dataON

Hyper-Converged Infrastructure Solutions for Windows Server 2016

Optimized Deployment of Hyper-V, Storage Spaces Direct and Azure Stack

Microsoft

dataON

Contents

dataON

Executive Summary

Hyper-Converged Infrastructure (HCI) is one of the hottest topics in data center architecture today. In the simplest terms, HCI combines computing, storage, and networking components into a single physical instance. HCI evolved into its current form from the blade servers of the last decade, and from the purpose-built data centers developed by hyperscale companies such as Google, Amazon Web Services (AWS), and Microsoft Azure. Today, nearly all of the leading information technology (IT) companies develop and sell HCI appliances, with acknowledged leaders including Nutanix, Simplivity, Dell/EMC, HPE, Cisco, Oracle, VMWare, and NetApp.

HCI originally developed around two use cases: remote office/branch office (ROBO), and "stand-up/ scale-out" data centers. In both cases, the key benefits of HCI were:

- The ability to quickly deploy IT resources, typically with low-skill personnel; and
- The ability to remotely manage/provision hardware/software ("lights-out" administration).

While these benefits are compelling for ROBO and "lights-out" data centers such as those operated by hyperscale companies, they also have many benefits for medium-sized enterprises who have limited IT staff. Moreover, HCI appliances make scaling of applications easier for medium-sized enterprises by simplifying the addition and provisioning of additional compute, storage, and networking resources, a task that can be challenging and labor-intensive under conventional IT architectures. The greatest challenge with HCI, especially for medium-sized enterprises, has often been finding a solution that is not too large (and expensive) for their needs. Many of the leading HCI solutions are rack-level products. Moreover, they often have expensive software licenses that often offset much of the CapEx and OpEx gains realized by utilizing HCI solutions over conventional IT architectures.

This is exactly the situation that EVGA, the leading hardware provider of high-end hardware components for PC gaming and video, found themselves in. They wanted to move away from their existing VMWare-based rack-level architecture, which consumed 26U of rack space in their data center, and move to a solution that provided more VMs and performance while reducing footprint, power, and costly software licenses. EVGA chose DataON's S2D-3110 platform, which combined Intel E5-2600-based servers, Microsoft Windows Hyper-V 2016, Microsoft Storage Spaces Direct, Intel NVMe-based Solid State Drives (SSDs), and DataON's Management Utility Software Tool (MUST) framework into a unified HCI solution. The result was more VMs, improved performance, a reduction of footprint from 26U to 7U (73% smaller), and an elimination of their VMWare license fees.

> ***"These new Hyper-Converged systems offer additional features and functions and most importantly increase reliability at an affordable cost,"*** said Joe Darwin, Chief Branding Officer for EVGA Corporation. ***"Working with companies like Microsoft, Intel, and DataON gives our IT team a good foundation to provide a solid infrastructure for future growth planning."***

As a leading provider of Hyper-Converged Cluster Appliances (HCCAs) for Microsoft environments, DataON is dedicated to providing the advantages enjoyed by EVGA to its customers. DataON HCCA solutions are built with the single purpose of rapidly and seamlessly deploying Microsoft applications, virtualization, data protection, and hybrid cloud services to meet the rapidly changing needs of our customers. See how DataON can help your organization or enterprise improve application performance while reducing CapEx and OpEx costs by visiting our website at www.dataonstorage.com.

dataON

Introduction: Hyper-Converged Infrastructure Solutions for Windows Computing Environments

This white paper explores the concept of the Hyper-Converged Infrastructure (HCI) computing model, and how the HCI model can be successfully deployed in Microsoft Windows computing environments. In particular, we investigate how DataON, Intel, and Microsoft partnered to implement a new class of Windows Server 2016 Certified solutions called Hyper-Converged Cluster Appliances (HCCAs). DataON HCCAs are designed and built with to provide a high-performance, easy-to-deploy, easy-to-manage scale-out solution for Microsoft environments. The paper includes an investigation of the following topics:

- Basic concepts of HCI, including adoption trends and challenges deploying HCI solutions
- How to go about choosing the right partners to architect and deploy an HCI solution
- The benefits that EVGA (a leading hardware provider of high-end hardware components for PC gaming and video) realized by deploying the DataON-Intel-Microsoft HCCA solution.
- The product portfolio that DataON, Intel, and Microsoft can apply to the HCI challenges of companies and organizations.

As a leading provider of Hyper-Converged Cluster Appliances (HCCAs) for Microsoft environments, DataON is dedicated to providing the advantages enjoyed by EVGA to its customers. Our solutions are built with the single purpose of rapidly and seamlessly deploying Microsoft applications, virtualization, data protection, and hybrid cloud services to meet the rapidly changing needs of our customers. See how DataON can help your organization or enterprise improve application performance while reducing CapEx and OpEx costs by visiting our website at www.dataonstorage.com.

Basic Concepts of Hyper-Converged Infrastructure

The Evolution of Data Center Computing Towards Hyper-Converged Infrastructure

The information technology (IT) industry has gone through several waves of computing and application architectures, and each new wave has resulted in significant changes in how compute, storage, and networking resources have been designed, utilized, and deployed. Table 1 below illustrates how these changes have affected data center architectures.

Decade	1970	1980	1990	2000	2010
Model	Centralized	Distributed	Shared	Virtualized	SDDC
Architecture	Mainframe/Mini	Client/Server	x86 Network	Cloud/HPC	Hyper-Converged
Applications	Centralized	Distributed	Shared	Disaggregated	Virtualized
Compute	Centralized	Distributed	Distributed	Application Aggregated	Software Defined Aggregation
Networking	Centralized	LAN	LAN SAN	LAN SAN HPC	Converged (LAN/SAN/HPC)
Storage	Centralized	Dedicated SCSI Storage	SAN NAS	SAN NAS DAS in Cloud	Software Defined Storage

Table 1: Trends in IT Architectures for Compute, Networking and Storage

In many respects, these changes can be summarized as an oscillation between integrated (or converged) and dis-integrated (or component-based) architecture across hardware, software, and geographical dimensions. Whether a given dimension was converged or component-based was often dictated by the capabilities of the hardware resources underlying the computing architecture, and the cost of the architecture relative to the value of the potential applications. For instance, the "mainframe era" saw highly mission-critical applications running on very expensive hardware (mainframe computers), while the 1990s saw a huge growth in applications that were available to enterprises (and users) due to the proliferation of low-cost x86 servers connected by dedicated local area network (LAN) hardware components (early switches and routers).

As the most recent incarnation of computing architectures, Hyper-Converged Infrastructure (HCI) is one of the hottest topics in data center architecture today. In the simplest terms, HCI combines computing, storage, and networking components into a single physical instance. These devices, often called "datacenters in a box" or "HCI appliances", are significantly easier and cheaper to deploy than standard data centers based on separate servers, storage arrays, and networking switches. HCI evolved into its current form from the blade servers of the last decade, and from the purpose-built data centers developed by hyperscale companies such as Google, Amazon Web Services (AWS), and Microsoft Azure. Today, nearly all of the leading information technology (IT) companies develop and sell HCI appliances, with acknowledged leaders including Nutanix, Simplivity, Dell/EMC, HPE, Cisco, Oracle, VMWare, and NetApp.

