



# Optimizing the Enterprise: *Using JBoss and RHEL Virtualization*

Isaac Christoffersen, Booz Allen Hamilton

# Virtualization is adding a new chapter to the history of the data center...

<b>Virtual Infrastructure</b> Hardware + Virtual Infrastructure Software	<b>Virtual Machines</b> Capabilities developed in more efficient virtual space, but developed with a traditional mind set
<b>Data Center</b> Physical space, power and cooling	<b>Systems</b> Hardware, software, etc.

# The demands on the data center are only going to increase...

Computing, network and storage resource needs will continue to increase...

“ In 2006, humans created 161 exabytes of data, 3 million times the information in all books ever written. ”  
- IDC US Green IT Survey

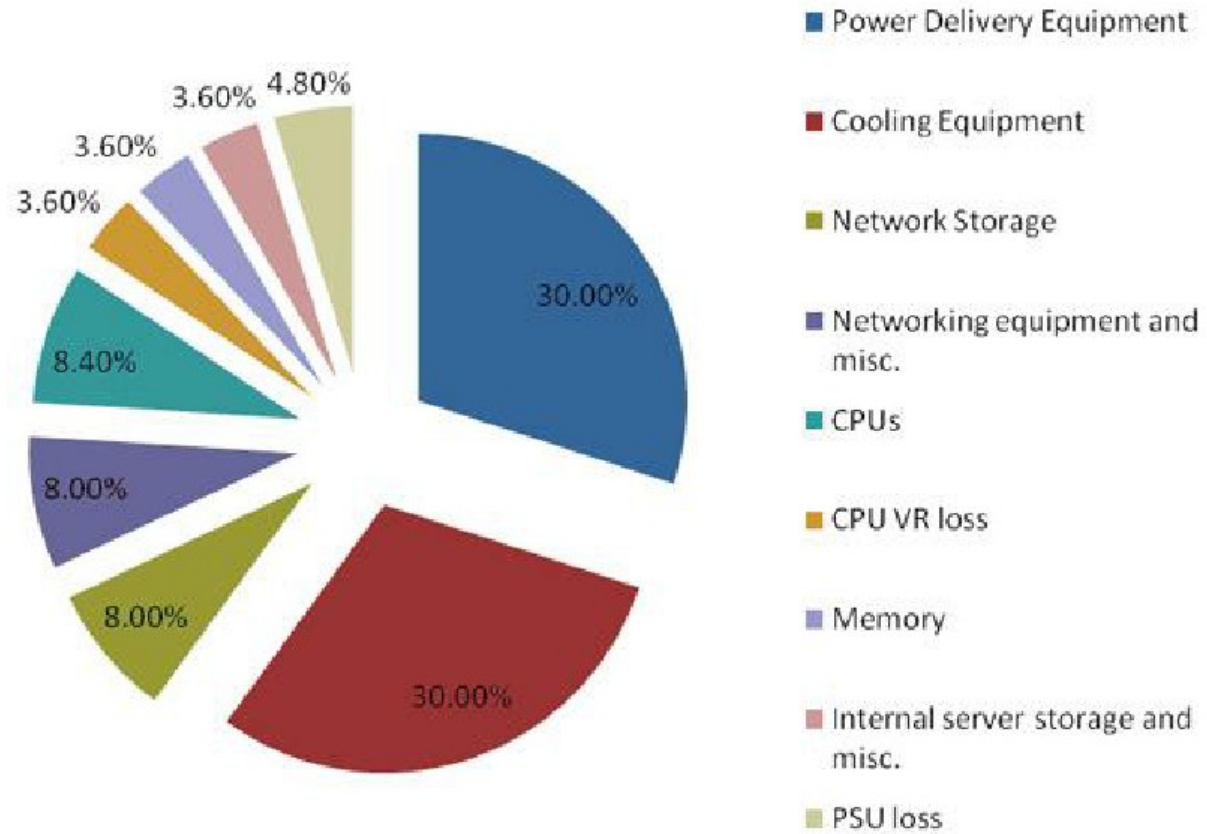
- Satisfying high availability and agility requirements
- Exponentially increasing data and rapidly growing user base
- Need to efficiently manage resources within flat IT budgets
- Integrating improved technology solutions
- Approaching size, weight and power (SWaP) limitations
- Costs increasing with user expectations and maintenance needs

# Power consumption is a growing concern for today's data centers...

Adding resources to meet these challenges is introducing accidental complexity...

