



# Lenovo ThinkAgile CP Storage Capacity Upgrade Guide



**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](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>

**First 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

<b>Figures</b> . . . . .	<b>.iii</b>	Installing the drives. . . . .	22
<b>Tables</b> . . . . .	<b>v</b>	Adding drives to the Cloud Controller . . . . .	25
<b>Chapter 1. ThinkAgile CP Overview</b> . . . . .	<b>1</b>	Install the ThinkAgile CP storage block. . . . .	37
ThinkAgile CP4000 series overview . . . . .	4	Install the storage slide rails into a round hole rack . . . . .	39
ThinkAgile CP6000 series overview . . . . .	7	Install the storage slide rails into a square hole rack . . . . .	41
ThinkAgile CP solution components . . . . .	10	Install the storage enclosure . . . . .	43
Storage block parts list . . . . .	10	<b>Chapter 4. Compute block capacity upgrade</b> . . . . .	<b>47</b>
Compute block parts list . . . . .	12	Install a compute node . . . . .	47
Power cords . . . . .	14	Power on the compute node . . . . .	48
<b>Chapter 2. Installation Guidelines</b> . . . . .	<b>15</b>	Install the ThinkAgile CP compute block . . . . .	49
System reliability guidelines . . . . .	16	Install the compute slide rails . . . . .	50
Working inside the solution with the power on . . . . .	16	Install the compute enclosure . . . . .	52
Handling static-sensitive devices . . . . .	17	<b>Chapter 5. Licensing considerations</b> . . . . .	<b>55</b>
<b>Chapter 3. Storage block capacity upgrade</b> . . . . .	<b>19</b>	<b>Chapter 6. Getting help and technical assistance</b> . . . . .	<b>57</b>
Storage capacity expansion . . . . .	19		
Hot-swap drive installation preparations . . . . .	20		
Preparing the drives . . . . .	21		



---

## Figures

1.	Storage block . . . . .	2	14.	Installing the enclosure into the rack . . . . .	44
2.	Compute block . . . . .	2	15.	Installing screws on front of enclosure . . . . .	44
3.	CP Interconnect . . . . .	3	16.	Reinstalling the side cover . . . . .	45
4.	Management switch . . . . .	3	17.	ThinkAgile CP compute slide rails box contents . . . . .	49
5.	Storage block components . . . . .	11	18.	Slide rails . . . . .	50
6.	Compute block components . . . . .	13	19.	Left side rear installation . . . . .	50
7.	Opening the drive tray handle . . . . .	22	20.	Left side front installation. . . . .	51
8.	Sliding the drive into the drive bay . . . . .	23	21.	Securing the compute block slide rails. . . . .	51
9.	ThinkAgile CP storage slide box contents . . . . .	37	22.	Front latches on slide rails . . . . .	52
10.	Screw pin types included with rail kit . . . . .	38	23.	ThinkAgile CP compute enclosure installation . . . . .	53
11.	Installing the rail in a round hole rack . . . . .	39	24.	Enclosure installation into rack. . . . .	53
12.	Installing the rail in a square hole rack . . . . .	41	25.	Compute enclosure thumbscrews . . . . .	54
13.	Removal of side cover . . . . .	43			





# Tables

1.	Storage block parts listing . . . . .	10	4.	Compute block capacity upgrade options . . . . .	47
2.	Compute block part listing . . . . .	12			
3.	Storage block capacity upgrade options . . . . .	19			





---

## 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](#)
- [Compute block](#)
- [Interconnect switch](#)
- [Management switch](#)
- [Services](#)
- [Software](#)

### 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](#)
- [Compute block](#)
- [Interconnect switch](#)
- [Management switch](#)
- [Services](#)
- [Software](#)

### 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. 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](#) and [Working inside the solution with the power on](#).
- 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 3. Storage block capacity upgrade

This section covers the procedures for installing additional storage drives or storage blocks to the ThinkAgile CP solution.

The following table shows the storage block capacity upgrade options for CP4000 and CP6000:

Table 3. Storage block capacity upgrade options

Model	ThinkAgile CP4000	ThinkAgile CP6000
Initial number of storage drives per storage block	8	
Maximum storage drives per storage block	24	
Initial number of storage blocks	1	
Maximum storage blocks	1	5

### Notes:

- Storage drives can be installed only in sets of eight drives, thus allowing for 8, 16 or 24 drives to be present at any time in a single storage block.
- The CP6000 can be upgraded with up to 120 storage drives (in sets of 8 drives), depending on the number of storage blocks installed.
- The maximum number of storage blocks in the CP6000 solution is limited by the number of available ports on the interconnect switches.

---

## Storage capacity expansion

This section covers the procedure to upgrade the storage block with additional hot-swap drives.

Use this information to install a set of eight hot-swap drives. You can install a hot-swap drive without turning off the storage, which helps you avoid significant interruption to system operation.

**Attention:** Before installing a hot-swap drive, read and follow the steps in the [Hot-swap drive installation preparations](#)

### Notes:

- The term “hot-swap drive” refers to all the supported hot-swap NVMe solid-state drives.
- When upgrading the storage drive capacity, hot-swap drives must be installed in sets of eight.
- 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 installation preparations

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

### FAQs

Before you begin the drive installation 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](#) section.

