

Technical Validation

Mission-critical Hyperconverged Workload Performance Testing on Cisco HyperFlex All NVMe with Intel Optane DC SSD

By Tony Palmer, Senior Validation Analyst

April 2019

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



Contents

- Introduction 3
 - Background 3
 - Powering Tier-1 workloads on HCI..... 3
 - All-flash Nodes Introduce Mission-critical Workloads to HCI 3
 - Utilizing NVMe to Enable More Mission-critical Workloads..... 4
 - Key Metrics to Consider when Evaluating HCI Solutions 4
 - Cisco’s Fully Engineered HCI Approach..... 5
 - Intel Xeon Scalable Processors and Intel Optane DC SSDs 7
- ESG Technical Validation..... 8
 - Mission-critical Hyperconverged Workload Testing..... 8
 - Aggregate Testing IOPS from the Vdbench Tool..... 9
 - ESG Testing..... 9
- The Bigger Truth..... 14

ESG Validation Reports

The goal of ESG Validation reports is to educate IT professionals about information technology solutions for companies of all types and sizes. ESG Validation 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 objectives are to explore some of the more valuable features and functions of IT solutions, show how they can be used to solve real customer problems, and identify any areas needing improvement. The ESG Validation Team’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 (HCI) performance testing, comparing the performance of the new Cisco HyperFlex All NVMe nodes utilizing Intel Optane Data Center (DC) SSDs to HyperFlex All Flash nodes, while servicing mission-critical workloads.

Background

Organizations today must be extremely agile and flexible in their ability to add applications and virtual machines (VMs) to mission-critical production environments quickly to handle the speed of business. This level of agility 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 solution with software-defined compute, network, and storage that is flexible, scalable, and easy to deploy.

Adoption of HCI has grown significantly since coming to market, and ESG research continues to confirm the popularity of HCI: In an ESG research study, 57% of respondents reported that they used or planned to use HCI solutions (see Figure 1). This is not surprising, given the factors driving them to consider HCI. Deployment drivers most cited by respondents include improved scalability (31%), total cost of ownership (28%), ease of deployment (26%), and simplified systems management (24%).¹

Figure 1. Organizations' Usage of Hyperconverged Infrastructure Solutions



Source: Enterprise Strategy Group

Powering Tier-1 workloads on HCI

All-flash Nodes Introduce Mission-critical Workloads to HCI

HCI solutions have matured and customers are leveraging HCI to modernize core data center operations. The introduction of all-flash nodes has led some organizations to evaluate or use HCI to power mission-critical workloads traditionally reserved for three-tiered architecture or converged infrastructure (CI) solutions. ESG has shown in previous reports that powering complex workloads can expose architectural deficiencies in an HCI solution not optimized to handle the workload

¹ Source: ESG Master Survey Results, [Converged and Hyperconverged Infrastructure Trends](#), October 2017.

requirements, and recommended that organizations looking to move mission-critical workloads give careful consideration to the solution they choose.² Predictable performance and low VM performance variability are critical to maximize end-user productivity across an organization, and in previous testing, Cisco HyperFlex All Flash solutions have proven to provide high IOPS and low read/write latency in a consistent, predictable manner.

Utilizing NVMe to Enable More Mission-critical Workloads

The performance gains introduced to HCI platforms using SATA- and SAS-based SSDs have opened the door to mission-critical workloads, and today, many organizations successfully run latency-sensitive workloads like Oracle, SQL Server, and mixed workload environments on their HCI clusters. However, protocol inefficiencies and the 6Gbps interface utilized by those drives limit overall performance, which can make scaling these workloads a delicate matter of balancing resources in a cluster and leaves more latency-sensitive workloads siloed in three-tiered architecture or converged infrastructure.

NVMe drives have yielded impressive performance gains over SATA and SAS SSDs, which has driven adoption into the high-performance server and storage market, but up until now, they have not been available in some well-known HCI solutions due to different hardware requirements that go beyond a simple drive qualification. To drive the next stage of workload adoption to HCI, Cisco has introduced NVMe drive technology to its HyperFlex platform to elevate performance in order to enable a greater number of VMs and more workloads. Intel and Cisco have completely qualified, validated, and engineered the entire stack—bios, driver, and controller—as one solution, which sets this offering apart from the pack.

