

Technical Review

Supporting Mission-critical Workloads at Scale With Dell EMC PowerFlex

Date: August 2021 **Author:** Tony Palmer, Senior Validation Analyst

Abstract

This ESG Technical Review documents hands-on evaluation of the Dell EMC PowerFlex software-defined infrastructure solution. Testing was designed to illustrate how the PowerFlex software-defined infrastructure platform achieves high levels of performance, scalability, availability, and manageability.

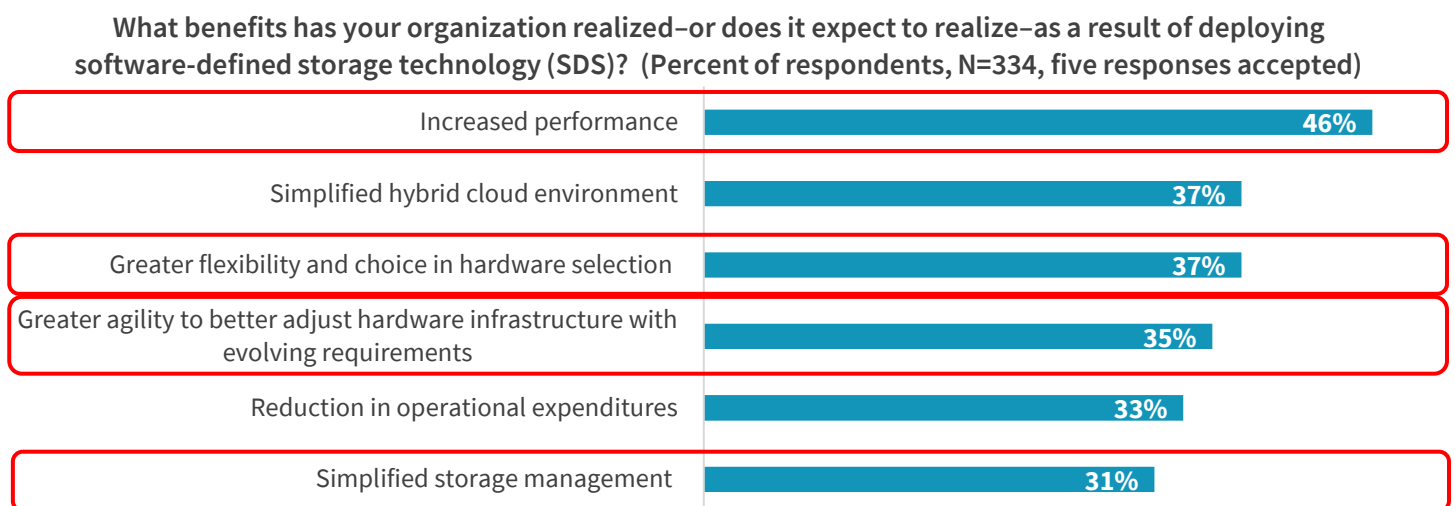
The Challenges

IT organizations are under pressure to deliver services at an unprecedented pace. To address these demands, many organizations are looking to modernize their data centers. Software-defined approaches provide a compelling choice to deliver organizational agility. They combine industry-standard hardware and software to pool and manage resources. When selecting a software-defined platform, there are some critical aspects to consider. The platform must offer extensive flexibility by supporting broad architectural choices and scaling needs. It must also deliver predictable performance to key workloads, and it must do so at scale.

An ESG research survey asked respondents about the benefits their organization has realized or expected to realize as a result of deploying software-defined infrastructure (SDI). Respondents cited increased performance (46%), greater flexibility and choice in hardware selection (37%), greater agility to better adjust hardware infrastructure with evolving requirements (35%), and simplified storage management (31%) (see Figure 1).¹

When asked about their organization’s data center storage technology spending over the next 24 months, 69% of respondents stated they expected to accelerate their spending on hyperconverged infrastructure (HCI) and 60% stated they expected to accelerate their spending on software-defined storage (SDS) technology.²

Figure 1. Top Six Benefits Realized or Expected to Be Realized as a Result of Deploying SDS



Source: Enterprise Strategy Group

¹ Source: ESG Master Survey Results, [2019 Data Storage Trends](#), November 2019.

² Ibid.

The Solution: Dell EMC PowerFlex

PowerFlex is designed to deliver flexibility, elasticity, and simplicity with predictable performance and resiliency at scale by combining compute and high-performance storage resources in a managed unified fabric. PowerFlex offers a simple and comprehensive toolset for IT operations and lifecycle management, helping to automate infrastructure workflows. It is ideal for high-value databases and workloads, agile cloud-native containerized workloads, and heterogeneous workload consolidation.

