

ESG Lab Validation

Cisco HyperFlex

Hyperconverged Infrastructure with Consistent High Performance for Virtual Machines

By Tony Palmer and Kerry Dolan, Senior Lab Analysts

March 2017

This ESG Lab Report was commissioned by Cisco and is distributed under license from ESG.

Contents

Introduction.....	3
Background.....	3
Cisco HyperFlex.....	4
ESG Lab Validation.....	6
Hybrid Testing.....	6
All-flash Testing.....	8
ESG Lab Validation Highlights.....	11
Issues to Consider.....	11
The Bigger Truth.....	12

ESG Lab Reports

The goal of ESG Lab reports is to educate IT professionals about data center technology products for companies of all types and sizes. ESG Lab reports are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objective is to go over some of the more valuable feature/functions of products, show how they can be used to solve real customer problems and identify any areas needing improvement. ESG Lab's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments.

Introduction

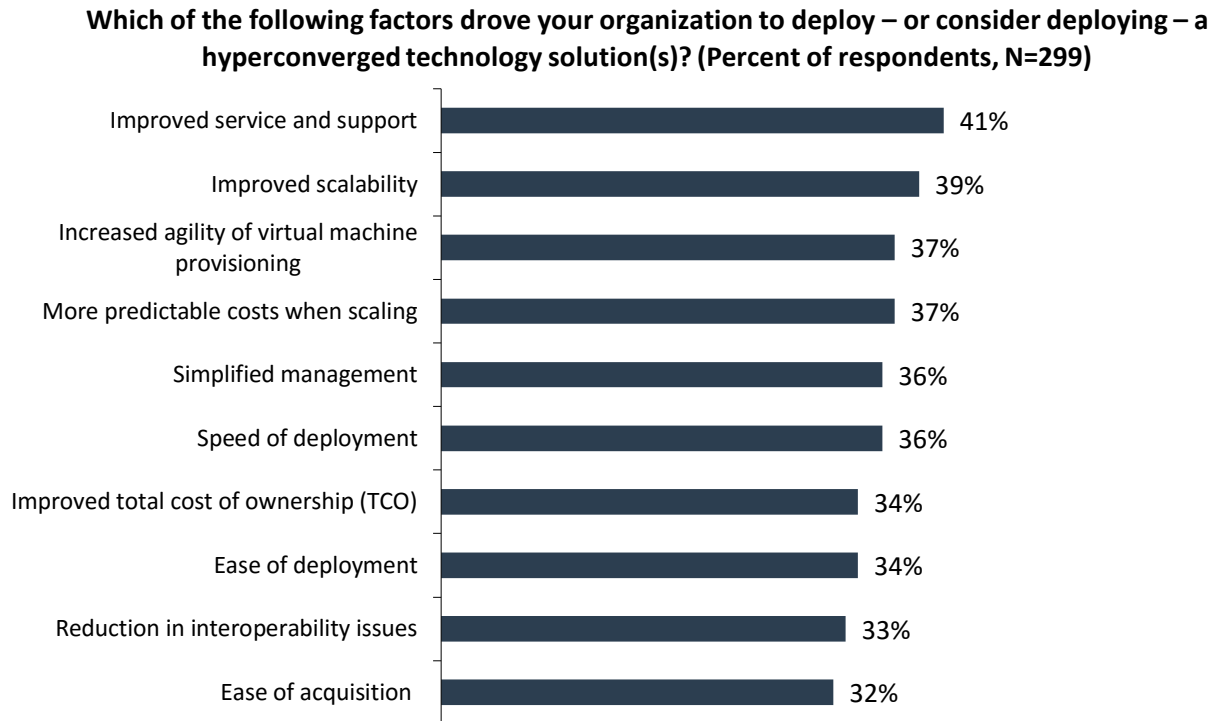
This report documents an ESG Lab audit and validation of Cisco HyperFlex hyperconverged infrastructure performance testing, which focused on comparisons of Cisco HyperFlex hybrid and all-flash solutions with anonymous competitive HCI solutions.

Background

Organizations today must be extremely flexible, with the ability to add applications and virtual machines (VMs) quickly to handle the speed of business. This is extremely difficult to achieve with silos of compute, network, and storage gear that are static and require individual management. This is one reason for the popularity of hyperconverged infrastructures (HCI). HCI offers a single, centrally managed unit with software-defined compute, network, and storage that is flexible, scalable, and easy to deploy.

ESG research confirms the popularity of HCI: in a recent study, 85% of respondents reported that they currently use or plan to use HCI solutions in the coming months.¹ This is not surprising given the factors driving them to consider HCI. As detailed in Figure 1, a wide range of capabilities have driven organizations to deploy or consider deploying hyperconverged technology solutions: improved service and support, scalability, agile VM provisioning, predictable costs, simplified management, fast deployment, better TCO, fewer interoperability problems, and ease of acquisition.² It almost sounds too good to be true.

