;login:
THE MAGAZINE OF USENIX & SAGE
August 2002 volume 27 • number 4

inside:

USENIX & SAGE
The Advanced Computing Systems Association &
The System Administrators Guild

high availability

by Hobbit

Hobbit espouses a straightforward, no-frills approach to infrastructure design and risk mitigation, which he has brought into environments spanning small home networks to large ASPs. He currently develops methodologies and tools for secure, scalable network and host deployment. He is perhaps best known as the author of Netcat, a useful tool that has found its way into many open-source operating system distributions.

hobbit@avian.org

EDITOR'S NOTE: THIS ARTICLE WAS FORWARDED TO ME, SO I ASKED HOBBIT IF WE MIGHT PUBLISH IT. I FOUND IT QUITE PROVOCATIVE.

I find myself wondering why so many organizations get their panties in such a bunch over high availability and failover. Many organizations that insist on it arguably don't need it, at least in most people's sense of HA as it's sold into the industry today. Your average small-to-medium business can get by just fine without it, which is a good argument for them to use something like IPF that doesn't do HA but solidly serves all of their other needs. A sensible management strategy will go a lot further than any of these vaunted HSRP/VRRP toys.

What does full HA buy you? In the *very rare* event of total firewall failure, session states aren't lost and the load quickly shifts to the standby box. Whoo, whoo. Most people are just surfing the Web, making new connections over and over, and wouldn't notice if they lost a TCP state or two. People who lose an SSH session will simply blame Sprint's half-million dollar piece of crap in Pennsauken and reconnect after the network starts working again. Streaming media will likely pick right back up again. In general, users are much more resilient than any of the HA proponents give them credit for. For this large community of users, some form of non-transparent failover is *fine*, and I really wish the HA vendors would get off their high horse about how *everyone* needs 5 nines of uptime. They don't. Get over it.

Besides, in the event of "seamless" failover, the firewall administrator is unlikely to realize that it has happened either, and won't know that box A had some problem until box B finally fails too. At which point everyone is *hosed* because now someone has to go and rebuild *both* firewalls and, well, the config probably never got properly backed up because the admin thought, "Oh, well, it's all this swoopy HA stuff, so I don't have to really worry about backups." Oops.

Some from the financial sector might point out how when their networks go down, they start losing millions of dollars per minute. (Losing to where, I wonder, if the bits representing value cannot be transferred in the first place?) Okay, so maybe they need seamless HA for their current business model, but let us also consider that if the financial industry has painted itself into a corner badly enough to wholly depend on continual transfer of data, there exists a much deeper problem beyond the scope of this rant. If I had a buck for every stupid, unrealistic SLA that has been drafted without taking such contingencies into account, I wouldn't have to worry much about finances anymore.

What is possibly a much more realistic (and cheaper) strategy is the warm or cold standby – a second box configured just like the first one, ready to go any time, with "failover" involving a junior NOC monkey walking out to the right rack and moving two or three clearly labeled wires. Then someone knows that box A has failed and can unrack it for repairs at leisure. Box B will likely hold up the load for the next year and a half without a hitch, so the priority of fixing box A becomes very flexible. Maybe it was time to upgrade to a faster machine for that spot anyways, eh?

I'm sure that more "automated" failover can be done with IPF and a little bit of scripting to run around and ifconfig interfaces and ping a few surrounding devices with the hope of forcing ARP and switching tables to update. The commercial HA products often have trouble convincing neighboring switches and things that no, this MAC address is over *here* now, really, at which point that multi-thousand-dollar HA setup still needs a swift kick. A successful semi-automated failover combined with user

resiliency gives you the same *long-term* net effect as the much more expensive "HA product." And a closer look inside a lot of HA products often reveals a nest of grotty shell scripts that do just that too, so don't be fooled.

So if you have an HA setup, go figure out how many times over the last year it has done its job in such a way that it even came close to paying for itself. Include all the factors such as administrative overhead, having to learn about the product, the engineering time spent integrating it into your network environment in a way that genuinely works, the effect of network design compromises you may have had to make to do so, etc. Don't forget to consider the increased *risk* that the lame JavaScript-laced HA management GUI has caused you to bear all that time, since it required you to open more avenues into the firewalls themselves. And were you ever able to get your run-of-the-mill NOC guy to understand the thing?

If your "HA" setup consists of swapping the hard drive and network cards into a new chassis and powering back up, be thankful and count those nice crisp hundreds you saved. And if you've read this far, go back up your nice simple little text-based rule set, if it's been a while, just in case it's the drive that craps out instead. But you might wait another two years for that to actually happen.

USENIX and SAGE Need You

People often ask how they can contribute to the our organizations. Here is a list of needs for which we hope to find volunteers (some contributions reap not only the rewards of fame and the good feeling of having helped the community, but authors also receive a small honorarium). Each issue we hope to have a list of openings and opportunities.

The SAGEwire and SAGEweb staff are seeking:

- Interview candidates
- Short article contributors (see *http://sagewire.sage.org*)
- White paper contributors (for topics like these):

Back-ups	Emerging technology	Privacy
Career development	User education/training	Product round-ups
Certification	Ethics	SAGEwire
Consulting	Great new products	Scaling
Culture	Group tools	Scripting
Databases	Networking	Security implementation
Displays	New challenges	Standards
E-mail	Performance analysis	Storage
Education	Politics and the sysadm	Tools: system

- Local user groups: If you have a local user group affiliated with USENIX or SAGE, please mail the particulars to *kolstad@sage.org* so they can be posted on the web-site.

;login: is seeking attendees of non-USENIX conferences who can write lucid conference summaries. Contact Tina Darmohray, *tmd@usenix.org* for eligibility and remuneration info. Conferences of interest include (but are not limited to): Interop, SOSP, O'Reilly Open Source Conference, Blackhat (multiple venues), SANS, and IEEE networking conferences. Contact *login@usenix.org*.

;login: always needs conference summarizers for USENIX conferences too! Contact Alain Hénon *ah@usenix.org* if you'd like to help.