AIX 6 AIX 7

Workload Partition (WPAR) Starter Pack

Nigel Griffiths, Advanced Technical Support, Power Systems, Europe

WPAR - 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

Google → AIX movie

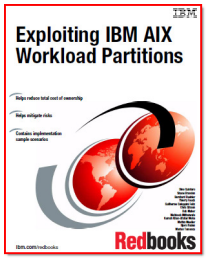

IBM Redbook
Exploiting IBM AIX Workload Partitions
www.redbooks.ibm.com **390 pages**

Workload Partitions History

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 = 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

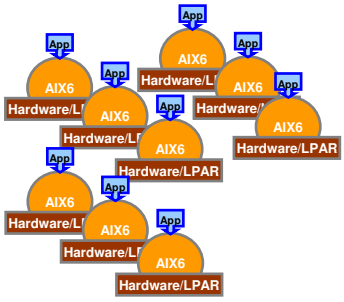

App = Application

Workload Partitions (WPAR) Concept

Consolidation of isolated workloads with a single AIX instance

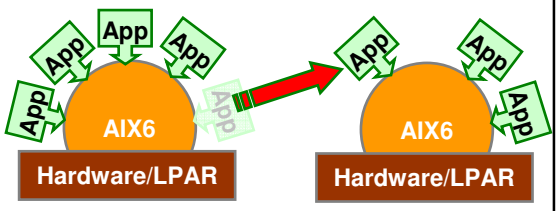

Live Application Mobility:
move WPAR between AIX instances
without restarting the WPAR applications

Six Reasons for Workload Partitions (WPAR)

1) Reduced AIX System Administration

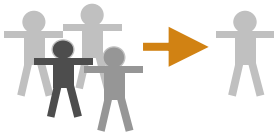

2) Application Encapsulation, monitoring and control

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

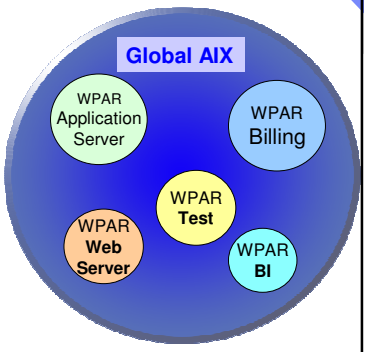

1 Global AIX kernel sharing
CPU, RAM, I/O between WPARs

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

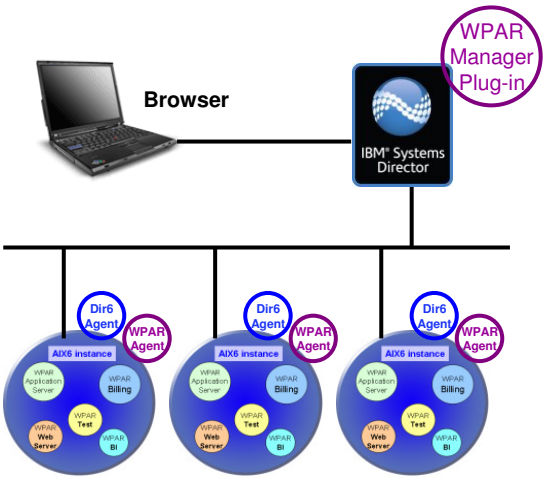

Functions
- Create & Remove
- Start & stop
- Hibernation & Cloning
- Monitoring & Reporting
- Manual Relocation
- Automated Relocation
- Policy driven change

AIX Systems
- Stand alone or LPARs

Two Workload Partitions Types

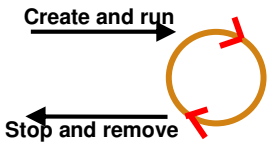

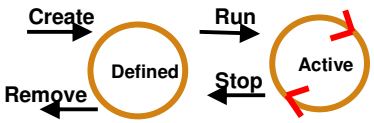

Application Workloads

- 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

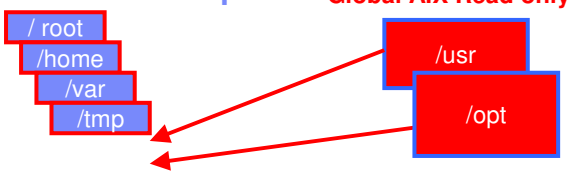

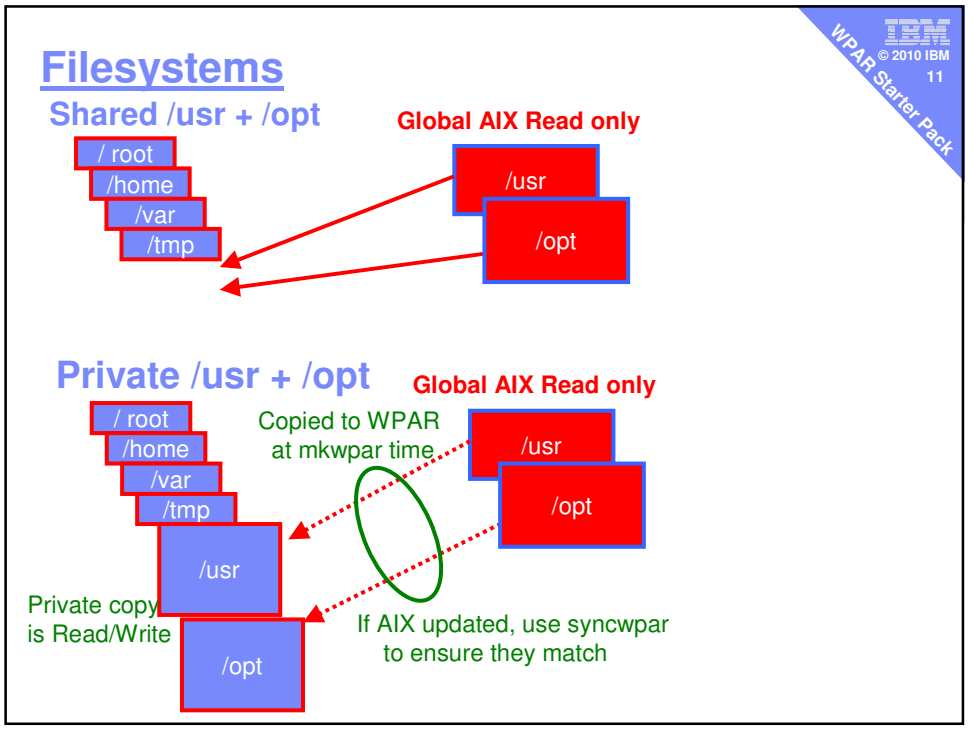
IBM
© 2010 IBM
11
WPAR Starter Pack
Filesystems
Shared /usr + /opt
Global AIX Read only
/ root
/home
/var
/tmp
/usr
/opt
Private /usr + /opt
Global AIX Read only
Copied to WPAR
at mkwpar time
/ root
/home
/var
/tmp
/usr
/opt
/usr
/opt
Private copy
is Read/Write
If AIX updated, use syncwpar
to ensure they match

