Felix Gessert

Baqend

Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG

Building a Shop with Sub-Second Page Loads

Lessons Learned

Hosting & Scalability

code.talks
commerce special by ABOUT YOU

About me

Felix Gessert

PhD Thesis & Research

Web & NoSQL Technologies:
- Web Performance
- Cloud Data Management

Universität Hamburg

CEO & Co-founder

Baqend:
High-Performance
Backend-as-a-Service

Baqend

www.baqend.com

Presentation
is loading

Why performance matters

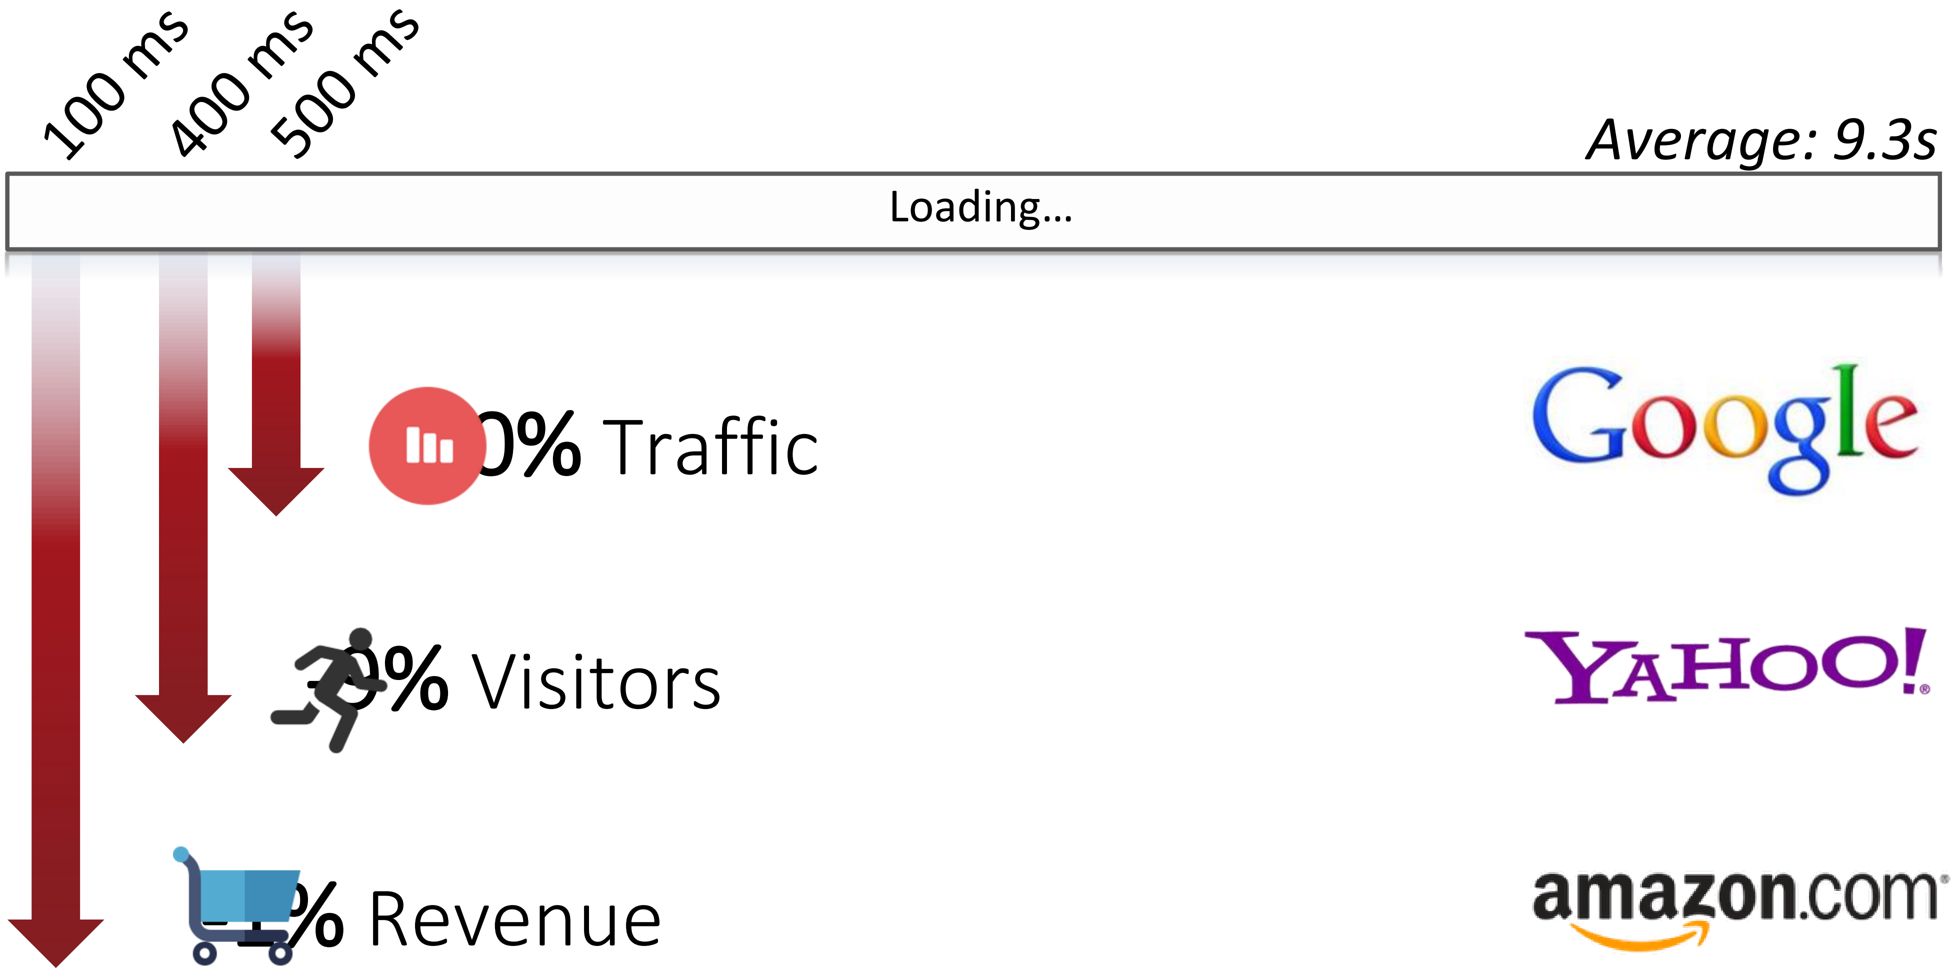

What is the goal?

Speed Perception

Delay	Perception
0 – 100 ms	Instant
100 – 300 ms	Small perceptible delays
300 – 1000 ms	Machine is working
1+ s	Mental context switch
10+ s	Task abandoned

Loading time < 1s

I. Grigorik, High performance browser networking. O'Reilly Media, 2013.

Concrete Example

A Scalable Webshop

THINKS

at

DIE HÖHLE DER LÖWEN

Expectations:

- **2.7** Million viewers
- **300.000** Visitors in **30** minutes
- **20.000** Requests per second
- **4 weeks** development and testing

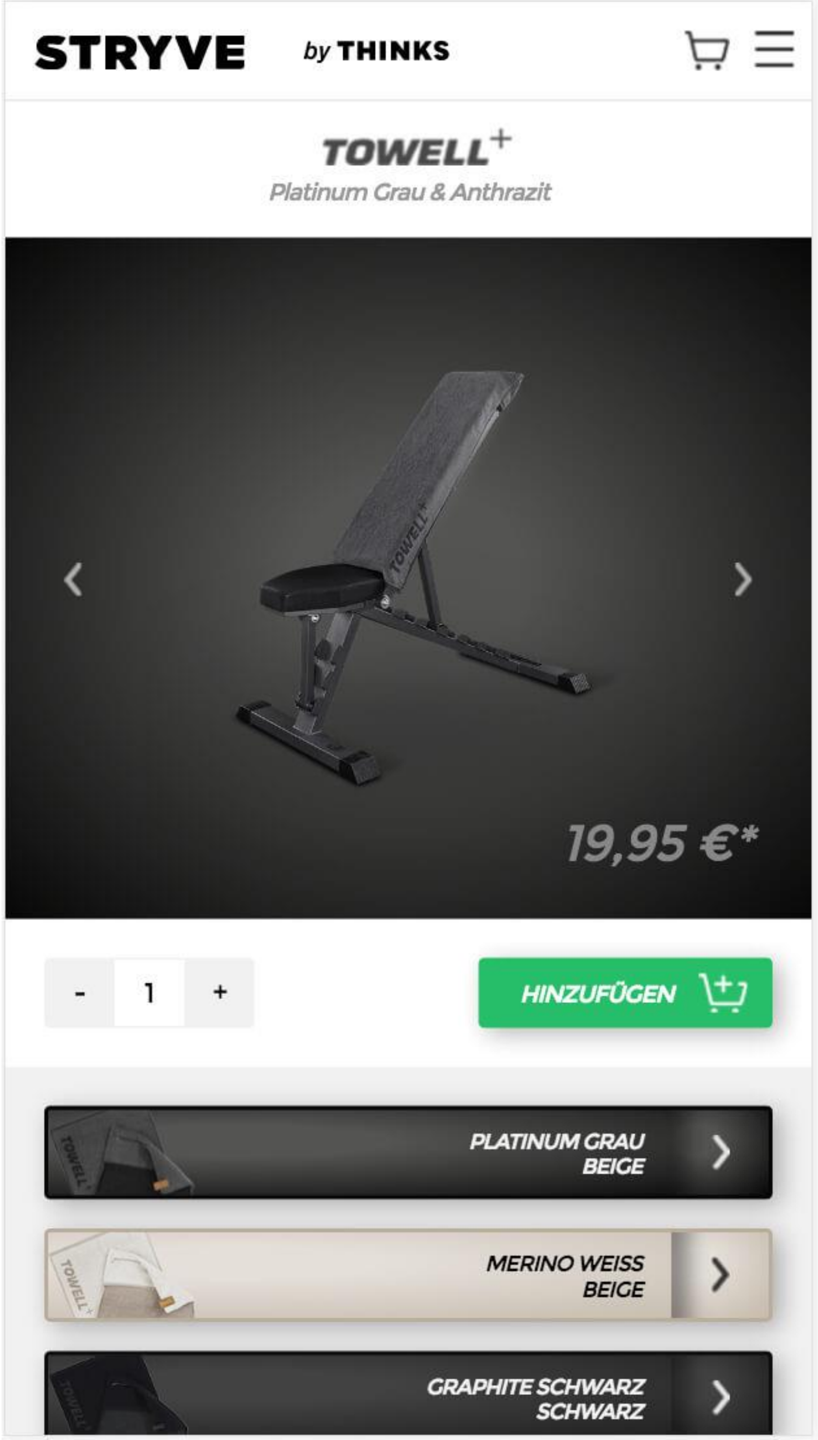

Goal:

Load time < **1s**

&

100% Availability

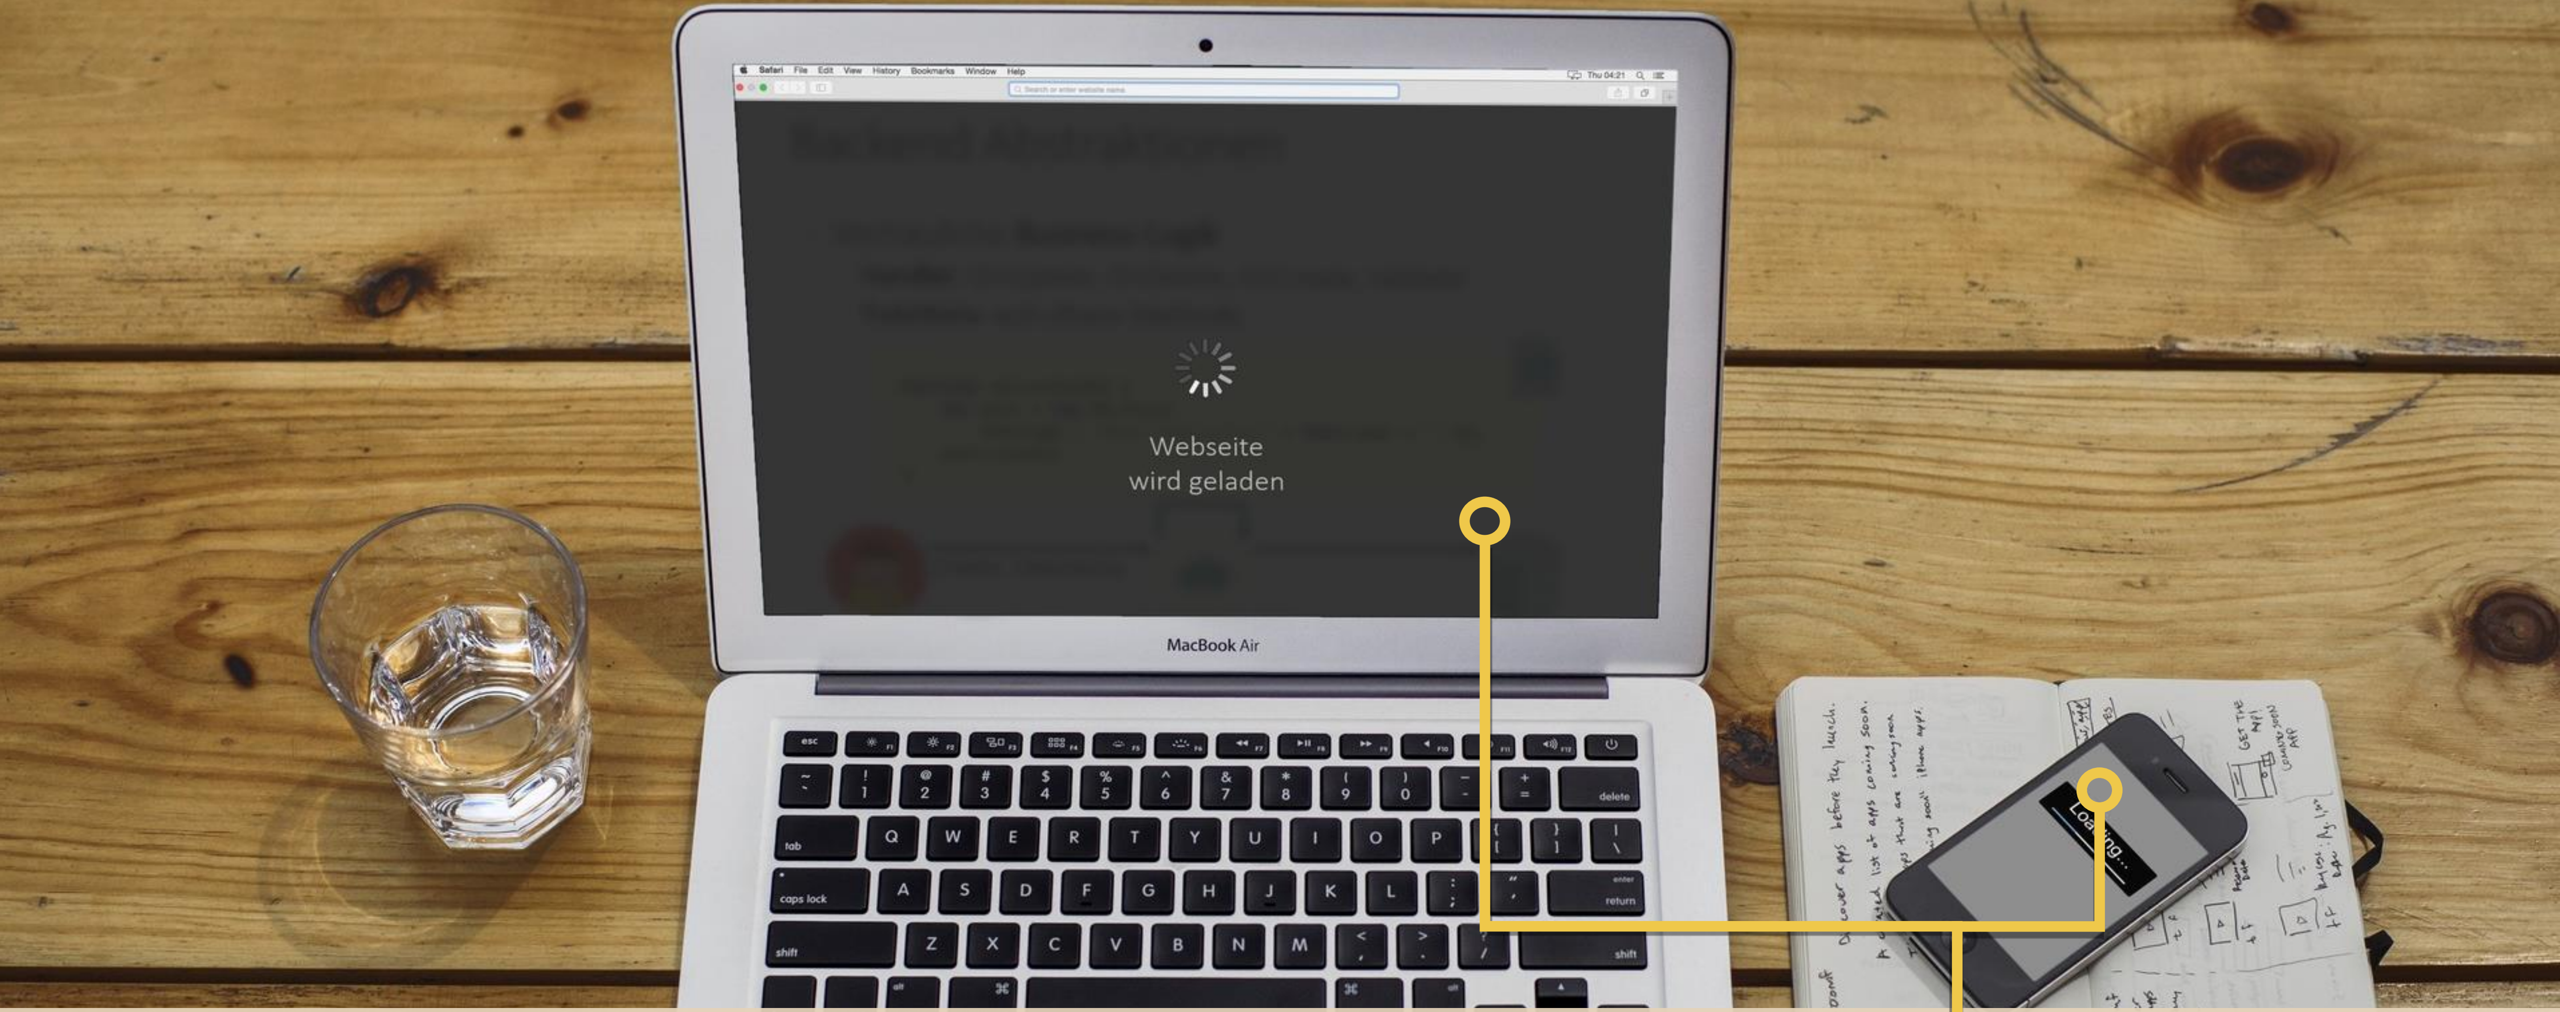

If performance is so important for

Business Success

...what causes slow page load times?

State of the Art

The Three Bottlenecks

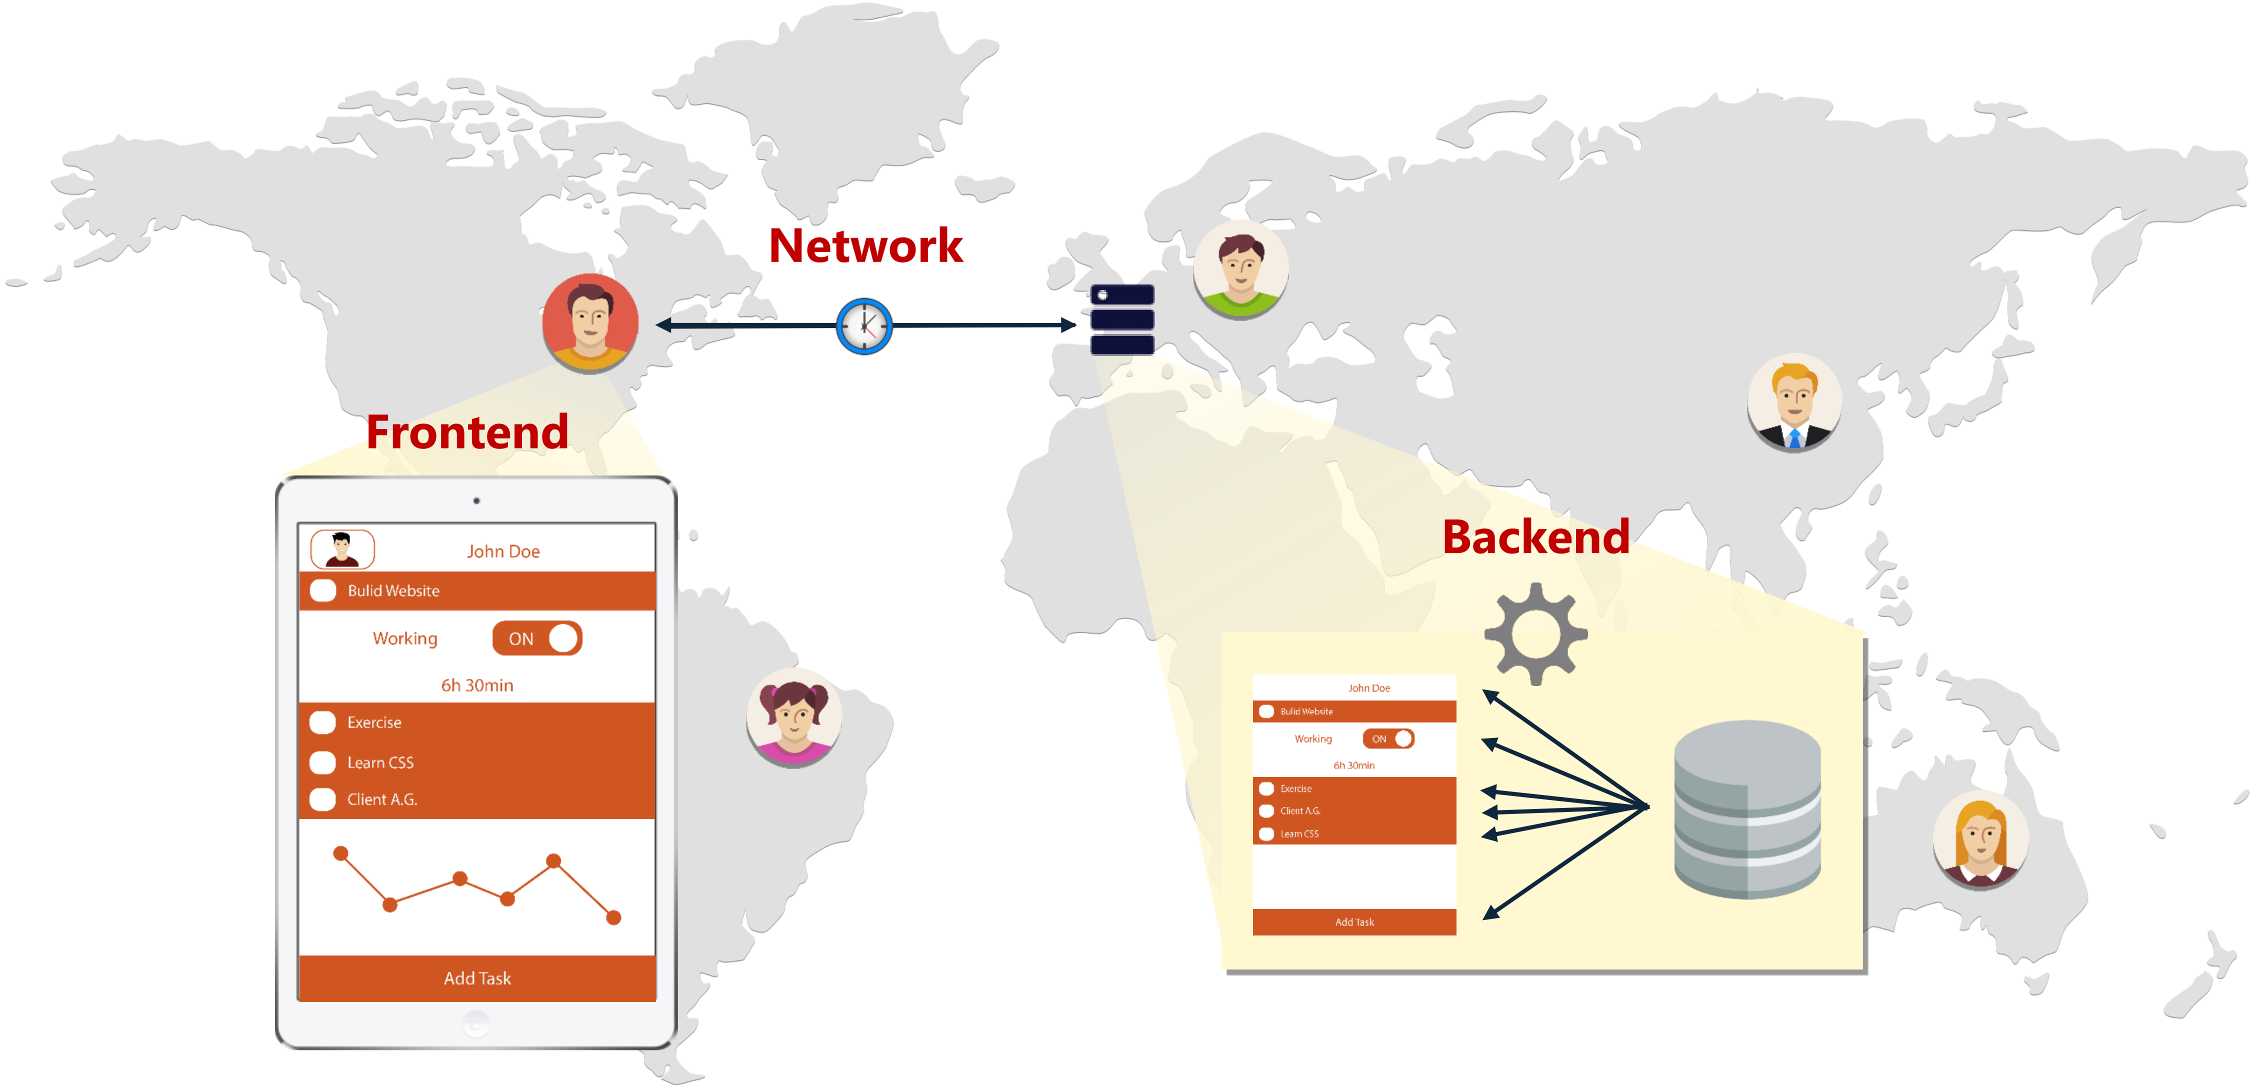

BOTTLENECK I

Optimizing **Frontend** Performance

Frontend Performance

The critical rendering path

```html
<!doctype html>
<title>Code Talks</title>
<link href=all.css rel=stylesheet />
<script src=app.css ></script>
<div>
  <h1>Web Performance</h1>
</div>
```

```html
<script>

   elem.style.width = "50px";
   document.write("JS is awesome!");

</script>
```

