POW3R
einfach. digital. agil.

PRÄSENTIERT

IPL in unter 30 Sekunden
NVMe

POW3R
einfach. digital. agil.

EIN BEITRAG VON

Peter Kindiger
System Engineer
Tech Data GmbH & Co. OHG
E-Mail: ibmsupport@techdata.de
Kistlerhofstraße 75
81379 München

Dr. Wolfgang Rother
IBM Power Platform Team DACH
IBM Deutschland GmbH
E-Mail: wrother@de.ibm.com
Nahmitzer Damm 12
12277 Berlin

Überblick

Was ist NVMe?

Vergleich NVMe vs.15k Platten

Ist NVMe für IBM i geeignet?

Preis Leistungsverhältnis

Anwenderbericht

What is NVMe? (Non-Volatile Memory Express)

- **NVMe is a communications interface and protocol**
  - NVMe is designed from the ground up to deliver high bandwidth and low latency storage access
  - NVM Express defines an efficient interface for host software to communicate with a non-volatile memory subsystem over PCI Express (NVMe over PCIe)
  - Functionally analogous to SAS and SATA but designed to reduce overhead from drivers, OS and application
  - Leverages poll-based I/O completions as opposed to interrupt-based completions
- **NVMe uses PCIe fabric**
  - Multiple devices today in the market
  - Multiple form factors including 2.5" drives
- **NVMe has been designed for high performance**
  - Increased IOPs, bandwidth and lower latency
  - Exploit Flash and next-generation Non-Volatile Memories
  - Leverage multicore environments, high I/O parallelism

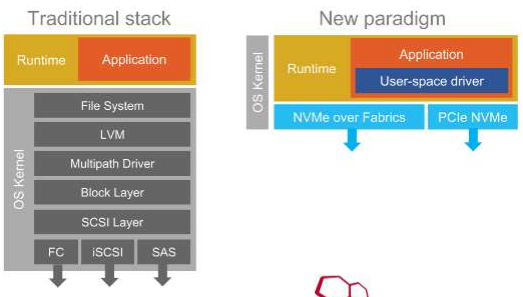

Where does NVMe performance come from ?

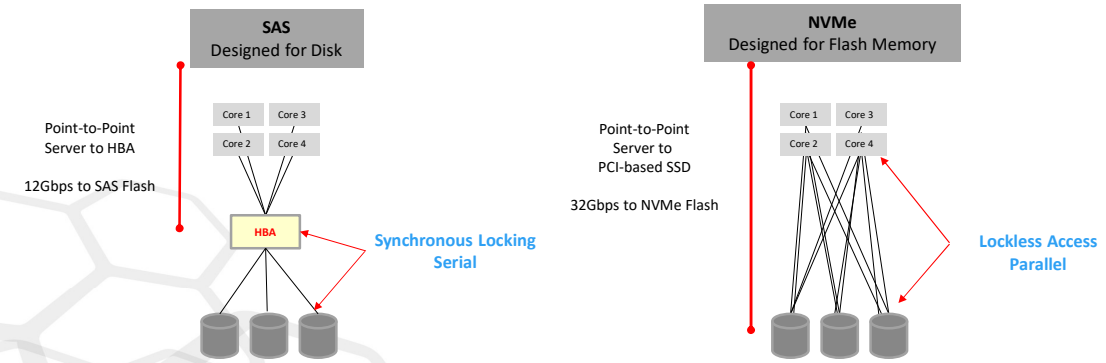

Serial Attached SCSI (SAS) Protocol

One Channel
One Queue
32 Commands
Software stack at each layer

Non-Volatile Memory (NVMe) Protocol

Up to 64,000 Queues
64,000 Commands per Queue
Each Core has dedicated queues per SSD

NVMe over PCIe reduces latency and increases throughput

5

POWER9 M.2 NVMe support (at time of announcement)

Power9 SCO models have two internal direct attached storage connectors
Connector support either:

- NVMe carrier card & attaches two 400 GB M.2 NVMe drives
- SAS controller requiring DASD backplane (like POWER8)

Note: not both at the same time

M.2 NVMe on Power9 SCO

- A maximum of four x M.2 NVMe drives
- Will be higher performance than SAS DASD in backplane
- Will not support concurrent maintenance (unlike SAS drives)
- Will have a write endurance of 1 drive write per day
- Intended primarily to store and boot OS (AIX / VIOS) images
- Each NVMe device→separate PCIe endpoint = assign to different LPARs
- NVMe drives may be assigned to the VIOS & virtualized to client OS

POWER9 U.2 NVMe support for E950 (at time of announcement)

E950 has four external storage bays
Independent of SAS controllers
    NVMe ~price of a SAS adapter
Concurrent 8 SAS and 4 NVMe drives
    SAS takes 1 or 2 PCIe slots

U.2 NVMe
- A maximum of four x U.2 NVMe drives
- Higher performance than SAS SSD
- Support concurrent maintenance (unlike Scale-Out S922/S924)
- Write endurance is 2.4 drive write per day DWPD (5 years)
- Intended primarily to store and boot OS (AIX / VIOS) images
- Each NVMe device is a separate PCIe endpoint = assign to different LPARs
    - On the HMC, it looks like each has own adapter slots
- NVMe drives may be assigned to the VIOS & virtualized to client OS
- Warranty: 5 years if not warn out. A tool to monitor wear is provided for AIX/Linux

Different to Scale-Out
→ **U.2, not M.2!**
→ these have higher DWPD
→ Better RAS

For IBM & Business Partner Education Use

POWER9 U.2 NVMe support for E950 (at time of announcement)

