Reliability Availability Serviceability on Itanium Servers

Ken Pomaranski: Itanium and HP9000 HA Architect, HP

Thane Larson: Itanium and HP9000 RASUM Architect for Entry Level Servers, HP

Mily Tsou: HP-UX SSHA Planning Mgr, HP

Jerry Chin: Itanium Fault Management Architect, HP

Introduction & Hardware Contribution to RAS

Ken Pomaranski

Itanium & HP9000 Server HA Architect
Hewlett-Packard

Agenda

- Integrity Server RAS Philosophy
- Hardware HA
  - Feature description and selected deep-dives
  - SSHA Future direction
  - Platform comparisons
- HP-UX HA
- Fault Management

Integrity ASURance & RAS Philosophy

Integrity ASURance is the unique value proposition of providing top Availability, Serviceability, Usability, and Reliability to the HP product line.

The RAS philosophy is HP's way of describing the reliability, availability, and serviceability aspects of a computer system and ties this to its higher-level value propositions of high availability (HA) and low total cost of ownership (TCO).

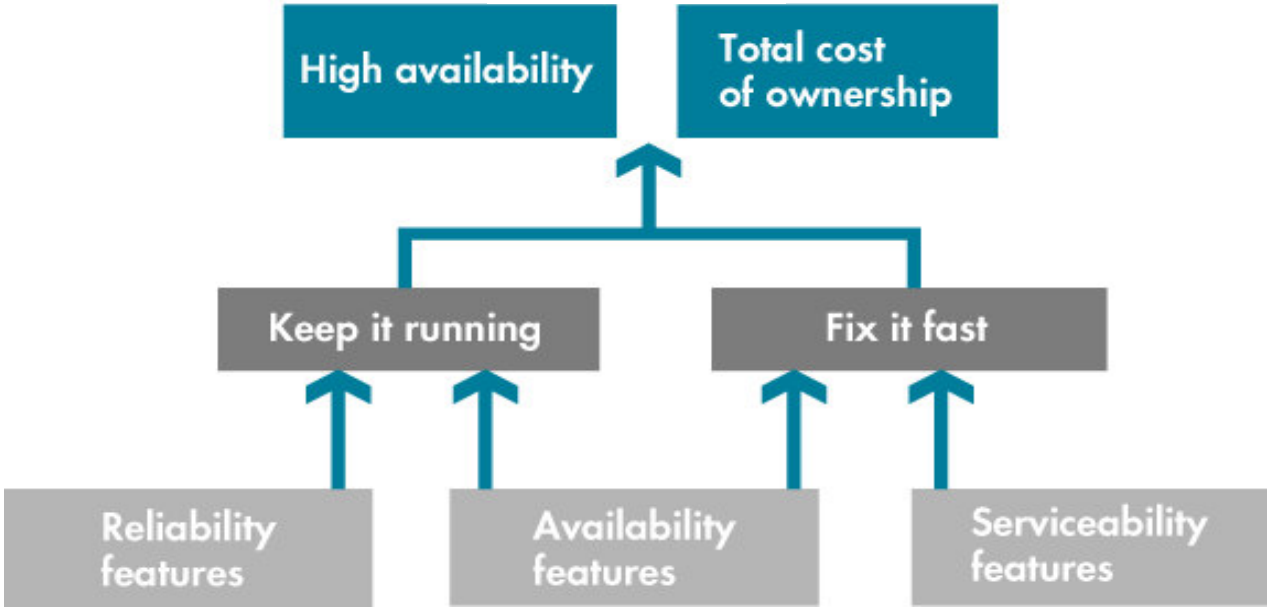

RAS on the Itanium platforms

HP has approached its Itanium-based Integrity platforms with the goals of building on our PA-RISC tradition of leadership in RAS and achieving industry leading RAS from the high-end to the entry-level.All systems – with only minor differences – give the same industry leading Reliability, Availability and Serviceability capabilities.

Note: Unless specifically mentioned throughout this presentation – RAS features on Itanium also exist on PA-RISC.

HA & Fault Management

The Fault Management and High Availability Design philosophy in HP is a coordinated and targeted approach to realizing our RAS objectives

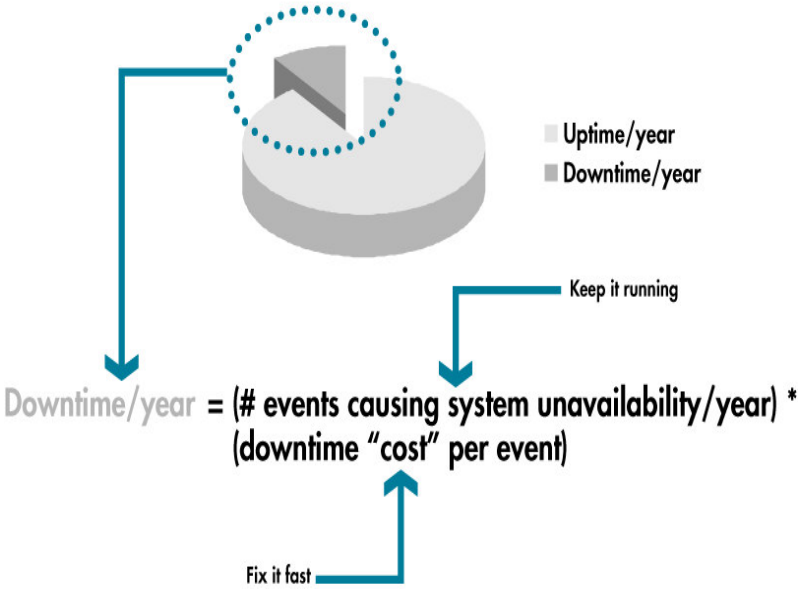

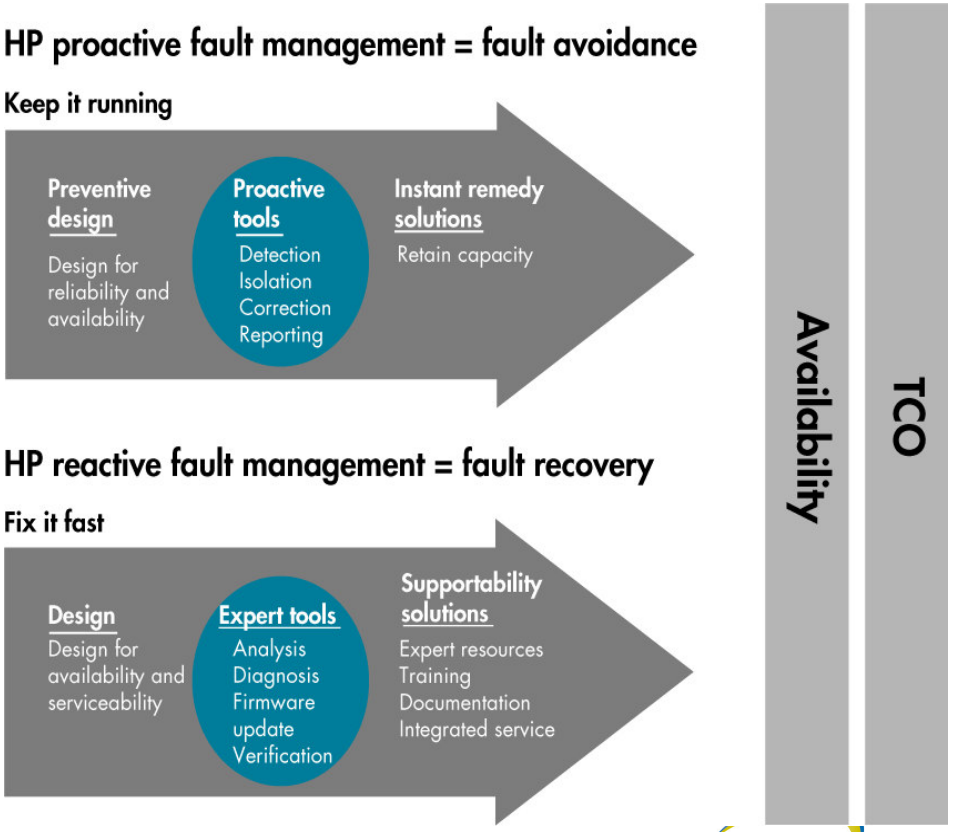

The High Availability Pyramid

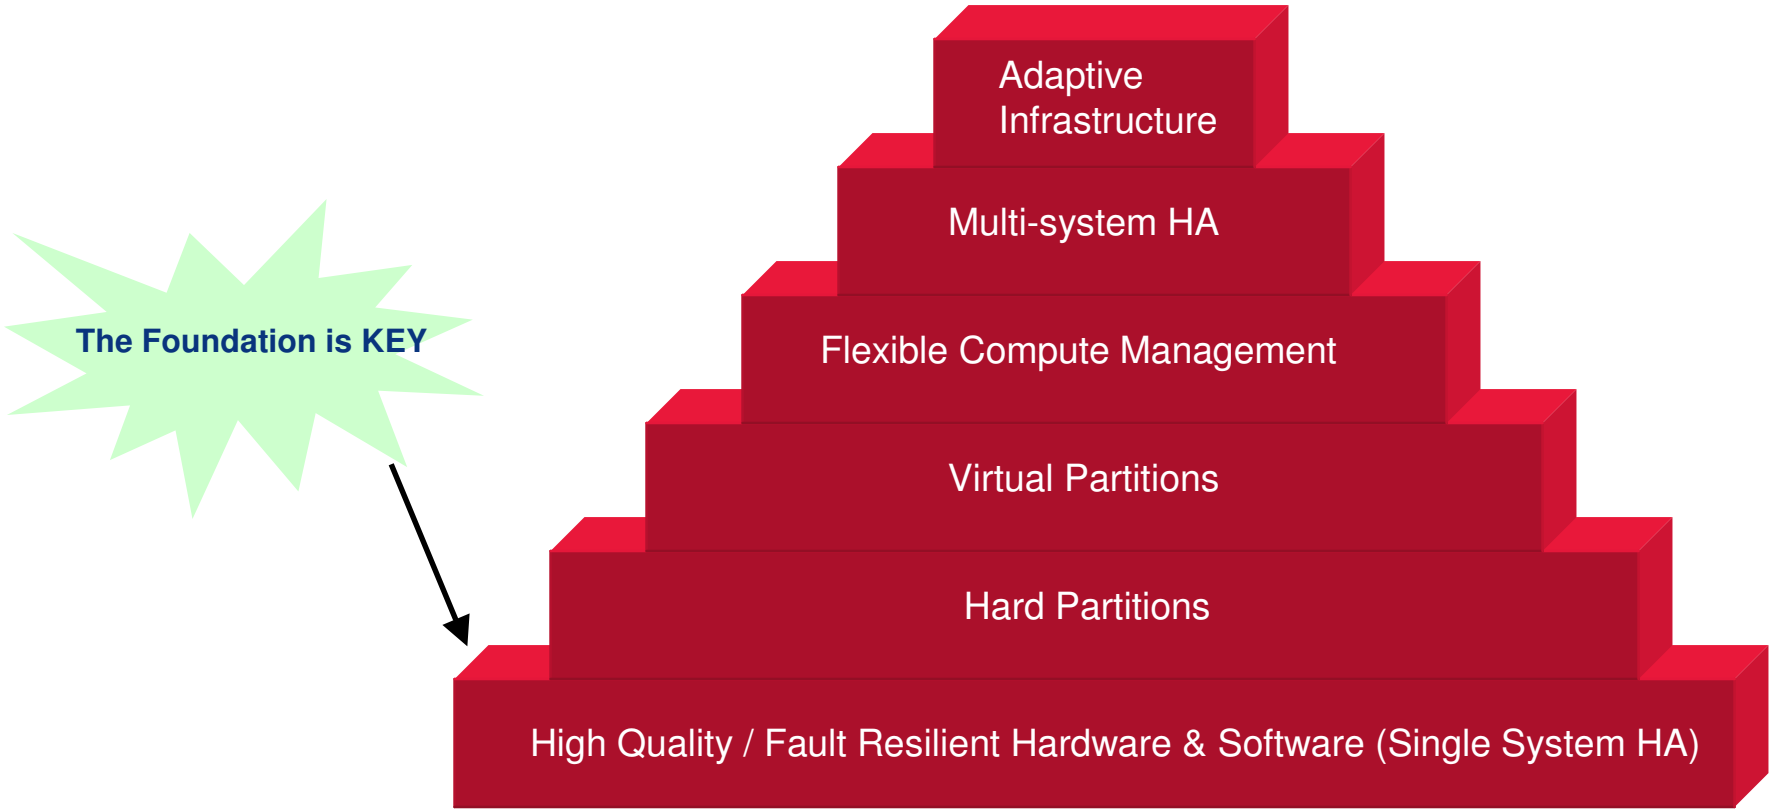

Hardware RAS Strategy

- Move to “Boot box once”: Resources shared by hard partitions should not go down, and the box should not need to go down to repair box-wide resources.

