

Workload Partition (WPAR) & Versioned WPARs for running AIX 5.2 & AIX 5.3



Nigel Griffiths
IBM Power Systems
Advanced Technology Support, Europe

© 2012 IBM Corporation

Test:



Who is installing AIX 6 or AIX 7
as the default OS version?



Who has tried WPAR?



Workload Partitions = NOT Bleeding Edge

AIX 6

- Release November 2007 = 5.5+ years ago
- Two key functions:
 1. Workload Partitions - WPAR
 2. Roll based Access Control - RBAC

WPAR updated with every Technology Level (TL)

- AIX 6.1 TL8 is out → Q4 2012
- So we are on roughly WPAR version 9

Plus WPAR for AIX7


- With Versioned WPAR to run AIX 5.2 + now AIX 5.3

WPAR - at the movies!



1. WPAR Theory and Background
2. WPAR Manager Introduction
3. WPAR Mobility/Relocation
4. Create WPAR Simple
5. Create WPAR Detailed
6. Power5 to Power6 Mobility
7. WPAR Full Priority
8. WPAR Command Line
9. Compare Global and WPAR Environment
10. WPAR Backup and Cloning
11. Faster Relocation
12. Static Relocation
13. Director based WPAR Mgr
14. Versioned WPAR

Total = 3 Hour Demo

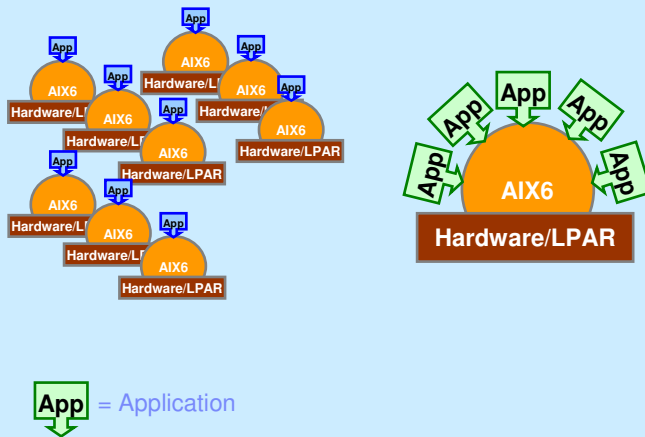


→ <http://tinyurl.com/newAIXmovies>
→ YouTube.com/user/nigelargriffiths

IBM Redbook
Exploiting IBM AIX Workload Partitions
www.redbooks.ibm.com **390 pages**


Workload Partitions (WPAR) Concept

Consolidation of isolated workloads with a single AIX instance




Workload Partitions (WPAR) Concept

Consolidation of isolated workloads with a single AIX instance

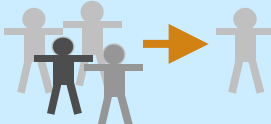







Live Application Mobility:
move WPAR between AIX instances
without restarting the WPAR applications



IBM
© 2012 IBM
7
Workload Partitions

Six Reasons for Workload Partitions (WPAR)

- 1) Reduced AIX System Administration
 
- 2) Application Encapsulation, monitoring and control
 
- 3) Rapid environment creation of a new application
 
- 4) Separated System Admin/Security at applications level
 
- 5) Live Application Mobility
Simple to move an application to a different machine for load balancing & evacuation
 
- 6) Reduced memory Use
Minimum WPAR = 65MB
 

7 of 58

IBM
© 2012 IBM
8
Workload Partitions

AIX Workload Partitions - in a Nutshell

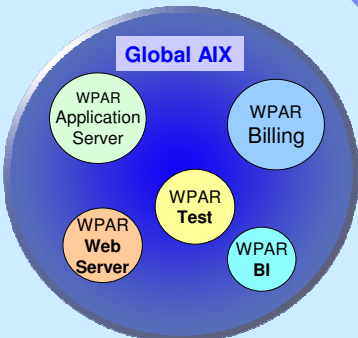
WPAR part of standard AIX 6 (2007)

- Pre-requisites = AIX 6=POWER4 to 7

1 Global AIX kernel sharing
CPU, RAM, I/O between WPARs

Each Workload Partition (WPAR)


- Independent WPAR start / stop
- Own network address (aliased)
- Isolated users, processes, IPC, cron, syslog
- Isolated filesystems: root, /tmp, /var & /home
- Optional read-only or separate /usr & /opt filesystems
- Full resource control – CPU, memory, paging space, disk



8 of 58

Workload Partitions Manager (WPAR)

- Cross System WPAR Management
- Live Application Mobility
- Automated, Policy-based Application Mobility




Functions

- Create & Remove
- Start & stop
- Hibernation & Cloning
- Monitoring & Reporting
- Manual Relocation
- Automated Relocation
- Policy driven change

AIX Systems
- Stand alone or LPARs

9 of 58

Two Workload Partitions Types



Application Workloads

- Isolate an individual application
- Light-weight, one process
 - can start further processes
- Created & started in seconds
- Starts when App created
- Automatically removed
 - when App stops
- Shares global file system
- Good for HPC
 - Long running applications

System Workloads

- Complete virtualized OS environment
 - Runs multiple services & applications
- Need to created - owns its filesystems
- Removed only when requested
- Like another AIX system
 - Own root user, users, and groups
 - Own services like inetd, cron, syslog
 - Can be stopped and restarted
- Does not share writable file systems with other workload partitions or the global level
- Integrated with RBAC - granular security controls
- Good for most purposes
 - Try this type first


