



**LEGAL NOTICE:**

© **Copyright 2008 to 2023 NVM Express®, Inc. ALL RIGHTS RESERVED.**

This Technical Proposal is proprietary to the NVM Express, Inc. (also referred to as “Company”) and/or its successors and assigns.

**NOTICE TO USERS WHO ARE NVM EXPRESS, INC. MEMBERS:** Members of NVM Express, Inc. have the right to use and implement this Technical Proposal subject, however, to the Member’s continued compliance with the Company’s Intellectual Property Policy and Bylaws and the Member’s Participation Agreement.

**NOTICE TO NON-MEMBERS OF NVM EXPRESS, INC.:** If you are not a Member of NVM Express, Inc. and you have obtained a copy of this document, you only have a right to review this document or make reference to or cite this document. Any such references or citations to this document must acknowledge NVM Express, Inc. copyright ownership of this document. The proper copyright citation or reference is as follows: “© 2008 to 2023 NVM Express, Inc. ALL RIGHTS RESERVED.” When making any such citations or references to this document you are not permitted to revise, alter, modify, make any derivatives of, or otherwise amend the referenced portion of this document in any way without the prior express written permission of NVM Express, Inc. Nothing contained in this document shall be deemed as granting you any kind of license to implement or use this document or the specification described therein, or any of its contents, either expressly or impliedly, or to any intellectual property owned or controlled by NVM Express, Inc., including, without limitation, any trademarks of NVM Express, Inc.

**LEGAL DISCLAIMER:**

THIS DOCUMENT AND THE INFORMATION CONTAINED HEREIN IS PROVIDED ON AN “**AS IS**” BASIS. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, NVM EXPRESS, INC. (ALONG WITH THE CONTRIBUTORS TO THIS DOCUMENT) HEREBY DISCLAIM ALL REPRESENTATIONS, WARRANTIES AND/OR COVENANTS, EITHER EXPRESS OR IMPLIED, STATUTORY OR AT COMMON LAW, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE, VALIDITY, AND/OR NONINFRINGEMENT.

All product names, trademarks, registered trademarks, and/or servicemarks may be claimed as the property of their respective owners.

The NVM Express® design mark is a registered trademark of NVM Express, Inc.

NVM Express Workgroup  
c/o VTM, Inc.  
3855 SW 153<sup>rd</sup> Drive  
Beaverton, OR 97003  
USA  
info@nvmexpress.org

## NVM Express® Technical Proposal (TP)

Technical Proposal ID	8025 NVMe-oF Security Configurations
Revision Date	2023.09.18
Builds on Specification(s)	NVM Express Base Specification 2.0c NVM Express TCP Transport Specification 1.0c
References	8019 – Authentication Verification Entity for DH-HMAC-CHAP 8018 – NVMe/TCP – TLS updates

### Technical Proposal Author(s)

Name	Company
Claudio DeSanti, David Black	Dell Technologies

### Technical Proposal Overview

Definition of a consistent way to express security protocol configurations of NVMe entities.

## Revision History

Revision Date	Change Description
2023.02.07	Initial draft
2023.02.14	Added priority requirement for TLS with configured PSK. Named the bits in TREQ and used the new names. Used the ASCR and ATR names for the AUTHREQ bits. Specified how to set the ASCR and ATR bits.
2023.02.21	Updated terminology in 8.13.4.1 and 8.13.1
2023.04.04	Synchronized with TP 8018
2023.04.11	Discussion on Discovery Log Page TSC bits
2023.04.18	Consider both TSC and SECTYPE in table NEW.2
2023.04.25	Ready for phase 2 exit
2023.05.16	Phase 3 clean up
2023.05.23	Clarified the scope of the term 'configuration' and added some examples
2023.05.25	Approved for 30-day member review
2023.07.11	Incorporated Mike Allison's editorial comments
2023.07.13	Integration ready
2023.09.18	Integrated

## Description for Changes Document for TP 8025

New Features/Feature Enhancements/Required Changes:

- Definition of a consistent way to express security protocol configurations of NVMe entities

### **Markup Conventions:**

Black:	Unchanged (however, hot links are removed)
<del>Red Strikethrough:</del>	Deleted
Blue:	New
Blue Highlighted:	TBD values, anchors, and links to be inserted in new text.
Orange Text:	Text from another TP or ECN
<Green Bracketed>:	Notes to editor

# Description of Specification Changes for NVM Express TCP Transport Specification 1.0c

**Update section 3.6.1.6 as shown below:**

## 3.6.1.6 TLS Implementations and Use Requirements

NVMe/TCP host and subsystem implementations shall not send or use 0-RTT data as it is subject to replay attacks (refer to Appendix E.5 of RFC 8446).

