Demartek

Dell EMC XC430 Xpress Hyper-Converged Solution Overview

Complete all-in-one solution for small and medium-sized businesses.

Executive Summary

The new Dell EMC XC430 Xpress appliance combines enterprise-class features of the XC Series appliances with unique, price-saving innovations to bring an affordable, all-in-one infrastructure solution to small and medium-sized businesses.

The XC430 Xpress incorporates the Intel® Xeon® E5 processors to create a powerful solution wrapped up in a small package. Each node consolidates servers and storage in a 1U rack mount chassis, so when deployed in a 3 or 4 node cluster smaller companies are about to dramatically reduce infrastructure footprint down to as little as 3U of rack space.

The XC430 Xpress ships from the factory with a complete and up-to-date software stack that is ready to install out of the box for easy deployment. In our experience, installation is a quick and painless process. It is easy enough that Dell EMC has given customers a choice to self-install their XC Xpress clusters and provides a variety of technical resources to support the customer throughout the deployment.

The XC430 Xpress is powered by Nutanix software and managed through the Nutanix Prism interface and customers have their choice of hypervisors.

For Microsoft users, XC430 Xpress has tight integration with Hyper-V and is able to support any virtualized workload. Also worth noting is that Dell EMC includes seamless integration with Azure Backup Service with these XC Xpress appliances. This is something new with the Dell EMC HCI solutions, and compliments the simplicity of the XC430 Xpress quite nicely. Users can now back up to the Azure cloud to and have a peace-of-mind that mission-critical data will be safely backed-up offsite. Billing and set-up are handled through Dell EMC, to make it very easy on the end user.

Key Findings

The Dell XC430 Xpress appliance offers multi-hypervisor support for Hyper-V, ESXi, and KVM (AHV) and ships complete with pre-installed hypervisor that requires no additional software or licensing.

> Installation of the system was lead through prompts that took us from “bare metal” to a working three-node cluster with minimal effort.

> In our tests, we ran a set of mixed workloads in 9 virtual machines on the XC430 Xpress appliance, including a mixture of Microsoft SQL Server, Microsoft Exchange Server, file and print server workloads and simulated web server workloads across the three-node cluster to mimick the day-to-day workload of a small to medium-sized company.

> The single XC430 Xpress cluster with the help of Intel® Xeon® E5 processors was able to handle all these workloads efficiently still leaving headroom for additional work.

Dell EMC XC430 Xpress
Hyper-converged Solution Overview

Installation

The first phase of the validation focused on the out-of-the box experience, including power-on and completing the initial configuration. We installed the XC430 Xpress cluster from a bare metal factory fresh state. Powering on the system brought us to a Microsoft Windows installer, which started the installation process. This installer loaded the Nutanix operating environment, and we were led through a series of prompts until a discovery launcher started and identified the three nodes for our cluster.

A cluster setup screen was shown, and we provided the name of our cluster, our local time zone, our cluster IP address and network gateway IP address. After the cluster setup screen was completed, we were brought to a node setup screen, where we provided the names and IP addresses of our three nodes and the type of hypervisor. We selected the single hypervisor option and chose Hyper-V for our hypervisor. Because the cluster ships with all the hardware, hypervisor and software components, the entire process of bringing up cluster from power-on took less than 30 minutes.

At this point, we created our virtual machines (VMs), loaded the operating systems (Windows Server 2012 R2) onto the VMs and then our applications. From this point, we ran our applications in the normal fashion.

Applications and Virtual Machines

We created nine virtual machines, three for each node and selected workloads that would be typical of a small or medium-sized business. The workloads were distributed across the nodes, one per VM, as follows:

Node 1

- > Exchange Jetstress 2013
- > File/Print workload
- > Web server simulation

Node 2

- > SQL Server 2016 OLTP workload
- > File/Print workload
- > Web server simulation

Node 3

- > SQL Server 2016 OLTP workload
- > File/Print workload
- > Web server simulation

The Exchange Jetstress workload running on node 1 was less CPU intensive but more I/O intensive that the other workloads. See the performance results for additional discussion of these workloads.

