NPIV
Functionality Profile

1. N_Port ID Virtualization

1.1. Overview

Clause 1 outlines the goals and implementation guidelines for Fibre Channel devices to support N_Port ID Virtualization (NPIV). NPIV provides a Fibre Channel facility for sharing a single physical N_Port among multiple N_Port IDs, thereby allowing multiple initiators, each with its own N_Port ID, to share the N_Port. Implementation guidelines are presented to ensure that all FC devices implement a common methodology. The purpose of this clause is to define the functionality that needs to be supported.

1.2. N_Port ID Virtualization Acquisition Procedure

This procedure defines the method by which an N_Port attached to a fabric acquires additional N_Port_IDs from the Fabric. N_Port ID Virtualization is not supported in an arbitrated Loop Topology. Figure 1 shows a flow diagram of the process.

NPIV
Functionality Profile

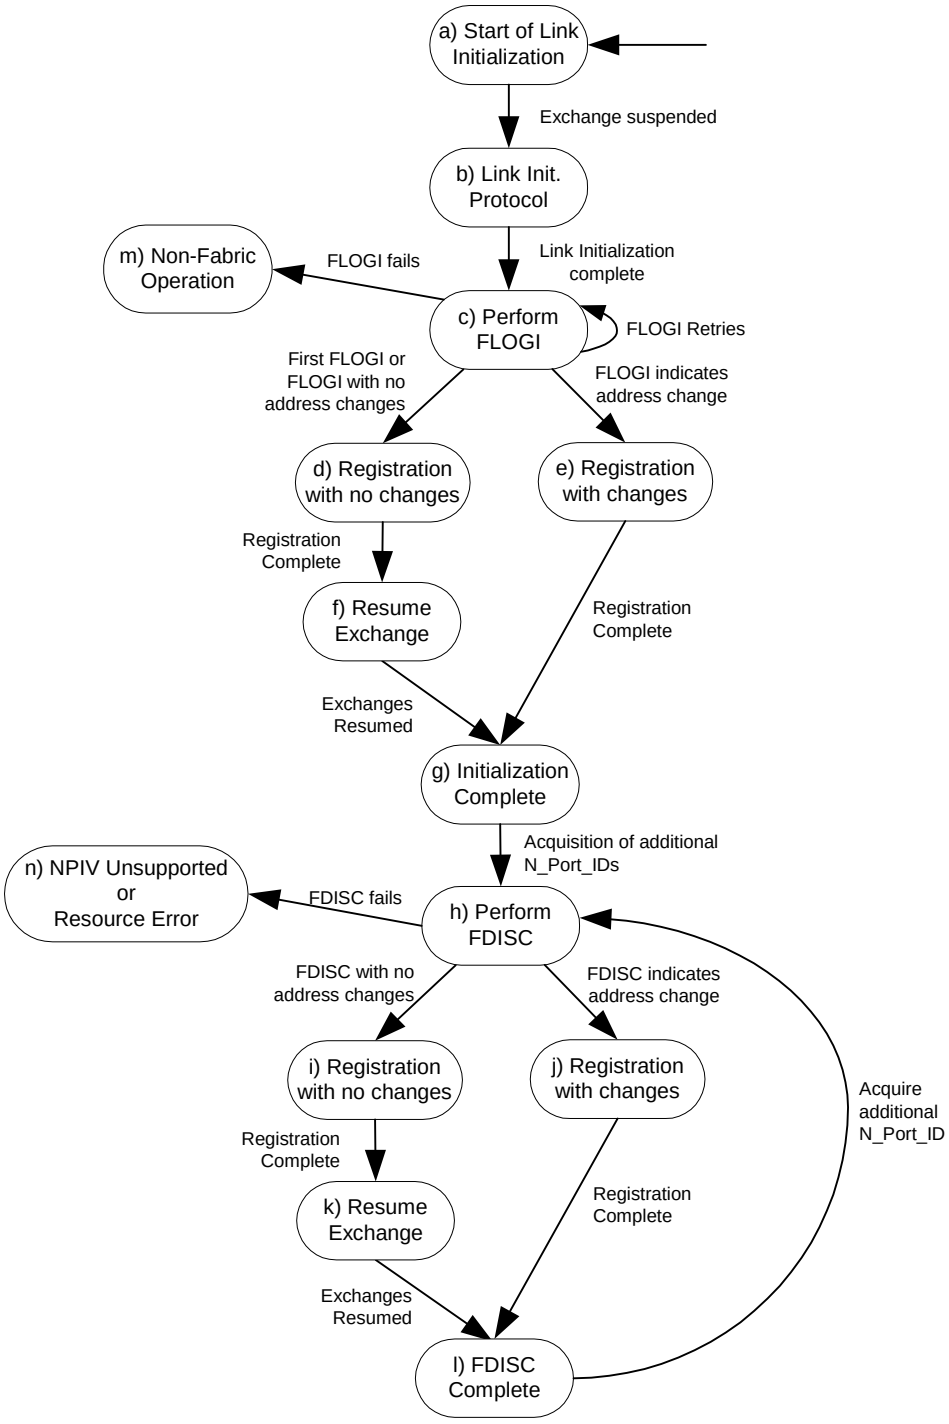

Figure 1 – N_Port_ID Acquisition Flow

NPIV
Functionality Profile

NOTE 1 – With the exception of the removal of Loop Behavior, the Link Initialization Process and FLOGI Procedure described below were extracted from the Nx_Port Initialization Procedure defined in FC-MI. Compliant implementations are required to implement the features defined in this procedure and, when applicable, FC-MI.

a) **Start of Link Initialization**. Whenever an N_Port receives an OLS, NOS, loss of synchronization for greater than R_T_TOV, or loss of signal, the N_Port shall begin the Link Initialization protocol as defined in FC-FS. The N_Port shall suspend all open Exchanges with all other Nx_Ports, implicitly logout with the Fabric, and shall proceed to step (b).
b) **Link Initialization Protocol.** Upon completion of the Link Initialization protocol, as defined in FC-FS, the N_Port shall proceed to step (c).
c) **Acquisition of FLOGI assigned N_Port_ID - Perform FLOGI**. The N_Port shall perform the following.
   A) The N_Port shall attempt to send FLOGI to the Fabric F_Port. The N_Port Originator shall set the D_ID to a well-known F_Port_ID (i.e., hex 'FF FF FE') and the S_ID to hex '000000'. The payload of the FLOGI Request Sequence contains an N_Port_Name and Node_Name associated with the physical N_Port initiating the login. The Common Service Parameters, Class 2 and 3 Service Parameters shall follow the rules as defined in FC-MI.
