HP-UX Workload Manager (WLM) and ISV Toolkits

Isom Crawford
Hewlett-Packard
3000 Waterview Parkway
M/S D1W43
Phone: (972) 497-4740
Fax: (972) 497-3123
isom_crawford@hp.com

Target Problem

- Handling Peak Demand for Critical Applications

Traditional Approach

Overprovisioning

- Lots of dedicated Unix servers
- Excess capacity on each

Drawbacks

- Cost of underutilized capacity
- Difficult to manage many systems

New Solutions

Application Consolidation

- **Run multiple workloads on a Unix box**

Spare Capacity Consolidation

- **Provide spare capacity for multiple apps on the same system or systems**

Capacity on Demand

- **iCOD**
- **vPar resource balancing**

HP Workbad Management Vision

Virtual Computing Resources

- Systems may run multiple apps
- Some apps span multiple systems
- Resources allocated as needed.

This Requires

- Vertical Scalability
- Horizontal Scalability
- New partitioning tools
- Resource management tools

Getting to the Vision

What we have today:

- **Process Resource Manager (PRM)**
- **HP-UX Workload Manager (WLM)**
- **Capacity on Demand**

Short to Medium Term Future:

- **Hardware Partitions**
- **Virtual Partitions**
- **Utility Computing**
- **Multi-system Workload Management**

HP-UX Process Resource Manager (PRM)

The Problem: Competition for resources

The Solution: Resource Partitioning with PRM

Process Resource Manager (PRM)

PRM is a partitioning tool

- Administrator defines:
  - Resource groups
  - Policies for putting processes into groups
  - Resource allocations for each group
- PRM controls:
  - CPU
  - Real memory
  - Disk bandwidth

HP-UX PRM

Key Features and Benefits of PRM

- Controls allocation of CPU, real-memory and disk I/O bandwidth based on user-specified policies.
- Applications do not require modification to work with PRM.
- PRM configuration can be changed at any time - even under load.
- Supports resource policies based on users and applications.

Enables running multiple, mission critical applications on a single system.

Key concept

Static resource partitions may not always provide the most efficient use of resources for all workloads.

HP-UX WLM gives you the flexibility you need

HP-UX Workload Manager (WLM)

HP-UX WLM is a dynamic front-end to PRM

It automatically adapts the PRM configuration based on what's happening on the server

WLM helps you fully achieve the benefits of consolidation

HP-UX WLM (cont.)

HP-UX WLM delivers the compute power when and where it's needed

- Fixed resource allocation
- Scheduled resources
- Resource on Demand
- Dynamic Resource policies
- Goal-based resource allocation

WLM ensures that priority workloads get the compute resources they need.

Dynamic Resource Policies

Dynamic resource policies use goals, constraints, and conditions.

A resource allocation policy consists of:

- **A workload (PRM group)**
- **Constraints (min, max cpu)**
- **A goal**
- **Priority**
- **Conditions (time of day, event, etc)**

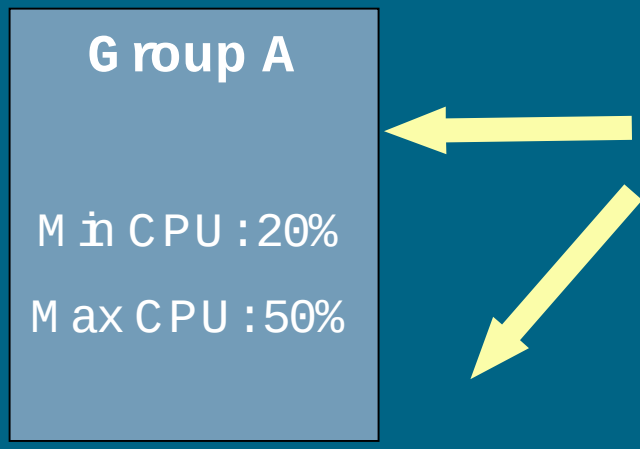

Group A receives 3 shares for each additional user.

