

Stockholm, Sweden
POWER & AIX Workshop
Nov 2012



Workload Partition (WPAR) & Versioned WPARs



Nigel Griffiths
IBM Power Systems
Advanced Technology Support, Europe

© 2012 IBM Corporation

Agenda

The basics and value add are covered quick in this session then its hands-on demo all the way.

First, the WPAR manager to show you how to


- Create and manage WPARs
- including backup, clone and recovery
- WPAR Relocation/Migrate
- Resource control and monitoring



We also cover the smart way to create WPARs and updating Global AIX and WPARs.


Also Versioned WPARs for AIX 5.2 & AIX 5.3


IBM
© 2012 IBM
3
Workload Partitions

Test:

 **Who is running AIX 6 or AIX 7?**


 **Who has tried WPAR?**



75% - 33%

IBM
© 2012 IBM
4
Workload Partitions

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/AIXmovies>

Exploiting IBM AIX Workload Partitions

Key features and benefits of WPAR

Key migration tools

Complete implementation and migration



IBM.com/redbooks

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

Workload Partitions

AIX 6

- Release November 2007 = 4+ 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 TL7 is just out
- So we are on version 8

Plus WPAR for AIX7

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

Workload Partitions (WPAR) Concept



Consolidation of isolated workloads with a single AIX instance



 = Application

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

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



AIX Workload Partitions - in a Nutshell

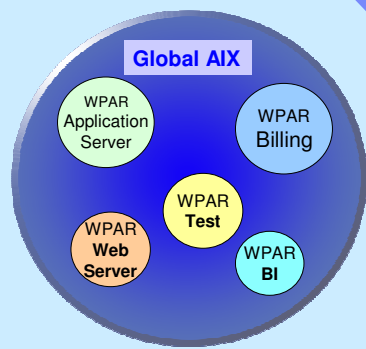
WPAR part of standard AIX 6 (2007)

- Pre-requisites = AIX 6=POWER4 to 7

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

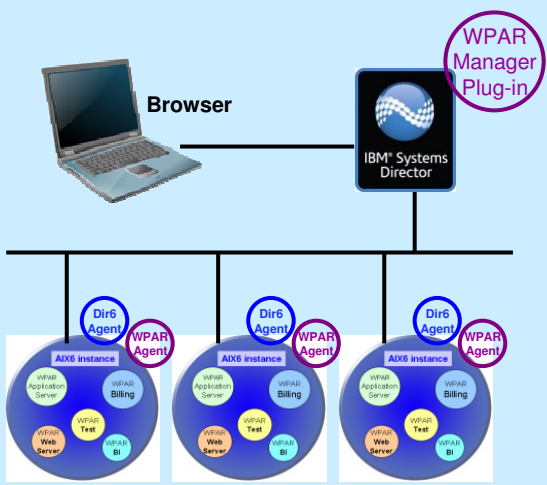
Each Workload Partition (WPAR)

- 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



Workload Partitions Manager (£\$€)