Key Metrics to Consider when Evaluating HCI Solutions

Simplicity is no longer the only priority; as adoption of latency-sensitive mission-critical workloads continues to grow, performance needs to be included as a key buying criterion for HCI solutions to enable the next generation of HCI-powered workloads. While first generation HCI architecture—consisting of software running on x86 servers connected through commodity grade switches—worked for early use cases, the mission-critical nature of tier-1 workloads requires a solution that can deliver trusted performance.

Input/output operations per second (IOPS)—Adoption of flash-based storage has greatly reduced I/O challenges in traditional shared-storage environments, but in a clustered environment like HCI, total IOPS can vary greatly depending on the network connection between nodes as well as the software layer powering the HCI solution. For HCI deployments, it's important to evaluate both the total number of IOPS delivered by the cluster as well as the IOPS consistency that is delivered. Consistent VM performance has been a challenge since the beginning of virtualized computing, but “noisy neighbor” VM performance can be even more pronounced with HCI deployments based on how the software layer writes data across the cluster.

Latency—While IOPS are an important performance indicator, latency as it relates to the application should also be considered when purchasing an HCI solution. Clustered environments like HCI can have multiple bottlenecks like storage performance, responsiveness, and network throughput, all of which can contribute to application latency. Increased latency means decreased responsiveness of applications for users.

- **Read latency**—The time required for the storage controller to find and deliver the proper data blocks. For flash storage as evaluated in this paper, this includes the time for the flash subsystem to find the required data blocks and prepare to transfer them, and the transit time through the network.

² Source: ESG Technical Validation, [Mission-critical Workload Performance Testing of Different Hyperconverged Approaches on the Cisco Unified Computing System Platform \(UCS\)](#), July 2018.


- **Write latency**—This is the time it takes for the storage controller to perform all the activities required to write data blocks, including determination of the proper location for the data and performance of overhead activities—block erase, copy, and “garbage collection,” then writing and acknowledging the write back to the host.
- **Total latency**—Total latency is simply a combination of the read and write latencies calculated using the ratio of reads and writes used by the application. For example, for a workload that consists of 70% reads and 30% writes, the total latency is the average of the read and write results, weighted according to the percentage of each.

Cisco’s Fully Engineered HCI Approach

Cisco HyperFlex is a fully engineered hyperconverged system that combines compute and software-defined storage as well as fully integrated networking optimized for the east-west traffic flow between nodes in an HCI platform. This fully integrated platform is 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 file system for hyperconvergence.

It supports edge to edge workloads from mission-critical core data center applications to remote locations. The latest HX 4.0 update adds an all-NVMe node to the HyperFlex line-up to enable broader support for mission-critical workloads. 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

HyperFlex HX-Series Nodes are engineered on the Cisco UCS platform and powered by the latest generation of Intel Xeon Scalable processors, and comprise:

- **Cisco HyperFlex HX Data Platform.** The core of any HCI solution is the software platform, and the HX Data Platform was engineered specifically for HCI 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 inline, 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 SSDs in parallel before the write is acknowledged. Reads are from local SSDs if possible or retrieved from remote SSDs.
- 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), SSD (all-flash), or all-NVMe persistent tiers. When data is de-staged to persistent tiers, a single sequential operation writes data to enhance performance. Inline deduplication and compression occur when data is de-staged; data is moved after the write is acknowledged so there is no performance impact.
- **Cisco UCS compute-only nodes.** Both UCS 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—compute nodes—to storage-intensive capacity nodes—HX nodes—so users can optimize the system as application needs shift. All-flash and hybrid nodes are available.
- **Cisco Unified Fabric—UCS 6200/6300/6400 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 enables HX clusters to scale easily and securely. The single hop architecture is designed to maximize the efficiency of the storage software to enhance overall 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 reroutes traffic around hotspots for optimal performance.
- **Choice of industry-leading hypervisors including VMware ESXi and vCenter as well as Microsoft Hyper-V.** The 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 Dynamic Data Distribution securely and evenly distributes data across all cluster nodes 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 Cisco HyperFlex Connect or Cisco Intersight, which eliminates separate management silos for compute and storage. HyperFlex Connect lets organizations manage and monitor clusters from anywhere and at any time with metrics and trends to support the entire management lifecycle. Intersight is an optional cloud-based platform that allows users to manage all their Cisco

