

Using the ThinkServer SA120 Disk Array

Lenovo Enterprise Product Group

Version 1.0

March 2014

©2014 Lenovo. All rights reserved.



LENOVO PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. This information could include technical inaccuracies or typographical errors. Changes may be made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any performance data contained herein was determined in a controlled environment; therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems, and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk.

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

Intel and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. Microsoft, Windows and Exchange Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Apache Hadoop and Hadoop are trademarks of the Apache Software Foundation. LSI and LSI Design logo are trademarks or registered trademarks of LSI Corporation.

Contents

1.0	Introduction.....	5
2.0	SA120 Overview	5
3.0	SA120 Configuration.....	7
3.1	Connectivity Considerations.....	7
3.2	Server and SA120 Connectivity Scenarios	8
3.3	LSI Management Tools	11
4.0	Sample Use Cases and Benefits.....	12
4.1	File Server / Unstructured Data.....	12
4.2	Backup to Disk	12
4.3	Database Acceleration.....	13
4.4	Building Block for Distributed Storage	13
5.0	Conclusion	15
6.0	References.....	16

List of Figures

Figure 1 – SA120 Front View	5
Figure 2 – SA120 Rear View with Optional Second PSU, SAS Controller and SSDs.....	6
Figure 3 – Basic SA120 Daisy-Chain	8
Figure 4 – High-Capacity SA120 Daisy-Chain.....	9
Figure 5 – Controller and Cable Redundancy	10
Figure 6 – Example Including SSD in First SA120.....	11

List of Tables

Table 1 – Server and Adapter Support	7
--	---

1.0 Introduction

There are many choices for data storage today. Some of the factors influencing storage decisions are performance, reliability, scalability and acquisition cost. In addition to these basic requirements, IT managers are also evaluating new storage strategies, such as distributed storage, that are expected to provide better flexibility and reduced total cost of ownership.

Traditionally, Direct Attached Storage (DAS) has been limited to a single server- single application use case. DAS is now also a critical part of distributed architectures.

With legendary “Think” quality and features that provide multiple levels of performance, the ThinkServer SA120 DAS array is the natural choice when expanding your ThinkServer system for a single server application, or as part of a larger distributed solution. This document provides sample use cases and configuration guidelines for the ThinkServer SA120, including connectivity diagrams and overview of configuration tools.

2.0 SA120 Overview

The ThinkServer SA120 is a 2U rack-mountable storage enclosure¹ that provides both 2.5 inch- and 3.5-inch drive bays in a single enclosure. The SA120 is unique in that twelve 3.5-inch Hard Disk Drives (HDDs) mount in the front while four 2.5-inch drives mount in the rear of the enclosure. The rear 2.5- inch bays are reserved exclusively for optional Intel Enterprise solid-state drives (SSDs), providing an optimal tiered storage platform in one dense enclosure.



Figure 1 – SA120 Front View

The SA120 supports direct attached 6Gbps SAS connectivity and integrates seamlessly with ThinkServer rack and tower models via supported ThinkServer LSI SAS and RAID adapters. Each

¹ Tower conversion kit planned for availability after initial SA120 launch

dual-port adapter supports cascading up to four SA120 systems per port, providing a maximum capacity of (96) 3.5-inch HDDs and (32) 2.5-inch SSDs on a single server.



Figure 2 – SA120 Rear View with Optional Second PSU, SAS Controller and SSDs

Servicing the SA120 is simple with tool-less components in addition to hot-swap disk drives, SAS Input/Output Controller Cards (IOCCs), fans and power supplies. Common power supplies and drives are shared with other ThinkServer systems, increasing convenience and reducing overall costs.

3.0 SA120 Configuration

3.1 Connectivity Considerations

The SA120 is supported with current and prior generation ThinkServer systems with Operating System support following the connected server². Three ThinkServer LSI adapters are available for connecting to the SA120:

- ThinkServer LSI9207-8e 6Gb SAS HBA³
 - PCIe Gen 3
 - No RAID, JBOD only
- ThinkServer LSI9280-8e 6Gb SAS RAID Adapter
 - PCIe Gen 2
 - RAID 0, 1, 10, 5, 50, 6, 60
 - 512MB cache
 - Battery Backup option
- ThinkServer LSI9286CV-8e 6Gb SAS RAID Adapter
 - PCIe Gen 3
 - RAID 0, 1, 10, 5, 50, 6, 60
 - 1GB cache
 - CacheVault option
 - CacheCade and FastPath I/O acceleration options

Each card has two external 6Gbps SAS ports (no internal ports). Table 1 shows supported adapter cards and servers.

