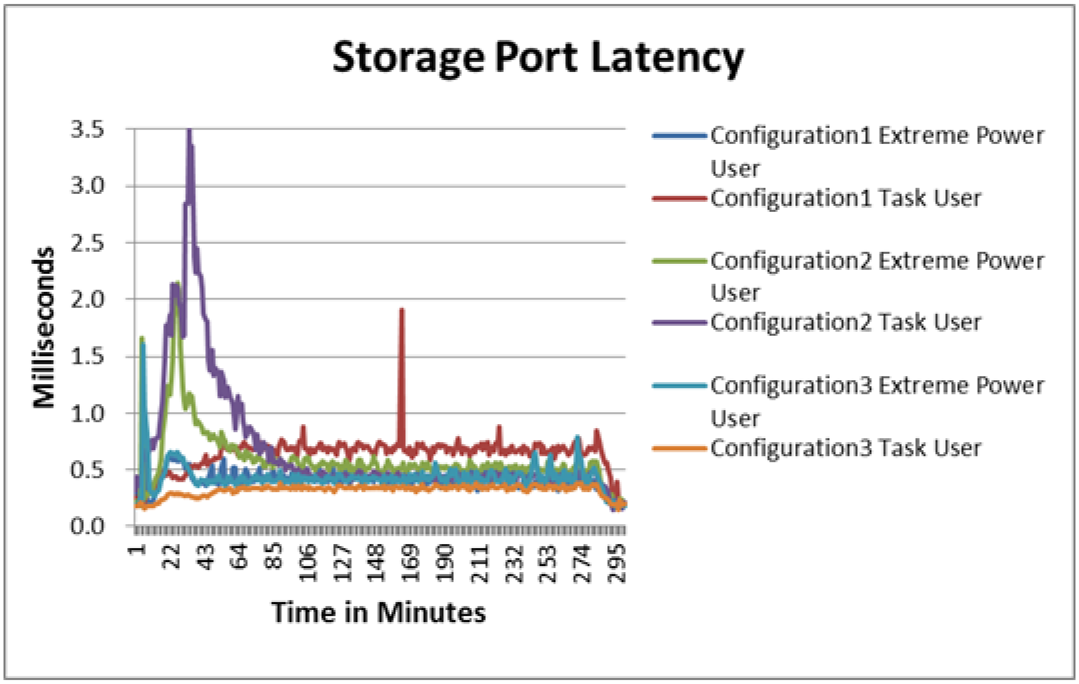
Storage Port Latency
Milliseconds
3.5
3.0
2.5
2.0
1.5
1.0
0.5
0.0
1
22
43
64
85
106
127
148
169
190
211
232
253
274
295
Time in Minutes
Configuration1 Extreme Power User
Configuration1 Task User
Configuration2 Extreme Power User
Configuration2 Task User
Configuration3 Extreme Power User
Configuration3 Task User

Figure 20

For More Information

Hitachi Data Systems Global Services offers experienced storage consultants, proven methodologies and a comprehensive services portfolio to assist you in implementing Hitachi products and solutions in your environment. For more information, see the Hitachi Data Systems Global Services website.

Live and recorded product demonstrations are available for many Hitachi products. To schedule a live demonstration, contact a sales representative. To view a recorded demonstration, see the Hitachi Data Systems Corporate Resources website. Click the **Product Demos** tab for a list of available recorded demonstrations.

Hitachi Data Systems Academy provides best-in-class training on Hitachi products, technology, solutions and certifications. Hitachi Data Systems Academy delivers on-demand web-based training (WBT), classroom-based instructor-led training (ILT) and virtual instructor-led training (vILT) courses. For more information, see the Hitachi Data Systems Services Education website.

For more information about Hitachi products and services, contact your sales representative or channel partner or visit the Hitachi Data Systems website.

Hitachi Data Systems

Corporate Headquarters
2845 Lafayette Street
Santa Clara, CA 96050-2639 USA
www.HDS.com community.HDS.com

Regional Contact Information
Americas: +1 408 970 1000 or info@hds.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hds.com
Asia Pacific: +852 3189 7900 or hds.marketing.apac@hds.com

AS-464-00. January 2016.