



Session ID O18

Linux Clusters in non HPC Applications

Dan Cummings, ATS Linux Team, dcumming@us.ibm.com

IBM @server xSeries
Technical Conference

Aug. 9 - 13, 2004

Chicago, IL

Topics in this Discussion

- A Broader Definition of Clustering
- The Value Proposition
- Some Low Hanging Fruit
- Possible Futures

A Broader Definition ...

A Cluster is any collection of systems that can be logically grouped together as a single machine either by task or to simplify administration.


A Broader Definition ...

- Scalable Servers
 - Replace discrete servers with clusters using a virtual IP address mechanism
 - F5 Hardware Load Balancing
 - WebSphere Edge Server
 - Linux Virtual Server
 - PolyServe Matrix
 - Others.....


A Broader Definition ...

- For Example --- Web Farms
 - Content Server Group
 - Application Server Group
 - DB Server Group
 - Spare Nodes Group
 - Proxy Server Group

Web Farm Example



Web Farm Cluster Example



Topics in this Discussion

- A Broader Definition of Clustering
- **The Value Proposition**
- Some Low Hanging Fruit
- Possible Futures

Value Proposition

- Scalability
- Increased Fault Tolerance
- The System Grows with the Workload
 - Lower incremental cost for power increase
- Cluster Management Tools Reduce Complexity of this Environment

Scalability – Network Solutions

- Blocking, Bandwidth and Latency
 - Interconnect must fit application
 - Generally, less latency means more scalability
 - Generally, non-blocking is better but also more expensive
 - Bandwidth is usually not the cause of performance problems
 - Tools for Performance Analysis Exist (Oprofile, Cricket)

Scalability – Fiberchannel Solutions

- Engage the TotalStorage Experts to do this
- Infiniband has fiberchannel interfaces and scales to fairly large node counts
- Large SAN switches and “federated” SAN switch configurations are possible.
- BladeCenters make this simple and relatively inexpensive even with fairly large node counts.

Fault Tolerance

- Clusters are redundant by definition
- Few or no single points of failure
- Virtual IP Capabilities with Heartbeats
 - Linux Virtual Server -- HA Linux
 - Steeleye Lifekeeper
 - Piranha
 - PolyServe Matrix

Grow with the Workload

- Grow as the company grows
 - Low incremental cost for growth
 - No need to obsolete older equipment to add new
- Dynamically add systems for transient changes
 - “Personality Servers”
 - Custom Solutions

Cluster Management Tools

- Reduce Administrative Complexity
- Enable Custom Response to Load / Failure Events
- CSM/Director

IBM Director

- When to use Director
 - When a GUI management interface is desired
 - When both IBM and non-IBM systems are to be managed
 - When extensive self discovery, capacity planning (xSeries only) or Application Workload management (xSeries only) are desired
 - Both Windows and Linux are to be managed

