VERITAS™

Veritas Access Appliance with IBM Spectrum® Protect

Backup, Archive and Long-Term Retention Solution

This whitepaper provides a technical overview of Veritas Access Appliance as a secondary storage solution with IBM Spectrum Protect for backup, archival and long-term retention. It highlights the overall solution architecture components, integration flow, best practices, sizing guidance, and deployment of Access Appliance with IBM Spectrum Protect.

IBM® Validated Solution

Contents

Revision History

Version	Date	Changes
1.00	3/05/2019	Initial Version (Internal only)
1.01	07/03/2019	Modification of Access version
2.0	06/02/2021	Updates related to Access release v7.4.3

INTRODUCTION

EXECUTIVE SUMMARY

As data increases at an accelerating pace, companies are striving to determine the best strategy in the management, preservation, and retention of their valued data. There are several challenges that come to mind when talking about backup, archival and long-term retention storage solutions which include cost, complexity, control, and visibility. Traditionally, the solution for backup, archival or long-term retention has been to send data to tape because of its low cost. However, the complexity in tape management in addition to the time to restore has been an issue. Recently companies have looked to the public cloud for a possible solution, however, issues in total cost of ownership and control become a concern.

To address these challenges, Veritas has designed the Access Appliance as a purpose-built, on-premises secondary storage appliance for backup, archival and long-term retention use cases. Together with IBM Spectrum® Protect, a comprehensive data protection, data resource and space management solution, the Access Appliance provides a resilient and cost-effective storage platform for the preservation of data backups that companies want to retain and have readily available. In addition, Veritas Access V7.4.2. plus patch is a validated cloud object storage device for IBM Spectrum Protect.

SCOPE

The purpose of this document is to provide technical details to assist in understanding the Access Appliance with IBM Spectrum Protect as a solution for backup, archival and long-term retention of data. It describes the components of this solution, its value, sizing guidance, and some best practices. It is advised to refer to Veritas and IBM product documentation for installation, configuration and administration of each of the products discussed in this whitepaper. NOTE: This document gets updated periodically and if you downloaded a local copy of this document, please get the latest from this link.

TARGET AUDIENCE

This document is for customers, partners, and Veritas field personnel interested in learning more about the Veritas Access Appliance with IBM Spectrum Protect for backup, archival and long-term retention. It provides a technical overview of this solution, guidance in sizing, and highlights some best practices.

SOLUTION VALUE

The Access Appliance acts as an on-premises storage platform for data that has been backed up and/or archived using IBM Spectrum Protect. The integration of these solutions provides a compelling offering for the backup, archival and long-term retention use cases.

Key values of utilizing Access Appliance with IBM Spectrum Protect include:

- **Minimize cost**—Access Appliance provides a low-cost, disk-based solution that is easy to manage. With IBM Spectrum Protect deduplication feature, the amount of storage space is reduced by saving only one copy of the data blocks and having the duplicates point to that one copy, thus providing a more storage efficient solution and reducing overall costs.
- **Increase visibility and control**—more and more companies would like to leverage the data that has been backed up, archived and retained for IT or business analysis and investigations so having the data on-premises under the company's control and visibility allows for quick restores.
- **Improve Recovery Point Objective (RPO)/Recovery Time Objective**—on-premises storage for backups and archived data improves service levels over tape or public cloud.

SOLUTION KEY FEATURES

There are certain key features that companies look for in a backup, archive and long-term retention solution product such as storage efficiency, security, and ease of management. The Access Appliance with IBM Spectrum Protect provides these features to assist customers in preserving their most valued data.

STORAGE EFFICIENCIES

Support for storage efficiency is one of the main factors when choosing and purchasing a backup, archival and long-term retention solution. The ability to maximize storage space assists in reducing overall cost. IBM Spectrum Protect has both compression and deduplication features. Any data compressed and deduplicated by IBM Spectrum Protect is preserved on the Access Appliance.

Compression improves storage utilization by reducing the number of bits required to represent data. IBM Spectrum Protect provides client side and server side compression prior to sending data to the Access Appliance. The difference between client side and server side compression is where the compression occurs. By compressing at the client side, it reduces network bandwidth since less bits or data is being sent to the server. IBM Spectrum Protect does examine the data prior to compression and it will not perform compression if the data is not a good candidate for compression. In either scenario, the Access Appliance stores the compressed format of the data conducted by IBM Spectrum Protect components.

Similarly with deduplication feature, IBM Spectrum Protect offers both client side and server side deduplication. There is a capability to do inline or post processing deduplication. For even better storage utilization, utilize compression with deduplication. When using compression, data is first deduplicated prior to being compressed. Refer to Strategies to Minimize the Use of Storage Space for Backups for more details on this feature.

SECURITY

For enhanced security, IBM Spectrum Protect offers encryption of data. Any encryption done by IBM Spectrum Protect is maintained on the Access Appliance. The Access Appliance also has encryption capabilities in conjunction with an external Key Management System (KMS). The appliance encrypts the volume that the "file system" resides on. An external KMS such as IBM KMS is required to create the keys for the encryption.

At a transport level, IBM Spectrum Protect sends data over dedicated network ports to Access. Additional security that is employed for this solution is the requirement to use Access user keys and credentials when configuring Access as a cloud storage destination. When SSL is enabled, certificates are generated on Access and placed in IBM Spectrum Protect repository of security certificates (Java Key Store) and data is sent via HTTPS. When Access is utilized as a Network Attached Storage (NAS) share, the ownership of the share belongs to an IBM Spectrum Protect administrator, hence, limiting the access control of the contents.

AUTOSUPPORT FEATURE

Veritas Access Appliance has the ability to call home if the health monitoring services observe hardware or software issues. Veritas AutoSupport service provides proactive monitoring, alerting 24x7, automated support case management, and guided workflows on the health of the appliances. This feature alerts customers and/or service engineers to quickly handle the issue and reduce further risks. Enabling this feature can be done simply by registering the appliance(s) at the Veritas MyAppliance portal as shown in Figure 1 and enabling the call-home functionality.

Figure 1- MyAppliance Portal View

MONITORING AND INTRUSION DETECTION

Available on the Access Appliance is Symantec Data Center Security (SDCS), an intrusion detection system. SDCS is a real-time monitoring and auditing software. It performs host intrusion detection, file integrity monitoring, configuration monitoring, user access tracking and monitoring, and produces logs and event reports. SDCS adds security hardening and monitoring for the Access Appliance to reduce security risks and attacks. For more information on the Access Appliance intrusion detection system, refer to the Access Appliance Initial Configuration and Administration Guide.

SOLUTION ARCHITECTURE

At a high level, sources are backed up or archived utilizing IBM Spectrum Protect and the Access Appliance is the target secondary storage for the backup and archived data as shown in Figure 2.

Figure 2 - Solution Overview

In order to get a better understanding of how the Access Appliance integrates with IBM Spectrum Protect, all involved solution components are explained in further detail in the following sections.

IBM SPECTRUM PROTECT

IBM Spectrum Protect provides protection for a variety of data and platforms such as operating systems, virtual systems, databases and applications, files, and other content. It has the capability to backup data to tape, storage area network (SAN), network attached Storage (NAS), public or private cloud. Schedules, retention periods, and the ability to backup and archive data to different types or tiers of storage are defined in policies.

A typical IBM Spectrum Protect environment consists of the following components:

- **Server**—manages and controls the backup and recovery activities, hosts the database, and maintains logs. Elements of the server include:
  - **Database**—contains information relating to the policies and schedules, metadata about the backups, archives, and migrations and server settings.
  - **Logs**—keeps records of database transaction in which the database utilizes to validate data consistency in the database. The logs include an active log, containing current transactions, and an archive log holding copies of log files that were closed. Optionally a log mirror, a copy of active logs and an archive failover log, a secondary copy of the archive log can be configured for redundancy.
  - **Storage Pool**—logical storage pools that map to varying types of storage such as local disks, tape, SAN, NAS, and cloud (on-premises and off-premises). The storage pool is the target for backup or archive data
- **Clients**—client components are installed on hosts that have the data to be backed up and responsible for sending and receiving data to and from IBM Spectrum Protect server for backup and recovery. There is also a backup and archive allowing end-users.
- **Operation Center**—a web-based console for managing, monitoring, and reporting of IBM Spectrum Protect servers and clients. The command-line interface (CLI) is also available within this console.

The components of IBM Spectrum Protect can be run on a single system. However, multiple systems running different instances of IBM Spectrum Protect handling different sources to backup or archive can be managed by a single administration console - Operations Center. Figure 3 illustrates a sample configuration of an IBM Spectrum Protect environment and its components. The clients communicate to the server over TCP/IP via a specified port (between 1000 to 32767, default 1500). Administration via the Operations Center is over HTTP(s).
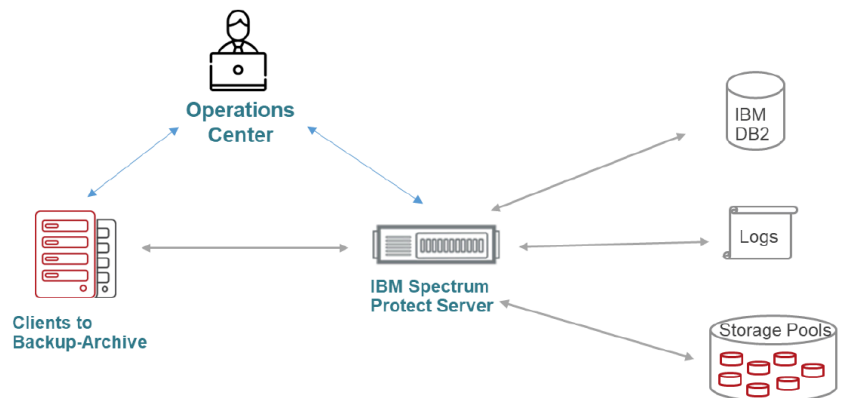

Figure 3 – An Example of an IBM Spectrum Protect Environment

ACCESS APPLIANCE 3340

IBM Spectrum Protect can send backup and archival images to various storage types (disk, tape, cloud, etc.). For those seeking an on-premises disk-based solution for faster recovery times, control and/or simplicity when compared to tape or cloud, Veritas has developed the Access Appliance for ease of acquisition, management, and support. Access Appliance is a cost-optimized, turn-key storage solution designed for high capacity, making it well suited for long-term retention.

Figure 4 - Access Appliance Rack Units

The Access Appliance model 3340 is comprised of two clustered nodes and one primary storage shelf and up to three additional expansion storage shelves. The appliance can scale up to 2,800 TB of usable space as can be seen in Figure 4.

Highlights of Access Appliance 3340 specifications are shown in Table 1. Refer to the Access Appliance datasheet for more detailed information.

