



T11 FC-BB-6

Distributed FCF Implicit Logout

Roger Hathorn
10/05/2011

Background


- When a VN_Port is logged out from a VF_Port on an FDF, the FDF sends a VNUN ILS to the controlling FCF to communicate unreachability of a VN_Port
- A VNUN received by a controlling FCF triggers a set of NPZD requests to the secondary controlling FCF and to all of the FDFs.
- There are events that occur that require a VF_Port to implicitly logout *all* of its logged in VN_Ports and de-instantiate all virtual links
 - **FLOGI from any VN_Port associated with a PPN**
 - **Physical or Virtual Link Failure**

FC-LS-2 Rule regarding reception of FLOGI

- Upon receipt of FLOGI from a PN_Port with at least one VN_Port previously logged in, a VF_Port shall;
 - implicit logout of all logged-in VN_Ports
 - disassociate the logged-out N_Port_IDs with PPN
 - re-assign one N_Port_ID to new VN_Port
 - Set PPN of new VN_Port to the F_Port_Name in the FLOGI LS_ACC.

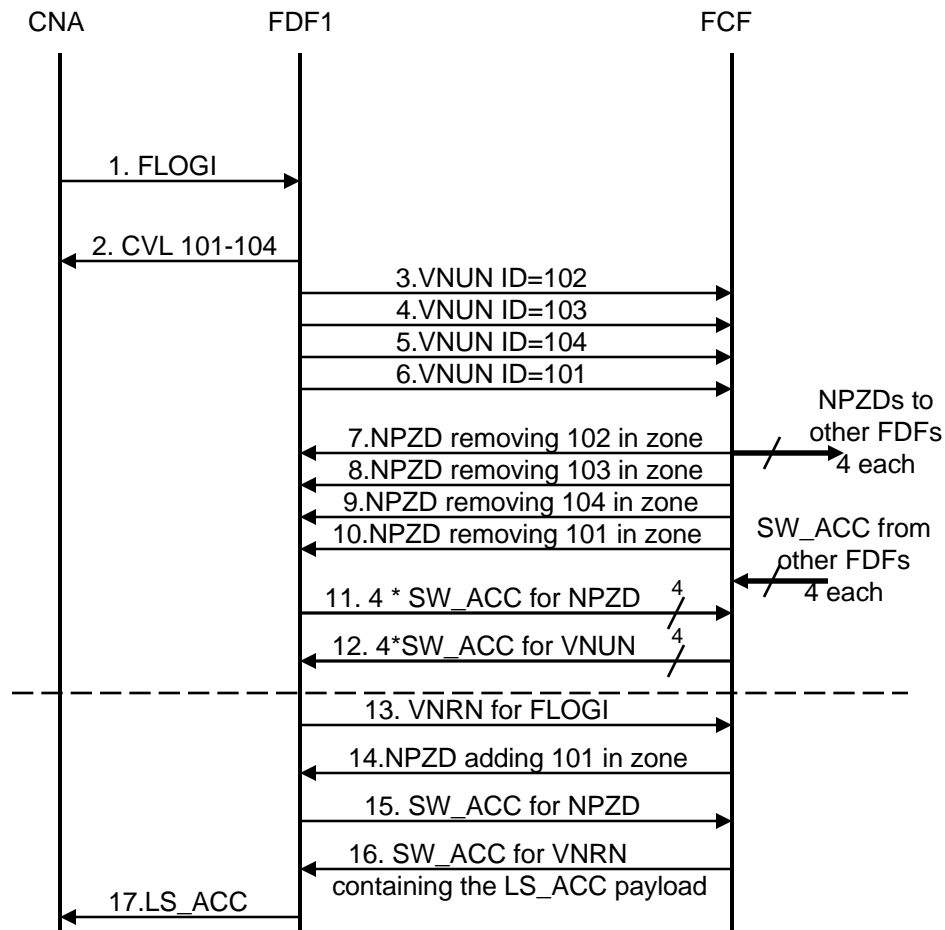
(From FC-GS-6: The Permanent Port Name (PPN) is a name identifier associated with a PN_Port and its N_Port_IDs and shall be the F_Port Name of the F_Port that is attached to the PN_Port.)

