

Tuning NGINX for high performance

Nick Shadrin

nick@nginx.com

All links on one page

shadrin.org/talks/

Twitter:

[@shadrin](https://twitter.com/shadrin)

[@nginx](https://twitter.com/nginx)

[@nginxorg](https://twitter.com/nginxorg)





About me

- Nick Shadrin
- Technical Solutions Architect with NGINX
- Based in San Francisco
- 15 years experience with web tech
- nick@nginx.com / [@shadrin](https://twitter.com/shadrin)



Agenda

- A basic NGINX configuration
- NGINX performance optimizations:
- Operating system-level optimizations
- Networking-level optimizations
- NGINX core optimizations
- Conclusions and questions

NGINX



“... when I started NGINX, I focused on a very specific problem – how to handle more customers per single server.”

- Igor Sysoev, NGINX creator & our founder

About NGINX, Inc.

- Founded in 2011, NGINX Plus first released in 2013
- VC-backed by enterprise software industry leaders
- Offices in San Francisco, London, and Moscow
- 750+ commercial customers
- 100+ employees






Web Scale

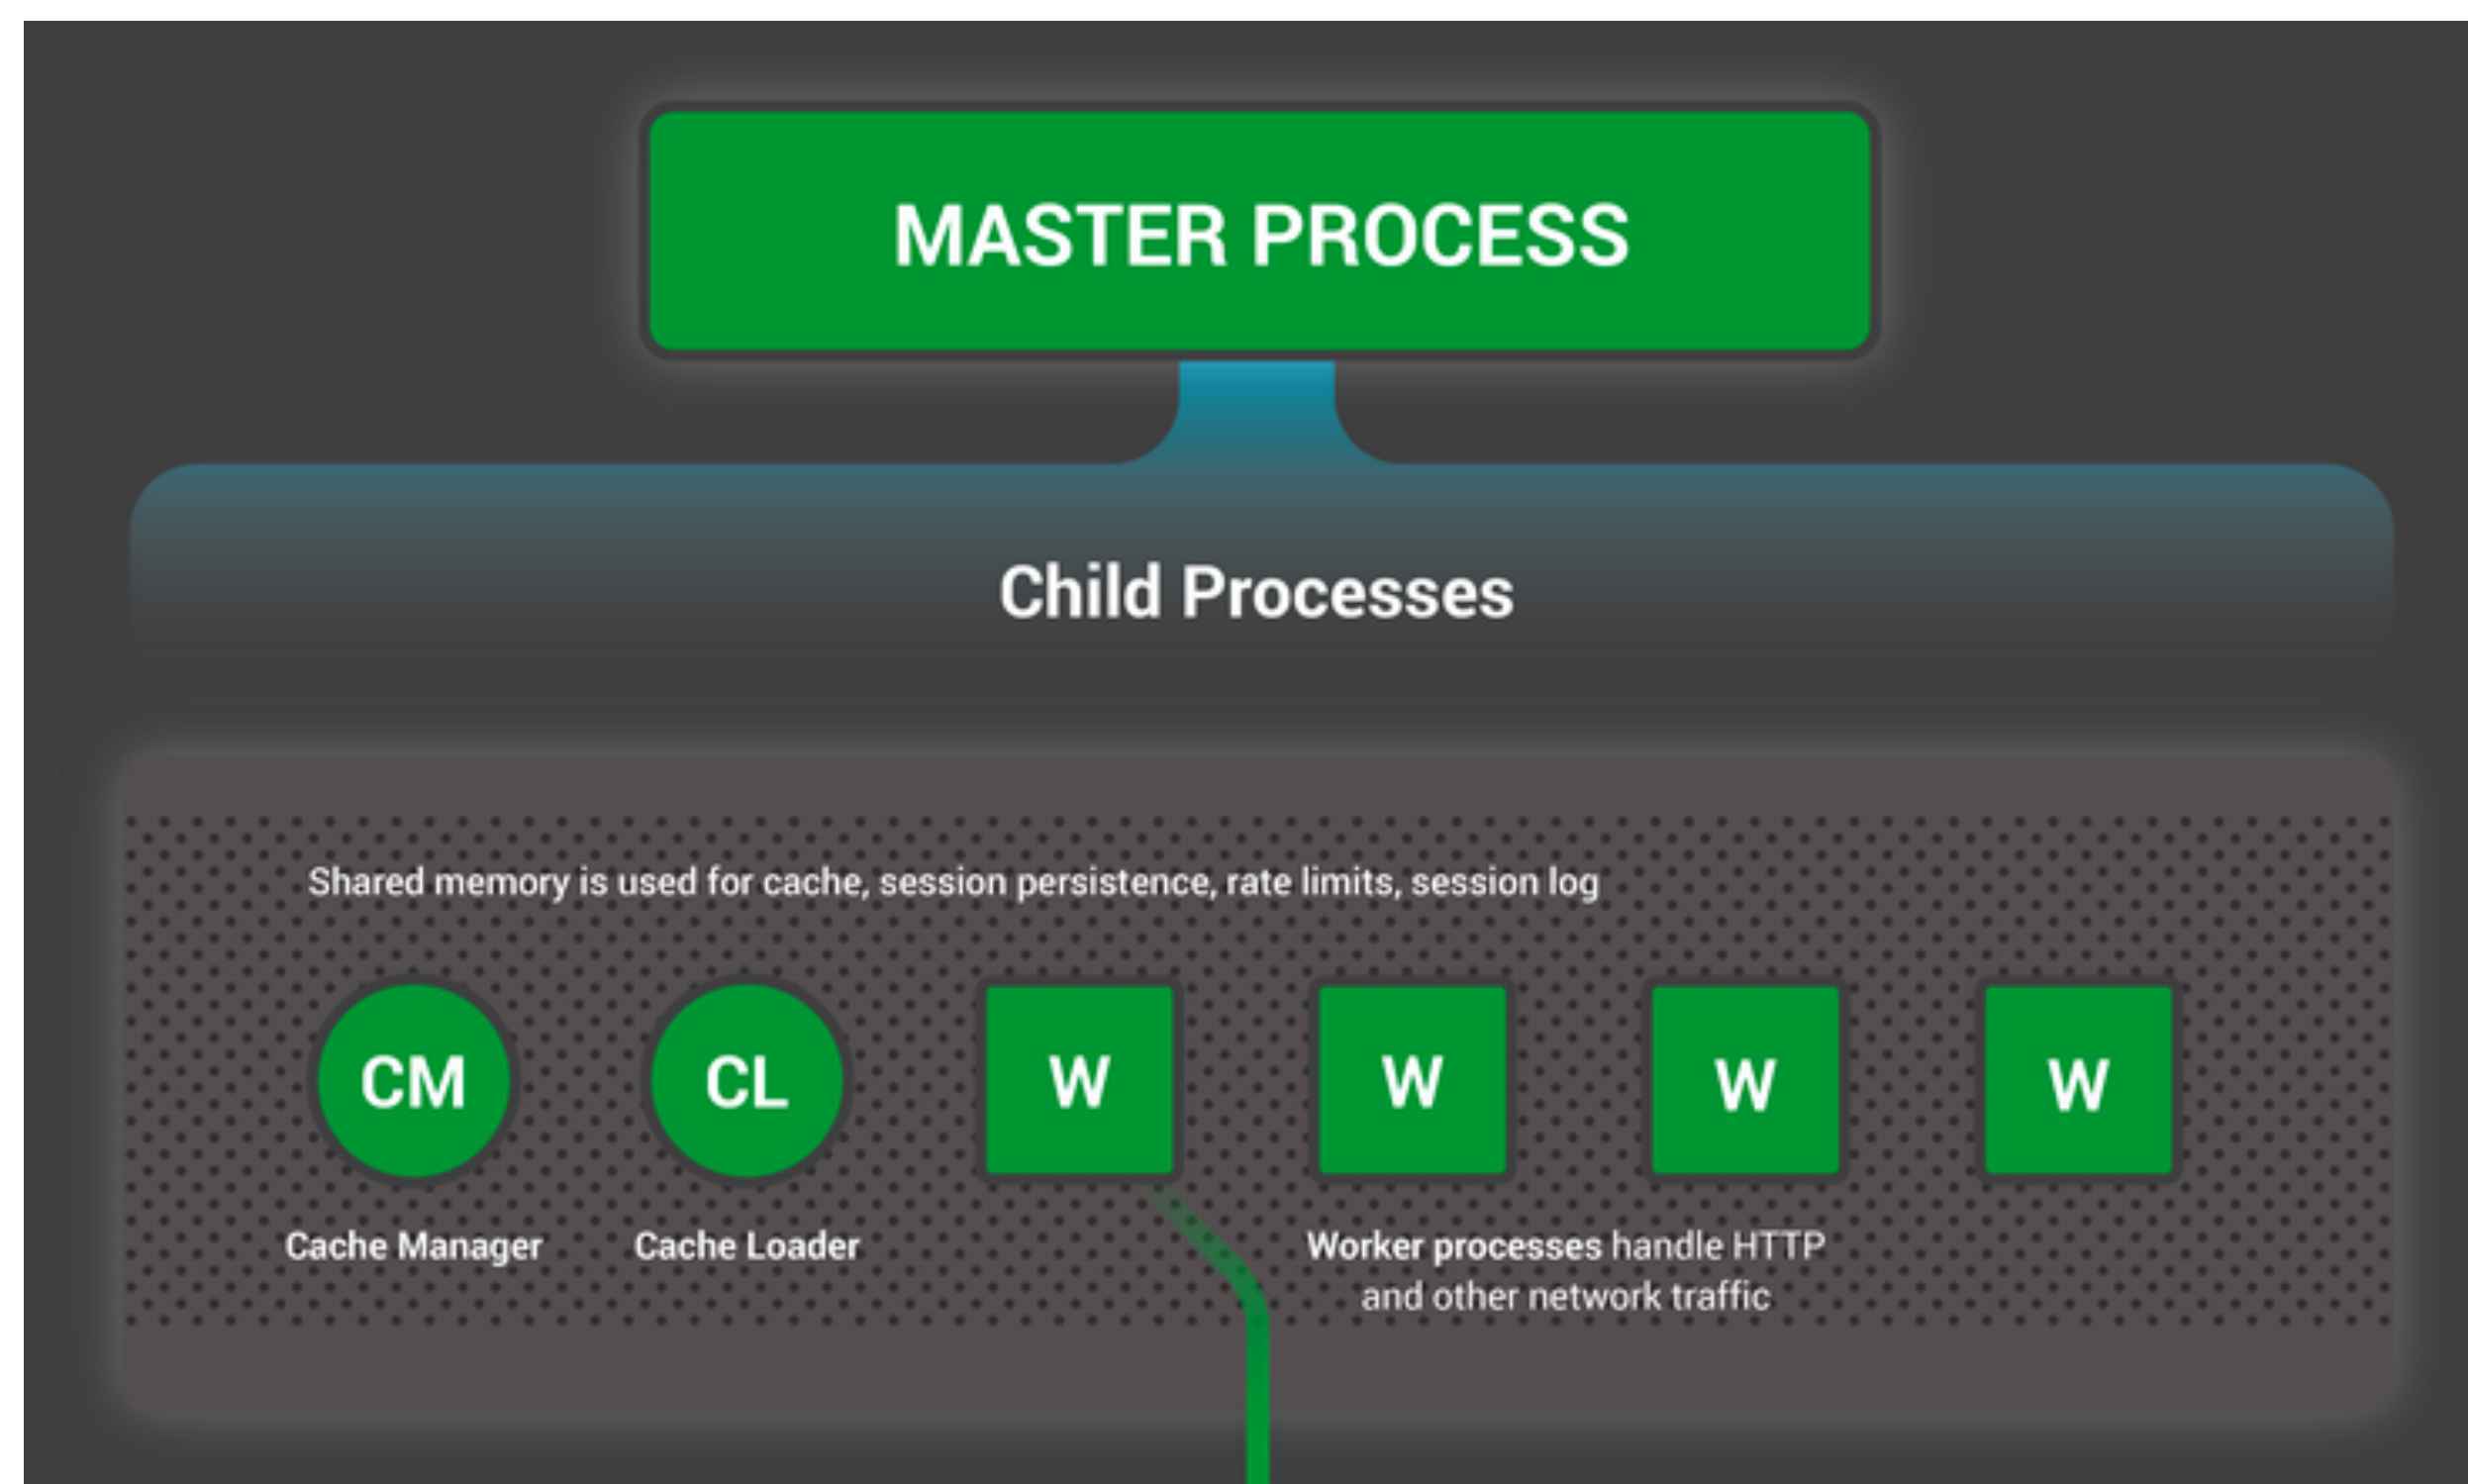
Architecture approach

- Design for scaling
- Segment microservices out
- Use caching and microcaching

Basic NGINX placement



Inside NGINX



<https://www.nginx.com/blog/inside-nginx-how-we-designed-for-performance-scale/>

<http://www.aosabook.org/en/nginx.html>

OS tuning

- `net.core.somaxconn`
- `net.core.netdev_max_backlog`
- `net.ipv4.ip_local_port_range`
- `sys.fs.file_max`
- `/etc/security/limits.conf`, nofile setting

