

Improving BIND Load Time

Neil Schelly

Director of Operations Improvement

Nameserver Startup Time


- Load times over 45 minutes
 - Roughly 110,000 primary zones
- Planning the large-scale ingestion of more
 - Up to 600,000 more zones
- Planning for future growth

Ideas for Improvement

- Mount configuration data with noatime
- Upgrade hardware
- Configuration and zone files on SSD
- BIND 9.8.1 Patch
 - <http://www.isc.org/community/blog/201107/major-improvement-bind-9-startup-performance>


Mount Option: noatime

- Reduced disk writes during configuration load




Hardware Platform

- Legacy servers
 - Dual Intel Xeon 5130 dual-core processors
 - 4GB memory
 - SATA RAID1 mirror for config/zone storage
- New Hardware
 - Dual Intel Xeon X5650 six-core processors

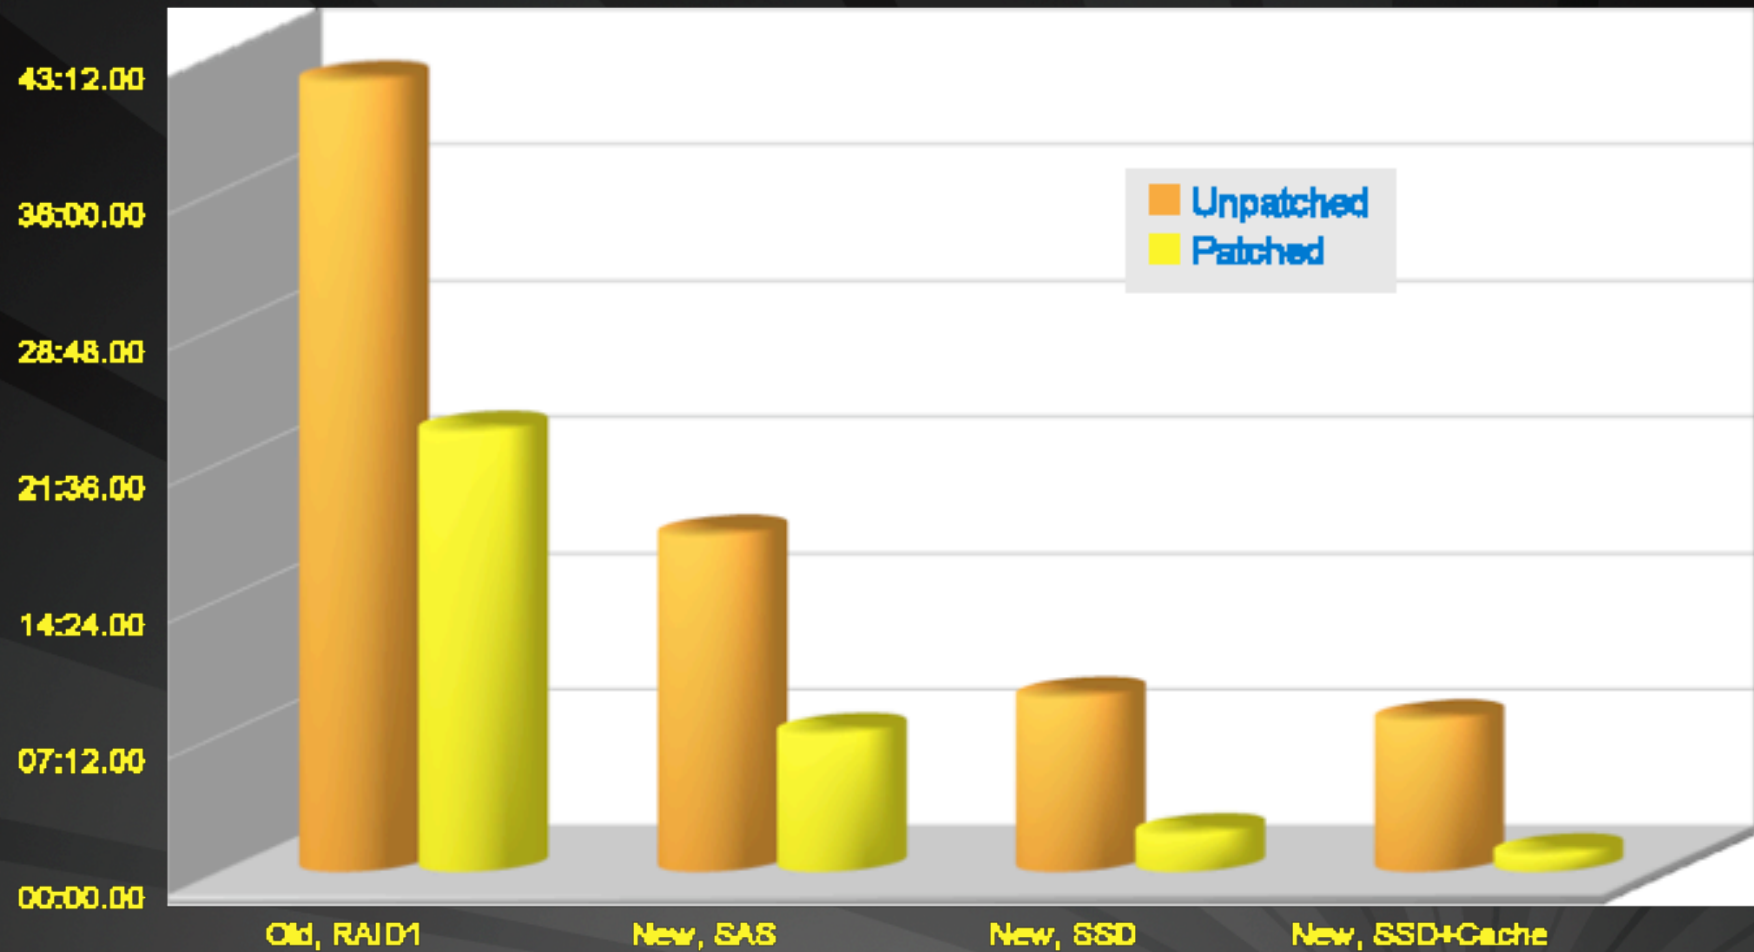


Spindles vs. SSD

- SSDs are fast.
- You should use SSDs wherever possible.
- I like SSDs.



BIND 9.8.1 Patch



Disk Cache on the New Hardware

- 12GB of memory enough for BIND and disk cache in memory
- Initial load time
 - Much slower, better measure of real world expectations.
- Successive/repeated load time
 - Better test nearly eliminating disk IO from limiting factors.
 - Best illustration of the BIND load time patch.

Conclusions

- Mount with noatime: 11% (25% w/ patch)
- Upgrade hardware: 58% (69% w/ patch)
- SSD storage: 48% (73% w/ patch)
- Overall: 96% improvement