



High Performance WebSocket

@wesleyhales

SH=PE
shapesecurity.com

This talk will cover...

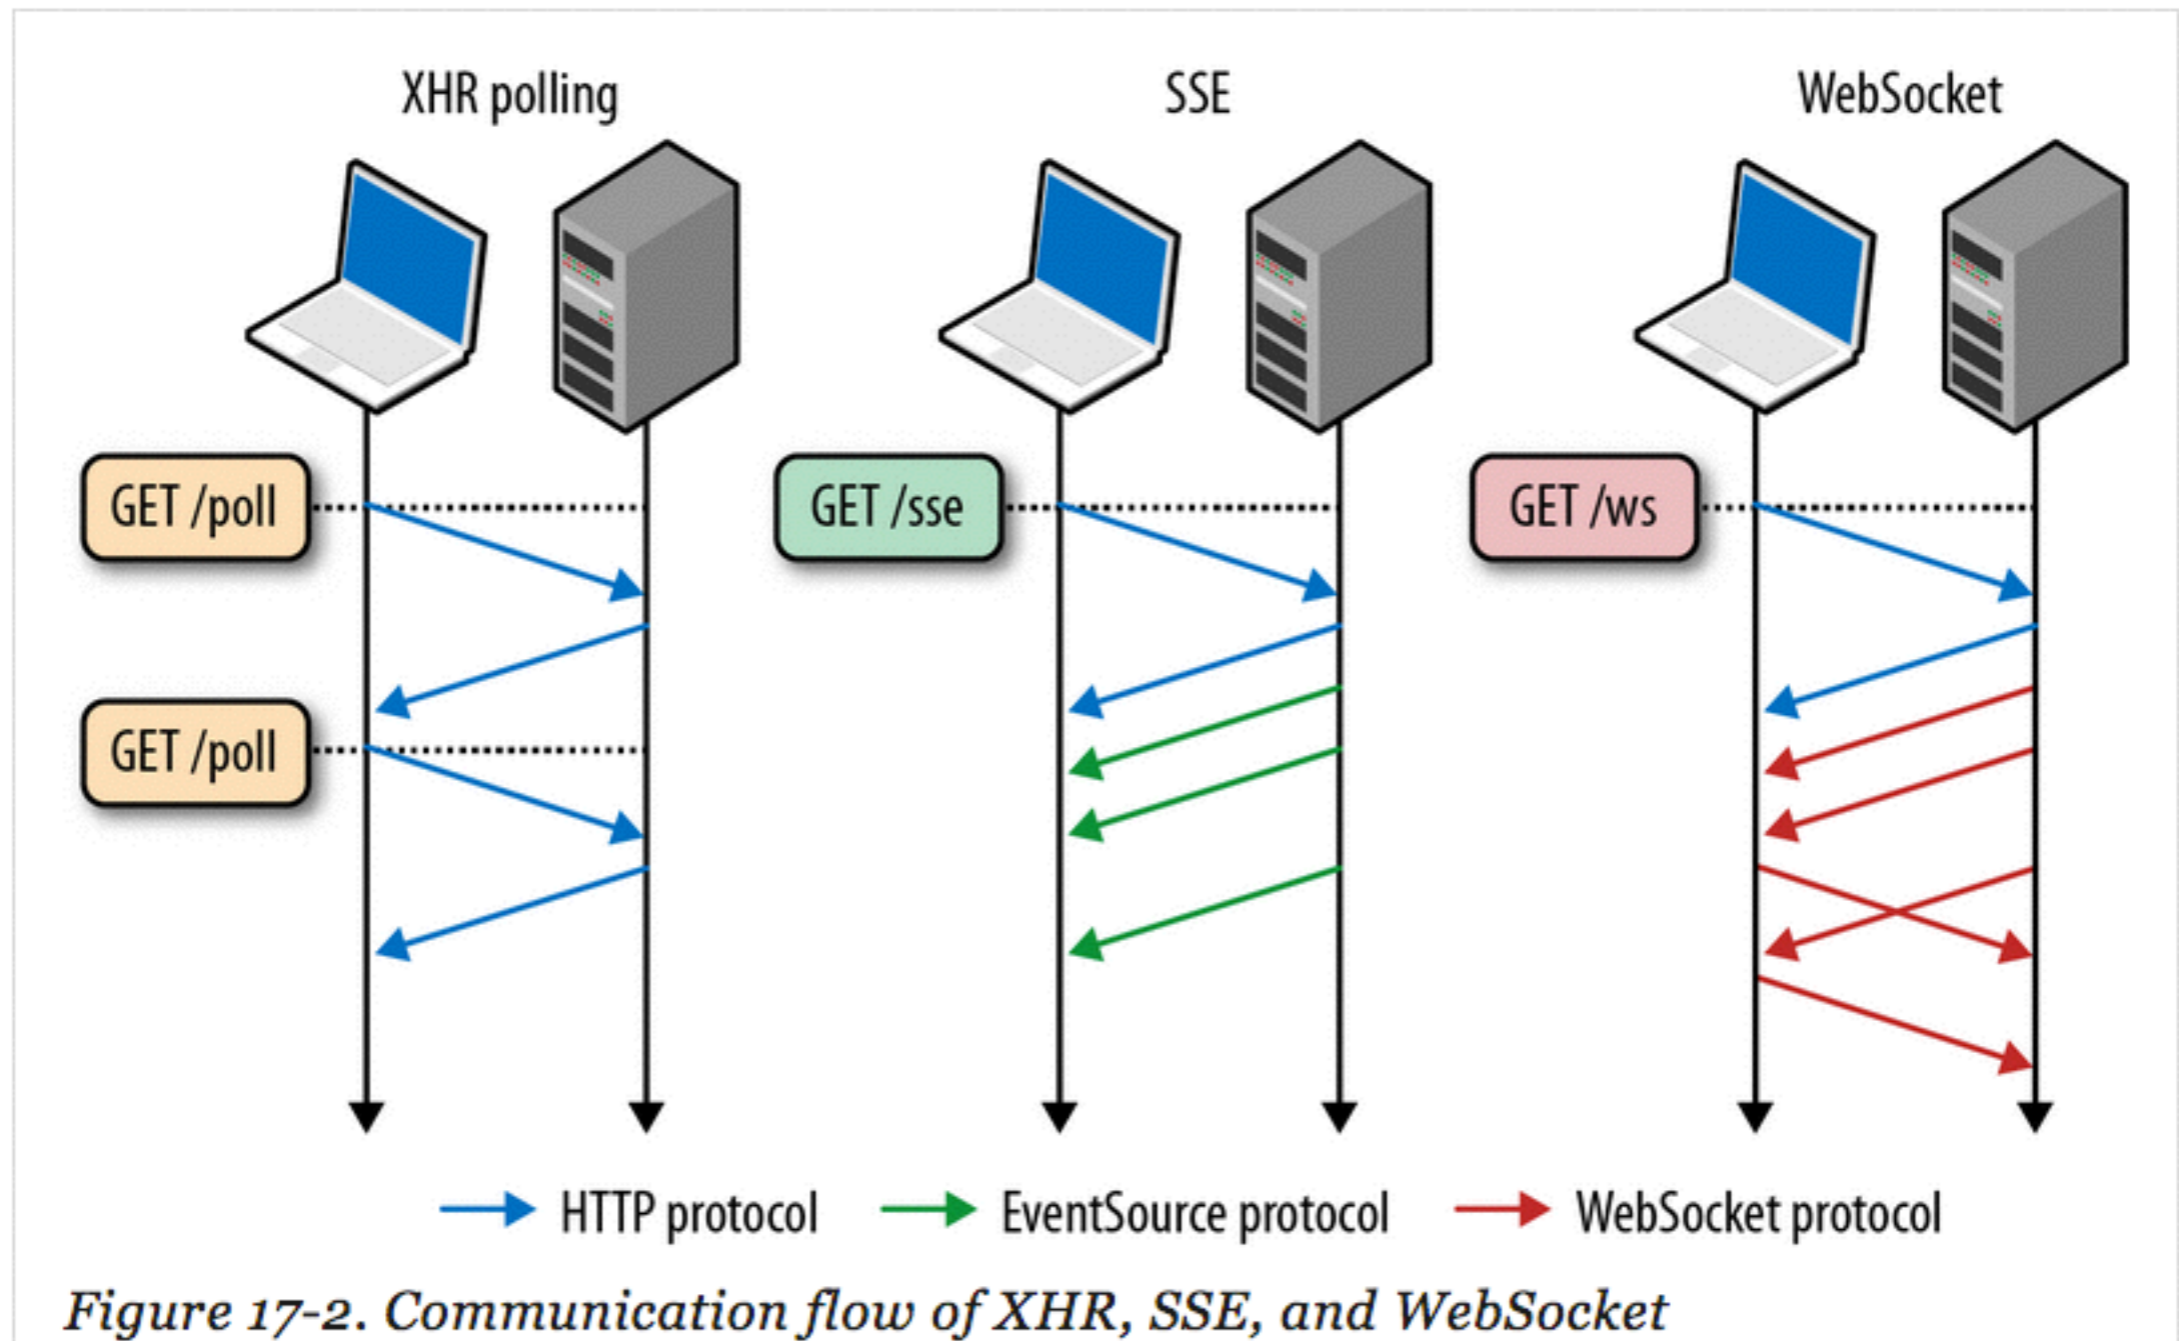
- Light overview of client-server comm.
- Case study of WSS+SPDY project
- Problems faced in collecting real time data
(Browser -> Server)
- Performance Checklist
- Animated GIFs

