

Dell EMC SC5020 Storage Array

Performance, Flexibility, and Integration for the Evolving On-premises Data Center

Date: November 2017 **Author:** Brian Garrett, Vice President ESG Lab

Abstract

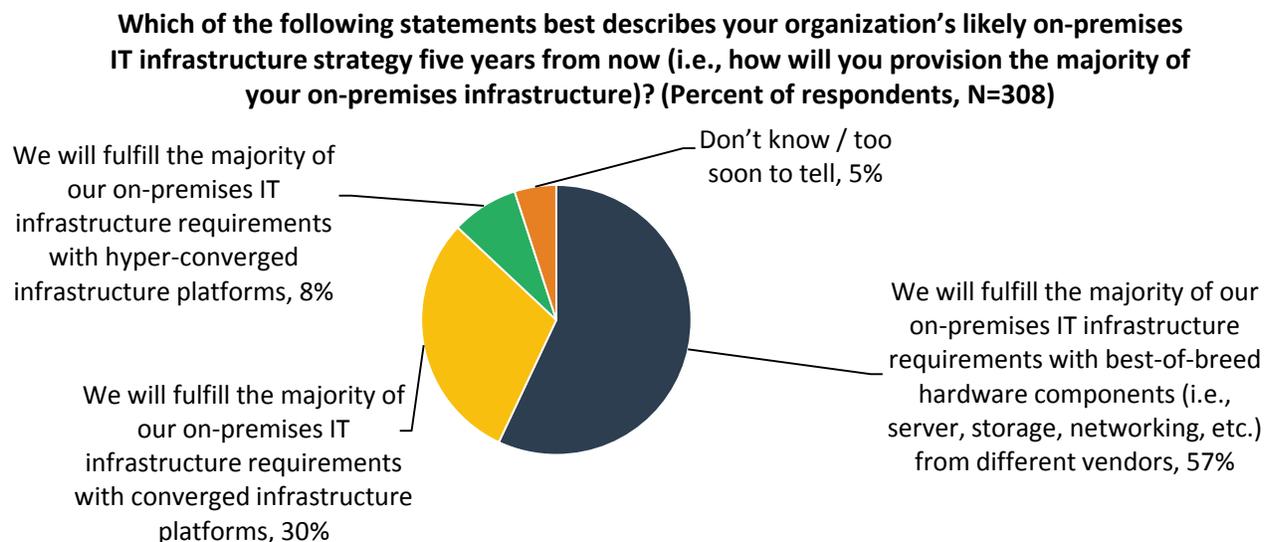
This ESG Lab Review documents validation testing of the Dell EMC SC5020 midrange block storage array with a focus on the all-flash performance of the recently updated hardware platform, the economic value of consolidating applications in a hybrid disk and flash configuration, and integration with EMC software. While this review focused specifically on the SC5020, the features tested in this review are representative of the entire SC Series.

The Challenges

The IT ecosystem is changing rapidly and the rate of change is accelerating. IT practitioners have more options to consider, from dramatic hardware improvements like flash storage, through new software and tool paradigms like DevOps, to entirely different deployment architectures like public, private, and hybrid clouds. However, ESG research shows that in the face of all this change, 57% of respondent organizations still plan to fulfill their on-premises IT infrastructure needs with best-of-breed hardware (see Figure 1).¹ Put simply, for the foreseeable future, most organizations will buy equipment to deploy into their own data centers and integrate with the other assets.

Additional ESG research shows that 32% of respondents believe that cost reduction initiatives will drive the most technology spending in their organizations, making it the second most-cited driver after cybersecurity initiatives.² This indicates that as organizations purchase equipment, they must not only keep up with the rate of change in the IT ecosystem, but also support the organization's overall cost reduction initiatives.

Figure 1. Expected On-premises IT Infrastructure Strategy



Source: Enterprise Strategy Group, 2017

¹ Source: ESG Research Report, [The Cloud Computing Spectrum: From Private to Hybrid](#), March 2016.

² Source: ESG Research Report, [2017 IT Spending Intentions Survey](#), March 2017.

The Solution: Dell EMC SC5020

The Dell EMC SC5020 is a midrange direct- or SAN-attached block storage array with the performance, front- and back-end flexibility, and software functionality that delivers enterprise-class storage services to resource-strapped organizations with midrange budgets. The Dell SC5020 is the first SC Series product that was released after Dell Computer and EMC Corp. merged in 2016. New enhancements include faster processors, more memory, and a simpler software model that eliminates cost and complexity with features enabled for the full array capacity.

The SC5020 hardware is a complete refresh from the SC4020, and now includes an eight-core Intel E5-2630v3 and quadruple the system memory. The back-end connectivity is twice as fast (12Gb SAS). New 12Gb SAS expansion chassis are available: 2U x 12 3.5" drives (targeted for high-performance HDDs), 2U x 24 2.5" drives (targeted for SSDs), and 4U x 60 3.5" drives (targeted for high-capacity HDDs).

Figure 2. Dell EMC SC5020



Source: Enterprise Strategy Group, 2017

The primary benefits of the SC5020's new hardware are improved performance and increased capacity. Dell EMC cites 3X throughput compared with the SC4020 and a 45% increase in IOPS. Total system drive count for the SC5020 is 222, up from 192, and the array's raw capacity has doubled to 2 PB. Additional spec improvements include twice the number of replays (a.k.a., snapshots) and nearly twice the number of replicated source and target volumes.

