Lenovo ThinkAgile MX Hardware Replacement Guide

Updated: May 2025

Provides guidance for replacing hardware components in ThinkAgile MX solutions

Applies to solutions running Azure Local or Windows Server Storage Spaces Direct

Dave Feisthammel
Wayne “Guy” Fusman
Hussein Jammal
David Ye

LENOVO PRESS

Table of Contents

1 Abstract

In order to ensure availability and optimal performance of a Windows Failover Cluster, certain considerations must be made when replacing a hardware component in a cluster node. This document provides an outline to these considerations and can be used as a checklist for any special steps or processes required, depending on which hardware component is replaced. This guide applies to clusters running Windows Server or Azure Local operating systems.

At Lenovo Press, we bring together experts to produce technical publications around topics of importance to you, providing information and best practices for using Lenovo products and solutions to solve IT challenges. See our publications at http://lenovopress.com.

Do you have the latest version? We update our papers from time to time, so check whether you have the latest version of this document by clicking the **Check for Updates** button on the front page of the PDF. Pressing this button will take you to a web page that will tell you if you are reading the latest version of the document and give you a link to the latest if needed. While you're there, you can also sign up to get notified via email whenever we make an update.

2 Introduction

In order to ensure availability and optimal performance of a Windows Failover Cluster, certain considerations must be made when replacing a hardware component in a cluster node. This document provides an outline to these considerations and can be used as a checklist for any special steps or processes required, depending on which hardware component is replaced.

Always make sure to place a node into Maintenance Mode before powering it off for any reason. Even for hot-swap components, the system might need to be powered off before being able to access the component. For example, even though system fans are hot-swappable, without cable management arms to allow the system to be pulled out of the rack without powering it off, the system will need to be shutdown before unplugging all cables from it to slide it out of the rack.

This document should server as a guideline for any special steps that need to be taken before and after replacing certain hardware in a cluster node, whether the hardware is replaced by the customer (for example, a failed hot-swap data drive) or by Lenovo Support (for example, a motherboard).

3 Motherboard

The system must be powered off before replacing the motherboard. Always make sure to place the node into Maintenance Mode before powering it off for any reason. Replacing the server motherboard has special requirements related to security. These include XCC and UEFI settings that must be configured properly after motherboard replacement and before the node is returned to the Azure Local instance.

In addition, if BitLocker volume encryption is enabled, the BitLocker Recovery Key will likely be needed to unlock encrypted volumes after motherboard replacement. The method used to retrieve BitLocker Recovery Keys is different based on whether the system is running Windows Server or Azure Local.

If running Windows Server:

To check BitLocker status, use the BitLocker Configuration Tool from the command prompt (`manage-bde`) or PowerShell (`Get-BitLockerVolume C: | fl`).

If running Azure Local:

To check BitLocker status, use the BitLocker Configuration Tool from the command prompt (`manage-bde c: -status`) or PowerShell (`Get-BitLockerVolume C: | fl`).

To retrieve BitLocker Recovery Keys for an Azure Local instance node, use the following PowerShell command:

```
get-asrecoverykeyinfo
```

To suspend BitLocker on an Azure Local instance node before powering off the server, use the following PowerShell command:

```
Suspend-ASBitlockerBootVolume
```

To resume BitLocker encryption on an Azure Local instance node after motherboard replacement, use the following PowerShell command:

```
Resume-ASBitlockerBootVolume
```

