



Storage at the Center: Giving Hyperconvergence a Boost

Cisco HyperFlex* All NVMe* reaches new performance and efficiency highs with optimized support for Intel® Optane™ DC SSDs and Intel® 3D NAND SSDs

This solution brief describes how to solve business challenges through investment in innovative technologies.

If you are responsible for...

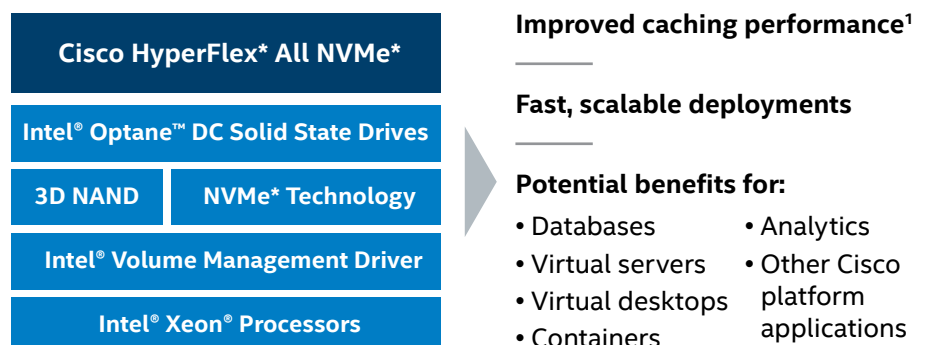
- **Business strategy:**
You will better understand how Cisco's HyperFlex* All NVMe* platform will enable you to successfully meet your business outcomes.
- **Technology decisions:**
You will learn how Cisco's HyperFlex* All NVMe* platform works to deliver IT and business value.



Executive Summary

The rising tide of virtualization helped to enable the converged architecture market, but that market remains under pressure from mounting performance and cost-efficiency demands. Specifically, the applications that enterprises are using to drive their business in such areas as artificial intelligence (AI), database, virtual desktop, and other mission-critical systems increasingly require more performance than what conventional hyperconverged infrastructure (HCI) solutions can provide. Storage platforms sourced from multiple vendors can produce bottlenecks and introduce costly complexity, especially when those platforms are software-only solutions.

Cisco's HyperFlex* infrastructure represents one of the latest evolutions in server-based efficiency for a wide range of applications, and it was designed for a data-driven landscape in which traditional data centers, emerging edge workloads, and cloud platforms must all cooperate seamlessly without impeding the others' performance. Cisco's new HyperFlex All NVMe* platform takes advantage of key features of Intel® Optane™ DC SSDs, Intel® Xeon® processors, and Intel® 3D NAND SSDs to deliver outstanding performance. HyperFlex support for Intel Optane DC SSDs complements an all-NVMe configuration, delivering significant improvements in caching performance¹ that will enhance a range of database, virtual server, virtual desktop, container, analytics, and other platform applications that enable businesses to help achieve greater cost and performance efficiencies² and create competitive difference through fast storage I/O, easy platform management, and scalability.



Author

Alessandro Goncalves
SSD Architect, Intel Corporation

Figure 1. Intel® architecture-based enhancements to Cisco HyperFlex* All NVMe* platform address many of the pain points and application needs of today's legacy-bound enterprise IT efforts.

Solution Benefits

- **Fast storage I/O.** Optimized for Intel® technologies—including Intel® Optane™ DC SSD and 3D NAND storage, as well as a strong emphasis on the network fabric—Cisco's HyperFlex* All NVMe* enables storage traffic to be placed on a lateral highway spanning the IT infrastructure rather than confining it within vertical silos.
- **Easy platform management.** Clusters can be spun up, self-configured, and ready to accept applications and workloads in under one hour.² Moreover, hyperconverged infrastructure (HCI) should make resource management easier for IT, but software-only approaches to HCI introduce complexity, stifle innovation, and boost overall operational costs. Similarly, multi-sourcing HCI components increases project risk and promotes cross-vendor blaming when service is required. The Cisco HyperFlex platform is a hardware-based, single-sourced solution that can remedy all these issues.
- **Scalability.** Moving to an all-NVME configuration helps HyperFlex to be "extremely well-suited to tier-1 production workloads."³ In part, this is because scale-out architecture can help every added storage node bolster performance and capacity for the entire cluster. Simultaneously, the HyperFlex All NVMe solution can lower solution costs. In the many cases where HyperFlex delivers higher performance, this translates into higher workload density, which reduces the need for additional nodes and can thereby lower total solution investment.

