GIS - Clustering Architectures

Raj Kumar
Integration Management
9/25/2008

Agenda

- What is Clustering
- Reasons to Cluster
- Benefits
- Perimeter Server Clustering
- Components of GIS Clustering
- Perimeter Server Clustering
- GIS Server Architectures
- GIS/SI Clustering Considerations
- Clustering Best Practices/Recommendations

What is Clustering?

- A clustered installation of Gentran Integration Suite connects one database to more than one installation of Gentran Integration Suite
- Clustering can be accomplished using any one of the following approaches:
  - Application Clustering – Also knows as software clustering, in which more than one servers can be turned into a cluster.
  - Hardware Clustering – Requires additional hardware to your existing infrastructure.
    - CPU Licensing, Storage space either File System or Database, Complexity
  - Database Clustering – 4.3 supports Database clustering
    - Scale out oracle instances to recover from database failures

Reasons to Cluster?

- Everybody else is doing it!!! If that's not enough then
- **Performance**
  - GIS System is constantly throwing "OOM" exceptions
  - Running out of threads and messages are being sent to the "Wait Queue" until resources become available.
  - For high availability and failover
  - Allows the use of all cluster nodes all of the time to handle peak processing volumes
- **Serviceability** - Service Level Agreements
- **Failover –** System that can be brought online with minimum efforts to recover from node down failures.
- **Sizing -** When considering these options, future growth and required service levels also must be evaluated
- **RAS** – Improvement in Reliability, Availability and Scalability

Benefits of Clustering

- Clustered installations can improve your operations of Gentran Integration Suite in the following areas
  - **Reliability**, which ensures correct operation.
    - For example: Critical Business Transactions such as Orders, Invoices and ASN's are received and processed in timely manner.
  - **Availability**, which minimizes scheduled and unscheduled downtime.
    - In order to achieve and maximize GIS High Availability hot-standby is required during scheduled and unscheduled maintenance.
  - **Scalability**, which optimizes GIS operational and technical processes to adapt/handle future growth in volume
    - Adapters/Services, Business Processes, Perimeter Servers