HyperFlex and Cisco Unified Computing System (Cisco UCS) infrastructure including traditional, hyperconverged, edge, and remote/branch offices through a single cloud-based GUI.

- **Independent scaling.** Different from other HCI systems, HyperFlex can independently scale compute and storage resources without the need to add full nodes to the cluster. Users can easily incorporate compute-only nodes with bare UCS servers through the Fabric Interconnects to add additional compute to the cluster, or, if more storage is needed, add individual drives to each node; data is automatically rebalanced. This provides the right resources for different application needs, instead of scaling in predefined node increments that also add additional software licensing costs.

Intel Xeon Scalable Processors and Intel Optane DC SSDs

The Intel Xeon Processor Scalable Family is engineered to deliver significant advances in performance and capabilities and designed to provide compelling benefits for a broad range of workloads across servers, networks, and storage. HCI solutions like HyperFlex can not only take advantage of more cores and memory bandwidth for workloads running in the hyperconverged architecture, but Xeon Scalable processors also power key storage technologies such as in-line compression and deduplication.

The Xeon Processor Scalable Family is complemented by the Intel SSD Data Center Family for PCIe:

- Intel Optane DC SSDs combine attributes of memory and storage to deliver a combination of low latency, high endurance, QoS, and high throughput, optimized to minimize storage bottlenecks.
- Intel's second generation of 3D NAND SSDs, including the Intel SSD DC P4500 Series, is optimized for reads to enable data centers to get more value out of servers and store more data. Designed for mixed workloads, the Intel SSD DC P4600 Series accelerates caching to enable more workloads per server.

Figure 3. The Intel SSD Data Center Family for PCIe



**Intel Optane SSD DC
P4800X**
PCIe 3.0 x4, NVMe



**Intel 3D NAND SSD DC
P4500, P4600**
PCIe 3.1 x4, NVMe

Source: Enterprise Strategy Group

ESG Technical Validation

Testing was conducted using industry-standard tools and methodologies and was focused on comparing the performance of Cisco’s fully engineered HyperFlex HCI solution in a traditional all-flash configuration with an all-NVMe configuration built on Intel Intel Optane DC SSD and Intel 3D NAND, with the latest generation of Intel Xeon Scalable processors. The bulk of the testing used HCI Bench and HXBench, tools designed to test the performance of HCI clusters running virtual machines. Both tools leverage Oracle’s Vdbench tool and automate 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 days and weeks, 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 storage across each cluster.³

Mission-critical Hyperconverged Workload Testing

The test bed included one four-node HyperFlex HX220c version 2.6 cluster and one four-node HyperFlex HX220c version 4.0 cluster. Configuration details are listed in Table 1.

Table 1. Tested HCI Configurations

Platform	Nodes	Processors/Cores Per Node	RAM Per Node	Cache Per Node	Storage Capacity Per Node	Hypervisor
Cisco HyperFlex All Flash	Four	2x Intel Xeon E5-2680 v4, 28 Cores	512GB	1x 800GB Enterprise Performance SAS SSD	6x 960GB SSD Enterprise Value SATA SSDs	VMware vSphere 6.5
Cisco HyperFlex All NVMe with Intel Optane DC SSD’s	Four	2x Intel Xeon Gold 6142, 32 Cores	384GB	1x 375GB Intel Optane DC P4800X SSD	6x 1TB Intel P4500 NVMe SSDs	VMware vSphere 6.5

