# **DOCEJS** Performance Case Study

@Fabian\_Frank
<fabian@pagefault.de>

Yahoo! Search, Engineer Youthmedia.eu, Volunteer



# A Dynamic Website



## self-contained App



## self-contained App

YAHOO!	beijing Search		🥥 Options 🗸	
	beijing	BEIJING 2008 OLYMPIC GAMES		
	beijing olympic 2008	TOP SITES		
FILTER BY TIME	beijing olympics	The Official Website of the Beijing 2008 Olympic		
Anytime	beijing weather	http://en.beijing2008.cn/		Sponsored Results
Past day	beijing china	2008 Summer Olympics - Wikipedia, the free W http://en.wikipedia.org/wiki/2008_Summer_Olympics		<u>∋rs</u> Products
Past week	beijing 2008 olympic games			
Past month	beijing airport	Beijing 2008 - Olympic Games		cturers Directory.
TRENDING SEARCHES	beijing time			
vale beijing	beijing hotels			
beijing capital airport	beijing map			
dulwich college beijing			beiling china	
hourly beijing			Jeijing onna	
air pollution in <b>beijing</b>			See your m	essage here
RELATED BOINTS OF	Beijing - Wikipedia, the free er	ncyclopedia		

RELATED POINTS OF

## self-contained App

### YAHOO! beiiina 0 Search Options beijing WEATHER BEIJING, CHINA beijing olympic 2008 FILTER BY TIME beijing olympics Anytime beijing weather Sponsored Results Past day beijing china ers TODAY TOMORROW THURSDAY Past week beijing 2008 olympic games Products. Few Snow Showers AM Snow Showers Sunny Past month cturers Directory. beijing airport 40°F 29°F 39°F 15°F 32°F 16°F m beijing time MORE WEATHER TRENDING SEARCHES beijing hotels Yahoo! Extended Forecast vale beijing beijing map beijing capital airport VAHOO! WEATHER program was starte, strate and write a strategy dulwich college beijing beijing china hourly beijing air pollution in beijing See your message here... Beijing - Wikipedia, the free encyclopedia

RELATED POINTS OF

### node v0.4.X

multi-core

http://developer.yahoo.com/blogs/ydn/posts/2010/07/multicore\_http\_server\_with\_nodejs/

# Time Sharing





### Event Queue

• does only one thing at a time

 events that occur are queued for processing

 after an event was processed the next one is fetched from the queue

## The difference

### <u>Workers</u>

- I. synchronous call
- 2. worker blocked
- 3. periodically check if worker can go ahead
- 4. call returns
- worker goes ahead on next check

### Event Queue

- I. asynchronous call that specifies event(s) that can occur (callback)
- 2. process next event
- 3. specified event occurs and is put on the queue
- 4. process goes ahead

## Another difference

### <u>Workers</u>

- I connection per worker
- N workers per CPU

### Event Queue

- N connections per process
- I process (per CPU)

# Regarding...

static file serving benchmarks

- are not relevant, unless you plan to do heavy duty static file serving (l'm not)
- hello world benchmarks
  - still ignore most of what matters, e.g. accessing a database or other back-ends

### Realistie

### Real World Benchmark

- a lot of variables that have to be controlled
- more complex, likely to surface weird bugs
- face unanswerable questions
- are expensive to do
- are easy to attack, but hard to defend

# My Reasoning

 "I can not compare everything out there against everything else. But I can compare what I want to use in the future against what I am using today."

# Search Case Study

• Right Panel retrieved using AJAX

• { 'html': '<div>...</div>', 'css': '...', 'js:' '...', ... }

beijing				
beijing beijing olympic 2008	BEIJING, CHINA 11:31 PM (CST) Current local time			
beijing olympics	TOP RATED THINGS TO DO	and the second second		
beijing <b>weather</b> beijing <b>china</b>	<ol> <li>Great Wall</li> <li>Yonghe Gong (Lama Temple)</li> <li>Forbidden City</li> </ol>	- Cones		
beijing 2008 olympic games	BEIJING OVERVIEW			
beijing airport	Hotels	Bei		
beijing hotels	Restaurant Guide			
beijing time	Flights	CEPTER		
beijing map				

# Refresh Right Panel

- I. receive a request
- 2. call a JSON API over HTTP
- 3. manipulate the data structure
- 4. render it as HTML using Mustache
- 5. write back wrapped in JSON



### Constraints

Network bandwidth and latency
 Gigabit and Ims

JSON API performance
 Ioad-balanced cluster with cache

### Implementations



• Node.js (+YUI)

# Apache + PHP

- works, in production for decades
- initial response times are "good enough"
- scales, but you need \$\$\$
- baseline to get a feeling for the numbers

## Apache + PHP

# 430req/s @ 60ms average latency 99% @ 114ms

# Node.js (+YUI)

- works, but very new and untested stack
- initial response times are very low
- scales extremely well
- YUI works well, but still has pitfalls

# Node.js Fun Facts

- >3100 requests per second
  - 100k requests in 32s
  - 9MByte/s network traffic

- serve >2k req/s at 20ms average latency
  - 99% @ 47ms

# Node.js



# Node.js

	Node.js+Mustache+HTTPClient		
<b>Concurrent Connections</b>	Requests served per second	Average response time (ms)	
10	690.93	14.47	
25	1587.43	15.75	
50	2389.20	20.93	
100	2961.82	33.76	
200	3107.91	64.35	

# Node.js Histogram ~1500req/s | 25 conc.



# Node.js Histogram ~2900/s | 100 conc.



# Node.js Histogram ~3100/s | 200 conc.



- works, but extremely new and untested stack
- initial response times are low (60% compared to Apache/PHP)
- scales well (twice as good as Apache/PHP)



	Cocktails		
<b>Concurrent Connections</b>	Requests served per second	Average response time (ms)	
10	422.01	23.70	
25	713.98	35.02	
50	863.30	57.92	
100	848.89	117.80	
200	840.34	238.00	

# Findings

- Both node.js implementations scale linearly before CPU usage hits 100%
- Node.js scales extremely well as proxy
- Manhattan and Mojito can perform better than Apache and PHP (for this use case)
- Node.js applications are very sensitive to memory leaks or complex/blocking code

# DNS in Node.js

- getaddrinfo() sys call is synchronous
  - gethostbyname() is, too

ares\_gethostbyname() is asynchronous
 no cache

node 0.6 uses a getaddrinfo() thread pool, but still relies on ares for other calls

### Cocktails

### • Manhattan



### Manhattan?

- Yahoo!'s Node.js cloud
- can run any node.js application
- allows calls to HTTP APIs (e.g.YQL)
- deployment of versioned apps
- apps can specify their environment

http://developer.yahoo.com/blogs/ydn/posts/2011/11/yahoo-announces-cocktails-%E2%80%93-shaken-not-stirred/

# Mojito?

- Yahoo!'s Node.js MVC framework
- provides identical similar runtime for your code on the client and server
- contains all server- and client-side code
- built on top of YUI
- will be open sourced

http://developer.yahoo.com/blogs/ydn/posts/2011/11/yahoo-announces-cocktails-%E2%80%93-shaken-not-stirred/

# The Cocktails way

### Mojito Application



### Conclusion

- Node.js can serve hundreds of concurrent requests quickly and reliable, "better" than Apache/PHP
- It scales extremely well for I/O bound use cases
- It tears down the client-server language barrier, opening new architectural possibilities

# Thank you!



### Picture: "Esther Müller" / www.jugendfotos.de, CC-License(by-nc)