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

No.

#### **How many drives I can install at a time?**

You can add as many drives as needed. The current capacity of the storage enclosure is 24 drives. The only requirement is that you must add drives in sets of eight (8).

#### **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
```

## Installing the drives

This section covers the procedure to physically install the set of eight drives to upgrade the storage capacity of the ThinkAgile CP solution.



Before installing the drive into the storage block, ensure that they are properly prepared according to [Preparing the drives](#) section.

**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.

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

Step 1. Check that the NVME namespace exists before installing the new drives:

```
# ls /dev/nvme*
```

Following is an example output:

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

Step 2. To open the drive tray handle, press the release tab.




Figure 7. Opening the drive tray handle

Step 3. Slide the drive into the drive bay until it snaps into position.




Figure 8. Sliding the drive into the drive bay

Step 4. Close the drive tray handle to lock the drive in place.

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. Continue to install the other hot-swap drives in the set until the drive bay is full.

Step 7. Check that the new drives are detected by the storage block:

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

```
# lsblk
```

Following is an example output:

NAME	MAJ:MIN	RM	SIZE	RO	TYPE	MOUNTPOINT
...						
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_71c510a8973349de9183b18f0acc4b3d	253:13	0	1G	0	dm	
├─SPECIAL_TEMPLATE_STORE_ef7543aa482e44f493c13db90b5796e8	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_71c510a8973349de9183b18f0acc4b3d	253:13	0	1G	0	dm	
├─SPECIAL_TEMPLATE_STORE_ef7543aa482e44f493c13db90b5796e8	253:14	0	43.3T	0	dm	

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	S3HCNX0JC02194	1	795.00 GB / 795.00 GB	512 B + 0 B	GPNA9B3Q
/dev/nvme1n1	S3HCNX0JC02193	1	795.00 GB / 795.00 GB	512 B + 0 B	GPNA9B3Q
/dev/nvme2n1	S3HCNX0JC02233	1	795.00 GB / 795.00 GB	512 B + 0 B	GPNA9B3Q
/dev/nvme3n1	S3HCNX0JC02191	1	795.00 GB / 795.00 GB	512 B + 0 B	GPNA9B3Q
/dev/nvme4n1	S3HCNX0JC02189	1	795.00 GB / 795.00 GB	512 B + 0 B	GPNA9B3Q
/dev/nvme5n1	S3HCNX0JC02198	1	795.00 GB / 795.00 GB	512 B + 0 B	GPNA9B3Q
/dev/nvme6n1	S3HCNX0JC02188	1	795.00 GB / 795.00 GB	512 B + 0 B	GPNA9B3Q
/dev/nvme7n1	S3HCNX0JC02195	1	795.00 GB / 795.00 GB	512 B + 0 B	GPNA9B3Q

Check that the NVME namespace exists before installing the new drives:

```
# ls /dev/nvme*
```

Following is an example output:

```
/dev/nvme0 /dev/nvme1 /dev/nvme2 /dev/nvme3 /dev/nvme4 /dev/nvme5 /dev/nvme6 /dev/nvme7  
/dev/nvme8 /dev/nvme9 /dev/nvme10 /dev/nvme11 /dev/nvme12 /dev/nvme13 /dev/nvme14 /dev/nvme15
```

The new drives are now added to the RAID array as spare devices.

The set of eight (8) drives is installed in the storage block.

The eight drives need to be added to the Cloud Controller. See [Adding drives to the Cloud Controller](#).



## Adding drives to the Cloud Controller

This section covers the procedures to add the newly installed storage drives to the Cloud Controller.

There are several steps required to add the drives the new drives to the Cloud Controller:

1. Add the drives to the RAID array. See [Adding the new drives to the RAID array](#),
2. Increase the size of the VDO data. See [Increasing the size of the VDO](#).
3. Increase the sizes of the CBS pool and CBS pool data. See [Increasing the size of the CBS pool](#).
4. Update the CBS superblock. See [Updating the CBS superblock](#).

### Adding the new drives to the RAID array

This section covers the steps required to add the new drives to the RAID array.

Ensure that the set of eight (8) drives is properly installed according to the [Installing the drives](#) section.

**Procedure performed by:** Customer or field service

To add the newly installed drives to the RAID array, follow these steps:

Step 1. Create the RAID5 device for the eight drives:

```
#!/usr/share/tacp/lenovo/add_disks_lenovo.py
```

Following is an example output:

```
all disks are attached correctly to this controller.  
Raid device created /dev/md123  
Congratulations. You have successfully created the raid device  
Please run the following command to continue: mdadm --grow /dev/md/md50 --raid-device=2 --add /dev/md123  
Success
```

The script provides a `-verbose` option if you want to know details of how the RAID5 device is created.

Step 2. Grow the RAID50 device

```
mdadm --grow /dev/md/md50 --raid-device=2 --add /dev/md123
```

Following is an example output:

```
mdadm: level of /dev/md/md50 changed to raid4  
mdadm: added /dev/md123
```

**Note:** This step is not scripted because this may take up to a few hours depending on how much data is stored in the RAID array.

Step 3. Check the RAID rebuild status:

```
cat /proc/mdstat
```

Following is an example output:

```
Personalities : [raid1] [raid6] [raid5] [raid4] [raid0]

md123 : active raid5 dm-22[7](S) dm-21[6] dm-20[5] dm-19[4] dm-18[3] dm-17[2] dm-16[1] dm-15[0]
4657379328 blocks super 1.2 level 5, 16k chunk, algorithm 2 [7/7] [UUUUUUU]
bitmap: 1/6 pages [4KB], 65536KB chunk

md124 : active raid4 md123[2] md125[0]
4657247232 blocks super 1.2 level 4, 512k chunk, algorithm 5 [3/2] [U_]
[>.....] reshape = 1.3% (62142976/4657247232) finish=424.9min speed=180210K/sec

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: 1/6 pages [4KB], 65536KB chunk

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

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

unused devices: <none>
```

Step 4. Check and calculate the completion time of the process.

As displayed in the example output from the previous step, the reshape process has a completion time of 424.9 minutes (finish=424.9min). In hours, this is a little over seven (7) hours (424.9 minutes / 60 = 7.08 hours).

Step 5. Wait until the completion time is up (in this example, more than seven (7) hours), and then check that the process has indeed finished:

```
cat /proc/mdstat
```

Following is an example output:

```
Personalities : [raid1] [raid6] [raid5] [raid4] [raid0]

md123 : active raid5 dm-22[7](S) dm-21[6] dm-20[5] dm-19[4] dm-18[3] dm-17[2] dm-16[1] dm-15[0]
4657379328 blocks super 1.2 level 5, 16k chunk, algorithm 2 [7/7] [UUUUUUU]
bitmap: 0/6 pages [0KB], 65536KB chunk

md124 : active raid0 md123[2] md125[0]
9314494464 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 sdb[1](S) sda[0](S)
6306 blocks super external:imsm

unused devices: <none>
```

Step 6. Calculate new size of the RAID50 device:

```
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/nvme10n1	S3HCNX0K100055	1	795.00 GB / 795.00 GB	512 B + 0 B	GPNA9B3Q
/dev/nvme11n1	S3HCNX0K100061	1	795.00 GB / 795.00 GB	512 B + 0 B	GPNA9B3Q
/dev/nvme12n1	S3HCNX0K100046	1	795.00 GB / 795.00 GB	512 B + 0 B	GPNA9B3Q
/dev/nvme13n1	S3HCNX0K100092	1	795.00 GB / 795.00 GB	512 B + 0 B	GPNA9B3Q
/dev/nvme14n1	S3HCNX0K100076	1	795.00 GB / 795.00 GB	512 B + 0 B	GPNA9B3Q
/dev/nvme15n1	S3HCNX0K100015	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
/dev/nvme8n1	S3HCNX0K100085	1	795.00 GB / 795.00 GB	512 B + 0 B	GPNA9B3Q
/dev/nvme9n1	S3HCNX0K100072	1	795.00 GB / 795.00 GB	512 B + 0 B	GPNA9B3Q

In the example output, 16 drives can be observed as part of RAID50 device. Each drive has a capacity of 795 GB.

The RAID setup is organized by grouping eight (8) drives in a RAID5 device. Then all RAID5 devices form a RAID0 device.

**Note:** For each RAID5 device there is a spare drive.

The total capacity of the RAID50 device is given by the following formula:

**total capacity** = (size of a single drive) X (number of drives in the RAID5 configuration – 2) X (number of RAID5 devices)

Replacing the terms in the formula with the values from our example, the total capacity of the RAID50 device is:

**total capacity** = 795 GB X (8 – 2) X 2 = 9540 GB

Step 7. Check the calculated size of the RAID50 device against the size shown by the following command:

```
mdadm --detail /dev/md/md50
```

Following is an example output:

```
/dev/md/md50:
  Version : 1.2
  Creation Time : Thu Aug 16 19:35:39 2018
  Raid Level : raid0
  Array Size : 9314494464 (8882.99 GiB 9538.04 GB)
  Raid Devices : 2
  Total Devices : 2
  Persistence : Superblock is persistent

  Update Time : Fri Aug 17 04:23:27 2018
  State : clean
  Active Devices : 2
  Working Devices : 2
  Failed Devices : 0
  Spare Devices : 0

  Chunk Size : 512K

  Consistency Policy : none

  Name : any:md50
  UUID : d35a4763:b4b9b490:b85614db:6aeb696a
  Events : 2986

Number Major Minor RaidDevice State
0         9      125    0        active sync /dev/md/md5
2         9      123    1        active sync /dev/md/md5_2
```

The calculated size is almost the size displayed in the example output: 9538.04 GB.

Step 8. Check the RAID level:

```
mdadm --detail /dev/md/md50
```

In the example output from the previous step, the RAID level is 0.

**Note:** If the RAID level is RAID4, then change the RAID4 level to RAID0:

```
mdadm --grow --level 0 --raid-devices=2 /dev/md/md50
```

```
mdadm: level of /dev/md/md50 changed to raid0
```

**Attention:**

- The new RAID0 array needs to be reshaped a second time. This can take up to several hours to complete, depending on the size of the array. For example, expanding from 16 TB to 32 TB takes 11-12 hours.
- Ensure the Array Size has increased to the correct amount and that all the RAID5 devices are included. If not, **stop** and do not continue! You may need to reshape to RAID4 and then back to RAID0.
- Each set of eight (8) drives has one (1) parity drive and one (1) spare drive.

