



# Reference Architecture: Lenovo Backup Solution with ThinkSystem DE/DM/DG Storage and Veeam for ThinkAgile and ThinkSystem Platforms

Last update: 06 September 2023  
Version 1.1

---

Reference Architecture for  
Veeam Backup solution on  
Lenovo servers and storage

---

Contains performance data and  
sizing recommendations for  
servers, storage, and networking

---

Describes deployment models  
for backup and recovery for  
ThinkAgile VX Hyperconverged  
Infrastructure

---

Contains detailed bill of materials  
for backup servers and storage

Chandrakandh Mouleeswaran  
Cristian Ghetau



# Table of Contents

---

<b>1</b>	<b>Introduction.....</b>	<b>4</b>
<b>2</b>	<b>Backup and Recovery Architecture.....</b>	<b>5</b>
2.1	Backup Technologies .....	6
2.2	Backup Methods .....	7
2.3	Immutable Backup .....	8
<b>3</b>	<b>Veeam Data Platform .....</b>	<b>9</b>
3.1	Veeam Backup and Recovery Components .....	9
3.2	Veeam Backup.....	11
3.2.1	Backup Methods and Retention Policies.....	11
3.3	Veeam Replication .....	13
3.4	Veeam Backup Copy .....	14
<b>4</b>	<b>Backup System Design .....</b>	<b>15</b>
4.1	Backup ThinkAgile VX Systems.....	15
4.2	Ransomware Protection.....	18
<b>5</b>	<b>Operational model .....</b>	<b>19</b>
5.1	Operational model scenarios .....	19
5.1.1	Primary Backup Storage with ThinkSystem DE/DM/DG Series .....	20
5.1.2	Dedicated Backup Server with ThinkSystem Servers .....	21
5.1.3	Shared Compute and Backup Server.....	22
5.1.4	Hyperconverged Backup Server.....	22
5.1.5	Host Attached Storage.....	22
5.2	Lenovo ThinkSystem DE/DM/DG Series Storage Arrays.....	22
5.2.1	Lenovo ThinkSystem DE Storage .....	22
5.2.2	Lenovo ThinkSystem DM Storage.....	23
5.2.3	Lenovo ThinkSystem DG Storage .....	23
5.2.4	IBM TS2270, TS2280 & TS2290 Tape Drives .....	24
5.2.5	IBM TS4300 Tape Library .....	25
5.2.6	Lenovo ThinkSystem FC SAN Switch .....	25
5.3	Backup Management Servers.....	25

5.4	Networking .....	26
5.5	Sizing Guidelines .....	26
5.5.1	Storage Capacity .....	26
5.5.2	Backup Window .....	28
5.5.3	Storage System Performance .....	29
5.5.4	Compute .....	30
5.5.5	Network Bandwidth .....	30
5.6	Backup and Recovery for VMware .....	31
5.6.1	Create Backup Job for VMware .....	31
5.6.2	Configure Virtual Machine Recovery .....	33
5.6.3	vCenter Integration .....	34
5.7	Storage Array Snapshot Integration .....	34
5.8	Deployment Model with ThinkSystem DE/DM/DG Storage .....	36
5.8.1	Onsite Backup with DE Storage and Shared ThinkAgile VX .....	36
5.8.2	Offsite Backup with DE Storage and Shared ThinkAgile VX .....	37
5.8.3	Offsite Backup with DE Storage and Dedicate ThinkSystem .....	38
<b>6</b>	<b>Appendix: Bill of materials .....</b>	<b>39</b>
6.1	BOM for Backup Storage .....	39
6.2	BOM for Backup Server .....	41
	<b>Resources .....</b>	<b>43</b>

# 1 Introduction

---

The intended audience for this document is technical IT architects, system administrators, and managers who are interested to design and deployment of backup and recovery solution for their infrastructure with Lenovo ThinkSystem and ThinkAgile hyperconverged servers and Lenovo ThinkSystem DE/DM/DG storage systems.

Ever growing infrastructure, devices, applications, technology, cloud, and data driven digital transformation journey for any business require robust data protection and disaster recovery strategy and solution. On-premises deployment is primary choice as it provides more security and control of data and infrastructure for mission critical applications. Backup and recovery bring many challenges in achieving Recovery Point Objectives (RPO) and Recovery Time Objectives (RTO) and it requires state of the art infrastructure and storage to achieve the goals without impacting business and operations. Enterprises are adopting new technologies, application deployment solutions and data science and AI/ML and it drives need to protect virtual machines, containers, applications, files and data from edge devices from ransomware and other cyber threats and avoid business disruptions.

Lenovo provides low cost and simplified deployment option to backup Lenovo ThinkSystem and ThinkAgile hyperconverged systems powered by ThinkSystem DE/DM/DG series storage and Veeam Backup and Recovery solution. This document gives backup capabilities of different Lenovo DE/DM/DG storage series and Tape systems and component model of Veeam Backup and Recovery suite. The document also provides the operational model for backup solution with Lenovo DE/DM/DG storage systems and Veeam Backup and Recovery to design and architect flexible data protection solution for Lenovo® hardware platforms such as ThinkSystem Servers and ThinkAgile VX/HX/MX hyperconverged servers.

The operational model presents performance measurements, considerations and sizing guidance, and some example deployment models. The last section contains detailed bill of material configurations for Lenovo storage and server to build backup infrastructure.

Developing right backup and testing strategy to attain higher level protection is a challenge and it requires proper planning for choosing right technology, backup methods, backup window and securing backup. This reference architecture considers many aspects and provides guidelines appropriately.

Hyperconverged systems are cost effective and provide enterprise class performance with different architecture like VMware vSAN Express Storage Architecture and protecting hyperconverged systems is similar to backup of traditional systems with shared storage and replicating to another hyperconverged system is not viable option due to underutilization of compute. Backup solution requires dedicated shared storage with Lenovo ThinkSystem DE/DM/DG series, and it provides flexible choices to achieve different service level objectives. This reference architecture currently focuses primarily on providing virtual machine backup solution for Lenovo ThinkAgile VX hyperconverged systems powered by VMware vSAN. The container backup with Kasten and integration with Veeam Backup Recovery is not in scope of this current version.

## 2 Backup and Recovery Architecture

Figure 1 shows common backup and recovery architecture spanning across multiple sites and different type of storage systems with Lenovo DE/DM/DG series storage systems and Veeam platform. The general recommended practice is to have more than one copy of backup to recovery from normal failover and disaster recovery scenarios. The backup storage architecture comprises of different media types all flash, hard drives and tape systems to provide flexible options and operationally cost efficient without impacting recovery point (RPO) and recovery time (RTO) objectives. The backup architecture can leverage either storage array-based replication or host based replication or combination of both for different applications backup performance requirements.

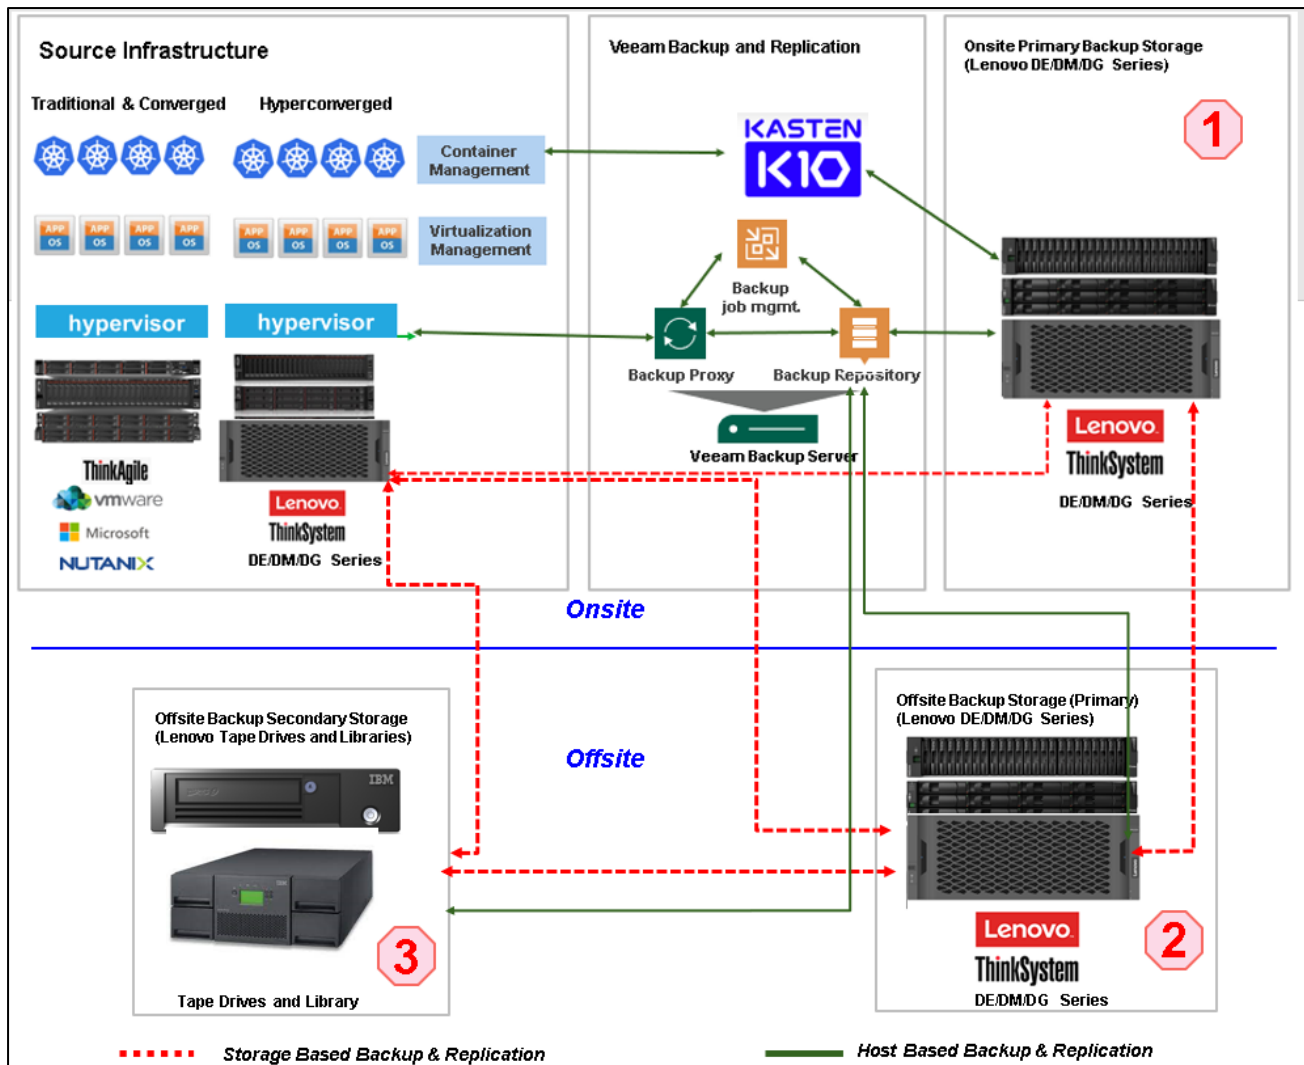


Figure 1: Onsite-Offsite Backup and Replication architecture with Lenovo DE/DM/DG Storage and Veeam

## 2.1 Backup Technologies

The source infrastructure can be comprised of converged systems, disaggregated systems and hyperconverged systems and different virtual machines, containers and applications do require either complete virtual disk backup or file backup or application level backup. The source infrastructure for different applications should consider different compatible backup technologies supported and the impact on their backup strategy to meet recovery point and time objectives. The table below explains different core backup technologies and pros and cons and these backup technologies would leverage server, storage and networking hardware capabilities to achieve efficiency. Table 1 shows different backup and replication technologies supported in Lenovo ThinkSystem storage systems and Veeam.

**Table 1: Lenovo ThinkSystem Storage and Veeam Backup and Replication technologies**

Technology	Description	Provided By	Use	Backup time / RPO	Recovery time / RTO
Storage Mirroring - Synchronous	The copy of source storage is continuously replicated to other storage and guarantees zero data loss.	Sync Mirror SnapMirror	Virtual machines, containers, applications, and databases	Immediate / Zero data loss	1-2 minutes
Storage Mirroring - Asynchronous	Create point in time copy and captures only the changes since last backup.	ASync Mirror	Virtual machines, containers, applications, and databases	Immediate / Depends on capture interval	1-2 minutes
Storage Volume Snapshots	Snapshot of a volume is taken in a particular point.	Storage Snapshot image	Virtual machines, containers, applications, and databases	Minimum 1 minute / Depends on capture interval	1-2 minutes
Virtual Machine Snapshots	It requires host intervention to take virtual machine snapshot	Veeam	Virtual machines, containers, applications, and databases	1-3 minutes / Depends on capture interval	1-3 minutes
Storage Snapshot replication - Asynchronous	Provide longer snapshot retention in an external pool for higher availability	ONTAP SnapVault	Files, Virtual Machines, containers, applications and databases.	Depends on replication interval	Minutes, depends on the application consistency validation

File Level Backup	Files within a virtual machine are backed up	Veeam	Files	Depends on capture interval	Depends on file size and network speed
Change Data Capture	It copies changes from application from the redo log before writing to disk	Application and database	databases	Immediate / zero data loss	1-2 minutes
Database Log Backup	Backup database transaction logs and recover from it	Veeam log backup	Databases	Immediate / zero data loss	1-2 minutes

## 2.2 Backup Methods

There are many different types of backup methods or types available and how to choose relevant to best protect data and recover appropriately during incidents and scenarios is a challenging one. Table 2 lists different backup methods followed in industry in general and Veeam does not support differential backup.

