

April 29, 2014

# IBM POWER8

**Erik Rex**

[rex@dk.ibm.com](mailto:rex@dk.ibm.com)

**Power S814**  
**Power S822**  
**Power S824**  
**Power S822L**

S = ScaleOut  
8 = POWER8  
x = Nbr of Sockets  
x = Nbr of Rack-U's  
L = Linux only



# Power Systems: Innovation to put data to work

## POWER8 Signature Innovation

### Processor and Memory

- Up to 12-core POWER8 processor card
- Simultaneous Multithread Thread (SMT) 8 per core
- Transactional Memory
- Java™ Code Optimization w/HW Assist

### I/O Improvement

- PCIe Gen 3
- Coherent Accelerator Processor Interface

### RAS

- RAID 0, 5, 6, 10 in the base, JBOD storage
- Concurrent maintenance PCIe Gen3 slots



## New naming



**Power Systems S824**

- 2-socket, 4U
- Up to 24 cores
- AIX, IBM i, Linux



**Power Systems S814**

- 1-socket, 4U
- Up to 8 cores
- AIX, IBM i, Linux



**Power Systems S824L**

- 2-socket, 4U
- Up to 24 cores
- Linux
- 2H14



**Power Systems S822**

- 2-socket, 2U
- Up to 20 cores
- AIX & Linux



**Power Systems S822L**

- 2-socket, 2U
- POWER8 processor
- Up to 24 cores
- Linux only



**Power Systems S812L**

- 1-socket, 2U
- POWER8 processor
- Up to 12 cores
- Linux only
- 2H14



1 & 2 Sockets




# Announce / Availability Plans

	<b>Annc</b>	<b>eConfig</b>	<b>GA</b>
Power S814	28 Apr	28 Apr	10 Jun
Power S822	28 Apr	28 Apr	10 Jun
Power S824	28 Apr	28 Apr	10 Jun
Power S822L	28 Apr	28 Apr	10 Jun
Add'l storage backplane option	28 Apr	17 July	Sept
SFF-3 146/139GB 15k & 300/283GB 10k HDD	28 Apr	27 May	25 July



**2U**



**POWER8**



**4U**

# Scale Out Systems - DCMs and POWER8 Chips

## 1S & 2S servers use DCM (Dual Chip Module)

- 1 DCM fills 1 socket .... Similar to POWER7+ 750 / 760
- 1 DCM has two Scale Out POWER8 chips
- 1 DCM can provide 6-core, 8-core, 10-core or 12-core sockets