- The total Array Size for eight (8) drives is:  
**total array size = 6 X (size of a single drive)**
- The total Array Size for 16 drives is:  
**total array size = (6+6) X (size of a single drive)**
- When the reshape process is completed it look similar to the following example:

```
mdadm --detail /dev/md/md50
```

```

/dev/md/md50:
  Version : 1.2
  Creation Time : Thu Aug 16 19:35:39 2018
  Raid Level : raid0
  Array Size : 9314494464 (8882.99 GiB 9538.04 GB)
  Raid Devices : 2
  Total Devices : 2
  Persistence : Superblock is persistent

  Update Time : Fri Aug 17 04:23:27 2018
  State : clean
  Active Devices : 2
  Working Devices : 2
  Failed Devices : 0
  Spare Devices : 0

  Chunk Size : 512K

  Consistency Policy : none

  Name : any:md50
  UUID : d35a4763:b4b9b490:b85614db:6aeb696a
  Events : 2986

  Number Major Minor RaidDevice State
  0          9     125    0          active sync /dev/md/md5
  2          9     123    1          active sync /dev/md/md5_2

```

Step 9. Check the size of the RAID50 device in 512-byte sectors:

```
blockdev --getsz /dev/md/md50
```

```
18,628,988,928
```

The first 2048 512- byte sectors are used to store the internal ThinkAgile CP VDO-specific metadata.

The RAID50 device now includes the new storage drives.

After including the new drives in the RAID array, increase the size of the VDO data. See [Increasing the size of the VDO](#).

## Increasing the size of the VDO

This section covers the steps required to increase the size of the VDO data after adding the new storage drives to the RAID array.

Ensure that the set of eight (8) drives are properly included in the RAID array according to the [Adding the new drives to the RAID array](#) section.

**Procedure performed by:** Customer or field service

To increase the size of the VDO, follow these steps:

Step 1. Calculate the new size of the VDO according to the following formula:

**new VDO size** = (size of the RAID50 device) — (size of the internal VDO-specific metadata)

**Note:** All the sizes are calculated in number of 512-byte sectors.

As calculated in our example in the [Adding the new drives to the RAID array](#) section:

- The size of the RAID50 device is 18,628,988,928
- The size of the internal ThinkAgile CP VDO-specific metadata is 2,048

Therefore, the new VDO size is 18,628,988,928 — 2,048 = **18,628,986,880** 512-byte sectors.

Step 2. Increase the size of the VDO data and reload the VDO:

```
dmsetup reload vdo_data --table '0 18628986880 linear /dev/md/md50 2048'
```

Step 3. Resume the VDO:

```
dmsetup resume vdo_data
```

Step 4. Check the new size of the VDO data:

```
dmsetup table vdo_data
```

```
0 18628986880 linear 9:50 2048
```

```
lsblk | grep vdo | head -1
```

```
vdo_data 253:9 0 8.7T 0 dm
```

Step 5. Check the storage drive configuration

```
# lsblk
```

Following is an example output:

NAME	MAJ:MIN	RM	SIZE	RO	TYPE	MOUNTPOINT
...						
nvme8n1	259:12	0	740.4G	0	disk	
├─sed8	253:15	0	740.4G	0	crypt	
│ └─md123	9:123	0	4.3T	0	raid5	
│ │ └─md124	9:124	0	8.7T	0	raid0	
│ │ │ └─vdo_data	253:9	0	8.7T	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_a0528f49b71948aa8be9d7f4d490b97f	253:13	0	1G	0	dm	
│ │ │ │ │ │ │ └─SPECIAL_TEMPLATE_STORE_0657a4c083b44d8bb0d22cb4e6d0d750	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_a0528f49b71948aa8be9d7f4d490b97f	253:13	0	1G	0	dm	
│ │ │ │ │ │ │ └─SPECIAL_TEMPLATE_STORE_0657a4c083b44d8bb0d22cb4e6d0d750	253:14	0	43.3T	0	dm	
nvme9n1	259:12	0	740.4G	0	disk	
├─sed9	253:15	0	740.4G	0	crypt	
│ └─md123	9:123	0	4.3T	0	raid5	
│ │ └─md124	9:124	0	8.7T	0	raid0	
│ │ │ └─vdo_data	253:9	0	8.7T	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_a0528f49b71948aa8be9d7f4d490b97f	253:13	0	1G	0	dm	
│ │ │ │ │ │ │ └─SPECIAL_TEMPLATE_STORE_0657a4c083b44d8bb0d22cb4e6d0d750	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_a0528f49b71948aa8be9d7f4d490b97f	253:13	0	1G	0	dm	
│ │ │ │ │ │ │ └─SPECIAL_TEMPLATE_STORE_0657a4c083b44d8bb0d22cb4e6d0d750	253:14	0	43.3T	0	dm	

Step 6. Increase the physical volume size of the VDO volume:

```
# vdo growPhysical --name=vdo
```