The following screenshots show examples of the above commands running on an Azure Local node:

```
PS C:\Users> Get-ASRecoveryKeyInfo

ComputerName PasswordId                            RecoveryKey
------------ ----------                            -----------
J30027HK     {049380F5-AF0C-4215-941A-E1D5AC766EB2} 601766-495759-615538-594407-490754-489896-166232-369369
J30027HK     {5AF9E1B5-76E9-462E-8B48-1FDE74956A3B} 587191-321519-254529-285274-614581-108075-220198-079046
J30027HK     {C4F7AB5A-6EE6-4079-8128-941E1BFC1689} 338327-139524-038038-561858-198792-226985-036641-123013


PS C:\Users> Suspend-ASBitlockerBootVolume -Force
PS C:\Users> manage-bde c: -status
BitLocker Drive Encryption: Configuration Tool version 10.0.25398
Copyright (C) 2013 Microsoft Corporation. All rights reserved.

Volume C: []
[OS Volume]

    Size:                 237.63 GB
    BitLocker Version:    2.0
    Conversion Status:    Used Space Only Encrypted
    Percentage Encrypted: 100.0%
    Encryption Method:    XTS-AES 256
    Protection Status:    Protection Off
    Lock Status:          Unlocked
    Identification Field: Unknown
    Key Protectors:
        TPM
        Numerical Password
```

```
PS C:\Users> Resume-ASBitlockerBootVolume
PS C:\Users> Get-ASBitlockerVolume

cmdlet Get-ASBitlockerVolume at command pipeline position 1
Supply values for the following parameters:
MountPoint:
PS C:\Users>
PS C:\Users> manage-bde c: -status
BitLocker Drive Encryption: Configuration Tool version 10.0.25398
Copyright (C) 2013 Microsoft Corporation. All rights reserved.

Volume C: []
[OS Volume]

    Size:                 237.63 GB
    BitLocker Version:    2.0
    Conversion Status:    Used Space Only Encrypted
    Percentage Encrypted: 100.0%
    Encryption Method:    XTS-AES 256
    Protection Status:    Protection On
    Lock Status:          Unlocked
    Identification Field: Unknown
    Key Protectors:
        TPM
        Numerical Password
```

For additional information about BitLocker, refer to the following Microsoft article:

https://learn.microsoft.com/en-us/windows/security/operating-system-security/data-protection/bitlocker/recovery-overview

To replace the server motherboard, follow these steps:

1) Place the node into Maintenance Mode, draining all Roles from it. Do not proceed until the node has completed this step.

2) Once the motherboard has been replaced, power on the server and ensure that the system machine type and serial number have been updated on the new motherboard to match the original values.
3) Update all system firmware (XCC, UEFI, LXPM) to the same Best Recipe level that is currently in use for the rest of the cluster nodes.

   The current ThinkAgile MX Best Recipes can be found at the following URL:

   https://support.lenovo.com/us/en/solutions/HT507406

   Full sets of firmware and device drivers for each ThinkAgile MX solution that comply with the last several Best Recipes can be downloaded from the following URL:

   https://thinkagile.lenovo.com/mx

4) Once the system firmware has been updated, certain XCC and UEFI settings must be configured to match the other nodes in the cluster. For details on OCC settings, refer to the following URL: https://pubs.lenovo.com/thinkagile-mx/mx_sbe_configure_xcc
5) The following UEFI settings must be restored after motherboard replacement:
   - Operating Mode = Maximum Performance
   - TPM = 2.0
   - Secure Boot = Enabled
   - Assert RPP (V1 and V2 systems)
6) Once the XCC and UEFI settings have been set properly, Resume the node and restore any cluster Roles that were drained when placing the system into Maintenance Mode.

4 Network Adapter

The system must be powered off before replacing a network adapter. Always make sure to place the node into Maintenance Mode before powering it off for any reason. To replace a network adapter, follow these steps:

1) Start by determining whether the network adapter being replaced is part of a Switch Embedded Team (SET) or Hyper-V Virtual Switch. It is also important to know if Network ATC has been implemented. If any of these are true, additional configuration steps will need to be taken after replacement.
2) Place the node into Maintenance Mode, draining all Roles from it. Do not proceed until the node has completed this step.
3) Replace the network adapter. Note that the replacement must be exactly the same model.
4) Power on the server and update the network adapter firmware to the same Best Recipe level that is currently in use for the rest of the cluster nodes.