- 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

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

Filesystems

Shared /usr + /opt

Global AIX Read only





Filesystems

Shared /usr + /opt




Global AIX Read only



Private /usr + /opt



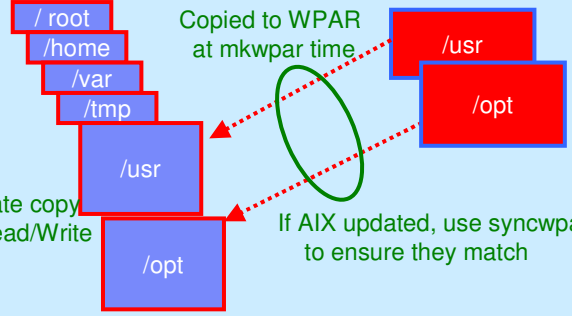
Global AIX Read only



Copied to WPAR at mkwpar time

Private copy is Read/Write

If AIX updated, use syncwpar to ensure they match





Filesystems

Shared /usr + /opt




Global AIX Read only



You can mount a filesystem or NFS into your WPAR

Private /usr + /opt

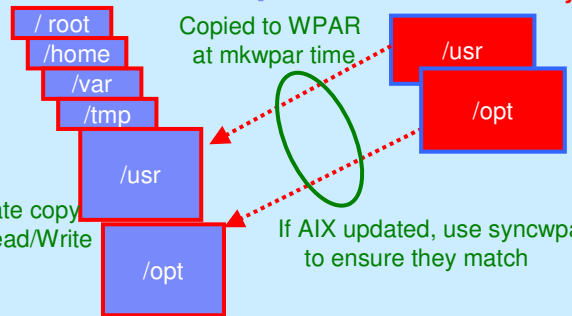
Global AIX Read only

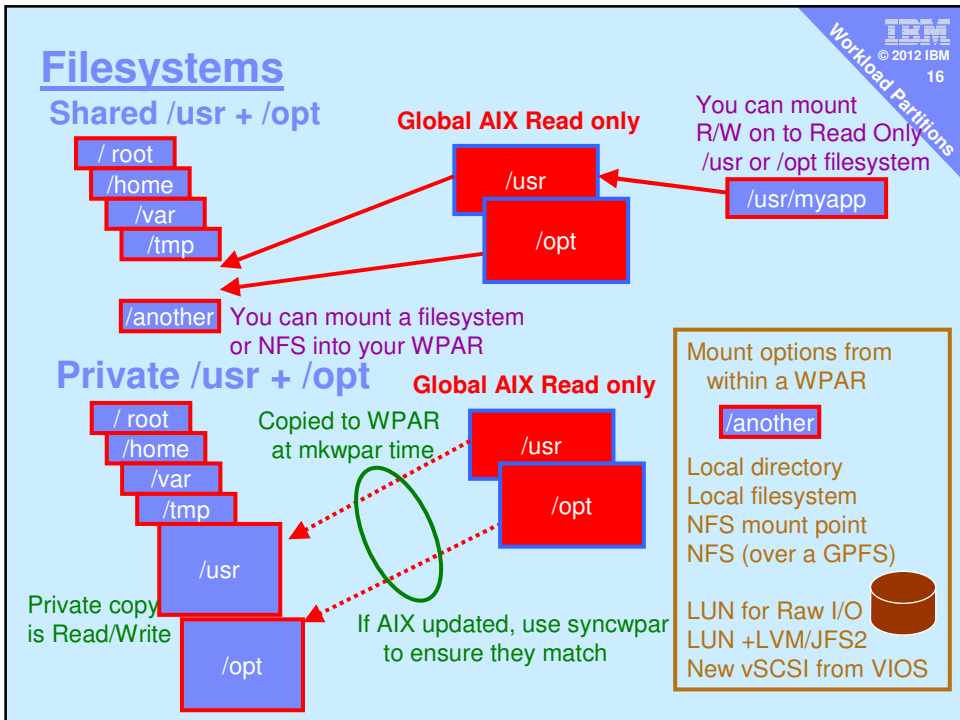
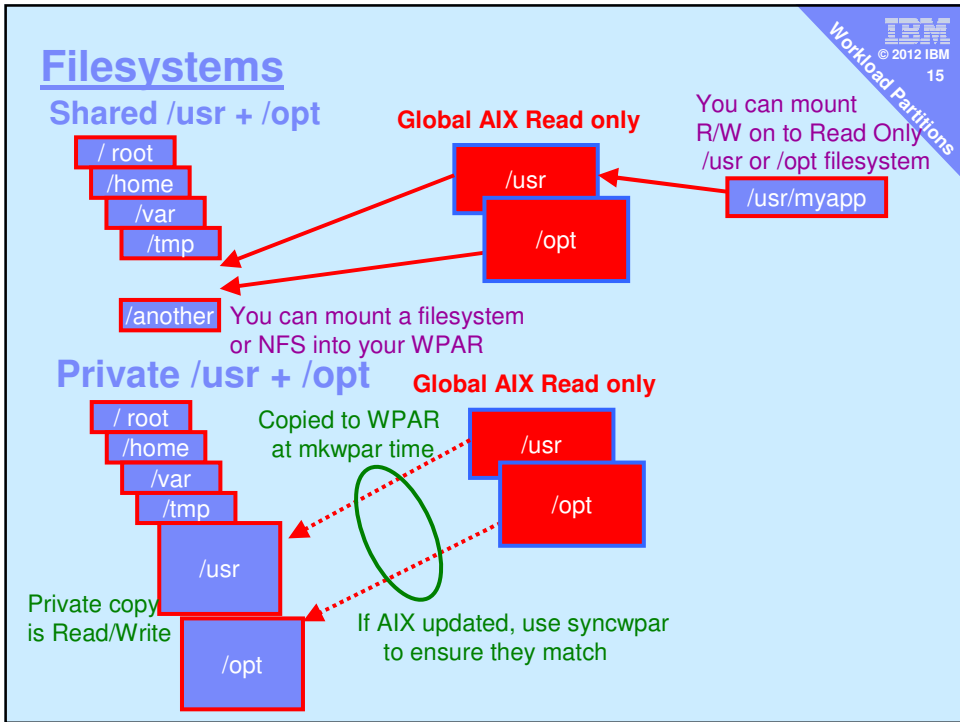


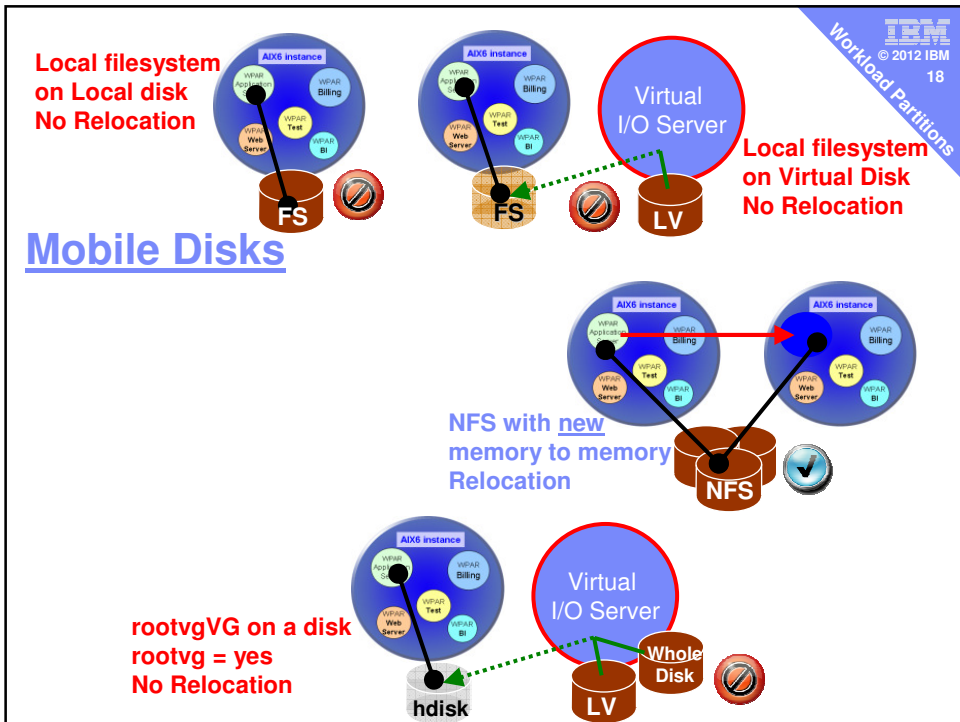
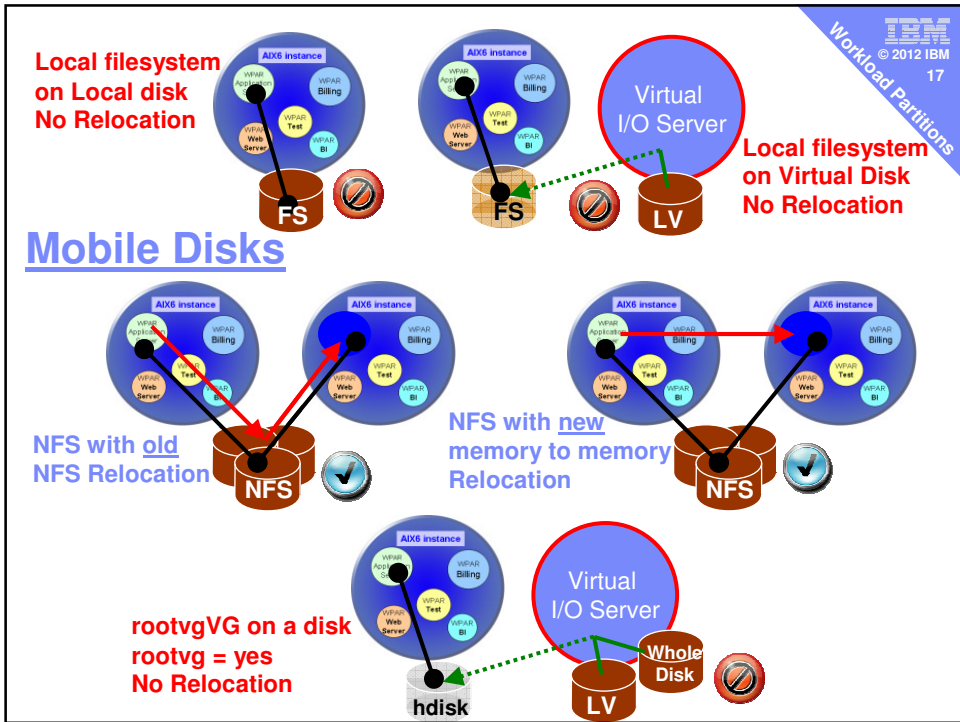
Copied to WPAR at mkwpar time