~~All NVMe/TCP host and subsystem implementations shall be configurable to require that all NVMe/TCP connections use TLS. If a host that supports TLS for NVMe/TCP receives a discovery log entry indicating that the NVM subsystem uses NVMe/TCP and does not support TLS, then the host should nonetheless attempt to establish an NVMe/TCP connection that uses TLS. This requirement applies independent of whether the host is configured to require use of TLS for all NVMe/TCP connections.~~

NVMe/TCP implementations that support TLS shall support disabling the following parameters, using a method outside the scope of this specification:

- each individual cipher suite;
- PSK-only authentication;
- each individual DH group;
- each individual ECDH group.

**Add section 3.6.1.NEW as shown below:**

## 3.6.1.NEW TLS Configuration for TCP Connection Initiation

NVMe/TCP implementations that support TLS shall be able to be provisioned with a configured PSK for each remote entity with which a TLS connection may be established.

This section uses the term configuration to indicate an NVMe/TCP internal state that controls the use of TLS. That internal state may or may not be externally configurable for a host or NVM subsystem. For an NVMe/TCP implementation that supports TLS, a TLS configuration may apply to:

- a single connection;
- a group of connections; or
- all connections.

NVMe/TCP implementations that support TLS shall support configuring use of TLS with TCP connection initiation using one or more of the configurations shown in **Figure NEW.1**.

**Figure NEW.1: TLS Configurations**

Configuration	Description
TLS Disabled	Only TCP connections without TLS with a remote entity are allowed
TLS Permitted	TCP connections with and without TLS with a remote entity are allowed
TLS Required	Only TCP connections with TLS with a remote entity are allowed

NVMe/TCP hosts that support TLS shall behave as shown in **Figure NEW.2** when establishing NVMe/TCP connections with an NVM subsystem.

**Figure NEW.2: Host TLS Behavior**

Host Configuration	Action
TLS Disabled	Do not initiate TCP connections with TLS.
TLS Permitted	<p>If the SECTYPE field in the TSAS field in the Discovery Log Page Entry for the remote entity is not cleared to 0h, then initiate TCP connections with TLS, irrespective of the value of the TSC field in that Discovery Log Page Entry. If establishing any TCP connection with TLS fails and the TSC field in that Discovery Log Page Entry is not set to 01b (i.e., Required), the host may fall back to initiate TCP connections without TLS.</p> <p>If the SECTYPE field in the TSAS field in the Discovery Log Page Entry for the remote entity is cleared to 0h and the TSC field is not set to 01b (i.e., Required), then initiate TCP connections without TLS. If the SECTYPE field in the TSAS field in the Discovery Log Page Entry for the remote entity is cleared to 0h and the TSC field is set to 01b (i.e., Required), then that Discovery Log Page Entry is inconsistent and TCP connections without TLS may or may not be initiated.</p> <p>If no Discovery Log Page Entry has been retrieved for the remote entity, then TCP connections with or without TLS may be initiated.</p>
TLS Required	Initiate TCP connections with TLS.

NVM subsystems that support TLS with NVMe/TCP shall behave as defined in [Figure NEW.3](#) when establishing NVMe/TCP connections with a host.

