

Technical Validation

Veritas NetBackup Flex Scale: Simple, Automated Data Protection Services

With a Scale-out, Hyperconverged, Containerized Platform

By Kerry Dolan, Sr. Validation Analyst; and Vinny Choinski, Sr. Validation Analyst

August 2021

This ESG Technical Validation was commissioned by Veritas and is distributed under license from ESG.

Contents

Introduction	3
Background	3
NetBackup Flex Scale	4
ESG Technical Validation	6
Ease of Deployment and Scaling	6
Initial Deployment.....	6
Add Node	7
Resiliency and Automation	9
Backups Continue During Failures	9
Replace Node	10
Performance	11
ESG Testing.....	11
The Bigger Truth.....	15

ESG Technical Validations

The goal of ESG Technical Validations is to educate IT professionals about information technology solutions for companies of all types and sizes. ESG Technical Validations 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

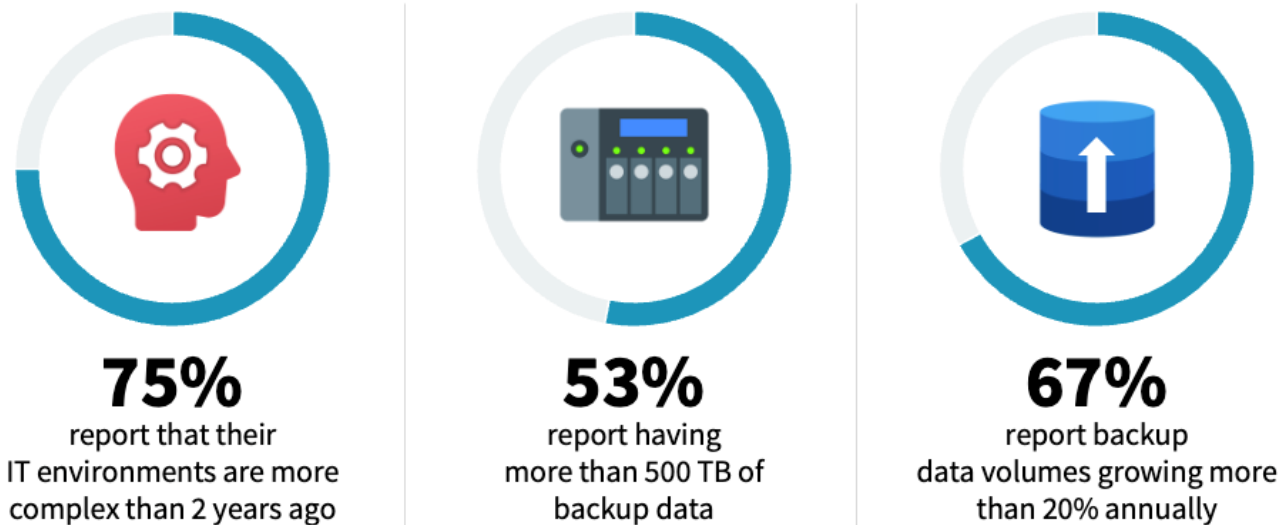
This report documents ESG's validation of Veritas NetBackup Flex Scale, a scale-out, hyperconverged data protection platform, with a focus on ease of use, scalability, and performance.

Background

While technology innovations are transforming IT in positive ways, organizations continue to struggle as IT becomes more complex. According to recent ESG research, 75% of survey respondents reported that their IT environments are more complex than two years ago.¹ Further complicating IT is the trend toward having more generalists on the IT staff handling a range of tasks, rather than experts dedicated to specific areas like data protection.

Storing and protecting the continually growing amounts of data also increase stress on both infrastructure and staff. According to recent ESG research, more than half (53%) of survey respondents reported having more than 500 TB of backup data (with 11% reporting having 10 PB or more), and 67% reported that backup data volumes were growing more than 20% annually (with 32% reporting that they were growing more than 50% annually).² Protecting this data is a critical assignment in today's world in which organizations depend on highly available applications and data. These realities have led to serious challenges for data protection.

Figure 1. IT and Data Protection Challenges



Source: Enterprise Strategy Group

To meet these complexity and growth challenges for a range of modern workloads, organizations need rock-solid data protection solutions that are simple to deploy and manage, highly scalable, high performing, and automated to relieve IT staff from management tasks.

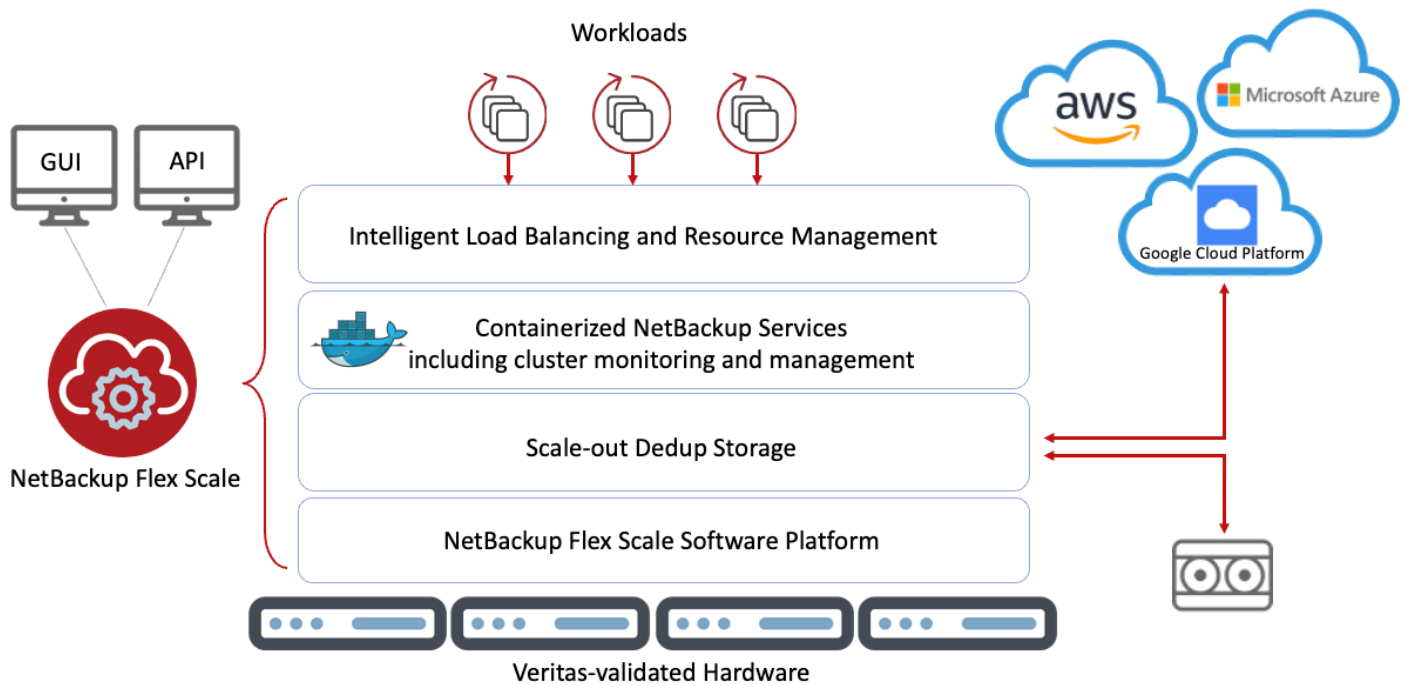
¹ Source: ESG Research Report, [2021 Technology Spending Intentions Survey](#), January 2021.