Dell Technologies engineered PowerFlex to deliver extreme performance at sub-millisecond response times. Aggregation of large sets of resources across many nodes is leveraged to deliver heterogeneous workload consolidation consistently and predictably. The platform can be deployed with as few as four nodes but can scale to thousands of nodes. A self-healing architecture is implemented to help organizations get closer to non-stop operations. Dell Technologies reports 99.9999% observed availability, a claim ESG examined for this report. PowerFlex offers flexible x86 node-based deployment configurations and scaling options with the ability of scaling compute and storage resources independently or together.

The PowerFlex platform is capable of hosting a variety of operating environments—ranging from bare-metal operating systems and multiple hypervisors to multiple container management platforms simultaneously—without creating silos of infrastructure. This helps organizations to consolidate and simplify disparate heterogeneous workloads and modernize applications at their pace on a common platform.


PowerFlex Manager offers tools for IT operations and lifecycle management that automate infrastructure workflows from BIOS and firmware to nodes, hypervisors, and networking. PowerFlex Manager manages all components and operations that support the PowerFlex deployment. The platform also provides open REST APIs that can help simplify application and DevOps workflows.

PowerFlex offers high availability with quick rebuilds, native data replication and snapshots, integrated hardware-based encryption, and data reduction. These services further simplify how administrators manage, protect, and secure data.

With the PowerFlex family, businesses have choice and flexibility in how they choose to consume the PowerFlex architecture:

- PowerFlex appliance allows customers the flexibility and savings to adapt their own networking. With PowerFlex appliance, customers can benefit from a smaller starting point with massive scale potential, without having to compromise on performance and resiliency.
- PowerFlex rack is a rack-scale engineered system with integrated networking, designed to simplify deployments and expedite time to value. A white glove deployment service ensures a complete turnkey experience while the Release Certification Matrix (RCM) further simplifies upgrades, keeps systems stabilized and optimized, and removes the challenge of self-testing all firmware and software.

Figure 2. Dell EMC PowerFlex Solution Overview



Source: Enterprise Strategy Group

ESG Tested

ESG completed a technical review of the Dell EMC PowerFlex solution with a focus on performance, scalability, availability, and manageability. This review was based on the following use cases:

- Databases – Traditional & modern/NoSQL
- Analytics – Traditional & cloud-native
- Container management and automation platforms

Performance Overview

PowerFlex is designed to deliver high performance with enterprise-class resiliency. PowerFlex delivers scale-out storage services by pooling resources from a large number of nodes. Data is distributed across all available nodes, with multiple high-performance ethernet connections to each of the nodes. With large pools of resources, uniform data distribution, and no network bottleneck, PowerFlex can provide IOPS and throughput performance that scales linearly with additional resources and nodes. This is essential for I/O-intensive- performance- and latency-sensitive workloads along with throughput-intensive analytics ingest workloads. PowerFlex also allows easy isolation of key workloads with multiple protection groups, which helps ensure predictable and uninterrupted performance of high-value, performance-hungry workloads while ensuring enterprise resiliency and availability.

ESG audited testing of PowerFlex in multiple application environments using benchmark tools designed to simulate real application workloads in those environments. Oracle RAC and Microsoft SQL Server were used to validate transactional database workloads using the HammerDB tool, Elastic Stack was used to validate big data analytics ingest and indexing workloads using the Elastic Rally tool, and Cassandra DB with Kubernetes was used to validate stateful cloud-native application workloads using the Cassandra stress tool.

High-level results of testing are summarized in Table 1. These tests were not designed to push PowerFlex to its limits, but to show achievable performance using realistic application workloads with sub-millisecond response times.

Table 1. Dell EMC PowerFlex Performance

Workload	Test Harness	Configuration	Performance
Oracle RAC	HammerDB	8-node cluster in 8RU	10 million TPM with average latency of <1 ms
Microsoft SQL Server 2019	HammerDB	8-node cluster in 8RU	7 million TPM with average latency of <1 ms
Elastic Stack	Elastic Rally	7-node cluster in 7RU	1 billion indexing events in 3 hours
MongoDB with Kubernetes	Yahoo! Cloud Serving Benchmark (YCSB)	8-node cluster in 8RU	2.2 Million Operations per second per node with average latency of <1 ms (95% reads)
Cassandra DB with Kubernetes	Cassandra stress tool	8-node cluster in 8RU	302,000 operations per second with average latency of <.5ms

Source: Enterprise Strategy Group

It is important to note that all of these results were obtained with small clusters of 6-8 nodes with basic configurations containing SAS SSDs and modest CPU and RAM, dedicated to the workload that was being tested. Most organizations deploying PowerFlex will have much larger clusters to support multiple workloads and PowerFlex supports higher performance Intel Optane and NVMe drives.