**Figure NEW.3: NVM Subsystem TLS Behavior**

Subsystem Configuration	Action
TLS Disabled	Close the TCP connection if a TLS handshake is initiated upon completion of the TCP handshake
TLS Permitted	Continue all TCP connections whether or not a TLS handshake is initiated upon completion of the TCP handshake
TLS Required	Close the TCP connection if a TLS handshake is not initiated upon completion of the TCP handshake

For example, consider a host that supports TLS configured with TLS Disabled and an NVM subsystem that supports TLS configured with TLS Permitted. As defined in [Figure NEW.2](#), the host initiates a TCP connection without TLS. As defined in [Figure NEW.3](#), the subsystem accepts the TCP connection.

## Description of Specification Changes for NVM Express Base Specification 2.0c

Update Figure 264 – Discovery Log Page Entry Data Structure as shown below:

Figure 264 - Discovery Log Page Entry Data Structure

Bytes	Description																			
03	<b>Transport Requirements (TREQ):</b> Indicates requirements for the NVMe Transport.																			
	<table><tr><th>Bits</th><th>Definition</th></tr><tr><td>7:36</td><td>Reserved</td></tr></table>	Bits	Definition	7:36	Reserved															
	Bits	Definition																		
	7:36	Reserved																		
	<table><tr><td rowspan="5">5:4</td><td><b>Transport Authentication and Secure Channel (TASC):</b> Indicates whether an authentication transaction (refer to section 8.13.2) or an authentication transaction concatenated to a secure channel establishment (refer to section 8.13.3) is required upon completion of the Connect command. These bits do not override the AUTHREQ bits in the Connect response; these bits apply only when the AUTHREQ bits specify no requirements (refer to figure 383).</td></tr><tr><td><table><tr><th>Value</th><th>Definition</th></tr><tr><td>00b</td><td>Not specified</td></tr><tr><td>01b</td><td>Authentication required</td></tr><tr><td>10b</td><td>Authentication concatenated to secure channel establishment required</td></tr><tr><td>11b</td><td>Reserved</td></tr></table></td></tr><tr><td></td><td></td></tr><tr><td></td><td></td></tr><tr><td></td><td></td></tr></table>	5:4	<b>Transport Authentication and Secure Channel (TASC):</b> Indicates whether an authentication transaction (refer to section 8.13.2) or an authentication transaction concatenated to a secure channel establishment (refer to section 8.13.3) is required upon completion of the Connect command. These bits do not override the AUTHREQ bits in the Connect response; these bits apply only when the AUTHREQ bits specify no requirements (refer to figure 383).	<table><tr><th>Value</th><th>Definition</th></tr><tr><td>00b</td><td>Not specified</td></tr><tr><td>01b</td><td>Authentication required</td></tr><tr><td>10b</td><td>Authentication concatenated to secure channel establishment required</td></tr><tr><td>11b</td><td>Reserved</td></tr></table>	Value	Definition	00b	Not specified	01b	Authentication required	10b	Authentication concatenated to secure channel establishment required	11b	Reserved						
	5:4		<b>Transport Authentication and Secure Channel (TASC):</b> Indicates whether an authentication transaction (refer to section 8.13.2) or an authentication transaction concatenated to a secure channel establishment (refer to section 8.13.3) is required upon completion of the Connect command. These bits do not override the AUTHREQ bits in the Connect response; these bits apply only when the AUTHREQ bits specify no requirements (refer to figure 383).																	
			<table><tr><th>Value</th><th>Definition</th></tr><tr><td>00b</td><td>Not specified</td></tr><tr><td>01b</td><td>Authentication required</td></tr><tr><td>10b</td><td>Authentication concatenated to secure channel establishment required</td></tr><tr><td>11b</td><td>Reserved</td></tr></table>	Value	Definition	00b	Not specified	01b	Authentication required	10b	Authentication concatenated to secure channel establishment required	11b	Reserved							
			Value	Definition																
00b			Not specified																	
01b		Authentication required																		
10b	Authentication concatenated to secure channel establishment required																			
11b	Reserved																			
3	<b>Zero Host ID Support (ZHIDS):</b> If set to '1' indicates that the controller supports a Host Identifier value of 0h in a Connect command. If cleared to '0' indicates that the controller does not support a Host Identifier value of 0h in a Connect command.																			
2	<b>SQ Flow Control Disable (SQFCD):</b> If set to '1' indicates that the controller is capable of disabling SQ flow control. A controller that is capable of disabling SQ flow control may accept or reject a host request to disable SQ flow control. If cleared to '0', then the controller requires use of SQ flow control.																			
1:0	<b>Transport Secure Channel (TSC):</b> Indicates whether connections shall be made over a fabric secure channel (which includes authentication) (refer to section 8.13.1). <table><tr><th>Value</th><th>Definition</th></tr><tr><td>00b</td><td>Not specified</td></tr><tr><td>01b</td><td>Required</td></tr><tr><td>10b</td><td>Not required</td></tr><tr><td>11b</td><td>Reserved</td></tr></table>	Value	Definition	00b	Not specified	01b	Required	10b	Not required	11b	Reserved									
Value	Definition																			
00b	Not specified																			
01b	Required																			
10b	Not required																			
11b	Reserved																			
0	Obsolete																			

