



Lenovo ThinkAgile CP Hardware Part Replacement and Component Maintenance Procedures



Models: CP4000, CP6000

Note

Before using this information and the product it supports, be sure to read and understand the safety information and the safety instructions, which are available at the following address:

http://thinksystem.lenovofiles.com/help/topic/safety_documentation/pdf_files.html

In addition, be sure that you are familiar with the terms and conditions of the Lenovo warranty for your solution, which can be found at the following address:

<http://datacentersupport.lenovo.com/warrantylookup>

Second Edition (March 2020)

© Copyright Lenovo 2018, 2020.

LIMITED AND RESTRICTED RIGHTS NOTICE: If data or software is delivered pursuant to a General Services Administration "GSA" contract, use, reproduction, or disclosure is subject to restrictions set forth in Contract No. GS-35F-05925.

Contents

Contents i

Figuresiii

Tables v

Chapter 1. ThinkAgile CP Overview. . . 1

ThinkAgile CP4000 series overview 4
ThinkAgile CP6000 series overview 7
ThinkAgile CP solution components 10
 Storage block parts list 10
 Compute block parts list 12
 Power cords 14

Chapter 2. Roles and responsibilities. 15

Chapter 3. Installation Guidelines. . . 19

System reliability guidelines 20
Working inside the solution with the power on . . . 20
Handling static-sensitive devices 21

Chapter 4. Managing Support Mode 23

Enable Support Mode in the Cloud Controller . . . 23
Access the customer's system via Support Mode 24
Place a support file on the storage or compute node 26
Disable Support Mode in the Cloud Controller . . . 28

Chapter 5. Part replacement and maintenance procedures 29

Storage block maintenance procedures 29
 Storage block component indicators. 29
 Shutting off a storage controller. 34
 Trigger manual storage HA. 34
 Shutting down the storage block 35
Replacing storage block hardware 36
 Replacing storage network cables or transceivers 36
 Hot-swap drive replacement 36
 Hot-swap power supply replacement 62
 Hot-swap storage controller replacement . . . 70

Enclosure replacement 78
 Complete the parts replacement 84
Compute block maintenance procedures 85
 Shutting off a compute node 85
 Migrating VMs from a compute node. 85
 Migrate a VM between compute nodes 86
Replacing compute block hardware 87
 Hardware replacement procedures 87
 Replace compute network cables 87
 Enable or Disable Maintenance Mode 88
 Hot-swap drive replacement 91
 Hot-swap power supply replacement 94
 Fan replacement 99
 EIOM replacement 107
 DIMM replacement. 110
 Compute node replacement 118
 Replace the SD530 system board 120
 Replace the compute enclosure. 120
Interconnect switch maintenance procedures . . . 121
 Power off the interconnect switch in an SS-V3 environment 121
 Make the secondary interconnect switch the primary interconnect 121
 Recovering from DNS misconfiguration. . . . 124
Interconnect parts replacement procedures 126
 Replace the interconnect switch 126
 Cabling a single-interconnect configuration 134
 Cabling a dual-interconnect configuration. . . 138
 Replacing a cable between two interconnect switches 142
 Hardware management network cable replacement 142
Management switch maintenance procedures (optional) 143
 Replacing the management switch (optional). 143
 Hardware management network cable replacement 149

Chapter 6. Licensing considerations151

Figures

1.	Storage block	2	42.	Drive installation	93
2.	Compute block	2	43.	Hot-swap power supply removal	96
3.	CP Interconnect	3	44.	Hot-swap power supply installation	99
4.	Management switch	3	45.	Fan fault LEDs	101
5.	Storage block components	11	46.	60x60x56mm fan removal	102
6.	Compute block components	13	47.	80x80x80mm fan removal	103
7.	Cloud Controller Support Mode menu item	23	48.	60x60x56mm fan installation	105
8.	Turn On Support Mode dialog box	24	49.	80x80x80mm fan installation	106
9.	Port number and password for Support Mode connection	24	50.	EIOM ports and corresponding node number	107
10.	Virtual Office sign-in page	24	51.	EIOM removal	108
11.	Command to connect to customer's interconnect	26	52.	EIOM installation	109
12.	Cloud Controller Support Mode menu item	28	53.	DIMM removal	111
13.	Turn Off Support Mode dialog box	28	54.	Location of the DIMM connectors on the system board	113
14.	Enclosure indicators	30	55.	Location of the DIMM connectors on the system board	115
15.	Drive indicators	31	56.	DIMM installation	116
16.	Storage controller indicators	32	57.	Recabling management switch from primary to secondary interconnect	123
17.	Storage controller port indicators	33	58.	Removing the interconnect from the rack	127
18.	Hot-swap drive removal	52	59.	Hardware management network connectivity in a single interconnect configuration	134
19.	Hot-swap drive removal	53	60.	Single—interconnect configuration — compute node and network interconnect connections	136
20.	Opening the drive tray handle	56	61.	Single—interconnect configuration — storage block and network interconnect connections	137
21.	Sliding the drive into the drive bay	57	62.	Hardware management network connectivity in a dual—interconnect configuration	138
22.	Hot-swap power supply removal (1/2)	64	63.	Dual—interconnect configuration — compute node and network interconnect connections	140
23.	Hot-swap power supply removal (2/2)	65	64.	Dual—interconnect configuration — storage block and network interconnect connections	141
24.	Hot-swap power supply installation (1/ 2)	68	65.	Dual—interconnect configuration — inter- interconnect link network connectivity	141
25.	Hot-swap power supply installation (2/ 2)	69	66.	Power-cord retention clip removal	144
26.	Hot-swap storage controller removal	70	67.	Releasing switch unit from the rack	145
27.	Hot-swap storage controller installation	71	68.	Removing mounting brackets from switch	145
28.	SSH sign-in to proxy server	75	69.	Attaching mounting brackets to switch	147
29.	Entering the proxy server port	76	70.	Securing switch to rack posts	147
30.	Entering the proxy server password	76	71.	Installing power cord retention clip	148
31.	Side cover removal	78			
32.	Screws removal	79			
33.	Enclosure removal from rack	79			
34.	Removal of side cover	81			
35.	Installing the enclosure into the rack	82			
36.	Installing screws on front of enclosure	82			
37.	Reinstalling the side cover	83			
38.	Enable Maintenance Mode for a Node	88			
39.	Maintenance Mode dialog box	89			
40.	Disable Maintenance Mode for a Node	90			
41.	Drive removal	92			

Tables

1.	Storage block parts listing	10	6.	DIMM installation sequence (mirror mode/lockstep mode)	114
2.	Compute block part listing	12	7.	DIMM installation sequence (mirror mode/lockstep mode for 3, 6, 9 and 12 identical DIMMs)	114
3.	Fan fault LEDs	101	8.	DIMM installation sequence (sparing mode)	114
4.	EIOM port corresponding node number	107			
5.	DIMM installation sequence (Independent mode/normal mode).	114			

Chapter 1. ThinkAgile CP Overview

This product overview provides technical information on the Lenovo ThinkAgile CP Series cloud platform, its key features and specifications, components and options, and configuration guidelines.

Lenovo ThinkAgile CP Series is an all-in-one, composable cloud with an integrated application marketplace (Lenovo Cloud Marketplace) and end-to-end automation, delivering a turn-key cloud experience in your own data center. The ThinkAgile CP Series uses modular compute, storage, and networking components paired with the cloud virtualization software to create pools of IT resources, independently scaling and allocating capacity, and automatically configuring resources to fulfill application requirements.

Due to its software-defined modular architecture, the ThinkAgile CP Series platform can be scaled easily by adding more compute and storage resources independently of each other as your needs grow. Suggested workloads for the ThinkAgile CP Series include web services, virtual desktop infrastructure (VDI), enterprise applications, OLTP and OLAP databases, data analytics, application development, cost-optimized virtualization, containers, and other back-office applications.

Key Features

Lenovo ThinkAgile CP Series offers the following key features:

- Modular and scalable ThinkAgile CP Series configurations of an on-premises cloud platform featuring the second generation of the Intel Xeon Processor Scalable Family that are designed to optimize your workload's performance and provide the IT agility for your business demands.
- Factory-preloaded ready-to-go cloud platform that is delivered with all the infrastructure you need for your workloads: Modular and scalable physical and virtual compute, network, and storage resources, and the cloud software.
- Lenovo deployment services that are included with the platform help get customers up and running quickly.
- The cloud software provides a scalable software-defined infrastructure (SDI) that simplifies cloud deployments with integrated compute, storage, networking, security, and management services that manage application infrastructure and automate and orchestrate workload provisioning.
- Extensive security features, such as data at rest encryption, virtualized network and VM-level firewalls, and two-factor authentication help customers meet the most stringent security requirements.
- Centralized cloud-based management automates discovery, deployment, and configuration of cloud resources, and provides automated, non-disruptive software updates.
- Lenovo ThinkAgile Advantage provides a single point of contact for all support issues and integrates support chat and virtual technical assistance into your management interface.

Components

The ThinkAgile CP Series consists of the following components:

- **Storage Blocks**

A Storage Block is a 2U storage enclosure with up to 24 PCIe NVMe SSDs and two controllers for high availability and redundancy. A storage block runs the cloud software that delivers data storage layer to the cloud. It simplifies storage management, provides enterprise-class storage functionality, and enables linear performance and capacity scaling.



Figure 1. Storage block

- **Compute Blocks**

A Compute Block is a modular 2U enclosure that contains up to four nodes, and it delivers processor and memory resources to the cloud. The compute nodes run the cloud hypervisor that combines open KVM-based virtualization software, hardware integration, and automation to orchestrate and deliver an end-to-end compute platform.




Figure 2. Compute block

- **CP Interconnect**

The Interconnect centralizes connectivity of your on-premises infrastructure to the Cloud Controller and acts as the entry point into your existing network, which connects the Cloud Controller to your environment. The CP Interconnect consists of one or two high-density, ultra-low-latency 10/40 GbE network devices and use specialized cloud software which provides the fabric for your composable environment, automating discovery, onboarding, and device management. Integrated network virtualization secures virtual data center and applications.




Figure 3. CP Interconnect

- **Management Switch**

The 1 GbE management switch is required to provide out-of-band (OOB) management for the Interconnect, Compute Blocks, and Storage Blocks. The switch can be supplied by Lenovo (ThinkSystem NE0152T RackSwitch), or customers may provide their own switch.



Figure 4. Management switch

- **Cloud Controller**

Cloud Controller orchestrates and manages the on-premises infrastructure and workloads; however, it resides in the cloud. It automates and orchestrates infrastructure provisioning in real time to achieve complete and secure vertical hardware, software and workload integration. It federates and abstracts all physical hardware into a private cloud service.

Cloud Controller provides a single point of management across an unlimited number of ThinkAgile CP stacks with role-based access control, two-factor authentication, and secure HTML and RESTful API interfaces.

ThinkAgile CP4000 series overview

The Lenovo ThinkAgile CP4000 series scales as an initial data-management footprint for managed service provider and cloud service provider environments that start on a smaller scale but are poised for growth. The CP4000 series also supports business continuity when deployed remotely as an off-site backup for disaster recovery.

The Lenovo ThinkAgile CP4000 series consists of the following components:


- “Storage block” on page 4
- “Compute block” on page 5
- “Interconnect switch” on page 5
- “Management switch” on page 5
- “Services” on page 6
- “Software” on page 6

Storage block




- 1 Storage Block (SB-S10)
- Dual controller - single CPU per controller
- 8x SFP Base-T SFP+ transceivers
- 8x CAT6 cables, 1.5m and other lengths permitted
- 2x CAT5e cables, 1.5m and other lengths permitted
- 32 GB RAM per controller
- 8, 16, 24 drives (8 drives per drive pack)
- 800 GB, 1.6 TB NVMe drives per controller
- Drives must be consistent within a storage block
- Dimensions: 87.6 x 446.4 x 857.2 mm (3.4 x 17.6 x 33.7 inches)
- Maximum weight: 42.0 kg (92.6 lb)

Compute block



- 1 D2 enclosure, fully populated PSU, fan, 220V with 10Gb SFP+ network module
- Minimum 2 SD530 nodes, maximum 4 nodes
- 128 GB, 256 GB RAM per node
- 8+ cores per CPU, 1-2 sockets filled AXQT, and others
- 2x SFP+ DAC cables per node, 1.5m and other lengths permitted
- Boot: 1x SATA 240 GB
- Dimensions: 41.0 x 562.0 x 222.0 mm (1.7 x 22.2 x 8.8 inches)
- Maximum weight: 55 kg (121.2 lb)

Interconnect switch



- 1 or 2 CP-I-10 interconnect switches
- 48x 10 Gb SFP+ fabric ports
- 6x 40 Gb QSFP+ fabric ports
- Dimensions: 438.4 x 473 x 43.4 mm (17.26 x 18.62 x 1.71 inches)
- Maximum weight: 8.95 kg (19.73 lb), with two installed PSUs

Management switch

Lenovo ThinkSystem NE0152T RackSwitch



Note: A dedicated out-of-band (OOB) management switch is required for this solution. If the Lenovo NE0152T is not used, you must provide an OOB management switch with the correct configuration.

Services

Lenovo Deployment Services are mandatory

Lenovo Hardware Installation Service is optional

Software

Software is preloaded on compute, storage, and switch.

ThinkAgile CP6000 series overview

The Lenovo ThinkAgile CP6000 series offering can be scaled up to 40 compute nodes to support mainstream enterprise workloads and mixed-workload environments.

The Lenovo ThinkAgile CP6000 series consists of the following components:


- “Storage block” on page 7
- “Compute block” on page 8
- “Interconnect switch” on page 8
- “Management switch” on page 8
- “Services” on page 9
- “Software” on page 9

Storage block




- 1 Storage Block - a maximum 5 blocks (SB-D20)
- Dual controller - 2CPU per controller
- 8x SFP Base-T SFP+ adapters
- 8x CAT6 cables per storage block, 1.5m and other lengths permitted
- 64 GB RAM per controller
- 8, 16, 24 drives (8 drives per drive pack)
- 1.6 TB, 3.2 TB, 6.4 TB NVMe drives per controller
- Drives must be consistent within a controller
- Dimensions: 87.6 x 446.4 x 857.2 mm (3.4 x 17.6 x 33.7 inches)
- Maximum weight: 42.0 kg (92.6 lb)

Compute block



- 1 D2 enclosure, fully populated PSU, fan, 220V with 10Gb SFP+ network module — maximum 10 enclosure
- Minimum 2 SD530 nodes, default 4, maximum 40 nodes
- 128 GB, 256 GB, 384 GB, 512 GB, 768 GB, 1024 GB RAM
- 8+ cores per CPU, 1-2 sockets filled AXQT, and others
- 2x SFP+ DAC cables per node, 1.5m and other lengths permitted
- Boot: 1x SATA 240 GB
- Dimensions: 41.0 x 562.0 x 222.0 mm (1.7 x 22.2 x 8.8 inches)
- Maximum weight: 55 kg (121.2 lbs)

Interconnect switch



- 2 CP-I-10 Interconnect switches
- 48x 10 Gb SFP+ fabric ports
- 6x 40 Gb QSFP+ fabric ports
- Dimensions: 438.4 x 473 x 43.4 mm (17.26 x 18.62 x 1.71 inches)
- Maximum weight: 8.95 kg (19.73 lb), with two installed PSUs

Management switch

Lenovo ThinkSystem NE0152T RackSwitch



Note: A dedicated out-of-band (OOB) management switch is required for this solution. If the Lenovo NE0152T is not used, you must provide an OOB management switch with the correct configuration.

Services

Lenovo Deployment Services are mandatory

Lenovo Hardware Installation Service is optional

Software

Software is preloaded on compute, storage, and switch.

ThinkAgile CP solution components

This section describes the storage and compute blocks components, and information about the available power cords.

Storage block parts list

This section describes the components of the ThinkAgile CP storage block.

The parts listed in the table below fall into these categories:

- **Tier 1 customer replaceable unit (CRU):** Replacing Tier 1 CRUs is your responsibility. If Lenovo installs a Tier 1 CRU at your request with no service agreement, you will be charged for the installation.
- **Tier 2 customer replaceable unit:** You may install a Tier 2 CRU yourself or request Lenovo to install it, at no additional charge, under the warranty service type designated for your storage.
- **Field replaceable unit (FRU):** FRUs must be installed only by trained service technicians.
- **Consumable and structural parts:** Purchase and replacement of consumable and structural parts (components such as a cover or bezel) is your responsibility. If Lenovo acquires or installs a structural component at your request, you will be charged for the service.

Table 1. Storage block parts listing

Index	Description	Tier 1 CRU	Tier 2 CRU	FRU	Consumable and structural parts
1	2 Storage controllers		✓		
2	Enclosure		✓		
3	Power supplies	✓			
4	Rail kit				✓
5	2.5-inch drive fillers				✓
6	2.5-inch hot swap drives	✓			




Figure 5. Storage block components

For more information about ordering the parts shown here, see the following website:

[Lenovo Data Center Support page \(7Y66\)](#)

Compute block parts list

This section describes the components of the ThinkAgile CP compute block.

The parts listed in the table below fall into these categories:

- **Tier 1 customer replaceable unit (CRU):** Replacing Tier 1 CRUs is your responsibility. If Lenovo installs a Tier 1 CRU at your request with no service agreement, you will be charged for the installation.
- **Tier 2 customer replaceable unit:** You may install a Tier 2 CRU yourself or request Lenovo to install it, at no additional charge, under the warranty service type designated for your compute.
- **Field replaceable unit (FRU):** FRUs must be installed only by trained service technicians.
- **Consumable and structural parts:** Purchase and replacement of consumable and structural parts (components such as a cover or bezel) is your responsibility. If Lenovo acquires or installs a structural component at your request, you will be charged for the service.

Table 2. Compute block part listing

Index	Description	Tier 1 CRU	Tier 2 CRU	FRU	Consumable and Structural part
1	PCIe adapter	✓			
2	Air baffle				✓
3	Processor and heat sink assembly (85mm heat sink)			✓	
4	Processor and heat sink assembly (108mm heat sink)			✓	
5	Processor and heat sink assembly (108mm heat sink)			✓	
6	Processor and heat sink assembly (T-shaped heat sink)			✓	
7	Trusted Cryptographic Module			✓	
8	M.2 backplane	✓			
9	DIMM	✓			
10	2.5-inch drive bay blank (for empty bays next to the backplane)				✓
11	2.5-inch drive bay blank panel (for drive bays on the backplane)	✓			
12	2.5-inch hot-swap drive	✓			
13	Compute node tray			✓	
14	KVM breakout module	✓			
15	2.5-inch 6-drive hot-swap SAS/SATA backplane			✓	
16	2.5-inch 6-drive hot-swap SAS/SATA/NVMe backplane			✓	
17	2.5-inch 4-drive hot-swap SAS/SATA backplane			✓	
18	2.5-inch 4-drive hot-swap NVMe backplane			✓	
19	Compute block cover	✓			




Figure 6. Compute block components

For more information about ordering the parts shown here, see the following website:

[Lenovo Data Center Support page \(7X20\)](#)

Power cords

Information about the power cords available for the ThinkAgile CP storage and compute blocks.

Several power cords are available, depending on the country and region where ThinkAgile CP is installed.

To view the power cords that are available for a component:

- Interconnect switch
 1. Go to: [Lenovo Data Center Support page \(7Y67\)](#)
 2. Click **Parts**
 3. Click **Select Commodity** and choose Line Cords
- Compute block
 1. Go to: [Lenovo Data Center Support page \(7X20\)](#)
 2. Click **Parts**
 3. Click **Select Commodity** and choose Line Cords
- Storage block
 1. Go to: [Lenovo Data Center Support page \(7Y66\)](#)
 2. Click **Parts**
 3. Click **Select Commodity** and choose Line Cords

Notes:

- For your safety, a power cord with a grounded attachment plug is provided to use with this product. To avoid electrical shock, always use the power cord and plug with a properly grounded outlet.
- Power cords for this product that are used in the United States and Canada are listed by Underwriter's Laboratories (UL) and certified by the Canadian Standards Association (CSA).
- For units intended to operate at 115 volts: Use a UL-listed and CSA-certified cord set consisting of a minimum 18 AWG, type SVT or SJT, three-conductor cord, a maximum of 15 feet (4.6 meters) in length and a parallel blade, grounding-type attachment plug rated 15 amperes, 125 volts.
- For units intended to operate at 230 volts (U.S. use): Use a UL-listed and CSA-certified cord set consisting of a minimum 18 AWG, Type SVT or SJT, three-conductor cord, a maximum of 15 feet (4.6 meters) in length and a tandem blade, grounding-type attachment plug rated 15 amperes, 250 volts.
- For units intended to operate at 230 volts (outside the U.S.): Use a cord set with a grounding-type attachment plug. The cord set should have the appropriate safety approvals for the country in which you install the equipment.
- Power cords for a specific country or region are usually available only in that country or region.

Chapter 2. Roles and responsibilities

This section covers information about the roles and responsibilities involved in the replacement of different ThinkAgile CP hardware components.

Each maintenance and replacement procedure in this guide is performed by one or a combination of the following actors:

- **Solution administrator:** May be customer, third party, or whoever has access to the solution for performing administrative procedures (Cloud Controller, infrastructure)
- **Customer or field service:** For physical replacement of CRUs/FRUs based on part type and entitlement level
- **Lenovo Support:** Works with solution administrator to perform remote support procedures

The following table illustrates which roles are responsible for each of the tasks in this guide. Some tasks can be completed by the customer alone, while other tasks require the participation of Lenovo Support.

Task	Actor		
	Solution administrator	Customer or field service	Lenovo Support
“Shutting off a compute node” on page 85	Shuts down compute node on Cloud Controller web portal		
“Migrating VMs from a compute node” on page 85	<ul style="list-style-type: none"> • Turns on Maintenance Mode • Uses Migrate tool OR <ul style="list-style-type: none"> • Enables Support Mode for ThinkAgile Advantage Support assistance 		Optional. If customer wants assistance: <ul style="list-style-type: none"> • Accesses node remotely via Support Mode • Turns on Maintenance Mode • Uses Migrate tool
“Migrate a VM between compute nodes” on page 86	<ul style="list-style-type: none"> • Adds and removes compute tags to destination compute node and VMs OR <ul style="list-style-type: none"> • (Optional) Enables Support Mode for ThinkAgile Advantage Support assistance 		Optional. If customer wants assistance: <ul style="list-style-type: none"> • Accesses node remotely via Support Mode • Adds and removes compute tags to destination compute node and VMs
“Replace compute network cables” on page 87		Physically replaces compute enclosure-to-interconnect cable	
“Remove a compute node” on page 118	<ul style="list-style-type: none"> • Shuts down and removes compute node • Enables Support Mode on interconnect 		<ul style="list-style-type: none"> • Signs into primary interconnect via Support Mode • Runs commands and gathers info • Requests hardware removal

Task	Actor		
	Solution administrator	Customer or field service	Lenovo Support
“Install a compute node” on page 118		Physically installs compute node in enclosure and powers on	
“Replace the compute enclosure” on page 120	<ul style="list-style-type: none"> Records compute enclosure serial number “Adds a compute node” on page 118 “Removes a compute node” on page 118 Contacts ThinkAgile Advantage Support 		Contacts Lenovosupport to request hardware removal
“Hot-swap drive replacement” on page 91		Performs all steps in task	
“Hot-swap power supply replacement” on page 94		Performs all steps in task	
“Fan replacement” on page 99		Performs all steps in task	
“EIOM replacement” on page 107		Performs all steps in task	
“DIMM replacement” on page 110		Performs all steps in task	
“Replace the SD530 system board” on page 120	Performs all steps in task		
“Shutting off a storage controller” on page 34	<ul style="list-style-type: none"> Physically connects to storage controller Enters shutdown commands in hardware console OR <ul style="list-style-type: none"> (Optional) Enables Support Mode for ThinkAgile Advantage Support assistance 		Optional. If customer wants assistance: <ul style="list-style-type: none"> Accesses storage controller remotely via Support Mode Enters shutdown commands in hardware console
“Trigger manual storage HA” on page 34	<ul style="list-style-type: none"> Connects via SSH to active storage controller Runs command to stop storage controller services OR <ul style="list-style-type: none"> (Optional) Enables Support Mode for ThinkAgile Advantage Support assistance 		Optional. If customer wants assistance: <ul style="list-style-type: none"> Accesses storage controller remotely via Support Mode Runs command to stop storage controller services

Task	Actor		
	Solution administrator	Customer or field service	Lenovo Support
“Shutting down the storage block” on page 35	Powers off the storage block (all storage controllers) in the Cloud Controller		
“Replacing storage network cables or transceivers” on page 36		Physically replaces storage controller-to-interconnect cable or transceiver	
“Remove a hot-swap storage controller” on page 70	<ul style="list-style-type: none"> Obtains serial number of storage controller to be removed “Shuts off storage controller” on page 34 Physically removes storage controller from rack Enables Support Mode 		<ul style="list-style-type: none"> Signs into primary interconnect Runs command and gathers information Contacts LenovoSupport and requests hardware removal
“Install a hot-swap storage controller” on page 71	<ul style="list-style-type: none"> Physically installs and powers on storage controller Enables Support Mode 		<ul style="list-style-type: none"> “Updates the storage node VPD” on page 74 Signs into primary interconnect and runs command Registers storage blocks
“Update the storage node VPD” on page 74		Performs all steps in task	
“Hot-swap drive replacement” on page 36		Performs all steps in task	
“Hot-swap power supply replacement” on page 62		Performs all steps in task	
“Hot-swap drive replacement” on page 36		Performs all steps in task	(Optional) Performs all steps in task at customer request. Charges may apply depending on the customer’s warranty terms.
“Hot-swap storage controller replacement” on page 70	<ul style="list-style-type: none"> Obtains and records storage enclosure serial number “Shuts off storage controller” on page 34 “Removes compute node” on page 118 Records exact drive, storage controller, and cable locations Enables Support Mode and contacts ThinkAgile Advantage Support 	<ul style="list-style-type: none"> “Shuts off storage controller” on page 34 Powers off and disconnects storage enclosure Replaces enclosure, repopulates storage controllers and drives, reconnects cables, and powers on enclosure 	<ul style="list-style-type: none"> Signs in to primary interconnect and runs command Contacts Lenovosupport to request hardware removal

Task	Actor		
	Solution administrator	Customer or field service	Lenovo Support
“Power off the interconnect switch in an SS-V3 environment” on page 121	Performs all steps in task		
“Replacing a cable between two interconnect switches” on page 142		Performs all steps in task	
“Make the secondary interconnect switch the primary interconnect” on page 121			Performs all steps in task
“Replace the secondary interconnect in a DS-V3 configuration” on page 133		<ul style="list-style-type: none"> Physically replaces and cables interconnect Powers on replacement interconnect 	<ul style="list-style-type: none"> Requests interconnect settings from Lenovo Configures to-be-replaced and replacement interconnects
“Hardware management network cable replacement” on page 142		Physically replaces cables connecting the hardware management network to other components	
“Replacing the management switch (optional)” on page 143		Physically replaces management switch in the rack	
“Removing the management switch (optional)” on page 143		Performs all steps in task	

Chapter 3. Installation Guidelines

This section covers the installation guidelines that need to be followed when installing components in your ThinkAgile CP solution.

