Windows Server Software-Defined Storage with DataON® S2D Hyper-Converged Infrastructure with Intel® Xeon® Scalable Processors and Intel® Optane™ Technology

DataON with Intel and Microsoft

August 2018

Contents

Abstract

Microsoft released Storage Spaces with Windows Server 2012 as the foundation of a Windows software-defined storage (SDS) stack. Microsoft continued the transition to software-defined data centers (SDDC) with Window Server 2016 Storage Spaces Direct, with infrastructures built on SDS, hyper-converged and Azure stack hybrid cloud.

Intel® Optane™ technology is a unique combination of Intel® 3D XPoint™ memory media with Intel-built advanced system memory controller, interface hardware, and software IP. Together, these building blocks deliver a game-changing technology that presents new computer architecture opportunities and computing possibilities for a breath of markets. By placing this new memory media between the processor and slower SATA-based storage devices (HDD, SSHD or SATA SSD), systems can store commonly used data and programs closer to the processor, which allows the system to access this information more quickly and improve overall system responsiveness.

DataON is a storage industry veteran and pioneer of cluster-aware storage that has emerged as the industry-leading provider of hyper-converged infrastructure (HCI) and storage systems optimized for Microsoft® Windows Server environments. DataON is exclusively focused on customers who have made the "Microsoft choice" to deploy Windows Sever-based storage solutions.

DataON's S2D family of hyper-converged infrastructure feature Intel® Xeon® Scalable Processors with Intel C620 Chipsets, NVM Express (NVMe™) and Optane™ technology to deliver incredible performance and responsiveness, with greater VM density. Able to achieve over 3.2M IOPS in a 4-node cluster, the S2D family is part of the first Intel Select Solution with Windows Server Software-Defined Storage. It has also been validated for Windows Server Software-Defined, making it the first solution to achieve both Intel and Microsoft certifications.

Mellanox is a leading supplier of end-to-end Ethernet intelligent interconnect solutions and services for servers, storage, and HCI. Mellanox intelligent interconnect solutions increase data center efficiency by providing the highest throughput and lowest latency, delivering data faster to applications and unlocking system performance.

Bradley is a law firm that provides a full suite of legal services to business clients around the world. It was looking for a solution that could replace its traditional storage area network (SAN) and eliminate costly virtualization licensing feeds. They chose Windows Server 2016 Storage Spaces Direct and DataON S2D-5224i HCI with Intel® Optane™ technology to power its data center.

Customer Background

Overview

Bradley is a national law firm that provides business clients around the world with a full suite of legal services in dozens of industries and practice areas. Bradley has offices in Alabama, Florida, Mississippi, North Carolina, Tennessee, Texas, and the District of Columbia.

IT Challenge: Replace a VMware and SAN storage with a Microsoft hyper-converged infrastructure (HCI) in its main data center

Bradley had been using VMware and a traditional storage area network (SAN) in its main data center. However, after seeing how transformative the hyper-converged model has been for other companies, Bradley started looking into alternative options for virtualization, storage, and compute.

Jeff Chase, Senior Enterprise Architect at Bradley, started looking at software-defined storage. With the introduction of Windows Server 2016, Chase started to think now would be the right time to move to a hyper-converged infrastructure.

Bradley's goals in updating its IT infrastructure included:

- Moving from a traditional SAN to a hyper-converged infrastructure
- Finding the best technology with innovative features that would enable them to utilize hybrid cloud in the future
- Eliminating expensive SAN maintenance and virtual server annual licensing fees
- Improving performance and reducing latency

Bradley Customer Story At-a-Glance

Customer	**Bradley** (www.bradley.com) Bradley is a national law firm that provides business clients around the world with a full suite of legal services in dozens of industries and practice areas. Bradley has offices in in Alabama, Florida, Mississippi, North Carolina, Tennessee, Texas, and the District of Columbia.

Products and Workloads	**Hardware:**
	• DataON S2D-5224i Hyper-Converged Infrastructure (optimized for IOPS & capacity)
	  o Validated Intel Select Solution for Windows Server software-defined storage
	  o Validated Microsoft Windows Server Software-Defined hyper-converged infrastructure premium solution
	  o 1-Node/2U Rack Server
	  o Intel® Xeon® Scalable Gold Processor with Intel C620 Chipsets
	  o Intel® S3520 240GB SATA SSDs
	  o Intel® DC S4500 6GB SATA 3.8TB SDDs
	  o Intel® Optane™ DC P4800X 375GB NVMe U.2 SSDs
	  o Mellanox MSN270-CS2FO Spectrum 100GbE 1U Open Switch
	  o Mellanox ConnectX-4 EN Single Port QSFP 28 40/56 GbE RDMA Adapters
	
	Workloads:
	• Windows Server 2016 Storage Spaces Direct
	• SQL Server
	• SharePoint
	• Office

Industry	Professional Services (Legal)
Organization Size	1001 to 5000 employees
Country	United States

Windows Server Software-Defined Storage

Overview

Windows Server 2016 is a cloud-ready operating system that delivers new layers of security and Microsoft Azure-inspired innovation for the applications and infrastructure that power your business. From a storage perspective, Windows Server 2016 includes new features and enhancements for software-defined storage, as well as for traditional file servers.

Storage Spaces Direct enables building highly available and scalable storage using servers with local storage, such as DataON S2D hyper-converged infrastructure. It simplifies the deployment and management of software-defined storage systems and unlocks use of new classes of disk devices, such as SATA SSD and NVMe disk devices.

Health Service improves the day-to-day monitoring, operations, and maintenance experience of cluster resources on a Storage Spaces Direct cluster.

Storage Spaces Direct

Overview

Storage Spaces Direct is a new feature in Windows Server 2016. It uses servers with locally-attached drives to create highly available, highly scalable software-defined storage at a fraction of the cost of traditional SAN or NAS arrays.

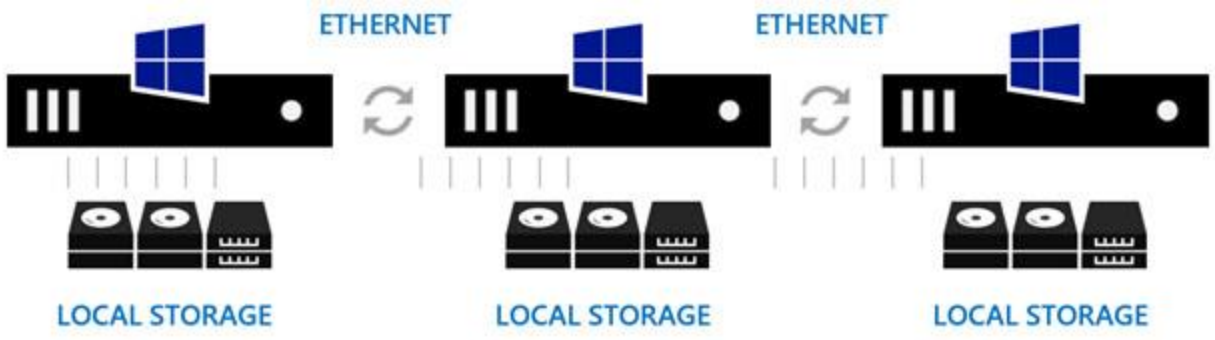

- Storage Spaces Direct utilizes NVMe storage devices for better performance and efficiency. These devices greatly reduce the I/O latency for storage, as well as reduce the CPU utilization to server storage.
- You can also use SATA storage devices to lower the cost of storage.
- In a DataON hyper-converged infrastructure, storage uses RDMA over Converged Ethernet (RoCE) as the storage fabric instead of having a shared physical storage fabric behind the servers.
- Storage Spaces Direct is included in Windows Server 2016 Datacenter edition.

Hardware

The most common configuration is 2-tiered physical storage, combining SSDs with traditional hard drives. You can also have an all-flash configuration with NVMe SSDs plus traditional SSDs where the NVMe SSDs are used for cache and the SATA SSDs are used for capacity.

In a 3-tiered physical storage configuration, NVMe SSDs are used for caching and both SATA SSDs and HDDs are used for additional tiering in the system. The SATA HDDs can be used to store the coldest data.

A traditional 10GbE or better Ethernet network can be used to connect servers/storage via RDMA. RDMA provides significant advantages because it lowers the latency of the storage I/O in the system and reduces the CPU utilization, resulting in higher IOPS performance in the system.

Feature Highlights

Built-in Always-on Cache – Storage Spaces Direct takes the fastest devices in the system and uses them as a write cache so the applications can continue immediately after writing data. It

also acts as a read cache for the most frequently read data from slower devices, for faster workloads. The cache automatically configures itself when you enable Storage Spaces Direct.

Single Pool of Storage – Storage Spaces Direct automatically groups the available storage devices into a single storage pool, eliminating the need for manual configuration or setting up multiple storage pools.

Scale from 2 to 16 nodes – Typically deployments consist of 2, 4, 8, 12 or 16 nodes.

Scales to 400+ drives – These servers can accommodate 26 drives per server, for a total of 416 drives. This provides petabytes of storage capacity in a fully scaled system.

Accelerated Erasure Coding – Storage Spaces Direct uses erasure coding for parity calculation in smaller deployments for the best possible efficiency and resiliency for two simultaneous failures. With larger clusters, storage efficiency is increased as there can be more data symbols without increasing the number of parity symbols. Developed in collaboration with Azure, erasure coding is a very efficient way to store data, driving storage efficiency of 50% or higher, compared to 33% for mirroring.

Deployment Options

Storage Spaces Direct is designed for two deployment options:

Converged – In this deployment option, storage and compute resources are in separate clusters. It layers a scale-out file server (SoFS) atop Storage Spaces Direct to provide network-attached storage over SMB3 file share. This allows for scaling compute/workload independently from the storage cluster, essential for larger-scale deployments such as Hyper-V IaaS (Infrastructure as a Service) for server providers and enterprises

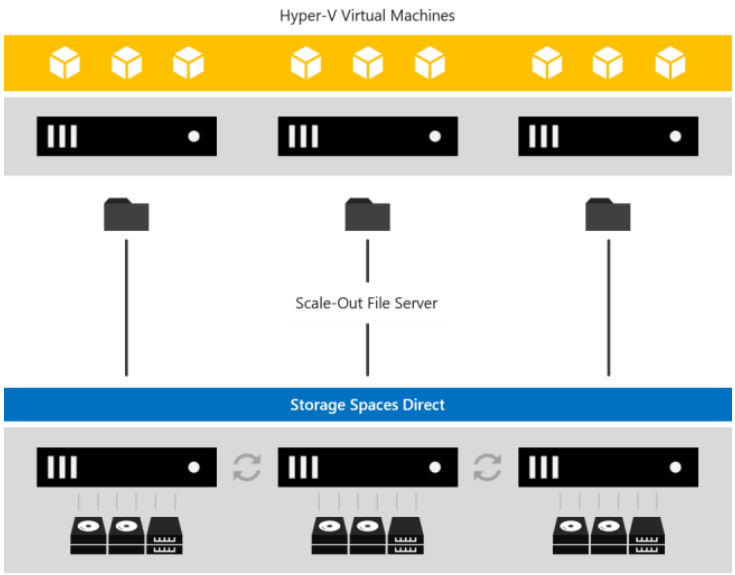

Hyper-Converged – This deployment option has one cluster for both compute and storage (such as DataON S2D HCI) and runs Hyper-V or SQL Server databases directly on the servers providing storage, storing their files on the local volumes. This eliminates the need to configure file server access and permissions and reduces hardware costs for small-to-medium business or remote office/brand office deployments. SQL Server runs natively on Storage Spaces Direct, providing outstanding IOPS and throughput for SQL database operations.

