

iscsi state diagrams

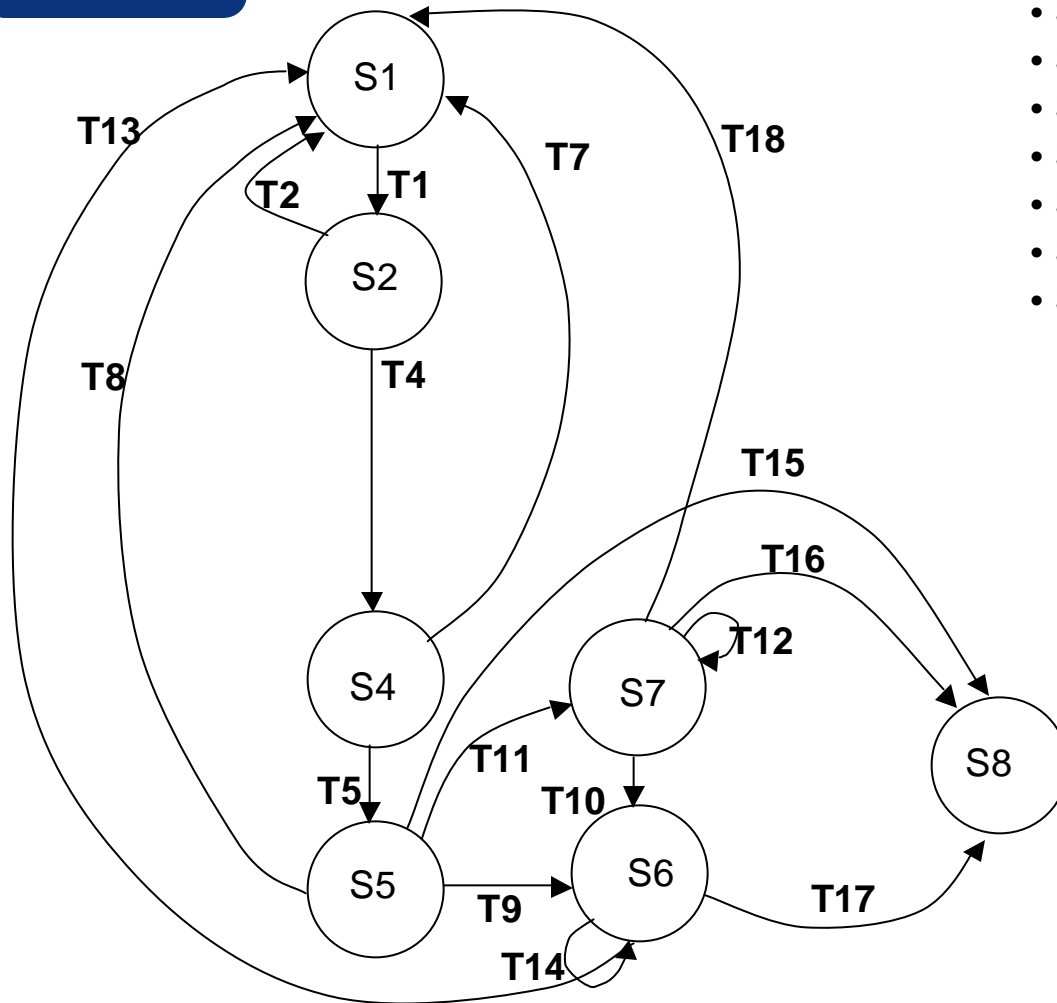
Design philosophy

- In an iSCSI connection closure, clear separation must be maintained between transition sequences leading to a properly shutdown iSCSI FFP connection (leading to S1 in the subsequent diagrams), and transition sequences that lead to an improperly shutdown iSCSI FFP connection (leading to S8 in the subsequent diagrams). Only the sequences terminating in S8 require the connection state machine to additionally traverse the “connection cleanup” transitions to eventually effect a proper closure of the state machine.
- The following four are considered the fundamental transport layer abstractions that the iSCSI state machine needs to operate on top of –
 - that the underlying transport is connection-oriented,
 - that a connect request-response (success/failure) API exists to the transport for an active open, and a connect indication-confirmation (success/failure) API exists for a passive open.
 - that there is a disconnection indication-response API exists to the transport,
 - transport can generate RESETs resulting in instantaneous termination of a connection.
- A TCP half-close does not have to be reflected in iSCSI state machines initiating transport connection closure. The half-close dynamics are an internal matter for the TCP software/hardware to handle.