² Source: ESG Research Report, [Tape's Place in an Increasingly Cloud-based IT Landscape](#), January 2021.

NetBackup Flex Scale

NetBackup Flex Scale is the latest of several options for deploying NetBackup, supporting the same 800+ workloads as other NetBackup deployments. NetBackup Flex Scale is a hyperconverged, scale-out solution that comes pre-installed with the NetBackup Flex Scale Software Platform on Veritas-validated commodity hardware. Backup services are containerized, with abstracted media and master servers. No external hardware components are involved. NetBackup Flex Scale offers simplicity, automation, scalability, and high performance; management is simple with the same NetBackup user interface (UI) to which customers are accustomed or with easy-to-integrate APIs. NetBackup customers have no learning curve to delay time to value, just an automated, scale-out, containerized architecture underneath what they already know. NetBackup Flex Scale also supports multiple clouds and tape.

Figure 2. NetBackup Flex Scale Architecture



Source: Enterprise Strategy Group

The foundation of NetBackup Flex Scale consists of the Veritas-validated hardware platform and NetBackup Flex Scale software platform. The current validated hardware is the Hewlett Packard Enterprise (HPE) ProLiant DL380 Gen10, a dense node that helps to reduce TCO. It includes:

- 2 x 16 Intel Xeon processors.
- 192GB of RAM.
- 2 x 7.68TB NVMe SSDs for catalog and metadata.
- 12 x hot-swappable, 14TB SAS HDDs for backup data.
- 2 x 10/25 GbE dual port adapters, with support for VLANs and optional bonding.



A private UDP network with two 25Gbps ports connects the nodes in the cluster, providing more bandwidth and performance than TCP.

The NetBackup Flex Scale software platform includes the customized RHEL OS, hardened with built-in security enhancements to minimize threats.

There are numerous features and automations built into NetBackup Flex Scale that contribute to its ease of management.

- *Intelligent load balancing and resource management.* Instead of using a round-robin system that doesn't consider resource usage, NetBackup Flex Scale has an intelligent, containerized load balancer that knows the load on each node and distributes tasks intelligently. When additional resources are needed, new nodes are added and NetBackup Flex Scale automatically spins up new containers and services. Administrators have no need to even know if they need more services; they are just created automatically. Should a container fail, NetBackup Flex Scale will automatically restart it in place; should a node fail, NetBackup Flex Scale will redistribute and restart services across available nodes. When nodes are added or replaced, the new nodes are automatically added to the resource pool. All of this happens non-disruptively.
- *Containerized NetBackup services.* NetBackup services, including media and master services, are all containerized using Docker. They are created automatically during the initial configuration and distributed across all nodes. There is no need for a direct connection to specific services; all clients connect to an alias for services, eliminating any client-side impacts for policy or hardware changes. Container snapshots make rolling upgrades simple and fast, and container management is built into NetBackup Flex Scale using the embedded Veritas InfoScale functionality.
- *Scale-out, deduplicated storage.* At initial configuration, all HDDs are aggregated into a single storage pool. Efficient data storage is guaranteed with NetBackup's built-in Media Server Dedup Storage Pool. As nodes are added, new HDDs are included in the pool.
 - Nodes can be added singly or in pairs, up to 16 nodes; each node adds 112 TB of usable capacity, for a total of 1.8 PB of usable storage in a single domain. Resiliency increases as nodes are added. NetBackup Flex Scale can also write directly to the cloud and out to tape.
 - The backup catalog is placed on SSDs in each node and protected with triple mirrors, as well as with default catalog backup and snapshot policies; replication options are also available. The catalog is automatically scaled as the cluster scales, eliminating the need for administrators to monitor or manually increase catalog capacity.
 - Data is protected with erasure coding 8:4, with each two-MB data slice divided into data fragments and parity fragments and spread across nodes and disks in the cluster. This offers the greatest balance of capacity, performance, and data durability; NetBackup Flex Scale can survive the loss of any four disks in the cluster with no data loss. Erasure coding is block-based, eliminating the performance impact of cluster management tasks related to garbage collection.

NetBackup Flex Scale also enables instant access to VMs so that organizations can resume operations quickly when critically needed without having to wait for the data to be restored.

The intuitive NetBackup Flex Scale UI and dashboard make it easy to view high-level overviews and details of the NetBackup Flex Scale deployment, including alerts; NetBackup services; and storage, catalog, infrastructure, and performance details.

ESG Technical Validation

ESG viewed several demos that showed how easy it is to deploy and scale NetBackup Flex Scale; its resiliency and ability to continue backup jobs during multiple node failures; and its automatic rebalancing of data when nodes are swapped out or added. We also viewed validated results that show NetBackup Flex Scale's enterprise performance and scalability.