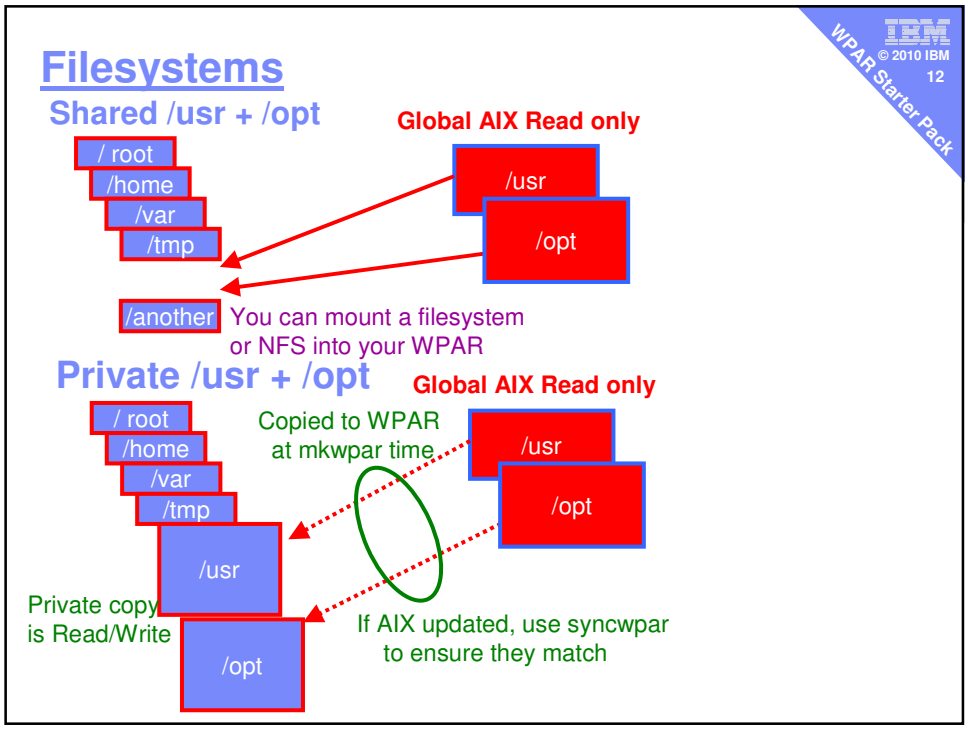
IBM
© 2010 IBM
12
WPAR Starter Pack
Filesystems
Shared /usr + /opt
Global AIX Read only
/ root
/home
/var
/tmp
/usr
/opt
/another
You can mount a filesystem
or NFS into your WPAR
Private /usr + /opt
Global AIX Read only
Copied to WPAR
at mkwpar time
/ root
/home
/var
/tmp
/usr
/opt
/usr
/opt
Private copy
is Read/Write
If AIX updated, use syncwpar
to ensure they match

