Scyld Software

Donald Becker, CTO
David Ingersoll, VP Sales

Highly SCYLD

AGENDA

- Introduction
- Brief History of Linux Clustering
- Traditional Linux Clustering
- Scyld ClusterWare Overview
- Case Study – Applied Biosystems
- Q&A

Corporate Introduction

- **Penguin Computing Founded May 1998**
- **Acquired Scyld Software 2003**
  - » Donald Becker – CEO Scyld Software, primary contributor to the Linux kernel and co-inventor of Beowulf – becomes Penguin's CTO
- **Over 2000+ Customers**
- **Aerospace, Defense, Government, Academic, Life Sciences, Finance, Electronics**

Target Customer Profiles

- **Users migrating to Linux Clusters from proprietary SMP machine environments.**
- **Cluster Users becoming Cluster Owners - users who may have shared time on public/private large clusters but now have budget/access their own clusters.**
- **OEM providers creating turn key clusters as appliances.**

Government, Research and Defense

AURA 50 1957•2007

LIGO

GATS, Inc.
Atmospheric Science

DISA
Defense Information Systems Agency
Department of Defense

Computational Infrastructure for Geodynamics

Federal Aviation Administration

BOEING

BAE SYSTEMS

CACI
EVER VIGILANT

L3 communications

ITT

FTT

NASA

GENERAL DYNAMICS
Strength On Your Side®

LOCKHEED MARTIN
We never forget who we're working for®

GENERAL ATOMICS

SAIC
From Science to Solutions

Jet Propulsion Laboratory
California Institute of Technology

MITRE

HR Textron
TEXTRON Systems

CSC
EXPERIENCE. RESULTS.

APPLIED RESEARCH ASSOCIATES, INC.

GE Aviation

NOAA

CATERPILLAR

Battelle
The Business of Innovation

Brookhaven National Laboratory
Office of Science / U.S. Dept. of Energy

INL
Idaho National Laboratory

Homeland Security

CDID Experimentation Division
Battle Command Battle Laboratory
Gordon Facility

Sandia National Laboratories

Argonne NATIONAL LABORATORY

NAVSEA
NAVAL SEA SYSTEMS COMMAND

NAVAIR
NAVAL AIR SYSTEMS COMMAND

www.penguin

Life Sciences

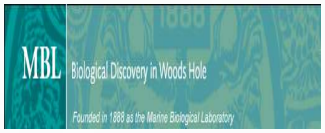

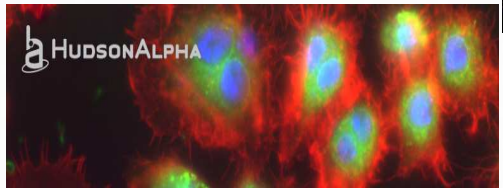

Partners

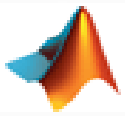

MATLAB®
Product Family

Beowulf Democratizes Supercomputing

- **Project conceived by Becker and Sterling in '93 and initiated at NASA in '94**
- **Objective: show that commodity clusters could solve some of the easier problems usually handled by $million supercomputers but at a fraction of the cost**
- **Build a system that scaled in all dimensions**
  - » Networking, bandwidth, disks, main memory, processing power
- **Initial prototype**
  - » 16 processors, Channel-bonded Ethernet, under $50K
  - » Matched performance of contemporary $1M machine
- **Idea spread quickly through NASA, research, academic communities**

HPC at a fraction of traditional cost

Early Beowulf Cluster

Lessons Learned: Original Beowulf

What did we learn?

- Complexity in Cluster Stack
- Extensive training to install, configure and use
- Long-term administration and updates were difficult
- Only “static scalability”