iSCSI standard connection state transitions

Symbolic Names for States

- S1: FREE
- S2: XPT_WAIT
- S4: IN_LOGIN
- S5: LOGGED_IN (full-feature phase)
- S6: IN_LOGOUT
- S7: LOGOUT_REQUESTED
- S8: CLEANUP_WAIT

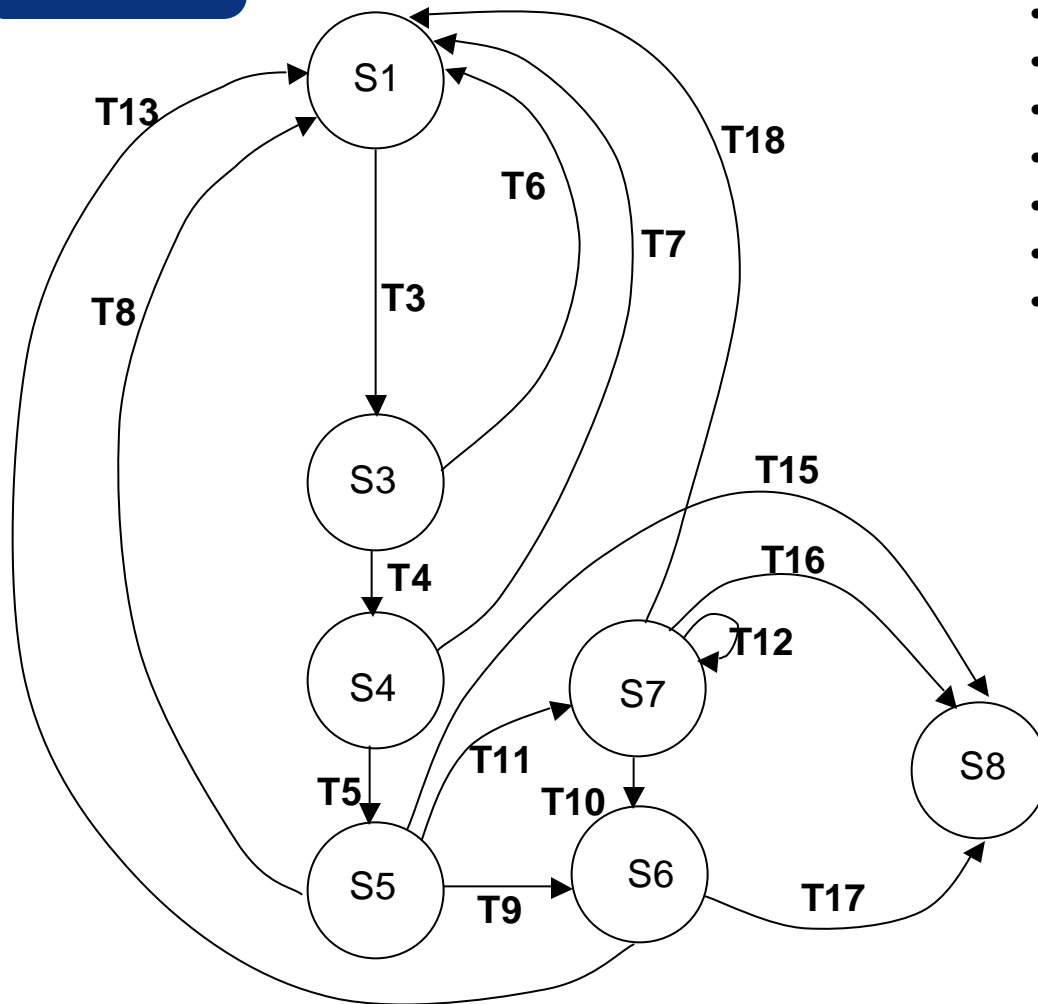


standard connection state diagram for an initiator

iSCSI standard connection state transitions (contd.)