See <https://www.nginx.com/blog/tuning-nginx/>



Overcoming ephemeral port exhaustion

- Increase local port range
- Split traffic across multiple IPs
- NGINX since 1.11.2 uses `IP_BIND_ADDRESS_NO_PORT` socket option when available

<https://www.nginx.com/blog/overcoming-ephemeral-port-exhaustion-nginx-plus/>

Minimal NGINX configuration

```
events {}

http {

    server {
        listen 80;
        location / {
            proxy_pass http://backend;
        }
    }

    upstream backend {
        server backend1.example.com:8080;
        server backend2.example.com:8080;
    }
}
```




NGINX Performance features



NGINX Core features

- Use correct number of **worker_processes**
 - auto
 - # of available CPU cores
- Increase **worker_connections**
- Increase **worker_rlimit_nofile**

NGINX Core Features (cont'd)

- Turn off accept_mutex:
accept_mutex off;
- Turn on Sendfile
sendfile on;
- Use thread pools if I/O needs offloading:
aio threads;

<https://www.nginx.com/blog/thread-pools-boost-performance-9x/>



HTTP Keep alive

- Keepalive connections allow to reuse the same TCP connection for multiple HTTP requests.
- For HTTP/1.1, no need to define anything, it's enabled by default on the frontend.
- Keepalives provide major performance benefit when used over SSL/TLS connections.




HTTP Keepalive: benchmark

- HTTPS with NO keepalive (worst setup)
- Plain HTTP
- HTTP/2 with SSL



HTTP/2 vs HTTP/1 vs HTTPS





HTTP Keepalive

- Keepalive on the Frontend:
keepalive_requests 100;
keepalive_timeout 75s;

HTTP Keepalive on the backend

Keepalive on the Backend:

```
server {  
    location / {  
        proxy_pass http://backend;  
        proxy_http_version 1.1;  
        proxy_set_header Connection "";  
    }  
}  
...  
upstream backend {  
    server example.com;  
    keepalive 32;  
}
```



HTTP Caching

- Microcaching with NGINX:
<https://www.nginx.com/blog/benefits-of-microcaching-nginx/>
- Cache placement strategies:
<https://www.nginx.com/blog/cache-placement-strategies-nginx-plus/>



HTTP/2


- Introduced in 2015 as a standard
- Based on Google's SPDY
- Includes major changes compared to HTTP/1:
 - Binary headers with HPACK
 - Multiple streams
 - Prioritization
 - Server Push