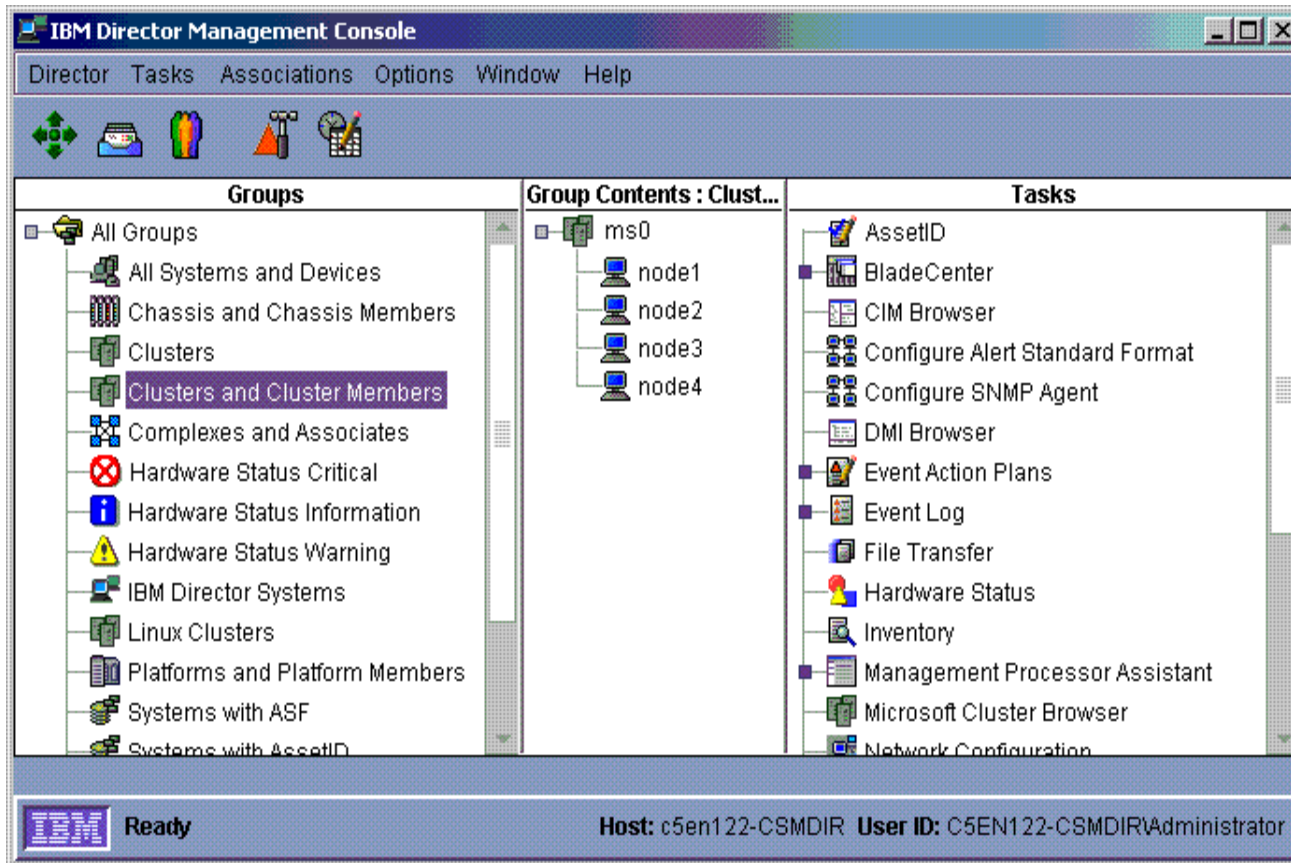
CSM – Cluster Systems Manager

- When to use CSM
 - A Command Line / Scriptable interface is desired
 - Linux only environment or Linux and AIX together
 - Want to manage pSeries and xSeries together
 - Custom integration into existing customer environment is desired.
- Soon There will be CSM Extensions to Director for best of both worlds

CSM and Director

- CSM currently provides a good initial level of integration with IBM Director so that the two products can be used together
 - The Linux CSM clusters are discovered and shown in the IBM Director Console
 - The event infrastructure of both CSM and Director are tied together
 - CSM Conditions are published as Director Event Types, which can then be tied to Event Action Plans
 - When events occur in the CSM cluster, the Director actions can then be run
 - CSM information is displayed on the Director Console Inventory section
- Additional integration between CSM and Director is being designed for follow-on releases

