

Product Overview Storage Products and new 400G NICs

Dietmar Hinze
FAE

23.10.2024

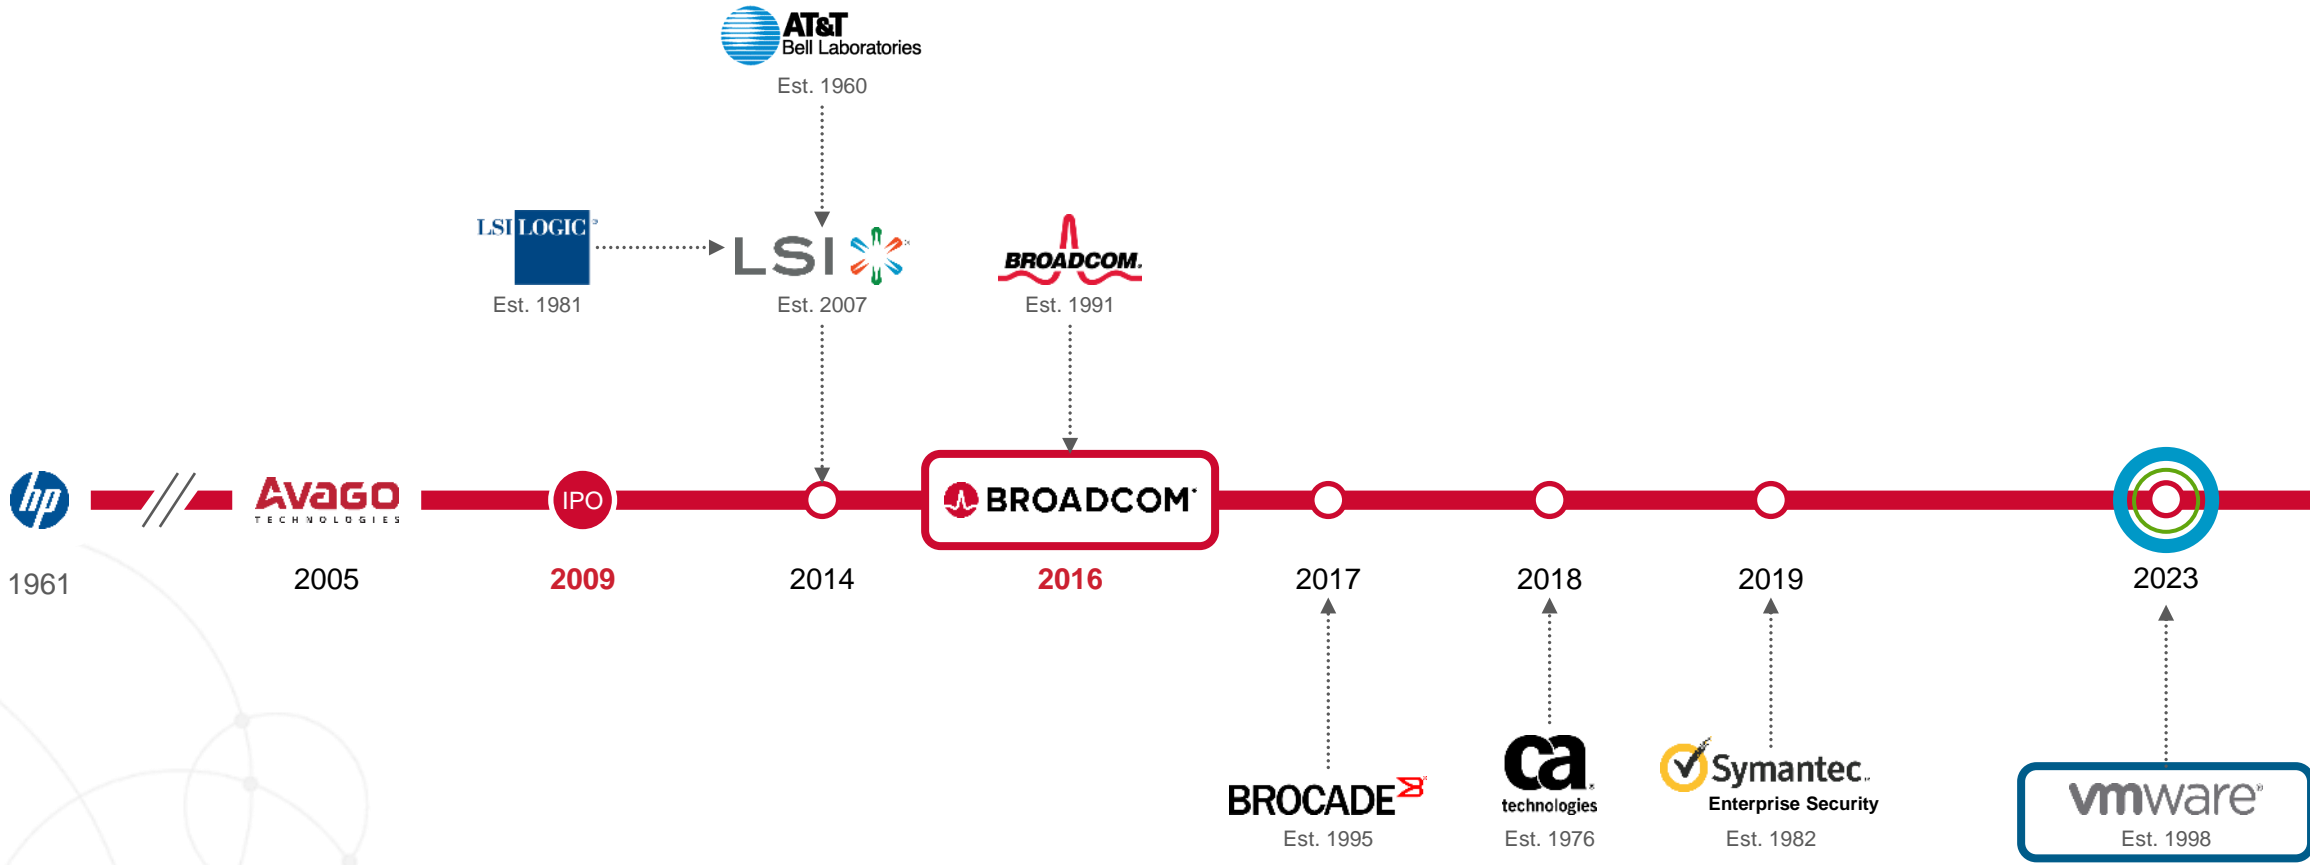
Agenda


- Overview of Storage Products
- Ethernet Technology Trends
- AI/ML NIC Requirements
- Ethernet Products
- 400G Ethernet NICs
- Q & A