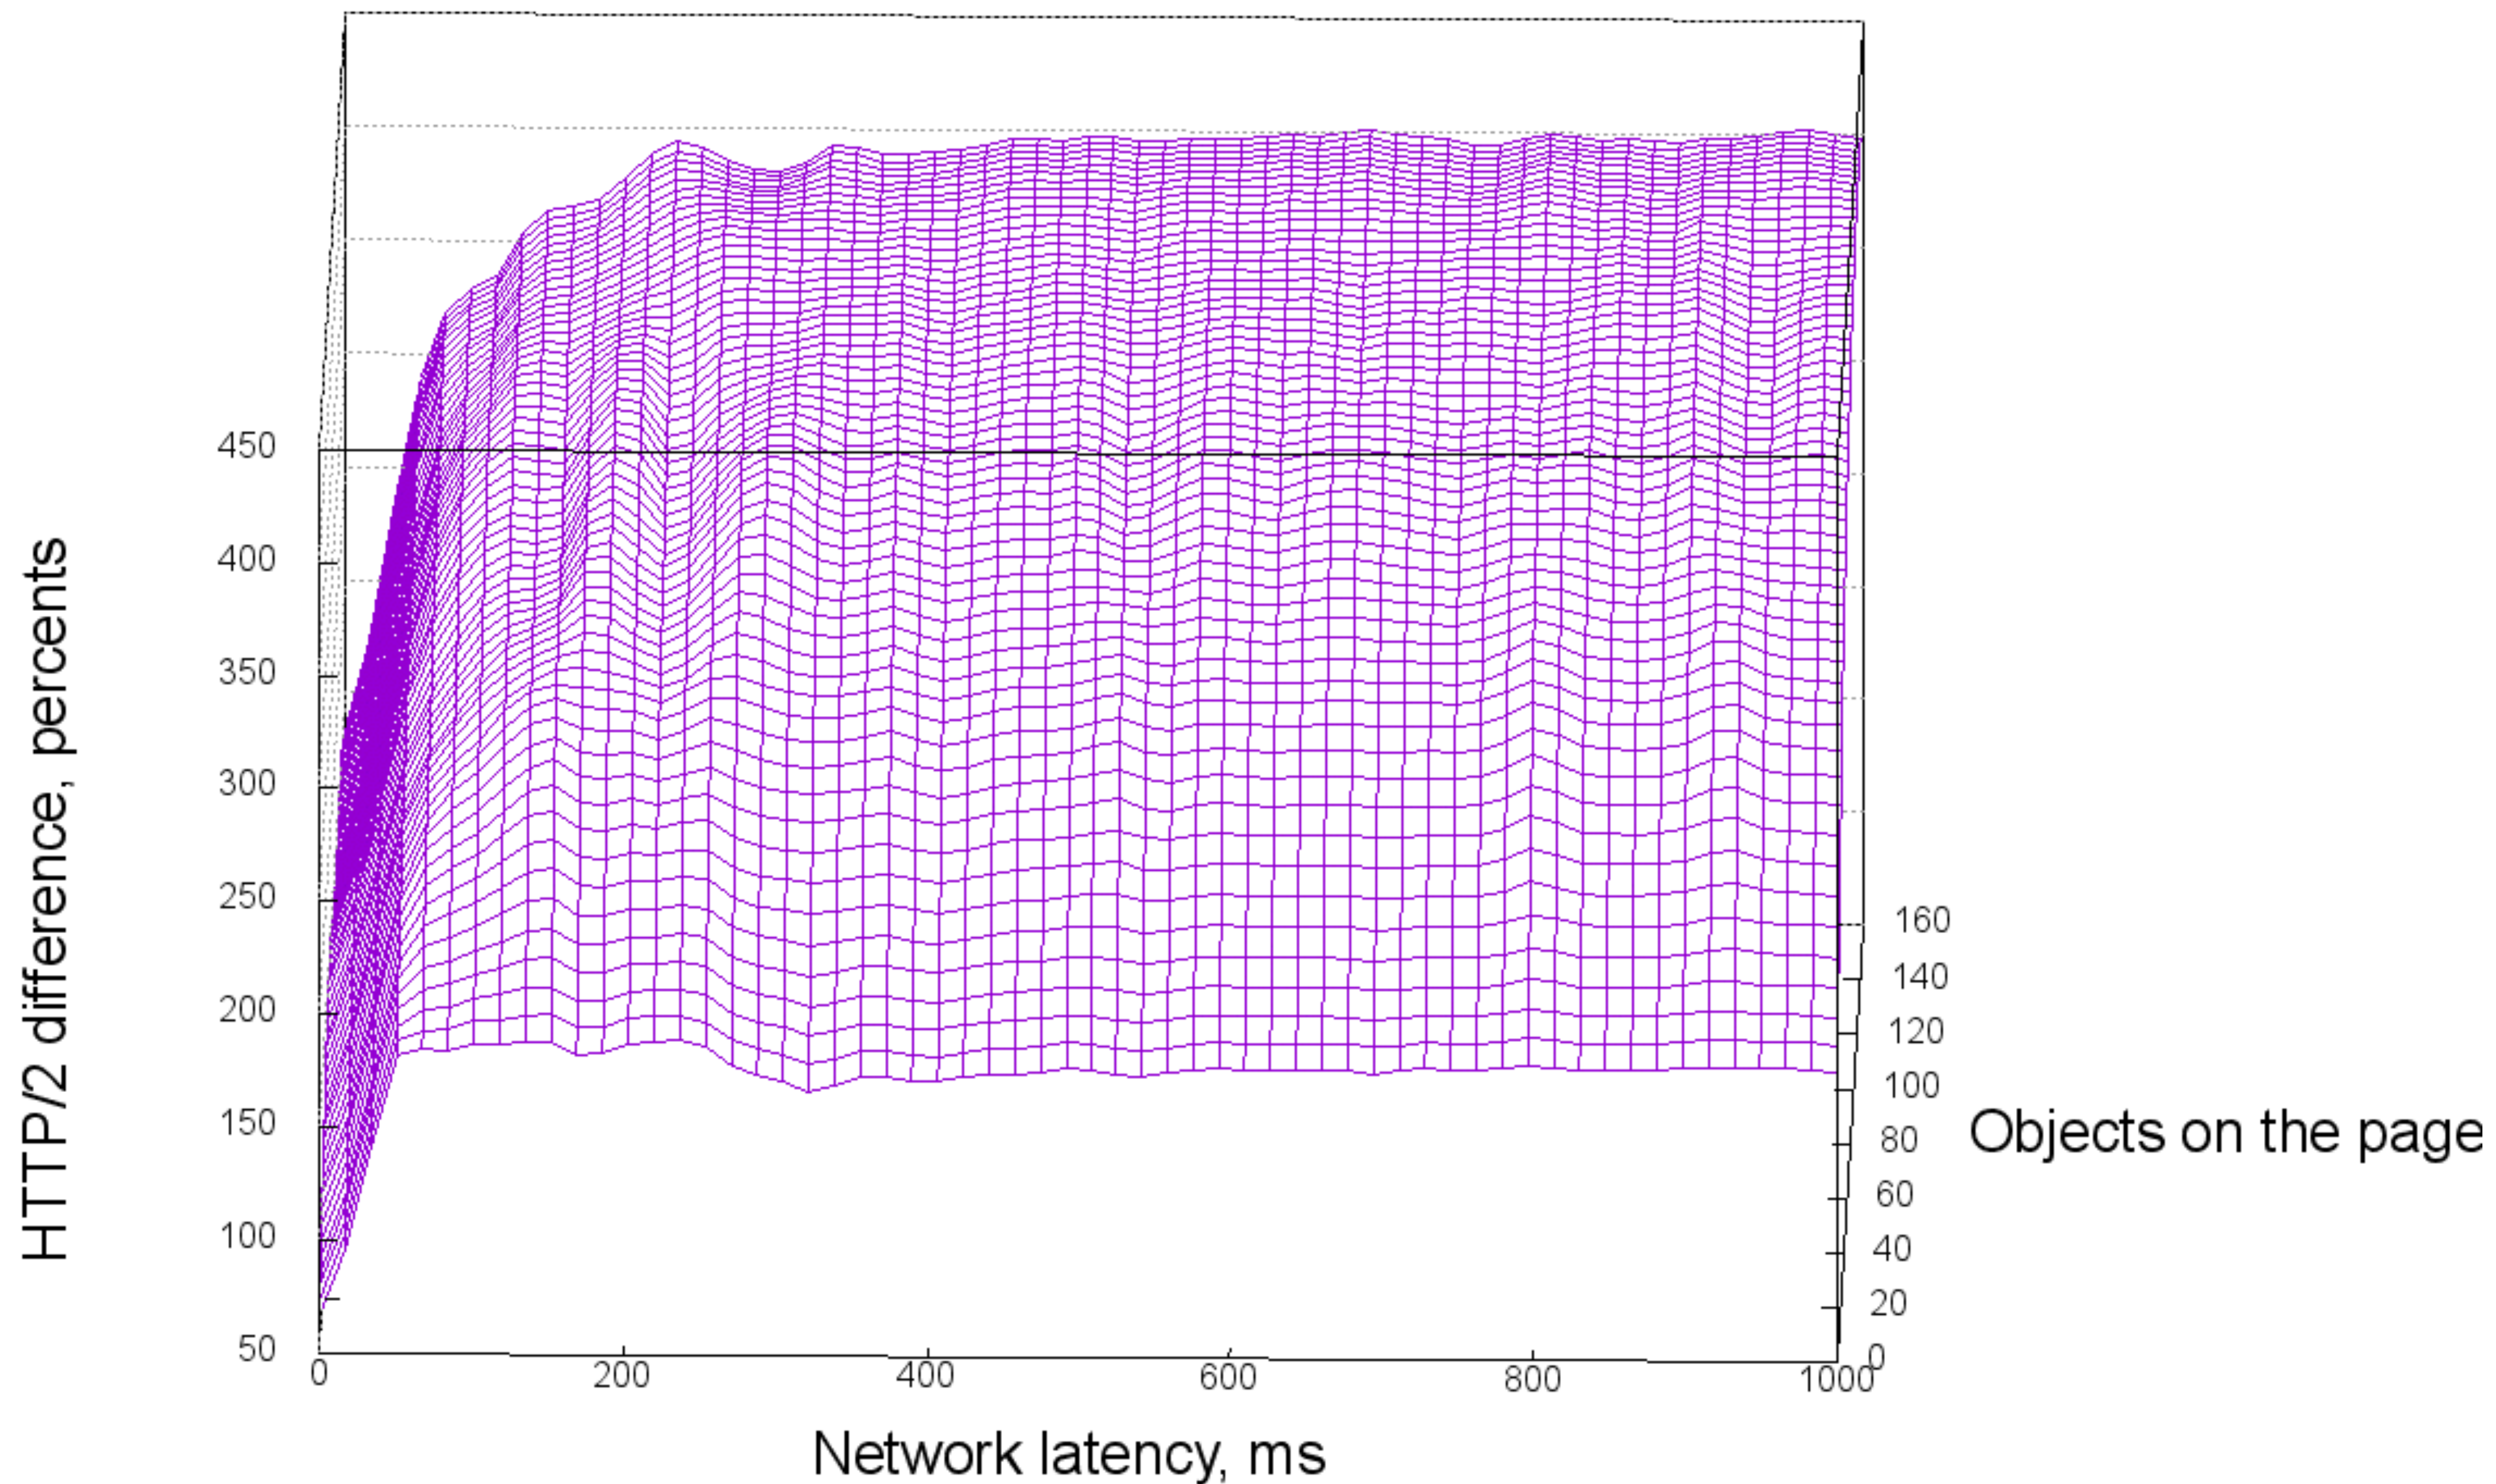
HTTP/2 benchmark

- NGINX 1.10.0
- Ubuntu 16.04
- Openssl 1.0.2
- Chrome Web browser
- Measuring full page reload



