

IBM Spectrum Protect Plus



Blueprints

- Reference architecture
- Detailed documentation
- Detailed test results
- Published for anyone
- Tool for capacity planning

Blueprints take into consideration your data protection needs in order to tell you:



Type of server
and storage



How much
storage



How to
configure

IBM Spectrum Protect Plus Blueprint

Version 10.1.12

November 2022

Jason Basler (jbasler@us.ibm.com)

Jim Smith (smithjp@us.ibm.com)

[e-mail all authors](#)

This edition applies to Version 10.1.12 and later of IBM Spectrum Protect Plus, and to all subsequent releases and modifications until otherwise indicated in new editions or technical newsletters.

© Copyright IBM Corporation 2022.

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

Table of Contents

Chapter 1. Introduction.....	6
Objective	6
Support for the IBM Spectrum Protect Plus Blueprint	6
What's New	7
Overview	8
Component Overview.....	11
Sample use case – deployment in two active locations	11
Sample use case – deployment in central and branch offices	12
Sample use case – copy to object storage for long-term data storage.....	13
Open Snap Store Manager (OSSM).....	14
Using the IBM Spectrum Protect family blueprints with the Open Snap Store Manager	14
Sites and SLA considerations with the Open Snap Store Manager	15
OSSM proxy agent considerations	15
Task overview	16
Chapter 2. Choosing the Appropriate Technologies.....	17
RAID	17
Compression	17
Deduplication.....	17
Deduplication performance considerations for vSnap servers	18
Additional deduplication information	18
Encryption (at-rest)	19
Physical or virtual vSnap server deployment	20
Disk technology.....	21
Sites and vSnap server distribution	21
Sites and vSnap replication partner relationships	22
SLA design recommendation.....	22
Replication considerations.....	24
Throttling network usage during replication and copy to object storage operations	25
VADP proxy considerations.....	26
Microsoft Hyper-V considerations.....	26

Chapter 3. How to Use the Sizing Tool	27
Introduction to the sizing tool.....	27
Specifying global options	27
Specifying VMware (and general Application) workload options.....	28
Reviewing the results	28
Example sizing exercise	29
Sizing example for the Primary Site	31
Sizing example for the Secondary Site.....	34
Calculating the final sizing.....	35
Sizing the IBM Spectrum Protect Plus server	36
Considerations when rebuilding the IBM Spectrum Protect server from a disaster recovery.....	36
Replication considerations when sizing the IBM Spectrum Protect Plus server.....	37
Sizing other applications	38
Using custom default values	38
Chapter 4. Server and Storage Configuration Blueprints	39
Hardware recommendations for vSnap servers	39
Adding storage for vSnap servers	39
Hardware recommendations for combined vSnap server and VADP proxy	42
Hardware recommendations for dedicated VADP proxy	44
Hardware recommendations for combined VADP and OSSM proxy	44
Blueprint for vSnap server	45
Hardware requirements for physical vSnap server with software defined RAID	45
Example hypervisor for running virtualized IBM Spectrum Protect Plus server, vSnap server, and VADP proxies.....	46
Example storage system to provide storage for vSnap server with storage hardware defined RAID	47
Chapter 5. IBM Spectrum Protect Plus Server.....	49
Adjusting IBM Spectrum Protect Plus server global settings.....	49
Changing the frequency of the Storage Server Inventory job.....	50
Chapter 6. vSnap Server Installation and Setup	51
Configuring a virtual vSnap server using storage hardware defined RAID	51
Configuring a physical vSnap server using storage software provided RAID	56
Configuring global options on vSnap server.....	62
Chapter 7. Configuring VADP Proxies	63
Installing VADP proxies.....	63
Setting the maximum number of VM's to process concurrently	64
Distributing workload to multiple VADP proxies.....	64
Chapter 8. Configuring Cloud Object Storage	66

Default cloud cache area67

Sizing the cloud cache.....67

Recommended cloud cache disk technology.....68

Expanding the cache area if it already exists.....69

Chapter 9. Conclusion 70

Appendix A. vSnap Server Maintenance 71

 Removing a disk71

 Erasing a disk71

 Setting the maximum number of replication streams71

 Setting the maximum number of cloud copy streams71

 Setting the maximum rate for copy to object storage72

 Checking file system integrity on vSnap pools72

Appendix B. Performance Information 74

Appendix C. vSnap Server Performance Test Tool 74

Appendix D. Protecting vSnap System Configuration 77

 Backing up vSnap System Configuration.....77

 Restoring vSnap System Configuration – Example Workflow78

Notices 79

Chapter 1. Introduction

Objective

This document provides guidance on how to build an IBM Spectrum Protect Plus™ solution. This document primarily focuses on how to properly size, build, and place storage components and data movement components for data protection in a VMware vSphere virtual machine environment. The guidance can be generalized for any application protection available with IBM Spectrum Protect Plus.

This guide does not provide instructions or sizing guidance for copying data to the IBM Spectrum Protect server. For information about this feature, see the *IBM Spectrum Protect Plus Installation and User's Guide* "Managing repository server storage"

Support for the IBM Spectrum Protect Plus Blueprint

The information in this document is distributed on an "as is" basis without any warranty that is either expressed or implied. Support assistance for the use of this material is limited to situations where IBM Spectrum Protect Plus support is entitled and where the issues are not specific to a blueprint implementation.

This document is intended to be used with IBM Spectrum Protect Plus Version 10.1.12 and is not applicable to prior releases of the product.

Note: Throughout this document these abbreviations are used: MB to indicate 1024^2 bytes, GB to indicate 1024^3 bytes, and TB to indicate 1024^4 bytes. The sizing tool that is provided in conjunction with this document uses the abbreviations MiB to indicate 1024^2 bytes, GiB to indicate 1024^3 bytes, and TiB to indicate 1024^4 bytes. Even though the abbreviations are used inconsistently between this document and in the sizing tool the units are the same; values can be transposed without regard to this inconsistency.

What's New

The following changes have been made to the IBM Spectrum Protect Plus Blueprints and accompanying sizing spreadsheet in Version 10.1.12 of these documents:

- Open Snap Store Manager (OSSM) sizing information

Overview

The IBM Spectrum Protect Plus solution is provided as a self-contained virtual appliance. While using the appliance as a self-contained solution will be suitable for certain workloads, this blueprint is based on deploying dedicated storage and data movement (proxy) components. Before getting into the details of how to size and build these components, it is necessary to understand the fundamental concepts and building blocks in an IBM Spectrum Protect Plus solution.

- **IBM Spectrum Protect Plus server** – This is the component of the infrastructure that manages and orchestrates the entire system. The server consists of several catalogs that track various system aspects such as recovery points, configuration, access, customizations, etc. A single IBM Spectrum Protect Plus server in a deployment is sufficient in most cases, even if the deployment is spread across multiple locations.
- **Site** – A *site* is an IBM Spectrum Protect policy construct which is used to manage data placement in the environment. A site can be physical (a data center location) or logical (a department or organization). IBM Spectrum Protect Plus components are assigned to sites to localize and optimize data paths. A deployment always has at least one site per physical location. Determining the appropriate number of sites for a deployment will be covered later in this document. The general philosophy is to localize data movement to the sites by placing vSnap or Open Snap Store Manger servers, and VADP proxies together in the sites. The placement of backup data to a site will be governed by the SLA policies.

Note: By default, IBM Spectrum Protect Plus defines two sites: Primary and Secondary.

- **vSnap server** – Also referred to in the interface as a *disk storage*. This is a pool of disk storage that receives data from production systems for the purposes of data protection or re-use. The vSnap server consists of one or more disks and can be scaled up (adding disks to increase capacity) or scaled out (introducing multiple vSnap servers to increase overall performance). There is always at least one vSnap server based on sizing needs. Determining the appropriate number of vSnap servers for a deployment is one of the basic questions addressed by this document.
- **vSnap pool** – Also referred to in this document as a *storage pool*. This is the logical organization of disks into a pool of storage which is consumed by the vSnap server component. This concept is not critical to understand in the scope of architecting a backup solution but will be used later in this document when the vSnap server configuration is discussed in more technical detail.

- **VADP proxy** – This is the component that is responsible for moving data from the vSphere datastores to provide protection for VMware virtual machines (VADP is the acronym for “VMware vSphere Storage APIs - Data Protection”). There is always at least one VADP proxy component for each site based on sizing needs which will be discussed later in this document.

- *Note:* You can install the **vSnap server** (disk storage) and **VADP proxy** on the same physical or virtual system. IBM Spectrum Protect Plus will optimize data movement by eliminating an NFS mount when these two systems are co-located

- **Open Snap Store Manager (OSSM)** – Open Snap Store Manager (OSSM) allows an IBM Spectrum Protect Plus user to store backup copies of VMware vSphere virtual machines directly into IBM Spectrum Protect storage without requiring a vSnap server. The Open Snap Store Manager is designed for users that already have deployed IBM Spectrum Protect directory container storage pool storage and want to use this storage for IBM Spectrum Protect Plus instead of creating new storage attached to vSnap servers.
- **Backup** – This term is used throughout this document to signify the movement of data from a production/host system into the vSnap server for the purposes of data protection.
- **Replication** – This term is used throughout this document to indicate the copying data from a vSnap server to another vSnap server at another site (or to manage replication between two Open Snap Store Manager repositories on two different Spectrum Protect servers) to create redundancy for disaster recovery.
- **Copy to standard object storage** – This term is used to indicate the copying data from a vSnap server to an IBM Spectrum Protect server instance or to cloud object storage to create additional (or unique) copies for offsite and/or long-term retention purposes. You can specify different and longer retentions in SLA policies for data copied to standard object storage to lower overall costs relative to keeping data in the vSnap server repositories.

Note: in previous version the term “offload” was used for this concept. This feature creates a copy of data from the vSnap server into object storage. This is not to be confused with a move or tiering operation.

Note: this feature is not available when using the Open Snap Store Manager repository.

- **Copy to archive object storage** – This term is used to indicate the copying of data from a vSnap server to an IBM Spectrum Protect server that stores the data on physical tape media or in a virtual tape library. Additionally, it can be used to copy snapshots to one of the following archive storage services: Amazon Glacier, IBM Cloud™ Object Storage Archive Tier, or Microsoft Azure Archive.

Note: in previous versions the term “archive” was used for this concept

Note: this feature creates a copy of data from the vSnap server into object storage. This is not to be confused with a move or tiering operation. The differences between copying data to *archive object storage* and *standard object storage* are:

- Copying to *archive object storage* always sends a full copy of the data; Copying to *standard object storage* only sends incremental changes since the last copy operation
- Recovering data from an *archive object storage* requires that the data be staged into an intermediate disk tier prior to accessing the data which affects time to first byte and ultimately recovery time objectives; data stored in *standard object storage* can be recovered directly from the object storage repository.

Component Overview

The following figures outline sample deployments of IBM Spectrum Protect Plus to demonstrate how the various components are deployed.

Sample use case – deployment in two active locations

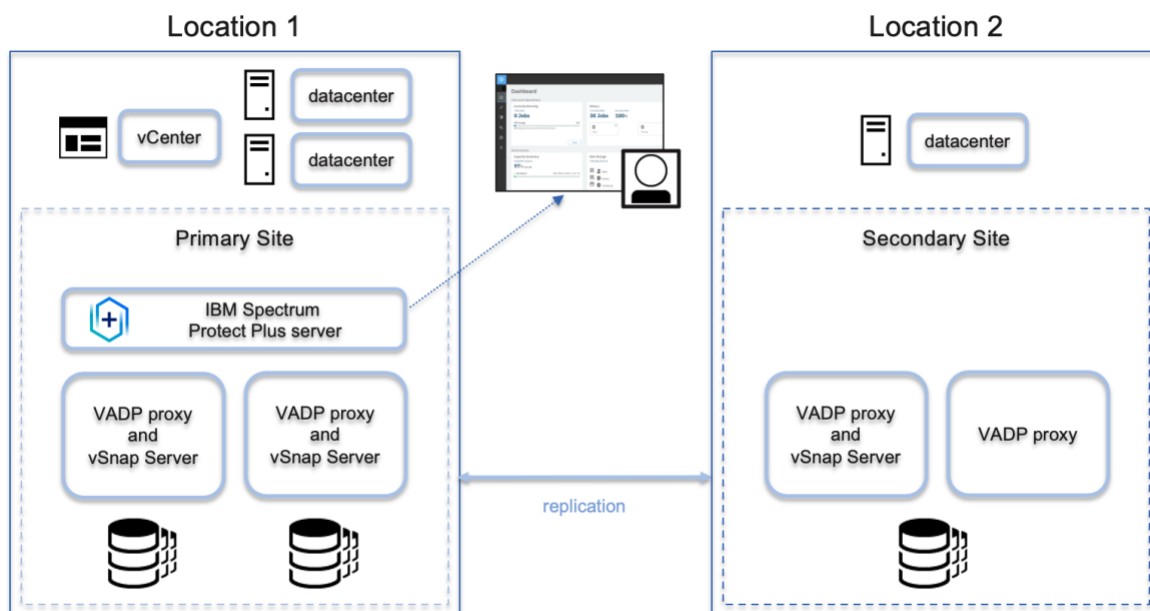


Figure 1 - IBM Spectrum Protect Plus deployed in two active locations

Figure 1 illustrates IBM Spectrum Protect Plus deployed in two active locations. Each location has inventory that needs to be protected. For example, *Location 1* has a vCenter server and two vSphere datacenters (and an inventory of virtual machines) and *Location 2* has a single datacenter (and a smaller inventory of virtual machines in the datacenter).

The IBM Spectrum Protect Plus server is only deployed in one of the sites. VADP proxies and vSnap servers (with their corresponding disks) are deployed in each site to localize data movement in the context of the protected vSphere resources. Note in this example the VADP proxies and vSnap servers have been located on the same systems in both locations and an additional VADP proxy has been installed at *Location 2*. An SLA site designation is given to each site, “Primary Site” and “Secondary Site” so that the backup workloads can be placed in the appropriate locations.

Bi-directional replication is configured to take place between the two sites.

Sample use case – deployment in central and branch offices

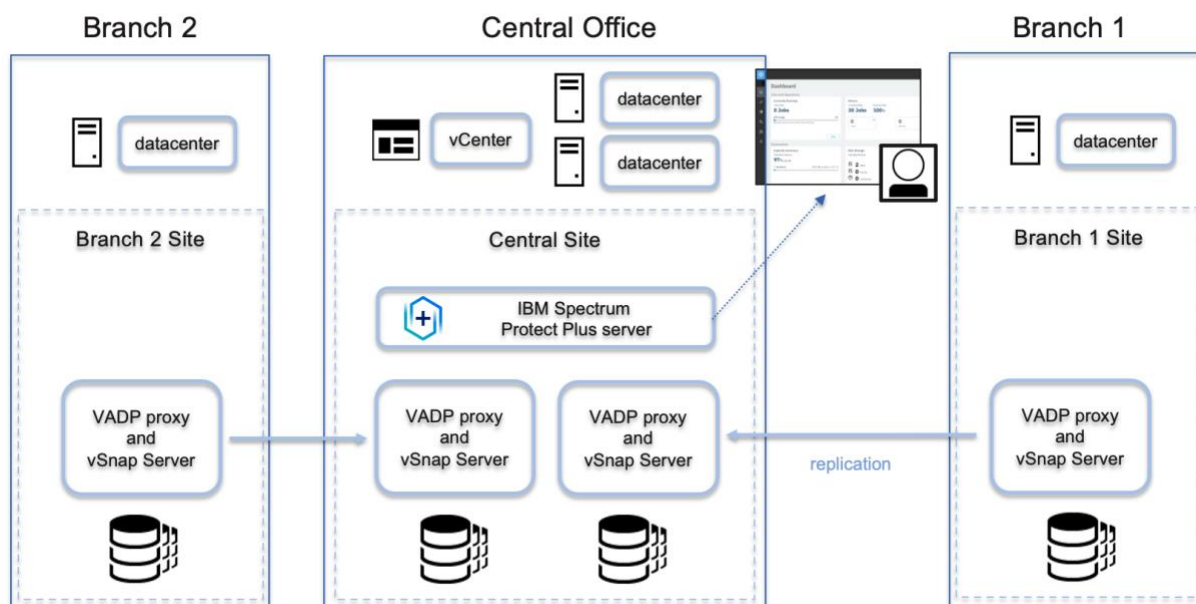


Figure 2 - IBM Spectrum Protect Plus deployed in a central location and two branch locations

Figure 2 illustrates IBM Spectrum Protect Plus deployed in a central location and two supported branch locations. Each location has inventory that needs to be protected. For example, *Central Office* has a vCenter server and two vSphere datacenters (and an inventory of virtual machines) and the branch locations each have a single datacenter (and a smaller inventory of virtual machines).

The IBM Spectrum Protect Plus server is only deployed in the *Central Office*. VADP proxies and vSnap servers (with their corresponding disks) are deployed in each site to localize data movement in the context of the protected vSphere resources. An SLA site designation is given to each site, *Central Site*, *Branch 1 Site*, and *Branch 2 Site*, so that the backup workloads can be placed in the appropriate locations.

Replication is configured so that each of the branch sites is replicating data to the Central Site. *Note:* The backups targeted for the Central Site should also be protected by replication, but this has been omitted to simplify the illustration.