E950 has four external storage bays
Independent of SAS controllers
    NVMe ~price of a SAS adapter
Concurrent 8 SAS and 4 NVMe drives
    SAS takes 1 or 2 PCIe slots

U.2 NVMe
- A maximum of four x U.2 NVMe drives
- Higher performance than SAS SSD
- Support concurrent maintenance (unlike Scale-Out S922/S924)
- Write endurance is 2.4 drive write per day DWPD (5 years)
- Intended primarily to store and boot OS (AIX / VIOS) images
- Each NVMe device is a separate PCIe endpoint = assign to different LPARs
    - On the HMC, it looks like each has own adapter slots
- NVMe drives may be assigned to the VIOS & virtualized to client OS
- Warranty: 5 years if not warn out. A tool to monitor wear is provided for AIX/Linux

Different to Scale-Out
→ **U.2, not M.2!**
→ these have higher DWPD
→ Better RAS

For IBM & Business Partner Education Use

POWER9 Scale-Out I/O Enhancements

July 2020

Scale-Out Gen4 Refresh System Boot Options

S924 and S914
- Up to 4 NVMe devices
- Up to 6 SAS plus up to 2 NVMe devices
- Up to 12 SAS devices + RDX
- Up to 18 SAS devices

S922
- Up to 4 NVMe devices
- Up to 8 SAS devices

All models
- Expansion capabilities for the EXP12SX or EXP24SX Drawer
- Concurrently maintainable PCIe Gen4 slots with no storage backplane required
  - NVMe
  - Fibre Channel
  - Ethernet

New for IBM i
- P05 tier configurations now support 2 x 1.6TB or **2 x 3.2TB** NVMe devices

Scale-Out Gen4 Refresh NVMe Options

- No M.2 NVMe options
- New U.2 front accessible NVMe options
  - 15mm Enterprise class devices in 1.6TB, 3.2TB and 6.4TB sizes

Scale-Out Gen4 Refresh New Fibre Channel Option

- Feature Codes #EN1G and #EN1H
- **VIOS:** 3.1.1.25 or later
- **AIX:** 7.2 TL4 with Service Pack 2-2028
- Supported on all S914, S922, S924, E950 & E980

This PCIe Gen3 16 Gigabit dual-port optical fibre channel (FC) adapter is a high-performance x8 short form adapter based on the Marvell QLE2692 PCIe host bus adapter.

Each port provides single initiator capability over a fibre link or with NPIV, multiple initiator capability is provided. The ports are SFP+ and include an optical SR transceiver. The ports have LC type connectors and utilize shortwave laser optics. The adapter operates at link speeds of 4, 8, and 16Gbps and will automatically negotiate to the highest speed possible. The adapter supports boot.

Scale Out Gen4 Refresh NVMe Options

- No M.2 NVMe options
- New U.2 front accessible NVMe options
  - 15mm Enterprise class devices in 1.6TB, 3.2TB and 6.4TB sizes
  - 7mm 800GB Datacenter class device in 15mm carrier
  - S922 requires EJ1T or EJ1U
  - S924/S914 requires EJ1V or EJ1W
- Continued support for PCIe AIC NVMe adapters

- New for IBM i
  - P05 tier configurations now support 2 x 1.6TB or **2 x 3.2TB** NVMe devices (available in eConfig as of 07/28)

NVMe Product Form Types used in POWER9 Systems

Add-in Card (AIC)

- x4 / x8 PCIe Gen3 or Gen4
- High performance and power (25-50 Watts)
- Focus on high capacity - up to 6.4TB currently
- Best performance and latency

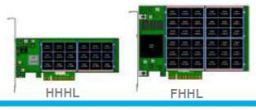

Feature Code	PCIe3 Card NVMe Flash Adapter 1.6TB	3.2TB	6.4TB
AIX/Linux (LP/FH)	#EC5G / #EC5B	#EC5C / #EC5D	#EC5E / #EC5F
IBM i (LP/FH)	#EC6U / #EC6V	#EC6W / #EC6X	#EC6Y / #EC6Z

NEW

Feature Code	PCIe4 Card NVMe Flash Adapter 1.6TB	3.2TB	6.4TB
AIX/Linux (LP/FH)	#EC7A / #EC5B	#EC7C / #EC7D	#EC7E / #EC7F
IBM i (LP/FH)	#EC7J / #EC7K	#EC7L / #EC7M	#EC7N / #EC7P

U.2

- x2 / x4 PCIe Gen3
- Higher power required to achieve max SSD performance
- Expect to be direct attached to CPU vs through a HBA

2.5" 7mm

E950/E980
800GB - #EC5J
1.6TB - #EC5K
3.2TB - #EC5L

NEW

S9xx G Models
800GB - #EC5X
1.6TB - #ES1F/ES1E
3.2TB - #ES1H/ES1G
6.4TB - #EC5W/EC5V

2.5" 15mm

M.2

- x2 / x4 PCIe Gen3
- Different physical sizes
- Power less than 9 watts
- Limitations: Minimal Connector Plug Capability (<10), No hot swap
- Targeted as a replacement for SATA SSDs

- Withdrawn from marketing 3/24/2020, announcement 920-065

Card #EC59 with 1-2 #ES14

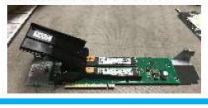

S9xx
400GB #ES14

NVMe on POWER9 Systems

- NVMe are only supported in the system unit of POWER9 servers. They are not supported in a PCIe Gen3 I/O drawer due to bandwidth and potential performance issues.
- 9009-41G support up to 11 NVMe. 9009-41A supports up to 3
- 9009-42G and 9009-22G support up to 14 NVMe (11 with only one processor module)
- 9009-42A and 9009-22A support up to 5 NVMe (3 with only one processor module)
- IBM i 4-core P05 system is limited to 2 x 1.6TB or 3.2TB devices only. No mixing of NVMe and SAS drives is allowed on the 4-core (ten maximum of SAS drives or two maximum of NVMe).
- E950 supports up to 9 NVMe PCIe addin cards and up to 4 U.2 7mm NVMe devices per server
- E980 supports up to 8 NVMe PCIe addin cards per drawer (6 first drawer, 8 each drawer 2, 3, and 4 for a maximum total of 30) and up to 4 U.2 7mm NVMe devices per drawer
- Note that using multiple PCIe card NVMe will affect available I/O slots which can affect system configurations with multiple partitions and using external drawers may affect the maximum amount of PCIe NVMe.