Business Challenge: Traditional Storage Can't Keep Today's Pace

Hyperconverged infrastructure (HCI) technologies help enable enterprises to eliminate resource silos. The ability to consolidate compute, storage, networking, and other IT resources under a single umbrella architecture delivers hardware cost savings, easier management, and greater scalability, particularly through virtualization. Data silos stem from infrastructure solutions that often inhibit scaling and isolate resource groups due to incompatibility issues.

Siloing also makes it more time consuming for IT to provision and maintain non-converged, legacy systems. Additionally, silos make aggregating data for analysis cumbersome. If performance benefits to maintaining legacy-rooted silos existed, that might justify perpetuating them, but such benefits remain scarce at best. Consider SATA storage technology as an example of one such legacy technology. With its long history and consumer-centric focus, SATA has the volume and market targeting to remain inexpensive, but SATA's real-world performance ceiling stands at 560 MB/s for sequential reads.⁴ This is trivial compared to NVM Express* (NVMe*) technology, which can quadruple SATA's top rate.⁵ Much of this bottleneck stems from SATA's inability to leverage the PCI Express* (PCIe*) bus's direct CPU connection and pipeline bandwidth — a

shortcoming that NVMe does not share. SATA is a 1:1 architecture, with each drive requiring its own controller port, which proves its limitations in terms of agility and scalability in enterprise environments. The PCIe bus can serve multiple end points simultaneously, and with much higher bandwidth, thus enabling far higher total performance with lower use of system resources in a hyperconverged, NVMe-based storage subsystem. The throughput bottlenecks caused by lower SATA bandwidth can propagate siloing. It can be tempting to maintain SATA-based legacy platforms due to their familiarity and sunk costs, but more evolved solutions, such as the Cisco HyperFlex* All NVMe platform, can greatly improve data processing performance and break down those silo walls.

Breaking down silos and improving solution performance can simultaneously improve operational efficiency, resulting in reduced overall ownership costs. Part of the business challenge lies in understanding the nature of present architecture bottlenecks and then helping to alleviate them through the adoption of new technologies such as Intel® Optane™ DC SSDs and other NVMe-based storage options.

Use Cases: Giving Major Enterprise Apps Free Rein

Cisco HyperFlex All NVMe with Intel Optane DC SSDs and Intel® 3D NAND SSDs offers an advantageous mix of performance-enhancing hardware, resource efficiency, and scalability. HyperFlex melds high-performance software-defined storage of the HyperFlex data platform with the Cisco Unified Computing System (UCS) platform, enabling the deployment of enterprise applications and application containers at a wide variety of scales within a spectrum of cloud models. This is no small order, and the advances enabled by Intel® technologies significantly contribute in helping to eliminate processing bottlenecks from several key use models, including:

- **Database access.** HyperFlex All NVMe serves equally well with traditional databases (i.e., SQL*, MySQL*, Oracle) and non-tabular alternatives, such as NoSQL*. Database applications rely on performance and scalability, both of which play to HyperFlex All NVMe's innovations.
- **Virtual desktop infrastructure (VDI).** VDI platforms, such as XenDesktop* and Horizon* 7, have remarkable potential for reducing unnecessary client costs and improving IT manageability. However, VDI deployment on traditional servers can be slow and complex. Ineffective scaling and spotty storage performance across network resources often plague VDI's efficacy. Evolved platforms, such as HyperFlex All NVMe, let IT allocate resources to accommodate VDI while optimizing capacity, simplifying deployment, and providing an easily scalable, centralized pool for compute, graphics, and storage resources.
- **Containers.** With the demand for faster deployment and higher agility, containers have stormed the enterprise world. The storage bandwidth supplied by Intel technologies coupled with the virtualization-optimized functionality of Intel® Xeon® Scalable processors allows customers to effectively run Cisco's Kubernetes* platform, enabling containers-as-a-service for multi-cloud environments.

