

Université IBM i 2017

17 et 18 mai – IBM Client Center de Bois-Colombes

S48 – IBM i et stockage SSD/Flash : retour d'expérience

Jeudi 18 mai – 15h15-16h45

Ludovic Ménard – IBM IT Specialist IBM i





IBM i - Montpellier Design Center Offering


Contact us

A Center of Competency committed to your success:

- ❑ Personalized workshops, all focused on **IBM i solutions**, including Live demonstrations, design sessions, testing & hands-on capabilities.
- ❑ Experience & Expertise from the Benchmark & Design Center.
- ❑ Innovation & hot topics around Digital Transformation: IBM i & Cloud, Analytics, Mobility... Always be at the edge of IBM latest technology.
- ❑ Focal point for ISVs and IBM Partners on IBM i.

Utilization terms

- ❑ In Center or on your site, for a design session or a large event. Contact us now!



Contact:

Benoit Marolleau – Power Systems Design Center

Benoit.marolleau@fr.ibm.com

+33 4 9958 0614 (Office)





Use IBM i Expertise from Montpellier

Bring your clients & convince them!



IBM i SYSTEM MANAGEMENT & PERFORMANCE




Topics covered look at system management & performance using iDoctor and Database monitoring, Java & Websphere best practices and System management solutions.

POWER SYSTEMS ROADMAP & INFRASTRUCTURE



Topics covered highlight the Power Systems technology and IBM i roadmap, Storage & Flash, and High Availability Solutions with IBM i.



DIGITAL TRANSFORMATION



Digital Transformation with IBM i, covering Hybrid Cloud, Analytics with Web Query & Cognos, Mobile Solutions on i based on MobileFirst & ISV solutions, New developments tools & modernization. Live demonstrations

VIRTUALIZATION & CONSOLIDATION



Topics covered address Power Virtualization with IBM i, PowerVC, VIOS Best Practices, Integration with Linux, iVirtualization. Live Demonstrations

CUSTOMIZABLE AGENDA



Build your own Agenda with topics taken from this 4 themes.



IBM i Modernization Workshops & Briefings

Digital Transformation / Modernization

Compose your Agenda!

- **Bluemix** for Developer productivity & Innovation - What is bluemix? includes a demo and/or lab
- **Mobility & Hybrid Cloud** : Mobile Applications on IBM i - Demonstrations with MobileFirst, IBM i and Bluemix
- Retail Demo - IBM i & Hybrid **Cloud** - Demonstrations with Node.js, Bluemix, Watson Services & IoT...
- **API Management & Integration Studio** - IBM i & Hybrid **Cloud** – Design , Create and Manage your API on IBM i.
- **Development** & modernization using RPG Free, Node.js & new languages, all on IBM i!
- **Security**: IBM **QRadar** & IBM i Security & Compliance solution (ISV Solution) - Includes a demo
- Analytics & Security: **Fraud Detection** Solution on Power Systems (AIX/Linux/IBM i) with IBM **ODM** , **Analytics** and ISV Solution. Includes a demo
- Power Systems & **Cloud Solutions** (private, public, hybrid) - Includes demos

Infrastructure Solutions

- Power Systems **Strategy**
- IBM i **Roadmap** , New Products & Beta Testing
- IBM i & **Storage** - Which Storage to choose? Feedback from Power/Storage Benchmark/Design Center
- **OpenPower**, Linux on Power, **PowerKVM**
- Solutions **Architecture & Virtualization** best practices: PowerVM & VIOS, iVirtualization, PowerVC , Live Partition Mobility.
- Storage & PowerHA with IBM i: SAN Storage & Flash Systems, IASP, HyperSwap...
- System **Performance** with iDoctor & Database monitoring, Java & WebSphere on IBM i Best Practices, System security & administration.
- **Analytics** Solutions with IBM i using IBM Cognos & DB2 Web Query
- Total Cost of Ownership – Economics Study



IBM i Benchmark center activities

IBM Systems Benchmark Centers have the expertise and tools to provide

- Proofs of concept (PoCs)
 - Demonstrate functions (flashcopy, brms)
- Proofs of technology (PoTs)
 - Demonstrate new hardware
- Performance benchmarks
 - Stress, tune, test workload and applications
- Pre-sales activities



Benchmark center portal

<http://www-03.ibm.com/systems/services/benchmarkcenter/>



2 Benchmarks SSD

2 Benchmarks Flash

Storage comparative



What is the problem and Why a Bench ?

Application changes

Cost reduction

Batch time reduction

Backup window

Business growth



What is the problem and Why a Bench ?

Cost reduction

Power 8 price

Storage price

Backup window

Flashcopy

BRMS

Application changes

RPG free form

New languages

SQL

Batch time reduction

CPU

Response time

Business growth

Consolidation of infrastructure

performance



Hardware and Software used





EJ14 card

- PCIe3 12 GB Cache RAID PLUS SAS Adapter Quad-port 6 Gb x8 (FC EJ14; CCIN 57B1); Adapter FRU number 01DH742PCIe3 x8
- Transfer speed of 6 Gbps
- 12 GB write cache
- Up to 1.6M read IOPS, Up to 360k write IOPS, Up to 878k mixed IOPs (70-30)
- One PCIe x8 slot per adapter
- Adapters are installed in pairs
- OS support: AIX, IBM i, and Linux operating systems
- VIOS supported

IBM i 7.3

IBM i 7.2 Technology Refresh 4

IBM i 7.1 Resave RS-710-S



EXP24SX or ESLS



- Up to 12Gb throughput
- 24 2.5-inch SFF SAS bays
- Designed to support future 12Gb SAS interfaces

| HDD 10k RPM | HDD 15k RPM | SSD |
|------------------------------|------------------------------|---------------------------------|
| 600/571 GB 4k (#ESEV, #ESEU) | 300/283 GB 4k (#ESEZ, #ESFA) | 387 GB 4k (#ES85, #ES86) |
| 1.2/1.1 TB 4k (#ESF3, #ESF2) | 600/571 GB 4k (#ESFP, #ESFN) | 775GB 4k (#ES8C, #ES8D) |
| 1.8/1.7 TB 4k (#ESFT, #ESFS) | | 1.55 TB 4k (#ES8F, ES8G) |
| | | 1.9 TB 4k RI (#ES80, #ES81) |

