

Fusion-io® Virtualization Reference Architecture

Deploying Server and Desktop Virtualization on Fusion ioControl® Hybrid Storage, Cisco® UCS®, and VMware®



A Principled Technologies reference architecture commissioned by Fusion-io, Inc. (Revised)

TABLE OF CONTENTS

Table of contents	2
Executive summary	3
Solution considerations	4
Summary of main findings	5
Solution components	5
Cisco UCS	5
VMware vSphere 5.1	6
VMware Horizon View	6
Fusion ioControl Hybrid Storage Appliance	7
Cisco UCS Fusion ioDrive2 785GB adapter and Fusion ioTurbine software	9
Workloads	11
Login VSI workload	11
OLTP (DVD Store) workload description	12
Exchange (Microsoft LoadGen)	13
Solution architecture	14
Setting up Fusion ioControl	15
Provisioning storage	17
Setting up the host network and storage adaptor	19
Setting up Fusion ioTurbine	24
Virtual machine architecture	30
Validation and test results	30
Test scenario 1 – ioControl supporting all application workloads ..	30
Test scenario 2 – Scaling performance with zero footprint with Cisco UCS Fusion ioDrive 2 and Fusion ioTurbine for VMware	36
Using the ioControl user interface	41
Summing it all up	43
Appendix A – Detailed system configuration	44
Appendix B – Detailed virtual machine configuration	47
About Principled Technologies	48

EXECUTIVE SUMMARY

Introduction

Virtualizing your x86 servers and desktops with VMware technologies has the potential to make your business more agile, efficient and profitable. Mid-size business and small-to-medium enterprises benefit most when they can host multiple, mixed workloads within a common infrastructure – compute, storage, networking and virtualization. To support this virtualized environment of multiple, mixed application workloads requires:

- Consistent, predictable performance for mission and business critical applications
- Management of resource contention amongst applications
- Responding to performance storms without impact to other applications
- Management of solution based on business priorities
- Scaling performance and capacity with no disruption to the business

Combining the Fusion ioControl Hybrid Storage Appliance and Cisco Unified Computing System™ (UCS) with VMware® virtualization technologies (vSphere® and Horizon View) can create a solution for your business that properly runs multiple virtualized applications, including business applications and virtual desktops, while delivering the consistent predictable performance needed in a cost-effective solution. In the labs at Principled Technologies, we configured the Fusion-io and Cisco UCS solution to run steady state business-critical applications, such as online transaction processing (OLTP) and Microsoft® Exchange, with a demanding VDI workload. This reference architecture shows you how to implement these Fusion-io, Cisco, and VMware components for your data center and validates how the solution performed in our testing.

We tested the reference architecture solution by running a mixed workload environment while conducting a 600-count boot storm, or influx of initiating VDI user sessions. For our testing, the measure of success depended on the responsiveness of OLTP and Exchange workloads without interruption or degradation related to I/O contention during the most demanding conditions, such as a virtual desktop boot storm.

Objective

The objective of this reference architecture (RA) is to validate and quantify the advantages and benefits of a solution based on Fusion ioControl Hybrid Storage with Cisco UCS/Nexus Switches and VMware software for a mixed workload environment

and how to extend storage performance without increasing spindle count or data center footprint.

This RA also documents configuration best practices for the Fusion ioControl Hybrid Storage Appliance in an environment running multiple virtualized applications. For the full report which includes detailed test infrastructure installation and deployment steps, see www.principledtechnologies.com/Fusion-io/ioControl_reference_architecture_unabridged_0714_v3.pdf.

Audience

The intended audience of this reference architecture is IT planners and decision-makers within medium-sized businesses and small-to-medium enterprises. Solution architects and value-added resellers may also find this reference architecture useful in deciding how to plan solutions for their customers. This reference architecture is not intended to be a comprehensive guide to deployment and configuration for every component in this solution.

SOLUTION CONSIDERATIONS

Proper design of the solution for a mixed application VMware environment is important because the server infrastructure needs abundant resources to handle the high density VDI user count in addition to running critical workloads, such as OLTP and Exchange. Your infrastructure should be provisioned to handle steady-state performance requirements along with bursts of data traffic. The shared storage needs to support the random IOPS that result from mixed applications along with surges of I/O traffic, such as VDI boot storm and virus scans.

The solution needs to meet the following requirements:

- Guaranteed I/O resources to mission-critical applications
- Support for a high density of virtual desktops per server
- Consistent and predictable user performance even during peak application workload conditions
- Simple configuration, management
- Out of the box reporting and monitoring
- Low latency and high bandwidth for clustering, provisioning, and storage interconnect networks

SUMMARY OF MAIN FINDINGS

We discovered multiple advantages in the reference architecture solution.

- **Guaranteed I/O performance:** By implementing ioControl QoS (Quality of Service) policies in a mixed workload environment, Cisco UCS and Fusion ioControl ensured mission-critical and business-critical storage tiers received guaranteed I/O under the most demanding portions of the mixed workload test. Our OLTP and Exchange workloads remained at optimum performance, including during a 600 virtual desktop boot storm.
- **Consolidation and VM density:** The reference architecture solution delivered performance without sacrificing density, supporting a mixed back-office workload and 600 concurrent virtual desktop users without session degradation.
- **Out-of-the-box reporting and monitoring:** ioControl came equipped with easy-to-use monitoring and reporting tools useful to a system administrator.
- **Simple deployment:** We unboxed, cabled, and configured the management interface of the Fusion ioControl storage appliance in less than half an hour.
- **Performance boost:** The solution provided uninterrupted performance and response times to all workloads tested. The addition of Fusion ioTurbine software for VMware along with Fusion ioMemory within UCS added power to our OLTP workload. This resulted in a 6X OLTP workload increase without any performance degradation or increase in storage footprint.
- **Flexible scalability:** We tested with an ioControl n5-150 appliance and a single chassis of eight UCS B-Series servers. Both solution components have additional available resources to scale that do not require a redesign of the solution or of the infrastructure.

SOLUTION COMPONENTS

Cisco UCS

The following Cisco components were utilized:

- Cisco UCS 6200 Series Fabric Interconnects
- Cisco UCS 2200 Series Fabric Extenders
- Cisco UCS 5108 Blade Server Chassis
- Cisco UCS B-series Blade Servers for applications and virtual desktop hosting

Careful planning and management are required to maintain a fully redundant network. Due to its unified fabric model, Cisco UCS is a solution that delivers automatic network failover right out of the box—you do not need to configure any hardware or software or purchase additional pieces.

You can manage your entire Cisco UCS infrastructure from anywhere with the Cisco UCS Manager software, which is a highly available, single location for systems management. UCS Manager can apply all firmware versions and updates to existing servers, new nodes, and other hardware with the click of a mouse. Cisco UCS Manager offers support for all current and previous generation UCS hardware models, including both blade and rack servers.

As your network scales out, the Cisco UCS Central software consolidates all management systems together into one dashboard. Regardless of the hardware in your UCS deployment, you can manage it the same way from one unified interface. This lets you optimize your infrastructure to match your workloads without sacrificing the management capabilities of UCS Manager or adding additional management products.

The Cisco UCS chassis contains eight half-height bays in only 6U, provides up to 160 Gbps to each chassis, and is capable of providing 80 Gbps of bandwidth to any blade. Cisco UCS allows customers to share bandwidth across the entire domain, ensuring that servers that need the bandwidth will get it without requiring switches in every chassis. The chassis requires no integrated switches, as switching is handled upstream by the redundant Fabric Interconnect modules for the entire domain.

Our configuration consisted of a Cisco UCS 5108 Blade Server Chassis, connected to a redundant pair of Fabric Interconnects, four links per fabric. This provided 40 Gbps per fabric, for a total of 80 Gbps to the chassis.

VMware vSphere 5.1

We used the data center virtualization software VMware vSphere 5.1 in our testing—the latest version of the application. VMware vSphere can virtualize CPU, RAM, hard disk, and network resources for fully functional virtual desktops that can run insulated and condensed operating systems and applications. The high-availability features of VMware vSphere coupled with VMware Distributed Resources Scheduler (DRS) and vMotion can allow for migration of virtual desktops from one VMware vSphere server to another with little or no effect on end-users.

VMware Horizon View

VMware Horizon View is a desktop virtualization solution built on and tightly integrated with VMware vSphere. This software allows customers to extend the value of

VMware infrastructure and its enterprise-class features, such as high availability, disaster recovery, and business continuity. VMware Horizon View includes enhancements to the end-user experience and IT control. Support for VMware vSphere uses the latest functions of the leading cloud infrastructure platform for highly available, scalable, and reliable desktop services.

Fusion ioControl Hybrid Storage Appliance

The Fusion ioControl Hybrid Storage Appliance (see Figure 1) combines PCIe flash memory with disk. ioControl's integrated storage QoS software allows you to provision, prioritize, and control shared storage performance.



Figure 1: The Fusion ioControl Hybrid Storage Appliance.

The ioControl n5 series storage appliances are available in several base models that differ in flash and disk capacity, all of which have the ability to add additional flash and disk to scale performance and capacity independently of one another. All ioControl n5 models come with software to manage and monitor the array. Fusion-io provided an n5-150 Hybrid Storage Appliance for this RA, with the following:

- 2.4TB ioMemory flash capacity and 48TB raw disk capacity
- A wide range of software to support array specific features, including Storage QoS, Dynamic Data Placement, Phased Data Reduction, Snapshot, and Replication

For more information on Fusion ioControl Hybrid Storage Appliances, see www.fusionio.com/products/iocontrol/.

Fusion ioControl can provision and prioritize storage resources with Quality of Service (QoS) performance policies that are grouped into three different service levels. QoS policies are assigned to each volume independently and define minimum IOPS, throughput and maximum latency for each application. A QoS policy is assigned to a

volume at creation and can be changed in real time as needs change. ioControl QoS was a key advantage in the deployment of the mixed workload reference architecture.

Figure 2 provides detailed information on QoS levels for the ioControl array and Figure 3 shows a representative mixed workload environment in which guaranteed service levels have been assigned.

	Service Level	IOPS Target	Bandwidth Target (MB/sec)	Latency Target (ms)
Policy 1	Mission-critical	75,000	750	10
Policy 2	Business-critical	30,000	375	20
Policy 3	Business-critical	15,000	150	40
Policy 4	Non-critical	7,500	75	100
Policy 5	Non-critical	1,500	37	250

Figure 2: ioControl QoS policy characteristics.

ioControl quality of service

Prioritized performance under severe load

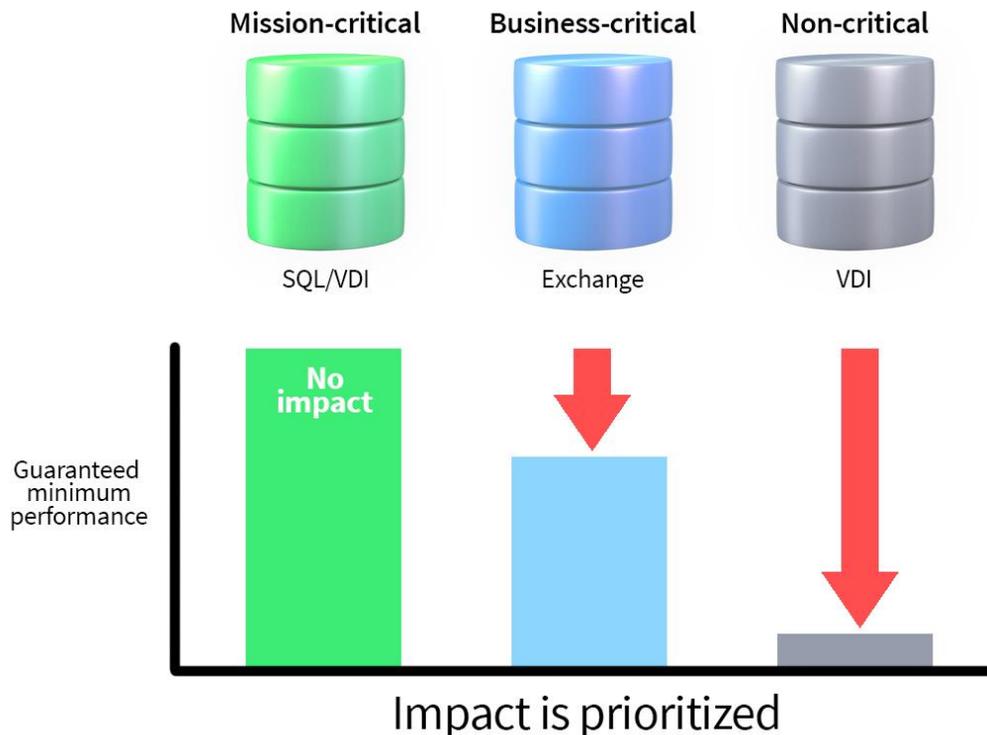


Figure 3: Guaranteed performance according to QoS service levels.

Data protection

According to Fusion-io, “ioControl provides a number of data protection tools built into their Hybrid Storage Appliance, such as snapshots, cloning, and asynchronous replication. ioControl snapshots allow for timely data backups and significantly reduce RPO and RTO times. In addition, ioControl asynchronous replication is a key component for data migrations and business continuity solutions. ioControl is supported with major backup and recovery products, including Veeam®. Fusion-io is a Veeam Technology Alliance Partner and ioControl has been certified with Veeam Backup and Replication.

“Veeam enables the Always-On Business™ by providing solutions that deliver Availability for the Modern Data Center™. Veeam recognizes the challenges in keeping a business up and running at all times and addresses them with solutions that provide high-speed recovery, data loss avoidance, verified protection, risk mitigation and complete visibility. Veeam Backup & Replication™ leverages technologies that enable the modern data center, including VMware vSphere, Microsoft Hyper-V, to help organizations meet recovery time and point objectives (RTPO™), save time, mitigate risks, and dramatically reduce capital and operational costs. Veeam Backup Management Suite™ provides all of the benefits and features of Veeam Backup & Replication along with advanced monitoring, reporting and capacity planning for the backup infrastructure.

“Together, Veeam and Fusion-io provide a cost-effective solution using the latest PCIe-based solid-state shared storage, backup, recovery and replication technologies for virtual workloads – guaranteeing predictable performance for business critical applications, Quality of Service (QoS) for VM datastores and faster backup and recovery times.”

For more information, visit go.veeam.com/fusion-and-veeam.html.

Cisco UCS Fusion ioDrive2 785GB adapter and Fusion ioTurbine software¹

Combining Fusion ioTurbine intelligent caching software with Fusion ioMemory on a UCS B-series Blade Server provided a performance extension to ioControl without increasing the datacenter footprint. This extension complements ioControl shared storage performance by applying a low latency, high-performance, high-capacity read-cache to Cisco UCS blade servers to immediately improve performance and relieve I/O

¹ ioControl SPX combines ioControl Hybrid Storage, ioTurbine for VMware, and ioMemory into a single solution sold and supported by Fusion-io under a single part number. See www.fusionio.com/data-sheets/iocontrol-n5-spx.

bottlenecks to the most performance intensive application in the mixed workload. This is an optional component that we chose to test as part of this reference architecture.

Storage management with ioControl vCenter plug-in

A plug-in is available with ioControl that allows VMware vCenter to manage ioControl. According to Fusion-io, the ioControl vCenter plug-in “simplifies the provisioning and management of Fusion ioControl shared storage.” For an image of the user interface, see Figure 4 below.

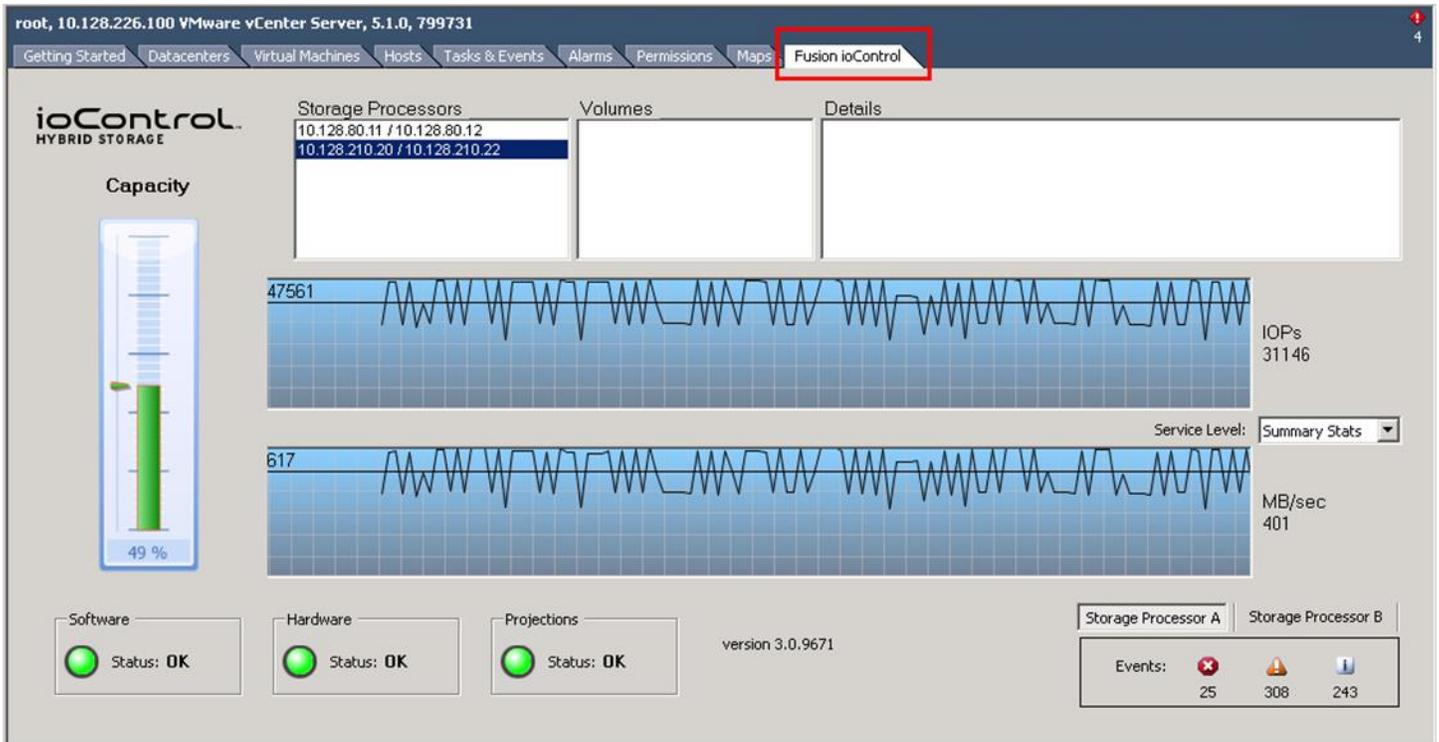


Figure 4: Storage management with the ioControl vCenter plug-in.

