

Fibre Channel over Ethernet Initialization Protocol

What You Will Learn

On June 3, 2009, INCITS T11 approved and released the final revision of FC-BB-5, the T11 standard covering Fibre Channel over Ethernet (FCoE). Cisco has been actively participating in the definition of FC-BB-5, and the Cisco Nexus™ 5000 Series Switches have offered an FCoE implementation since the time the T11 committee released a draft of FC-BB-5 that finalized the FCoE encapsulation (frame format) and data plane aspects. Since then, the T11 committee has focused on the FCoE control plane functions, and specifically on the virtual link establishment and management functions known as FCoE Initialization Protocol (FIP).

The Cisco Nexus 5000 Series FCoE implementation could establish point-to-point FCoE links to first-generation converged network adapters (CNAs) without FIP, but FIP is required to build more complex topologies in which FCoE links can map to an Ethernet path that consists of more than one Ethernet link segment. The final revision of FC-BB-5 completed the definition of FIP and was soon followed by Cisco® NX-OS Software Release 4.1(3)N1(1), the first software release for the Cisco Nexus 5000 Series that offers FIP capabilities.

This document describes FIP as it is defined in FC-BB-5 and shows how FIP and FC-BB-5 map to the Cisco Nexus 5000 Series software implementation.

Background

Fibre Channel has established itself in the storage industry as the best-performing and most reliable transport for block storage. This role is the result of years of innovation and tuning that have built a solid ecosystem of end-to-end services and tools to satisfy the storage needs of mission-critical applications.

FCoE offers the capability to transport Fibre Channel payloads on top of an Ethernet network. FCoE originates from the simple observation that as long as the same ecosystem of end-to-end services and tools remains available, and as long as performance is not affected, users do not have a preference for a physical representation of data on a wire. By offering equivalent quality-of-service (QoS) semantics, and by presenting the same data-link-layer payloads to upper-layer services, an Ethernet link can effectively replace a Fibre Channel link at the physical layer.

If the technology offered only this, however, FCoE would not be very interesting. With Fibre Channel frames carried on an Ethernet cable, the next step is to take these frames and mix them with regular Ethernet frames on the same wire. By enabling the same Ethernet link to carry both Fibre Channel and Ethernet payloads, FCoE increases utilization of the physical infrastructure and therefore reduces the total number of network ports that must exist overall. In this basic model, the Fibre Channel physical layer is replaced by an Ethernet physical layer, Fibre Channel and Ethernet frames share the same wire, and each switch forwards each individual frame from one ingress port to one egress port using the forwarding rules that are appropriate: Fibre Channel rules (based on T11 FC-SW) for Fibre Channel frames, and Ethernet rules (based on IEEE 802.1D) for Ethernet frames.

This model is the approach that the Cisco Nexus 5000 Series followed in incorporating FCoE when the product line was first released. The one-to-one mapping between a Fibre Channel link and an Ethernet link is a simple way to build an FCoE network that can potentially extend end to end, from initiators to targets, through a number of switching points. The switching points can perform switching with Fibre Channel rules and can implement the entire stack of Fibre Channel services (starting from FC-2). The fabric is just a Fibre Channel fabric that happens to use Ethernet to encode signals from one point to the next.

This simplicity is a strength of FCoE, but it limits end-to-end deployments to a class of devices that understand both the Ethernet and the Fibre Channel stacks (dual-stack switches called FCoE forwarders [FCFs] in T11 parlance). How can this limitation be overcome? From among the few ways to allow frames with different data-link-layer structures to coexist on the same transmission line, the T11 standard committee chose for FCoE the way that is least disruptive to the existing Ethernet definition: encapsulation. Encapsulating Fibre Channel frames with an additional Ethernet header transforms the Fibre Channel frame into a standard legal Ethernet frame—one that can potentially be forwarded by any standard Ethernet switching implementation. This property of encapsulation can be used to interleave dual-stack switches with regular Ethernet switches that appear as transparent points (FCoE passthrough switches) to the Fibre Channel stack.


However, while Fibre Channel protocols and Fibre Channel switches are designed to deal only with point-to-point connectivity, FCoE passthrough switches introduce shared-media (point-multipoint) semantics. In other words, an FCoE passthrough switch is invisible to the Fibre Channel stack, but it acts as a concentrator of flows from multiple servers into the same port of a dual-stack switch, and thus it creates a shared-medium hanging off the dual-stack switch port. A typical Fibre Channel stack is not prepared to address shared-media links, since they do not exist in native Fibre Channel; these new semantics must therefore be handled by FCoE without affecting the Fibre Channel stack.

