




Optimize your Oracle workload thanks to HP VSE

Yann Allandit, HP
Oracle pre-sales consultant



HP & Oracle Alliance

ORACLE



- Over 25 years of partnership
- Over 140,000 joint customers
- Leading Emea Market share on Oracle with 49% installations on HP
- Executive alignment
- Joint engineering and product development
- HP is Oracle's largest customer
- Oracle tests its products on HP systems
- 13 worldwide joint Solution Centers
- HP is Oracle's first major server vendor to offer a joint solution offering



HP/Oracle Cooperative Technology Center



- Located at HP in Germany, France & Oracle UK
- 7 HP & Oracle employees in one team
- Delivering services to HP, Oracle & Partners
- Founded in spring 1994
 - Technical pre-sales assistance/consulting for partners & customers
 - customer workshops/demos, proof of concept projects, benchmarks, RFX contributions, architecture consulting, complex sizing and configuration assistance, customer presentations,
 - Know-How Transfer & Country Enablement
 - Evaluation and tests of new products/features
 - Technology transfer to and from US labs

Agenda

- Why HP virtualization?
- Benefits of Oracle in HP VSE
 - Implementation case
- HP VSE components
 - G/T/iCAP,
 - nPARs,
 - vPARs,
 - HPVM,
 - gWLM





Virtualization challenge

Customers pain points

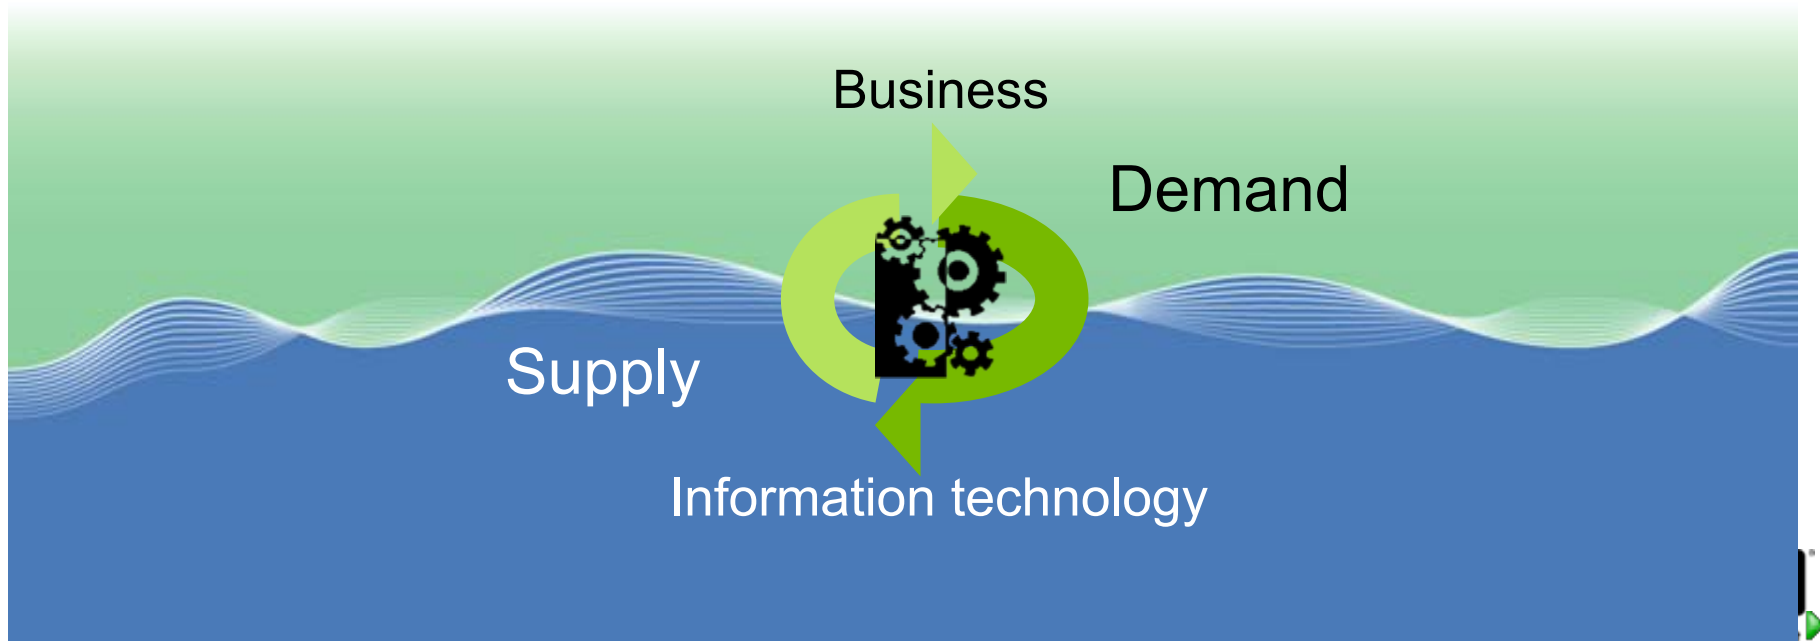
- “We’ve got too many applications, too much customization and too many underutilized servers.”
- “We struggle to meet service level agreements and fast response times for critical workloads.”
- “We can’t implement new projects fast enough.”
- “We’re asked to reduce headcount every year but the work never seems to go away.”

“Our IT environment is too expensive to manage and maintain.”



Definition of Virtualization

An approach to IT that pools and shares resources
so utilization is optimized and
supply automatically meets demand