Policy applies 9am to 5pm AND when ServiceGuard Package XYZ

Consolidation example with HP-UX WLM

Stretch Goals

Group B

(Funded 30%)

CPU Allocation: 30%

Min CPU: 15%

Max CPU: 60%

- If everyone is busy, give them what they paid for
- Keep a minimum reserved for each group
- If a group is busy, and there are idle cycles, let it "borrow" from another group (up to its maximum)

Policy 1

Goal: If I'm using my share, give me what I paid for

Priority: 1

Min CPU: 15%

Max CPU: 30%

Policy 2

Goal: If I'm really busy, request more CPU

Priority: 2

Min CPU: 30%

Max CPU: 60%

Response-time goals

A Policy (SLO) can contain a response-time goal

Example: Operation "add-to-shopping-cart" must complete in less than 1 second

Data collected via Application Response Measurement (ARM) API used by WLM

HP-UX WLM Process Flow Diagram

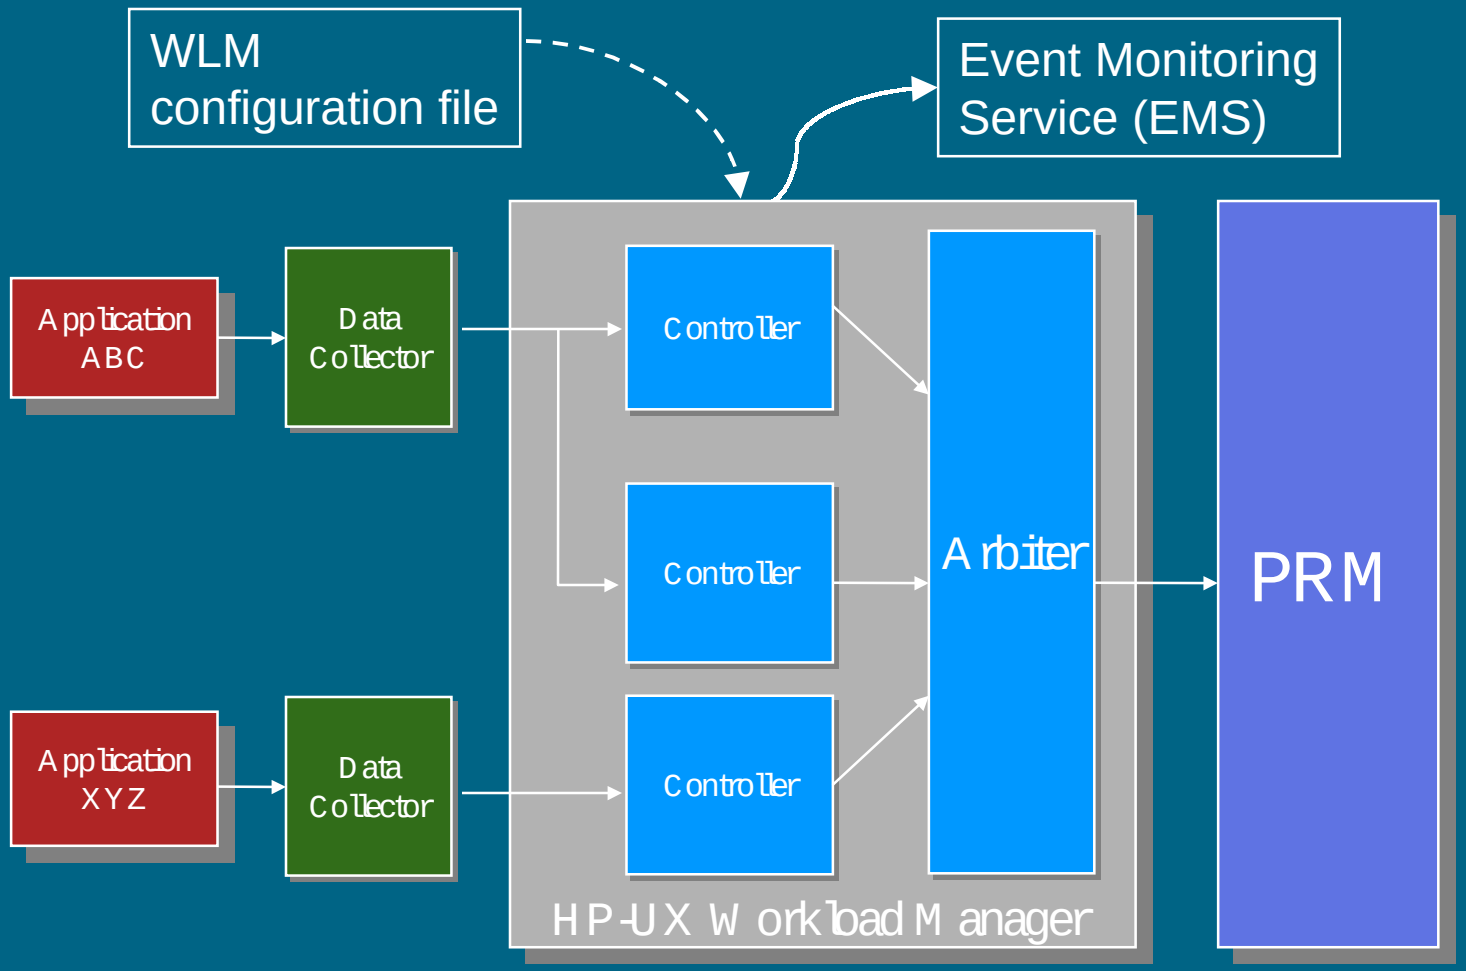

WLM Data Collector Interfaces

- WLM API
- Scripting Interface
- GlancePlus Toolkit
- ARM Toolkit
- ISV Toolkits

ServiceGuard integration

HP-UX Workbad manager is ideal in a ServiceGuard Cluster

Example: If ServiceGuard Package ABC is running on the box, give it 20% of the CPU cycles

WLM Benefits

Maximized utilization of system resources while maintaining performance goals of highest priority applications

IT Cost Containment Through:

- Prioritized sharing of system resources among various applications
- Reduction in the requirement for separate servers for every application
- Reduction in the requirement for dedicated spare systems

WLM Futures:

- Managing Virtual Partitions
- Managing SuperDome Hardware Partitions
- Capacity extension
  - iCOD
  - Utility Computing
- Multi-resource (CPU, Memory, Disk I/O)