Model	CPU Processor	RAM	Ports		Capacity	Rack Units
			1 GbE	10 GbE		
3340 (2 nodes)	2 x Xeon® 4108 (1.8 GHz) per node Total: 16 core per node	384 GB per node	4 per node	2 per node	280 TB – 2800 TB (254 TiB – 2544 TiB)	Server: 2U Storage per Shelf: 5U

Table 1 - Highlights of Access Appliance Specifications

Note: TB - Capacity values are calculated using Base 10; TiB - Capacity values are calculated using Base 2.

The two nodes are clustered in active/active configuration such that each node can handle I/O requests. Storage shelves are connected to each node and configured with dynamic multipathing so I/O can be sent to either node for performance and availability. The redundant hardware RAID controller in the primary storage shelf configures and presents the shelves' physical disks into disk groups (volumes) protected by a RAID 6 storage layout. With a RAID 6 configuration, data with dual parity is striped across the configured volumes (5 volumes per storage shelf with each volume containing 16 disks and each shelf with 2 hot spares). Each data volume can remain operational despite two concurrent disk failures.

The nodes run RHEL 7.4 or later as the operating system platform and Access software version 7.4.2 or later. The Access Appliance is a scale-up NAS platform that supports multiple protocols, including NFS, SMB, FTP, iSCSI and S3. With IBM Spectrum Protect, the Access Appliance is seen as an S3, NFS, or SMB target.

When using Access as an S3 target, data written to Access from IBM Spectrum Protect is placed in an S3 bucket. A bucket maps to a file system of type "cluster filesystem". The Access Appliance supports a maximum usable capacity of 2.8 PB and thus the maximum size of an S3 bucket in an appliance is 2.8 PB. When using the S3 protocol to backup or archive data, the S3 object URL, access key and secret key are presented to clients. The IBM Spectrum Protect server utilizes this URL as the S3 endpoint for reading and writing to the Access bucket. A dedicated S3 communication port, 8143, is required for both HTTP and HTTPS and thus firewalls have to keep this port open.

Similarly, Access can also be an NFS or SMB target. A "cluster filesystem" type is created where shares on that filesystem can be exported and mounted or mapped by the IBM Spectrum Protect server. The shares should be exported at the minimum with "rw" and sync option set on Access. After mounting or mapping share on the server, permissions of IBM Spectrum "administrator" user would need to be set appropriately for the share exported.

For management, the appliance can be managed by the command-line shell referred to as the CLISH and/or a web-based graphical user interface (GUI) where one can provision storage pools, create filesystem and provision Access as an S3, NFS, or SMB target.

NOTE: For an example of how to deploy and configure the Access Appliance as an S3 target with IBM Spectrum Protect, refer to the Appendix section of this whitepaper.

SOLUTION INTEGRATION

This section explores how all these components integrate together and how data flows through each component. The policies configured in IBM Spectrum Protect define where the client's backup and archive data will be stored and how long to retain it. A client is associated or bound to one active policy set within IBM Spectrum Protect. The Access Appliance acts as a storage platform for the data being backed up or archived by IBM Spectrum Protect.

IBM Spectrum Protect maps various type of storage platforms such as NAS, SAN, tape or cloud to a logical construct referred to as storage pool. In context with the Access Appliance, the types of IBM Spectrum storage pools that can be configured include:

- Container Storage Pool—automatically deduplicates data inline or as it is ingested. No device class or volumes need to be defined. There are two types of container storage pool that include:

  - Directory Container—file based storage using filesystem directories. If multiple directories are specified within a directory container, the data is distributed across the available directories.

  - Cloud Container—data is sent in object format using TCP/IP connection and target storage is an on-premises cloud object storage.

- File Device Class Storage Pool—legacy storage pools in which a device class of type file would need to first be defined. Data is written to this storage pool sequentially as in tapes. The File device class storage pool only supports post-processing deduplication which is data is deduplicated after it has been ingested.

There is no option to disable the deduplication on container storage pools and thus if your data are not good candidates for deduplication, such as encrypted or compressed data, it is recommended that the legacy FILE device class storage pool be utilized. The next sections provide a more detail view of Access as a cloud container, directory container and FILE device class storage pool for IBM Spectrum Protect. However, highlights of the differences are shown in Table 2.

Feature	Storage Pool Type		
	Cloud Container	Directory Container	File Device Class
Protocol	S3	NFS/SMB	NFS/SMB
Storage Pool	One bucket per storage pool. More than one storage pool can point to the same bucket.	Can specify multiple directories per storage pool	Can specify multiple directories per storage pool
Deduplication	Automatic Inline	Automatic Inline	Post Processing
Read Sizes	Range reads of 10 KB – 100 KB	Range reads of 10 KB – 100 KB (avg. 256 KB)	256 KB
Write Sizes	S3 Multi-part uploads (100 MB default).	50 KB – 4 MB (avg. 256 KB)	256 KB
File Sizes	Varies but maximum 1 GB (Default)	Varies but maximum 10 GB (Default)	Varies but maximum 2 GB (Default)
Ports	8143	NFS: 2049, 111, 4001, 4045 SMB: 139, 445	NFS: 2049, 111, 4001, 4045 SMB: 139, 445
Transmission	Asynchronous	Synchronous	Synchronous

Table 2- Highlights of the Differences in the IBM Spectrum Protect Storage Pool Types for Access

CLOUD CONTAINER STORAGE POOL

Configuring Access as a cloud container, deduplicated data is sent asynchronously to the Access Appliance via the S3 protocol from a local disk cache on the IBM Spectrum Protect server. To reduce network contention and for performance, a local disk cache is required on the IBM Spectrum Protect server to initially stage the data prior to sending to an Access S3 bucket as shown in Figure 5. Per IBM Spectrum Protect documentation, it is recommended to have at least 3 TB of local cache or enough cache to handle a single days' worth of backup data. If there is a fast network connection, then a smaller staging area can be used. If compression and encryption is enabled, data is first deduplicated, compressed and then encrypted prior to sending to a cloud container. Multiple threads can process multiple containers at a time. Thus, backups involve disk writes to a cache and as the containers are filled, data is sent to Access. Transfers to Access are conducted using large I/O sizes and multi-part uploads. Once transferred, data is removed on local cache. The Access object URL, access and secret keys, and bucket name are required when configuring Access as cloud container storage pool. Multiple storage pools can be configured to point to a single bucket or each storage pool has its own bucket, but multiple buckets cannot be assigned to a single storage pool.
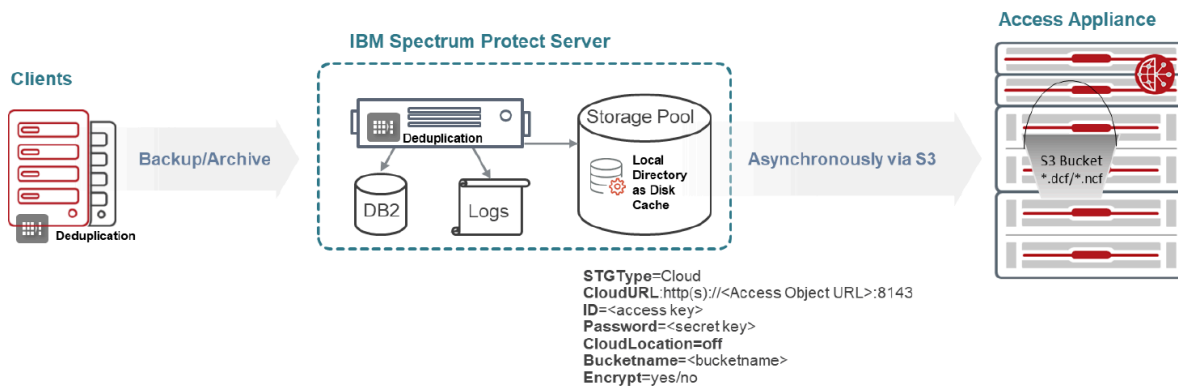
Figure 5 – Access as a Cloud Container Storage Pool

When using a cloud container, both client side and server side deduplication is supported. In client-side deduplication, the data is deduplicated prior to being sent to IBM Spectrum Protect whereas server side deduplication is conducted on the IBM Spectrum Protect server. Decision-making on where deduplication occurs highly depends on network bandwidth and compute resources available on clients and server.

IBM Spectrum Protect deduplication algorithm involves using a variable chunk segment size for analysis of duplicate data within the container pool. Only unique data is stored and a reference pointer is used when duplicate data is identified. IBM Spectrum Protect data or metadata inside container storage pools are stored into parts referred to as extents. The extents can range in size from 50 KB to 4 MB with an average of 256 KB. Any data smaller than 2 KB or data that cannot be deduplicated, such as encrypted or compressed, are not deduplicated. Data that are deduplicated are stored on Access as files with extension names *.dcf (deduplicated container file) or with *.ncf (non-deduplicated container file) extensions. A

Figure 6- View of Container Files on Access

sample view of a container file is shown in Figure 6. By default, these files are 1 GB in size and are configurable using the parameter CloudTransferContainerSize server option (i.e. specifying in dsmserv.opt or using the "setopt" server command). These files are transferred to Access using S3 multipart-upload. With 1 GB default file size, the default part size that the file is broken up to is 100 MB. This value is configurable using the server parameter, CloudMinUploadPartSize. For restores, IBM Spectrum Protect does range reads in smaller sizes of 10 KB – 100KB.

Note: Access Appliance v7.4.2 plus patch was utilized to conduct the S3 validation with IBM Spectrum Protect v8.1.5.

DIRECTORY CONTAINER STORAGE POOL

Using a directory container pool, deduplicated data is sent to the Access Appliance synchronously via NFS (Linux environment) or SMB (Windows environment) protocol as pictured in Figure 7. The file system created on the Access Appliance is exported as a read/write and synchronous share and is either mounted to a directory or mapped to a drive letter on the IBM Spectrum Protect server. As previously mentioned, the data within container storage pools are automatically deduplicated inline prior to being sent to the storage platform and both client and server-side deduplication is supported.

