

IBM Spectrum Protect

Cloud Blueprint for Google Cloud

Document version 2.0.1

James Damgar

IBM Spectrum Protect Performance Evaluation



© Copyright International Business Machines Corporation 2018, 2023

US Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents.....	3
List of Figures.....	5
List of Tables	6
Introduction.....	6
1.1 Purpose of this Paper.....	6
1.2 Considerations for Disk-to-Cloud Tiering Versus Direct-to-Cloud Data Movement.....	7
1.2.1 Cloud Accelerator Cache Considerations.....	8
1.2.2 Workload Limitations and Considerations with Tiering	9
1.3 Cloud Deployment Patterns.....	12
1.4 Cloud Environment Considerations	13
1.4.1 Importance of Adequate Sizing	13
1.4.2 Linux Logical Volume Manager (LVM)	14
1.5 References to Physical IBM Spectrum Protect Blueprints.....	14
1.6 Database Backup to Object Storage	15
1.6.1 Tuning Database Backup Operations to Object Storage	16
1.7 Server Maintenance Scheduling Considerations	18
1.8 Session Scalability by Blueprint Size	19
1.9 IBM Spectrum Protect Security Concepts and Guidelines.....	20
1.9.1 Identity Security in the IBM Spectrum Protect Environment.....	22
1.9.2 IBM Spectrum Protect Administrative Security.....	23
1.9.3 IBM Spectrum Protect Data Security and Policy Based Management	24
1.9.4 IBM Spectrum Protect Bit-Level Security.....	25
1.10 Optimizing Restore Performance from Cloud-Container Storage Pools	26
1.10.1 Enabling the Cloud-container Storage Pool Cloud Read Cache.....	26
1.10.2 Increasing the Minimum Extent Size for Client Nodes	28
Google Cloud Platform Configurations.....	28
2.1 Compute, Disk, and File System Setup Instructions for Linux-based Systems	34
2.1.1 Linux Operating System Packages and General Guidance	34
2.1.2 Large Blueprint Instructions.....	36
2.1.2.1 Disk Setup Instructions	36
2.1.3 Medium Blueprint Instructions.....	38
2.1.3.1 Disk Setup Instructions	38
2.1.4 Small Blueprint Instructions.....	41
2.1.4.1 Disk Setup Instructions	41
2.1.5 Extra-Small Blueprint Instructions	42
2.1.5.1 Disk Setup Instructions	42
2.1.6 IBM Spectrum Protect Format and Server Options.....	44

2.2	Design Considerations for Google Cloud Platform Compute Engine Instances	46
2.2.1	Considerations for Direct-to-Cloud Architectures	47
2.2.2	Sizing the Cloud Accelerator Cache	48
	Appendix	49
	Disk Benchmarking	49
	Object Storage Benchmarking	52
	Benchmarking Performance Goals	56
	Instance and Object Storage: Navigating the Google Cloud Platform Console User Interface	56
	References	67
	Notices	69
	Trademarks	70

LIST OF FIGURES

Figure 1: Disk-to-cloud tiering, before and after	10
Figure 2: Disk-to-cloud tiering, after tiering.....	11
Figure 3: Deployment patterns	12
Figure 4: Layered security	22
Figure 5: Credential security	23
Figure 6: Container pool security.....	26
Figure 4 : Sizing the cloud accelerator cache for Google Cloud Platform.....	49

LIST OF TABLES

Table 1: IBM Spectrum Protect physical Blueprint targets (Linux x86)	15
Table 2: Preferred ranges of maximum values for client session counts	19
Table 3: Google Cloud Platform, large configuration.....	29
Table 5: Google Cloud Platform, medium configuration.....	31
Table 6: Google Cloud Platform, small configuration	32
Table 7: Google Cloud Platform, extra-small configuration	33

Introduction



1.1 Purpose of this Paper

This document introduces possibilities for integrating an IBM Spectrum® Protect server with **Google Cloud Platform (GCP)** infrastructure as a service (IaaS). You can use the configurations as starting points to deploy a large, medium, or small system (as defined in the [IBM Spectrum Protect Blueprints](#)) or an extra-small system. With the goal of achieving a target daily ingestion rate (corresponding to a large, medium, small, or extra-small deployment), configuration possibilities are offered so that you can get a sense of the relative CPU, memory, disk, and network capabilities that are needed to satisfy requirements. In addition, a survey of options for fulfilling these needs is provided. Certain cloud instance and disk types offered by providers might be sufficient in some areas while lacking in others. You must recognize where system bottlenecks might arise that could limit IBM Spectrum Protect capability.

Use this paper as a **starting point for guidance** about where to deploy an instance of the IBM Spectrum Protect server within a **Google Cloud Platform Compute Engine** virtual machine instance with the goal to eventually store the bulk of primary backup and archive data on cost-effective object storage. In the case of **Google Cloud Platform**, this means **Cloud Storage** object storage. This goal can be accomplished by configuring an IBM Spectrum Protect cloud-container storage pool with (block disk-based) accelerator cache. Two approaches are generally available for storing data in an IBM Spectrum Protect cloud-container storage pool: a direct-to-cloud approach or a disk-to-cloud tiering approach.

For **direct-to-cloud** architectures, backup data is ingested directly into a **cloud-container storage pool** with a performant accelerator cache disk location tuned for a system's ingestion workload as the initial "landing spot" (for more information, see [Sizing the Cloud Accelerator Cache](#)). Data is then immediately, asynchronously transferred to object storage while further data is also ingested into the disk staging area (also known as *overlapped I/O*). The key consideration here is to determine the performance characteristics that the disk staging area must provide to allow for this mixed write-and-read behavior to ensure that ingestion targets are met. A *Cloud Cache and Object Storage Benchmarking* guide and "Cloud benchmarking tools" packages are provided along with this paper to assist in benchmarking both the cloud accelerator cache and object storage system from a prospective host server.

By contrast, **disk-to-cloud** tiering architectures make use of IBM Spectrum Protect **storage rules** to demote data from one or more (usually small) directory-container storage pools (a disk tier) to a cloud-container storage pool. Backup data is initially ingested into a directory-container storage pool and later a portion of this data is moved asynchronously to a cloud-container storage pool. The data can be moved with either age-based tiering (available as of IBM Spectrum Protect Version 8.1.3) or tiering by backup state (available as of IBM Spectrum Protect V8.1.6). With IBM Spectrum Protect V8.1.6, a combination of storage rules and storage subrules can be used to facilitate more granular tiering behavior between the disk and object storage tiers, allowing for flexibility and filtering, for example, by node and node file space. This paper provides some guidance to determine whether tiering is suitable for your workload demands and characteristics. (For an introduction to the tiering features that are available in IBM Spectrum Protect and for detailed guidance, see [Tiering data to cloud or tape storage](#) in the online product documentation.)

With **Google Cloud Platform**, solutions involving disk-to-cloud tiering can be cost prohibitive due to the higher cost of block storage disk, which might be required for the disk tier, in contrast to the relatively lower cost of **Cloud Storage** object storage. However, the requirements for restoring data might be such that having operational recovery data on a fast-performing disk tier is worth this additional cost. Further guidance is provided in the following section regarding considerations for using a direct-to-cloud or tiering approach. The architectures referenced within this paper use a direct-to-cloud approach but can be adjusted to conform to the requirements of a tiering architecture.

1.2 Considerations for Disk-to-Cloud Tiering Versus Direct-to-Cloud Data Movement

The primary advantage of the **tiering** model is that operational recovery data can be preserved on a localized, fast disk tier for rapid recovery while older copies of data or data intended for long-term retention can be demoted to object storage, which is typically more affordable. The tiering model can also be used as an alternative to the direct-to-cloud model with a relatively small disk tier footprint (not strictly for operational recovery purposes). When the `TIERDELAY` parameter is set to 0, age-based tiering can be used to tier each day's worth of ingested client data. In this case, potentially less expensive disk can be provisioned for use by the small directory-container storage pool tier because no ingestion and cloud transfer input/output (I/O) operations occur in parallel. Tiering can be run serially after the completion of data ingestion during scheduled windows with less or no contention for this disk; the disk area can be cleared in preparation for the next day's ingestion.

The same ingestion targets can be satisfied with the disk-to-cloud tiering model as with the direct-to-cloud model, assuming that the direct-to-cloud approach makes use of an accelerator cache and overlapped data ingestion.

Restriction: To implement cloud tiering, you must provision enough disk space to hold a full day's worth of ingested data (plus some buffer) to avoid failed backup operations. The same underlying disk technology can be used in both cases. If, however, you plan to use disk-to-cloud tiering to hold one or more days' worth of operational recovery data within a directory-container storage pool disk tier, the instance disk capacity might have to be much

greater, with the caveat that a slower-performing disk might be sufficient for this case. In all cases, you must understand the ingestion targets (after data deduplication and compression) to determine a daily disk capacity for a transient disk case. Meanwhile, operational recovery requirements in terms of the number of days' worth of recovery data (after data deduplication and compression) should be determined to further size a directory-container storage pool with tiering to cloud if necessary.

With the **direct-to-cloud** model, you can minimize local block storage capacity. This is an advantage because local block storage can be cost prohibitive in cloud-hosted environments.

Note that pursuing a tiering model with Google Cloud Platform Compute Engine virtual instances may be cost prohibitive in terms of the amount of block storage disk required for a directory-container storage pool “disk tier”. As such, only direct-to-cloud model options are presented within this paper.

1.2.1 Cloud Accelerator Cache Considerations

Beginning with IBM Spectrum Protect V8.1.2, data ingestion from clients is throttled if the accelerator cache area is near capacity. This feature makes it possible for this disk cache location to be **underprovisioned** from a capacity standpoint in that the disk cache location does not have to be sized large enough to hold a full day's worth of deduplicated and compressed data. However, the accelerator disk still must be performant enough in terms of input/output operations per second (IOPS) so that client data ingestion and replication target activity can be completed in a timely manner. In the end, you have to compare costs to determine whether larger capacity, less-expensive disk with tiering has an advantage over a direct-to-cloud cache model for a given environment, ingestion rate, and recovery objective.

Restriction: If you plan to use the direct-to-cloud ingestion model, the cloud accelerator cache should be sized large enough to hold at least two times the largest front-end object being ingested. For example, if a 512 GB object is to be ingested directly into a cloud-container storage pool, the cloud accelerator cache should be at least 1 TB in size. Similarly, if 5 client sessions will be backing up 100 GB files each at the same time, the cloud accelerator cache should be sized to at least 1000 GB (5 clients x 100 GB files x 2). This is because the IBM Spectrum Protect server will attempt to “reserve” space in the cloud accelerator cache for in-flight ingestion until ingestion is completed for those objects and their database transactions are committed to the server. By default, this processing assumes no savings from data deduplication or compression and attempts to reserve the total front-end amount of data to ensure sufficient storage capacity.

Beginning with IBM Spectrum Protect V8.1.6, a server option can be used to influence this behavior. The undocumented server option `PreallocReductionRate` can be used to give the server a “hint” about the expected reduction ratio for ingested data and cause the server to reserve less physical space in the directory-container storage pool. For example, setting this option to 5 will cause the server to assume a 5:1 data reduction rate for front-end to back-end data so that only 1 unit of back-end space will be reserved for 5 units of front-end protected data. This option can range from 1 (the default, no reduction) to 25 (a 25:1 assumed reduction). Use this option only when a smaller cloud accelerator cache is

desired and data reduction rates are certain. If the storage pool has inadequate space, backup failures can occur.

Beginning with IBM Spectrum Protect V8.1.10, the undocumented server option `UsePreallocReduction` was introduced to enable automatic preallocation reduction in V8.1.10 and later servers. By default, this option is disabled. If set to YES, the IBM Spectrum Protect server automatically determines an accurate data reduction rate based on existing storage pool statistics and applies this reduction when preallocating space within storage pool directories (or the cloud accelerator cache in the case of a cloud-container storage pool). Use this option only when a smaller cloud accelerator cache is configured and data reduction rates are consistent between backup operations. If the storage pool has inadequate space, backup operations might fail. Also beginning with IBM Spectrum Protect V8.1.10, the undocumented server option `PreallocReductionPadding` can be used to add “padding” to the preallocated space. Specify this option as a percentage. The default value is 5. The following example illustrates this feature. For example, assume that the `UsePreallocReduction` option is enabled, the `PreallocReductionPadding` value is set to 5, and storage pool statistics show a 4:1 data reduction with deduplication and compression (75%). When a 100 MB object is ingested, the IBM Spectrum Protect server will attempt to allocate space in the cloud accelerator cache based on the following formula:

$$((100\% - 75\%) + 5\%) \times 100 \text{ MB} = 30 \text{ MB}$$

Use the `PreallocReductionPadding` option in combination with the `UsePreallocReduction` option if data reduction rates are inconsistent between backups and an accurate padding value can be determined.

The `PreallocReductionRate` option takes precedence over the `UsePreallocReduction` option.

1.2.2 Workload Limitations and Considerations with Tiering

Not all client workloads are suitable for a disk-to-cloud tiering model. Tiering by age (as of IBM Spectrum Protect V8.1.3) allows for the demotion of *backup* objects that have reached the specified age threshold. Inactive backup generations that are older than the specified age are transitioned to object storage. Tiering by state (as of IBM Spectrum Protect V8.1.6) allows for the demotion of backup objects that have reached the specified age threshold and are *inactive* within the server. Active backup objects are preserved on the disk tier, while inactive copies of backup objects are transitioned to object storage.

Disk-to-cloud tiering is **suitable** for client workloads that have **low** data deduplication rates (backup generations differ greatly). In this case, data is highly unique between backup operations. When a backup generation is tiered to object storage, the deduplicated extents (chunks) that make up that object have their references decremented on the source directory-container storage pool. In this case, reference counts are likely to be low and more deduplicated extents are likely to be removed from the disk tier as objects are tiered and space is released.

Disk-to-cloud tiering **might not be suitable** for client workloads that have a **high** data deduplication rate. In this case, data is not very unique between backup generations and many shared deduplicated extents are referenced by multiple object generations. Even

though an object can be tiered by a storage tiering rule, because the object shares many extents with other objects (which might still be active), a large proportion of the object's data will not be removed from the disk tier (although it will be copied to the object storage tier).

The following figures illustrate how data movement with disk-to-cloud tiering can occur. [Figure 1](#) depicts a scenario in which multiple versions of three backup objects (A, B, and C) have been ingested and are stored in a directory-container storage pool on disk. Dotted lines represent references to deduplicated extents (colored, numbered boxes). With the tier-by-state option, the inactive object copies (shown in the gray rectangle) would be tiered to a cloud-container storage pool.

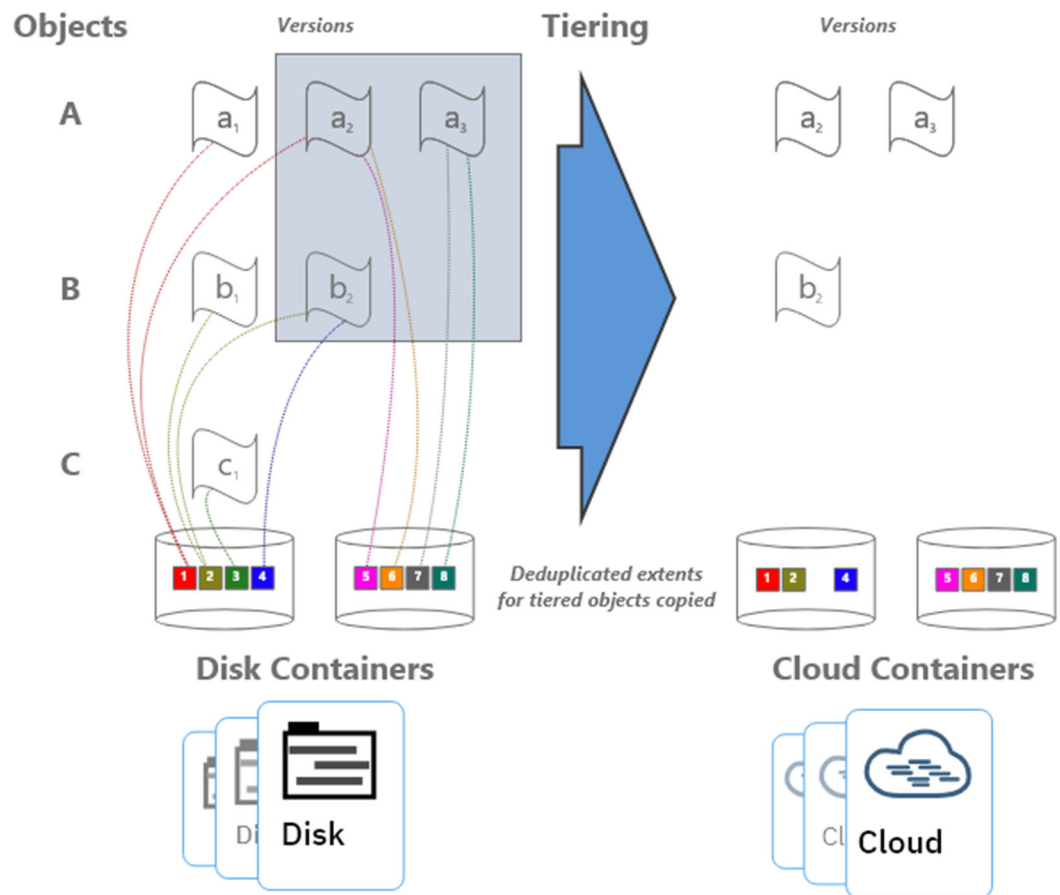


Figure 1: Disk-to-cloud tiering, before and after

[Figure 2](#): Disk-to-cloud tiering, after tiering depicts the situation after tiering is completed and the `REUSEDDELAY` parameter value of the source directory-container storage pool is exceeded (so that deduplicated extent removal for extents with zero reference count can occur).

