# Are Blade Servers Right For HEP ?

**Rochelle Lauer**
**Yale University Physics Department**
**rochelle.lauer@yale.edu**

## Outline

- Blade Server Evaluation – Why and How
- The HP BL Blade Servers
  - The Blade Servers
  - The Infrastructure
  - The Software
- My Questions
- Answers At the "Solutions Center"
- Our Plans
- Your Comments and Questions

## Blade Server Evaluation – Why

Our Linux Environment

- ∼15 (And Growing) Linux Desktops(EVOs)
- Proliant Four Way Running Linux (for CDF)
  - Compile and Compute
  - Some NFS Serving
  - NIS Master
- SANs with ∼4 TerabytesServed By A TruCluster

---

⇒ Need For A Group Independent Server
⇒ Need For Easily Expandable Linux Batch Farm
⇒ (Need To Replace Our Terminal Server)

## Blade Server Evaluation – Why

Our Problems

- Linux Management
- Power and Cooling
- Growth Without a Plan

$\Rightarrow$ Are Blade Servers The Answer ?

---

To Find Out:

1. Do Some Research
2. Plan A Trip To The Solutions Center
3. (Perform Benchmarks)

## The Research : What Is A Blade Server?

### From ZDnet

‘‘... systems-on-card offer key advantages over traditional servers’’

‘‘... savings in power,space,ease of operation and scalability ...’’

| Before Merger<br>After Merger | HP and Compaq Blade Servers<br>? |
|---|---|

$\Uparrow$ **It Doesn't Matter**

---

$\Rightarrow$ **An Example : The Proliant BL Line**

## The Research : What Is An HP BL Blade Server ?

Basic Features

- Power Efficient, Ultra Dense
- Hot Pluggable
- Multiple NICs
- Supports Windows 2000 Server, Windows .NET Server, RedHat and SUSE Linux

$\Rightarrow$ Hardware To Support Remote Management

## The Research

**Highlights**

| BL10e | BL20P |
|---|---|
| 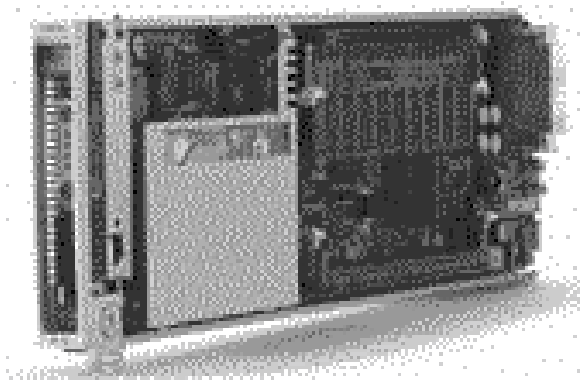 |  |
| 1 Processor<br>2 Network Connections<br>800Mhz PIII<br>No Slots<br>1U X 3U | up to 2 Processors<br>4 Network Connections<br>1.4GHz PIII<br>No Slots<br>1U X 6U |

- Cost : Comparable to a Server, $>$ Than A Desktop

$\Rightarrow$ Nothing Special (?)

$\Rightarrow$ Blade Server $+$ Infrastructure $=$ Something Special

## The Research : Blade Server Infrastructure

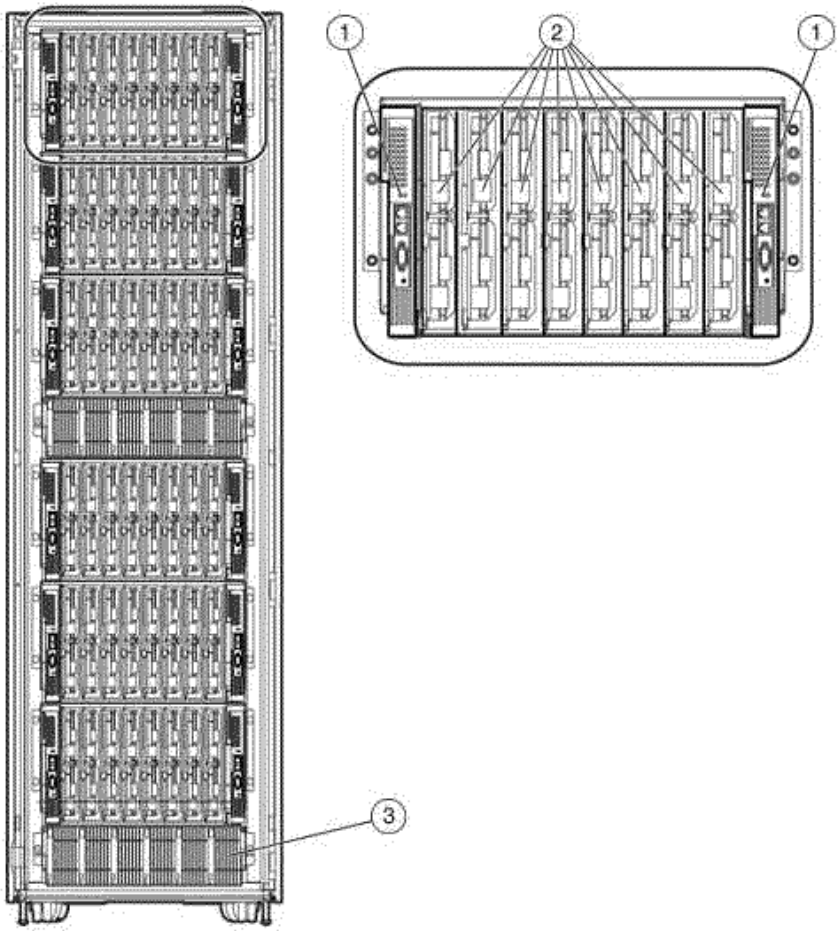


1. Blade Interconnects
2. Server Blades
3. Power Enclosure, Hot Plug Power Supplies

## The Research : Blade Server + Infrastructure

**A Scalable Solution**

- Can Add Processors to the Initial Enclosure(s)
- Can Add Enclosures to the Rack
- Can Upgrade Network Connections