Linux Clustering Ecosystem Today

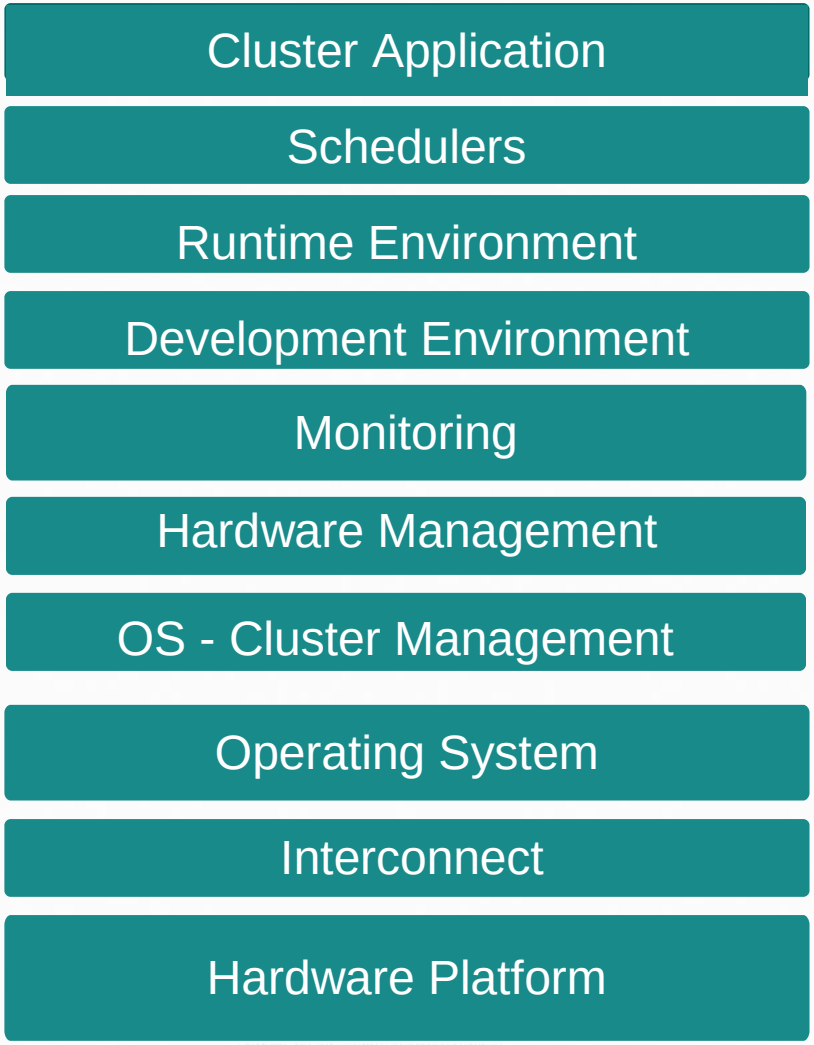

Traditional Linux Cluster Management

Manual management of multiple, individual servers

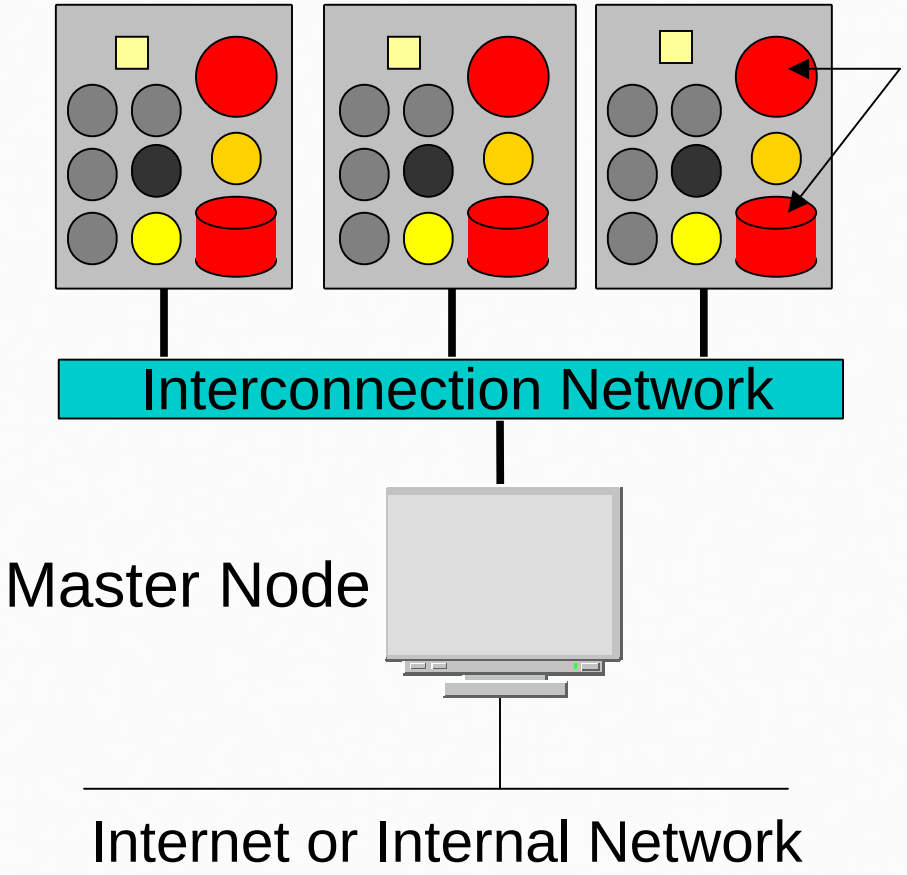

- Full Linux install to Local disk
  - Very Manual & Slow; 5-30 minutes
- Full set of disparate daemons, services, user/password, & host access setup
- Crude parallel shell with complex glue scripts for administration
- Monitoring & management mapped to inefficient underlying glue to N nodes