WORKLOADS

For the first test, we set up a mixed workload consisting of common virtualized business applications and VDI to run in concert over a 90-minute period. We ran OLTP (SQL) and Microsoft LoadGen (Exchange) in a steady state condition and then kicked off a 600 virtual desktop boot storm. After the 600 virtual desktops reached an idle state we started the VDI workload.

For the second test, we configured ioTurbine intelligent caching software with a Cisco UCS Fusion ioDrive2 mezzanine card to cache the SQL database volumes and repeated the first test with a heavy OLTP workload. Specific parameters for each workload are as follows:

Login VSI workload

We used Login VSI to measure the Fusion ioControl and Cisco UCS common virtualized business applications and VDI workload capabilities. The Medium workload is the standard work load on the Login VSI v4 benchmark; it simulates a typical business user's work behavior by using productivity-related applications. We added users incrementally every 20 seconds until all users were logged in. Login VSI measures response time for several items, such as how long it takes to open a text file or compress a file. These response times degrade as the number of active users attach to the pool of desktops.

From the response times collected, Login VSI v4 initially calculates a system baseline called VSIBase, which determines—in milliseconds—the response time on a system while there is little to no load. The VSIBase is later used to calculate the VSImax v4 threshold, which is the point the system under test will be considered saturated and subsequently degrade end-user performance and experience. Once this point of saturation is reached, Login VSI determines the VSImax v4 result, which is the maximum number of active sessions on a system before saturation, degrading end-user experience. A successful test will not reach VSImax v4.

For more information, see section 1.2.2 (Workloads: Medium workload) at www.loginvsi.com/documentation/index.php?title=All_settings.

VM parameters

- 24GB master desktop image
 - Virtual desktop OS: Windows 7 Enterprise edition
 - 1GB RAM
 - 1 vCPU
- Desktop application environment
 - Microsoft Office 2010
 - Internet Explorer
 - Adobe Reader
 - Adobe Flash

- Java/FreeMind
- VMware View 4GB change disk
- About 10 to 15 IOPS per desktop at steady state, 100 percent concurrency
 - Number of sessions: 600
 - Launch interval: 20 seconds
 - Login VSI Workload: Medium
 - Number of launchers: 24

Output

- Login VSI response time, depicted in graphs (VSI_{max} = pass/fail)

OLTP (DVD Store) workload description

To simulate an OLTP workload, we used DVD Store, which is an open-source application that models an online DVD store where customers log in, search for titles, and purchase movies. Higher thread count and lower think times increase CPU utilization and IOPS. For more information on DVD Store, please see en.community.dell.com/techcenter/extras/w/wiki/dvd-store.aspx.

Medium OLTP (no ioTurbine)

Testing parameters

- Database size: 20 GB
- Number of databases: one per VM
- Number of threads per VM: 12
- Actual test run time: 90 minutes
- Think time: 100 milliseconds
- Warm-up time: 15 minutes
- Ramp rate: 10 minutes

Additional VM parameters

- See Figure 32 for VM parameters
- Each SQL VM was built on a single VMDK file with volume segregation
- Each SQL VM was configured as per best practices (C:\binaries, E:\SQL Database, L:\SQL Logs)
- ioTurbine was configured to utilize guest-based caching
- ioTurbine guest-based caching was applied to the E:\ drive on each of the four VMs, which contains both the 50 GB test database plus the tempdb database
- ioTurbine guest caching was not applied to the SQL .ldf files in the L:\ drive

Output

- Orders per minute (OPM)

Heavy OLTP (with ioTurbine)

The medium OLTP workload was repeated with the think time reduced to 0ms.

Exchange (Microsoft LoadGen)

The Microsoft Load Generator 2010 (LoadGen) benchmark performs tasks to simulate a typical user generating mail activity on Microsoft Exchange Server 2010. We divided our users into four mailbox databases, each with 100 users, for a total of 400 users/mailboxes. We ensured that performance was within acceptable RPC average latency thresholds. For more information on Microsoft Load Generator 2010, please see www.microsoft.com/downloads/details.aspx?FamilyId=DDEC1642-F6E3-4D66-A82F-8D3062C6FA98&displaylang=en.

Parameters

- Simulation: Two days
- Time: 90-minute duration
- 100 users per database (400 total)
- Client: Outlook 2007 Online
- Action profile: Outlook_200
- Mailbox profile: 100 MB

VM parameters

- See Figure 32 for VM parameters

Output

- RPC latency²

² Microsoft states that successful testing requires RPC latency to remain below 50 ms. For more information, see [technet.microsoft.com/en-us/library/bb508871\(v=exchg.80\).aspx](http://technet.microsoft.com/en-us/library/bb508871(v=exchg.80).aspx).

SOLUTION ARCHITECTURE

Figure 5 presents an overview of the components from our solution and how they connected to one another.

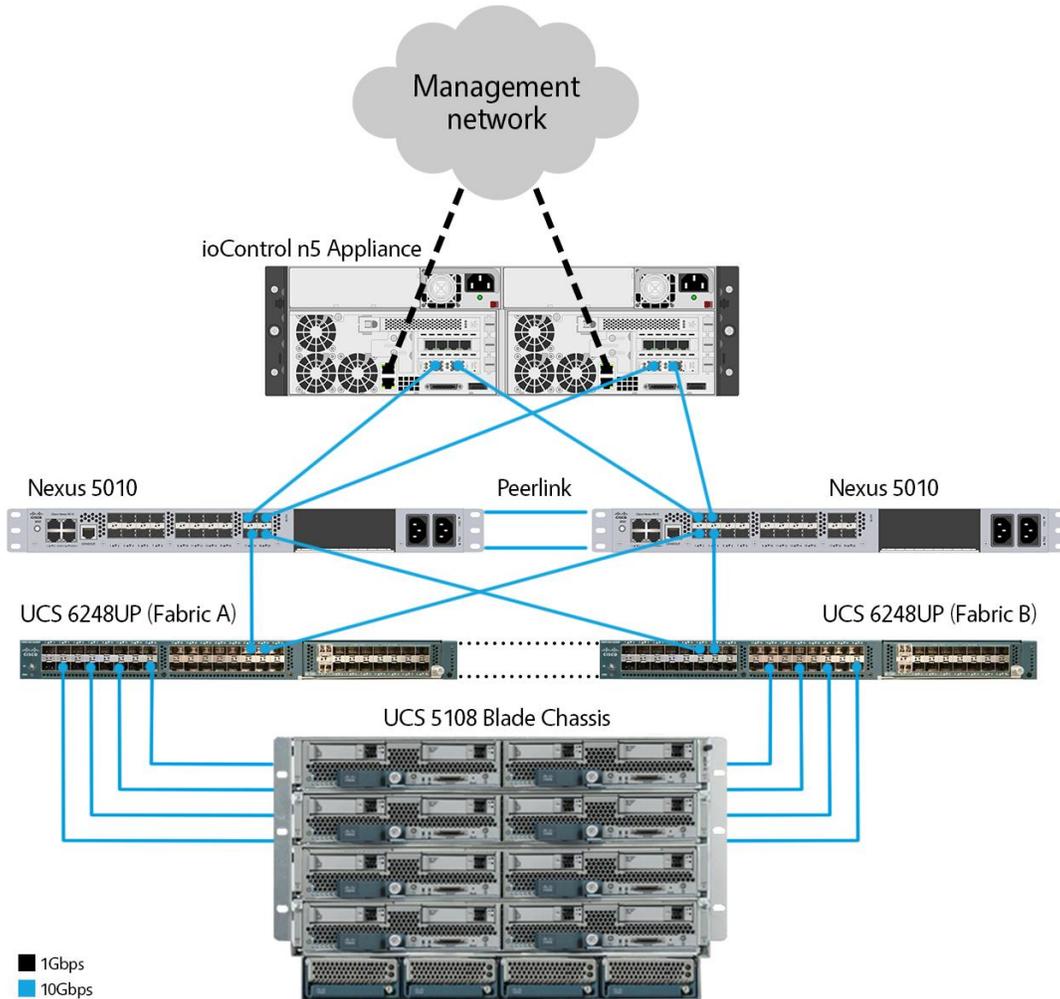


Figure 5: Overview of solution components.

We installed and configured VMware vSphere 5.1 on all eight of the servers. The three infrastructure servers were Cisco B200 M2 blades and the remaining five application servers were Cisco B200 M3 blades. In the SQL/Exchange host c1-s7, we installed one Cisco UCS Fusion ioDrive2 mezzanine card. Figure 6 details configuration for each server.

Server name	Model	Processors	Memory (GB)	Adaptors	Functional role
Infra1	Cisco B200 M2	Intel Xeon X5670	96	M171KR-Q QLogic CNA	Infrastructure host
Infra2	Cisco B200 M2	Intel Xeon X5670	96	M171KR-Q QLogic CNA	Infrastructure host
Infra3	Cisco B200 M2	Intel Xeon X5670	96	M171KR-Q QLogic CNA	Infrastructure host
c1-s4	Cisco B200 M3	Intel Xeon E5-2690	256	Cisco UCS VIC 1240	VDI host
c1-s5	Cisco B200 M3	Intel Xeon E5-2690	256	Cisco UCS VIC 1240	VDI host
c1-s6	Cisco B200 M3	Intel Xeon E5-2690	256	Cisco UCS VIC 1240	VDI host
c1-s7	Cisco B200 M3	Intel Xeon E5-2690	256	Cisco UCS VIC 1240, FIO ioDrive2	SQL and Exchange host
c1-s8	Cisco B200 M3	Intel Xeon E5-2690	256	Cisco UCS VIC 1240	VDI host

Figure 6: Detailed server configuration.

Setting up Fusion ioControl

We powered on ioControl and connected a notebook to the management network interface on storage processor A. The factory default address is 192.168.100.100/24. After logging on, we edited the iSCSI and management addresses. Figure 7 shows the ioControl user interface (UI) as we configured network information.

Welcome, admin | [Log Out](#) System Health OK | [Help](#)

ioControl Settings

System Info | Event Log | Host Access Groups | Long Running Tasks | **Network Addressing** | Network Services | Power ▾

System Network Interfaces

Name	State	Storage Processor	Mode	Address	Speed	
Management 1	Up	B	Static	172.16.0.16	1 Gb	Edit
Management 2	Up	B	Static	192.168.100.200	1 Gb	Edit
Slot 1 Port 1	Down	B	Disabled		1 Gb	Edit
Slot 1 Port 2	Down	B	Disabled		1 Gb	Edit
Slot 1 Port 3	Down	B	Disabled		1 Gb	Edit
Slot 1 Port 4	Down	B	Disabled		1 Gb	Edit
Slot 2 Port 1	Up	B	Static	192.168.0.102	10 Gb	Edit
Slot 2 Port 2	Up	B	Static	192.168.0.103	10 Gb	Edit
Management 1	Up	A	DHCP	172.16.12.210	1 Gb	Edit
Management 2	Up	A	Static	192.168.100.100	1 Gb	Edit
Slot 1 Port 1	Down	A	Disabled		1 Gb	Edit
Slot 1 Port 2	Down	A	Disabled		1 Gb	Edit
Slot 1 Port 3	Down	A	Disabled		1 Gb	Edit
Slot 1 Port 4	Down	A	Disabled		1 Gb	Edit
Slot 2 Port 1	Up	A	Static	192.168.0.100	10 Gb	Edit
Slot 2 Port 2	Up	A	Static	192.168.0.101	10 Gb	Edit

iSCSI targets

FUSION-IO
© 2013 Fusion-io, Inc

Figure 7: Locating system network interface and addresses with the ioControl UI.

Figure 8 shows where the specific network configuration occurs in the ioControl User Interface.

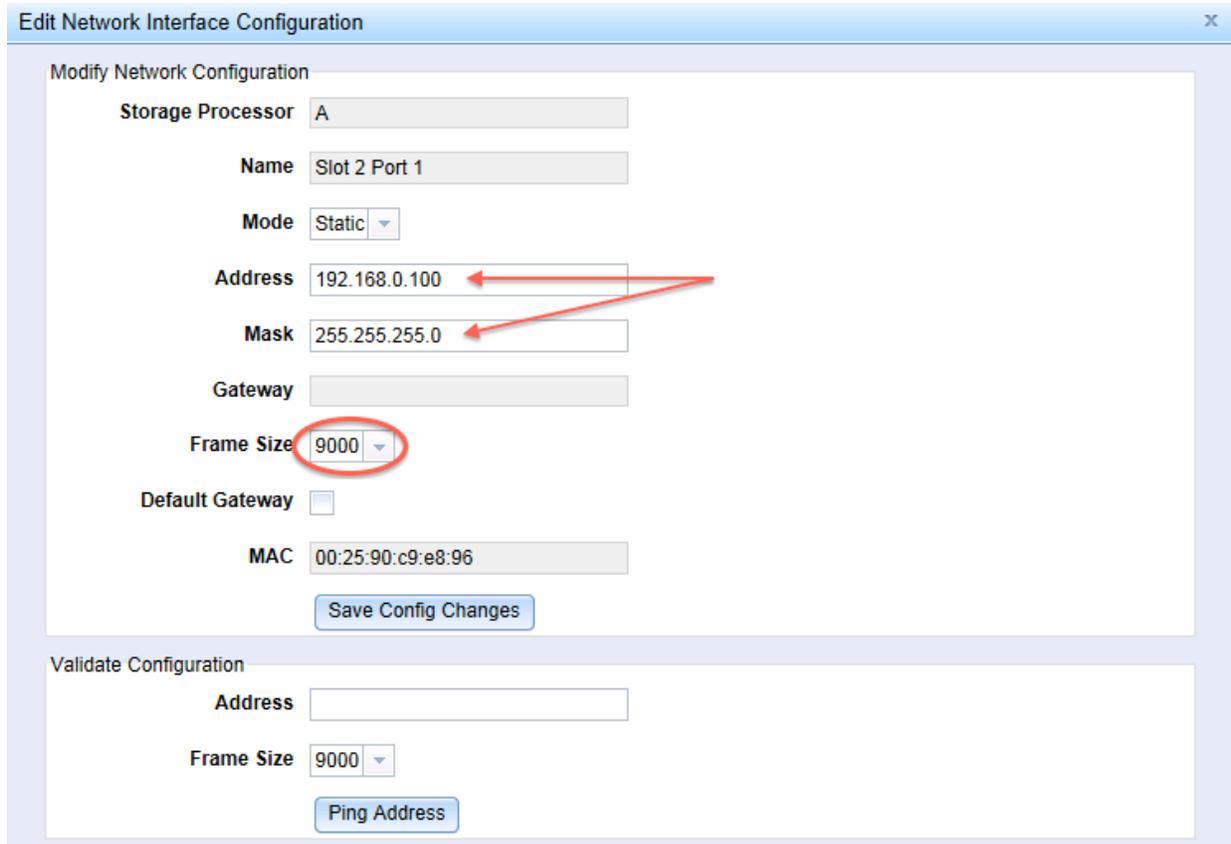


Figure 8: Editing network interface configuration with ioControl UI.

We then connected the four 10Gbps links to our Cisco Nexus 5010 Switches and two 1Gbps RJ-45 uplinks to our management network.

Provisioning storage

ioControl hosted all workload VM datastores and we sized each virtual desktop LUN as per VMware best practices.³ We started by classifying our applications as Mission-Critical, Business-Critical and Non-Critical based on ioControl QoS guaranteed minimum IOPS, throughput, and not-to-exceed latency policy values. Based on a typical mixed-workload business infrastructure, we assigned Microsoft SQL and 150 high-end VDI users as Mission-Critical, Microsoft Exchange and 300 mid-range VDI users as Business-Critical and remaining 150 VDI users as Non-Critical. We provisioned LUNs with the ioControl GUI. See Figure 9 for detailed storage information.

³ www.vmware.com/files/pdf/view/Server-Storage-Sizing-Guide-Windows-7-TN.pdf

LUN/NAS Name	Size (GB)	Host access	Policy	Hosting role	Storage processor
Exchange	450	c1-s7	Business-critical: 30,000 IOPS, 375 MB/sec bandwidth, 20ms latency	LoadGen Exchange Servers	A
DVD-SQL	500	c1-s7	Mission-critical: 75,000 IOPS, 750 MB/sec bandwidth, 10ms latency	OLTP SQL servers	A
VDT-A	400	c1-s4, c1-s5, c1-s6, c1-s8	Mission-critical: 75,000 IOPS, 750 MB/sec bandwidth, 10ms latency	150 virtual desktops	B
VDT-B	400	c1-s4, c1-s5, c1-s6, c1-s8	Business-critical: 30,000 IOPS, 375MB/Sec bandwidth, 20ms latency	150 virtual desktops	B
VDT-C	400	c1-s4, c1-s5, c1-s6, c1-s8	Business-critical: 30,000 IOPS, 375MB/Sec bandwidth, 20ms latency	150 virtual desktops	B
VDT_D	400	c1-s4, c1-s5, c1-s6, c1-s8	Non-critical: 7,500 IOPS, 75MB/sec bandwidth, 100ms latency	150 virtual desktops	A
NAS01	1 TB	Infra1-3	N/A	Test harness datastore, generic NAS	N/A

Figure 9: Detailed storage configuration.