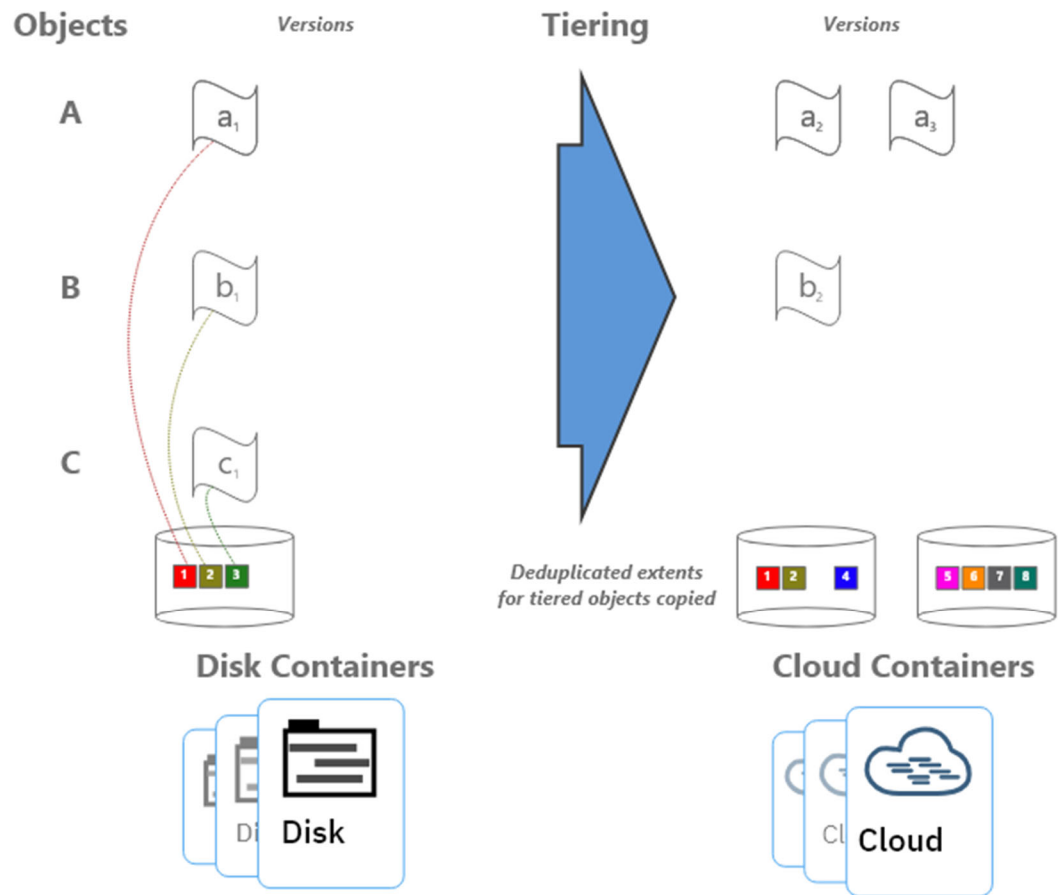


Figure 2: Disk-to-cloud tiering, after tiering

Notice that deduplicated extents 1 and 2 remain on disk even after tiering and extent cleanup have occurred. This is due to the fact that those extents are shared between the active and inactive backup copies. If many deduplicated extents are shared by objects (a high duplicate data rate with high data deduplication ratios), it is more likely that data will remain on disk, even after backup objects have been tiered at an IBM Spectrum Protect inventory level. Keep this factor in mind when you consider a disk-to-cloud tiering model and when you size an environment.

For workloads that deduplicate well from day to day, there will be many shared extents across backup and archive generations and a smaller capacity footprint on tiered object storage as a result because these backup and archive generations will also share many extents in the cloud-container storage pool. For workloads that deduplicate poorly day to day (highly unique data change each day), there will be few shared extents across backup and archive generations and potentially a larger capacity footprint on tiered object storage because these backup and archive generations will each point to (more) unique data in the cloud-container storage pool.

If the primary motivation for using disk-to-cloud tiering is rapid recovery of operational data, a tiering model might provide the best approach. You must understand the nature of the client workload to accurately size the directory-container storage pool on disk.

1.3 Cloud Deployment Patterns

The described configurations can be used as starting points in situations where the IBM Spectrum Protect cloud instance will be a **primary server and in situations where it is used as a replication target**. In scenarios where the cloud-based instance is a replication target, adequate “public” network capability might be necessary to satisfy replication throughput requirements. **Google Cloud Platform Interconnect** can be used to establish a dedicated link of either 10 or 100 Gbps from an on-premises data center to Google Cloud private and public resources to facilitate efficient IBM Spectrum Protect replication or backup processing from peer servers or clients outside of the Google Cloud infrastructure.

Generally, IBM Spectrum Protect deployments making use of cloud-based object storage will align with one of the following three patterns:

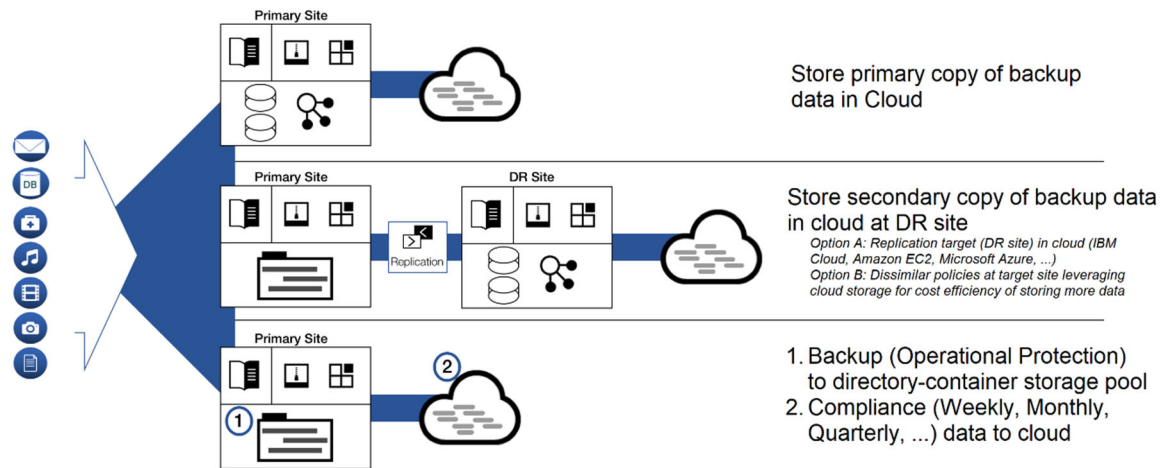


Figure 3: Deployment patterns

In the figure, the first deployment pattern could involve an IBM Spectrum Protect server that is installed on premises or on a Google Cloud Platform Compute Engine virtual machine system, with primary backup and archive data landing in object storage immediately. The positioning of the IBM Spectrum Protect server in relationship to clients could be one critical decision point when you consider whether to have a server instance on premises or within Google Cloud Platform. This pattern could involve use of a direct-to-cloud architecture with accelerator cache or a small directory-container storage pool with immediate tiering to a second cloud-container storage pool without accelerator cache.

The second deployment pattern would make use of cloud-based Google Cloud object storage at the secondary disaster recovery (DR) site. This DR server could be installed at an on-premises site or on a Google Cloud Platform Compute Engine system. In the latter case, sufficient wide area network (WAN) bandwidth between the primary and secondary sites is required for acceptable performance. Much like the first deployment pattern, here the IBM Spectrum Protect server at the DR site could make use of a direct-to-cloud topology with a cloud-container storage pool featuring accelerator cache, or it could use a small directory-container storage pool landing spot with immediate tiering to a cloud-container storage pool backed by object storage.

The third deployment pattern features specific use of disk-to-cloud tiering, available with IBM Spectrum Protect V8.1.3 and later, to allow for operational recovery data to reside on faster performing disk storage. Data that is older, archived, or both would be tiered to Google Cloud object storage after a specified number of days. This deployment could also be performed at an on-premises site or within a Google Cloud Platform Compute Engine instance. However, the additional cost of having a larger capacity directory-container storage pool should be factored into cost estimates with an in-the-cloud solution.

A **combination** of approaches is also possible within the same deployment. For example, a cloud-container storage pool could be configured with accelerator cache disk and made to store long-term retention or compliance archives. A directory-container storage pool could be configured as a disk tier for normal backups, and a tiering relationship could be set up so that operational recovery data (for example, backups from the previous 7 days) is kept on this disk tier, while older data is demoted to the same cloud-container storage pool. The same cloud-container storage pool can be a direct backup target and a tiering target. However, if the pool is a direct target of a backup-archive client, the pool must be configured with accelerator cache disk.

1.4 Cloud Environment Considerations

For Google Cloud Platform Compute Engine, you can choose one of multiple potentially suitable instance types to satisfy an IBM Spectrum Protect server requirement. You do not have complete flexibility in terms of the type and proportion of CPU and memory resources available for an instance. However, standard instance types are available, which can roughly match recommended server requirements. Take care to select instances that have the required **CPU and memory** resources to support inline server data deduplication, compression, encryption, and cloud API operations.

With Google Cloud Platform Compute Engine instances, the possibility exists for you to alter the instance type later, when resource requirements might be clearer. The process is fairly straightforward via the Google Cloud Platform console web portal for most instance types. However, there are some caveats to keep in mind. For more information, see Google's documentation in [References](#) [3].

1.4.1 Importance of Adequate Sizing

Ingested backup data reaching the cloud accelerator cache or the initial directory-container storage pool tier requires the use of block storage allocated to the Google Cloud Platform cloud server. IBM Spectrum Protect database activity also uses some level of throughput and elevated I/O operations during workload processing. Therefore, disk **I/O capability** and instance-to-disk throughput considerations must be evaluated when choosing and configuring an in-the-cloud IBM Spectrum Protect instance with Google Cloud. With a final ingestion point on Google Cloud object storage via a cloud-container storage pool, the **Ethernet capability** of the instance and the nature of the network between the instance and the Google Cloud object storage endpoint (in addition to an understanding of how a front-end client data workload deduplicates and compresses into a back-end stored quantity) should be kept in mind.

During the planning phase, consider how the ingested data will be **reduced via data deduplication and compression** in the back-end storage location. These factors will help you estimate how much back-end data must be moved within a certain time window (measured in hours) and can help predict the throughput (megabytes per second or terabytes per hour) that the Ethernet network and object storage endpoint require to satisfy ingestion requirements. Generally, **10 Gbps** Ethernet capability to private Google Cloud Platform object storage endpoints is required for large, medium, or small Blueprint ingestion targets, while 1 Gbps is sufficient for extra-small targets.

Beginning with IBM Spectrum Protect V8.1.3, the server automatically throttles client backup operations if the cloud accelerator cache portion of a cloud-container storage pool is nearing full capacity. As a result, it is not mandatory to configure cloud accelerator disk cache space that would be large enough to hold a full day's worth of backups (after data deduplication and compression). However, disk benchmarks should be run to ensure that the anticipated back-end workload that an IBM Spectrum Protect server is expected to support will not result in this disk location being the primary bottleneck of the system (see [Disk Benchmarking](#)). In practice, any planned deployment should be validated to ensure that it will meet performance requirements.

1.4.2 Linux Logical Volume Manager (LVM)

The described reference architectures use the Red Hat Enterprise Linux (RHEL) operating system. For a lower-cost alternative, a supported version of the Ubuntu Linux operating system may be used. To deploy IBM Spectrum Protect on a Linux based operating system, the preferred method is to use the **Linux Logical Volume Manager (LVM)** for the cloud accelerator cache disk, Db2 database, and Db2 active log disks. The remaining IBM Spectrum Protect disk components can be satisfied with file systems formatted on directly mounted Google Cloud Platform block disks. The overlapped I/O pattern experienced with concurrent backup ingestion activity and transfer of data to object storage can lead to “hot spots” on disk when more than one storage pool directory is defined for a cloud-container storage pool as accelerator cache. To help avoid a throughput bottleneck, you can configure a single logical volume to span all physical volumes assigned for use as cloud accelerator cache. Furthermore, the preferred method is to use a stripe size of **256 KiBytes** for the single logical volume and ensure that the number of stripes matches the number of physical disks. For guidance about specific commands to use when setting up Linux based instances, see the **Compute, Disk, and File System Setup Instructions for Linux-based Systems** section of this document. The use of Linux LVM to logically stripe across several Google Cloud Platform block disks should not be depended upon to increase the durability of the underlying storage, as Google Cloud Platform block disks are already redundant within one or more datacenters and so do not benefit from the recovery characteristics of LVM striping or RAID.

1.5 References to Physical IBM Spectrum Protect Blueprints

Throughout this paper, the server specifications in the *IBM Spectrum Protect Blueprint and Server Automated Configuration for Linux x86* document (also known as an IBM Spectrum Protect Blueprint) are referenced as targets for CPU and memory configurations matching small, medium, and large server builds. For more information about the Blueprints, see [References](#) [1]. The intention with the server builds outlined here is to provide systems

capable enough from a CPU, memory, disk, and Ethernet point of view to approach Blueprint-level ingest capability. Although different instance types can be used to satisfy the same requirements, the disk specifications in particular should be noted in this document as guidance for those deploying environments of their own.

As a reference, the following table indicates the throughput, capacity, CPU, and memory targets for each of the referenced Blueprints. The values for total managed data and daily ingested data are for the block storage Blueprints. These ingestion targets assume an 8-hour backup window.

Table 1: IBM Spectrum Protect physical Blueprint targets (Linux x86)

Sizing category	CPU	RAM memory	Total managed data (front end)	Daily ingested data (front end)
Small	16 cores	64 GB	60 TB – 240 TB	Up to 10 TB per day
Medium	20 cores	128 GB	360 TB – 1440 TB	10 – 30 TB per day
Large	44 cores	384 GB	1000 TB – 4000 TB	20 – 100 TB per day

Although not defined explicitly in the physical Blueprints, the extra-small cloud Blueprint systems target up to 10 TB or more of total managed (front-end) data with a daily ingestion rate of up to 1 TB, or more, per day.

1.6 Database Backup to Object Storage

Beginning with IBM Spectrum Protect V8.1.10, you can back up the server database to object storage for disaster recovery purposes. The configurations in this document are based on the assumption that the server database is backed up to object storage. An alternative approach for cloud-hosted servers is to use the FILE device class and run backup operations to provisioned disk storage. Previous versions of the Cloud Blueprints included FILE device class configurations that offered approximately two days' worth of full database backups in the worst case (with the Db2 database consuming close to its maximum resident capacity). However, this approach can be cost-prohibitive in some cloud-hosted environments such as Google Cloud Platform, where disk storage is typically more expensive than object storage.

Another advantage of using object storage for IBM Spectrum Protect database backups is that object storage pricing with Google Cloud is based on the amount of used storage, while disk storage pricing is based on the amount of storage space provisioned, even if a portion is unused. Not only is unused provisioned disk space a deterrent to cost savings, the actual rate charged for this space is much more than object storage considering that the data involved (database backups) is archive-like in nature. Static provisioning of disk storage is no longer required and the amount of storage consumed for database backup can better match the requirements of the environment. By taking advantage of this pricing model, you can enjoy greater freedom in choosing and changing retention policies for

database backups to achieve the required recovery window. For example, you can transition from 2 days' worth of full database backups to 7 days without having to re-provision and configure disk storage.