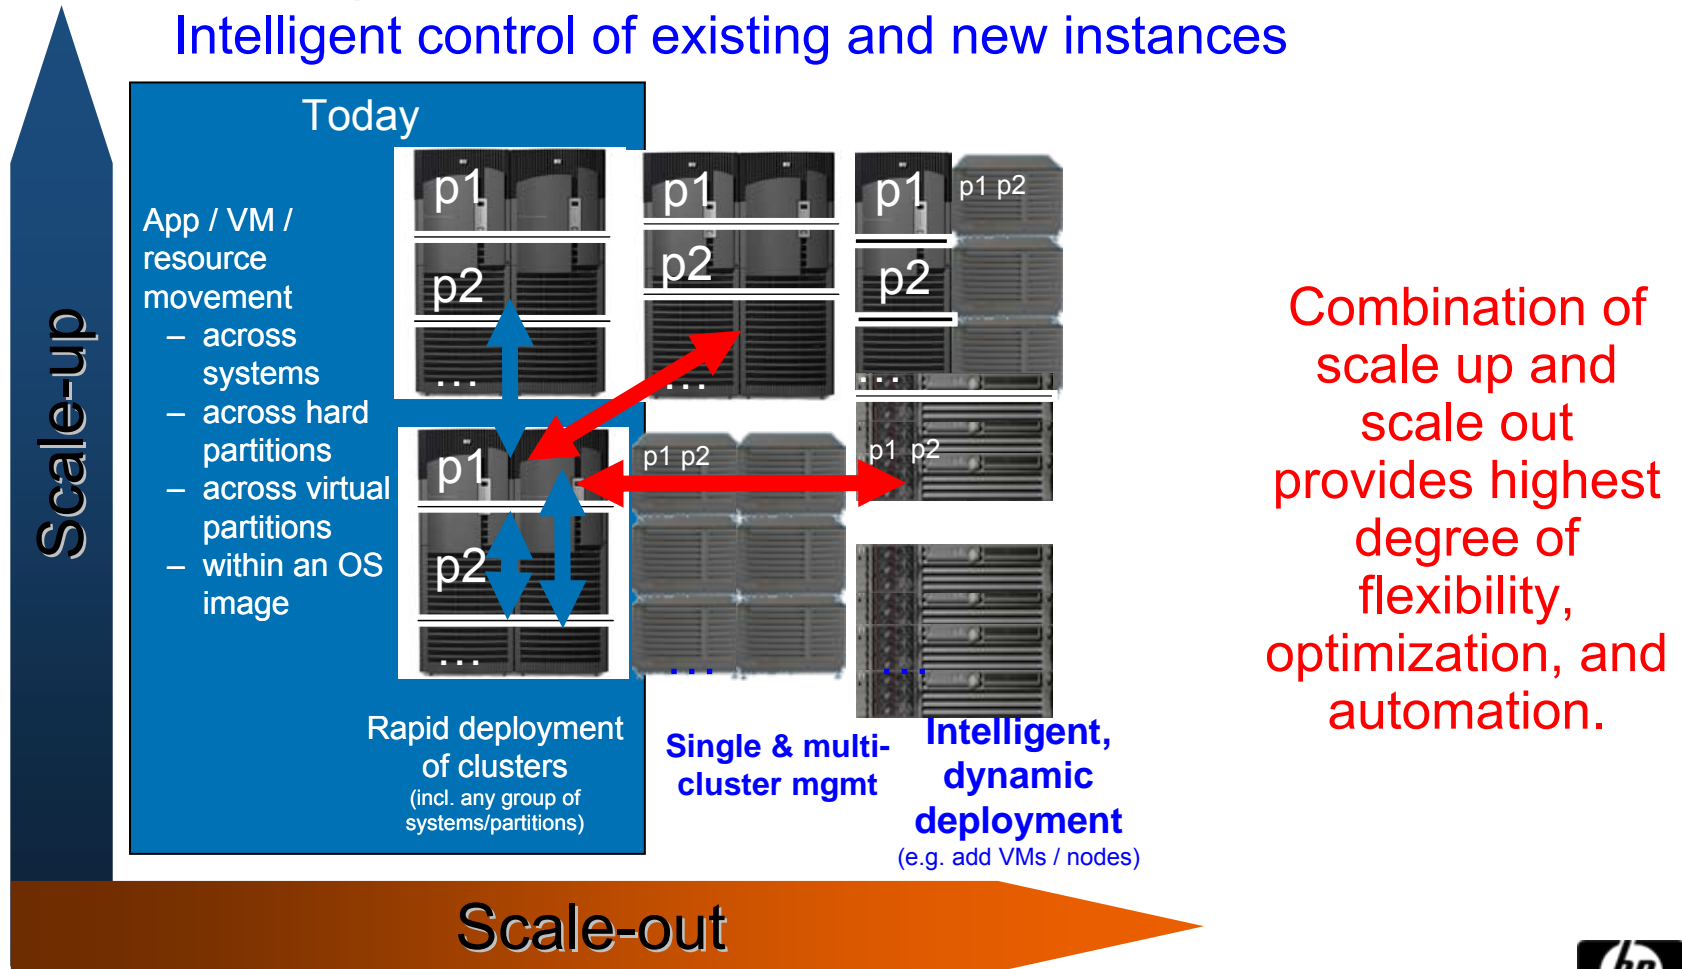


HP Virtual Server Environment

Resource optimization in real time for scale up and scale out on the road to the complete IT utility

Scale up and scale out:

Intelligent control of existing and new instances



Partitioning Continuum

Single Physical Node

Single OS image per node within a cluster

nPartitions

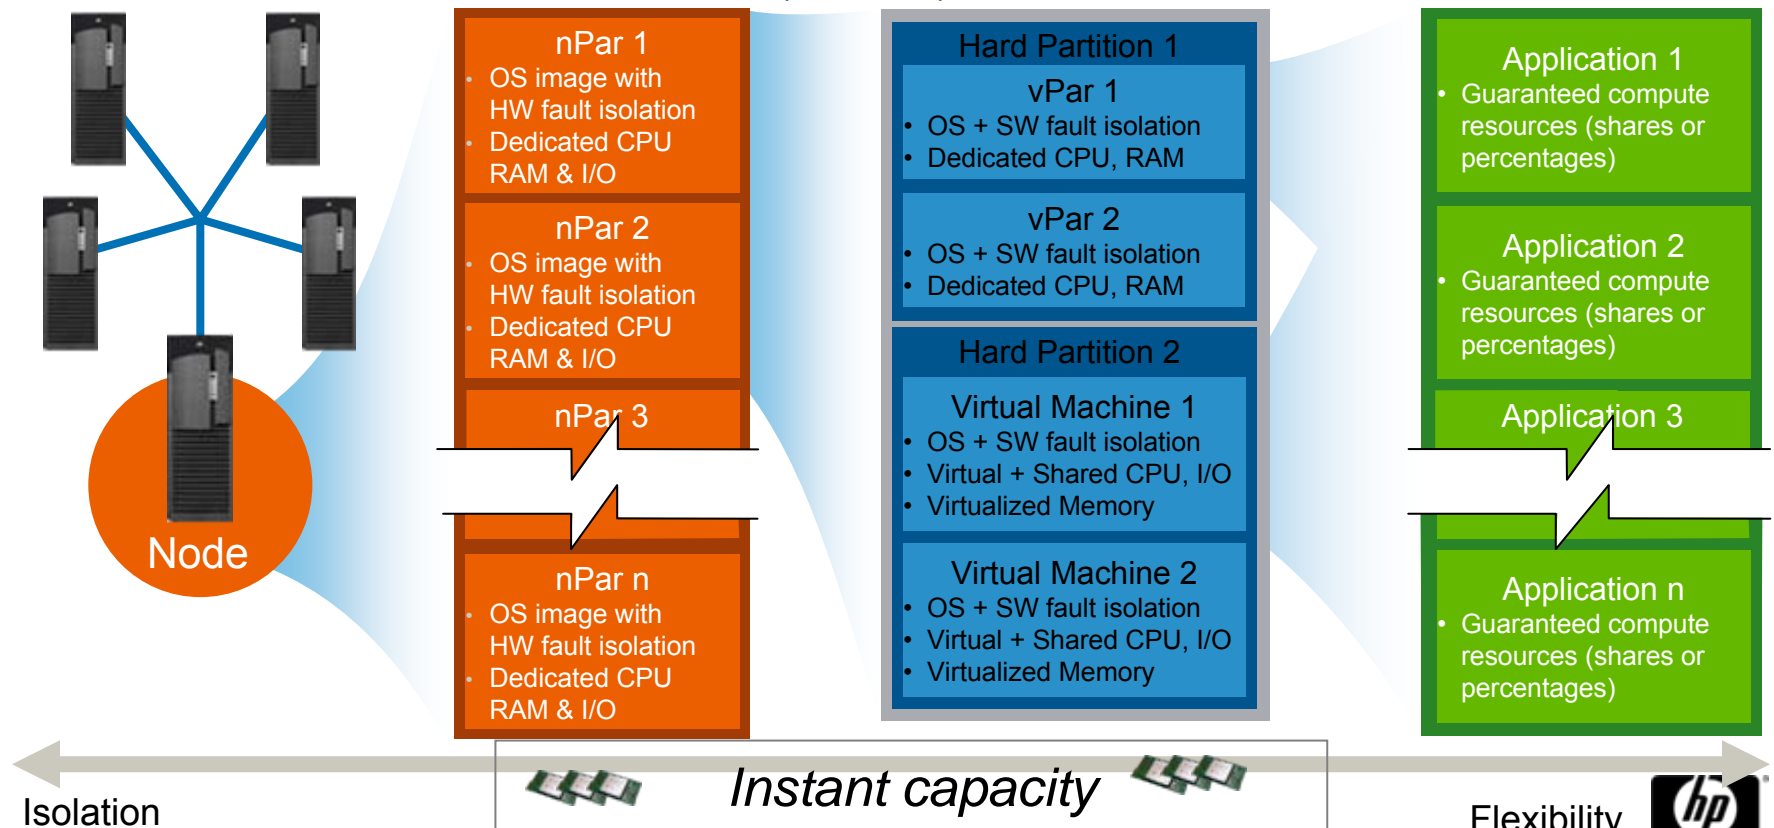
Hard partitions within a node

Virtual Partitions & HP Integrity Virtual Machines

Within a hard partition (or server)

Secure Resource Partitions

-- Partitions within a single OS image with security containment





hp

Implementation Case

Business Case Scenario – Resource Flexing

Temporary spike in application demand results in poor service response times.... poor response means lost revenue!