### 6-core Processor Chip

- 362 mm<sup>2</sup>
- 22nm SOI w/ eDRAM

### Strengthen Cores

- 8 Threads per Core


### Caches

- D Cache: 64KB
- L2: 512KB
- L3: 8 MB per Region      Total: 48MB

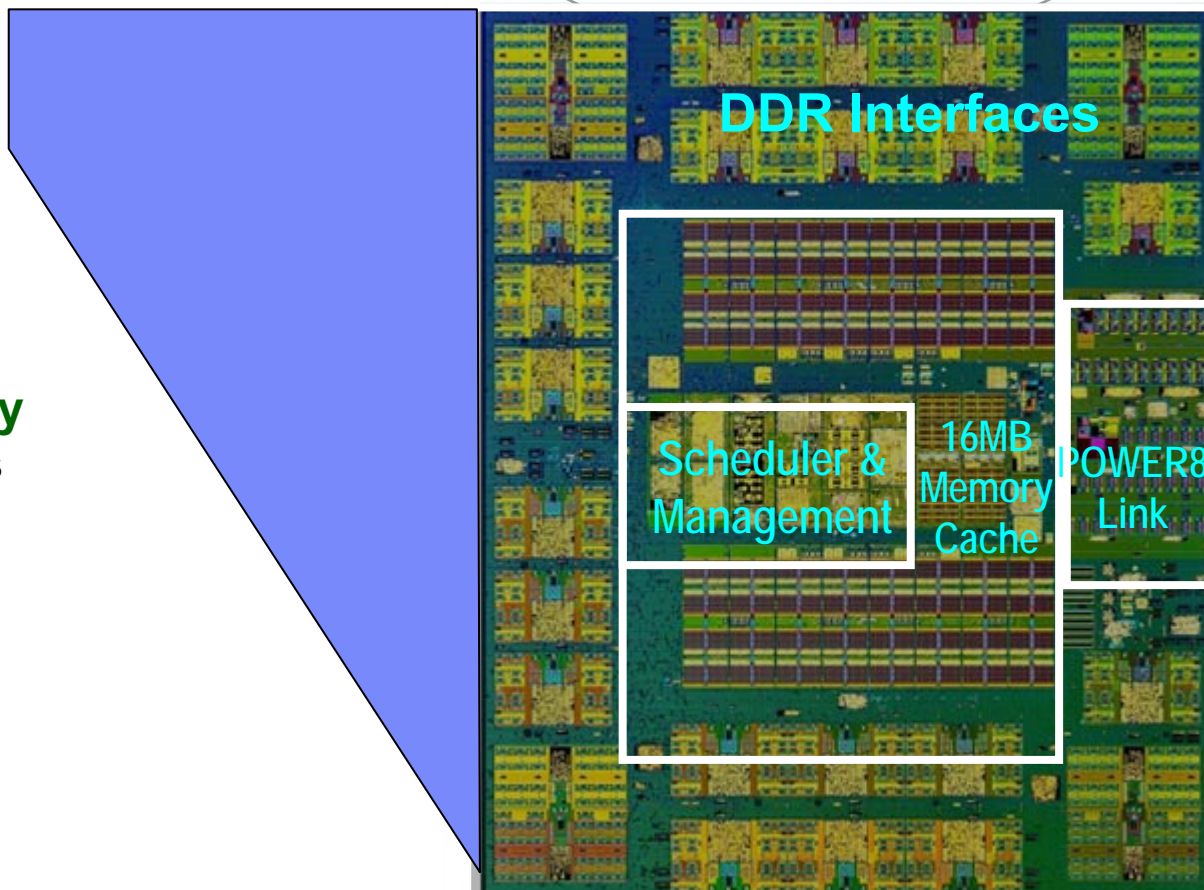


### Fine Grained Power Management

- On Chip power management

### Excellent I/O bandwidth per socket



# POWER8 Memory Buffer Chip



## Intelligence Moved into Memory

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

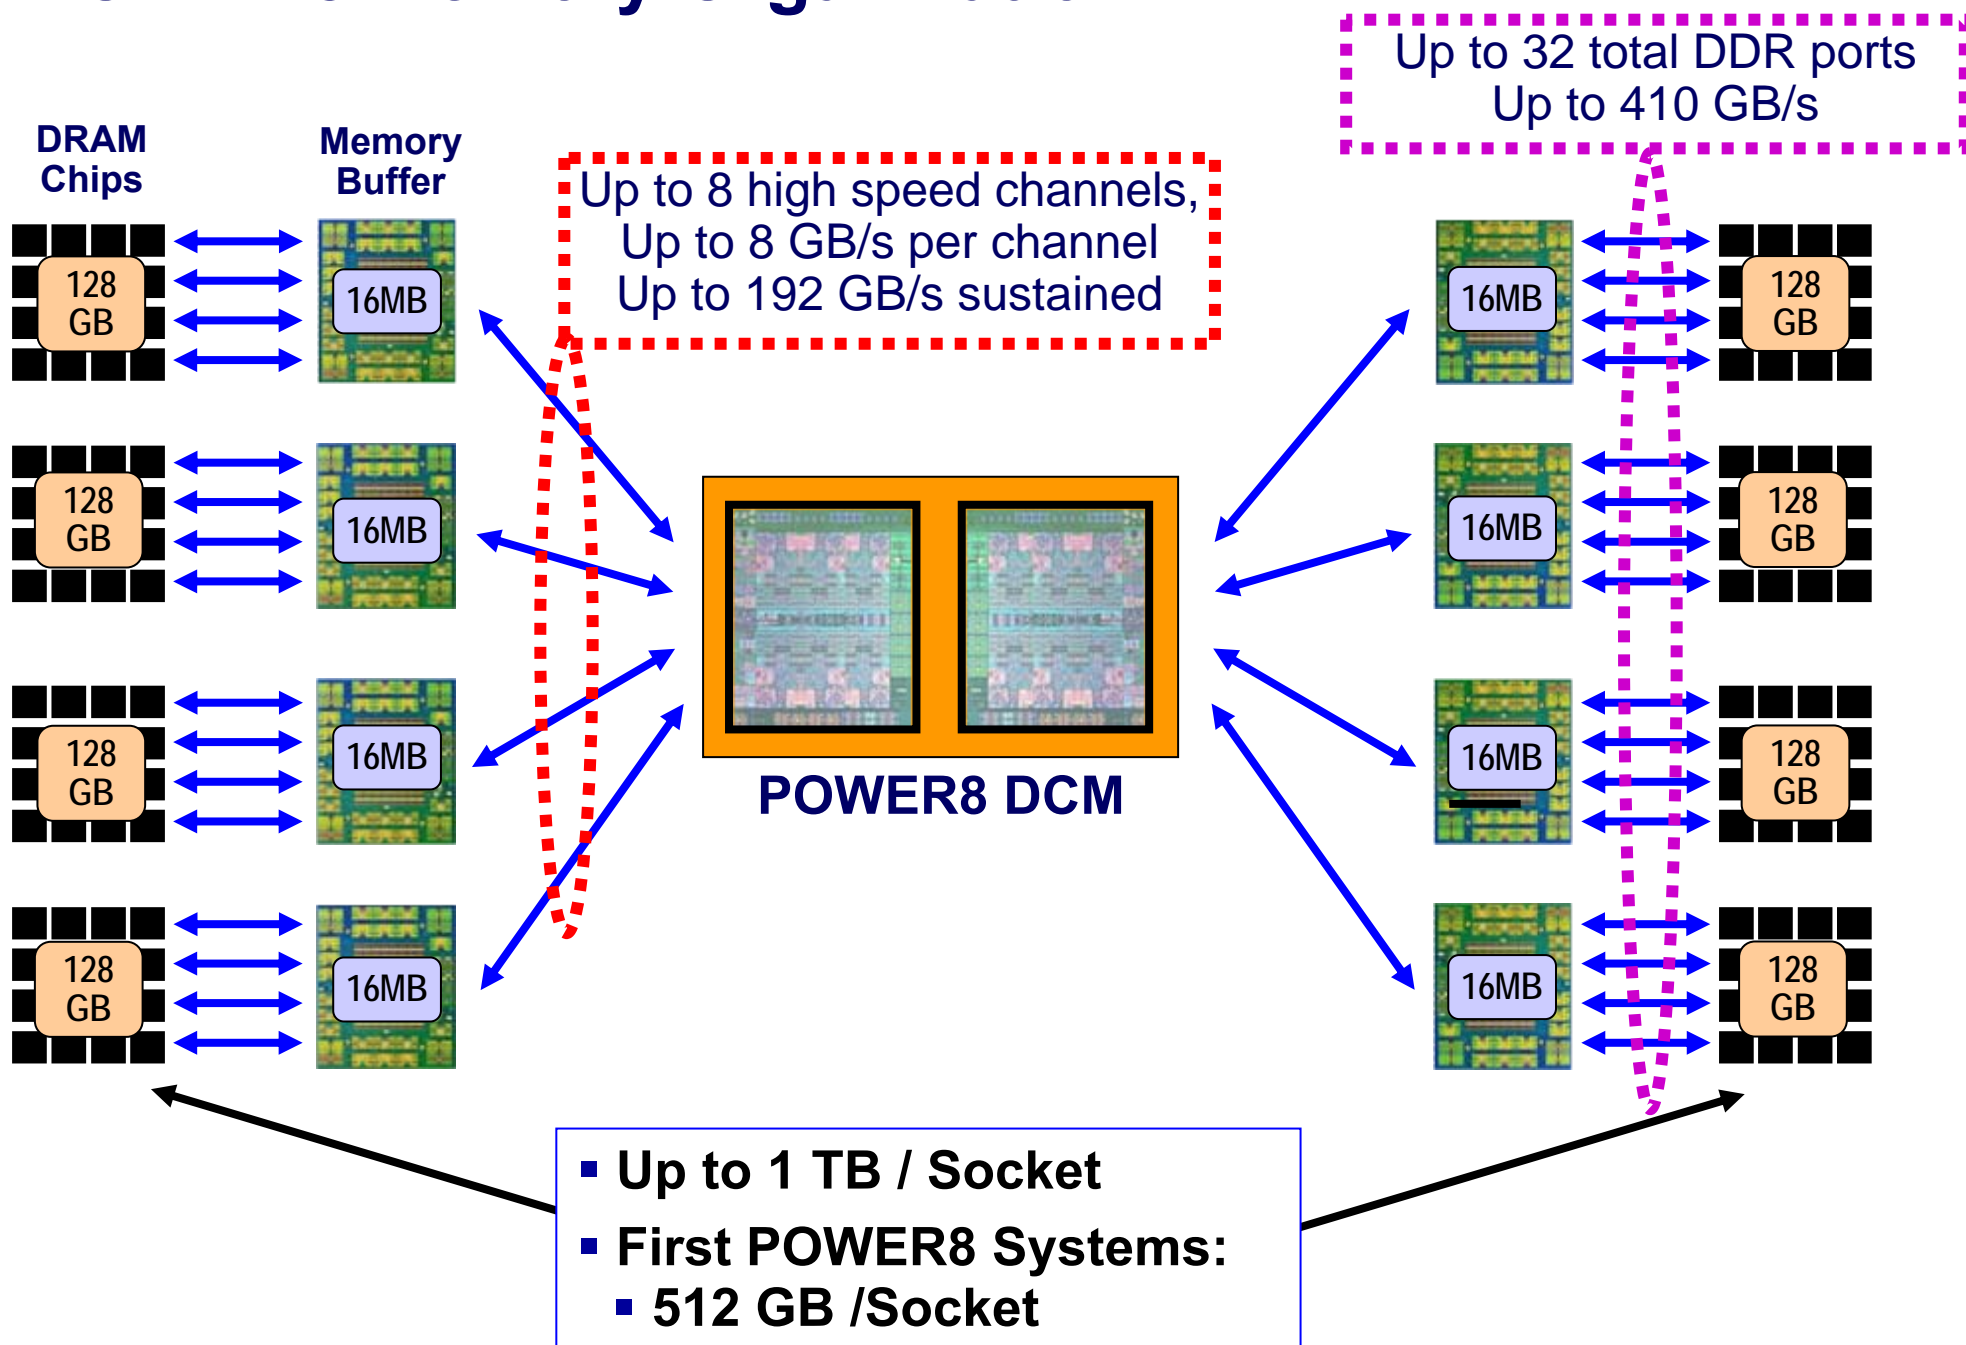
## Processor Interface

- 9.6 GB/s high speed interface
- More robust RAS
- “On-the-fly” lane isolation/repair

## Performance Value


- End-to-end fastpath and data retry (latency)
- Cache → latency/bandwidth, partial updates
- Cache → write scheduling, prefetch, energy

# POWER8 Memory Organization



# POWER8 Integrated PCI Gen 3

## POWER7



GX Bus

I/O Bridge

PCIe Gen2

PCI Devices

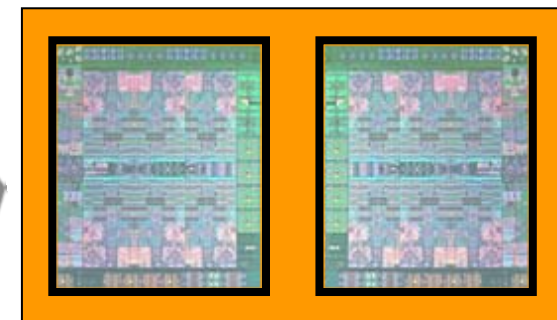
## Native PCIe Gen 3 Support

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

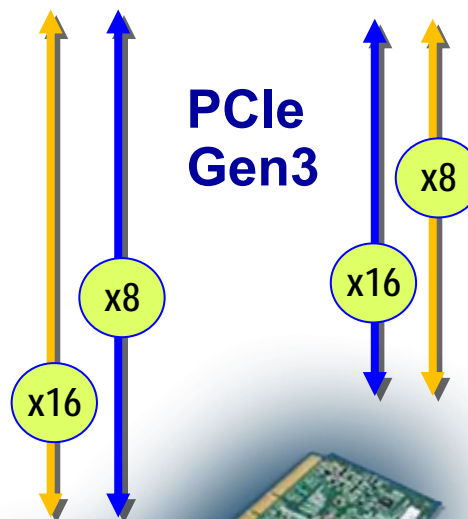
## Transport Layer for CAPI Protocol

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

## POWER8



PCIe Gen3



PCI Devices



# CAPI

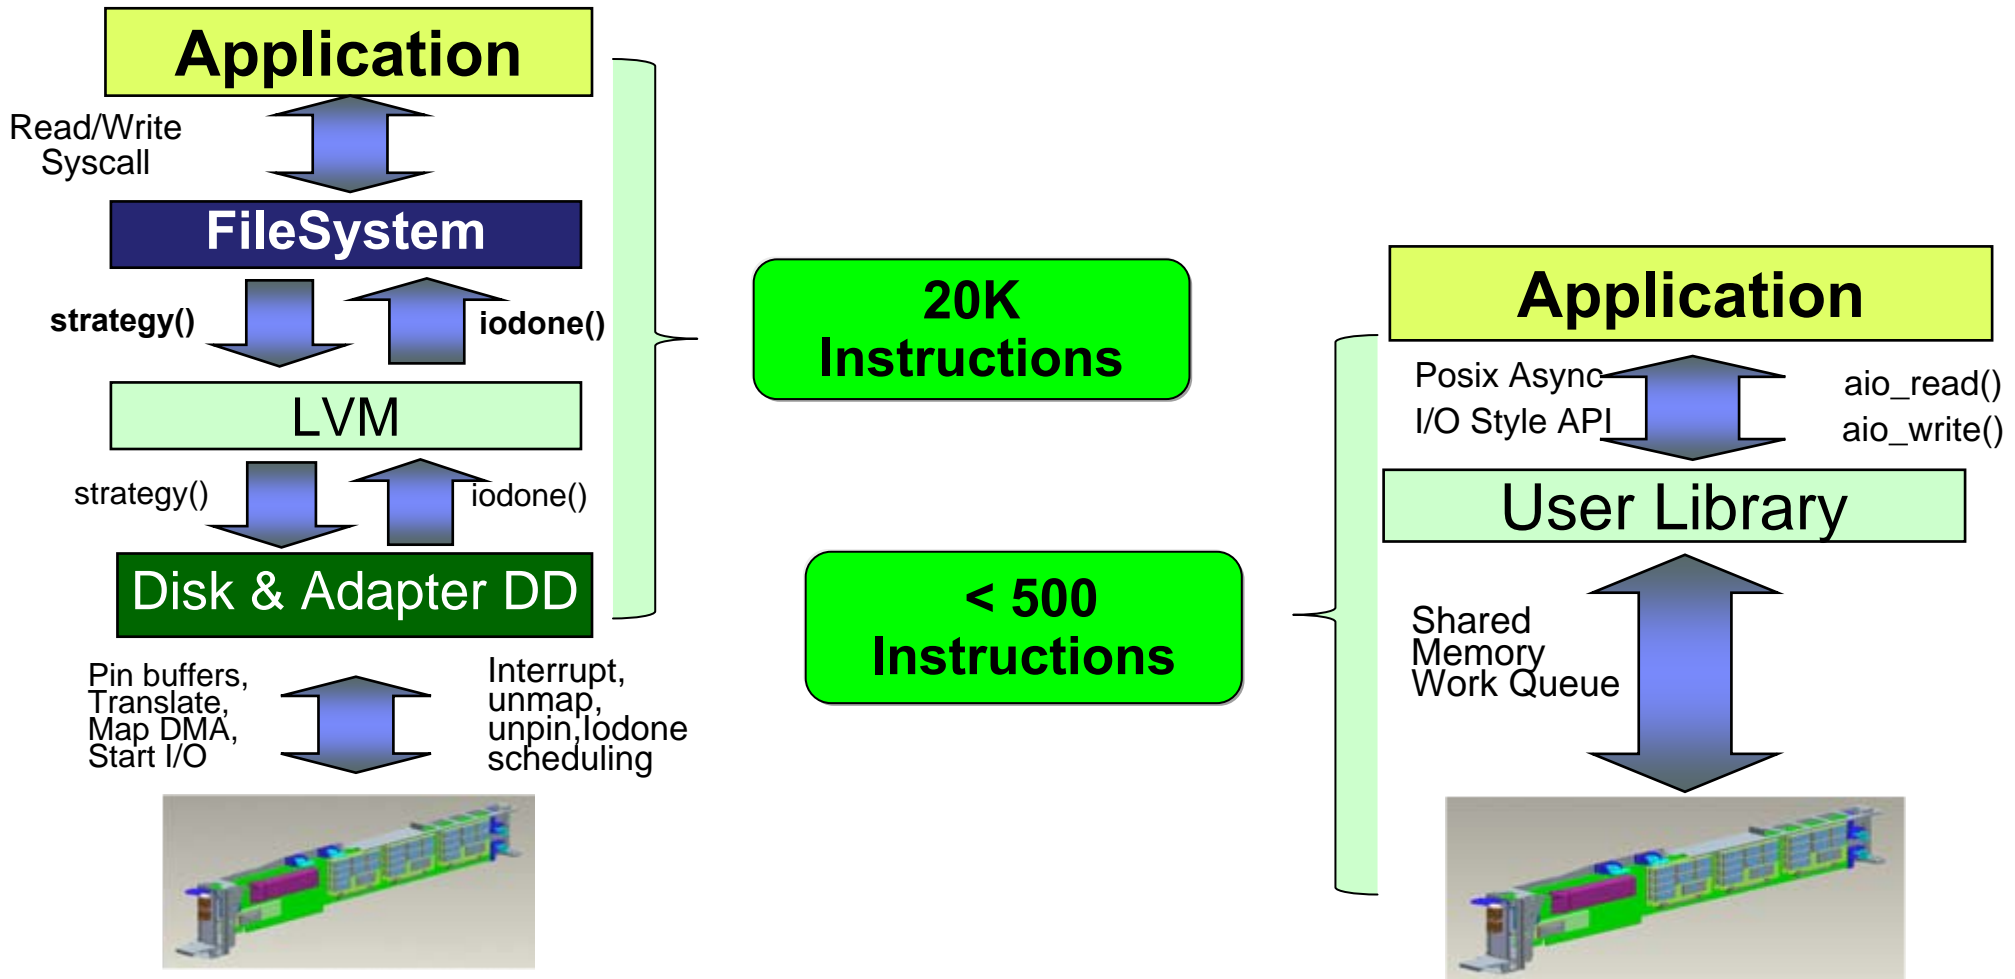
Coherent Accelerator Processor Interface

