




Flash Memory Summit

Handling Common Challenges When Designing with NVMe-over-Fabrics

Nikhil Jain

Member Consulting Staff, Mentor Graphics

Nikhil_jain@mentor.com



Discussion topics



- What is Verification IP?
- NVMe-OF World
- NVMe-OF Hurdles
- Summary



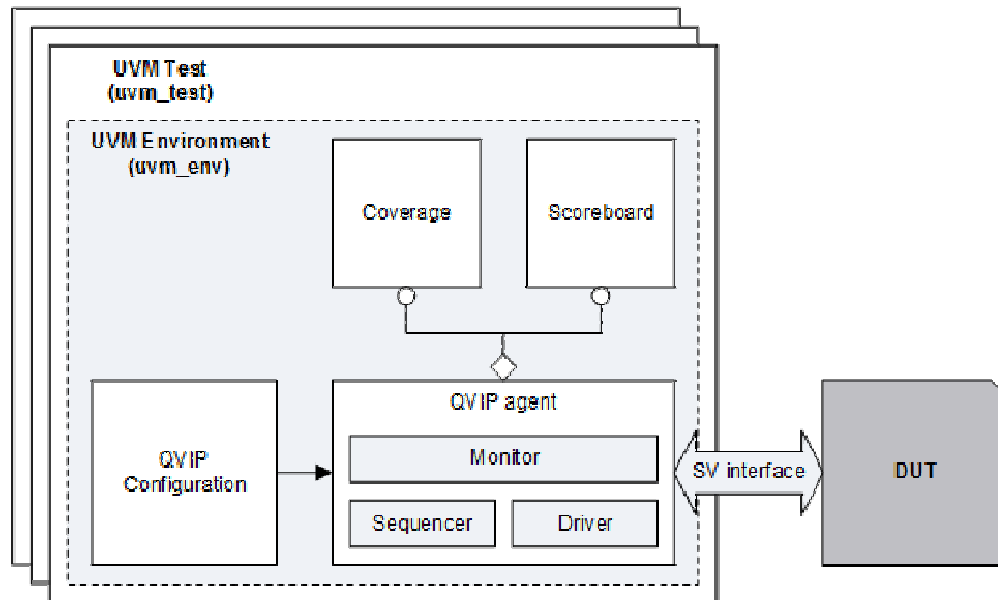
Verification IP



- Verification IP
 - What
 - a library of reusable components
 - Supports industry standards
 - Built-In Protocol Checks and Coverage
 - Why
 - To ensure IP interoperability and system behavior.
 - To improve quality
 - Saves engineering resources to do verification on own
 - Reduces schedule time