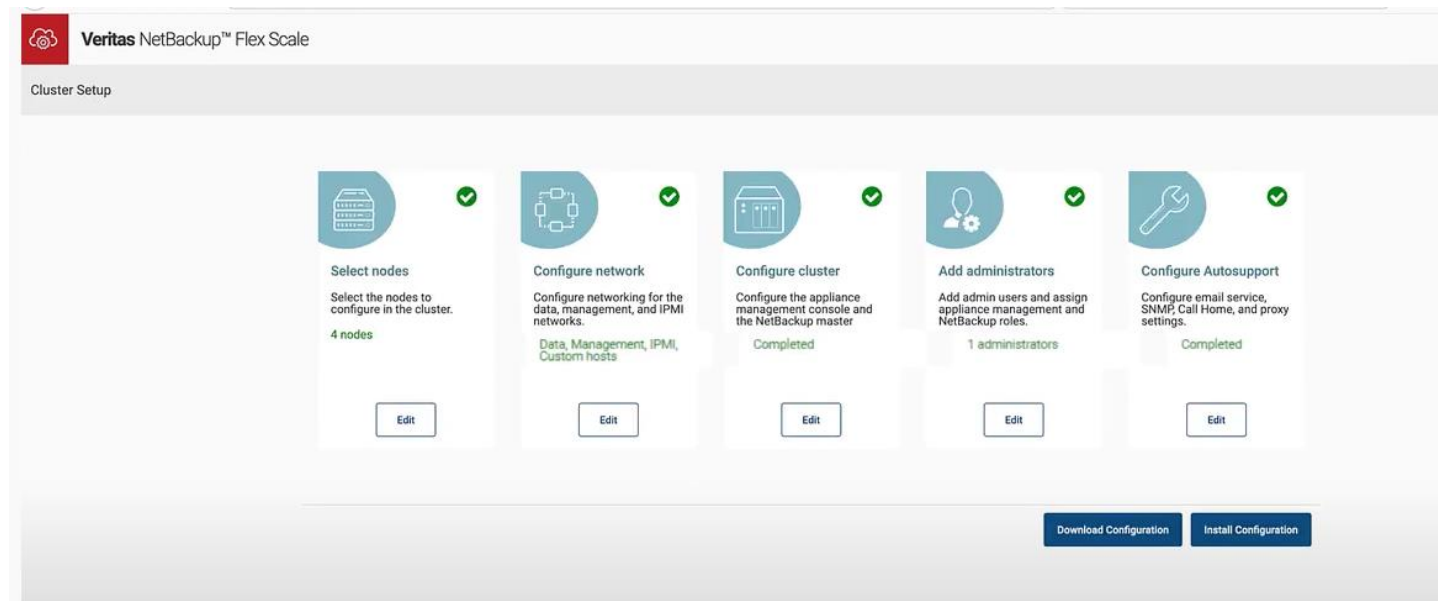
Ease of Deployment and Scaling

ESG viewed demos showing the simplicity and speed of deploying NetBackup Flex Scale, which can be done with the set-up wizard or API; a YAML file can be used to reduce manual entry errors. IT administrators only rack, cable, and power on the hardware; the software is pre-loaded. They log in and use the cluster set-up wizard to add a few details, which takes less than five minutes; then the software automatically completes all the configuration on its own.

Initial Deployment

The deployment began with a four-node cluster, racked, cabled, and powered on. First, we logged into one of the NetBackup Flex Scale nodes and added an IP address to configure the management port; from here, a connection was made to a built-in webservice to complete the cluster configuration. After logging into the NetBackup Flex Scale application and viewing prerequisites, the five-step cluster set-up wizard appeared.

Figure 3. Completed Set-up Wizard



Source: Enterprise Strategy Group

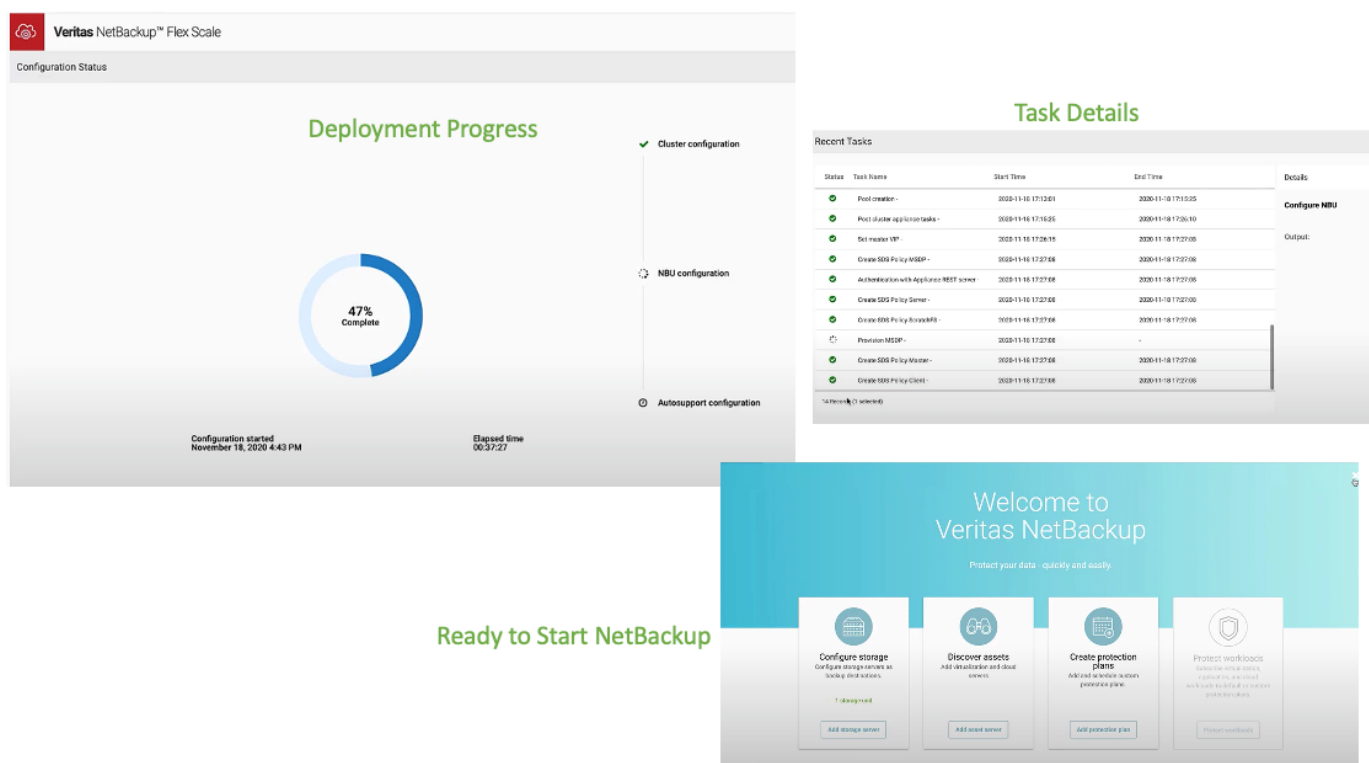
- After clicking **Start** under **Select Nodes**, NetBackup Flex Scale detected and listed all four nodes in the cluster, which we then selected. We browsed to and opened a YAML file containing configuration details to avoid manual entry that can be error prone. All configuration details were then pre-populated, so we only needed to confirm them in subsequent screens. For example, the domain, cluster, and individual node names were then visible on the **Select Node** screen.
- We confirmed network configurations next, including IP addresses for the data network for NetBackup services, including media servers; the management network, including nodes and DNS; and the optional IPMI network and custom host files.

