



Reference Architecture: Lenovo ThinkSystem Entry Level Compute and Storage Solution for AI Training Workloads

Last update: 14 July 2020

Version 1.0

Reference Architecture for AI Training Workloads

Contains performance data for AI training using shared storage

Designed for small and medium data science teams

Contains Bill of Materials for compute servers and storage

Miro Hodak, Lenovo
David Arnette, NetApp



Table of Contents

1	Introduction	1
1.1	Target Audience	1
1.2	Solution Architecture	1
2	Technology Overview	3
2.1	Lenovo ThinkSystem DM Series System	3
2.2	ONTAP 9	3
2.2.1	Simplify Data Management	3
2.2.2	Accelerate and Protect Data	4
2.2.3	Future-Proof Infrastructure	4
2.3	FlexGroup Volumes	4
2.4	Lenovo ThinkSystem Server Portfolio	5
2.5	Lenovo ThinkSystem SR670 Rack Server	5
2.6	MLPerf	6
3	Test Overview	7
4	Test Configuration	8
4.1	Solution Architecture	8
5	Test Procedure	10
5.1	Data Read Speed: Local vs Network Storage	10
5.2	AI Training using ResNet 50v1.5 with Network Data	10
6	AI Training Results	12
6.1	Shared Use Case: Multiple Independent Single Node Jobs	12
6.2	Exclusive Use Case: Dedicated Training Model Scale Out	13
6.3	Exclusive Use Case: Scale Out with a Faster Compute Network	15
7	Architecture Adjustments	18
7.1	Compute Scale Out	18

7.2	CPU Adjustments.....	18
7.3	GPU Adjustments.....	18
7.4	Storage Capacity Increase.....	18
7.5	Other ThinkSystem Servers	18
7.6	Other DM Series Storage Options	18
8	Deployment considerations.....	19
9	Conclusion	20
10	Appendix: Lenovo Bill of materials.....	21
10.1	BOM for compute servers	21
10.2	BOM for Storage Servers.....	22
	Resources	25
	Document history	26

1 Introduction

This solution describes an entry-level, clustered architecture using Lenovo ThinkSystem compute servers and Lenovo DM Series storage systems optimized for Artificial Intelligence (AI) training workflows accelerated by GPUs. The architecture is targeted for enabling small and medium sized teams where most compute jobs are single node (single or multi-GPU) or distributed over a few computational nodes. This is not a major limitation as most day-to-day AI training jobs are single node.

This document covers testing and validation of a compute/storage configuration consisting of four accelerated Lenovo ThinkSystem SR670 servers and an entry-level 10 GbE-connected Lenovo DM storage system, providing an efficient and cost-effective solution for small and medium-sized organizations starting out with AI that require the enterprise-grade capabilities of ONTAP® cloud-connected data storage available with Lenovo DM-Series storage.

1.1 Target Audience

This document is intended for the following audiences:

- Data scientists, data engineers, administrators, and developers of AI systems.
- Enterprise architects who design solutions for the development of AI models and software.
- Data scientists and data engineers who are looking for efficient ways to achieve deep learning (DL) and machine learning (ML) development goals.
- IT decision makers and business leaders who want to achieve the fastest time to market possible from AI initiatives.

1.2 Solution Architecture

This Lenovo ThinkSystem server and DM Series storage solution is designed to handle AI training on large datasets using the processing power of GPUs alongside traditional CPUs. This validation demonstrates high performance and optimal data management with a scale-out architecture that uses either one, two, or four ThinkSystem SR670 servers alongside a DM Series storage system. The scheme of the architecture is given in Figure 1.

This compute/storage solution provides the following key benefits:

- Highly efficient and cost-effective performance when executing multiple training jobs in parallel
- Scalable performance based on different numbers of ThinkSystem servers
- Robust data protection to meet low recovery point objectives (RPOs) and recovery time objectives (RTOs) with no data loss
- Optimized data management with snapshots and clones to streamline development workflows




Figure 1: Architecture Overview

2 Technology Overview

This Section describes key technologies leveraged in this solution.

2.1 Lenovo ThinkSystem DM Series System

State-of-the-art Lenovo DM storage systems enable IT departments to meet enterprise storage requirements with industry-leading performance, superior flexibility, cloud integration, and best-in-class data management. Designed specifically for flash, DM storage systems help accelerate, manage, and protect business-critical data.