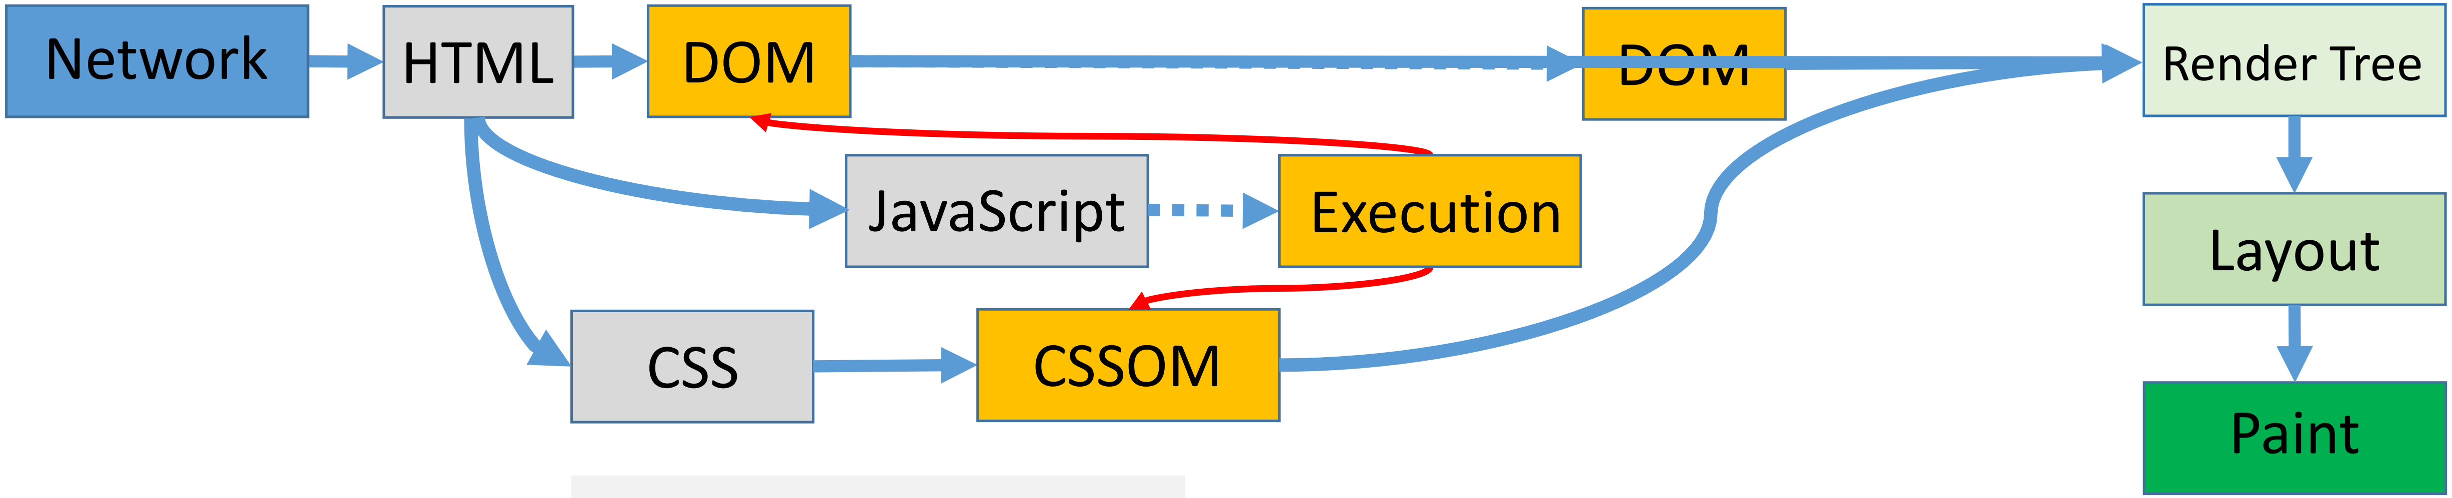


```css
div   { color: green; }
h1    { padding: 10px; }
```

Inlining cricital CSS and JS “above the fold”

Critical

nd

concat

JS

UglifyJs & cssmin

Single page application

handlebars

F ssive

ering

processhtml

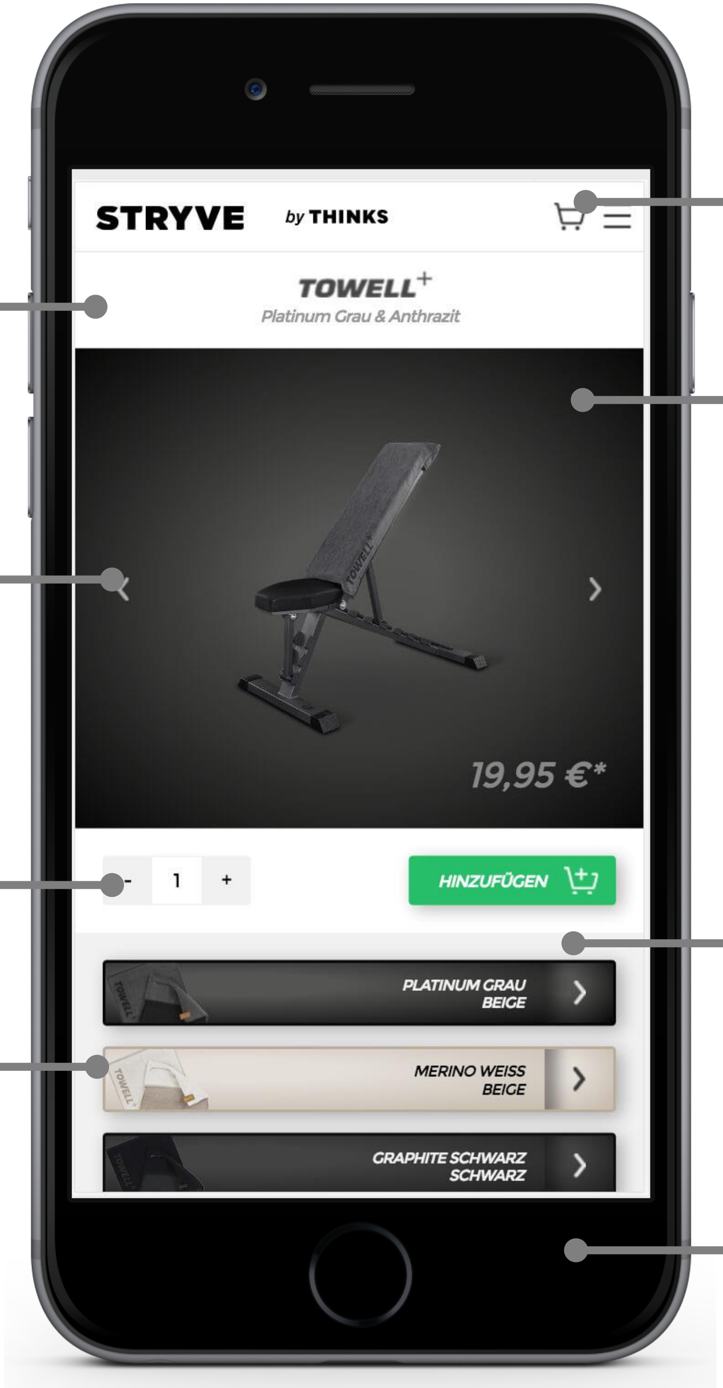

PostCSS

Co ess images

tiny png

GTmetrix

Test your performance

WEBPAGETEST

Load non-critical CSS and JS asynchronously

PageSpeed Insights

Frontend Performance
Applied in the example

Load time:

767 ms

Size:

565.9 KB

Requests:

24

THINKS Towell+ Sporthandtuch – TV: DHDL | STRYVE

GET www.thinks.com	200 OK	thinks.com	7.1 KB	26ms
GET 01.png	200 OK	thinks.com	86.4 KB	20ms
GET founders.png	200 OK	thinks.com	93.3 KB	31ms
GET jquery.min.js	200 OK	ajax.googleap	28.6 KB	145ms
GET app.js	200 OK	thinks.com	86.6 KB	33ms
GET app.css	200 OK	thinks.com	25.9 KB	24ms
GET sprite.png	200 OK	thinks.com	34.3 KB	30ms
GET bg.jpg	200 OK	thinks.com	103.6 K	48ms
GET bg.jpg	200 OK	thinks.com	31.3 KB	29ms
GET zhcz-_WihjSQC0oHJ9T	200 OK	fonts.gstatic.c	13.2 KB	276ms
GET lQHow_FEYlDC4Gzy_m	200 OK	fonts.gstatic.c	13.2 KB	281ms
GET connect	200 OK	thinks.com	3.6 KB	166ms
GET analytics.js	200 OK	google-analyti	11.3 KB	199ms
GET logo-stryve-inv.svg	200 OK	thinks.com	939 B	7ms
GET logo-thinks-inv.svg	200 OK	thinks.com	1.2 KB	8ms
GET 10.png	200 OK	thinks.com	21.4 KB	7ms
GET alert_message	200 OK	thinks.com	281 B	7ms
GET grau_anthrazit	200 OK	thinks.com	511 B	7ms
GET merino_taupe	200 OK	thinks.com	514 B	6ms
GET graphit_anthrazit	200 OK	thinks.com	534 B	8ms
GET platinum_black	200 OK	thinks.com	572 B	6ms
GET neonblue_black	200 OK	thinks.com	504 B	7ms
GET neonred_black	200 OK	thinks.com	563 B	6ms
GET neongreen_black	200 OK	thinks.com	630 B	14ms
24 Requests			**565.9 KB**	**693ms (onload: 767ms)**

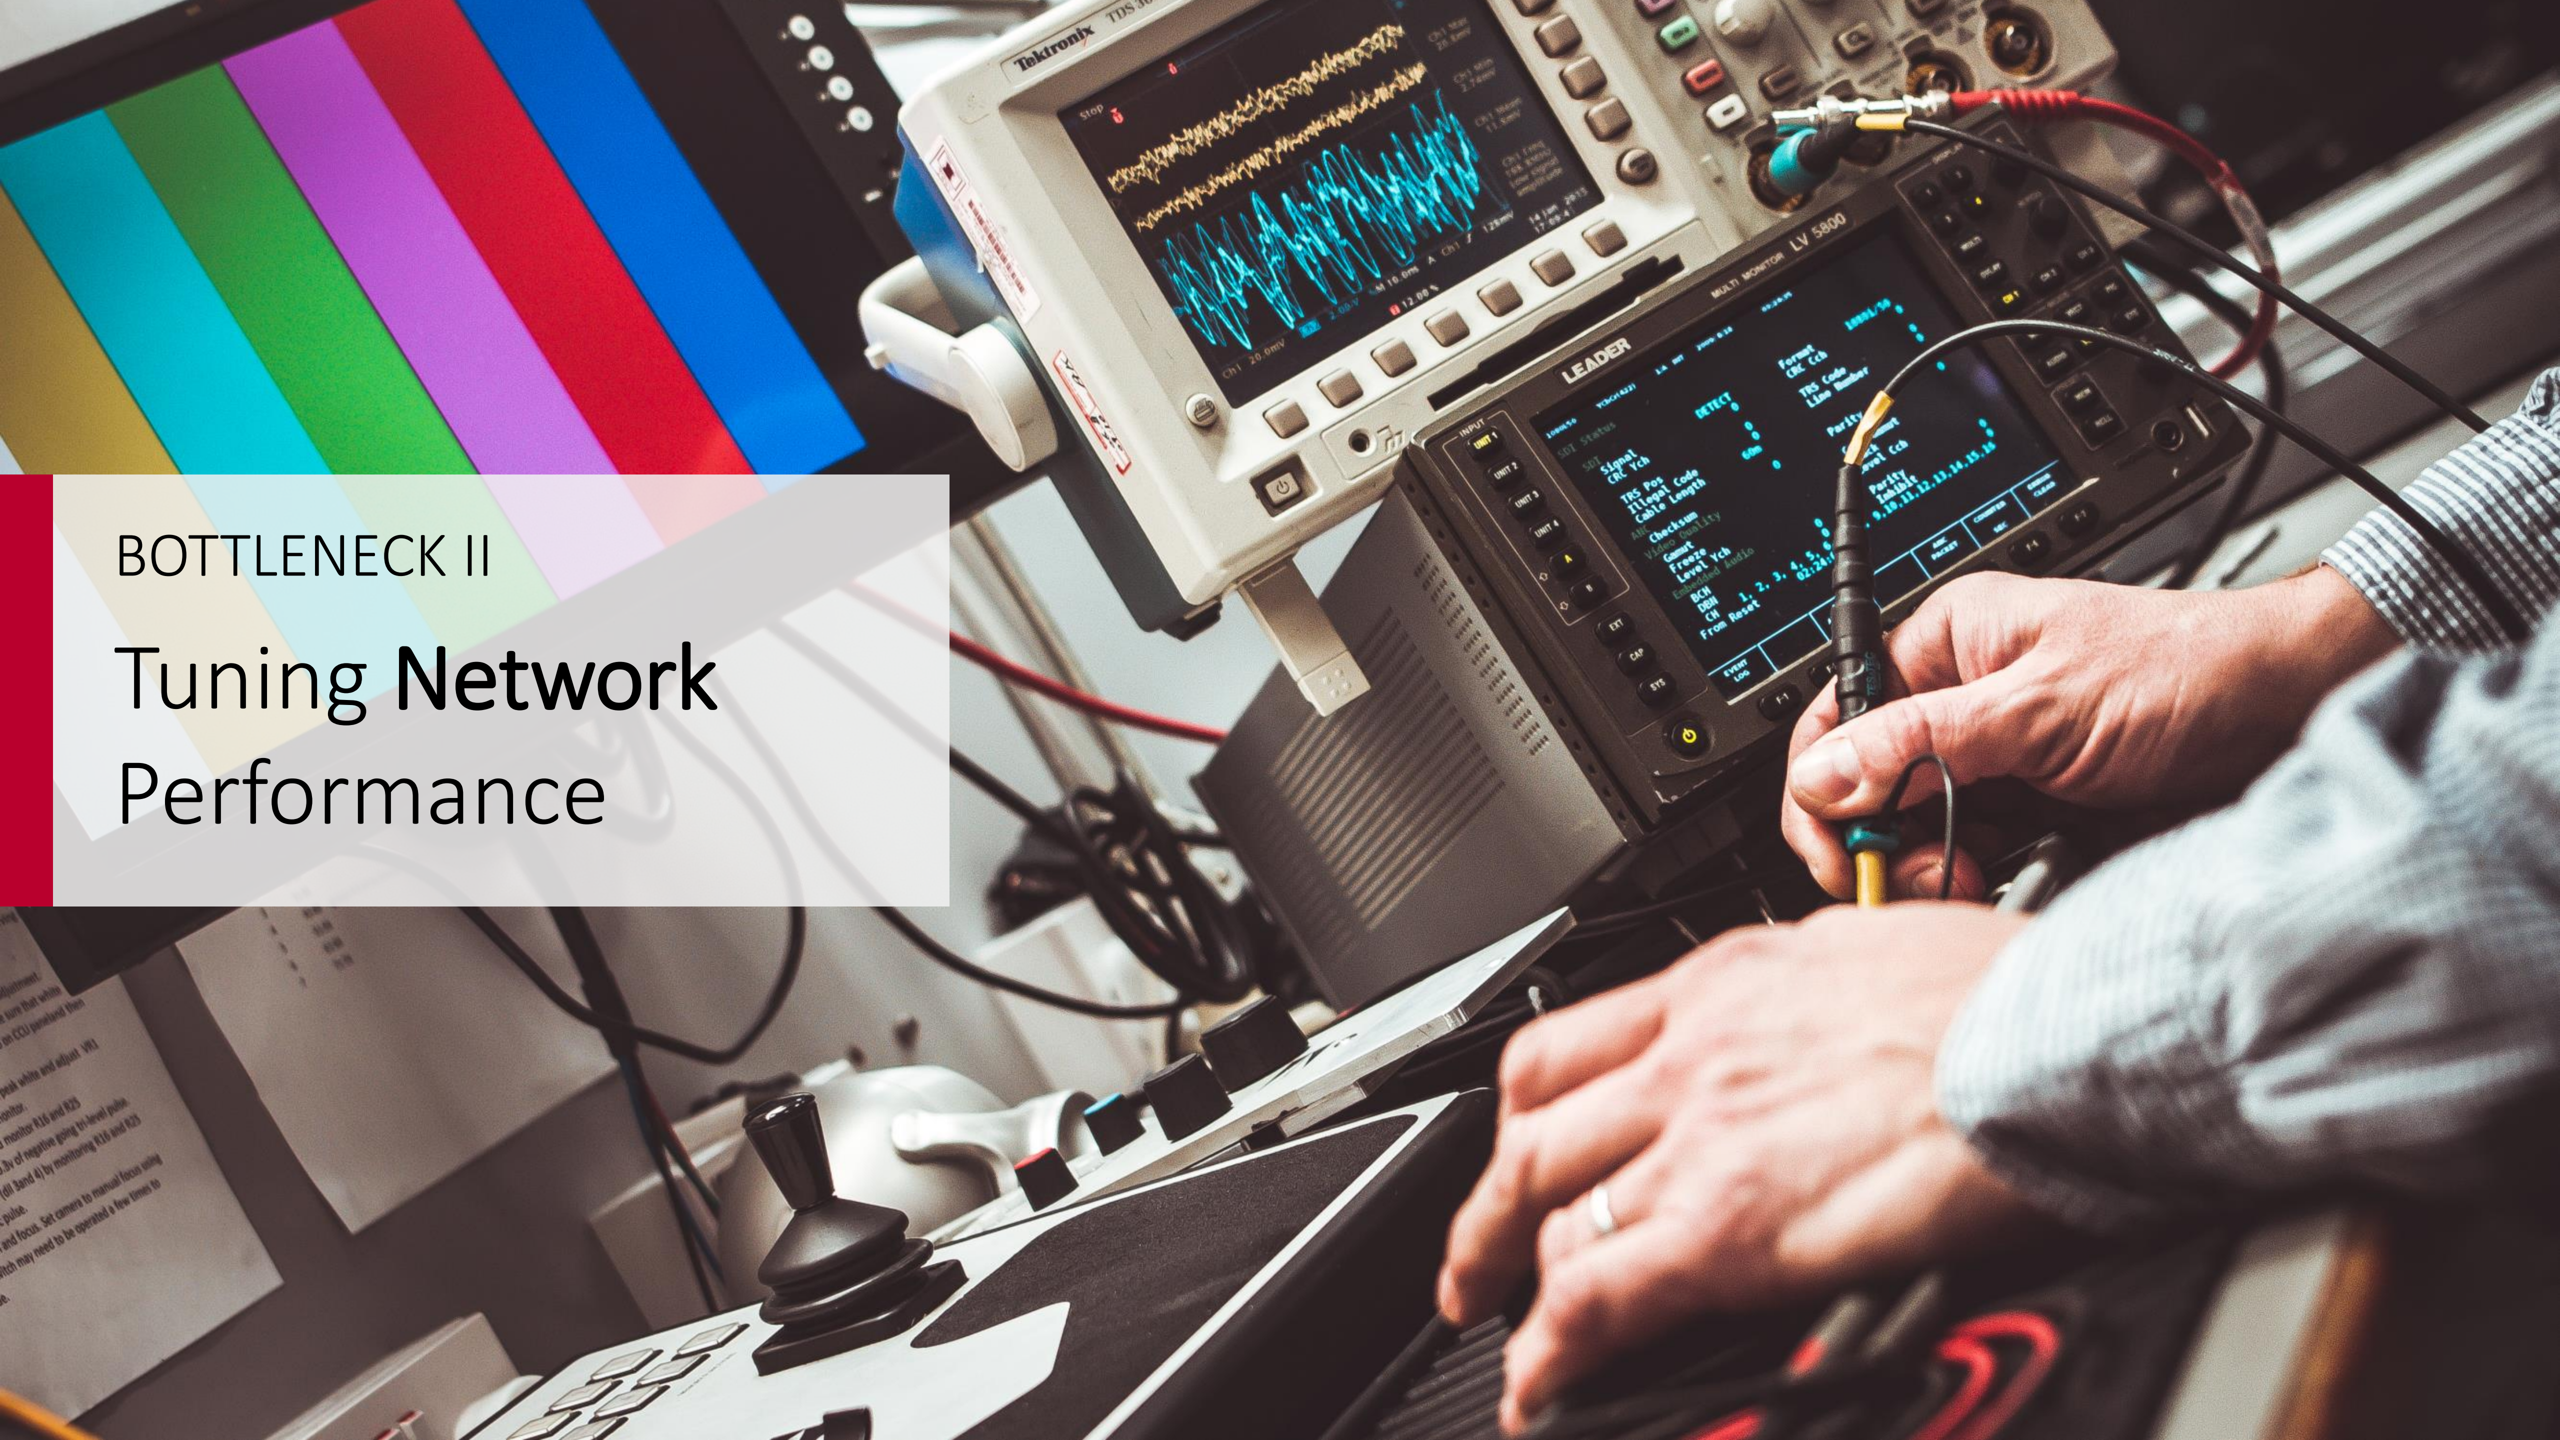

BOTTLENECK II

Tuning **Network** Performance

Network Performance

Break down of a single resource load

Maximum 6 parallel connections

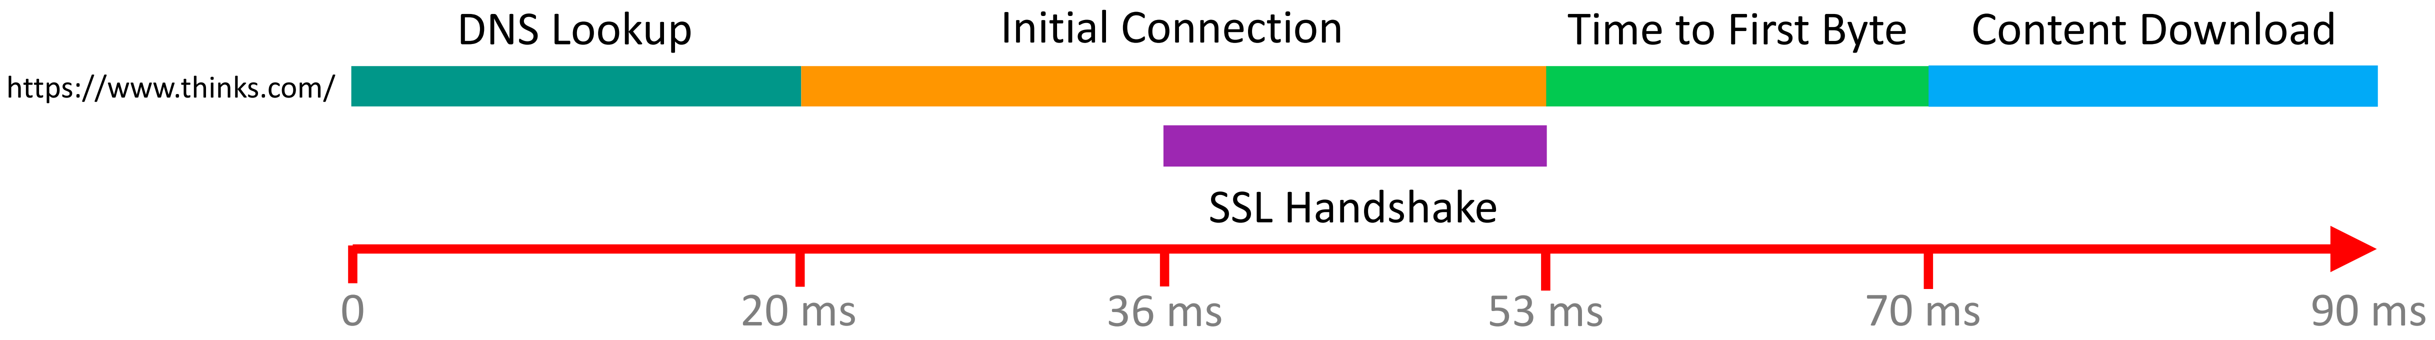

DNS Lookup

- Every domain has its own **DNS lookup**

Initial connection

- TCP makes a **three-way handshake** → 2 roundtrips (1 with TCP Fast Open)
- **SSL** connections have a more complex handshake → +2 roundtrips (only 1 with TLS False Start or Session Resumption)

Time to First Byte

- Depends heavily on the **distance** between client and the backend
- Includes the time the backend needs to **render**
  - ➢ Session lookups, Database Queries, …

Content Download

- New connections are slow (**initial congestion window** is small) → many roundtrips

