American National Standard
for Information Technology —

Fibre Channel —
NVMe (FC-NVMe)

4.3.1 NVMeoFC association overview

An NVMeoFC association is an NVMeoFC layer abstraction for an exclusive communication relationship that is established between a particular NVMe host, connected via a particular initiator NVMe_Port, and a particular NVMe controller in an NVM subsystem connected via a particular target NVMe_Port. The association encompasses the controller, its state and properties, its Admin Queue, and all I/O Queues of that controller (see NVMe over Fabrics).

The NVMeoFC association is created by transmitting a Create_Association NVMe_LS request (see 8.3.3). If the target NVMe_Port and NVM subsystem allow the communication relationship to be created, the target NVMe_Port transmits a Create_Association NVMe_LS accept payload (see 8.3.3) to the initiator NVMe_Port. The Create_Association accept payload contains an Association Identifier that shall be used by the NVMeoFC layer on the initiator NVMe_Port to refer to the NVMeoFC association in subsequent Fabric traffic transmitted to the target NVMe_Port. If the NVMeoFC association cannot be created, the target NVMe_Port shall transmit an NVMe_RJT (see 8.3.1) to the initiator NVMe_Port with the reason code and reason code explanation set to an appropriate value.

An NVMe over Fabrics association is established when the NVMe-oF Connect command (see NVMe over Fabrics) to create the Admin Queue is issued on the NVMeoFC association's Admin Queue connection. Refer to NVMe over Fabrics for additional requirements and behaviors of NVMe over Fabrics associations and Admin Queue creation.

An active NVMeoFC association shall be terminated if:

a) any NVMeoFC connection (see 4.4) for the NVMeoFC association is terminated;
b) a Disconnect NVMe_LS request (see 8.3.5) is received with an Association Identifier descriptor containing the Association Identifier of the NVMeoFC association;
c) one of the clearing effects (see 4.15) occur which terminates the NVMeoFC association; or
d) the NVMe host or NVM subsystem detects a condition which causes it to terminate the corresponding NVMe-oF association.

An initiator NVMe_Port or a target NVMe_Port may terminate an NVMeoFC association. While an active NVMeoFC association is being terminated, the NVMe_Port shall:

a) not transmit NVMe_CMND IUs or NVMe_LS requests other than a Disconnect NVMe_LS, that contain the Association Identifier or Connection Identifiers that correspond to the NVMeoFC association being terminated;
b) discard any received NVMe_CMND IUs or NVMe_LS requests other than a Disconnect NVMe_LS, that contain the Association Identifier or Connection Identifiers that correspond to the NVMeoFC association being terminated;

c) if the NVMe_Port is an initiator NVMe_Port, then perform the initiator NVMe_Port NVMeoFC association termination process (see 4.3.2); and
d) if the NVMe_Port is a target NVMe_Port, then perform the target NVMe_Port NVMeoFC association termination process (see 0.0.3).

If an NVMeoFC association is terminated, the NVMeoFC layer on the initiator NVMe_Port and target NVMe_Port shall implicitly terminate all Admin Queue and I/O Queue connections for the association.

4.3.2 Initiator NVMe_Port NVMeoFC association termination process

The initiator NVMe_Port NVMeoFC association termination process is performed by an initiator NVMe_Port to terminate the NVMeoFC association and associated NVMeoFC connections, and recover all outstanding exchange resources that are associated with the NVMeoFC association to be terminated.

An initiator NVMe_Port shall perform the following steps to terminate an NVMeoFC association:

1) transmit an ABTS-LS (see 11.3) for all open Exchanges for the NVMeoFC association being terminated, except for a Disconnect NVMe_LS Exchange;
2) transmit a Disconnect NVMe_LS request to the corresponding target NVMe_Port, unless there is no valid login with the associated NVMe_Port. The Disconnect NVMe_LS request, if transmitted, shall contain an Association Identifier descriptor with the Association Identifier of the NVMeoFC association being terminated; and
3) await recovery of all open Exchange resources as described in 4.3.3 for the NVMeoFC association being terminated.

If the initiator NVMe_Port has received or receives a Disconnect NVMe_LS request for the NVMeoFC association being terminated, then the initiator NVMe_Port shall not transmit the Disconnect NVMe_LS response until the initiator NVMe_Port has completed steps 1 and 2 of the initiator NVMe_Port NVMeoFC association termination process for the association to be terminated.

If the Disconnect NVMe_LS response to the Disconnect NVMe_LS request, if transmitted as described in step 2 in this subclause, is received after the NVMeoFC association is terminated, then the response shall be ignored by the initiator NVMe_Port.