Table 1 – Server and Adapter Support

		9207-8e	9280-8e	9286CV-8e
3 rd Generation	RD630	Yes	Yes	No
	RD530	Yes	Yes	No
	RD430	Yes	Yes	No
	RD330	Yes	Yes	No
	TS430	Yes	Yes	No
4 th Generation	RD640	Yes	Yes	Yes
	RD540	Yes	Yes	Yes
	RD440	Yes	Yes	Yes
	RD340	Yes	Yes	Yes
	TD340 ³	Yes	Yes	Yes
	TS440	Yes	Yes	Yes

² For supported operating systems on ThinkServer models, see the link in Section 6: References.

³ Available in May 2014.

When choosing an adapter card, consider these factors:

- Server model
- If hardware RAID is required
- Performance requirements
- Project budget

SAS cables are available in multiple lengths. Care should be taken when selecting cables to ensure the length is adequate, as the SA120 and connected server may not be adjacent in the rack.

3.2 Server and SA120 Connectivity Scenarios

Now, with an understanding of what is required in the connected ThinkServer, this section explores supported connectivity scenarios.

3.2.1 Basic Daisy Chain / Cascade

This is the most basic connectivity scenario. Figure 3 shows the maximum number of SA120 enclosures on a single adapter port, although one, two, three or four enclosures may be used. Only one SAS IOCC is required per SA120. The advantage of this configuration is its low cost compared to other connectivity scenarios. Note that the Figure 3 example has no optional SSD in the rear, although SSD may be used in this basic configuration.

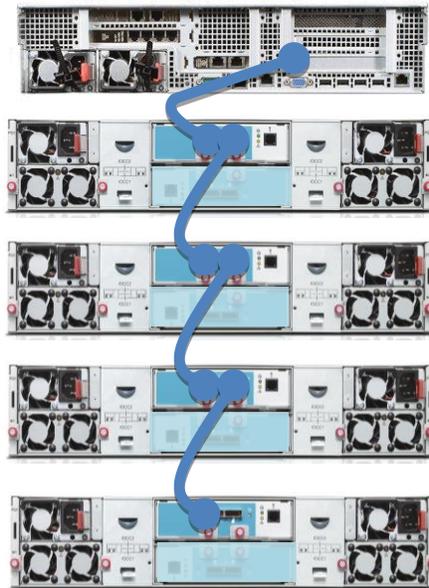


Figure 3 – Basic SA120 Daisy-Chain

Figure 4 shows a variation on the basic daisy-chain scenario, with four SA120 enclosures connected to one port of the adapter and four SA120 enclosures connected to the other adapter port. This depicts the maximum scalability of the SA120 solution.

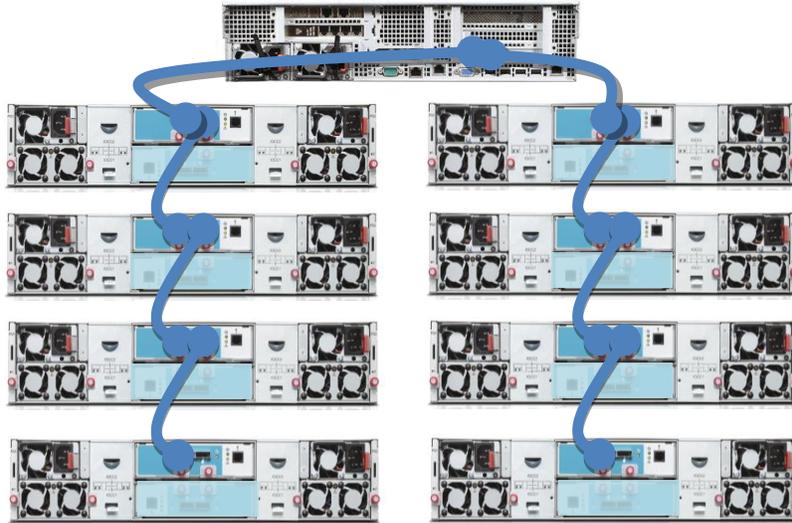


Figure 4 – High-Capacity SA120 Daisy-Chain

3.2.2 Controller and Cable Redundancy

This connectivity method requires two SAS IOCCs and two cables per SA120, and ensures the application continues if an IOCC or attached cable were to fail. Maximum expansion when using redundant paths is four SA120 enclosures.