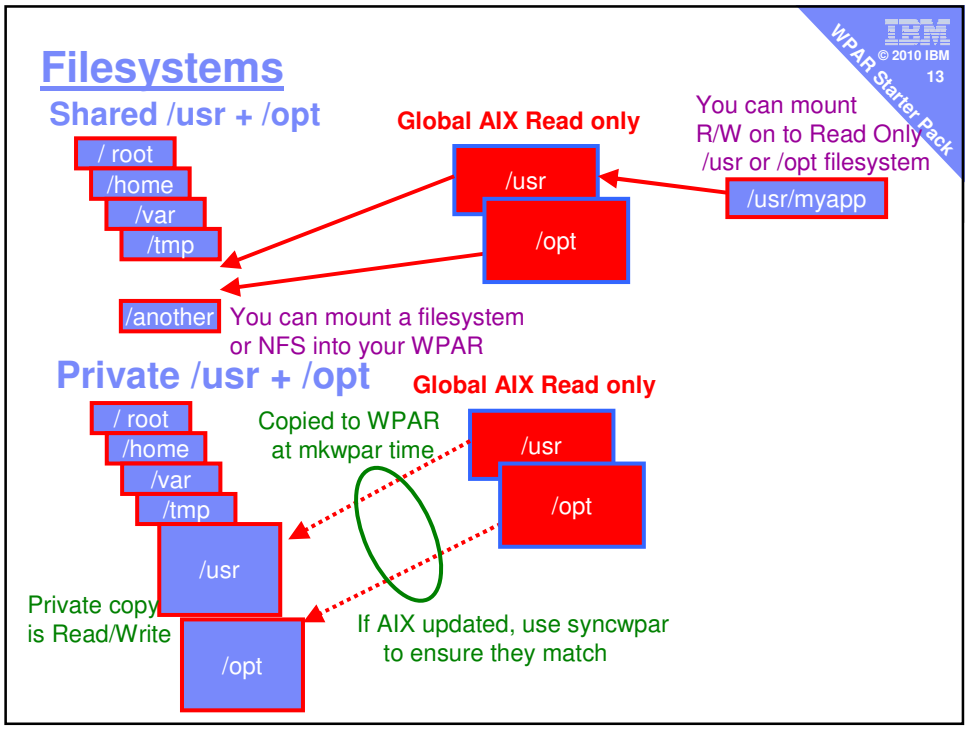
Filesystems
Shared /usr + /opt
Global AIX Read only
You can mount
R/W on to Read Only
/usr or /opt filesystem
/ root
/home
/var
/tmp
/usr
/opt
/usr/myapp
/another
You can mount a filesystem
or NFS into your WPAR
Private /usr + /opt
Global AIX Read only
Copied to WPAR
at mkwpar time
/ root
/home
/var
/tmp
/usr
/opt
Private copy
is Read/Write
If AIX updated, use syncwpar
to ensure they match
WPAR Starter Pack
© 2010 IBM
13

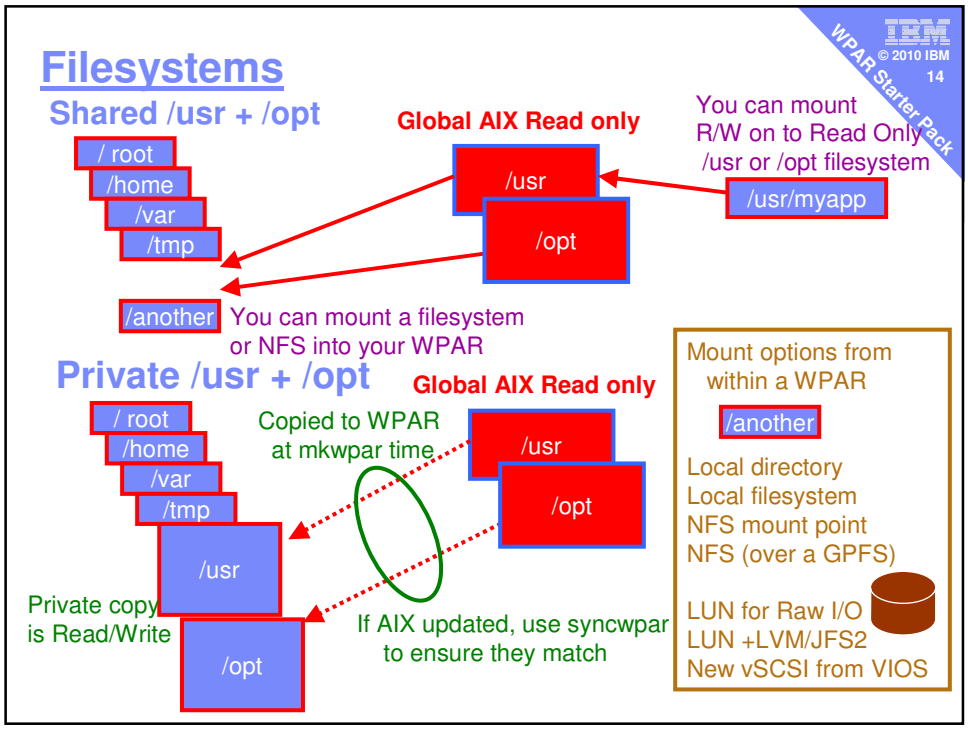
Filesystems
Shared /usr + /opt
Global AIX Read only
You can mount
R/W on to Read Only
/usr or /opt filesystem
/ root
/home
/var
/tmp
/usr
/opt
/usr/myapp
/another
You can mount a filesystem
or NFS into your WPAR
Private /usr + /opt
Global AIX Read only
Copied to WPAR
at mkwpar time
/ root
/home
/var
/tmp
/usr
/opt
Private copy
is Read/Write
If AIX updated, use syncwpar
to ensure they match
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
WPAR Starter Pack
© 2010 IBM
14