- System crash rate and planned downtime continues to go down aggressively over time:
  - Continuous improvement for each subsystem
  - HA features NOT ‘for show’.. attack true problem areas

- Evolve resiliency into self-healing
  - resiliency implies planned downtime to ‘fix’..
  - self-healing ‘erases’ error from the system, mitigating need for planned downtime. Improves TCE and warranty costs

Itanium hardware RAS features at-a-glance

Memory

- Dynamic memory resiliency with reactive memory scrubbing
- Parity protected ECC 'chip spare'
- ✓address control bus parity

Processors

- Dynamic processor resiliency
- 1-bit ECC & 2-bit detection on L1&L2 processor caches
- 1-bit ECC on all processor addr/data paths
- ✓ Dynamic processor resiliency & auto replace with spare
- ✓ Per Processor power supplies for de-allocation (Integrity Superdome)
- ✓ Per processor 'smart fans' for increased processor reliability (rx8620, rx7620)

Infrastructure

- Redundant and hot-swap power supplies (Uptime Institute)
- Redundant and hot-swap system fans
- Out-of-band management, Firmware updates, remote power control, health monitoring, remote console, and event logging
- Color coded cables and latches for faster upgrades and repairs
- ✓ Explicit hardware support for hard partitioning
- ✓ Hot-swap service processor (SP) & system tolerance to SP failures (Integrity Superdome)
- ✓ Redundant crossbar backplane DC converters (hot-swappable for Integrity Superdome)
- ✓ 'Spare Wire' recovery on communication links (Cell-to-cell, Cell-to-IO)
- ✓ Redundant DC power conversion for key subsystems

I/O

- Isolated I/O buses give error containment
- hot-swap disks with optional mirroring in HPUX
- ✓ Dual-path I/O with automated switchover

HP's proprietary CECs has the highest performance and reliable any CECs on the market

✓ Cellular System Advanced RAS Features

DMR & DPR Deep Dive

- Goal of DMR / DPR is to provide system 'self-healing'
  - CPU errors are common, due to chip complexity
  - Memory errors are common, due to sheer array size
- Smart, accurate fault management of processor / memory errors is critical to meet mission critical objectives:
  - Key differentiator, for both HW & SW
  - Most HW & SW do not deal with such errors correctly

  HP's field data has shown that DMR/DPR virtually eliminates system failures due to memory problems

DMR: Dynamic Memory Resiliency

Main memory failures are demonstrated to be a significant cause of HW downtime. Great care has been taken to address this failure mode on Itanium with these specific features:

- **Memory 'chip spare':** the ability of the system to continue to run in the face of any single or multi-bit chip error on a DRAM.
- **Dynamic memory resiliency (DMR)**: is the system's ability to de-allocate failed memory pages *online*. It works similar to Dynamic Processor Resiliency in that if a location in memory proves to be 'questionable' (i.e., exhibits persistent errors), that memory will be de-allocated online, with no customer visible impact. The number of spares is NOT limited by hardware, like in other industry available systems.
- **HW memory scrubbing:** The HW feature that automatically removes single bit errors (SBE) that reside in main memory. This is far superior to SW scrubbing mechanisms that skip locked memory locations.
- **Address Control Parity:** A DIMM set is only de-allocated if a fatal error that effects the quad has been detected. (ie, address / control parity error or multi-path data error). This is done to protect against data corruption. **HP is the only vendor to offer parity checking on the memory address / control Bus**

The combination of these features have proven to nearly eliminate memory as a cause of downtime in HP systems.

DPR: Dynamic Processor Resiliency Hot Spare CPU

dynamic processor resiliency

HP proven technology

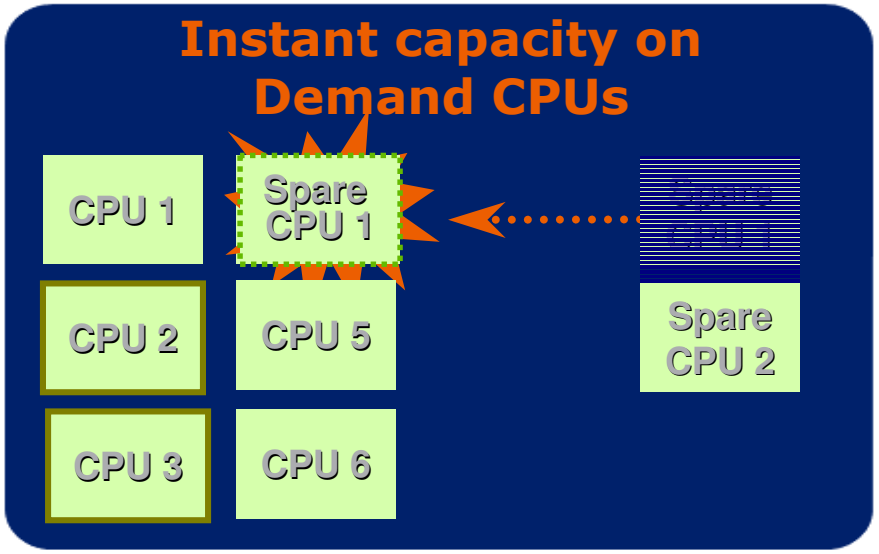

no system crash
no performance loss
no resource loss

completely transparent to the end-users

Inside DPR

1. An error occurs
2. Hardware detects error, corrects it, then generates an interrupt to the firmware

3. Firmware gathers error data, then vectors the interrupt to an OS level handler

4. OS level handler passes data to the diagnostic logging driver

5. The diagnostic logging daemon passes error data to the EMS monitor

6. EMS tracks thresholds and generates an event

* If exceeded, monitor sends requests to scheduler and firmware driver

    1. Scheduler stops further application on faulty CPU and **fails over** to **hot spare**

    2. Firmware driver sends request back to firmware to deallocate or 'fix' (Repair on Reboot – ROR) faulty CPU on next system boot.

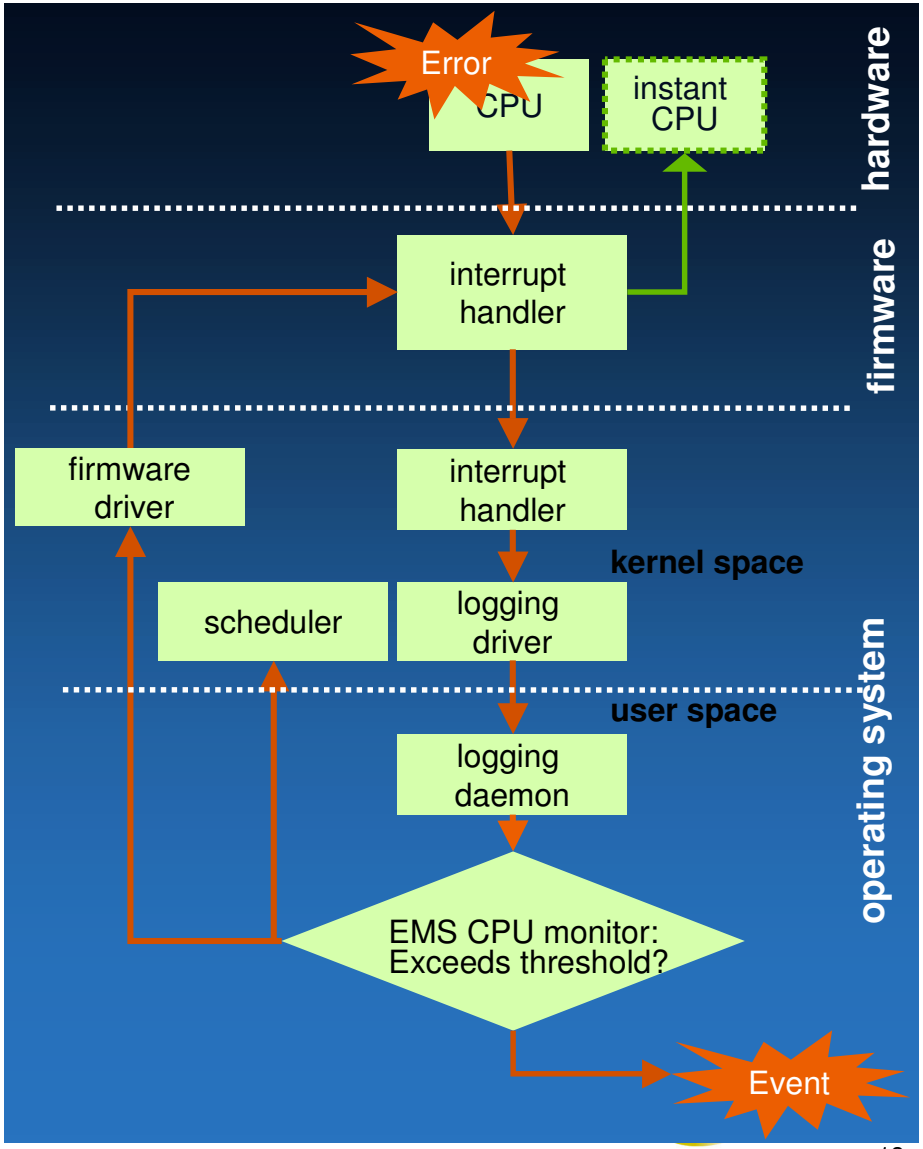

Hard Partitioning

- Only vendor with true hardware and software partition isolation:
  - HP's nPAR isolation for Superdome is field demonstrated to be greater than 100 year MTBF
  - Soft partition isolation (e.g. IBM's Lpars) tend to run about 5-10 year MTBF
- Enables a true server consolidation solution
  - Can physically service any nPAR while other nPARs are booted and running. (includes power up / down, HW addition and removal, etc..)
  - Superior to other partitioning solutions (IBM p690 requires all Lpars to come down to replace a failed system fan, for example)
- Great use models:
  - Production & pre-production machines in same box
  - Ability to size and re-size partitions based on need