Storage Health Service Integration

Microsoft believes that partners such as DataON are best positioned to deliver an end-to-end solution user experience for customers. Customers have many different deployment models and use different kinds of hardware, so Microsoft provides partners an API foundation to build monitoring and management experiences.

There are two major components of the Microsoft Health Service API:

- **Alerts** – An event-driven model that identifies and surfaces hardware and software problems.
- **Metrics** – Hooks into the core stack in Windows to gather, synchronize, and aggregate statistics.

DataON's exclusive MUST tool is the first to market that leverages the Health Services API to provide visibility, monitoring, and management for Windows Storage deployments.

Windows Admin Center

Windows Admin Center is a new, locally-deployed, browser-based management tool set that lets you manage your Windows Servers with no Azure or cloud dependency. Windows Admin Center gives you full control over all aspects of your server infrastructure and is particularly useful for managing servers on private networks that are not connected to the Internet.

Windows Admin Center is the modern evolution of "in-box" management tools, like Server Manager and MMC. It complements System Center and Operations Management Suite – it's not a replacement.

Windows Admin Center runs in a web browser and manages Windows Server 2016, Windows Server 2012 R2, Windows Server 2012, Windows 10, and more through the Windows Admin Center gateway installed on Windows Server 2016 or Windows 10. The gateway manages servers by using Remote PowerShell and WMI over WinRM. The gateway is included with Windows Admin Center in a single lightweight .msi package that you can download.

The Windows Admin Center gateway, when published to DNS and given access through corresponding corporate firewalls, lets you securely connect to, and manage, your servers from anywhere with Microsoft Edge or Google Chrome.

- **Familiar functionality** – Windows Admin Center is the evolution of long-standing, well known management platforms like Microsoft Management Console (MMC), built from the ground up for the way systems are built and managed today. Windows Admin Center contains many of the familiar tools you currently use to manage Windows Servers and clients.

- **Easy to install and use** – Install on a Windows 10 computer, and start managing in minutes, or install on a Windows 2016 server acting as a gateway to enable your entire organization to manage computers from their web browser.

- **Complements existing solutions** – Windows Admin Center works with solutions like System Center and Azure management and security, adding to their capabilities to perform detailed, single-machine management tasks.

- **Manage from anywhere** – Publish your Windows Admin Center gateway server to the public Internet, then you can connect to and manage your servers from anywhere, all in a secure manner.

- **Enhanced security for your management platform** – Windows Admin Center has many enhancements that make your management platform more secure. Role-based access control lets you fine-tune which administrators have access to which management features. Gateway authentication options include local groups, local domain-based Active Directory, and cloud-based Azure Active Directory. Also, gain insight into management actions performed in your environment.

- **Azure integration** – Windows Admin Center has many points of integration with Azure services, including Azure Active Directory, Azure Backup, Azure Site Recovery, and more.

- **Manage hyper-converged clusters** – Windows Admin Center offers the best experience for managing hyper-converged clusters - including virtualized compute, storage, and networking components.

- **Extensibility** – Windows Admin Center was built with extensibility in mind from the beginning, with the ability for Microsoft and 3rd party developers to build tools and solutions. DataON's MUST visibility, monitoring, and management tool is available as an extension for use with Windows Admin Center.

DataON MUST provides a high level of infrastructure visibility, monitor, and management for Windows Server-based systems, networking and storage. MUST simplifies data center management helps enterprise customers transition from traditional SANs to Windows Server-based hyper-converged infrastructure.

MUST can be use through its standalone console or can be used within Windows Admin Center, allowing customers to use both Windows Admin Center an MUST through a single pane of glass.

MUST provides expanded functionality, including:

- **Historic data reporting** – Provides real-time and monthly dashboards of your system performance data including IOPS, latency, throughput on your cluster, storage pool, volume, and nodes.
- **Disk mapping** – MUST displays the device types and components in each of the nodes, providing a clear disk map of your entire node. It shows the number of disks, disk type, location and slot of each drive, and disk health status.
- **System alerts** – Leverages Windows Health Service faults to identify hardware failures, configuration issues, and resource saturation. It also provides a multi-level assessment of specific locations, fault descriptions, and recovery actions. You can also leverage third-party SNMP monitoring traps to alert you when you need disk or hardware replacements.
- **SAN-like call home service** – Prompted by system alerts, administrators can have automated email alerts sent to key contacts.

Intel Technologies for Windows Server 2016

Intel Data Center Blocks for Cloud and Microsoft Windows Server 2016

Intel Server Systems are engineered from the CPU out to help meet a wide range of business needs, from virtualization deployments to high-performance computing (HPC) infrastructure. Each design is built to a high specification, delivering server products with maximum processing power, great flexibility, excellent manageability, and high reliability. Stringent design and manufacturing practices, rigorous validation and testing, and excellent warranty and technical support ensure you receive incredible value.

The Intel Data Center Blocks for Cloud (Intel DCB for Cloud) and Microsoft Windows Server 2016 include both single node and multi-node server systems. Server systems within this product family were specifically created to offer Intel customers with pre-configured systems that are Microsoft Windows Server 2016 certified. Intel has extensively tested these systems to ensure best operation and reliability with the Microsoft Windows operating environment.

Intel® Xeon® Scalable Processors with Intel C620 Chipsets

Intel® Xeon® Scalable Processors with Intel® C620 Chipsets, formerly Purley (Skylake-SP and Lewisburg), provides the foundation for a powerful data center platform. Disruptive by design, this innovative processor sets a new level of platform convergence and capabilities across compute, storage, memory, network, and security. Enterprises and cloud and communications

service providers can now drive forward their most ambitious digital initiatives with a feature-rich, highly versatile, and more secure platform.

Enabling Greater Efficiencies and Lower TCO – Systems built on the Intel® Xeon® Scalable platform are design to deliver agile services and reduce TCO up to 65 percent due to lower software and OS licensing fees, and acquisition, maintenance, and infrastructure costs.

Supports More VMs – Intel® Xeon® Scalable platform enables 4.2x more VMs per server compared to earlier generations, allowing IT to increase their consolidation of more services on less hardware.

Pervasive, Breakthrough Performance – From its new Intel Mesh Architecture and widely expanded resources to its hardware-accelerating and newly integrated technologies, the Intel® Xeon® Scalable platform enables a new level of consistent, pervasive, and breakthrough performance.

Higher-Per-Core Performance – Up to 28 cores, delivering high performance and scalability for compute-intensive workloads across compute, storage, and network usages.

Greater Memory Bandwidth/Capacity – 50 percent increased memory bandwidth and capacity. Xix memory channels versus four memory channels of previous generation for memory-intensive workloads.

Expanded I/O – 48 lanes of PCIe 3.0 bandwidth and throughput for demanding I/O-intensive workloads.

Integrated Intel Ethernet with iWARP RDMA – Provides up to four 10GbE ports for high data throughput and low-latency workloads. Ideal for software-defined storage solutions, NVM Express over Fabric solutions, and virtual machine migration. Integrated in the chipset.

Support for Intel® Optane™ SSDs and Intel 3D NAND Solid State Drives – Delivers industry-leading combination of high throughput, low latency, high QoS, and ultra-high endurance to break through data access bottlenecks.

Deploy next generation storage with confidence with Intel Volume Management Devices (Intel VMD) – Enables hot swapping of NVMe SSDs from the PCIe bus without shutting down the system. Brings enterprise reliability, availability, and serviceability (RAS) features to NVMe SSDs, enabling deployment of next-generation storage with confidence.

Intel® Optane™ Technology

Intel® Optane™ memory is a system acceleration solution for new 7th Gen Intel® Core™ processor platforms. This solution comes in a module or DDR format and uses Intel® Optane™ technology, based on 3D XPoint™ memory media along with the Intel® Rapid Storage Technology (Intel® RST) driver. By placing this new memory media between the processor and slower SATA-based storage devices (HDD, SSHD or SATA SSD), Optane™ can help store commonly used data and programs closer to the processor, which allows the system to access this information more quickly and improve overall system responsiveness.

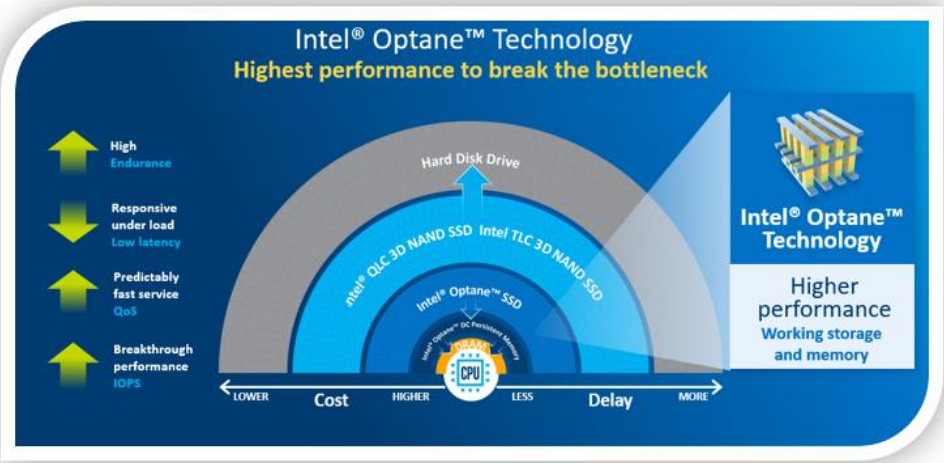

Intel® Optane™ Technology Benefits[1]

With broad availability in 2019, Intel® Optane™ DC Persistent Memory, when paired with Intel® Xeon® Scalable processors, revolutionizes the data center memory-storage hierarchy of the past and brings massive data sets closer to the CPU for faster time to insight. It also fundamentally changes data center resiliency, taking database restart times from hours down to minutes or seconds, and delivering an affordable memory footprint to scale system memory capacity to unprecedented level.

[1] Intel® Optane™ Technology Presentation, November 2018

Intel® Optane™ SSD DC P4800 Series

The Intel® Optane™ SSD DC P4800X is the first product to combine the attributes of memory and storage. With an industry-leading combination of high throughput, low latency, high QoS and ultra-high endurance, this innovative solution is optimized to break through data access bottlenecks by providing a new data storage tier. The DC P4800X accelerates applications for fast caching and fast storage to increase scale per server and reduce transaction costs for latency sensitive workloads. In addition, the DC P4800X enables data centers to deploy bigger and more affordable datasets to gain new insights from large memory pools. Features include:

- High throughput for breakthrough performance
- Low latency: responsive under load
- Predictably fast service: QoS
- Ultra-high endurance

In data centers, Intel DC P4800X can be used for fast storage or cache, or extended memory. Fast storage or cache refers to the tiering and layering which enable a better memory-to-storage hierarchy. The DC P4800X provides a new storage tier that breaks through the bottlenecks of traditional NAND storage to accelerate applications and enable more work to get done per server.

Used as extended memory, the Intel® Optane™ SSD works within a shared memory pool with DRAM at either the OS or application level enabling bigger memory or more affordable memory. Bigger memory dramatically increases the size of 'working sets' to enable new insights from data in growing segments such as scientific computing, healthcare and autonomous driving. More affordable memory means data centers can use Intel® Optane™ SSDs to displace some DRAM.