**Table 2: Backup Methods**

Method	Description	RPO	RTO	Efficiency
<b>Mirroring</b>	It requires disks and storage system and can be done at complete storage level or at volume level. It always has the latest copy of data and versioning is difficult.	Complete backup would be available and zero data loss.	Immediate	More storage space is required. No compute resources require from host side.
<b>Full Backup</b>	It can be done at virtual machine level or host level and complete copy of current state is backed up	Complete disk of a virtual machine is backed up.	Immediate	Takes more time to do full backup and requires more storage.
<b>Incremental Backup</b>	One full backup is required and then changes are incrementally backed up from last backup.	Due to less size, backup can be done often to increase RPO.	Takes time to rebuild image with all restore points	Need more compute resource to rebuild.

<b>Differential backups</b>	It considers full back up as the base and considers the changes every time from the previous full backup.	Backup time varies depending on change rate and it increases over a period time.	Better than incremental backup	Need more compute resource to rebuild.
<b>Synthetic Full backup</b>	It reconstructs the full backup image using all required incremental backups.			Reduces restoration time.

## 2.3 Immutable Backup

Protecting backup data from altering, deleting, changing, and corrupting from infrastructure and users is an important consideration in the solution. Immutable backup is critical for any organization and any loophole in the strategy will results severe loss to data and business as well. The storage system and backup software should be capable of protecting ransomware attacks or any source from being tampering or corrupting recovery points.



## 3 Veeam Data Platform

Veeam Data Platform is a collection of products and services that work together to provide a comprehensive data protection and management solution.

### 3.1 Veeam Backup and Recovery Components

Figure 2 is a layered view of the Veeam Backup and Replication solution that is mapped to the VMware virtualization infrastructure.

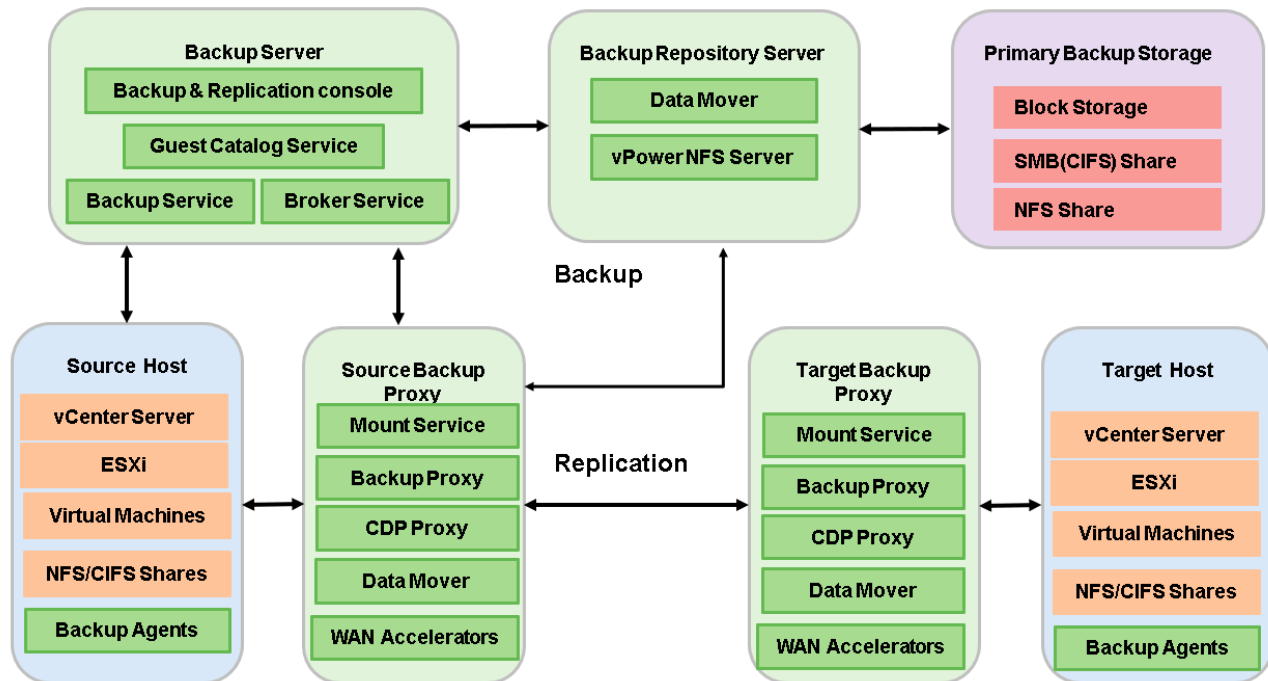


Figure 2: Component model with Veeam Backup and Replication

Veeam Backup and Replication solution infrastructure comprises the following components:

## **Veeam Backup and Replication Server**

Coordinates backup and replication tasks, controls resource allocation and replica job scheduling. It contains the following services and components.

**Backup Service** is a Windows service that coordinates all operations performed by Veeam Backup & Replication such as backup, replication, recovery verification and restore tasks

**Broker Service** interacts with the virtual infrastructure to collect and cache the virtual infrastructure topology

**Guest Catalog Service** manages guest OS file system indexing for VMs and replicates system index data files to enable search through guest OS files.

**Mount Service** mounts backups and replicas for file-level access, browsing the VM guest file system and restoring VM guest OS files and application items to the original location

**CDP Coordinator Service** communicates with vCenter, assigns continuous data protection (CDP) tasks and manages the infrastructure components involved

**Configuration Database** stores data about the backup infrastructure, jobs, sessions and other configuration data. The database instance can be located on a SQL Server

**Backup & Replication Console** provides the application user interface and allows user access to the application functionality.

**PowerShell Module** is an extension for Microsoft Windows PowerShell that adds a set of cmdlets to allow users to perform backup, replication and recovery tasks through the command-line interface to run custom scripts to fully automate operation

## **Veeam Backup Proxy**

It collects, transforms and transports VM data during the backup and replication process. It can be deployed on physical or virtual machines. It uses different services to perform the end-end operations.

**Data Mover:** It copies VM data from the source storage in one of transport modes. It filters out zero data blocks, blocks of swap files and blocks of excluded VM guest OS files, compresses and deduplicates VM data blocks and moves them to the target-side Veeam Data Mover.

## **Veeam Backup Repository**

The backup repository stores replica metadata that contains information on the read data blocks. Also stores backup metadata along with backup data.

## **Source**

The sources are VMware vCenter, ESXi, Physical servers, virtual machines and databases.

## **Target**

Backup repository for backup data and the host where virtual machine replicas will be created and maintained in the ready-to-start state

### Backup Storage

Block, file and object storage systems are used for backup.

### WAN Accelerators

WAN accelerators are optional components in the backup infrastructure. You can use WAN accelerators if you replicate VMs over a slow connection or over WAN.

## 3.2 Veeam Backup

Veeam Backup & Replication(VBR) uses image level backup at block level and copies complete virtual machine object and restore can be done at virtual machine level, virtual machine files and file level recovery within virtual machine. Veeam Backup & Replication creates backup jobs which connects with Veeam Backup Manager to establish connection with Data Movers running on the backup proxies and backup repositories. Data movers uses snapshots from vCenter server or ESXi host to transfer to backup storage and it uses Changed Block tracking(CBT) to retrieve changes only from the previous session. Figure 3 shows backup workflow in any Veeam deployment.

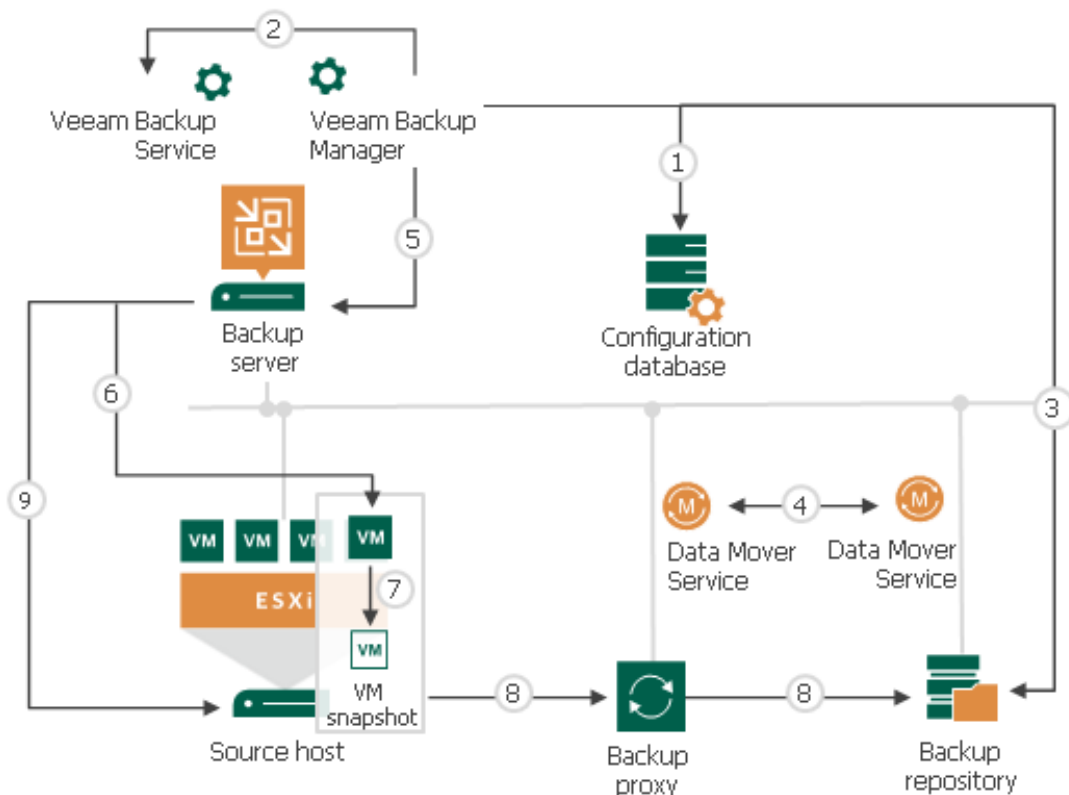


Figure 3: VM backup with Veeam Backup and Replication

### 3.2.1 Backup Methods and Retention Policies

Veeam supports different backup methods and retention policies to address all sort of backup requirements by efficiently planning for storage capacity and fast recovery. Table 3 lists different backup types supported by Veeam platform.

**Table 3: Veeam supported backup methods**

Backup Method	Description
<p>Forever forward incremental (FFI)</p> 	<p>The backup chain consists of the first full backup file (VBK) and a set of forward incremental backup files (VIBs). It stores only one full backup file and removes incremental backup files once the retention period is exceeded. It is the fastest compared to other methods because the first available restore point is always a full backup.</p>
<p>Forward incremental (FI)</p> 	<p>The backup chain consists of multiple full backup files (VBKs) and sets of forward incremental backup files (VIBs) following each full backup file. It is split into shorter series so this lowers the chances of losing the backup chain completely. Restore from backup files created using the FI method is the most optimal in time compared to other methods</p>
<p>Reverse Incremental Backup*</p> 	<p>The backup chain consists of the last full backup file (VBK) and a set of reverse incremental backup files (VRB) preceding it. The most recent restore point in the backup chain is always a full backup and it makes restoring virtual machine immediately without any additional processing.</p>
<p>Active Full Backup</p> 	<p>It compresses and deduplicates whole VM and stores it to the full backup file (vbk). The latest full backup file is new starting point for all incremental backups.</p>
<p>Synthetic Full Backup</p>	<p>It is identical to a regular full backup and it synthesizes a full back up from data you already have in the backup repository rather than from source VM</p>

\* Reverse Incremental Backup is supported technically in Veeam currently, but it may not be available in future releases.

