



## Quick Dip

1. Active Systems Optimizer
2. Active Memory Deduplication
3. Performance Advisors



Nigel Griffiths  
IBM Power Systems  
Advanced Technology Support, Europe

© 2012 IBM Corporation



## Active Systems Optimizer (ASO) and Dynamic Systems Optimizer (DSO)

© 2012 IBM Corporation


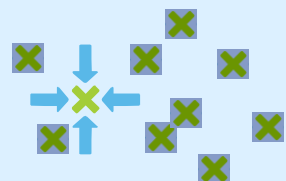
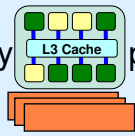
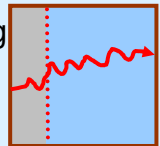

## ASO Pre-Requisites

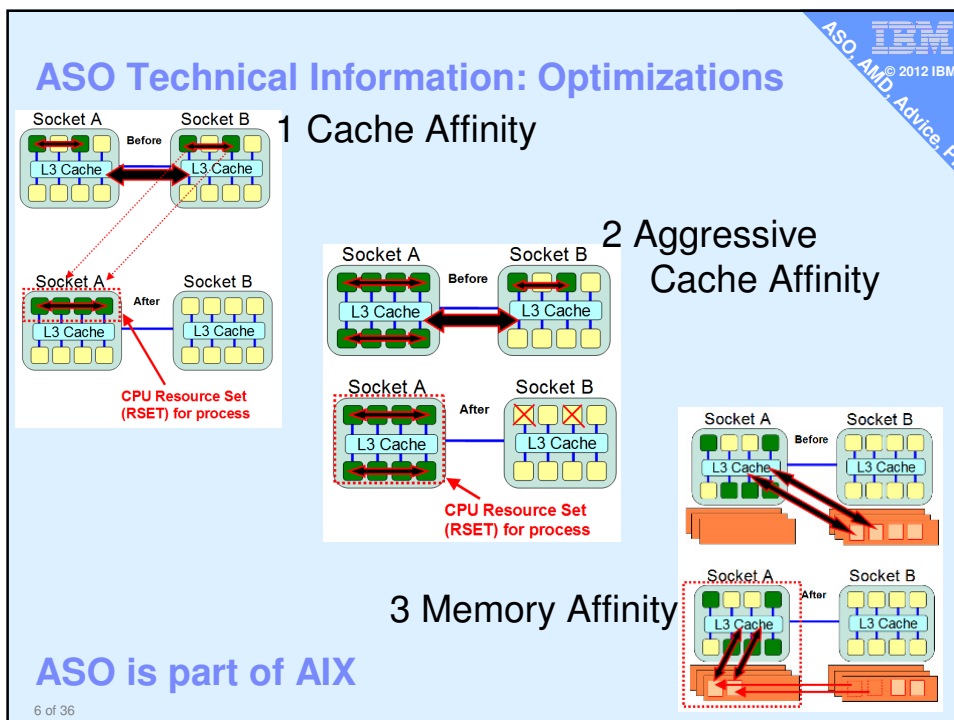
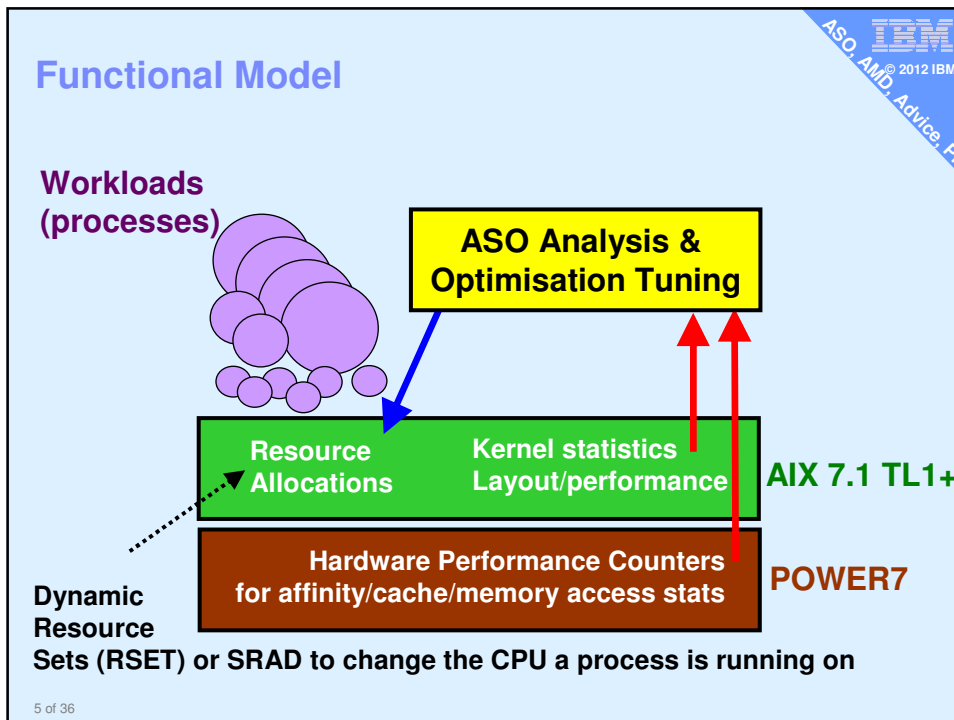
- Only AIX7.1 TL01+ on POWER7 or later