- “We struggle to meet service level agreements and fast response times for critical workloads...”
- Capability sets needed for efficient response:
 - Integrated Oracle and HP Management framework
 - HP Virtual Server Environment
 - Global Workload Manager and Partitioning Continuum

A photograph of a flight information display board. The board shows flight status for various destinations, including Buenos Aires, Amsterdam, London, and Madrid. The board is dark with white text, and the word 'DESTINOS / DEPART' is visible at the top.

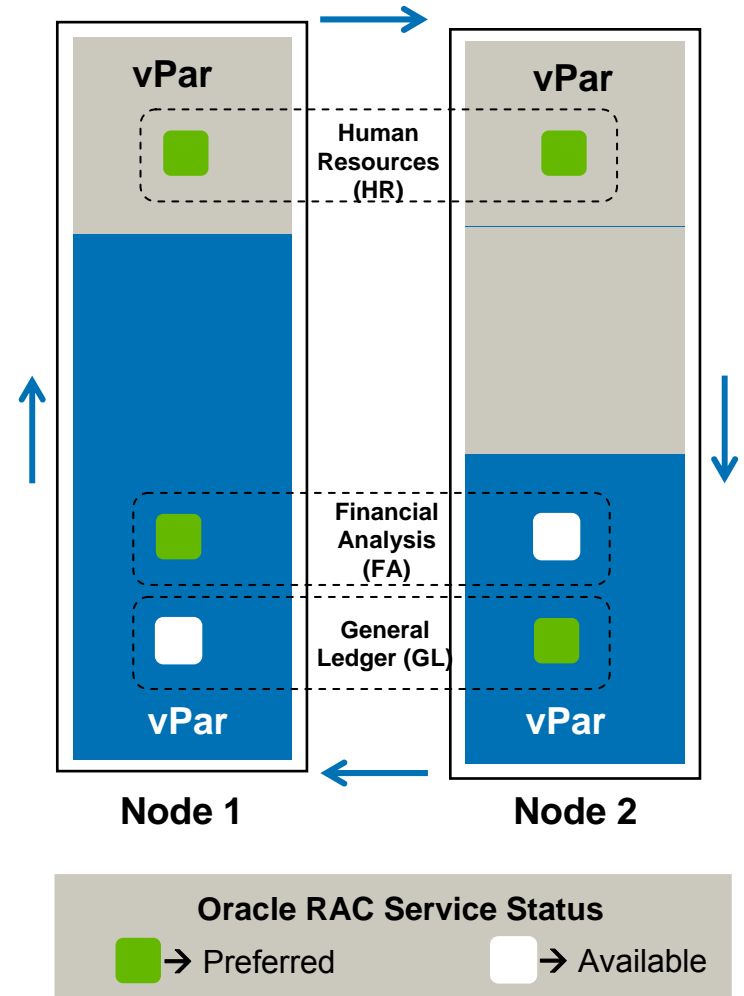
Dynamic Resource Flexing with HP VSE and Oracle RAC

Database Load

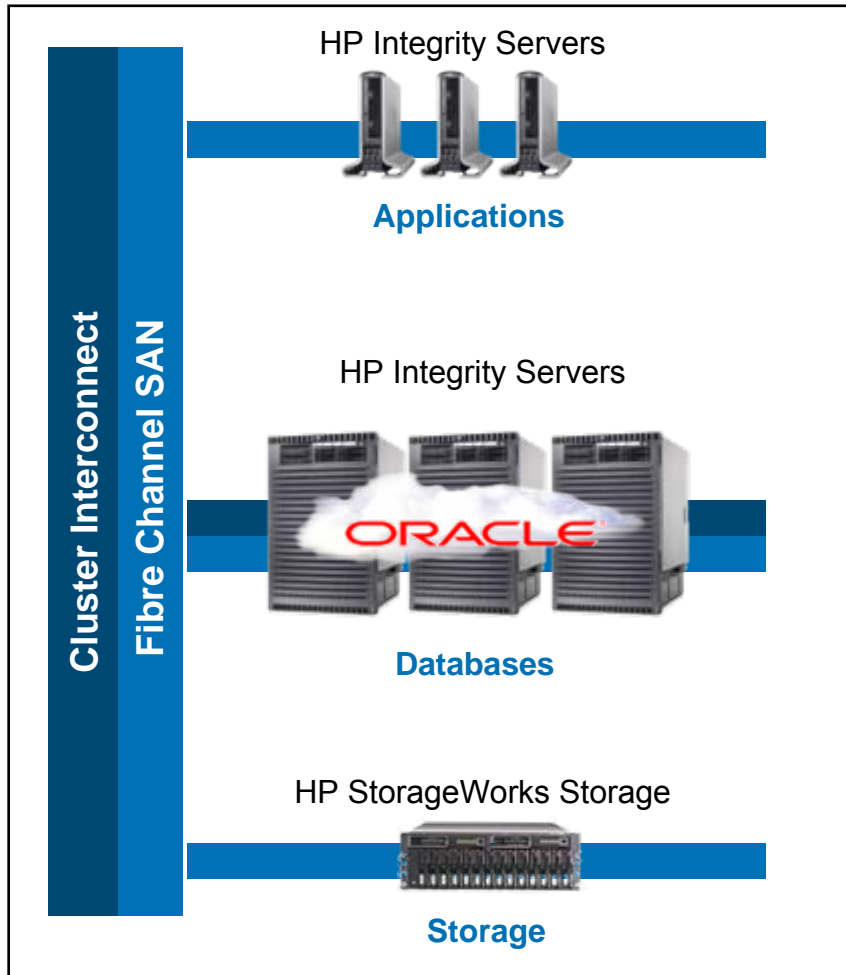
- Increase in database workload as sales reporting begins
- Continued increase in database workload as sales reporting continues
- Sales reporting activity subsides, database workload returns to normal levels

VSE Response

- Sense service level breached and **scales up** – moves CPUs from HR to FA on Node 1
- Sense capacity exhausted on Node 1 and **scales out** – activates FA on Node 2
- Sense reduced demand and **scales back** – deactivates FA on Node 2 and moves CPUs back from FA to HR on Node 1



The HP VSE Reference Architecture for Oracle RAC on HP-UX 11i



- Application Components
 - Oracle 10g RAC
 - Oracle 10g Application Server
 - Oracle Enterprise Manager 10g Database Control
- Availability
 - HP Serviceguard Cluster File System for RAC
- Partitioning
 - HP Virtual Partitions - database
 - HP Resource Partitions - apps
- Intelligent Control
 - HP Systems Insight Manager
 - HP Integrity Essentials Virtualization Manager



HP Adaptive Infrastructure
components

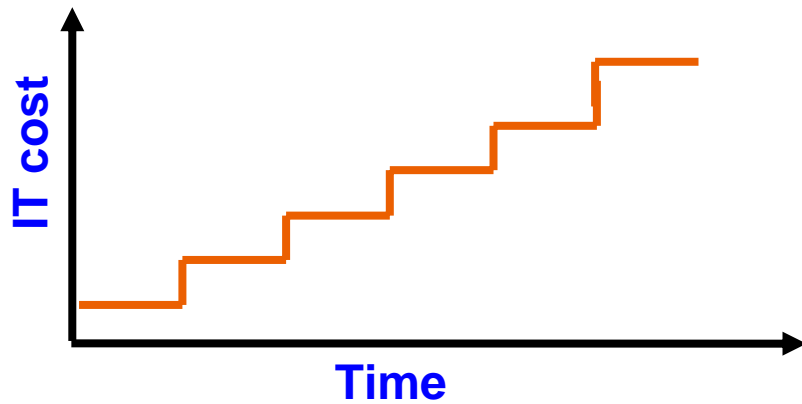
What are the virtualization challenges for your Oracle environment