- **Machine learning/artificial intelligence (AI).** The inferencing operations that drive machine learning and AI involve copious amounts of sequential calculations. This requires ample memory bandwidth and processing performance. “Training” processes for machine learning run volumes of parallel operations, requiring highly parallelized compute resources.⁶ HyperFlex All NVMe, bolstered by the latest Intel® storage technologies, excels at scaling such resource bandwidth.
- **Mission-critical applications.** The potential failure of an application can pose a serious or existential risk to the enterprise, making the ability to rapidly route new resources into action crucial. A properly provisioned HCI platform makes remediation nearly seamless and aids in proactively preventing such failures. HyperFlex addresses this through logical availability zones, which mitigate the risk of increases in maximum cluster size that may lead to cluster shut-down. Logical availability zones partition the cluster into multiple logical zones, and HyperFlex helps to ensure there is never more than one data copy in every zone. While containing excessive data redundancy, this also reduces the risk of multiple failures in a single zone, which may result in cluster failure.

Solution Value: HyperFlex* Efficiency Excels With Intel® Technology

As noted previously, HCI can address many IT needs. According to research from Enterprise Strategy Group (ESG), those interested in adopting HCI do so for improved scalability (31 percent), total cost of ownership (28 percent), ease of deployment (26 percent), and simplified systems management (24 percent).⁷ Cisco’s HyperFlex platform has been named a leader in multiple prominent analyst reports centered on the industry’s HCI solutions, and the new optimizations for Intel storage technologies in HyperFlex All NVMe solutions further extend HyperFlex’s benefits. By designing HyperFlex as a fully integrated solution that weaves together hardware, software, and networking into one product rather than a collection of multi-sourced pieces, total solution complexity is reduced. This results in less deployment intricacy, less management, and higher performance.

HyperFlex support for All NVMe with Intel Optane DC SSD configurations also plays a central role in performance. Figure 2 illustrates ESG’s findings⁷ in comparing two HyperFlex clusters with similar CPU and node count configurations:

- The first all-flash cluster used Cisco HyperFlex 2.6 with one 800 GB SAS SSD for caching and six 960 GB enterprise SATA SSDs for storage.
- The second cluster used HyperFlex All NVMe configuration with one 375 GB Intel® Optane™ SSD DC P4800X for caching and six 1 TB NVMe Intel® SSD DC P4500 for storage.

Tests included OLTP workloads for SQL Server*, virtual application server, and desktop activity. The results show a 64 percent average performance benefit from adopting an all-NVMe storage architecture under a platform that has been optimized to leverage Intel storage technologies (see Figure 2). The HyperFlex All NVMe configuration also showed a 35.5 percent average drop in latency.⁸

Moreover, ESG ran a different set of HCI bench tests that employed two workloads, one with a read/write ratio of 70/30 and another with a read/write ratio of 50/50. Each node in the four-node cluster ran 35 VMs. All told, the HyperFlex All NVMe cluster realized 63 percent greater IOPS and 37 percent improved latency in the 70/30 workload tests.⁹ Perhaps most interesting of all, ESG noted that the Intel-based configurations “showed little variation across all 140 VMs — aggregate testing IOPS stayed close to the targets of 600 for all-flash and 900 for all-NVMe. Software-only HCI vendors previously tested by ESG showed a wide variation, with much higher standard deviation.”¹⁰ Performance consistency is often key to the success of high-load, virtualized platforms.

In addition to its performance benefits, an all-NVMe platform such as HyperFlex can also provide:

- Fast deployment — HyperFlex All NVMe clusters can be installed in under one hour, saving on deployment costs²
- Easy cluster management
- Inherently high shareability
- Massive scalability
- High VM density at an attractive price
- High reliability, particularly with the high endurance of Intel Optane DC SSDs
- High overall value, especially when resource consolidation and productivity gains are considered
- Excellent levels of platform reliability as a result of collaboration between Cisco and Intel

Last but far from least, Cisco HyperFlex All NVMe can help organizations move closer to eliminating their data silos.

