kaminario.

WHITE PAPER

Performance Testing of VMware View VDI on Kaminario K2 All-Flash Storage Array

August 2016

TABLE OF CONTENTS

Executive Summary

Virtual Desktop Infrastructure, or VDI, enables enterprise IT groups to easily provision, manage and secure the growing proliferation of desktop computing. In addition to being subject to frequent security risks and often requiring onsite helpdesk visits, physical desktop computers typically run at only a few percent utilization, but still need to be provisioned for peak demands. Virtualizing these desktops has become practical using mature virtualization suites like VMware's Horizon View along with recent dramatic improvements in server and CPU performance. VDI enables users to access their own Windows desktop from anywhere while providing numerous benefits for IT groups, including greatly simplified desktop computing management, monitoring and security.

However, consolidating hundreds or thousands of virtualized desktop computing sessions onto shared servers introduces new demands on the storage infrastructure supporting the virtual drives. These large populations of virtualized desktop operating systems and applications generate a random blend of I/O requests that can heavily stress both VDI servers and their storage infrastructure.

While most of user activity in virtual desktops is randomly distributed over time, there are periods during the day when large groups of users perform the same action. For example: login at the start of the workday, suspending when leaving for lunch, and suspending or shutting down at the end of the workday. These commonly timed activities create peak loads or storms of I/O requests on the underlying storage of the VDI system. In addition, scheduled, automated processes in the Windows desktop operating system such as virus scans or patching of Windows and application code can also induce storms of peak I/O demand.

The usual method of addressing these peaks of random I/O demand is to over-provision storage capacity and performance, greatly inflating both purchase and operating costs. However, even that approach may not be effective for large VDI deployments. A better alternative is to move to an All-Flash Array solution instead, ensuring that the selected array can scale up as well as out in order to handle growth and new demands from VDI users. With many flash array systems, though, this approach might succeed from a technical standpoint, but will be too costly to be practical with regard to large VDI deployments.

This technical report covers detailed performance and efficiency tests for a Kaminario K2 array supporting VMware Horizon View virtual desktops that quantify its ability to support a range of critical VDI workloads:

- **Provision 1,000 Linked Clone Virtual Desktops** – For the duration of the 70 minute test, throughput peaked above 1,000 MB/sec while IOPS averaged around 30,000, with frequent peaks as high as 60,000 – 70,000 IOPS and latency remaining below 1.0 ms.
- **Provision 1,000 Full Clone Virtual Desktops** – Average throughput was 2,600 MB/sec, while IOPS peaked at 20,000 and averaged around 10,000 IOPS.
- **Boot 1,500 Mixed Virtual Desktops** – During the eight minute test Throughput peaked at over 3,000 MB/s while IOPS peaked at over 170,000.

- **Power Off 1,500 Mixed Virtual Desktops** – During the three minute test Throughput peaked at over 800 MB/ sec with IOPS peaking at 75,000. Latency remained well below 0.8 ms.

- **Data Reduction for Virtual Desktops** – K2's native deduplication and compression reduced the storage required for 1,000 virtual desktops by over 98%, a 68:1 ratio. More complex View Planner virtual desktops containing numerous applications and test datasets were reduced by over 97%, or 36:1.

- **Desktop and Application Performance on 1,250 Virtual Desktops** – Using 1,250 virtual desktops created by VMware's View Planner showed throughput peaking at over 800 MB/sec and IOPS with peaks of 24,500 IOPS. Latency was around 0.5 ms throughout the test. Mean application response times were 0.38 seconds on average, well below View Planner's passing mark of 1.0 seconds.

Performance test results showed that Kaminario's fifth-generation K2 flash array can support even the worst-case workloads of provisioning and shutdown/boot storms, and can do so in large VDI deployments that would overwhelm other legacy or hybrid storage arrays. The K2 array was able to handle all of these strenuous tasks with no performance degradation while still being highly cost-effective, with capacity costing below $1 per effective GB after deduplication and compression.

Introduction to Kaminario K2

Kaminario is leading the revolution in enterprise flash storage by creating the industry's most scalable, intelligent and cost-effective all-flash storage solution on the market. Built from the ground up to take advantage of the most modern flash SSD capabilities, the fifth-generation K2 All-Flash Storage Array is the only product to feature a true scale-out and scale-up architecture that allows organizations to grow capacity and performance based on their needs. This architecture ensures both data availability and a consistent level of high throughput, IOPS and low latencies needed to support the demanding random I/O generated by business-critical systems like virtualized desktops.