- Installed by default with AIX
  - Don't forget the mandatory Service Packs
- Warning: Any older AIX release or hardware!
  - **NOT supported ...**
  - May start but will do nothing


## ASO in Operation Overview

1. Once activated  requires no user involvement
2. Identifies & optimizes suitable workloads 
3. Improves cache & memory  performance
4. Performs pre- & post-optimization monitoring 
5. Hibernates when not busy 



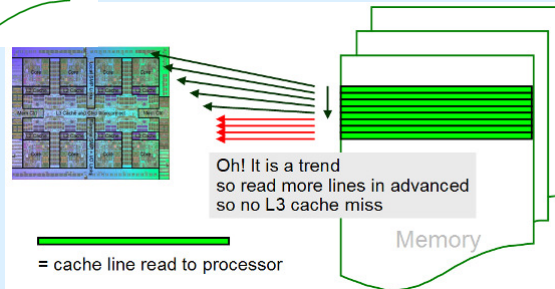
## DSO 4 Technical Information: Optimizations

### 4 Dynamic migration to Large Pages (16MB)



AIX default action:  
Merges 16 contiguous 4KB pages  
into one 64KB page or vice versa  
Lower HPT & higher lookup ratios


### 5 Data Stream Pre-fetch



DSO separate package (€\$£)  
using the ASO framework

## Technical Information: Eligible Workloads

- Multi-threaded workloads with periods of stability
  - CPU Utilization, Load and Latency must be stable for a period of time
- Minimum utilization = machine must be busy
  - Higher for aggressive cache optimization
- Minimum lifetime
  - 10 seconds (5 minutes for memory affinity)
- Not manually tuned
  - If too much of the system load is manually tuned, ASO hibernates
- Not explicitly marked as unoptimizable



8 of 36

## 1) and 2) ASO Start service and Activate

```
topas_nmon Host=purple7 Refresh=2 secs 16:15.15
Top-Processes-(262) Mode=4 [1=Basic 2=CPU 3=Perf 4=Size 5=I/O 6=Cmnds]
PID %CPU Size Res Res Res Char RAM Paging Command
Used KB Set Text Data I/O Use io other repage
.....
| 4981098 0.0 3712 3912 572 3340 0 0% 0 0 0 aso |

topas_nmon Host=purple7 Refresh=2 secs 08:58.53
Top-Processes-(125) Mode=1 [1=Basic 2=CPU 3=Perf 4=Size 5=I/O 6=Cmnds]
PID Parent User Proc Nice Pri Status proc-Flag Thrs Files Command
PID Id PID Id Group ority Foreground=F
4981098 2752544 root none 41 32 Running 0x00240103 14 336 aso
```

6 - 14 threads

From my experience on 16 CPU VM  
- Not seen aso use any CPU time  
= Less than 0.1% of one CPU

Much larger virtual machines might see some

9 of 36

## ASO in Practice

```
# startsrc -s aso
# asoo -op aso_active=1
# tail -f /var/log/aso/aso_process.log
```

Workloads running ....

