

IBM Spectrum Protect

In-the-Cloud Deployment Guidelines with Microsoft Azure


Document version 1.3

James Damgar

Daniel Benton

Jason Basler

IBM Spectrum Protect Performance Evaluation



© Copyright International Business Machines Corporation 2018, 2020

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

Contents.....	3
List of Figures	4
List of Tables.....	5
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.....	15
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
Microsoft Azure Configurations.....	21
2.1 Design Considerations for Microsoft Azure Instances.....	28
2.1.1 Considerations for Direct-to-Cloud Architectures	30
2.1.2 Sizing the Cloud Accelerator Cache	31
2.1.3 Microsoft Azure: Large Instance Considerations.....	32
2.1.4 Microsoft Azure: Medium and Small Instance Considerations	33
Throughput Measurements and Results.....	33
3.1 Dataset Descriptions.....	33
3.2 Backup and Restore Measurements	35
3.2.1 Microsoft Azure: Large Instance Measurements	35
Appendix	38
Disk Benchmarking.....	38
Object Storage Benchmarking.....	42
Instance and Object Storage: Navigating the Microsoft Azure Portal	46
References.....	57
Notices	58
Trademarks.....	59

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 : Sizing the cloud accelerator cache for Microsoft Azure	32
Figure 5: Microsoft Azure large configuration; database volume average throughput; 8 KiByte random writes/reads.....	41
Figure 6: Microsoft Azure large configuration; database volume average IOPS; 8 KiByte random writes/reads	41
Figure 7: Microsoft Azure large configuration; cloud cache volume average throughput; mixed 256 KiByte writes and reads.....	42

LIST OF TABLES

Table 1: IBM Spectrum Protect physical Blueprint targets (V4.2, Linux x86)	15
Table 2: Preferred ranges of maximum values for client session counts	20
Table 3: Microsoft Azure, large configuration	22
Table 4: Microsoft Azure, alternative large configuration	23
Table 5: Microsoft Azure, medium configuration	24
Table 6: Microsoft Azure, small configuration	26
Table 7: Microsoft Azure, extra-small configuration	27
Table 8: Microsoft Azure instance options	29
Table 9: Throughput measurement datasets	33
Table 10: Microsoft Azure, large configuration, 128 MiByte VE-like dataset backup results	35
Table 11: Microsoft Azure, large configuration, 128 MiByte VE-like dataset restore results	36
Table 12: Microsoft Azure, large configuration, 1 GiByte dataset backup results	36
Table 13: Microsoft Azure, large configuration, 1 GiByte dataset restore results	36
Table 14: Microsoft Azure, large configuration, 128 KiByte dataset backup results	36
Table 15: Microsoft Azure, large configuration, 128 KiByte dataset restore results	37

Introduction



1.1 Purpose of this Paper

This document introduces possibilities for integrating an IBM Spectrum® Protect server with **Microsoft Azure** 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 Microsoft Azure dedicated or shared compute instance with the goal to eventually store the bulk of primary backup and archive data on cost-effective object storage. In the case of Microsoft, this means Azure Blob 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 Microsoft Azure, 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 Microsoft Azure Blob 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.

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](#) depicts the situation after tiering is completed and the `REUSEDELAY` parameter value of the source directory-container storage pool is exceeded (so that deduplicated extent removal for extents with zero reference count can occur).

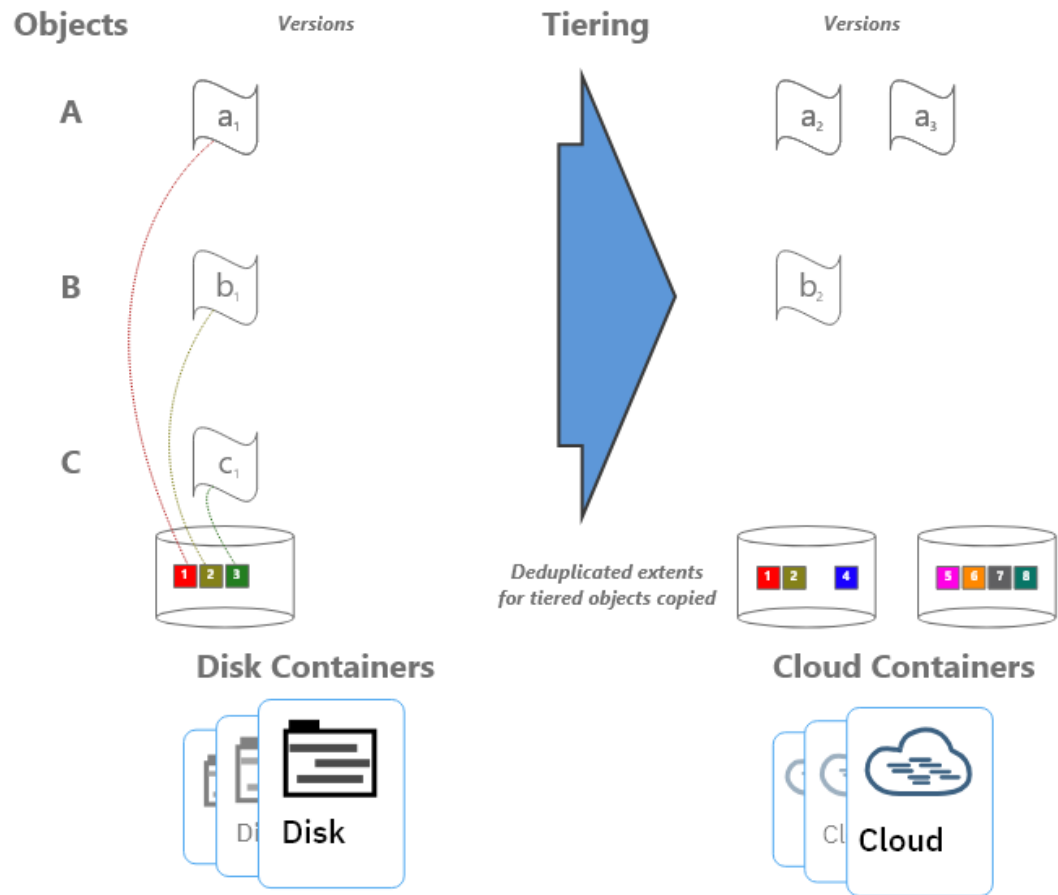


Figure 2: Disk-to-cloud tiering, after tiering

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. Microsoft Azure ExpressRoute can be used to establish a dedicated link ranging from 50 Mbps to 10 Gbps from an on-premises data center to Microsoft Azure private and public resources to facilitate efficient IBM Spectrum Protect replication or backup processing from peer servers or clients outside of the Microsoft Azure infrastructure.