K2 is the only All-Flash Array with the architecture that can scale up, scale out or both. K2 can scale performance or capacity without impacting latency, with zero hot spots and no system tuning required. Each K-Block includes two K-Nodes (storage controllers), up to two shelves of 24 SSDs (scale-up), and either Fibre Channel or iSCSI host connectivity. For increased performance, K-Blocks can be added (scale-out) as shown in Figure 1 on the next page.

Key attributes of Kaminario's new fifth-generation K2 All-Flash Storage Array are:

- **Low Deployment and Operating Costs** – K2's CAPEX and OPEX costs can be two-thirds of legacy or hybrid storage arrays. The K2 has the unique ability to start small and scale according to storage requirements. After savings from deduplication and compression, K2's effective costs can be below $1 per GB or under $20 per virtual desktop.

- **No Single Point of Failure** –A shared-nothing architecture and automated recovery ensure the array can keep running at peak performance even with concurrent hardware failures of up to two SSDs per shelf.

- **Consistent High Performance** – Intelligent hardware and software ensure the K2 array delivers low latency even under peak workloads.
- **Non-Disruptive Upgrades and Expansion** – Software and hardware upgrades and expansion of the array require no downtime and manual tuning.
- **Simple Deployment and Management** – No disk groups to plan, configure or manage; everything is managed via a simple browser-based GUI. VMware integration simplifies management of K2 arrays as part of a virtualized server infrastructure, with a vCenter plugin and support for the VMware VAAI hardware offload/acceleration API.
- **High-Efficiency RAID** – K2's dual-protected K-RAID™ is highly efficient with 87.5% of the physical capacity available for data storage and consistent high performance even during rebuilds.
- **Advanced Data Reduction** – Thin provisioning, native inline compression and global selective inline deduplication provide savings of up to 70-95% without impacting host VDI servers. K2's deduplication can reduce the capacity required for VMware Horizon View virtual desktops by 90% or more.
- **Advanced Features** – Consolidated writes, adaptive block sizes and writeable snapshots simplify management, optimize performance, and save capacity.

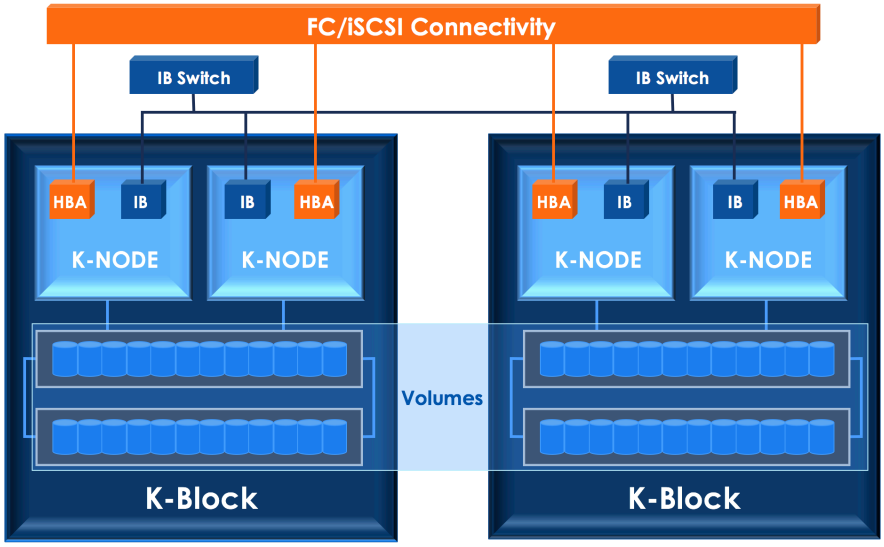

Figure 1: K2 Architecture Diagram

More information on the features and cost effectiveness of the fifth generation of Kaminario's K2 All-Flash Storage Array can be found on the **Kaminario website** or the **K2 Technical Architecture White Paper**.