10 of 58

Filesystems

Shared /usr + /opt



Global AIX Read only




Filesystems

Shared /usr + /opt



Global AIX Read only




Private /usr + /opt


Global AIX Read only



Copied to WPAR
at mkwpar time




Private copy
is Read/Write



If AIX updated, use syncwpar
to ensure they match

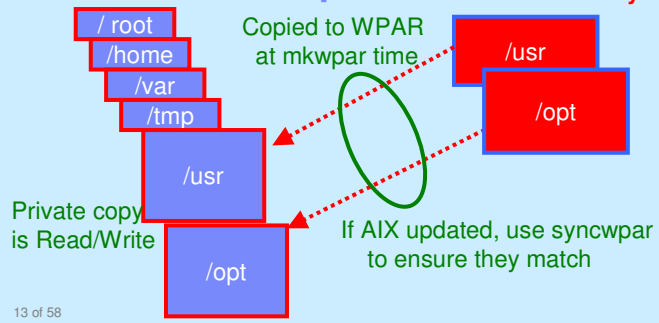
Filesystems

Shared /usr + /opt




You can mount a filesystem or NFS into your WPAR

Private /usr + /opt



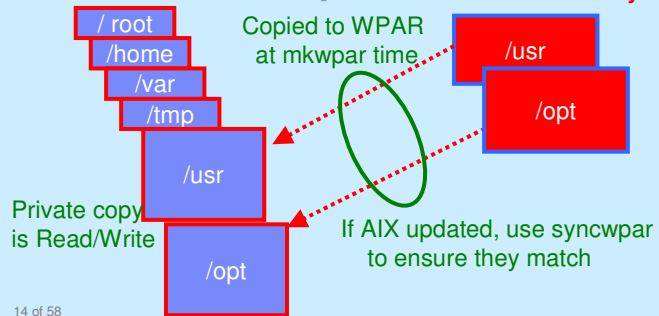
Filesystems

Shared /usr + /opt




You can mount a filesystem or NFS into your WPAR

Private /usr + /opt

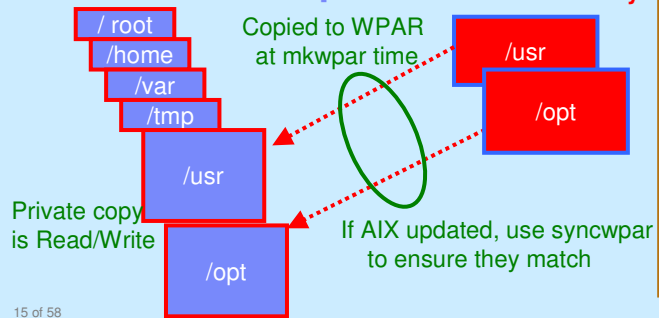


Filesystems

Shared /usr + /opt




Private /usr + /opt




- Mount options from within a WPAR
- `/another`
 - Local directory
 - Local filesystem
 - NFS mount point
 - NFS (over a GPFS)
 - LUN for Raw I/O
 - LUN +LVM/JFS2
 - New vSCSI from VIOS

Local filesystem on Local disk No Relocation




Local filesystem on Virtual Disk No Relocation




Mobile Disks


NFS with old NFS Relocation

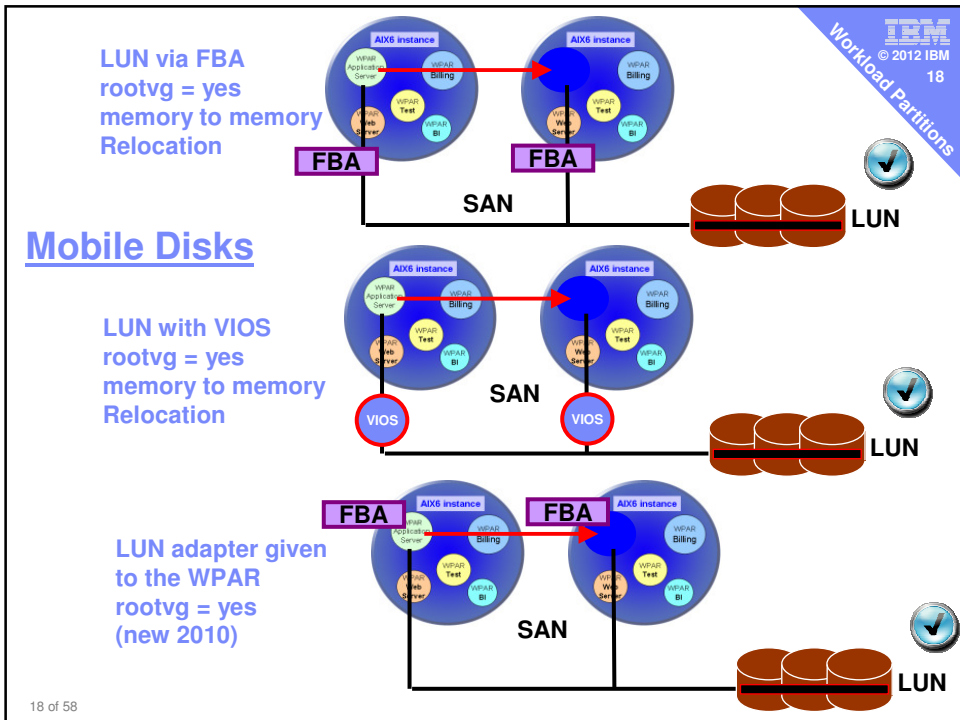
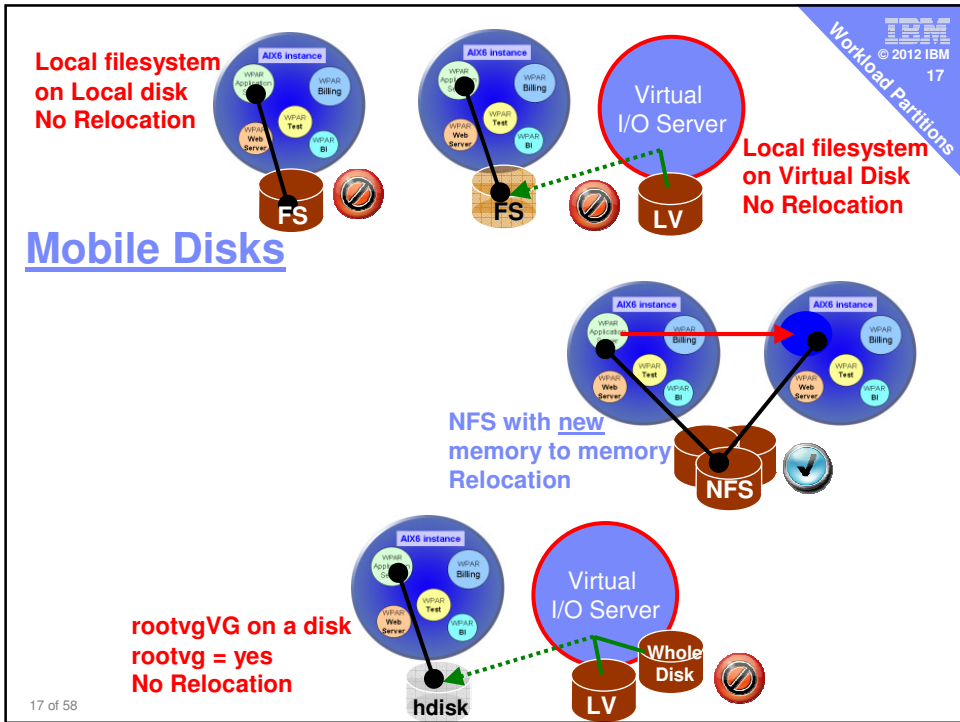