Network

Bandwidth vs. Latency

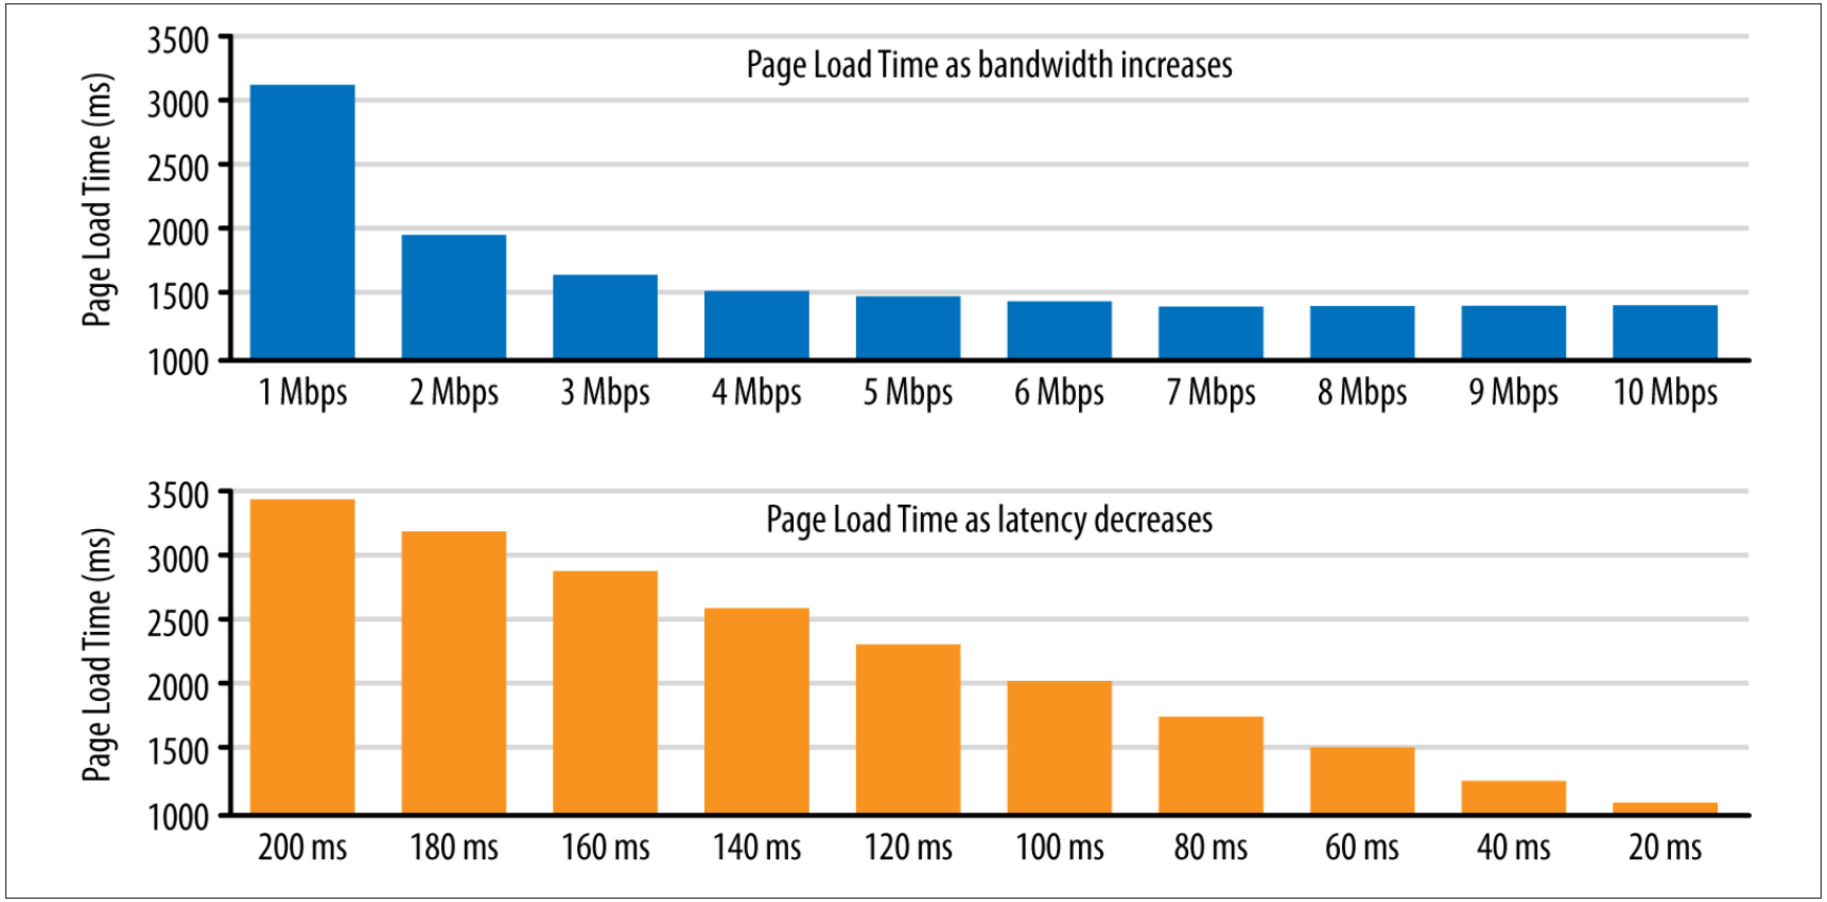

I. Grigorik, High performance browser networking. O'Reilly Media, 2013.

Netzwerk

Bandbreite vs. Latenz

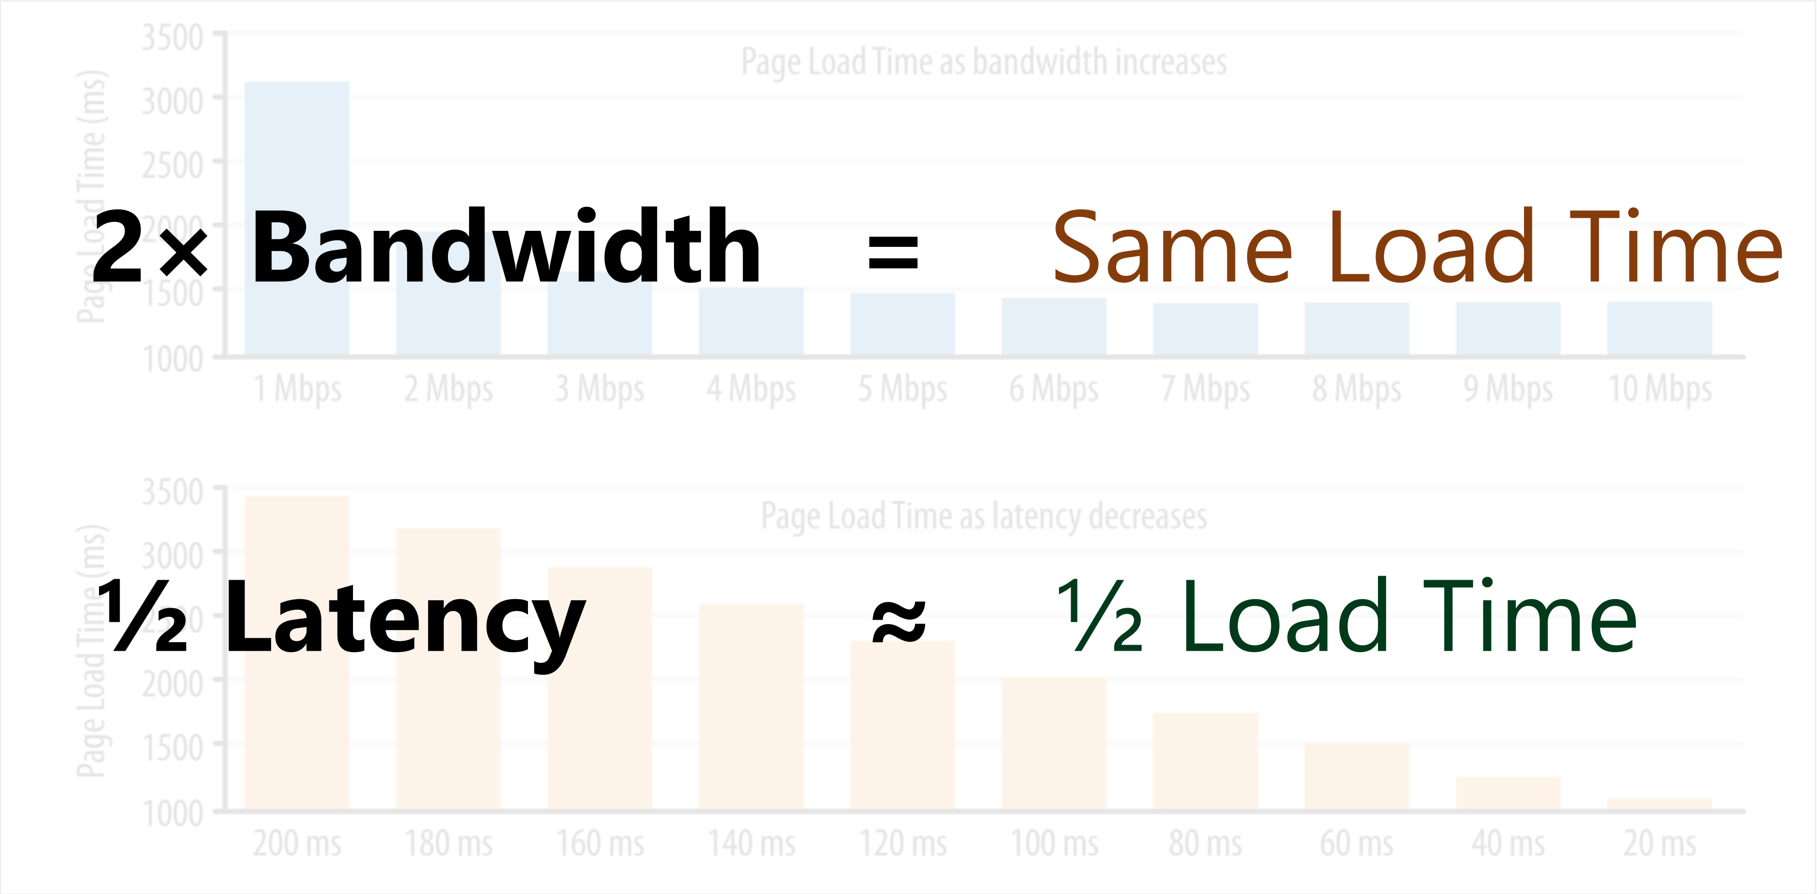

2× Bandwidth = Same Load Time

½ Latency ≈ ½ Load Time

Network Performance

Common Tuning Knobs

- **Persistent** connections
- Avoid **redirects**
- Set **caching headers** (no heuristic caching)
- Content **D**elivery **N**etworks
  - Lower Latency
  - Cache images, CSS, JS
  - Terminate SSL early and optimized
- Single **P**age **A**pps:
  - Small initial page with JS
  - Loads data lazily

HTTP/2:

- **Multiplexing** over 1 TCP connection (no head-of-line blocking)
- Request **Pipelining**
- Server **Push**
- **Header Compression**

Why HTTP/2 Matters

Typical Improvement

HTTP	HTTPS
3,22s	4,03s

HTTP	HTTPS
with CDN	with CDN and h2
0,44s	0,35s

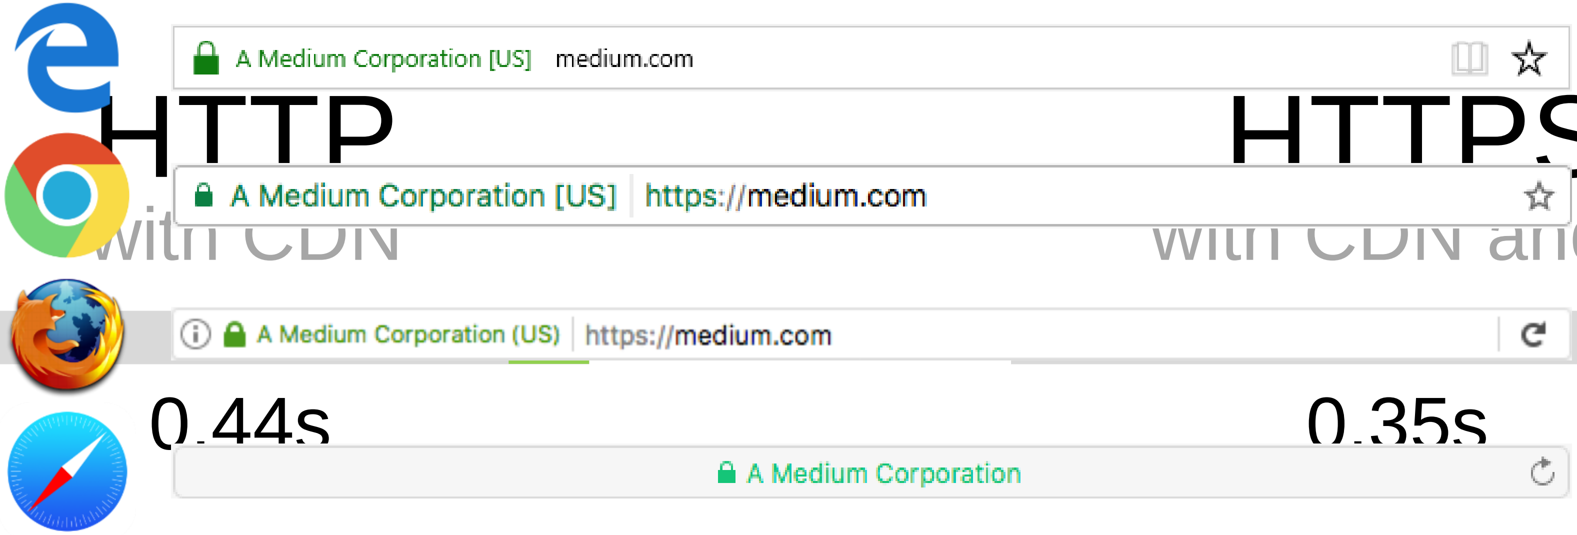

HTTP/1.1 vs. HTTP/2

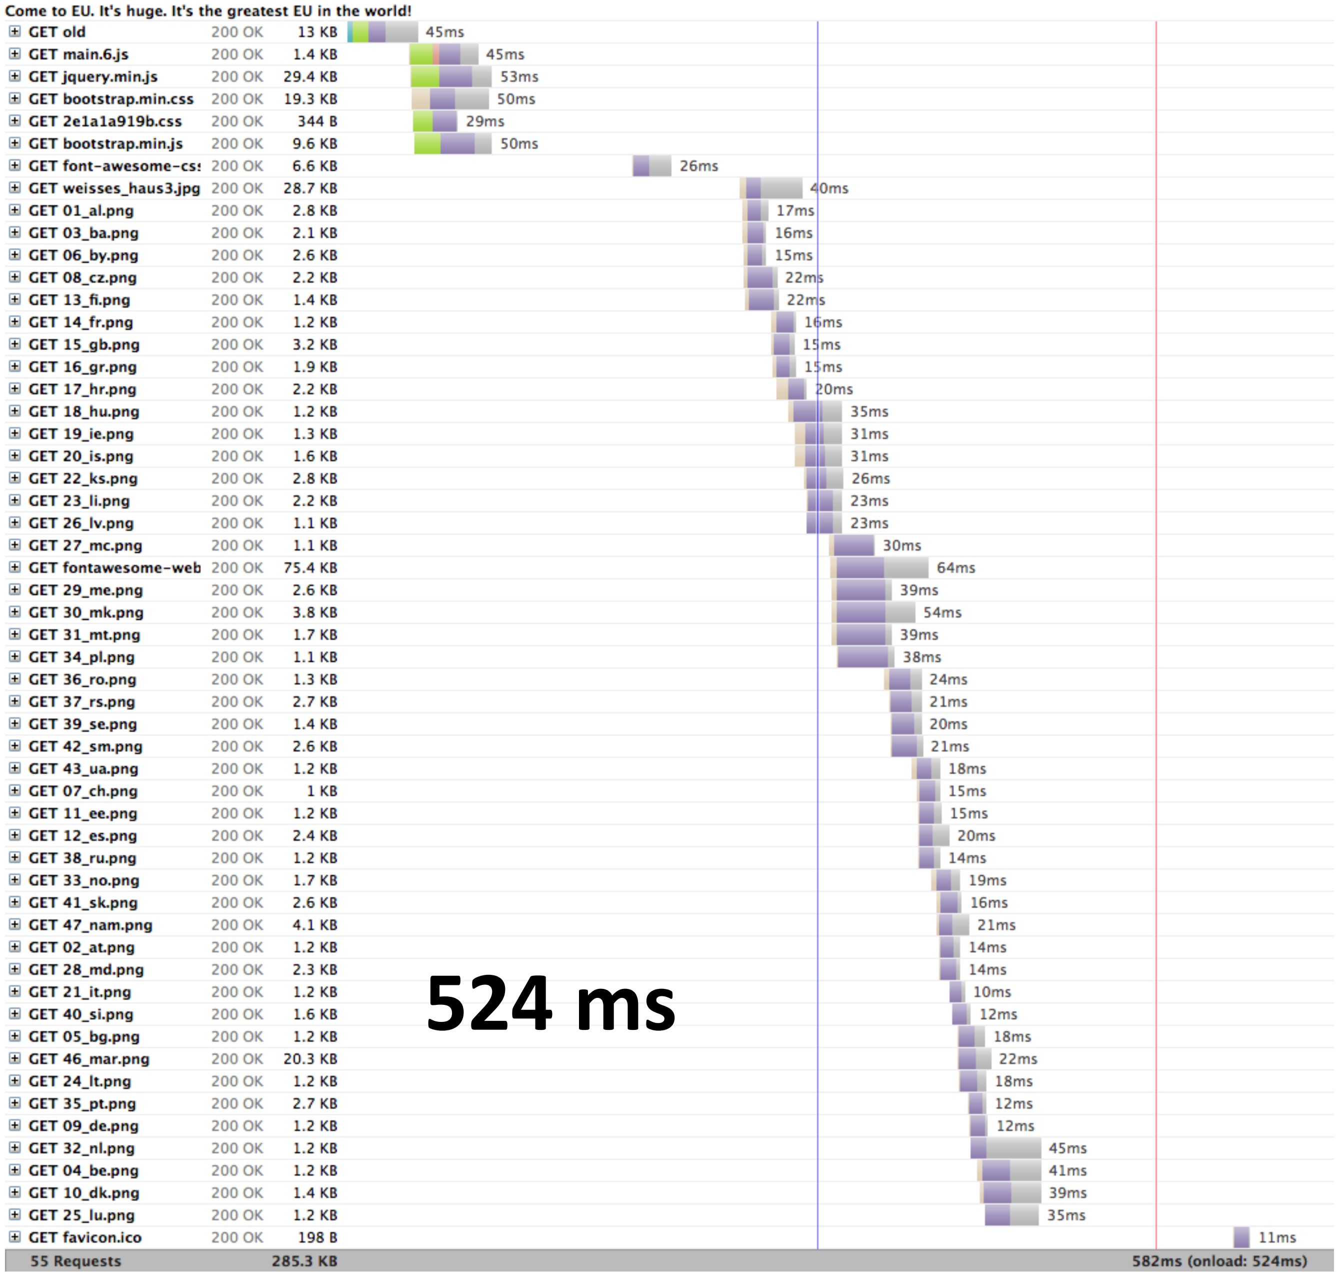

524 ms

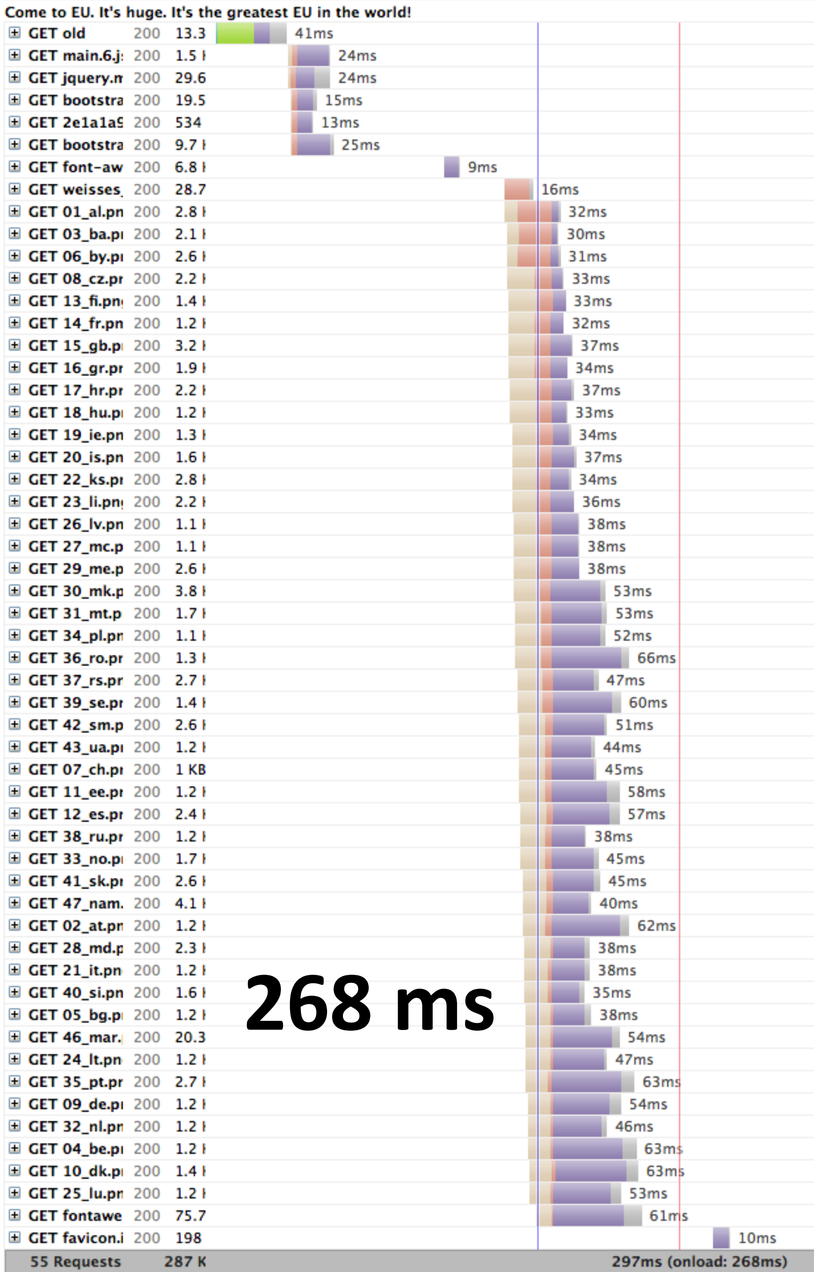

268 ms

Why CDNs Matter

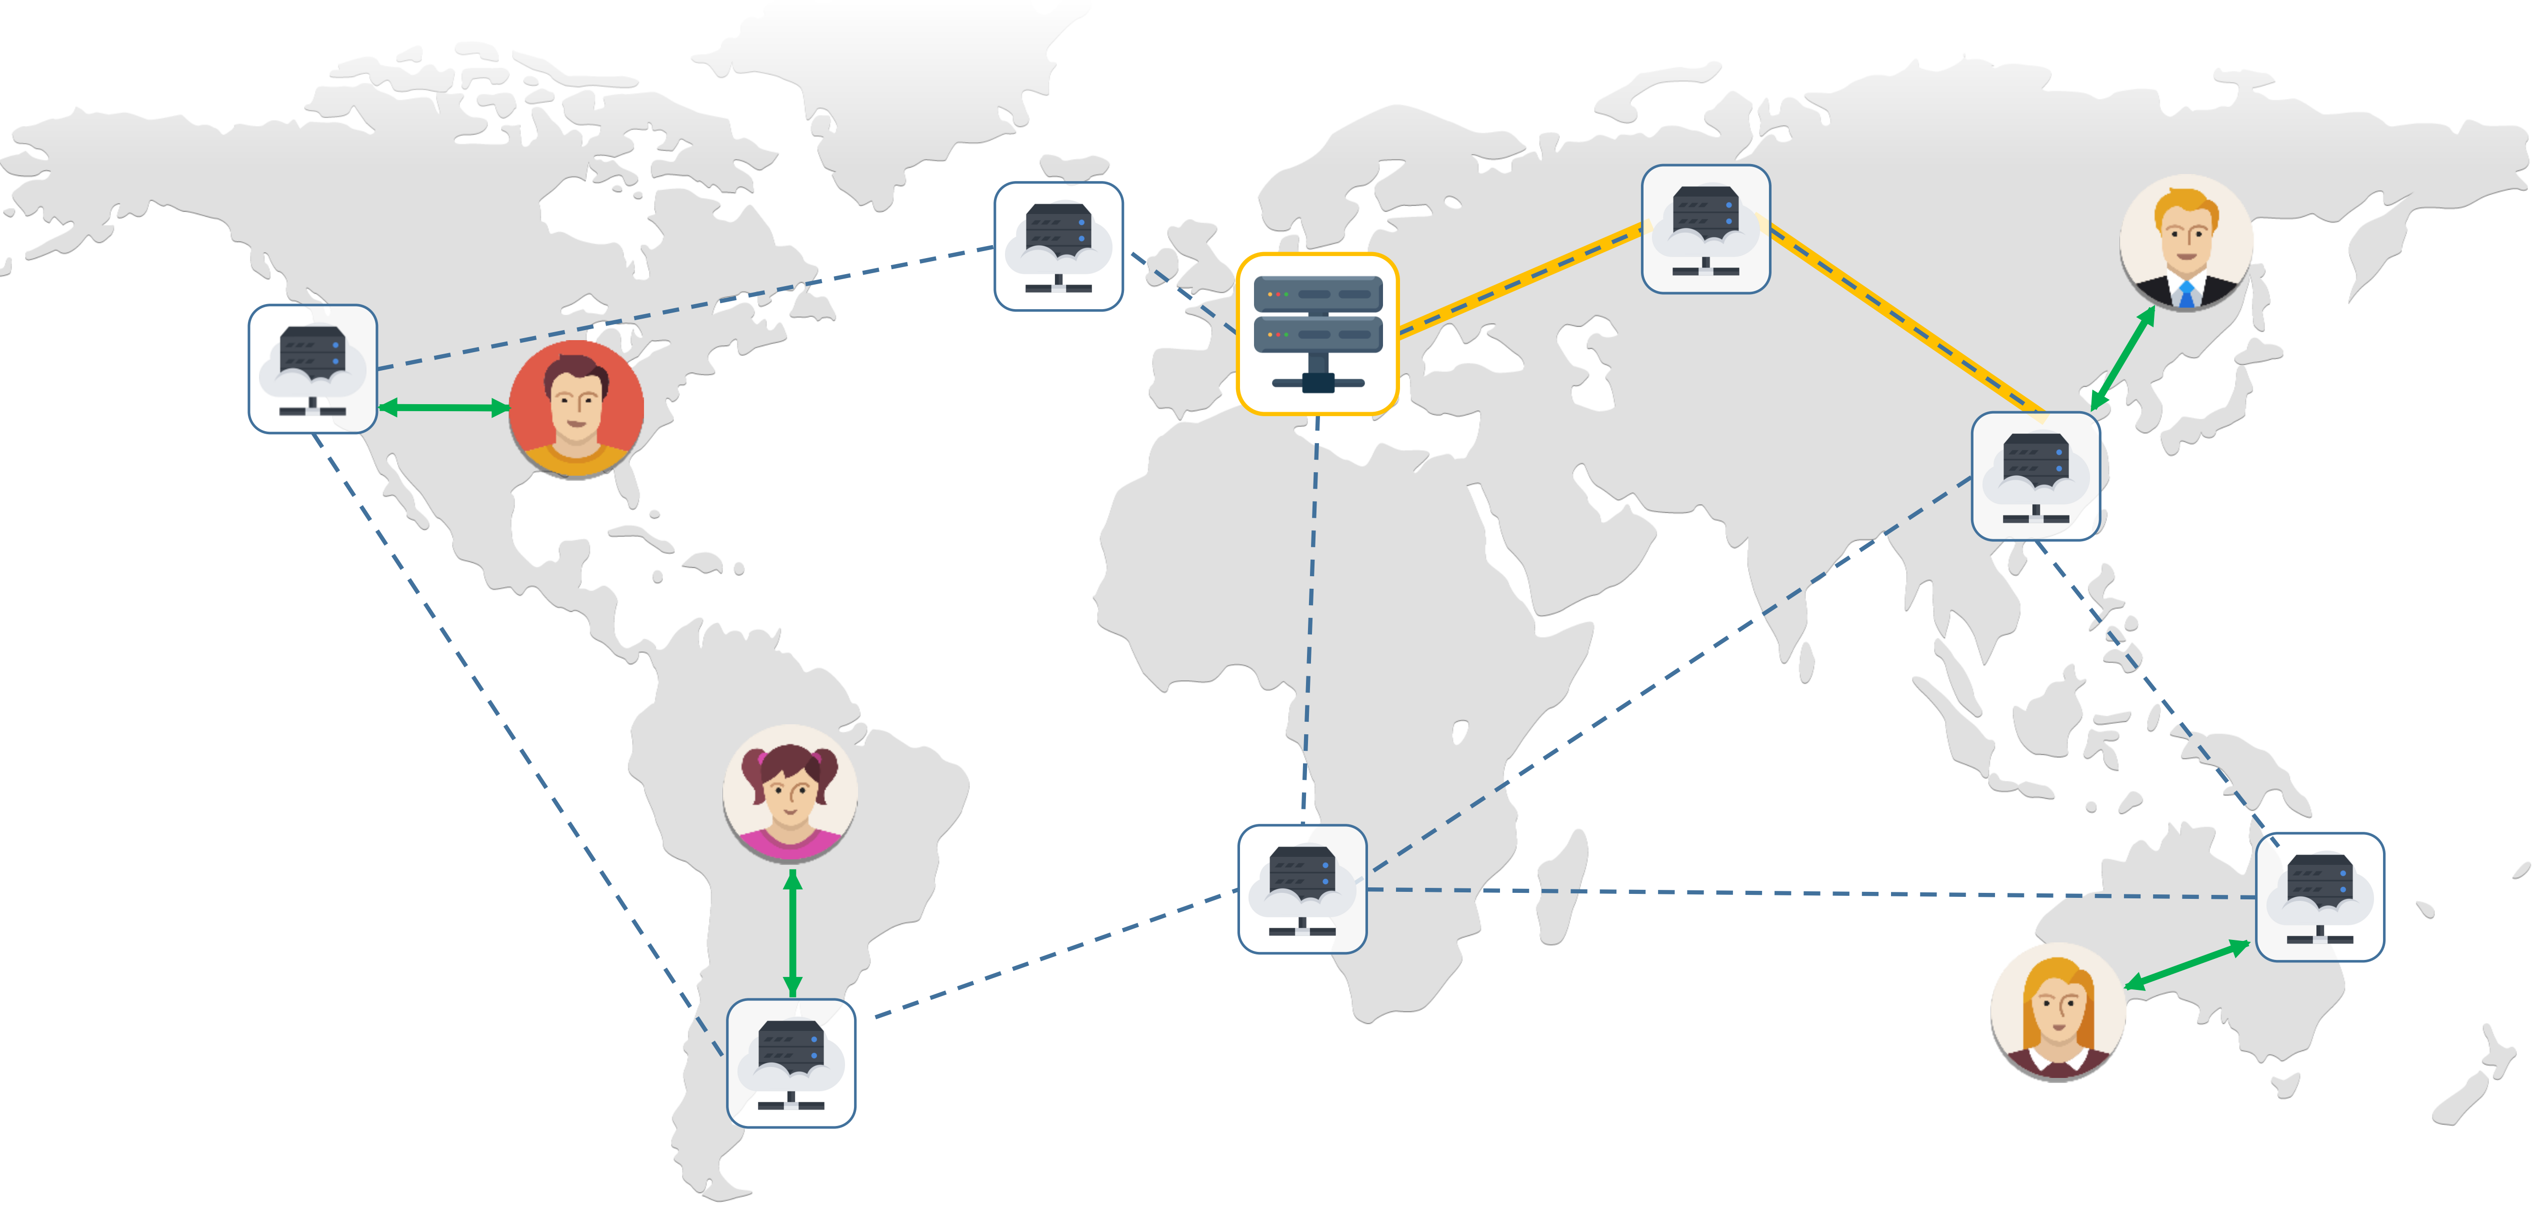

Why CDNs Matter

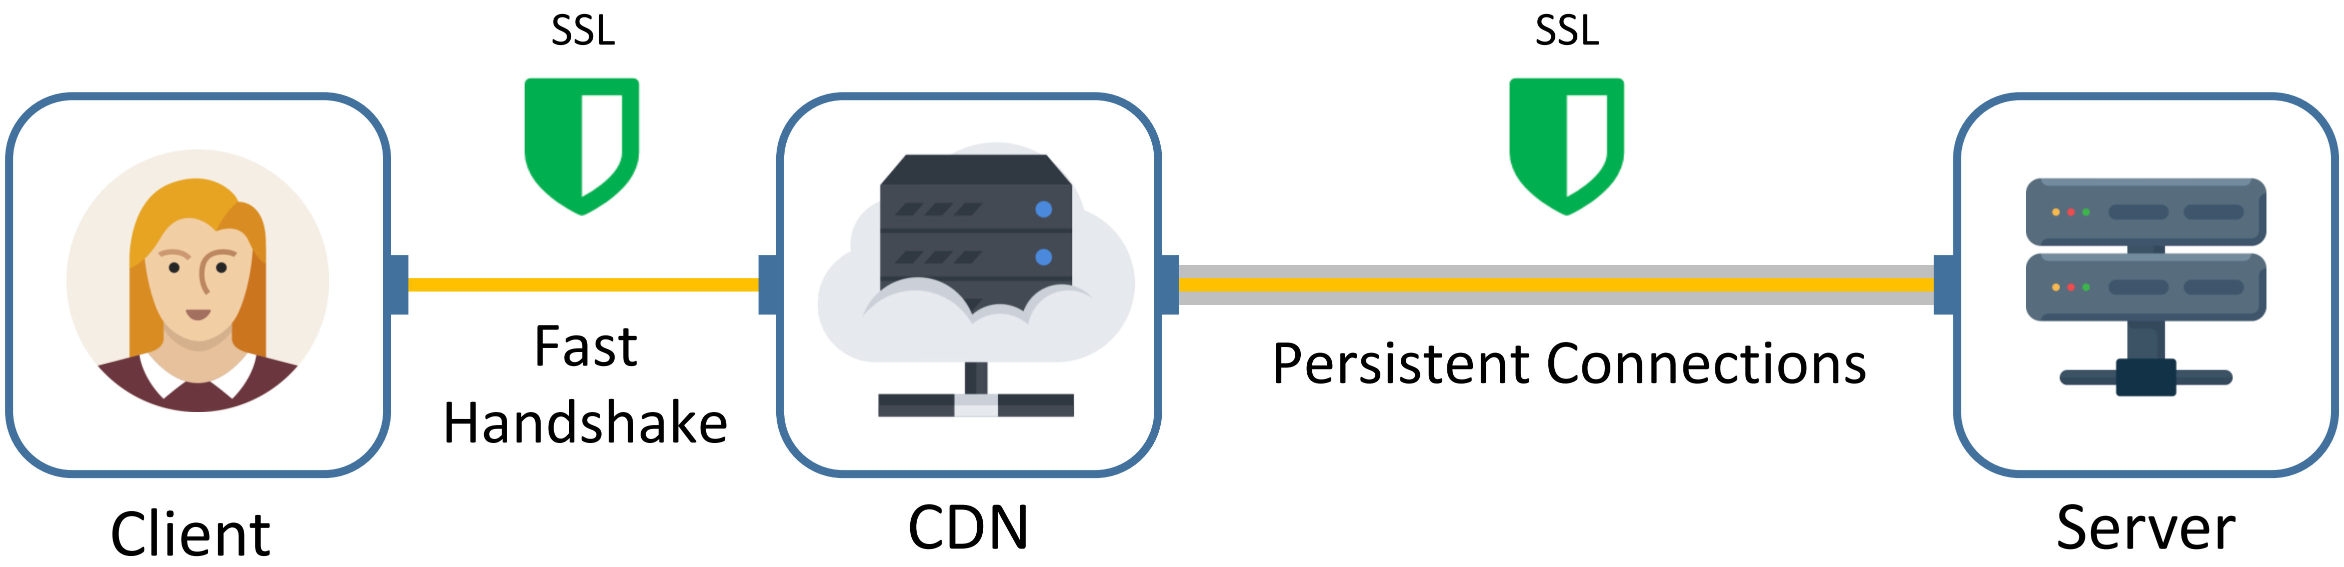

- Lower **latency** to client
- **Caching** on the edge
- **DDOS** protection
- **Failover** & Stale-on-error
- **Scalable**
- Warm **backend connections**