Update section 8.13.1 as shown below:

### 8.13.1 Fabric Secure Channel

...

An NVM subsystem that requires use of a fabric secure channel (i.e., as indicated by the TREQ TSC field in the associated Discovery Log Page Entry) shall not allow capsules to be transferred until a secure channel has been established for the NVMe Transport connection.

...

**Update section 8.13.4.1 as shown below:**

## 8.13.4 Common Authentication Messages

### 8.13.4.1 AUTH\_Negotiate Message

...

Upon receiving an AUTH\_Negotiate message, if the SC\_C value indicated by the host does not satisfy the security requirements of the controller (e.g., the host did not request secure channel concatenation, but the controller's security ~~policy~~ **configuration** requires secure channel concatenation), then the controller shall:

- reply to the AUTH\_Negotiate message with an AUTH\_Failure1 message having reason code 'Authentication failure' and reason code explanation 'Secure channel concatenation mismatch'; and
- disconnect the NVMe over Fabrics connection upon transmitting the AUTH\_Failure1 message.

...

**Add section 8.13.5.NEW as shown below:**

## 8.13.5 DH-HMAC-CHAP Protocol

...

### 8.13.5.NEW DH-HMAC-CHAP Configuration

DH-HMAC-CHAP implementations that do not use an AVE (refer to section **8.13.TBD** <in TP 8019>) shall be able to be provisioned with their own DH-HMAC-CHAP secret and with a verification secret per each remote entity that is able to be authenticated. DH-HMAC-CHAP implementations that use an AVE shall be able to be provisioned with their own DH-HMAC-CHAP secret and the parameters for accessing the AVE (refer to section **8.13.TBD.2** <in TP 8019>).

This section uses the term configuration to indicate an NVMe internal state that controls the use of DH-HMAC-CHAP. That internal state may or may not be externally configurable for a host or NVM subsystem. For an NVMe/TCP implementation that supports DH-HMAC-CHAP, a DH-HMAC-CHAP configuration may apply to:

- a single connection;
- a group of connections; or
- all connections.

DH-HMAC-CHAP implementations that do not support secure channel concatenation with the TLS protocol shall support configuring use of DH-HMAC-CHAP using one or more of the configurations shown in **Figure NEW.4**.

**Figure NEW.4: DH-HMAC-CHAP Configurations**

Configuration	Description
Authentication Disabled	Authentication of a remote entity not allowed
Authentication Permitted	Authentication of a remote entity allowed
Authentication Required	Authentication of a remote entity required

NVM subsystems that do not support secure channel concatenation with the TLS protocol shall set the ATR and ASCR bits in the AUTHREQ field in the response of a successful Connect command (refer to **Figure 383**) as defined in **Figure NEW.5**.



**Figure NEW.5: NVM Subsystem AUTHREQ Settings for DH-HMAC-CHAP**

Subsystem Configuration	ASCR	ATR
Authentication Disabled	0	0
Authentication Permitted	0	0
Authentication Required	0	1

Hosts that do not support secure channel concatenation with the TLS protocol shall behave as defined in Figure NEW.6, according to the ATR and ASCR bits in the AUTHREQ field received in the response of a successful Connect command (refer to Figure 383).

