

## 1 FCoE Security Recommendations (*Informative*)

### 1.1 Introduction

During the development of FC-BB-5, a detailed threat analysis was performed to ensure that FCoE networks may be practically deployed while maintaining security characteristics comparable to native Fibre Channel (see T11/08-532v0 et. seq.). As a result of this analysis, a series of deployment recommendations were developed which are presented in this Annex. Many of the recommendations are included in other normative sections of this specification as they are required for correct operation of FCoE in addition to providing a level of security comparable to Fibre Channel. In the remaining recommendations, their implementation is not technically required for correct operation; however, to ensure security characteristics of a FCoE fabric are comparable to a native Fibre Channel fabric, these recommendations should be considered requirements.

### 1.2 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 gain 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, however, may come at the cost of limiting fabric scalability, its ability to autonomously respond to fabric faults, and increased administrative complexity. These trade-offs should be considered during fabric design and deployment.

Finally, FCoE fabrics may be subjected to various forms of catastrophic failure if duplication VN\_Port MAC addresses occur. These failures may include failure of the fabric to provide service, undetected data interception, and undetected data corruption. Addressing this issue differs depending on the addressing mode (FPMA or SPMA) being used. This is discussed in more detail in section 1.6.

### 1.3 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 1 – 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 FC-BB-5. 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 failure 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.

### 1.4 Bridge Recommendations

- 1) All Bridge Ports, except those known to be connected to FCFs, or 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) Discard all frames with Ethertype = FCoE\_Type except those:
  - i) Containing a source MAC address currently assigned by an FCF to a VN\_Port (Note: currently assigned MAC address refers to one that has been assigned by the FCF via an FLOGI/FDISC accept and has not been unassigned by virtue of a corresponding FIP LOGO, FIP Clear Virtual Link, or the expiration of associated FIP Keep Alive timers); and
  - ii) Containing a Destination MAC address equal to the MAC address of the FCF that assigned the source MAC address; and
  - iii) Received on the bridge port through which the FLOGI/FDISC accept was forwarded to the VN\_Port that assigned the VN\_Port's MAC address.

NOTE 2 – Annex TBD provides a description of one method (the use of Access Control Lists) to satisfy this recommendation.

NOTE 3 – To ensure security characteristics comparable to native Fibre Channel, the above recommendation should be applied to all bridge ports directly connected to ENodes. If this is not possible (e.g. the bridges connected to the EPorts do not provide this filtering capability), enhanced security may still be obtained by applying this filtering to upstream bridge ports; however, the security characteristics in this case are significantly weaker.

NOTE 4 – Implementation of the ingress filtering in this recommendation on bridge ports connected to other bridge ports in many deployments may have undesirable consequences with respect to scalability, ability of the fabric 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, for 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 not addressed to FCFs; and
    - ii) Discards that FCoE frames are not addressed to FCFs; and
    - iii) Discard FCoE frames from one FCF destined to another FCF
  - II) If it is known that the port is to receive frames from FCFs destined to ENodes but not other FCFs:
    - i) Discards all FIP frames not sourced from an FCFs; and

- ii) Discards all FIP frames that are not destined to another FCF; and
  - iii) Discards all FCoE frames are not sourced from an FCF; and
  - iv) Discards all FCoE frames destined for another FCF
- III) If it is known that the port may receive frames from ENodes and FCFs that may be destined to both FCFs and ENodes:
- i) Discards all FIP frames that are not either sourced by or destined to an FCF; and
  - i) Discards all FCoE frames or are not either sourced by are destined to an FCF

NOTE 5 – These recommendations provide a more scalable and resilient alternative to the 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.

NOTE 6 – Annex TBD provides a description of one method (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/FIP traffic should ensure that any VLAN carrying FCoE/FIP traffic for a given virtual fabric is in an independent learning set relative to all other VLANs carrying FCoE/FIP traffic for other virtual fabrics.
- 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.