IBM i 7.3 Technology Refresh 1
IBM i 7.2 Technology Refresh 5

V9000



Encryption
Compression

2 * 6 ports 16GB used
Software level 7.8.1.0



IBM i 7.1 Technology Refresh 8
IBM i 7.2
IBM i 7.3

Power 8



S824 – 8286-42A
16 cores
512 GB MEM

E880 – 9119-MHE
192 cores
16 TB MEM



IBM i 7.1 Technology Refresh 8
IBM i 7.2
IBM i 7.3



iDoctor for IBM i

- ❑ Suite of performance tools
- ❑ Investigate and analyse performance data
- ❑ Monitor overall system health




Case N°1

EJ14 – EXP24SX – 17 SSD


Power 8 S824

IBM i 7.2





Client configuration



- **Power**

- Power 720
- 6 cores 3,6GHz
- 256GB memory

- **Storage**

- 72 disks HDD 139 GB 15K Raid5 protection
- 4 x 5887 disks drawer EXP24S
- 4 x PCI gen2 1.8GB Cache RAID SAS Adapter Tri-port 6Gb
- ERP M3 Application response time : 1ms to 4ms

- **IBMi partition**

- IBM i V6.1
- 4 cores
- 249 GB memory
- 68 disks HDD 139 Go 15K Raid5 protection
- Disk capacity available 8,7 TB, 70% used



Objectives




The customer's objective is to test the performance of a set of batch jobs of the M3 version 10.1 ERP in an IBM i V7.2.

Power System S824 and internal SSD disks, similar to that envisaged to replace the client infrastructure.



Benchmark Results - System activity



Batch time : 60 minutes vs 2h - 53% gain


Journaling

Dispatched CPU





Benchmark Results - Disk activity



- Writes per second
- Reads per second
- Average response time
- Average service time





Conclusion

- Reduce number of disk
 - 68 disks HDD to 17 disks SSD
- IBM i V6R1 to V7R2 M3 validate
- Same number of cores
 - 4 cores
 - Batch time divides per 2 (53% gain)
 - ERP M3 Interactive query divides per 10
 - LTO 6 Backup window reduction from 4h to 1h40
- Response time
 - Close to 0,1 ms
 - divides per 10
- Customer hardware proposition
 - S814 - 8 cores 'internal 24 SSD' buy.
 - It's on hold





Case N°2

EJ14 – EXP24SX – SSD

End of Month Batch Job time between 8h-9h

Benchmark time 4h32



Client configuration

- **Power**

- Power 720
- 6 cores 3 GHz
- 80 GB memory

- **Storage**

- 39 disks HDD 139 Go 15K Raid5 protection
- Response time : 4ms to 10ms

- **IBMi partitions**

- IBM i V7.1
 - 1,5 cores
 - 64 GB memory
 - 32 disks HDD 139 GB 15K Raid5 protection
 - Disk capacity available 4 TB, 68% used
-
- IBM i V7.1
 - 0;25 core
 - 13 GB memory
 - 7 disks HDD 139 GB 15K Raid5 protection
 - Disk capacity available 837 GB, 81% used





Objectives

The client is currently migrating data from the second partition to the first partition. In the end, there should be only one partition.

IBM proposed to client a new IBM POWER System :

- S814 - 2 active cores on 6-core 3.02 GHz POWER8)
- 128 GB MEM
- 2 Cards EJ14 in pair (12Gb cache)
- 11 disks of 775 GB SSD*

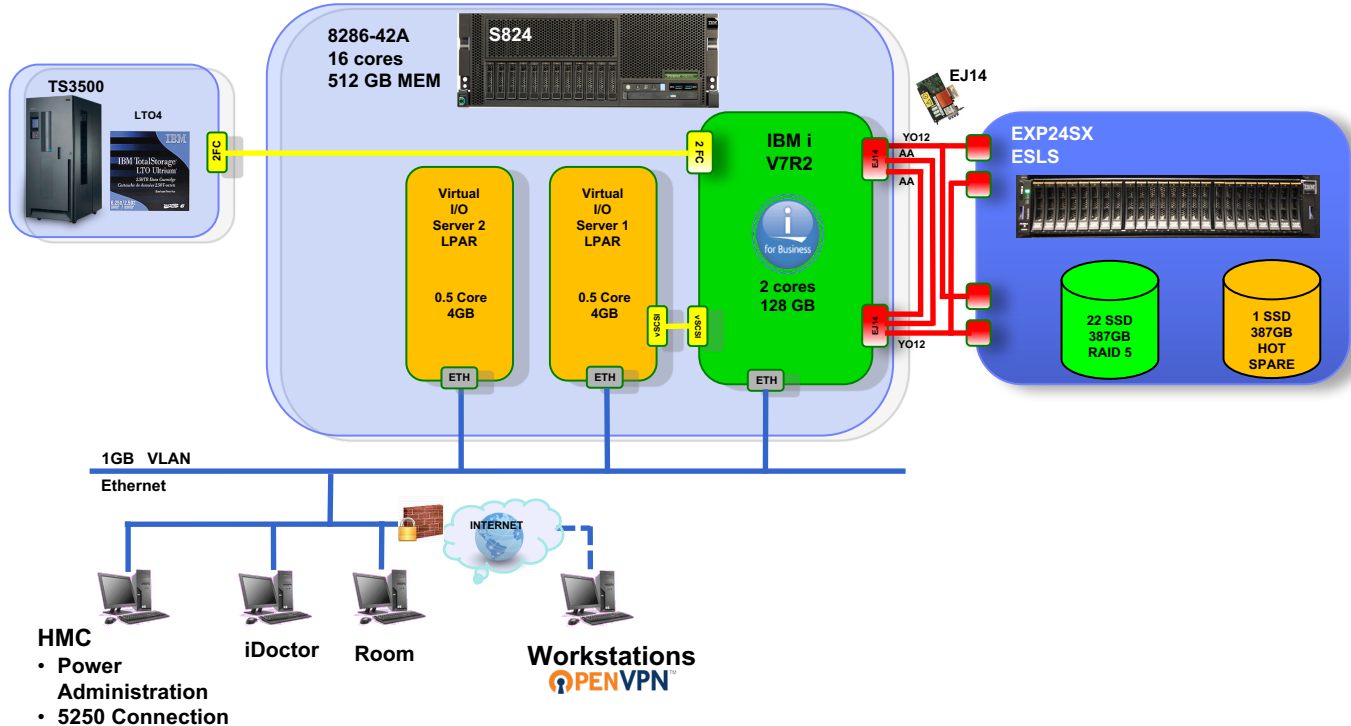


**Client wants the batch duration divided by at least 2.
The original batch duration is 9h.**

* 22 disks 387 GB used




Benchmark Infrastructure





Benchmark Results - System activity

Collection overview time signature



Interval end date and time (10 minute intervals)

