



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

[Main](#)

11/09/2010

FIP, FIP Snooping Bridges, and FCFs (Part 1 - “FIP”, the FCoE Initialization Protocol)

Background:

I had an interesting conversation the other day with a colleague in another engineering group at EMC. He was in the process of getting ready to order a 10 GbE switch to support some NAS testing. He also mentioned that he would eventually like to test FCoE. Although I won't mention it by name, the switch he was considering (mainly due to its low cost) was going to be problematic since it didn't support FIP Snooping. When I pointed this out he responded that the product's literature stated it supported FCoE and the vendor had apparently received an industry award for this. My response was something along the lines of:

They can claim whatever they want :-) However, unless they implement FIP Snooping, E-Lab would not qualify or support it.

To provide a bit more detail, currently there are two different types of “FCoE-capable” switches: FIP Snooping Bridges (FSBs) and Fibre Channel Forwarders (FCFs).

An FSB makes use of DCBX, PFC, ETS, and Dynamic ACLs (as defined in FC-BB-5). The last requirement, Dynamic ACLs, implies that the device must be FIP Aware (a.k.a., FIP Snooping). Currently, an FSB must connect to an FCF in order for FCoE to function. This is due to the fact that there is no support for the PT2PT or VN2VN protocols, both of which are FC-BB-6 features, in any of the ENodes are available today.

An FCF is an FSB that also provides FC Services (Name Server, Fabric Server, etc.) as well as FC/FCoE encapsulation and de-encapsulation. FCFs have to be compliant with FC-BB-5, FC-SW, FC-GS, FC-LS, FC-FS, etc. Because of these requirements, only Brocade, Cisco, and QLogic provide FCFs today.

Finally, since FSBs must be connected to an FCF to support FCoE, you are also going to have to buy an FCF to support your FCoE testing which may negate the cost savings you are trying to achieve. In addition, if you want to



[Erik Smith](#)

[Follow](#)

[ARCHIVES](#)

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.

use an FSB, you will have to use a Cisco FCF, since currently Brocade FCFs will not work with FIP Snooping Bridges. This is mainly due to the fact that Brocade only supports a single FCoE login per physical interface.

After I responded, I got the sinking feeling that I am going to be answering this question multiple times if 10 GbE switch vendors all start claiming support for FCoE. So, I decided to create a post on FIP, FIP Snooping, and FCFs. Ideally, I'd like to arm you with enough information to make an informed decision about a vendor's claim of FCoE support. As I started writing it, I quickly realized that I'm going to have to create multiple installments due to the amount of information I am going to have to cover. This particular installment discusses the FCoE Initialization Protocol, or “FIP”. I decided to start with FIP because you need to understand what it is and how it works to appreciate the difference between a 10 GbE switch, an FSB, and an FCF.

FIP, the FCoE Initialization Protocol - Why is it needed?

Before I describe what FIP is and how it works, it would probably help if I described why it's needed. To do this, it's important to back up and look at some of the network features that the Fibre Channel Protocol relies on to function properly.

1. Fibre Channel relies on the fact that there is a physical link between the two ports being attached.
2. Due to the use of Well Known Addresses, FC N_Ports know which address to log in with.
3. FC assumes a point-to-point connection exists between an N_Port and the fabric being logged into. (BTW, notice I said N_Port not NL_Port...)
4. Devices are implicitly logged out from the fabric should the physical link be lost between an N_Port and the switch it connects to. This is an extremely important point because this allows for things like RSCNs to be generated and for the distributed Name Server to be kept in sync.
5. Due to the point-to-point nature of FC, there is implicit security since a man-in-the-middle attack is very difficult to accomplish in this topology.

When you start to think about running FC over Ethernet without FIP, the potential topologies that could result would violate some of the expectations of the FC protocol. With Ethernet:

- There is not necessarily a single physical link between an ENode and its FCF; as a result there is not always an implicit logout when a physical link is lost, which negatively impacts RSCN generation as well as the distributed name server.
- ENodes do not know the address of the FCF that will allow login, or if one even exists in the network.
- There is no one-to-one relationship between an ENode and a fabric, but as I will describe later on, this could be a good thing.
- There is no implicit security so a man-in-the-middle attack is theoretically possible.

So, you may be asking how all of this relates to FIP. The answer is: **FIP bridges the gap between the expectations of the FC protocol and the reality of an**


FCoE topology.

FIP does this by:

- Allowing ENodes to discover who to log in with.
- Enabling a single ENode to communicate with multiple different FC fabrics and as a result a one-to-many relationship is built in. For me, this is one of the more exciting use cases for FCoE. I mean, imagine being able to log in to multiple FC fabrics with a single HBA and not have to use some form of Layer 3 FC Routing such as IVR, FCR or IR!
- Including Link Keep Alive and Clear Virtual Link functions to allow a loss of a physical link or logical connectivity to be detected and for both ends of the virtual link to be notified when this happens. This allows RSCN to function properly and for the distributed name server to remain in sync.
- Reducing the security concerns with FCoE provided that FIP snooping and dynamic ACLs are implemented.

FIP, the FCoE Initialization Protocol - How does it work?

To describe how the FCoE Initialization Protocol (FIP) works, the following topology diagram will be used.



The diagram consists of a DCB Cloud containing a single DCB-capable Lossless Ethernet switch. Attached to the Lossless Ethernet switch is a Converged Network Adapter (CNA) as well as two Fibre Channel Forwarders (FCFs). The FCFs are then connected to a single FC fabric. A couple of important points to note are: the CNA has three MAC Addresses. As shown in the diagram, they are the Universal, ENode, and VN_Port MAC Addresses. The Universal MAC is used for non-FCoE traffic, IP, etc. As you will see, the ENode MAC is used in the FCoE Initialization Protocol and the VN_Port MAC is assigned by the fabric to the CNA and is used for FCoE frames. Other important points are each FCF has been assigned a priority and finally that the fabric has a WWNN associated with it and this is typically the WWN of the principal switch in the FC fabric. Again, more on how these are used later.