© Copyright IBM Corporation 2019

NVMe on POWER9 Systems

- Internal HDDs/SSDs are allowed in the system unit “with” NVMe (except the S914 4-core with IBM i).

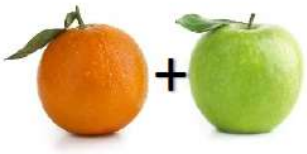

- A storage backplane is “not” required when PCIe card NVMe adapters but may be nice to have for future storage expansion.
- A Scale out NVMe backplane is not required if no U.2 devices are to be populated

POW3R einfach. digital. agil.

POWER9 Enterprise Scale-Out Family

S914 9009-41G	S922 9009-22G	S924 9009-42G	H922 9223-22H	H924 9223-42H	L922 9008-22L
• 1-socket, 4U Rack & Tower • 4,6,8 cores/ socket • 16 IS RDIMM slots • 1TB memory • 2 CAPI 2.0 Slots • Internal RDX Media • 8 PCIe adapters slots • **11 NVMe devices ***	• 1,2-socket, 2U • 1, 4, 8,10, 11 cores/ socket • 32 IS RDIMM slots • 4TB memory • 4 CAPI 2.0 Slots • 11 PCIe adapters slots • **14 NVMe devices**	• 1,2-socket, 4U • 8,10, 11, 12 cores/ socket • 32 IS RDIMM slots • 4TB memory • 4 CAPI 2.0 slots • Internal RDX Media • 11 PCIe adapters slots • **14 NVMe devices**	• 1,2-socket, 2U • 4, 8,10 cores/ socket • 32 IS RDIMM slots • 4TB memory • 4 CAPI 2.0 Slots • 9 PCIe adapters slots • **5 NVMe devices** • AIX, IBM i up to 25%	• 1,2-socket, 4U • 8,10,12 cores/ socket • 32 IS RDIMM slots • 4TB memory • 4 CAPI 2.0 slots • Internal RDX Media • 11 PCIe adapters slots • **5 NVMe devices** • AIX, IBM i up to 25%	• 1,2-socket, 2U • 8,10,12 cores/ socket • 32 IS RDIMM slots • 4TB memory • 4 CAPI 2.0 Slots • 9 PCIe adapters slots • **5 NVMe devices**
***4-core S914 IBM i P05 2 x 1.6TB or 3.2TB only**		AIX IBM i Linux + PowerVM			Linux + PowerVM

Technology Leadership

- PowerVM hypervisor is built in with no license charge, so every workload is virtualized, mobile, and cloud ready
- Up to 4TB in 2 sockets - DDR4 Industry Standard memory RDIMMs
- High Speed 25Gb/s external ports – one per socket
- Embedded Analytics and Algorithms on the chip help run POWER9 at an always optimized frequency
- G-models: 2x more bandwidth and up to 180% more Enterprise NVMe capacity for a seamless integration between on-prem and public clouds with a full PCIe Gen4 architecture

Vergleich NVMe vs. 15k Platten

U.2 Modules (Storage -> Hard Drives) with a maximum of 4 (all models)

- #EC5X - Mainstream (2.4 DWPD) 800 GB SSD *PCIe3* NVMe U.2 7mm (in 15mm carrier) module for AIX/Linux
- #ES1E / #ES1F (IBM i) - Enterprise (3 DWPD) 1.6 TB SSD **PCIe4** NVMe U.2 15mm module
- #ES1G / #ES1H (IBM i) - Enterprise (3 DWPD) 3.2 TB SSD **PCIe4** NVMe U.2 15mm module
- #EC5V / #EC5W (IBM i) - Enterprise (3 DWPD) 6.4 TB SSD **PCIe4** NVMe U.2 15mm module

PCIe3 NVMe Adapter (Adapters -> NVMe Adapter -> Adapters) with a maximum of 7 (S914) or 10 (S924, S922) / No MEX I/O drawer support

- #EC5G (LP) / #EC5B (FH) / #EC6U (LP IBM i) / #EC6V (FH IBM i)) - 5 DWPD *PCIe3* x8 1.6 TB NVMe Flash Adapter
- #EC5C (LP) / #EC5D (FH) / #EC6W (LP IBM i) / #EC6X (FH IBM i)) - 5 DWPD *PCIe3* x8 3.2 TB NVMe Flash Adapter
- #EC5E (LP) / #EC5F (FH) / #EC6Y (LP IBM i) / #EC6Z (FH IBM i)) - 5 DWPD *PCIe3* x8 6.4 TB NVMe Flash Adapter

Keine M.2 Module!

POW3R einfach. digital. agil.

POW3R einfach. digital. agil.

PRÄSENTIERT

NVMe Performance

POW3R einfach. digital. agil.

Is NVMe Fast?

