SUN'S PARTITIONING ADVANTAGE

Ben Lenail

Technical Lead

Competition Office

www.sun.com

Sun microsystems

We make the net work.

Reasons for Partitioning

- Fault containment
- Security Isolation
- Workload separation
- Also
  - Flexibility
  - Manageability
  - Serviceability
  - Reconfigurability
  - Appropriate Granularity
  - Low overhead

UNIX Partitioning Experience

	Sun Dynamic System Domains	IBM pSeries LPARs	HP nPars	HP vPars
Static Partitions	1994	2001	2001	2001
Dynamic Partitions	1997	2002	Not available	CPUs Only
Automated, Load-Balanced Capacity management	2000	Not available	Not available	Not available

Partitioning Technologies

Soft ⟷ Hard

Solaris Service Containers	IBM LPAR Partitions	Dynamic System Domains
SW Fault Containment	SW Fault Containment	HW & SW Fault Containment
Highly Flexible	Limited HW Fault Containment	Flexible
Security Isolation	Flexible	High Security Isolation
Workload Separation	Security Isolation	High Workload Separation
Fully Dynamic	Workload Separation	Fully Dynamic
Failover Mechanisms	Partially Dynamic (AIX 5.2)	Allows Online Serviceability
Management Framework		

IBM LPARs:

- **Share Critical Components**
- **Lack Hardware Fault Isolation**
- **Cannot be Serviced Online**

Partitioning Technologies

Soft — Hard

Solaris Service Containers	HP Virtual Partitions	Dynamic System Domains
SW Fault Containment	Fairly Flexible	HW & SW Fault Containment
Highly Flexible	Workload Separation	Flexible
Security Isolation	Poor Security and Fault Isolation	High Security Isolation
Workload Separation	Partially Dynamic (HP-UX shortcomings)	High Workload Separation
Fully Dynamic		Fully Dynamic
Failover Mechanisms		Allows Online Serviceability
Management Framework		

HP vPars:

- **Share Critical Components**
- **Lack Good Security and Fault Isolation**
- **Come with Increased Overhead and Decreased Performance**

Partition Feature/Functionality

FEATURE	Sun Fire 4800-15K DSD	IBM p670/p690 LPAR	HP nPar	HP vPar
Dynamic CPUs	Yes	AIX 5.2 Only	No	Yes
Dynamic Memory	Yes	AIX 5.2 Only	No	No
Dynamic I/O	Yes	AIX 5.2 Only	Yes	No
Automated DR	Yes	No	No	No
DR Certified Apps	Yes	No	No	No
ISV DR API	Yes	No	No	No
Overhead	Low	Med	Low	High
Security	High	Med	High	Low

Online Serviceability Comparison

FEATURE	Sun Fire 4800-15K	IBM pSeries	HP PA-RISC
Hot Swap CPUs	Yes	No	No
Hot Swap Memory	Yes	No	No
Hot Swap I/O	Yes	PCI Cards Only	PCI Cards Only
Hot Swap Service Processors	Yes	No	No
Hot Swap Power Supplies	Yes	Yes	Yes
Hot Swap Fans	Yes	Partial	Yes
Serviceable Partitions	Yes	No	Partial
Isolated Hardware Diagnostics	Yes	No	Yes
Dynamically Resize Partitions	Yes	Yes	No
File System Snapshots	Standard	Standard	Optional

Partition Granularity

FEATURE	Sun Fire 4800-15K DSD	IBM p670/690 LPAR	HP nPar	HP vPar
Min CPUs	2	1	2	1
Min Rec CPUs	2	3	2	2
Min Memory	2 GB	256 MB	2 GB	512 MB
Min Rec Memory	2 GB	4 GB	2 GB	1 GB
Min I/O Slots	4 or 8	1	7/11	2
Min Rec I/O Slots	4 or 8	10	7/11	2
Min CPUs Increment	2	1	2	1
Min Memory Increment	2 GB	256 MB	2 GB	64 MB
Min I/O Slot Increment	4 or 8	1	8/12	1
Min/Rec I/O Slot Inc	4 or 8	10	7	1

Partition RAS

FEATURE	Sun Fire 4800-15K DSD	IBM p670/p690 LPAR	HP nPar	HP vPar
App SW Fault Isolation	Yes	Yes	Yes	Yes
OS SW Fault Isolation	Yes	Yes	Yes	Yes
CPU Fault Isolation	Yes	Yes	Yes	No
Cache Fault Isolation	Yes	No	Yes	No
Memory Fault Isolation	Yes	No	Yes	No
I/O Subsystem Isolation	Yes	No	Yes	No
I/O Bus Isolation	Yes	No	Yes	No
I/O Slot Isolation	Yes	Yes	Yes	No
I/O Driver Isolation	Yes	No	Yes	??
Partition Isolated Diagnostics	Yes	No	Yes	No
Partition Isolated Servicing	Yes	No	Yes	No
Partition Online Service	Yes	No	No	No
Partition Online Upgrade	Yes	No	No	No
Dedicated Mgt Net	F12K/F15K	No	N/A	No