Heritage of Innovation

Semiconductors

Infrastructure Software





Overview of Storage Product



Product Brands Storage



**NVMe/SAS/SATA
Host Bus Adapter & RAID Controller**



**Fibre Channel
Host Bus Adapter**



PCIe Ethernet Adapters



Fibre Channel Switches

Storage/Server Connectivity: On-Premises and Cloud

Enterprise



Hyperscale Data Center




Fibre Channel HBAs



Fibre Channel SAN Fabric



RAID Adapters



SSD/HDD Controllers / Preamp

Broadcom Storage Controller Portfolio

Controllers



Adapters/CVPM05



Cables/Connectors



30 years of server leadership, Industry first PCIe4 Tri-mode technology

Reduces latency upto 50%, increases throughput upto 60%, supports NVMe, SAS or SATA storage devices

Broadcom DCSG Gen 4 Tri-Mode IT Adapters

	9500-16i	9500-8i	9500-16e	9500-8e	9600-24i	9600-16i	9600-8i8e	9600-16e	9600W-16e
PCIe Host	x8, Gen 4				x8, Gen 4				x16, Gen 4
SAS	12Gb/s				24G				
PCIe/NVMe	Gen 4				Gen 4				
IOC	SAS3816 (Aero IOC)	SAS3808 (Aero IOC)	SAS3816 (Aero IOC)	SAS3808 (Aero IOC)	SAS4024 (Avenger IOC)	SAS4016 (Avenger IOC)	SAS4016 (Avenger IOC)	SAS4016 (Avenger IOC)	SAS4016W (Avenger IOC)
Form Factor	LP	LP	LP	LP	LP	LP	LP	LP	LP
IOPS (4K RR)	3.0M	3.0M	3.0M	3.0M	3.4M	3.4M	3.4M	3.4M	6.4M
Bandwidth (256K Seq Rd)	13,700 MBs	13,700 MBs	13,700 MBs	13,700 MBs	13,700 MBs	13,700 MBs	13,700 MBs	13,700 MBs	27,900 MBs
Power (Typical)	9W	6W	9W	6W	17W	15W	15W	15W	18W
Thermal Requirements	150 LFM @ 55C	150 LFM @ 55C	150 LFM @ 55C	150 LFM @ 55C	150 LFM @ 55C	150 LFM @ 55C	200 LFM @ 55C	200 LFM @ 55C	250 LFM @ 55C
Security	HW Secure Boot	HW Secure Boot	HW Secure Boot	HW Secure Boot	HW Secure Boot Attestation	HW Secure Boot Attestation	HW Secure Boot Attestation	HW Secure Boot Attestation	HW Secure Boot Attestation
Connectors	2 x8 SFF-8654	2 x8 SFF-8654	4 x4 SFF-8644	2 x4 SFF-8644	3 x8 SFF-8654	2 x8 SFF-8654	1 x8 SFF-8654 2 x4 SFF-8674	4 x4 SFF-8674	4 x4 SFF-8674
FW Stack	SAS 3.5	SAS 3.5	SAS 3.5	SAS 3.5	MR 8 eHBA (Unified Driver)	MR 8 eHBA (Unified Driver)	MR 8 eHBA (Unified Driver)	MR 8 eHBA (Unified Driver)	MR 8 eHBA (Unified Driver)
Availability (GCA)	Production	Production	Production	Production	Production	Production	Production	Production	Production



Broadcom DCSG MegaRAID Adapters

	9540-8i	9540-2M2	9560-8i	9560-16i	9580-8i8e	9620-16i	9660-16i	9670-24i	9670W-16i
PCIe Host	x8, Gen 4					x8, Gen 4			x16, Gen 4
SAS	12Gb/s					24Gb/s			
PCIe/NVMe	Gen 4					Gen 4			
ROC/IOC	SAS3808 (Aero IOC)	SAS3816 (Aero IOC)	SAS3908 (Aero ROC)	SAS3916 (Aero ROC)	SAS3916 (Aero ROC)	SAS4016 (Avenger 1 IOC)	SAS4116 (Avenger 1 ROC)	SAS4124 (Avenger 1 ROC)	SAS4116W (Avenger 1 ROC)
Form Factor	LP	LP	LP	LP	LP	LP	LP	FH	FH
Bandwidth MBs (large seq rd)	13,700	13,700	13,700	13,700	13,700	14,900	14,900	14,900	29,500
IOPS (4K RR)	3.0M	3.0M	3.0M	3.0M	3.0M	3.45M	3.45M	3.45M	7M
RAID 5 IOPS (4K RW)	-	-	240K	240K	240K	N/A	500K	1.1M	1.1M
R5 Wr BW (MBs)	-	-	4,860	4,860	4,860	N/A	6,600	13,100	13,100
R5 WR Latency (us)	-	-	>200	>200	>200	N/A	8	8	8
R5 Rebuild Rate (Min/TB)	-	-	~35 Min/TB	~35 Min/TB	~35 Min/TB	N/A	~32 Min/TB	~16 Min/TB	~16 Min/TB
Perf Under Rebuild (R5) (4K RR IOPs)	-	-	43K	43K	43K	-	2,900K	2,900K	2,900K
Typical Power	6W	6W (without drives)	12W	14W	14W	15W	18W	24W	25W
Thermal Reqt's LFM @55C	150	TBD	200	200	200	150	250	200	200
Security	HW Secure Boot	HW Secure Boot	HW Secure Boot	HW Secure Boot	HW Secure Boot	HW SB Attestation	HW SB Attestation	HW SB Attestation	HW SB Attestation
Connectors	1 x8 SFF-8654	2 M.2 On-board	1 x8 SFF-8654	2 x8 SFF-8654	2 x8 SFF-8654 2 x4 SFF-8674	2 x8 SFF-8654	2 x8 SFF-8654	3 x8 SFF-8654	2 x8 SFF-8654
Energy Backup	-	-	CVPM05	CVPM05	CVPM05	N/A	CVPM05	CVPM05	CVPM05
FW Stack	MR 7.x (iMR)	MR 7.x (iMR)	MR 7.x	MR 7.x	MR 7.x	MR 8.x /EVP (Unified Driver)	MR 8.x (Unified Driver)	MR 8.x (Unified Driver)	MR 8.x (Unified Driver)
Availability	Production	Production	Production	Production	Production	Production	Production	Production	Production