Dell EMC XC430 Xpress
Hyper-converged Solution Overview

Demartek
June 2017

Solution Reference Architecture

Technical Requirements

As the use of clustered servers grows in HCI and similar environments, we find certain common technical characteristics in these types of deployments.

Processor and Memory

We recommend any of the one or two-socket server processors, such as the Intel® Xeon® E5, depending on the desired workload levels and/or price points. We tested on servers with low-end, entry-level single-processors. The 128 GB of memory was more than enough to run our test workloads.

Storage

For business applications, we have found it very advantageous to have the operating environment and applications installed on some form of non-volatile storage, such as SSDs. In our opinion, this has become a minimum requirement for server workloads.

Networking

In our opinion, another minimum requirement for today's business workloads is a 10Gbps or faster network, especially for server clusters. For resiliency, we prefer to connect application servers to two different switches.

Technical Configuration Details

Each of the three nodes of the cluster included Dell EMC XC430 Xpress servers with the following:

- One Intel® Xeon® E5-2620 v4, 2.1 GHz, 8 cores, 16 threads
- 128 GB RAM
- One 960 GB SSD
- Three 6 TB HDD

In addition, a separate virtual machine running on separate server was used as the Domain Controller, DNS server and NTP (time) server.

Two Dell EMC S4048-ON 10GbE switches were used to connect the servers of the cluster to each other. Each server node was connected to both switches for redundancy purposes. In the event that one switch fails, or a connection to one switch fails, the cluster remains up and running.

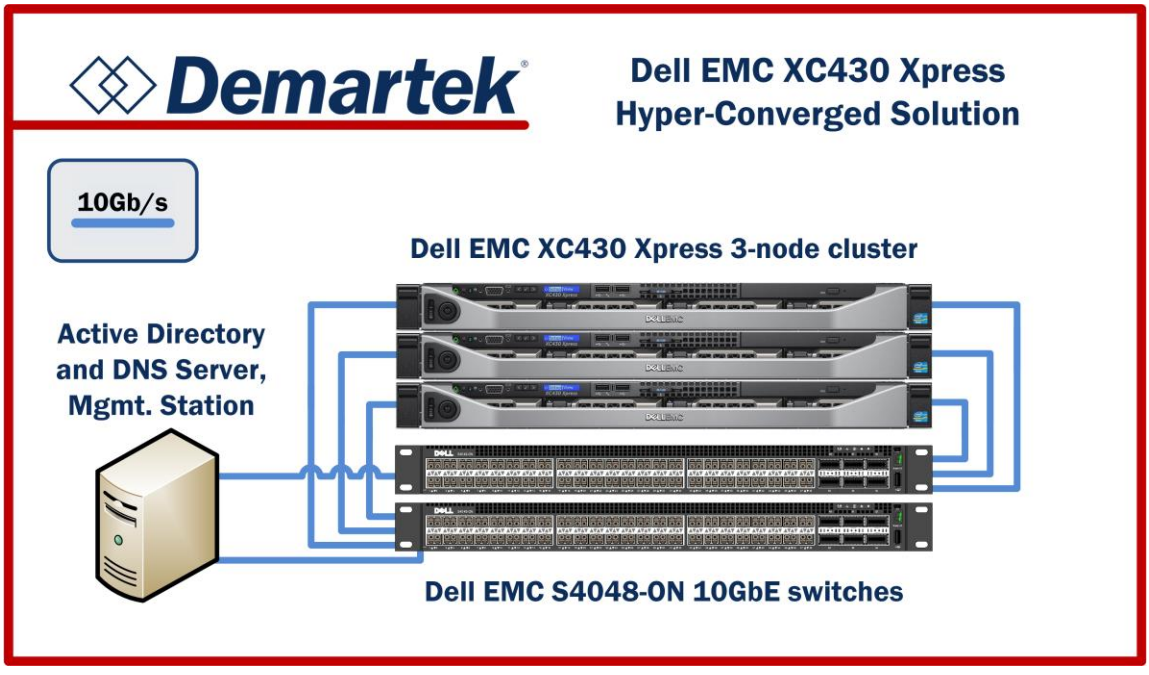

Dell EMC XC430 Xpress Hyper-converged Solution Overview

Performance Results

CPU Utilization

CPU utilization is a measure of the amount of CPU processing that occurs as workloads run on a system. Mixed workloads running at the same time on the same physical server put different loads on the CPU at different times and these fluctuating loads are reflected in the graphs.

