



Brace for impact

@bephpug 2012

Christoph Lühr

@chluehr



basilicom

A word cloud featuring various terms related to digital marketing and technology. The words are arranged in a roughly circular shape, with 'facebook' being the largest and most central. Other prominent words include 'apps', 'websites', 'cms', 'strategy', and 'shop'. Smaller words include 'mobile', 'concept', 'social games', 'e-commerce', 'newsletter', 'design', 'marketing', and 'banner'.

apps
shop
banner
strategy
mobile
concept
social games
facebook
e-commerce
websites
newsletter
design
cms
marketing

(short term)

High Load Scenarios

DISCLAIMER

~~**PHP**~~



==

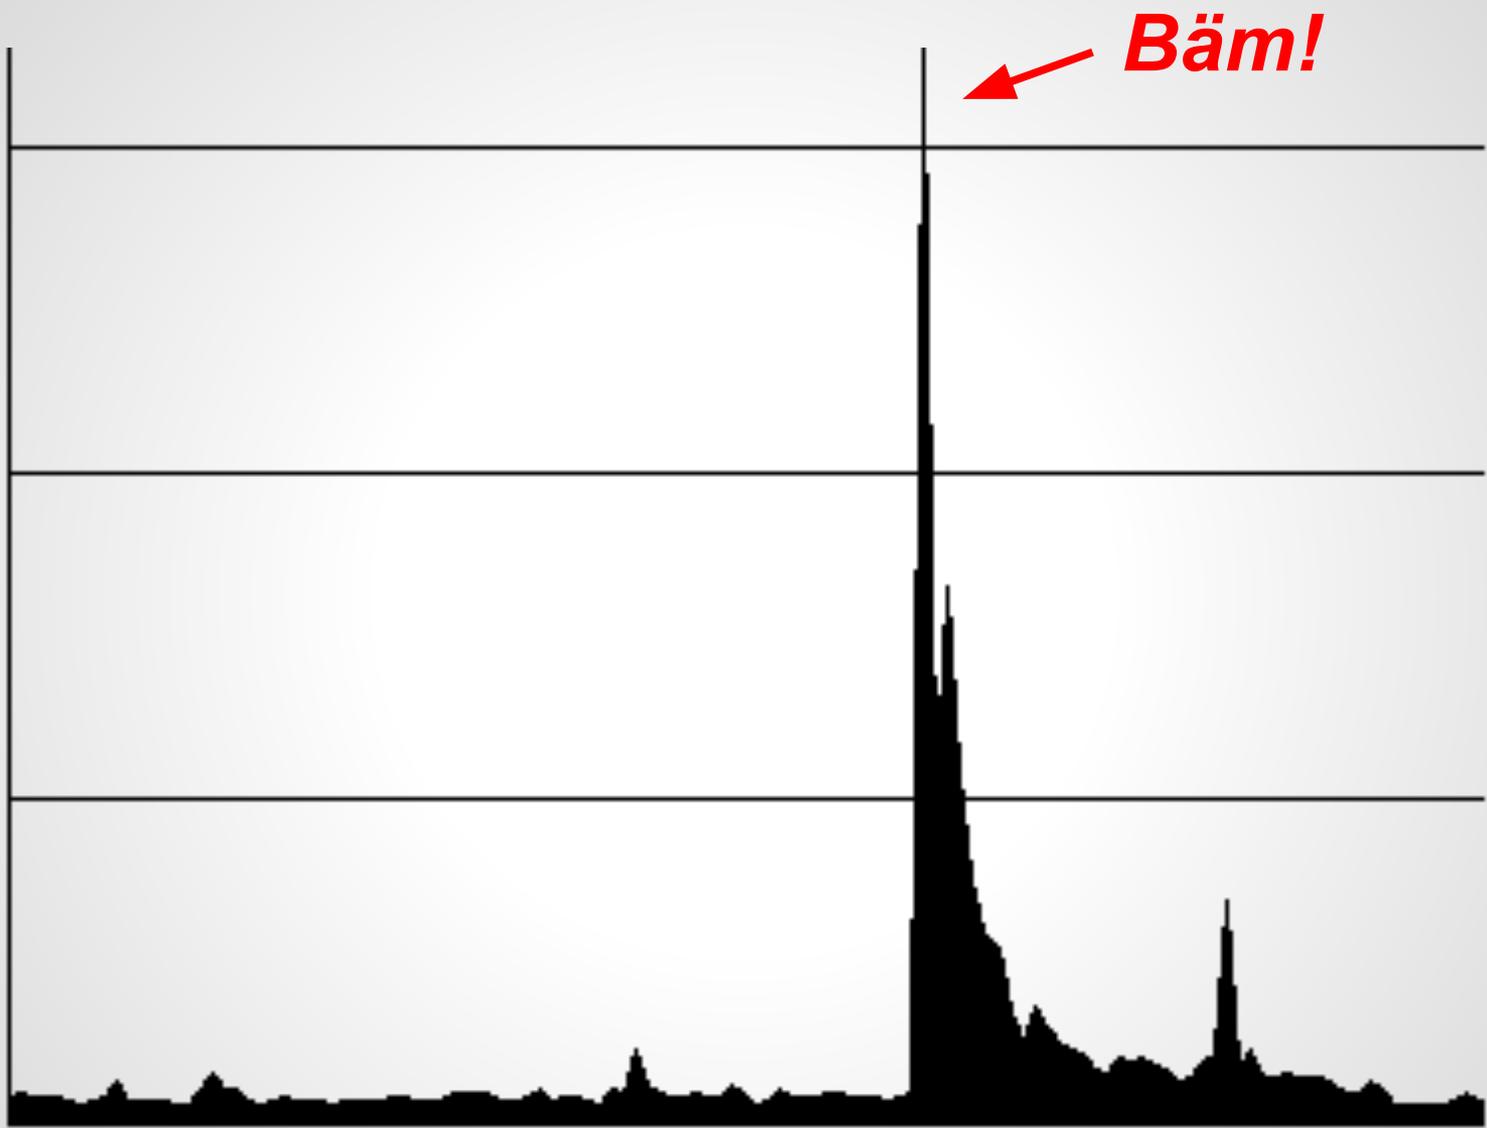


(if successful)



"Slashdot Effect"

(now: Facebook, Bild.de, ...)



Why



care?

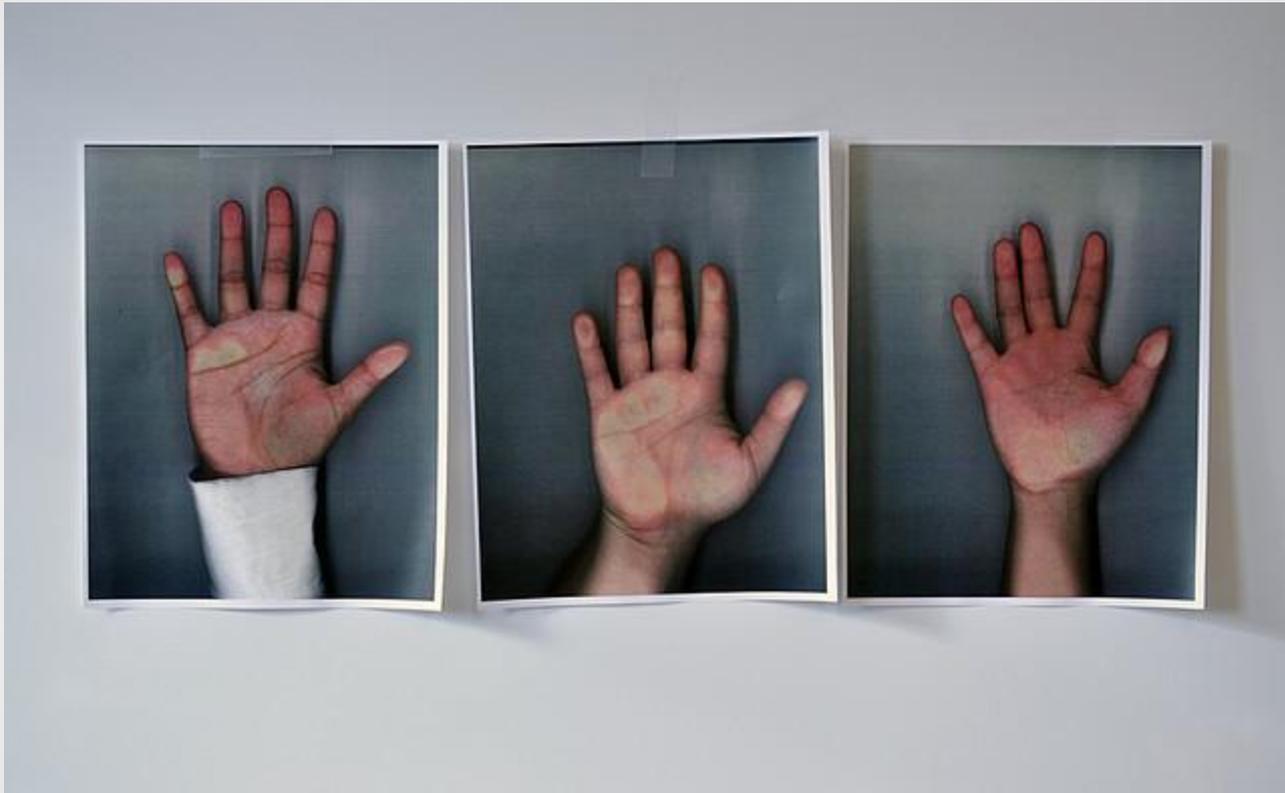
\$\$\$

TV ads are expensive
downtime == wasted money

Why me?

5 yrs++

Deutsche Bahn, Ferrero, Peugeot
BMW, Sony, UMG, studiVZ, ...