HTTP/2 vs HTTP/1/SSL, percentage performance increase



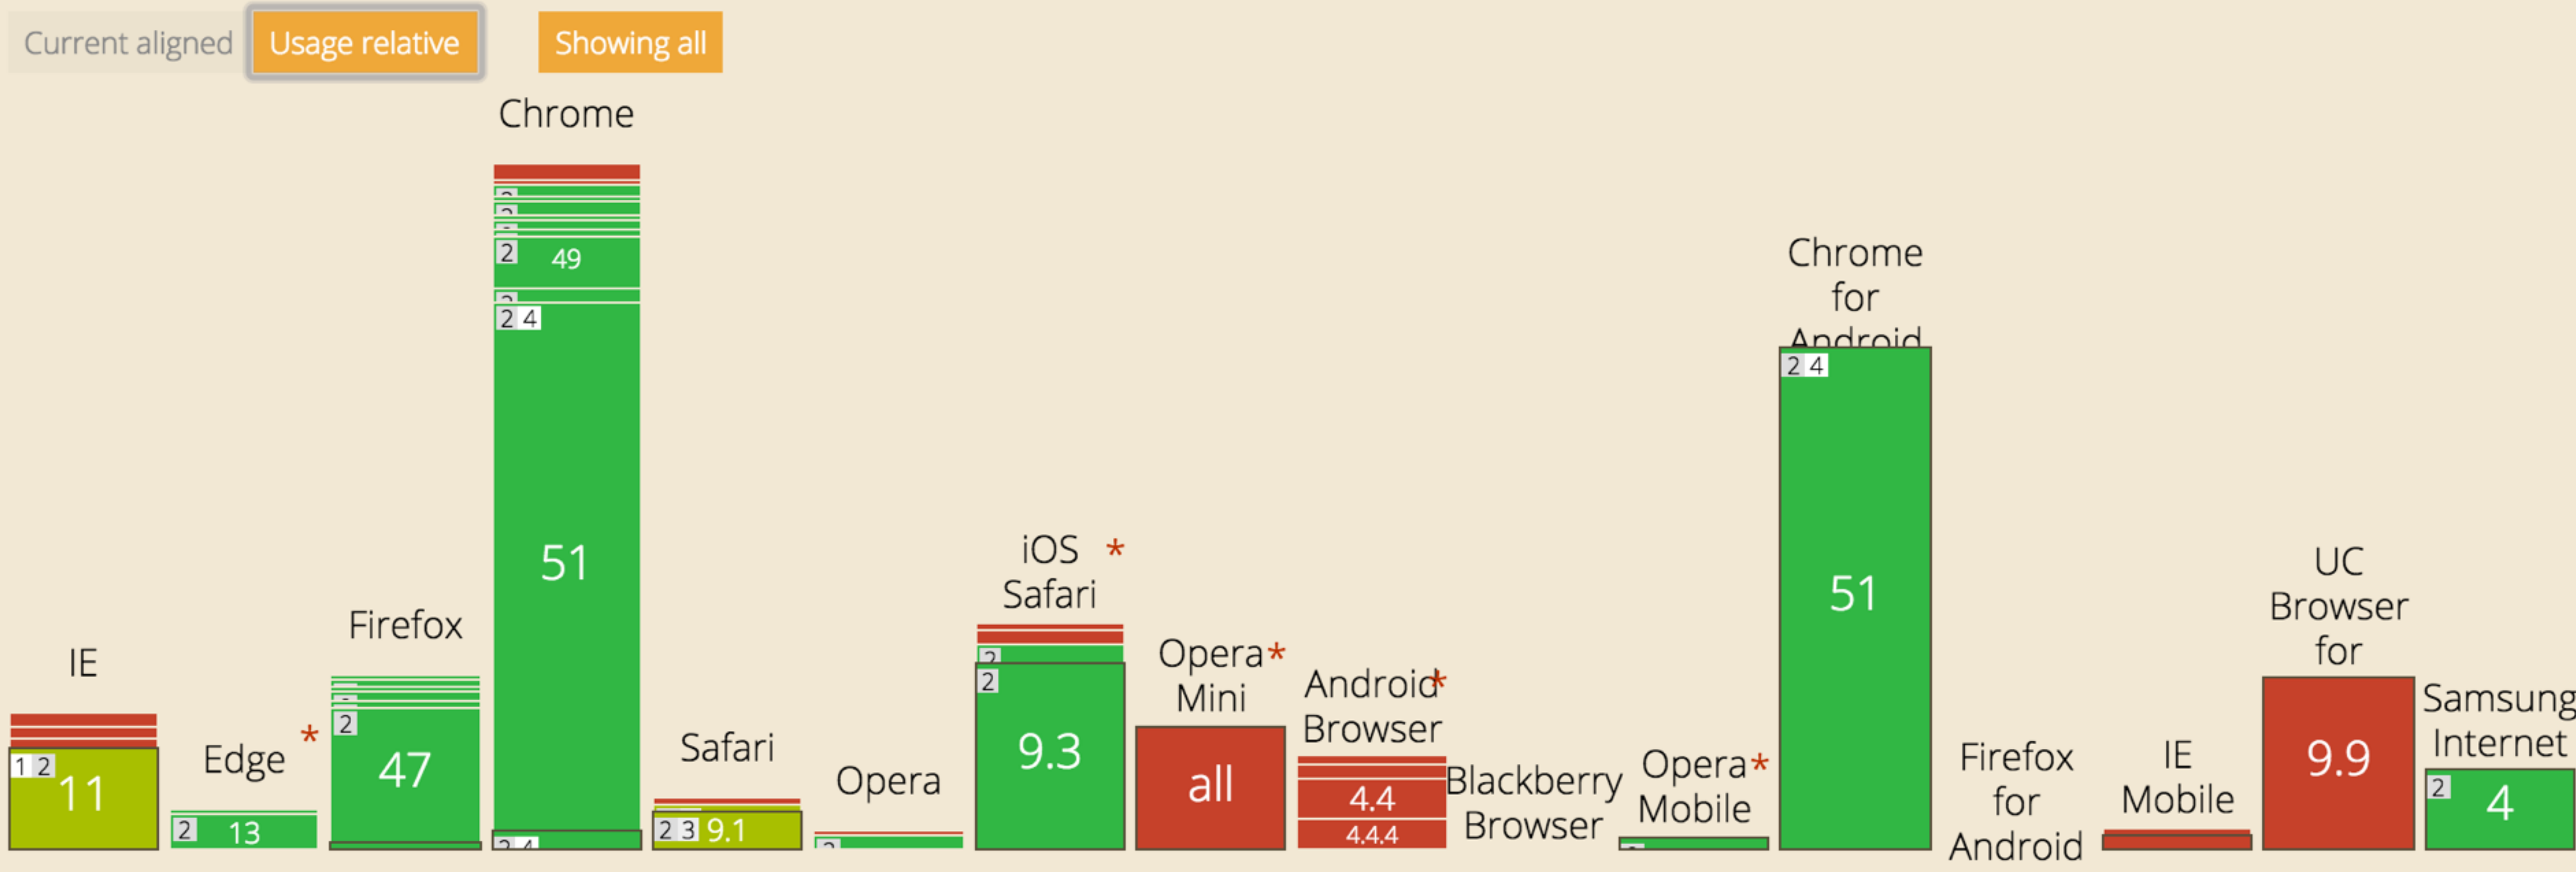
HTTP/2 latency divided by HTTP/1 latency



Some numbers

- 40ms / 50 objects:
HTTP/1: **510ms**
HTTP/2: **250ms**  **~2 times faster**
- 200ms / 100 objects:
HTTP/1: **4.0s**
HTTP/2: **1.1s**  **~4 times faster**

Networking protocol for low-latency transport of content over the web. Originally started out from the SPDY protocol, now standardized as HTTP version 2.



Screenshot: 2016-08-23, caniuse.com

Measure your results

- NGINX provides extensive logs with custom variables. Configure **log_format** with:
\$upstream_response_time
\$request_time
\$upstream_cache_status
- NGINX has simple set of metrics with stub_status module. Configure **stub_status**
- **NGINX Plus** provides more extensive metrics with Extended Status module
- **NGINX Amplify** is a free monitoring SaaS solution.

Sign up at amplify.nginx.com

The screenshot displays the NGINX Amplify dashboard interface. At the top, the navigation bar includes 'NGINX Amplify beta', 'Graphs', 'Dashboards', 'Reports', and 'Alerts'. The main content area is divided into three sections: 'SYSTEMS', 'PREVIEW', and 'GRAPH FEED'.