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

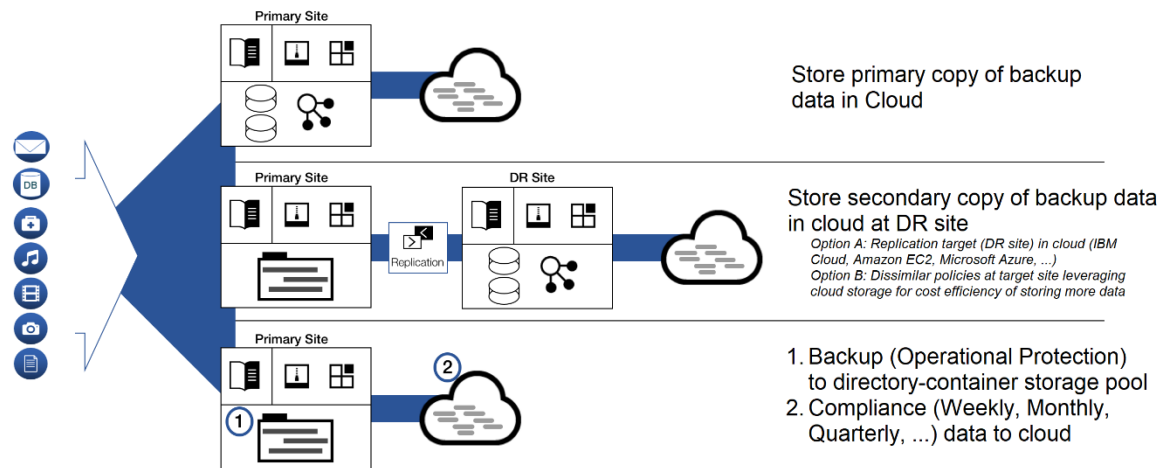


Figure 3: Deployment patterns

In the figure, the first deployment pattern could involve an IBM Spectrum Protect server that is installed on premises or on a Microsoft Azure 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 Microsoft Azure. 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 Azure Blob object storage at the secondary disaster recovery (DR) site. This DR server could be installed at an on-premises site or on a Microsoft Azure 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 Microsoft Azure Blob object storage after a specified number of days. This deployment could also be performed at an on-premises site or within a Microsoft Azure 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 Microsoft Azure, 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. The instances also must have the associated Microsoft Azure Virtual Machine Disks and network (Ethernet) capability to support the desired workload.

With Microsoft Azure 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 Microsoft Azure web portal for most instance types. However, there are some caveats to keep in mind. For more information, see Microsoft'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 Microsoft Azure 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 Microsoft Azure. With a final ingestion point on Azure Blob storage via a cloud-container storage pool, the **Ethernet capability** of the instance and the nature of the network between the instance and the Azure Blob 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.

Certain Microsoft instances might or might not have access to dedicated bandwidth to attached disks (Microsoft managed disks). A lack of access can create a bottleneck in the database operations of the IBM Spectrum Protect server. Certain instances might have limited throughput over Ethernet, and this limitation could hamper ingestion and restore throughput with object storage. 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 Microsoft Azure Blob 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 and, optionally, the IBM Spectrum Protect database archive log disk (when more than one physical disk is utilized). The remaining IBM Spectrum Protect disk components can be satisfied with file systems formatted on directly mounted Microsoft Azure 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 **16 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 Disk Setup Commands for Linux Deployments section of the *IBM Spectrum Protect In-the-Cloud with IBM Cloud* or *IBM Spectrum Protect In-the-Cloud with Amazon Web Services* papers.

The use of Linux LVM to logically stripe across several Microsoft Azure block disks should not be depended upon to increase the durability of the underlying storage, as Microsoft Azure 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 Version 4.2* 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 (V4.2, 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 Microsoft Azure, where disk storage is typically more expensive than object storage. For example, with a large Blueprint cloud system, an approximate monthly savings of 40-50% can be achieved when using 16 TB of Standard,

Hot access tier, LRS Blob (object) storage for database backup purposes compared to statically provisioning 16 TB of Standard class block disk storage.

Another advantage of using Blob object storage for IBM Spectrum Protect database backups is that Blob object storage pricing with Microsoft Azure 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 Blob object storage is that increased data redundancy, availability, and durability can be achieved by using a Blob object storage account with different data redundancy settings. Locally redundant storage (LRS) is the most cost-efficient option, where data is copied synchronously three times within a single physical location in the primary Microsoft Azure region. Zone-redundant storage (ZRS) copies data across three availability zones (data centers) in the primary region synchronously with each write operation and can be used to protect against the outage of a single availability zone. Greater availability and durability can be achieved by using Geo-redundant storage (GRS) or Geo-zone-redundant storage (GZRS) to replicate data from the primary Microsoft Azure region to a secondary region. Both of these options copy data to a single physical location in the secondary region (as with LRS), but differ in how they copy data in the primary region. As with LRS, GRS makes three copies in a single physical location at the primary region while GZRS, as with ZRS, copies data to three availability zones in the primary region. In the case of GRS and GZRS, data is copied to the additional Microsoft Azure region asynchronously with a recovery point objective (RPO) of approximately 15 minutes or less (although with no guaranteed service level agreement, SLA).

You can use the same Microsoft Azure Blob object storage account 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 an availability zone within a Microsoft Azure 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 Microsoft Azure server instance located within the same region. For more information about Microsoft Azure redundancy options, see [References](#) [8]. 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 Blob 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 Microsoft Azure Blob 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 Microsoft Azure Blob 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 Blob 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 V4.2* 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 Microsoft Azure 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 Microsoft Azure Blob object storage.

For an IBM Spectrum Protect server running in Microsoft Azure 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 V4.2** 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. For more information about these data types, see [Throughput Measurements and Results](#).

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

Cloud system size	Large object, poor data deduplication and compression¹	Large object, good data deduplication and compression²	Large object, small extent size, good data deduplication and compression³	Small object, poor data deduplication and 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.

Microsoft Azure Configurations

To access this cloud computing service, go to the [Microsoft Azure](#) portal.

The following configurations represent preferred IBM Spectrum Protect deployments within a Microsoft Azure infrastructure. These deployments make use of Microsoft Azure Blob object storage. Data is ingested into a cloud-container storage pool backed by a **standard, hot access tier, locally redundant storage (LRS) storage account** container in the same Azure region, using the HTTPS protocol.

Blob Storage requires a standard type storage account. The Blob Storage account should be configured as **LRS** both to ensure the most performant object storage behavior and to avoid using the zone-redundant and geo-redundant Azure features that IBM Spectrum Protect does not currently support. A storage account for Blob Storage in the same region as the instance should be chosen to minimize latency and to take advantage of the Microsoft internal network for endpoint communication, where possible.