Figure 1. Top Ten Factors Driving Deployment of Hyperconverged Infrastructure



Source: Enterprise Strategy Group, 2017

And in many cases, it *is* too good to be true. First, because the initial generation of HCI solutions, using software on generic x86 servers, focused on simplicity and getting to market quickly. In doing so, they traded off features that are essential for

¹ Source: ESG Research Report, [The Cloud Computing Spectrum, from Private to Hybrid](#), March 2016.

² Source: Ibid.

the speed and agility required today, such as network automation, independent resource scaling, and high performance. In addition, they often retain separate management silos that reduce the simplicity benefit.

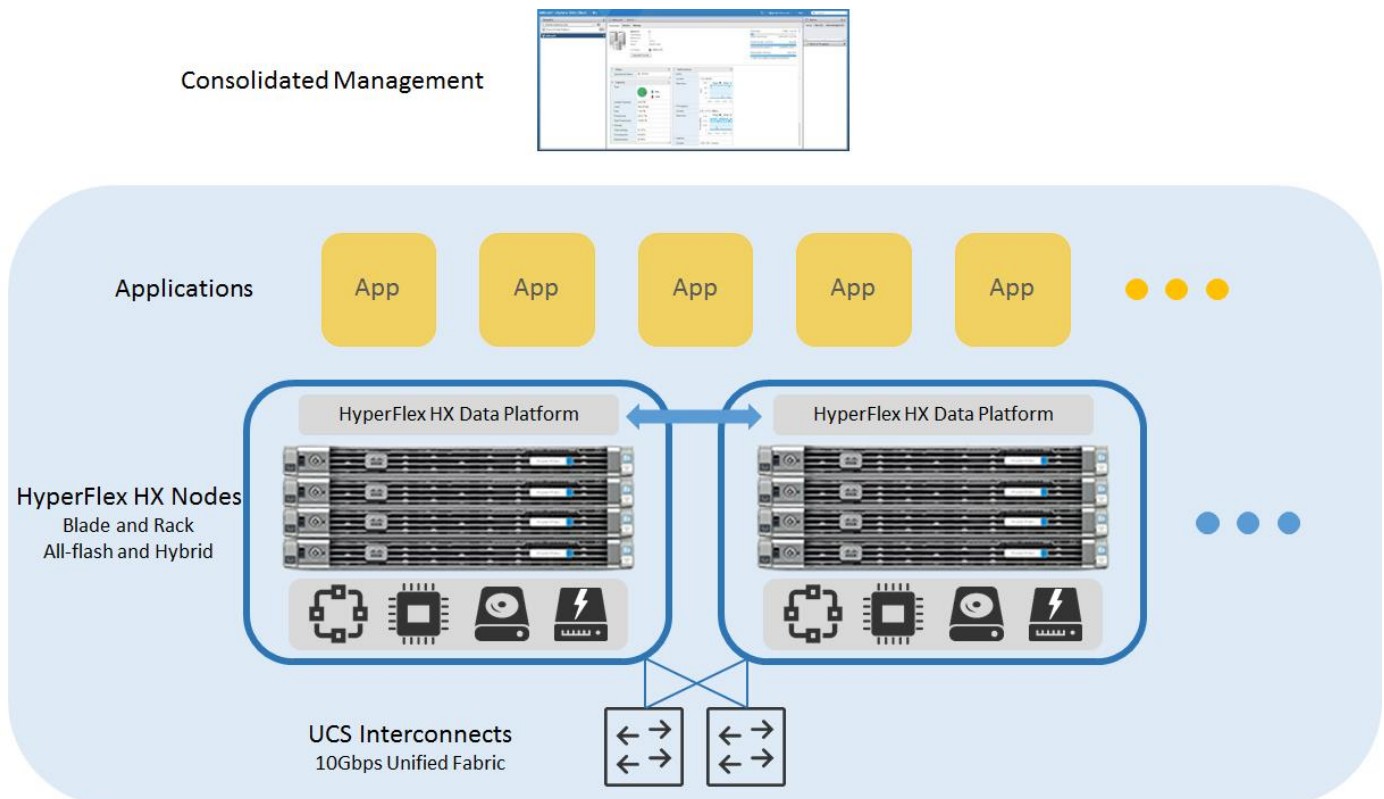
Another reason some organizations have resisted HCI is that many solutions cannot deliver the consistent high performance that mission-critical workloads demand. Simplicity is no longer the only priority; as more HCI solutions have come to market, the key buying criteria have expanded to include performance.

Cisco HyperFlex

Cisco HyperFlex systems combine compute, network, and storage in a fully integrated, fully engineered platform designed to scale resources independently and deliver consistent high performance. Cisco HyperFlex is engineered on Cisco UCS, combining the benefits of the UCS platform (such as policy-based automation for servers and networking) with those of the HX Data Platform's distributed filesystem for hyperconvergence.