- Journaling
- Dispatched CPU


- Average partition CPU utilization
- Maximum partition CPU utilization
- Average batch CPU utilization
- CPU power-savings rate(Scaled CPU: Nominal CPU)
- VCPU delays a percentage of Dispatched CPU





Benchmark Results - Disk activity

Read/write rates for ASP 1



Writes per second

Reads per second

Average read response time (ms)

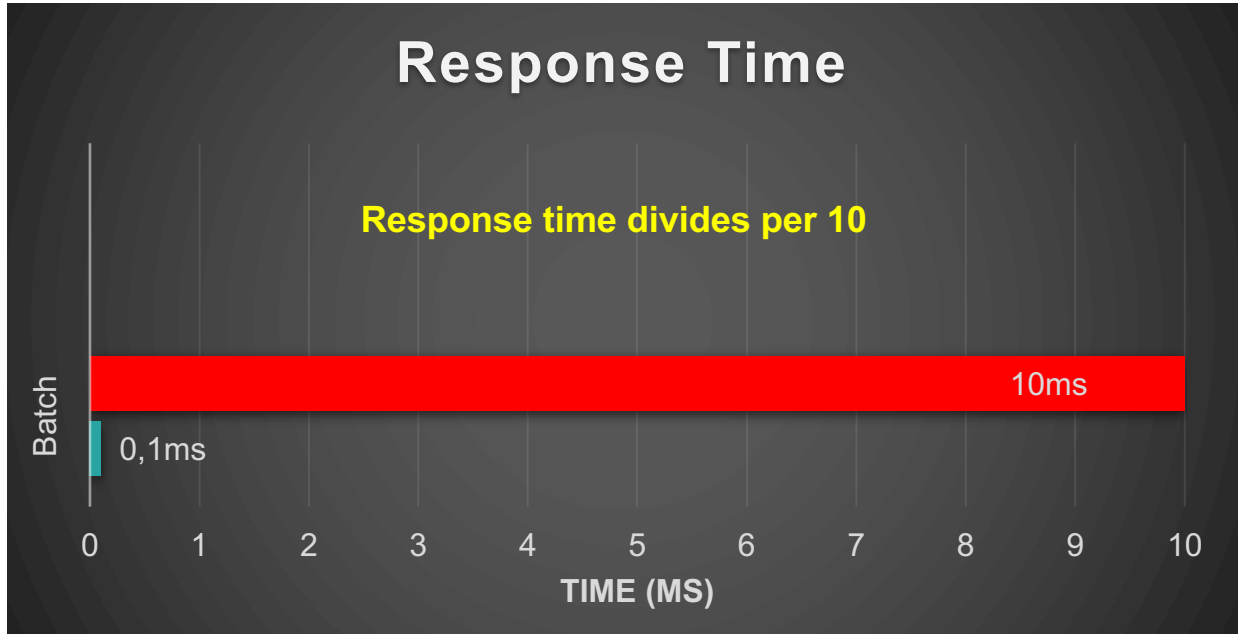
Average service time (ms)

Average response time (ms)

Average write response time (ms)

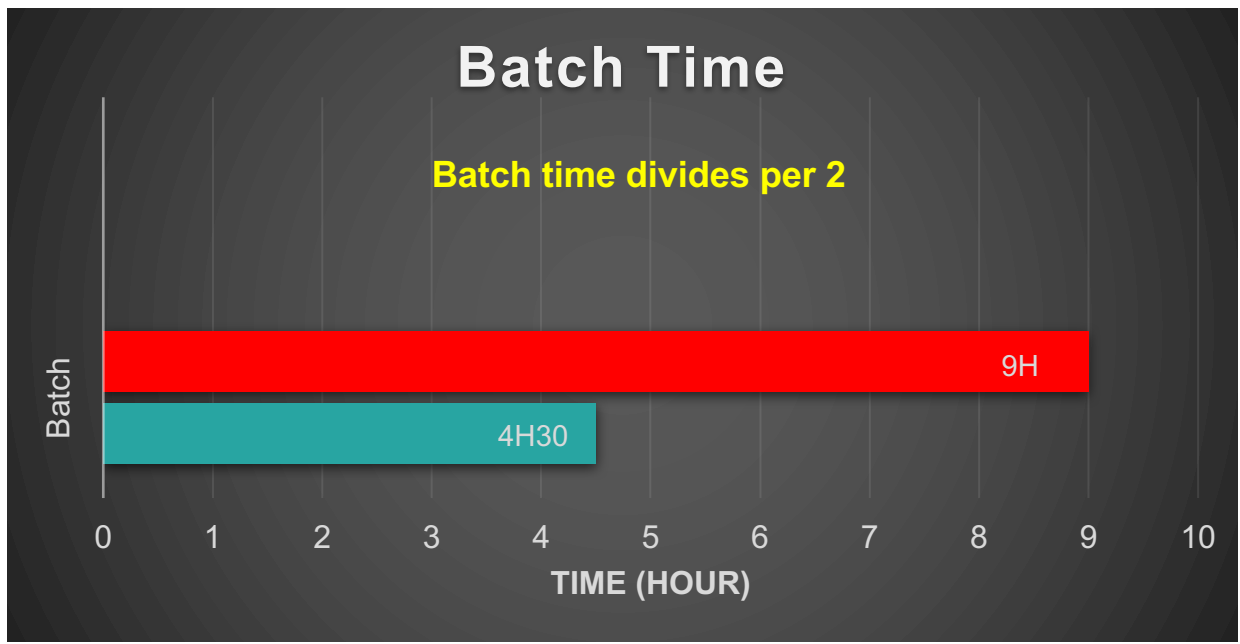


Conclusion



Project is hold on

Conclusion




Project is hold on

Case N°3

V9000 FlashSystem

VIOS - NPIV

IBM i V7R2





Client configuration

- **Power**
 - Power 740
 - 6 cores 4,2 GHz
 - 64 GB memory
- **Storage**
 - Storwize V7000 Gen 1
 - 84 disks HDD 146 GB 15K Raid5 protection
- **IBMi partitions**
 - IBM i V7.2
 - 2 cores
 - 20 GB memory
 - 32 LUNs 139 GB
 - Disk capacity available 2,8TB, 75% used
 - Redundant VIOS vSCSI connection