Stmt of Direction





# 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**



# 2U Server


# POWER8 2S2U Layout



# POWER8 2S2U Scale-Out System

## Power S822


- **Single Socket populated**
  - **Cores:** 6 or 10
  - **Memory:** Up to 512 GB
  - **PCIe Slots:** 6 PCIe Gen3 LP (Hotplug)
- **Both Sockets populated**
  - **Cores:** 12 or 20
  - **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 (not IBM i)



**3 Yr Warranty**


# POWER8 4U Scale-out Comparison – S814

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



# 4U Servers


# System 1S4U Layout



# POWER8 2S4U Scale-Out System


## Power S824

- **Single Socket populated**
  - Cores: 6 or 8
  - Memory: Up to 512 GB
  - Slots: 7 PCIe Gen3 full-high (Hotplug)
- **Both Sockets populated**
  - Cores: 12, 16, or 24
  - Memory: Up to 1 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
  - 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, Linux**



**3 Yr Warranty**







# Performance

# Power 740+ / POWER8 S824 rPerf Comparisons


### Performance per Core



### Performance per Socket




### Performance per KW




# Power 740 / POWER8 S824 CPW Comparisons

Performance per Core



Performance Dual Socket



## IBM i - CPW

720 POWER7+ (1 socket)		
■ 4-core	3.6 GHz	28,400
■ 6-core	3.6 GHz	42,400
■ 8-core	3.6 GHz	56,300


S814 (1 socket)		
■ 4-core	xxxxx	xxxxxx
■ 6-core	3.0 GHz	59,500
■ 8-core	3.7 GHz	85,500

740 POWER7+ (1 or 2 socket)		
■ 6-core	4.2 GHz	49,000
■ 12-core	4.2 GHz	91,700
■ 8-core	3.6 GHz	56,300
■ 16-core	3.6 GHz	106,500
■ 8-core	4.2 GHz	64,500
■ 16-core	4.2 GHz	120,000

S824 (1 or 2 socket)		
■ 6-core	3.8 GHz	72,000
■ 12-core	3.8 GHz	130,000
■ 8-core	4.1 GHz	94,500
■ 16-core	4.1 GHz	173,500
■ 12-core	1-socket not offered	
■ 24-core	3.5 GHz	230,500


# Power 750+ vs POWER8 4U2S

## Performance




POWER8 ≥  
Performance

## Total # Cores




Fewer cores means  
Lower Software Costs  
Does not reflect lower  
PVU savings

## TCA



Lower investment  
costs


## Performance per \$



Better performance  
per dollar spent

# POWER8 Delivers **Over TWICE** the Throughput Compared To Ivy Bridge-EP at **47% Lower Cost**

## Web Application



Online Banking Workload v3.6

**183,800** User Interactions per second

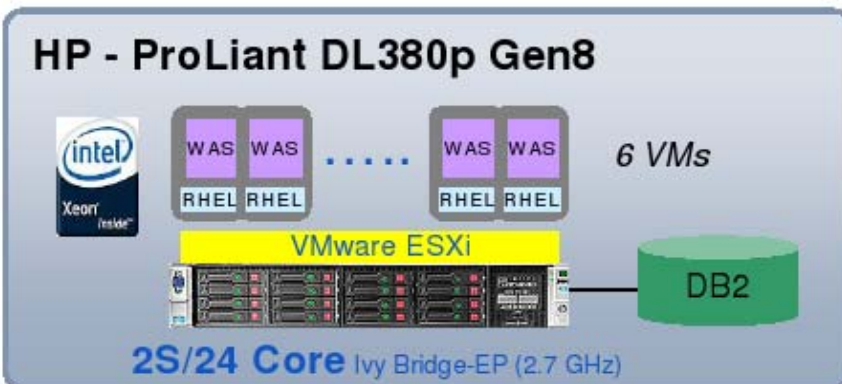
**\$3.09** per UI per sec

WebSphere on platform  
Database off platform

**2.1x**  
Faster

**47%**  
Lower cost  
per UI per sec

Both Servers configured to achieve maximum throughput



**85,939** User Interactions per second

**\$5.84** per UI per sec

WebSphere on platform  
Database off platform

This is an IBM internal study designed to replicate a typical IBM customer workload usage in the marketplace. The results were obtained under laboratory conditions, and not in an actual customer environment. IBM's internal workload studies are not benchmark applications, nor are they based on any benchmark standard. As such, customer applications, differences in the stack deployed, and other systems variations or testing conditions may produce different results and may vary based on actual configuration, applications, specific queries and other variables in a production environment. Prices, where applicable, are based on published US list prices for both IBM and competitor, and the cost calculation compares the cost per request for the 3yr life of the machine. 3 year total cost of acquisition comparisons are based on similar expected hardware, software, service & support offerings

# DB2 BLU on POWER8 is **2.7x Faster** than Oracle DB on Exadata and **80% Lower Cost**

## OLTP (Brokerage) Workload

### DB2 AESE v10.5

Power S824  
with 24 cores  
AIX 7.1, 64-bit  
  
FlashSystem 840



System Cost  
3 year TCA  
**\$2,094,115**

**1,543** Trade Completions  
per second (TPS)  
**\$1,357** per TPS

**2.7x**  
Faster

**80%**  
Lower cost  
per Trade per  
second

### Exadata X3-2 Quarter Rack

Intel-based Processors  
¼ Rack  
32 database cores  
2 Server Nodes  
3 Storage Nodes




System Cost  
3 year TCA  
**\$3,835,481**

**561** Trade Completions  
per second (TPS)  
**\$6,837** per TPS

Based on IBM internal tests comparing the IBM system with a comparably priced, comparably tuned competitor configuration (version available as of 01/01/2013) executing a materially identical online transaction processing workload in a controlled laboratory environment. Tests measured Trade Completion throughput rates to execute identical SQL query workloads. More Trade Completions is indicated by higher Trade Completions/second. Competitor configuration: ¼ Unit (usable uncompressed capacity = 9.5TB) including competitor recommended software options and features. Results may not be typical and will vary based on actual workload, configuration, applications, queries and other variables in a production environment. Users of this document should verify the applicable data for their specific environment.


# POWER8 Turbo

12 Core Sockets



- Configure via ASMI menu
- Potential increase in processor frequency: ~ 11+%
- Requires firmware 810 ( POWER8 support)
- CPW & rPerf measured using nominal





# PCIe Adapter

# PCIe Slots - High Level

	4U		2U	
	1S 4U	2S 4U	1S 2U	2S 2U
<b>Total PCIe slots (all hot swap)</b>	<b>7</b>	<b>11</b>	<b>6</b>	<b>9</b>
<b>Required* LAN adapter (available for client use)</b>	<b>1</b>	<b>1</b>	<b>1</b>	<b>1</b>
<b>PCIe slots after required* LAN adapter</b>	<b>6</b>	<b>10</b>	<b>5</b>	<b>8</b>
<b>However if use high performance, expanded function backplane</b>	<b>-1</b>	<b>-1</b>	<b>-1</b>	<b>-1</b>
<b>PCIe slots after required* LAN and if using high performance backplane</b>	<b>5</b>	<b>9</b>	<b>4</b>	<b>7</b>

\* required for IBM Manufacturing

- **PCIe slots are all Gen3 slots**
- **2U are all low profile and 4U are all full high**
- **There is no PCI expansion drawer announced. There is an SOD.**



# Storage Backplanes

# Storage Backplanes

2U server example  
12 SAS bays base & split



## Backplanes provide

- High performance integrated SAS controller(s) built on IBM industry leading PCIe Gen3 SAS adapter technology
  - ❖ All include RAID 0, 1, 5, 6, 10. Plus hot spare capability
  - ❖ Split backplane option with zero-write cache controllers
  - ❖ Easy Tier® function\*
- 8-18\*\* SAS bays for 2.5-inch (SFF) HDD or SSD
- 6-8 SAS bays for 1.8-inch SSD\*
- One DVD bay
- Option for attaching one EXP24S drawer of HDD or SSD\*

\* With dual IOA, expanded function backplane with write cache (not IBM i)

\*\* number varies based on 2U/4U and backplane option selected

# POWER8 4U Storage Backplane Options

Must select one →

#EJ0N

#EJ0N+EJ0S

#EJ0P \*

**12 SFF SAS bays**  
**1 SAS controller**  
**No write cache**  
**DVD bay**

*Staged  
availability*

**6+6 SFF SAS bays**  
**2 SAS controllers**  
**No write cache**  
**DVD bay**

**18 SFF SAS bays**  
**Dual SAS controllers**  
**7.2\*\*GB cache**  
**DVD bay**  
**8-bay SSD cage\*\*\***  
**External SAS ports**

Note that no HH tape bay is provided – different than POWER7 720/740

AIX / IBM i / Linux	yes		yes		yes
<i>Easy Tier Function</i>	no		no		yes
JBOD	yes		yes		no
RAID 0	yes		yes		yes
RAID 10	yes		Yes		yes
RAID 5/6	Yes		Yes		yes
Split backplane	no		yes		no