The base unit is a 3U chassis with 30 2.5" drive slots. The drives may be high-performance SAS HDDs, high-capacity nearline SAS HDDs, SSDs, and/or self-encrypting drives (SEDs) for data-at-rest security. The SC5020 features an onboard 10Gb Ethernet mezzanine card, which can be configured with SFP+ or RJ45 connections. Each controller has a single PCIe expansion slot that can support host connections for 12Gb SAS, 16Gb or 32Gb FibreChannel, or additional iSCSI ports. The SC5020 can support a wide range of protocol connectivity including single (FC, SAS, and iSCSI) and dual (SAS+iSCSI or FC+iSCSI) with additional choices for interface speed, e.g., 16Gb/32Gb FC, and connection type, e.g., SFP+/BaseT.

The SC5020 continues the SC Series' full-featured software offering including remote replication, Intelligent Deduplication and Compression on both SSDs and HDDs (even in hybrid configurations), and Data Progression auto-tiering. Live Migrate enables multi-array federations with seamless data mobility and load balancing between arrays, and Live Volume provides auto-failover capability for an affordably priced business continuity solution.

The SC5020 also introduces a much simpler software licensing model with no capacity tiers. All licenses are unlimited, up to the maximum number of drives supported on the array. Optional software is now packaged as feature bundles to simplify ordering and reduce costs. The SC5020 and the SC Series includes integration with management, mobility, and data protection tools familiar to EMC users including PowerPath, ViPR, VPLEX, RecoverPoint for Virtual Machines, Connectrix, and Data Domain.

Post-validation Feature Updates

Since ESG completed the validation, Dell EMC has added a new SC Series all-flash line. This line brings all the same SC Series functionality, with the added benefit of all-inclusive software licensing. In addition, Dell EMC has created a Future-Proof Storage Loyalty Program; a three-year satisfaction guarantee; 4:1 storage efficiency guarantee; and hardware investment protection. Customers are covered by this program when they purchase a new SC Series all-flash array with a maintenance contract.

ESG Lab Tested

The ESG Lab Validation began with an audit of the performance capabilities of the recently announced SC5020 hardware platform.

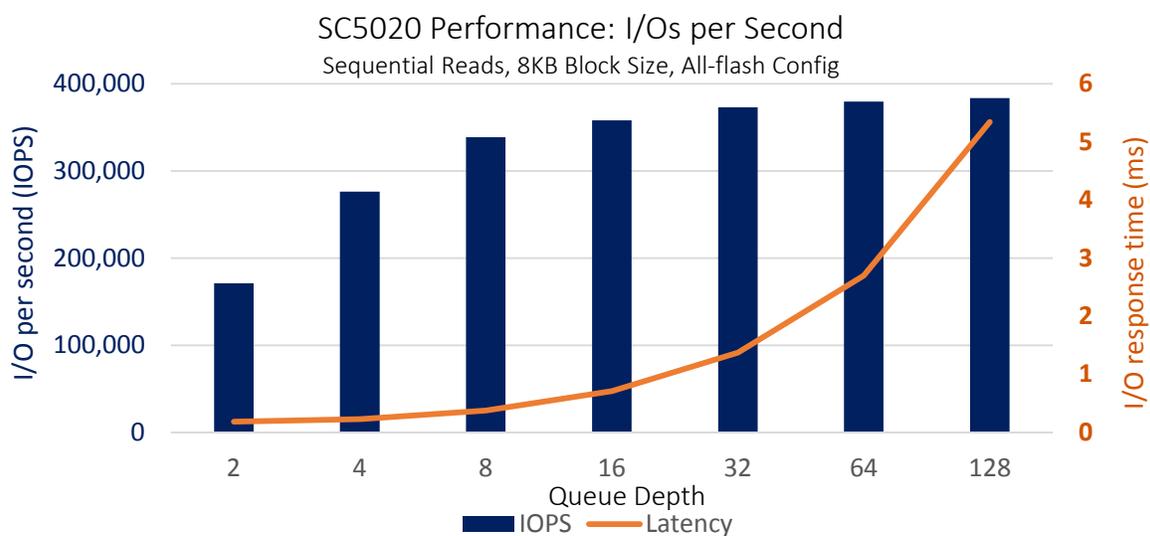
Performance

The primary driver of the improved performance of the SC5020, over the previous generation SC4020, is its enhanced hardware. In each storage controller, the CPU is now an Intel E5 2630 v3 processor with eight cores versus a four-core E3-1265L v2 processor in the SC4020. Controller memory was increased 4X to 64 GB per controller and updated to DDR4 technology. Back-end bandwidth to the drives was doubled with the use of 12Gb SAS for drive shelf connectivity. Front-end connectivity and bandwidth were also improved with up to a total of eight ports of 32Gb/s FibreChannel for the SC5020 versus four ports of 16Gb/s FibreChannel for the SC4020.

ESG Lab tested the performance of an all-flash configuration of the SC5020 to characterize upper performance limits and to determine performance in a typical application scenario. The test harness³ raised the I/O load to increase the number of outstanding I/Os at the array, known as queue depth. Array performance increased as more I/O was presented to the array.

Small block performance is indicative of a read in an operational database, like doing a query to set up a transaction. For small block reads, the SC5020 peaked at 383,511 IOPS and achieved about 365,000 IOPS at 1ms latency (see Figure 3).

