



Broadcom 9400 Tri-Mode Controller, Anwendungsgebiete und Tipps zum NVMe-Einsatz




Dietmar Hinze

Field Application Engineer

20.02.2019

Overview

- Broadcom Overview (in short)
- Tri-mode Controller (Overview)
- Eco-System Update



Broadcom Overview



Global Technology Leadership

Data Center




- Enterprise
- Cloud
- Telecom

Broadband



- DSL
- Cable
- PON

Mobile



- Wi-Fi/Bluetooth
- GNSS
- RF Filters


Industrial



- Automotive
- Factory automation
- Alternative energy

Data Center: Networking & Storage Connectivity






Tri-Mode Controller




PCIe/SAS/SATA Host Bus & MegaRAID Adapter



Application

General purpose Servers
High Storage Demand
High Data Protection Demand



Application

Software defined Storage
Attaching external RAID Devices
Attaching Tape Storage Drives

Board Name Decoder

Brand Name

MegaRAID SAS = MegaRAID Entry, Value, Feature
 BCL SAS = Host Bus Adapters

Connectors

"X"i = # of Internal ports
 "X"e = # of External ports
 "X"i"X"e = # of Internal and External ports

MegaRAID SAS

94 6 0 - 16i

Technology

9 = SATA+SAS+PCIe

Generation

	Host Interface	Storage
1	PCIe Gen 1	3 Gb/s SAS / SATA
2	PCIe Gen 2/3	6 Gb/s SAS / SATA
3	PCIe Gen 3	12 Gb/s SAS / SATA
4	PCIe Gen 3	Tri-mode 12 Gb/s SAS/ SATA/ PCIe 3
5	PCIe Gen 4	Tri-mode 12 Gb/s SAS/ SATA/ PCIe 4
6	PCIe Gen 4	Tri-mode 24 Gb/s SAS/ SATA/ PCIe 4

Subcategory

0 = HBA Initiator / Target
 1 = HBA Integrated RAID
 4 = MegaRAID Entry (iMR)
 6 = MegaRAID Value
 8 = MegaRAID Feature

Linear Variant Digit

Starts at "0"; This digit is used for changes that are meant show succession from an older version to a newer version.

Introducing 9400 Family HBA & MegaRAID



- New Chipdesign utilizing ARM A15 1.2Ghz architecture
- Larger L2 Cache and OCM
- TRI-Mode SerDes supporting SAS/SATA/NVMe Media



- Updated Memory Controller with 16% increased bandwidth
- Higher Memory Capacity with up to 4GB of DRAM
- Onboard flash memory for automated Cache offload



- Improvements to Fastpath
- Minimized Firmware based IO processing
- Parallel IO Processing

HBA 9400 Series Tri-Mode Storage Adapters

9400-16i, 9400-16e, 9400-8i, 9400-8e, 9405W-16i, 9405W-16e



16-port Tri-Mode HBAs
9405W-16i (05-50047-00)
9405W-16e (05-50044-00)



16-port Tri-Mode HBAs
9400-16i (05-50008-00)
9400-16e (05-50013-00)
9400-8i8e(05-50031-02)



8-port Tri-Mode HBAs
9400-8i (05-50008-01)
9400-8e (05-50013-01)

Applications

- High-port count SAS/SATA/NVMe adapters for direct attached high connectivity applications
- Tri-Mode connectivity enabling maximum data center flexibility
- Flexible solutions for cloud computing
- External storage requiring high connectivity SAS/SATA interface for host or drive side connect

Key Features

- Tri-Mode Storage Interface Ports
 - SFF-8680 Bay
 - x1 SAS
 - x1 SATA
 - x2 SAS (Multi Link)
 - Two x1 SAS (Dual port usingMPIO)
 - SFF-8639 (U.2) Bay
 - x2, x4 NVMe
- Supports 12, 6, and 3Gb/s SAS and 6, 3Gb/s SATA data transfer rates
- Up to 8 storage interface PCIe links.
- Each link supporting x4 or x2 link widths up to 8.0 GT/s (PCIe Gen3) per lane
- PCIe Gen3 x16 Host Interface for max. throughput

MegaRAID 9400 Series Tri-Mode Storage Adapters

9460-16i, 9460-8i, 9440-8i, 9480-8i8e



16-port Tri-Mode RAID
9460-16i (05-500011-00)



8-port Tri-Mode RAID/iMR
9460-8i (05-50011-02) & 9440-8i (05-50008-02)