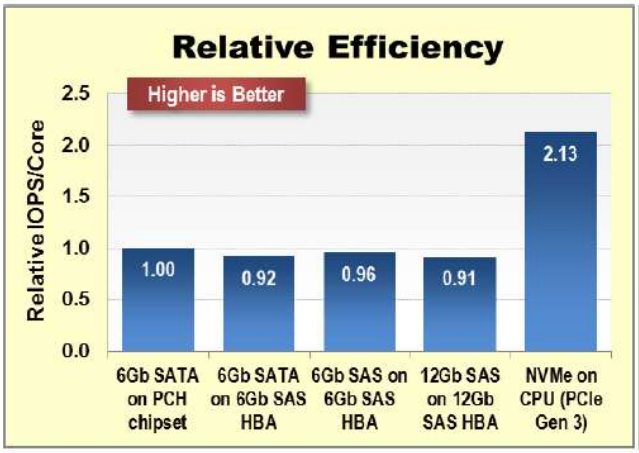

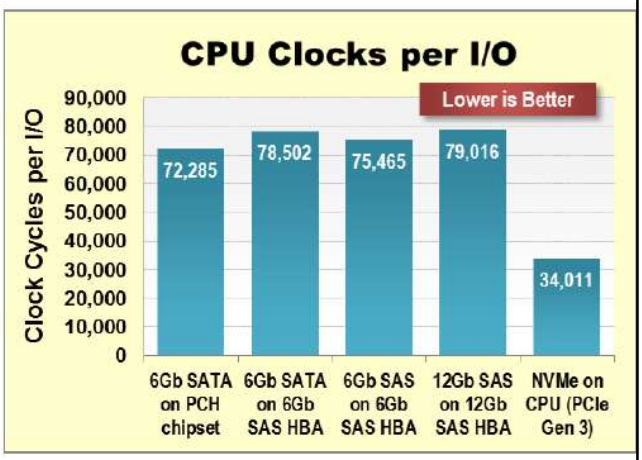

Vergleich NVMe vs. 15k Platten

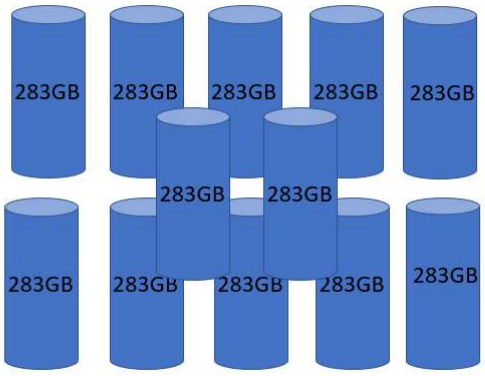

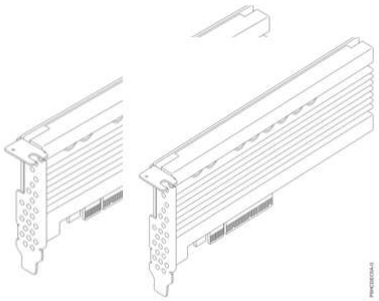

Ca. 2.400 IOPS

Ca. 320.000 IOPS

POW3R
einfach. digital. agil.

NVMe I/O Performance

Workload	Target (1.6 TB)	Target (3.2/6.4)
Read (IOPS)	**700K**	**910K**
Write (IOPS)	**100K**	**170K**
Mixed R/W (70/30)	**250K**	**320K**
Read Data Tp (GB/s)	**4.7**	**6.0**
Write Data Tp (GB/s)	**1.9**	**3.0**
Read Latency (us)	**110**	**110**
Write Latency (us)	**30**	**30**

POW3R einfach. digital. agil.

NVMe Performance Specs

FC	Capacity	Read BW GB/s	Write BW GB/s	Read IOPS	Write IOPS	DWPD (5 years)	PCIe Gen	Form Factor
EC5J, EC5X	800 GB	1.6	1.05	380K	60K	2.4	3	U.2
EC5G, EC5B EC6U, EC6V	1.6 TB	4.7	1.9	700K	100K	5	3	AIC
EC5C, EC5D EC6W, EC6X	3.2 TB	6.0	3.0	910K	170K	5	3	AIC
EC5E, EC5F EC6Y, EC6Z	6.4 TB	6.0	3.0	910K	170K	5	3	AIC
ES1E, ES1F	1.6 TB	7.0	2.5	1000K	220K	3	4	U.2
ES1G, ES1H	3.2 TB	7.4	3.4	1500K	250K	3	4	U.2
EC5V, EC5W	6.4 TB	7.2	3.8	1500K	250K	3	4	U.2
EC7A, EC7B EC7J, EC7K	1.6 TB	8.0	2.5	1000K	250K	3	4	AIC
EC7C, EC7D EC7L, EC7M	3.2 TB	8.0	3.8	1500K	250K	3	4	AIC
EC7E, EC7F EC7N, EC7P	6.4 TB	8.0	3.8	1500K	250K	3	4	AIC

NVMe vs SAS Performance Comparison (IBM i Workload)

