OpenPOWER™

IBM Platform Computing : infrastructure management for HPC solutions on OpenPOWER

Jing Li, Software Development Manager
IBM

Scale-out and Cloud Infrastructure Management Needs

Traditional HPC cluster	• **Provisioning & monitoring system for scale-out computing infrastructures**
Multi-tenants HPC environment	• **Technical computing capacity to multiple departments, projects, and users (multi-tenants)** • **Self service capability**
Data Management	• **Elastic Storage Management and monitoring** • **Elastic Storage Appliance management and monitoring**
Cloud Infrastructure	•VM, virtual network, & virtual storage management •Underlying infrastructure (undercloud) management •Configuration management
Capacity overflow	•Automated cloud bursting capabilities •Flexibility and benefits of VMs
Big data clusters	•Multi-tenancy •Enterprise class system management capabilities (audit, Role Based Access Control etc.)

OpenPOWER™

Architecture Overview

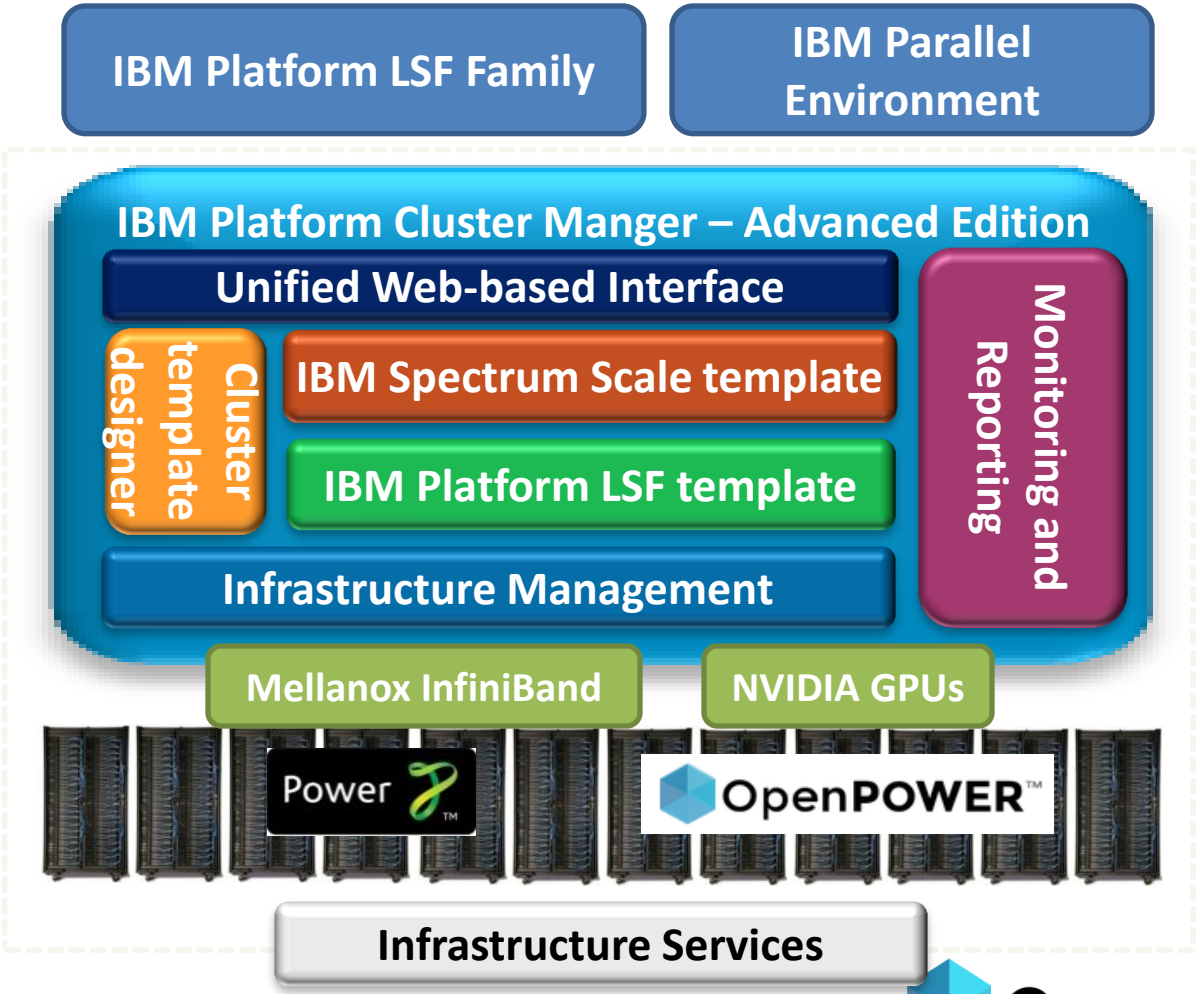

IBM Platform Cluster Manager

Overview

Powerful lifecycle management for scale-out cluster environments

Key Capabilities

- Simplified management with cluster template designer
- Scales from single clusters to complex multi-team environments
- Robust, scalable alerting and reporting
- Automated infrastructure management – one-click cluster deployment
- Spectrum Scale cluster support

Benefits

- Faster time to cluster readiness
- Unified interface for management and monitoring
- Increased administrator productivity
- Single infrastructure supporting multiple business needs