It supports a broad range of applications and workloads for data centers and remote locations. VMware environments are currently supported, with support for other hypervisors plus bare metal and containerized environments on the roadmap. HyperFlex deployments require a minimum three-node cluster for high availability, with data replicated across at least two nodes, and a third node to protect against single-node failure.

Figure 2. Cisco HyperFlex Hyperconverged Infrastructure



Source: Enterprise Strategy Group, 2017

HyperFlex HX-Series Nodes are powered by Intel Xeon processors, and comprise:

- **Cisco UCS servers.** Both blade and rack servers can be combined in the cluster, with a single network hop between any two nodes for maximum east-west bandwidth and low latency. HyperFlex lets you alter the ratio of CPU-intensive blades to storage-intensive capacity nodes so users can optimize the system as application needs shift. All-flash and

hybrid nodes are available. UCS management is accessed via a VMware vCenter plug in, a web-based GUI, CLI, or XML API.

- **Cisco HyperFlex HX Data Platform** for software-defined storage. Operating as a controller on each node, the HX Data Platform is a high-performance, distributed file system that combines all SSD and HDD capacity across the cluster into a distributed, multi-tier, object-based data store, striping data evenly across the cluster. It also delivers enterprise data services such as snapshots, thin provisioning, and instant clones. Policy-based data replication across the cluster ensures high availability. Dynamic data placement in memory, cache, and capacity tiers optimize application performance, while in-line, always-on deduplication and compression optimize capacity.
 - The HX Data Platform handles all read and write requests for volumes accessed by the hypervisor. By striping data evenly across the cluster, network and storage hotspots are avoided, and VMs enjoy optimal I/O performance regardless of location. Writes go to local SSD cache, and are replicated to remote SSD in parallel before the write is acknowledged. Reads are from local SSD if possible, or retrieved from remote SSD.
 - The log-structured file system is a distributed object store that uses a configurable SSD cache to speed reads and writes, with capacity in HDD (hybrid) or larger SSD (all-flash) persistent tiers. When data is de-staged to persistent tiers, a single sequential operation that writes a large amount of data enhances performance. In-line deduplication and compression occur when data is de-staged; data movement happens after the write is acknowledged so there is no performance impact.
- **Cisco Unified Fabric /UCS 6200 Fabric Interconnects** enable software-defined networking. High bandwidth, low latency, and 40Gbps and 10Gbps connectivity in the fabric enable high availability as data is securely distributed and replicated across the cluster. The network scales easily, and each connection is fully secure. The single hop architecture enhances cluster performance.
- **Cisco Application Centric Infrastructure (ACI)** for automated provisioning. ACI enables automation of network deployment, application services, security policies, and workload placement per defined service profiles. This provides faster, more accurate, more secure, lower cost deployments. ACI automatically routes traffic to optimize performance and resource utilization, and re-routes traffic around hotspots for optimal performance.
- **VMware ESXi and vCenter.** The VMware hypervisor and management application come pre-installed, providing a familiar management interface for all hardware and software.

Cisco HyperFlex delivers numerous benefits, including:

- **High performance.** In addition to performance features mentioned above, HyperFlex securely distributes data across servers and storage in the cluster to reduce bottlenecks.
- **Fast, easy deployment.** This pre-integrated cluster can be deployed just by plugging into the network and applying power. Node configuration and connection is handled through Cisco UCS service profiles. Cisco says that customers report typical deployment times of less than one hour.
- **Consolidated management.** Systems are monitored and managed through VMware vCenter, eliminating the separate management silos for compute and storage. Provisioning, cloning, and snapshots are offloaded to vSphere using VAAI. APIs support cloud-native data types.
- **Independent scaling.** Different from other HCI systems, HyperFlex can independently scale compute and storage by adding or subtracting either servers or individual drives; data is automatically rebalanced. This provides the right resources for different application needs, instead of scaling in pre-defined increments.

ESG Lab Validation

Testing was conducted using industry-standard tools and methodologies, and was focused on HyperFlex hybrid and all-flash performance with comparisons to unnamed alternative solutions. These solutions included two “software-only” systems from leading vendors that leveraged standard x86-based servers, and a proprietary system from a single vendor based on its own hardware and partially integrated with its own software.

The bulk of the testing used HCIBench, an industry-standard tool designed to test the performance of HCI clusters running virtual machines. HCIBench leverages Oracle’s Vdbench tool and automates the end-to-end process that includes deploying test VMs, coordinating workload runs, aggregating test results, and collecting data.