Network
Thinks

Avoid **redirects**, when necessary serve from CDN

Heavy **browser and CDN caching**

Persistent backend connections and **IP anycasting**

HTTP/2 with optimized **SSL and TCP**

Gzip compression for text-based files

Minimize **DNS lookups**

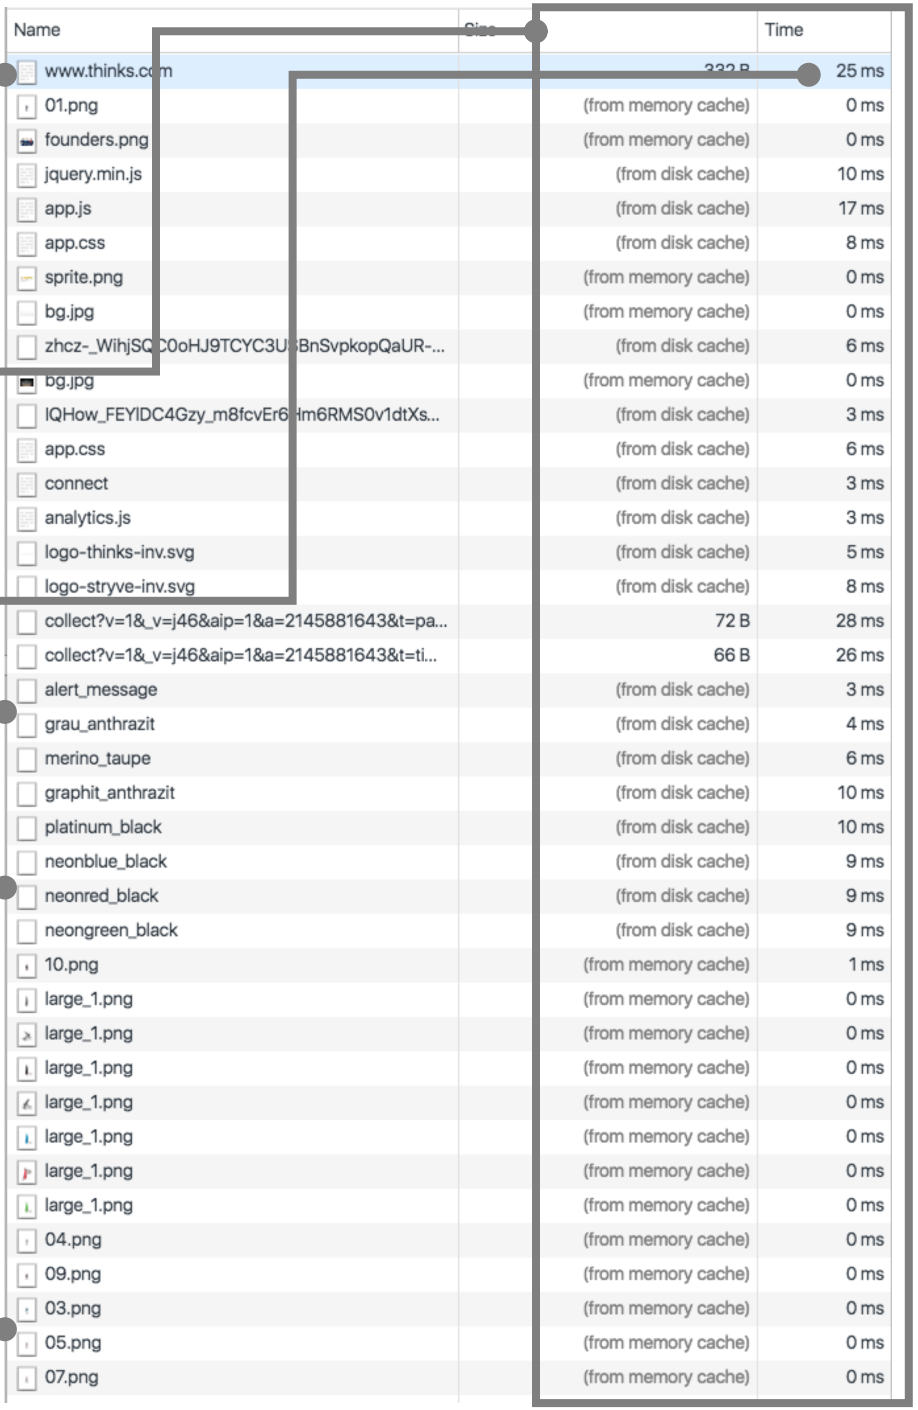

Name	Size	Time
www.thinks.com	332 B	25 ms
01.png	(from memory cache)	0 ms
founders.png	(from memory cache)	0 ms
jquery.min.js	(from disk cache)	10 ms
app.js	(from disk cache)	17 ms
app.css	(from disk cache)	8 ms
sprite.png	(from memory cache)	0 ms
bg.jpg	(from memory cache)	0 ms
zhcz-_WihjSQC0oHJ9TCYC3UBnSvpkopQaUR-...	(from disk cache)	6 ms
bg.jpg	(from memory cache)	0 ms
IQHow_FEYlDC4Gzy_m8fcvEr6Hm6RMS0v1dtXs...	(from disk cache)	3 ms
app.css	(from disk cache)	6 ms
connect	(from disk cache)	3 ms
analytics.js	(from disk cache)	3 ms
logo-thinks-inv.svg	(from disk cache)	5 ms
logo-stryve-inv.svg	(from disk cache)	8 ms
collect?v=1&_v=j46&aip=1&a=2145881643&t=pa...	72 B	28 ms
collect?v=1&_v=j46&aip=1&a=2145881643&t=ti...	66 B	26 ms
alert_message	(from disk cache)	3 ms
grau_anthrazit	(from disk cache)	4 ms
merino_taupe	(from disk cache)	6 ms
graphit_anthrazit	(from disk cache)	10 ms
platinum_black	(from disk cache)	10 ms
neonblue_black	(from disk cache)	9 ms
neonred_black	(from disk cache)	9 ms
neongreen_black	(from disk cache)	9 ms
10.png	(from memory cache)	1 ms
large_1.png	(from memory cache)	0 ms
large_1.png	(from memory cache)	0 ms
large_1.png	(from memory cache)	0 ms
large_1.png	(from memory cache)	0 ms
large_1.png	(from memory cache)	0 ms
large_1.png	(from memory cache)	0 ms
large_1.png	(from memory cache)	0 ms
04.png	(from memory cache)	0 ms
09.png	(from memory cache)	0 ms
03.png	(from memory cache)	0 ms
05.png	(from memory cache)	0 ms
07.png	(from memory cache)	0 ms

Network

Thinks

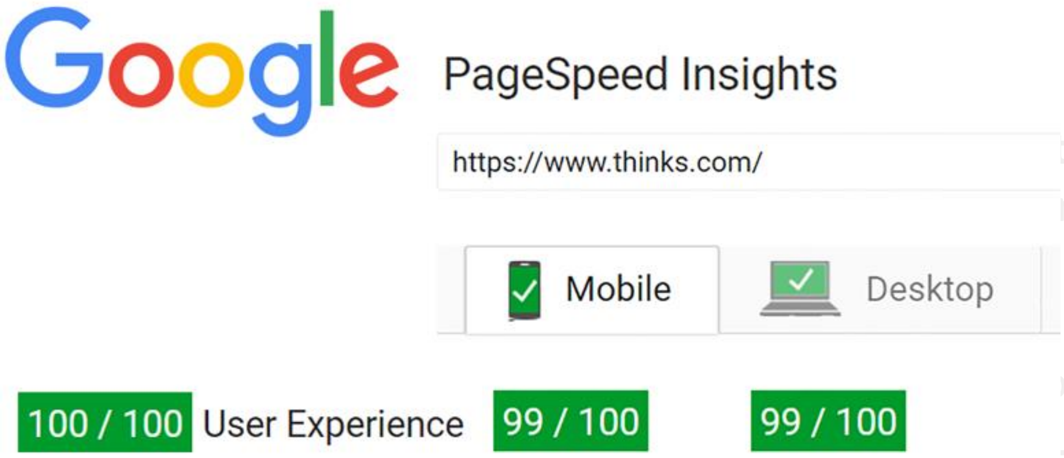

Avoid **redirects**, when necessary serve from CDN

Heavy **browser and CDN caching**

Gzip compression for text-based files

Minimize **DNS lookups**

Minus 1 point:

Google Analytics TTL too short

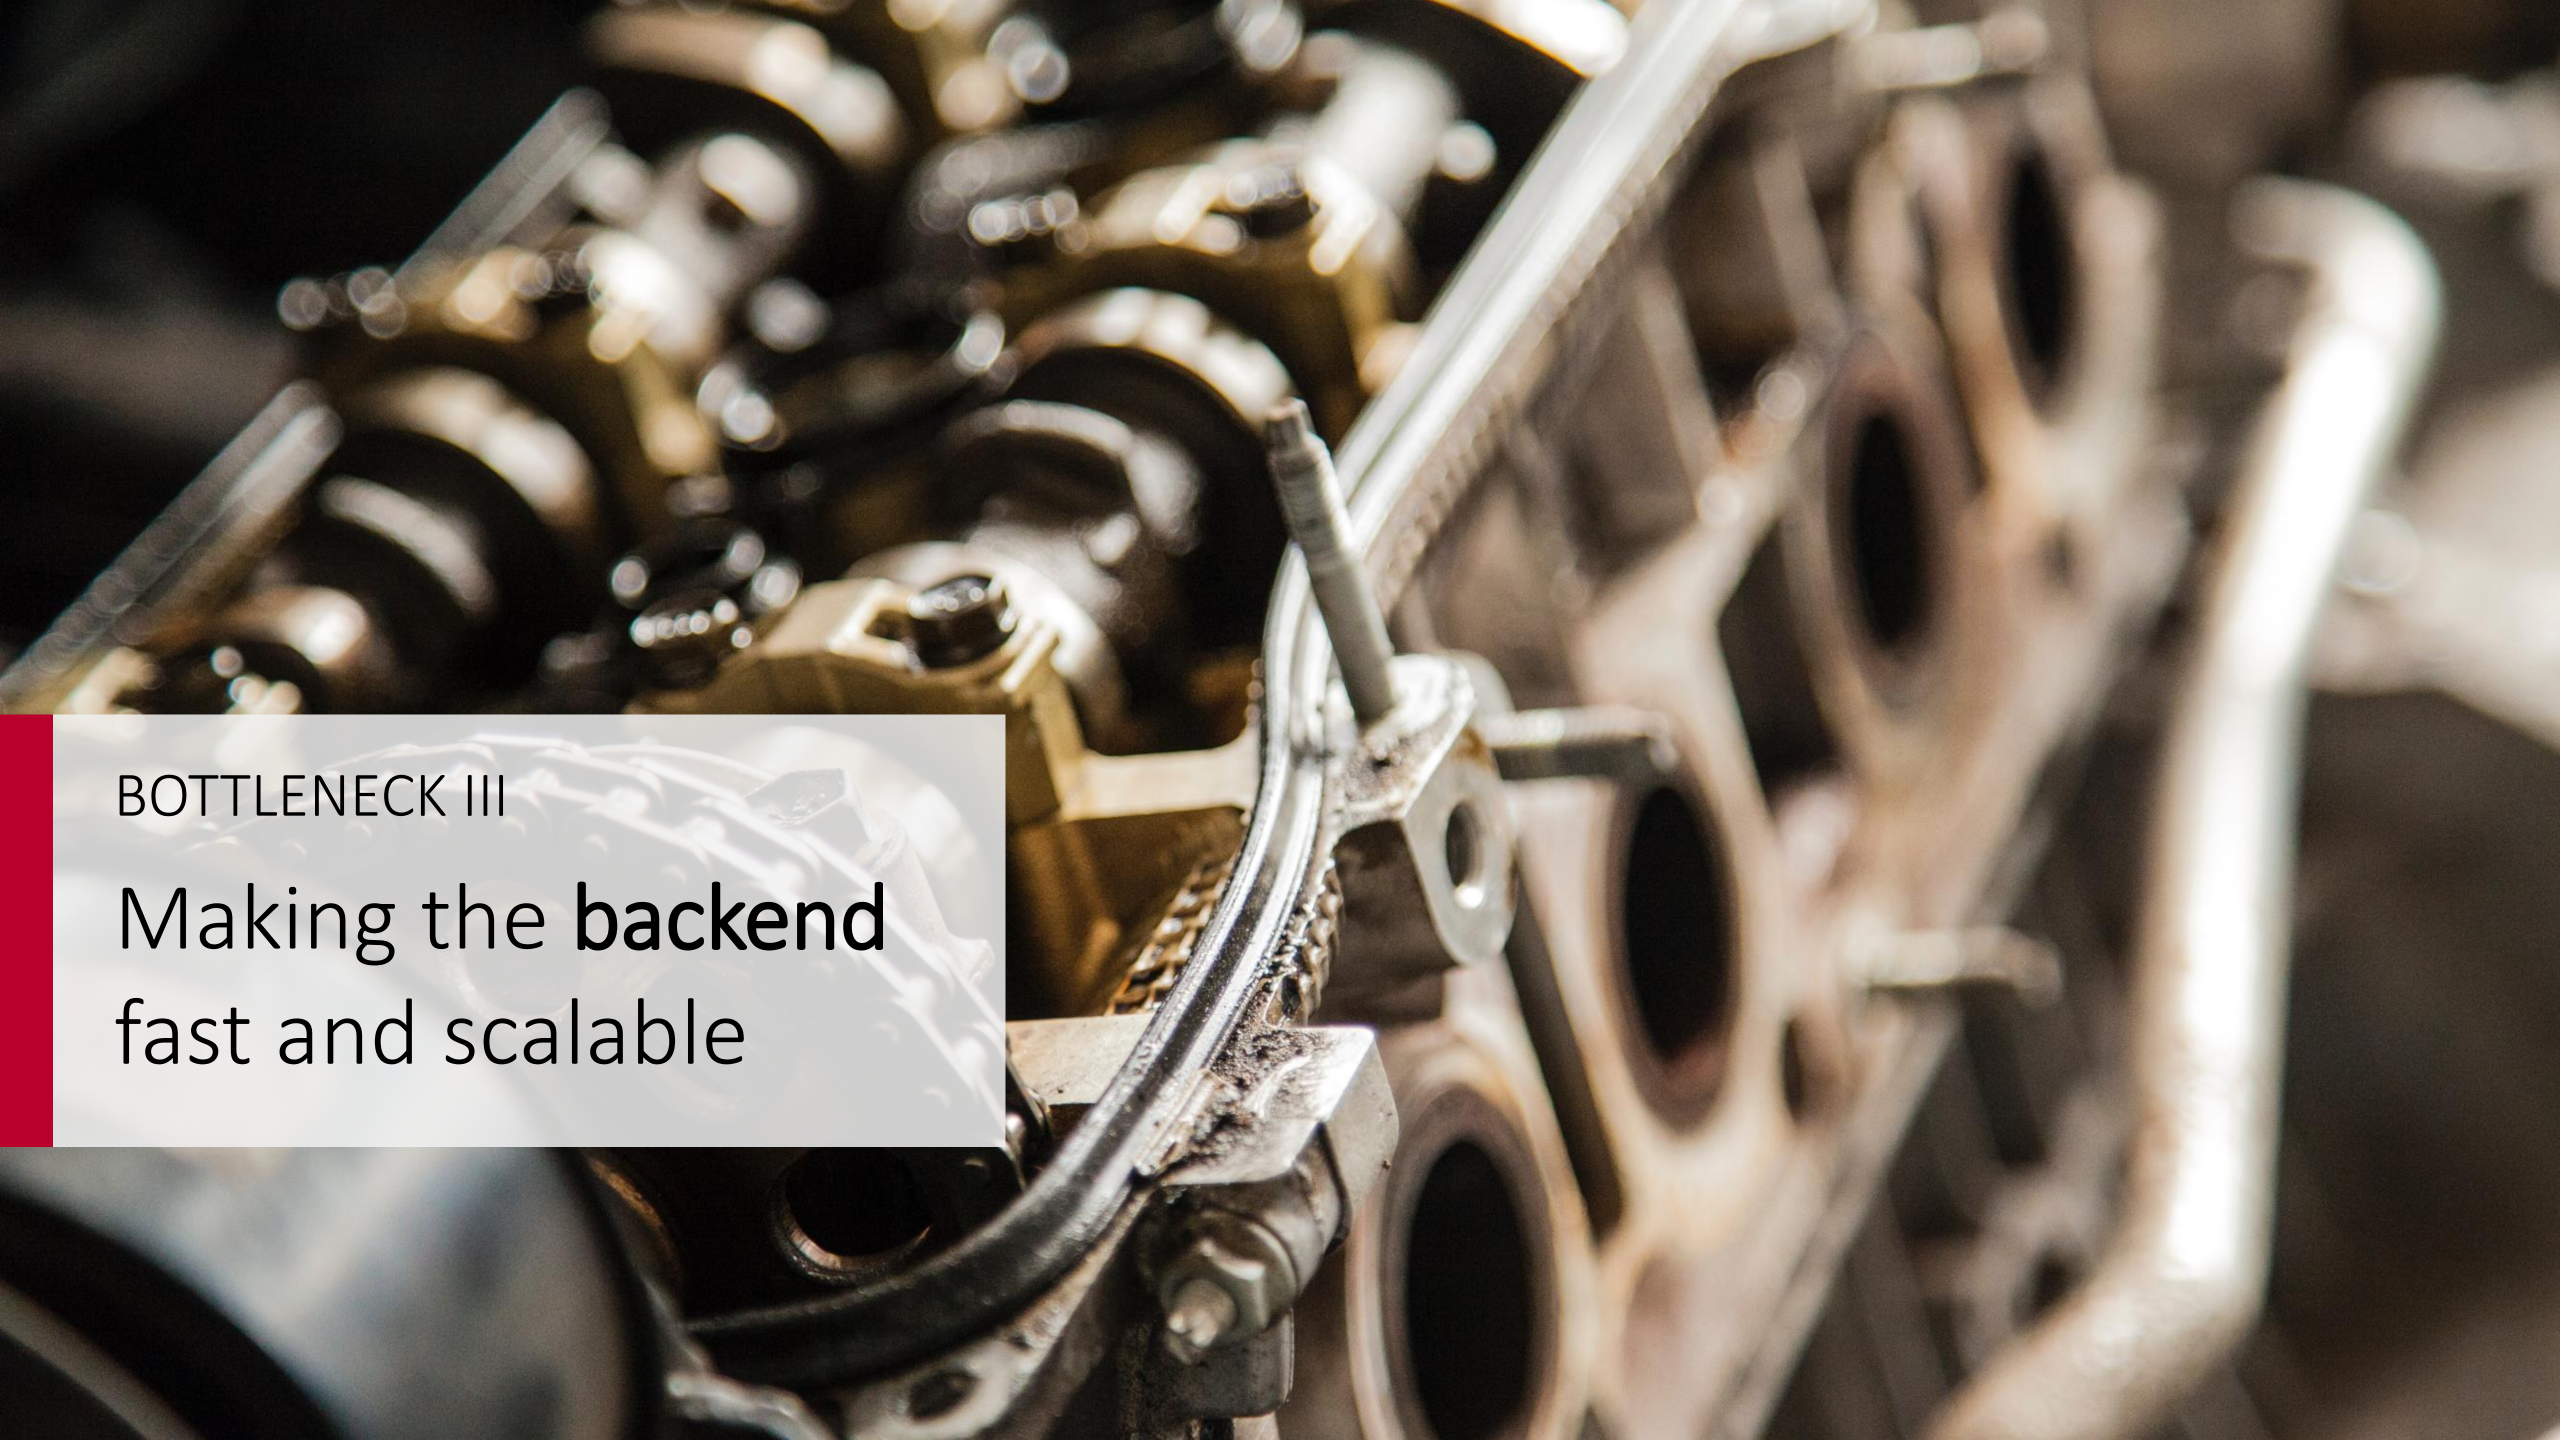

BOTTLENECK III

Making the **backend** fast and scalable

Backend Performance

Overview

Load Balancer | **Application Server** | **Database**

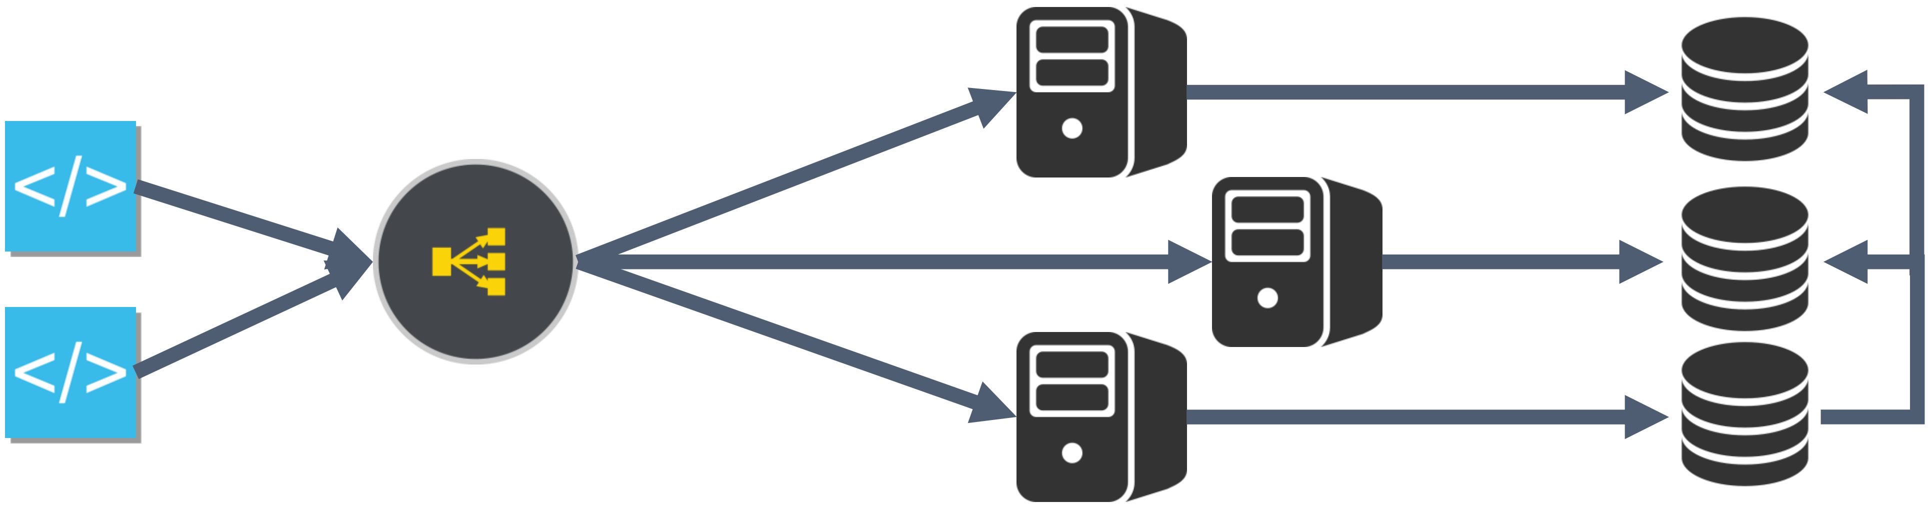

Load Balancer

- Load Balancing
- Auto-scaling
- Failover

Application Server

- Stateless Sessions
- Minimize shared state
- Efficient Code & IO

Database

- Horizontally scalable databases (e.g. “NoSQL”)
  - Replication
  - Sharding
  - Failover

