



Overview of Librato Technology

Joseph Ruscio
Chief Software Architect

HP & Intel EDA IT Forum 2009
April 7, 2009

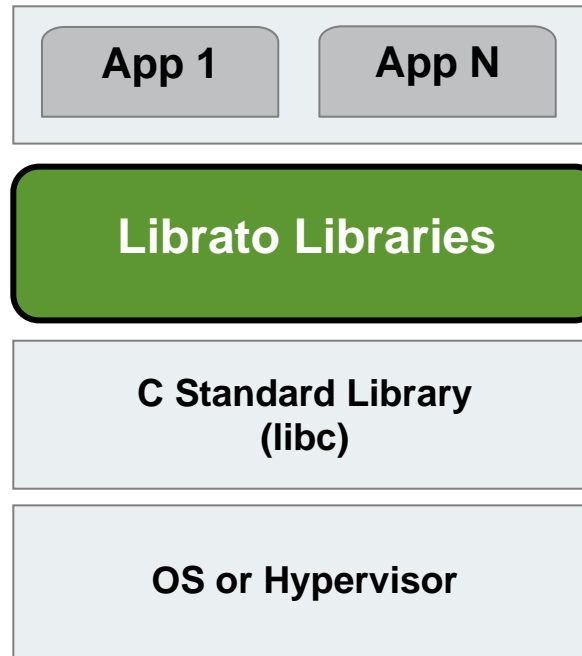
Agenda

- Librato Overview
 - Company Information
 - Core Technology and Design Philosophy
 - Librato Solutions
- Librato Smart Suspend
- Discussion / Q&A

Librato Overview

- Founded in August 2005 as Evergrid
- Well funded by top tier VCs
- Headquartered in Santa Clara, CA
 - R&D in Blacksburg, VA and Pune, India
- Company Focus
 - Improve data center efficiency
 - Application focused, easy to deploy solutions
- 10 patents --1 granted, 8 in process, 1 provisional
- Multiple partners in technology and distribution channels

Design Philosophy



- User space technology
 - Completely transparent to both application and operating system
 - No recompiling or re-linking of applications, no kernel modifications
 - Very low overhead

Librato Solutions

- Smart Suspend
 - Improve efficiency by safely suspending applications and immediately releasing resources such as CPU, memory, I/O bandwidth and application licenses
- Availability Services
 - Increase throughput with checkpoint/restart, providing fault tolerance, preemption and migration capabilities
- Load Manager
 - Guarantee resource allocations to processes or applications to simplify server consolidation and maximize utilization
- Load Monitor
 - Gain precise understanding of the demand for and current usage of resources such as memory, CPU, network I/O and storage I/O