### Short Term Retention Policy

Whenever the number of restore points or days is exceeded, the backup chain makes room for most recent restore point. The full backup also moves one step forward with the next incremental backup in the backup

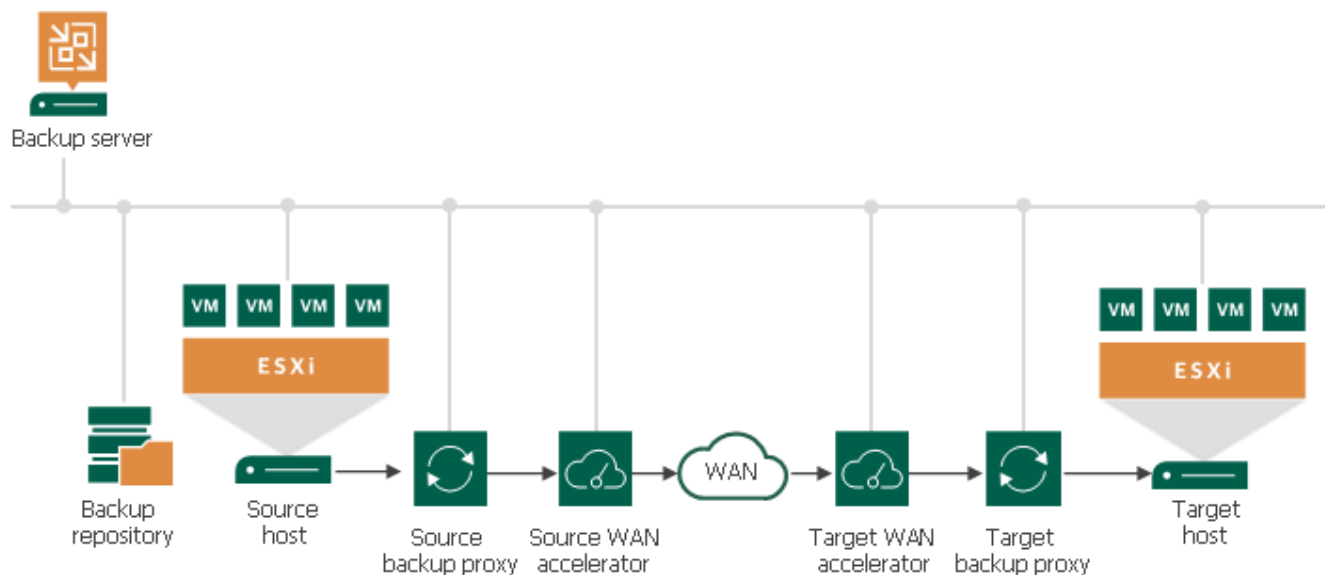
chain. The earliest incremental backup file is deleted immediately when it is outdated, and the full backup chain will be removed only after the last incremental backup file in the chain becomes outdated. All the incremental backup methods apply relevant short term retention policy based on the restore points and full backup requirements.

### Long Term Retention Policy (GFS)

The long-term or Grandfather-Father-Son (GFS) retention policy allows to store for longer period and backup can be done weekly, monthly, and yearly. Veeam Backup and Replication follows different algorithms based on different type of GFS policies to take full backup based on the retention period.

## 3.3 Veeam Replication

The replication creates exact copy of the source virtual machine on the target host. The virtual machine replica are in ready to start state and it helps to achieve minimum recovery point objective(RPO). Veeam Backup and Replication leverages vSphere snapshots to replicate virtual machines and incrementally replicate without compression for subsequent changes. The replication can be done at onsite for high availability and offsite for disaster recovery scenarios. Figure 4 shows typical replication setup with Veeam backup and replication.



**Figure 4: VM Replication with Veeam Backup and Replication**

**Continuous data protection (CDP)** is a technology to protect VMware virtual machines when data loss for seconds or minutes is unacceptable. CDP does not use vSphere snapshots and uses vSphere APIs for I/O filtering (VAIO) to create short term restore points which helps achieve zero loss.

## 3.4 Veeam Backup Copy

Backup copy or secondary backup storage helps to recover from disaster recovery scenarios and when any issue with primary backup storage. Veeam Backup copy contains the same data and format as the primary backup, and it is created using jobs and retention settings can be applied with desired number of restore points. All backup jobs can be linked to backup copy jobs to store on secondary backup storage.

## 4 Backup System Design

Designing robust backup solution is based on the Recovery Time Objectives (RPO) and Recovery Time Objectives(RTO) and backup storage infrastructure must be sized to meet these objectives. These objectives are purely depending on applications or workloads and compute and storage need to be assessed appropriately. The following steps will provide guidance to make decision and refer operational model for detailed information.

1. Decide RPO and RTO for applications and virtual machines
2. Decide number of onsite and offsite copies
3. Group applications based on RPO/RTO
4. Decide Backup Method and Backup Window
5. Decide Total storage capacity and Retention Period
6. Choose VM snapshot vs Storage Snapshot
7. Choose storage system and drives to meet RPO/RTO
8. Define volumes for different backup groups and RAID levels
9. Define number of backup jobs
10. Evaluate sharing compute and network from source infrastructure
11. Decide compute for Veeam backup components

### 4.1 Backup ThinkAgile VX Systems

Lenovo ThinkAgile VX Integrated Systems and Certified nodes are hyperconverged systems powered by VMware vSAN. VMware vSAN supports two different types of architectures to meet different performance requirements and both do not require external storage and leverages ethernet connectivity across nodes in the cluster. vSAN supports RAID1(mirroring) and RAID-5/6 erasure coding protection level and vSAN storage can be accessed from host outside cluster using iSCSI and file shares (SMB and NFS).vSAN uses FTT(Failure to Tolerate) mechanism to provide protection and the usable capacity varies based on this value. The usable capacity of vSAN storage depends on the combination of RAID level and FTT values. It means that the backup storage requirement needs to be estimated based on the virtual machine actual disk requirements but should not be based on vSAN total capacity.

**Table 4: VMware vSAN Architectures**

Type	Ethernet	Storage	Maximum Capacity per node
Original Storage Architecture	10GbE/25GbE	All Flash - Cache and capacity tier with Flash  Hybrid - Cache tier is with flash and capacity tier with hard drives	35x7.68 TB
Express Storage Architecture	25GbE/100GbE	Single tier with NVMe drives	24x6.68TB

Lenovo provides ThinkAgile VX systems in different sizes to support wide variety of workloads and scalability. The latest system does support fibre channel to use for backup with Lenovo DE/DM systems and since vSAN cluster nodes are with ethernet connectivity, the backup can be done leveraging iSCSI without investing additional hardware. With the latest generation processors, the backup server and services can be run on the same vSAN cluster without impacting application workload performance. Lenovo ThinkSystem DE/DM storages support direct host attach connectivity and vSAN nodes can be connected to storage system directly for onsite backup. Table 5 shows list of ThinkAgile VX systems and capacity supported in the source infrastructure.

**Table 5: Lenovo ThinkAgile VX system models**

<b>VX System</b>	<b>vSAN Support</b>	<b>Base System</b>	<b>CPU</b>	<b>Max drives</b>	<b>Max Possible Capacity</b>	<b>Fibre Channel</b>
<b>VX2330 Appliance</b>	Hybrid All Flash	ThinkSystem SR630 V2	Intel Xeon 3rd Gen	4	1 disk group 3 drives	No
<b>VX3330 Appliance</b>	Hybrid All Flash	ThinkSystem SR630 V2	Intel Xeon 3rd Gen	12	2 disk group 10 drives	No
<b>VX7330-N Appliance</b>	Hybrid All Flash	ThinkSystem SR630 V2	Intel Xeon 3rd Gen	12	2 disk group 10 drives	No
<b>VX3331 Certified Node</b>	Hybrid All Flash	ThinkSystem SR630 V2	Intel Xeon 3rd Gen	12	2 disk group 10 drives	No
<b>VX2375 Integrated System</b>	Hybrid All Flash	ThinkSystem SR645	AMD EPYC 3rd Gen	4	1 disk group 3 drives	No
<b>VX3375 Integrated System</b>	Hybrid All Flash	ThinkSystem SR645	AMD EPYC 3rd Gen	12	2 disk group 10 drives	No
<b>VX7375-N Integrated System</b>	All Flash	ThinkSystem SR645	AMD EPYC 3rd Gen	12	2 disk group 10 drives	No
<b>VX3376 Certified Node</b>	Hybrid All Flash	ThinkSystem SR645	AMD EPYC 3rd Gen	12	2 disk group 10 drives	No
<b>VX3530-G Appliance</b>	Hybrid All Flash	ThinkSystem SR650 V2	Intel Xeon 3rd Gen	24	3 disk group 21 drives	No



<b>VX5530 Appliance</b>	Hybrid All Flash	ThinkSystem SR650 V2	Intel Xeon 3rd Gen	16	2 disk group 14 drives	No
<b>VX7530 Appliance</b>	Hybrid All Flash	ThinkSystem SR650 V2	Intel Xeon 3rd Gen	40	5 disk group 35 drives	No
<b>VX7531 Certified Node</b>	Hybrid All Flash	ThinkSystem SR650 V2	Intel Xeon 3rd Gen	40	5 disk group 35 drives	No
<b>VX3575-G Integrated System</b>	Hybrid All Flash	ThinkSystem SR665	AMD EPYC 3rd Gen	24	3 disk group 21 drives	No
<b>VX5575 Integrated System</b>	Hybrid All Flash	ThinkSystem SR665	AMD EPYC 3rd Gen	16	2 disk group 14 drives	No
<b>VX7575 Integrated System</b>	Hybrid All Flash	ThinkSystem SR665	AMD EPYC 3rd Gen	35	5 disk group 30 drives	No
<b>VX7576 Certified Node</b>	Hybrid All Flash	ThinkSystem SR665	AMD EPYC 3rd Gen	35	5 disk group 30 drives	No
<b>VX630 V3 IS</b>	Hybrid All Flash	ThinkSystem SR630 V3	Intel Xeon 4th Gen	12	2 disk group 10 drives	Yes
<b>VX630 V3 CN</b>	Hybrid All Flash	ThinkSystem SR630 V3	Intel Xeon 4th Gen	12	2 disk group 10 drives	Yes
<b>VX650 V3 IS</b>	Hybrid All Flash	ThinkSystem SR650 V3	Intel Xeon 4th Gen	32	4 disk group 28 drives	Yes
<b>VX650 V3 CN</b>	Hybrid All Flash	ThinkSystem SR650 V3	Intel Xeon 4th Gen	32	4 disk group 28 drives	Yes
<b>VX650V3 DPU IS</b>	Hybrid All Flash	ThinkSystem SR650 V3	Intel Xeon 4th Gen	32	4 disk group 28 drives	Yes
<b>VX650V3 DPU CN</b>	Hybrid All Flash	ThinkSystem SR650 V3	Intel Xeon 4th Gen	32	4 disk group 28 drives	Yes
<b>VX655 V3 Integrated System</b>	Hybrid All Flash	ThinkSystem SR655 V3	AMD EPYC 4th Gen	40	5 disk group 35 drives	Yes

<b>VX655 V3 Certified Node</b>	Hybrid All Flash	ThinkSystem SR655 V3	AMD EPYC 4th Gen	40	5 disk group 35 drives	Yes
<b>VX665 V3 Integrated System</b>	Hybrid All Flash	ThinkSystem SR665 V3	AMD EPYC 4th Gen	40	4 disk group 28 drives	Yes
<b>VX665 V3 Certified Node</b>	Hybrid All Flash	ThinkSystem SR665 V3	AMD EPYC 4th Gen	40	4 disk group 28 drives	Yes

## 4.2 Ransomware Protection

Lenovo ThinkSystem DE series and Veeam Hardened Linux Repository together can be leveraged to implement native Linux XFS filesystem capabilities for disk-based immutable backups for primary backup storage and prevent from malicious or accidental backup deletion. The solution meets the requirements for non-rewritable, non-erasable storage as specified by SEC 17a-4(f), FINRA 4511(c) and CFTC 1.31(c)-(d) regulations. Figure 5 shows how backup files are protected from outside world.

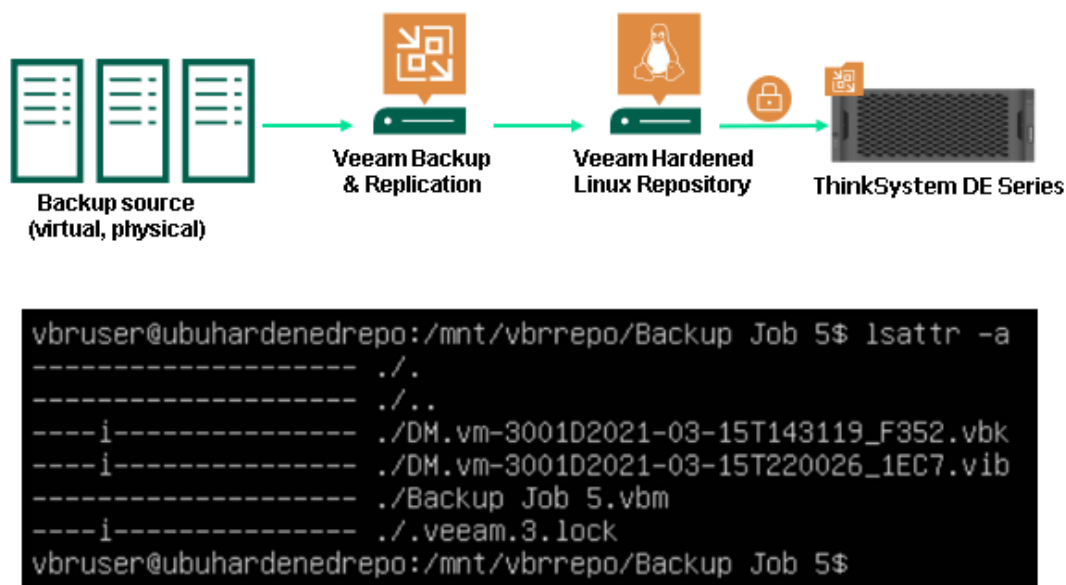


Figure 5: Immutable backup with Linux hardened repository

## 5 Operational model

This section describes the options for mapping the logical components of backup and replication solution onto hardware and software. The “Operational model scenarios” section gives an overview of the available mappings and has pointers into the other sections for the related hardware. Each subsection contains recommendations on how to size for that particular hardware, and a pointer to the BOM configurations that are described in section 6 on page 39. The last part of this section contains some deployment models for example customer scenarios.

