

# Varnish

---

- the http accelerator

Poul-Henning Kamp

[phk@FreeBSD.org](mailto:phk@FreeBSD.org)

# CMS systems

---

- **GET / HTTP/1.1**
  - Hang on, I need to look up a few hundred things in my database and then do a lot of editing with some badly written software.

# CMS systems

---

- **GET / HTTP/1.1**
  - Hang on, I need to look up a few hundred things in my database and then do a lot of editing with some badly written software.
- **HEAD / HTTP/1.1**
  - Hang on, I need to look up a few hundred things in my database and then do a lot of editing with some badly written software, *and then I will throw the result away.*

# CMS systems are SLOW!

---

- Complex content generation process
- Single database prevents clustering
- Expensive software ditto.
  
- => Need for server side caching.
  - Apache
  - Squid
  - \$CALL

# Apache as cache

---

- Not what Apache is built for.
- Not what Apache is good at.
- Square peg, round hole.

# Squid is a cache...

---

- ...built for client side caching.
- Lots of unwanted functionality.
  - Authentication.
  - FTP
- 1980 software design.
- Fragile and buggy.
- Wrong configuration features.
- Wrong policy decisions.

# \$CALL

- 
- Akamai will happily take your money.
  - But you loose control.
  - You're stuck with their service.

# dictionary:Varnish

---

- tr. v. var·nished, var·nish·ing, var·nish·es
  - 1. To cover with varnish.
  - 2. To give a smooth and glossy finish to.
  - 3. **To give a deceptively attractive appearance to.**



# Varnish Cheat-Sheet

---

- Web-acceleration for slow CMS systems
- Narrow focus on server side speedup
  - No FTP etc.
  - Content provider features
- High Performance
  - 32 & 64bit, large RAM, sendfile, accept filters
  - SMP/Multicore friendly architecture
  - 2006 software design
  - 11 syscall + 7 locks per cache hit

# Performance Design

---

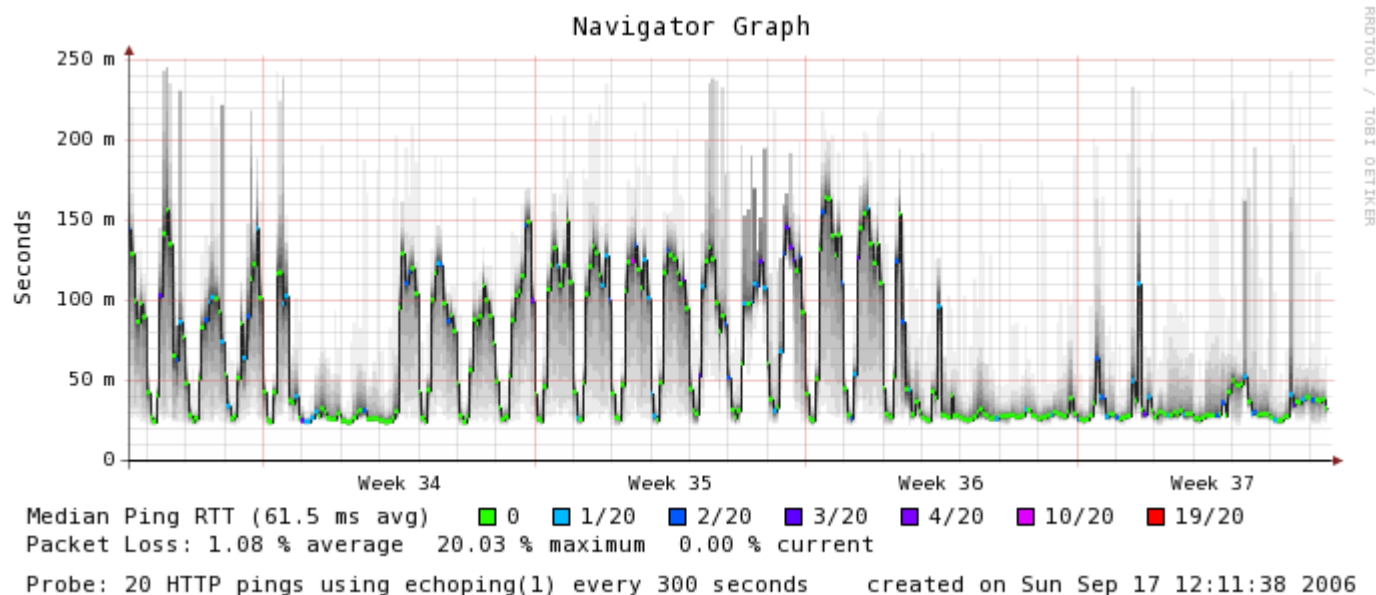
- Don't copy data if we can avoid it
  - Avoid text-processing headers
- Convert Chunked encoding to Direct
- Also cache "cannot be cached" info
- Maximize session usage
  - Pass-through mode understands chunked encoding etc.
  - Pipe mode for weird stuff (selected by VCL)