Cisco HyperFlex* Platform-Enabled All-NVMe* Cluster Performance

Vdbench SQL Server* Curve Test, 140 VMs, 70/30 Mixed Workload

- HyperFlex All Flash 2.6 (SAS, SATA)
- HyperFlex All NVMe with Intel® Optane™ SSD DC P4800X

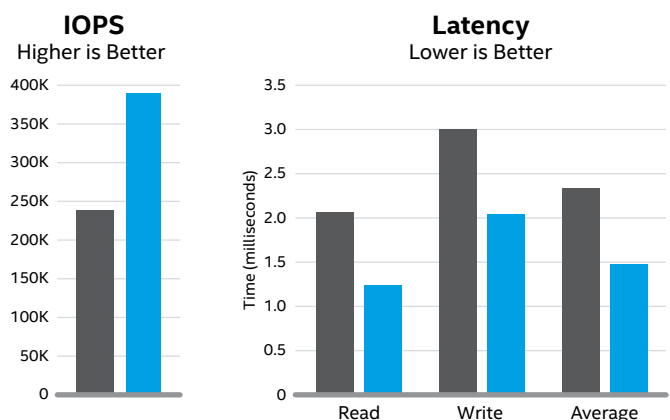


Figure 2. All-NVMe* cluster performance under the Vdbench SQL Server* Curve Test shows a clear advantage⁸ for Cisco’s HyperFlex* platform, especially with write latency.

With operating environments now dispersed across localized data centers, edge networks, and cloud platforms, the likelihood of data sets becoming isolated increases. HCIs can seamlessly bridge these data congregation zones, but not all HCI platforms function equally. Performance limitations can constrict the ability to run all applications on the HCI, which then leaves the applications and their attendant data to run in off-platform, tiered solutions. This separation silos data and impairs what gains might be made by the HCI. HyperFlex All NVMe provides the performance necessary to run these applications, preserve the organization's HCI strategy, and keep those silos from reforming.

Solution Architecture: Evolving Hyperconverged Infrastructure

Among the many innovations in HyperFlex, this iteration marks the first time that Cisco has adopted an all-NVMe node capability. HyperFlex All NVMe offers performance advantages over legacy Advanced Host Controller Interface (AHCI)-based models. The inherent limitations of AHCI are due to its decades-old technology designed for legacy hard drive systems. AHCI incorporates protocol layers and overhead processing that are no longer needed. SATA SSDs require a storage controller to act as the interpretive bridge between the CPU and storage. PCIe-based storage systems, including NVMe, dispense with that controller, saving formidable numbers of instruction cycles. The more SATA bandwidth there is in an HCI implementation, the more legacy bandwidth bottlenecks hinder total platform performance (see Figure 3).

Despite the benefits of running NVMe over PCIe, one key omission in its specification had the potential to hamper HyperFlex's utility for enterprises. Currently, NVMe does not provide native support for drive hot swapping – an essential need in data centers, where there's no time or budget for downtime in drive servicing. To rectify this, Intel worked with Cisco to integrate the Intel® Volume Management Device (Intel® VMD) capability into HyperFlex that allows NVMe to hot swap.

Hot-swapping capability is part of a broader discussion about reliability, availability, and serviceability (RAS) that dominates enterprise storage concerns. In SATA platforms, the storage controller provides RAS functionality, but this is absent in NVMe implementations because NVMe eliminates the storage controller. Thus, the sudden physical removal of an NVMe drive may crash the system, and the NVMe protocol lacks a hotplug specification to prevent this. Intel teams joined with Cisco to craft enhancements within the Intel Xeon processor's PCIe root complex that links to an OS driver. With these Intel VMD advances (see Figure 4), HyperFlex helps remedy NVMe hotplug and drives removal needs to improve manageability. Intel VMD also allows HyperFlex to supply a software interface for the set up and control of storage domains, including enumeration, event/error handling, LED configuration, firmware management, and more.




Figure 3. This comparison shows the architectural differences in a virtualized storage platform based on SATA and PCIe*/NVMe* resources. Link bandwidth makes a significant difference in performance.

