Lab Validation Report

IBM Spectrum Protect

Simplified Data Protection for the Modern Enterprise

By Vinny Choinski, Senior ESG Lab; and Kerry Dolan, Lab Analyst

December 2015

Contents

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. This ESG Lab report was sponsored by IBM.

Introduction

This ESG Lab Validation documents remote hands-on testing of IBM Spectrum Protect. Testing focused on how IBM is helping its customers confidently protect their data with Spectrum Protect by simplifying deployment and management, building prescriptive guidance into the solution, delivering data protection-as-a-service, and improving overall performance to help manage today's data protection workloads.

Background

In spite of the fact that backup and recovery are longstanding IT disciplines, they remain challenging and problematic to the point that only 3% of organizations surveyed by ESG reported experiencing no challenges of any kind. It is also worth noting that respondents clearly acknowledge the diversity of ongoing challenges. When asked about current data protection difficulties, those prioritized by respondents as *primary* challenges are closely associated with data growth and complexity, including cost(s), keeping pace with the capacity of data to protect, and managing the complexity of data protection environments (see Figure 1).[1]

Figure 1. Top 10 Primary Data Protection Challenges

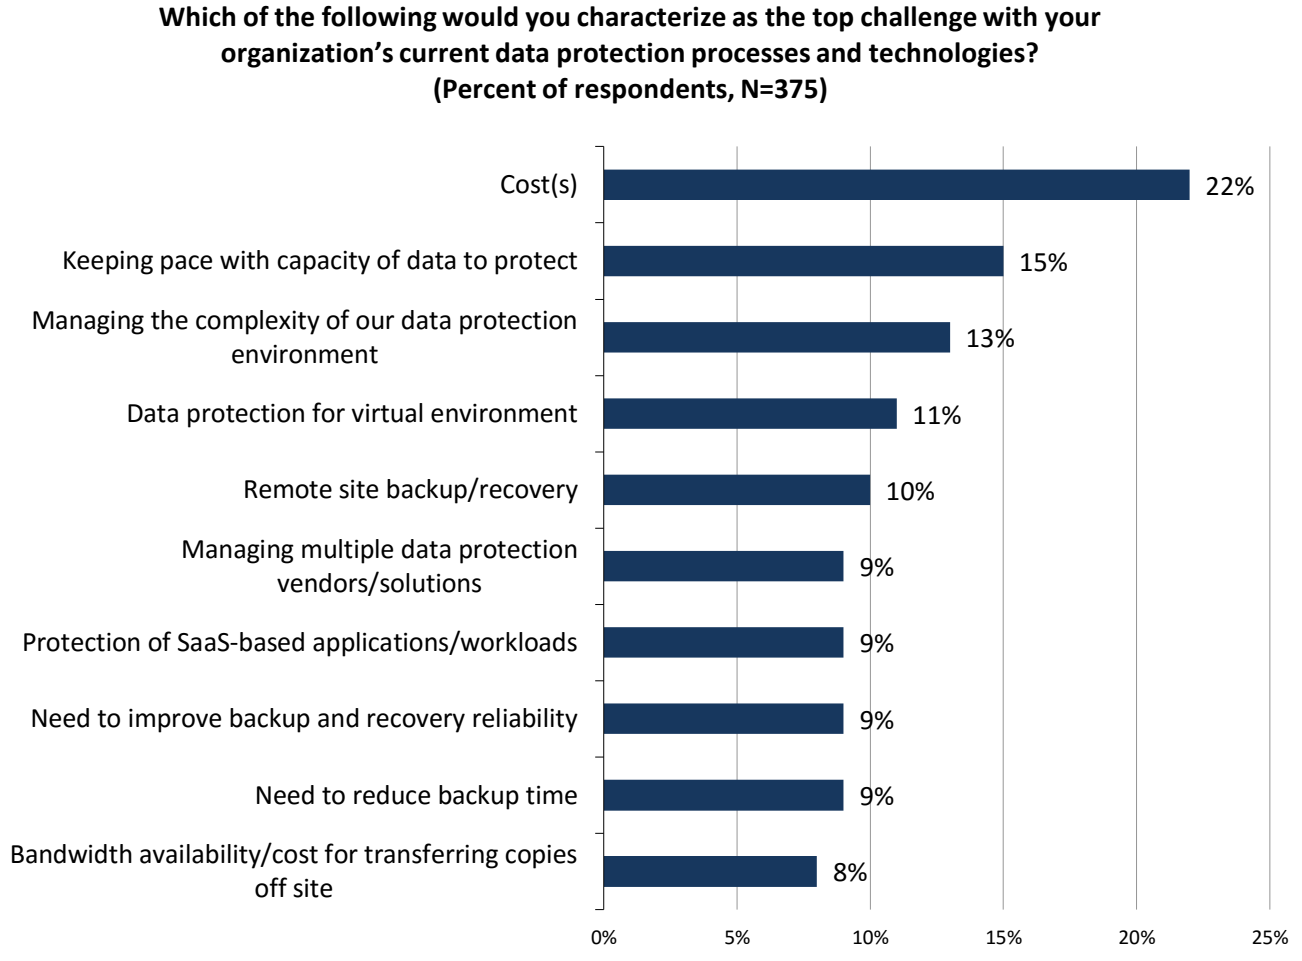

Source: Enterprise Strategy Group, 2015.

[1] Source: ESG Research Report, *2015 Trends in Data Protection Modernization,* September 2015

The Solution: IBM Spectrum Protect

IBM Spectrum Protect, formerly Tivoli Storage Manager, is a member of the IBM Spectrum Storage family. It enables advanced data backup and data recovery for virtual, physical, cloud, and software-defined environments—as well as core applications and remote facilities.

