

HEPiX Fall 2013 Workshop

Grid Engine: One Roadmap

Cameron Brunner
Director of Engineering
brunner@univa.com

The UNIVA logo is a red rectangular button with rounded corners and a slight 3D effect. The word "UNIVA" is written in white, bold, sans-serif capital letters across the center of the button.

UNIVA

Agenda



- **Grid Engine History**
- **Univa Acquisition of Grid Engine Assets**
- **What Does Univa Offer Our Customers**
- **Grid Engine 8.0 and 8.1**
- **Grid Engine 8.2 Features**
- **Case Studies**

20 Years of History



- **1992: Initial developments @ Genias in Regensburg**
- **1993: First Customer Shipment (as CODINE)**
- **1996: Addition of GRD policy module**
 - Collaboration with Raytheon & Instrumental Inc
- **1999: Merger with Chord Systems into Gridware**
- **2000: Acquisition through Sun**
 - Re-launch as Sun Grid Engine
- **2001: Open Sourcing**
- **Until 2010: Massive growth in adoption (>10,000 sites)**
- **2010: Acquisition through Oracle**
 - Open Source gets orphaned
- **2011: Key engineering team joins Univa**
- **2013: Grid Engine Assets join team at Univa**

Univa Acquisition of GE Assets



- **Univa acquires Grid Engine software IP**
 - Copyright
 - Trademarks
- **Univa assumes support for Oracle customers under support for Grid Engine software**
 - Version 6.2u6, u7 or u8 binaries
- **Customers:**
 - Can elect to easily, economically upgrade to Univa Grid Engine
 - Services available to simplify upgrade process



What does this mean?

UNIVA

- **Univa is the Home of Grid Engine**
 - Re-unites intellectual property with original developers
 - Fritz Ferstl and his original team
- **Consolidation eliminates confusion**
 - There is one Grid Engine software
 - There is one roadmap



What Univa Offers Grid Engine Users



- **Univa Grid Engine Customer Program**

- Univa Grid Engine is a **drop in replacement** for Sun and Oracle Grid Engine users
 - **Migration is an upgrade**
- Technical Support from experienced team
 - 7 Days a week
 - Severity Level Response Times
 - Severity Level 1 Response Time 4 Working Hours
 - Severity Level 2 Response Time Next Working Day
 - Severity Level 3 Response Time 2 Working Days
 - **Hot fixes, configuration help, tuning guidance**
- Continued Product Development
 - New features and improvements
 - Support for new hardware (GPU, Phi) and Operating Systems
 - Access to add-on products – example: UniSight and Hadoop Integrations
 - Access to improved documentation
 - Security and maintenance patches
- New Products
 - UniCloud
 - License Orchestrator
 - Native Windows implementation
- Indemnification from open source technologies used in Grid Engine

Since SGE 6.2u5



- **28 Releases & Updates until 8.1.6**
 - >500 fixes and enhancements
- **Prior to 8.1.x:**
 - **8.0.0: Released May 11 2011 (open sourced)**
 - Stabilize code base from what Oracle has left behind
 - New documentation
 - New multi-threaded interactive job support
 - **8.0.1: Released Oct 03 2011**
 - Improved core binding
 - Support for NVIDIA GPUs
 - Univa UniSight 2.0 packaged with Univa Grid Engine
 - New Job Submission Verifier Extensions

- **Current version of Univa Grid Engine**

- Version 8.1 – available since Aug 2012
 - Update 8.1.6 released in Sep 2013
- Objective of this release: decrease TCO of Grid Engine at scale
- High Availability:
 - New support for Postgres database job spooling balances speed of submission with reliability in high volume clusters with lots of small jobs
- Performance:
 - Processor core and NUMA memory binding for jobs enables applications to run consistently and over 10% faster. Also increases utilization. Now fully integrated with the scheduler. Binding is guaranteed.
 - Resource maps define how hardware and software resources are ordered and used in the cluster helping to improve throughput and utilization of the cluster
 - Fair Urgency allow for the sharing of the urgency value assigned to valuable resources. Each additional request of a resource receives only a portion of the maximum resource value allowing for intelligent interleaving of multiple high value resources.
- Save Time:
 - Job Classes describe how applications run in a cluster helping to slash the time to onboard and manage workflow
 - Improved Job Debugging and Diagnostics lets administrators discover issues in less time
 - Documented and tested integrations with common MPI environments save valuable time since we did the work

