

Annex D: FCoE Security Recommendations (Informative)

D.1 Overview

During the development of this standard, a detailed threat analysis was performed to ensure that FCoE networks may be practically deployed while maintaining security characteristics comparable to native Fibre Channel. As a result of this analysis, the deployment recommendations presented in this annex were developed. Many of the recommendations are included in other normative clauses of this standard as they are required for correct operation of FCoE in addition to providing a level of security comparable to Fibre Channel. The implementation of the remaining recommendations is not technically required for correct operation. However, to ensure the security characteristics of a FCoE network are comparable to the ones of a native Fibre Channel Fabric, these recommendations should be considered requirements.

D.2 Considerations

FCFs are assumed to be trusted devices. Therefore, a bridge known to be directly connected to an FCF is not required to perform any verification of frames received from the FCF.

For the purposes of the annex, “known” refers to knowledge gained through administrative action or from a trusted management application.

Furthermore, bridges that fully implement these recommendations provide a defensive perimeter. Therefore bridges known to be directly connected to bridges forming a defensive perimeter are not required to perform verification of frames received from the bridges forming the defensive perimeter. However, performing such verification enhances the security characteristics of the network. Doing so may come at the cost of limiting the scalability of the network and its ability to autonomously respond to faults, as well as in increased administrative complexity. These trade-offs should be considered during Fabric design and deployment.

Finally, FCoE networks may be subjected to various forms of catastrophic failures if duplication of VN_Port MAC addresses occur. These failures may include failure of the network to provide service, undetected data interception, and undetected data corruption. Addressing this issue differs depending on the addressing mode (i.e., FPMA or SPMA) being used. This is discussed in more detail in D.6.

D.3 General deployment recommendations

- 1) No VLAN should carry more than one Fibre Channel Virtual Fabric (applies to the LAN if VLANs are not in use).

NOTE 29 – The possibility of carrying more than one Virtual Fabric on a given VLAN was beyond the scope of the threat analysis performed during the development of this standard. Therefore, if more than one Virtual Fabric is deployed on a VLAN, this standard provides no assurance with respect to the security characteristics of the Fabric. Furthermore, it was observed that such a deployment greatly increases the possibility of a MAC address being assigned to multiple VN_Ports which may result in catastrophic Fabric failures and undetected data corruption. The probability of duplicate VN_Port MAC addresses is greatly exacerbated by the use of FPMA, but the concern applies to both FPMA and SPMA. Therefore, the deployment of multiple Virtual Fabrics on a single VLAN is strongly discouraged. This concern does not apply to the deployment of multiple Virtual Fabrics with each being deployed on an independent VLAN.

D.4 Bridge recommendations

- 1) All Bridge ports, except those known to be connected to FCFs, or those that are to be explicitly prohibited from carrying FCoE/FIP traffic, should implement ingress filtering that:
 - I) discards all frames with a source MAC address matching that of any FCF;
 - II) discards all frames with Ethertype = FIP_TYPE except those:
 - i) addressed to the All-FCF-MACs group address; or
 - ii) addressed to any FCF-MAC address;
 - III) discards all frames with Ethertype = FCoE_TYPE except those:
 - i) containing a source MAC address currently assigned by an FCF to a VN_Port;

NOTE 30 – A currently assigned MAC address refers to the MAC address that has been assigned by the FCF via a FIP FLOGI / FIP NPIV FDISC LS_ACC and has not been unassigned by virtue of a corresponding FIP LOGO, FIP Clear Virtual Links, or the expiration of associated FIP Keep Alive timers.

- ii) containing a Destination MAC address equal to the FCF-MAC address of the FCF that assigned the source MAC address; and
- iii) received on the bridge port through which the FIP FLOGI / FIP NPIV FDISC LS_ACC was forwarded to the VN_Port that has been assigned the VN_Port MAC address.

Annex C provides a description of a method (i.e., the use of Access Control Lists) to satisfy this recommendation.