Performance Tests for VDI

To quantify the ability of Kaminario's K2 array to support VMware Horizon View-based VDI servers, a set of performance tests were run on identical configurations of host server and storage hardware supporting large populations of desktop virtual machines (DVMs). Detailed hardware and software configurations used for testing are in the appendix of this document. In most tests, the K2 array was not the limiting factor in VDI performance, with throughput being limited by the server hardware as well as by the desktop operating system and virtualization software.

Provision 1,000 Linked Clone Virtual Desktops:

Challenge: The initial test was to provision 1,000 linked clones of DVMs. Linked clones are copies of DVMs that share virtual disks with a parent DVM in an ongoing manner. This conserves disk space and allows multiple DVMs to use the same software installation.

During the test, the K2 array's throughput, IOPS and latency were measured using the K2 GUI performance analysis tab for the duration of the main portion of the test period.

Results: Provisioning of 1,000 linked clone desktop virtual machines ran for a total of 70 minutes. During the test, throughput of the K2 array peaked above 1,000 MB/sec. IOPS averaged around 30,000 with frequent peaks as high as 60,000 – 70,000 IOPS. Latency remained below 1.0 ms during the entire test, even during IOPS peaks. Details of these results can be seen in Figure 2.

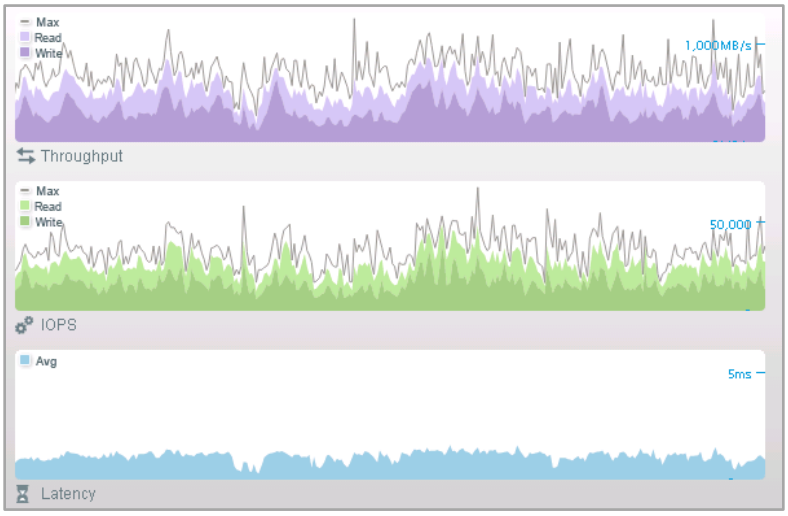

Figure 2: K2 GUI during 1,000 Linked Clone Provisioning Test

Provision 1,000 Full Clone Virtual Desktops

Challenge: The next test provisioned 1,000 full clones of DVMs. Full clones are copied as independent virtual machines with no need to access the parent DVM from which they were cloned. The cloned virtual disk is a complete copy of the parent's virtual disk at the time of cloning, but any subsequent updates to either's drives are independent. This makes them more flexible than linked clones but requires much more I/O for cloning and raises storage space requirements significantly before K2's deduplication is applied.

During the test, the K2 array's throughput, IOPS and latency were measured using the K2 GUI performance analysis tab for the duration of the test period. In addition, the ability of K2 deduplication to greatly reduce the storage capacity needed to store the DVMs was measured with the provisioned capacity compared to allocated and the actual physical capacity used on the array. K2 supports VMware's VAAI API allowing the desktop cloning operations to be offloaded to the storage array where they can be handled very efficiently and run extremely fast.

Results: The test's provisioning of 1,000 full clone virtual machines showed average throughput of 2,600 MB/sec. Measured IOPS peaked at 20,000 and averaged around 10,000.

In addition, the storage capacity needed for these 1,000 full clone DVMs was reduced from an allocated total of 14.1 TB to only 211 GB of physical storage on the array. K2's native inline deduplication and compression achieved an impressive reduction ratio of 68:1 between allocated and actual physical storage used, as shown in Figure 3 below. This highly efficient 98% data reduction by the K2 array makes it possible to get the flexibility of full clones with the same or better space savings of linked clones.