Symbolic Names for States

- S1: FREE
- S3: XPT_UP
- S4: IN_LOGIN
- S5: LOGGED_IN (full-feature phase)
- S6: IN_LOGOUT
- S7: LOGOUT_REQUESTED
- S8: CLEANUP_WAIT



standard connection state diagram for a target

state transition table for an initiator

		T O						
		S1	S2	S4	S5	S6	S7	S8
FROM	S1	-	T1	-	-	-	-	-
	S2	T2	-	T4	-	-	-	-
	S4	T7	-	-	T5	-	-	-
	S5	T8	-	-	-	T9	T11	T15
	S6	T13	-	-	-	T14	-	T17
	S7	T18	-	-	-	T10	T12	T16
	S8	-	-	-	-	-	-	-

state transition table for a target

		S1	S3	S4	S5	S6	S7	S8
FROM	S1	-	T3	-	-	-	-	-
	S3	T6	-	T4	-	-	-	-
	S4	T7	-	-	T5	-	-	-
	S5	T8	-	-	-	T9	T11	T15
	S6	T13	-	-	-	-	-	T17
	S7	T18	-	-	-	T10	T12	T16
	S8	-	-	-	-	-	-	-

Note: “-” stands for “undefined”.

State descriptions for the standard connection state diagram

- S1: FREE
 - initiator: state on instantiation, or after successful connection closure.
 - target: state on instantiation, or after successful connection closure.
- S2: XPT_WAIT
 - initiator: waiting for a response to its transport connection establishment request.
 - target: illegal
- S3: XPT_UP
 - initiator: illegal
 - target: waiting for the Login process to commence.
- S4: IN_LOGIN
 - initiator: waiting for the Login process to conclude, possibly involving several PDU exchanges.
 - target: waiting for the Login process to conclude, possibly involving several PDU exchanges.
- S5: LOGGED_IN
 - initiator: in full-feature phase, waiting for all internal, iSCSI and transport events.
 - target: in full-feature phase, waiting for all internal, iSCSI and transport events.
- S6: IN_LOGOUT
 - initiator: waiting for a Logout response.
 - target: waiting for an internal event signalling completion of logout processing.
- S7: LOGOUT_REQUESTED
 - initiator: waiting for an internal event signalling readiness to proceed with Logout.
 - target: waiting for the Logout process to start after having requested a Logout via an Async Message.

State descriptions for the standard connection state diagram (contd.)

- S8: CLEANUP_WAIT
 - initiator: waiting for an internal event signalling readiness to proceed with connection cleanup.
 - target: waiting for the cleanup process to start for this CSM.

State transition descriptions for the standard connection state diagram

- T1:
 - initiator: Transport connect request was made (ex: TCP SYN sent).
 - target: illegal
- T2:
 - initiator: Transport connection request timed out, or failed.
 - target: illegal
- T3:
 - initiator: illegal
 - target: Received a valid transport connection request, establishing the transport connection.
- T4:
 - initiator: Transport connection established, prompting the initiator to start the iSCSI Login.
 - target: Initial iSCSI Login command was received.
- T5:
 - initiator: The final iSCSI Login response with a status class of zero was received.
 - target: The final iSCSI Login command to conclude the Login phase was received, prompting the target to send the final iSCSI Login response with a status class of zero.
- T6:
 - initiator: illegal
 - target: Timed out waiting for an iSCSI Login, or transport disconnect indication was received, or transport reset was received, or an internal event indicating a transport timeout was received, or an internal event of sending a Logout response (success) on another connection for a “close the session” Logout command was received. In all these cases, the connection is to be closed.

State transition descriptions for the standard connection state diagram (contd.)

- T7:
 - initiator: The final iSCSI Login response was received with a non-zero status class, or Login timed out, or transport disconnect indication was received, or transport reset was received, or an internal event indicating a transport timeout was received, or an internal event of receiving a Logout response (success) on another connection for a “close the session” Logout command was received. In all these cases, the transport connection is closed.
 - target: The final iSCSI Login command to conclude the Login phase was received, prompting the target to send the final iSCSI Login response with a non-zero status class, or Login timed out, or transport disconnect indication was received, or transport reset was received, or an internal event indicating a transport timeout was received, or an internal event of sending a Logout response (success) on another connection for a “close the session” Logout command was received. In all these cases, the connection is to be closed.
- T8:
 - initiator: An internal event of receiving a Logout response (success) on a nother connection for a “close the session” Logout command was received, thus closing this connection requiring no further cleanup.
 - target: An internal event of sending a Logout response (success) on another connection for a “close the session” Logout command was received, thus prompting the target to close this connection cleanly.
- T9, T10:
 - initiator: An internal event indicating the readiness to start the Logout process was received, prompting an iSCSI Logout to be sent by the initiator.
 - target: An iSCSI Logout command was received.
- T11, T12:
 - initiator: Async PDU with iSCSI event “Request Logout” was received.
 - target: An internal event requiring decommissioning of the connection is received, causing an Async PDU with iSCSI event “Request Logout” to be sent.

State transition descriptions for the standard connection state diagram (contd.)

