

Slide 1: Availability

I'm Keith White, Director of Availability Engineering, here to talk to you about availability everywhere across your systems.

Slide 2: Availability Everywhere

By everywhere, I mean network availability, data availability, and application availability. Let's talk about each one of those individually.

Slide 3: Simple Cluster

A simple cluster, in particular a Sun Cluster, is composed of redundancy across all pieces of the environment. That means the servers are redundant, as in two servers, two pieces of storage, two interconnects between the servers themselves, and two connections or shared connections to the storage. The application is available to start on either of these systems and we provide application management; so, starting, stopping, monitoring, failover, and restarting of an application across this environment.

Slide 4: Network Availability

For the network we provide a global network service; that is, the IP address of the service is a single IP address in the event of any failure. So we provide redundant connections between the servers themselves to the network; so, the interface to the outside clients accessing this service, as well as redundant connections in between all of the hosts and storage in the network. So, any failure of the networking doesn't cause any failure in the system. It just passes back the same IP address or it looks like the response and requests go to the same IP address no matter what the failure of the networking side of the solution. In addition, the systems come in through that single request and get farmed out across the system, so you have load balancing and you can provide scalable services. So services that can take advantage of multiple instances running at the same time can be load balanced across that.

Slide 5: Data Availability

For data availability -- like I said earlier, we have redundant connections and shared connections for the storage, so all devices have access to the storage. So any failure of a storage device does not provide failure to the service because the machines have access to the storage through the other devices that are connected to it. So we have global data services that way. In addition, a file system -- looks like a global file system or a shared file system because each of the machines or the path to that storage is shared across these systems. So the service itself responds back even with any of the failures to that storage as if there was no failure at all. In addition to a shared file system, we have failover file systems. So, should one instance of the file system, the other one starts up, brings up the service; you can't tell the difference.

Slide 6: Application Availability

For application availability -- so your application, should it fail, we have the failover service that restarts the application on another system, brings it up, looks like the service has a brief pause, and then restarts just like that so you can't tell the difference, that the service was down. It was never down; it was just briefly interrupted. And, as I stated previously, a scalable service -- in the case previously for file system -- a scalable service allows multiple instances to run across all of the systems in the cluster and be load balanced or share all the load across those systems. As part of the Sun Cluster solution, we have a resource manager that manages interdependencies between resources. So, if your service is dependent on a database or some other file system or something else in the system to be there, you can set up those dependencies so that the service does not start until those things are available.

Slide 7: Application Availability

So, our integration with Solaris containers -- everybody's aware of the Solaris Container concept. We have the ability in the 3.1 U4 product to manage a whole container failover to whole other container. So if a container or part of that container goes away, the whole container is brought up on another system or another instance of the container is brought up and the service continues over there. In the next revision of the product, we actually have services that will failover between containers. So the whole container does not need to be brought up, but the service itself can just failover from one container to another much like between systems of that service. So, as you can see, with all the parts of a service, from the storage, the server, the interconnects, to the application, all parts are covered by the Sun Cluster software. Any failure does not result in service failure. Sun Cluster continues the service indefinitely.