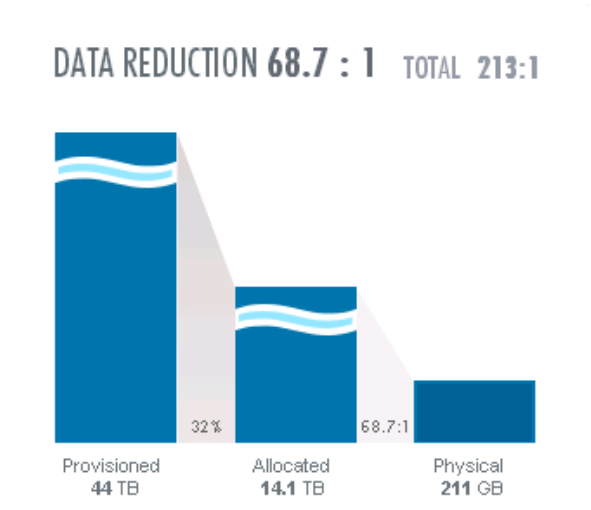

Figure 3: Data Reduction during 1,000 Full Clone Provisioning Test

Boot 1,500 Mixed Virtual Desktops

Challenge: The next test measured the load on the array when 1,500 virtual desktops with a 50/50 mix of linked clones and full clones were booted simultaneously. This test simulated the boot storms that can be seen in real-world VDI deployments when large numbers of users start up the virtual desktops at the beginning of a business day.

Results: Booting 1,500 mixed DVMs caused a series of I/O peaks during the eight minute test. Throughput varied with peaks of over 3,000 MB/s while IOPS peaked at over 170,000. Details of these results over the entire 10 minute test period, including the steady-state reading prior to the start of the test, can be seen in Figure 4.

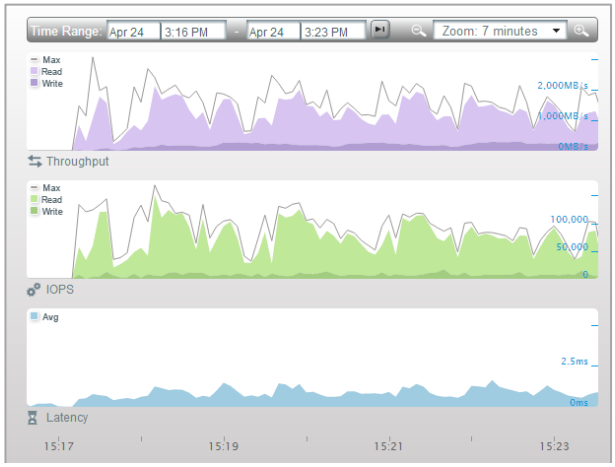

Figure 4: K2 GUI during 1,500 Virtual Desktops Boot Test

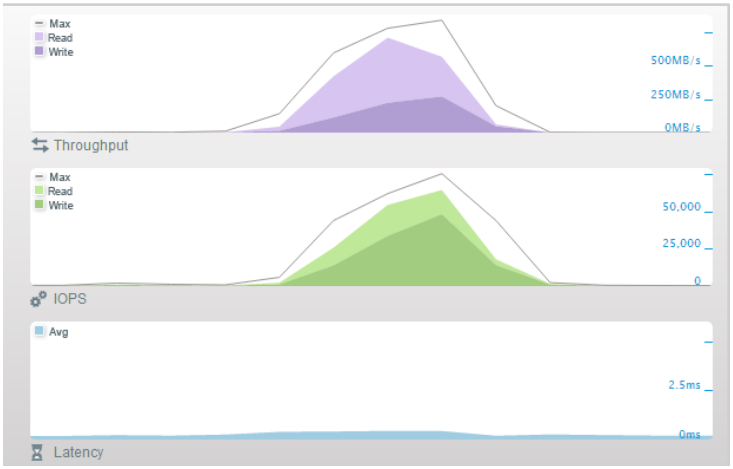

Figure 5: K2 GUI during 1,500 Virtual Desktops Power-Off Test

Power Off 1,500 Mixed Virtual Desktops

Challenge: The next test measured the load on the array when 1,500 virtual desktops with a 50/50 mix of linked clones and full clones were shut down using a Power Off command. Shutting down a DVM causes a peak of I/O requests as disk buffers are flushed to the virtual drive and applications are closed down. This test measured the blend of random I/O seen in real-world VDI deployments when virtual desktops are rebooted after a critical patch is applied or at the end of a business day when workers simultaneously shut down their virtual desktops.