Shop Backend

On Baqend Cloud

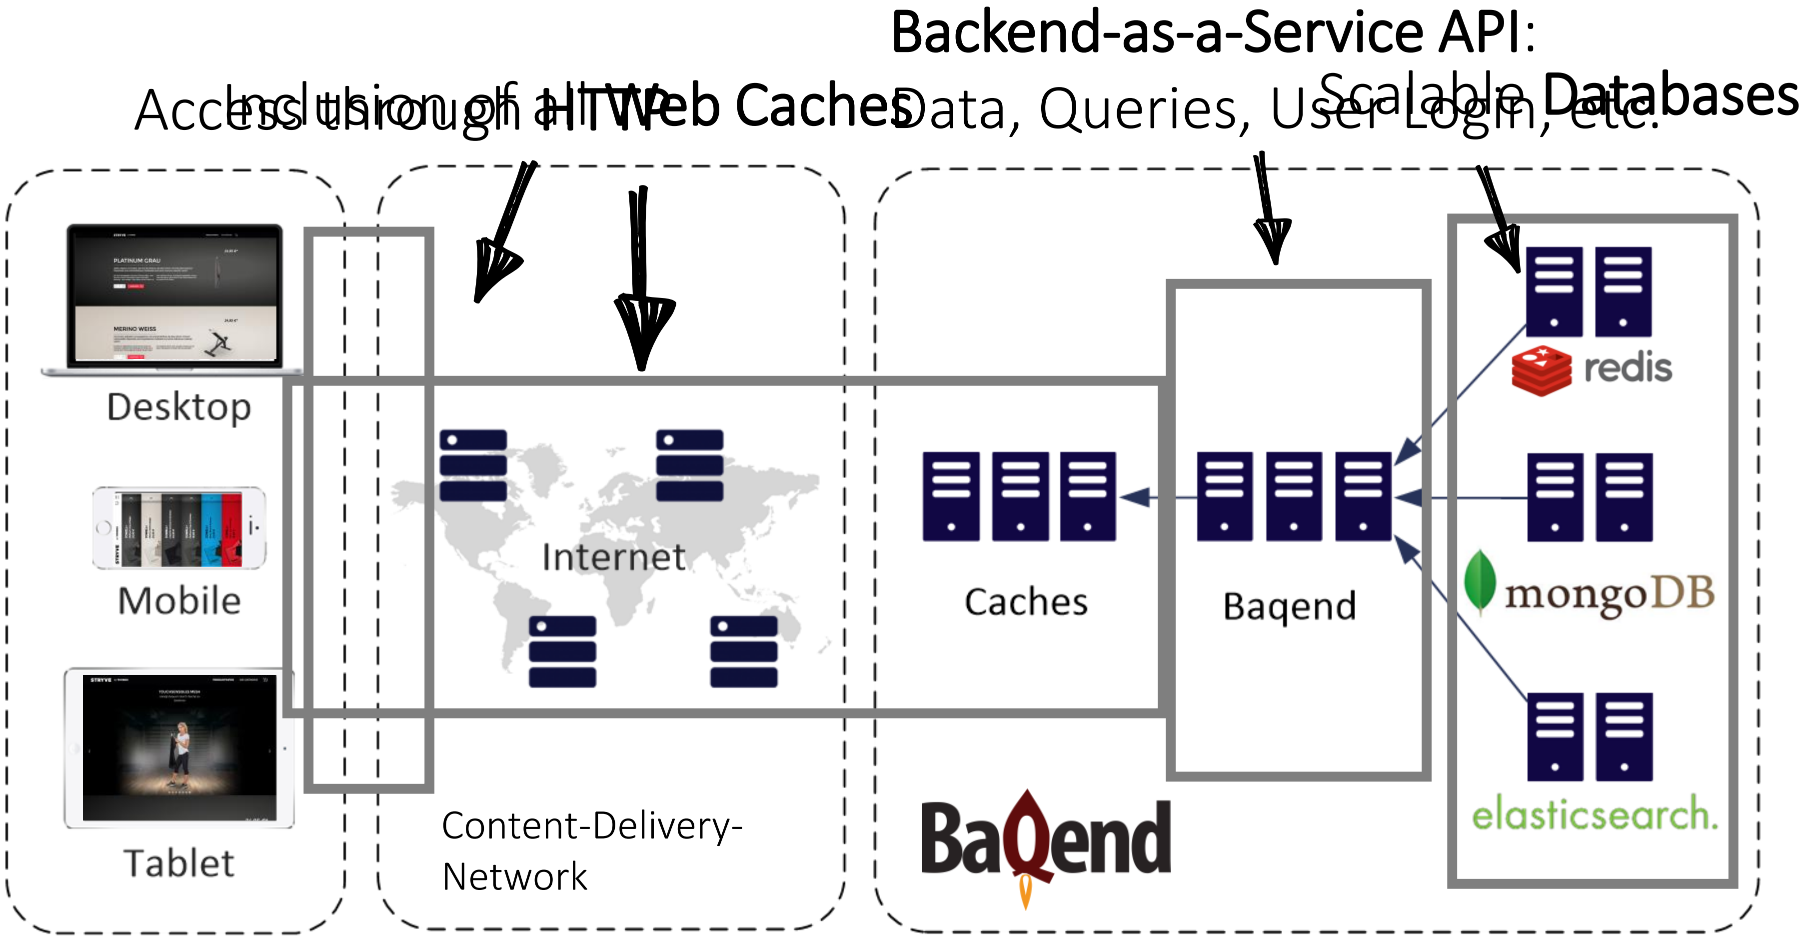

Shop Backend

On Baqend Cloud

Load Testing

```
summary + 286400 in 00:00:30 = 9543.8/s Avg:     5 Min:     1 Max:  5121 Err:     0 (0.00%) Active: 70 Started: 62370 Finished: 62351
summary = 2595820 in 00:04:33 = 9502.9/s Avg:     5 Min:     1 Max:  5201 Err:     1 (0.00%)
summary + 283600 in 00:00:30 = 9455.2/s Avg:     5 Min:     1 Max:  5110 Err:     0 (0.00%) Active: 131 Started: 69245 Finished: 69165
summary = 2879420 in 00:05:03 = 9498.2/s Avg:     5 Min:     1 Max:  5201 Err:     1 (0.00%)
summary + 288500 in 00:00:30 = 9619.6/s Avg:     7 Min:     1 Max:  5150 Err:     0 (0.00%) Active: 66 Started: 76115 Finished: 76100
summary = 3167920 in 00:05:33 = 9509.1/s Avg:     5 Min:     1 Max:  5201 Err:     1 (0.00%)
summary + 287000 in 00:00:30 = 9565.7/s Avg:     5 Min:     1 Max:  5126 Err:     0 (0.00%) Active: 52 Started: 82990 Finished: 82989
summary = 3454920 in 00:06:03 = 9513.8/s Avg:     5 Min:     1 Max:  5201 Err:     1 (0.00%)
summary + 285100 in 00:00:30 = 9502.4/s Avg:     5 Min:     1 Max:  5142 Err:     0 (0.00%) Active: 65 Started: 89861 Finished: 89847
summary = 3740020 in 00:06:33 = 9512.9/s Avg:     5 Min:     1 Max:  5201 Err:     1 (0.00%)
summary + 287100 in 00:00:30 = 9563.9/s Avg:     5 Min:     1 Max:  5083 Err:     0 (0.00%) Active: 46 Started: 96737 Finished: 96742
summary = 4027120 in 00:07:03 = 9516.6/s Avg:     5 Min:     1 Max:  5201 Err:     1 (0.00%)
summary + 285200 in 00:00:30 =                                                   0 (0.00%) Active: 47 Started: 103604 Finished: 103608
summary = 4312320 in 00:07:33 =                                                   1 (0.00%)
summary + 285400 in 00:00:30 =                                                   1 (0.00%) Active: 70 Started: 110479 Finished: 110460
summary = 4597720 in 00:08:03 =                                                   2 (0.00%)
summary + 286000 in 00:00:30 =                                                   0 (0.00%) Active: 47 Started: 117350 Finished: 117354
summary = 4883720 in 00:08:33 =                                                   2 (0.00%)
summary + 285600 in 00:00:30 =                                                   0 (0.00%) Active: 46 Started: 124222 Finished: 124227
summary = 5169320 in 00:09:03 =                                                   2 (0.00%)
summary + 286200 in 00:00:30 =                                                   0 (0.00%) Active: 44 Started: 131091 Finished: 131098
summary = 5455520 in 00:09:33 =                                                   2 (0.00%)
summary + 286100 in 00:00:30 =                                                   0 (0.00%) Active: 48 Started: 137969 Finished: 137972
summary = 5741620 in 00:10:03 =                                                   2 (0.00%)
summary + 285700 in 00:00:30 =                                                   0 (0.00%) Active: 52 Started: 144838 Finished: 144837
summary = 6027320 in 00:10:33 =                                                   2 (0.00%)
summary + 285500 in 00:00:30 =                                                   0 (0.00%) Active: 47 Started: 151709 Finished: 151713
summary = 6312820 in 00:11:03 =                                                   2 (0.00%)
summary + 285300 in 00:00:30 = 9504.3/s              Min:     1 Max:  5169 Err:     0 (0.00%) Active: 46 Started: 158579 Finished: 158584
summary = 6598120 in 00:11:33 = 9518.9/s Avg:         Min:     1 Max:  5230 Err:     2 (0.00%)
summary + 285900 in 00:00:30 = 9533.5/s Avg:          Min:     1 Max:  1002 Err:     0 (0.00%) Active: 47 Started: 165449 Finished: 165453
summary = 6884020 in 00:12:03 = 9519.5/s Avg:     5 Min:     1 Max:  5230 Err:     2 (0.00%)
```

6.8 M Requests

10K/s

Latency 5ms

Performance: State of the Art

Summarized

Frontend | **Network** | **Backend**

- Doable with **best practices** and support through **build tools**
- Caching and CDNs help, but only for **static content**
- Many frameworks and platforms, but horizontal **scalability** is very difficult

https://www.ampproject.org/docs/reference/spec.html

Accelerated Mobile Pages

Google's Approach for a Faster Web

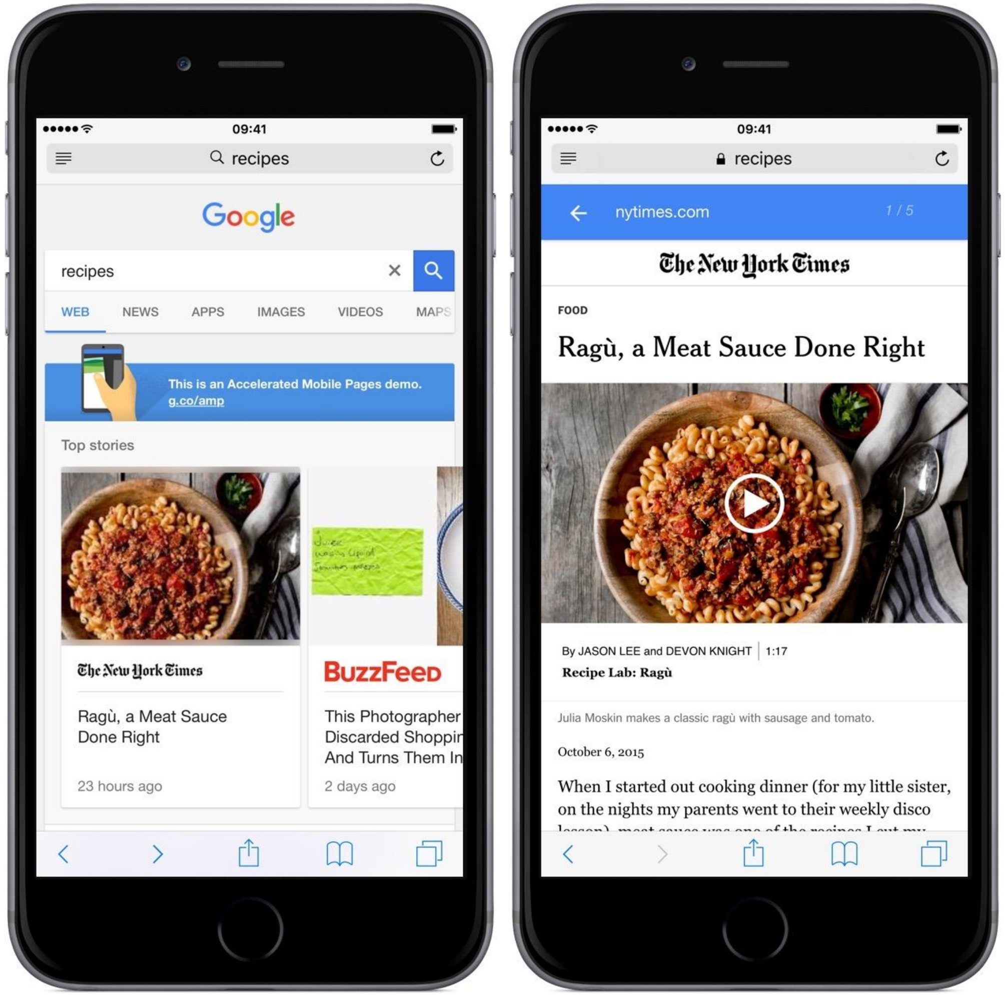

How AMP works:

- **Stripped down** HTML + AMP tags (e.g img) → rendered asynchronously by AMP runtime
- CSS must be **inlined + <50 KB +**
- No **custom JS** (except in iframes)
- Only static sizes → **no repaints**
- **Cached** in Google CDN, as long as it is crawled the next time → only suited for **static** media, e.g. news

How to apply these techniques for ***any* website**?

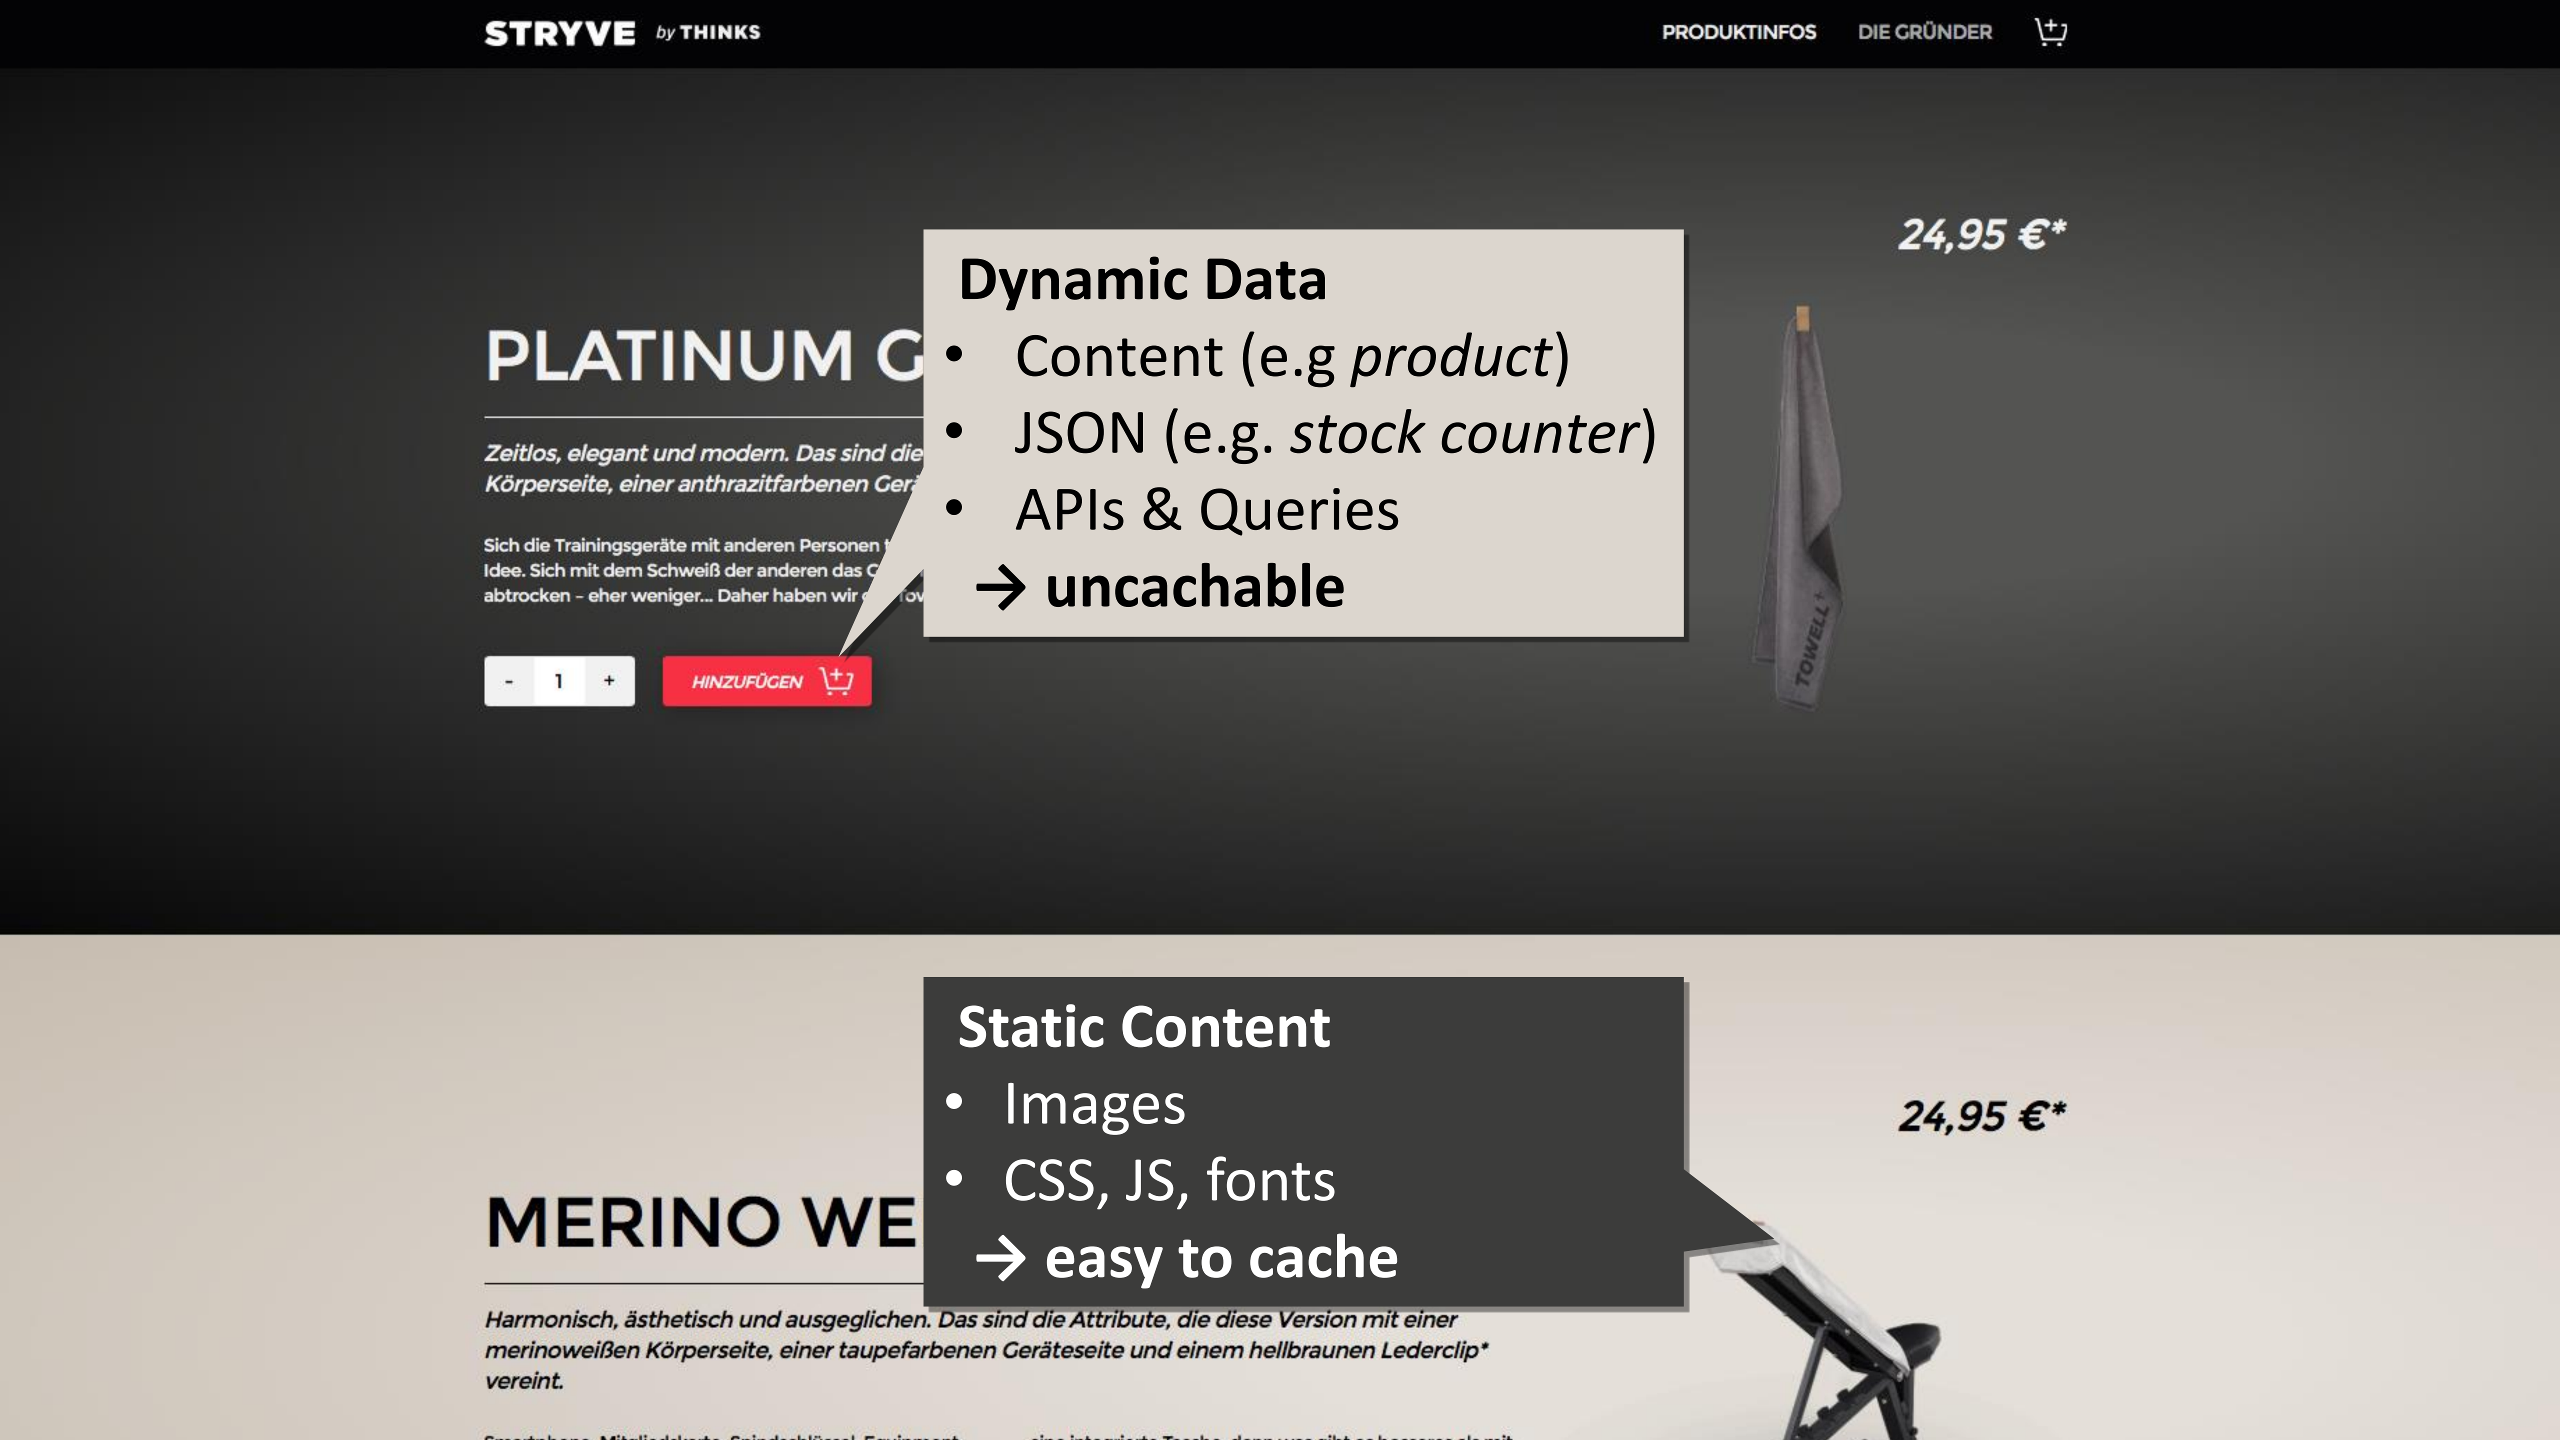

Dynamic Data

- Content (e.g *product*)
- JSON (e.g. *stock counter*)
- APIs & Queries

→ **uncachable**

Static Content

- Images
- CSS, JS, fonts

→ **easy to cache**

Goal: Low-Latency for Dynamic Content

By Serving Data from Ubiquitous Web Caches

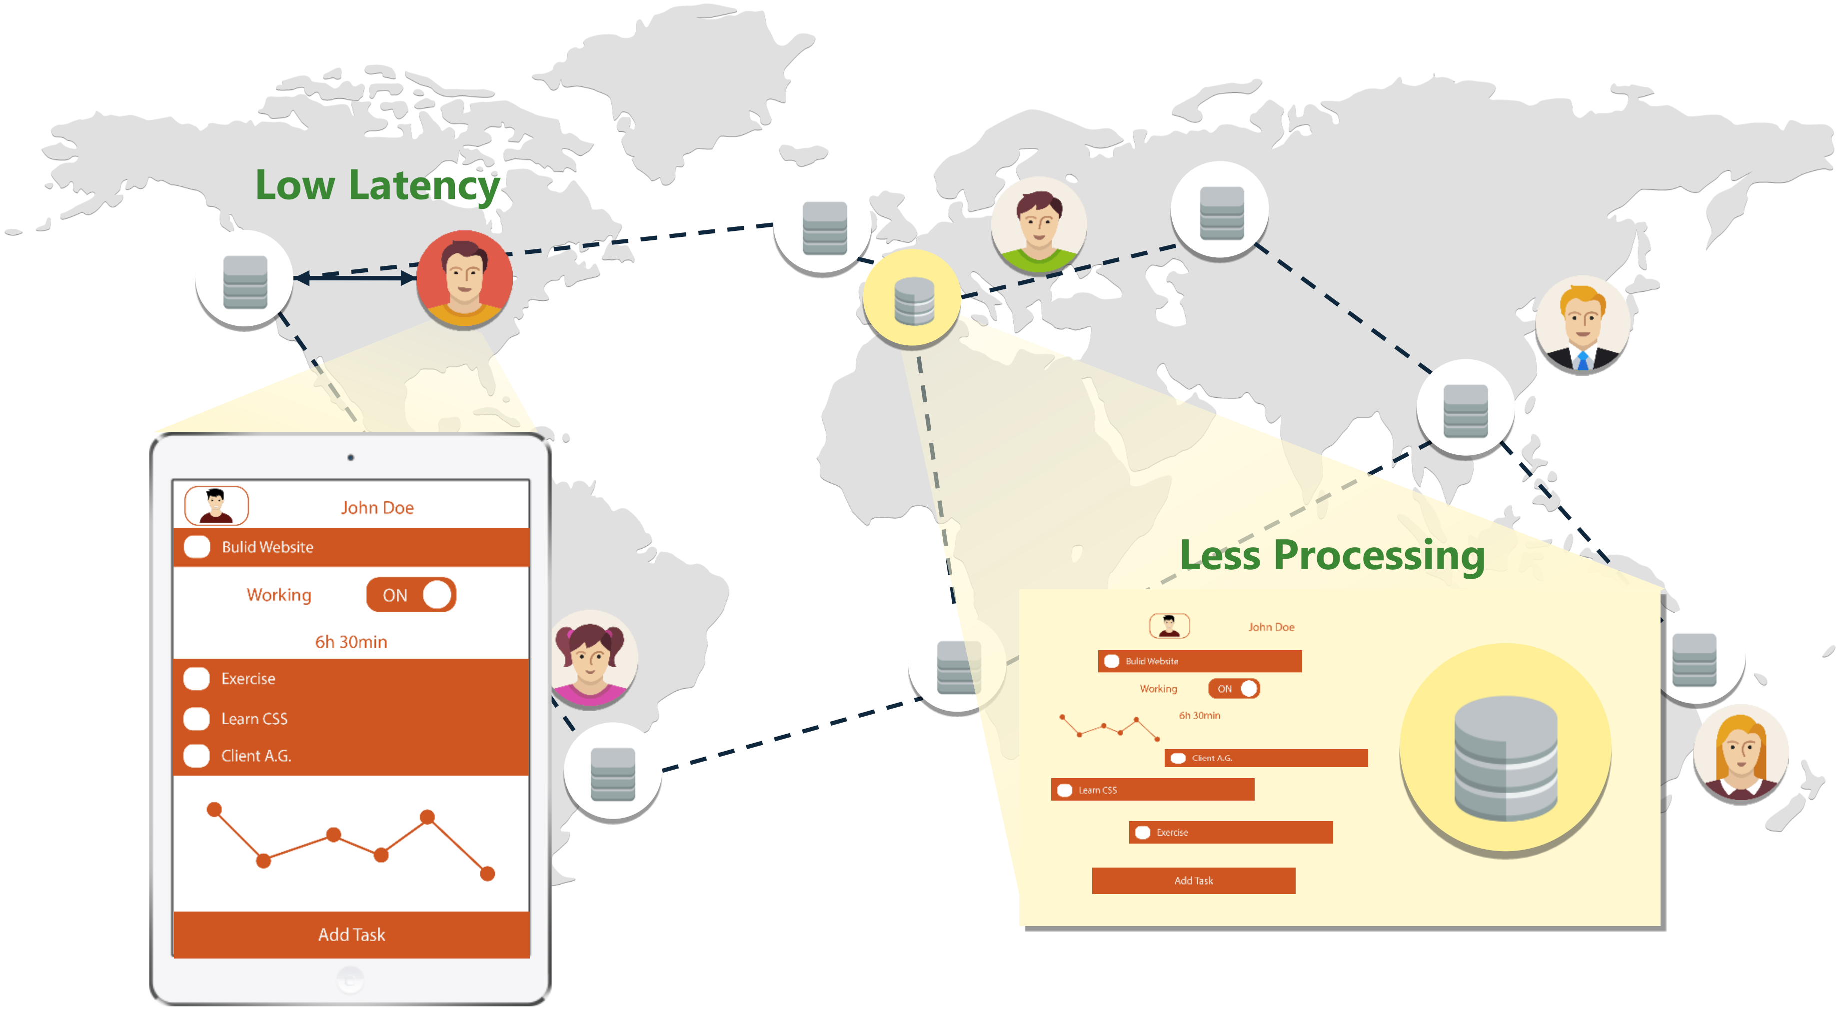

How does this work?