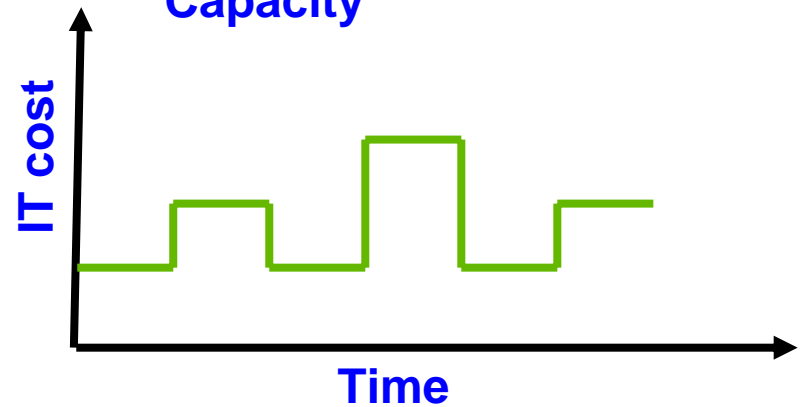
- Manage your CPU capacity as a pool of resources
- Define the right granularity for your Oracle Compartment
- Automatically meet with your SLO
- Provide High Availability

HP offers different utility pricing options for different capacity needs

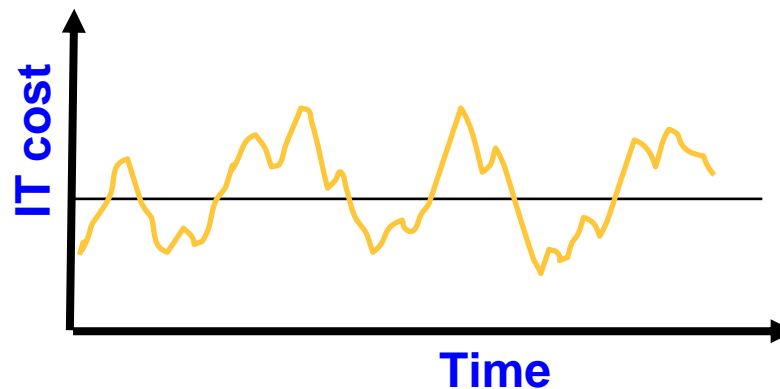
Instant Capacity



Temporary Instant Capacity

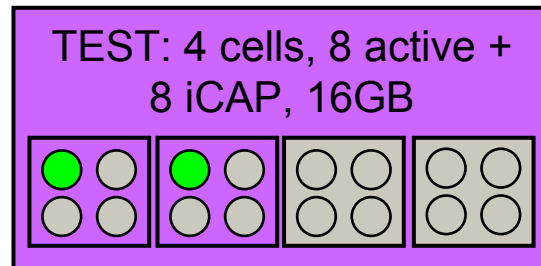
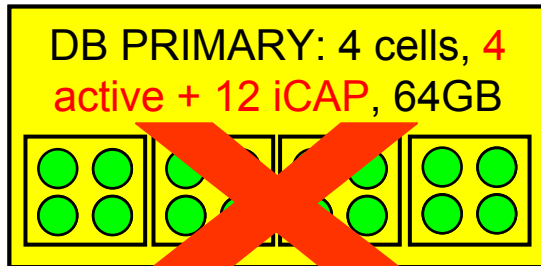


Pay per use

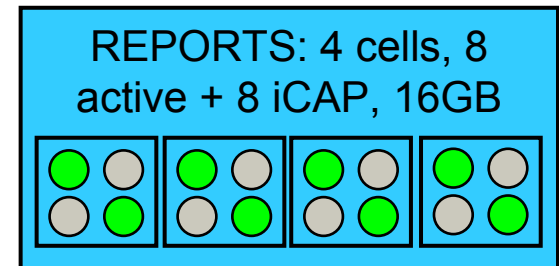
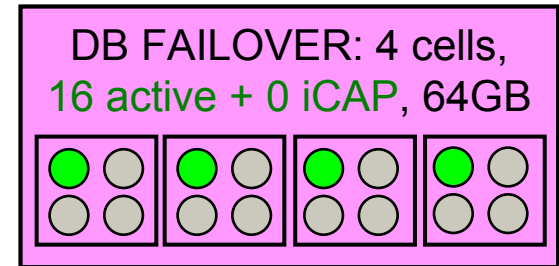


(G/T)iCAP :Exchanging RTUs at no cost

GiCAP Group Manager



London



Reading

Instant Capacity

- Instant Capacity (iCAP)
 - CPUs can be activated on-line – no reboot required
 - Activation of entire Cell Boards requires reboot of partition that needs to be enlarged (work without reboot in HP-UX V3)
 - CPUs / Cell Boards are paid for when they are activated
 - Price paid is current price when activated
 - CPUs can be deactivated in one nPar/vPAR and activated in another
- Temporary Capacity Instant (TiCAP)
 - Capacity is purchased in 30 Day increments
 - Activating iCAP CPU`s deduct minutes from the “bank”
 - Deactivating the iCAP CPU`s stops the deductions
 - Any number of iCAP CPUs can be activated
 - Excellent solution for:
 - Short term peaks in application load
 - Lower cost failover server
- Global iCAP (GiCAP) pushes the envelope across servers and even across sites
 - Excellent for failover / disaster recovery iCAP pools: in case of an outage, the number of defective CPUs can be activated on any server out of the pool
 - Shift resources between production and non-production, even if they are on different servers or datacenters

What are the virtualization challenges for your Oracle environment

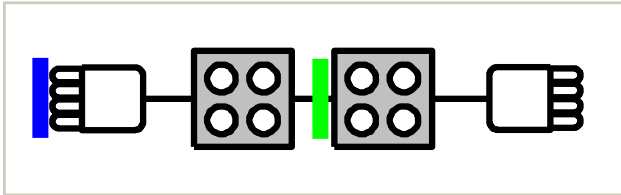
- Manage your CPU capacity as a pool of resources
- Define the right granularity for your Oracle Compartment
- Automatically meet with your SLO
- Provide High Availability

Hard partitions aka nPars

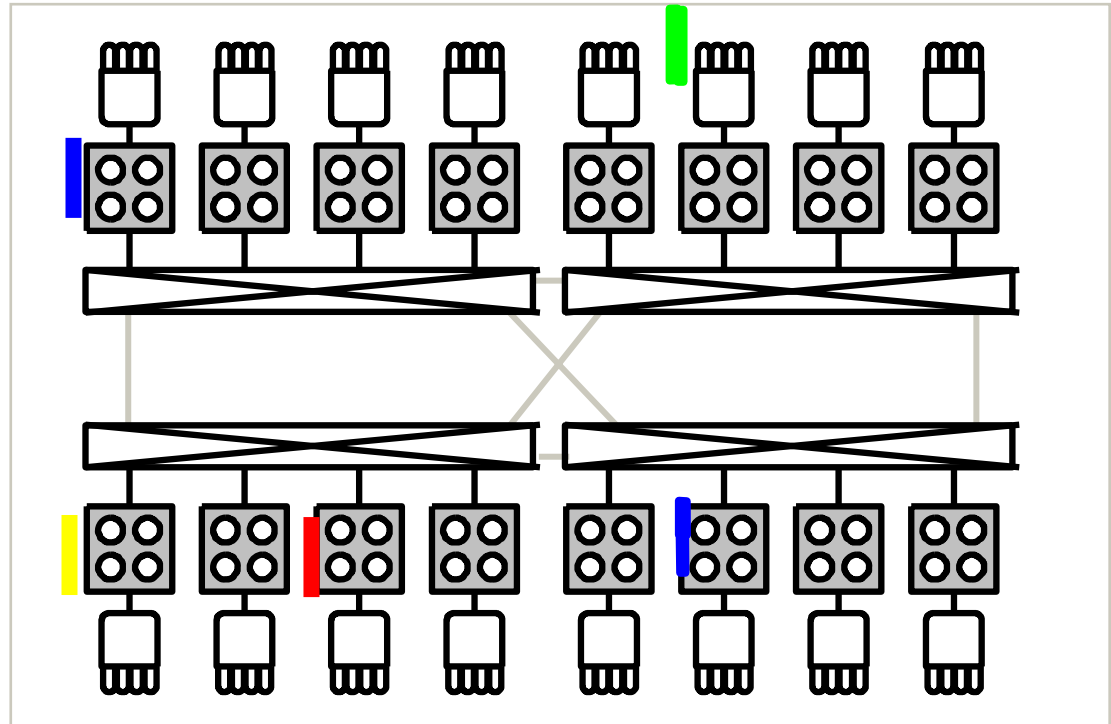
Full electrical isolation between partitions