Sun's Dynamic System Domains Minimize Impact of Failures

Dynamic System Domains isolate HW & SW failures

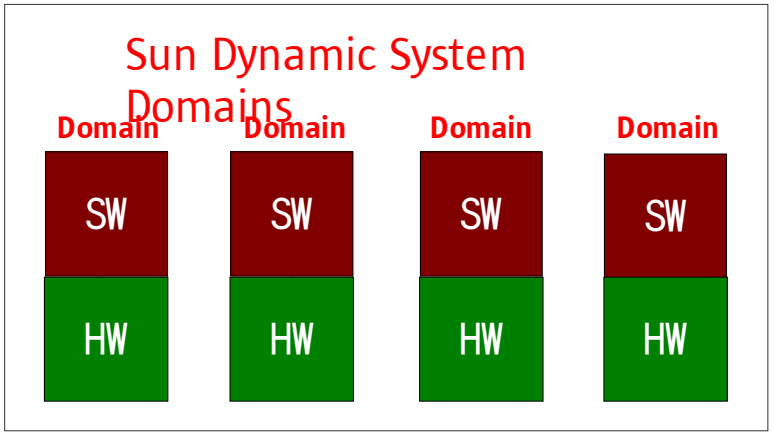

IBM pSeries server failures can impact multiple LPARs *(interconnect, I/O hypervisor, clock and memory)*

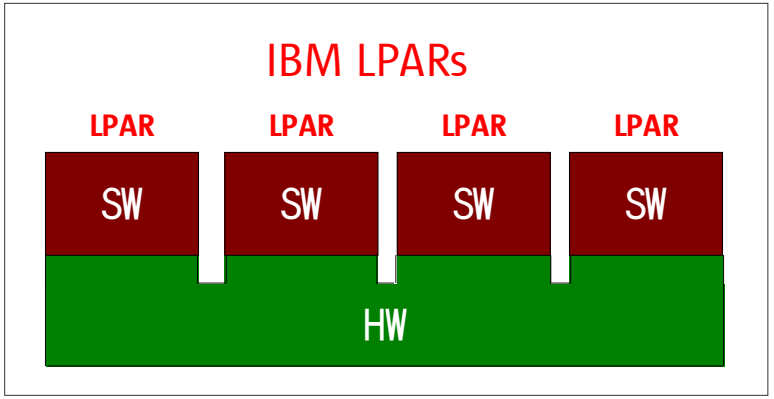

LPAR Shortcomings

- LPARs lack fault and security isolation; lack online serviceability of CPU and memory
- Small servers (p630) require additional I/O systems, and complex configurations to support LPARs
- Dynamic LPARs require AIX5L 5.2, ISV apps must be recertified
- IBM recommend larger LPARs; granularity claims are marketing hype
  - AIX requires 3 CPUs for CPU deallocation to work
  - 4 GB RAM minimum recommended LPAR
  - IBM recommends 10-slot PCI planars be dedicated to LPARs

HP nPartitions - Inflexible and Less Available

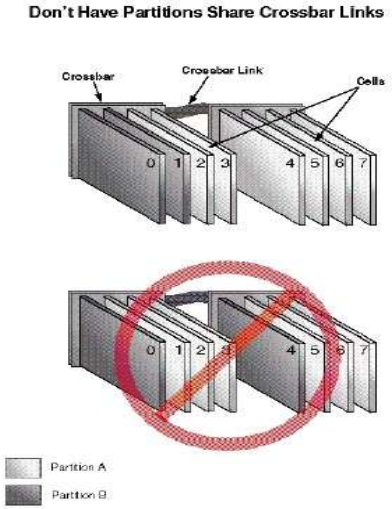

Limited #: 8 nPars

Max. 8 I/O drawers

"Each partition should consist of at least two cells.

"Each partition should contain at least two viable core cells."

Managing Superdome Complexes:

A Guide for HP-UX System Admins.

Inflexible

I/O tied to CPU/Memory Cell boards (14 slots per drawer)

Boards and resources cannot be moved without a reboot

CPUs w/in a partition must run at same speed

Less Available

Reboot required to move I/O drawer & resize partitions

HP vPars - Limited Security, Flexibility, and Availability

- Low Security
  - A root user on any vPar can run vPar commands for its vPar or any other vPar ('root or nothing')
- Poor Availability
  - Limited isolation from hardware failures or hardware resources bound to other vPars
- Inflexible
  - Memory and I/O cannot be moved between vPars without a reboot of both affected partitions
  - Vpars require that each "Top-Level" PCI bus adapter be assigned [fixed] to one vPar
  - Does not work with iCOD
  - Not all CPUs can be moved dynamically between vPars.

“For those who value the fault isolation that physical partitions bring, Sun's are the most established and the most dynamic.”

Illuminata, November 2002

“The ultimate in server consolidation is the ability to run multiple workloads on the same system under a single copy of the OS.”

Gartner, May 2002

Sun's Clear Advantage

- Best hardware partitions
  - Most flexible, dynamic, reliable
- Best software partitioning approach
  - Based on one OS instance: no overhead
  - Excellent resource flexibility
  - Mapped to “services”
  - Strong Failover and Management Framework

Sun offers the best of both worlds!

Ben Lenail

ben.lenail@sun.com

This box provides space for call to action text, URLs,

or any relevant info

Sun microsystems

We make the net work.