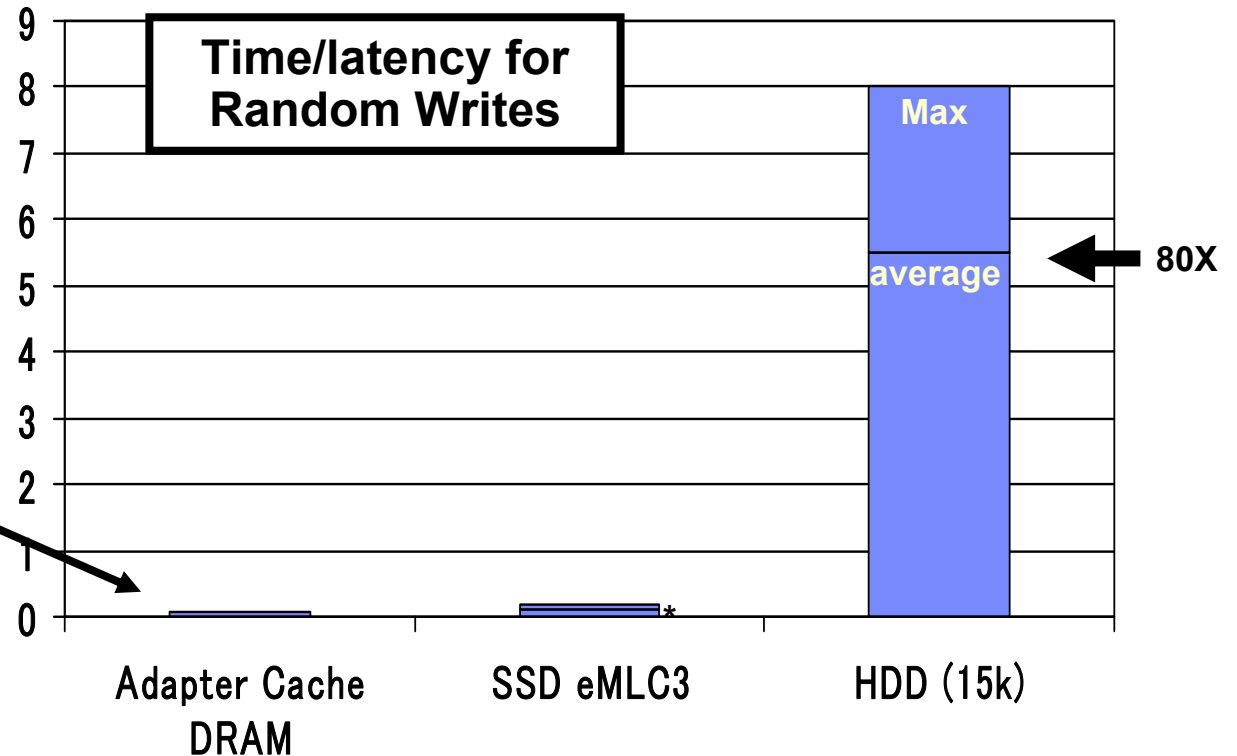
\* Uses one x8 PCIe slot (space taken up by #EJ0Z SAS ports/cabling – EJ0Z is mandatory with EJ0P)

\*\* 1.8GB physical write cache provides up to 7.2GB effectively with compression

\*\*\* 8-bay 1.8-inch SSD cage #EJTM NOTE: Not available on mdl 41A, Required on mdl 42A with #EJ0P

# Controller Write Cache Value

Milliseconds



**Controller cache is**

- Up to 1.5 to 2.5X faster than SSD
- Up to **80X** faster than HDD average (Up to 115 X HDD max )

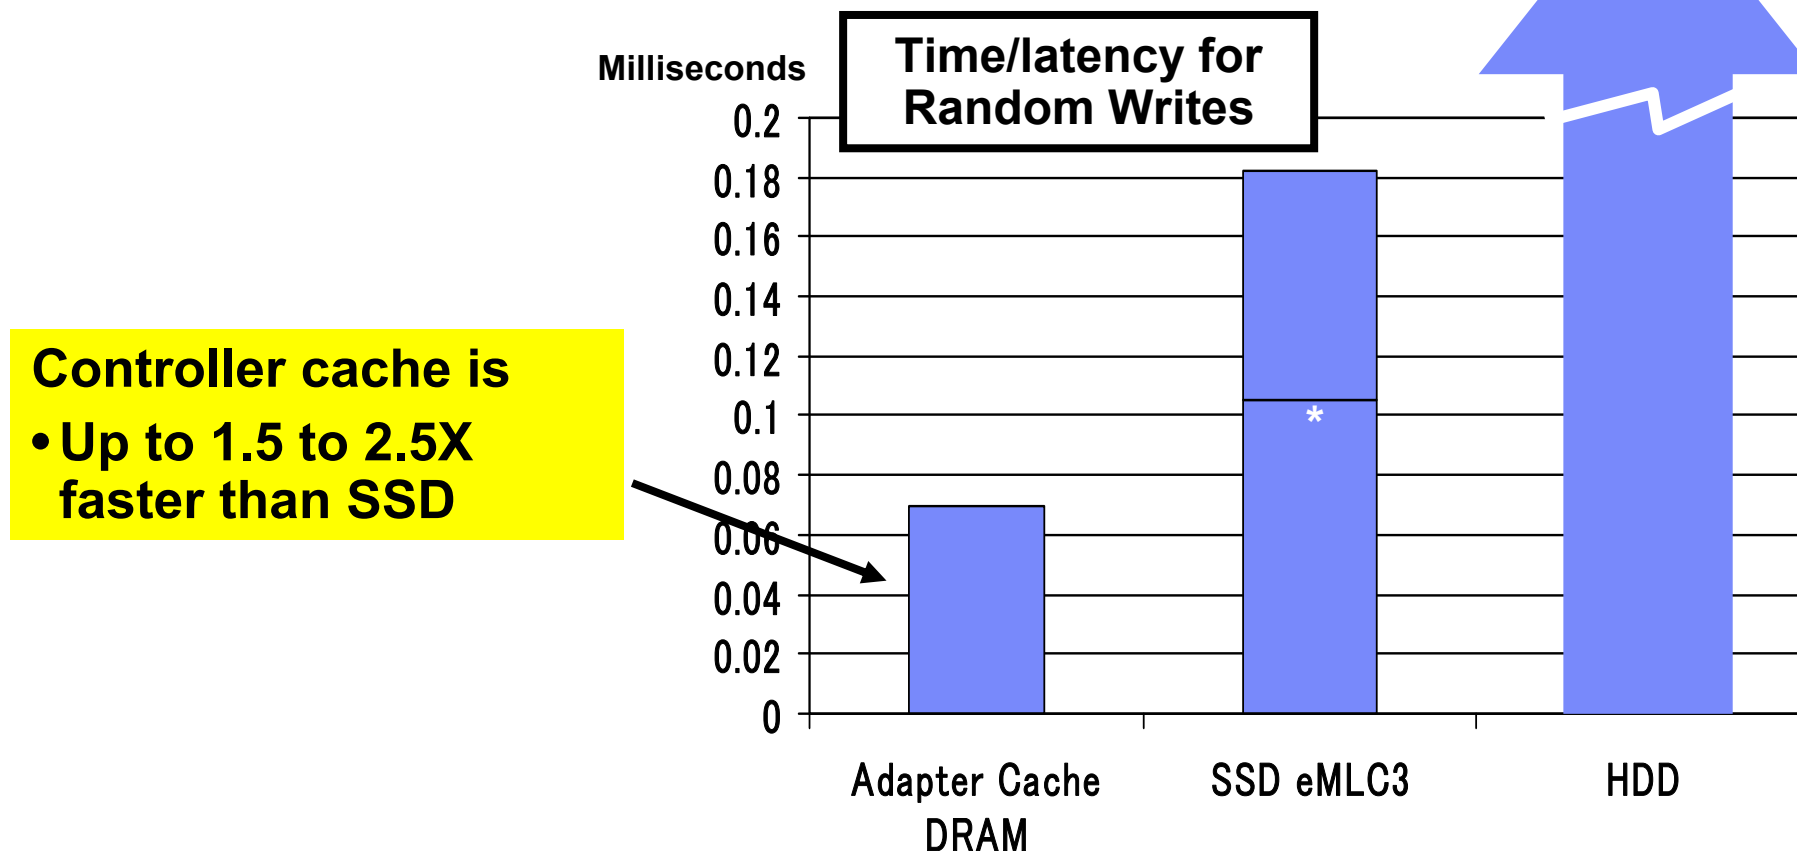
Value depends on the amount/percentage of Writes

Valuable for SSD, even more valuable for HDD

Can even improve “reads” if application using recently written data still sitting in cache

Graph is a simplification. All performance discussions start with the words “it depends”. HDD 15k Max ms shows typical maximum rotational delay and arm movement. 10k HDD is about 1 ms slower. Non-random work will have better HDD measurements. Actual HDD performance varies from HDD to HDD. Adapter write cache can also speed reads, but value of write cache for reads is highly application dependent. The bottom line (\*) of the SSD is obtained when the DRAM write cache integrated into SSD can handle the write and with a low queue depth. The higher SSD value is with a higher queue depth and/or when the SSD write cache is not able to keep up with a stream of writes and the write is occurring to the NAND flash memory.

# Controller Write Cache Value



Value depends on the amount/percentage of Writes

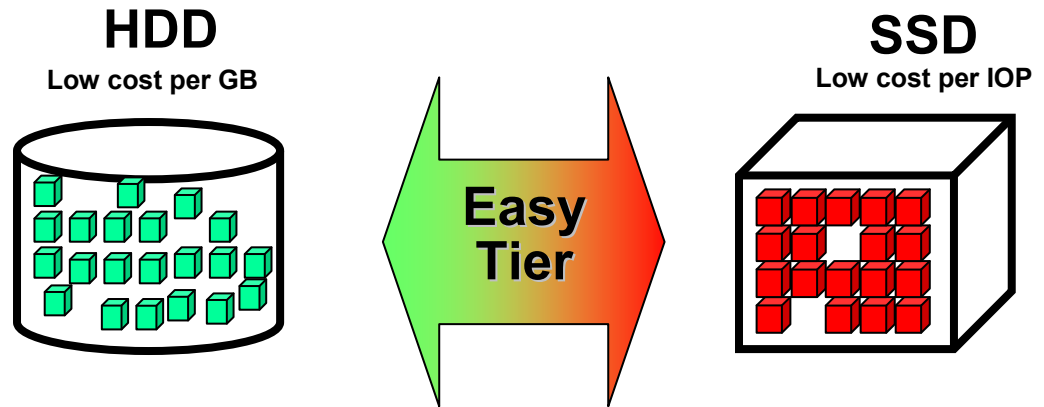
Valuable for SSD, even more valuable for HDD

Can even improve “reads” if application using recently written data still sitting in cache

Graph is a simplification. All performance discussions start with the words “it depends”. HDD 15k Max ms shows typical maximum rotational delay and arm movement. 10k HDD is about 1 ms slower. Non-random work will have better HDD measurements. Actual HDD performance varies from HDD to HDD. Adapter write cache can also speed reads, but value of write cache for reads is highly application dependent. The bottom line (\*) of the SSD is obtained when the DRAM write cache integrated into SSD can handle the write and with a low queue depth. The higher SSD value is with a higher queue depth and/or when the SSD write cache is not able to keep up with a stream of writes and the write is occurring to the NAND flash memory.