Intel® Optane™ DC Persistent Memory

Intel® Optane™ DC persistent memory represents a groundbreaking technology innovation. Delivered with the next-generation Intel® Xeon® Scalable processor, this technology will transform critical data workloads – from cloud and databases, to in-memory analytics, and content delivery networks.

Microsoft features unique capabilities delivered by Intel® Optane™ DC Persistent Memory through two special operating modes. Applications that have been specifically tuned can take advantage of App Direct mode, delivering the full value of the product's native persistence and

larger capacity. In Memory Mode, applications running in a supported OS or VM environment will be able to use the product as volatile memory, taking advantage of the additional system capacity made possible from module sizes up to 512 GB without needing to rewrite software.

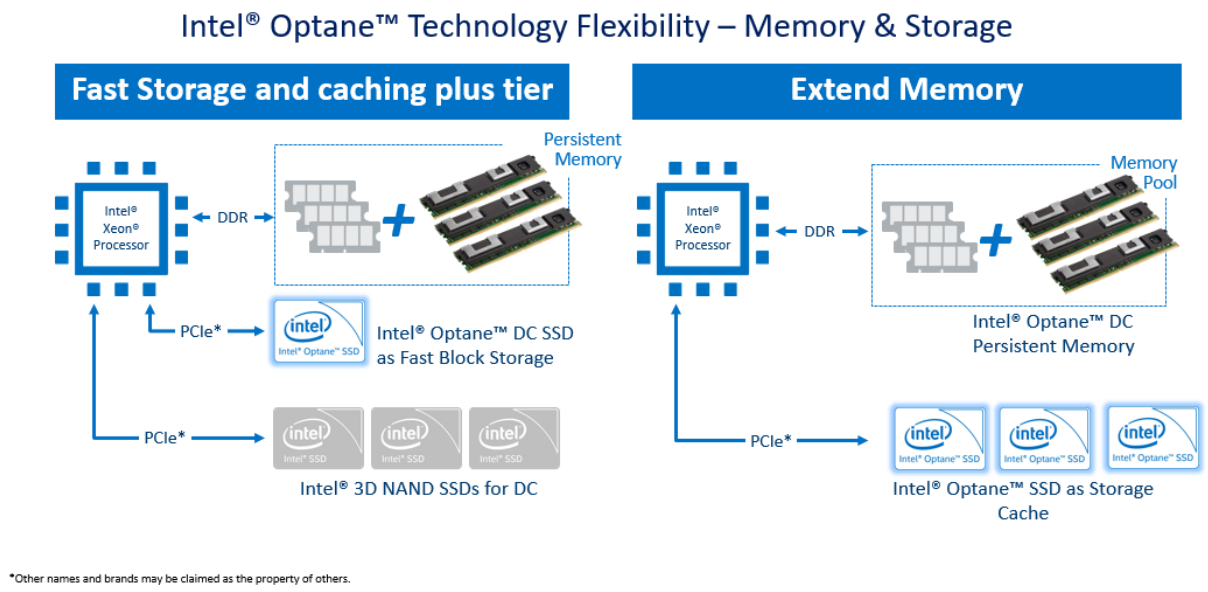

Intel® Optane™ Technology Configuration Options[2]

Intel® Optane™ DC Persistent Memory has slightly higher latency but offers affordable capacity and data persistence. Running Windows Server 2019 on a future Intel® Xeon® Scalable Cascade Lake processor configured with Intel® Optane™ Persistent Memory, Microsoft benchmarked a groundbreaking 13.7 million IOPS on 25% fewer servers.

[2] Intel® Optane™ Technology Presentation, November 2018

Intel Select Solutions

Overview

Intel Select Solutions are workload-optimized data center solutions that simplify and accelerate the process of selecting and deploying the hardware and software needed for today's workloads and applications. Intel Select Solutions:

- Simplify evaluation and eliminate guesswork via tightly specified hardware and software components
- Have pre-defined settings and system-wide tuning, enabling smooth deployment
- Are designed and benchmarked to perform optimally for specific workloads

Based on Intel® Xeon® Scalable processors, Intel Select Solutions enable enterprises to make confident choices in deploying data center infrastructure. The Intel Select Solution for Windows Server Software-Defined Storage features:

- Intel® Xeon® Scalable processors with Intel® C620 Series chipsets
- All-flash in the capacity tier with Intel 3D NAND SSDs using the NVMe or SATA interface

Intel Select Solutions are rigorously tested and benchmarked at a system-wide level. Well-known third party and industry benchmarking tools that simulate real-world loads are used to identify and eliminate pressure points.

Intel Select Solution for Windows Server Software-Defined Storage

The DataON S2D-5000 family is part of the first Intel Select Solution for Windows Server Software-Defined Storage and the first platform featuring Intel® Xeon® Scalable processors with Intel® C620 Series chipsets to achieve this certification. In addition to the Intel® Xeon® Scalable Processors Family, it also leverages NVM Express (NVMe™) and Intel® Optane™ memory to deliver breakthrough performance and incredible responsiveness, with greater VM density. Features include:

- Breakthrough performance and dramatically reduced disk latency with greater IOPS performance (over the previous generation) for write-intensive workloads by leveraging NVMe SSDs for the fast cache tier in Window Server 2016
- Greater VM density per cluster with Intel® Xeon® Scalable Processors, delivering up to 28 cores per socket, 18 DIMMS per CPU memory density, and 224 physical cores for a 4-node Storage Spaces Direct HCI deployment
- Highest networking efficiency and throughput with SMB3 RDMA fabric delivering 2x throughput compared to TCP/IP, less than 1µsec latency from VM-to-VM communication, and fewer CPU cycles per I/O with better core utilization
- Quicker, more responsive data access to slower SATA-based storage devices (HDD, SSHD or SATA SSD) with Intel® Optane™ Memory, system acceleration solution for new 7th Gen Intel® Core™ processor platforms

The Intel Select Solution has also achieved Windows Server 2016 Logo certification and Windows Server Software-Defined program certification. These certifications give customers confidence that this pre-configured and stress-tested solution follows Microsoft's requirements and best practices for a software-defined data center.

Microsoft Windows Server Software-Defined (WSSD)

The S2D-5000 family has also achieved Windows Server Software-Defined program certification. The goal of the Microsoft WSSD program is to ensure that customers have a seamless deployment and steady-state operational experience on validated hardware.

Based on Storage Spaces Direct technology, the WSSD program provides guidance for a Windows Server Software-Defined offering. It enables solution vendors to design and validate hardware and deploy a Windows Server 2016 software-defined infrastructure in a prescriptive manner that follows Microsoft's requirements and best practices.

Mellanox RDMA over Converged Ethernet (RoCE) Solutions for Windows Server 2016

Mellanox Low Latency 10/25/40/50/100 Gigabit Ethernet End-to-End Solutions

Mellanox offers complete 10/25/40/50/100GbE solutions for Windows Server 2016 data centers with end-to-end (RDMA over Converged Ethernet (RoCE) support. These end-to-end solutions deliver high bandwidth and low latency to I/O intensive applications and fast flash storage, enabling data centers to operate with high performance and efficiency. By supporting speeds from 10GbE to 100GbE, Mellanox Spectrum® switches and ConnectX-4® network adapter cards give IT managers flexibility in how they deploy higher bandwidth to the servers, providing simple upgrade paths from 10GbE to 40GbE, or to 25/100GbE.

Mellanox's networking solutions based on Ethernet or RoCE provide outstanding price, performance, and power value proposition for networking and storage I/O processing capacities. Data accessed over RDMA using file-based protocols such as Microsoft SMB Direct (SMB over RDMA) enables:

Increased throughput – Leverages the full throughput of high-speed networks in which the network adapters coordinate the transfer of large amounts of data at line speed.

Low latency – Provides extremely fast responses to network requests, and, as a result, makes remote file storage feel as if it is directly attached block storage.

Low CPU utilization – Uses fewer CPU cycles when transferring data over the network, which leaves more power available to server applications.

Mellanox 10/25/40/50/56/100 Gigabit Ethernet Switching

Mellanox MSN2700-CS2FO Spectrum 100GbE 1U Open Switch

Keeping with the Mellanox tradition of setting performance record switch systems, the MSN2700 introduces the world's lowest latency for a 100GbE switching and routing element and does so while having the lowest power consumption in the market. With the MSN2700, the use of 25, 40, 50 and 100GbE in large scale is enabled without changing power infrastructure facilities.

The MSN2700 is part of Mellanox's complete end-to-end solution which provides 10GbE through 100GbE interconnectivity within the data center. Other devices in this solution include the ConnectX® family of network interface cards, and LinkX® copper or fiber cabling. This end-to-end solution is topped with Mellanox NEO™, a management application that relieves some of the major obstacles when deploying a network. NEO enables a fully certified and interoperable design and speeds up time to service.

The MSN2700 switch provides the most predictable, highest density 100GbE switching platform for the growing demands of today's data centers. The MSN2700 switch is an ONIE (Open Network Install Environment) based platform for allowing a multitude of operating systems to be mounted on it and utilizing the advantages of Open Networking and the capabilities of the Mellanox Spectrum ASIC.

The MSN2700 switch is an ideal spine and top-of-rack (ToR) solution, allowing maximum flexibility, with port max speeds spanning from 10Gb/s to 100Gb/s per port and port density that enables full rack connectivity to any server at any max speed. The uplink ports allow a variety of blocking ratios that suit any application requirement.

Mellanox 10/25/40/50/100 Gigabit Ethernet Adapter Cards

ConnectX-4 10/25/40/50/100 Gigabit Ethernet adapters provide exceptional high performance for the most demanding data applications. ConnectX-4 adapters support RoCE specifications delivering low-latency and high-performance over Ethernet networks. Leveraging data center bridging (DCB) capabilities as well as advanced congestion control hardware mechanisms, ConnectX-4 RoCE provides efficient low-latency RDMA services over Layer 2 and Layer 3 networks. In addition to RoCE offload, ConnectX-4 adapters implement a rich set of features, from network function offload (e.g., LSO/LRO/RSS/TSS), SR-IOV support for I/O virtualization and Accelerated Switching and Packet Processing (ASAP2) technology to perform OVS data plane and VTEP functions.

With RDMA/RoCE offload in the ConnectX-4 adapters and predictable high performance by the Spectrum switches, the Mellanox Ethernet solution accelerates Microsoft Server 2016 Storage Spaces Direct, unleashes the power of faster storage devices such as NVMe SSDs, and greatly improves server CPU and application efficiency.

Mellanox 40/56 Gigabit ConnectX-4 EN Ethernet Adapter Cards

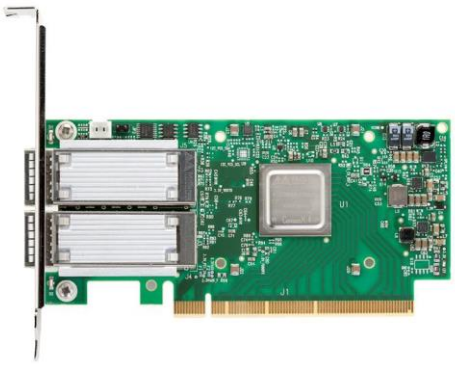

Mellanox ConnectX-4 40/56 GbE Adapters

ConnectX-4 EN network controller with 40/56 Gb/s Ethernet connectivity, provide the highest performance and most flexible solution for high performance, Web 2.0, Cloud, data analytics, database, and storage platforms.