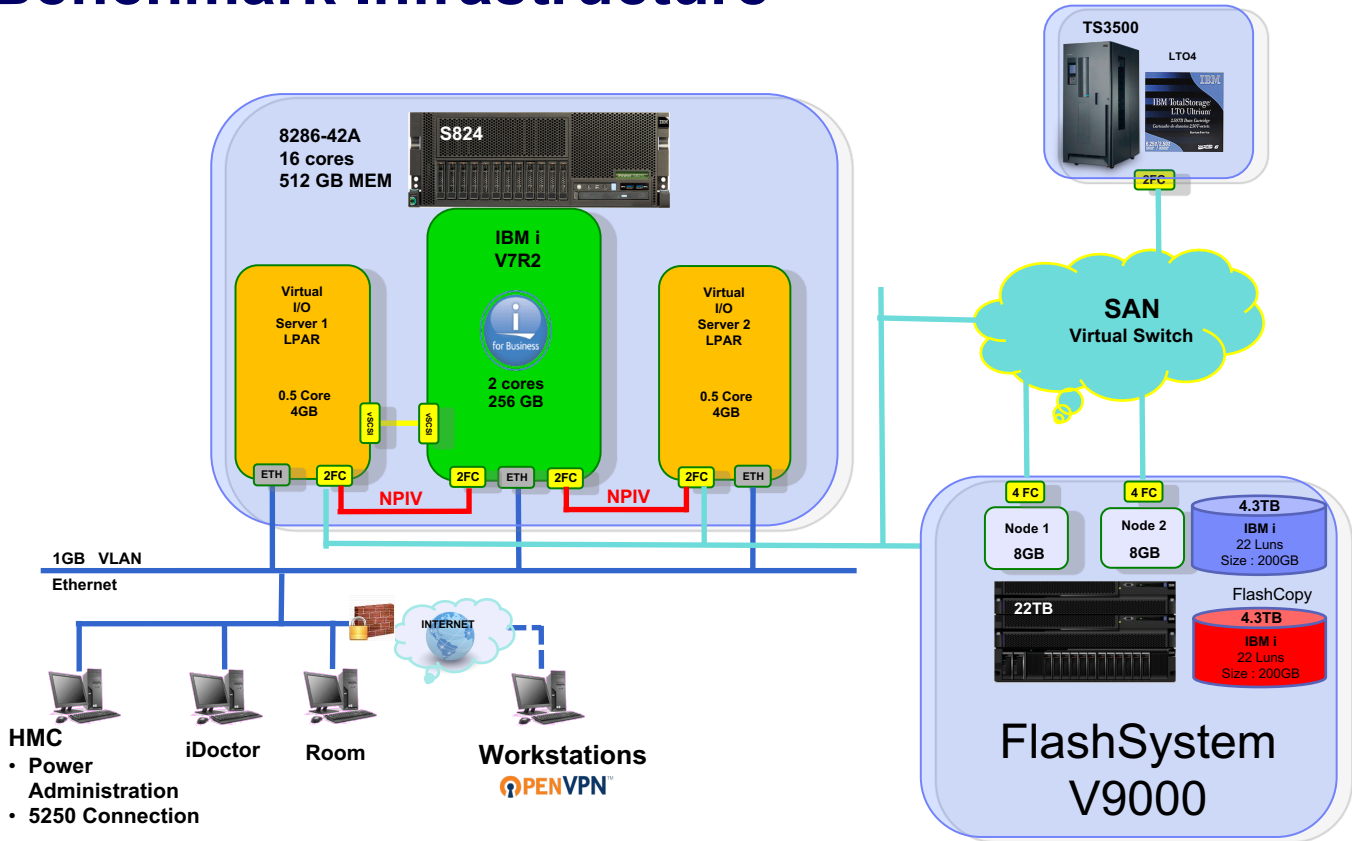
Objectives



- The objective of this benchmark is to test the performance of a batch job in an IBM i V7.2
- Power System S824 and FlashSystem V9000, similar to that envisaged in replacement of the existing infrastructure
- This batch job populates a Business Intelligence (BI) database from IBM i DB2.




Benchmark Infrastructure





TEST 1 : system activity




- Disk non fault reads
- Dispatched CPU
- Average partition CPU utilization
- Maximum partition CPU utilization
- Average batch CPU utilization
- CPU power-savings rate(Scaled CPU: Nominal CPU)
- VCPU delays a percentage of Dispatched CPU





TEST 1 : disk activity

Read/write size totals for ASP 1



Writes per second
Reads per second


Average read response time (ms)
Average service time (ms)
Average response time (ms)
Average write response time (ms)