RAID-5 vs. Mirrored | Mirrored vs. Mirrored

SAS vs NVMe Performance Comparisons

Protection	Parameter RAID Type	Units 0,1,5,6,10	SAS RAID 5	NVMe OS Mirror	Difference	SAS RAID 10	NVMe OS Mirror	Difference
Application Speed	Random 255B Reads	Trans/s	273,000	390,000	42.0%	276,185	387,284	40.2%
	Random 255B Writes	Trans/s	2,898,000	5,824,000	101.0%	6,674,419	5,785,859	-13.3%
	Sequential 32KB Reads	GB/s	3.51	8.22	134.2%	3.33	8.22	146.8%
	Sequential 32KB Writes	GB/s	1.01	2.46	143.6%	1.79	2.48	38.5%
	Random 4KB Reads	io/s	218,000	368,000	68.8%	217,629	385,410	77.1%
	Random 4KB Writes	io/s	167,000	690,000	313.2%	385,088	666,441	73.1%
Storage Subsystem	Sequential 32KB Reads	GB/s	3.48	8.05	131.3%	3.22	7.588	135.4%
	Sequential 32KB Writes	GB/s	1.00	4.56	356.0%	1.80	4.969	175.4%
	Random 4KB Reads	us/io	123	99	-19.5%	133	99	-25.6%
	Random 4KB Writes	us/io	59	24	-59.3%	34	28	-17.6%

Conclusions

- 2 NVMe devices consistently outperform 18 SAS SSDs
- Random 255B Writes see a slower NVMe performance because the Mirror code is offloaded to the CPUs instead of the SAS Controllers

Color Key	Meaning
Green	Advantage NVMe
Yellow	Within +/-10%
Red	Advantage SAS

Checking the amount of remaining life in NVMe devices

https://www.ibm.com/support/knowledgecenter/9223-42H/p9hak/pxhak_nvme_remaining_life_kickoff.htm

Learn how to use operating system commands to find the amount of remaining life in an NVMe device.

Using the operating system commands

The AIX®, IBM® i, and Linux® operating systems have commands that you can use to determine the amount of life that is remaining in an NVMe device.

- **Running the AIX command to check the amount of remaining life in NVMe devices**
  Find information about using the AIX operating system to find the amount of remaining life in an NVMe device.
- **Running the IBM i command to check the amount of remaining life in NVMe devices**
  Find information about using the IBM i operating system to find the amount of remaining life in an NVMe device.
- **Running the Linux smart-log command to check the amount of remaining life in NVMe devices**
  Find information about using the Linux operating system smart-log command to find the amount of remaining life in an NVMe device.

Rate this content

Parent topic:
→ Related procedures for installing and removing adapters

Parent topic:
→ Related procedures for installing and removing adapters

Parent topic:
→ Related procedures for installing adapters

POW3R einfach. digital. agil.

Checking the amount of remaining life in NVMe devices

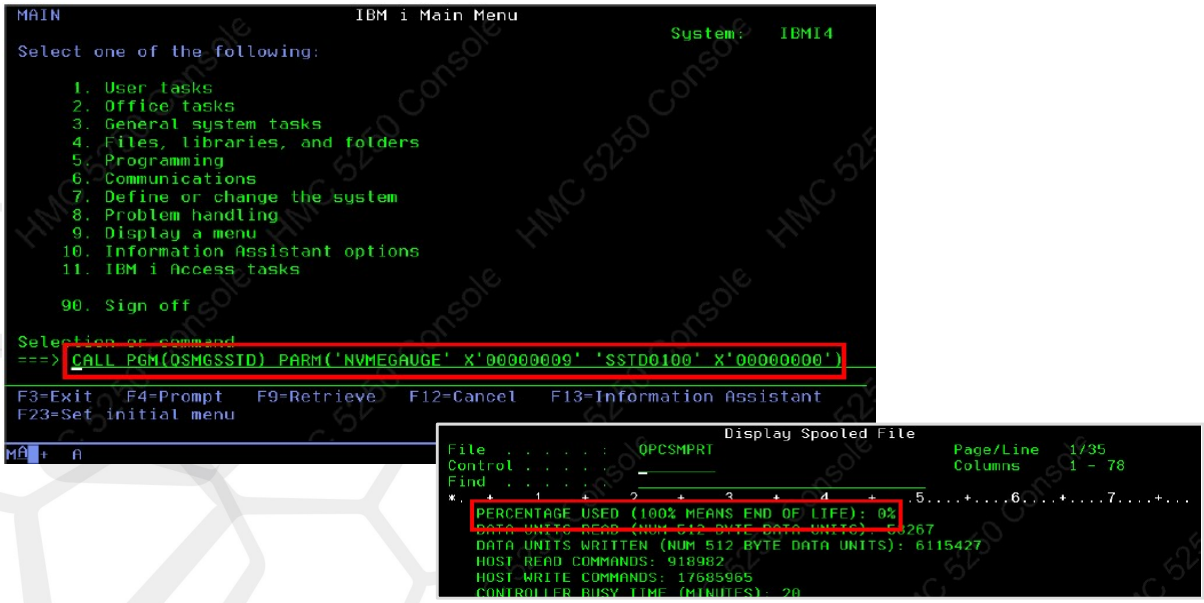

NVMe für IBM i?

POW3R
einfach. digital. agil.

Preis Leistungsverhältnis

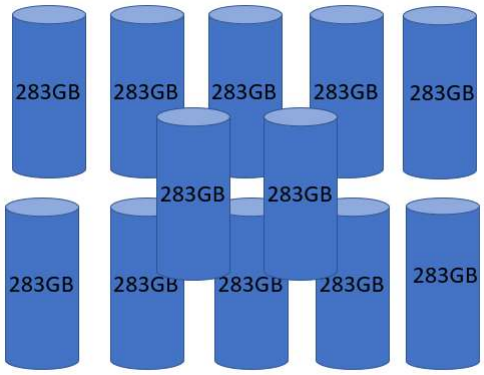

12 x Feature ESNJ 283GB Festplatte 15k = 3396GB
= ca. Liste 11.000€

Ca. 2.400 IOPS

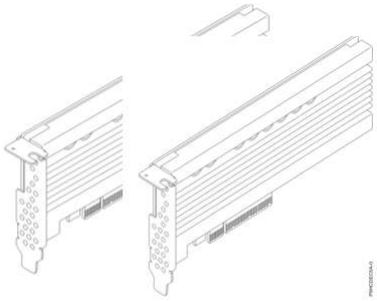

2 x Feature EC6X 3,2TB NVMe PCIe3 = 6,4 TB
= ca. Liste 7.600€

Ca. 320.000 IOPS

POW3R
einfach. digital. agil.