Hard Partition Error Containment

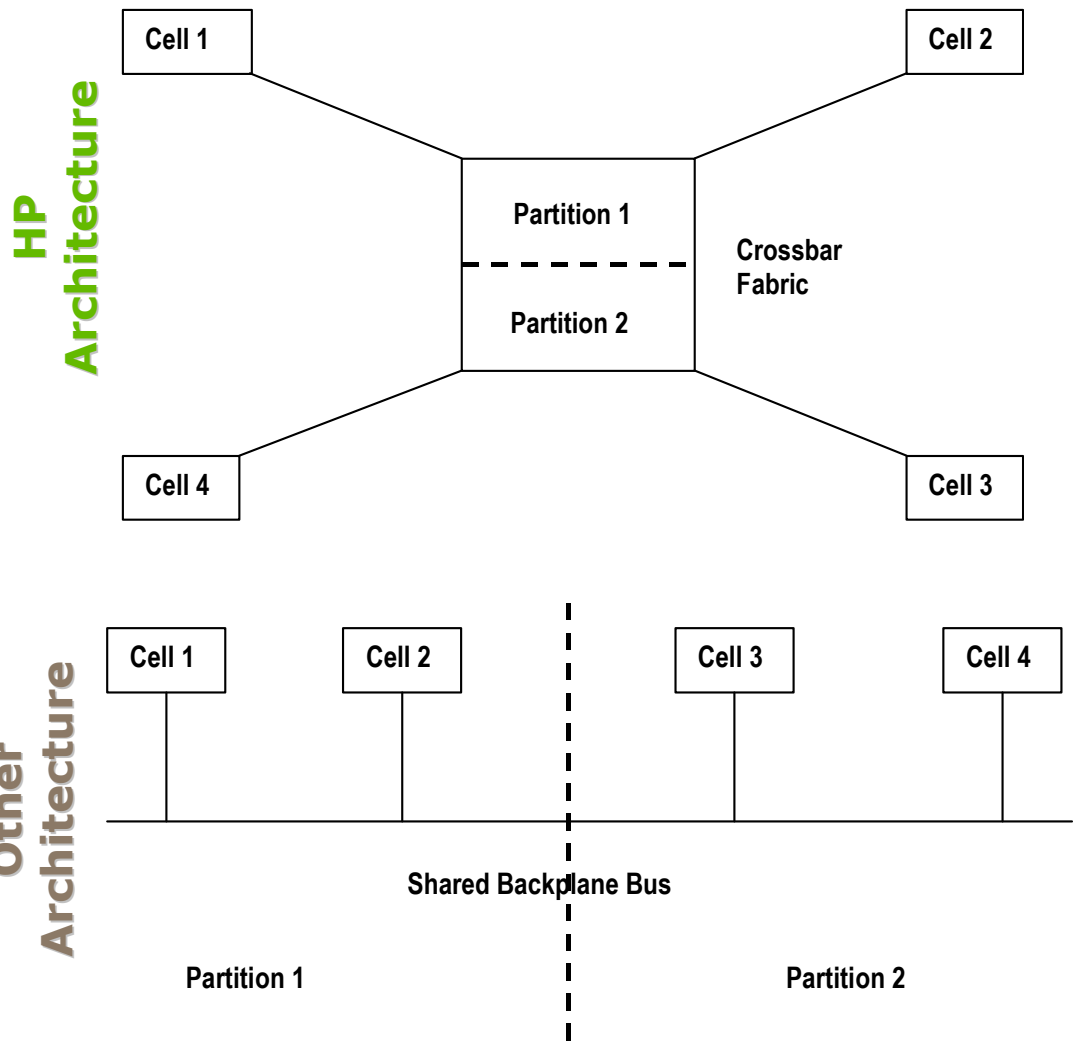

On the HP system, the crossbar logically separates the two physical partitions to provide performance and isolation.

A shared backplane has all its cells competing for the same electrical bus. In this design, a snoopy bus-coherency scheme requires all transactions to be broadcast to and processed by all system cells. The high-queuing delays and saturation of the shared backplane bus can limit performance scaling, and results in many shared failure modes

Itanium HA Features, Platforms, and Operating Systems

Thane Larson

RASUM Architect for Entry Level Itanium and HP9000 Servers

Hewlett-Packard

HP Delivering HA Solutions from Scale-up to Scale-out

- **Cellular Systems**
  - Optimum scale-up from 1 to 128 processors
  - Multiple hard and virtual partitions
  - Minimum single points of failure
  - Advanced DPR and iCOD
  - In the box scalability
- **Non-Cellular Systems**
  - Optimum scale-out with multiple systems and HA clusterware like Serviceguard and Oracle 9iRAC
  - High single-system reliability
  - Unlimited scalability
  - Flexible serviceability models
- **Both can be run single-system or clustered**

HP is uniquely positioned to deliver more

- **Seamless 32/64-bit choice, innovation and value**
- **Driving the future of standards across OS, application & platform**
- **Delivering the best choice of price:performance from scale-up to scale-out**

Leading adaptive management	End-to-end, common storage integration	Proven enterprise solutions & partners	Unmatched 32/64-bit expertise, service and support

ProLiant

Opteron performance price

Leading x86 value and HA

1-8P
Industry Standard
Windows, Linux, NetWare

Integrity

Ultimate scalable performance and reliability

1-128P
Industry Standard Linux, Windows
& Mission Critical HP-UX, OpenVMS
99.95% -> 99.999% Availability

NonStop

Bulletproof data integrity and instant data processing

Fault Tolerant NSK
99.999%+ Availability

Choosing a System

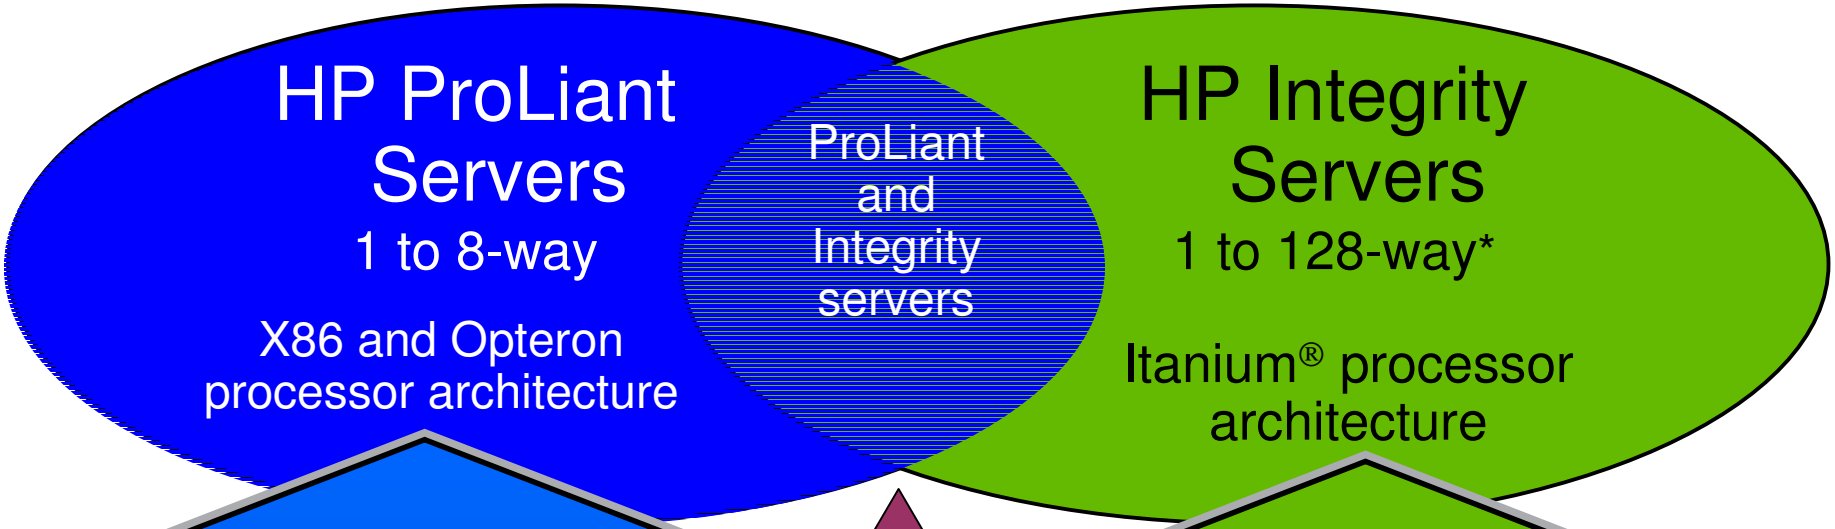

- Small to medium scale application and databases
- Well-defined, less-complex workloads
- Primarily front-end/network edge & application tier
- Scale out and small to mid-size scale up

- Large scale applications and databases
- Complex workloads – technical and commercial
- Primarily back end DB & application tier
- Enterprise scale up and scale out
- Server consolidation

Customer-specific needs driven

Processor RAS Differences

	Itanium*	HP PA-RISC*	Opteron*	Xeon*	*Sun Sparc*	*IBM POWER4*
Chip thermal sensors & management	**YES**	NO	**YES**	**YES**	NO	Limited
Full Cache Parity / ECC	**YES**	**YES**	SOME	SOME	SOME	**YES**
CPU Data Bus ECC	**YES**	**YES**	NO	NO	**YES**	**YES**
†Data Poisoning/Signaling for Error Recovery	**YES**	NO	NO	NO	NO	**YES**
‡Enhanced MCA handling & Error Logging	**YES**	Proprietary Method	Limited	Limited	Limited	Proprietary Method
Dynamic Processor Resiliency	**YES**	**YES**	NO	NO	Limited	**YES**

*Processors used in HP Products

†To prevent an unrecoverable error from propagating and corrupting data, the “Data is Poisoned” in a way that marks it as permanently bad so that the system will either force a reread of that data, or in the worst case forces a crash.

‡Enhanced MCA handling and Error Logging are tools provided to allow for more granularity of error containment. This allows errors to not propagate and to let recovery to have the least impact as possible.

Memory RAS Differences

	Itanium	HP PA-RISC*	Opteron*	Xeon*	*Sun Sparc*	*IBM POWER4*
Data Bus ECC protection	**YES**	**YES**	**YES**	**YES**	**YES**	**YES**
Address Bus Parity protection	†**YES**	†**YES**	NO	NO	NO	NO
Advanced ECC / Chip Spare / Chip Kill	**YES**	**YES**	NO	**YES**	NO	**YES**
Page Deallocation	**YES**	**YES**	NO	**YES**	NO	NO
Dynamic Memory Resiliency (DMR)	**YES**	**YES**	NO	NO	NO	NO
‡Mirroring / RAID	NO	NO	NO	**YES**	NO	**YES**

*Processors used in HP Products

†Cellular systems only with proprietary DIMMs

‡Memory mirroring/RAID is an expensive way (in terms of $ and performance) to accomplish the same protection as DMR

HP-UX, Linux, Windows RAS Differences

Feature	HP-UX	Linux	Windows
Dynamic Processor Resiliency (DPR)	●		
iCOD	●		
Superior Online Diagnostics, STM	●		
Online Burn-in Tools, Exercisers	●		
USB devices hot-pluggable		●	●
Predictive Monitoring	●		●
Serviceguard	●	●	
Dynamic Memory Resiliency (DMR)	●	●	●
Partitions in Cellular Systems	●	● *	● *

*Hard Partitions Only

HP-UX contribution to RAS

Itanium supports HP-UX, Windows, Linux and OpenVMS. The next portion of the presentation emphasizes HP-UX's contribution.