- Next, we validated cluster configuration details, including the management console, NetBackup Master server, and time zone. License keys are added here as well.
- Then, we added the usernames of appliance and NetBackup administrators with defined roles, and finally, we verified the optional Autosupport details, including email addresses, SNMP contacts, and call home and proxy settings.

In less than five minutes, all five steps were completed (with a green check mark indicating completion); we simply clicked **Install Configuration**, and NetBackup Flex Scale completed all node, network, and NetBackup configurations on its own. This included setting up the clustered file system, spinning up NetBackup media and master service containers, configuring the storage unit, creating catalog protection policies and snapshot schedules, and setting up Autosupport. Progress was tracked graphically, and we could view details at any time with a single click. In this example, the automated process was completed in about an hour and a quarter.

Finally, we logged into NetBackup, showing the already configured storage unit created from the disks in the cluster for storing local backups. Now, we were ready to discover assets and create workload protection plans using NetBackup. Figure 4 shows the deployment progress graphic, details of tasks completed, and the NetBackup welcome screen.

Figure 4. NetBackup Flex Scale Automatic Deployment



Source: Enterprise Strategy Group

Add Node

ESG also viewed demonstrations of how easy it is to add and replace nodes in a NetBackup Flex Scale cluster. Administrators can add nodes individually or in groups. Each node includes the complete hardware and software pre-installed. With a four-node cluster running numerous NetBackup jobs, we watched as a new node was powered on and was then automatically discovered by NetBackup Flex Scale. We selected the new node and clicked **Add nodes to the**

cluster. NetBackup Flex Scale then displayed the before and after details, including the number of nodes and TBs of capacity.

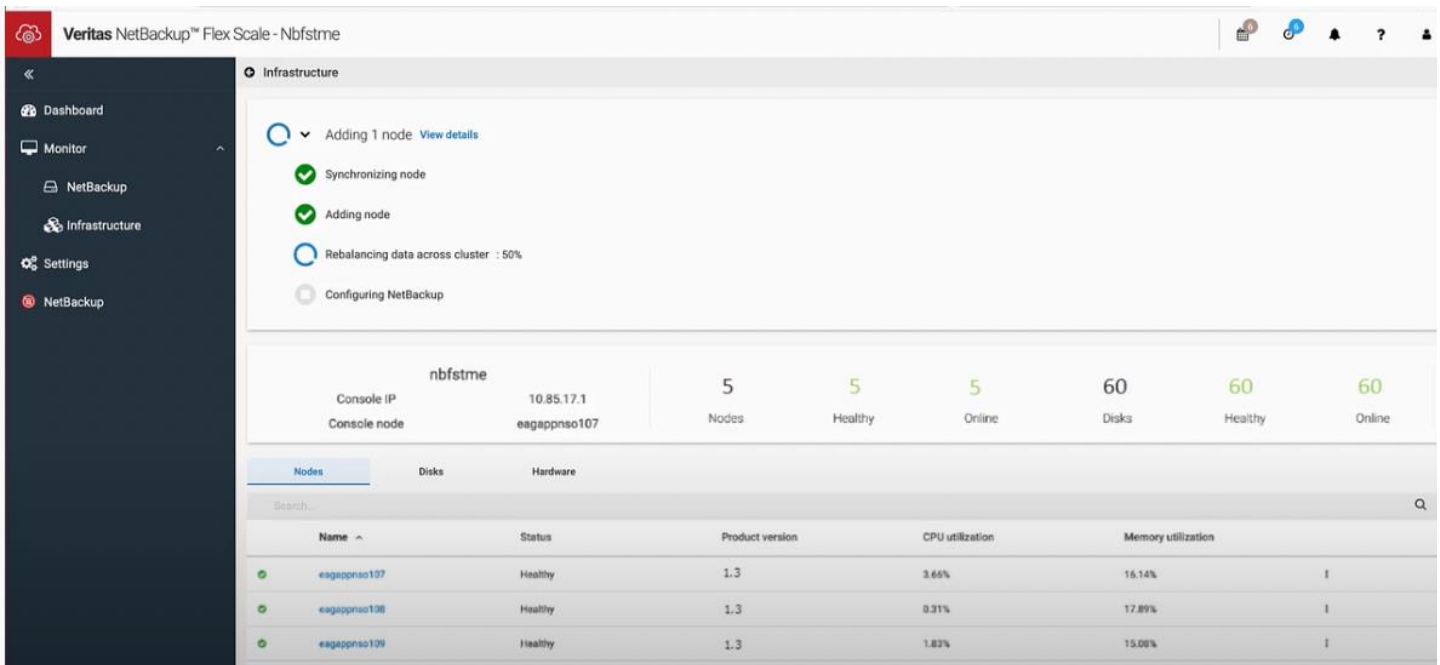
Next, we could select the workload priority, maximizing either the overall system performance (which would prioritize backup/recovery jobs) or faster reconfiguration (which would prioritize node addition and cluster rebalancing), depending on the business need at the time. NetBackup jobs continue uninterrupted in both cases.

We added the node name and the names and IP addresses of the data, management, and optional IPMI networks, and the reconfiguration began automatically. NetBackup Flex Scale then added the node, configured its software and network, scaled the clustered file system, and rebalanced data across all nodes.

Importantly, during the data rebalancing, NetBackup Flex Scale moves the minimum number of blocks to optimize performance and resiliency. We could view task details with a single click and view the NetBackup jobs continuing uninterrupted. Finally, NetBackup Flex Scale automatically scaled the NetBackup services by adding the required containers on the new node and connecting that node into the cluster.