Scalability Overview


PowerFlex systems can be deployed and scaled utilizing storage, compute, or HCI nodes, allowing enterprises to scale compute and storage resources together or independently without restrictions. The nodes are available in broad resource configuration options, helping IT organizations meet the needs of a wide range of workloads. The flexibility to utilize and mix compute, storage, and HCI nodes enables organizations to meet the most stringent performance requirements while minimizing software licensing expenses that are often tied to CPU cores.

With PowerFlex, resources such as storage and compute can be scaled together or separately, non-disruptively, and in small increments. The system can scale from just a few nodes to hundreds in a cluster, linearly scaling I/O performance and throughput.

To test the linear scalability of the platform, a series of benchmark workloads were run against clusters of increasing size. We used nodes with 48 CPU cores, 384GB RAM, four 7.5TB NVMe SSDs, and two 25GbE network ports each. This configuration ensures that the nodes are not CPU, disk, or network bound. At larger block sizes, the network can become a bottleneck, so PowerFlex nodes now have *four* 25GbE ports. For IOPS, a random workload with 4KB block size was used. For bandwidth, we used a sequential workload with a 256KB block size.

Figure 3 and Figure 4 show 4KB read and write IOPS performance as the cluster was scaled from eight nodes to 128 nodes. In the following charts, the percentage above the columns is the scaling factor, which was calculated as the percentage of perfect scaling (100%) as node count was doubled.


Figure 3. Dell EMC PowerFlex Scalability—Read IOPS



Source: Enterprise Strategy Group

Both read and write IOPS scaled nearly linearly from eight to 128 nodes, with no indication of flattening or falling off, and exceeding 100% in one test run.


Figure 4. Dell EMC PowerFlex Scalability—Write IOPS



Source: Enterprise Strategy Group

Figure 5 and Figure 6 show read and write throughput performance as the cluster was scaled from eight nodes to 128 nodes.


Figure 5. Dell EMC PowerFlex Scalability—Read Throughput



Source: Enterprise Strategy Group

Again, the percentage above the columns is the scaling factor, which was calculated as the percentage of perfect scaling (100%) as node count was doubled. As Figure 5 and Figure 6 show, the system scaled read and write throughput nearly linearly from eight to 128 nodes, with no indication of flattening or falling off, exceeding 100% in two test runs.