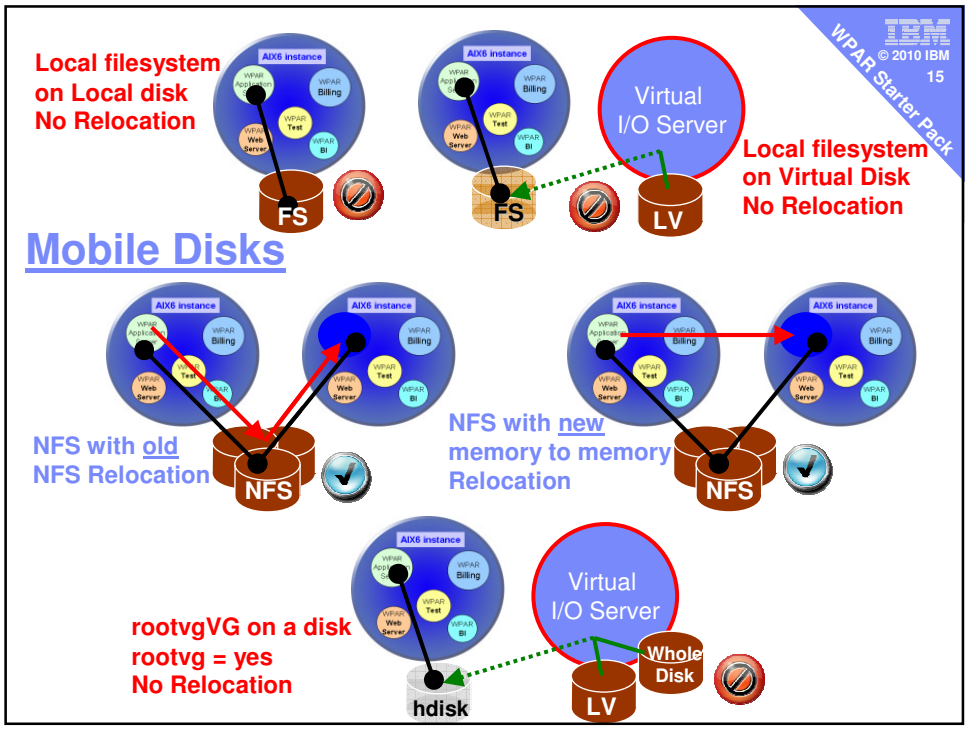
IBM
© 2010 IBM
15
WPAR Starter Pack
Local filesystem
on Local disk
No Relocation
AIX6 instance
WPAR Billing
WPAR Test
WPAR Web Server
WPAR BI
FS
Virtual
I/O Server
LV
Local filesystem
on Virtual Disk
No Relocation
Mobile Disks
NFS with old
NFS Relocation
NFS
NFS with new
memory to memory
Relocation
rootvgVG on a disk
rootvg = yes
No Relocation
hdisk
Whole
Disk

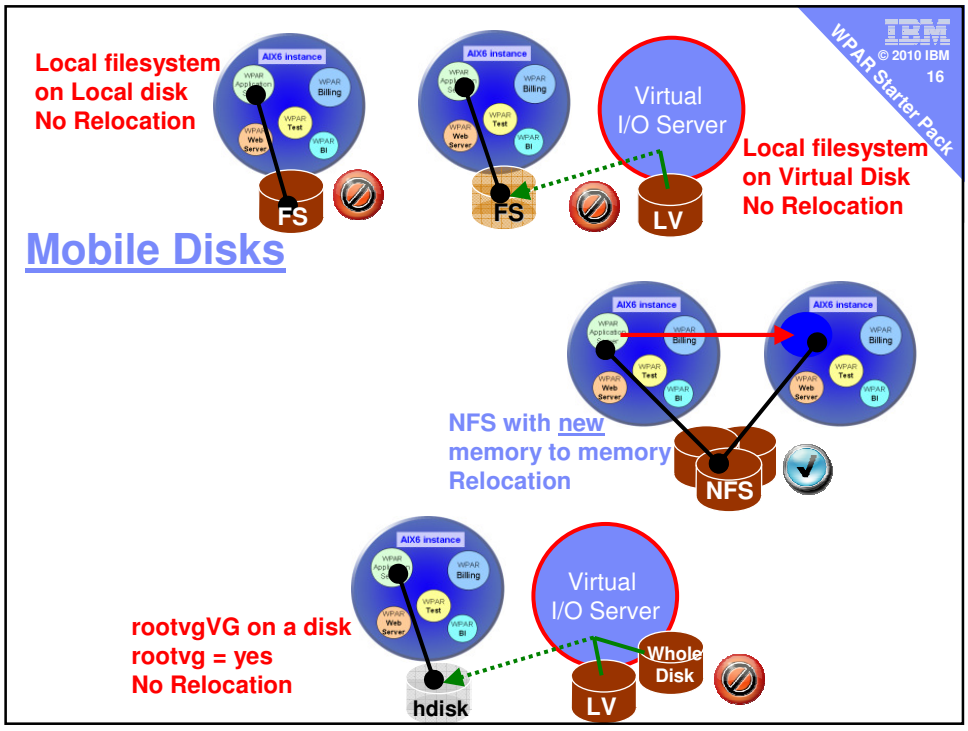
IBM
© 2010 IBM
16
WPAR Starter Pack
Local filesystem
on Local disk
No Relocation
AIX6 instance
WPAR Billing
WPAR Test
WPAR Web Server
WPAR BI
FS
Virtual
I/O Server
LV
Local filesystem
on Virtual Disk
No Relocation
Mobile Disks
NFS with new
memory to memory
Relocation
NFS
rootvgVG on a disk
rootvg = yes
No Relocation
hdisk
Whole
Disk

Mobile Disks

LUN via FBA
rootvg = yes
memory to memory
Relocation

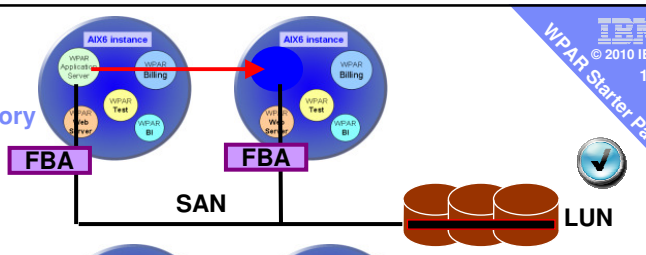

LUN with VIOS
rootvg = yes
memory to memory
Relocation

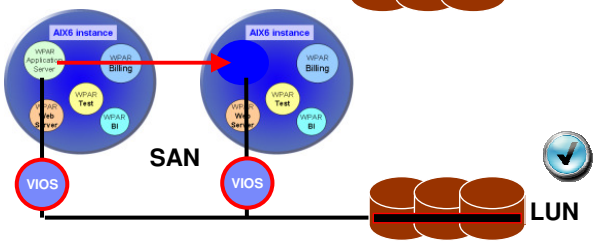

LUN adapter given
to the WPAR
rootvg = yes
(new 2010)