This Reference architecture is based on ThinkSystem DM5000F storage. For NAS workloads, this single entry-level storage system supports throughput of 5GBps for sequential reads and 148K IOPS for small random reads at latencies of 1ms or less. Entry-level AFF storage systems support the following features:

- Scalable throughput of up to 120GBps and 3.55 million IOPS in a 24-node cluster
- 10 Gigabit Ethernet (GbE) and FC connectivity
- Up to 15TB solid-state drives (SSDs) with multi-stream write
- ONTAP 9, with a complete suite of data protection and replication features for industry-leading data management

Lenovo offers other storage systems, such as the DM7000F and DM7100F, that offer higher performance and scalability for larger-scale deployments.

2.2 ONTAP 9

ONTAP 9 is the latest generation of storage management software that enables businesses to modernize infrastructure and transition to a cloud-ready data center. Leveraging industry-leading data management capabilities, ONTAP enables the management and protection of data with a single set of tools, regardless of where that data resides. Data can also be moved freely to wherever it's needed—the edge, the core, or the cloud. ONTAP 9 includes numerous features that simplify data management, accelerate and protect critical data, and future-proof infrastructure across hybrid cloud architectures.

2.2.1 Simplify Data Management

Data management is crucial to enterprise IT operations so that appropriate resources are used for applications and datasets. ONTAP includes the following features to streamline and simplify operations and reduce the total cost of operation:

- **Inline data compaction and expanded deduplication.** Data compaction reduces wasted space inside storage blocks, and deduplication significantly increases effective capacity. This applies to data stored locally and data tiered to the cloud.
- **Minimum, maximum, and adaptive quality of service (QoS).** Granular QoS controls help maintain performance levels for critical applications in highly shared environments.
- **ONTAP FabricPool.** This feature provides automatic tiering of cold data to public and private cloud storage options, including Amazon Web Services (AWS) and Azure.

2.2.2 Accelerate and Protect Data

ONTAP delivers superior levels of performance and data protection and extends these capabilities as follows:

- **Performance and lower latency.** ONTAP offers the highest possible throughput at the lowest possible latency.
- **Data protection.** ONTAP provides built-in data protection capabilities with common management across all platforms.
- **Volume Encryption.** ONTAP offers native volume-level encryption with both onboard and external key management support.

2.2.3 Future-Proof Infrastructure

ONTAP 9 helps meet demanding and constantly changing business needs:

- **Seamless scaling and nondisruptive operations.** ONTAP supports the nondisruptive addition of capacity to existing controllers as well as to scale-out clusters. Customers can upgrade to the latest technologies such as NVMe and 32Gb FC without costly data migrations or outages.
- **Cloud connection.** ONTAP is the most cloud-connected storage management software, with options for software-defined storage (ONTAP Select) and cloud-native instances (Cloud Volumes Service) in all public clouds.
- **Integration with emerging applications.** ONTAP offers enterprise-grade data services for next-generation platforms and applications such as OpenStack, Hadoop, and MongoDB by using the same infrastructure that supports existing enterprise apps.

2.3 FlexGroup Volumes

Training datasets are typically a collection of potentially billions of files. Files can include text, audio, video, and other forms of unstructured data that must be stored and processed to be read in parallel. The storage system must store many small files and must read those files in parallel for sequential and random I/O.

A FlexGroup volume (Figure 3) is a single namespace made up of multiple constituent member volumes that is managed and acts like a FlexVol[®] volume to storage administrators. Files in a FlexGroup volume are allocated to individual member volumes and are not striped across volumes or nodes. They enable the following capabilities:

- Up to 20 petabytes of capacity and predictable low latency for high-metadata workloads.
- Up to 400 billion files in the same namespace.
- Parallelized operations in NAS workloads across CPUs, nodes, aggregates, and constituent FlexVol volumes.




Figure 2 FlexGroup Volumes

2.4 Lenovo ThinkSystem Server Portfolio

Lenovo ThinkSystem servers feature innovative hardware, software and services that solve customer challenges today and deliver an evolutionary fit-for-purpose, modular design approach to address tomorrow's challenges. These servers capitalize on best-in-class, industry-standard technologies coupled with differentiated Lenovo innovations to provide the greatest possible flexibility in x86 servers.