Figure 6. Dell EMC PowerFlex Scalability—Write Throughput

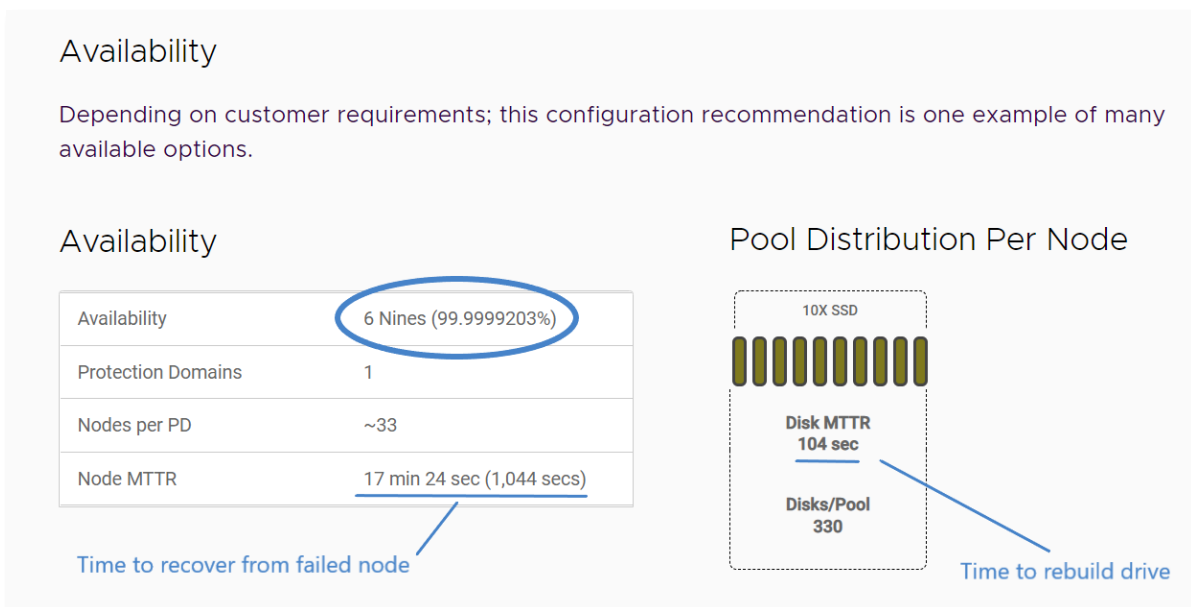


Source: Enterprise Strategy Group

Availability Overview

PowerFlex offers flexible availability options based on the configuration of the system. Mission-critical and high-value applications can be afforded extremely high availability far beyond the traditional metric of five nines—which is equivalent to eight minutes of downtime per year. Organizations commonly configure PowerFlex for six nines—31.5 seconds of downtime per year—as seen in Figure 7, but PowerFlex can be configured for even higher levels, with commensurately lower downtime per year. Eight nines, or 99.999999%, for example, equals just 315 *milliseconds* of downtime per year. PowerFlex has built-in redundancy in the following areas: POUs, aggregation switches, access switches, network connections, management control plane, node components (power supplies, fans, NIC cards, ports on the NICs), and data mirroring.

Figure 7. Dell EMC PowerFlex Availability Configuration




Source: Enterprise Strategy Group

PowerFlex is a distributed storage system, so availability calculations can be quite complex. The following is a high-level description of how it works. Dell Technologies takes into consideration multiple factors, including the mean time to recover (MTTR), write endurance, the throughput of the drives being used, the number of drives participating, the reliability of the server nodes housing those drives, the number of server nodes contributing, and the throughput of the network over which the nodes are communicating. The time to rebuild data in the cluster can either be disk limited or network limited, which is why it's important to factor in the throughput capabilities of both.

Parallel rebuild time is a significant part of Dell Technologies' availability and reliability considerations. For example, a 3.84TB drive in a 33-node pool rebuilds in just one minute 44 seconds, which means the data on the drive will likely be rebuilt before an admin has a chance to react to the failure. Dell EMC PowerFlex distributed mesh mirroring is unlike RAID protection schemes. As soon as the data is re-protected in both copies, the system data is back to full health. The protection scheme does not require replacing the drive or node to sustain additional drive failures, although organizations will want to replace failed drives or nodes to maintain capacity and I/O.

Full-node rebuild testing was performed on systems ranging from eight to 128 nodes. Each node had a 9TB volume mapped to it, filled with data, and protected with distributed mesh mirroring. The full node rebuild numbers in Figure 8 show that PowerFlex rebuild times scale with cluster size.

Figure 8. Dell EMC PowerFlex Availability Scaling—Full Node Rebuild Rate



Source: Dell Technologies

The full node rebuild times observed by ESG show that while the 18 terabytes of combined primary and distributed mesh mirrored data rebuilt in 1,349 seconds—22.5 minutes—on an eight-node cluster, it took less than two minutes—117 seconds—on a 128-node cluster.

Table 2. Dell EMC PowerFlex Performance

Number of Nodes	Rebuild Time in Seconds (Minutes)	Throughput in GB/Sec
8	1,349 (22.5)	13.7
16	641 (10.7)	22.8
32	338 (5.6)	54.5
64	185 (3.1)	99.6
128	117 (1.95)	157.5

Source: Dell Technologies

Manageability Overview

PowerFlex Manager provides deployment capabilities through the use of standardized templates aligned with hyperconverged, storage-only, and compute-only nodes. Health status and alerting is achieved via Secure Remote Services (SRS) or configuring an alert connector to send email alerts. PowerFlex Manager uses the Release Certification Matrix (RCM) or Intelligent Catalog (IC) to track and remediate drift management for better compliance. PowerFlex has maintenance modes that allow for servicing the system online. The system can also be expanded by duplicating existing services, matching the existing configuration. PowerFlex Manager automates complex processes, such as initial platform deployment. To quantify this, ESG compared the administrator time of manual and automated deployments of a six-node PowerFlex cluster with three compute-only nodes and three storage-only nodes. Administrator time refers to the periods in which the admin is required to be present for screen input and process management.

Table 3. Dell EMC PowerFlex Manager Deployment Efficiency


Task	Manual Deployment Administrator Time	Automated Deployment Administrator Time
Initial Storage Node Setup	1 hour, 25 minutes	1 hour, 30 minutes
Create Three-node Storage Service	5 hours, 25 minutes	10 minutes
Post-installation Process (Storage)	50 Minutes	45 minutes
Initial Compute Node Setup	1 hour	0 minutes
Create Three-node Storage Service	4 hours, 45 minutes	10 minutes
Post-installation Process	10 minutes	20 minutes
Total	13 hours 35 minutes	2 hours 55 minutes

Source: Dell Technologies

It's important to note that while administrator time in front of a keyboard is reduced by more than 10 hours in his scenario, the total time for the entire process is also reduced—by more than four hours—thanks to the fact that PowerFlex manager can complete the automated tasks faster than a human admin could.