For circumstances in which the IBM Spectrum Protect server is provisioned within a Microsoft Azure virtual machine compute 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 Microsoft Azure ExpressRoute to establish a dedicated, high bandwidth link from your premises to the Azure-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. Microsoft Azure ExpressRoute offers monthly data plan bandwidth options in the range of 50 Mbps up to 10 Gbps with the additional benefit of inbound data transfer bound for Azure resources being free of charge. See [References](#) [6]. For each of the instances, isolated throughput measurements should be conducted with an Azure storage API tool to determine maximum throughput to collocated Microsoft Azure Blob vault storage within that region.

Table 3: Microsoft Azure, large configuration

Cloud component	Microsoft Azure component	Detailed description	Quantity
Server and network	GS5 Microsoft Azure instance (dedicated)	32-core Intel Xeon CPU E5-2698B v3 @ 2.00 GHz	1
		448 GB RAM	
		896 GB local SSD, 64 maximum attached disks for this instance	
		16 Gbit Ethernet connectivity to Blob Storage	
	Operating system	RHEL or Ubuntu Linux server	1
Block storage	128 GB premium SSD P10 (managed disk)	Operating system disk	1
	128 GB premium SSD P10 (managed disk)	IBM Spectrum Protect instance disk	1
	1024 GB premium SSD P30 (managed disk)	IBM Spectrum Protect database disk	8
	512 GB premium SSD P20 (managed disk)	IBM Spectrum Protect database active log disk	1
	1024 GB standard HDD S30 (managed disk)	IBM Spectrum Protect database archive log disk	4

Cloud component	Microsoft Azure component	Detailed description	Quantity
	1024 GB premium SSD P30 (managed disk)	IBM Spectrum Protect cloud cache disk	8
Object storage	Microsoft Azure Storage Account	IBM Spectrum Protect Azure Storage Account created and used in same Azure Region as instance	1

The following environments were not tested, but can be used as starting points for medium, small, and extra-small Microsoft Azure based instances that could satisfy CPU, memory, ingest throughput, and server database requirements, among other concerns. Also included is an alternative large system that would be **more price-competitive**.

Table 4: Microsoft Azure, alternative large configuration

Cloud component	Microsoft Azure component	Detailed description	Quantity
Server and network	E32s v4 Microsoft Azure instance (dedicated or shared)	32-core Intel Xeon CPU Platinum 8272CL@ 2.50GHz	1
		256 GB RAM	
		32 maximum attached disks for this instance	
		“Extremely high” Ethernet connectivity to Blob Storage	
	Operating system	RHEL or Ubuntu Linux server	1

Cloud component	Microsoft Azure component	Detailed description	Quantity
Block storage	128 GB premium SSD P10 (managed disk)	Operating system disk	1
	128 GB premium SSD P10 (managed disk)	IBM Spectrum Protect instance disk	1
	1024 GB premium SSD P30 (managed disk)	IBM Spectrum Protect database disk	8
	512 GB premium SSD P20 (managed disk)	IBM Spectrum Protect database active log disk	1
	1024 GB standard HDD S30 (managed disk)	IBM Spectrum Protect database archive log disk	4
	1024 GB premium SSD P30 (managed disk)	IBM Spectrum Protect cloud cache disk	8
Object storage	Microsoft Azure Storage Account	IBM Spectrum Protect Azure Storage Account created and used in the same Azure region as the instance	1

Table 5: Microsoft Azure, medium configuration

Cloud component	Microsoft Azure component	Detailed description	Quantity
Server and network	E20s v4 Microsoft Azure instance (dedicated or shared)	20-core Intel Xeon CPU Platinum 8272CL@ 2.50GHz	1

Cloud component	Microsoft Azure component	Detailed description	Quantity
		160 GB RAM	
		32 maximum attached disks for this instance	
		<i>Shared</i> 8-bit Ethernet connectivity to Blob Storage (~4 Gbit Ethernet)	
	Operating system	RHEL or Ubuntu Linux server	1
Block storage	128 GB premium SSD P10 (managed disk)	Operating system disk	1
	128 GB premium SSD P10 (managed disk)	IBM Spectrum Protect instance disk	1
	1024 GB premium SSD P30 (managed disk)	IBM Spectrum Protect database disk	2
	256 GB premium SSD P15 (managed disk)	IBM Spectrum Protect database active log disk	1
	1024 GB standard S30 (managed disk)	IBM Spectrum Protect database archive log disk	3
	1024 GB premium SSD P30 (managed disk)	IBM Spectrum Protect cloud cache disk	6
Object storage	Microsoft Azure Storage Account	IBM Spectrum Protect Azure Storage Account created and used in the same Azure	1

Cloud component	Microsoft Azure component	Detailed description	Quantity
		region as the instance	

Table 6: Microsoft Azure, small configuration

Cloud component	Microsoft Azure component	Detailed description	Quantity
Server and network	E16s v4 Microsoft Azure instance (shared or shared)	16-core Intel Xeon CPU Platinum 8272CL@ 2.50GHz	1
		128 GB RAM	
		32 maximum attached disks for this instance	
		Shared 8 Gbit Ethernet connectivity to Blob Storage (~4 Gbit for Ethernet)	
	Operating system	RHEL or Ubuntu Linux server	1
Block storage	128 GB premium SSD P10 (managed disk)	Operating system disk	1
	128 GB premium SSD P10 (managed disk)	IBM Spectrum Protect instance disk	1
	256 GB premium SSD P15 (managed disk)	IBM Spectrum Protect database disk	4
	128 GB premium SSD P10 (managed disk)	IBM Spectrum Protect database active log disk	1

	1024 GB standard HDD S30 (managed disk)	IBM Spectrum Protect database archive log disk	1
	1024 GB premium SSD P30 (managed disk)	IBM Spectrum Protect cloud cache disk	3
Object storage	Microsoft Azure Storage Account	IBM Spectrum Protect Azure Storage Account created and used in same Azure Region as instance	1

Table 7: Microsoft Azure, extra-small configuration

Cloud component	Microsoft Azure component	Detailed description	Quantity
Server and network	D4s v4 Microsoft Azure instance (shared)	4-core Intel Xeon CPU Platinum 8272CL@ 2.50GHz	1
		16 GB RAM	
8 maximum attached disks for this instance			
Shared 8 Gbit Ethernet connectivity to Blob Storage (~2 Gbit for Ethernet)			
	Operating system	RHEL or Ubuntu Linux server	1
Block storage	64 GB premium SSD P6 (managed disk)	Operating system disk and IBM Spectrum Protect instance disk	1

Cloud component	Microsoft Azure component	Detailed description	Quantity
	64 GB premium SSD P6 (managed disk)	IBM Spectrum Protect database and active log disk	5 (5x file systems created with Linux LVM)
	256 GB standard HDD S20 (managed disk)	IBM Spectrum Protect archive log disk	1
	1024 GB premium SSD P30 (managed disk)	IBM Spectrum Protect cloud cache disk	1
Object storage	Microsoft Azure Storage Account	IBM Spectrum Protect Azure Storage Account created and used in the same Azure Region as the instance	1

2.1 Design Considerations for Microsoft Azure Instances

A range of Microsoft Azure 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.