Figure 7 - Access as a Directory Container Storage Pool

As in cloud container storage pools, the same algorithm and method is used for deduplication and the files on an NFS/SMB share would have files with extensions of *.dcf and *.ncf for deduplicated and non-deduplicated data respectively. The size of the container or files can vary depending on the amount of free space that is available in the storage pool directories and how many concurrent containers or files that can be opened to avoid using all the free space. However, the default container or file size is 10 GB. This size can be changed with the ContainerSize server option, specified in units of megabytes. Figure 8 provides an example view of the directory container storage pool contents of an NFS share. When multiple directories for the directory container are specified, data distributes the write across the directories improving performance, assuming that different volumes, disks or filesystems are assigned to each directory.

```
[root@tmesavm8157 00]# ls -h
000000000000022.ncf  000000000000024.dcf  000000000000026.dcf  000000000000028.dcf  00000000000002a.dcf
000000000000023.dcf  000000000000025.dcf  000000000000027.dcf  000000000000029.dcf
[root@tmesavm8157 00]# ls -lh
total 624M
-rw-------. 1 tsmuser tsmgrp 4.6K Oct 24 13:27 000000000000022.ncf
-rw-------. 1 tsmuser tsmgrp 100M Oct 24 13:27 000000000000023.dcf
-rw-------. 1 tsmuser tsmgrp  53M Oct 24 13:29 000000000000024.dcf
-rw-------. 1 tsmuser tsmgrp 100M Oct 24 13:28 000000000000025.dcf
-rw-------. 1 tsmuser tsmgrp 100M Oct 24 13:28 000000000000026.dcf
-rw-------. 1 tsmuser tsmgrp 100M Oct 24 13:29 000000000000027.dcf
-rw-------. 1 tsmuser tsmgrp 8.7M Oct 24 13:29 000000000000028.dcf
-rw-------. 1 tsmuser tsmgrp 100M Oct 24 13:29 000000000000029.dcf
-rw-------. 1 tsmuser tsmgrp  64M Oct 24 13:29 00000000000002a.dcf
[root@tmesavm8157 00]#
```

Figure 8 - Example Contents of a Directory Container Storage Pool on Access

FILE DEVICE CLASS STORAGE POOL

Implementing a FILE device class storage pool stores data in a sequential fashion as on tape, however, data is stored in files. Files are viewed by IBM Spectrum Protect as volumes. Exported Access NFS or SMB shares are mounted to a directory or mapped to a drive on the IBM Spectrum Server and then the directory or mapped drive is specified when defining the file device class (i.e. def devclass nasclass1 devtype=file mountlimit=20 maxcapacity=20g directory='/mnt/dir'). As for other considerations when configuring FILE device class, refer to "Defining Sequential-Access Disk (FILE) Device Classes". Data can be sent synchronously or asynchronously to the Access Appliance via NFS or SMB protocol as shown in Figure 9. However, IBM recommends to mount share synchronously.

Figure 9 - Access as a File Device Class Storage Pool

The size of each file is defined by the parameter MAXCAPACITY where default is 2 GB. This value should not exceed the maximum size of file supported on the Access Appliance which is the maximum usable space of 2.8PB. For more details on this parameter, refer to "Optimal number and size of volumes for storage pools" in IBM product documentation.

```
[root@tmesavm8157 accessfs1]# ls -lh
total 3.9G
-rw-------. 1 tsmuser tsmgrp 3.9G Oct 24 11:47 000000d4.bfs
drwxr-xr-x. 2 root   root     96 Oct 17 10:35 lost+found
```

Figure 10 - Example Content a FILE Device Class Storage Pool on Access

DISASTER RECOVERY

Having a disaster protection plan is imperative for business continuity. IBM Spectrum Protect offers node replication where copies of the data are incrementally copied automatically to a remote or off-site IBM Spectrum Protect server. Client restores failovers to target if primary server is down. An Access Appliance can be deployed at the target side as the storage platform to store the replicated data. For information on how to best protect an IBM Spectrum Protect environment and node replication, please refer to IBM Product Documentation and Strategies for Disaster Recovery with IBM Spectrum Protect and Node Replication Guidelines. **Note**: A cloud container storage pool can NOT be a source target pool, however, it can be a destination pool as shown in Figure 11. Also, the use of PROTECT STGPOOL command which allows for only the data to be replicated without the associated metadata is not supported for cloud container storage pools.

Figure 11 - Example of Access as an S3 Target for Replication

BEST PRACTICES AND RECOMMENDATIONS

Following best practices is important in creating an optimum deployment. This section covers some best practices relating to the Access Appliance as a backup, archival and long-term retention storage for IBM Spectrum Protect.

DATA LAYOUT ON ACCESS APPLIANCE

The appliance contains hardware RAID 6 controllers and inherently does striping with dual parity across disks on the storage shelves for high performance and data durability. Selecting other layouts such as mirrored or erasure coding layout for data protection is not necessary. As a best practice, for backup, archival and long-term retention use case, it is recommended to configure the Access Appliance using defaults such as clustered file system, simple layout, and block size of 8 KB. For additional performance, the layout can be configured to be stripe. NOTE: To maintain the stripe performance, when growing the storage pool, the volumes must be added in multiples of the stripe columns (the entire shelf), and thus, it is advisable to plan or size the system appropriately.

ACCESS PROTOCOLS

IBM Spectrum Protect storage pools can map to Access as an S3 bucket, an NFS share or SMB share. IBM has technical validations and certification processes for when sending data to a cloud container storage device via S3 protocol. After completing this validation and certification process using the Access Appliance, Access v7.4.2 plus patch is now supported as an IBM Spectrum Protect cloud object storage device. Refer to the following IBM web page indicating the support: https://www-01.ibm.com/support/docview. wss?uid=swg22000915. As previously mentioned, when conducting restores from a cloud object storage device, IBM Spectrum Protect does range reads of 10 KB to 100 KB in size. If there is a requirement for a faster RTO then it is recommended to make use of disk-to-cloud tiering using "storage rules" which was introduced in IBM Spectrum Protect 8.1.3. For some additional recommendations when using Access as a cloud object storage, refer to IBM documentation relating to Optimizing Performance for Cloud Object Storage.

There is no formal validation or certification when configuring Access as an NFS target for IBM Spectrum Protect, however, IBM provides some recommendations in the below links relating to when using NFS:

- http://www-01.ibm.com/support/docview.wss?uid=swg21470193
- http://www-01.ibm.com/support/docview.wss?uid=swg22005129

In general, it is recommended to mount the NFS share with "sync" option, not to use a device type of disk (i.e. devtype=disk) and to conduct a proof of concept or tests to validate that NFS provides the Service Level Agreement (Recovery Point Objective/Recovery Time Objective) required by the customer in addition to running periodic audits on the storage pool volume on an NFS mount. Similar recommendations should be followed when using Access SMB protocol.

DEDUPLICATION AND COMPRESSION

Deduplication is done inline when defining a container pool such as directory or cloud container storage pool. As a FILE device class, deduplication can be enabled or disabled. When using container storage pools or enabling deduplication, more server resources are needed such as DB capacity, CPU and memory and should be sized appropriately. Notable recommendations stated in Effective Planning and Use of IBM Spectrum Protect Container Storage Pools and Data Deduplication include:

- Deduplication is recommended for backup data that is less than 4 PB total and limited to 100 TB per day for each server instance.
- For the Database, use high performance SSD or flash disk.
- Use only for data that can benefit from deduplication. Data such as encrypted or does not have much change will not benefit from deduplication.

For compression, it is an option and best to be done before data is deduplicated. Thus, if using client-side compression, client-side deduplication must also be enabled.

ACCESS S$_3$ TUNABLES FOR PERFORMANCE

When utilizing Access as an S3 target for IBM Spectrum Protect, there are tunables on Access that is recommended to be modified for improved performance. The following tunables can only be set using Access command-line interface and would require a restart of the object services on Access:

- cf_max_s3_threads – controls the number of threads to handle S3 requests
- cf_mp_calc_md5 - MD5 internal calculation at multipart merge

It has been observed that increasing the number of cf_max_s3_threads to 128 and disabling of cf_mp_calc_md5 provided the optimum performance for IBM Spectrum Protect workload.

NETWORK CONNECTIVITY

The Access Appliance has two 10 GbE uplinks per node. Each physical port maps to a virtual IP. Thus, there are four virtual IP addresses. Always present the virtual IP to clients or client applications so it will automatically transition to the other node if one node fails or the physical links on one node fails or is unreachable. For instance, map one of the virtual IPs to the S3 object URL when using the S3 protocol or use the virtual IP when mounting the NFS share on the server.

Bonding is an option on Access Appliance. Joining or bonding multiple network interfaces on the Access appliances into a single interface improves the bandwidth and network throughput through the combined single interface. Bonding is only configurable via the Access command-line interface. As a best practice, the switch that the uplinks of the Access Appliance are connected to should be configured appropriately for the link aggregation.

LOAD BALANCING

There are two nodes on the Access Appliance configured as active/active. As a best practice, balancing the load across nodes on Access is recommended. Load balancing can be achieved using any of the following techniques:

- **External load balancing**—using an external load balancer such as HAProxy or F5, allows for more algorithms to distribute load across nodes such as least connections or weights. It also frees the Access nodes from the proxy handling and balances the network traffic between the nodes.
- **Manual load balancing**—virtual IP addresses of the nodes can be manually assigned to applications in a distributed manner. The disadvantage of this approach is that even distribution might be difficult to gauge since applications are not all equal in sense of workload.
- **DNS load balancing**—the S3 object URL name for an Access S3 bucket, s3.<clustername> is created in DNS and includes all the virtual IP addresses of the nodes. DNS round-robins through the virtual IP addresses. The disadvantage of using DNS is in case of connectivity issues, the virtual IP is still in rotation until it is manually removed.

MONITORING

It is important to monitor or be aware of the alerts especially storage utilization warnings and hardware critical alerts. The AutoSupport features assists in this manner, but as a best practice, it is advisable to be pro-active instead of re-active. For instance, once the capacity reaches 60%, it might be a good time to revisit the storage utilization or plan for growth.

SIZING GUIDANCE

In planning for data protection, two considerations come to mind: recovery point objectives (RPO) and recovery time objectives (RTO). From a backup and recovery standpoint, the RPO and RTO determine which policies are implemented, and therefore the resources required by an IBM Spectrum Protect deployment in terms of the necessary amount of systems, appliances, and storage. Other considerations include the number of users and applications, amount of data that is backed up, the frequency, and how long to keep the data. When planning for a long-term retention storage solution for backup data there are two factors:

- ▪ Capacity - the amount of data that can be stored.
- ▪ Performance – how much workload (concurrent streams and bandwidth) the storage platform can handle.

The Access Appliance can store up to 2.8 PB worth of backup, archival and long term retention data. It is best to confer and work with both IBM and Veritas personnel for sizing of the IBM Spectrum Protect resources and the Access Appliance. Some parameters that might enter in the equation when estimating backup, archival and long-term storage requirements include:

1. Volume of source data.
2. Daily data change ratio
3. Annual storage growth
4. Data retention for daily incremental.
5. Retention for weekly, monthly and yearly full backups.
6. Estimated deduplication ratio for initial backup and daily incremental.
7. Estimated deduplication ratio for weekly, monthly, and yearly full backups.
8. Compression/Encryption enabled or disabled, client-side or server sided.
9. Performance and/or service level requirements.

There is an IBM Spectrum Protect blueprint that provides examples of a small, medium and large environments and has some guidance on how to configure IBM Spectrum Protect servers. There is also IBM Spectrum Protect Optimizing Performance document available for reference.

