

ESG Lab Validation

Dell EMC Isilon All-Flash

Scale-out All-flash Storage for Demanding Unstructured Data Workloads

By Tony Palmer, Senior Lab Analyst

June 2017

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

Contents

Introduction	3
Background	3
Dell EMC Isilon F800 All-Flash Overview	4
ESG Lab Validation Testing	5
Performance and Scalability	5
The Bigger Truth.....	9
Appendix	10

ESG Lab Reports

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

Introduction

This Lab Validation Report documents the results of recent hands-on testing of [Dell EMC Isilon All-Flash](#) storage. Testing focused on the platform’s performance and scalability and how it enables organizations to support more applications and workloads with flash based on their ever-increasing volumes of unstructured data.

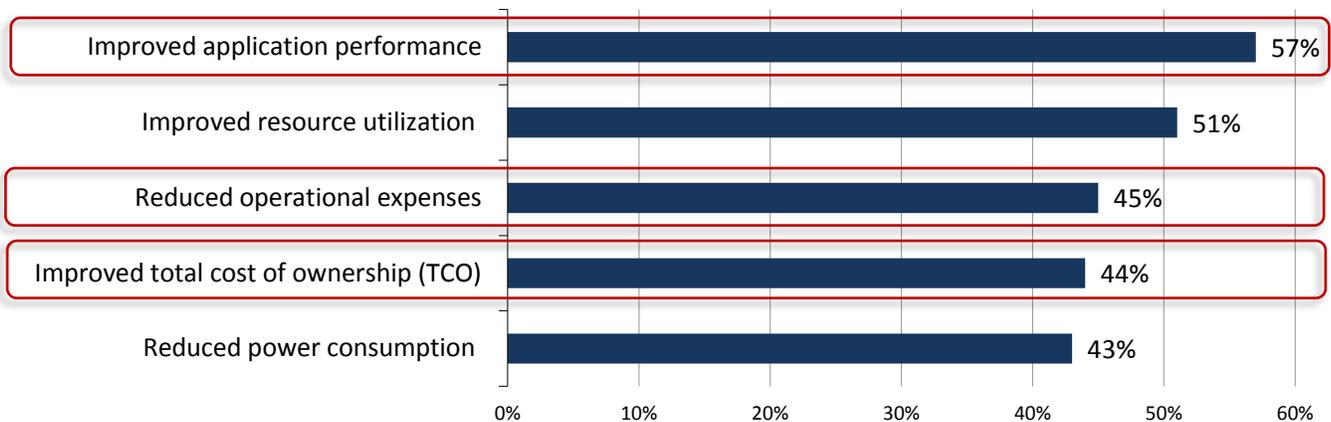
Background

With the emergence of flash storage, IT infrastructure no longer must depend on mechanical, spinning disks for its storage needs. IT organizations can now enjoy a wealth of benefits that help them keep pace with the ever-increasing demands being placed on the data sets that they store and manage.

In ESG’s *2017 IT Spending Intentions Survey*, IT professionals noted that they expect to make significant investments in solid-state/flash storage technology during the current year. As a matter of fact, solid-state/flash storage technology is among the top five areas in which IT professionals said they will make the most significant investments as part of the effort to modernize their data centers.¹ Another ESG research survey asked about the benefits that IT professionals realized from deploying solid-state/flash storage technology. Improved application performance, reduced operational expenses (OpEx), and improved total cost of ownership (TCO) were all in the top five most-cited responses, as seen in Figure 1.²

Figure 1. Top Five Benefits of Deploying Solid-state Storage

Which of the following benefits has your organization realized as the result of deploying solid-state storage technology? (Percent of respondents, N=181, multiple responses accepted)



Source: Enterprise Strategy Group, 2017

Over 50% of respondents noted that they want to increase their application performance, which is not surprising given that the performance advantages of flash storage are well known. However, the industry has enjoyed a substantial decline in flash storage prices recently, making the decision to invest in flash storage much easier. Additionally, the impact that flash storage has on the surrounding data center ecosystem is possibly more profound. Since the data storage component is eliminated as the performance bottleneck, the remaining ecosystem elements spend less time waiting for the data storage to respond. Reducing wait times enables the applications, servers, and network all to become more efficient. This increase to overall efficiency reduces spending on other IT ecosystem components (such as server hardware or software licenses) while achieving the same performance, further reducing the TCO of managing the data center.