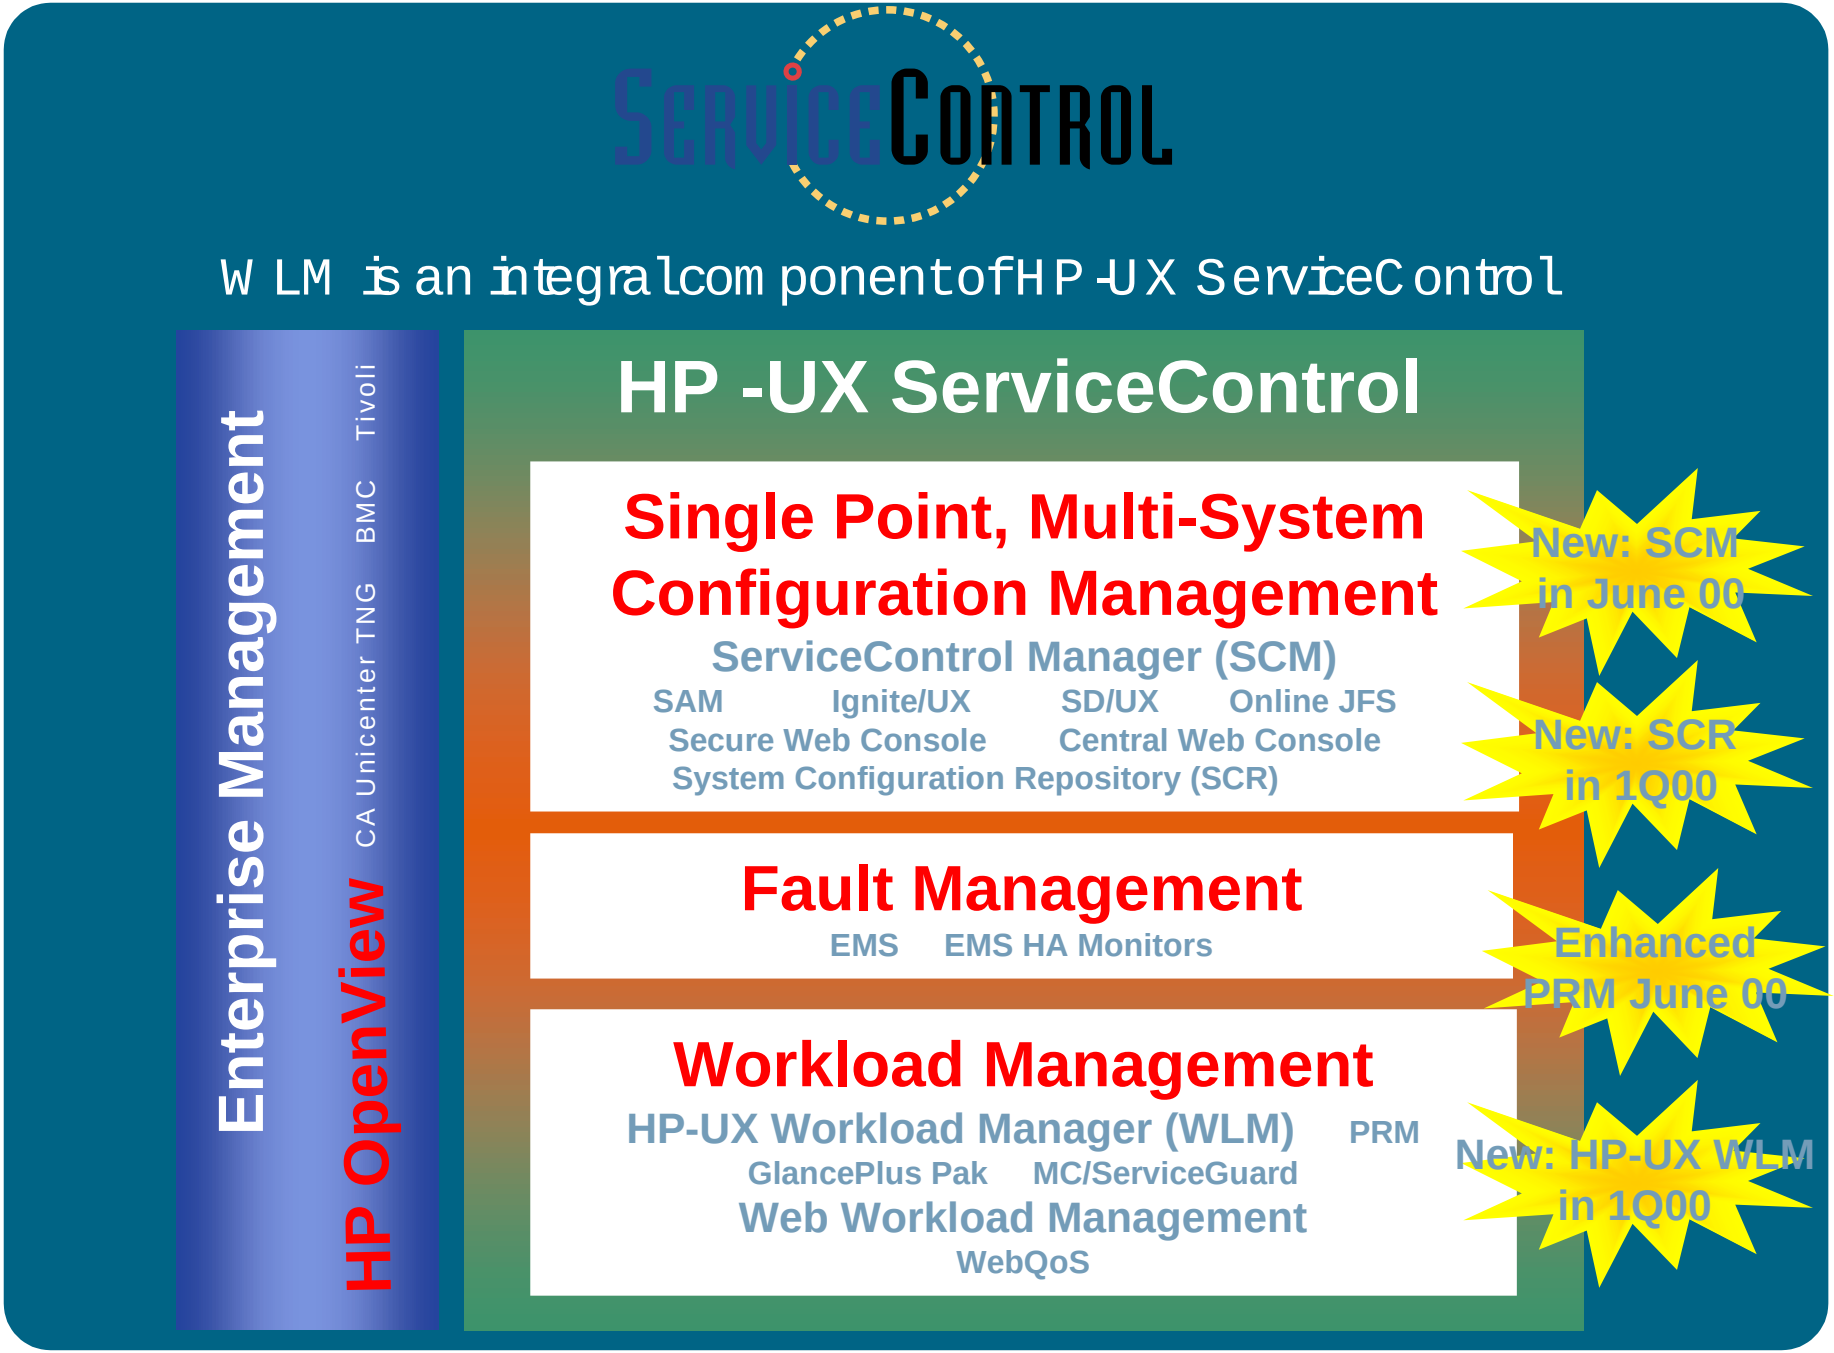
ServiceControl
WLM is an integral component of HP-UX ServiceControl
Enterprise Management
HP OpenView
CA Unicenter TNG
BMC
Tivoli
HP-UX ServiceControl
Single Point, Multi-System Configuration Management
ServiceControl Manager (SCM)
SAM
Ignite/UX
SD/UX
Online JFS
Secure Web Console
Central Web Console
System Configuration Repository (SCR)
Fault Management
EMS
EMS HA Monitors
Workload Management
HP-UX Workload Manager (WLM)
PRM
GlancePlus Pak
MC/ServiceGuard
Web Workload Management
WebQoS
New: SCM in June 00
New: SCR in 1Q00
Enhanced PRM June 00
New: HP-UX WLM in 1Q00