   B) If the N_Port receives a PLOGI Request, it may either delay reply until the FLOGI is complete or respond to the PLOGI Request Sequences with an LS_RJT Reply Sequence, with a reason code of "Unable to perform command request."
   C) If the N_Port detects that the FLOGI Request attempt fails (e.g., receiving an LS_RJT), or if the FLOGI ACC Sequence is not received within 2xR_A_TOV of the FLOGI Request, it may either:
      aa) retry the attempt in step (A) after R_A_TOV; or,
      bb) proceed to step (m).
   D) If the FLOGI is completed successfully; the first FLOGI initiated by the N_Port or the N_Port had previously completed a FLOGI; the N_Port determines that the N_Port has the same N_Port_ID, N_Port_Name and Node_Name that the N_Port had before link initialization; and the N_Port determines that the F_Port has the same F_Port_Name and Fabric_Name that the F_Port had before link initialization, the N_Port shall proceed to step (d).
   E) If the FLOGI is completed successfully and any of the conditions in step (D) are not met, the N_Port shall proceed to step (e).

NOTE 2 – It is possible for an N_Port to receive ELS frames before any PLOGI frames are received.

d) **Registration with no address change.** The N_Port may optionally login (PLOGI) with the Directory Server (hex 'FFFFFC') and perform registration, deregistration and queries with the Name Server. The Name Server requests may be performed in any order. The N_Port shall proceed to step (f).
e) **Registration with address change**. An N_Port performing this step has determined that its own addressing information and/or that of the F_Port have changed, or that it did not have complete addressing information, prior to this initialization. The N_Port shall perform the following.
   A) The N_Port shall discard all suspended and queued Exchanges in a manner consistent with the ULP, implicitly LOGO with all other Nx_Ports, and perform ULP-specific actions to clear pending tasks.
   B) The N_Port shall perform PLOGI explicitly with the Directory Server (hex 'FFFFFC') and shall perform registration with the Name Server. The Name Server Registration requests may be performed in any order.

