



Technical Report

NetApp ONTAP and Lenovo ThinkSystem SR670 for AI and ML Model Training Workloads

Karthikeyan Nagalingam, NetApp
Miroslav Hodak, Lenovo

February 2020 | TR-4810

Abstract

This solution document describes a cost-effective, entry-level compute and storage architecture to deploy GPU-based artificial intelligence (AI) training on NetApp® storage controllers and Lenovo® ThinkSystem™ servers. The setup is designed as a shared resource for small to medium-sized teams running multiple training jobs in parallel.

This document provides performance data for the industry-standard MLPerf benchmark evaluating image classification training with TensorFlow on V100 GPUs. To measure performance, we used ResNet50 with the ImageNet dataset, a batch size of 512, half precision, CUDA, and cuDNN. We performed this analysis using four-GPU SR670 servers and an entry-level NetApp storage system. The results show highly efficient performance across the multiple use cases tested here—shared, multiuser, multijob cases, with individual jobs scaling up to four servers. Large scale-out jobs were less efficient but still feasible.

TABLE OF CONTENTS

1	Introduction	4
1.1	Target Audience	4
1.2	Solution Architecture	4
2	Technology Overview	5
2.1	NetApp AFF Systems	5
2.2	NetApp ONTAP 9	5
2.3	NetApp FlexGroup Volumes	6
2.4	Lenovo Overview	6
3	Test Plan Summary	7
4	Test Results Summary	7
5	Test Configuration Details	8
5.1	Solution Architecture	8
6	Test Procedure and Detailed Results	9
6.1	Image Recognition Training Using ResNet in ONTAP	9
6.2	Data Read Speed: Local vs Network Storage	10
6.3	Shared use case: Multiple simultaneous jobs	10
6.4	Exclusive use case: Scale out	11
6.5	Exclusive use case: Scale out with a faster compute network	12
7	Architecture Adjustments	14
7.1	CPU Adjustments	14
7.2	Storage Capacity Increase	14
8	Conclusion	14
	Acknowledgments	14
	Where to Find Additional Information	14
	Version History	15

LIST OF TABLES

Table 1)	Test result summary	7
Table 2)	Base components for the solution architecture	8
Table 3)	Storage configuration	9

LIST OF FIGURES

Figure 1)	Architecture overview	4
Figure 2)	NetApp FlexGroup volumes	6
Figure 3)	Network topology of tested configuration	9

Figure 4) Aggregate Images per second on concurrent training models 10
Figure 5) Runtime in minutes on concurrent training models 11
Figure 6) ResNet50 scale out results for the runtime in minutes. 11
Figure 7) ResNet50 scale-out results in images per second. 12
Figure 8) Lenovo servers are connected through 100 GbE compute network..... 12
Figure 9) ResNet50 scale-out results in images per second for 100 GbE compute network..... 13
Figure 10) ResNet50 scale-out results runtime for 100 GbE compute network..... 13

1 Introduction

This solution document focuses on an entry-level, clustered architecture using NetApp storage and Lenovo servers optimized for artificial intelligence (AI) workflows. It is meant for 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, because most day-to-day AI training jobs are single node.

This document covers testing and validation of a compute and storage configuration consisting of four-GPU Lenovo SR670 servers and an entry-level NetApp AFF storage system such as an A200 or A220. This architecture provides an efficient and cost-effective solution for small and medium-sized organizations just starting out with AI that require the enterprise-grade capabilities of NetApp ONTAP® cloud-connected data 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 possible time to market from AI initiatives

1.2 Solution Architecture

This Lenovo ThinkSystem server and NetApp ONTAP 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 Lenovo SR670 servers alongside a single NetApp storage system.

Figure 1) Architecture overview.



This NetApp and Lenovo solution offers the following key benefits:

- Highly efficient and cost-effective performance when executing multiple training jobs in parallel
- Scalable performance based on different numbers of Lenovo servers and different models of NetApp storage controllers
- 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

2 Technology Overview

2.1 NetApp AFF Systems

State-of-the-art NetApp AFF 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, AFF systems help accelerate, manage, and protect business-critical data.