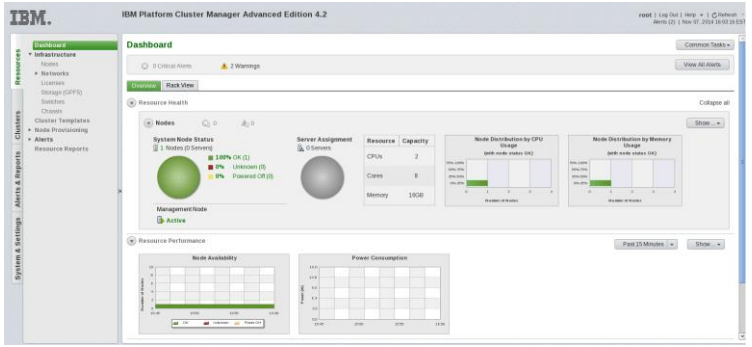

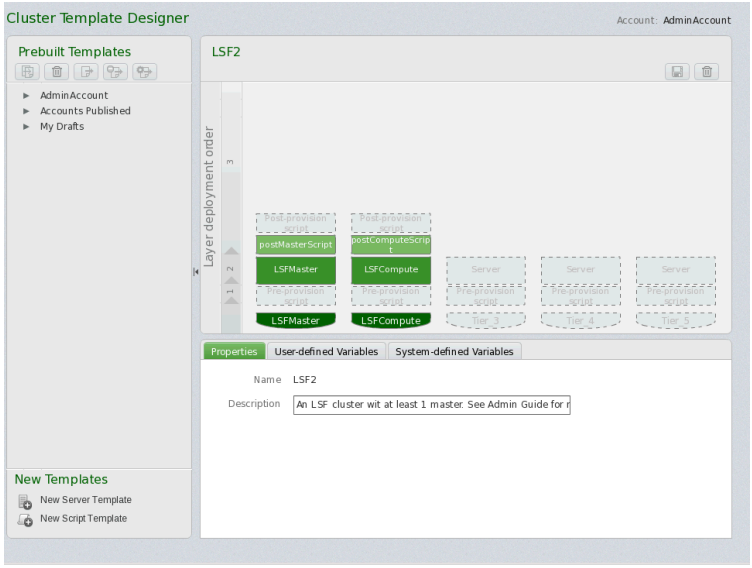

OpenPOWER™

Infrastructure at a glance

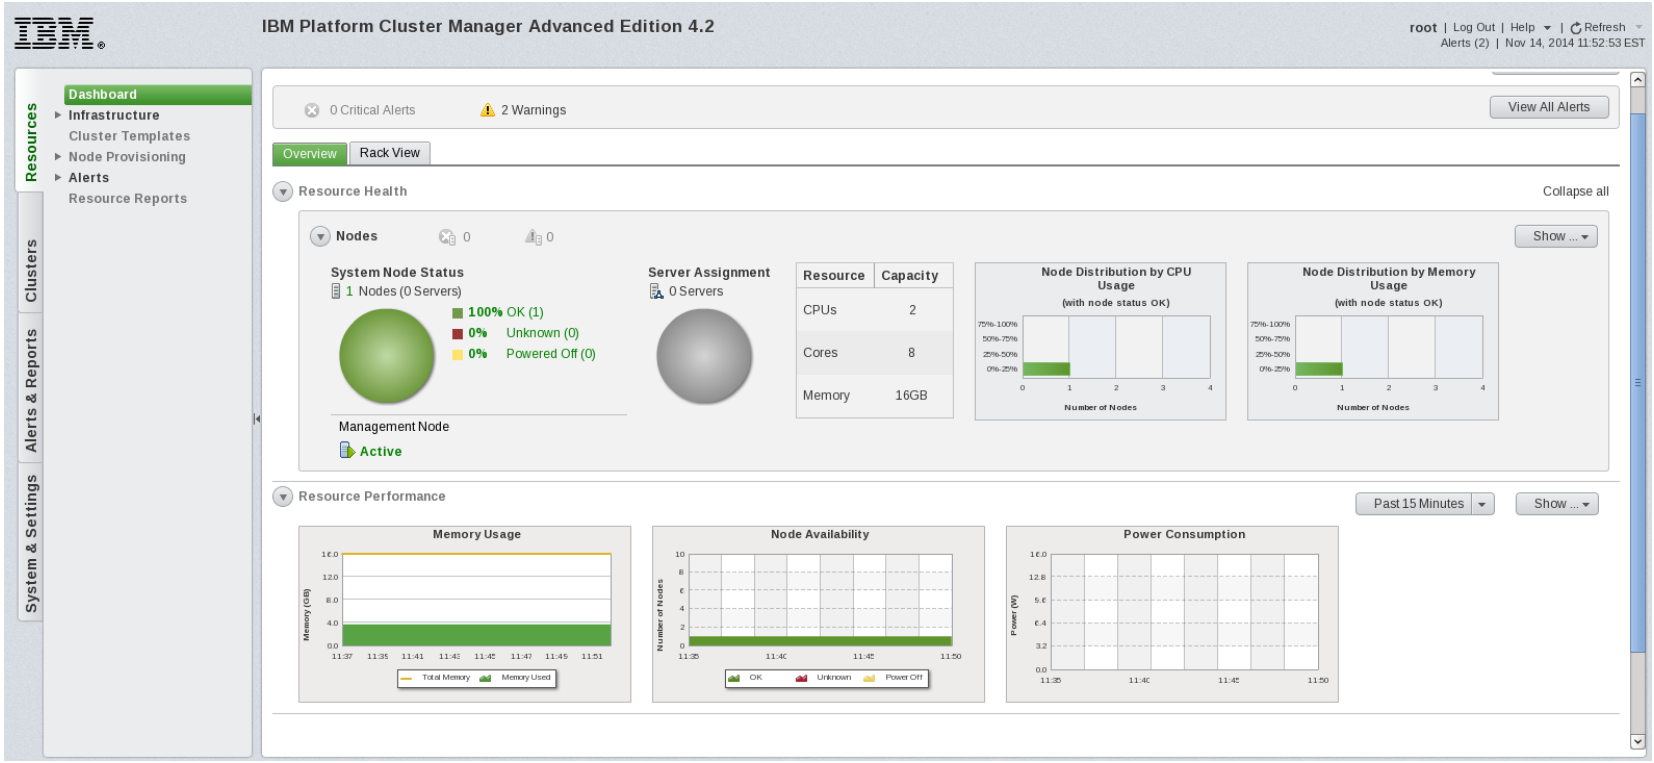

- Single interface for management and monitoring of multiple clusters
- Dashboard provides overview of resources, allocations

OpenPOWER™

Rackview – graphical cluster overview

- 2D representation of the data center (racks, nodes)
- It allows administrator to quickly examine the status of individual nodes
- Admin can drill into node details by clicking the node
- Chassis console can be launched from the rack view
- Overview of entire cluster at a glance