To ensure security characteristics comparable to native Fibre Channel, this recommendation should be applied to all bridge ports directly connected to ENodes. If this is not possible (e.g., the bridges connected to ENodes do not provide this filtering capability), enhanced security may still be obtained by applying this filtering to upstream bridge ports, but the security characteristics in this case are significantly weaker.

In many deployments, implementation of ingress filtering according to this recommendation on bridge ports connected to other bridge ports may have undesirable consequences with respect to scalability, ability of the network to autonomously respond to faults, and administrative complexity. If care is taken to ensure that the far-end bridge implements the recommendations in this annex, then this ingress filtering is not strictly required on the near-end bridge port. However, implementation of this recommendation provides additional protection (e.g., against configuration errors). Careful consideration should be given to the advantages and disadvantages of implementing this filtering on bridge ports connected to other bridges.

- 2) As an alternative to recommendation 1, all bridge ports known to be connected to other bridge ports should implement filtering that:
 - I) if it is known that the port is to receive frames from ENodes but not FCFs:
 - i) discards FIP frames that are not addressed to FCFs; and
 - ii) discards FCoE frames that are not addressed to FCFs; and

- iii) discards FCoE frames that are transmitted from one FCF to another FCF.
- II) if it is known that the port is to receive frames from FCFs addressed to ENodes but not other FCFs:
- i) discards all FIP frames that are not transmitted from an FCF;
 - ii) discards all FIP frames that are addressed to another FCF;
 - iii) discards all FCoE frames that are not transmitted from an FCF; and
 - iv) discards all FCoE frames that are addressed to another FCF.
- III) if it is known that the port may receive frames from ENodes and FCFs that may be addressed to both FCFs and ENodes:
- i) discards all FIP frames that are not either sourced by or addressed to an FCF; and
 - i) discards all FCoE frames that are not either sourced by or addressed to an FCF.

This recommendation provides a more scalable and resilient alternative to recommendation 1 for bridge-to-bridge ports. However, the security characteristics of this recommendation are weaker than that of recommendation 1. Consideration should be given to which recommendation best meets the overall needs of a given deployment.

Annex C provides a description of a method (i.e., the use of Access Control Lists) to satisfy this recommendation.

- 3) Bridges should not perform any learning function based on the source address of a frame that was discarded by recommendation 1 or recommendation 2.
- 4) All ethernet bridges carrying FCoE and FIP traffic should ensure that any VLAN carrying FCoE or FIP traffic for a given Virtual Fabric is in an independent VLAN learning set relative to all other VLANs.
- 5) Bridge ports intended to specifically exclude ingress of FIP and FCoE traffic should implement ingress filtering that discards all frames with Ethertype equal to FIP_TYPE or FCoE_TYPE.
- 6) Bridges intended to transport FIP and FCoE traffic should not discard FIP and FCoE frames due to congestion.

Ethernet bridges, unlike Fibre Channel switches, do not by default provide a flow control mechanism. Therefore, ethernet bridges discard frames when congested. Such discarding of frames can result in significant performance degradation of Fibre Channel traffic. This may be avoided by deploying mechanisms to prevent this type of packet discard. Two possible mechanisms of accomplishing this are the use of the Pause mechanism (see IEEE 802.3-2008) or of the Priority-based Flow Control mechanism (see IEEE 802.1Qbb) within the ethernet bridges.

D.5 ENode and FCF recommendations

All of the recommendations in this subclause, except for recommendation 12, appear as normative requirements in this standard, and in some cases, appropriate ethernet standards. They are repeated here to highlight their applicability to Fabric security considerations.

- 1) ENodes discard all received frames with an Ethertype equal to FIP_TYPE except:
 - I) those that contain a Destination MAC address equal to All-ENode-MACs; and
 - II) those that contain a Destination MAC address that equals a source MAC address used in a FIP Discovery Solicitation from the ENode.
- 2) ENodes discard all received frames with an Ethertype equal to FCoE_TYPE that:
 - I) contain a destination MAC address / destination N_Port_ID pair that was not assigned by an FCF to one of the VN_Ports on the ENode; or
 - II) contain a source MAC address that does not match the MAC address of the FCF that assigned the corresponding VN_Port MAC address.