Entry-level NetApp AFF storage systems are based on FAS2650 hardware and SSD flash media. For NAS workloads, a single entry-level AFF 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 10Gb FC connectivity
- Up to 15TB SSDs with multistream write
- NetApp ONTAP 9, with a complete suite of data protection and replication features for industry-leading data management

NetApp also offers other storage systems, such as the AFF A800, AFF A700, and AFF A320 that provide higher performance and scalability for larger-scale deployments.

2.2 NetApp ONTAP 9

ONTAP 9, the latest generation of storage management software from NetApp, 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.

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 automatically tiers cold data to public and private cloud storage options, including Amazon Web Services (AWS), Azure, and NetApp StorageGRID® object storage.

Accelerate and Protect Data

ONTAP delivers superior levels of performance and data protection and extends these capabilities in the following ways:

- **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.
- **NetApp Volume Encryption.** ONTAP offers native volume-level encryption with both onboard and external key management support.

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 (NetApp 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 NetApp 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 2) is a single namespace made up of multiple constituent member volumes that is managed and acts like a NetApp 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) NetApp FlexGroup volumes.



2.4 Lenovo ThinkSystem Portfolio

Lenovo ThinkSystem servers feature innovative hardware, software, and services that solve customers' 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 to fully understand the value for their particular use case. To improve time to value, this customer-centric approach gives customers proofs of concept for solution development platforms that are ready to use and optimized for AI.

2.5 Lenovo SR670

The Lenovo ThinkSystem SR670 delivers optimal performance for accelerated AI and high-performance computing (HPC). Supporting up to either four large or eight small GPUs per 2U node, the SR670 is suited for the computationally intensive workload requirements of ML, DL, and inference.

With the latest scalable Intel Xeon CPUs that support high-end GPUs, including the NVIDIA Tesla V100 and T4, the ThinkSystem SR670 delivers optimized accelerated performance for AI and HPC workloads.

Because more workloads 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 is an optimized, enterprise-grade solution for deploying accelerated HPC and AI workloads in production, maximizing system performance while maintaining data center density.

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 [TF cnn benchmarks](#) training script was used to drive the AI training. The script contains implementations of several popular conventional models and is designed to be as fast as possible. It can be run on a single machine or run in distributed mode across multiple hosts.

3 Test Plan

In this validation, we performed image recognition training as specified by MLPerf v0.5. Specifically, we trained the ResNet v1.5 model with the ImageNet dataset until we reached an 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.

4 Test Results

Table 1 summarizes the results for all tests performed for this solution.

Table 1) Test results 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 NetApp 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 NetApp storage system delivered training performance comparable to local SSD storage while enabling easy sharing of data between servers.
--	--

5 Test Configuration

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

5.1 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
Lenovo ThinkSystem servers	<ul style="list-style-type: none"> 4 SR670 servers each with 4 NVIDIA 32GB V100 GPU cards Each server contains 2 Intel Xeon Gold 6142 CPUs (16 physical cores) and 384GB RAM
Linux (CentOS – 7.6 with CUDA 10.0)	
Entry-level NetApp AFF storage system (HA pair)	<ul style="list-style-type: none"> NetApp ONTAP 9 software 24 x 960GB SSDs NFS protocol 1 interface group (ifgrp) per controller, with 4 logical IP addresses for mount points

In this validation, we used ResNet v1.5 with the ImageNet basis set as specified by MLPerf v0.5. The dataset is stored in a NetApp AFF storage system with the NFS protocol. The SR670s were connected to the NetApp AFF storage system 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 TF. The TFRecord data is packaged into 1024 separate files, each approximately 136MB in size.

Figure 3) Network topology of tested configuration.




Table 3) Storage configuration.

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

The /netappLenovo_AI_fg folder contains the dataset used for ResNet validation.

6 Test Procedure and Detailed Results

6.1 Image Recognition Training Using ResNet in ONTAP

We ran the ResNet50 benchmark with one, two, and four SR670 servers. This test used the [TensorFlow benchmark script](#) and TensorFlow container version 1.13 from DockerHub to run the training.

We used the following test procedure in this validation.

1. 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
```

2. We ran the benchmark script with the ImageNet dataset in server storage (local SSD storage) as well as on the NetApp AFF storage system.
3. We validated network and local storage performance using the `dd` command.
4. For the single-node run, we used the following command:

```

--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. We used four parameter servers per node, 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 comparison between single-node and distributed cases.