New with HPUX V3 : full cell OLAD

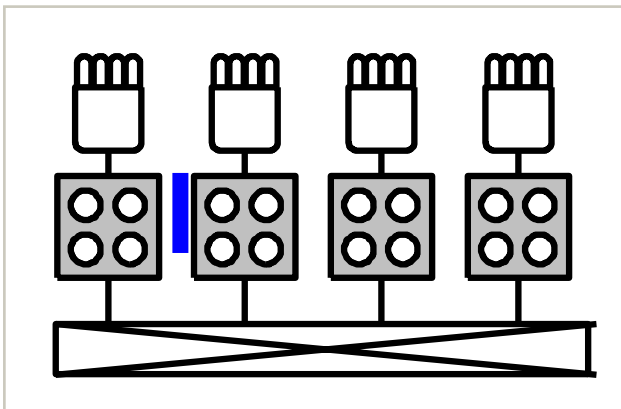
8 Socket system



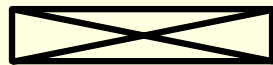
64 Socket system



16 Socket system



Legend



Two crossbar switches



4 Socket Cell



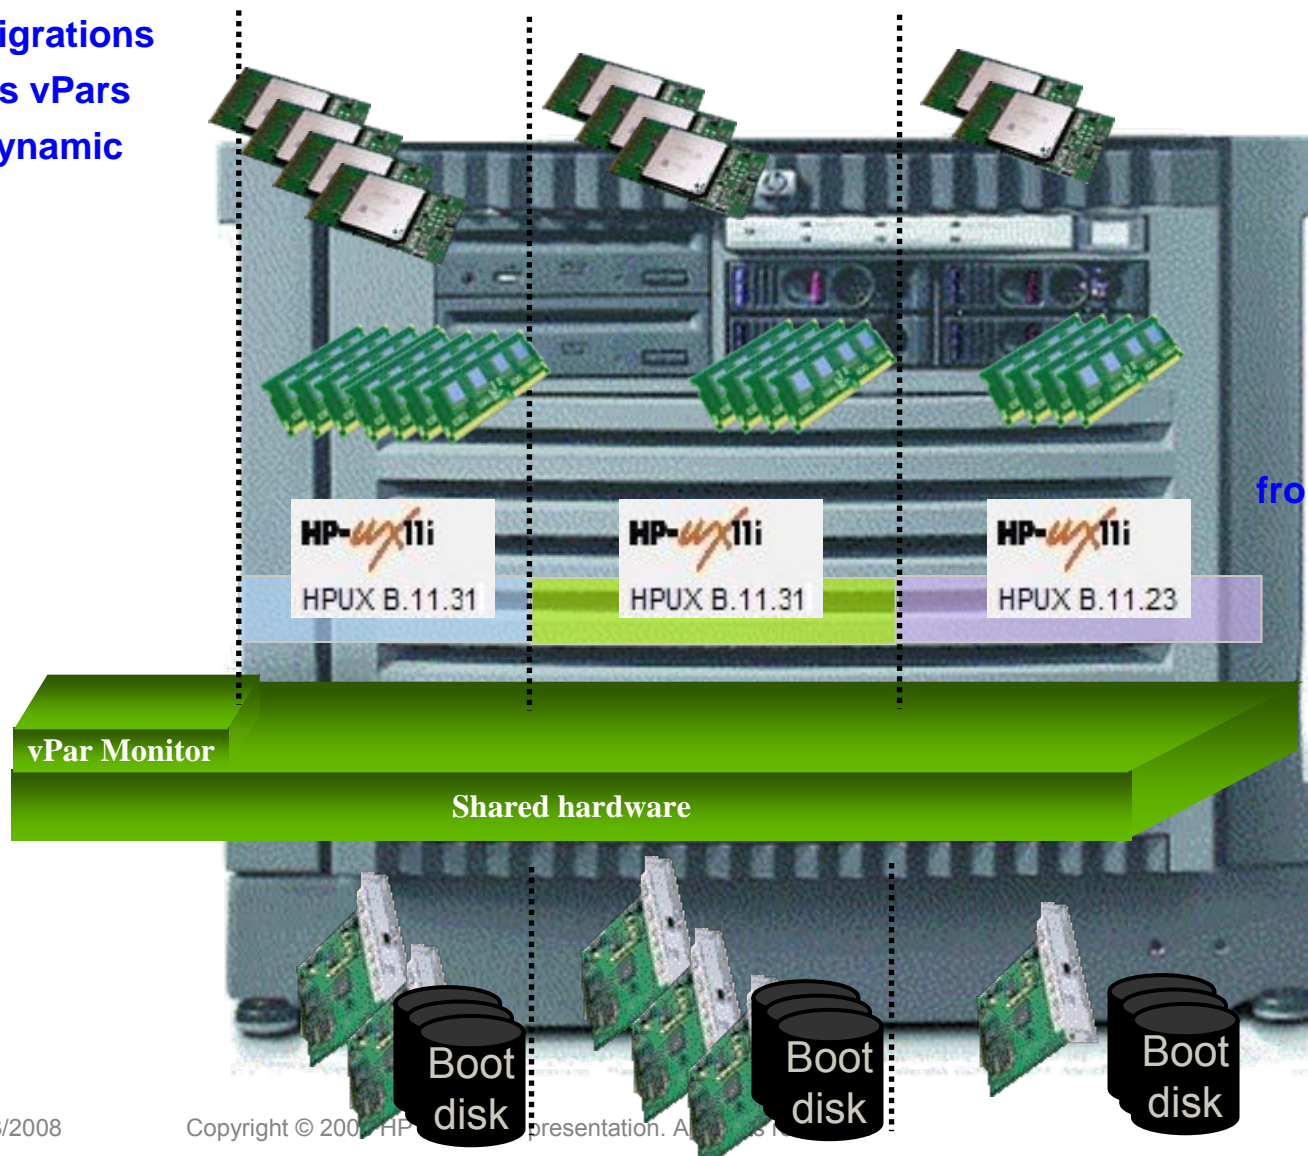
I/O Backplane

Sweet Spots for nPars

1. Create an nPar for each Mission Critical Application to provide complete hw fault and OS isolation
2. Each nPar can run a unique instance of an operating system (HP-UX, Linux, Windows, OpenVMS)
3. nPars are supported on both PA-RISC and Integrity
4. nPars can be either PA-RISC or Integrity on Superdomes (excellent for mixed environments or during transitions)
5. nPars can be further sub-partitioned with vPars, Integrity Virtual Machines, or Secure Resource Partitions

vPar V5 : Illusion of separate hardware platform and dynamic CPU and Memory movement