“ 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)

# Common misconceptions are adding to the overall complexity...

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

- 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 homogeneous
- 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

*Peter Deutsch's 8 Fallacies of Distributed Computing*

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

# Tomorrow's systems need to be able to expand and contract with demand...

<b>Grid Enabled Services Platform</b> Data Center design and planning (floor space, power and cooling), Hardware, Software, grid storage container, clustered applications server, integrated data source abstraction, grid database container, Directory Server and other capabilities as they are developed and integrated		<b>Services</b> Services Delta
<b>Virtual Infrastructure</b> Hardware + Virtual Infrastructure Software	<b>Virtual Machines</b> Capabilities developed in more efficient virtual space, but developed with a traditional mind set	
<b>Data Center</b> Physical space, power and cooling	<b>Systems</b> Hardware, software, etc.	

# ...and infrastructure needs to shift from a dedicated state to a dynamic one

How does a Grid Enabled Services Infrastructure (GESI) resemble a Service Oriented Infrastructure (SOI)?

“SOI delivers bottom-line benefits to the enterprise. It provides the basis for greater IT automation which results in higher IT productivity and lower operational costs.”

*SOA Practitioner's Guide, part 2*

Addresses all the aspects of the infrastructure

- Networks, servers, data centers, and firewalls
- Application infrastructure, security, monitoring, middleware, etc.

Key elements


- Life cycle support to manage the deployment of SOI components
- Virtualization of infrastructure resources to SOI users
- Service management to assure the SOI solution provides the required service characteristics

Combines architecture building blocks and role-based portals

- High reuse of common services
- Reuse of infrastructure and foundational components
- Reduction in time needed to deliver capabilities

# Infrastructure is still an important component of service oriented architectures

Important business benefits can be realized with good enterprise architecture...



*A more efficient IT operation:*

- Lower software development, support, and maintenance costs
- Increased portability of applications
- Improved interoperability and easier system and network management
- Easier upgrade and exchange of system components

*Better return on existing investment, reduced risk for future investment:*

- Reduced complexity in IT infrastructure
- The flexibility to make, buy, or out-source IT solutions
- Reduced risk overall in new investment, and the costs of IT ownership

*Faster, simpler, and cheaper procurement:*

- Buying decisions are simpler, because the information governing procurement is readily available in a coherent plan.
- The procurement process is faster - maximizing procurement speed and flexibility without sacrificing architectural coherence.

# There are many useful architectural and system design patterns for creating dynamic infrastructures...

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
- Service Oriented Architecture
- Repeatabe Build
- Clustering
  - Operating System
  - Application
  - File System
- Open Architecture

# 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.”

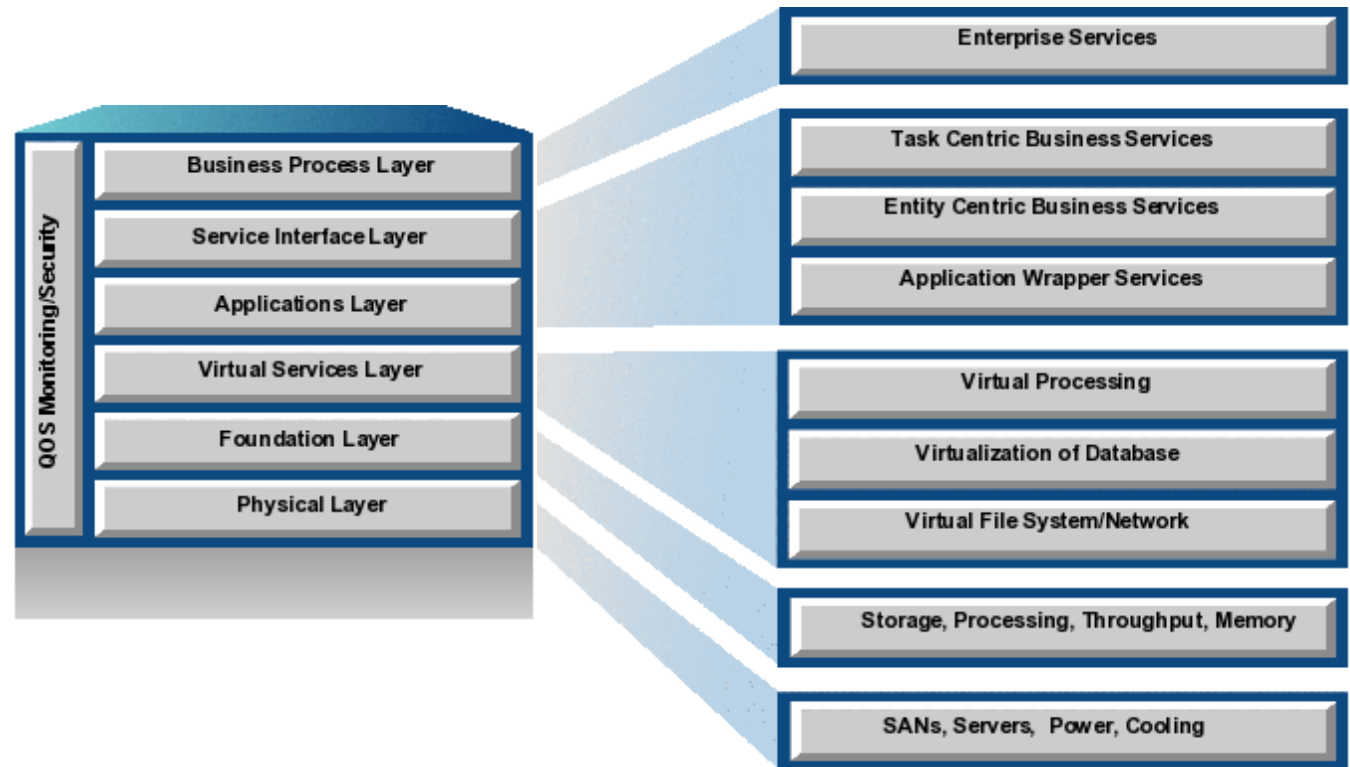
- Paris Avgeriou & Uwe Zdun, *Architectural Patterns Revisited – A Pattern Language*