OCP Storage Cards



9562-16i: Aero Based OCP 3.0 SFF RAID Card

- Aero based OCP RAID card concept drawing
- 4G Cache; 16 Ports
- Two low profile vertical SFF-8654 connectors
- CVPM05 (FBU345) for energy backup




	9562-16i
PCIe Host	x8, Gen 4
SAS	12Gb/s
PCIe/NVMe	Gen 4
ROC/IOC	SAS3916 (Aero ROC)
Form Factor	OCP 3.0 SFF
Bandwith MBs (large seq rd)	13,700
IOPS (4K RR)	3.0M
RAID 5 IOPS (4K RW)	240K
R5 Wr BW (MBs)	4,860
R5 WR Latency (us)	>200
R5 Rebuild Rate (Min/TB)	~35 Min/TB
Perf Under Rebuild (4K RR IOPs)	43K
Typical Power*	14W
Thermal Reqt's LFM @55*	250 OCP Tier 5
Security	HW Secure Boot
Connectors	2 x8 SFF-8654
Energy Backup	CVPM05 (FBU345)
FW Stack	MR 7.x

*Projections based on current 9560-16i card

9502-16i: Aero Based OCP 3.0 SFF HBA Card


- Aero based OCP IT RAID card concept drawing
- 16 internal ports
- Two low profile vertical SFF-8654 connectors
- Current schedule assumes single spin
 - EVT to PVT



	9502-16i
PCIe Host	x8, Gen 4
SAS	12Gb/s
PCIe/NVMe	Gen 4
ROC/IOC	SAS3816 (Aero IOC)
Form Factor	OCP 3.0 SFF
Bandwith MBs (large seq rd)	13,700
IOPS (4K RR)	3.0M
Typical Power*	10W
Thermal Req't's * LFM @55	200 OCP Tier 4
Security	HW Secure Boot
Connectors	2 x8 SFF-8654
Energy Backup	-
FW Stack	SAS 3.5

9602W-16e: Avenger Based OCP 16e HBA

- OCP 3.0 SFF external HBA card based on the Avenger 4016W IOC
- Will use a TSFF bracket meaning that unlike the 9502-16e card in previous slide, it will fit in a compliant OCP system with a TSFF opening



	9602W-16e
PCIe Host	x16, Gen4
SAS	24Gb/s
PCIe/NVMe	Gen 4
ROC/IOC	SAS4016W (Avenger IOC)
Form Factor	OCP 3.0 SFF with TSFF bracket
Bandwith MBs (large seq rd)	27,900
IOPS (4K RR)	6.7M
Typical Power	18W
Thermal Req't's LFM @55	TBD
Security	HW Secure Boot Attestation
Connectors	4 x4 SFF-8674
FW Stack	MR 8.x /eHBA (Unified Driver)

**Estimates based on current standard FF Aero based cards.*

MegaRAID Entry Level Products



Broadcom DCSG Entry RAID Adapters


	9540-8i	9540-2M2	9520-8i	9520-2M2	9620-16i
PCIe Host	x8, Gen 4		x8, Gen 4		x8, Gen 4
SAS	12Gb/s		12Gb/s		24G
PCIe/NVMe	Gen 4		Gen 4		Gen 4
ROC/IOC	SAS3808 (Aero IOC)	SAS3808 (Aero IOC)	SAS3808 (Aero IOC)	SAS3808 (Aero IOC)	SAS4016 (Avenger 1 IOC)
Form Factor	LP	LP	LP	LP	LP
Bandwidth MBs (large seq rd)	13,700	13,700	13,700	13,700	14,900
IOPS (4K RR)	3.0M	3.0M	3.0M	3.0M	3.45M
Memory Requirement	Host Memory	Host Memory	OCM	OCM	OCM
RAID Levels	0/1	0/1	0/1	0/1	0/1
Max expanders	4	0	2	0	2
Max PCIe switches	0	0	0	0	0
Max SAS/SATA drives	63 (single path); 32(dual path)	2	16	2	8
Max NVMe drives	4	2	2	2	2
Typical Power	6W	6W (without drives)	6W	6W (without drives)	15W
Thermal Reqt's LFM @55	150	TBD	150	TBD	150
Security	HW Secure Boot	HW Secure Boot	HW Secure Boot	HW Secure Boot	HW SB Attestation
Connectors	1 x8 SFF-8654	2 M.2 On-board	1 x8 SFF-8654	2 M.2 On-board	2 x8 SFF-8654
Energy Backup	-	-	-	-	N/A
FW Stack	MR 7.x (iMR)	MR 7.x (iMR)	MR 7.x (iMR)	MR 7.x (iMR)	MR 8.x /EVP (Unified Driver)
Availability	Production	Production FW: MR 7.24+	July (MR 7.30)	July (MR 7.30)	Production