With the exponential growth of data being shared and stored by applications and social networks, the need for high-speed and high-performance compute and storage data centers is skyrocketing.

ConnectX-4 EN provides exceptional high performance for the most demanding data centers, public and private clouds, Web2.0 and Big Data applications, and storage systems, enabling today's corporations to meet the demands of the data explosion.

ConnectX-4 EN provides an unmatched combination of 40/56 Gb/s bandwidth in a single port, the lowest available latency, and specific hardware offloads, addressing both today's and the next generation's compute and storage data center demands.

DataON: The Certified Microsoft Platform

The DataON Difference

DataON is exclusively focused on customers who have made the "Microsoft choice" to deploy a Windows Server-based storage solution. Our team of Microsoft Server experts know how to design, deploy and support Windows Server storage and will work with you to performance tune your workloads. DataON storage solutions are:

- Certified for Windows Server 2012 R2, 2016 SDDC, Windows Server Software-Defined, and Intel Select Solution for Windows Server Software-Defined Storage
- Customer-proven with over 600 enterprise installations and 120PB of DataON S2D HCI deployments
- Optimized by our team of Microsoft experts to ensure successful deployments into your IT environment, tuned to your workloads

DataON S2D Family of Hyper-Converged Infrastructure

The DataON S2D family of hyper-converged infrastructure is built to optimize the full stack of Microsoft Storage Spaces Direct in a hyper-converged platform. It is built with integrated compute, network and storage infrastructure with near-linear scalability to simplify and maximize the deployment of Microsoft applications, virtualization, data protection and hybrid cloud services. Each pre-configured cluster can support 40 Hyper-V VMs per node, for expanded capacity and operational flexibility.

From scale-out file server (SoFS) and software storage bus to storage and networking hardware, this appliance runs on the cluster Shared Volumes Resilient File System (ReFS) and uses high performance Intel® Optane™ SSDs, NVMe SSDs with SMB3 networking to maximize performance and capacity.

DataON S2D-5224i Hyper-Converged Infrastructure

The DataON S2D-5224i is an Intel-based server system that provides scale-out and scale-up infrastructure and management services for deploying Microsoft Windows Server 2016. It features Intel® Xeon® Scalable Processors with Intel® C620 Chipsets and Intel® Optane™ SSDs and NVMe SSDs. It is optimized for balanced IOPS and performance for the delivery of core Microsoft services and enterprise applications. The S2D-5224i is designed on three core principles:

- Scale-out hyper-converged cluster
- Integrated software-defined services
- Complete visibility and management of the storage infrastructure via DataON's exclusive MUST infrastructure and management software tool

DataON MUST™ (Management Utility Software Tool) Visibility and Management Tool

The DataON S2D hyper-converged solutions are integrated with the exclusive DataON MUST visibility and management tool. It provides infrastructure visibility and management for Microsoft's new suite of software-defined storage stack technologies like Storage Spaces Direct, Storage Replica, and storage quality of service (QoS) which, based on your policy, monitors hardware and software storage infrastructure to identify potential problems. Using an event-driven model for rapid detection with minimal overhead, MUST also provides on-demand access to curated collections of hyper-converged and converged clusters, storage performance, and capacity metrics. The MUST dashboard display is designed to efficiently and dynamically connect the dots to help provide root cause analysis.

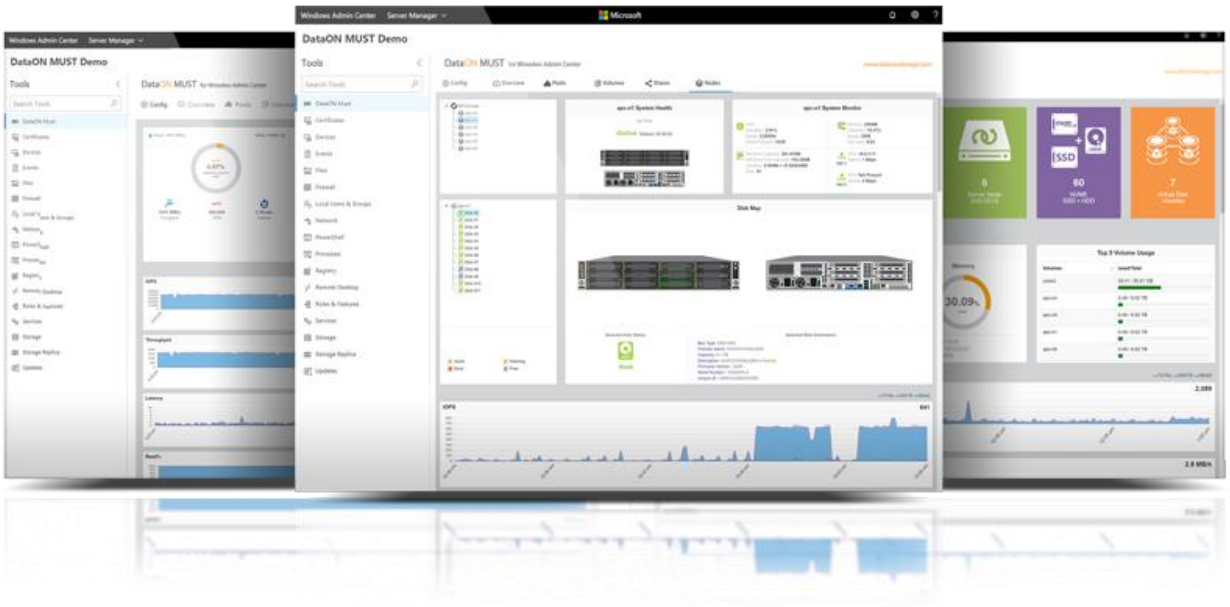

MUST offers complete integration with Microsoft Storage Health Services for Windows Server 2016. It is built to provide multiple tiers of storage visibility and monitoring.

SDDC & Hyper-Converged Infrastructure Tier – Provides system-level information on performance, capacity, and hardware inventory, as well as faults and alerts. MUST gives you a dashboard-level view of your operations, analytics, infrastructure health management, storage systems metrics and even event logging insights.

Systems and Storage Services Audit Log Tier – Provides detailed logging-level visibility for events, so you can perform root cause analysis and export source data for analytics.

Hyper-Converged Cluster Appliance (HCCA)/Node Tier – Displays pool, volume and device-level performance, health and operational analytics for your cluster. This enables you to proactively perform systems maintenance and better understand your requirements for workload migrations.

SAN-like Call Home Service Support – Leveraging the Health Services Faults in Windows Server 2016, administrators can have automated email alerts sent to key contacts. You can also leverage third party SNMP monitoring traps to alert you when you need disk or hardware replacements.

MUST now integrates with Windows Admin Center, a locally-deployed, browser-based tool that gives you full control over all aspects of your server infrastrcuture. Windows Admin Center is the modern evolution of 'in-box' managemnt tool, like Server Manager and MMC. It complements System Center and Operations Management Suite but it's not a replacement.

MUST is available as an extension to Windows Admin Center, providing additional functionality, including:

- **Historic data reporting** – Provides real-time and monthly dashboards of your system performance data including IOPS, latency, throughput on your cluster, storage pool, volume, and nodes.
- **Disk mapping** – MUST displays the device types and components in each of the nodes, providing a clear disk map of your entire node. It shows the number of disks, disk type, location and slot of each drive, and disk health status.
- **System alerts** – Leverages Windows Health Service faults to identify hardware failures, configuration issues, and resource saturation. It also provides a multi-level assessment of specific locations, fault descriptions, and recovery actions. You can also leverage third-party SNMP monitoring traps to alert you when you need disk or hardware replacements.
- **SAN-like call home service** – Prompted by system alerts, administrators can have automated email alerts sent to key contacts.

Deploying a Windows Server 2016 Software-Defined Storage Solution with Bradley

Finding the Solution

Jeff Chase is a regular follower of tech news and first heard about Windows Server and DataON from Aidan Finn, a long-time Microsoft MVP. He monitored the evolution of Windows Server for several years and when Microsoft announced Windows Server 2016, he felt then might be the right time to try something different.

"I really liked the features that were included with Storage Spaces Direct and the pace of innovation from Microsoft in storage was impressive," said Chase. "In addition, the roadmap for Microsoft's solution lines up perfectly with where the industry is going with hybrid cloud. And the features for Hyper-V are starting to get better than VMware." Another advantage with a Microsoft solution is that Bradley is already a Microsoft shop. It already uses SQL Server, SharePoint, and Office so integration would not be an issue.

From his research, he knew that there were very few vendors that specialized in Microsoft solutions. With endorsements from Finn and Microsoft, DataON was Chase's first choice for a vendor to work with.

Chase did an initial small HCI deployment with Storage Spaces Direct in a firm office and was very happy with the performance and how well it performed. As a result, he was very comfortable with recommending the change and Bradley purchased a DataON S2D-5224i hyper-converged infrastructure with all-flash Intel NVMe SSDs and Intel® Optane™ SSDs. The firm also purchased 100GbE Mellanox networking switches to plan for faster servers, faster storage, and greater bandwidth needs in the future.

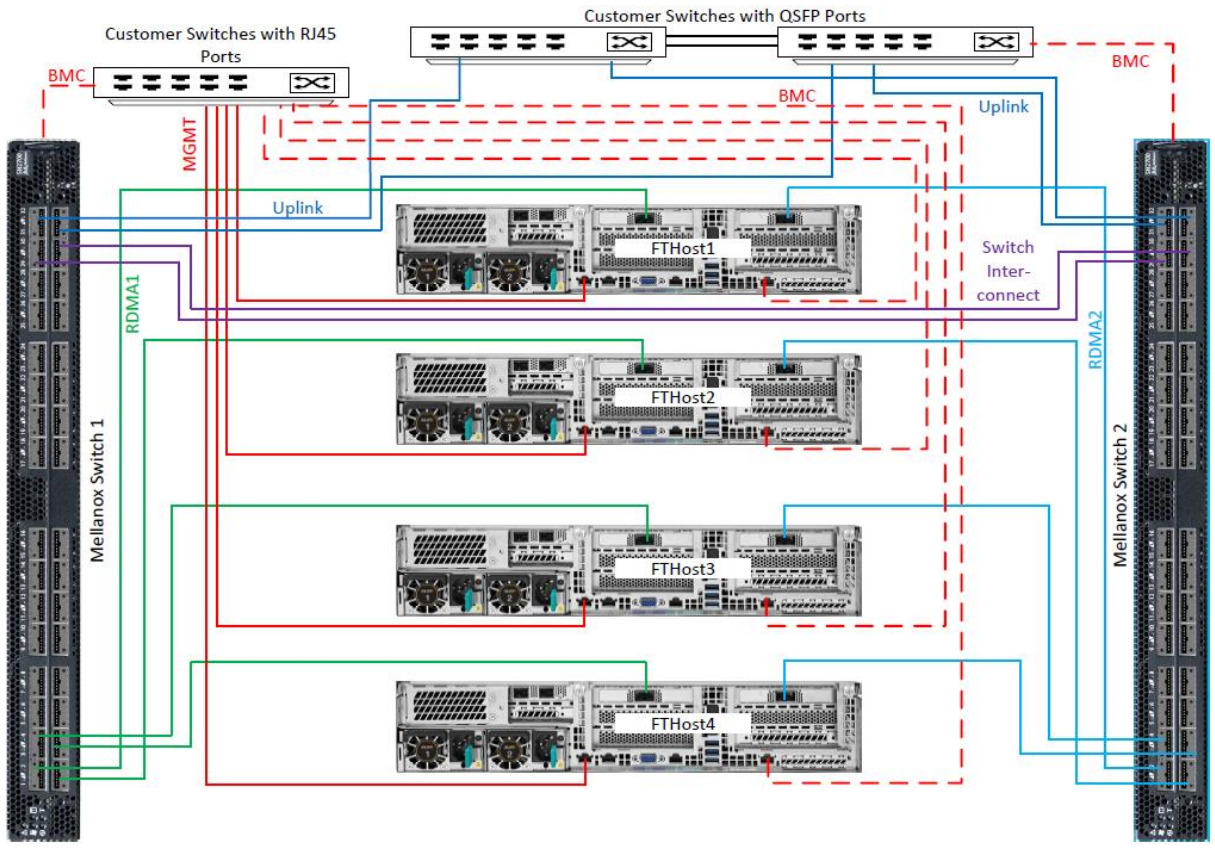

Benchmarks and Results

"Storage Spaces Direct is a really good product and I really like the DataON hardware. I'm glad we went down this path," said Chase.