Mily Tsou

HP-UX SSHA Planning Mgr

Hewlett-Packard

HP's Vision for the Adaptive Enterprise

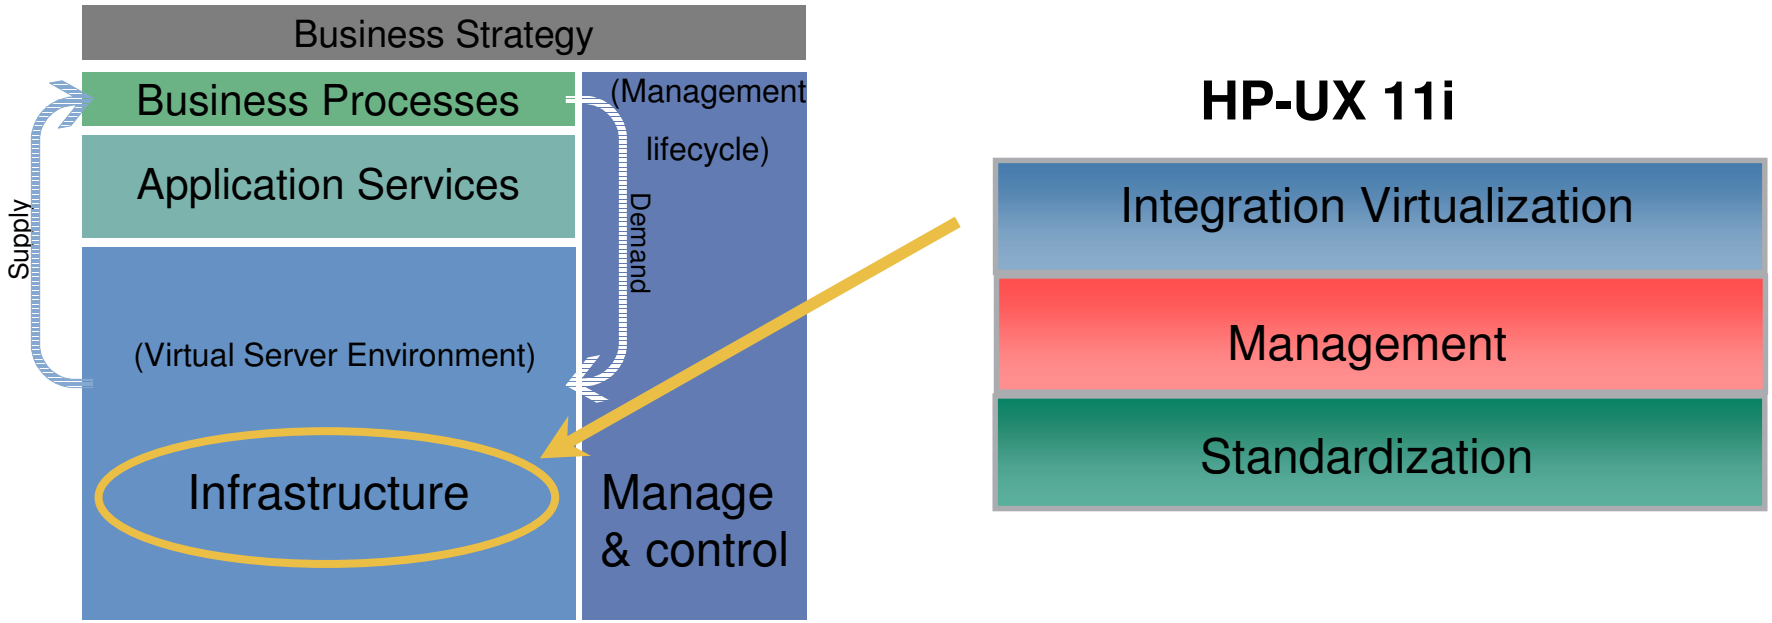

HP-UX 11i – the foundation of adaptive infrastructure

react quickly to dynamic business demands with virtualization

HP-UX 11i A Solid Foundation
Delivered Today for the Adaptive Enterprise

Best-in-class O/S infrastructure

self-tuning

self-healing

fundamentals

Integration Virtualization

load balancing
- icod/ppu
- process resource manager (PRM)
- UDC multi-OS capacity metering and mgmt

dynamic tuning
- dynamic kernel tunables
- variable page sizing

partitions
- nPars, vPars
- Psets
- memory resource groups (MRG)

software deployment
- igniteUX
- software distributor (SD)
- Fine granularity OE Installation configuration

Management

fault detection and isolation
- Event monitoring services (EMS)
- FRU level fault isolation
- High Availability Observatory (HAO)
- System Configuration Repository (SCR)/System Inventory Manager (SIM)

fault correction
- Automated Analysis
- Intelligent recommended actions
- Diagnostics tools (ODE & STM)
- HP MC analysis tools
- OnlineJFS
- VxVM

goal based workload management
- Workload Manager WLM

intelligent configuration
- SCM and SAM
- Partition Manager with Big System Views

Standardization

reliability
- hw & sw quality
- patch management

hw/sw resiliency
- Dynamic Processor Resilience
- Dynamic Memory Resilience
- PCI and PCI-X OL AR
- multi-path I/O
- APA
- HA TCP

recovery
- root disk journaling
- Fast boot /reboot via parallel ioscan
- Serviceguard

secure
- intrusion detection
- IPFilter & Bastille lock-down tool
- IPSec/Kerberos Authentication
- Install-time security hardening

Connectivity

- IPv6, mobile IPv4

Kernel Configuration

better application performance via dynamic kernel tuning

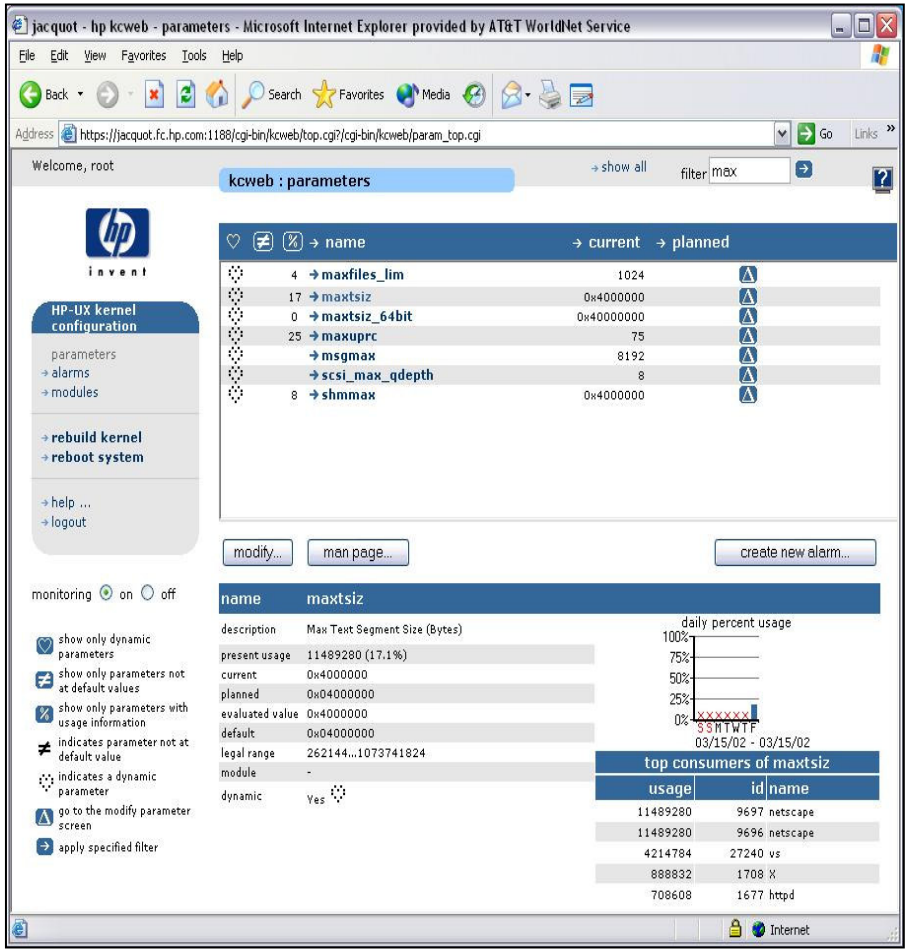

what does it do?

key kernel parameters can be dynamically tuned based upon changing system conditions

- key kernel parameter changes do not require a reboot
- reduced planned downtime
- intuitive graphs of parameters
- notification of potential issues to monitor kernel resource usage
- create/import configuration file
- CLI preview

increased automation

HP-UX Dynamic Tunables in HP-UX 11i

hp-ux 11i v1 **PLUS** hp-ux 11i v1.6 **PLUS** hp-ux 11i v2

- core_addshmem_read
- core_addshmem_write
- maxfiles_lim
- maxtsiz
- maxtsiz_64bit
- maxuprc
- msgmax
- msgmnb
- scsi_max_qdepth
- semmsl
- shmmax
- shmseg

5%

- ksi_alloc_max
- maxdsiz
- maxdsiz_64bit
- maxssiz
- maxssiz_64bit
- maxswapchunks
- max_thread_proc
- nkthread
- nproc
- secure_sid_scripts
- shmmni

40%

- dbc_max_pct
- dbc_min_pct
- maxfiles
- nfile
- nflocks
- semmns
- semmni

50%

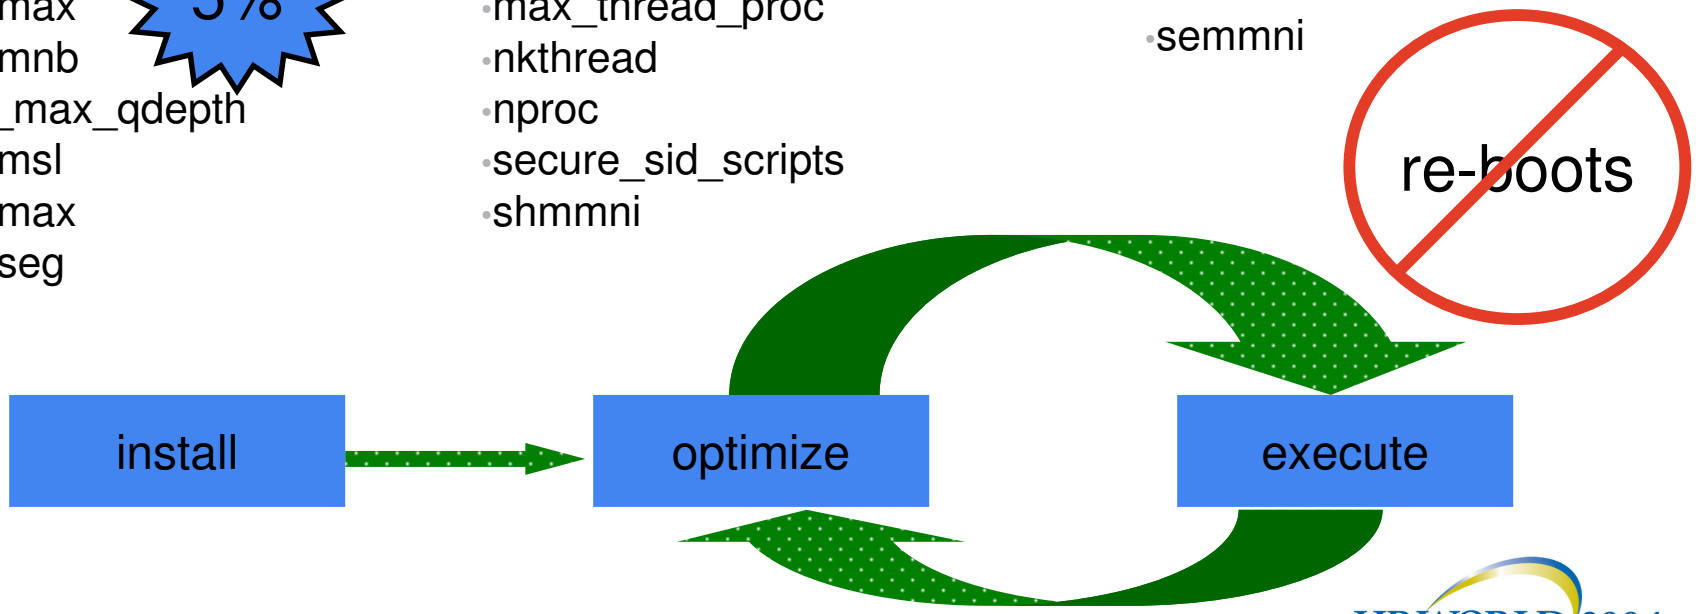

HA Observatory Estimate

Peripheral Devices (pdweb)