6.2 Data Read Speed: Local Versus Network Storage

The read speed was tested by using the `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 can deliver data at a rate essentially equal to that of local storage.

6.3 Shared Use Case: Multiple Independent Simultaneous Jobs

This test simulated the expected use case for this solution—multijob, multiuser AI training. Each node ran its own training while using the shared network storage. The results are displayed in Figure 4, which shows that the solution case provided excellent performance with all jobs running at essentially the same speed as individual jobs. The total throughput scaled linearly with the number of nodes.

Figure 4) Aggregate images per second on concurrent training models.


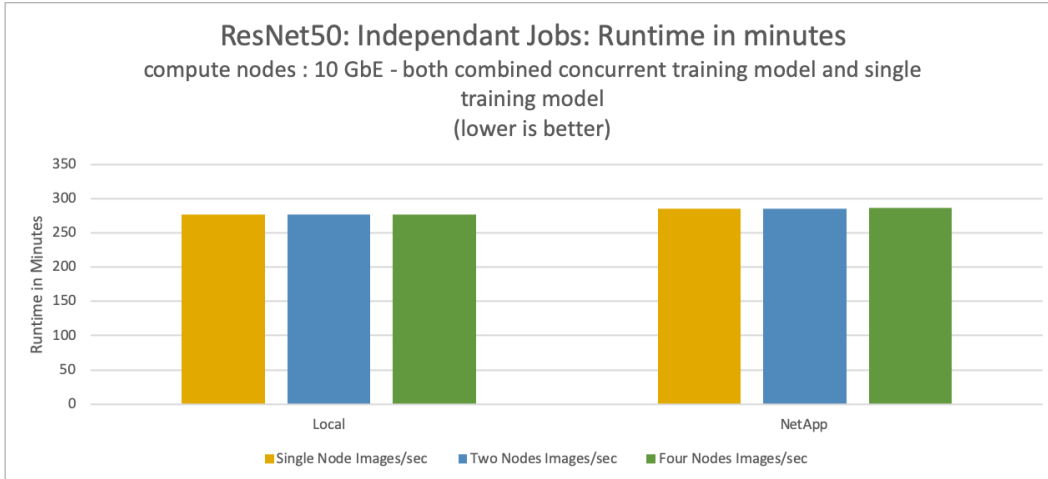


Figure 5) Runtime in minutes on concurrent training models.



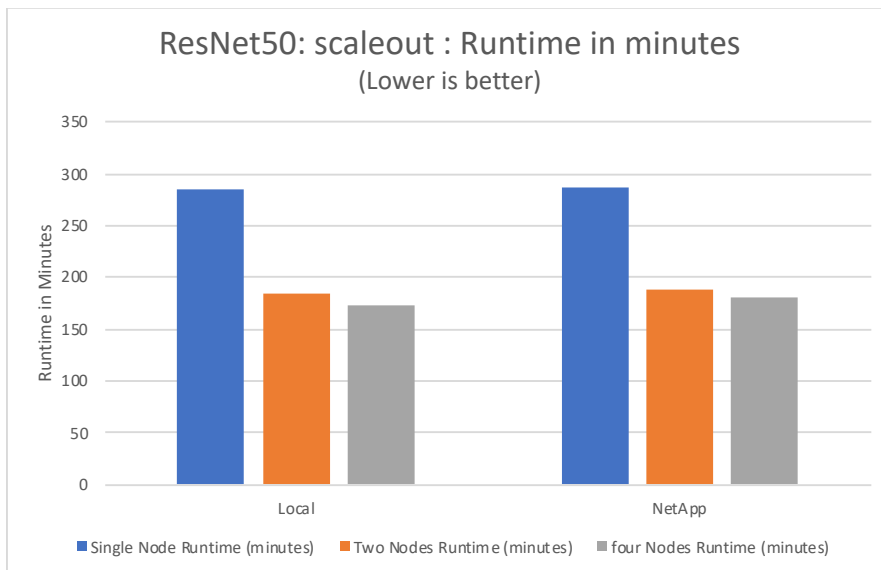
Four Independent training models used four GPUs from each server. The average runtime for the training model was 286 minutes and 20 seconds. The individual runtimes were 285 minutes and 35 seconds, 286 minutes and 11 seconds, 286 minutes and 27 seconds, and 287 minutes and 7 seconds. The average images per second for the training model were 4,600, and the individual images per second were 4,615, 4,607, 4,596, and 4,582.