Before installing optional devices, read the following notices carefully:

Attention: Exposure to static electricity may lead to system halt and loss of data. Prevent this problem by keeping static-sensitive components in their static-protective packages until installation, and handling these devices with an electrostatic-discharge wrist strap or other grounding system.

- Read the safety information and guidelines to ensure that you work safely.
 - A complete list of safety information for all products is available: [Safety Information](#)
 - The following guidelines are available as well: “[Handling static-sensitive devices](#)” on page 21 and “[Working inside the solution with the power on](#)” on page 20.
- Make sure the components that you are installing are supported by the solution.
- When you install a new solution, apply the latest firmware. This will ensure that any known issues are addressed, and that your solution is ready to work with optimal performance. Contact Lenovo Support for firmware updates.

Important: Some cluster solutions require specific code levels or coordinated code updates. If the component is part of a cluster solution, verify that the latest level of code is supported for the cluster solution before you update the code.

- It is good practice to make sure that the solution is working correctly before you install an optional component.
- Keep the working area clean, and place removed components on a flat and smooth surface that does not shake or tilt.
- Do not attempt to lift an object that may be too heavy for you. If you have to lift a heavy object, read the following precautions carefully:
 - Make sure that you can stand steadily without slipping.
 - Distribute the weight of the object equally between your feet.
 - Use a slow lifting force. Never move suddenly or twist when you lift a heavy object.
 - To avoid straining the muscles in your back, lift by standing or by pushing up with your leg muscles.
- Back up all important data before you make changes related to the disk drives.
- Have a small flat-blade screwdriver, a small Phillips screwdriver, and a T8 torx screwdriver available.
- To view the error LEDs on the system board and internal components, leave the power on.
- You do not have to turn off the solution to remove or install hot-swap power supplies, hot-swap fans, or hot-plug USB devices. However, you must turn off the solution before you perform any steps that involve removing or installing adapter cables, and you must disconnect the power source from the solution before you perform any steps that involve removing or installing a riser card.
- Blue on a component indicates touch points, where you can grip to remove a component from or install it in the solution, open or close a latch, and so on.
- The Red strip on the drives, adjacent to the release latch, indicates that the drive can be hot-swapped if the solution and operating system support hot-swap capability. This means that you can remove or install the drive while the solution is still running.

Note: See the system specific instructions for removing or installing a hot-swap drive for any additional procedures that you might need to perform before you remove or install the drive.

- After finishing working on the solution, make sure you reinstall all safety shields, guards, labels, and ground wires.

System reliability guidelines

Review the system reliability guidelines to ensure proper system cooling and reliability.

Ensure that the following requirements are met:

- When the server comes with redundant power, a power supply must be installed in each power-supply bay.
- Adequate space around the server must be spared to allow server cooling system to work properly. Leave approximately 50 mm (2.0 in.) of open space around the front and rear of the server. Do not place any object in front of the fans.
- For proper cooling and airflow, refit the server cover before you turn the power on. Do not operate the server for more than 30 minutes with the server cover removed, for it might damage server components.
- Cabling instructions that come with optional components must be followed.
- A failed fan must be replaced within 48 hours since malfunction.
- A removed hot-swap fan must be replaced within 30 seconds after removal.
- A removed hot-swap drive must be replaced within two minutes after removal.
- A removed hot-swap power supply must be replaced within two minutes after removal.
- Every air baffle that comes with the server must be installed when the server starts (some servers might come with more than one air baffle). Operating the server with a missing air baffle might damage the processor.
- All processor sockets must contain either a socket cover or a processor with heat sink.
- When more than one processor is installed, fan population rules for each server must be strictly followed.
- Do not operate the enclosure without the System Management Module (SMM) assembly installed. Operating the solution without the SMM assembly might cause the system to fail. Replace the SMM assembly as soon as possible after removal to ensure proper operation of the system.

Working inside the solution with the power on

You may need to keep the power on with the solution cover removed to look at system information on the display panel or to replace hot-swap components. Review these guidelines before doing this.

Attention: The solution may stop and loss of data may occur when internal solution components are exposed to static electricity. To avoid this problem, always use an electrostatic-discharge wrist strap or other grounding systems when working inside the solution with the power on.

- Avoid loose-fitting clothing, particularly around your forearms. Button or roll up long sleeves before working inside the solution.
- Prevent your necktie, scarf, badge rope, or long hair from dangling into the solution.
- Remove jewelry, such as bracelets, necklaces, rings, cuff links, and wristwatches.
- Remove items from your shirt pocket, such as pens and pencils, in case they fall into the solution as you lean over it.
- Avoid dropping any metallic objects, such as paper clips, hairpins, or screws into the solution.

Handling static-sensitive devices

To reduce the possibility of damage from electrostatic discharge, review these guidelines before you handle static-sensitive devices.

Attention: Exposure to static electricity may lead to system halt and loss of data. Prevent this problem by keeping static-sensitive components in their static-protective packages until installation, and handling these devices with an electrostatic-discharge wrist strap or other grounding system.

- Limit your movement to prevent building up static electricity around you.
- Take additional care when handling devices during cold weather, as heating reduces indoor humidity and increases static electricity.
- Always use an electrostatic-discharge wrist strap or other grounding system, especially when working inside the solution with the power on.
- While the device is still in its static-protective package, touch it to an unpainted metal surface on the outside of the solution for at least two seconds. This drains static electricity from the package and your body.
- Remove the device from the package and install it directly into the solution without putting it down. If you must put the device down, put it back into the static-protective package. Never place the device on the solution or on any metal surface.
- When handling a device, carefully hold it by the edges or the frame.
- Do not touch solder joints, pins, or exposed circuitry.
- Keep the device out of others' reach to prevent damage.

Chapter 4. Managing Support Mode

For a Lenovo support agent to access a customer's ThinkAgile CP environment and provide assistance with its configuration, the customer must enable Support Mode.

These topics provide information about the modifications required to ensure that Support Mode in the ThinkAgile CP interconnect switch can be enabled or disabled, as well as accessing the customer's environment and placing support files on the customer's compute and storage nodes. A filtering or firewall rule must be applied to prevent any SSH outbound connection on port 22. By default, SSH outbound is disabled or prevented.

Note: The commands described here must be run on all interconnect switches in your ThinkAgile CP solution.

Enable Support Mode in the Cloud Controller

Complete this procedure to enable Support Mode on the Cloud Controller.

Procedure performed by: Customer

Execution: Web Portal only

To enable Support Mode, follow these steps:

- Step 1. In the Cloud Controller, click the  **Support** button in the top right corner of the toolbar, and then click **Turn on Support Mode** from the drop-down menu.




Figure 7. Cloud Controller Support Mode menu item

Step 2. A dialog box appears on screen. Click the **Turn On Support Mode** button.

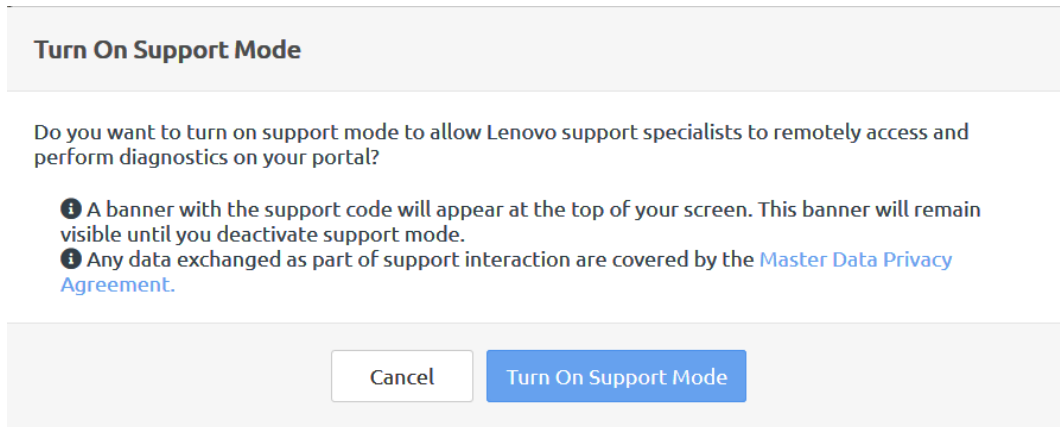


Figure 8. Turn On Support Mode dialog box

Step 3. After a few seconds, the interconnect switch port number and password information displays at the top of the Cloud Controller screen. The Lenovo ThinkAgile Advantage Support agent records the port number and password to use in the remote connection.

Access for ThinkAgile CP Network Interconnect IC 1 (Paired): 11781 n322rz32 ■ Port number
Access for ThinkAgile CP Network Interconnect IC 2 (Paired): 4195 bz4rhg2h ■ Password

Figure 9. Port number and password for Support Mode connection

Access the customer's system via Support Mode

Follow this procedure to gain access to the customer's system via Support Mode.

Step 1. SSH into the proxy server (tacp-proxy.cp.lenovo.com) using a valid Support ID.

For example:

Note: marc.lavergne is an *example* Support ID in the following:

```
~$ ssh marc.lavergne@tacp-proxy.cp.lenovo.com
marc.lavergne@tacp-proxy.cp.lenovo.com's password:
Last login: Thu Jan 2 09:14:16 2020 from 216.85.170.150
[marc.lavergne@ip-10-251-6-28 ~]$
```

Figure 10. Virtual Office sign-in page

Step 2. The SSH session CLI displays. Enter the following command:

```
ssh -p [user's port number] cldtx_support@localhost
```

where the user's port number is the number shown in the user's Cloud Controller Support Mode screen.

In this example, the user's switch port number is 11076.

```
[marc.lavergne@ip-10-251-6-28 ~]$ ssh -p 11076 cldtx_support@localhost -o StrictHostKeyChecking=no
cldtx_support@localhost's password:

This ThinkAgile CP Interconnect may include software made publicly available by Lenovo, including software licensed under
the General Public License and/or the Lesser General Public License (the open source software). Visit /usr/share/cloudi
stics/lenovo/legal directory to view additional information regarding licenses, acknowledgments and required copyright no
tices for the open source software used on this CP Interconnect.

The open source software is distributed in hope it will be useful, but WITHOUT ANY WARRANTY; without even the implied war
ranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See for example the GNU General Public License and/or the L
esser General Public License for more information
cldtx_support@R23-1-Primary-IC:~$
```

Figure 11. Command to connect to customer's interconnect

Step 3. Access a storage or compute node by first looking for the IP address of the node under the hardware section on the Cloud Controller or at `/var/lib/dhcp/dhcpd.leases` via the CLI. Then use the following command to connect to the node:

```
ssh -i ~/.ssh/<pri_file> cldtx_support@<ip>
```

where `<ip>` is the IP address of the node you are trying to connect to and `<pri_file>` is the private key file.

Step 4. If the configuration has dual switches, `-b` needs to be added to the node connect command, thus making it:

```
ssh -b 192.168.1.1 -i .ssh/<pri_file> cldtx_support@<ip> #primary
```

```
ssh -b 192.168.1.2 -i .ssh/<pri_file> cldtx_support@<ip> #secondary
```

Rather than looking up the Multi-enclosure Link Aggregation Group (MLAG) IP for the switch, you can run the following commands:

```
source /etc/tacp/tacp_switch_resource.sh
echo $MLAG_VLAN_IP
ssh -b $MLAG_VLAN_IP -i .ssh/<pri_file> cldtx_support@<ip>
```

Place a support file on the storage or compute node

This section covers the procedure to place a support file on a storage or compute node.

Follow these steps to place a support file on a storage or compute node:

Step 1. Establish an SSH session with the proxy server:

```
ssh firstname.lastname@tacp-proxy.cp.lenovo.com
```

where:

firstname.lastname is your Support user ID.

Step 2. On the proxy server, create a temporary directory and copy the support file there.

Step 3. Log out of the proxy server.

Step 4. Use SCP to copy the firmware packages (and LLDP64e) from your laptop to the directory on the proxy server:

```

~$ ssh marc.lavergne@tacp-proxy.cp.lenovo.com
marc.lavergne@tacp-proxy.cp.lenovo.com's password:
Last login: Thu Jan 2 09:14:16 2020 from 216.85.170.150
[marc.lavergne@ip-10-251-6-28 ~]$

```

- Step 5. Establish an SSH session with the proxy server again.
- Step 6. Make sure that the customer has enabled Support Mode.
- Step 7. Use scp to copy the files from the SSH proxy to the primary interconnect switch (in the example, the files are copied to a /tmp directory)

```

[firstname.lastname@ib-172-31-23-104 BR1.51]$ scp -P 34165 * cldtx_support@localhost:/tmp
cldtx_support@localhost's password:
intc-lnvgy_dd_nic_23.2-c_rhel7_x86-64.tgz          100% 7866KB   2.9MB/s   00:02
intc-lnvgy_fw_nic_6.80-4.00-sd530-e_linux_x86-64.bin 100% 6021KB   3.4MB/s   00:01
LLDP64e                                           100% 2964KB   2.9MB/s   00:01
lnvgy_fw_drvln_pdl218z-1.50_anyos_noarch.uxz      100% 150MB    4.2MB/s   00:36
lnvgy_fw_drvwn_pdl318w-1.50_anyos_noarch.uxz      100% 111MB    3.1MB/s   00:35
lnvgy_fw_lxpm_pdl118z-1.50_anyos_noarch.uxz      100% 60MB     4.0MB/s   00:15
lnvgy_fw_smm_tesm14f-1.10_anyos_noarch.zip       100% 19MB     4.0MB/s   00:04
lnvgy_fw_uefi_tee138c-2.20_anyos_32-64.uxz       100% 13MB     3.5MB/s   00:03
lnvgy_fw_xcc_tei344c-2.42_anyos_noarch.uxz       100% 124MB    4.4MB/s   00:28

```

Note: The port number used to connect to the interconnect switch is provided when the customer enables Support Mode.

- Step 8. Establish an SSH session with the primary interconnect switch:

```
ssh -p<port_number>cldtx_support@localhost -o StrictHostKeyChecking=no
```

- Step 9. Verify that the files that you copied are in the temporary directory.

```

[marc.lavergne@ip-10-251-6-28 BR-5.0.0]$
[marc.lavergne@ip-10-251-6-28 BR-5.0.0]$ scp -P 11076 * cldtx_support@localhost:/tmp
cldtx_support@localhost's password:
intc-lnvgy_dd_nic_23.2-c_rhel7_x86-64.tgz          100% 7866KB  715.0KB/s   00:11
intc-lnvgy_fw_nic_6.80-4.00-sd530-e_linux_x86-64.bin 100% 6021KB  752.5KB/s   00:08
LLDP64e                                           100% 2964KB  741.0KB/s   00:04
lnvgy_fw_drvln_pdl222d-1.70_anyos_noarch.uxz      100% 85MB    685.6KB/s   02:07
lnvgy_fw_lxpm_pdl122h-1.70_anyos_noarch.uxz      100% 60MB    675.5KB/s   01:31
lnvgy_fw_smm_tesm14h-1.12_anyos_noarch.zip       100% 19MB    688.1KB/s   00:28
lnvgy_fw_uefi_tee148m-2.41_anyos_32-64.uxz       100% 13MB    714.2KB/s   00:18
lnvgy_fw_xcc_tei392m-3.00_anyos_noarch.uxz       100% 125MB   681.7KB/s   03:07
[marc.lavergne@ip-10-251-6-28 BR-5.0.0]$
[marc.lavergne@ip-10-251-6-28 BR-5.0.0]$

```

Disable Support Mode in the Cloud Controller

Complete this procedure to disable Support Mode on the Cloud Controller.

Procedure performed by: Customer

Execution: Web Portal only

To disable Support Mode, follow these steps:

- Step 1. In the Cloud Controller, click the  **Support** button in the top right corner of the toolbar, and then click **Turn off Support Mode** from the drop-down menu.

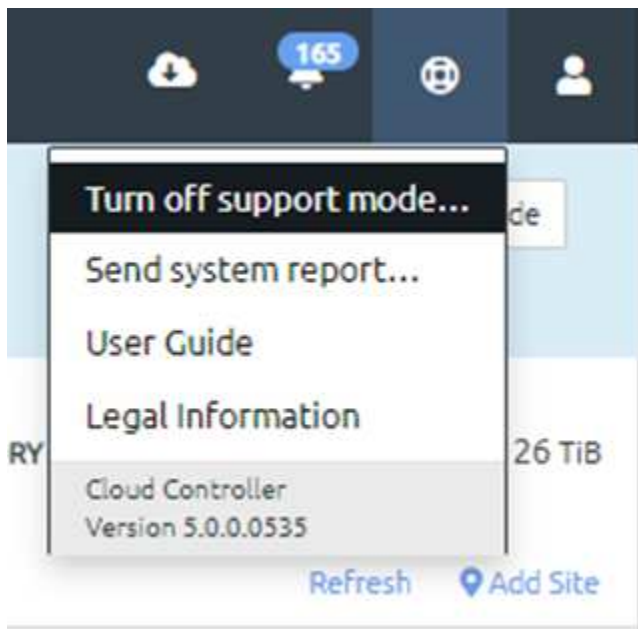


Figure 12. Cloud Controller Support Mode menu item

- Step 2. A dialog box appears on screen. Click the **Turn Off Support Mode** button.

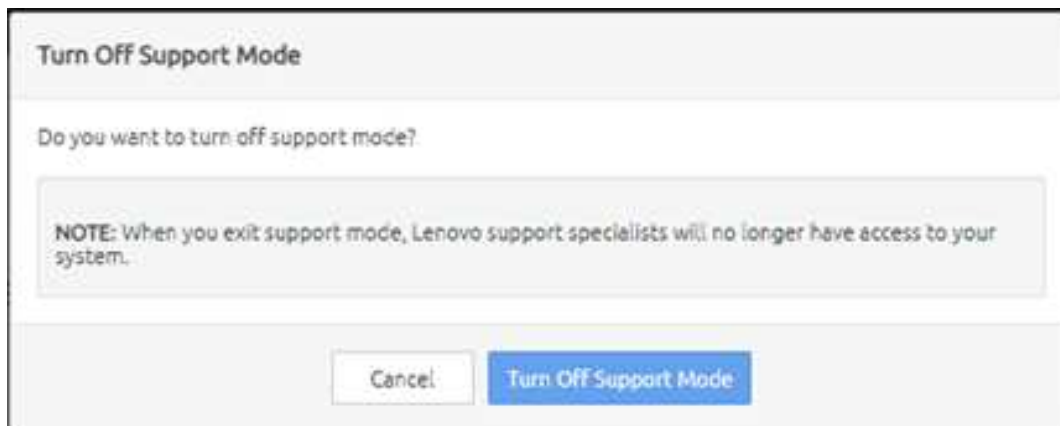


Figure 13. Turn Off Support Mode dialog box

Chapter 5. Part replacement and maintenance procedures

This section covers the procedures required to maintain and replace ThinkAgile CP hardware components.

Storage block maintenance procedures

This section covers the procedures required to perform storage block maintenance.

Refer to [Chapter 2 “Roles and responsibilities” on page 15](#) to determine who can perform these procedures.

Storage block component indicators

This section describes the ThinkAgile CP storage block component indicators.

Before performing maintenance procedures on your ThinkAgile CP storage controller, familiarize yourself with the following indicators to identify the storage block components' operating states.

Enclosure indicators

The LED indicators for the storage enclosure are located on the left side of the front of the storage block.

The enclosure has three LED indicators:

- Operation LED indicator:
 - On: the enclosure is receiving power
 - Off: the enclosure is not receiving power
- Identification LED indicator:
 - On: the enclosure is identified
 - Off: the enclosure is not identified
- Error/Warning LED indicator:
 - On: fault conditions are present
 - Off: the enclosure is operating under normal conditions




Figure 14. Enclosure indicators

Drive indicators

Storage drives are located in the front of the storage block.

Each drive has two LED indicators:

- Drive activity LED indicator:
 - On: drive activity is present
 - Off: no drive activity
- Drive error/warning LED indicator:
 - Constant On: fault conditions are present
 - 4 Hz blink rate (four times a second): locating the drive
 - 1 Hz blink rate (once a second): rebuilding the drive
 - Constant Off: the drive is operating under normal conditions

Each drive also has the following:

- drive tray label
- drive bay number written on the enclosure
- drive tray latch release

Drive fault LED indicator:

Constant On - drive fault

4 Hz blink rate - drive locate

1 Hz blink rate - drive rebuild

Constant Off - normal operation




Figure 15. Drive indicators

Storage controller indicators

The storage controller LED indicators are located on the rear of the storage block. Each block has two controllers. The top storage controller has its LED indicators on the top right side of the block, while the bottom controller has its LED indicators on the bottom left side of the block.

Each controller has three LED indicators:

- Operation LED indicator:
 - On: the controller is receiving power
 - Off: the controller is not receiving power
- Identification LED indicator:
 - Constant On: the controller is identified
 - 1 Hz blink rate (once a second): the controller is not identified
 - Constant Off: the controller is operating under normal conditions
- Error/Warning LED indicator:
 - On: fault conditions are present
 - Off: the controller is operating under normal conditions




Figure 16. Storage controller indicators

Storage controller port indicators

The Ethernet port LED indicators for the storage controllers are found on the rear of the storage block.

Each Ethernet port has two LED indicators:

- Activity LED indicator:
 - On: traffic is present on the port
 - Off: no traffic is present
- Error/Warning LED indicator:
 - On: fault conditions are present
 - Off: the port is operating under normal conditions

Activity LED indicator:

On - traffic is traversing the port

Off - no traffic




Figure 17. Storage controller port indicators

Shutting off a storage controller

This procedure describes how to shut off a storage controller.

Procedure performed by: Solution administrator, or ThinkAgile Advantage Support (optional, via [Chapter 4 “Support Mode” on page 23](#)).

Execution: On hardware

Storage controllers must be shut off in the storage controller itself, not in the Cloud Controller. To connect to the OS running on the storage controller, connect your laptop to the storage controller and use the Hardware Console available on the storage controller action menu in the Cloud Controller.

While connecting to the storage controller, enter the username `manager` and the password set during installation. If you have forgotten the password, reset it in the Cloud Controller.

Notes: Shutting off a passive storage controller is non-disruptive, but shutting down the active storage controller disconnects all exported drives for the VMs. If you are shutting down the active storage controller, either:

- fail over to the passive storage controller, or
- move the VMs to another storage node

After a successful sign-in, execute the following commands to shut down the storage controller at the node-level OS:

```
> sudo /usr/bin/tacp-storage-controller-safe-stop  
> sudo shutdown -h now
```

When you initiate shutoff of the entire storage block, the Cloud Controller lists the VMs that have disks exported to the storage block.

At this point, the shut-down storage controller can be pulled and serviced.

Trigger manual storage HA

This procedure describes how to trigger a storage HA, moving disk export control from the active storage controller to the passive storage controller.

Procedure performed by: Solution administrator, or ThinkAgile Advantage Support (optional, via [Chapter 4 “Support Mode” on page 23](#) to ThinkAgile Advantage).

Execution: On hardware only

To transfer disk export control from the active storage controller to the passive storage controller, you must simulate an active storage controller failure, so the system safely generates a storage HA event. To do this, you must simulate storage HA heartbeats between the Storage Controller service running on the active storage controller.

To trigger a manual storage HA event, follow these steps:

- Step 1. Connect to the active storage controller via SSH. On the Hardware page, the active storage controller is labeled.
- Step 2. In the active storage controller OS, run the following command to stop the Storage Controller Services:

```
> sudo systemctl stop tacp-storage-controller
```

The passive storage controller becomes active in the Cloud Controller.

Shutting down the storage block

Procedure to power off the ThinkAgile CP storage block.

Procedure performed by: Solution administrator

Execution: Web Portal only

You shut down a storage block (shutting down all storage controllers) directly in the Cloud Controller.

In the Hardware page, in the storage block Action menu , select **Power off...**

At that point, if there are VMs running and using disks exported by the storage block, the system requests that the VMs be shut down. This shutdown can be done in the dialog box.

For best results, VM owners should perform a proper shutdown within the VMs based on their application-specific shutdown procedures.

Replacing storage block hardware

Procedures for replacing the Lenovo ThinkAgile CP storage block hardware.

This section provides installation and removal procedures for all serviceable ThinkAgile CP system components. Each component replacement procedure references any tasks that need to be performed to gain access to the component being replaced.

Refer to [Chapter 2 “Roles and responsibilities” on page 15](#) to determine who can perform these procedures.

For more information about ordering parts, go to: <http://datacentersupport.lenovo.com/us/en/products/storages/thinksystem/sr630/7x01/parts>.

Replacing storage network cables or transceivers

This section describes how to replace storage network cables or transceivers.

Procedure performed by: Customer or field service

Replacing a cable or transceiver that connects the storage controller to an interconnect can be done live, without any specific preparation steps.

When a cable or transceiver is disconnected, the storage controller associated with the port will decrease in bandwidth by up to 25%. Disconnecting all four cables results in zero bandwidth to the controller, which may trigger failover to the other controller. If there is only one interconnect in the chassis, disconnecting all cables results in taking the storage offline as there is no failover option.

Hot-swap drive replacement

This section covers information about replacing a hot-swap drive.

Use this information to remove and install a hot-swap drive. You can remove or install a hot-swap drive without turning off the storage, which helps you avoid significant interruption to system operation.

Attention: Before replacing a hot-swap drive, read and follow the steps in the [“Hot-swap drive replacement preparations” on page 37](#).

Notes:

- The term “hot-swap drive” refers to all the supported hot-swap NVMe solid-state drives.
- Use any documentation that comes with the drive and follow those instructions in addition to the instructions in these topics. Ensure that you have all the cables and other equipment that come with the drive, as specified in its documentation.
- The electromagnetic interference (EMI) integrity and cooling of the storage are protected by having all drive bays covered or occupied. The vacant bays are either covered by an EMI-protective panel or occupied by drive fillers. When installing a drive, save the removed drive filler in case you later remove the drive and need the drive filler to cover the space.
- *SED* = Self-Encrypting Drive
- *NVME* = Non-Volatile Memory Express

Hot-swap drive replacement preparations

This section covers information about the preparations required before replacing a storage block hot-swap drive.

FAQs

Before you begin the drive replacement procedure, it is very important to read the following list of questions and answers.

Attention: Do not skip this section, even if you have read it earlier, as answers may change or more questions may be added to the list.

Do the drives need to be prepared before they can be added to the storage enclosure?

Yes. Refer to the [“Preparing the drives” on page 38](#) section.

Can I prepare the drives on the storage enclosure where I want to add the disks?

No.

How many drives I can replace at a time?