For **proof-of-concept and test purposes**, the GS5 instance was chosen to validate Microsoft Azure object storage capability at scale with IBM Spectrum Protect. Although this instance type is prohibitively expensive for use as a production IBM Spectrum Protect server, the instance type supported an unconstrained environment in which to set an upper bound for IBM Spectrum Protect workload capability to Azure Blob Storage in a cloud setting. In practice, a **less expensive instance** type could be chosen to satisfy IBM Spectrum Protect server requirements just as well.

The **ESv4-series** of Azure instances provides the vCPU and memory resources necessary to build small, medium, and large Blueprint-level IBM Spectrum Protect servers while also featuring the availability of Azure **premium managed disks**, which have the benefit of **persistently existing** independent of any instance and can be attached or detached on demand to an instance. These disks are necessary for forming IBM Spectrum Protect database, active log, and cloud accelerator cache volumes. The ESv4 line of instances are also economically viable and compare favorably to other instance line candidates, such as the DS line. For these instances, be sure to choose the **“s” variants** to ensure access to

Azure premium storage. These new instance types represent capable and economical choices for a large IBM Spectrum Protect build that were not previously available.

The E16s v4 instance type provides 16 vCPU cores and 128 GB of memory, aligning with a small Blueprint system (with additional RAM). The E20s v4 instance provides 20 vCPU cores and 160 GB of memory, covering the CPU and memory requirements of a medium build. The E20s v4 might be able to support the lower end of a large Blueprint workload as well. In general, the use of any transient storage **should be avoided** for IBM Spectrum Protect storage pool or database data due to the potential for data loss.

For an extra-small Microsoft Azure IBM Spectrum Protect deployment, the D4s v4 instance type represents a low-cost instance with the vCPU, memory, disk, and object storage throughput characteristics that are required for a smaller server instance.

Table 8: Microsoft Azure instance options

Sizing category	Instance line	Instance type	CPU	Memory	Blueprint CPU	Blueprint memory
Small	D2-64s v4	D16s v4	16	64 GB		
	E2-64s v4	E16s v4	16	128 GB	12	64 GB
	DS v2	DS14 v2	16	112 GB		
Medium	D2-64s v4	D16s v4	16	64 GB		
	E2-64s v4	E16s v4	16	128 GB		
	E2-64s v4	E20s v4	20	160 GB	16	128 GB
	DS v2	DS15 v2	20	140 GB		
Large	DS v2	DS15 v2	20	140 GB		
	D2-64s v4	D32s v4	32	128 GB		
	E2-64s v4	E32s v4	32	256 GB		
	D2-64s v4	D64s v4	64	256 GB	32	192 GB
	E2-64s v4	E64s v4	64	432 GB		
	GS	GS5	32	448 GB		

The above table is sorted from **lowest to highest** cost instance in each category. Costs are estimated by using the Microsoft Azure calculator based on a RHEL instance in the West US 2 Region with hourly pricing. As pricing is subject to change, actual figures are not included here.

When possible, enable and use Microsoft Azure **Accelerated Networking** for supported Azure instance types. Instances with Azure Accelerated Networking feature the ability to offload computationally expensive network policy enforcement (such as with network security groups, access control lists, isolation, and other network virtualized serves to

network traffic) to hardware. This permits the bypassing of the Azure virtual switch layer and can help facilitate more efficient Azure Blob storage and IBM Spectrum Protect backup client communication. See [References](#) [7].

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](#)). In the previous table, the EC2 instances were carefully selected to ensure that they have the dedicated EBS bandwidth and disk resources that are necessary for the assigned role. In practice, any planned deployment should be validated to ensure that it will meet performance requirements.

The Microsoft Azure Blob storage account underlying an IBM Spectrum Protect cloud-container storage pool can be configured as a **hot access tier** or as a **cool access tier** storage account. The Azure storage archive access tier is currently not supported by IBM Spectrum Protect. The performance characteristics of the hot and cool storage account types should be identical. Hot access tier storage accounts feature higher storage costs than the cool access tier but with lower access costs. 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 cool access tier Azure blob storage account 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, as IBM Spectrum Protect data deduplication takes place at 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.

2.1.1 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. 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 tebibytes 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 Microsoft Azure standard HDD magnetic 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, Microsoft Azure block disks provide acceptable durability.

2.1.2 Sizing the Cloud Accelerator Cache

[Figure 4](#) 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, Microsoft Azure Blob object storage is reachable over the same LAN (for example, within the same Microsoft Azure 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).

④ Sizing the Accelerator Cache

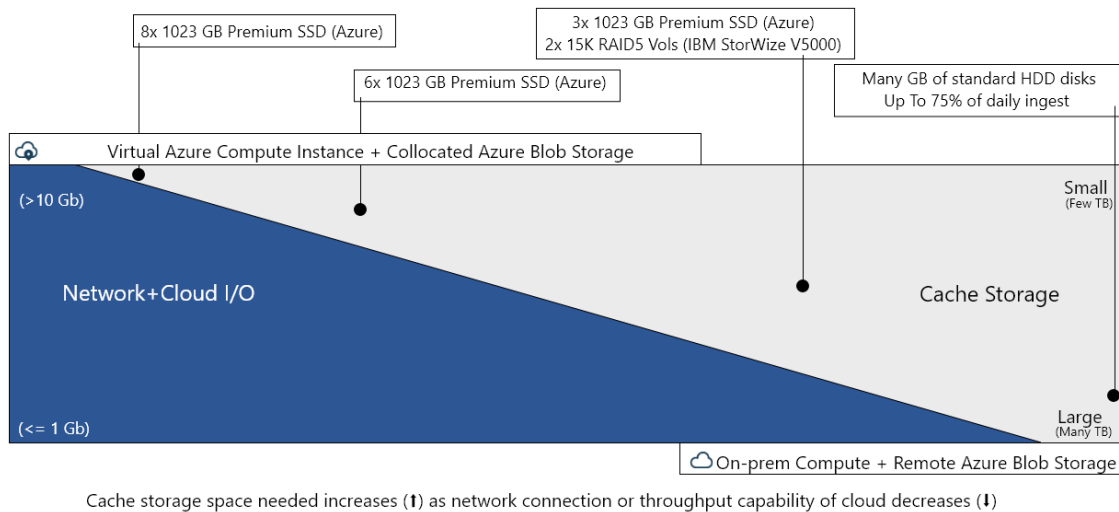


Figure 4 : Sizing the cloud accelerator cache for Microsoft Azure

2.1.3 Microsoft Azure: Large Instance Considerations

For the large Microsoft build, the GS5 instance type was chosen to provide for a completely unconstrained environment in which to stress Azure object storage to capacity with an IBM Spectrum Protect workload from a single instance. This instance type met the specification for a large x86 Linux 64-bit physical system while also providing an excess of memory at 448 GB of RAM. This instance type was primarily chosen for its advertised 16 Gbit Ethernet capability to provide the best link possible to Azure Blob Storage. Whereas a different instance type could be chosen in practice for a large Azure build, the **outlined disk configuration** should be referenced to provide for the disk IOPS and throughput necessary to support a large workload, regardless of the instance type.