Figure 2. Spectrum Protect Solution Overview

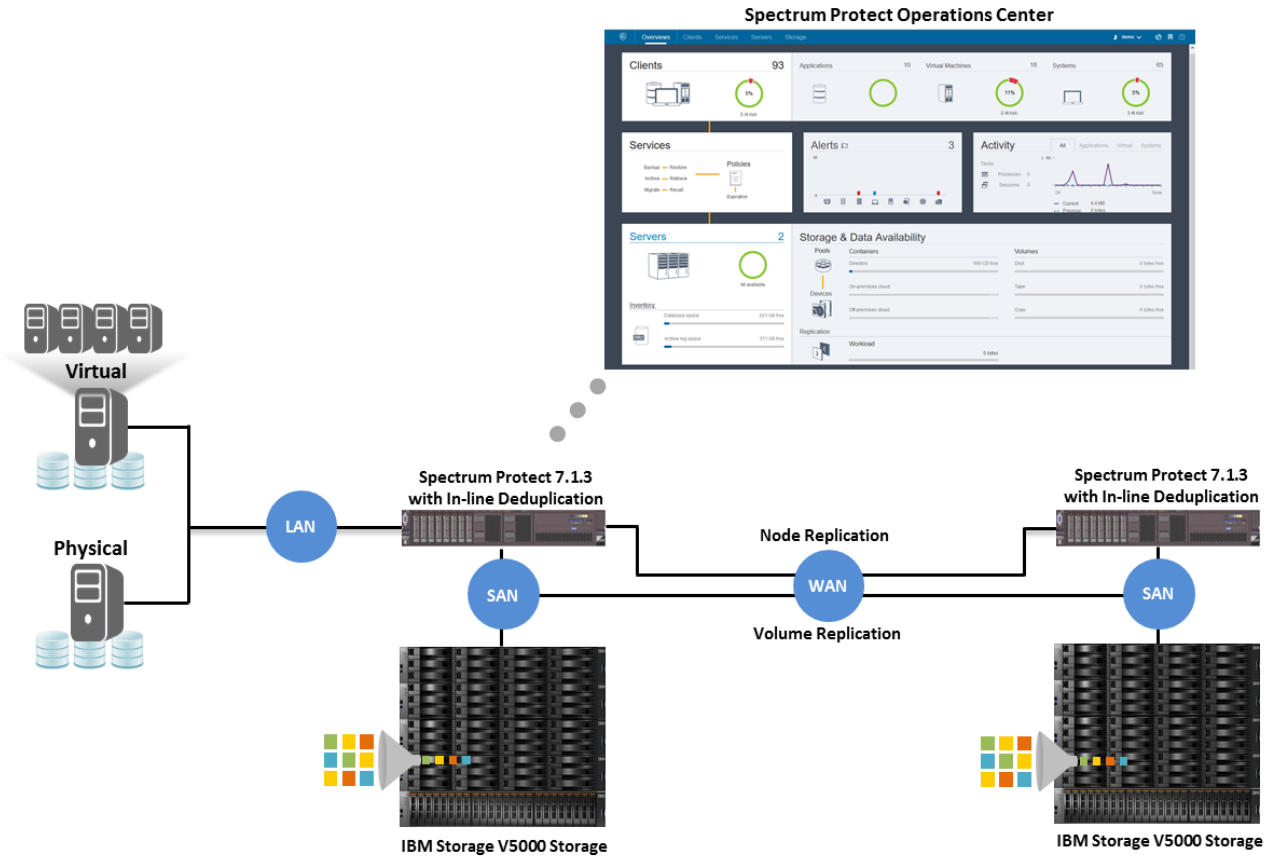

IBM Spectrum Protect solutions:

- **Enable software-defined storage environments,** by delivering file, block, and object data protection for IBM Spectrum Storage and other SDS environments.
- **Enable cloud data protection** with OpenStack and REST support, cloud portal for multi-tenancy, IBM SoftLayer integration, and hybrid cloud optimization.
- **Integrate with VMware and Hyper-V** for efficient virtual machine data protection that includes hardware-assisted snapshots, copy management, incremental "forever" backups, and new self-service restore portal, initially for VMware.
- **Mitigate the risk of data loss** with frequent snapshots, multi-site replication, and disaster recovery management.
- **Reduce the total cost of data protection** with built-in software-defined deduplication.
- **Can leverage IBM Spectrum Protect Blueprints** for controlled implementations designed to eliminate risk and speed deployment. The blueprints are offered in small, medium, and large configurations depending on data protection workloads.

ESG Lab Validation

ESG Lab performed remote hands-on evaluation and testing of IBM Spectrum Protect via an IBM facility in Tucson, Arizona. Using industry standard tools and methodologies, testing was designed to demonstrate deployment automation, simplified management, improved solution performance, and prescriptive guidance built into the user experience via the new blueprints, sample policies, and configuration wizards. Also of interest was how IBM is leveraging Frontsafe software with Spectrum Protect to deliver data protection-as-a-service.

Deployment and Usability

The deployment and usability section of this validation report documents the IBM Spectrum Protect blueprint deployment process, integrating a new Spectrum Protect server into the Spectrum Protect Operations Center, using Frontsafe software to deliver Spectrum Protect capabilities as services, and using the self-service web interface to conduct end-user restores.

ESG Lab Testing

ESG Lab started testing by conducting an automated small blueprint deployment of IBM Spectrum Protect. As shown in Figure 3, ESG leveraged an existing Lenovo x3650 server that was SAN-attached to a Storwize V3700 array. The server was running the Windows operating system and the storage was configured and provisioned to the server as described in the blueprint details. ESG Lab downloaded and extracted the Spectrum Protect blueprint configuration on the server.[2] We then ran the ***TSMserverconfig.pl*** script with the solution-specific parameters file as a command option. The parameter file defined configuration settings such as system name, users, passwords, and the location of the Spectrum Protect database and backup pools. After launching the automated configuration script, it took approximately eight minutes to configure the system and prepare it for use.

Figure 3. Small Blueprint Deployment

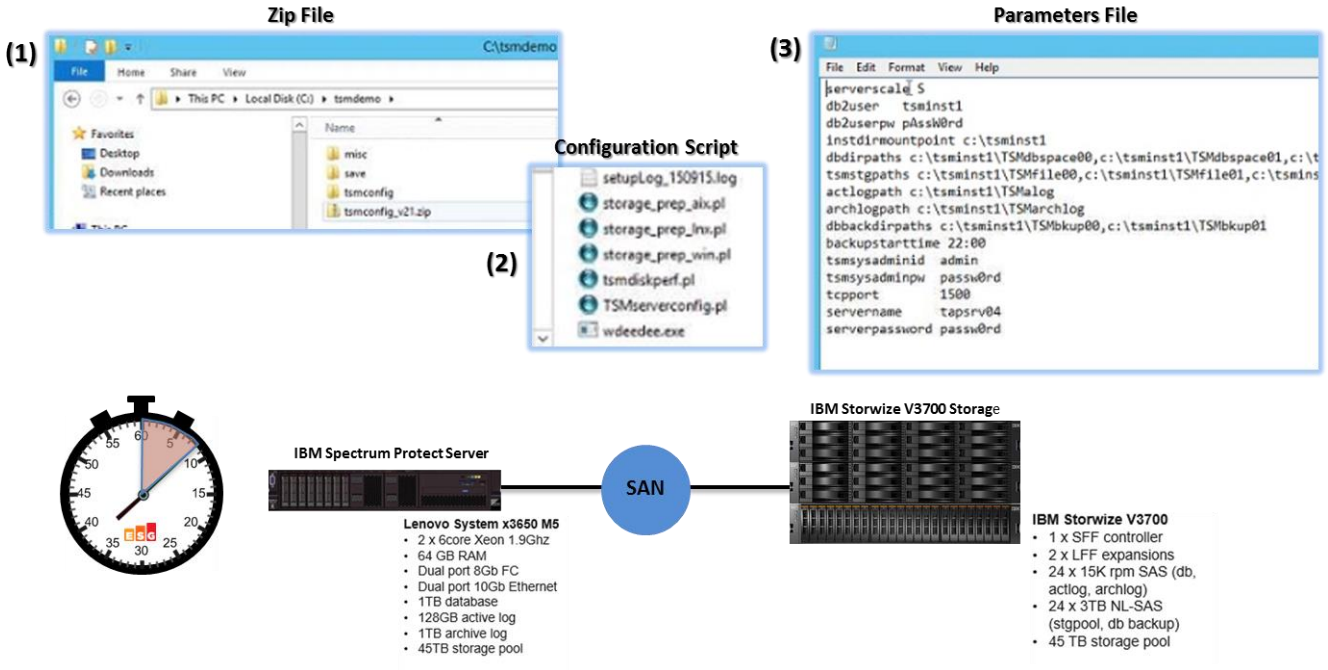

Next, ESG Lab logged into the Spectrum Protect Operations Center server in the test environment and navigated to the server tab. As shown in Figure 4, we used the ***Connect Spoke Server*** wizard to add the newly provisioned Spectrum Protect server to the hub server environment.

