



[Home](#) [Archives](#) [Profile](#) [Subscribe](#)

[Main](#)

11/04/2010

Nexus 5000 / 2232PP configuration

As I've mentioned in my previous posts, I work for the Connectivity team in EMC's E-Lab. Our primary responsibility is to set up configurations that our customers are likely to use, enable many of the more useful features and then see what kind of trouble we can cause by performing disruptive tests like cable pulls, inserting errors with a jammer, switch reboots during I/O, etc. After I complete a round of testing, I know I've done a good job if I've found an error condition that a device can't recover from.

Once a configuration is set up and we're satisfied that it's operating correctly, we document exactly what we tested as well as any test results and caveats. If the configuration contains a new piece of technology, we create a new case study that will be published in the [FCoE Tech book](#). We publish this information so our customers will be able to get a particular piece of equipment into production with as little frustration as possible. The idea is, if a customer has a miserable time installing something or installs it incorrectly, all of the work done to ensure that the device is as robust as possible won't provide much value.

I mention all of this because we recently added a new case study that shows how to incorporate the Nexus 2232PP into an FCoE environment. The configuration that was set up and tested is shown below.



[Erik Smith](#)

[ARCHIVES](#)

[September 2022](#)

[August 2020](#)

[April 2019](#)

[December 2017](#)

[November 2017](#)

[July 2017](#)

[April 2017](#)

[September 2016](#)

[July 2016](#)

[February 2016](#)

SEARCH


Search

[CATEGORIES](#)

[Subscribe to this blog's feed](#)

DISCLAIMER

This is not an official EMC blog. The opinions expressed here are my personal opinions. Content published here is not read or approved in advance by EMC and does not necessarily reflect the views and opinions of EMC nor does it constitute any official communication of EMC.



The step-by-step instructions for setting up this topology are documented in the EMC [FCoE tech book](#), so I will not repeat them here. However, it may be helpful to provide a brief overview of the topology components, summarize the setup steps and then discuss some of the caveats specific to this topology.

Topology elements:

The topology consists of (2) Nexus 5020s, (2) Nexus 2232PP Fabric Extenders, a host with a dual ported CNA, (2) different FC SANs and a LAN.

Nexus 5020:

The Nexus 5020 is Cisco's 10GbE switch that also happens to contain a Fibre Channel Forwarder (FCF). As you would expect, there is a ton of information on Cisco's website that covers every aspect of the 5020 in detail. One of the more interesting documents covers the [Nexus 5000 architecture](#). In the above topology, each Nexus 5020 is connected to a different SAN in order to maintain EMC's SAN A / SAN B best practice for redundancy. The link between one of the Nexus 5020s and the Cisco SAN consists of TE_Ports and these ports are part of a SAN Port Channel. The 5020 connected to the Brocade switch is in NPV mode. Each 5020 is also connected to the LAN allowing for the non-FCoE frames to be forwarded to their ultimate destination.

Nexus 2232PP:

The Nexus 2232PP is Cisco's Fabric Extender (FEX) product. Again there is a ton of information available on Cisco's website for this product. A few interesting points about the 2232PP include:

- It's managed from the Nexus 5000 that it's connected to. In other words interface display, configuration, etc is done from the 5000.
- Firmware is automatically downloaded to it from the 5000.
- It does not perform any local switching. All frames are forwarded to the Nexus 5000 it is attached to before a forwarding decision is made. Given the actual topologies that the Nexus 2232PP is used in, I don't see this as being a huge problem. (Note: I plan to cover realistic FCoE topologies in the near future)
- Worst case oversubscription is 4:1
- It uses NIV. [Scott Lowe has done an outstanding job](#) of describing what this technology is and how it works, so I won't dive into those details here.
- It is capable of using lower cost FET (Fabric Extender Transceivers) for connectivity to the Nexus 5000.
- It can be simultaneously connected to two Nexus 5000s to create an Active-Active topology but not if you want to use FCoE!

Host and CNA:

The CNA connections are spread across both 2232PPs to allow for redundancy. In our testing, PowerPath was used for failover and load balancing.

