Fencing
: To ensure reliable fencing when using **qdiskd**, use power fencing. While other types of fencing (such as watchdog timers and software-based solutions to reboot a node internally) can be reliable for clusters not configured with **qdiskd**, they are not reliable for a cluster configured with **qdiskd**.

Maximum nodes
: A cluster configured with **qdiskd** supports a maximum of 16 nodes. The reason for the limit is because of scalability; increasing the node count increases the amount of synchronous I/O contention on the shared quorum disk device.

Quorum disk device
: A quorum disk device should be a shared block device with concurrent read/write access by all nodes in a cluster. The minimum size of the block device is 10 Megabytes. Examples of shared block devices that can be used by **qdiskd** are a multi-port SCSI RAID array, a Fibre Channel RAID SAN, or a RAID-configured iSCSI target. You can create a quorum disk device with **mkqdisk**, the Cluster Quorum Disk Utility. For information about using the utility refer to the mkqdisk(8) man page.

Note

Using JBOD as a quorum disk is not recommended. A JBOD cannot provide dependable performance and therefore may not allow a node to write to it quickly enough. If a node is unable to write to a quorum disk device quickly enough, the node is falsely evicted from a cluster.

2.7. Red Hat Cluster Suite and SELinux

Red Hat Cluster Suite for Red Hat Enterprise Linux 5 requires that SELinux be disabled. Before configuring a Red Hat cluster, make sure to disable SELinux. For example, you can disable SELinux upon installation of Red Hat Enterprise Linux 5 or you can specify **SELINUX=disabled** in the **/etc/selinux/config** file.

2.8. Multicast Addresses

Red Hat Cluster nodes communicate among each other using multicast addresses. Therefore, each network switch and associated networking equipment in a Red Hat Cluster must be configured to enable multicast addresses and support IGMP (Internet Group Management Protocol). Ensure that each network switch and associated networking equipment in a Red Hat Cluster are capable of supporting multicast addresses and IGMP; if they are, ensure that multicast addressing and IGMP are enabled. Without multicast and IGMP, not all nodes can participate in a cluster, causing the cluster to fail.

Note

Procedures for configuring network switches and associated networking equipment vary according each product. Refer to the appropriate vendor documentation or other information about configuring network switches and associated networking equipment to enable multicast addresses and IGMP.

Note

IPV6 is not supported for Cluster Suite in Red Hat Enterprise Linux 5.

2.9. Considerations for Using Conga

When using **Conga** to configure and manage your Red Hat Cluster, make sure that each computer running **luci** (the **Conga** user interface server) is running on the same network that the cluster is using for cluster communication. Otherwise, **luci** cannot configure the nodes to communicate on the right network. If the computer running **luci** is on another network (for example, a public network rather than a private network that the cluster is communicating on), contact an authorized Red Hat support representative to make sure that the appropriate host name is configured for each cluster node.

2.10. General Configuration Considerations

You can configure a Red Hat Cluster in a variety of ways to suit your needs. Take into account the following considerations when you plan, configure, and implement your Red Hat Cluster.

No-single-point-of-failure hardware configuration

Clusters can include a dual-controller RAID array, multiple bonded network channels, multiple paths between cluster members and storage, and redundant un-interruptible power supply (UPS) systems to ensure that no single failure results in application down time or loss of data.

Alternatively, a low-cost cluster can be set up to provide less availability than a no-single-point-of-failure cluster. For example, you can set up a cluster with a single-controller RAID array and only a single Ethernet channel.

Certain low-cost alternatives, such as host RAID controllers, software RAID without cluster support, and multi-initiator parallel SCSI configurations are not compatible or appropriate for use as shared cluster storage.

Data integrity assurance

To ensure data integrity, only one node can run a cluster service and access cluster-service data at a time. The use of power switches in the cluster hardware configuration enables a node to power-cycle another node before restarting that node's HA services during a failover process. This prevents two nodes from simultaneously accessing the same data and corrupting it. It is strongly recommended that *fence devices* (hardware or software solutions that remotely power, shutdown, and reboot cluster nodes) are used to guarantee data integrity under all failure conditions. Watchdog timers provide an alternative way to to ensure correct operation of HA service failover.

Ethernet channel bonding

Cluster quorum and node health is determined by communication of messages among cluster nodes via Ethernet. In addition, cluster nodes use Ethernet for a variety of other critical cluster functions (for example, fencing). With Ethernet channel bonding, multiple Ethernet interfaces are configured to behave as one, reducing the risk of a single-point-of-failure in the typical switched Ethernet connection among cluster nodes and other cluster hardware.