easy three-tier packaging

mission-critical environment

mid-tier application server

bundled with every hp 9000 server

servicecontrol mission-critical

HP-UX workload manager (*), MC/serviceguard, ECM (enterprise cluster master) toolkit

servicecontrol enterprise

process resource manager, hp online JFS, glancepluspak, EMS HA monitors

servicecontrol essential

servicecontrol manager, WebQoS peak, Ignite/UX, SAM, SD/UX, EMS, system configuration repository, virtualpartitions core (L,N,Superdome), secure web console (for A,L,N-Class)

recommended add-on products: central web console, openview vantagepoint

(*) HP-UX WLM includes PRM

Toolkits

We are developing toolkits for databases and major applications

These will allow customers to quickly implement policies most appropriate to each

Current Toolkits:

- Oracle Database
- SAS – Under Development

WLM Toolkit Objectives

- Leverage capabilities of WLM:
  - Goal based resource management
  - Dynamic allocation of system resources
- Out-of-the-box integration of WLM with ISV Application
  - Create a toolkit to make it easier to deploy WLM
  - Enable customers to extract performance & other pertinent information from the application
  - Eliminate end user development
- Enhance the Total Customer Experience

Oracle Database Toolkit Features

- Timed SQL – Run a set of SQL statements and pass the execution time into WLM.
- Value SQL – Run a set of SQL statements to extract performance information and pass the resultant value into WLM
- Examples - SQL and WLM config files