For the IBM Spectrum Protect **database volume disks**, eight **1024 GB premium SSD P30 managed disks** were configured. According to the manufacturer, these disks have a 5000 IOPS limit per disk. In total, they provided for 8 TB database capacity and the IOPS necessary to support a large, data-deduplicated workload. A 512 GB premium SSD P20 managed disk was also configured for the **active log disk**. According to the manufacturer, a premium SSD disk of this size has a 2300 IOPS capability. A volume would have to be configured to 1024 GB or larger to reach the next Azure threshold of 5000 IOPS per disk capability.

Azure **standard managed disks** were configured for the server archive log disk. Premium SSD disks are not required for this purpose.

For the **cloud accelerator cache**, eight 1024 GB premium SSD P30 managed disks were configured. Premium SSD disks were required to provide the mixed write/read IOPS capability needed for overlapped client ingestion and disk-to-cloud transfer. Eight disks were the maximum that could be configured in a Linux LVM volume group. This provided for a cache disk volume configuration permitting 8 TB of capacity (while staying within the 32 attached disk limit for the alternative E32s v4 large build).

2.1.4 Microsoft Azure: Medium and Small Instance Considerations

The medium, small, and extra-small Microsoft Azure instances as defined in this document were never built nor tested. However, the specifications in the table above provide guidance for the relative instance types and disk configurations that would be necessary to meet the requirements for ingestion throughput on medium and small systems. Instances must possess the necessary throughput capability to the managed Azure disk layer, including premium SSD disks. Additionally, the instances must have the needed Ethernet network capability to transmit data at an acceptable rate to meet ingestion needs.

Premium SSD managed disks proved to be necessary for the IBM Spectrum Protect **database, active log, and cloud accelerator cache** disk as the only realistic Azure disk type capable of satisfying IOPS and throughput needs for those roles. Managed disks as opposed to unmanaged disks should be chosen in all cases to preserve and persist the state of the IBM Spectrum Protect server data independent of the instance state.

Throughput Measurements and Results

Throughput measurements for backup and restore operations are provided with representative datasets for selected configurations. All throughput measurements were collected with IBM Spectrum Protect servers configured in a direct-to-cloud, accelerator cache-backed setup, with primary backup data stored to object storage. Throughput measurements with a disk-to-cloud tiering configuration are not included in this document.

For each configuration tested, the following preferred settings were adopted for the cloud-container storage pool. To optimize performance when you deploy an IBM Spectrum Protect server instance in the cloud, use the following settings:

- Storage pool compression enabled
- Storage pool encryption enabled
- Storage pool configured as an off-premises cloud-container storage pool
- HTTPS URLs specified for all cloud endpoints

3.1 Dataset Descriptions

For the performance tests that were conducted within these environments, all or a subset of the following datasets were used:

Table 9: Throughput measurement datasets

Front-end object size	Average duplication extent size	Duplicate data percentage	Data compressibility	Notes
128 MiByte	~150-200 KiByte	~70%	~50%	VE-like, favorable extent size

128 MiByte	~200-300 KiByte	~0%	~0%	Random data, large extent size
1 GiByte	~60-100 KiByte	~50%	~50%	DB-like, small extent size
128 KiByte	~128 KiByte	~0%	~95%	Small file, high overhead

The **128 MiByte, VE-like** front-end dataset represents a relatively large object size that aligns with the IBM Spectrum Protect for Virtual Environments: Data Protection for VMware API client's **VE megablock** size for virtual machine disk backups. The large object size and **relatively large, though realistic, deduplication extent** size represents a favorable profile for the IBM Spectrum Protect server's ingestion engine to achieve good performance. A duplicate data rate of 70% combined with a compressibility rate of 50% for this dataset yields an **85% total data reduction** from front-end data as compared with data that is actually stored to the (cloud-accelerator cache and object storage) back-end after data deduplication, compression, and encryption processing. Although this workload does not qualify as a "best case," it does represent a realistic, favorable scenario in which to model top-end throughput capability of an IBM Spectrum Protect system without overestimating throughput performance.

The **128 MiByte, random** front-end dataset represents a larger object size with a large, favorable deduplication extent size. However, the random nature of the data ensures that it does not deduplicate well with existing storage pool data or compress well. This dataset is included to represent a workload that is throughput intensive from the perspective of storage pool disk and object storage network load. Full backups of large objects containing relatively random data content would be modeled well by this dataset.

The **1 GiByte** front-end dataset represents a model of **structured, database-like data** possessing a relatively small deduplication extent size relative to the front-end object size. Such a workload is representative of what might be experienced with an IBM Spectrum Protect for Databases: Data Protection for Oracle backup environment protecting production databases. The smaller extent size causes additional strain and overhead for the IBM Spectrum Protect ingestion engine and typically results in less throughput than the 128 MiByte dataset. A duplicate data rate of 50% and compressibility of 50% yield a 75% overall front-end to back-end reduction for this workload, with a **4:1 ratio** reduction, which approaches what is seen for this type of data in the field.

The **128 KiByte** front-end dataset is used here as a relative "worst case" workload in regard to throughput efficiency for the IBM Spectrum Protect server. This "**small file**" **workload** will stress the data deduplication, compression, encryption, and object storage transfer components of the IBM Spectrum Protect server to a higher degree relative to the amount of data that is actually protected and transferred to object storage. This high overhead dataset allows for predicting a lower estimate on performance of an IBM Spectrum Protect environment within these cloud solutions.

3.2 Backup and Restore Measurements

The following sections outline the backup and restore throughput results that were experienced with the previously mentioned datasets in the built cloud environments.

Prior to conducting backup and restore tests on the IBM Spectrum protect environments, a **load phase** was conducted whereby the servers were initially loaded with a set of deduplicated 128 MiByte front-end data in order to populate the server database tables to provide for a more realistic customer configuration. IBM Spectrum Protect database queries can change their behavior based on the size and layout of server database tables. This load phase was performed to bring behavior in line with real environment expectations.

For each dataset, up to 50 IBM Spectrum Protect client backup sessions were initiated in parallel to the server for the large Microsoft configuration. The results presented here for backup represent the maximum front-end throughput experienced with the largest number of sessions tested against that system.

For each dataset on restore, between 1 and 40 client restore sessions were initiated for the large Microsoft system. Results presented here include the intermediate session count values to give an idea on how restore throughput scales with the number of restore sessions involved for datasets similar to these types.