# Varnish is fast!

- 1000 cache hits @ 100 Mbit/sec:

Sample	Min	Max	Median	Average	Stddev
Full	12,9	3001	16,2	71,1	338,5
90% fractile	12,8	26	15,9	16,3	1,7

(all times are in microseconds)



# Content Provider Features

---

- Instant URL invalidation
  - Regexp matching
  - Easy Integration to CMS system
- Expiry time control (via VCL)
- Load/Situation mitigation (via VCL)

# OaM

---

- Cluster controller (version 2 feature)
  - Invalidate URL on all caches instantly
  - Change VCL on all caches simultaneously.
  - Aggregate statistics
- Command Line Interface
  - Remote access via ssh/mobile phone
  - Interrogation/configuration/investigation

# Varnish Config Language

---

- Simple domain specific language
  - Compiled via C language to binary
    - Transparantly!
  - Dynamically loaded
- Multiple configs loaded concurrently
  - Instant switch from one VCL to another.

# VCL example

---

```
if (client.ip in 10.0.0.0/8) {
    pass;
}

if (req.url.host ~ "cnn.no$") {
    rewrite req.url.host "cnn.no" "vg.no"
}

if (!backend.up) {
    if (obj.exist) {
        set obj.ttl += 10m;
        deliver;
    }
    switch_config "ohhshit";
}
```

# Logging

---

- Logging to shared memory
  - No slowdown for real workload
- Daemons tailing shm generate "real" output:
  - Apache format
  - Custom format
  - Realtime views



# “Varnishtop”

---

- Shows an real-time view of log records.
- “varnishtop -i rxurl”

```
342.49 /gfk/front/tipsvg.png
340.41 /sistenytt.html
326.53 /gfk/ann/n.gif
307.61 /annonser/telenor/menyknapp/svart/vg115x160_1.html
297.70 /include/global/art.js
276.73 /gfk/ann/ng.gif
273.84 /css/hoved.css
243.01 /
241.85 /annonser/telenor/menyknapp/svart/pil.gif
226.04 /css/kjaerevg.css
```

# "varnishhist"

- Shows response time as real-time histogram:



# Statistics

---

- Stored in shared memory
- Multiple programs can monitor, react to and & present stats data.
  - At practically no cost in cpu time.

# Varnish 1.0

---

- Sponsored by Verdens Gang
  - Norwegian Newspaper
- 500.000 Kr (\$85K)
- 22000 lines of code
  - 22 Kr/line (\$3.80/line)
- 1000 checkins
  - 500 Kr/checkin (\$85/checkin)

# Varnish 1.0 Cheatsheet

---

- [WWW.varnish-cache.org](http://WWW.varnish-cache.org)
  - Yes, we need to work on the web-page.
- Coded by: Poul-Henning Kamp
  - [phk@FreeBSD.org](mailto:phk@FreeBSD.org)
- Runs on: Any resonable modern UNIX
  - FreeBSD, Linux etc.
- Open Source
  - BSD license