

# Optimizing the Enterprise: Using JBoss, Red Hat and Virtualization Technologies Within a SOA

Isaac Christoffersen  
Booz Allen Hamilton  
Associate

Dan Santillo  
Booz Allen Hamilton  
Associate

14 February 2008

# The demand for computing, network and storage resources shows no signs of decreasing...

What are your current technical challenges?

- Approaching power/AC/space limitations
- Handling increasing data and growing user base
- Satisfying stability, availability, and agility requirements
- Updating antiquated technology solutions
- Paying unreasonable maintenance costs, which tie up budgets
- IT professionals / expertise unavailable

“ In 2006, humans created 161 exabytes of data, 3 million times the information in all books ever written. ”

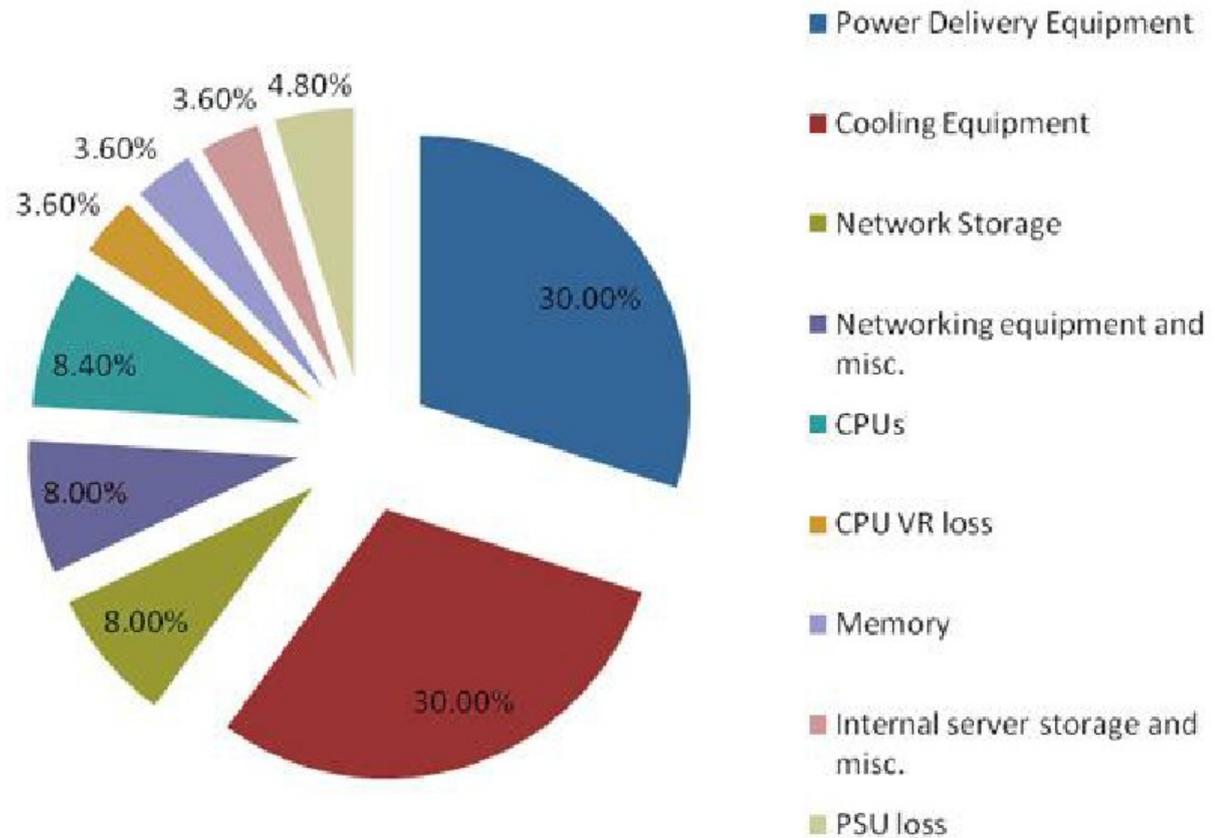
- IDC US Green IT Survey

# Adding resources to meet these challenges will only increase power consumption demands...

What is consuming the most power in your environment?