Re-login – example scenario



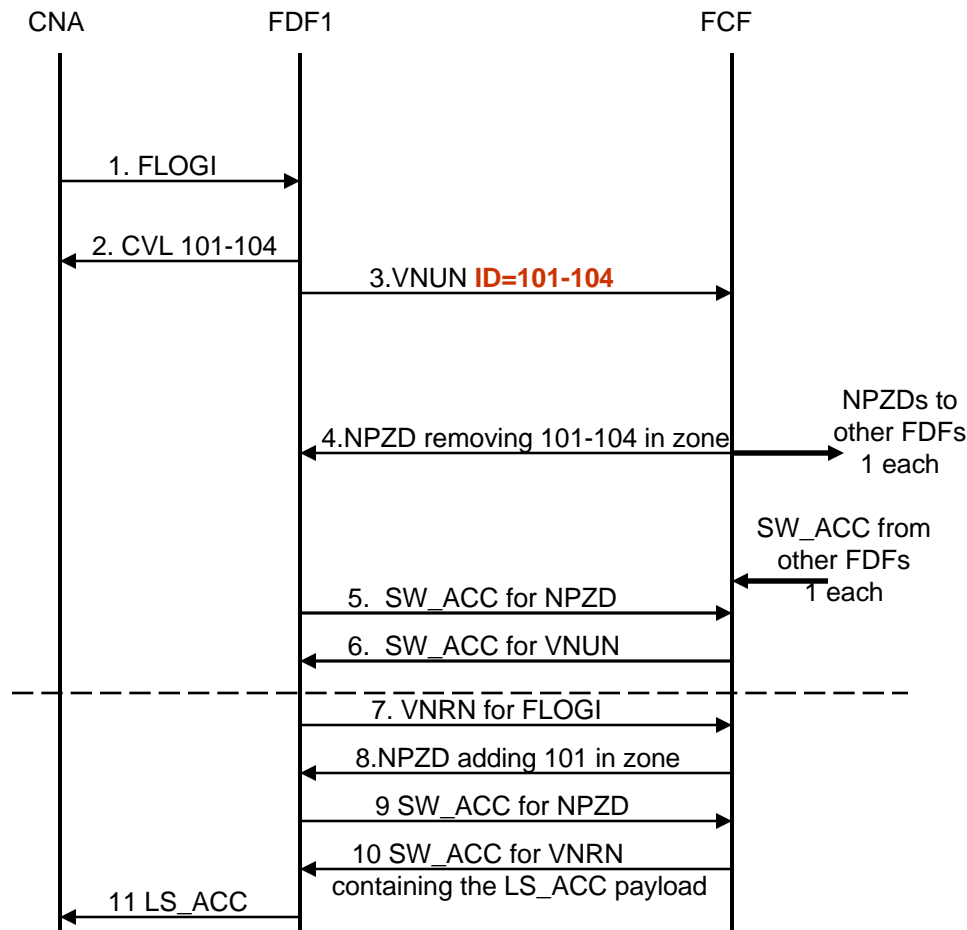
- An ENode may detect a temporary physical or virtual link failure upon which implicit logout for all VN_Ports is performed at the ENode
- Recovery requires a new Fabric Login (FLOGI) to log back into the fabric.
- The VF_Port may or may not have detected the temporary link failure (Bridges may exist)
- The VF_Port will receive FLOGI while VN_Ports from the same PN_Port are still logged in.
- Implicit logout of the previously logged in VN_Ports is performed and the PN_Port is assigned a new N_Port ID.
- This will result in many VNUN notifications to the primary cFCF and even more NPZDs from the primary cFCF to the secondary cFCF and all of the FDFs in the distributed switch.
- The new N_PORT ID assignment requires a VNRN to the primary cFCF and its associated NPZDs to complete before the FCF can respond LS_ACC to the FLOGI.
- The NPZDs for the VNRN request have to be serialized with the NPZDs for all of the VNUN requests.