Results: Powering down the DVMs generated a high I/O load during the three minute shutdown. Throughput peaked at over 800 MB/ sec while IOPS reached a peak of roughly 75,000. Latencies remained well below 0.8 ms during the entire test. Details of these results can be seen in Figure 5.

View Planner On 1,250 Mixed Virtual Desktops

Challenge: This test used the VMware View Planner 3.0.1 software to simulate application activity within a set of 1,250 View Planner virtual desktops, with a 50/50 mix of 625 linked clones and 625 full clones. During the test, throughput, IOPS and latency were measured using the K2 GUI.

Results: Results for the View Planner test showed several peaks of I/O demands on the K2 array as applications were loaded and simulated workloads with roughly 46% read and 54% write activity were run. Throughput varied with peaks of over 800 MB/sec. IOPS varied as well with peaks of 24,500 IOPS. Latency remained around 0.5 ms throughout the test period, as shown in Figure 6.

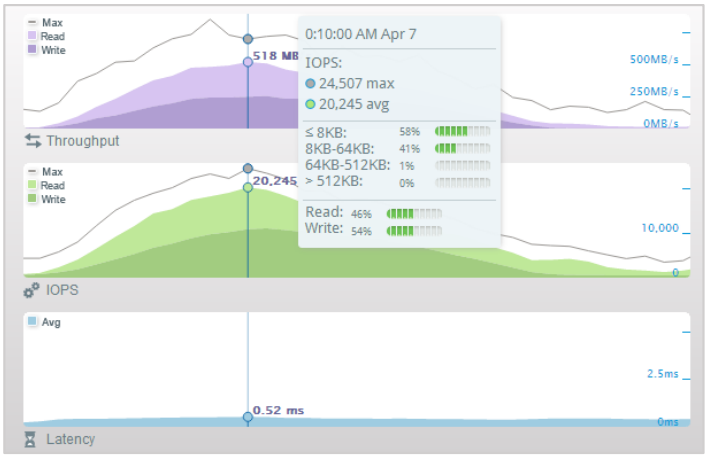

Figure 6: K2 GUI during View Planner Test with 1,250 Virtual Desktops

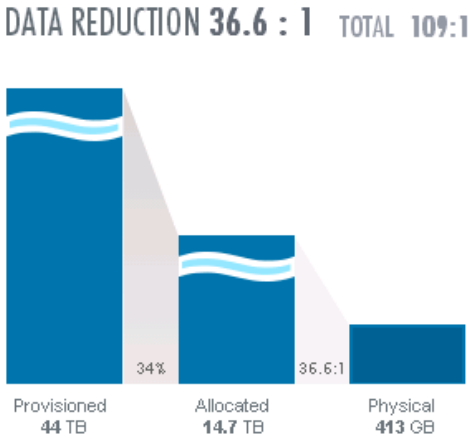

Figure 7: Data Reduction during View Planner 1,250 Virtual Desktops Test

The test DVMs created by View Planner also showed the benefits of data reduction by the K2 array's native compression and deduplication. The thinly provisioned size of 44 TB dropped to an allocated size of just 14.7 TB. More importantly, the actual physical capacity used for the 1,250 mixed virtual desktops was only 413 GB, a savings of over 97%, or over 36:1, as shown in Figure 7.

As the DVMs created by View Planner are designed to simulate normal activity, the data reduction from the K2 array's native compression and deduplication will reach a steady state as the virtual desktops are used. Over time, the steady-state data reduction ratio is expected to be in the range of 10-20:1.

View Planner Mean Application Response Times

Challenge: The View Planner benchmark tool supplied a template DVM preinstalled with a range of desktop applications. These desktops provide a simulated workload reflecting typical enterprise desktop computing environments. Applications included in the View Planner desktops were Adobe Reader, Microsoft Office (Excel, Outlook, PowerPoint, and Word), the Internet Explorer web browser and Mozilla's Firefox web browser. View Planner recorded response times for typical tasks in each application. More information on VMware View Planner can be found at: **www.vmware.com/products/view-planner.**

Results: Reported application mean response times for View Planner's Group A activity subset had an average of 0.38 seconds, which were well below View Planner's passing mark of 1.0 seconds, providing a crisp response for virtual desktop users. Detailed results for the measured application activities are shown in Figure 8.