C) If the Clean Address bit in the FLOGI Accept (ACC) is set to one, the N_Port shall proceed to step (g); otherwise, the N_Port shall wait for R_A_TOV before originating any new Exchanges. During this time, the N_Port shall discard all frames received except for the PLOGI Request Sequences and LOGO ACC reply Sequences. The N_Port shall respond to all PLOGI Request Sequences with an LS_RJT Reply Sequence, with a reason code of "Unable to perform command request." The N_Port shall proceed to step (g).

f) **Resume Exchanges**. The N_Port shall resume all suspended Exchanges (if any) with other Nx_Ports. The N_Port shall proceed to step (g).

g) **Completion of N_Port initialization**. This completes initialization for the N_Port. The N_Port may proceed to originate and respond to Exchanges and login with other Nx_Ports as needed, or attempt to acquire additional N_Port_IDs. If the N_Port attempts to acquire additional N_Port_IDs, the N_Port shall proceed to step (h).

h) **Acquisition of FDISC assigned N_Port_ID - Perform FDISC**. Upon successful completion of the FLOGI Request Sequence the N_Ports shall perform the following.

A) The N_Port shall attempt to send FDISC to the Fabric F_Port. The N_Port Originator shall set the D_ID to a well-known F_Port_ID (i.e., hex 'FF FF FE') and the S_ID to hex '000000'. The payload of the FDISC Request contains an N_Port_Name and Node_Name associated with the source of the FDISC Request. The Service Parameters should be identical to the parameters defined in the last FLOGI payload originated by the N_Port.

NOTE 3 – When the S_ID of the FDISC ELS is set to a previously assigned N_Port_ID, with corresponding N_Port_Name and Node_Name, the FDISC ELS shall provide the means for exchanging Service Parameters between the N_Port and F_Port. The interchange of FDISC information shall not modify the operating environment or Service Parameters between the two ports.

B) The F_Port shall either reply with a LS_ACC or LS_RJT reply Sequence. The D_ID field of the LS_ACC shall be set to the additional N_Port_ID. The Service Parameters should be identical to the parameters defined in the FDISC Request originated by the N_Port. However, the Service Parameters contained in the FDISC Request shall be ignored and no error condition shall be reported regardless of whether they are the same as the service parameters contained in the last FLOGI Request. An LS_RJT reply Sequence has the following implications:

aa) a reason code of "Command not supported" and a reason code explanation of "Request not supported" indicate that the Fabric does not support NPIV;

bb) a reason code of "Unable to perform command request" and a reason code explanation of "Insufficient resources to support Login" indicates that the Fabric is unable to support additional FDISC assigned N_Port_IDs; or,

cc) a reason code of "Unable to perform command request" and a reason code explanation of "Login required" indicates that there are no N_Port_IDs assigned to the N_Port, and a FLOGI is required.

C) If the N_Port detects that the FDISC Request attempt fails (e.g., receiving an LS_RJT), or if the LS_ACC Sequence is not received within 2xR_A_TOV of the FDISC Request, it may either:

aa) retry the attempt in step (A) after R_A_TOV; or,

bb) proceed to step (n).

D) If the FDISC is completed successfully; the N_Port determines that the FDISC assigned N_Port_ID has the same associated N_Port_Name and Node_Name that it had after FLOGI or FDISC; and the N_Port determines that the F_Port has the same F_Port_Name and Fabric_Name that the F_Port had after FLOGI, the N_Port shall proceed to step (A).

E) If the FDISC is completed successfully and any of the conditions in step (D) are not met, the N_Port shall proceed to step (j).

i) **Registration with no Address change.** The N_Port, associated with the FDISC assigned N_Port_ID, may optionally login (PLOGI) with the Directory Server (hex 'FFFFFC') and perform registration, deregistration and queries with the Name Server. The Name Server requests may be performed in any order. The N_Port shall proceed to step (k).

j) **Registration with Address change**. An N_Port performing this step has determined that the associated N_Port_ID information and/or that of the F_Port have changed after FLOGI or FDISC. The N_Port shall perform the following.
   A) The N_Port shall discard all suspended and queued Exchanges, associated with the FDISC assigned N_Port_ID, in a manner consistent with the ULP, perform ULP-specific actions to clear pending tasks and implicitly logout with all the Nx_Ports associated with the N_Port_ID.
   B) The FDISC assigned N_Port_ID shall perform PLOGI explicitly with the Directory Server (hex 'FFFFFC') and shall perform registration with the Name Server. The Name Server Registration requests may be performed in any order.
   C) If the Clean Address bit in the FDISC Accept (ACC) is set to one, the N_Port shall proceed to step (l); otherwise, the N_Port shall wait for R_A_TOV before originating any new Exchanges associated with the FDISC assigned N_Port_ID. During this time, all frames received with a destination address of the FDISC assigned N_Port_ID shall be discarded except for the PLOGI Request Sequences and LOGO ACC reply Sequences. The N_Port shall respond to all PLOGI Request Sequences with a reason code of "Unable to perform command request." The N_Port shall proceed to step (l).