expect ASO to monitor workloads for a few minutes

Note:

Log format is not documented but fairly readable

Some guess work in the following example logs

... your mileages will vary as every workload is different

10 of 36

## Performance Benefits

- Out of box performance boost for many workloads
  - Multi-threaded, memory / cache intensive, poor scaling
- Example workloads
  - SpecJBB – multi-threaded JVM benchmark
    - From 16 cores (2 sockets) up to 72 cores (9 sockets)
  - Daytrader – Websphere (java) + DB2
    - 16 / 32 cores (2 / 4 sockets)
  - Websphere Message Broker (WMB)
    - 16 cores
  - COPR – large DB2 benchmark
    - 64 cores (8 sockets)

Benchmark	SpecJBB	Daytrader	WMB	COPR
ASO	<b>Banned</b>			
Hand Tuned				

## Active System Optimizer Summary

1. AIX 7.1 TL01+ on POWER7 or later  
= “Set & forget”
2. Advanced Autonomic Affinity Tuning
  - Low CPU impact with zero negative effects
  - High performance impact
3. Particularly good for
  - Complex, multi-threaded, long running processes
  - Large CPU + RAM LPARs on larger machines


- Dynamic Systems Optimiser (DSO)
- Not free but uses same framework
  - AIX 6 TL8 and AIX 7 TL2

## Dynamic Platform Optimizer (DPO)

POWER Virtual Machines before DPO →



After DPO →



Cool right 😊


13 of 36

## Dynamic Platform Optimizer (DPO)

Firmware 760+ mandatory

- Currently

- POWER7+ Power 770/780 D
- POWER7+ lower End Power 710-740
- POWER7+ Mid-range Power 750/760
- POWER7 Power 795
- Possibly older POWER7 770/780 models ... have to "wait & see"



- Is my machine capable?

- Machine Properties → Capabilities →

Dynamic Platform Optimization Capable True

- This is a chargeable feature

- DPO requires AIX 6.1 TL8 or AIX 7.1 TL2

14 of 36

## Dynamic Platform Optimizer (DPO)



### HMC command: `optmem` & `lsmemopt`

1. `lsmemopt -m $MACHINE -o currscore curr_sys_score=37`
  2. `lsmemopt -m $MACHINE -o calcscore curr_sys_score=37,predicted_sys_score=92, .....`
  3. `optmem -m $MACHINE -o start -t affinity`  
... takes a lot of time here
  4. `optmem -m $MACHINE -o stop -t affinity`
- Can also DPO 1 LPAR or a set of LPARs or all except certain LPARs
  - Hint to get the LPARNAME use: `Issyscfg -r sys -F name`

15 of 36

IBM

# Active Memory Deduplication



Nigel Griffiths  
IBM Power Systems  
Advanced Technology Support, Europe


© 2012 IBM Corporation



## Power Systems Memory Deduplication



- **Redpaper** 98 pages (80 really)
- Pretty good content & easy reading
- Content
  - Concepts
  - Planning & Set-up
  - Monitoring commands
  - Tuning
  - Worked examples & Best Practice
  - Trouble shooting



17 of 36

## Pre-Requisites

1. **POWER7** only
2. PowerVM **Enterprise** Edition
  - HMC → Server → Properties → Capabilities: "Active Memory Capable" True
3. System **Firmware** level **740**
  - HMC → Update panel "EC Number"=01A\*740
  - **Power7xx C models introduced in Oct 2011 only**
4. **HMC** level **7.7.4**
5. **Operating Systems**
  - AIX Version 6: AIX 6.1 TL7, or later
  - AIX Version 7: AIX 7.1 TL1 SP1, or later
  - IBM i: 7.14 or 7.2, or later
  - SLES 11 SP2, or later and RHEL 6.2, or later
6. **Virtual I/O Server 2.1.1.10** (FP21) or later
  - Use VIOS ioslevel command
  - AMD uses VIOS CPU cycles via the Hypervisor code but not VIOS/AIX code = so no dependency