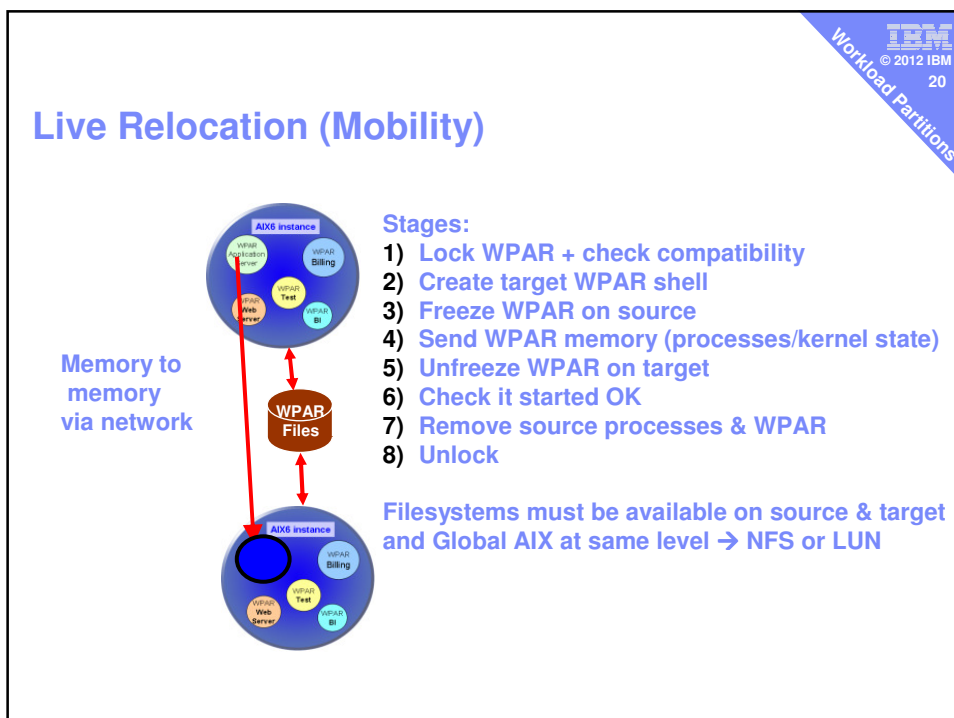
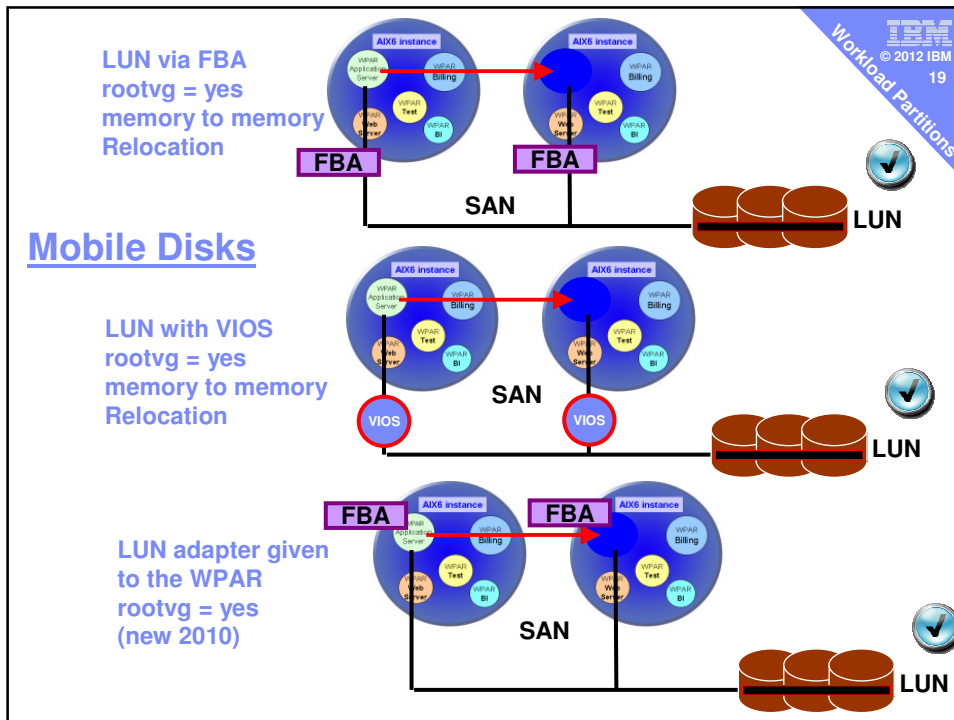
Private copy is Read/Write

If AIX updated, use syncwpar to ensure they match









Static Relocation (Mobility)

Local Disk so
no Relocation

Backup
WPAR

NFS

WPAR
Restore

NFS
Temporarily
used - for just
a few minutes



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?