Figure 3. SC5020 Performance: IOPS



Source: Enterprise Strategy Group, 2017

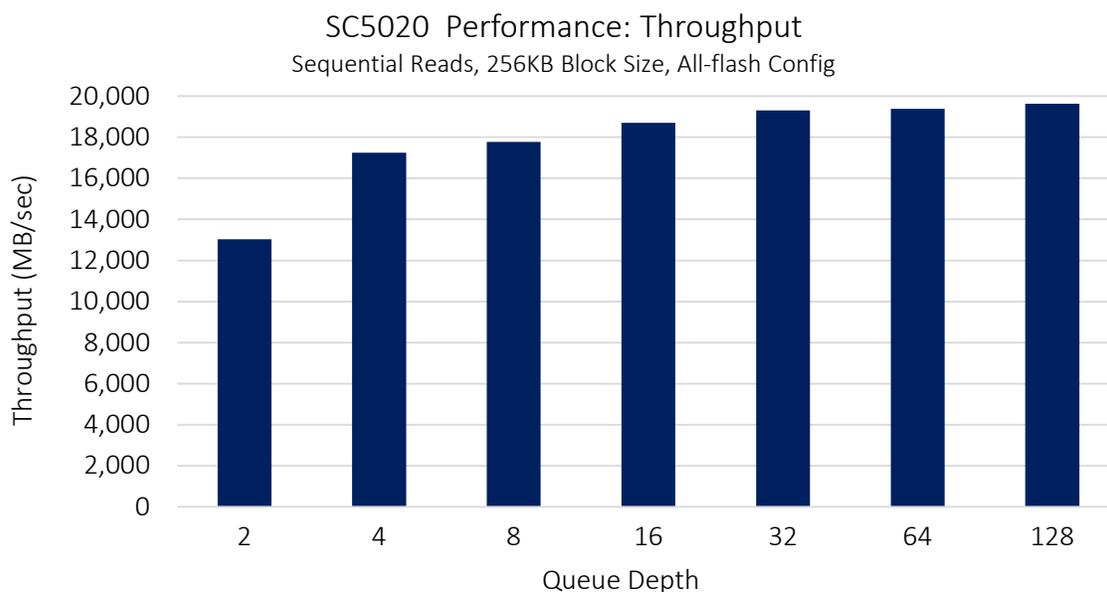
³ The test harness consisted of a single Dell R730 server, each with two ports of 16Gbs FibreChannel connections. The servers ran the Windows Server 2016 operating system and the HammerDB workload generator to generate loads.

Table 1. SC5020 Performance: IOPS

Queue depth	2	4	8	16	32	64	128
IOPS	171,157	266,676	338,596	358,050	373,020	379,546	383,511
Latency (ms)	0.19	0.23	0.38	0.71	1.37	2.70	5.34

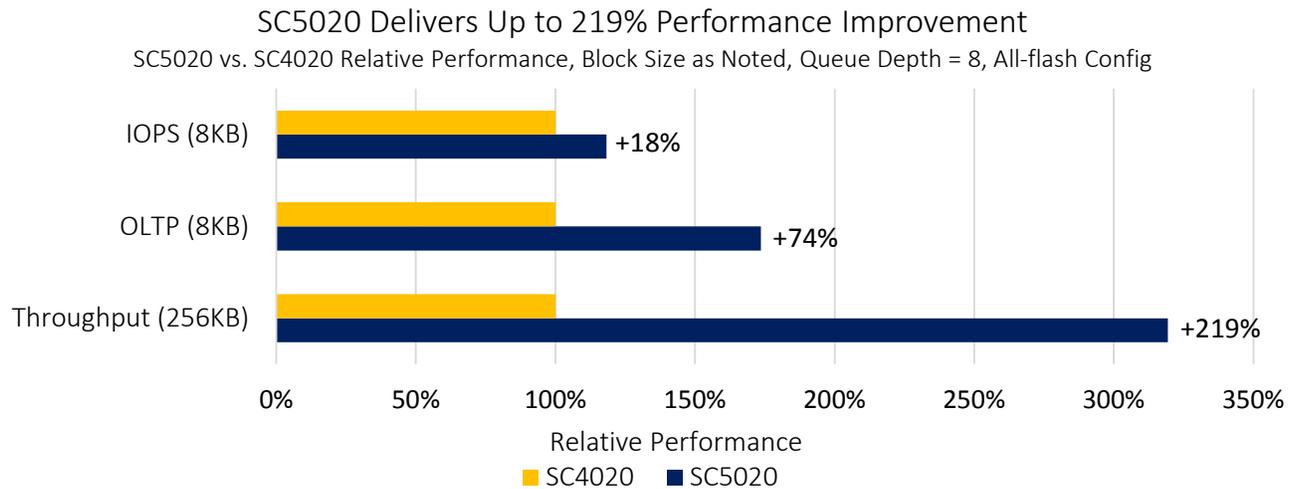
Large block sequential reads are typical of operations like bulk reading data structures to prepare a backup, or doing a full table scan of an online database to populate a data warehouse. In ESG Lab’s testing, the SC5020 delivered up to 19,265 MB/s throughput (see Figure 4).

Figure 4. SC5020 Performance: Throughput



Source: Enterprise Strategy Group, 2017

Finally, ESG compared the performance of the SC5020 with the previous generation SC4020. We compared the two arrays with an IOPS-intensive workload, a 70/30 read/write OLTP workload, and a throughput-intensive workload. To get a representative point of comparison across the workloads, we chose a middle-of-the-range queue depth of eight. As shown in Figure 5, the SC5020 outperformed its predecessor up to 219%.

Figure 5. SC5020 Delivers Up to 219% Performance Improvement

Source: Enterprise Strategy Group, 2017

What the Numbers Mean

In the all-flash configuration tested, SC5020 performance compares favorably to a state-of-the-market standalone all-flash array. A maximum of 385,311 IOPS, ~365,000 IOPS at one millisecond latency, and latencies below 250 microseconds for IOPS below 100,000 are comparable with market-leading all-flash arrays. Throughput of up to 19,265 MB/s exceeds most of today's all-flash arrays due to the SC5020 architecture with FC SAN support.