You can replace one drive at a time.

Can the drives be of different capacity than the drives already present?

No.

Should all drives have the same firmware level?

Yes.

Note: To check the drive firmware level, use the following command and check the values in the FW Rev column:

```
$ nvme list
```

How do I upgrade the firmware?

Contact Lenovo Support to verify the necessary firmware levels or information about updating firmware for the ThinkAgile CP solution.

Can drives from different vendors be mixed?

No.

Preparing the drives

This section covers the steps required for preparing new drives to be added to the storage enclosure.

Ensure that the size of the drive and its firmware level match with the drives that are already part of the target storage enclosure.

Note: To check the drive firmware level, use the following command and check the values in the FW Rev column:

```
$ nvme list
```

Node	SN	Model	Namespace	Usage	Format	FW Rev
-----	-----	-----	-----	-----	-----	-----

You must destroy the existing namespace on the drive and rescan using one of the following options:

Option 1

Use a new storage enclosure.

CAUTION:

Do not attempt this option on the enclosure where you are adding or removing disks. Use a new enclosure that is not currently in use.

If the disks are part of an existing ThinkAgile CP storage block array, follow these steps to destroy the namespace:

1. Execute the following command:

```
# /usr/share/tacp/encryption/ignite_build_sed -f -d -v
```

2. Rescan the PCI bus by executing the following commands:

```
modprobe -r nvme:echo 1 > /sys/bus/pci/rescan:modprobe nvme
```

Option 2

Use this option if you do not have access to another storage enclosure.

1. List all the namespaces on the drive

```
nvme list-ns /dev/nvme0
```

2. Delete the namespaces

For example, to delete the first namespace use the following command:

```
nvme delete-ns /dev/nvme0 -n 1
```

Attention: You must delete all namespaces on all the drives you are installing.

3. Verify that no namespaces exist on the drives:

```
$ nvme list
```

Example of a clean drive configuration

This section covers an example of a clean drive configuration.

To check the storage drive configuration, use the following command:

```
# lsblk
```

Following is an example output:

NAME	MAJ:MIN	RM	SIZE	RO	TYPE	MOUNTPOINT
sda	8:0	0	119.2G	0	disk	
├─sda1	8:1	0	200M	0	part	
├─sda2	8:2	0	1G	0	part	
├─sda3	8:3	0	2G	0	part	
├─sda4	8:4	0	110.1G	0	part	
└─md126	9:126	0	113.3G	0	raid1	
├─md126p1	259:0	0	200M	0	md	/boot/efi
├─md126p2	259:1	0	1G	0	md	/boot
├─md126p3	259:2	0	2G	0	md	[SWAP]
└─md126p4	259:3	0	110.1G	0	md	/
sdb	8:16	0	119.2G	0	disk	
├─sdb1	8:17	0	200M	0	part	
├─sdb2	8:18	0	1G	0	part	
├─sdb3	8:19	0	2G	0	part	
├─sdb4	8:20	0	110.1G	0	part	
└─md126	9:126	0	113.3G	0	raid1	
├─md126p1	259:0	0	200M	0	md	/boot/efi
├─md126p2	259:1	0	1G	0	md	/boot
├─md126p3	259:2	0	2G	0	md	[SWAP]
└─md126p4	259:3	0	110.1G	0	md	/
sr0	11:0	1	1024M	0	rom	
nvme0n1	259:4	0	740.4G	0	disk	
├─sed0	253:0	0	740.4G	0	crypt	
├─md125	9:125	0	4.3T	0	raid5	
├─md124	9:124	0	4.3T	0	raid0	
├─vdo_data	253:9	0	4.3T	0	dm	
├─vdo	253:10	0	43.4T	0	vdo	
├─CBS_POOL_data	253:8	0	43.3T	0	dm	
├─CBS_POOL	253:12	0	43.3T	0	dm	
├─SPECIAL_METADATA_STORE_4bd8d3ce75b1465082619c1f3fc15fb2	253:13	0	1G	0	dm	
├─SPECIAL_TEMPLATE_STORE_cd6d34c281054ef690a89030bef839ec	253:14	0	43.3T	0	dm	
├─CBS_POOL_meta	253:11	0	16G	0	dm	
├─CBS_POOL	253:12	0	43.3T	0	dm	
├─SPECIAL_METADATA_STORE_4bd8d3ce75b1465082619c1f3fc15fb2	253:13	0	1G	0	dm	
└─SPECIAL_TEMPLATE_STORE_cd6d34c281054ef690a89030bef839ec	253:14	0	43.3T	0	dm	
nvme1n1	259:10	0	740.4G	0	disk	
├─sed1	253:1	0	740.4G	0	crypt	
├─md125	9:125	0	4.3T	0	raid5	
├─md124	9:124	0	4.3T	0	raid0	
├─vdo_data	253:9	0	4.3T	0	dm	
├─vdo	253:10	0	43.4T	0	vdo	
├─CBS_POOL_data	253:8	0	43.3T	0	dm	
├─CBS_POOL	253:12	0	43.3T	0	dm	
├─SPECIAL_METADATA_STORE_4bd8d3ce75b1465082619c1f3fc15fb2	253:13	0	1G	0	dm	
├─SPECIAL_TEMPLATE_STORE_cd6d34c281054ef690a89030bef839ec	253:14	0	43.3T	0	dm	
├─CBS_POOL_meta	253:11	0	16G	0	dm	
├─CBS_POOL	253:12	0	43.3T	0	dm	
├─SPECIAL_METADATA_STORE_4bd8d3ce75b1465082619c1f3fc15fb2	253:13	0	1G	0	dm	
└─SPECIAL_TEMPLATE_STORE_cd6d34c281054ef690a89030bef839ec	253:14	0	43.3T	0	dm	

nvme2n1	259:9	0	740.4G	0	disk
└─sed2	253:2	0	740.4G	0	crypt
└─┬─md125	9:125	0	4.3T	0	raid5
└─┬─md124	9:124	0	4.3T	0	raid0
└─┬─vdo_data	253:9	0	4.3T	0	dm
└─┬─vdo	253:10	0	43.4T	0	vdo
└─┬─CBS_POOL_data	253:8	0	43.3T	0	dm
└─┬─CBS_POOL	253:12	0	43.3T	0	dm
└─┬─SPECIAL_METADATA_STORE_4bd8d3ce75b1465082619c1f3fc15fb2	253:13	0	1G	0	dm
└─┬─SPECIAL_TEMPLATE_STORE_cd6d34c281054ef690a89030bef839ec	253:14	0	43.3T	0	dm
└─┬─CBS_POOL_meta	253:11	0	16G	0	dm
└─┬─CBS_POOL	253:12	0	43.3T	0	dm
└─┬─SPECIAL_METADATA_STORE_4bd8d3ce75b1465082619c1f3fc15fb2	253:13	0	1G	0	dm
└─┬─SPECIAL_TEMPLATE_STORE_cd6d34c281054ef690a89030bef839ec	253:14	0	43.3T	0	dm
nvme3n1	259:6	0	740.4G	0	disk
└─sed3	253:3	0	740.4G	0	crypt
└─┬─md125	9:125	0	4.3T	0	raid5
└─┬─md124	9:124	0	4.3T	0	raid0
└─┬─vdo_data	253:9	0	4.3T	0	dm
└─┬─vdo	253:10	0	43.4T	0	vdo
└─┬─CBS_POOL_data	253:8	0	43.3T	0	dm
└─┬─CBS_POOL	253:12	0	43.3T	0	dm
└─┬─SPECIAL_METADATA_STORE_4bd8d3ce75b1465082619c1f3fc15fb2	253:13	0	1G	0	dm
└─┬─SPECIAL_TEMPLATE_STORE_cd6d34c281054ef690a89030bef839ec	253:14	0	43.3T	0	dm
└─┬─CBS_POOL_meta	253:11	0	16G	0	dm
└─┬─CBS_POOL	253:12	0	43.3T	0	dm
└─┬─SPECIAL_METADATA_STORE_4bd8d3ce75b1465082619c1f3fc15fb2	253:13	0	1G	0	dm
└─┬─SPECIAL_TEMPLATE_STORE_cd6d34c281054ef690a89030bef839ec	253:14	0	43.3T	0	dm
nvme4n1	259:8	0	740.4G	0	disk
└─sed4	253:4	0	740.4G	0	crypt
└─┬─md125	9:125	0	4.3T	0	raid5
└─┬─md124	9:124	0	4.3T	0	raid0
└─┬─vdo_data	253:9	0	4.3T	0	dm
└─┬─vdo	253:10	0	43.4T	0	vdo
└─┬─CBS_POOL_data	253:8	0	43.3T	0	dm
└─┬─CBS_POOL	253:12	0	43.3T	0	dm
└─┬─SPECIAL_METADATA_STORE_4bd8d3ce75b1465082619c1f3fc15fb2	253:13	0	1G	0	dm
└─┬─SPECIAL_TEMPLATE_STORE_cd6d34c281054ef690a89030bef839ec	253:14	0	43.3T	0	dm
└─┬─CBS_POOL_meta	253:11	0	16G	0	dm
└─┬─CBS_POOL	253:12	0	43.3T	0	dm
└─┬─SPECIAL_METADATA_STORE_4bd8d3ce75b1465082619c1f3fc15fb2	253:13	0	1G	0	dm
└─┬─SPECIAL_TEMPLATE_STORE_cd6d34c281054ef690a89030bef839ec	253:14	0	43.3T	0	dm
nvme5n1	259:5	0	740.4G	0	disk
└─sed5	253:5	0	740.4G	0	crypt
└─┬─md125	9:125	0	4.3T	0	raid5
└─┬─md124	9:124	0	4.3T	0	raid0
└─┬─vdo_data	253:9	0	4.3T	0	dm
└─┬─vdo	253:10	0	43.4T	0	vdo
└─┬─CBS_POOL_data	253:8	0	43.3T	0	dm
└─┬─CBS_POOL	253:12	0	43.3T	0	dm
└─┬─SPECIAL_METADATA_STORE_4bd8d3ce75b1465082619c1f3fc15fb2	253:13	0	1G	0	dm
└─┬─SPECIAL_TEMPLATE_STORE_cd6d34c281054ef690a89030bef839ec	253:14	0	43.3T	0	dm
└─┬─CBS_POOL_meta	253:11	0	16G	0	dm
└─┬─CBS_POOL	253:12	0	43.3T	0	dm
└─┬─SPECIAL_METADATA_STORE_4bd8d3ce75b1465082619c1f3fc15fb2	253:13	0	1G	0	dm
└─┬─SPECIAL_TEMPLATE_STORE_cd6d34c281054ef690a89030bef839ec	253:14	0	43.3T	0	dm

nvme6n1	259:11	0	740.4G	0	disk
└─sed6	253:6	0	740.4G	0	crypt
└─md125	9:125	0	4.3T	0	raid5
└─md124	9:124	0	4.3T	0	raid0
└─vdo_data	253:9	0	4.3T	0	dm
└─vdo	253:10	0	43.4T	0	vdo
└─CBS_POOL_data	253:8	0	43.3T	0	dm
└─CBS_POOL	253:12	0	43.3T	0	dm
└─SPECIAL_METADATA_STORE_4bd8d3ce75b1465082619c1f3fc15fb2	253:13	0	1G	0	dm
└─SPECIAL_TEMPLATE_STORE_cd6d34c281054ef690a89030bef839ec	253:14	0	43.3T	0	dm
└─CBS_POOL_meta	253:11	0	16G	0	dm
└─CBS_POOL	253:12	0	43.3T	0	dm
└─SPECIAL_METADATA_STORE_4bd8d3ce75b1465082619c1f3fc15fb2	253:13	0	1G	0	dm
└─SPECIAL_TEMPLATE_STORE_cd6d34c281054ef690a89030bef839ec	253:14	0	43.3T	0	dm
nvme7n1	259:7	0	740.4G	0	disk
└─sed7	253:7	0	740.4G	0	crypt
└─md125	9:125	0	4.3T	0	raid5
└─md124	9:124	0	4.3T	0	raid0
└─vdo_data	253:9	0	4.3T	0	dm
└─vdo	253:10	0	43.4T	0	vdo
└─CBS_POOL_data	253:8	0	43.3T	0	dm
└─CBS_POOL	253:12	0	43.3T	0	dm
└─SPECIAL_METADATA_STORE_4bd8d3ce75b1465082619c1f3fc15fb2	253:13	0	1G	0	dm
└─SPECIAL_TEMPLATE_STORE_cd6d34c281054ef690a89030bef839ec	253:14	0	43.3T	0	dm
└─CBS_POOL_meta	253:11	0	16G	0	dm
└─CBS_POOL	253:12	0	43.3T	0	dm
└─SPECIAL_METADATA_STORE_4bd8d3ce75b1465082619c1f3fc15fb2	253:13	0	1G	0	dm
└─SPECIAL_TEMPLATE_STORE_cd6d34c281054ef690a89030bef839ec	253:14	0	43.3T	0	dm

To check the status of the RAID array, use the following command:

```
# cat /proc/mdstat
```

Following is an example output:

```
Personalities : [raid1] [raid6] [raid5] [raid4] [raid0]
md124 : active raid0 md125[0]
4657247232 blocks super 1.2 512k chunks

md125 : active raid5 dm-7[7](S) dm-6[6] dm-5[5] dm-4[4] dm-3[3] dm-2[2] dm-1[1] dm-0[0]
4657379328 blocks super 1.2 level 5, 16k chunk, algorithm 2 [7/7] [UUUUUUU]
bitmap: 0/6 pages [0KB], 65536KB chunk

md126 : active raid1 sda[1] sdb[0]
118778880 blocks super external:/md127/0 [2/2] [UU]

md127 : inactive sda[1](S) sdb[0](S)
6306 blocks super external:imsm

unused devices: <none>
```

To view detailed information about the status of the RAID array, use the following command:

```
# mdadm --detail /dev/md/md5
```

Following is an example output:

```
/dev/md125:
Version : 1.2
Creation Time : Mon Aug 20 11:03:31 2018
Raid Level : raid5
Array Size : 4657379328 (4441.62 GiB 4769.16 GB)
Used Dev Size : 776229888 (740.27 GiB 794.86 GB)
Raid Devices : 7
Total Devices : 8
Persistence : Superblock is persistent

Intent Bitmap : Internal

Update Time : Mon Aug 20 11:04:34 2018
State : clean
Active Devices : 7
Working Devices : 8
Failed Devices : 0
Spare Devices : 1

Layout : left-symmetric
Chunk Size : 16K

Consistency Policy : bitmap
Name : any:md5
UUID : 5f7b873c:16d6d418:7b5f6fde:247566a7
Events : 2

Number Major Minor RaidDevice State
0      253    0      0      active sync /dev/dm-0
1      253    1      1      active sync /dev/dm-1
2      253    2      2      active sync /dev/dm-2
3      253    3      3      active sync /dev/dm-3
4      253    4      4      active sync /dev/dm-4
5      253    5      5      active sync /dev/dm-5
6      253    6      6      active sync /dev/dm-6
7      253    7      -      spare /dev/dm-7
```

To check the SED mappings, use the following command:

```
# ls -l /dev/mapper/sed*
```

Following is an example output:

```
lrwxrwxrwx. 1 root root 7 Aug 20 11:03 /dev/mapper/sed0 -> ../dm-0
lrwxrwxrwx. 1 root root 7 Aug 20 11:03 /dev/mapper/sed1 -> ../dm-1
lrwxrwxrwx. 1 root root 7 Aug 20 11:03 /dev/mapper/sed2 -> ../dm-2
lrwxrwxrwx. 1 root root 7 Aug 20 11:03 /dev/mapper/sed3 -> ../dm-3
lrwxrwxrwx. 1 root root 7 Aug 20 11:03 /dev/mapper/sed4 -> ../dm-4
lrwxrwxrwx. 1 root root 7 Aug 20 11:03 /dev/mapper/sed5 -> ../dm-5
lrwxrwxrwx. 1 root root 7 Aug 20 11:03 /dev/mapper/sed6 -> ../dm-6
lrwxrwxrwx. 1 root root 7 Aug 20 11:03 /dev/mapper/sed7 -> ../dm-7
```

Use the following command to list the NVME namespaces:

```
# ls /dev/nvme?
```

Following is an example output:

```
/dev/nvme0 /dev/nvme1 /dev/nvme2 /dev/nvme3 /dev/nvme4 /dev/nvme5 /dev/nvme6 /dev/nvme7
```

To check the drive firmware level, use the following command and check the values in the FW Rev column:

```
$ nvme list
```

Following is an example output:

Node	SN	...	Namespace Usage	Format	FW Rev
/dev/nvme0n1	S3HCNX0JCO2194	1	795.00 GB / 795.00 GB	512 B + 0 B	GPNA9B3Q
/dev/nvme1n1	S3HCNX0JCO2193	1	795.00 GB / 795.00 GB	512 B + 0 B	GPNA9B3Q
/dev/nvme2n1	S3HCNX0JCO2233	1	795.00 GB / 795.00 GB	512 B + 0 B	GPNA9B3Q
/dev/nvme3n1	S3HCNX0JCO2191	1	795.00 GB / 795.00 GB	512 B + 0 B	GPNA9B3Q
/dev/nvme4n1	S3HCNX0JCO2189	1	795.00 GB / 795.00 GB	512 B + 0 B	GPNA9B3Q
/dev/nvme5n1	S3HCNX0JCO2198	1	795.00 GB / 795.00 GB	512 B + 0 B	GPNA9B3Q
/dev/nvme6n1	S3HCNX0JCO2188	1	795.00 GB / 795.00 GB	512 B + 0 B	GPNA9B3Q
/dev/nvme7n1	S3HCNX0JCO2195	1	795.00 GB / 795.00 GB	512 B + 0 B	GPNA9B3Q

To display the current RAID configuration, use the following command:

```
# /usr/share/tacp/lenovo/tacp-nvme-control.py -display
```

Following is an example output:

```
discovering NVMe disks. this operation will take a few seconds.
Chassis UUID: 500e0eca08057b00
Number of canisters: 2
This is the bottom (primary) canister and has the controller id 33
The top (secondary) canister has the controller id 34
This chassis has the following controller ids: 33 , 34

NVMe devices: 8
-----

NVMe control device: /dev/nvme0
Slot: 8
Manufacturer: SAMSUNG MZWL800HEHP-00003
Serial number: S3HCNX0JCO2194
Firmware: GPNA9B3Q
Capacity: 800166076416 bytes 800.166076416 GB
Drive Temp: 41 C
Drive Health: Ok
Total namespaces: 1
Namespace id(s): 1
Namespace 1 size: 794996768768 bytes, 794.996768768 GB
Namespace 1 is attached to controller id(s): 33
```

```
NVMe control device: /dev/nvme1
Slot: 7
Manufacturer: SAMSUNG MZWL800HEHP-00003
Serial number: S3HCNX0JC02193
Firmware: GPNA9B3Q
Capacity: 800166076416 bytes 800.166076416 GB
Drive Temp: 41 C
Drive Health: Ok
Total namespaces: 1
Namespace id(s): 1
Namespace 1 size: 794996768768 bytes, 794.996768768 GB
Namespace 1 is attached to controller id(s): 33

NVMe control device: /dev/nvme2
Slot: 6
Manufacturer: SAMSUNG MZWL800HEHP-00003
Serial number: S3HCNX0JC02233
Firmware: GPNA9B3Q
Capacity: 800166076416 bytes 800.166076416 GB
Drive Temp: 40 C
Drive Health: Ok
Total namespaces: 1
Namespace id(s): 1
Namespace 1 size: 794996768768 bytes, 794.996768768 GB
Namespace 1 is attached to controller id(s): 33

NVMe control device: /dev/nvme3
Slot: 5
Manufacturer: SAMSUNG MZWL800HEHP-00003
Serial number: S3HCNX0JC02191
Firmware: GPNA9B3Q
Capacity: 800166076416 bytes 800.166076416 GB
Drive Temp: 40 C
Drive Health: Ok
Total namespaces: 1
Namespace id(s): 1
Namespace 1 size: 794996768768 bytes, 794.996768768 GB
Namespace 1 is attached to controller id(s): 33

NVMe control device: /dev/nvme4
Slot: 4
Manufacturer: SAMSUNG MZWL800HEHP-00003
Serial number: S3HCNX0JC02189
Firmware: GPNA9B3Q
Capacity: 800166076416 bytes 800.166076416 GB
Drive Temp: 39 C
Drive Health: Ok
Total namespaces: 1
Namespace id(s): 1
Namespace 1 size: 794996768768 bytes, 794.996768768 GB
Namespace 1 is attached to controller id(s): 33

NVMe control device: /dev/nvme5
Slot: 3
Manufacturer: SAMSUNG MZWL800HEHP-00003
Serial number: S3HCNX0JC02198
Firmware: GPNA9B3Q
Capacity: 800166076416 bytes 800.166076416 GB
Drive Temp: 39 C
Drive Health: Ok
Total namespaces: 1
Namespace id(s): 1
Namespace 1 size: 794996768768 bytes, 794.996768768 GB
Namespace 1 is attached to controller id(s): 33
```

```
NVMe control device: /dev/nvme6
Slot: 2
Manufacturer: SAMSUNG MZWL800HEHP-00003
Serial number: S3HCNX0JC02188
Firmware: GPNA9B3Q
Capacity: 800166076416 bytes 800.166076416 GB
Drive Temp: 39 C
Drive Health: Ok
Total namespaces: 1
Namespace id(s): 1
Namespace 1 size: 794996768768 bytes, 794.996768768 GB
Namespace 1 is attached to controller id(s): 33

NVMe control device: /dev/nvme7
Slot: 1
Manufacturer: SAMSUNG MZWL800HEHP-00003
Serial number: S3HCNX0JC02195
Firmware: GPNA9B3Q
Capacity: 800166076416 bytes 800.166076416 GB
Drive Temp: 39 C
Drive Health: Ok
Total namespaces: 1
Namespace id(s): 1
Namespace 1 size: 794996768768 bytes, 794.996768768 GB
Namespace 1 is attached to controller id(s): 33
```

Removing the drive from the RAID array

This section covers the steps required to remove a storage drive from the RAID array, before physically removing the drive from the storage enclosure.

In the example, we assume that a drive has failed.

To check the status of the RAID array, use the following command:

```
# cat /proc/mdstat
```

Following is an example output:

```
Personalities : [raid1] [raid6] [raid5] [raid4] [raid0]
md124 : active raid0 md125[0]
4657247232 blocks super 1.2 512k chunks

md125 : active raid5 dm-7[7] dm-6[6] dm-5[5] dm-4[4] dm-3[3](F) dm-2[2] dm-1[1] dm-0[0]
4657379328 blocks super 1.2 level 5, 16k chunk, algorithm 2 [7/6] [UUU_UUU]
[>.....] recovery = 1.5% (12246008/776229888) finish=62.7min speed=202807K/sec
bitmap: 0/6 pages [0KB], 65536KB chunk

md126 : active raid1 sda[1] sdb[0]
118778880 blocks super external:/md127/0 [2/2] [UU]

md127 : inactive sda[1](S) sdb[0](S)
6306 blocks super external:imsm

unused devices: <none>
```

To view detailed information about the status of the RAID array, use the following command:

```
# mdadm --detail /dev/md/md5
```

Following is an example output:

```
/dev/md/md5:
Version : 1.2
Creation Time : Mon Aug 20 11:03:31 2018
Raid Level : raid5
Array Size : 4657379328 (4441.62 GiB 4769.16 GB)
Used Dev Size : 776229888 (740.27 GiB 794.86 GB)
Raid Devices : 7
Total Devices : 8
Persistence : Superblock is persistent

Intent Bitmap : Internal

Update Time : Mon Aug 20 11:50:30 2018
State : clean, degraded, recovering
Active Devices : 6
Working Devices : 7
Failed Devices : 1
Spare Devices : 1

Layout : left-symmetric
Chunk Size : 16K

Consistency Policy : bitmap

Rebuild Status : 20% complete
```

```
Name : any:md5
UUID : 5f7b873c:16d6d418:7b5f6fde:247566a7
Events : 162
```

Number	Major	Minor	RaidDevice	State
0	253	0	0	active sync /dev/dm-0
1	253	1	1	active sync /dev/dm-1
2	253	2	2	active sync /dev/dm-2
7	253	7	3	spare rebuilding /dev/dm-7
4	253	4	4	active sync /dev/dm-4
5	253	5	5	active sync /dev/dm-5
6	253	6	6	active sync /dev/dm-6
3	253	3	-	faulty /dev/dm-3

In this example, dm-3 is marked as failed (dm-3[3](F)).

To remove a drive from the RAID array, follow these steps:

Step 1. Determine the SED mapping for dm-3:

```
# ls -l /dev/mapper/sed*
```

Following is an example output:

```
lrwxrwxrwx. 1 root root 7 Aug 20 11:03 /dev/mapper/sed0 -> ../dm-0
lrwxrwxrwx. 1 root root 7 Aug 20 11:03 /dev/mapper/sed1 -> ../dm-1
lrwxrwxrwx. 1 root root 7 Aug 20 11:03 /dev/mapper/sed2 -> ../dm-2
lrwxrwxrwx. 1 root root 7 Aug 20 11:03 /dev/mapper/sed3 -> ../dm-3
lrwxrwxrwx. 1 root root 7 Aug 20 11:03 /dev/mapper/sed4 -> ../dm-4
lrwxrwxrwx. 1 root root 7 Aug 20 11:03 /dev/mapper/sed5 -> ../dm-5
lrwxrwxrwx. 1 root root 7 Aug 20 11:03 /dev/mapper/sed6 -> ../dm-6
lrwxrwxrwx. 1 root root 7 Aug 20 11:03 /dev/mapper/sed7 -> ../dm-7
```

In this example, dm-3 is mapped to sed3.