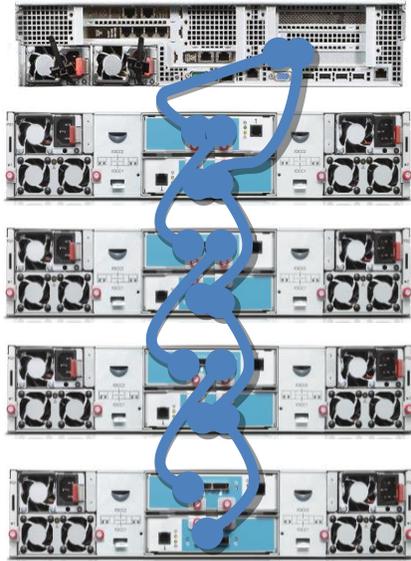


Figure 5 – Controller and Cable Redundancy

3.2.3 Data Acceleration with SSD

Because the SA120 has the unique ability to store two different tiers of data with 12 large-capacity drives in the front and four fast SSD in the rear, it can easily become part of an application data acceleration solution.

Many workload intensive applications are constrained by the performance limitations of the physical spinning media within a HDD. Lenovo ThinkServer enterprise SSDs from Intel provide consistently high I/O performance that is generally 150 times faster than 15K RPM HDDs, and have lower and more consistent access times.

Although the price per gigabyte of SSDs continues to decline, they are still considerably more expensive relative to larger and slower HDDs. A cost effective strategy is to mix SSD with the HDD in a hybrid approach for the application's data. When applying this strategy, a traditional method is to manually and selectively place the more frequently accessed hot data on the SSD and the remaining data on the HDD. While effective for increasing performance, this can be time consuming and require further administration when the application or data changes. The optional CacheCade software for the ThinkServer LSI9286CV-8e 6Gb SAS RAID Adapter intelligently and dynamically handles the data placement without the need to monitor and configure applications. When reading data, CacheCade automatically copies the most frequently accessed data on HDD volumes to a much faster tier of SSD cache, removing many of the latency bottlenecks associated with traditional HDD-based data volumes.

The example shown in Figure 6 includes two SA120 enclosures. Similar to non-SSD configurations, up to eight SA120 enclosures may be connected to one server in a single path

configuration or four in a redundant path configuration (redundant paths shown). This example shows the first SA120 with four SSDs and the second with none. When choosing the number and size of SSDs for CacheCade, it is important to recognize that the maximum CacheCade cache pool volume size is 512GB. In addition, the CacheCade volume supports RAID levels 0 or 1. CacheCade can be configured as read cache only, or read and write cache together. If write cache is enabled, it is recommended to configure the cache pool as a RAID 1 volume when protection of write-cached data is desired.

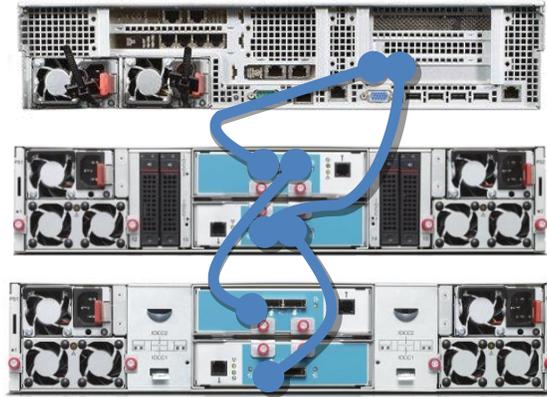


Figure 6 – Example Including SSD in First SA120

3.3 LSI Management Tools

Similar to internal ThinkServer RAID adapters based on LSI technology, the ThinkServer external SAS and RAID adapters may also be managed and configured with LSI tools.

Because the ThinkServer LSI9207-8e 6Gb SAS HBA is not a RAID adapter, no configuration is required. This adapter presents each attached drive to the operating system independently. It is up to the system administrator to create appropriate aggregated (or RAID-protected) volumes with operating system or software-based tools. The 9207-8e does provide a pre-boot firmware utility where specifications and status of the adapter and attached drives may be viewed; this interface is accessed during boot by pressing Ctrl-C when the adapter is initializing.