A comparison with the previous generation SC4020 shows the benefit of improved hardware, especially the increased bandwidth of the front-end FC and back-end SAS connections. The performance improvements vary depending on the application workload and I/O patterns. For example, ESG measured a 74% improvement for an OLTP workload due primarily to CPU and memory improvements, and a 219% improvement due primarily to bandwidth improvements.



Why This Matters

Storage performance correlates directly to organizational performance—increases in revenue, productivity, and often customer and user satisfaction ride on the back of increased performance of the IT infrastructure, including storage.

The SC5020 delivers enterprise-class performance at midrange storage prices. During ESG Lab testing, SC5020 performance topped out at 385,311 IOPS and 19,265MB/s throughput. For workloads below 100,000 IOPS, the SC5020 maintained fast all-flash I/O response times of less than 250 microseconds.

Data Progression

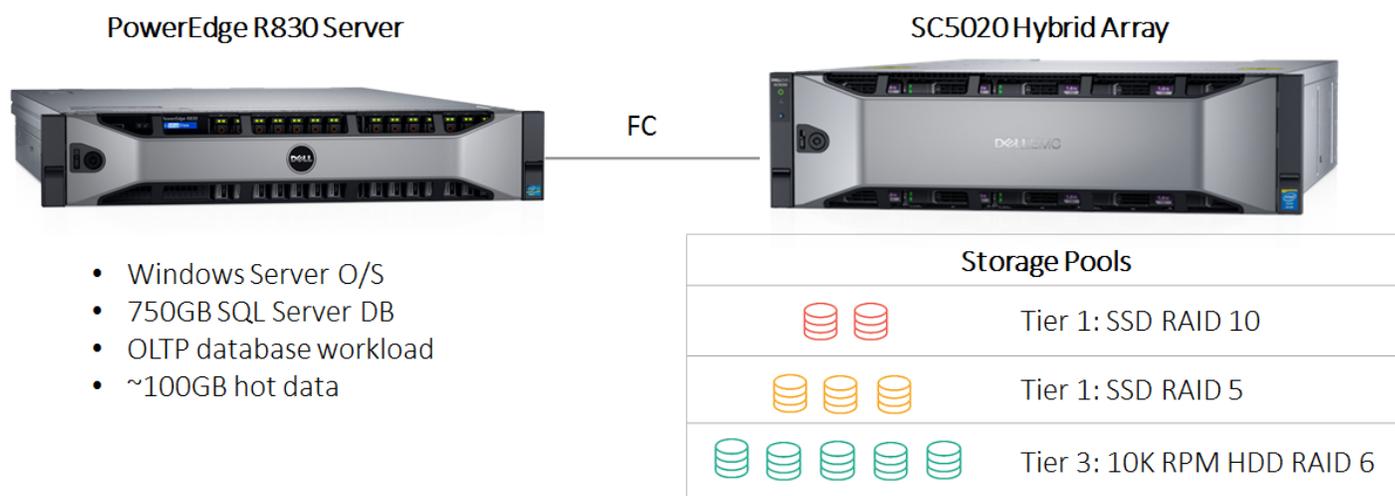
Data Progression, the SC Series' auto-tiering method, is key to many inherent advantages of the platform. Data Progression began as a way to store frequently accessed data on high-performance disk drives, or the outer tracks of high-performance disk drives, and less important data to higher capacity disk drives that are more affordable. The technology has now evolved to become a way to optimize performance, capacity, and cost across multiple types and configurations of HDDs and SSDs in a single storage array.

The SC Series' "0-100% flash" architecture lets customers target specific price/performance ratios for their initial configuration, deploying as much or as little flash as needed as performance and capacity needs change.

Within a single SC Series array, different media types can be deployed in different data protection configurations, for example, write-intensive SSDs in RAID 10 for the maximum performance of hot data, read-intensive SSDs in RAID 5 to lower cost while maintaining read performance, and HDDs of varying speeds—7.2K RPM, 10K RPM, and 15K RPM—in RAID 6 for lowest cost to hold cold data. The SC Series can support up to three different storage tiers. By default, all writes go to the highest performing drives and then Data Progression moves the data to more space-efficient tiers for subsequent reads. Data Progression enables an SC array to maintain performance while decreasing an application’s \$/GB over time.

ESG Lab tested the Dell EMC SC5020’s Data Progression capability with a SQL Server database as show in Figure 6. The array was configured with a hybrid mix of six SSDs and 12 10K RPM HDDs. Tier 1 was five SSDs containing RAID 10 and RAID 5, and Tier 3 was 11 HDDs in RAID 6. A Dell EMC PowerEdge R830 server hosted a SQL Server database and a HammerDB OLTP database workload generator that simulates order and payment processing, and stock checking, at a wholesale supplier. SQL Server ran on Windows Server on bare metal and was configured for “maximum server memory” of 100 GB. The database contained 750 GB of data. Approximately 20% of database capacity was actively being referenced and updated; the remainder was inactive, but accessible.