Some workloads are more CPU intensive than others, while some are more I/O intensive. Looking at the first chart on the right, the overall CPU utilization from node 1 appears to be different than the other two. This is because the Exchange Server workload is more I/O intensive than the SQL Server workload running on the other two nodes. The average peak workload on each node was approximately 20% CPU utilization.

In order to see the cumulative effect of all the workloads running across all the nodes, we provide the cumulative utilization in the second graph. It should be noted that in theory, the cumulative utilization graph could potentially rise to 300% because there are three nodes involved.

The relatively low CPU utilization rates shown in the graphs indicate that the XC430 Xpress cluster was able to support all workloads efficiently with plenty of room left for future growth.

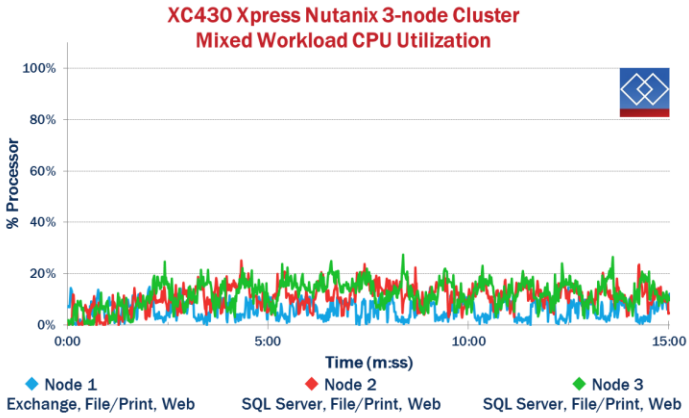

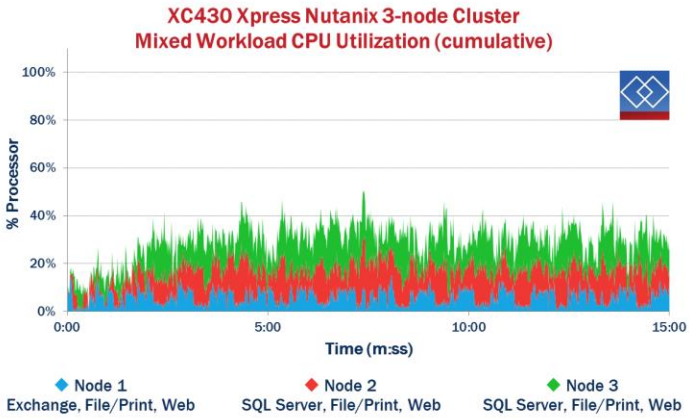

Memory Utilization

Memory utilization is a measure of the amount of memory used by the applications running. Because the applications were started approximately simultaneously, they acquired and held memory for the duration of the tests, so the graphs are flat.

Because identical workloads were running on nodes 2 and 3, the amount of memory used was the same. As a result, the red and green lines are at the same level.