Based on our validation, one independent training model with a NetApp data runtime was 285 minutes and 19 seconds with 4,611 images/sec. One independent training model with a local data (DAS) runtime was 277 minutes and 44 seconds with 4,724 images/sec. During those runs the average GPU utilization was 96%, as reported by nvidia-smi. Note that this average includes the testing phase, during which GPUs were not used, while CPU utilization was 40% as measured by mpstat. This demonstrates that the data delivery rate is sufficient in each case.

6.4 Exclusive Use Case: Dedicated Training Model Scale Out

Figures 6 and 7 show the scale-out results from our lab validation. Figure 6 shows the test runtime in minutes when using one, two, or four Lenovo SR670 servers.

Figure 6) ResNet50 scale-out results for runtime in minutes.

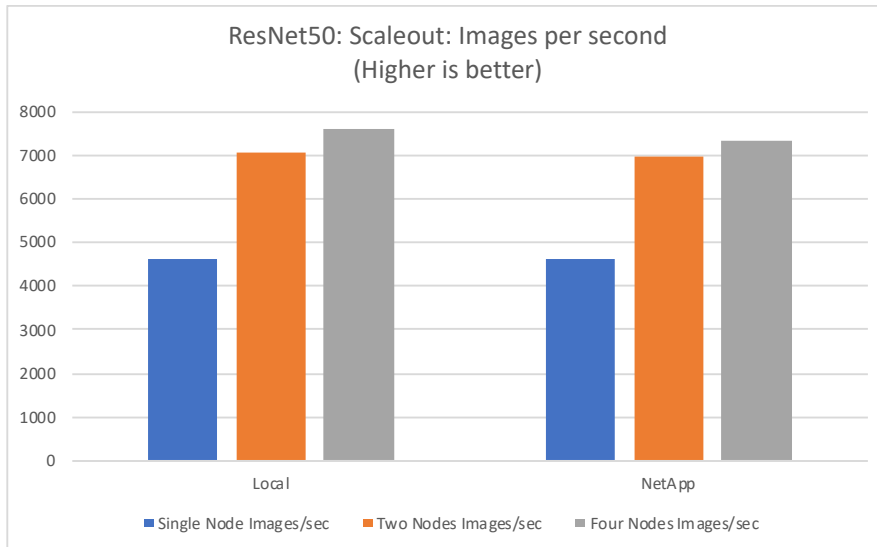


Note the following details about Figure 6:

- The runtime results for one, two, and four nodes in both the NetApp storage controller and local DAS storage.

- The NetApp AFF storage system provided performance similar to the local server SSD storage. Lower numbers are better, indicating how quickly the benchmark completed the task.

Figure 7) ResNet50 scale-out results in images per second.

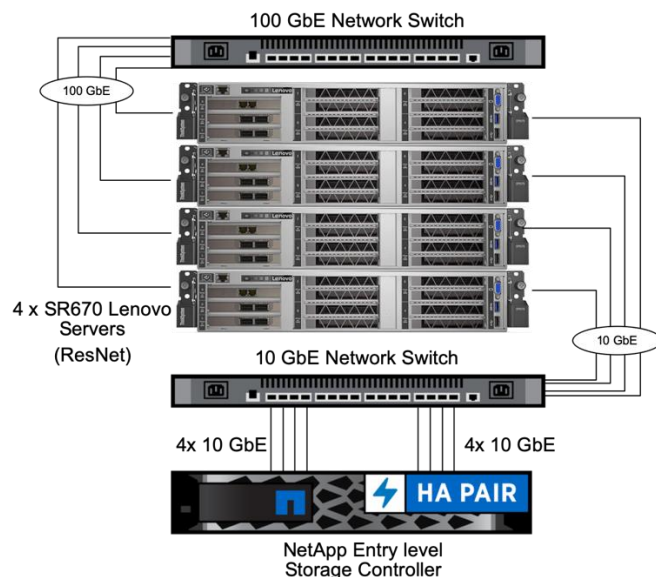


- In Figure 7, a higher number is better, indicating that the benchmark processed more images per second.
- Based on this validation with a single NetApp NFS export path, the average throughput was 4,603 images/sec, and the runtime was 288 minutes for 61 epochs.
- The less-than-linear scaling was because the 10GbE network is insufficient for node-to-node communication. Much better scaling can be obtained with a faster network, as demonstrated in section 6.5. Note that the NetApp storage connected over a 10GbE network provides enough data bandwidth for the cluster.