9520-2M2: Features

- Broadcom Data Protection
- Broadcom Unified Tool Suite
- High Performance
 - Separates OS array from primary data storage
 - Increases overall bandwidth
 - Rebuilds on one does not impact the other



Feature	9520-2M2
PCIe Host	x8 Gen 4
Drive Form Factor	M.2 22110, 2280, 2230
Drive Interface	NVMe x1, x2, x4; SATA
RAID Levels	R0, R1
Power Loss Protection	Yes
Stripe Size	64K
Maximum Number of Drives	2
Maximum Supported Virtual Disk	2
Management	Unified Broadcom Tool Suite
Typical Power (no drives)	7W
Airflow	150LFM @ 55C(without SSDs)
Available	Now



Ethernet Technology Trends



PCI Technology Trends

Interface		PCIe 1.0	PCIe 2.0	PCIe 3.0	PCIe 4.0	PCIe 5.0	PCIe 6.0	PCIe 7.0
Interface		PCIe 1.0	PCIe 2.0	PCIe 3.0	PCIe 4.0	PCIe 5.0	PCIe 6.0	PCIe 7.0
Date of Release		2003	2007	2010	2017	2019	2021	2025(planned)
Transfer rate		2.5 GT/s	5 GT/s	8 GT/s	16 GT/s	32 GT/s	64 GT/s	128 GT/s
Bandwidth	x1	250 MB/s	500 MB/s	0.99 GB/s	1.97 GB/s	3.9 GB/s	7.5 GB/s	15.1 GB/s
	x8	2 GB/s	4 GB/s	7.88 GB/s	15.75 GB/s	31.5 GB/s	60.5 GB/s	121 GB/s
	x16	4 GB/s	8 GB/s	15.75 GB/s	31.5 GB/s	63 GB/s	121 GB/s	242 GB/s
Encoding		NRZ				PAM-4 FEC		
		8b/10b		128b/130b		1b/1b 242B/256B FLIT		
Error Correction		8b/10b		no correction		CRC+FEC		

CRC - Cyclic Redundancy Check

FEC - Forward Error Correction

FLIC - Flow Control Unit (a 256B FLIT carries 242B of data)

Ethernet Technology Trends

Transceiver Technology

Technology	Speed	Features	Modulation
SFP+, SFP28	10G, 25G	SFP28 backwards compatible to SFP+	NRZ
QSFP+, QSFP28	10G, 25G	4x 10G, 4x 25G, QSFP28 backwards compatible to QSFP+	NRZ
SFP56, QSFP56 SFP112, QSFP112	Up to 100G per port Up to 400G	Single, quad port, backward compatible to SFP28/QSFP28	PAM4/NRZ
SFP-DD, QSFP-DD	Up to 100G per port	Dual/Octal(8) port, SFP28/QSFP28 should work Standard for 400G	PAM4/NRZ
DSFP	up to 50G per port 100G??	Dual port, backwards compatible to SFP28	PAM4/NRZ
LPO (Linear Plugable Optics)	Up to 400G	Quad(4) port compatible to QSFP-112	PAM4/NRZ
OSFP	Up to 400G	Octal(8) port not compatible to QSFP-DD, compatible via adapter with QSFP28 Connector doesn't fit on OCP3 nor standard PCIe form factor	PAM4/NRZ
Plugable Interconnects	Up to 400G per port	Internal Fibre to module or chip	PAM4/NRZ

Modulations

Technology	Speed	Features
NRZ	10G, 25G	None-Return-to-Zero
PAM4	50G,100G,200G, 400G	Pulse Amplitude Modulation (PAM4 will also be used in PCIe gen 6) also 56G,112G,224G

Optics and Copper Trends

Technology	Metrics	10G	25G	100G	200G	400G
Optics (Xcvr)	Power	1W	2W	4W	10W	20W
	Cost	\$40	\$80	\$200	\$900	\$1,200
Optics (AOC)	Power	1W	2W	4W	10W	20W
	Cost	\$30	\$50	\$160	\$800	\$1,300
Passive Cu (DAC)	Power	0.2W	0.5W	0.5W	0.5W	0.5W
	Cost	\$15	\$40	\$50	\$130	\$200

High-speed optics → Expensive and high-power

Ethernet Transfer Rates

Transfer rates

Frequency	Transfer Rate	Modulation	# of Ports	Technology
11.625GHz	10G	NRZ	1	SFP+
28GHz	25G	NRZ	1	SFP28
28GHz	50G	PAM4	1	SFP56
56GHz	100G	PAM4	1	SFP112
11.625GHz	40G	NRZ	4	QSFP+
28GHz	100G	NRZ	4	QSFP28
28GHz	200G	PAM4	4	QSFP56
56GHz	400G	PAM4	4	QSFP112
112GHz	800G	PAM4	4	QSFP224
224GHz	1.6T	PAM4	4	QSFP448