Key advantages of deploying Lenovo ThinkSystem servers include:

- Highly scalable, modular designs to grow with your business
- Industry-leading resilience to save hours of costly unscheduled downtime
- Fast flash technologies for lower latencies, quicker response times and smarter data management in real-time

In the AI area, Lenovo is taking a practical approach to helping enterprises understand and adopt the benefits of ML and AI for their workloads. Lenovo customers can explore and evaluate Lenovo AI offerings in Lenovo AI Innovation Centers so that they can better understand the value for their particular use case. To improve time to value, this customer-centric approach provides customers with proofs of concept for solution development platforms that are ready-to-use and optimized for AI.

2.5 Lenovo ThinkSystem SR670 Rack Server

Lenovo ThinkSystem SR670 delivers high performance for accelerated AI and High-Performance Computing (HPC) workloads. Supporting the latest scalable Intel Xeon CPUs and up to either four large or eight small GPUs per 2U node, the SR670 is designed for the computationally intensive demands of Machine Learning (ML), Deep Learning (DL), inference and other accelerated applications. Supported accelerators include the NVIDIA Tesla V100, V100S, and T4, and with integrated modularity within the system further optimization for specific workload characteristics can be achieved.

As more workloads have been designed to leverage the performance of accelerators, the demand for GPU density has increased. Industries such as retail, financial services, energy, and healthcare are leveraging GPUs to extract greater insights and drive innovation with ML, DL, and inference techniques. The ThinkSystem SR670 provides an optimized enterprise-grade solution for deploying accelerated HPC and AI

workloads in production, maximizing system performance while maintaining data center density.



Figure 3 ThinkSystem SR670

2.6 MLPerf

MLPerf is the industry leading benchmark suite for evaluating AI performance. In this validation, we used its image classification benchmark with TensorFlow, currently the most popular AI framework. The Tensorflow benchmark script, ([github:tensorflow/benchmarks](https://github.com/tensorflow/benchmarks)), was used to drive AI training in this work. The script contains implementations of several popular convolutional neural network (CNN) models and is designed to be as fast as possible. It can be run on both a single machine and in a distributed mode across multiple hosts.

3 Test Overview

In this validation, we performed image recognition training as specified by MLPerf v0.5. Specifically, we trained the ResNet v1.5 model with ImageNet dataset until we reached accuracy of 74.9%. The main metric is the time to reach the desired accuracy. We also report training bandwidth in images per second to better judge scale out efficiency.

The primary test case evaluated multiple independent training processes (one per node) running concurrently. This simulates the main use case: a shared system used by multiple data scientists. The second test case evaluated scale-out efficiency.

Table 1 lists a summary of results for all tests performed for this solution.

Table 1) Test result summary.

Test Description	Results Summary
Image recognition training: Multiple concurrent jobs	Highly efficient performance. All the jobs ran at full speed even when the cluster was fully used. The DM storage systems delivered training performance comparable to local SSD storage while enabling easy sharing of data between servers.
Image recognition training: Scale out	Highly efficient up to four nodes. At that point, scale out was less efficient but still feasible. Using a higher speed computational network improves scalability. The DM Series storage system delivered training performance comparable to local SSD storage while enabling easy sharing of data between servers.

4 Test Configuration

This section describes tested configurations, network infrastructure, SR670 server, and storage provisioning details.

4.1 Solution Architecture

We used the solution components listed in Table 2 for the validation.

Table 2) Base components for the solution architecture.

Solution Components	Details
ThinkSystem servers running Linux (CentOS 7.6 with CUDA 10.1)	<ul style="list-style-type: none">• Four SR670 servers each with 4 x NVIDIA 32GB V100 GPU cards• Each server contains two Intel Xeon Gold 6142 CPUs (16 physical cores) and 384GB RAM
DM Series Storage DM5000F (HA pair)	<ul style="list-style-type: none">• ONTAP 9 software• 24 x 960GB SSDs• NFS protocol• 1 interface group (ifgrp) per controller, with 4 logical IP addresses for mount points

An earlier version of the storage system was leveraged for solution architecture and performance validation. The currently available DM5000F system should meet or exceed the performance results described in this document.

In this validation, we used ResNet v1.5 with the ImageNet basis set as specified by MLPerf Training v0.5. The dataset is stored in a DM Series storage system accessed over the NFS protocol. The SR670s were connected to the storage over a 10GbE switch.