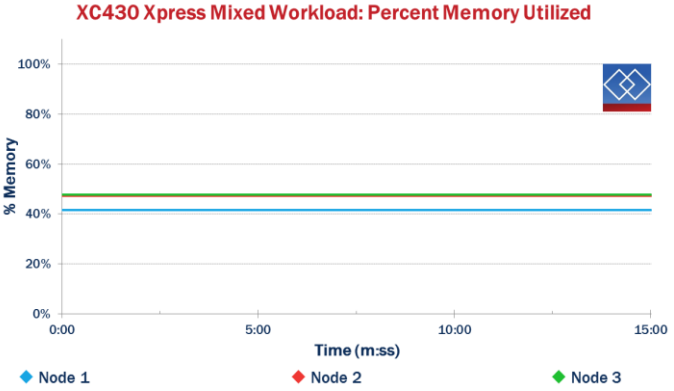

Dell EMC XC430 Xpress Hyper-converged Solution Overview

Read and Write Throughput

Throughput is a measure of the amount of data transferred to or from a storage device or set of storage devices, per second. This is typically measured in kilobytes per second (KBPS), megabytes per second (MBPS) or gigabytes per second (GBPS). Storage performance metrics are measured in bytes per second. By contrast, network performance is measured in bits per second.

Storage performance also frequently separately identifies reads and writes. Reads are requests from hosts to retrieve data from storage and writes are requests by hosts to send data to storage. Many workloads are asymmetrical with respect to reads and writes – they retrieve different amounts of data from storage than they send to storage.

Looking carefully at the ***Read Throughput*** chart, it is apparent that the blue workload, on average, read slightly more data than the red and green workloads. This was due to the Exchange Jetstress workload which is different than the SQL Server OLTP workloads running on nodes 2 and 3.

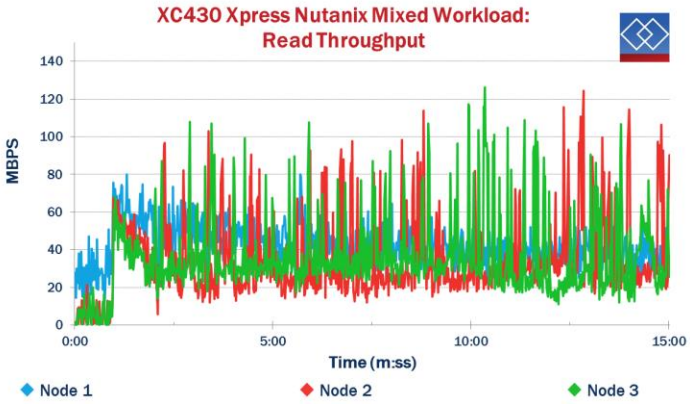

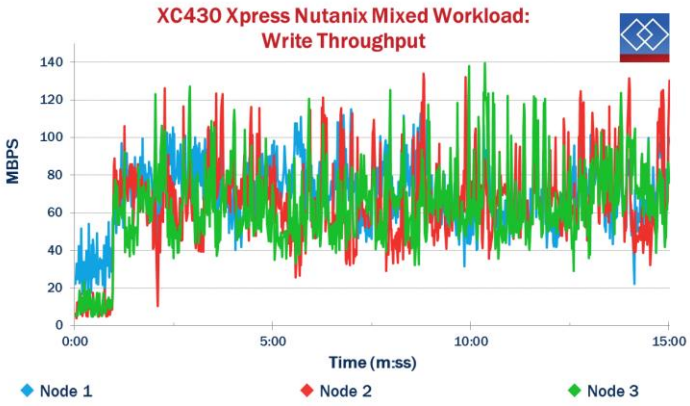

Dell EMC XC430 Xpress Hyper-converged Solution Overview

Reads and Writes per second

Another common storage performance metric is reads and writes per second, or I/Os per second (IOPS). This is a measure of the number of I/O requests issued by a host server per second.

Workloads that are transaction oriented tend to have higher IOPS rates than bulk transfer types of applications such as backup or video streaming.

As is typical for real-world workloads, IOPS rates are not flat but fluctuate, sometimes with occasional bursts. In our testing, depending on the combination of workloads we ran, we observed IOPS rates in the tens, hundreds and thousands on this XC Xpress cluster. The graphs below are typical for all the combinations we ran.