[2] Spectrum Protect Blueprints: http://ibm.biz/IBMSpectrumProtectBlueprints

As shown in the red callout box in Figure 4, after walking through the simple configuration steps in the wizard, the new spoke server began to synchronize. Once fully synced, the status icon will turn green, indicating that the Spectrum Protect server is ready to be managed and monitored.

Figure 4. Connect to Operations Center

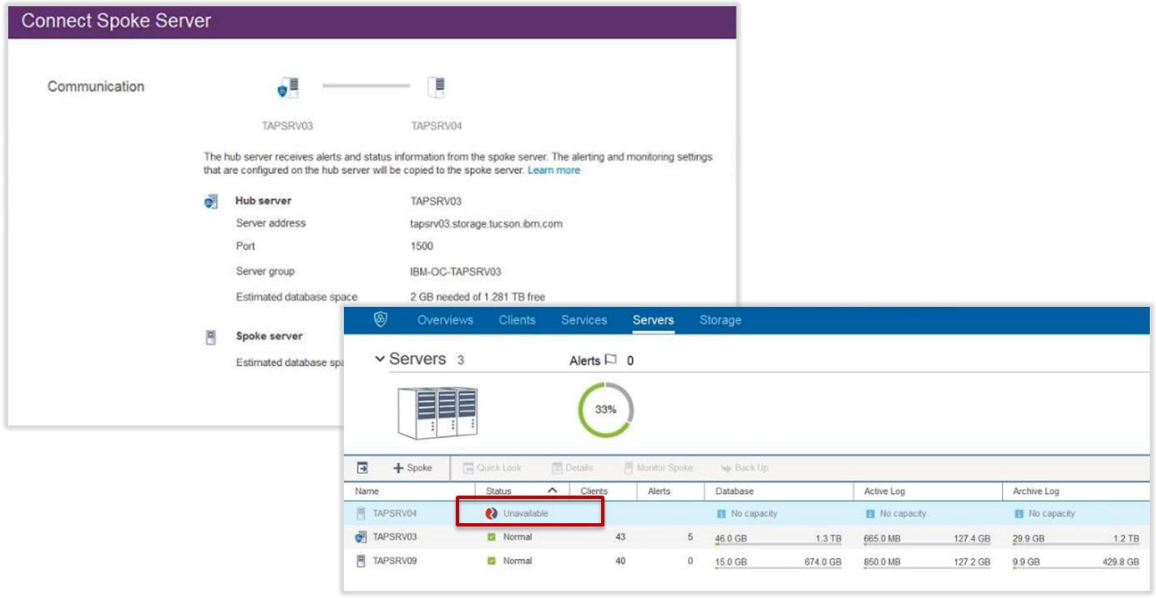

Next, ESG Lab connected to a Frontsafe demo environment. Frontsafe software provides a web-based interface that enables users to productize Spectrum Protect capabilities and deliver them as a service. Figure 5 shows the Global Enterprise settings view from the user interface in a distributed Spectrum Protect model environment.

Figure 5. Frontsafe Software Overview

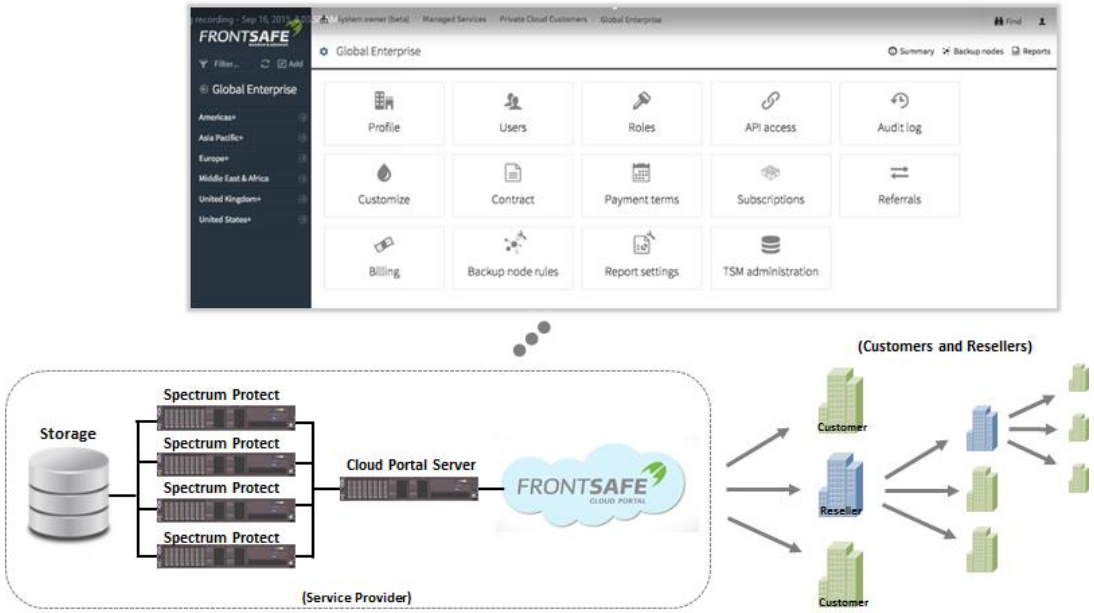

Frontsafe Cloud Portal provides very granular and comprehensive capabilities for creating, delivering, and managing service-based data protection products. It enables complete separation of backup data for groups with additional privacy or accounting requirements, with a multi-tenancy model originally developed for managed service providers. As shown in Figure 5, there is access to billing automation, backup provisioning, reporting, user management, self-service, and RESTful API integration on all levels of the multi-tenant model.

Finally, ESG Lab used the web-based self-service portal to restore a deleted file. As shown in Figure 6, ESG logged into the portal by selecting a virtual machine from a dropdown list, chose the backup time and date, browsed to the deleted file, and restored it to the original location to recover the file.