Source: Enterprise Strategy Group

OLTP tests were run with four VMs and a 3.2TB working set, while the mixed workload test used 140 VMs (35 VMs per node), each with 4 vCPUs, 4 GB RAM, and one 40GB disk, and running RHEL version 7.2. The working set size was 5.6 TB. Tests were run for a minimum of one hour and up to five hours, with a five-minute ramp-up before each test and a minimum one-hour cool-down between tests. Before every test was run, each VM was primed with written data by the test tool. This ensures that the test is reading “real” data and writing over existing blocks and not simply returning null or zero values directly from memory. This happens when data is not primed so it is an important step to ensure that the test accurately reflects how data is read and written in an application environment. Priming of this large working set can take many hours to complete but is a wise investment in time to get more accurate performance results.

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.

Testing was performed using the I/O profiles of benchmarks designed to emulate complex, mission-critical workloads, including OLTP using Oracle and SQL Server back-ends, as well as virtual application server and desktop activity. Block sizes were assigned according to the applications being emulated, with 100% random data access. VMs by nature generate random I/O by combining I/O from multiple applications and workloads. It is important to note that all tests were run with compression and deduplication active on the Cisco HX cluster.

Aggregate Testing IOPS from the Vdbench Tool

The Vdbench tool uses a specific methodology to derive an aggregated IOPS result during benchmark testing. Aggregate testing IOPS are calculated by taking the average IOPS delivered to test virtual machines (VMs) at various workload levels—12 curves ranging from 20% to 100% loads. The average IOPS of each test VM are then aggregated to derive the aggregated testing IOPS in each test—for example, aggregated IOPS from four test VMs and each of their 12 load curves.


Note: Aggregate testing IOPS cannot be used to size workloads for specific applications.

ESG Testing

First, ESG Lab looked at an OLTP workload designed to emulate an Oracle environment.⁴ 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 3 with a unit size of 4 KB and the compressibility ratio also set to 3. The test was run with four virtual machines.

Over the course of the four-hour test, the HyperFlex All NVMe System was able to aggregate 722,187 testing IOPS in Vdbench with a total response time of just 2.8 ms, a 71% improvement over all-flash, as seen in Figure 4.


Figure 4. Oracle OLTP Workload—Aggregate Testing IOPS



Source: Enterprise Strategy Group


Response times improved considerably, as well, with the Hyperflex All NVMe system averaging 37% lower latency overall. Compression and deduplication were active on all systems.

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

Figure 5. Oracle OLTP Workload—Response Times


Source: Enterprise Strategy Group


Next, we looked at an OLTP workload designed to emulate a Microsoft SQL Server environment.⁵ There are subtle but potentially significant differences that warranted testing both Oracle and SQL workloads. 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. Again, the test was run with four virtual machines.

Figure 6. SQL Server OLTP Workload—Aggregate Testing IOPS


Source: Enterprise Strategy Group

As Figure 6 shows, the HyperFlex All NVMe cluster serviced 57% more testing IOPS than HyperFlex All Flash.


⁵ 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 7. SQL Server OLTP Workload—Response Times


Source: Enterprise Strategy Group


HyperFlex All NVMe posted an average response time of 2.895 ms. This was 34% lower than the average response time of HyperFlex All Flash, at 4.41 ms.

Next, we looked at a mixed workload designed to emulate a virtualized environment with multiple VMs running different applications. Vdbench was used to create a workload that exercised transfer sizes from 4 KB to 64 KB. We ran two sets of tests, with a read/write ratio of 70/30 and with a read/write ratio of 50/50. These tests were run using HCIBench against 140 VMs in each cluster—35 per node, emulating a mixed workload environment with many virtual machines running a variety of applications. 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 8. 70/30 Mixed Workload—Aggregate Testing IOPS


Source: Enterprise Strategy Group


As Figure 8 shows, the HyperFlex All NVMe cluster sustained 63% more aggregate testing IOPS over the five-hour test than the HyperFlex All Flash System.

Figure 9. 70/30 Mixed Workload—Response Times


Source: Enterprise Strategy Group

HyperFlex All NVMe posted an average response time of 1.482 ms. By way of comparison, that’s a 37% improvement over the HyperFlex All Flash result of 2.342 ms.