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

To improve network performance and thus scaling performance, we connected the compute nodes to a faster network—Mellanox HDR InfiniBand in IPoIB mode. In the current generation of ThinkSystem servers, this network provides equivalent bandwidth and lower latency than 100GbE (Figure 8).

Figure 8) Lenovo servers connected through 100GbE compute network.



The performance tests (Figure 9 and Figure 10) were run in the same way as in the previous section. The results show that the faster compute network improves scale out, providing near linear scaling. Network storage shows performance equivalent to that of the local storage.

Figure 9) ResNet50 scale-out results in images per second for 100GbE compute network.

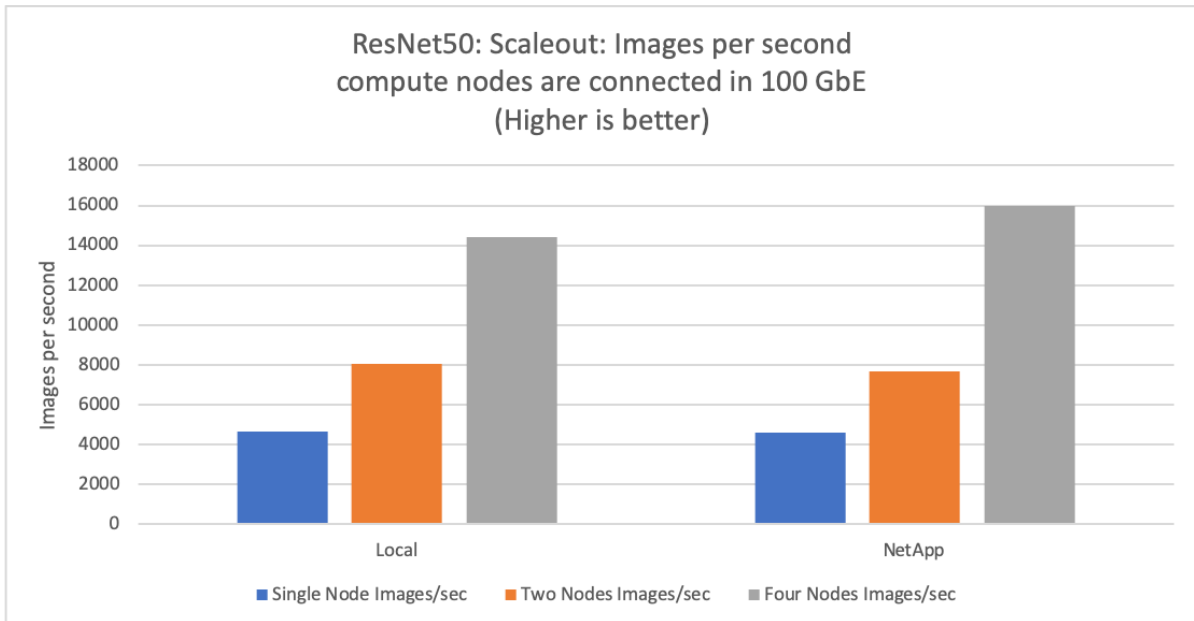
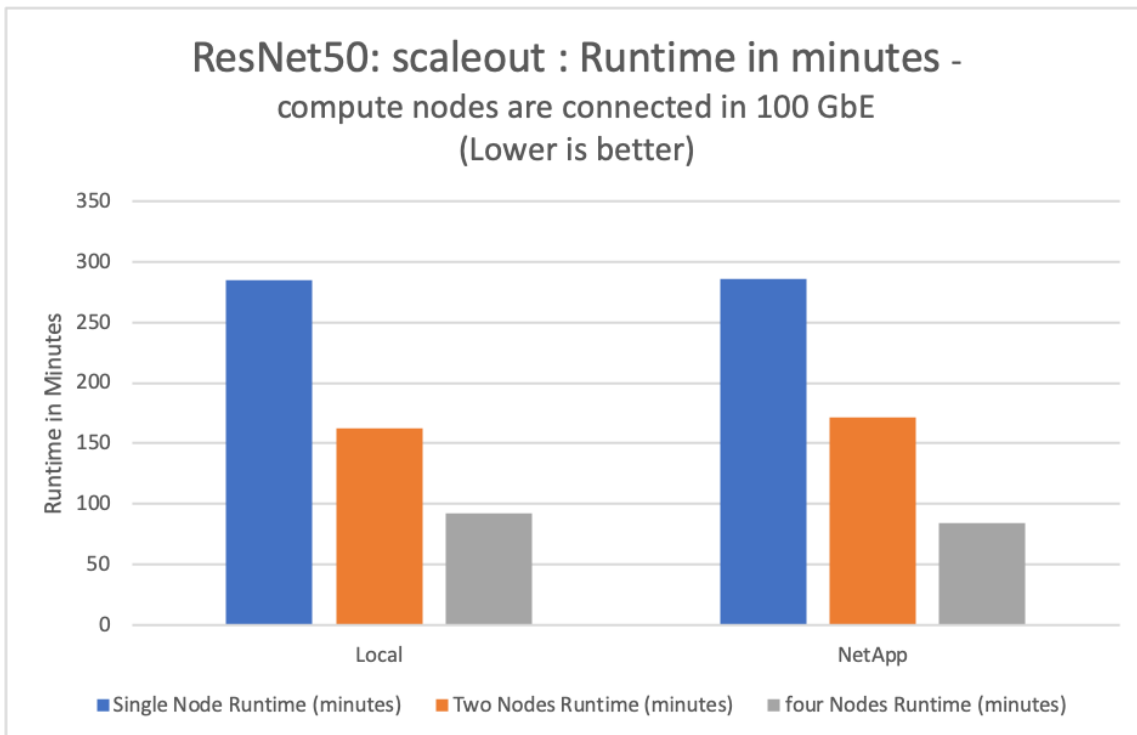