Figure 6. Self Service-Restore

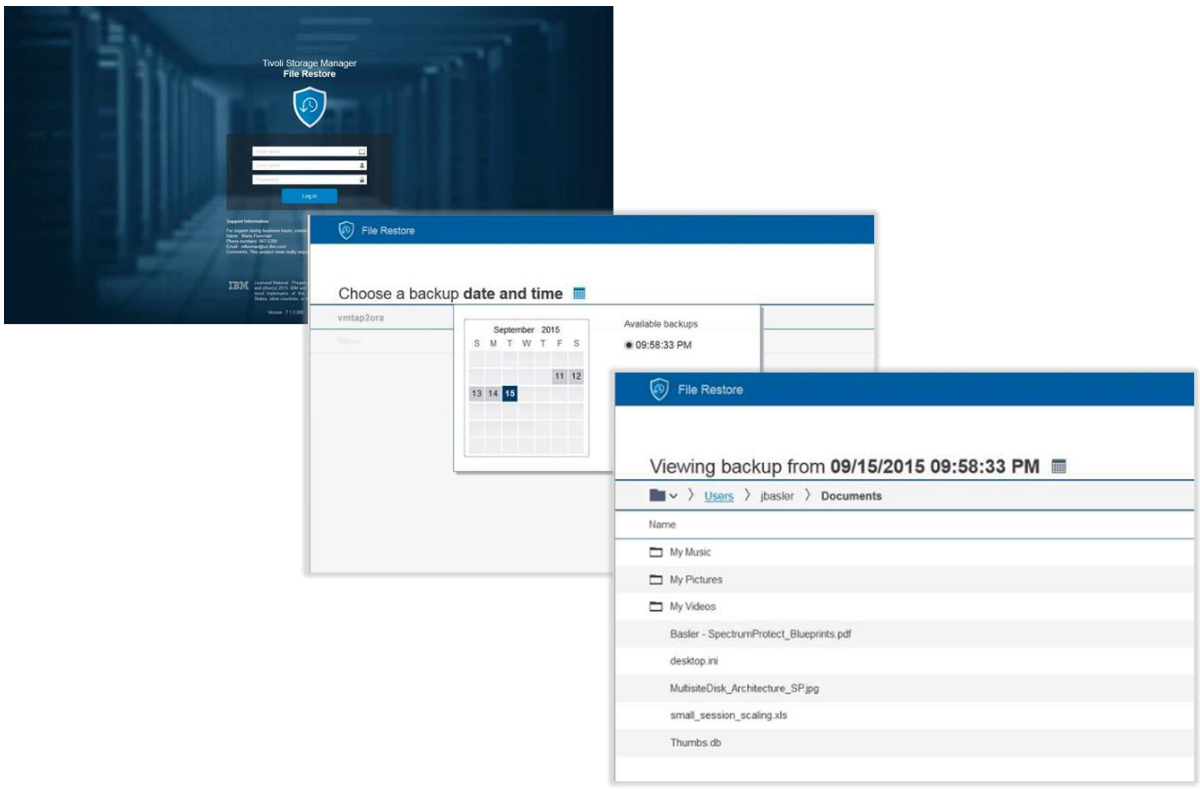

Why This Matters

According to ESG research, managing cost is the most-cited challenge that organizations reported facing with current data protection processes and technologies.[3]

ESG Lab confirmed that IBM helps customers address this challenge with Spectrum Protect blueprints that speed deployments and reduce risk by automatically implementing best-practice configurations and taking the guesswork out of performance optimization. In addition, the Frontsafe software and the self-service portal simplify solution management. With Spectrum Protect, each of these components ultimately helps reduce OpEx by minimizing expensive administration time.

[3] Source: ESG Research Report, *2015 Trends in Data Protection Modernization*, September 2015.

Data Protection Management

In this section of the validation, ESG Lab explores the Spectrum Protect administrator experience when managing different data protection concepts, including setting up replication, monitoring workload status, configuring storage pools, and decommissioning obsolete clients.

ESG Lab Testing

ESG Lab started its data protection management testing by exploring the replication configuration process. As shown in Figure 7, replication is now a two-phase process that consists of node replication and protect pool replication. The upper left side of the figure shows the initial setup steps. First, we selected the source and target Spectrum Protect servers. Then, as shown by the blue pie wedge in the ***Add Server Pair*** wizard initial setup screen, ESG created a schedule window for the replication job. We were also offered the opportunity to adjust the transfer rate of the replication job based on an estimate of the time needed to synchronize the data between the two servers. Next, as shown in the upper right side of Figure 7, ESG configured pool-level replication between the source and target Spectrum Protect servers, which initiates chunk-level replication between storage volumes. The orange pie wedges in the wizard view represent the scheduled ***Protect Pool*** replication. It should be noted that node-level replication maintains a complete inventory of the protected client's metadata and is fully integrated with ***Protect Pool*** replication, so only unique data is replicated with any schedule (e.g., node-level or protect pool-level schedules). The Add Server Pair wizard walked us through the three screens to set up replication according to best practices. Then, as shown in the bottom of Figure 7, we used the replication screen to review the configuration and monitor status.

