

Cluster Availability

May 2007

In our companion article, [Active/Active versus Clusters](#), we hinted at an availability difference between active/active systems and clusters. Let us look at that difference in more detail.

Failures and Failovers

There are two predominant reasons for a redundant system such as a cluster or an active/active system to be unavailable to its users:

- *System Failure* - The system is down.
- *Server Failover* - One server in the system is down and is failing over to a backup system.

The analysis of the availability impact of these two failure modes is documented in great detail in our previous articles.¹ We review this analysis briefly below and apply it to a cluster and an equivalent active/active system.

System Failure

Let us consider the single-spared case. That is, the cluster or the active/active system has a single spare node. If one node should fail, the system survives. However, if two nodes should fail, the system is down.

We define the following parameters:

- n is the number of nodes in the system.
- a is the availability of a node (the percent of time that it is up).
- f is the probability of failure of a node (the percent of time that it is down).
- F_d is the probability of failure of the system or the probability that the system will be down.
- F is the probability that the system will be down.
- A is the availability of the system.

a is the probability that a node is up. Since a node is either up or down, the probability that the node will be down, f , is $f = (1-a)$.

¹ *Calculating Availability – Redundant Systems*, Availability Digest, October, 2006, and *Calculating Availability – Failover*, Availability Digest, February, 2007.

The probability that two specific nodes will be down at the same time is the probability that one node is down and the probability that a second node is down. This is $F^2 = (1-a)^2$.

The system will be down if any two nodes in the system fail. If there are n nodes in the system, there are $n(n-1)/2$ ways that two nodes can fail. That is, one out of n nodes can first fail followed by one of the $(n-1)$ remaining nodes for a total of $n(n-1)$ combinations of node failures. But this has counted each node pair twice, i.e., a node 3 failure followed by a node 2 failure and a node 2 failure followed by a node 3 failure. Thus, there are only $n(n-1)/2$ unique combinations of a dual node failure.

Therefore, the probability of a system failure due to a dual node failure, F_d , is

$$F_d = \frac{n(n-1)}{2}(1-a)^2 \quad (1)$$

Server Failover

If a single node fails, the system survives, However, during the time that the applications (in the case of a cluster) or the users (in the case of an active/active system) are being switched over, users are not being serviced. So far as they are concerned, the system is unavailable.

Let

- F_f be the probability that the system is down during a failover.
- MTFO be the mean time to failover (that is, the average failover time).
- mtbf be the mean time between failures of a node.

When one node fails, service to the users will be down for a period of MTFO. Since this will happen on an average every mtbf period of time, the proportion of time that the system is down during failover of a particular node is MTFO/mtbf.

However, there are n nodes in the system, the failure of any one of which will cause a failover outage. Therefore, the probability of a failover outage for the system, F_f , is

$$F_f = n \frac{\text{MTFO}}{\text{mtbf}} \quad (2)$$

Availability

The probability of failure of the system, F , is the probability, F_d , that the system will be down due to a dual node failure plus the probability, F_f , that the system will be down due to a failover:

$$F = F_d + F_f = \frac{n(n-1)}{2}(1-a)^2 + n \frac{\text{MTFO}}{\text{mtbf}} \quad (3)$$

The availability, A , of the system is

$$A = 1 - F \quad (4)$$

Clusters

Let us apply these equations to a typical cluster to calculate its availability. We will consider a two-node cluster comprising industry standard servers, each with an availability of three 9s and an mtbf of 4,000 hours (about six months). Cluster failover time is five minutes, or 300 seconds

(bringing up the applications, recovering the database, connecting to the database, and switching IP addresses). Thus,

$$\begin{aligned} n &= 2 \\ a &= .999 \\ \text{MTFO} &= 300 \text{ seconds} \\ \text{mtbf} &= 4,000 \text{ hours} \end{aligned}$$

For these parameters, Equations (1) through (4) yield

$$\begin{aligned} F_d &.000001 (1 \times 10^{-6}) \\ F_f &.000042 (42 \times 10^{-6}) \\ F &.000043 (43 \times 10^{-6}) \\ A &.999957 (4+ 9s) \end{aligned}$$

Clearly, in this case, failover time is the major factor in system availability and yields an availability halfway between four 9s and five 9s.

Active/Active Systems

Let us make this same availability calculation for an active/active system, except that the failover time is three seconds (switching over users):

$$\begin{aligned} n &= 2 \\ a &= .999 \\ \text{MTFO} &= 3 \text{ seconds} \\ \text{mtbf} &= 4,000 \text{ hours} \end{aligned}$$

But first, we must make an adjustment to Equation (2). Unlike a cluster, in which all users are down during failover, in an active/active system only the users at the failed node are down during failover. Thus, only $1/n$ users are affected. Since we count downtime based on users and not on a server basis, the probability of a failover fault, F_f , as given by Equation (2), must be divided by n :

$$F_f = \frac{1}{n} \frac{\text{MTFO}}{\text{mtbf}} = \frac{\text{MTFO}}{\text{mtbf}} \quad \text{for active/active systems} \quad (2a)$$

In this case, the results are

$$\begin{aligned} F_d &.000001 (1 \times 10^{-6}) \\ F_f &.00000021 (.21 \times 10^{-6}) \\ F &.00000121 (1.21 \times 10^{-6}) \\ A &.99999879 (6- 9s) \end{aligned}$$

Thus, the availability of an active/active system, using the same systems as the clustered system, is almost six 9s.

In this case, failover time for the active/active system is 1% of the cluster failover time. Though failover time is the predominant availability component for the cluster, it is not very significant for the active/active system. Because of the much shorter failover time and the fact that only a portion of users suffer a failover outage when a node in an active/active system fails, the active/active configuration enjoys an availability that is more than an order of magnitude greater than that of a cluster.

It is instructive to consider the availability of the same active/active system if a fault-tolerant system such as HP's NonStop server is used. These systems have proven availabilities of about four 9s, about ten times that of an industry standard server. If an industry standard server has an mtbf of 4,000 hours (about six months), a fault-tolerant server will have an mtbf of 40,000 hours (about five years). The availability parameters for our example then become

| | |
|------|----------------|
| n | = 2 |
| a | = .9999 |
| MTFO | = 3 seconds |
| mtbf | = 40,000 hours |

The resulting failure probabilities and system availability are

| | |
|-------|-------------------------------------|
| F_d | .00000001 (1×10^{-8}) |
| F_f | .000000021 (2.1×10^{-8}) |
| F | .000000031 (3.1×10^{-8}) |
| A | .999999969 (7+ 9s) |

In this case, failover times are much more significant in the active/active system. However, with the improved nodal availability, a system availability of over seven 9s can be achieved.

Summary

Though cluster and active/active architectures have a great deal in common, there is a tremendous difference in the availability that they offer. An active/active system can be an order of magnitude more reliable than a cluster. This is primarily due to the rapid failover times that can be achieved by active/active systems as compared to clusters.

A second contributor is that only the users at a failed node in an active/active system are affected by a failover, but all of the users in a cluster are affected by a failover.

In our example using industry standard servers, a cluster can achieve almost five 9s availability (about five minutes of downtime a year on the average). An active/active system using the same hardware can achieve almost six 9s availability (about 30 seconds of downtime per year on the average). If fault-tolerant nodes are used, an active/active system can achieve over seven 9s of availability (about three seconds of downtime per year on the average).