¹ Source: ESG Research Report, [2017 IT Spending Intentions Survey](#), March 2017.

² Source: ESG Research Report, [2015 Data Storage Market Trends](#), October 2015.

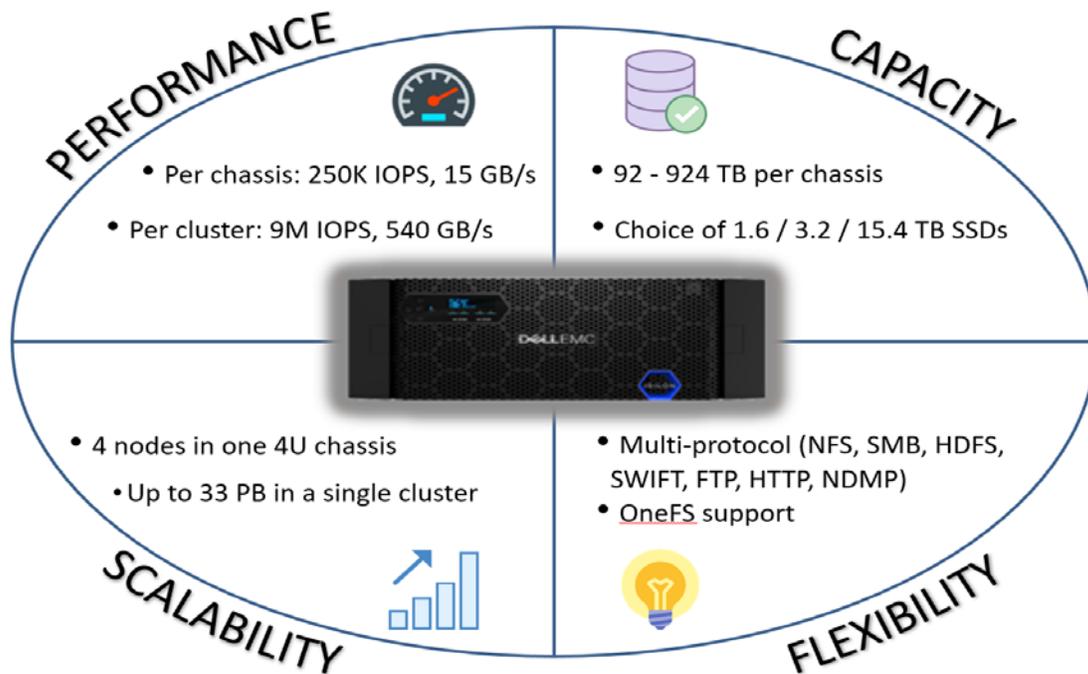
Recently, businesses have been finding new opportunities for high-performance data access. Organizations want to leverage all-flash technology to garner near real-time insights from unstructured data sets with business intelligence or big data analytics, but this requires a system designed specifically for the needs of unstructured data. Similar demand for low-latency access to unstructured content is also seen across industries such as media and entertainment, electronic design automation (EDA), and life sciences.

To effectively serve unstructured data-based workloads, storage systems must deliver substantial capacity scaling to keep pace with the rapid rate of unstructured data growth. Despite the TCO savings resulting from flash storage deployments mentioned earlier, deploying large storage capacities with all-flash arrays designed for structured or block-based data workloads will most likely require substantial upfront capital costs; doing so is simply not feasible for many organizations. In response to this growing demand for high-performance access to unstructured or file-based content, Dell EMC offers the Isilon F800 All-Flash platform. The storage system is designed to leverage the declining costs of flash storage components and architected to deliver the enterprise file system functionality and scale necessary to serve high-performance access to large capacity file storage pools.