Sample use case – copy to object storage for long-term data storage

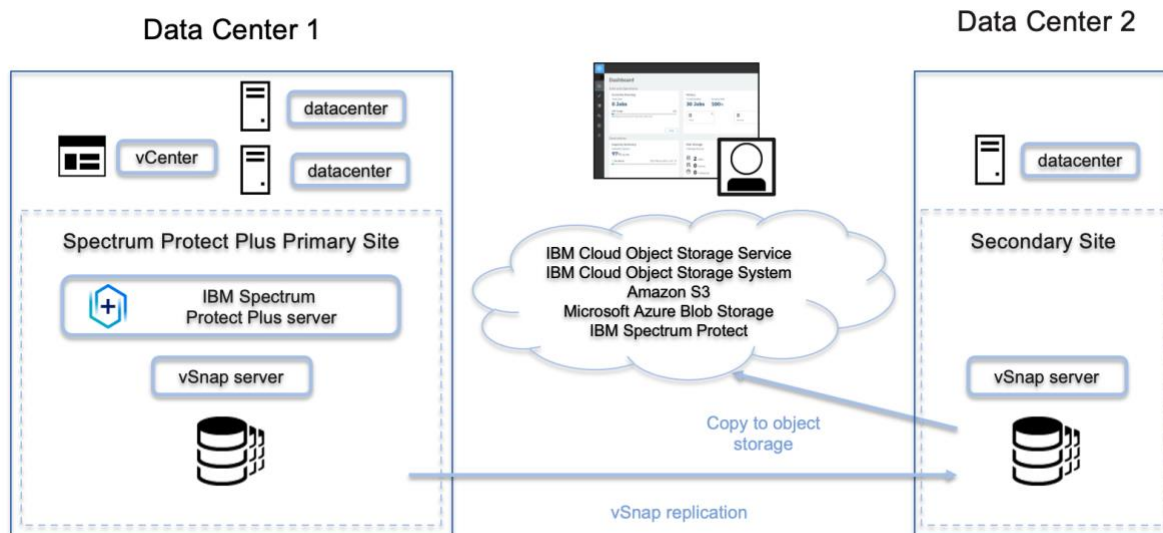


Figure 3 - data copy to cloud object storage or IBM Spectrum Protect

Figure 3 illustrates the data copy capability that is ideal for long-term data storage and data compliance and provides an attractive storage option for disaster recovery and cyber resiliency.

In this figure replication is used to protect the data stored in the Primary Site by copying data from the Primary Site or the Secondary site to cloud object storage targets (IBM Cloud Object Storage, Amazon S3, and Microsoft Azure Blob storage) or to a Spectrum Protect server.

Open Snap Store Manager (OSSM)

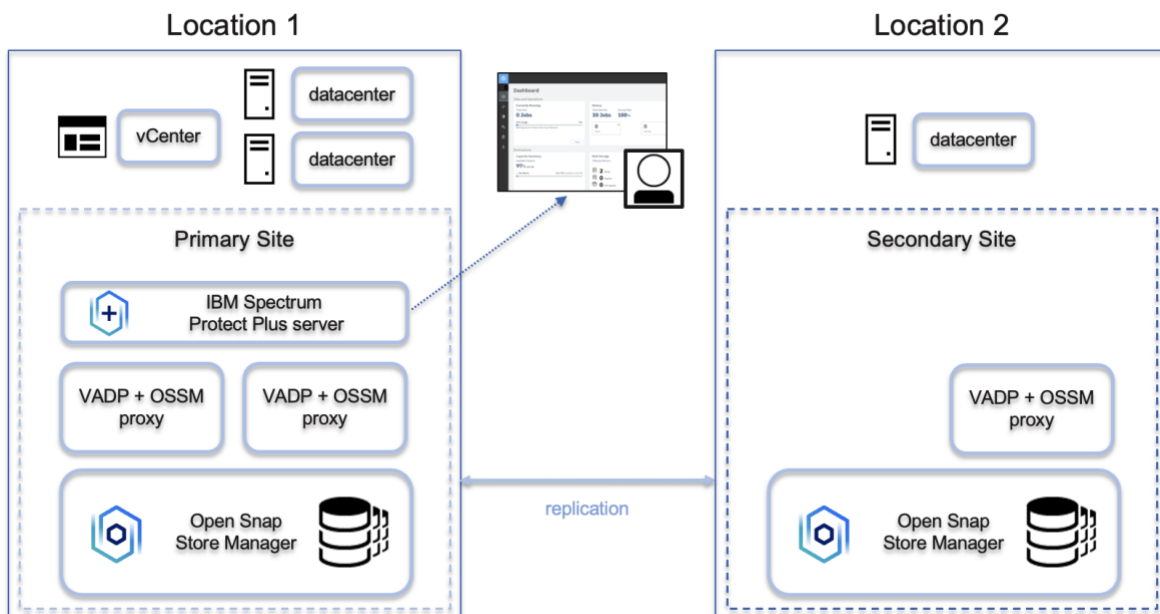


Figure 4 - IBM Spectrum Protect Plus deployed with the Open Snap Store Manager

Open Snap Store Manager (OSSM) allows an IBM Spectrum Protect Plus user to store backup copies of VMware vSphere virtual machines directly into IBM Spectrum Protect storage without requiring a vSnap server. The Open Snap Store Manager is designed for users that already have deployed IBM Spectrum Protect directory container storage pool storage and want to use this storage for IBM Spectrum Protect Plus instead of creating new storage attached to vSnap servers.

Using the IBM Spectrum Protect family blueprints with the Open Snap Store Manager

To size a workload which will utilize the Open Snap Store Manager as the backup storage, you will need to use both the IBM Spectrum Protect Plus and the [IBM Spectrum Protect blueprints](#).

1. The first step is to input the workload inputs into the IBM Spectrum Protect Plus blueprints sizing spreadsheet. This is a similar exercise to creating a sizing with a vSnap server with the following exceptions. If any of these conditions is not met the sizing spreadsheet will not produce a result for the Open Snap Store Manager:
 - a. You can only use VMware vSphere workloads as input
 - b. You cannot specify an SLA with a copy to object storage
 - c. You must specify both compression and deduplication inputs on the Start Here worksheet

2. The Spectrum Protect Plus blueprints sizing spreadsheet will indicate the appropriate [IBM Spectrum Protect blueprints](#) sizing reference of a small, medium, or large Spectrum Protect deployment.

Sites and SLA considerations with the Open Snap Store Manager

When using the Open Snap Store Manager with IBM Spectrum Protect Plus, keep in mind the following when planning SLAs and Sites:

1. Each Site can contain only one Open Snap Store Manager instance. You cannot mix instances of Open Snap Store Managers and vSnap servers within a site.
2. When creating SLAs in IBM Spectrum Protect Plus, you must create separate SLAs for backups targeting Open Snap Store Manager storage and vSnap servers. You cannot create an SLA that can targets both Open Snap Store Manager and a vSnap server.
3. When using Open Snap Store Manager and replication, you can use each Open Snap Store Manager repository as either a backup target or replication target but not both simultaneously.
4. It is not possible to have more than 1 replication job running which has the same Open Snap Store Manager source and target destinations. For smaller virtual machine counts, consider having a single SLA for backup and replication operations.
5. You can reduce the number of proxies by running SLAs at different times. For example, if you have two SLAs with high virtual machine counts you can run the SLAs at separate times in the day if your 24-hour cycle allows for it.

OSSM proxy agent considerations

When using the Open Snap Store Manager with IBM Spectrum Protect Plus, there is a new component called the OSSM proxy agent which orchestrates communication between the VADP proxy and the Open Snap Store Manager storage. These two components (the OSSM proxy and the VADP proxy) are always installed together on the same machine. The IBM Spectrum Protect Plus blueprints sizing will include the number and specifications of these combined proxy machines. For reference hardware recommendations are published in this document in the section [Hardware recommendations for combined VADP and OSSM proxy](#).

Task overview

The document will help you create a suitable IBM Spectrum Protect Plus deployment for your environment by guiding you through the following tasks:

1. Refer to the *IBM Spectrum Protect Plus Installation and User's Guide* chapter "Installation and Setup" for guidance on how to install the IBM Spectrum Protect Plus server. The details of the server installation are not covered in this document as they are well documented in the *User's Guide*. The rest of the steps listed assume that you have deployed the IBM Spectrum Protect Plus server.
2. Now that you understand the basic components and infrastructure of an IBM Spectrum Protect Plus solution, you will need to make decisions about the technology which will directly affect how the solution will be deployed. These choices will include discussion on RAID, data compression and deduplication, and other technology choices. Chapter 2 walks you through the various decision points and advises you on making appropriate choices.
3. The next task is to answer the question "How many components do I need to deploy?" Chapter 3 will show you how to use the sizing tool that is provided in conjunction with this document to appropriately size your environment. This includes sizing not only initial ingest (backup) data but also replication and additional copies of data.
4. Now that you know how many components you need to deploy; Chapter 4 provides templates of systems that you can use to build the solution. These templates are the "blueprints" of the building block you will use to architect your system.
5. Chapter 5 provides additional information on sizing the IBM Spectrum Protect Plus server component.
6. Once the appropriate number of hardware components are deployed, you will need to configure the vSnap server or Open Snap Store Manager components so that they can be used. Chapter 6 walks you through a few different vSnap server configuration scenarios based on the technology choices that were made in Chapter 2.
7. You will need to configure one or more VADP proxies to protect a VMware vSphere environment. Chapter 7 walks you through this final configuration task so that you can begin managing and monitoring IBM Spectrum Protect Plus.
8. Finally, if you are planning on creating backup copies for long-term retention in cloud object storage or IBM Spectrum Protect, Chapter 8 describes how you appropriately size the cloud cache on the vSnap server which is necessary to enable data protection operations to these endpoints.

Chapter 2. Choosing the Appropriate Technologies

Before you start sizing your environment, you will have to understand different technology choices available to you as these decisions will directly affect how your environment is sized and architected. This chapter will help give you guidance on choosing the appropriate technologies to suit your desired outcomes.

RAID

Recommendation: Use RAID 6 technology to protect the vSnap server from disk failures.

Ultimately the decision to use software RAID or storage hardware RAID is up to the user. These are the two options presented in this blueprint:

- If you have pre-existing investment in storage that provides storage hardware RAID such as the IBM® FlashSystem® 5035, you should consider using RAID 6 provided by the storage hardware combined with vSnap RAID 0.
- If storage hardware RAID is not available in the storage system (for example, a JBOD configuration), you should use the vSnap server provisioned RAID 6. Another advantage of this configuration is that the vSnap server provisioned software RAID can detect and correct data corruption.

Compression

Recommendation: Use compression with vSnap Servers and Open Snap Store Manager.

The vSnap server provides a suitable compression algorithm for most workloads. If you have already made the decision to enable compression in the hardware, do not enable compression on the vSnap server.

<i>Note: Throughout this document, all examples are given with vSnap compression enabled.</i>

Deduplication

Recommendation: This is a user decision balancing potential storage savings against the costs of sizing a system that is optimized for data deduplication. The use of data deduplication requires strict adherence to the sizing specifications outlined in this document. It is strongly recommended that you use Open Snap Store Manager storage if you are interested in taking advantage of deduplication.

Deduplication needs to be considered as part of a total data reduction strategy. Not only should you consider whether to use deduplication technologies, but you must also determine the best place in the stack to perform deduplication. Ultimately the decision comes down to weighing the potential storage space savings against the costs of sizing a system that is optimized for data deduplication.

Deduplication performance considerations for vSnap servers

Note: deduplication considerations are specific to vSnap servers. If you are using the Open Snap Store Manager, deduplication is always enabled.

Data deduplication requires additional system resources (primarily CPU and memory) to sustain acceptable backup performance. Some points to consider when evaluating resource requirements for data deduplication:

- The additional memory specified for deduplication is critical for optimizing deduplication extent lookups during backups. If insufficient CPU or memory is allocated, backup performance can be severely degraded.
- Even with adequate resources, backup performance will be slower than if using only compression.
- You can achieve better overall backup performance with deduplication by increasing the number of vSnap server targets. The number of vSnap server targets can be controlled in the sizing tool by modifying the capacity of each vSnap.
- Following a system restart and there may be a period of degraded backup performance while deduplication tables are repopulated into memory.

There are protections built into IBM Spectrum Protect Plus to automatically disable deduplication in cases when deduplication tables become too large for available memory. See [Applying global preferences](#) for more information on the options which control this behavior.

Additional deduplication information

Data deduplication identifies blocks of storage which are identical (redundant) and intelligently manages the backup system by storing only one copy of the redundant block. Data deduplication is most effective if there are large amounts of redundant data in an ecosystem.

From an overall data reduction standpoint, IBM Spectrum Protect Plus uses a full-once, incremental forever backup technology which minimizes the amount of redundant data that the system must process. Contrast this to a backup system which periodically requires a full backup of data that mostly has not changed. In such a backup system, there is naturally more potential for data deduplication as there is more redundant data.

Note: In IBM Spectrum Protect Plus, the scope of data deduplication is within a single vSnap server repository.

While it is hard to predict the amount of potential data redundancy that can be eliminated by data deduplication, consider these general guidelines:

- Virtual machine and virtual desktop environments can have a high rate of data redundancy, especially if the virtual machines are using similar deployments (similar operating system and patch levels). In addition, many virtual machines or virtual

desktops deployed in an environment will usually translate into a higher amount of redundant data

- Unstructured data usually contains more redundant data than structured data. File servers (unstructured) typically show better deduplication results than Oracle or Microsoft SQL servers (structured).

“Pros” of data deduplication

- If there is enough data redundancy in a system, data deduplication can reduce the storage costs more efficiently than other technologies such as compression.

“Cons” of data deduplication

- Performance is slower as compared to systems using compression only.
- The cost of a system optimized for data deduplication in terms of server resources will be higher. See section **Hardware recommendations for vSnap servers** for more information about required system resources.

Encryption (at-rest)

Recommendation: This is a user decision. The most important factor is that you size the vSnap server resources (memory, CPU, etc.) appropriately and set-up sites to appropriately manage encryption in the context of SLA policies.

IBM Spectrum Protect Plus offers the option to encrypt data using AES 256-bit encryption in the vSnap server at-rest. Data ingested during a backup or replication operation to a vSnap server can be encrypted after the data is compressed and/or deduplicated. The data remains encrypted in the vSnap server until it is read for a restore, re-use, or replication operation at which time it will be decrypted.

Note: vSnap disk encryption is based on the dm-crypt disk encryption subsystem utilizing the LUKS extension. The cipher used is aes-xts-plain64 with a 256-bit key size.

Note: When replicating data, the replication process protects the data in-flight, but you will need to ensure that the target vSnap server is also encrypted if you intend to store the replicated data in an encrypted format. When copy data to object storage the data will be protected in-flight, but you will need to configure the object storage system to provide at-rest encryption if you intend to protect the copied data at-rest.

Encryption is predominantly a function of CPU. The sizing information presented in this document and the accompanying sizing tool assume that encryption will be used. Encryption will consume up to 5-10% extra CPU resource depending on the environment.

Encryption is enabled on each vSnap server. If you chose to use encryption globally, ensure that you enable encryption on all vSnap servers in the environment. If you plan to only encrypt some of the data, it is recommended that you group vSnap servers into sites so that you can manage

encryption through SLAs. Refer to the section [Sites and vSnap server distribution](#) for more information on this topic.

When using the Open Snap Store Manager, you can enable encryption on the directory container storage pool using the IBM Spectrum Protect administrative tools. The IBM Spectrum Protect Plus interface will report on the status of encryption but will not allow you to manage encryption settings directly. Also note that the option to “Only use encrypted disk storage” when defining an SLA is not available when defining an Open Snap Store Manager SLA.

Note: At-rest encryption cannot be used in conjunction with IBM Spectrum Scale integration.

Physical or virtual vSnap server deployment

Recommendation: This is a user decision. The most important factor is that you size the vSnap server resources (memory, CPU, etc.) appropriately.

Physical vSnap servers can be preferable because the backup infrastructure is not tied to the virtualization platform, there are no additional virtual licensing costs incurred, and no dependency on the availability of vCenter.

Virtual vSnap servers have the advantage of having a single management interface for both the vSnap servers and virtual infrastructure you are protecting. It is also easy to deploy additional components through deployment of virtual appliances, and you can take advantage of high-availability facilities in the virtual infrastructure. If you use a virtual vSnap server, consider dedicating an ESXi host to manage the vSnap server or servers. For additional security, consider managing the IBM Spectrum Protect Plus virtual infrastructure from a dedicated vCenter or with vCenter permissions restricted to a unique set of users that differs from the users with permissions for the production systems you are protecting.

Operating system maintenance is another factor to consider. If you are building a vSnap server from an existing operating system (either physical or virtual) the operating system maintenance is the responsibility of the user. If you are using the vSnap server virtual appliance, operating system updates are provided as part of the IBM Spectrum Protect Plus software updates.

Note: Software updates for vSnap are provided in the form of a `.run` file which should be applied to both physical and virtual vSnap implementations.

Note: Additional operating system updates for a vSnap virtual appliance can be applied by copying the file `spp_with_os.iso` into the `/tmp` directory of the vSnap system before installing the vSnap `.run` file update.

Note: The VADP proxy software is included in the vSnap virtual appliance for convenience when the virtual appliance will be used for the combined roles of VADP proxy and vSnap server.

If the VADP proxy role is not going to be used, you can uninstall this component or disable the *remote_vadp* service to save system resources.

If you choose to use a virtual vSnap server, you can use pRDM storage, virtual disks, or you can present the storage as an NFS datastore (for example if you have a NAS device or an IBM Elastic Storage Server). If you are using an NFS datastore, the storage must be supported by both VMware and by the vendor for use as an NFS datastore.

If you use virtual disks, it is important to avoid situations which can cause out-of-space conditions in the datastore by taking these steps:

1. Allocate all of the space for the storage at time of creation by using thick disk provisioning. Do not use thin disk provisioning.
2. Do not take vSphere snapshots of the virtual vSnap server. As a safeguard, consider making the virtual machine disk mode as *Independent - Persistent* so snapshots cannot be performed against the disks.
3. Datastores should be dedicated for vSnap servers. You can have multiple vSnap servers on a datastore, but you should not have data from other applications comingled on the datastore with the vSnap server data.
4. Follow VMware recommendations for maintaining sufficient datastore free space. A minimum reserve datastore space of at least 1GB is recommended since you are not taking snapshots in the dedicated vSnap server datastore. See this [VMware article](#) for more information about VMware's minimum free space requirements.