NVMe Namespaces als „Plattenarme“

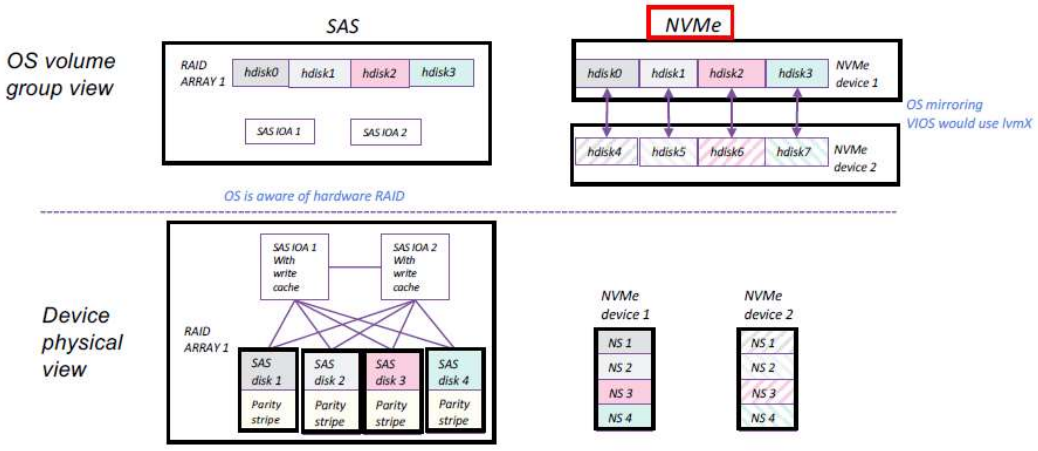

SAS Versus PCIe Card NVMe Storage with IBM i

IBM i PCIe Card NVMe Mirroring (3 NVMe, only one Option)

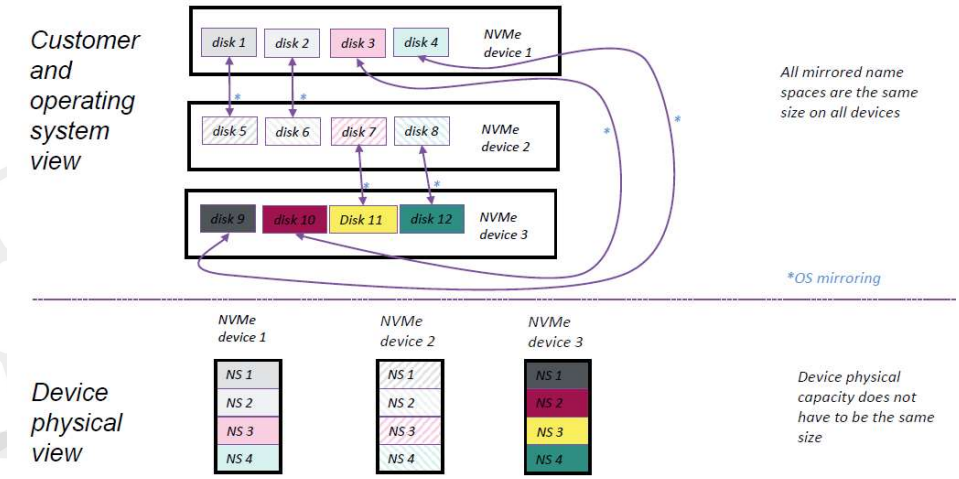

VIOS and NVMe - Single VIOS and 2 NVMe

- Create Namespace NS1 on NVMe1 which VIOS sees as hdisk0
- Create Namespace NS1 on NVMe2 which VIOS sees as hdisk4
- Add hdisk0 and hdisk4 to volume group nvmevg1 (optionally setup mirror in VIOS)
- Create logical volume nvmelv1
- Create vscsi device vnvme1 with nvmelv1 as the backing target
- Assign vnvme1 to vhost of IBM i guest partition

- You decide where to mirror (VIOS or OS) and the lv & vg configs

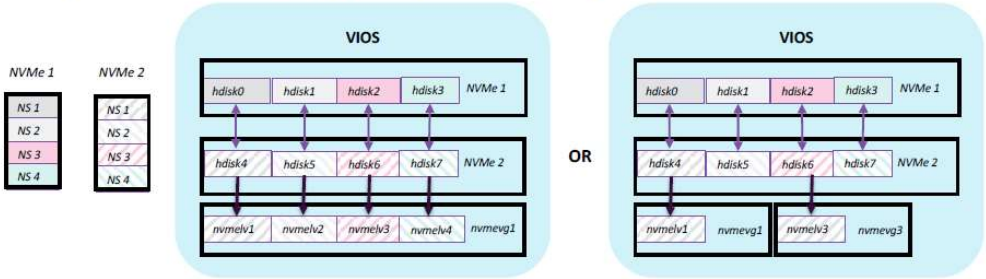

POWER
einfach. digital. agil.