Dell EMC Isilon F800 All-Flash Overview

In late 2016, Dell EMC introduced Isilon All-Flash, designed as scale-out NAS storage for customers seeking high levels of performance, capacity, and scalability. This storage targets customers that seek to scale out storage to meet high-performance demands, specifically for applications that deliver near real-time insights, yet does so cost-effectively and as needed with data growth. Figure 2 highlights some of the key features that Dell EMC has designed into the Isilon F800 All-Flash scale-out NAS storage system.

Figure 2. Dell EMC Isilon F800 All-Flash Overview



Source: Enterprise Strategy Group, 2017

The key features to note about Dell EMC Isilon All-Flash include:

- **Support for High-performance Workloads:**

Dell EMC has designed its Isilon F800 All-Flash NAS storage to support customers in storing—and processing in near real time—their growing pools of unstructured data. Each chassis (four nodes) serve up to 250,000 IOPS, as well as more than 15 GB/sec of throughput, and the cluster can scale to 33 PB of solid-state storage.

- **High Capacity:**

To ensure that customers can take advantage of large capacity storage without adding onto the additional burden of infrastructure—power, cooling, and floor space costs—a customer can begin with 92 TB in a single 4U chassis. The high density enables organizations to reap the benefits of all-flash storage while significantly reducing the footprint required for the initial deployment. With the choice of 1.6, 3.2, and 15.4TB SSDs available, one four-node chassis can support up to 924 TB, enabling organizations to serve large data sets from a single Isilon F800 All-Flash chassis, without compromising on performance.

- **Granular Scalability:**

Although a single chassis of Isilon F800 All-Flash can support up to 924 TB, a customer can begin with as little as 96 TB. The Isilon F800 All-Flash architecture supports the ability to scale the system in increments of four-node chassis today. Each chassis can support from 96 to 924 TB. This granularity reduces the initial capital cost required to start leveraging flash storage and eases the ability to scale over time as incremental capacity can be deployed at much smaller capacity points, reducing overall TCO as the system scales.

- **Flexibility:**

Dell EMC enables customers to support multiple file protocols to carry the transfer of many types of unstructured data to and from storage, from NAS protocols like NFS and SMB, to FTP, NDMP, Hadoop Distributed File System (HDFS), and object protocols like REST and OpenStack Swift. In addition, Isilon F800 All-Flash integrates with existing Isilon clusters and automatically tiers data across Isilon F800 All-Flash nodes, Isilon SAS HDD nodes, Isilon SATA HDD nodes, Elastic Cloud Storage (ECS) object storage, and public cloud.

Isilon F800 All-Flash enjoys the same benefits of running OneFS as Isilon's disk-based offerings. Customers are assured that all enterprise features also apply to data stored in flash, including encryption and security features, automated migration across storage tiers and the cloud, and a full suite of resiliency features. Additionally, customers can move data to the Isilon F800 All-Flash and transition its storage infrastructure as they see fit. Potential savings are substantial due to the complexity and cost associated with migrating large capacity file storage environments.

