

Technical Review

Efficiently Protect Virtual Environments with Integrated Data Protection Appliances from Dell EMC

Date: July 2020 **Author:** Vinny Choinski, Senior Validation Analyst; and Jack Poller, Senior Analyst

Abstract

This ESG Technical Review documents the results of extensive performance testing conducted on the [Dell EMC Integrated Data Protection Appliance \(IDPA\)](#) and a competitive clustered server appliance technology solution. Testing focused on the ability of each solution to efficiently protect virtual environments at scale.


The Challenges

IT is growing more demanding every year, driven by more data, more applications, more devices, more users, and the transformation of IT architecture to hybrid and multi-cloud environments. Indeed, according to ESG research, two-thirds (66%) of organizations say that IT has gotten more complex over the last two years.¹

Increasing complexity is impacting many aspects of IT operations, and data protection is no exception. When respondents to a separate ESG research survey were asked about their top current data protection challenges, the most cited response—by nearly half of organizations (49%)—was improving SLAs, RPOs, and RTOs (see Figure 1).²

Figure 1. Current Data Protection Challenges

Which of the following would you characterize as the top challenges with your organization’s current data protection processes and technologies? (Percent of respondents, N=320, three responses accepted)



Source: Enterprise Strategy Group

¹ Source: ESG Master Survey Results, [2019 Technology Spending Intentions Survey](#), March 2019.

² Source: ESG Master Survey Results, [2018 Data Protection Landscape Survey](#), November 2018.

Protecting digital assets and ensuring their availability requires proven and comprehensive storage and software tools to work in unison. Missing SLA commitments can cause significant downtime and loss of data, putting the business in jeopardy. The risk to the business is driving the demand for enhancements in data protection—48% of organizations say their IT leadership is mandating improvements in backup and recovery SLAs, RPOs, and RTOs.³

Dell EMC Integrated Data Protection Appliance

Dell EMC designed the Integrated Data Protection Appliance (IDPA) as a comprehensive backup, replication, recovery, and cloud-ready solution for physical and virtual workloads. The IDPA converges data protection storage, protection software, search, and analytics into an integrated solution.



The IDPA is available in different configurations, meeting the requirements of SMB to mid-size enterprises (DP4400 and DP5800), and large size enterprises (DP8300 and DP8800). Organizations deploying IDPA benefit from:

- Accelerated time to protect with fast deployment options.
- Integrated backup, replication, recovery, and deduplication in a single appliance.
- Private and public cloud data protection with long-term retention to cloud and cloud disaster recovery.
- VMware file-level recovery, dynamic policies for VMs, and application-consistent VMware image backups.
- Single pane-of-glass management console; management integration with management tools, such as VMware and Oracle.
- Monitoring, analytics, and reporting.
- Accelerated performance and instant recoverability with flash storage.
- Client-side deduplication for efficient use of storage.
- Scalability from 8 TB to 150 PB logical capacity and up to 1 PB of physical capacity.

Dell EMC also offers the [Future-Proof Loyalty Program](#), which provides investment protection through guaranteed data protection deduplication rates, satisfaction guarantees, cloud-enabled consumption, and predictable support pricing.

ESG Tested

ESG compared the performance of the IDPA against a competitive appliance technology solution. The Dell EMC solution used in this review was an IDPA DP5800. The competitive solution used in this review, Vendor X, was a hyperconverged appliance with a proprietary file system that enabled backup data, metadata, and protection tasks to be distributed across each server node in the architecture. Table 1 details the configuration of each solution.

Table 1. Backup Appliance Configuration

	Dell EMC IDPA DP5800	Competitor Backup Appliance (Vendor X)
Nodes	3 x Dell servers	6 nodes
Physical storage devices	42 x 4 TB 7200 RPM SAS	18 x 8 TB 7200 RPM SAS
Network	4 x 10GbE	8 x 10GbE
Backup streams	144 8 streams/proxy 18 proxies	144 24 streams/node