Both the 9280-8e and 9286CV-8e RAID Adapters utilize the pre-boot WebBIOS Configuration Utility or MegaRAID Storage Manager (MSM) application for configuration of drives and volumes in the external SA120 enclosure. MSM also provides monitoring and alerting for the SA120 and associated drives. The WebBIOS interface is accessed during boot by pressing Ctrl-H when the adapter is initializing. MSM can be downloaded from <http://support.lenovo.com>, and requires installation at the operating-system level of the server. If the server already has an internal LSI-

based RAID adapter and MSM is installed, the same instance of MSM will display and manage both internal and external adapters, in addition to the SA120 enclosure. MSM will also display, and provide alerts from, the ThinkServer LSI9207-8e, but no configuration is required for this SAS HBA.

The 9280-8e adapter is very similar to the ThinkServer RAID 700 adapter for internal drives, and the 9286CV-8e is similar to the ThinkServer RAID 710 adapter. For further information on configuring RAID and CacheCade, see links to separate guides *Configuring ThinkServer RAID 700 Adapters* and *Configuring ThinkServer RAID 710 Adapters* in the reference section of this document.

4.0 Sample Use Cases and Benefits

ThinkServer SA120 fits well into many storage expansion scenarios, because of its simplicity, scalability and flexible performance choices.

This section explores a few sample use cases, and the associated benefits of the ThinkServer SA120.

4.1 File Server / Unstructured Data

Unstructured data is data that is not easily organized or categorized, such as documents, web pages and media files. Structured data, in contrast, typically refers to data in a database that is well organized and has defined fields that can be sorted and searched. Analysts have estimated that up to 80 percent of corporate data is unstructured.

In an IT infrastructure, the typical repository for centralized unstructured data is a file server. Because of the large amounts of unstructured data created daily, one of the basic uses of the ThinkServer SA120 is also very common: expanding file server storage.

Benefits of the SA120 in this application are:

- Ease of integration with existing ThinkServer systems
- Up to 384TB of external storage from one file server⁴

4.2 Backup to Disk

In addition to finding new ways for the initial storage of the massive amounts of data continuously created, strategies for backing up that data are changing to allow much faster backup (and restore). In some traditional IT environments, tape is still used for backup. However, in many cases, systems administrators have moved to using disk as a backup repository because larger datasets can no longer be backed up to slower tape drives in the time allotted for the backup.

⁴ With currently available 4TB drives.

Speed is even more important when considering recovery of backed-up data. In today's always-on commerce environment, businesses cannot afford the loss of revenue while waiting for IT systems to be brought back online with recovered data. Benefits of using the ThinkServer SA120 in backup-to-disk applications are similar to the basic file server expansion advantages — simple integration and massive scalability.

4.3 Database Acceleration

As previously described, a database contains structured data. Because of a database's ability to search and sort data, Database Management Systems (DBMS) have become what companies rely on to run their businesses. Many DBMS', such as Microsoft SQL or Oracle Relational DBMS, use transaction logs or redo logs to record every transaction made to the database. This log is a file (or files), used to keep an audit trail, so that in the event of a system failure or database corruption, the log can be played back with DBMS tools and the database can be recovered to a known good state.

As you might imagine, the log file is typically the bottleneck of a DBMS. This file is even more of a performance inhibitor when the DBMS is used for Online Transaction Processing (OLTP) applications. These applications may involve hundreds of transactions per second, with each transaction recorded in the log file.

The log files are typically much smaller than the database files themselves, so a reasonable and cost effective solution is to place the log file on fast SSD and leave the remaining data on spinning disk. Because the ThinkServer SA120 offers both SSDs and larger HDDs in the same easily managed enclosure, it is a perfect choice for a DBMS.

As described in the configuration section, further benefits are available when combining CacheCade with the SA120 and SSDs. If the database has other hot spots in addition to the log files, there is no need for time-consuming analysis or manual data placement. CacheCade and SSDs can be added to an existing LSI9286CV-8e RAID adapter and connected to the SA120 with no impact to existing data.

4.4 Building Block for Distributed Storage

Distributed storage is the opposite of centralized storage. Centralized storage is the traditional way multiple users or servers have accessed persistent data. The term SAN (Storage Area Network) is sometimes used for centralized storage in an enterprise IT environment. The SAN storage device is typically a large monolithic array that requires a significant upfront expense, and is much more difficult and costly to scale when the storage frame has been completely filled with drives.

Many distributed storage architectures work best with repeatable server and storage building blocks. One of the major benefits of these architectures is that they are easy to scale in a well-defined way, and allow IT shops to start small and pay as they grow, reducing the need for a large upfront investment.