Bradley is already seeing the increased performance of their new hyper-converged infrastructure over their previous all-flash NAS. "Everything is so much faster. Reboots are almost instantaneous. Latency has dropped to less than 1 microsecond," said Chase.

Chase continued, "The hyper-converged solution with Intel® Optane™ NVMe SSDs has delivered a vast performance improvement in both latency and IOPS."

Bradley is also very happy with the cost savings from the Storage Spaces Direct and DataON solution. "Storage Spaces Direct is a breath of fresh air from a cost perspective," said Chase.

"In addition, DataON was very competitive on their hardware price compared to our previous blades and racks. It was a no-brainer to switch to DataON," said Chase.

The Bradley team is also using Windows Admin Center, Microsoft's new browser-based management tool. It complements System Center and Operations Management Suite and allows

users to monitor and manage Hyper-V, Azure and Storage Spaces Direct from a single console. "It works really well and we're very happy with it," said Chase.

In addition, the team is also using DataON's MUST management and utility tool extension for Windows Admin Center to monitor and manage their Windows Server cluster. MUST fully integrates with Windows Admin Center and provides historical data reporting, disk mapping and system alerts and call home serves for when problems arise.

VM Fleet Testing

Multi-resilient volumes for maximum capacity with all-flash NVMe SSDs.

Hardware Information / per node		Windows Setup	
CPU	Intel Gold 6132 x2 (2.6 Ghz, 28 cores, 56 logical processors)	**Storage pool**	1
Memory	768GB	**Number of virtual disk**	6
RDMA	Mellanox ConnectX-4 40/56Gb Single Port Adapter x2	**Number of columns**	8
NVMe	Intel P4800 375GB NVMe SSD x4	**Interleave size**	256KB (default)
SAS/ SATA	Intel S4500 3.8TB SATA SSD x20		
HDD	N/A		

Volume Information

Volume	File System	Capacity	Resiliency	Size (Mirror)	Size (Parity)	Footprint	Efficiency
Volume1	ReFS	32TB	MRV (10+90)	2TB	30TB	67TB	48%
Volume2	ReFS	32TB	MRV (10+90)	2TB	30TB	67TB	48%

Volume3	ReFS	32TB	MRV (10+90)	2TB	30TB	67TB	48%
Volume4	ReFS	32TB	MRV (10+90)	2TB	30TB	67TB	48%
Volume5	ReFS	32TB	MRV (10+90)	2TB	30TB	67TB	48%
Volume6	ReFS	32TB	MRV (10+90)	2TB	30TB	67TB	48%

Random Reads and Writes

Block size 4Kb, 8 threads, 8 outstanding I/O (100% read / 0% read)

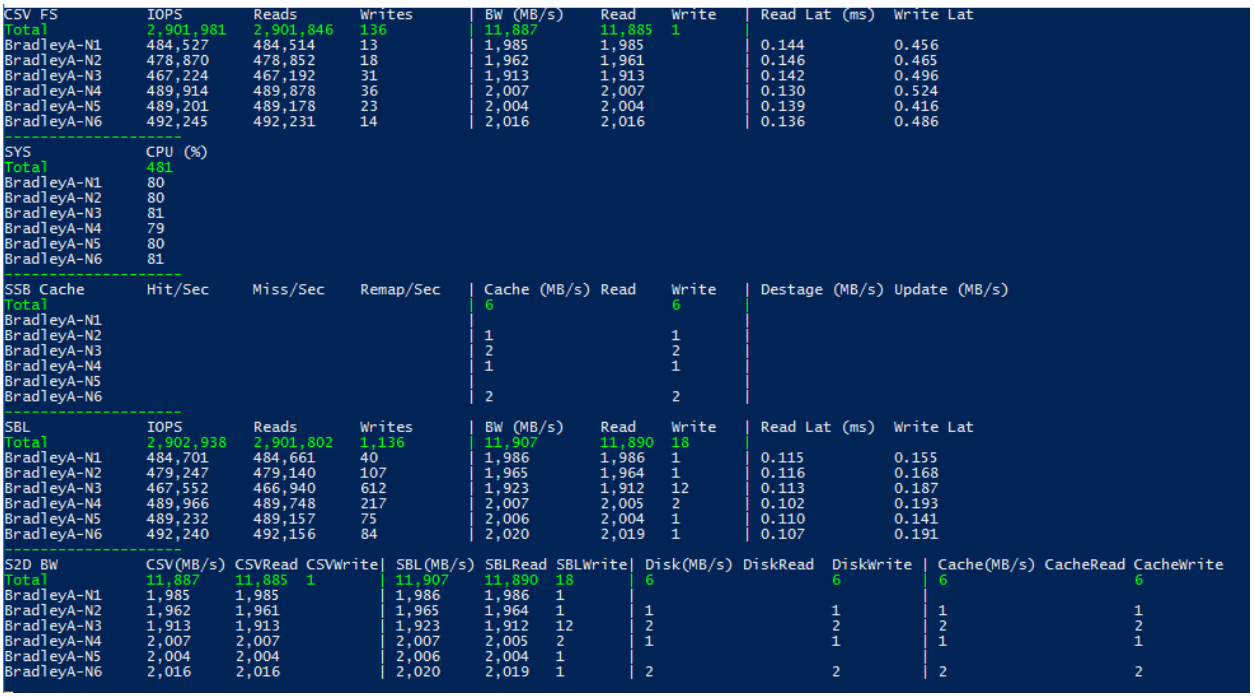

CSV FS	IOPS	Reads	Writes	BW (MB/s)	Read	Write	Read Lat (ms)	Write Lat
Total	2,901,981	2,901,846	136	11,887	11,885	1		
BradleyA-N1	484,527	484,514	13	1,985	1,985		0.144	0.456
BradleyA-N2	478,870	478,852	18	1,962	1,961		0.146	0.465
BradleyA-N3	467,224	467,192	31	1,913	1,913		0.142	0.496
BradleyA-N4	489,914	489,878	36	2,007	2,007		0.130	0.524
BradleyA-N5	489,201	489,178	23	2,004	2,004		0.139	0.416
BradleyA-N6	492,245	492,231	14	2,016	2,016		0.136	0.486

SYS	CPU (%)
Total	481
BradleyA-N1	80
BradleyA-N2	80
BradleyA-N3	81
BradleyA-N4	79
BradleyA-N5	80
BradleyA-N6	81

SSB Cache	Hit/Sec	Miss/Sec	Remap/Sec	Cache (MB/s) Read	Write	Destage (MB/s)	Update (MB/s)
Total				6	6		
BradleyA-N1							
BradleyA-N2				1	1		
BradleyA-N3				2	2		
BradleyA-N4				1	1		
BradleyA-N5							
BradleyA-N6				2	2		

SBL	IOPS	Reads	Writes	BW (MB/s)	Read	Write	Read Lat (ms)	Write Lat
Total	2,902,938	2,901,802	1,136	11,907	11,890	18		
BradleyA-N1	484,701	484,661	40	1,986	1,986	1	0.115	0.155
BradleyA-N2	479,247	479,140	107	1,965	1,964	1	0.116	0.168
BradleyA-N3	467,552	466,940	612	1,923	1,912	12	0.113	0.187
BradleyA-N4	489,966	489,748	217	2,007	2,005	2	0.102	0.193
BradleyA-N5	489,232	489,157	75	2,006	2,004	1	0.110	0.141
BradleyA-N6	492,240	492,156	84	2,020	2,019	1	0.107	0.191

S2D BW	CSV(MB/s)	CSVRead	CSVWrite	SBL(MB/s)	SBLRead	SBLWrite	Disk(MB/s)	DiskRead	DiskWrite	Cache(MB/s)	CacheRead	CacheWrite
Total	11,887	11,885	1	11,907	11,890	18	6		6	6		6
BradleyA-N1	1,985	1,985		1,986	1,986	1						
BradleyA-N2	1,962	1,961		1,965	1,964	1	1		1	1		1
BradleyA-N3	1,913	1,913		1,923	1,912	12	2		2	2		2
BradleyA-N4	2,007	2,007		2,007	2,005	2	1		1	1		1
BradleyA-N5	2,004	2,004		2,006	2,004	1						
BradleyA-N6	2,016	2,016		2,020	2,019	1	2		2	2		2

Block size 4Kb, 8 threads, 8 outstanding I/O (0% read / 100% write)