NFS with new memory to memory Relocation

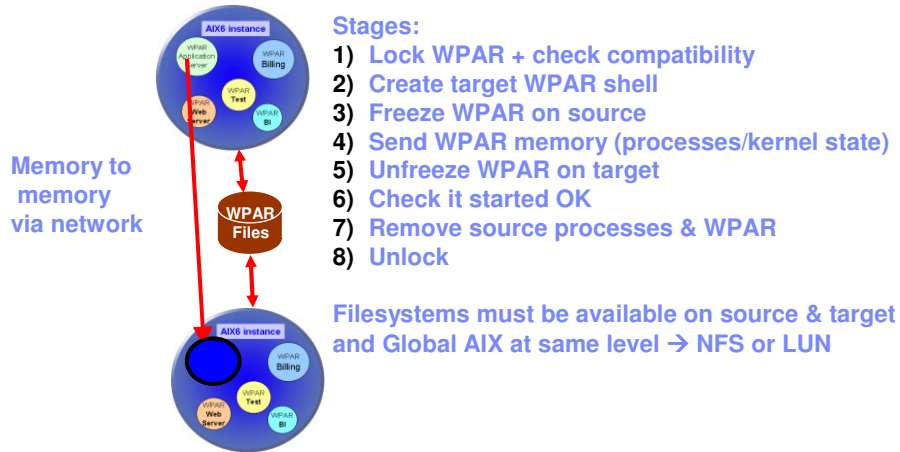


rootvgVG on a disk rootvg = yes No Relocation





Live Relocation (Mobility)

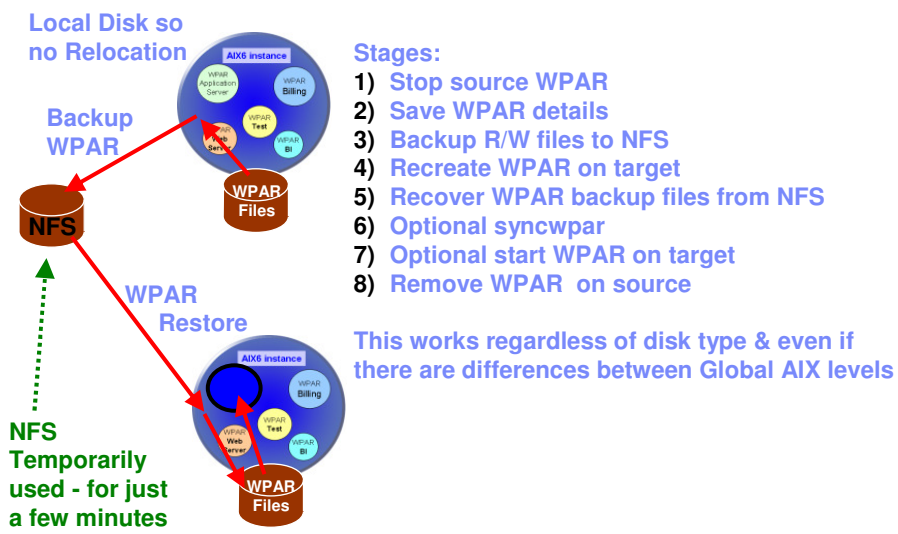


Stages:

- 1) Lock WPAR + check compatibility
- 2) Create target WPAR shell
- 3) Freeze WPAR on source
- 4) Send WPAR memory (processes/kernel state)
- 5) Unfreeze WPAR on target
- 6) Check it started OK
- 7) Remove source processes & WPAR
- 8) Unlock

Filesystems must be available on source & target and Global AIX at same level → NFS or LUN

Static Relocation (Mobility)



Stages:

- 1) Stop source WPAR
- 2) Save WPAR details
- 3) Backup R/W files to NFS
- 4) Recreate WPAR on target
- 5) Recover WPAR backup files from NFS
- 6) Optional syncwpar
- 7) Optional start WPAR on target
- 8) Remove WPAR on source

This works regardless of disk type & even if there are differences between Global AIX levels

Creating a WPAR is easy ~3 minutes

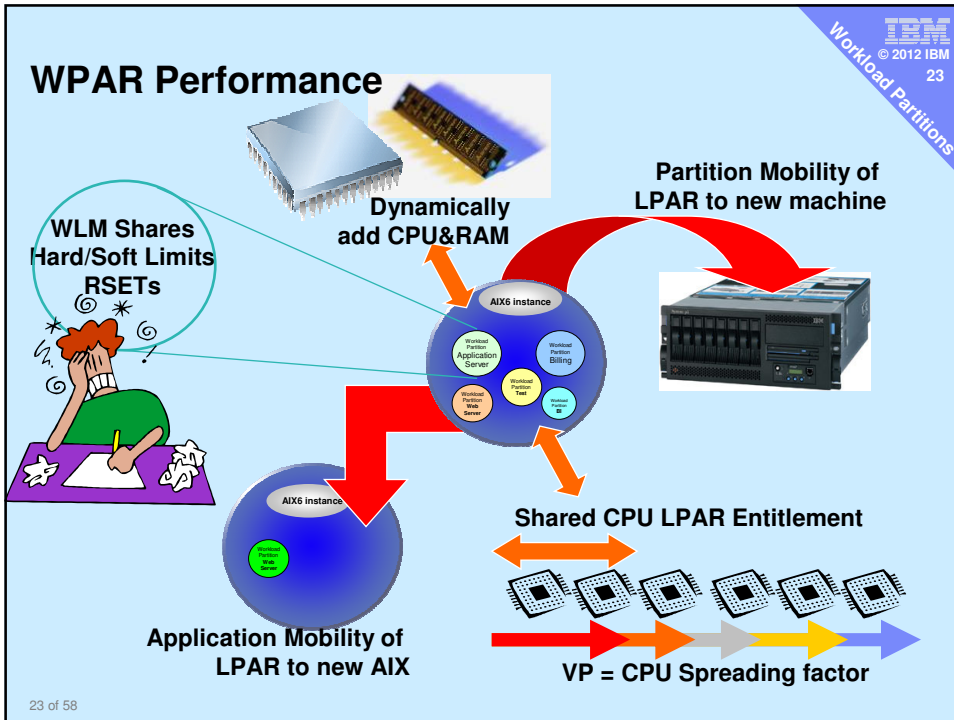
GUI → Use Systems Director WPAR Manager?

– Get it wrong by 1 character & you have to start again!

Script → wp13, network + DNS, using NFS, relocatable:

```
# mkwpar \  
-n wp13 \  
-h wp13 \  
-N netmask=255.255.255.0 address=9.69.44.123 \  
-r \  
-c \  
-M directory=/ vfs=nfs host=my nfs.ibm.com dev=/nfs/wp13root \  
-M directory=/home vfs=nfs host=my nfs.ibm.com dev=/nfs/wp13home \  
-M directory=/tmp vfs=nfs host=my nfs.ibm.com dev=/nfs/wp13tmp \  
-M directory=/var vfs=nfs host=my nfs.ibm.com dev=/nfs/wp13var \  
-P  
# startwpar wp13
```

