






2 Mars 2017


Agenda

- ✓ L'IBM i à travers les siècles (20^{eme} et 21^{eme}... !!)
- ✓ Technologie POWER... du silicium à la virtualisation
- ✓ Virtualisation dans un contexte IBM i..
- ✓ POWER8 Architecture et performances
- ✓ Stockage interne ou externe: les critères de choix dans un contexte de performance, de fonctionnalités et de haute disponibilité
- ✓ Les sauvegardes... bande ou pas bande ??

Un peu d'histoire...



Un peu plus de 35 ans d'histoire ...



S/38
1978



Dr F.G Soltis



AS/400
1988



Series/1
1976



S/36
1983



64 Bit
1995



POWER8
2014



POWER7
2010



POWER5
2004




POWER6
2007

Architecture intégrée




Microsoft



Système classique



Système IBM i



Plus de 40 ans d'innovations et d'investissements ...




L'offre Power Systems en 2012

Un dénominateur commun: la virtualisation



* Pas de support IBM i

L'offre POWER8 en Octobre 2017



POWER et Virtualisation







La saga de la virtualisation



PowerVM : l'aboutissement de 40 ans de virtualisation

... La Virtualisation permet d'avoir la bonne ressource au bon endroit au bon moment, ou comment faire plus avec moins !!





La virtualisation sur POWER

Les composants de PowerVM



Le matériel POWER intègre ses propres fonctions de virtualisation:

- ✓ Partitionnement logique et micro partitionnement
- ✓ Processeur dédié ou processeur partagé (*capping/uncapping*)
- ✓ Processeurs virtuels
- ✓ Groupes de Processeurs Partagés Multiples
- ✓ Partitionnement Logique Dynamique (DLPAR)
- ✓ Adaptateurs virtuels:
 - adaptateurs virtuels SCSI (*Serveur et Client*)
 - adaptateur virtuel « Fiber Channel »
 - adaptateurs virtuels Ethernet (*entre partitions*)
 - adaptateur série
 - consoles virtuelles (*IVM*)




Un grand ordonnateur ... l'HyperViseur

Un ange gardien ... le Flexible Service Processor



POWER Hypervisor (PHYP)

- Couche logicielle entre le matériel et les systèmes d'exploitation des partitions
- Toujours présent et actif, quelque soit la configuration du serveur POWER
- L'hyperviseur n'a pas de ressource spécifique assignée (*processeur et mémoire*)
- C'est grâce à l'hyperviseur que les fonctions embarquées dans le processeur POWER sont disponibles:
 - micro partitionnement
 - processeurs virtuels
 - LAN virtuel et adaptateurs virtuels Ethernet (IVE)
 - adaptateurs virtuels SCSI
 - adaptateurs virtuels Fiber Channel
 - consoles virtuelles (IVM)
- L'hyperviseur renforce l'intégrité de chaque partition en offrant une couche de sécurité entre celles-ci et contrôle les niveaux d'interruptions d'E/S pour les différentes partitions






Ressources Physiques du Serveur





Flexible Service Processor (FSP)


- Dispositif matériel et micro-logiciel qui supportent:
 - les diagnostics lors de la mise en route du matériel (*initial boot tests*)
 - initialisation et configuration des partitions
 - détection et correction d'erreur lors du fonctionnement du serveur
 - charge le PHYP en mémoire du serveur à partir de la RAM Flash, initialise et maintient les tables de contrôle de l'Hyperviseur
 - dialogue avec la HMC (*interface web ASMI*) et le RCP (*Real Control Panel*)
 - Toujours présent et actif dès la mise sous tension du serveur administré,
 - Le FSP a ses propres ressources processeur et mémoire (*carte FSP ou module ASIC sur la carte mère*)







Flexible Service Processor

C'est le cœur du serveur.




A mettre à jour régulièrement !!



Attention aux dépendances HMC / FSP

Architecture de virtualisation PowerVM





La virtualisation et PowerVM





PowerVM et ses différentes éditions

👉 C'est le composant de virtualisation sur plateforme POWER


Un socle technique unique, 3 éditions ...

- ✓ Une technologie unifiée quelque soit le serveur Power et le système d'exploitation AIX, i et Linux
- ✓ Basé sur 40 ans d'excellence technique dans le domaine de la virtualisation
- ✓ Quasi 100% des serveurs Power sont commandés avec PowerVM
- ✓ HMC/IVM/VHMC obligatoire pour virtualiser la gamme des P6, P7 et P8
- ✓ Express réservée aux serveurs Power Express. Standard ou Enterprise : tout modèle
- ✓ AMS.. P6 + VIOS + Shared Processor Pool (*Active Memory Sharing*)

Editions PowerVM

PowerVM Edition :	Express	Standard	Enterprise
Maximum LPARs	3 /Server *	10 / Core	10 / Core
Management	IVM	IVM, HMC	IVM, HMC
Virtual I/O Server	✓	✓	✓
Lx86	✓	✓	✓
Multiple Shared Processor Pools		✓	✓
Active Memory Sharing			✓
Live Partition Mobility			✓

* VIOS + 2 partitions



Les applications de la virtualisation sur IBM i ...






Virtual I/O Server (VIOS)



Le VIOS d'un coup d'œil ...

- C'est une machine virtuelle à usage particulier résidante dans une partition spécifique...
 - Assure *la virtualisation des E/S*
(disque, bande, DVD et réseau)
 - *Supporte les fonctions avancées de cette virtualisation (AMS sauf IBM i, LPM ...)*
- Première version de l'appliance en 2004
- Construite sur une base AIX, mais ce n'est pas une partition AIX
- S'installe sur POWER5 (AIX), 6, 7 et 8
- IBM i 6.1 est la première version supportant le VIOS en 2008
- La licence VIOS fait partie intégrante de PowerVM
(composant optionnel)
- Supporte des partitions clientes AIX, IBM i et Linux



Niveau du VIOS en Novembre 2016 ..
VIOS 2.2.5.10 Sans Interim Fix



Virtual I/O Server (VIOS)


Virtualisation du stockage



IBM i : *partitions et virtualisation*

■ Partitionnement indépendant


- Chaque partition possède son stockage
- Les partitions sont totalement indépendantes les unes des autres



Attachment IXA et IXS via iSCSI et NWSD (serveurs Windows et VMware)

■ Partitionnement par virtualisation (V6R1)


- La partition « hôte » assure le stockage par mutualisation (disques internes ou baie)
- Les partitions sont dépendantes du fonctionnement de la partition hôte



Attachment IXA et IXS via iSCSI et NWSD (serveurs Windows et VMware)

■ Partitionnement et VIOS (V6R1)


- Le VIOS va permettre une virtualisation des ressources dont le stockage
- Les partitions dépendent du VIOS, le stockage est interne ou externe et peut utiliser des cartes PCIe FC 8/16Gbp en NPIV




Le VIOS et l'IBM i 6.1 hébergeant une partition IBM i 6.1 nécessite un serveur POWER6 à minima

Virtualisation des E/S sur POWER

Virtualisation des Bus E/S avec des adaptateurs dédiés



Virtualisation des adaptateurs E/S via VIO Server



IBM i et « Open Storage »


Jusqu'à présent, les baies IBM DS6800, DS8100, DS8300 et DS8700/8800 implémaient le support du formatage de disque en 520 octets par secteur pour supporter l'i5/OS ou l'IBM i.