Scyld Software - Platform

Scyld ClusterWare™

Highly SCYLD

The Scyld Platform

Fully Integrated Software Solution for Linux Clusters

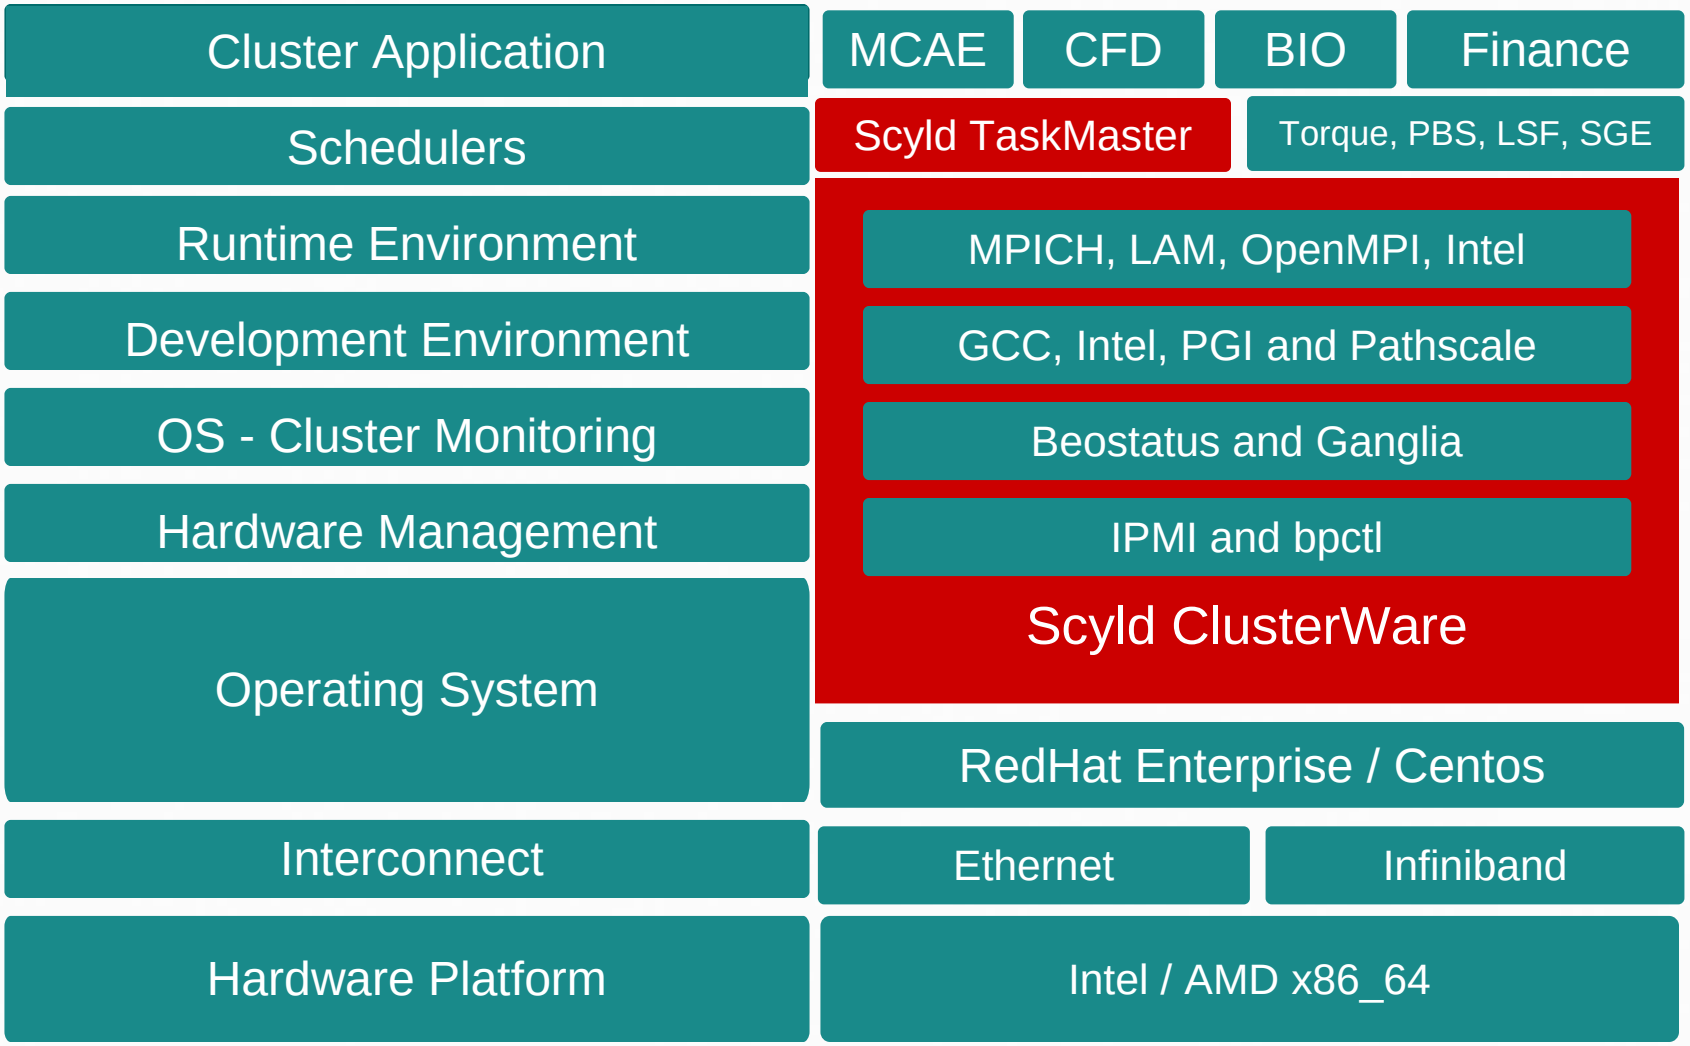

Scyld Overview

- **Scyld ClusterWare provides a single point for cluster installation, administration, security and monitoring.**
- **The unique architecture of Scyld is based on three fundamental concepts:**
  - » *The operating environment deployed to the compute nodes is provisioned "stateless", directly to memory.*
  - » *The compute node environment is lightweight, stripped of unnecessary software, overhead and vulnerabilities.*
  - » *A simple Linux extension managers the cluster into a pool of processors operating like a single virtual machine*

