

Dell EMC XC Series Appliances – A Winning VDI Solution with Scalable Infrastructure

The linear scalability of the Dell EMC XC series appliances powered by Nutanix for VDI deployments.

Dell EMC Engineering
February 2017

Revisions

Date	Description
November 2016	Initial release
February 2017	Update branding

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

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

Copyright © 2017 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA. [2/2/2017] [00005]

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

Table of contents

Revisions.....	2
Acknowledgments.....	4
Executive summary.....	4
1 Introduction.....	5
1.1 Audience.....	5
2 Product architecture and infrastructure overview.....	6
2.1 Nutanix architecture.....	6
3 Infrastructure and test configuration.....	8
3.1 Hardware.....	8
3.2 Software.....	8
3.3 LoginVSI.....	9
3.4 Solution components.....	9
3.5 VDI I/O workload information.....	9
4 Test plan.....	11
5 Results and observations.....	12
5.1 Average IOPS (cluster-level analysis).....	12
5.2 Average IO latency (cluster-level analysis).....	14
5.3 User-level analysis.....	15
5.4 Node level analysis.....	16
6 Conclusion.....	18
7 References.....	19

Acknowledgments

Munir Ahmad is a Solution Engineer in the Dell EMC Converged Platforms & Solutions Division team. He contributes in planning all solution level reference architectures and best practices. In addition, Munir assists in pre-sales support, where he provides guidance on configuration and sizing for workload, and assists with proof of concept.

David Adams is a Storage Development Engineer in the Dell EMC Converged Platforms & Solutions Division team delivering the XC Series. His areas of interest include all things XC Series. He contributes to the development of the XC Series factory installation, testing, deployment, recovery and serviceability.

German Florez-Larrahondo is a Storage Architect in the Dell EMC Converged Platforms & Solutions Division team delivering the XC Series. His areas of interest include software defined storage and distributed systems. He contributes to the architecture and performance analysis of the XC appliances as well as other Software Defined Solutions within DELL EMC.

Executive summary

Virtual Desktop Infrastructure (VDI) enables servers to store and run desktop workloads, such as a Windows client operating system and associated applications. Users interact with the desktop presented on a user device to experience seamless access to virtual desktop and applications.

It's commonly known that many VDI deployments by different solution providers fail to scale when the number of desktops increase. Most of these VDI deployments use standard servers with SANs, which could result in performance bottlenecks. In contrast, the DELL EMC XC Web-Scale Hyper-converged Appliance uses local storage to keep compute right next to storage—the ideal solution for delivering incremental scalability.

The Dell EMC Converged Platforms & Solutions Division in Austin, TX, conducted a test to demonstrate such scalability of the XC Series appliances. The experiment explored the impact to user performance as the number of VDI users scaled in a real life scenario by adding more nodes to the cluster, keeping the number of Virtual Desktops (VDs) for each node constant.

By performing experiments on 4, 8, and 16 node clusters we concluded that the DELL EMC XC appliance supports the desired objectives of linear scalability for end-user performance.

1 Introduction

Desktop virtualization is an important strategy for organizations seeking to reduce the cost and complexity of managing an expanding variety of client desktops, laptops, and mobile handheld devices. VDI offers an opportunity not only to reduce the operational expenses for desktop management and provisioning, but also to improve user mobility and data security.

A VDI deployment places high performance and capacity demands on the storage platform. Many VDI solution vendors use a centralized infrastructure with centrally managed shared storage for VDI deployments. However, the Nutanix™ powered DELL EMC XC Web-Scale Hyper-converged Appliance infrastructure has a unique architecture, which lets enterprises incrementally scale their virtual desktops one node at a time starting from a couple hundred desktops to tens of thousands of desktops in a linear fashion, providing customers with a simple path to enterprise deployment. The DELL EMC XC Series system is a web-scale hyper-converged appliance based on the Dell PowerEdge server.

1.1 Purpose

This white paper demonstrates the linear scalability of desktop VDI architecture using DELL EMC XC Web-Scale Hyper-converged Appliance.

1.1 Audience

This document is intended for IT professionals, system administrators and solution architects who are interested in managing the cost of storage for their enterprise VDI deployments.