57600 – Linear Pluggable Optics (LPO) Support



Lower Cost

Up to 30% less Cost per Transceiver

Lower Power

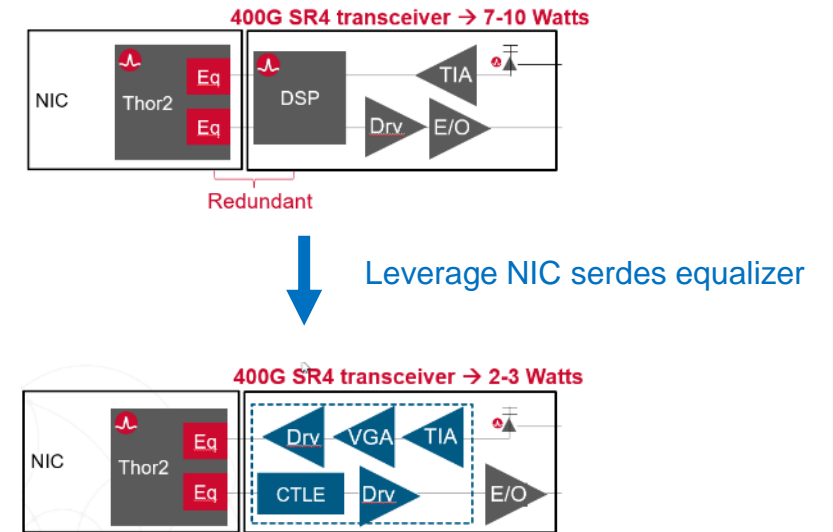
>50% lower Power per Transceiver

Improved Thermals

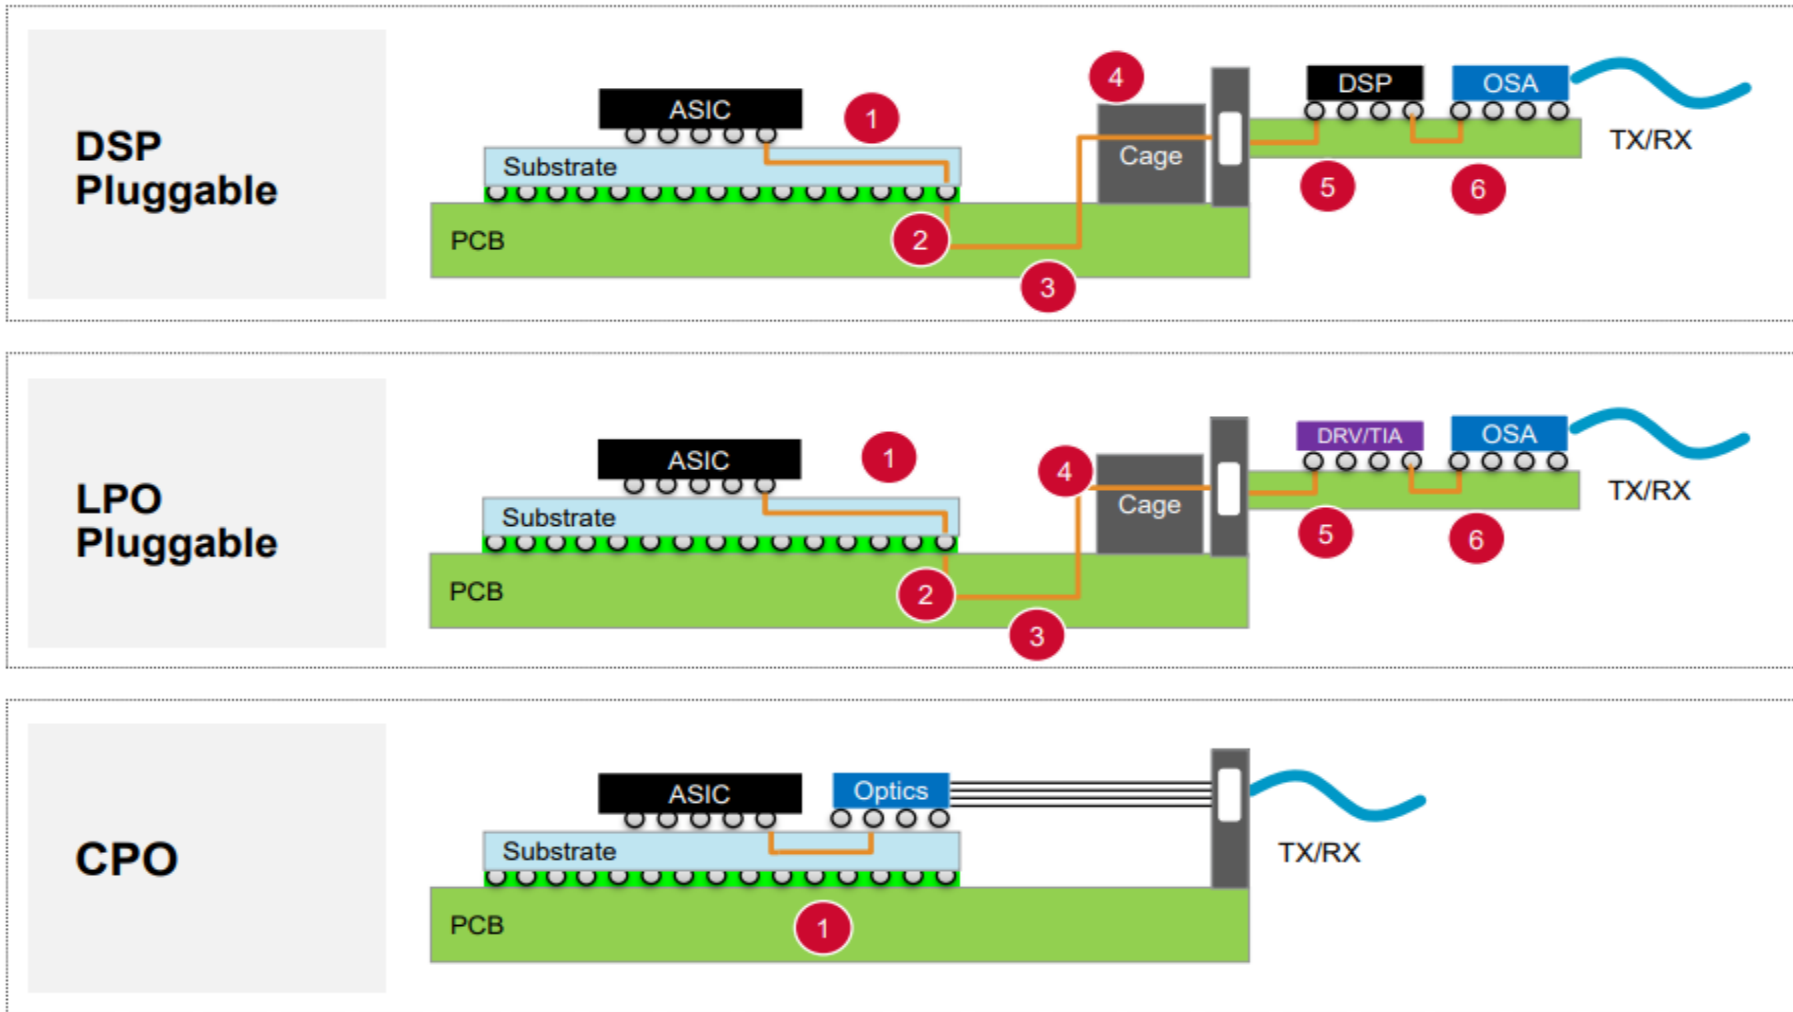
Less airflow required to cool

Lower Latency

>80% lower latency per Transceiver



CPO Insertion Loss Savings vs. Pluggables



\$

>30%
Cost Savings



>80%
Insertion Loss savings

- 1 ASIC through substrate
- 2 Via Losses
- 3 PCB Trace Loss
- 4 Connector Losses
- 5 Module PCB Losses
- 6 OSA Interconnect loss