<i>Note:</i> The information presented in this document is based on the use of pRDM storage.
--

Disk technology

Recommendation: In this document we have selected disk technology ideally suited for backup and restore streaming at a low-cost point. For some data re-use scenarios, higher performing disk systems such as flash might be preferable. This can be accomplished by having an SLA policy configured to place data in the higher performing disk system.

Sites and vSnap server distribution

Recommendation:

- Always have at least one site per physical location
- Always have at least one vSnap server per site
- Use the sizing tool that is provided in conjunction with this document to determine the appropriate number of vSnap servers for each site

As noted in the introduction, the general philosophy is to localize data movement to the sites by placing vSnap servers and VADP proxies together in the sites. The placement of backup data to a site will be governed by the SLA policies. In general, you will define a site for each physical location but there might be cases where you want to define multiple sites in a physical location

if you want to achieve the outcomes listed below:

- Maintain data separation (tenancy) between groups such as organizations or departments
- Maintain data separation between primary backup data and data replicated from another vSnap server. For example, if you have two locations and each location uses the other location as a replication target. You may choose to separate the data at each location so that the primary backup data at the location is separated from the replication data at the location.
- Provide different SLAs based on class-of-disk. For example, an SLA for faster backup and recovery based on flash storage and an SLA for standard backup and recovery based on standard spinning disk technology.
- Provide different SLAs based on encrypting data at-rest in the vSnap server repositories. For example, an SLA for encrypted backup storage and an SLA which does not encrypt backup storage.

In each of the cases above it is necessary to have at least two sites to create and manage the appropriate SLA policies.

Sites and vSnap replication partner relationships

Recommendation:

- *Each vSnap server acting as a source for vSnap replication should have a partner relationship with a single vSnap server acting as a target for vSnap replication*

To enable vSnap replication between sites, each vSnap server must have a defined partnership with at least one other vSnap server in the desired replication site. For example, if you are replicating between two sites, and each site includes two vSnap servers, you have to define partner relationships for each of the individual vSnap servers. While IBM Spectrum Protect Plus does not have any limits on how many partner relationships can be defined between vSnap servers, you should design the system so that each vSnap server acting as a replication source only has one defined partner acting as a replication target. You can have a vSnap server acting as a replication target for multiple source vSnap servers assuming that the vSnap target has been appropriately sized for replication.

When configuring replication with the Open Snap Store Manager, you must have two sites with Open Snap Store Manager repositories configured. The partnership must be defined in the IBM Spectrum Protect Plus user interface between the two repositories.

SLA design recommendation

Recommendation:

- *Do not create SLAs with more than 1500 VMs under protection.*

- *If you are protecting more than 1500 VMs consider creating multiple SLAs with no more than 1500 VMs in a single SLA.*
- *Consider the Site which each SLA is targeting. You can have each SLA target a unique Site or you can have multiple SLAs targeting the same site in which case you will need to ensure that enough vSnap servers are deployed within the Site to avoid having more than 1500 VMs stored on any individual vSnap server.*
- *Avoid mapping VMs from the same host and hostcluster groups across multiple SLAs; this will help avoid scenarios where the backup schedules might overlap and compete for resource*
- *Avoid situations where multiple SLAs are executing at the same time within the same Site; this will help avoid resource contention*
- *Use the sizing tool that is provided in conjunction with this document to determine the appropriate number of vSnap servers for each site*

Replication considerations

Replication between the vSnap servers is bi-directional. In other words, a vSnap server that is used to backup data can serve both as a replication source and a replication target. More specifically, in each location, each site acts as a backup target (receiving data from protected applications and having its dedicated VADP proxies), a replication source, and a replication target, receiving replication data from the partner site. Consider the following illustration in which each location has a single site which acts as both a replication source and target.

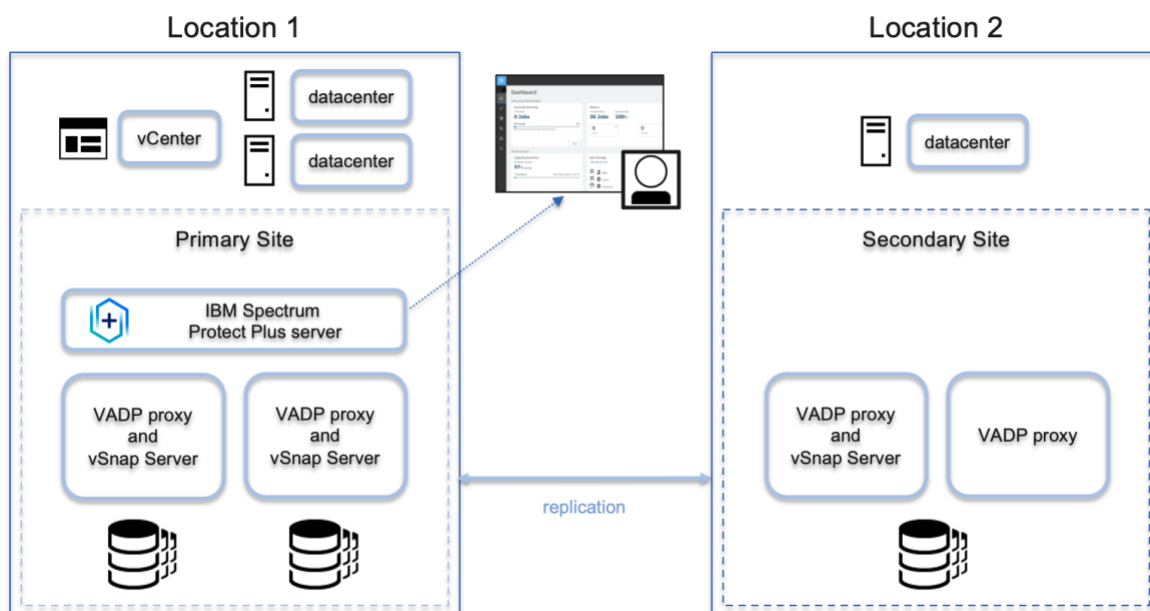


Figure 5 - IBM Spectrum Protect Plus deployed with bi-directional replication between sites

Another option is to dedicate a vSnap server to a single replication function, either a replication source or a replication target with the goal being to more easily manage the separation of data between the vSnap servers. To do this, you must have two sites defined at each location: one site is defined as the backup site (and, therefore, the replication source), and the other site is defined as the replication target. Each site will have a partner site at the second location as illustrated in the following figure.

In *Location 1* there are two sites: The *Primary Backup Site* and the *Primary Rep-Tgt Site*. *Location 2* has a similar configuration.

The *Primary Backup Site* receives backup data for *Location 1* (and therefore has other components such as VADP proxies deployed) and replicates the data to the *Secondary Rep-Tgt Site* at *Location 2* which acts as a replication target. Because the site only receives replication data, there are no other components such as VADP proxies assigned to the site.

Note: Data that is replicated between two vSnap servers is protected by SSH encryption.

Note: The IBM Spectrum Protect Plus server and protected applications have been omitted from this illustration for simplicity.

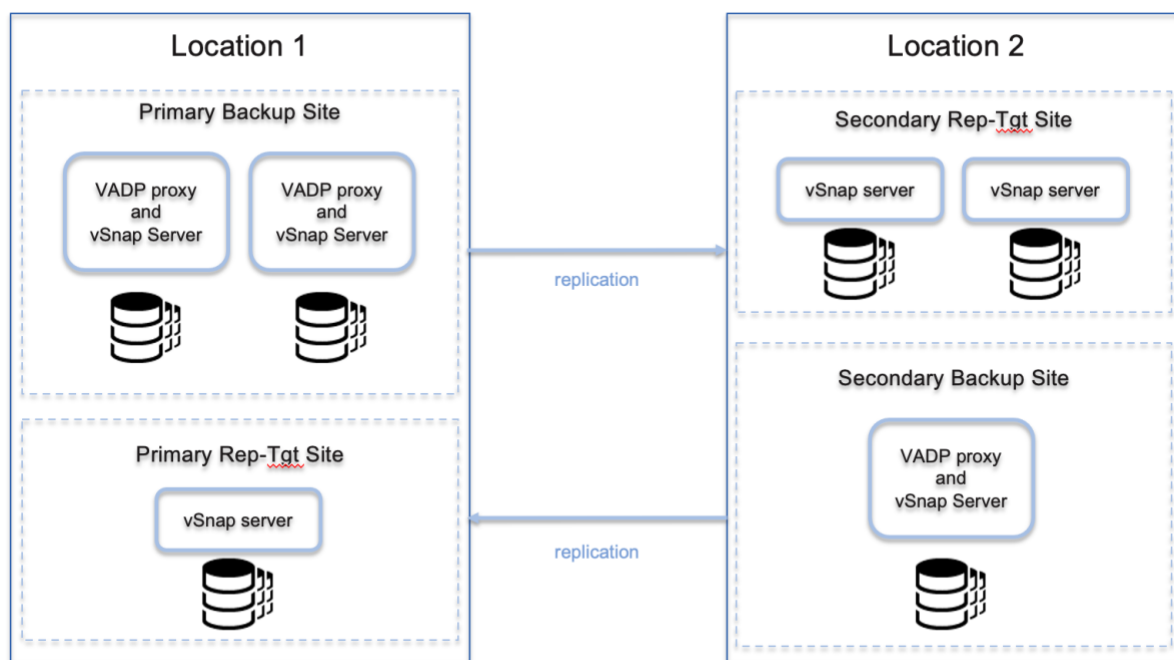


Figure 6 - IBM Spectrum Protect Plus deployed with dedicated replication sites.

There are no recommendations for how to configure sites for replication. Some users will be more comfortable with the simplicity of the first model (each site acting as a replication source and target) while other users will want to have the replication data cleanly separated from the backup data and choose the second model. In either case it will be important to properly size the site for the replication workloads.

Throttling network usage during replication and copy to object storage operations

You can limit the bandwidth usage by using site throttling between vSnap repositories. By throttling workloads, you can help to avoid impact to any critical operations that use the same networks. For more information about defining throttling rates, refer to the IBM Knowledge Center topic [Changing throughput rate](#).

For additional tuning of replication and copy to object storage operations see [Appendix A. vSnap Server Maintenance](#).

VADP proxy considerations

These are some general considerations when planning for VADP proxy deployment.

1. As noted in the introduction, the general philosophy is to localize data movement to the sites by placing VADP proxies in each site that you are trying to protect. Each site should have at least one VADP proxy.
2. Physical vs. virtual proxy? The first question that most users have is whether to deploy a physical or virtual proxy. If the proxy has appropriate resources (memory, CPU, network) there is generally no difference in physical vs. virtual VADP proxies. Usually, it will be easier to deploy a virtual proxy since there is already an existing vSphere deployment. Additionally, the virtual appliance for vSnap includes the VADP proxy software and can be used as a method of deploying VADP proxies even if the vSnap component is not going to be used. Refer to the [IBM Spectrum Protect Plus – All Requirements Doc](#) for more information about VADP system requirements.
3. As noted previously, you can optimize backup performance by collocating the VADP proxy on the same machine as the vSnap server(s) which will eliminate a network hop when writing backup data to the vSnap server. This is supported on Linux distributions which support both the vSnap server and the VADP proxy. If you choose to collocate these two components make sure the system is sized appropriately for both workloads, taking the sum of the CPU, memory, and storage requirements of the two components. When using the Open Snap Store Manager, the OSSM proxy is always collocated with the VADP proxy. There is no option for a stand-alone VADP proxy with the Open Snap Store Manager.
4. If you only have a 1 Gbps network, consider having two adapters on the VADP proxy. One adapter should be dedicated to reading data from the vSphere datastores (transport) and the other should be dedicated to sending the data to the vSnap server.
5. The choice of transport (the data path between the VADP proxy and the vSphere datastore) is based on the VADP proxy type.
 - a. For virtual VADP proxies, the transport can be HotAdd, NBDSSL, or NBD
 - b. For physical VADP proxies, the transport can be SAN, NBDSSL, or NBD

Microsoft Hyper-V considerations

While this document is written with a primary focus on IBM Spectrum Protect Plus deployed in a VMware vSphere environment, the information in this document is also applicable to IBM Spectrum Protect Plus deployed in a Microsoft Hyper-V environment.

1. The main difference between in these environments is that there is no VADP Proxy component in Microsoft Hyper-V environments. This will be reflected in the sizing tool when you size Microsoft Hyper-V workloads.
2. Test results from our lab have shown that performance is much better for a physical vSnap server deployment than a virtual vSnap server deployed as a Microsoft Hyper-V virtual machine.

Chapter 3. How to Use the Sizing Tool

Introduction to the sizing tool

The “IBM Spectrum Protect Plus vSnap Sizer” is a Microsoft Excel spreadsheet which is designed to provide an estimated number and size of vSnap server and VADP Proxies that you will need to deploy to optimally use IBM Spectrum Protect Plus to protect your environment. This chapter demonstrates how to use the sizing tool.

Note: You will have to use this tool to size each site in your environment by creating a separate spreadsheet instance for each site.

The general workflow with the sizing tool will be covered in detail in this chapter. Below is a summary of the tasks to be accomplished:

1. For each site you need to create, you will go through the exercise in the spreadsheet
2. Start by specifying global options for the site which include compression, deduplication, and other estimated global values
3. Fill in the amount of protected data and policy information for each application (the example below will be for VMware vSphere data).
4. Review the results which will indicate how many vSnap servers and VADP proxies you will need to deploy at each site.

Note: The workbook contains macros. You must click *Enable Macros* to use this tool.

Note: Throughout this document these abbreviations are used: MB to indicate 1024^2 bytes, GB to indicate 1024^3 bytes, and TB to indicate 1024^4 bytes. The sizing tool that is provided in conjunction with this document uses the abbreviations MiB to indicate 1024^2 bytes, GiB to indicate 1024^3 bytes, and TiB to indicate 1024^4 bytes. Even though the abbreviations are used inconsistently between this document and in the sizing tool the units are the same; values can be transposed without regard to this inconsistency.

Specifying global options

Start with the “Start Here” sheet to set global options:

1. “vSnap enable compression” - Enable compression based on the choice you made in the previous chapter.
2. “Compression estimate” - Leave this value at “2:1” unless you have site specific information about your compression ratios.
3. “vSnap enable deduplication” - Enable data deduplication based on the choice you made in the previous chapter.

4. “Deduplication estimate” – Leave this value at the default for a vSnap server “1.5:1” unless you have site specific information about your deduplication ratios. If you are using the Open Snap Store Manager change this value to “2:1”
5. “VMware host to VADP proxy network (Gbps)” – Choose the network bandwidth of the network that connects the VMware ESXi hosts to the VADP proxy, either 1, 2, 4, or 10 Gbps.
6. “vSnap network (Gbps)” – Choose the network bandwidth of the network that connects the VADP proxy to the vSnap servers, either 1, 2, 4, or 10 Gbps.

Specifying VMware (and general Application) workload options

Select the “VMware” sheet to provide information about your protected virtual machines.

“Daily change rate” – This is the expected daily change rate of the virtual machine storage. Daily change rates can vary based on workloads. Use the default value unless you have site-specific information about your change rates. Some general guidelines are:

- Larger databases (structured databases) will tend to have lower change rates such as 2.5%.
- Smaller file and web servers will tend to have larger change rates such as 10%.

“Annual growth” – This is the annual expected growth of virtual machine storage. Use the default value unless you have site-specific information about your expected annual storage growth.

For each SLA policy that you intend to use, fill in the estimated front-end capacity of the virtual machines that you will protect with the policy. Fill-in the backup frequency, expected backup duration (backup window), retention, and whether you intend to replicate the data to another vSnap server.

Reviewing the results

Select the “Sizing Results” sheet to view the results of the sizing. Note the results for both the Primary Copy and the Replication Copy. The results will indicate how many vSnap servers and VADP Proxies you will need at each site.

Note: There is now a new “OSSM Sizing Results” sheet that will provide an Open Snap Store Manager sizing information.

Please pay attention to the following points when reviewing results:

1. Validate that the network between the VADP proxies and vSnap server can meet the “Peak vSnap front-end rate”. If the network is not capable of meeting this rate you will need to either increase the backup window, upgrade the network, or use more vSnap servers.

2. The “vSnap Quantity” is based on the “vSnap physical capacity estimate (TiB)” divided by the value of “Preferred vSnap Size (TiB)” on the *Sizing Results* sheet.
3. “VADP quantity” is based on various factors including front-end data rate.
4. The “vSnap quantity” for the replication copy is shown separately. The replication capacity needs to be added to a different spreadsheet for the site that will receive the replication. An example is provided below to illustrate this point

Example sizing exercise

The following example sizing exercise illustrates how to use the tool. Consider the example in the figure below.