2 Product architecture and infrastructure overview

2.1 Nutanix architecture

The Nutanix web-scale hyper-converged infrastructure is a scale-out cluster of high-performance cluster nodes. Each node runs a standard hypervisor and contains processors, memory, and local storage consisting of high performance solid-state drives (SSD) and high capacity hard disk drives (HDD). Each node runs workloads in virtual machines on industry standard hypervisors like VMware ESXi or Microsoft Windows Hyper-v.

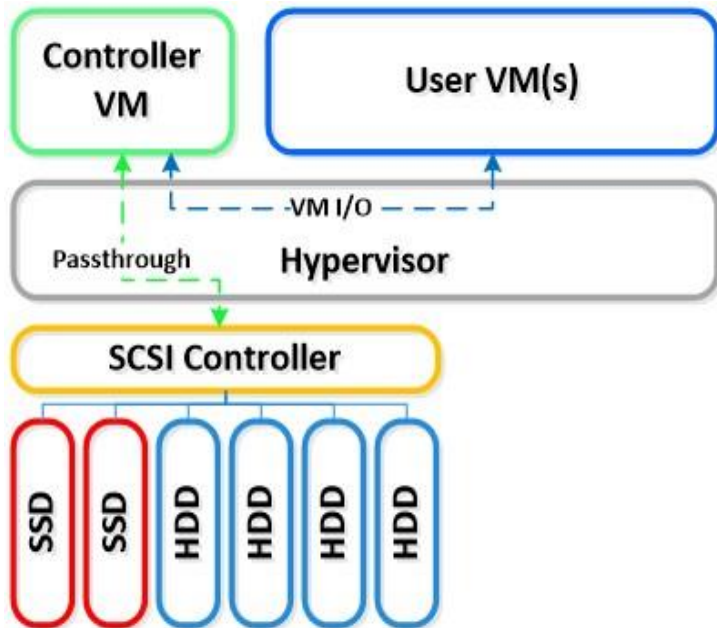


Figure 1 Nutanix web-scale hyper-converged infrastructure.

In addition, local storage from all nodes is virtualized into a unified pool by the Nutanix Acropolis Distributed Storage Fabric (DSF). In effect, DSF uses local SSDs and HDDs from all nodes to store virtual machine data. Virtual machines running on the cluster write data to DSF as if they were writing to shared storage.

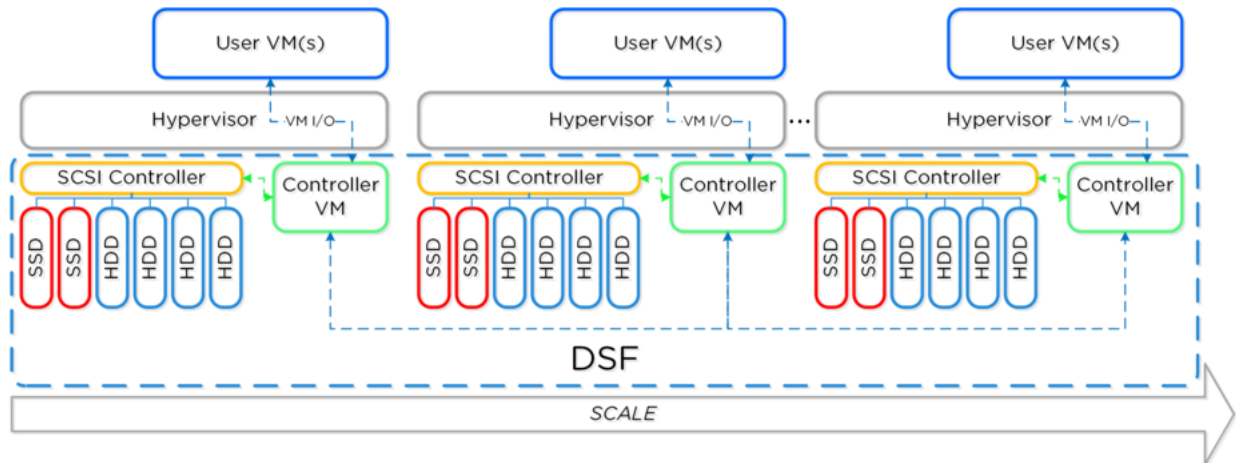


Figure 2 Nutanix Distributed Storage Fabric.