BROADCOM

BROADCOM



AI/ML NIC Requirements

AI – Artificial Intelligence
ML – Machine Learning

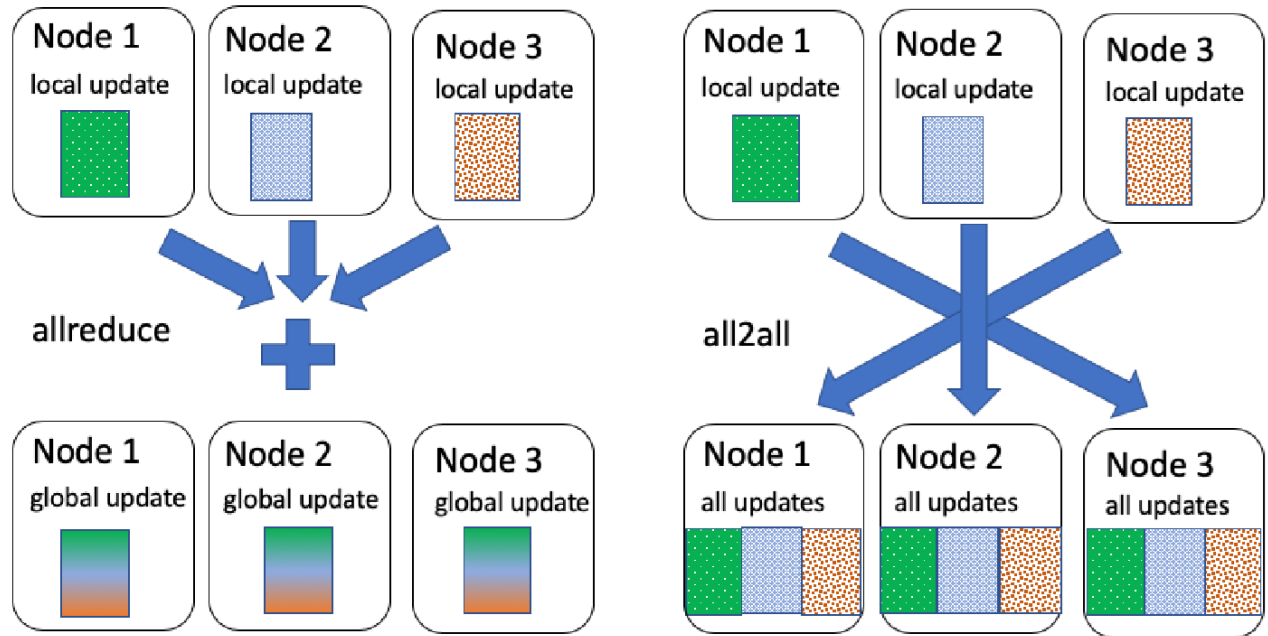


AI/ML Lifecycle



Job Completion Time (JCT) is an importance performance metric for AI/ML Ops

AI/ML Compute with Large Datasets

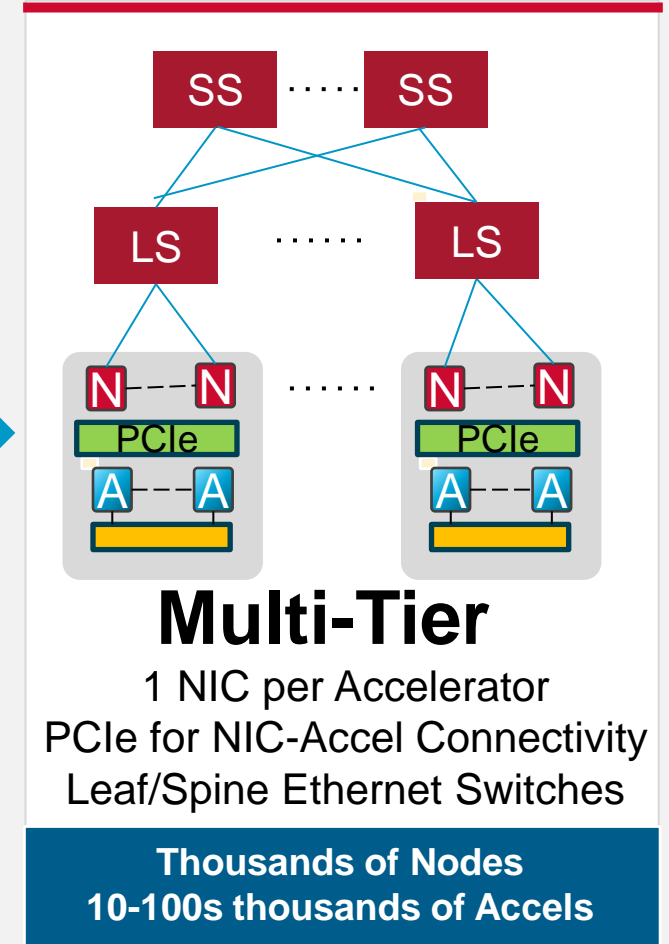
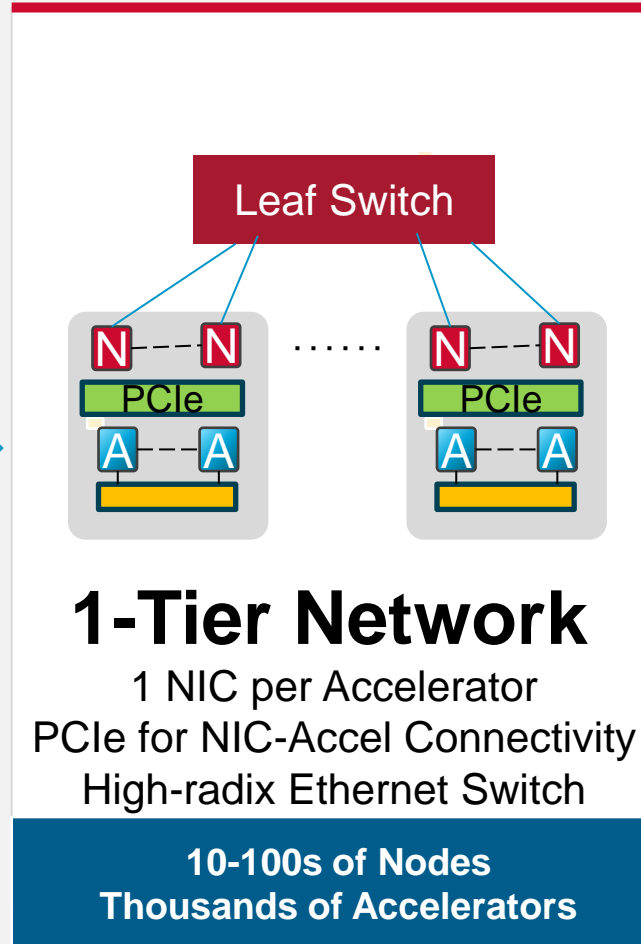
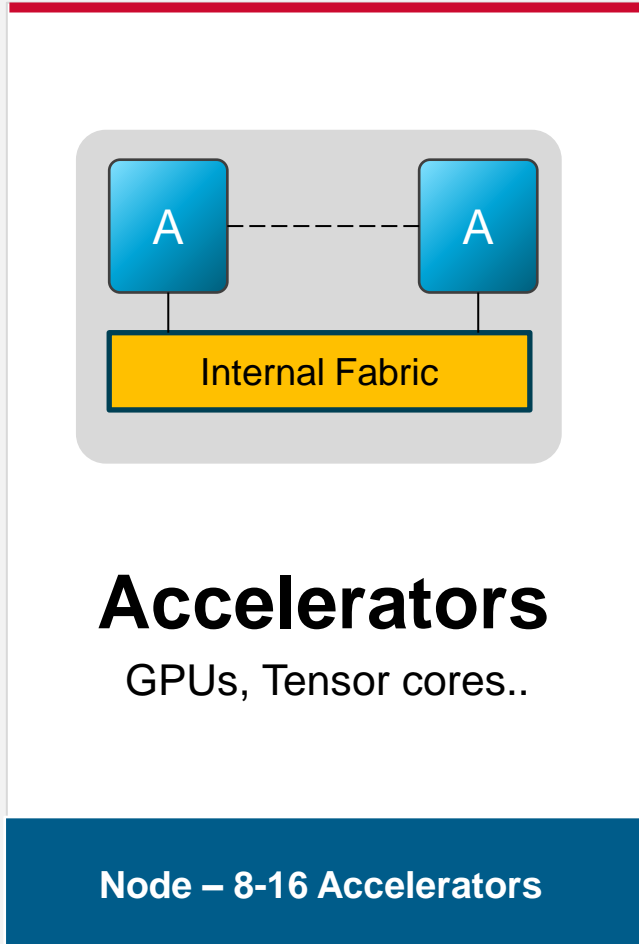


Training Requires
Distribution of Model
across multiple devices

- **Data Parallelism:** Aggregation of local parameters require *allreduce* communication
- **Model Parallelism:** Synchronization across all nodes require *all2all* communication
- **Tensor and pipeline parallelisms** within parallel models require *all-gather* and *send/recv*

High-Performance Networking is Essential to AI/ML Job Completion Time (JCT)

AI/ML Networking Scale




AI/ML Networking

Requirement	Today	Future
Scale (# of Nodes)	100-10000	10000-1000000
Network Port Speed	200 Gbps-400 Gbps	800 Gbps-1.6 Tbps
Transport/Network	Ethernet/InfiniBand	Ethernet




Ethernet Products




Broadcom Ethernet NIC Portfolio

Controllers




Adapters




Connectors/Transceivers

Fiber



Driving roadmap investment, most advanced and secure Ethernet solutions
Low Power Consumption, Multihost, Security, RoCE

400G Ethernet NICs



Broadcom 400G AI NICs (57600 Family)

6th Generation Hardened NIC Architecture



1x400G / 2x200G NIC



Lowest Power 400G NIC (17W)