   The current ThinkAgile MX Best Recipes can be found at the following URL:

   https://support.lenovo.com/us/en/solutions/HT507406

   Full sets of firmware and device drivers for each ThinkAgile MX solution that comply with the last several Best Recipes can be downloaded from the following URL:

   https://thinkagile.lenovo.com/mx

5) Once the firmware for the new network adapter has been updated, check the network adapter device driver in Device Manager to ensure that the proper driver is shown and that an inbox driver is not being used.
6) If the network adapter being replaced was part of a team, hosted a Hyper-V switch, or was configured via intent-based network ATC, complete the appropriate step below:
   a) If the network adapter was part of a SET, which is common for Storage traffic, its port(s) will need to be removed from the team and the new port(s) will need to be added to the team using the same team name(s) and IP address(es). For more information, refer to the following Microsoft articles.

      Remove team member: https://learn.microsoft.com/en-us/powershell/module/hyper-v/remove-vmswitchteammember?view=windowsserver2022-ps

      Add team member: https://learn.microsoft.com/en-us/powershell/module/hyper-v/add-vmswitchteammember?view=windowsserver2022-ps

   b) If the network adapter hosted a Hyper-V virtual switch, which is required for Compute traffic, the virtual switch will need to be recreated using the same name as the previous virtual switch.
   c) If the network adapter was configured via Network ATC, ensure that the intent for the replaced adapter matches the original. For more information, refer to the following Microsoft article:

      https://learn.microsoft.com/en-us/azure-stack/hci/manage/manage-network-atc?tabs=22H2

7) Resume the node and restore any cluster Roles that were drained when placing the system into Maintenance Mode.

5 RAID Adapter for boot volume

The system must be powered off before replacing a RAID adapter. Always make sure to place the node into Maintenance Mode before powering it off for any reason. The RAID adapter present in any ThinkAgile MX solution supports only the OS boot volume and not the storage devices that make up the solution storage pool.

If BitLocker volume encryption is enabled for the boot drive, it must be suspended before replacing a failed RAID adapter and then resumed after hardware replacement. The method used to retrieve BitLocker Recovery Keys is different based on whether the system is running Windows Server or Azure Local.

If running Windows Server:

To check BitLocker status, use the BitLocker Configuration Tool from the command prompt (`manage-bde`) or PowerShell (`Get-BitLockerVolume C: | fl`).

If running Azure Local:

To check BitLocker status, use the BitLocker Configuration Tool from the command prompt (`manage-bde c: -status`) or PowerShell (`Get-BitLockerVolume C: | fl`).

To retrieve BitLocker Recovery Keys for an Azure Local instance node, use the following PowerShell command:

```
get-asrecoverykeyinfo
```

To suspend BitLocker on an Azure Local instance node before powering off the server, use the following PowerShell command:

```
Suspend-ASBitlockerBootVolume
```

To resume BitLocker encryption on an Azure Local instance node after motherboard replacement, use the following PowerShell command:

```
Resume-ASBitlockerBootVolume
```

The following screenshots show examples of the above commands running on an Azure Local node:

```
PS C:\Users> Get-ASRecoveryKeyInfo

ComputerName PasswordId                           RecoveryKey
------------ ----------                           -----------
J30027HK     {049380F5-AF0C-4215-941A-E1D5AC766EB2} 601766-495759-615538-594407-490754-489896-166232-369369
J30027HK     {5AF9E1B5-76E9-462E-8B48-1FDE74956A3B} 587191-321519-254529-285274-614581-108075-220198-079046
J30027HK     {C4F7AB5A-6EE6-4079-8128-941E1BFC1689} 338327-139524-038038-561858-198792-226985-036641-123013


PS C:\Users> Suspend-ASBitlockerBootVolume -Force
PS C:\Users> manage-bde c: -status
BitLocker Drive Encryption: Configuration Tool version 10.0.25398
Copyright (C) 2013 Microsoft Corporation. All rights reserved.

Volume C: []
[OS Volume]

    Size:                 237.63 GB
    BitLocker Version:    2.0
    Conversion Status:    Used Space Only Encrypted
    Percentage Encrypted: 100.0%
    Encryption Method:    XTS-AES 256
    Protection Status:    Protection Off
    Lock Status:          Unlocked
    Identification Field: Unknown
    Key Protectors:
        TPM
        Numerical Password
```

```
PS C:\Users> Resume-ASBitlockerBootVolume
PS C:\Users> Get-ASBitlockerVolume

cmdlet Get-ASBitlockerVolume at command pipeline position 1
Supply values for the following parameters:
MountPoint:
PS C:\Users>
PS C:\Users> manage-bde c: -status
BitLocker Drive Encryption: Configuration Tool version 10.0.25398
Copyright (C) 2013 Microsoft Corporation. All rights reserved.

Volume C: []
[OS Volume]

    Size:                 237.63 GB
    BitLocker Version:    2.0
    Conversion Status:    Used Space Only Encrypted
    Percentage Encrypted: 100.0%
    Encryption Method:    XTS-AES 256
    Protection Status:    Protection On
    Lock Status:          Unlocked
    Identification Field: Unknown
    Key Protectors:
        TPM
        Numerical Password
```