“ By 2008, half of data centers will have insufficient power and cooling. By 2010, half of all data centers will have to move or outsource.”

- Gartner Symposium,  
Orlando, October 2007



source: "Improving Data Center Energy Efficiency: A Holistic Approach", Andrew Kurz, Burton Group (September 2007)

# Several fallacies and facts have steered us towards our current path...

How did we get to this point?

“The significant problems we face cannot be solved by the same level of thinking that created them.”  
- Albert Einstein

## Fallacies

- The network is reliable
- Latency is zero
- Bandwidth is infinite
- The network is secure
- Topology doesn't change
- There is one administrator
- Transport cost is zero
- The network is homogenous

*Peter Deutsch's 8 Fallacies of Distributed Computing*

## Facts

- Individual components are not expensive enough... yet:
  - Compute Power, Storage, Energy
- Hardware used to compensate for poor software design
- Demand for Information Technology is still growing
- Effectiveness  $\neq$  Efficiency

*source: "Improving Data Center Energy Efficiency: A Holistic Approach", Andrew Kurz, Burton Group (September 2007)*

# There are many useful architectural and system design patterns to help in overcoming these challenges...

What are the “silver-bullet” answers that will solve these problems?

*(Hint: Trick Question)*

“ There is no single development, in either technology or in management technique, that by itself promises even one order-of-magnitude improvement in productivity, in reliability, in simplicity.”

- Frederick Brooks, *No Silver Bullet: Essence and Accidents of Software Engineering*

- Architecture Layers
- Virtualization
  - Server
  - Network
  - Storage
- Shared Storage
- Cycle Harvesting
- Repeatable Build
- Clustering
  - Operating System
  - Application
  - File System

# A layered view of an enterprise system shows complex heterogeneous entities...

How do you decouple the system components in your Enterprise Architecture?

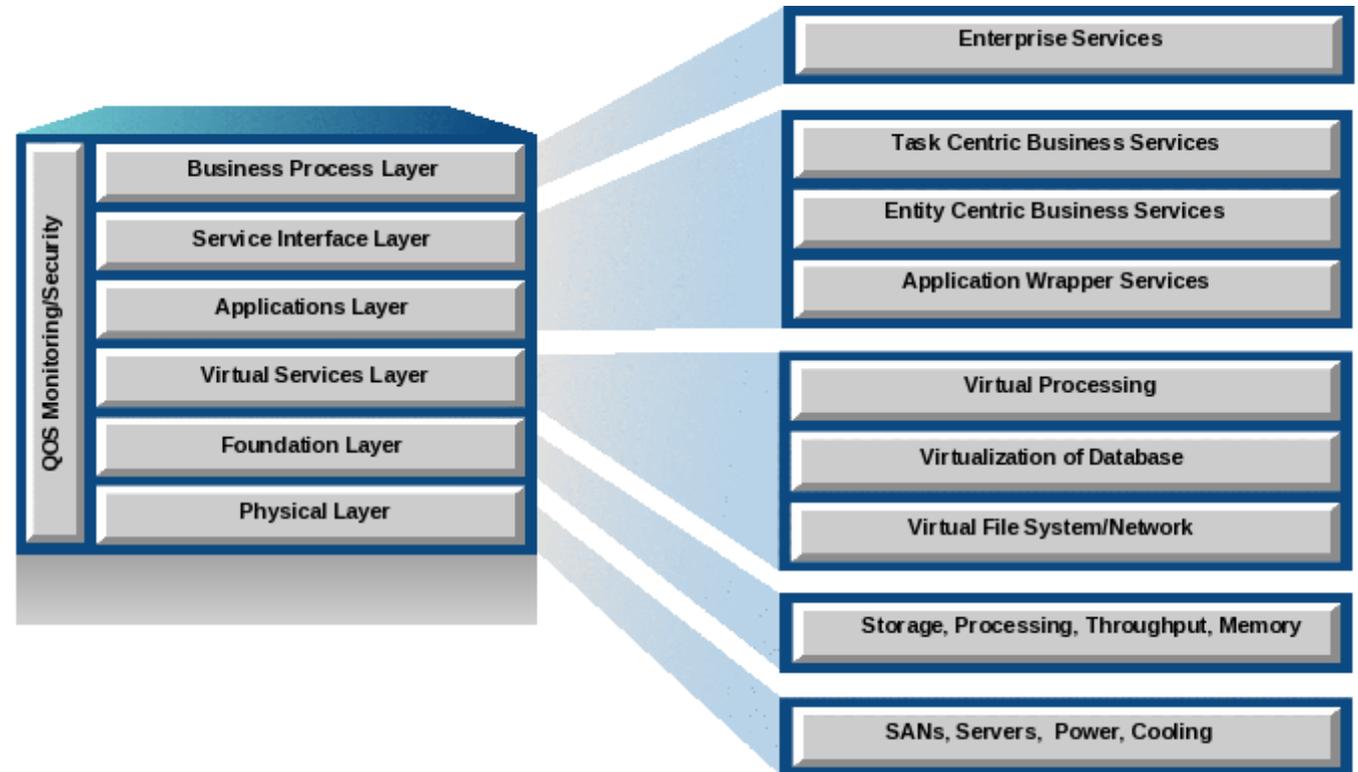
- What are the components that make up the whole system?
- How do these components interact with each other?
- Are the components loosely coupled?
- How are the quality attributes of modifiability, scalability, and integrability supported?