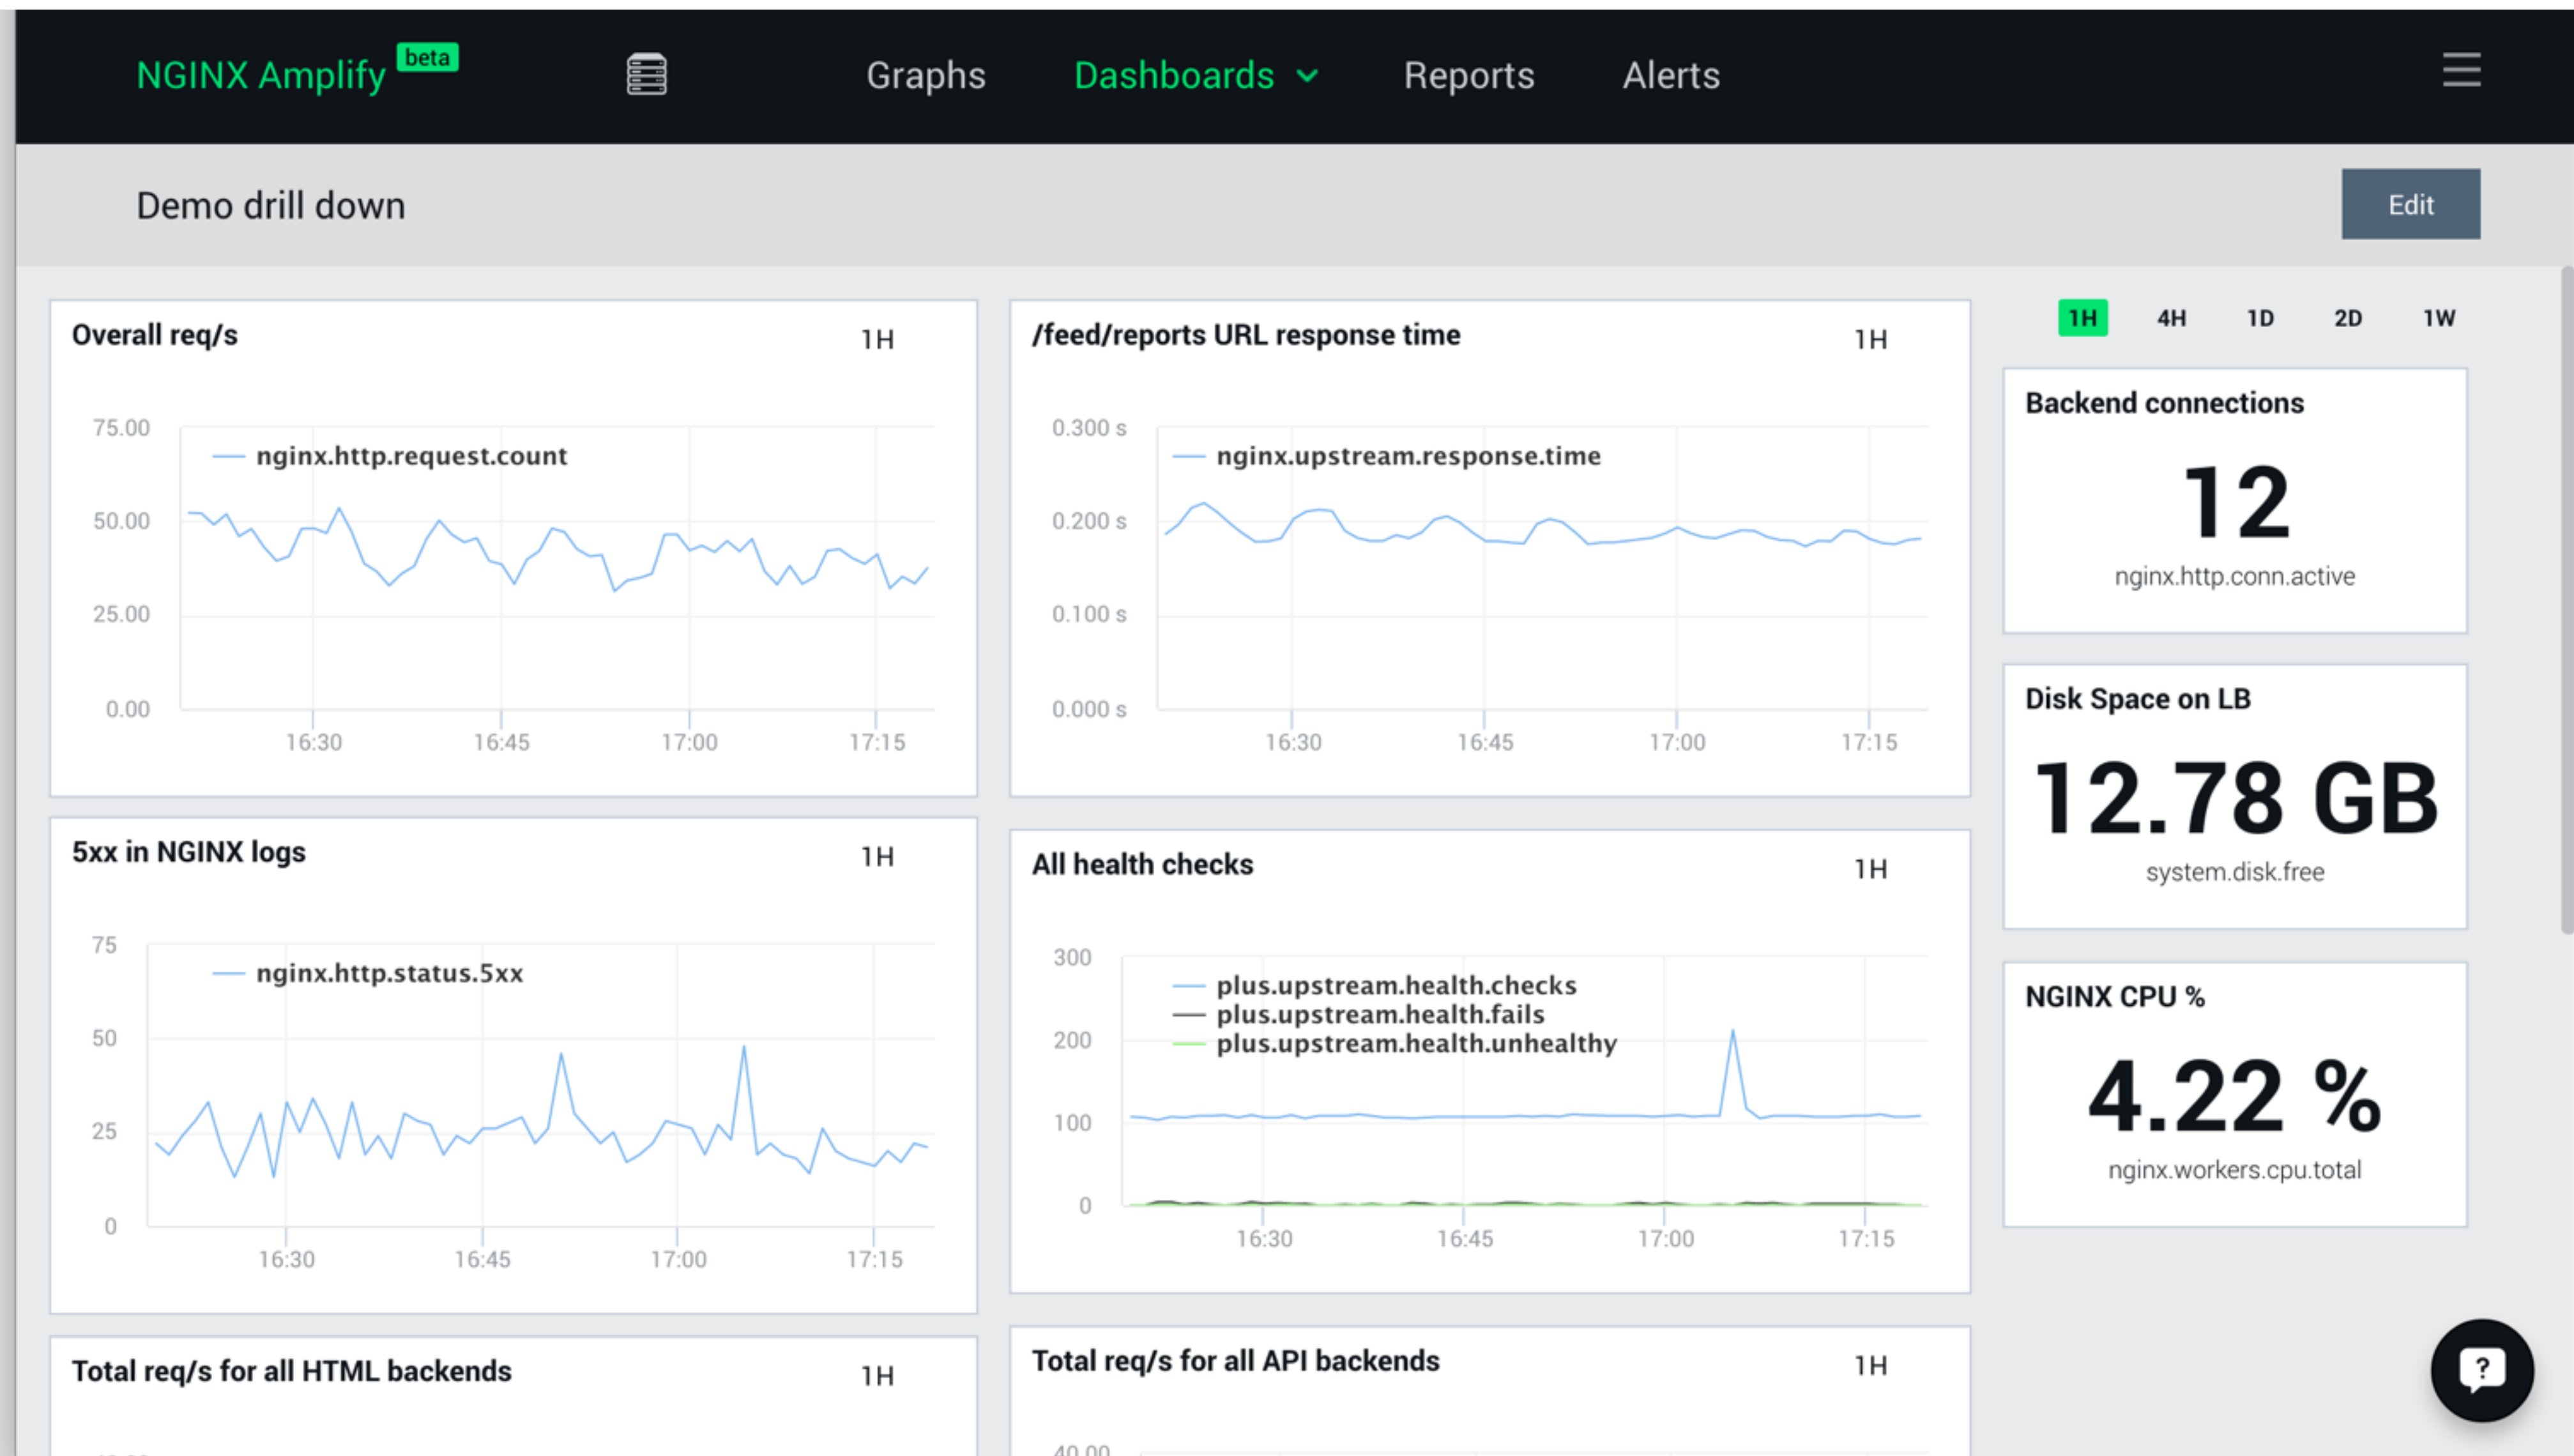
SYSTEMS: A list of systems is shown on the left, with 'prod-nginxplus-lb01' (nginx-plus-extras-r10-pre1) highlighted with a green border. Other systems include dev-nodejs-api01-03, experimental-api-v2, prod-rails-web01-02, and a 'New System' button.