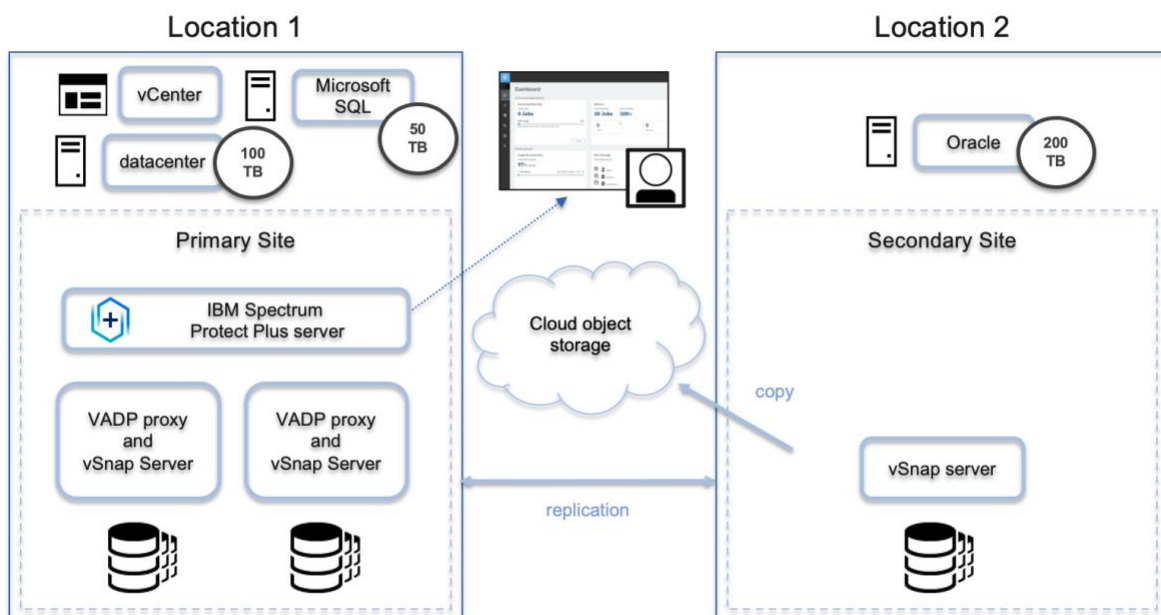


Figure 7 – Example environment

In this example a company has two datacenters in two locations. *Location 1* contains a vSphere vCenter and datacenter with 100 TB of deployed capacity and a Microsoft SQL deployment of 50 TB. *Location 2* has an Oracle database deployment of 200 TB.

The business has provided the following data protection requirements and statistics

- All virtual machine and database data have to be protected once a day with a 30-day retention
- Virtual machines require an additional backup copy with a 180-day retention. It is preferable to use lower cost storage for these additional backup copies.
- All data must be replicated to the non-hosting site (for example, virtual machine data at *Location 1* needs to be replicated to *Location 2*). All replicated data needs to be retained with the same retention as the source data.

- Storage optimization is preferable.
- Daily change rate for VMware vSphere data has been observed to be higher than normal at about 7% with an annual growth rate of 20%.
- Daily change rate for the Oracle databases has been observed to be lower than normal at about 2.5%.
- The data protection ecosystem needs to be sized for three years of protection.

Before using the sizing tool, note some of the following aspects of the system.

- Two sites will be created in the IBM Spectrum Protect Plus server so that the data can be placed appropriately (VMware vSphere data and Microsoft SQL data placed in the *Primary Site* in *Location 1* and Oracle database data placed in the *Secondary Site* at *Location 2*).
- The default *Silver* SLA can be used in both sites to meet the requirement for daily backup with 30-day retention
- The new data copy to cloud object storage feature will be used to create weekly backup copies with 180-day retention. The data copies can be generated from the replicated data in the *Secondary Site*.
- Since storage optimization is preferable, deduplication and compression will be used.
- Since there is no stated requirement for segregation of backup and replication data, replication can be set-up as bi-directional between the sites, in other word each site will act as a replication source and replication target.
- Since there is no native VMware vSphere data to back up at *Location 2*, there will not be a need to deploy VADP proxies at the *Secondary Site*. (Note: even though VMware data will be copied from the replicated copy at the *Secondary Site*, the copy operation does not require a VADP proxy).

To size this environment, you will have to use the tool to calculate the sizing of the components at both the *Primary Site* and the *Secondary Site*, considering that each site acts as a replication target and must consider the replication data that will be received.

Sizing example for the Primary Site

We will start with the *Primary Site* which contains the VMware vSphere and Microsoft SQL data.

1. Save a copy of the sizing spreadsheet named for the *Primary Site* so that you do not accidentally overwrite your worksheet.
2. Select **Enable Macros** to enable macros in the spreadsheet.
3. Select the **Start Here** worksheet.
 - a. Select **Click to reset** the data in the spreadsheet.
 - b. Set the **vSnap enable deduplication** cell to “Yes”.
 - c. Confirm the **Deduplication estimate ratio x:1** is set to “1.5:1”
4. Select the **VMware** worksheet.
 - a. Enter “7%” for **Daily change rate**.
 - b. Enter “20%” for **Annual growth**.
 - c. In the **Silver Policy** column, enter and/or verify the following values in the **Backup frequency** box
 - i. “100” for **Front-end (TiB)**
 - ii. “1” “Day(s)” for **Backup frequency**
 - iii. “8” for **Backup duration hr(s)**
 - iv. “30” for **Retention day(s)**
 - d. In the **Silver Policy** column, enter and/or verify the following values in the **Replication policy** box
 - i. “Yes” for **Replication policy?**
 - ii. “12” for **Daily replication duration hr(s)**
 - iii. “30” for **Replication retention day(s)**
 - e. In the **Silver Policy** column, enter and/or verify the following values in the **Copy to standard object storage policy** box
 - i. “Yes – Repl vSnap” for **Copy to standard object storage policy?**
 - ii. “1” “Week(s)” for **Copy to std. object storage frequency**
 - iii. “8” for **Copy to std. object storage duration hr(s)**
 - iv. “180” for **Copy to std. object storage retention day(s)**
5. Select the **Application-1** worksheet to fill-in the information for Microsoft SQL
 - a. In the **Silver Policy** column, enter and/or verify the following values in the **Backup frequency** box
 - i. “50” for **Front-end (TiB)**
 - ii. “1” “Day(s)” for **Backup frequency**
 - iii. “8” for **Backup duration hr(s)**
 - iv. “30” for **Retention day(s)**
 - v. “Yes” for **Backup logs?**
 - b. In the **Silver Policy** column, enter and/or verify the following values in the **Replication Policy** box
 - i. “Yes” for **Replication policy?**
 - ii. “12” for **Daily replication duration hr(s)**
 - iii. “30” for **Replication retention day(s)**

- c. In the **Silver Policy** column, enter and/or verify the following values in the **Copy to standard object storage policy** box
 - i. “No” for **Copy to standard object storage policy** ?
- 6. Select the **Sizing Results** worksheet.
 - a. For **Preferred vSnap size (TiB)** select “Automatic”

We will now check the sizing results for the *Primary Site* as calculated on the **Sizing Results** worksheet. Since we are sizing for a three-year period, we will be looking at the column labeled **3 year projection**.

1. Note the following sizing results for the **Primary Copy** as you will need them for the final sizing
 - a. The **vSnap quantity** is calculated to 3 x 100 TB vSnap servers with a capacity of 276.3 TiB.
 - b. The **VADP proxy resides on vSnap server?** Is set to “Yes” by default in the **Sizing Results** worksheet. The VADP Quantity is the included in the 3 vSnap servers.
 - c. Note the recommended **VADP Softcap** setting. This value needs to be set on each VADP proxy. See [Configuring VADP Proxies](#) for more information on setting the **VADP Softcap** setting.
 - d. The **Peak vSnap front-end rate** is listed at 4 Gbps. Considering that the vSnap server has a 10 Gbps network, the network should be sufficient for the estimated rates.
2. Note the following sizing results for the **Replication Copy** as you will need them for the final sizing
 - a. The **vSnap quantity** is calculated to 3 x 100 TB vSnap servers with a capacity of 276.3 TiB.
 - b. The **Peak vSnap front-end rate** is listed at 1.3 Gbps.
3. Note the following sizing results for the **Copy to standard object storage copy**
 - a. The **Object storage capacity estimate** represents the estimated amount of cloud object storage (or Spectrum Protect storage) that will be consumed for the long-term copies. For this exercise that number should be 487.5 TiB
 - b. The **Copy to object storage rate** is listed at 5.1 Gbps.
 - c. A cloud cache of 1536 GiB (1.5 TiB) will be required on the vSnap server in the *Secondary Site*.
4. Note the “vSnap reserve %” – This is the amount of reserve space you will plan for your vSnap servers to avoid space constraints. Leave this value at the default which is calculated based on the **vSnap pool size**. In this example the **vSnap reserve %** will default to 15%.

Space needs to be reserved in the vSnap repositories to allow for data re-use scenarios, snapshot management, and for overhead in a newly created pool which is referred to as *slop space*. The amount of reserved slop space will be at least 1/32 of the total space and will be more when software RAID 6 is used. Note the recommended settings for **Target free space error (percentage)** and **Target free space warning (percentage)**. You

will need to configure these values in the **Global Preferences** section in the IBM Spectrum Protect Plus user interface so that the system will take appropriate action as the vSnap server free space approaches the reserve space value:

- a. The **Target free space error (percentage)** should be set to the same value as the **vSnap reserve %**. When the vSnap server free space reaches this value new backup operations will fail.
 - b. The **Target free space warning (percentage)** should be set to 10% less than the **vSnap reserve %** so that the IBM Spectrum Protect Plus administrator can be warned when the vSnap server free space is approaching the reserve space.
5. Validate that the storage system that will be used is capable of meeting the **Peak vSnap back-end rate** for large sequential write workloads.

Sizing example for the Secondary Site

Now we will size the *Secondary Site* which contains the Oracle database data.

1. Save a copy of the sizing spreadsheet named for the *Secondary Site* so that you do not accidentally overwrite your worksheet.
2. Select **Enable Macros** to enable macros in the spreadsheet.
3. Select the **Start Here** worksheet.
 - a. Select **Click to reset** the data in the spreadsheet.
 - b. Set the **vSnap enable deduplication** cell to “Yes”.
 - c. Confirm the **Deduplication estimate ratio x:1** is set to “1.5:1”
4. Select the **Application-1** worksheet to fill-in the information for Oracle
 - a. Enter “2.5%” for **Daily change rate**.
 - b. In the **Silver Policy** column, enter and/or verify the following values in the **Backup frequency** box
 - i. “200” for **Front-end (TiB)**
 - ii. “1” “Day(s)” for **Backup frequency**
 - iii. “8” for **Backup duration hr(s)**
 - iv. “30” for **Retention day(s)**
 - v. “Yes” for **Backup logs?**
 - c. In the **Silver Policy** column, enter and/or verify the following values in the **Replication policy** box
 - i. “Yes” for **Replication policy?**
 - ii. “12” for **Daily replication duration hr(s)**
 - iii. “30” for **Replication retention day(s)**
 - d. In the **Silver Policy** column, enter and/or verify the following values in the **Copy to standard object storage policy** box
 - i. “No” for **Copy to standard object storage policy?**
5. Select the **Sizing Results** worksheet.
 - a. For **Preferred vSnap Size (TiB)** select “Automatic ”

We will now check the sizing results for the *Secondary Site* as calculated on the Sizing Results worksheet. Since we are sizing for a three-year period, we will be looking at the column labeled **3 year projection**.

1. Note the following sizing results for the **Primary Copy** as you will need them for the final sizing
 - a. The **vSnap quantity** is calculated to 3 x 100 TB vSnap servers with a capacity of 264.4 TB.
 - b. The **VADP quantity** is calculated to 0 VADP proxies since there is no VMware vSphere data being protected at this site.
 - a. The **Peak vSnap front-end rate** is listed at 1.8 Gbps. Considering the vSnap servers has a 10 Gbps network, the network should be sufficient for the estimated rates.
2. Note the following sizing results for the **Replication Copy** as you will need them for the final sizing

- a. The **vSnap quantity** is calculated to 3 x 100 TB vSnap servers with a capacity of 264.4 TB.
- b. The **Peak vSnap front-end rate** is listed at 0.6 Gbps. Considering the vSnap servers have a 10 Gbps network, the network should be sufficient for the estimated rates.

Calculating the final sizing

The last step is to calculate the final sizing of the environment based on the individual sizing results for each site.

The *Primary Site* requires 3 vSnap servers (100 TB capacity) to manage a total capacity of 251.1 TB for the primary backup data plus an additional 3 vSnap servers (100 TB capacity) with a total capacity of 264.4 TB for the replicated data received from the *Secondary Site*. It also requires 3 VADP proxies which are collocated on the same systems as the vSnap servers. This is a total of 6 vSnap servers with a total capacity of 540.6 TB capacity. The *Primary Site* will have a peak vSnap front-end rate of 4.0 Gbps for the primary backup data plus 0.6 Gbps for the data replicated from the *Secondary Site*.

The *Secondary Site* requires 3 vSnap servers (100 TB capacity) with a total capacity of 264.4 TB for the primary backup data plus an additional 3 vSnap servers (100 TB capacity) with a total capacity of 251.1 TB for the replicated data received from the *Primary Site*. This is a total of 6 vSnap servers with a total capacity of 540.6 TB capacity. The *Secondary Site* will have a peak vSnap front-end rate of 1.8 Gbps for the primary backup data plus 1.3 Gbps for the data replicated from the *Primary Site*.

The WAN bandwidth required between the sites for replication is 1.9 Gbps (the sum of 0.6 Gbps and 1.3 Gbps).

<i>Site</i>	vSnap Calculated Capacity	vSnap Quantity	VADP Proxies
<i>Primary Site</i>	540.6 TB	6 x 100 TB	3 (These are deployed on the same systems as one of the vSnap servers)
<i>Secondary Site</i>	540.6 TB	6 x 100 TB	0

Table 1: Example sizing results

Sizing the IBM Spectrum Protect Plus server

The sizing spreadsheet also includes the sizing of the IBM Spectrum Protect Plus server, including the compute requirements (CPU and memory) and the storage requirements for the two main catalogs.

You will only use one set of sizing results for the IBM Spectrum Protect Plus server if you are sizing multiple sites. You will only need to use the results from the Primary Site, but you will have to account for replicated workloads that are being managed on other sites. Examples of sizing replicated workloads are included below.

The IBM Spectrum Protect Plus server sizing is based on the number of VMware virtual machines that are being protected. The number of vms is determined by either:

1. The total of the VMware front-end workload divided by the average vm size. The average vm size default is 250 GiB. You can change this by entering a different value in the **Average vm size (GiB)** cell in the **VMware** worksheet.
2. If you know the initial number of vms being protected, you can enter this directly into the **Specify initial vm count instead of using average size** cell in the **VMware** worksheet.

Note the following sizing results for the IBM Spectrum Protect Plus server quantity in the sizing spreadsheet that was created for the *Primary Site*. We will be looking at the column labeled **3 year projection**.

1. The **Spectrum Protect Plus server quantity** is calculated to 1 IBM Spectrum Protect Plus server.
2. The **server CPU cores** is calculated to 8 (this is the default).
3. The **server memory** is calculated to 48 GB (this is the default). The server memory is based on protecting up to 1000 virtual machines and is increased 16 GB for each additional 1000 virtual machines under protection.
4. The **Configuration catalog** is calculated to 50 GiB (this is the default). The configuration is based on protecting up to 1000 virtual machines and is increased 50 GiB for each additional 1000 virtual machines under protection. For information regarding how to expand the configuration catalog please refer to IBM Spectrum Protect Plus documentation.
5. The **Recovery catalog** is calculated to 50 GiB (this is the default).

Considerations when rebuilding the IBM Spectrum Protect server from a disaster recovery

If you are rebuilding the IBM Spectrum Protect Plus server from a disaster recovery you will need to reapply any CPU, memory, and disk capacities to the IBM Spectrum Protect Plus hosting environment before restoring the server catalog.

As an example, imagine that you sized an environment with 2000 virtual machines under protection. The sizing spreadsheet would recommend that you increase the server memory

from 48 GB to 64 GB and the recovery catalog from 50 GiB to 100 GiB. After deploying the properly sized IBM Spectrum Protect Plus server and taking backups of the IBM Spectrum Protect Plus server catalog, you lose your IBM Spectrum Protect Plus server and need to restore it from the catalog backup. When deploying the new IBM Spectrum Protect Plus server, the server memory will be set at the default 48 GB and the recovery catalog will be set at the default 50 GiB. You will need to increase the server memory to 64 GB and the recovery catalog to 100 GiB before you initiate a restore of the IBM Spectrum Protect Plus catalog.

Replication considerations when sizing the IBM Spectrum Protect Plus server

If you are replicating VMware workloads from other sites, you will need to include the number of virtual machines replicated from other sites in order to accurately calculate the size of the IBM Spectrum Protect Plus server.

Note you would not have to do this in the current sizing example because there is no additional virtual machine workload being replicated into the environment; in this example only Oracle data was replicated into the Primary Site, so no additional calculations have to be made. Imagine however that you were replicating in an additional 100.00 TiB of VMware data from the *Secondary Site*. In this case you will need to enter the number of virtual machines under protection in the *Secondary Site* into the sizing spreadsheet for the *Primary Site*.

1. Locate the initial number of virtual machines that were calculated for the sizing exercise from the *Secondary Site*. This number is the **Number of VMs (VMware only)** in the **Sizing Results** worksheet. You will use the number in the **1 year projection** column. Enter this number in the *Primary Site Start Here* worksheet in the **Initial number of VMs (VMware only)** cell. If there are additional sites in the sizing exercise you will have to take the sum of all the virtual machines in all the secondary sites.
2. Locate the **Annual growth rate % (VMware)** in the **Sizing Results** worksheet from the *Secondary Site*. Enter this number in the *Primary Site Start Here* worksheet in the **Replication annual growth rate % (VMware)** cell.

The final sizing for the IBM Spectrum Protect Plus server in the Sizing Results worksheet for the Primary Site will now accurately reflect all of the virtual machines under protection across all of the sites.