Scyld ClusterWare Architecture

Manage the Cluster as a single system

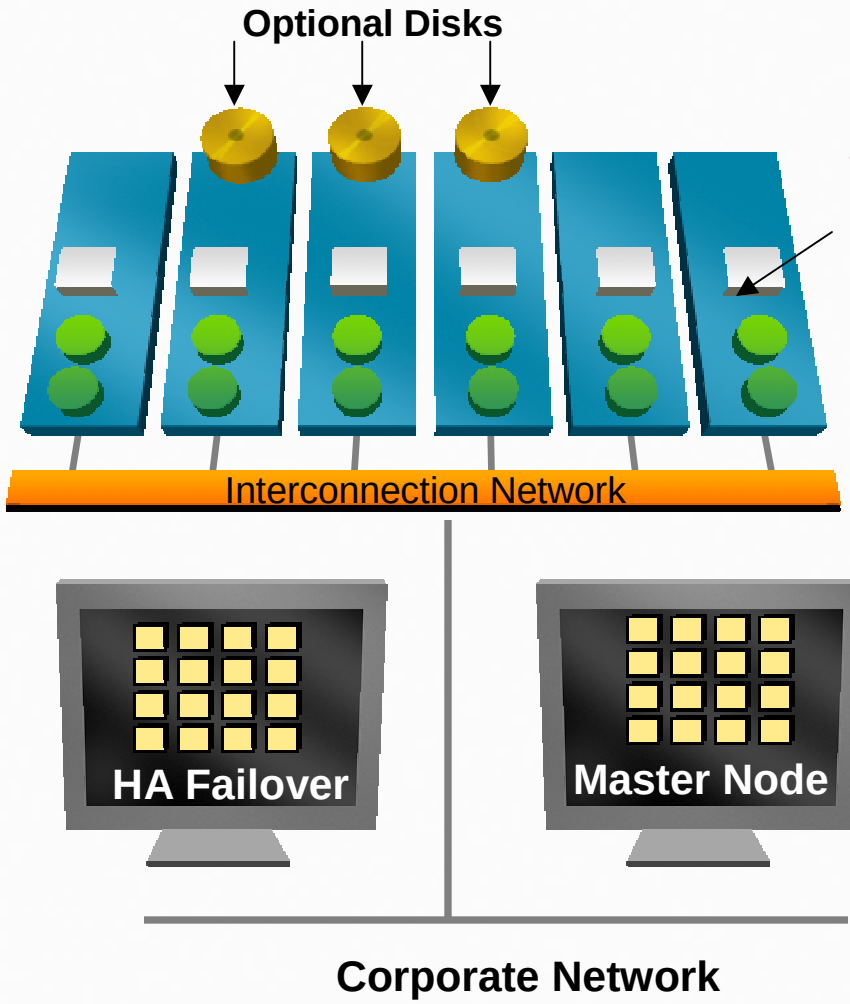

- Minimal in-memory OS with single daemon rapidly deployed in seconds – optional local disks
- Unified Process Space with Single Sign-On
- Seamless Integration with site wide Directory Services, such as, LDAP, ADS, NIS
- Monitor & Manage efficiently from the Master
- High Availability - Active/Passive or Active/Active

Scyld ClusterWare - Provisioning

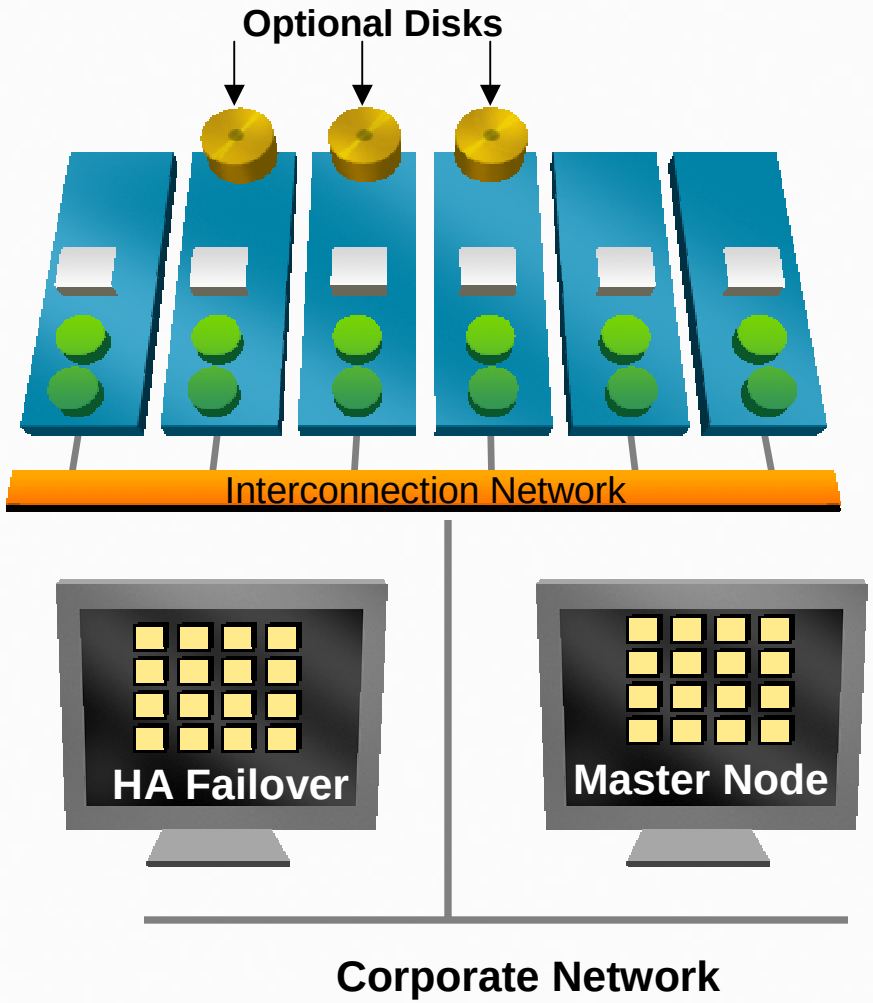

Traditional full standard installation on Masters

- Dynamically built "distribution" on other nodes
- Core system is "stateless" non-persistent
- Provided by the boot server as ramdisks
- No local storage or network file system used
- Additional elements provided under centralized master control

Scyld ClusterWare - Consistency

- **Traditional Approach: Replicated local full OS Installs**
  - » Time consuming
  - » Error prone

- **Scyld ClusterWare – Consistency by Design**
  - » Remote execution produces same results as local execution
    - Same executable (including version!)
    - Same libraries, including library linkage order
    - Same parameters and environment