Step 2. Determine the NVME namespace — SDE drive mapping for sd3:

```
# cryptsetup status /dev/mapper/sed3
```

Following is an example output:

```
/dev/mapper/sed3 is active and is in use.
type:          LUKS1
cipher:        aes-xts-plain64
keysize:       512 bits
device:        /dev/nvme3n1
offset:        4096 sectors
size:          1552723968 sectors
mode:          read/write
```

The output displays the following information:

- sed3 is an encrypted device mapped to /dev/nvme3n1
- /dev/nvme3n1 is a namespace associated with /dev/nvme3

Step 3. Check the current RAID configuration:

```
# /usr/share/tacp/lenovo/tacp-nvme-control.py -display
```

Following is an example output:

```
discovering NVMe disks. this operation will take a few seconds.
Chassis UUID: 500e0eca08057b00
Number of canisters: 2
This is the bottom (primary) canister and has the controller id 33
The top (secondary) canister has the controller id 34
This chassis has the following controller ids: 33 , 34

NVMe devices: 8
-----

NVMe control device: /dev/nvme0
Slot: 8
Manufacturer: SAMSUNG MZWL800HEHP-00003
Serial number: S3HCNX0JC02194
Firmware: GPNA9B3Q
Capacity: 800166076416 bytes 800.166076416 GB
Drive Temp: 42 C
Drive Health: Ok
Total namespaces: 1
Namespace id(s): 1
Namespace 1 size: 794996768768 bytes, 794.996768768 GB
Namespace 1 is attached to controller id(s): 33

NVMe control device: /dev/nvme1
Slot: 7
Manufacturer: SAMSUNG MZWL800HEHP-00003
Serial number: S3HCNX0JC02193
Firmware: GPNA9B3Q
Capacity: 800166076416 bytes 800.166076416 GB
Drive Temp: 41 C
Drive Health: Ok
Total namespaces: 1
Namespace id(s): 1
Namespace 1 size: 794996768768 bytes, 794.996768768 GB
Namespace 1 is attached to controller id(s): 33

NVMe control device: /dev/nvme2
Slot: 6
Manufacturer: SAMSUNG MZWL800HEHP-00003
Serial number: S3HCNX0JC02233
Firmware: GPNA9B3Q
Capacity: 800166076416 bytes 800.166076416 GB
Drive Temp: 41 C
Drive Health: Ok
Total namespaces: 1
Namespace id(s): 1
Namespace 1 size: 794996768768 bytes, 794.996768768 GB
Namespace 1 is attached to controller id(s): 33

NVMe control device: /dev/nvme3
Slot: 5
Manufacturer: SAMSUNG MZWL800HEHP-00003
Serial number: S3HCNX0JC02191
Firmware: GPNA9B3Q
Capacity: 800166076416 bytes 800.166076416 GB
Drive Temp: 40 C
Drive Health: Ok
Total namespaces: 1
Namespace id(s): 1
Namespace 1 size: 794996768768 bytes, 794.996768768 GB
Namespace 1 is attached to controller id(s): 33
```



```
NVMe control device: /dev/nvme4
Slot: 4
Manufacturer: SAMSUNG MZWL800HEHP-00003
Serial number: S3HCNX0JC02189
Firmware: GPNA9B3Q
Capacity: 800166076416 bytes 800.166076416 GB
Drive Temp: 40 C
Drive Health: Ok
Total namespaces: 1
Namespace id(s): 1
Namespace 1 size: 794996768768 bytes, 794.996768768 GB
Namespace 1 is attached to controller id(s): 33

NVMe control device: /dev/nvme5
Slot: 3
Manufacturer: SAMSUNG MZWL800HEHP-00003
Serial number: S3HCNX0JC02198
Firmware: GPNA9B3Q
Capacity: 800166076416 bytes 800.166076416 GB
Drive Temp: 40 C
Drive Health: Ok
Total namespaces: 1
Namespace id(s): 1
Namespace 1 size: 794996768768 bytes, 794.996768768 GB
Namespace 1 is attached to controller id(s): 33

NVMe control device: /dev/nvme6
Slot: 1
Manufacturer: SAMSUNG MZWL800HEHP-00003
Serial number: S3HCNX0JC02188
Firmware: GPNA9B3Q
Capacity: 800166076416 bytes 800.166076416 GB
Drive Temp: 39 C
Drive Health: Ok
Total namespaces: 1
Namespace id(s): 1
Namespace 1 size: 794996768768 bytes, 794.996768768 GB
Namespace 1 is attached to controller id(s): 33

NVMe control device: /dev/nvme7
Slot: 2
Manufacturer: SAMSUNG MZWL800HEHP-00003
Serial number: S3HCNX0JC02195
Firmware: GPNA9B3Q
Capacity: 800166076416 bytes 800.166076416 GB
Drive Temp: 40 C
Drive Health: Ok
Total namespaces: 1
Namespace id(s): 1
Namespace 1 size: 794996768768 bytes, 794.996768768 GB
Namespace 1 is attached to controller id(s): 33
```

The output shows that drive /dev/nvme3 is in slot 5.

Step 4. Remove the sed3 drive from the RAID array:

```
# mdadm --manage /dev/md/md5 --remove /dev/mapper/sed3
```

Following is an example output:

```
Output: mdadm: hot removed /dev/mapper/sed3 from /dev/md/md5
```

The sed3 drive is nor removed from the RAID array.

Re-check the status of the RAID array to verify that sed3 drive has been removed:

```
# cat /proc/mdstat
```

Following is an example output:

```
Personalities : [raid1] [raid6] [raid5] [raid4] [raid0]
md124 : active raid0 md125[0]
4657247232 blocks super 1.2 512k chunks

md125 : active raid5 dm-7[7] dm-6[6] dm-5[5] dm-4[4] dm-2[2] dm-1[1] dm-0[0]
4657379328 blocks super 1.2 level 5, 16k chunk, algorithm 2 [7/6] [UUU_UUU]
[=====>.....] recovery = 58.5% (454378348/776229888) finish=26.4min speed=202487K/sec
bitmap: 0/6 pages [0KB], 65536KB chunk

md126 : active raid1 sda[1] sdb[0]
118778880 blocks super external:/md127/0 [2/2] [UU]

md127 : inactive sda[1](S) sdb[0](S)
6306 blocks super external:imsm

unused devices: <none>
```

To view detailed information about the status of the RAID array, use the following command:

```
# mdadm --detail /dev/md/md5
```

Following is an example output:

```
/dev/md/md5:
Version : 1.2
Creation Time : Mon Aug 20 11:03:31 2018
Raid Level : raid5
Array Size : 4657379328 (4441.62 GiB 4769.16 GB)
Used Dev Size : 776229888 (740.27 GiB 794.86 GB)
Raid Devices : 7
Total Devices : 7
Persistence : Superblock is persistent

Intent Bitmap : Internal

Update Time : Mon Aug 20 12:18:05 2018
State : clean, degraded, recovering
Active Devices : 6
Working Devices : 7
Failed Devices : 0
Spare Devices : 1

Layout : left-symmetric
Chunk Size : 16K

Consistency Policy : bitmap

Rebuild Status : 62% complete
```

```
Name : any:md5
UUID : 5f7b873c:16d6d418:7b5f6fde:247566a7
Events : 480

Number Major Minor RaidDevice State
0      253   0     0      active sync /dev/dm-0
1      253   1     1      active sync /dev/dm-1
2      253   2     2      active sync /dev/dm-2
7      253   7     3      spare rebuilding /dev/dm-7
4      253   4     4      active sync /dev/dm-4
5      253   5     5      active sync /dev/dm-5
6      253   6     6      active sync /dev/dm-6
```

Step 5. Close the SED:


```
# cryptsetup close sed3
```

The drive is now removed from the RAID array.

The drive can now be physically removed from storage block slot 5. Follow the steps in [“Removing a hot-swap drive” on page 52](#)

Removing a hot-swap drive

Use this information to remove a hot-swap drive.



Attention:

- To ensure that there is adequate system cooling, do not operate the solution for more than two minutes without either a drive or a filler panel installed in each bay.
- Before you make changes to the drives, back up all important data stored on them.
- If one or more NVMe solid-state drives need to be removed, disable them in the OS before removing them.

Procedure performed by: customer or field service

To remove a hot-swap drive, complete the following steps:

Step 1. To open the drive tray handle, press the tab.




Figure 18. Hot-swap drive removal

Step 2. Grasp the handle and slide the drive out of the drive bay.




Figure 19. Hot-swap drive removal

Step 3. Check the status of the RAID array to verify that slot 5 is empty:

```
# /usr/share/tacp/lenovo/tacp-nvme-control.py -display
```

Following is an example output:

```
discovering NVMe disks. this operation will take a few seconds.
Chassis UUID: 500e0eca08057b00
Number of canisters: 2
This is the bottom (primary) canister and has the controller id 33
The top (secondary) canister has the controller id 34
This chassis has the following controller ids: 33 , 34

NVMe devices: 7
-----

NVMe control device: /dev/nvme0
Slot: 8
Manufacturer: SAMSUNG MZWL800HEHP-00003
Serial number: S3HCNX0JC02194
Firmware: GPNA9B3Q
Capacity: 800166076416 bytes 800.166076416 GB
Drive Temp: 42 C
Drive Health: Ok
Total namespaces: 1
Namespace id(s): 1
Namespace 1 size: 794996768768 bytes, 794.996768768 GB
Namespace 1 is attached to controller id(s): 33
```

```
NVMe control device: /dev/nvme1
Slot: 7
Manufacturer: SAMSUNG MZWL800HEHP-00003
Serial number: S3HCNX0JC02193
Firmware: GPNA9B3Q
Capacity: 800166076416 bytes 800.166076416 GB
Drive Temp: 42 C
Drive Health: Ok
Total namespaces: 1
Namespace id(s): 1
Namespace 1 size: 794996768768 bytes, 794.996768768 GB
Namespace 1 is attached to controller id(s): 33

NVMe control device: /dev/nvme2
Slot: 6
Manufacturer: SAMSUNG MZWL800HEHP-00003
Serial number: S3HCNX0JC02233
Firmware: GPNA9B3Q
Capacity: 800166076416 bytes 800.166076416 GB
Drive Temp: 41 C
Drive Health: Ok
Total namespaces: 1
Namespace id(s): 1
Namespace 1 size: 794996768768 bytes, 794.996768768 GB
Namespace 1 is attached to controller id(s): 33

NVMe control device: /dev/nvme4
Slot: 4
Manufacturer: SAMSUNG MZWL800HEHP-00003
Serial number: S3HCNX0JC02189
Firmware: GPNA9B3Q
Capacity: 800166076416 bytes 800.166076416 GB
Drive Temp: 40 C
Drive Health: Ok
Total namespaces: 1
Namespace id(s): 1
Namespace 1 size: 794996768768 bytes, 794.996768768 GB
Namespace 1 is attached to controller id(s): 33

NVMe control device: /dev/nvme5
Slot: 3
Manufacturer: SAMSUNG MZWL800HEHP-00003
Serial number: S3HCNX0JC02198
Firmware: GPNA9B3Q
Capacity: 800166076416 bytes 800.166076416 GB
Drive Temp: 40 C
Drive Health: Ok
Total namespaces: 1
Namespace id(s): 1
Namespace 1 size: 794996768768 bytes, 794.996768768 GB
Namespace 1 is attached to controller id(s): 33
```

```

NVMe control device: /dev/nvme6
Slot: 1
Manufacturer: SAMSUNG MZWL800HEHP-00003
Serial number: S3HCNX0JC02188
Firmware: GPNA9B3Q
Capacity: 800166076416 bytes 800.166076416 GB
Drive Temp: 40 C
Drive Health: Ok
Total namespaces: 1
Namespace id(s): 1
Namespace 1 size: 794996768768 bytes, 794.996768768 GB
Namespace 1 is attached to controller id(s): 33

NVMe control device: /dev/nvme7
Slot: 2
Manufacturer: SAMSUNG MZWL800HEHP-00003
Serial number: S3HCNX0JC02195
Firmware: GPNA9B3Q
Capacity: 800166076416 bytes 800.166076416 GB
Drive Temp: 41 C
Drive Health: Ok
Total namespaces: 1
Namespace id(s): 1
Namespace 1 size: 794996768768 bytes, 794.996768768 GB
Namespace 1 is attached to controller id(s): 33

```

The output shows that there is no drive in slot 5.

Step 4. Check that the NVME namespace to verify that nvme3 is deleted:

```
# ls -l /dev/nvme
```

Following is an example output:

```

crw----- 1 root root 243, 0 Aug 20 11:02 /dev/nvme0
crw----- 1 root root 243, 1 Aug 20 11:02 /dev/nvme1
crw----- 1 root root 243, 2 Aug 20 11:02 /dev/nvme2
crw----- 1 root root 243, 4 Aug 20 11:02 /dev/nvme4
crw----- 1 root root 243, 5 Aug 20 11:02 /dev/nvme5
crw----- 1 root root 243, 6 Aug 20 11:02 /dev/nvme6
crw----- 1 root root 243, 7 Aug 20 11:02 /dev/nvme7

```

The output shows that there is no nvme3.


The removal of the hot-swap drive is now complete.

After removing the hot-swap drive:

- Install a drive filler or a new drive to cover the drive bay. See [“Installing a hot-swap drive”](#) on page 56.
- If you are instructed to return the component or optional device, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Installing a hot-swap drive

Use this information to install a hot-swap drive.



Procedure performed by: customer or field service

Notes:

- The drive bays are numbered to indicate the installation order (starting from number “01”). Follow this installation order when you install a drive.
- You cannot mix drives of different capacities in one system.
- In this procedure, dm-3 and nvme3 are used as examples. Replace these examples with the appropriate values for the storage drive you are installing.

To install a hot-swap drive, complete the following steps:

Step 1. To open the drive tray handle, press the release tab.




Figure 20. Opening the drive tray handle

Step 2. Slide the drive into the drive bay until it snaps into position.




Figure 21. Sliding the drive into the drive bay

Step 3. Close the drive tray handle to lock the drive in place.

Step 4. Check the drive status LED to verify that the drive is operating correctly.

- If the yellow drive status LED is lit continuously, that drive is faulty and must be replaced.
- If the green drive activity LED is flashing, the drive is being accessed.

Step 5. Continue to install additional hot-swap drives if necessary.

Step 6. Check that the new drives are detected by the storage block:

Check that the NVME namespace to verify that `nvme3` appears:

```
# ls -l /dev/nvme
```

Following is an example output:

```
crw-----. 1 root root 243, 0 Aug 20 11:02 /dev/nvme0
crw-----. 1 root root 243, 1 Aug 20 11:02 /dev/nvme1
crw-----. 1 root root 243, 2 Aug 20 11:02 /dev/nvme2
crw-----. 1 root root 243, 3 Aug 20 12:32 /dev/nvme3
crw-----. 1 root root 243, 4 Aug 20 11:02 /dev/nvme4
crw-----. 1 root root 243, 5 Aug 20 11:02 /dev/nvme5
crw-----. 1 root root 243, 6 Aug 20 11:02 /dev/nvme6
crw-----. 1 root root 243, 7 Aug 20 11:02 /dev/nvme7
```

The output show that `nvme3` is present in the NVME namespace.

Check the status of the RAID array:

```
# /usr/share/tacp/lenovo/tacp-nvme-control.py -display
```

Following is an example output:

```
discovering NVMe disks. this operation will take a few seconds.
Chassis UUID: 500e0eca08057b00
Number of canisters: 2
This is the bottom (primary) canister and has the controller id 33
The top (secondary) canister has the controller id 34
This chassis has the following controller ids: 33 , 34

NVMe devices: 8
-----

NVMe control device: /dev/nvme0
Slot: 8
Manufacturer: SAMSUNG MZWL800HEHP-00003
Serial number: S3HCNX0JC02194
Firmware: GPNA9B3Q
Capacity: 800166076416 bytes 800.166076416 GB
Drive Temp: 42 C
Drive Health: Ok
Total namespaces: 1
Namespace id(s): 1
Namespace 1 size: 794996768768 bytes, 794.996768768 GB
Namespace 1 is attached to controller id(s): 33

NVMe control device: /dev/nvme1
Slot: 7
Manufacturer: SAMSUNG MZWL800HEHP-00003
Serial number: S3HCNX0JC02193
Firmware: GPNA9B3Q
Capacity: 800166076416 bytes 800.166076416 GB
Drive Temp: 42 C
Drive Health: Ok
Total namespaces: 1
Namespace id(s): 1
Namespace 1 size: 794996768768 bytes, 794.996768768 GB
Namespace 1 is attached to controller id(s): 33

NVMe control device: /dev/nvme2
Slot: 6
Manufacturer: SAMSUNG MZWL800HEHP-00003
Serial number: S3HCNX0JC02233
Firmware: GPNA9B3Q
Capacity: 800166076416 bytes 800.166076416 GB
Drive Temp: 41 C
Drive Health: Ok
Total namespaces: 1
Namespace id(s): 1
Namespace 1 size: 794996768768 bytes, 794.996768768 GB
Namespace 1 is attached to controller id(s): 33

NVMe control device: /dev/nvme3
Slot: 5
Manufacturer: SAMSUNG MZWL1T6HEHP-00003
Serial number: S3HDNX0JB00078
Firmware: GPNA9B3Q
Capacity: 1600321314816 bytes 1600.32131482 GB
Drive Temp: 39 C
Drive Health: Ok
Total namespaces: 0
```

```
NVMe control device: /dev/nvme4
Slot: 4
Manufacturer: SAMSUNG MZWL800HEHP-00003
Serial number: S3HCNX0JC02189
Firmware: GPNA9B3Q
Capacity: 800166076416 bytes 800.166076416 GB
Drive Temp: 40 C
Drive Health: Ok
Total namespaces: 1
Namespace id(s): 1
Namespace 1 size: 794996768768 bytes, 794.996768768 GB
Namespace 1 is attached to controller id(s): 33

NVMe control device: /dev/nvme5
Slot: 3
Manufacturer: SAMSUNG MZWL800HEHP-00003
Serial number: S3HCNX0JC02198
Firmware: GPNA9B3Q
Capacity: 800166076416 bytes 800.166076416 GB
Drive Temp: 40 C
Drive Health: Ok
Total namespaces: 1
Namespace id(s): 1
Namespace 1 size: 794996768768 bytes, 794.996768768 GB
Namespace 1 is attached to controller id(s): 33

NVMe control device: /dev/nvme6
Slot: 1
Manufacturer: SAMSUNG MZWL800HEHP-00003
Serial number: S3HCNX0JC02188
Firmware: GPNA9B3Q
Capacity: 800166076416 bytes 800.166076416 GB
Drive Temp: 39 C
Drive Health: Ok
Total namespaces: 1
Namespace id(s): 1
Namespace 1 size: 794996768768 bytes, 794.996768768 GB
Namespace 1 is attached to controller id(s): 33

NVMe control device: /dev/nvme7
Slot: 2
Manufacturer: SAMSUNG MZWL800HEHP-00003
Serial number: S3HCNX0JC02195
Firmware: GPNA9B3Q
Capacity: 800166076416 bytes 800.166076416 GB
Drive Temp: 40 C
Drive Health: Ok
Total namespaces: 1
Namespace id(s): 1
Namespace 1 size: 794996768768 bytes, 794.996768768 GB
Namespace 1 is attached to controller id(s): 33
```

The new drive now appears in slot 5.

Step 7. Prepare the replacement drive by running script `add_disks_lenovo.py`. The script creates namespaces and layout encryption.

```
# /usr/share/tacp/lenovo/add_disks_lenovo.py --setup_replacement_disk
```

Following is an example output:

```
Encrypted device /dev/mapper/sed3 /dev/dm-3  
Please continue with the instructions to add the disk
```

Step 8. Grow the RAID array:

```
# mdadm /dev/md/md5 --add /dev/dm-3
```

Following is an example output:

```
Output: mdadm: added /dev/dm-3
```

Step 9. Check the status of the RAID array:

```
# cat /proc/mdstat
```

Following is an example output:

```
Personalities : [raid1] [raid6] [raid5] [raid4] [raid0]  
md124 : active raid0 md125[0]  
4657247232 blocks super 1.2 512k chunks  
  
md125 : active raid5 dm-3[8](S) dm-7[7] dm-6[6] dm-5[5] dm-4[4] dm-2[2] dm-1[1] dm-0[0]  
4657379328 blocks super 1.2 level 5, 16k chunk, algorithm 2 [7/7] [UUUUUUU]  
bitmap: 0/6 pages [0KB], 65536KB chunk  
  
md126 : active raid1 sda[1] sdb[0]  
118778880 blocks super external:/md127/0 [2/2] [UU]  
  
md127 : inactive sda[1](S) sdb[0](S)  
6306 blocks super external:imsm  
  
unused devices: <none>
```

To view detailed information about the status of the RAID array, use the following command:

```
# mdadm --detail /dev/md/md5
```

Following is an example output:

```
/dev/md/md5:
Version : 1.2
Creation Time : Mon Aug 20 11:03:31 2018
Raid Level : raid5
Array Size : 4657379328 (4441.62 GiB 4769.16 GB)
Used Dev Size : 776229888 (740.27 GiB 794.86 GB)
Raid Devices : 7
Total Devices : 8
Persistence : Superblock is persistent

Intent Bitmap : Internal

Update Time : Mon Aug 20 15:27:26 2018
State : clean
Active Devices : 7
Working Devices : 8
Failed Devices : 0
Spare Devices : 1

Layout : left-symmetric
Chunk Size : 16K

Consistency Policy : bitmap

Name : any:md5
UUID : 5f7b873c:16d6d418:7b5f6fde:247566a7
Events : 749

Number Major Minor RaidDevice State
0      253  0    0      active sync /dev/dm-0
1      253  1    1      active sync /dev/dm-1
2      253  2    2      active sync /dev/dm-2
7      253  7    3      active sync /dev/dm-7
4      253  4    4      active sync /dev/dm-4
5      253  5    5      active sync /dev/dm-5
6      253  6    6      active sync /dev/dm-6

8      253  3    -      spare /dev/dm-3
```

The new drive is now added to the RAID array as a spare device.

The installation of a new hot-swap drive is now complete.


Hot-swap power supply replacement

Use this information to remove and install a hot-swap power supply.

Note: Use any documentation that comes with the power supply and follow those instructions in addition to the instructions in these topics. Ensure that you have all the cables and other equipment that come with the power supply, as specified in its documentation.

Remove a hot-swap power supply

Use this information to replace a hot-swap power supply.



S035



CAUTION:

Never remove the cover on a power supply or any part that has this label attached.

Hazardous voltage, current, and energy levels are present inside any component that has this label attached. There are no serviceable parts inside these components. If you suspect a problem with one of these parts, contact a service technician.

S002



CAUTION:

The power control button on the device and the power switch on the power supply do not turn off the electrical current that is supplied to the device. The device also might have more than one power cord. To remove all electrical current from the device, ensure that all power cords are disconnected from the power source.

S001



 **DANGER**

**Electrical current from power, telephone, and communication cables is hazardous.
To avoid a shock hazard:**

- **Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.**
- **Connect all power cords to a properly wired and grounded electrical outlet.**
- **Connect to properly wired outlets any equipment that will be attached to this product.**
- **When possible, use one hand only to connect or disconnect signal cables.**
- **Never turn on any equipment when there is evidence of fire, water, or structural damage.**
- **Disconnect the attached power cords, telecommunications systems, networks, and modems before you open the device covers, unless instructed otherwise in the installation and configuration procedures.**
- **Connect and disconnect cables as described in the following table when installing, moving, or opening covers on this product or attached devices.**

To Connect:

1. Turn everything OFF.
2. First, attach all cables to devices.
3. Attach signal cables to connectors.
4. Attach power cords to outlet.
5. Turn device ON.

To Disconnect:

1. Turn everything OFF.
2. First, remove power cords from outlet.
3. Remove signal cables from connectors.
4. Remove all cables from devices.

Procedure performed by: customer or field service

Important:

- To ensure proper cooling during normal storage operation, both of the power supply bays must be occupied. Each bay must have a power supply installed.
- If you are replacing two power supplies, replace the power supplies one by one to ensure that the power supply to the storage is not interrupted. Do not disconnect the power cord from the second-replaced power supply until the power output LED for the first-replaced power supply is lit.

To remove a hot-swap power supply, complete the following steps:

- Step 1. Disconnect the power cord from the hot-swap power supply.
- Step 2. Open the handle.




Figure 22. Hot-swap power supply removal (1/2)

Step 3. Press the release tab down, and pull the handle at the same time to slide the hot-swap power supply out of the enclosure.




Figure 23. Hot-swap power supply removal (2/2)


The removal of the hot-swap power supply is now complete.

After removing the hot-swap drive:

- Install a new power supply unit. See [“Install a hot-swap power supply” on page 66](#).
- If you are instructed to return the component or optional device, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Install a hot-swap power supply

Use this information to install a hot-swap power supply.



The following tips describe the type of power supply that the storage supports and other information that you must consider when you install a power supply:

S035



CAUTION:

Never remove the cover on a power supply or any part that has this label attached.

Hazardous voltage, current, and energy levels are present inside any component that has this label attached. There are no serviceable parts inside these components. If you suspect a problem with one of these parts, contact a service technician.

S002



CAUTION:

The power control button on the device and the power switch on the power supply do not turn off the electrical current that is supplied to the device. The device also might have more than one power cord. To remove all electrical current from the device, ensure that all power cords are disconnected from the power source.

S001



 **DANGER**

**Electrical current from power, telephone, and communication cables is hazardous.
To avoid a shock hazard:**

- **Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.**
- **Connect all power cords to a properly wired and grounded electrical outlet.**
- **Connect to properly wired outlets any equipment that will be attached to this product.**
- **When possible, use one hand only to connect or disconnect signal cables.**
- **Never turn on any equipment when there is evidence of fire, water, or structural damage.**
- **Disconnect the attached power cords, telecommunications systems, networks, and modems before you open the device covers, unless instructed otherwise in the installation and configuration procedures.**
- **Connect and disconnect cables as described in the following table when installing, moving, or opening covers on this product or attached devices.**

To Connect:

1. Turn everything OFF.
2. First, attach all cables to devices.
3. Attach signal cables to connectors.
4. Attach power cords to outlet.
5. Turn device ON.

To Disconnect:

1. Turn everything OFF.
2. First, remove power cords from outlet.
3. Remove signal cables from connectors.
4. Remove all cables from devices.

Before installing a hot-swap power supply:

1. Touch the static-protective package containing the new hot-swap power supply to any unpainted surface on the outside of the storage.
2. Take the new hot-swap power supply out of the package and place it on a static-protective surface.

Procedure performed by: customer or field service

To install a hot-swap power supply, complete the following steps:

- Step 1. Ensure that the handle is in the open position.
- Step 2. Press the release tab down, and slide the new hot-swap power supply into the bay at the same time until the release latch clicks into place.




Figure 24. Hot-swap power supply installation (1/2)

Step 3. Close the handle.




Figure 25. Hot-swap power supply installation (2/2)

The installation of a hot-swap power supply is now complete.

After installing the power supply:

- Connect one end of the power cord to the new power supply connector. Then, connect the other end of the power cord into a properly grounded electrical outlet.
- Ensure that the power supply LED is lit, indicating that the power supply is operating properly.

Hot-swap storage controller replacement

Use this information to remove and install a storage controller. You can remove or install a storage controller without turning off the storage, which helps you avoid significant interruption to the operation of the system.

Remove a hot-swap storage controller

This procedure describes how to remove a storage controller from a ThinkAgile CP configuration.

Procedure performed by: customer or field service

To remove a hot-swap storage controller, complete the following steps:

The solution administrator completes the following steps:

- Step 1. Obtain the serial number of the storage controller to be removed. You can find the serial number on the Hardware page of the Web Portal.
- Step 2. If the storage controller is running and is active, perform a Storage HA event by following the procedure in [“Trigger manual storage HA” on page 34](#).
- Step 3. Shut down the storage controller by following the procedure in [“Shutting off a storage controller” on page 34](#).