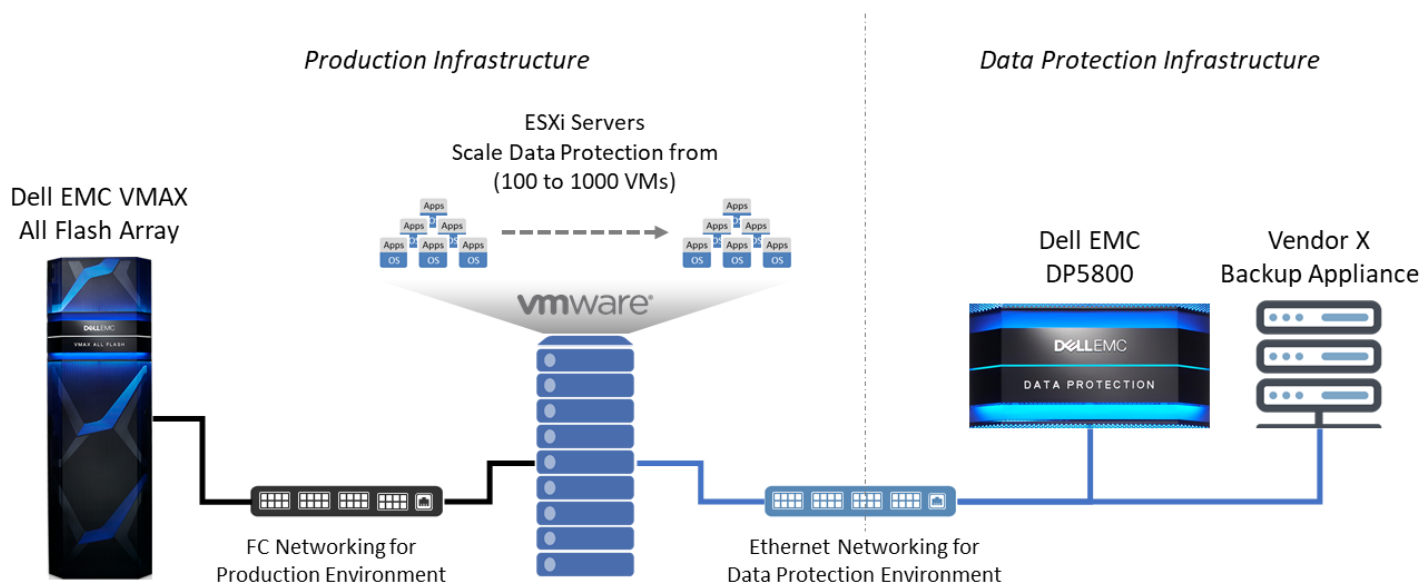
Source: Enterprise Strategy Group

ESG leveraged Dell EMC’s test environment shown in Figure 2. The production infrastructure to be protected consisted of 28 ESXi servers attached via multiple 8G FC links to a Dell EMC VMAX All-Flash Array, functioning as primary storage. The

³ *ibid.*

IDPA used four 10GbE connections and Vendor X used six 10GbE connections to connect to a dedicated network switch. The switch used four 10GbE links to connect the data protection infrastructure to the simulated virtual production environment.

Figure 2. Test Bed



Source: Enterprise Strategy Group

To demonstrate performance capabilities for data protection tasks critical to the protection of virtual environments, ESG configured the IDPA infrastructure to match the Vendor X solution as closely as possible, including throttling down the IDPA to use 144 streams to match Vendor X.