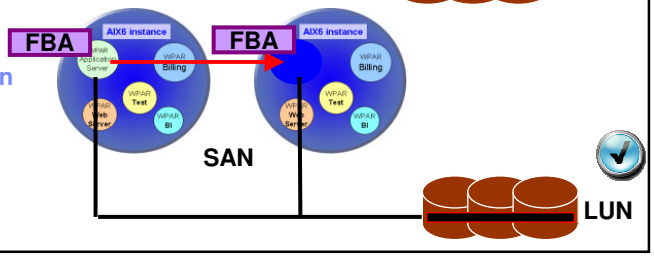

Live Relocation (Mobility)

Memory to
memory
via network

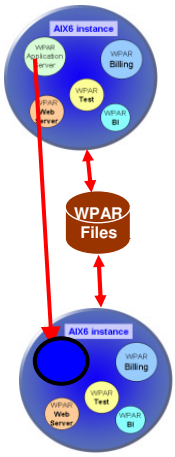

Stages:
1) Lock WPAR + check compatibility
2) Create target WPAR shell
3) Freeze WPAR on source
4) Send WPAR memory (processes/kernel state)
5) Unfreeze WPAR on target
6) Check it started OK
7) Remove source processes & WPAR
8) Unlock

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

Static Relocation (Mobility)

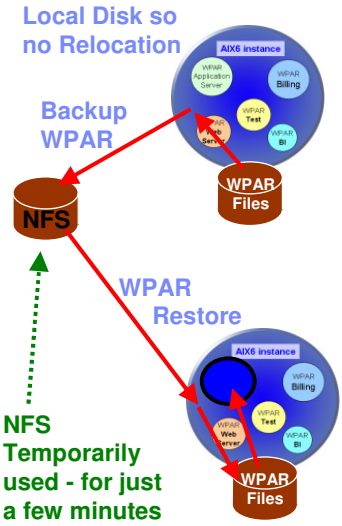

Stages:

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=mynfs.ibm.com dev=/nfs/wp13root \
-M directory=/home vfs=nfs host=mynfs.ibm.com dev=/nfs/wp13home \
-M directory=/tmp  vfs=nfs host=mynfs.ibm.com dev=/nfs/wp13tmp  \
-M directory=/var  vfs=nfs host=mynfs.ibm.com dev=/nfs/wp13var
-P
# startwpar wp13
```

How does it work?

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

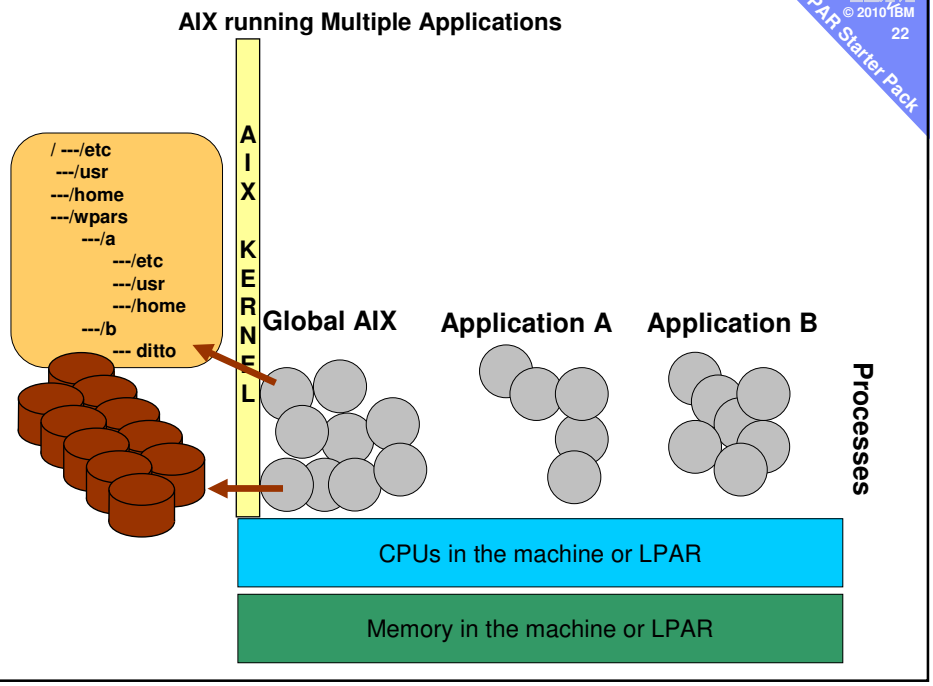

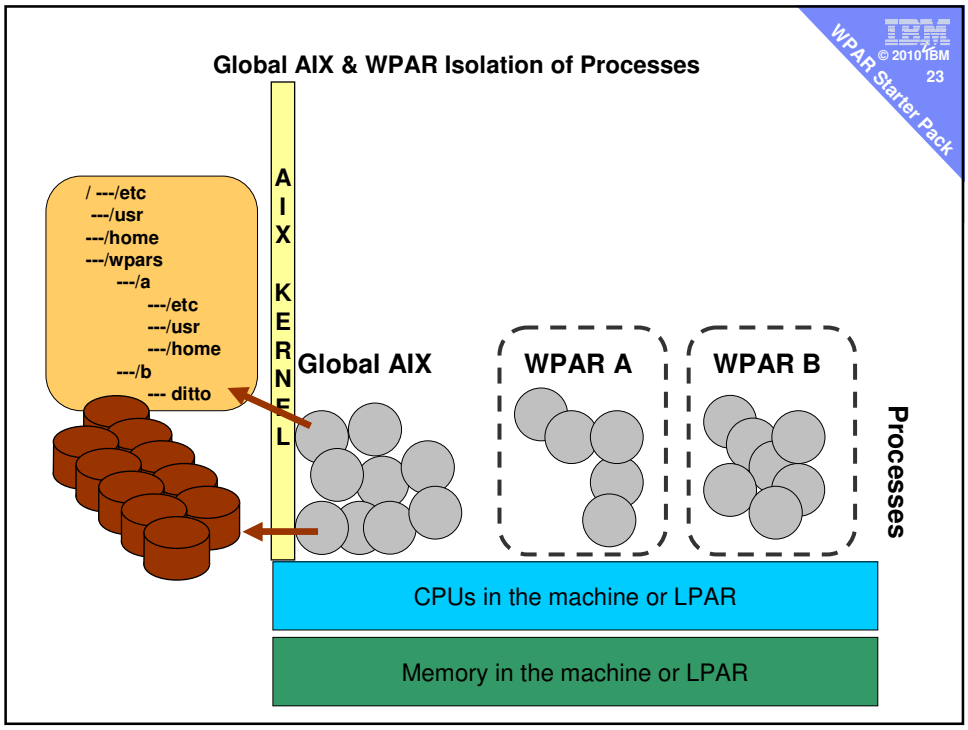
Global AIX & WPAR Isolation of Processes
WPAR Starter Pack
© 2010 IBM
23
/ ---/etc
---/usr
---/home
---/wpars
---/a
---/etc
---/usr
---/home
---/b
--- ditto
AIX KERNEL
Global AIX
WPAR A
WPAR B
Processes
CPUs in the machine or LPAR
Memory in the machine or LPAR

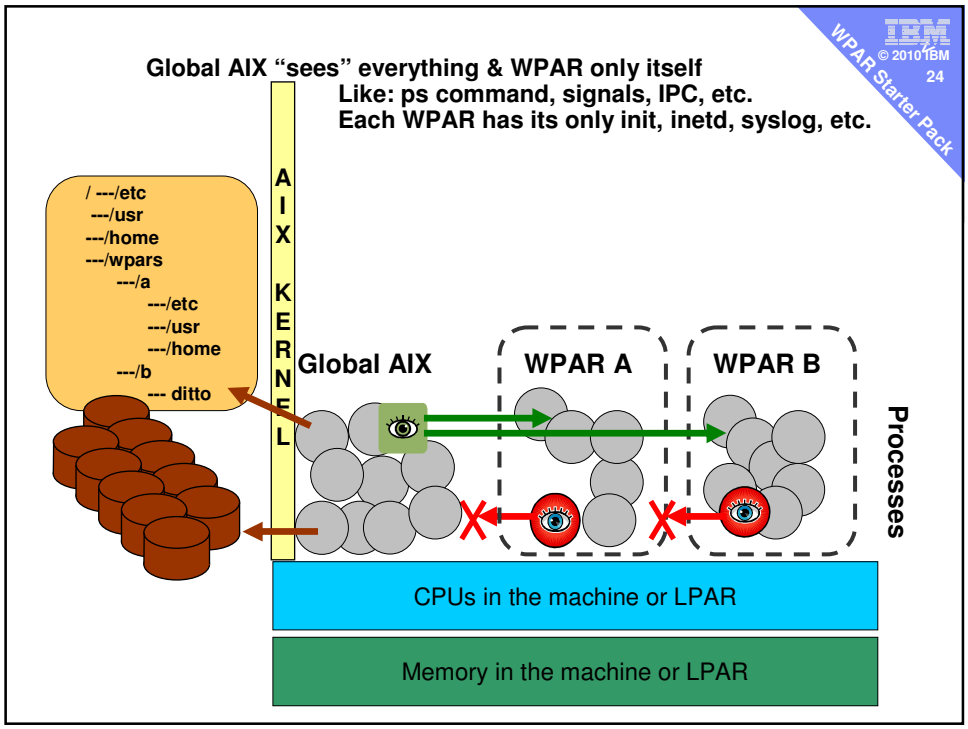
Global AIX "sees" everything & WPAR only itself
Like: ps command, signals, IPC, etc.
Each WPAR has its only init, inetd, syslog, etc.
WPAR Starter Pack
© 2010 IBM
24
/ ---/etc
---/usr
---/home
---/wpars
---/a
---/etc
---/usr
---/home
---/b
--- ditto
AIX KERNEL
Global AIX
WPAR A
WPAR B
Processes
CPUs in the machine or LPAR
Memory in the machine or LPAR