The Customer or field service completes the following steps:

- Step 4. Open the release latches on the storage controller outwards.
- Step 5. Pull the storage controller to slide it out of the enclosure.




Figure 26. Hot-swap storage controller removal

The solution administrator completes the following step:

- Step 6. Enable [Chapter 4 “Support Mode” on page 23](#).

ThinkAgile Advantage completes the following step:

Step 7. ThinkAgile Advantage signs in to the primary interconnect. The username is `admin` and the password is the one you set during the installation.

The solution administrator completes the following steps:

Step 8. Run the following command on the primary interconnect:

```
> sudo service tacp-network-controller restart
```

ThinkAgile Advantage completes the following step:

Step 9. ThinkAgile Advantage gathers the relevant information, contacts support, and requests a hardware removal from the database by providing the serial number.

The removal of the storage controller is now complete.

After removing the storage controller:

- Install a new storage controller. See [“Install a hot-swap storage controller” on page 71](#)
- If you are instructed to return the component or optional device, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Install a hot-swap storage controller

This procedure explains how to add a storage controller to the ThinkAgile CP configuration, assuming a storage controller has been previously removed.

Procedure performed by: customer or field service

To install a hot-swap storage controller, complete the following steps:

The customer or field service preforms the following steps:

- Step 1. Ensure that the release latch is in the open position. Slide the canister all the way into the enclosure.
- Step 2. Secure the storage controller by rotating the release latch until it clicks into the closed position.




Figure 27. Hot-swap storage controller installation

Notes:

- The electromagnetic interference (EMI) integrity and cooling of the storage are protected by having all bays covered or occupied. When replacing a storage controller, minimize the amount of time of the empty bay to avoid overheating or abnormal operation.
- The storage controller on the upper bay is upside down. Ensure that the storage controller is in the right position (the release latch is on top) before installing.

ThinkAgile Advantage Support performs the following step:

Step 3. ThinkAgile Advantage Support updates the VPD to maintain SN/UUID for entitlement. See [“Update the storage node VPD” on page 74](#).

The solution administrator performs the following steps:

Step 4. Sign in to the primary interconnect with [Chapter 4 “Support Mode” on page 23](#). For the username enter `admin` and enter the password set during the solution deployment.

Step 5. Run the following command on the primary interconnect:

```
> sudo service tacp-network-controller restart
```

Step 6. In the Cloud Controller dashboard, select **Hardware** from the left menu.

Step 7. On the node that you want to register, click **Register**. The Register Node page displays.

Step 8. In the **Storage Block Name** field, enter a unique name for the storage block.

Step 9. In the **Administrator Password** field, enter a password for the manager account on the storage block.

Step 10. In the **Storage Pool** drop-down, assign the storage block to a storage pool from the list.

In the Cloud Controller, storage resources (storage blocks) are organized into storage pools. A storage block must be added to a storage pool for it to store application instances.

To create a storage pool:

- a. In the **Storage Pool** drop-down, select **New storage pool**
- b. Enter a name for the storage pool.
- c. Click **Create**.

Step 11. Click **Register Storage Block** to complete registration.

Step 12. Repeat this process to register all remaining storage blocks.

The installation of a hot-swap storage controller is now complete.

Update the storage node VPD

This section covers the procedure to update the ThinkAgile CP storage node VPD.

Prior to updating the storage node VPD, perform the following tasks:

- “Enable Support Mode in the Cloud Controller” on page 23
- “Access the customer’s system via Support Mode” on page 24
- “Place a support file on the storage or compute node” on page 26

Check if the SES driver is present by entering the command:

```
lsmod | grep cls
```

If the driver is not present, insert it with the following command:

```
insmod /usr/share/tacp/lenovo/cls.ko
```

Procedure performed by: ThinkAgile Advantage Support

To pull current FRU information before beginning the VPD reprogram process (also logged during reprogram), use the **check** option on the ThinkAgile CP-SB VPD tool (`./thinkagile_cpsb_vpd_tool.sh check`).

To update the storage node VPD, follow these steps:

Step 1. Download the VPD tool bundle from <https://support.lenovo.com/us/en/downloads/ds506277>.

Step 2. Copy the VPD tool bundle (`thinkagile_cpsb_vpd_tool_vX.X.tgz`) onto the storage controller. See “Place a support file on the storage or compute node” on page 26.

Step 3. Untar the VPD tool and navigate to the tool’s root directory (`thinkagile_cpsb_vpd_tool`).

Step 4. Execute the following command to bring up the parameters for the programmable fields:

```
./thinkagile_cpsb_vpd_tool.sh help
```

Step 5. Execute the following command to begin the reprogramming process for specified FRU fields:

```
./thinkagile_cpsb_vpd_tool.sh -p <fields>
```

Note: Be mindful of maximum lengths as outlined in included readme document.

Step 6. The VPD tool prompts you for new values for each specified field and will reprogram with the new values. Enter the new values as directed.

Step 7. The VPD tool verifies that the programming process has completed. Confirm that the values are correct.

Step 8. Remove all artifacts from the VPD tool with the following commands:

```
rm -rf thinkagile_cpsb_vpd_tool
rm -rf thinkagile_cpsb_vpd_tool_vX.X.tgz
```

Logs are stored under the `log/` directory and contains all prior FRU information as well as the fields programmed. The previous values for programmed fields are listed under `src/previous_FRU_info.txt`.

Refer to the included readme document for more info on VPD tool.

Note: If you are reprogramming the controller model (canister board product): if they are mismatched on the two storage controllers, ESM B (the top controller) enters lockdown mode. To remedy this, ensure that the canister model on both storage controllers is the same. As necessary (if B is current active controller), failover all traffic and applications to ESM A (bottom controller) and then hard-reboot ESM B.

Access the storage controller using Support Mode

This section covers the procedure to access the storage controller using Support Mode.

Performed By: Lenovo Support

To access the storage controller using Support Mode, follow these steps:

- Step 1. Have the customer or administrator access the Cloud Controller and enable Support Mode (see [Chapter 4 “Managing Support Mode” on page 23](#)).
- Step 2. In the Cloud Controller, navigate to the Organizational Viewer topic and select the customer organization.

A temporary port and password are generated to access the customer environment using a proxy server.

- Step 3. Sign in to the customer environment using SSH. Use the temporary port and password generated in [Step 2 step 2 on page 75](#) to sign in to the proxy server..

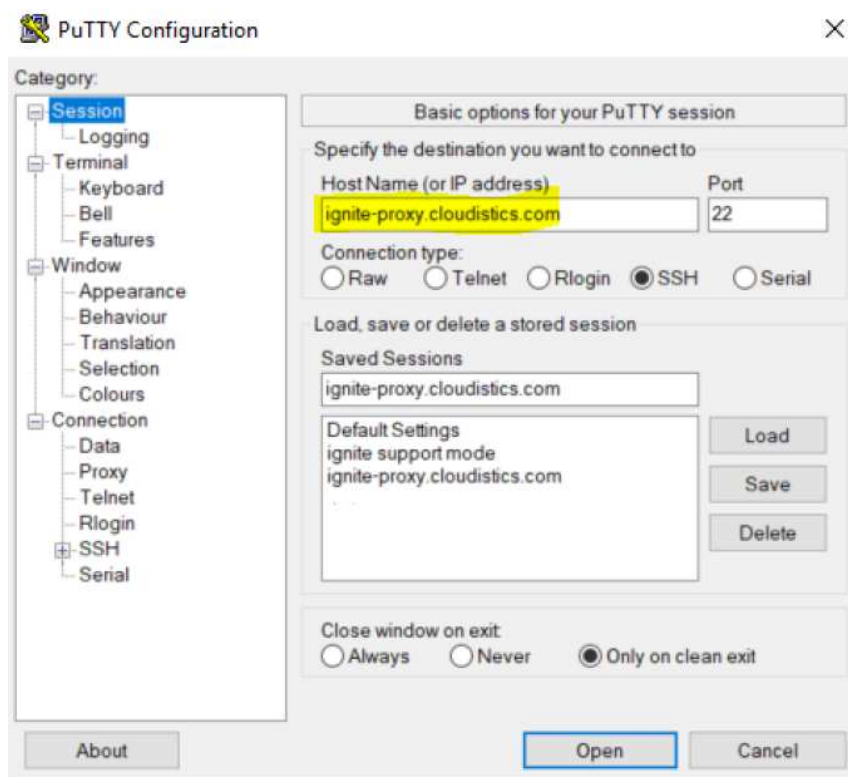
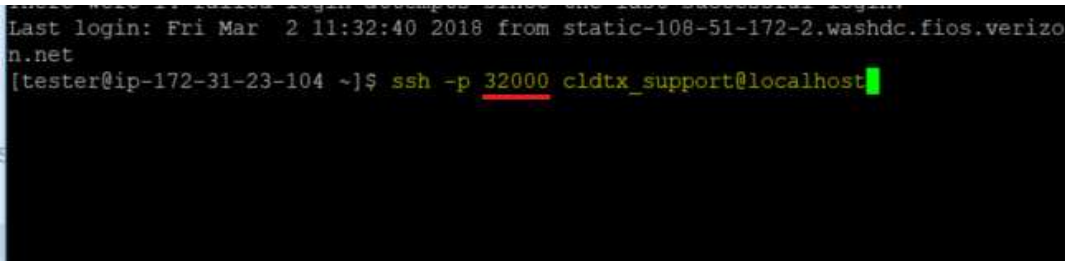


Figure 28. SSH sign-in to proxy server

Step 4. From the proxy server, access the port as found in the Cloud Controller by entering following command at the prompt:

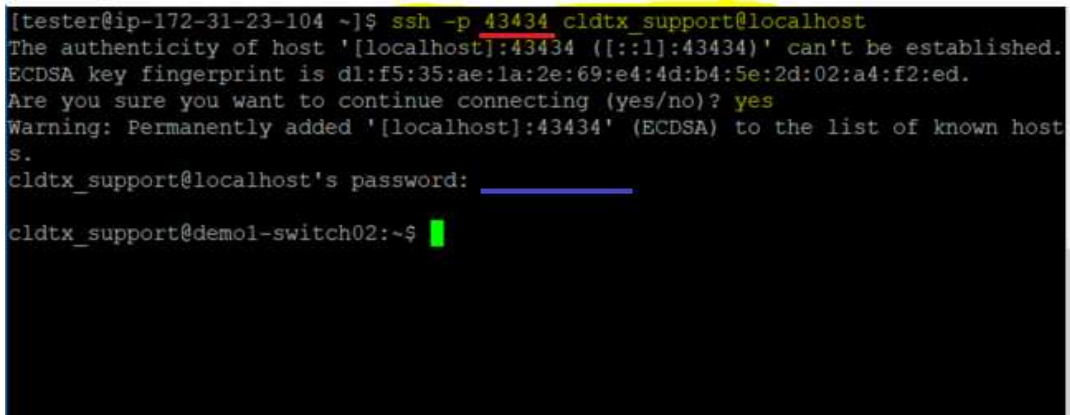
```
ssh -p <port> cldtx_support@localhost
```



```
Last login: Fri Mar 2 11:32:40 2018 from static-108-51-172-2.washdc.fios.verizon.net
[tester@ip-172-31-23-104 ~]$ ssh -p 32000 cldtx_support@localhost
```

Figure 29. Entering the proxy server port

Enter *yes* at the prompt. The password is located next to the port number in Cloud Controller.



```
[tester@ip-172-31-23-104 ~]$ ssh -p 43434 cldtx_support@localhost
The authenticity of host '[localhost]:43434 ([::1]:43434)' can't be established.
ECDSA key fingerprint is dl:f5:35:ae:1a:2e:69:e4:4d:b4:5e:2d:02:a4:f2:ed.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '[localhost]:43434' (ECDSA) to the list of known hosts.
cldtx_support@localhost's password: _____
cldtx_support@demo1-switch02:~$
```

Figure 30. Entering the proxy server password

Obtain storage controller IP addresses

This section covers the procedure to obtain the storage controller IP addresses.

To obtain the storage controller IP addresses, follow these steps:

Step 1. To obtain the storage controller IP addresses from the Cloud Controller:

- a. Navigate to the **Organization Viewer** topic.
- b. Locate the rows of storage controller data.
- c. Locate the **IP** column for the storage controller IP addresses.

Step 2. To obtain the storage controller IP addresses from the CLI:

- a. Sign in to the interconnect as in [“Access the storage controller using Support Mode” on page 75](#).
- b. Enter the command:

```
cat /var/lib/dhcp/dhcpd.leases | grep storage -B 10 | grep lease
```

IP addresses you can use to sign in to the storage controllers display.

Step 3. To obtain the IP address of the storage controller’s Baseboard Management Controller (BMC), enter the command:

```
cat /var/lib/dhcp/dhcpd.leases | grep Athena -B 10 | grep lease
```

Use these IP addresses to connect to the Web interface using HTTPS.

Transfer the VPD utility to the storage controller

This section covers the procedure to transfer the VPD utility to the storage controller.

To transfer the VPD utility to the storage controller, follow these steps:

Step 1. In the CLI, SSH into one of the storage controllers by entering:

```
ssh -I .ssh/cldt [tab to get support controller ID] cldtx_support@<Storage IP>
```

Note: Use one of the storage controller IP addresses from [“Obtain storage controller IP addresses” on page 77](#).

Step 2. If there is an interconnect switch link, add the following to the command:

```
-b <primary or secondary interconnect IP>
```

This changes the command to:

```
ssh -I .ssh/cldt [tab to get support controller ID] -b <primary or secondary interconnect IP>  
cldtx_support@[Storage IP]
```

To check for the switch IP, execute `ip a` and look for IPs under VLAN 4 or 5. An address ending in 1 will be the primary interconnect switch, while a 2 designates the secondary interconnect switch.

Step 3. Create a directory named `thinkagile_cpsb_vpd_tool` on the proxy server to copy the `thinkagile_cpsb_vpd_tool_vX.X.tgz` utility bundle from your laptop to the proxy server first.

Once the bundle is copied onto the proxy server, enter the following command to copy the utility to the storage controller:


```
scp thinkagile_cpsb_vpd_tool cldtx_support@<storage IP>:/home/manager/
```

Enclosure replacement

Use this information to replace a storage enclosure.

Remove the enclosure

Use this information to remove the storage enclosure.



Before removing the enclosure:

1. Shut down the storage.
 - a. Select **Power off** from the drop-down menu.
 - b. Review the list of application instances that will be affected when the storage pool is powered off. Select **Shut down** or **Force shut down** for each instance.
 - c. Click **Continue**.
2. Locate and disconnect all the cables connected to the storage.

Attention: Disengage all latches, cable clips, release tabs, or locks on cable connectors before proceeding. Failing to release these items before removing the cables damages the cable connectors on the system board. Any damage to the cable connectors may require replacing the system board.

Procedure performed by: Solution administrator, customer or field service, ThinkAgile Advantage Support

To remove the enclosure, complete the following steps:

- Step 1. Remove all the hot-swap drives from the enclosure and put them in a safe, static-protective place. See [“Removing a hot-swap drive” on page 52](#).
- Step 2. Remove both side covers on the enclosure by grasping and pulling out the bottom of the side cover, and then pulling it straight up.




Figure 31. Side cover removal

- Step 3. Remove the M5 screws and thumbscrews on both sides used to fasten the storage enclosure. Keep the M5 screws for installing the new enclosure.




Figure 32. Screws removal

- Step 4. Slide the storage enclosure out of the rack along the rails and put it in a safe, static-protective place.




Figure 33. Enclosure removal from rack

Step 5. Remove any of the following components that are installed on the enclosure and put them in a safe, static-protective place:

- [“Hot-swap power supply” on page 62](#)
- [“Storage controllers” on page 70](#)


The removal of the storage enclosure is now complete.

After removing the enclosure:

- Install a new enclosure. See [“Install the enclosure” on page 81](#).
- If you are instructed to return the component or optional device, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Install the enclosure

Use this information to install the enclosure.



Before installing the enclosure, touch the static-protective package that contains the new enclosure to any unpainted surface on the outside of the storage. Then, take the new enclosure out of the package and place it in a safe, static-protective place.

Procedure performed by: customer or field service

To install the enclosure, complete the following steps:

Step 1. Install any of the following components that are previously removed back to the new enclosure:

- “Hot-swap power supply” on page 66
- “Storage controllers” on page 71

To install the ThinkAgile CP storage enclosure on the slide rails, do the following:

Step 2. Remove both of the side covers on the enclosure by grasping and pulling out the bottom of each side cover, and then pulling it straight up.




Figure 34. Removal of side cover

Step 3. Lift the enclosure and slowly place the enclosure on the rail.

Step 4. Carefully slide the enclosure into the rack along the rails until the enclosure is fully inserted.




Figure 35. Installing the enclosure into the rack

Step 5. Secure the enclosure with the thumbscrews and the M5 screws on the front.




Figure 36. Installing screws on front of enclosure

Step 6. On each side of the storage enclosure, fit the hook on the top of the side cover and push the bottom of it inward to reinstall the side cover to the storage enclosure.




Figure 37. Reinstalling the side cover

Step 7. Install the hot-swap drives that were previously removed. See [“Installing a hot-swap drive” on page 56](#).

Step 8. Connect all the cables and start the enclosure.

The installation of the storage enclosure is now complete.

After installing the enclosure:

- Connect the cables and start the storage.
- Configure the storage.

Complete the parts replacement

Use this information to complete the parts replacement.

To complete the parts replacement, follow these steps:

- Step 1. Ensure that all components have been reassembled correctly and that no tools or loose screws have been left inside your storage.
- Step 2. Properly install and secure the canisters or power supply in the storage.
- Step 3. Reconnect external cables and power cords to the storage.

Attention: To avoid component damage, connect the power cords last.

- Step 4. Update the storage configuration if necessary.
 - a. Download and install the latest device drivers: <https://datacentersupport.lenovo.com/products/solutions-and-software/thinkagile-cp/>
 - b. Update the system firmware.

Compute block maintenance procedures

This section covers the procedures required to perform compute block maintenance.

Refer to [Chapter 2 “Roles and responsibilities” on page 15](#) to determine who can perform these procedures.

Shutting off a compute node

This procedure describes how to shut off a compute node and how to handle the VMs running on the compute node.

Procedure performed by: Solution administrator

Execution: Web portal only

Compute node shutdown is performed on the Cloud Controller’s web portal in the compute node’s action menu (...).

If you shut down a compute node with VMs running, prior to shutdown the system provides options to either migrate the VMs to another compute node (according to the migration zone, categories, tags, and available resources), or shut down the VM.

For best results, the VM owners should perform a proper shutdown within the VMs based on their application-specific shutdown procedure.

You can find the VMs running in a specific compute node by going in the compute node from the Hardware page and clicking **View All** in the Application Instances section.

Migrating VMs from a compute node

This procedure explains how to migrate VMs running on a compute node and prevent the compute node from hosting new VMs.

Procedure performed by: Solution administrator; ThinkAgile Advantage Support (optional)

Execution: Web portal only

This procedure uses an in-product functionality called Maintenance Mode.

In a compute node action on the Hardware page in Cloud Controller, click **Turn on Maintenance Mode**. The system performs a sanity check and presents the VMs running on the compute node that need immediate action. Select **Migrate**. If a VM does not satisfy the requirements to move to another node (Migration Zone, Compute Categories and Tags, resources, etc.), migration is not possible and the restriction (tags) can be relaxed, or the VM must be shut down before being migrated to another node.

For best results and because of space limitations, VM owners should perform a proper shutdown within the VMs based on their application-specific shutdown procedures.

Migrate a VM between compute nodes

This section covers the procedure required to migrate a VM from one compute node to another compute node.

Execution: Web Portal only

Procedure performed by: Solution administrator, or remotely by ThinkAgile Advantage Support (optional)

The compute node has no direct capability to migrate a VM between nodes, but you can use compute tags for VM migration.

Compute tags allow the system to stick a VM based on compute tag requirements. If a VM has a tag, it can only run on a compute node with this tag. When tags are added to VMs, the system automatically migrates VMs to the proper compute node.

To move one VM to a specific compute node:

Step 1. Add a new compute tag to the destination compute node. Do not reuse an existing tag. The best practice is to either use a random number or something like “temp_migrate_host_1.”

Step 2. Add the same tag to the VM you are migrating.

Step 3. Wait until the VM is running on the destination compute node. You can find the VMs running in a specific compute node by going in the compute node from the Hardware page and clicking **View All** in the Application Instances section.

Step 4. Remove the tag from the migrated VM.

Important: This removal must be done before proceeding to the next step.

Step 5. Remove the tag from the compute node.

Replacing compute block hardware

Procedures for replacing the Lenovo ThinkAgile CP compute block hardware.

This section contains detailed procedures for replacing hardware parts in the ThinkAgile CP compute block. If you are instructed to return a part, follow all packaging instructions, and use any packaging materials supplied to you for shipping.

The types of replaceable parts include:

- **Structural parts:** Purchase and replacement of structural parts such as enclosure assembly, top cover, and bezel is your responsibility. If Lenovo acquires or installs a structural component at your request, you will be charged for the service.
- **Customer-replaceable unit (CRU):** Replacing CRUs is your responsibility. You can request a Lenovo field representative to install a CRU, but you will be charged for the service.
- **Field-replaceable unit (FRU):** You may install a FRU yourself, or request a Lenovo to install it at no additional charge while the appliance is under warranty.

Refer to [Chapter 2 “Roles and responsibilities” on page 15](#) to determine who can perform these procedures.

For information about the terms of the warranty, visit <https://datacentersupport.lenovo.com/us/en/warrantylookup>.

Hardware replacement procedures

This section provides installation and removal procedures for all serviceable system components. Each component replacement procedure references any tasks that need to be performed to gain access to the component being replaced.

The ThinkAgile CP appliances are based on the ThinkSystem SD530 Compute Node and ThinkSystem D2 Enclosure. For hardware replacement procedures not listed in this section, refer to the replacement instructions here:

http://thinksystem.lenovofiles.com/help/topic/7X21/maintenance_manual_hardware_replacement_procedures.html

For more information on ordering parts:

<https://datacentersupport.lenovo.com/us/en/products/solutions-and-software/thinkagile-vx/vx3720/7y12/parts>

Note: If you replace a part, such as an adapter, that contains firmware, you may also need to update the firmware for that part. Contact Lenovo Support for firmware updates.

Replace compute network cables

This section describes how to replace compute network cables.

Procedure performed by: Customer or field service

Replacing a cable that connects the compute enclosure to an interconnect can be done live, without any specific preparation steps.

When a cable is disconnected, the compute node associated with the port decreases in bandwidth by 50%.

Enable or Disable Maintenance Mode

This section covers the procedures to enable or disable maintenance mode for compute blocks.



Attention: Before engaging in any part replacement procedures involving a compute block, you must turn on Maintenance Mode for that block.

After finishing the part replacement procedure, turn off Maintenance Mode for the involved block.

Enable Maintenance Mode

This section covers the procedure to enable maintenance mode for compute blocks.

Complete the following steps to enable Maintenance Mode for a compute block:

- Step 1. Log into your ThinkAgile CP portal.
- Step 2. From the left menu select **Hardware**.
- Step 3. Navigate to the Compute Block section.
- Step 4. Click the **...** button next to the desired node and select **Turn on maintenance mode....**

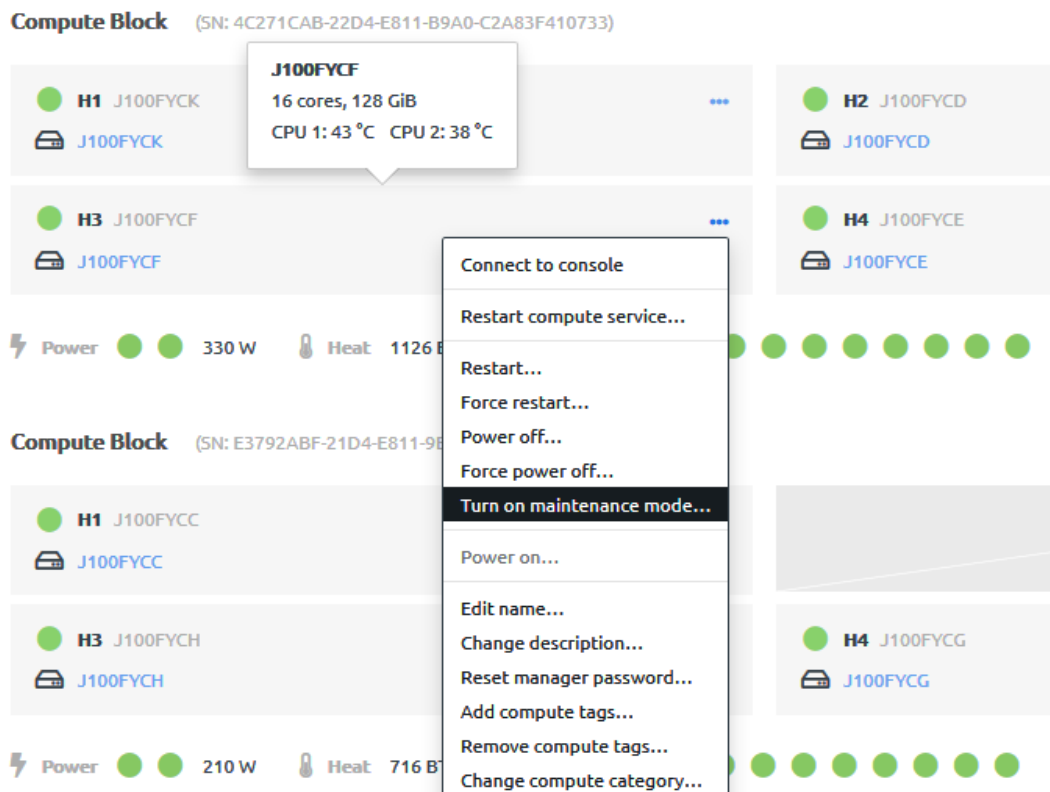


Figure 38. Enable Maintenance Mode for a Node

Step 5. If there are any application instances currently running on the compute node, you must select an action for each instance from the **Action Plan** drop-down menu and then click **Continue**:

- Migrate
- Remove constraints and migrate
- Shut down
- Force shutdown

Turn on Maintenance Mode

Are you sure you want to turn on maintenance mode for node J100FYCK?

Running Application Instances
Putting a node in maintenance mode will migrate all running application instances off the node. Please select an action plan for each instance.

Note: No new instances will be allowed to use this node as long as it is in maintenance mode.

Application Instance	Virtual Datacenter	Compute Constraints	Action Plan
RHEL 7.4 (Minimal) - Lenovo Tem...	ACME VDC 1	None	Migrate
TurnKey - Wordpress - Lenovo Te...	ACME VDC 2	None	Migrate
Windows 7 Pro SP1 (64-bit) - Len...	ACME VDC 2	None	Migrate

Cancel Continue

Figure 39. Maintenance Mode dialog box

Step 6. If the system encounters any errors when resolving the selected actions, a pop-up box appears with information about the errors.