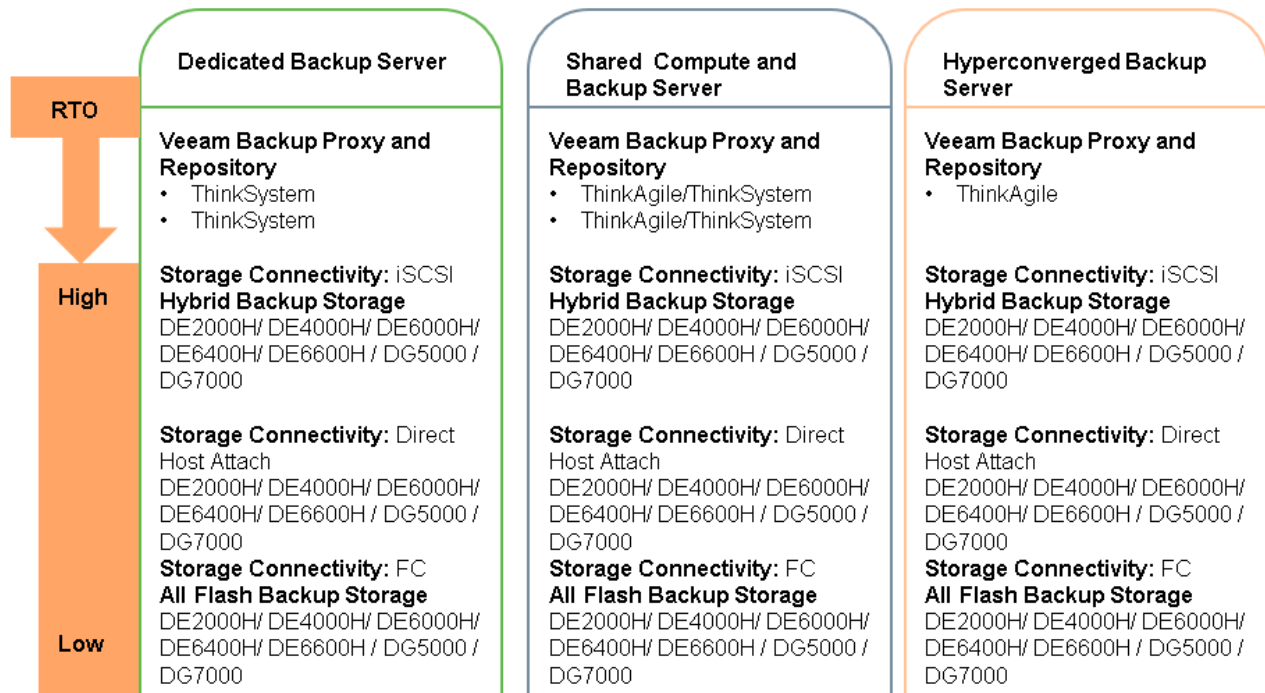
### 5.1 Operational model scenarios

Figure 6 shows the following operational models (solutions) in Lenovo Backup solution with Veeam for enterprise and small-medium business (SMB) with dedicated and shared deployment and different performance classes. Backup is necessary for mission critical applications for any business and having more copies on different locations and different media prevents data loss from circumstances ranges from planned maintenance to disaster scenarios. Having more copies are good for large business and it brings additional operational costs for backup storage and compute depends on RPO and RTO objectives. Table 6 shows typical backup requirements for SMB and enterprises.

**Table 6: Backup for different business sizes**

Type of Business Size	Total Backup Size	Number of Mission Critical Applications	Number of Sites	Backup Model
Small	Up to 100 TB	1-2	1	1 Onsite copy + 1 Archival copy
Medium	101-500TB	2-5	1-2	1 Onsite copy + 1 Offsite copy + 1 Archival Copy
Enterprise	500TB+	5-10 or more	3+	1 Onsite + 1 or more Offsite copies + 1 Archival Copy

The lower RTO improves availability by providing faster recovery and it does require more compute, faster connectivity, and all flash storage systems. RTO is a critical factor for mission critical applications irrespective of company size. Dedicated compute for backup server for small deployments is not required for small and medium deployment and the source infrastructure should be able to accommodate and scale to support backup software computing requirements. The performance of the source storage systems also important to achieve lower RTO. Figure 6 shows operation model with different Lenovo servers and storage systems to meet backup RTO objectives.



**Figure 6: Backup Operational model scenarios**

### 5.1.1 Primary Backup Storage with ThinkSystem DE/DM/DG Series

Lenovo ThinkSystem DE/DM/DG series provide various options to support backup storage requirements. Usually, HDDs are recommended for backup due to the cost factor. However, SSDs and NVMe drives can be used for deployments which require shorter recovery time. The primary backup storage can coexist on the same data center, but it is recommended to keep one backup copy on separate physical location than the source environment. Table 7 below lists Lenovo storage systems support for backup solutions.

**Table 7: Lenovo ThinkSystem Storage for SMB Backup**

Backup Storage	All Flash	SnapMirror	Protocol	Expansion	Connectivity	Maximum Capacity
DE2000H 2U 12 LFF	No	No	Block		Fibre Channel, iSCSI	1.47 TB
DE4000H 2U 12 LFF / 4U 60LFF	No	No	Block	DE120S 2U12 LFF, DE600S 4U60	Fibre Channel, iSCSI	3.07 PB
DE6000H 2U 12 LFF / 4U 60 LFF	No	No	Block	DE120S 2U12 LFF, DE600S 4U60	Fibre Channel, iSCSI	4.2 PB

<b>Backup Storage</b>	<b>All Flash</b>	<b>SnapMirror</b>	<b>Protocol</b>	<b>Expansion</b>	<b>Connectivity</b>	<b>Maximum Capacity</b>
DE2000H 2U 12 LFF	No	No	Block		Fibre Channel, iSCSI	1.47 TB
DE4000H 2U 12 LFF / 4U 60LFF	No	No	Block	DE120S 2U12 LFF, DE600S 4U60	Fibre Channel, iSCSI	3.07 PB
DM3000H 2U 12LFF	No	Yes	Block/NFS		Fibre Channel, iSCSI, Ethernet	2.3 PB Cluster (27.6 PB NAS or 13.8 PB SAN)
DE4000F 2U24 SFF	Yes	No	Block		Fibre Channel, iSCSI, Ethernet	1.440 PB
DE6000F 2U24 SFF	Yes	No	Block		Fibre Channel, iSCSI, Ethernet	1.84 PB
DM5000F 2U24 SFF	Yes	Yes	Block		Fibre Channel, iSCSI, Ethernet	2.2 PB Cluster (26.5 PB NAS or 13.2 PB SAN)
DG5000	QLC	Yes	Block/NFS	DM240N	Fibre Channel, iSCSI, Ethernet	8.8 PB
DG7000	QLC	Yes	Block/NFS	DM240N	Fibre Channel, iSCSI, Ethernet	17.6 TB

### 5.1.2 Dedicated Backup Server with ThinkSystem Servers

For the enterprise performance and availability requirements, dedicated Lenovo ThinkSystem SR630 V2 nodes are recommended to host Veeam backup proxy and repository server components. In many circumstance the source rack might not have enough space and also source virtual machines and applications need to reserve for compute overhead and deploying Veeam components brings additional management complexities during maintenance and upgrade operations. So, it is recommended to separate Veeam components to dedicated servers to isolate from source and to backup different workloads running across converged and hyperconverged compute clusters in the data center. The ThinkSystem servers require connectivity to external storage for Veeam management components and it can leverage backup storage as well.

### **5.1.3 Shared Compute and Backup Server**

For small and medium size source environments, it is expensive to have dedicated servers to host Veeam backup components. Also, many of the source workloads are being deployed on hyperconverged servers and having separate shared storage for only Veeam management components is not a viable solution. The Veeam backup components can coexist on the source cluster without impacting source applications. Deploying Veeam proxy and backup server closer to source virtual machines provide better performance. This deployment is well suited for storage array-based snapshots where hosts are freed from resources required for virtual machine based snapshots.

### **5.1.4 Hyperconverged Backup Server**

Veeam backup components require considerable amount of compute, memory and storage and it increases when the backup storage and number of backup jobs increase. Veeam provides additional features such as orchestrator, monitoring platform and container backup support and it requires additional resources. For 200TB backup requirements, the complete Veeam management stack mentioned above requires 32 cores, 64 GB and 3TB storage and it can be managed with either 2 node ThinkAgile VX 2000 VMware vSAN systems or ThinkAgile HX 2000 Nutanix series or ThinkAgile MX 1021 AzureStack HCI systems are perfect fit and these system can also host other infrastructure management software such as vCenter, vCloud Foundation and others.

### **5.1.5 Host Attached Storage**

Lenovo ThinkSystem DE/DM series supports connecting storages to local port of the hosts. It reduces additional switches and hops in the network and it is an ideal solution for onsite backup model. This design is perfect fit for dedicated backup server with ThinkSystem servers and hyperconverged systems and DE/DM storage can be placed on the same rack or within reasonable distance within data center if feasible. The support is available for 1GbE/10GbE/100GbE and also 4/8/16/32 Gb FC connectivity depends on the storage systems.

## **5.2 Lenovo ThinkSystem DE/DM/DG Series Storage Arrays**

In this section we provide an overview of the newly introduced Lenovo DE and DM series storage arrays. For backup requirements the DE/DM series storage provides a wide range of options, supporting high-capacity, high performance and scale-out options as well as rich set of software capabilities such as storage management, block/file level support, encryption, deduplication, and so forth.

### **5.2.1 Lenovo ThinkSystem DE Storage**

Lenovo ThinkSystem DE Series Storage Arrays are SAN storage systems that are designed to provide performance, simplicity, capacity, security, and high availability for medium to large businesses. ThinkSystem DE Series Storage Arrays deliver hybrid and all flash storage with enterprise-class storage management capabilities and a wide choice of host connectivity options, flexible drive configurations, and enhanced data management features.



**Figure 7: Lenovo ThinkSystem DE Storage**

### 5.2.2 Lenovo ThinkSystem DM Storage

The Lenovo ThinkSystem DM storage arrays provide a unified storage solution to manage all your block-and-file workloads on one array. DM series comes in both Hybrid Flash and All Flash configurations. NVMe support drives to achieve superior performance for mission critical workloads.

DM Series Hybrid Flash systems simplify the task of managing growth and complexity by delivering high performance, supporting a broad range of unified workloads, and seamlessly scaling of performance and capacity. For growing organizations that are concerned about maximizing performance and tiering capabilities to meet challenging IT needs, ThinkSystem DM Series Hybrid Flash systems are the perfect choice.



**Figure 8: Lenovo ThinkSystem DM Storage**

### 5.2.3 Lenovo ThinkSystem DG Storage

The Lenovo ThinkSystem DG series supports high capacity all flash QLC storage to provide increased performance over HDD arrays with reduced cost, rack space and power consumption. Lenovo ThinkSystem DG system supports file, block and object storage and includes advanced features for hybrid cloud management and ransomware protection. DG Series delivers optimal price per performance for verity of workloads and seamlessly scaling of capacity without compromising performance.

Lenovo ThinkSystem DG5000 models are 2U rack-mount controller enclosures that include two controllers, 128 GB RAM and 16 GB battery-backed NVRAM (64 GB RAM and 8 GB NVRAM per controller), and 24 SFF hot-swap drive bays (2U24 form factor). A single ThinkSystem DG5000 Storage Array scales out to 48

QLC SSDs with the attachment of one Lenovo ThinkSystem DM240N 2U24 SFF Expansion Enclosure. Up to 12x DG5000 Storage Arrays can be combined into a clustered system in a NAS environment, or up to 6x DG5000 Storage Arrays can be combined into a clustered system in a SAN environment.



**Figure 9: Lenovo ThinkSystem DG5000 Storage**

A single ThinkSystem DG7000 Storage Array consists of the 4U rack-mount controller enclosure and one or more expansion enclosures. The controller enclosure includes two controllers, 256 GB RAM (128 GB RAM per controller), and 32 GB battery-backed NVRAM (16 GB NVRAM per controller). The attachment of the Lenovo ThinkSystem DM240N 2U24 SFF Expansion Enclosures to the controller enclosure provides scalability up to 96 QLC NVMe.



**Figure 10: Lenovo ThinkSystem DG7000 Storage**

For more information and specifications of the Lenovo ThinkSystem DE,DM and DG series storage arrays, see the following Lenovo Press publications:

[Lenovo ThinkSystem DM Series Unified Storage Arrays](#)

[Lenovo ThinkSystem DE Series Storage Arrays](#)

[Lenovo ThinkSystem DG Series](#)

#### **5.2.4 IBM TS2270, TS2280 & TS2290 Tape Drives**

The IBM TS2270, TS2280, and TS2290 Tape Drives are excellent tape storage solutions for small and midsized businesses that require backup and low-cost archival data storage. These tape drives incorporate the latest generation of industry leading LTO technology that helps organizations to handle the growing data demands for backup and archival. Refer Lenovo datasheet for more details [IBM TS2270, TS2280 & TS2290 Tape Drives](#)



The TS2270 Tape Drive provides a physical capacity of 6TB of native storage capacity (15TB with 2.5:1 compression) and a transfer rate of up to 300MBps with 6Gbps SAS interface connectivity. The TS2280 Tape Drive claims a physical capacity of 12TB (up to 30TB with 2.5:1 compression).