An interesting observation was made during mixed workload testing of both Hyperflex All Flash and Hyperflex All NVMe systems. Cisco HyperFlex showed little variation across all 140 VMs—aggregate testing IOPS stayed close to the targets of 600 for Hyperflex All Flash and 900 for All NVMe—see Figure 10. Software-only HCI vendors previously tested by ESG varied wildly, with much higher standard deviation.

Figure 10. Mixed Workload, 70% Read, 100% Random—140 Virtual Machines


Source: Enterprise Strategy Group

Inconsistency of performance across VMs 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 of resources so others are not starved.



Why This Matters

ESG research asked 306 IT managers and executives what benefits their organizations have realized as a result of deploying a hyperconverged infrastructure technology solution, and the top two most-cited reasons were improved scalability and improved total cost of ownership.⁶ Executives want IT to purchase new technologies to modernize their infrastructures and meet business requirements, but they prefer to not spend a lot to do so.

ESG previously validated that Cisco HyperFlex All Flash systems delivered higher, more consistent performance than other similarly configured HCI solutions using simulated OLTP, SQL, and mixed workloads.⁷ Cisco HyperFlex All NVMe with Intel Optane DC SSDs has widened the gap, increasing performance and reducing latency across the board. This translates directly to lower upfront and ongoing costs because a given workload can potentially be serviced by an even smaller number of Cisco HyperFlex nodes.

⁶ Source: ESG Master Survey Results, [Converged and Hyperconverged Infrastructure Trends](#), October 2017.

⁷ Source: ESG Technical Validation, [Mission-critical Workload Performance Testing of Different Hyperconverged Approaches on the Cisco Unified Computing System Platform \(UCS\)](#), July 2018.

The Bigger Truth

Hyperconverged infrastructures, while becoming mainstream, have long been considered more appropriate for tier-2 workloads. When asked in 2016 why they would choose converged infrastructure over hyperconverged, ESG research survey respondents' most-often-cited (54%) response was better performance. In addition, 32% of respondents believed converged, i.e., loosely integrated independent components packaged together, was better for mission-critical workloads.⁸

Fast forward to the present, and the picture has shifted, with only 24% of respondents citing performance as a reason to choose converged, while just 22% believe converged is better suited to tier-1 workloads.⁹

Cisco HyperFlex provides the typical benefits of HCI—it is cost-effective and simple to manage, and lets organizations start small and scale. Cisco HyperFlex All NVMe with Intel Optane DC SSDs provides the high performance and low latency 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 the latest generation of Intel Xeon Scalable processors that provide pre-integrated clusters that include the network fabric, data optimization, unified servers, and choice of hypervisor including VMware ESXi/vSphere and Microsoft Hyper-V, enabling fast deployment. This makes them simple to manage and scale. ESG has previously validated that HyperFlex provides consistent high performance for VMware environments running mission-critical workloads, outpacing multiple competitive solutions with higher IOPS, lower latency, and better consistency over time and across VMs. HyperFlex All NVMe has raised the bar, increasing performance by up to 64%, even while reducing latency across the board.

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.

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. HyperFlex supports VMware and Microsoft on-premises virtualized environments, and expansion to bare metal, containerized, and multi-cloud environments.

HCI solutions have been focused on second tier workloads, but the consistent high performance offered by Cisco HyperFlex All NVMe further validates HyperFlex as extremely well-suited to tier-1 production workloads. Organizations seeking cost-effective, scalable, high-performance infrastructure solutions for mission-critical workloads would be smart to take a close look at Cisco HyperFlex All NVMe with Intel Optane DC SSDs.

⁸ Source: ESG Research Report, [The Cloud Computing Spectrum, from Private to Hybrid](#), March 2016.

⁹ Source: ESG Master Survey Results, [Converged and Hyperconverged Infrastructure Trends](#), October 2017.

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.

© 2019 by The Enterprise Strategy Group, Inc. All Rights Reserved.