This extensive testing was executed using a stringent methodology including many months of baselining and iterative testing. While it is often easier to generate good performance numbers with a short test, benchmarks were run for long periods of time to observe performance as it would occur in a customer’s environment. In addition, tests were run many times, never back-to-back but separated by hours and days, and the results averaged. These efforts add credibility by reducing the chances that results were influenced by chance circumstances. Also, testing was conducted using data sets large enough to ensure that data did not remain in cache, but leveraged the back-end disk across each cluster.³

Hybrid Testing

Testing of hybrid solutions included both SSD and HDD. The hybrid test bed included a four-node HyperFlex HX220c cluster with one 480GB SSD for cache and six 1.2TB SAS HDDs for capacity. Tests were run with 140 VMs (35 VMs per node), each with 4 vCPUs, 4 GB RAM, one 20GB disk, and running RHEL version 7.2. The working set size was 2.8 TB. Tests were run for a minimum of one hour, with a five-minute ramp-up before each test and a minimum one-hour cool-down between tests.

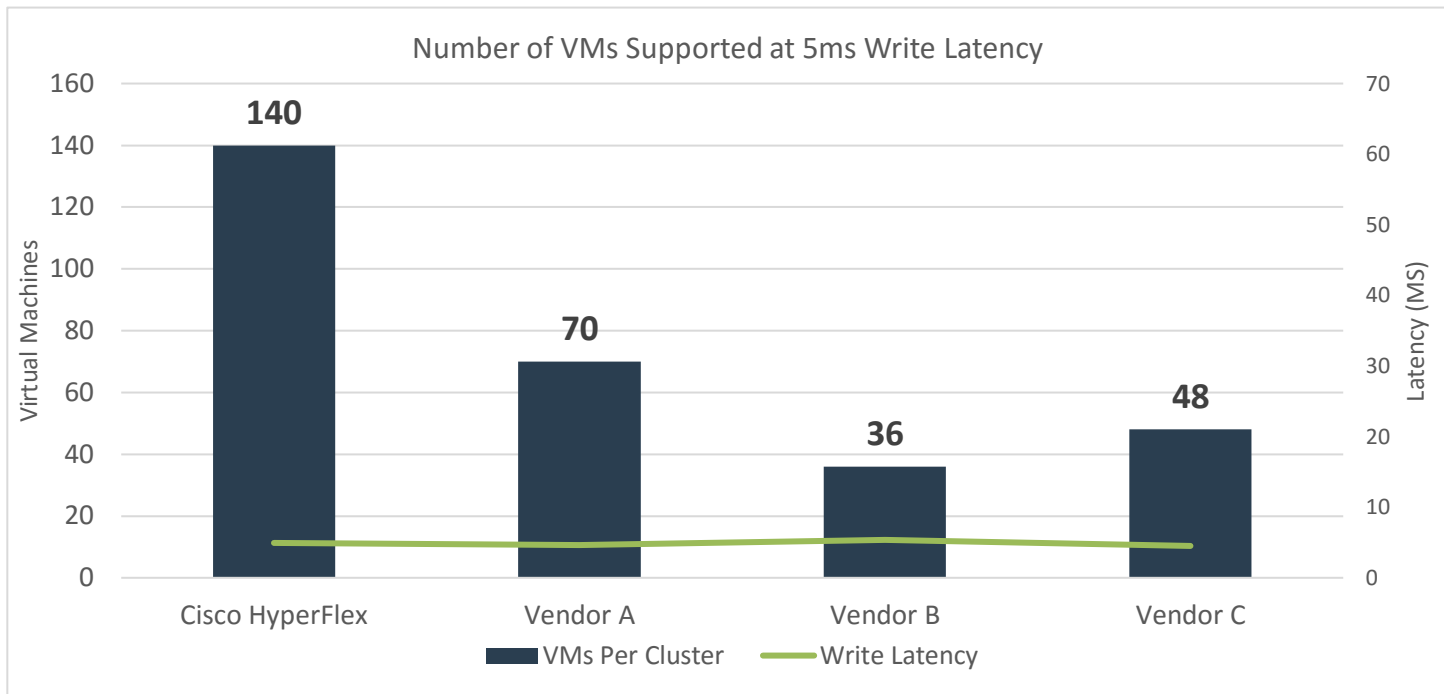
Comparative HCI solutions were also 2U, four-node systems with similar configurations, although all used two cache SSDs while HyperFlex used only one. Vendor A used two 400GB SSDs and four 1TB SATA HDDs; Vendor B used two 400GB SSDs and 12 1.2TB SAS HDDs; Vendor C used four 480GB SSDs and 12 900GB SAS HDDs.

Testing was performed using various read/write profiles and block sizes, with 100% random data. VMs by nature generate random I/O by combining I/O from multiple applications and workloads. ESG Lab focused on results obtained using workloads designed to simulate real-world applications such as a 4KB and 8KB OLTP and SQL Server.