3 Infrastructure and test configuration

Technology has evolved to the place where servers with multiple cores can now keep up with the CPU demand for large numbers of VDI sessions. Also, servers can now hold enough RAM to support large numbers of VDI sessions as individual Virtual Machines (VMs). However, as the number of VDI sessions increase, the storage stack can become a bottleneck, resulting a poor desktop and application performance. Hence, optimal input/output operations per seconds (IOPs) and low IO latency for satisfactory desktop experience are cornerstones of successful VDI deployments.

The test in this white paper was conducted to investigate the performance of the VDI solution as it scales while keeping the number of users per node constant using three different cluster configurations. This document provides a clear analysis of the storage stack, such as IOPS and latency for each user, which affects real-world implementations. The test was run in a multi-run fashion in identical LoginVSI tests (125 users/node) on 4, 8, and 16 node clusters. The test also ran the Dell Performance Analysis Collection Kit (DPACK2) simultaneously during LoginVSI runs to provide extended host level statistics. The test also gathered Nutanix performance logs for storage statistics.

3.1 Hardware

Table 1 Hardware configuration for each DELL EMC XC630-10 node

Component	Type
Host CPUs	2x Intel® E5-2690 v3 (Advanced 2.6GHz, 3.5GHz Turbo, 135W TDP, 12C)
Host memory	384 GB memory @ 1866MHz
VDI cache and Data SSD	4 x 200 GB Intel SSDSC2BA20
VDI Data HDD	4 x 1TB 7.2k Seagate ST91000640SS
Network Daughter Card (NDC)	Intel X520 (Dual Port 10Gb SFP+ and Dual Port 1Gb BaseT) rNDC

3.2 Software

Table 2 Software components

Component	Type
Hypervisor	ESXi 5.5
Nutanix	NOS 4.1.2.1
Citrix XenDesktop	7.6 Embedded Database
LoginVSI	4.1.2
Microsoft Office (Gold image)	2010
MS Windows Active Directory	2012 R2

3.3 LoginVSI

LoginVSI is the industry standard tool of choice for testing VDI environments and server-based computing or terminal services environments. LoginVSI installs a standard collection of desktop application software on each VDI desktop (for example: Microsoft Office, Adobe Acrobat Reader and more). LoginVSI then uses launcher systems to connect a specified number of users to available desktops within the environment. After the user connects, the workload starts by running a logon script. This process starts the test script after the login script configures the user environment. Each launcher system can launch connections to a number of *target* machines (for example: VDI desktops) with a centralized management console managing the launchers while also configuring and managing the LoginVSI environment.

Our intent in this paper is to highlight the linear scalability of the DELL EMC XC architecture in VDI environment, we will discuss the cluster level, node level and user level performance as we increase the number of nodes in a cluster. As pointed in Section 4.0, Login VSI results for a similar configuration are presented in the TechCenter blog post [Dell Solution Center Engagement Update – 2,000 seat VDI on 16-node Nutanix, found in Dell Tech Center](#).

3.4 Solution components

Table 3 Appliance settings.

Component	Type
LoginVSI Launchers	110, ~20 sessions for each launcher
File servers	6 file servers; Windows shares; 10 GB
VMware vCenter	1 Windows VM, embedded database
Microsoft Office (Gold image)	2010
Gold Image OS	Windows 7 SP1 64-bit
Citrix XenDesktop Provisioning Mode	Machine creation services (MCS)

3.5 VDI I/O workload information

The challenge with VDI benchmarking is providing a representation or simulation of actual VDI users in a real-world scenario as there is a broad spectrum of VDI users. At one end of spectrum, the user may be using simple Microsoft Office applications at a relatively moderate speed. At the other end of spectrum, the user may be running multimedia applications and switching between many applications at a much faster rate. In order to simplify the analysis, the DELL EMC XC platform solution splits VDI user workloads into three categories shown here.

Table 4 Customer user groups.