Grâce à la virtualisation, quand l'environnement IBM i est virtualisé et s'adresse à une partition VIOS, son stockage est hébergé sur des secteurs disques formatés en 512 octets (VIOS=AIX).

PowerVM contient un composant nommé POWER Hypervisor. Ce dernier va transmettre à la partition IBM i 6.1 ou 7.1 neuf secteurs de 512 octets, les huit premiers contenant les données et le neuvième contenant les « headers » de ces huit secteurs (soit 64 octets). **Ce sera au microcode IBM i de se charger de reconstituer le tout.**


En final, cette technique, certes un peu coûteuse en CPU, permet le support de l'IBM i sur des baies de moyenne gamme IBM. Il s'agit des DS3200, DS3400, DS4700, DS4800, DS5020, DS5100 et DS5300, Storwize V7000.

Attention au dimensionnement de la baie: cette fonction nécessite **une augmentation de 12.5% de l'espace utile (8/9).**




La mécanique inverse est utilisée par la fonction « Header Strip/Merge » pour augmenter le débit d'E/S vers le stockage externe. Cette fonction est prise en charge par l'adaptateur E/S et n'est supportée que par les Smart IOA

Solution classique: IBM i + Virtual SCSI



Virtualisation du stockage IBM i avec PowerVM_VIOS_VSCSI









* Multi-path IBM i (i 6.1.1 et plus) ou miroir pour la redondance du Virtual I/O Server

(**) *Subsystem Device Driver Path Controller Module*


<http://www-01.ibm.com/support/docview.wss?uid=ssg1S4000201>

VIOS ... Multi-path et Miroir – Options de configuration




IBM i + VSCSI Storage Virtualizer

- Le stockage interne est associé à la partition du VIOS
- A partir du VIOS, on mappe les LUNs (hdisk) au vhost correspondant à la partition cliente
- L'allocation de la gestion du stockage est effectuée à la fois à partir des disques internes et du VIOS
- La taille des équivalents disques est variable – jusqu'à 2To, formatés en 512 octets par secteur
- Maximum 16 disques par adaptateur vscsi
- Les cartes #5913 et suivantes ne sont pas recommandées dans cette configuration en IBM i





IBM i + Bande et unité optique en VSCSI

- Ces unités sont associées à la partition du VIOS
- A partir du VIOS, on va mapper la bande physique ou l'unité optique au vhost correspondant à la partition cliente
- *Attention:* pas de support de bandothèque via l'adaptateur VSCI.
Nécessité d'utiliser un adaptateur de type FC attaché au VIOS dans le cas de la virtualisation d'une bandothèque LTO.



IBM i + N_PIV (Virtual Fiber Chanel)

Stockage externe




- L'Hyperviseur assigne 2 WWPNs uniques à chaque Virtual Fiber Adaptateur
- Un "Hostconnect" est créé en tant que "hosttype" iSeries,
- Dans le cas d'une DS8K connectée au "hostconnect" iSeries, il faudra formater les LUNs qui lui sont associés en FB 520 octets par secteur
- Possibilité de migrer les LUNS existants si connectés en direct
- Sont supportées les DS8xx0 et la famille Storwize

Note: Un switch compatible N_PIV (N_port) est obligatoire pour connecter le VIOS à la baie

Seules les cartes FC 8Gb et 16Gb sont compatibles N_PIV

Virtualisation du stockage IBM i avec PowerVM VIOS N_PIV



Virtualisation N_Port ID (NPIV) :

Le Virtual I/O Server se comporte comme un path-through FC.

Les unités logiques bandothèques et disques DS8000 sont vues par la partition cliente IBM I comme des unités natives de type 3592 (bandes) et 2107 (disques)

- ✓ Seules les cartes FC 8Gb PCIe Gen2 / 2-ports et 4-ports sont compatibles N_PIV (#5735 et #5729) ainsi que les suivantes

* Multi-path IBM i (i 6.1.1 et plus) ou miroir pour la redondance du Virtual I/O Server


Virtualisation des unités de sauvegarde

Mutualisation des ressources de sauvegarde (bandothèques) via VIOS et N_PIV



- Pour les bandothèques :

- ✓ 3573 (TS3100/TS3200) avec LTO3/LTO4/LTO5
- ✓ 3576 (TS3310) avec LTO3/LTO4/LTO5
- ✓ 3577 (TS3400) avec TS1120/TS1130 (*)
- ✓ 3584 (TS3500) avec LTO3/LTO4/LTO5, TS1120/TS1130 (*), 3592-J1A
- ✓ VTL ProtecTIER TS7610, TS7600A (*Appliance*) et TS7650G (*Gateway*)

(*) lecteurs Jaguar2 et Jaguar3



En synthèse VSCSI et support NPIV



- Nombreux sous-systèmes de stockage supportés
- Stockage affecté au VIOS en premier, puis virtualisé pour les partitions IBM i
- Configuration possible avec HMC, VHMC ou IVM

- Support du DS8xx0, Storwize et de certaines bandothèques
- Stockage affecté directement à l'adaptateur virtuel FC dans la partition IBM i,
- Le VIOS est utilisé en mode « pass through »

(*) La partition cliente identifie le type de disque




Virtual I/O Server (VIOS)

Virtualisation du Réseau




Configuration réseau : pont Ethernet virtuel



```
$ mkvdev -sea ent0 -vadapter ent1 -default ent1 -defaultid 10
```

³ent2 available


Configuration réseau : agrégat (LACP)*



\$ mkvdev -lnagg ent0,ent1
38 ent3 available

(*) Link Aggregation Control Protocol 802.3ad au niveau 2 du modèle OSI


Configuration réseau : agrégat + backup



```
$ mkvdev -lnagg ent0,ent1 -attr backup_adapter=ent2
```

³⁰ent4 available

Dual VIOS : cartes Ethernet






Toujours plus de virtualisation ...



Virtualisation ... PowerVM et IBM i

- Complète les briques pour l'implémentation du Cloud en IBM i
- Augmente l'utilisation et la productivité des systèmes
- Performance et Résilience
- Chaque Technology Refresh amène de nouvelles fonctions de virtualisation




Coup d'œil sur l'IBM i






IBM i Version.Release avec Technology Refresh



Les Technology Refresh intermédiaires apportent des nouvelles fonctionnalités d'entrées/sorties simples à installer dans la version actuelle du serveur



Evolution de l'IBM i ...



- ✓ Mise à niveau intermédiaire grâce au Technology Refresh pour l' IBM i 7.x (par le biais de PTFs)
SF99707 pour V7R1, SF99717 pour V7R2 et SF99727 pour V7R3
- ✓ Avantages des Technology Refreshes intermédiaires ...
 - Amènent le support des nouvelles fonctions et des nouveaux dispositifs E/S
 - Simples à installer sur la version en cours, ne nécessitant pas un arrêt de production.

✓ Pas de nouvelle version ni en 2012, ni en 2013.... mais en 2014 et 2016 !!