## Easy Tier Function Advantage

Optional function with expanded-function, high-performance storage backplane




- Automatically moves high activity (hot) data to SSD and low activity (cold) data to HDD
- Function handled totally by POWER8 integrated SAS controllers. No application coding. No SAN, just internal SAS drives.
- For AIX/Linux/VIOS. Just configure as a new type RAID array
- IBM i essentially already has same kind of hot/cold function in OS for all SAS adapters, but could this function via VIOS



# POWER8 Memory Card

	4U		2U	
	1S 4U	2S 4U	1S 2U	2S 2U
Max number DIMMs	8	16	8	16
Min number DIMMs	1	2	1	2
Max GB memory	512GB	1TB	512GB	1TB



## POWER8 Memory Card

- Capacity: 16GB / 32 GB / 64 GB
- 1600 MHz
- Memory Sparing - RAS improvement
- 8 Cards per socket (Scale-Out Systems)
- Low Profile

# Power System Software

Power **VM**

Power **KVM**

Power **VC**

Power **VP**

Power **HA**

Power **SC**



[ibm.com/systems/software](http://ibm.com/systems/software)

# PowerKVM

## Open Virtualization for POWER8 Linux Scale-out Servers

Kernel-Based Virtual Machine(KVM)  
virtualization targets new Linux  
workloads

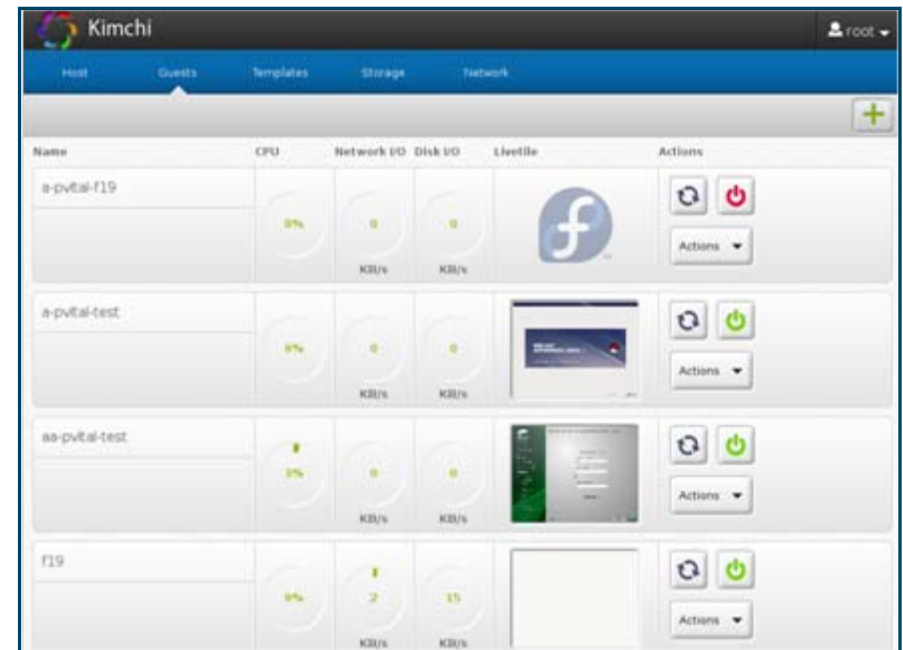
Provides simplicity and familiarity for  
VMware and KVM Intel Linux  
admins

Accelerates adoption of Power Linux  
with Linux oriented virtualization  
offering

Allows cloud providers to integrate  
Power Linux into OpenStack  
environments

Managed by PowerVC or open source  
tools such as Kimchi

Exploits POWER8 micro-threading and  
supports NFS and iSCSI



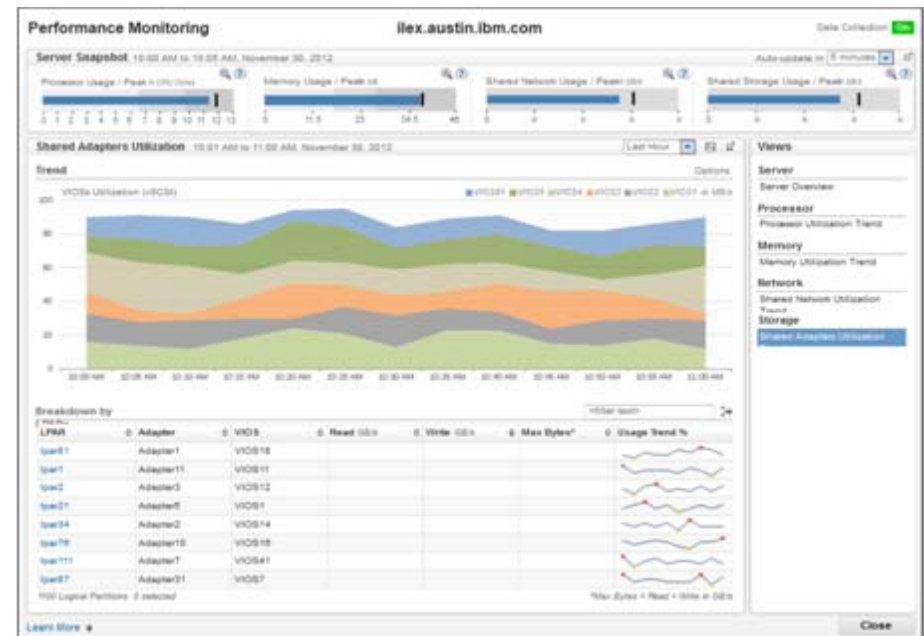
PowerKVM



# PowerVM

## Virtualization without Limits

- Major enhancements to user experience simplify administration and management through HMC
- No touch management and one touch deployment simplifies administration and accelerates deployment of VIOS
- System and partition templates follow best practices for repeatable deployments with low risk
- New performance monitoring and helps to simplify performance and capacity management
- Support for latest POWER8 capabilities
- SRIOV NIC support for enhanced network virtualization



PowerVM



# PowerVP

## Virtualization Performance Intelligence

Provides performance information to help optimize virtualized systems

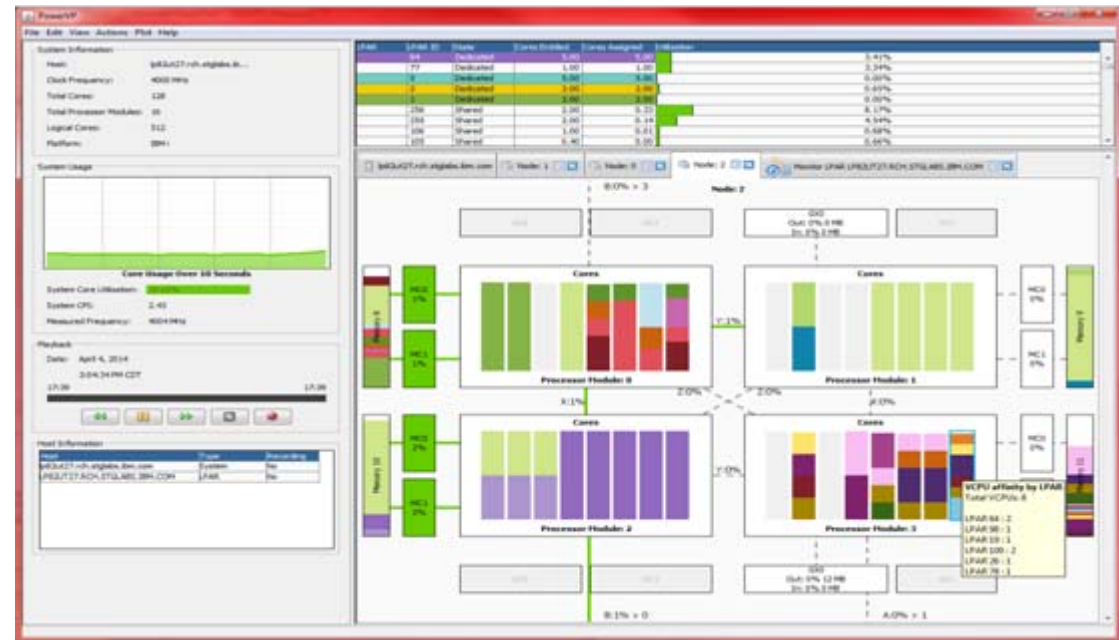
Understand the mapping of virtualized workloads to physical hardware

Accelerates identifying system performance bottlenecks

Improved memory affinity information allows optimization for virtualized workloads

Improved shared processor affinity information allows optimization of shared processor configurations

Included in PowerVM Enterprise Edition



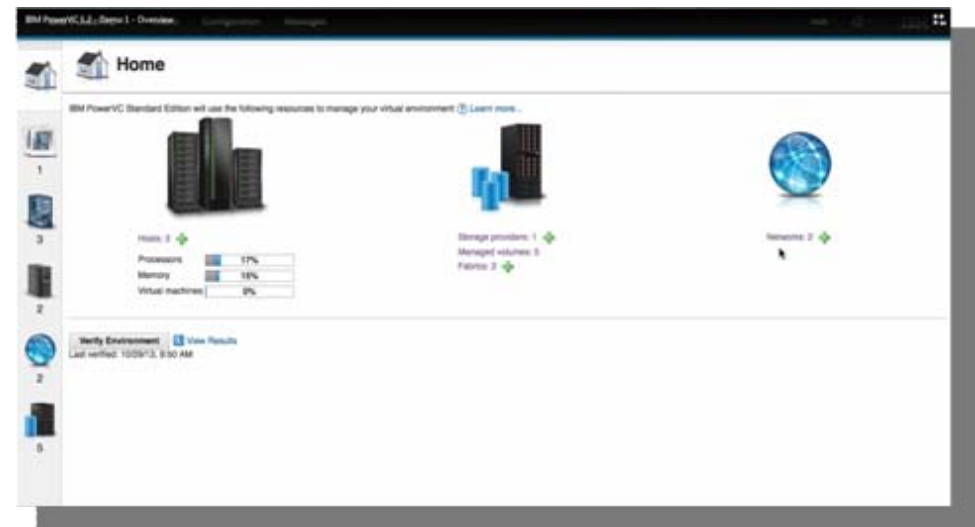
PowerVP

# IBM Systems Director Roadmap



# PowerVC Virtualization Center

- Leadership solution for managing PowerVM and PowerKVM
- Virtual Image management and deployment
- Resource pooling and dynamic VM placement
- New support for **PowerKVM**
- New Shared Storage Pool support enables support for a broader set of storage and clients with EMC and Hitachi
- Leverage the performance of new POWER8 hardware
- Expanded scale to manage larger enterprise environments