Note: The errors must be resolved to turn on Maintenance Mode.

Maintenance Mode for the selected compute node is now enabled.

Disable Maintenance Mode

This section covers the procedure to disable maintenance mode for compute blocks.

Complete the following steps to disable Maintenance Mode for a compute block:

- Step 1. Log into your ThinkAgile CP portal.
- Step 2. From the left menu select **Hardware**.
- Step 3. Navigate to the Compute Block section.
- Step 4. Click the  button next to the desired node and select **Turn off maintenance mode...**

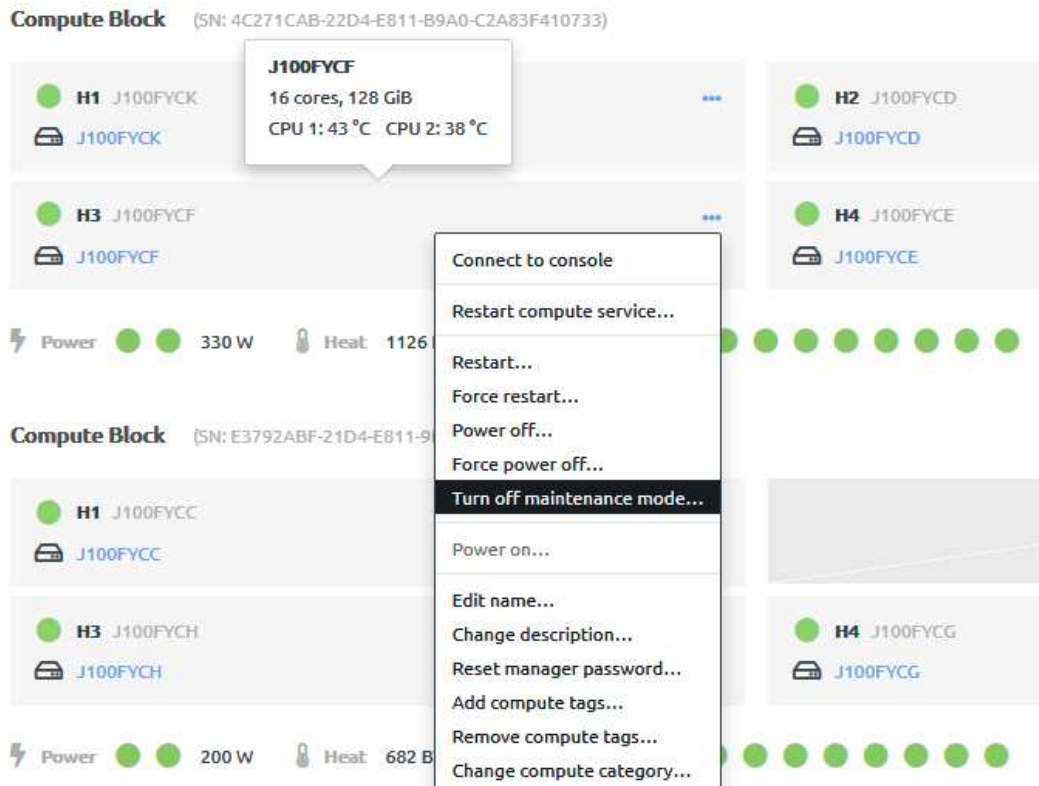


Figure 40. Disable Maintenance Mode for a Node

Maintenance Mode for the compute node is now disabled.

Hot-swap drive replacement

Follow these instructions to replace a front SSD in the ThinkAgile CP compute block appliance.


Each ThinkAgile CP compute block appliance supports Intel 240GB solid-state drives (SSDs). These drives are used for installing the operating system and no user data are stored on these drives. These drives are used in RAID 1 mode.

Always replace failed data drives as soon as possible. The following indicate a failed front SSD:

- On cloud controller or portal
- The yellow status LED on the node front panel is illuminated.

Remove a hot-swap drive

Use this procedure to remove a drive.



Before you remove a drive, ensure that you save the data on your drive (especially if it is part of a RAID array) before you remove it from the node.

Procedure performed by: customer or field service

Attention:

- To ensure that there is adequate system cooling, do not operate the solution for more than 2 minutes without either a drive or a filler panel installed in each bay.
- Before you make changes to drives, drive controllers (including controllers integrated on the system board), drive backplanes, or drive cables, back up all important data stored on drives.
- Before you remove any RAID array components, back up all RAID configuration information.
- If one or more NVMe solid-state drives are to be removed, disable them in OS before removal.

Complete the following steps to remove a drive:

Step 1. Slide the latch to unlock the drive handle.

Step 2. Grasp the handle and slide the drive out of the drive bay.




Figure 41. Drive removal


The removal of the compute hard drive is now complete.

After you remove the drive:

- Install either a drive filler or a replacement drive to cover the drive bay. See [“Install a hot-swap drive” on page 92](#)
- If you are instructed to return the component or optional device, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Install a hot-swap drive

Use this procedure to install a drive.



Before installing a hot-swap drive:

1. Remove the drive filler from the empty drive bay. Keep the filler panel in a safe place.
2. Touch the static-protective package that contains the drive to any unpainted metal surface on the solution. Then, remove the drive from the package and place it on a static-protective surface.

Procedure performed by: customer or field service

The following notes describe the drive that the node supports, and other information that you must consider when you install a drive.

- Supported drive: ThinkSystem 2.5" Intel S4500 240GB Entry SATA 6GB Hot-swap SSD.

- Locate the documentation that comes with the drive and follow those instructions in addition to the instructions in this chapter.
- You can install up to two hot-swap SATA 2.5-inch drives for each node.
- The electromagnetic interference (EMI) integrity and cooling of the solution are protected by having all bays and PCI and PCI Express slots covered or occupied. When you install a drive, PCI, or PCI Express adapter, save the EMC shield and filler panel from the bay or PCI or PCI Express adapter slot cover in the event that you later remove the device.
- For a complete list of supported optional devices for the node, see the [Lenovo ServerProven website](#).

Complete the following steps to install a drive:

Note: If you have only one drive, you must install it in the bay 0 (upper-left).

- Step 1. Ensure that the tray handle is in the open (unlocked) position.
- Step 2. Align the drive with the guide rails in the bay.
- Step 3. Gently push the drive into the bay until the drive stops.
- Step 4. Rotate the tray handle to the closed (locked) position until you hear a click.




Figure 42. Drive installation

- Step 5. Check the drive status LED to verify that the drive is operating correctly.
 - If the yellow drive status LED is lit continuously, that drive is faulty and must be replaced.
 - If the green drive activity LED is flashing, the drive is being accessed.
- Step 6. Repeat steps [Step 1 1 on page 93](#) to [Step 5 5 on page 93](#) if installing multiple drives.

The installation of a hot-swap drive is now complete.

After you install all drives:

- If the node is configured for RAID 1 operation, you may need to re-configure your disk arrays after you install drives. Software RAID is used in the solution.
- If you are instructed to return the component or optional device, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.



Hot-swap power supply replacement

Use this information to remove and install a hot-swap power supply.

Note: Use any documentation that comes with the power supply and follow those instructions in addition to the instructions in these topics. Ensure that you have all the cables and other equipment that come with the power supply, as specified in its documentation.

Remove a hot-swap power supply

Use this information to remove a hot-swap power supply.

 <p>Chapter 3 “Read the Installation Guidelines” on page 19</p>	 <p>“ATTENTION: Static Sensitive Device Ground package before opening” on page 21</p>
---	---

To avoid possible danger, read and follow the following safety statement.

S001



 **DANGER**

Electrical current from power, telephone, and communication cables is hazardous.

To avoid a shock hazard:

- **Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.**
- **Connect all power cords to a properly wired and grounded electrical outlet.**
- **Connect to properly wired outlets any equipment that will be attached to this product.**
- **When possible, use one hand only to connect or disconnect signal cables.**
- **Never turn on any equipment when there is evidence of fire, water, or structural damage.**
- **Disconnect the attached power cords, telecommunications systems, networks, and modems before you open the device covers, unless instructed otherwise in the installation and configuration procedures.**
- **Connect and disconnect cables as described in the following table when installing, moving, or opening covers on this product or attached devices.**

To Connect:

1. Turn everything OFF.
2. First, attach all cables to devices.
3. Attach signal cables to connectors.
4. Attach power cords to outlet.
5. Turn device ON.

To Disconnect:

1. Turn everything OFF.
2. First, remove power cords from outlet.
3. Remove signal cables from connectors.
4. Remove all cables from devices.

S002



CAUTION:

The power control button on the device and the power switch on the power supply do not turn off the electrical current that is supplied to the device. The device also might have more than one power cord. To remove all electrical current from the device, ensure that all power cords are disconnected from the power source.

S019



CAUTION:

The power-control button on the device does not turn off the electrical current supplied to the device. The device also might have more than one connection to dc power. To remove all electrical current from the device, ensure that all connections to dc power are disconnected at the dc power input terminals.

S035



CAUTION:

Never remove the cover on a power supply or any part that has this label attached.

Hazardous voltage, current, and energy levels are present inside any component that has this label attached. There are no serviceable parts inside these components. If you suspect a problem with one of these parts, contact a service technician.

Before you remove a hot-swap power supply:

1. Disconnect the power cords and all external cables.
2. Identify the failed power supply. ThinkAgile CP Compute Block appliances contain two power supplies.

Note: If the amber warning LED on the node front panel is illuminated, then a power supply unit has failed.

Procedure performed by: customer or field service

Complete the following steps to remove a hot-swap power supply:

- Step 1. Disconnect the power cord from the connector on the back of the power supply.
- Step 2. Press and hold the orange release tab to the left.
- Step 3. Grasp the handle and pull the power supply out of the bay.




Figure 43. Hot-swap power supply removal



The removal of the hot-swap power supply is now complete.

After you remove the power supply:

- Install a new power supply unit. See [“Install a hot-swap power supply” on page 97](#)
- If you are instructed to return the component or optional device, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Install a hot-swap power supply

Use this information to install a hot-swap power supply.

 <p>Chapter 3 “Read the Installation Guidelines” on page 19</p>	 <p>“ATTENTION: Static Sensitive Device Ground package before opening” on page 21</p>
---	---

To avoid possible danger, read and follow the following safety statement.

S001



Electrical current from power, telephone, and communication cables is hazardous.
To avoid a shock hazard:

- **Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.**
- **Connect all power cords to a properly wired and grounded electrical outlet.**
- **Connect to properly wired outlets any equipment that will be attached to this product.**
- **When possible, use one hand only to connect or disconnect signal cables.**
- **Never turn on any equipment when there is evidence of fire, water, or structural damage.**
- **Disconnect the attached power cords, telecommunications systems, networks, and modems before you open the device covers, unless instructed otherwise in the installation and configuration procedures.**
- **Connect and disconnect cables as described in the following table when installing, moving, or opening covers on this product or attached devices.**


To Connect:

1. Turn everything OFF.
2. First, attach all cables to devices.
3. Attach signal cables to connectors.
4. Attach power cords to outlet.
5. Turn device ON.

To Disconnect:

1. Turn everything OFF.
2. First, remove power cords from outlet.
3. Remove signal cables from connectors.
4. Remove all cables from devices.

S035



CAUTION:

Never remove the cover on a power supply or any part that has this label attached.

Hazardous voltage, current, and energy levels are present inside any component that has this label attached. There are no serviceable parts inside these components. If you suspect a problem with one of these parts, contact a service technician.

Before you install a hot-swap power supply:

Notes:

1. Ensure that the devices you are installing are supported. For a list of supported optional devices for the solution, see <http://www.lenovo.com/us/en/serverproven/>.
2. Do not install two power supply units with different wattages. Related information is available from the following:
 - Read the label on the top cover for maximum wattage output of installed power supply units. Only replace the existing units with those having the same wattage as marked on the label.
 - Check the rear of the node to ensure that there is no length difference between the two installed units. A visible difference in length means that the two units have different wattages, and one of them must be replaced.

Procedure performed by: customer or field service

Complete the following steps to install a hot-swap power supply.

Step 1. Slide the hot-swap power supply into the bay until the release latch clicks into place.

Important: During normal operation, each power-supply bay must contain either a power supply or power-supply filler panel for proper cooling.

Step 2. Connect one end of the power cord for the new power supply into the AC connector on the back of the power supply. Then, connect the other end of the power cord into a properly grounded electrical outlet.




Figure 44. Hot-swap power supply installation

Note: Connect the power cord to the power supply unit. Ensure that the cord is properly connected to the power supply.

The installation of a hot-swap power supply is now complete.

After installing the hot-swap power supply:

1. Reconnect any cables that you removed.
2. If the node is turned off, turn it on and ensure that the power supply is replaced successfully:
 - The amber warning LED on node front panel is off.
 - The green indicator light on each power supply is illuminated.
3. If you are instructed to return the component or optional device, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Fan replacement


Use the following procedures to remove and install a compute enclosure fan.

Failed or failing system fans can cause the system to overheat and shut down. Replace a failed system fan as soon as possible.

Attention: If the amber warning LED on the node front panel is illuminated, this means that a enclosure fan has failed.

Remove a fan

Use this information to remove a fan from the compute enclosure.



To avoid possible danger, read and follow the following safety statement.

S017



CAUTION:
Hazardous moving fan blades nearby.

S033



CAUTION:
Hazardous energy present. Voltages with hazardous energy might cause heating when shorted with metal, which might result in spattered metal, burns, or both.

Before removing a fan:

1. Remove the fan cover (see http://thinksystem.lenovofiles.com/help/topic/7X21/remove_the_fan_cover.html).

Attention:

- Static electricity released into internal solution components when the solution is powered on may cause the solution to halt, which may result in data loss. To avoid this problem, always use an electrostatic-discharge wrist strap or other grounding system when you work inside the solution with the power on.
 - Replace the failed fan and reinstall the fan cover in three minutes.
2. Identify the failed fan. The following illustration shows locations of fan fault LEDs. When the LED is lit, it indicates the fan is having operational issues.




Figure 45. Fan fault LEDs

Table 3. Fan fault LEDs

1 80x80x80mm fan fault LEDs
2 60x60x56mm fan fault LEDs

Remove a 60x60x56 mm Fan

Procedure for removing a 60x60x56 mm fan from the compute enclosure.

Procedure performed by: customer or field service

Complete the following steps to remove a fan:

To remove a 60x60x56 mm fan, follow these steps:

Step 1. Squeeze the fan release latches together.

Step 2. Lift the fan out of the enclosure.




Figure 46. 60x60x56mm fan removal

The removal of the fan is now complete.

After the removal of the fan:

- Install a new fan. See [“Install a fan” on page 104](#)
- If you are instructed to return the component or optional device, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Remove a 80x80x80 mm Fan

Procedure for removing a 80x80x80 mm fan from the compute enclosure.

Procedure performed by: customer or field service

Complete the following steps to remove a fan:

To remove a 80x80x80 mm fan, follow these steps:

- Step 1. Carefully pull the cable out from underneath the sheet metal flange.
- Step 2. Disconnect the cable.
- Step 3. Grasp the fan and lift it out of the enclosure.




Figure 47. 80x80x80mm fan removal


The removal of the fan is now complete.

After the removal of the fan:

- Install a new fan. See [“Install a fan” on page 104](#)
- If you are instructed to return the component or optional device, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Install a fan

Use this information to install a fan in the compute enclosure.



To avoid possible danger, read and follow the following safety statement.

- **S017**



CAUTION:
Hazardous moving fan blades nearby.

- **S033**



CAUTION:
Hazardous energy present. Voltages with hazardous energy might cause heating when shorted with metal, which might result in spattered metal, burns, or both.

Before you install a fan, remove the fan cover (see http://thinksystem.lenovofiles.com/help/topic/7X21/remove_the_fan_cover.html).

Attention:

- Static electricity released into internal solution components when the solution is powered on may cause the solution to halt, which may result in data loss. To avoid this problem, always use an electrostatic-discharge wrist strap or other grounding system when you work inside the solution with the power on.
- Replace the failed fan and reinstall the fan cover in three minutes.

Install a 60x60x56 mm Fan

Use this information to install a 60x60x56 mm fan in the compute enclosure.

Procedure performed by: customer or field service

Complete the following steps to install a fan.

To install a 60x60x56mm fan, follow these steps:

- Step 1. Lower the fan into the socket.
- Step 2. Push it downward until it clicks into place.




Figure 48. 60x60x56mm fan installation

The installation of the fan is now complete.

After installing a fan, complete the following steps.

1. Reinstall the fan cover (see http://thinksystem.lenovofiles.com/help/topic/7X21/install_the_fan_cover.html).
2. Verify that the system fan is replaced successfully:
 - The error LED for the replaced system fan is off.
 - The amber warning LED on the node front panel is off.
3. If you are instructed to return the component or optional device, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Install a 80x80x80 mm Fan

Use this information to install a 80x80x80 mm fan in the compute enclosure.

Procedure performed by: customer or field service

Complete the following steps to install a fan.

To install a 80x80x80 mm fan, follow these steps:

- Step 1. Lower the fan into the socket, and push it downward until it clicks into place.
- Step 2. Connect the power cable.
- Step 3. Carefully route the cable underneath the flange and ensure that the cable is routed through the notch.




Figure 49. 80x80x80mm fan installation

Note: Ensure that the cable is routed in the proper location and that no wire is stuck in the flange.

The installation of the fan is now complete.

After installing a fan, complete the following steps.

1. Reinstall the fan cover (see http://thinksystem.lenovofiles.com/help/topic/7X21/install_the_fan_cover.html).
2. Verify that the system fan is replaced successfully:
 - The error LED for the replaced system fan is off.
 - The amber warning LED on the node front panel is off.
3. If you are instructed to return the component or optional device, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

EIOM replacement

Use the following procedures to remove and install the Ethernet I/O module (EIOM).

The ThinkAgile CP Compute Block appliances offer connectivity to the LOM (LAN on motherboard) of each node via ports on the EIOM located at the rear of the enclosure. The top Phy maps to port one of the system LOM, the bottom Phy maps to port two of the system LOM.

The ports are labeled sequentially from left to right.





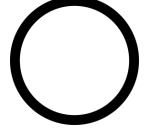

Figure 50. EIOM ports and corresponding node number

Table 4. EIOM port corresponding node number

1 Ports to node 1
2 Ports to node 2
3 Ports to node 3
4 Ports to node 4

Remove the EIOM

Use this information to remove the EIOM.

 <p>Chapter 3 “Read the Installation Guidelines” on page 19</p>	 <p>“Power off the compute node” on page 119</p>	 <p>“ATTENTION: Static Sensitive Device Ground package before opening” on page 21</p>
--	---	---

Before you remove the EIOM:

1. Place the node in maintenance mode.
2. Power off all compute nodes and peripheral devices.
3. Disengage all the compute nodes from the enclosure (see http://thinksystem.lenovofiles.com/help/topic/7X21/remove_a_compute_node_from_the_chassis.html).
4. Disconnect the power cords and all external cables from the rear of the enclosure.
5. If the cable management arm is installed, remove it (see http://thinksystem.lenovofiles.com/help/topic/7X21/remove_the_cable_management_arm.html).
6. Remove the shuttle (see http://thinksystem.lenovofiles.com/help/topic/7X21/remove_the_shuttle.html) and place it on a stable work surface.

Procedure performed by: customer or field service

Complete the following steps to remove the 10GbE cage (SFP+) model EIOM:

Step 1. Disconnect two cables from the EIOM.

Note: Ensure that you push the release latch only when disconnecting the signal cable.

Step 2. Turn the thumbscrews counterclockwise.

Step 3. Grasp and push the EIOM slightly towards the front side of the shuttle, and then lift the EIOM up to remove the EIOM from the shuttle.




Figure 51. EIOM removal


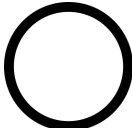
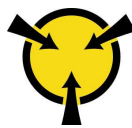
The removal of the EIOM is now complete.

After the EIOM removal:

- Install a new EIOM. See [“Install the EIOM” on page 108](#)
- If you are instructed to return the component or optional device, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Install the EIOM

Use this information to install the EIOM.

 <p>Chapter 3 “Read the Installation Guidelines” on page 19</p>	 <p>“Power off the compute node” on page 119</p>	 <p>“ATTENTION: Static Sensitive Device Ground package before opening” on page 21</p>
---	--	--

Before you install the EIOM, touch the static-protective package that contains the new EIOM to any unpainted surface on the outside of the server. Then, take the new EIOM out of the package and place it on a static-protective surface.

Note: The minimum networking speed requirement for the EIOM is 1Gbps.

Procedure performed by: customer or field service

Complete the following steps to install the 10GbE cage (SFP+) model EIOM.

- Step 1. Grasp the EIOM and align the four EIOM tabs with the slots in the shuttle. Then, lower the EIOM into the slots, and then pull the EIOM slightly towards the rear side of the shuttle.
- Step 2. Connect required cables to the EIOM.
- Step 3. Turn the thumbscrews clockwise.




Figure 52. EIOM installation

The installation of the EIOM is now complete.

After you install the EIOM, complete the following steps:

1. Reinstall the shuttle (see http://thinksystem.lenovofiles.com/help/topic/7X21/install_the_shuttle.html).
2. If the cable management arm is removed, install it (see http://thinksystem.lenovofiles.com/help/topic/7X21/install_the_cable_management_arm.html).
3. Reconnect the power cords and any cables that you removed.
4. Push all compute nodes back into the enclosure (see http://thinksystem.lenovofiles.com/help/topic/7X21/install_a_compute_node_in_the_chassis.html).
5. Turn on all compute nodes.
6. Exit maintenance mode.
7. Check the following to verify whether the EOIM has been successfully replaced:
 - The amber warning light on the node front panel is off.
 - All the network ports on the EIOM are fully functional.
8. If you are instructed to return the component or optional device, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

DIMM replacement

Use the following procedures to remove and install a DIMM.

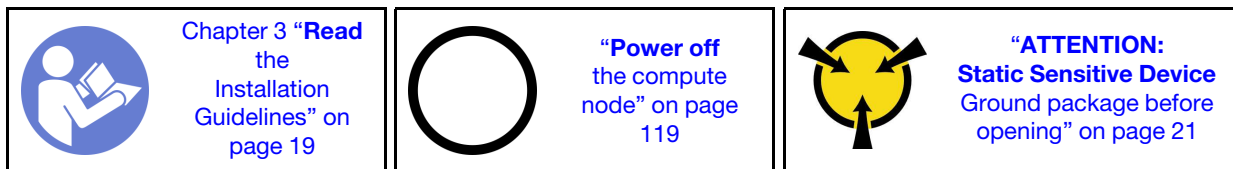
A node might be able to self-correct for certain memory errors. However, failed memory can lead to system degradation and should be replaced quickly.

Indications of a failed DIMM are:

- A post error on boot.
- An error in the Cloud Controller interface.
- Not all memory is detected. For example, the appliance should have 256 GB per node and the host only shows 240 GB

Remove DIMM

Use this information to remove a memory module.



Attention: Memory modules are sensitive to static discharge and require special handling. In addition to the standard guidelines for ["Handling static-sensitive devices" on page 21](#):

- Always wear an electrostatic-discharge strap when removing or installing memory modules. Electrostatic-discharge gloves can also be used.
- Never hold two or more memory modules together so that they touch. Do not stack memory modules directly on top of each other during storage.
- Never touch the gold memory module connector contacts or allow these contacts to touch the outside of the memory-module connector housing.
- Handle memory modules with care. Never bend, twist, or drop a memory module.

Before you remove a DIMM:

1. Identify the failed DIMM by either checking the event logs in UEFI setup or use the Cloud Controller to browse the post event log. The event log contains information about the location of the DIMM fault.
2. Place the node in maintenance mode.
3. Power off the node and if the cable management arm is not in place, disconnect all the cables.
4. Remove the compute node from the enclosure (see http://thinksystem.lenovofiles.com/help/topic/7X21/remove_a_compute_node_from_the_chassis.html).
5. Remove the compute node cover (see http://thinksystem.lenovofiles.com/help/topic/7X21/remove_the_compute_node_cover.html).
6. Remove the air baffle (see http://thinksystem.lenovofiles.com/help/topic/7X21/remove_the_air_baffle.html).

Procedure performed by: Customer or field service

Complete the following steps to remove a DIMM:

Step 1. Carefully press down the retaining clips on each end of the DIMM connector.

Attention: To avoid breaking the retaining clips or damaging the DIMM connectors, open and close the clips gently.

Step 2. Lift the DIMM from its slot.




Figure 53. DIMM removal

The removal of the DIMM is now complete.

After removal of the DIMM:

- Install a new DIMM. See [“Memory module installation” on page 112](#)
- If you are instructed to return the component or optional device, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Memory module installation

The following notes describe the types of DIMMs that the node supports and other considerations when you install DIMMs.

- Confirm that the node supports the DIMM that you are installing (see <http://www.lenovo.com/us/en/serverproven/>).
- When you install or remove DIMMs, the node configuration information changes. When you restart the node, the system displays a message that indicates that the memory configuration has changed. You can use the Setup utility to view the node configuration information.
- Install higher-capacity (ranked) DIMMs first, following the population sequence for the memory mode being used.
- The node supports only industry-standard double-data-rate 4 (DDR4), 2666 MT/s, PC4-21300 (single-rank or dual-rank), unbuffered or synchronous dynamic random-access memory (SDRAM) dual inline memory modules (DIMMs) with error correcting code (ECC).
- Do not mix RDIMMs, LRDIMMs and 3DS DIMMs in the same node.
- The maximum operating speed of the node is determined by the slowest DIMM in the node.
- If you install a pair of DIMMs in DIMM connectors 1 and 3, the size and speed of the DIMMs that you install in DIMM connectors 1 and 3 must match each other. However, they do not have to be the same size and speed as the DIMMs that are installed in DIMM connectors 2 and 4.
- You can use compatible DIMMs from various manufacturers in the same pair.
- The specifications of a DDR4 DIMM are on a label on the DIMM, in the following format.
- *gggGBpheRxff PC4-wwwwaa-mccd-bb*

where:

- *gggGB* is the total capacity, in gigabytes, for primary bus (ECC not counted) 4GB, 8GB, 16GB, etc. (no space between digits and units)
- *pheR* is the number of package ranks of memory installed and number of logical ranks per package rank
 - *p* =
 - 1 = 1 package rank of SDRAMs installed
 - 2 = 2 package ranks of SDRAMs installed
 - 3 = 3 package ranks of SDRAMs installed
 - 4 = 4 package ranks of SDRAMs installed
 - *he* = blank for monolithic DRAMs, else for modules using stacked DRAM:
 - *h* = DRAM package type
 - D = multi-load DRAM stacking (DDP)
 - Q = multi-load DRAM stacking (QDP)
 - S = single load DRAM stacking (3DS)
 - *e* = blank for SDP, DDP and QDP, else modules using 3DS stacks, logical ranks per package rank
 - 2 = 2 logical ranks in each package rank
 - 4 = 4 logical ranks in each package rank
 - 8 = 8 logical ranks in each package rank

- R = rank(s)
- xff = Device organization (data bit width) of SDRAMs used on this assembly
 - x4 = x4 organization (4 DQ lines per SDRAM)
 - x8 = x8 organization
 - x16 = x16 organization
- $wwwww$ is the DIMM bandwidth, in MBps: 2133, 2400, 2666, 2933, 3200
- aa is the SDRAM speed grade
- m is the DIMM type
 - E = Unbuffered DIMM (UDIMM), x64 primary + 8 bit ECC module data bus
 - L = Load Reduced DIMM (LRDIMM), x64 primary + 8 bit ECC module data bus
 - R = Registered DIMM (RDIMM), x64 primary + 8 bit ECC module data bus
 - U = Unbuffered DIMM (UDIMM) with no ECC (x64-bit primary data bus)
- cc is the reference design file used for this design
- d is the revision number of the reference design used
- bb is the JEDEC SPD Revision Encoding and Additions level used on this DIMM

The following illustration shows the location of the DIMM connectors on the system board.