HCI represents an architecture where nearly all of the hardware components (compute, networking, and storage) are integrated into unified building blocks, which are referred to as HCCAs. In most cases, these HCCAs run as virtualized environments which utilize ***hypervisors*** to provide security and execution isolation between the various applications running on the HCCA compute resources. The HCI concept is largely an evolutionary combination of the ***bladed server concept*** and the scale-out architectures of hyperscale companies (Google, Amazon Web Services, and Microsoft Azure). Figure 1 below illustrates the topographical difference between the HCI approach (shown on the right), and the architecture characteristic of virtualized computing environments today.

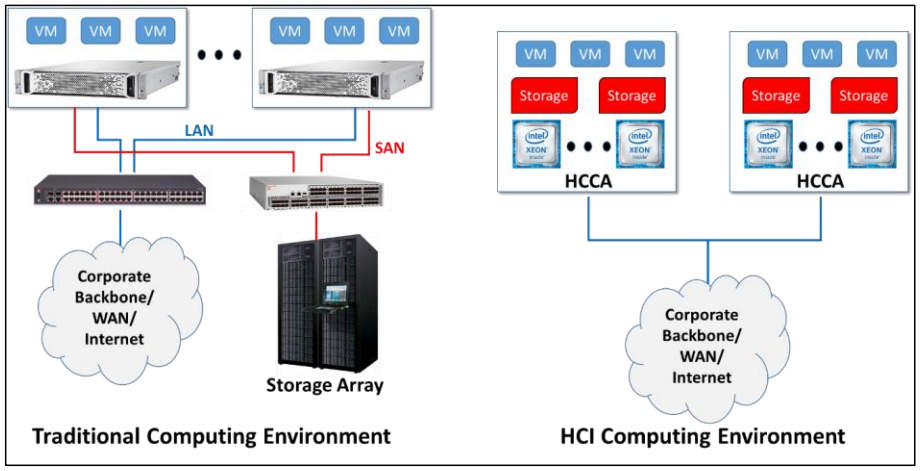

Figure 1: Comparison of HCI and Traditional Computing Architecture

HCI originally developed around two specific use cases:

- **Remote Office/Branch Office (ROBO)**, where large enterprises needed the ability to deploy computing and storage resources that were literally "plug and play", and could be remotely monitored, managed, and configured; and
- **"Stand-Up/Scale-Out" Data Centers**, where global enterprises such as hyperscale companies, social media companies (Facebook), or network/telecom service providers (AT&T, Verizon, China Mobile, or Google Fiber) needed to be able to stand up large data centers quickly with low-tech staff, and operate these with few or no local staff ("lights out" operation).

Hyper-Converged Infrastructure Adoption

Arguments about when HCI became a distinct computer architecture abound, but most industry analysts place this between the advent of Cisco's UCS server architecture, and the emergence of the Nutanix HCI platform. In any case, by the start of this decade HCI emerged as a significant force in data center computing. As with all technologies, the time line between invention and adoption is modulated by application development, CAPEX retirement lifetimes, and market segmentation, and this is no different for HCI. That said, HCI sales are expected to reach $5B annually in 2019 according to Gartner.

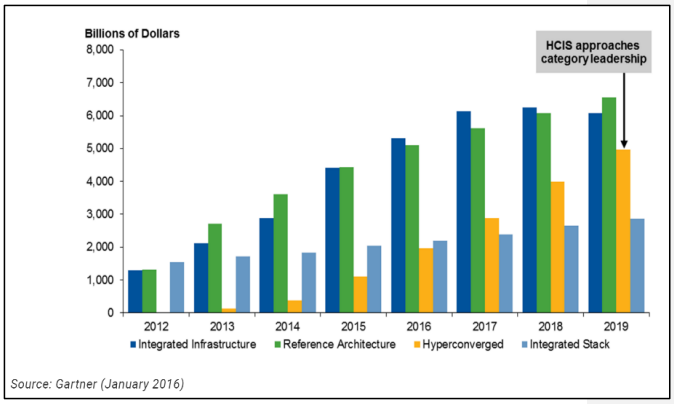

Source: Gartner (January 2016)

One of the reasons that HCI is growing rapidly is that its value proposition extends beyond the original use cases. The main reasons why HCI is being adopted rapidly outside of the ROBO and stand-up/scale-out datacenter use cases include:

1. **Business Focused Efficiency**: In most organizations, IT departments are cost centers that exist to help the organization run their businesses, and as such are under constant pressure to reduce costs. HCI makes IT more efficient and effective by lowering the operational expenses (OpEx) of deploying, provisioning, and managing compute resources.
2. **Quicker time to Production:** HCI solutions are easy to deploy, and require less specialized technology skills from the IT teams that are deploying them. This allows IT teams to stand up new resource faster, and allows IT for focus on business solutions rather than specific technologies.
3. **Scalability** – One of the main drivers for HCI is the ability to scale resources on-premises just like a cloud. HCI gives you the modularity to scale-out and scale-up concurrently.

Adoption of every new generation of technology occurs in waves, first starting with specific use cases, and spreading to more general application. The timeline and steps for HCI have been studied by Gartner, and are shown in Figure 2 below as HCI evolved from the blade server architecture. HCI provided three core

capabilities that blade servers lacked (and which prevented blade servers from achieving significant market penetration. Those three capabilities include:

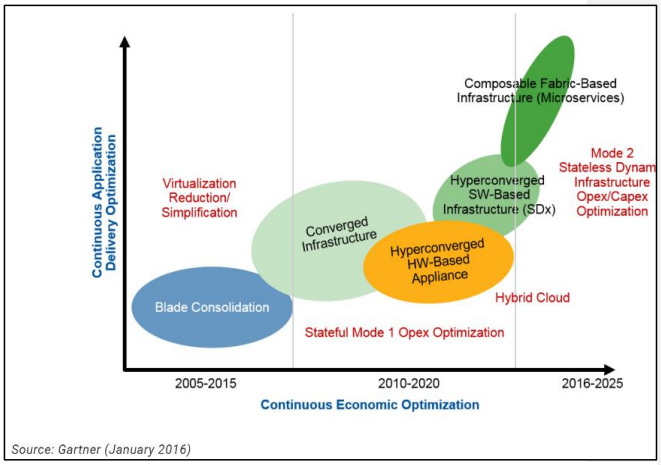

Figure 2: HCI Adoption "Waves" (Gartner, 2016)

1. **Shared Compute, Network and Storage:** While many blade servers contained internal switches to connect the blades together, they were designed from the start to connect to enterprise LANs and SANs. More importantly, few contained any concept of shared storage. In contract, HCI was designed from day one to have integrated networking and storage resources, and are capable of operating as a standalone system.
2. **Inherent Scale-out Capabilities:** In essence, blade servers are simply rackmount server designs with better packaging and (sometimes) lower power consumption. Chassis of blade servers were "islands" in the same way that rackmount servers were; they were not designed to share resources across multiple blade server chassis. In contrast, HCI is inherently designed with the scale-out model in mind. Each node is aware of the other nodes, and they all have the ability to share and combine resources.
3. **Software Defined Data Center:** By combining dynamically reconfigurable resources, HCI inherently embodies the capabilities required to implement the Software Defined Data Center (SDDC) concept. Since each element of the HCI architecture (processing, networking, and storage) can be virtualized and dynamically allocated, they can be reconfigured to meet changing demands, such as those encountered in private cloud deployments. The combination of the HCI architecture and SDDC makes this architecture particularly suited to today's compute models.