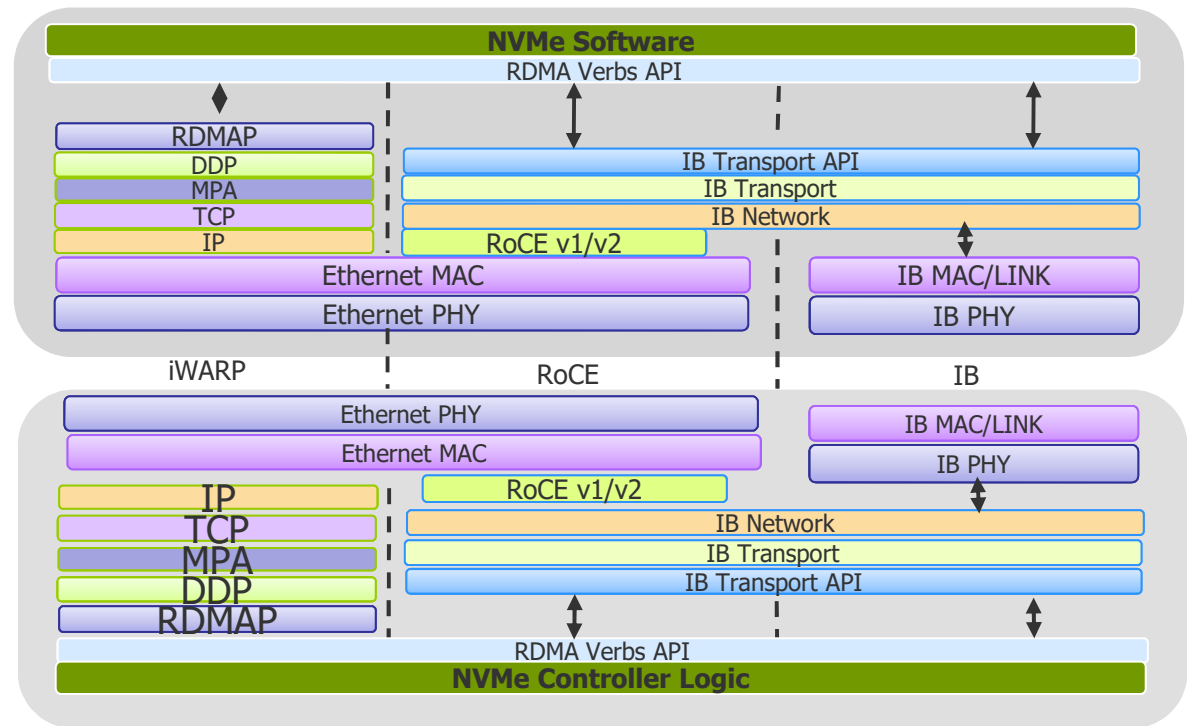
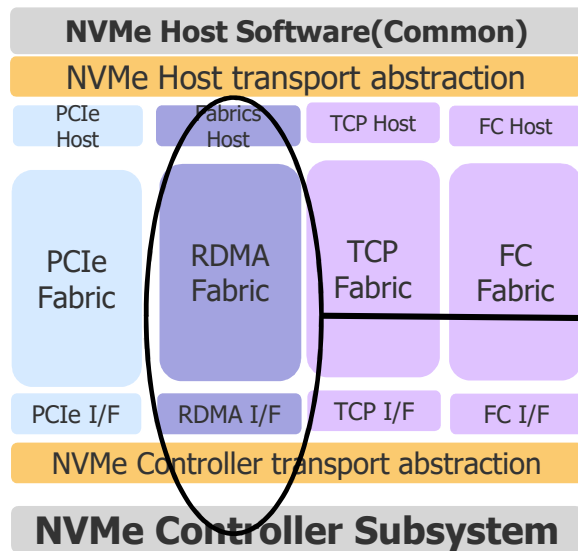
Symbolic VIP environment



- Configuration
 - To Configure host/controller
- Stimulus
 - To execute scenarios
- Assertions
 - To check protocol validity
- Coverage
 - To observe execution of test plan



NVMe-OF World





NVMe-OF hurdles



Standard in its infancy.



Transport Hurdles-1



- Transport Connect for Discovery service at initial stage
 - Hurdle
 - Which transport to use to connect to discovery service?
 - Solution/Clarification
 - Administrator configures every node with transport type
 - Manually configures each host's discovery client with the transport address of the Discovery Service.
 - How VIP helps
 - **Configuration** - Configurable transport type of each port
 - **Stimulus** - generation as per transport requirement



Transport Hurdles-2



- Max RNR NAK handling if requester exhausts
 - Hurdle
 - RNR Retry limit exhausted, how host knows
 - Solution/Clarification
 - Treat this as FATAL
 - start teardown process
 - How VIP helps
 - **Assertion** – for max retry limit and others
 - **Configuration** – FATAL severity can be changed
 - **Coverage** - all possible scenarios covered



Transport Hurdles-3



- Queue deletion at Transport Level
 - Hurdle
 - No Disconnect process definition
 - What if queue deletion process at transport fails?
 - Solution
 - Host will initiate disconnect process through pre defined verbs and allows it to completed before initiating any new process
 - If disconnect fails transport will report the same to local NVMe layer through verbs
 - How VIP helps
 - **Configuration** – Dynamic deletion of Queues and its process
 - **Stimulus** – to start deletion process
 - **Assertion** – to report errors for any violation