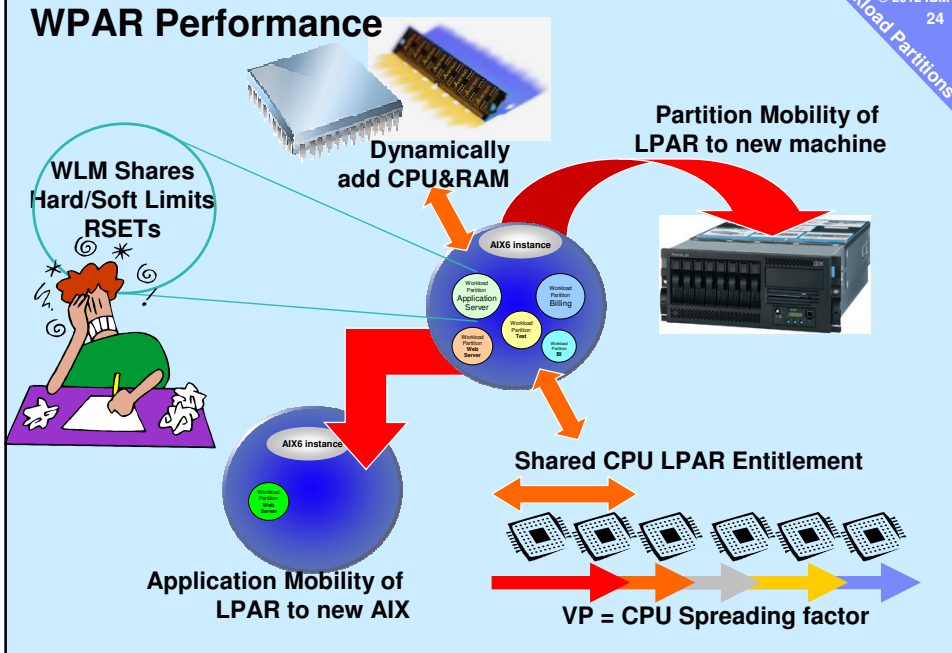
– Get it wrong by 1 character & you have to start again!

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

WPAR Performance



IBM PowerVM Workload Partitions Manager for AIX

Create, manage and relocate workload partitions (WPARs). Discover systems capable of supporting WPARs.

Workload Partitions Resource Status

10 Workload partitions and their problem severity

Manage Resources

10 Workload partitions (WPARs)
10 System WPARs
0 Application WPARs

Workload Partitions and Hosts

Name	WPAR Type	WPAR State	Access	CPU Utilizab...	Memory Utiliz...
wp07	System	Active	OK	5.98	78.0
wp10	System	Active	OK	0	11.1
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	-	-
red2.axncc.uk.ibm...			OK	10.5	63.4
red3.axncc.uk.ibm...			OK	98.85	75.48
wp06	System	Active	OK	91.75	2.22

Modify Workload Partition

Name: wp10
Type: System
State: Active
Host: bronze6.axncc.uk.ibm.com

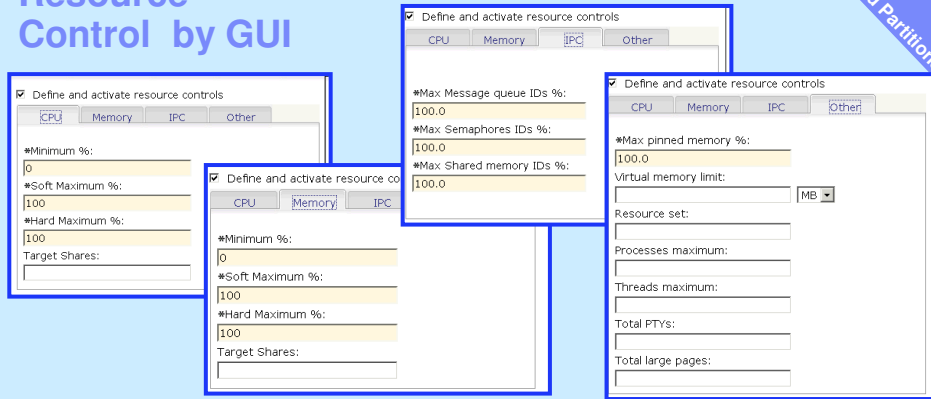
Filesystems

Select	Directory	Device	Filesystem type	NFS Host	Mount options
<input type="checkbox"/>	/	/nfs/wp10root	nfs	nfs.axncc.uk.ibm.c...	bg,intr
<input type="checkbox"/>	/home	/nfs/wp10home	nfs	nfs.axncc.uk.ibm.c...	bg,intr
<input type="checkbox"/>	/opt	/opt	namefs		ro
<input type="checkbox"/>	/tmp	/nfs/wp10tmp	nfs	nfs.axncc.uk.ibm.c...	bg,intr
<input type="checkbox"/>	/usr	/usr	namefs		ro
<input type="checkbox"/>	/var	/nfs/wp10var	nfs	nfs.axncc.uk.ibm.c...	bg,intr

Network

Select	IP address	Mask/Prefix length	Broadcast	Network interface
<input type="checkbox"/>	9.69.44.120	255.255.255.0	9.69.44.255	eno

Resource Control by GUI



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

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?

https://orange_lpar7.aixncc.uk.ibm.com:14443/ibm/console

WPAR versus LPAR – from my AIXpert Blog

Creation

1. Create WPAR in seconds & LPARs minutes
2. LPAR=VIOS LV/disk/LUNs. WPARs also has NFS or Global AIX directory
3. Rapid cloning is easy = "disposable images" - simply create, use & delete

Less memory

4. AIX boot memory LPAR = 1000 MB but WPAR ~60 MB
5. Share application code. WPARs = Lower man-power, disk space & RAM

Less updates = man-power

6. Maintain 1 Global AIX plus syncwpar command not dozens of AIX's
7. Run WPAR buckets: run Global AIX's at different AIX levels. WPAR upgrade means simple WPAR move + syncwpar (see later slide)

