

Linpro AS

The leading Norwegian provider of
Linux and OSS solutions and services



- Norway's leading provider of solutions and services based on open source software and Linux
- More than 10 years of experience in consulting, development, integration and IT operations
- Over 50 courses, including certification courses from Red Hat, Novell and MySQL
- Linpro hosts and / or administrates over 120 Escenic sites, including over 50 Norwegian daily newspapers

- Senior Software Developer
 - Multiframe, Varnish
 - SIG for C, C++, kernel development and embedded systems
- FreeBSD committer since 1998
 - OpenPAM, OpenSSH, libfetch, pseudofs...

Turbo-charging your web site with Varnish

Dag-Erling Smørgrav
Senior Software Developer
Linpro AS

2007 Escenic Customer & Partner Days

- HTTP accelerator, aka. reverse proxy
 - Place in front of your CMS to improve performance
- Designed for modern computers, modern operating systems and modern workloads
- Very flexible configuration system
 - Cache policy defined by VCL (Varnish Configuration Language) scripts
- Runs on FreeBSD, Linux
 - Soon also Solaris
- Open Source

- Developed by Linpro and Danish independent consultant Poul-Henning Kamp
- Jointly owned by Linpro and VG
 - an independent foundation is being set up
- Initial development phase (2006) financed by VG
- Current development phase (2007) sponsored by Escenic and API
- Naturally, Linpro offers Varnish-related services



- Optimized for 64-bit
 - Also supports 32-bit
- Optimized for multi-core / multi-processor
 - Also supports single-core / single-processor
- Work *with* the kernel, not *against* it
 - Squid tries to second-guess the VM system
- Take advantage of advanced I/O facilities
 - epoll, kqueue, accept filters...
- Minimize the amount of work done in the fast path
 - Delegate non-essential work to support applications

How fast do you think Varnish is?

- Short answer: very
- Long answer: we don't really know, but:
 - < 20 system calls per request on average
 - < 20 μ s to handle a cache hit on reasonable hardware
 - Reached 700 Mbps on a single server (dual Opteron blade with GigE)
 - Never been Slashdotted yet...

1000 req/s @ 100 Mbps

Sample	Min	Max	Median	Average	Std. dev.
Full	12.9	3001.0	16.2	71.1	338.5
90% fractile	12.9	26.0	15.9	16.3	1.7

All numbers are in microseconds

- “Preview release” on the way to 2.0
- Bug fixes
 - HTTP protocol compliance
 - Stability and error handling
- Minor features
 - Support for content negotiation
 - Header / URL rewriting

- Partial ESI support
- Load balancing
- Administration and monitoring tools
 - Webmin, Nagios, Munin
- Escenic integration

- More control over expiry policy
 - Keep expired content in case the backend fails
 - Prefetch frequently requested objects
- Traffic replay tool
- Improved HTTP conformance
- Regression test suite
- More (and better) documentation

- Compression
- Cache persistence
- Multiple cache files
- Raw devices

And perhaps



- Direct SRun protocol support
- Direct AJP protocol support
- SSL / TLS termination
- Your feature here

and, hopefully, some answers as well

- <http://varnish.projects.linpro.no/>
- <http://www.varnish-cache.org/>
- http://en.wikipedia.org/wiki/Varnish_cache
- <http://people.freebsd.org/~phk/>
- <http://www.linpro.no/>
- <http://www.vg.no/>
- <http://www.api.no/>
- <http://www.escenic.no/>