“ Decoupling the components in a vertical manner is crucial in order to support modifiability, portability, and reusability. ”

- Paris Avgeriou & Uwe Zdun, *Architectural Patterns Revisited – A Pattern Language*

# ...decomposed into their simpler, individual components and responsibilities...

How do you decouple the system components in your Enterprise Architecture?

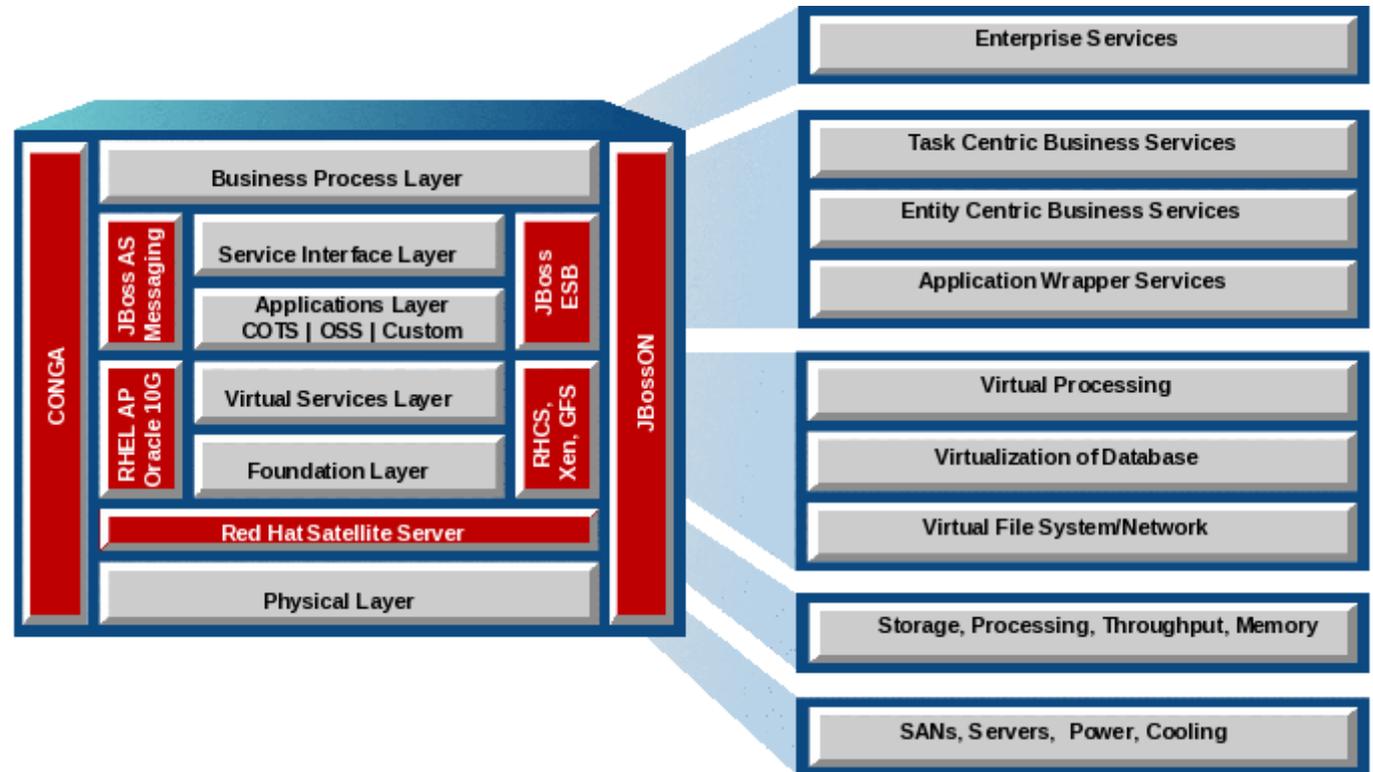


“ First and foremost, a loosely coupled architecture allows you to replace components, or change components, without having to make reflective changes to other components in the architecture/systems.”

- David Linthicum, *SOA Magazine*, October 2007

# ...decomposed into their simpler, individual components and responsibilities...

How do you decouple the system components in your Enterprise Architecture?



“ First and foremost, a loosely coupled architecture allows you to replace components, or change components, without having to make reflective changes to other components in the architecture/systems.”

- David Linthicum, *SOA Magazine*, October 2007

# Server Virtualization allows multiple operating systems to run on one physical server...

Why should you add virtualization to your enterprise?

Virtualization offers benefits such as consistent deployment environments, dynamic provisioning, real-time migration, high availability and load balancing.

## Reasons to Virtualize Servers

|   |     |
|---|-----|
| Cut costs via server consolidation                      | 81% |
| Improve disaster recovery and backup plans              | 63% |
| Provision computing resources to end users more quickly | 55% |
| Offer more flexibility to the business                  | 53% |