Figure 10) ResNet50 scale-out results runtime for 100GbE compute network.



Based on our validation for four nodes with NetApp shared data, the training model's runtime was 91 minutes and 54 seconds with 15,948 images/sec. For the same four nodes, the results for training models with local data (DAS) were 83 minutes and 55 seconds with 14,394 images/sec. Similarly, the results for two nodes with NetApp shared data demonstrated a training model runtime of 162 minutes and 41 seconds with 8,073 images/sec. For the same two-node training models with local data (DAS), we recovered a runtime of 170 minutes and 57 seconds with 7,680 images/sec.

7 Architecture Adjustments

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

7.1 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. Furthermore, NetApp recommends the Intel Xeon Silver 4216 processor as the best value for this setup. If maximum performance is needed, NetApp recommends the Intel Xeon Platinum 8280 processor.

7.2 Storage Capacity Increase

Based on your storage capacity needs, you can increase the share storage (NFS volume) on demand by using the command `volume size -vserver <vserver name> -volume <volumename> -new-size +<size>GB`. You can do this from the CLI or from the NetApp web interface of the storage controller as admin user.

8 Conclusion

This NetApp and Lenovo solution is a flexible scale-out architecture that is ideal for entry-level enterprise AI. NetApp storage delivers the same or better performance as local SSD storage and offers 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 data sharing 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.

Acknowledgments

The authors gratefully acknowledge the contributions that were made to this technical report by our esteemed colleagues from Lenovo, J.J. Falkanger and Doug Toth, DCG.

The authors would also like to acknowledge the contributions of key NetApp team members Sung-Han Lin, David Arnette, Mike McNamara, and Erik Mulder.

Where to Find Additional Information

To learn more about the information described in this document, refer to the following documents and/or websites:

- NetApp All Flash Arrays product page
<https://www.netapp.com/us/products/storage-systems/all-flash-array/aff-a-series.aspx>
- NetApp ONTAP data management software product page
<http://www.netapp.com/us/products/data-management-software/ontap.aspx>
- MLPerf
<https://mlperf.org>
- TensorFlow benchmark
<https://github.com/tensorflow/benchmarks>

Version History

Version	Date	Document Version History
Version 1.0	February 2020	Initial release

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 2020 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc.

United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

Lenovo, the Lenovo logo, and ThinkServer are trademarks or registered trademarks of Lenovo. Other company, product, and service names may be trademarks or service marks of others.