Step 7. Increase the logical volume size to ten times the size of the physical volume:

a. Check the new size of the VDO data:

```
# dmsetup table vdo_data
```

```
0 18628986880 linear 9:124 2048
```

b. Calculate the new logical volume size according the following formula:

**new logical volume size** = (size of VDO data in 512-byte sectors) **X** (number bytes per sector)  
**X** (relative size of the logical volume compared to the physical volume)

In our example, this becomes:

**new logical volume size** = 18,628,986,880 **X** 512 **X** 10 = 88,829,931,640,625 bytes or approximately 88,829 GB.

- c. Increase the size of the logical volume:

```
# vdo growLogical --name=vdo --vdoLogicalSize=88829G
```

**Note:** The maximum size of the logical volume is 2,013,265 GB (1.92 PB). If the number you previously calculated is greater than that, use the maximum of 2,013,265 GB instead.

- d. Check the new size of the VDO data:

```
# lsblk | tail -10
```

NAME	MAJ:MIN	RM	SIZE	RO	TYPE	MOUNTPOINT
└─vdo_data	253:9	0	8.7T	0	dm	
└─┬─vdo	253:10	0	86.8T	0	vdo	
│ └─CBS_POOL_data	253:8	0	43.3T	0	dm	
│ └─CBS_POOL	253:12	0	43.3T	0	dm	
│ │ └─SPECIAL_METADATA_STORE_a0528f49b71948aa8be9d7f4d490b97f	253:13	0	1G	0	dm	
│ │ └─SPECIAL_TEMPLATE_STORE_0657a4c083b44d8bb0d22cb4e6d0d750	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_a0528f49b71948aa8be9d7f4d490b97f	253:13	0	1G	0	dm	
│ └─SPECIAL_TEMPLATE_STORE_0657a4c083b44d8bb0d22cb4e6d0d750	253:14	0	43.3T	0	dm	

- Step 8. Save the contents of the new VDO configuration to the RAID50 device:

```
cat /etc/vdoconf.yml > /dev/md/md50
```

After increasing the size of the VDO data, increase the sizes of the CBS pool and CBS pool data. See [Increasing the size of the CBS pool](#).



## Increasing the size of the CBS pool

This section covers the steps required to increase the sizes of the CBS pool and of the CBS pool data.

Ensure that the size of the VDO data is properly increased according to the [Increasing the size of the VDO](#) section.

**Procedure performed by:** Customer or field service

To increase the sizes of the CBS pool and the CBS pool data, follow these steps:

Step 1. Get the current size of the CBS pool and CBS pool data:

```
# dmsetup table CBS_POOL
```

```
0 93066911448 thin-pool 253:11 253:8 8192 0 0
```

```
# dmsetup table CBS_POOL_data
```

```
0 93066911448 linear 253:10 67608872
```

Step 2. Get the size of the VDO data:

```
blockdev --getsize /dev/mapper/vdo
```

```
186287915008
```

The size of the ThinkAgile CP superblock is 67,608,872.

Step 3. Calculate the new size for the CBS pool and CBS pool data according to the following formula:

**new CBS pool and CBS pool data size** = (size of VDO data) – (size of ThinkAgile CP superblock)

In our example, this becomes:

**new CBS pool and CBS pool data size** = 186,287,915,008 – 67,608,872 = 186,220,306,136

Step 4. Increase the size of the CBS pool data:

**Note:** All CBS pool input/output halts while the CBS pool data device is suspended until the device is resumed, so make sure these operations are as fast as possible. You can do this by stringing the commands together with semicolons and then executing them as one.

```
# dmsetup suspend CBS_POOL_data
# dmsetup reload CBS_POOL_data --table '0 186220306136 linear /dev/mapper/vdo
67608872'
# dmsetup resume CBS_POOL_data
```

Step 5. Check the size of the CBS pool data:

```
# lsblk /dev/md/md50 | head -n 100 | grep -v DISK
```

NAME	MAJ:MIN	RM	SIZE	RO	TYPE	MOUNTPOINT
└─md124	9:124	0	8.7T	0	raid0	
├─vdo_data	253:9	0	8.7T	0	dm	
├─vdo	253:10	0	86.8T	0	vdo	
├─CBS_POOL_data	253:8	0	86.7T	0	dm	
├─CBS_POOL	253:12	0	43.3T	0	dm	
├─SPECIAL_METADATA_STORE_a0528f49b71948aa8be9d7f4d490b97f	253:13	0	1G	0	dm	
├─SPECIAL_TEMPLATE_STORE_0657a4c083b44d8bb0d22cb4e6d0d750	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_a0528f49b71948aa8be9d7f4d490b97f	253:13	0	1G	0	dm	
└─SPECIAL_TEMPLATE_STORE_0657a4c083b44d8bb0d22cb4e6d0d750	253:14	0	43.3T	0	dm	

```
# dmsetup table CBS_POOL_data
```

0 186220306136 linear 253:10 67608872
---------------------------------------

Step 6. Increase the size of the CBS pool:

**Note:** All CBS pool input/output halts while the CBS pool device is suspended until the device is resumed, so make sure these operations are as fast as possible. You can do this by stringing the commands together with semicolons and then executing them as one.

```
# dmsetup suspend CBS_POOL
# dmsetup reload CBS_POOL --table '0 186220306136 thin-pool /dev/mapper/CBS_POOL_meta
/dev/mapper/CBS_POOL_data 8192 0 0'
# dmsetup resume CBS_POOL
```