Once this process was complete, NetBackup Flex Scale began intelligently balancing NetBackup jobs across the expanded cluster, with no changes to policy or client configurations. Figure 5 shows the node addition and data rebalancing in process.

Figure 5. Automated Node Addition



Source: Enterprise Strategy Group



Why This Matters

Growing data volumes and management complexity complicate data protection, often resulting in data vulnerability.

ESG validated that NetBackup Flex Scale, with its containerized, hyperconverged, scale-out architecture, offers simple, automated deployment and scalability. Using the easy set-up wizard, deployment took less than five minutes of administrator time, after which NetBackup Flex Scale completed the installation on its own. Adding nodes to the cluster was equally easy, requiring just a few clicks. NetBackup Flex Scale automatically added the nodes, scaled the clustered file system, and rebalanced data across the expanded cluster, with no changes to policies or clients.

Resiliency and Automation

NetBackup Flex Scale is built for resiliency, with erasure coding 8:4 built in to protect against disk failures. NetBackup Flex Scale can lose up to four disks or an entire node without interruption. As the cluster expands, resiliency increases, as shown in Table 1.

Table 1. NetBackup Flex Scale Resiliency Increases with Additional Nodes

Cluster Size	Maximum # of Supported Failures		
	4-5 Nodes	Any 4 HDDs, or	1 Node + Any 1 HDD, or
6-12 Nodes	Any 4 HDDs, or	1 Node + Any 2 HDDs, or	2 Nodes
12+ Nodes	Any 4 HDDs, or	1 Node + Any 3 HDDs, or	2 Nodes

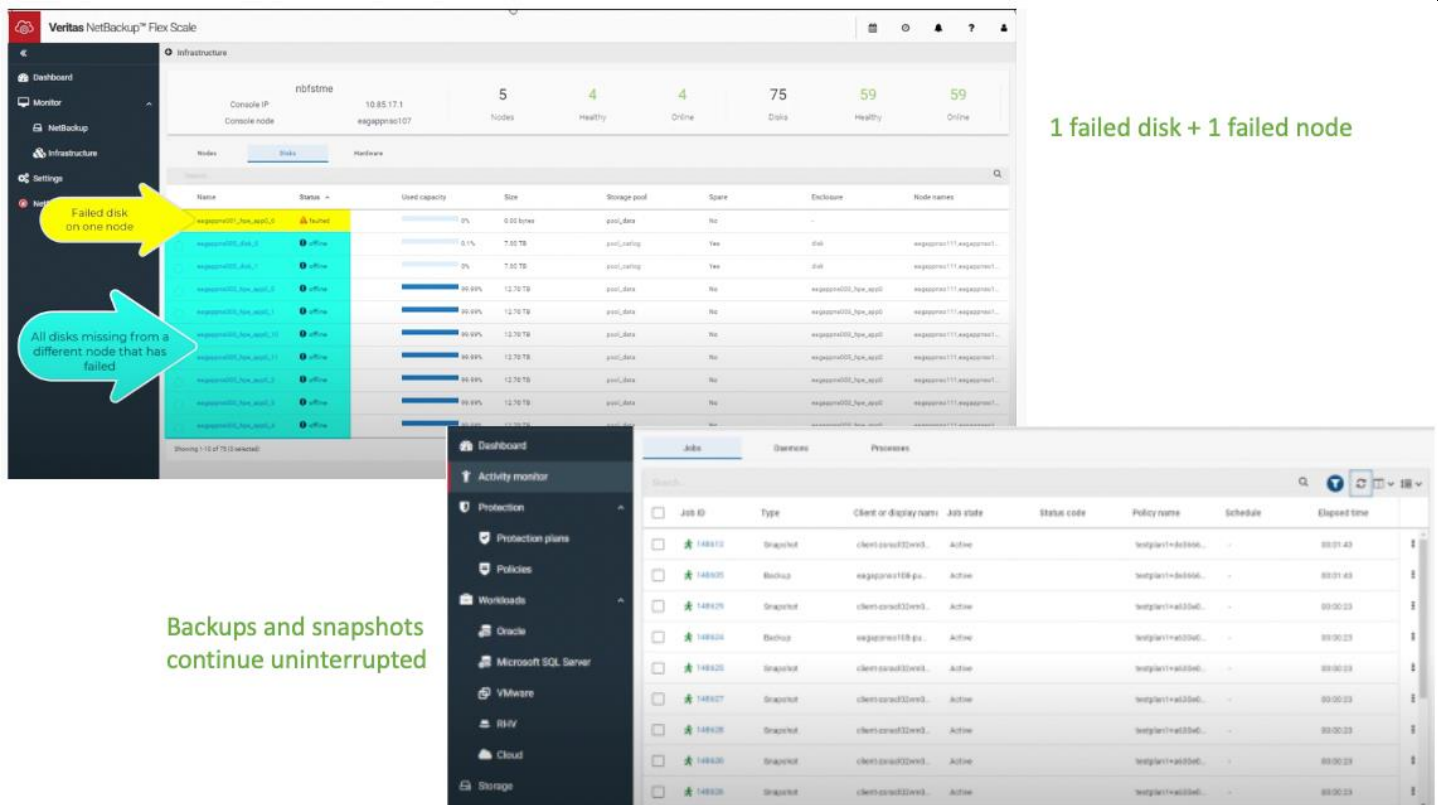
It also supports single- and dual-domain replication for disaster recovery. For single site domains, a single API call will promote the remote site and reverse catalog replication, ensuring optimal protection and simplicity. Node upgrades are also non-disruptive.

Backups Continue During Failures

ESG viewed multiple demos showing the continuation of backup jobs during disk failures, as well as the intelligent resumption of NetBackup services.

We began with a five-node cluster. Drilling into the Infrastructure tab, we could see one failed disk on one node, plus a failed node (showing all its disks offline). All data remained intact, and when we viewed the NetBackup job screen (see Figure 6), all jobs were continuing to run uninterrupted.

Figure 6. Node Failure with Uninterrupted NetBackup Jobs



Source: Enterprise Strategy Group

Replace Node

In addition, we added a node after a failure. Once the new node was connected to the network, we selected the failed node and clicked **Replace Node**. After we selected the workload priority option, NetBackup Flex Scale was able to rebuild the data from the failed node using the erasure-coded fragments striped across the rest of the cluster, rebuilding that data quickly. Next, on the new node, NetBackup Flex Scale automatically re-created the configuration from the failed node and spun up NetBackup services as needed. Once that was complete, the intelligent load balancer began re-optimizing data across the available resources.