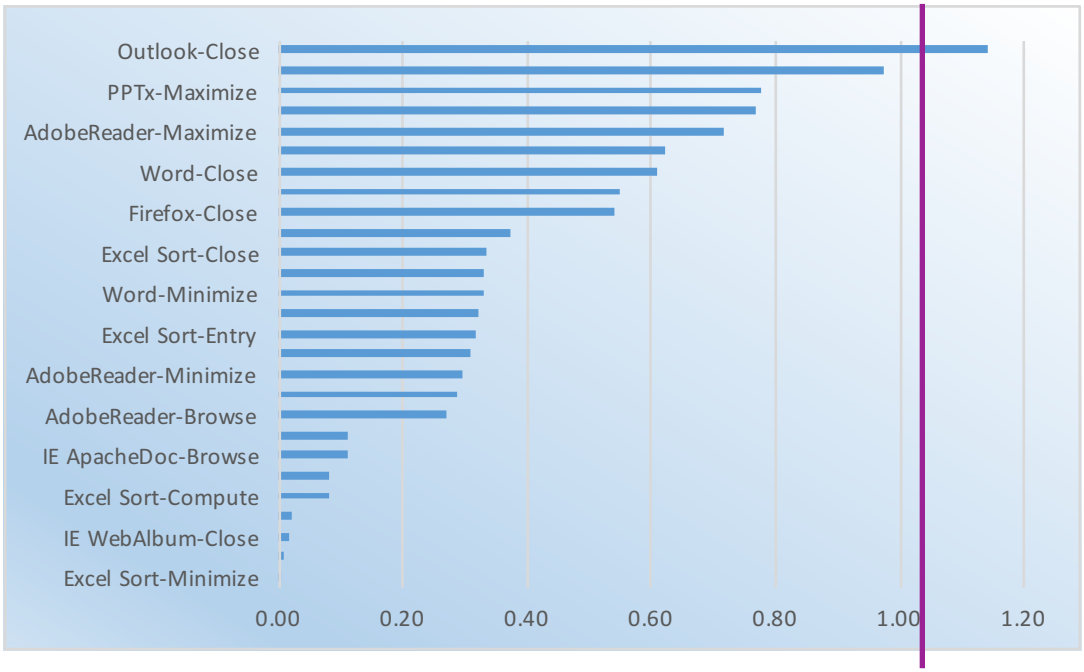

Figure 8: Detailed Mean Application Response Times with 1,250 Desktops

Summary

VDI provides IT groups with the ability to greatly simplify the deployment, management and support of desktop computing throughout the enterprise. However, by concentrating and blending the I/O activity of hundreds or thousands of formerly isolated PC hard drives, VDI places huge demands on the supporting storage systems. This concentrated demand is not steady, peaking into storms when groups of users perform the same activity such as logging in, leading to dissatisfaction and lower productivity. An all-flash storage architecture can fill VDI's need for performance, but only if is also cost-effective, robust and simple to operate.

The ability of Kaminario's new fifth-generation K2 All-Flash Storage Array to support large VDI deployments, even when faced with peak workloads from simultaneous provisioning, boot up and shutdown, was demonstrated in this series of performance tests. K2 is easily managed via a common GUI and provides robust availability with a shared-nothing architecture and highly efficient RAID implementation. K2 also uniquely meets requirements for cost-effectiveness, effectively costing below $1 per GB or $20 per virtual desktop after capacity requirements are reduced by its highly efficient inline deduplication and compression.

Appendix: Test Configuration Details

Hardware Configurations – Arrays, Servers, Connectivity

For testing, a K2 array with two K-Blocks was configured with different sets of VMware Horizon View desktop and management virtual disks. Each of the two K Blocks consisted of one SSD shelf holding 24 SSDs each.

An out-of-box configuration was used for the K2 array and no tuning was performed for any of the specific tests. The volumes were configured in a deduplicated volume-group and mapped to the ESX hosts.

The test servers included one ESX host server used to hold all the management virtual machines (vCenter, VMware Horizon View, VMware View Composer, Microsoft Active Directory), and 16 ESX host servers for the desktop virtual machines. All the hosts used in this test were running ESX version 5.5.0.