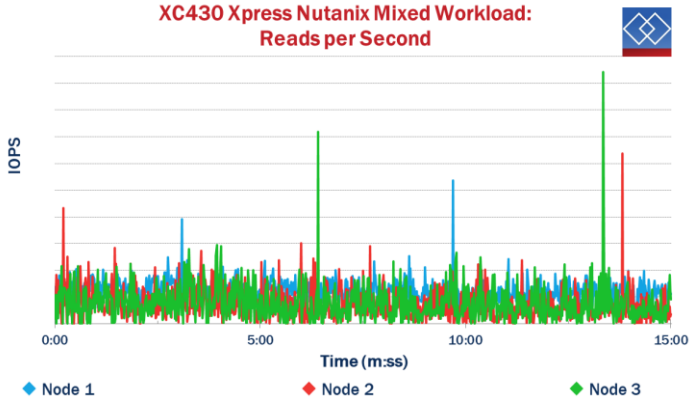

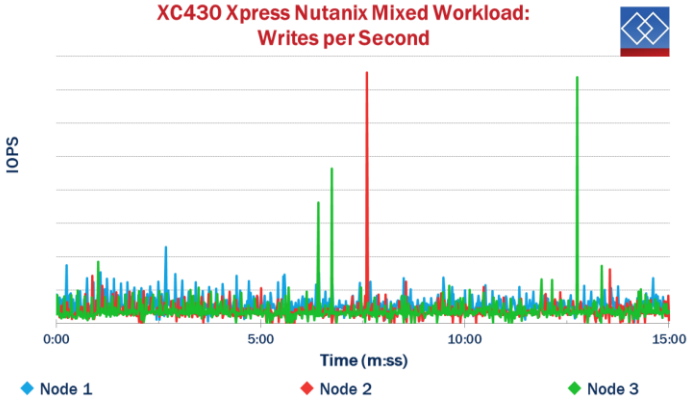

Read and Write Queue Length

The queue lengths are a measure of the number of read and write requests that get queued waiting for resources to become available.

The queue lengths can rise either because the storage is getting busy handling many requests or the host processor is extremely busy. The workloads running on node 1 had higher write queue lengths because of the nature of the Exchange Jetstress workload.

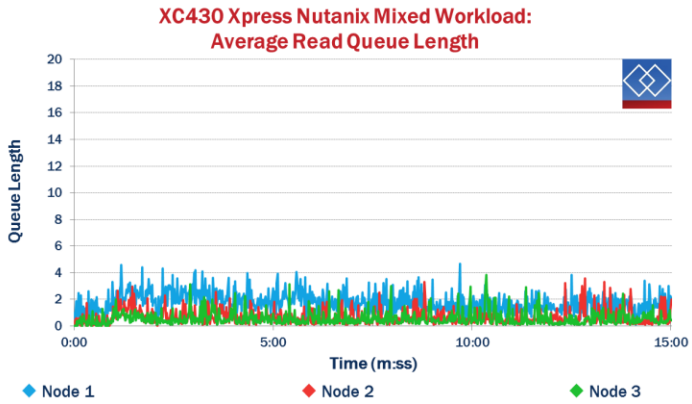

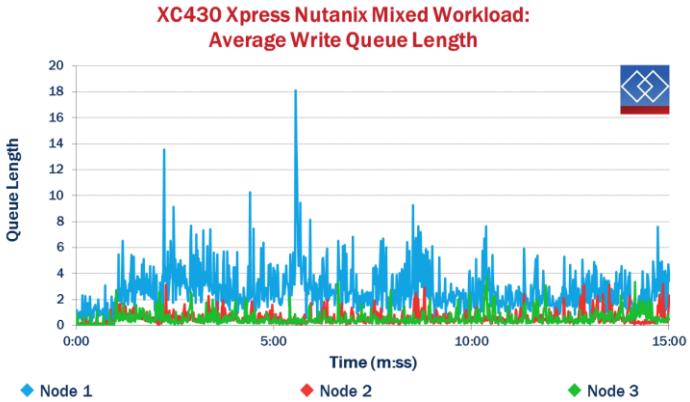

Dell EMC XC430 Xpress Hyper-converged Solution Overview

Read and Write Latency

Latency is a measure of the time taken to complete an I/O request, also known as response time. Latency is frequently measured in milliseconds (ms), which is thousandths of a second.