Figure 7. Replication Configuration

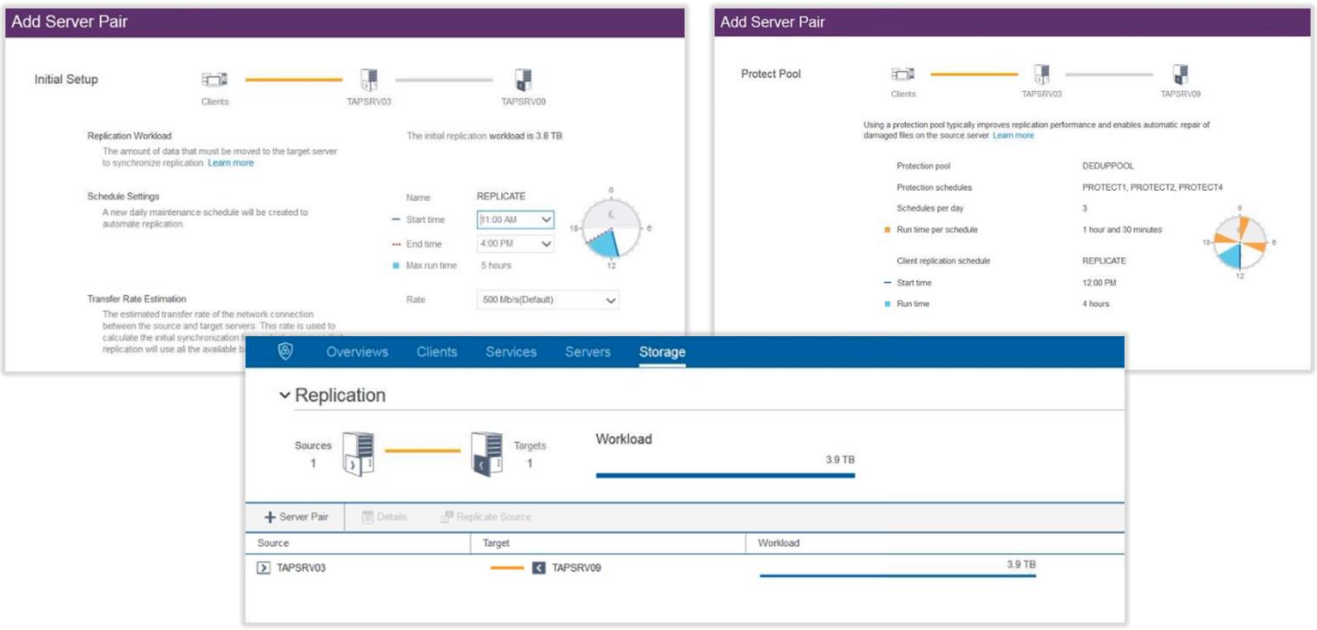

Next, as shown in Figure 8, ESG Lab used the workload view to review the status of the newly configured replication job. The right side of Figure 8 shows key aspects of replication. The light blue bar represents the amount of new data stored on the server and eligible for replication. The amount of data already replicated is shown in purple. The total replication workload is shown in dark blue. This is the total amount of data that still needs to be replicated. In a steady state, we would expect the amount backed up and the amount replicated to match and the total workload to be very small. Figure 8 depicts a backlog for the total workload because replication was just enabled for the demonstration environment and existing data on the source server needed to be replicated to the target. The red X on the lower right side of the figure indicates that workload processing is falling behind. This was the case because the recently configured replication hadn't yet had enough time to catch up. The ideal view would be to have each

of the workloads relatively balanced. This is expected to happen over time if the solution is properly designed to handle the expected workloads. It should be noted that, at the completion of the replication wizard configuration process, an estimate of the time required to complete replication based on the schedule, data set, and bandwidth is provided to the administrator.

Figure 8. Workload Overview

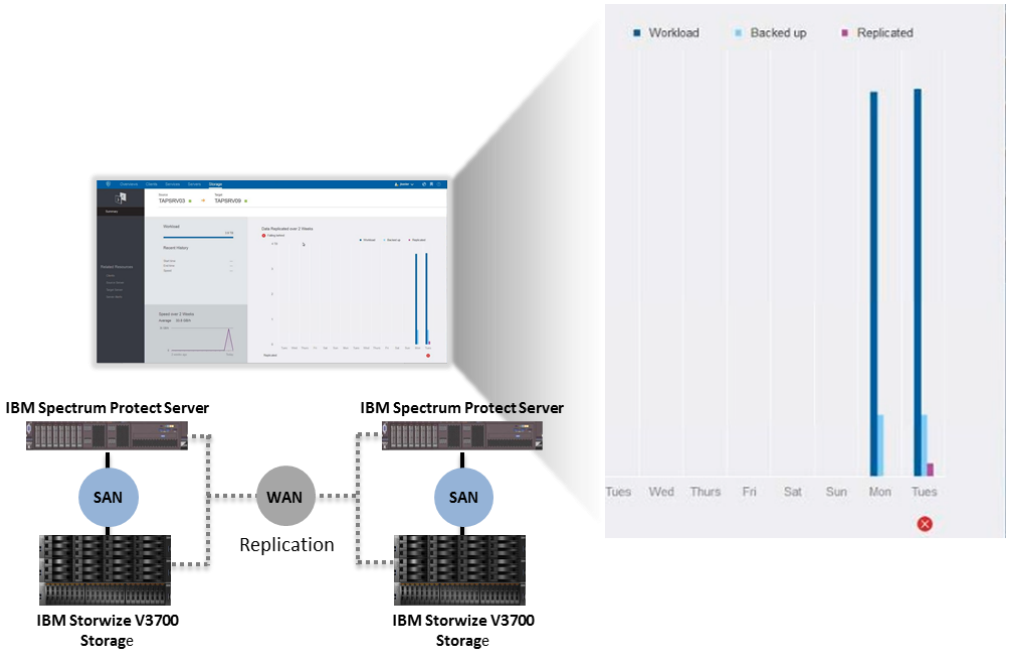

Next, as shown in Figure 9, ESG Lab used the ***Add Storage Pool*** wizard to configure an on-premises cloud pool. In this case, we used an existing Openstack Swift implementation in the test environment.

Figure 9. Cloud Pool

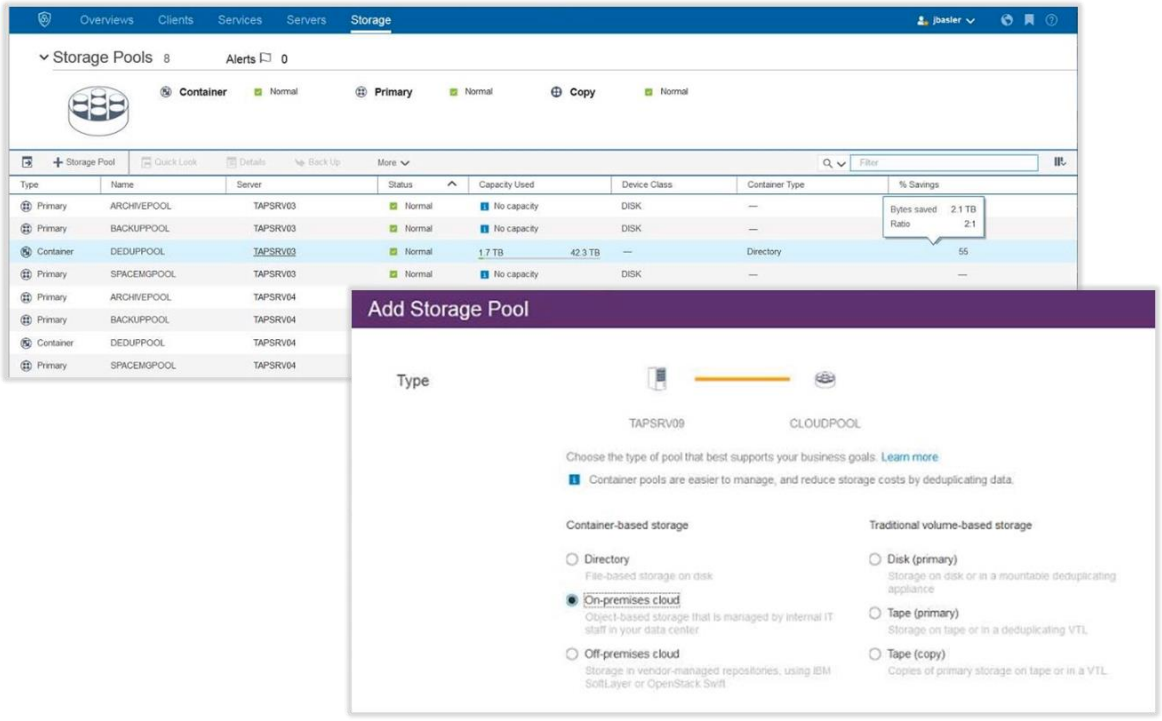

Creating the pool was simple, and required only that we add the URL to the storage with the login credentials. Once the pool was created, we updated an existing policy and client to start using the new configuration. Each step was wizard-based and was accompanied with prescriptive guidance such as encryption suggestions for on-premises and off-premises cloud pool creation or the direction to update policy destination settings to leverage the newly created storage pool.