The solution in FCoE is the concept of virtual Fibre Channel interfaces and the introduction of a small set of control plane functions to build virtual links among pairs of virtual FC interfaces. Multiple virtual FC interfaces can be presented on the same physical Ethernet interface of a dual-stack switch. After the virtual links are established among pairs of virtual FC interfaces, the result is presented to the Fibre Channel stack as a set of point-to-point connections that behave exactly like regular Fibre Channel connections, regardless of the number of FCoE passthrough switches that are physically inserted between the actual Fibre Channel endpoints. FIP is the set of control plane functions that enable discovery of FCoE-capable devices across FCoE passthrough switches and establishment of legal combinations of virtual links. Both FCoE and FIP are described in T11 FC-BB-5.

FCoE Architectural Model

T11 FC-BB-5 FCoE defines two types of endpoints for the Ethernet encapsulation of Fibre Channel frames: FCoE nodes (ENodes) and FCoE forwarders (FCFs). Figure 1 shows the FCoE architecture.

Figure 1. FCoE Architectural Model



ENodes are the combination of FCoE termination functions and Fibre Channel stack on the CNAs, and in that sense they are equivalent to host bus adapters (HBAs) in native Fibre Channel networks. FCFs are the combination of FCoE termination functions and Fibre Channel stack on Ethernet switches (dual-stack switches) and are therefore equivalent to Fibre Channel switches in native Fibre Channel networks.

ENodes present virtual FC interfaces in the form of VN_Ports, which can establish FCoE virtual links with FCFs' VF_Ports (Figure 2). FCFs present virtual FC interfaces in the forms of VF_Ports or VE_Ports; a VF_Port establishes FCoE virtual links with a CNA's VN_Port, and VE_Ports enable FCFs to establish FCoE virtual links with one other (Figure 3). These interface types have their equivalents in native Fibre Channel's N_Ports, F_Ports, and E_Ports.

Figure 2. ENode Model in FC-BB-5

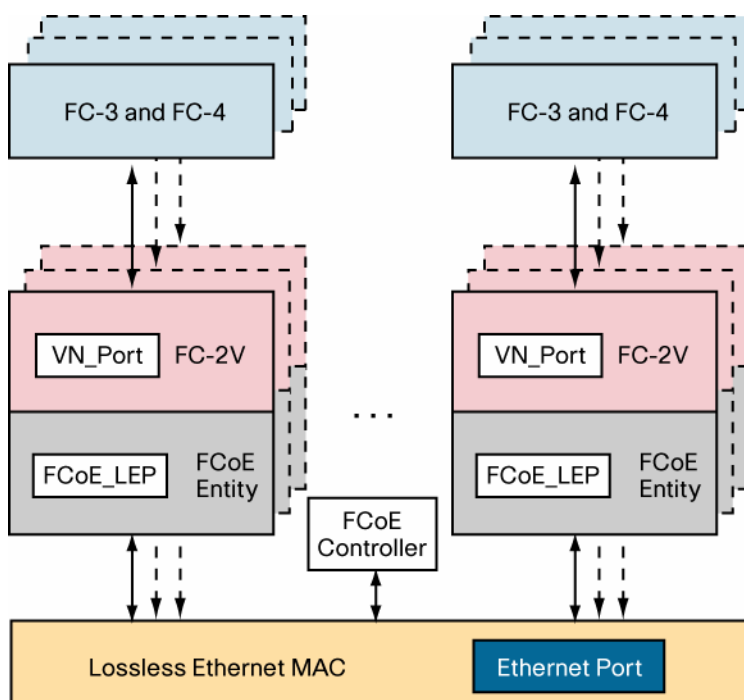
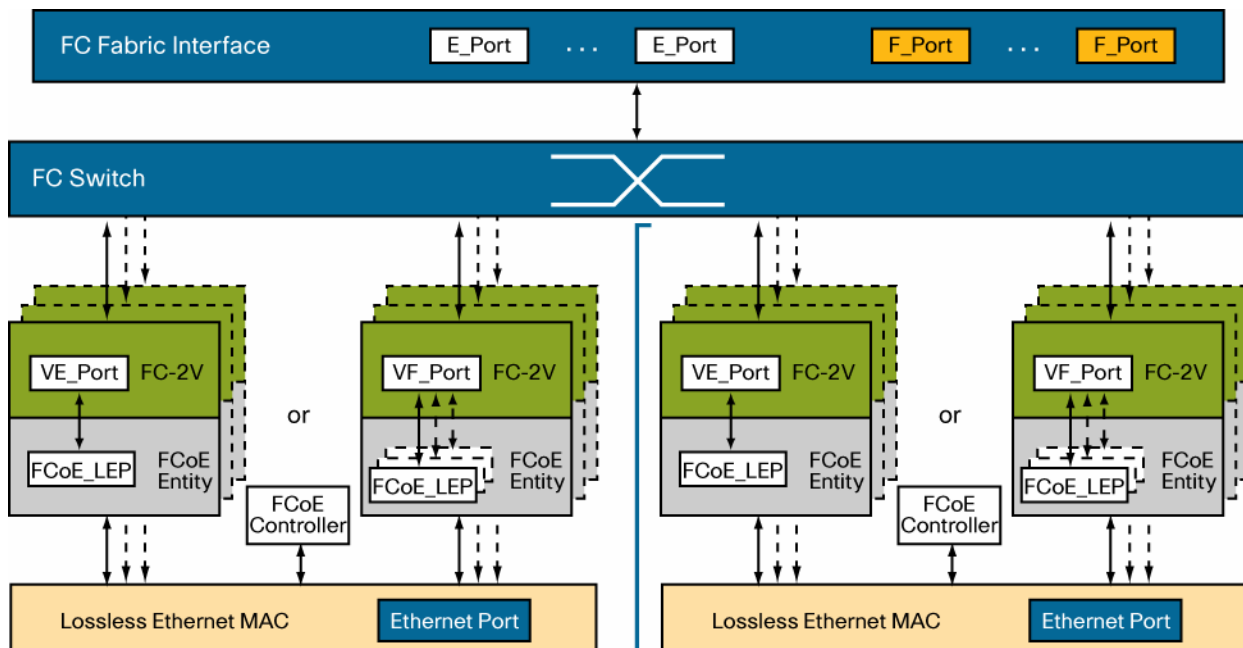


