Troubleshoot upgrade failure due to vMotion error for VMs with vGPU

Contents

Introduction

This document describes how to troubleshoot upgrade failure due to vMotion error for VMs with vGPU.

Prerequisites

Hyperflex/Esxi cluster with inconsistent ECC configuration for Nvidia GPUs.

> **Note**: Confirm system is not affected by [CSCvp47724](#)

Requirements

vCenter cluster with Nvidia GPU enabled for VMs.

Components Used

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

- Intersight (SaaS in this case)
- HyperFlex 5.0(2a)
- Nvidia GPU (Tesla T4)

Background information

In this specific scenario, vMotion was failing because of issues with inconsistent ECC configuration for Nvidia GPUs, which was causing error while upgrading HyperFlex.
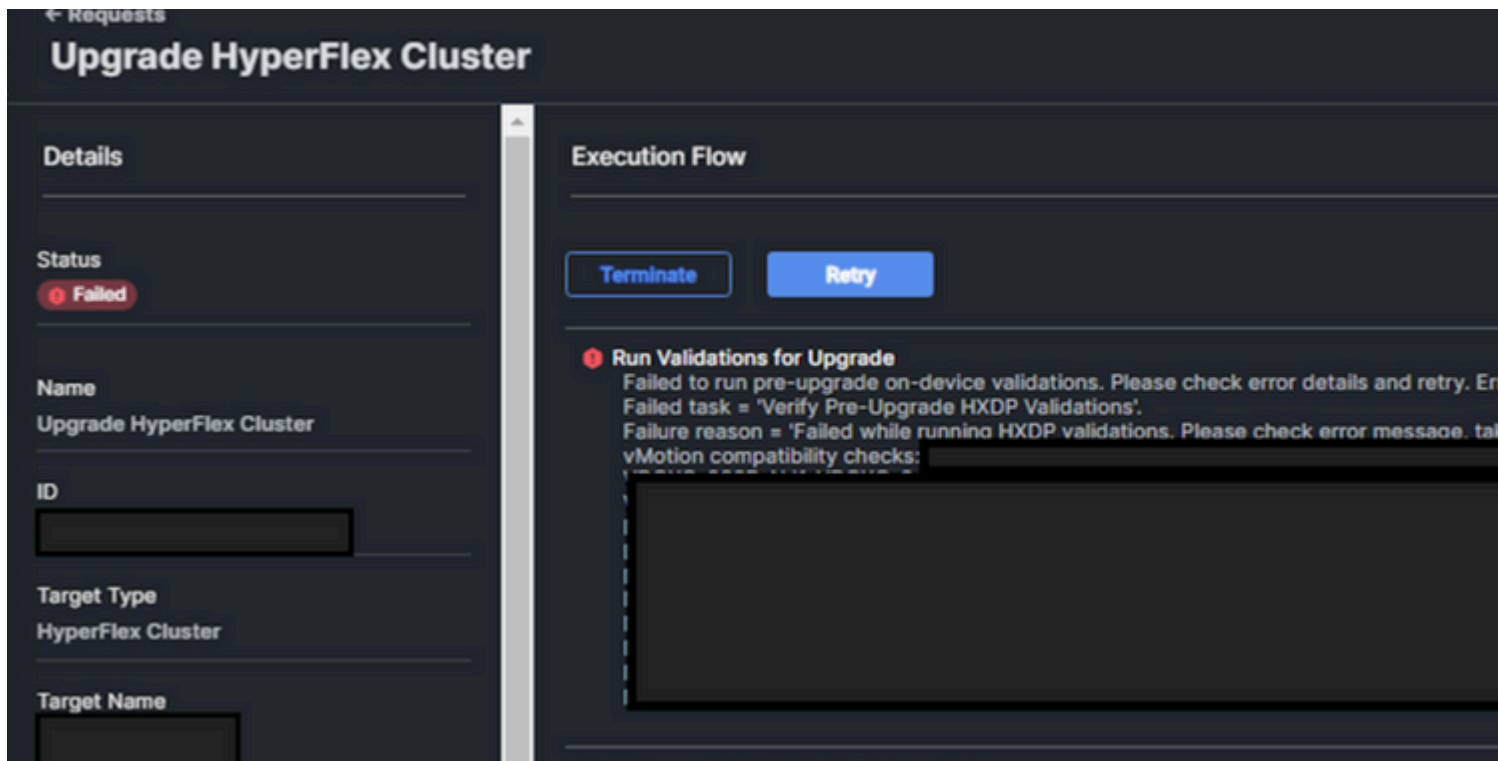
> **Note**: NVIDIA GPU Cards that use the Pascal architecture, such as Tesla V100, P100, P40, as well as the Tesla M6 and M60 GPUs, support ECC memory for improved data integrity. However, the

NVIDIA vGPU software does not support ECC. You must therefore ensure that ECC memory is disabled on all GPUs when using NVIDIA vGPU.

Problem

vMotion failure identified due to inconsistent ECC configuration on the Nvidia GPU.

Failed task : 'Verify Pre-Upgrade HXDP Validations'.

Solution

Disable ECC mode on the affected node.

How was the affected node identified?

Manual vMotion will display error for the node *"Error while migrating VMs to this node : "One or more devices (pciPassthru0) required by VM XXXX are not available on host XXXX"*

Steps:

- **List the VIB that was installed in the ESXi Hypervisor**

```
# esxcli software vib list |grep â€"i NVIDIA
```

- **Check the NVIDIA Driver Operation**

```
[root@hxesxi:~] nvidia-smi
Sat Jul 22 01:31:42 2023
+-----------------------------------------------------------------------------+
| NVIDIA-SMI 470.182.02   Driver Version: 470.182.02   CUDA Version: N/A      |
|-------------------------------+----------------------+----------------------+
| GPU  Name        Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|         Memory-Usage | GPU-Util  Compute M. |
|                               |                      |               MIG M. |
|===============================+======================+======================|
|   0  Tesla T4            On   | 00000000:AF:00.0 Off |                    0 |
| N/A   35C    P8    16W /  70W |   1971MiB / 15359MiB |      0%      Default |
|                               |                      |                  N/A |
+-------------------------------+----------------------+----------------------+
```

- **Found ECC (Error Correcting Code) mode enabled on the affected node.**

```
# nvidia-smi â€"q
ECC Mode
Current                    : Enabled
Pending                    : Enabled
```

- **Disable Error Correcting Code (ECC)**

```
# nvidia-smi â€"e 0
Disabled ECC support for GPU 0000â€¦.
All done.
```

- **Check that the ECC mode is disabled:**

```
# nvidia-smi â€"q
ECC Mode
Current                    : Disabled
Pending                    : Disabled
```

After disabling ECC mode and rebooting the node, vMotion will succeed and upgrade will progress.

Related information

[Nvidia - Insufficient resources. One or more devices](#)

[VMware - Using GPUs with Virtual Machines on vSphere](#)