We configured LUNS, assigned QoS policies, and configured iSCSI access policies in the ioControl UI. Figure 10 shows LUN configuration with ioControl.

ioControl Provision

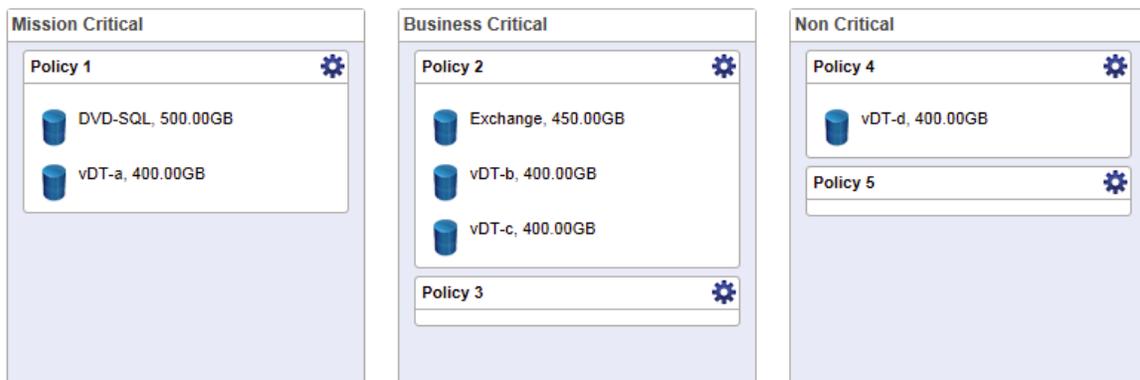


Figure 10: Configuring LUNs with the ioControl UI.

Figure 11 shows configured access groups with ioControl.

ioControl Settings

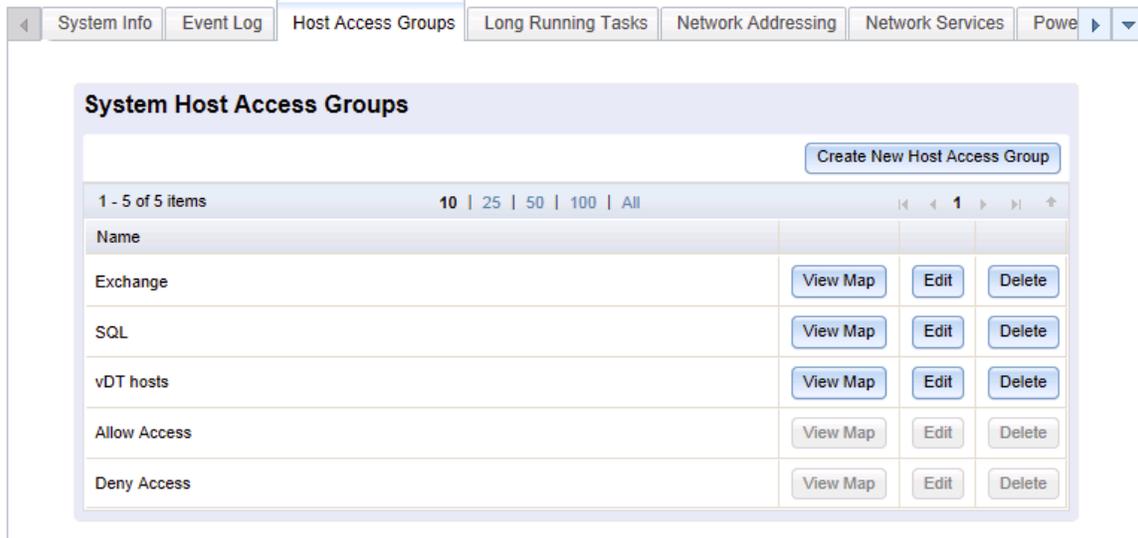


Figure 11: Configuring storage host access groups with ioControl UI.

Figure 12 shows how we configured initiators in a storage access group.

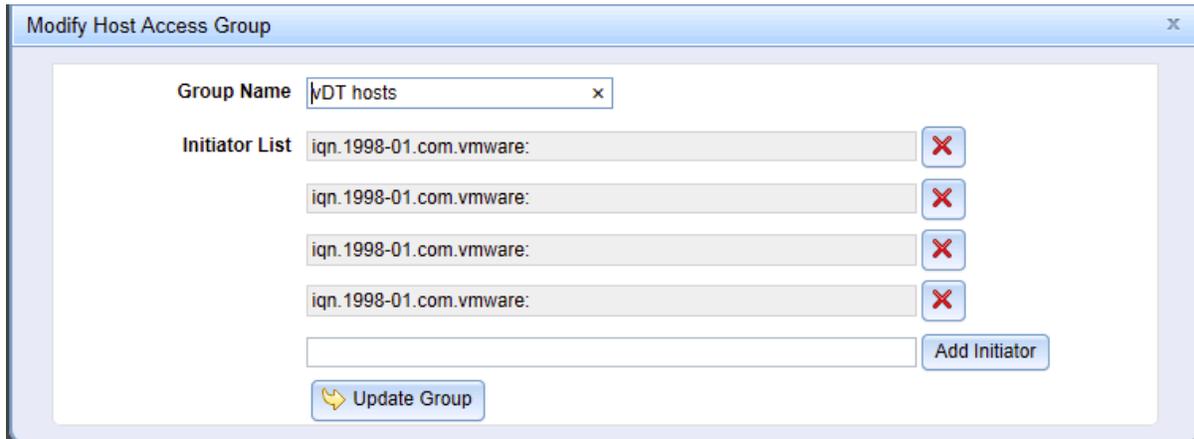


Figure 12: ioControl GUI: Storage host access groups

Setting up the host network and storage adaptor

Each Cisco UCS B200 M3 server has two 20Gbps vNICs; one 20Gbps link is on Fabric A and the other is connected to Fabric B. We set up a VMware VMKernel interface on each fabric to ensure redundancy and multipathing across both fabrics. To accomplish this, we needed to specify the override vSwitch setting and specifically bind iSCSI-a to vNIC0, and iSCSI-B to vNIC1. We set up Jumbo frames on both vNICs and

tested connectivity. Figure 13 shows how we accessed the VMware host network configuration with the UI.

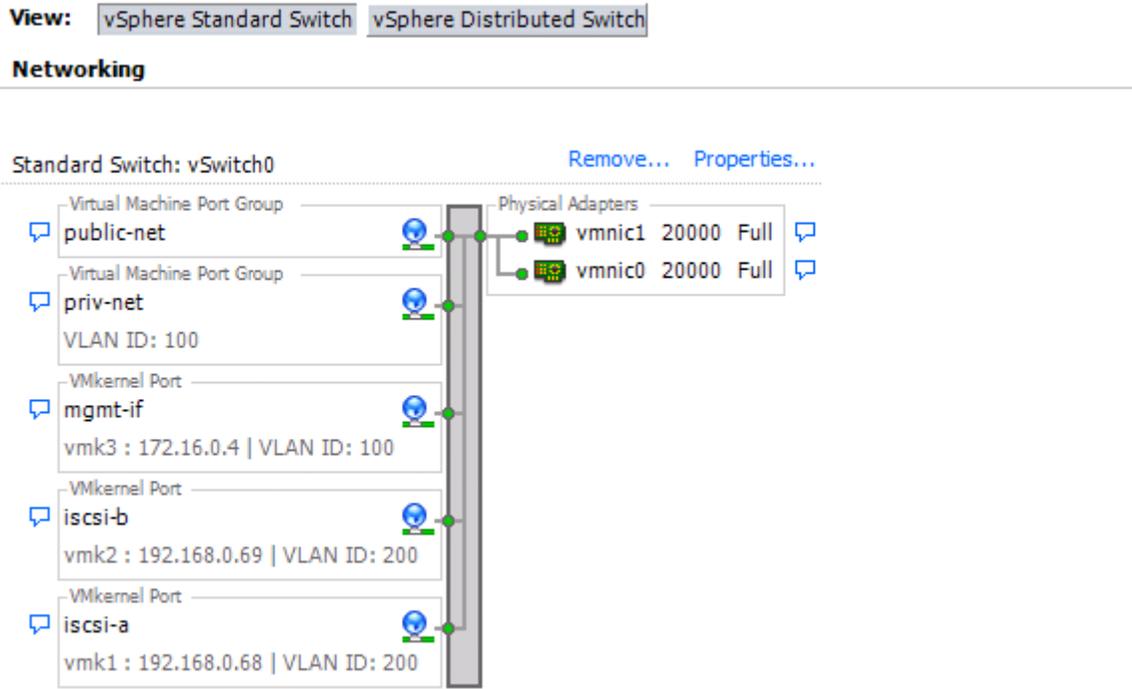


Figure 13: Configuring VMware host network.

Figures 14 and 15 show where we configured the iSCSI settings.

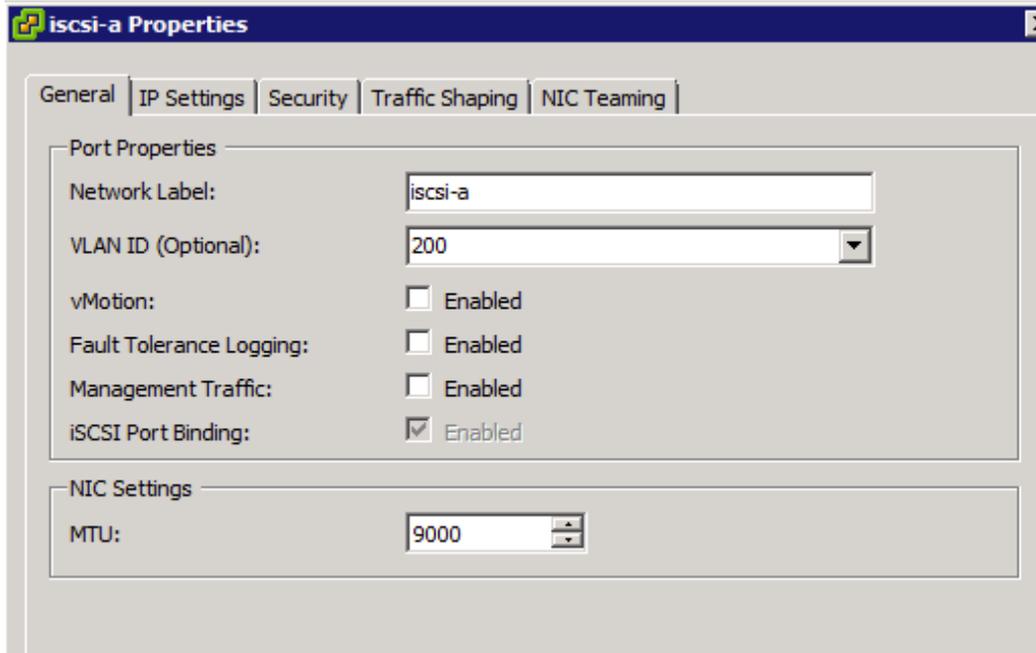


Figure 14: Configuring iSCSI-a with the VMKernel interface.

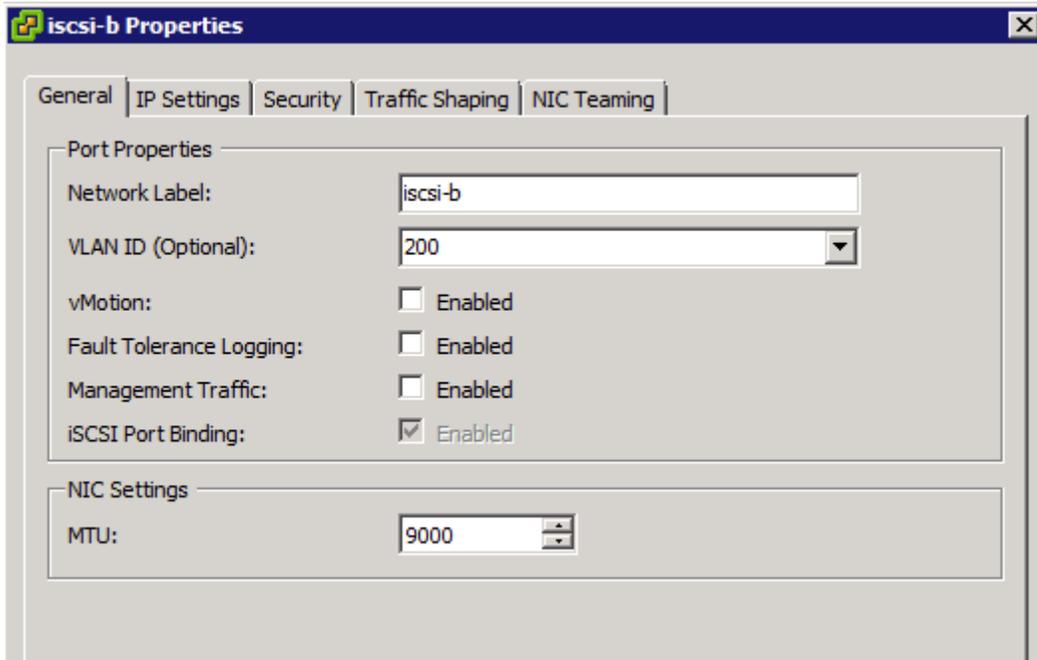


Figure 15: Configuring iSCSI-b with the VMKernel interface.

Figure 16 and 17 show where to override NIC teaming for iSCSI-a and iSCSI-b.

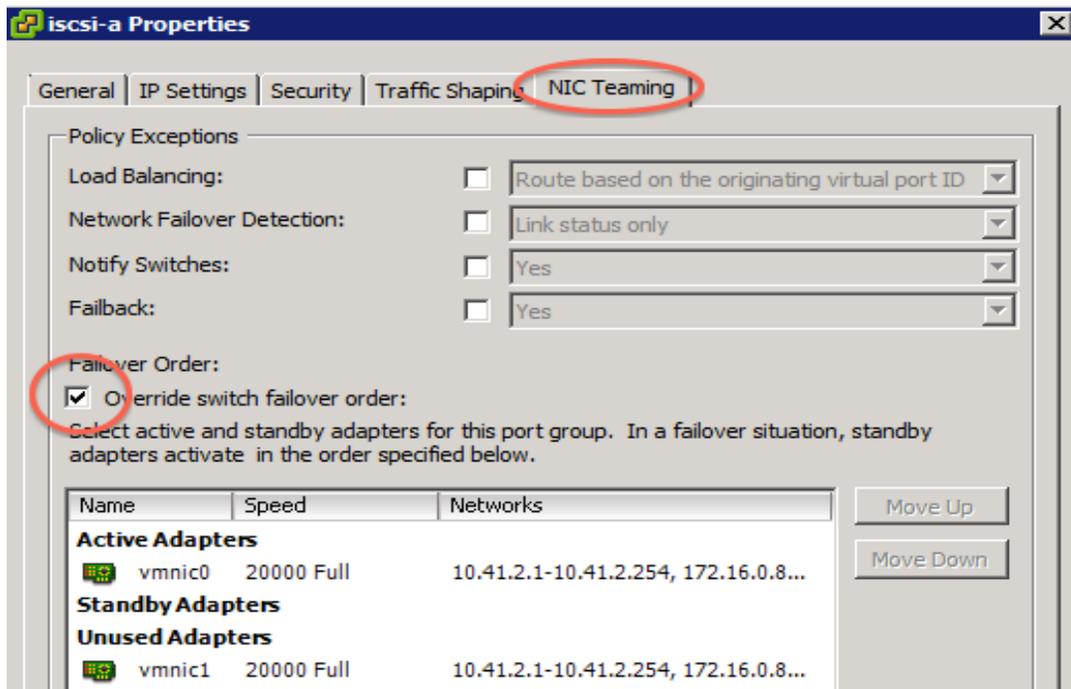


Figure 16: Overriding NIC teaming for iSCSI-a.