All throughput values represent front-end, “protected data” values, before inline data deduplication, compression, and encryption. These are the data rates experienced by a **client** that is backing up data to or restoring data from the IBM Spectrum Protect server. The rates are similar to what customers would likely describe as their performance experience with the product. On ingestion, the actual quantity of data that makes it to accelerator cache disk and onwards to object storage will be less, depending on the data deduplication and compression rate. On restore, all individual extents comprising a front-end object will be restored using HTTP GET calls from the object storage device. However, the built-in caching within the IBM Spectrum Protect server’s restore engine might reduce the number of restore operations needed if a workload contains duplicate data.

3.2.1 Microsoft Azure: Large Instance Measurements

Tip: None of the backup or restore throughput rates presented here with the Microsoft Azure large instance build were necessarily vCPU or memory bound. A more cost-effective alternative to the GS5 instance that still provides for roughly 10 Gbit Ethernet connectivity and access to Azure “premium” managed disks should provide for analogous performance to the results seen here with this more expensive instance type.

128 MiByte VE-like front-end dataset results:

Table 10: Microsoft Azure, large configuration, 128 MiByte VE-like dataset backup results

Sessions	Data deduplication %	MiBytes/s	TiBytes/hour	TiBytes per 8 hours	TiBytes per 10 hours
50	70%	1434.5	4.9	39.4	49.3

50	0%	696.9	2.4	19.1	23.9
----	----	-------	-----	------	------

Ingestion throughput for the large Microsoft Azure build with the favorable 128 MiByte dataset was well within the target throughput range of a large Blueprint system (20 – 100 TiBytes per day).

Table 11: Microsoft Azure, large configuration, 128 MiByte VE-like dataset restore results

Sessions	GiBytes/hour
Per session	120
Top-end	483

1 GiByte front-end dataset results:

Table 12: Microsoft Azure, large configuration, 1 GiByte dataset backup results

Sessions	Data deduplication %	MiBytes/s	TiBytes/hour	TiBytes per 8 hours	TiBytes per 10 hours
50	50%	532.0	1.8	14.6	18.3

Ingestion rates for the smaller extent size 1 GiByte object workload at 50 sessions resulted in an aggregate throughput rate that was slightly less than the lower end of the large Blueprint target.

Table 13: Microsoft Azure, large configuration, 1 GiByte dataset restore results

Sessions	GiBytes/hour
Per session	38.9
Top-end	389

128 KiByte front-end dataset results:

Table 14: Microsoft Azure, large configuration, 128 KiByte dataset backup results

Sessions	Data deduplication %	MiBytes/s	TiBytes/hour	TiBytes per 8 hours	TiBytes per 10 hours
50	0%	517.5	1.8	14.2	17.8

Table 15: Microsoft Azure, large configuration, 128 KiByte dataset restore results

Sessions	GiBytes/hour
Per session	66
Top-end	664

Disk Benchmarking

As a part of vetting a Microsoft Azure 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-separate `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.

Microsoft Azure Large Configuration




Figure 5: Microsoft Azure large configuration; database volume average throughput; 8 KiByte random writes/reads




Figure 6: Microsoft Azure large configuration; database volume average IOPS; 8 KiByte random writes/reads




Figure 7: Microsoft Azure large configuration; cloud cache volume average throughput; mixed 256 KiByte writes and reads

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 Microsoft Azure Java 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=type endpoints=endpoint user="user"
pass="pass" bucket=bucket min=1 max=100 step=10 flist=
comma_delimited_source_files_list
```

where:

- *type* should be `azure` for Microsoft Azure Blob object storage.
- *endpoint* specifies a comma-delimited list of IP addresses or URLs for the object storage endpoints. With Microsoft Azure, this value will be a single URL that should be accessed over HTTPS (for security) and that represents the Microsoft Azure provided endpoint URL for the Azure Blob Storage Account.
- For Microsoft Azure, the *user* should be the Azure Account Name.
- For Microsoft Azure, the *pass* should be a SAS token that was configured for the Azure Storage Account in that specific Azure region. This user must have valid Azure credentials to create containers (buckets) and PUT and GET Blob objects in the region indicated by the *endpoint* URL. These values align with those that are used to define an IBM Spectrum Protect cloud-container storage pool, either via the Operations Center or the command line.
- The *bucket* value should be a Microsoft Azure container 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 Microsoft Azure based endpoint in the West US 2 (Oregon) Region, using 100 upload threads with an existing test container:

```
perl tsmobjperf.pl type=azure endpoints=
https://spobjpvthot.blob.core.windows.net/ user="spobjpvthot"
pass="SASTOKENSTRING" bucket=testcontainer min=1 max=100 step=10
flist=
/mnt/ramdisk/file.1,/mnt/ramdisk/file.2,/mnt/ramdisk/file.3,/mnt
/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 container 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:          azure
: Endpoints:     https://spobjpvthot.blob.core.windows.net/
: User:         spobjpvthot
: Pass:         SASTOKENSTRING
: Test Bucket:  testcontainer
: 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	rcv	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.


Instance and Object Storage: Navigating the Microsoft Azure Portal

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

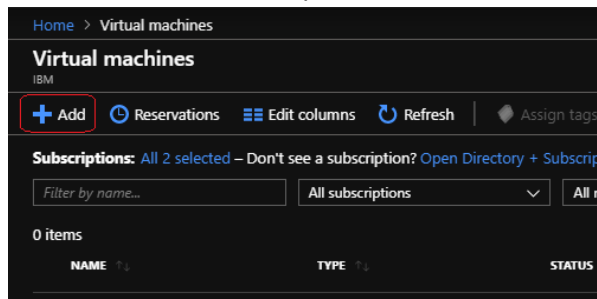
Begin by navigating to the [Microsoft Azure Portal](#) home page and sign in with the appropriate user credentials.

Microsoft Azure Instances

1. In the navigation pane on the left side of the window, either click **Create a resource** and search for **Virtual machines** in the search bar or click **Virtual machines** if it appears in the list.



2. In the **Virtual machines** pane, click **Add**.



3. On the **Basics** tab, in the **Subscription** and **Resource group** sections, specify appropriate settings.
4. In the **Virtual machine name** field, enter a name for the virtual machine.
5. In the **Region** and **Availability options** sections, specify appropriate settings. In the **Image** section, specify an image, for example, Red Hat Enterprise Linux 7.6. Then, click **Select size**.

Basics Disks Networking Management Advanced Tags Review + create

Create a virtual machine that runs Linux or Windows. Select an image from Azure marketplace or use your own customized image. Complete the Basics tab then Review + create to provision a virtual machine with default parameters or review each tab for full customization. [Learn more](#)

Project details

Select the subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.