Latency can be critical even though it doesn't necessarily have a direct effect on IOPS and throughput. It does have a very significant effect on application performance and user experience. Unlike IOPS and throughput, where more is better, with latency, the goal is to keep it as low as possible. The impacts of latency vary with the workload deployed. Some applications have a greater tolerance, while other applications are negatively impacted by even small increases in latency.

For systems with hard disk drives, which this XC430 Xpress cluster has, these latencies are reasonable. Typically, if lower latency is important, moving from hard disk drives for data storage to SSDs for storage often provides significant improvements (lower latency).

The workloads run on node 1, specifically the Exchange Jetstress workload, had higher read and write latencies because of the nature of this workload. In particular, this workload uses a larger block size, which typically takes a bit longer to satisfy.

By contrast, the nodes running the SQL Server OLTP workloads (red and green lines on the chart) had very low latency due to the transactional nature of these workloads.

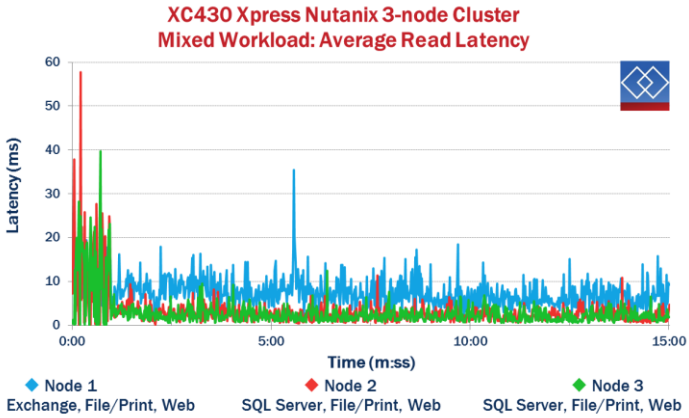

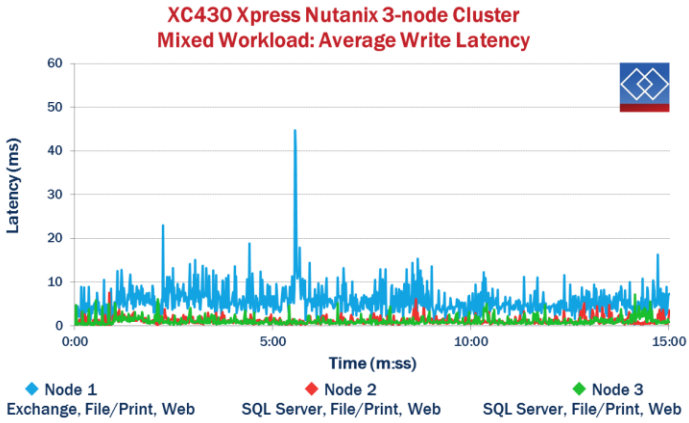

Dell EMC XC430 Xpress Hyper-converged Solution Overview

Demartek
June 2017

Microsoft Azure Services

Dell EMC is enhancing the robustness of Dell EMC XC430 Xpress Hyper-converged appliance infrastructure (HCI) environments by integrating with Microsoft Azure backup solutions. This solution helps administrators reduce the complexity of their data protection environment while adhering to data protection best practices.

Dell EMC provides streamlined on-boarding and set-up procedures for Microsoft Azure Cloud Services, including Azure Backup. These services require no capital expenses, but are pay-as-you-go, no contract, monthly services based on consumption.

Two of the Azure backup services include:

> Microsoft Azure Recovery Services (MARS)
> Microsoft Azure Backup Server (MABS)

Microsoft Azure Recovery Services (MARS)

The Microsoft Azure Recovery Services Agent is installed in the individual virtual machines or on the parent partition and can protect application files and folders. This service backs up the data directly to the Microsoft Azure cloud.