PowerFlex can support a broad set of operating environments and application architecture—ranging from bare-metal workloads to virtualized applications on multiple hypervisors and containerized cloud-native applications—all on a single platform simultaneously. Leveraging PowerFlex, organizations can modernize all their applications, support all workloads, and evolve application architecture in a way that makes sense to an organization's goals and priorities.

Figure 9. Dell EMC PowerFlex Flexibility



Source: Enterprise Strategy Group

When performing maintenance based on RCMs or ICs, PowerFlex Manager supports upgrading all major components, including BIOS, firmware, drivers, SAN software (NX-OS), VMware ESXi, PowerFlex software, and CloudLink.

i Why This Matters

According to ESG research, respondents cited that hardware costs (30%); data protection (27%); and management, optimization, and automation of data placement (24%) were the most common challenges in their organization’s on-premises block storage environment.³

ESG validated that the Dell EMC PowerFlex solution delivers high performance with enterprise-class resiliency for databases, analytics, and containerized applications. Testing revealed that entry-level PowerFlex clusters delivered millions of IOPS and transactions per minute running Oracle RAC, Microsoft SQL Server, Elastic Stack, and Cassandra on Kubernetes with average latency below 1 millisecond across the board. With PowerFlex, resources such as storage and compute can be scaled together or separately, non-disruptively, and in small increments. The system can scale from just a few nodes to hundreds of nodes in a cluster, linearly scaling I/O performance and throughput. PowerFlex was shown to deliver configurable availability for high-value workloads far beyond the traditional measure of “five nines” with multiple protection groups and fast rebuilds. PowerFlex Manager provides simplified storage management, including deployment capabilities through the use of standardized templates aligned with hyperconverged, storage-only, and compute-only offerings.

While no system ESG has tested has achieved perfect linear scalability, PowerFlex comes quite close and exceeded 100% in more than one instance for both IOPS and throughput. Leveraging software-defined approaches to deliver performant, scalable, and highly available infrastructure for high-value workloads continues to be an IT priority. ESG validated that Dell EMC PowerFlex SDI (software-defined infrastructure) delivers flexibility, performance, scalability, availability, and manageability—from just a few nodes to massive scale.

³ Source: ESG Research Report, [Data Storage Trends in an Increasingly Hybrid Cloud World](#), March 2020.

The Bigger Truth

Many IT organizations continue to modernize their data centers. Software-defined approaches are a popular choice, but based on ESG research, users expect software-defined technology to realize increased performance, greater flexibility and choice in hardware selection, greater agility to better adjust hardware infrastructure with evolving requirements, and simplified storage management.

Dell EMC PowerFlex is designed to deliver flexibility, elasticity, and simplicity with predictable performance and resiliency at scale. PowerFlex offers a simple and comprehensive management toolset to help automate infrastructure workflows. This makes it a great fit for high-value databases and workloads, agile containerized cloud-native deployments, and heterogeneous workload consolidation. In an ESG research survey, organizations were asked to identify the benefits they realized or expected to realize as a result of deploying software-defined infrastructure. The most cited benefits/expectations included performance (46%), simplification of their hybrid cloud environment (37%), greater flexibility and choice in hardware (37%), and greater agility to better adjust infrastructure with evolving requirements (35%).⁴

ESG's analysis revealed that PowerFlex provides these benefits and more—it is highly performant, scales nearly linearly, is simple to manage, and lets organizations start small and scale as needed. On top of this, PowerFlex also provides the availability that mission-critical storage environments demand.

ESG validated extremely high performance—millions of IOPS and transactions per minute at sub millisecond response times, massive scalability, and mission-critical availability appropriate for environments running high-value transactional and analytics workloads on Oracle, Microsoft SQL Server, Elastic Stack, MongoDB, and Cassandra.

The consistency of performance and scalability across multiple applications and data protection activity was particularly notable. In addition, its flexibility enables organizations to adapt quickly to changing requirements, as today's environments demand.

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 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.

ESG's technical analysis revealed that the Dell EMC PowerFlex software-defined solution achieves consistently high levels of performance, scalability, availability, and simple manageability from small deployments to hundreds of nodes. If your organization is looking for a flexible software-defined solution that delivers predictable performance and simple manageability to large-scale, mission-critical workloads, ESG recommends considering Dell EMC PowerFlex.

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.

⁴ Source: ESG Master Survey Results, [Data Storage Trends](#), November 2019.