Dynamic Caching in Detail

Has **Time-to-Live** (expiration)

False-Positive Rate: $f \approx \left(1 - e^{-\frac{kn}{m}}\right)^k$

Hash-Functions: $k = \left\lceil \ln(2) \cdot \left(\frac{n}{m}\right) \right\rceil$

With 20.000 entries and a 5% false positive rate: **11 Kbyte**

Consistency: Δ-Atomicity, Read-Your-Writes, Monotonic Reads, Monotonic Writes, Causal Consistency

How does this work?

Dynamic Caching in Detail

F. Gessert, F. Bücklers, und N. Ritter, „ORESTES: a Scalable Database-as-a-Service Architecture for Low Latency“, in *CloudDB 2014*, 2014.

F. Gessert, S. Friedrich, W. Wingerath, M. Schaarschmidt, und N. Ritter, „Towards a Scalable and Unified REST API for Cloud Data Stores“, in *44. Jahrestagung der GI*, Bd. 232, S. 723–734.

F. Gessert und F. Bücklers, „ORESTES: e
Cloud-Datenbanken“, in Informatiktage

ath, S. Friedrich, und N. Ritter, „The
Caching in the Age of Cloud Data

F. Gessert und F. Bücklers, *Performanz-*
vermittels der Web-Caching-Hierarchie

eb-Caching von Datenbankobjekten im

M. Schaarschmidt, F. Gessert, und N. R
Persistence“, in BTW 2015.

t, „Who Watches the Watchmen? On
rking“, in BTW 2015.

S. Friedrich, W. Wingerath, F. Gessert,
Survey“, in *44. Jahrestagung der Gesell*
704.

d-Datenbanken in Forschung und

7 Years
Research & Development

Universität Hamburg

W. Wingerath, F. Gessert, S. Friedrich, N. Ritter „Real-time stream processing for Big Data“, *Big Data Analytics it - Information Technology*, 2016

F. Gessert, N. Ritter „Scalable Data Management: NoSQL Data Stores in Research and Practice“, *32nd IEEE International Conference on Data Engineering, ICDE*, 2016

F. Gessert, W. Wingerath, S. Friedrich, N. Ritter “NoSQL Database Systems: A Survey and Decision Guidance”, *Computer Science - Research and Development*, 2016

F. Gessert, N. Ritter „Polyglot Persistence“, *Datenbank Spektrum*, 2016.

From Objects to Queries

Query Caching

Real-Time
Queries
(*Websockets*)

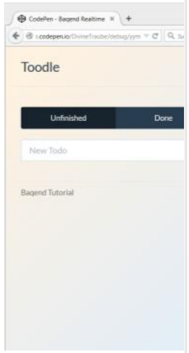

How to **cache queries**:
„Give me the most popular products that are in stock.“

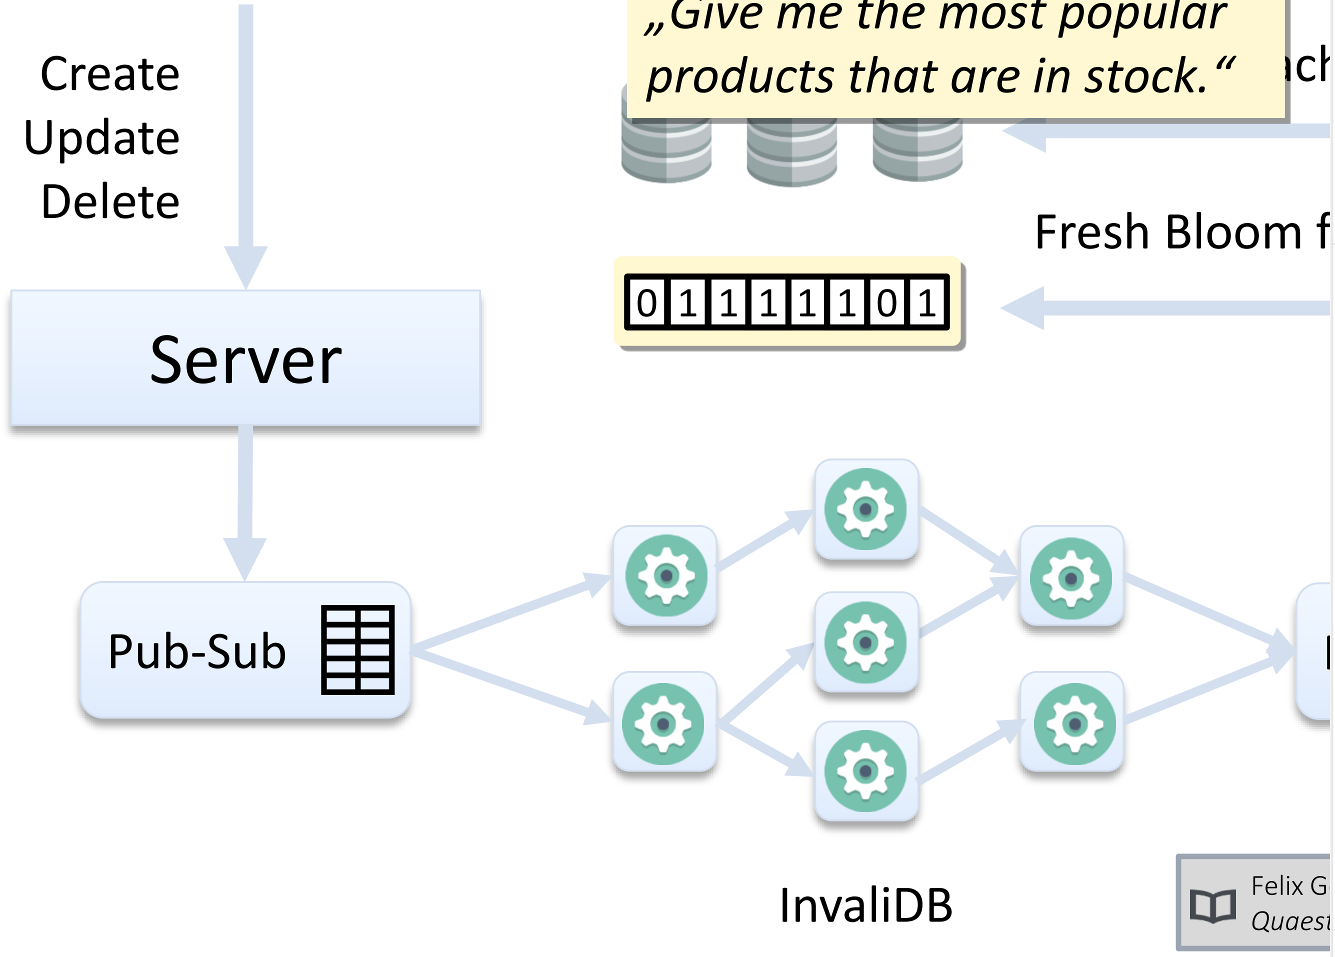

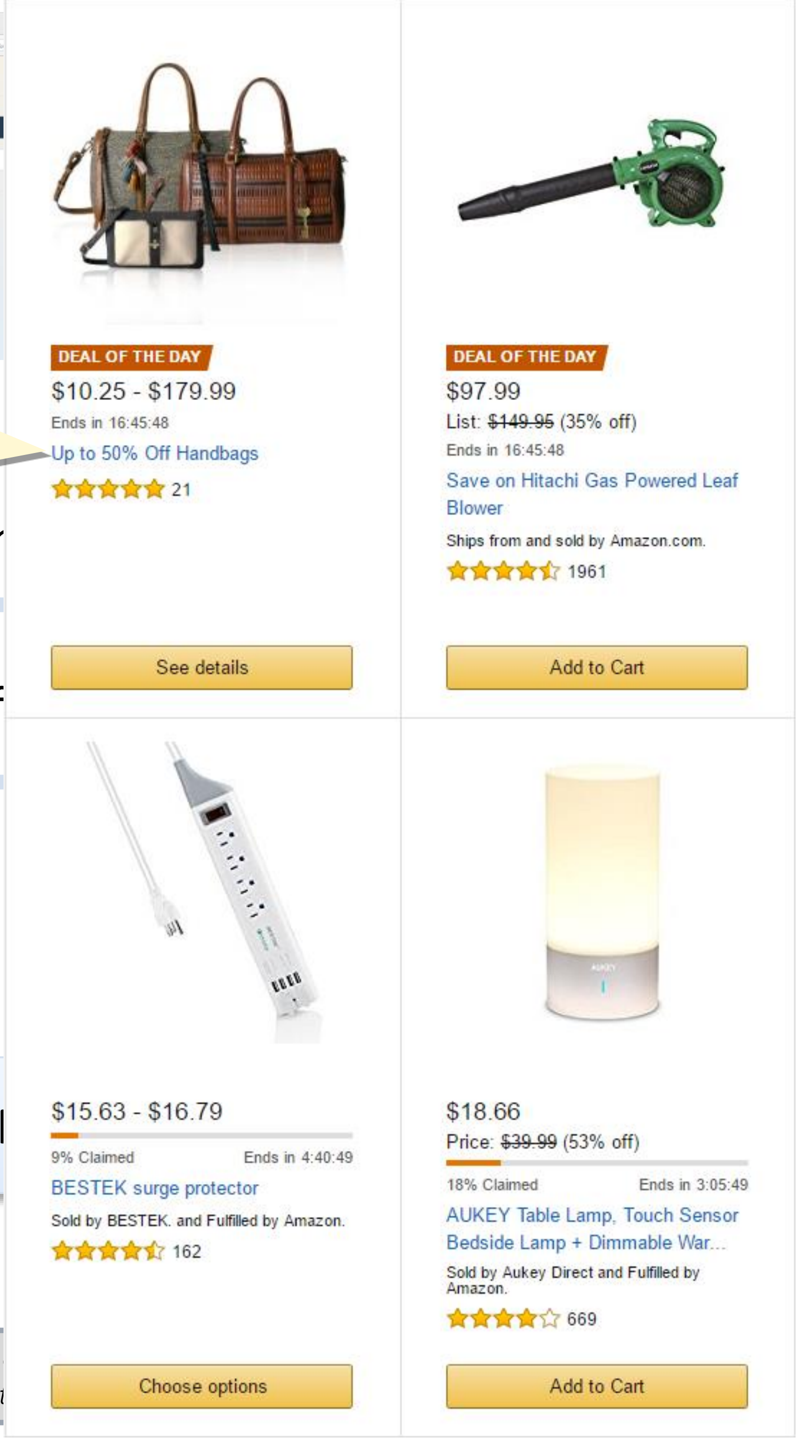

Developing a Shop

On Baqend

Dashboard

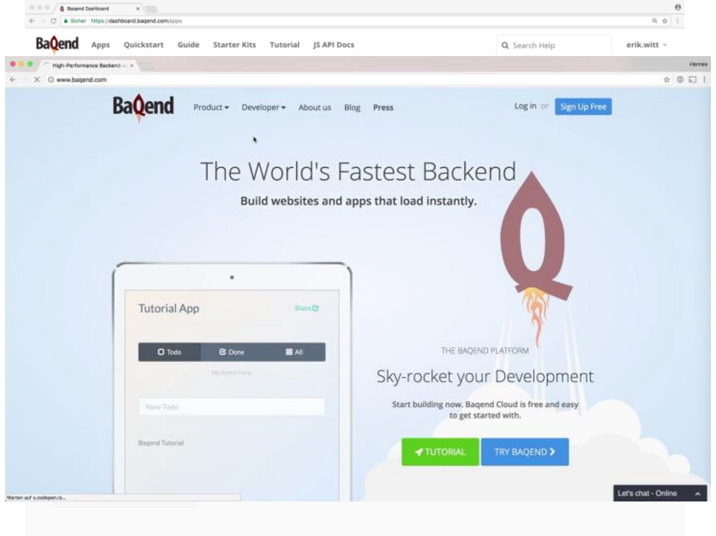

- Schema (e.g. products)
- Data (e.g. orders)
- Users

CLI

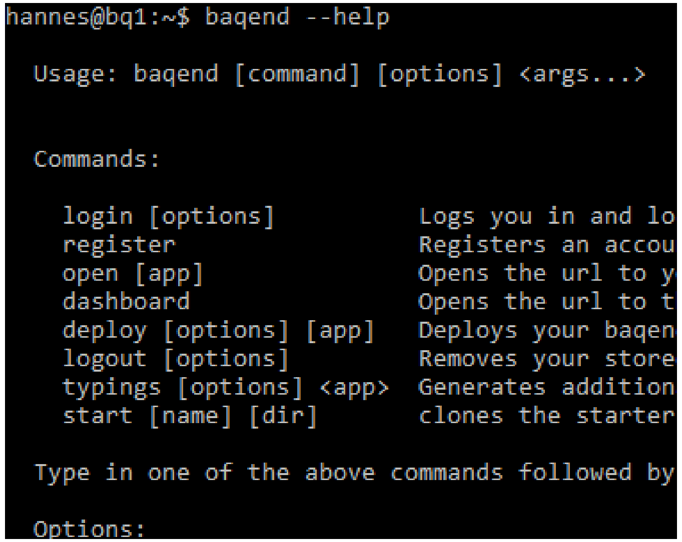

- Deploy shop frontend and backend code
- Run tests & staging

REST & SDK

- Load products
- Call checkout

E-COMMERCE

How to make **legacy shops** fast?

New Standard: Service Workers

What they do:

- **Proxy** any HTTP request
- **Offline Cache** for „progessive web apps“

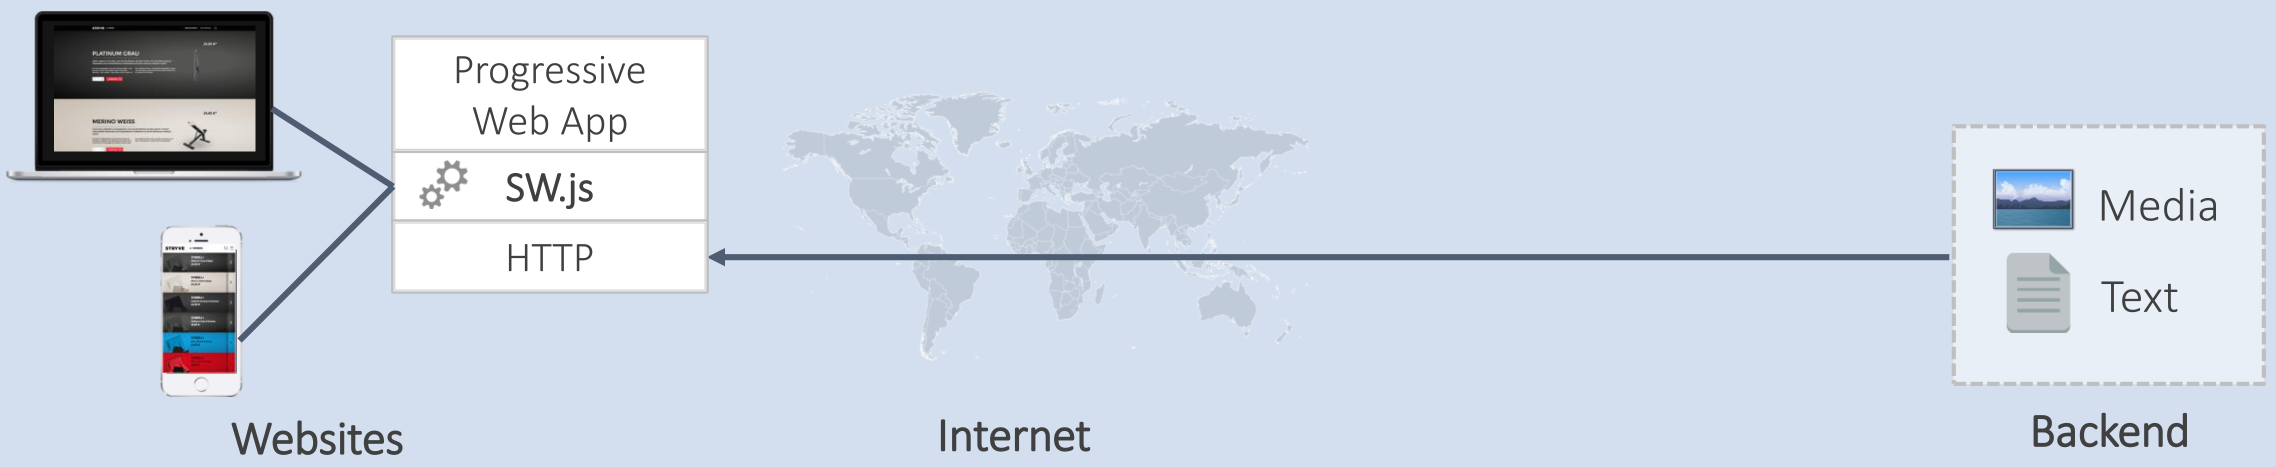

Service Workers + Baqend Caching

Baqend Worker:

– **Redirect** requests to Baqend for faster delivery by including **a snippet**
– **Update** of cached data: refresh of stored data against origin

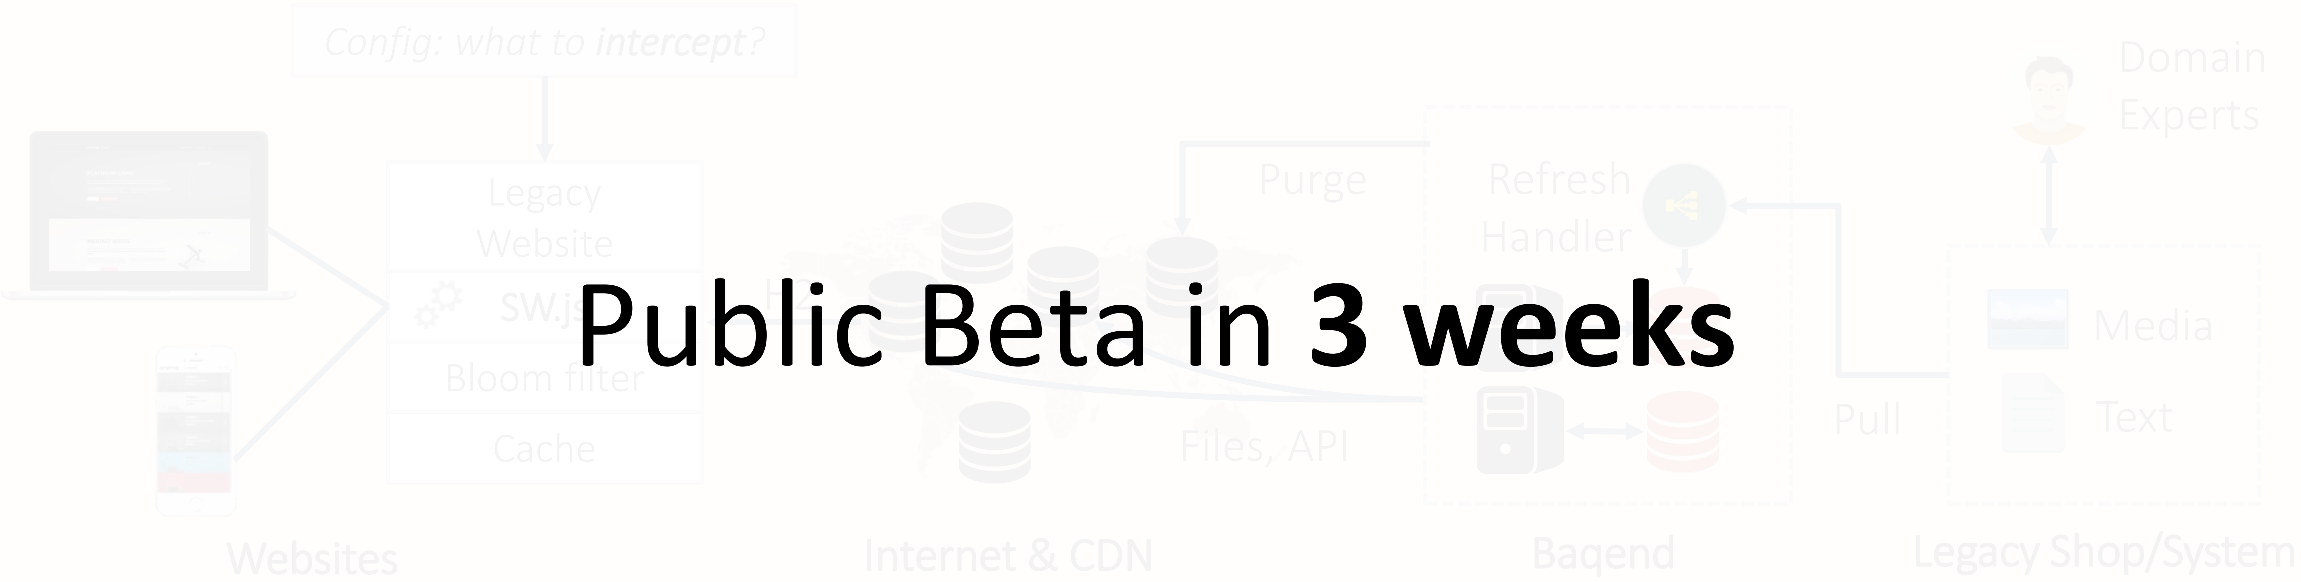

Public Beta in 3 weeks

commerce.codetalks.de
(Rails)

commerce.codetalks.de
Loaded Through
Baqend Worker

Try it: codetalks.baqend.com
(with offline mode)