ImageNet is a frequently used image dataset. It contains almost 1.3 million images for a total size of 144GB. The average image size is 108KB. The dataset is stored in a TFRecord format, which is the recommended way for data input into TensorFlow. The TFRecord data is packaged in 1024 separate files each approximately 136MB in size.

Figure 4) Network topology of tested configuration.




Table 3) Storage configuration.

Controller	Aggregate	FlexGroup Volume	Aggregate Size	Volume Size	Operating System Mountpoint
Controller1	Aggr1	/netapplenovo_AI_fg	6.9TB		/netapp_lenovo_fg
				1TB	
Controller2	Aggr2		6.9TB		

The /netappLenovo_AI_fg folder has the dataset used for ResNet AI Training validation.

5 Test Procedure

Several types of tests were performed in this work: network read speed, AI training using multiple independent jobs, and AI training using a single scale-out job. Details are given below.

5.1 Data Read Speed: Local vs Network Storage

The read speed was tested by using `dd` command on one of the TFRecord files for the ImageNet dataset. Specifically, we ran the following commands for both local and network data:

```
sync ; echo 3 > /proc/sys/vm/drop_caches
dd if=/mnt/netapp/train-00001-of-01024 of=/dev/null bs=512k count=2048
Results (average of 5 runs):
Local storage: 516 MB/s
Network storage: 510 MB/s.
```

Both values are very similar demonstrating that the network storage is capable of delivering data at a rate essentially equal to that of local storage.

5.2 AI Training using ResNet 50v1.5 with Network Data

We ran the ResNet50 benchmark with one, two, or four SR670 servers. This test used the [TensorFlow benchmark script](#) and TensorFlow container version 1.13.1 from DockerHub to run the training.

We used the following test procedure in this validation:

We cleared the host cache before running the script to make sure that data was not already cached.

```
sync ; echo 3 > /proc/sys/vm/drop_caches
```

We ran the benchmark script with the ImageNet dataset in server storage (local SSD storage) as well as with the networked DM.

We validated network and local storage performance using the `dd` command.

For the single-node run, we passed the following options to the training script:

```
--num_gpus=4 --batch_size=512 --variable_update=replicated --all_reduce_spec=nccl --use_fp16=True --
data_format=NCHW --model=resnet50_v1.5 --optimizer=momentum --weight_decay=1e-4 --
num_warmup_batches=0 --nodistortions --gradient_repacking=2 --local_parameter_device=gpu --
display_every=100 --
eval_during_training_at_specified_epochs='1,5,9,13,17,21,25,29,33,37,41,45,49,53,57,61,62,63' --
num_eval_batches=25 --eval_batch_size=250 --loss_type_to_report=base_loss --single_l2_loss_op --
compute_lr_on_cpu --datasets_parallel_interleave_cycle_length=20 --
datasets_sloppy_parallel_interleave --stop_at_top_1_accuracy=0.749 --xla_compile=True --
resnet_base_lr=0.06 --ml_perf --ml_perf_compliance_logging=True --data_dir=/mnt/netapp/imagenet/ --
num_epochs=64
```

For the distributed runs, we used the parameter server's parallelization model. Four parameter servers per node were used, and we set the number of epochs to be the same as for the single node run, that is 61. We

did this because distributed training often takes more epochs due to imperfect synchronization between processes. The different number of epochs can skew runtime comparison between single-node and distributed cases.

6 AI Training Results

This Section contains test results for two use cases: Multiple independent single node jobs and a single job running on the entire cluster. The former is the expected use case, the latter shows that scale out jobs are also possible with this architecture.

6.1 Shared Use Case: Multiple Independent Single Node Jobs

This test simulated the main use case for this solution—multi-job, multi-user AI training. Each node ran its own training - each training utilized all four GPUs in a server - while using the shared network storage. The results are displayed in Figures 5 and 6, which show that the solution provides excellent performance with all jobs running at maximum speed and no slowdown due to multiple workloads is observed. The total throughput scaled linearly with the number of nodes.

Figure 5) Aggregate images per second on concurrent training models.





Figure 6) Average runtime in minutes for concurrent training models.



The average runtime for the network training model was 286 minutes and 20 seconds, and the individual runtimes were 285 minute and 35 seconds, 286 minute and 11 seconds, 286 minute and 27 seconds, and 287 minutes and 7 seconds. The average images per second for the training model were 4600, and the individual images per second were 4615, 4607, 4596, and 4582.