WLM Database Toolkit for Oracle Architecture

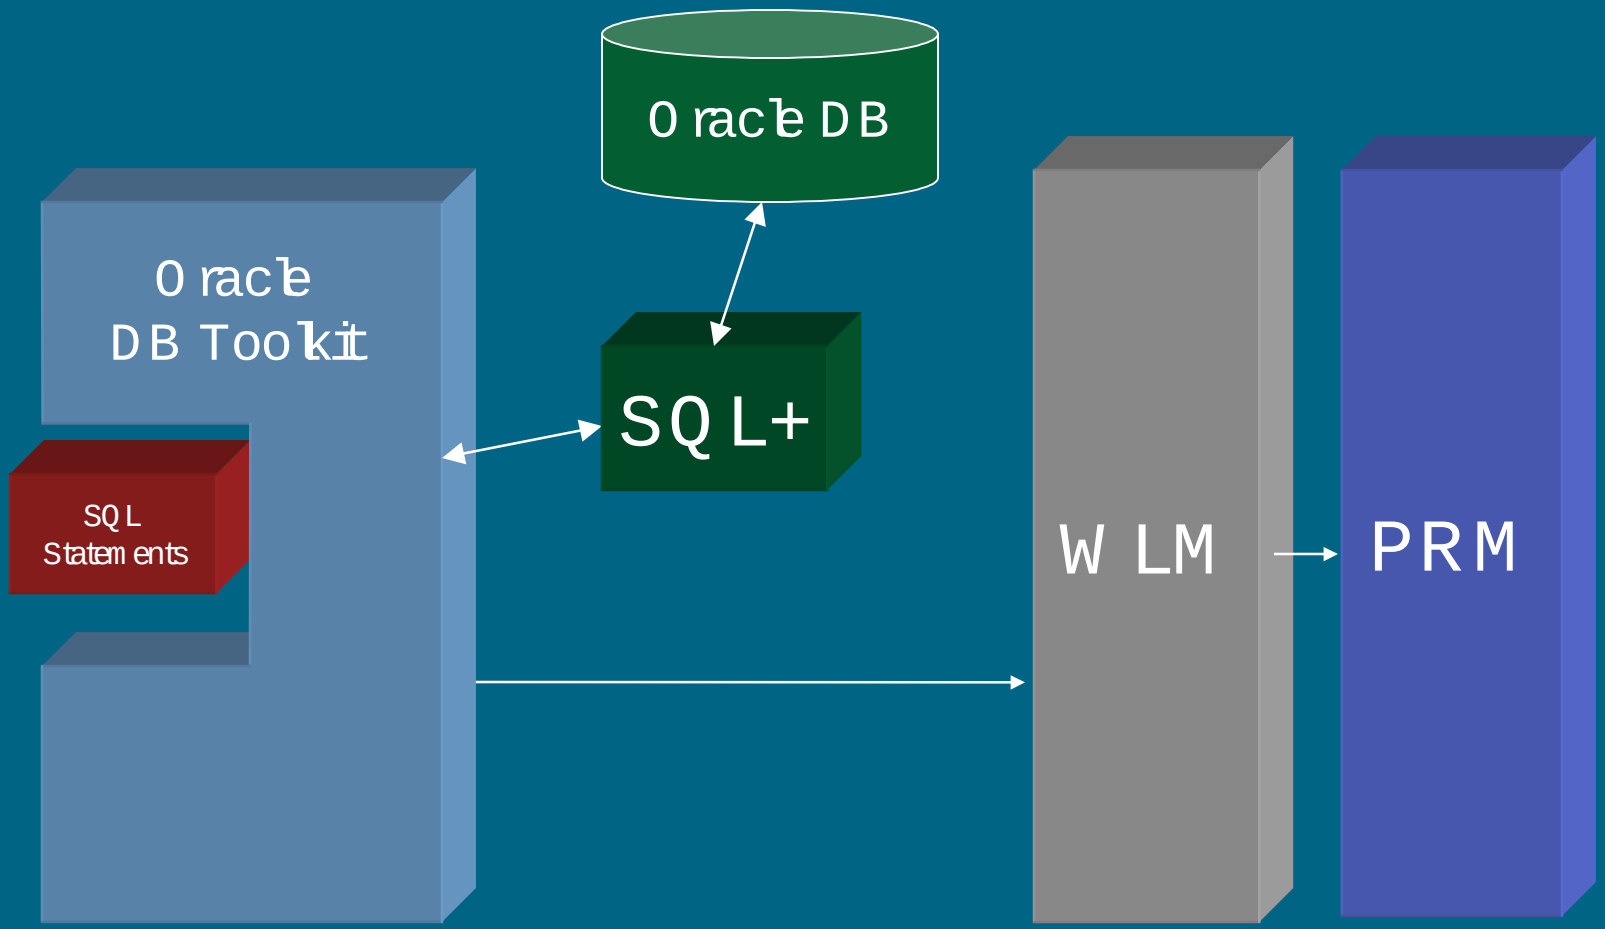

WLM/Oracle use cases

Policy based

- Oracle internal metrics drive entitlement (ex. Num users)
- 'boost' if user X connected
- 'boost' if particular SQL statements or functions

Goal based

- simulated transactions response time drive entitlements
- DB batch jobs

Dynamic Resource Policies using HP-UX WLM

Example: Allocate CPU based on number of users

WLM SAS Toolkit Features

- Facilitate the Moving of Processes into the Correct Resource Groups
- Facilitate Goal-based Service Level Objectives
  - SAS ARM Toolkit
  - Job Duration Goals
- Provide "Express Lane" for Top Priority Jobs
- Examples - SAS and WLM config files to simplify initial setup

WLM SAS Toolkit Architecture (Preliminary)