PowerVC



# SmartCloud Entry

Use self-service interface  
for simpler cloud delivery

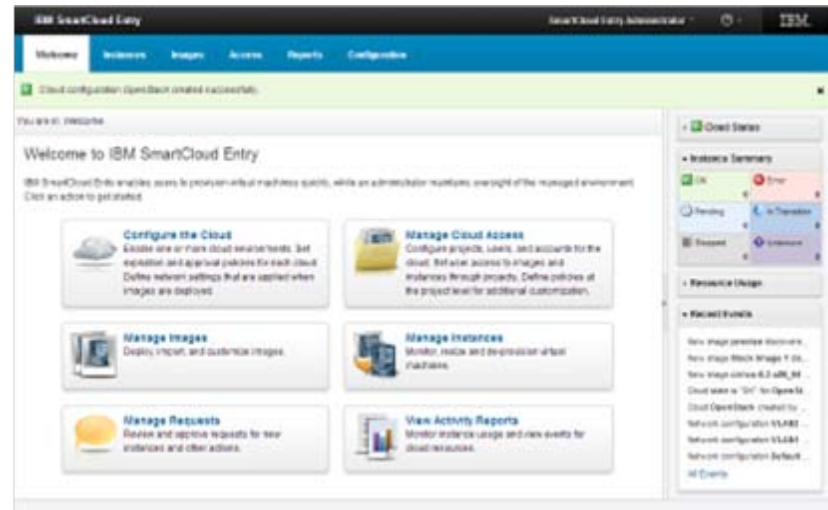

Build on community supported  
OpenStack for maximum  
flexibility

Accelerate infrastructure delivery  
and improve quality with  
automation

Enforce resource consumption  
oversight and accountability

Save time and resources on  
manual tasks with process  
automation

Support hypervisors across  
Power,  
x86 and z



Support for PowerVC  
enhancements

Simpler installation with Chef  
scripts

Expanded OpenStack driver  
support with components such  
as Cinder

VM placement policies such as  
packing, striping, load balance

More hypervisor choices,  
including PowerKVM and z/VM

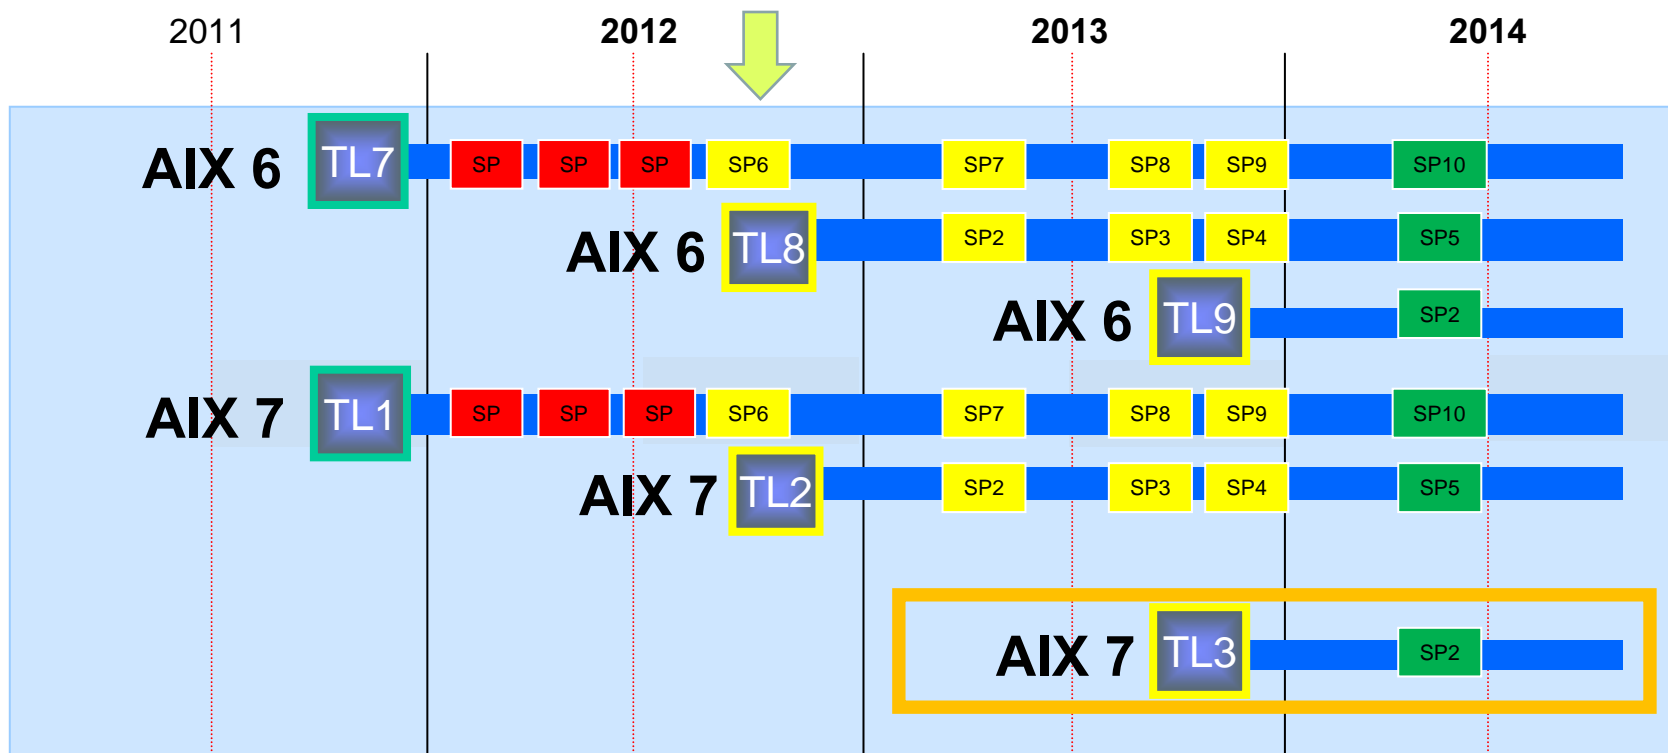


# Simplified UNIX Client Transition to POWER8



- **Goal: NO AIX UPDATE REQUIRED** to live relocate to a **POWER8** system with **Virtual I/O** only in **P7, P6+ or P6** mode
  
- **3 flexible options for AIX clients migrating to POWER8**
  - 1. Simply relocate to POWER8 with LPM and Virtual I/O**
    - AIX 6.1 or AIX 7.1 with supported AIX TL and Service Pack
    - LPAR is Live Partition Mobility using virtual I/O
  
  - 2. Leverage full I/O support on POWER8**
    - AIX 6.1 or 7.1 supported TL and latest Service Pack
    - Exploits POWER6/6+ or POWER7 modes
  
  - 3. Fully exploit latest POWER8 mode capabilities**
    - AIX 7.1 with TL 3 and SP2
    - Leverages latest POWER8 enhancements such as increased threads

# Planned POWER8 support built into current AIX TLs



**SPx - Full I/O & POWER8 mode**

**SPx - Full I/O support on POWER8**

**SPx - Can live relocate to POWER8 with Virtual I/O**

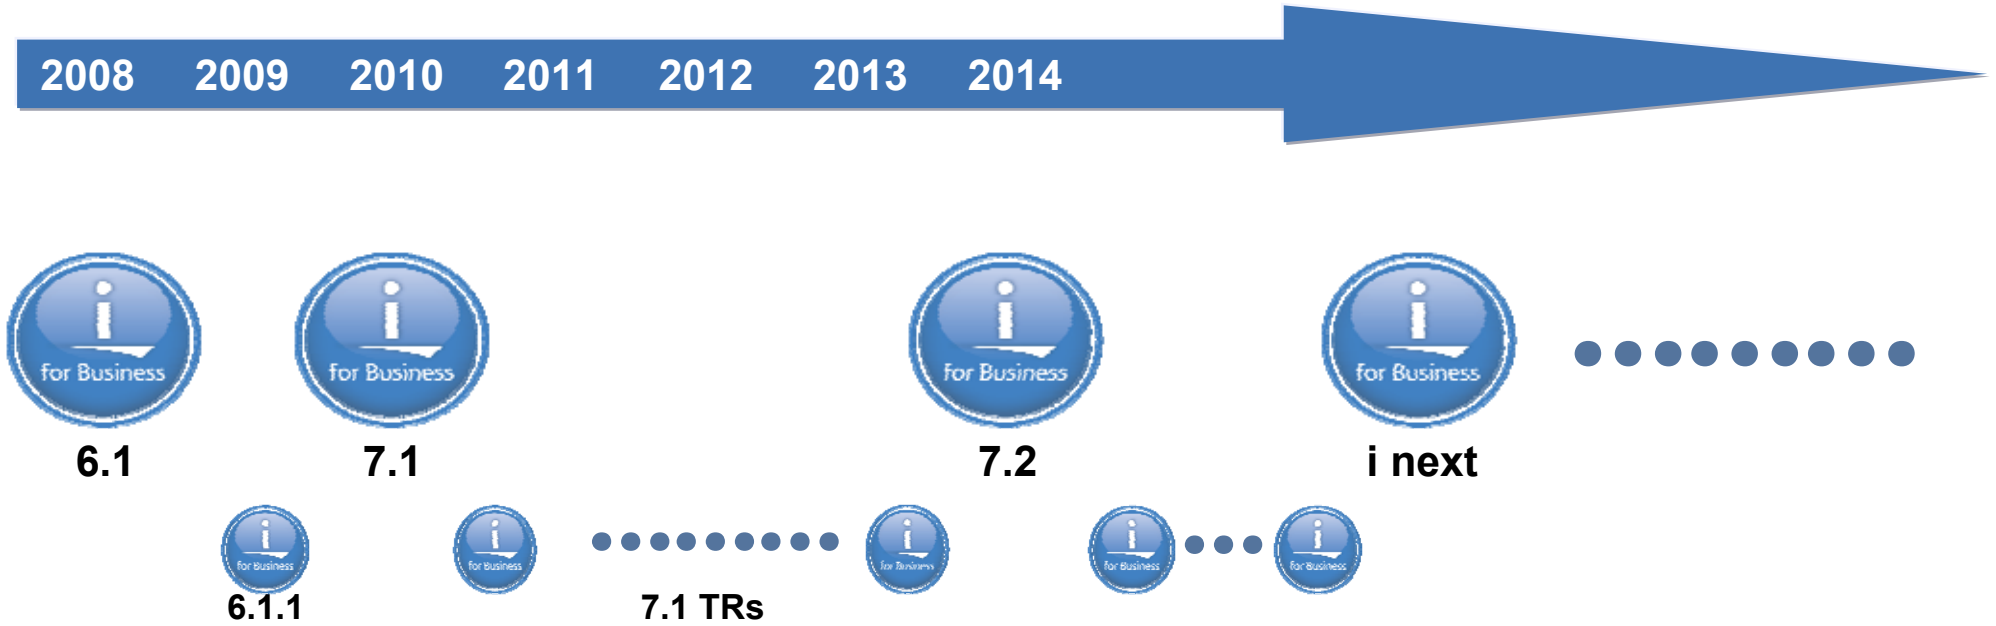
**SPx - Not supported on POWER8**



# Linux Plans

	1H / 2014	2H / 2014
RHEL6	RHEL 6.5 P7 Mode in P8	RHEL 6.6 P7 mode in P8
RHEL 7	RHEL 7.0 POWER8 Support	
SLES 11	SLES 11 + SP3 P7 Mode in P8	
SLES 12		SLES 12 (LE) POWER8 Support
Ubuntu (LE)	14.04.00/01 P8 Support	14.04.00/02
Debian	LE Introduction POWER8 Support	LE Update
PowerVM	2.2.3.3	2.2.3.X