In the case of SPMAs, the MAC address assigned by an FCF refers to the MAC address approved by the FCF during the FIP FLOGI / FIP NPIV FDISC process.

Using SPMAs, it is possible that multiple VN_Ports are assigned the same MAC address by one or more FCFs. This recommendation ensures that the frame is addressed from a VF_Port to its corresponding VN_Port. FPMAs uniquely address all VN_Ports.

- 3) FCFs discard all frames received with an Ethertype = FCoE_TYPE that:
 - I) contain a destination MAC address that does not match the MAC address of one of the FCF's VE_Ports or VF_Ports; or
 - II) contain the source MAC address that does not match the MAC addresses that the FCF has assigned to the corresponding VN_Port or was established for the corresponding VE_Port; or
 - III) in the case of a VN_Port, contains a Fibre Channel source address that does not match the one assigned to the VN_Port by the FCF.
- 4) On transmission, VN_Ports construct all frames with:
 - I) the Destination MAC address set to the MAC address of the FCF that it successfully performed a FIP FLOGI or FIP NPIV FDISC with; and
 - II) the Source MAC address set to the MAC address assigned to the VN_Port by the FCF as a result of the FIP FLOGI or FIP NPIV FDISC.
- 5) On transmission, VF_Ports construct all frames with:
 - I) the destination MAC address set to the MAC address of the VN_Port as assigned by the transmitting FCF during FIP FLOGI / FIP NPIV FDISC; and
 - II) the source MAC address set to the MAC address of the VF_Port (i.e., that of the FCF).
- 6) On transmission, VE_Ports construct all frames with:
 - I) the source MAC address of the transmitting VE_Port; and
 - II) the destination MAC address of the remote VE_Port.

- 7) The MAC Client within a FCF does not deliver:
 - I) to a VE_Port or VF_Port, any frame whose Ethertype is not equal to FCoE_TYPE; and
 - II) to the FCoE controller, any frame whose Ethertype is not equal to FIP_TYPE; or
 - III) alternatively, VE_Ports, VF_Ports, and FCoE Controllers discard all frames that do not contain an Ethertype of FCoE_TYPE, FCoE_TYPE, and FIP_TYPE, respectively.
- 8) FCF ports that implement multiple port types (i.e., VF_Port and VE_Port) do not use the same MAC address for different port types.
- 9) ENodes may choose to transmit a FIP FLOGI / FIP NPIV FDISC to any FCF(s).
- 10) While processing a FIP FLOGI or FIP NPIV FDISC, an FCF either rejects the request or ensures that the MAC address assigned to the requesting ENode:
 - I) complies with local administrative policy; and
 - II) in the case of FPMA, the 24 most significant bits contain the Fabric's FC-MAP and the 24 least significant bits equal that of the assigned Fibre Channel address identifier.

For FPMAs, the fact that the assigned MAC address contains a Fabric wide unique Fibre Channel address identifier provides assurance that the MAC address itself is unique Fabric wide.

- 11) FCFs may choose to create or not create VE_Ports with other FCFs based on local policy information (e.g., the MAC address of other FCFs).
- 12) All source MAC addresses used in FIP should be globally assigned (see IEEE 802-2001 for a description of globally assigned MAC addresses).

D.6 Additional threat isolation using FPMAs

There is a class of threats related to the misuse of MAC addresses assigned to VN_Ports. If multiple VN_Ports utilize the same MAC address (e.g., through mis-configuration or other network issues), catastrophic network failures may occur including undetected corruption of data. In addition, the use of MAC addresses assigned to VN_Ports by malicious stations provides a number of attack possibilities that include denial of service attacks and undetected data interception. In addition, threats exist related to using an Fibre Channel address identifier that is not associated with the MAC address assigned to a given VN_Port.