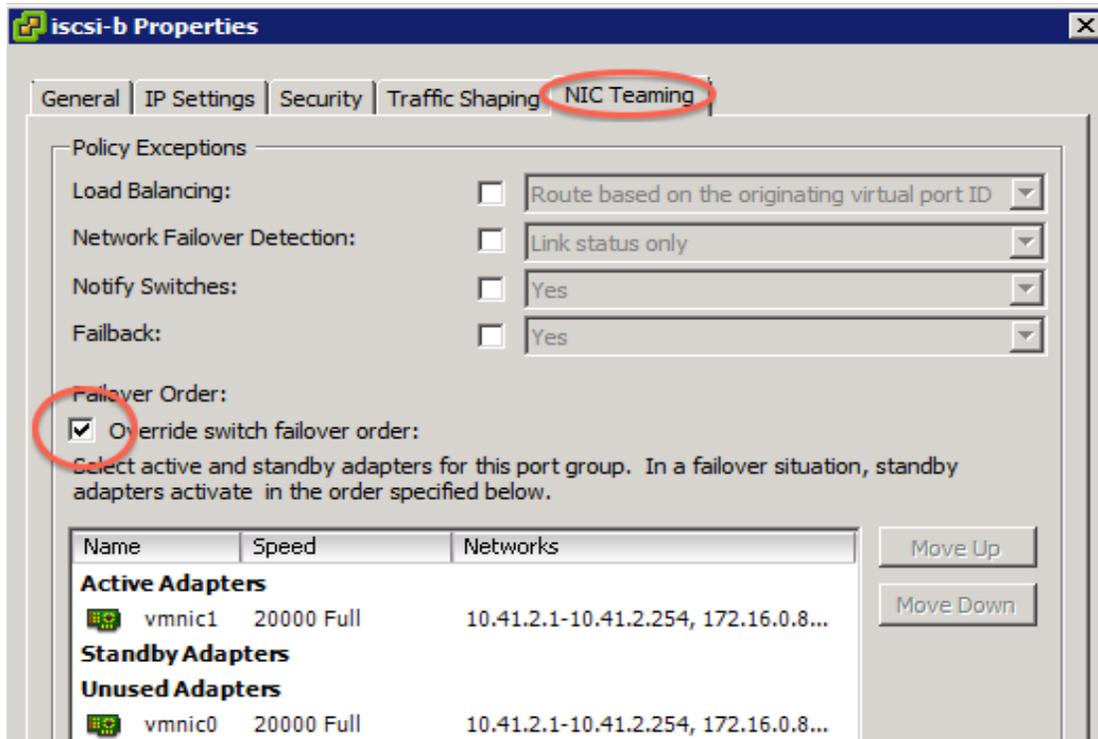


Figure 17: Overriding NIC teaming for iSCSI-b.

Figure 18 shows the VMKernel port bindings with the host iSCSI Initiator interface.

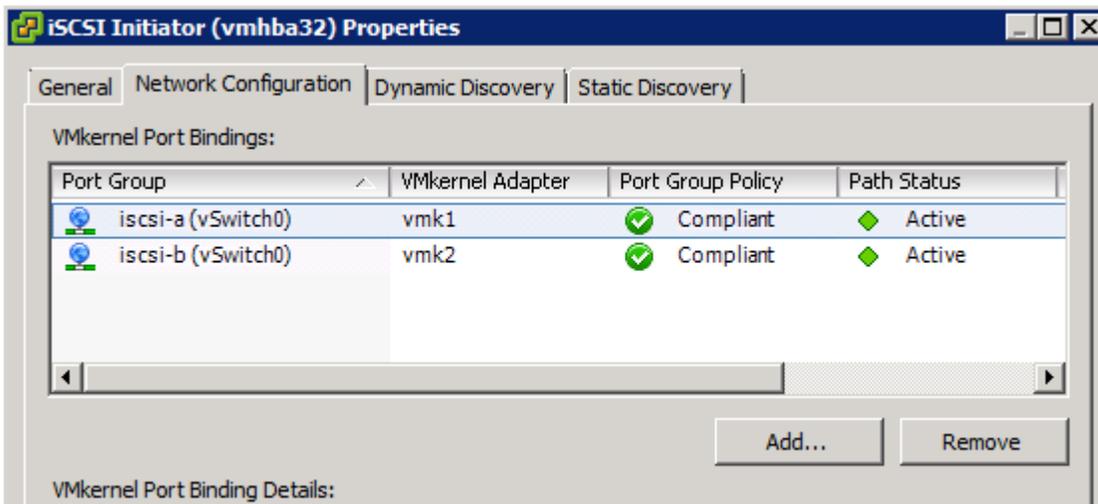


Figure 18: Using the host iSCSI Initiator interface to confirm VMKernel port binding.

Figure 19 shows how we discovered iSCSI targets with the host iSCSI Initiator.

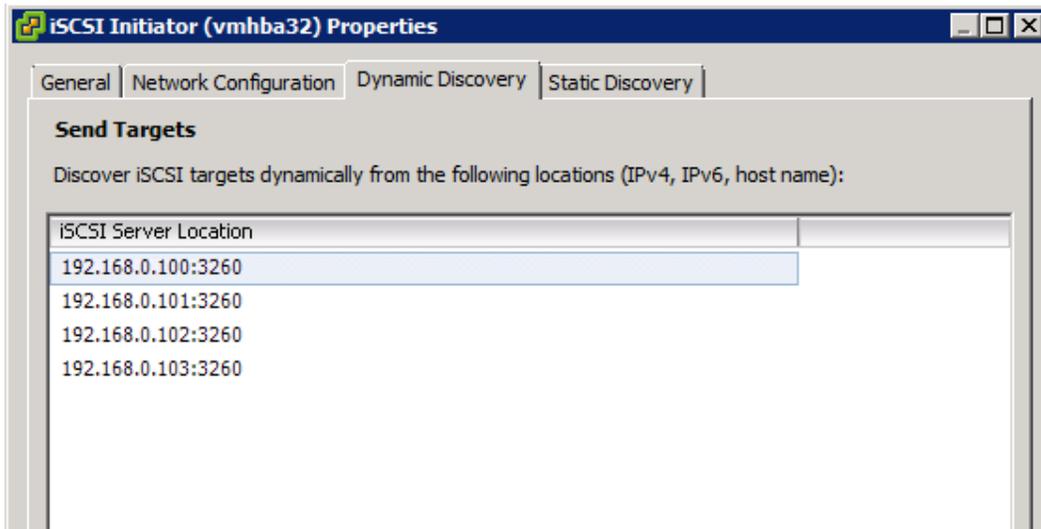


Figure 19: Using the host iSCSI Initiator to discover ioControl targets.

Each LUN was attached to VMWare hosts. We changed the default path selection from default fixed to round robin, enabling all four 10Gbps interfaces on the array to be utilized. Figure 20 shows where to manage the paths in host datastore properties with the VMKernel interface.

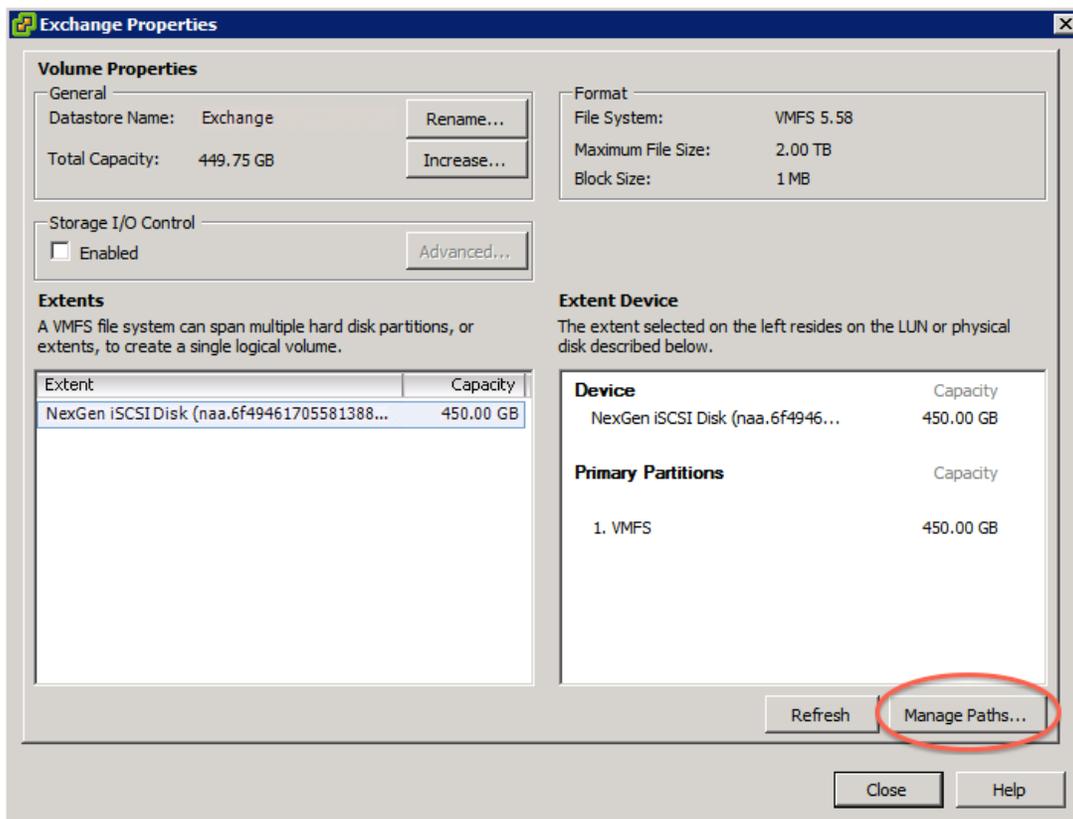


Figure 20: Managing paths in host datastore properties.

Figure 21 shows where to select the path setting.

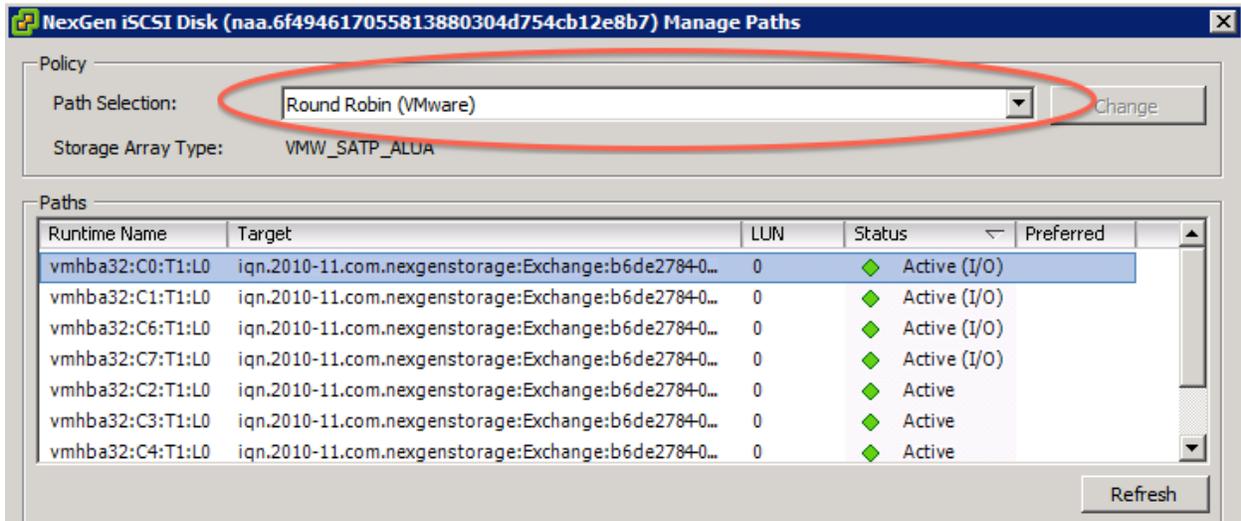


Figure 21: Selecting the multipath policy in host datastore properties.

Setting up Fusion ioTurbine

Fusion ioTurbine has two components: the ioSphere management appliance and a Fusion ioDrive2 mezzanine card. By using the PCI-e flash memory as a read cache, ioTurbine enabled our solution to handle large I/O demands without taxing the shared storage. The cache can be configured for the entire ESXi host or at the guest level. As Figure 22 shows, we enabled the cache at the guest level and instructed it to cache only the database volume (E:\) for each SQL server in our testing.

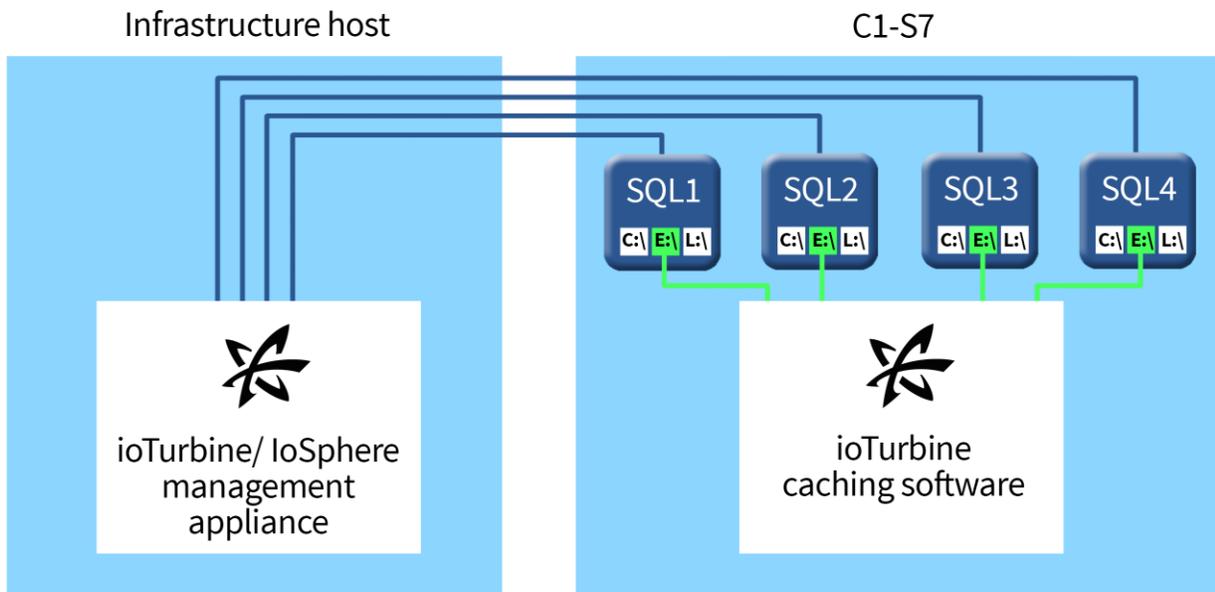


Figure 22: ioTurbine overview.

We installed the ioDrive2 mezzanine card into server c1-s7, which was the SQL and Exchange host. The ioTurbine management server is a virtual appliance that is deployed on an infrastructure ESXi host and managed by vCenter. The management server is deployed into the VMware environment as an OVA file. See the full report for detailed test infrastructure installation and deployment steps.⁴

Figure 23 shows how to register vCenter in ioSphere.

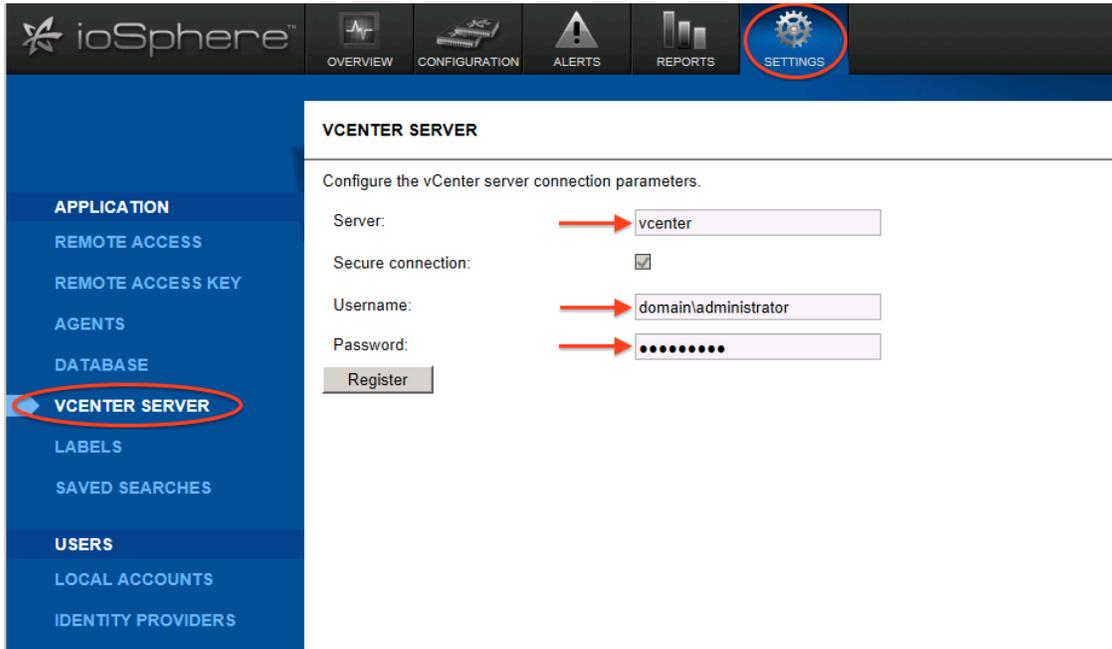


Figure 23: Registering vCenter in ioSphere.

Figure 24 shows where to find the Fusion ioSphere tab in vCenter.

⁴ www.principledtechnologies.com/Fusion-io/ioControl_reference_architecture_unabridged_0714_v3.pdf

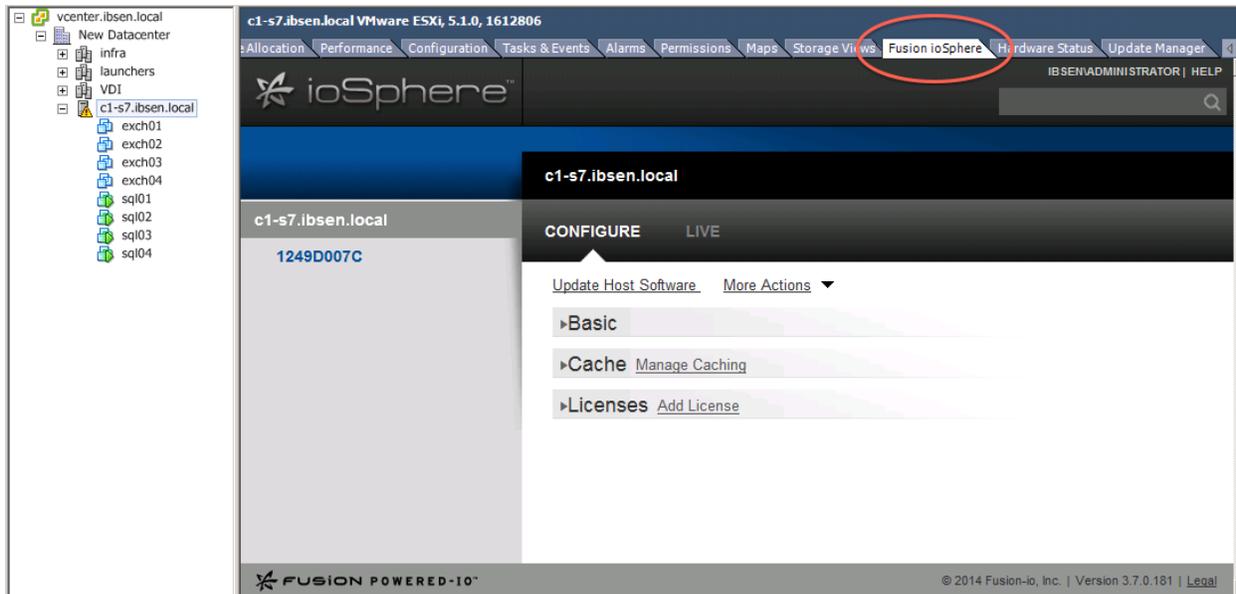


Figure 24: Locating the ioSphere tab for vCenter.