User Profile	VM vCPU	VM Memory	OS image notes
Office worker	2 vCPU	2 GB RAM	This user workload leverages a shared desktop image and emulates a task worker. Only two applications are open simultaneously and session idle time is approximately one hour and 45 minutes.
Knowledge worker	2 vCPU	2 GB RAM	This user workload leverages a shared desktop image and emulates a medium knowledge worker. Up to five applications are open simultaneously and session idle time is approximately two minutes.
Power worker	2 vCPU	4 GB RAM	This user workload leverages a shared desktop image and emulates a high-level knowledge worker. More than five applications are opened simultaneously and session idle time is two minutes.

In this experiment, we used Knowledge worker as our VDI user profile. The Knowledge worker user profile consists of email, typical office productivity applications and web browsing for research or training. The workload requirement for a Knowledge worker is moderate and most closely matches the majority of office worker profiles in terms of CPU, memory, network and Disk I/O.

4 Test plan

Tests were run on a 4, 8, and 16 node clusters with the same VDI configuration in a 3-run per cluster fashion.

The base storage performance was validated first for each of these 4, 8, and 16 node clusters. To keep the workload the same each node was assigned 125 VDI sessions. In a previous experiment performed by the Dell EMC Converged Platforms & Solutions Division on a 16 node Nutanix cluster, we have seen that 125 VDI sessions per node with similar hardware and software configuration results an optimal usage of about 90 percent for CPU and memory. For more information, see [Dell Solution Center Engagement Update – 2,000 seat VDI on 16-node Nutanix, found in DellTech Center](#). For each cluster variation, three LoginVSI runs were performed with two hour idle times between each run.

5 Results and observations

The test cases were designed to simulate a real-life scenario, where a system is busy for few minutes and then the system experiences idle time after that. Test data was sampled and collected at every 30 second mark. As such, most of the data pattern was repeated bi-modal as data was mixed with idle and active states. There were distinct peaks during active states and distinct valleys during idle states. This type of performance over a time series graph is classified as multimodal distribution.

Due to the nature of multimodal distribution of the data, in most cases, the density function is plotted for the analysis. For simplicity, the following colors were chosen for the charts throughout this paper, where applicable.

- Black – 4 node cluster
- Red – 8 node cluster
- Green – 16 node cluster

Typically, when measuring VDI performance, the main focus are the input/output (IOPS) operations per second, since a good VDI solution must handle a high number of IOPS in a real life application. However, IOPS alone do not tell the whole story. In this paper, we looked at both IOPS and latency parameters in detail. IOPS and latency parameters were taken from PRISM UI (a graphical interface developed by Nutanix to manage virtual environments).

5.1 Average IOPS (cluster-level analysis)

The graphs below show the number of IOPS for various clusters configurations (4, 8, and 16 node) in a scatterplot where IOPS were plotted against time. They show that sustained IOPS practically doubled between the 4 node and 8 node, and again between the 8 node and 16 node cluster.