supports hot-pluggable and hot-swappable I/O cards

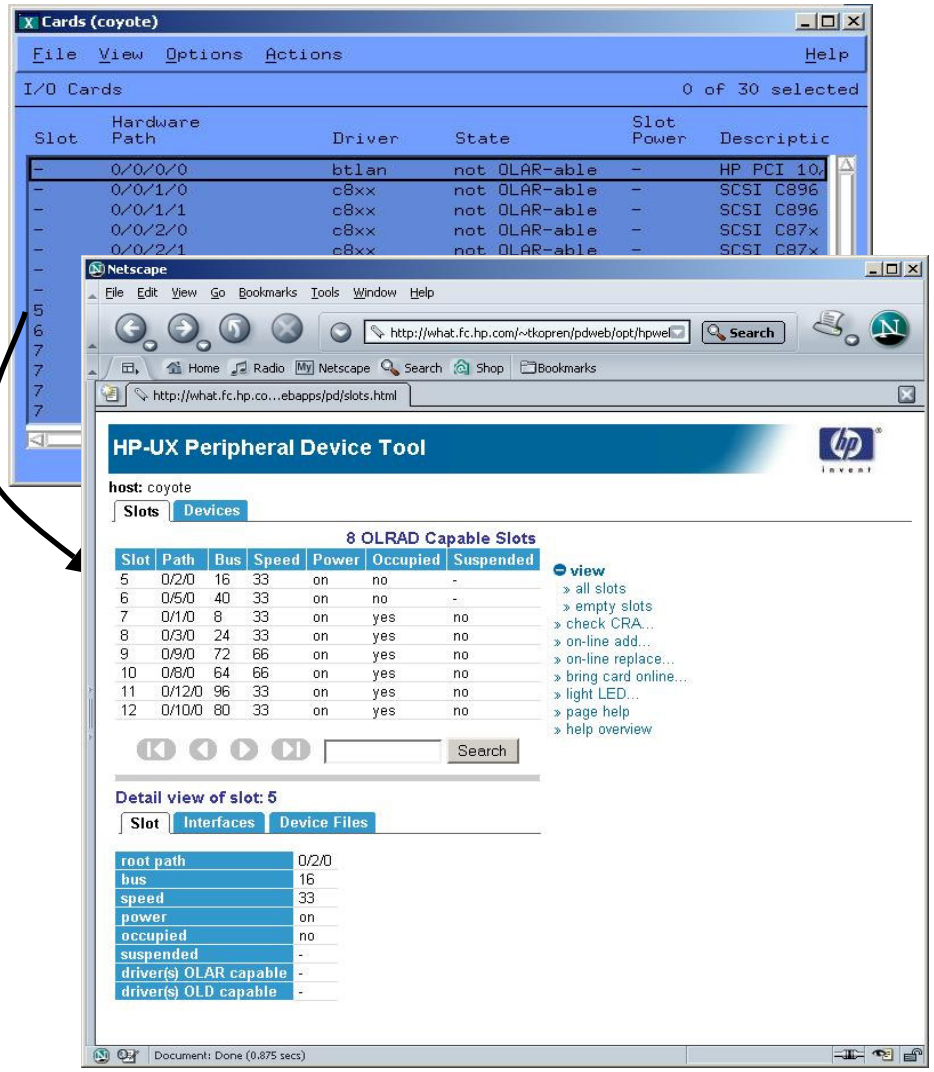

- Same functionality as in SAM PD (replacement); SAM and SCM 3.0 launch
- PCI and PCI-X OLAR support to reduce planned downtime
- Displays slots and devices
- Check CRA (critical resource analysis)
- Improved functionality in CLI
- Command preview
- Creates device files
- Web-based interface
- GUI, messages, help localized