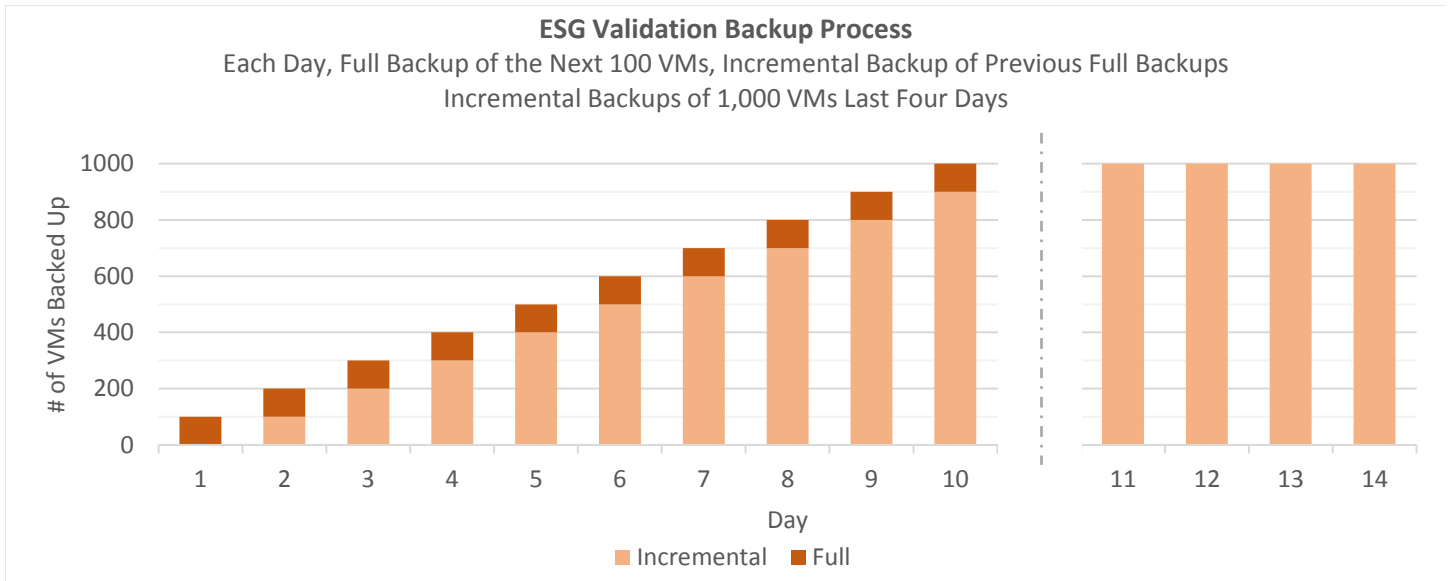
Testing Methodology

First, ESG configured the environment to be protected. We configured 1,000 virtual machines in the production environment, divided into 10 groups of 100 VMs. Each group of 100 VMs comprised 80 Windows VMs, 10 Linux VMs, and 10 MSSQL VMs. A file generation utility was used to create a 3% daily data change rate on the Windows and Linux file systems. The DBgen utility was used to create a 5% daily data change rate on each database VM.⁴ The VM image-level backups were configured to be application consistent, and the indexing policy was turned on.

Next, we devised a 14-day test strategy. On the first day, we performed a full backup of the first group of 100 VMs. On the second day, we performed an incremental backup of the first group of 100 VMs, and a full backup of the second group of VMs. On the third day, we performed an incremental backup of the first two groups of VMs, and a full backup of the third group of VMs. This sequence progressed until day 10, when we performed a full backup of the last group of VMs and an incremental backup of the first 9 groups of VMs. For each of the last four days of the test sequence, we performed an incremental backup of all 10 groups of VMs (1,000 VMs). This sequence is shown graphically in Figure 3.

⁴ DBgen is a database population program that ESG downloaded from here: <http://www.tpc.org/tpch/>