# IBM i Roadmap



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

# IBM i System Support

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

Servers	IBM i 5.41	IBM i 6.1	IBM i 7.1	IBM i 7.2
POWER8			✓ 3	✓
POWER7/7+ PS700/701/702/730/704, Power 710, 720, 730, 740, 750, 760, 770, 780, 795, Pureflex p260/460		✓ 2	✓	✓
POWER6 JS12, 22, 23/43, 550* 560		✓	✓	✓ 4
POWER6 520, 550*, 570, 595	✓	✓	✓	
POWER5/5+ 515, 520, 525, 550, 570, 595	✓	✓	✓	
800, 810, 825, 870, 890	✓	✓		
270, 820, 830, 840	✓			

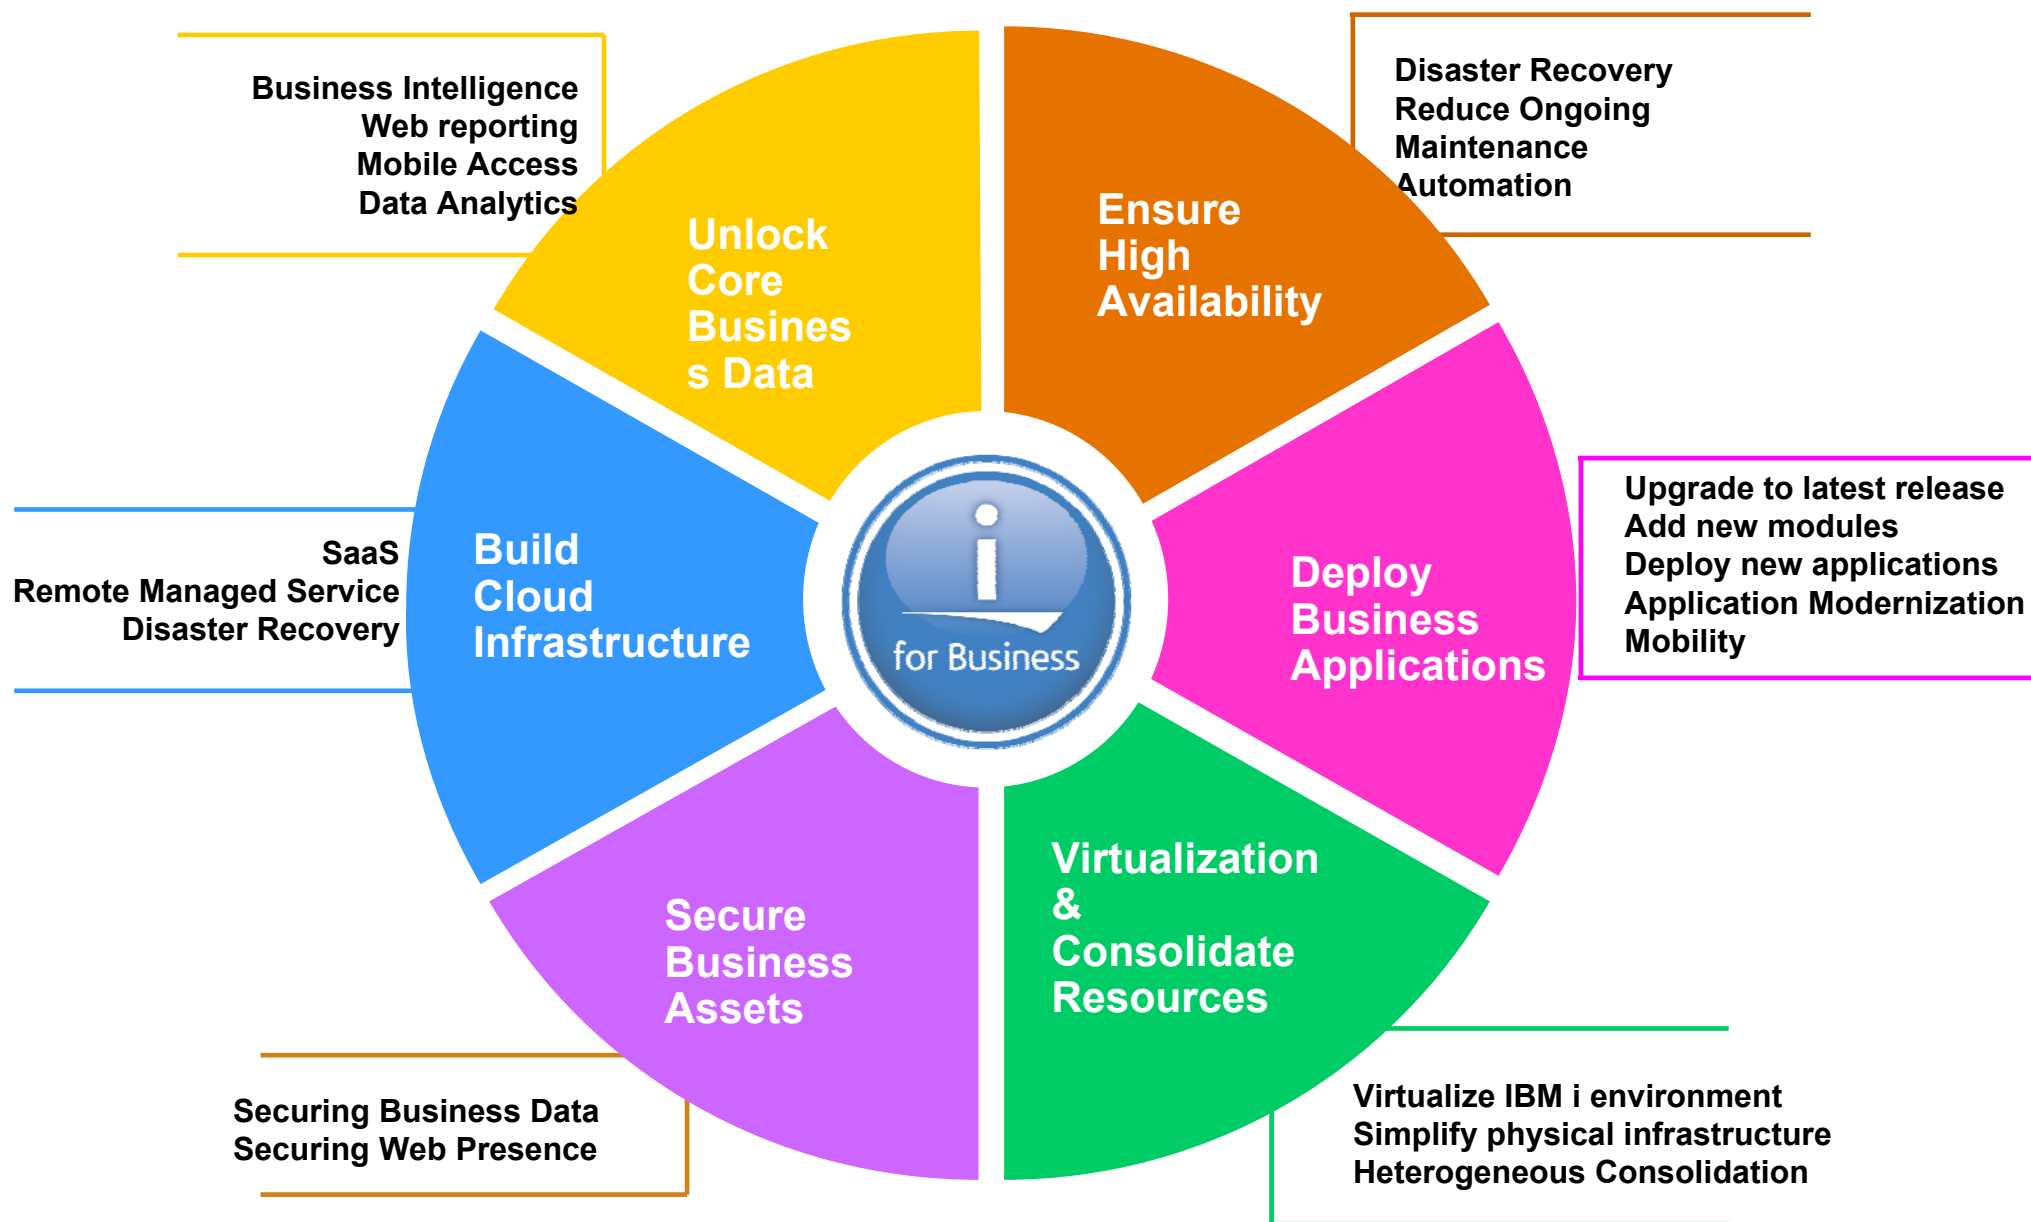
1 – IBM i V5R4 is no longer marketed or supported other than through extended service contracts

2 – POWER7+ 750/760 do not support native I/O. 6.1 in PureFlex must be client of 7.1 or later