Fabric measures to prevent such attacks using SPMA are generally not practical. Doing so requires that all bridges with edge ports have knowledge of all MAC addresses being used for VN_Ports. The standards do not provide a mechanism where this knowledge may be reasonably obtained. Furthermore, even if these MAC addresses were known, implementations would face challenging scalability issues. Consequently, protection against these failures and attacks are accomplished by other means (e.g., careful network configuration, enforcement of strict physical security measures). Specific recommendations for protection against these failures and attacks when using SPMA are beyond the scope of this annex.

With FPMAs it is possible to test all MAC addresses assigned to VN_Ports, both now and in the future. This may be accomplished by specifically allowing the valid VN_Port MAC addresses on a given port, which is learned by examining FIP FLOGI / FIP NPIV FDISC LS_ACC messages egressing the bridge port, and discarding all other frames in which the 24 most significant bits of the source MAC

address match the Fabric's FC-MAP. Furthermore, it is trivial to ensure the proper association between a VN_Port MAC address and its associated Fibre Channel address identifier (i.e., the 24 least significant bits of the VN_Port MAC address matches the VN_Port Fibre Channel address identifier).

To obtain the additional security capability provided by FPMA, the following recommendations are provided in addition to those previously discussed:

- 1) Bridge ports, except those known to be connected to FCFs, those known to be connected to other bridge ports, and those that are to be explicitly prohibited from carrying FCoE/FIP traffic, should implement ingress filtering that discards all frames containing a source MAC address in which the 24 most significant bits match the FCoE Fabric's FC-MAP and in which the source MAC address does not match a valid VN_Port MAC address assigned by an FCF to a device connected to that port. This requirement applies regardless of the Ethertype.

The MAC addresses currently assigned to VN_Ports that are reached through the bridge port may be obtained by examining the FIP FLOGI / FIP NPIV FDISC LS_ACC, FIP LOGO, FIP Clear Virtual Links, and FIP Keep Alive messages.

- 2) Bridge ports known to be connected to other bridge ports should:
 - I) if it is known that the bridge port is to receive frames from ENodes but not FCFs:
 - i) discard FCoE frames where the 24 most significant bits of the source MAC address do not match the Fabric's FC-MAP.
 - II) if it is known that the bridge port is to receive frames from FCFs addressed to ENodes but not other FCFs:
 - i) discard all FCoE frames where the 24 most significant bits of the destination MAC address do not match the Fabric's FC-MAP.
 - III) if it is known that the bridge port may receive frames from ENodes and FCFs that may be addressed to both FCFs and ENodes:
 - i) discard all FCoE frames that do not contain either a source MAC address matching that of an FCF or a source MAC address where the 24 most significant bits match the Fabric's FC-MAP; and
 - ii) discard all FCoE frames that do not contain either a destination MAC address matching that of an FCF or a destination MAC address where the 24 most significant bits match the Fabric's FC-MAP.
- 3) Bridges should not perform any address learning function based on the source MAC address of a frame that was discarded by recommendation 1 or 2.
- 4) Bridge ports intended to specifically exclude ingress FIP and FCoE frames should implement ingress filtering that discards all frames with a source MAC address where the 24 most significant bits match the FCoE Fabric's FC-MAP.

This recommendation prevents a malicious host from injecting a packet utilizing a victim's source MAC address in an attempt to alter the bridge learning tables such that it may intercept the data destined to the victim.

- 5) On reception, VN_Ports verify that the destination Fibre Channel address identifier matches the 24 least significant bits of the destination MAC address.
- 6) On reception, VF_Ports verify that the source Fibre Channel address identifier matches the 24 least significant bits of the source MAC address.
- 7) On transmission, VN_Ports construct all frames such that the source Fibre Channel address identifier matches the 24 least significant bits of the source MAC address.
- 8) On transmission, VF_Ports construct all frames such that the destination Fibre Channel address identifier matches the 24 least significant bits of the destination MAC address.