Step 7. Check the size of the CBS pool:

```
# lsblk /dev/md/md50 | head -n 100 | grep -v DISK
```

NAME	MAJ:MIN	RM	SIZE	RO	TYPE	MOUNTPOINT
└─md124	9:124	0	8.7T	0	raid0	
├─vdo_data	253:9	0	8.7T	0	dm	
├─vdo	253:10	0	86.8T	0	vdo	
├─CBS_POOL_data	253:8	0	86.7T	0	dm	
├─CBS_POOL	253:12	0	86.7T	0	dm	
├─SPECIAL_METADATA_STORE_a0528f49b71948aa8be9d7f4d490b97f	253:13	0	1G	0	dm	
├─SPECIAL_TEMPLATE_STORE_0657a4c083b44d8bb0d22cb4e6d0d750	253:14	0	43.3T	0	dm	
├─CBS_POOL_meta	253:11	0	16G	0	dm	
├─CBS_POOL	253:12	0	86.7T	0	dm	
├─SPECIAL_METADATA_STORE_a0528f49b71948aa8be9d7f4d490b97f	253:13	0	1G	0	dm	
└─SPECIAL_TEMPLATE_STORE_0657a4c083b44d8bb0d22cb4e6d0d750	253:14	0	43.3T	0	dm	

```
# dmsetup table CBS_POOL
```

0 186220306136 thin-pool 253:11 253:8 8192 0 0
--

The sizes of the CBS pool and CBS pool data are now increased.

After increasing the sizes of the CBS pool and CBS pool data, update the CBS superblock. See [Updating the CBS superblock](#).

## Updating the CBS superblock

This section covers the steps required to update the CBS superblock.

Ensure that the sizes of the CBS pool and CBS pool data are properly increased according to the [Increasing the size of the CBS pool](#) section.

**Procedure performed by:** Customer or field service

To update the CBS superblock, follow these steps:

Step 1. Copy the current CBS superblock to a temporary file:

```
head -c 4096 /dev/mapper/vdo > /tmp/superblock.tmp
```

Step 2. Check the current values in the previously created temporary file:

```
# cat /tmp/superblock.tmp
```

```
{ "special_thin_blocks": { "TEMPLATE": { "size_sectors": "93066911448", "thinid": 2, "uuid":  
"0657a4c083b44d8bb0d22cb4e6d0d750"}, "METADATA": { "size_sectors": "2097152", "thinid": 1, "uuid":  
"a0528f49b71948aa8be9d7f4d490b97f"}}, "thin_block_tracking_enabled": true, "name": "CBS Superblock",  
"pool_device": { "size_sectors": "93066911448", "block_size_sectors": 8192}, "data_device":  
[{"offset_sectors": "67608872", "start_offset_sectors": "0", "size_sectors": "93066911448"}],  
"metadata_device": { "offset_sectors": "34054440", "size_sectors": "33554432"}, "thin_block_metadata":  
{ "record_size_bytes": 256, "record_count": 1000000, "superblock_offset_bytes": 4096}, "version": 2,  
"metadata_backup_device": { "offset_sectors": "500008", "size_sectors": "33554432"}}
```

Step 3. Get the new number of 512-byte sectors:

```
dmsetup table CBS_POOL
```

```
0 186220306136 thin-pool 253:11 253:8 8192 0 0
```

There are now 186,220,306,136 512-byte sectors.

Step 4. Update the `pool_device : size_sectors` and `data_device : size_sectors` in `/tmp/superblock.tmp`

a. Edit `superblock.tmp`.

```
vi /tmp/superblock.tmp
```

b. Type `i` to input data.

c. Replace the values of the `pool_device : size_sectors` and `data_device : size_sectors` with the number of 512-byte sectors from step 3.

In our example, this value is 186,220,306,136 512-byte sectors.

- d. Press ESC and enter :wq to save the changes and close the superblock.tmp file.

**Attention:**

- Do not modify the size\_sectors in the TEMPLATE section as this remains unchanged since the filesystem is not resized.
- Do not add extra characters or new lines as these may cause problems, such as the storage stack failing to reassemble the next time it is rebooted or HA is performed.

```
{ "special_thin_blocks": { "TEMPLATE": { "size_sectors": "93066911448", "thinid": 2, "uuid": "0657a4c083b44d8bb0d22cb4e6d0d750"}, "METADATA": { "size_sectors": "2097152", "thinid": 1, "uuid": "a0528f49b71948aa8be9d7f4d490b97f"}}, "thin_block_tracking_enabled": true, "name": "CBS Superblock", "pool_device": { "size_sectors": "186220306136", "block_size_sectors": 8192, "data_device": [{"offset_sectors": "67608872", "start_offset_sectors": "0", "size_sectors": "186220306136"}], "metadata_device": { "offset_sectors": "34054440", "size_sectors": "33554432"}, "thin_block_metadata": { "record_size_bytes": 256, "record_count": 1000000, "superblock_offset_bytes": 4096, "version": 2, "metadata_backup_device": { "offset_sectors": "500008", "size_sectors": "33554432"}}
```

- Step 5. Copy the configuration to the CBS superblock header:

```
cat /tmp/superblock.tmp > /dev/mapper/vdo
```

- Step 6. Restart the storage controller:

```
# tacp-storage-controller-safe-stop  
# systemctl start tacp-storage-controller
```

**Note:** This forces the storage controller to send a new registration message to the Management Layer, which updates the controller with the new available capacity.