Performance Control



IBM Systems Director

WPAR Manager

IBM PowerVM Workload Partitions Manager for AIX

Workload Partitions Resource Status

Manage Resources

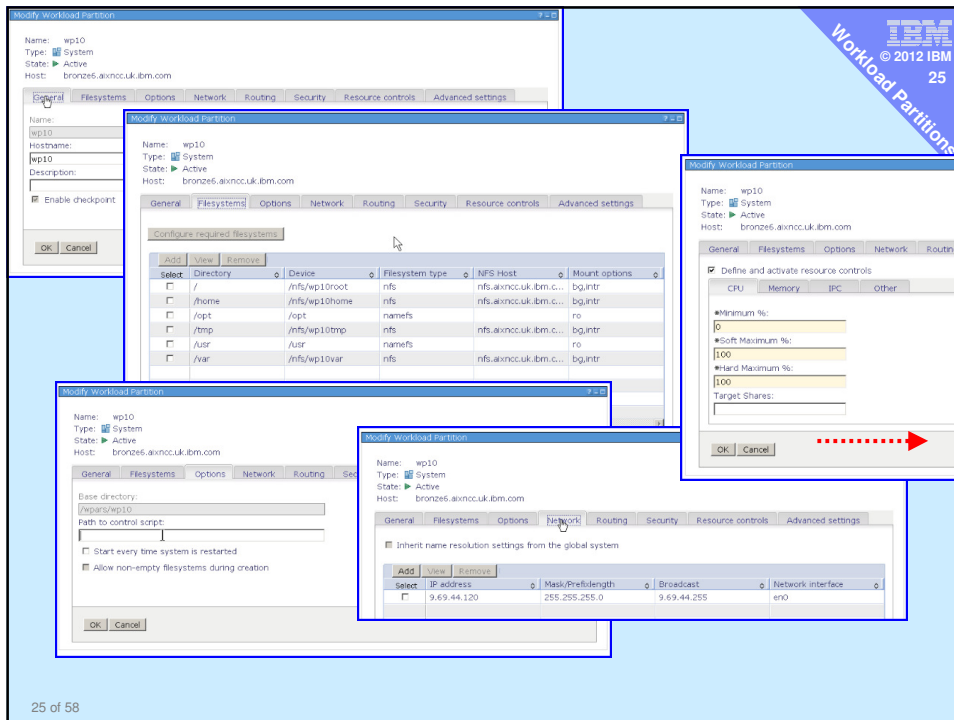
License Information

Workload Partitions and Hosts

Name	WPARType	WPARState	Access	CPU Utilizatio...	Memory Utiliz...
qbronz6.akmcc.uk.ibm...	System	Defined	OK	5.98	78.0
wp07	System	Active	OK	0	1.1
wp11	System	Defined	OK	-	-
wp12	System	Defined	OK	-	-
wp14	System	Defined	OK	-	-
wp15	System	Defined	OK	-	-
wp19	System	Defined	OK	-	-
wp20	System	Defined	OK	-	-
wp21	System	Defined	OK	-	-
red4.akmcc.uk.ibm...	System	Defined	OK	10.5	63.4
red3.akmcc.uk.ibm...	System	Defined	OK	98.85	73.48
wp06	System	Active	OK	91.75	2.22

24 of 58

© 2012 IBM
Workload Partitions 24



25 of 58

Resource Control by GUI

- CLI easy to use but:
 - Removing RSET is no obvious → use CLI -K option:
`chwp4 -K -R rset wp04`
`chwp4 -K -R shares_CPU wp04`
 - Then flip “Resource Control” off and back to “Activate”

26 of 58

Resource Control by Command Line

```
chwparr -R
  active=yes           active=no
  shares_CPU=n        CPU=min%-soft%,hard%
  shares_memory=n     memory=min%-soft%,hard%
  totalProcesses=n    totalThreads=n
  rset=rset           procVirtMem=n[M|G|T]
```

Examples:

- Switched on control (the default) `chwparr -R active=yes wp04`
- 200 CPU shares `chwparr -R shares_CPU=200 wp04`
- Set min,softmax & hardmax `chwparr -R CPU=10-50,75 wp04`
- Use only CPU four `chwparr -R rset=sys/cpu.00004 wp04`

CLI easy to use but Removing RSET is not obvious → use CLI -K option:

Examples: `chwparr -K -R rset wp04 --` or `-- chwparr -K -R shares_CPU wp04`
 - Then flip "Resource Control" off & back to "Activate" (-R active=no/yes)

Advanced WPARs: How does it work?

Which is Faster WPAR versus LPAR ?