Nigel suggests: latest VIOS 2.2.1.3 = FP25 Oct 2011 or at least 2.2.something

18 of 36

## Pre-Requisites

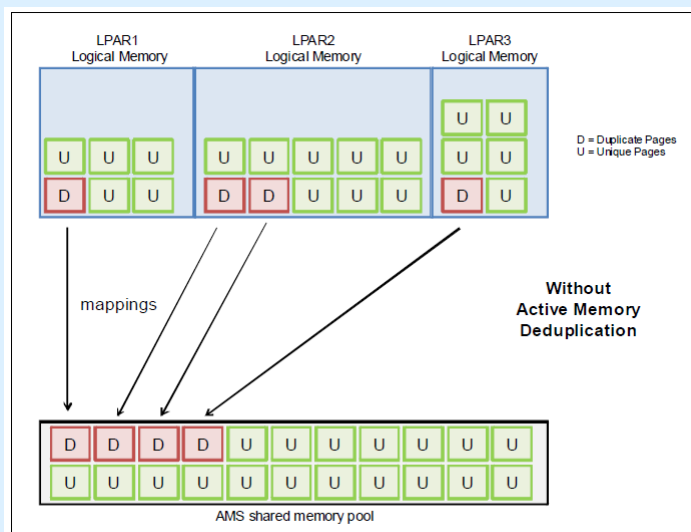


### 7. AMS virtual machine requirements

- Deduplication is ONLY for Active Memory Sharing virtual machines (LPARs), so AMS pre-reqs apply
- Shared CPU only (no dedicated CPUs)
- Shared I/O only (no dedicated adapters)
- No 16 MB pages (used by some HPC codes)
- LPAR needs restarting in AMS mode
- Only one pool = single set of co-operating VMs

19 of 36

## Old School regular Active Memory Sharing

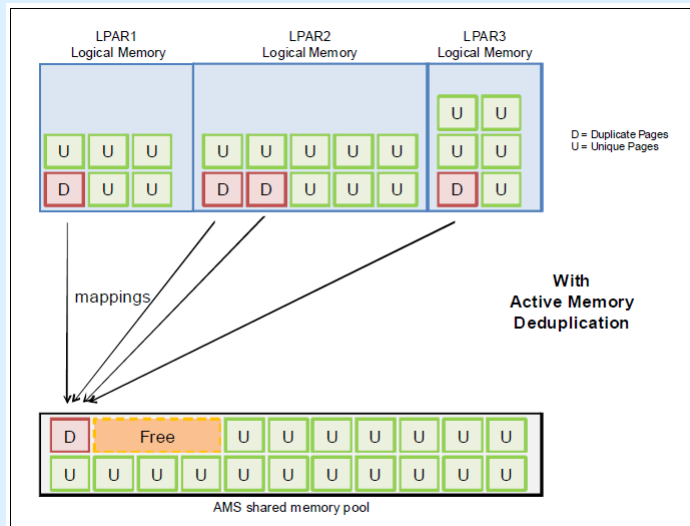


20 of 36

Figure 3-13 AMS shared memory pool without AMD enabled

Diagram from the Redbook

## Active Memory Sharing with new Deduplication



21 of 36

Figure 3-14 Identical memory pages mapped to a single physical memory page with Active Memory Duplication enabled

Diagram from the **Redbook**

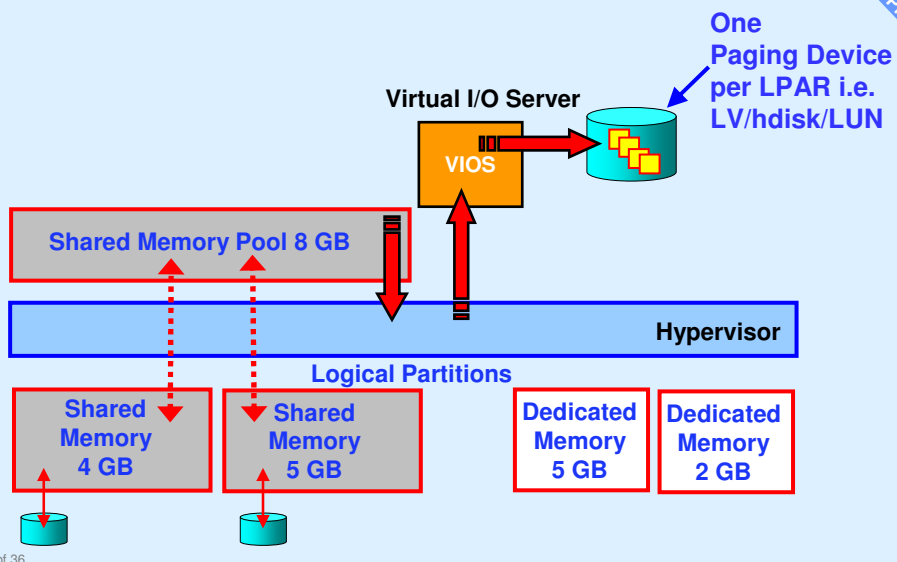
## Active Memory Sharing (AMS)