8-internal / 8-external -port Tri-Mode RAID
9480-8i8e (05-500031-00)


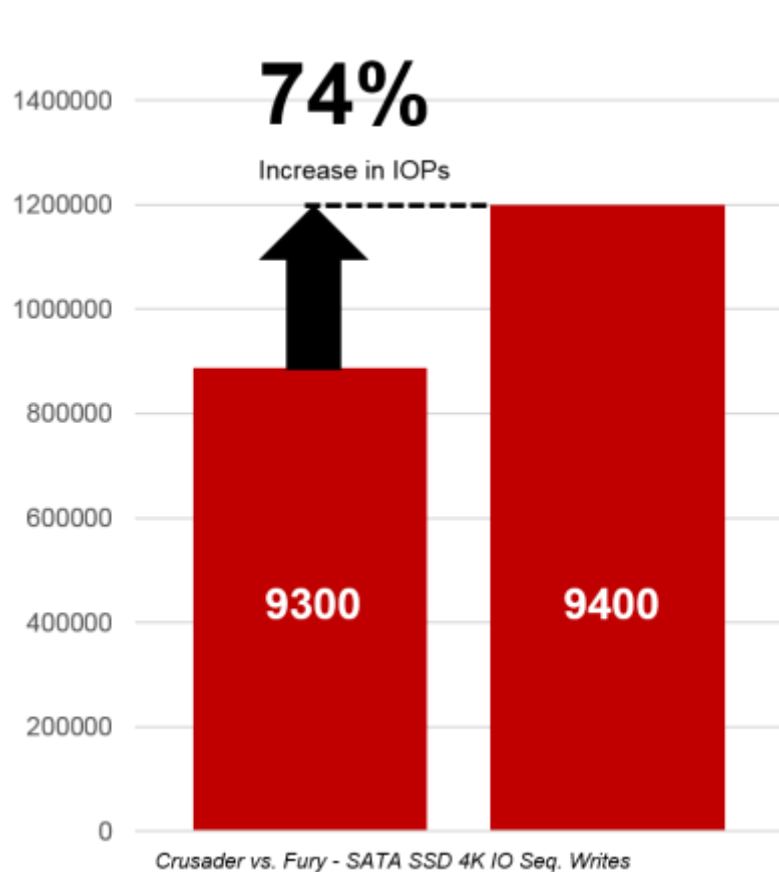
Applications

- High-port count SAS/SATA/NVMe controllers for direct attached high connectivity applications
- Tri-Mode connectivity enabling maximum data center flexibility
- Flexible solutions for cloud computing
- External storage requiring high connectivity SAS/SATA interface for host or drive side connect

Key Features

- Tri-Mode Storage Interface Ports
 - SFF-8680 Bay
 - x1 SAS
 - x1 SATA
 - x2 SAS (Multi Link)
 - Two x1 SAS (Dual port using MPIO)
 - SFF-8639 (U.2) Bay
 - x2, x4 NVMe
- Supports 12, 6, and 3Gb/s SAS and 6, 3Gb/s SATA data transfer rates
- Up to 8 storage interface PCIe links.
- Each link supporting x4 or x2 link widths up to 8.0 GT/s (PCIe Gen3) per lane