Millions of process switch/sec
Hot cache

WPARs are FASTER than LPARs

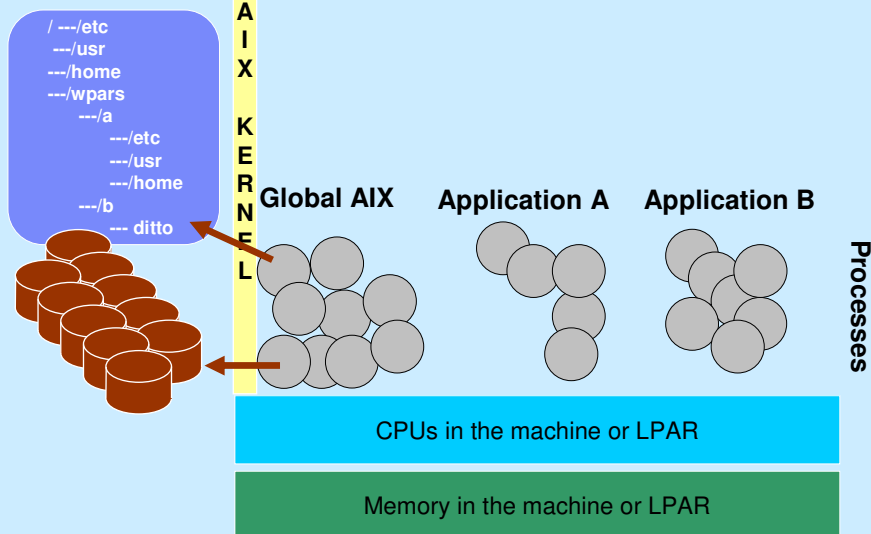
See SAS benchmark white paper

Thousands of LPAR switch/sec
Warm cache

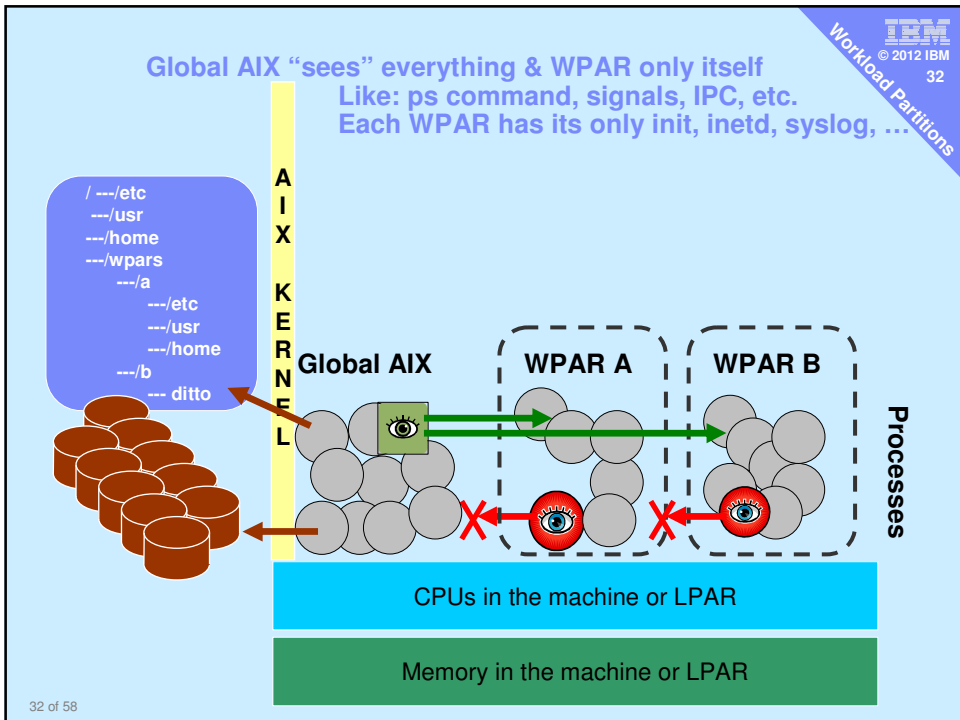
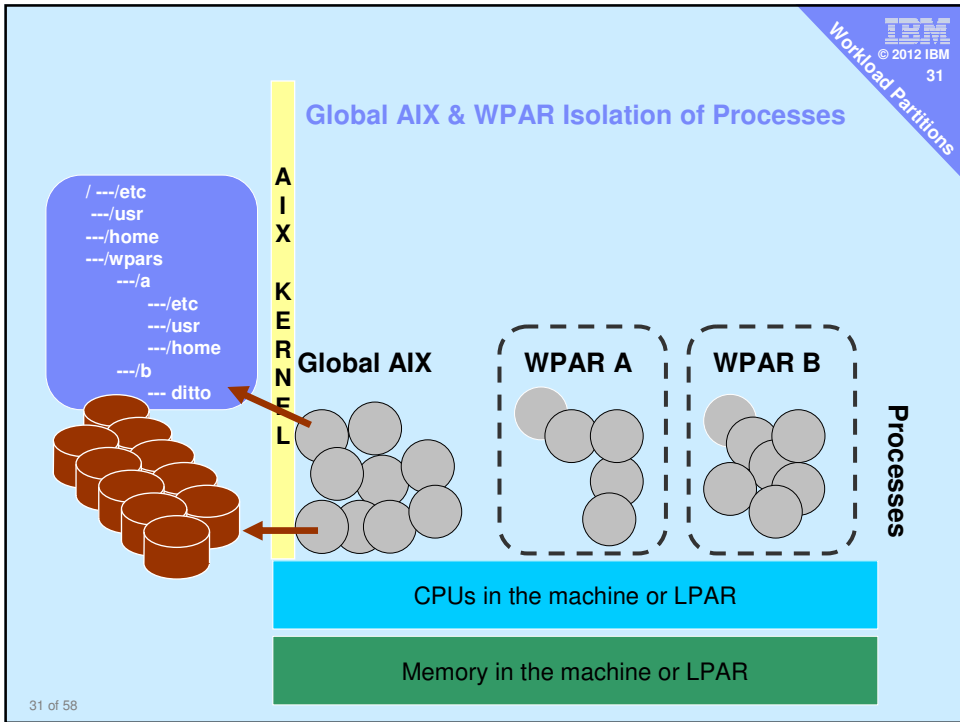