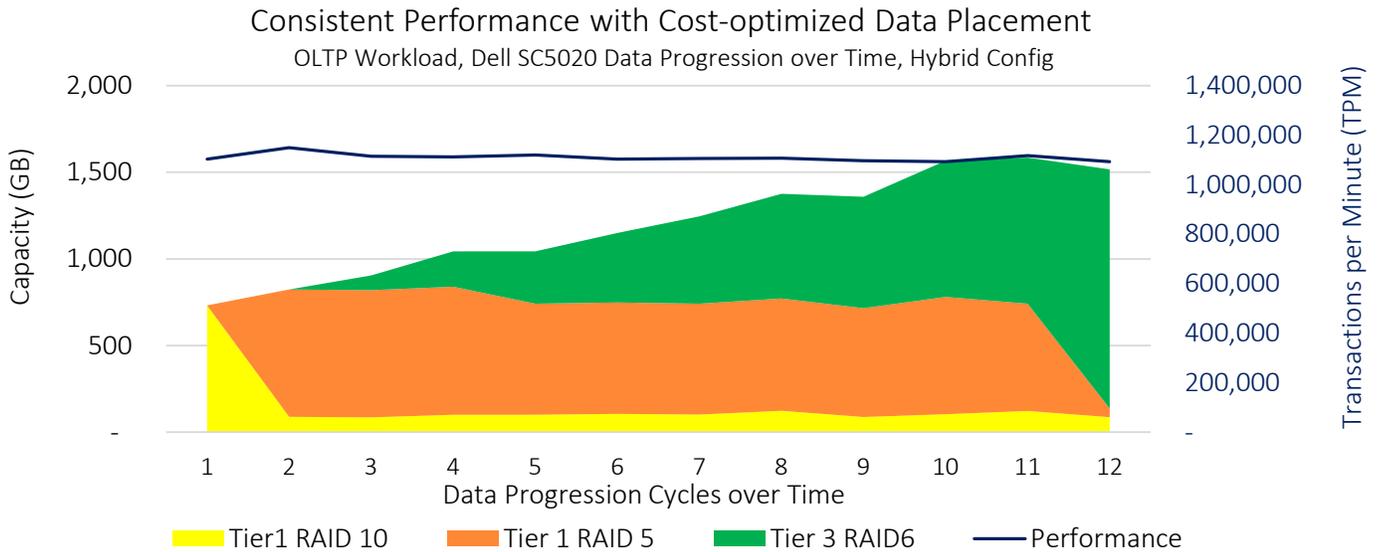
Figure 6. SC5020 Data Progression Test Bed



Source: Enterprise Strategy Group, 2017

Testing included regular snapshots using the SC5020’s thin snapshot capability, which records changes only. Each test run was a performance load test and snapshot, repeated until all of the cold data was moved to the HDD Tier 3. The movement of the data between tiers is shown in Figure 7. At test run 2, the snapshot data moves to Tier 1 RAID 5 and at test run 3, the snapshot data begins to move to Tier 3. By test run 12, all of the snapshot and cold data moved to Tier 3. Note **how** application performance remained essentially unchanged at 1.1M transactions per minute due to the fact that approximately 100 GB of hot data lived on the highest performing Tier 1 storage for the duration of the test.

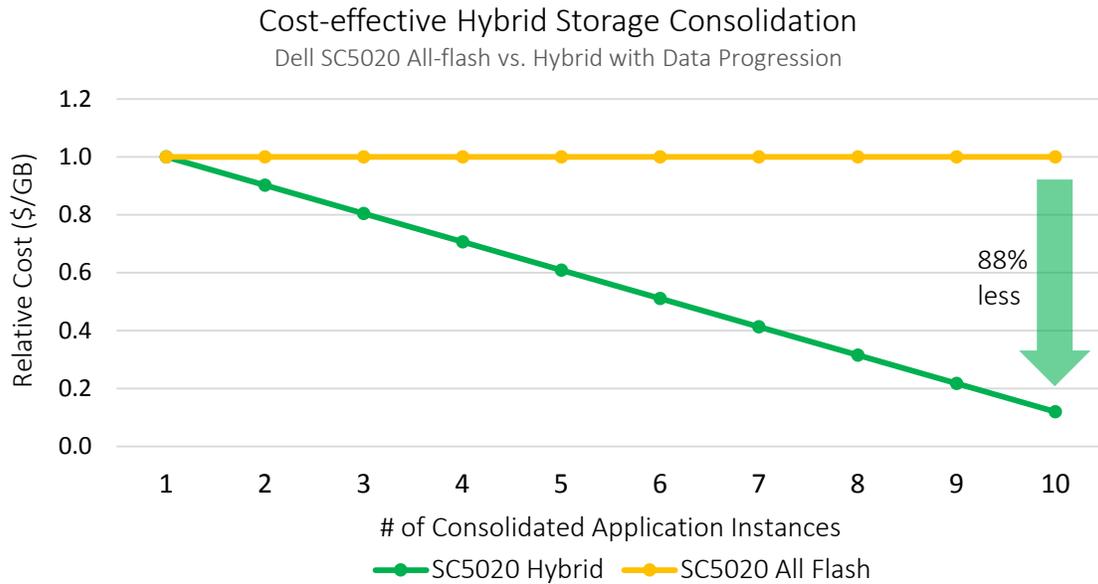
Figure 7. Consistent Performance with Cost-optimized Data Placement



Source: Enterprise Strategy Group, 2017

A straightforward calculation based on the results of this test can be used to illustrate the economic power of SC5020 Data Progression. The total required storage capacity for the 750GB database settled at almost 1,600 GB as the snapshot activity approached steady state. However, by moving snapshot and infrequently accessed cold data to a more cost-effective RAID 5 SSD configuration, then to an even more cost-effective 10K RPM HDD RAID 6 configuration, the effective \$/GB for storage media for the application was reduced by 88%. Based on the results of this test and the all-flash performance results presented earlier in this report, ESG is confident that the combination of Data Progression and the turbocharged performance of SC5020 hardware platform can be used to cost-effectively consolidate more applications on the same hybrid storage array. Specifically, for typical multiuser applications with hot spots that fit comfortably in Tier 1 all-flash, ESG Lab has confirmed that a cost-effective SC5020 hybrid configuration can handle up to ten application instances at an effective \$/GB that's 88% less than an all-flash array.