TEST 2 : system activity

Collection overview time signature



- Disk non fault reads
- Dispatched CPU

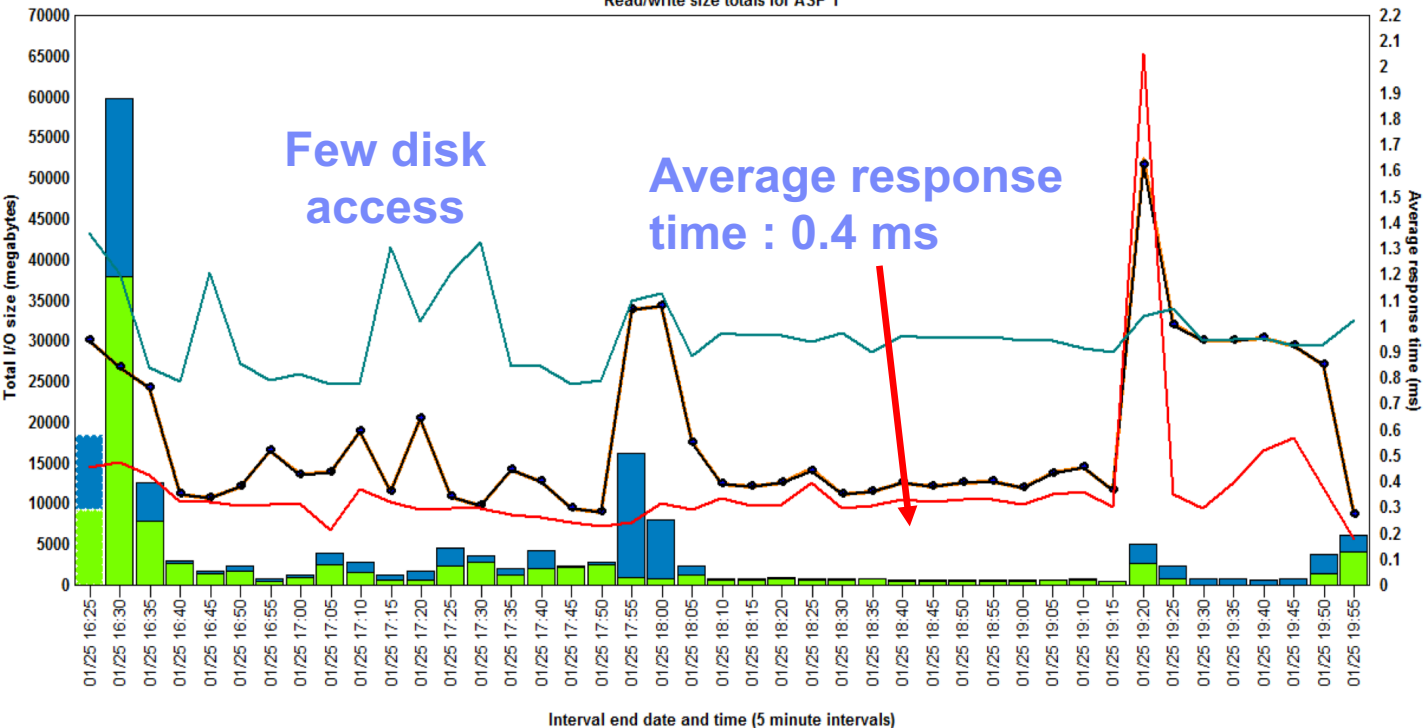
- Average partition CPU utilization
- Maximum partition CPU utilization
- Average batch CPU utilization
- CPU power-savings rate(Scaled CPU; Nominal CPU)
- VCPU delays a percentage of Dispatched CPU





TEST 2 : disk activity

Read/write size totals for ASP 1



Writes per second
Reads per second

Average read response time (ms)
Average service time (ms)
Average response time (ms)
Average write response time (ms)





Conclusion

- The Power System S824 and FlashSystem V9000 architecture improve BI batch run time by 18%.
- The first test shows a CPU usage of less than 30%, a memory usage of up to 50 GB, and moderate disk activity with an average response time of 0.5 ms.
- The second test (activation of parallelization of DB2 queries) gives similar results, which means that this function does not improve (the ONPQRYF function is only slightly used).
- The performance of this batch processing can not be improved beyond this result without studying an optimization of the code.

A Storwize V7000 Gen2 full SSD would provide an equivalent performance gain compared to this batch process.



Case N°4

FS900 – SVC - VIOS - NPIV

Encrypted Compressed

Encrypted Uncompressed

IBM i V7R1



Client configuration

• Power

- 4 Power 770
 - 2 systems with 48 cores and 1TB of memory
 - 2 systems with 48 cores and 1,25TB of memory
- 97 partitions – production
- 71 partitions – non production and DR

• Storage

- 2 DS8880 with 2 expansions by DS8880
- Prod DS8880
 - 608 disks 146GB 15K
 - 62,8 TB usable
- Non prod DS8880
 - 1056 disks 146/300GB 15K
 - 126 TB usable

• IBMi

- 2 x Partitions @ 3TB Disk, 4GB Memory and 0.33 CPU (1 virtual) uncapped
- 2 x Partitions @ 1.5TB Disk, 13GB Memory and 0.21 CPU (1 virtual) uncapped