VIOS and NVMe - Single VIOS and 2 NVMe

- Provides redundant devices and VIOS instances
- Like single VIOS create namespace, volume group logical volume on 1 NVMe device from each VIOS instances and assign both to 1 common IBM i partition
- Setup IBM i OS level mirroring

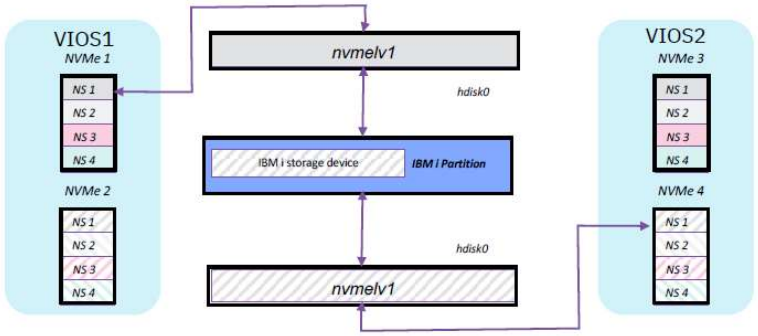

POW3R
einfach. digital. agil.

PCIe Card NVMe Namespaces and IBM i

- NVMe architecture can use namespaces which is a collection of logical blocks whose logical block addresses range from 0 to the size of the namespace. A namespace ID (NSID) is an identifier used by a controller to provide access to a namespace.
- With NVMe , an ‘arm’ (logical unit) is a namespace. A namespace is a logical chunk of a physical NVMe device and multiple namespaces are allowed on one NVMe device.
- IBM i is the management interface used by a customer to create and manage namespaces
- IBM i’s use of NVMe architected multiple namespaces provides for many ‘arms’ on a small number of high capacity NVMe physical devices
- IBM i can use a NVMe device (up to 16 TB) with only a single namespace for the whole device. However, for almost all customers, this will cause sub optimum performance since more (and smaller) ‘arms’ (logical units) are better than fewer and larger.
- Note that the word “namespace” is used in the industry and by IBM i in different ways and in different contexts, so for example, a NVMe namespace should not be confused with and has nothing to do with ASP namespaces.

Recommended PCIe Gen3 Card NVMe Namespace Sizes

First generation NVMe devices have a hardware boundary of 16 GB for name spaces. Device capacity can be wasted/lost if name spaces are not multiples of 16. The maximum number of namespaces on a device is 32.

IBM i screens show Capacities in 'GB' (1000**3 (GB), not 1024**3 (GiB))

IBM recommends namespaces of 188 GB or 393 GB (the e-Config default)

Consider only using 393 GB. On the 6.4 TB device there is a lot of unused space when using 188 because the maximum number of namespaces is 32. The choice of namespace size is a balance between number of "arms" for (storage management) performance and other individual customer factors such as scaling/growth of the system.

Device Nominal Size	Device Actual Size	Number of Namespaces	Namespace Size	Total User Capacity Used By Namespace	Remaining Space on the Device (unallocated)
1.6TB	1575	8	188	1448	87
3.2TB	3151	16	188	2977	174
6.4TB	6364	32	188	6016	348
1.6TB	1575	4	393	1556	19
3.2TB	3151	8	393	3112	38
6.4TB	6364	16	393	6288	76

POW3R einfach. digital. agil.

Recommended PCIe Gen4 NVMe Namespace Sizes

- New PCIe Gen4 NVMe devices have a hardware boundary of 8 KB for name spaces, and up to 64 name spaces per device
- IBM recommends namespaces of 200 GB or 400 GB (the e-Config default)
- U.2 NVMe and the new PCIe card NVMe devices are both Gen4 and use the same IBM i load source name space sizes:
  - #ENSA 200GB IBM i NVMe LOAD SOURCE NAME SPACE SIZE
  - #ENSB 400GB IBM i NVMe LOAD SOURCE NAME SPACE SIZE
- Note that PCIe Gen3 card NVMe can only mirror to PCIe PCIe Gen3 card NVMe
- U.2 NVMe and new PCIe card NVMe devices are both Gen4 and can mirror each other if using the same namespace sizes

POW3R einfach. digital. agil.

PCIe Card NVMe IBM i DST/SST Screenshots

IBM i and NVMe – Working with SQL

Thank you to **Scott Forstie @Forstie_IBMi** - [i202758] SQL for Every IBM i User

- The QSYS2.SYSDISKSTAT view contains information about disks
- The view is enhanced to recognize NVMe devices

```
select disk_model, disk_type,
       case when unit_type = 1 then 'SSD' else
       'Spinning'  end as "Type of disk",
       case when unit_nvme = 1 then 'NVMe' else 'Not NVMe'
       end as "NVMe indicator", percent_used from sysdiskstat;
```

- Mixed SSD and NVMe example

DISK_MODEL	DISK_TYPE	Type of disk	NVMe indicator	PERCENT_USED
0099	198E	Spinning	Not NVMe	7.848
0050	198C	Spinning	Not NVMe	0.006
0205	6B7D	SSD	NVMe	0.006
0205	6B7D	SSD	NVMe	0.006
0205	6B7D	SSD	NVMe	0.006

POW3R
einfach. digital. agil.

PRÄSENTIERT

NVMe Erfahrungen von ersten IBM i Installationen

POW3R einfach. digital. agil.