Flight Plan

1. MAKE A PLAN

2. MAKE IT SCALE?

3. MAKE IT FAST!

4. [TEST, SIZE]

5. WATCH IT!

6. BONUS: EXAMPLES



Architecture

Brace for impact

Plan
~~Brace~~ **for impact**

Knowledge is Power

Media plans

TV schedules



**TELEVISION
LIES**

Planning

- just **Peaks** matter
- Media / marketing exposure?
- Target group? Click Rate?
- Max. daily visits? Top hour?
- Pages ^ visit, resources ^ page?
- Do not forget storage / traffic requirements.

Planning

Metric	High-Case	Realistic-Case
TV Potential Users	7270000	7270000
Percentage of participating Users	4,5	0,8
Participating Users Total	327150	58160
Print Potential Users	660000	660000
Percentage of participating Users	3,5	0,7
Participating Users Total	23100	4620
Online Potential Users	53000000	53000000
Percentage of participating Users	0,6	0,2
Participating Users Total	318000	106000
Total Promo Users	668250	168780

Planning

Metric	High-Case		Realistic-Case	
Potential Users Total	670000	#	170000	#
Percentage of participating Users	100	%	100	%
Participating Users Total	670000	#	170000	#
Number of Campaign Days	14	#	14	#
Number of Visits per User during Campaign	1	#	1,2	#
Number of additional Viral User Visits per User	0,5	#	0,1	#
Visits total	1005000	#	221000	#
Percentage of Visits during Peak Day	20	%	15	%
Potential Peak DAUs / Visits per Day	201000	#	33150	#
Percentage of Peak Day DAUs in Top Hour	30	%/h	20	%/h
Potential Peak Day DAUs / Visits per Top Hour	60300	#/h	6630	#/h
HTTP Requests per PI	4	#/PI	3	#/PI
Initial HTTP Requests per Visit	20	#/V	20	#/V
Average Initial PI Size (KB)	700	kB	600	kB
Average consecutive PI Size (KB)	200	kB	150	kB
Traffic per Visit (KB)	1900	kB/V	1350	kB/V
Traffic Total (GB)	1821	gB/T	285	gB/T
Average PIs per Visit	6	#/V	5	#/V
Requests per Visit	44	#/V	35	#/V
Requests total in top Hour	2653200	#/Th	232050	#/Th
Average HTTP Requests per Second in top Hour	737	#/sTh	64	#/sTh

Binford ***Tools***

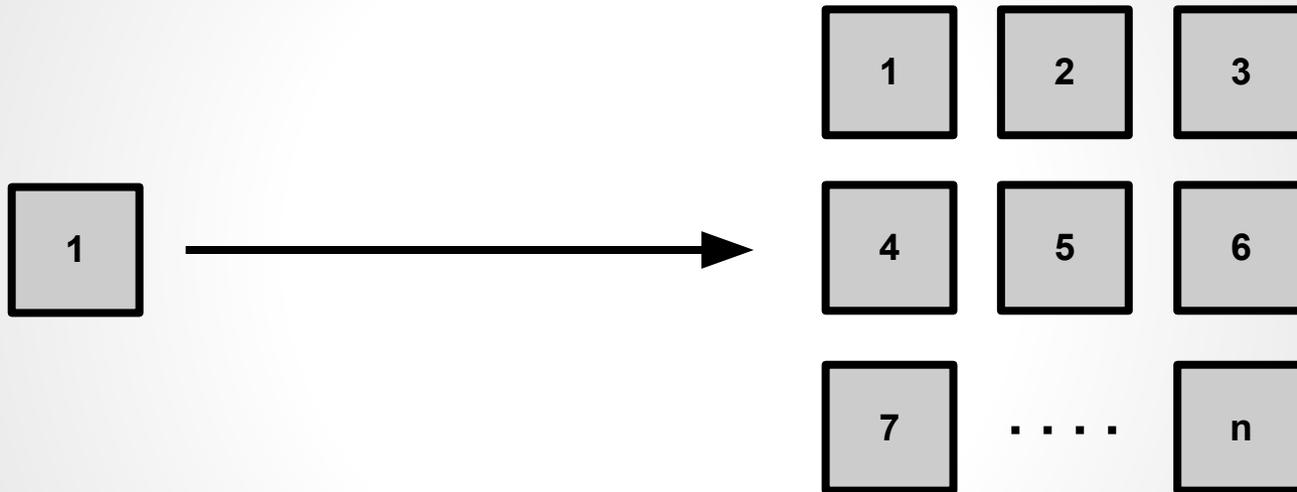
Architecture

(& Tool Time)

Fancy Solution

Elasticity
Scaling

"just add more boxes.."