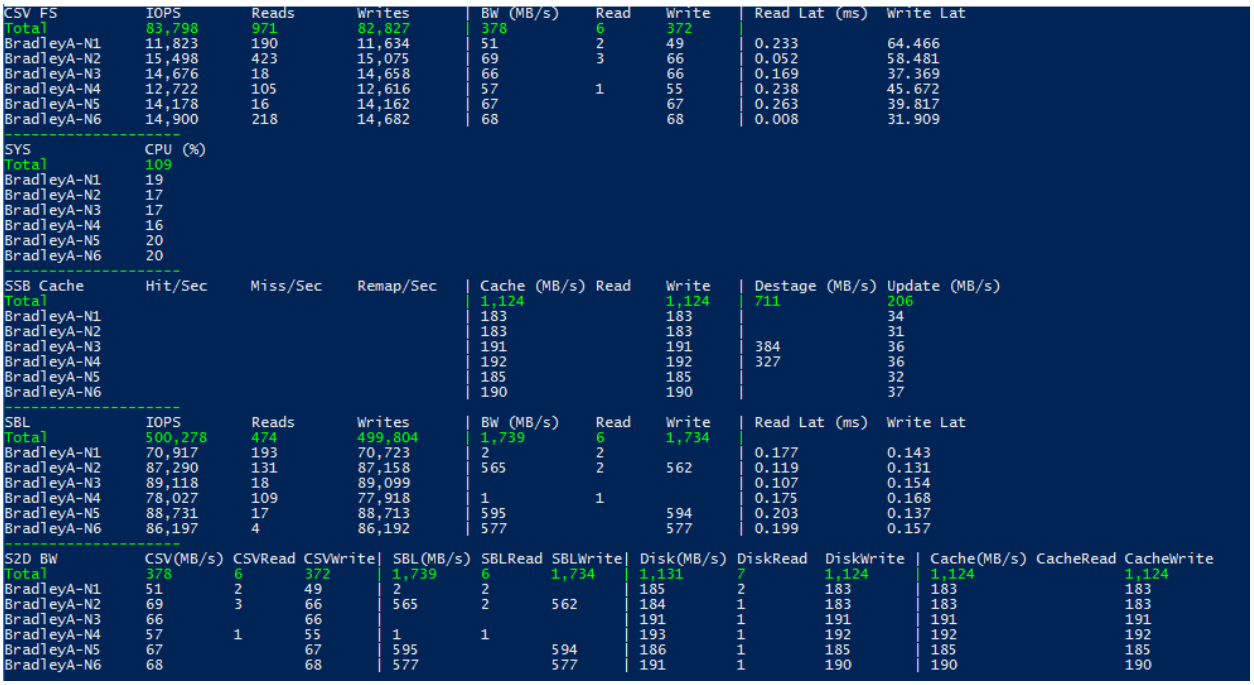

CSV FS	IOPS	Reads	Writes	BW (MB/s)	Read	Write	Read Lat (ms)	Write Lat
Total	83,798	971	82,827	378	6	372		
BradleyA-N1	11,823	190	11,634	51	2	49	0.233	64.466
BradleyA-N2	15,498	423	15,075	69	3	66	0.052	58.481
BradleyA-N3	14,676	18	14,658	66		66	0.169	37.369
BradleyA-N4	12,722	105	12,616	57	1	55	0.238	45.672
BradleyA-N5	14,178	16	14,162	67		67	0.263	39.817
BradleyA-N6	14,900	218	14,682	68		68	0.008	31.909

SYS	CPU (%)
Total	109
BradleyA-N1	19
BradleyA-N2	17
BradleyA-N3	17
BradleyA-N4	16
BradleyA-N5	20
BradleyA-N6	20

SSB Cache	Hit/Sec	Miss/Sec	Remap/Sec	Cache (MB/s) Read	Write	Destage (MB/s)	Update (MB/s)
Total				1,124	1,124	711	206
BradleyA-N1				183	183		34
BradleyA-N2				183	183		31
BradleyA-N3				191	191	384	36
BradleyA-N4				192	192	327	36
BradleyA-N5				185	185		32
BradleyA-N6				190	190		37

SBL	IOPS	Reads	Writes	BW (MB/s)	Read	Write	Read Lat (ms)	Write Lat
Total	500,278	474	499,804	1,739	6	1,734		
BradleyA-N1	70,917	193	70,723	2	2		0.177	0.143
BradleyA-N2	87,290	131	87,158	565	2	562	0.119	0.131
BradleyA-N3	89,118	18	89,099				0.107	0.154
BradleyA-N4	78,027	109	77,918	1	1		0.175	0.168
BradleyA-N5	88,731	17	88,713	595		594	0.203	0.137
BradleyA-N6	86,197	4	86,192	577		577	0.199	0.157

S2D BW	CSV(MB/s)	CSVRead	CSVWrite	SBL(MB/s)	SBLRead	SBLWrite	Disk(MB/s)	DiskRead	DiskWrite	Cache(MB/s)	CacheRead	CacheWrite
Total	378	6	372	1,739	6	1,734	1,131	7	1,124	1,124		1,124
BradleyA-N1	51	2	49	2	2		185	2	183	183		183
BradleyA-N2	69	3	66	565	2	562	184	1	183	183		183
BradleyA-N3	66		66				191	1	191	191		191
BradleyA-N4	57	1	55	1	1		193	1	192	192		192
BradleyA-N5	67		67	595		594	186	1	185	185		185
BradleyA-N6	68		68	577		577	191	1	190	190		190

Block size 4Kb, 8 threads, 8 outstanding I/O (70% read / 30% write)

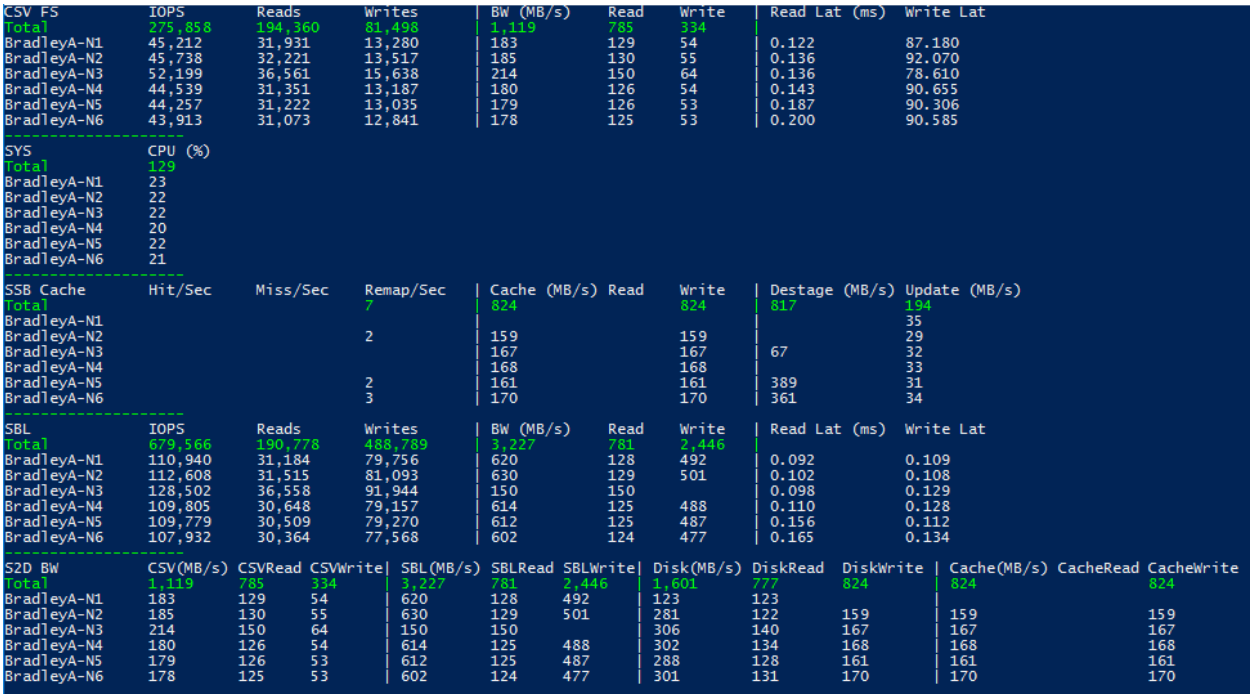

CSV FS	IOPS	Reads	Writes	BW (MB/s)	Read	Write	Read Lat (ms)	Write Lat
Total	275,858	194,360	81,498	1,119	785	334		
BradleyA-N1	45,212	31,931	13,280	183	129	54	0.122	87.180
BradleyA-N2	45,738	32,221	13,517	185	130	55	0.136	92.070
BradleyA-N3	52,199	36,561	15,638	214	150	64	0.136	78.610
BradleyA-N4	44,539	31,351	13,187	180	126	54	0.143	90.655
BradleyA-N5	44,257	31,222	13,035	179	126	53	0.187	90.306
BradleyA-N6	43,913	31,073	12,841	178	125	53	0.200	90.585

SYS	CPU (%)
Total	129
BradleyA-N1	23
BradleyA-N2	22
BradleyA-N3	22
BradleyA-N4	20
BradleyA-N5	22
BradleyA-N6	21

SSB Cache	Hit/Sec	Miss/Sec	Remap/Sec	Cache (MB/s) Read	Write	Destage (MB/s)	Update (MB/s)
Total			7	824	824	817	194
BradleyA-N1							35
BradleyA-N2			2	159	159		29
BradleyA-N3				167	167	67	32
BradleyA-N4				168	168		33
BradleyA-N5			2	161	161	389	31
BradleyA-N6			3	170	170	361	34

SBL	IOPS	Reads	Writes	BW (MB/s)	Read	Write	Read Lat (ms)	Write Lat
Total	679,566	190,778	488,789	3,227	781	2,446		
BradleyA-N1	110,940	31,184	79,756	620	128	492	0.092	0.109
BradleyA-N2	112,608	31,515	81,093	630	129	501	0.102	0.108
BradleyA-N3	128,502	36,558	91,944	150	150		0.098	0.129
BradleyA-N4	109,805	30,648	79,157	614	125	488	0.110	0.128
BradleyA-N5	109,779	30,509	79,270	612	125	487	0.156	0.112
BradleyA-N6	107,932	30,364	77,568	602	124	477	0.165	0.134

S2D BW	CSV(MB/s)	CSVRead	CSVWrite	SBL(MB/s)	SBLRead	SBLWrite	Disk(MB/s)	DiskRead	DiskWrite	Cache(MB/s)	CacheRead	CacheWrite
Total	1,119	785	334	3,227	781	2,446	1,601	777	824	824		824
BradleyA-N1	183	129	54	620	128	492	123	123				
BradleyA-N2	185	130	55	630	129	501	281	122	159	159		159
BradleyA-N3	214	150	64	150	150		306	140	167	167		167
BradleyA-N4	180	126	54	614	125	488	302	134	168	168		168
BradleyA-N5	179	126	53	612	125	487	288	128	161	161		161
BradleyA-N6	178	125	53	602	124	477	301	131	170	170		170

Block size 4Kb, 8 threads, 8 outstanding I/O (50% read / 50% write)

CSV FS	IOPS	Reads	Writes	BW (MB/s)	Read	Write	Read Lat (ms)	Write Lat
Total	175,494	87,483	88,011	720	359	361		
BradleyA-N1	30,324	15,135	15,189	124	62	62	0.173	61.415
BradleyA-N2	28,544	14,210	14,334	117	58	59	0.155	79.875
BradleyA-N3	31,056	15,494	15,562	128	64	64	0.154	57.845
BradleyA-N4	28,642	14,292	14,349	117	59	59	0.186	87.940
BradleyA-N5	28,650	14,260	14,390	118	59	59	0.218	64.362
BradleyA-N6	28,278	14,092	14,186	116	58	58	0.213	89.053

SYS	CPU (%)
Total	127
BradleyA-N1	25
BradleyA-N2	21
BradleyA-N3	19
BradleyA-N4	20
BradleyA-N5	22
BradleyA-N6	21

SSB Cache	Hit/Sec	Miss/Sec	Remap/Sec	Cache (MB/s) Read	Write	Destage (MB/s)	Update (MB/s)
Total				1,049	1,049	1,075	200
BradleyA-N1				171	171	358	35
BradleyA-N2				170	170		30
BradleyA-N3				175	175	211	33
BradleyA-N4				180	180	439	36
BradleyA-N5				172	172	67	31
BradleyA-N6				181	181		35

SBL	IOPS	Reads	Writes	BW (MB/s)	Read	Write	Read Lat (ms)	Write Lat
Total	604,089	87,450	516,640	3,035	359	2,676		
BradleyA-N1	105,342	15,135	90,207	619	62	557	0.121	0.121
BradleyA-N2	97,172	14,210	82,962	570	58	512	0.111	0.116
BradleyA-N3	105,957	15,494	90,462	622	64	558	0.110	0.137
BradleyA-N4	98,745	14,297	84,448	583	59	525	0.142	0.137
BradleyA-N5	99,736	14,260	85,476	583	59	525	0.171	0.125
BradleyA-N6	97,138	14,053	83,084	58	58		0.166	0.140