Les nouvelles fonctionnalités ont été incluses dans les différents Technology Refreshes disponibles depuis Octobre 2010

- ✓ La version IBM i 7.2 intègre les différents TR de la version 7.1 et ajoute de nouvelles fonctions
- ✓ La version IBM i 7.3 est en partie réécrite et ajoute de nouvelles fonctions et de nouveaux produits

** All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

HMC level: V8.R8.4.0 SP2 + eFix MH01680

IBM i System Support

<http://www-947.ibm.com/systems/support/i/planning/upgrade/osmapping.html>

Servers	IBM i 5.4 ¹	IBM i 6.1	IBM i 7.1	IBM i 7.2	IBM i 7.3
POWER8 S814, S822(VIOS only), S824, E870, E880	✗	✗	✓ ⁴	✓	✓
POWER7/7+ Power 710, 720, 730, 740, 750, 760, 770, 780, 795	✗	✓ ^{2 / 3}	✓	✓	✓
POWER7/7+ Blade/Compute Node PS700/701/702/730/704 PureFlex p260/460	✗	✓	✓	✓	✗
POWER6 JS12, 22, 23/43, 550* 560	✗	✓	✓	✓ ⁵	✗
POWER6 520, 550*, 570, 595	✓	✓	✓	✗	✗
POWER5/5+ 515, 520, 525, 550, 570, 595	✓	✓	✓	✗	✗

1 – IBM i V5R4 is no longer marketed or supported other than through extended service contracts

2 – POWER7+ 750/760 do not support native I/O.


3 - IBM i 6.1 in PureFlex must be client of 7.1 or later

4 – Requires Technology Refresh 8

5 – no IOP or HSL support 5

Support de l'IBM i V5R4 / V6R1Mx

- ✓ 3 Versions actuellement disponibles, fonction « support et service » normale



Java

- Built by IBM to Exploit POWER8
- IBM 32-bit & 64-bit Java™ Virtual Machine
 - Technology shared with AIX, IBM i and Linux
 - Improves performance & portability
- Exceptional Java environment and performance
- **Support for Java 5 and Java 1.4.2 is ending**
 - We no longer will be able to provide fixes for these older versions
- To be compliant, need to be on the latest version of Java



© 2015 IBM Corporation

Le POWER8



Scale Out / Scale In definitions ...

- To *scale horizontally* (or *scale out*) means to add more nodes to a system, such as adding a new computer to a distributed software application. An example might be scaling out from one Web server system to three.
- To *scale vertically* (or *scale up*) means to add resources to a single node in a system, typically involving the addition of CPUs or memory to a single computer. Such vertical scaling of existing systems also enables them to use virtualization technology more effectively, as it provides more resources for the hosted set of operating system and application modules to share.

1 & 2 Socket Servers

- New Scale-Out Servers with POWER8 technology
 - 1 socket : 4U **S814**
 - 2 socket: 2U and 4U **S822** and **S824**
- Linux-only Power Systems (*Not called “PowerLinux”*)
 - 1 socket: 2U **S812L**
 - 2 socket: 2U **S822L**

Marketing name = “socket” then “number EIA/U”

For example:


S 824





S for Scale Out

Pour les serveurs E850, E870 et E880 ... E pour Entreprise Serveur

Power Processor Technology Roadmap



Next Power Processor Technology Roadmap



Power Systems is investing for the future with enhanced POWER8 and POWER9 Processor


POWER7	POWER8	POWER9	POWER10
2010 POWER7 8 cores 45nm New Micro-Architecture New Process Technology	2012 POWER7+ 8 cores 32nm Enhanced Micro-Architecture New Process Technology	2014 POWER8 12 cores 22nm New Micro-Architecture New Process Technology	2016 POWER8 w/ NVLink Enhanced Micro-Architecture With NVLink – 5X Bandwidth New Micro-Architecture New Process Technology
Large eDRAM L3 Cache Optimized VSX Enhanced Memory Subsystem	Optimized for Data-Centric Workloads CAPI Acceleration / I/O	Datacenter TCO Optimization Acceleration Enhancements to CAPI and NVLINK	2020+ POWER10 10nm New Micro-Architecture New Process Technology Extreme Analytics Optimization

Des microchips tridimensionnels pour l'ordinateur de demain

Une nouvelle technologie devrait permettre de multiplier par 10 la puissance des micro-ordinateurs. Développés à l'EPFL en collaboration avec l'EPFZ et IBM, les premiers processeurs tridimensionnels devraient faire leur entrée sur le marché en 2015. Pour la version munie d'un système de refroidissement interne, il faudra tout de même patienter jusqu'en 2020.

Un microchip actuel

Cœurs - Aujourd'hui, les processeurs de nos ordinateurs comptent souvent plusieurs coeurs de 2 à 16.




Câblage - Les coeurs sont disposés côté à côté et communiquent grâce à un câblage.

- Le système de câblage consomme beaucoup d'énergie.
- Le système de câblage génère une importante chaleur qui peut endommager le microchip.
- Chaleur et consommation d'énergie limitent l'évolution de cette technologie.

Un microchip tridimensionnel

Cœurs - Les coeurs ne sont plus disposés côté à côté mais empilés les uns sur les autres.




+ 100 connexions au mm² relient les coeurs sur toute leur surface. Avantages:


- L'échange de données entre les coeurs est beaucoup plus rapide.
- Le processeur consomme beaucoup moins d'énergie.
- Le processeur génère peu de chaleur.

Canaux - Aussi fin que des cheveux (50 microns), des canaux remplis d'eau ou de liquide réfrigérant traversent le microchip pour le maintenir à une température idéale.

POWER8 Processor



POWER8



- ***Significant Performance at Thread, Core, and System***
- ***Optimization for VM Density & Efficiency***
- ***Strong Enablement of Autonomic System Optimization***
- ***Excellent Big Data Analytics Capability***

POWER8 Highlights

... *MEMORY and BANDWIDTH*


Linley Group microprocessor report: “The Power8 specs are mind boggling... IBM’s newest server processor will smash existing performance records, particularly for memory-intensive applications”

■ Caching Structure

- L1 to L4 cache with NUCA Non-Uniform Cache Architecture
- 4 TB/s L2 bandwidth* per chip (4GHz 12core)
- 3 TB/s L3 bandwidth* per chip (4GHz 12core)

■ Memory Subsystem

- 196-230 GB/s sustained* bandwidth per chip
- 410 GB/s bandwidth* at DRAM level per chip
- 1 TB memory per socket (e.g. 4sockets = 4 TB)
- Transactional Memory




* approximate

POWER8 Highlights

... *I/O, BANDWIDTH, VALUE*

TechInvestor/VentureBeat: “IBM preps its massive 12-headed Power 8 chip”

- Balance I/O capability
 - PCIe Gen3 **on-chip** I/O connectivity and protocol (Low latency)
 - I/O bandwidth 96 GB/s per socket
 - Flexible chip interface
- Energy 3X capacity per watt improvement*
- Improved RAS
- P6/P7/P8 modes and Live Partition Mobility
- Workload density / LPAR density



POWER8 Processor

Cores


- 6 cores (SMT8)
- 8 dispatch, 10 issue, 16 exec pipe
- 2X internal data flows/queues
- Enhanced prefetching
- 64K data cache, 32K instruction cache

Accelerators

- Crypto & memory expansion
- Transactional Memory
- VMM assist
- Data Move / VM Mobility

Technology


22nm SOI, eDRAM, 15 ML 650mm²



Energy Management


- On-chip Power Management Micro-controller
- Integrated Per-core VRM
- Critical Path Monitors