Today, nearly every major IT hardware vendor offers an HCI solution. The popularity of HCI is illustrated by Gartner's recent Magic Quadrant for Integrated Systems, which predicted that approximately 30% of global storage array capacity installed in enterprise data centers will be deployed on SDS or HCI systems by 2019 (a six-fold increase from today's 5% rate). Additionally, Gartner predicts that 20% of mission-critical applications will transition to HCI systems by 2020 (Gartner, October 2016).

Top Application for HCI Deployments

IT Brand Pulse® recently conducted a survey of IT professionals that have deployed or are considering deploying HCI in the next 12 months. The top applications and reasons that the participants identified for adopting HCI are shown in Figure 3:

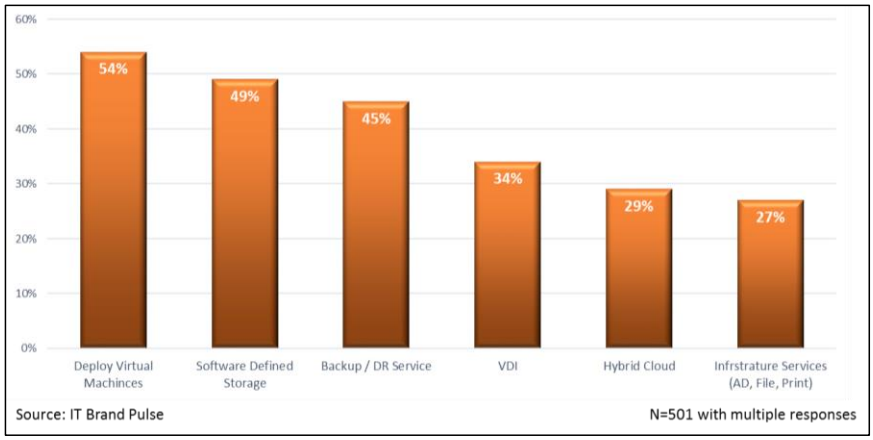

Figure 3: Factors Driving HCI Adoption (IT Brand Pulse, 2016)

Clearly, virtualized environments, whether for servers, VDI, or private/hybrid clouds, are perhaps the largest driver of HCI adoption in large enterprises, an insight supported by the fact that software-defined storage (also highly correlated with cloud implementations) was the second-most sited factor. These trends reflect the fact that HCI is an infrastructure implementation choice more than it is an architecture targeted at specific applications. Large scale-out application deployments such as big data (Hadoop, Spark, etc.) are particularly well suited to HCI architectures. Private/hybrid cloud deployments are also strong use cases for HCI, as the ability to rapidly add, repurpose, and (eventually) retire hardware components with minimal impact to the overall system is critical to their resilience in the face of changing workloads.

Trends and Issues with the Adoption of HCI Outside of Large Enterprises

For medium-sized enterprises who have limited IT staff and cannot move to cloud-based architectures due to business constraints/needs, the advantages provided by HCA are incredibly compelling. The elimination of the need for highly-technical storage and network engineers, coupled with the ease of deploying additional HCI components, make on-premises computing attractive again for medium enterprise customers. However, most mainstream HCI solutions have adoption challenges for medium enterprises. These include:

- The ability to select and tailor a solution that is not too large (and too expensive) for the needs of medium-sized enterprises. Many of the leading HCI solutions are rack-level products, with expensive software licenses and maintenance agreements. These often offset much of the CapEx and OpEx gains realized by utilizing HCI solutions over conventional IT architectures.
- Finding a way to "get the attention" of large HCI vendors, or conversely finding channel resellers that have the expertise to help medium-sized enterprises choose the right solutions. This is especially true of HCI solutions that are implemented as "reference architectures", where the reseller (or even the customer) must integrate hardware and software, a job that is typically far beyond the capabilities of medium enterprises.

Overcoming these factors (size, cost, implementation partner) is critical if medium enterprises are to take advantage of the benefits of HCI.

dataON

DataON, Intel, and Microsoft – Choosing the Right Partners to Deliver High-Performance HCI for Window Environments

HCI hardware platforms by themselves are not a solution; the platform needs to incorporate storage components, a hypervisor, and a management framework. Moreover, the companies providing these building blocks must work together to deliver not only the hardware, but also deployment help, training, and support to the customer. If done correctly, the result is a seamless integration that works "out of the box" and meets the customer's business needs.

To meet these demands, DataON, Intel and Microsoft partnered to develop DataON's ClusterBlock Architecture™. DataON further extended the HCCA concept with its Management Utility Software Tool (MUST) framework, which provides for seamless management of the nodes within an HCA cluster. These technologies form the nucleus of DataON's Hyper-Converged Cluster Appliance (HCCA) product line. The HCCA product line utilizes the following key components (shown in the figure below):

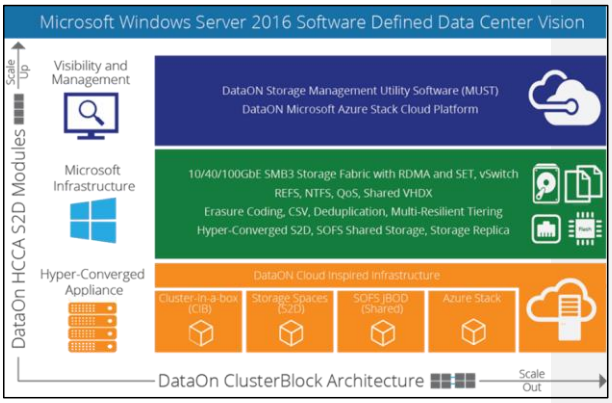

- Microsoft's Windows Server 2016 operating system, which includes Microsoft Windows Hyper-V hypervisor for server virtualization, Microsoft Storage Spaces for software-defined storage, and the Microsoft Azure Stack to support hybrid cloud storage configurations.
- Intel E5 CPUs to provide the optimum balance of compute power and low power consumption necessary for high-density HCI platforms;
- Intel Non-Volatile Memory Express® (NVMe™) 3D NAND solid state drives (SSDs) for performance intensive application needs, and Intel capacity-optimized NAND SSDs for data protection and cost-sensitive services such as file/print server, Active Directory, DNS and other similar operational services.
- DataON MUST framework, which provides a seamless tool to manage all of a customer's hyper-converged cluster appliances.