Example: Semiconductor Company

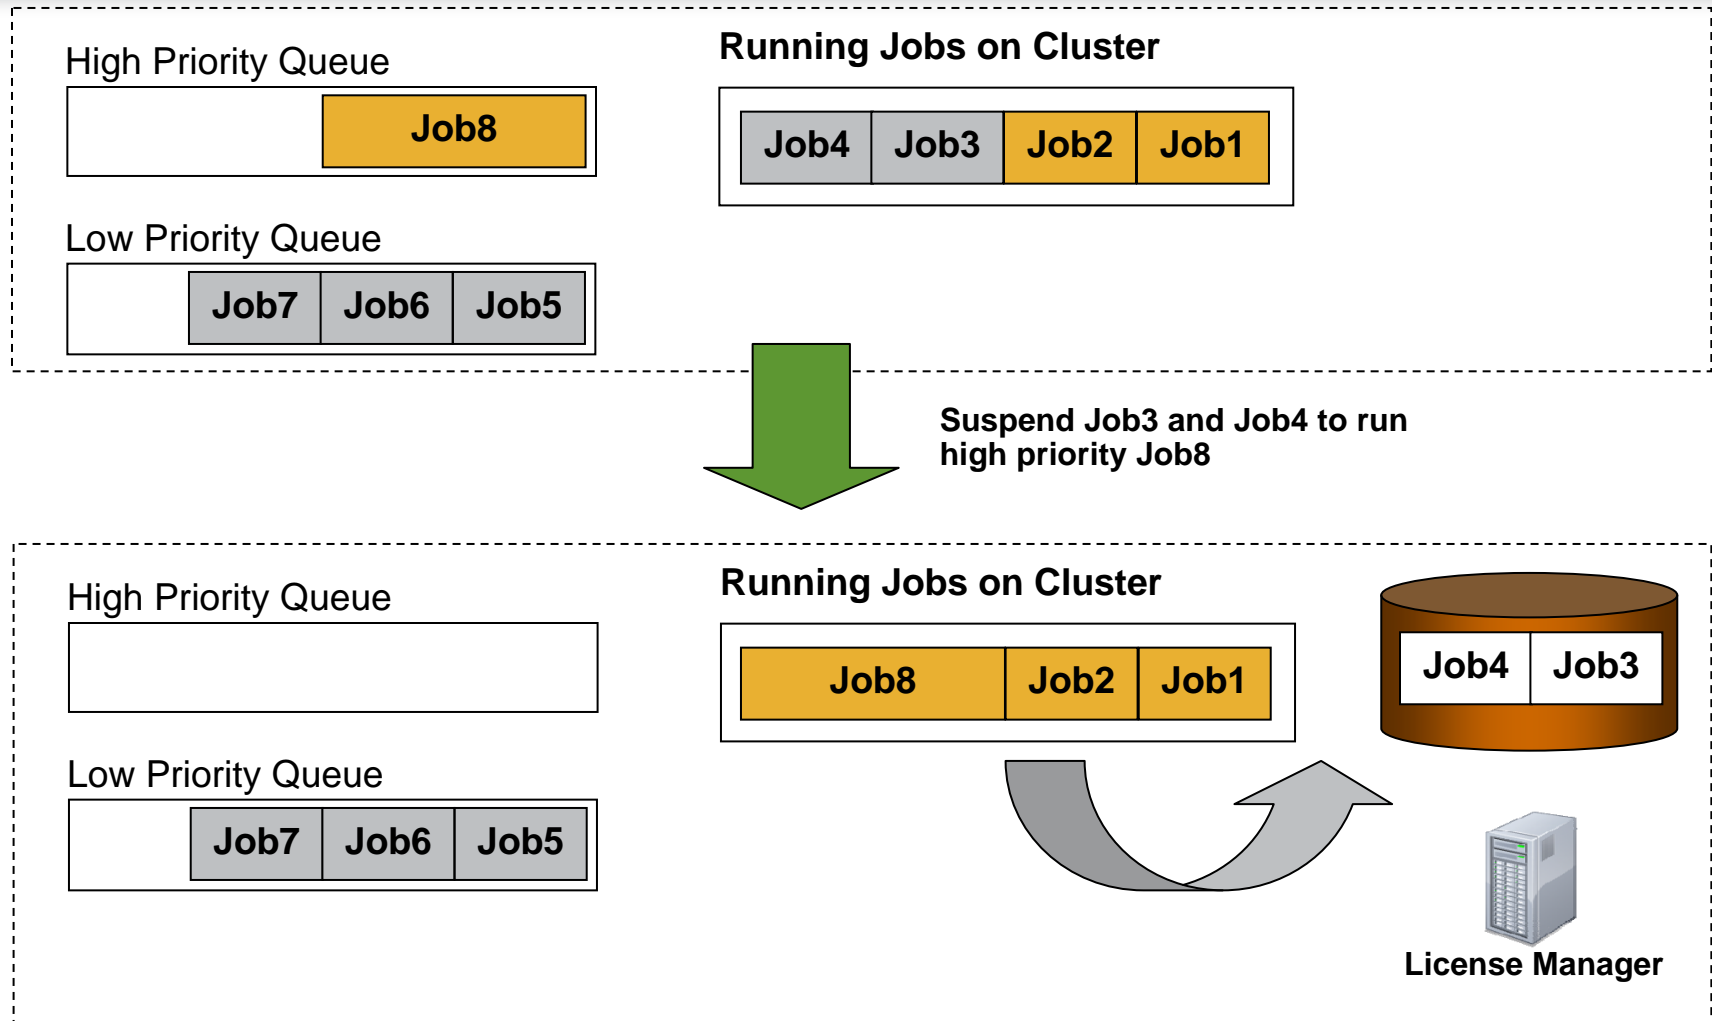
- Goal: Increase server utilization in 1000-node grid
 - Grid handles interactive (design) and background (chip lay-out) workloads
 - Current inefficiencies
 - Preemptive scheduling requires lower priority jobs to be aborted
 - Interactive workloads have very low utilization (<10%)
 - Background workloads do not optimally use resources (e.g. processors idle when waiting for I/O)
- Server virtualization not an attractive solution
 - Too many OS images to manage, performance overhead, complexity
- Chose Librato Smart Suspend and Load Manager
 - Smart Suspend enables preemptive scheduling by suspending jobs, temporarily freeing all resources and software licenses
 - Load Manager allows background jobs to “sponge up” all resources not used by interactive workload
 - Load Manager allows multiple background jobs to coexist on same “slot”
- **Result: Improve resource utilization by >30%**