NOTE 7 – Ethernet bridges, unlike Fibre Channel switches, do not by default provide a flow control mechanism. Therefore, Ethernet bridges will discard frames when congested. Such discard of frames can result in significant performance degradation of Fibre Channel. This may be avoided by deploying mechanisms to prevent this packet discard. Two possible mechanisms of accomplishing this is the activation of PAUSE or Priority Flow Control within the Ethernet bridges. As of the writing of this annex, the PAUSE mechanism is fully defined by Ethernet standards and the Priority Flow Control Ethernet standard was under development.

## 1.5 ENode and FCF Recommendations

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

**Editor's Note:** We (including our esteemed editor of FC-BB-5) need to verify that these recommendations are indeed normative in the main body of the specification (or we have made a conscious decision to not make individual recommendations normative so that the above statement may be updated).

- 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 solicitation from the ENode.
- 2) ENodes discard all received frames with an Ethertype equal to FCoE\_Type that:
- I) contain a destination MAC address 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 F\_Port within the FCF that assigned the corresponding VN\_Port MAC address.

NOTE 8 – In the case of SPMA, “MAC address...assigned by an FCF” refers to the MAC address approved by the FCF during the FLOGI / FDISC process.

NOTE 9 – Using SPMA it is possible that multiple VN\_Ports are assigned the same MAC address by one or more FCFs. The above recommendation ensures that the frame is addressed from an F\_Port to its corresponding VN\_Port. FPMA addresses 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 FC-SID 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 on which it successfully performed an FLOGI or FDISC; and
  - II) the Source MAC address set to the MAC address assigned to the VN\_Port by the FCF as a result of the FLOGI or 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 FLOGI/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;
  - 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 any FCF(s) to which to attempt an FLOGI / FDISC.
- 10) During FLOGI or 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 ID.
- NOTE 10 – With FPMA, the fact that the assigned MAC address contains a fabric wide unique Fibre Channel ID provides assurance that the MAC address itself is unique fabric wide.
- 11) FCFs may choose to form or not to form 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 Std 802-2001, Overview and Architecture, for a description of globally assigned MAC addresses).

## 1.6 Additional threat isolation using FPMA

There exists a class of threats related to the misuse of MAC addresses assigned to VN\_Ports. If multiple VN\_Ports utilize the same MAC Address (through mis-configuration or other fabric issues), catastrophic fabric failures may occur including undetected corruption of data. In addition, if the use of a 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 FCID 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 by which this knowledge may be reasonably obtained. Furthermore, even if these addresses were known, implementations would face challenging scalability issues. Consequently, protection against these failures and attacks must be accomplished by other means which may include careful fabric configuration, enforcement of strict physical security measures, etc. Specific recommendations for protection against these failures and attacks when using SPMA are beyond the scope of this annex.

With FPMA it is possible to test for all addresses assigned (both now and in the future) to VN\_Ports. This may be accomplished by specifically allowing the valid VN\_Port addresses on a given port (which is knowable by examining FIP FLOGI/FDISC accept messages egressing the bridge port) and discarding all other frames in which the 24 most significant bits of the source address match the fabric's FC\_MAP. Furthermore, it is trivial to ensure the proper association between a VN\_Port's MAC address and its associated FCID: the 24 least significant bits of the VN\_Port MAC address matches the VN\_Port's FCID.

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, to other bridge ports, or 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 do not match the FCoE fabric's FC\_MAP (regardless of Ethertype)

NOTE 11 – The MAC addresses currently assigned to VN\_Ports that are reached through the bridge port may be obtained by examining the FIP FLOGI/FDISC Accept, FIP LOGO, FIP Clear Virtual Link, and FIP keep alive messages.

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

NOTE 12 – 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) VN\_Ports verify that the destination FCID matches the 24 least significant bits of the destination MAC address.
- 6) VF\_Ports verify that the source FCID matches the 24 least significant bits of the source MAC address.

- 7) On transmission, VN\_Ports construct all frames such that the source FCID matches the 24 least significant bits of the source MAC address.
- 8) On transmission, VF\_Ports construct all frames such that the destination FCID matches the 24 least significant bits of the destination MAC address