| Provide competitive advantage                           | 13% |

*(Respondents chose up to three)*  
*SOURCE: CIO Research*

# Storage Virtualization abstracts logical storage from physical storage...

## How do we reclaim storage space?

---

With storage needs growing faster than computing power needs, consolidated storage reduces the number of individual storage devices and reclaims the underutilized capacity.

- Benefits
  - Abstracts the physical location of data
  - Improves utilization
  - Increases scalability
- Risks
  - Some performance overhead

# Shared Storage allows multiple systems/services to utilize the same storage resources...

## How do we maximize our storage resources?

---

With tera-, exa-, and peta-bytes of data being processed, pooling resources and reducing data duplication frees up storage space, decreases network load, and increases scaling flexibility.

- Benefits
  - Maximize storage resources and minimize costs
  - Reduce data duplication
  - Scale clusters on the fly
  - Common file system
- Risks
  - Can become a single point of failure if incorrectly implemented

# Cycle Harvesting fully utilizes the available CPU resources of the members in the grid...

How do we improve CPU usage?

---

Cycle harvesting makes the best use of the CPU resources available at any given time.

- With Cycle Harvesting, we can:
  - Perform the same work using processing resources from an available pool.
  - Perform more work with an improved utilization.
- Utilizing intelligent management software, idle resources can be allocated as needed.

# Consistent, repeatable builds encourage fast deployment and consistently identical systems...

How do we efficiently deploy solutions?