Figure 54. Location of the DIMM connectors on the system board

Installation order

Memory modules must be installed in a specific order based on the memory configuration that you implement on your node.

The following memory configurations are available:

Independent memory mode population sequence

Table 5. DIMM installation sequence (Independent mode/normal mode)

Number of processor	Installation sequence (connectors)
Processor 1 installed	6, 3, 7, 2, 8, 1, 5, 4
Processor 1 and 2 installed	6, 14, 3, 11, 7, 15, 2, 10, 8, 16, 1, 9, 5, 13, 4, 12

Memory mirroring population sequence

Table 6. DIMM installation sequence (mirror mode/lockstep mode)

Number of processors	Installation sequence (connectors)
Processor 1 installed	(6, 7), (2, 3), (8, 1)
Processor 1 and 2 installed	(6, 7, 14, 15), (2, 3), (10, 11), (1,8), (9, 16)

If you are installing 3, 6, 9 or 12 identical DIMMs for the mirroring mode, comply with the following installation sequence to achieve the best performance.

Table 7. DIMM installation sequence (mirror mode/lockstep mode for 3, 6, 9 and 12 identical DIMMs)

Number of processors	Installation sequence (connectors)
Processor 1 installed	(6, 7, 8), (1, 2, 3)
Processor 1 and 2 installed	(6, 7, 8), (14, 15, 16), (1, 2, 3), (9, 10, 11)

Memory rank sparing population sequence

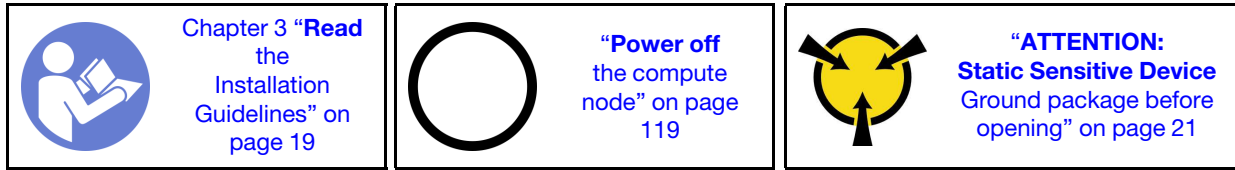
Table 8. DIMM installation sequence (sparing mode)

Number of processors	Installation sequence (connectors)
Processor 1 installed	6, 3, 7, 2, 8, 1, 5, 4
Processor 1 and 2 installed	6, 14, 3, 11, 7, 15, 2, 10, 8, 16, 1, 9, 5, 13, 4, 12

Note: Single-rank RDIMM is not supported by sparing. If you install single-rank RDIMM, it switches to the independent mode automatically.

Install DIMM

Use this information to install a DIMM.



Attention: Memory modules are sensitive to static discharge and require special handling. In addition to the standard guidelines for “[Handling static-sensitive devices](#)” on page 21:

- Always wear an electrostatic-discharge strap when removing or installing memory modules. Electrostatic-discharge gloves can also be used.
- Never hold two or more memory modules together so that they touch. Do not stack memory modules directly on top of each other during storage.
- Never touch the gold memory module connector contacts or allow these contacts to touch the outside of the memory-module connector housing.
- Handle memory modules with care: never bend, twist, or drop a memory module.

Procedure performed by: customer or field service

The following illustration shows the location of the DIMM connectors on the system board. The following illustration shows the location of the DIMM connectors on the system board.




Figure 55. Location of the DIMM connectors on the system board

Complete the following steps to install a DIMM:

Important: Before installing a memory module, make sure that you understand the required installation order, depending on whether you are implementing memory mirroring, memory rank sparing, or independent memory mode. See [“Installation order” on page 114](#) for the required installation order.

Step 1. Open the retaining clip on each end of the DIMM connector.

Attention:

- DIMMs are static-sensitive devices. The package must be grounded before it is opened.
- To avoid breaking the retaining clips or damaging the DIMM connectors, open and close the clips gently.

Step 2. Touch the static-protective package that contains the DIMM to any unpainted metal surface on the outside of the node. Then, remove the DIMM from the package.

Step 3. Turn the DIMM so that the alignment slot align correctly with the alignment tab.

Step 4. Insert the DIMM into the connector by aligning the edges of the DIMM with the slots at the ends of the DIMM connector.

Step 5. Firmly press the DIMM straight down into the connector by applying pressure on both ends of the DIMM simultaneously. The retaining clips snap into the locked position when the DIMM is firmly seated in the connector.




Figure 56. DIMM installation

Note: If there is a gap between the DIMM and the retaining clips, the DIMM has not been correctly inserted; open the retaining clips, remove the DIMM, and then reinsert it.

Step 6. Reconnect any cable that you removed.

The installation of the DIMM is now complete.

After you install a DIMM, complete the following steps:

1. Reinstall the air baffle (see http://thinksystem.lenovofiles.com/help/topic/7X21/install_the_air_baffle.html).
2. Reinstall the compute node cover (see http://thinksystem.lenovofiles.com/help/topic/7X21/install_the_compute_node_cover.html).
3. Reinstall the compute node (see http://thinksystem.lenovofiles.com/help/topic/7X21/install_a_compute_node_in_the_chassis.html).
4. Check the power LED to make sure it transitions between fast blink and slow blink to indicate the node is ready to be powered on.
5. Power on the node.
6. Exit maintenance mode.
7. Verify that the DIMM memory failure is resolved by reloading the system. If successful, no POST errors appear during the booting process.
8. If you are instructed to return the component or optional device, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Compute node replacement

Use the following procedures to remove and install the compute node from the ThinkSystem D2 Enclosure.

Remove a compute node

This procedure describes how to remove a compute node from your ThinkAgile CP configuration.

Procedure performed by: Solution administrator, ThinkAgile Advantage Support

To remove a compute node from your ThinkAgile CP configuration, perform the following steps:

Step 1. Obtain the serial number from the web portal's Hardware page.

Note: This step requires administrator-level access.

Step 2. Shut down the compute node. See [“Shutting off a compute node” on page 85](#).

Step 3. Physically remove the compute node.

Step 4. Sign in to the primary interconnect from Support Mode (see [Chapter 4 “Managing Support Mode” on page 23](#)).

The username is admin and the password is the one set during the installation.

Step 5. Run the following command on the primary interconnect:

```
> sudo service tacp-network-controller restart
```

Step 6. Contact Lenovo Support and request a hardware removal from the database by providing the serial number and the region the hardware is (US or EU).

Install a compute node

This procedure describes how to add a compute node to your ThinkAgile CP configuration.

Procedure performed by: Customer or field service

To add a compute node to your ThinkAgile CP configuration, perform the following steps:

Step 1. Physically install the compute node in the compute enclosure and power it on.

Step 2. Launch the Cloud Controller and register the compute nodes. All compute and storage hardware assets connected to the network interconnect(s) are displayed on the Hardware page in the Cloud Controller and are ready for registration.

Step 3. In the Cloud Controller dashboard, select **Hardware** from the left menu.

Step 4. On the node that you want to register, click **Register**. The Register Node page displays.

Step 5. In the **Node Name** field, enter a unique name for the compute node.

Step 6. In the **Administrator Password** field, enter a password for the compute node's manager account.

Step 7. In the **Compute Category** drop-down, select a compute category from the list.

A compute category is a descriptor of nodes. Each node must belong to exactly one compute category.

Note: If you do **not** intend to create multiple compute categories, you should use the provided **Default** category for all nodes.

Step 8. (Optional) In the **Compute Tags** field, specify compute tags.

Compute tags are additional information that can be applied to nodes. While each node must belong to exactly one category, a node may have any number of tags. The Cloud Controller uses tags to match application resource constraints and available nodes. An application instance with compute tag constraints can only run on nodes that have all those compute tags.

Step 9. In the **Migration Zone** drop-down, assign the node to a migration zone from the list.

In the Cloud Controller, compute resources (nodes) are organized into migration zones. A node must be added to a migration zone for it to host application instances. Migration zones act as boundaries around a set of nodes. Applications can migrate among nodes within a migration zone, but cannot migrate to nodes outside of the migration zone.

To create a migration zone:

- a. In the **Migration Zone** drop-down, select **New migration zone....**
- b. Enter a name for the migration zone.
- c. Click **Create**.

Step 10. Click **Register Node** to complete registration.

Step 11. Repeat this process to register all remaining nodes.

Power off the compute node

The compute node remains in the standby state when connected to a power source, while Cloud Controller is allowed to respond to remote power-on requests. To completely power off the compute node (power status LED off), you must disconnect all power cables.

To power off the compute node that is in a standby state (power status LED flashes once per second):

- Start an orderly shutdown using the operating system (if supported by your operating system), or
- Press the power button to start an orderly shutdown (if supported by your operating system), or
- Press and hold the power button for more than 4 seconds to force a shutdown.

Note: The Cloud Controller can place the compute node in a standby state as an automatic response to a critical system failure.

In the standby state, the compute node can respond to remote power-on requests sent to the Cloud Controller. For information about powering on the compute node, see [“Power on the compute node” on page 119](#).

Power on the compute node

After the compute node performs a short self-test (power status LED flashes rapidly) when connected to input power, it enters the standby state (power status LED flashes once per second).

A compute node can be turned on (power LED on) in any of the following ways:

- You can press the power button, or
- The compute node can restart automatically after a power interruption, or
- The compute node can respond to remote power-on requests sent to the Cloud Controller.

Replace the SD530 system board

Follow this procedure to replace the SD530 system board.

Procedure performed by: Solution administrator

To replace the SD530 system board from the ThinkAgile CP, perform these steps.

- Step 1. Get the serial number and (optional) UUID information from Lenovo Cloud Controller.
- Step 2. Remove the compute node from the D2 enclosure.
- Step 3. Sign in to the primary interconnect through Support Mode or, if at a data center, using a method provided by the manufacturer. Sign in with user name **admin** and the password set during installation.
- Step 4. Run the following command on the primary interconnect: `sudo service tacp-network-controller restart`.
- Step 5. Follow the procedure at http://thinksystem.lenovofiles.com/help/index.jsp?topic=%2F7X21%2Fmaintenance_manual_hardware_replacement_procedures.html
- Step 6. Insert the compute node back into the D2 enclosure.
- Step 7. Wait 10 minutes and verify that the compute node is visible again in the web portal.

Replace the compute enclosure

This procedure will explain how to replace a compute enclosure from the system, assuming a compute node has been removed before.

Procedure performed by: Solution administrator, ThinkAgile Advantage Support

- Step 1. Record the serial number of the removed compute enclosure.
- Step 2. [“Remove a compute node” on page 118](#) for every compute node, except step 6; there’s no need for a hardware removal from the database for the compute nodes.
- Step 3. [“Add a compute node” on page 118](#) for every compute node in the new compute enclosure.
- Step 4. Contact ThinkAgile Advantage, who then contacts Lenovo Support to request a hardware removal from the database by providing the compute enclosure serial number.
- Step 5. Restore the enclosure to the rack.

Interconnect switch maintenance procedures

This section covers the ThinkAgile CP interconnect switch maintenance procedures.

Refer to [Chapter 2 “Roles and responsibilities” on page 15](#) to determine who can perform these procedures.

Power off the interconnect switch in an SS-V3 environment

This section covers the procedure to power off the interconnect switch in an SS-V3 environment.

Procedure performed by: Solution administrator

Execution: On hardware only (with exceptions for procedures dependencies)

Since inter-component communication is not possible while the interconnect is not operational, all hardware components and VMs must be shut down.

To power off the interconnect switch in an SS-V3 environment, follow these steps:

- Step 1. Sign in as an infrastructure administration into the Cloud Controller portal.
- Step 2. Navigate to **Infrastructure → Hardware**.
- Step 3. Select the stack in the top tab. It lists all hardware connected to the interconnect switch.
- Step 4. Perform the [“Shutting down the storage block” on page 35](#) procedure. This shuts off the storage block in the stack that you have selected, and will ask to shut down the VMs as well.
- Step 5. Perform the [“Shutting off a compute node” on page 85](#) procedure to shut off the compute node in the selected stack.
- Step 6. Sign in to the primary interconnect via Support Mode or, if at the data center, with a method provided by the manufacturer. The username is `admin` and the password is the one you set during the installation
- Step 7. Execute the following command:

```
> sudo shutdown -h now
```

The interconnect switch can now be unplugged safely.

Make the secondary interconnect switch the primary interconnect

Procedure for making the secondary ThinkAgile CP interconnect switch into the primary interconnect.

Attention: You have between 6 and 12 hours to fix the primary interconnect before you notice network connectivity problems, as that is the default lease time of the DHCP server. If you cannot replace the interconnect or DHCP server within 6 to 12 hours, you should attempt to make existing DHCP users static IPs.

- If the primary interconnect is still accessible, back up `/etc/dhcp/dhcp-failover.conf`.
- If the primary interconnect is still accessible, back up `/var/lib/dhcp/dhcpd.leases` and hardware management network settings in `/etc/dhcp/vlan-mgmt.conf` and `/etc/dhcp/vlan-main.conf`.
- Obtain the hardware management network settings (`BMC_DHCP_START`, `BMC_DHCP_END`, `BMC_ROUTER_IP`, `BMC_PREFIX`) from the primary network interconnect (if still accessible) or from the customer install journal. These settings are not backed up in the Cloud Controller. If these settings are not available, pick a new range for this network (for example, `10.65.0.0/24`) since it is internal to the network stack.

- Obtain the hypervisor network settings (DHCP_RANGE_START, DHCP_RANGE_END, DHCP_ROUTER_IP, DHCP_NETMASK, PRIMARY_DNS, SECONDARY_DNS) from the primary interconnect. This information comes from the configuration file and is located in the configuration spreadsheet provided before installation.

Important: This procedure can be performed only under supervision of or with help from Lenovo Support.

Execution: On hardware only

Performed by: Lenovo Support

Background: Dual interconnect configuration (DS-V3) is active-active. A DHCP server runs on both interconnect switches, and the DHCP lease file is synchronized between the two interconnect. In the event that the primary interconnect switch goes down, the DHCP server on the secondary interconnect begins to service IP requests.

Attention:

- This only works on the Hypervisor Network. DHCP High Availability does not function on the Management Network.
- This procedure requires physical intervention on the hardware.

Note: Ensure that the primary interconnect is shut down when running these steps. Having two DHCP servers in split-brain configuration causes severe disruptions on the stack network.

To make the secondary interconnect switch into the primary interconnect, follow these steps:

Step 1. (Optional) If the primary interconnect is not shut down, **sign in to the primary interconnect in Support Mode** or, if at the data center, using a method provided by the manufacturer. The username is `admin` and the password is the one you set during installation.

Step 2. (Optional) If the primary interconnect is not shut down, run the following commands on the primary interconnect to disable the DHCP service:

```
> sudo service isc-dhcp-server stop

> sudo service isc-dhcp-server disable
```

Step 3. Sign into the second interconnect using the same procedure.

- Edit `/etc/tacp/tacp_switch_resource.sh` with the hypervisor network settings (DHCP_RANGE_START, DHCP_RANGE_END, DHCP_ROUTER_IP, DHCP_NETMASK, PRIMARY_DNS, SECONDARY_DNS) from the primary interconnect.
- Copy the backed-up lease database `/var/lib/dhcp/dhcpd.leases` from the primary interconnect to the corresponding location on the secondary interconnect, if available.

Note: If the backed-up lease database is not available, you need to configure the secondary interconnect as new. ThinkAgile Advantage assists you with this procedure.

- Run the following command to create `/etc/dhcp/vlan-mgmt.conf` with the hardware management network settings (BMC_DHCP_START, BMC_DHCP_END, BMC_ROUTER_IP, BMC_PREFIX) from the primary interconnect.

```
> sudo /usr/share/tacp/reset_bmc_network.sh
```

Step 4. If you were unable to copy `/etc/dhcp/dhcp-failover.conf` from the primary interconnect, create it with the following content.

```
failover peer "failover-partner" {
    primary;
    address 100.64.253.1;
```

```

port 647;
peer address 100.64.253.2;
peer port 647;
max-response-delay 60;
max-unacked-updates 10;
auto-partner-down 1;
mclt 120;
split 255;
load balance max seconds 3;
}

```

- Step 5. Disconnect the hardware management network cable from the former primary interconnect and connect it to the current primary (former secondary) interconnect.




Figure 57. Recabling management switch from primary to secondary interconnect

- Step 6. Start the DHCP server in the new primary interconnect with the following command:

```
> sudo service isc-dhcp-server start
```

Recovering from DNS misconfiguration

When specifying DNS settings for the first time, specifying incorrect DNS settings can cause issues with DNS resolution

If this happens, do not attempt to rerun the script `v3/tacp_switch_config_v3`.

Important: This procedure can be performed only by Lenovo Professional Services.

To recover from DNS misconfiguration, follow these steps:

Step 1. Update `/etc/tacp/tacp_switch_resource.sh` with the correct DNS IP address on both interconnect switches:

```
### Hardware management network DHCP configuration ###

BMC_DHCP_START=10.103.31.2
BMC_DHCP_END=10.103.31.253
BMC_ROUTER_IP=$BMC_VLAN_IP
BMC_PREFIX=$BMC_VLAN_IP_PREFIX
DNS=10.0.64.10

#SWITCH_HA="DS-V3"|"SS-V3"

SWITCH_HA="DS-V3"
```

Step 2. Edit `/etc/dhcp/vlan-mgmt.conf` on the primary interconnect switch:

This file is used to specify the hardware management network setting.

```
subnet 10.103.31.0 netmask 255.255.255.0 {
    option domain-name-servers 10.0.64.10;
    option routers 10.103.31.1;
    default-lease-time 43200;
    max-lease-time 43200;

    pool {
        range 10.103.31.2 10.103.31.253;
    }
}
```

Step 3. Edit `/etc/dhcp/vlan-main.conf` on the primary interconnect switch:

```
subnet 10.102.31.0 netmask 255.255.255.0 {
    option domain-name-servers 10.0.64.10;
    option routers 10.102.31.254;
    default-lease-time 43200;
    max-lease-time 43200;

    #Disable DHCP-HA until off rules are fixed
    #failover peer "failover-partner";

    pool {
        range 10.102.31.2 10.102.31.253;
    }
}
```

Step 4. Restart the DHCP server on the primary interconnect switch:

```
> sudo service isc-dhcp-server stop  
> sudo service isc-dhcp-server start
```

Interconnect parts replacement procedures

Procedures for replacing the Lenovo ThinkAgile CP interconnect switch.

Important: Interconnect switch maintenance and replacement procedures should be performed only by Lenovo Support.

For information about the terms of the warranty, visit <https://datacentersupport.lenovo.com/us/en/warrantylookup>.

Replace the interconnect switch

Procedures for replacing the Lenovo ThinkAgile CP interconnect.

To replace an interconnect switch, first remove the unit from the rack and then install the new interconnect switch as described in the following sections.

Remove an interconnect switch

Procedure for removing the Lenovo ThinkAgile CP interconnect.

Procedure performed by: customer or field service

To remove the ThinkAgile CP interconnect from a standard rack, complete the following steps:

- Step 1. Disconnect all external cables on power and Ethernet ports.
- Step 2. Loosen and remove the screws (1) and clip nuts (or cage nuts) to release the switch and the rail from the rack.
- Step 3. Remove the switch from the rail.
- Step 4. Remove the two screws (2) used to secure each of the rear-post brackets at the midpoint on the sides of the rails.
- Step 5. Remove the four screws (3) on the front- and rear-post brackets.




Figure 58. Removing the interconnect from the rack

Replace the interconnect switch by following the steps in [“Install the ThinkAgile CP interconnect switch”](#) on [page 128](#).

Install the ThinkAgile CP interconnect switch

Procedures for installing ThinkAgile CP interconnect switch in a rack.

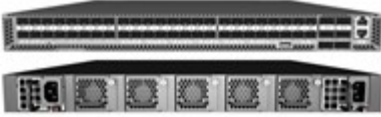
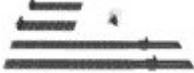

Attention:

- The switch drawings in this topic are for illustration only and may not match your specific switch model.
- If you are using a non-Lenovo switch as the management switch, keep 1U of empty space between the interconnect switches and the management switch so that power or network cables can use this space.

To install the ThinkAgile CP interconnect, do the following:

Note: When installing two interconnects, make sure that you install the secondary interconnect lower in the rack than the primary interconnect. The secondary interconnect is labeled with a rack location label that is lower than the rack location label of the primary interconnect.


Step 1. Unpack the switch and check the contents.

Component	Description
	7Y67 switch
	Rack mounting kit <ul style="list-style-type: none">• 2 front-post brackets• 2 rear-post brackets• 20 screws• 2 ear-locking screws
	Console cable, RJ-45 to DB-9

CAUTION:

The switch includes plug-in power supply (PSU) and fan tray modules that are installed into its chassis. All installed modules must have a matching airflow direction. That is, if the installed power modules have a front-to-back (F2B) airflow direction, all the installed fan tray modules must also have an F2B airflow direction.

Step 2. Attach the brackets.




CAUTION:

Installing the switch in a rack requires two people. One person should position the interconnect switch in the rack, while the other secures it using the rack screws.

- a. Attach each of the rails to the interconnect switch using four of the included rail screws.
- b. Use an additional two screws to secure each of the rails at the midpoint on the sides of the interconnect switch.
- c. Use the screws and cage nuts supplied with the rack to secure the switch in the rack.


Note: Make sure that you attach the cage nuts to the rack before installing the interconnect switch.

Step 3. Adjust the rear-post rail flange and then lock the position of the rear-post rail flange using the included position-locking screws.



Note: You can also adjust the rear-post rail flange to fit different rack depths from 56 cm to 75 cm.

Step 4. Ground the interconnect switch.




- a. Ensure the rack is properly grounded and in compliance with ETSI ETS 300 253. Verify that there is a good electrical connection to the grounding point on the rack (no paint or isolating surface treatment).
- b. Attach the grounding wire #14 AWG to the grounding point on the switch rear panel. Then connect the other end of the wire to the rack ground.

CAUTION:


The earth connection must not be removed unless all supply connections have been disconnected.

Step 5. Connect the power.




- a. Install one or two AC power modules in the switch. The switch supports up to two PSUs that must have the same matching airflow direction as the installed fan tray.
- b. Connect an external AC power source to the modules.

Step 6. Verify basic switch operation by checking the system LEDs. When operating normally, the PSU1/ PSU2, Diag, and Fan LEDs should all be green.



Step 7. Connect the network cables.




- a. For the RJ-45 Management port, connect 100-ohm Category 5, 5e or better twisted-pair cable.
- b. Connect DAC cables to the SFP+/QSFP+ slots. Or first install SFP+/QSFP+ transceivers and then connect fiber optic cabling to the transceiver ports.

The following transceivers are supported:

- 40G QSFP+ SR4 Transceiver 90Y3505
- 10G BASE-T (RJ45) SFP Transceiver SBB7A02090
- 10G SFP+ SR Transceiver 46C3446

Note: As connections are made, check the port status LEDs to be sure the links are valid.

Step 8. View the product label. The switch product label is located below SFP+ ports 7–12 on left side of the front panel. Pull the label out to view the product information.



The installation of the interconnect switch is now complete.

Replace the secondary interconnect in a DS-V3 configuration

Complete these steps to replace the secondary interconnect in a DS-V3 configuration.

Before you begin:

ThinkAgile Advantage Support: Request Support to retrieve the following interconnect settings from the Cloud Controller database:

- appliance_serial_number, controller_uuid, hypervisor network settings, uplink network settings, and the Network Controller Services version running on the interconnect
- the reserved VLAN IDs (Hardware management network, MLAG/ISL, Hypervisor, Uplink)
- the peerServiceTag of the primary interconnect

Procedure performed by: customer or field service, Lenovo Support

Step 1. Configure the interconnect you are replacing outside of the stack environment.

The default credentials are:

- Username: manager
- Password: ThinkAgileCP

Step 2. On the replacement interconnect, perform the following steps in Cloud Controller:

- a. Edit `/etc/tacp/tacp_switch_resource.sh` to populate the entries based on settings gathered above (see Installation documents).
- b. Replace `/etc/tacp/controller_uuid` with the value from the database.
- c. Update `/usr/share/tacp/datastore.sqlite` using `sqlite3` commands, with `<DHCP_RANGE_START>` and `<DHCP_RANGE_END>` representing values from the database:

```
INSERT INTO property (name,value) VALUES ('DHCP_RANGE_START', '<DHCP_RANGE_START>');
INSERT INTO property (name,value) VALUES ('DHCP_RANGE_END', '<DHCP_RANGE_END>');
```

Step 3. Redo any per customer custom configuration from the install journal.

Step 4. On the primary interconnect switch:

- a. Edit `/etc/tacp/tacp-network-controller.conf` and set `peerServiceTag` to the replacement interconnect's serial number.
- b. Restart the Network Controller Services using the following command: `> sudo service tacp-network-controller restart`

Step 5. Replace the failed interconnect switch with the new interconnect, including the cabling. See [“Cabling a dual-interconnect configuration” on page 138](#).

Step 6. Power on the replacement interconnect switch.

Cabling a single-interconnect configuration

Procedures for cabling the ThinkAgile CP single-interconnect switch configuration.

Complete these procedures to cable a single-interconnect configuration.