Details on the host servers were:

- For management virtual machines one Dell PowerEdge R810 with four 1.994 GHz Intel® Xeon® E7-4850 CPUs for a total of 80 hyperthreaded cores, 256 GB of memory, and two Dual Port 8 GB Qlogic FC HBA used to connect to the K2 array.
- For each of the sixteen M610 Dell Blades used for desktop virtual machines:
  - Two Intel Xeon E5649 @ 2.53 GHz CPUs – total of 24 cores per server
  - 96 GB of memory
  - Two Dual Port 8 GB Qlogic FC HBA to connect to the K2 array

Two 1 GbE networks were used to connect virtual machines during the testing while a separate Fibre Channel fabric was used to connect the K2 array to the host servers.

One network was used for external communication with the VDI host servers and the server hosting the management virtual machines, as shown in Figure 9 on the next page.

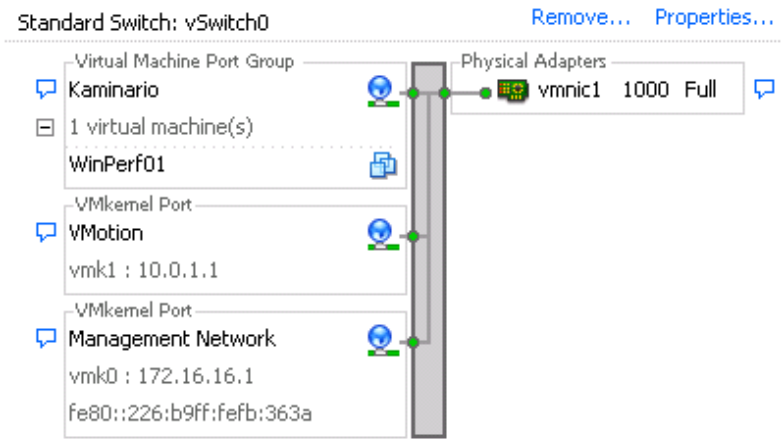

Figure 9: Network Diagram for External Network

A second internal network was defined to connect the Active Directory Domain Controller to the virtual desktops, as shown in Figure 10.

A third SAN network was used to connect the ESX host servers and the K2 array via an 8 Gb Fibre Channel (FC) fabric. Each host was connected to the fabric with four FC paths and the K2 was connected to the fabric with eight FC paths. All HBAs used for the testing were 8 Gb Fibre Channel from QLogic. The following configurations were set for the Qlogic FC driver:

- ql2xoperationmode=6
- ql2xintrdelaytimer=1

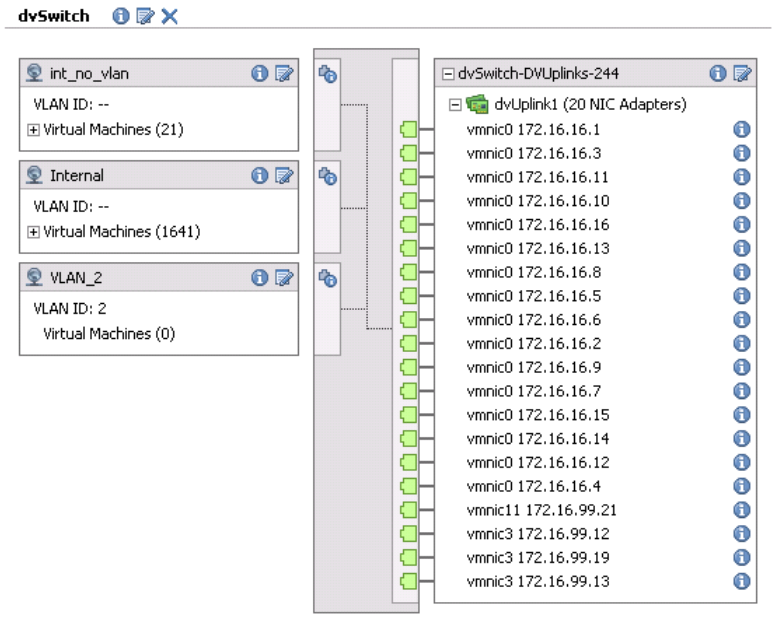

Figure 10: Network Diagram for Internal Network

VMware Configurations – Desktop Virtual Machines

For the provisioning tests, two pools of desktop virtual machines were deployed. Both the set of 1000 linked clones and the set of 1000 full clones were assigned to the same 10 datastores and the settings shown in Table 1, on the next page, were used.