A further benefit of database backup operations to object storage is that increased data redundancy, availability, and durability can be achieved by using an object storage bucket with dual-region or multi-region location settings. "Region" is the most cost-efficient option since it offers the lowest latency, where data is present within a single Google Cloud Region. "Dual Region" copies data across two different Google Cloud Platform regions asynchronously, providing higher availability in the event of a regional outage. Finally, "Multi-Region" offers the highest availability with data copied across multiple regions in a geographic area (for example, "us" for the United States).

You can use the same Google Cloud Platform object storage bucket for database backups and the cloud-container storage pool of the IBM Spectrum Protect server to ensure matching redundancy, availability, and durability attributes for database metadata and storage pool data. In the case of an outage of a zone within a Google Cloud Platform region, an IBM Spectrum Protect server instance can be recovered via a database restore operation and by using the cloud-container storage pool resident data that is accessed by a different Google Cloud Platform Compute Engine virtual machine server instance located within the same region. For more information about Google Cloud Platform redundancy options, see [References](#) [7]. For detailed guidance about setting up database backup operations to object storage, see [References](#) [4].

1.6.1 Tuning Database Backup Operations to Object Storage

When a CLOUD device class is used for IBM Spectrum Protect database backup operations, the following files are copied to object storage:

- Database volumes
- The device configuration file
- The volume history file
- The master encryption key for the server

Large items, such as database volumes, are copied to object storage by using multipart upload. By specifying multiple, concurrent data streams during database backup, you can reduce the time that is required to back up the database. The number of data streams that are used for the database backup operation is the same as the number of data streams that are required for any subsequent database restore operation. The number of data streams affects throughput. Each database backup operation uses the following separate resources:

- A session connection from Db2 to the IBM Spectrum Protect server
- A server thread that sends data from the server to object storage

Several performance factors affect operations for backing up the server database to object storage, for example:

- Database disk performance (256 - 512 KB sequential input/output operations)
-

- Object storage system performance
- Network performance to the object storage system

When using database backup to Google Cloud Platform object storage with the systems presented here, consider the maximum network throughput that is required to complete database backup operations on schedule to meet service level commitments. A common expectation is that a daily full backup of the IBM Spectrum Protect server database can be completed in 2 - 3 hours (or less). A 1 Gbit Ethernet link provides approximately 100 MB/s of throughput while a 10 Gbit Ethernet link provides roughly 1000 MB/s of throughput. A full backup of an 8 TB database (at the larger end of a large Blueprint system) would take more than 20 hours on a 1 Gbit connection and approximately 2 - 3 hours on a 10 Gbit connection. These estimates assume that these network links are hardly utilized otherwise. The relative load on these networks should be considered when scheduling database backup operations and when selecting which network links to provision and configure for cloud compute instances. The health of the network link should also be evaluated. TCP/IP packet loss of as little as 2% can cause a large degradation in throughput for database backup and restore operations from object storage, jeopardizing the daily database backup.

Db2 database encryption is used by default for database backup operations to cloud device classes to provide additional data security for database data. You can specify encryption or compression for a database backup operation to cloud, but not both. If you specify compression for a database backup operation to cloud, encryption is disabled. Compression impacts backup performance and limits front-end throughput to approximately 0.5 TB per hour, or less, and so is not typically suitable for larger server environments unless a longer database backup window can be tolerated. Compression can, however, result in a smaller data footprint in Blob object storage and slightly improve database restore performance. For smaller IBM Spectrum Protect servers with smaller databases (such as the extra-small and small configurations shown here) use compression when the following conditions are met:

- The network link to object storage is 1 Gbit or less.
- Database encryption is not necessary.
- Compression savings are required.

Depending on the IBM Spectrum Protect server size, use the following stream quantities as starting points for optimal performance for database backup operations for extra-small, small, medium, and large Blueprint systems:

- Extra-small system: 1-5 streams
- Small system: 10 streams
- Medium system: 25 streams
- Large system: 50 streams

Then, adjust the number of data streams until you achieve optimal throughput on a consistent basis over time. Each data stream uses approximately 20 MB of memory on the

IBM Spectrum Protect server. For example, a 50-stream database backup operation consumes approximately 1000 MB of memory on the server.

A beneficial step when building a cloud-based IBM Spectrum Protect server is to benchmark the components involved in the solution to prove that the resources available to the server are sufficient to meet performance demands. For database backup operations to Google Cloud Platform object storage, this means benchmarking the sequential read throughput of the IBM Spectrum Protect database disks and benchmarking the throughput capability of the link to object storage from the server instance. By ensuring that these components perform adequately, you can achieve a higher level of confidence that database backup (and restore) operations will perform as expected within the allotted time. For instructions about how to benchmark performance and interpret the results, see [References](#) [5].

1.7 Server Maintenance Scheduling Considerations

The *IBM Spectrum Protect Blueprint and Server Automated Configuration for Linux x86* document provides a detailed breakdown of the procedure for setting up IBM Spectrum Protect server maintenance schedules (see [References](#) [1], Chapter 5). Use this information as a reference for establishing a maintenance schedule on cloud-hosted servers.

For an IBM Spectrum Protect server in Google Cloud Platform that is serving as a replication target, a replication window and schedule might have to be established. For servers using the direct-to-cloud model, where primary backup data is ingested directly into a cloud-container storage pool, a replication window might not be required if this server is not a replication target server because a cloud-container storage pool cannot be used as a replication source. In this case, redundancy requirements for the ingested client data can be met by the inherit redundancy of Google Cloud Platform object storage.

For an IBM Spectrum Protect server running in Google Cloud Platform that is using the disk-to-cloud tiering model, a replication source strategy might be required. Replication can help to protect client data objects in the disk directory-container storage pool that have not yet been tiered (demoted) to object storage because only one copy of that data is present. To prevent excess data from being stored (pinned) to the disk tier, verify the following items:

- The source replication server (used for disk-to-cloud tiering) should be configured with a longer retention policy than the target replication server. In other words, data should be retained for a longer period on the source replication server.
 - The retention policy that affects client node data on the target replication server should match the value of the `TIERDELAY` parameter of the storage rule that is responsible for tiering the same client node data on the source server.
-

In general, the server that is used for disk-to-cloud tiering (whether it be the source replication server or the target replication server) should be the server with the longer retention policy for the client nodes that are affected by the tiering storage rule.

1.8 Session Scalability by Blueprint Size

The **IBM Spectrum Protect Blueprint and Server Automated Configuration for Linux x86** document describes how to set the IBM Spectrum Protect server option `MAXSESSIONS`, based on Blueprint system size:

- Small system: 250 maximum simultaneous client sessions
- Medium system: 500 maximum simultaneous client sessions
- Large system: 1000 maximum simultaneous client sessions

(For more information about the Blueprint configurations, see [References](#) [1].)

The actual throughput scalability of a cloud-based solution depends on many factors, including the configured disk capability and capacity of the system, the amount of CPU and memory resources available on the system, and the relative rate of data deduplication and compression for the dataset that is ingested into the server. Larger objects, which feature a larger deduplicated extent size (for example, 250 - 350 KiBytes, or more) and which do not deduplicate or compress well (for example, less than 10%), will result in less database and computation (CPU) overhead, but will utilize more disk and network bandwidth. The logical reduction of front-end client data to the physical back-end data (which is actually written out and stored to disk and object storage) means that the disk, network, and object storage components will be stressed to a higher degree as client/server session counts increase. Memory usage by the IBM Spectrum Protect server might also be greater. As session counts increase, these components are likely to become a system bottleneck, limiting front-end throughput.

Objects that feature smaller, deduplicated extent sizes (for example, 60 - 100 KiBytes or similar) and that deduplicate and compress well (for example, 50% data deduplication with 50% compressibility) will result in less network, disk, and object storage bandwidth used, but will lead to more database and computation overhead to facilitate these data reduction operations. As session counts increase, CPU and database-related memory are likely to first become limiting factors for these data types. In general, the more successfully data can be deduplicated and compressed (and therefore the greater the data reduction from front-end to back-end data), the greater the number of feasible client sessions. The following table indicates a reasonable range of client session counts based on system size and data type, as well as the likely limiting factor for the system as the high end of the range is approached.

Table 2: Preferred ranges of maximum values for client session counts

Cloud system size	Large object, poor data deduplication and	Large object, good data deduplication	Large object, small extent size, good data deduplication	Small object, poor data deduplication and

	compression¹	and compression²	and compression³	compression⁴
Extra small	10 – 50	25 – 50	10 – 50	10 – 50
Small	50 - 100	100 - 200	50 - 100	50 - 100
Medium	100 - 200	200 - 400	100 - 150	100 - 150
Large	300 - 400	400 - 500	150 - 200	150 - 200
Limiting factor at scale	Network, disk, object storage bandwidth, memory	CPU, memory or network, disk, object storage bandwidth	CPU, memory	CPU, memory

¹ This model uses 128 MiByte objects, 250 - 350 KiByte extents, and <10% data deduplication and compressibility. Full backup operations are used with pseudo random data or data that cannot be easily deduplicated or compressed. For example, this model can be applied to encrypted data.

² This model uses 128 MiByte objects, 150 - 200 KiByte extents, and 50% data deduplication and compressibility. For example, this model can be applied to virtual machine backups.

³ This model uses 1 GiByte objects, 60 - 100 KiByte extents, and 50% data deduplication and compressibility. For example, this model can be applied to database image backups.

⁴ This model uses 128 KiByte objects and <10% data deduplication and compressibility. For example, this model can be applied to file server data and other small files or objects.

Often, a diminishing rate of return in regard to throughput is experienced when 50 - 100 total client sessions are exceeded, regardless of data type. More sessions are possible and might be warranted, given the client schedule or other requirements. However, aggregate gains in total throughput of a single IBM Spectrum Protect instance might not be substantial past this point.

1.9 IBM Spectrum Protect Security Concepts and Guidelines

We live in an increasingly hostile environment in terms of data protection and security. Internal and external malicious actors target organizations and institutions for political or social reasons. Many seek financial gain through the use of *ransomware*, a type of malware that encrypts client production data or backup data for the purpose of obtaining a payment (ransom) to decrypt that data. Often, ransomware attacks include the following operations:

1. Infiltrate the environment
2. Obfuscate data by using encryption
3. Export the data as an additional blackmail tactic

Beyond ransomware, other forms of malware and computer viruses can disrupt production workloads and lead to business downtime. Organizations must develop and understand their security posture to combat these and other evolving threats. IBM Spectrum Protect supports a layered approach to addressing security concerns within a data protection environment. Primarily, this approach is based on secure communication protocols, secure user passwords, and different access levels for IBM Spectrum Protect administrators. Through the use of policy-based management, IBM Spectrum Protect administrators can ensure that a sufficient number data versions are available to support "freedom of action" for recovering data at different points in time within the environment. IBM Spectrum Protect also provides for bit-level security, anomaly detection, and audit mechanisms to help secure an environment.

For an overview and primer about securing an IBM Spectrum Protect environment, see [References](#) [15] and [16].

Overall, a layered approach to data security and protection should be pursued with IBM Spectrum Protect in both on-premises and in-the-cloud deployments. The layered approach is based on the following guidelines:

- The set of people with physical and virtual access to administrative functions with the data protection environment should be limited and evaluated regularly.
- The set of administrators should be limited and access restricted based on business function and a need-to-know basis.
- Data deletion authority should be limited and deletion-prevention controls should be implemented, including the IBM Spectrum Protect command approval feature.
- Redundant copies of data should be made, if possible.
- Protected data should be verified for integrity.
- Recovery procedures should be validated through regular audits and by testing restore operations and disaster recovery (DR) procedures. The goal is to ensure that recovery time objective (RTO) and recovery point objective (RPO) targets can be achieved.

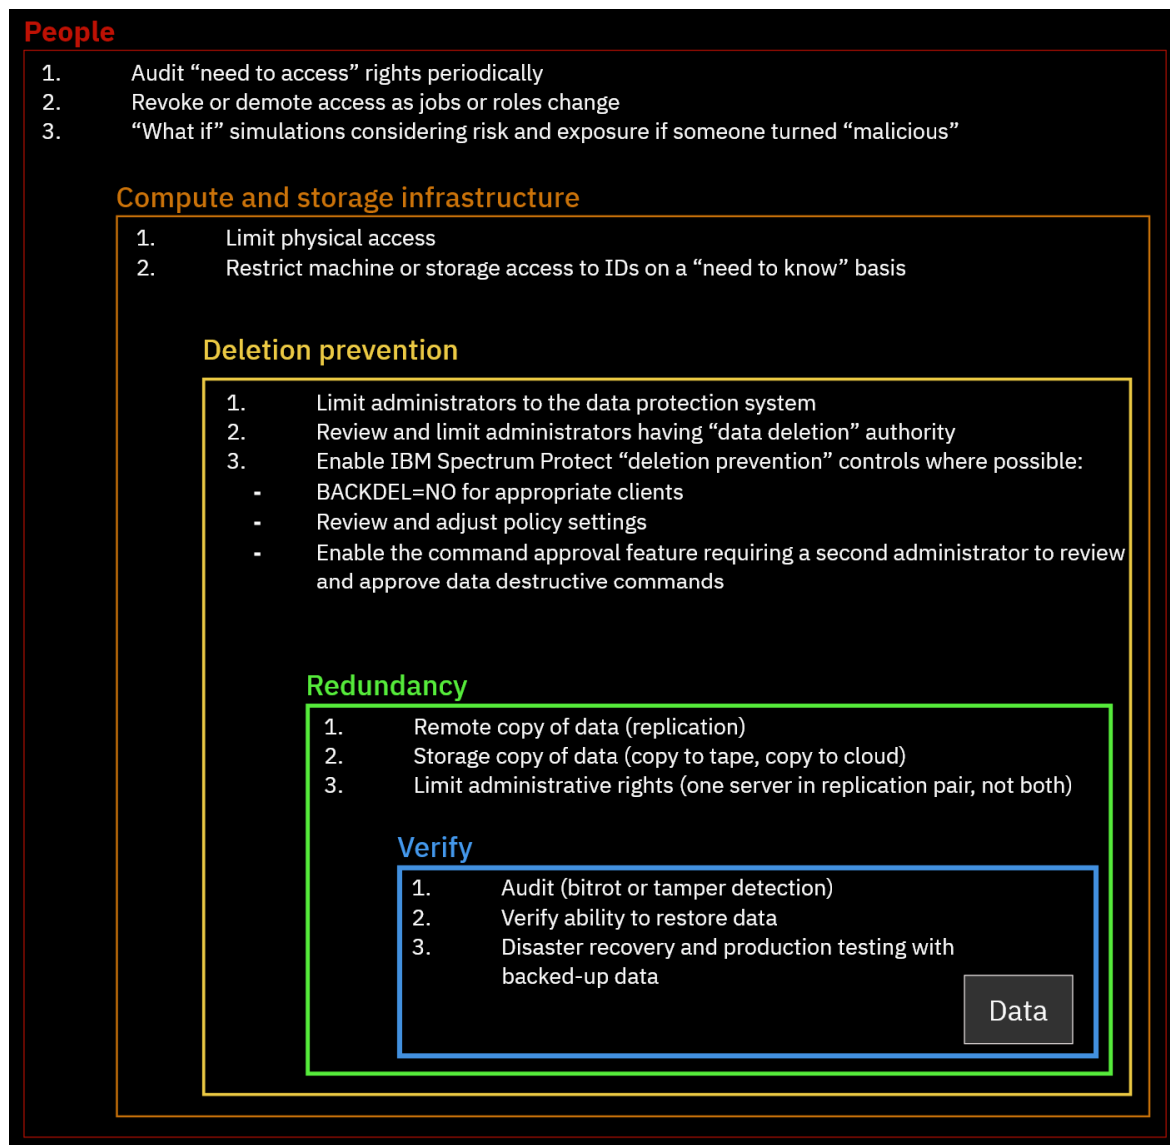


Figure 4: Layered security

1.9.1 Identity Security in the IBM Spectrum Protect Environment

IBM Spectrum Protect uses a novel *verb protocol* communication mechanism for client/server and server/server communication. This verb protocol is an API protocol in which each side of the communication channel initiates requests or replies with responses that drive behavior within the IBM Spectrum Protect client or server. When first establishing communication between two endpoints, the client and server or two servers authenticate with one another. With recent versions of IBM Spectrum Protect, Transport Layer Security (TLS) over TCP/IP is used by default as a security protocol for communication. IBM Spectrum Protect offers industry-standard local credential management with passwords that are encrypted at rest and decrypted at user sign-on and authentication. During the verb protocol sign-on procedure, TLS is used for authentication (certificate and authentication exchange).