- Available since 2009
  - On POWER6 with VIOS 2.1
- **It is assumed you ALL know AMS well ... right?**
  - If not ask & we can run an AMS session again
  - Or read the **Redbook**
  - Or watch the AMS move at <http://tinyurl.com/AIXmovies>
- Three slide reminder ... next

22 of 36

# How is it set up?

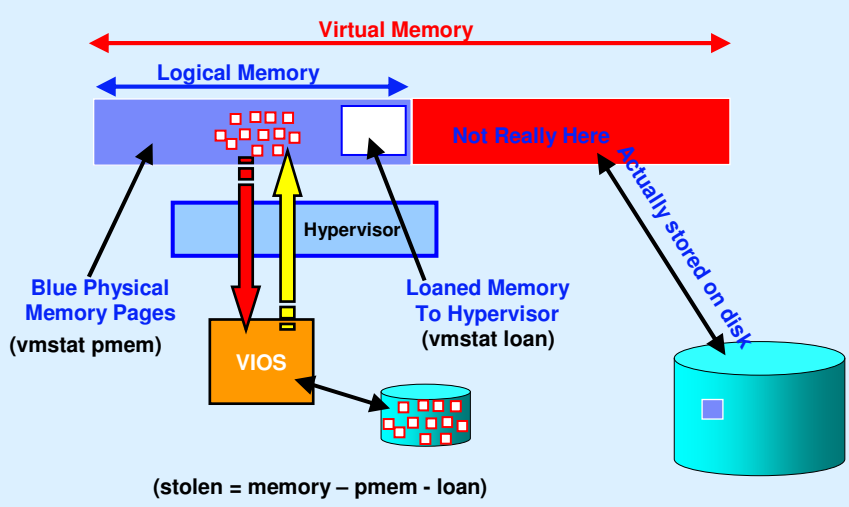
ASO, AMD, Advice, PI  
 IBM  
 © 2012 IBM



23 of 36

# Active Shared Virtual Memory (LPAR)

ASO, AMD, Advice, PI  
 IBM  
 © 2012 IBM



24 of 36

## Who is providing the Dedup function?

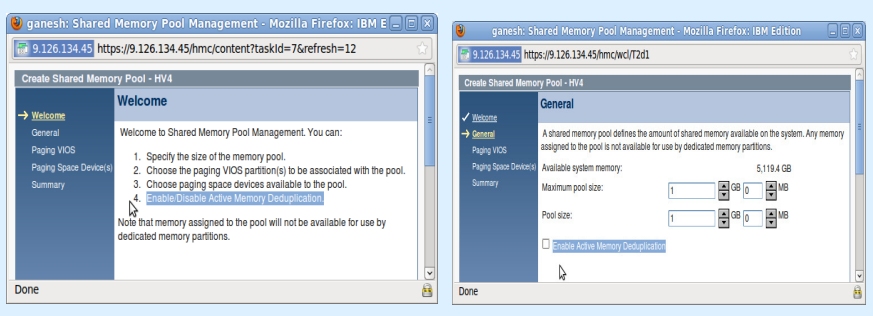
1. Hypervisor finding duplicates is not a high priority task
2. Hypervisor uses non-busy VIOS CPU cycles
3. Uses an in-memory cache of light weight fingerprints
4. On a match full byte check is made
5. If identical it is deduplicated & spare page on the free list
6. When a shared page is modified → a duplicate is made