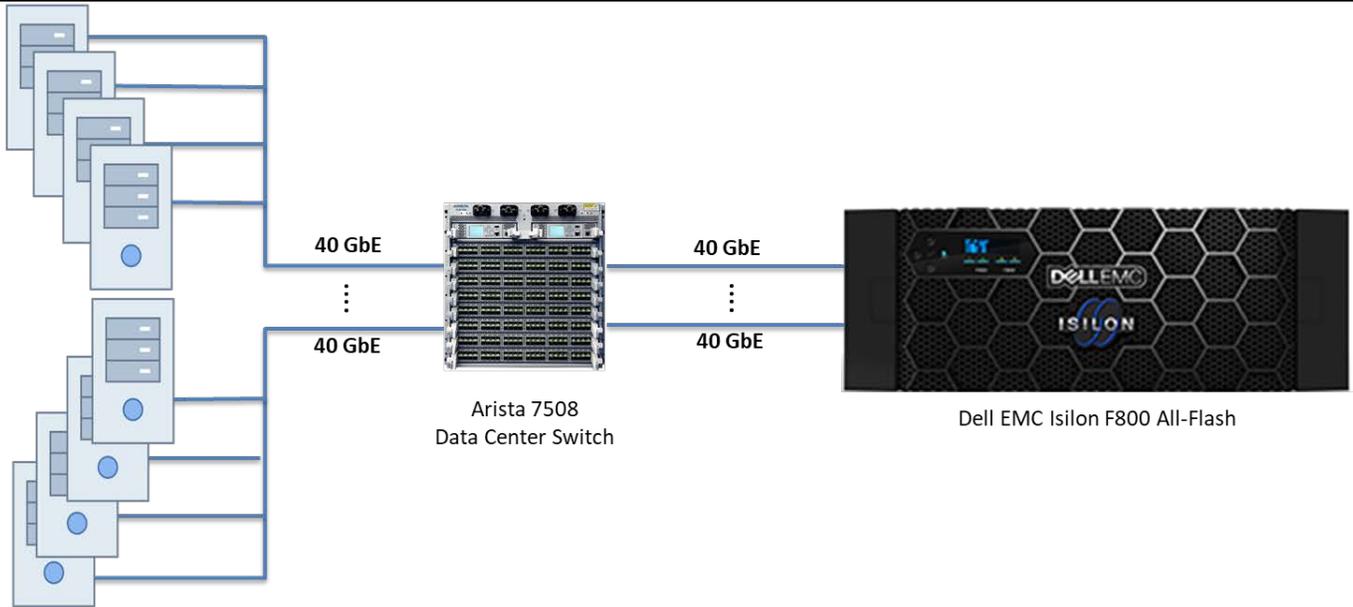
ESG Lab Validation Testing

ESG Lab performed hands-on evaluation and testing of the Isilon F800 All-Flash at a Dell EMC facility. Testing was designed to validate the performance and analyze the scalability of Isilon F800 All-Flash in unstructured data environments.

Performance and Scalability

ESG Lab began with a test bed built around a Dell EMC Isilon F800 All-Flash four-node chassis running OneFS 8.1.0, as seen in Figure 3.

Figure 3. The ESG Lab Validation Test Bed



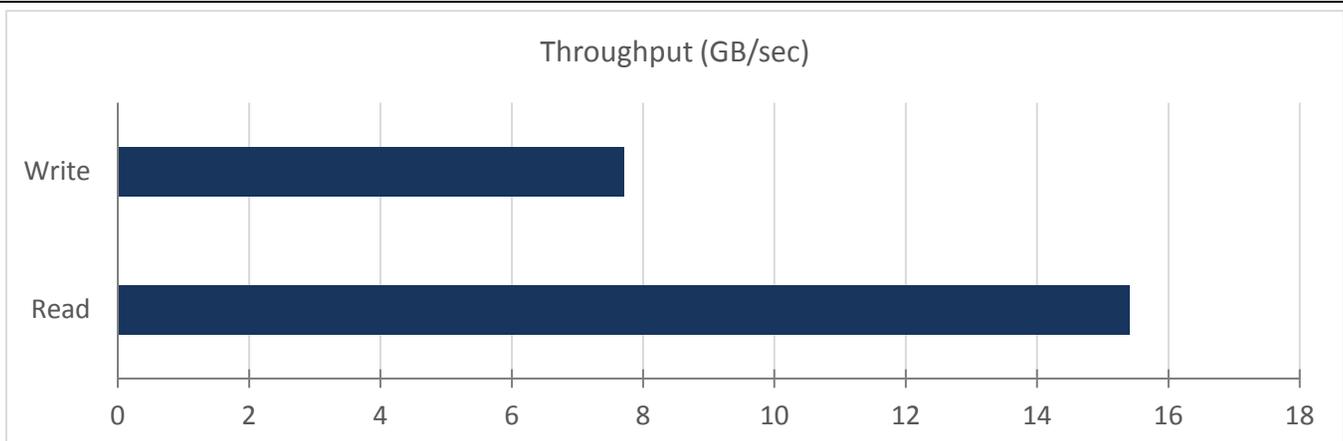
Source: Enterprise Strategy Group, 2017

Eight clients were attached to the Isilon F800, each with a 40GbE link via an Arista 7508 director-class switch. All clients communicated with the Isilon F800 over NFSv3. Each Isilon F800 node was populated with 15 3.2TB SSDs, with 60 total SSDs in the cluster.