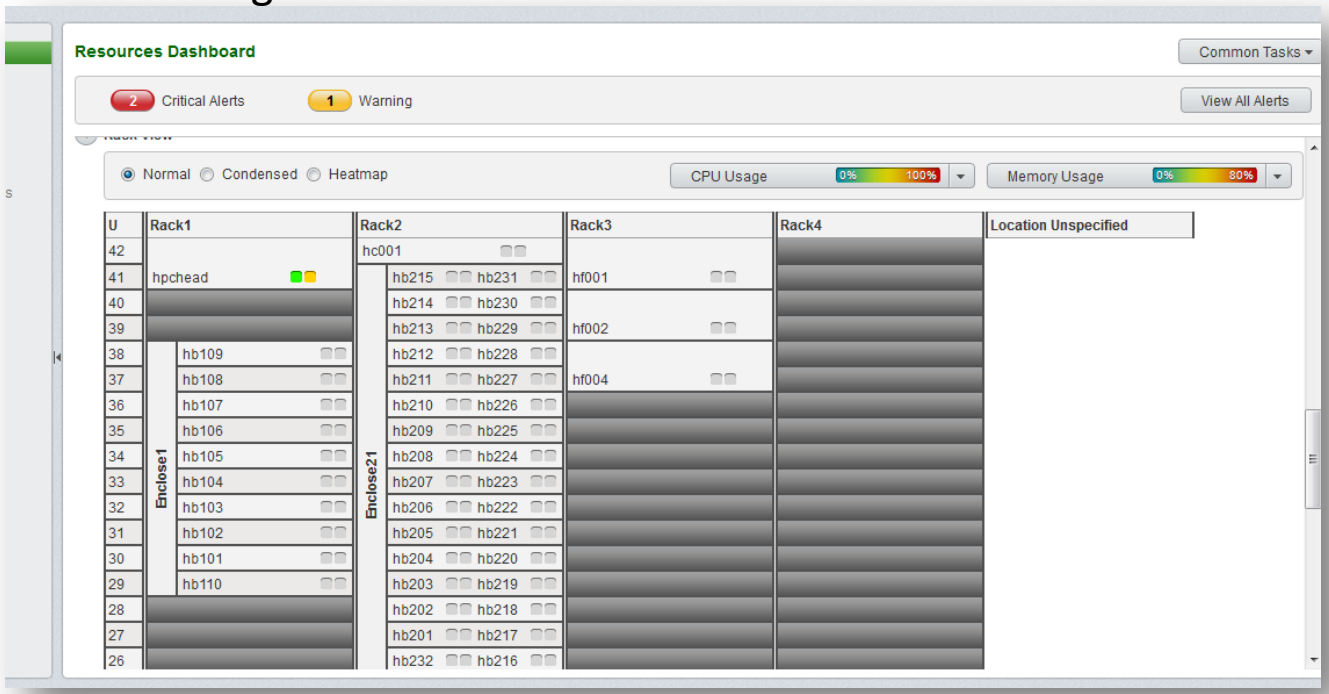

Drag and drop cluster builder

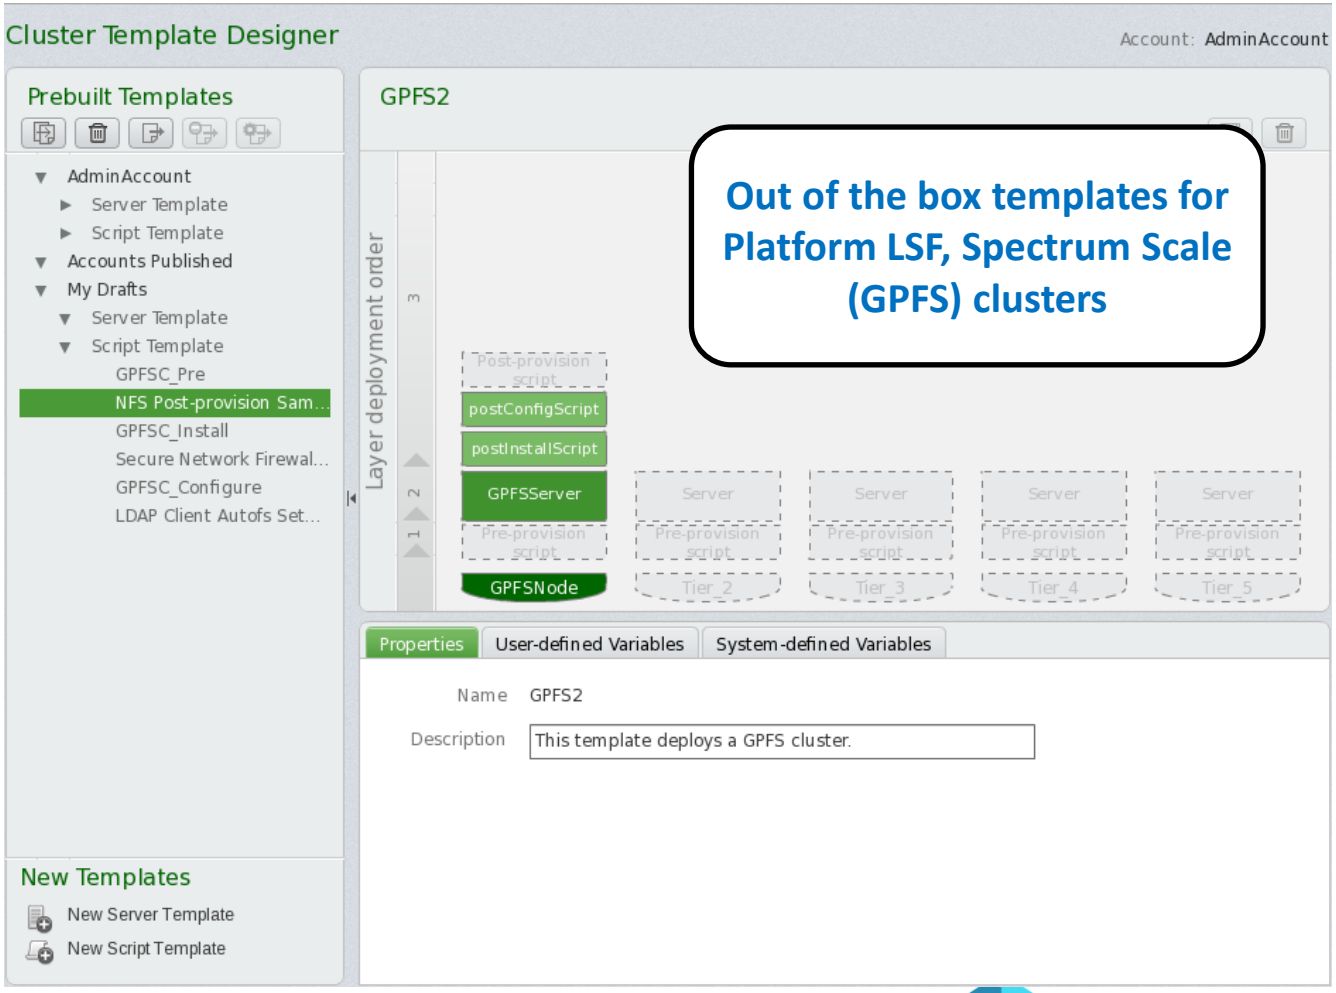

Managing node personalities

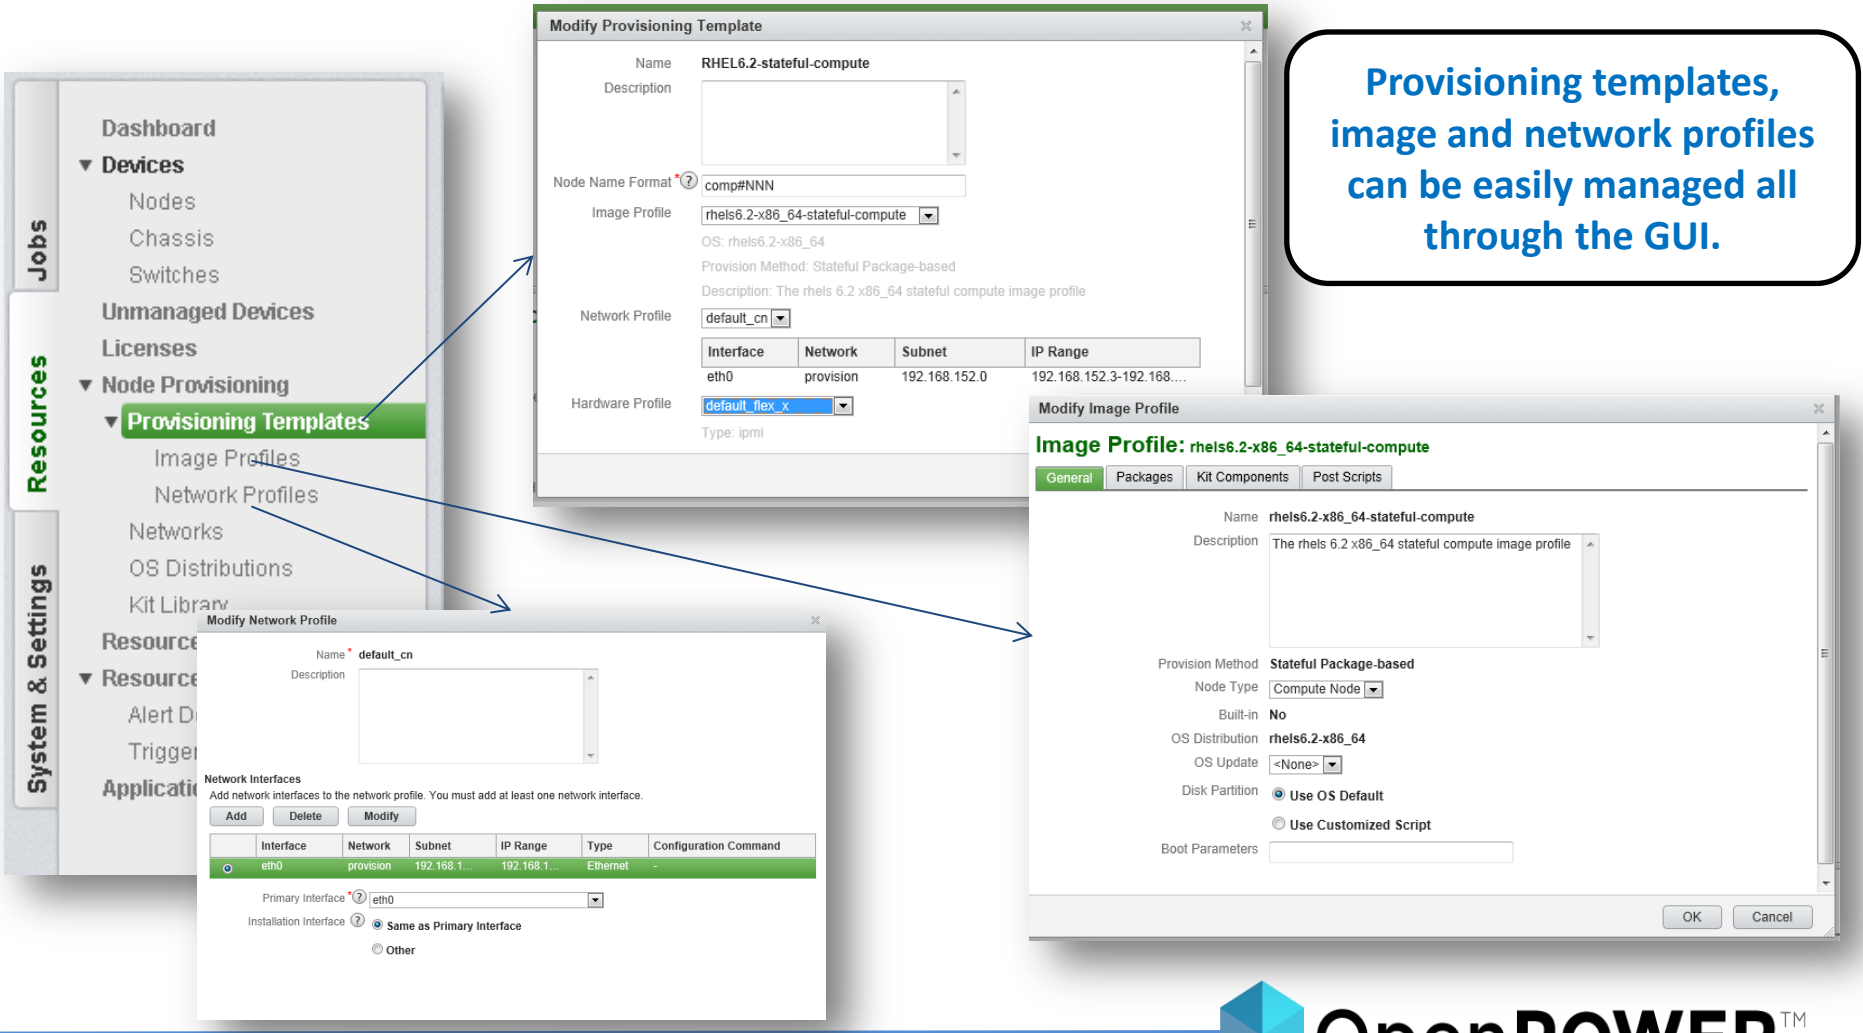

Provisioning templates, image and network profiles can be easily managed all through the GUI.

OpenPOWER™

Hardware management and monitoring

- Node detail is monitored with system & performance
- Management capabilities include: power cycling, firmware updates, OS reboot, reprovision, synchronize, node LED control, BMC, SSH, VNC console
- Integrated network switch and chassis monitoring
- Integrated Spectrum Scale monitoring