Subscription *

Resource group * [Create new](#)

Instance details

Virtual machine name *

Region *

Availability options

Image * [Browse all public and private images](#)

Azure Spot instance Yes No

Size * [Select size](#)

- In the **Select a VM size** pane, search for an appropriate virtual machine instance type, for example, e32s for a large Blueprint. Click on the instance type and then click the **Select** button at the bottom.

Select a VM size

Display cost: **Monthly** vCPUs: **All** RAM (GiB): **All** [Add filter](#)

Showing 2 of 348 VM sizes. | Subscription: Microsoft Azure Sponsorship 2 | Region: West US | Current size: Standard_D2s_v3 | Image: Red Hat Enterprise Linux 7.6 | [Learn more about VM sizes](#)

VM Size	Family	vCPUs	RAM (GiB)	Data disks	Max IOPS	Temp storage (GiB)	Premium disk
E-Series v4 The latest generation E family sizes for your high memory needs							
E32s_v4	Memory optimized	32	256	32	51200	0	Supported
E-Series v3 The 3rd generation E family sizes for your high memory needs							

- In the **Username** and **SSH public key** sections, specify settings for initial instance authentication.
- Optionally, specify settings in the **Inbound port rules** section. Click the **Next: Disks** button.

Size * ⓘ Standard_E32s_v4 - (Loading price...) ▼
[Select size](#)

Administrator account

Authentication type ⓘ SSH public key Password

ℹ Azure now automatically generates an SSH key pair for you and allows you to store it for future use. It is a fast, simple, and secure way to connect to your virtual machine.

Username * ⓘ azureuser ✓

SSH public key source Generate new key pair ▼

Key pair name * sp-prod01_key ✓

Inbound port rules

Select which virtual machine network ports are accessible from the public internet. You can specify more limited or granular network access on the Networking tab.

Public inbound ports * ⓘ None Allow selected ports

Select inbound ports * SSH (22) ▼

⚠ This will allow all IP addresses to access your virtual machine. This is only recommended for testing. Use the Advanced controls in the Networking tab to create rules to limit inbound traffic to known IP addresses.

- On the **Disks** tab, add IBM Spectrum Protect premium SSD and standard HDD block disks that are appropriate for the planned cloud Blueprint size. For each disk added, click the **Change size** link to set the appropriate size for the disk.

Warning: Microsoft Azure bills for disk usage based on the value in the **Disk Tier** field. For each disk tier, the maximum size of a disk is listed. If a custom disk size is used that is larger than the size specified, the next disk tier up is used, and usage is billed for this next size. For example, if a 200 GiB disk is chosen, you are billed for a 256 GiB P15 tier disk. Guidance in this paper is to use disks only of the appropriate disk tier size for optimal capacity and billing.

Home > Virtual machines > Create a virtual machine > Create a new disk

Create a new disk

Create a new disk to store applications and data on your VM. Disk pricing varies based on factors including disk size, storage type, and number of transactions. [Learn more about Azure Managed Disks](#)

* Name sp-prod01_DataDisk_0









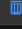



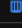



* Source type ⓘ None (empty disk) ▼

* Size ⓘ 1023 GiB
 Premium SSD
[Change size](#)

- When all disks are added for the appropriate cloud Blueprint size, click **Next: Networking**.

Data disks

You can add and configure additional data disks for your virtual machine or attach existing disks. This VM also comes with a temporary disk.

LUN	NAME	SIZE (GIB)	DISK TYPE	HOST CACHING	
0	instance	64	Premium SSD	None	 
1	db01	64	Premium SSD	None	 
2	db02	64	Premium SSD	None	 
3	db03	64	Premium SSD	None	 
4	db04	64	Premium SSD	None	 
5	db05	64	Premium SSD	None	 
6	dbb_archlog	512	Standard H...	Read-only	 
7	cloudcache	1023	Premium SSD	None	 

[Create and attach a new disk](#) [Attach an existing disk](#)

Advanced

[Review + create](#) [< Previous](#) [Next : Networking >](#)

11. On the **Networking** tab, in the **Virtual network** section, create a virtual network or place the virtual machine into an existing network. Select appropriate subnet and public IP values for the virtual machine.

12. Optionally, set public inbound ports and other options. Click **Next: Management**.

Tip: Place the IBM Spectrum Protect server virtual machine in the same virtual network as the client systems that are being protected. In this way, you can help to ensure optimal performance and minimize ingress and egress charges.

Basics Disks **Networking** Management Advanced Tags Review + create

Define network connectivity for your virtual machine by configuring network interface card (NIC) settings. You can control ports, inbound and outbound connectivity with security group rules, or place behind an existing load balancing solution. [Learn more](#)

Network interface

When creating a virtual machine, a network interface will be created for you.

* Virtual network ⓘ (new) spectrumprotecttest-vnet

* Subnet ⓘ (new) default (10.0.0.0/24)

Public IP ⓘ (new) sp-prod01-ip

NIC network security group ⓘ None Basic Advanced

* Public inbound ports ⓘ None Allow selected ports

Select inbound ports

ⓘ All traffic from the internet will be blocked by default. You will be able to change inbound port rules in the VM > Networking page.

Accelerated networking ⓘ On Off

Load balancing

You can place this virtual machine in the backend pool of an existing Azure load balancing solution. [Learn more](#)

Place this virtual machine behind an existing load balancing solution? Yes No

13. On the **Management** tab, add or customize the settings in the **Azure Security Center** and **Azure Active Directory** sections for the virtual machine. Click **Next: Advanced**.
14. On the **Advanced** tab, add any extensions or customizations that you require. Click **Next: Tags**.
15. On the **Tags** tab, add any tags that you require for this virtual machine. Click **Next: Review + create**.
16. On the **Review + create** tab, ensure that the virtual machine passes the validation test. To create the virtual machine, click **Create**.

Basics Disks Networking Management Advanced Tags Review + create

PRODUCT DETAILS

Standard E32s v3
by Microsoft
[Terms of use](#) | [Privacy policy](#)

Subscription credits apply ⓘ
2.2400 USD/hr
[Pricing for other VM sizes](#)

TERMS

By clicking "Create", I (a) agree to the legal terms and privacy statement(s) associated with the Marketplace offering(s) listed above; (b) authorize Microsoft to bill my current payment method for the fees associated with the offering(s), with the same billing frequency as my Azure subscription; and (c) agree that Microsoft may share my contact, usage and transactional information with the provider(s) of the offering(s) for support, billing and other transactional activities. Microsoft does not provide rights for third-party offerings. See the [Azure Marketplace Terms](#) for additional details.

Name

* Preferred e-mail address

* Preferred phone number

Basics

Subscription	Microsoft Azure Sponsorship 2
Resource group	spectrumprotecttest
Virtual machine name	sp-prod01
Region	(US) West US
Availability options	No infrastructure redundancy required
Authentication type	Password
Username	<input type="text"/>
Public inbound ports	None


Disks

OS disk type	Premium SSD
--------------	-------------


[Download a template for automation](#)

Microsoft Azure Blob Object Storage

1. In the navigation pane on the left side of the window, either click **Create a resource** and search for **Storage accounts** in the search bar or click **Storage accounts** if it appears in the list.