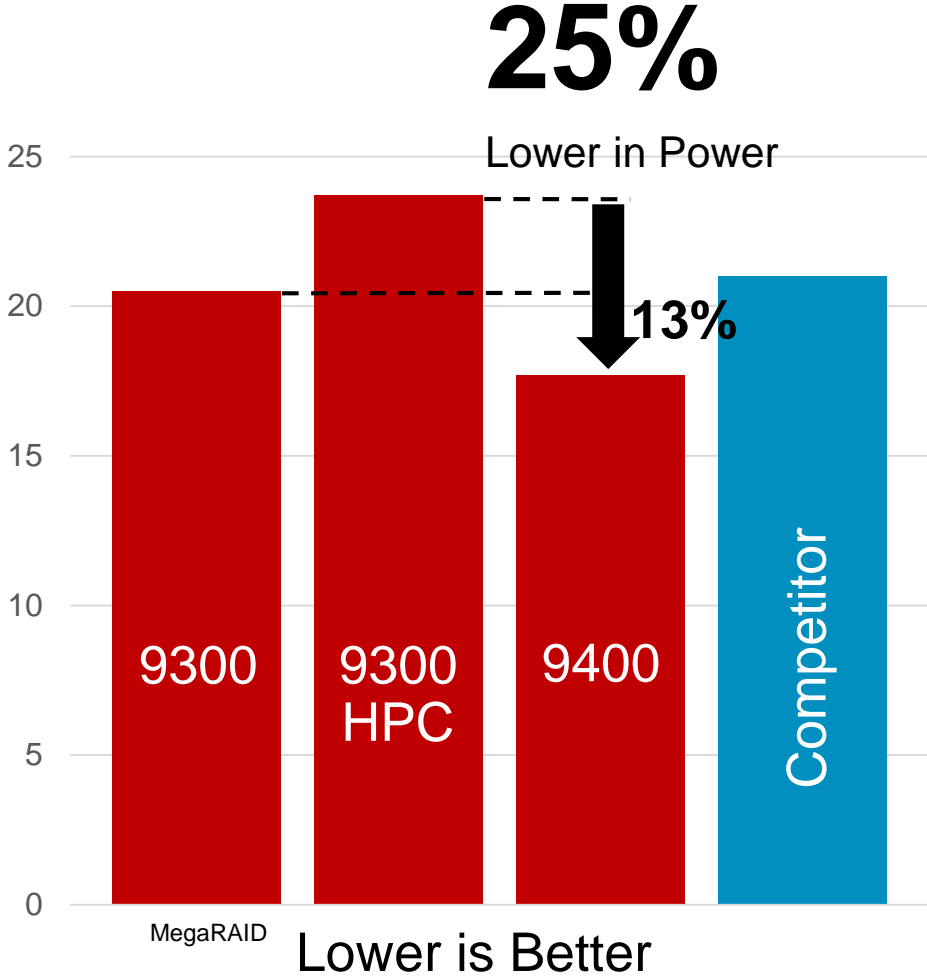
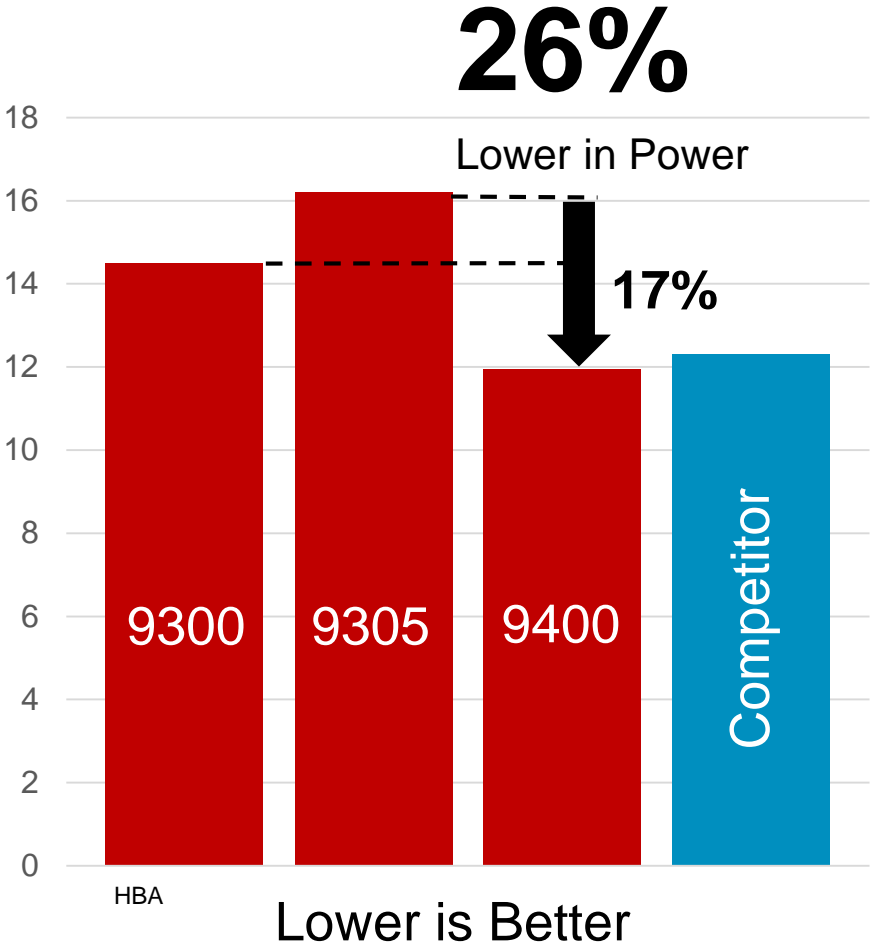
Let SATA SSDs rock!