For additional information about BitLocker, refer to the following Microsoft article:

https://learn.microsoft.com/en-us/windows/security/operating-system-security/data-protection/bitlocker/recovery-overview

After suspending BitLocker volume encryption, follow these steps to replace a RAID adapter:

1) Place the node into Maintenance Mode, draining all Roles from it. Do not proceed until the node has completed this step.

2) Replace the RAID adapter. Note that the replacement must be exactly the same model.
3) Power on the server and update the RAID adapter firmware to the same Best Recipe level that is currently in use for the rest of the cluster nodes.

   The current ThinkAgile MX Best Recipes can be found at the following URL:

   https://support.lenovo.com/us/en/solutions/HT507406

   Full sets of firmware and device drivers for each ThinkAgile MX solution that comply with the last several Best Recipes can be downloaded from the following URL:

   https://thinkagile.lenovo.com/mx

4) Once the firmware for the new RAID adapter has been updated, Resume the node and restore any cluster Roles that were drained when placing the system into Maintenance Mode.

6 Boot drive

Most ThinkAgile MX solutions have boot drives that are internal to the server and, therefore, are not hot-swappable. Even if your solution uses externally accessible hot-swap boot drives, the server must be rebooted to reestablish the RAID-1 mirror across the two boot drives. Always make sure to place the node into Maintenance Mode before powering it off for any reason.

If BitLocker volume encryption is enabled for the boot drive, it must be suspended before replacing a failed boot drive and then resumed after drive replacement. The method used to retrieve BitLocker Recovery Keys is different based on whether the system is running Windows Server or Azure Local.

If running Windows Server:

To check BitLocker status, use the BitLocker Configuration Tool from the command prompt (`manage-bde`) or PowerShell (`Get-BitLockerVolume C: | fl`).

If running Azure Local:

To check BitLocker status, use the BitLocker Configuration Tool from the command prompt (`manage-bde c: -status`) or PowerShell (`Get-BitLockerVolume C: | fl`).

To retrieve BitLocker Recovery Keys for an Azure Local instance node, use the following PowerShell command:

```
get-asrecoverykeyinfo
```

To suspend BitLocker on an Azure Local instance node before powering off the server, use the following PowerShell command:

```
Suspend-ASBitlockerBootVolume
```

To resume BitLocker encryption on an Azure Local instance node after motherboard replacement, use the following PowerShell command:

```
Resume-ASBitlockerBootVolume
```