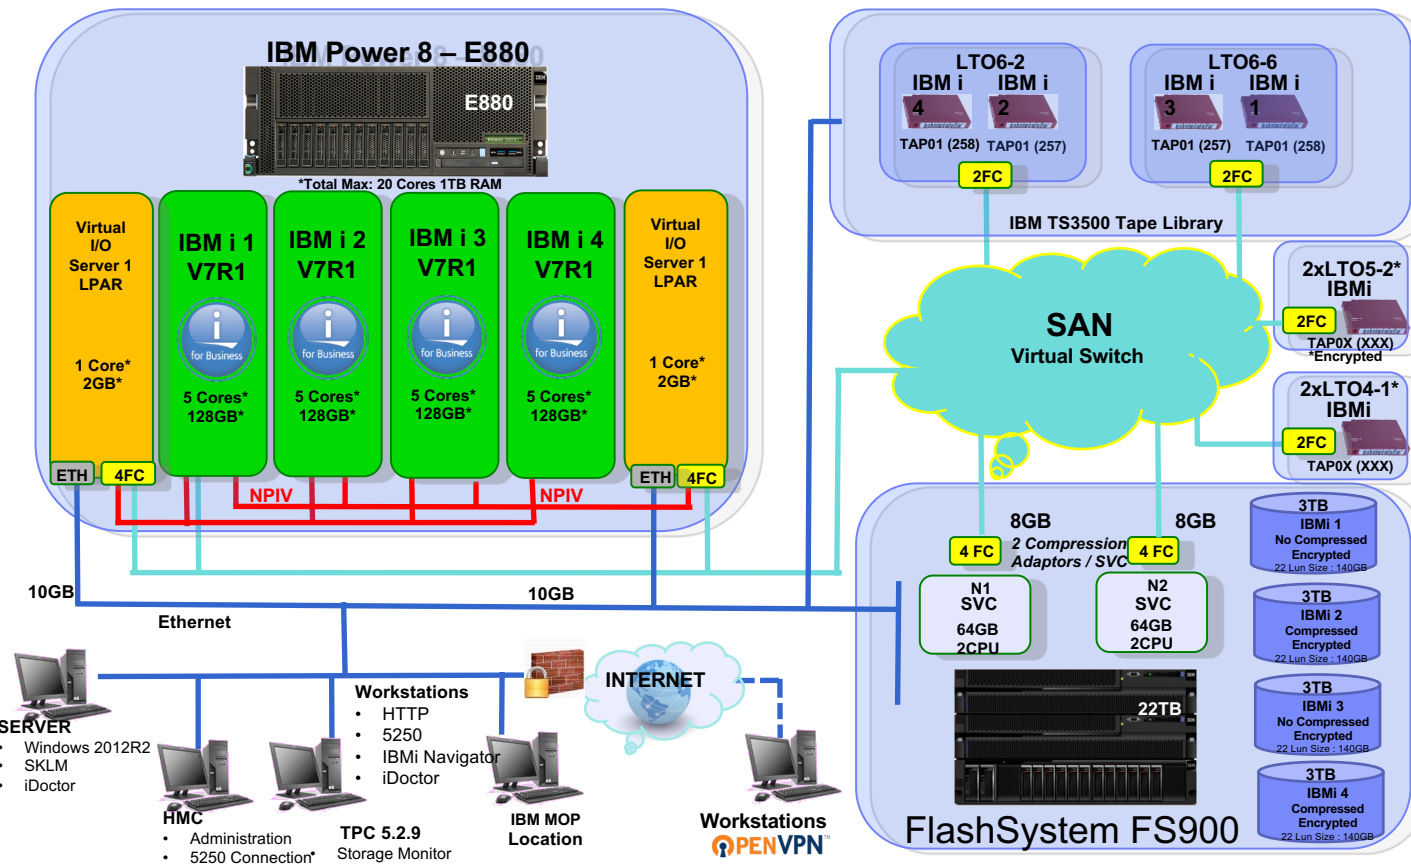
Objectives

- Client wants to drive the cost per TB down as it currently is just over 10k EUR per TB usable using 15k RAID5 ranks on IBM DS8800 with same performance.
- Compare performances data running batchs Power 7 and DS8880 versus Power 8 and FlashSystem.
- IBM i environments in the following configurations:
 - Non-compressed but encrypted disk with encrypted tape drives
 - Compressed and encrypted disk with encrypted tape drives





Benchmark Infrastructure




FlashSystem Hardware Encryption Activated





Benchmark Results - System activity



**IBM i
uncompressed**


- Average partition CPU utilization
- Maximum partition CPU utilization
- Average batch CPU utilization
- - - CPU power-savings rate(Scaled CPU; Nominal CPU)
- VCPU delays a percentage of Dispatched CPU

- Disk writes
- Dispatched CPU
- Disk page faults

**IBM i
compressed**

70 % disk space saving


Δ 19 minutes





Benchmark Results - Disk activity

Read/write rates for ASP 1




**IBM i
uncompressed**

Writes per second
Reads per second

**IBM i
compressed**

70 % disk space saving


Read/write rates for ASP 1





Storage Overall I/O rate & Latency : TPC

**IBM i
uncompressed**




**IBM i
compressed**
70 % disk space saving



Conclusion

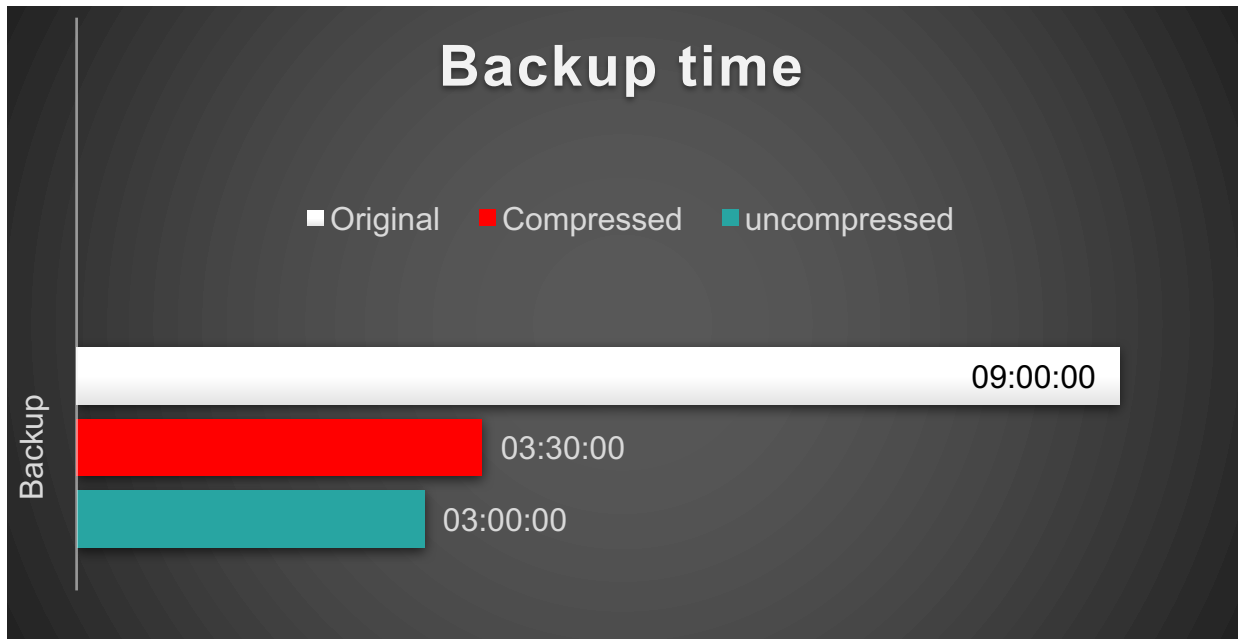
- **Original Daily batch run : 2h**
 - compressed : **1h41**
 - uncompressed : **1h22**
- **Original Weekly backup : 9h**
 - DS8K/LTO3 vs Flash/LTO6
 - compressed : **3h30**
 - uncompressed : **3h**

Conclusion



Project is hold on

Conclusion



Project is hold on



COMPARATIVE STORAGE






Comparative storage

Description of the workload

- **Writefile (WF):** RPG programs that write to 24 files at the same time, 70 million records to each file in 24 jobs. The record in each file is 75 packed decimal. OVRDBF with FRCRATIO(1000) is used for each file.
- **Readfile (RF):** RPG programs that sequentially read the files with 70 million records that were written with WF in 24 jobs. Each file is read 16 times. OVRDBF with SEQONLY(*YES 2000) is used on each file.