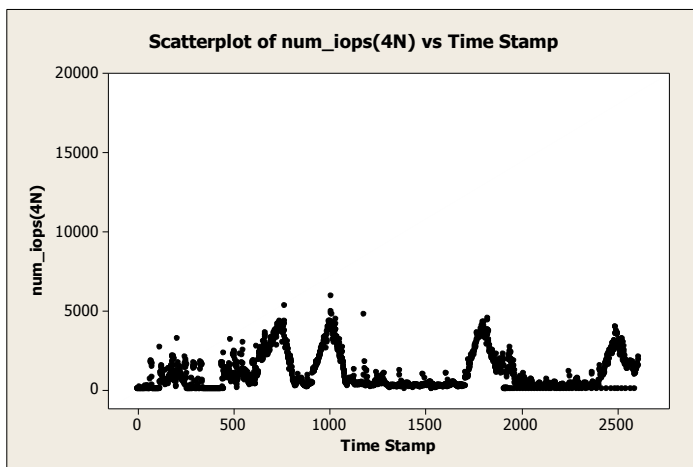


Figure 3 Average IOPS for 4-node cluster configuration

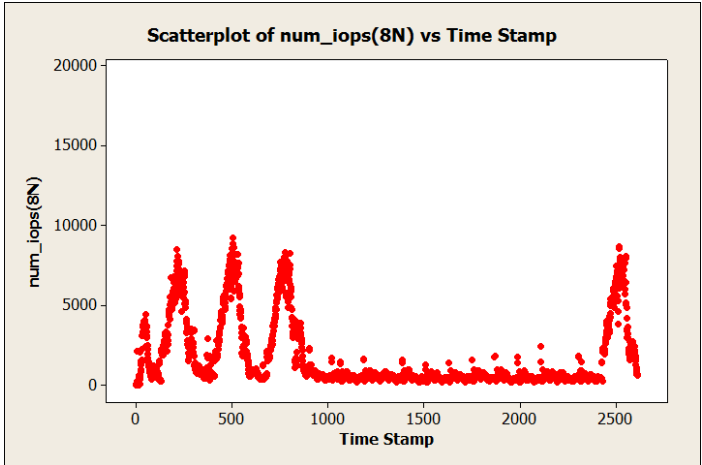


Figure 4 Average IOPS for 8-node cluster configuration




Figure 5 Average IOPS for 16-node cluster configuration

To look at these sustained IOPS in more detail, and to provide statistical significance, the density function over time (idle and active state) is plotted in Figure 6. The graph here shows that the number of IOPS nearly doubles as the number of nodes doubles. We observed that a 16-node cluster is performing slightly better than expected in terms of IOPS. For example, the 8-node cluster IOPS are less than ~ 8000 or below 99 percent of the time. For a 16-node cluster, the number of IOPS is less than ~17000 or below 99 percent of the time. Note that the means and standard deviation of the distribution are not meaningful for these analysis as we are dealing with a multimodal distribution.




Figure 6 Density function for IOPS at cluster for various cluster configuration

5.2 Average IO latency (cluster-level analysis)

The scatterplot here shows the average IO latency for three different cluster configurations in a time series fashion. The graph shows that the average IO latency for a 16-node cluster has a much tighter distribution than the 8-node cluster, followed by the 4-node cluster. Note that there were a few high-latency data points or outliers that are more prominent for the 4-node cluster.

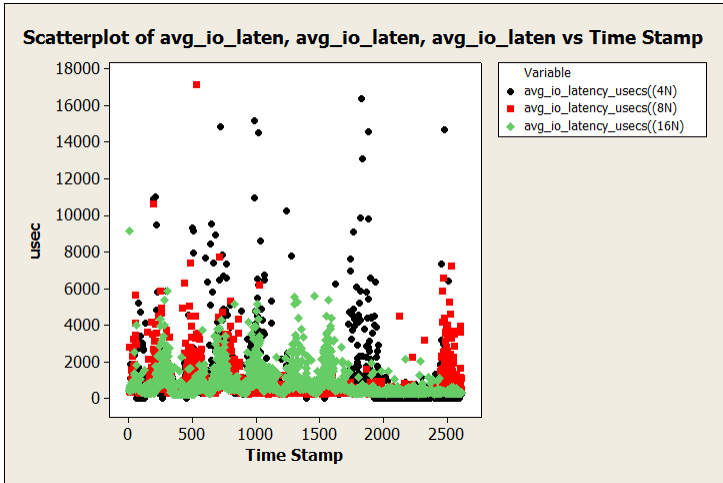


Figure 7 IO latency scatterplot for various cluster configurations

The density chart in Figure 8 shows overall cluster latency distribution with the x-axis showing latency in microseconds. Using a 95 percent mark as an arbitrary checkpoint, all 3 cluster configurations have similar latency. However the latency distribution shows that number of occurrences (a small percentage) of high latency IO operations are higher for a 4 node cluster when comparing with 8 node and 16 node cluster.