Figure 3. Validation Backup Process



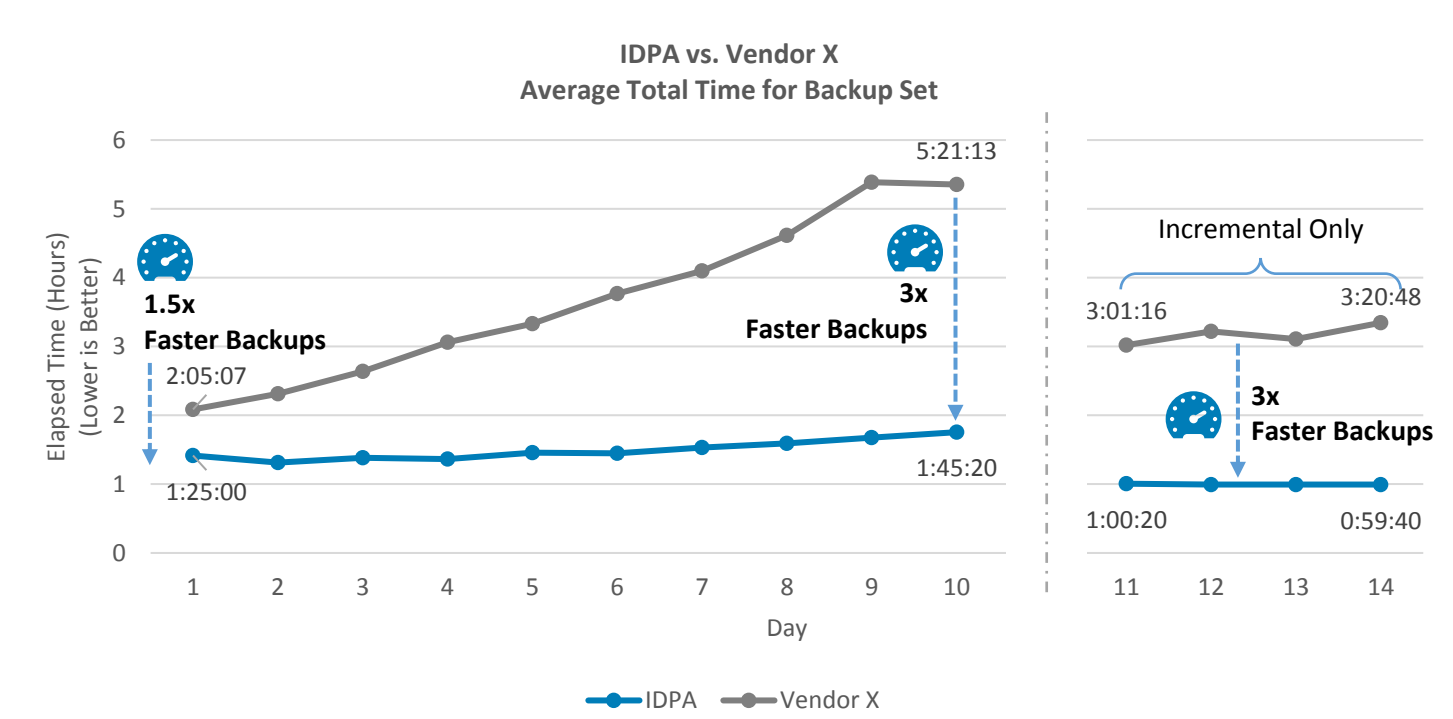
Source: Enterprise Strategy Group

The 14-day test sequence was run three times (a total of 42 days). A full set of performance metrics was collected and averaged across the three 14-day test sequences.

Backup Performance

ESG started its performance analysis by auditing the overall backup times for each solution. Figure 4 shows the average time (across the three runs) required for each day (day 1: 100 full, 0 incremental; day 2: 100 full, 100 incremental; and day 10: 100 full, 900 incremental).

Figure 4. Backup Performance



Source: Enterprise Strategy Group

What the numbers mean

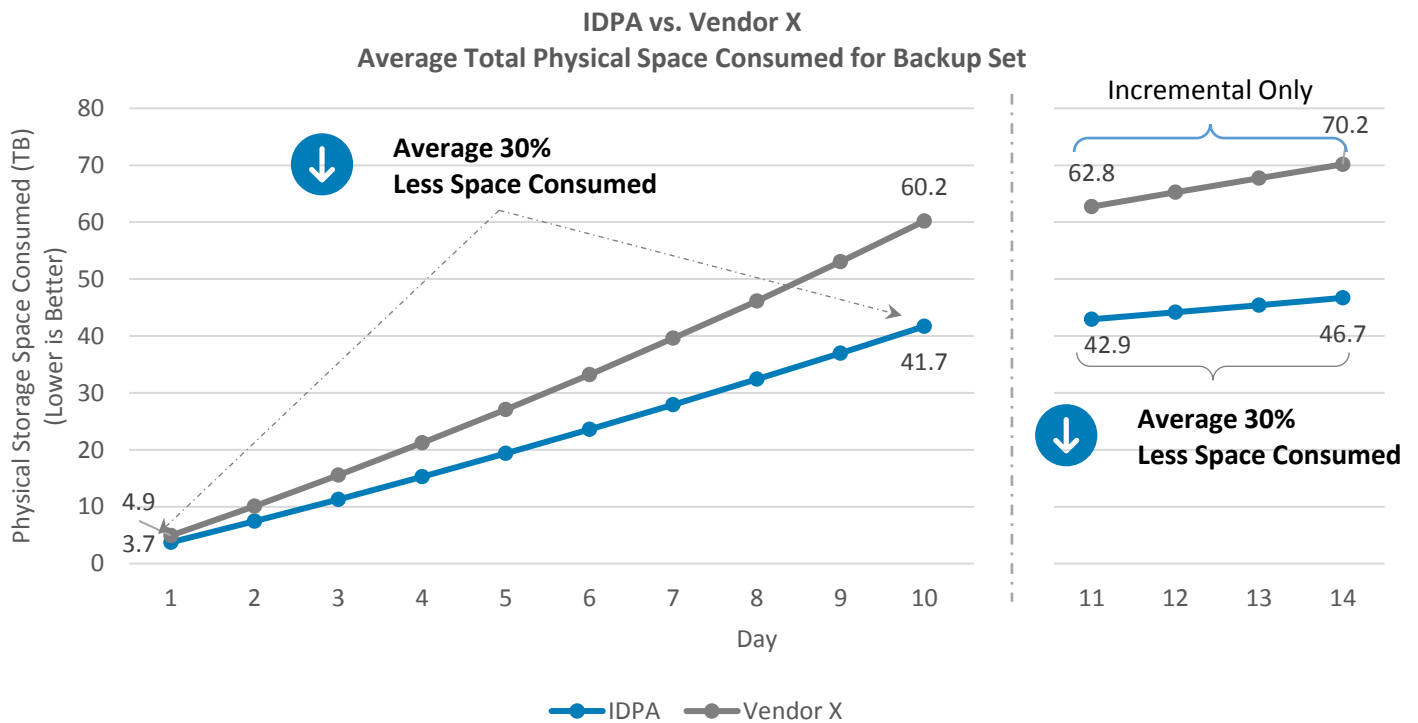
- For 100 full backups, the IDPA completed backups approximately 1.5x faster than Vendor X.
- When scaling the workload by adding incremental backups, the time to complete backups for IDPA grew 24% while the time to complete backups for Vendor X grew by 157%.
- IDPA was approximately 3x faster than Vendor X when completing the heaviest workload of 100 full and 900 incremental backups.
- IDPA was approximately 3x faster than Vendor X when completing 1,000 incremental backups.