POWER8 DCM



Caches

- 512 KB SRAM L2 / core
- 96 MB eDRAM shared L3
- Up to 128 MB eDRAM L4 (off-chip)



Memory

- Up to 230 GB/s sustained bandwidth
- Up to 1TB per socket

Bus Interfaces


- Durable open memory attach interface
- Integrated PCIe Gen3
- SMP Interconnect
- CAPI (Coherent Accelerator Processor Interface)




CAPI

POWER8 Multi-threading Options

- **SMT1:** Largest unit of execution work
- **SMT2:** Smaller unit of work, but provides greater amount of execution work per cycle
- **SMT4:** Smaller unit of work, but provides greater amount of execution work per cycle
- **SMT8:** Smallest unit of work, but provides the maximum amount of execution work per cycle
- Can **dynamical** shift between modes as required: SMT1 / SMT2 / SMT4 / SMT8
- Mixed SMT modes supported within same LPAR
 - ☞ Requires use of “Resource Groups”




POWER8 I/O Bandwidth



Memory Buffer Chip *with 16MB Cache !!!*

DDR3 or DDR4 CDIMMs



« Centaur L4 buffer chip »

Intelligence Moved into Memory


- Scheduling logic, caching structures
- Energy Mgmt, RAS decision point
 - Formerly on Processor
 - Moved to Memory Buffer

Processor Interface


- 9.6 GB/s high speed interface
- More robust RAS
- “On-the-fly” lane isolation/repair
- Extensible for innovation build-out

Performance Value


- End-to-end fastpath and data retry (latency)
- Cache → latency/bandwidth, partial updates
- Cache → write scheduling, prefetch, energy
- 22nm SOI for optimal performance / energy
- 15 metal levels (latency, bandwidth)




POWER8 Memory Organization (Max Config shown)



POWER8 Memory Organization



POWER8 Integrated PCI Gen 3




Native PCIe Gen 3 Support

- Direct processor integration
- Replaces proprietary GX/Bridge
- Low latency
- Gen3 x16 bandwidth (32 GB/s)

Transport Layer for CAPI Protocol

- Coherently Attach Devices connect to processor via PCIe
- Protocol encapsulated in PCIe

POWER8



CAPI



CAPI Enables off-chip accelerator to act as an on-chip processor

- **How do we do it ?**


- Unit on CPU that extends on chip coherence to attached accelerator

- **Efficient model for accelerator**

- Eliminates the data copy required for other accelerators like GPUs
 - Low latency accesses possible since host CPU not involved

- **Flexible**

- Can choose from a large mix of FPGA devices
 - Enables correct mix of CPU and FPGA performance
 - Any PCIe device can be enabled




POWER8 CAPI (Coherent Accelerator Processor Interface)

Virtual Addressing

- Accelerator can work with same memory addresses that the processors use
- Pointers de-referenced same as the host application
- Removes OS & device driver overhead

Hardware Managed Cache Coherence

- Enables the accelerator to participate in “Locks” as a normal thread
- Lower Latency over IO communication model



PCIe Gen3
Transport for encapsulated messages

Processor Service Layer (PSL)


- Present robust, durable interfaces to applications
- Offload complexity / content from CAPP

Customizable Hardware Application Accelerator

- Specific system SW, middleware, or user application
- Written to durable interface provided by PSL

(*) Field Programmable Gate Array
Application-Specific Integrated Circuit

Possible Example: CAPI Attached Flash Optimization




- Attach flash memory to POWER8 via CAPI coherent Attach
- Issues Read/Write Commands from applications to eliminate 97% of instruction path length
 - CAPI Flash controller Operates in User Space
- Saves 10 Cores per 1M IOPs

Serveurs POWER8



L'offre POWER8 en Février 2017



Power Systems Express

Processor and Memory


- Up to 12-core POWER8 processor card
- Up to 8 Simultaneous Multi Thread per core
- Transactional Memory
- Java Code Optimization w/HW Assist

I/O innovation

- PCIe Gen3
- Coherent Accelerator Processor Interface

RAS

- RAID 0, 5, 6, 10 in the base
- Concurrent maintenance PCIe Gen3 slots
- Standard redundant power/cooling



POWER Scale-out Systems - Created for New Workloads & Enhanced Core Values

- Provide efficient consolidation of business workloads
- Designed for demanding environments where uptime is critical and downtime is costly
- Deliver a higher value, more secure, and more reliable platform than competitors



**Power Systems
S824**

- 2-socket, 4U
- Up to 24 cores POWER8
- 2 TB memory
- 11 PCIe Gen 3
- AIX, IBM i, Linux
- Up to 4 CAPI support
- PowerVM



**Power Systems
S814**

- 1-socket, 4U
- Up to 8 cores POWER8
- 1 TB memory
- 7 PCIe Gen 3
- AIX, IBM i, Linux
- Up to 2 CAPI support
- PowerVM

**Power Systems
S822**

- 2-socket, 2U
- Up to 20 cores POWER8
- 1 TB memory
- 9 PCIe Gen 3
- AIX, IBM i & Linux
- Up to 4 CAPI support
- PowerVM



**Power Systems
S822L**

- 1-socket, 2U
- Up to 12 cores POWER8
- 512 GB memory
- 6 PCI Gen3 slot
- Linux only
- Up to One CAPI support
- PowerVM or PowerKVM

**Power Systems
S812L**

- 1-socket, 2U
- Up to 12 cores POWER8
- 512 GB memory
- 6 PCI Gen3 slot
- Linux only
- Up to One CAPI support
- PowerVM or PowerKVM





1 & 2 Sockets


Power**KVM** Power**VP** Power**VM**
Power**HA** Power**VC** Power**SC**



2U Servers



Power S822 Layout



Power S822 Scale-Out System



Single Socket populated

- Cores: 6 (3.8 GHz) or 10 (3.4 GHz)
- Memory: Up to 512 GB
- PCIe Slots: 6 PCIe Gen3 LP (Hotplug)

Both Sockets populated


- Cores: 12 (3.8 GHz) or 20 (3.4 GHz)
- Memory: Up to 1 TB
- PCIe Slots: 9 PCIe Gen3 LP (Hotplug)

Ethernet: Quad 1 Gbt / (x8 slot)

Integrated ports: USB (4), Serial (2), HMC (2)

Internal Storage

- DVD
- 12 SFF Bays -- Split Backplane: 6 + 6
- or 8 SFF Bays & 6 1.8" SSD Bays with Easy Tier with 7GB write cache



Hypervisor: PowerVM

OS: AIX, Linux, **IBM i with maxi 4 cores by partition**

3 years Warranty



Power S822 & IBM i

IBM i 7.2 TR3 ou IBM i 7.1 TR11



Conditions d'éligibilité (Special terms/conditions)

- **Groupe P10** (vs *P20 pour le S824 2-socket*)
- **S822 éligible à l'offre PurePower**
- **Maximum de 4 cœurs par partition -- Plusieurs partitions par serveur possibles**
- **Tous les I/O IBM i doivent être virtualisés à travers un VIOS.**
 - Pas de cartes PCIe Async/Bisync
 - Pas de support pour les cartes Crypto
 - Si utilisation de disque/SSD dans le server ou dans un tiroir EXP24S ou SX, les secteurs doivent être de 4k octets, et non pas de 5xx octets
 - **CBU feature #0444 not announced Oct 2015**




Plus normal 2U limitations of smaller physical server vs 4U server -- smaller memory, fewer PCIe slots, fewer PCIe I/O drawers, subset of PCIe I/O card as full-high cards won't fit including any PCIe SAS adapters with write cache.