It is important to note that the containerization of services makes it very easy and fast to recover from failures. If the load-balancing or master service container had been on the failed node, it would have been created again on another node automatically. No data migration, manual reconfiguration, changes on the client side, or professional services are needed; NetBackup Flex Scale handles it alone.

Why This Matters

Modern businesses demand high availability of applications and data to remain competitive and compliant.

With block-based erasure coding 8:4, the minimum NetBackup Flex Scale cluster can survive losing four disks, with resiliency increasing as nodes are added. Data replication is also supported, as are non-disruptive upgrades and node replacements, and the NetBackup catalog is automatically protected with triple mirrors, snapshots, and scheduled backups. ESG validated that NetBackup Flex Scale continued running backups and snapshots during multiple disk failures; we also validated NetBackup Flex Scale's fast, automated node replacement, with automatic data rebuilding and cluster re-balancing. NetBackup Flex Scale's extensive automation saves administrator time and reduces errors, ensuring high availability of data.

Performance

ESG audited performance results of Veritas testing. Testing was focused on backup performance, restore performance, and instant VM access.

ESG Testing

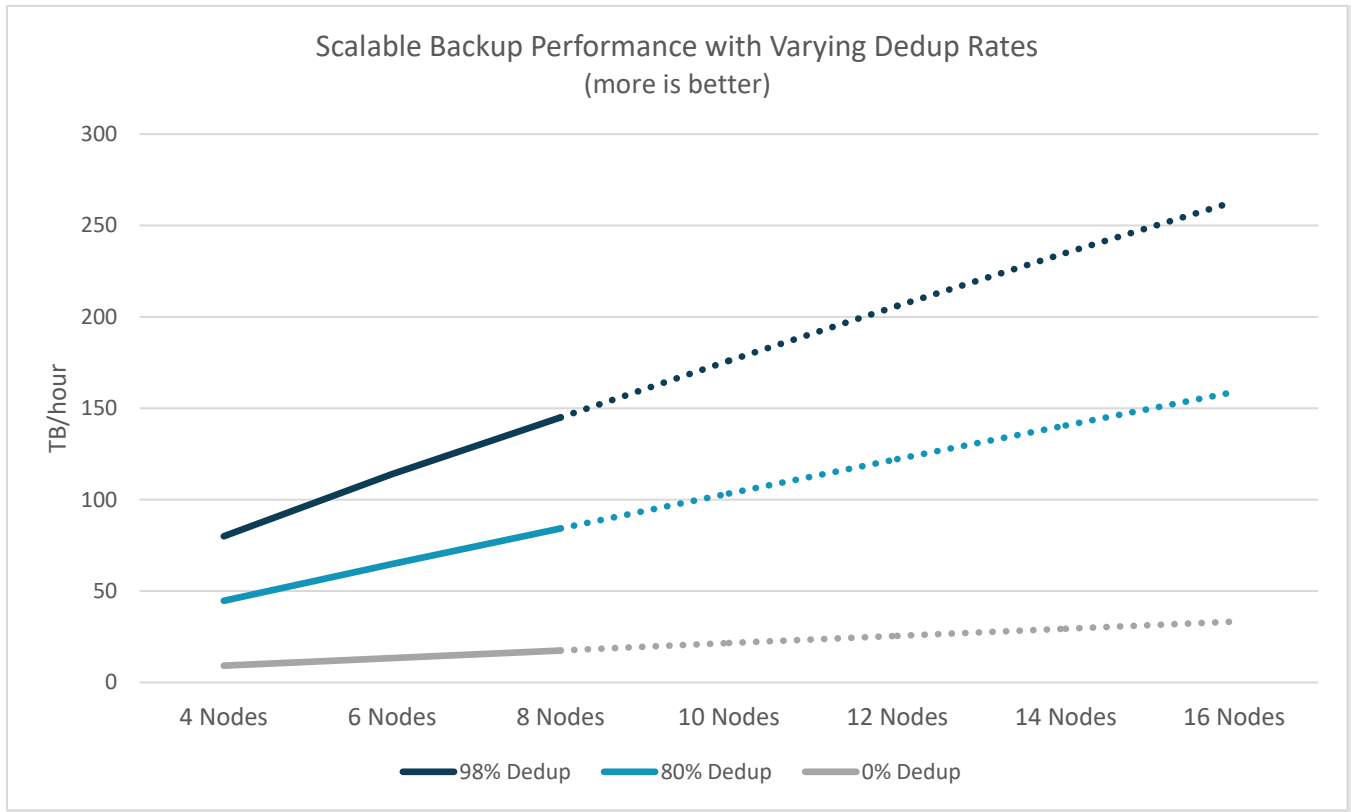
NetBackup Flex Scale's enterprise-class backup and restore performance come from numerous sources.

- Because it is a scale-out system, as nodes are added, additional disk, CPU, and network resources are added to serve I/O requests; this enhances performance by spreading the workload among more resources.
- Data is deduplicated, speeding backup performance; the higher the dedupe rate, the faster backups can be completed.
- NetBackup Flex Scale's parity calculation is fast because it uses a version of Reed-Solomon code that leverages processor optimization. As a result, only a single operation is required per processor.
- Because NetBackup Flex Scale uses block-based erasure coding (versus file- or object-based), there is no need to consume cycles and slow performance while the system reclaims space from temporary replicas or landing zones.
- Instead of using TCP, NetBackup Flex Scale uses UDP for the cluster interconnect among the commodity hardware nodes, delivering 50% more IOPS, according to Veritas. Because NetBackup Flex Scale handles network management across the cluster, there is also no need for NIC bonding, removing another performance obstacle.

ESG audited Veritas performance testing focused on backup and restore performance. The test bed included 10 RHEL clients backing up over 25 GbE to 4-8 standard NetBackup Flex Scale nodes. Veritas' Gendata tool was used to send increasing streams of 30 GB, 60GB, and 100GB files with varying deduplication rates. Testing with additional nodes was conducted, but due to lack of sufficient servers to exhaust the cluster, the 10-16-node results were extrapolated.