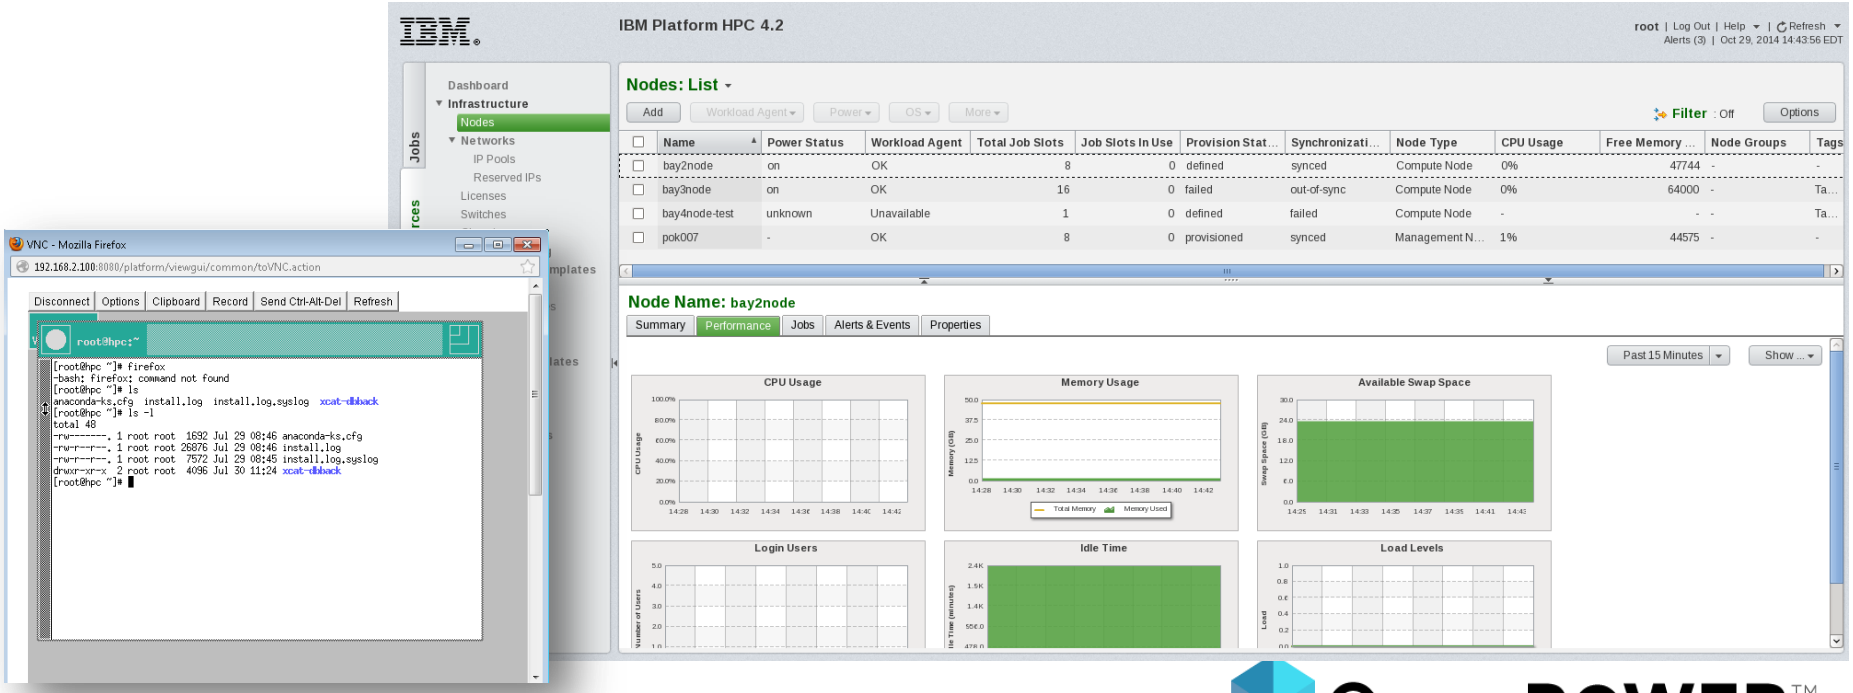

OpenPOWER™

Alerts highlight potential issues in the cluster

- Fully customizable; alerts can be defined using any monitored metrics
- Alert can trigger an automated pre-defined action
- Alert history shows the detail of the triggered alert