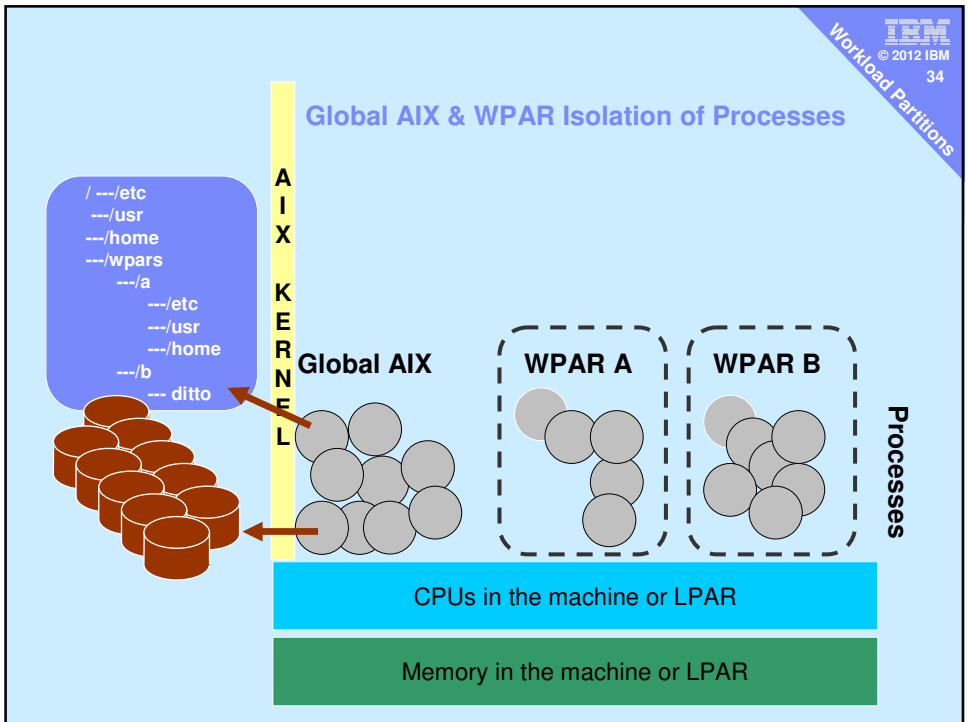
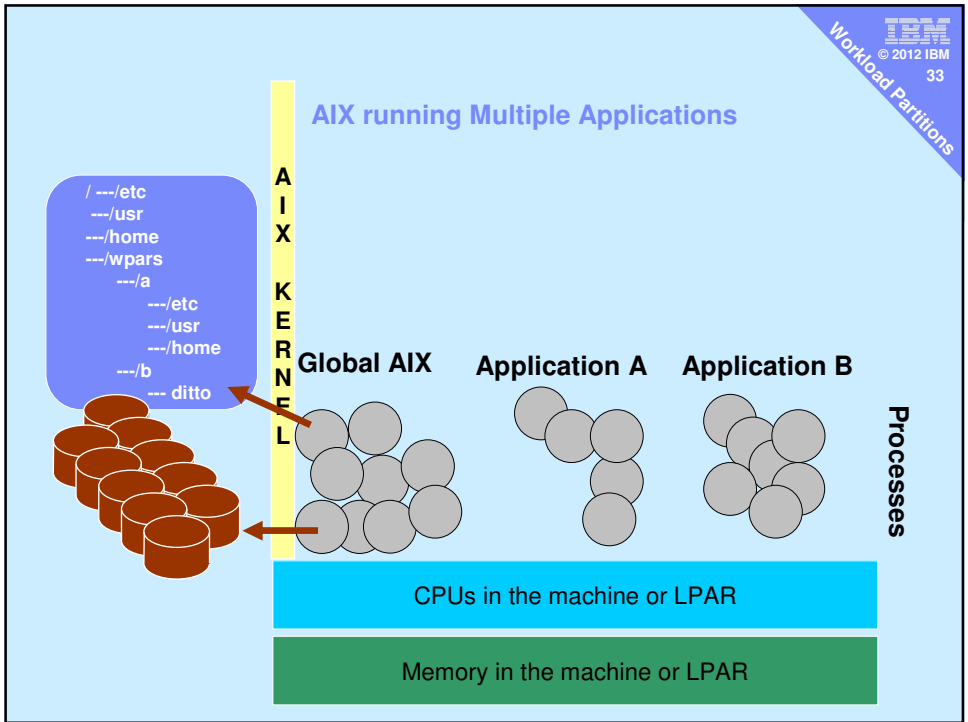
More mobile

