

Sep 13
2011

File Servers – The Business Case for High Availability

Introduction

You have probably heard of high availability transaction processing servers. You have most likely read about the sophisticated systems used by the airlines to sell tickets online. They have to be non-stop because downtime translates to lost orders and revenue. In this article I will discuss the economics of using non-stop technologies for everyday applications. I will show that even ordinary file sharing applications can benefit from inexpensive Linux based Pacemaker clustering technology.

Availability Goal

What is our availability goal? Our goal should be to take prudent and cost effective measures to reduce computer downtime to nil in the required service window. I'm not talking about 99.999 % (five 9s) up time. This is the popular (and very expensive) claim made by high availability vendors. I'm talking about maintaining enough up time to service the application. Take a simple example, for office document preparation the service time window is office hours (9-5). The rest of the time the desktop PCs can be turned off, nobody is there to operate them anyway. You only need the PCs for 5 days a week for 8 hours a day or for 2080 hours per desktop PC per year. This translates into an up time requirement of 24 percent. Ideally you want the desktop PCs to be available all the time during office hours but are willing to give up availability for routine maintenance and for the infrequent breakdowns that may occur only once per workstation every five years or so. Perhaps you have two spare desktop PC workstations for every 100. This extra capacity allows your office workers to resume their work on a spare while their workstation is being repaired. In this example the cost of maintaining adequate availability is the cost of maintaining two spare desktop PCs. You might adjust this cost to account for real world conditions at the work site. Wide swings in operating temperature or poor quality electricity supply, might dictate that you increase the number of spare PCs. Sounds like a low stress, straightforward availability solution.

Network Effects

The problem gets more complicated when the desktop PCs are networked together and all the documents are stored on a central file server rather than on each workstation's hard drive. There is a multiplicative effect. If the file server is not available then all 100 document processing PCs are rendered unavailable. Then you have 98 (remember the 2 spares from above) workers being paid but not producing documents. A failure during office hours can become expensive. One hour of downtime can cost as much as \$1500 in lost worker wages. A day of downtime can cost \$12,000 of lost worker wages. How long will it take for a hardware repair person to travel to your site? How long will it take for spare parts to arrive? How long will it take the repair person to replace the parts? How long will it take for damaged files to be replaced from backup by your own people? A serious but not unlikely failure can take several days to be completely resolved. Its not unreasonable to assume that such a \$24,000 failure can occur once every 5 years. This is a very simple example. We are not talking about a complicated order-entry or inventory control system.

We are talking about 98 office workers saving files to a central file share so that they can be indexed and backed up.

The Effects of Time

I'm going to add another wrinkle to our office document processing example. This file sharing setup has been in use for 4 years. Time flies. The hardware is getting old faster than you realize. Old hardware is more likely to fail. It has been through more thunderstorms, more A/C breakdowns, people knocking the server by accident and all that. You've been noticing that your hardware maintenance plan is costing more every year. How long is the hardware vendor going to stock spare parts for your obsolete office equipment? Please forgive me for playing on your paranoia but the real world can be rude.

Time for an Upgrade

In this scenario you conclude that you are going to have to replace that file server soon. Its going to be a pain to migrate all the files to a new unit. I am going to have to upgrade to a new version of Windows server. How much is that going to cost? How much has Windows changed? If I am going to have to go to all this trouble, why not get some new improvement out of it. I know I can get bigger disks and more RAM (random access memory) for less money than I paid for the old server. Whoops. Windows is going to cost more. I have to pay a charge for every workstation attached to it. That CAL (Client Access License) price has gone up. I read something about high availability clustering in Windows. Enterprise Server does that. Wow. Look at the price of that! Remember that \$12,000 per day of downtime cost overhang? It's more of an issue now that you are dealing with an old system.

A Debian Cluster Solution

Enough of this already. Since I asked so many questions and raised so many doubts, I owe you, the reader, some answers. Debian Linux provides a very nice high availability solution for file servers. You need two servers with directly attached storage and also a third little server that can be little more than a glorified workstation. You need Debian Squeeze 64 bit edition that has the Pacemaker, Corosync, drbd and Samba packages installed for each server. The software is free. You pay for the hardware and a trustworthy Linux consultant who can set everything up for you. What you get is a fully redundant quorum cluster with fully redundant storage, multiple CPU cores on each node, much more RAM than you had before and much more storage capacity.

Here are hardware price estimates:

Tie breaker node: Two hard drives, 512MB RAM	\$500.00
Name brand file server node: 8 2TB SATA drives, 24GB RAM, 1 4 core CPU chip, 3 year on site parts and labor warranty.	\$6,000.00
Second file server node like above.	\$6,000.00
Misc parts for storage and control networks.	\$200.00
Total:	\$12,700.00

Each file server node has software RAID 5 and each node holds 14 terabytes of disk storage. Because it is completely redundant across nodes, total cluster storage capacity is 14 terabytes. Performance of this unit will be much better than the old unit. It effectively has 4 CPUs per file

storage node and much more RAM for file buffering. Software updates from Debian are free. You just need someone to apply the security patches and version upgrades.

The best feature is complete redundancy for file processing. In our file server example, any one of the nodes can completely fail and file server processing will continue. Based on the lost labor time cost estimates above, this system pays for itself if it eliminates 1 day of downtime in a five year period. You also have hardware maintenance savings of whatever the yearly charge is for your old system times 3 years because you get 3 years of warranty coverage on the new hardware. You have the consultant's charges for converting to the new system, but remember, you were going to have to pay that fee for a new Windows system as well.

Conclusion

I hope I have stirred your interest in Linux Pacemaker based clusters. I have shown a file server upgrade that pays for itself by reducing downtime. You also upgrade your file server's performance while reducing out of pocket expenses for software and hardware maintenance. Not a bad deal.

Posted by Laird Hariu in Debian, Eternally Regenerative Software Administration, Systems Management, Virtualization, 0 comments