HP Partitioning Continuum Strategy for HP-UX

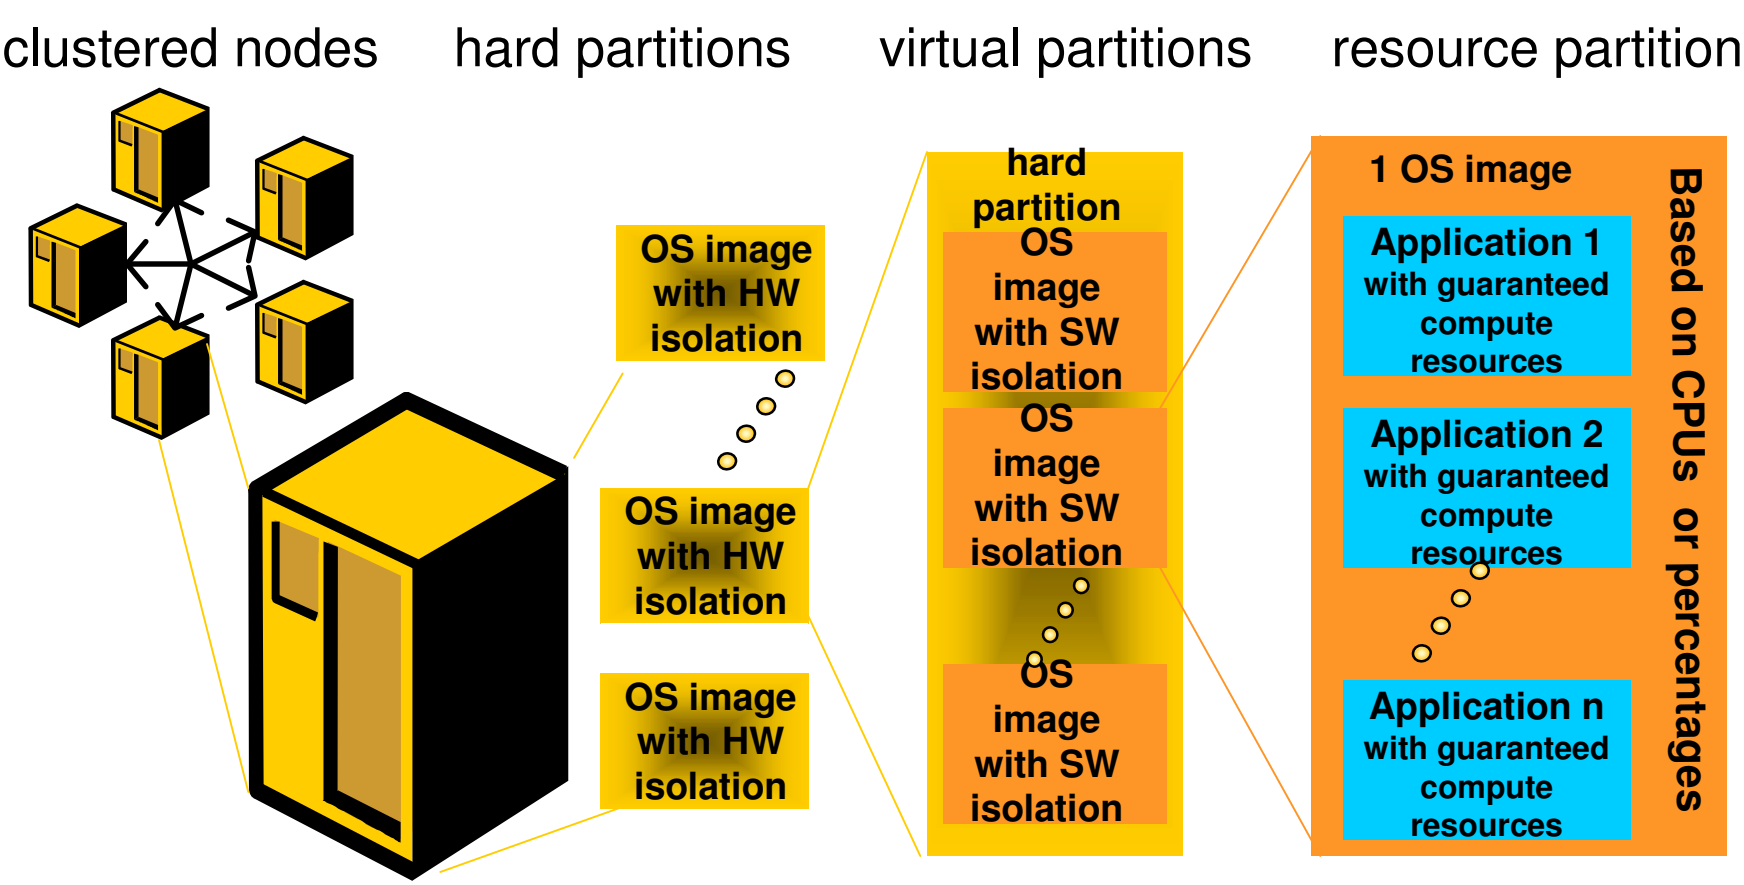

HP Partitioning Continuum for HP-UX 11i

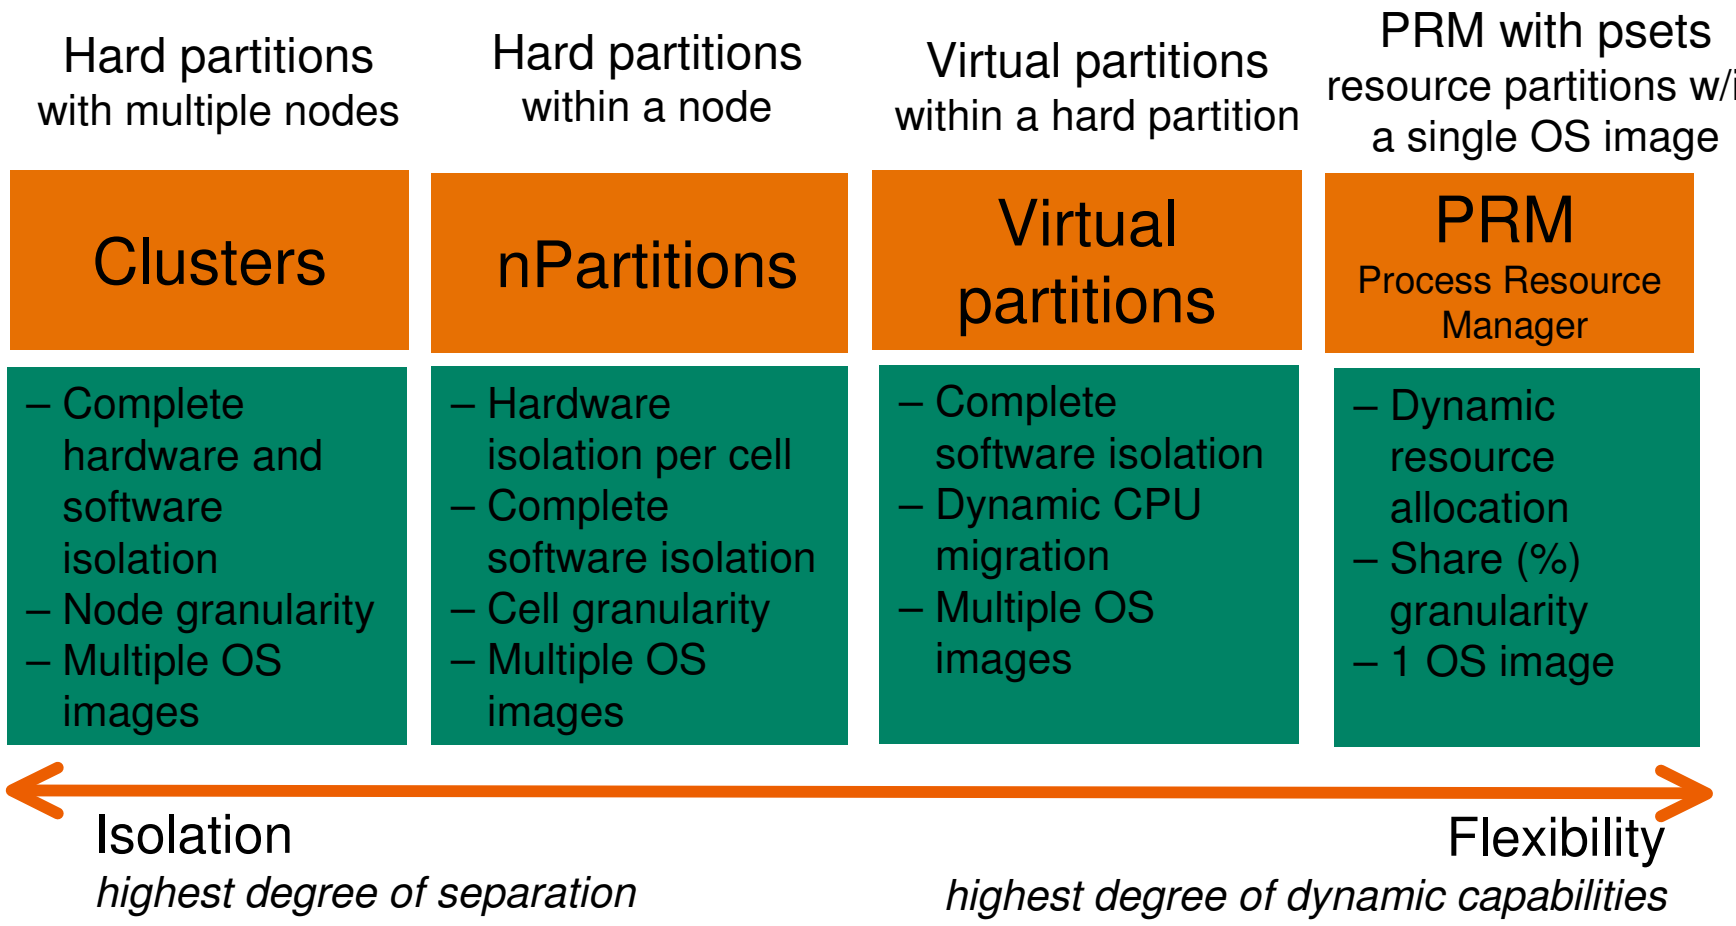

New Partition Manager Benefits

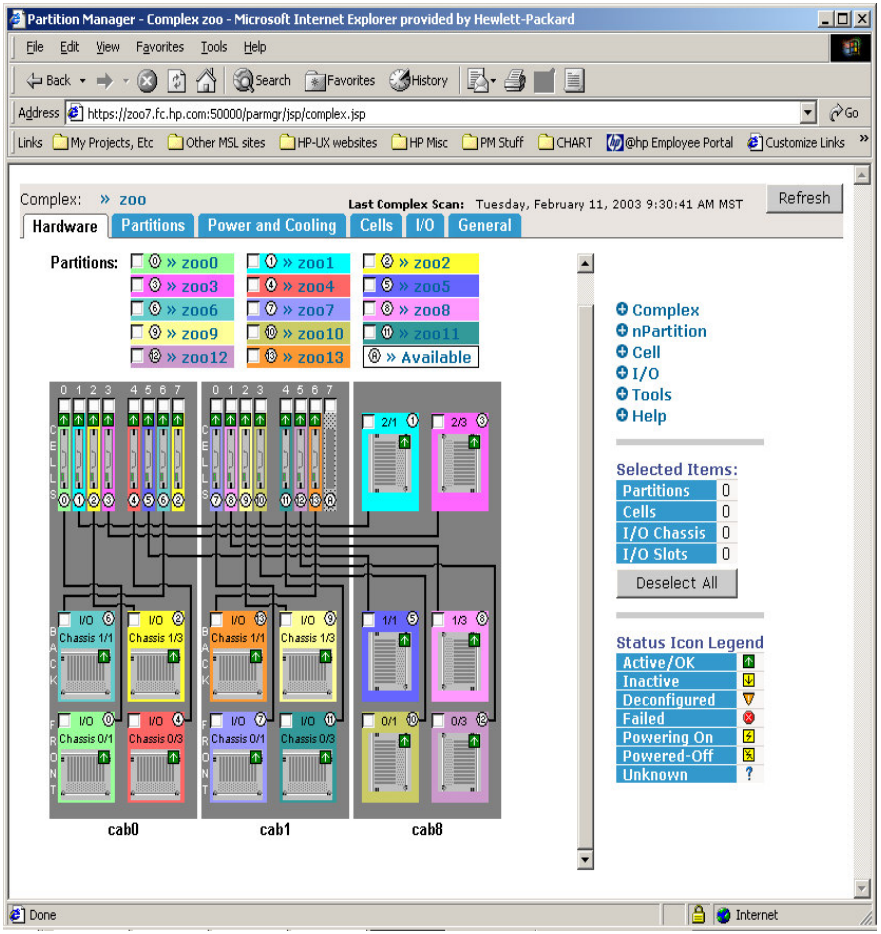

Today's HP Partition Manager features **PLUS**…

- New intuitive web interface
- Graphical "big picture" views & configuration of complexes including:
  - nPars
  - I/O
  - Cells
  - Power & cooling
- Hardware component status
- Remote administration capability
- Manage Itanium servers (rx7610 & above) & new PA servers (rp7420 and above) from the Integrity system
- Preserve compatibility with iCOD/PPU
- Non-root user read-only access

Integrated Virtualization: HP Virtual Server Environment for HP-UX 11i

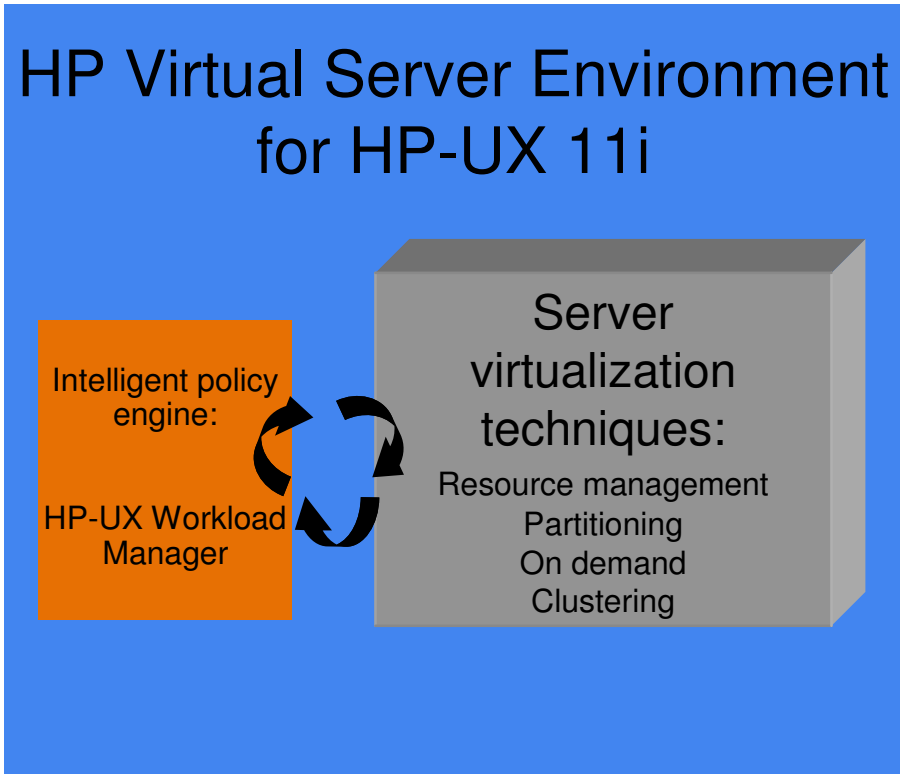

- Automates the virtualized environment
- Goal-based or policy-based resource management
- Exclusive integration:
  - CPU resource allocation
    - –within and across partitions
    - –in between multiple apps in a single OS image
  - Automatic reallocation of resources upon Serviceguard package activation
- Application transparent
- Application-specific toolkits