...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

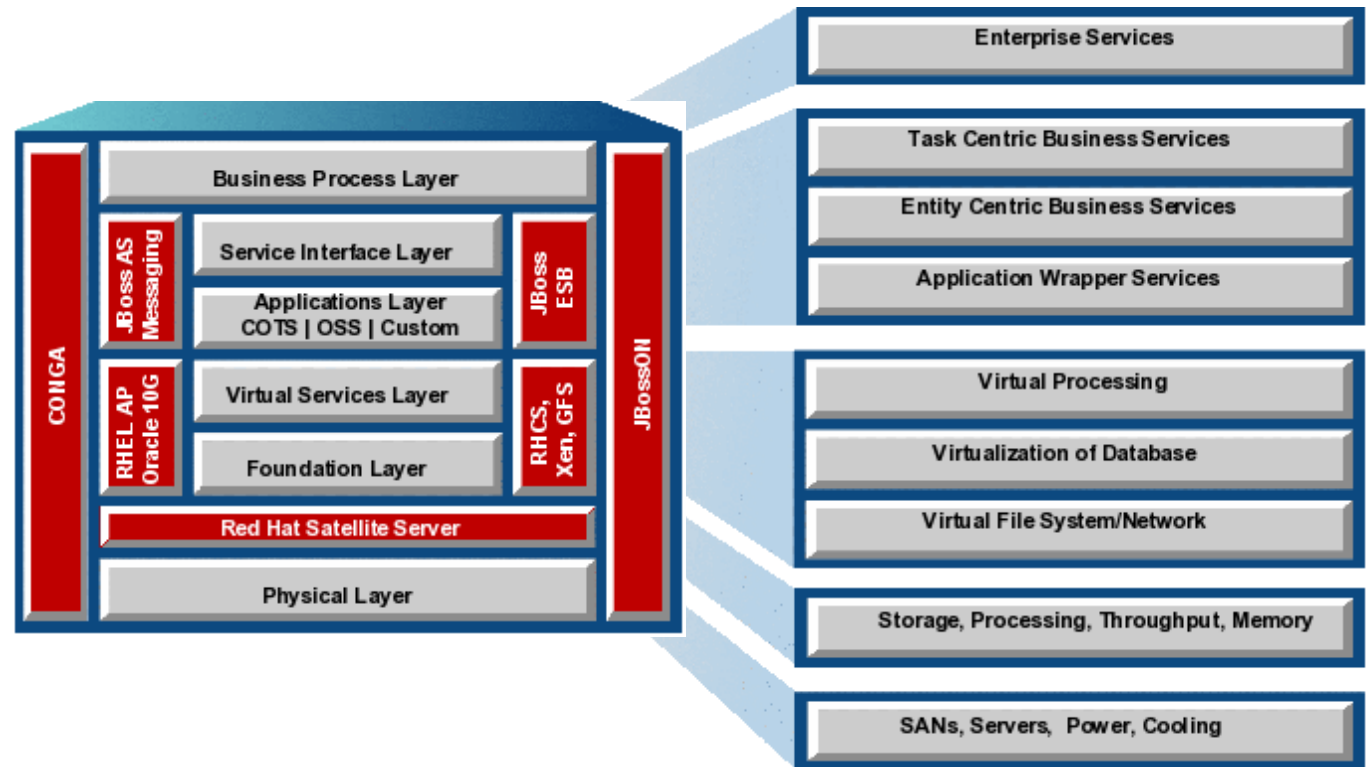


...decomposed into their simpler, individual components and responsibilities...