Figure 8. Cost Advantage of Consolidating More Application Instances with SC5020 Data Progression



Source: Enterprise Strategy Group, 2017

i Why This Matters

IT is always pressured between the demands of higher performance and the constraint to lower cost. ESG research shows that while 57% of respondents name performance as a key factor for deploying solid-state storage, those same respondents most frequently ranked hardware costs as the primary storage challenge for their organization.⁴

In ESG Lab’s testing, the SC5020’s Data Progression moved infrequently accessed data from the highest-performing Tier 1 SSD RAID 10 storage to Tier 3 HDD RAID 6 storage. The effective storage medium \$/GB cost for the test application was reduced by 88% and freed up SSD capacity so the array under test could handle up to ten instances of the test application at this lower effective storage cost.

Tools Integration

Dell EMC has an extensive storage hardware and software product portfolio. Since the merger of Dell and EMC, the combination of the companies’ capabilities enables Dell EMC to offer customers new combinations of best-of-breed data management and storage solutions. Concurrent with the merger of Dell and EMC, the newly combined company provided support for legacy EMC tools with SC Series arrays, including PowerPath, ViPR, VPLEX, RecoverPoint for Virtual Machines, Connectrix, and Data Domain. ESG Lab tested PowerPath, ViPR, and RecoverPoint for Virtual Machines operation in conjunction with the SC Series.

⁴ Source: ESG Research Report, [2015 Data Storage Market Trends](#), October 2015.

PowerPath

PowerPath is a server-resident software solution designed to enhance performance and application availability. PowerPath enhances application availability by providing load balancing, automatic path failover, and recovery functionality.

PowerPath Migration Enabler

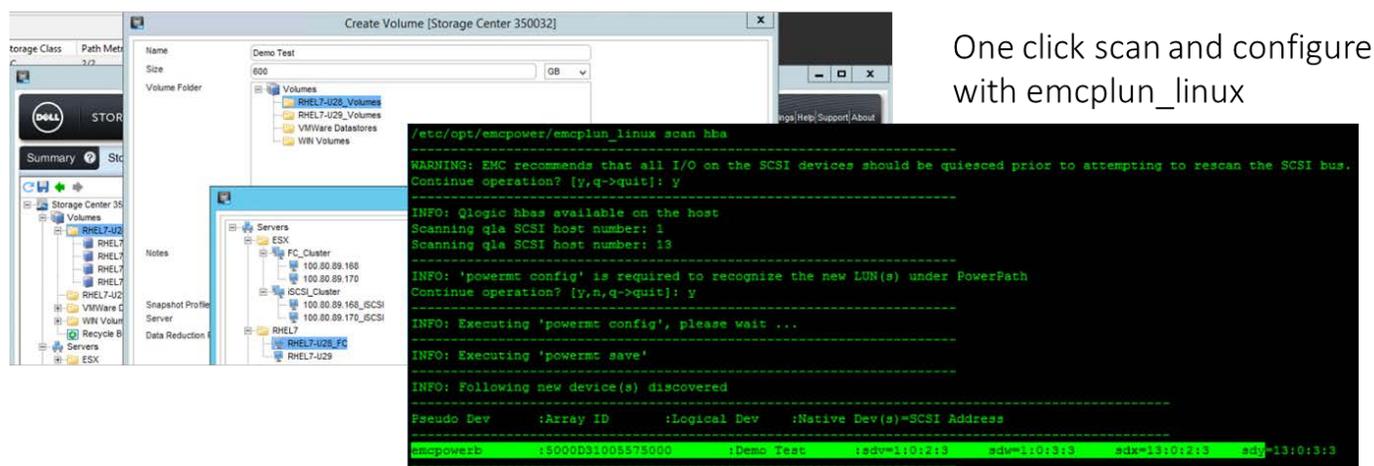
PowerPath Migration Enabler is a host-based migration tool that allows you to migrate data between supported storage systems. Migration Enabler uses the PowerPath filter driver, in conjunction with an underlying technology, to provide nondisruptive, or minimally disruptive, data migrations.

PowerPath optimizes the data path between hosts and storage arrays. PowerPath ensures fault tolerance of the storage network by handling failover among multiple data paths between storage and host. PowerPath optimizes performance in the storage network by using proprietary multipathing algorithms on the traffic between hosts and storage volumes. For current or future SC5020 users, PowerPath offers a common path management tool across Dell EMC midrange storage products. The PowerPath host utilities also simplify path management tasks that are not addressed natively by the host operating system.

ESG Lab used PowerPath's host utilities to configure the path between an SC volume and an RHEL 7 Linux host (see Figure 9). First, we used Dell Storage Manager to create a volume named *Demo Test*. Next, we used PowerPath's native Linux utility, `emcplun_linux`, to rescan the storage connections, identify the new LUN, and connect it to the host. We used PowerPath's `powermt` command on the host to verify new SC Series storage capacity available to the host over four different FibreChannel paths.

Figure 9. Connecting an SC Volume to Linux with PowerPath

Create Volume on SC5020 with Dell Storage Manager



The image shows a screenshot of the Dell Storage Manager interface for creating a volume named 'Demo Test' on a Storage Center 350032. The interface includes fields for Name, Size (600 GB), and Volume Folder. Below the interface, a terminal window displays the output of the `emcplun_linux scan hba` command. The terminal output shows a warning to quiesce I/O, followed by scanning for SCSI hosts and LUNs, and successful execution of `powermt config` and `powermt save`. The final output shows the discovery of new devices with their respective pseudo, array, logical, and native device identifiers.