Finally, as shown in Figure 10, we decommissioned obsolete virtual machines from the Spectrum Protect environment. The integration between vCenter and Spectrum Protect allowed us to easily identify deleted VMs and reclaim valuable space for the data protection environment.

Figure 10. Decommission

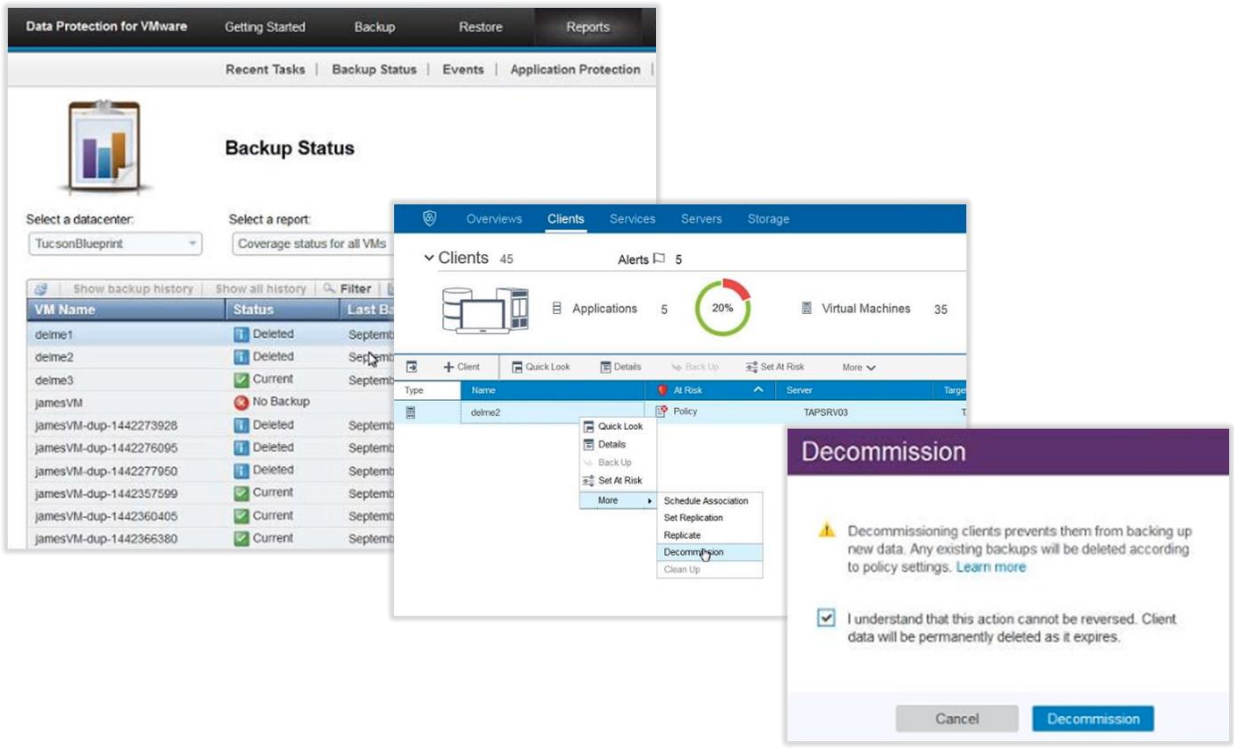

Why This Matters

Organizations are constantly looking for ways to simplify management and improve efficiency. Continually evolving and growing production resources demand that today's data protection solutions be flexible, easy to manage, and even able to address multiple business workloads in a single solution.

ESG Lab validated that Operations Center makes Spectrum Protect extremely flexible and easy to manage. The solution is designed to manage modern data protection workloads while providing visibility and insight. ESG was able to set up replication jobs with prescriptive guidance, update an existing policy with a new cloud storage pool, understand the impact to existing backup jobs, reconcile obsolete virtual machines between VMware and Spectrum Protect, and decommission those VMs. It is a solution that provides local and remote data protection and disaster recovery with the visibility needed for efficient data protection management.

Enhanced Performance

Performance is an important component of any data protection solution, and the latest version of Spectrum Protect offers its customers significant performance improvements. ESG Lab reviewed the technology changes that enable the new performance capabilities, audited the results of test data captured during the development process, and explored the ability to deploy the solution leveraging Spectrum Protect Blueprints.

ESG Lab Testing

ESG Lab started performance testing by reviewing the blueprint methodology, which is designed with prescriptive hardware and build detail to handle defined workloads with build automation. As shown in Figure 11, they are designed around small, medium, or large workloads with each blueprint server and storage architecture optimized for deduplication.