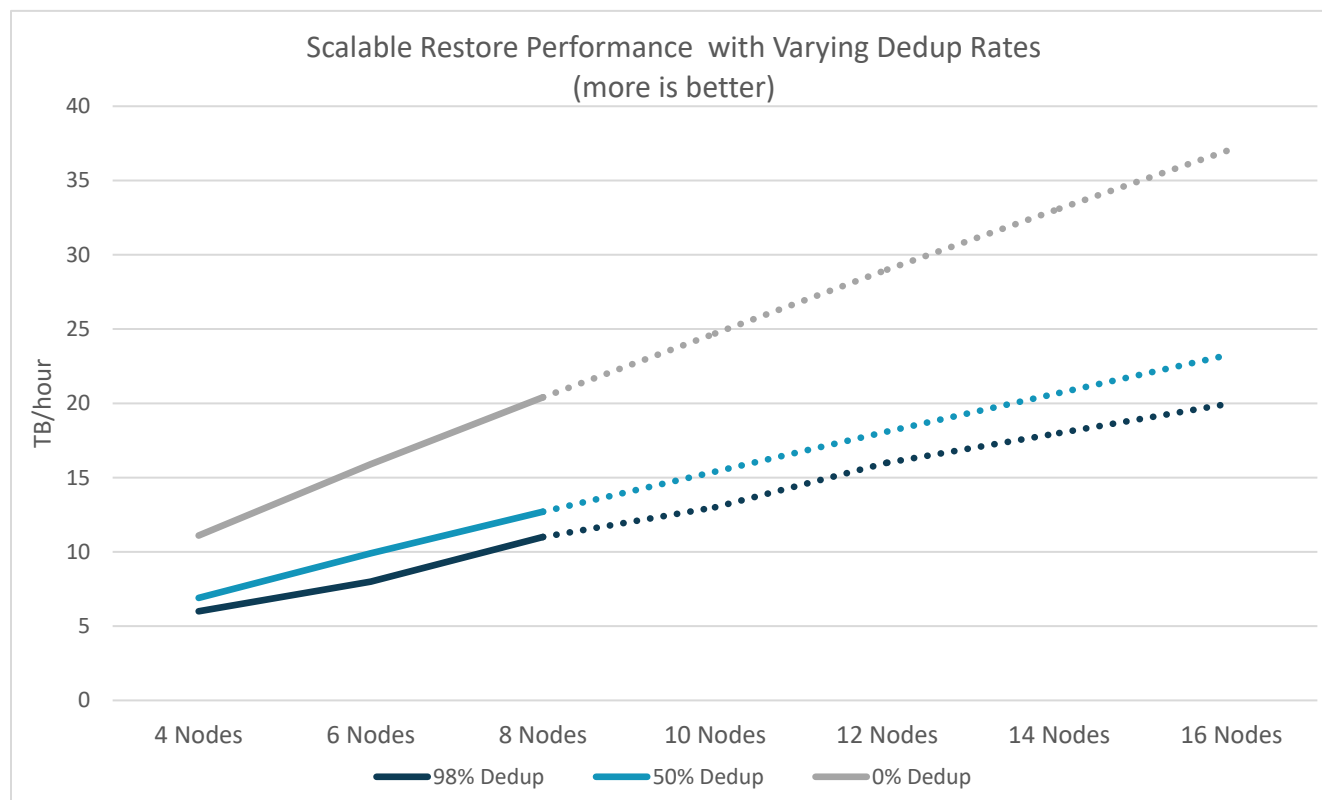
First, ESG reviewed backup performance results. For backup, the more deduplicated the data, the faster the performance. ESG validated that NetBackup Flex Scale backup performance scaled almost linearly as nodes were added. With 98% deduplicated data, performance scaled from 80 TB/hour with four nodes to 145 TB/hour with eight nodes (and up to 263 TB/hour with 16 nodes) (see Figure 7, extrapolated data in dotted lines). Even with no deduplication, performance scaled from 9 TB/hour with four nodes to 17.5 TB/hour with eight nodes (and to more than 33 TB/hour with 16 nodes).

Figure 7. Scalability of Backup Performance



Source: Enterprise Strategy Group

Next, we looked at restore performance. Restore is the reason for backup, and the faster data can be restored or made usable, the faster organizations can get back to productivity. ESG validated that NetBackup Flex Scale restore performance scaled almost linearly as the cluster was expanded. As expected, restore performance was fastest with no deduplication, ranging from 11 TB/hour with four nodes to 20.4 TB per hour with eight nodes; additional data points were extrapolated, resulting in up to 37 TB/hour with 16 nodes (see Figure 8). Deduplication takes some cycles from any system during restore, but even with data that was 98% deduplicated, NetBackup Flex Scale delivered 6 TB/hour with four nodes, scaling to 11 TB/hour with eight nodes (20 TB/hour with 16 nodes). NetBackup Flex Scale clearly delivers enterprise-class restore performance for data at all levels of deduplication.