S2D BW	CSV(MB/s)	CSVRead	CSVWrite	SBL(MB/s)	SBLRead	SBLWrite	Disk(MB/s)	DiskRead	DiskWrite	Cache(MB/s)	CacheRead	CacheWrite
Total	720	359	361	3,035	359	2,676	1,404	354	1,049	1,049		1,049
BradleyA-N1	124	62	62	619	62	557	227	56	171	171		171
BradleyA-N2	117	58	59	570	58	512	226	55	170	170		170
BradleyA-N3	128	64	64	622	64	558	236	61	175	175		175
BradleyA-N4	117	59	59	583	59	525	242	62	180	180		180
BradleyA-N5	118	59	59	583	59	525	232	59	172	172		172
BradleyA-N6	116	58	58	58	58		241	60	181	181		181

Block size 4Kb, 8 threads, 1 outstanding I/O (100% read / 0% write)

CSV FS	IOPS	Reads	Writes	BW (MB/s)	Read	Write	Read Lat (ms)	Write Lat
Total	1,839,036	1,836,755	2,282	7,567	7,548	19		
BradleyA-N1	303,669	303,293	376	1,250	1,247	3	0.090	0.283
BradleyA-N2	302,813	302,398	414	1,247	1,243	4	0.094	0.266
BradleyA-N3	243,045	242,740	306	998	995	3	0.087	0.277
BradleyA-N4	329,468	329,025	443	1,356	1,352	3	0.083	0.266
BradleyA-N5	323,053	322,664	389	1,330	1,327	3	0.087	0.249
BradleyA-N6	336,989	336,635	354	1,386	1,383	3	0.086	0.272

SYS	CPU (%)
Total	227
BradleyA-N1	38
BradleyA-N2	38
BradleyA-N3	35
BradleyA-N4	39
BradleyA-N5	38
BradleyA-N6	40

SSB Cache	Hit/Sec	Miss/Sec	Remap/Sec	Cache (MB/s) Read	Write	Destage (MB/s)	Update (MB/s)
Total				66	66		1
BradleyA-N1				11	11		
BradleyA-N2				12	12		
BradleyA-N3				11	11		
BradleyA-N4				11	11		
BradleyA-N5				11	11		
BradleyA-N6				11	11		

SBL	IOPS	Reads	Writes	BW (MB/s)	Read	Write	Read Lat (ms)	Write Lat
Total	1,849,493	1,835,452	14,040	7,692	7,546	145		
BradleyA-N1	305,994	303,394	2,600	1,274	1,248	26	0.069	0.094
BradleyA-N2	304,586	302,148	2,437	1,270	1,243	27	0.073	0.094
BradleyA-N3	243,591	242,014	1,578	1,013	995	19	0.067	0.112
BradleyA-N4	331,300	329,005	2,295	1,375	1,353	22	0.062	0.108
BradleyA-N5	325,453	322,669	2,784	1,353	1,327	26	0.067	0.093
BradleyA-N6	338,568	336,222	2,346	1,407	1,381	25	0.065	0.106

S2D BW	CSV(MB/s)	CSVRead	CSVWrite	SBL(MB/s)	SBLRead	SBLWrite	Disk(MB/s)	DiskRead	DiskWrite	Cache(MB/s)	CacheRead	CacheWrite
Total	7,567	7,548	19	7,692	7,546	145	5,133	5,067	66	66		66
BradleyA-N1	1,250	1,247	3	1,274	1,248	26	11		11	11		11
BradleyA-N2	1,247	1,243	4	1,270	1,243	27	1,181	1,169	12	12		12
BradleyA-N3	998	995	3	1,013	995	19	1,183	1,172	11	11		11
BradleyA-N4	1,356	1,352	3	1,375	1,353	22	1,373	1,362	11	11		11
BradleyA-N5	1,330	1,327	3	1,353	1,327	26	11		11	11		11
BradleyA-N6	1,386	1,383	3	1,407	1,381	25	1,374	1,363	11	11		11

Block size 4Kb, 8 threads, 1 outstanding I/O (0% read / 100% write)

CSV FS	IOPS	Reads	Writes	BW (MB/s)	Read	Write	Read Lat (ms)	Write Lat
Total	97,646	20	97,626	400		400		
BradleyA-N1	15,286		15,286	63		63	0.000	8.764
BradleyA-N2	17,941		17,941	74		74	0.000	8.402
BradleyA-N3	15,971		15,971	65		65	0.000	8.039
BradleyA-N4	17,161		17,161	70		70	0.000	8.844
BradleyA-N5	16,508		16,508	68		68	0.000	9.257
BradleyA-N6	14,778	20	14,758	60		60	0.002	7.639

SYS	CPU (%)
Total	122
BradleyA-N1	23
BradleyA-N2	20
BradleyA-N3	19
BradleyA-N4	20
BradleyA-N5	20
BradleyA-N6	20

SSB Cache	Hit/Sec	Miss/Sec	Remap/Sec	Cache (MB/s) Read	Write	Destage (MB/s)	Update (MB/s)
Total				999	999	867	57
BradleyA-N1				204	204		10
BradleyA-N2				194	194		9
BradleyA-N3				201	201	6	9
BradleyA-N4				205	205	195	9
BradleyA-N5				196	196		8
BradleyA-N6						667	12

SBL	IOPS	Reads	Writes	BW (MB/s)	Read	Write	Read Lat (ms)	Write Lat
Total	586,258		586,258	1,771		1,771		
BradleyA-N1	91,722		91,722	567		567	0.000	0.122
BradleyA-N2	107,477		107,477				0.000	0.115
BradleyA-N3	95,864		95,864	592		592	0.000	0.138
BradleyA-N4	103,396		103,396				0.000	0.133
BradleyA-N5	99,198		99,198	613		613	0.000	0.114
BradleyA-N6	88,601		88,601				0.000	0.134

S2D BW	CSV(MB/s)	CSVRead	CSVWrite	SBL(MB/s)	SBLRead	SBLWrite	Disk(MB/s)	DiskRead	DiskWrite	Cache(MB/s)	CacheRead	CacheWrite
Total	400		400	1,771		1,771	999		999	999		999
BradleyA-N1	63		63	567		567	204		204	204		204
BradleyA-N2	74		74				194		194	194		194
BradleyA-N3	65		65	592		592	201		201	201		201
BradleyA-N4	70		70				205		205	205		205
BradleyA-N5	68		68	613		613	196		196	196		196
BradleyA-N6	60		60									

Block size 4Kb, 8 threads, 1 outstanding I/O (70% read / 30% write)

CSV FS	IOPS	Reads	Writes	BW (MB/s)	Read	Write	Read Lat (ms)	Write Lat
Total	315,303	220,979	94,324	1,317	915	402		
BradleyA-N1	49,453	34,602	14,850	207	143	64	0.130	9.771
BradleyA-N2	52,420	36,688	15,732	219	151	67	0.133	8.685
BradleyA-N3	54,168	38,233	15,936	229	161	67	0.133	7.995
BradleyA-N4	54,493	38,141	16,352	226	157	69	0.130	8.544
BradleyA-N5	54,490	38,143	16,347	226	157	69	0.132	8.145
BradleyA-N6	50,280	35,173	15,107	210	146	64	0.128	8.846

SYS	CPU (%)
Total	152
BradleyA-N1	26
BradleyA-N2	24
BradleyA-N3	25
BradleyA-N4	24
BradleyA-N5	26
BradleyA-N6	26

SSB Cache	Hit/Sec	Miss/Sec	Remap/Sec	Cache (MB/s) Read	Write	Destage (MB/s)	Update (MB/s)
Total				1,218	1,218		53
BradleyA-N1				206	206		9
BradleyA-N2				197	197		8
BradleyA-N3				203	203		9
BradleyA-N4				207	207		8
BradleyA-N5				197	197		8
BradleyA-N6				209	209		11

SBL	IOPS	Reads	Writes	BW (MB/s)	Read	Write	Read Lat (ms)	Write Lat
Total	788,483	220,954	567,529	2,658	914	1,744		
BradleyA-N1	124,147	34,595	89,551	708	143	566	0.098	0.110
BradleyA-N2	130,842	36,738	94,104	152	152		0.099	0.108
BradleyA-N3	134,302	38,243	96,059	763	161	602	0.097	0.130
BradleyA-N4	135,695	38,095	97,600	157	157		0.092	0.127
BradleyA-N5	135,880	38,143	97,737	157	157		0.097	0.110
BradleyA-N6	127,617	35,139	92,478	721	145	576	0.094	0.123

S2D BW	CSV(MB/s)	CSVRead	CSVWrite	SBL(MB/s)	SBLRead	SBLWrite	Disk(MB/s)	DiskRead	DiskWrite	Cache(MB/s)	CacheRead	CacheWrite
Total	1,317	915	402	2,658	914	1,744	1,976	758	1,218	1,218		1,218
BradleyA-N1	207	143	64	708	143	566	348	142	206	206		206
BradleyA-N2	219	151	67	152	152		339	142	197	197		197
BradleyA-N3	229	161	67	763	161	602	203		203	203		203
BradleyA-N4	226	157	69	157	157		369	162	207	207		207
BradleyA-N5	226	157	69	157	157		353	156	197	197		197
BradleyA-N6	210	146	64	721	145	576	364	156	209	209		209

Sequential Reads and Writes

Block size 4Kb, 1 thread, 4 outstanding I/O (100% read / 0% write)

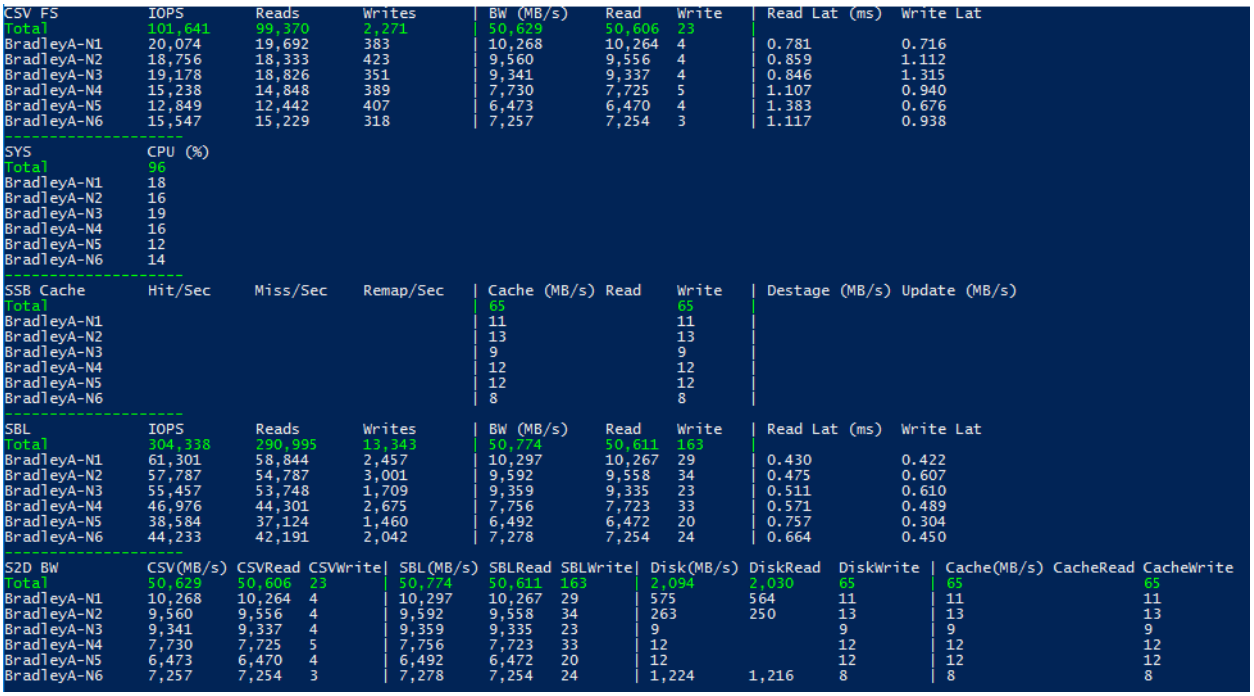

CSV FS	IOPS	Reads	Writes	BW (MB/s)	Read	Write	Read Lat (ms)	Write Lat
Total	101,641	99,370	2,271	50,629	50,606	23		
BradleyA-N1	20,074	19,692	383	10,268	10,264	4	0.781	0.716
BradleyA-N2	18,756	18,333	423	9,560	9,556	4	0.859	1.112
BradleyA-N3	19,178	18,826	351	9,341	9,337	4	0.846	1.315
BradleyA-N4	15,238	14,848	389	7,730	7,725	5	1.107	0.940
BradleyA-N5	12,849	12,442	407	6,473	6,470	4	1.383	0.676
BradleyA-N6	15,547	15,229	318	7,257	7,254	3	1.117	0.938