**Figure NEW.6: DH-HMAC-CHAP Host Behavior**

Host Configuration	ASCR	ATR	Action
Authentication Disabled	0	0	Do not begin an authentication transaction
	0	1	Disconnect from the NVM subsystem
	1	any	
Authentication Permitted	0	0	If the TASC field in the Discovery Log Page Entry for the remote entity is set to 01b or set to 10b, then begin an authentication transaction with the SC_C field set to NOSC.
			If the TASC field in the Discovery Log Page Entry for the remote entity is cleared to 00b or set to 11b or if no Discovery Log Page Entry for the remote entity has been retrieved, then do not begin an authentication transaction.
	0	1	Begin an authentication transaction with the SC_C field set to NOSC
	1	any	Disconnect from the NVM subsystem
Authentication Required	0	any	Begin an authentication transaction with the SC_C field set to NOSC
	1	any	Disconnect from the NVM subsystem

NVM subsystems that do not support secure channel concatenation with the TLS protocol shall behave as defined in [Figure NEW.7](#), according to the SC\_C field in a received AUTH\_Negotiate message (refer to section [8.13.4.1](#)).

**Figure NEW.7: DH-HMAC-CHAP NVM Subsystem Behavior**

Subsystem Configuration	SC_C Value	Action
Authentication Disabled	any	Disconnect from the host
Authentication Permitted	NOSC	Participate in the authentication transaction
	NEWTLSPPSK or REPLACETLSPPSK	Disconnect from the host
Authentication Required	NOSC	Participate in the authentication transaction
	NEWTLSPPSK or REPLACETLSPPSK	Disconnect from the host

For example, consider a host that supports DH-HMAC-CHAP without secure channel concatenation with the TLS protocol configured with Authentication Permitted and an NVM subsystem that supports DH-HMAC-CHAP without secure channel concatenation with the TLS protocol configured with Authentication Required. As defined in [Figure NEW.5](#), the NVM subsystem clears the ASCR bit to '0' and sets the ATR bit to '1' in the Connect response. As defined in [Figure NEW.6](#), upon receiving the Connect response the host begins an authentication transaction with the SC\_C field set to NOSC. As defined in [Figure NEW.7](#), the NVM subsystem participates in the authentication transaction.

DH-HMAC-CHAP implementations that support secure channel concatenation with the TLS protocol shall support configuring use of DH-HMAC-CHAP with secure channel concatenation with the TLS protocol using one or more of the configurations shown in [Figure NEW.8](#).

**Figure NEW.8: DH-HMAC-CHAP with TLS Concatenation Configurations**

Configuration		Description
Authentication Disabled	TLS Disabled	Authentication and set up of a secure channel with a remote entity not allowed
Authentication Permitted	TLS Disabled	Authentication of a remote entity allowed, set up of a secure channel not allowed
	TLS Permitted	Authentication and set up of a secure channel with a remote entity allowed
Authentication Required	TLS Disabled	Authentication of a remote entity required, set up of a secure channel not allowed
	TLS Permitted	Authentication of a remote entity required, set up of a secure channel allowed
	TLS Required	Authentication and set up of a secure channel with a remote entity required

NVM subsystems that support secure channel concatenation with the TLS protocol shall set the ATR and ASCR bits in the AUTHREQ field in the response of a successful Connect command on an Admin Queue over a TCP channel without TLS (refer to [Figure 383](#)) as defined in [Figure NEW.9](#).

**Figure NEW.9: NVM Subsystem AUTHREQ Settings for DH-HMAC-CHAP with TLS Concatenation**

Subsystem Configuration		ASCR	ATR
Authentication Disabled	TLS Disabled	0	0
Authentication Permitted	TLS Disabled	0	0
	TLS Permitted	0	0
Authentication Required	TLS Disabled	0	1
	TLS Permitted	0	1
	TLS Required	1	0

Hosts that support secure channel concatenation with the TLS protocol shall behave as defined in [Figure NEW.10](#), according to the ATR and ASCR bits in the AUTHREQ field received in the response of a successful Connect command on an Admin Queue over a TCP channel without TLS (refer to [Figure 383](#)).