Software Defined Storage (SDS) and storage virtualization are other trends to consider when contemplating distributed storage. The basic idea behind SDS is that the storage infrastructure is separated into two components: software to control and manage, plus simple cost-effective hardware for where the data resides. SDS and storage virtualization bring similar benefits to data storage as those we have seen in the virtualization of compute infrastructure:

- Easier to scale
- Increased availability
- Faster deployment and time to value
- Simplified hardware refreshes / repurposing of hardware
- Aggregation (or federation) of disparate resources
- Reduced vendor lock-in

Here are examples of applications using distributed storage where the unique features of the SA120 are beneficial.

4.4.1 Microsoft Windows Server 2012 Storage Spaces

Microsoft describes Storage Spaces as a new virtualization capability in Windows 8 and Windows Server 2012, which enables users to dramatically reduce the cost of highly available storage for virtualized or physical deployments, while also providing high resiliency and operational simplicity. Storage Spaces fits well into the general definition of SDS in that it provides a robust control plane for the SA120 data plane.

Some of the features of Storage Spaces are:

- Resiliency with RAID
- Automatic tiering between disk with different performance levels
- Thin provisioning
- Failover clustering support

4.4.2 Microsoft Exchange Server for Email

Microsoft introduced the Database Availability Group (DAG) feature in Microsoft Exchange Server 2010, and added improvements in Exchange Server 2013. This feature lets the administrator create an active mailbox database volume on one server, and passive mailbox databases on one or more additional servers. Exchange Server software then replicates the database from server to server. The DAG ensures high availability of the mailbox database without resorting to complicated and expensive shared storage (also known as SAN).

4.4.3 Data Analytics / Hadoop

Data analytics is the process of examining data with the purpose of drawing conclusions about the information, and is used by businesses to make better decisions on customers or products. “Big data” is a subset of data analytics and the term is most often invoked when the size of the

data set to be examined is beyond the typical amount of data the organization has dealt with in the past. Because the data is sometimes as large as multiple petabytes (thousands of terabytes), conventional centralized shared storage architectures are too costly to implement and scale.

There are other reasons that practitioners of big data analytics prefer server DAS over shared storage. To reduce the large amount of time required to process massive data sets, applications have evolved to allow a cost-effective single compute and storage node to process a smaller chunk of data, and then provide the ability to reassemble the analyzed chunks into a recognizable response. In addition, having the storage directly connected to the server allows consistent high performance as opposed to the sometimes unpredictable behavior of a SAN with many servers accessing the same central storage.

There are many applications available for big data. One of the most popular is the open-source framework called Hadoop. Prominent users of Hadoop include Yahoo (where it was originally developed) and Facebook. Some have estimated Hadoop is used by over 50 percent of Fortune 500 companies today. Lenovo collaborates with many providers of supported Hadoop distributions, including Intel and Cloudera.

5.0 Conclusion

The ThinkServer SA120 DAS array addresses a wide range of workloads, including fileserver, backup and database acceleration. Because of its low complexity and cost, it fits well into distributed storage architectures, yet can still provide multiple levels of performance for changing application needs. The SA120 integrates seamlessly into existing ThinkServer environments, and uses familiar ThinkServer tools for configuration and monitoring.

6.0 References

SA120 Datasheet:

http://www.lenovo.com/images/products/server/pdfs/datasheets/thinkserver_sa120_ds_ww.pdf

User Guide and Hardware Maintenance Manual (English):

http://download.lenovo.com/ibmdl/pub/pc/pccbbs/thinkservers/sa120_hmm_ug_en.pdf

MegaRAID SAS Software User Guide:

http://download.lenovo.com/ibmdl/pub/pc/pccbbs/thinkservers/megaraid_swug_en.pdf

9207-8e User Guide:

http://www.lsi.com/downloads/Public/Host%20Bus%20Adapters/LSI_UG_SAS9207-8e.pdf

9280-8e and 9286CV-8e User Guide:

http://www.lsi.com/downloads/Public/MegaRAID%20SAS/41450-04_RevA.pdf

Configuring ThinkServer RAID 700 Adapters (for 9280-8e):

http://www.lenovo.com/images/products/server/pdfs/tech_resources/thinkserver_configuring_raid_500_700_tr.pdf

Configuring ThinkServer RAID 710 Adapters (for 9286CV-8e):

http://www.lenovo.com/images/products/server/pdfs/tech_resources/thinkserver_configuring_raid_710_tr.pdf

ThinkServer Supported Operating Systems:

http://www.lenovo.com/images/products/server/pdfs/datasheets/thinkserver_options_os_virtualization_ds.pdf