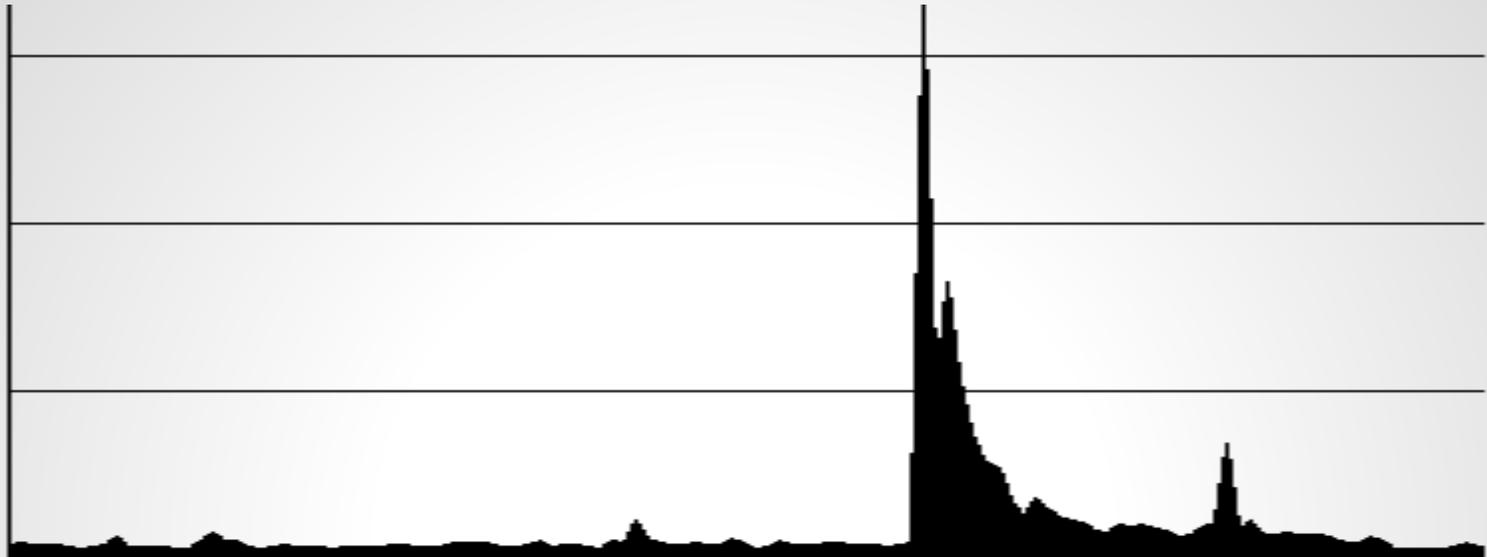
really?

System complexity ++

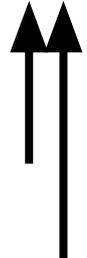
Database?

Time to spawn?

.....



TV-Ad



10 sec. later

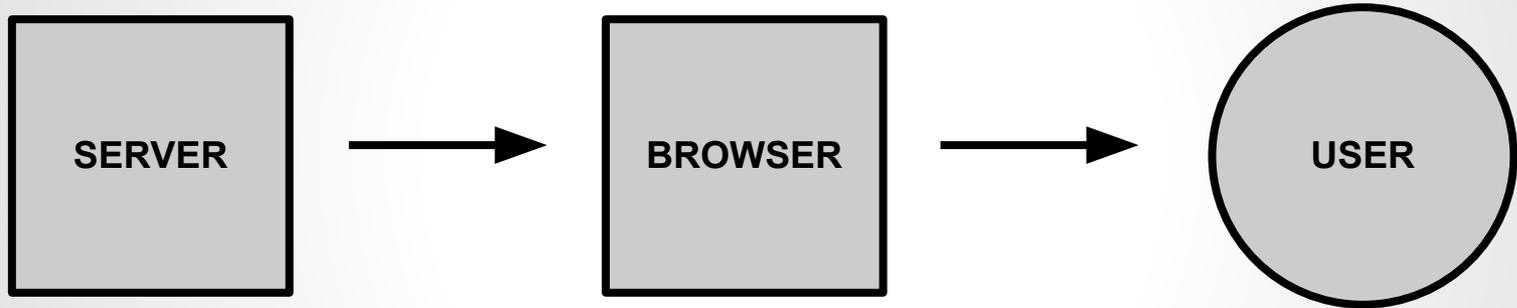


10 min. later

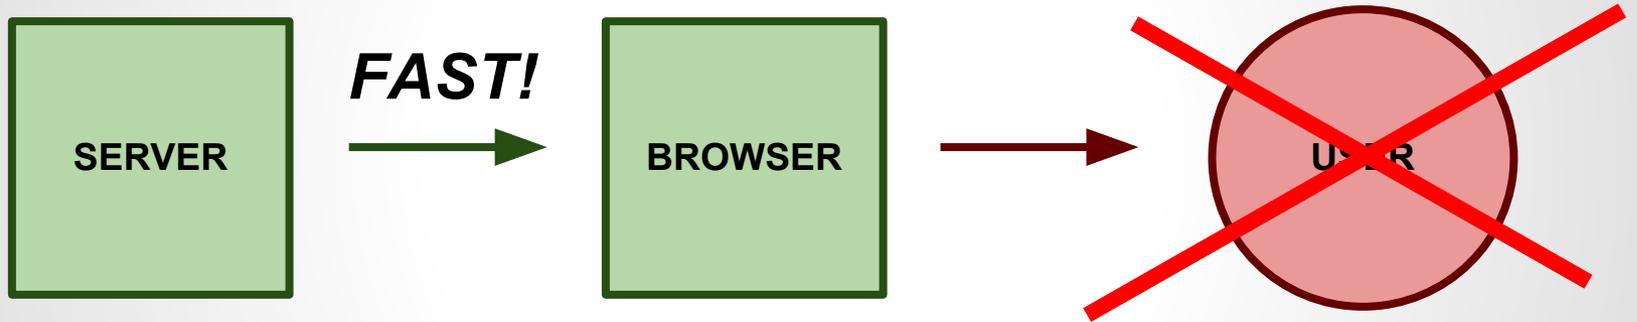
Traditional Solution

Performance ++

Performance



Performance





Performance Rules

Rule #1

Don't get hit.