**Figure NEW.10: DH-HMAC-CHAP with TLS Concatenation Host Behavior**

Host Configuration		ASCR	ATR	Action
Authentication Disabled	TLS Disabled	0	0	Do not begin an authentication transaction
		0	1	Disconnect from the NVM subsystem
		1	any	
Authentication Permitted	TLS Disabled	0	0	If the TASC field in the Discovery Log Page Entry for the remote entity is set to 01b or set to 10b, then begin an authentication transaction with the SC_C field set to NOSC.
				If the TASC field in the Discovery Log Page Entry for the remote entity is cleared to 00b or set to 11b or if no Discovery Log Page Entry for the remote entity has been retrieved, then do not begin an authentication transaction.
		0	1	Begin an authentication transaction with the SC_C field set to NOSC
		1	any	Disconnect from the NVM subsystem
	TLS Permitted	0	0	If the TASC field in the Discovery Log Page Entry for the remote entity is set to 01b, then begin an authentication transaction with the SC_C field set to NOSC.
				If the TASC field in the Discovery Log Page Entry for the remote entity is set to 10b, then begin an authentication transaction with the SC_C field set to NEWTLSPSK.
		0	1	If the TASC field in the Discovery Log Page Entry for the remote entity is cleared to 00b or set to 11b or if no Discovery Log Page Entry has been retrieved for the remote entity, then do not begin an authentication transaction.
				Begin an authentication transaction with the SC_C field set to NOSC
Authentication Required	TLS Disabled	0	any	Begin an authentication transaction with the SC_C field set to NOSC
		1	any	Disconnect from the NVM subsystem
	TLS Permitted	0	any	Begin an authentication transaction with the SC_C field set to NOSC
		1	any	Begin an authentication transaction with the SC_C field set to NEWTLSPSK
	TLS Required	any	any	Begin an authentication transaction with the SC_C field set to NEWTLSPSK
		any	any	Begin an authentication transaction with the SC_C field set to NEWTLSPSK

NVM subsystems that support secure channel concatenation with the TLS protocol shall behave as defined in [Figure NEW.11](#), according to the SC\_C field in a received AUTH\_Negotiate message on an Admin Queue over a TCP channel without TLS (refer to section [8.13.4.1](#)).

**Figure NEW.11: DH-HMAC-CHAP with TLS Concatenation NVM Subsystem Behavior**

Subsystem Configuration		SC_C Value	Action
Authentication Disabled	TLS Disabled	any	Disconnect from the host
Authentication Permitted	TLS Disabled	NOSC	Participate in the authentication transaction
		NEWTLSPSK or REPLACETLSPSK	Disconnect from the host
	TLS Permitted	NOSC	Participate in the authentication transaction
		NEWTLSPSK	Participate in the authentication transaction and establish a TLS channel as described in section 8.13.3
		REPLACETLSPSK	Disconnect from the host
Authentication Required	TLS Disabled	NOSC	Participate in the authentication transaction
		NEWTLSPSK or REPLACETLSPSK	Disconnect from the host
	TLS Permitted	NOSC	Participate in the authentication transaction
		NEWTLSPSK	Participate in the authentication transaction and establish a TLS channel as described in section 8.13.3
		REPLACETLSPSK	Disconnect from the host
	TLS Required	NOSC or REPLACETLSPSK	Disconnect from the host
		NEWTLSPSK	Participate in the authentication transaction and establish a TLS channel as described in section 8.13.3

As an example, consider a host that supports DH-HMAC-CHAP with secure channel concatenation with the TLS protocol configured with Authentication Permitted, TLS Permitted and an NVM subsystem that supports DH-HMAC-CHAP with secure channel concatenation with the TLS protocol configured with Authentication Required, TLS Required. As defined in Figure NEW.9, the NVM subsystem sets the ASCR bit to '1' and clears the ATR bit to '0' in the Connect response. As defined in Figure NEW.10, upon receiving the Connect response the host begins an authentication transaction with the SC\_C field set to NEWTLSPSK. As defined in Figure NEW.11, the NVM subsystem participates in the authentication transaction and establishes a TLS channel.