First, ESG Lab looked at overall cluster scalability. The test began with a synthetic workload designed to emulate a typical OLTP I/O mix, 70% read, 100% random with a per-VM target of 800 IOPS. The test was run across 140 VMs in each cluster for three to four hours with a goal of remaining at or below 5ms write latency. As shown in Figure 3, HyperFlex was the only platform to complete this test with 140 VMs and stay below 5ms (4.95ms). For each of the other clusters, the test was re-run against decreasing numbers of virtual machines until write latency of 5ms was achieved. Vendor A successfully supported 70 VMs at 4.65ms average response time, Vendor B passed running 36 VMs with 5.37ms average response times, and Vendor C supported 48 VMs at sub 5.02ms response times.

³ When evaluating technology solutions, customers would be wise to understand the details behind vendor testing. Timing of test runs, volumes of data, and other details will impact performance results; these results may or may not be relevant to the customer environment.

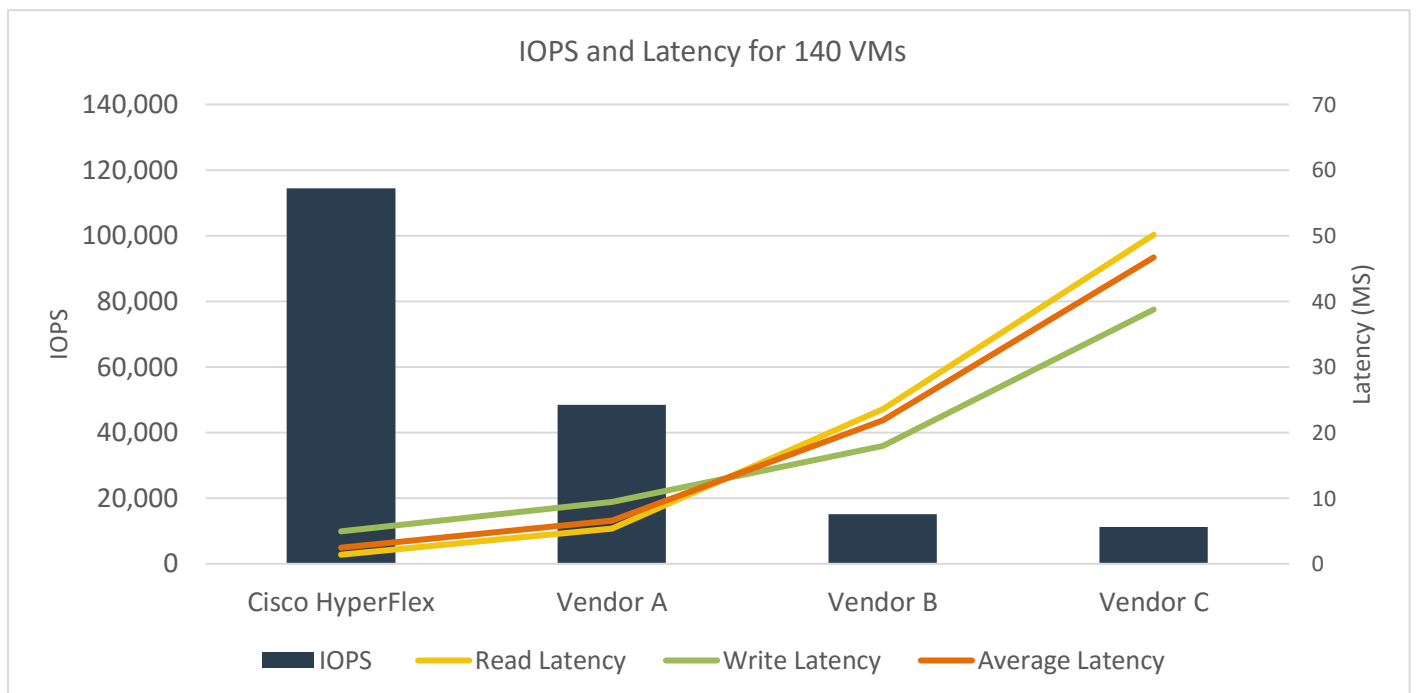
Figure 3. Hybrid Cluster Scalability



Source: Enterprise Strategy Group, 2017

Next, ESG Lab examined the same synthetic workload against 140 virtual machines to measure the latency of each cluster against IOPS. As seen in Figure 4, the Cisco HyperFlex cluster more than doubled the IOPS of vendor A and supported nearly eight times the IOPS of Vendor B and Vendor C with an average response time of 2.46ms. In comparison, Vendor A’s average response time was 6.61 ms, Vendor B’s was 21.88ms, and vendor C’s was the highest, at 44.45ms.