Optimized utilization with service levels and agility

HP-UX VSE for the Adaptive Enterprise – in production use today

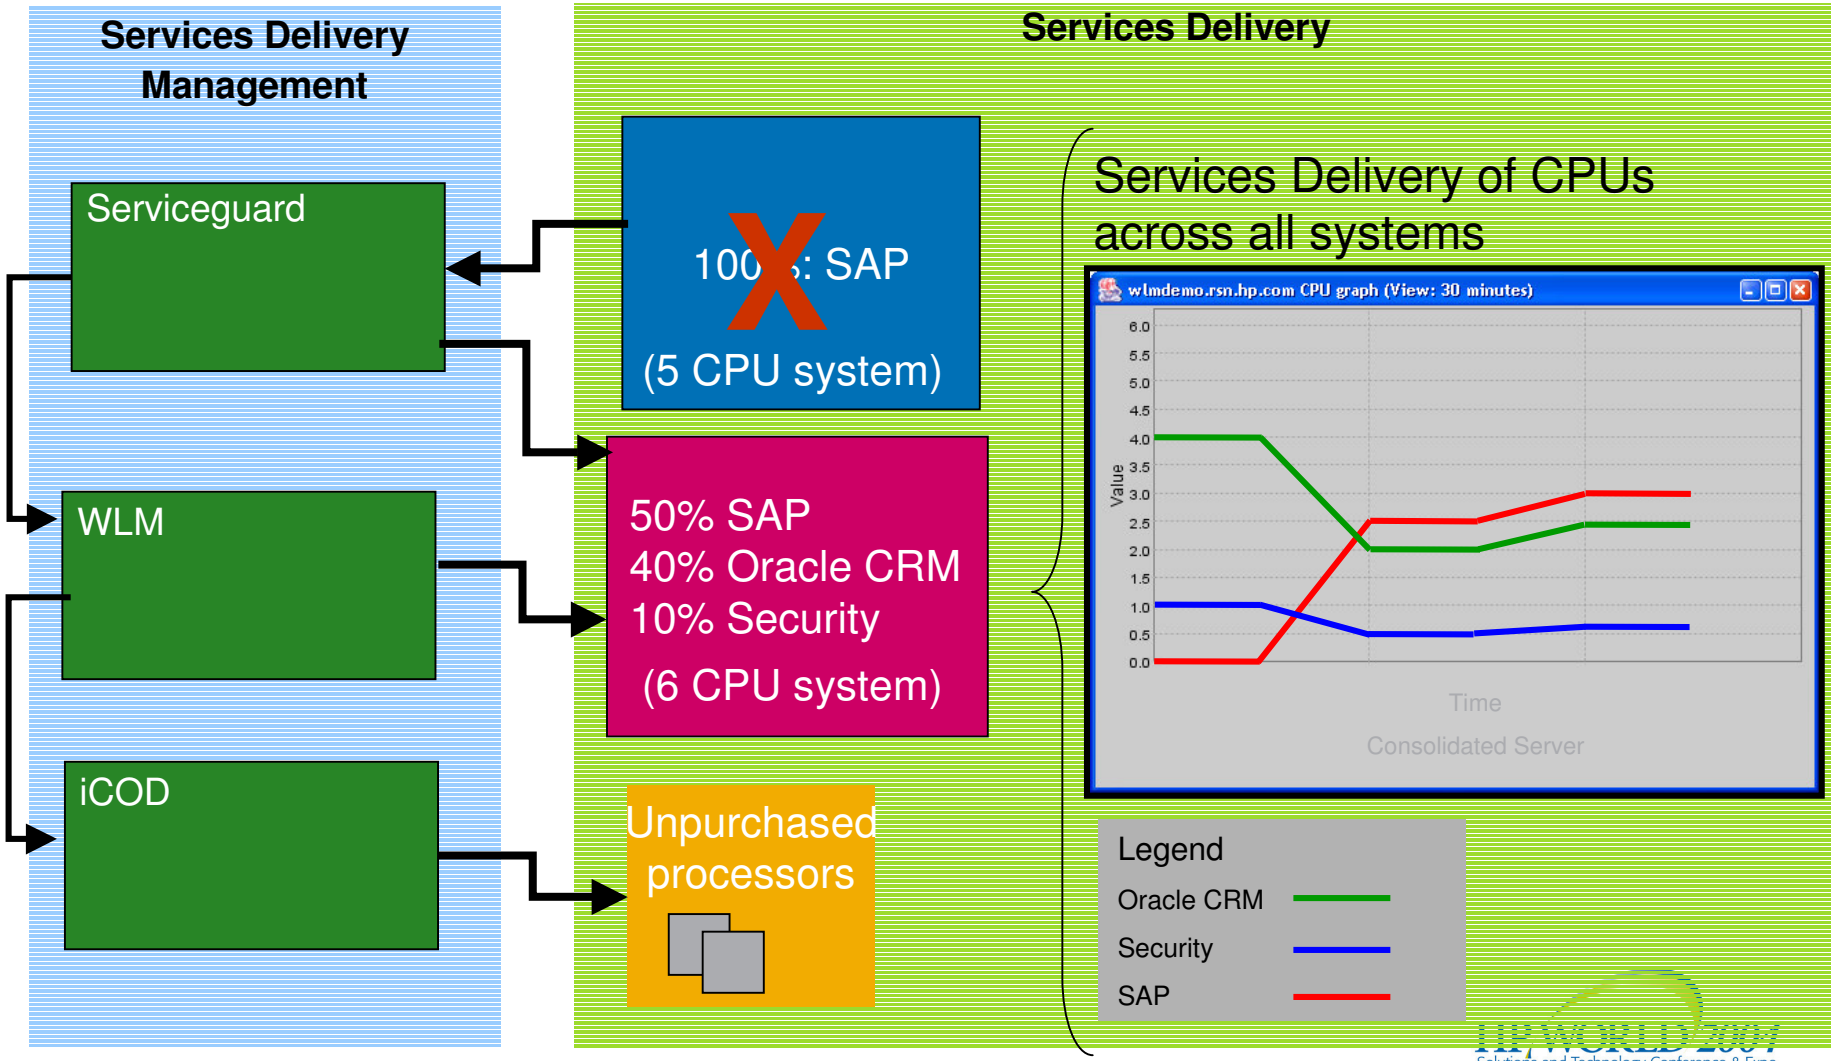

Enhancing HP VSE for HP-UX 11i – proven foundation for the Adaptive Enterprise

HP Virtual Server Environment (VSE)

Control

	Availability	Partitioning	Utility Pricing
Existing components	Workload Manager Systems Insight Manager (SIM)	Serviceguard SGeRAC	nPars vPars Process Resource Manager / pSets	Instant Capacity Temporary Instant Capacity Pay Per Use
New components	Global Workload Manager	Enhanced SGeRAC Serviceguard Fast Failover Mixed clusters Enhanced Serviceguard Manager	Secure resource partitions vPars on Integrity HP Integrity Virtual Machines with sub-CPU and shared I/O	
New packaging	HP VSE Suites for HP-UX 11i			

HP Virtual Server Environment: innovation based on standards

		HP MC VSE Suite for HP-UX 11i	IBM VE Server Suite for pSeries
Control	Goal-based workload management	YES	NO
	Multi OS management with provisioning	YES	multi OS planned
Availability	Integrated management console for virtualization	YES	NO Director? VE Console? HMC?
	Integrated mission-critical virtualization	YES	NO HA not integrated, HACMP not yet supported on AIX 5.3
Partitioning	Hard Partitions	YES	NO
	Soft Partitions	YES	YES
	Sub CPU Partitioning	planned	YES
	Secure resource partitions	planned	NO
	Resource Partitioning (PRM, pSets)	YES	YES
Utility Pricing	Integrated Utility Pricing with virtualization	YES	NO

High Availability product portfolio

Availability through manageability:
- Mainframe capabilities
- Monitoring and management
- Ease of use
  - Serviceguard manager
  - EMS HA monitors
  - OpenView/Operations & NNM
  - Systems Insight manager
  - Servicecontrol manager

Extended Campus Cluster, Metrocluster, Continentalclusters, StorageWorks Cluster Extension XP

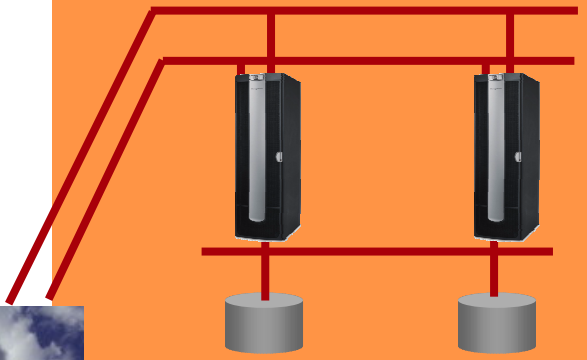

Business Continuity technologies

Single system availability
- Servers
- Disk arrays
- Network

Multi-system availability
- Serviceguard (HP-UX & Linux)
- SGeRAC

Network web application database

Solution enablement: Improve end-to-end availability including application stack
- HA Toolkits
  - SGeSAP, NFS, ECMT
- HP C&I; HP Cluster Consistency Service

The Fault Management contribution to RAS

The Itanium platform is designed with the OS to ensure industry leading RAS. The Following portion of the presentation highlights the fault management and serviceability capabilities.