This backup service is not application aware, but provides protection at the file, folder and volume level only.

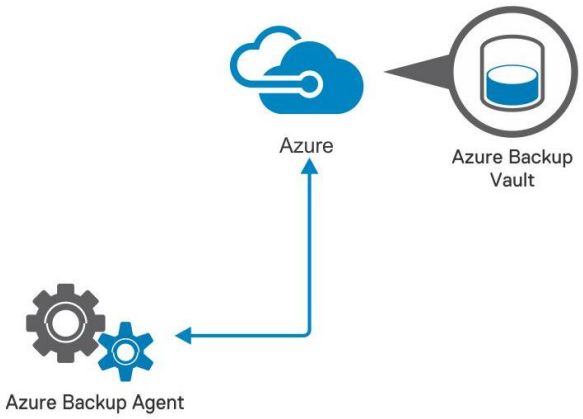

Microsoft Azure Backup Server (MABS)

The Microsoft Azure Backup Server (MABS) is a separate server installed on-premise that can protect application workloads along with files, folders and volumes. This backup data is stored on a locally attached disk or volume. The backed up data can also be sent to the cloud for additional protection.

The MABS server must be joined to the domain of the XC430 Xpress cluster.

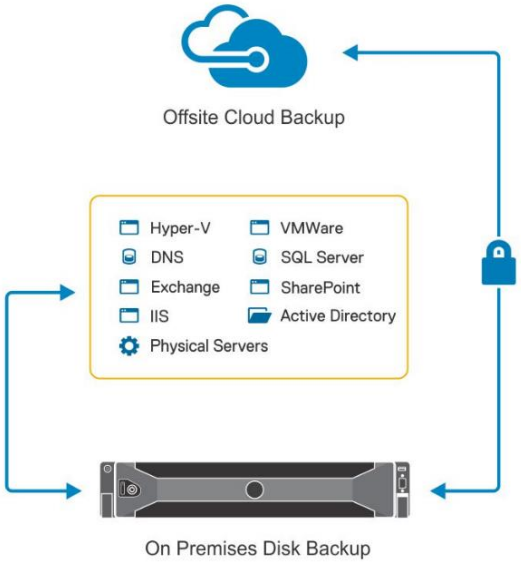

MARS & MABS Workloads

We installed the MARS agent and configured a separate MABS server to test backing up data to the Microsoft Azure cloud.

Using these tools, we were able to backup files, folders, SQL databases and Exchange databases from the nodes of the XC430 Xpress cluster to the Azure cloud, and were able to restore them. The speed of the backup and restore depends on the speed of the Internet connection to the cloud.

Dell EMC XC430 Xpress Hyper-converged Solution Overview

Demartek
June 2017

Summary and Conclusion

The Dell XC430 Xpress, powered by Intel® Xeon® E5 processors and Nutanix software, was simple to install and get up and running. Once up and running, it was simple to manage the VMs and get a view of the entire system, including activity on each of the nodes, free storage space and several other cluster-level criteria using PRISM.

We ran a set of mixed workloads in nine virtual machines spread across three nodes of the XC430 Xpress cluster. These workloads were typical of small and medium-sized businesses, including file and print, web server, SQL Server and Exchange Server.

The single XC430 Xpress cluster with the help of Intel® Xeon® E5 processors was able to handle all these workloads efficiently still leaving headroom for additional work.

In our opinion, the XC430 Xpress Hyper-converged Solution should be considered for its ease of setup, choice of hypervisors and ability to handle a mixed workload. Ready to deploy out of the box, the XC430 Xpress is a complete all-in-one solution that consolidates server, storage, and virtualization in a single cluster designed to simplify IT infrastructure for small to mid-sized companies. Utilizing proven enterprise class features that have been optimized for smaller environments, the XC430 Xpress provides smaller customers a simple, reliable, and affordable onsite infrastructure solution.

The most current version of this report is available at www.demartek.com/Dell-XC430-Xpress on the Demartek website.