Figure 3. FCF Model in FC-BB-5



FCoE virtual links replace the physical Fibre Channel links by encapsulating Fibre Channel frames in Ethernet frames, and an FCoE virtual link is identified by the MAC addresses of the two FCoE endpoints. T11 FC-BB-5 FCoE mandates lossless characteristics for the Ethernet segments traversed by encapsulated Fibre Channel frames, and although it does not specify the exact means to obtain such QoS, IEEE 802.1 Data Center Bridging (DCB), and specifically 802.1Qbb, is assumed for this purpose on the Ethernet segments.

Communication from one endpoint to the other in an FCoE virtual link uses Ethernet forwarding semantics, augmented by IEEE 802.1 DCB functions. Obviously, if no FCoE passthrough switches exist in the path between the two FCoE endpoints (FCoE direct attach), no Ethernet switching is required, but IEEE 802.1 DCB semantics still apply on the physical Ethernet link. Switching from ingress to egress virtual interface of an FCF requires Fibre Channel forwarding semantics.

A Fibre Channel virtual link between two FCoE devices, A and B, is uniquely identified by this triplet:

[MAC address of FCoE device A, MAC address of FCoE device B, FCoE VLAN ID]

Given a physical Ethernet interface on an FCoE device (or, more accurately, what FC-BB-5 calls an FCoE controller behind that interface), the FCoE device must establish all the virtual links for which that physical interface is configured. Typically, the configuration is represented in the form of one or more virtual FC interfaces (what FC-BB-5 calls FCoE logical endpoints [FCoE-LEPs]) mapped on top of the physical Ethernet interface.

The information for the triplets that characterize the virtual links can be statically configured on the FCoE device, but FC-BB-5 defines FIP to dynamically discover all the information required. In other words, for every virtual FC interface configured on top of a physical Ethernet interface, FIP is responsible for discovering a pairing virtual FC interface somewhere else in the Ethernet network and establishing a virtual link between the two virtual interfaces in a specific FCoE VLAN.

FCoE VLANs

FCoE packets must be exchanged in a VLAN. For FCFs with a Fibre Channel stack that includes multi-VSAN capabilities like the Cisco Nexus 5000 Series, the FCoE traffic belonging to different VSANs must remain separated by different VLANs on the Ethernet plane. This choice simplifies the implementation, since it removes the necessity to include both a VLAN and a VSAN header in each FCoE packet: the VLAN is assumed to be a proxy for a VSAN.

For this reason, the Cisco Nexus 5000 Series software introduced the concept of a mapping table between FCoE VLANs and Fibre Channel VSANs. For each Fibre Channel VSAN used in the Fibre Channel fabric, the administrator associates one and only one unique FCoE VLAN; all FCoE packets tagged with that VLAN ID are then assumed for Fibre Channel control plane and forwarding purposes to belong to the corresponding Fibre Channel VSAN.

The Cisco Nexus 5000 Series implementation also expects all FCoE VLANs to be used exclusively for FCoE traffic and never shared by other traditional Ethernet payloads (for example, IP traffic).