k) **Resume Exchanges**. The N_Port shall resume all suspended Exchanges (if any) with other Nx_Ports. The N_Port shall proceed to step (l).

l) **Completion of FDISC**. This completes the acquisition of an additional FDISC assigned N_Port_ID. The FDISC assigned N_Port_ID may proceed to originate and respond to Exchanges and login with other Nx_Ports as needed, or attempt to acquire additional N_Port_IDs. If the N_Port attempts to acquire additional N_Port_IDs, the N_Port shall proceed to step (h).

m) **Fall-back to non-Fabric operation**. The N_Port is not connected to a Fabric.

n) **NPIV Unsupported or Resource Error.** If it is determined that the Fabric does not support NPIV or a resource error occurred due to resource limitations, the N_Port may proceed to originate and respond to Exchanges, and login with other Nx_Ports as needed. If FDISC failed due to an incomplete FLOGI, the N_Port may login with the Fabric and acquire additional N_Port_IDs (if supported).

NOTE 4 – Error detection/recovery (i.e., Timeouts) should be implemented as defined in FC-FS.

1.3. Buffer-to-buffer flow control management

N_Port ID Virtualization shall support Buffer-to-buffer flow control as specified in FC-FS to include:

a) Buffer-to-buffer control occurs in both directions between the local N_Port and local F_Port.
b) BB_Credit represents the number of receive buffers supported by an N_Port for receiving Fibre Channel frames. The total number of receive buffers is mutually inclusive of all the Sequence Initiators and Recipients associated with all N_Port_IDs.
c) BB_Credit_CNT is defined as the number of unacknowledged or outstanding frames awaiting R_RDY responses from the directly attached Port. As defined in FC-FS, the N_Port increments BB_Credit_CNT by one for each frame transmitted and decrements by one for each R_RDY received. The total number of transmitted frames, for which R_RDY responses

are outstanding, is mutually inclusive of all the Sequence Initiators and Recipients associated with all N_Port_IDs.

1.4. Priority

When supported, Priority allows end devices and/or Switches in a Fabric to resolve resource contention or to determine the order in which to deliver frames. The Priority field in the Frame Header indicates the Priority assigned to the frame. Compliant devices shall implement Priority (if supported) as specified in FC-FS with the following restrictions:

a) The Sequence Initiator shall set the Priority, for the duration of the Sequence Initiative, to the same value for all frames.
b) If multiple sequences occur, it is the responsibility of the Sequence Initiator to set the Priority of each frame in each Sequence to the same value.

1.5. Logout Procedure

1.5.1. Overview

The destination Logout procedure, as defined in FC-FS, provides a method for removing service between an N_Port and Nx_Ports or for removing an N_Port_ID which was previously assigned by the Fabric. Logout releases resources associated with maintaining Service with a Fabric and/or destination Nx_Port.

1.5.2. Explicit N_Port_ID Logout

Explicit logout is accomplished by transmitting a Logout (LOGO) Extended Link Service (ELS) request Sequence to the Fabric (well-known address hex ‘FF FF FE’) or a destination Nx_Port. The Logout procedure is complete when the responding Fabric or Nx_Port transmits a LS_ACC Link Service reply Sequence.

The LOGO ELS, when sent to the Fabric, shall request the removal and logout of the previously assigned N_Port_ID specified in the S_ID field. Logout releases resources associated with maintaining Service with Fabric.

NOTE 6 - The name server may generate an explicit LOGO if it logs a Port Identifier out due to resource limitations.

The LOGO ELS, when sent to the destination Nx_Port, shall request the removal of resources associated with a particular N_Port_ID, which was previously assigned by an F_Port. The N_Port_ID and N_Port_Name of the N_Port requesting Logout are identified in the Payload. Both the source N_Port and the destination Nx_Port of the Logout Request Sequence shall abnormally terminate all open Exchanges that used the N_Port_ID indicated in the Payload of the Logout Request Sequence. Fabric Login shall be required following Explicit Logout of all the N_Port_IDs associated with an N_Port. Communication with other Nx_Ports shall not be accepted until the Fabric Login procedure is complete.

1.5.3. Implicit N_Port_ID Logout

If an N_Port receives or transmits the NOS or OLS Primitive Sequence, the N_Port and all its associated N_Port_IDs (if any) shall be implicitly logged out from the Fabric. Fabric Login shall be required following implicit Logout. Communication with other Nx_Ports shall not be accepted until the Fabric Login procedure is complete.