Before diving into the details of how FIP works, keep in mind that FIP allows an ENode to:

- Perform VLAN and FCF discovery
- Ensure that the Layer 2 network is capable of supporting mini-jumbo frames
- Perform fabric login
- Use LKA (Link Keep Alive) to maintain the virtual link with the FCF

All of these were identified previously as important things to do when trying to run the FC protocol over an Ethernet topology.

FIP VLAN Request


Once DCBX information has been exchanged, the CNA will transmit a FIP VLAN Request in order to determine which VLAN provides support for FCoE.



The FIP VLAN Request is multicast to the destination MAC Address of ALL-FCF-MACs. The Source Address for the VLAN Request is the ENode MAC and it is important to note that the frame is transmitted without an 802.1Q (VLAN) tag. As a result, the VLAN Request will use the default VLAN for the switch port that it is being attached to. Another important point to note is that earlier in this post I mentioned the concept of Dynamic ACLs. In order for Dynamic ACL updates to work properly, these FIP VLAN Requests must be allowed on a switch port by default.

FIP VLAN Notification

Since the VLAN Request was multicast, any device listening for frames destined to ALL-FCF-MACs, such as an FCF, may transmit a FIP VLAN Notification to the ENode MAC.




A couple of important points to note are the number of notifications that an ENode can receive will be equivalent to the number of FCFs listening on the default VLAN. This will include FCFs that are providing FCoE services on different VLANs, so the ENode can get Notifications for multiple FCoE VLANs. The ENode determines the VLAN by examining the value in the FCoE VID field.

Finally, the notification will include the FCFs MAC Address, which will be used for subsequent FIP Frames.

FIP Discovery Solicitation

Once the ENode has discovered which FCFs exist, it needs to determine which ones it can log in with. It does this by transmitting a multicast Solicitation to ALL-FCF-MACs.




Notice that the Solicitation is tagged with a VLAN ID. If FCoE is being supported on multiple VLANs, there would be multiple Solicitations, one sent to each VLAN where FCoE is being supported. Another important piece of information contained in the solicitation is the MAX_FCOE_SIZE value. This value will be used by the FCF in the response to the Solicitation.

FIP Discovery Advertisement

In response to the Solicitation, each FCF will respond with a unicast FIP Advertisement that will contain:


- The priority of the FCF. This is used by the ENode to help determine which FCF should be logged into if there are multiple FCFs connected to the same fabric.
- The name of the fabric the FCF is attached to. This allows the ENode to determine if two FCFs are attached to the same fabric.
- In addition, a pad field that inflates the Advertisement to be the size of a full FCoE Frame. The purpose of the pad field is to validate that the path between the FCF and the ENode is capable of transporting FCoE frames that are the maximum size supported by the ENode.



Another function of the Advertisement is to act as a catalyst for updating the dynamic ACL. The reception of this frame would cause FCoE frames to be allowed to be sent from this ENode to the DA of the responding FCFs.


FIP Virtual Link Instantiation Request (FLOGI)

As the Advertisements are received from the FCFs, the CNA will store them in an “Internal FCF list” and then select one from that list to perform fabric login with. In this example, the CNA will log in to the FCF with the lower priority.



FIP Virtual Link Instantiation Response (FLOGI ACC)

Finally, the FCF will respond to the Fabric Login with a Fabric Login accept. This is the step where the CNA is provided with its FPMA and it will start to use actual FCoE frames to complete the rest of the login process.



Once the FIP FLOGI ACC has been received, the ENode will continue the login process by transmitting a PLOGI to the Name Server. This PLOGI is encapsulated in an FCoE frame, not a FIP Frame.

Well, it still ended up being a long post but I hope you have a better understanding of why FIP is needed and how it works. In part two, I will focus on the differences between 10 GbE switches, FSBs, and FCFs.

Thanks for reading!

Erik

Posted at 04:44 AM | [Permalink](#)

Reblog (0)

Comments

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



Erik,

This is a fantastic introduction to FCoE login processing. However, the basic question remains: why would you mandate that a DCB-enabled bridge supports FIP snooping?

FIP snooping is a nice security feature, but it's not mandatory to make FCoE work ... and you should always put FCoE traffic in a separate VLAN anyway.

Looking forward to future posts!

Ivan

Posted by: [Ivan Pepelnjak](#) | [11/30/2010 at 06:38 AM](#)



Hi Ivan, thanks! My next post should be ready in a day or so and it will explain exactly why I think FIP Snooping is required.

Erik

Posted by: [Erik Smith](#) | [11/30/2010 at 07:02 AM](#)




Thank you, Erik, for the topic. I read FC-BB-6 at this moment (FCoE) and try to understand FIP procedures. I think that you describe basic routine very well.

Posted by: Streif Vest | [12/27/2010 at 06:22 AM](#)



Excellent write up on FIP. Thanks very much.
Just one question, why there must be a VLAN associated with every VSAN? Just I could not understand this? regards. Mohan

Posted by: Mohan | [03/03/2012 at 10:30 AM](#)



Hi Mohan, thanks!

In regards to your question, there can only be one VSAN per VLAN due to a requirement in the FCoE standard (FC-BB-5). Essentially, if you have more than one VSAN per VLAN, it would be possible to have duplicate FPMAs (Fabric Provided MAC Addresses) and this can lead to data corruption under the right circumstances.

Posted by: [Erik Smith](#) | [03/04/2012 at 08:17 AM](#)

The comments to this entry are closed.

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