



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

[Main](#)

03/08/2011

PT2PT and VN2VN - Directly connecting an FCoE initiator to an FCoE target (Part 1 - Overview)

Recently, I received a few questions asking “how do you connect an FCoE initiator directly to an FCoE target without using an FCF?” The short answer is, “with the FC-BB-5 version of FCoE you can’t!” The reason is, this functionality was intentionally deferred to FC-BB-6.

To help clarify what is, and what is not supported today, I will describe the basic connectivity options supported by FC-BB-5 and then describe the basic features of the PT2PT and VN2VN protocols being worked on in FC-BB-6. In the near future, I will provide a part 2 and 3 to this post that will explain the PT2PT and VN2VN modes of operation in more detail.

FCoE connectivity options

The relationship between the connectivity options supported in FC and those supported in FC-BB-5 and FC-BB-6 are shown in the following diagram:




Figure 1

As shown, Fibre Channel Switched Fabric (FC-SW) maps directly onto FC-BB-5. When I say FC-SW, I’m referring to environments that utilize an FC switch for connectivity, as opposed to directly connecting an HBA to a storage port or utilizing a hub for the same purpose. The reason FC-BB-5 tackled this particular connectivity option first is that the overwhelming majority of data centers utilize FC-SW.



[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.

Fibre Channel Arbitrated Loop (FC-AL) maps somewhat more loosely to the proposed VN2VN protocol being developed in FC-BB-6. VN2VN and FC-AL are similar in that they both allow for many devices to connect to one another (up to 127 with FC-AL and at least that many with VN2VN) and have similar issues with sharing a common address space. As you will see, special care had to be taken when the protocol was being developed to prevent 24-bit address conflicts.

Fibre Channel point-to-point also maps somewhat loosely onto the proposed PT2PT protocol. PT2PT and FC point-to-point are similar in that they both expect only one device to be present at the other end of the link and bad things will happen if this turns out not to be the case.

FC-BB-5 connectivity options

With FC-BB-5, the FCoE initiators and targets need to be connected to an FCF either directly (figure 2) or through a FIP Snooping Bridge (figure 3).




Figure 2




Figure 3

As I described in my previous post [“FIP, FIP Snooping Bridges, and FCFs \(Part 1 - “FIP”, the FCoE Initialization Protocol\)”](#), the FIP login process relies on the use of FIP frames and on the fact that an entity exists somewhere in the network that is listening for frames destined to the Destination Address (DA) of ALL-FCF-MACs. If there is no such entity, FIP is unable to determine which VLAN FCoE is supported on and is also unable to determine the DA that the FIP FLOGI should be transmitted to. In addition, in both of the cases, it is important to note that all FCoE frames are transmitted with a DA of the FCF-MAC that they logged in with. The point is, without an FCF in the network, the login process would not complete successfully and FCoE frames will have no valid destination that they can be transmitted to. As a result, with FC-BB-5, the following topologies cannot be supported:




Figure 4




Figure 5

As you may have noticed, in the previous topology diagrams I've included a "DCB Cloud". This probably didn't look too strange until the diagram showing the FCoE initiator and target directly connected by a physical cable. The reason I am including the DCB cloud has to do with the fact that with FCoE, a multi-access link is assumed and this had an enormous impact on the development of the PT2PT and VN2VN protocols.

For the sake of completeness, let me also point out that FCF to FCF connections (VE_Port to VE_Port) are supported with FC-BB-5. This allows for the creation of FCoE ISLs and the creation of an all FCoE Fabric should you want to create one.

FC-BB-6 connectivity options

FC-BB-6 will support all of the FC-BB-5 connectivity options (shown above) as well as some new protocols and a new network element that will allow for much more flexibility. The new protocols, PT2PT and VN2VN, will be described in detail in the next posts (part 2 and 3). The new network element, referred to as an FDF (FCoE Data Forwarder), will also be described in detail in a future post (once it is finalized). For now, just think of an FDF as an enhanced FSB or a simplified FCF.

PT2PT

The FCoE PT2PT (Point to Point) protocol is useful if you want to connect two (and ONLY two) FCoE ENodes directly together. The ENodes can either be physically connected with a cable or connected through a Lossless Ethernet Bridge (see figures 4 and 5 above).

When physically connecting two ENodes together, you can either use twinax (copper) or optical fiber. However, as I mentioned in my previous “optical or twinax” post, the downside to using copper is the need to verify that both ENodes support the same twinax cable type. If you use an optical fiber, it is more expensive but you will not have to worry about physical layer interoperability.