Based on our validation, one independent training model with network data runtime was 285 minutes and 19 seconds with 4611 images/sec. One independent training model with a local data (DAS) runtime was 277 minute and 44 seconds with 4724 images/sec.

During those runs the average GPU utilization was 96% as reported by *nvidia-smi* - note that this includes the testing phase, during which GPUs are not used, while CPU utilization was 40% as measured by *mpstat*. This demonstrates that the data delivery rate is sufficient in each case.

6.2 Exclusive Use Case: Dedicated Training Model Scale Out

The following figures show the scale out results from our lab validation. The tests used 1, 2 and 4 SR670 nodes. Comparison between local SSD and network storage again shows a very similar level of performance.

Figure 7) ResNet50 scale-out results for the runtime in minutes.





Figure 8) ResNet50 scale-out results in images per second.



In addition, note the following:

- In Figure 7, a higher number is better, while in Figure 8, the lower value indicates a better performance
- For the network storage case using a single NFS export path, the average throughput was 4603 images/sec, and the runtime was 288 minutes for 61 epochs.


- The less-than-linear scaling was because the 10GbE network is insufficient for node-to-node communication. Much better scaling can be obtained with a faster network as demonstrated in the next section. Note that the DM storage connected over a 10GbE network provides enough data bandwidth for the cluster.

6.3 Exclusive Use Case: Scale Out with a Faster Compute Network

Network

To improve network performance and thus scaling performance, we connected the compute nodes to a faster network - Mellanox HDR InfiniBand in in IPoIB mode, which in the current generation of ThinkSystem servers provides equivalent bandwidth and lower latency than 100GbE). This setup is shown in Figure 9.

Figure 9) Compute nodes connected through 100 GbE compute network.



The performance tests were run in the same way as in the previous section and are shown in Figures 10 and 11. The results show that the faster compute network improves scale out, providing near linear scaling. Network storage provides performance close to that of the local storage.

Figure 10) ResNet50 scale-out results in images per second for 100GbE compute network.





Figure 11) ResNet50 scale-out results runtime for 100GbE compute network.



For the DM Series storage four node training took 90 minutes with image throughput being 14894 images/sec. The local storage provides similar, albeit somewhat better, performance with runtime being 84 minutes and throughput of 15948 images/sec. Similarly, the results for two nodes with shared network data demonstrated a training model runtime of 170 minutes and 57 seconds with 7680 images/sec. For the same two-node training models with local data (DAS), we recovered a runtime of 162 minutes and 41 seconds with 8073 images/sec.

7 Architecture Adjustments

The setup used for the validation can be adjusted to fit other use cases.

7.1 Compute Scale Out

This validation used 4 SR670 servers, but the underlying architecture is flexible and adding more computational servers is allowed. The performance in the main use case – shared usage with training jobs taking all the GPUs in a server – should be maintained even when the computational resources are expanded. The reason for this is that this validation utilized only a single NFS mount point, while the DM Series storage server used here provide four of such access points. This can be also leveraged in the scale out case and we expect that the compute-compute communication would become a bottleneck rather than data access.

7.2 CPU Adjustments

We used a Skylake Intel Xeon Gold 6142 processor for this validation. The equivalent Cascade Lake CPU is an Intel Xeon Gold 6242 processor, which we expect would deliver the same performance because this workload is not CPU bound. The best value CPU for this setup is Intel Xeon Silver 4216, which offers 16 cores while using only 100W. This should be sufficient for driving common AI training workloads. On the other hand, if CPU performance is an important factor, we recommend upgrading to Intel Xeon Platinum 8280.

7.3 GPU Adjustments

This validation used 4 V100 GPUs per server, but 8 T4 GPUs can be used instead. Such architecture would be more applicable to the inference workloads although AI training would still be feasible. In terms of data access, this would require less data bandwidth because combined performance of 8 T4 cards is less than that of 4 V100 GPU. Thus, the storage system designed in this work is expected to perform well even for T4 GPUs.

7.4 Storage Capacity Increase

DM5000F can be scaled out up to 12 High Availability pairs to 24 PB of total storage.

7.5 Other ThinkSystem Servers

This validation was based on ThinkSystem SR670, which is currently provides highest GPU density in Lenovo server portfolio a total of 4 GPUs for a 2 V100 GPUs/ 1U ratio. Other GPU servers have lower ratios along with less total number of GPUs. This means that their data requirements would be less than that of a SR670.