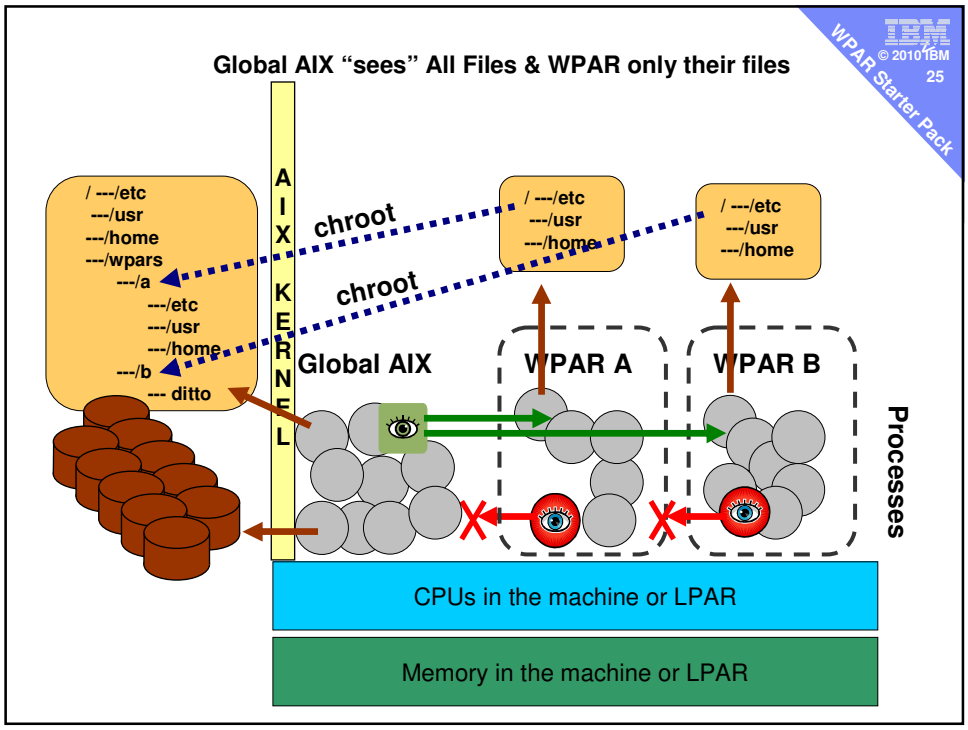
Global AIX "sees" All Files & WPAR only their files
© 2010 IBM
25
WPAR Starter Pack
/ ---/etc
---/usr
---/home
---/wpars
---/a
---/etc
---/usr
---/home
---/b
--- ditto
AIX KERNEL
chroot
chroot
/ ---/etc
---/usr
---/home
/ ---/etc
---/usr
---/home
Global AIX
WPAR A
WPAR B
Processes
CPUs in the machine or LPAR
Memory in the machine or LPAR

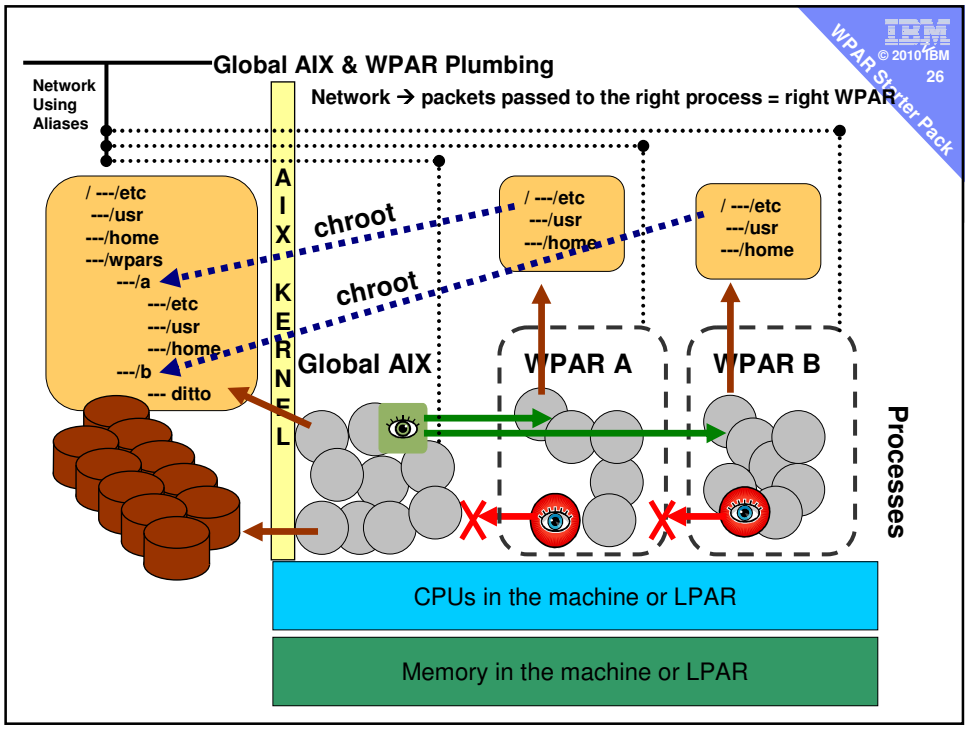
Global AIX & WPAR Plumbing
© 2010 IBM
26
WPAR Starter Pack
Network
Using
Aliases
Network → packets passed to the right process = right WPAR
/ ---/etc
---/usr
---/home
---/wpars
---/a
---/etc
---/usr
---/home
---/b
--- ditto
AIX KERNEL
chroot
chroot
/ ---/etc
---/usr
---/home
/ ---/etc
---/usr
---/home
Global AIX
WPAR A
WPAR B
Processes
CPUs in the machine or LPAR
Memory in the machine or LPAR

Performance Control

WPAR Performance

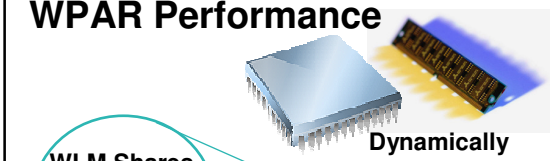

Dynamically add CPU&RAM

Partition Mobility of LPAR to new machine

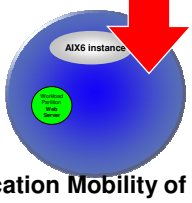

Application Mobility of LPAR to new AIX

Shared CPU LPAR Entitlement

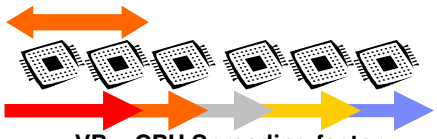

VP = CPU Spreading factor