Minimum code levels required: IBM i 7.2 TR 3 or IBM i 7.1 TR 11, VIOS 2.2.3.52, FW 840 for 2 cores per partition, FW860 for 4 cores per partition (announced in October 2016)

If using HMC then HMC V8R8.4.0 minimum code level (*HMC V8R8.6.0 available*)



Power S822 Layout




Power S822 ... Planar Layout



4U Servers



Power S814 Layout



Power S814 Scale-Out System


Form Factor: 4U or Tower

- Single Socket


- Cores: 6 (3.0 GHz) or 8 (3.7 GHz)
- Memory: Up to 1 TB
- Slots: 7 PCIe Gen3 Full-high (Hotplug)



- Ethernet: Quad 1 Gbt / (x8 slot)
- Integrated ports: USB (4/5), Serial (2), HMC (2)
- Internal Storage
 - DVD
 - **No Tape !!**
 - 12 SFF Bays -- Split Backplane: 6 + 6
 - or 18 SFF Bays with Easy Tier with 7GB write cache
- Hypervisor: PowerVM
- OS: AIX, IBM i (P10), Linux
- 3 years Warranty



Power S814 Planar Layout



Power S814 Comparison vs p720 – 4-core

	Power 720	Power System S814
Processor	POWER7+	POWER8
Sockets	1	1
Cores	4	4
Maximum Memory*	64 GB @ 1066 MHz	64 GB @ 1600 MHz
Memory Cache	No	Yes
Memory DRAM Spare	No	Yes
System unit PCIe slots	6 PCIe Gen2 FH Opt 4 PCIe Gen2 LP	7 PCIe Gen3 FH
PCIe Hot Plug Support	No	Yes
Ethernet ports in required PCIe adapter	4-port 1 Gbt (in x4 slot)	4-port 1 Gbt (x8 Slot) or 2-port 10/1 Gbt (x8 slot) *
SAS bays in system unit	6 or 8 SFF-1 bays	10** SFF-3 bays
Integrated write cache	Optional 175 MB	Optional effectively 7GB
Easy Tier Support	No (IBM i can use OS integrated capability)	Yes (AIX/Linux) IBM i can use OS integrated capability
Integrated split backplane*	Yes (3 + 3) But not supported by IBM i	Yes (5 + 5) note zero write cache for SAS controller
Bay for SAS tape drive in System unit	1 HH bay	No HH bay


* with latest microcode level in adapter available late August 2015

** Physically more than 10 SAS bays, but a max of 10 drives supported. Could split 4+6 if desired.

Power S814 Comparison vs p720 – 6 or 8-core

	Power 720	Power System S814
Processor	POWER7+	POWER8
Sockets	1	1
Cores	4 / 6 / 8	6 / 8
Maximum Memory	512 GB @ 1066 MHz	1 TB @ 1600 MHz
Memory Cache	No	Yes
Memory Bandwidth	136 GB/sec	192 GB/sec
Memory DRAM Spare	No	Yes
System unit PCIe slots	6 PCIe Gen2 FH Opt 4 PCIe Gen2 LP	7 PCIe Gen3 FH
CAPI (Capable slots)	N / A	One
PCIe Hot Plug Support	No	Yes
IO bandwidth	40 GB/sec	96 GB/sec
Ethernet ports	Quad 1 Gbt (x4 slot)	Quad 1 Gbt (x8 Slot)
SAS bays in system unit	6 or 8 SFF-1 bays	12 SFF-3 bays Or 18 SFF-3 bays
Integrated write cache	Optional 175 MB	Optional effectively 7GB
Easy Tier Support	No	Yes
Integrated split backplane	Yes (3 + 3)	Yes (6 + 6)
Service Processor	Generation 1	Generation 2

Power S824 Layout



Power S824 Scale-Out System

Single Socket populated

- Cores: 6 (3.8 GHz) or 8 (4.1 GHz)
- Memory: Up to 1 TB
- Slots: 7 PCIe Gen3 full-high (Hotplug)

Both Sockets populated

- Cores: 12 (3.8 GHz), 16 (4.1 GHz), or 24 (3.5 GHz)
- Memory: Up to 2 TB
- Slots: 11 PCIe Gen3 full-high (Hotplug)

Ethernet: Quad 1 Gbt / (x8 slot)

Integrated ports: USB (4/5), Serial (2), HMC (2)


Internal Storage

- DVD
- **No Tape !!**
- 12 SFF Bays -- Split Backplane: 6 + 6
- or 18 SFF bays & 8 SSD bays with Easy Tier with 7GB write cache


Hypervisor: PowerVM

OS: AIX, IBM i (P20), Linux


3 years Warranty



Top View Power S824 System




Power S824 Planar Layout



Power S824 Scale-Out Comparison vs p740

	Power 740	Power System S824
Processor	POWER7+	POWER8
Sockets	1 (upgradeable) / 2	1 (upgradeable) / 2
Max Cores	8 / 16	8 / 24
Maximum Memory	512 GB / 1TB @ 1066 MHz	Up to 1TB* / Up to 2TB* @ 1600 MHz
Memory Cache	No	Yes
Memory Bandwidth	136 GB/sec	384 GB/sec
Memory DRAM Spare	No	Yes
System unit PCIe slots	6 PCIe Gen2 FH Opt 4 PCIe Gen2 LP	7 / 11 PCIe Gen3 FH
CAPI (Capable slots)	N / A	Two
PCIe Hot Plug Support	No	Yes
IO bandwidth	60 GB/sec	192 GB/sec
Ethernet ports	Quad 1 Gbt (x4 slot)	Quad 1 Gbt (x8 Slot)
SAS bays in system unit	6 or 8 SFF-1	12 SFF-3 bays Or 18 SFF-3 + 8 SSD bays
Integrated write cache	Optional 175 MB	Optional effectively 7GB
Easy Tier Support	No	Yes
Integrated split backplane	Yes (3 + 3)	Yes (6 + 6)
Service Processor	Generation 1	Generation 2
Software group IBM i	P20	P20

High Performance RAID Feature



S814 (1 socket)

- 4-core 3.02 GHz 39,500
- 6-core 3.02 GHz 59,500
- 8-core 3.72 GHz 85,500

IBM i P05



IBM i P10

S824 (1 or 2 socket)

- 6-core 3.89 GHz 72,000
- 12-core 3.89 GHz 130,000

- 8-core 4.15 GHz 94,500
- 16-core 4.15 GHz 173,500

- 12-core 1-socket not offered
- 24-core 3.52 GHz 230,500

IBM i P20

Measured using
SMT8SMT4 would be
somewhat lower**S822 (1 or 2 socket)**

- 6/12-core 3.89 GHz 23,000 per 2 core
- 8/16-core 4.15 GHz 27,000 per 2 core
- 10/20-core 3.42 GHz 25,500 per 2 core