The following screenshots show examples of the above commands running on an Azure Local node:

```
PS C:\Users> Get-ASRecoveryKeyInfo

ComputerName PasswordId                           RecoveryKey
------------ ----------                           -----------
J30027HK     {049380F5-AF0C-4215-941A-E1D5AC766EB2} 601766-495759-615538-594407-490754-489896-166232-369369
J30027HK     {5AF9E1B5-76E9-462E-8B48-1FDE74956A3B} 587191-321519-254529-285274-614581-108075-220198-079046
J30027HK     {C4F7AB5A-6EE6-4079-8128-941E1BFC1689} 338327-139524-038038-561858-198792-226985-036641-123013


PS C:\Users> Suspend-ASBitlockerBootVolume -Force
PS C:\Users> manage-bde c: -status
BitLocker Drive Encryption: Configuration Tool version 10.0.25398
Copyright (C) 2013 Microsoft Corporation. All rights reserved.

Volume C: []
[OS Volume]

    Size:                 237.63 GB
    BitLocker Version:    2.0
    Conversion Status:    Used Space Only Encrypted
    Percentage Encrypted: 100.0%
    Encryption Method:    XTS-AES 256
    Protection Status:    Protection Off
    Lock Status:          Unlocked
    Identification Field: Unknown
    Key Protectors:
        TPM
        Numerical Password
```

```
PS C:\Users> Resume-ASBitlockerBootVolume
PS C:\Users> Get-ASBitlockerVolume

cmdlet Get-ASBitlockerVolume at command pipeline position 1
Supply values for the following parameters:
MountPoint:
PS C:\Users>
PS C:\Users> manage-bde c: -status
BitLocker Drive Encryption: Configuration Tool version 10.0.25398
Copyright (C) 2013 Microsoft Corporation. All rights reserved.

Volume C: []
[OS Volume]

    Size:                 237.63 GB
    BitLocker Version:    2.0
    Conversion Status:    Used Space Only Encrypted
    Percentage Encrypted: 100.0%
    Encryption Method:    XTS-AES 256
    Protection Status:    Protection On
    Lock Status:          Unlocked
    Identification Field: Unknown
    Key Protectors:
        TPM
        Numerical Password
```

For additional information about BitLocker, refer to the following Microsoft article:

https://learn.microsoft.com/en-us/windows/security/operating-system-security/data-protection/bitlocker/recovery-overview

After suspending BitLocker volume encryption, follow these steps to replace a boot drive:

1) Place the node into Maintenance Mode, draining all Roles from it. Do not proceed until the node has completed this step.

2) Replace the boot drive. Note that the replacement must be exactly the same model.
3) Reboot the server and update the storage device firmware to the same Best Recipe level that is currently in use for the rest of the cluster nodes.

   The current ThinkAgile MX Best Recipes can be found at the following URL:

   https://support.lenovo.com/us/en/solutions/HT507406

   Full sets of firmware and device drivers for each ThinkAgile MX solution that comply with the last several Best Recipes can be downloaded from the following URL:

   https://thinkagile.lenovo.com/mx

4) Once the firmware for the new storage device has been updated, Resume the node and restore any cluster Roles that were drained when placing the system into Maintenance Mode.

7 HBA

The system must be powered off before replacing an HBA. Always make sure to place the node into Maintenance Mode before powering it off for any reason. The HBA present in any ThinkAgile MX solution supports the storage devices that make up the solution storage pool. To replace an HBA, follow these steps:

5) Place the node into Maintenance Mode, draining all Roles from it. Do not proceed until the node has completed this step.
6) Replace the HBA. Note that the replacement must be exactly the same model.
7) Power on the server and update the HBA firmware to the same Best Recipe level that is currently in use for the rest of the cluster nodes.

   The current ThinkAgile MX Best Recipes can be found at the following URL:

   https://support.lenovo.com/us/en/solutions/HT507406

   Full sets of firmware and device drivers for each ThinkAgile MX solution that comply with the last several Best Recipes can be downloaded from the following URL:

   https://thinkagile.lenovo.com/mx

8) Once the firmware for the new HBA has been updated, Resume the node and restore any cluster Roles that were drained when placing the system into Maintenance Mode.

8 Data (pool) drive

Most ThinkAgile MX solutions use externally accessible hot-swap data drives. However, for the following solutions, all storage devices are internal to the server and are not hot-swappable:

- ThinkAgile MX1020 on SE350
- ThinkAgile MX1021 on SE350
- ThinkAgile MX450 Edge IS

For these solutions, contact Lenovo Support to have the failed device replaced. The node must be placed into Maintenance Mode before it is powered off to replace the storage device.

To learn about the various health and operational states of storage pools, virtual disks, and physical drives, refer to the following Microsoft article:

https://learn.microsoft.com/en-us/windows-server/storage/storage-spaces/storage-spaces-states

Once the determination has been made that a physical drive must be replaced, use the following process to locate the physical location of the failed data drive.