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 instances 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*

# Shared Storage allows multiple 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

# 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

# Service Oriented Architecture uniquely addresses a rapidly changing environment



## Agility

- ▶ Focus more on core competencies and missions by creating a network of producers-suppliers with intense interactions
- ▶ Improve access to information to enable faster cycle times
- ▶ Enable enterprises to be more agile and respond quickly to business needs

## Process

- ▶ Increase business flexibility through plug-and-play architecture and re-use of existing services
- ▶ Ensure system change is not a constraint on business or mission change
- ▶ Allow inter-operation with other systems & partners without customization

## Interoperability

- ▶ Facilitate integration with multiple solutions via open IT standards
- ▶ Remain platform, language, and vendor independent to remove IT barriers for using best-of-breed software packages

## Costs

- ▶ Reduce development costs by acquiring pre-built capabilities
- ▶ Leverage previous IT investments through re-use of assets
- ▶ Lower maintenance costs and TCO through fewer "instances" of a function, and fewer software licenses

IT alignment with an organization's mission

Improved agility, focus on core competencies, IT efficiencies, and ROI for IT assets

# Consistent, repeatable builds lead to 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
- Reduce data duplication
- Scale clusters on the fly
- Common File System

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



“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...*

- Commodity based approach
  - Low Start-Up and Maintenance Costs
  - Lower Technology Refresh Expenses
- Open Source Promotes Faster Deployment
- High Availability
- Energy Efficient System Scaling
- Encourages Standardization
- Open Standards = Business Process Agility
- 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

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

- Operating System Clustering



- Red Hat Cluster Management

- Virtualization

- Server
- Network



- Xen Virtualization

- Application Server Clustering



- JBoss EAP

- Storage

- Virtualization
- File System Clustering
- File Replication



- Global File System (GFS)
- Globus Toolkit

# Virtualization alone cannot improve the effectiveness of the system...


“Grids enable more flexible, dynamic resource sharing among diverse physical or logical computing environments.”

- Burton Group Grid Services, “Pooling Distributed Resources for Virtual Supercomputing”


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

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

## 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

# Systems built today need to display key characteristics...

- Accredited Platform
- Scalable (Elastic)
  - Vertical
  - Horizontal
- More Efficient
  - Power
  - Cooling
  - Floor Space
- More Effective
  - New Service Delivery
  - Leverage existing capabilities and knowledge
- Agile
  - Infrastructure
  - Software Development
- Cost Effective
- Maintainable
  - Highly available
  - Continual refresh
  - Hardware independent
  - Redundant
- Establish SLO/SLA (Active...Deployable)
  - Availability
  - Scalability
  - Performance

# Grid Enabled Services Infrastructures are dynamic infrastructures...

## What are the benefits of a GESI solution?

- Improve performance, availability, and cost-effectiveness of compute- and data-intensive applications
- Run growing volume of complex, resource-intensive, high-performance computing jobs within existing distributed infrastructure
- Provision additional capacity dynamically as it becomes available, and failover gracefully around unavailable capacity, without interrupting jobs in progress
- Reduces total cost of ownership of ongoing information technology (IT) operations by making more efficient use of available computing, storage, and network capacity
- Postpone the need for deployment of additional capacity to support growing transaction loads
- Provision, scale, and reconfigure virtualized computing resources within a service-oriented environment

# Questions?

Booz | Allen | Hamilton

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

**Associate**

**Christopher Dale**  
[dale\\_christopher@bah.com](mailto:dale_christopher@bah.com)

**Associate**

**Henry Robertson**  
[robertson\\_henry@bah.com](mailto:robertson_henry@bah.com)

**Consultant**

BACKUP SLIDES



**The Art of Data Center Feng Shui:**  
*Achieving Enterprise Harmony Using JBoss, Red Hat and  
Virtualization Technologies*

Isaac Christoffersen, Booz Allen Hamilton

# REFERENCE

What are the key characteristics of today's enterprise applications?

“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
- 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
- DataGuard supports DR Integration
- 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