Figure 11. Reduced Backup Duration

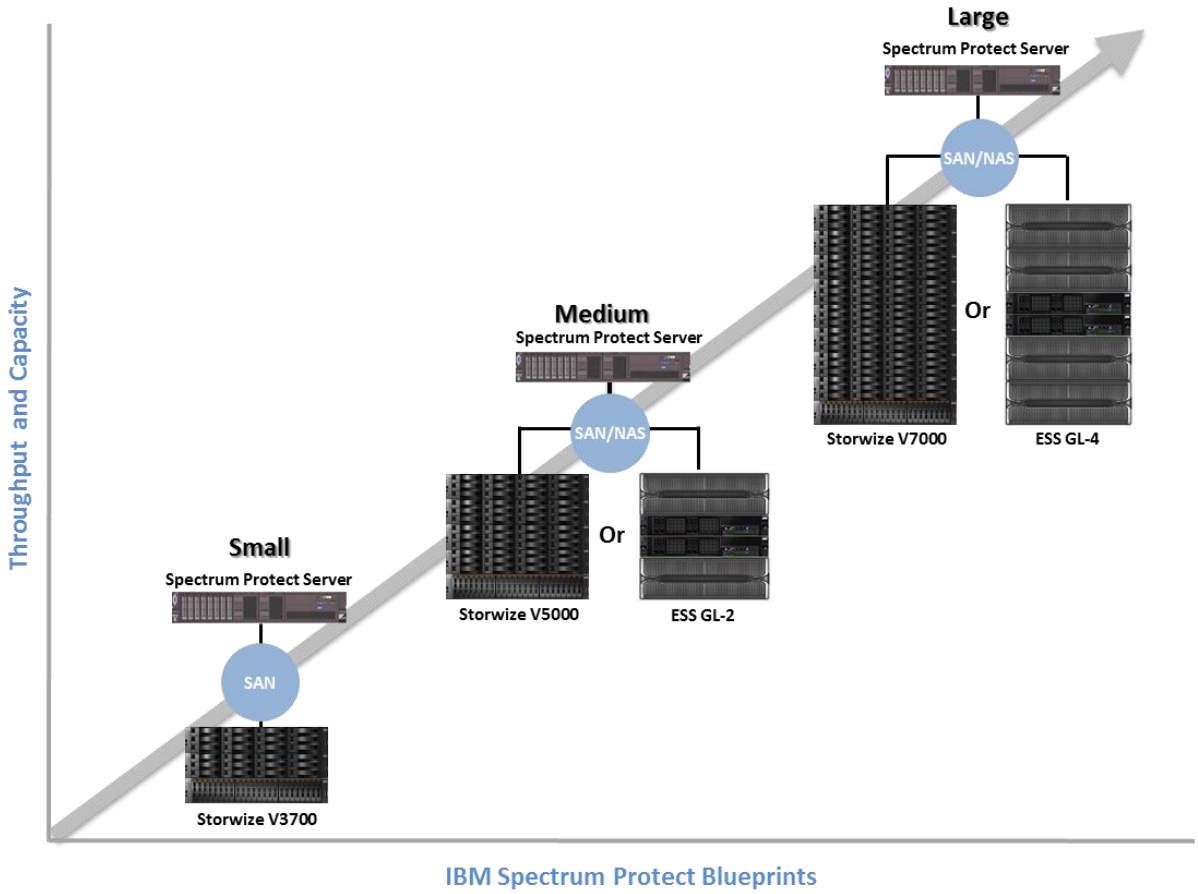

Design elements include automated validation of the hardware, file system, and operating system setup, and automated configuration of the IBM Spectrum Protect database, storage pool, policies, and schedules. Solution blueprints include:

- **Small Blueprint:** designed to process up to 6 TB of backup data per day, for environments with 45 to 180 TB of managed data. It can be deployed on Lenovo x3650 M5 or Power8 S822 servers and uses IBM Storwize V3700 storage.
- **Medium Blueprint:** designed to process 6 to 20 TB of backup data per day, for environments with 200 to 800 TB of managed data. It can be deployed on Lenovo x3650 M5 or Power8 S822 servers and uses IBM Storwize V5000 storage or IBM ESS GL-2.

- **Larger Blueprint:** designed to process 20 to 100 TB of backup data per day, for environments with 1,000 to 4,000 TB of managed data. It can be deployed on Lenovo x3650 M5 or Power8 S822 servers and uses IBM Storwize V7000 storage or IBM ESS GL-4.

Next, ESG Lab validated performance by auditing the results of IBM data ingest testing. Testing leveraged a large SAN-based blueprint configuration and a consistent workload from release to release with randomized backup data and randomized daily change patterns to simulate real-world workloads. Figure 12 shows improvements to backup duration between version 7.1.1 (using Volume Pools and post-process dedupe) and version 7.1.3 of Spectrum Protect. Version 7.1.3 introduced Container Pools and in-line deduplication that not only deliver faster ingest and replication rates, but also improve overall performance by completely eliminating the reclamation process, resulting in a significant reduction in deletion time when container pools are used.

Figure 12. Reduced Backup Duration

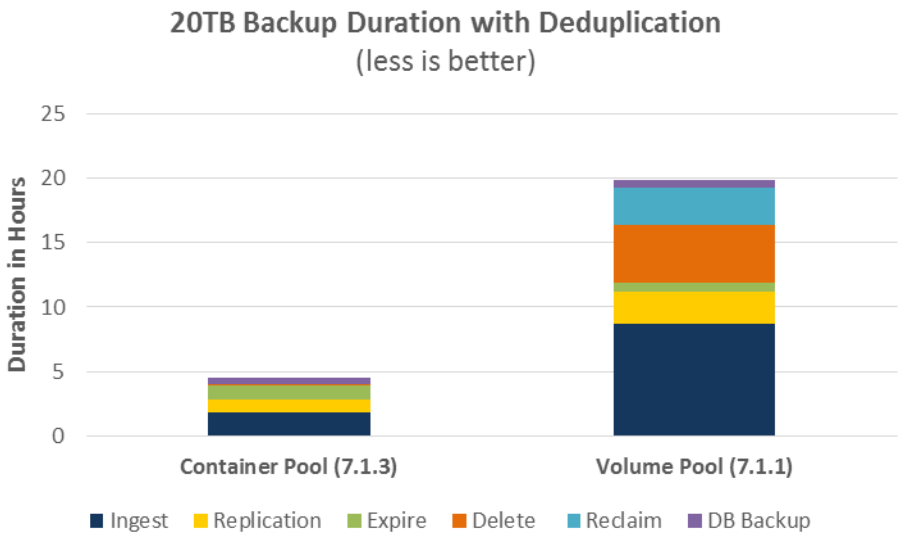

What the Numbers Mean

- Spectrum Protect 7.1.3 was ***4.7 times faster*** than 7.1.1 in completing backup operations.
- The introduction of the new container pool code in 7.1.3 enabled the performance gains.
- Enables IBM to hit the small, medium, and large blueprints performance targets.

Finally, ESG Lab confirmed the impact of in-line deduplication on the ingest portion of the backup process. Figure 13 shows the improvement in ingest throughput rates for container pool backups.

Figure 13. Improved Throughput

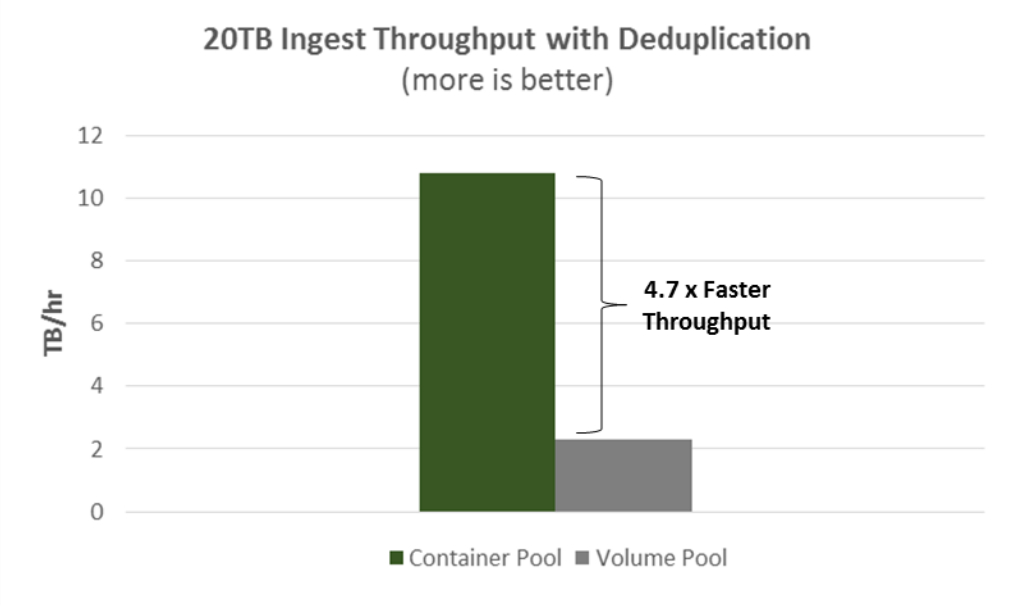

What the Numbers Mean

- The container pool with in-line deduplication delivered ***4.7 times faster*** throughput during data ingest.
- The new inline-deduplication process delivers ***10 times more*** usable storage capacity.
- Both container and traditional volumes pools are available in version 7.1.3 of Spectrum Protect.
    - The version upgrade does not force conversion to container pools.
    - This enables customers to better manage their migration process.