- Step 7. Check that the new storage capacity is displayed in the Cloud Controller Web Portal.

The storage capacity upgrade is now complete. The new additional storage space can be used immediately.

---

## Install the ThinkAgile CP storage block

Procedures for installing ThinkAgile CP storage block in a rack.

To install the ThinkAgile CP storage block in a rack, you must do the following:

1. Unpack the storage block box and check the contents
2. Install the slide rails in the rack
3. Install the storage enclosure

Firstly, unpack the storage block box and check the contents:

- One storage block
- One storage slide rails box:
  - One left slide rail
  - One right slide rail
  - Six M5 screws
  - Eight screws for square hole rack
- Two Category 5e patch cables
- Eight Category 6 patch cables
- Eight SFP+ 10 GbE transceivers

**Note:** If any of the slide box items are missing or are damaged, contact your place of purchase.




Figure 9. ThinkAgile CP storage slide box contents

**Attention:** You can install the rails into the racks with round holes or square holes. Each rail comes preinstalled with four screw pins used for round holes. The screw pins used for square holes are provided separately. Ensure that you use the pins that are appropriate for the mounting holes in your rack. Refer to the following figure.




Figure 10. Screw pin types included with rail kit

## Install the storage slide rails into a round hole rack

Procedures for installing ThinkAgile CP storage block slide rails in a round hole rack.

To install the ThinkAgile CP storage block slide rails in a round hole rack, do the following:

**Attention:** Ensure that you install the rear mounting rail brackets before installing the front mounting brackets.




Figure 11. Installing the rail in a round hole rack

- Step 1. Open the hinge bracket on the rear of the rail.
- Step 2. Line up the pins on the rear of the rail with the corresponding holes in the rear-mounting flange of the rack, and then insert the pins into the holes in the rack flange.
- Step 3. Close the hinge bracket to secure the rail to the rack flange.

- Step 4. Secure the rear mounting bracket of the rail with two M5 screws.
- Step 5. Extend the rail until the front mounting bracket of the rail reaches the front rack.
- Step 6. Repeat steps **1** through **3** to install the front rail mounting bracket to the front mounting flange.

**Note:** Do not secure the front mounting bracket of the rail with screws until you install the storage enclosure into the rack.

- Step 7. Repeat steps **1** through **6** to install the opposite rail in the rack.

The installation of the storage block slide rails in a round hole rack is now complete.

To continue with the installation of the ThinkAgile CP storage block, follow these steps: [Install the storage enclosure](#)



## Install the storage slide rails into a square hole rack

Procedures for installing ThinkAgile CP storage block slide rails in a square hole rack.

To install the ThinkAgile CP storage block slide rails in a square hole rack, do the following:

**Attention:** Ensure that you install the rear mounting rail brackets before installing the front mounting brackets.




Figure 12. Installing the rail in a square hole rack

- Step 1. Open the hinge bracket on the rear of the rail.
- Step 2. Remove the preinstalled screw pins for round hole rack.
- Step 3. Line up the holes on the rear of the rail with the corresponding holes in the rear-mounting flange of the rack.
- Step 4. Close the hinge bracket to secure the rail to the rack flange.
- Step 5. Secure the rear mounting bracket of the rail with two screws for square hole rack.

- Step 6. Secure the rear mounting bracket of the rail with two M5 screws.
- Step 7. Extend the rail until the front mounting bracket of the rail reaches the front rack.
- Step 8. Repeat steps **1** through **5** to install the front rail mounting bracket to the front mounting flange.

**Attention:**

- For the front mounting bracket, the screws for square hole rack need to be fastened in the same holes as the preinstalled screws for round hole rack. Otherwise, the storage enclosure cannot be secured to the front mounting bracket.
- After fastening the two screws for square hole rack, do not further secure the front mounting bracket of the rail with screws until you install the storage enclosure into the rack.

- Step 9. Repeat steps **1** through **8** to install the opposite rail in the rack.

The installation of the storage block slide rails in a square hole rack is now complete.

The next step is to install the storage enclosure: [Install the storage enclosure](#)

## Install the storage enclosure


Procedures for installing ThinkAgile CP storage enclosure in a rack.

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

### CAUTION:

**Ensure that at least two people lift and install the enclosure together.**

Step 1. 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 13. Removal of side cover*

Step 2. Lift the enclosure and slowly place the enclosure on the rail.

Step 3. Carefully slide the enclosure into the rack along the rails until the enclosure is fully inserted.




Figure 14. Installing the enclosure into the rack

Step 4. Secure the enclosure with the thumbscrews and the M5 screws on the front.





Figure 15. Installing screws on front of enclosure

Step 5. 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 16. Reinstalling the side cover*

The installation of the storage block in a rack is now complete.