Sizing other applications

You can use the Application-1 through Application-4 tabs in the sizing spreadsheet to size most of the applications that are currently supported by IBM Spectrum Protect Plus. Here are some considerations when sizing other applications:

1. IBM Db2, Oracle, Microsoft Exchange, and Microsoft SQL workloads will normally include log backups. Make sure to select “Yes” for **Backup logs?** when sizing these applications. The other applications (MongoDB, Microsoft 365, Microsoft Windows, Kubernetes, and OpenShift) do not have log backups.
2. Kubernetes and OpenShift protection based on block volume storage do not employ incremental backups when storing data in the vSnap repository. When sizing Kubernetes or OpenShift workloads set **Daily change rate** to “100%” if using block volume storage. If you are using file system storage you can take advantage of incremental backups and set the **Daily change rate** appropriately.
3. Amazon EC2 workloads do not store data in the vSnap repository. Do not use the sizing spreadsheet for Amazon EC2 workloads.
4. For more information about sizing Microsoft Hyper-V workloads see **Microsoft Hyper-V considerations**.
5. Each Application tab in the sizing spreadsheet has an **Application name label**; this is only for convenience if you are sizing different workloads and doesn’t not affect the sizing

Using custom default values

You can use custom default values in the sizing spreadsheet to change the behavior of the **Click to reset** button on the **Start Here** sheet. This can be helpful if you want to maintain consistency and eliminate potential errors when using the sizing spreadsheet multiple times. For example, the customer might have a policy to always use encryption. To ensure that encryption is always selected, you can change the **Click to reset** button so that the **Enable vSnap encryption** setting is always reset to “yes”. To use custom default values:

1. Ensure that your workbook allows you to show the **Developer** tab. You can search the web for “EXCEL SHOW DEVELOPER TAB” for instructions on how to do this.
2. From the **Developer** tab, select the **Macro** icon and select the **Unhide_Defaults** macro.
3. Navigate to the **Defaults** worksheet. Information on default values that are customizable is included in the worksheet.

Chapter 4. Server and Storage Configuration Blueprints

This chapter examines the hardware references that have been tested in the lab. You might be able to substitute hardware from other vendors with equivalent specifications.

Note: IBM does not support running any other software packages on the same machine that is hosting a vSnap server. This includes but is not limited to anti-virus tools, reporting tools, security hardening tools, or co-hosting a Spectrum Protect server on the same machine.

Note: If you are using the Open Snap Store Manager, please refer to the [IBM Spectrum Protect Blueprints](#) for blueprints for the small, medium, and large storage references.

Hardware recommendations for vSnap servers

The following table shows recommended CPU, memory, and disk configurations for use with compression and deduplication. This table will be used as a reference in this chapter.

In general, you should round-up the results from the sizing tool based on the following table. For example, if the sizing tool indicates that you need 30 TB of capacity, you should build to the 50 TB configuration. If you are using compression only and the sizing tool indicates that you need > 200 TB of capacity, you should build based on increments of 200 TB. For example, if the sizing tool indicates you need 300 TB of capacity you should build 2 x 200 TB vSnap servers. It is also recommended that you provision all vSnap servers within the same site with identical sizes, so they will be optimized for load balancing. You must use the same capacity when building multiple vSnap servers in a single site. For example, if you are going to build a site with a total of 200 TB using two vSnap servers, each vSnap server should be 100 TB.

Note: It is recommended that you manage a maximum of 1500 resources (virtual machines, applications, databases, etc.) per instance of a vSnap server. The sizing spreadsheet will take the number of VMware virtual machines into account when calculating the sizing for the vSnap servers.

When creating a new vSnap server, it's recommended to use uniform LUN sizes. For example, if you are creating a 75 TB vSnap server, you should use 5 x 15 TB LUNs and not a combination of 10 TB and 5 TB LUNs.

When creating a new Site, it is recommended that the **VMware VM allocation** slider be set to 50 in the **Site Properties** section.

Adding storage for vSnap servers

If you need to add additional storage to your IBM Spectrum Protect Plus infrastructure there are three general ways of adding storage depending on your desired outcome:

1. Adding a vSnap server and a new site:

Use this for cases where you are introducing a new location for replication or for a new workload. For example, if you have established backups at Primary Site in Location 1, and you want to start backups of new workloads in Location 2, it is recommended to add a vSnap server (with its required storage as calculated by the sizing spreadsheet) to the Secondary Site in Location 2.

2. Adding a vSnap server to an existing site:

Use this for cases where you are introducing new workloads to an existing location. For example, if you have virtual machines being protected by Primary Site in Location 1, and you have additional virtual machines (for example, a new datacenter) that you need to start protecting, it is recommended to add a vSnap server (with its required storage as calculated by the sizing tool) to the Primary Site.

3. Expanding vSnap server pool by adding additional LUNs:

Use this for cases where you need additional space in your vSnap server for existing workloads that grew beyond what was originally calculated by the sizing tool. Expanding the vSnap server pool by expanding the existing LUNs is not supported by IBM. You must add additional LUNs if you wish to increase the size of an existing vSnap server.

Adherence to uniform LUN sizes is not necessary when expanding a vSnap server. In the case of expansion, it is ok to use a LUN size that is different than the original LUN size, but you should use a uniform LUN size for all newly added disks. For example, if you have a 75 TB vSnap server consisting of 5 x 15 TB LUNs and want to expand your vSnap server by 50 TB, it is recommended that you use 5 x 10 TB LUNs and not a combination of 15 TB and 5 TB LUNs to provide the additional space. For more information about expanding the vSnap server storage refer to the “Expanding a vSnap storage pool” section of the IBM Spectrum Protect Plus [Installation and User’s Guide](#).

If the site has more than one vSnap server and you need to expand the storage in the site to protect new workload, all the vSnap servers need to be expanded by the same amount so that they will all have identical capacities after expansion.

If the site has more than one vSnap server and you need to add space because one (or more) of the vSnap servers is about to run out of space because of improper vSnap storage balancing, follow these steps:

1. Disable new allocations on the vSnap servers that are space constrained. You can locate the disk in the **Backup Storage** tab of the user interface, select the vSnap server disk, and check *Disable New Allocation* in the **Advanced Options** tab.
2. Add capacity to vSnap servers that are space constrained.

Log and cache can help optimize general vSnap server performance in specific instances:

- The log is used to optimize write performance for database application log backups and for data re-use operations. Starting in IBM Spectrum Protect Version 10.1.3 the log is no longer used during backup operations. The log is considered optional.
- The cache is used to optimize backups using deduplication and performance for data re-use scenarios. The cache is considered optional.

The purpose of the cloud cache is reviewed in [Chapter 7: Configuring Cloud Object Storage](#)

Note: It is recommended to limit each vSnap server with the default recommended usable capacity of 200 TB (compression only) or 100 TB (deduplication and compression) to optimize performance and not to exceed usable capacity of 600 TB (compression only) or 200 TB (deduplication and compression).

Hardware recommendations for combined vSnap server and VADP proxy

The following hardware recommendations are for deployments using the same system to host the vSnap server and VADP proxy components. These are the general guidelines that are used in the sizing tool. The sizing tool will provide you with a more precise recommendation based on your environment.

Capacity	CPU ¹	Memory	Storage Pool ²	Log ³ (Optional)	Cache ³ (Optional)	Cloud Cache ^{3,4}	VADP soft cap	vSnap streams ⁵
1 TB	8 cores	40 GB	2 x 500 GB	none	none	128 GB	4	4
10 TB	8 cores	48 GB	5 x 2 TB	10 GB	none	500 GB	4	8
25 TB	10 cores	60 GB	5 x 5 TB	10 GB	none	750 GB	4	12
50 TB	16 cores	80 GB	5 x 10 TB	10 GB	100 GB	1 TB	8	25
75 TB	18 cores	100 GB	5 x 15 TB	10 GB	150 GB	1.25 TB	12	35
100 TB	20 cores	120 GB	10 x 10 TB	10 GB	200 GB	1.5 TB	24	50
200 TB	28 cores	220 GB	10 x 20 TB	10 GB	500 GB	1.5 TB	36	50

Table 2 - Recommendations when using vSnap deduplication, compression, and encryption – combined vSnap server and VADP proxy

Capacity	CPU ¹	Memory	Storage Pool ²	Log ³ (Optional)	Cache ³ (Optional)	Cloud Cache ^{3,4}	VADP soft cap	vSnap streams ⁵
1 TB	8 cores	40 GB	2 x 500 GB	None	none	128 GB	4	4
10 TB	8 cores	40 GB	5 x 2 TB	10 GB	none	500 GB	4	8
25 TB	8 cores	40 GB	5 x 5 TB	10 GB	none	750 GB	4	12
50 TB	14 cores	40 GB	5 x 10 TB	10 GB	100 GB	1 TB	8	25
75 TB	14 cores	40 GB	5 x 15 TB	10 GB	150 GB	1.25 TB	12	35
100 TB	16 cores	40 GB	10 x 10 TB	10 GB	200 GB	1.5 TB	24	50
200 TB	20 cores	60 GB	10 x 20 TB	10 GB	500 GB	1.5 TB	36	50
400 TB	20 cores	60 GB	20 x 20 TB	10 GB	500 GB	1.5 TB	36	50
600 TB	24 cores	76 GB	30 x 20 TB	10 GB	750 GB	2 TB	36	50

Table 3 - Recommendations when using vSnap compression and encryption – combined vSnap server and VADP proxy

¹ As an example of CPUs tested see tables listed later in this chapter.

² If you have LUN sizes that are not the same as in the table, you can use other sizes.

Approximate the total size of the storage pool you are allocating and use at least 2 or 3 LUNs. For example, if the table indicates 5 x 10 TB LUNs it can be acceptable to use a configuration that provides approximately the same amount of total storage such as 6 x 8 TB LUNs.

³ Log, cache and cloud cache should be placed on SSD Flash or NVMe to optimize performance.

⁴ For more information about the cloud object storage cache see **Chapter 8. Configuring Cloud Object Storage**

⁵ The vSnap streams settings do not have to be set on vSnap servers that are exclusively vSnap replication targets

Hardware recommendations for dedicated vSnap server

The following hardware recommendations are for vSnap servers in deployments using different systems to host the vSnap server and VADP proxy components. These are the general guidelines that are used in the sizing tool. The sizing tool will provide you with a more precise recommendation based on your environment.

Capacity	CPU ¹	Memory	Storage Pool ²	Log ³ (Optional)	Cache ³ (Optional)	Cloud Cache ^{3,4}	vSnap streams
1 TB	4 cores	32 GB	2 x 500 GB	none	none	128 GB	4
10 TB	4 cores	40 GB	5 x 2 TB	10 GB	none	500 GB	8
25 TB	6 cores	52 GB	5 x 5 TB	10 GB	none	750 GB	12
50 TB	8 cores	72 GB	5 x 10 TB	10 GB	100 GB	1 TB	25
75 TB	10 cores	92 GB	5 x 15 TB	10 GB	150 GB	1.25 TB	35
100 TB	10 cores	112 GB	10 x 10 TB	10 GB	200 GB	1.5 TB	50
200 TB	16 cores	208 GB	10 x 20 TB	10 GB	500 GB	1.5 TB	50

Table 4 - Recommendations when using vSnap deduplication, compression, and encryption – vSnap server only

Capacity	CPU ¹	Memory	Storage Pool ²	Log ³ (Optional)	Cache ³ (Optional)	Cloud Cache ^{3,4}	vSnap streams
1 TB	4 cores	32 GB	2 x 500 GB	none	none	128 GB	4
10 TB	4 cores	32 GB	5 x 2 TB	10 GB	none	500 GB	8
25 TB	4 cores	32 GB	5 x 5 TB	10 GB	none	750 GB	12
50 TB	6 cores	32 GB	5 x 10 TB	10 GB	100 GB	1 TB	25
75 TB	6 cores	32 GB	5 x 15 TB	10 GB	150 GB	1.25 TB	35
100 TB	6 cores	32 GB	10 x 10 TB	10 GB	200 GB	1.5 TB	50
200 TB	8 cores	48 GB	10 x 20 TB	10 GB	500 GB	1.5 TB	50
400 TB	8 cores	48 GB	20 x 20 TB	10 GB	500 GB	1.5 TB	50
600 TB	12 cores	64 GB	30 x 20 TB	10 GB	750 GB	2 TB	50

Table 5 - Recommendations when using vSnap compression and encryption – vSnap server only

¹ As an example of CPUs tested see tables listed later in this chapter.

² If you have LUN sizes that are not the same as in the table, you can use other sizes.

Approximate the total size of the storage pool you are allocating and use at least 2 or 3 LUNs. For example, if the table indicates 5 x 10 TB LUNs it can be acceptable to use a configuration that provides approximately the same amount of total storage such as 6 x 8 TB LUNs. When using larger LUN sizes ensure you have enough spindles for desired performance characteristics. Drive re-build times will also be longer with larger LUN sizes.

³ Log, cache and cloud cache should be placed on SSD Flash or NVMe to optimize performance.

⁴ For more information about the cloud object storage cache see [Chapter 8: Configuring Cloud Object Storage](#)

Hardware recommendations for dedicated VADP proxy

The following hardware recommendations are for VADP proxies in deployments using different systems to host the vSnap server and VADP proxy components. These recommendations are based on the [IBM Spectrum Protect Plus system requirements](#).

Note: It is recommended to use a minimum of one VADP proxy for each 350 virtual machines under protection. For example, if you are protecting 500 virtual machines you should use 2 VADP proxies. In general provisioning an extra VADP proxy can provide an extra level of availability if one of the VADP proxies becomes unavailable.

The sizing tool provides more specific resource recommendations based on the workload.

vSnap Capacity	CPU¹	Memory	VADP soft cap
200, 400, 600 TB	12 cores	12 GB	36
100 TB	10 cores	8 GB	24
75 TB	8 cores	8 GB	12
50 TB	8 cores	8 GB	8
1, 10, 25 TB	4 cores	8 GB	4

Table 6 – Dedicated VADP proxy recommendations

Hardware recommendations for combined VADP and OSSM proxy

The following hardware recommendations are for Open Snap Store Manager (OSSM) deployments. The OSSM and VADP proxies are installed on the same machine. The sizing spreadsheet uses the *Medium* size *VADP + OSSM proxy* size as a basis for its calculations. Use this table as a guideline if you wish to provide more compute and storage for the OSSM and VADP proxy.

VADP + OSSM proxy size	CPU	Memory	VADP soft cap	/ossm capacity	Maximum vms managed
Small	8 cores	8 GB	8	200 GB	350
Medium	10 cores	16 GB	16	500 GB	500
Large	12 cores	24 GB	24	750 GB	500
Extra-large	14 cores	32 GB	32	1 TB	500

Table 7 – Dedicated VADP proxy recommendations

The following sections provide detailed specifications for storage layout. Use them as a map when you set up and configure your hardware. If you deviate from the specifications listed in this section, you must account for any changes when you configure your storage.

Blueprint for vSnap server

This section provides a blueprint for a dedicated physical vSnap server using storage internal to the server.

Hardware requirements for physical vSnap server with software defined RAID

The following hardware was used for a physical vSnap server using redundancy provided by the vSnap software. Customization, for example using drive sizes smaller than 10TB, can be made for smaller environments. 10TB drives were used to optimize capacity.

Blueprint component	Requirements	Detailed description	Quantity	Part number
Supermicro SuperStorage Server used as a vSnap server	<ul style="list-style-type: none"> 16 CPU cores, 2.2 GHz or faster 384 GB RAM 10Gb Ethernet 36 x NL-SAS drives 2 x 1.6 TB NVMe 2 x 480GB SAS OS drives 	Supermicro SuperStorage Server	1	6049P-E1 CR36H
		Intel Xeon Gold 6134 CPU (8 cores, 3.2 GHz, 25 MB cache)	2	P4X-SKL6134-SR3
		RDIMM DRx4 32 GB 2666	12	MEM-DR432L-SL02-ER26 Samsung M393A4K40BB2-CTD Samsung K4A8G045WB-BCTD 2 RoHS Low Profile
		10 TB NL-SAS 3.5", SAS12Gb/s, 7.2K RPM drives	36	HDD-A10T-SM0F27576
		Micron 9200 MAX 1.6TB NVMe PCIe3.0 3D TLC 2.5" 15mm	2	HDS-2VD-MTFDHAL1T6TCU1AR
		Micron 5100 PRO 2.5"480GB, SATA,6Gb/s,	2	HDS-2TD-MTFDDAK480TCB1AR
		Dual port Intel 10GbE with SFP	1	AOC-STGN-I2S
		Avago MegaRAID adapter	1	

Table 8 - Supermicro SuperStorage Server reference

Blueprints for multi-purpose storage and hypervisor

This section provides blueprints for multi-purpose storage servers. These servers can be used for the vSnap servers, VADP proxies, and for the IBM Spectrum Protect Plus server virtual appliance.

The first blueprint demonstrates using a Supermicro Superserver as a hypervisor to host the IBM Spectrum Protect Plus components. The second blueprint demonstrates using the IBM FlashSystem 5035 to provide storage for the IBM Spectrum Protect Plus components.

Example hypervisor for running virtualized IBM Spectrum Protect Plus server, vSnap server, and VADP proxies

The following is the specification used during testing for the hypervisor hosting virtual machines for the IBM Spectrum Protect Plus server, virtual vSnap storage servers, and VADP proxy servers. This example is given to use as a reference of a configuration proven to be suitable for this purpose rather than as exact minimum specifications. The requirements listed below will be adequate to host the IBM Spectrum Protect Plus server and several vSnap servers and VADP proxies.

Note: Care must be taken not to overprovision the CPU and memory resources assigned to virtual machines relative to the actual physical resources in the hypervisor. When the total resources consumed across virtual machines exceeds the physical capabilities of the hypervisor, failures can occur as the result of unresponsive virtual machines.