Figure 4. Hybrid Cluster Performance—4 KB I/O, 70% Read, 100% Random



Source: Enterprise Strategy Group, 2017

Next, ESG Lab looked at a synthetic workload designed to simulate SQL Server I/O patterns.⁴ Vdbench was used to create a synthetic workload that exercised different transfer sizes and read/write ratios. In the Vdbench profile the deduplication ratio was set to 2 with a unit size of 4 KB and the compressibility ratio also set to 2. Again, the test was run with 140 virtual machines.

Figure 5. Hybrid Cluster Performance-Vdbench SQL Server Curve Test



Source: Enterprise Strategy Group, 2017

As Figure 5 shows, the Cisco HyperFlex cluster nearly doubled the IOPS of both Vendor A and Vendor B and was more than five times the IOPS of Vendor C. Cisco HyperFlex posted an average response time of 8.2ms. By way of comparison, Vendor A’s average response time was 30.6ms, Vendor B’s was 12.8ms, and Vendor C’s was 10.33ms.

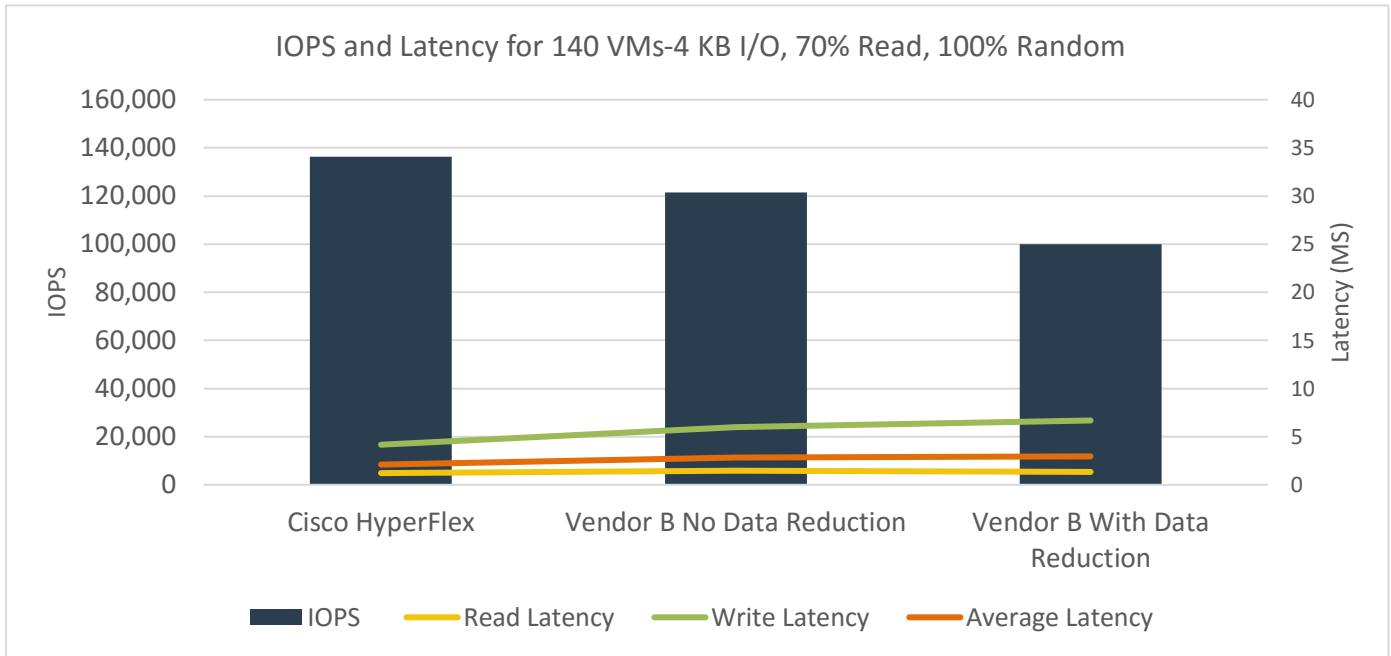
All-flash Testing

ESG Lab also looked at performance of all-flash configurations of Cisco HyperFlex and Vendor B, a software-based HCI offering running Cisco C240 M4 Rack Servers. All-flash testing used a four-node, Cisco HyperFlex 220C cluster with one 400GB SSD and six 960GB SSDs. The comparative four-node cluster used twice the cache—two 400GB SSDs—and the same number (six) of 960GB SSDs. It’s important to note that Vendor B’s system was configured with the same CPU and memory configuration as in the Cisco HyperFlex 220C cluster.

Testing again used 140 VMs per cluster (35 per node). Each VM, running RHEL 7.2, leveraged four vCPUs, 4 GB RAM, a 16GB local disk, and one 40GB raw disk. The working set was 5.6 TB, and I/O was 100% random; tests were run with a five-minute warmup, a one-hour test run, and one-hour cluster cool-down between tests. While deduplication and compression are always enabled on the Cisco HyperFlex cluster, tests were run against Vendor B with deduplication and compression set to 50%, and again with both disabled. As shown in Figure 6, the Cisco HyperFlex cluster supported more IOPS at lower latency than Vendor B with or without deduplication enabled.