---

## Chapter 4. Compute block capacity upgrade

This section covers the procedures for installing additional compute nodes or compute blocks to the ThinkAgile CP solution.

The following table shows the compute block capacity upgrade options for CP4000 and CP6000:

Table 4. Compute block capacity upgrade options

Model	ThinkAgile CP4000	ThinkAgile CP6000
Initial number of compute nodes per compute block	2	4
Maximum compute nodes per compute block	4	
Initial number of compute blocks	1	
Maximum compute blocks	1	10

### Notes:

- The CP6000 can be upgraded with up to 40 compute nodes, depending on the number of compute blocks installed.
- The maximum number of compute blocks in the CP6000 solution is limited by the number of available ports on the interconnect switches.

---

## 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 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.



---

## Install the ThinkAgile CP compute block

Procedures for installing ThinkAgile CP compute block in a rack.

To install the ThinkAgile CP compute block in a rack, do the following:

1. Unpack the compute block box and check the contents
2. Install the slide rails into the rack
3. Install the enclosure on the slide rails
4. Slide the compute enclosure into the rack

Firstly, unpack the storage block box and check the contents:

- One compute block
- One compute slide rails box:
  - One left slide rail
  - One right slide rail
  - Two M5 screws
  - Two washers
- One Category 5e patch cable
- Two power cords
- Two SFP+ 10 GbE Direct Attach Cables (DACs)

**Note:** If any of the slide box items are missing or are damaged, contact your place of purchase.




Figure 17. ThinkAgile CP compute slide rails box contents

## Install the compute slide rails

Procedures for installing ThinkAgile CP compute block slide rails in a rack.

The slide rails are extensible as shown in the following illustration:




Figure 18. Slide rails

To install the ThinkAgile CP compute block slide rails, do the following:

- Step 1. Ensure that the slide rail is shortened to its shortest position.
- Step 2. From the front of the rack cabinet, line up the pins on the rear of the slide rail with the flange holes at the rear of the rack.
- Step 3. Push the slide rail so that the pins and the latch snap into place.

**Note:** The holes on the rack are either be square-shaped or round-shaped.




Figure 19. Left slide rear installation

Step 4. Pull the slide forward and insert the pins and the latch into the flange holes at the front of the rack.




Figure 20. Left slide front installation

Step 5. Repeat steps 1 through 4 to install the other slide rail. Ensure that each slide rail is securely engaged in the flange by:

1. Inspecting to confirm that the hook has caught in the rack.
2. Moving the slide rail backward and forward to ensure that it does not pop out.

Step 6. Insert a M5 screw and washer into the rear bracket of each slide rail, and then fasten the two M5 screws to secure the slide rails to the rack flanges.




Figure 21. Securing the compute block slide rails

The installation of the compute block slide rails in a rack is now complete.

To continue with the installation of the ThinkAgile CP compute block, follow these steps: [Install the compute enclosure](#)

## Install the compute enclosure

Procedures for installing ThinkAgile CP compute enclosure in a rack.

### CAUTION:

Ensure that at least two people lift and install the enclosure together.




Figure 22. Front latches on slide rails

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

- Step 1. Pull the slide rails forward until the rails click into place.
- Step 2. Push up the front latches.
- Step 3. Pull the slide rails all the way out.

**Attention:** You can only install the compute block enclosure successfully when the slide rails are fully extended.

- Step 4. Remove all compute nodes and the shuttle from the enclosure.
- Step 5. Carefully lift the enclosure.
- Step 6. Tilt and lower the enclosure; then, push the slide rails toward the enclosure and make sure the farthest nailheads on each side go into the slots on the slide rails.

Step 7. Slowly lower the enclosure and ensure that the other three nailheads on each side slip into the slots.




Figure 23. ThinkAgile CP compute enclosure installation

Step 8. To slide the ThinkAgile CP compute enclosure into the rack, first push up the rear latches on the slide rails.

Step 9. Push the enclosure until the slide rails click into place.

Step 10. Push up the front latches on the slide rails.

Step 11. Push the enclosure all the way into the rack.





Figure 24. Enclosure installation into rack

Step 12. Fasten the two thumbscrews located on the front of the compute enclosure.



*Figure 25. Compute enclosure thumbscrews*

Step 13. Insert back into the enclosure the shuttle and all the previously removed compute nodes.

The installation of the compute block in a rack is now complete.

---

## Chapter 5. Licensing considerations

The following table explains the details of the ThinkAgile CP Guardian edition.

	<b>Guardian</b>
Operating System (for compute nodes and storage controllers)	RHEL 7
Network Operating System for Interconnect Switches	Pica8 OS
At-Rest Encryption	✓
<b>Compliance</b>	
KMIP Compliant	✓
FIPS 140-2	✓
Common Criteria	✓
USGv6 (DOD IPv6)	✓
STIG	✓





---

## Chapter 6. Getting help and technical assistance

This section covers how to get help or technical assistance for the ThinkAgile CP solution.

If you need help, service, or technical assistance or just want more information about Lenovo products use the Lenovo Support Plan for ThinkAgile CP.

Details for the ThinkAgile CP Support Plan are available at the following site:

<https://datacentersupport.lenovo.com/products/solutions-and-software/thinkagile-cp/>





**Lenovo**