Setup steps:

Starting from scratch, the setup steps used to create this topology are summarized as follows:

On each 5000 in the topology:

1. Configure the switch management interface, basic settings and password
2. Enable the FCoE, LACP, FEX, LLDP and NPV (if applicable) features
3. Create the appropriate VSANs
4. Create the appropriate VLANs
5. Create and configure a FEX instance
6. Create and configure the FEX EtherChannel
7. Create vFC interfaces and bind them to the ENode MAC of the CNA
8. Assign each vFC to the appropriate VSAN
9. Configure Ethernet uplink settings
10. Add FC interfaces to the appropriate VSAN (if applicable)
11. Connect cables and verify login
12. Save the configuration
13. Perform Fabric zoning
14. Provision storage

Again, for more details refer to the "Nexus 2000 Supported features and topologies" as well as the "Nexus 50x0/Nexus 2232PP Topology" sections of the

[FCoE Tech book](#).

Caveats:

Overall, I found that this configuration works well. The complaints I do have center around the need to bind the vFC to the ENode MAC address of the CNA. With some CNAs, this can be a relatively straight-forward process that involves opening the CNA Management tool and locating the ENode MAC. On others, you will need to take note of both MAC addresses when you are installing the CNA into the server and then bind one of these to the vFC. The trick is choosing the correct MAC address because choosing the wrong one will prevent the CNA from successfully logging in! I've found a couple of different ways to determine the ENode MAC (described below) that work well in the lab, but I would really like it if Cisco would provide a "show mac enode" command to make things simpler.

How to find the ENode MAC:

1. As I mentioned above, the most reliable way is via the CNAs management tool. However, this information is not always available via this method.
2. Another method is to create a vFC and bind it to a physical interface on the 5000 (i.e. interface ethernet 1/1). Next, attach the CNA in question to the interface on the 5000 and use "show fcoe database". Assuming that the CNA logs in, the ENode MAC will be displayed under the MAC Address column. You can then bind this MAC Address to the vFC that you actually want to use.
3. If you took note of both MAC Addresses when you installed the CNA into the host, you can use "show mac address-table" on the 5000. The ENode MAC is the MAC Address that is **NOT** shown in this table. For example, let's say that during the installation of the CNA, you record MAC Address A and B. Let's also assume that when you use the "show mac address-table" command on the 5000, MAC Address A is present in the list. If this is the case, then MAC Address B is the ENode MAC.
4. If you didn't take note of the MAC Address on the CNA before you installed it, and you are unable or unwilling to power down the server, pull the CNA and look at the label, you can always take an educated guess. This can be done because all of the CNAs that I am familiar with have the ENode MAC set to be one or two numbers higher or lower than the MAC Address that is displayed via "show mac address-table". For example let's say that "show mac address-table" shows you a MAC Address of "0000.976f.c16a", the ENode MAC is going to be one of the following "0010.186f.c168", "0010.186f.c169", "0010.186f.c16b" or "0010.186f.c16c". With this knowledge in hand, you can experiment and try each one until you see a valid FCoE login. It isn't pretty but it works in a pinch. Also, I advise trying the MAC Address that is one number higher (i.e. 0010.186f.c16b) first.

If you made it this far, thanks for reading! Hopefully you've found it informative and ideally helpful. In early December, look for a similar post on Cisco's vPC feature that provides an alternative to spanning tree and improves both bandwidth and link failover facilities.

Posted at 03:29 PM | [Permalink](#)



Comments

 You can follow this conversation by subscribing to the [comment feed](#) for this post.



Quick question Erik, do you know when we will be able to uplink the 5020 to an MDS using NPV and port channels?

I notice in your example you have used TE-ports, but I'd like to have a trunked NP-port instead.

Posted by: Anthony Padula | [11/05/2010 at 09:41 PM](#)



Hi Anthony, great question. Cisco is referring to this as F_Port Channeling and it is supported as of NX-OS 4.2(1)N1(1). I'll update the topology and add the setup steps to the FCoE tech book. Hopefully I'll be able to get it into the December release.

Erik

Posted by: Erik Smith | [11/06/2010 at 12:03 AM](#)



That's great, do you know if the same thing applies to the UCS 61x0XP's?

Posted by: Anthony Padula | [11/07/2010 at 07:42 PM](#)

The comments to this entry are closed.

[Brass Tacks](#) [Powered by TypePad](#)