## Good memory targets

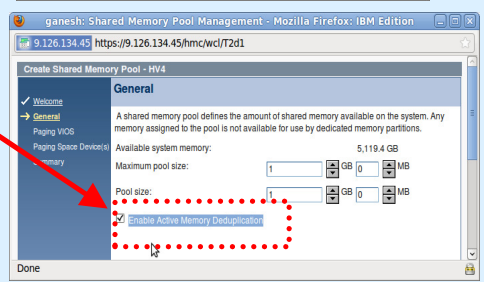
- Zero filled memory (perfect!)
- Partly used DB blocks
- Common read-only program code & static data
- Anything used by Java ☺

## A) Active Memory Sharing Pool Setup



**Deduplication flag**  
 "Enable Active Memory Duplication"

Pictures from the **Redbook**



## Summary: Active Memory Deduplication

- Largest pre-req = Oct 2011 C models firmware
- Good **Red**book(s) & simple to understand
- Very simple to implement (once AMS set-up)
- Very low CPU impact (VIOS idle time)
- High gains in memory use
- Set & forget
- Some internal benchmarks yield **very good results**
  - 46 Java VMs - dedicated RAM to AMS+AMD ~40% less RAM needed
  - 70 WAS VMs - dedicated RAM to AMS+AMD ~65% less RAM needed
  - Your mileage will vary & real workloads are more complex but significant memory reduction = cost is possible

27 of 36

IBM  
ASO, AMD, Advice, PI  
© 2012 IBM

IBM


# Power Advisors

**Would you like a AIX developer to look around your machines & look for common areas for improvement?**

**Well, now you can!**

© 2012 IBM Corporation

## Would you like some advice on:

- A. VIOS Performance
  - B. Java on POWER7 Performance
  - C. PowerVM (LPAR) Performance
- 

### Questions like

- Have I missed something?
- Am I missing a trick?
- Is it about right?
- Is there a bottleneck?
- Are there spare resources?

29 of 36

## Advisors on IBM DeveloperWorks

- <http://tinyurl.com/PowerAdvisor>  
– <https://www.ibm.com/developerworks/mydeveloperworks/wikis/home?lang=en#/wiki/Power%20Systems/page/Performance%20Tools>
- 1. Virtual I/O Server Performance Advisor
  - a. Downloadable
  - b. New “part” VIOS command
- 2. Java Performance Advisor
- 3. PowerVM Virtualization Performance Advisor  
(some times called the LPAR Performance Advisor)
  - a. LPAR mode = Runs on AIX
  - b. System mode = Runs on AIX but will include all LPARs

30 of 36

## Advisors on IBM DeveloperWorks/YouTube


- Briefly:
  - Popular - 1000's of downloads in a few months
  - Being productised !
  - Actively being development by the AIX development performance team
  - Two movies from me
    - <http://tinyurl.com/newAIXmovies>
  - Or on YouTube
    - <http://www.youtube.com/user/nigelargriffiths>

31 of 36

## Advisors on IBM DeveloperWorks



Download small file from DeveloperWorks



FTP a script to your targets

FTP back a .XML file



Run the Data Collector

Run the reporter which outputs to a browser

Note: LPAR Advisor needs SSH access to your HMC



## Virtualisation Advisor

- A bit more complicated
- From one LPAR you can get
  - Reports on other LPARs or
  - Whole machine reports
- You need to set-up ssh connection to your HMC
  - So it can extract LPAR config and performance data
- The it is pretty easy

## Java Advisor

- One Java process per run
- Mostly interesting for Java programmers  
but does give them hints on areas to look at

33 of 36

## Example

PowerVM™ VIOS Performance Advisor

The ratings and recommendations in the table below were chosen with the following information:  
 PartfileID: 1  
 Monitoring Start Time : 08/17 13:14:23  
 Monitoring Stop Time : 08/17 13:19:23 Duration : 5 min  
 VIOS sizing tool WLE (WorkLoad Estimator) link: <http://www-912.ibm.com/estimator>