Storage Efficiency

ESG reviewed the storage efficiency of both solutions. Figure 5 shows the amount of storage consumed by each solution after each backup run during the complete testing sequence. Here, we define storage consumed as the amount of physical storage capacity that’s used for storing backup data including protection overhead, deduplication, and compression.

ESG found that IDPA outperformed the alternative solution, consuming an average of 30% less capacity regardless of the mix of full and incremental backups. It should be noted, the IDPA solution also uses client-side deduplication so only data that has changed since the last backup is transferred over the network to the backup target to improve overall network transfer efficiency

Figure 5. Storage Efficiency: Total Space Consumed



Source: Enterprise Strategy Group

What the numbers mean

- Both solutions processed approximately 1,200 TB of backup data during testing.
- The IDPA solution consumed approximately 47 TB of capacity to store the backup data.
- The Vendor X solution consumed approximately 70 TB of capacity to store the backup data.
- The IDPA consumed an average of 30% less capacity than Vendor X.

Restore Performance

Dell EMC IDPA provides four recovery and restore options while Vendor X provides two options, as shown in Table 1.

Table 2. Restore and Recover Options

	IDPA	Vendor X
Restore	Restore VM to original VM on primary storage	Boot VM from Backup Appliance storage; move to primary storage with VMware Storage vMotion
	Restore VM to existing VM on primary storage	
	Restore VM to new VM on primary storage	
Recovery (Instant Access)	Boot VM from IDPA storage	Boot VM from Backup Appliance storage


Source: Enterprise Strategy Group

In some situations, such as a ransomware attack, an organization may need to recover a substantial portion of its environment to a point in time before the environment was compromised. ESG reviewed restore performance, comparing an IDPA restore of the original VMs to a point in time on primary storage against a Vendor X instant mass restore, which boots VMs and uses VMware Storage vMotion to move the VM to the primary storage system. ESG found that IDPA outperformed Vendor X, with all VMs restored to primary storage almost 3x faster when restoring 32 VMs, and more than 10x faster when restoring all 1,000, as shown in Figure 6.

Two factors drive Dell EMC IDPA's performance advantage. First, Dell EMC uses change block tracking (CBT) during restore. Rather than restoring the entire VM, CBT operates like snapshots, recovering just the blocks that have changed since the last backup. Restore times are accelerated since less data is transferred between the backup appliance and primary storage.

Second, Vendor X uses Storage vMotion to transfer data. Storage vMotion was designed for VM portability rather than mass VM recovery; Storage vMotion can consume significant ESXi resources, and VMware enforces limits on the number of simultaneous Storage vMotion operations.