One click scan and configure with `emcplun_linux`

```

/etc/opt/emcpower/emcplun_linux scan hba
-----
WARNING: EMC recommends that all I/O on the SCSI devices should be quiesced prior to attempting to rescan the SCSI bus.
Continue operation? [y,n,q->quit]: y
-----
INFO: Qlogic hbas available on the host
Scanning qla SCSI host number: 1
Scanning qla SCSI host number: 13
-----
INFO: 'powermt config' is required to recognize the new LUN(s) under PowerPath
Continue operation? [y,n,q->quit]: y
-----
INFO: Executing 'powermt config', please wait ...
-----
INFO: Executing 'powermt save'
-----
INFO: Following new device(s) discovered
-----
Pseudo Dev      :Array ID      :Logical Dev   :Native Dev(s)=SCSI Address
-----
emcpowerb      :15000D31003571000 :!Demo Test   :!sdvr=1:0:2:13  sdw=1:0:3:13  sdw=13:0:2:3  sdv=13:0:3:3

```

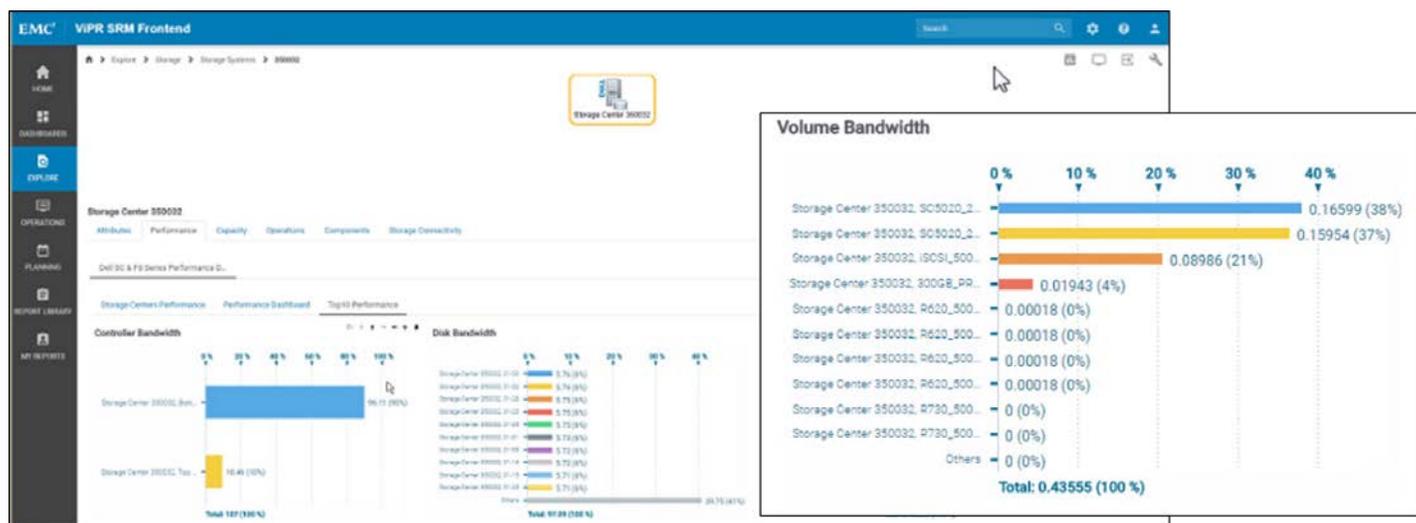
Source: Enterprise Strategy Group, 2017

ViPR

ViPR is a storage virtualization appliance (ViPR Controller) and storage resource management application (ViPR SRM) that provides virtualized multi-tenant storage administration and reporting, with service management capabilities including chargeback, cancellation, and departmental control. The ViPR Controller can aggregate capacity from a wide variety of SAN-connected back-end storage arrays into a single logical pool that can then be administered by ViPR SRM. ViPR SRM gathers raw data from the Dell Storage Manager Data Collector to create reports about SC Series arrays. For SC Series users, ViPR provides multitenant management and reporting and its own value-added reporting.

ESG Lab used ViPR to create capacity and manage performance on SC arrays. Using ViPR Controller's capacity creation process was as easy as shopping on the Internet—with menus of available services and a single screen from which to order storage capacity. In the ViPR SRM front-end, we selected **Explore** from the left rail navigation. For the SC5020 array, we selected the **Performance** tab and **Top 10 Performance** metrics. ViPR displayed the bandwidth usage for the top 10 controller, disk, and volume resources (see Figure 10).

Figure 10. ViPR SRM Top 10 Performance Display for a Dell SC5020 Storage Array



Source: Enterprise Strategy Group, 2017

RecoverPoint for Virtual Machines

RecoverPoint for Virtual Machines is an appliance- or software-based storage replication product. RecoverPoint for Virtual Machines can mirror between dissimilar storage array types. For SC5020 users, RecoverPoint for Virtual Machines allows disaster recovery configurations with other types of storage arrays. RecoverPoint can mirror from one SC Series system to a different storage array. The SC5020's flash performance and cost-effective hybrid capability make it a compelling disaster recovery target.

EMC Lab tested RecoverPoint for Virtual Machines mirroring data from an SC array to a Dell EMC Unity array. For this test, the SC array was primary, at the *Datacenter*, also known as *Site A*. The Unity array was the target, at the *DR_Datacenter*, also known as *Site B*. First, we logged into the VMware vSphere Web Client and used the left rail to navigate to the RecoverPoint for VMs plugin where the **Overall Results** tab showed the protected VMs. We then identified the SC array between the primary and disaster recovery site by reviewing the RecoverPoint for Virtual Machines Wizard to highlight the data flow between both locations. The RecoverPoint Wizard at both *Site A* and *Site B* established the connection between the sites.