Components of GIS Clustering

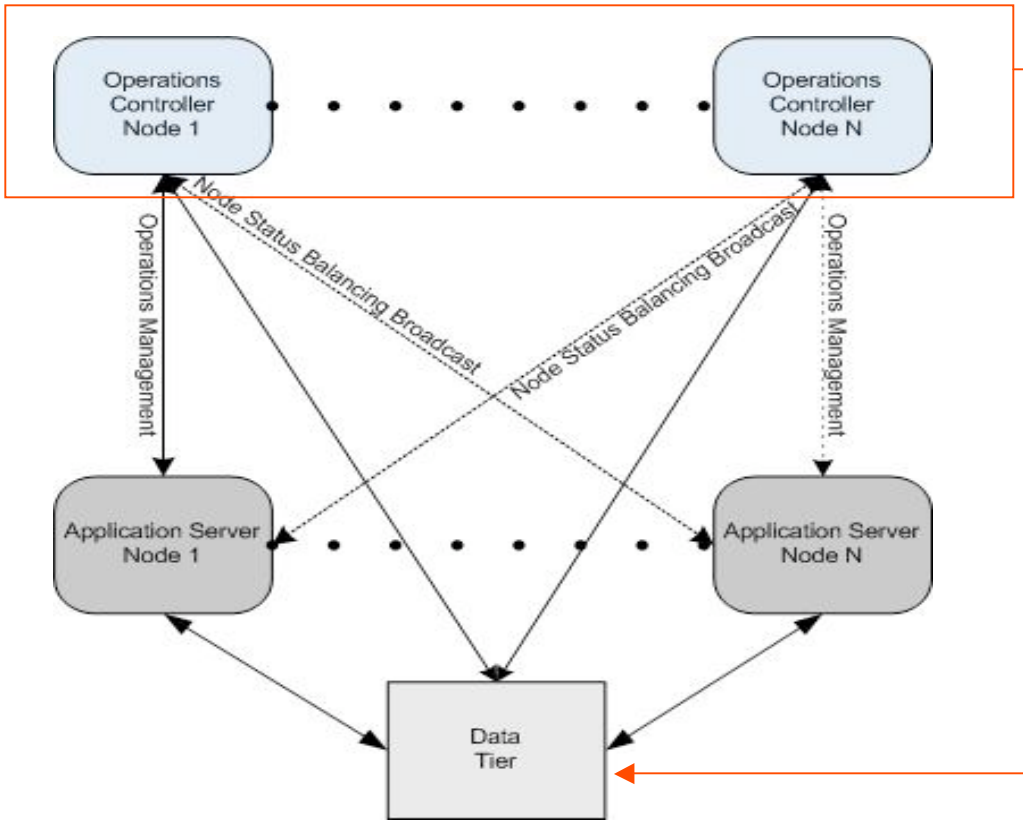

Operations Controller
- Communication

Database Layer

GIS Server Architecture

- **Operations Controller**
  - Manages JVM (Java Virtual Machine) resources.
    - Heap sizes, minimum and maximum stack sizes.
      - These are critical while installing GIS application and configuring adapters such as Command Line Adapter
      - These values can be viewed after the application starts up
  - Each GIS node has an Operations Controller
    - which is responsible for listening across all queues
      - Ops Controller at all time keeps tabs on resources that have allocated to any given node in a cluster.
    - The Operations Controller uses JNDI (Java Naming Directory Interface) technology to communicate directly with other Operations Controllers in the cluster.

GIS Server Architecture

- **Service Controller**
  - Responsible for managing, configuring, querying and caching information for all services and adapters.
    - Includes - Properties files, Process Data I/O.

GIS Server Architecture

- **Central Operations Controller**
  - Communicates with the Operations Controllers
  - heartbeat mechanism for the entire cluster.
    - It runs on a same JVM
    - Each node has its own CentralOps Controller. During GIS startup the Central Operations Controller coordinates the local cluster node startup operations and its corresponding components, such as the Service Controller and the Scheduler.
    - Token Nodes- Allows individual nodes to assume the role of being the Central Operations Controller.

Clustering Configurations

- **Active- Active**
  - This is the preferred clustering configuration.
  - All nodes are available and working all the time, sharing the workload equally

Clustering Configurations

- **Active-Passive**
  - Active-Passive setup includes a second GIS/hardware node that's configured exactly like the first and if the first instance fails, the second instance is ready to take over.
- **Hybrid**
  - An Active-Active cluster is backed up by a passive node that can replace either of the active nodes in the event of hardware failure.

Perimeter Server Clustering

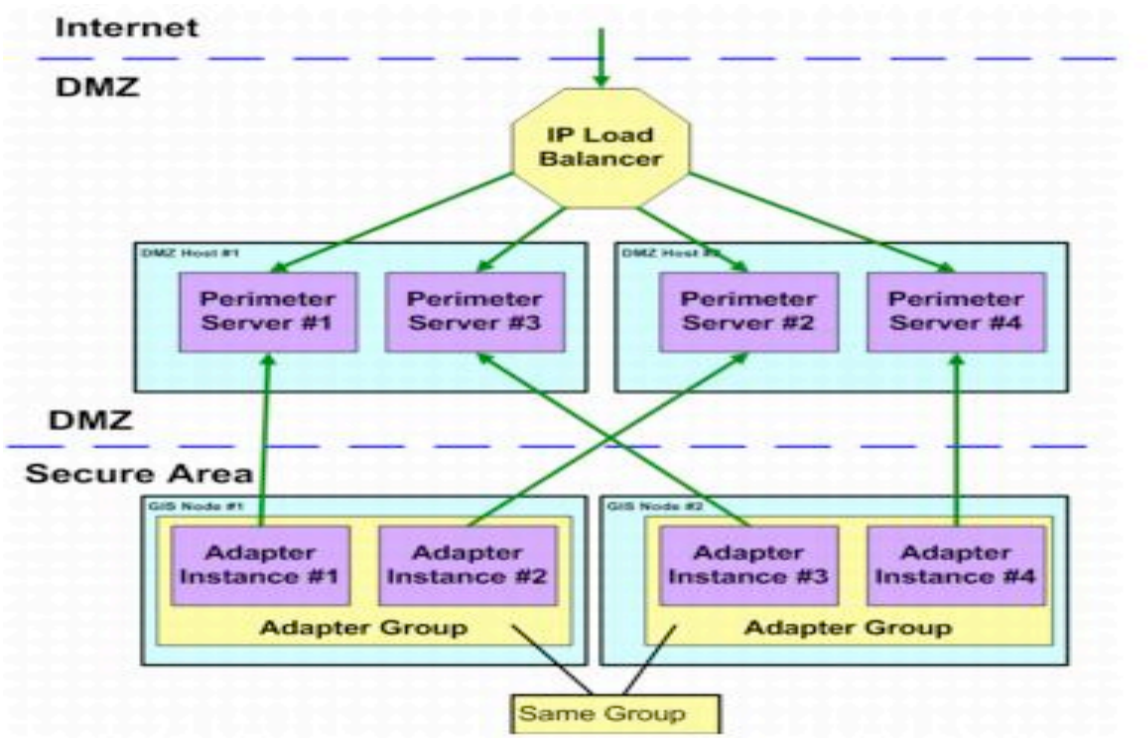

Clustering Server Architectures

- GIS clustered installations can be optimized based on the business priorities
  - **Vertical Clustering** – Increase resources on the existing GIS Server by adding CPU, Memory and hard disk.
  - **Horizontal Clustering** - Adding multiple machines to the GIS architecture.