Re-login as proposed in 11-026v3



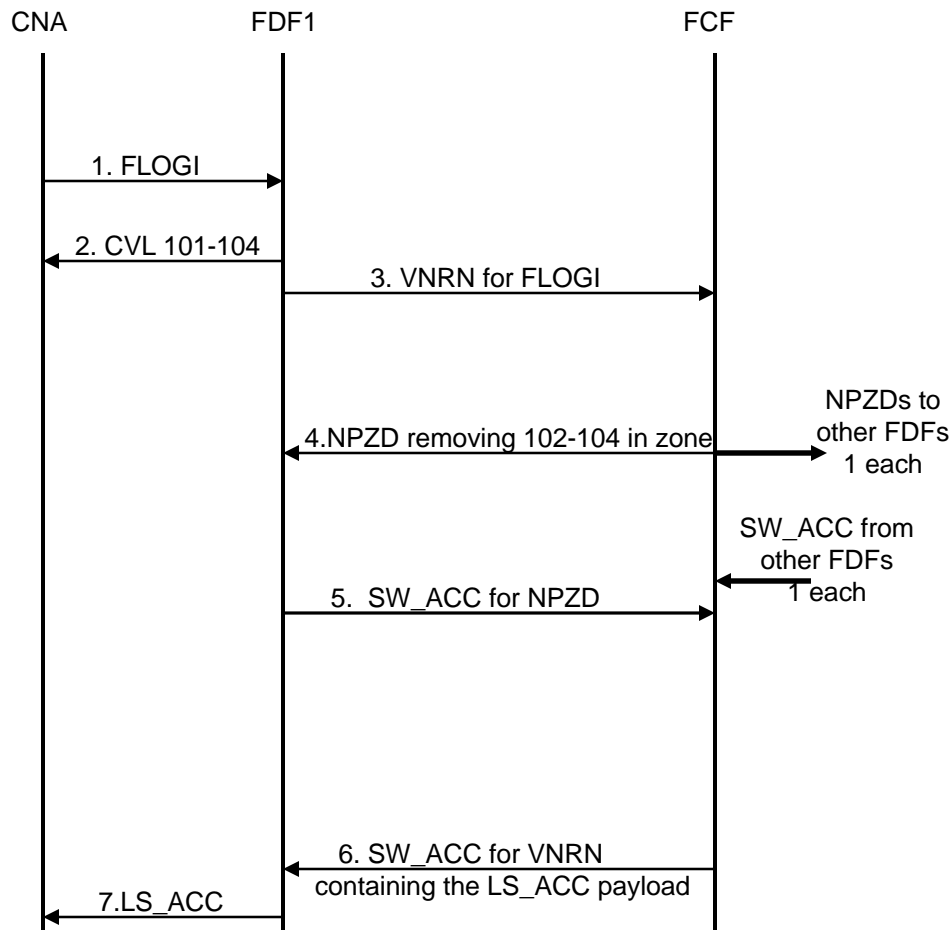
- Precondition: CNA is logged in with one FLOGI (assigned fc_id 0x000101) followed by three FDISCs (assigned fc_id IDs 0x000102, 0x000103, 0x000104)
- FLOGI is received again and triggers the “re-FLOGI” state transition defined in FC-LS-2, clause 6.4.5
 - Message 2 de-instantiates previous virtual links (depending on outcome of T11/11-409v0)
 - Messages 3-12 perform the “implicit” logout
 - Message 13-17 are login sequence for FLOGI
- Problems in this design:
 - In the current description NPZDs are initiated one-by-one for each VNUN, Calculating the NPZD zoning data is not an inexpensive operation

Proposed Re-login Process – VNUN with ID List



- **Precondition:**CNA is logged in with one FLOGI (assigned fc_id 0x000101) followed by three FDISCs (assigned fc_id IDs 0x000102, 0x000103, 0x000104)
- **FLOGI is received again and triggers the “re-FLOGI” state transition defined in FC-LS-2, chapter 6.4.5**
 - FDF1 sends the complete list of depending IDs to FCF in a single VNUN message 3
 - Therefore FCF has the complete data to create a single message 4 and only compute zoning update 1 time.
- **Proposal is to allow VNUN to specify a list of N_Port IDs, to be able to specify in ranges as in NPRD, and to remove the N_Port Name from VNUN.**

Proposed Optimized Re-login Process – Skip VNUN



- Precondition: CNA is logged in with one FLOGI (assigned `fc_id` 0x000101) followed by three FDISCs (assigned `fc_id` IDs 0x000102, 0x000103, 0x000104)
- FLOGI is received again and triggers the “re-FLOGI” state transition defined in `fc-ls-2`, chapter 6.4.5
 - **FDF1 forwards the FLOGI as defined in 11-026v2**
 - **FCF tracks the interdependencies of 101-104, registers that this is a re-FLOGI**
 - **FCF sends a single enhanced NPZD message which removes the old 101-104 and adds the new 101 at the same time**
- Note: the capability to send multiple `FC_IDs` in a VNRN should be preserved, it simplifies other implicit logout cases (such as virtual link failures).

