Table of Contents

High Availability Portfolio **1**
Client Access 1
Storage 1

High Availability Portfolio

This page describes potential options for improving the availability of a product without requiring a change of the application code. While the best high availability is that achieved at the application layer, it is often not possible to implement a full solution in a cost effective manner.

These solutions below can be used to reduce the impact of an outage

- Transparently switching the service from one instance to another
- Assisting the change by providing a script which could be executed by the operators without assistance or surveillance
- Recovery via a set of remote actions by an administrator

The aim of all of these options is to allow production services to be recovered as soon as possible in the event of failures of hardware, fabric components (operating system, disk, network) or application components.

High Availability requires two basic components

- The client software needs to be able to find the service regardless of who is the current master.
- The application servers need to provide a consistent view of the data regardless of who was the master

Client Access

The following client access methods are available and will be investigated for each of the Grid services

- List of servers in the client configuration. When the first server in the list is down, the second server is tried. This requires that the client code supports the concept of multiple servers and that the application does not require an extended session which has state but can be retried in the event of failure.

- Round Robin DNS. The DNS provides a list of hosts which provide the service and the client selects one to talk with. The Round Robin can provide both availability (chose machines which are running ok) and load balancing (i.e. chose the best/fastest machine). The constraints of this approach are in the frequency of update of the DNS since expriration can take several hours. This is currently used extensively at CERN.

- IP Address takeover. A standby machine can take the IP address of a production machine in the event of failure. The standby machine detects the problem, adopts the address of the master machine and then responds to requests. This requires that either the requests are stateless or that the state data is switched along with the network address. Software such as Linux-HA provide this function and it is used for the DNS servers at CERN.

Storage

The storage contains the state of the open transactions. The following storage approaches may be considered

- Database. Using an Oracle RAC cluster, a high availability backend can be constructed. These services are already running extensively across CERN.

- SAN attached storage, cold standby. All application state data is stored on external disks connected by fibre channel. In the event of a failure of the master machine, a second machine can be made active and access the physical shared disk. An fsck operation may be required to get consistent data for the file system.

- A distributed file system such as GFS can be used to provide a coherent file system. These share physical disks attached to SANs or IP networks such that any machine can access the data. A distributed lock manager handles any conflucts.

-- TimBell - 06 Sep 2005

This topic: LCG > ScFourHighAvailabilityPortfolio
Topic revision: r2 - 2005-09-07 - TimBell