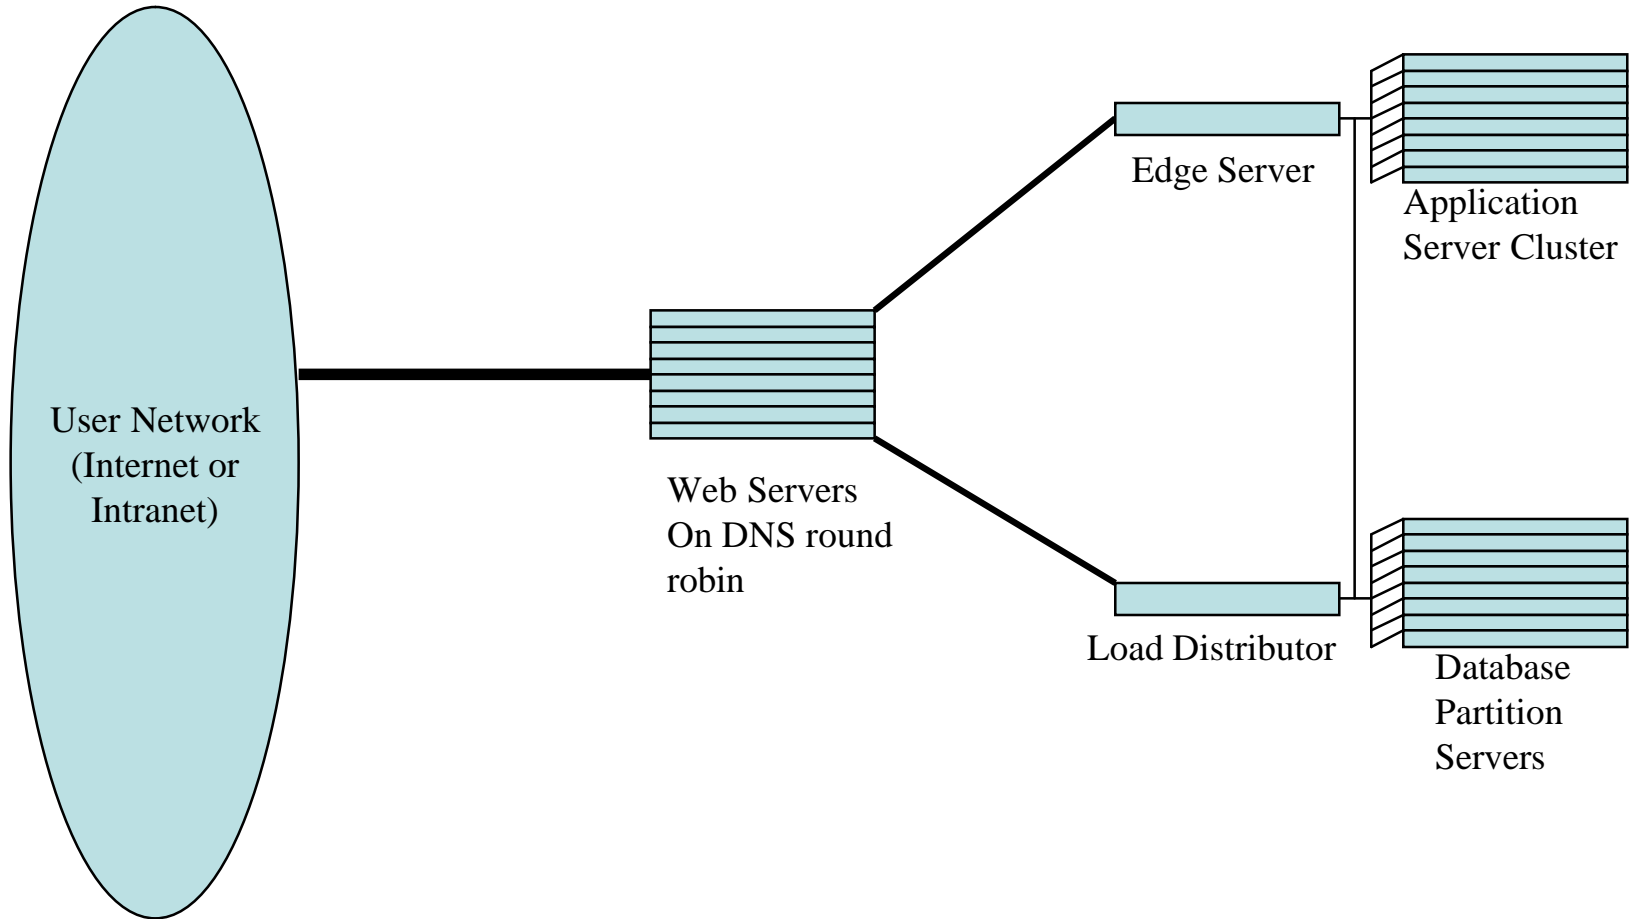
CSM and Director



Topics in this Discussion

- A Broader Definition of Clustering
- The Value Proposition
- **Some Low Hanging Fruit**
- Possible Futures

Web Farm Cluster Example



Web Servers

- Can use DNS Round Robin for simple load balancing
- Since disk access is read-only sharing of disk farm is straight forward
- Smooth growth curve from very small to large and complex as growth / application require

Database Clusters

- Two ways to do DB Clusters
 - Parallel databases
 - HA with failover
- Commercial Products Support This
 - DB2/EEE
 - Oracle 9i RAC