WebSocket

WebSocket provides low latency delivery of text and binary application data in both directions.

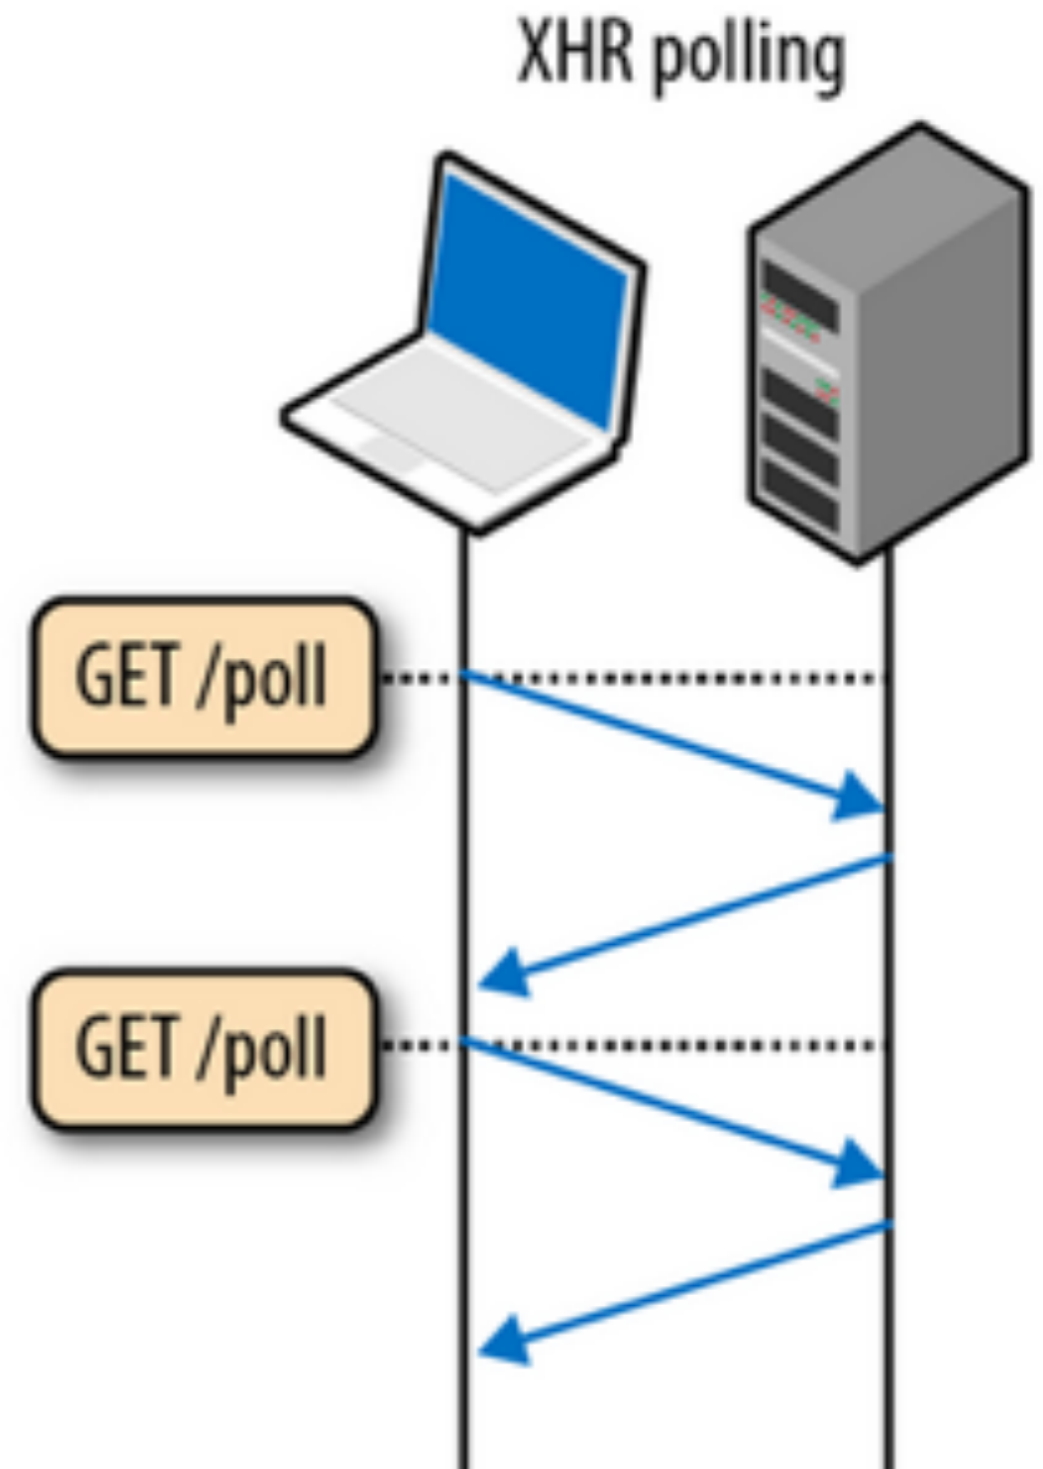
Not a replacement for XHR or SSE

XHR, SSE, and WS




XHR is....

"transactional"
request-response
communication



SSE is....