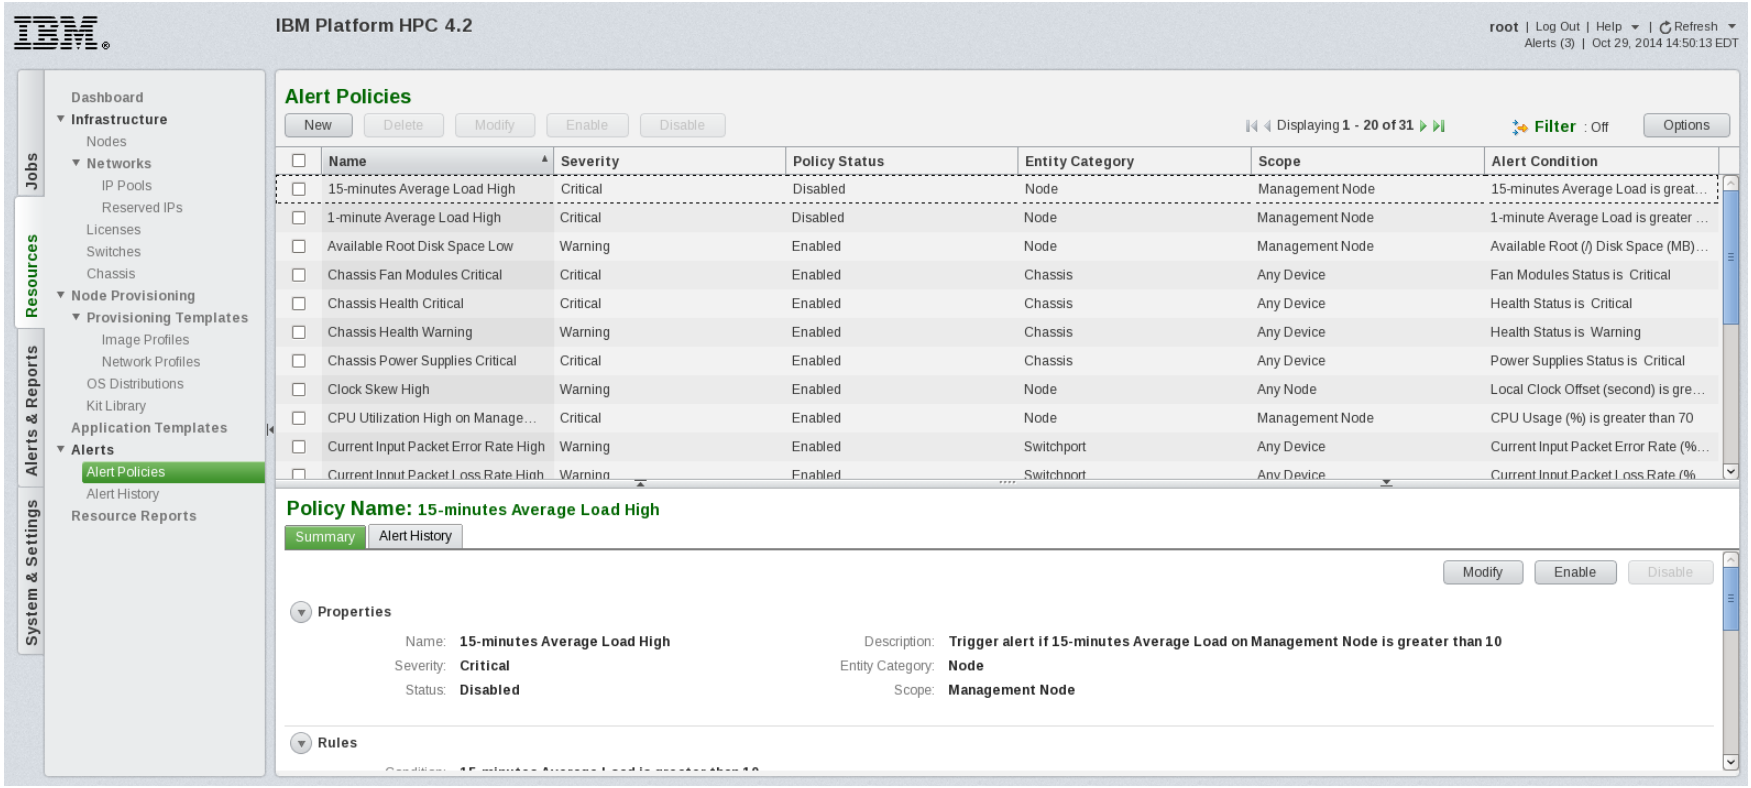

OpenPOWER™

Resource Reporting – Gauge utilization of the cluster

- Historical reports can be generated for
  - Cluster availability
  - Cluster performance and usage
  - Free application licenses

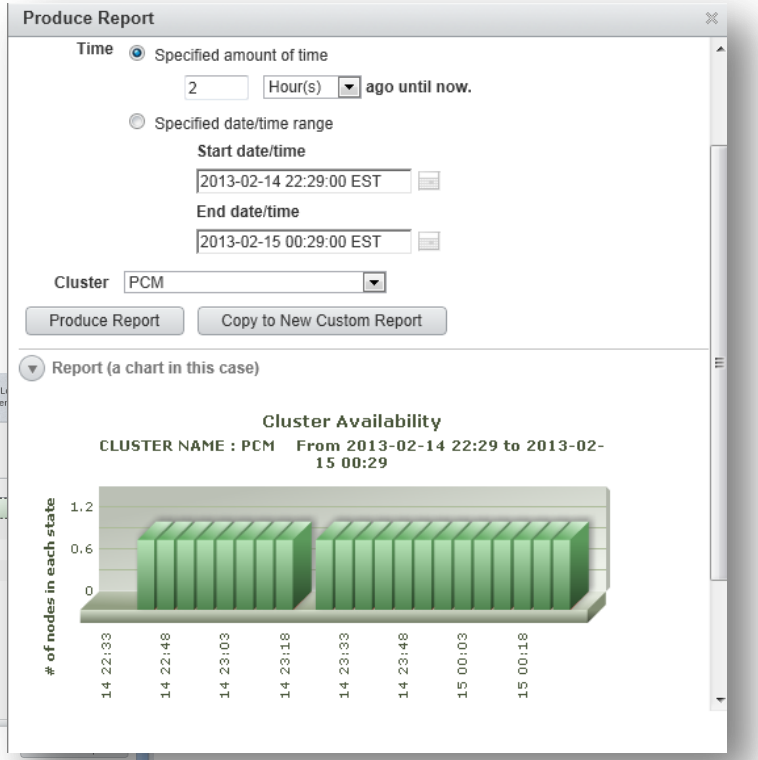

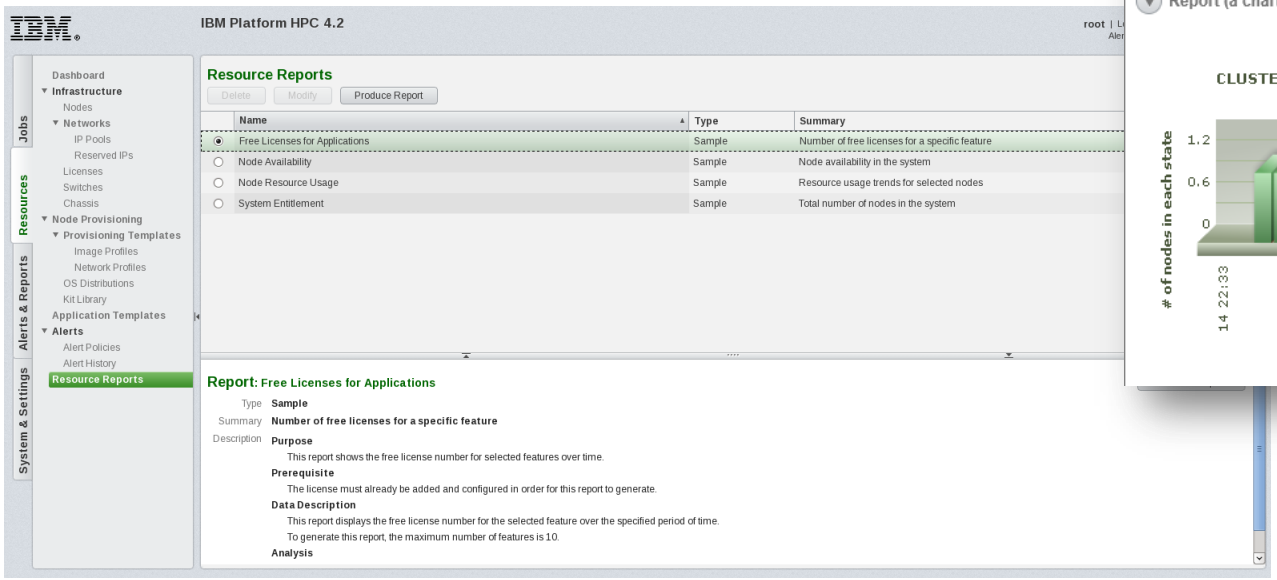

OpenPOWER™

Comprehensive HPC Software Stack

	Products	Client Benefits
Systems Management	Platform Cluster Manager – Advanced Edition	•Ease of Use: web portal •Customizable: admin productivity •Faster time to system productivity •Robust monitoring
Application Runtime	PE Runtime Ed ESSL / PESSL	•Optimized Parallel Runtime •Optimized LAPACK and ScaLPACK libraries •User controlled workflow support
Development Productivity	PE Developer Ed XL Compiler	•Modern application development environment using Eclipse •Performance analysis tools to help analyze applications •Optimized compiler for Power
Workload Management	Platform LSF	•Optimized utilization of resources •Policy, energy and resource aware scheduling •Robust add-on features
Data Management	Spectrum Scale HPSS TSM	•Scalable/reliable storage for parallel filesystem (GSS) •ILM for transparent migration of data from storage to tape and back
Application Environment	Platform Application Center Platform Process Manager	•Simplify job submission for repeatable workload: customization •Customizable •Faster time to system productivity