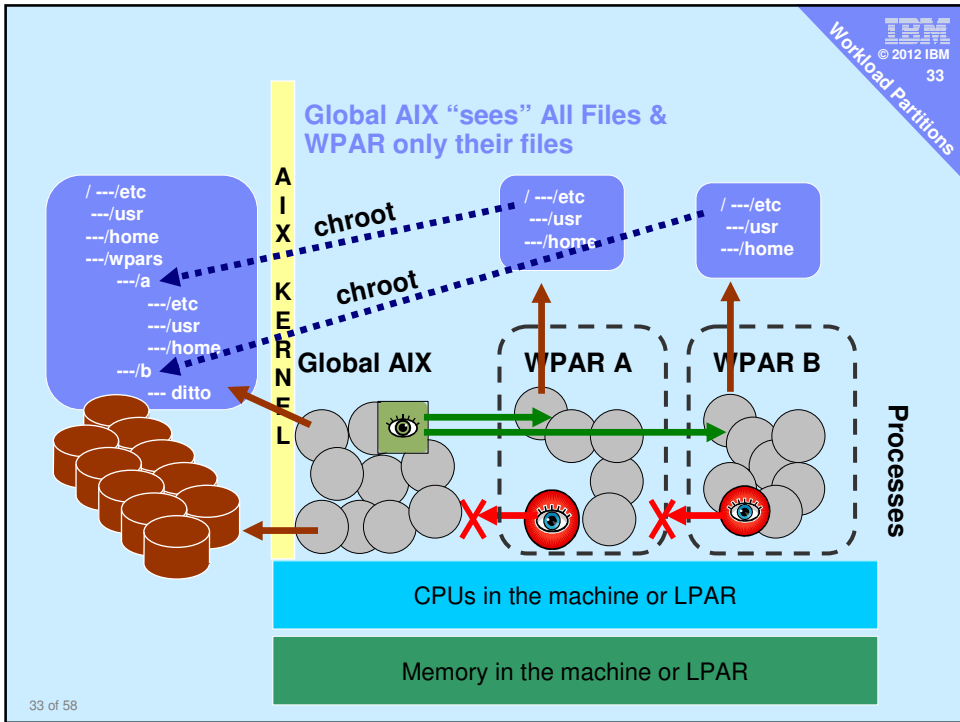
29 of 58

AIX running Multiple Applications

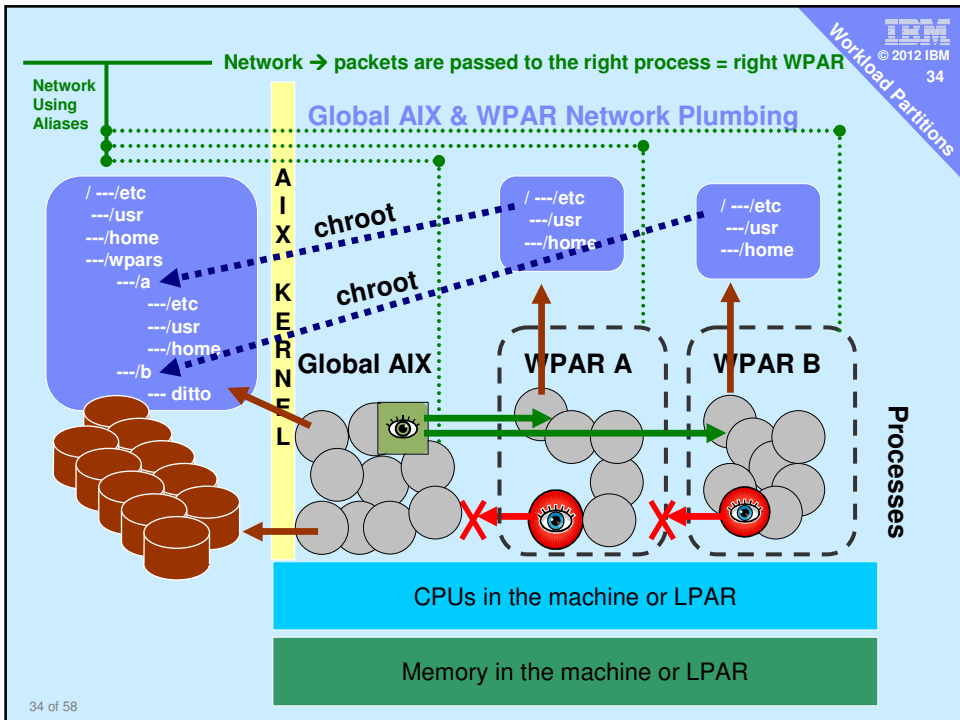


30 of 58






33 of 58



34 of 58

Backup & Restore – WPAR Mgr or simple command





Backup (savewpar)

- 1) Save WPAR details
 - 2) Save R/W files
- Minimum Size 60MB
= 2 minute

Restore (restwpar)

- Global AIX=easy
- Recreate WPAR
 - Put back R/W files



Backup/Restore

Same but option to backup files on the NFS server

WPAR Clone Example


```
restwpar -n wp14 -h wp14 -d /wpars/wp14 -r -U \
-M "-N address=9.69.44.124 \
-M directory=/      vfs=nfs host=abc.ibm.com dev=/nfs/wp14root \
-M directory=/home  vfs=nfs host=abc.ibm.com dev=/nfs/wp14home \
-M directory=/tmp    vfs=nfs host=abc.ibm.com dev=/nfs/wp14tmp \
-M directory=/var    vfs=nfs host=abc.ibm.com dev=/nfs/wp14var" \
-f /backup/wp18.bk
```

Notes: " and "

Same time as the restore ~ default WPAR = >2 minutes
WPAR Manager GUI clone has limitations

Syncwpar

- Update Global AIX 6
- If shared /usr & /opt then small syncwpar
- If private /usr & /opt then large syncwpar
- Shared Recommended:
 - WPAR 2 GB smaller
 - Uses less memory
 - Sync quicker
 - Smaller backups
- Use Private /usr&/opt only for badly behaving Applications (writing to many places)
- Can mount on to the read only /usr/&opt in sub-dir



Don't make this classic mistake ...

Don't just add WPARs to your LPAR mix

Do use large Global AIX LPARs to reduce your LPAR/AIX count for the WPAR benefits
= reduced admin & freeing up resources

WPAR+LPAR **Current LPAR Centric**
Lots of small LPARs = lots of system Admin effort

64 copies of AIX to maintain
 Manpower ~1:16→4 people

64 AIX ~64 GB RAM++

IBM © 2012 IBM 39

WPAR+LPAR **Current LPAR Centric with scattered WPARs=Bad**
Lots of small LPARs = lots of system Admin effort

64 copies of AIX to maintain
 Manpower ~1:16→4 people

64 AIX ~64 GB RAM++

IBM © 2012 IBM 40

Mostly WPAR **WPAR Pool Centric few larger LPARs = Good Production LPAR & everything else WPARs**

10 copies of AIX to maintain
Manpower ~1:16→1 person
Saving = 3 peoples time

10 AIX ~10 GB RAM
Saving 54 GB RAM=\$50K
[prices subject to change]

IBM
© 2012 IBM
41
Workload Partitions



41 of 58

Long term WPAR Maintenance
.... Bucket Concept

Making AIX updates really simple with WPARs

IBM
© 2012 IBM
42
Workload Partitions

42 of 58



WPAR All LPARs AIX 6.1

© 2012 IBM 45

Workload TL7

TL5 TL6 TL7

TL5 TL6 TL7

TL5 TL6 TL7

TL5 TL6 TL7

Empty TL2 so retire the LPAR

45 of 58

WPAR All LPARs AIX 6.1

© 2012 IBM 46

Workload TL7

TL6 TL7

TL6 TL7

TL6 TL7

TL6 TL7

Empty TL2 so retire the LPAR

46 of 58

WPAR All LPARs AIX 6.1

TL6 TL7 TL8

TL6 TL7 TL8

TL6 TL7 TL8

TL6 TL7 TL8

Move up to newer AIX 6 TLxx

47 of 58

Running old AIX 5.2 + 5.3 versions in a WPAR on AIX 7


AIX L
AIX 5.2
AIX 5.3

POWER7
BUILT ON
Power

AIX 7

48 of 58

Lets go back in time



49 of 58

AIX 5.2 – Shock Horror

AIX 5.2 features

→ Think “POWER3, RS64, POWER4”

- No shared CPU or virtual CPU
- Only dedicated adapters
- No VIOS support (net or disks)
- No Simultaneous Multi-Threading
- No DVD or no virtual optical



50 of 58

History Lesson



AIX 5.2

[my records]

- Oct 2002 - Released
- Mid 2005 - “functionally stabilised” ← my guess
- Mid 2007 - last TL (TL10) released
- Mid 2008 - withdrawn from marketing
- Mid 2009 - Last Service Pack 8

AIX 5.3

- Aug 2004 - Released for POWER5
- Mid 2008 - “functionally stabilised” ← my guess
- Mid 2011 - Last Technology Level 12
- April 2012 - Extended Support (price hike)

AIX 6

- Nov 2007 - Released



Why still running on 4 to 10 years old boxes?