CONCLUSION

The addition of the Access Appliance in the Veritas appliance portfolio provides a competitive disk-based solution for backup, archive and long-term retention of data. Integrated with IBM Spectrum Protect, the Access Appliance becomes a great option in data protection, and preservation of critical digital assets. Implementing the Access Appliance with IBM Spectrum Protect as a backup, archival and long-term retention solution minimizes costs, improves control and visibility. CONCLUSION

REFERENCES

- IBM Spectrum Protect
  - Product Documentation
    - https://www.ibm.com/support/knowledgecenter/SSEQVQ/landing/welcome_sseqvq.html
  - Data Storage Concepts in IBM Spectrum Protect
    - https://www.ibm.com/support/knowledgecenter/en/SSEQVQ_8.1.5/srv.solutions/c_stg_overview.html
  - IBM Spectrum Protect Blueprints
    - Blueprints –
      - https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli%20Storage%20Manager/page/IBM%20Spectrum%20Protect%20Blueprints
    - Spectrum Protect Blueprint Sizer (may require login)
      - https://www-356.ibm.com/partnerworld/wps/servlet/mem/ContentHandler/TSJ03927USEN/lc=en_ALL_ZZ?mhq=IBM%20Spectrum%20Protect%20Blueprints&mhsrc=ibmsearch_a
    - Cloud Blueprints –
      - https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli%20Storage%20Manager/page/IBM%20Spectrum%20Protect%20Cloud%20Blueprints
  - Cloud Container Storage Pool FAQ
    - https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli%20Storage%20Manager/page/Cloud-container%20storage%20pools%20FAQs?section=cloudenv
  - Directory Container Storage Pool FAQ
    - https://www.ibm.com/developerworks/community/wikis/home?lang=en&_escaped_fragment_=/wiki/Tivoli%2520Storage%2520Manager/page/Directory-container%2520storage%2520pools%2520FAQs
  - Global Solutions Directory
    - https://www-356.ibm.com/partnerworld/gsd/solutiondetails.do?&solution=55426&lc=en
- Access Appliance
  - Product Documentation
    - 7.4.2 - https://sort.veritas.com/documents/doc_details/AAPP/7.4.2/Appliance%203340/ProductGuides/
    - Installation and Configuration - https://sort.veritas.com/DocPortal/pdf/129305376-133498283-1
    - 7.4.2 plus patch - https://www.veritas.com/support/en_US/article.100045858.html
    - 7.4.3 - https://sort.veritas.com/documents/doc_details/AAPP/7.4.3/Veritas%203340/Documentation/

APPENDIX

This section has only sample configurations and readers are expected to refer to the Veritas Access Product Documentation and IBM Spectrum Protect documentation for definitive and specific installation, administration and configuration details.

The deployment example walks through the configuration of the Access Appliance as an S3 target with IBM Spectrum Protect include as shown in Figure 12.

Figure 12 - Configuration Used in this Example Deployments

Highlights of the steps involved in this example include:

- Configuration and storage provisioning of the Access Appliances as an S3 target.
- Configuration of the Access S3 bucket as a cloud container pool in IBM Spectrum Protect
- Modification of the STANDARD domain policy to utilize the storage pool mapped to an Access S3 bucket
- Validation of configuration by conducting an archive and retrieval using the Backup-Archive client tool
- SSL configuration example

ACCESS APPLIANCE STORAGE CONFIGURATION AND PROVISIONING

In this example, the Access Appliance graphical user interface and command-line is utilized to provision the storage. It is assumed that the Access Appliance has already been installed and connected to the same network as the IBM Spectrum server. Highlights of steps to configure and provision Access Appliance as an S3 target involves the following:

a. Configure the storage disk pools

b. Set the default S3 parameters (i.e. layout, size, filesystem type, storage pool, etc.), start the S3 service and generate the S3 keys for the user,

c. Create a bucket via Access CLI where a filesystem will manually be created and map a bucket to the filesystem.

CONFIGURE STORAGE POOLS

Step 1) Using a web browser, login to the Access Appliance graphical user interface: https://<consoleIP>:14161. Click on "Go to HomePage".
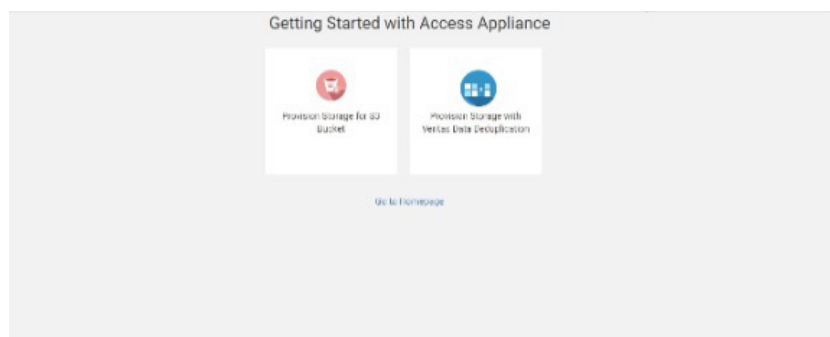
Step 2) Click on **Infrastructure** on left pane. Select the number of disks or volumes to be added to storage pool and click "**Add to Storage Pool**".
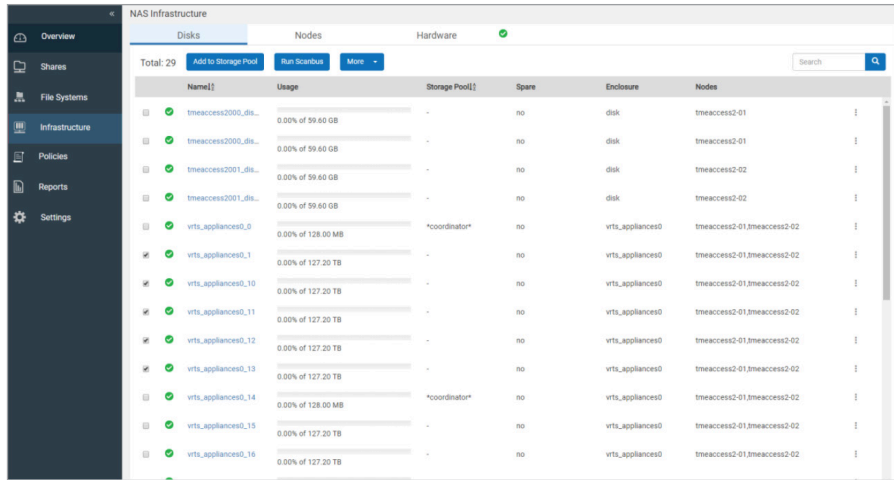
Step 3) Select "**Add to new storage pool**" and enter a **Storage Pool Name**, i.e. tsmpool and click **Next**. On next screen, click **Finish**.
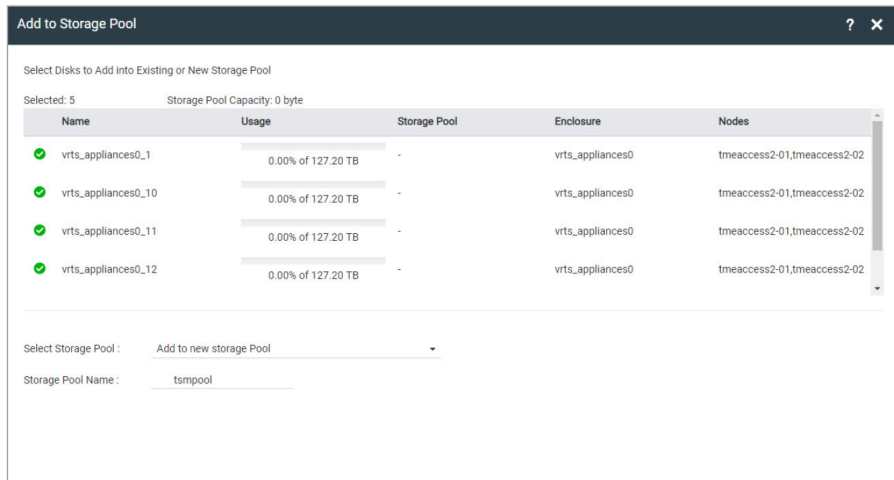
Step 4) Check the Activity icon, the clock on top and click on "Show All Recent Activities". Wait until pool creation succeeds.
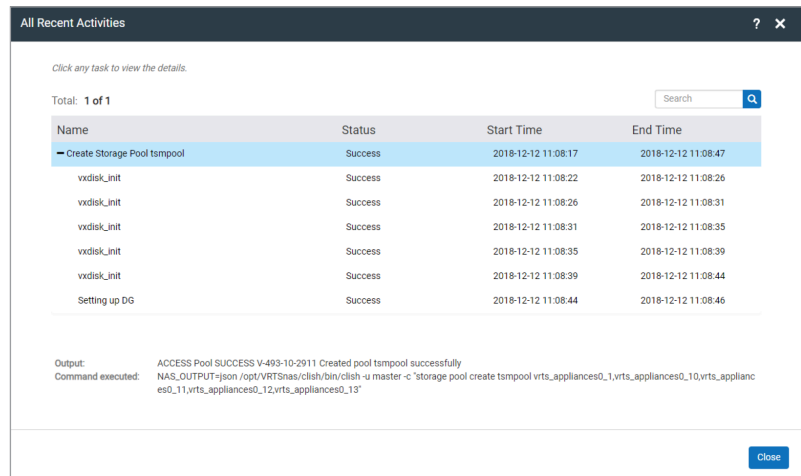
Step 5) Next step is to set the S3 parameters when creating a bucket on Access. Click on **Settings** on left pane and click on "**S3 Management**".
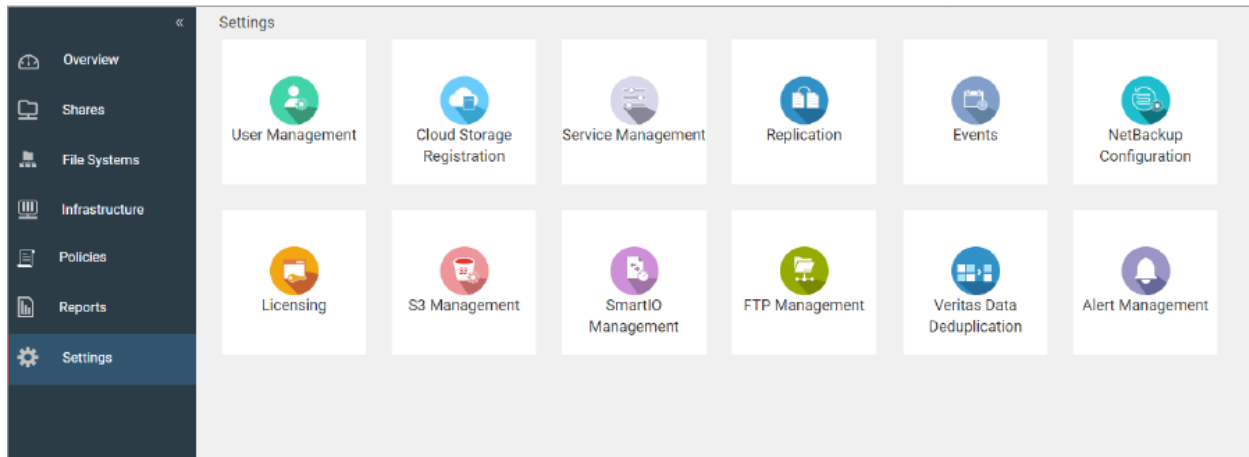
Step 6) Click on **Edit** on the Default parameters for S3 buckets.