Platform Computing

Linux

Power Systems™ S824L

Power Systems

NVIDIA

Mellanox TECHNOLOGIES

Power Systems S822L

OpenPOWER™

IBM Parallel Environment

High Performance Execution Environment to Take Full Advantage of Scalable Compute Resources

- Parallel Operating Environment (POE) for submitting and managing jobs.
- IBM's MPI and PAMI libraries for communication between parallel tasks.
- A parallel debugger (pdb) for debugging parallel programs.
- IBM High Performance Computing Toolkit for analyzing performance of parallel and serial applications.
- Integrated with LSF to assist in resource management, job submission and node allocation

What's New:

- Ubuntu 14.04.1 Little Endian NV (Non Virtulaized)
- MPICH as BASE and collective performance improvements
- MPI 3.0 (via MPICH)
- MPI I/O Improved Performance

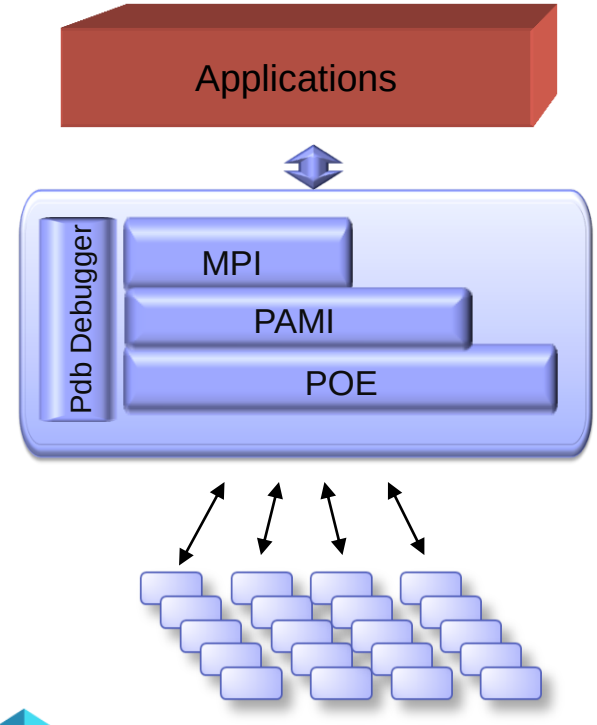

OpenPOWER™

IBM Platform LSF

Most Complete

- Advanced, feature-rich workload scheduling
- Robust set of add-on features
- Integrated application support

Most Powerful

- Policy & resource-aware scheduling
- Resource consolidation for max performance
- Advanced self-management

Most Scalable

- Thousands of concurrent users & jobs
- Virtualized pool of shared resources
- Flexible control, multiple policies

Best TCO

- Optimal utilization, less infrastructure cost
- Better productivity, faster time to result
- Robust capabilities, administrative productivity

Linux on POWER8 LE support

IBM Platform Application Center

- Increase user productivity with browser based access for job submission, management
- Capture best practices with guided submission (templates)
- Enable access via mobile devices with web services
- Support for 2D/3D remote visualization

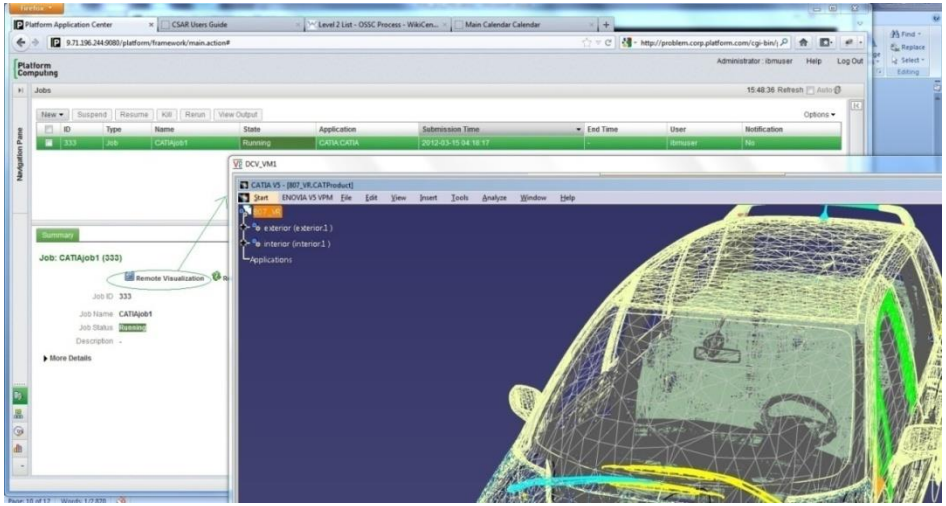

Linux on POWER8 LE support

IBM Platform Process Manager

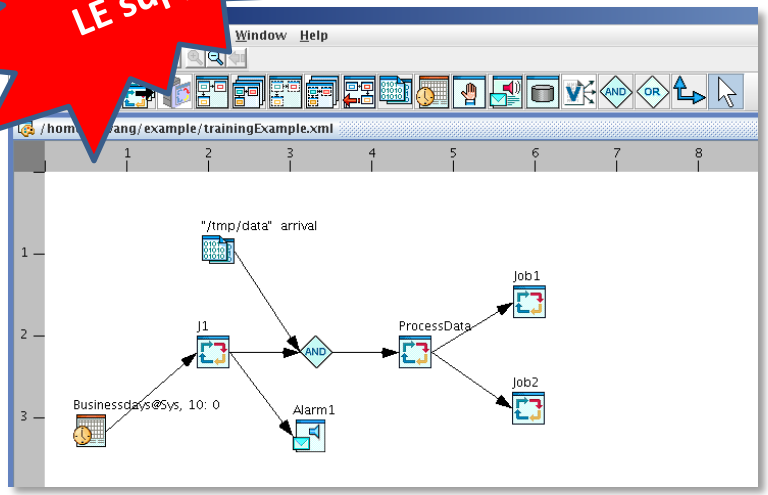

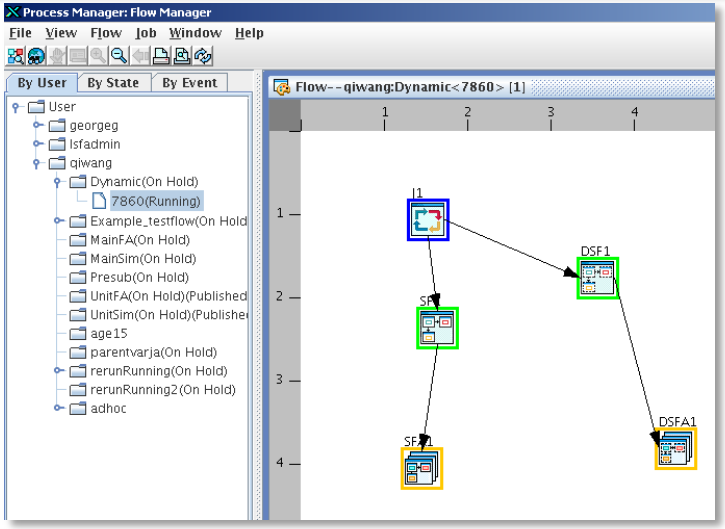

Platform Process Manager Flow Editor

- Intuitive drag-and-drop interface
- Creates self-documenting flows
- Support for sub-flows, job arrays
- Rich error-handling / retry capability
- Save workflows in XML format
- Publish flows directly to Flow Manager