4th generation RoCEv2



Advanced Congestion Control and Load Balancing



End-to-End Network Telemetry



5 Meter DAC and Linear Pluggable Optics



Broadcom DCSG Ethernet NIC Adapters: 57600

Product	N1400GD	N2200G	P1400GD	P2200G
Form Factor	OCP 3.0 SFF		PCIe LP	
PCIe Host	x16 PCIe Gen5 112G PAM4			
Network SerDes				
Ethernet Ports	1	2	1	2
Linear Drive Capable	Yes	Yes	Yes	Yes
Bandwidth	400G	400G	400G	400G
Bandwidth/Port	400G	200G	400G	200G
Adapter Power (Typical/Passive)	17W	17W	17W	17W
Thermal Level Passive	Tier 5 (250 @ 55C)	Tier 7 (350 @ 55C)	Level 3 (200 @ 55C)	Level 4 (350 @ 55C)
Connectors	QSFP112-DD*	QSFP112*	QSFP112-DD*	QSFP112*
Multi-Host	Yes		N/A	
Security	HW Secure Boot, Attestation			
Availability	Now	Samples: Now Production: Sept'24	Samples: Now Production: Aug'24	Now
EOL*	2040	2040	2040	2040
MPN	BCM957608-N1400GDP00	BCM957608-N2200GQP00	BCM957608-P1400GDF00	BCM957608-P2200GQF00

*Cable Guide available on our website
under Application Note
„BCM957608 Cable Solutions Guide“
and
under User Guide
„Ethernet NIC Supported Cables“



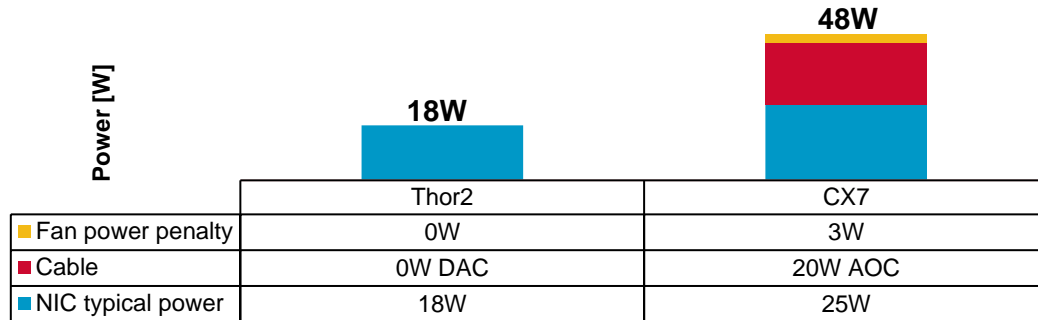
Broadcom Ethernet Controller Portfolio

NIC Bandwidth	2G / 4G	50G	200G	400G
Product (Chip Family)	5719 / 5720	57410	57500	57600
Ethernet ports	4 / 2	2	4	4
1GBASE-T / 10GBASE-T	✓	✓	-	-
Primary Link-Speeds	1G (x4 / x2)	2x 25G 2x 10GBASE-T	4x 50G 2x 100G, 1x 200G	4x 100G 2x 200G, 1x 400G
Network Serdes	1G	25G NRZ	56G PAM4	112G PAM4
PCIe Host Interface	Gen2	Gen3 x8	Gen4 x16	Gen5 x16
Adapter Power (Typical)	3W	10W	16W	17W
Min cooling @55C (LFM)	n/a	100	250	250
Multi-Host			x2 / x4	x2 / x4
RoCE v2		DCQCN	DCQCN	DCQCN / DCN
Latency (OSU)		4,000 ns	1,200 ns	1,100 ns
Silicon Root of Trust		✓	✓	✓
Attestation			✓	✓
IEEE1588 PTP	-	T-TSC	T-TSC / T-BC / T-GM	T-TSC / T-BC / T-GM
Synchronous Ethernet			✓	✓
TLS / kTLS HW offload				✓
Availability	In Production	In Production	In Production	Sampling Now Production June '24

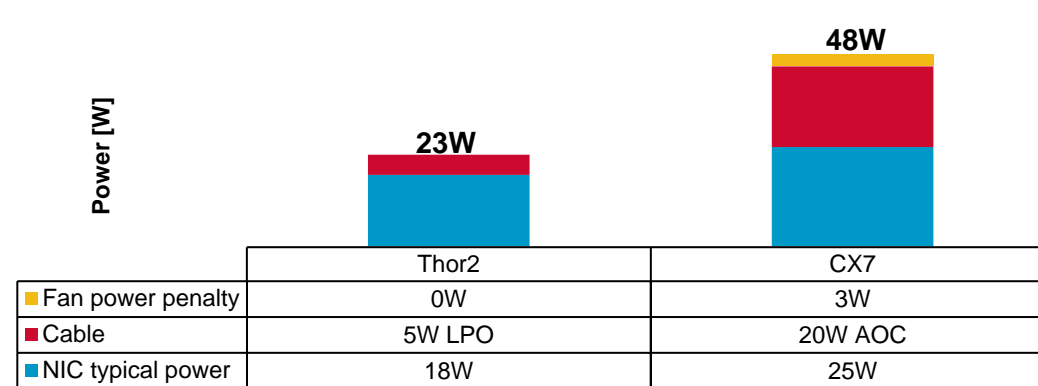
112G SerDes Power Advantage

Cable Length vs. Gauge

400G Intra-Rack Cabling (3m)



400G Rack to Rack Cabling (30m)



57600 – Tomahawk 5 passive cable gauge vs reach

Cable	1.5m	2m	2.5m	3m	3.5m	4m	5m
24 AWG							
26 AWG							
28 AWG							
30 AWG							
32 AWG							

Competition

Cable	1.5m	2m	2.5m	3m	3.5m	4m	5m
24 AWG	?	?	?				
26 AWG	?	?	?				
28 AWG							
30 AWG							
32 AWG							

Passive Copper Cables up to 4-5 meters

Referent

- Ausbildungen im Fernmeldebereich und in der Avionik
 - Studium der Nachrichtentechnik an der FH Augsburg und University of Ulster in Nordirland
 - Storage Engineer seit 1995
 - FAE im Storagebereich seit 2000
-
- Kontakt:
 - Tel +49 89 41352 0121
 - Email: dietmar.hinze@broadcom.com
 - Support:
 - Email: tech.support@broadcom.com





Thank you!



BROADCOM[®]

connecting everything[®]