Subject: Re: DRBD?
Posted by cdevidal on Tue, 01 Aug 2006 12:40:36 GMT
View Forum Message <> Reply to Message

wfischer wrote on Tue, 01 August 2006 04:24The second reason is that in a active-passive configuration you can get aware of performance bottlenecks soon enough. We had a for example a cluster running, that ran Apache on node1 and MySQL on node2 (without any virtualization). When we started the project, every machine had 1,5 GB RAM. Apache needed about 500 MB, and also MySQL needed about 500 MB. After some time we discovered that Apache now needs 1 GB, and also MySQL consumes 1 GB of RAM - so if a failover would have happened the remaining cluster node would have started swapping and get very slow (in fact so slow, that it would have seemed that the cluster is down)
When you run all services on only one node, you can sooner discover those performance bottlenecks (actually before a failover happens) - and enlarge e.g. RAM like in this case.
Evan Marcus and Hal Stern have a very interesting discussion about why to use active/passive and what to answer to management when they ask: "how can I use the standby server?"

I'm glad I talked to someone with experience

Good info, thank you.

Two questions:
1.) Did you actually perform a failover and observe it to be so slow because it was swapping?
2.) So then am I to understand that in theory load balancing and high availability aren't in conflict but in practice they are? For example OpenSSI, which gives you high availability + load balancing, but if you just have two nodes and every service fails over to the first node it gets to be so slow you might as well not have anything at all.

In other words, are load balancing and high availability mutually exclusive not in theory but in practice, at least for two nodes?