Efficient, low-latency
server-to-client
streaming of text-
based data



Experiment onslyde



<https://www.youtube.com/watch?v=uuHZ-d0Stsg#t=710>

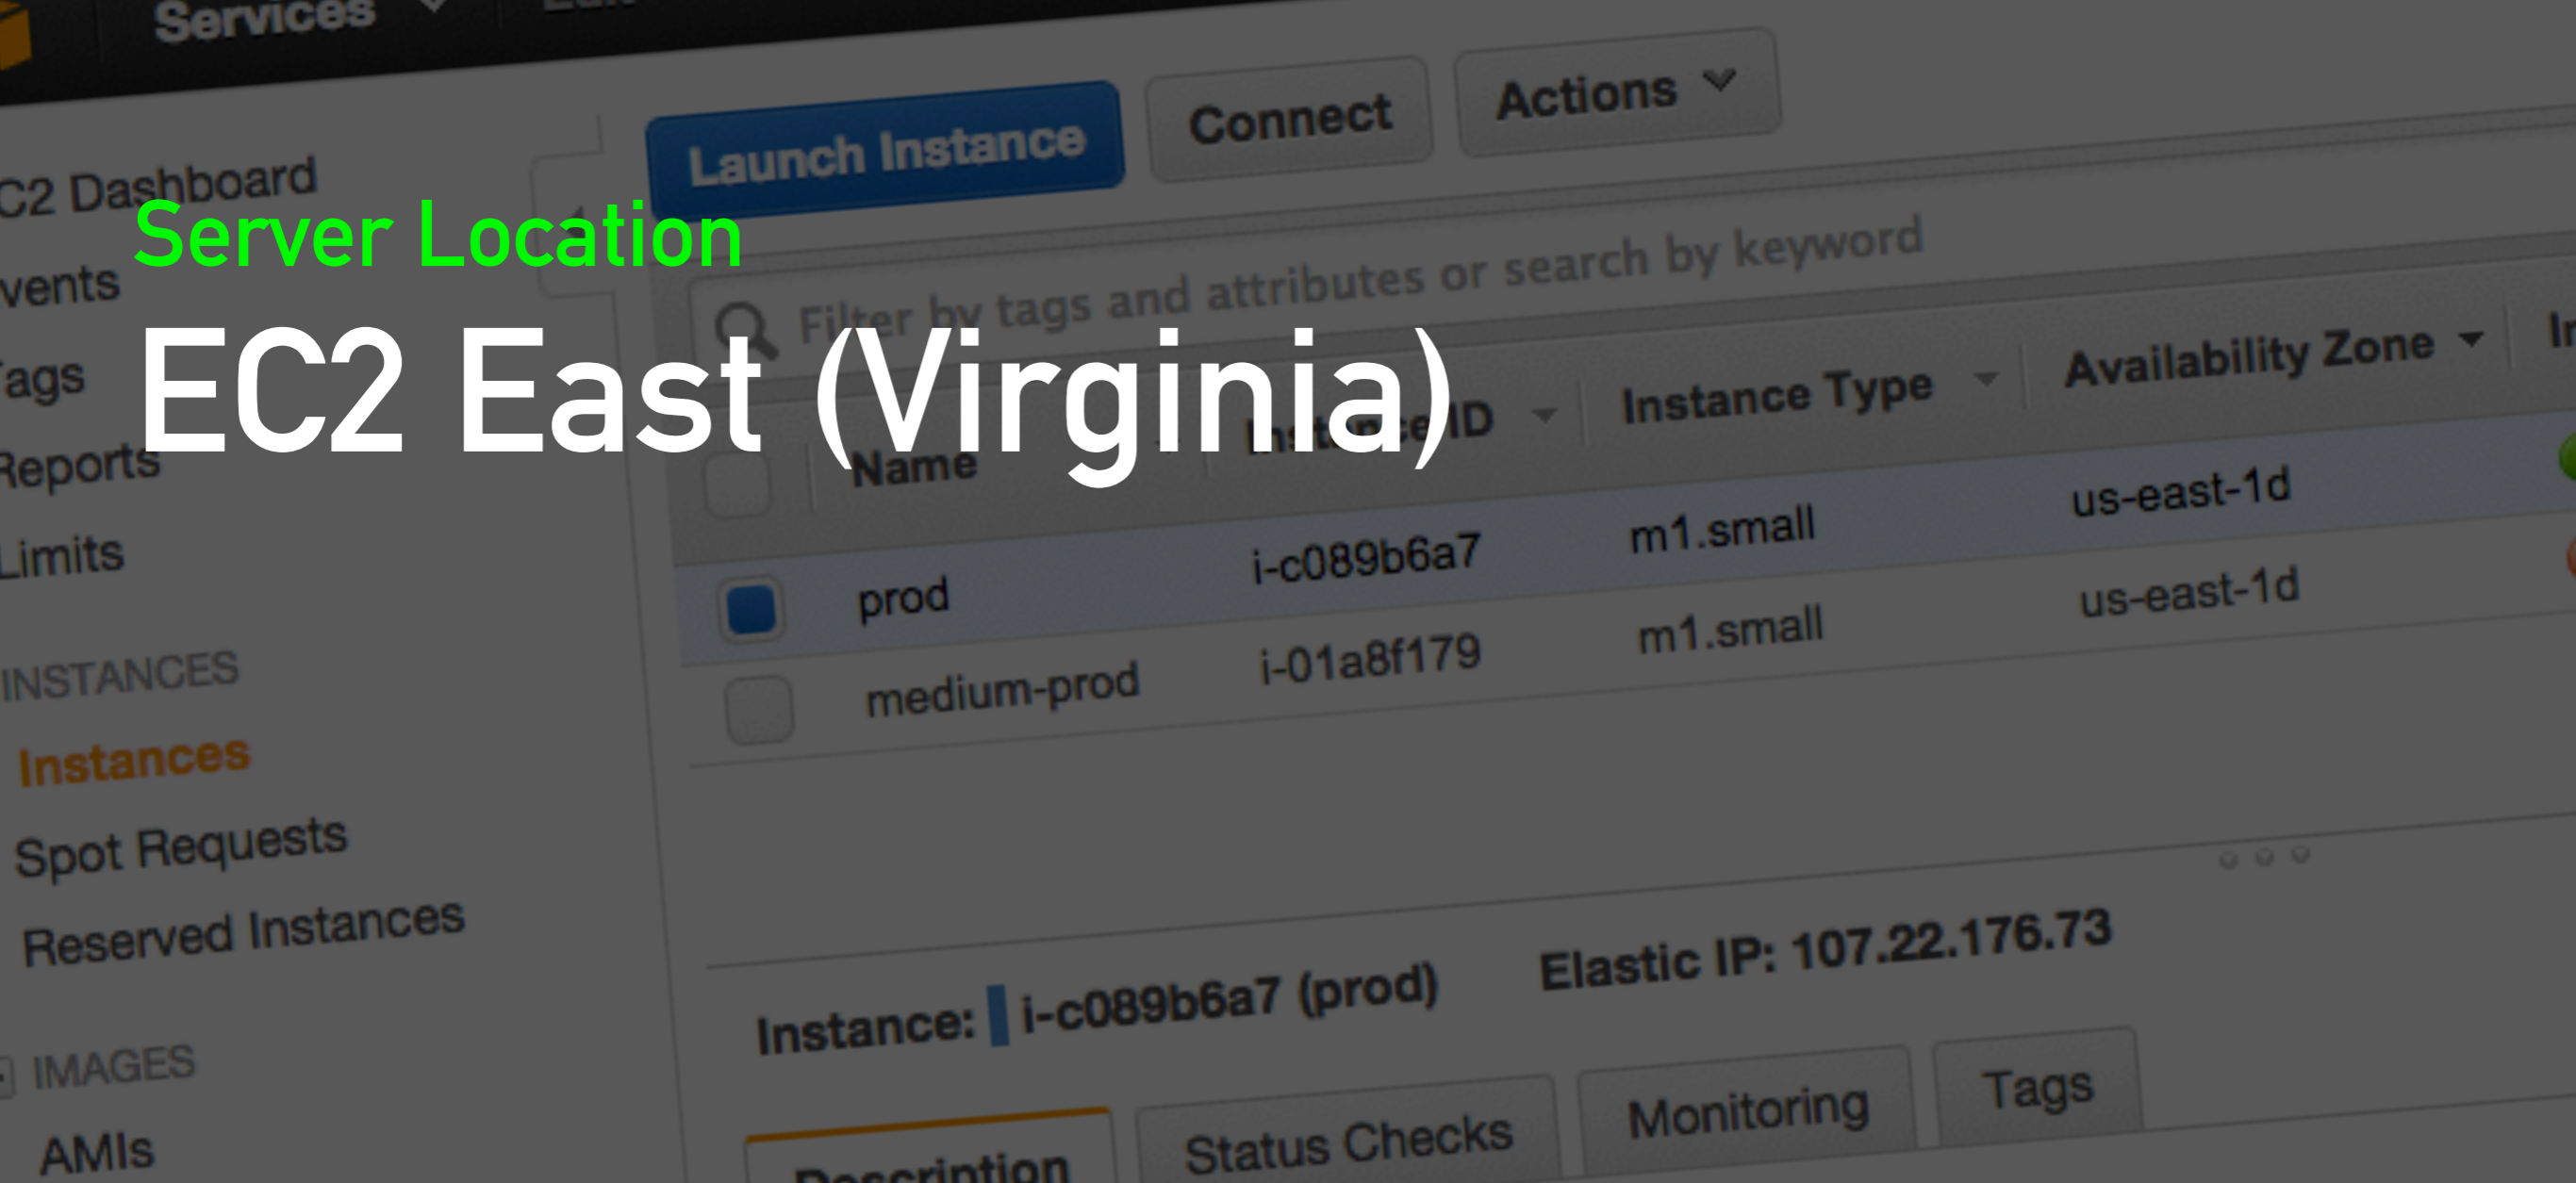


onslyde details

- Platform for real time feedback
- Started in early 2012
- Custom fallbacks no JS framework
- github.com/onslyde