Power 8 Configuration

VIOS



OLERON

Data Collection 1

Data Collected: 05-11-2017, 15:15:00 PM

Processor Usage / Peak

4%

Installed: 16.00 Average: 0.58
Activated: 16.00 Peak: 5.08
Available: 8.00

Memory Allocated / Available

52%

Installed: 512.0 Average: 264.0
Activated: 512.0 Peak: 264.0
Available: 235.0

Network I/O Usage

0%

Sent/Received: 0.000 TB/s
Sent/Received: 0.000 packets/s

Storage I/O Usage

35%

Written: 12,387 KB/s
Read: 0.067 KB/s

[View Performance Dashboard](#)

i | ↻

Operating


⚠ Attention LED

2%

52%

0%

36%




D-OLERON-VIOS1

i | ↻

⏻ Not activated

0.50 PU 4 VP

4.00 GB Allocated



D-OLERON-VIOS2

i | ↻

⏻ Running


0.4%

4 GB Allocated

12%

9%

IBM i partitions



DIRECT_IBMi

i | ↻

⏻ Running


⚠ Attention LED

9%

128 GB Allocated

0%

0%




EJ14-IBMi

i | ↻

⏻ Not activated

0.00 Processors Allocated

0.00 GB Allocated



VIOS_IBMi



i | ↻

⏻ Not activated

1.00 Processors Allocated

128.00 GB Allocated

Storage Configuration

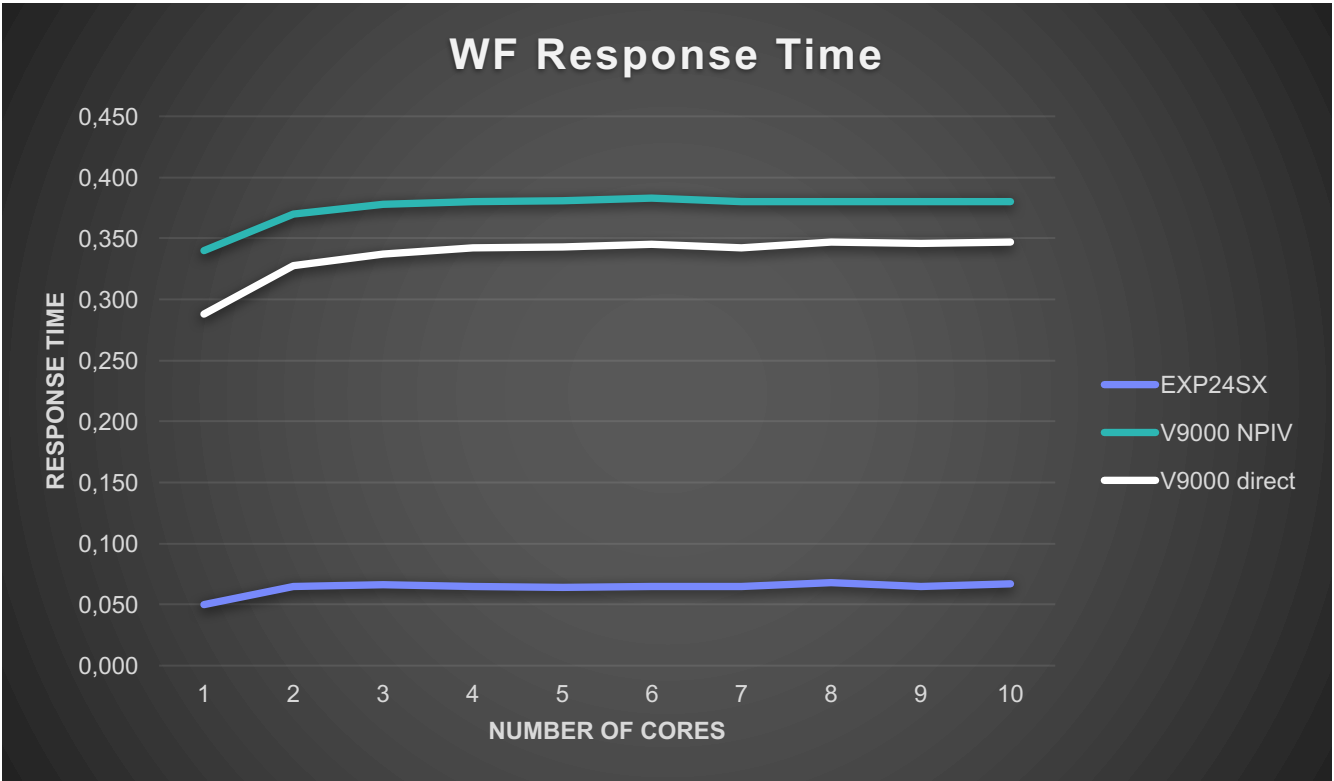


 NPIV

 Direct Attach

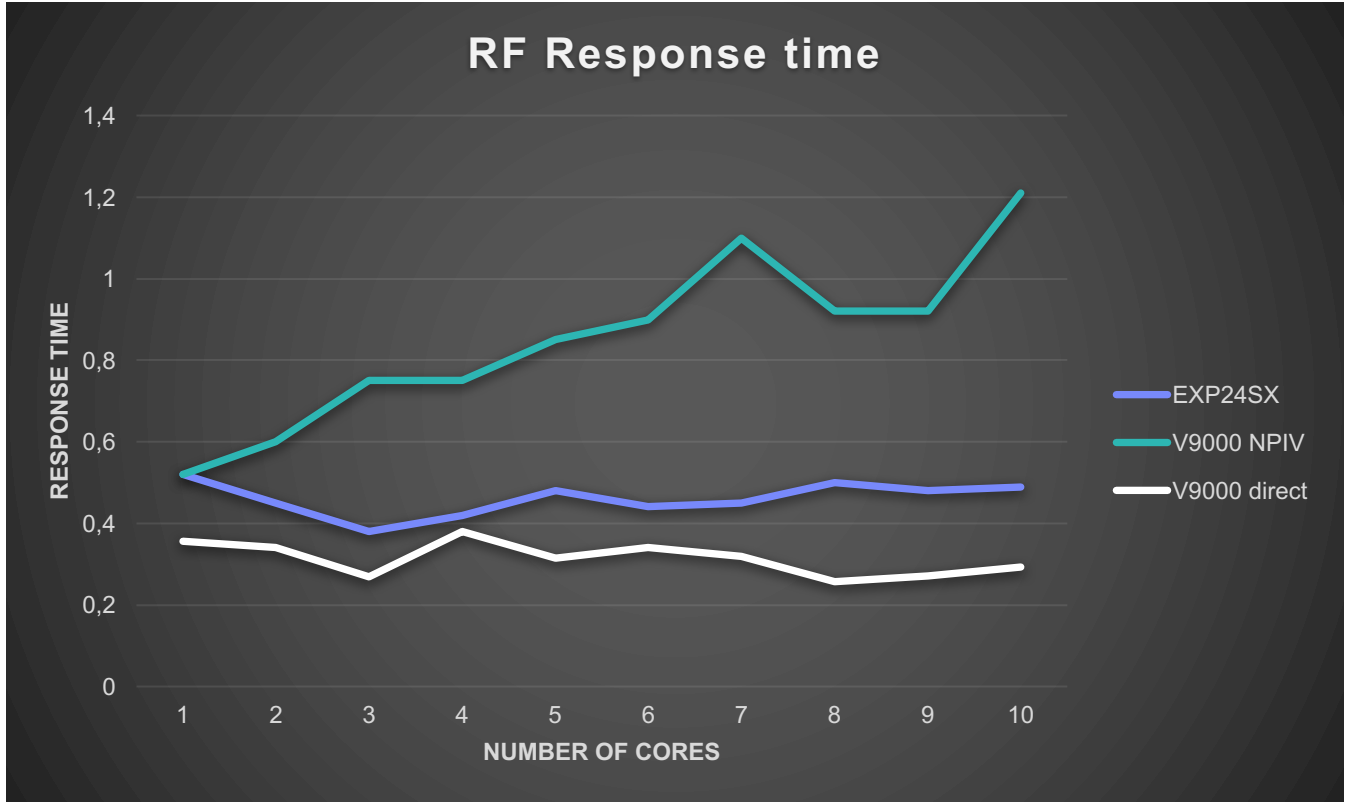


WF Response Time

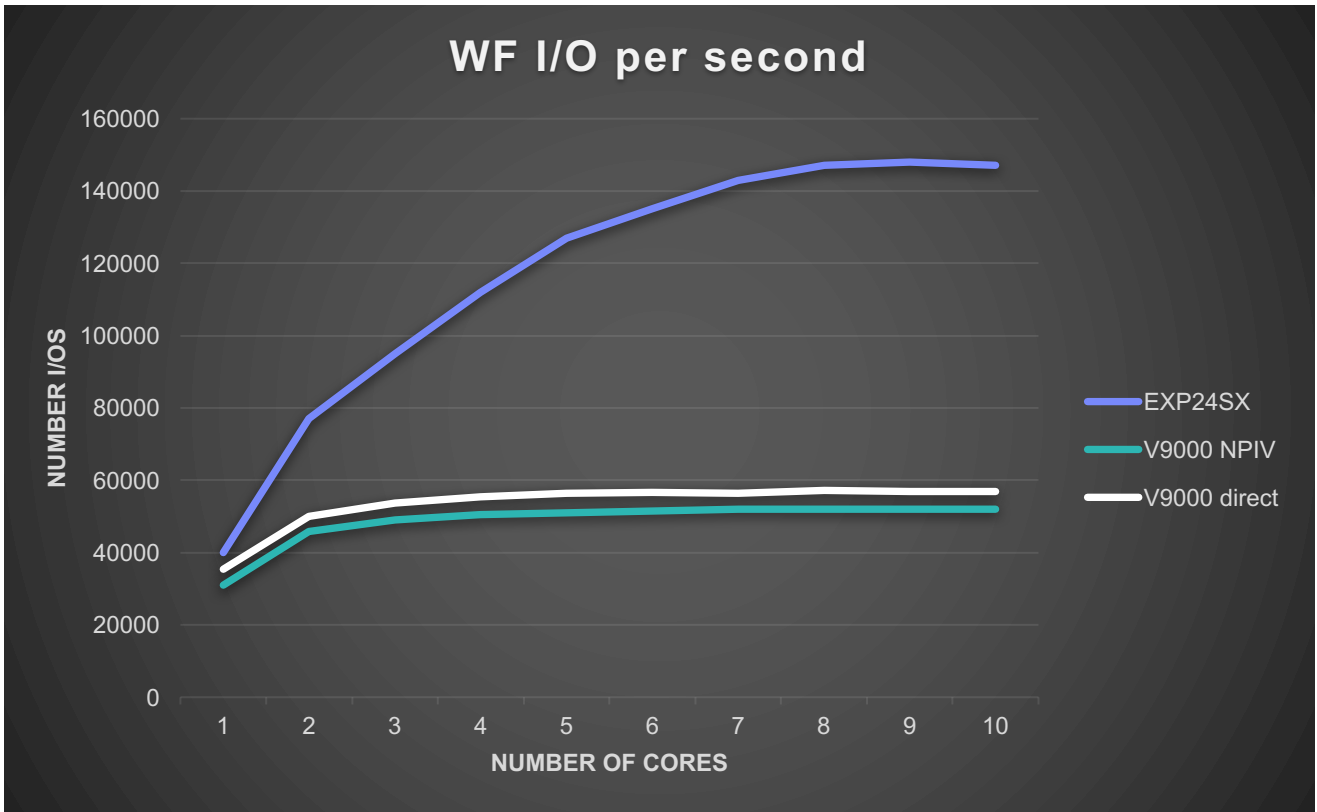




RF Response Time




WF I/O per second





RF I/O per second






The key point is

- EJ14/EXP24SX + SSD Best performance
- Flashsystem and EJ14/EXP24SX Read intensive same performance
- External storage : FlashCopy, Windows reduction, best way to manage LUNs

- IBM i workload : 80-20 (80% writes – 20% reads)
- It depends on the workload



■ Questions ?