Step 7) Modify the default parameters such as **file system size, file system layout, block size, and select the storage pool to use, ssl**, etc. These default parameters are used when a bucket is created by other applications, e.g. S3 Browser. In this example, a bucket is created with a cluster filesystem of size 5 TB with a simple layout inside storage pool tsmpool with SSL not enabled. Click **Next**.

Step 8) Click on **Activity** icon and click on "**Show All Recent Activities**" and wait until operations succeeds.

Step 9) By default, the S3 URL is s3.<clustername>. However, in this example, one of the virtual IP addresses will be used instead. Thus, to modify the S3 URL to use the IP addresses on each node prior to starting the S3 Service, **ssh as admin to the Access CLI** (i.e. ssh admin@<IP of console>). Enter in objectaccess and then use the "**set data_endpoints**" command to set the new name or IP. As shown in the example below, "set data_endpoints 10.182.81.86" is specified. More than one endpoint or virtual IP address separated by a comma can be specified, for instance, "set data_endpoints 10.182.81.86,10.182.81.87". **NOTE**: For all the aliases to be picked up, one would need to specify all of the aliases in a single set command. Also, any changes made to the endpoints would require the objectaccess server to be restarted via GUI or CLI if already started.

```
tmeaccess2>objectaccess
tmeaccess2.ObjectAccess> set data_endpoints 10.182.81.86
ACCESS ObjectAccess SUCCESS V-493-10-0 Data endpoints set successfully.
tmeaccess2.ObjectAccess> show
Name              Value
========================  ============================
Server Status      Enabled
Admin_URL          http://admin.tmeaccess2:8144
S3_URL             http://10.182.81.86:8143
SSL certificate validity  -
admin_port         8144
s3_port            8143
ssl          no
pools        tsmpool
fs_size       5T
fs_type       simple
fs_ncolumns       5
fs_stripeunit      512
fs_blksize       8192
fs_pdirenable      yes
fs_encrypt       off
```

Step 10) On the Access GUI, enable the S3 Service by clicking on the "**S3 server status**" slider and moving it to the right.
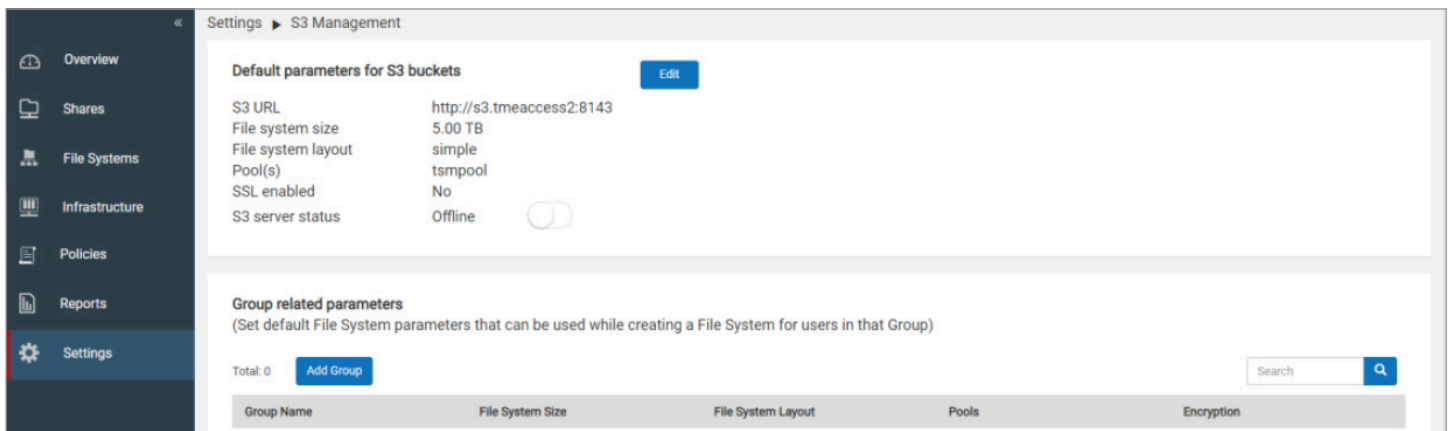
Step 11) Click **Start**.
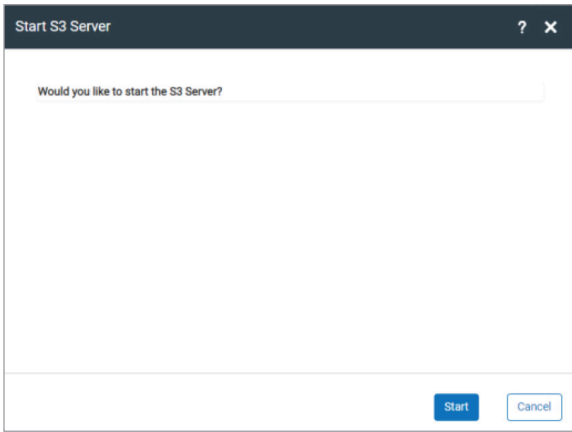
Step 12) Monitor the **activity** and wait until complete.
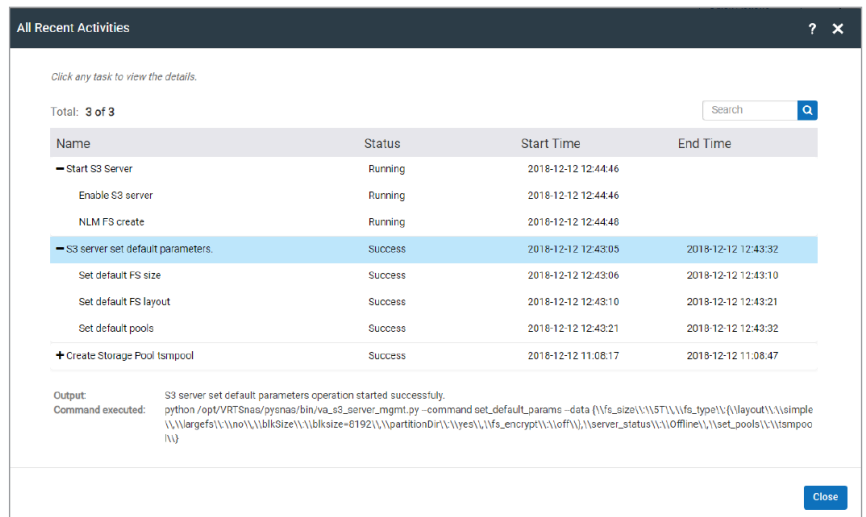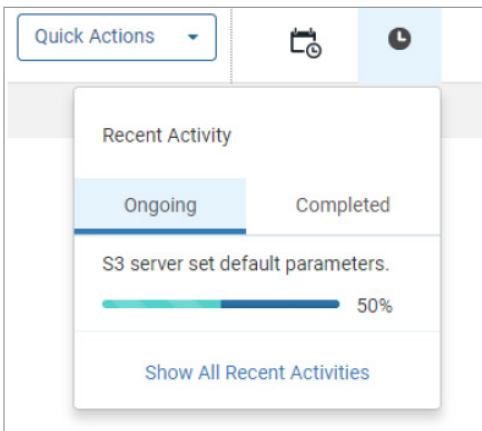
Step 13) Under the S3 Management page at the bottom, click on **Create Keys**.
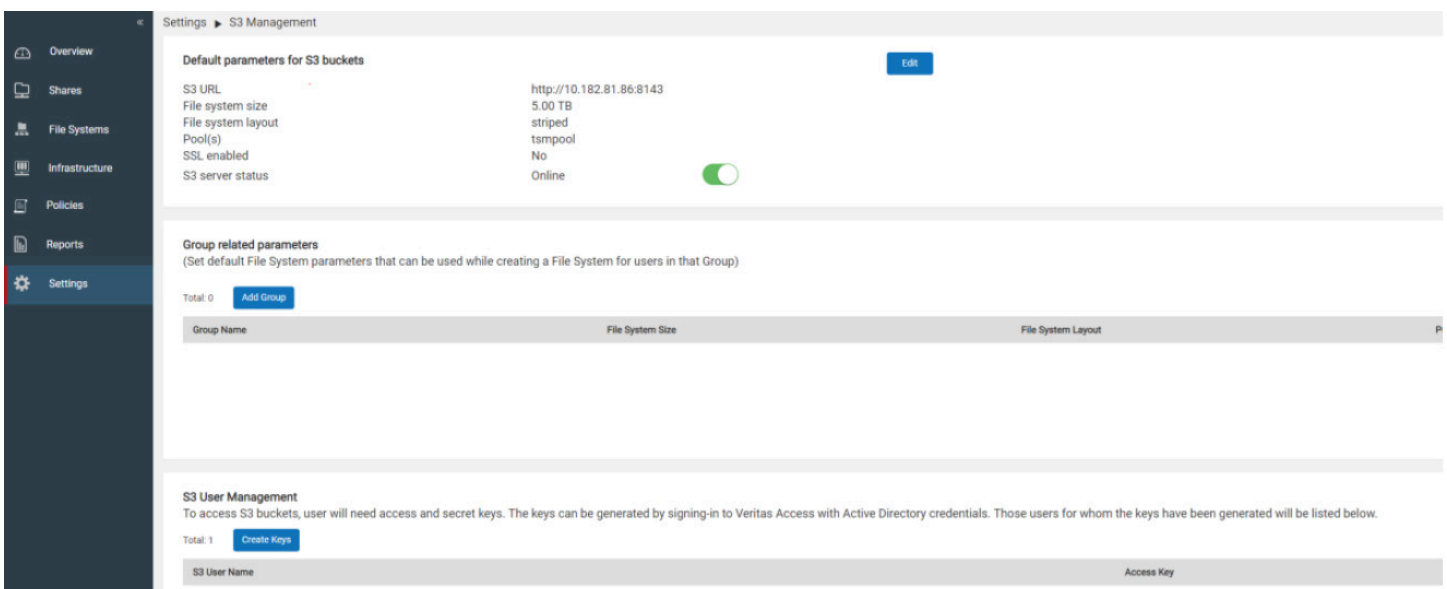
Step 14) Enter the user name and password of user. Click Next. Save the Access Key and Secret Key. Click Finish after saving the keys. These keys would be used to connect to Access S3 Bucket on the application.

Create S3 Keys ? X	Create S3 Keys ? X
S3 User Name admin S3 User Password •••••••• [Next] [Cancel]	✅ Create S3 Keys successful Access Key MzUxMDNlZmYxM2U4OGF Secret Key NjRlNjhhMDJhOWRlNjlwYmE0Y 2E3NDBjYzU3NzE [Finish]

CREATION OF BUCKET ON ACCESS APPLIANCE

At this point, there are two ways to create an S3 bucket. One can use an S3 Browser or similar applications to create bucket by providing the S3 URL(<hostname or IP:8143), access key and secret key or it can be done via the command line on Access. In either methods, a filesystem is first created and then a bucket is created within the filesystem. In this example, creation of bucket is done via the Access CLI using the virtual IP of the console.