Figure 6. Restore Performance: Time to Restore to Primary Storage



Source: Enterprise Strategy Group


What the numbers mean

- IDPA restored 32 VMs to primary storage approximately 2.9x faster than Vendor X.
- When scaling the workload by increasing the number of VMs to restore, IDPA’s performance advantage increased. IDPA restored 100 VMs 9x faster than Vendor X and restored 1,000 VMs 10x faster.

To recover VMs back to production storage, Vendor X requires that you first boot the VMs on the backup appliance. ESG measured the elapsed time for Vendor X to boot VMs from the backup appliance storage.

Using IDPA, organizations can skip the intermediate step of first booting VMs before restoring to primary storage. As shown in Figure 7, the elapsed time for Vendor X *just* to boot VMs from the backup appliance was comparable to the elapsed time for IDPA to *fully restore* the VMs to primary storage.

Figure 7. Restore Performance: IDPA Time to Restore versus Vendor X Time to Boot



Source: Enterprise Strategy Group



Why This Matters

Backup and recovery performance can affect how quickly an organization can minimize disruption and return to normal business operations. Efficient, fast, and scalable data protection solutions are necessary to meet stringent SLAs, RPOs, and RTOs.

ESG validated that the Dell EMC IDPA performed better than Vendor X’s solution when backing up and recovering VMs. IDPA backed up 100 to 1,000 VMs 1.4 to 3x faster than Vendor X. Likewise, IDPA restored VMs to primary storage 3x to 10x faster than Vendor X. IDPA restored all 1,000 VMs to primary storage in approximately the same amount of time that Vendor X was able to *just* boot 1,000 VMs on the backup appliance. IDPA enables administrators to rapidly recover and boot VMs, rather than boot VMs in a shared-resource environment with degraded performance while relying on storage vMotion to move VMs to primary storage. In addition, IDPA backups were more efficient, consuming an average of 30% less physical storage than Vendor X regardless of the workload.

The Dell EMC Integrated Data Protection Appliance’s ability to do both source and target-side deduplication enables end-users to access data faster, allowing them to run applications even while backups are streaming in the background. Dell EMC shows its proven technology can enable large enterprises to back up data while minimizing disruption to ongoing business operations.

The Bigger Truth

Enterprise IT is expected to be a well-oiled machine, providing compute, network, storage, and cloud services without failures. The reality is that it's difficult to keep up with the diverse workloads and datasets in modern IT environments. Enterprises face tremendous challenges to establish effective and efficient data protection that meets stringent SLAs, RPOs, and RTOs, and scales with the demands of modern organizations.

The Dell EMC Integrated Data Protection Appliance provides a unified solution for backup, replication, recovery, and cloud readiness for physical and virtual workloads. IDPA incorporates cloud, storage, data protection software, search, and analytics to provide rapid deployment and scalability for small, medium, and large enterprises.

ESG validated the performance of the Dell EMC IDPA solution, comparing an IDPA DP5800 with a leading competitive hyperconverged data protection appliance solution from Vendor X. Despite throttling down the IDPA to use 144 streams to match Vendor X, performance testing revealed:

- Dell EMC IDPA completed backups 1.5 to 3x faster than Vendor X.
- As workload increased with additional incremental backups, the time to complete backups increased 24% for the IDPA compared with 157% for Vendor X.
- IDPA completed the heaviest workload of 100 full and 900 incremental backups 3x faster.
- IDPA completed 1,000 incremental backups 3x faster than Vendor X.
- IDPA consumed an average of 30% less capacity than Vendor X.
- IDPA restored VMs to primary storage 3x to 10x faster than Vendor X.
- Source-side and storage-side deduplication, and change block tracking contribute significantly to the IDPA advantage.

We would expect IDPA to outperform Vendor X by larger amounts without throttling.

The results presented in this review are based on testing in a controlled environment. Due to the many variables in production environments, it is important to perform planning and testing in your own environment to validate the viability and efficacy of any solution. Organizations in need of a large-scale data protection solution should thoroughly test the efficacy, functionality, and operational capabilities before purchasing or deploying any data protection solution.

ESG's testing demonstrates that the IDPA outperforms a comparable hyperconverged data protection appliance when it comes to the critical attributes of performance, scalability, and efficiency. If your organization is looking for data protection that can meet your stringent SLAs, RPOs, and RTOs, ESG believes that you should seriously consider the performance and scalability of the [Dell EMC Integrated Data Protection Appliance \(IDPA\)](#).

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. 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.

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.