- “If it ain’t broke, don’t fix it!”
- “Please, AIX upgrade to 6 or 7?”
“No way man!!”
= retest of new Middle-ware & new App Version & new OS

Versioned WPAR offers

- **Same** AIX, **same** Middle-ware & **same** Apps
on a new hardware base with AIX7 features too

Plus a hidden bonus pack ... next slide



Remove older machine & speed boost

Remove the older inefficient machine

1. Reduced electricity & cooling costs
2. Free up rack space
3. Stop hardware maintenance
4. Some old hardware parts are valuable (see eBay!)



5. Boost performance with SMT=4
Assumes thread count can use them
May reduce CPU count = reduced software licences
6. Access to AIX6/7 features
Kernel security, RBAC, WPARs(!),



Workload Partitions – Packaging

1. Base WPAR included with AIX 6 / AIX 7

- Basic WPAR functions
- AIX command line access only
- No Live Application Mobility



2. PowerVM Workload Partition (WPAR) Manager

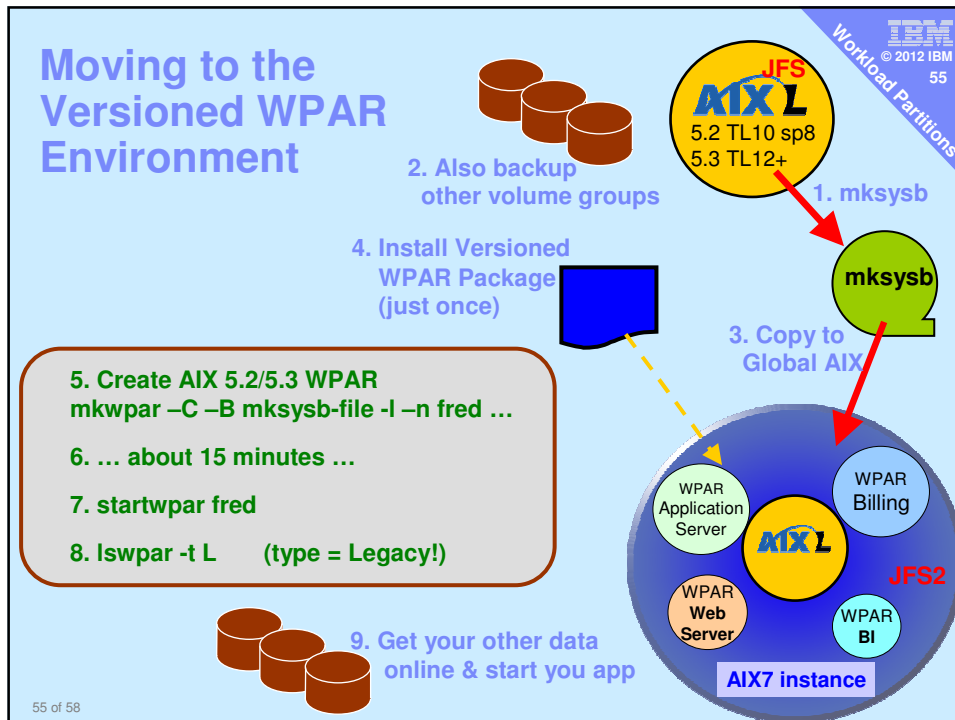
- FC 5765-G83
- Systems Director plug-in
- 60 day trial then buy a licence
- Includes Live Application Mobility (Relocation)
- Version 2.2.1 (10th Sept 2010) understands the below



3. AIX 5.2 and 5.3 Workload Partitions for AIX 7

- Extra product/LPP at a cost
 - AIX 5.2 FC 5765-H38 with limited AIX 5.2 phone + defect support
 - AIX 5.3 FC 5765-WP7 with limited AIX 5.3 phone + defect support
- AIX 5.2 media is **not provided & not available**
- AIX 5.2 TL10 SP8 Update is on Fix-Central = 1.6 GB
- AIX 5.3 TL12 Update is on Fix-Central
 - Earlier version may run but no support
- Supported on **AIX 7 & POWER7** only





Creating a Versioned WPARs – command line

Versioned WPAR AIX 5.x mksysb includes /usr & /opt

- So mandatory to have **private /usr & /opt** filesystems
- It is easy to forget the `mkwpar -l` option (lowercase L)
- Default creates local filesystem based on mksysb sizes
- Example: using local disks = no Relocation

```
mkwpar \  
-n wp08 \  
-h wp08 \  
-N interface=en0 netmask=255.255.255.0 address=9.3.2.8 \  
-r \  
-C -B my_mksysb_from_AIX53_TL12 \  
-l  
# New options for Versioned WPAR & mksysb file  
# Additions for private /usr + /opt which are mandatory with versioned WPAR
```

■ Recommend: redirect output to a file then use: `tail -f outfile`

IBM
© 2012 IBM
56
Workload Partitions

56 of 58