CPU migrations across vPars are dynamic



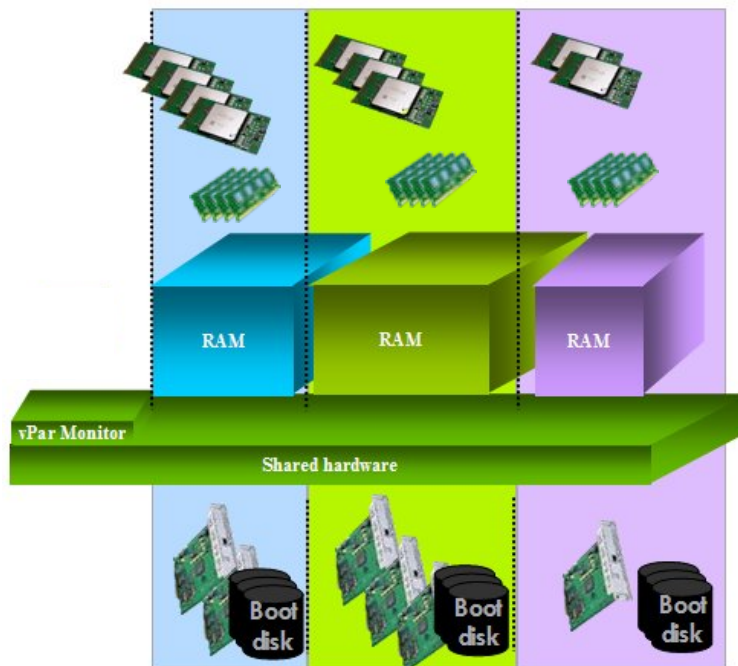
Memory can move from 11iV3 vPar to 11iV vPar

Each vPar owns HBAs and at least a boot disk

HP-UX Virtual Partitions aka vPars

Dynamic granularity is one CPU

Multiple HP-UX instances
running on the same system
or in the same nPar



Increased system utilization

- partitioning a single physical server or hard partition into multiple virtual partitions

Increased Flexibility

- multiple independent instances of HP-UX
- dynamic CPU migration across virtual partitions
- dynamic memory movement between HPUX V3 virtual partitions

Increased Isolation

- application isolation across virtual partitions
- OS isolation, HPUX V2 and HPUX V3
- individual reconfiguration and reboot

Sweet spot

Linked to a SAN with snapshot/snapclone facility

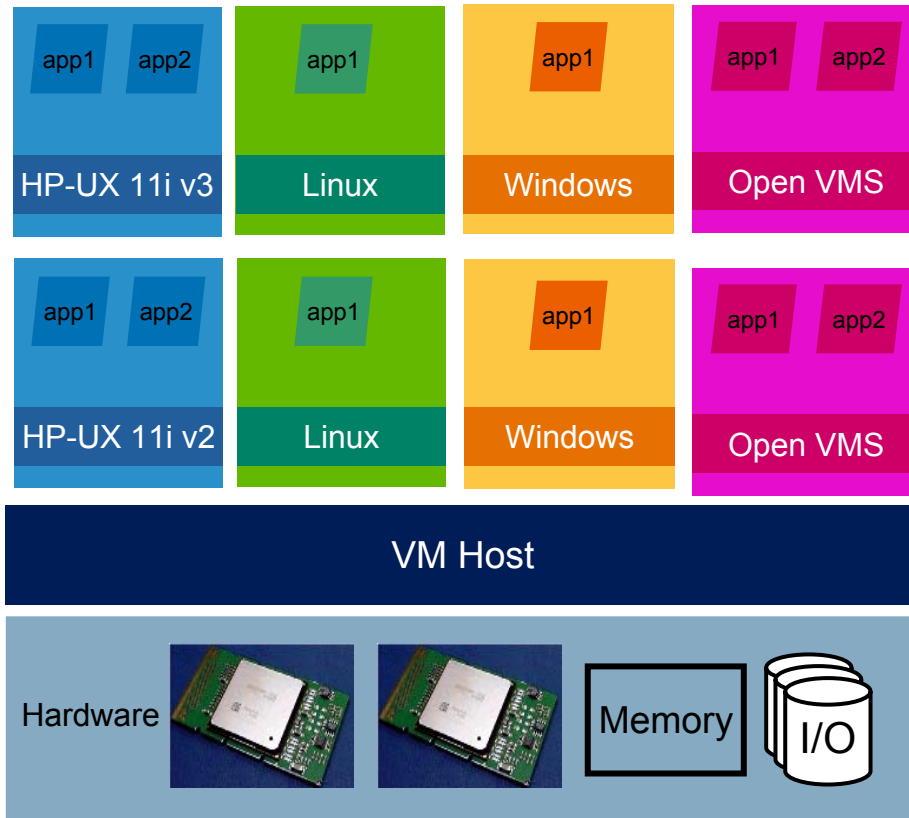
- One or two production virtual partitions
- One development vPar
- One test vPar

Sweet Spots for vPars

1. If you want to share CPU resources between workloads in different partitions (ex : 1 Production vPar, 2 dev & test vPars)
 - a. Allows active CPUs and memory to be dynamically reallocated
 - b. iCAP CPUs can be activated (either permanently or with TiCAP) in any vPar within the same nPar
2. If you require finer granularity than nPars (i.e. an entire cell not needed for the workload), and...
 - a. Need OS isolation or different versions of the OS
 - b. Have I/O intensive applications

HP Integrity Virtual Machines (VM)

Optimum utilization across multiple OS



HP-UX 11i software can be licensed by virtual machine!

- Sub CPU virtual machines with shared I/O
- Runs on a server or within an nPar
- Dynamic resource allocation built in
- Resource guarantees as low as 5% CPU granularity
- OS fault and security isolation
- Supports all (current and future) HP Integrity servers
- Designed for multi OS
 - HP-UX 11i v2 & V3 guests
 - Linux guest support
 - Windows guest support
 - OpenVMS guests in future
- Integrated with VSE

How Oracle can benefit of Integrity VM

1. Easy, efficient means of HW consolidation
 - HW sharing increases utilization
 - OS isolation eliminates effort and complexity of combining workloads on same application stack
 - Fast system deployment and provisioning
2. Applications that don't normally need dedicated hardware (or a whole CPU) but do need OS isolation, different OS versions, different Operating Systems, and/or different application stack versions
3. Good choice for smaller non-cell based systems
4. Features for Oracle
 - Sub cpu partitioning
 - Workload management
 - HPVM entitlement
 - Integration with gWLM for SLO purpose
 - Dynamic memory allocation

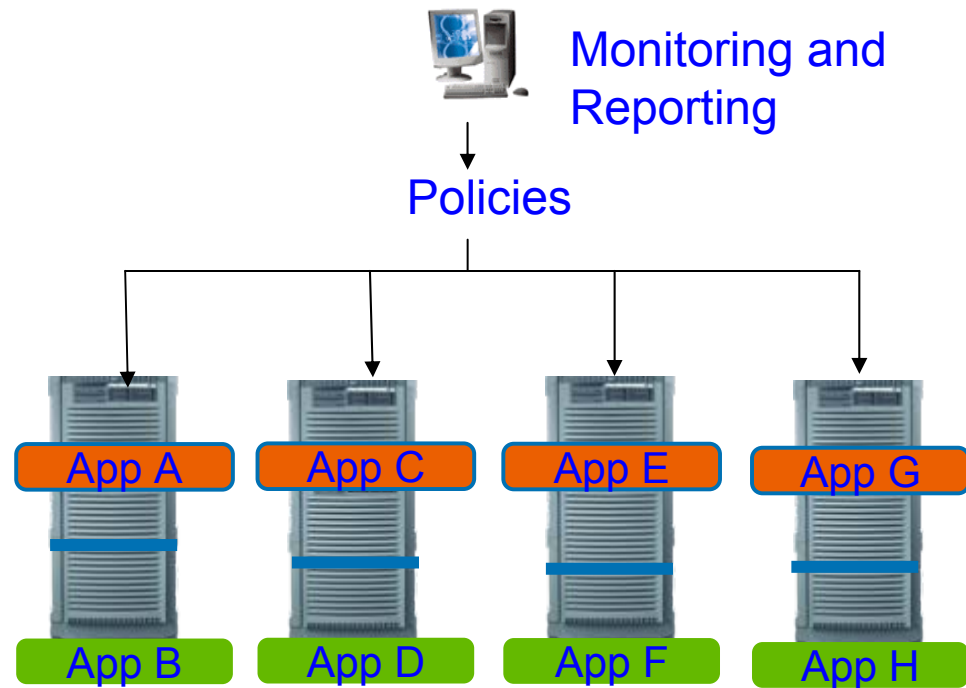
What are the virtualization challenges for your Oracle environment