- T13:
 - initiator: An iSCSI Logout response (success) was received, or an internal event of receiving a Logout response (success) on another connection for a “close the session” Logout command was received.
 - target: An internal event indicating successful processing of the Logout was received, prompting an iSCSI Logout response (success) to be sent, or an internal event of sending a Logout response (success) on another connection for a “close the session” Logout command was received. In both these cases, the transport connection is closed.
- T14:
 - initiator: Async PDU with iSCSI event “Request Logout” was received again.
 - target: illegal.
- T15, T16:
 - initiator: One or more of the following events caused this transition –
 - internal event indicating a transport connection timeout was received prompting transport RESET or transport connection closure,
 - a transport RESET,
 - a transport disconnect indication.
 - Async PDU with iSCSI event “Drop connection” (for this CID),
 - Async PDU with iSCSI event “Drop all connections”.

State transition descriptions for the standard connection state diagram (contd.)

- target: One or more of the following events caused this transition –
 - internal event indicating a transport connection timeout was received prompting transport RESET or transport connection closure,
 - a transport RESET,
 - a transport disconnect indication,
 - internal emergency cleanup event was received prompting an Async PDU with iSCSI event “Drop connection” (for this CID), or event “Drop all connections”.


- T17:
 - initiator: One or more of the following events caused this transition –
 - Logout response (failure, i.e. a non-zero status) was received,
 - any of the events specified for T15 and T16.
 - target: One or more of the following events caused this transition –
 - internal event indicating a failure of the Logout processing was received, prompting a Logout response (failure, i.e. a non-zero status) to be sent,
 - any of the events specified for T15 and T16.

iSCSI connection cleanup state transitions

Whenever a connection state machine (say, CSM-C) enters the CLEANUP_WAIT state (S8), it must go through the state transitions additionally described in the connection cleanup state diagram either a) using a separate full-feature phase connection (let us call it CSM-E) in the LOGGED_IN state in the same session, or b) using a new transport connection (let us call it CSM-I) in the FREE state that is to be added to the same session. In the CSM-E case, an explicit logout for the CID corresponding to CSM-C (either as a connection or session logout) needs to be performed to complete the cleanup. In the CSM-I case, an implicit logout for the CID corresponding to CSM-C needs to be performed by way of connection reinstatement (Refer the Login PDU section) for that CID. In either case, the protocol exchanges on CSM-E or CSM-I determine the state transitions for CSM-C. This cleanup state diagram hence is applicable only to the instance of the connection in cleanup, i.e. CSM-C. In the case of an implicit logout for example, CSM-C reaches FREE (R3) at the time CSM-I reaches LOGGED_IN. In the case of an explicit logout, CSM-C reaches FREE (R3) when CSM-E receives a successful logout response while continuing to be in the LOGGED_IN state.

Symbolic names & state descriptions for the connection cleanup state diagram :

- R1: CLEANUP_WAIT (Same as S8)
 - initiator: waiting for the internal event to initiate the cleanup processing for CSM-C.
 - target: waiting for the cleanup process to start for CSM-C.
- R2: IN_CLEANUP
 - initiator: waiting for the connection cleanup process to conclude for CSM-C.
 - target: waiting for the connection cleanup process to conclude for CSM-C.
- R3:FREE (Same as S1)
 - initiator: end state for CSM-C.
 - target: end state for CSM-C.



connection cleanup state diagram
for both
initiators and targets

iscsi state diagrams

State transition descriptions for the connection cleanup state diagram


- M1: One or more of the following events was received -
 - initiator:
 - an internal event indicating connection state timeout,
 - a successful Logout response was received on a different connection for a “close the session” Logout.
 - target:
 - an internal event indicating connection state timeout,
 - an internal event of sending a Logout response (success) on a different connection for a “close the session” Logout command.
- M2: An implicit /explicit logout process was initiated by the initiator.
 - In CSM-I usage:
 - initiator: An internal event requesting the CID reinstatement was received prompting a connection reinstatement Login to be sent transitioning to state IN_LOGIN.
 - target: A connection reinstatement Login was received while in state XPT_UP.
 - In CSM-E usage:
 - initiator: An internal event indicating that an explicit logout was sent for this CID in state LOGGED_IN.
 - target: An explicit logout was received for this CID in state LOGGED_IN.
- M3: Logout failure detected
 - In CSM-I usage:
 - initiator: CSM-I failed to reach LOGGED_IN, instead arrived into FREE.
 - target: CSM-I failed to reach LOGGED_IN, instead arrived into FREE.
 - In CSM-E usage:
 - initiator: CSM-E either moved out of LOGGED_IN, or Logout timed out and/or aborted, or Logout response (failure) was received.
 - target: CSM-E either moved out of LOGGED_IN, or Logout timed out and/or aborted, or an internal event indicating a failed Logout processing was received. A Logout response (failure) was sent in the last case.

