the Availability Digest

What's That Nerd Logo?

October 2006

It's the availability equation, of course, though admittedly somewhat stylized. It gives the probability of failure of a redundant system with one spare. The system will survive the failure of any one component; but if two components fail, the system is down.

The first "f" represents the number of ways that two components can fail. The second "f" represents the probability that any one of the components will fail. Therefore, the probability that two specific components will fail is f^2. The probability that the system will fail is the probability that two specific components will fail multiplied by the number of ways that two components can fail, or ff^2, if you will forgive our stylization.

For instance, if a cluster has four nodes and can survive the failure of any one node, there are six ways that two nodes can fail (count them). If the probability of failure of a node is 0.5%, the probability of failure of the cluster is 6 x .005 x .005. This works out to be .00015, or .015% (a little less than four 9s for those of you who are into 9s).

We thought that this would be a handy reminder of the search for continuous availability. We want the probability of failure to be as close to zero as possible. With today's active/active technology, six to seven 9s is being attained in real practice. This equates to average downtimes of 3 to 30 seconds per year.

© 2006 Sombers Associates, Inc., and W. H. Highleyman
www.availabilitydigest.com