Commented [SW1]: Fix DataOn lower case

Together, the combination of DataON, Microsoft, and Intel provide a compelling HCI platform for medium-sized enterprises. Some of the key advantages of this platform for customers include:

- **High Performance Storage Where/When it is Needed:** Windows Server 2016 Storage Spaces Direct allows you to create highly available, highly scalable software-defined storage at a fraction of the cost of traditional SAN or NAS arrays. Hyper-Converged architecture radically simplifies procurement and deployment, easily scaling up and while features like caching, storage tiers, and erasure coding, together with the latest hardware innovation like RDMA networking and NVMe drives, deliver unrivaled efficiency and performance. All of this is available with any Windows Server 2016 Datacenter license. Furthermore, the ability to utilize Microsoft's multi-tier storage

architecture to combine high-performance NVMe SSDs, capacity-optimized SSDs allows customers to "dial in" the performance they need for a given workload, while at the same time optimizing the overall HCI platform cost. This results in performance that is similar to that provided by all-flash arrays, but at a fraction of the cost.

- **A High-Performance Hypervisor that is Optimized for the Windows Environment:** By utilizing Microsoft Windows Server Hyper-V, the DataON HCCA platform provide a "ready to roll" hypervisor that is fully integrated with Microsoft Windows applications. While other mature hypervisors that support the Microsoft environment exist (for instance, VMWare, KVM, and Linux distros), these hypervisors either have significant licensing/support costs (VMWare, Linux distros), or force the user to perform their own deployment and integration (KVM). Only Hyper-V provides a free hypervisor that is fully integrated and tested with Windows Server 2016.
- **A No-Cost HCI Management Framework:** DataON's MUST is the only free system management tool for the Windows Server Environment. By going with DataON, customers avoid the need to buy expensive 3rd party tools to manage their HCI environment.

These advantages combine to make the DataON HCCA product line a very compelling offering for medium-sized enterprises that want to take advantage of the reduced footprint, lower operational costs, and seamless scale-out capabilities provide by HCI platforms, but without the expensive license costs and/or having to buy solutions that are too large for their needs.

EVGA Case Study

EVGA is one of the top NVIDIA authorized partners in channel sales throughout North America. The company has been recognized by dozens of magazines and product reviewers for their superior products (http://www.evga.com/about/awards). EVGA differentiates their solutions by offering the best technology, 24/7 tech support, a 90-day step-up program that gamers and video connoisseurs value. With headquarters in Brea, CA, EVGA's global coverage includes sites in Munich (EMEA), Miami (for LATAM), and Hong Kong (APAC). Online gaming is a 24/7 business and EVGA needs an IT infrastructure that can scale with their growing installed base, during high demand periods when new PC video game are released, or during major holiday rushes.

EVGA's Goals for their Hyper-Converged Infrastructure Deployment

As EVGA continued to grow their business, they needed to update their IT infrastructure so that it could be resilient, scale on demand, and reduce operational costs. EVGA's last IT overhaul program was a critical component in helping them become one of the top NVIDIA suppliers, but that infrastructure is now several years old, and contained several high-maintenance-cost subsystems, including a storage system built around an expensive and proprietary NetApp filer and associated storage area network (SAN). In addition, the configuration contained network switches and several rackmount compute servers, and the overall size of the configuration was 26U (65% of a standard 40U rack). VMWare ESXi was being used as the hypervisor for this configuration, but its expensive software licenses and ongoing maintenance fees had a significant OpEx impact on EVGA's overall IT budget.

As leaders in a highly competitive market segment, EVGA knew that they would have to upgrade their IT infrastructure to keep pace with their business growth, to provide the extremely high level of customer support that EVGA is known for, and to keep their IT costs in line. EVGA's goal for this upgrade to move to

a solution utilizing Microsoft Windows server software, and which utilized the Hyper-Converged Infrastructure (HCI) compute model to enable future growth without forklift upgrades. The goals of this upgrade program included the following:

- Improve the likelihood that EVGA could meet the performance demand of their users and IT enterprise applications as customer and business demands continued to grow.
- Leverage their existing investments in the Windows Server Platform to reduce deployment times and IT staff training costs.
- Reduce VMWare software licensing cost and maintenance cost with proprietary SAN.
- Migrate to SQL 2016, Exchange, and DNS servers hosted on an HCI platform to simplify future growth.
- Simplify the overall IT Infrastructure, and increase the ease of deployment and migration of VMs.
- Provide "bullet-proof" data level protection and future storage replication capabilities.

Choosing an Implementation: Windows Server 2016 and Intel-Based HCI

Why did EVGA choose to move to Windows Server 2016 and DataON's CBA architecture to improve its IT operations? The answer is simple: they are ready to build a scalable, cost effective, cloud ready IT infrastructure that will scale as their compute and storage needs grow in the future. The DataON S2D-3110 HCI implementation will enable them to deploy more resources quickly and at a low cost, scale their resources to a greater extent, buying fewer software licenses, accelerate VM deployments and migrations, and reduce staff workloads. The result is greater IT capabilities, with significantly lower CapEx and OpEx. Some measures of these advantages include the following:

1. **Consolidation of Infrastructure**: We have typically seen about a 3:1 reduction in the number of infrastructure chassis' required to support HCI versus traditional infrastructure. This consolidation is across servers, switches, and storage arrays. This will save on the absolute number of chassis' deployed, reduce complexity, reduce cabling plants, and reduce maintenance costs.
2. **Infrastructure Licensing and Support**: Every extra chassis comes with annual support costs, software management licenses, and annual software support costs. The reduction of these extra should scale at the same 3:1 ratio as you reduce the infrastructure.
3. **Scale-Out Flexibility and Capabilities**: The core value proposition for HCI is the ability to scale-out IT resources and the flexibility to reallocate those resources on demand. The ability to scale in a near-linear fashion will enable you to better plan CAPEX, improves SLAs, and reduce down time.
4. **Windows Hyper-V vs VMWare ESXi**: By moving to Windows Server 2016, EVGA saves on the cost of VMware ESXi, and reduces IT staff management and team training demands. These savings can be significant, as shown by the TCO chart below.
5. **External SAN vs Storage Spaces Direct and Replica** – Utilizing Storage Spaces Direct as part and HCI implementation not only saves on hardware costs (HBA/CNAs, switches, and array hardaware), but also on the cost of backup software. Using Microsoft Replica, the cost of backup software (and the management tools for it) are eliminated, as Replica is included in Windows Server 2016 and is managed through the DataON MUST Management Framework.

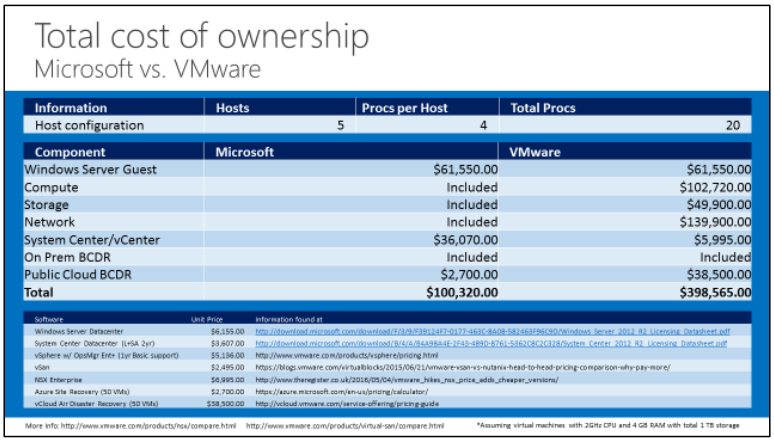

6. **Accelerating Applications with Intel NVMe 3D NAND SSDs with 2.4M IOPS** – One of the biggest innovations for accelerating applications is the advent of NVMe (Non-Volatile Memory Express) based SSD storage that can improve read performance by 3X and write performance by 1.5X vs 12Gb SAS SSD[1].