- landing page on a different server.

Rule #2

Let **others** do the work

- use a Content Delivery Network (CloudFlare is free).

Rule #3

Serve static content

SERVE

STATIC

CONTENT!

Rule #3

Serve static content

- use a caching reverse proxy:
Varnish
- **NGINX** is fast, too
- "static" can even mean only 5 mins.

Rule #4

Serve static content once

- use correct cache headers
- leverage browser caches
- allow external proxies to cache your content.

Rule #5

Minimize requests

- compress HTML
- combine / minify JS, CSS
- inline images
- == get a good SpeedRanking.

Rule #6

Avoid *dynamic* requests

- load data upfront
- stay in browser, validate via JS first
- collect data / batch requests
- load only on action (captcha, ...).

Rule #7

small **is BIG**

- don't load a framework for an insert()
- just write a logfile?
- use nodeJS.

Rule #8

De-couple / postpone

- use queuing and finite # of workers
- use external services
(transloadit, youtube, ...).

Rule #9

Degrade gracefully

- skip doing hard work (thumbnails,...)
- hide expensive functionality (edit)
- ... to keep the site online.

Secret Rule #10

Cheat

- discard requests
- stats & voting: sample only 10%
- simulate server calls (spinner).

Final Rule #11

Don't go offline.

- put campaign on different server
- limit resources, don't kill your server
- if all fails, be sure to display a notice.

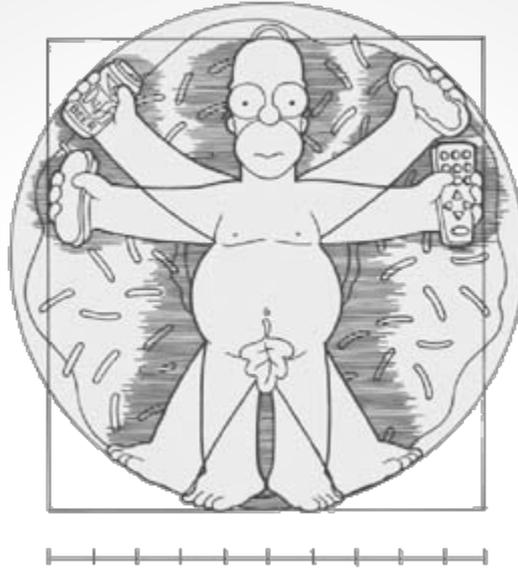


TESTING

I FIND YOUR LACK OF TESTS DISTURBING.

Testing

- Goal: Determine app limits/threshold
- Testing is **HARD**
- Test the production environment
- Monitor **performance & resources**
- httpperf, apache bench, siege, JMeter.



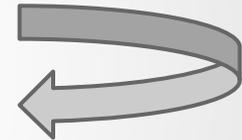
Sizing & Tuning

Sizing & Tuning

- RAM
- CPU cores / CPU speed
- IO: Storage Type / Storage Size
- Cache memory (Varnish)
- MySQL tuning (Caches, ...)
- System/Network tuning (max_open_files, ...)
- Apache (max. childs, ...)



Rinse, repeat.



(building, testing, sizing)



Online:
Monitoring

Monitoring

- **Access** (*Google Analytics*)
- **External Availability** (*Pingdom, ...*)
- **Speed, Latency & Resources** (*New Relic, ...*)
- **Application** (*# participants, ... StatHat*)
- *OR homegrown (statsd, etc.).*

Examples

Forms

Signup Form

← → × ↶

Name:

Captcha:

Photo Upload

SUBMIT

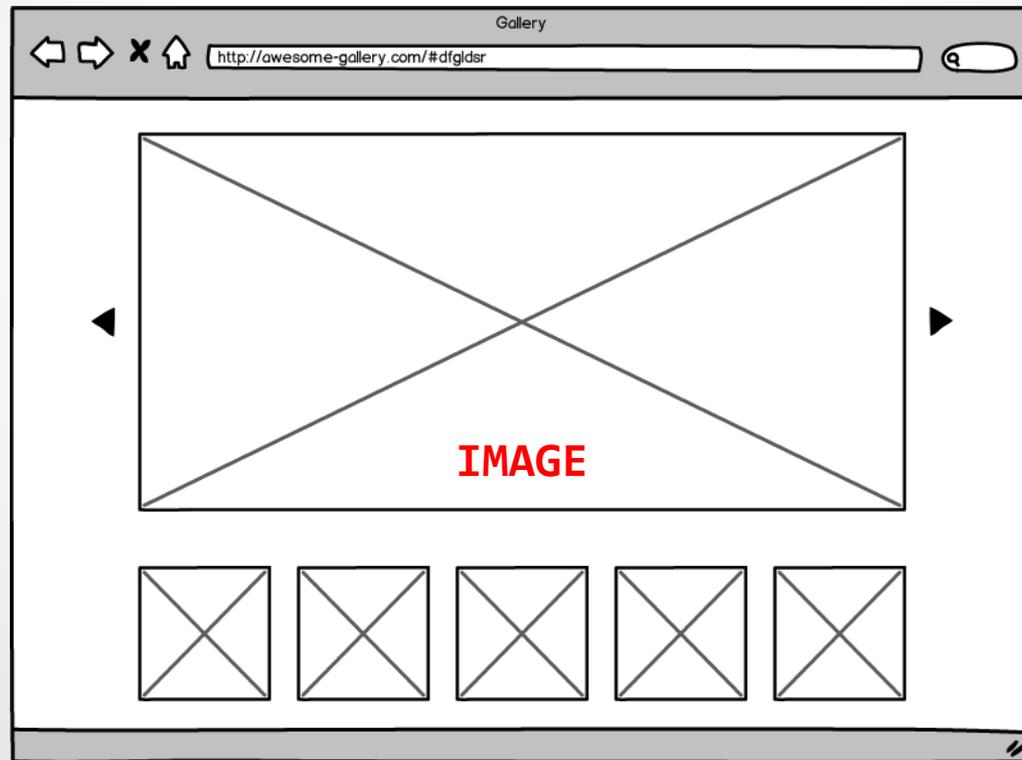
Forms

The image shows a browser window titled "Signup Form" with a navigation bar containing back, forward, close, and home icons, and a search bar. The main content area contains a registration form with the following elements:

- Name:** A text input field with a yellow circle labeled "1" at its right end.
- Captcha:** A captcha image showing the characters "m9|b" with a green circle labeled "2" at its right end, positioned above a corresponding text input field.
- Photo Upload:** A square area with a large "X" and the text "Photo Upload" in the center, with a red circle labeled "3" at its top-right corner.
- SUBMIT:** A button labeled "SUBMIT" located at the bottom right of the form.