SYSTEM - CONFIGURATION		VIOS - CPU						
Name	Value	Name	Measured Value	Recommended Value	First Observed	Last Observed	Risk 1:lowest 5:highest	Impact 1:lowest 5:highest
Processor Family	POWER7	CPU Capacity	4.0 ent	-	08/17 13:14:23	-	n/a	n/a
Server Model	IBM9117-MMC	CPU Consumption	avg:28.0% (cores:1.1) high:26.4% (cores:1.1)	-	-	-	n/a	n/a
Server Frequency	3.920 GHz	Processing Mode	Shared CPU (UnCapped)	-	08/17 13:14:23	-	n/a	n/a
Server - Online CPUs	16 cores	Variable Capacity Weight	120	120-255	08/17 13:14:23	-	1	5
Server - Maximum Supported CPUs	64 cores	Virtual Processors	4	-	08/17 13:14:23	-	n/a	n/a
VIOS Level	2.3.1.0	SMT Mode	SMT4	-	08/17 13:14:23	-	n/a	n/a

VIOS - I/O ACTIVITY		SYSTEM - SHARED PROCESSING POOL						
Name	Value	Name	Measured Value	Recommended Value	First Observed	Last Observed	Risk 1:lowest 5:highest	Impact 1:lowest 5:highest
Disk I/O Activity	avg: 1180 iops @ 111KB peak: 1217 iops @ 35KB	Shared Pool Monitoring	enabled	-	08/17 13:14:23	-	n/a	n/a
Network I/O Activity	[ avgSend: 9442 iops 0.5Mbps , avgRcv: 73811 iops 10.8Mbps ] [ peakSend: 9949 iops 0.6Mbps , peakRcv: 78453 iops 112.2Mbps ]	Shared Processing Pool Capacity	16.0 ent.	-	08/17 13:14:23	-	n/a	n/a
		Free CPU Capacity	avg_free:15.0 ent. lowest_free:14.9 ent.	-	-	-	n/a	n/a

VIOS - DISK ADAPTERS		VIOS - MEMORY					
Name	Measured Value	Recommended Value	First Observed	Last Observed	Risk 1:lowest 5:highest	Impact 1:lowest 5:highest	
FC Adapter Count	2	-	08/17 13:14:23	-	n/a	n/a	
FC Avg Iops	avg: 179 iops @ 3KB	-	08/17 13:14:23	08/17 13:19:23	n/a	n/a	
FC Idle Ports (Total)	idle	-	08/17 13:14:23	08/17 13:19:23	4	4	
FC Adapter Utilization	pass	-	-	-	n/a	n/a	
FC Port Speeds	running at speed	-	-	-	n/a	n/a	
Name	Measured Value	Recommended Value	First Observed	Last Observed	Risk 1:lowest 5:highest	Impact 1:lowest 5:highest	
Real Memory	4,000 GB	7,000 GB	08/17 13:14:23	-	1	5	
Available Memory	0.317 GB	1.5 GB Avail.	08/17 13:16:14	08/17 13:18:13	n/a	n/a	
Paging Rate	138.0 MB/s pg rate	No Paging	08/17 13:14:43	08/17 13:19:10	n/a	n/a	

■ C:/users/IBM/My documents/Advisors/part\_VIOS222/orangevios1\_vios\_advisor\_report.xml

## VIOS, Java, PowerVM Advisor Summary

1. Perfect cost:benefit ratio
2. Regular new releases
3. Simple to use and understand
4. Very low impact on systems
5. Valuable Advice to help you make improvements

My systems are small & simple

- Output for larger, complex, busy productions systems  
→ more interesting & worth some time investment

In the future: POWER/AIX support may ask you to run an Advisor during problem determinations as a quick check

- If that helps, next time, you can run it yourself

35 of 36



## What next?

- Give it a try today:
  - <http://tinyurl.com/PowerAdvisor>
  - <http://tinyurl.com/PowerAdvisor>
- VIOS Advisor & Java Advisors videos
  - <http://www.youtube.com/user/nigelargriffiths>
  - <http://www.youtube.com/user/nigelargriffiths>

### Other Performance info worth knowing about

- Java Performance on Power7 – Best Practice
  - <http://public.dhe.ibm.com/common/ssi/ecm/en/pow03066usen/POW03066USEN.PDF>
- Oracle Architecture and Tuning on AIX v2.20
  - <http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100883>
- Power7 Virtualization Best Practice Guide
  - [https://www.ibm.com/developerworks/wikis/download/attachments/53871915/P7\\_virtualization\\_bestpractice.doc](https://www.ibm.com/developerworks/wikis/download/attachments/53871915/P7_virtualization_bestpractice.doc)

36 of 36