Because of this, it is expected that this architecture would provide equivalent (or better) performance for other GPU-enabled servers such as SR650, SR655, or SR665.

7.6 Other DM Series Storage Options

This validation used DM5000F storage server, which is the lowest storage option available in ThinkSystem portfolio. **Other DM Series options, DM7000F and DN7100F, provide better performance and more features and thus it is expected that they would provide the same or better performance as the one obtained in this work.**

8 Deployment considerations

This entry-level architecture was designed to be small scale and cost effective. As such, it is assumed that deployment reuses existing data center resources such as rack and network switching and thus these items are not covered in this document.

Additionally, as pointed out in previous section, adjustment to this architecture are possible and for many cases we expect the same or better performance as in this work.

9 Conclusion

This solution provides a flexible scale-out architecture that is ideal for entry-level enterprise AI. DM Series storage delivers the same or better performance as local SSD storage and provides the following benefits to data scientists, data engineers, and IT decision makers:

- Effortless sharing of data between AI systems, analytics, and other critical business systems. This reduces infrastructure overhead, improves performance, and streamlines data management across the enterprise.
- Independently scalable compute and storage to minimize costs and improve resource utilization.
- Streamlined development and deployment workflows using integrated snapshots and clones for instantaneous and space-efficient user workspaces, integrated version control, and automated deployment.
- Enterprise-grade data protection for disaster recovery and business continuance.

10 Appendix: Lenovo Bill of materials

This appendix contains the bill of materials (BOMs) for computational servers and a storage server.

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.

Within a specific BOM section, optional items are numbered with alternatives shown as lower-case letters. For example, a Fibre Channel adapter for a compute server is only needed for shared storage connected through a SAN.

10.1 BOM for compute servers

Part #	Description	Quantity
	SR670 4xV100S HPC/AI Training	4
7Y37CTO1WW	SR670 4xV100S HPC/AI Training : ThinkSystem SR670 - 3yr Warranty	4
B3XX	ThinkSystem SR670 2.5" Chassis - 8 Drive, 4 GPU	4
B4HG	Intel Xeon Gold 6242 16C 150W 2.8GHz Processor	8
B4H3	ThinkSystem 32GB TruDDR4 2933MHz (2Rx4 1.2V) RDIMM	48
B42K	1U 2.5" SATA/SAS 8-Bay Backplane	4
5977	Select Storage devices - no configured RAID required	4
B5MR	ThinkSystem SR670 430-8i SAS/SATA HBA	4
B4Y7	ThinkSystem 2.5" SS530 3.2TB Performance SAS 12Gb Hot Swap SSD	32
AUKX	ThinkSystem Intel X710-DA2 PCIe 10Gb 2-Port SFP+ Ethernet Adapter	4
AUZY	ThinkSystem I350-T2 PCIe 1Gb 2-Port RJ45 Ethernet Adapter	4
B34S	ThinkSystem NVIDIA Tesla V100 32GB PCIe Passive GPU	16
B3YC	2000W Platinum PSU	8
B0N4	2.0m, 10A/100-250V, C13 to C14 Jumper Cord	8
AVUT	ThinkSystem XClarity Controller Standard to Advanced Upgrade	4
B47V	ThinkSystem SR670 Slide Rail Kit	4
A51N	1.5m Passive DAC SFP+ Cable	4
B0MK	Enable TPM 2.0	4
B7Y0	Enable IPMI-over-LAN	4
B3Y8	ThinkSystem SR670 1-3 Slot PCIe x16 FHFL Riser Kit	12
B3Y5	ThinkSystem SR670 WW Lenovo LPK	4