Gallery detail pages

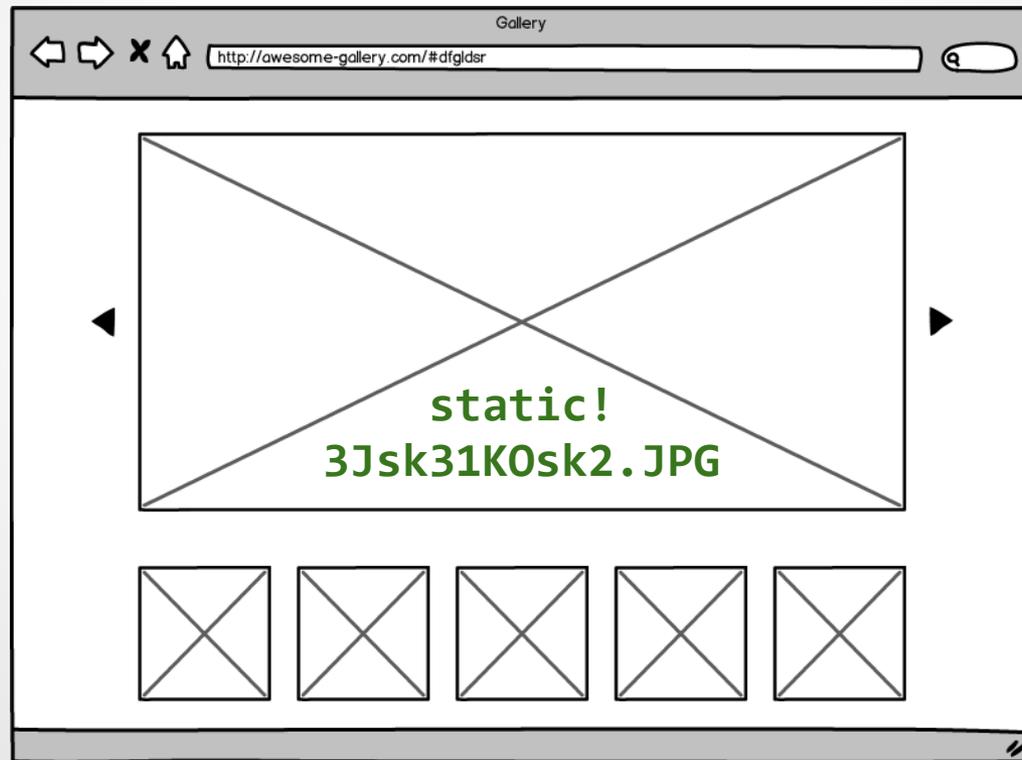
<http://foo.com/view/?key=3Jsk31K0sk2>



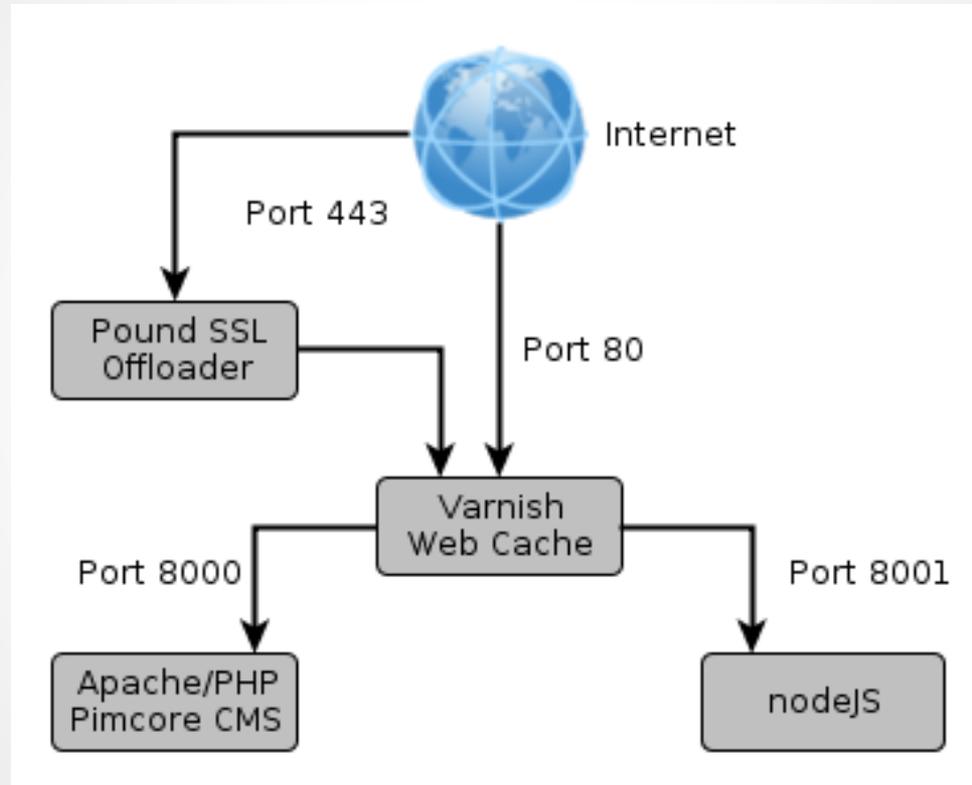
Gallery detail pages

`http://foo.com/view/?key=3Jsk31K0sk2` ==> :-)

`http://foo.com/view/#3Jsk31K0sk2` ==> :-)



