

T11/08-547v0 Addressing the Threat Model

****Preliminary****

The standard ACL:

For FPMA:

```
SA = {FCFs}, deny;  
DA = ALL_FCF_MACs, Type = FIP_TYPE, permit;  
DA = {FCFs}, Type = FIP_TYPE, permit;  
Type = FIP_TYPE, deny;  
Type = FCoE_Type, deny;  
SApre = FC_MAP, deny;  
Any non-FCoE related ACEs
```

For SPMA:

```
SA = {FCFs}, deny;  
DA = ALL_FCF_MACs, Type = FIP_TYPE, permit;  
DA = {FCFs}, Type = FIP_TYPE, permit;  
Type = FIP_TYPE, deny;  
Type = FCoE_Type, deny;  
SA = {AllSPMAMACs}, deny;  
Any non-FCoE related ACEs.
```

For each successful FLOGI (or FDISC), an ACE is added (i.e., prior to Type=FCoE_Type, deny;) of the form:

SA = FCF assigned MAC address, DA = FCF MAC address, Type = FCoE, permit;

Assumptions:

1. All devices known to be FCFs are trusted and therefore no ACEs are required.
2. All bridges known to carry FIP/FCoE traffic and known to provide perimeter protection may be trusted and therefore no ACEs are required (it is possible to assume that no bridges are trusted, however, configuring a bridge to be trusted greatly enhances scalability).
3. No other bridges are trusted and therefore the port connected to it must either implement the "standard" ACEs, or the ACEs in rule 19 that prohibit all FCoE/FIP traffic, as appropriate.
4. All FCFs must be directly connected to a bridge that implements these rules

Note: It is not assumed that all ENodes are directly connected to bridges implementing these rules; however, full protection can only be achieved if they are.

Rules (do not apply to ENodes or Bridges that do not carry FIP/FCoE traffic)

1. All Bridge Ports, except those known to be connected to FCFs, known to be connected to other bridges that comply with these rules, or are to be explicitly prohibited from carrying FCoE/FIP traffic, must implement ingress filtering (e.g. ACEs) that:
 - 1.a. Discards all frames with an SA matching that of any FCF
 - 1.b. Discards all frames with Ethertype = FIP_TYPE except those:
 - 1.b.i. Addressed to the ALL_FCF_MACs group address; or

- 1.b.ii. Addressed to any FCF MAC address
 - 1.c. Discard all frames with Ethertype = FCoE_Type except those:
 - 1.c.i. Containing an Source MAC address currently assigned by an FCF to a VN_Port connected to the physical bridge port receiving the frame (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 Keel Alive timers); and
 - 1.c.ii. Containing a Destination MAC address equal to the MAC address of the FCF that assigned the source MAC address
 - 1.d. Discard all frames containing a Source MAC address assigned by any FCF in the fabric to any VN_Port connected to any other physical bridge port or FCF on the fabric (regardless of Ethertype)
2. Bridges shall not perform any learning function based on the source address of a frame that was required to be discarded by any of the rules in 1.
3. All ENodes shall discard all received frames with an Ethertype = FIP_Type except:
 - 3.a. those that contain a Destination MAC address equal to ALL_ENODE_MACs; and
 - 3.b. those that contain a Destination MAC address and Source MAC address equal to the Source MAC address and Destination MAC address (respectively) used in a FIP solicitation from the ENode
4. All ENodes shall discard all received frames with an Ethertype = FCoE_Type except those that contain an Destination Mac address that was assigned by an FCF to one of the VN_Ports on the ENode and contain an Source MAC address equal to that of the FCF that assigned the VN_Port MAC address
5. All bridges transporting FCoE traffic shall not discard frames due to congestion (e.g. through the implementation of PAUSE or Priority Flow Control)
6. A physical port that implements multiple port types (i.e. VN_Port, VF_Port, and/or VN_Port) shall not use the same MAC address for different port types.
7. All FCFs shall discard all frames received with an Ethertype = FCoE_Type except those that contain a Destination MAC address equal to the MAC address of one of the FCF's VE_ports or VF_Ports and contain the Source MAC address equal to the MAC addresses the FCF has assigned to the corresponding VN_Port or was established for the corresponding E_Port
8. No VLAN shall carry more than one Fibre Channel virtual fabric (applies to the LAN if VLANs are not in use).

9. On transmission, a VN_Port shall construct all frames containing:
 - 9.a. A Destination MAC address containing the MAC address of the FCF on which it successfully performed an FLOGI or FDISC; and
 - 9.b. An Source MAC address containing the MAC address assigned to the VN_Port by the FCF as a result of the FLOGI or FDISC
10. On transmission, a VF_Port shall construct all frames containing:
 - 10.a. A Destination MAC address containing the MAC address of the VF_Port (i.e. that of the FCF); and
 - 10.b. A Source MAC address containing the MAC address of the VN_Port as assigned by the transmitting FCF during FLOGI/FDISC.
11. During FLOGI or FDISC, an FCF shall either reject the request or ensure that the MAC address to be assigned to the requesting ENode:
 - 11.a. In the case of FPMA, the 24 MSBs contains the fabric's FC_MAP and the 24 least significant bits are unique among all devices that are currently logged into the FCF
 - 11.b. In the case of SPMA, the entire MAC address is unique among all devices on the entire fabric regardless of whether those devices are executing FIP/FCoE or not.
 - 11.c. In either case, the address to be assigned complies with local administrative policy.
12. The MAC Client within a bridge shall not deliver to a VE_Port or VF_Port any frame whose Ethertype is not equal to FIP_TYPE or FCoE_Type.
13. A VE_Port, upon frame reception, shall discard frames in which the Source MAC address is not that of the VE_Port on the remote end of the link.
14. On transmission, a VE_Port shall construct the frames containing:
 - 14.a. A source address of the transmitting VE_Port; and
 - 14.b. A destination address of the remote VE_Port
15. An ENode may choose any FCF(s) to which to attempt an FLOGI / FDISC
16. All source addresses used in FIP shall be globally assigned.
17. FCFs may chose to form or not to form VE_Ports with other FCFs based on local policy information including, at a minimum, the MAC address of other FCFs.
18. All Ethernet bridges carrying FCoE/FIP traffic shall 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.
19. On ports explicitly configured to exclude FIP/FCoE, the port shall implement ingress filtering that:

- 19.a. Discards all frames with Ethertype equal to FIP_Type or FCoE_Type; and
 - 19.b. Discards all frames containing a SA assigned by any FCF in the fabric to any VN_Port connected to any other physical bridge port or FCF on the fabric (regardless of Ethertype)
- 20.