Advantages of Vertical Clustering

- **Concurrency**
  - More CPUs allow concurrent threads to run simultaneously
  - Increases throughput
- **BP/Services Execution Performance**
  - Faster CPUs add horsepower to GIS resulting efficient BP and services execution
- **Overall Performance**
  - Adding memory to GIS improves the overall health of GIS. Memory is critical when caching objects and parameters at run-time

Advantages of Horizontal Clustering

- **High Availability - No Outage**
  - Nodes can be added to a running clustering without any outage
  - Maintains “No-Down-Time” Service Level Agreements with your Trading Partner
- **Failover**
  - Processes can be recovered efficiently
- **Disaster Recovery**
- **Quick Build – Very less configuration required**

GIS/SI Clustering Considerations

- Disk space for locally stored documents
  - File Systems
- Disk space in the database
  - Document storage, Document Tracking and Persistence
- Database processing power
  - Database connections (along with memory and other database resources)
  - Database layer must be comparable to your application cluster infrastructure for cluster to work efficiently.

GIS/SI Clustering Considerations

- GIS node memory
  - Currently 2GB is supported per node until 64 bit JVM is supported.
- Network bandwidth
  - To and from GIS adapters
  - To and from database(s)
  - Between cluster nodes

GIS/SI Clustering Considerations

- Multicast vs. JGroups
  - Methods in which GIS nodes communicate with each other in a cluster.
  - Multicast not supported by some network configurations and is a bottleneck when stretching the cluster across a WAN.
- Data Storage
  - If using File System Storage – a common/shared storage location will be necessary
- Perimeter Servers
  - To maintain Cluster standards, at least two Perimeter Server instances will need to be created for each GIS Node in the cluster

GIS/SI Clustering Considerations

- Non-Clusterable Adapters
  - Some Adapters/Services cannot apply to all the nodes in a Cluster. You will need to configure these adapters for each node that applies.
    - Connect Enterprise for Unix adapter – CEU
  - Common/Similar configurations can be grouped together in Service Groups. These Service Groups can be called from Business Processes instead of individual service configurations.
    - HTTP, B2B, FTP and SFTP Client adapters can be grouped together to achieve load balancing at run time.
  - An adapter/service configuration will be required for each Perimeter Server instance in the DMZ (GIS/SI nodes X 2)

GIS/SI Clustering Considerations

Examples of Non-Clusterable Adapters

- FTP server adapter
- HTTP server adapter
- FTP client adapter
- HTTP client adapter
- Connect:Direct adapter
- Connect:Enterprise adapter
- B2B Communication adapter
- CLA2 Adapter
- File System Adapter

GIS/SI Clustering Considerations

IP Load Balancing

- IP Traffic to non-clustered adapters will require the use of an external Load Balancer to split the load incoming traffic between the Perimeter Server (or Local Node) Instances.
- Options include:
  - DNS – Multiple IPs to one URL
  - Proxy – Particularly for HTTP
  - IP Balancing – Hardware or Software that controls incoming sessions and hands them off to the most available Perimeter/Node

Clustering Best Practices/Recommendations

Network Level

- Strive for low latency between servers.
  - If your latency is greater than 10 milliseconds – you may experience issues like nodes not recognizing one another.
    - Messages such as “Node 1 went down” are thrown
- In GIS 4.3 and later use JGroups verses Multicast in cluster configurations
  - JGroups does not flood the system with global packets, just the specified network segment.
  - Multicast will send information to the entire subnet as oppose to targeted GIS hosts.

Clustering Best Practices/Recommendations

Application Level

- Node-to-Node communication:
  - If installing multiple nodes on the same machine, go to the install_dir/properties directory of nodes 2 and higher and change the mcast_port property in jgroups_cluster.properties.in to point to the value of node 1's mcast_port property in jgroups_cluster.properties.

Clustering Best Practices/Recommendations

Application Level (cont)

- In a cluster environment, the service group does a round robin of all adapters on a per node basis. Each node handles its own load balancing across all adapters belonging to the service group.
- If a service group detects failure of a specific adapter, it goes out to get a good one, thus providing not only load balancing but also failover.

Clustering Best Practices/Recommendations

Application Level (cont)

- Scanner BP can help identify disabled or adapter failures and redirect the functionality to an active node
  - Overhead of Business Process Development.
- In GIS version 4.3 intelligence exists to move the BP back to correct nodes for communication.
- Pin a Adapter to a node and then put adapters in a service group and use Service Group in the BP

Q&A