Step 1) ssh onto the Access CLI (i.e. ssh admin@<console IP>) and then do the following on the CLI:

a) Go into the storage mode and create a file system by entering "**fs create <layout> <name> <size> <storage pool name created in previous step>**". Then exit out of the storage mode.

b) Enter the objectaccess mode and then map the filesystem created in previous step to a bucket by the following command: "**map / vx/<fs name>/<new bucket name>**"

c) Validate the bucket has been created by entering "**bucket show**" or you can validate using the Access GUI.

```
tmeaccess2> storage
tmeaccess2.Storage> fs create simple s3tsmfs 5t tsmpool
100% [#] Creating simple filesystem
ACCESS fs SUCCESS V-493-10-2095 Created simple file system s3tsm

tmeaccess2.Storage> fs list
FS          STATUS SIZE  LAYOUT  MIRRORS  COLUMNS USE%  USED    NFS SHARED  CIFS SHARED  FTP SHARED  SECONDARY TIER
============== ====== ===== ======= ======= ======= ===== ======= ========== =========== ========== ==============
s3tsm        online 5.00T simple  -       -       0.02% 801.88M no      no          no          no

tmeaccess2.Storage> exit
tmeaccess2> objectaccess
Entering Object Access Service configuration mode...
tmeaccess2.ObjectAccess> map /vx/s3tsmfs/tbucket admin
ACCESS ObjectAccess SUCCESS V-493-10-4 Successfully mapped bucket tbucket to admin.
tmeaccess2.ObjectAccess> bucket show
Bucket Name FileSystem    Pool(s)  Owner
=========== ============== ======= =====
tbucket    s3tsmfs        tsmpool admin
```

CONFIGURATION OF ACCESS WITH IBM SPECTRUM PROTECT

To configure the Access Appliance onto IBM Spectrum Protect as a cloud container storage pool, a directory on the IBM Spectrum Protect server is required to act as a local disk cache for staging of the backups prior to sending the container files to the Access Appliance. In addition, the following information is required:

- Access and secret keys saved previously.
- Access Object S3 URL (in this example it would be: http://10.182.81.86)
- The bucket name created in previous step (i.e. tbucket)

CONFIGURE AS A CLOUD CONTAINER STORAGE POOL

Step 1) Create a directory to act as temporary disk cache for data to back up or archive on the server running IBM Spectrum Protect. NOTE: Make sure the directory created is owned by the IBM Spectrum Protect administrator.

```
[root@scltselnx27 ~]# cd /home/
[root@scltselnx27 home]# ls
astgpooldir cstgpooldir dstgpooldir fpooldir inst install key.txt root self tsmNotes2 tsmuser tushar vncuser1
[root@scltselnx27 home]# su - tsmuser
Last login: Thu Jan 10 03:25:44 PST 2019 on pts/3
[tsmuser@scltselnx27 ~]$ cd /home
[tsmuser@scltselnx27 home]$ ls
astgpooldir cstgpooldir dstgpooldir fpooldir inst install key.txt root self tsmNotes2 tsmuser tushar vncuser1
[tsmuser@scltselnx27 home]$ mkdir tstgpooldir
[tsmuser@scltselnx27 home]$ ls -ld .
drwxrwxrwx. 14 root root 286 Jan 10 16:40 .
[tsmuser@scltselnx27 home]$ ls -ld tstgpooldir
drwxr-xr-x. 2 tsmuser tsmgrp 6 Jan 10 16:40 tstgpooldir
[tsmuser@scltselnx27 home]$
```

Step 2) Using a web browser, log onto the Operations Center (i.e. https://localhost:11090/oc) and click on menu bar indicated by and then click on Storage Pools on left pane.

Step 3) Click on the "**+ Storage Pool**" to add a storage pool.

Step 4) Enter **name of pool**, for example, "tpool" and then click **Next**.

Step 5) Select "**On-premise cloud**" and then **Next**.

Step 6) Enter the **access key, secret key, bucket name, object URL** of the Access Appliance. Click **Next**.

Step 7) Enter the **directory** of the local disk cache created in the previous step. Then, click "**Add Storage Pool**".

Step 8) The next step is to create policies thus, click Close and View Policies. The steps to modify existing domain policies are described in the next section.

MODIFY DOMAIN POLICY

The policies within IBM Spectrum Protect specifies which storage pool is utilized as the backup and archive destination. Clients are bound to one active policy at a time. Multiple client nodes can be associated to the same active domain policy. Thus, backups and archives from multiple clients can be directed to a single storage pool. The steps in this section define how to set up policies using the storage pool created in previous section. If you did not click on "Close and View Policies" in the previous section, then click on the Menu and click on "Services".

Step 1) In this example, the default STANDARD domain policy is modified to use the storage pool just created in the previous step that maps to the Access Appliance. **Double click on the policy domain STANDARD**.

Step 2) Click the "**Configure**" slider and slide the button to right to modify the Archive Destination to the storage pool TPOOL.
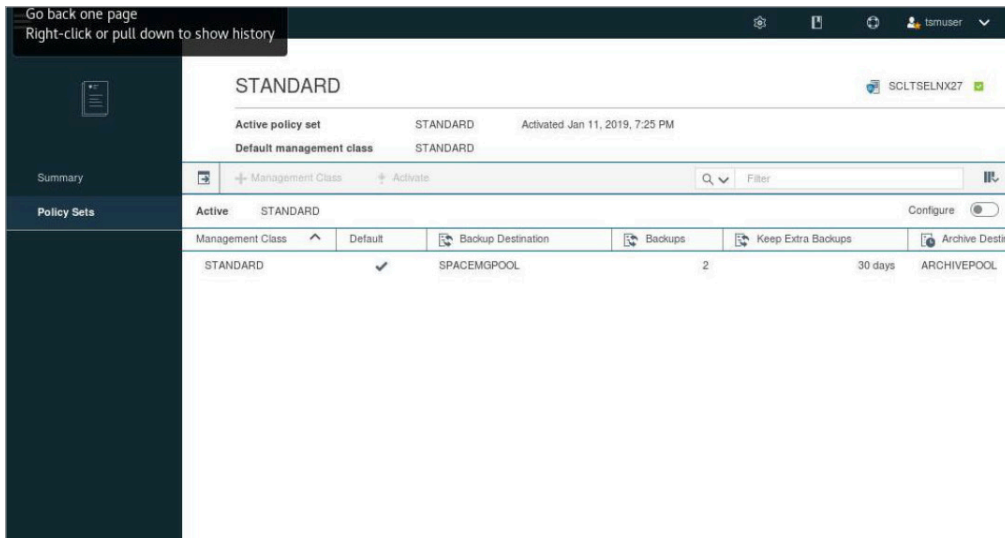
Step 3) Select **TPOOL** as the new archive destination from the drop down.

Step 4) Click **Save** and confirm in the following pop-up window. After it succeeds, close the pop-up window.
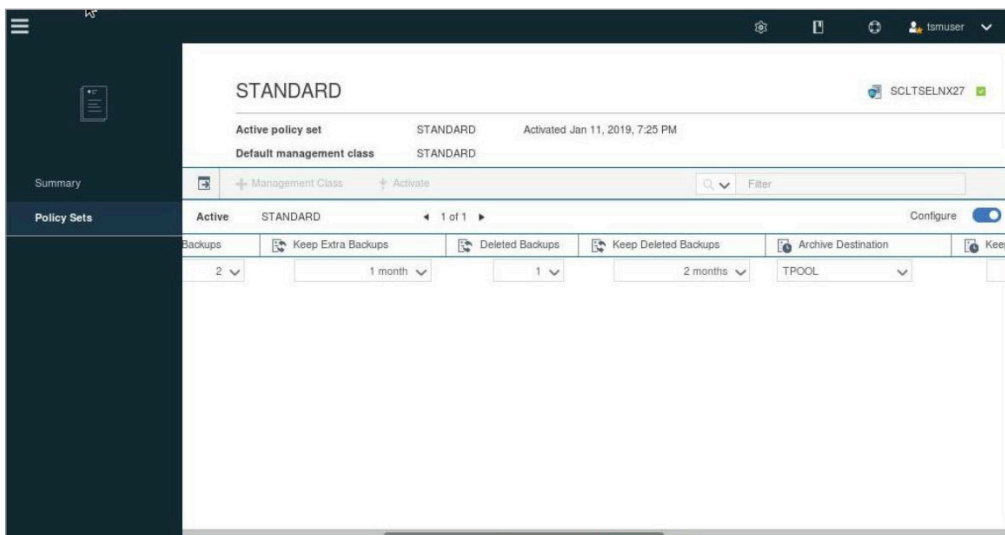
Step 5) The policy would need to be activated to put into the affect the changes. Thus, click on "**Activate**".

Step 6) Acknowledge by **checking the box at the bottom**, that you understand that changing the policy can cause data loss or deletion. After it succeeds, close the pop-up window.
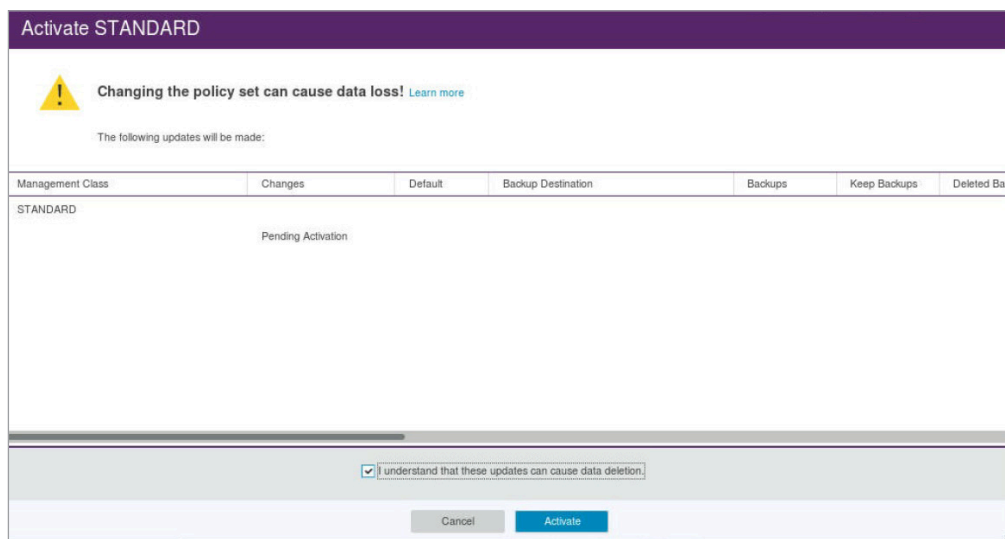
Step 7) The STANDARD policy would now indicate that the Archive Destination is TPOOL.
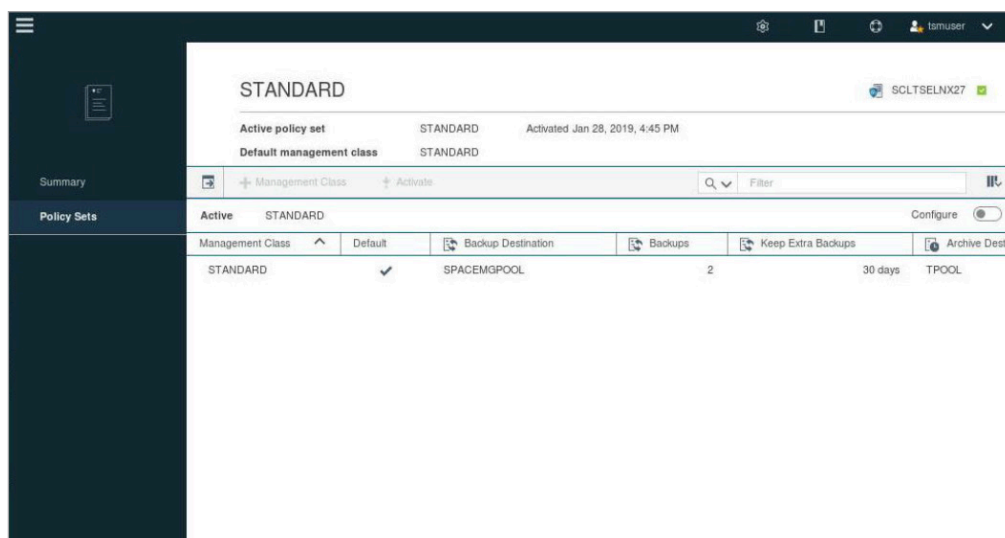
REGISTER A CLIENT NODE AND ASSOCIATE WITH A DOMAIN POLICY