Proposed VNUN Changes

- Remove the Unreachable N_Port Name
- Add N_Port ID Range Entries

Current VNUN Request Payload

SW_ILS Code = XX00 0002h
Originating FDF/FCDF Switch_Name
Destination Controlling FCF/Switch Switch_Name
Unreachable N_Port_Name
Unreachable N_Port_ID



Proposed Request Payload

SW_ILS Code = XX00 0002h
Originating FDF/FCDF Switch_Name
Destination Controlling FCF/Switch Switch_Name
Number of N_Port_ID Range Entries (n)
N_Port_ID Range Entry #1
N_Port_ID Range Entry #2
...
N_Port_ID Range Entry #n

NPZD with remove multiple and add capability

Item	Size (bytes)
SW_ILS Code = XX00 0006h	4
Originating Controlling FCF/Switch Switch_Name	8
Destination FDF/FCDF or Controlling FCF/Switch Switch_Name	8
Flags	4
Allocated / Deallocated N_Port_ID	4
N_Port_Name associated with the Allocated/Deallocated N_Port_ID	8
Switch_Name of the FDF/FCDF associated with the Allocated/Deallocated N_Port_ID	8
FLOGI / NPIV FDISC LS_ACC Parameters	116
Number of Peering Entries (h)	4
Peering Entry #1	see table 15
Peering Entry #2	see table 15
...	
Peering Entry #h	see table 15

Repeat this (as TLV) for each N_Port ID allocated or deallocated.

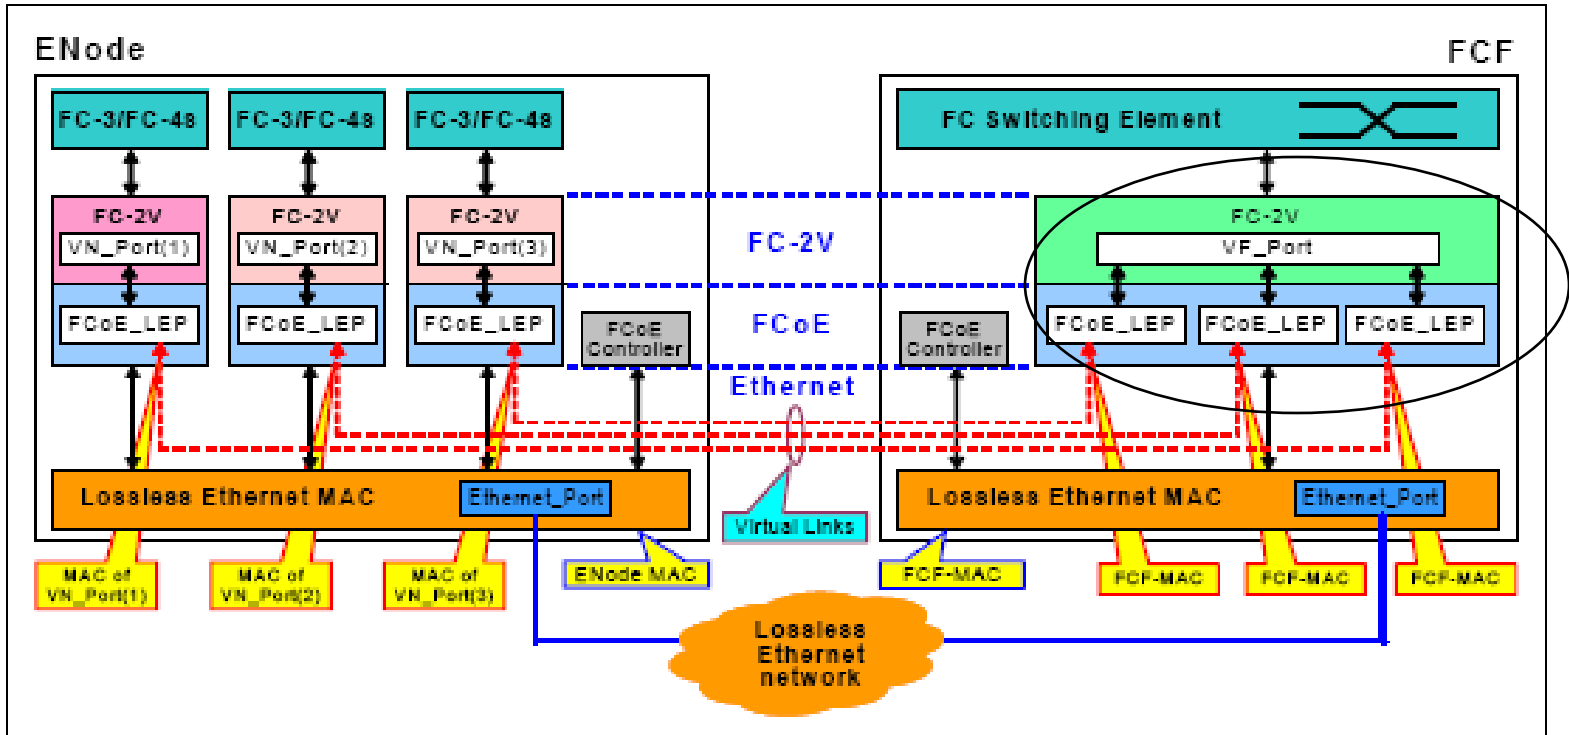
Limits on size of ILSs?