IBM i P10

Limitation à 4 cœurs
par partition

PCIe Gen3 I/O Expansion Drawer



PCIe Gen3 I/O Expansion Drawer



Rear view

Fan-out Module
6 PCIe Gen3 Slots
4 x8 & 2 x16

Feat #EMXF

Fan-out Module
6 PCIe Gen3 Slots
4 x8 & 2 x16

Feat #EMX0



Front view


- ✓ 12 PCIe Gen3 slots
- ✓ 4U drawer
- ✓ Full high PCIe slots
- ✓ Hot plug PCIe slots
- ✓ Modules not hot plug
- ✓ 8.3 firmware required

Scale-out

- Redondance sur tous les câbles
- Maintenance concurrente des cartes

Use same Blind Swap Cassette (BSC) as used
in #5802/5877/5803/5873 I/O drawer

PCIe Gen3 I/O Drawer




- ✓ Each fan-out module has two CXP ports
- ✓ One active optical cable (AOC) connects to each port
- ✓ Order one pair of cables (one feat code) per fan-out module:
 - #ECC6 2 meter optical cable pair
 - #ECC8 10 meter optical cable pair

System Node to PCIe Gen3 I/O Drawer


PCIe3 Optical Cable Adapter (#EJ07)
• One #EJ07 per fan-out module
• Can be in any of node's x16 PCIe slots

AOC cable pairs :
#ECC6 2 meter length
#ECC8 10 meter length
One feat code ships two identical cables


- ✓ Connect top CXP port of #EJ07 to top CXP port of fan-out module. Likewise connect bottom port to bottom port. Do NOT reverse !!
- ✓ Do NOT mix lengths of AOC cables for the same fan-out module
- ✓ Do NOT cross cables connecting one fan-out module to two different #EJ07 adapters




POWER & STOCKAGE



IBM i sur Power Systems ... un délicat équilibre




Performances du stockage à échelle humaine



DAS & SAN : les deux options sont possibles




- Les deux options sont importantes
- Les deux options ont leurs avantages
- Possibilité d'utiliser les deux options sur le même serveur




- La solution la plus rapide (*latence faible*)
 - Typiquement la moins chère (*hardware/software*)
 - Souvent la configuration la plus simple
 - Fonction Easy Tier
- Rapide
 - Mutualisation entre serveurs
 - Fonctions avancées:
Flash Copy, Metro/Global Mirror,
Live Partition Mobility, Easy Tiering
 - ...

La stratégie du stockage en IBM i: Aujourd'hui ...



External disk options for IBM i




Les différents paramètres à prendre en compte ...

- Projets et Budgets versus Coût
- Plateformes concernées...
 - type de serveur et système d'exploitation / virtualisation
 - mutualisation / consolidation
- Stockage interne ou stockage externe
- Virtualisation ou attachement direct
- Choix des matériels
- Installation physique ...
 - encombrement / charge au sol
 - câblage et équipements / alimentation électrique (tension et courant)
 - climatisation
- Type de protection ...
 - ✓ RAID-1 / RAID-10 ou RAID-5 / RAID-6
- Disponibilité nécessaire (RAS)
- Solutions HA/DR associées
- Fonctions évoluées (*Copy Services, LPM, Flash Copy, Provisioning, RTC....*)
- Performances attendues
- Compétences et Assistance




IBM i... Resiliency and Replication Solutions



PowerHA SystemMirror for i


An IBM delivered, end-to-end solution for HA, DR and on-line backups



Flash Copy Services




FlashCopy Overview



- FlashCopy is a point-in-time-copy of external storage logical volumes that can be established very quickly and with minimal or no disruption or resource on the production LPAR
- IASP based or Full System solution
- FlashCopy options – full copy, no copy
- Space Efficient FlashCopy volumes can reduce FlashCopy storage by 70-80%
- Use with Quiesce to eliminate IASP vary off or LPAR shut down, journal for object and data integrity of flashcopy
- Automate with *IASP Copy Services Manager for PowerHA on i* or with *Full System Copy Services Manager* (FSCS) from IBM Systems Lab Services
- Integration with BRMS
- Ideal for off-line backup solutions

IBM i Full System FlashCopy Process



Managing LPAR with IP connection to production/backup LPARs, HMC and DS8K or V7K
HMC is key to manage the full system FlashCopy process as FlashCopy cannot be
invoked from a quiesced or powered off production LPAR

POWER & SAUVEGARDES



Les différentes sauvegardes en IBM i

- ✓ **SAVF** ... sauvegardes sur disque
- ✓ **Bandes virtuelles dans l'IFS** ... création de cartouches virtuelles dans l'IFS

1/ Création d'une unité bande virtuelle

```
CRTDEV TAP DEVD(TAP02) RSRCNAME(*VRT) ASSIGN(*YES) TEXT('Unité Bande virtuelle')
VRYCFG CFGOBJ(TAP02) CFGTYPE(*DEV) STATUS(*ON)
```

2/ Création du catalogue d'images, et du fichier IFS.

```
CRTIMGCLG IMGCLG(TAPVRT) DIR('/imgCatalog/tapcatalog') TYPE(*TAP) CRTDIR(*YES | *NO)
ADDIMGCLGE IMGCLG(TAPVRT) FROMFILE(*new) TOFILE(volume01) IMGSIZ(*IMGCLGTYPE) VOLNAM(VOL01) .... la taille par défaut est de 1 Go
puis LODIMGCLG IMGCLG(TAPVRT) DEV(TAP02) OPTION(*LOAD)
```

- ✓ **Dérouleur de bande LTOx unitaire** ... unités autonomes SAS TS22x0



- ✓ **Bandothèques TS3100, TS3200, TS3310, TS3500, TS4500**

les différentes bandothèques disponibles dans l'offre IBM

- ✓ **VTL (Virtualisation, déduplication, transferts d'image de cartouche entre VTL)** ...


système de stockage informatique incluant un serveur, une grappe de disques et un logiciel capable d'émuler cet espace disque en bandes magnétiques.