The TS2290 Tape Drive provides even higher capacity than the previous generation, with a native storage capacity of 18TB (up to 45TB per cartridge with 2.5:1 compression).

### 5.2.5 IBM TS4300 Tape Library

IBM TS4300 Tape Library is a high-density, highly scalable, easy-to-manage solution designed to keep data securely stored long-term, while helping reduce the costs associated with data center space and utilities. The supported tape drives are IBM LTO Ultrium 9, 8, 7, or 6 Tape Drives with SAS or Fibre Channel host connectivity interfaces and with LTO Ultrium 9 technology, the TS4300 Tape Library provides up to 5.04 PB of native backup storage. Refer Lenovo press paper for more details [IBM TS4300 Tap Library](#)

### 5.2.6 Lenovo ThinkSystem FC SAN Switch

Lenovo ThinkSystem Gen6/Gen7 FC SAN switches are designed with storage connectivity technologies from Brocade and these switches are ultra-dense, 1U switch and supports 4, 8, 16, 32 and 64 Gbit/sec capabilities. Fiber channel connectivity provides better performance for backup and recovery and dedicated switches are recommended for large environments and the ports can be shared with source infrastructure as well. Table 8 shows different models available in Lenovo FC SAN switch portfolio.

**Table 8: Lenovo ThinkSystem FC SAN Switches**

Switch Model	Ports and Bandwidth
<a href="#">Lenovo ThinkSystem DB610S FC SAN Switch</a>	24 ports 4/8/16/32 Gbps
<a href="#">Lenovo ThinkSystem DB620S FC SAN Switch</a>	24/48 ports 4/8/10/16/32 Gbps
<a href="#">Lenovo ThinkSystem DB630S FC SAN Switch</a>	48/96 ports 4/8/10/16/32 Gbps
<a href="#">Lenovo ThinkSystem DB720S FC SAN Switch</a>	24 ports 4/8/10/16/32/64 Gbps
<a href="#">Lenovo ThinkSystem DB730S FC SAN Switch</a>	48 ports 4/8/10/16/32/64 Gbps

## 5.3 Backup Management Servers

For medium to large environment, the backup management components should have dedicated compute servers so that it will not impact application workloads performance if it is sharing the compute. However, if there are many clusters to be backed up and backup window is not impacting application performance, then shared compute can be used and it is recommended to keep proxy closer to applications which require back up. Veeam Backup and Replication server components can run on ESXi along with virtualization and cloud management software such as vCenter server and vCloud Foundation. The table below shows minimal

configuration requirements for Veeam components. Table 9 shows the minimal system requirement to install Veeam backup and server components.

**Table 9: Veeam Management components system requirements**

Management VM	Virtual processors	System memory	Storage	OS	HA needed
Veeam Backup and Replication Server	8	8 GB	60 GB	Windows Server	No
Backup Proxy	8	24 GB	60 GB	Windows Server or Linux	Yes
Backup Repository Server	12	24 GB	100 GB	Windows Server or Linux	No
Veeam ONE	4	8 GB	50 GB	Windows Server	No
Veeam Recovery Orchestrator	16	16 GB	275 GB	Windows Server or Linux	No

## 5.4 Networking

The networking for backup can use 8 or 16 Gbps FC, or 10 GbE iSCSI or Ethernet connections. For offsite backup it is highly recommended to have high speed connectivity to avoid latency. It is recommended to isolate backup traffic, Veeam management and application data traffic. When Veeam components are co-hosted on the source infrastructure, the networking bandwidth need to be considered for backup and it should not impact the application performance.

## 5.5 Sizing Guidelines

The backup infrastructure consists of primary and secondary backup storage systems and compute required for running Veeam backup components.

### 5.5.1 Storage Capacity

The factors below are considered to estimate total back up storage consumption.

**Source Data:** The source can be virtual machines disks, images, files, and databases and the actual size of each protected component needs to be considered for full backup. In hyperconverged environments, the usable capacity is calculated and varies based on RAID and FTT value (in VMware vSAN) and it is

recommended to estimate based on the virtual machine actual disk requirements instead of considering total storage on the source environment.

**Number of Virtual Machines:** The size of virtual machines cannot be expected same in large environments. However, it can be grouped based on the same disk size to estimate target backup capacity easily.

**Daily Change Rate:** Estimation daily change rate for each protected source helps to size storage appropriately. It can be started with 5-10% conservatively for any applications and then it must be monitored and revise sizing appropriately. The more the change rate, it does require frequent backup to avoid data loss.

**Retention Scheme:** How many copies are retained for how long an import factor. Many backup can be designed to have latest copy but many critical applications require more retention period based on industry and regulatory requirements. It is recommended to have multiple retention points created daily, weekly, monthly, and yearly.

Compression Ratio: The backup storage footprint can be reduced by increasing compression ratio and it requires processing capacity based on the compression level chosen.

Table 10 lists the backup storage needed for different size of source environment. The sizing is based on forward incremental with weekly full backup method with 50% compression ratio.

**Table 10: Primary Backup Storage Sizing - Forwards Incremental with Weekly Full**

Source Capacity	Daily change rate (%)	Yearly growth (%)	Growth Scope	Retention Points	Reduction(%)	Estimated Backup Storage	Estimated Backup Copy
10 TB	10%	10%	3 years	Daily 12 Weekly 0 Monthly 0 Yearly 0	50%	20.39 TB	20.39 TB
10 TB	10%	10%	3 years	Daily 12 Weekly 2 Monthly 2 Yearly 2	50%	41.05 TB	41.05 TB
20 TB	10%	10%	3 years	Daily 12 Weekly 2 Monthly 2 Yearly 1	50%	80.14 TB	80.14 TB
50 TB	10%	10%	3 years	Daily 12 Weekly 2 Monthly 2 Yearly 1	50%	192.68 TB	192.68 TB

100 TB	10%	10%	3 years	Daily 12 Weekly 2 Monthly 2 Yearly 1	50%	379.28 TB	379.28 TB
250 TB	10%	10%	3 years	Daily 12 Weekly 2 Monthly 2 Yearly 1	50%	924.8 TB	924.8 TB

### 5.5.2 Backup Window

Choosing right backup window is critical for any environment as it is tightly coupled with recovery point and time objectives and it can impact running applications if it is not planned well. It has be designed considering RPO/RTO objectives for each application and virtual machines. Each backup method has it own advantages and drawbacks and we need to design purely based on the recovery objectives and tolerance factors. Many businesses can follow 4-8 hours night backup window for normal applications, but critical applications require continuous backup along with multiple backup copies across different locations. Table 11 shows different design factors which should be taken care while planning backup window.

**Table 11: Backup Window influencing factors**

Factors	Design Considerations
Full Backup	Full back up takes more time, and it depends on number of full backup required Daily/Weekly/Monthly. It will occupy more storage space as well. Many critical applications require full back up during weekdays. If the source is large, then backup window should be increased to fit.
Incremental Backup	Since full backup takes more time and it is advisable to use combination of full and incremental back up. Incremental back up require less space and consumes less resources and can be taken frequently.
Recovery Point Objective (RPO)	Having frequent backup and restore points help to improve RPO. Restoring from full back is quicker than recovering from incremental backups. During recovery time the RPO would reduce as backup cannot be done when restore is in progress and restore takes priority than backup.
Recovery Time Objective (RTO)	Restoring backup is dependant on the backup size and the storage media and network speed. Storage array-based backups and restore

	are faster than software-based recovery which involves traverse through backup chain to rebuild required restore point.
Parallel Backup Proxies	Configuring more concurrent backup proxies to run backup for different application would reduce the backup time and hence backup window.
Storage Array based replication	Storage array based replication provide faster backup and recovery and it can reduce backup window.
Storage System Performance	Storage systems maximum throughput capacity needs to be considered. Either we have to scale out storage systems to support within backup window or adjust backup window to fit storage system performance.

### 5.5.3 Storage System Performance

The selection of storage model, drive, host connectivity and storage controller should be based on RPO and RTO objectives and backup window. If RTO is high, then it is recommended to use SSD and Fibre Channel connectivity and storage array-based snapshots for faster backup and recovery. The backup NAS storage is recommended for applications and workloads require to access large number of files. The storage system performance may degrade after 80% utilization and the storage capacity should be sized considering headroom and performance limitations.

The storage selection should address RPO/RTO objectives, recovery speed and backup window. The storage system should support enough concurrent backup jobs without impacting backup window. The backup operation is sequential and full backup takes longer time based on the disk size of source virtual machine or applications. All Flash is good for mission critical applications which require RTO in few minutes range. Storage array-based snapshot and replication is recommended in these scenarios as it avoids additional CPU cycle for source system and network latencies. Also, onsite backup provides better recovery time than offsite backup due to network latency. Figure 11 shows different DE systems supported throughput.

Operation	DE200H	DE4000H	DE6000H	DE6400H	DE6600H	DE4000F	DE6000F	DE6400F	DE6600F
Sequential Read	3 GB/s	9.2 GB/s	21 GB/s	20 GB/s	44 GB/s	9.2 GB/s	21 GB/s	20 GB/s	44 GB/s
Sequential Write	0.9 GB/s	2.7 GB/s	7 GB/s	7 GB/s	12.5 GB/s	2.7 GB/s	9 GB/s	7 GB/s	12.5 GB/s

**Figure 11: Lenovo DE storage system throughput**

Lenovo ThinkSystem DG series are with QLC all flash drives and it provides better price, performance and power consumption over HDD based storage systems. DG5000 and DG7000 systems are optimal solution for backup to satisfy RPO and RTO objectives and capacity requirements. Table 12 below shows comparison of Vdbench performance on DM3000H and DG5000 with 512k and 1MB block sizes. The DG5000 provides 4:1 compression ratio and can reduce backup storage footprints with improved throughput and much reduced sub seconds latency.

**Table 12: Vdbench performance results for DG5000 and DM3000H**

Storage	DG5000		DM3000H	
Number of Drives	24x 15.3TB QLC SSD		12x 4TB HDDs + DM600S 120x 10TB HDDs	
Total Capacity	367.2 TB		1248 TB	
Number of LUNs	10		10	
Host Connectivity	4 x 10 Gb iSCSI		4 x 10 Gb iSCSI	
Compression Ratio	4:1		N/A	
Que Depth	100/LUN		100/LUN	
Block Size	1M Seq Write	512k Seq Write	1M Seq Write	512k Seq Write
Throughput	1198 MB/s	1195 MB/s	1123 MB/s	1122 MB/s
Latency	1.4ms	1.1ms	98ms	62ms
Total Watts	645	644	2118	2109

## 5.5.4 Compute

Veeam backup components consumes considerable resources, and these should be sized and scaled appropriately to meet availability and recovery objectives. These components are sized based on the number of backup jobs and backup storage requirements. It is recommended to use dedicated servers to avoid resource contention in the source environment. However if the source is with latest generation high performance processors and it can be run along with workloads.

### Backup Proxy

The backup proxy transfer data from source to Repository server and it should have enough cores to process. It is recommended to allocate at least 1 core per 100MB/s transfer and split backup to multiple tasks to increase concurrency and reduce time. Each proxy task require 2GB memory and proxies can be distributed evenly across nodes in the source cluster to balance performance and prevent it from reserving compute for source workloads. More sizing examples can be found here [Veeam Backup Proxy for VMware](#)

**Backup Repository Server:** The backup repository server is sized based on the number of proxy servers and backup copy requirements. For each proxy, the repository server requires one third processing power required for the proxy. The additional processing power required for secondary backup is depends on the total backup storage capacity. Also, it is recommended to keep repository server closer to proxy servers to reduce latency and improve backup and recovery time. Refer design guidelines here [Veeam Backup Repositories](#)

## 5.5.5 Network Bandwidth

The network requirement needs be sized based on recovery speed, storage connectivity options, storage system throughput limit and backup window requirements. It is recommended to use dedicated 10Gbps/25Gbps links for backup traffic. The higher backup and recovery speed provides better RPO/RTO, and it should be at least 10Gbps. If the backup storage is connected through iSCSI or NAS, it is

recommended to use more parallel links between proxies and repository servers and between repository server to backup storage systems due to speed constraints.

Fibre Channel connections provide better speed than iSCSI and Ethernet and it is recommended to use for critical applications and also if storage array-based replication is chosen instead of Veeam proxy based replication.

## 5.6 Backup and Recovery for VMware

Veeam backup server supports adding vSphere hosts and vCenter servers as source environment and create backup jobs from the wizard. It also provides instant recovery wizard through which virtual machines can be restored from the backup easily.


### 5.6.1 Create Backup Job for VMware

To add vCenter, open the Inventory view, in the inventory pane select the VMware vSphere node and click Add Server, then specify vCenter server name, specify credentials and apply settings.