Simplified setup

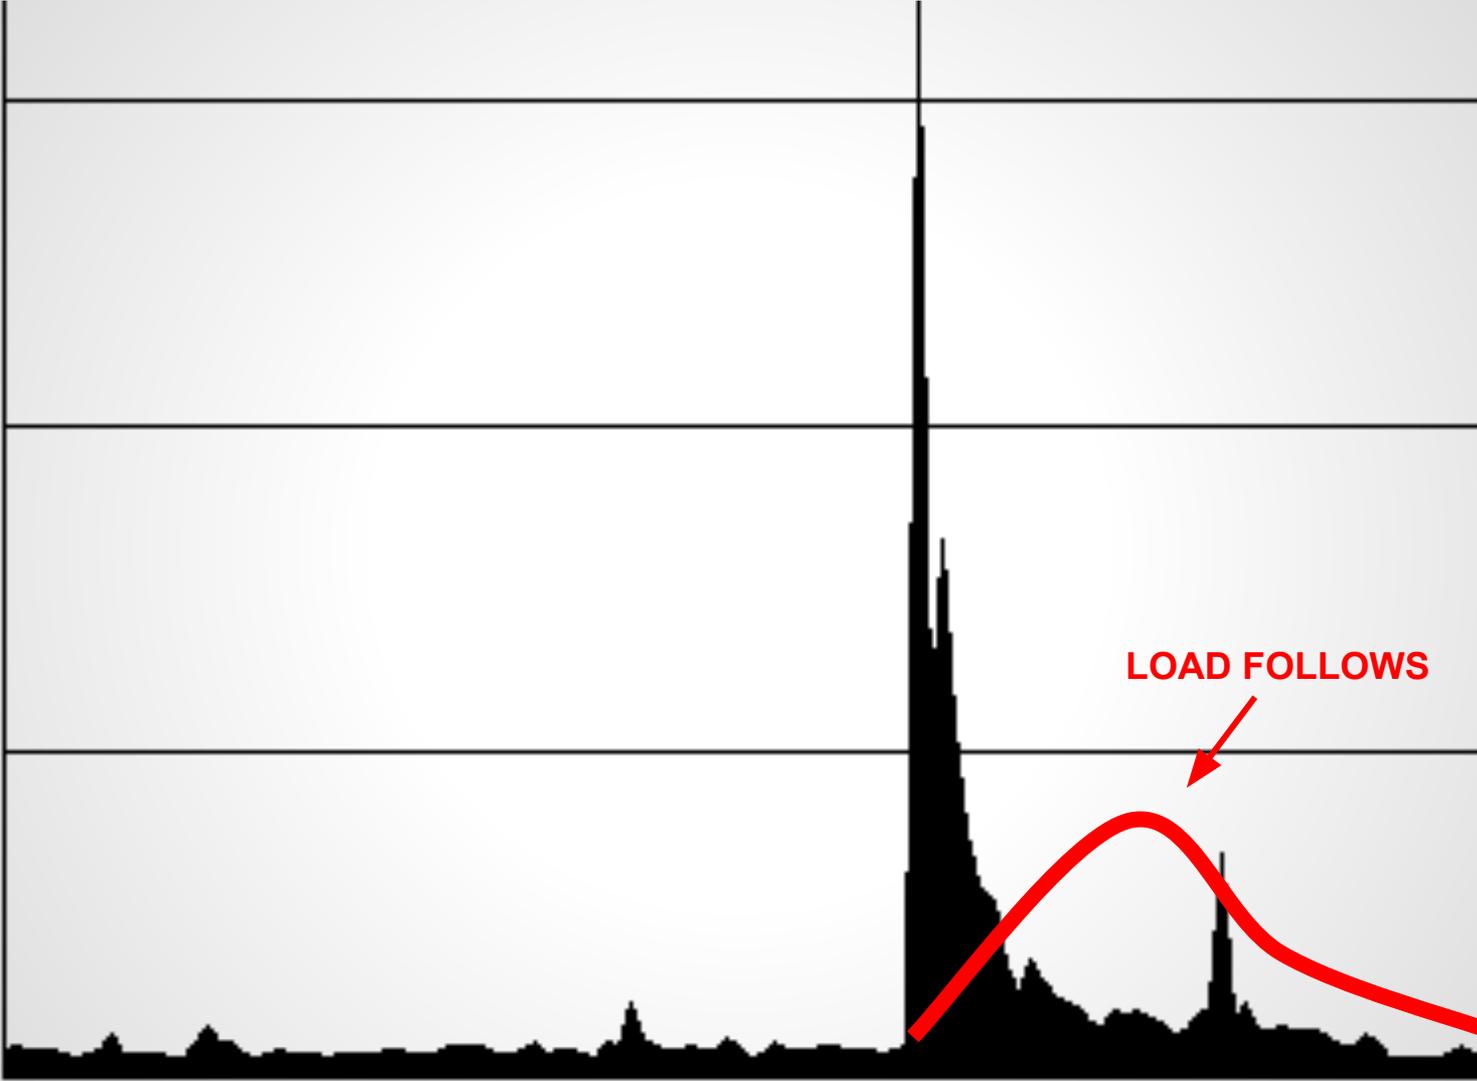


Summary

SERVE

STATIC

CONTENT!



LOAD FOLLOWS

CRITICAL ↔

Questions?

Suggestions?

Thanks

... and always have a safe landing!

Contact

Christoph Lühr

eMail: luehr@r-pentomino.de, christoph.luehr@basilicom.de

Twitter: [@chluehr](https://twitter.com/chluehr)

Slides license

Attribution-NonCommercial-ShareAlike 3.0

<http://creativecommons.org/licenses/by-nc-sa/3.0/>



Links

- **Google PageSpeed Insights**
<https://developers.google.com/speed/pagespeed/insights>
- **Transloadit: Image / Video uploads / conversion**
<https://transloadit.com/>
- **Pingdom Monitoring & Tools**
<http://tools.pingdom.com/fpt/>
- **New Relic application & server monitoring**
<http://newrelic.com/>
- **CloudFlare CDN**
<https://www.cloudflare.com/>

Links

- **Continuous Performance Testing**

[http://qafoo.](http://qafoo.com/talks/12_05_ipc_se_continuous_performance_tests.pdf)

[com/talks/12_05_ipc_se_continuous_performance_tests.pdf](http://qafoo.com/talks/12_05_ipc_se_continuous_performance_tests.pdf)

- **Varnish**

<https://www.varnish-cache.org/>

- **Varnish Speed**

[http://kly.](http://kly.no/posts/2010_10_23__High_End_Varnish___275_thousand_requests_per_second__.html)

[no/posts/2010_10_23__High_End_Varnish___275_thousand_requests_per_second__.html](http://kly.no/posts/2010_10_23__High_End_Varnish___275_thousand_requests_per_second__.html)

Links

- **NGINX**

<http://www.slideshare.net/Edorian/nginx-php-fpm-the-webserver-you-might-actually-like-php-usergroup-berlin>

- **High Performance Web Sites - Steve Souders**

<http://shop.oreilly.com/product/9780596529307.do>

- **Pound Load Balancer, Reverse Proxy & SSL Offloader**

<http://www.apsis.ch/pound>