Part #	Description	Quantity
B47U	ThinkSystem SR670 Service Label	4
AUSF	Lenovo ThinkSystem 2U MS CPU Performance Heatsink	8
B7HC	ThinkSystem SR670 CLX 2-CPU, 24 DIMM System Board	4
B3Y2	GPU Cage1 Cable Pair	4
B3Y3	GPU Cage2 Cable Pair	4
AURS	Lenovo ThinkSystem Memory Dummy	48
B4EA	Model Number Label	4
B3Y4	SR670 Agency Label	4
B4E7	EIA NamePlate	4
AWF9	ThinkSystem Response time Service Label LI	4
AUTV	ThinkSystem large Label for non-24x2.5"/12x3.5"/10x2.5"	4
B5DM	SR670 RAID SAS Cable	4
B3YB	SR670 System Packaging	4
AUTJ	ThinkSystem common Intel Label	4
AUTA	XCC Network Access Label	4
B4RM	ThinkSystem SR670 Front IO Label	4
B4RN	Blaze Jumper Setting	4
AVJ2	ThinkSystem 4R CPU HS Clip	8
B0ML	Feature Enable TPM on MB	4
5AS7A02045	Hardware Installation Server (Business Hours)	4
7S09CTO1WW	Lenovo HPC AI LiCO Software	4
B1YA	Lenovo HPC AI LiCO Software w/3Yr S&S	24
7S0FCTO1WW	Red Hat Linux w/Lenovo Support	4
S0N6	RHEL Server Physical or Virtual Node, 2 Skt Standard Subscription w/Lenovo Support 3Yr	4

10.2 BOM for Storage Servers

Part #	Description	Quantity
	Controller	1
7Y41CTOLWW	Controller : Lenovo ThinkSystem DM5000F All Flash Array	1

Part #	Description	Quantity
B38L	Lenovo ThinkSystem Storage 2U24 Chassis	1
B5RJ	DM Series Premium Offering	1
B39G	Lenovo ThinkSystem DM Series DM3000/DM5000 Cntr, 10G BaseT	2
B65T	Lenovo ThinkSystem 46.1TB (6x 7.68TB, 2.5", SSD) Drive Pack for DM5000F	2
A3RG	0.5m Passive DAC SFP+ Cable	2
B4BP	Lenovo ThinkSystem Storage USB Cable, Micro-USB	1
6311	2.8m, 10A/100-250V, C13 to C14 Jumper Cord	2
B79W	Lenovo ThinkSystem DM Series ONTAP 9.6 Encryption	1
B0W1	3 Years	1
B46X	Essential Service	1
B472	Configured with Lenovo ThinkSystem DM5000F	1
B4SF	DM Series CIFS Protocol License	2
B4SG	DM Series NFS Protocol License	2
B4SH	DM Series iSCSI Protocol License	2
B4SJ	DM Series FCP Protocol License	2
B4SK	DM Series SnapMirror License	2
B4SL	DM Series SnapRestore License	2
B4SM	DM Series FlexClone License	2
B4SN	DM Series Software Encryption License	2
B4SP	DM Series SnapManager License	2
B4SU	TPM	2
B5AZ	DM Series SnapVault License	2
B7AQ	SnapMirror Synchronous	2
B38Y	Lenovo ThinkSystem Storage Rack Mount Kit 2U24/4U60	1
B4CX	Lenovo ThinkSystem DM Series 2U Accessory	1
B39L	Lenovo ThinkSystem DM Series 2U24 Bezel	1
B38Z	Lenovo ThinkSystem Storage SFF Drive Filler	12
B4BG	Lenovo ThinkSystem Storage 2U24 System Label	1
B396	Lenovo ThinkSystem DM5000F Product Label	1

Part #	Description	Quantity
B4AW	Lenovo ThinkSystem Storage Packaging 2U	1
B39C	Lenovo ThinkSystem DM Series Ship Kit (RoW)	1
5PS7A18279	Essential- 3Y 24x7x4 ThinkSystem DM5000F AFA	1
5WS7A32554	Essential- 3Y 24x7x4 DM5000F 92TB (12x 7.68TB SSD) Pack	1
5AS7A02067	Hardware Installation Storage (Business Hours)	1
5MS7A79862	ThinkSystem DM Remote Deployment	1
	Auto-Derived Part Items	
AU16	0.5m External MiniSAS HD 8644/MiniSAS HD 8644 Cable	2

Resources

[ThinkSystem SR670 Rack Server](#)

[ThinkSystem DM Series Storage](#)

[Lenovo/NetApp Compute Storage Technical Report](#)

Document history

Version 1.0

July 2020

First version validated on SR670 and DM5000F

Trademarks and special notices

© Copyright Lenovo 2019.

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

Lenovo, the Lenovo logo, ThinkSystem, ThinkCentre, ThinkVision, ThinkVantage, ThinkPlus and Rescue and Recovery are trademarks of Lenovo.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel 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.