Throughput was tested by generating sequential I/O with all eight clients using the industry-standard, open source [FIO](#) benchmark and stress testing tool. 128KB requests were used for the read workload, while 512KB requests were specified in the write test.

As seen in Figure 4, the single Isilon F800 chassis sustained throughput of 7.7 GB/sec for writes and 15.4 GB/sec for reads. A storage solution with high levels of aggregate throughput is ideally suited for bandwidth-intensive workloads including high-performance computing, big data analytics, video editing, and content delivery. Considering that a single 4K video stream can consume up to 15 Mb/sec of bandwidth, the single Isilon F800 chassis we tested would be able to support nearly 9,000 simultaneous 4K streams.

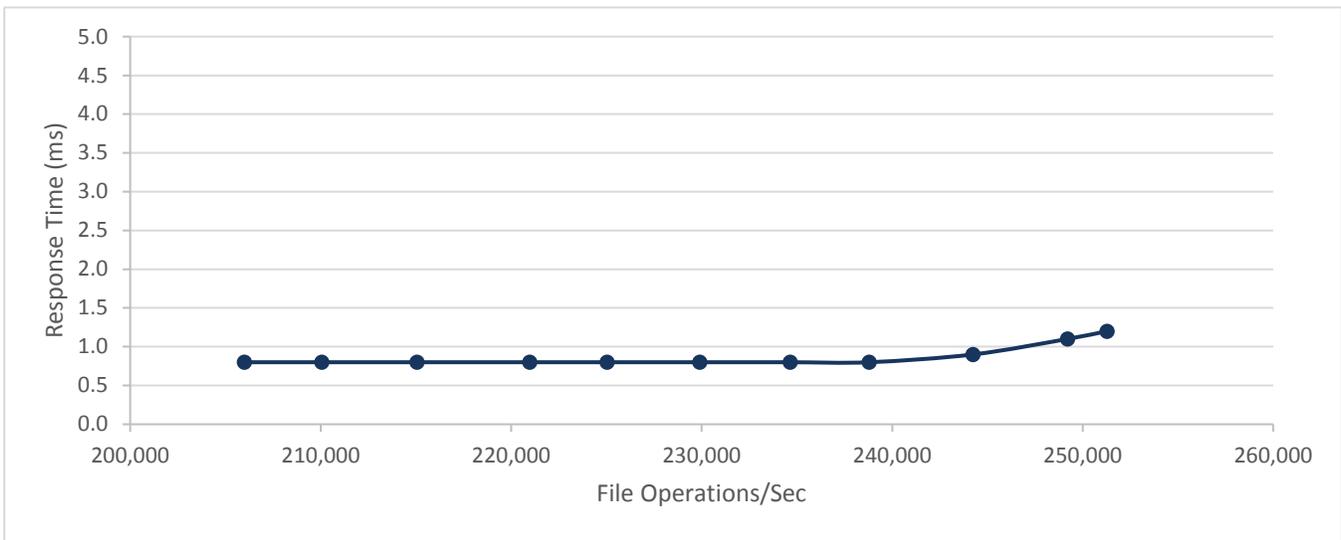
Figure 4. Results of Performance Testing—Four-node Cluster Aggregate Throughput



Source: Enterprise Strategy Group, 2017

For testing of file operations and response time, ESG Lab employed a methodology designed to simulate large numbers of users and applications performing multiple file system specific activities, including data-specific and metadata-specific activities common in file serving environments. Each of the eight clients generated an equal portion of the overall requested load. ESG Lab measured the number of operations per second (ops/sec) the system could sustain while maintaining acceptable response times. We measured response time as traffic was increased, and Figure 5 shows the results of our testing.