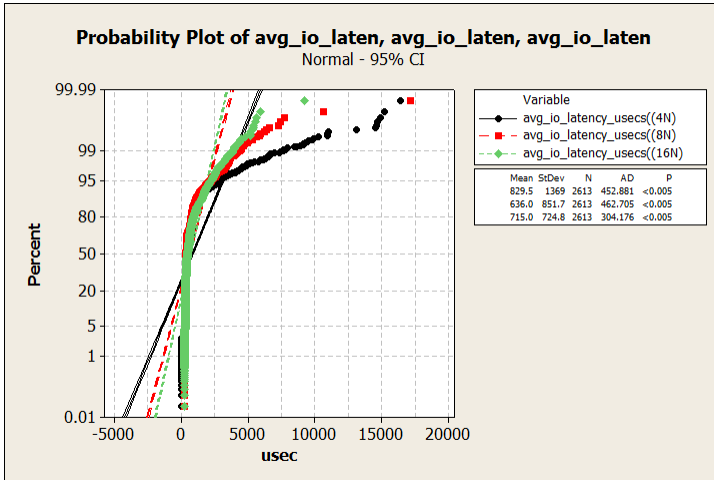


Figure 8 IO latency density plot for various cluster configurations

5.3 User-level analysis

As described before, in our tests the workload for each node never changes regardless of the cluster configuration (each node was assigned 125 VDI sessions with the same user profiles). Therefore we hypothesized seeing uniform and harmonic distribution of performance data from a user perspective. Figure 9 shows the scatterplot and density plot of average latency from a user viewpoint for three different clusters (4, 8, and 16 node clusters). Figure 9 shows no statistical significant variation between each series of data confirming our hypothesis. Similarly Figure 10 shows no statistical significant variation in IO bandwidth from a user point of view.




Figure 9 Scatterplot and density function of average IO latency for various clusters from user point of view

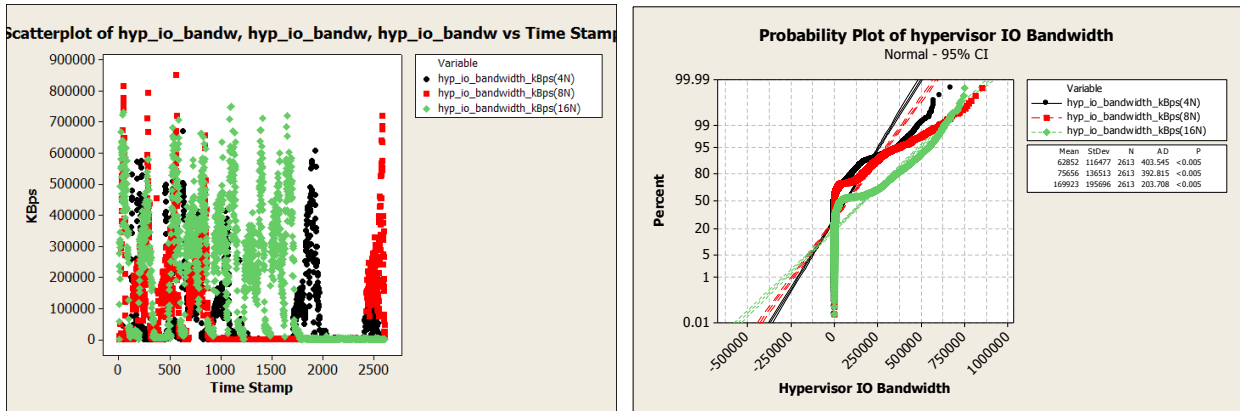


Figure 10 Scatterplot and Density Function of average bandwidth for various clusters from user point of view

5.4 Node level analysis

Since in our tests the workload for each node never changes regardless of the cluster configuration, we also theorized that node-to-node storage performance would be similar. The graph below shows the density plot of average IO latency for three different randomly selected nodes (one node from each cluster configuration). As shown here, no significant variation is found in IO latency while comparing these random nodes.