DataON's S2D-3110 HCCA – Delivering Scalability and 2.4M IOPs

Performance was one of EVGA's biggest concerns in their new IT deployment. To determine if the DataON S2D-3110 HCCA cluster could provide the performance needed, EVGA tested the production configuration using VMfleet. The figure to the right represents the production configuration in the EVGA data center. EVGA's S2D-3110 HCCA cluster achieved more than 2.4 million IOPS when running Microsoft Server 2016 Storage Spaces Direct, as shown below. This provided EVGA with the confidence that the DataON solution would meet its performance needs, today and in the future.

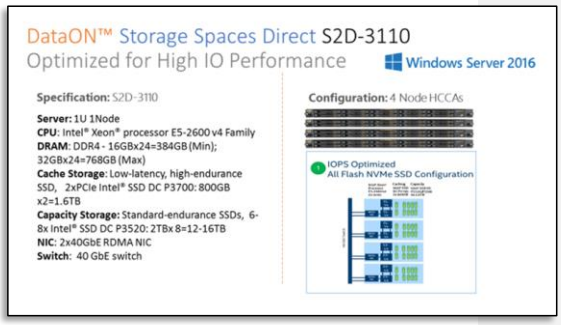

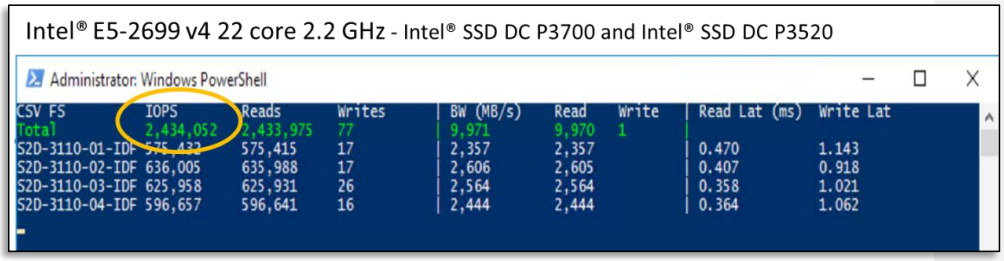

[1] http://blogs.cisco.com/datacenter/nvme-for-absolute-beginners

dataON

Results: DataON, Intel, Microsoft Exceed EVGA's Goals for Their New HCI Infrastructure

EVGA is in a highly competitive business, and its customers are always looking for ultimate gaming experience. These customers expect product ordering and technical support experiences that matches the performance of EVGA's products. To meet consumer expectations, EVGA updated their data center configuration to DataON HCCAs, Intel NVMe SSDs, and Windows Server 2016. This is one more example of EVGA's strategic leadership, their dedication to stay ahead in the market, and their commitment to improve the experience for their tech-savvy customers.

EVGA's deployment of DataON's S2D-3110s, with Intel processors and NVMe SSDs and the Microsoft Windows Server 2016 software infrastructure, more than met EVGA's expectations for this deployment. Here are how the results compared to EVGA's original goals for this deployment:

Requirement	Results
Improve the likelihood that EVGA could meet performance demands of their users/applications as customer and business demands grow.	The performance of over 2.4M IOPS by the DataON S2D-3110 ensured EVGA that this new solution would be able to successfully host their current and future applications without issues.
Leverage their existing investments in the Windows Server Platform to reduce deployment times and IT staff training costs.	Since EVGA's IT staff already knew the Microsoft Windows Server platform, deploying and maintaining the DataON S2D-3110 configuration will easy for them. Moreover, DataON's MUST management framework further simplifies overall system and component management, allowing EVGA's IT staff to put their energy into growing their business.
Reduce VMWare software licensing cost and maintenance cost with proprietary SAN.	Since Hyper-V is included at no cost with Windows Server 2016, the software licensing and maintenance costs associated with VMWare ESXi are no longer an issue for EVGA. Moreover, the use of DataON's HCI-based S2D-3110 eliminated the need for (and the cost of) the proprietary SAN.
Migrate to SQL 2016, Exchange, and DNS servers hosted on an HCI platform to simplify future growth.	The DataON S2D-3110's storage and deployment architecture provided the ideal platform to run Microsoft SQL Server 2016, Microsoft Exchange, and EVGA's DNS servers. The automated out-of-box workflows of the S2D-3110 also accelerate the deployment times for Windows Server 2016, Storage Spaces Direct, and Storage Replica environments.
Simplify the overall IT Infrastructure, and increase the ease of deployment and migration of VMs.	With an ability to support up to 40+ Hyper-V virtual machines (VMs) per node, the DataON S2D-3110 HCCA significantly increased EVGA's VM deployment capabilities. The configuration of the S2D-3110 also simplifies VM migration by utilizing SMB 3.0 over a high-speed RDMA-based Ethernet network, increasing CPU efficiency and eliminating multiple-copy latencies.
Provide "bullet-proof" data level protection and future storage replication capabilities.	The use of Microsoft Storage Replica on the DataON S2D-3110 ensures that EVGA's critical data is always available and always protected. The high speed SMB Direct "storage bus" within the S2D-3110 cluster also ensures that if a node restore is required, the node will be back on line quickly.

According to Joe Darwin, EVGA's chief branding officer, "At EVGA we became the #1 NVIDIA provider in the market because we are gamers who build gaming systems because we want the same experience, expertise and support our customers want. We don't just sell gaming, we live it." EVGA knows that the right systems create the best experiences for their customers and they know the same rule applies to their business and DataON, Intel, and Microsoft are thrilled to be part of the EVGA success story.

"Building on our mission to use the latest technology, EVGA continues to invest its IT Data infrastructure using faster hardware with the latest Microsoft software. These new Hyper-Converged systems offer additional features and functions and most importantly increase reliability at an affordable cost," said Joe Darwin, Chief Branding Officer for EVGA Corporation. ***"Working with companies like Microsoft, Intel, and DataON gives our IT team a good foundation to provide a solid infrastructure for future growth planning."***

Best Practices in HCI Programs: How EVGA and DataON Deployed the New HCI Solution

The deployment of EVGA's HCI solution was a four-step process, which included the deployment of the DataON S2D-3110 and Windows Server 2016, setting up the server virtualization environment using Windows Hyper-V, configuring Windows Storage Spaces Direct, and configuring Windows Storage Replica. Each of these steps are described below.