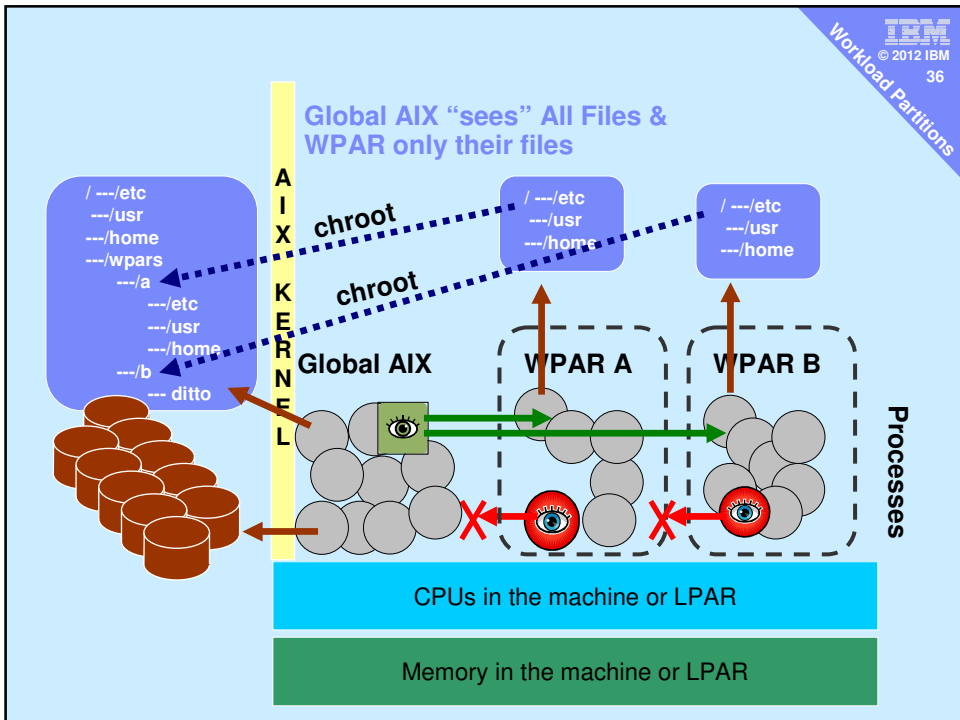
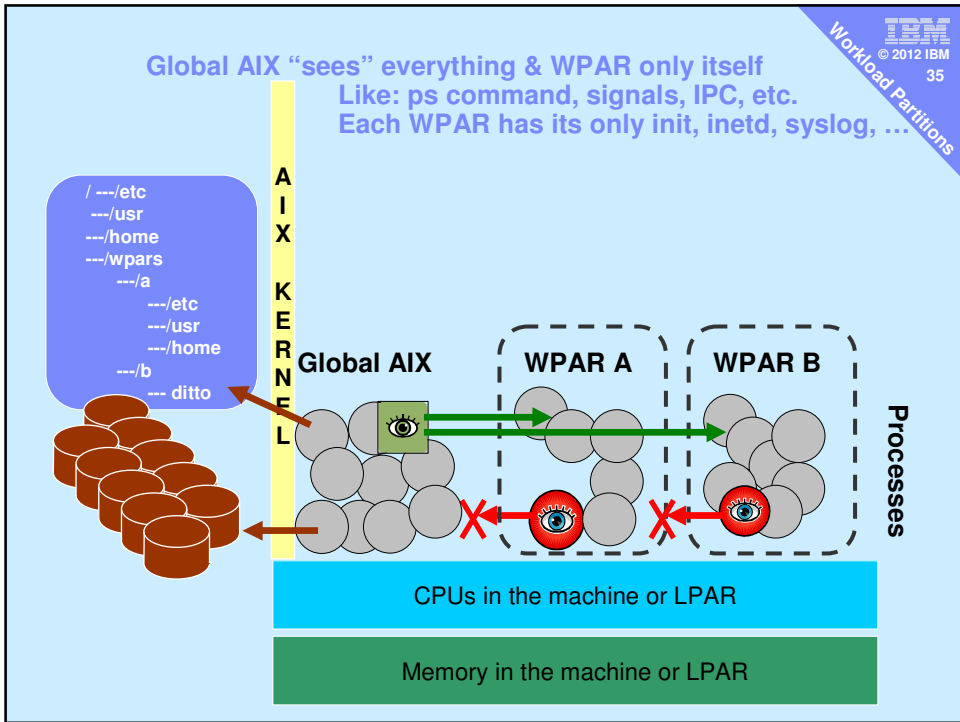
8. Application mobility much simpler to organise than LPM

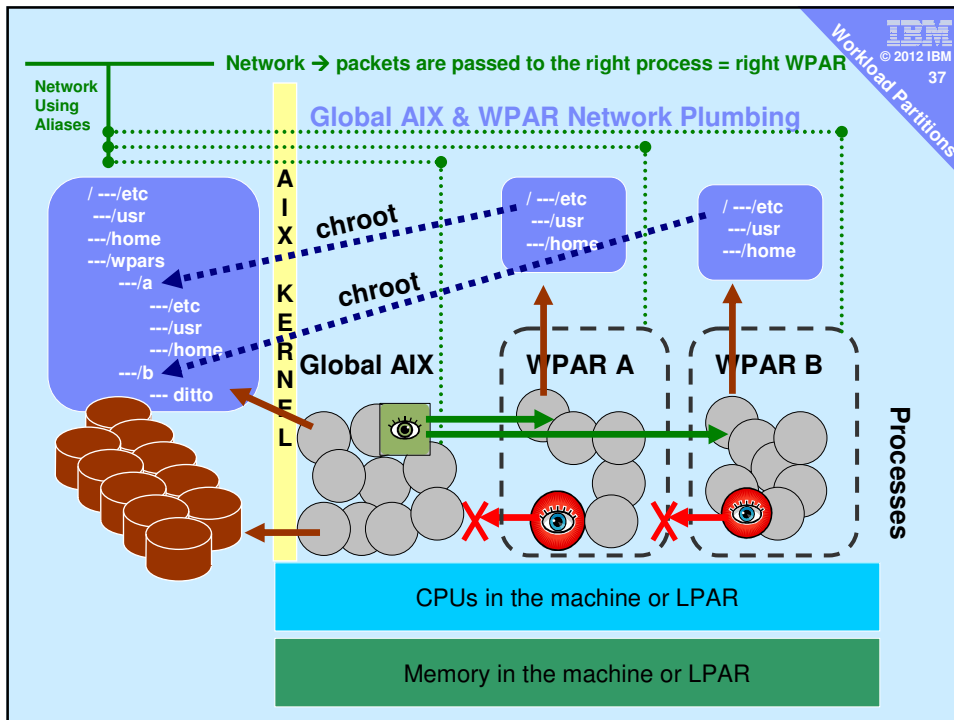
Less admin = man-power

9. Global AIX admin can access ALL WPAR files
10. Global AIX simpler Performance Monitoring than lots of AIX LPARs
11. WPAR mess up! Use Global AIX to fix it!
12. Backups are easy & small - default WPAR is 75 MB file