Note that virtual links exist in one and only one FCoE VLAN. VSAN trunking capabilities (as defined in FC-SW-5, a standard for native Fibre Channel switching fabrics) are not explicitly described by FC-BB-5; however, equivalent semantics can easily be obtained by defining multiple virtual FC interfaces on the same physical interface: one per FCoE VLAN.

FIP

FCoE Initialization Protocol (FIP) is the FCoE control protocol responsible for establishing and maintaining Fibre Channel virtual links between pairs of FCoE devices (ENodes or FCFs). During the virtual link establishment phase, FIP first discovers FCoE VLANs and remote virtual FC interfaces; then it performs virtual link initialization functions (fabric login [FLOGI] and fabric discovery [FDISC], or exchange link parameters [ELP]) similar to their native Fibre Channel equivalents. After the virtual link is established, Fibre Channel payloads can be exchanged on the virtual link, and FIP remains in the background to perform virtual link maintenance functions; it continuously verifies reachability between the two virtual FC interfaces on the Ethernet network, and it offers primitives to delete the virtual link in response to administrative actions to that effect. This document does not describe the virtual link maintenance functions of FIP.

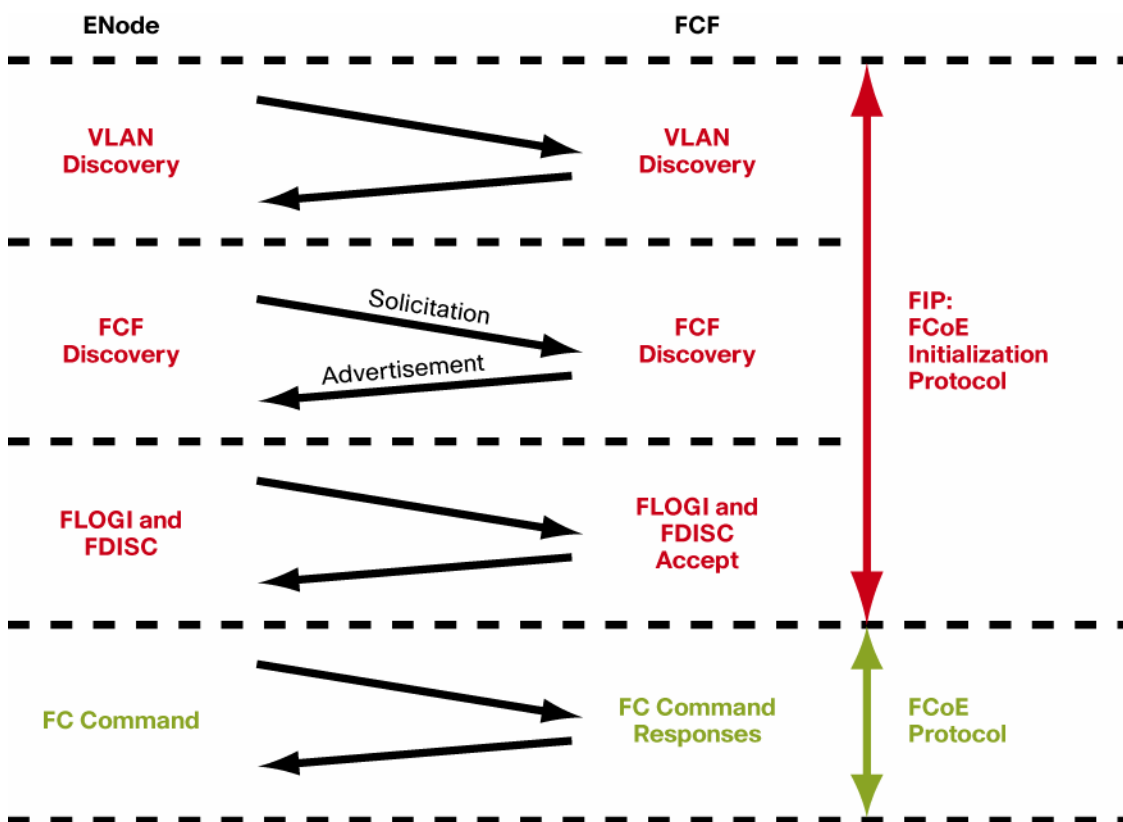
FIP aims to establish virtual FC links between VN_Ports and VF_Ports (ENode to FCF), as well as between pairs of VE_Ports (FCF to FCF), since these are the only legal combinations supported by native Fibre Channel fabrics. Standards-compliant implementations are not required to support both forms of virtual FC links, and Cisco has decided to focus initially on implementing FIP only between ENodes and FCFs. FCF-to-FCF connectivity is considered a strategic direction for end-to-end FCoE deployments, but the short-term urgency is for FCoE adoption between CNAs and the Fibre Channel fabric perimeter, where unified fabric can offer the greatest capital expenditure (CapEx) savings today. For this reason, the remainder of this document focuses on FIP in the context of virtual FC links between VN_Ports and VF_Ports. This capability is available on the Cisco Nexus 5000 Series products starting with Cisco NX-OS Software Release 4.1(3)N1(1).