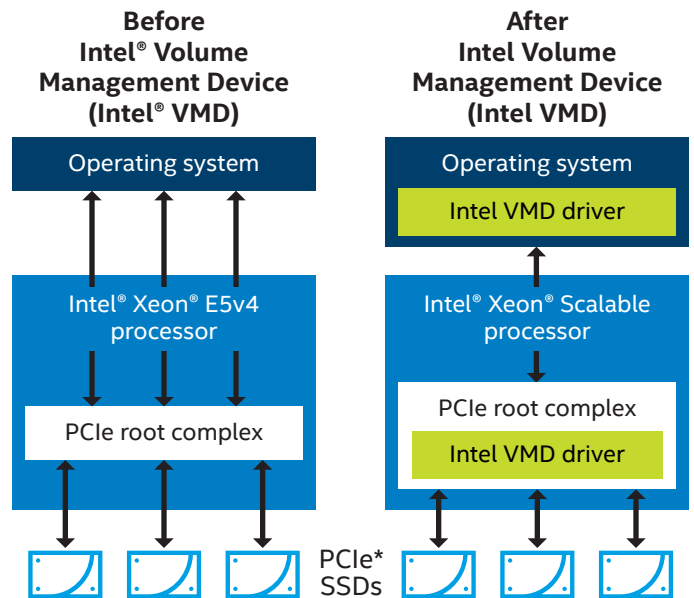


Figure 4. Intel worked with Cisco to develop its Intel® Volume Management Device (Intel® VMD) for Cisco HyperFlex* All NVMe*. This required enhancements to select Intel® Xeon® processor models as well as a novel, validated driver for supported operating systems.

The Cisco HyperFlex All NVMe solution contains the following key ingredients:

- Cisco HX220c* All NVMe M5 node
- Intel Optane DC SSDs
- Intel 3D NAND SSDs
- Intel Xeon Scalable processor-based servers
- Intel VMD Drivers

Collectively, these ingredients are designed to reach the optimal balance of areal storage density and processor performance. They reflect the need to tackle sharply rising data volumes in the smallest possible footprint, emphasizing low latency, reduced processor request bottlenecks, and higher per-server work output. Cisco's incorporation of these ingredients eliminates many of the integration and deployment headaches that so often plague multi-vendor solutions.

Note that HyperFlex contains deep optimizations for Intel Optane DC SSDs. In the context of HCI, Intel Optane DC SSDs can be deployed at the caching storage tier. In the prior section, we discussed Intel Optane DC SSD performance improvements, including its endurance, which

stems from the properties of its Intel® Optane™ memory media. Compared to the Intel® SSD DC P4600, the Intel Optane SSD DC P4800X specifies 60 drive writes per day rather than three – a huge benefit in high-volume caching applications – and ultimately yields a cost per terabyte written of \$0.026 versus \$0.168.¹¹

In practice, HyperFlex All NVMe excels when placing Intel Optane DC SSDs in the caching tier and Intel 3D NAND SSDs in the capacity tier. This is an evolutionary step from traditional HCI architectures, which often place NAND SSDs in the caching role and hard disks in capacity. Intel Optane DC SSDs provide high performance, and NAND supplies high density for low cost.¹¹ Putting Intel Optane DC SSDs in a caching role leverages the exceptionally low-latency characteristics of Intel Optane memory media. This increased cache speed will facilitate more data requests in a shorter time, alleviating the need to maintain as much mass storage and further lowering total platform costs.

Also, Intel 3D NAND-based storage is ideally suited to HyperFlex's capacity role. A complete overview of how these storage options fit within HyperFlex is shown in Figure 5.