---

Fast, repeatable builds provide efficient deployment of new systems and expansion of existing systems.

- Utilizing automated and scripted installs, we can:
  - Create new systems quickly with little manual interaction.
  - Have a setup for repeating the build process in an efficient and stable manner.
- And these builds can:
  - Use the most current version of software
  - Configure the desired server functionality (i.e. Web, SMTP, DHCP server)
  - Execute all relevant tests **successfully.**

# Clustering provides redundancy and enhanced communication between services...

How do we maximize uptime?

---

Utilizing clustering provides fail-over, load-balancing, and up-scaling benefits to the enterprise.

- Operating System Clustering
  - Application/Service fail-over
- Application Server Clustering
  - Fail-over
  - Load-balancing
  - Distributed Deployment
- File System Clustering
  - Maximize storage resources
  - Minimize Costs
  - Reduce data duplication
  - Scale clusters on the fly
  - Common File System

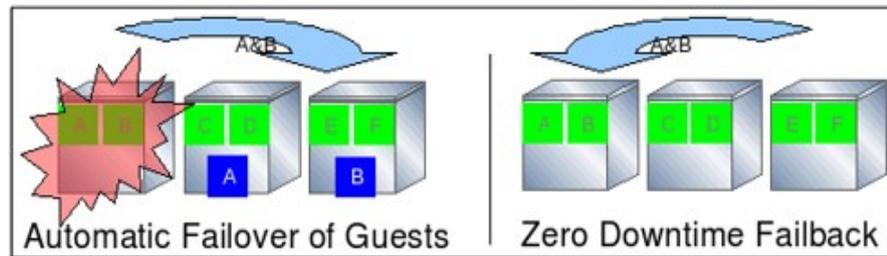
# Leveraging these patterns can be done today using widely-available open source offerings...

- Operating System Clustering 
- Virtualization
  - Server 
  - Network
- Application Server Clustering 
- Storage
  - Virtualization 
  - Shared
  - File System Clustering
  - File Replication
- Red Hat Cluster Management
- Xen Virtualization
- JBoss
- Global File System (GFS)
- Globus Toolkit

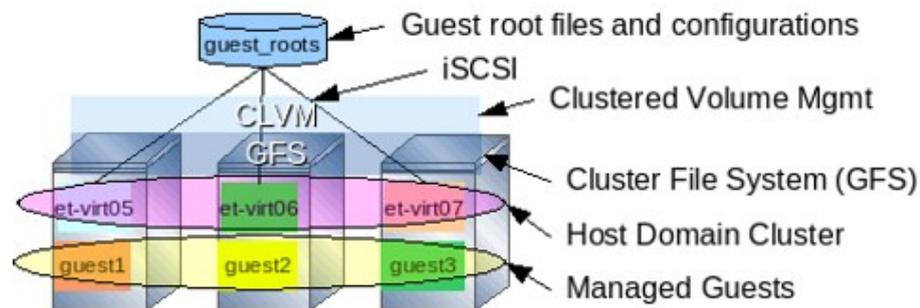
# Virtualization can reduce cost while also improving flexibility...

What new challenges will server consolidation introduce?