For the sake of completeness, note that virtual Fibre Channel links between a pair of VN_Ports, that is, direct connectivity of end nodes (for instance, a server to a disk) without a Fibre Channel fabric, is not supported in FC-BB-5, but it may be included in a future revision of the standard.

Discovery and Virtual Link Establishment

FIP defines two discovery protocols as well as a protocol to establish virtual links between VN_Ports and VF_Ports. Figure 4 shows a typical FIP protocol exchange resulting in the establishment of a virtual link between an ENode's VN_Port and an FCF's VF_Port.


Figure 4. FIP Virtual Link Establishment



All the protocols are usually initiated by ENodes, although FCFs can generate unsolicited FIP advertisements, as discussed later in this section. Note that the FIP frames at the top and the FCoE frames at the bottom of Figure 4 use different EtherTypes and encapsulations, since the FCoE frames encapsulate native Fibre Channel payloads, whereas FIP frames describe a new set of protocols that have no reason to exist in native Fibre Channel definitions. Among the differences, note that ENodes use different source MAC addresses for FIP and FCoE encapsulation. FIP

packets are built using a globally unique MAC address assigned to the CNA at manufacturing (called the ENode MAC address), whereas FCoE packets are encapsulated using a locally unique MAC address (that is, unique only within the boundaries of the local Ethernet subnet) dynamically assigned to the ENode by the FCF as part of the FIP virtual link establishment process (a fabric-provided MAC address [FPMA]). The definition of the FPMA is shown in Figure 5.

Figure 5. FPMA



FPMA uses the 24-bit-wide Fibre Channel ID (FC_ID) assigned to the CNA during the FIP FLOGI and FDISC exchange, and therefore they cannot be available to the CNA before the fabric login has occurred. The FPMA is built by appending the FC_ID to a 24-bit quantity called the FCoE MAC address prefix (FC-MAP). FC-BB-5 defined a range of 256 FC-MAPs to facilitate FCoE deployments. Cisco has established very simple best practices (see “FCoE VLANs” earlier in this document) that make the manipulation of FC-MAPs unnecessary, and most users should find the default FC-MAP value 0E-FC-00 sufficient. The 256 different FC-MAPs make available to users up to 256 pools of locally unique MAC addresses. The pools are useful when the FC_IDs are not unique on an Ethernet VLAN; for instance, when different Fibre Channel fabrics or different VSANs are encapsulated in the same Ethernet VLAN, the ranges of FC_IDs assigned in each Fibre Channel fabric may overlap. Cisco strongly recommends that you never attempt to map multiple Fibre Channel fabrics onto the same Ethernet VLAN. Most users will not ever need to map multiple Fibre Channel fabrics onto the same physical Ethernet network, but if such a need arises, each Fibre Channel fabric should be encapsulated in a separate VLAN.

FIP VLAN Discovery

FIP VLAN discovery discovers the FCoE VLAN that will be used by all other FIP protocols as well as by the FCoE encapsulation for Fibre Channel payloads on the established virtual link. One of the goals of FC-BB-5 was to be as nonintrusive as possible on initiators and targets, and therefore FIP VLAN discovery occurs in the native VLAN used by the initiator or target to exchange Ethernet traffic. The FIP VLAN discovery protocol is the only FIP protocol running on the native VLAN; all other FIP protocols run on the discovered FCoE VLANs.

The ENode sends a FIP VLAN discovery request to a multicast MAC address called All-FCF-MACs, which is a multicast MAC address to which all FCFs listen. All FCFs that can be reached in the native VLAN of the ENode are expected to respond on the same VLAN with a response that lists one or more FCoE VLANs that are available for the ENode’s VN_Port login. This protocol has the sole purpose of allowing the ENode to discover all the available FCoE VLANs, and it does not cause the ENode to select an FCF.

FIP VLAN discovery is an optional protocol in FC-BB-5. An ENode implementation can choose to offer only manual configuration for FCoE VLANs, and therefore choose not to perform FIP VLAN discovery. It is commonly assumed that such implementation will default to VLAN 1002 for its FCoE VLAN. The Cisco Nexus 5000 Series supports FIP VLAN discovery, and it will respond to any ENode that performs a query. The contents of the response depend on how the virtual FC interface is configured on the Cisco Nexus 5000 Series Switch, as discussed later in this document.