Table 1: Settings for Linked Clone and Full Clone Virtual Desktops

Pool Parameter	Setting
Pool Type	Automated Pool
User Assignment	Floating
Desktop Power Policy	Take no action
Automatic Logoff	Never
Display Protocol	PCoIP
Monitors	2
Monitor Resolution	1920 x 1200
Adobe Flash Quality	Do Not Control
Adobe Flash Throttling	Disabled
Disposable Disk	None
Number of Desktops Powered On	1,000
Provision All Desktops Up Front?	Yes
Linked Clone Overcommit Policy	Aggressive
Dedicated Replica Storage	None (replica in each clone datastore)

Virtual desktops used in the View Planner tests were based on the View Planner 3.0 requirements and used the settings shown in Table 2.

Table 2: Settings for View Planner Virtual Desktops

Attribute	Specification
Desktop OS	Windows 7 Enterprise, SP1
Hardware	VMware Virtual Hardware version 7
vCPU	1
vMemory	1,024 MB
OS Pagefile	1.5 GB starting & maximum
vNICs	1
Virtual Network Adapter 1	VMXNet3 Adapter
Virtual SCSI Controller 0	LSI Logic Parallel
Virtual disk – VMDK	20 GB
Virtual CD/DVD/Floppy Drives	Removed
VMware View Agent	VMware View Agent 5.3
Installed Applications	Internet Explorer 9.0, Microsoft Office 2007

VMware Configurations – Management Virtual Machines

The management virtual machines were configured as follows:

- **vCenter** - The vCenter Appliance version 5.5 with an internal database was used with a configuration of eight vCPUs and 32 GB of memory.
- **View Manager 5.3** – The virtual machine was installed with Windows Server 2008 R2 and configured with eight vCPUs and 8 GB of memory.
- **View Composer 5.3** – The virtual machine was installed with Windows Server 2008 R2 and configured with eight vCPUs and 8 GB of memory.
- **Active Directory Domain Controller** – The virtual machine was installed on a Windows Server 2008 R2 running the Active Directory, a DHCP server and a DNS server and was configured with four vCPUs and 8 GB memory.
- **View Planner 3.0.1** – The virtual machine was configured with one vCPU and 3 GB of memory.

Software Configurations – Horizon View Virtual Disks

For testing purposes, 11 datastores were used, each of which was 4 TB in size. The following configuration was used for all disks:

- Disk.SchedQuantum = 64
- Disk.SchedNumReqOutstanding = 256
- CBRC was enabled on all hosts with CBRC.DCacheMemReserved = 2048

VStorage API for Array Integration (VAAI) Settings

All VMware vSphere Storage APIs – Array Integration (VAAI) primitives were supported by the K2 Array. All the ESX hosts were configured to use the following VAAI primitives:

- DataMover.HardwareAcceleratedMove
- DataMover.HardwareAcceleratedInit
- VMFS3.HardwareAcceleratedLocking

Native Vsphere Storage Multipathing

The Native Multipath was used with a custom Kaminario SATP (Storage Array Type Plug-in) rule. For all volumes, the "Round Robin" Policy was used with a path alteration policy setting of one I/O.

kaminario.

Contact

Contact a business development representative to answer any questions you may have.

Schedule a Demo

Schedule a demo with an engineer and learn if Kaminario's solution works for you.

Request a Quote

Request a quote for your application from our business development team.

About Kaminario

Kaminario, the leading all-flash storage company, is redefining the future of modern data centers. Its unique solution enables organizations to succeed in today's on-demand world and prepares them to seamlessly handle tomorrow's innovations. Only Kaminario K2 delivers the agility, scalability, performance and economics a data center requires to deal with today's cloud-first, dynamic world and provide real-time data access -- anywhere, anytime. Hundreds of customers rely on the Kaminario K2 all-flash array to power their mission critical applications and safeguard their digital ecosystem. Headquartered in Needham, MA, Kaminario works with an extensive network of resellers and distributors, globally.

For more information, visit www.kaminario.com