If a response to the Disconnect NVMe_LS request, if transmitted as described in step 2 in this subclause, is not received within two times R_A_TOV, then the initiator NVMe_Port shall:

a) send an ABTS-LS to recover the Exchange resources for the Disconnect NVMe_LS request; and
b) if the Exchange resources for all open Exchanges associated with the NVMeoFC association to be terminated have not been recovered, the initiator NVMe_Port shall transmit a second Disconnect NVMe_LS request containing the same request payload in a separate Exchange, or transmit a LOGO ELS to the corresponding target NVMe_Port. If a response was not received for the second Disconnect NVMe_LS request, if transmitted, after two times R_A_TOV, then the initiator NVMe_Port shall transmit a LOGO ELS to the corresponding target NVMe_Port.

4.3.3 Initator NVMe_Port Exchange recovery

For an NVMeoFC association being terminated, an initiator NVMe_Port shall consider an open Exchange resource recovered:

a) upon the reception of an ABTS-LS from the corresponding target NVMe_Port for the Exchange;
b) upon the reception of a BA_ACC or BA_RJT to the ABTS-LS issued for the Exchange as described in step 1 in 4.3.2 ;
c) after a delay of R_A_TOV after the reception of a Disconnect NVMe_LS request that contains the Association Identifier descriptor with the Association Identifier set to the identifier of the NVMeoFC association being terminated;
d) after a delay of R_A_TOV after the reception of the Disconnect NVMe_LS response to the Disconnect NVMe_LS request transmitted as described in step 2 in 4.3.2 ; or
e) upon a logout or process logout with the target NVMe_Port.

4.3.4 Target NVMe_Port NVMeoFC association termination process

The target NVMe_Port NVMeoFC association termination process is performed by a target NVMe_Port to terminate the NVMeoFC association and associated NVMeoFC connections, recover all outstanding exchange resources that are associated with the NVMeoFC association to be terminated, and recover the Association Identifier and Connection Identifiers corresponding to the NVMeoFC association to be terminated.

A target NVMe_Port shall perform the following steps to terminate an NVMeoFC association:

1) transmit an ABTS-LS (see 11.3) for each open Exchange for the NVMeoFC association being terminated, except for a Disconnect NVMe_LS Exchange;
2) transmit a Disconnect NVMe_LS request to the corresponding initiator NVMe_Port unless there is no valid login with the associated NVMe_Port. The Disconnect NVMe_LS request, if transmitted, shall contain an Association Identifier descriptor with the Association Identifier set to the identifier of the NVMeoFC association being terminated;
3) wait for:
   A) reception of a Disconnect NVMe_LS request from the intiator NVMe_Port that contains the Association Identifier descriptor with the Association Identifier set to the identifier of the NVMeoFC association being terminated;
   B) reception of the Disconnect NVMe_LS response to the Disconnect NVMe_LS request, if transmitted, as described in step 2 in this subclause. If that response is an NVMe_RJT and the Link Service Reject descriptor has the reason code set to a value other than 40h (i.e., Invalid Association Identifier), then the target NVMe_Port shall wait for the condition described in step 3)A) in this subclause or for the condition described in step 3)C) in this subclause. If that response is an NVMe_ACC or that response is an NVMe_RJT and the Link Service Reject descriptor is set to a reason code of 40h (i.e, Invalid Association Identifier), the target NVMe_Port may proceed with step 4 in this subclause; or
   C) a delay of four times R_A_TOV. After that delay, the target NVMe_Port shall transmit a LOGO ELS to the corresponding initiator NVMe_Port.
4) await recovery of all open Exchange resources as described in 4.3.5 for the NVMeoFC association being terminated; and
5) make the Association Identifier and Connection Identifiers corresponding to the terminating NVMeoFC association available for re-use:
   A) after a R_A_TOV delay; or
   B) upon a logout or Process Logout with the initiator NVMe_Port.

If the target NVMe_Port has received or receives a Disconnect NVMe_LS request for the NVMeoFC association being terminated, the target NVMe_Port shall not transmit the Disconnect NVMe_LS response until the target NVMe_Port has completed steps 1 and 2 of the target NVMe_Port NVMeoFC association termination process for the association to be terminated.

If the Disconnect NVMe_LS response to the Disconnect NVMe_LS request, if transmitted as described in step 2 in this subclause, is received after the NVMeoFC association is terminated, then the response shall be ignored by the target NVMe_Port.