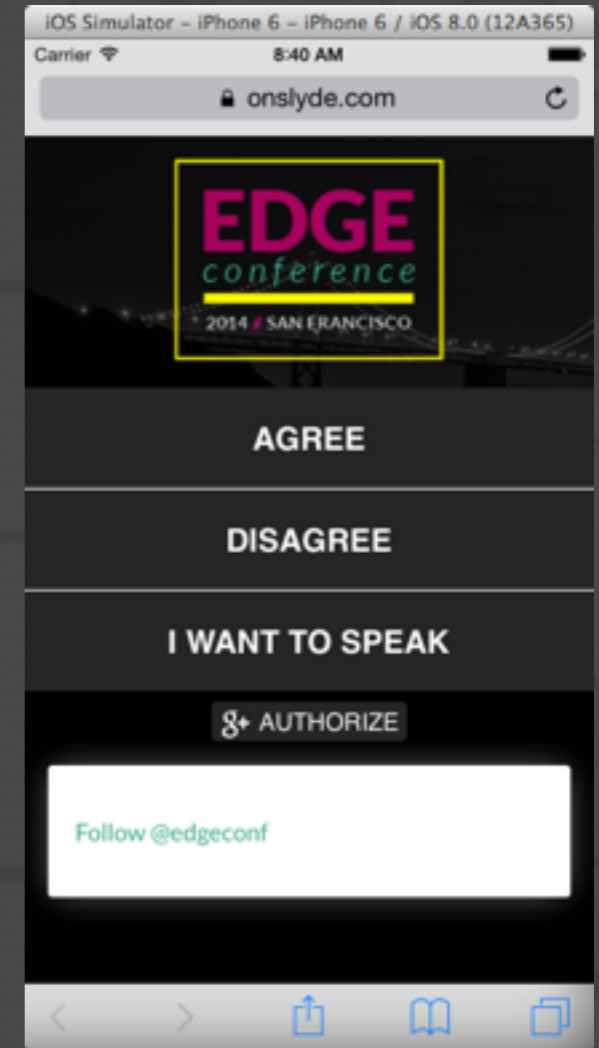
Server Location

EC2 East (Virginia)



Browser Locations

Live Event Based






Live Event Based

Mobile Device Info	Sessions	Service Provider	Sessions
Apple iPhone	140	tw telecom holdings inc.	282
LG Nexus 5	62	service provider corporation	52
(not set)	26	(not set)	37
Apple iPad	14	t-mobile usa inc.	13
Mozilla Firefox for Android	12	cellco partnership dba verizon wireless	6
BlackBerry KBD	9	softlayer technologies inc.	5
Google Nexus 4	9	sprint nextel corporation	5
Mozilla Firefox OS	7	comcast cable communications holdings	4
Motorola XT1053 Moto X	6	agile-inap	3
LG LGMS323 Optimus L70	5	vodafone limited	2

Battery Consumption

Visits and Pages / Visit by Mobile Device Info

Mobile Device Info		Visits	Pages / Visit
Apple iPhone		252	3.57
Google Nexus 4		35	2.37
Apple iPad		29	2.97
Samsung Galaxy Nexus		29	2.10
Google Nexus 7		26	2.81
(not set)		18	1.39
RIM Z10		11	1.18
Samsung GT-I9300 Galaxy SIII		8	2.62
Samsung GT-N7100 Galaxy Note II		8	1.12
Samsung GT-I9505 Galaxy S IV		7	1.86

Visits by Browser

Browser

Chrome

Safari

Firefox

Android Brows

BlackBerry

Safari (in-app)

Avg. Visit Dur

● Avg. Visit D

00:21:40


@wesleyhailes

Server Stack

Supports the latest WebSocket
& SPDY protocols

SPDY, WSS, and 443


- Layering is not currently supported in browsers
- Why do it?



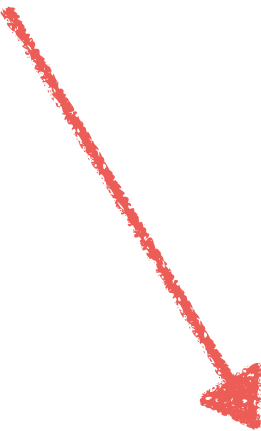
SPDY Goodness



Pre-Concat over SPDY3




	/foundation-4.3.1/js/foundation	GET
	angular.js /js/libs/angular-1.2.0-rc.2	GET
	angular-resource.js /js/libs/angular-1.2.0-rc.2	GET
57 requests 126 KB transferred 4.26 s (load: 2.86 s, DOMContentLoaded: 2.81 s)		

Post-Concat over SPDY3




	index.min.js	GET
	entypo.woff	GET
38 requests 121 KB transferred 4.24 s (load: 3.01 s, DOMContentLoaded: 2.97 s)		

SPDY Goodness



SPDY Badness



SPDY and HTTP 2

“HTTP/2 isn’t magic Web performance pixie dust; you can’t drop it in and expect your page load times to decrease by 50%.”

- Mark Nottingham (chair)

About WS Data

× Headers Frames Cookies		
Data	Length	Time
{"onslydeEvent":{"sessionId":"619","fire":function(){window.eventObjb = document.createEvent('Event');eventO...	244	10:17:04 AM
{"onslydeEvent":{"sessionId":"619","fire":function(){window.eventObjg = document.createEvent('Event');eventO...	317	10:17:04 AM
activeOptions:null,null,0:0	27	10:17:04 AM
Binary Frame (Opcode 2, mask)	0	10:17:04 AM
{"remoteMarkup":"%3Ch2 Binary Frame (Opcode 2, mask) %3EMy%20awesome%20presentation%3C%2Fh2%3E"}	86	10:17:04 AM
{"onslydeEvent":{"sessionId":"619","fire":function(){window.eventObjg = document.createEvent('Event');eventO...	358	10:17:03 AM
{"onslydeEvent":{"sessionId":"619","fire":function(){window.eventObjb = document.createEvent('Event');eventO...	244	10:17:03 AM
::connect::	11	10:17:03 AM

What is “High Performance” WebSocket?



Performance

32 sides to the perf coin:

- Server
- **Connections**
- **Client side**

FACE_KRISH
0 DISAGREE
3 AGREE

onslyde

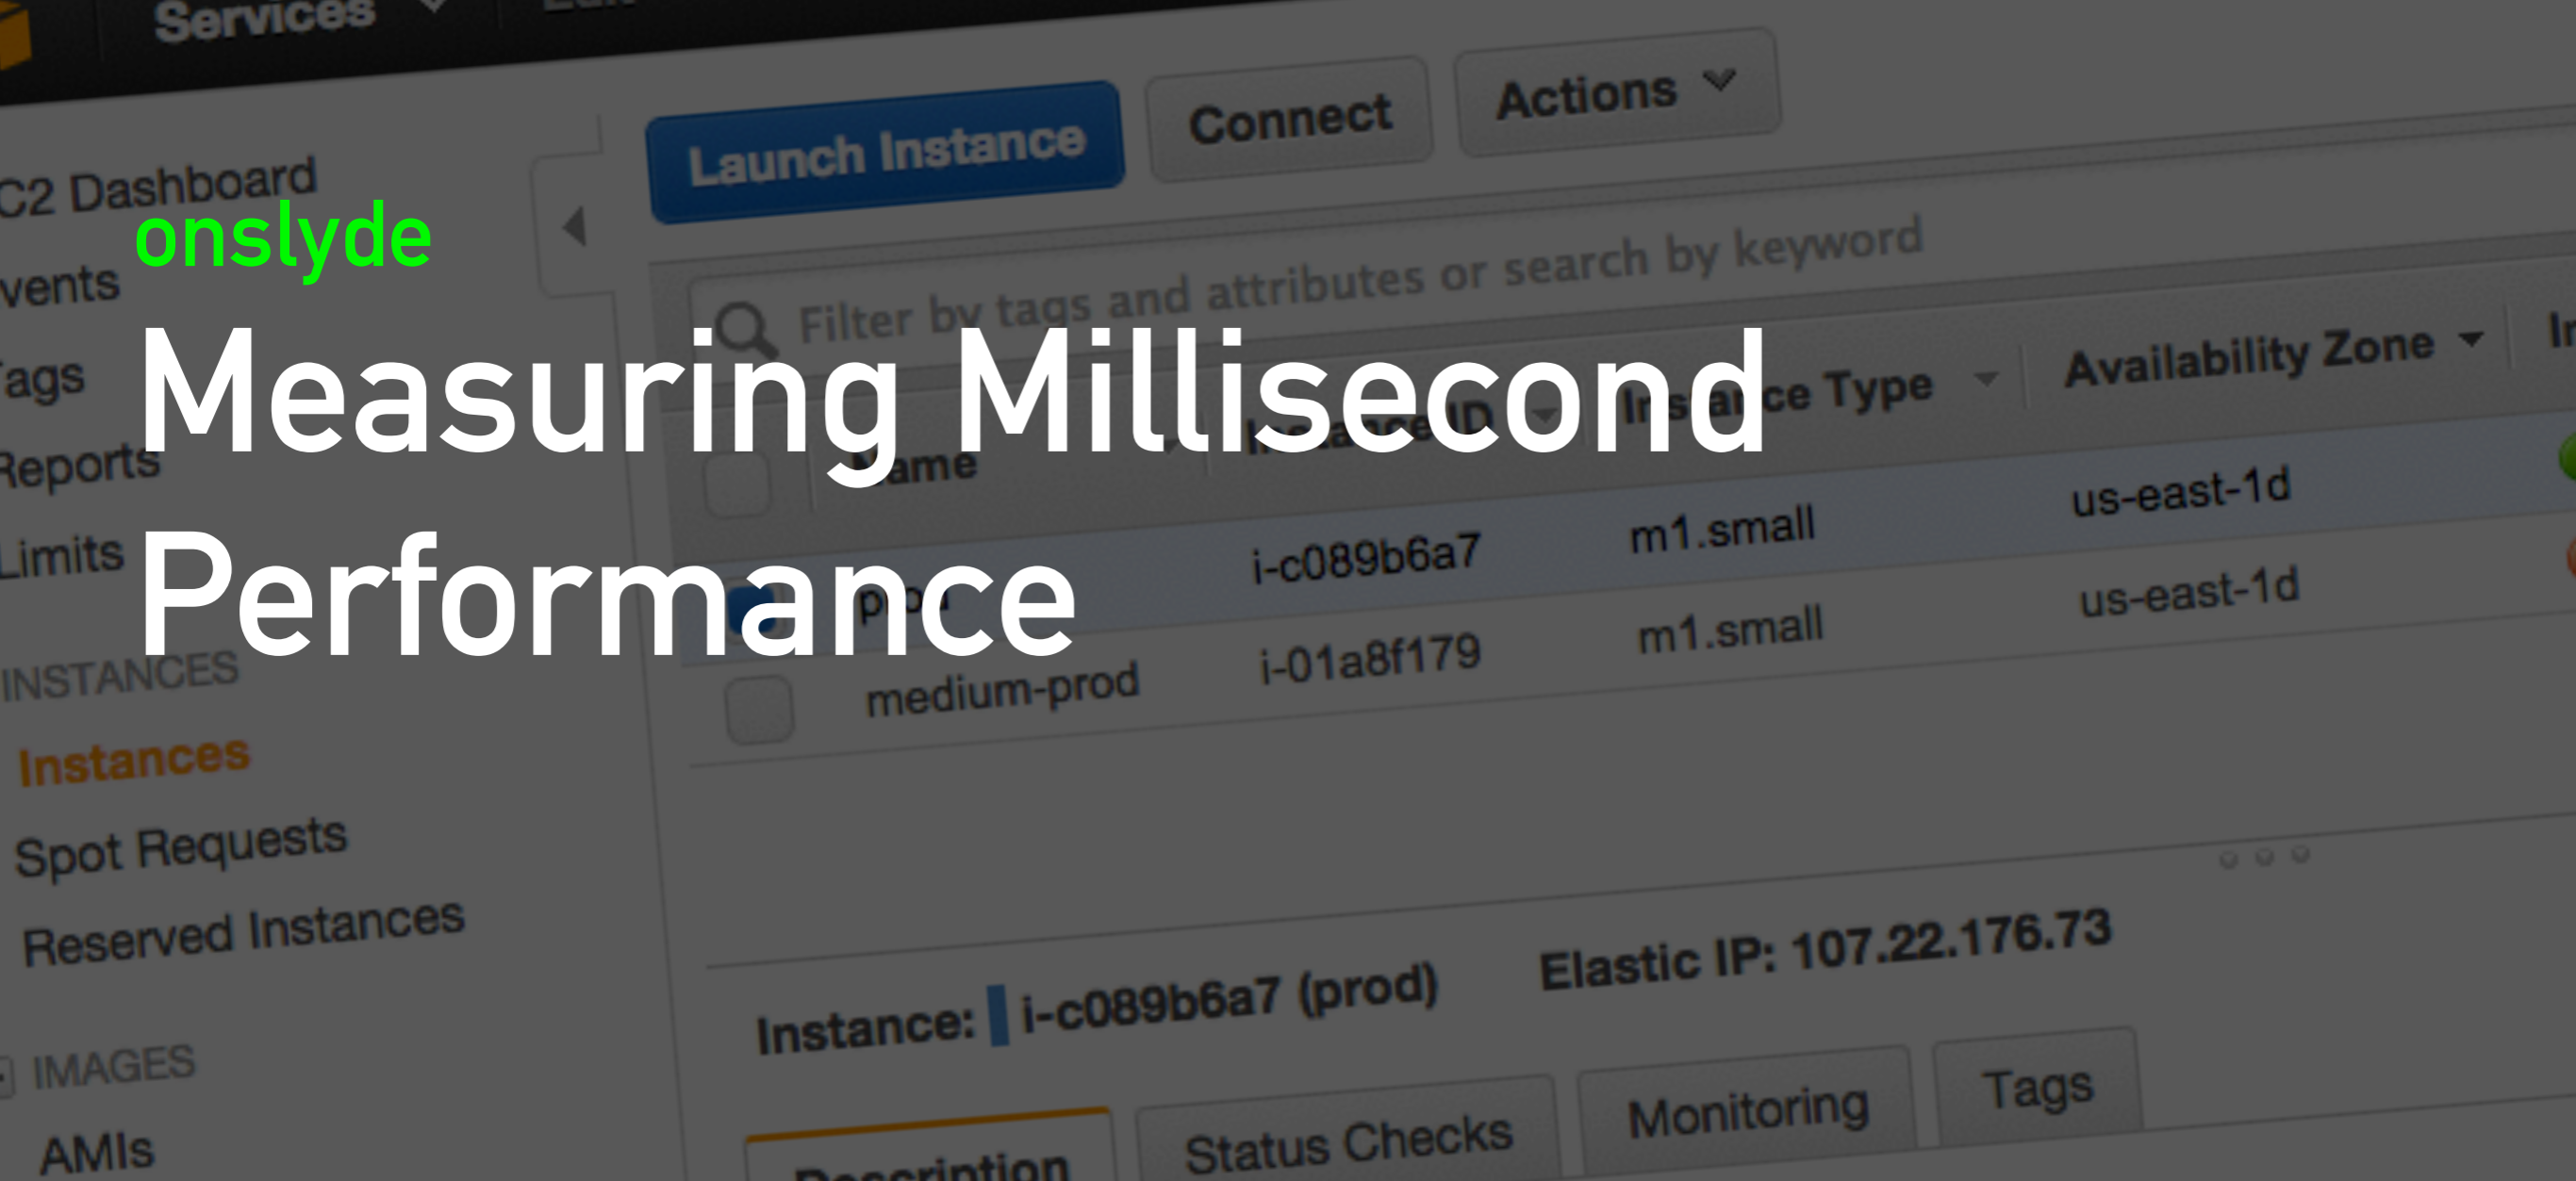
Demo



disag

14:03

Measuring Millisecond Performance



Getting Server Time

Headers vs. Content rewrite

- Headers
- Content Rewrite

Sync Client with Server Time

Unfortunately, there isn't an API to give you the HTTP response headers for your initial page request.