Scyld ClusterWare - Monitoring

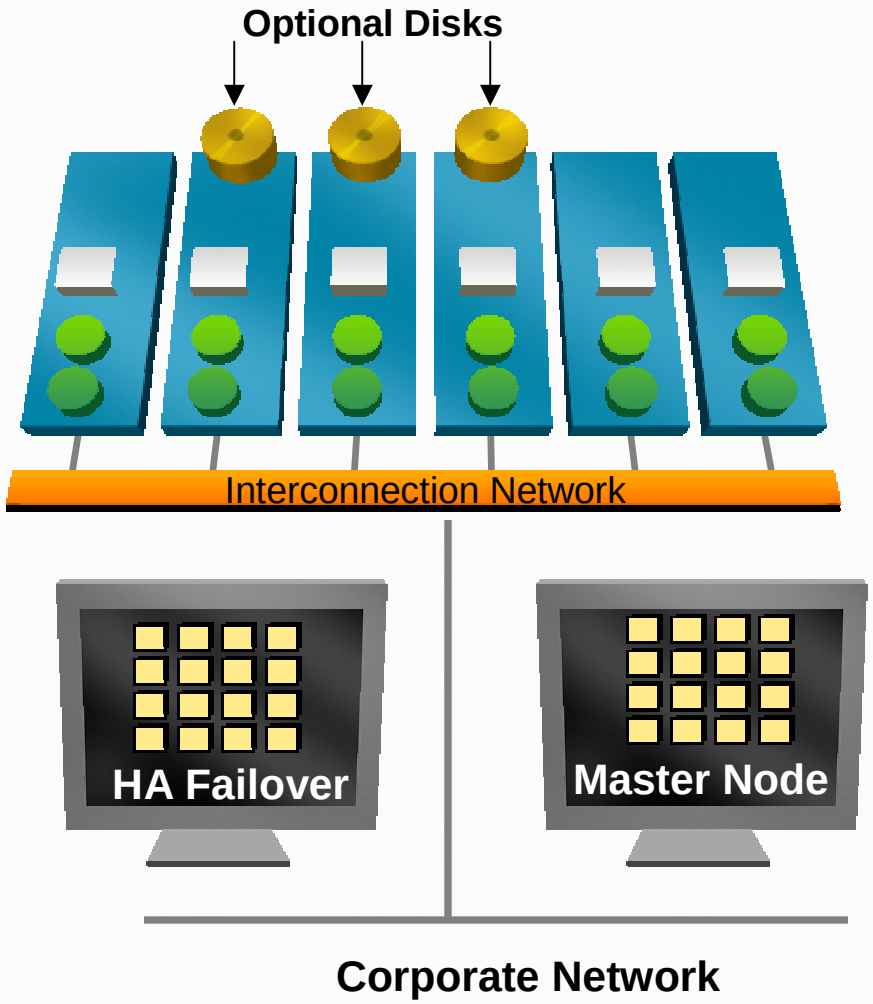

- **Single Point of Monitoring**
  - BeoStat
  - Unified Process Space
  - Ganglia