Figures 25 and 26 show how we enacted a low-level format with Fusion ioSphere.

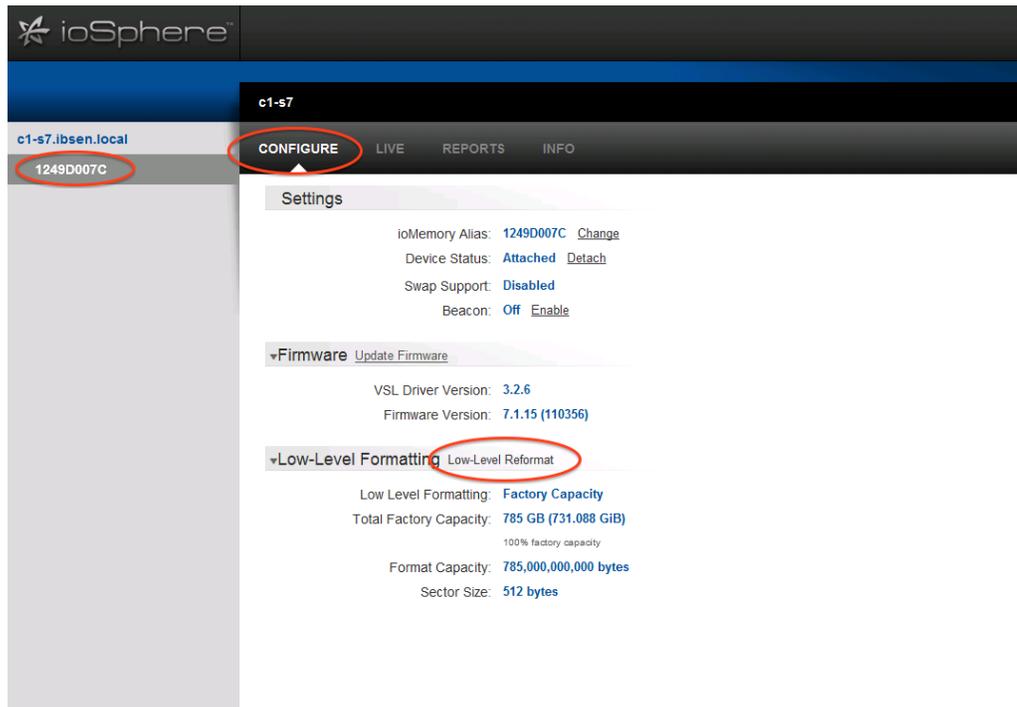


Figure 25: Selecting the low level reformat option in ioSphere.

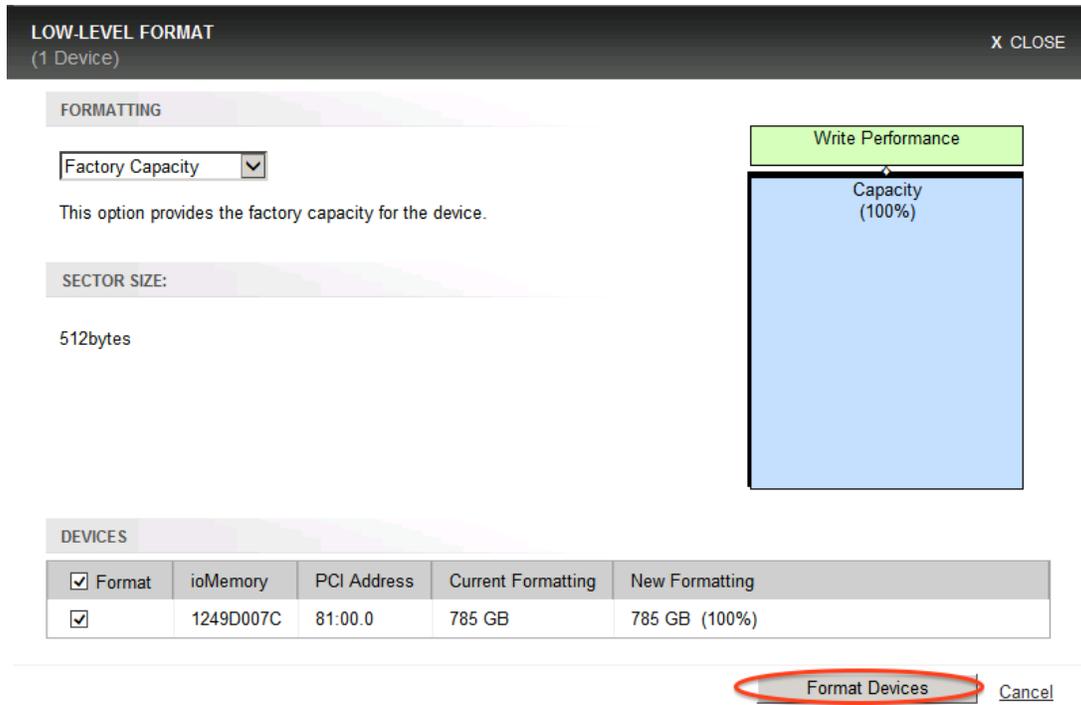


Figure 26: Performing a low-level format.

Figures 27 and 28 show how we enabled guest-based caching.

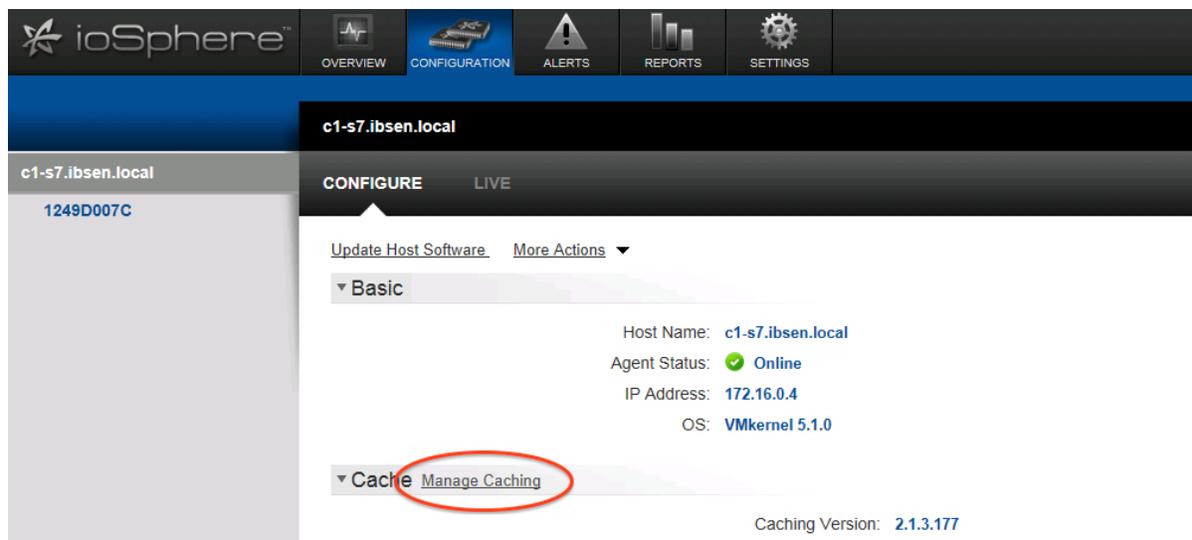


Figure 27: Locating the cache settings in vSphere.

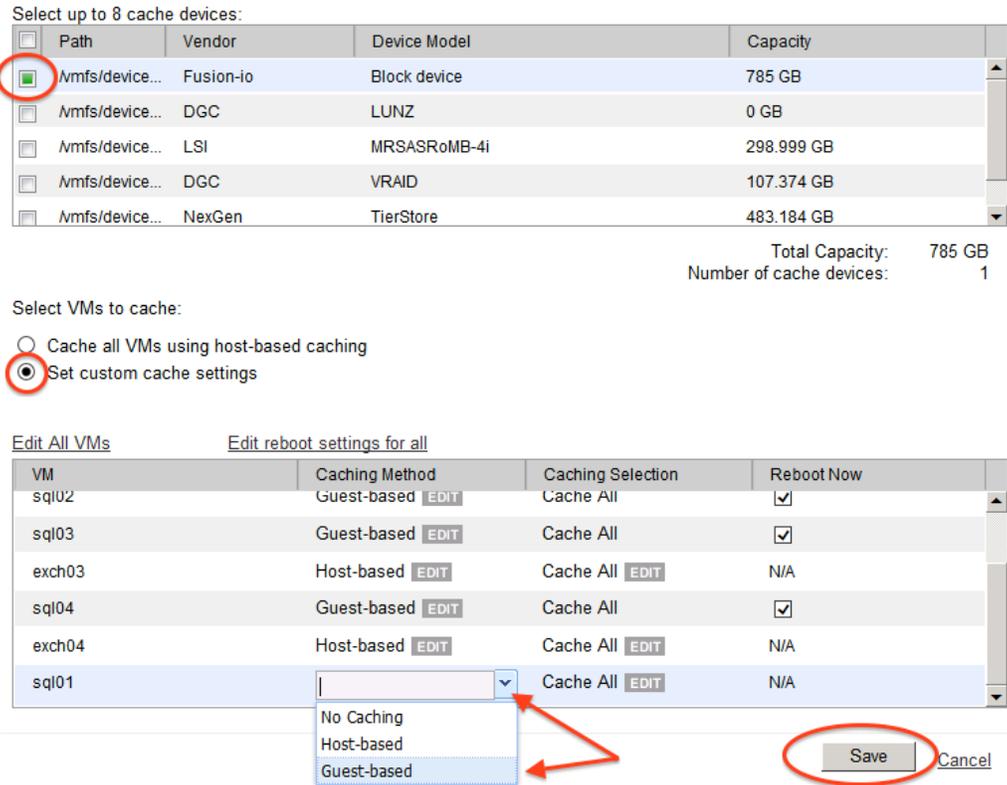


Figure 28: Choosing the correct options to enable guest-based caching.

Figure 29 shows where we entered guest credentials. This action is required to enable guest-based caching.

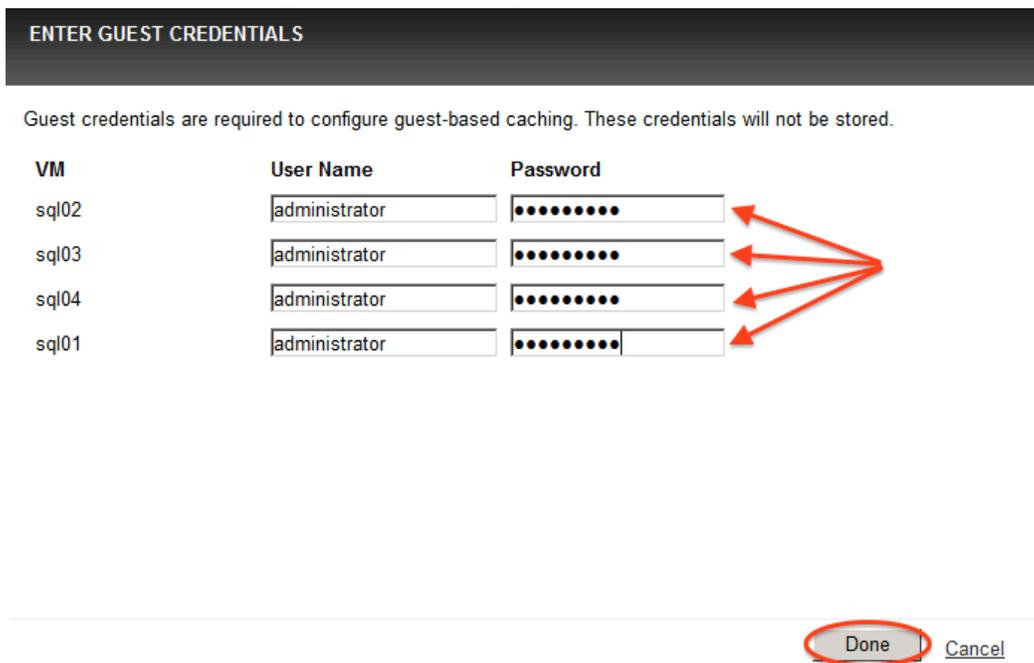


Figure 29: Entering guest credentials to enable guest-based caching in ioSphere.

Figure 30 shows where to change cache settings.

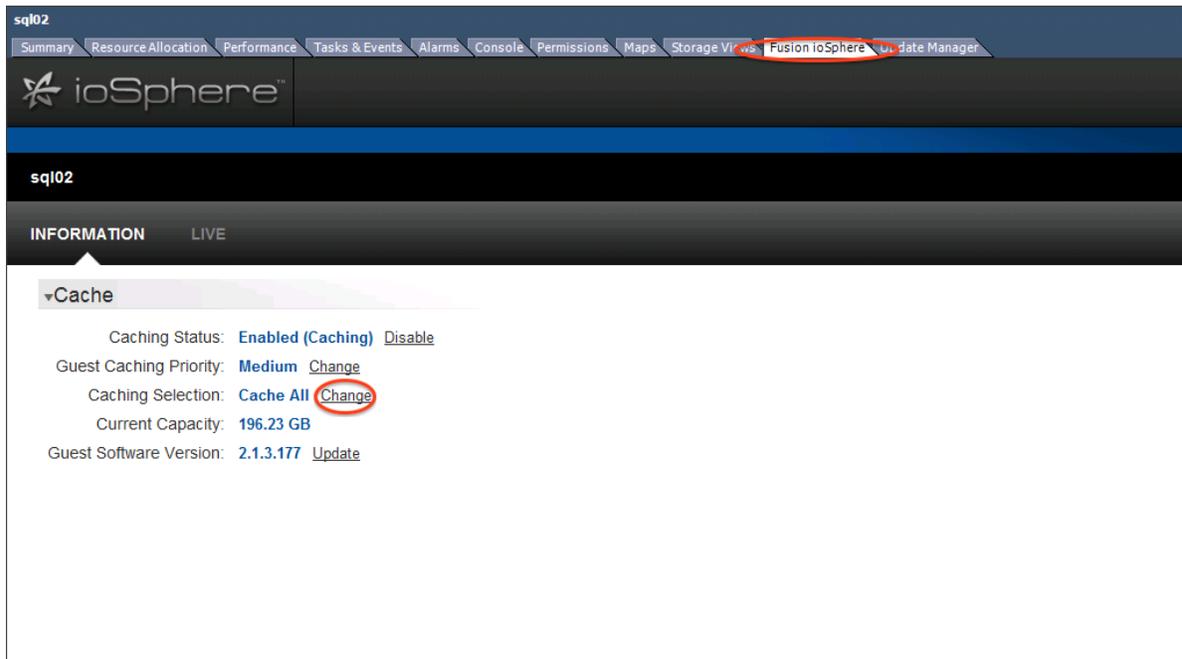


Figure 30: Editing cache information in ioSphere.

Figure 31 shows how we changed the cache volume.

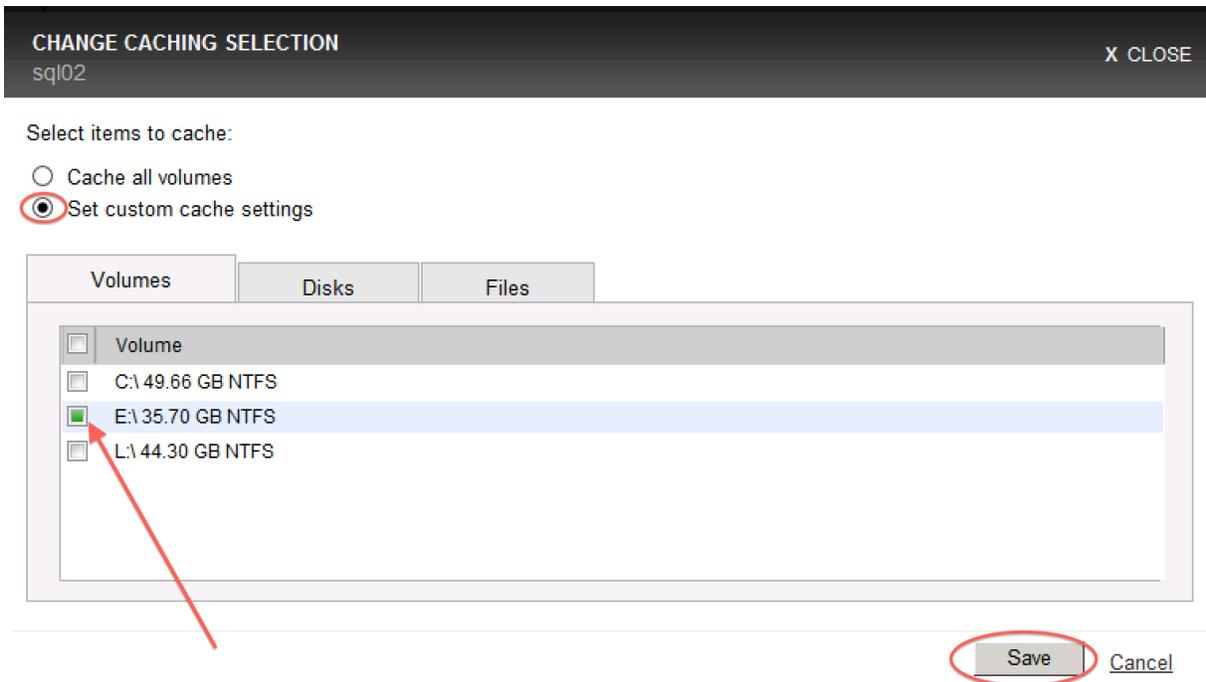


Figure 31: Changing the cache volume in ioSphere.