Connect Hurdles-4



- CID and SQHD for Connect commands
 - Hurdle
 - SQHD and CID for Connect Command?
 - Solution/Clarification
 - Zero ! Initially for non created queues
 - “Command sequence error” for created queues ,so valid CQID and SQHD must be present
 - How VIP helps
 - **Configuration**
 - Configurable SQID and CID fields for any command
 - Configurable status messages
 - **Stimulus** – to insert any user defined values in command fields
 - **Assertion** – to report error for violation of any field values



Connect Hurdles-5




- Connect Commands:- Status “invalid Field in command”
 - Hurdle
 - What status for invalid fields in connect command?
 - Solution/Clarification
 - “Connect Invalid Parameter”
 - How VIP helps
 - **Configuration** - Configurable status messages
 - **Stimulus** – to insert any user defined values in command fields
 - **Assertion** – to report error for violation of any field values

Protocol Hurdles-6

- Mix SGLs type in Single Capsule

- Hurdle
 - Assorted SGLs in single command capsule?



- Solution/Clarification
 - Assorted Data Segments Not Allowed

- How VIP helps

- Stimulus** – to execute commands contain assorted SGL types
- Assertion** – to report error for violation such behavior
- Coverage** – to cover all SGL types



Miscellaneous Hurdles-7



- Discovery Log Entry
 - Hurdle
 - Assorted Log Entries?
- PSDT
 - Hurdle
 - Reserved in NVMe-OF means 2'b00 or 2'b10?
- Command Sequence Error for create IO Queue
 - Hurdle
 - Can host send IO queue connect before enabling controller
- TRADDR and IP Address mapping
 - Hurdle
 - TRADDR – Network Address DGID etc. mapping not defined



Summary



- The transport hurdles can be avoided by proper communication channel
- connect commands not interpreted by Controller i.e. fabric specific
- NVMe-Of standard will get hand-2-hand with NVMe standard
- Above all can be easily resolved by using QVIP in any defined/generic behavior



Flash Memory Summit

Mentor QVIP portfolio



Questa VIP Protocols

AMBA® Family	PCIe® Family	Ethernet Family	USB Family	MIPI® Family	Serial Family	Display Family	Automotive Family	Mil-Aero Family
ACE	NVMe	200/400G	USB 3.1	I3C	JTAG	HDMI 2.1	CAN	Spacewire
AXI4	AHCI	100G	USB 3.0	UFS	SmartCard	HDMI 1.4	CAN-FD	1553b
AXI3	PCIe 4.0	25/50G	USB 2.0	Unipro	I2C	DisplayPort	LIN	PCI
AHB5	PCIe 3.1	40G	USB 1.1	LLI	I3C	eDP		
AHB	PCIe 2.1	10G	USB PD	CSI-2 / CSI-3	I2S	V-by-One	5G Family	AMBA 5 CHI
APB3	PCIe 1.1	2.5/5G	oHCI	DSI	SPI	CEC	CPRI	CHI 5
		1G	xHCI	DigRF	SPI 4.2	HDCP	JESD204B	
		100M	eHCI	HSI	UART			NVMe over Fabric
		10M		MPHY				
		Automotive		CPHY			Storage Family	
		Interlaken		DPHY			SATA	

Questa Memory Library

DRAM Family	Flash Family
LPDDR4	SDCard 4.2
LPDDR3	SDIO 4.1
LPDDR2	eMMC 5.1
DDR4	ONFI 4.0
DDR3	Toggle
DDR2	UFS
HBM2	Parallel NOR
DIMM	Serial NOR



Thank You



- Please visit us at our Booth 136 in the Exhibition area for a demo
 - **Mentor Graphics Corporation**
 - 8005 SW Boeckman Road
Wilsonville, OR 97070
 - Phone: [800-592-2210](tel:800-592-2210) or [503-685-7000](tel:503-685-7000)
 - **Product and Sales**
 - Phone: [800-547-3000](tel:800-547-3000) or [503-685-8000](tel:503-685-8000)
Email: sales_info@mentor.com