Scyld ClusterWare - Monitoring

Basic Monitoring

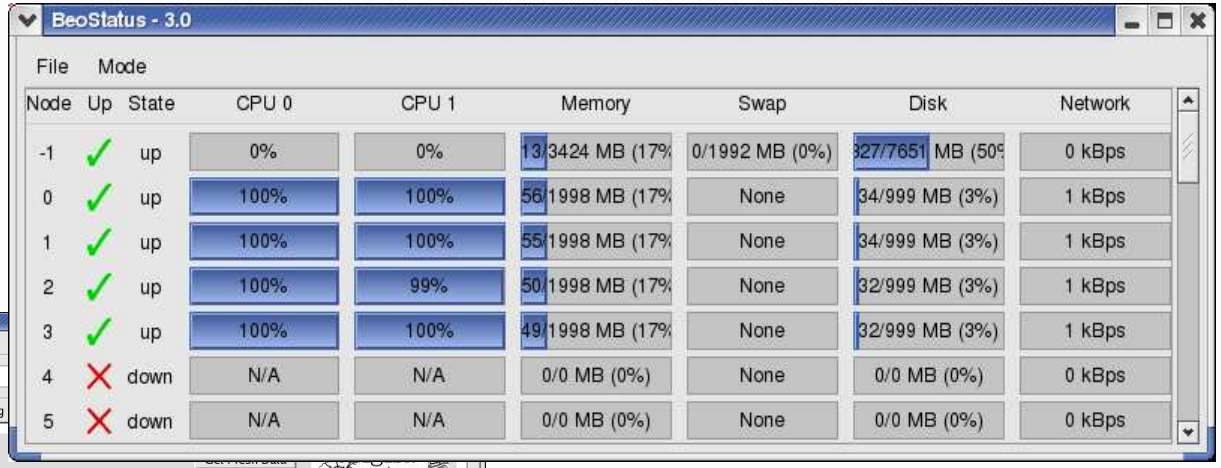

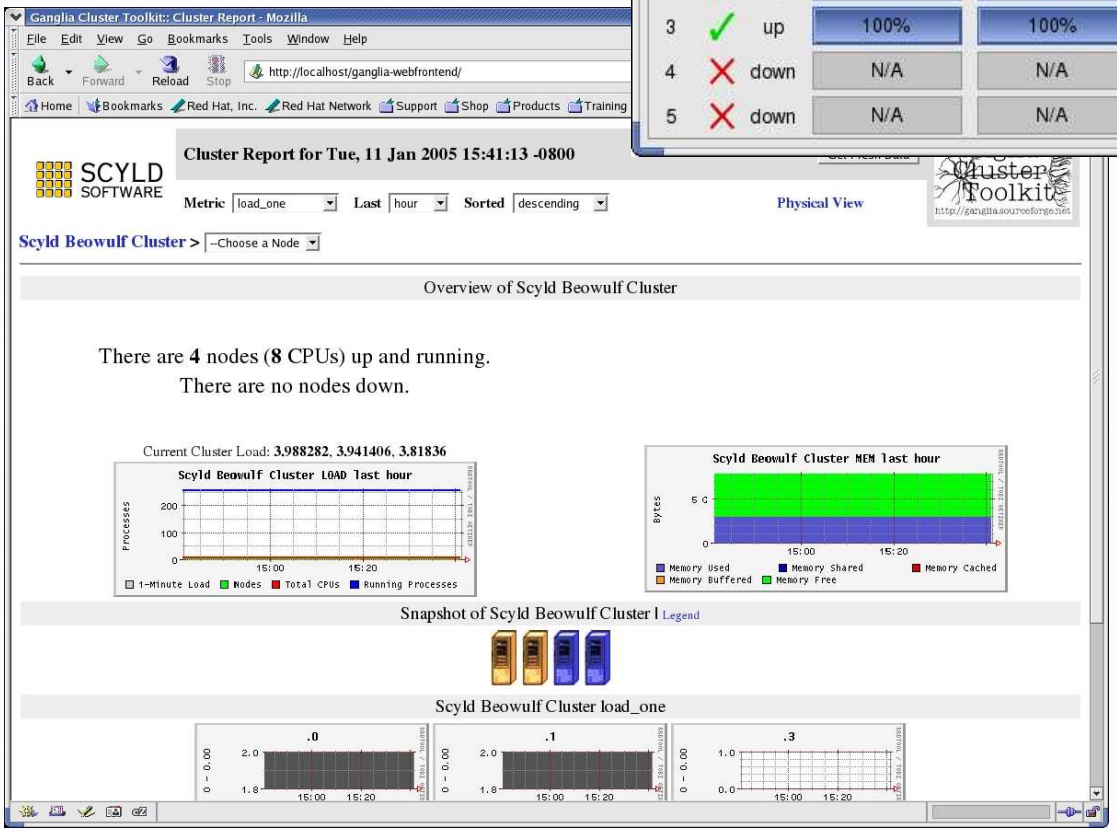

Generally “status” only:
Hardware focused

Enhanced Monitoring
with Ganglia

Scyld Software

Scyld TaskMaster™

Highly SCYLD

The Scyld Platform

Fully Integrated Software Solution for Linux Clusters

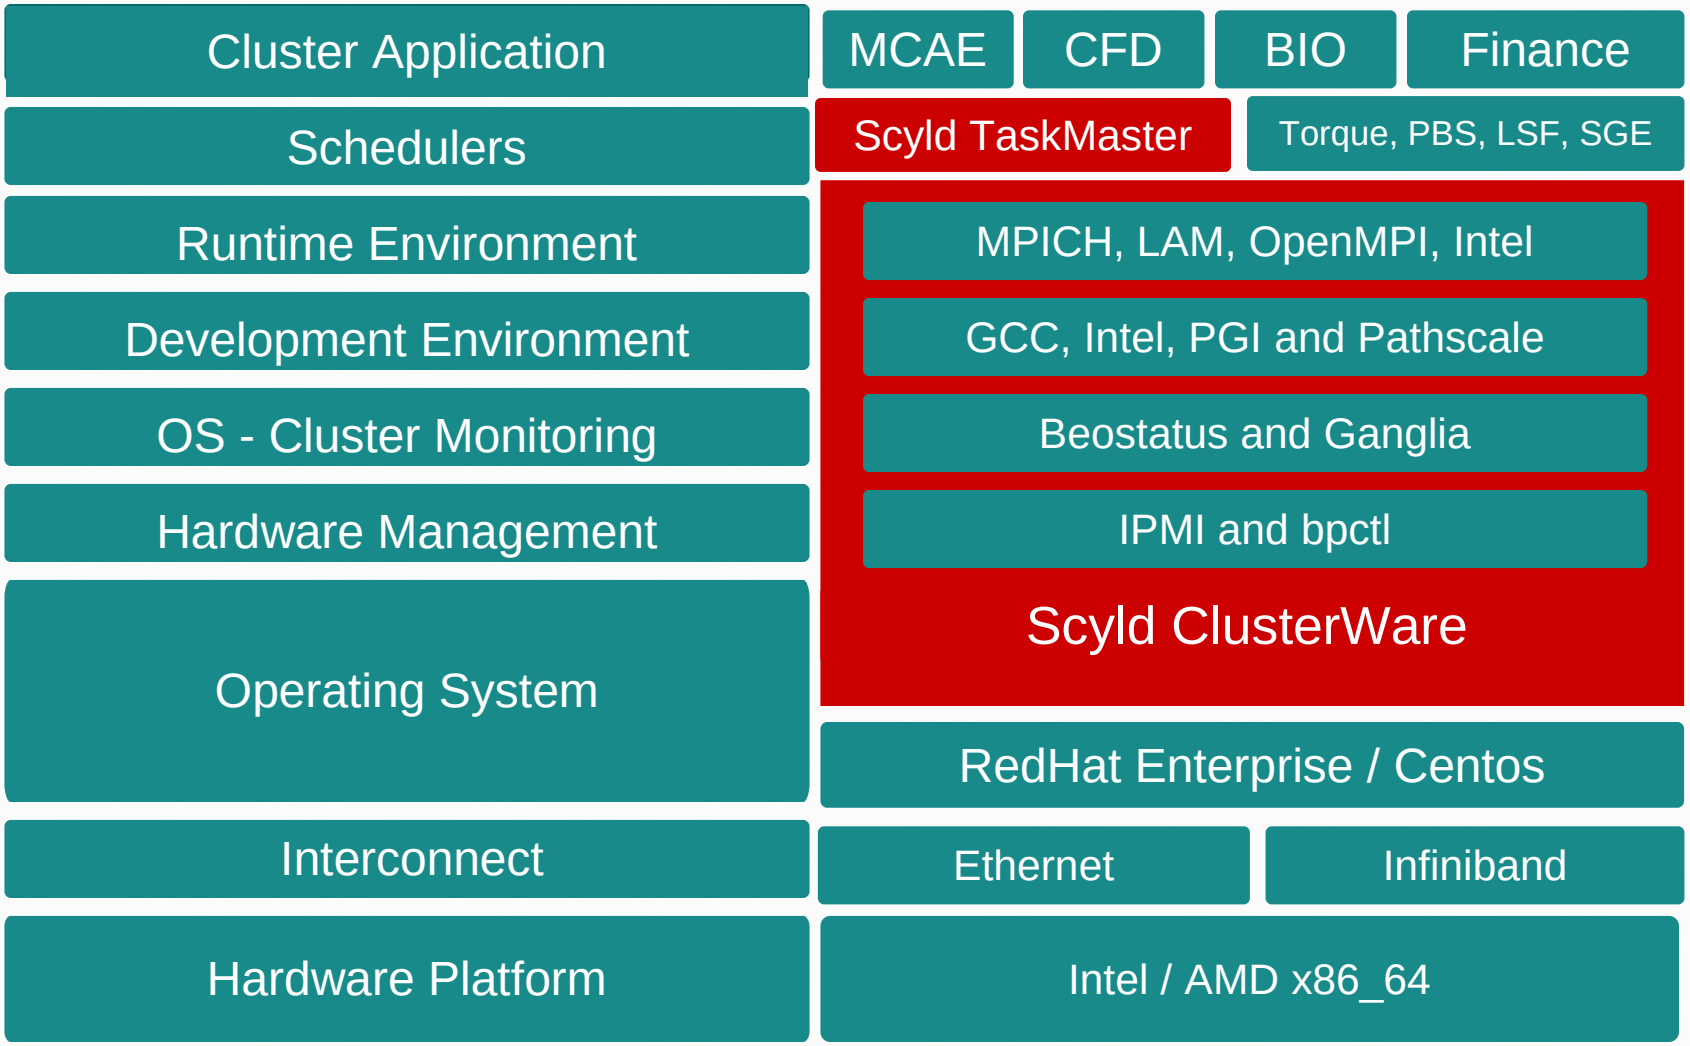

Scyld TaskMaster Suite

Scyld TaskMaster Features	Customer Benefits
✓Intuitive GUI- local/Web portal	▪ Very intuitive UI maximizes user productivity
✓Policy-based job scheduler & event engine	▪ Increase cluster utilization to 90%-99%
✓User Reservations/Triggers	▪ Prioritized workloads ensure right resources are available at right time
✓Graphical reporting- live & historical	▪ Graphical reporting helps plan current & future workloads
✓Built-in accounting for shared usage	