Virtual machine architecture

We installed all of our VMs on a flat network. We installed all of our infrastructure VMs on generic NAS and they were hosted on a separate infrastructure server. Figure 32 shows detailed information of the VMs.

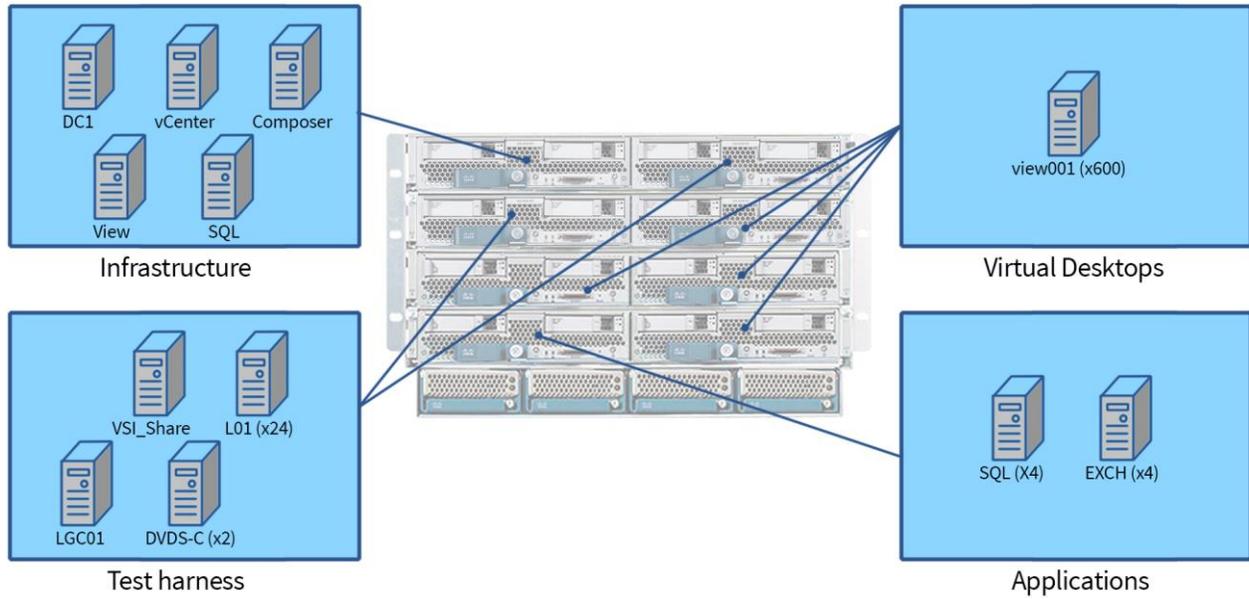


Figure 32: Detailed VM information and locations.

VALIDATION AND TEST RESULTS

Test scenario 1 – ioControl supporting all application workloads

- SQL/OLTP: Medium—about 6,500 OPM
- LoadGen/Exchange profile: Outlook_200
- VDI: 600 users, boot storm and Login VSI Medium workload (with flash), benchmark mode

Test scenario 1 showed that the four VDI servers and one ESXi server hosting OLTP and Exchange workloads pushed I/O requirements to burst over 20,000 IOPS and sustained a peak of over 6,000 IOPS. Storage response time remained low in spite of the demand placed on the array by the UCS Blade Intel E5-2690 processors. There was no I/O contention detected between the three workloads as VSI_{max} was not reached on the virtual desktop workload, RPC latency was low on the Exchange workload, and operations per minute (OPM) was consistent for our SQL workload. Figure 34 shows the system total for read/write IOPS.

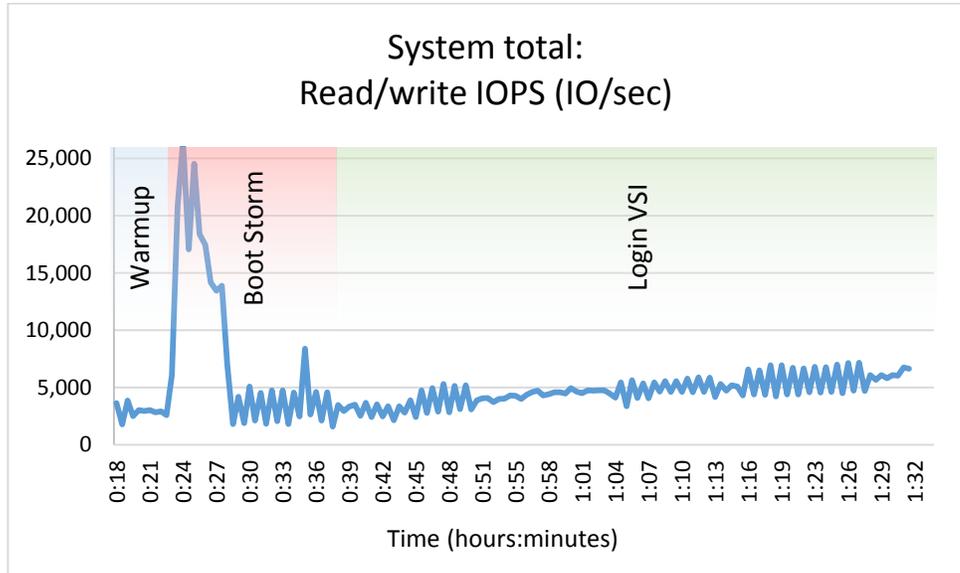


Figure 33: Test scenario 1: I/O footprint. The boot storm completed in about five minutes.

Figure 34 shows the system total read/write response time.

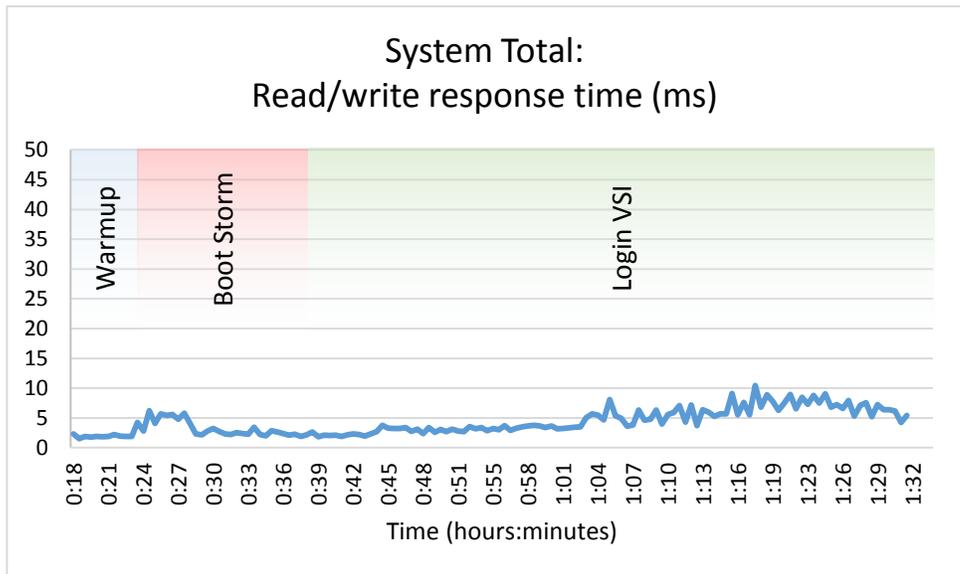


Figure 34: In test scenario 1, I/O response time remained low through all phases of testing.

Figure 35 shows the overall average OPM. The SQL server was not in contention for I/O and OPM remained flat throughout all testing.

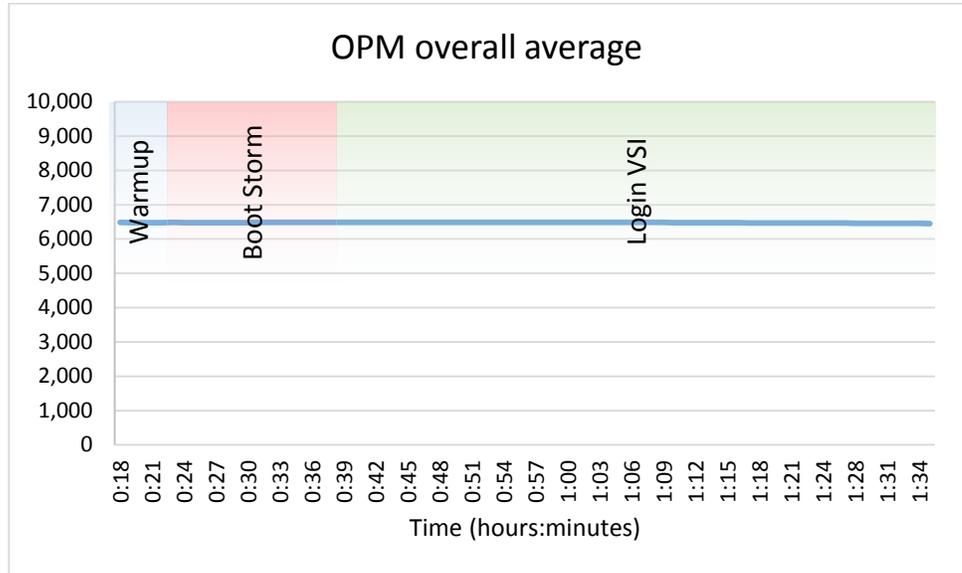


Figure 35: Test scenario 1: OLTP OPM.

Figure 36 shows the average Exchange RPC latency. At no point in testing did RPC latency reflect I/O contention. Note: Microsoft guidance states that anything less than 50ms is acceptable RPC latency.

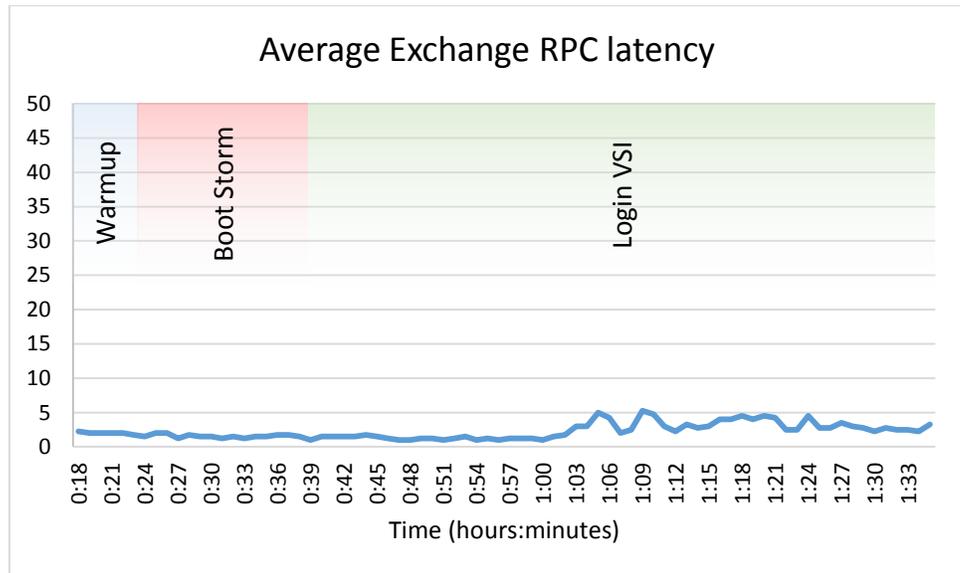


Figure 36: Test scenario 1: LoadGen RPC latency.

Looking at IOPS and response time from within the ioControl interface shows the additional IOPS being generated during VDI boot storm and the ability of the ioControl QoS engine to manage I/O prioritization for all three service levels. Figure 37 shows the IOPS and Figure 38 shows the associated response times.



Figure 37: ioControl interface showing IOPS generated during testing.



Figure 38: ioControl interface showing response times during testing.

Figure 39 shows Login VSI v4 output. At 600 users, VSI_{max} was not reached. The increase in response times was due to processor saturation. Figure 40 shows processor utilization on host c1-s4 throughout the testing.

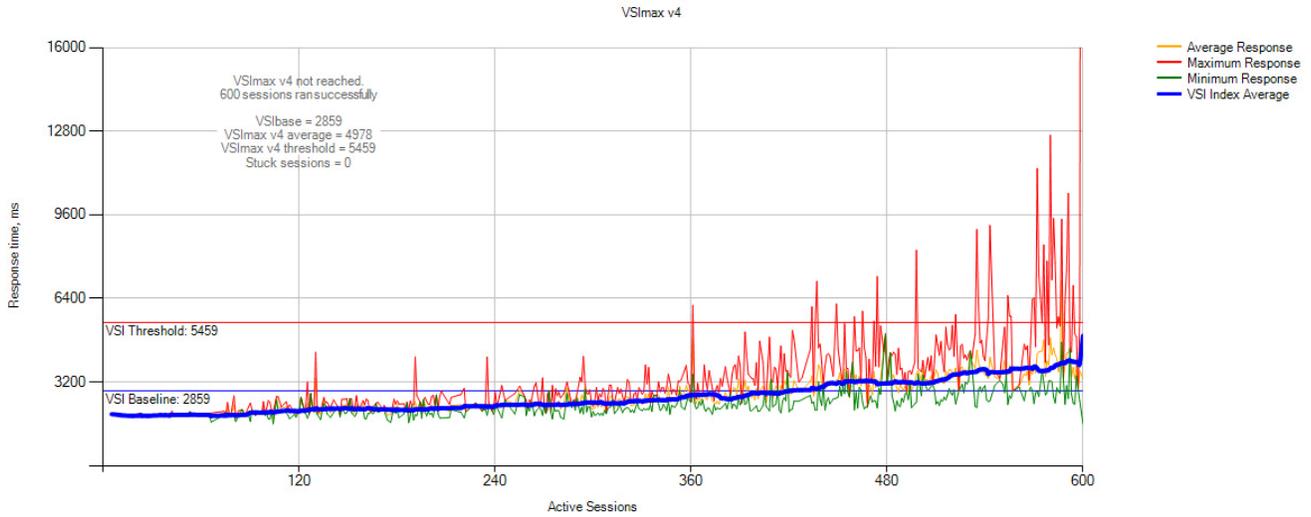


Figure 39: Test scenario 1: Login VSI v4 output.

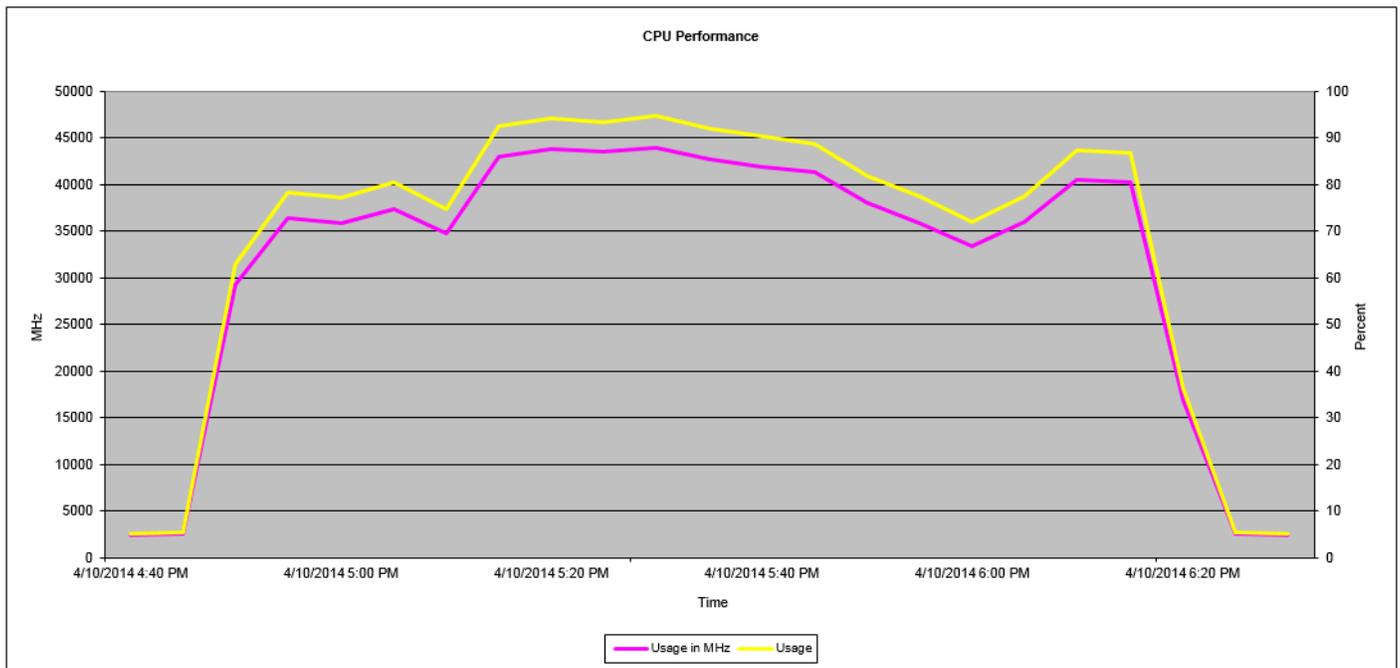


Figure 40: C1-S4 VDI processor utilization throughout the test.

