

SSS Build and Configuration Management Update 1/2005

Narayan Desai
desai@mcs.anl.gov



Overview

- Infrastructure
 - LRS support
- System Management
 - Prepare for BG/L
 - Other MPPs?
 - Deploy diagnostics with OSCAR release

Infrastructure

- LRS support
- SSS/SDK
 - Match support
 - Function call support
 - Return control
- New components
 - service directory
 - more to come
- New schemas
 - More stringent

System Management

- LRS conversions
 - NSM with system diagnostics
 - Build System
 - Hardware Manager
- Integrated NSM with system diagnostics into OSCAR release

BG/L

- System arrived 1/05
- Not functional/accepted from IBM yet
- Plan to run SSS software on it
 - Possible with large amount of reuse
 - System architecture requires different component instances (in some cases)
 - Many components remain the same

BG/L Architecture

- Single process/user per compute node
- System diversity
 - 4 (!) different kernels (3 Linux/1 LWK)
 - 3 different hardware architectures
- JTAG interface to compute nodes
- Nodes allocated in 32/256/512/1024/+ partitioning scheme
- RAS hardware features

BG/L Implications for SSS

- Single process per node
 - Implications for Process Manager, System Monitor, System Management tools
 - No non-application endpoint on compute nodes
- No direct TCP support
 - System calls forwarded to IO nodes
 - 64 way overcommitted
- Allocation granularity issues
 - Process Manager works differently
- Hardware interface to nodes
 - RAS data
 - Performance countersp

Development/Deployment Plan

- Chiba resource management incorporated (as-is)
 - Scheduler
 - Queue Manager
 - Allocation Manager
- New implementation of Process Manager
 - System partitioning
 - BG/L executable/kernel loading scheme
 - New PMI implementation
- New configuration management components
 - Different diagnostic model
 - What are software configurations on BG/L?
 - Verify suitability of BCM interfaces for non-cluster machines