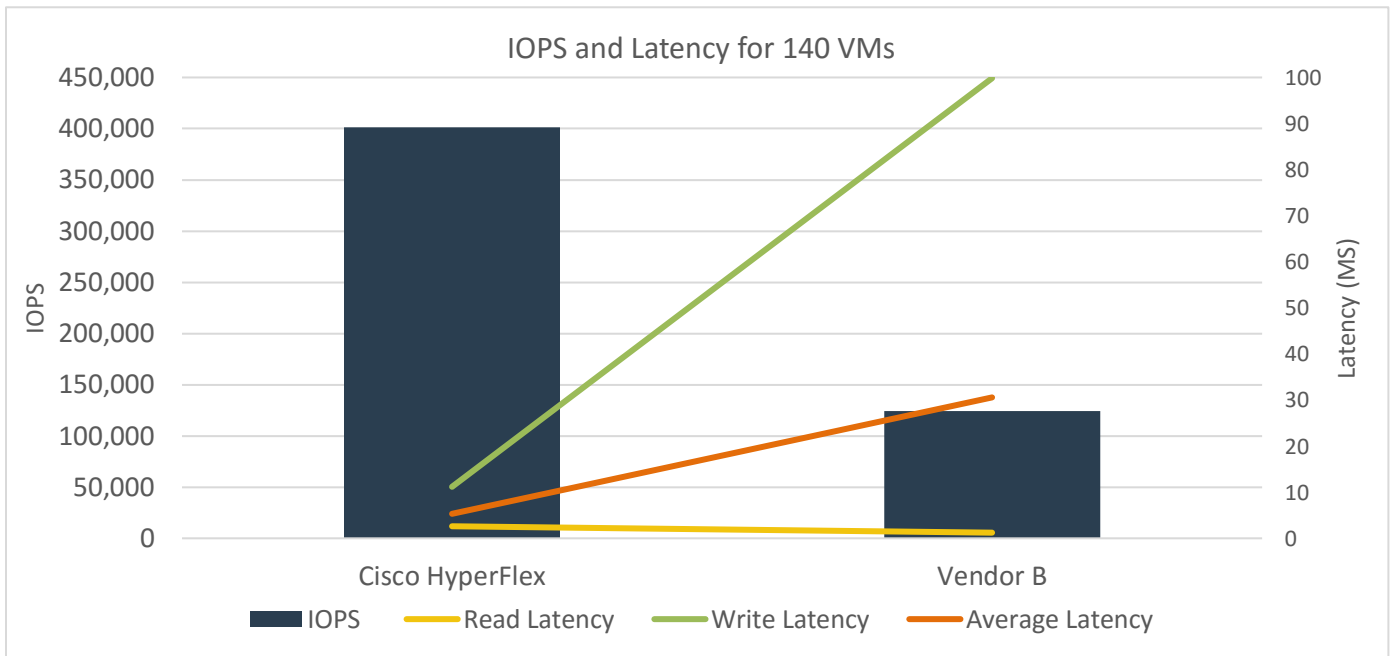
⁴ A publicly available Vdbench profile was used to simulate the I/O and data patterns produced by SQL Server and these results should not be interpreted as SQL application measurements.

Figure 6. All-flash Cluster Performance-4 KB I/O, 70% Read, 100% Random



Next, ESG Lab looked at a synthetic workload designed to simulate SQL Server I/O against all-flash configurations of Cisco HyperFlex and Vendor B. Vdbench was used to create a workload that exercised different transfer sizes and read/write ratios. In the Vdbench profile the deduplication ratio was set to 2 with a unit size of 4 KB and the compressibility ratio also set to 2.


Figure 7. All-flash Cluster Performance-Vdbench SQL Server Curve Test



As Figure 7 shows, the Cisco HyperFlex cluster more than tripled the IOPS of Vendor B with an average response time of 5.3ms. Vendor B’s average response time was 30.58ms, due to an extremely high write response time of 99.84ms throughout the test. This test was run several times on multiple days with consistent results.

An interesting observation was made during all-flash testing. Vendor B showed considerable variability in performance from VM to VM. This test was run using HClBench against 140 VMs in each cluster. While Cisco HyperFlex showed little variation across all 140 VMs—IOPS stayed very close to 600—Vendor B IOPS varied wildly, from a low of 64 to a high of 1024 IOPS.