So let's try AJAX...

Date Header

```
var oReq = new XMLHttpRequest();
oReq.onload = function(){
    var dateStr = oReq.getResponseHeader('Date');
    var serverTimeMillisUTC =
        new Date(Date.parse(dateStr)).getTime();
    var localMillisUTC =
        new Date().getTime();
    console.log(new Date(serverTimeMillisUTC),
        new Date(localMillisUTC))
};
oReq.open("HEAD", "/", false);
oReq.send();
```


Date Header - Round Trip

```
var requestStart = new Date();
var oReq = new XMLHttpRequest();
oReq.onload = function(){
    var dateStr = oReq.getResponseHeader('Date');
    var serverTimeMillisUTC =
        new Date(Date.parse(dateStr)).getTime();
    var localMillisUTC =
        new Date().getTime();
    reqRespOffset = localMillisUTC - requestStart;
    offset =
        (serverTimeMillisUTC - localMillisUTC) + responseOffset;
};
oReq.open("HEAD", "/", false);
oReq.send();
```

Date Header and Trip

```
var requestDate = new Date();
var oReq = new XMLHttpRequest();
oReq.onreadystatechange = function(){
    var dateStr = oReq.getResponseHeader('Date');
    var serverTimeMillisUTC =
        new Date(Date.parse(dateStr)).getTime();
    var localTimeMillisUTC =
        new Date().getTime();
    reqResponseOffset = localTimeMillisUTC - requestStartTime;
    offset = (serverTimeMillisUTC - localTimeMillisUTC) - reqResponseOffset;
};
oReq.open("HEAD", "http://www.example.com", false);
oReq.send();
```