Blueprint Component	Requirements	Detailed description	Quantity	Part number
Supermicro Superserver used as a hypervisor	• 44 CPU cores	Supermicro Superserver 2U	1	2028U-E1CNRT+-OTO-22
	• 768 GB RAM	Intel Xeon E5-2699 v4, 22C, 2.2 Ghz, 55MB Cache	2	P4X-DPE52699V4-SR2JS
	• 10Gb Ethernet	32 GB DDR4-2400 ECC REG DIMM	24	MEM-DR432L-CL02-ER24
	• 16Gb Fibre channel	Seagate 2.5" 1TB SAS3 12Gb/s 7.2K RPM	2	HDD-2A1000-ST1000NX0453
	• 2 x 1 TB SAS OS drives	Dual port Intel 10GbE with SFP	1	AOC-STGN-I2S
		Emulex LPe16002B-M6 16G Fibre Channel HBA	1	AOC-LPE16002B-M6-O

Table 9 - Supermicro SuperStorage Server hypervisor reference

Example storage system to provide storage for vSnap server with storage hardware defined RAID

The following IBM FlashSystem 5035 system is well suited to providing storage for various IBM Spectrum Protect Plus components such as storage capacity for vSnap storage servers, as well as for the operating systems disks for the IBM Spectrum Protect Plus server, virtual vSnap storage servers, and VADP proxy systems. This example is given to use as a reference of a configuration proven to be suitable for this purpose rather than as exact minimum specifications.

This example gives enough capacity for pRDM LUNs for several vSnap servers as well as providing storage for datastores to host virtual machines. The expectation is that the flash storage will be created in a separate array from which volumes will be assigned to the hypervisors and formatted as datastores. The IBM Spectrum Protect Plus server and vSnap virtual appliances should be deployed into these flash datastores so that their operating system virtual disks use flash storage. The NL-SAS storage can then be used with a separate array from which to create volumes to assign to the vSnap virtual machines as pRDM disks. Customization, for example using drive sizes smaller than 6TB, or using 4 Flash modules instead of 6, can be made for smaller environments. If larger drive sizes are desired, make sure that there are enough spindles for performance considerations and be aware of larger rebuild times in case of failures.

Note: Use thick provisioned LUNs.

Blueprint Component	Requirements	Detailed description	Quantity	Part number
IBM FlashSystem 5035 disk system used as storage for IBM Spectrum Protect Plus	<ul style="list-style-type: none">IBM FlashSystem 5035 SFF Control16Gb FC host ports6 x 1.9 TB Flash drives92 drive LFF expansion92 x 6TB NL-SAS	5035 SFF Control	1	2072-3N4
		16Gb FC Adapter Pair	1	ALBB
		1.9TB 2.5 Inch Flash Drive	6	AL80
		5000 HD LFF Expansion	1	2072-92G
		6TB 7.2K 3.5 Inch NL HDD	92	AL47
		3m 12Gb SAS Cable	2	ACUC

Table 10 - IBM FlashSystem 5035 reference

Storage configuration for the IBM FlashSystem 5035

The following configuration was used for configuring arrays on the FlashSystem 5035, using distributed arrays. The flash storage was used for providing log and cache storage for vSnap servers, as well as for creating datastores to use with VMware for providing virtual disk capacity for operating system disks. The NL-SAS storage was used across four DRAID6 distributed arrays shared in a single virtualized pool to provide capacity for vSnap pool storage.

Server Storage requirement	Disk type	Disk quantity	Hot Spare coverage	RAID type	RAID Array quantity	Usable size	Suggested storage pool and MDisk names
LOG, cache, and OS datastores	1.9 TB Tier1 flash	6	1 Rebuild-areas=1	DRAID6	1 6 DDM	5.16TB	Pool = SSD / ssd_array
vSnap pool capacity	6 TB NL-SAS	92	4 Rebuild-areas=2 per array	DRAID6	2 46 DDM	197.91 TB each 395.82 TB total	Pool = NLSAS Array1 Array2

Table 11 - IBM FlashSystem MDisk configuration

Use the table above for guidance on creating arrays and adding storage to create MDisk and storage pools. You will then need to create volumes based on the quantities and sizes listed in the beginning of this chapter. After volume creation, assign the volume to the VMware vSphere hosts, physical machines or both depending on choice of physical or virtual server.

Note: When you create volumes don't use thin provisioned volumes. (Capacity savings = None).

For example, to host the IBM Spectrum Protect Plus virtual appliance you would first create a 400 GB volume from the SSD pool. After the volume is created, you will assign the volume to the VMware vSphere host or hostcluster that will be hosting the virtual appliance.

Chapter 5. IBM Spectrum Protect Plus Server

The sizing spreadsheet also includes the sizing of the IBM Spectrum Protect Plus server, including the compute requirements (CPU and memory) and the storage requirements for the two main catalogs. While the default provisioned compute and storage sizes will be adequate for most workloads, sizing for larger VMware workloads will require additional compute and storage for the IBM Spectrum Protect Plus server. For more information about the default IBM Spectrum Protect Plus server system requirements refer to the [IBM Spectrum Protect Plus – All Requirements Doc](#).

Note: it is recommended to place the IBM Spectrum Protect Plus server appliance on high performing, flash storage to optimize catalog performance.

Note: the current sizing recommendations for the IBM Spectrum Protect Plus server appliance is only based on VMware workload at this time.

The following is an overview of the system requirements that will be recommended by the sizing spreadsheet and what you can expect from the sizing exercise:

- CPU requirements – the recommendation will always be to use the default 8 CPU
- Memory requirements – recommendation to increase memory in increments of 16 GB for larger workloads due to larger catalogs
- Configuration catalog – recommendation to increase the catalog size (/data directory) in increments of 50 GiB for larger workloads. Information about expanding virtual disks can be found in the “Adding virtual disks” chapter of the IBM Spectrum Protect Plus [documentation](#).
- Recovery catalog – the recommendation will always be to use the default 50 GiB
- File catalog – optional for catalog file metadata; the sizing of the file catalog is not in the scope of the Blueprints exercise. Information about sizing the file catalog can be found in the [File indexing and restore requirements](#) documentation.

Note: If you plan on using the file catalog feature, please review the information in this [IBM Support document](#) regarding adjusting the available memory in the Virgo component of the IBM Spectrum Protect Plus server.

Adjusting IBM Spectrum Protect Plus server global settings

There are a set of global settings that should be adjusted for the IBM Spectrum Protect Plus server based on the results of the sizing exercise. Global preferences can be set in the **System Configuration -> Global Preferences** panel in the IBM Spectrum Protect Plus user interface. All of the recommended settings can be found on the **Sizing Results** worksheet of the sizing spreadsheet.

- Target free space error (percentage) – This is the value that is used to determine when error messages will be issued based on available free space in the vSnap server repository. This value is dependent on the size of the vSnap server.
- Target free space warning (percentage) – This is the value that is used to determine when warning messages will be issued based on available free space in the vSnap server repository. This value is dependent on the size of the vSnap server.
- Group VMs by: Number of VMs in group – this is the number of vms that will be used in atomic operations on a vSnap server such as creating a recovery point (snapshot creation) or expiring a recovery point (snapshot delete). The default value is 5 and it is not recommended to lower this value.

Changing the frequency of the Storage Server Inventory job

It is recommended that you change the frequency of the Storage Server Inventory job on the IBM Spectrum Protect Plus server as the job is no longer needed to be executed on a daily basis. You can navigate to the **Jobs and Operations** panel and select the **Calendar** icon to change the frequency of the *Storage Server Inventory* job.

Chapter 6. vSnap Server Installation and Setup

Configuring a virtual vSnap server using storage hardware defined RAID

This example illustrates how to implement the vSnap server component within a virtual machine where the RAID protection is being provided by the storage hardware. Storage is assigned to the virtual machine as a physical RDM type disk. The exception to this is the operating system disk which is deployed as a virtual disk into an existing datastore as part of the virtual appliance deployment.

Note: Use vSnap RAID 0 when the storage LUNs being used already have RAID redundancy provisioned by the storage hardware.

Note: The virtual vSnap appliance includes one virtual SCSI adapter of type *VMware Paravirtual*. Do not modify the type of virtual SCSI adapter or the system can fail to boot.

Note: the single virtual SCSI adapter provides sufficient SCSI target slots for thirteen additional disks to be attached. In cases where more disks will be needed for vSnap storage volumes, or in cases where the VADP proxy component will be used with the HotAdd transport, additional virtual SCSI adapters will need to be added to provide additional slots. The additional adapters must use the same *VMware Paravirtual* type.

Use the following steps to configure a virtual vSnap server using storage hardware provided RAID:

1. Deploy the vSnap virtual appliance (OVT template installation file) as documented in the *IBM Spectrum Protect Plus Installation and User's Guide*.
2. Prior to powering on the newly deployed virtual machine, edit the hardware settings for the machine as follows using recommendations from the previous table:
 - a. Adjust the number of virtual CPUs according to the planned vSnap capacity. For our testing, we allocated CPU cores across two sockets.
 - b. Adjust the amount of virtual RAM according to the planned vSnap capacity.
 - c. Delete the 100 GB virtual disk that was created during virtual machine import as that storage will not be used. This disk is reported as *Hard disk 2*.

Note: After deleting this disk, the virtual machine should now contain only a 50 GB *Hard disk 1* and a 128 GB *Hard disk 2*.

- d. If you will be performing data copy to cloud object storage (***Hard disk 2***) or the Spectrum Protect server, increase the size of the 128 GB virtual disk if you have determined that more capacity is required for the cloud cache

- e. Attach all required storage pool, log, and cache LUNs as physical RDM disks using the recommended size and quantity according to the preceding table. For example, the following is recommended for a vSnap with 50 TB of capacity:
 - i. 5 x 10 TB storage pools (*Hard disk 3, 4, 5, 6, and 7*)
 - ii. 2 x 20 GB log (*Hard disk 8, 9*)
 - iii. 100 GB cache (*Hard disk 10*)
3. Power on the vSnap virtual machine. After the system boots, connect to the vSnap system using *ssh* and log in as *serveradmin* using the default password of *sppDP758-SysXyz*. You will be prompted to change the default password to a secure password. Confirm that network connectivity is set up correctly by issuing a ping to other systems on the network. If there is a problem with network connectivity, it can be corrected using the *nmtui* program.
4. Complete the configuration of a vSnap storage pool by following these steps from the *ssh* terminal:
 - a. Initialize vSnap:

```
vsnap system init --skip_pool
```

- b. Past versions of this document instructed you to create a vSnap server **Username** and **Password**. The *serveradmin* username is already configured with the necessary permissions for registering the vSnap server to the IBM Spectrum Protect Plus server and for optionally registering the VADP proxy. There is no longer any need to create a new username.
- c. List the disks assigned to vSnap and carefully review the list to identify the disks that were assigned as pRDM disks. Note the *NAME* field reported for each of these disks to be used in the vSnap pool commands that follow.

Note: The rescan command is usually not necessary, but it is helpful in cases where additional disks are assigned to the system after it has been powered on.

Note: You can typically identify the disks by their sizes.

ATTENTION: Other disks such as operating system disks are listed in this output, and care must be taken not to include those disks in any commands which follow.

```
vsnap disk rescan  
vsnap disk show
```

The following is the expected output for a vSnap server with 50 TB of capacity.

The disks must show as unused to be used with a new vSnap pool. If the disks are not listed as unused, then there is the possibility that some other operating system feature is using them, or they were used previously for some other purpose. There is a *vsnap support* command that can clear a disk in cases where you are absolutely certain it is safe to do so. Refer to the section [Erasing a disk](#) in the appendix.

UUID	< ... >	TYPE	VENDOR	MODEL	SIZE	USED AS	NAME
36000c295a4e<...>c5579		SCSI	VMware	Virtual disk	50.00GB	xfs	/dev/sda
36000c29e938<...>1f415		SCSI	VMware	Virtual disk	120.00GB	LVM_member	/dev/sdb
360050763808<...>0004d		SCSI	IBM	2145	10.00TB	unused	/dev/sdc
360050763808<...>0004e		SCSI	IBM	2145	10.00TB	unused	/dev/sdd
360050763808<...>0004f		SCSI	IBM	2145	10.00TB	unused	/dev/sde
360050763808<...>00050		SCSI	IBM	2145	10.00TB	unused	/dev/sdf
360050763808<...>00051		SCSI	IBM	2145	10.00TB	unused	/dev/sdg
360050763808<...>00065		SCSI	IBM	2145	10.00GB	unused	/dev/sdh
360050763808<...>00065		SCSI	IBM	2145	100.00GB	unused	/dev/sdi

- d. Create the vSnap storage pool with compression enabled and optionally deduplication enabled based on the technology choice you made earlier in this document. The disks used for the disk list might need to be adjusted based on what was reported in the output of the command in the previous step.

Note: You can only create a single pool named *primary*

Compression only:

```
vsnap pool create --name primary --disk_list
/dev/sdc,/dev/sdd,/dev/sde,/dev/sdf,/dev/sdg --
pool_type raid0 --deduplication off --compression on
```

Compression, deduplication, and encryption

```
vsnap pool create --name primary --disk_list
/dev/sdc,/dev/sdd,/dev/sde,/dev/sdf,/dev/sdg --
pool_type raid0 --deduplication on --compression on
--encryption on --enc_type disk
```

Note: You can choose not to deduplicate, compress, and/or encrypt data.

Change the parameters of the deduplication, compression, or encryption options to “off” where applicable, for example: `--encryption off`

Note: You can only specify the use of encryption when a vsnap pool is created. Compression and deduplication can be changed on a vsnap pool at any time.

- e. For configurations that will use a log on high-performing disk, run the following command. Optionally, you can specify two devices for a mirrored log.

```
vsnap pool addlog --id 1 --disk_list=/dev/sdh
```

- f. For configurations that will use a cache on high-performing disk, run the following command:

```
vsnap pool addcache --id 1 --disk_list /dev/sdi
```

g. Run the following command to verify that your pool was created: successfully.

```
vsnap pool show
```

The following example is for a pool with 50 TB of capacity:

```
TOTAL: 1
ID: 1
NAME: primary
POOL TYPE: raid0
STATUS: ONLINE
HEALTH: 100
COMPRESSION: Yes
COMPRESSION RATIO: 3.10
DEDUPLICATION: Yes
DEDUPLICATION RATIO: 1.15
ENCRYPTION:
    ENABLED: Yes
    TYPE: disk

SYNC WRITE: Yes
TOTAL SPACE: 48.13TB
FREE SPACE: 47.97TB
USED SPACE: 163.81GB
DATA SIZE BEFORE DEDUPLICATION: 188.37GB
DATA SIZE BEFORE COMPRESSION: 581.63GB
CREATED: 2018-03-13 00:07:33 UTC
UPDATED: 2018-03-13 00:08:36 UTC
DISKS PER RAID GROUP: 1
DISKS IN POOL:
    CACHE:
        /dev/sdj1
    LOG:
        /dev/sdh1
        /dev/sdi1
    RAID0:
        /dev/sdc1
        /dev/sdd1
        /dev/sde1
        /dev/sdf1
        /dev/sdg1
```

Configuring a physical vSnap server using storage software provided RAID

This example illustrates how to implement the vSnap server component within a physical server using the software RAID 6 protection being provided by the vSnap server. Storage is provided by near line (NL) SAS drives internal to the vSnap server that are presented from the SAS controller as a JBOD. The physical server will be provisioned with Linux during the workflow.

In this example we will demonstrate building a vSnap server with approximately 200 TB of usable storage.

Use the following steps to configure a physical vSnap server using the software provided RAID from the vSnap server:

1. Configure the system integrated management module that makes it simpler to remotely perform tasks such as configuring the system BIOS as demonstrated in the next step.
2. Prepare the BIOS settings. The following settings require adjustment:
 - a. Create a RAID1 mirror of the two operating system boot disks using the integrated *Intel® RSTe sSATA Controller*.
 - b. Set the system to use UEFI.
3. Install a Linux operating system that is supported for the vSnap component. Testing was performed using Red Hat Enterprise Linux. The following installation choices were made:
 - a. Software selection is *Server with GUI*, and an add-on including *Development Tools*.
 - b. The installation target was the 424.77 GiB BIOS RAID set (mirror). Default partitioning was adjusted to provision at least 300 GB to the root file system (/).
 - c. The recommended Swap size of 16 GB was used
 - d. Network settings and default passwords were selected.

Note: It is recommended that you allocate at least 300 GB to the root file system (/) and at least 16 GB to Swap.

4. Perform additional operating system configuration after the system successfully boots for the first time.
 - a. Enable an update subscription so that the operating system can be updated to the latest updates using the *yum* update command. The vSnap software cannot be installed without completing this step.
 - b. Install the device driver for the Avago MegaRAID adapter and the MegaRAID storage manager software:

(*Note:* Ensure that the MegaRAID storage manager software is current).

<ftp://ftp.supermicro.com/driver/SAS/LSI/3108/>

```
gzip -d megaraid_sas-07.702.06.00-src.tar.gz
tar -xvf megaraid_sas-07.702.06.00-src.tar
cd megaraid_sas-07.702.06.00/
./compile.sh
modinfo megaraid_sas

tar -xvf MSM_linux_x64_installer-17.05.00-02.tar
cd disk
./RunRPM.sh
```

- c. Put the NL-SAS disks in JBOD mode:

```
/opt/MegaRAID/storcli/storcli64 /c0 set jbod=on
```

5. Install the vSnap server software:

```
chmod +x vsnap-dist-<build>.run
./vsnap-dist-1<build>.run
```

6. Complete the configuration of a vSnap storage pool by following these steps:

- a. Initialize vSnap:

```
vsnap system init --skip_pool
```

- b. Past versions of this document instructed you to create a vSnap server **Username** and **Password**. The *serveradmin* username is already configured with the necessary permissions for registering the vSnap server to the IBM Spectrum Protect Plus server and for optionally registering the VADP proxy. There is no longer any need to create a new username.
- c. List the disks assigned to vSnap and carefully review the list of disks to identify the disks that were assigned as physical RDM disks. Note the *NAME* field reported for each of these disks to be used in the *vsnap pool* commands that follow.

Note: The rescan command is usually not necessary, but it is helpful in cases where additional disks are assigned to the system after it has been powered on.

Note: You can typically identify the disks by their sizes.