Figure 8. All-flash Cluster Performance—4 KB I/O, 70% Read, 100% Random



It’s important to note that this variability was observed in every iteration of testing, and that no form of storage QoS was used during these test runs on either of the clusters. Network QoS was used for both systems. Inconsistency like this could be quite problematic for administrators, who would likely need to use some form of QoS (if available from the HCI vendor) to attempt to control the VMs that are consuming more than their share, so others are not starved.



Why This Matters

A common complaint about HCI systems has been performance. HCI customers have been more focused on cost-efficiency and simpler management, often relegating HCI to tier-two workloads. IT departments are unlikely to saddle their tier-one production applications with high latency and inconsistent, “noisy neighbor” VM performance that some HCI solutions offer.

ESG Lab validated that Cisco HyperFlex hybrid and all-flash systems delivered higher, more consistent performance than other similarly configured HCI solutions using simulated OLTP and SQL workloads. For hybrid clusters, HyperFlex not only consistently outpaced competitors in terms of IOPS and latency, it supported more than twice the number of VMs than both software-based and engineered proprietary systems while maintaining high performance.

The HyperFlex all-flash cluster, with always-on deduplication and compression, delivered higher IOPS and lower latency than a competitor with and without data reduction turned on. Equally important, HyperFlex all-flash performance was consistent across all VMs in the cluster, eliminating the need for storage QoS to ensure user satisfaction. In contrast, individual VMs in the competitive cluster received widely varying IOPS, indicating significantly better performance for some VMs than others.

ESG Lab Validation Highlights

- ☑ ESG Lab was impressed with the HyperFlex hybrid cluster's ability to support more than twice the number of VMs as competitors while maintaining low latency, and to deliver 2X-8X the IOPS for 140 VMs in a cluster, using an OLTP workload.
- ☑ With a SQL workload, the hybrid HyperFlex also delivered significantly more IOPS and lower latency than other solutions.
- ☑ For all-flash testing, HyperFlex delivered higher IOPS and lower latency, but even more impressive was the consistent high performance across all VMs that can ensure user satisfaction without extra management.

Issues to Consider

- ☑ HyperFlex currently supports VMware environments, an essential segment of the market. ESG looks forward to Cisco expanding to serve additional use cases such as other hypervisors, bare metal, and containerized environments.
- ☑ The test results presented in this report are based on applications and benchmarks deployed in a controlled environment with industry-standard testing tools. Due to the many variables in each production data center environment, capacity planning and testing in your own environment are recommended. While the methodology in these tests was more stringent than most, customers are well advised to always explore the details behind any vendor testing to understand the relevance to your environment.

The Bigger Truth

Hyperconverged infrastructures, while becoming mainstream, have long been considered more appropriate for tier-two workloads. When asked why they would choose converged infrastructure over hyperconverged, ESG research survey respondents' most-often-cited response was "better performance." Other responses revealed that respondents believed converged, i.e., loosely integrated independent components packaged together, was better for mission-critical workloads, and that it was available from more established players.⁵

Cisco—clearly an "established player"—has an answer to those deficiencies. HyperFlex provides the typical benefits of HCI—it is cost-effective, simple to manage, and lets organizations start small and scale. But it also provides the performance that mission-critical, virtualized workloads demand. The *consistency* of performance over time and across all VMs in a cluster was particularly notable. In addition, its independent resource scalability enables organizations to adapt quickly to changing requirements, as today's environments demand.

Cisco HyperFlex HCI solutions are highly integrated, fully engineered systems powered by Intel Xeon processors, and provide pre-integrated clusters that include the network fabric, data optimization, unified servers, and VMware ESXi/vSphere, enabling fast deployment. This makes them simple to manage and scale. ESG Lab validated that HyperFlex provides consistent high performance for VMware environments, across hybrid and all-flash clusters. HyperFlex outpaced multiple anonymous competitive solutions with higher IOPS, lower latency, and better consistency over time and across VMs.

When market evolution changes the buying criteria in an industry, there is often a mismatch between what customers want and what they can get. Vendors that can see what's missing and fill the void gain an advantage. Cisco delivers an HCI solution that provides the essential simplicity and cost-efficiency features of HCI, but also the consistent high performance that has been missing—and that customers need for mission-critical workloads. Currently, HyperFlex only supports VMware environments, and expansion to other hypervisors, bare metal, and containerized environments will be important additions.

HCI solutions have been focused on second tier workloads, but with the consistent, high performance offered by Cisco HyperFlex, there is no reason HCI cannot support tier-one production workloads. Cisco HyperFlex could be the right solution at the right time, for organizations seeking cost-effective, scalable, high performance infrastructure solutions.

⁵ Source: ESG Research Report, [The Cloud Computing Spectrum, from Private to Hybrid](#), March 2016

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.



Enterprise Strategy Group is an IT analyst, research, validation, and strategy firm that provides market intelligence and actionable insight to the global IT community.

© 2017 by The Enterprise Strategy Group, Inc. All Rights Reserved.