INSIC 2015-2025 Tape Roadmap

Parameter/Year	2015	2017	2019	2021	2023	2025		
1. Capacity (TB)	8	16	32	63	125	248	41.00%	per year
2. Data rate per channel (MB/sec)	10.0	13.2	17.5	23.1	30.6	40.5	15.00%	per year
3. Total data rate (MB/sec)	320.0	480.2	720.6	1081.4	1622.7	2435.1	22.50%	per year
4. FC Speed Roadmap (MB/sec)**	3200	6400	12800	12800	25600	25600		
5. Number of channels	32	36	41	47	53	60	6.52%	per year
6. Tape thickness (um)	5.20	4.79	4.42	4.07	3.75	3.46	-4.00%	per year
7. Data capacity reserve	3.0%	3.0%	3.0%	3.0%	3.0%	3.0%		defect reserve
8. Tape length that is recordable (meters)***	969	1051	1141	1238	1343	1457	3.90%	winding reserve
9. Tape length total (meters)****	1041	1129	1225	1330	1443	1565	4.17%	per year
10. Track density (TPI)	10,686	15,652	22,868	33,543	49,372	72,891	21.17%	per year
track pitch = 25400/tpi (um)	2.38	1.62	1.11	0.76	0.51	0.35		
11. Linear bit density (kfcf)*****	480	581	703	850	1029	1245	10.00%	per year
fcmm = kfci/0.0254	18,898	22,866	27,668	33,478	40,509	49,016		
12. Areal density (Gbits/inch ²)	5.13	9.09	16.07	28.52	50.80	90.75	33.28%	per year
13. Tape speed (m/sec)	5.4	5.8	6.2	6.7	7.2	7.7	3.61%	per year
14. Tape width in mm	12.65	12.65	12.65	12.65	12.65	12.65		
15. ECC and formatting overhead	22.00%	20.28%	18.69%	17.22%	15.87%	14.63%	-4.00%	per year
16. Servo track and layout overhead *****	15.80%	14.84%	13.46%	12.24%	11.17%	10.22%	-6.00%	per year
17. Number of passes to write a tape	140	183	239	314	412	541		
18. Number of passes to end-of-life (media)	27200	29194	31333	33630	36095	38741	3.6%	per year
19. Time to fill a tape in mins	417	552	731	969	1284	1701	15.10%	per year
20. Number of data tracks	4,481	6,639	9,856	14,660	21,842	32,593	21.95%	per year
21. Number of data bands	4	5	7	9	12	16	15.00%	per year
overall head span (um)	3,000	2,268	1,715	1,297	981	742		
22. Tape Dimensional Stability (ppm)	317	286	259	234	210	188	-5.09%	per year
23. Bit Aspect Ratio (BAR)	53	44	36	29	23	19	-9.80%	per year
24. Bit Error Rate	1.00E-19	5.00E-20	2.50E-20	1.25E-20	6.25E-21	3.12E-21	-29.29%	per year

Site IBM pour trouver les produits de stockage IBM:



<http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=TSO00364USEN>



HMC



HMC Model 7042-CR8 / 7042-CR9



- **7042-CR8 is System x model 7914 Ref3**

- Based on 7042-CR7
- 2GHz Ivy Bridge Intel Proc*
- 8GB DDR3 Mem (1x8GB DIMM)
- Dual 500GB HDD (RAID1)
- Quad GigE (First eth port can be shared IMM)
- Dedicated IMM port

- **UEFI, IMM, DSA is similar to CR7 (same for all 7914)**
- **FW Level: V8.R8.4.0 SP2 + eFix MH01680 (au 02/17)**

POWER8 Systems Management Requirements

Minimum Requirement:

HMC FW	V8.R1.0 or Later
HMC HW	7042-CR5 / 7042-C08 or Later
HMC Memory	2GB or More
Power HW	POWER6 and Later
PowerVM/VIOS	V2.2.3

Enhanced HMC Management Capabilities

Virtual Environment Deployment

Create System Template that meets specific configuration requirements

- § Edit System Template & Capture Configuration as Template

Deploy Virtualization Configuration via “Deploy System from Template”

Create Partition Template that meets specific workload requirements

Deploy Partition

- § Create Partition from Template and Install Operating System

Virtual Environment Management

1. Configure Virtual Network via “Manage PowerVM”
2. View Virtual Storage Adapter Mappings via “Manage PowerVM”
3. Add storage to a partition via “Manage Partition: Virtual Storage”
4. Connect partition to a virtual network via “Manage Partition: Virtual Network”

Performance Monitoring

1. View System Performance Dashboard
2. View Resource Utilization Trends (Processor, Memory, Network, Storage)

New Server-level Tasks

The screenshot shows the Hardware Management Console Workplace (V7R7.6.0.0) running in Mozilla Firefox: IBM Edition. The main window title is "Hardware Management Console Workplace (V7R7.6.0.0) - Mozilla Firefox: IBM Edition". The URL in the address bar is <https://.austin.ibm.com/hmc/connects/mainuiFrameset.jsp>.

The left sidebar navigation includes:

- Welcome
- Systems Management
 - Servers
 - 9.
 - ilex.
 - zh29ce-sflp
 - Custom Groups
- System Plans
- HMC Management
- Service Management
- Updates

The main content area displays the "Systems Management > Servers" list. The table columns are:

Select	Name	Status	Available Process... Units	Available Memory (...)	Reference Code
<input type="checkbox"/>	zh29ce-sflp	Power Off	0	0	
<input checked="" type="checkbox"/>	ilex.aust...	Operating	0.3	21.875	
<input type="checkbox"/>		No Connection	0	0	Connecting 0000-0000-000000

A context menu is open over the selected server "ilex.austin.ibm.com". The menu items are:


- Properties
- Manage PowerVM
- Templates
- Operations
- Configuration
- Connections
- Hardware Information
- Updates
- Serviceability
- Capacity On Demand (CoD)
- Performance

The "Properties" and "Performance" items are highlighted with green boxes.

The status bar at the bottom shows "Status: Attentions and Events" and "Transferring data from .austin.ibm.com...".

System Configuration / Deployment – Using Templates

1) Capture Configuration as Template



2) View & Edit Templates

Template Library

Fugit definiebas ne sea, et vis sapientem consulatu, maiorum corpora maiestatis nam ea. Lucretius molestaie consequat et pri, an magna impetus accusata usu. Utinam aliquip maiestatis ei vix, quot eius dicunt his et, et vis patroisque interpretaris. Ea nam vivendo tractatos. Cu modo saperit incorrupte pro, pro ne odio similique.

Template Name	Description	Partition Type	Processors	Network	Storage	Features
Template A	Template A Description	IBM i	Dedicated	Both	Virtual	LPM
Template B	Template B Description	AIX/Linux	Shared	Dedicated	Both	None
Template C	Template C Description	AIX/Linux	Shared	Virtual	Virtual	LPM

Virtual Networks

Virtual Network Name PVID (PowerVM) VLAN ID (External) Virtual Switch Network Bridge Load Balance Group

Production	1	vSwitch0	Bridge 1	1
Test		vSwitch0	Bridge 1	2
Development	3	vSwitch0	Bridge 1	2
Security	100	vSwitch1	Bridge 2	
Intranet1	10	vSwitch1		

Virtual Switches

Virtual Switch	Mode
vSwitch0	VEB
vSwitch1	VEB

Network Bridges

Network Bridge Name	Failover	Load Balance	Virtual IO Servers
Bridge 1	Yes	Yes	VIOS1, VIOS2
Bridge 2	No	No	VIOS3

[Learn more →](#) [Close](#)

3) Deploy Template

Deploy System from Template


Select a System Template

Target System: [System Name]

Template Name	Description	Number of VIOS	Number of Virtual Networks	Physical I/O Information
Template A	Lorem ipsum dolor sit amet, consectetur adipiscing elit. Sed sit amet neque in nisi aliquam imperdiet. Maecenas blandit ultrices mollis. Nam id augue est. Mauris leo tortor, malesuada sed hendrerit ac, ultricies vitae nunc. Aenean lacus mi, sodales sit amet luctus vel, vestibulum vel massa.	2	3	No
Template B	Lorem ipsum dolor sit amet, consectetur adipiscing elit.	2	4	No
Template C	Lorem ipsum dolor sit amet, consectetur adipiscing elit.	1	2	Yes
Template D	Lorem ipsum dolor sit amet, consectetur adipiscing elit.	0	0	No

[About this wizard →](#) [Back](#) [Next >](#) [Finish](#) [Cancel](#)

PowerVM Management GUI



- Single launch point for all PowerVM configuration
- Management of all VIOS function
- Fast, scalable management of VIOS via a hardened HMC<->VIOS API
- Simplified virtual network & storage model
- Manage VIOS, Networks, and Storage

This screenshot shows the Power VM Configuration interface. It features three main panels: 'Virtual Networks' (listing Production, Test, Development, Security, and Internal networks with their respective VNI and VLAN ID mappings), 'Virtual Switches' (listing vSwitch1 and vSwitch2 with Mode set to VBS), and 'Network Bridges' (listing Network Bridge 1 and Network Bridge 2 with Tethered set to Yes and Virtual I/O Servers assigned to VIOS 1 and VIOS 2). A 'Learn more' link is at the bottom of each panel.