IBM Spectrum Protect client node and administrator password security can be further enhanced by using Lightweight Directory Access Protocol (LDAP) to apply stricter requirements for passwords. This approach includes the integration of Microsoft Active Directory. With LDAP, the server receives credentials from the client over TLS, and the server “binds” to the LDAP directory service or completes an operation that is equivalent to a bind. For instructions about integrating IBM Spectrum Protect with LDAP, see [References](#) [17].

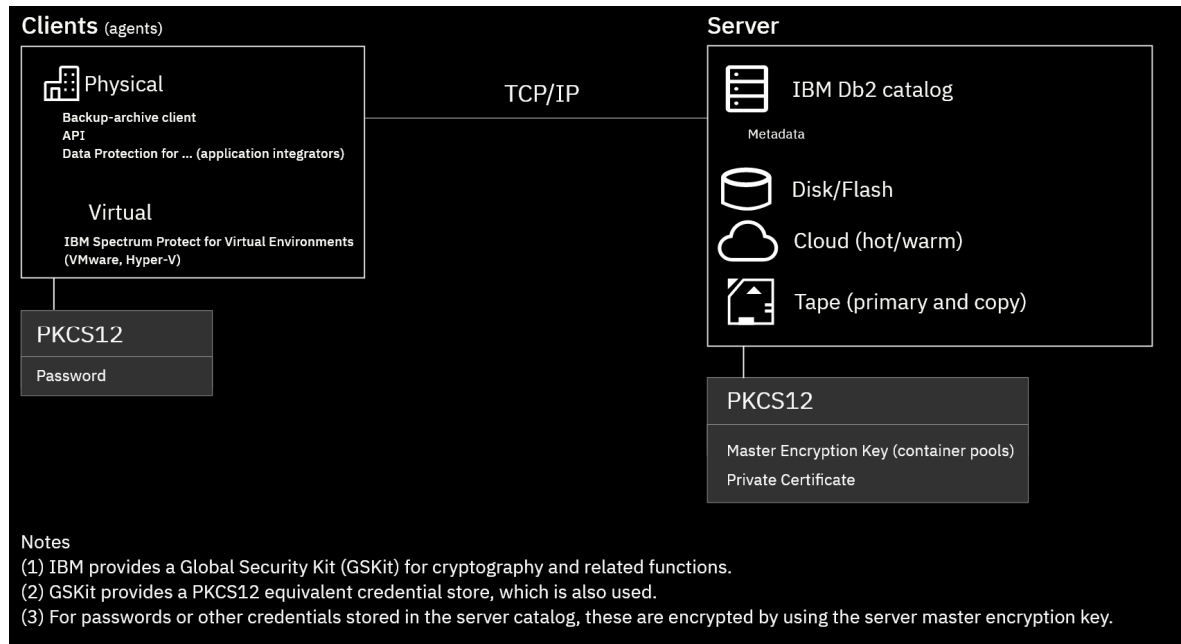


Figure 5: Credential security

The preferred method of identity security is to maintain at least one local administrator account that does not depend on LDAP or Active Directory in case the system that hosts LDAP or Active Directory is compromised by ransomware or other malware or is otherwise unavailable.

IBM Spectrum Protect TLS certificates can be self-signed or certificate authority (CA) signed. By default, the server generates a unique self-signed certificate during server formatting (initialization). Alternatively, a trusted CA-signed certificate provided by an organization can be installed for use by the server. For instructions about installing a trusted CA-signed certificate on the server, see [References](#) [18].

1.9.2 IBM Spectrum Protect Administrative Security

IBM Spectrum Protect offers different authority classes for administrators. The classes restrict or enable capability within the product. Administrative network traffic can also be directed over a server TCP/IP port that is separate from data traffic. In this way, a networking team can monitor different types of traffic with different tools or with different levels of granularity. IBM Spectrum Protect offers an audit trail in terms of activity log entries stored in the server's database for each operation an administrator performs. The IBM Spectrum Protect Operations Center can be used to specify certain audit-able messages as triggering "alerts" that can be configured to notify specific users via email.

"Command approval" is a mechanism that can be enabled within the server to ensure that an additional, second administrator must approve any potentially destructive server command, including data policy changes which might affect data retention. This helps to ensure a "four eyes" approach so that a single malicious, careless, or uninformed administrator cannot cause damage within an environment. For information setting up administrator command approval, see [References](#) [19].

1.9.3 IBM Spectrum Protect Data Security and Policy Based Management

Client data versioning and retention represent a critical aspect to a data security posture. IBM Spectrum Protect tracks every version of a backed-up file separately in its inventory. The files are not altered after they are stored. Rather, each version is treated as a separate object in the server's metadata catalog. IBM Spectrum Protect uses *policies* to control how the server stores and manages data objects through the hierarchy of policy domains, policy sets, management classes, and copy groups. You can specify IBM Spectrum Protect policy settings so that many versions of client data are protected for a longer period. In this way, you grant yourself more "freedom of action" in terms of having more backup versions, originating from a longer period of time, to provide a wider selection of backup versions for recovery. Even if a security event or malware infiltration is not caught promptly, having several versions of data stored within IBM Spectrum Protect (for example, 30, 45, or 60 days of data versions) helps to ensure recovery flexibility. Policy settings of interest include the `VEREXISTS` (versions existing), `VERDELETED` (versions deleted), `RETEXTRA` (retained extra), and `RETONLY` (retained only) settings.

To learn more about IBM Spectrum Protect policy-based data management, see [References](#) [20].

To further enhance data security, consider configuring single or multi-target IBM Spectrum Protect replication, where the operating systems of each IBM Spectrum Protect server differ. For example, the Microsoft Windows operating system is often an attack vector for malware. By using a different operating system platform for each of two IBM Spectrum Protect servers in a replication pair relationship, you might help to protect the system against malware that is designed to target one platform. Increasing the number of "attack vectors" in this way can build another barrier against malicious actors.

IBM Spectrum Protect can be configured with different retention policy settings at different ends of a replication relationship. For example, the source replication server can store data to local disk storage with a directory-container storage pool while the target replication server can store data to a cloud-container storage pool in the cloud. With more cost-effective object storage on the target, a longer retention period can be configured to protect more versions of client data for a longer time. Alternatively, IBM Spectrum Protect *tiering storage rules* can be used with a directory-container storage pool at the target site to tier older or inactive versions of data from disk to object storage for cost-effectiveness. In this case, a longer policy data retention period should be configured at the site from which data is tiered. In both cases, having more versions of data available helps to increase the "freedom of action" within the data protection system.

IBM Spectrum Protect clients can restore data from either end of a replication pair. For more information about setting up replication, see [References](#) [21].

1.9.4 IBM Spectrum Protect Bit-Level Security

IBM Spectrum Protect provides data stream security in various ways. During data transfer, a data stream can be encrypted by using Transport Layer Security (TLS). The TLS protocol performs data content digest and validation functions during transfer and provides for implicit tamper protection. TCP/IP packets are rejected if they appear to be tampered with. During security negotiation, the server prioritizes TLS 1.3 ciphers. Many cipher suites are supported.

The IBM Spectrum Protect Operations Center makes use of an anomaly detection mechanism that detects unexpected changes in client backups. For example, if data is successfully deduplicated for a workload, but then the data reduction rate drops significantly (to near 0%), this change could indicate that ransomware has encrypted the data. In general, an IBM Spectrum Protect administrator should pay attention to client deduplication and data reduction rates over time. Most client workloads make use of *incremental forever* backup operations. During an incremental backup operation, only new or changed data is backed up. If ransomware encrypts client data, the next incremental backup could be much larger than normal, indicating a possible problem.

IBM Spectrum Protect encrypts data at rest by using Galois/Counter Mode (GCM) signing and AES-256 encryption. These security mechanisms are applied to both directory-container and cloud-container storage pools. Each unique data extent that is identified by the data deduplication algorithm is encrypted separately by using distinct encryption keys after the extent is optionally compressed by using LZ4 compression.

Tip: A *data extent* is the part of a file that is created during data deduplication. Extents are compared with other file extents to identify duplicates.

These data extents are then further encrypted by using a server-level *master encryption key* that is protected as part of the IBM Spectrum Protect database backup process.

The validity of client objects is ensured in multiple ways. Each deduplicated extent has a Secure Hash Algorithm (SHA-1) value that is computed and stored in the server database. The value is used during data validation. An end-to-end object-level hash is also computed by using the xxHash (formerly MD5) hashing algorithm to ensure object-level integrity.

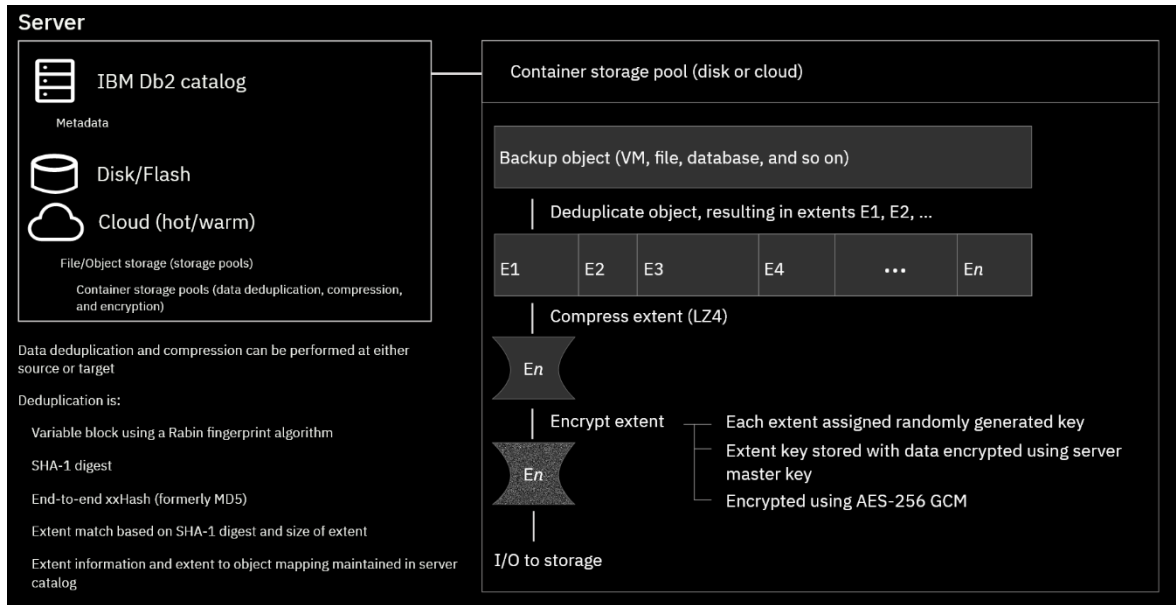


Figure 6: Container pool security

IBM Spectrum Protect offers explicit audit capabilities that can be enabled by using the **AUDIT CONTAINER** command for directory and cloud-container storage pools. This command can be run ad hoc or be scheduled to run regularly by using storage rules. The audit can help to detect tampering or *bit rot*, which is the slow deterioration of data. For cloud-container storage pools, the audit can be configured to validate object metadata entity tags, known as ETags. You can issue the `AUDIT CONTAINER` command with the `VALIDATECLOUDEXTENTS=NO` parameter setting to audit only ETags. Or you can issue the `AUDIT CONTAINER` command with the `VALIDATECLOUDEXTENTS=YES` parameter setting to enable a more thorough audit that inspects data content. This latter form should be used only when a problem with the data content is suspected because this type of audit can lead to a large amount of data egress and HTTP GET operations from the object storage system. For more information, see [AUDIT CONTAINER](#).

1.10 Optimizing Restore Performance from Cloud-Container Storage Pools

1.10.1 Enabling the Cloud-container Storage Pool Cloud Read Cache

Starting with IBM Spectrum Protect Version 8.1.12, a read cache for cloud-container storage pools can be enabled to help improve the performance of large restore operations. When this feature is enabled for a cloud-container storage pool, container objects are cached on disk within the "cloud cache" (storage pool directories defined to the cloud-container storage pool) during restore operations, which can in some cases improve total throughput.

You can enable a cloud read cache for a cloud-container storage pool by using the `CLOUDREADCACHE` parameter when you define or update a cloud-container storage pool.

When the `CLOUDREADCACHE` parameter is enabled for a cloud-container storage pool, the IBM Spectrum Protect server analyzes container data read requests for the storage pool. Read requests are initiated by client restore and retrieve activity. If significant read activity occurs within a short time for a cloud-container object, the server downloads the requested data from object storage, and then stages it as a copy to the cloud cache. The cloud cache is in local disk storage pool directories. The system uses the read cache data that was downloaded to the cloud cache to meet further client read requests for that cloud-container object. By using a cloud cache, you can improve the performance of restore operations in some circumstances because reading data from local disk storage pool directories can be much faster than sending small read requests to object storage over the Ethernet network.

Enabling the `CLOUDREADCACHE` parameter might be useful in the following circumstances:

- The IBM Spectrum Protect server is connected to the object storage system on a network with high bandwidth. For example, the network bandwidth is 5 Gbps, 10 Gbps, or 40 Gbps.
- The IBM Spectrum Protect server is connected to the object storage system on a network with high latency. For example, the object storage system might be an off-premises public cloud with 10 milliseconds of round-trip latency.
- You must restore more than 10 GB of data.
- You have an extent workload that is predominantly composed of small files, is stored with a low ratio of data deduplication, or both. For example, you have one of the following types of workload:
 - IBM Spectrum Protect backup-archive client, file server type workloads
 - Oracle database workloads
- You would like to restore data multiple times to different clients or for different purposes. For example, you might want to recover a virtual machine or database multiple times when setting up test environments.

Enabling the `CLOUDREADCACHE` parameter might not be useful in the following circumstances:

- The IBM Spectrum Protect server is connected to the object storage system on a network with low bandwidth. For example, the network bandwidth is 1 Gbps or 100 Mbps bandwidth.
- You must restore only a small amount of data (less than 10 GB).
- You have a large deduplicated extent workload and the IBM Spectrum Protect server is connected to the object storage system on a network with low latency. For example, the object storage system might be on the same local area network (LAN) as the cloud-container storage pool, and the round-trip latency might be less than 10 milliseconds.

In cases where the total network bandwidth to the object storage device is limited (that is, 1 Gbps or 100 Mbps), you might get better performance by leaving the `CLOUDREADCACHE` parameter disabled. In these circumstances, the default IBM Spectrum Protect restore

procedure from object storage might result in better performance by using small read requests for each deduplicated extent. However, if you repeatedly restore data over many hours or days, it might still be advantageous to enable the `CLOUDREADCACHE` parameter so that data is staged on higher-performing, local disk storage for subsequent restore activity.

For the latest guidance in the IBM Spectrum Protect documentation for recommendations on implementing and tuning this feature, see References [12].

1.10.2 Increasing the Minimum Extent Size for Client Nodes

Starting with IBM Spectrum Protect Version 8.1.12, it is possible to increase the minimum size of extents that are used for data deduplication operations at the client node level. In some circumstances, specifying a value that is higher than the default can improve the performance of backup, expiration, and restore operations. A larger extent size might also help to prevent unnecessary database growth.

During restore operations with cloud-container storage pools, data is fetched from object storage using individual HTTP GET requests for each deduplicated extent. Using a larger minimum extent size for a node will likely lead to larger HTTP GET requests from object storage, which can benefit restore throughput. The exception to this is when the cloud read cache is enabled for the storage pool, in which case larger read requests from object storage may occur as a part of container caching. The cloud read cache and larger minimum extent sizes can be used together to improve throughput performance for cloud-container storage pool restore operations.

In most system environments, the default extent size of 50 KB is appropriate. However, if you plan to deduplicate data from an Oracle or SAP database, consider increasing the extent size as these data types can often deduplicate with extents that are very small (less than 100 KB on average). To set the extent size, use the `REGISTER NODE` or `UPDATE NODE` command and specify the `MINIMUMEXTENTSIZE` parameter. Please see the latest guidance for this parameter in the `REGISTER NODE` and `UPDATE NODE` documentation to help evaluate your storage environment and determine whether to use a non-default value for a node.

One consideration is for environments that already have significant data stored for a client node. If the `MINIMUMEXTENTSIZE` parameter is changed for a node, then it is likely that that node's new backups will not deduplicate very well with older backups ingested with the default or other setting. Consider opening case with IBM Support if you are considering changing this value for an existing node with a large amount of stored data to evaluate this impact.

Google Cloud Platform Configurations

To access this cloud computing service, go to the [Google Cloud Platform](#) portal.

The following configurations represent preferred IBM Spectrum Protect deployments within Google Cloud Platform infrastructure. These deployments make use of Google Cloud Platform object storage. Data is ingested into a cloud-container storage pool backed by a **standard storage class, regional object storage bucket**, using the HTTPS protocol.