- Manage your CPU capacity as a pool of resources
- Define the right granularity for your Oracle Compartment
- Automatically meet with your SLO
- Provide High Availability

HP Global Workload Manager to manage large, multi-system VSE

Dynamic, multi-system workload management
for HP-UX 11i and Linux

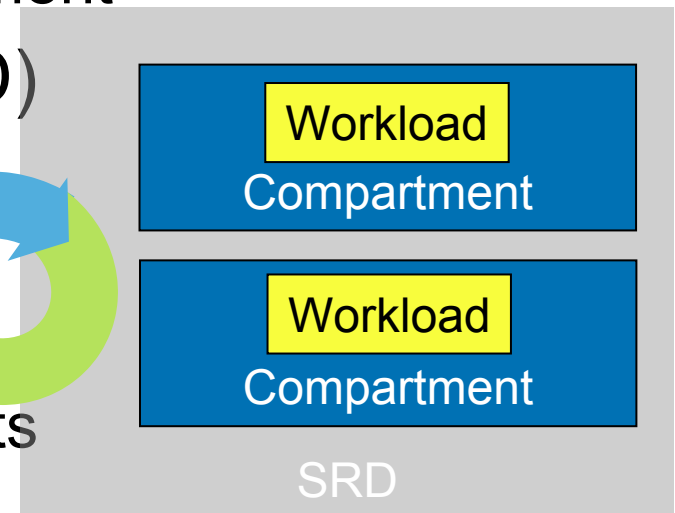
- Goal-based policy engine for managing workloads across multiple systems simultaneously
- Easy to use management integrated with HP Systems Insight Manager
- Enables central IT to deliver an IT utility supporting multiple LOBs
Resources can be assigned to LOB based on:
 - Own/borrow/lend model
 - Fixed entitlement model
 - CPU utilization model
 - Service Level Objectives
- Based on HP-UX Workload Manager, the only goal-based policy engine for UNIX



Dynamic resource and workload
reallocation

gWLM terms

- Workload
 - Collection of processes that are a set of executables
- Compartment
 - It is a system division as a nPar, a vPar, a HP VM guest, a PSET or a FSS group
 - The workload runs in the compartment
- Shared Resource Domain (SRD)
 - A collection of compartments that share physical resources
- Policy
 - A collection of settings that instructs gWLM how to manage workload resources within the SRD



HP Global Workload Manager

Centralized Workload Management

HP Systems Insight Manager - Microsoft Internet Explorer

Address: <https://gaudix.etc.vbe.cpqcorp.net:50000/mxportal/home/MxPortalFrames.jsp>

User: root
[Home](#) | [Sign Out](#)

VSE Management: Global Workload Manager

Tools | Deploy | Configure | Diagnose | Optimize | Reports | Tasks & Logs | Options | Help

System | Workload | Shared Resource Domain | Capacity Advisor

Tools | Create | Modify | Delete | Policy | Report

Show: Deployed SRDs Undeployed SRDs [Refresh Data](#)

<input type="checkbox"/>	SRD Name	Policy Status	Mode	TICAP	State	Size	CPU Utilization	Last Update
<input type="checkbox"/>	titane.srd	✓	Managed		Deployed	4	72.294%	Mar 01, 2007 15:28:30

Details for SRD: titane.srd

Shared Resource Domain: [titane.srd \(Modify SRD\)](#)

Mode: Managed

State: [Deployed \(Undeploy SRD\)](#)

<input type="checkbox"/>	Workload	CPU Utilization	Type	Policy	Status	Request	Allocation	Size	Hostname
<input type="checkbox"/>	hpux.2	73.959%	hpvm	own_5_3_P150	✓	2.28	2.31	2.31	titane.etc.vbe.cpqcorp.net
<input type="checkbox"/>	linux.2	0%	hpvm	Owns 0.1 CPU-Max 1 CPU	⚡	0	0	0	titane.etc.vbe.cpqcorp.net
<input type="checkbox"/>	oracle.2	33.376%	hpvm	Owns 0.1 CPU-Max 1 CPU	✓	0.07	0.14	0.14	titane.etc.vbe.cpqcorp.net
<input type="checkbox"/>	titane.OTHER	11.846%	hpvm	Owns 0.1 CPU-Max 1 CPU	✓	0.05	0.15	0.15	titane.etc.vbe.cpqcorp.net
<input type="checkbox"/>	windows.2	79.914%	hpvm	zCustom tivmwin_va	✓	1.08	1.4	1.4	titane.etc.vbe.cpqcorp.net
						Totals:	Req = 3.48	Alloc = 4	Size = 4

Unknown - Data is not yet available. Data should appear on the next screen refresh. Causes could be a recent configuration change that has not yet completed, or a transient network issue. You may use the [Refresh Data](#) button for an immediate refresh.

System Overview

- All Systems
- All Events
- Systems
 - Private
 - All INO Complex
 - All TITANE
 - Shared
 - Systems by Type
 - All Systems
 - All Servers
 - All VSE Resources
 - All nPartition Servers
 - All HP Integrity Virtual M
 - All Virtual Partition Serv
 - All Resource Partitions
 - All Shared Resource Do
 - All HP Serviceguard Clu
 - All Standalone Servers

System Status
Legend... Customize...
Updated: Thu, 3/1/2007, 3:26:02 PM CET
9 41 1 50 Uncleared Event Status

Search
Search
Advanced Search...

Systems and Events
System Overview
All Systems
All Events

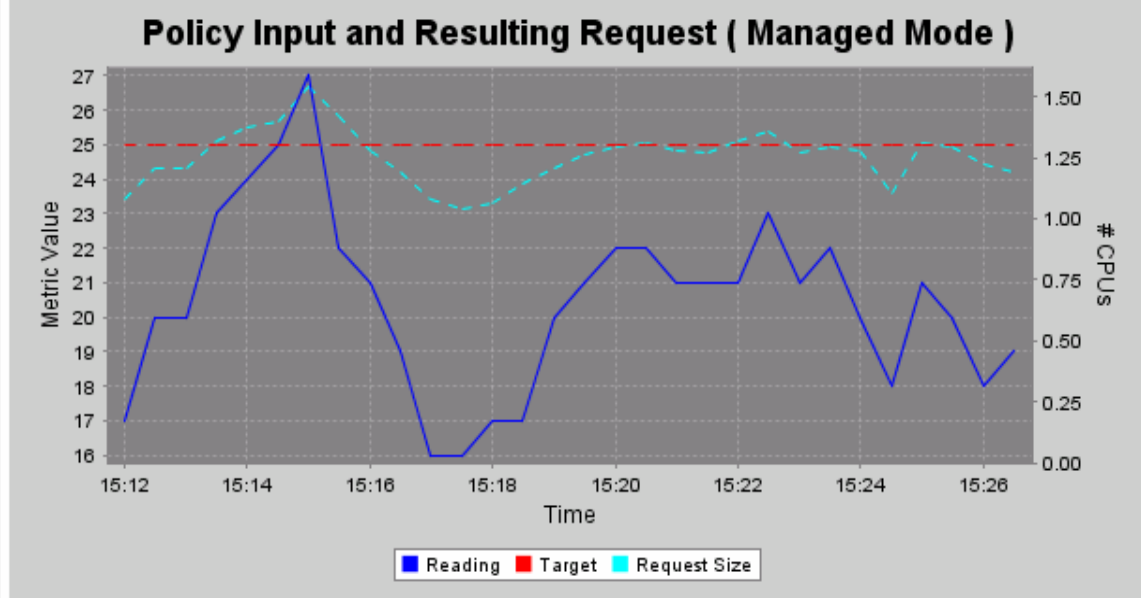
- Systems
 - Private
 - All INO Complex
 - All TITANE
 - Shared
 - Systems by Type
 - All Systems
 - All Servers
 - All VSE Resources
 - All nPartition Servers
 - All HP Integrity Virtual M
 - All Virtual Partition Ser
 - All Resource Partitions
 - All Shared Resource Do
 - All HP Serviceguard Clu
 - All Standalone Servers