Improved Fastpath engine shows best SSD performance

- 9400 HBA Family performs 74% higher in Seq. Writes
- 9400 HBA Family increases bandwidth by up to 7% in Writes
- 9400 MegaRAID Family performs up to 44% higher in Seq. Writes (RAID 5)
- 9400 MegaRAID Family increases bandwidth by up to 13% in Writes (RAID 5)

Reduced Power




Place it right



SAS 9300 family

- Any standard 8 Port Opportunity
 - Backup/Tape
 - External SAS Storage



SAS 9305 family

- Need for high direct attached possibility
 - Software Defined Storage
 - Block Storage
 - Low Latency
 - Multiple Expander Systems



SAS 9400 family

- High Performance needs
- Low Power / Heat needs
- NVMe Connection

Major differentiators - HBAs

	9300 family (Fury)	9305 family (Cutlass)	9400 family (Tomcat/Crusader)
Port Count	8	16-24	8-16
Connectivity	SAS/SATA	SAS/SATA	SAS/SATA/NVMe
IR support	Yes	No	No
Performance	1.45M IOPs / 6.000 MB/s	1.5M IOPs / 6.400 MB/s	1.7M IOPs / 6.840 MB/s
Power Consumption	14.5W max.	16.2W max.	11.95W max
Cooling Requirements	200 LFM @ 55°C	200 LFM @ 55°C	200 LFM @ 55°C
Price	\$	\$\$	\$\$\$

Place it right




MegaRAID 9300 family

- Any standard 8 Port Opportunity
- SSD Caching



MegaRAID 9300-16i/24i/8i8e family

- Need for high direct attached possibility
- Eliminating Expander issues
 - Connectivity issues
 - Latency
- SSD Caching



MegaRAID 9400 family

- High Performance needs
- Low Power / Heat needs
- NVMe Connection
- Flash/SSDs as Primary Storage

Major differentiators - MegaRAID

	9300 family (Invader)	9300 family (Intruder)	9400 family (Harpoon/Ventura)
Port Count	8	16-24	8-16
Connectivity	SAS/SATA	SAS/SATA	SAS/SATA/NVMe
Cache	1-2GB	2-4GB	2-4GB
CacheVault onboard	No	Yes (Supercap option)	Yes (Supercap option)
Cache Cade support	Yes	Yes	No
LSA support (WebGUI)	Yes	Yes	Yes
Performance	0.9M IOPs / 6.000 MB/s	1M IOPs / 6.400 MB/s	1.4M IOPs / 6.840 MB/s
Power Consumption	20.5W max.	23.7W max.	17.7W max
Cooling Requirements	200 LFM @ 55°C	300 LFM @ 55°C	250 LFM @ 55°C
Price	\$	\$\$	\$\$\$

Introducing next Generation Management

OS GUI - LSI Storage Authority

- Webbased GUI
- Same Feature-Set like MSM
- No client side installation needed
- Compare to MSM
 - Significantly better performance
 - Highly secure (as secure as browser)
 - Less backend traffic to server
 - Can be scripted using ReST
 - Has potential to run on BMC



Command Line - Unified StorCLI

- Merge of SASxFlash/SASxIRCU & StorCLI
- Single Point of Management
- Use of common StorCLI command set
- Enhanced command set

PreBoot - MegaRAID Human Interface Infrastructure (HII)

- Integration in servers UEFI BIOS
- Single Point of Configuration
- Additional functions (e.g. FW update)



SAS / SATA – Just Better!



Lower Power

- Up to 25% lower in power consumption
- Reduced cooling requirement



High Performance


- Up to 247% performance improvement over previous Generation
- Best solution to work with SSD



Updated Management & Monitoring

- Unified StorCLI combines CLI for HBA and MegaRAID
- Webbased GUI LSA

backwards compatible for generations!