13 AIX7 upgrade







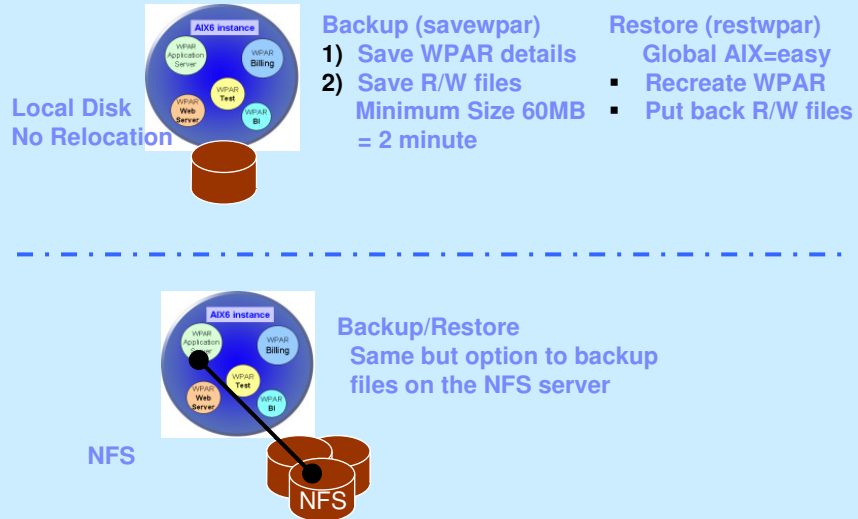
Single filesystem NFS server (not 4)

Directory: Specifies that the directory that is specified by the directory attribute is added to the file system structure of the workload partition. No file system is created. Use this attribute to reduce the number of file systems to manage in a workload partition, such as by eliminating the separate file systems for /tmp & /var

- Classic techie speak!it means in the mkwpar options:
 - M directory=/ vfs=nfs host=abc.ibm.com dev=/nfs/wp01root \
 - M directory=/home vfs=nfs host=abc.ibm.com dev=/nfs/wp01home \
 - M directory=/tmp vfs=nfs host=abc.ibm.com dev=/nfs/wp01tmp \
 - M directory=/var vfs=nfs host=abc.ibm.com dev=/nfs/wp01var
- If changed to:
 - M directory=/ vfs=nfs host=abc.ibm.com dev=/nfs/wp01root \
 - M directory=/home vfs=nfs host=abc.ibm.com dev=/nfs/wp01root/home \
 - M directory=/tmp vfs=nfs host=abc.ibm.com dev=/nfs/wp01root/tmp \
 - M directory=/var vfs=nfs host=abc.ibm.com dev=/nfs/wp01root/var
- Now only uses one NFS mount point

The IBM logo, '© 2012 IBM', and '38' are in the top right corner, along with the text 'Workload Partitions'.