1) The process to locate a failed drive is different for HDD/SDD devices than for NVMe devices.
   a) For a failed HDD or SSD device, determine which storage device must be replaced and blink the physical disk locator LED to aid in locating the correct device in a node. This can be achieved using the following PowerShell commands running as Administrator.

      To find a failed drive on a specific cluster node:

```
Get-StorageScaleUnit -FriendlyName <NodeName> | Get-PhysicalDisk
```

```
PS C:\> Get-StorageScaleUnit -FriendlyName Lenovo-S1-N01 | Get-PhysicalDisk
Number FriendlyName           SerialNumber MediaType CanPool OperationalStatus  HealthStatus Usage          Size
------ ------------           ------------ --------- ------- -----------------  ------------ -----          ----
4011   ATA HUH721212ALE600    8DHH8N0H     HDD       FALSE   OK                 Healthy      Auto-Select 10.91 TB
4009   ATA HUH721212ALE600    8DHE7DRH     HDD       FALSE   OK                 Healthy      Auto-Select 10.91 TB
4007   LENOVO HUSMM3216ASS200 4GV063LX     SSD       FALSE   OK                 Healthy      Journal      1.46 TB
4014   ATA HUH721212ALE600    8DHG4ZEH     HDD       FALSE   IO error           Warning      Auto-Select 10.91 TB
4001   LENOVO HUSMM3216ASS200 4GV06AVX     SSD       FALSE   OK                 Healthy      Journal      1.46 TB
4002   ATA HUH721212ALE600    8DHHL75H     HDD       FALSE   OK                 Healthy      Auto-Select 10.91 TB
4005   ATA HUH721212ALE600    8DHG4STH     HDD       FALSE   OK                 Healthy      Auto-Select 10.91 TB
4012   ATA HUH721212ALE600    8DHH1RKH     HDD       FALSE   OK                 Healthy      Auto-Select 10.91 TB
4010   LENOVO HUSMM3216ASS200 4GV03Y7X     SSD       FALSE   OK                 Healthy      Journal      1.46 TB
4003   ATA HUH721212ALE600    8DHHJELH     HDD       FALSE   OK                 Healthy      Auto-Select 10.91 TB
4006   ATA HUH721212ALE600    8DHH7EXH     HDD       FALSE   OK                 Healthy      Auto-Select 10.91 TB
4013   ATA HUH721212ALE600    8DHHH9ZH     HDD       FALSE   OK                 Healthy      Auto-Select 10.91 TB
4008   ATA HUH721212ALE600    8DHHNAKH     HDD       FALSE   OK                 Healthy      Auto-Select 10.91 TB
4004   LENOVO HUSMM3216ASS200 AGV063VX     SSD       FALSE   OK                 Healthy      Journal      1.46 TB
```

      If the name of the cluster node with the failed drive is not known, remove the “-FriendlyName” parameter from the command above. This will result in a list of all drives on all nodes in the cluster. Once the serial number of the failed drive has been obtained, use it to determine which node contains the failed drive using the following command:

```
PS C:\> Get-PhysicalDisk -SerialNumber <SerialNumber> | Get-StorageScaleUnit
```

```
PS C:\> Get-PhysicalDisk -SerialNumber 8DHG4ZEH | Get-StorageScaleUnit

Type             FriendlyName       SerialNumber       PhysicalLocation       HealthStatus   OperationalStatus
----             ------------       ------------       ----------------       ------------   -----------------
StorageScaleUnit Lenovo-S1-N01      J10002E8                                  Warning        IO error
```