NVMe Ecosystem Update





U.2 Enabler Cable

- x4 NVMe drive connect requires a special cable
- Cable re-orders lanes and effectively de-swizzles the adapter layout
- U.2 Enabler Cable also necessary for maximizing card flexibility when connecting to x2 and x1 NVMe drives



These cables are for NVMe usage only – improper cabling can cause damage of Controller and Media



Adapter Cable Information

Ventura/Crusader Adapter Cables


MPN	Description	Where Used	Status
05-50061-00	Cable, U.2 Enabler, HD to HD(W) 1M	Intel Grantley or backplanes using legacy PCIe Mini SAS HD wiring. Also works with SMC backplanes, but without LED functionality.	Production
05-50062-00	Cable, U.2 Enabler, HD to OCuLink 1M	Intel, SMC Oculink based backplanes	Production
05-50063-00	Cable, U.2 Enabler, HD to SlimLine 1M	Backplanes with 2x4 Simline connectors	Production
05-50065-00	Cable, U.2 Enabler, HD to SFF8639 0.5M	Direct connect to SFF-8639 drives. 0.5M cable.	Production
05-50064-00	Cable, U.2 Enabler, HD to SFF8639 1M	Direct connect to SFF-8639 drives. 1M cable.	Production
05-60008-00	Cable, U.2 Enabler, HD to HD(W), SMC 1M	Targeted for SMC Purley backplanes and supports LED functionality.	Prototype

Usage Scenarios – NVMe getting started

- Single 16 Port MegaRAID / HBA
 - 2 pcs. x4 NVMe Media
 - 8 pcs. SAS/SATA Media

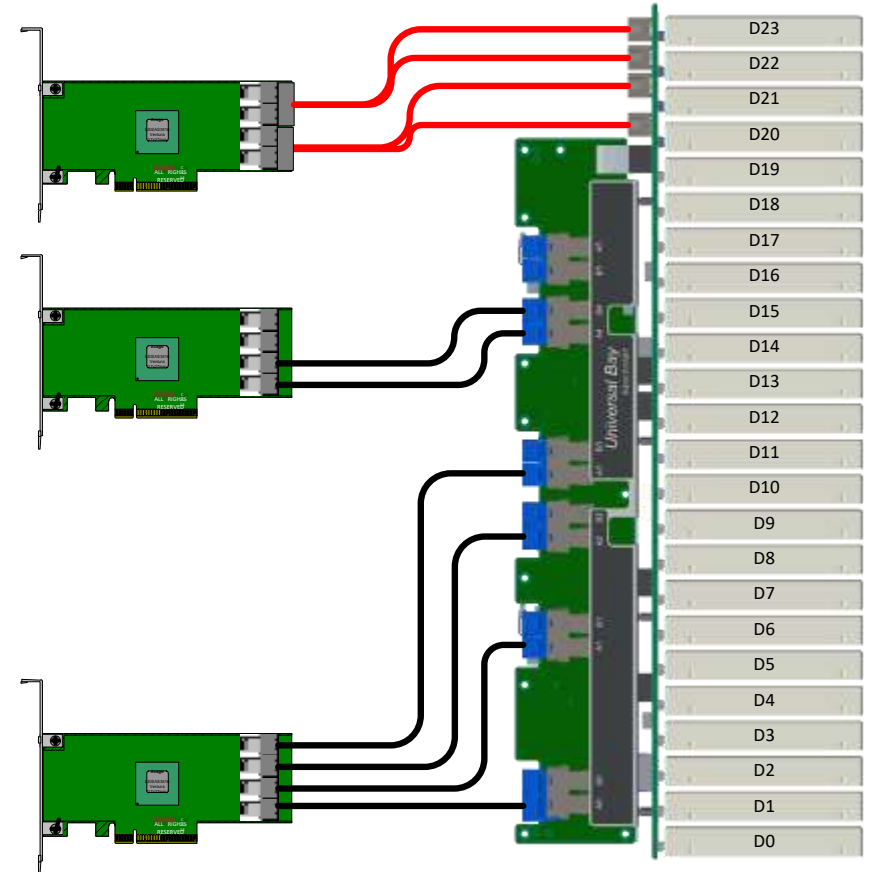
- ✓ Best of both Worlds – Performance & Capacity
- ✓ Single Monitor and Management
- ✓ Saving PCIe Lanes by using single Controller
- ✓ NVMe Data Protection utilizing HW MegaRAID

- ✗ PCIe x8 Host Connection potential
Performance bottleneck



Usage Scenarios – Get the max out of it

- 1x 16 Port MegaRAID / HBA
 - 4 pcs. x4 NVMe Media
- 1x 8/16 Port MegaRAID / HBA
 - 8 pcs. SAS/SATA or
 - 20 pcs. SAS/SATA Media
- ✓ Max of both worlds
- ✓ Single Monitor and Management
- ✓ NVMe Data Protection utilizing HW MegaRAID
- ✓ Bandwidth optimization by using multiple PCIe x8 Slots
- ✗ Multiple Controllers have to be managed
- ✗ Higher PCIe lane usage