The Google Cloud Platform bucket that is used for IBM Spectrum Protect should be created in the same or similar region as the instance to minimize latency and to take advantage of the Google Cloud Platform internal network for endpoint communication, where possible.

For circumstances in which the IBM Spectrum Protect server is provisioned within a Google Cloud Platform Compute Engine virtual machine instance and this server is the target server for storage pool protection and node replication operations from a source server located at an on-premises data center, additional network bandwidth might be needed to satisfy replication requirements. Consider using Google Cloud Platform Interconnect to establish a dedicated, high bandwidth link from your premises to the Google Cloud Platform-based IBM Spectrum Protect instance running within the remote virtual machine. Consider this option if your anticipated back-end replication workload will cause the needed throughput to exceed the bandwidth capacity between the replication pair. For example, if an IBM Spectrum Protect source server were to ingest 40 TB of front-end data that reduced to 10 TB after (2:1) data deduplication and (2:1) data compression, then that 10 TB of back-end data would need to be replicated during the daily storage pool protection and node replication window. A 1 Gbps dedicated link can support approximately 3.5 TB per hour of throughput. Google Cloud Platform Interconnect offers bandwidth options of 10 or 100 Gbps.. For more information about Google Cloud Platform Interconnect, see [References](#) [6]. For each of the instances, isolated throughput measurements should be conducted with a Google Cloud Platform storage API tool to determine maximum throughput to collocated Google Cloud Platform object storage within that region.

In many cases, the provisioned disk sizes chosen would appear to be too large for the IBM Spectrum Protect purpose they are dedicated to. The reason for this is because Google Cloud Platform block disk devices make use of “IOPs density” – IOPs and throughput capability increase with the capacity of the block disk device. Larger disk devices than needed were chosen in the configurations below to satisfy throughput and IOPs needs.

For each virtual machine instance type shown, there are Intel and AMD-based processor variations. For example, “n2-standard-64” is an Intel-based instance while “n2d-standard-64” is an AMD-based instance. Depending on the Google Cloud Platform region, Intel-based instances might be more common. However, AMD-based instances are generally more cost-effective and offer similar performance to their Intel-based counterparts. Both variations are presented as options here.

Table 3: Google Cloud Platform, large configuration

Cloud component	Google Cloud Platform component	Detailed description	Quantity
-----------------	---------------------------------	----------------------	----------

Cloud component	Google Cloud Platform component	Detailed description	Quantity
Server and network	nd(d)-standard-64	64-core Intel or AMD	1
		256 GB RAM	
		10 Gbit or more Ethernet connectivity to object storage	
	Operating system	RHEL or Ubuntu Linux server	1
Block storage	100 GB "standard" disk	Operating system disk	1
	100 GB "standard" disk	IBM Spectrum Protect instance disk	1
	1250 GB "ssd" disk	IBM Spectrum Protect database and active log disks	8
	8 TB "standard" disk	IBM Spectrum Protect database archive log disk	1
	10 TB "balanced" disk	IBM Spectrum Protect cloud cache disk	1
Object storage	Google Cloud Platform standard storage class bucket	IBM Spectrum Protect object storage bucket created and used in same Google Cloud Platform region	1

Cloud component	Google Cloud Platform component	Detailed description	Quantity
		as instance	

Table 4: Google Cloud Platform, medium configuration

Cloud component	Google Cloud Platform component	Detailed description	Quantity
Server and network	n2(d)-standard-32	32-core Intel or AMD	1
		128 GB RAM	
		10 Gbit or more Ethernet connectivity to object storage	
	Operating system	RHEL or Ubuntu Linux server	1
Block storage	100 GB “standard” disk	Operating system disk	1
	100 GB “standard” disk	IBM Spectrum Protect instance disk	1
	1 TB “ssd” disk	IBM Spectrum Protect database and active log disks	5
	4 TB “standard” disk	IBM Spectrum Protect database archive log disk	1
	4 TB “balanced” disk	IBM Spectrum Protect cloud cache disk	1

Cloud component	Google Cloud Platform component	Detailed description	Quantity
Object storage	Google Cloud Platform standard storage class bucket	IBM Spectrum Protect object storage bucket created and used in same Google Cloud Platform region as instance	1

Table 5: Google Cloud Platform, small configuration

Cloud component	Google Cloud Platform component	Detailed description	Quantity
Server and network	n2(d)-standard-16	16-core Intel or AMD	1
		64 GB RAM	
		10 Gbit or more Ethernet connectivity to object storage	
	Operating system	RHEL or Ubuntu Linux server	1
Block storage	100 GB "standard" disk	Operating system disk	1
	100 GB "standard" disk	IBM Spectrum Protect instance disk	1
	1 TB "ssd" disk	IBM Spectrum Protect database and active log disks	3

	1 TB “standard” disk	IBM Spectrum Protect database archive log disk	1
	3300 GB “balanced” disk	IBM Spectrum Protect cloud cache disk	1
Object storage	Google Cloud Platform standard storage class bucket	IBM Spectrum Protect object storage bucket created and used in same Google Cloud Platform region as instance	1

Table 6: Google Cloud Platform, extra-small configuration

Cloud component	Google Cloud Platform component	Detailed description	Quantity
Server and network	n2(d)-standard-4	4-core Intel or AMD	1
		16 GB RAM	
		1 Gbit or more Ethernet connectivity to object storage	
	Operating system	RHEL or Ubuntu Linux server	1
Block storage	100 GB “standard” disk	Operating system disk	1
	50 GB “standard” disk	IBM Spectrum Protect instance disk	1
	64 GB “ssd” disk	IBM Spectrum Protect database	5

Cloud component	Google Cloud Platform component	Detailed description	Quantity
		disks and active log disks	
	512 GB “standard”	IBM Spectrum Protect archive log disk	1
	1200 GB “balanced” disk	IBM Spectrum Protect cloud cache disk	1
Object storage	Google Cloud Platform standard storage class bucket	IBM Spectrum Protect object storage bucket created and used in same Google Cloud Platform region as instance	1

2.1 Compute, Disk, and File System Setup Instructions for Linux-based Systems

The following guidance assumes that IBM Cloud compute, block disk, and object storage resources specified in the above tables were deployed for a chosen Blueprint size configuration.

For IBM Spectrum Protect deployments on Google Cloud Platform computing systems, the preferred operating system is Linux, either the latest IBM Spectrum Protect supported Red Hat Enterprise Linux (RHEL) or Ubuntu Linux. With Ubuntu Linux, care should be taken to ensure that all required Linux packages are installed to enable Linux Logical Volume Manager (LVM) functionality. For more information about the operating systems, see the IBM Spectrum Protect technote ([References](#) [8]).

2.1.1 Linux Operating System Packages and General Guidance

The following actions represent the best guidance for your Linux operating system packages. Ensure that the following Linux Operating System packages are installed:

- ftp
 - perl
 - libnsl
 - nss
-

- ncurses-compat-libs
- device-mapper-multipath
- ksh (x86_64)
- lvm2

Ensure that the following Linux kernel parameters are set to their preferred value:

Parameter	Description	Preferred value
<code>kernel.randomize_va_space</code>	Configured the use of memory address space layout randomization (ASLR) for the kernel. If you set this value to 0, you disable ASLR. To learn more about Linux ASLR and Db2, see technote 1365583.	Set this parameter value to 2, which is the default value for the operating system. Later, if you decide not to use ASLR, you can reset the value to 0.
<code>vm.swappiness</code>	Defines whether the kernel can swap application memory out of physical random access memory (RAM). For more information about kernel parameters, see the Db2 product information.	5
<code>vm.overcommit_memory</code>	Influences how much virtual memory can be allocated, based on kernel parameter settings. For more information about kernel parameters, see Db2 product information.	0

Enable and start the multipath daemon:

```
mpathconf -enable -with_multipathd y
```

Disable SELinux prior to IBM Spectrum Protect server installation by editing the `/etc/sysconfig/selinux` file and setting **SELINUX** to “permissive” or “disabled” and restarting the operating system.

Create the following directory paths:

```
mkdir /sp
mkdir /sp/tsminst1
```

```
mkdir /sp/sp_db1
mkdir /sp/sp_db2
mkdir /sp/sp_db3
mkdir /sp/sp_db4
mkdir /sp/sp_alog
mkdir /sp/sp_archlog
mkdir /sp/sp_cc
```

For Medium or Large Blueprint systems, also create the following directory paths:

```
mkdir /sp/sp_db5
mkdir /sp/sp_db6
mkdir /sp/sp_db7
mkdir /sp/sp_db8
```

2.1.2 Large Blueprint Instructions

2.1.2.1 Disk Setup Instructions

IBM Spectrum Protect / Db2 Instance Directory

Using the `lsblk` command, determine which disk device was specified for the instance disk and then use the disk device name in the following commands in place of {Instance Disk Device}:

```
mkfs.xfs /dev/{Instance Disk Device}
mount /dev/{Instance Disk Device} /sp/tsminst1
```

IBM Spectrum Protect / Db2 Database and Active Log Directories

Using the `lsblk` command, determine which disk device was specified for the database disks and then use the disk device names in the following commands in place of {DB Disk 1-8}:

```
pvcreate /dev/{DB Disk 1}
pvcreate /dev/{DB Disk 2}
pvcreate /dev/{DB Disk 3}
pvcreate /dev/{DB Disk 4}
pvcreate /dev/{DB Disk 5}
pvcreate /dev/{DB Disk 6}
pvcreate /dev/{DB Disk 7}
pvcreate /dev/{DB Disk 8}
```

```
vgcreate sp_db /dev/{DB Disk 1} /dev/{DB Disk 2} /dev/{DB Disk 3}  
/dev/{DB Disk 4} /dev/{DB Disk 5} /dev/{DB Disk 6} /dev/{DB Disk  
7} /dev/{DB Disk 8}
```

```
lvcreate --stripes 8 --stripesize 16 --extents 304000 --name  
sp_db_db1 sp_db
```

```
lvcreate --stripes 8 --stripesize 16 --extents 304000 --name  
sp_db_db2 sp_db
```

```
lvcreate --stripes 8 --stripesize 16 --extents 304000 --name  
sp_db_db3 sp_db
```

```
lvcreate --stripes 8 --stripesize 16 --extents 304000 --name  
sp_db_db4 sp_db
```

```
lvcreate --stripes 8 --stripesize 16 --extents 304000 --name  
sp_db_db5 sp_db
```

```
lvcreate --stripes 8 --stripesize 16 --extents 304000 --name  
sp_db_db6 sp_db
```

```
lvcreate --stripes 8 --stripesize 16 --extents 304000 --name  
sp_db_db7 sp_db
```

```
lvcreate --stripes 8 --stripesize 16 --extents 304000 --name  
sp_db_db8 sp_db
```

```
lvcreate --stripes 8 --stripesize 16 --extents 100%FREE --name  
sp_alog sp_db
```

```
mkfs.ext4 /dev/mapper/sp_db-sp_db_db1
```

```
mkfs.ext4 /dev/mapper/sp_db-sp_db_db2
```

```
mkfs.ext4 /dev/mapper/sp_db-sp_db_db3
```

```
mkfs.ext4 /dev/mapper/sp_db-sp_db_db4
```

```
mkfs.ext4 /dev/mapper/sp_db-sp_db_db5
```

```
mkfs.ext4 /dev/mapper/sp_db-sp_db_db6
```

```
mkfs.ext4 /dev/mapper/sp_db-sp_db_db7
```

```
mkfs.ext4 /dev/mapper/sp_db-sp_db_db8
```

```
mkfs.ext4 /dev/mapper/sp_db-sp_alog
```

```
mount /dev/mapper/sp_db-sp_db_db1 /sp/sp_db1
```

```
mount /dev/mapper/sp_db-sp_db_db2 /sp/sp_db2
```

```
mount /dev/mapper/sp_db-sp_db_db3 /sp/sp_db3
```

```
mount /dev/mapper/sp_db-sp_db_db4 /sp/sp_db4
```

```
mount /dev/mapper/sp_db-sp_db_db5 /sp/sp_db5
mount /dev/mapper/sp_db-sp_db_db6 /sp/sp_db6
mount /dev/mapper/sp_db-sp_db_db7 /sp/sp_db7
mount /dev/mapper/sp_db-sp_db_db8 /sp/sp_db8
mount /dev/mapper/sp_db-sp_alog /sp/sp_alog
```

IBM Spectrum Protect / Db2 Archive Log Directory

Using the `lsblk` command, determine which disk device was specified for the Db2 archive log disk and then use the disk device name in the following commands in place of {Archive Log Disk Device}:

```
mkfs.ext4 /dev/{Archive Log Disk Device}
mount /dev/{Archive Log Disk Device} /sp/sp_archlog
```

IBM Spectrum Protect Cloud-container Storage Pool Cloud Cache

Using the `lsblk` command, determine which disk device was specified for the cloud-container storage pool cloud cache disks and then use the disk device names in the following commands in place of {Cloud Cache Disk 1-2}:

```
pvcreate /dev/{Cloud Cache Disk 1}
pvcreate /dev/{Cloud Cache Disk 2}
vgcreate sp_cc /dev/{Cloud Cache Disk 1} /dev/{Cloud Cache Disk 2}
lvcreate --stripes 2 --stripesize 256 --extents 100%FREE --name
sp_cc1 sp_cc

mkfs.xfs /dev/mapper/sp_cc-sp_cc1
mount /dev/mapper/sp_cc-sp_cc1 /sp/sp_cc
```

2.1.3 Medium Blueprint Instructions

2.1.3.1 Disk Setup Instructions

IBM Spectrum Protect / Db2 Instance Directory

Using the `lsblk` command, determine which disk device was specified for the instance disk and then use the disk device name in the following commands in place of {Instance Disk Device}:

```
mkfs.xfs /dev/{Instance Disk Device}
mount /dev/{Instance Disk Device} /sp/tsminst1
```

IBM Spectrum Protect / Db2 Database and Active Log Directories

Using the `lsblk` command, determine which disk device was specified for the database disks and then use the disk device names in the following commands in place of {DB Disk 1-5}:

```
pvcreate /dev/{DB Disk 1}
pvcreate /dev/{DB Disk 2}
pvcreate /dev/{DB Disk 3}
pvcreate /dev/{DB Disk 4}
pvcreate /dev/{DB Disk 5}

vgcreate sp_db /dev/{DB Disk 1} /dev/{DB Disk 2} /dev/{DB Disk 3}
/dev/{DB Disk 4} /dev/{DB Disk 5}

lvcreate --stripes 5 --stripesize 16 --extents 147460 --name
sp_db_db1 sp_db

lvcreate --stripes 5 --stripesize 16 --extents 147460 --name
sp_db_db2 sp_db

lvcreate --stripes 5 --stripesize 16 --extents 147460 --name
sp_db_db3 sp_db

lvcreate --stripes 5 --stripesize 16 --extents 147460 --name
sp_db_db4 sp_db

lvcreate --stripes 5 --stripesize 16 --extents 147460 --name
sp_db_db5 sp_db

lvcreate --stripes 5 --stripesize 16 --extents 147460 --name
sp_db_db6 sp_db

lvcreate --stripes 5 --stripesize 16 --extents 147460 --name
sp_db_db7 sp_db

lvcreate --stripes 5 --stripesize 16 --extents 147460 --name
sp_db_db8 sp_db

lvcreate --stripes 5 --stripesize 16 --extents 100%FREE --name
sp_alog sp_db

mkfs.ext4 /dev/mapper/sp_db-sp_db_db1
mkfs.ext4 /dev/mapper/sp_db-sp_db_db2
mkfs.ext4 /dev/mapper/sp_db-sp_db_db3
```

```
mkfs.ext4 /dev/mapper/sp_db-sp_db_db4
mkfs.ext4 /dev/mapper/sp_db-sp_db_db5
mkfs.ext4 /dev/mapper/sp_db-sp_db_db6
mkfs.ext4 /dev/mapper/sp_db-sp_db_db7
mkfs.ext4 /dev/mapper/sp_db-sp_db_db8
mkfs.ext4 /dev/mapper/sp_db-sp_alog
```

```
mount /dev/mapper/sp_db-sp_db_db1 /sp/sp_db1
mount /dev/mapper/sp_db-sp_db_db2 /sp/sp_db2
mount /dev/mapper/sp_db-sp_db_db3 /sp/sp_db3
mount /dev/mapper/sp_db-sp_db_db4 /sp/sp_db4
mount /dev/mapper/sp_db-sp_db_db5 /sp/sp_db5
mount /dev/mapper/sp_db-sp_db_db6 /sp/sp_db6
mount /dev/mapper/sp_db-sp_db_db7 /sp/sp_db7
mount /dev/mapper/sp_db-sp_db_db8 /sp/sp_db8
mount /dev/mapper/sp_db-sp_alog /sp/sp_alog
```