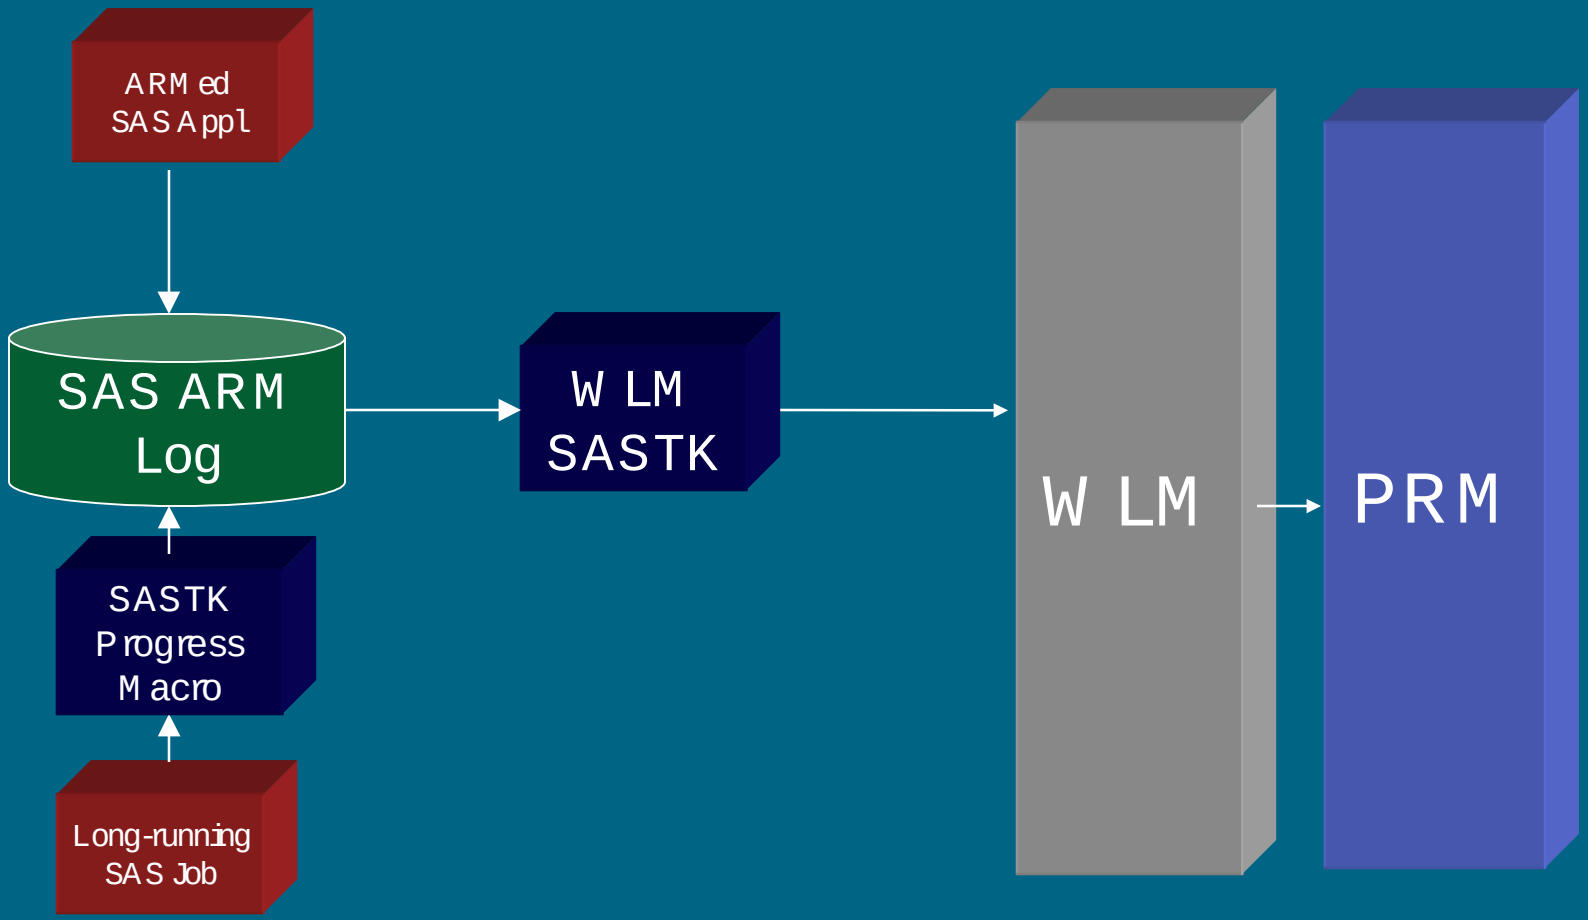

WLM/SASTK Use cases

Policy based

- Ad hoc high priority jobs run as fast as possible
- SAS jobs can move between groups based on resource requirements (CPU vs. I/O)

Goal based

- ARMed apps (SAS or HP implementation) can pass info into WLM
- Can set job duration goals for SAS programs

Summary

HP-UX WLM delivers compute resources when and where they're needed

It has the power and flexibility needed when static resource allocation solutions don't meet the requirements

More information

On the web:

- www.hp.com/go/wlm

or email:

- wlmfeedback@rsn.hp.com


hp
invent