Figure 5. Results of Performance Testing—Operations per Second and Response Time

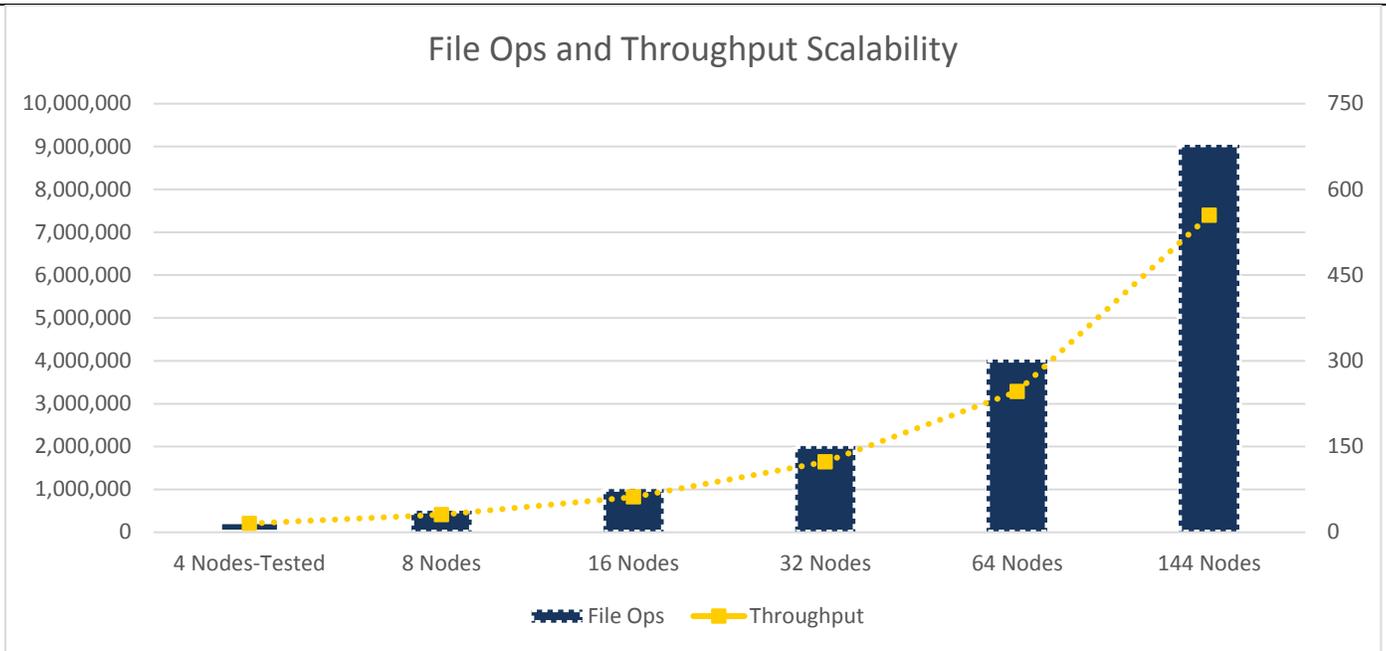


Source: Enterprise Strategy Group, 2017

As the test was running, ESG Lab observed consistently low response times, starting at 0.8 milliseconds and increasing to just 1.2 milliseconds as workload increased from 205,984 to 251,263 file operations per second. For comparison, the Isilon F800 All-Flash posted more than 62,800 file ops per node, which is roughly six times the performance of the disk-based S200 we tested in 2011.