FIP FCF Discovery

FIP FCF discovery is the protocol used by ENodes to discover FCFs that can accept logins. FCFs periodically send FIP FCF discovery advertisement messages on each configured FCoE VLAN; these messages are destined for the multicast MAC address All-ENode-MACs, a multicast MAC address to which all ENodes listen. The FIP FCF discovery advertisement is used by the FCF to inform any potential ENode in the VLAN that FCF VF_Ports are available for virtual link establishment with ENodes' VN_Ports. The advertisement includes the MAC address of the FCF as well as other parameters useful for tuning the characteristics of the virtual link (FIP timeout values, FCF priority, etc.).

Given the periodic nature of the advertisements, new ENodes joining the network will typically not want to wait to collect multicast FIP FCF discovery advertisements from all FCFs, and therefore FC-BB-5 allows ENodes to solicit unicast advertisements by sending a FIP FCF discovery solicitation to the All-FCF-MACs multicast MAC address. FCFs receiving the solicitation can generate a unicast FIP FCF discovery advertisement addressed to the requesting ENode. Upon collection of these advertisements, the ENode can make the final decision as to which FCF to contact for the establishment of a virtual link with its VN_Port.

FIP FLOGI and FDISC


After the ENode has discovered all FCFs and selected one for login, the last step is to inform the selected FCF of the intention to create a virtual link with its VF_Port. After this step, Fibre Channel payloads (encapsulated in FCoE frames) can start being exchanged on the new virtual link just established. On any native Fibre Channel link between an N_Port and an F_Port, the first protocol exchange performed as part of activating the data-link layer is the fabric login, or FLOGI, which results in the assignment of an FC_ID to the N_Port. In designing FIP, the T11 committee decided to merge the logical step of FCF selection by an ENode in FIP with the native Fibre Channel fabric login exchange. The result of this optimization is a single FIP exchange that serves both purposes of FCF selection, as well as fabric login and FC_ID allocation. This optimization is not only convenient; it is a requirement for obtaining an appropriate FPMA for the ENode to use in the subsequent FCoE encapsulated frames.

FIP FLOGI and FDISC are unicast frames almost identical to the native Fibre Channel FLOGI and FDISC frames they replace. The VN_Port sends an FLOGI or an FDISC request, followed by the corresponding FLOGI or FDISC accept payload from the FCF. Completion of this exchange terminates the FIP virtual link establishment phase.

FCoE Security Model

Direct-Attach FCoE


When an ENode is directly attached to an FCF, the FIP virtual link establishment described previously is a set of point-to-point exchanges that map one or more virtual links on top of a physical Ethernet link (Figure 6). In this simple scenario, a Fibre Channel stack exists at both ends of the physical wire. Although on the physical wire (at Layer 1) all packets are encoded using Ethernet encodings, the Fibre Channel stack assumes the responsibility of forwarding FCoE frames as if they were native Fibre Channel frames.

Figure 6. Direct-Attach Connection Between ENodes and FCFs

In the direct-attach model, Fibre Channel semantics apply fully to all FCoE frames from the data-link layer up. Forwarding, management, troubleshooting, link events: everything looks exactly alike between a native Fibre Channel physical link and a virtual link using FCoE. Obviously, this approach implies that even from a security perspective, FCoE does not inherently add any more risk than a native Fibre Channel segment in a Fibre Channel fabric. This property of direct-attach FCoE makes the direct-attach model highly desirable, since it offers SAN administrators the opportunity to start using FCoE without really having to go beyond their domain of expertise. In a way, even FIP can be seen as just a formality in this context, and first-generation CNAs prove this point, since first-generation CNAs do not support FIP but nonetheless enable FCoE in direct-attach topologies.

Remote-Attach FCoE

If direct-attach FCoE topologies are the most desirable way to enable FCoE, why would anyone ever want to use another approach? The quick answer is complexity: switches with FCF capabilities are more complex for vendors to implement than regular Ethernet switches, and therefore in certain contexts users will face a trade-off between the convenience of a direct-attach model and the lack of a strong portfolio of FCF offerings to meet their needs. One typical scenario that is often cited is that of blade switches. A blade switch is typically a simpler device than a traditional access-layer switch, and time-to-market or return-on-investment (ROI) concerns may cause blade switch vendors to choose to have their Ethernet blade switches behave as FCoE passthrough switches, rather than implement a full FCF. Under those conditions, users have no choice but to forego direct-attach FCoE and use remote-attach FCoE solutions. In a remote-attach FCoE solution, the virtual link established between a VN_Port and a VF_Port maps to an Ethernet path that includes one or more FCoE passthrough switches (Figure 7). Each FCoE passthrough switch does not have a Fibre Channel stack and therefore makes forwarding decisions based purely on Ethernet semantics. This type of environment is where FIP adds the most value to the FCoE solution.