DataON S2D-3110 and Windows Server 2016 Deployment

DataON's HCCA provide a fully optimized platform for deploying Windows Server 2016. Roll out of Windows Server 2016 on the DataON S2D-3110 is painless and fully automated, simplifying the overall system deployment.

Hyper-V Virtualization Configuration

Server virtualization divides up the resources of a physical server to create multiple isolated virtual machines or VMs. DataON's HCCA was designed specifically for running Hyper-V on Hyper-Converged infrastructure. The HCCA provides high-availability via the HCI design, no additional configuration and no additional software purchases. When a node fails, VMs are automatically restarted on other nodes in the cluster via the native Live Migration capability and it does this without disrupting application activity. Windows Server 2016 Hyper-V is included as part of the Windows enterprise licensing costs.

Storage Spaces Direct Deployment

DataON's S2D-3110 HCCA provides fully automated deployment of the Storage Spaces Direct high-availability scale-out clustered storage systems. Storage Spaces Direct is a distributed file system option in Windows Server 2016, along with the traditional clustered Storage Space Shared file system from Windows Server 2012. Storage Spaces Direct adds support for new classes of storage devices, such as NVDIMM, SATA and NVMe disk devices, that were previously not possible with clustered Storage Spaces with shared disks.

Data Protection via Replica for Windows Server 2016

The complexity and cost of data protection is a particular challenge for many IT teams who often make due with a variety of backup solutions with offsite tape storage for disaster recovery. Windows Server 2016 adds a new data protection solution to the tool belt of the IT teams called Microsoft Storage Replica. This tool provides a complete way to backup and replicate data locally, to Azure cloud, or to other DR sites. Storage Replica enables storage-agnostic, block-level, synchronous replication between clusters or servers for disaster preparedness and recovery. Storage replica also enables the "stretching" of a failover

cluster across sites for high availability and to ensure zero data loss at the file system level. The 4U S2D-3202 HCCA and its JBOD expansion chassis provides an ideal platform to leverage the latest Windows Server 2016 data protection technologies.

Building Blocks to Make Your HCI Deployment Successful

As is illustrated by the EVGA case study, Microsoft Windows 2016 in and of itself does not guarantee a successful HCI deployment. There are many platforms and components available to implement Windows 2016 HCI deployments, many from leading IT server suppliers and vendors. DataON HCI systems and Intel processors and SSDs can help ensure that your Windows HCI implementation meets your expectations around performance, footprint/power savings, and ease of management and maintenance.

DataON: "The Certified Microsoft Platform"

At DataON, we do one thing: design, build, and deploy Microsoft-certified Hyper-Converged Cluster Appliances (HCCAs) and storage systems for Windows compute environments. Our systems are architected and built with a singular purpose: to optimize their performance for Microsoft's Windows Server 2016 software-defined data center (SDDC) and Microsoft Azure cloud-inspired deployments. We have built multiple generations of platforms for each evolution in Microsoft's enterprise product line, including Windows Server, Azure, Hyper-V, Storage Spaces Direct, and Microsoft Management Systems. In the design of each generation of our products, we have collaborated with Intel and Microsoft to ensure that our solutions are able to achieve the following:

- Scale out to multiple nodes and systems, without affecting overall system performance.
- Result in lower CapEx and OpEx costs, both in hardware and licensing costs, and in operational and maintenance cost.
- Simplify system deployment, enabling customers to utilize DataON platforms in a variety of environments and use cases.

Our comprehensive certifications by Microsoft continue to show why DataON systems has earned the label "The Certified Microsoft Platform."

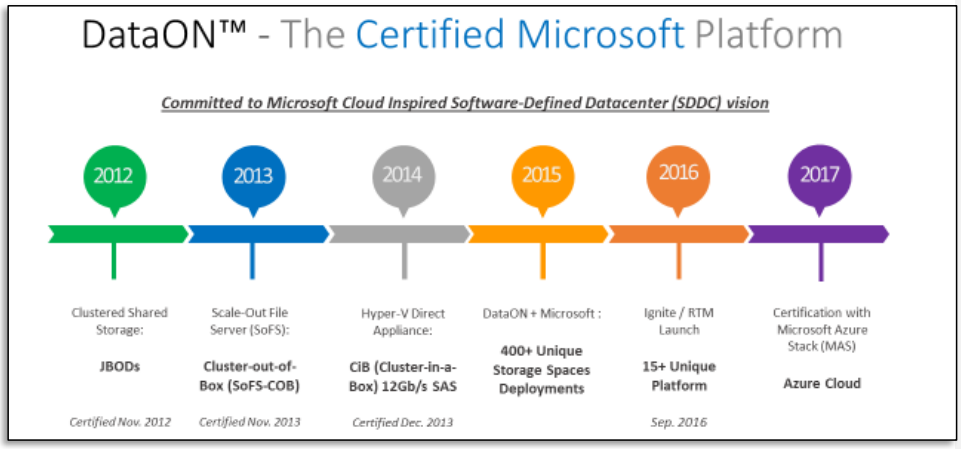

DataON ClusterBlock Architecture

DataON's ClusterBlock Architecture (CBA) was created in partnership with Intel and Microsoft to provide unprecedented scale-out options and an easy-to-use management capability for the deployment of Microsoft's SDDC platform. The DataON HCCA, which is based on the CBA design, optimizes deployment of Windows Server 21016, Hyper-V, Storage Spaces Direct, and Azure cloud solutions, and enables DataON to configure HCCAs that can be performance, cost or capacity focused. This enables IT teams to select HCCAs that are purpose built for specific application workloads, such as:

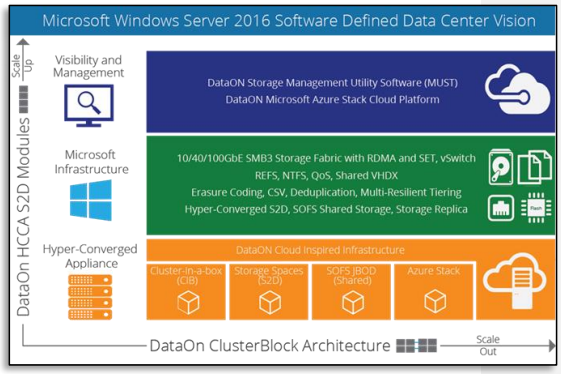

- Highly-scalable Hyper-V server virtualization deployments.
- Microsoft Azure-based hybrid cloud solutions.
- High-reliability file and storage servers utilizing Microsoft Storage Spaces Direct and/or Replica.
- Scalable database systems based on Microsoft SQL, Hekaton™.
- High-availability Microsoft Exchange, SharePoint™, and Dynamics ERP clusters.

DataON platforms can be highly optimized for these core IT use cases, while leverage the cost and simplicity benefits of HCI.

DataON S2D-3000 Family of Hyper-Converged Cluster Appliances

The DataON S2D-3000 family of Hyper-Converged Cluster Appliances for Windows Server 21016 and Cluster-in-Box (CiB) Hyper-Converged Cluster Appliances provide a diverse set of HCI and storage systems to support Hyper-V, S2D, Azure and Data Protection.