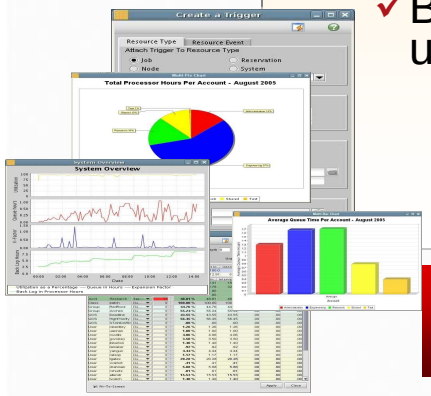

Flexible & cost effective workload management solution

Scyld TaskMaster Workload Management

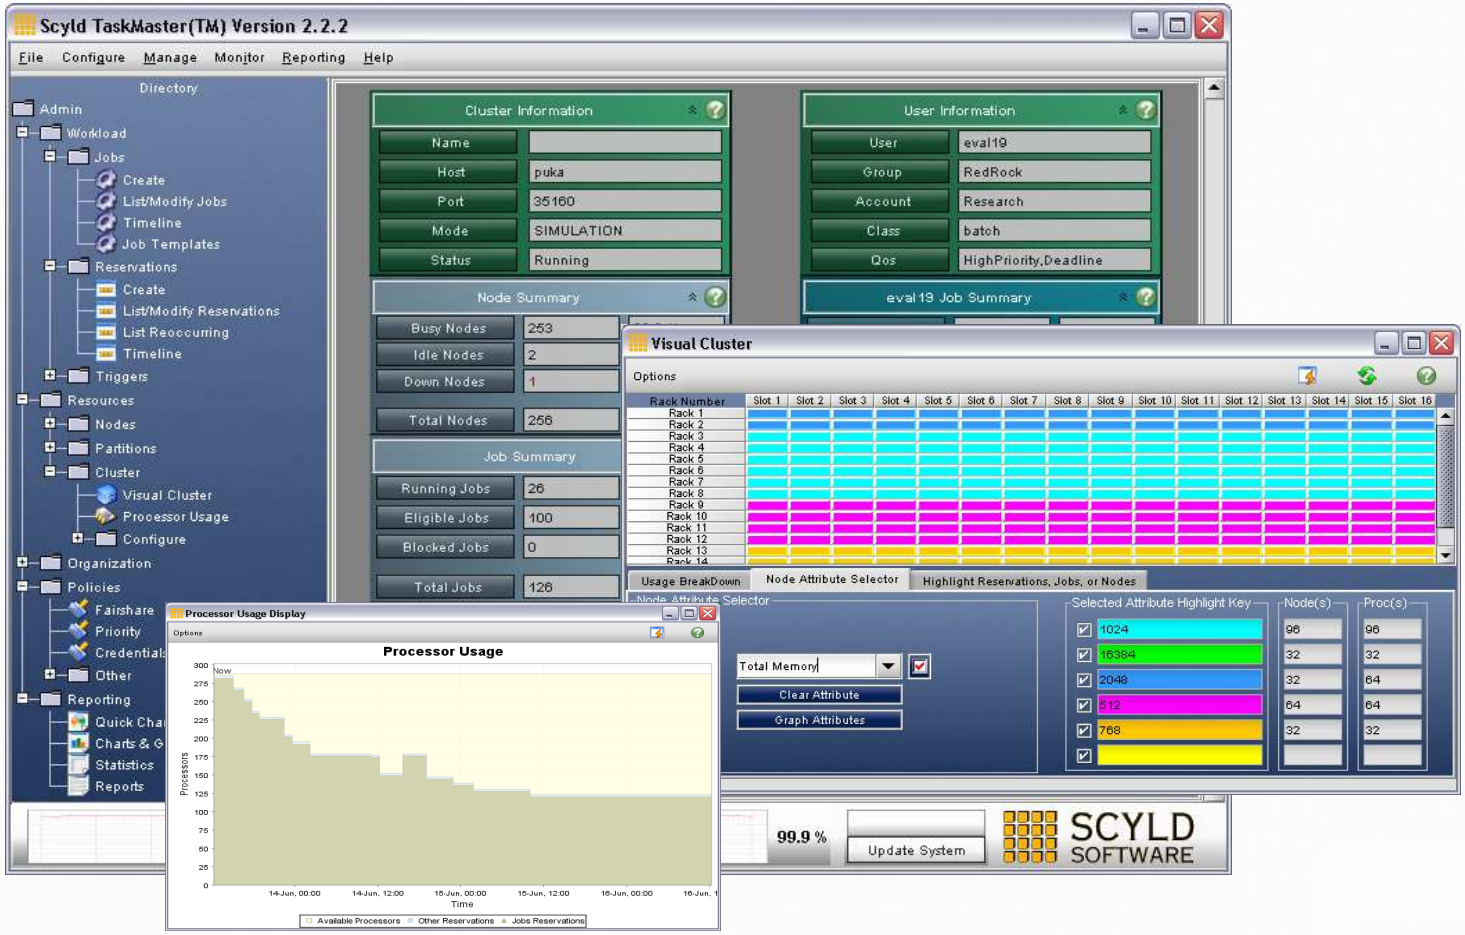

Maximizing cluster utilization - mapped to business priorities

Demonstration
Scyld ClusterWare

Highly SCYLD

Case Study - Applied Biosystems

PENGUIN COMPUTING

SOLiD™

Applied Biosystems

The Next Generation Has Arrived
SOLiD™ System

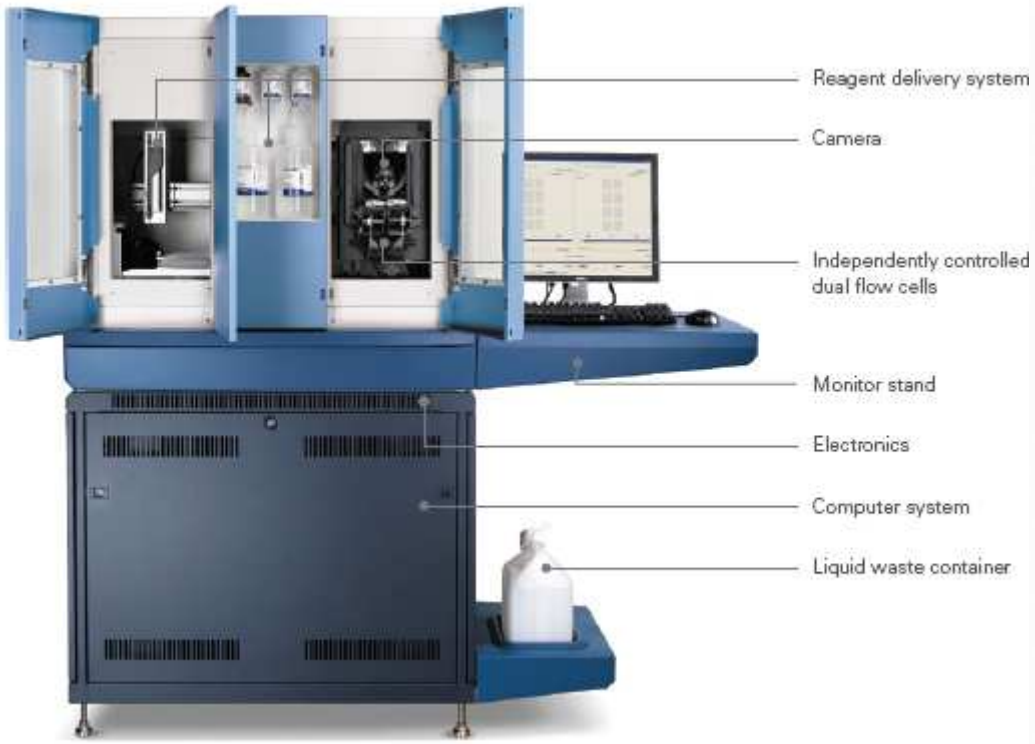

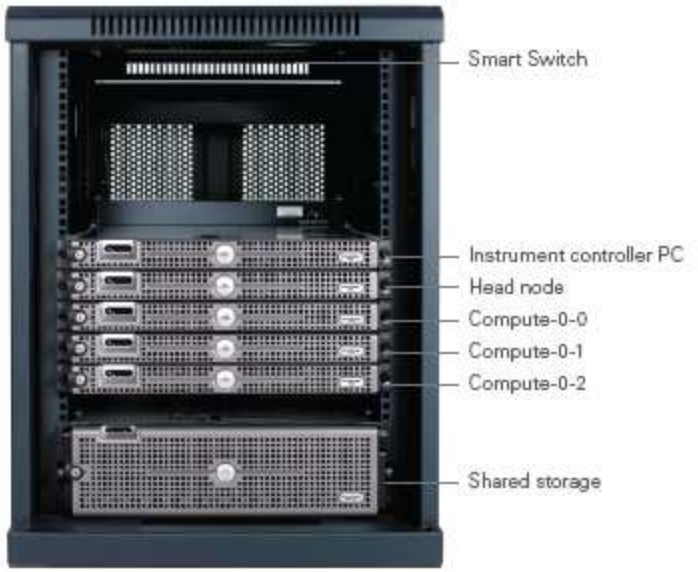

com

PENGUIN COMPUTING

Question & Answers

Highly SCYLD

Backup

Highly SCYLD

How we measure up: HPC

Feature/ Competitor	Scyld ClusterWare	IBM	HP (XC)	Open Source Rocks/Oscar/Warewulf	Linux Networx
Fast, Master- Only Install	+ +	-	-	-	-
Fast In Memory Provisioning	+ +	+	-	- (warewulf+)	-
Virtual Single System Arch.	+ +	-	-	-	-
Lightweight, Secure OS- JIT Provisioning	+ +	-	-	-	-
SW Consistency Model	++	-	-	-	-
Commercial Support	++	IBM++	++	-	+
Cost	$$	IBM: $$	$$	Free(Platform Rocks: $)	$
Custom Hardware for full capability	No	Yes	Yes	+	Yes

Cluster v Grid

- **Basic differences**
  - » Clusters generally defined as:
    - Dedicated purpose machines
    - Homogeneous environment (i.e. same OS)
    - LAN-based
  - » Grids generally defined as:
    - Multi-purpose machines (user WS by day; grid node by night)
    - Geographically disparate

Cluster v Grid

- **Other Grid differences**
  - » Typically already statically provisioned
    - Windows, Unix, Linux with infrastructure to operate as a Grid node
    - Heterogeneous environment
    - (Active Directory- Windows, Globus Linux/cross)
    - Or dual-bootable machines in some cases
  - » Can interface to dedicated clusters as pools
  - » Inter-org/facility challenges enormous
  - » Ambitious consortium efforts to address these challenges
    - Globus Alliance working on elaborate “toolkit”.

Cluster Evolution