Now use the failed drive serial number to blink the physical disk locator LED to aid in locating the correct device in the node using the following command:

```
Get-PhysicalDisk -SerialNumber <SerialNumber> | Enable-PhysicalDiskIdentification
```

b) It is not possible to blink a locator LED for NVMe devices. In the case of a failed NVMe device or for ThinkAgile MX Edge Solutions based on SE350 and SE450 servers, which use internal non hot-swap drives, the same PowerShell commands are used to determine the serial number of the device that has failed. However, in the case of NVMe devices, the “-SerialNumber” parameter returns a GUID-like string assigned to each device by the operating system. To find the actual physical device serial number, use the following command:

```
Get-PhysicalDisk | FT FriendlyName,Usage,MediaType,SerialNumber,AdapterSerialNumber
```

```
PS C:\> Get-PhysicalDisk | FT FriendlyName,Usage,MediaType,SerialNumber,AdapterSerialNumber

FriendlyName              Usage       MediaType SerialNumber                                AdapterSerialNumber
------------              -----       --------- ------------                                -------------------
Micron_7450_MTFDKCB3T8TFR Auto-Select SSD       0000_0000_0000_0001_00A0_7524_471C_DEBC.    2408471CDEBC
Micron_7450_MTFDKCB3T8TFR Auto-Select SSD       0000 0000 0000 0001 00A0 7524 471C DD69.    2408471CDD69
Micron_7450_MTFDKCB3T8TFR Auto-Select SSD       0000_0000_0000_0001_00A0_7524_471C_DDB2.    2408471CDDB2
Micron_7450_MTFDKCB3T8TFR Auto-Select SSD       0000_0000_0000_0001_00A0_7524_471C_DC86.    2408471CDC86
BROADCOM RAID B540i-2i    Auto-Select SSD       00b52e639093a9562e00700d0c2ed245
Micron_7450_MTFDKCB3T8TFR Auto-Select SSD       0000_0000_0000_0001_00A0_7524_471C_F9F6.    2408471CF9F6
Micron_7450_MTFDKCB3T8TFR Auto-Select SSD       0000_0000_0000_0001_00A0_7524_471C_FA8C.    2408471CFA8C
```

In the example above, the green box highlights the device serial number from the OS and the yellow box contains the actual physical device serial number assigned by the manufacturer. This is the serial number that is used to determine the physical location of the failed device using the XCC browser interface. To do this, open the XCC browser interface to the node containing the failed device. In the left navigation pane click Inventory and on the right side of the interface, select DISK. This will show all storage devices installed in the node. Find the failed device using the “AdapterSerialNumber” parameter from the command above and note the Bay or Slot number of the device.