PREVIEW: The 'prod-nginxplus-lb01 - System' section contains eight monitoring graphs for the time period 14:00 to 17:00. The graphs are: CPU USAGE %, MEMORY, DISK I/O (VDA1), DISK LATENCY (VDA1), NET. TRAFFIC (ETH0), LOAD AVERAGE, SWAP, and DISK USAGE.

GRAPH FEED: This section shows a scrollable list of graphs for the selected system. The top graph is 'nginx-plus-extras-r10-pre1 @ prod-nginxplus-lb01 - nginx', titled 'NGINX HTTP ERRORS, COUNT'. It shows three data series: nginx.http.status.4xx (blue), nginx.http.status.5xx (black), and nginx.http.status.discarded (green). The y-axis ranges from 0 to 500. The second graph is titled 'NGINX UPSTREAM ERRORS, COUNT' and shows nginx.upstream.request.failed (blue) and nginx.upstream.response.failed (black) series, with a y-axis from 0 to 40. The time range is set to 4H.

Copyright © 2015 NGINX, Inc. Terms of Service.

Sign up at amplify.nginx.com





Conclusions

- Plan for scalability early
- Tune low level operating system
- Configure Keepalive
- Configure caching
- Enable HTTP/2
- Measure your results



How to Contribute

- hg.nginx.org
- github.com/nginx
- wiki.nginx.org
- nginx.org/mailman

Thank You

All links in one page:
<https://shadrin.org/talks>

Twitter: @shadrin

nick@nginx.com