Figure 7. Remote-Attach Connection Between ENodes and FCFs

Obviously, if a virtual link is built on top of one or more Ethernet devices that are transparent to the Fibre Channel fabric, the native Fibre Channel security model will not be sufficient to address security concerns on those devices. For this reason, FC-BB-5 includes an informative addition, Annex D, to offer a solution to the potential Fibre Channel security threats on FCoE passthrough switches. Implementation of the recommendations in FC-BB-5 (Recommendation D.4.1, which is more stringent, and Recommendation D.4.2, which is less stringent) is optional for vendors providing FCoE passthrough switches. Absence of these functions in an FCoE passthrough switch does not necessarily make FCoE undeployable, but it leaves the deployment susceptible to the security threats described in Annex D of FC-BB-5. Some administrators may not find these threats applicable to their environments and so choose to deploy FCoE passthrough switches that do not offer these functions (or choose to keep them disabled), but Cisco recommends enabling these functions whenever possible.

Since for the foreseeable future the Cisco Nexus 5000 Series is expected to be deployed as an FCF and not as an FCoE passthrough switch, these recommendations are not available on the Cisco Nexus 5000 Series as of Cisco NX-OS Software Release 4.1(3)N1(1).

FIP Snooping

In the industry, FCoE passthrough switches that implement either Recommendation D.4.1 or D.4.2 are commonly referred to as FIP snooping switches, to distinguish them from FCoE passthrough switches that simply implement IEEE DCB and do not offer any additional FCoE security capability on top of that. The term “FIP snooping” is not defined in FC-BB-5, and it is slightly misleading, since under certain conditions Recommendation D.4.1 or D.4.2 could be implemented without snooping any packet. This document avoids use of this definition, but it is reported here for completeness.

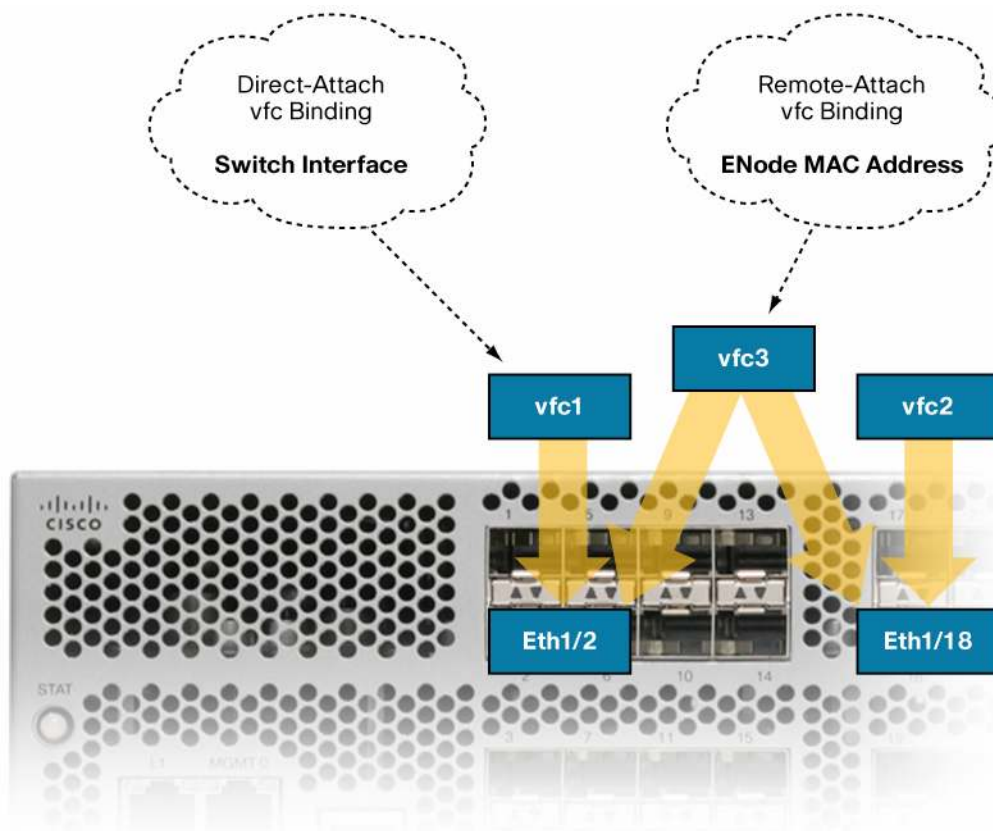
Cisco Nexus 5000 Series Virtual Fibre Channel Interfaces

The Cisco Nexus 5000 Series software uses the virtual Fibre Channel interface construct to infer the number of VF_Ports that need to be exposed by FIP in each FCoE VLAN. Like any interface in Cisco NX-OS, virtual FC interfaces (called `vfc` in the CLI) are manipulable objects with properties such as configuration and state. Virtual FC interfaces behave like native Fibre Channel interfaces, except for the lack of VE_port capabilities as of Cisco NX-OS Software Release 4.1(3)N1(1). Therefore, virtual FC interfaces get assigned to a VSAN, and the combination of the

vfc VSAN assignment and the global VLAN-to-VSAN mapping table enables the Cisco Nexus 5000 Series FIP implementation to choose the appropriate VLAN for a VF_Port.