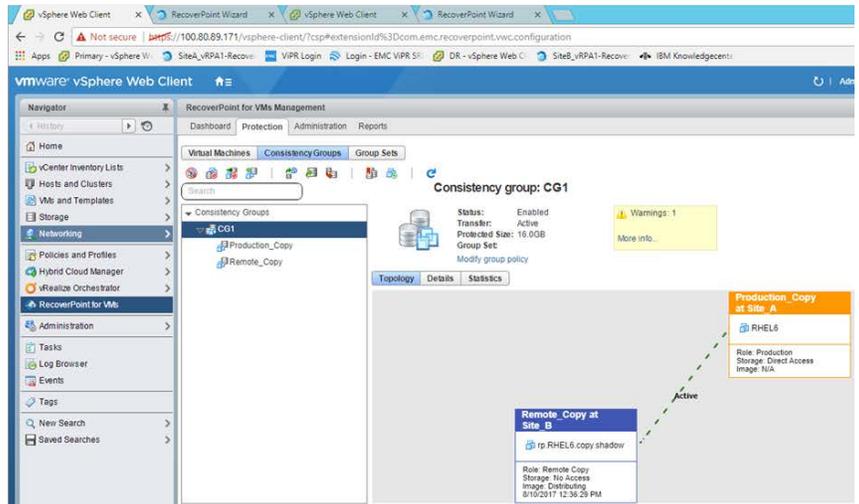
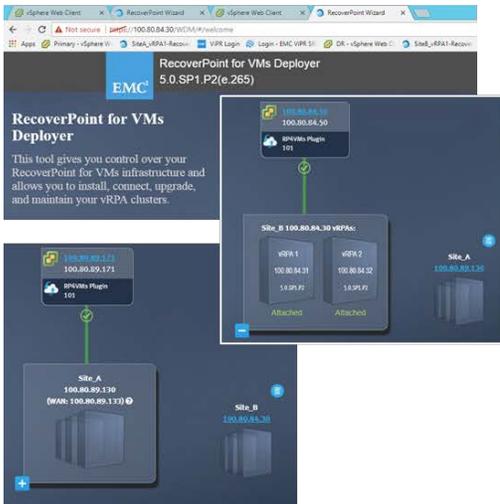
VMs running on SC have been protected by RecoverPoint for Virtual Machines to *Site B* by right clicking the protection option directly through the vSphere Web Client.

Protection status can be monitored within the RecoverPoint plugin. Under protection, the **Consistency Groups** tab gives a visual confirmation of data flow from *Datacenter Site A* to *DR_Datacenter Site B* through an animated display (see Figure 11).

Figure 11. SC Array with RecoverPoint for VMs

Wizard, Sites A&B

Data Flow, Site A->B



Source: Enterprise Strategy Group, 2017

Why This Matters

Recent ESG research shows that the most important considerations in justifying IT investments include return on investment, cited by 29% of respondents, and reduced operational expenditures, cited by 25% of respondents.⁵ Purchases that can leverage value from other purchases, especially by reducing complexity, strengthen the balance sheet.

ESG Lab confirmed that SC5020 and the SC Series has been integrated with legacy EMC software tools including PowerPath, ViPR, and RecoverPoint. Installing an SC5020 where these products already exist increases the value of an SC5020 investment. Adding one or more of these tools to your next SC Series purchase magnifies the value of your investment.

⁵ Source: ESG Research Report, [2017 IT Spending Intentions Survey](#), March 2017.

The Bigger Truth

The users of midrange storage are uniquely beset with the challenge to “do more with less.” These users are usually small to medium businesses or smaller business units and departments within larger enterprises. Because they are smaller organizations, their budgets are smaller than their large enterprise counterparts. However, these users still have to handle the global tectonic shifts to support more users (e.g., mobile workers), more data (e.g., big data), and online performance expectations. And on-premises compute and storage is still a reality in their data centers.

The merger of Dell and EMC in late 2016 offers new combinations of capabilities for midrange storage buyers. Dell, with its SC Series arrays, brings a high-performance array with hybrid hardware and software capabilities to address midrange users’ cost concerns. EMC brings a set of tools—appliances and software products—that enhance the capability of any storage array and reduce operational complexity through fewer panes of glass. The Dell EMC SC5020 is the first SC Series array launched following the merger of the two companies. New enhancements in the SC5020 include improved controller hardware, a modular chassis from the SC7020, and with the all-flash models, a simpler, all-inclusive software model.

While this lab focused on the SC5020, the functionality applies to the SC Series. Dell EMC recently released the SCv3000, the entry model in the SC Series, as well as an all-flash line with all-inclusive software, and an extended storage loyalty program. While not tested by ESG Lab, we expect these new models to continue the SC Series value proposition.

EMC Lab tested the Dell EMC SC5020 for performance, its ability to manage data across multiple tiers of storage media, and integration with legacy EMC tools. The performance of the recently updated SC5020 hardware platform was comparable to industry-leading all-flash arrays (up to 385,311 IOPS and 19,000 MB/sec). Data Progression testing with a cost-effective mix of flash and disk demonstrated how the storage costs for a consolidated pool of applications can be reduced by as much as 88%. ESG believes that the combination of the new SC5020 hardware platform, hybrid disk and flash support with advanced data tiering, and integration with EMC software is ideally suited for organizations with enterprise-class storage aspirations and midrange budgets.

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 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.