3 – Requires Technology Refresh 8

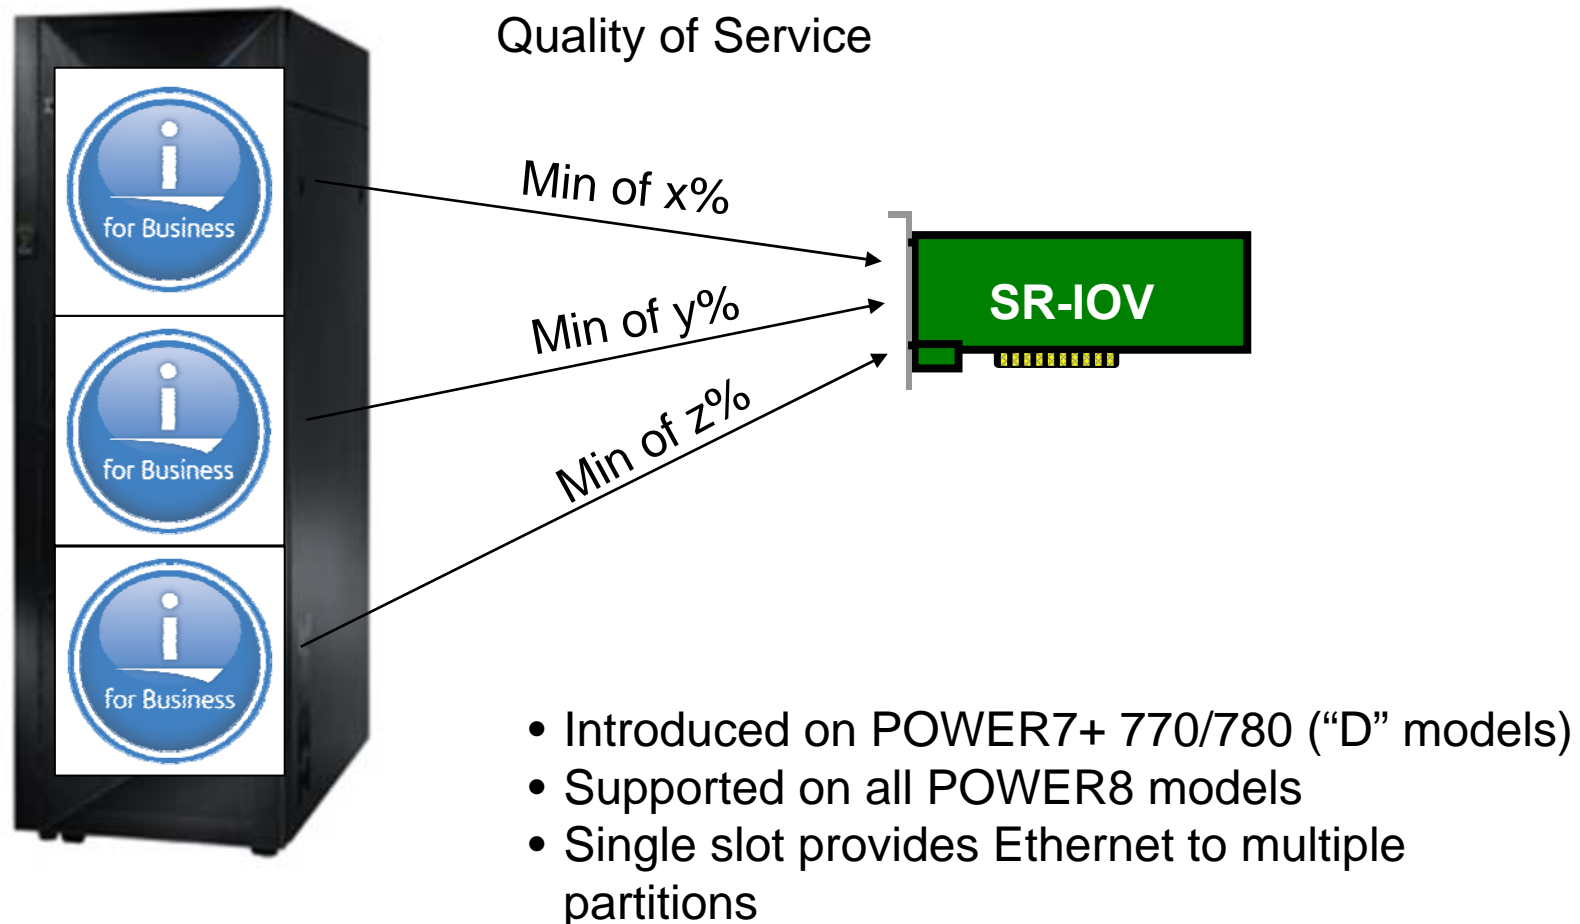
4 – no IOP or HSL support

# Top IBM i Client Projects



# SRIOV (Single Root I/O Virtualization) for Ethernet

- Simple virtualization with or without VIOS
- With quality of service controls to specify minimum bandwidth per partition







# IBM Navigator for i – new functions

## Performance & Usability

- New browsers
- My Favorites

## PTF management

- Installed PTF Navigation
- Load/apply PTF on single or group systems

## Message Queue monitor

- Based on the same CIM infrastructure
- Filter messages based on rules
- Trigger action to handle messages automatically

## System monitor

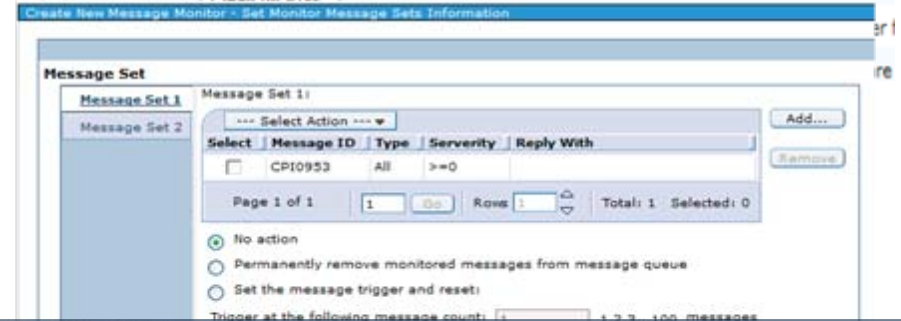
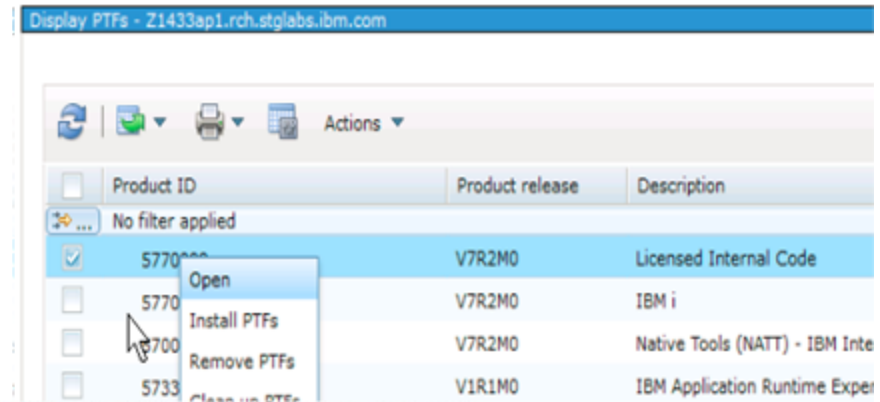
- User defined IBM i OS metrics monitor

## Graphical displays of performance trend

- User defined event automation

## Database

- Create new function and procedure
- DB Performance Metrics



# Collection Services & Performance Data Investigator – 7.2

## Significant enhancements to PDI

- New Collection Services Metrics
- Batch Model
- System Monitoring
- New PDI Perspectives
- Enhanced Left Hand Navigation

The screenshot displays the IBM Navigator for i Performance Data Investigator (PDI) interface. On the left, a navigation pane is highlighted with a red box, containing a tree view of various performance metrics and collection services. A callout box with an arrow points to this pane, containing the text: "Select package and perspective from left hand".

The main content area shows the "Investigate Data - Performance Data Investigator" window. It features a "Selection" section with "Name" (Database Locks Overview) and "Description" (This chart shows the database record lock contention time for all contributing jobs and tasks over time for the selected collection.). Below this is a "View List" section with "Database Locks Overview".

The "Collection" section includes a "Library" dropdown set to "QPFRDATA" and a "Collection Name" dropdown set to "Most Recent". A "Display" button and a "Search" input field are also present. Below these controls is a list of collection names and their timestamps:

Collection Name	Timestamp
Q141132820 (*CSFILE)	May 21, 2013 1:28:20 PM
Q141164450 (*CSFILE)	May 21, 2013 4:44:50 PM
Q142000003 (*CSFILE)	May 22, 2013 12:00:03 AM
Q143000002 (*CSFILE)	May 23, 2013 12:00:02 AM

On the right side of the interface, another navigation pane is visible, also highlighted with a red box. It contains a tree view of performance-related tasks and reports, including "Batch Model" and "Sizing".

# RPG Significantly Enhanced October 2013

## New Free Format RPG

- New Syntax; new style
- Modern behavior

## Conversion tool changes old RPG to RPG Free Format

- Arcad Converter
- Linoma Software

Rational Developer for i provided support in December 2013.

```

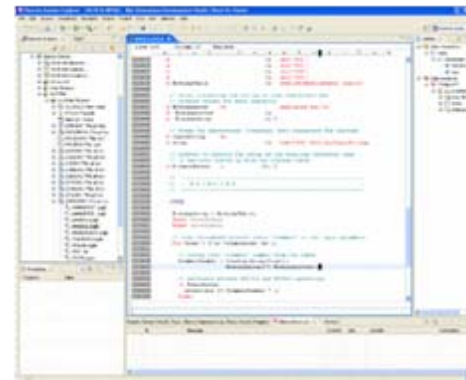

ctl-opt bnmdir('ACCRCV');

dcl-f custfile usage(*update);
dcl-ds custDs likerec(custRec);
dcl-f report printer;

read custfile custDs;
dow not %eof;
  if dueDate > %date(); // overdue?
    sendOverdueNotice();
    write reportFmt;
    exec sql insert :name, :duedate into
      mylib/myfile;
  endif;
  read custfile custDs;
enddo;
*inlr = '1';

dcl-proc sendOverdueNotice;
  sendInvoice (custDs : %date());
end-proc;

```




# IBM i Application Development Strategy

**RPG**



# PowerHA SystemMirror for i Express Edition – Hyperswap

- Provides ability to ‘instantly’ switch access from production IBM i DS8000 instance to remote DS8000
- Switch can happen automatically in case of DS8000 failure
- Switch can be manually triggered for planned maintenance
- Affinity can be defined so an LPM switch triggers a DS8000 switch
- DS8000 storage servers only
- IASP based replication not yet supported



# PowerHA 7.2 Enhancements

## SYSBAS Replication Enhancements

- Object Authority and Ownership can now be replicated with Administrative Domain
- Increase the 25,000 Administrative Domain monitored resource entry (MRE) limit to 45,000 MREs



## Reduced Downtime

- DSPASPSTS improvements for monitoring vary-on time
- Reduce downtime by shortening UID/GID processing time during vary-on

## Management Improvements

### Independent ASP Assignment

- Enables use of one partition to save multiple production cluster environments
- Allows attachment of an IASP to a partition not in the cluster device domain
- Only one IASP can be attached to the partition at a time

**IBM** has  
the technology  
you need Today



**IBM** is built  
for your future  
and Beyond



**Power Systems**

**The End...**

**Thanks !**