HP gWLM: View Real-time Reports

Generate and view gWLM real-time workload and policy reports
 Request Size Allocation Size Owned Effective PolicyMin Effective PolicyMax

Refresh Graph Graph range: 15 mins Real-time cache: 20 mins

Workload windows.2 Affected by Policy zCustom_tivmwin_ya



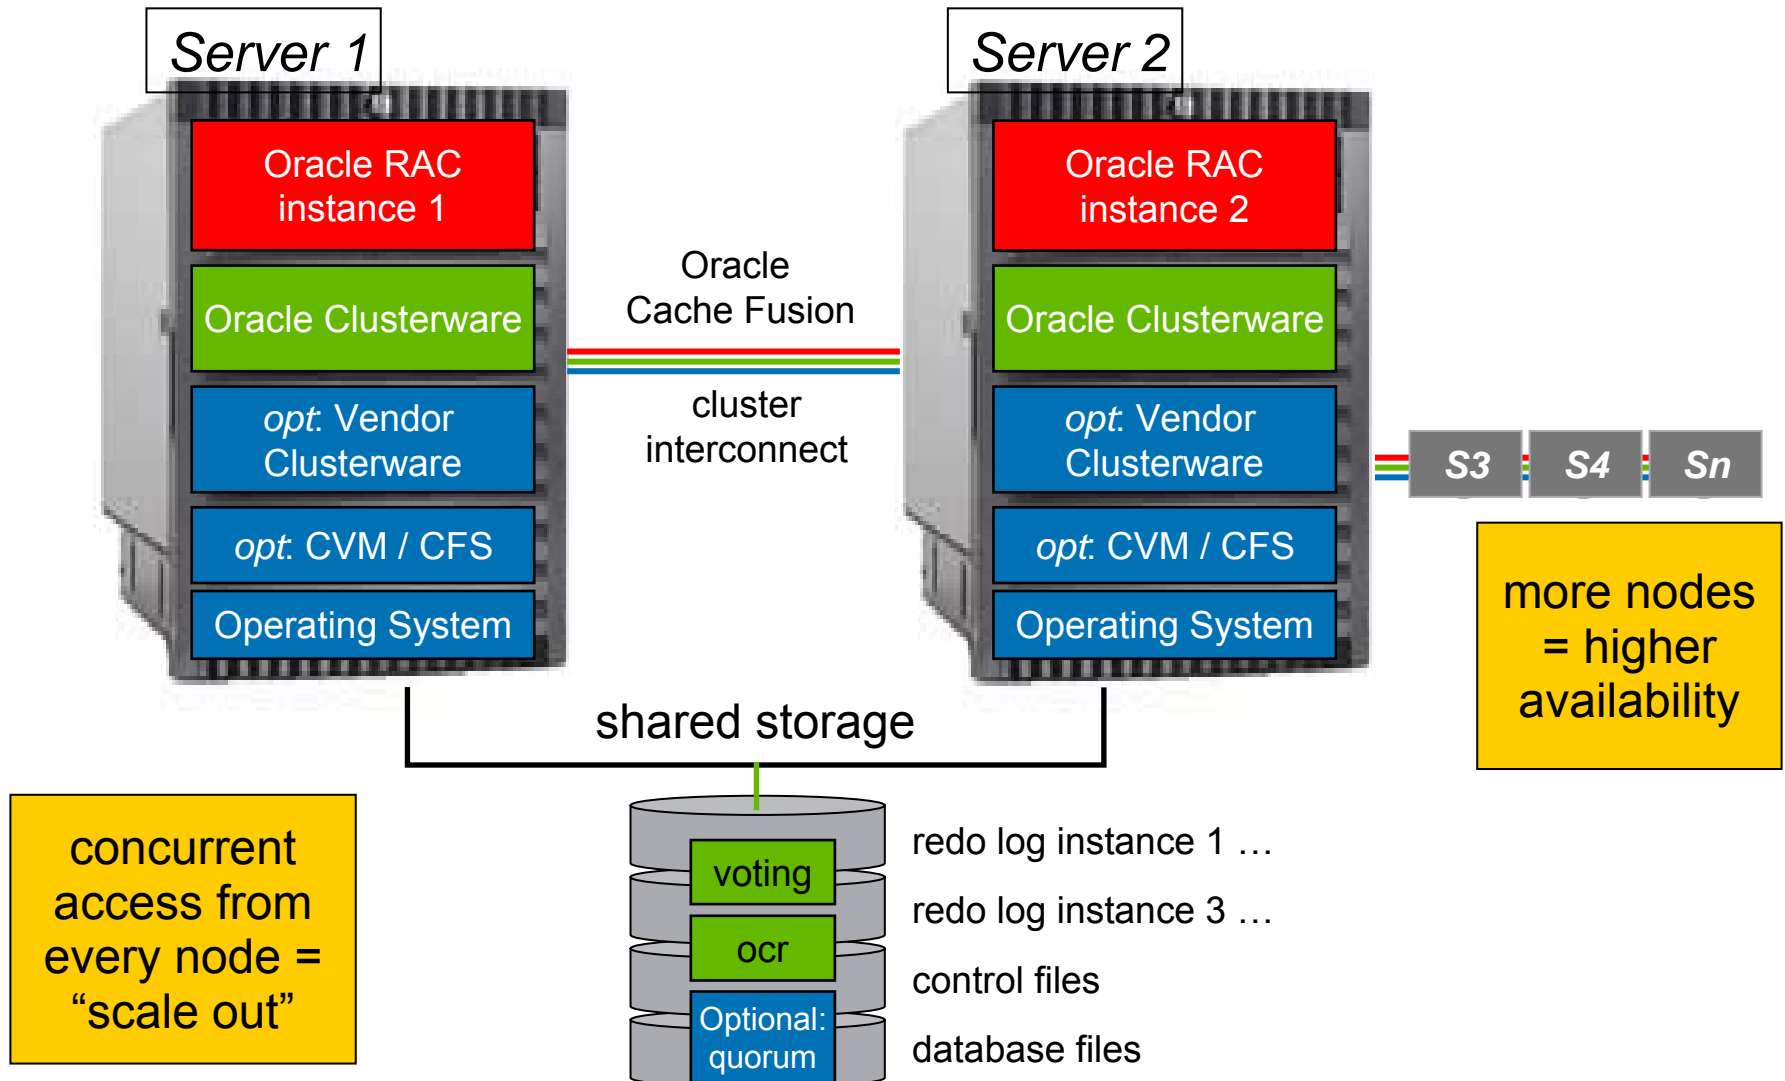
Reading Target Request Size (right axis)
Refresh Graph Graph range: 15 mins Real-time cache: 20 mins

What are the virtualization challenges for your Oracle environment



- Manage your CPU capacity as a pool of resources
- Define the right granularity for your Oracle Compartment
- Automatically meet with your SLO
- Provide High Availability

Oracle RAC Architecture



Conclusion

1. HP VSE automatically allocate the right amount of resources based on SLO requirements
2. Various level of granularity. From sub-cpu to cell.
3. Multi-OS support (Linux, HP-UX, Windows, Open VMS)
4. HP VSE is fully integrated with Oracle RAC
5. Disaster tolerant solution

Q & A