Platform Process Manager Flow Manager

- Manages multiple flows for multiple users and groups simultaneously
- Monitor workflow execution graphically
- Trigger flows automatically through calendar events, the flow manager or the command line.

Genomic medicine – reference architecture

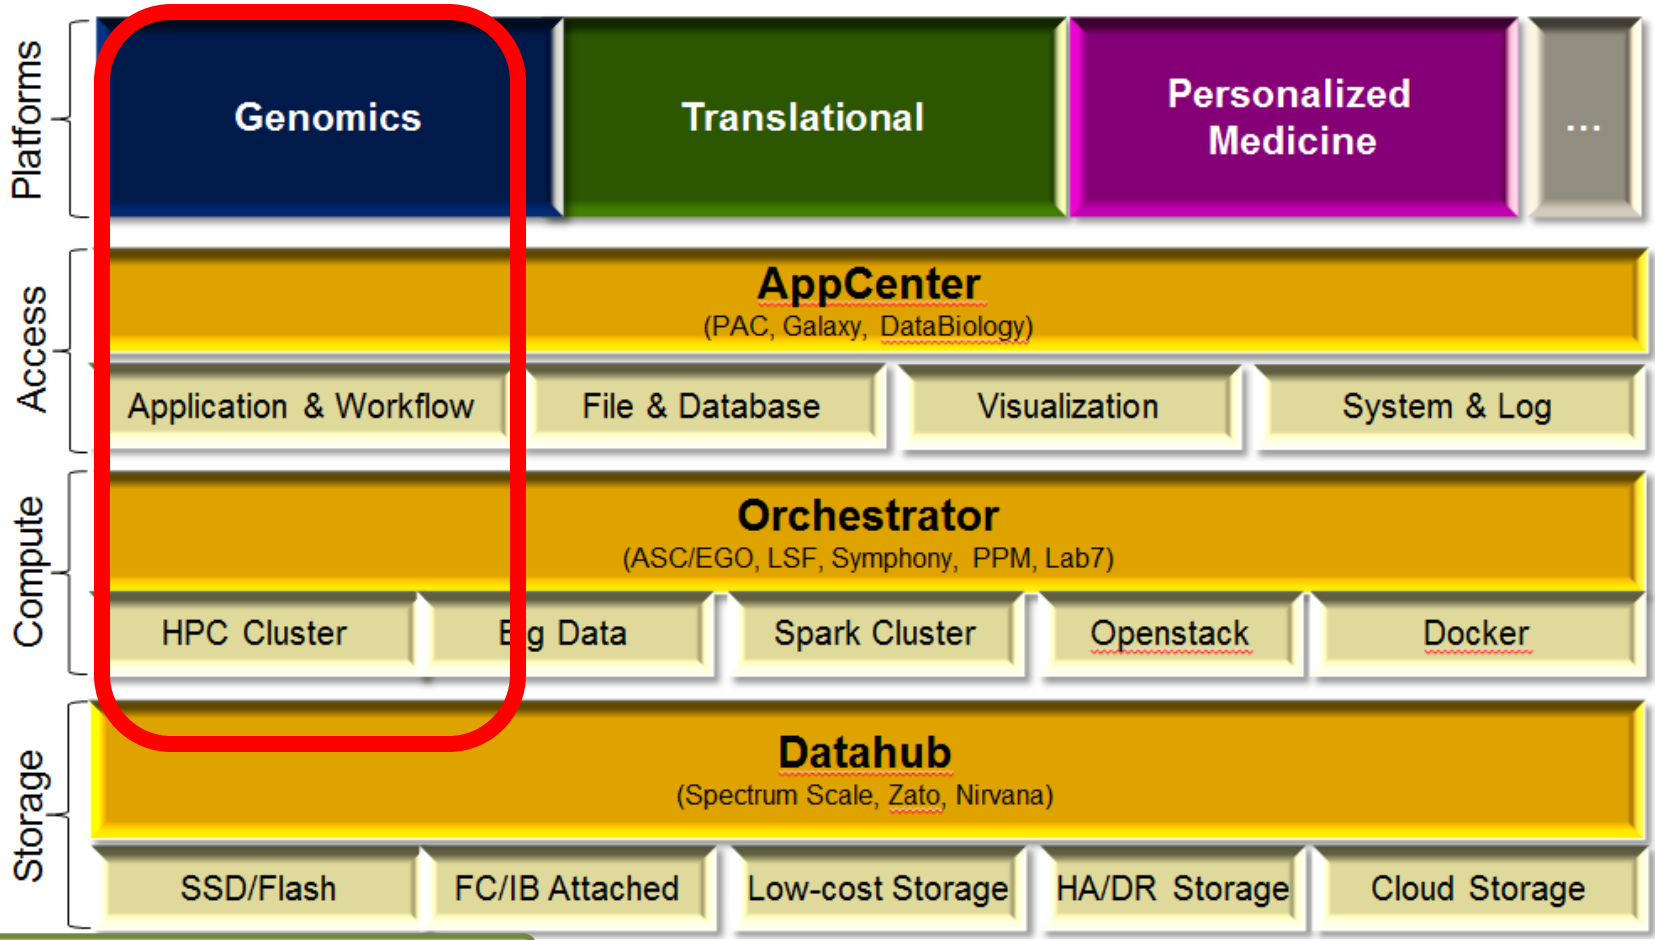

http://www.powergene.net

Links

- IBM Platform Computing product information (https://ibm.biz/BdXBDR)
- Service Management Connect – Technical Computing Community (https://ibm.biz/BdFr8R)
- IBM Knowledge Center (https://ibm.biz/BdXBDX)

Contacts

- Development Manager: Jing Li (jingili@cn.ibm.com)
- Product Management: Mehdi Bozzo-Rey (mbozzore@ca.ibm.com)
- Product Marketing: Gabor Samu (gsamu@ca.ibm.com)

Q&A