May need to send multiple NPZD

Thank you

Backups

Multiple Virtual Links between ENode and VF_Port



One
VF_Port
Multiple
VN_Ports
logged in

Figure 32 – VN_Port to VF_Port Virtual Links example

From FC-BB-5

E_Node Rules for De-instantiation of Virtual Links From FC-BB-5

- Explicit VN_Port to VF_Port Virtual Link de-instantiation is performed by an ENode MAC by performing Fabric logout. Fabric logout (i.e., Fabric LOGO) shall be performed by an ENode using FIP frames (see table 24) and the associated FIP descriptor type (see table 29).
- To deal with local physical layer faults, an ENode MAC shall de-instantiate all its VN_Ports to VF_Port Virtual Links upon detecting that its physical layer is not operational. This condition shall be handled as an implicit Fabric logout (see FC-LS-2) for the involved VN_Ports.
- If unsolicited multicast Discovery Advertisements are not received within $2.5 * FKA_ADV_PERIOD$, all the VN_Port to VF_Port Virtual Links with that VF_Port shall be implicitly de-instantiated. This condition should be counted as a Virtual Link failure and shall be handled as an implicit Fabric logout (see FC-LS-2) for the involved VN_Ports.

VF_Port Rules for De-instantiation of Virtual Links From FC-BB-5

- A VF_Port capable FCF-MAC shall de-instantiate **all its VF_Ports** upon detecting that its physical layer is not operational
- If VN_Port FIP Keep Alive frames are not received within $2.5 * FKA_VN_PERIOD$, the associated VN_Port to VF_Port Virtual Link shall be explicitly de-instantiated (i.e., a FIP Clear Virtual Links frame listing the unreachable VN_Port shall be generated). This condition shall be handled as an implicit Fabric logout (see FC-LS-2) for the involved VN_Port
- If ENode FIP Keep Alive frames are not received within $2.5 * FKA_ADV_PERIOD$, **all associated VN_Port to VF_Port Virtual Links shall be explicitly de-instantiated** (i.e., a FIP Clear Virtual Links frame listing all the unreachable VN_Ports shall be generated). This condition shall be handled as an implicit Fabric logout (see FC-LS-2) for the involved VN_Ports.
- $FKA_ADV_PERIOD < FKA_VN_PERIOD$
 - ENode FIP keep alives are used for link failure detection, VN_Port keep alives are used to detect dead VN_Ports (sent at a much slower rate)
- Explicit VN_Port to VF_Port Virtual Link de-instantiation is invoked by a VF_Port capable FCF-MAC by transmitting a FIP Clear Virtual Links frame (see 7.8.7.6). **A FIP Clear Virtual Links frame transmitted to an ENode MAC with logged in VN_Ports provides the list of VN_Ports to be removed.**

Currently Proposed FLOGI Process – with previous logins