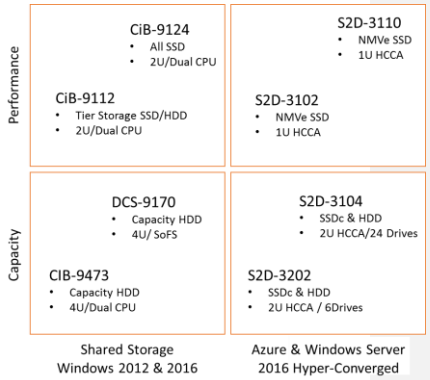

The first two products to ship in the S2D-3000 family are the S2D-3110, a high-performance platform designed to maximize input/output (I/O) transactions per second (IOPS) for scaling Hyper-V virtualization, SQL, Exchange, SharePoint, Dynamics ERP, Microsoft Hekaton™, SharePoint™ and the S2D-3202 a capacity optimized HCCA ideal for Data Protection and other core Windows Services (AD, File, Print, DNS). The S2D was first demonstrated at the IDF 2016 tradeshow in August 2016 and the S2D-3202 was debuted at the Microsoft Ignite at the Windows Server 2016 announcement.

Model Number	S2D-3110	S2D-3102	S2D-3202	S2D-3204
Form Factor Server Nodes	1U Single (1)	1U Single (1)	2U Four (4)	2U Single (1)
Drive Configuration	10x2.5" NVMe SSD drives per Server Node	4x2.5" NVMe + 8x2.5" SSD drives per Server Node	2x2.5" NVMe + 4x2.5" SSD drives per server blade times 4-nodes	4x2.5" NVMe + 20x2.5" SSD drives per Server Node
Network Adapter	10/40GbE RDMA NIC Adapter	10/40GbE RDMA NIC Adapter	10/40GbE RDMA NIC Adapter	10/40GbE RDMA NIC Adapter
Storage Network Adapter	12Gb SAS HBA	12Gb SAS HBA	12Gb SAS HBA	12Gb SAS HBA
Storage Capacity (x4 Nodes): Performance	8xNVMe SSD (400Gb-2TB per SSD)	12xNVMe SSD (400Gb-2TB per SSD)	8xNVMe SSD (400Gb-2TB per SSD)	16xNVMe SSD (400Gb-2TB per SSD)
Storage Capacity (x4 Nodes): Capacity	32xNVMe SSD (400Gb-2TB per SSD)	32x2.5" SSD (400Gb-2TB per SSD)	24x2.5" SSD (400Gb-2TB per SSD)	80x2.5" SSD (400Gb-2TB per SSD)
SMB3 Storage Fabric	40/56GbE DCB Switch	40/56GbE DCB Switch	40/56GbE DCB Switch	40/56GbE DCB Switch
Storage Management	DataON MUST	DataON MUST	DataON MUST	DataON MUST
Windows Server 2016 and Azure Stack	Datacenter Edition Automation / Backup	Datacenter Edition Automation / Backup	Datacenter Edition Automation / Backup	Datacenter Edition Automation / Backup
Hyper-converged Platform	IOPs Optimized	Balanced IOPs / Capacity Optimized	Balanced IOPs / Density Optimized	Throughput / Capacity Optimized

DataON's S2D-3000 family of products is able to achieve such a wide range of capacity, scalability, and performance optimization because of its use of the Cluster Block Architecture, which provides the product line with three main capabilities:

- **Scale-Out Services:** These services provide the ability to support Windows Server 2016's SDDC, S2D cloud-inspired infrastructure, and Azure Stack. As IT managers deploy each HCCA they are adding incremental compute, networking, and storage resources that provide near-linear scalability. This includes the **Scale-out File System**, which enables HCCAs to scale across racks and into the cloud.
- **Microsoft-Centric Infrastructure:** The CBA is designed with the single purpose to enable, simplify, scale, and optimize deployment of Microsoft environments. This means that no compromises or tradeoffs are made in order to be able to support other environments. The DataON infrastructure which supports deployments in Microsoft Windows environments includes: SMB Direct Networking**,** including RDMA (Remote Direct Memory Access) capable NICs with speeds up to 100Gb/s to support Microsoft's Scale out File System; Intel NVMe Based SSDs, which eliminate the latencies experienced with SAS and SATA SSD; 12Gb SAS connectivity to support scale-up JBODs for backup, disaster recovery and business continuance; and backwards compatibility and migration capabilities for existing and new Microsoft file systems, including NTVS and ReFS.

DataON MUST Management Framework

DataON's MUST™ (Management Utilization Services Tools) framework delivers predictive insight, monitoring, planning and management capabilities across an entire DataON deployment. MUST extends the management services currently provided Windows Server 2012 or 2016 at the following levels:

- **HCCA Platform:** MUST allows IT managers to gain visibility and actionable insight into the health and usage statistics for the HCCA platform's CPU, memory and networking, internal connectivity and storage devices. The MUST framework provides predictive analysis based on historical data to understand how peak usage is affecting service level agreements (SLAs), as shown below.

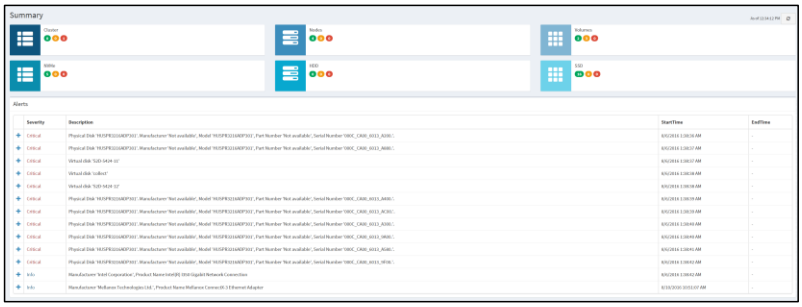

- **HCCA Storage Cluster:** The MUST framework provides detailed understanding of the S2D cluster volume including storage pool usage/allocation, resiliency storage commitments and total pool capacity, CPU utilization, memory utilizations, I/O latency, IOPs and bandwidth performance.

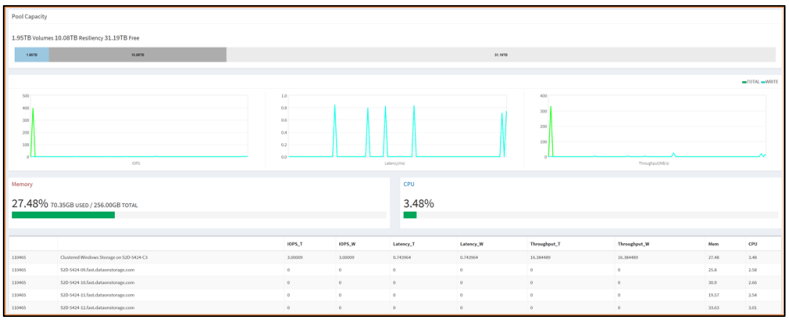

- **HCCA S2D Volumes** – MUST enables you to drill down into the specific virtual volumes to see volume-specific I/O latency, IOPs, and bandwidth performance.