**Provides For Technology Integration**

- Can Mix and Match Old and New Technology in the Same Rack

$\Rightarrow$ Can Mix and Match Old and New Technology in the Same Enclosure

**Simplifies The Environment**

- Small Footprint For a Large Number of Servers
- Consolidated Power, Cooling and Network
- Minimal Power and Cooling Requirements

**Provides Remote Management**

- At The Hardware Level
- At The Software Level

## Questions

⇒ **Will A Stack Of Desktops Or Servers Work As Well ?**

**Hardware Questions**

- **Why Are The Processors More Costly ?**
- **Are The Processors Limited to Stable Technology?**
- **Why Do I Need An Enclosure ?**
  **Will a Rack Work As Well ?**
- **Why Do I Need A Network Interconnect ?**

## Software Questions

- Why Do I Need Hardware Management Software?
  (Remote Power On/Off, Remote Console)
- Why Do I Need Deployment Software ?
  - The Tools To Deploy Are Freely Available
    (e.g. PXE, Kickstart), Why Not Create My Own ?
  - We Have Created Our Own
    A 2001 HEPiX Talk: `www.slac.stanford.edu/ alfw/talks/pxe.pdf`

---

⇒ Does It Work ?

⇒ Is It Too GUI and Windows Oriented ?

## What I Learned At The Solutions Center (And At HPETS2002)

At The Solutions Center

- Two Hour Demo
- 5 Hours To Work By Myself (With Help)
- One Hour Concall with the Remote Deployment Software Developer

---

Highlights of What I Learned

- The Hardware Is Designed For Upgradeability
- The Software Will Soon Be Linux Based
- The Software Actually Does What I Need It to Do
- There Are Options For Stand Alone Maintenance of The Blades
- It Really Is Plug and Play

$\Rightarrow$ My Opinion : Well Thought Out

## What I Learned : Some Specifics
## The Processors

More Costly ?
Remember : More Than A Desktop; Comparable To A Server

- Management ROM Provides Value
- Pre-Boot Execution Environment (PXE) Enabled NIC Provides Value
- Multiple NICs Provide Bandwidth and Redundancy
- Dense Configuration With Low Cooling and Power Requirements
  - Provides Small Footprint
  - Allows For Power and Cooling Consolidation

Stable Technology ?

- Isn't That What You Want ?
- Dense Configuration With Low Cooling and Power Requirements
- Provides A More Stable Software Environment
  - More Likely To Work

$\Rightarrow$ Very Close to What You Might See on The Desktop

## What I Learned : Some Specifics

### The Enclosure

⇒ **Plug and Play Backplane Provides For Hot Swappable Modules**

- **Backplane Consolidates Network Connections**
  - **Contributes To Plug and Play Capability**
- **Connected to a Server Blade Management Module**
  - **Polls Blades for thermal,power, etc. events**

## What I Learned : Some Specifics
## The Network Interconnects

- Provides Redundancy
- Option For Gbit on the Blades
  - Plans for Future Blade Interconnects
- Option for Switch to Consolidate Uplinks
- Necessary For Plug and Play

---

The Total Hardware Package

- Small Footprint, Highly Reliable, Manageable, Scalable Set Of Servers
- At What Price ?
  $\Rightarrow$ About $10,000 US

## What I Learned : Some Software Specifics

HP's iLO

⇒ Remote Power Off and On

- Failure Reporting
- Remote Access To Consoles

⇒ Managed as One System Via A Web Interface

---

⇒ Reduces Staff Requirements

⇒ Provides A Consistent Methodology

## What I Learned : Remote Deployment Software

**Highlights**

- Organizes The Available Tools (PXE, Kickstart, etc.)
- Provides a Consistent Methodology Through The GUI
- Provides For Disk Imaging and Restore
- Provides For Parallel Boot
  - Allows For Multicast To Reduce Network Bandwidth Requirements
- Provides Scripting Capabilities To Tailor The Environment

## What I Learned : Remote Deployment Software

- Provides A Consistent Interface
  - Pulls All The Pieces Together – Provides The Glue
- Allows For Management As A Single System
  Or Groups Of Systems
  $\Rightarrow$ Maintains Software Consistency
- Maintains the Bookkeeping
- Provides High Availability

## Our Requirements and Plans

- Provide For Batch Processing (Many Linux Nodes)
  - On Short Notice
- Remove Linux Server Functions From The Proliant
- (Eventually) Move Server Functions From Tru64 To Linux
- Find a Replacement (Hardware or Functionality)
  For Our NT Terminal Server

## Our General Methodology

In the Past : Purchase Infrastructure

- A Shell To Provide For Future Growth

$\Rightarrow$ Right Now : A Hard Sell

## What Will We Do ?

Hardware Versus Staff Costs – Hardware Wins

Likely Scenario

- Purchase Rack + 1 Enclosure + 1 Dual Processor
- Decide On Type Of Network Interconnect And Include In Initial Purchase

## Possible Show Stoppers

- For Us : Software Licensing Costs

  $\Rightarrow$ Not A Problem

## Conclusions

- Blade Servers Look Like A Good Solution For Our Group
- Blade Servers Look Like A Good Solution for Other Farms

$\Rightarrow$ Why Don't I See Other Groups Using This Technology ?

# Thank You

# Comments and Questions