Performance Monitoring – Metrics & Dashboard

Performance metric indicators & utilization dashboard

- Processor, memory & I/O


Server & LPAR level information

Basic trend data collection and visualization


- Identify bottlenecks
- Early problem detection

REST based API to access:

- All platform (PHYP & VIOS) metrics for Tivoli
- Third Party tools



Performance Monitoring – Processor Utilization



Power Systems Management Requirements


Minimum Requirement:

HMC FW	V8.R1.0 or Later
HMC HW	7042-CR5 / 7042-C08 or Later
HMC Memory	2GB or More
Power HW	POWER6 and Later
PowerVM/VIOS	V2.2.3

En synthèse



Le premier processeur conçu pour Big Data





Documentations



L'espace IBM i sur DeveloperWorks

- ✓ Nouvelle ressource pour la communauté IBM i
- ✓ Un large éventail d'informations techniques
- ✓ Organisé par sujet
- ✓ Topics
 - Les nouveautés en IBM i
 - Téléchargement et Outils
 - Bibliothèque technique avec des exemples !
 - Communauté et Forums
 - Les services autour de l'IBM i
 - Evénements
 - Wiki sur les nouveautés technologiques
- ✓ S'enregister pour recevoir les notifications de mises à jour
- ✓ Postez vos commentaires !!!

The screenshot shows the IBM i developerWorks homepage. At the top, there's a navigation bar with links for English, Sign in (or register), developerWorks, Technical topics, Evaluation software, Community, and Events. A search bar is also present. The main content area has a sidebar on the left with links for IBM i, New to IBM i, Downloads & tools, Technical library, Technology updates, Community & forums, Services, and Events. Below the sidebar, there's a "Related links" section with a "Training" link. The main content area features a large image of a computer monitor displaying a database interface. To the right of the image, there's a summary of an article titled "Searching Spool Files and IFS Stream Files". Further down, there are sections for "Featured content" (listing articles like "Invoking web services using RPG proxy" and "Next generation RPG documentation") and "Key topics" (listing topics like "MERGE ahead: Introducing the DB2 for i SQL MERGE statement" and "IBM i Performance Data Investigator: Browser-based viewing of performance data"). On the right side of the page, there are additional sections for "Update My dw interests", "Seeking nominations", "IBMCHAMPION", and "Spotlight".

ibm.com/developerworks/ibmi

**Info Center reste le site de référence pour la documentation technique, toutefois certaines informations peuvent apparaitre sur dW en premier.

IBM i maintenant sur developerWorks!!

developerWorks

Technical topics Evaluation software Community Events Search developerWorks

IBM products: Technologies: Solutions:

AIX and UNIX Java technology Cloud computing
IBM i Linux Industries
Information Management Open source Integrated Service Management
Lotus SOA and web services
Rational Web development Technical library
Tivoli XML Feeds and syndication
WebSphere Newsletters

→ More about Technical topics Close [x]

Innovate 2011: Join the conversation

The smarter and more agile way to follow Rational's premier developer conference, the Innovate 2011 social media center provides tweets, blogs, photos, videos, and live broadcasts to keep you in touch with people and announcements all week long.

Featured content Downloads Popular forums

31 May 2011 — Show descriptions | Hide descriptions

- HTML5 form controls for better, bolder user interfaces (Web development)
- Access HTTP and RESTful services from DB2 (Information Management)
- Functional programming: Knowing when to let go (Java technology)
- Create custom FTP commands for WebSphere Message Broker (WebSphere)
- Configure OpenAFS to support Kerberos 5 authentication (Open source)
- HTML5's new approach to structuring documents (Web development)

Spotlight

- News: Meet the IBM Champions with Rational expertise
- Evaluate: Lotus Symphony
- Free webcast: Building mobile applications with quality and testing (02 Jun)

Seeking nominations

IBM i Technology Updates wiki

- www.ibm.com/developerworks/ibmi/techupdates
- Starting point to learn about IBM i enhancements
- Content organized for easy navigation and subscription
- Timely updates from IBM with all the details

IBM i Technology Updates - by IBM i product or subject matter

Backup Recovery and Media Services (BRMS)

DB2 for i (Database)

General IBM i operating system

Hardware and Firmware (including Technology Refresh content)

IBM i Access Client Solutions

IBM Integrated Web Services for i

Integration with BladeCenter and System x

Java on IBM i


Performance Tools

PowerHA SystemMirror for i

Web Integration on i

New topics

#ibmi
On Twitter



DawnMayiCan
On Twitter



Steve_Will_IBMi
On Twitter

IBM Systems Magazine blogs

- “i Can” by Dawn May
 - http://bit.ly/i_can
 - To share the “hidden gems” within IBM i.
- “You and i” by Steve Will
 - http://bit.ly/You_and_i/
 - Directly communicate with the i community
 - Strategy, architecture, announcements, and news.

<https://www.ibm.com/developerworks/ibmi/>

- DeveloperWorks IBM i Zone

<http://www.redbooks.ibm.com/redbooks/pdfs/sg247858.pdf>

Other Social Sites - Increasing Awareness

1. Power of i website
2. You and i blog
3. i Can blog
4. DB2 for i Blog
5. IBM i Community Page and Forums
6. LinkedIn
7. Twitter
8. Rational Cafes

The screenshot shows the 'i Can' website. At the top, there are three main sections: 'Develop' (with a blue background), 'You and i' (with a blue background), and 'Blogger' (with a white background). The 'Blogger' section features a photo of Mike Cain, his name, and a link to his blog 'DB2 for i'. Below this, there is an 'About me' section with another photo of Mike Cain.

The screenshot shows the LinkedIn homepage with a list of user groups. The groups listed are:

- COMMON A Users Group
- COMMON IT Execs
- I5/AS400 Application Development and Support Network
- QRPGLESRC
- IBM i Applications and Technologies
- IBM i Manifest Americas
- IBM i professionals

The screenshot shows the developerWorks RPG Cafe page. It includes a navigation bar with links for 'Public Groups', 'My Groups', 'Help', and 'This Group'. Below the navigation is a section titled 'Overview' with a large graphic of three stylized people. The text in the 'Overview' section reads:

This is the place for RPG developers and IBM developers to mingle and discuss our favorite language. We'll talk about the language itself, its history, and its future. We will discuss the tools that support it and the directions they are taking. We hope this is a lively area of the Web.


This is a great place for all things related to RPG, to expound on proven techniques, to post interesting bits of code, and to act as a jumping off point into other areas of the IBM Web for support and information.



Questions ??



Merci !!!



ibm.com/systems/fr/power/