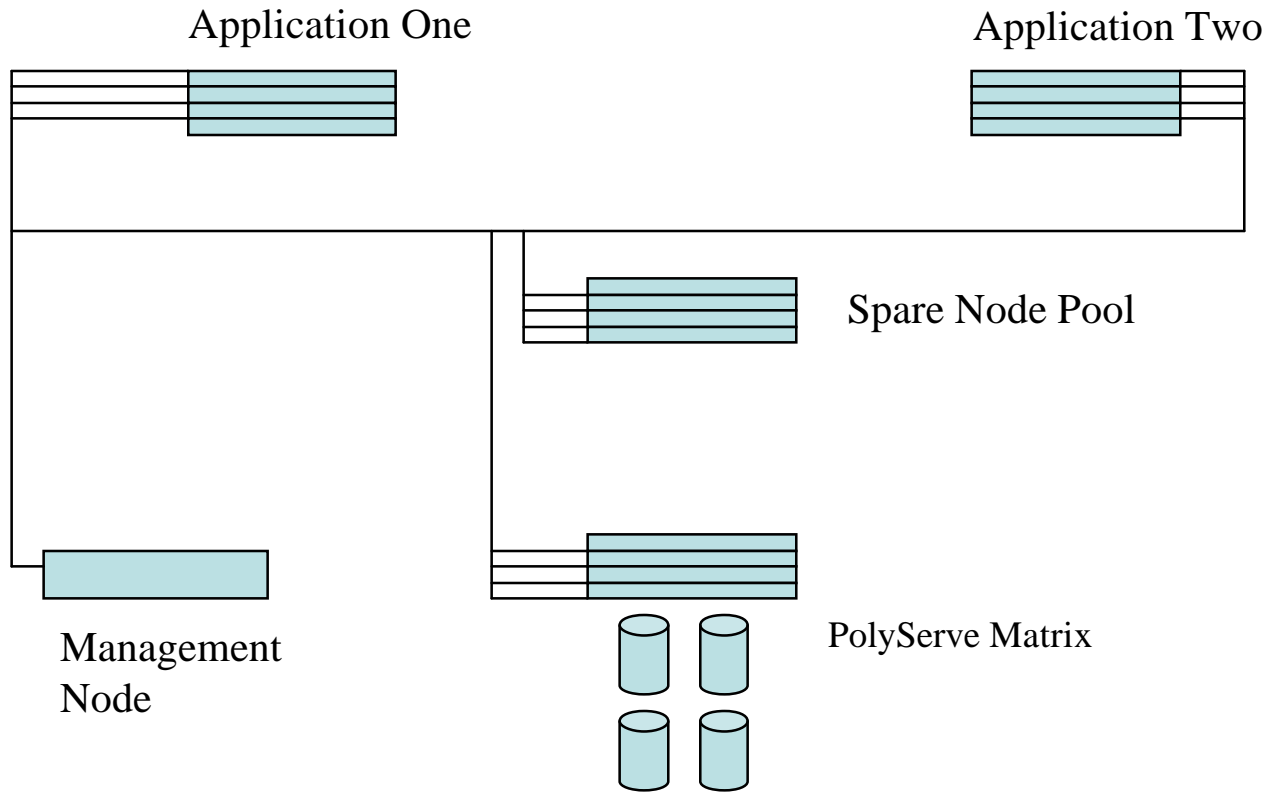
Application Server Clusters

- Can be done with Open Source
 - Apache / Tomcat
 - JBOSS
 - Linux Virtual Server
- Can be done with Commercial Tools
 - WebSphere & WebSphere Edge Server
 - Weblogic

Topics in this Discussion

- A Broader Definition of Clustering
- The Value Proposition
- Some Low Hanging Fruit
- Possible Futures

Do this with monolithic servers



Possible Today

- Management Node runs CSM and Director
- Nodes run CSM, Linux, Applications
- File Service is PolyServe Matrix
- Except for PolyServe and Management node, all nodes are installed the same
- “Personality” is applied by mounts from PolyServe scalable NFS service.
- Systems dynamically applied to applications from spares in under 1 minute automatically.

Summary

- Clusters are for more than supercomputing
- Clusters solve scaling and load balancing problems
- New cluster management tools make clusters simpler to deploy and manage.
- Linux Clusters are Easy