IBM Spectrum Protect / Db2 Archive Log Directory

Using the **lsblk** command, determine which disk device was specified for the Db2 archive log disk and then use the disk device name in the following commands in place of {Archive Log Disk Device}:

```
mkfs.ext4 /dev/{Archive Log Disk Device}
mount /dev/{Archive Log Disk Device} /sp/sp_archlog
```

IBM Spectrum Protect Cloud-container Storage Pool Cloud Cache

Using the **lsblk** command, determine which disk device was specified for the cloud-container storage pool cloud cache disk and then use the disk device name in the following commands in place of {Cloud Cache Disk 1}:

```
pvcreate /dev/{Cloud Cache Disk 1}
vgcreate sp_cc /dev/{Cloud Cache Disk 1}
lvcreate --stripes 1 --stripesize 256 --extents 100%FREE --name
sp_cc1 sp_cc

mkfs.xfs /dev/mapper/sp_cc-sp_cc1
```

```
mount /dev/mapper/sp_cc-sp_cc1 /sp/sp_cc
```

2.1.4 Small Blueprint Instructions

2.1.4.1 Disk Setup Instructions

IBM Spectrum Protect / Db2 Instance Directory

Using the `lsblk` command, determine which disk device was specified for the instance disk and then use the disk device name in the following commands in place of {Instance Disk Device}:

```
mkfs.xfs /dev/{Instance Disk Device}
mount /dev/{Instance Disk Device} /sp/tsminst1
```

IBM Spectrum Protect / Db2 Database and Active Log Directories

Using the `lsblk` command, determine which disk device was specified for the database disks and then use the disk device names in the following commands in place of {DB Disk 1-3}:

```
pvcreate /dev/{DB Disk 1}
pvcreate /dev/{DB Disk 2}
pvcreate /dev/{DB Disk 3}
vgcreate sp_db /dev/{DB Disk 1} /dev/{DB Disk 2} /dev/{DB Disk 3}
lvcreate --stripes 3 --stripesize 16 --extents 180225 --name
sp_db_db1 sp_db
lvcreate --stripes 3 --stripesize 16 --extents 180225 --name
sp_db_db2 sp_db
lvcreate --stripes 3 --stripesize 16 --extents 180225 --name
sp_db_db3 sp_db
lvcreate --stripes 3 --stripesize 16 --extents 180225 --name
sp_db_db4 sp_db
lvcreate --stripes 3 --stripesize 16 --extents 100%FREE --name
sp_alog sp_db

mkfs.ext4 /dev/mapper/sp_db-sp_db_db1
mkfs.ext4 /dev/mapper/sp_db-sp_db_db2
mkfs.ext4 /dev/mapper/sp_db-sp_db_db3
mkfs.ext4 /dev/mapper/sp_db-sp_db_db4
mkfs.ext4 /dev/mapper/sp_db-sp_alog
```

```
mount /dev/mapper/sp_db-sp_db_db1 /sp/sp_db1
mount /dev/mapper/sp_db-sp_db_db2 /sp/sp_db2
mount /dev/mapper/sp_db-sp_db_db3 /sp/sp_db3
mount /dev/mapper/sp_db-sp_db_db4 /sp/sp_db4
mount /dev/mapper/sp_db-sp_alog /sp/sp_alog
```

IBM Spectrum Protect / Db2 Archive Log Directory

Using the `lsblk` command, determine which disk device was specified for the Db2 archive log disk and then use the disk device name in the following commands in place of {Archive Log Disk Device}:

```
mkfs.ext4 /dev/{Archive Log Disk Device}
mount /dev/{Archive Log Disk Device} /sp/sp_archlog
```

IBM Spectrum Protect Cloud-container Storage Pool Cloud Cache

Using the `lsblk` command, determine which disk device was specified for the cloud-container storage pool cloud cache disk and then use the disk device name in the following commands in place of {Cloud Cache Disk 1}:

```
pvcreate /dev/{Cloud Cache Disk 1}
vgcreate sp_cc /dev/{Cloud Cache Disk 1}
lvcreate --stripes 1 --stripesize 256 --extents 100%FREE --name
sp_cc1 sp_cc

mkfs.xfs /dev/mapper/sp_cc-sp_cc1
mount /dev/mapper/sp_cc-sp_cc1 /sp/sp_cc
```

2.1.5 Extra-Small Blueprint Instructions

2.1.5.1 Disk Setup Instructions

IBM Spectrum Protect / Db2 Instance Directory

Using the `lsblk` command, determine which disk device was specified for the instance disk and then use the disk device name in the following commands in place of {Instance Disk Device}:

```
mkfs.xfs /dev/{Instance Disk Device}
mount /dev/{Instance Disk Device} /sp/tsminst1
```

IBM Spectrum Protect / Db2 Database and Active Log Directories

Using the `lsblk` command, determine which disk device was specified for the database disks and then use the disk device names in the following commands in place of {DB Disk 1-5}:

```
pvcreate /dev/{DB Disk 1}
pvcreate /dev/{DB Disk 2}
pvcreate /dev/{DB Disk 3}
pvcreate /dev/{DB Disk 4}
pvcreate /dev/{DB Disk 5}

vgcreate sp_db /dev/{DB Disk 1} /dev/{DB Disk 2} /dev/{DB Disk 3}
/dev/{DB Disk 4} /dev/{DB Disk 5}

lvcreate --stripes 5 --stripesize 16 --extents 14080 --name
sp_db_db1 sp_db

lvcreate --stripes 5 --stripesize 16 --extents 14080 --name
sp_db_db2 sp_db

lvcreate --stripes 5 --stripesize 16 --extents 14080 --name
sp_db_db3 sp_db

lvcreate --stripes 5 --stripesize 16 --extents 14080 --name
sp_db_db4 sp_db

lvcreate --stripes 5 --stripesize 16 --extents 100%FREE --name
sp_alog sp_db

mkfs.ext4 /dev/mapper/sp_db-sp_db_db1
mkfs.ext4 /dev/mapper/sp_db-sp_db_db2
mkfs.ext4 /dev/mapper/sp_db-sp_db_db3
mkfs.ext4 /dev/mapper/sp_db-sp_db_db4
mkfs.ext4 /dev/mapper/sp_db-sp_alog

mount /dev/mapper/sp_db-sp_db_db1 /sp/sp_db1
mount /dev/mapper/sp_db-sp_db_db2 /sp/sp_db2
mount /dev/mapper/sp_db-sp_db_db3 /sp/sp_db3
mount /dev/mapper/sp_db-sp_db_db4 /sp/sp_db4
mount /dev/mapper/sp_db-sp_alog /sp/sp_alog
```

IBM Spectrum Protect / Db2 Archive Log Directory

Using the `lsblk` command, determine which disk device was specified for the Db2 archive log disk and then use the disk device name in the following commands in place of {Archive Log Disk Device}:

```
mkfs.ext4 /dev/{Archive Log Disk Device}
mount /dev/{Archive Log Disk Device} /sp/sp_archlog
```

IBM Spectrum Protect Cloud-container Storage Pool Cloud Cache

Using the `lsblk` command, determine which disk device was specified for the cloud-container storage pool cloud cache disk and then use the disk device name in the following commands in place of {Cloud Cache Disk 1}:

```
pvcreate /dev/{Cloud Cache Disk 1}

vgcreate sp_cc /dev/{Cloud Cache Disk 1}
lvcreate --stripes 1 --stripesize 256 --extents 100%FREE --name
sp_cc1 sp_cc

mkfs.xfs /dev/mapper/sp_cc-sp_cc1
mount /dev/mapper/sp_cc-sp_cc1 /sp/sp_cc
```

2.1.6 IBM Spectrum Protect Format and Server Options

First, create a Linux group and user for use with IBM Spectrum Protect and Db2:

```
groupadd tsmsrvrs
useradd -g tsmsrvrs -m tsminst1
```

Change ownership of all of the file system directories associated with IBM Spectrum Protect to this new user:

```
chown -R tsminst1:tsmsrvrs /sp/
```

As the `root` user, create the Db2 instance for IBM Spectrum Protect:

```
/opt/tivoli/tsm/db2/instance/db2icrt -u tsminst1 tsminst1
```

As the `tsminst1` user, format the IBM Spectrum Protect server instance using the following command, depending on the Blueprint size:

Large Blueprint:

```

dmserv -i /sp/tsminst1 format
dbdir=/sp/sp_db1,/sp/sp_db2,/sp/sp_db3,/sp/sp_db4,/sp/sp_db5,
/sp/sp_db6,/sp/sp_db7,/sp/sp_db8 activelogsiz=450000
activelogdir=/sp/sp_alog archlogdir=/sp/sp_archlog

```

Medium Blueprint:

```

dmserv -i /sp/tsminst1 format
dbdir=/sp/sp_db1,/sp/sp_db2,/sp/sp_db3,/sp/sp_db4,/sp/sp_db5,
/sp/sp_db6,/sp/sp_db7,/sp/sp_db8 activelogsiz=450000
activelogdir=/sp/sp_alog archlogdir=/sp/sp_archlog

```

Small Blueprint:

```

dmserv -i /sp/tsminst1 format
dbdir=/sp/sp_db1,/sp/sp_db2,/sp/sp_db3,/sp/sp_db4
activelogsiz=200000 activelogdir=/sp/sp_alog
archlogdir=/sp/sp_archlog

```

Extra-Small Blueprint:

```

dmserv -i /sp/tsminst1 format
dbdir=/sp/sp_db1,/sp/sp_db2,/sp/sp_db3,/sp/sp_db4
activelogsiz=90000 activelogdir=/sp/sp_alog
archlogdir=/sp/sp_archlog

```

Ensure that the dmserv.opt IBM Spectrum Protect server options file contains the following pieces of information:

```

ACTIVELOGSize           {Active Log Size}
ACTIVELOGDirectory      /sp/sp_alog
ARCHLOGDirectory        /sp/sp_archlog

```

Where {Active Log Size} matches the following, depending on Blueprint size:

```

Large Blueprint:        450000
Medium Blueprint:       450000
Small Blueprint:        200000
Extra-Small Blueprint:  90000

```

2.2 Design Considerations for Google Cloud Platform Compute Engine Instances

A range of Google Cloud Platform Compute Engine instances are available that provide copious options for choosing vCPU and memory combinations, along with specific tuning for compute, memory, storage, or machine learning and analytics intensive workloads. The N2 series “standard” class of Google Cloud Platform Compute Engine “general-purpose” instances provide the best mix of vCPU and memory resources necessary to build extra-small, small, medium, and large Blueprint-level IBM Spectrum Protect servers while also featuring the availability of sufficient Ethernet and block storage connectivity. These instances provide an economically viable and well-balanced instance family suited for use with IBM Spectrum Protect.

Beginning with IBM Spectrum Protect V8.1.3, the server automatically throttles client backup operations if the cloud accelerator cache portion of a direct-to-cloud storage pool is nearing full capacity. As a result, it is not mandatory to configure cloud accelerator disk cache space that would be large enough to hold a full day’s worth of backups (after data deduplication and compression). However, disk benchmarks should be run to ensure that the anticipated back-end workload that an IBM Spectrum Protect server is expected to support will not result in this disk location being the primary bottleneck of the system (see [Disk Benchmarking](#)).

Google Cloud Platform enforces various quotas and limits on Cloud Storage. This includes limits on object storage bucket name lengths, object creation and deletion rates, maximum object sizes, naming conventions, etc. IBM Spectrum Protect is implemented so as not to violate these restrictions. For network bandwidth, there is a limit of 200 Gbps per Google Cloud Platform “project”. If multiple IBM Spectrum Protect server instances are being configured within the same project, or the project containing the IBM Spectrum Protect server(s) is used for other purposes, consider this limit during your deployment and/or request a quota increase. Note that this only needs to be considered if Google services are being read out of one region and into another (not a very common configuration). For more information on these limits and quotas, see References [13].

Google Cloud Platform block disk storage volumes can be configured as “standard”, “ssd”, “balanced”, or “extreme” disk types. In general, the IBM Spectrum Protect system operating system disk, Db2 instance disk, and archive log perform acceptably with “standard” type disks. The Db2 database and active log disks perform best when configured with “ssd” type disks. “Extreme” disks can provide greater IOPs performance, however this is unnecessary to satisfy IBM Spectrum Protect database requirements. The IBM Spectrum Protect cloud-container storage pool cloud cache is a best match with “balanced” type disks, which provide an economical balance between IOPs/throughput performance and cost while still satisfying the overlapped I/O demands of the cloud cache disk.

Google Cloud Platform block disk storage volumes are encrypted by default. Disk encryption can be managed using Google-managed, customer-managed (CMEK), or customer-supplied (CSEK) encryption keys. With Google-managed encryption keys, no additional configuration or management is required when ordering virtual machine instance

disks. With CMEK keys, the Google Cloud Key Management Service can be used to manage encryption keys. With CSEK keys, the keys themselves are managed outside of the Google Cloud Platform and are provided by the customer at time of need. When configuring disk encryption for IBM Spectrum Protect server instance disks, choose the management strategy that best aligns with your security posture.

Google Cloud Platform object storage buckets can be configured with the following storage classes: standard, nearline, coldline, or archive. The standard storage class provides for the cheapest egress cost but greatest resident storage cost. For cloud-container storage pools that predominantly store primary (and/or “active”) data, the standard storage class should be used since background processes such as cloud reclamation with the IBM Spectrum Protect server and recovery activities can drive egress activity against object storage. For data that will be retained for at least 30 days and that will be restored infrequently (for example, long-term or compliance retention data and archive data), consider directing this workload to an IBM Spectrum Protect cloud-container storage pool with an underlying Google Cloud Platform object storage bucket with a storage class other than standard to save on costs. However, be aware that segmenting your data ingestion workload into more than one cloud-container storage pool could potentially reduce the overall data deduplication efficiency of the solution, because IBM Spectrum Protect data de-duplication takes place within the granularity of the container storage pool. Consider balancing the cost savings from segmenting longer-term retention workloads with potential losses in data deduplication efficiency and a greater data storage footprint. For the most comprehensive list of storage class support for the various cloud vendors, see References [14].

Object storage buckets can also be configured as regional, multi-regional, or dual-regional. In general, the best performance can be experienced with regional object storage buckets. If greater data resiliency is desired, IBM Spectrum Protect can support dual-regional or multi-region buckets for all feature purposes. IBM Spectrum Protect does not currently internally take advantage of buckets that span multiple regions.

2.2.1 Considerations for Direct-to-Cloud Architectures

The **direct-to-cloud** configurations discussed here are architected such that the disk of the cloud accelerator cache is on fast-performing SSD (Google Cloud Platform “balanced” block disk storage). In general, the limiting factor for end-to-end ingestion throughput for an IBM Spectrum Protect server using object storage is the network to the object storage system or the object storage system itself. For this reason, you must understand the throughput capability requirements of the planned system. Those requirements will drive your decisions about assigning servers to Ethernet networks. Next, the ceiling for daily ingestion throughput (in terms of mebibytes per second or terabytes per day) must be determined at the object storage end by using **object storage benchmarking** tests. The test results will establish what the overlapped I/O capability would have to be for the cloud accelerator cache disk location for maximum back-end throughput. These two factors will help you select a disk technology to sustain daily backup ingestion requirements while working within object storage limits. After you select a disk technology, run **disk benchmarking** tests to verify throughput capability (see [Disk Benchmarking](#)).

Tip: In this paper, the abbreviation MiB is used for mebibytes, the abbreviation TiB is used for tebibytes, the abbreviation KiB is used for kibibytes, and the abbreviation GiB is used for gibibytes.