Jerry Chin

Fault Management Architect for HP's Itanium Products and Services

HP Server Fault Management in the Adaptive Enterprise

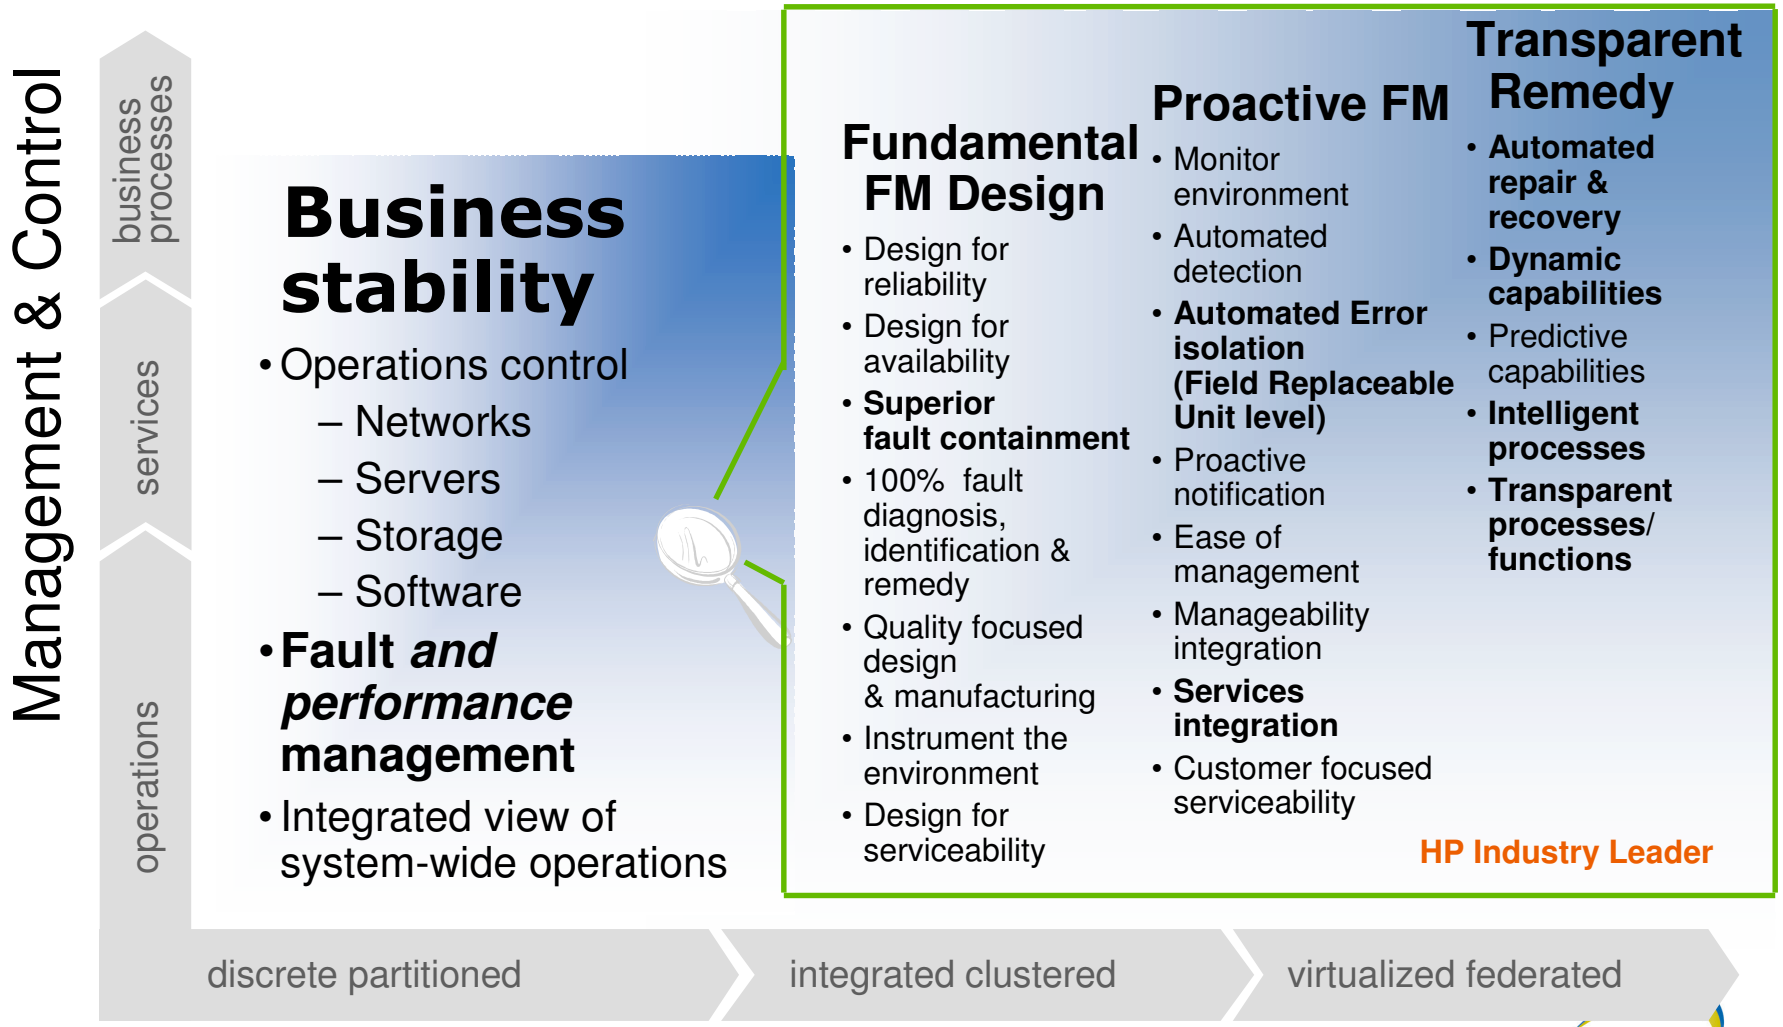

IPF Fault Management Architecture

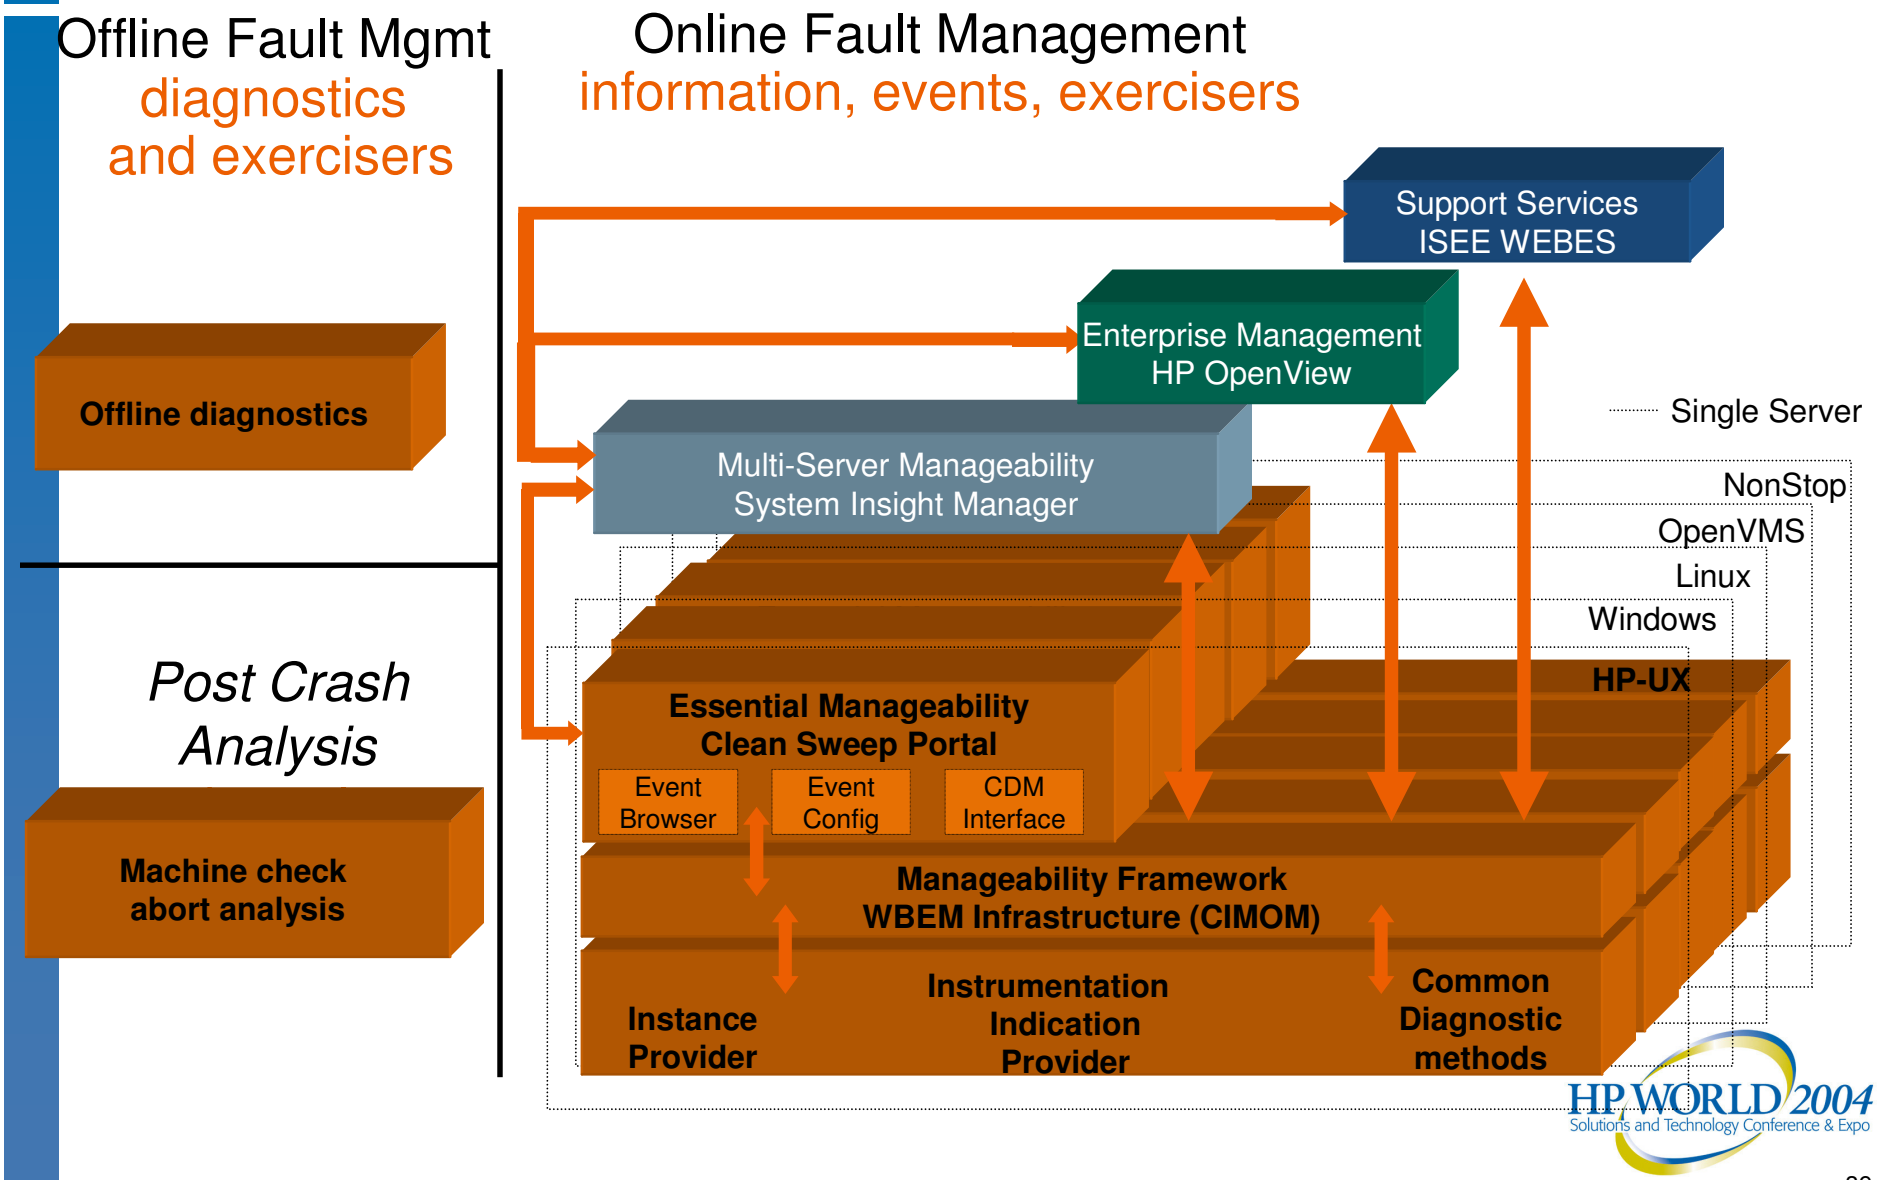

Serviceability example - rx4640

Offering the most effective HW problem determination & repair tools

- ✓**Quick Find Diagnostic Panel, hardware-based, external visual health indicator/status of subsystems**
- ✓**Modular FRUs**
  - **-highest number of POP parts, 90% of entire FRU list**
  - **-tool-less repair**
- ✓**Color Coded Labels and Latches**
- ✓**Redundant hot-swappable fans**
- ✓**Redundant, N+1 PSUs**
- ✓**Hot-swap HDDs**
  - **-2 internal, RAID capable**
- ✓**OL***
  - **-PCI-x slots**
  - **-Doorbells and latches**

- ✓**internally Traditional PA-RISC troubleshooting tools now available for IPF**
- **- diagnostics, offline & online (hp-ux)**
- **- exercisers**
- **- monitors (all OS's), that integrate into management tools and into HP Services**
- **- MCA analyzer (field tool)**
- ✓**Remote Connectivity**
  - **-LAN, http included**
  - **-Serial**
- ✓**Customer installable FW**
- ✓**Cable management**
  - **-Minimized**
  - **-Arm for external rack mounts**

The Bottom Line

HP's RAS features:

- Dynamic Processor Resiliency
- Dynamic Memory Resiliency
- PCI/PCI-X OLAR
- Dynamic kernel tunables
- Partition fault (hardware and software) isolation
- I/O Error Containment
- Uptime Institute Certification
- Hardware memory scrubbing
- Address parity checking
- Wire sparing
- Redundant/hot-swappable fans and power supplies
- Fault Management/Isolation/Recovery
- World-class clustering solutions (Serviceguard)
- … and future HA enhancements

result in → Higher Single System Resiliency + Improved Recovery Times = Higher levels of availability!

HP WORLD 2004
Solutions and Technology Conference & Expo

HP WORLD 2004

Solutions and Technology Conference & Expo

Co-produced by:

interex

shared knowledge • shared power

encompass

AN HP USER GROUP

RECOMMENDED TRAINING VENUE FOR THE

HP Certified Professional