- Cluster multiple physical servers to gain failover capabilities



- Share resources across physical servers

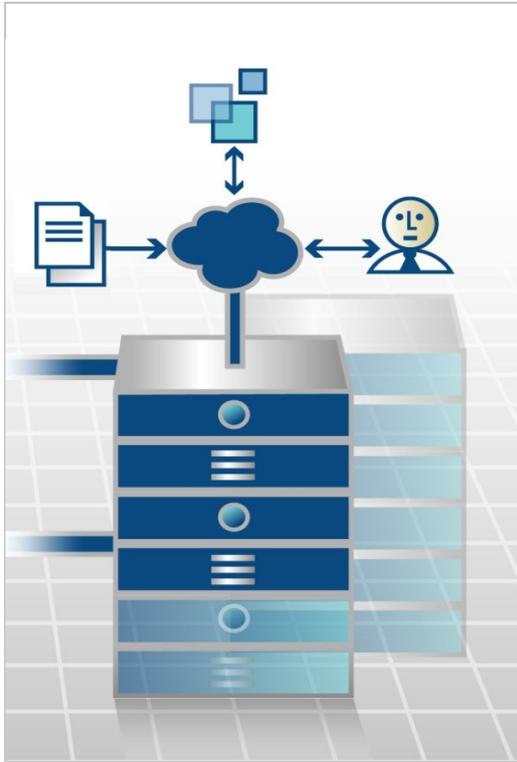


source: "Automated failover and recovery of virtualized guests in Advanced Platform", Rob Kenna, Redhat Magazine

“Virtualization success depends on tight cooperation between server, storage, network and security teams.”

Chris Wolf, Burton Group  
CIO Magazine,  
“Virtualization in the Enterprise Survey”

# An open architecture allows for the ability to add, upgrade and swap components....



- Commodity based approach
  - Low Start-Up and Maintenance Costs
  - Lower Technology Refresh Expenses
- Open Source Promotes Faster Deployment
- High Availability
- Energy Efficient System Scaling – “Green”
- Encourages Standardization
- Open Standards = Business Process Agility
- Reuse of Business Components
- Agile Infrastructure
  - Virtualized Pools of Resources
  - Transparent Sharing of Distributed Resources by Multiple Clients
- Realign current infrastructure
  - Reduce current footprint by powering down excess equipment
  - Reclaim previously wasted resources
- Distributed Processing / Storage

“ Service Oriented practices help create more flexible, adaptable enterprise solutions and encourage standards-based service sharing, reuse, and integration.”

- From the FBI OCIO, February 2005,  
*Building an Enterprise-Class IT Capability...*

# A holistic and integrated approach to Enterprise Architecture must address four dimensions of change...

How do we ensure success?

The People, Process, Technology, and Physical Infrastructure dimensions are the key enablers of the capabilities that an organization uses to fulfill its mission.

## People

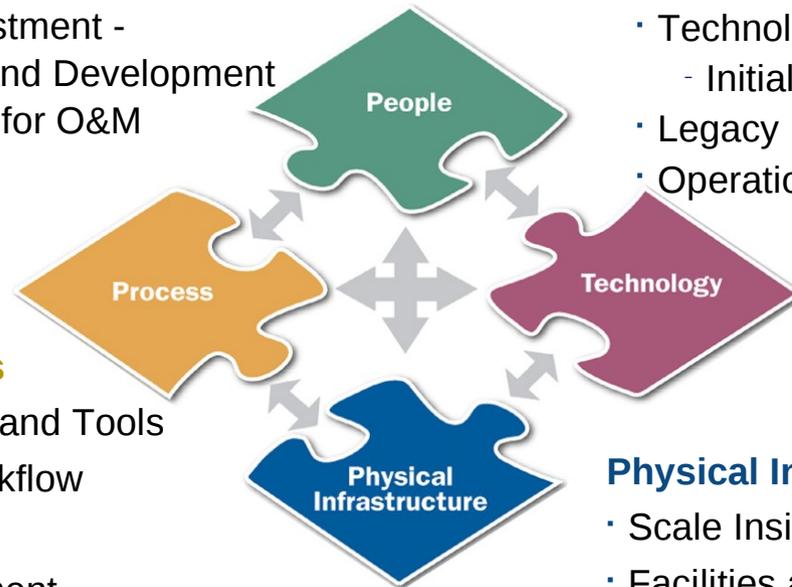
- Communication
- Qualified Staff
- People Investment -
  - Training and Development
  - Coaching for O&M

## Processes

- Activities and Tools
- Data Workflow
- Scope Management
  - Governed Decision Making
  - Support Inherent Agility
- Legacy Capabilities

## Technology

- Data
- Applications
- Technology Infrastructure
  - Initial & Refresh
- Legacy Databases
- Operations & Maintenance