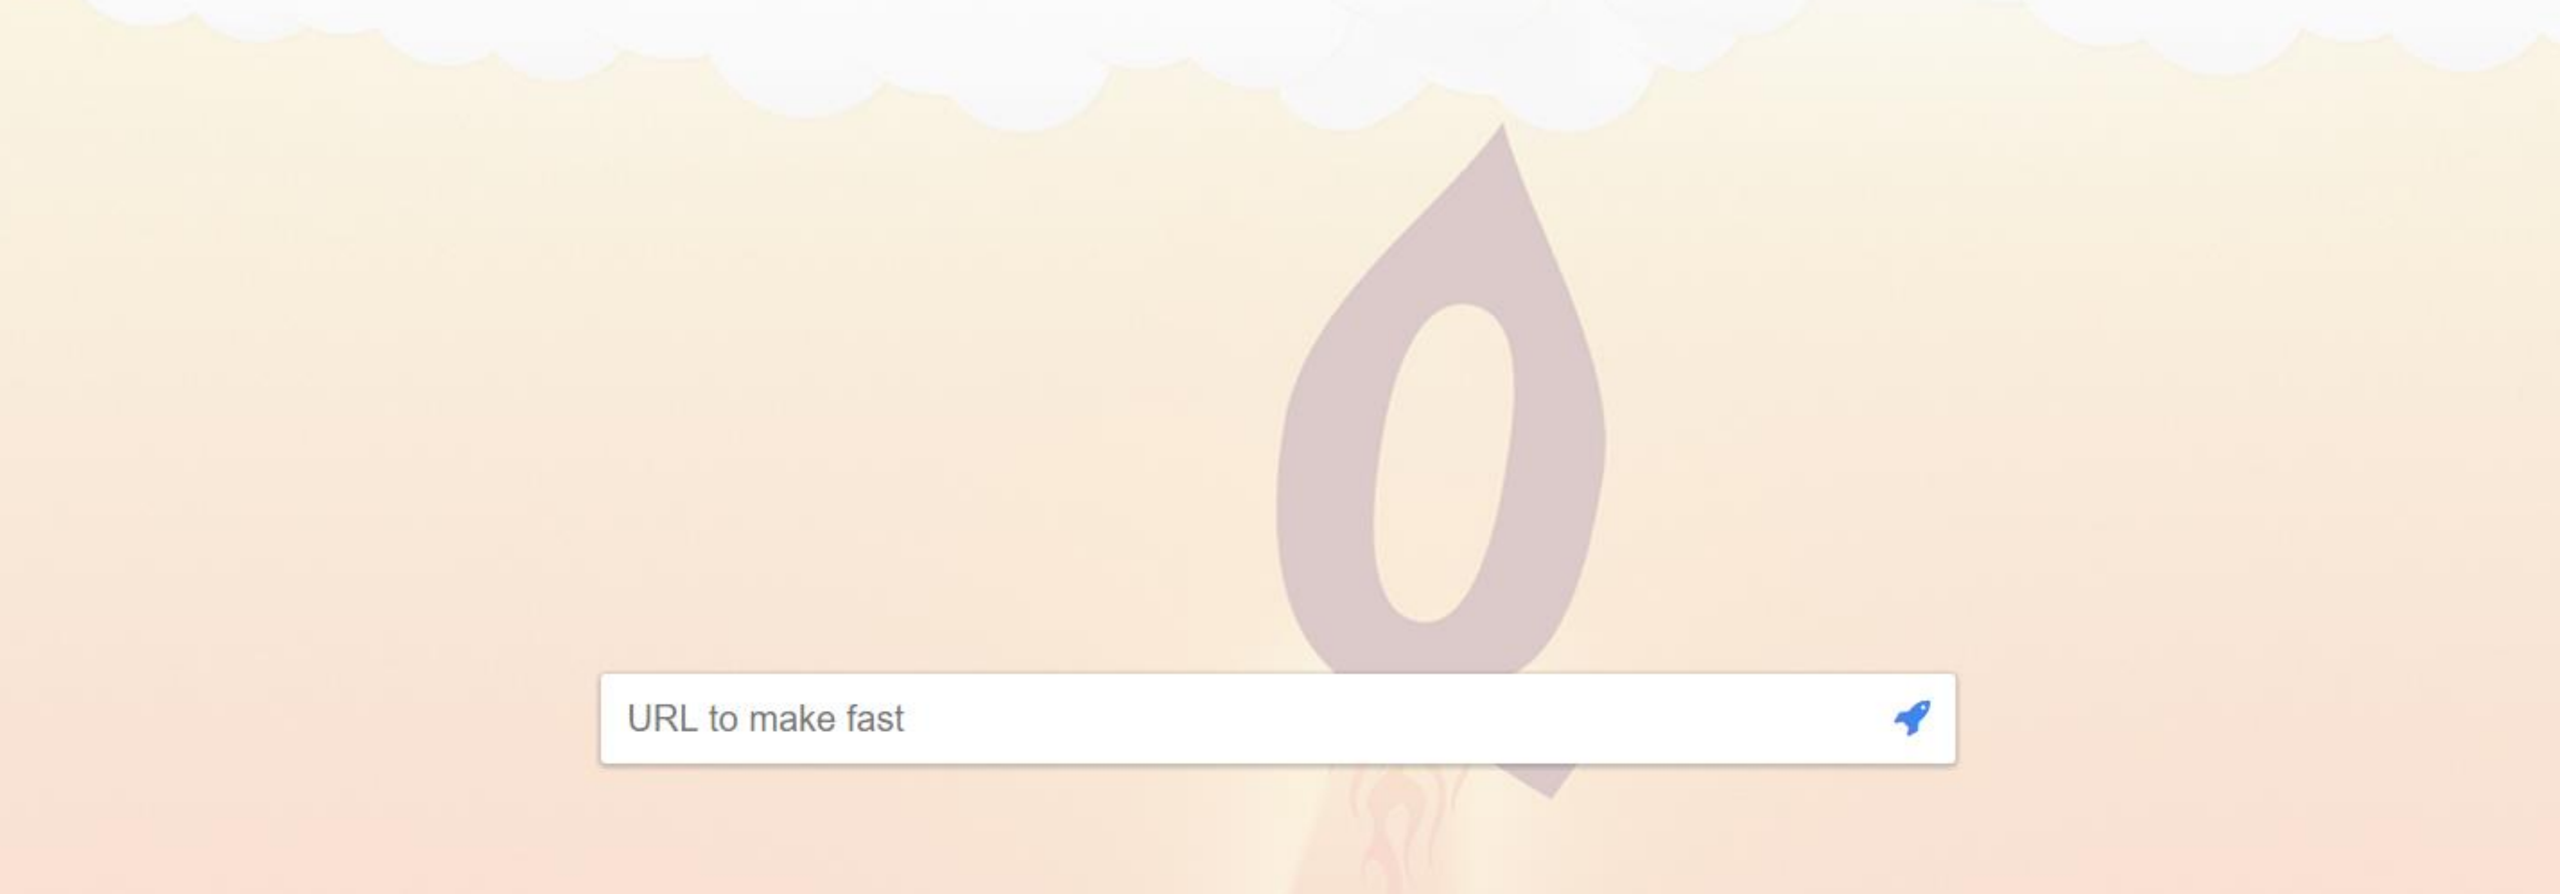

DEMO

In action: makefast.baqend.com

Typical Speedup: 15x

Impact of Global Caching

TRY THIS

benchmark.baqend.com

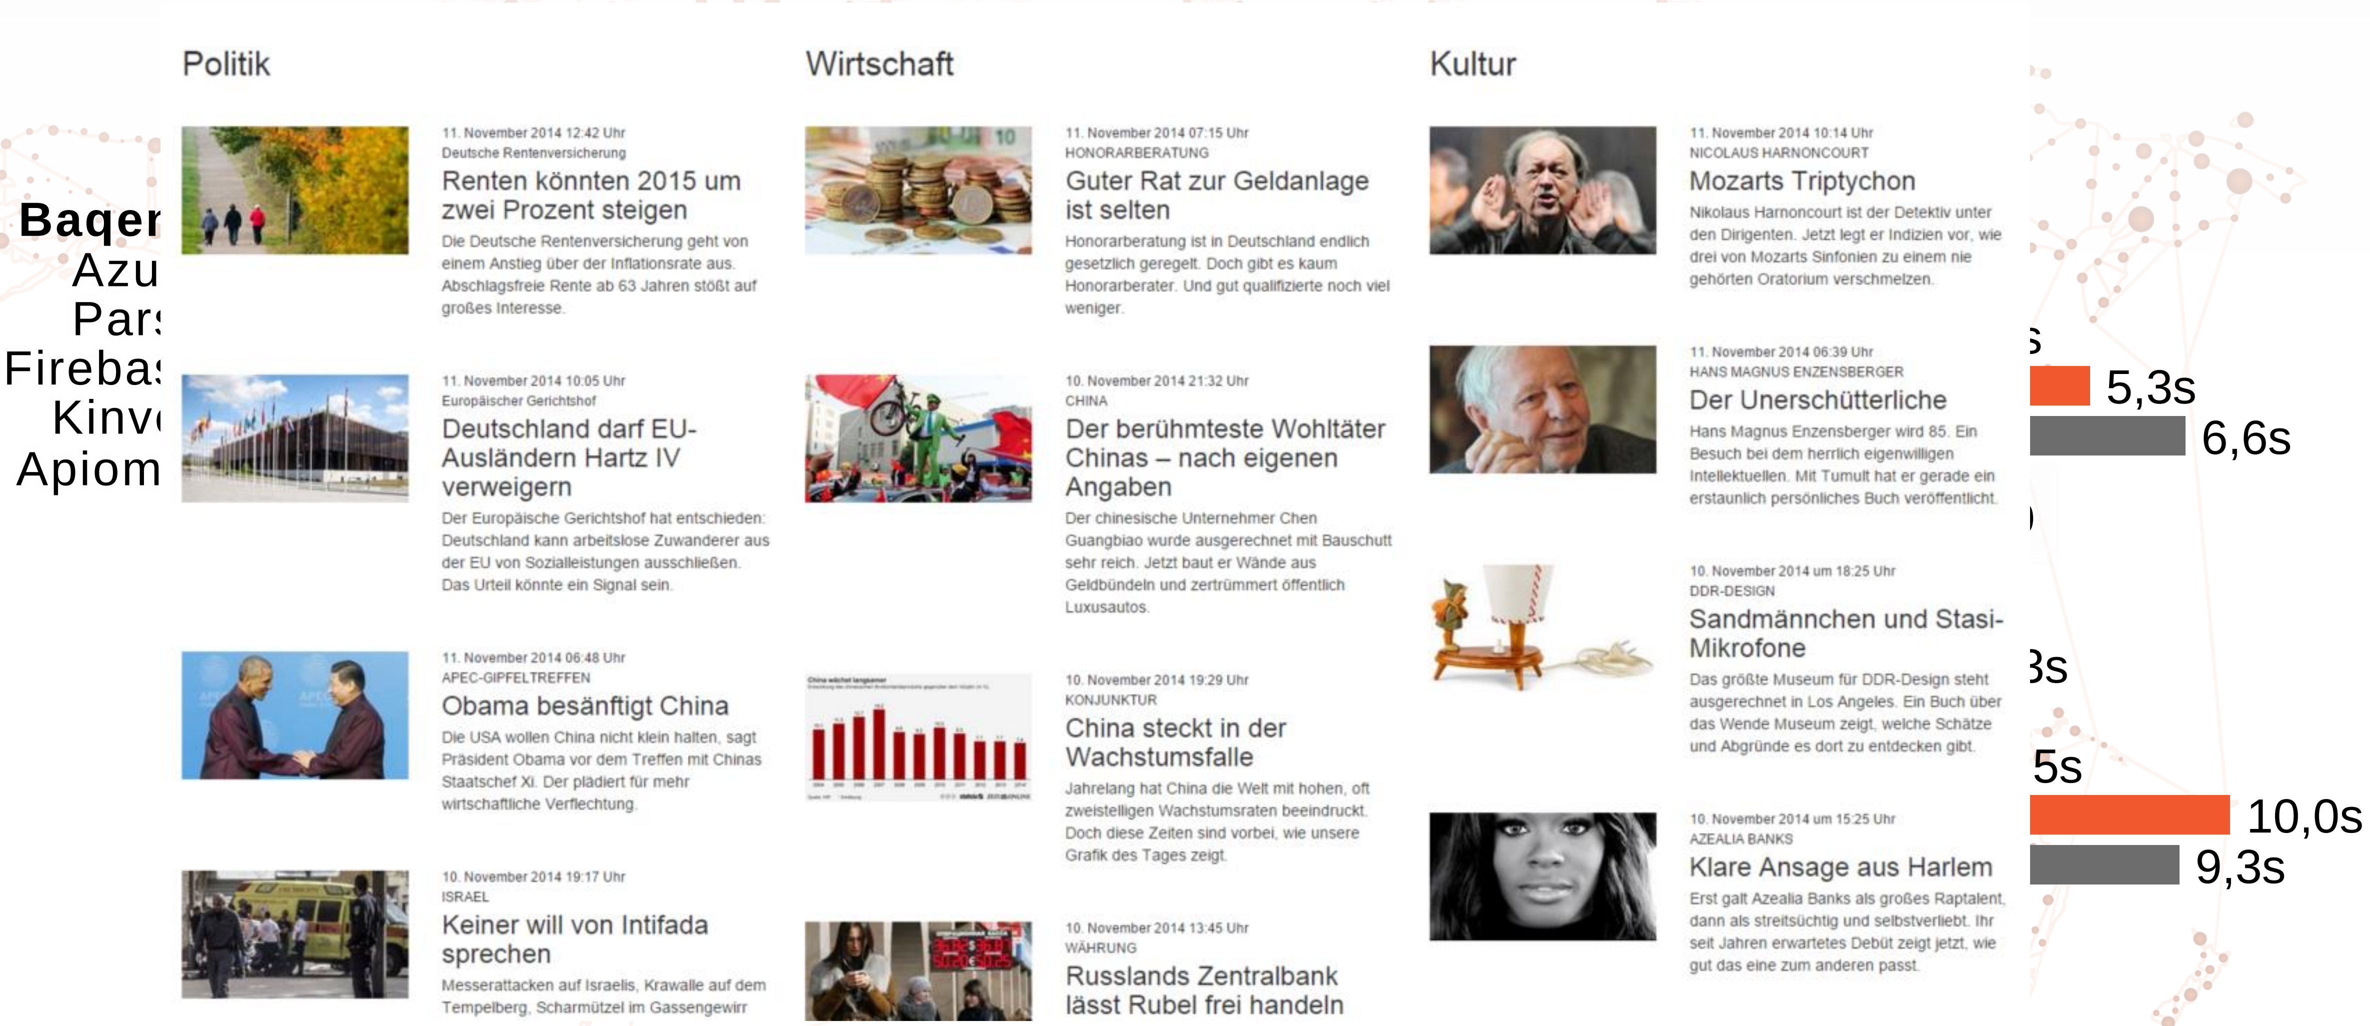

Politik

11. November 2014 12:42 Uhr
Deutsche Rentenversicherung
Renten könnten 2015 um zwei Prozent steigen
Die Deutsche Rentenversicherung geht von einem Anstieg über der Inflationsrate aus. Abschlagsfreie Rente ab 63 Jahren stößt auf großes Interesse.

11. November 2014 10:05 Uhr
Europäischer Gerichtshof
Deutschland darf EU-Ausländern Hartz IV verweigern
Der Europäische Gerichtshof hat entschieden: Deutschland kann arbeitslose Zuwanderer aus der EU von Sozialleistungen ausschließen. Das Urteil könnte ein Signal sein.

11. November 2014 06:48 Uhr
APEC-GIPFELTREFFEN
Obama besänftigt China
Die USA wollen China nicht klein halten, sagt Präsident Obama vor dem Treffen mit Chinas Staatschef Xi. Der plädiert für mehr wirtschaftliche Verflechtung.

10. November 2014 19:17 Uhr
ISRAEL
Keiner will von Intifada sprechen
Messerattacken auf Israelis, Krawalle auf dem Tempelberg, Scharmützel im Gassengewirr

Wirtschaft

11. November 2014 07:15 Uhr
HONORARBERATUNG
Guter Rat zur Geldanlage ist selten
Honorarberatung ist in Deutschland endlich gesetzlich geregelt. Doch gibt es kaum Honorarberater. Und gut qualifizierte noch viel weniger.

10. November 2014 21:32 Uhr
CHINA
Der berühmteste Wohltäter Chinas – nach eigenen Angaben
Der chinesische Unternehmer Chen Guangbiao wurde ausgerechnet mit Bauschutt sehr reich. Jetzt baut er Wände aus Geldbündeln und zertrümmert öffentlich Luxusautos.

10. November 2014 19:29 Uhr
KONJUNKTUR
China steckt in der Wachstumsfalle
Jahrelang hat China die Welt mit hohen, oft zweistelligen Wachstumsraten beeindruckt. Doch diese Zeiten sind vorbei, wie unsere Grafik des Tages zeigt.

10. November 2014 13:45 Uhr
WÄHRUNG
Russlands Zentralbank lässt Rubel frei handeln

Kultur

11. November 2014 10:14 Uhr
NICOLAUS HARNONCOURT
Mozarts Triptychon
Nikolaus Harnoncourt ist der Detektiv unter den Dirigenten. Jetzt legt er Indizien vor, wie drei von Mozarts Sinfonien zu einem nie gehörten Oratorium verschmelzen.

11. November 2014 06:39 Uhr
HANS MAGNUS ENZENSBERGER
Der Unerschütterliche
Hans Magnus Enzensberger wird 85. Ein Besuch bei dem herrlich eigenwilligen Intellektuellen. Mit Tumult hat er gerade ein erstaunlich persönliches Buch veröffentlicht.

10. November 2014 um 18:25 Uhr
DDR-DESIGN
Sandmännchen und Stasi-Mikrofone
Das größte Museum für DDR-Design steht ausgerechnet in Los Angeles. Ein Buch über das Wende Museum zeigt, welche Schätze und Abgründe es dort zu entdecken gibt.

10. November 2014 um 15:25 Uhr
AZEALIA BANKS
Klare Ansage aus Harlem
Erst galt Azealia Banks als großes Raptalent, dann als streitsüchtig und selbstverliebt. Ihr seit Jahren erwartetes Debüt zeigt jetzt, wie gut das eine zum anderen passt.

THINKS USE CASE

How do these techniques work for a high-traffic shop?

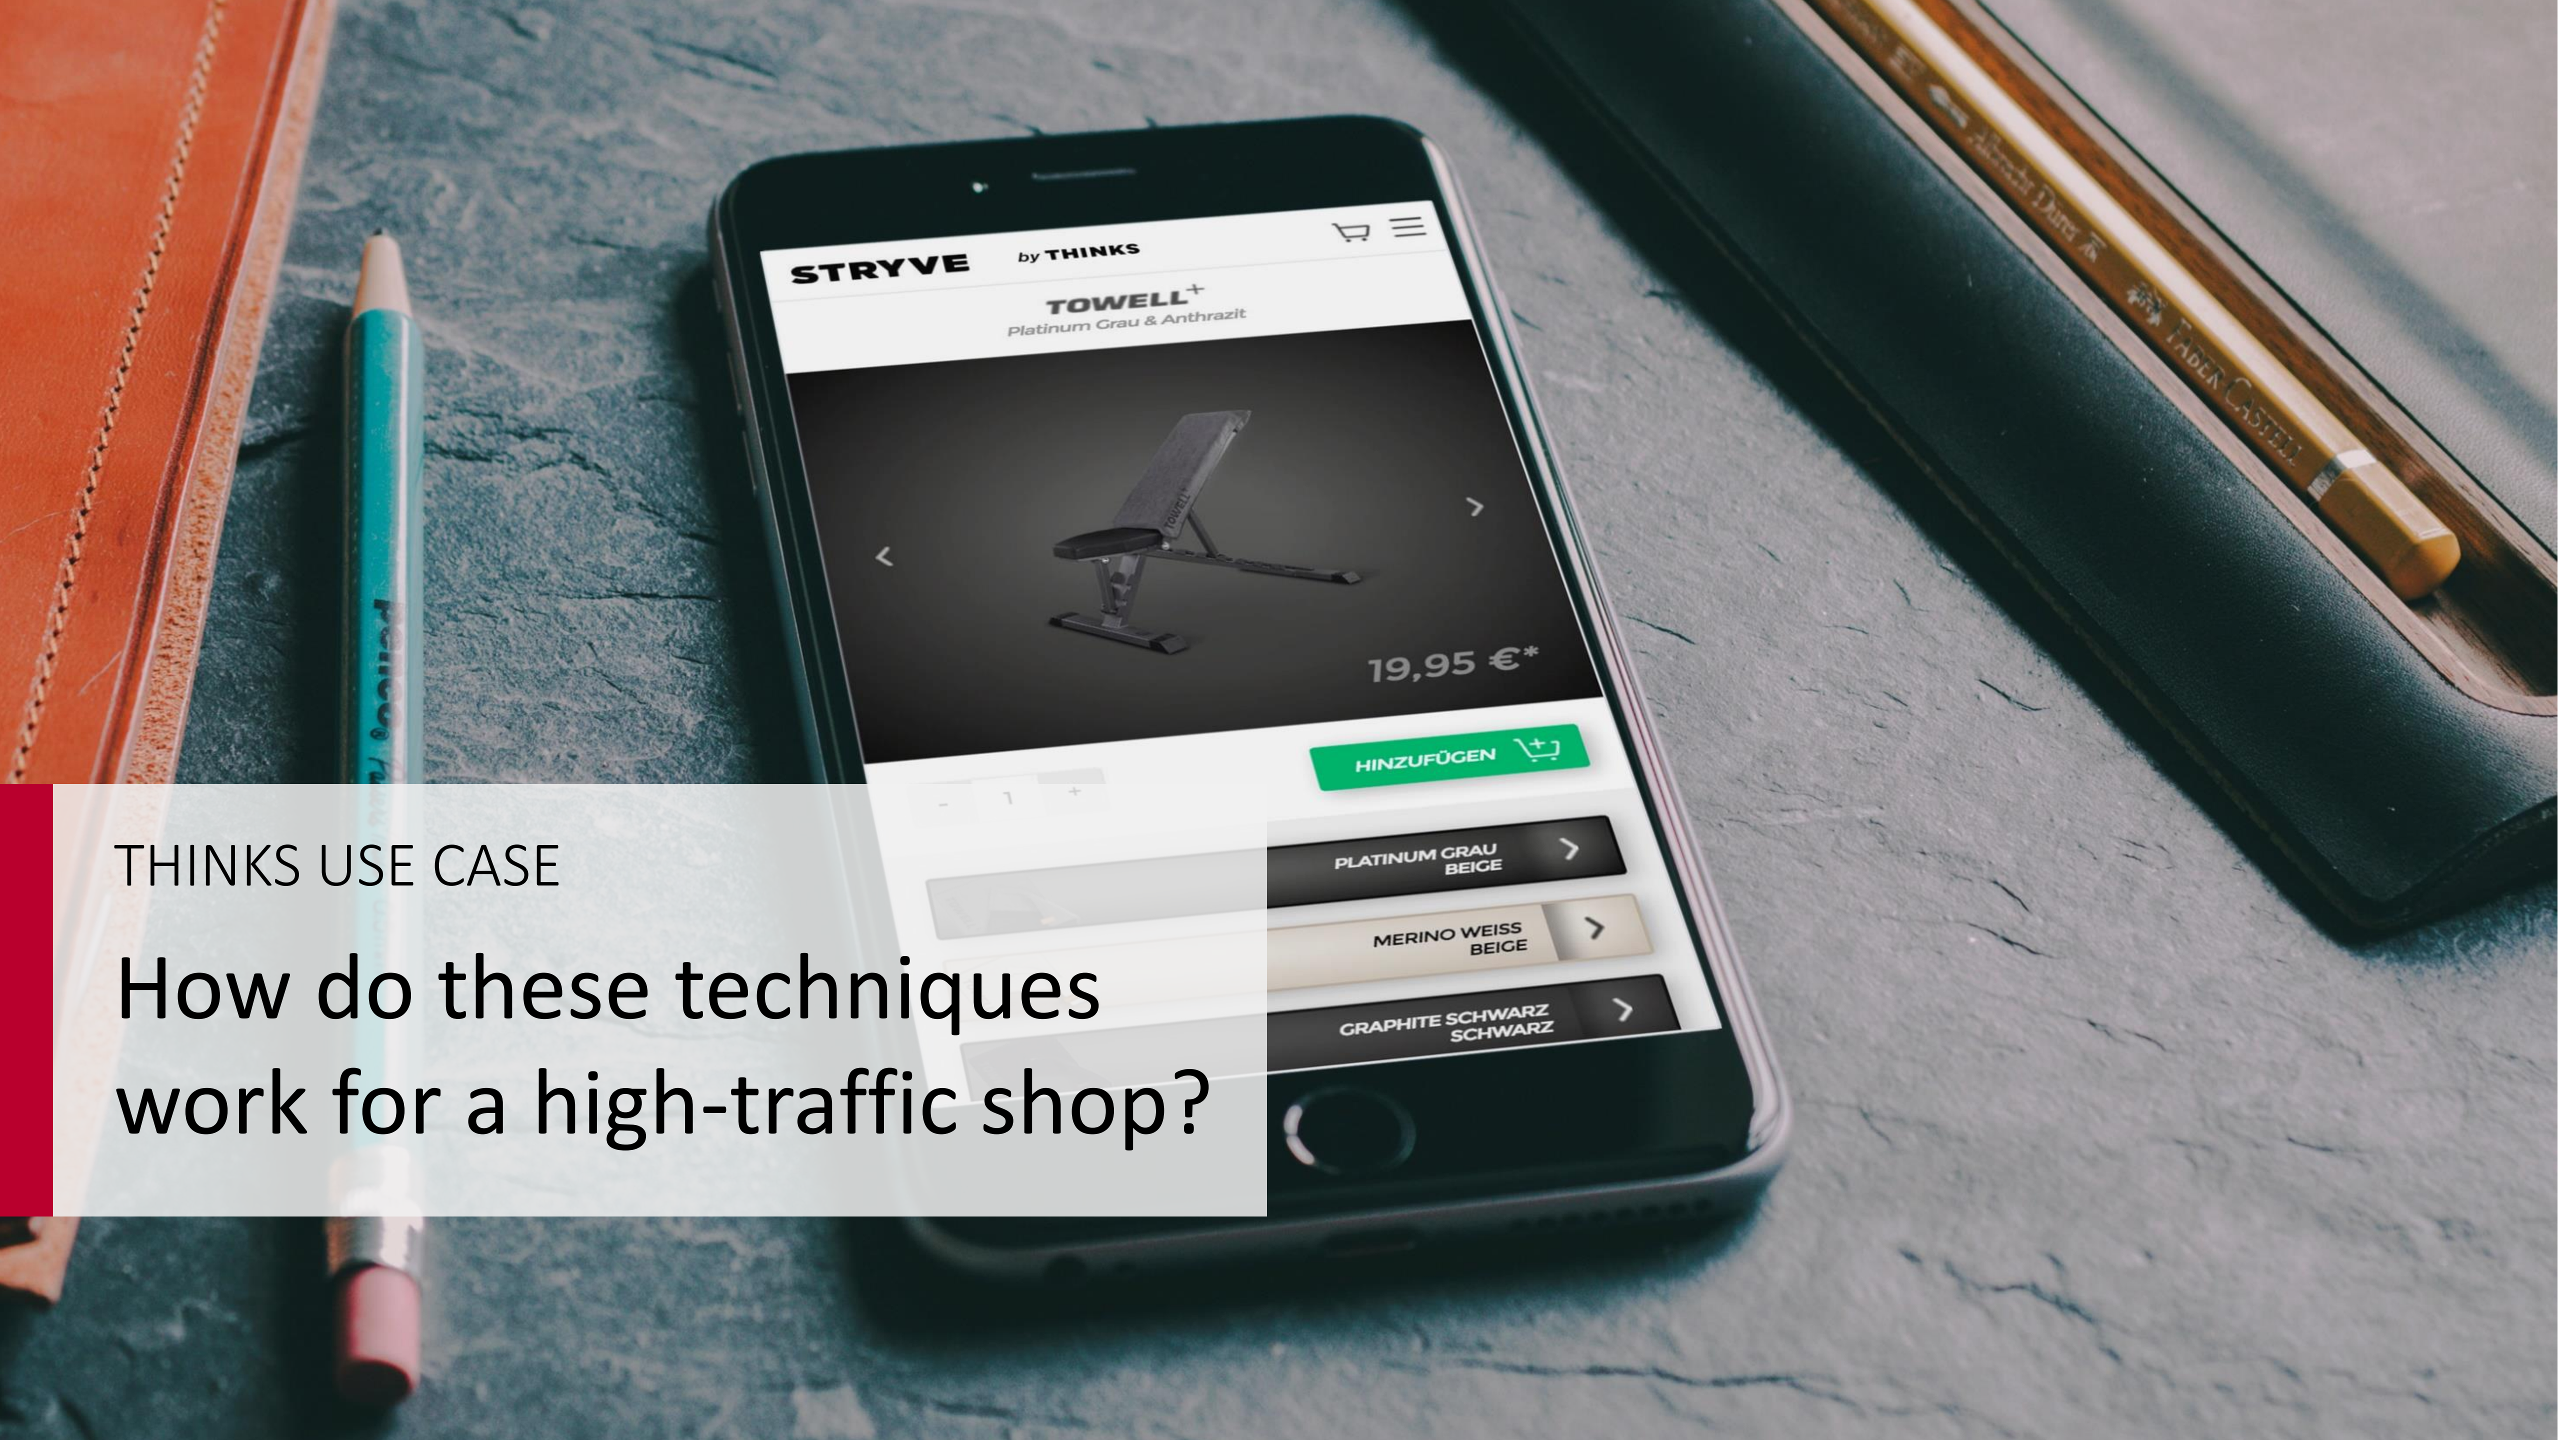

BaQend

STRYVE by THINKS Shop

< 1 second
Page Load Time

7.8%
Conversion Rate

Simultaneous Users

3%
Server Usage

Shops in "Die Höhle der Löwen"