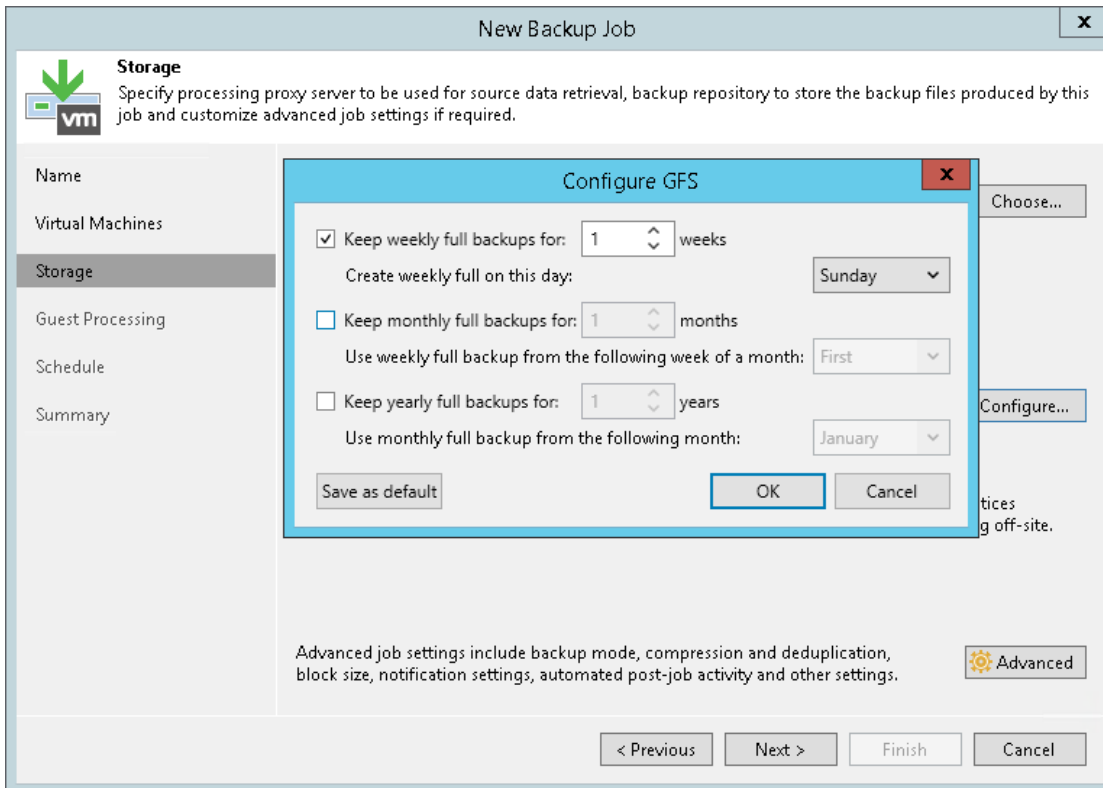
To create backup jobs,

- Right click jobs and navigate to Backup->Virtual Machines->VMware vSphere.
- Provide name for a backup job, select virtual machines for backup from the host, create exclusion list if required.
- Define backup order for the selected virtual machines.
- Choose backup proxy and primary backup storage.
- Choose advanced settings for vSphere Changed block tracking (CBT) and VMware tools quiescence.
- Select secondary target storage if required.
- Specify "Guest Processing Settings" for application aware processing and guest file system indexing.
- Create "Backup Schedule" based on your backup window and Finish backup job wizard.

Refer Veeam documentation for detailed steps [Create Backup Job](#)



**Figure 12: Add source objects from vCenter for backup**



**Figure 13: Choose backup methods**



## 5.6.2 Configure Virtual Machine Recovery

Virtual machines can be recovered through Veeam Instant recovery wizard. Refer Veeam documentation for detailed steps [VMware Virtual Machine Recovery](#)

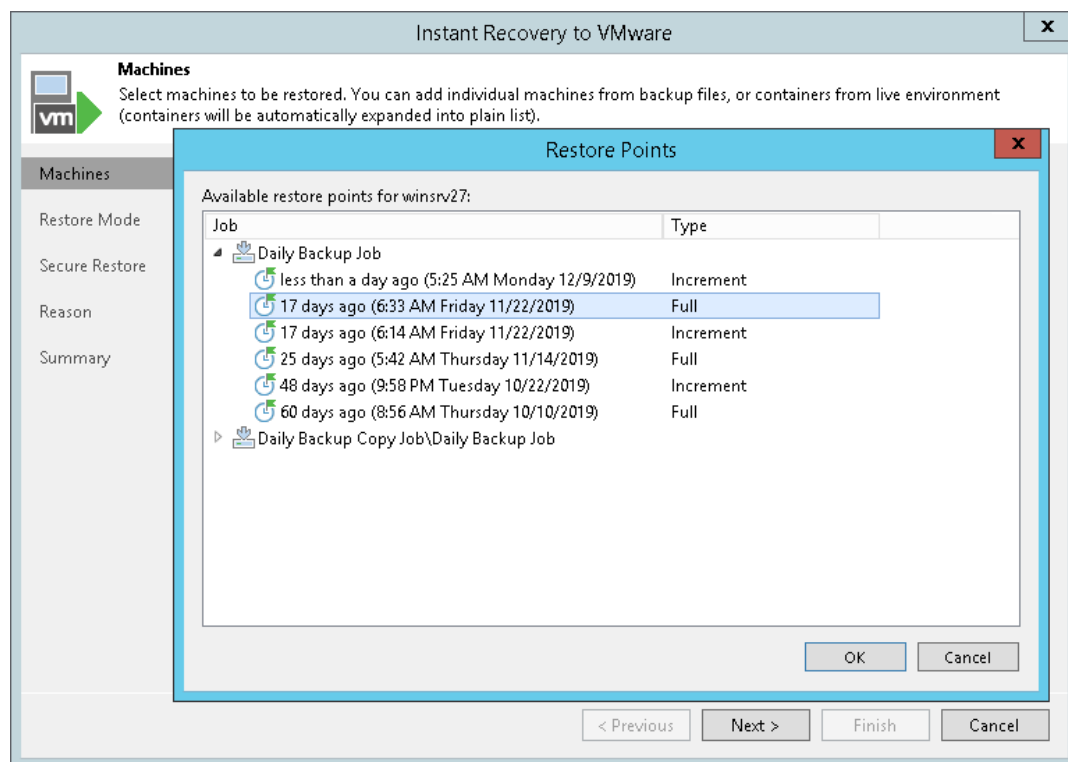


Figure 14: Recovering vSphere virtual machines

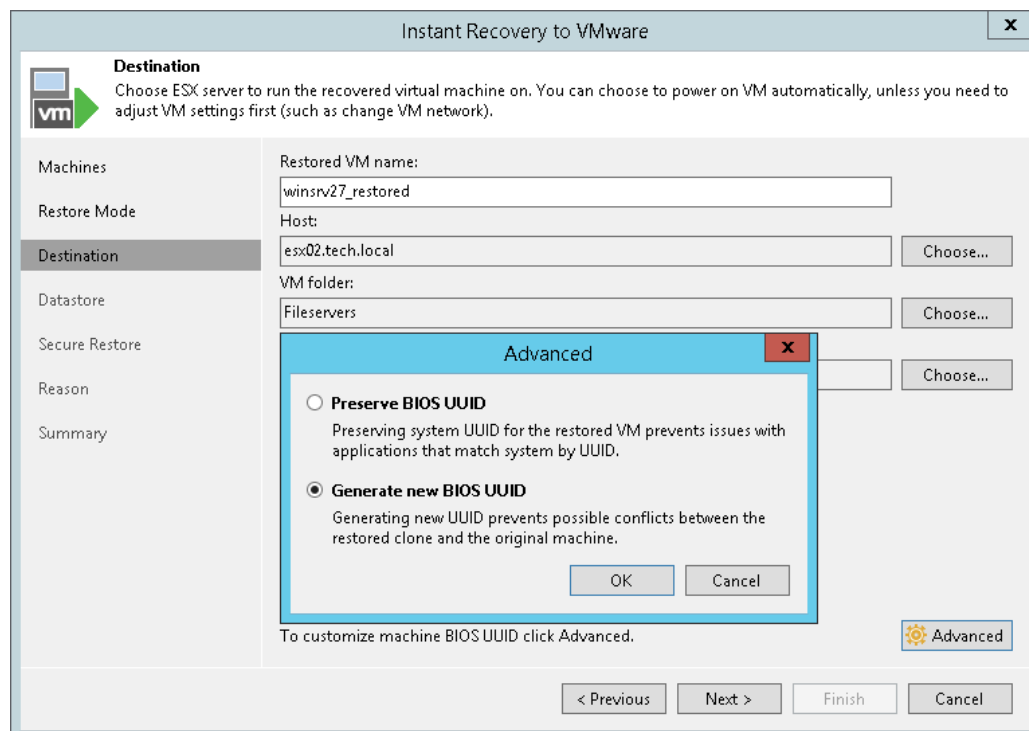


Figure 15: Recovery on destination hosts

### 5.6.3 vCenter Integration

Veeam Backup provides vCenter plugin to monitor backup jobs from the vCenter console. Veeam backup server manages backup jobs and connects with vCenter to backup hosts and virtual machines by leveraging snapshots.

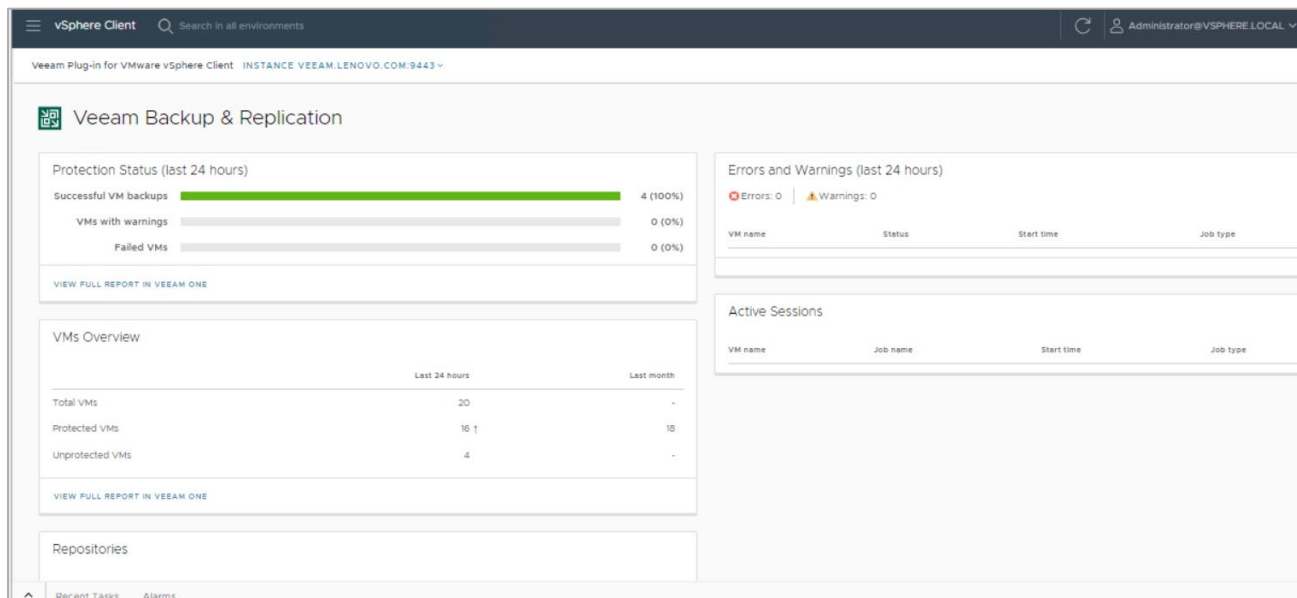


Figure 16: vCenter integration

## 5.7 Storage Array Snapshot Integration

Veeam Backup support leveraging Lenovo DM/DG storage snapshot feature and it enables faster backup and recovery than virtual machine snapshots. In cases where primary VM storage is replicated to a secondary DM/DG, Veeam can perform backup from the secondary DM and eliminate backup load from the production environment. Storage snapshots fill the gap in recovery points as most backups scheduled during backup window.


