

Troubleshoot Temporary Storage Path Loss on ESXi 8.0 with NFNIC Driver

Contents

[Introduction](#)

[Problem](#)

[Solution](#)

[Workaround](#)

[Cisco's response](#)

Introduction

This document address the number of increased cases logged at both Cisco and Broadcom related to Cisco nfnic driver behavior and Broadcom new FPIN (Fabric Performance Impact Notifications) architecture in release 8.0, this article is written to address concerns.

Problem

FPIN (Fabric Performance Impact Notifications) capability was added to ESXi 8.0 U2 to be able to better understand fabric related issues. Due to a bug in the StorageFPIN code, when FPIN tries to allocate memory and is unable to, it can hold onto a reference count on the paths which prevents the Cisco NFNIC driver from being able to allocate new paths or re-establish existing ones.

Reference:

See [Broadcom KB](#)

FPIN (Fabric Performance Impact Notifications) capability was added to ESXi 8.0 to be able to better understand fabric related issues. Due to a bug in the StorageFPIN code, when FPIN tries to allocate memory and is unable to, it can hold onto a reference count on the paths which prevents the Cisco NFNIC driver from being able to allocate new paths or re-establish existing ones.

This is a known issue with both FPIN as well as how the Cisco NFNIC driver is coded to behave when there are path losses. The NFNIC driver does not save storage port bindings so when a storage path re-establishes after an outage or path loss, it simply create brand new paths and increment target numbers. Because of the bug with FPIN keeping a reference count on those paths, the Cisco NFNIC driver is eventually unable to establish new paths.

A code fix to alter the FPIN open reference count behavior is going to be available in an upcoming ESXi 8.x release.

Solution

Refer to [Broadcom KB](#) article for the workaround fix. And when the ESXi patch is available, apply that patch as the solution for long term fix.

Workaround


To workaround this issue, it is recommended to disable FPIN on ESXi 8.0 hosts, especially when using Cisco UCS and NFNIC:

esxcli storage fpin info set -e false

To confirm the setting:

esxcli storage fpin info get

Aside from this Broadcom recommended change, reboot the host recover all storage paths if storage is behaving properly.



Note: This change does not require a reboot on its own. However, if an ESXi host is already in a memory heap exhaustion state for storageFPINHeap, then rebooting the host is required after this setting change.


Cisco's response

Our nfnic driver has always incremented target ID number on every target disconnect/connect. This incrementing target ID number on current and prior NFNIC driver versions is what exposed the memory leak condition in the new ESXi FPIN feature.

Additionally, the issue mentioned in the article is an ESXi OS bug which going to be fixed in an upcoming ESXI release. The article also mentions [Cisco bug ID CSCwn00553](#) which tracks a different issue and the nfnic driver fix to [Cisco bug ID CSCwn00553](#) is not be recommended to resolve the ESXi issue mentioned in the Broadcom KB article.

The VMware KB article is indicating that a Cisco bug fix is required as well as their FPIN fix. This is incorrect and this additional statement can be provided.

Broadcom is going to deliver a fix for the FPIN issue which is going to be available in the upcoming release of a 8.0.U3 patch. Once Broadcom releases the FPIN fix, the current VIC drivers work for FPIN.



Note: Meanwhile, NFNIC driver, and its behavior around creation of target-ID. This implementation on NFNIC with respect to target-ID has been VIC day one behavior and a change

in this behavior is not required for the FPIN functionality once VMware fix is available.

Reference [Cisco bug ID CSCwn00553](#)