Configuration considerations

Item	Answers
Determine what type of NVMe enabled system/backplane will be used.	<ul style="list-style-type: none">• Has the backplane a SAS expander and NVMe slots
Determine how to properly connect the adapter to NVMe drives.	<ul style="list-style-type: none">• Must use U.2 Enabler Cable• Option without backplane - direct attached cables
Determine how to set the MegaRAID or IT Profiles to enable NVMe.	<ul style="list-style-type: none">• SAS/SATA Only Mode(default)• NVMe Only Mode• Mixed Mode
Status LED's on SMC slots	<ul style="list-style-type: none">• SMC issues with LED control (on white MiniSAS HD backplanes)• FW fix for OCulink backplanes is available• Backplane drives activity LED
How will the OS show the NVMe devices	<ul style="list-style-type: none">• NVMe will show up as SCSI device• No differences for the tools (LSA)

NVMe / Mixed Mode support

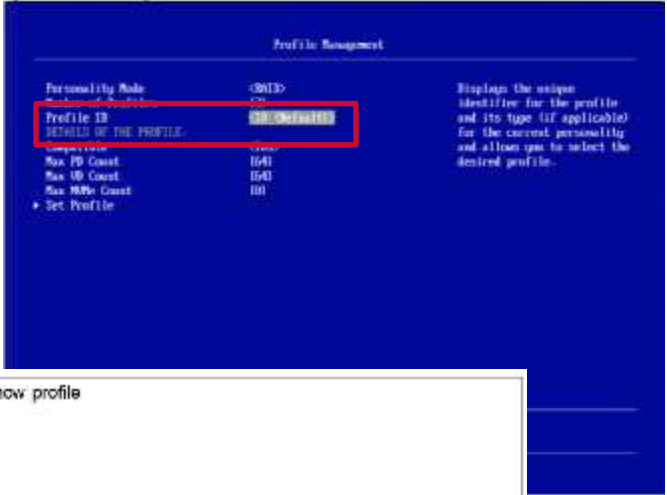
	SAS / SATA only	NVMe only	Mixed Mode
MegaRAID 9440-8i	✓	✓	✗ using U.2/NVMe x4
MegaRAID 9460-8i	✓	✓	✗ using U.2/NVMe x4
MegaRAID 9460-16i	✓	✓	✓
MegaRAID 9480-8i8e	✓	✓ For internal Ports only	✓ NVMe supported on internal Ports only
SAS 9400-8i	✓	✓	✗ using U.2/NVMe x4
SAS 9400-16i	✓	✓	✓
SAS 9400-8e	✓	✗	✗
SAS 9400-16e	✓	✗	✗
SAS 9405W-16e	✓	✗	✗
SAS 9405W-16i	✓	✓	✓

Considerations

- U.2 Enable Cable occupies 2 Wide Ports (unswizzle)
- NVMe is not supported for external connectivity (today)

NVMe Deployment Model

- Profile driven
 - MR: Profiles are hold in FW and need to be activated
 - HBA: Profiles requires flashing the proper FW
 - Possible Profiles:
 - SAS/SATA only
 - NVMe only
 - SAS / SATA / NVMe mixed
- Set profile either by
 - StorCLI
 - HII (EFI BIOS)
- We ship with SAS/SATA profile as default
 - Customer can enable NVMe mixed mode with
 - Storcli (profile 13)



```
C:\Ventura>storcli64.exe /c0 show profile
Controller = 0
Status = Success
Description = None

Controller Properties :
=====
Mode ProfID MaxPhyDrv MaxLD MaxNVMeDev MaxAHCIDev isDefault isCurrent
-----
RAID 0b      4   64      4      0 No      Yes

ProfID=Profile ID|MaxPhyDrv=Maximum Physical Drives
MaxLD=Maximum Logical Devices|MaxNVMeDev=Maximum NVMe Devices
MaxAHCIDev=Maximum AHCI Devices

C:\Ventura>storcli64.exe /c0 set profile profileID=0a
Controller = 0
Status = Success
Description = ProfileID a has been set. Please reboot the system for the changes to take effect
```



BROADCOM®

connecting everything®