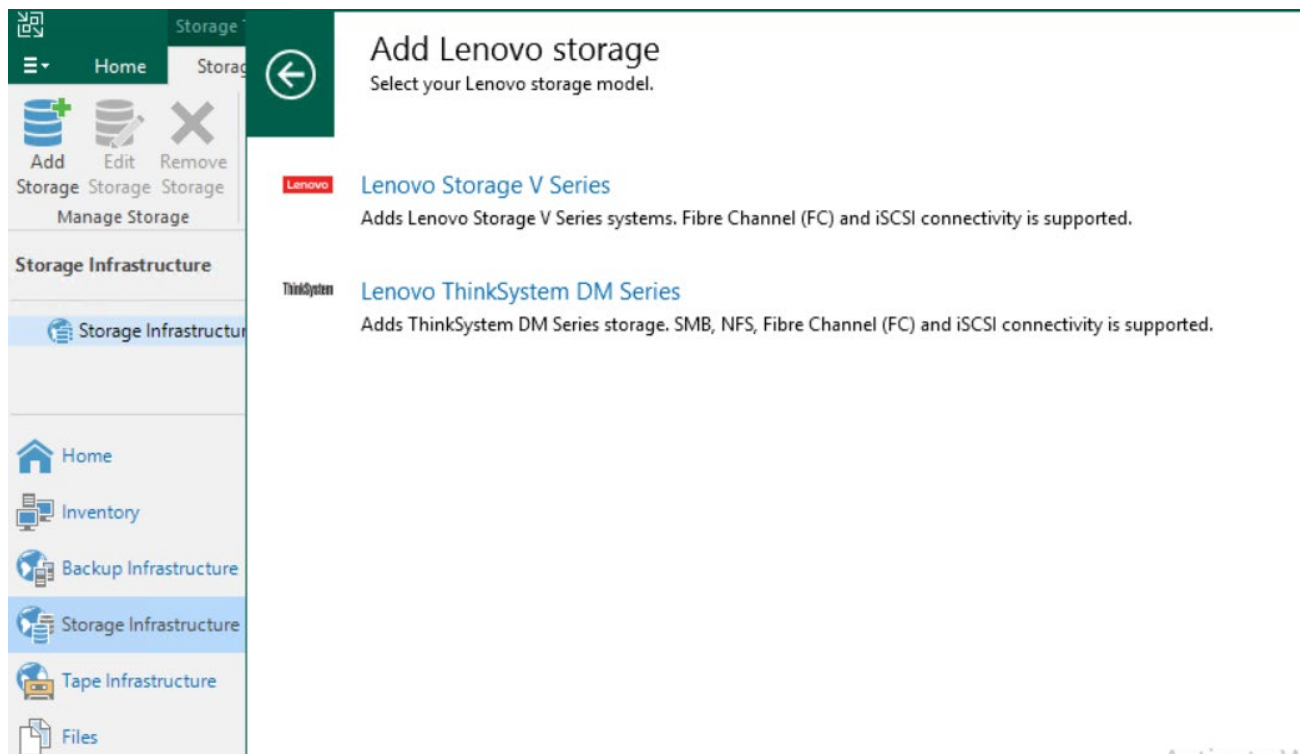


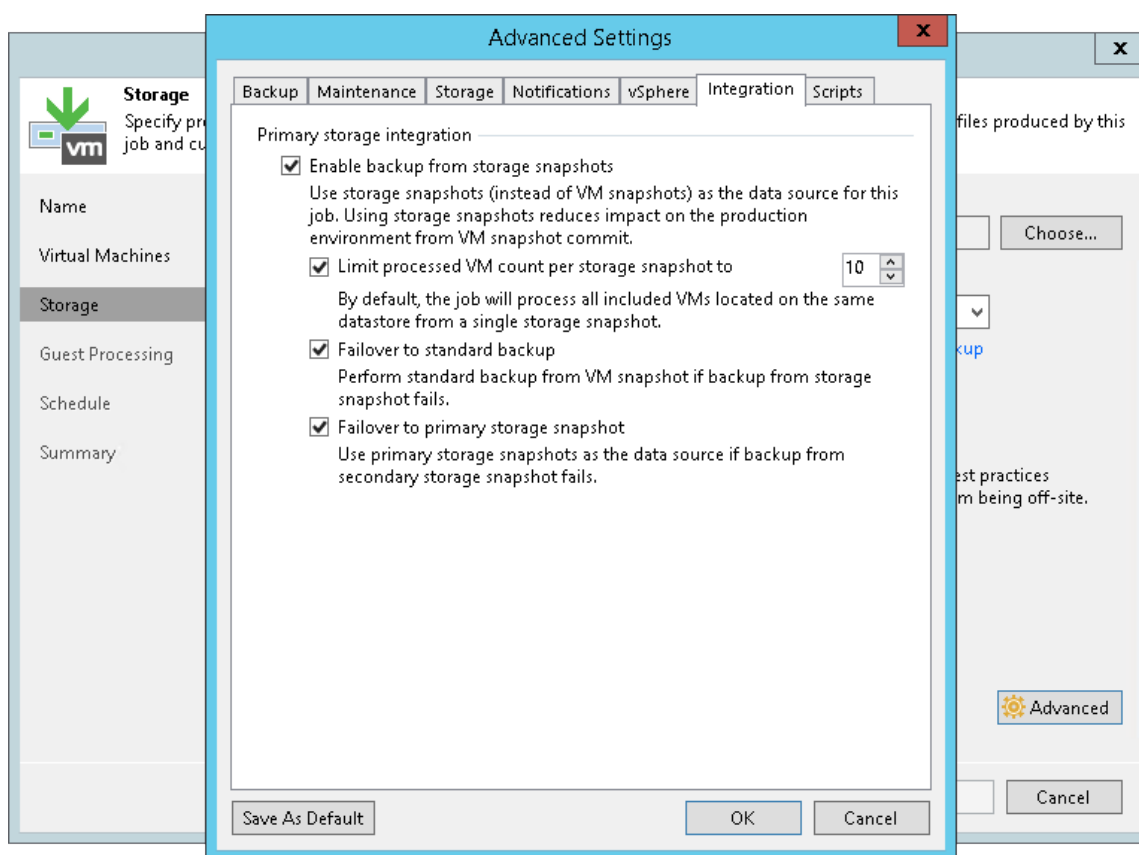
Figure 17: DM SnapMirror integrations

Veeam Backup is compatible Lenovo ThinkSystem DE/DG storage and integrated with Lenovo DM systems



**Figure 18: Offsite Backup with DE and ThinkSystem**

If the source environment is with DM storage and it can be configured in advance settings when creating backup job.



**Figure 19: Lenovo storage systems feature integration**

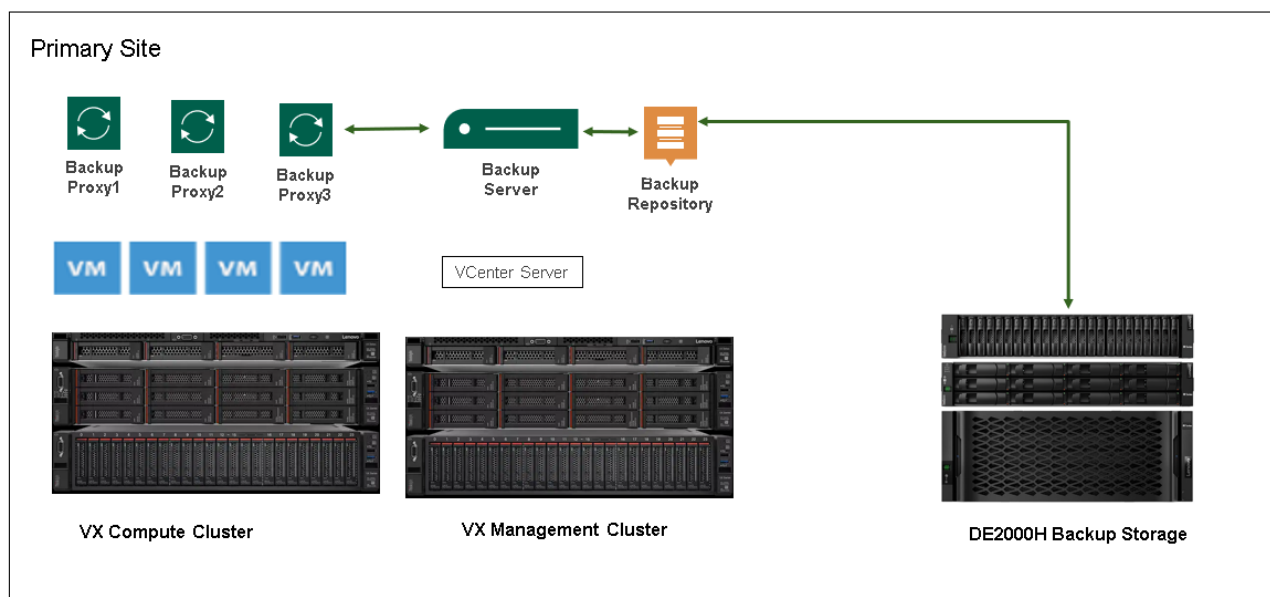
## 5.8 Deployment Model with ThinkSystem DE/DM/DG Storage

Veeam Backup and Replication components can be deployed either on dedicated servers or can share the compute with application workloads and infrastructure management components. The deployment model in this reference architecture uses dedicated primary backup storage which can be deployed at onsite along with source or can be deployed offsite. The backup management components can access backup storage from either source site or from backup site. This reference architecture emphasises on deploying all backup management components at the source site closer to virtual machines and applications which require backup.

### 5.8.1 Onsite Backup with DE Storage and Shared ThinkAgile VX

In this deployment model the source environment is with 8 node ThinkAgile VX compute cluster running different workloads and another 4 node ThinkAgile VX management cluster. The backup storage is DE2000H which co-exist on the same primary data center.

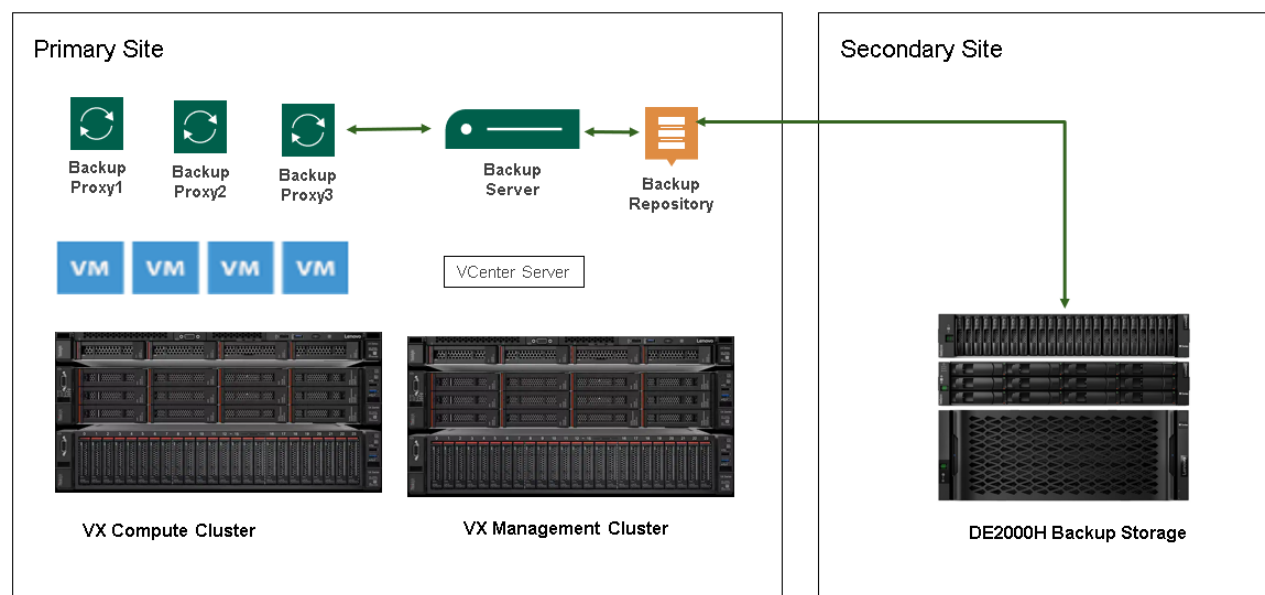
In this deployment model, the backup proxies are kept closer to the source virtual machines on the VX compute cluster and Veeam backup server and repository servers are running on the VX management cluster.



**Figure 20: Onsite Backup with DE and VX**

### 5.8.2 Offsite Backup with DE Storage and Shared ThinkAgile VX

In this deployment model the backup storage DE2000H is deployed on the secondary site and all the Veeam backup management components are deployed on the VX clusters at the source site. The Veeam backup repository uses dedicated high-speed connection to access the storage at the remote site.



**Figure 21: Offsite Backup with DE and Shared VX**

### 5.8.3 Offsite Backup with DE Storage and Dedicate ThinkSystem

In this deployment model the backup storage DE2000H and Veeam Repository server is deployed at the secondary site. All other Veeam backup management components deployed on the source site. The backup repository server uses dedicated ThinkSystem servers.