Backup & Restore – WPAR Mgr or simple command



WPAR Clone Example

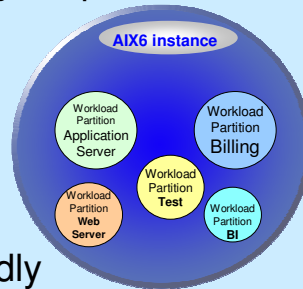
```
restwpar -n wp14 -h wp14 -d /wpars/wp14 -r -U \
-M "-N address=9.69.44.124 \
-M directory=/      vfs=nfs host=abc.ibm.com dev=/nfs/wp14root \
-M directory=/home  vfs=nfs host=abc.ibm.com dev=/nfs/wp14home \
-M directory=/tmp    vfs=nfs host=abc.ibm.com dev=/nfs/wp14tmp \
-M directory=/var    vfs=nfs 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
 43
 Workload Pa

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
 44
 Workload Pa

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


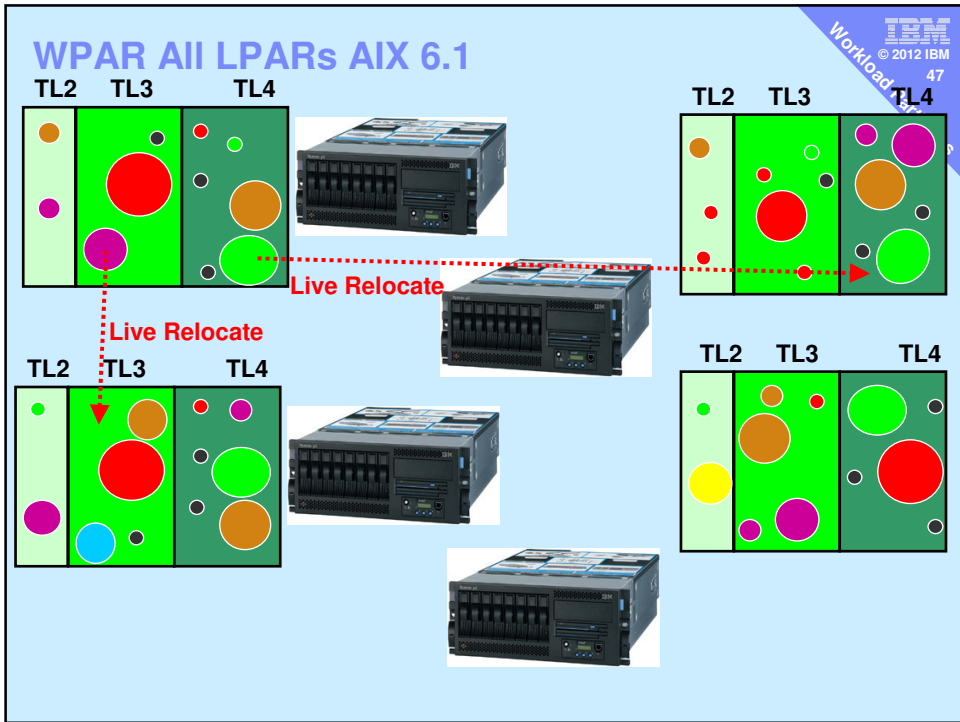
10 AIX ~10 GB RAM
Saving 54 GB RAM=\$50K
[prices subject to change]

IBM
© 2012 IBM
45
Workload Partitions

Long term WPAR Maintenance
.... Bucket Concept

Making AIX updates really simple with WPARs

IBM
© 2012 IBM
46
Workload Partitions



WPAR All LPARs AIX 6.1

© 2012 IBM 49

Workload

TL2 TL3 TL4

TL2 TL3 TL4

TL2 TL3 TL4

TL2 TL3 TL4

Empty TL2 so retire the LPAR

WPAR All LPARs AIX 6.1

© 2012 IBM 50

Workload

TL3 TL4

TL3 TL4

TL3 TL4

TL3 TL4

Empty TL2 so retire the LPAR

WPAR All LPARs AIX 6.1

© 2012 IBM
51
Workload Partitions

TL3 TL4 TL5

TL3 TL4 TL5

TL3 TL4 TL5

TL3 TL4 TL5

Move up to newer AIX 6 TLxx

Six Reasons for Workload Partitions (WPAR)

© 2012 IBM
52
Workload Partitions



- 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

Running old AIX 5.2 + 5.3 versions in a WPAR

IBM
© 2012 IBM
53

Workload Partitions

AIX L
AIX 5.2
AIX 5.3



AIX 5.2 – Shock Horror

IBM
© 2012 IBM
54

Workload Partitions

AIX 5.2 features

→ Think “POWER3, RS64, POWER4”

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



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”
- AIX upgrade to 6 or 7?
“No way man!!”
= retest of newer Middle-ware & App Version & 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

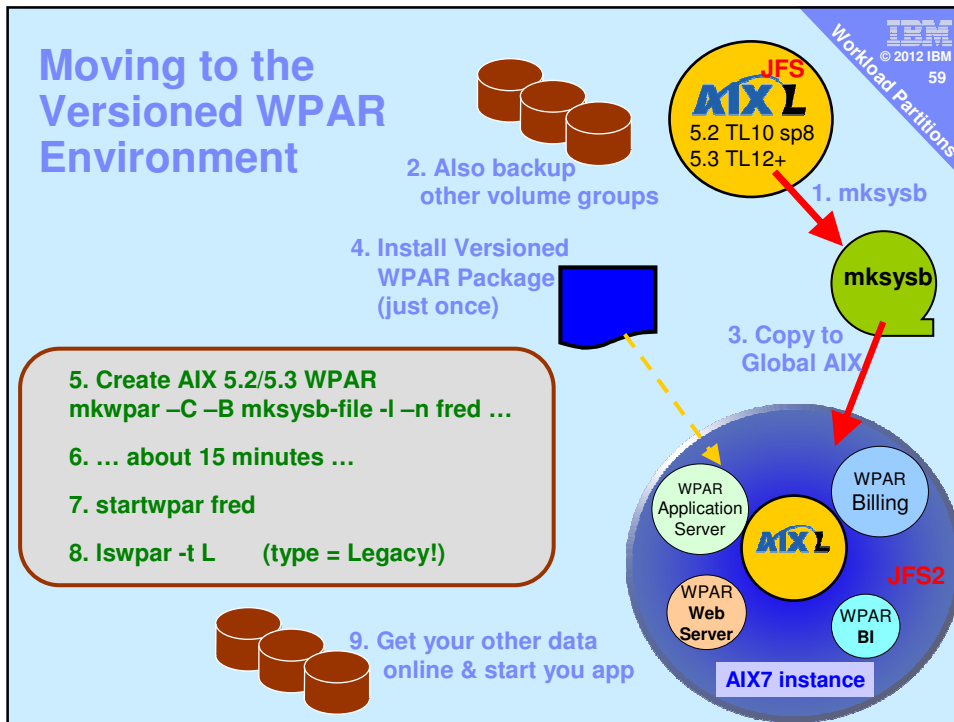
- FC 5765-G83
- Systems Director plug-in
- 60 day trial then buy a licence
- Includes Live Application Mobility (Relocation)
- Version 2.2.1 (10th Sept 2010) understands the below




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






AIX 5.2/5.3 WPAR → How do they do that?



AIX 5.2/5.3 WPAR → How do they do that?

300 AIX commands overlay to use AIX7 versions hence private /usr + /opt with overlay like
uname, perfstat, fs, vmstat, lparstat, topas_nmon ...
Fewer for AIX 5.3



AIX 5.2/3 Apps have no clue its changed

Systems Calls adjusted AIX 5.2=13 & AIX 5.3=5

/dev/kmem 25 symbols support performance tools get stats

Can load some kernel extensions into WPAR

- Also possible to install some
- AIX 5.2 or 5.3 PTFs
 - Interim fixes



Versioned WPARs - Costs

Versioned WPARs are relatively inexpensive

▪ Ballpark cost?

- GB Pounds list price at the time of creating this slide
- For small machine model = Power 750 & below
- Per core = £55 + £285 three year SWMA**

- Charged per core actually running Versioned WPAR includes limited AIX 5.2 & AIX 5.3 support



**Prices change so ask for a local quote in your currency

4

Don't forget to first get to the right TL levels

Trying a Versioned WPAR is simple

So you have AIX 5.2 or AIX 5.3 running ... just "give it a try"!

1. Probably have mksysb backups
 - If not: NFS mount Global AIX 7 disk space & run mksysb
 2. Do you know all the server connections?
 - What other servers it "talks" too or supplies data too/from?
 3. Decide disk devices and data protection:
 - Local disks, NFS, VIOS vSCSI, filesystem on LUN, rootvg LUN
 - AIX Mirror, VIOS Mirror, RAID5 in H/W, disk subsystem based
 4. WPAR will have a different IP address = safe
 5. Start it up & run a quick Regression Test
 - Does the middleware & applications start OK? Can you connect?
 - Remote connections, NFS, data load etc.
 - Check admin support: backup, monitoring, scripts, escalation, logs etc.
- Issues like missing 64 bit kernel extensions will be obvious

7

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

What you need to give it a go ...

1. mksysb for AIX 5.3 TL12 from customer
 - Suggest a USB memory key
 - Root password, knowledge of any extra security settings
 - List of applications to try and how
 - Regression test details
2. AIX 5.3 Versioned WPAR product
 - Tricky to get a copy
3. Known good AIX 5.3 mksysb
 - Time consuming to create
4. POWER7 + AIX 7 test machine
 - 1+ CPUs
 - 4GB+ memory
 - 20 GB disk

**Power ATS
Can Help**



Sharing Directories