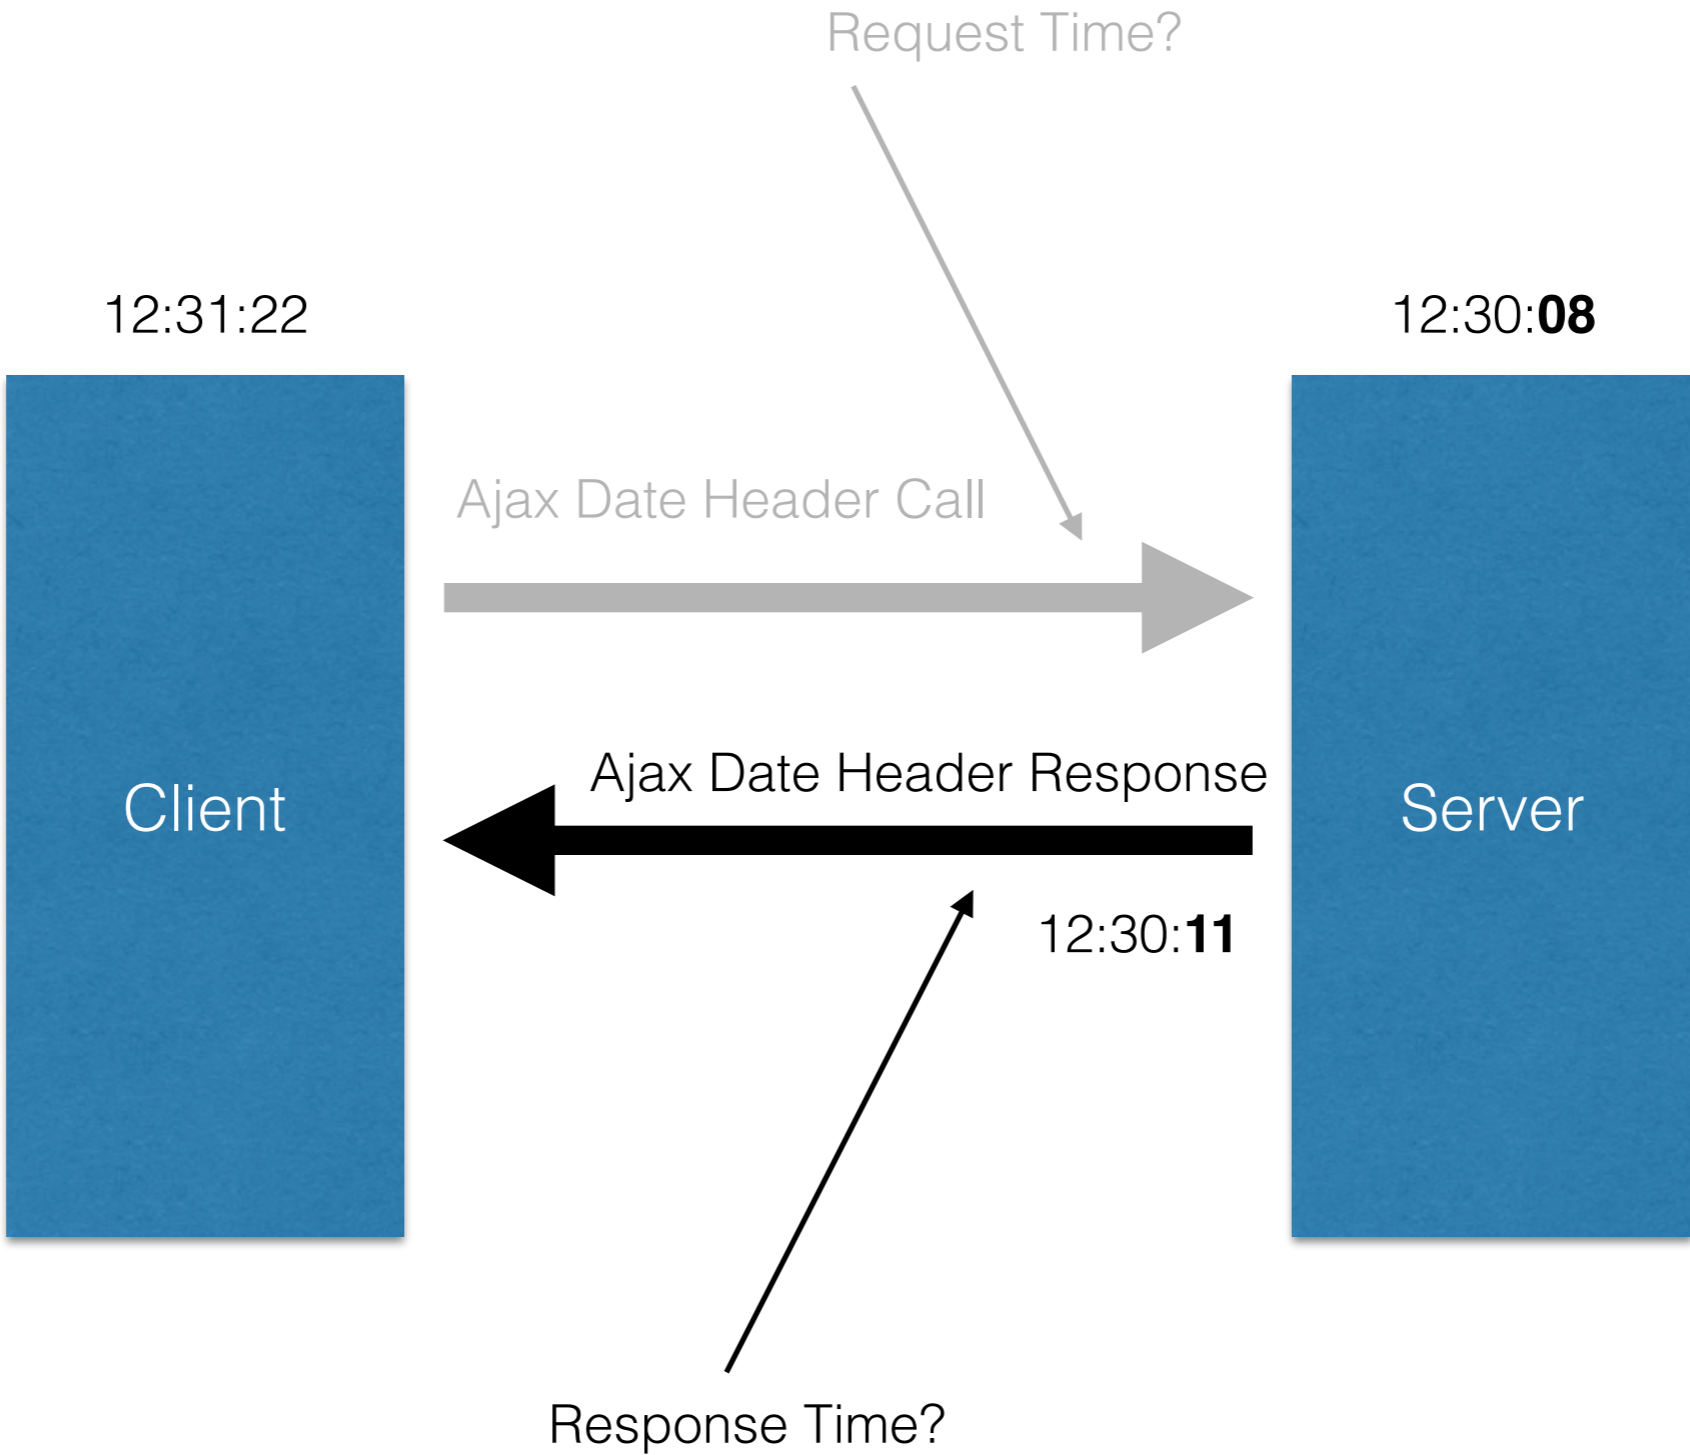
Instrumentation



Response Latency

- **First request (out of sync with server) has to be measured.**
- **Add total response time to offset.**

Instrumentation



Navigation Timing

Why can't we just pull the date header from the initial GET request and calculate with Navigation Timing API?

Perf Timing - Response Only

```
var requestStart = performance.timing.now;
var oReq = new XMLHttpRequest();
oReq.onload = function(){
  var dateStr = oReq.getResponseHeader('Date');
  var serverTimeMillisUTC =
    Date.parse(new Date(Date.parse(dateStr)).toUTCString());
  var localMillisUTC =
    Date.parse(new Date().toUTCString());
  responseOffset =
    performance.timing.responseEnd - performance.timing.responseStart;
  offset = (serverTimeMillisUTC - localMillisUTC) + responseOffset;
};
oReq.open("HEAD", "/", false);
oReq.send();

function getServerTime() {
  var date = new Date();
  date.setTime(date.getTime() + offset);
  return date;
}
```

Perf Timing Response Only

```
var requestStart = performance.timing.now
var oReq = new XMLHttpRequest();
oReq.onload = function(){
  var date = oReq.getResponseHeader('Date');
  var serverTimeMillisUTC =
    Date.parse(new Date(Date.parse(date)).toUTCString());
  var localTimeMillisUTC =
    Date.parse(new Date().toUTCString());
  responseTime =
    performance.timing.responseStart - performance.timing.requestStart;
  offset = (serverTimeMillisUTC - localTimeMillisUTC) + responseTime;
};
oReq.open("GET", "/");
oReq.send();

function getServerTime(){
  var date = new Date();
  date.setTime(date.getTime() + offset);
  return date;
}
```

Nav Timing - Response Only

AJAX requests don't have the Performance Timing API.

Solution (Hack)

- **Write the epoch timestamp to the content.**
- **Use the navigation timing API with that timestamp to calculate the offset.**

Perf Timing - Response Only

```
var serverTimeMillisUTC =  
    new Date(document.getElementById('time-holder'));  
var localMillisUTC = new Date().getTime;  
responseOffset =  
    performance.timing.responseEnd -  
    performance.timing.responseStart;  
var offset =  
    (serverTimeMillisUTC - localMillisUTC) + responseOffset;  
  
function getServerTime() {  
    var date = new Date();  
    date.setTime(date.getTime() + offset);  
    return date;  
}
```

Solution (Legit)

- **Use Resource Timing**

Resource Timing

- Browser Support

IE	Firefox	Chrome	Safari	Opera	iOS Safari *	Opera Mini *	Android Browser *	Chrome for Android
		31						
		33						
8		35	5.1				4.1	
9	31	36	7		7.1		4.3	
10	32	37	7.1		8		4.4	
11	33	38	8	25	8.1	8	4.4.4	38
	34	39		26			37	
	35	40		27				
	36	41						