Why This Matters

What happens when new application data in a highly virtualized production environment turns out to be far larger than initially planned, and starts to push the performance and capacity capabilities of the backup environment to its limits? Unfortunately, for many data protection implementations this type of event often leads to a large and immediate investment in additional hardware.

ESG Lab validated that with each major release IBM invests significant effort in improving performance and efficiency. The addition of the new container pool technology helps reduce backup times and increase throughput while software-defined capabilities such as deduplication and replication can help reduce costs. This type of upgrade can often help eliminate immediate hardware purchases and can make CapEx easier to manage and control.

ESG Lab Validation Highlights

- ☑ ESG Lab was very impressed with the blueprint deployment process. It starts by giving customers three different options, small, medium, and large, depending on the size of their environment. It provides server and storage design guidance and even provides a script to test performance before the automated configuration is initiated.
- ☑ ESG validated that enhancements to the Operations Center make Spectrum Protect easy to manage and use. The GUI provides guided wizards for most administrator tasks and in most cases explains the impact of a configuration change before the administrator is allowed to commit the changes.
- ☑ ESG Lab was pleased to see a workload concept integrated in the management interface. We leveraged the workload view to display the impact of newly configured replication jobs on the backup process.
- ☑ We confirmed the self-service capabilities of the solution by conducting file-level restores with the new web-based interface. It provides an intuitive recovery interface for end-users that are not familiar with the Spectrum Protect application enabling unassisted end-user recovery.
- ☑ ESG Lab validated that Frontsafe is an effective application for delivering Spectrum Protect capabilities as a service. The Frontsafe application provides a feature-rich solution to help highly distributed organizations and managed service providers implement retail-style delivery of Spectrum Protect.

Issues to Consider

- ☑ At the time this report was written, Spectrum Protect cloud support was limited to SoftLayer. ESG Lab would like to see that support extended to more public cloud providers like Google, Amazon, and Azure to give Spectrum Protect customers even more infrastructure flexibility.
- ☑ ESG Lab would also like to see the IBM Spectrum Protect blueprints packaged as virtual appliances to make it just as easy to deploy the solution in the public cloud. ESG believes this would greatly extend the disaster recovery and business continuance capabilities of the solution.

The Bigger Truth

In large part, the software that handles your banking transactions is the same software that has been running them for 30 years. But today, you don't have to physically go into the bank to transact business—you can get cash from ATMs all over the world, scan checks and deposit them online, pay bills electronically, and transfer between accounts with the click of a mouse. The basic "crown jewels" of code that drive banking applications have remained intact, while the external wrappers have changed dramatically to accommodate the demands of today's users, leveraging today's advanced technologies.

IBM Spectrum Protect (formerly Tivoli Storage Manager, or TSM) has gone through a similar metamorphosis, shedding its mantle of cumbersome, time-consuming execution to become a fully modernized data protection solution: not just software-defined, integrated with virtualization, and cloud-capable, but dramatically simpler and faster to deploy and use. I can personally attest to the change. In a previous job some years ago I was responsible for setting up TSM solutions before delivery to customers—if all went well, it took about a week to get that job done. Today, as ESG Lab validated, we set up a Spectrum Protect server using a blueprint in about eight minutes. IBM data protection has always been good, but it has not always been easy—until now.

Organizations have always been confident trusting their data to IBM, with good reason. Spectrum Protect comes from a solid data management engine born in the mainframe space, where losing data is not an option. In this validation, ESG Lab was impressed with the modernization aspects, such as the blueprints and automation that make it very simple and fast to deploy. Ongoing management is equally modern: integration with Operations Center includes wizard-based tools, prescriptive guidance such as context-driven alerts and suggestions, and policy templates built from customer feedback. Frontsafe delivers a retail-like interface for management and billing from which enterprises and service providers benefit, and user self-service restore maximizes productivity for IT and end-users. These features all deliver better cost efficiency by reducing the management time and effort required for initial deployment and day-to-day management.

ESG Lab also validated key additions to support modern data centers such as cloud pools, simpler decommissioning to reclaim space quickly and easily, and workflow visibility for replication. Administrators can see when replication has fallen behind, ensuring that they are fully aware of the protected state of their environment. More effective deduplication and faster ingest were also validated, ensuring minimum storage capacity and bandwidth needs.

IBM Spectrum Protect today offers simpler management to the backup administrator, and has added user capabilities that enhance the solution in productivity and cost-effectiveness. In addition, today's Spectrum Protect enables people without data protection expertise to protect their data—small organizations such as physical practices can gain the full protection that the solution delivers without the management headache. That is a true demonstration of the metamorphosis that has taken place for this solution.

Kudos to IBM for the remarkable transformation they have accomplished with this product, keeping the crown jewels in place while delivering an agile, automated, modern protection solution that is easy to deploy, manage, and use. Spectrum Protect is a backup solution that really *supports* an organization—instead of being a ball and chain that drags you down.

Appendix

Table 1. ESG Lab Test Bed

Software	
IBM Spectrum Protect	Version: 7.1.3
IBM Spectrum Protect Operations Center	Version: 7.1.3
Frontsafe Cloud Portal	Version: 16.2.0.0
Microsoft Windows Server	Version: 2012
Linux Red Hat Enterprise	Version: 7.1
Servers	
Lenovo (small blueprint)	Model: x3650 M5 CPU: 2 x 6 core Xeon 2.4Ghz Memory: 64GB RAM Dual port 8Gb FC Dual port 10Gb Ethernet
Lenovo (large blueprint for performance results)	Model: x3650 M5 CPU: 2 x 16 core Xeon 2.3Ghz Memory: 192GB RAM 2 x dual port 8Gb FC 2 x dual port 10Gb Ethernet
Storage	
IBM Storwize (small blueprint)	Model: V3700 1 x SFF controller 2 x LFF expansion 24 x 15K rpm SAS 24 x 3TB NL-SAS
IBM Storwize (large blueprint for performance results)	Model: V7000 gen2 1 x SFF controller 19 x LFF expansion 8 x 800GB SSD 16 x 1.8TB 10K SAS 228 x 6TB NL-SAS

Enterprise Strategy Group | **Getting to the bigger truth.**

20 Asylum Street | Milford, MA 01757 | Tel: 508.482.0188 Fax: 508.482.0218 | www.esg-global.com