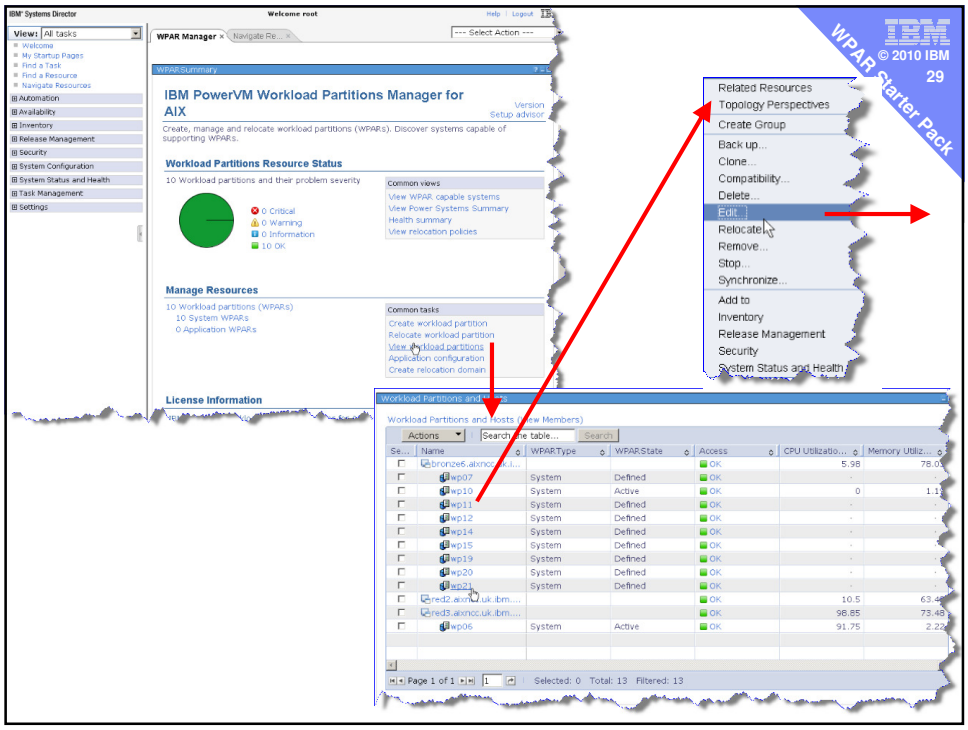
IBM PowerVM Workload Partitions Manager for AIX
Workload Partitions Resource Status
Manage Resources
License Information
Related Resources
Topology Perspectives
Create Group
Back up...
Clone...
Compatibility...
Delete...
Edit
Relocate...
Remove...
Stop...
Synchronize...
Add to
Inventory
Release Management
Security
System Status and Health
Workload Partitions and Hosts
WPAR Starter Pack
© 2010 IBM
29

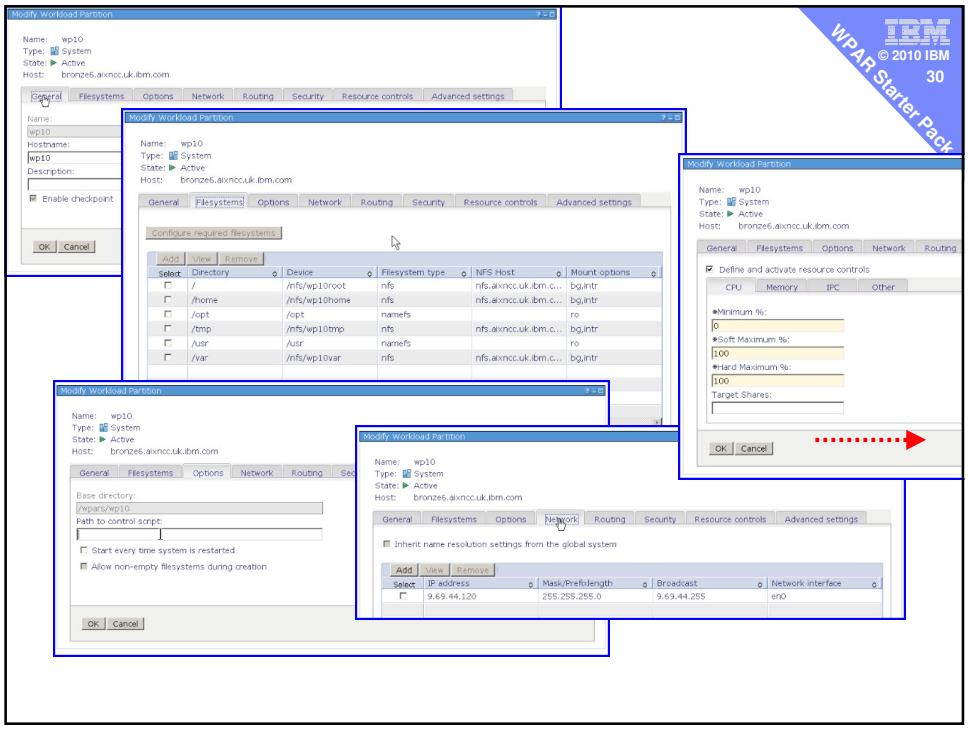
Modify Workload Partition
WPAR Starter Pack
© 2010 IBM
30

Resource Control by GUI

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

Resource Control by Command Line

chwpar –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)** **chwpar -R active=yes wp04**
- **200 CPU shares** **chwpar -R shares_CPU=200 wp04**
- **Set min,softmax & hardmax** **chwpar -R CPU=10-50,75 wp04**
- **Use only CPU four** **chwpar -R rset=sys/cpu.00004 wp04**

CLI easy to use but Removing RSET is not obvious → use CLI –K option:
Examples: chwpar –K -R rset wp04 -- or -- chwpar –K –R shares_CPU wp04

- Then flip “Resource Control” off & back to “Activate” (-R active=no/yes)