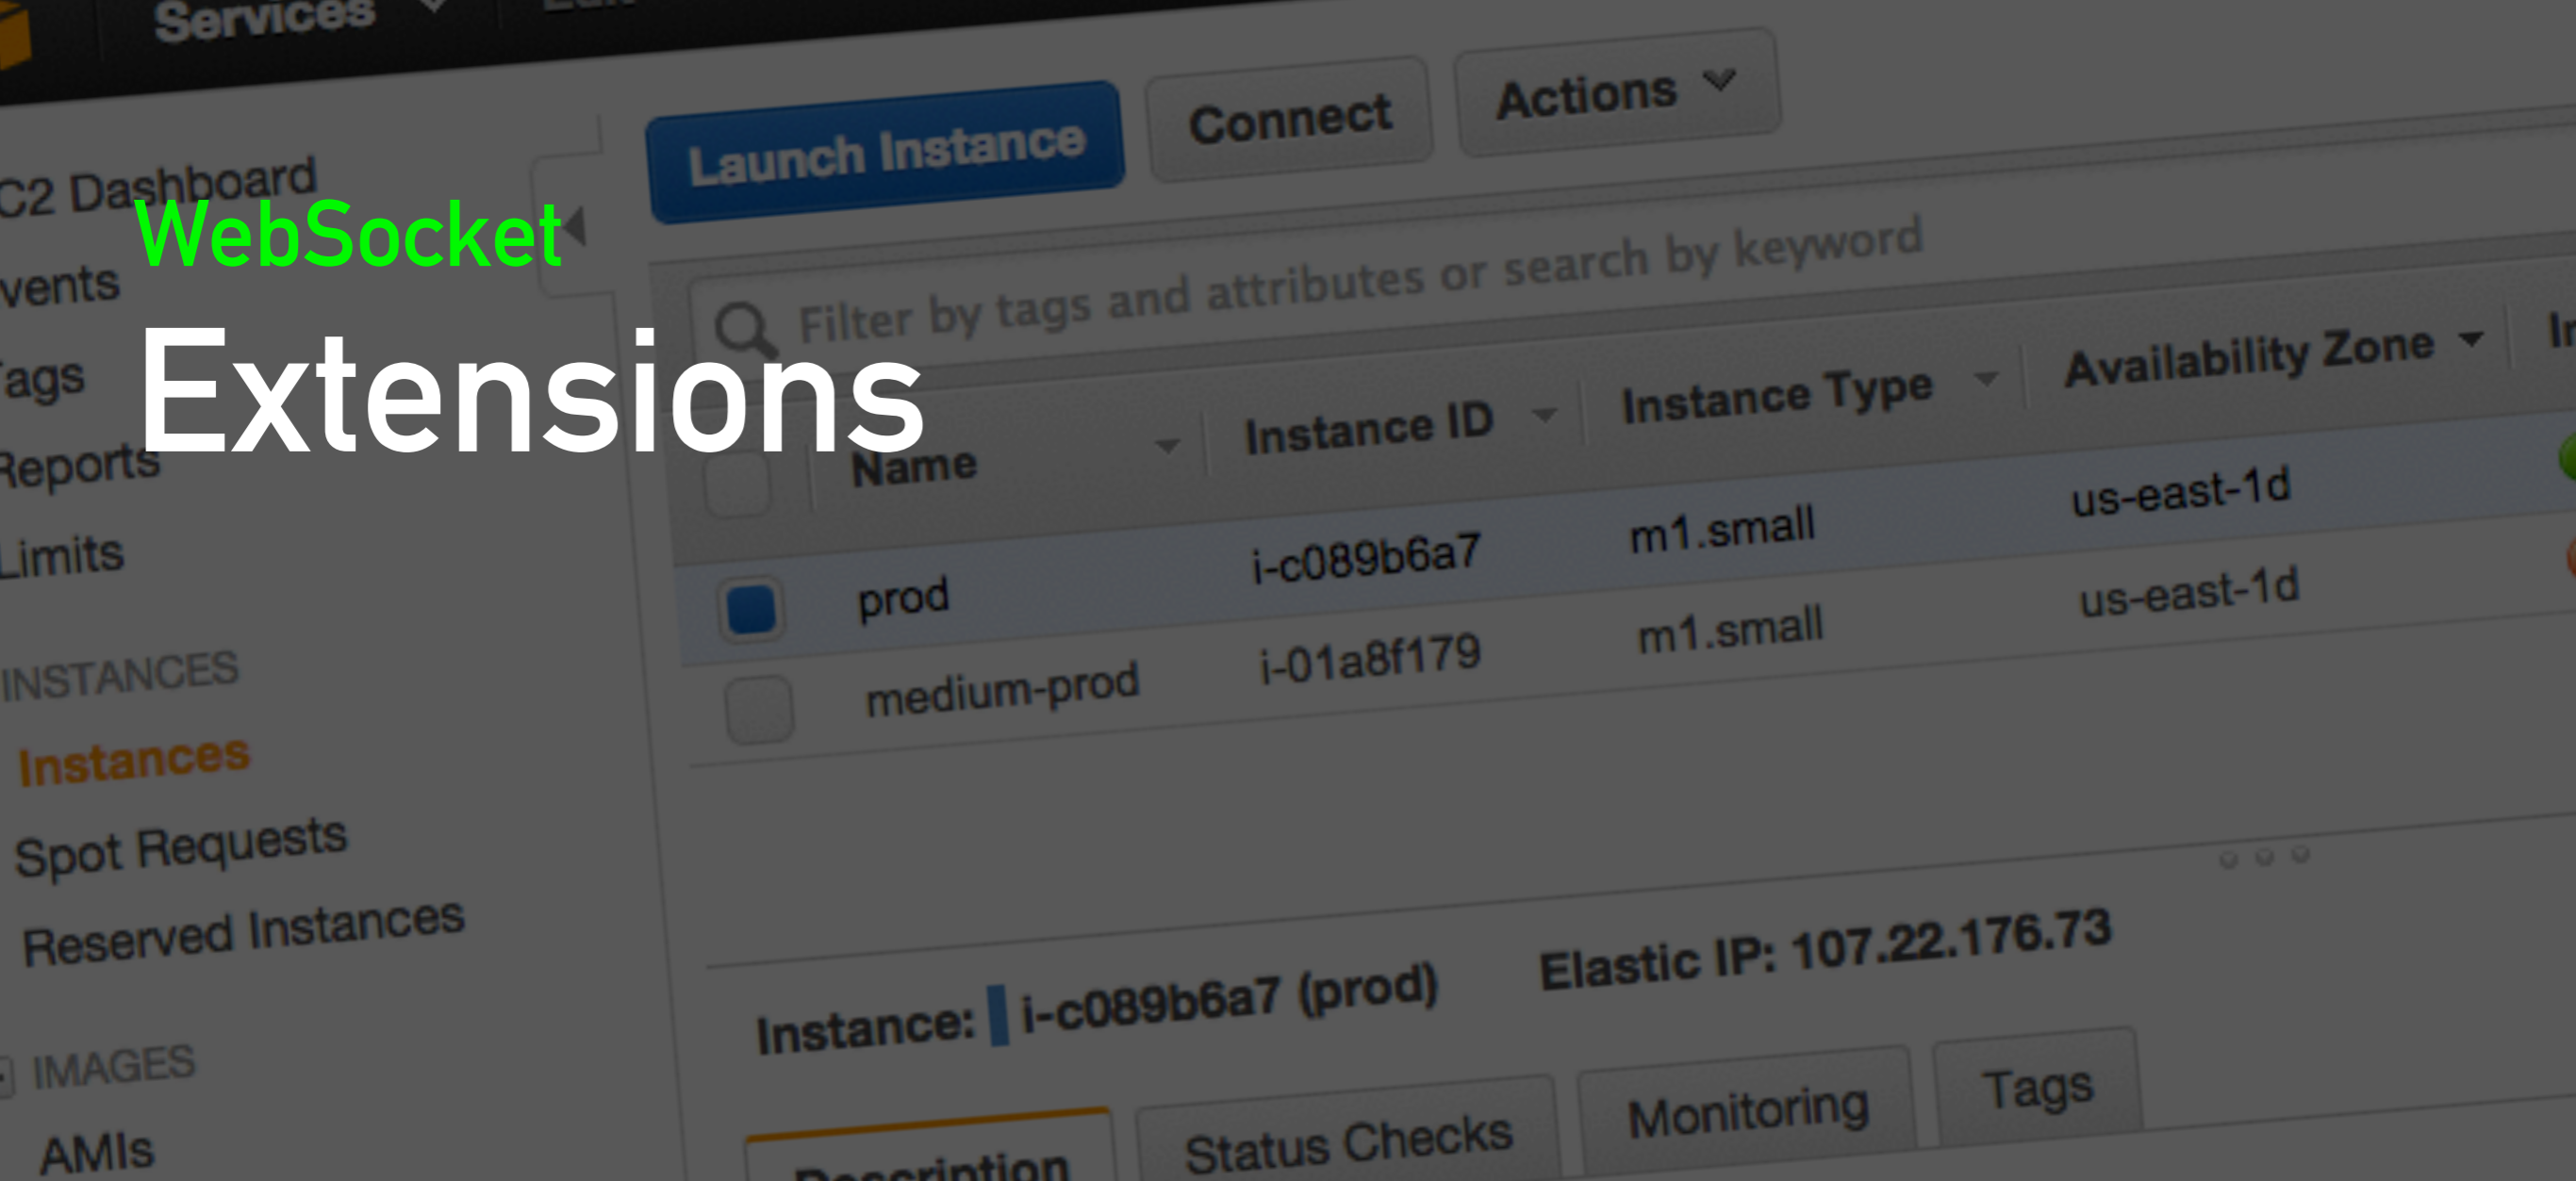
Resource Timing on AJAX

```
var serverTimeMillisUTC;
var oReq = new XMLHttpRequest();
oReq.onload = function(){
    var dateStr = oReq.getResponseHeader('Date');
    serverTimeMillisUTC =
        new Date(Date.parse(dateStr)).getTime();
};
oReq.open("HEAD", "/?foo123", false);
oReq.send();

function getServerTime() {
    window.performance.getEntriesByType("resource").
        reduceRight(function(previousValue, currentValue, index, array) {
            if(currentValue.name.indexOf('foo123') > 0){
                return serverTimeMillisUTC + (currentValue.responseEnd -
                                                currentValue.responseStart);
            }
        })
}
```

WebSocket

Extensions



WS Multiplexing and Head-of-Line Blocking

- **Frames from different messages can't be interleaved.**
- **A large message, even when split into multiple WebSocket frames, will block the delivery of frames associated with other messages.**

WebSocket Extensions

- **WebSocket specification allows for protocol extensions**
- **"Multiplexing Extension for WebSocket"**
- **"Compression Extensions for WebSocket"**

WebSocket

Performance Checklist



**Use secure WebSocket
(WSS over TLS) for
reliable deployments.**

Use compression.

```
Sec-WebSocket-Extensions: permessage-deflate; client_max_window_bits  
Sec-WebSocket-Key: /LqLDugB0+EK9Bnp6d3gqw==  
Sec-WebSocket-Version: 13
```

**Pay close attention to
polyfill performance (if
necessary).**

**Optimize binary payloads
to minimize transfer size.**

<https://github.com/onslyde/onslyde/blob/master/js/deck/onslyde-1.0.0.deck.js#L1581>

**Consider compressing
UTF-8 content to minimize
transfer size.**

<https://www.igvita.com/2013/11/27/configuring-and-optimizing-websocket-compression/>

**Split large application
messages to avoid head-
of-line blocking.**

Use subprotocol negotiation

http://chimera.labs.oreilly.com/books/1230000000545/ch17.html#_subprotocol_negotiation

Parting Words

HTTP2, SPDY, SSE, and WebSocket are all different. They could be combined.

Are you willing to polyfill your WebSocket Implementation?

Thanks!

@wesleyhales

SH-PE

shapesecurity.com