Clients that require backup and archive would need to register with the IBM Spectrum Protect server managing the backup and archival as well as the target storage platform for the backups and/or archives. After client is registered then a domain policy is associated with the client. A client is associated with only one active domain policy. This section describes the steps to register a client and associate that client with a domain policy.

Step 1) Click on Menu indicated by ☰ and then click on Clients.

Step 2) Click on "+ Client" to register a client.
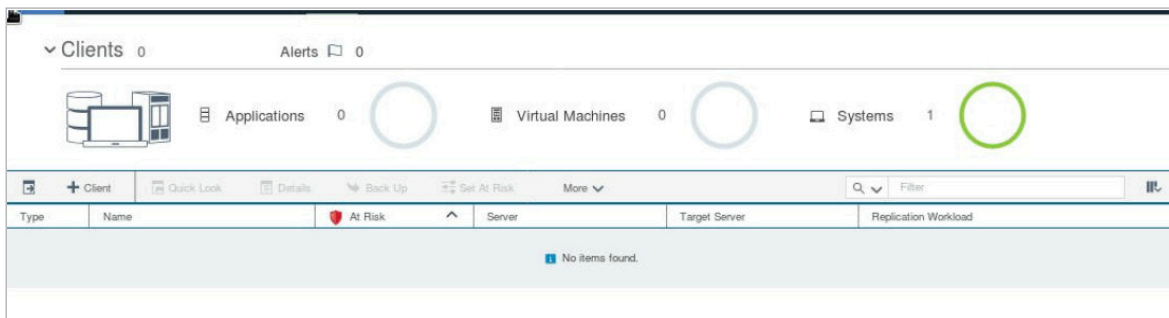
Step 3) Select the name of the IBM Spectrum Protect Server managing this client.

Step 4) Enter in **client name, client password and verify password**.

Step 5) Review configuration information and click "**Next**".

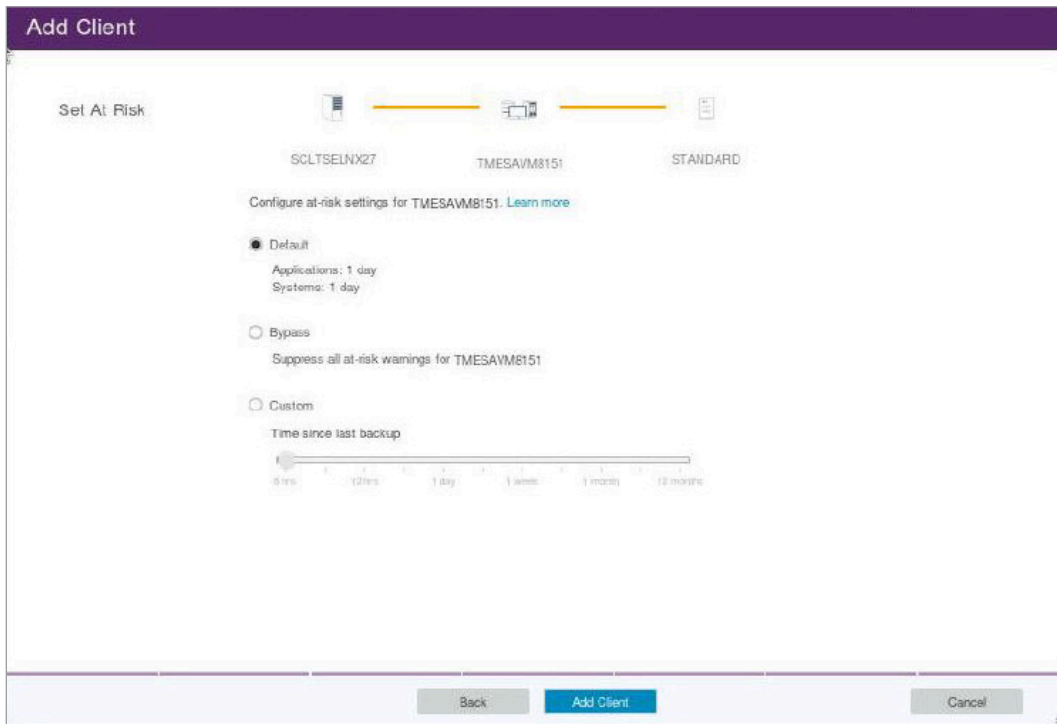Step 6) Associated with an existing domain policy. In this example, it is **STANDARD**.

Step 7) Click **Next**. In this example, no existing schedule is defined since we are doing the backup and restore manually via the Backup-Archive client tool.

Step 8) Just selecting Default for now as the at-risk setting, then click "**Add Client**" and "**Close**" after it succeeds.

Step 9) Review the client page and validation that the client has been registered.
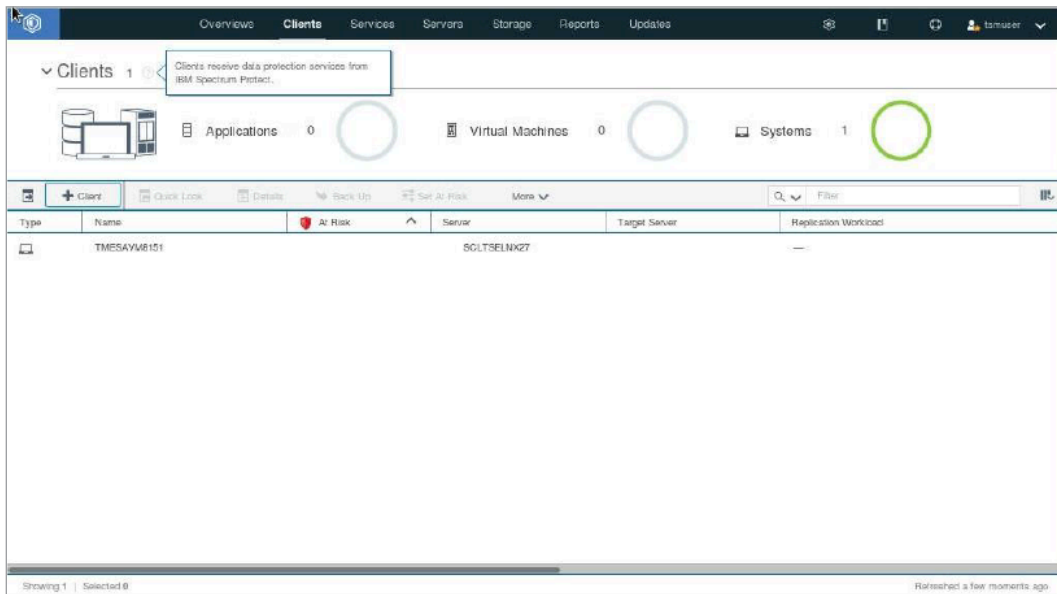
VALIDATION OF THE SETUP

To validate the configuration of the Access Appliance as an S3 target for IBM Spectrum Protect, the backup and archive client is installed on the system registered as the client in the previous section. An archive of several directories and files in the registered client is conducted. Afterwards a restore to temporary location is done to validate restores.

Archive

Step 1) Start up the backup and archive client by entering in **dsmj** on an open terminal. To validate that this client is using the Standard domain policy, click on **Utilities** at the top pane and select "**View Policy Information**".


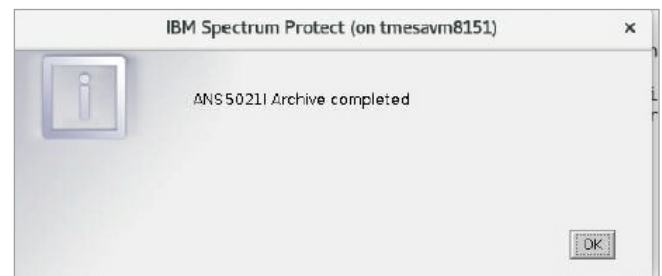Step 2) **Validate** that TPOOL is the **Copy Destination**.

Step 3) Click on **Archive**.

Step 4) Select the files or directories to archive. In this example, the **binaries** directory is archived.

Step 5) Wait until the archive is complete.