Cisco HyperFlex* Platform Solution Architecture

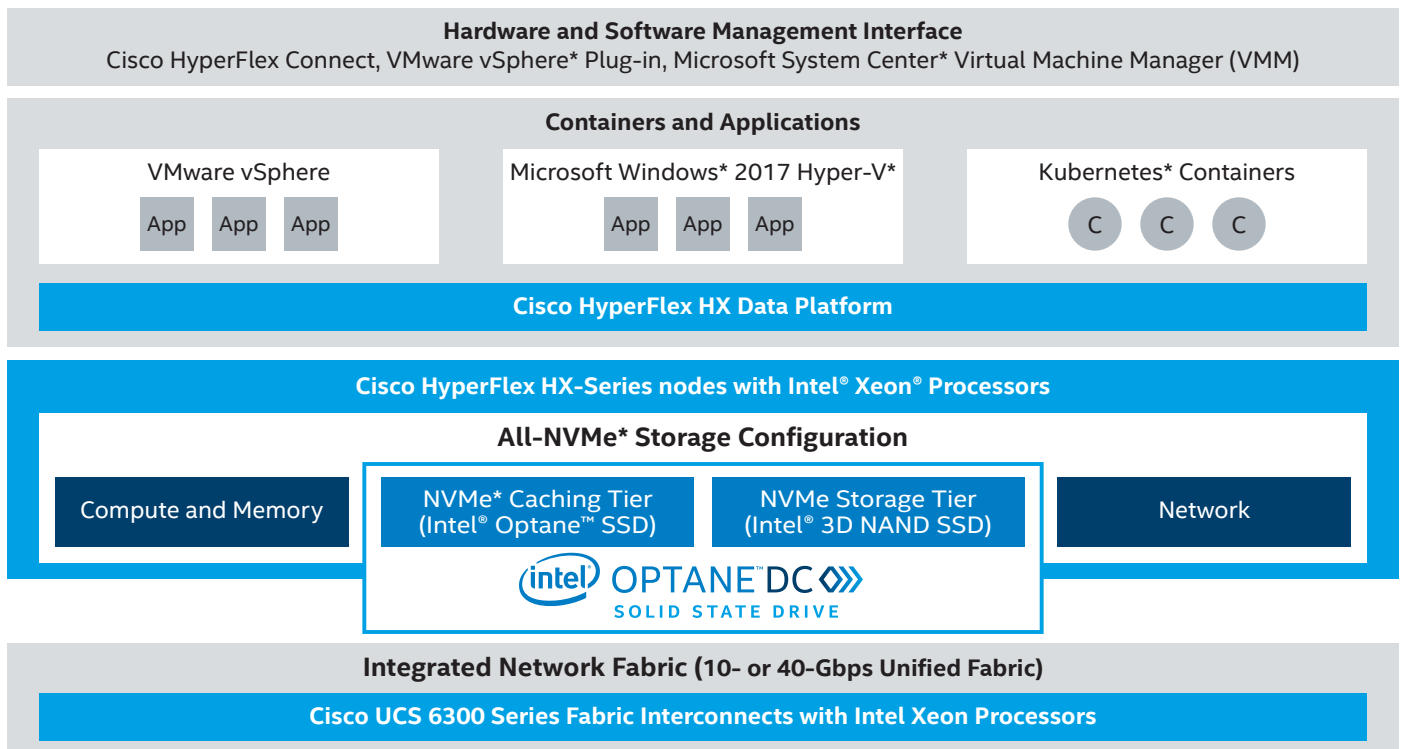


Figure 5. This illustration of Cisco's overall HyperFlex* platform shows the interplay between storage resources, virtualized systems, and network fabric. Storage and compute performance are pivotal in the HyperFlex ecosystem.

Conclusion

The demands on data centers continue to climb. HCI has proven its ability to help businesses scale their large computing efforts. Advances and cumulative benefits in Intel Optane DC SSDs, Intel Xeon processors, and Intel 3D NAND SSDs usher in a new wave of performance and benefit possibilities for forward-looking HCI deployments, particularly when integrated with the bandwidth-expanding benefits of a storage platform such as Cisco HyperFlex All NVMe. Cisco has worked closely with Intel to enable these processor and storage innovations in its new solution. The platform enables significant improvements in storage performance and latency compared to previous, non-all-NVMe platforms while simultaneously enhancing platform RAS. Beyond fast storage I/O, HyperFlex also delivers easy platform management and high scalability. Fueled by these advances, Cisco can now deliver a HyperFlex solution to users with faster data handling, easier platform management, and greater scalability across demanding workloads, including AI, machine learning, and mission-critical applications.

Find the right solution for your organization. Contact your Intel representative or visit cisco.com/c/en/us/products/hyperconverged-infrastructure/hyperflex-solutions.html.