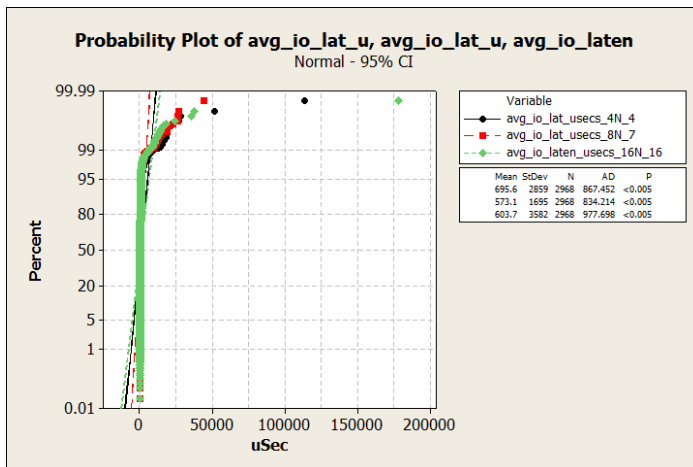


Figure 11 Density plot of average IO latency for three different randomly selected nodes

We also plotted the data of average node IOPS for all three clusters. The graph here shows the density plot of average sustained IOPS per nodes for three cluster configuration. As shown in Figure 12, no significant variation is found in average IOPS while comparing nodes for different cluster configurations.

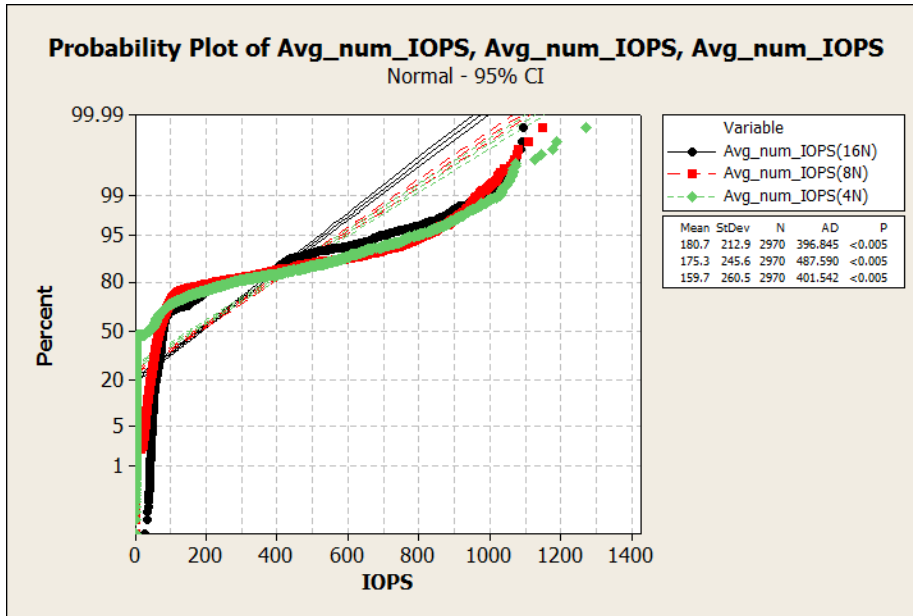


Figure 12 Density Function of average IOPS for 3 different randomly selected node

6 Conclusion

Large and small enterprise companies are looking to deploy virtual desktop infrastructure in the most cost-effective and flexible manner possible. Many IT organizations have delivered successful VDI projects; however, as they scale the storage infrastructure the IO performance degrades. This causes poor desktop experiences. Our tests have shown that the DELL EMC XC Series solutions provide linear scalability in the performance of the VDI deployment as the number of server's increases to satisfy more virtual desktop sessions. Therefore, the DELL EMC XC Series VDI solution lets IT independently and non-disruptively scale VDI linearly, which can consistently deliver a high performance desktop experience.

7

References

- Dell Storage with Microsoft Storage Spaces SOFS for Wyse vWorkspace VDI (http://en.community.dell.com/techcenter/extras/m/white_papers/20441734)
- Desktop Virtualization and Healthcare
- <http://www.nutanix.com/resources/white-papers-ebooks/>
- Top Challenges & Best Practices for Virtualizing Citrix XenApp & XenDesktop on VMware vSphere
- <http://content.eginnovations.com/top-challenges-best-practices-for-virtualizing-citrix-xenapp-xendesktop>
- Enterprise Virtual Desktop Infrastructure: Design for Performance and Reliability
- <http://www.emc.com/collateral/hardware/white-papers/white-paper-vdi.pdf>