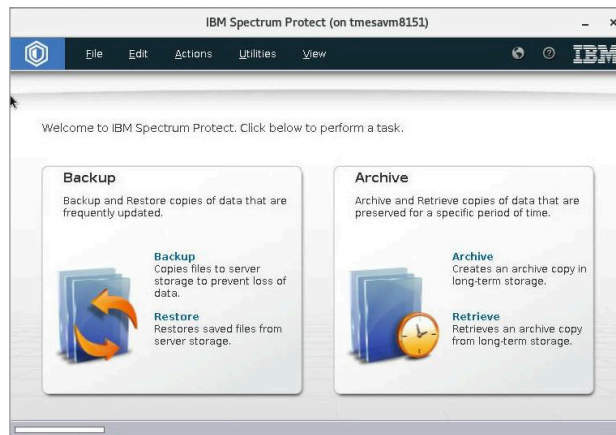Step 6) Using an S3 Browser, validate the contents of the bucket. Contents of the bucket are directories with *.dcf or *.ncf files.
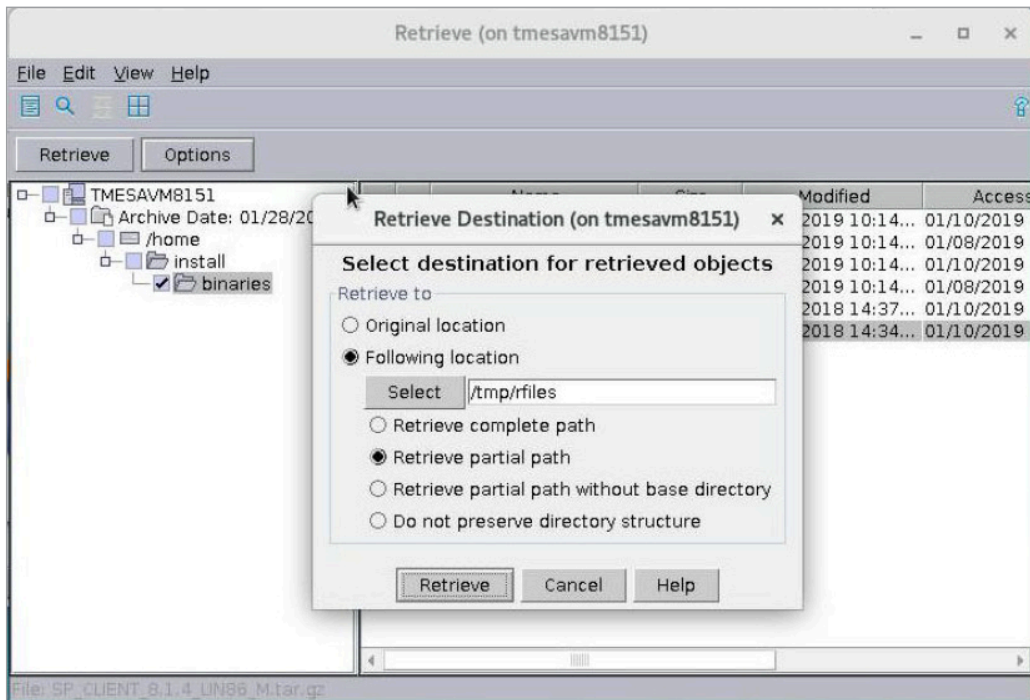

Restore

Step 1) Click on Retrieve in the Archive box.

Step 2) Expand the Archive Date and then select the desired directory archived in previous section to retrieve.

Step 3) Select a destination location to place the retrieved files.

Step 4) Wait for the retrieval to complete.

Step 5) Open a terminal and validate that the directory and its contents have been retrieved.

```
root@tmesavm8151:~                                    _  □  ✕

File   Edit   View   Search   Terminal   Help
[root@tmesavm8151 ~]# ls /tmp/rfiles
binaries
[root@tmesavm8151 ~]# ls /tmp/rfiles/binaries/
citScanOutput.xml   dsminstr.log         sesstest_v12.tar.gz
dsmerror.log        dsminstr.log.lock   SP_CLIENT_8.1.4_LIN86_M.tar.gz
[root@tmesavm8151 ~]# █
```

USING SSL

When using SSL, Access certificate would need to be added the IBM Spectrum Protect keystore. The certificates get generated when SSL is first enabled on Access. **NOTE**: Any changes to the endpoint in the portald.conf would require a re-generation of the certificate ticket using the "certificate renew" from the objectaccess mode on the command-line interface (CLI). SSL can be enabled using GUI or CLI. In this example, SSL is enabled using the CLI.

Step 1) Connect to the Access CLI using ssh (i.e. ssh admin@<console IP>).

Step 2) Enter the objectaccess mode and enable SSL and then stop and start the object services.

```
tmeaccess2> objectaccess
tmeaccess2.ObjectAccess> set ssl_enabled yes
tmeaccess2.ObjectAccess> server stop
ACCESS ObjectAccess SUCCESS V-493-10-4 ObjectAccess stopped successfully.
tmeaccess2.ObjectAccess> server start
ACCESS ObjectAccess SUCCESS V-493-10-4 ObjectAccess started successfully.
```

Step 3) Get the certificate and save the -----BEGIN/END---- clause inside a temporary file (i.e. cert.pem) inside the IBM Spectrum Protect server.

```
tmeaccess2.ObjectAccess> certificate show
Type            Self-signed
CA certificate expiry     Jan 31 21:19:05 2022 GMT
Server certificate expiry  Feb  1 21:46:57 2020 GMT
CA certificate          -----BEGIN CERTIFICATE-----
MIIDhTCCAm2gAwIBAgIJAPJ30XIJwsMqMA0GCSqGSIb3DQEBCwUAMFkxCzAJBgNV
BAYTAlVTMRMwEQYDVQQIDApDYWxpZm9ybmlhMRAwDgYDVQQKDAdWZXJpdGFzMQ8w
DQYDVQQLDAZBY2Nlc3MxEjAQBgNVBAMMCXMzLmFjY2VzczAeFw0xOTAyMDEyMTE5
MDVaFw0yMjAxMzEyMTE5MDVaMFkxCzAJBgNVBAYTAlVTMRMwEQYDVQQIDApDYWxp
Zm9ybmlhMRAwDgYDVQQKDAdWZXJpdGFzMQ8wDQYDVQQLDAZBY2Nlc3MxEjAQBgNV
BAMMCXMzLmFjY2VzczCCASIwDQYJKoZIhvcNAQEBBQADggEPADCCAQoCggEBANTD
6Nf4n9sxIR7ZgXezLX7I8mZS+Js+XiedyLBcrhwdduziLvQ/zrfouwCbKYIu+wW4
/3Xv3D40PBJ6FOLaddS5TDNo5JSaIPKV+6dkAzHfBpio1OlJiwX0jW2+6iW8rj9k
tOhrerWAyLiy3x3GLb9y5E0JH+NP6XbIAQP9V07D4GoZ1pXnhNeWe/31Ng1odDo4
wDqQgJpqN0sbhayzdA9IZByYbMtHPq2XuIUezYHosiBs21RYyasAbpnvhsZclR6B
w32MxPVXIUrsjsLSMTT0I+Yiemmw8smB52C3sv/7TQqUbaPsXiTmxUc2+j7/Bil9
6ixY4bQh9fBfrRLXgd8CAwEAAaNQME4wHQYDVR0OBBYEFJmCJklQXc2KwsAnRNID
8V66EqUFMB8GA1UdIwQYMBaAFJmCJklQXc2KwsAnRNID8V66EqUFMAwGA1UdEwQF
MAMBAf8wDQYJKoZIhvcNAQELBQADggEBABJ7U3HmZP6hQ647ZDYDKw/dOt4mA1m4
pN0iJ8pRW3Do57k1fAzyclG7iavFv/pQMPGPHCWmhGtt0h6jzGpTHnOvgzaE4fYC
ZzViHi4M3RIoQLgFJiXTt71ghj13uuYa+87H5WclbJyHmru8e/KlKVxUovai1Ovm
ZBu378GIZjKp1sj/IzyGkPFLDrD8ftpYovKkhV4nm03nKbKcwtQIHklnww8m0wzG
108jN+KvIF3SS1QM6xs+QH+KCtAKE8M+SUu9H/qlXHHK7OiOFHrSGjBGTD2es4py
A4emuO19jA3joRBOAWz+BHkW06PQXaJpuEW+b+Yww4XcnFRhV6Q4VAE=
-----END CERTIFICATE-----
```

Step 4) On the IBM Spectrum Protect Server, use a terminal window to convert the certificate file *.pem (a Base64 encoded ASCII files) to a *.der (the binary form of the certificate that does not contain "BEGIN CERTIFICATE/END CERTIFICATE" statements.

```
# openssl x509 -outform der -in cert.pem —out /tmp/cert1.der
```

Step 5) Save the old keystores in the <IBM Spectrum Protect directory>/jre/bin (i.e. /opt/tivoli/tsm/jre/bin)

```
# cd /opt/tivoli/tsm/jre/bin
# cp ../lib/security/cacerts ../lib/security/cacerts.original
```

Step 6) Import the certificate into the keystore using the keytool utility in the jre/bin directory.

> ./keytool -import -keystore ../lib/security/cacerts -alias <somealias> -file <yourfile> where <some alias> is a unique alias you chose to identify this certificate in the keystore. This is important if you have more than one certificate and <yourfile> is the path and file name of the certificate created in previous step.

If a password is asked, use "changeit" and enter yes to trust the certificate.

```
# ./keytool -import -keystore ../lib/security/cacerts -alias acert -file /tmp/cert1.der
```

Step 7) Stop the IBM Spectrum Protect Server. To halt IBM Spectrum Protect, ssh onto the client 10.182.81.51 as root/P@ssw0rd and enter into the IBM Spectrum Protect CLI and issue a halt command. This can also be done via the IBM Spectrum Protect GUI using the Command builder.

```
[root@tmesavm8151 ~]# dsmadmc -id=tsmuser -pa=P@ssword

IBM Spectrum Protect

Command Line Administrative Interface - Version 8, Release 1, Level 4.0

(c) Copyright by IBM Corporation and other(s) 1990, 2017. All Rights Reserved.

 Session established with server SCLTSELNX27: Linux/x86_64

 Server Version 8, Release 1, Level 5.000

 Server date/time: 01/09/2019 17:46:45  Last access: 01/09/2019 17:45:46

Protect: SCLTSELNX27> halt
```

Step 8) Verify IBM Spectrum Protect is halted by checking the services (e.g. dsmserv) has stopped by doing a "ps –eaf | more".

Step 9) Restart the server, on the Spectrum Protect server by doing an ssh onto the IBM Spectrum Protect server as the administrator (e.g. tsmuser/P@ssw0rd) and source the db2profile and run dsmserv.

```
# . ~/sqllib/db2profile

# dsmserv
```

NOTE: It takes a while for Spectrum Protect to come up. The way to verify it is fully up is to logon to the client and run the "dsmadmc -id=tsmuser -pa=P@ssw0rd" to get into the server CLI. If you can get in, then it is up.

ABOUT VERITAS

Veritas Technologies is a global leader in data protection and availability. Over 50,000 enterprises—including 87 percent of the Fortune Global 500—rely on us to abstract IT complexity and simplify data management. The Veritas Enterprise Data Services Platform automates the protection and orchestrates the recovery of data everywhere it lives, ensures 24/7 availability of business-critical applications, and provides enterprises with the insights they need to comply with evolving data regulations. With a reputation for reliability at scale and a deployment model to fit any need, Veritas Enterprise Data Services Platform supports more than 800 different data sources, over 100 different operating systems, more than 1,400 storage targets, and more than 60 different cloud platforms. Learn more at www.veritas.com. Follow us on Twitter at @veritastechllc.

2625 Augustine Drive, Santa Clara, CA 95054
+1 (866) 837 4827
www.veritas.com

For specific country offices and contact numbers, please visit our website.
www.veritas.com/company/contact

VERITAS™