The Google Page Speed Scores
for Season 3, 09/06/2016

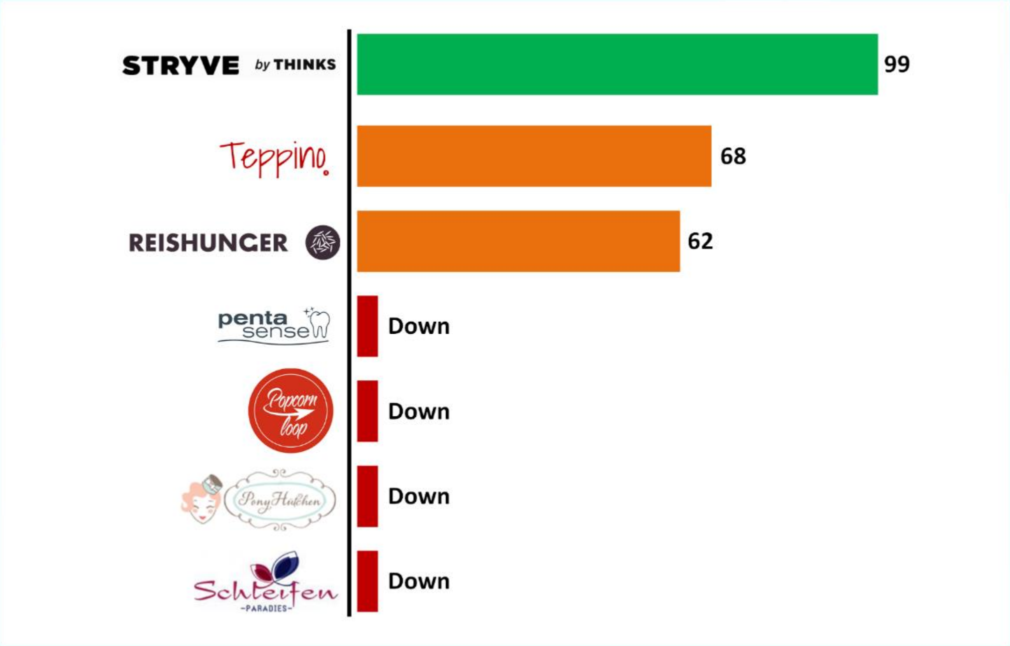

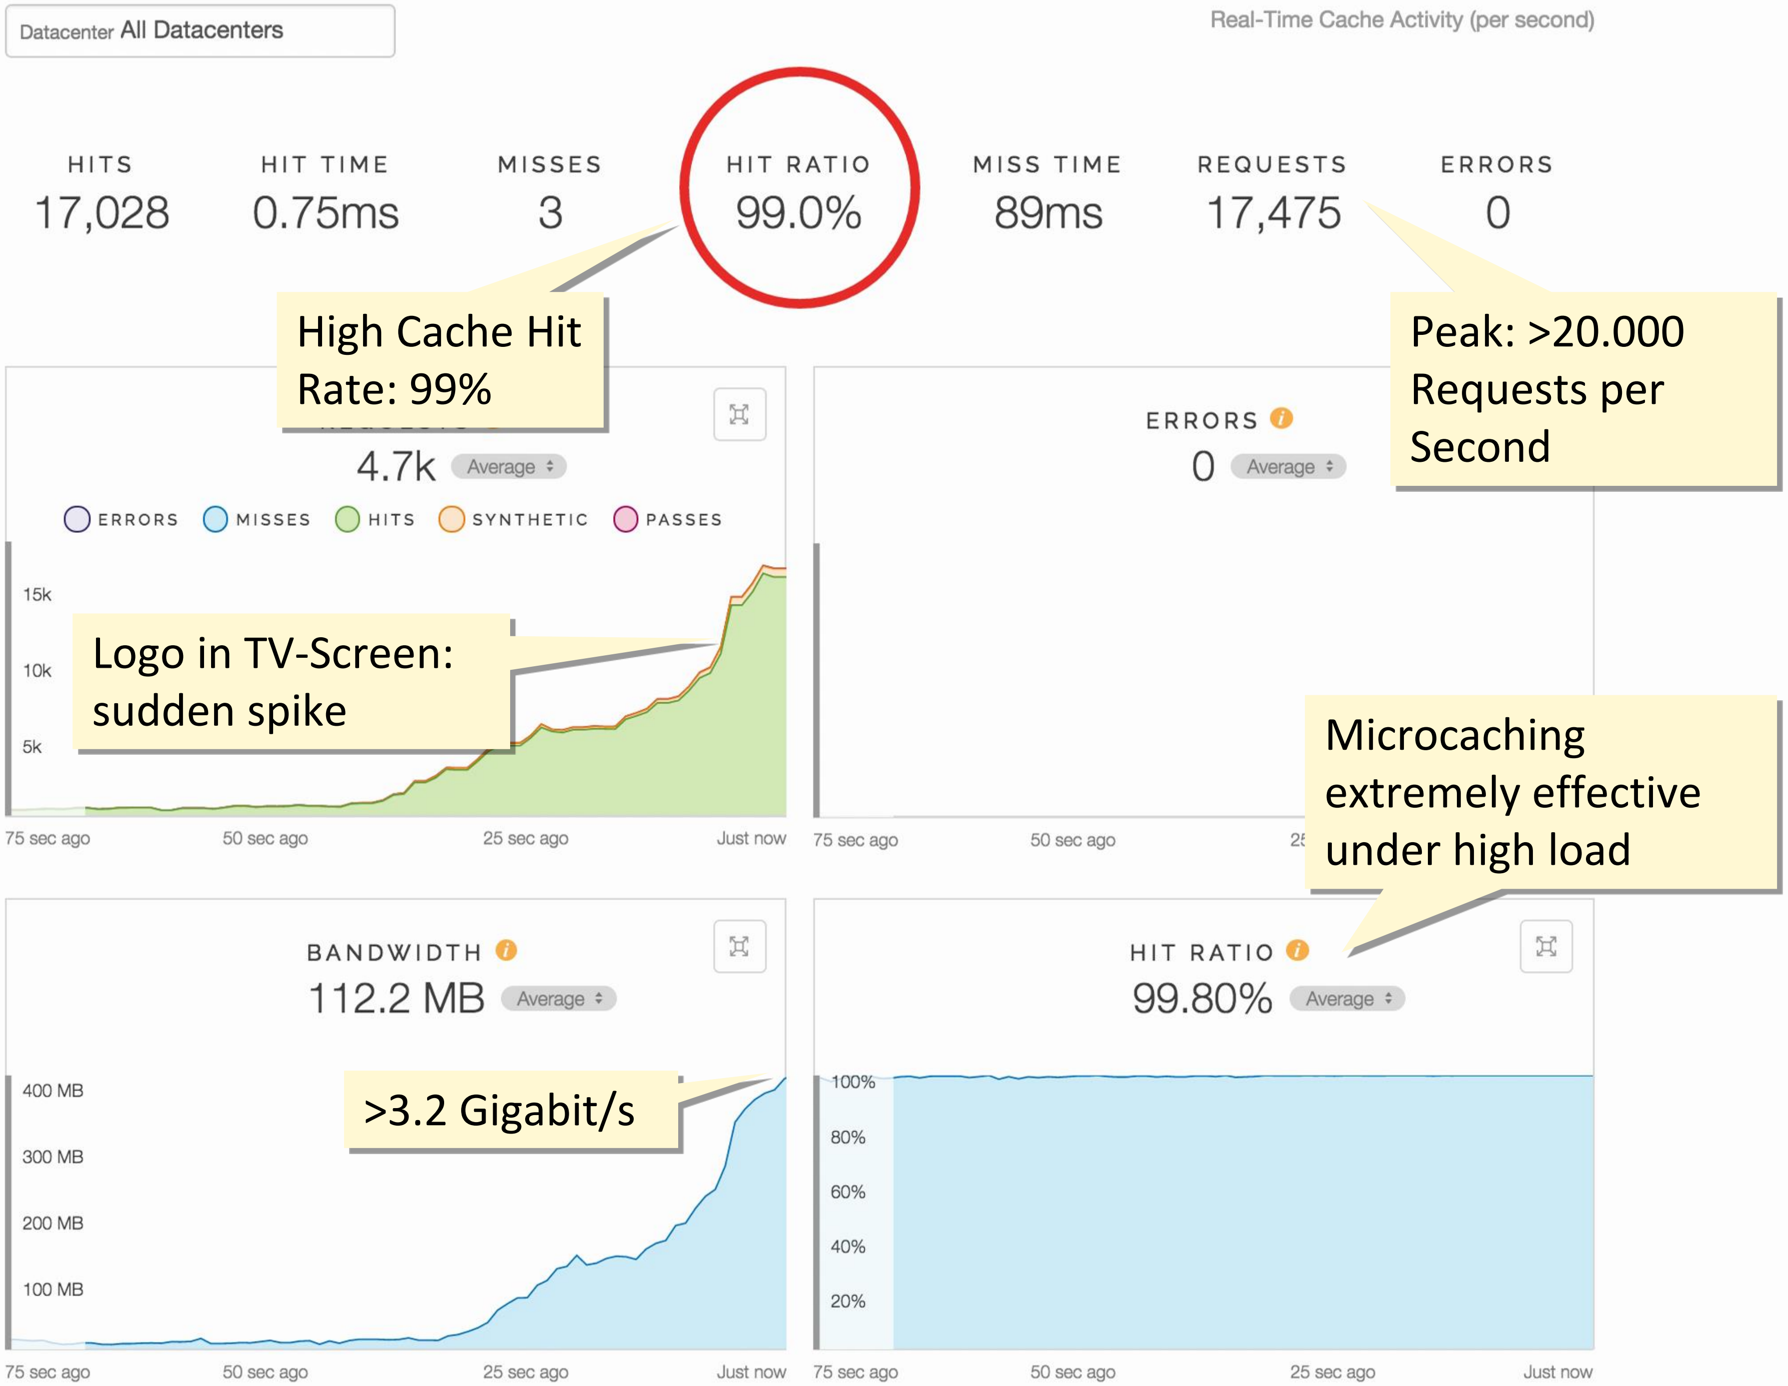
Datacenter All Datacenters
Real-Time Cache Activity (per second)
HITS 17,028
HIT TIME 0.75ms
MISSES 3
HIT RATIO 99.0%
MISS TIME 89ms
REQUESTS 17,475
ERRORS 0
High Cache Hit Rate: 99%
Peak: >20.000 Requests per Second
4.7k Average
ERRORS MISSES HITS SYNTHETIC PASSES
15k
10k
5k
Logo in TV-Screen: sudden spike
75 sec ago
50 sec ago
25 sec ago
Just now
ERRORS
0 Average
Microcaching extremely effective under high load
BANDWIDTH
112.2 MB Average
400 MB
300 MB
200 MB
100 MB
>3.2 Gigabit/s
HIT RATIO
99.80% Average
100%
80%
60%
40%
20%

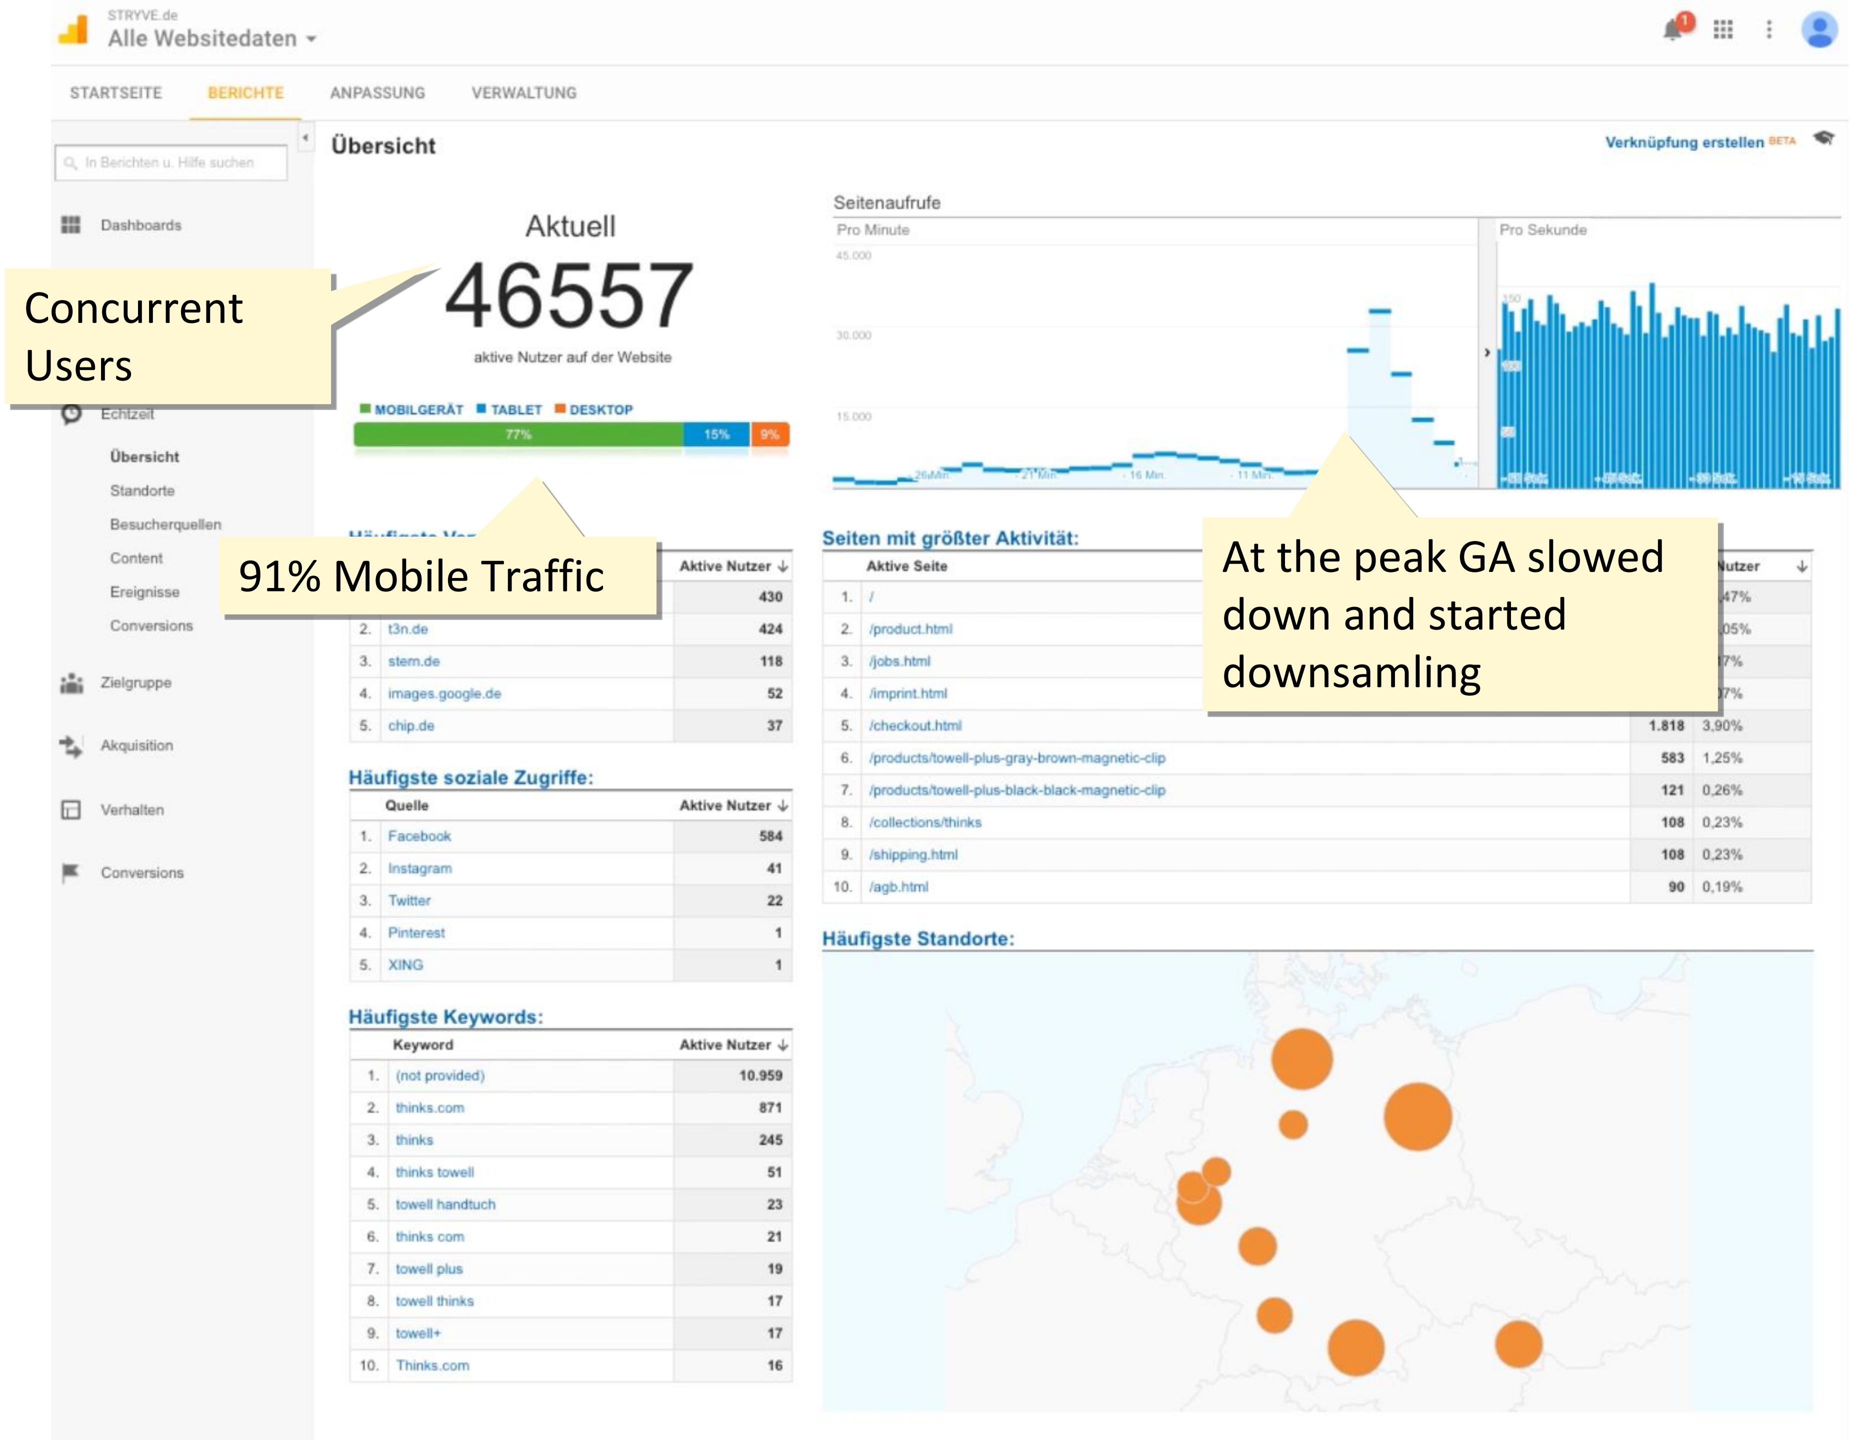
STRYVE.de
Alle Websitedaten
STARTSEITE
BERICHTE
ANPASSUNG
VERWALTUNG
Übersicht
Verknüpfung erstellen BETA
In Berichten u. Hilfe suchen
Dashboards
Concurrent Users
Echtzeit
Übersicht
Standorte
Besucherquellen
Content
Ereignisse
Conversions
Zielgruppe
Akquisition
Verhalten
Conversions
Aktuell
46557
aktive Nutzer auf der Website
MOBILGERÄT
TABLET
DESKTOP
77%
15%
9%
91% Mobile Traffic
Aktive Nutzer
430
2. t3n.de 424
3. stern.de 118
4. images.google.de 52
5. chip.de 37
Häufigste soziale Zugriffe:
Quelle
Aktive Nutzer
1. Facebook 584
2. Instagram 41
3. Twitter 22
4. Pinterest 1
5. XING 1
Häufigste Keywords:
Keyword
Aktive Nutzer
1. (not provided) 10.959
2. thinks.com 871
3. thinks 245
4. thinks towell 51
5. towell handtuch 23
6. thinks com 21
7. towell plus 19
8. towell thinks 17
9. towell+ 17
10. Thinks.com 16
Seitenaufrufe
Pro Minute
Pro Sekunde
45.000
30.000
15.000
Seiten mit größter Aktivität:
Aktive Seite
1. /
2. /product.html
3. /jobs.html
4. /imprint.html
5. /checkout.html 1.818 3,90%
6. /products/towell-plus-gray-brown-magnetic-clip 583 1,25%
7. /products/towell-plus-black-black-magnetic-clip 121 0,26%
8. /collections/thinks 108 0,23%
9. /shipping.html 108 0,23%
10. /agb.html 90 0,19%
At the peak GA slowed down and started downsamling
Häufigste Standorte:

Lessons Learned

Frontend

- **Simple** frontend means good performance
- **Single-page application**
- **Tooling** for optimizations (Inlining, Above-the-fold, Minification, etc.)

Network

- **Caching** in the CDN and browser
- **Dynamic data** should also be cached
- **Minimize latency**, SSL and HTTP tuning

Backend

- Horizontal **scaling** with stateless web servers
- **NoSQL** databases
- **Cloud**-hosted
- **Load-Tests** important
- **Failover** and **autoscaling**

fg@baqend.com @baqendcom

READ MORE

Written version:

blog.baqend.com

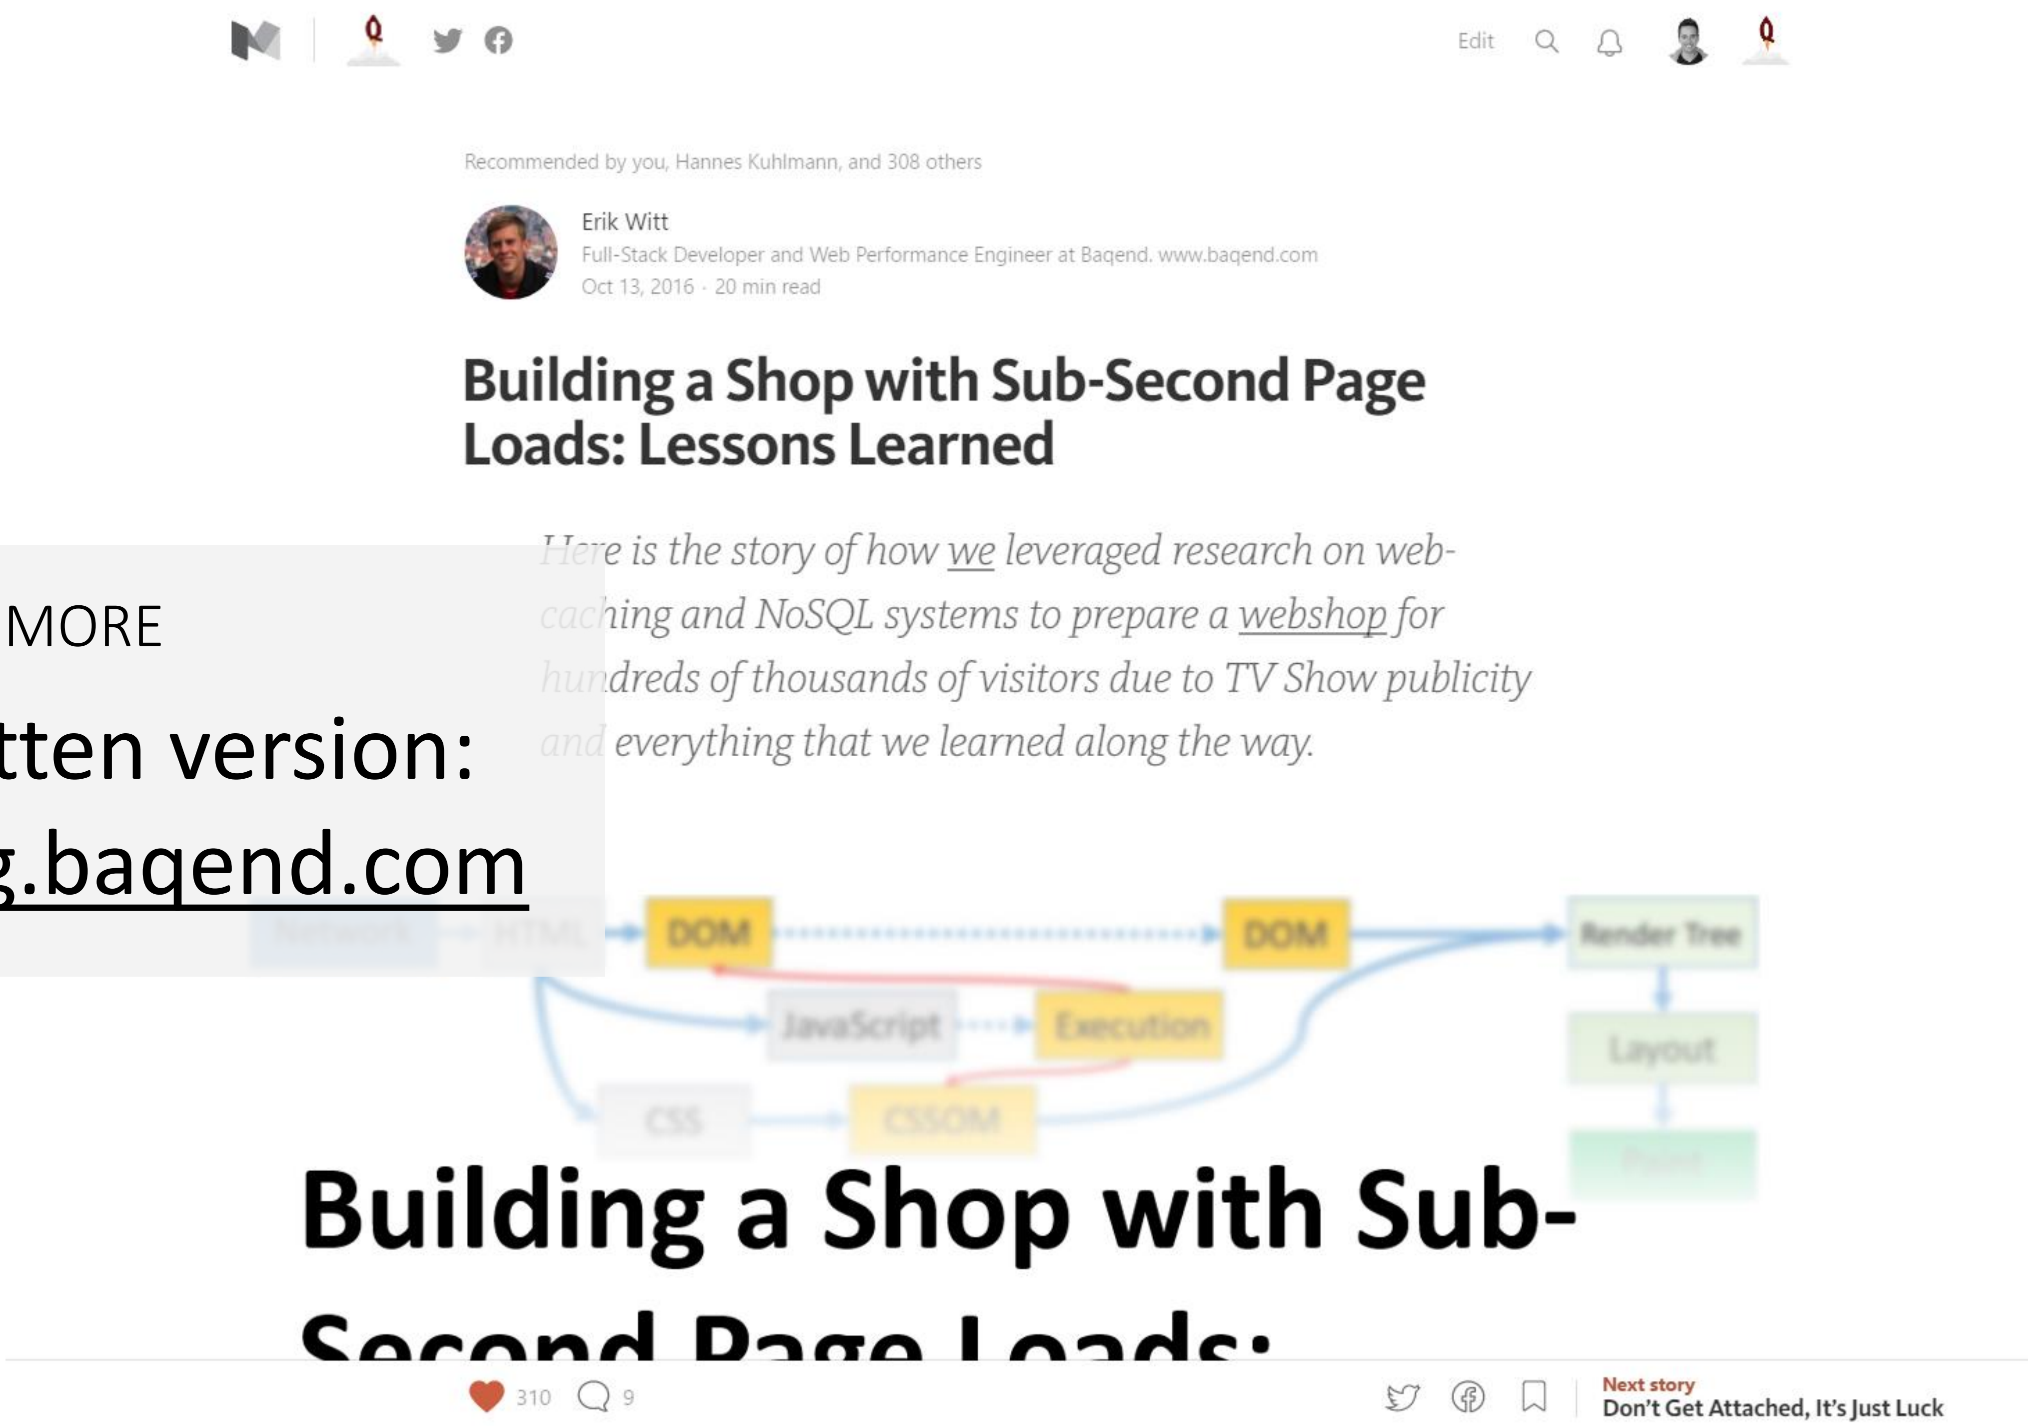
Edit

Recommended by you, Hannes Kuhlmann, and 308 others

Erik Witt
Full-Stack Developer and Web Performance Engineer at Baqend. www.baqend.com
Oct 13, 2016 · 20 min read

Building a Shop with Sub-Second Page Loads: Lessons Learned

Here is the story of how we leveraged research on web-caching and NoSQL systems to prepare a webshop for hundreds of thousands of visitors due to TV Show publicity and everything that we learned along the way.

Building a Shop with Sub-Second Page Loads:

310 9

Next story
Don't Get Attached, It's Just Luck

Recommendations

Literature and Tools

Good Ressources:

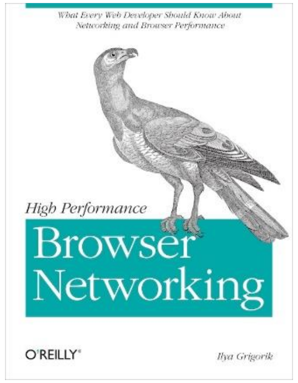

https://hpbn.co/

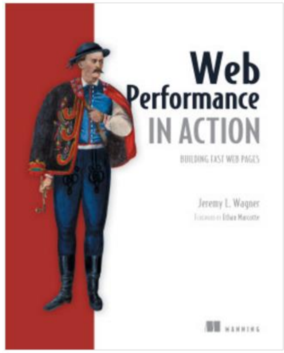

https://www.manning.com/books/web-performance-in-action

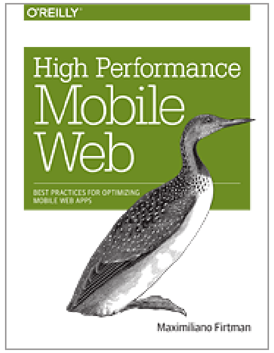

http://shop.oreilly.com/product/0636920035060.do

Google Developers Performance

https://developers.google.com/web/fundamentals/performance/?hl=en

Website Performance Optimization

The Critical Rendering Path

https://www.udacity.com/course/website-performance-optimization--ud884

Baqend Blog

On Building a Faster Web

https://medium.baqend.com/

Good Tools:

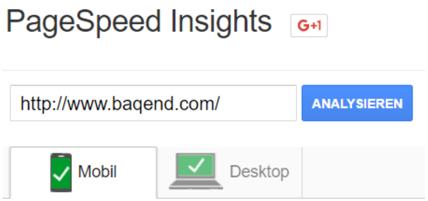

https://developers.google.com/speed/pagespeed/

GTmetrix

https://gtmetrix.com

Baqend

https://www.baqend.com/

WEBPAGETEST

http://www.webpagetest.org/