ATTENTION: Other disks such as operating system disks are listed in this output,

and care must be taken not to include those disks in any commands which follow:

```
vsnap disk rescan
vsnap disk show
```

The following is the expected output for the case of a vSnap server with a 205 TB capacity. The disks must show as unused to be used with a new vSnap pool. If the disks are not listed as unused then there is the possibility some other operating system feature is using them, or they were used previously for some other purpose. There is a *vsnap support* command that can clear a disk in cases where you are absolutely certain it is safe to do so. Refer to the section [Erasing a disk](#) in the appendix.

UUID	TYPE	VENDOR	MODEL	SIZE	USED AS	NAME
35000cca26c25cc44	SCSI	HGST	HUH721010AL5200	9.10TB	unused	/dev/sda
35000cca26c255288	SCSI	HGST	HUH721010AL5200	9.10TB	unused	/dev/sdaa
35000cca26c25aab4	SCSI	HGST	HUH721010AL5200	9.10TB	unused	/dev/sdab
35000cca26c25aa58	SCSI	HGST	HUH721010AL5200	9.10TB	unused	/dev/sdac
35000cca26c24c928	SCSI	HGST	HUH721010AL5200	9.10TB	unused	/dev/sdad
35000cca26c25aae8	SCSI	HGST	HUH721010AL5200	9.10TB	unused	/dev/sdae
35000cca26c25a3fc	SCSI	HGST	HUH721010AL5200	9.10TB	unused	/dev/sdaf
35000cca26b6063d4	SCSI	HGST	HUH721010AL5200	9.10TB	unused	/dev/sdag
35000cca26c25a32c	SCSI	HGST	HUH721010AL5200	9.10TB	unused	/dev/sdah
35000cca26c25cdb0	SCSI	HGST	HUH721010AL5200	9.10TB	unused	/dev/sdai
35000cca26c2551ec	SCSI	HGST	HUH721010AL5200	9.10TB	unused	/dev/sdaj
Micron_5100_MT<...>E	ATA	ATA	Micron_5100_MTFD	447.13GB	vfat	/dev/sdak
Micron_5100_MT<...>B	ATA	ATA	Micron_5100_MTFD	447.13GB	vfat	/dev/sdal
35000cca26c259a54	SCSI	HGST	HUH721010AL5200	9.10TB	unused	/dev/sdb
35000cca26b951104	SCSI	HGST	HUH721010AL5200	9.10TB	unused	/dev/sdc
35000cca26c245844	SCSI	HGST	HUH721010AL5200	9.10TB	unused	/dev/sdd
35000cca26b9df564	SCSI	HGST	HUH721010AL5200	9.10TB	unused	/dev/sde
35000cca26c223b60	SCSI	HGST	HUH721010AL5200	9.10TB	unused	/dev/sdf
35000cca26b9cabdc	SCSI	HGST	HUH721010AL5200	9.10TB	unused	/dev/sdg
35000cca26c24ca14	SCSI	HGST	HUH721010AL5200	9.10TB	unused	/dev/sdh
35000cca26c245724	SCSI	HGST	HUH721010AL5200	9.10TB	unused	/dev/sdi
35000cca26c25d6f8	SCSI	HGST	HUH721010AL5200	9.10TB	unused	/dev/sdj
35000cca26c25a804	SCSI	HGST	HUH721010AL5200	9.10TB	unused	/dev/sdk
35000cca26c257a10	SCSI	HGST	HUH721010AL5200	9.10TB	unused	/dev/sdl
35000cca26c25a10c	SCSI	HGST	HUH721010AL5200	9.10TB	unused	/dev/sdm
35000cca26c25cf0c	SCSI	HGST	HUH721010AL5200	9.10TB	unused	/dev/sdn
35000cca26c25cd9c	SCSI	HGST	HUH721010AL5200	9.10TB	unused	/dev/sdo
35000cca26c25d0b8	SCSI	HGST	HUH721010AL5200	9.10TB	unused	/dev/sdp
35000cca26c24c8cc	SCSI	HGST	HUH721010AL5200	9.10TB	unused	/dev/sdq
35000cca26b606708	SCSI	HGST	HUH721010AL5200	9.10TB	unused	/dev/sdr

- d. Create a Linux *Multiple Device* driver (md) device to store the vSnap storage cache, cloud cache and logs. Two NVMe devices will be used with software RAID 10 to back the md device.
 - i. Create the raid10 md device

```
mdadm --create /dev/md0 --level=raid10
--raid-devices=2 /dev/nvme0n1 /dev/nvme1n1
```

- ii. Monitor progress of the initialization of the raid. You will see output like below, where md0 is the newly created device, and md126 is the mirrored OS boot drives

```
cat /proc/mdstat
```

```
Personalities : [raid1] [raid10]
md0 : active raid10 nvme0n1[0] nvme1n1[1]
      937560064 blocks super 1.2 2 near-copies [2/2] [UU]
      bitmap: 0/7 pages [0KB], 65536KB chunk

md126 : active raid1 sdak[1] sda1[0]
      445431808 blocks super external:/md127/0 [2/2] [UU]
```

- iii. Create three partitions on the md device: a log partition, a cache partition and a cloud cache partition.

```
fdisk /dev/md0
```

- iv. Set the new md device to automatically start during system boot
(Note: Do not to run this command more than one time)

```
mdadm --detail --scan | grep /dev/md0 >> /etc/mdadm.conf
```

- e. Create the vSnap storage pool with compression enabled and optionally deduplication enabled using 35 of the 10TB NL-SAS drives. The disks used for the disk list might need to be adjusted based on what was reported in the output of the command in the previous step. The expected result is that a pool consisting of 5 separate arrays each with 7 disks will be created.

Note: You can only create a single pool named **primary**

Compression, deduplication, and encryption

```
vsnap pool create --name primary --disk_list
/dev/sda,/dev/sdb,/dev/sdc,/dev/sdd,/dev/sde,/dev/sdf,
/dev/sdg,/dev/sdh,/dev/sdi,/dev/sdj,/dev/sdk,/dev/sdl,
/dev/sdm,/dev/sdn,/dev/sdo,/dev/sdp,/dev/sdq,/dev/sdr,
/dev/sds,/dev/sdt,/dev/sdu,/dev/sdv,/dev/sdw,/dev/sdx,
/dev/sdy,/dev/sdz,/dev/sdaa,/dev/sdab,/dev/sdac,/dev/s
dad,/dev/sdae,/dev/sdaf,/dev/sdag,/dev/sdah,/dev/sdai
--pool_type raid6
--compression on --deduplication on
--encryption on --enc_type disk
```

Note: You can choose not to deduplicate, compress, and/or encrypt data.

Change the parameters of the deduplication, compression, or encryption options to “off” where applicable, for example: `--encryption off`

Note: You can only specify the use of encryption when a vsnap pool is created.

Compression and deduplication can be changed on a vsnap pool at any time.

- f. Define the last 10 TB NL-SAS drive as a spare:

```
vsnap pool addspare --id 1 --disk_list /dev/sdaj
```

- g. Optionally define the first NVMe drive as the log:

```
vsnap pool addlog --id 1 --disk_list /dev/md0p1 --force
```

- h. Expand the cloud cache if the size of the cache virtual disk was increased in an earlier step. For more information about expanding the size of the cloud cache refer to [Expanding the cache area if it already exists](#).

- i. Optionally define the second NVMe drive as a cache:

```
vsnap pool addcache --id 1 --disk_list /dev/md0p2 --force
```

- j. Verify that your pool was created successfully:

```
vsnap pool show
```

```
ID: 1
NAME: primary
POOL TYPE: raid6
STATUS: ONLINE
HEALTH: 100
COMPRESSION: Yes
COMPRESSION RATIO: 1.00
DEDUPLICATION: Yes
DEDUPLICATION RATIO: 1.00
ENCRYPTION:
    ENABLED: Yes
    TYPE: disk

SYNC WRITE: Yes
TOTAL SPACE: 205.11TB
FREE SPACE: 204.85TB
USED SPACE: 2.18MB
DATA SIZE BEFORE DEDUPLICATION: 935.09KB
DATA SIZE BEFORE COMPRESSION: 101.50KB
CREATED: 2018-03-02 01:28:36 UTC
UPDATED: 2018-03-08 23:28:39 UTC
DISKS PER RAID GROUP: 7
DISKS IN POOL:
```

```
CACHE :  
    /dev/md0p2  
LOG :  
    /dev/md0p1  
RAID6 GROUP1 :  
    /dev/sda1  
    /dev/sdb1  
    /dev/sdc1  
    /dev/sdd1  
    /dev/sde1  
    /dev/sdf1  
    /dev/sdg1  
RAID6 GROUP2 :  
    /dev/sdh1  
    /dev/sdi1  
    /dev/sdj1  
    /dev/sdk1  
    /dev/sdl1  
    /dev/sdm1  
    /dev/sdn1  
RAID6 GROUP3 :  
    /dev/sdo1  
    /dev/sdp1  
    /dev/sdq1  
    /dev/sdr1  
    /dev/sds1  
    /dev/sdt1  
    /dev/sdu1  
RAID6 GROUP4 :  
    /dev/sdv1  
    /dev/sdw1  
    /dev/sdx1  
    /dev/sdy1  
    /dev/sdz1  
    /dev/sdaa1  
    /dev/sdab1  
RAID6 GROUP5 :  
    /dev/sdac1  
    /dev/sdad1  
    /dev/sdae1  
    /dev/sdaf1  
    /dev/sdag1  
    /dev/sdah1  
    /dev/sdai1  
SPARE :  
    /dev/sdaj1
```

Adding disk storage

Now that you have configured the vSnap server, you will need to register each instance of a vSnap server to the IBM Spectrum Protect Plus server from the server user interface.

Procedure

1. If you are certain that no data has been stored in the default vSnap server that is part of the IBM Spectrum Protect Plus server virtual appliance, unregister the default disk storage that comes with the appliance. From the navigation menu, click **Backup Storage**. Locate the **localhost** disk storage and click the “**x**” to unregister the localhost disk storage. You will be prompted to confirm the action.
2. Register the disk storage that was installed in the previous chapter. From the navigation menu, click **Backup Storage** and click **Add**. You will need to provide the hostname/IP, credentials (**UserId** *serveradmin* and **Password** created in prior steps), and register the disk storage to the desired site.

Configuring global options on vSnap server

There are a set of global settings that should be adjusted for each of the vSnap servers based on the results of the sizing exercise. Global preferences can be set in the **Advanced Options settings** for each vSnap server. All of the recommended settings can be found on the **Sizing Results** worksheet of the sizing spreadsheet.

- Interval in seconds between volume/snapshot deletions – this is the amount of time used between snapshot deletions when IBM Spectrum Protect Plus is expiring inventory from the catalog
- Concurrent backup: Limit VM or Database streams – this setting limits concurrent backups for virtual machines or databases so that the resources on a vSnap server will not be overwhelmed.

Chapter 7. Configuring VADP Proxies

After the vSnap servers have been installed and the disk storage has been registered, the final step is to install the VADP proxies that were sized in **Chapter 3. How to Use the Sizing Tool**.

Note: You can install the **vSnap server** (disk storage) and **VADP proxy** on the same physical or virtual system. IBM Spectrum Protect Plus will optimize data movement by eliminating an NFS mount when these two systems are co-located.

Note: It is recommended to use a minimum of one VADP proxy for each 350 virtual machines under protection. For example, if you are protecting 500 virtual machines you should use 2 VADP proxies. In general provisioning an extra VADP proxy can provide an extra level of availability if one of the VADP proxies becomes unavailable.

Installing VADP proxies


1. Install a supported operating system and then deploy the VADP proxy software on this system from the IBM Spectrum Protect Plus server. You must first install a supported operating system.
2. Suspend the default proxy that comes with the virtual appliance. This can be done from the IBM Spectrum Protect Plus interface. Expand the **System** menu and select **VADP Proxy**. Locate the localhost entry in the table of VADP Proxies, expand the **Actions** button and select **Suspend**.
3. Install the VADP proxy. This can be done from the IBM Spectrum Protect Plus interface. Expand the **System** menu and select **VADP Proxy**. Select the “+” button to add a proxy. Provide the hostname and credentials and select **Install**.
4. Register the VADP proxy. After the installation, you will need to register the VADP proxy. Locate the VADP proxy that you just installed on the panel. Select the appropriate site from the pulldown menu and select **Register**.
5. Configure the VADP proxy options. Expand the **Actions** button next to the VADP proxy that was added in the last step and select **Set Options**.
 - a. *Transport Modes* – Remove transport modes that the proxy is not capable of using. For example, if you are using a virtual VADP Proxy, remove the **san** transport. If you are using a physical VADP proxy, remove the **hotadd** transport. You should also order the remaining transports in order of preferred usage. Also consider removing the **hotadd** transport after the initial full backups are complete and the system is in steady-state where it is predominantly taking incremental backups and full backups of only new machines because the HotAdd transport is better suited for full backup workloads than incremental workloads.
 - b. Change the default *Enable NBDSSL Compression* to one of the listed compression algorithms (**libz**, **fastlz**, or **skipz**). If you have a 1 GbE network between the VMware

- host and VADP Proxy. If you have a 10 GbE network between the VMware host and the VADP proxy use the default of **disabled** for *Enable NBDSSL Compression*.
- c. Use the default values for *Log retention in days* and *Buffer Size in bytes*.
 - d. The **Softcap task limit** as noted below.

Setting the maximum number of VM's to process concurrently

You can set the maximum number of VM's to process concurrently. The default value is 3 concurrent VM's per ESXi host. This value is configured as part of the vCenter management in the IBM Spectrum Protect Plus server user interface. The default value will be sufficient for most workloads.

To change the maximum number of VM's to process concurrently:

1. From the navigation menu, expand **Hypervisor**, expand **VMware**, and select **Backup**.
2. Select **Manage vCenter**
3. Select the vCenter from the menu and click on the pencil icon () to change the settings
4. Expand the **Options** setting and you will see the option to set this value.

Distributing workload to multiple VADP proxies

If you need to use multiple VADP proxies to protect your environment, it is important to understand how workload is distributed among multiple VADP proxies.

1. It is recommended that you *Suspend* the default proxy that comes with the virtual appliance as noted above. This will guarantee that the VADP proxy is not placing additional workload on your IBM Spectrum Protect Plus server host.
2. To ensure multiple VADP proxies are optimally utilized to maximize data throughput, IBM Spectrum Protect Plus will automatically distribute each virtual machine backup task to the VADP proxy which has the most available memory and “free tasks” available on the system to process VADP proxy requests. This load balancing is referred to as “throttling” in IBM Spectrum Protect Plus documentation.
3. “Free tasks” are determined by the number of free CPU cores that the system has designated to process VADP proxy requests. This number is governed by the **Softcap task limit** that has been configured on the “Set VADP Proxy Options” panel on the VADP Proxy page. The **Softcap task limit** should be set based on whether the VADP proxy is a dedicated system or is collocated on a system with the vSnap server as follows:
 - a. If the VADP proxy is a dedicated system, the **Softcap task limit** should be set to **36** to enable the VADP proxy to maximize the CPU usage on the system.
 - b. If the VADP proxy is collocated on a system with the vSnap server, the Softcap task limit should be set according to the results in the sizing spreadsheet.

Note: **The Softcap task** limit cannot guarantee that the number of CPUs used by a VADP proxy will always honor the user specified value.

Chapter 8. Configuring Cloud Object Storage

For all copy operations (backup, recovery, data reuse) to cloud object storage or Spectrum Protect each vSnap server requires a disk cache area (referred to as the *cloud cache*) to perform the following functions:

- During copy operations as a temporary staging area for objects that are pending upload to the cloud object storage endpoint.
- During restore operations to cache downloaded objects as well as to store any temporary data that may be written into the restore volume.

Note: There is no disk cache needed for the archive operation.

Note: To use Spectrum Protect for copy to object storage or archive operations, refer to IBM Spectrum Protect documentation to learn how to configure storage pools for these operations.

Most of the cache space is freed up at the end of each copy or restore, but a small amount may continue to be used to cache metadata that will be used to speed up subsequent operations.

The cache area must be configured in the form of an XFS filesystem mounted at `/opt/vsnap-data` on the vSnap server. If this mount point is not configured, copy or restore jobs will fail with the error: "Cloud functionality disabled: Data disk `/opt/vsnap-data` is not configured."

Note: Do not unmount or manipulate files under `/opt/vsnap-data` while any copy or restore jobs are active. Once you have ensured that no jobs are active, it is safe to perform any maintenance activities such as unmounting and reconfiguring the cache area.

The data stored under `/opt/vsnap-data` is also safe to delete as long as no copy or restore jobs are active, although deleting this data may result in vSnap having to re-download data from the cloud object storage endpoint during the next copy or restore operation which in turn may introduce a small delay during the job. You should only delete this data if directed to perform this action by IBM support.

Default cloud cache area

Depending on the installation method and version of IBM Spectrum Protect Plus that was initially deployed, a default cloud cache may or may not already be present on the system.

For new installations of a vSnap server:

- When the vSnap server is deployed as a virtual appliance, the cloud cache area is already present as a pre-configured 128 GB data disk mounted at `/opt/vsnap-data`.
- When the vSnap server is installed on a custom server, the cloud cache area must be configured manually.

Use the "df" command on the vSnap server to confirm the presence of the mount point `/opt/vsnap-data`. If the mount point is not present, it must be configured manually.

Sizing the cloud cache

Although the cloud cache area is sized at 128 GB as a starting point, it must be expanded based on the size of the vSnap pool on that system. The table below shows some general recommendations for sizing of the cache area.

Size of pool	Size of cache area
1 TB	128 GB
10 TB	500 GB
25 TB	750 GB
50 TB	1 TB
75 TB	1.25 TB
100 TB	1.5 TB
200 TB or above	1.5 TB

Table 13 - Cloud cache sizing table

Recommended cloud cache disk technology

It is recommended to use flash storage for the cloud cache to improve copy to object storage and restore performance. For some workloads SAS 10K drives might be acceptable based on performance objectives.