iscsi state diagrams

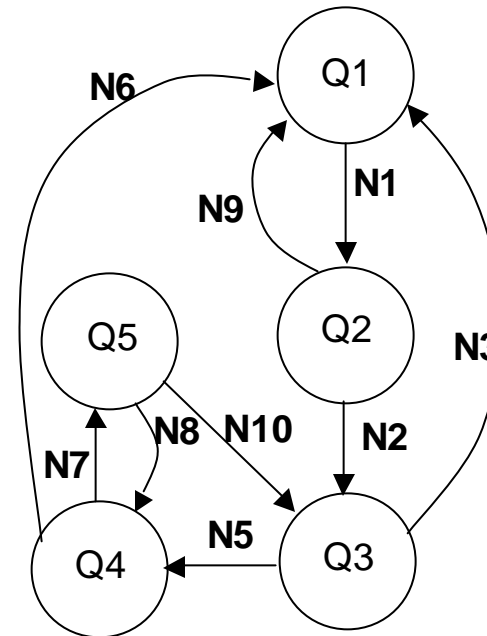
State transition descriptions for the connection cleanup state diagram (contd.)

- M4: Successful implicit/explicit logout was performed.
 - In CSM-I usage:
 - initiator: CSM-I reached state LOGGED_IN.
 - target: CSM-I reached state LOGGED_IN.
 - In CSM-E usage:
 - initiator: CSM-E stayed in LOGGED_IN, and received a Logout response (success).
 - target: CSM-E stayed in LOGGED_IN, and an internal event indicating a successful Logout processing was received.

iSCSI session state transitions



session state diagram for an initiator



session state diagram for a target

iSCSI session state transitions

Session continuation is the process by which the state of a pre-existing session is continued to be in use by either connection reinstatement (section 3. 12.1 of rev08 draft), or by adding a connection with a new CID. Either of these actions associates the new transport connection with the pre-existing session state.

Symbolic Names & state descriptions for the session state diagram

•Q1: FREE

- initiator: state on instantiation or after cleanup.
- target: state on instantiation or after cleanup.

•Q2: ACTIVE

- initiator: illegal
- target: The first iSCSI connection in the session transitioned to IN_LOGIN, waiting for it to complete the login process.

•Q3: LOGGED_IN

- initiator: waiting for all session events.
- target: waiting for all session events.

•Q4: FAILED

- initiator: waiting for session recovery or session continuation.
- target: waiting for session recovery or session continuation.

Q5: IN_CONTINUE

- initiator: illegal
- target: waiting for session continuation attempt to reach a conclusion.

iSCSI session state transitions

State transition descriptions for the session state diagram

•N1:

- initiator: At least one iSCSI connection reached the LOGGED_IN state.
- target: The first iSCSI connection in the session had reached the IN_LOGIN state.

•N2:

- initiator: illegal
- target: At least one iSCSI connection reached the LOGGED_IN state.

•N3:

- initiator: Graceful closing of the session was completed - either via a “close the session” Logout, or last successful “close the connection” Logout.
- target: Graceful closing of the session was completed - either via a “close the session” Logout, or last successful “close the connection” Logout.

•N4:

- initiator: A session continuation attempt has succeeded.
- target: illegal

•N5:

- initiator: Session failure occurred (all connections reported CLEANUP_WAIT).
- target: Session failure occurred (all connections reported CLEANUP_WAIT).

•N6:

- initiator: Session state timeout had happened, or an implicit session logout by reuse of ISID with TSID=0 cleared this session instance. This results in all associated resources to be freed and the session state discarded.
- target: Session state timeout had happened, or an implicit session logout by reuse of ISID with TSID=0 cleared this session instance. This results in all associated resources to be freed and the session state discarded.

iSCSI session state transitions

State transition descriptions for the session state diagram (contd.)

•N7:

- initiator: illegal
- target: A session continuation attempt is initiated.

•N8:

- initiator: illegal
- target: A session continuation attempt failed.

•N9:

- initiator: illegal
- target: Login attempt on the leading connection failed.

•N10:

- initiator: illegal
- target: A session continuation attempt succeeded.