## Physical Infrastructure

- Scale Inside the Box
- Facilities at a Given Location
- Workplace Environment
- Space and Utilization Targets
- Power / Cooling

# Questions?

Isaac Christoffersen

[christoffersen\\_isaac@bah.com](mailto:christoffersen_isaac@bah.com)

Booz | Allen | Hamilton  
**Associate**

*Isaac Christoffersen has over ten years of experience in system integration and software solutions development for commercial and government clients. Currently he is involved in the design, development, and implementation of mission-critical systems for the U.S. Federal Government using the latest SOA, JEE, and open source standards and technologies.*

Dan Santillo

[santillo\\_daniel@bah.com](mailto:santillo_daniel@bah.com)

Booz | Allen | Hamilton  
**Associate**

*Dan Santillo is a senior software engineer with over a decade of experience in software development, software and system testing, network and system administration, and system integration. He is currently supporting a Federal Government client in the design, development, and implementation of a grid-enabled services infrastructure.*

BACKUP SLIDES

# REFERENCE

## What is Service-Oriented Architecture (SOA)?

“ The key value proposition of a service-based environment is to enable greater and more flexible information and capability sharing across the DoD, IC, the rest of the Federal Government, and other mission partners.”  
- "DoD and IC Commitment to an Interoperable Services-Based Environment" dated July 13, 2007

### Services

- interface contracts that are platform independent
- dynamically discoverable and invoked
- self-contained, maintaining their own state
- loosely-coupled, n-Tier friendly

### Messages

- standards based (XML/XSD) conduits
- queuing technologies can provide ESB backbone and are platform independent

### Dynamic Discovery

- scalable, add horsepower as needed
- decouple the producer and consumer
- ease integration and upgrade

### Web Service

- built on well known platform independent protocols (HTTP, SOAP, WSDL, XML, UDDI...)
- lend to agility of business process changes and technology changes
- facilitate distributed processing with distributed storage

### Orchestration

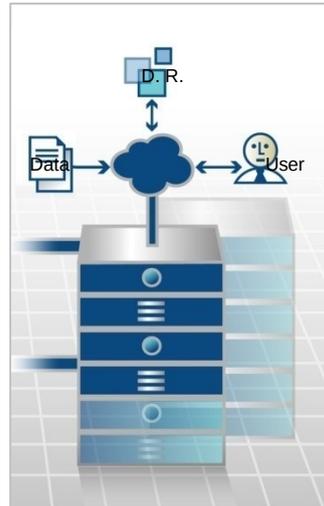
- compose-able workflows for data and users
- BPEL not yet standardized but maturing and workable

### Cluster

- service redundancy
- extensibility
- reliability

# REFERENCE

Technologies  
implemented within the  
**Service Grid**



## RHEL 5

- Xen (Virtualization)
- GFS (Clustered File System)
- Grid Management (Conga)
- Red Hat Cluster Suite

## JBoss

- ESB (+Orchestration)
- High Availability Messaging (JMS, SOAP)
- JBoss Operation Network

## Open Source Tools

- Integrator (Not Developer)

## Web Server

- Compose-able Application
- Java Portlet (e.g., JSF UI)

# REFERENCE

Technologies  
implemented within the

## Database Grid



### Oracle 10G

- RAC (High Availability)/Dynamic Provisioning
- Rapid, Cost Effective Integration of Hardware Advancements
- DataGuard supports DR Integration
- Grid Control
  - Balancing/Availability (TAF/FAN)
  - Maintenance
  - Extensible/Scalable
  - Automatic Storage Management (ASM)
  - Dynamic Load Balancing and Optimization of I/O Throughput continuously without Third Party License
  - Remote monitoring
- Data Protection
  - RMAN
    - Backup/recovery
    - Database Duplication
  - Flash Recovery
  - DataGuard (Disaster Recovery)
- Text – Index
  - Multi-language Content Search
  - Unicode
  - Stored within Database Content
- RHEL 4