SYS	CPU (%)
Total	96
BradleyA-N1	18
BradleyA-N2	16
BradleyA-N3	19
BradleyA-N4	16
BradleyA-N5	12
BradleyA-N6	14

SSB Cache	Hit/Sec	Miss/Sec	Remap/Sec	Cache (MB/s) Read	Write	Destage (MB/s)	Update (MB/s)
Total				65	65		
BradleyA-N1				11	11		
BradleyA-N2				13	13		
BradleyA-N3				9	9		
BradleyA-N4				12	12		
BradleyA-N5				12	12		
BradleyA-N6				8	8		

SBL	IOPS	Reads	Writes	BW (MB/s)	Read	Write	Read Lat (ms)	Write Lat
Total	304,338	290,995	13,343	50,774	50,611	163		
BradleyA-N1	61,301	58,844	2,457	10,297	10,267	29	0.430	0.422
BradleyA-N2	57,787	54,787	3,001	9,592	9,558	34	0.475	0.607
BradleyA-N3	55,457	53,748	1,709	9,359	9,335	23	0.511	0.610
BradleyA-N4	46,976	44,301	2,675	7,756	7,723	33	0.571	0.489
BradleyA-N5	38,584	37,124	1,460	6,492	6,472	20	0.757	0.304
BradleyA-N6	44,233	42,191	2,042	7,278	7,254	24	0.664	0.450

S2D BW	CSV(MB/s)	CSVRead	CSVWrite	SBL(MB/s)	SBLRead	SBLWrite	Disk(MB/s)	DiskRead	DiskWrite	Cache(MB/s)	CacheRead	CacheWrite
Total	50,629	50,606	23	50,774	50,611	163	2,094	2,030	65	65		65
BradleyA-N1	10,268	10,264	4	10,297	10,267	29	575	564	11	11		11
BradleyA-N2	9,560	9,556	4	9,592	9,558	34	263	250	13	13		13
BradleyA-N3	9,341	9,337	4	9,359	9,335	23	9		9	9		9
BradleyA-N4	7,730	7,725	5	7,756	7,723	33	12		12	12		12
BradleyA-N5	6,473	6,470	4	6,492	6,472	20	12		12	12		12
BradleyA-N6	7,257	7,254	3	7,278	7,254	24	1,224	1,216	8	8		8

Block size 4Kb, 1 threads, 4 outstanding I/O (0% read / 100% write)

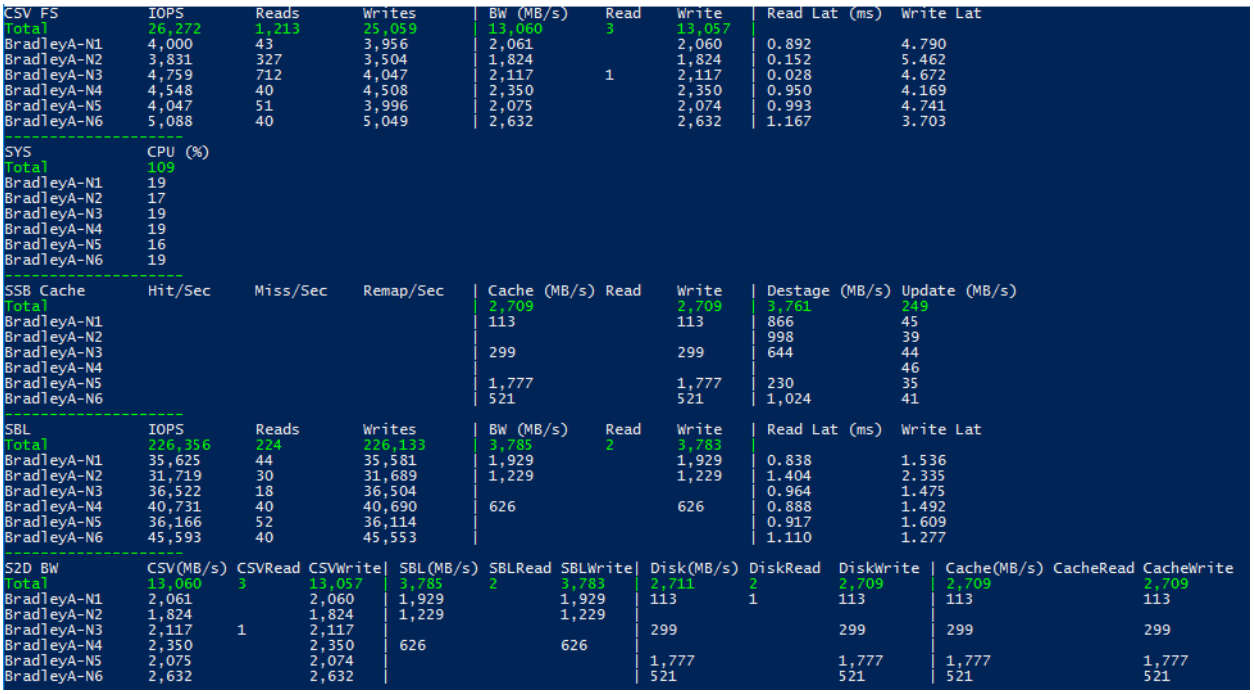

CSV FS	IOPS	Reads	Writes	BW (MB/s)	Read	Write	Read Lat (ms)	Write Lat
Total	26,272	1,213	25,059	13,060	3	13,057		
BradleyA-N1	4,000	43	3,956	2,061		2,060	0.892	4.790
BradleyA-N2	3,831	327	3,504	1,824		1,824	0.152	5.462
BradleyA-N3	4,759	712	4,047	2,117	1	2,117	0.028	4.672
BradleyA-N4	4,548	40	4,508	2,350		2,350	0.950	4.169
BradleyA-N5	4,047	51	3,996	2,075		2,074	0.993	4.741
BradleyA-N6	5,088	40	5,049	2,632		2,632	1.167	3.703

SYS	CPU (%)
Total	109
BradleyA-N1	19
BradleyA-N2	17
BradleyA-N3	19
BradleyA-N4	19
BradleyA-N5	16
BradleyA-N6	19

SSB Cache	Hit/Sec	Miss/Sec	Remap/Sec	Cache (MB/s) Read	Write	Destage (MB/s)	Update (MB/s)
Total				2,709	2,709	3,761	249
BradleyA-N1				113	113	866	45
BradleyA-N2						998	39
BradleyA-N3				299	299	644	44
BradleyA-N4							46
BradleyA-N5				1,777	1,777	230	35
BradleyA-N6				521	521	1,024	41

SBL	IOPS	Reads	Writes	BW (MB/s)	Read	Write	Read Lat (ms)	Write Lat
Total	226,356	224	226,133	3,785	2	3,783		
BradleyA-N1	35,625	44	35,581	1,929		1,929	0.838	1.536
BradleyA-N2	31,719	30	31,689	1,229		1,229	1.404	2.335
BradleyA-N3	36,522	18	36,504				0.964	1.475
BradleyA-N4	40,731	40	40,690	626		626	0.888	1.492
BradleyA-N5	36,166	52	36,114				0.917	1.609
BradleyA-N6	45,593	40	45,553				1.110	1.277

S2D BW	CSV(MB/s)	CSVRead	CSVWrite	SBL(MB/s)	SBLRead	SBLWrite	Disk(MB/s)	DiskRead	DiskWrite	Cache(MB/s)	CacheRead	CacheWrite
Total	13,060	3	13,057	3,785	2	3,783	2,711	2	2,709	2,709		2,709
BradleyA-N1	2,061		2,060	1,929		1,929	113	1	113	113		113
BradleyA-N2	1,824		1,824	1,229		1,229						
BradleyA-N3	2,117	1	2,117				299		299	299		299
BradleyA-N4	2,350		2,350	626		626						
BradleyA-N5	2,075		2,074				1,777		1,777	1,777		1,777
BradleyA-N6	2,632		2,632				521		521	521		521

Summary

DataON's S2D-5224i HCI solution with Intel® Xeon® processors, Intel® Optane™ technology and Mellanox 100GbE network infrastructure with Windows Server 2016 Storage Spaces Direct turned out to be the ideal solution to replace a traditional SAN at Bradley. It was able to greatly improve performance and reduce latency, while eliminating expensive SAN maintenance and virtual machine annual licensing fees. In addition, the new Storage Spaces Direct solution provided them with innovative features that sets them up to utilized hybrid cloud in the future.

The DataON all-flash platform, coupled with a low latency cache tier, now provides performance in the millions of IOPS. By including a 100GbE network switch in the configuration, Bradley is now well positioned to handle the firm's future bandwidth needs.

Trademarks

Appendix

System Specifications

- DataON S2D-5224i Hyper-Converged Infrastructure
    - Intel® Xeon® Scalable™ Gold 6132 2.6 GHz, 14-Core (two CPUs per node)
    - Mellanox® ConnectX-4 EN Single Port QSFP28 40/56 GbE RDMA Cards
    - Intel® S3520 240GB SATA M.2 2280 3D1 MLC Boot SSDs
    - Intel® Optane™ DC P4800X 375GB NVMe U.2 SSDs, 3D XPoint™, Read/Write Cache Tier (four per node)
- Intel® DC S4500™ 6G SATA 3.8TB 2.5" 3D1 TLC SSDs
- Mellanox Spectrum™ based 100GbE**,** 1U Open Ethernet Switch with MLNX-OS, 32 QSFP28 ports

Collateral & Resources

Collateral

Microsoft Windows Server 2016 datasheet

DataON S2D-5000 Family datasheet

DataON MUST datasheet

Intel® Optane™ datasheet

Intel® Optane™ DC P4800X SSDs datasheet

Videos

Storage Spaces Direct in Windows Server 2016 presentation

DataON MUST demo video