Virtual Fibre Channel interfaces can be bound directly on top of a physical Ethernet interface, or they can be bound to an ENode MAC address, as shown in Figure 8.

Figure 8. vFC Binding Options



Virtual FC interfaces bound to a physical Ethernet interface on the Cisco Nexus 5000 Series signal to the software that a direct-attach model is used for that interface. The switch then expects the CNA to be at the other end of the physical Ethernet cable, and it accepts one FLOGI from that one CNA, regardless of its ENode MAC address; moving the CNA to a different physical Ethernet interface on the switch will cause the CNA to become subject to the configuration on the vfc attached to the other interface (if any). After the FLOGI, more FDISC operations can follow if N_port ID virtualization (NPIV) is enabled, with exactly the same semantics as native Fibre Channel interfaces. FIP VLAN discovery advertises only the FCoE VLAN that maps to the VSAN configured on the vfc. Only zero or one vfc's can be bound to a specific physical Ethernet interface at any point in time.

Virtual FC interfaces bound to an ENode MAC address signal to the software that a remote-attach model is used to create a virtual link with that specific ENode. FIP FCF advertisements are exchanged on all physical Ethernet interfaces that are not subject to direct-attach bindings, and the ENode can create a virtual link through any of those interfaces, but regardless of the ingress interface, it will always be subject to the vfc that is bound to its MAC address.

You should bind virtual FC interfaces to physical Ethernet interfaces whenever possible, and vfc binding to MAC addresses is offered in anticipation of blade switches that will implement only FCoE passthrough switch capabilities. First-generation CNAs that are not capable of exchanging FIP frames will continue to work in a direct-attach configuration with the Cisco Nexus 5000 Series, since the Cisco Nexus 5000 Series Switch will recognize the condition and revert to the pre-FIP exchanges that CNAs expect for FCoE.

Conclusion

T11 FC-BB-5 enables an FCoE topology of virtual FC links to be built on top of a physical Ethernet network. Ethernet switches in the physical topology participate to the FCoE overlay as either FCFs that terminate the virtual links and offer Fibre Channel forwarding, or as FCoE passthrough switches that are invisible to the FCoE overlay and use Ethernet forwarding to preserve the abstraction of virtual FC links across them. Initiators and targets participate in the FCoE overlay topology with CNAs (ENodes) that terminate the virtual FC links on the other side of the FCFs. FIP is the protocol that enables ENodes and FCFs to establish and maintain virtual FC links across the transparent FCoE passthrough switches.

The Cisco Nexus 5000 Series is the first Ethernet switch in the industry that offers FCF capabilities that include FIP, and that enables first-generation CNAs to create virtual FC links in direct-attach configurations without FIP, and second-generation CNAs to create virtual FC links in both direct-attach and remote-attach configurations with FIP.

For More Information

For more information about the Cisco Nexus 5000 Series, please visit <http://www.cisco.com/go/nexus5000> or contact your local account representative.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

CCDE, CCENT, CCSE, Cisco Eos, Cisco HealthPresence, Cisco IronPort, the Cisco logo, Cisco Lumina, Cisco Nexus, Cisco Nitro Connect, Cisco Pulse, Cisco StackPower, Cisco StadiumVision, Cisco TelePresence, Cisco Unified Computing System, Cisco WebEx, DCE, Flip Channels, Flip for Good, Flip Mini, FlipShare (Design), Flip Ultra, Flip Video, Flip Video (Design), Instant Broadband, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn, Cisco Capital, Cisco Capital (Design), CiscoFinanced (Stylized), Cisco Store, and Flip Gift Card are service marks; and Access Register, Aironet, AllTouch, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCR, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, Connum, EtherFast, EtherSwitch, Event Center, Explorer, Fast Step, Follow Me Browsing, FormShare, GainMaker, GigaDrive, HomeLink, ILYN, Internet Quotient, IOS, IPPhone, iQuickStudy, IronPort, the IronPort logo, Laser Link, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerKEY, PowerPanel, PowerTV, PowerTV (Design), PowerVu, Prime, ProConnect, ROSA, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TennaPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0908R)