Rainer Ross

Performance Vergleich:

- Datei mit 150.000 Hotels mit Volltext indiziert mit OmniFind
- 1000 Zugriffe mit SQL unter ähnlichen Bedingungen auf beiden Maschinen

„Alte“ **Power7+** Maschine:

28,2 Sekunden für 1000 Zugriffe
= **0,028 s/Satz**

Neue **Power9** Maschine:

8,3 Sekunden für 1000 Zugriffe
= **0,008 s/Satz**

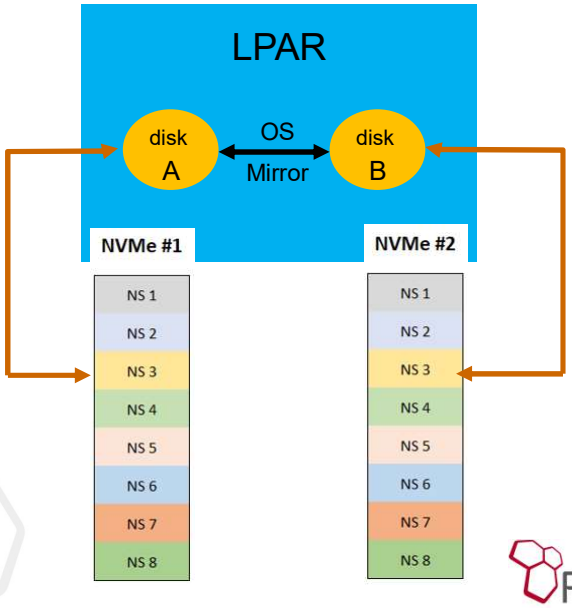

POW3R einfach. digital. agil.

NVMe Virtualisierung

Cognitive Systems

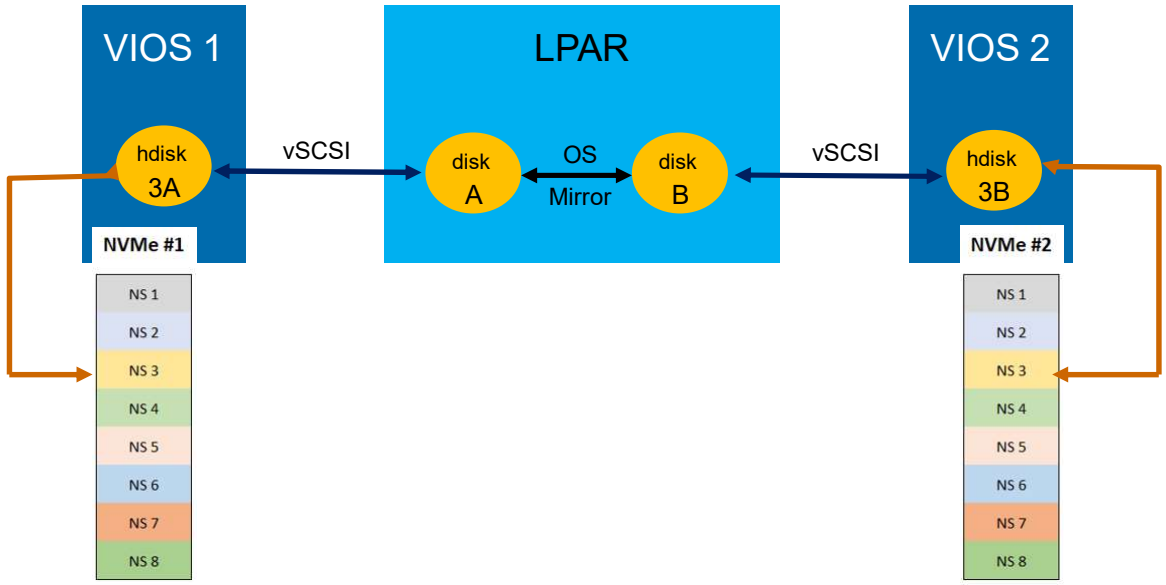

https://developer.ibm.com/technologies/systems/articles/au-aix-virtualization-nvme/

Anwenderbericht

Uwe Kohlberg
Maxim Quint

NAVO GmbH
Vorsprung durch Kompetenz

Hans-Böckler-Straße 21
65468 Trebur

tel:+49 (0)6147/44188-71
web: https://www.navo.de
mail: uwe.kohlberg@navo.de
mail: maxim.quint@navo.de

„**Wir haben in einer neuen IBM Power 9 9009-41A zwei NVMe Adapter mit jeweils 6,4TB Flash und 2x 300 GB Platte eingerichtet.** Auf den beiden 300 GB Platten wurde der VIOS V3 gespiegelt eingerichtet. Dadurch können alle anderen Adapter virtualisiert an die IBM i Partitionen konfiguriert werden. Anschliessend wurden über die vHMC die 3 Power i Partitionen konfiguriert, wo wir den Partitionen die definierten vscsi Platten konfiguriert haben (gespiegelt auf die beiden NVMe Adapter). Die Installation der verschiedenen Betriebssysteme erfolgte einfach und sehr schnell über das Netzwerk. Nachdem auch die aktuelle PTF Pakete installiert wurden erfolgten die erstens Test. Was hier bemerkenswert war ist das **neustarten einer Partition in ca. 15 sec**. So etwas hatten wir selbst bei sehr großen IBM Power System noch nicht gesehen. Auch die weiteren Test verliefen extrem schnell. Durch die Flexibilität der Ressourcenzuweisung z.B. Prozessor, Hauptspeicher oder auch des internen RDX Bandlaufwerkes ist das System gerade für kleine Unternehmen oder Softwareentwickler sehr interessant. Die Leistungsdaten der neuen IBM NVMe Adapter ist hier schon mehr wie bemerkenswert, eine echte Alternative für interne Platten.“

„Würde auf jeden Fall separate Platten bestellen, um VIOS OS von den LPAR Daten zu trennen (Thema: VIOS Update etc….)“

POW3R
einfach. digital. agil.

Fragen / Kontakt

POW3R
einfach. digital. agil.

ENDE/PAUSE