If the Disconnect NVMe_LS response to the Disconnect NVMe_LS request, if transmitted as described in step 2 in this subclause, is not received within two times R_A_TOV, then the target NVMe_Port shall:

a) send an ABTS-LS to recover the Exchange resources for the Disconnect NVMe_LS request; and
b) if the Exchange resources for all open Exchanges associated with the NVMeoFC association to be terminated have not been recovered, then the target NVMe_Port shall transmit a second Disconnect NVMe_LS request containing the same request payload in a separate Exchange, or transmit a LOGO ELS to the corresponding initiator NVMe_Port. If a response was not received for the second Disconnect NVMe_LS request, if transmitted, after two times R_A_TOV, then the target NVMe_Port shall transmit a LOGO ELS to the corresponding initiator NVMe_Port.

4.3.5 Target NVMe_Port Exchange recovery

For an NVMeoFC association being terminated, a target NVMe_Port shall consider an open Exchange resource recovered upon:

a) the reception of an ABTS-LS from the corresponding initiator NVMe_Port for the Exchange;
b) the reception of a BA_ACC or BA_RJT to the ABTS-LS issued for the Exchange as described in step 1 in 0.0.3;
c) the reception of an NVMe_ACC response to the Disconnect NVMe_LS request, if transmitted, as described in step 2 in 0.0.3;
d) the reception of an NVMe_RJT response to the Disconnect NVMe_LS request, if transmitted as described in step 2 in 0.0.3, and that response has the Link Service Reject descriptor set to a reason code 40h (i.e., Invalid Association Identifier).
e) the reception of a Disconnect NVMe_LS request that contains the Association Identifier descriptor with the Association Identifier set to the identifier of the NVMeoFC association being terminated; or
f) a logout or Process Logout with the initiator NVMe_Port.

9.6 NVMe_ERSP IU format

Table 36 – ERSP Result field values

Value	Name	Description
00h	SUCCESS	No status.
01h	INVALID FIELD	NVMe_CMND IU field is invalid.
02h	-	Reserved
03h	ILLEGAL CONNECT PARAMETERS	Connect command parameters are invalid for the NVMeoFC connection.
Others	-	Reserved

11.3.2 Initiating NVMe_Port Exchange termination

The NVMe_Port terminating the Exchange shall transmit an ABTS-LS to the D_ID of the corresponding NVMe_Port of the Exchange being terminated. The ABTS-LS shall be generated using the OX_ID field and RX_ID field values of the Exchange to be aborted. FC-FS-5 allows ABTS-LS to be transmitted by an Nx_Port regardless of whether or not it has Sequence Initiative. Following the transmission of ABTS-LS, any Device_Data Frames received for the Exchange being terminated shall be discarded until the BA_ACC or BA_RJT with the F_CTL field Last_Sequence bit set to one (i.e., last Sequence of the Exchange) is received from the corresponding NVMe_Port.

The NVMe_Port terminating the Exchange may reclaim the Exchange resources without requiring the reception of a BA_ACC or BA_RJT or performing second level error recovery (see 11.4) if the NVMe_Port is an initiator NVMe_Port and one of the Exchange recovery options are met as described in 4.3.3, or if the NVMe_Port is a target NVMe_Port and one of the Exchange recovery options are met as described in 4.3.5.

If the NVMe_Port is an initiator NVMe_Port and a BA_ACC or BA_RJT response is not received from the corresponding NVMe_Port within two times R_A_TOV, and the Exchange resource has not been recovered by other options as described in 4.3.3, then second level error recovery (see 11.4) shall be performed.

If the NVMe_Port is a target NVMe_Port and a BA_ACC or BA_RJT response is not received from the corresponding NVMe_Port within two times R_A_TOV, and the Exchange resource has not been recovered by other options as described in 4.3.5, then second level error recovery (see 11.4) shall be performed.

11.4.1 ABTS-LS error recovery

If a response to an ABTS-LS is not received within two times R_A_TOV, and the Exchange resource has not been recovered by other options as described in 4.3.3 for an initiator NVMe_Port and in 4.3.5 for a target NVMe_Port, then the NVMe_Port may transmit the ABTS-LS again, attempt other retry operations allowed by FC-FS-5, or explicitly logout the corresponding NVMe_Port. If those retry operations attempted are unsuccessful, then the NVMe_Port shall explicitly logout (i.e., transmit a LOGO ELS) the corresponding NVMe_Port. All outstanding Exchanges, as well as all NVMeoFC connections and NVMeoFC associations with the corresponding NVMe_Port, shall be terminated at the NVMe_Port.

11.6.6 Clearing effects of Disconnect NVMe_LS

A Disconnect NVMe_LS causes an NVMeoFC association to be terminated. Transmission or reception of a Disconnect NVMe_LS also cause all related Exchanges and connections to be terminated (see 4.3.2 and 4.3.4).