Using a Lossless Ethernet bridge to connect two PT2PT ENodes together may prove to be useful if you want to use twinax, but the ENodes do not support the same twinax cable type. For example, ENode A only supports active twinax and ENode B only supports optical. A more realistic example would be that ENode A checks to make sure that a particular brand of twinax is in use and ENode B doesn't support that cable type. When using PT2PT over a Lossless Ethernet Bridge, you are going to need to ensure that either there are only two FCoE PT2PT ENodes in that broadcast domain, or create a VLAN and put only the two FCoE PT2PT ENodes that should be connected into it. In terms of supported L2 Ethernet topologies, PT2PT will operate just like any other Ethernet-based protocol and will not be limited to connectivity between ports on the same bridge. It has been designed to support connectivity between any two ports in a L2 segment and can operate independently of the type of Link Aggregation in use.

One of the strengths of the PT2PT protocol is that it allows for the PT2PT link to initialize very quickly (< 2 seconds). This is of particular importance to mainframe environments and, in fact, a major manufacturer of mainframe products was the driving force behind the creation of a PT2PT type of protocol.

PT2PT limitations

In terms of drawbacks to the PT2PT protocol, since PT2PT is a peer-to-peer type of protocol and may be used when an FCF is not present, the protocol was designed to operate independently of the FC Fabric Services. As a result, there will be no protection provided by FC zoning and no notification via RSCN if the other device should disappear or re-appear for some reason. Instead, the ENodes will need to rely on FIP Link Keep Alive (LKA) Frames to ensure that the other ENode is still present and Advertisements to discover when another PT2PT is added to the network.

VN2VN

The FCoE VN2VN (VN_Port to VN_Port) protocol allows an arbitrary number of VN_Ports to communicate with one another in a peer-to-peer fashion. As with PT2PT, the ENodes can either be physically connected together with a cable or through a Lossless Ethernet Bridge (see figures 4 and 5 above). Also, since VN2VN allows for an arbitrary number of VN_Ports, the topology shown in figure 6 is also allowed.




Figure 6

One of the strengths of the VN2VN protocol is that it allows multiple VN_Ports to communicate with one another without the presence of an FCF. As with PT2PT, VN2VN will operate just like any other Ethernet-based protocol and will not be limited to connectivity between ports on the same bridge. It has been designed to support connectivity between any two ports in a L2 segment and operates independently of the type of Link Aggregation in use.

VN2VN limitations

As with PT2PT, since VN2VN is a peer-to-peer type of protocol and may be used when an FCF is not present, the protocol was designed to operate independently of the FC Fabric Services. As a result, there will be no protection provided by FC zoning and no notification via RSCN if the other device should disappear or re-appear for some reason. Instead, the ENodes will need to rely on FIP Link Keep Alive (LKA) Frames to ensure that the other ENode is still present and Advertisements to discover when another PT2PT is added to the network.


In addition, the VN2VN protocol places some additional requirements on the VN_Ports. The biggest of these is the requirement that each VN_Port keep track of all VN2VN ENodes that are on the same network segment. This effectively means that each ENode will need to maintain a local copy of a Name Server. As I'll show in part 3 of this post, there are very good reasons for why this was necessary, but it may prove to be a rather heavy burden for ENode implementations to bear.

One other point worth mentioning about VN2VN is this: a corner case with a network join scenario could result in duplicate addresses (FPMAs) for a brief period of time. The solution to this will be to implement some form of dynamic ACLs. As a result, when E-Lab eventually qualifies a VN2VN solution, we will probably require the use of an FSB or impose a default ACL requirement. I'll provide more detail in Part 3.

With all of this in mind, even though I was a proponent of the PT2PT protocol and presented my own version of the PT2PT protocol in FC-BB-6, my personal favorite is now VN2VN. It is much more flexible than PT2PT and will make our customers' lives easier.

Thanks for reading!

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



Comments

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



hi,
In FC-BB-5 document i found that FC-BB_GFPT is a point-to-point topology that is supported by the standard.
did I get it wrong?
its on page 23 in the following link

Posted by: [Anton](#) | [11/22/2012 at 04:03 AM](#)



Hi Anton, FCoE is defined in the FC-BB_E section of FC-BB-5. FC-BB_GFPT is something else..

Posted by: [Erik Smith](#) | [11/24/2012 at 07:46 PM](#)

The comments to this entry are closed.

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