Learn More

You may also find the following resources useful:

- **Cisco HyperFlex Solutions**
cisco.com/c/en/us/products/hyperconverged-infrastructure/hyperflex-solutions
- **ESG Lab Validation of Cisco HyperFlex**
esg-global.com/cisco-hyperflex
- **HyperFlex and SQL Server**
cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/hyperflex-le-64801.pdf
- **HyperFlex and Accelerating VDI Deployment**
cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/le-hyperflex-system.pdf
- **Cisco Container Platform**
cisco.com/c/en/us/products/cloud-systems-management/container-platform
- **Cisco At a Glance: HyperFlex All NVMe**
cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/le-69503-aag-all-nvme.pdf
- **Intel Select Solutions for Hyper-Convergence via VMware vSAN**
intel.com/content/dam/www/public/us/en/documents/solution-briefs/select-vsan-solution-brief.pdf

Solution Provided By:

- ¹ "Mission-critical Hyperconverged Workload Performance Testing on Cisco HyperFlex All NVMe with Intel Optane DC SSD," April 2019, pp. 9-12 under ESG Testing; [cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/dc-ssd-esg.pdf](https://www.cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/dc-ssd-esg.pdf)
- ² "Mission-critical Hyperconverged Workload Performance Testing on Cisco HyperFlex All NVMe with Intel Optane DC SSD," April 2019, p. 5 under Cisco's Fully Engineered HCI Approach; [cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/dc-ssd-esg.pdf](https://www.cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/dc-ssd-esg.pdf)
- ³ "Mission-critical Hyperconverged Workload Performance Testing on Cisco HyperFlex All NVMe with Intel Optane DC SSD," April 2019, p. 14 under The Bigger Truth; [cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/dc-ssd-esg.pdf](https://www.cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/dc-ssd-esg.pdf)
- ⁴ For an example of a drive that hits the ceiling of SATA 3.0's 6 Gb/s bandwidth, see the Intel® SSD D3-S4510 Series, ark.intel.com/content/www/us/en/ark/products/134912/intel-ssd-d3-s4510-series-960gb-2-5in-sata-6gb-s-3d2-tlc.html
- ⁵ Intel® Optane™ SSDs showcase NVMe's bandwidth capabilities, as with the Intel Optane SSD 905P Series, ark.intel.com/content/www/us/en/ark/products/129834/intel-optane-ssd-905p-series-960gb-1-2-height-pcie-x4-20nm-3d-xpoint.html
- ⁶ "Technology Requirements for Deep and Machine Learning," nextplatform.com/2017/07/14/technology-requirements-deep-machine-learning
- ⁷ "Mission-critical Hyperconverged Workload Performance Testing on Cisco HyperFlex All NVMe with Intel Optane DC SSD," April 2019, p. 3 under Introduction; [cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/dc-ssd-esg.pdf](https://www.cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/dc-ssd-esg.pdf)
- ⁸ "Mission-critical Hyperconverged Workload Performance Testing on Cisco HyperFlex All NVMe with Intel Optane DC SSD," April 2019, pp. 9-11; [cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/dc-ssd-esg.pdf](https://www.cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/dc-ssd-esg.pdf)
- ⁹ "Mission-critical Hyperconverged Workload Performance Testing on Cisco HyperFlex All NVMe with Intel Optane DC SSD," April 2019, pp. 11-12 under ESG Testing; [cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/dc-ssd-esg.pdf](https://www.cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/dc-ssd-esg.pdf)
- ¹⁰ "Mission-critical Hyperconverged Workload Performance Testing on Cisco HyperFlex All NVMe with Intel Optane DC SSD," April 2019, p. 12 under ESG Testing; [cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/dc-ssd-esg.pdf](https://www.cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/dc-ssd-esg.pdf)
- ¹¹ Based on pricing from Newegg.com fetched in early March 2019.

All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software, or service activation. Performance varies depending on system configuration. Check with your system manufacturer or retailer, or learn more at [intel.com](https://www.intel.com).

Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information, visit www.intel.com/benchmarks.

Performance results are based on testing as of February 2019 and April 2019 and may not reflect all publicly available security updates. All testing was conducted by The Enterprise Strategy Group, Inc. For details, see [cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/dc-ssd-esg.pdf](https://www.cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/dc-ssd-esg.pdf). No product or component can be absolutely secure.

Intel, the Intel logo, Optane, and Xeon are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others. © Intel Corporation 0619/JGAL/KC/PDF 337467-001US