Creating the cloud cache area

Note: The sample commands below assume they are being run as the user *serveradmin*. If running as root, the sudo prefix can be omitted.

- Attach one or more disks to the vSnap system. The cumulative size of the disk(s) should be based on the sizing guidelines described above.
- On the vSnap console, rescan to discover newly attached disk(s), then list them and identify the name(s) e.g. /dev/sdx, /dev/sdy.

```
$ vsnap disk rescan
$ vsnap disk show
OR
$ sudo lsblk
```
- Create a Physical Volume on each disk, then create a Volume Group named *vsnapdata* that spans all the disks, and create a Logical Volume named *vsnapdata1v*.

```
$ sudo pvcreate /dev/sdx
$ sudo pvcreate /dev/sdy
$ sudo vgcreate vsnapdata /dev/sdx /dev/sdy
$ sudo lvcreate -l 100%VG -n vsnapdata1v vsnapdata
```
- Create an XFS partition on the Logical Volume, create the mount point directory, and mount the volume.

```
$ sudo mkfs.xfs /dev/mapper/vsnapdata-vsnapdata1v
$ sudo mkdir -p /opt/vsnap-data
$ sudo mount /dev/mapper/vsnapdata-vsnapdata1v /opt/vsnap-data
```
- To ensure that the volume is remounted on reboot, edit the file */etc/fstab* and append the following line to the end of the file:

```
/dev/mapper/vsnapdata-vsnapdata1v /opt/vsnap-
data xfs defaults 0 0
```
- Run "df -h" and verify that the volume */opt/vsnap-data* is mounted and has the desired size.

Expanding the cache area if it already exists

Note: The sample commands below assume they are being run as the user *serveradmin*. If running as "root", the "sudo" prefix can be omitted.

- Attach one or more disks to the vSnap system. The cumulative size of the disk(s) should be based on the amount of space you want to add to the existing cache area. Use the "df -h" command to view the existing size of the /opt/vsnap-data mount point.
- On the vSnap console, rescan to discover newly attached disks, then list them and identify the names e.g. /dev/sdx, /dev/sdy.
\$ vsnap disk rescan
\$ vsnap disk show
OR
\$ sudo lsblk
- Create a Physical Volume on each disk, then add them to the existing Volume Group named "vsnapdata" to expand it, and then expand the existing Logical Volume named "vsnapdatalv".
\$ sudo pvcreate /dev/sdx
\$ sudo pvcreate /dev/sdy
\$ sudo vgextend vsnapdata /dev/sdx /dev/sdy
\$ sudo lvextend -l 100%VG /dev/mapper/vsnapdata-vsnapdatalv
- Extend the XFS partition to fully occupy the newly expanded Logical Volume.
\$ sudo xfs_growfs /dev/mapper/vsnapdata-vsnapdatalv
- Run "df -h" and verify that the volume /opt/vsnap-data is mounted and has the desired new size.

Chapter 9. Conclusion

Congratulations! You are now ready to start using IBM Spectrum Protect Plus. Please refer to the *IBM Spectrum Protect Plus Installation and User's Guide* for additional information on product usage.

Appendix A. vSnap Server Maintenance

Removing a disk

If are considering removing a disk (LUN) from the vSnap server please note the following:

- You can safely remove LUNs associated with vSnap cache and log if you no longer need to use the cache and/or log.
- Removing of LUNs hosting backup or replication data is not supported. Contact IBM Support if you need to remove LUNs hosting backup or replication data.

Erasing a disk

If you need to erase a disk during the initial configuration you can destroy all data on the disk and make the disk reusable for vSnap pool selection and creation. Use this command with extreme caution because it will overwrite the disk header and destroy all data on the disk.

```
vsnap disk wipe <diskid>
```

Setting the maximum number of replication streams

If you need check or change the maximum number of streams used for replications you can use the following commands.

To check the current number of streams used for replication:

```
vsnap system pref get
```

The number of streams used for replication is reported by the **replMaxStreams** value. (The default value is **5**)

To change the number of streams used for replication:

```
vsnap system pref set --name replMaxStreams --value <number>
```

Where **number** is the number of streams that you want to use.

Setting the maximum number of cloud copy streams

If you need check or change the maximum number of streams used for data copy to object storage you can use the following commands.

To check the current number of streams used for copy to object storage:

```
vsnap system pref get
```

The number of streams used for data copy to object storage is reported by the **cloudMaxStreams** value (the default value is **5**).

To change the number of streams used for data copy to object storage:

```
vsnap system pref set --name cloudMaxStreams --value <number>
```

Where **number** is the number of streams that you want to use.

Setting the maximum rate for copy to object storage

If you need check or change the maximum rate used for data copy to object storage you can use the following commands.

To check the current maximum rate used for data copy to object storage:

```
vsnap system pref get
```

The maximum copy rate is reported by the **cloudOffloadRate** value (the default value is **536870912** bytes/second which is 512 Mb/sec).

To change the number of streams used for copy to object storage:

```
vsnap system pref set --name cloudOffloadRate --value <number>
```

Where **number** is the maximum copy rate in bytes/second.

Checking file system integrity on vSnap pools

When using the vSnap server provisioned RAID 6 as discussed in [Choosing the Appropriate Technologies: RAID](#) the vSnap server provisioned software RAID can detect and correct data corruption. This is done through the underlying file system checksum mechanism. A checksum operation is performed for every block that is written to the vSnap pool. The checksum of each block is transparently validated when the block is read, allowing the detection silent corruption. In the event that there is a checksum mismatch on a read, the underlying file system will attempt to recover the data from any available redundant copy of the data.

When recovering data from a redundant copy of the data, it will be more likely to be successful the sooner that data corruption can be detected. Having a strategy to periodically read the data from the vSnap server pools will be advantageous in detecting this type of data corruption sooner than later. There are several ways to read data from the vSnap server storage pools (which in turn will validate the data checksums):

1. Replication or copy of data in the vSnap server to another vSnap server or to a cloud object storage repository. Note that this method only will read the most recently changed blocks due to the incremental forever nature of the replication and copy to object storage features.
2. Periodic recovery of resources (virtual machines, applications, databases) into production or clone mode. Note that this method will only read the selected snapshot of the resource.

3. Periodically reading all the data in the vSnap server pools using the ***vsnap pool scrub*** command

To start a scrub of a vSnap pool, issue the command

```
vsnap pool scrub start --id=1
```

To see the status of the scrub process, issue the command

```
vsnap pool scrub status --id=1
```

If you need to stop the scrub process, issue the command

```
vsnap pool scrub stop --id=1
```

Appendix B. Performance Information

You can compare IBM system performance results against your IBM Spectrum Protect Plus storage configuration as a reference for expected performance.

Observed results are based on measurements that were taken in a test lab environment. Test systems were configured according to the blueprints in this document. Backups of VMware systems were taken across a 10 Gigabit Ethernet connection to the IBM Spectrum Protect Plus vSnap server. Because many variables can influence throughput in a system configuration, do not expect to see exact matches with the results shown in this appendix.

The following typical factors can cause variations in actual performance:

- The read performance of the source disk during backup and write performance of the target disk during restore.
- The number of client sessions that are used in your environment.
- The amount of duplicate data and the response to compression.

This information is provided to serve only as a reference. The following measurements were taken using the physical Supermicro vSnap server with software provided RAID using both deduplication and compression (except where compression-only is noted). The vSnap log was not mirrored.

Metric		Limit or range	Notes
Daily amount of new backup data		8 – 20 TB/day	Amount before data reduction
Backup ingestion rate	Deduplication and compression	1.4 TB/hr	
	Compression only	2.8 TB/hr	

Table 14 - Data intake

Metric	Range	Notes
Total managed data	205 TB – 820 TB	Total managed data is the volume of data stored (size before data reduction)

Table 13 - Protected data

Metric	Number of vm's restored	Rate
Throughput of VMware clone restore	1	337.6 GB/hr
	2	675.3 GB/hr
	4	988.9 GB/hr
	8	969.6 GB/hr

Table 15 - Data recovery

Appendix C. vSnap Server Performance Test Tool

The vSnap server performance test tool (vsnapperf.pl) is useful for evaluating the performance the storage used by the vSnap server. The tool can be used prior to using IBM Spectrum Protect Plus in a production environment.

The tool generates pseudo-random data that is useful for performance measurements when the vSnap server is using compression and deduplication.

The vSnap server software must be installed and a vSnap pool must be created and mounted before using this tool.

The readme for this tool is provided here for reference:

```

#*****
# IBM Spectrum Protect Plus vSnap disk performance test tool
#
# name:      README.txt
# version: 1.4c
#
# Notice: This program is provided as a tool intended for use by IBM Internal,
#         IBM Business Partners, and IBM Customers. This program is provided as is,
#         without support, and without warranty of any kind expressed or implied.
#
# (C) Copyright International Business Machines Corp. 2013, 2021
#*****

```

OVERVIEW

The vsnapperf.pl tool is useful for evaluating the performance of storage used by vSnap prior to going into production. The tool builds around the ldeedee tool which is able to generate pseudo-random data that is useful for more accurate performance measurements when vSnap is using compression and deduplication.

As a pre-requisite, the vSnap software must be installed and a vSnap pool must be created and mounted. For more information on creating a vSnap pool, consult the IBM Spectrum Protect Plus blueprint which are available at <http://ibm.biz/IBMSpectrumProtectPlusBlueprints>

There are currently two workloads implemented, seqwrite and seqread, which are intended to simulate the large seq write and read workload driven by backup and restore. In addition, a third cleanup worklod can be used to clean up the data created by the tool. See below for more details on cleanup.

INSTRUCTIONS

Common use cases for the tool are best conveyed with a series of examples. View the tools online help with:

```
perl vsnapperf.pl -help
```

1) Perform a sequential write of a 10GiB file using a single thread.

```
perl vsnapperf.pl seqwrite -filesize 10g
```

2) Sequential write of two 10GiB files using two threads. The -overwrite option is required to overwrite the file left behind from the previous test. Do not take a snapshot of the volume after the write completes.

```
perl vsnapperf.pl seqwrite -filesize 10g -numVol 2 -snapshot no -overwrite=yes
```

3) Perform five iterations of five threads writing 10GiB files. Take a snapshot of all five volumes at the end of each iteration.

```
perl vsnapperf.pl seqwrite -filesize 10g -numVol 5 -iterations 5 -snapshot yes -overwrite=yes
```

4) Perform two iterations of five thread reads of the 10 GiB files left from the previous test #3.

```
perl vsnapperf.pl seqread -filesize 10g -numVol 5 -iterations 2
```

5) Resume another five iterations five thread 10GiB writes picking up where test #3 left off.

```
perl vsnapperf.pl seqwrite -filesize 10g -numVol 5 -iterations 10 -skipNum 5 -snapshot yes -overwrite=yes
```

6) Cleanup the volumes left behind from previous tests.

```
perl vsnapperf.pl cleanup -numVol 5
```

CLEANUP

The tool will automatically cleanup the files and volumes that are created in cases where a single iteration is run. In cases where more than one iteration is performed, cleanup is not performed to facilitate restore testing. The default cleanup behavior can be modified by specifying the -cleanup=yes|no option.

Also, clean can be perform later by running vsnapperf.pl with the cleanup action.

EXAMPLE

The following example was taken from the SPP blueprint Physical Supermicro reference.

```
# perl vsnapperf.pl seqwrite -fileSize 10g -numVol 5 -iterations 10 -skipNum 5
-snapshot yes -overwrite=yes

=====
: IBM Spectrum Protect Plus vSnap storage performance test (Program version 1.4b)
:
: Running with action seqwrite
:
: Number of iterations: 10
: Skipping iterations: 5
: Number of volumes: 5
:
: Block size: 256k
: File size: 10g
: Syncwrite: disabled
: Snapshots: yes
: Cleanup: no
: Random: /dev/randhigh
: Volumes: /vsnap/vpool14/SPPvSnapPerf1 /vsnap/vpool14/SPPvSnapPerf2
/vsnap/vpool14/SPPvSnapPerf3
/vsnap/vpool14/SPPvSnapPerf4 /vsnap/vpool14/SPPvSnapPerf5
:
=====
: Beginning I/O test.
: Run 6. Output: vsnapperf6.txt Status: COMPLETE Seconds: 97.10 Total GiB: 50.0 MiB/sec:
527.29 IOPS: 2109.16
: Run 7. Output: vsnapperf7.txt Status: COMPLETE Seconds: 97.81 Total GiB: 50.0 MiB/sec:
523.49 IOPS: 2093.95
: Run 8. Output: vsnapperf8.txt Status: COMPLETE Seconds: 97.49 Total GiB: 50.0 MiB/sec:
525.17 IOPS: 2100.68
: Run 9. Output: vsnapperf9.txt Status: COMPLETE Seconds: 97.62 Total GiB: 50.0 MiB/sec:
524.50 IOPS: 2097.99
: Run 10. Output: vsnapperf10.txt Status: COMPLETE Seconds: 98.03 Total GiB: 50.0 MiB/sec:
522.31 IOPS: 2089.22
: All iterations complete!
:
: Total processed (GiB): 250.0
: Minimum throughput (MiB/sec): 522.31
: Maximum throughput (MiB/sec): 527.29
: Average throughput (MiB/sec): 524.55

=====
===
CHANGE HISTORY
v1.2 Introduced cleanup functionality.
v1.4 Introduced improved pseudo-random data generation targeting closer to 2 to 1
compression with randmed.
v1.4b Bugfix for excessive dedup savings resulting from the improved pseudo-random data
generation and improved logging.
v1.4c Change default -bSize blocksize parameter from 256K to 128K
```

Appendix D. Protecting vSnap System Configuration

Note: To repair a vSnap server that has failed and must be replaced, please refer to the IBM Support document [How do I repair a failed vSnap server in an IBM Spectrum Protect Plus environment?](#)

A vSnap server consists of two sets of data

1. The vSnap pool (or *storage pool*) which is the logical organization of disks into a pool of storage which is consumed by the vSnap server component.
2. Configuration and metadata information which resides in the `/etc/vsnap` directory of the vSnap server

You can back up the configuration and metadata information for use cases where the vSnap pool data is intact and valid, but the configuration or metadata information is lost or not available. This can occur in these situations:

1. The vSnap server compute environment is lost but the storage is not. An example is a vSnap server that is running as a virtual machine and the storage backing the vSnap pool is on a pRDM disk. In this case the vSnap server virtual machine is lost but the data on the pRDM disk is still valid.
2. The vSnap server compute environment needs to be changed but the storage does not. An example is a vSnap server that is running in a virtual machine and the storage backing the vSnap pool is on a pRDM disk. In this case the vSnap server needs to be rebuilt or recreated on a new virtual machine but the data on the pRDM disk is still valid.

Backing up vSnap System Configuration

The backup procedure is based on the `vsnap system config backup` command which creates a tar (gzip compressed) file. The resulting file can then be securely copied to a central location, such as the IBM Spectrum Protect Plus server.

To back up the vSnap system configuration, issue the command

```
vsnap system config backup --outfile backup.tgz
```

(alternatively you can append a date stamp)

```
vsnap system config backup --outfile backup."$(date +%Y%m%d-%H%M%S)".tgz
```

Restoring vSnap System Configuration – Example Workflow

In this example, assume that IBM Spectrum Protect Plus was configured with a virtual vSnap server which was backed by a pRDM disk and that the virtual machine hosting the vSnap server was lost or damaged but that the data on the pRDM is still intact. Also assume that the administrator was periodically backing up the vSnap system configuration as outlined above. To recover the vSnap server you would have to follow these steps:

1. Deploy a new virtual vSnap server with the same configuration as the original vSnap server. For example, if you had deleted the 100 GB virtual disk on the original vSnap server virtual machine, you should also delete the same virtual disk on the new virtual machine.
2. Attach the original pRDM disk to the virtual machine.
3. Before powering on virtual machine it is critical to pause all job operations so that the new vSnap server is not contacted for any type of maintenance or workload ingest until it is recovered to its original state. In the **Schedule** tab of the **Jobs and Operations** panel, select the *Pause Schedule* action for each listed job
4. Power on the virtual machine and change the initial *serveradmin* password.
5. Copy the latest vSnap system configuration backup file from the central location where it is stored. For example, if you were backing up the vSnap system configuration files and copying the files to a location on the IBM Spectrum Protect Plus server.
6. Restore the vSnap system configuration using the ***vsnap system restore*** command:

```
vsnap system config restore --file <backup_file_name>
```

This command will perform all the necessary steps to recover the system configuration: initializing the vSnap server, stopping and starting the vSnap services, extracting the backup data to the appropriate location, recover encryption keys, etc. At the end of the restore command, you will see a message indicating “CONFIG RESTORE SUCCEEDED”.

7. Generate new SSH keys for each of the recovered vSnap server partners so that vSnap to vSnap operations (such as replication) can continue
 - a. From the newly recovered vSnap server issue the command: `vsnap partner show`
 - b. For each vSnap server partner listed in the output, issue the command:
`vsnap partner add --remote_addr <partner_ip> --local_addr <local_ip>`
8. It is now safe to restart all the scheduled operations. In the **Schedule** tab of the **Jobs and Operations** panel, select the *Release Schedule* action for each listed job
9. The vSnap server system configuration has now been recovered and the vSnap server will continue to function normally.

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 supply in any way it believes appropriate without incurring any obligation to you.

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

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

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

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

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

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products. This information 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 the names and addresses used by an actual business enterprise 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, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

VMware, VMware vCenter Server, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

Terms and conditions for product documentation Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein. IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed. You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations. IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> in the section entitled "Cookies, Web Beacons and Other Technologies," and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.