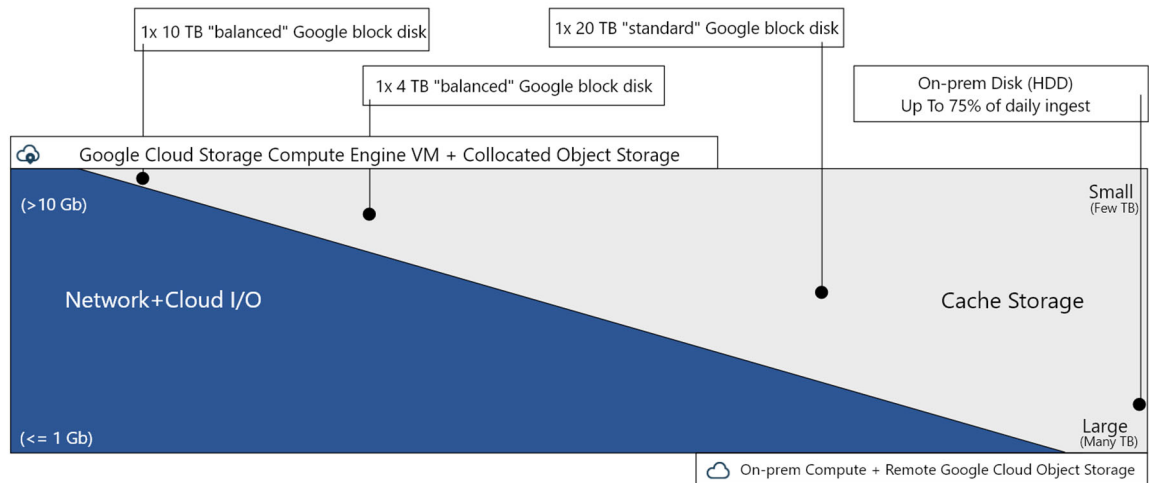
Example

If an object storage link is capable of 10 Gbps, this data transfer speed equals about 1000 MiB/s after packet overhead and other efficiency loss. In order to saturate this link for long periods, the cloud accelerator cache disk location must be capable of taking in client ingestion data (writes) at 1000 MiB/s and transmitting staged data to object storage (reads) at a similar speed, 1000 MiB/s (~128-256 KiB I/O size). This capacity ensures that the cloud accelerator cache disk can remain as small as possible while sustaining maximum throughput. **Alternatively**, a larger capacity, slower disk technology (such as Google Cloud Platform “standard” block storage disks) can be used such that the client ingestion data that has been staged to accelerator disk cache can be transmitted to object storage over a longer period of the day (extending past the backup window). However, be aware that data residing only in the cloud accelerator cache is unprotected in the sense that only a single copy of the data exists. The redundancy protection inherent in cloud object storage is available only if the data is transmitted to object storage. Generally, Google Cloud Platform block disks provide acceptable durability.

2.2.2 Sizing the Cloud Accelerator Cache

[Error! Reference source not found.](#) can be used as a rough guide for the appropriate disk technology to use based on object storage and object storage network capability. At the top left, Google Cloud Platform object storage is reachable over the same LAN (for example, within the same Google Cloud Platform region). As we move from top to bottom in the figure, the network capability becomes slower (10 Gbps to 1 Gbps), while the storage capacity requirements increase to store data that is queued up in the accelerator disk cache awaiting transfer to object storage. In the case of slower network-to-object storage, it is more likely that (asynchronous) data ingestion from local client systems can run at a faster rate than cloud transfer. In such a scenario, client ingestion data begins to fill the cache disk location faster than the data can be transferred to object storage and cleared from the cache. As of IBM Spectrum Protect V8.1.2, an internal **throttling** mechanism is in place to slow client ingestion speeds if the cloud accelerator cache disk area begins nearing capacity. However, to avoid slowing client ingestion in cases where ingestion exceeds the cloud transfer rate (which might not be desired), the accelerator cache should be sized with a larger capacity, perhaps up to a single day’s worth of back-end client ingestion (after data deduplication and compression).

⊕ Sizing the Accelerator Cache



Cache storage space needed increases (↑) as network connection or throughput capability of cloud decreases (↓)

Figure 7 : Sizing the cloud accelerator cache for Google Cloud Platform

APPENDIX

Disk Benchmarking

As a part of vetting a Google Cloud Platform test configuration, disk benchmark tests were performed to validate the capability of the disk volumes underlying the IBM Spectrum Protect database and cloud accelerator cache. From a database point of view, this vetting was done to ensure that the volumes were sufficiently capable from an IOPS perspective to support the 8 KiByte random mixed write and read workload that a busy Blueprint-level system would demand. From a cloud cache standpoint, the vetting was performed to ensure that overlapped 128-256 KiByte write and read throughput could achieve a rate high enough such that the server's bottleneck for IO would be at the instance-to-object storage network level and not the disk level. The goal was to ensure that the disk could perform at a rate such that the IBM Spectrum Protect server could utilize it during overlapped ingest and be able to stress the network link layer simultaneously.

Disk benchmarking was performed by using the `tsmdiskperf.pl` Perl script, provided as a part of the Blueprint configuration scripts package found on the IBM Spectrum Protect Blueprints page ([References \[1\]](#)). Execution of the script was performed as follows:

```
perl tsmdiskperf.pl workload=stgpool fslist=directory_list
perl tsmdiskperf.pl workload=db fslist=directory_list
```

With a `stgpool` workload specification, the script drives a 256 KiByte IO pattern, whereas with a `db` workload specification, the script drives 8 KiByte operations. For each directory location provided as a value to the comma-separated `fslist`, a pair of IO processes is created to perform writes and reads to test files that are generated in that directory.

Typical script output for a `stgpool` workload run resembles the following example:

```
=====
: IBM Spectrum Protect disk performance test      (Program version 3.1b)
:
: Workload type:                                stgpool
: Number of filesystems:                        1
: Mode:                                         readwrite
: Files to write per fs:                        5
: File size:                                    2 GB
:
=====
:
: Beginning I/O test.
: The test can take upwards of ten minutes, please be patient ...
: Starting write thread ID: 1 on filesystem /sp/sp_cc/1
: Starting read thread ID: 2 on filesystem /sp/sp_cc/1
: All threads are finished. Stopping iostat process with id 111519
=====
: RESULTS:
: Devices reported on from output:
: dm-2
:
: Average R Throughput (KB/sec):                19473.92
: Average W Throughput (KB/sec):                19377.45
: Avg Combined Throughput (MB/sec):             37.94
: Max Combined Throughput (MB/sec):            160.57
:
: Average IOPS:                                 464.63
: Peak IOPS:                                    2154.57 at 11/10/2017 04:22:32
:
: Total elapsed time (seconds):                  443
=====
```

The value that was extracted for the purposes of comparison and validation for `stgpool` workloads was Avg Combined Throughput (MB/sec). The goal was to determine the largest aggregate average throughput for writes and reads to the accelerator cache disk such that overlapped backup ingest and transfer to object storage will not be constrained by disk capability.

When running the tool in `db` workload mode, output should appear similar to the following example:

```
=====
: IBM Spectrum Protect disk performance test      (Program version 3.1b)
:
: Workload type:                                db
: Number of filesystems:                        1
: Mode:                                         readwrite
: Thread count per FS:                          2
: I/Os per thread:                             500000
: File size:                                   10 GB
:
=====

:
: Creating files of 10 GB to simulate IBM Spectrum Protect DB.
: Issuing command ./ldeedee if=/dev/zero of=/sp/sp_db1/1/tsmdiskperf_1 bs=1024k
count=10240 dio=2 > /dev/null 2>&1
: Issuing command ./ldeedee if=/dev/zero of=/sp/sp_db1/1/tsmdiskperf_2 bs=1024k
count=10240 dio=2 > /dev/null 2>&1
:
: Beginning I/O test.
: The test can take upwards of ten minutes, please be patient ...
: All threads are finished. Stopping iostat process with id 111978
=====

: RESULTS:
: Devices reported on from output:
: dm-6
:
: Average R Throughput (KB/sec):                12907.53
: Average W Throughput (KB/sec):                12707.15
: Avg Combined Throughput (MB/sec):             25.01
: Max Combined Throughput (MB/sec):             42.70
:
: Average IOPS:                                 3202.28
```

```
: Peak IOPS:                               5465.86 at 11/10/2017 04:31:47
:
: Total elapsed time (seconds):             30
```

=====
For the db workload tests, the Avg Combined Throughput (MB/sec) and Average IOPS metrics are significant for evaluating database disk capability. Here, the small random IOPS capability of the underlying disk that is used for the IBM Spectrum Protect Db2 database is of interest.

To conduct measurements of your own, increase the number of write/read threads pairs (and directories) by 1 for each test until the average throughput, the average IOPS, or both stabilize (level off). Benchmark test results are provided here as a reference for those who want to build systems resembling those laid out in this document and who want to validate that their system is capable of supporting the described level of ingestion. For each graph, the horizontal axis represents the quantity of write/read thread pairs (and the number of directory locations used with `fslist`). For each successive bar to the right, the thread count affecting the disk is increased by 2 (1 write thread, 1 read thread, and adding a directory location). The vertical axis represents total average throughput in MiBytes/s.

Object Storage Benchmarking

Another important step in validating the capability of an IBM Spectrum Protect in-the-cloud solution is to benchmark the throughput of the server to the object storage system with a workload that is typical of IBM Spectrum Protect. Ideally, any in-the-cloud IBM Spectrum Protect solution should be network bound in terms of its connection to object storage. Post-inline data deduplication, compression, and encryption, and the back-end ingestion rate over HTTPS should dictate an upper bound for the daily ingestion performance of a system.

To help facilitate this test activity in-house, a Java program was developed by the IBM Spectrum Protect test team to emulate the behavior of the server's use of the Google Cloud Storage API. The tool can be used to drive various backup and restore-type activities to object storage, including direct HTTP PUT, GET, multipart file upload, and range-read restore behavior with a variable number of threads. This tool, known as `SPObjBench.jar`, is included with the Benchmarking package provided with this document.

Also included in the Benchmarking package is a Perl script, `tsmobjperf.pl`, which can be used to automate execution of the `SPObjBench.jar` file with several thread counts to measure ingest (PUT) and restore (GET) scalability.

On the normal ingestion path within the scope of a direct-to-cloud with accelerator cache architecture, the IBM Spectrum Protect server attempts to upload up to 100 1 GB disk container files from the accelerator cache in parallel by using multipart upload. Within a production environment, this work would occur in conjunction with overlapped ingestion to

another set of container files within the same storage pool directory locations on accelerator cache disk storage.

To attempt to isolate and gauge the instance-to-object storage ingestion capability of a system, you can complete the following steps:

1. Populate a set of 10 1 GB files in a memory-mapped file system location to use as source data for ingestion. The use of memory-mapped locations (such as `tmpfs` on Linux) is preferred to eliminate source disk bottlenecks. For a Linux system with at least 11 GB of free RAM, run the following commands:

```
mkdir /mnt/ramdisk

mount -t tmpfs -o size=11g tmpfs /mnt/ramdisk

for I in `seq 10`; do dd if=/dev/urandom
of=/mnt/ramdisk/file.$I bs=1048576 count=1024; done
```

2. To run a set of automated tests scaling from 1 to 100 threads, run the `tsmobjperf.pl` tool by using the recently created RAM disk files as source files to upload. If more threads are specified than files are present in the source list, the tool completes a round-robin action over these source files. Because all activity is read-only, using separate file handles from memory-mapped sources, multiple threads sharing the same file is not a concern. To test with 1, 10, 20, 30, 40, 50, 60, 70, 80, 90, and 100 threads, run the tool as follows, specifying the arguments as needed:

```
perl tsmobjperf.pl type=google gkey=google_key_file
bucket=bucket min=1 max=100 step=10 flist=
comma_delimited_source_files_list
```

where:

- *type* should be `google` for Google Cloud Platform object storage.
- *gkey* should be a Google Cloud Platform key file in JSON format. This key file can be downloaded for a credentialed user from the Google Cloud Platform console UI.
- The *bucket* value should be a Google Cloud Platform object storage bucket name that the credentialed user has create/PUT/GET access to and that exists in the object storage system.
- The *min* and *max* values should indicate the minimum and maximum thread counts to test.
- The *step* value should indicate the increase in thread count from test to test.
- The *flist* parameter should include a comma-delimited list of source files to be used for multipart upload. These files should be the same as those created earlier in the memory-mapped file system.

The following example is for execution of a Google Cloud Platform-based test, using 100 upload threads with an existing test bucket:

```
perl tsmobjperf.pl type=google gkey=jdangar.key.json
bucket=testbucket min=1 max=100 step=10 flist=
/mnt/ramdisk/file.1,/mnt/ramdisk/file.2,/mnt/ramdisk/file.3,/mnt
/ramdisk/file.4,/mnt/ramdisk/file.5,/mnt/ramdisk/file.6,/mnt/ram
disk/file.7,/mnt/ramdisk/file.8,/mnt/ramdisk/file.9
,/mnt/ramdisk/file.10
```

Each thread count test (for 1, 10, 20, or more threads) uploads 10 x 1 GB objects per thread. The previous example would result in a total of 5510 GB of data being stored to the test bucket after all thread tests are completed. The tool does not remove objects that are created. You must remove the objects manually after test completion.

Upon completion, the tool generates aggregate throughput metrics that can be used to estimate practical instance-to-object storage performance rates for IBM Spectrum Protect. Data is provided in comma-separated-value format (CSV) and the output of the SPObjBench.jar tool can be inspected upon completion as well:

```
=====
: IBM Spectrum Protect object storage test
:
: Test Mode:      write
: Type:          google
: Google Key File:  jdangar.key.json: Min Threads:  1
: Max Threads:   100
: Thread Step:   10
: File List:
/mnt/ramdisk/file.1,/mnt/ramdisk/file.2,/mnt/ramdisk/file.3,/mnt/ramdisk/file.4,
/mnt/ramdisk/file.5,/mnt/ramdisk/file.6,/mnt/ramdisk/file.7,/mnt/ramdisk/file.8,
/mnt/ramdisk/file.9 ,/mnt/ramdisk/file.10
: Using java:    java
=====
```

SPObjBench.jar output being captured to file: tsmobjperf.1540336631.out

```
=====
: Test Results
Thread Count, Write Throughput (MB/s), Read Throughput (MB/s)
1, XXX, YYY
10, XXX, YYY
20, XXX, YYY
30, XXX, YYY
40, XXX, YYY
50, XXX, YYY
60, XXX, YYY
```

```
70, XXX, YYY
80, XXX, YYY
90, XXX, YYY
100, XXX, YYY
```

=====
It can be beneficial to monitor network transmission rates externally from the tool, as well, to validate the absolute throughput rate that is experienced to object storage over the (Ethernet) network. The tool reports an aggregate rate that can include build-up and tear-down overhead associated with the tool. Calculating an actual transmission rate from the instance-to-object storage while the test is running can give an indication of the throughput limits of the environment. On Linux, for example, the `dstat` utility can be used to monitor several system metrics at once, including network interface send and receive statistics, by using the basic command:

```
% dstat
You did not select any stats, using -cdngy by default.
----total-cpu-usage---- -dsk/total- -net/total- ---paging-- ---system--
usr sys idl wai hiq siq| read  writ| recv  send|  in   out | int   csw
  0  0 100  0  0  0| 60B 2511B|  0    0 |  0    0 | 76    71
 15  1  84  0  0  1|  0   24k|1674k 58M|  0    0 | 42k 2785
 15  1  83  0  0  1|  0    0 |1838k 62M|  0    0 | 46k 2969
 16  1  82  0  0  1|  0    0 |1832k 61M|  0    0 | 45k 3127
 15  1  84  0  0  1|  0    0 |1753k 61M|  0    0 | 44k 2822
 16  1  83  0  0  1|  0    0 |1811k 62M|  0    0 | 45k 3001
 15  1  83  0  0  1|  0    0 |1778k 62M|  0    0 | 45k 3068
 16  1  82  0  0  1|  0    0 |1870k 63M|  0    0 | 46k 3068
 16  1  82  0  0  1|  0    0 |1933k 64M|  0    0 | 46k 3011
 15  1  83  0  0  1|  0    0 |1834k 63M|  0    0 | 46k 2974
```

The `dstat` tool outputs a new line of metrics at a configured interval, much like the standard `iostat` and `netstat` utilities. For the execution above, the `net/total send` column is of greatest interest, here reported in MiBytes, as an indication of how quickly data could be sent to the object storage endpoint from the server.