Summary of test scenario 1

As evidenced by the test results outlined above, in spite of a heavy mixed workload that taxed our infrastructure, ioControl demonstrated its ability to support multiple high performance applications by utilizing real time data placement along with QoS policies to guarantee predictable performance levels.

Test scenario 2 – Scaling performance with zero footprint with Cisco UCS Fusion ioDrive 2 and Fusion ioTurbine for VMware

- SQL/OLTP: Heavy—35,000 to 40,000 OPM
- LoadGen/Exchange profile: Outlook_200
- VDI: 600 users, boot storm and login VSI medium workload (with flash), benchmark mode
- ioTurbine: Guest-based cache enabled

Test scenario 2 showed that, similar to test scenario 1, the four VDI servers and one ESXi server hosting OLTP and Exchange workloads pushed I/O requirements to burst over 20,000 IOPS and pushed a sustained peak of over 7,000 IOPS. The difference in this test came when ioTurbine was enabled and serviced about 10,000 IOPS on average, with bursts into the 15,000 IOPS range and allowing for an increase of 40,000 OPM. Storage response time again remained low. There was no I/O contention detected between the three workloads as VSI_{max} was not reached on the virtual desktop workload. RPC latency was low on the Exchange workload and OPM was consistent for our SQL workload. Figure 41 shows the system total read/write IOPS for test scenario 2.

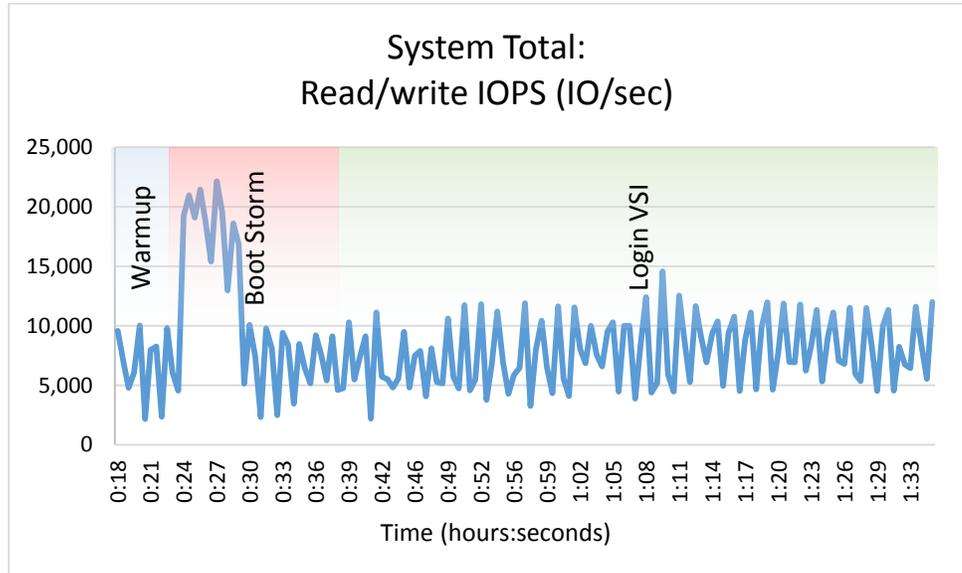


Figure 41: Test scenario 2: I/O footprint. The boot storm completed in about five minutes.

Figure 42 shows the read cache serviced as many as 17,000 read operations per second (OPS).



Figure 42: Test scenario 2: ioTurbine.

Figure 43 shows the system total read/write response time for test scenario 2.

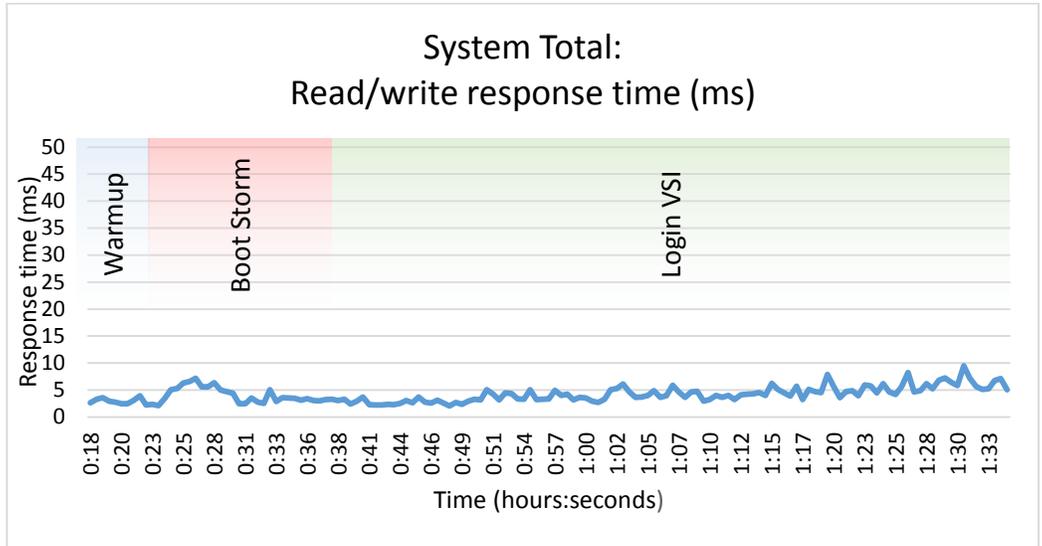


Figure 43: Test scenario 2: I/O footprint. The boot storm completed in about five minutes.

Figure 44 shows the OPM overall average in test scenario 2. The SQL workload was not in contention for I/O and OPM remained flat throughout all testing.

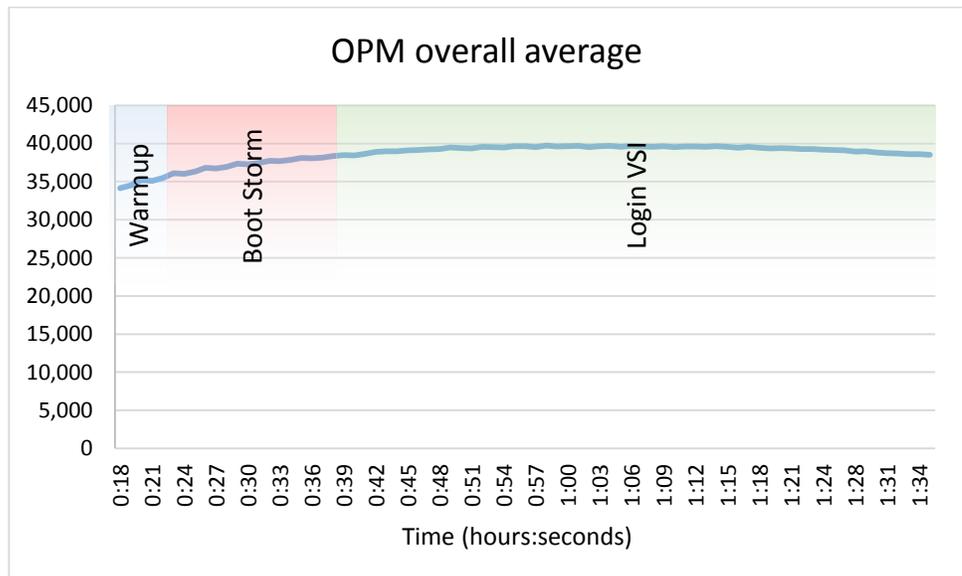


Figure 44: Test scenario 2: OLTP OPM.

Figure 45 shows the average Exchange RPC latency in test scenario 2. At no point in testing did RPC latency reflect I/O contention. Note: Microsoft guidance states that anything less than 50ms is acceptable RPC latency.

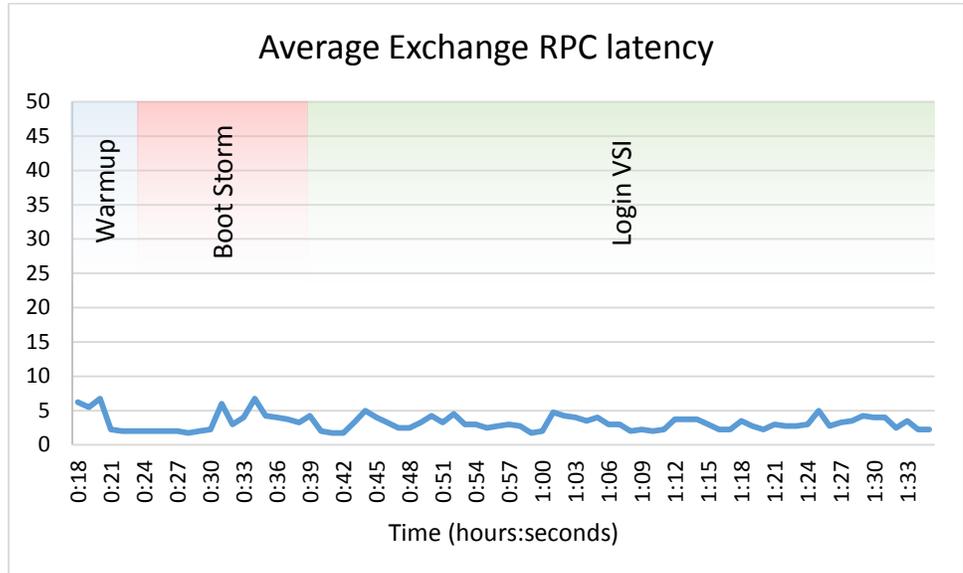


Figure 45: Test scenario 2: LoadGen RPC latency.

Looking at IOPS from within the ioControl interface shows only a minor increase in Mission Critical IOPS serviced by ioControl, in addition to again demonstrating the IO prioritization managed by the ioControl QoS engine (see Figure 46).

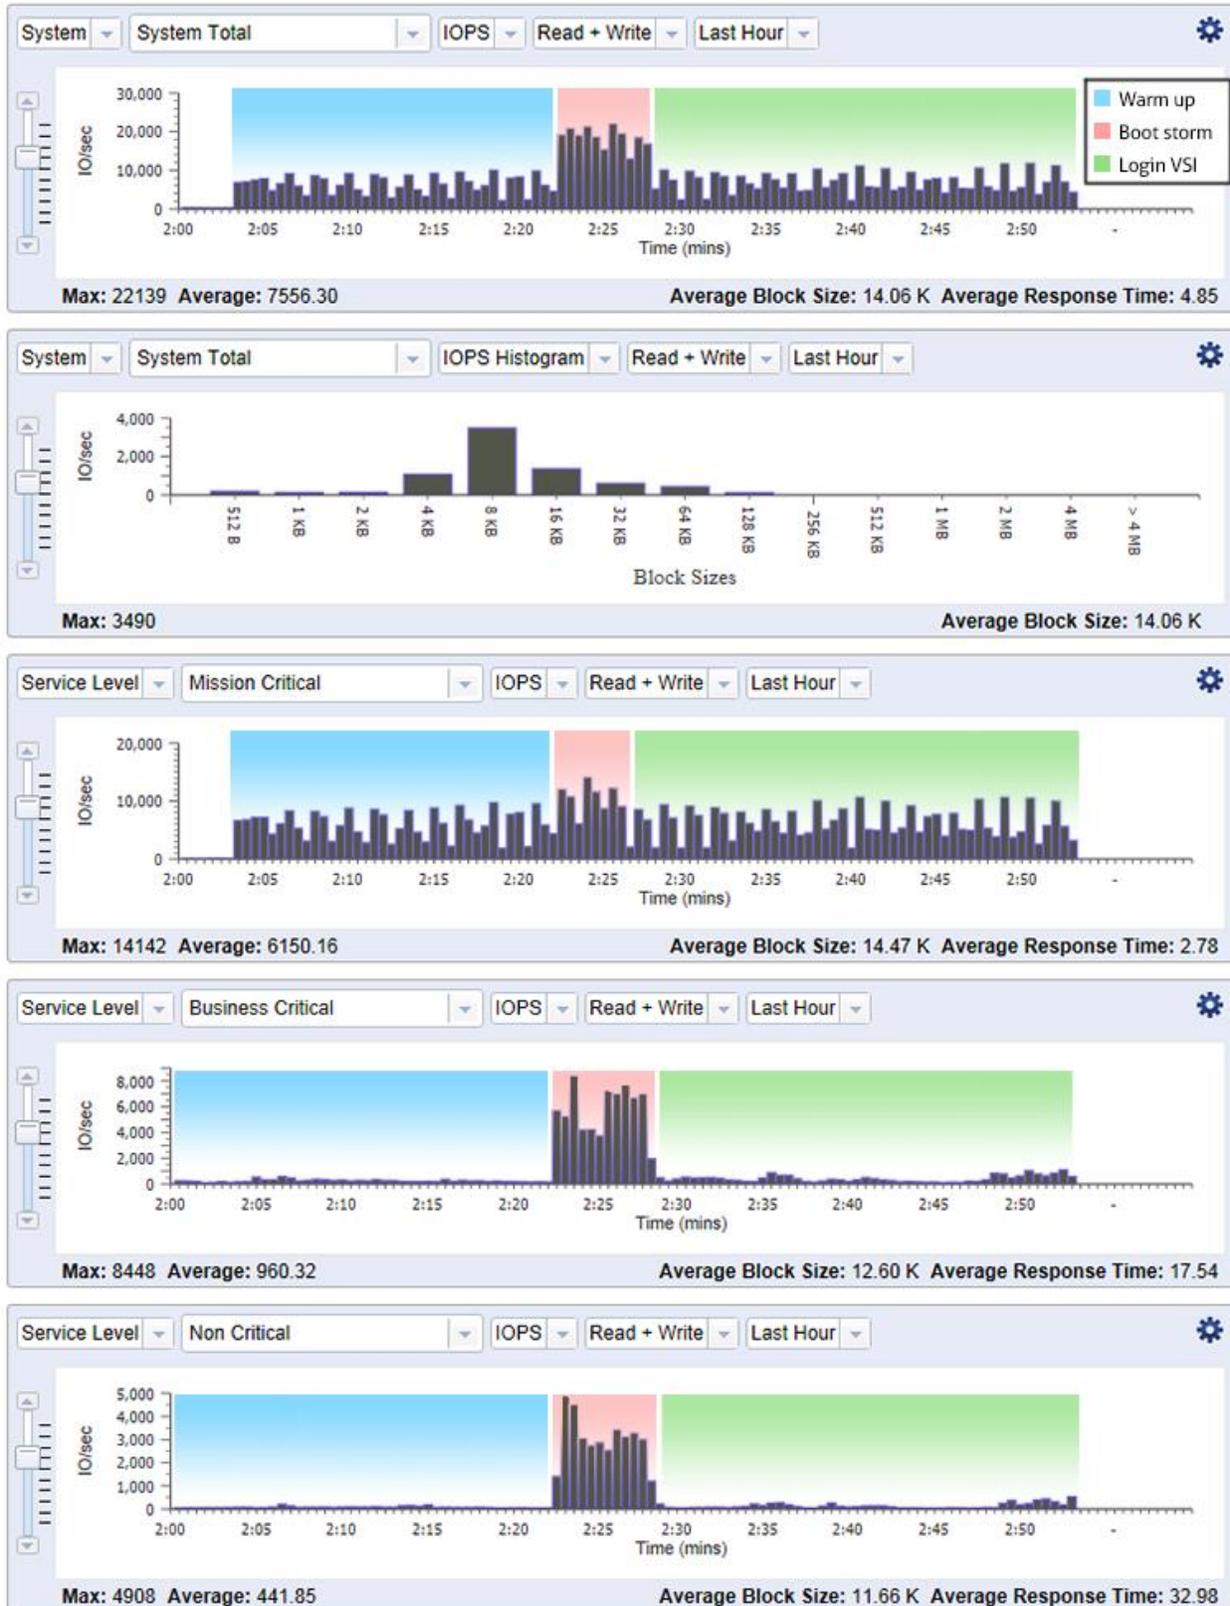


Figure 46: ioControl interface showing IOPS generated during testing.

Figure 47 shows the Login VSI v4 output. At 600 users, VSI_{max} was not reached. The increase in response times was due to processor saturation.