ESG Lab next projected the file operations and throughput that a full 144-node cluster consisting of 36 Isilon F800 chassis could support, based on previous testing of large Isilon clusters. In every iteration of testing we have executed on the platform, we have observed near-perfect linear scaling of performance as nodes are added to the cluster, up to 144 nodes in our last encounter. In our example, we began with results from the Isilon F800 four-node cluster we tested and projected out to 144 Isilon F800 All-Flash nodes in 36 chassis. Assuming each chassis was configured identically to the Isilon F800 in our tests, a 144-node cluster would support more than 9 million file ops and 550 GB/sec.

Figure 6. Projected Scalability—File Operations and Throughput



Source: Enterprise Strategy Group, 2017



Why This Matters

The data organizations generate and manage continues to grow at a rapid pace with 35% of respondents to an ESG survey reporting greater than 40% year-over-year growth.³ In the same survey, respondents identified business intelligence and data analytics, video surveillance, and other generators/consumers of unstructured data as applications driving their storage growth.

To deliver products and services relying on such applications to their customers, organizations need the ability to access this data quickly and efficiently. Performance issues can drive negative results such as missed SLAs, increased costs, or decreased revenue. Organizations need all-flash storage that can start small without sacrificing performance, and scale out granularly and easily as needed.

ESG Lab confirmed that the Dell EMC Isilon F800 All-Flash can support extremely high levels of read and write throughput in an entry-level four-node cluster. The Isilon F800 under test sustained throughput of 7.7 GB/sec for writes and 15.4 GB/sec for reads, and more than 250,000 file operations at 1.2ms response time. Using projections based on previous testing of large Isilon clusters, the near linear scalability of the platform should enable a single Isilon F800 All-Flash file system to scale to 9 million file operations and 550 GB/sec of throughput.

³ Source: ESG Research Report, [2015 Data Storage Market Trends](#), October 2015.

The Bigger Truth

The downward trend of the cost of flash storage has made the benefits of flash more available than ever before. At the same time, unstructured data sets, which have historically leveraged performance-based infrastructure to a lesser extent, are increasingly demanding lower latency technology to drive faster business outcomes. Applications associated with unstructured data—big data analytics, high-performance computing, video processing, and life sciences, to name a few—are especially impacted as the amount of data they need to process is growing rapidly.

Anticipating the need for high-performance and cost-effective scalability for these application types, Dell EMC designed the Isilon F800 All-Flash, a scale-out NAS storage solution that offers high performance, capacity, and scalability. With capacities ranging from 92 terabytes to 33 *petabytes*, Isilon F800 All-Flash allows customers to take advantage of declining prices in solid-state technology without sacrificing performance. One 4U chassis can be deployed as a self-contained four-node cluster, which ESG Lab testing verified supports up to 250,000 file operations per second with all-flash response time of just over one millisecond. ESG Lab projected the file operations and throughput that a full 144-node cluster consisting of 36 Isilon F800 chassis could support, based on previous testing of large Isilon clusters. Based on ESG's calculations, a 144-node Isilon F800 cluster should be able to sustain 9 million file operations and throughput of up to 550 GB/sec. Dell EMC is enabling its customers to take advantage of all-flash technology for unstructured data-based applications that need extremely low latency. The Isilon F800 All-Flash runs the same OneFS as the rest of the Isilon line, providing the same enterprise functionality and integrating seamlessly with disk-based Isilon platforms with the ability to tier data across flash, disk, and the cloud.

ESG Lab has confirmed that Dell EMC Isilon F800 All-Flash is remarkably well-suited to support I/O-intensive unstructured data workloads. That this is achieved using a modular architecture that can be used to cost-effectively meet the performance needs of a wide variety of workloads using flash, disk, and the cloud, all from a single file system architecture, is especially impressive.

Appendix

TABLE 1. The ESG Lab Test Bed

Hardware	
Dell EMC Isilon F800 All-Flash	8 40GbE frontend ports 8 40GbE backend ports OneFS 8.1.0
Cisco C220 M4S Blade Server	128 GB RAM CentOS 7.3 40GbE NICs
Arista 7508R Modular Switch	75Tbps switching capacity 288 48GbE ports
Software	
FIO Utility	Version 2.1.10

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



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

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