Benchmarking Performance Goals

For each IBM Spectrum Protect Blueprint size, run benchmarks to ensure that the following input output operations per second (IOPs) and throughput thresholds and/or ranges are met for the IBM Spectrum Protect database disks and connectivity to object storage:

IBM Spectrum Protect Database Disks

Blueprint Size	Large	Medium	Small	Extra-Small
IOPs	35,000-50,000 8 KB IOPs	20,000-35,000 8 KB IOPs	10,000+ 8 KB IOPs	1,000+ 8 KB IOPs

IBM Spectrum Protect Bandwidth and Throughput to Object Storage

Blueprint Size	Large	Medium	Small	Extra-Small
Throughput	800 - 1000 MB/s	500 – 800 MB/s	250+ MB/s	50+ MB/s

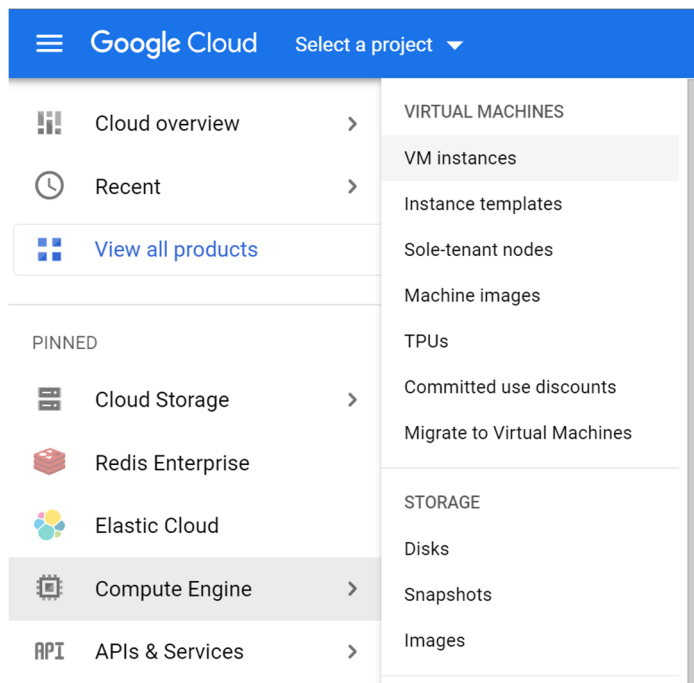
Instance and Object Storage: Navigating the Google Cloud Platform Console User Interface

When deploying a new IBM Spectrum Protect environment based on Google Cloud Platform IaaS, you must navigate to the console web portal and create the relevant cloud resources. Use the following steps as a starting point for creating Google Cloud Platform resources. An important consideration when building an IBM Spectrum Protect server on the Google Cloud Platform is to correctly configure host servers and object storage such that the private (internal Google) network link between these resources is efficient.

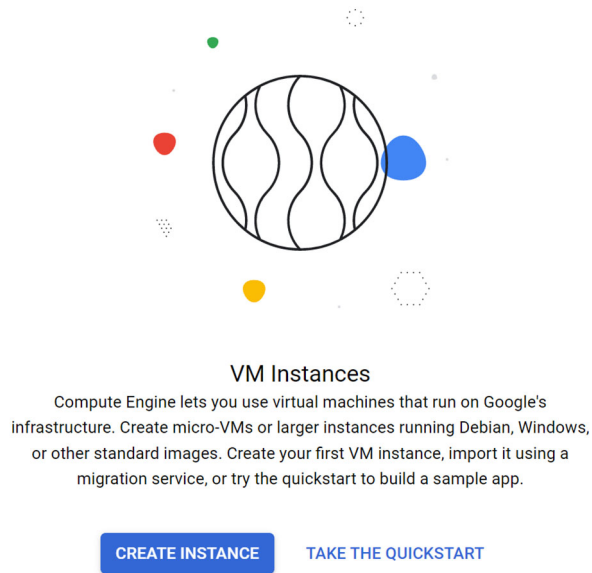
Begin by navigating to the [Google Cloud Platform portal](#) home page and sign in with the appropriate user credentials.

Google Cloud Platform Compute Engine Instances

1. In the navigation pane on the left side of the window, click on the horizontal-lined “hamburger” icon to expand the resource list. Then click on **Compute Engine**, and then **VM instances**. If Compute Engine is not shown, then first click on the **View all products** link.



2. Within the VM instances panel, click the **CREATE INSTANCE** button.



3. From the **Create an instance** wizard, give the instance a **Name**, define an optional set of **Labels**, and select the desired GCP **Region** and **Zone** for the instance. Then select the desired **Machine family Series** and **Machine type** that matches the chosen blueprint size. Optionally, choose to enable a **Display device** and/or enable
-

Confidential Computing,

Name *
production-1 ?

Labels ?

[+ ADD LABELS](#)

Region *
us-central1 (Iowa) ?
Region is permanent

Zone *
us-central1-a ?
Zone is permanent

Machine configuration

Machine family

[GENERAL-PURPOSE](#) [COMPUTE-OPTIMIZED](#) [MEMORY-OPTIMIZED](#) [GPU](#)

Machine types for common workloads, optimized for cost and flexibility

Series
N2 ▼

Powered by Intel Cascade Lake and Ice Lake CPU platforms

Machine type
n2-standard-32 (32 vCPU, 128 GB memory) ▼



vCPU
32

Memory
128 GB

- Click the **CHANGE** button below the **Boot disk** section and select a Linux-based boot disk **Operating system** that is supported by IBM Spectrum Protect. The preferred operating system is the latest supported version of Red Hat Enterprise Linux or Ubuntu Linux. Set the **Size** of the boot disk to match the specified blueprint configuration. Optionally, configure advanced featured under the **SHOW ADVANCED CONFIGURATION** section. Then click the **SELECT** button on the panel to complete

the wizard.

Boot disk

Select an image or snapshot to create a boot disk; or attach an existing disk. Can't find what you're looking for? Explore hundreds of VM solutions in [Marketplace](#)

PUBLIC IMAGES

CUSTOM IMAGES

SNAPSHOTS

ARCHIVE SNAPSHOTS

EXISTING DISKS


Operating system

Red Hat Enterprise Linux

Version *

Red Hat Enterprise Linux 8

x86/64, x86_64 built on 20220822, supports Shielded VM features

License type  *

Select the type of license to use with this OS image.

Pay-as-you-go (PAYG)

Google Cloud charges a license fee for usage of this image. [Pricing](#) varies based on the operating system.

Bring your own subscription or license (BYOS or BYOL)

Google Cloud doesn't charge you a license fee to use this OS image. With this [option](#), you use the license subscription from your OS vendor.

Boot disk type *

Balanced persistent disk

Size (GB) *

100

[SHOW ADVANCED CONFIGURATION](#)

SELECT

CANCEL

5. Expand the **Advanced options** section. Under the **Networking** section, enter a custom **Hostname** of the instance and optional **Network tags**. Optionally add

additional **Network interfaces** if more than one is required.

Advanced options

Networking

Hostname and network interfaces

Network tags 

Hostname 

Set a custom hostname for this instance or leave it default. Choice is permanent

IP forwarding

Enable

Network performance configuration

Network interface card

— 


Network bandwidth

Increase total egress bandwidth

Maximum outbound network bandwidth: 32Gbps

Network interfaces

Network interface is permanent

default default (10.128.0.0/20) 

ADD NETWORK INTERFACE

 To create another network interface you need to have a new network first.

6. Expand the **Disks** section. Click the **ADD NEW DISK** button for each disk specified in the blueprint configuration and configure the disk to the required specifications. Give each disk an informative **Name** and optional **Description**. Choose “Blank disk” for **Disk source type** and choose the appropriate **Disk type** and **Size** for the disk's

purpose. Set **Encryption** settings as needed.

Add new disk ✕

Name *
db-disk-01 ?
Name is permanent

Description
 G

Source
Create a blank disk, apply a bootable disk image, or restore a snapshot of another disk in this project.

Disk source type *
Blank disk ▼

Disk settings

Disk type *
SSD persistent disk ▼ ?

[COMPARE DISK TYPES](#)

Size *
1024 GB ?
Provision between 10 and 65,536 GB

Snapshot schedule (Recommended)
Use snapshot schedules to automate disk backups. [Learn more](#)

Select a snapshot schedule ▼

Encryption
Data is encrypted automatically. Select an encryption key management solution.

Google-managed encryption key
No configuration required

Customer-managed encryption key (CMEK)
Manage via Google Cloud Key Management Service

Customer-supplied encryption key (CSEK)
Manage outside of Google Cloud

Labels ?

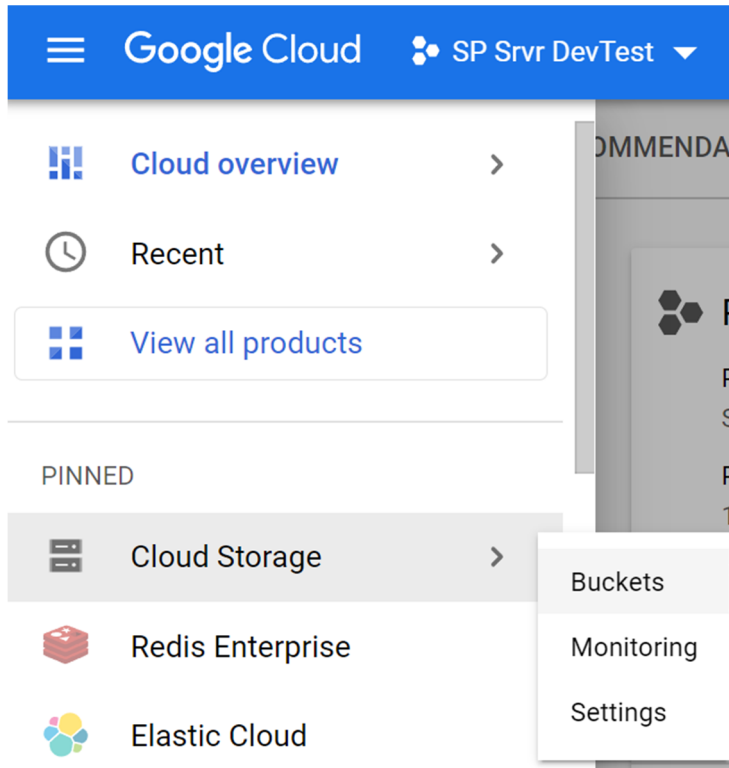
[+ ADD LABEL](#)

SAVE CANCEL


7. After specifying all the needed block storage disks, complete the **Create an instance** wizard by filling in other desired, optional options. These include enabling **deletion protection** of the instance, associating **automation**, specifying **metadata**, and other options. When finished, click the **CREATE** button to create the instance.

Google Cloud Platform Object Storage

1. In the navigation pane on the left side of the window, click on the horizontal-lined “hamburger” icon to expand the resource list. Then click **Cloud Storage**, then **Buckets**. If Cloud Storage is not shown, then first click on the **View all products** link.



2. On the **Buckets** panel, click the **CREATE** button to create a new bucket.



3. In the **Create a bucket** wizard, under the **Name your bucket** section, enter a globally unique **Name** for the new bucket. Optionally, add key/value **Labels** as metadata. Then

click the **CONTINUE** button.

← Create a bucket

✓ **Name your bucket**

Pick a globally unique, permanent name. [Naming guidelines](#)

testbucket-2022

Tip: Don't include any sensitive information

Labels (optional)

Labels are key:value pairs that allow you to group related buckets together or with other Cloud Platform resources. [Learn more](#)

Key 1 Value 1

+ ADD LABEL

^ SHOW LESS

CONTINUE

- Under the **Choose where to store your data** section, select a **Location type** for the bucket. The best-performing location type is generally **Region**. If greater bucket resiliency is desired, choose one of the other Location types. Then click the **CONTINUE** button.

- **Choose where to store your data**

This permanent choice defines the geographic placement of your data and affects cost, performance, and availability. [Learn more](#)

Location type

- Multi-region
Highest availability across largest area

us (multiple regions in United States) ▼

- Dual-region
High availability and low latency across 2 regions
- Region
Lowest latency within a single region


CONTINUE

- Under the **Choose a default storage class for your data** section, choose a storage class. Generally, the **Standard** storage class should be chosen for cloud-container storage pool data. For data that is stored with longer retention – such as periodic full backups or archive data – and for IBM Spectrum Protect database backup to object storage, one of the other storage classes (**Nearline, Coldline, or Archive**) could be
-

chosen. When finished, click the **CONTINUE** button.

- **Choose a default storage class for your data**

A storage class sets costs for storage, retrieval, and operations. Pick a default storage class based on how long you plan to store your data and how often it will be accessed. [Learn more](#)

- Standard** 
Best for short-term storage and frequently accessed data
- Nearline**
Best for backups and data accessed less than once a month
- Coldline**
Best for disaster recovery and data accessed less than once a quarter
- Archive**
Best for long-term digital preservation of data accessed less than once a year

CONTINUE

6. Under the **Choose how to control access to objects** section, optionally select **Enforce public access prevention** on this bucket to harden the security of the bucket. Select either **Uniform** or **Fine-grained Access control** to comply with your desired IAM policy for bucket access. When finished, click the **CONTINUE** button.

- **Choose how to control access to objects**

Prevent public access

Restrict data from being publicly accessible via the internet. Will prevent this bucket from being used for web hosting. [Learn more](#)

- Enforce public access prevention on this bucket

Access control

- Uniform**
Ensure uniform access to all objects in the bucket by using only bucket-level permissions (IAM). This option becomes permanent after 90 days. [Learn more](#)
- Fine-grained**
Specify access to individual objects by using object-level permissions (ACLs) in addition to your bucket-level permissions (IAM). [Learn more](#)

CONTINUE

7. Finally, under the **Choose how to protect object data** section, optionally select **Protection tools**. When finished, click the **CREATE** button to create the bucket.

- **Choose how to protect object data**

Your data is always protected with Cloud Storage but you can also choose from these additional data protection options to prevent data loss. Note that object versioning and retention policies cannot be used together.

Protection tools

- None
- Object versioning (best for data recovery)
For restoring deleted or overwritten objects. To minimize the cost of storing versions, we recommend limiting the number of noncurrent versions per object and scheduling them to expire after a number of days. [Learn more](#)
- Retention policy (best for compliance)
For preventing the deletion or modification of the bucket's objects for a specified minimum duration of time after being uploaded. [Learn more](#)

▼ **DATA ENCRYPTION**

CREATE

CANCEL

REFERENCES

- [1] **IBM Spectrum Protect Blueprints:**
<https://www.ibm.com/support/pages/ibm-spectrum-protect-blueprints>
- [2] **Google Cloud Platform Console:**
<https://console.cloud.google.com/getting-started>
- [3] **Google Cloud Platform, “Changing the machine type of a VM instance”:**
<https://cloud.google.com/compute/docs/instances/changing-machine-type-of-stopped-instance>
- [4] **Backing up a server database to cloud storage:**
<https://www.ibm.com/docs/en/spectrum-protect/8.1.16?topic=storage-backing-up-server-database-cloud>
- [5] **Optimizing database backup operations to cloud object storage:**
<https://www.ibm.com/docs/en/spectrum-protect/8.1.16?topic=storage-optimizing-database-backup-operations-cloud-object>
- [6] **Google Cloud Platform “Cloud Interconnect documentation”:**
<https://cloud.google.com/network-connectivity/docs/interconnect>
- [7] **Google Cloud Platform “Storage classes”:**
<https://cloud.google.com/storage/docs/storage-classes>
- [8] **Overview - IBM Spectrum Protect Supported Operating Systems:**
<http://www.ibm.com/support/docview.wss?uid=swg21243309>
- [9] **Directory-container storage pool FAQs:**
<https://www.ibm.com/support/pages/directory-container-storage-pools-faqs>
- [10] **Cloud-container storage pool FAQs:**
<https://www.ibm.com/support/pages/cloud-container-storage-pools-faqs>
- [11] **Container Storage Pools Best Practices:**
<https://www.ibm.com/support/pages/container-storage-pools-best-practices>
- [12] **Guidelines for caching cloud-container storage pool data to improve the performance of restore operations**

<https://www.ibm.com/docs/en/spectrum-protect/8.1.12?topic=tcspv-guidelines-caching-cloud-container-storage-pool-data-improve-performance-restore-operations>
- [13] **Google Cloud Platform “Quotas & limits”:**
<https://cloud.google.com/storage/quotas>
- [14] **IBM Spectrum Protect object storage support**
<https://www.ibm.com/support/pages/ibm-spectrum-protect-object-storage-support>
- [15] **Securing the IBM Spectrum Protect server**

<https://www.ibm.com/docs/en/spectrum-protect/8.1.16?topic=servers-securing-server>

[16] Security concepts

<https://www.ibm.com/docs/en/spectrum-protect/8.1.16?topic=ss-security-concepts-1>

[17] Authenticating IBM Spectrum Protect users by using an LDAP server

<https://www.ibm.com/docs/en/spectrum-protect/8.1.16?topic=server-authenticating-users-by-using-ldap>

[18] Installing a trusted certificate on an IBM Spectrum Protect server

<https://www.ibm.com/docs/en/spectrum-protect/8.1.16?topic=directory-installing-trusted-certificate-spectrum-protect-server>

[19] SET COMMANDAPPROVAL (Specifies whether command approval is required)

<https://www.ibm.com/docs/en/spectrum-protect/8.1.16?topic=sc-set-commandapproval-specifies-whether-command-approval-is-required>

[20] Processes for managing data protection with IBM Spectrum Protect

<https://www.ibm.com/docs/en/spectrum-protect/8.1.16?topic=overview-data-protection-management-processes>

[21] Managing replication

<https://www.ibm.com/docs/en/spectrum-protect/8.1.16?topic=managing-replication>

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Microsoft is a trademark of Microsoft Corporation in the United States, other countries, or both.

Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

VMware is a registered trademark of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.