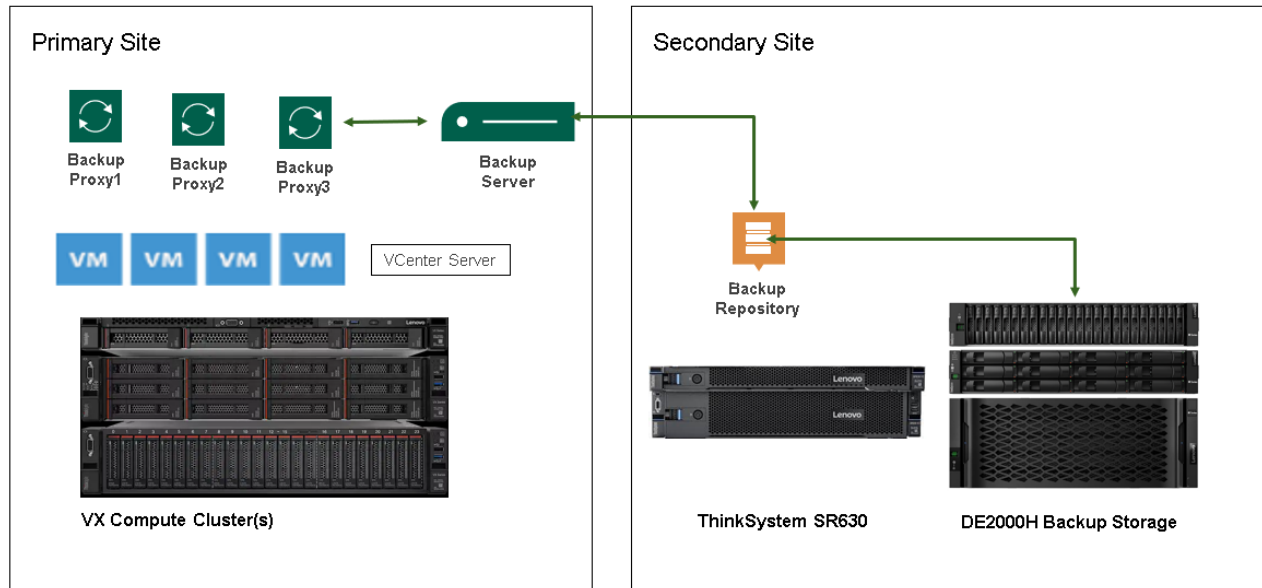


Figure 22: Offsite Backup with DE and ThinkSystem

## 6 Appendix: Bill of materials

This appendix contains the bill of materials (BOMs) for different configurations of hardware. There are sections for user servers, management servers, storage, networking switches, chassis, and racks that are orderable from Lenovo. The last section is for hardware orderable from an OEM.

The BOM lists in this appendix are not meant to be exhaustive and must always be double-checked with the configuration tools. Any discussion of pricing, support, and maintenance options is outside the scope of this document.

### 6.1 BOM for Backup Storage

This section contains the bill of materials for enterprise and SMB compute servers. Table 13 to Table 17 list the BOM for a Backup storage with ThinkSystem DE/DM/DG storage systems.

**Table 13: ThinkSystem DE2000H - 50TB Backup Storage**

Part number	Product Description	Qty
7Y70A00YWW	Storage : DE Controller DE2000H 8G HICless 2U12 LFF, Gen2	1
4C57A14369	Lenovo ThinkSystem DE2000/4000 HIC, 10/25GbE iSCSI,4-ports	2
4M17A13529	Lenovo 10/25GbE iSCSI SFP28 Module	2
4XB7A74942	Lenovo ThinkSystem DE Series 18TB 7.2K 3.5" HDD 2U12	5
00WE139	3m Green Cat6 Cable	2
7Z57A03558	Lenovo 3m Passive 25G SFP28 DAC Cable	4
5PS7A21506	Premier Essential - 3Yr 24x7 4Hr Resp + YDYD DE2000H 2U12	1
7S0L002EWW	Veeam Backup & Replication Universal License. Includes Enterprise Plus Edition features. - 1 Year Subscription Upfront Billing & Production (24/7) Support	2

**Table 14: ThinkSystem DE2000H - 200TB Backup Storage**

Part number	Product Description	Qty
7Y70CTO2WW	Storage : Lenovo ThinkSystem DE2000H Hybrid Flash Array LFF Gen2	1
BEY6	Lenovo ThinkSystem Storage Series 2U12 Chassis	1
B4DC	iSCSI (Optical)	1
BQ9Y	Lenovo ThinkSystem DE2000 Controller 8GB Gen2	2
BHDP	Lenovo ThinkSystem DE Series 18TB 7.2K 3.5" HDD 2U12	12
A1PK	5m Passive DAC SFP+ Cable	4
B4BP	Lenovo ThinkSystem Storage USB Cable, Micro-USB	1
6201	1.5m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	2
BQH6	Lenovo ThinkSystem DE Series Ship Kit (RoW) 2U, Gen2	1
B38X	Lenovo ThinkSystem Storage Rack Mount Kit 2U12	1
BU1E	SMB Express Flag	1
B4MC	Lenovo ThinkSystem DE2000H SMID Controller Base Setting	1
B4Q4	Lenovo ThinkSystem DE2000H Base Bundle	1
B4B0	Lenovo ThinkSystem DE Series 2U12 End Cap Kit (Pair)	1
B4BJ	Lenovo ThinkSystem Storage 2U12 System Label	1

Part number	Product Description	Qty
B4BK	Lenovo ThinkSystem DESeries DE2000H Product Label	1
B4AW	Lenovo ThinkSystem Storage Packaging 2U	1
B4BA	Lenovo ThinkSystem DE2000/4000 HIC, 10/25GbE iSCSI,4-ports	2
7S0LCTO2WW	Veeam Backup & Replication Universal License. Includes Enterprise Plus Edition features. - 3 Years Subscription Upfront Billing & Production (24/7) Support	1

**Table 15: ThinkSystem DE4000H - 300TB Backup Storage**

Part number	Product Description	Qty
7Y77A00EWW	Storage : DE Controller DE4000H 8G HICless 4U60 LFF, Gen2	1
4C57A14369	Lenovo ThinkSystem DE2000/4000 HIC, 10/25GbE iSCSI,4-ports	2
4M17A13529	Lenovo 10/25GbE iSCSI SFP28 Module	4
4XB7A74943	Lenovo ThinkSystem DE Series 18TB 7.2K 3.5" HDD 4U60	20
00WE139	3m Green Cat6 Cable	2
5AS7A83028	Hardware Installation (Business Hours) for DE4000H 4U60	1
5PS7A21008	Premier Essential - 3Yr 24x7 4Hr Resp + YDYD DE4000H 4U60	1
7S0L002GWW	Veeam Backup & Replication Universal License. Includes Enterprise Plus Edition features. - 3 Years Subscription Upfront Billing & Production (24/7) Support	2

**Table 16: ThinkSystem DE6000H - 300TB Backup Storage**

Part number	Product Description	Qty
7Y80A00NWW	Storage : DE Controller DE6000H 32G 10/25Gb ISCSI 4U60 LFF, Gen2	1
4XB7A74943	Lenovo ThinkSystem DE Series 18TB 7.2K 3.5" HDD 4U60	20
90Y9433	5m Passive DAC SFP+ Cable	8
5AS7A83034	Hardware Installation (Business Hours) for DE6000H 4U60	1
5PS7A21681	Premier Essential - 3Yr 24x7 4Hr Resp + YDYD DE6000H 4U60	1
7S0L002LWW	Veeam Backup & Replication Universal License. Includes Enterprise Plus Edition features. - Subscription Upfront Billing & Production (24/7) Support- Monthly Coterm	2

**Table 17: ThinkSystem DG5000 - 300TB Backup Storage**

Part number	Product Description	Qty
7DE4CTO1W	Controller : ThinkSystem DG5000 QLC All Flash Array	1
BF3C	Lenovo ThinkSystem Storage 2U NVMe Chassis	1
BQHN	Lenovo ThinkSystem 2U NVMe Controller with Titanium PSU	2
BXG8	Lenovo ThinkSystem 30.7TB (2x 15.36TB QLC NVMe SSD) Drive Pack for DG5000 - Unified Complete Bundle	10



BEVQ	Lenovo ThinkSystem Storage HIC, 10/25Gb iSCSI,4-ports	2
AVG0	3m Green Cat6 Cable	2
AV1W	Lenovo 1m Passive 25G SFP28 DAC Cable	2
B4BP	Lenovo ThinkSystem Storage USB Cable, Micro-USB	1
BX90	Lenovo ThinkSystem Storage ONTAP 9.12 Software Encryption - IPAv2	1
BY3B	SnapMirror License Bundle	1
5PS7B27398	Premier Essential 3Y 24x7x4+YDYD ThinkSystem DG5000 Unified Complete	1
5WS7B27496	Premier Essential 3Y 24x7x4 DG5000 307TB (20x15.36TB QLC NVMe)Pack Unified Complete	1
7S0L002EWW	Veeam Backup & Replication Universal License. Includes Enterprise Plus Edition features. - 1 Year Subscription Upfront Billing & Production (24/7) Support	2

## 6.2 BOM for Backup Server

This section contains the bill of materials for compute servers for Veeam Backup Repository server

**Table 18: ThinkSystem SR630 V2**

Part number	Product Description	Qty
7Z71CTO1WW	Server : ThinkSystem SR630 V2-3yr Warranty	1
BH9Q	ThinkSystem 1U 2.5" Chassis with 8 or 10 Bays	1
BFYE	Operating mode selection for: "Efficiency - Favoring Performance Mode"	1
BB2Z	Intel Xeon Silver 4314 16C 135W 2.4GHz Processor	2
B965	ThinkSystem 32GB TruDDR4 3200 MHz (2Rx8 1.2V) RDIMM	6
5977	Select Storage devices - no configured RAID required	1
BHTU	ThinkSystem V2 750W(230V/115V) Platinum Hot-Swap Power Supply v2	1
6311	2.8m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	1
4XC7A08237	ThinkSystem Broadcom 57414 10/25GbE SFP28 2-port OCP Ethernet Adapter	1
AUPW	ThinkSystem XClarity Controller Standard to Enterprise Upgrade	1
BH9M	ThinkSystem V3 1U Performance Fan Option Kit v2	6
B8LA	ThinkSystem Toolless Slide Rail Kit v2	1
BA2Y	ThinkSystem 1U Front VGA Cable	1
B0MK	Enable TPM 2.0	1
BH9R	10x2.5" Media Bay w/ Cable	1
BHJS	1U MB PSU Airduct for CPU>125W	1
B978	ThinkSystem SR630/SR850/SR860 V2 Standard Heatsink	2
B8NB	ThinkSystem 1U MS LP Riser Filler	1
B8NA	ThinkSystem 1U MS LP+LP Riser Filler	1
B5WJ	ThinkSystem OCP3 Filler	1
AVKJ	ThinkSystem 2x2 Quad Bay Gen4 2.5" HDD Filler	2
AVEN	ThinkSystem 1x1 2.5" HDD Filler	2
5641PX3	XClarity Pro, Per Endpoint w/3 Yr SW S&S	1

Part number	Product Description	Qty
1340	Lenovo XClarity Pro, Per Managed Endpoint w/3 Yr SW S&S	1
7Q13CTOSWW	SERVER INSTALLATION	1
QA51	SR630 V2 3Y STANDARD	1
7Q01CTO9WW	SERVER PREMIER ESSENTIAL 24x7x4	1

# Resources

---

For more information, see the following resources:

Lenovo ThinkSystem DE Series Storage Arrays

<https://lenovopress.lenovo.com/lp0940-lenovo-thinksystem-de-series-storage-arrays>

ThinkSystem DE Series Hybrid Flash Array

<https://lenovopress.lenovo.com/datasheet/ds0049-lenovo-thinksystem-de-series-hybrid-flash-array>

ThinkSystem DE Series All-Flash Array

<https://lenovopress.lenovo.com/datasheet/ds0051-lenovo-thinksystem-de-series-all-flash-array>

ThinkSystem DM Series Hybrid Flash

<https://lenovopress.lenovo.com/datasheet/ds0048-lenovo-thinksystem-dm-series-hybrid-flash>

ThinkSystem DG5000 and DG7000

<https://lenovopress.lenovo.com/lp1754-thinksystem-dg5000>

<https://lenovopress.lenovo.com/lp1755-thinksystem-dg7000>

Veeam Data Platform

<https://www.veeam.com/products/data-platform-premium.html>

ThinkAgile VX Series

<https://lenovopress.lenovo.com/ds0104>

VMware Hyperconverged Infrastructure - vSAN

<https://www.vmware.com/in/products/vsan.html>

# Document History

Version 1.0	30 June 2023	<ul style="list-style-type: none"><li>Initial version</li></ul>
Version 1.1	06 September 2023	<ul style="list-style-type: none"><li>Sections 1,2, 5.1, 5.2, 5.5 - ThinkSystem DG storage system features, configuration and performance data are included.</li><li>Section 6.1 - BOM is updated with DG5000 storage</li></ul>

# Trademarks and special notices

---

© Copyright Lenovo 2024

References in this document to Lenovo products or services do not imply that Lenovo intends to make them available in every country.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®  
ThinkAgile®  
ThinkSystem®  
TruDDR4  
XClarity®

The following terms are trademarks of other companies:

Intel® and Xeon® are trademarks of Intel Corporation or its subsidiaries.

Linux® is the trademark of Linus Torvalds in the U.S. and other countries.

Microsoft®, PowerShell, SQL Server®, Windows PowerShell®, Windows Server®, and Windows® are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used Lenovo products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-Lenovo products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by Lenovo. Sources for non-Lenovo list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. Lenovo has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-Lenovo products. Questions on the capability of non-Lenovo products should be addressed to the supplier of those products.

All statements regarding Lenovo future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local Lenovo office or Lenovo authorized reseller for the full text of the specific Statement of Direction.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in Lenovo product announcements. The information is presented here to communicate Lenovo's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard Lenovo benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models.

Any references in this information to non-Lenovo 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 Lenovo product and use of those websites is at your own risk.