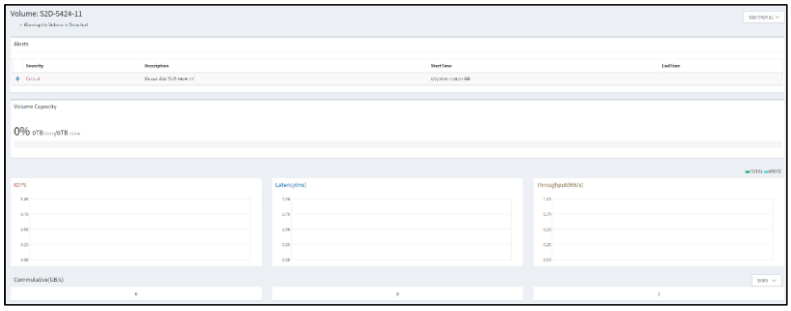

Built On Intel® Technologies

Intel® Xeon® Processor E5-2600 v4 Family

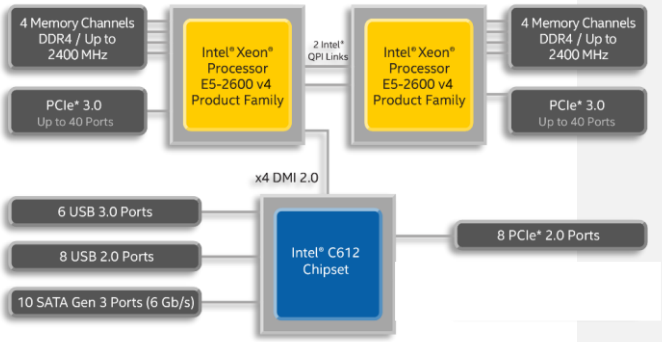

The Intel® Xeon® processor E5-2600 v4 product family, manufactured on the latest 14 nm process technology, delivers the high performance and increased memory bandwidth for the next generation of hyper-converged infrastructure (HCI) platforms. This new processor family provides the capabilities to meet the needs of a wide range of compute-intensive applications including video analytics, high-end medical imaging, and security applications. The enhanced security capabilities of the E5-2600 v4 family enable improved data encryption and deliver higher protection from attacks on data moving from devices at the edge, over the network to the cloud. Some of the key features of the E5-2400 v4 family include:

- **High performance for the broadest range of applications**: Core architectural advancements, along with improved memory DDR4 speed (up to 2400 MHz) and high I/O bandwidth, provide higher performance for the broadest range of applications.
- **Better manage shared platform resources**: The new Intel® Resource Director Technology (Intel® RDT) provides a mechanism to better manage shared platform resources such as cache and memory, provide deterministic performance and a guaranteed resource for high-priority workloads. Intel® RDT also allows service providers to control and manage shared resource utilization to match service level agreements.
- **Boost performance for multi-threaded workloads**: Intel® Transactional Synchronization Extensions (Intel® TSX) exposes hidden parallelism to help boost performance for multi-threaded workloads that are currently slowed by memory locking.
- **Virtualization technology enhancements**: With the rapid adoption of virtualization technologies across multiple industries, and implementation of multiple virtual machines (VM) on a single server platform, features like Posted Interrupts help reduce the average latency for servicing I/O interrupts by mapping multiple interrupts from a single VM into a single APIC/vector.
- **Workload acceleration**: Intel® QuickAssist Adapter 8950 offers customers a variety of capabilities such as hardware acceleration and offload to assist with the performance demands of compute-intensive security and compression operations, thereby reserving processor cycles for application and control processing.

- **Abundant integrated I/O Options**: Abundant integrated I/O's available on platform including up to 80 PCIe* Gen3 lanes, up to 10 SATA Gen3 ports, and up to 14 USB ports to fulfill the design needs.

Manufactured using Intel's advanced 14nm processes, the E5-2600 v4 family provides key features for HCI applications such as Intel® Advanced Vector Extensions 2 (Intel® AVX2), Intel® Resource Director Technology (Intel® RDT), Cache Monitoring Technology (CMT), Cache Allocation Technology (CAT), Memory Bandwidth Monitoring (MBM), Intel® Transactional Synchronization Extensions (Intel® TSX), Intel® Virtualization Technology (Intel® VT), Intel® Advanced Encryption Standard New Instructions (Intel® AES-NI), and Intel® Secure Key. These capabilities optimize the performance of Windows Server 2016, while reducing overall power consumption.

Intel® Solid State Drive Data Center Family

Intel's 1000+ researchers lead the industry in developing new memory technology breakthroughs, including creating the first commercial solid state drive, and now, building the world's fastest SSD. Discover the story behind Intel's two latest disruptive technologies, 3D NAND and 3D XPoint™ memory. Intel has a wide variety of solid date drives (SSDs) and PCIe storage devices for the enterprise data center. Intel SSDs have several significant innovations that make them particularly suited to use for critical storage needs in the enterprise data center. These include:

- **Lasting Integrity**: Intel SSD are built with proven end-to-end data protection capabilities and the highest level of drive reliability. Intel SSDs protect data through millions of interaction (read or write) cycles, stand up to more than 2 million PLI cycles and self-tests without losing data (even during power losses), and are 100 times more reliable at preventing Silent Data Corruption (SDC), to ensure that your enterprise's files and information are there when you need them.
- **Reliably Effective Performance**: When you invest in an SSD, you want to ensure that it will perform effectively for years to come, no matter how much you stress it. Intel SSDs are tested more rigorously than the competition, with thousands of drives in thousands of different configurations, over one million power cycles. As a result, Intel SSDs provide highly consistent performance in the enterprise data center.
- **Platform Confidence**: Intel SSDs interoperate seamlessly with other Intel components such as processors, chipsets, network interfaces, firmware, and drivers, eliminating potential interoperability issues.

DataON's HCCAs use the Intel® SSD DC P3700, Intel® SSD DC P3520, and Intel® DC P3500. These Intel NVMe/PCIe-based SSDs offer incredible performance and enhanced capabilities, with advanced capacity and performance coming in future Intel® Optane™ SSDs based on 3D XPoint™ Technology. These drives ensure extreme performance for demanding applications, lower cost of ownership, and proven reliability for your business's critical enterprise applications and data.

Conclusion

DataON, Intel, and Microsoft are leaders in the HCI market, and are committed to provide the highest performance and easily deployable solutions for Windows Server 2016 and Microsoft's SDDC platform. HCI will undoubtedly be a major force in private enterprise, hybrid cloud, and public hosting environments. As evidenced by the EVGA case study in this paper, the move to HCI provides measurable returns on investment (ROI) and strategic business advantages over traditional IT architectures. For Microsoft environments, the combination of DataON's Cluster Block Architecture embodied in DataON's HCCA products, Intel's latest-generation processors which are optimized for scale-out systems, and Intel's high-performance NVMe SSD allow customers to deploy solutions that provide high performance, reduced physical footprint and power consumption, and lower CapEx and OpEx, while providing a scale-out capability that lets these systems easily grow to meet a customer's evolving needs.