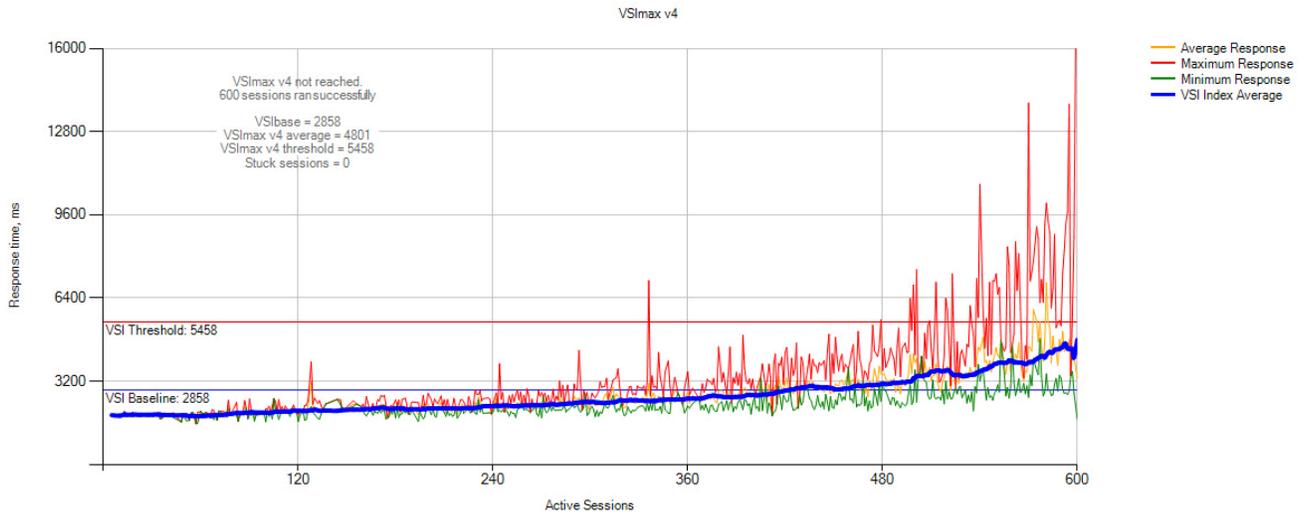


Figure 47: Test scenario 2: Login VSI v4 output.

Summary of test scenario 2

As can be seen from the test results, utilizing the ioDrive2 and ioTurbine software along with Fusion ioControl increased application performance while adding no additional footprint to the reference architecture. This resulted in a SQL OPM increase over six times the OPM from test scenario 1 with no performance degradation, while at the same time resulting in only a small increase in ioControl Mission-critical array IOPS.

Using the ioControl user interface

The ioControl UI let us view and record both real time and historical statistics on the ioControl appliance throughout all testing. We could view statistics as they pertain to QoS policy, storage processor, or the latency on SAS disk drives versus flash memory. Figure 48 shows the main dashboard, which provides an overview of the state of the appliance, and Figure 49 provides a view into the monitoring interface.

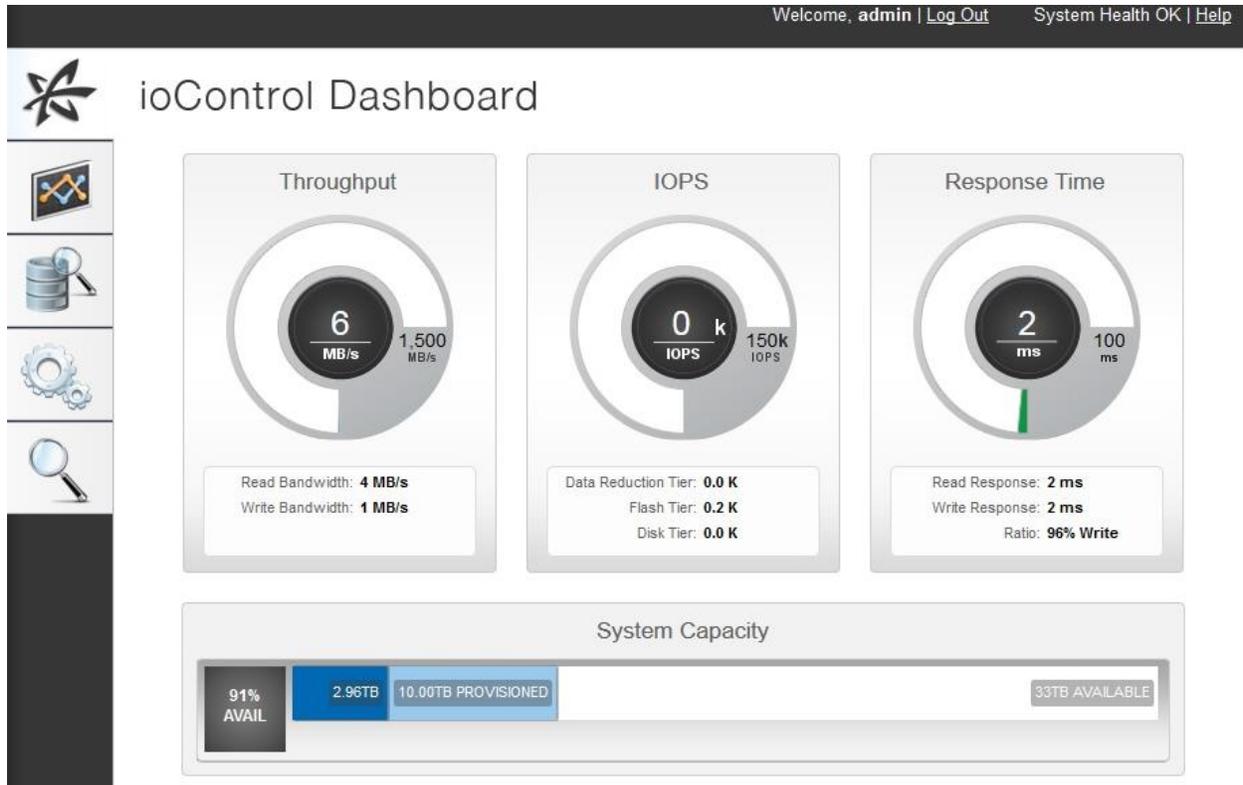


Figure 48: The ioControl Dashboard.

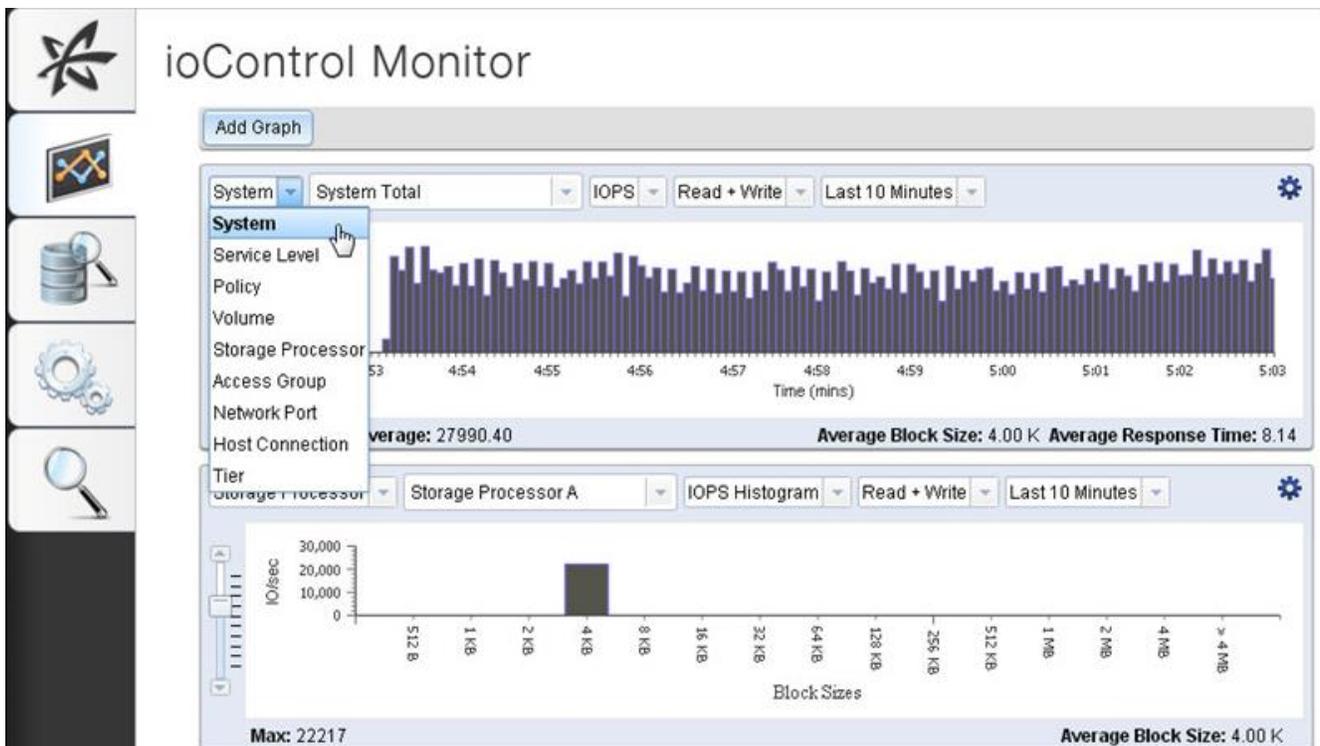


Figure 49: ioControl Monitoring.

SUMMING IT ALL UP

In our testing, Fusion ioControl, Cisco UCS, and VMware technologies created a high-performance, easy-to-configure and manage solution that could handle a demanding mixed application environment consisting of common virtualized business applications and VDI. By prioritizing performance resources with ioControl QoS service levels and policies, we demonstrated that mission-critical and business-critical applications will run at optimum performance levels in a 600 virtual desktop boot storm, or an OLTP heavy-workload use case. By adding Fusion ioTurbine software and Cisco Fusion ioDrive2 into the UCS Blade Server, we further extended OLTP performance with no additional solution footprint. With the Fusion ioControl UI and ioSphere UI integration with vCenter Server, you can manage all aspects of storage, both ioControl and UCS-side Fusion ioDrive flash, from within vCenter console. Fusion ioControl Hybrid Storage and Cisco UCS is a strong solution to meet today's demanding mixed workloads in your VMware virtualized data center.

APPENDIX A – DETAILED SYSTEM CONFIGURATION

Figure 50 shows the detailed information of the servers we used in testing.

System	Cisco UCS B200 M3 server	Cisco UCS B200 M2 server
General		
Number of processor packages	2	2
Number of cores per processor	8	6
Number of hardware threads per core	2	2
System power management policy	OS Control	OS Control
CPUs		
Vendor	Intel	Intel
Name	Xeon	Xeon
Model number	E5-2690	X5670
Stepping	6	CO
Socket type	LGA2011	LGA 1366
Core frequency (GHz)	2.90	2.93
Bus frequency	8.0 GT/s	6.4
L1 cache	32 KB + 32 KB	32 KB+ 32 KB (per core)
L2 cache	256 KB (per core)	256 KB (per core)
L3 cache	20 MB	12 MB
Platform		
Vendor and model number	Cisco UCS B200 M3	Cisco UCS B200 M2 Blade Server
Motherboard model number	Cisco FCH153271DA	N20-B6625-1
BIOS name and version	Cisco B200M3.2.1.1a.0.121720121447	Cisco S5500.2.1.1.0.1005201213338
BIOS settings	Default	Default
Memory module(s)		
Total RAM in system (GB)	256	96
Vendor and model number	Samsung M393B2G70BH0-YK0	Samsung M393B5170FH0-YH9
Type	PC3L-12800R	DDR3 PC3-10600
Speed (MHz)	1,600	1,333
Speed running in the system (MHz)	1,333	1,333
Size (GB)	16	8
Number of RAM module(s)	16	12
Chip organization	Double-sided	Double-sided
Rank	Dual	Dual
Hard disk		
Vendor and model number	Seagate ST9146803SS	Seagate ST9146803SS
Number of disks in system	2	2
Size (GB)	146	146
RPM	15,000	10,000
Type	SAS	SAS

System	Cisco UCS B200 M3 server	Cisco UCS B200 M2 server
RAID controller		
Vendor and model	LSI MegaRAID SAS 2004	LSI Logic® SAS 1064E
Controller firmware	20.10.1-0100	01.32.04.00
Operating system		
Name	VMware vSphere 5 VMware ESXi 5.1	VMware vSphere 5.1
Build number	1612806	1612806
Language	English	English
Operating system power profile	Maximum Performance	Maximum Performance
Adaptor(s)		
Vendor and model number	Cisco UCSB-MLOM-40G-01	UCS M71KR-Q QLogic® Converged Network Adapter
Type	mLOM	mLOM
Vendor and model number	Fusion IO iodrive2, 768 GB	N/A
Type	Mezzanine	N/A

Figure 50: Detailed configuration information for the servers.

Figure 51 shows detailed information for the storage we used in testing.

Storage array	Fusion ioControl n5-150 array
Firmware revision	
Model	Fusion-io n5-150
Part Number	XN50-2448-11E
ioMemory Capacity	2400 GB
Storage Processors	Dual Active-Active
RAM	48 GB
Network Interfaces	Data (4) 10GbE + (8) 1 GbE, Management : (4) 1GbE
Storage Protocol	iSCSI
Management	HTTP
Power	11.5 AMP
Dimensions	5.2" × 17.2" × 25.5"
Weight	75 lbs
Tray 1	
Number of disks	16
Disk vendor and model number	Toshiba MG03SCA300
Disk size (GB)	3 TB
Disk buffer size (MB)	64
Disk RPM	7200
Disk type	SAS-2, 6 Gbps
Firmware revision	0108 Rev a1

Storage array	Fusion ioControl n5-150 array
PCIe card	ioDrive II
Model	Fusion ioDrive II
Interface	PCIe

Figure 51: Detailed configuration information for the storage array.

APPENDIX B – DETAILED VIRTUAL MACHINE CONFIGURATION

Figure 53 shows the detailed information of the virtual machine we used in testing.

VM name	Qty	OS	Role	Host(s)	LUN	vCPUs	Mem (GB)	vDisk (GB)
DC1	1	Windows 2008 R2	Active Directory, DNS, DHCP	Infra1-3	NAS01	2	4	100
VSI_Share	1	Windows 2008 R2	Login VSI share	Infra1-3	NAS01	2	2	100
ioSphere VM	1	Virtual appliance	Fusion-io IO mgmt server	Infra1-3	NAS01	2	4	60
vCenter	1	Windows 2008 R2	VMware vCenter	Infra1-3	NAS01	2	8	100
Composer	1	Windows 2008 R2	VMware Composer	Infra1-3	NAS01	2	2	100
View	1	Windows 2008 R2	VMware View server	Infra1-3	NAS01	2	4	100
SQL	1	Windows 2008 R2	vCenter, and Composer DB server	Infra1-3	NAS01	2	4	100
L01-I24	24	Windows 7 (x64)	Login VSI Launcher	Infra1-3	NAS01	2	8	40
LGC01	1	Windows Server 2012	Exchange client	Infra1-3	NAS01	1	4	40
DVDS-C1,2	2	Windows Server 2012	OLTP clients	Infra1-3	NAS01	2	4	40
View001-600	600	Windows 7 (x86)	VMware Virtual Desktop	c1-s5-8	VDT A-D	1	1	24
SQL01-04	4	Windows Server 2012	OLTP, SQL server	c1-s4	DVD-SQL	4	12	130
EXC01-04	4	Windows Server 2012	LoadGen Exchange server	c1-s4	Exchange	4	4	90

Figure 53: Detailed information for our virtual machines.

ABOUT PRINCIPLED TECHNOLOGIES



Principled Technologies, Inc.
1007 Slater Road, Suite 300
Durham, NC, 27703
www.principledtechnologies.com

We provide industry-leading technology assessment and fact-based marketing services. We bring to every assignment extensive experience with and expertise in all aspects of technology testing and analysis, from researching new technologies, to developing new methodologies, to testing with existing and new tools.

When the assessment is complete, we know how to present the results to a broad range of target audiences. We provide our clients with the materials they need, from market-focused data to use in their own collateral to custom sales aids, such as test reports, performance assessments, and white papers. Every document reflects the results of our trusted independent analysis.

We provide customized services that focus on our clients' individual requirements. Whether the technology involves hardware, software, Web sites, or services, we offer the experience, expertise, and tools to help our clients assess how it will fare against its competition, its performance, its market readiness, and its quality and reliability.

Our founders, Mark L. Van Name and Bill Catchings, have worked together in technology assessment for over 20 years. As journalists, they published over a thousand articles on a wide array of technology subjects. They created and led the Ziff-Davis Benchmark Operation, which developed such industry-standard benchmarks as Ziff Davis Media's Winstone and WebBench. They founded and led eTesting Labs, and after the acquisition of that company by Lionbridge Technologies were the head and CTO of VeriTest.

Principled Technologies is a registered trademark of Principled Technologies, Inc.
All other product names are the trademarks of their respective owners.

Disclaimer of Warranties; Limitation of Liability:

PRINCIPLED TECHNOLOGIES, INC. HAS MADE REASONABLE EFFORTS TO ENSURE THE ACCURACY AND VALIDITY OF ITS TESTING, HOWEVER, PRINCIPLED TECHNOLOGIES, INC. SPECIFICALLY DISCLAIMS ANY WARRANTY, EXPRESSED OR IMPLIED, RELATING TO THE TEST RESULTS AND ANALYSIS, THEIR ACCURACY, COMPLETENESS OR QUALITY, INCLUDING ANY IMPLIED WARRANTY OF FITNESS FOR ANY PARTICULAR PURPOSE. ALL PERSONS OR ENTITIES RELYING ON THE RESULTS OF ANY TESTING DO SO AT THEIR OWN RISK, AND AGREE THAT PRINCIPLED TECHNOLOGIES, INC., ITS EMPLOYEES AND ITS SUBCONTRACTORS SHALL HAVE NO LIABILITY WHATSOEVER FROM ANY CLAIM OF LOSS OR DAMAGE ON ACCOUNT OF ANY ALLEGED ERROR OR DEFECT IN ANY TESTING PROCEDURE OR RESULT.

IN NO EVENT SHALL PRINCIPLED TECHNOLOGIES, INC. BE LIABLE FOR INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH ITS TESTING, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. IN NO EVENT SHALL PRINCIPLED TECHNOLOGIES, INC.'S LIABILITY, INCLUDING FOR DIRECT DAMAGES, EXCEED THE AMOUNTS PAID IN CONNECTION WITH PRINCIPLED TECHNOLOGIES, INC.'S TESTING. CUSTOMER'S SOLE AND EXCLUSIVE REMEDIES ARE AS SET FORTH HEREIN.