Hardware management network connections

Connect the following ports to the management switch:

- Interconnect port 47 using a SFP+ Direct Attach Cable (DAC)
- Compute block port (C.M.) using a RJ45 patch cable
- Upper storage block port (S.U.M.) using a RJ45 patch cable
- Lower storage block port (S.L.M.) using a RJ45 patch cable

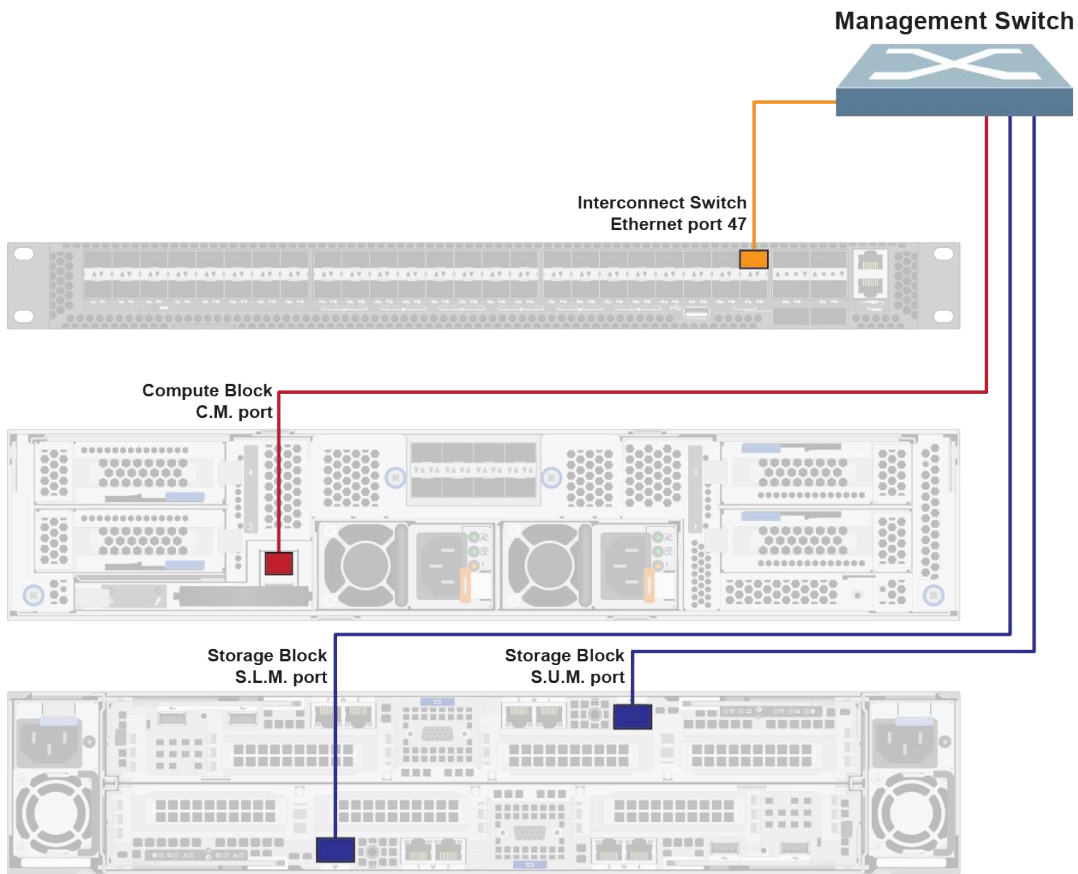
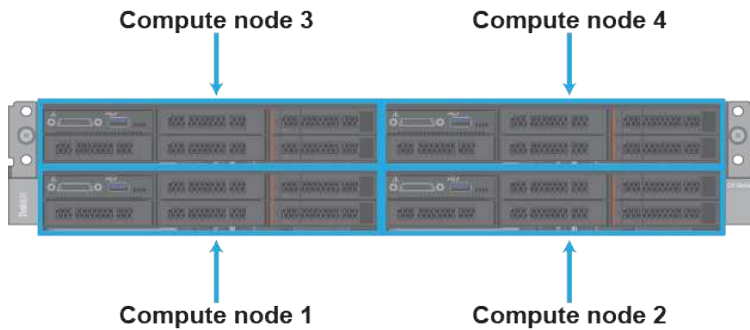


Figure 59. Hardware management network connectivity in a single interconnect configuration

Note: If not using the Lenovo ThinkSystem NE0152T RackSwitch as the management switch, you may require different cables for the management network connections.

Compute block connections (single-interconnect configuration)

The compute block consists of an enclosure with up to four compute nodes, which are numbered as follows:



Connect the compute nodes according to the populated compute node location and corresponding port mapping. The cables to be used are clearly labeled; make sure that you use the correct cables.

You need two SFP+ Direct Attach Cables (DACs) for each populated compute node.

The compute block port mapping is as follows:

Front compute node location	Back network port mapping	Network interconnect port
Top left	Compute node 3, port 1 (C.3.1)	Interconnect, port 5
	Compute node 3, port 2 (C.3.2)	Interconnect, port 6
Top right	Compute node 4, port 1 (C.4.1)	Interconnect, port 7
	Compute node 4, port 2 (C.4.2)	Interconnect, port 8
Bottom left	Compute node 1, port 1 (C.1.1)	Interconnect, port 1
	Compute node 1, port 2 (C.1.2)	Interconnect, port 2
Bottom right	Compute node 2, port 1 (C.2.1)	Interconnect, port 3
	Compute node 2, port 2 (C.2.2)	Interconnect, port 4

Note: You only need to cable the number of compute nodes that are installed in the compute enclosure.




Figure 60. Single—interconnect configuration — compute node and network interconnect connections

Storage block connections (single-interconnect configuration)

Connect the storage controllers according to the populated controller location and corresponding port mapping. The cables to be used are clearly labeled; make sure that you use the correct cables.

You need eight Category 6 patch cables and eight SFP+ transceivers. Attach a single SFP+ transceiver to each patch cable.

The storage block port mapping is as follows:

Front storage controller location	Back network port mapping	Network interconnect port
Upper	Upper storage controller, port 1 (S. U.1)	Interconnect, port 17
	Upper storage controller, port 2 (S. U.2)	Interconnect, port 20
	Upper storage controller, port 3 (S. U.3)	Interconnect, port 21
	Upper storage controller, port 4 (S. U.4)	Interconnect, port 24
Lower	Lower storage controller, port 1 (S. L.1)	Interconnect, port 25
	Lower storage controller, port 2 (S. L.2)	Interconnect, port 28
	Lower storage controller, port 3 (S. L.3)	Interconnect, port 29
	Lower storage controller, port 4 (S. L.4)	Interconnect, port 32

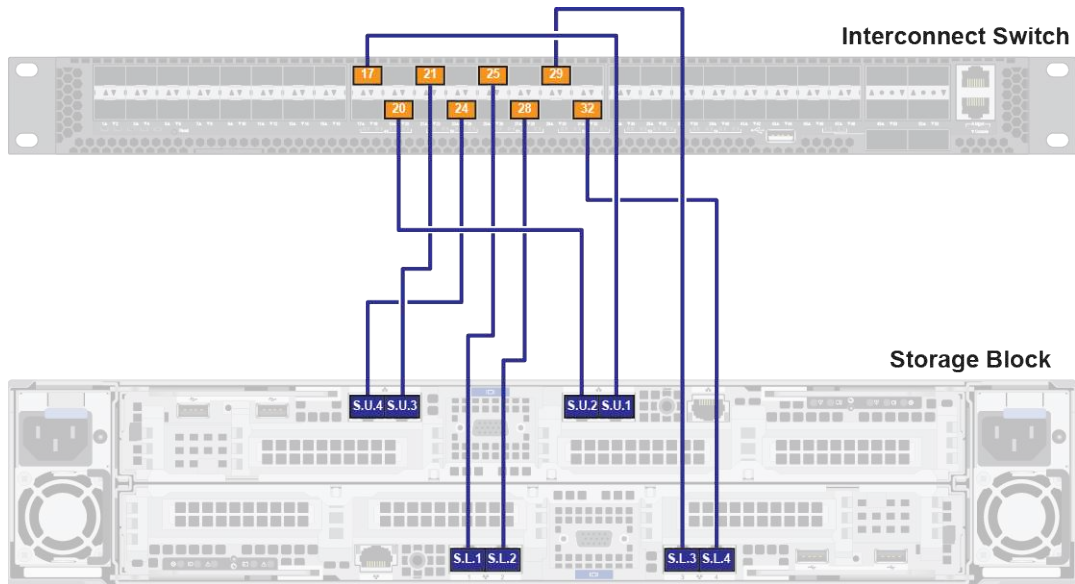


Figure 61. Single—interconnect configuration — storage block and network interconnect connections

Cabling a dual-interconnect configuration

Procedures for cabling the ThinkAgile CP dual-interconnect switch configuration.

Hardware management network connections (dual-interconnect configuration)

Connect the following ports to the management switch:

- Primary interconnect port 47 using a SFP+ Direct Attach Cable (DAC)
- Compute block port (C.M.) using a RJ45 patch cable
- Upper storage block port (S.U.M.) using a RJ45 patch cable
- Lower storage block port (S.L.M.) using a RJ45 patch cable

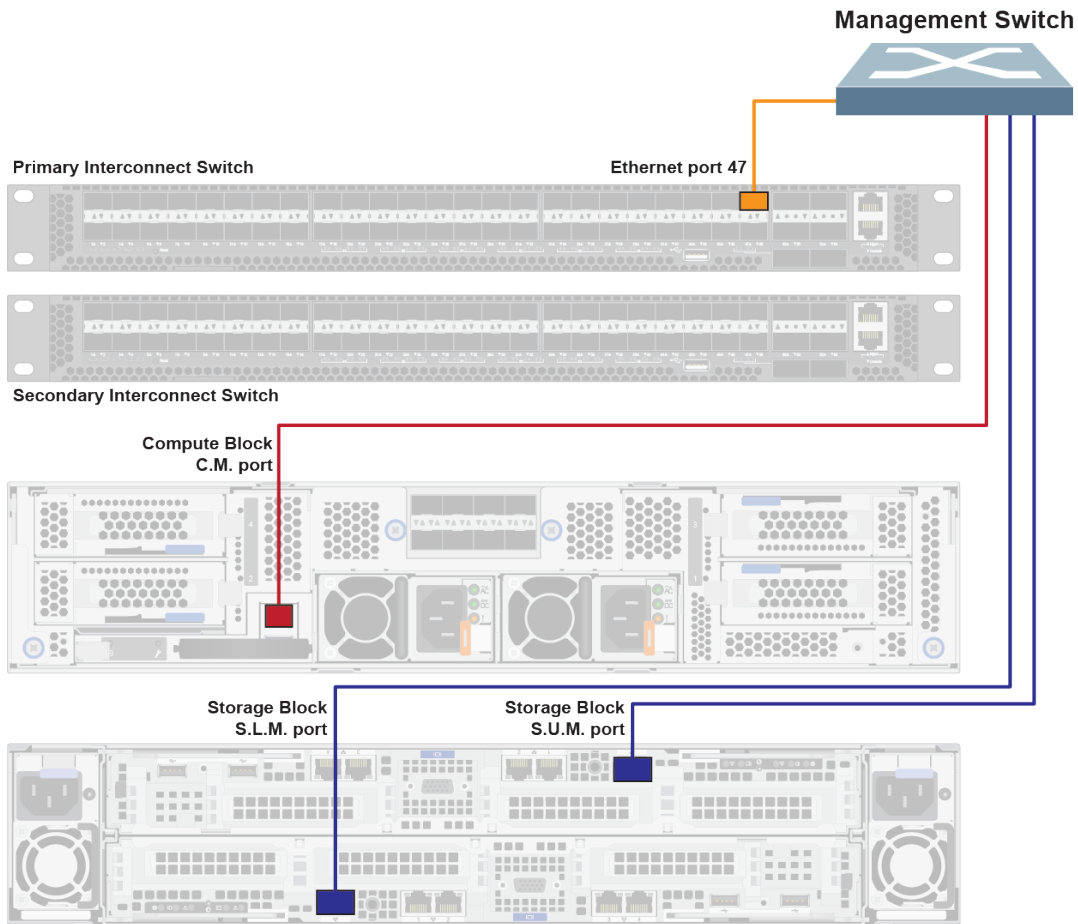
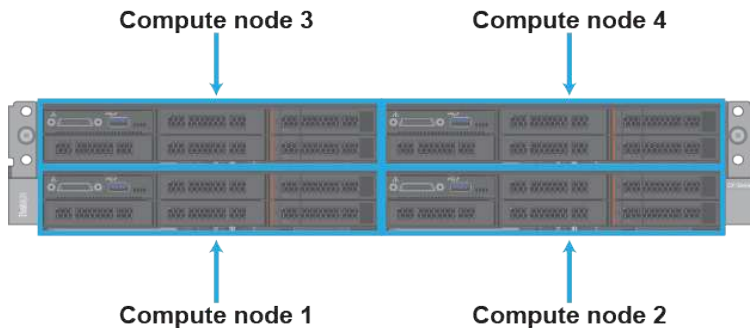


Figure 62. Hardware management network connectivity in a dual—interconnect configuration

Note: If not using the Lenovo ThinkSystem NE0152T RackSwitch as the management switch, you may require different cables for the management network connections.

Compute block connections (dual-interconnect configurations)

The compute block consists of an enclosure with up to four compute nodes, which are numbered as follows:



Connect the compute nodes according to the populated compute node location and the corresponding port mapping. The cables to be used are clearly labeled; make sure that you use the correct cables.

You need two SFP+ Direct Attach Cables (DACs) for each populated compute node.

The compute block port mapping is as follows:

Front compute node location	Back network port mapping	Network interconnect port
Top left	Compute node 3, port 1 (C.3.1)	Primary interconnect, port 3
	Compute node 3, port 2 (C.3.2)	Secondary interconnect, port 3
Top right	Compute node 4, port 1 (C.4.1)	Primary interconnect, port 4
	Compute node 4, port 2 (C.4.2)	Secondary interconnect, port 4
Bottom left	Compute node 1, port 1 (C.1.1)	Primary interconnect, port 1
	Compute node 1, port 2 (C.1.2)	Secondary interconnect, port 1
Bottom right	Compute node 2, port 1 (C.2.1)	Primary interconnect, port 2
	Compute node 2, port 2 (C.2.2)	Secondary interconnect, port 2

Note: You only need to cable the number of compute nodes that are installed in the compute enclosure.

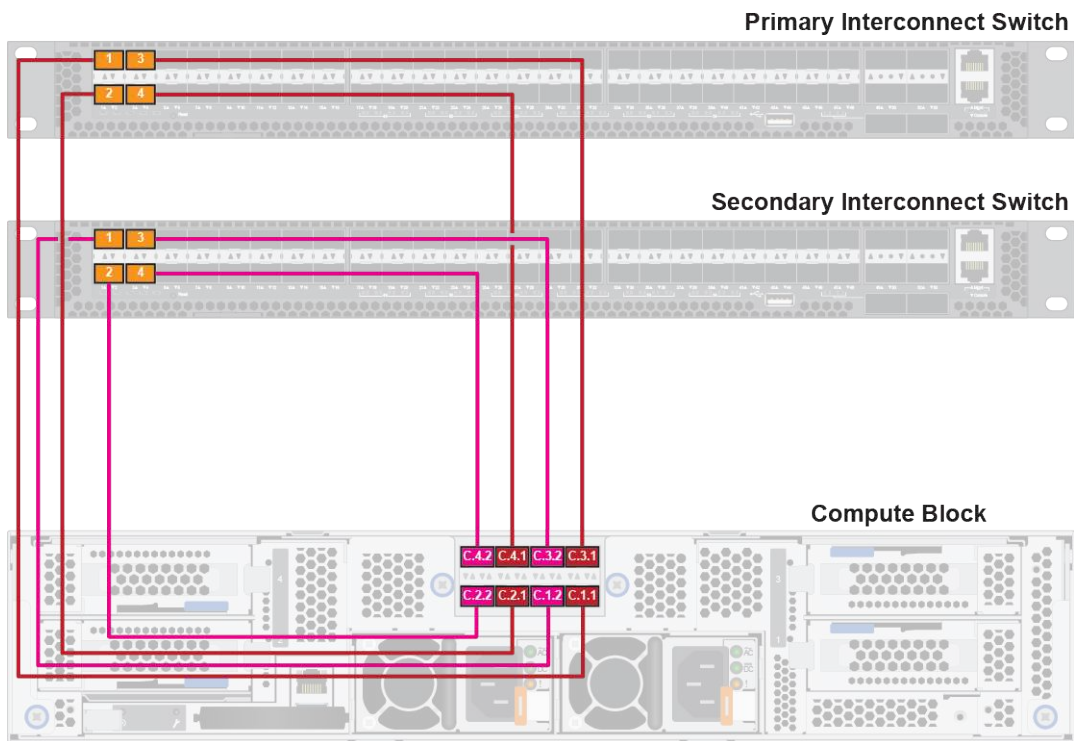


Figure 63. Dual—interconnect configuration — compute node and network interconnect connections

Storage block connections (dual-interconnect configurations)

Connect the storage controllers according to the populated controller location and corresponding port mapping and interconnect. The cables to be used are clearly labeled; make sure that you use the correct cables.

You need eight Category 6 patch cables and eight SFP+ transceivers. Attach a single SFP+ transceiver to each patch cable.

The storage block port mapping is as follows:

Front storage controller location	Back network port mapping	Network interconnect port
Upper	Upper storage controller, port 1 (S. U.1)	Primary interconnect, port 17
	Upper storage controller, port 2 (S. U.2)	Primary interconnect, port 20
	Upper storage controller, port 3 (S. U.3)	Secondary interconnect, port 17
	Upper storage controller, port 4 (S. U.4)	Secondary interconnect, port 20
Lower	Lower storage controller, port 1 (S. L.1)	Primary interconnect, port 21
	Lower storage controller, port 2 (S. L.2)	Primary interconnect, port 24
	Lower storage controller, port 3 (S. L.3)	Secondary interconnect, port 21

Front storage controller location	Back network port mapping	Network interconnect port
	Lower storage controller, port 4 (S.L.4)	Secondary interconnect, port 24

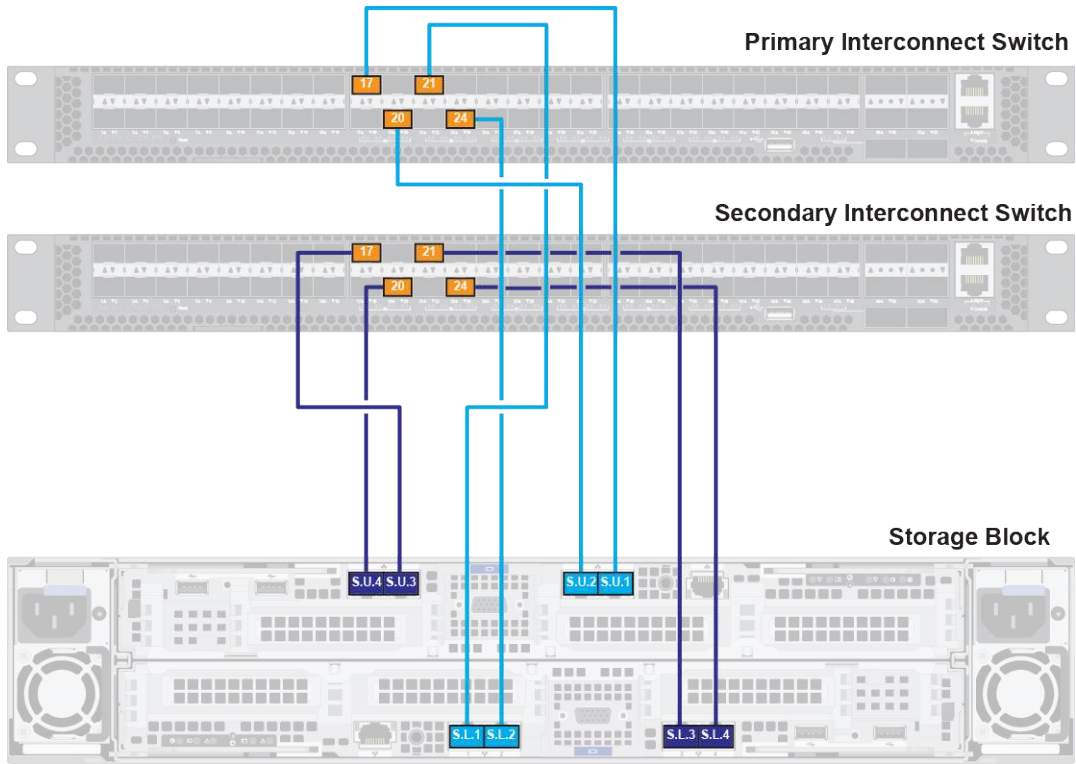




Figure 64. Dual—interconnect configuration — storage block and network interconnect connections

Peer-link connection (dual-interconnect configuration)

Set up the peer link between the interconnects using the QSFP+ Direct Attach Cables (DACs) that are provided with the configuration.

1. Connect port 49 on the primary interconnect to port 49 on the secondary interconnect.
2. Connect port 50 on the primary interconnect to port 50 on the secondary interconnect.

Primary Interconnect Switch



Secondary Interconnect Switch

Figure 65. Dual-interconnect configuration — inter-interconnect link network connectivity

Replacing a cable between two interconnect switches

Information on replacing a cable that connects two interconnects.

Procedure performed by: Customer or field service

Replacing one cable that connects two interconnects can be done live without any specific preparations steps.

When a cable is disconnected, the product data uplink may lose a portion of its uplink, based on the installation network topology.

Hardware management network cable replacement

This section describes how to replace a hardware management network cable.

Procedure performed by: Customer or field service

When replacing cables that connect any Baseboard Management Controller (BMC) from the compute enclosure, storage controllers, or the primary interconnect-dedicated port to the management switch, you are required to unplug all hardware management network cables and then plug back into the management switch to initiate a DHCP renewal. This operation can be done live without any specific pre- or post-steps.

During this operation, the hardware management network, the storage HA feature, and hardware detection functionality are unavailable. Hardware-related statistics also cannot be retrieved.

Management switch maintenance procedures (optional)

Procedures for replacing the Lenovo ThinkAgile CP management switch.

Note: A dedicated out-of-band management switch is required for this solution. If the Lenovo ThinkSystem NE0152T RackSwitch is not used, you must provide an out-of-band management switch with the correct configuration. Disregard this procedure if you are not including the Lenovo NE0152T as the management switch in your solution.

Refer to [Chapter 2 “Roles and responsibilities” on page 15](#) to determine who can perform these procedures.

For information about the terms of the warranty, visit <https://datacentersupport.lenovo.com/us/en/warrantylookup>.

Replacing the management switch (optional)

This section covers the procedures required to replace an existing management switch.

Note: A dedicated out-of-band management switch is required for this solution. If the Lenovo ThinkSystem NE0152T RackSwitch is not used, you must provide an out-of-band management switch with the correct configuration. Disregard this procedure if you are not including the NE0152T as the management switch in your solution.

Removing the management switch (optional)

This section describes how to remove the ThinkSystem NE0152T RackSwitch from a standard 19-inch equipment rack. For information about removing the switch from other supported racks, see the appropriate section in this chapter.

Note: A dedicated out-of-band management switch is required for this solution. If the Lenovo ThinkSystem NE0152T RackSwitch is not used, you must provide an out-of-band management switch with the correct configuration. Disregard this procedure if you are not including the NE0152T as the management switch in your solution.

This section describes how to remove the NE0152T from a standard 19-inch equipment rack.

Procedure performed by: customer or field service

To replace the NE0152T from a standard rack, complete the following steps:

Step 1. If the power cord retention clip is installed, remove it:




Figure 66. Power-cord retention clip removal

Step 2. Disconnect all external cables.

Step 3. Loosen and remove M6 screws, washers, and clip nuts (or cage nuts) to release the switch unit from the rack.





Figure 67. Releasing switch unit from the rack

Step 4. Slide the switch unit out of the rack.

Step 5. Loosen and remove the M4 screws attaching the mounting bracket on each side of the switch.

Figure 68. Removing mounting brackets from switch



Install the replacement NE0152T unit by completing the steps in [“Install the management switch \(optional\)”](#) on page 146.

Install the management switch (optional)

This section covers the procedures to install the ThinkSystem NE0152T RackSwitch as the management switch in a ThinkAgile CP solution.

Note: A dedicated out-of-band (OOB) management switch is required for this solution. If the Lenovo ThinkSystem NE0152T RackSwitch is not used, you must provide an OOB management switch with the correct configuration. Disregard this procedure if you are not including the NE0152T as the management switch in your solution.

This section describes how to install the NE0152T switch in a standard 19-inch equipment rack using the mounting kit included with the switch.

The following parts are included in the standard mounting kit:

Item number	Description	Quantity
1	Mounting brackets	2
2	M4 screws	8
3	M6 screws	4
4	M6 clip nuts	4
5	M6 cage nuts	4
6	M6 locking washers	4

Attention:

- The rack-mounting frame may be unable to support the weight of the switch with only the front post mounting brackets (2-post application). If the switch has an undesirable amount of sag, it is recommended to use a 4-post mounting kit.
- For earthquake stability, mount the switch in a 4-post rack.
- If you are using a non-Lenovo switch as the management switch, keep 1U of empty space between the interconnect switches and the management switch so that power or network cables can use this space.

Statement 26

CAUTION:



Do not place any object on top of rack-mounted devices.

To install the ThinkSystem NE0152T RackSwitch in a standard equipment rack, complete the following steps:

- Step 1. Locate, record, and retain the product switch information in order to configure and register your product.
- Step 2. Attach a mounting bracket (item 1 in the illustration) to each side of the switch with M4 screws (item 2). Torque the screws to approximately 2.0 newton-meters (Nm) \pm 0.1 Nm (17.7 inch-pounds).




Figure 69. Attaching mounting brackets to switch

- Step 3. From the front, slide the switch into the rack at the desired height.
- Step 4. Secure the switch unit to the rack posts with M6 screws (item 3 in illustration), washers (item 6), and either clip nuts (item 4) or cage nuts (item 5). Torque the screws to approximately 5.7 Nm \pm 0.1 Nm (50 inch-pounds).




Figure 70. Securing switch to rack posts

Step 5. Connect all external cables.

Step 6. If not using the 1U air-duct option, install the power cord retention clip.

Note: The power cord retention clip is not used when the 1U air-duct option is installed.




Figure 71. Installing power cord retention clip

The installation of the management switch is now complete.

Hardware management network cable replacement

This section describes how to replace a hardware management network cable.

Procedure performed by: Customer or field service

When replacing cables that connect any Baseboard Management Controller (BMC) from the compute enclosure, storage controllers, or the primary interconnect-dedicated port to the management switch, you are required to unplug all hardware management network cables and then plug back into the management switch to initiate a DHCP renewal. This operation can be done live without any specific pre- or post-steps.

During this operation, the hardware management network, the storage HA feature, and hardware detection functionality are unavailable. Hardware-related statistics also cannot be retrieved.

Chapter 6. Licensing considerations

The following table explains the details of the ThinkAgile CP Guardian edition.

	Guardian
Operating System (for compute nodes and storage controllers)	RHEL 7
Network Operating System for Interconnect Switches	Pica8 OS
At-Rest Encryption	✓
Compliance	
KMIP Compliant	✓
FIPS 140-2	✓
Common Criteria	✓
USGv6 (DOD IPv6)	✓
STIG	✓

Lenovo