Librato Smart Suspend

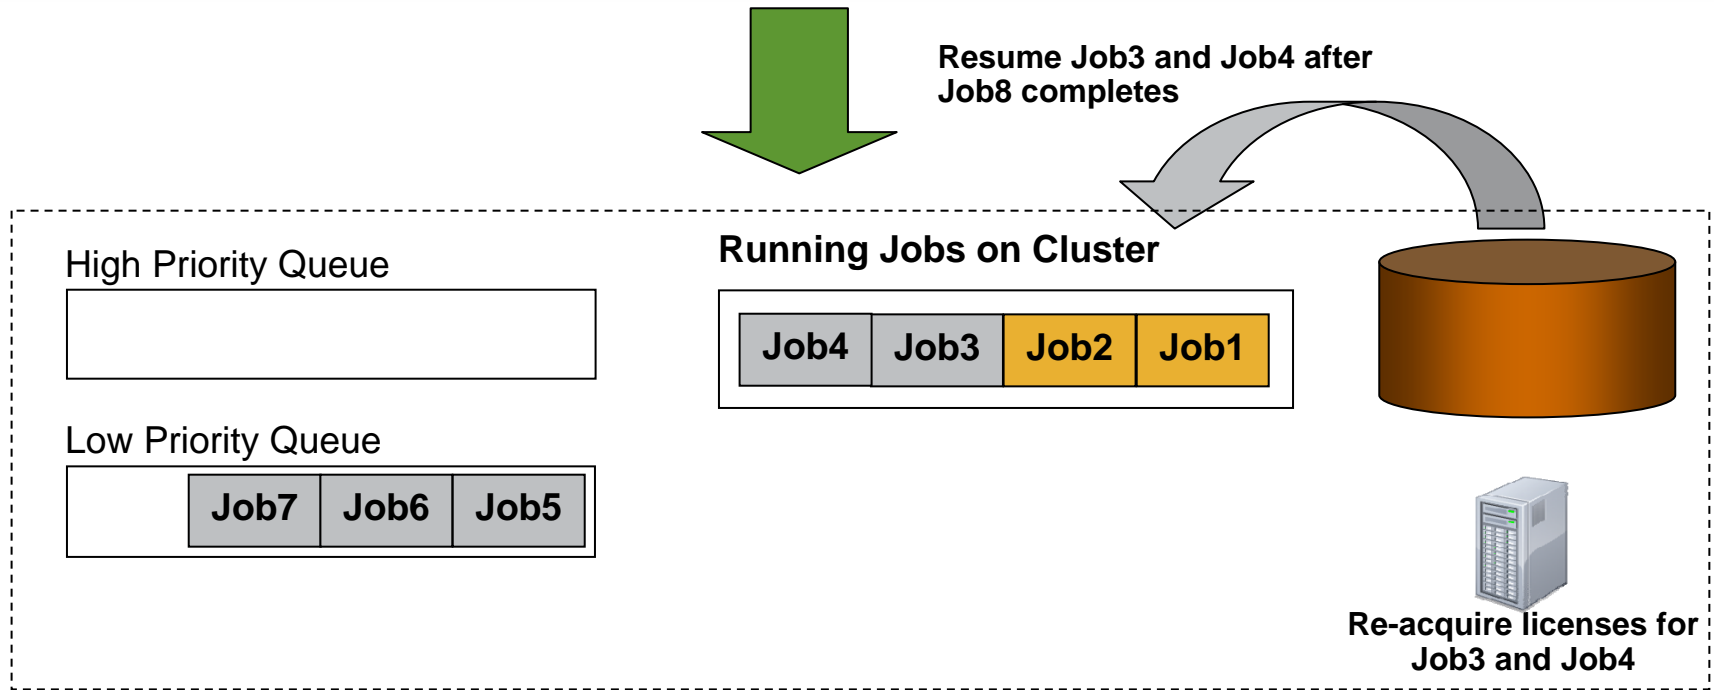
Today's Challenges

- Quality of Service vs. resource utilization
 - Manage long running jobs with non-deterministic run-times
 - Guarantee SLAs for high priority jobs while managing low priority jobs
 - Maximize resource utilization
- Simple to deploy

Preemptive Scheduling



Preemptive Scheduling (cont)



Smart Suspend Benefits

- Asynchronously suspends arbitrarily complex applications
 - Preserves completed work
- Immediately reassigns resources to new high priority jobs
 - No impact to overall utilization
- Suspends interactive applications during off-work hours
 - Immediate and significant ROI
- Installs effortlessly into existing environments
 - Seamlessly integrates with market leading queueing systems
 - Supports entire cluster or single job granularity
 - Supports all job types (Serial, SMP, Parallel)

Smart Suspend: Under the Covers

- Immediately releases resources
 - Correctly records all memory state to stable storage
 - Supports diskless nodes
 - Suspends all CPU and I/O activities
 - Releases expensive software licenses
- Abstracts complex groups of processes as single application
 - All child processes inherit the Smart Suspend configuration
 - Suspends/Resumes entire group of processes by default
 - Also enables selective suspension of single process
- Keeps processes intact to ensure the application resumes smoothly
 - Socket and network connections are maintained
 - Applications maintain their process IDs, file descriptors, etc

Smart Suspend Advantages

- Elegant and uniform process control: suspend or resume all processes of a job on all hosts with one command
- No modifications to applications or the operating system
- Memory
 - Writes out memory on demand
 - Write memory to any storage: Works for diskless nodes, not swap-space dependent
- Licenses
 - Uniform license handling for all applications
 - License immediately made available after suspend
 - Not susceptible to the `lrmremove` waiting period

Thank You

Librato Contact: Jack Hanna (jhanna@librato.com)