2. In the **Storage accounts** pane, click **Add**.



3. On the **Basics** tab, in the **Subscription** and **Resource group** sections, select appropriate settings. Specify settings in the **Storage account name** and **Location** fields. For the performance level, select **Standard**. For the account kind, select **Blob Storage**. For the replication setting, select the appropriate redundancy setting for your Blob storage. The **Locally-redundant storage (LRS)** selection provides the lowest cost option. For the access tier, select **Hot**. Click **Next: Advanced**.

Project details

Select the subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.

* Subscription

* Resource group [Create new](#)

Instance details

The default deployment model is Resource Manager, which supports the latest Azure features. You may choose to deploy using the classic deployment model instead. [Choose classic deployment model](#)

* Storage account name ✓

* Location

Performance Standard Premium

Account kind

Replication
Accounts with the selected kind, replication and performance type only support block and append blobs. Page blobs, file shares, tables, and queues will not be available.


Access tier (default) Cool Hot

[Review + create](#) [< Previous](#) [Next : Advanced >](#)

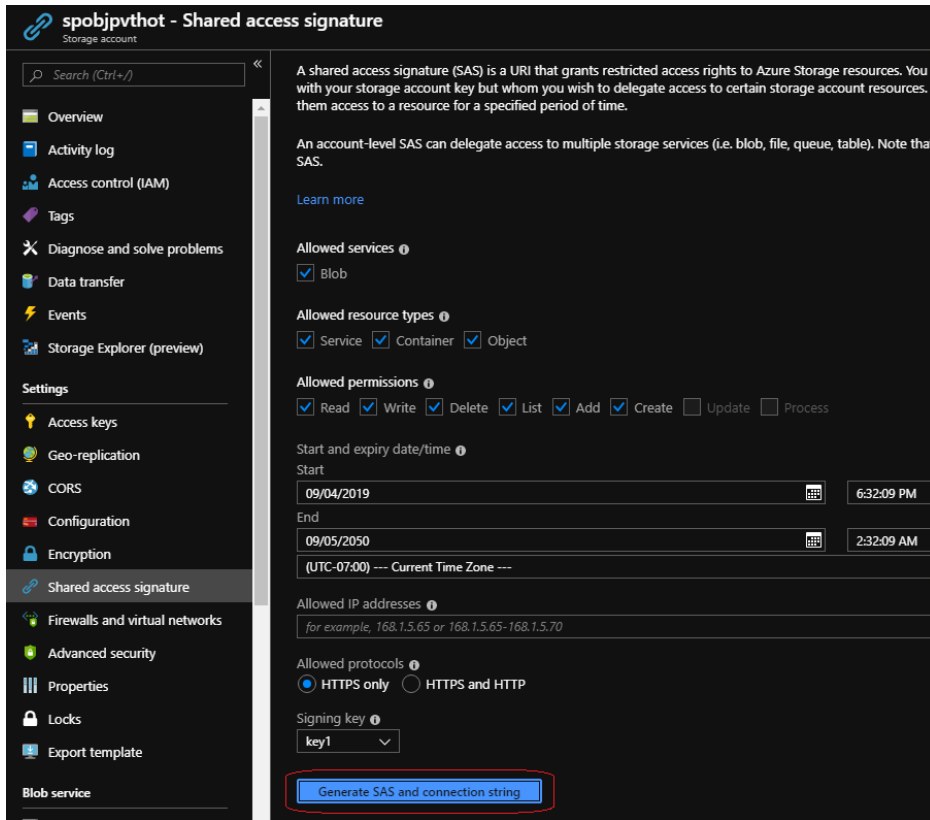
4. On the **Advanced** tab, specify whether to require secure data transfer (HTTPS with TLS) and other options. Click **Next: Tags**.
5. On the **Tags** tab, add any tags that you require for this storage account. Click **Next: Review + create**.
6. On the **Review + create** tab, ensure that the storage account passes the validation test. To create the storage account, click **Create**.

After a Blob storage account is created, a shared access signature (SAS) can be created along with a SAS token to use with IBM Spectrum Protect cloud-container storage pools.

7. To create the SAS and SAS token, click the relevant storage account on the **Storage accounts** pane. Click the **Shared access signature link** in the navigation pane for the Blob storage account.



8. In the **Shared access signature** view, select an appropriate start and end range for the SAS token based on your security needs. After a SAS token expires, another one must be created, and the cloud-container storage pool must be updated with this new token value. Read, Write, Delete, List, Add, and Create permissions are required for cloud container storage pool access. Optionally, restrict access to certain IP addresses and for the HTTPS protocol only. In the **Signing key** section, select an appropriate signing key. Click **Generate SAS and connection string**.



9. Copy the SAS token entry that appears below the **Generate SAS and connection string** button and use this for IBM Spectrum Protect cloud-container storage pools.

REFERENCES

- [1] **IBM Spectrum Protect Blueprints:**
<https://www.ibm.com/support/pages/ibm-spectrum-protect-blueprints>
- [2] **Microsoft Azure:**
<https://portal.azure.com>
- [3] **Microsoft Azure, “Resize virtual machines” topic:**
<https://azure.microsoft.com/en-us/blog/resize-virtual-machines>
- [4] **Backing up a server database to cloud storage:**
http://www.ibm.com/support/knowledgecenter/SSEQVQ_8.1.11/srv.admin/t_backupdb_cloud_object_storage.html
- [5] **Optimizing database backup operations to cloud object storage:**
http://www.ibm.com/support/knowledgecenter/SSEQVQ_8.1.11/perf/t_perf_db_bup_cloud.html
- [6] **Microsoft Azure ExpressRoute:**
<https://azure.microsoft.com/en-us/services/expressroute>
- [7] **Microsoft Azure, “Creating a Windows VM with accelerated networking using Azure PowerShell” topic:**
<https://docs.microsoft.com/en-us/azure/virtual-network/create-vm-accelerated-networking-powershell>
- [8] **Azure Storage redundancy:**
<https://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy>
- [9] **Overview - IBM Spectrum Protect Supported Operating Systems:**
<http://www.ibm.com/support/docview.wss?uid=swg21243309>
- [10] **Directory-container storage pool FAQs:**
<https://www.ibm.com/support/pages/directory-container-storage-pools-faqs>
- [11] **Cloud-container storage pool FAQs:**
<https://www.ibm.com/support/pages/cloud-container-storage-pools-faqs>
- [12] **Container Storage Pools Best Practices:**
<https://www.ibm.com/support/pages/container-storage-pools-best-practices>

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.

Licenses 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.