Case Study: Oil & Gas

Decision support for petroleum and gas exploration

UNIVA



أرامكو السعودية
Saudi Aramco



CHALLENGE

It's a multi-million dollar mistake to drill a dry hole (up to \$100M in deep water)

Oil and gas exploration driven to ever more difficult, previously neglected sources

Reduce margin of error → drastic increase in computation environment for decision support (100,000 cores and more)

SOLUTION

Univa Grid Engine enables the scalability and provides the workflow and application support

Policies and access control to manage project priorities and resource entitlements

UniSight reporting and analytics for accountability and capacity planning

CUSTOMER VALUE

- Optimized ROI for infrastructure spending
- Accelerated time to results
- Avoided staff increases by leveraging shared infrastructure

Case Study: Chip Design

UNIVA

Mentor
Graphics®

SAMSUNG

SYNOPSYS®

Panasonic
ideas for life

Google™
MOTOROLA

CHALLENGE

Complete reliance on computerized simulation in chip design cycle for test and verification of integrated circuits up until silicon production

Typical environments have tens of thousands of cores and process 30M jobs per month

Need to manage SW licenses costing \$500K per seat across projects and departments

SOLUTION

Univa Grid Engine for the required workload and resource management capabilities including sophisticated prioritization policies

Provide reporting, auto-scaling and monitoring to gain efficiencies in existing cluster

Univa License Orchestrator for the crucial license management requirements

CUSTOMER VALUE

- Reduced SW license costs
- Accelerated time to market and increase quality of product
- Avoided staff increases by leveraging shared infrastructure

Archimedes Shared Infrastructure

UNIVA



Archimedes Inc. is a healthcare modeling organization. Its core technology - the Archimedes Model - is a clinically realistic, mathematical model of human physiology, diseases, interventions and healthcare systems.

CHALLENGE

Need to run new data analytical models on Hadoop framework

Limited budget to buy new compute resources

Time to operationalize mission critical model

SOLUTION

Univa Big Data Platform with Hadoop automation that supports multiple applications on shared infrastructure

Provide reporting, auto-scaling and monitoring to gain efficiencies in existing cluster

Univa integrated solution enabled Archimedes Big Data solution to “go live” in dramatically less time

CUSTOMER VALUE

- Reduced infrastructure capital expense by 50%
- Accelerated time to market
- Avoided staff increases by leveraging shared infrastructure

Grid Engine Roadmap



- **Univa Grid Engine 8.1.5-8.1.6**
 - 8.1.5 Released July 2013
 - 8.1.6 Release September 2013
 - 8.1.7 Scheduled for December 2013
- **Univa Grid Engine 8.2.0**
 - Scheduled for Q1, 2014

Upcoming Features



- **Native Windows**
 - No more SUA/SFU or cygwin requirements
- **CGROUPS Support**
 - Likely pre-released with 8.1.7
 - Better control over execution environment
- **Read-Only Qmaster Thread**
 - Improved qstat performance in large environments
- **DRMAA v2**
- **Performance Optimizations**
 - Improved qsub -sync performance
 - Analysis spooling performance at job start and finish
 - General performance improvements

Thank You!

Cameron Brunner
brunner@univa.com

The UNIVA logo is a red rounded rectangle with the word "UNIVA" in white, uppercase, sans-serif font. It is positioned in the center of a black horizontal bar at the bottom of the slide.

UNIVA