Figure 8. Scalability of Restore Performance


Source: Enterprise Strategy Group

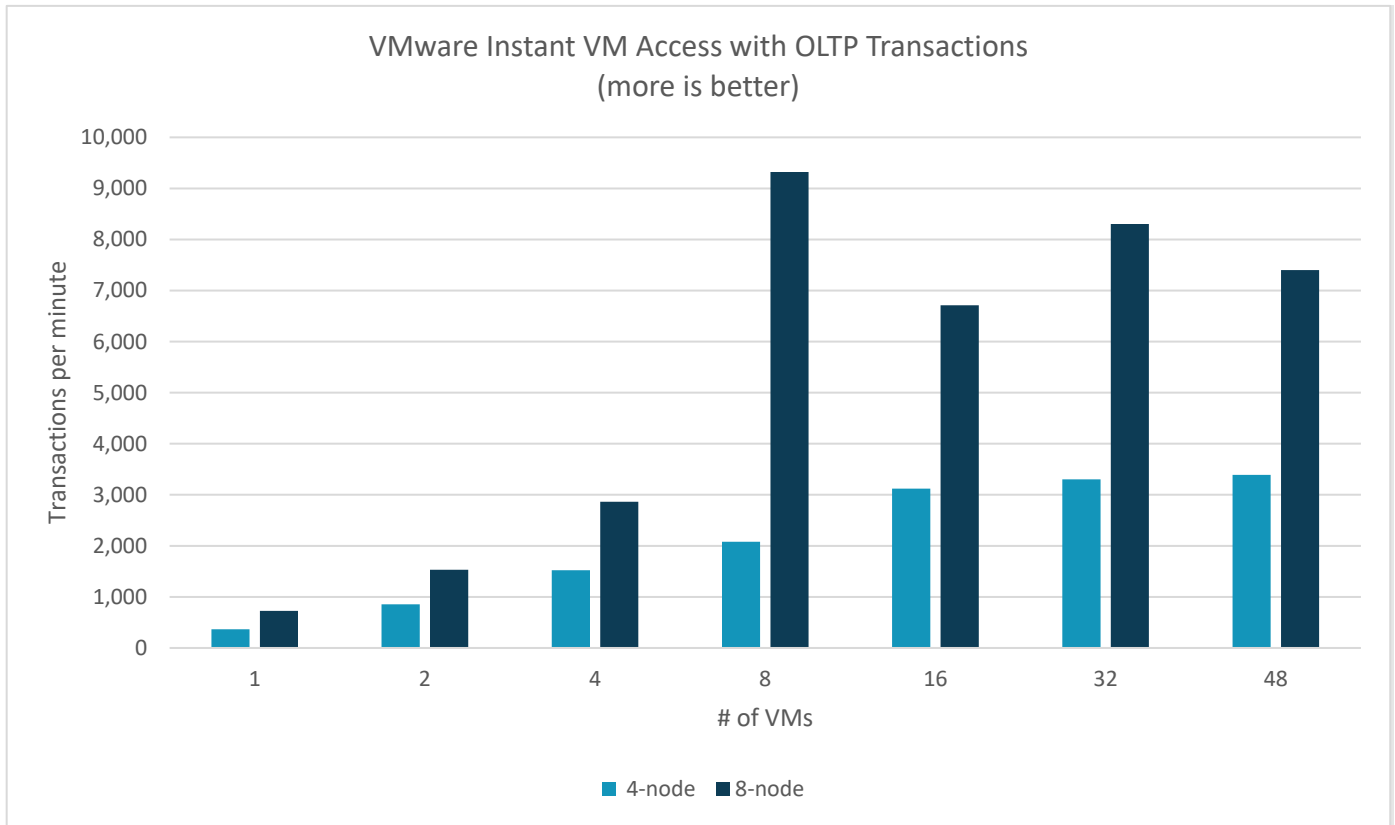
ESG also reviewed performance related to instant access virtual machines. With instant access, organizations can start up VMs on NetBackup Flex Scale to get back in operation immediately, instead of waiting for the restore to take place. This is an important feature for critical production VMs.

ESG validated Veritas testing that used 16 VMware ESXi servers with 10 GbE connections and up to eight NetBackup Flex Scale nodes. NetBackup Flex Scale was able to run VMs on the NetBackup Flex Scale node; each VM included 12 vCPUs, 24 GB memory, and was running a TPC-C-like OLTP workload.

Running on four NetBackup Flex Scale nodes, 48 instant access VMs achieved 3,389 transactions per minute (TPM); running on eight NetBackup Flex Scale nodes, 48 VMs achieved 7,401 TPM (see Figure 9). The eight-node configuration was able to deliver an average 2.66x TPM compared with the four-node.³

³ It should be noted running these tests with additional threads is likely to increase the number of instant access VMs that can be supported. This was not possible in the test bed that was used.

Figure 9. NetBackup Flex Scale High-performance Instant Access VMs



Source: Enterprise Strategy Group



Why This Matters

Simplicity and automation are important features in a growing environment, but a data protection solution also must deliver backup and restore performance at scale to keep applications both protected and available.

NetBackup Flex Scale was built to scale out, adding disk, CPU, and network resources as nodes are added. ESG validated high performance for both backup and restores, with near-linear scalability as nodes were added regardless of deduplication rates. With eight NetBackup Flex Scale nodes, ESG validated up to 145 TB/hour for backup (37 TB/hour extrapolating to 16 nodes) and up to 11 TB/hour for restores (263 TB/hour extrapolating to 16 nodes). In addition, ESG validated the ability of NetBackup Flex Scale to run 48 instant access VMs with 3,389 TPM on four nodes and 7,401 TPM on eight nodes.

The Bigger Truth

Veritas NetBackup has a long history of not only delivering rock-solid data protection, but also implementing changes along the way to optimally serve customers. With more than 800 workloads supported, NetBackup is a leader in the industry. The latest incarnation of NetBackup is NetBackup Flex Scale, a scale-out, hyperconverged appliance with all NetBackup services containerized using Docker. It is focused on simplicity and automation for all tasks, including backup, restore, and scale.

NetBackup Flex Scale automatically scales with each node addition, adding deduplicated storage capacity and increasing catalog capacity and NetBackup services, all without intervention. The intelligent load balancer ensures resource optimization across the cluster. If a node fails, the required services are automatically transferred to other nodes, and business continues uninterrupted. Erasure coding provides high availability, and reliability increases as nodes are added.

ESG validated NetBackup Flex Scale's simplicity, scalability, resiliency and automation, and performance.

- *Deployment.* Administrators only rack, cable, and power on the appliance, add a few IP addresses, and the NetBackup Flex Scale software does the rest.
- *Scalability.* Nodes are automatically discovered; when adding nodes, administrators can choose to optimize resources for backups and restores or for faster reconfiguration. NetBackup Flex Scale adds nodes and rebalances data automatically and non-disruptively.
- *Resiliency.* Backups continue during multiple disk failures; node upgrades and replacements are automated and non-disruptive, with automatic data rebuilding and cluster rebalancing.
- *Performance.* Backup and restore displayed enterprise-class performance with near linear scalability at every deduplication rate, with backup reaching 145 TB/hour with eight nodes and restore reaching 20.4 TB/hour.

Currently, the HPE ProLiant hardware is the only validated platform for NetBackup Flex Scale, offering high reliability and scale.

Veritas NetBackup Flex Scale is just what organizations need as IT becomes more complex, data volumes continue to grow, and application availability remains paramount. NetBackup Flex Scale is extremely simple, but without compromising features. It is a high-performance, enterprise-class data protection solution that delivers rock-solid protection at scale without increasing management effort or cost. With NetBackup Flex Scale, Veritas has added to its portfolio of enterprise data protection solutions, with greater flexibility and choice for customers.

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.

© 2021 by The Enterprise Strategy Group, Inc. All Rights Reserved.



www.esg-global.com



contact@esg-global.com



508.482.0188