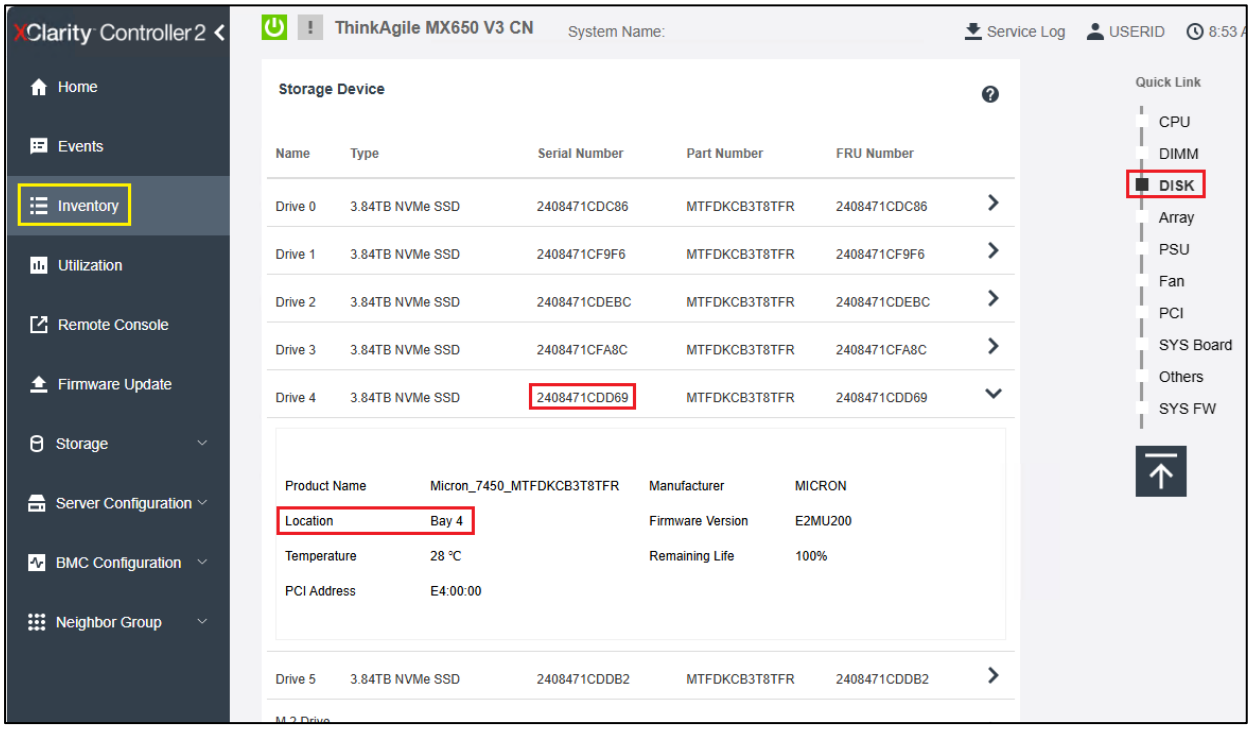

2) The failed drive can now be replaced. The storage recovery process is automatic and should begin immediately. Although no further manual intervention is required, it is best practice to use the following commands to retire and remove the failed drive from the storage pool. Otherwise, the failed device could continue to be shown in the pool. Note that the storage pool name shown on the second line below should be “*SU1*” for an Azure Local instance or “*S2D*” for an S2D cluster.

```
$Disk = Get-PhysicalDisk | ? HealthStatus -EQ “Unhealthy”
$Pool = Get-StoragePool *SU1*
Set-PhysicalDisk -InputObject $Disk -Usage “Retired”
Remove-PhysicalDisk -StoragePool $Pool -PhysicalDisks $Disk
```

You can check the status of the storage rebalance process using the following PowerShell command:

```
Get-StorageJob
```

3) To ensure the replacement drive gets added into the storage pool, use the following commands. Again, for an Azure Local instance, use “*SU1*” for the storage pool name and for an S2D cluster, use “*S2D*” in the second command below.

```
$addedDisk = Get-PhysicalDisk |? canpool -like true
Get-StoragePool *SU1* | add-PhysicalDisk -PhysicalDisks $addedDisk
```

You can also check the progress of any virtual disk repairs that are required using the following PowerShell command:

```
Get-VirtualDisk
```

The replacement process is complete. If the node had to be powered off to get to the drives, Resume the node and restore any cluster Roles that were drained when placing the system into Maintenance Mode.

9 Power supply

For solutions that use rack servers (for example, the SR650 V3) that are configured with redundant power, there is no need to slide the server out of the rack to replace a power supply. Therefore, the process can be done without powering off the server. Simply remove the power cord from the impacted power supply, replace the power supply, and reconnect the power cord.

For rack servers that are not configured with redundant power and for solutions that use edge servers (for example, the SE350), the system must be powered off to replace a power supply. In this scenario the node should be placed into Maintenance Mode before powering off to replace a power supply. No solution-specific configuration is required after replacement.

10 CPU or Memory DIMM

The system must be powered off before replacing a CPU or memory DIMM. Always make sure to place the node into Maintenance Mode before powering it off for any reason. To replace a CPU or memory DIMM, follow these steps:

1) Place the node into Maintenance Mode, draining all Roles from it. Do not proceed until the node has completed this step.
2) Replace the CPU or memory DIMM. Note that the replacement must be exactly the same model.
3) Power on the server and Resume the node, returning any Roles that were drained when placing the system into Maintenance Mode. No solution-specific configuration is required after replacement of a CPU or memory DIMM.

11 Fan

For system fans, although most are hot-swappable, if cable management arms are not used, the system will need to be shutdown in order to slide it out of the rack to replace a fan. Always make sure to place the node into Maintenance Mode before powering it off for any reason. No solution-specific configuration is required after replacement.