Improving the Performance and Security of AJAX Web Applications

Ben Livshits
Microsoft Research
Redmond, WA

http://research.microsoft.com/~livshits/

Web 2.0 is Upon Us

Source:
flickr.com

Web 1.0 → Web 2.0

Server-side

Advantage of the AJAX model:

greater application responsiveness

Client-side rendering

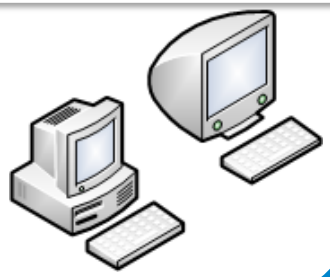

Outline of the Talk

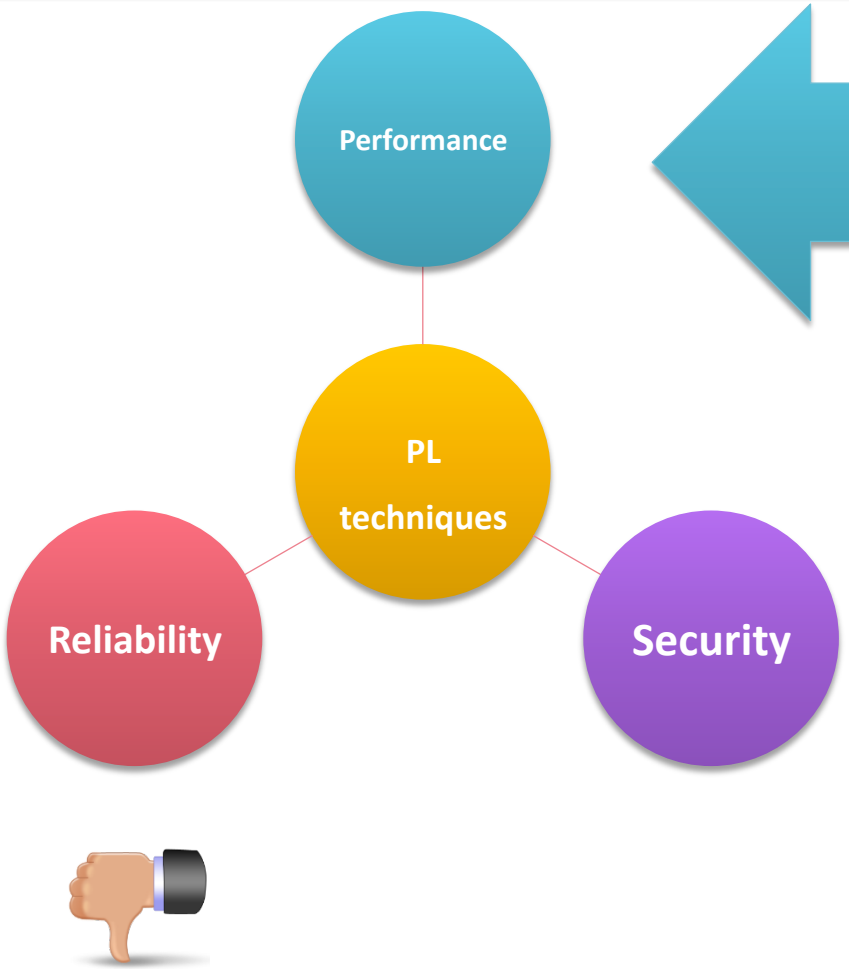

Doloto

Ripley

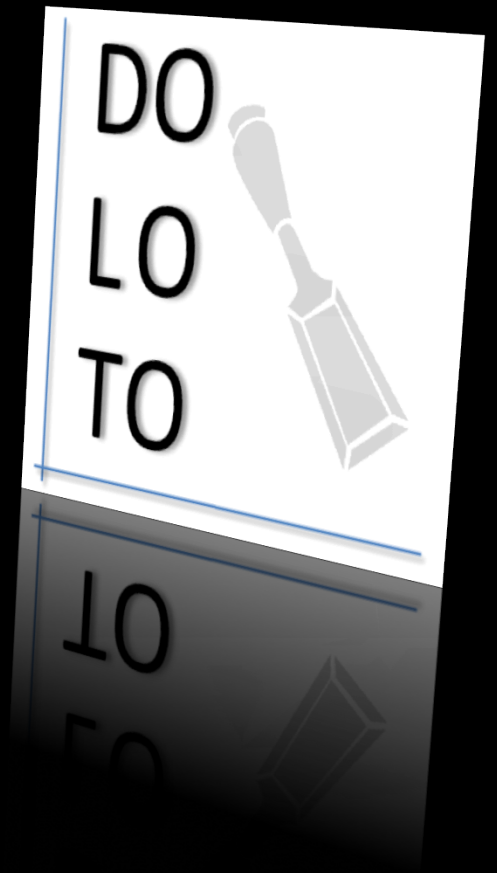

Doloto

Code Splitting for AJAX Applications

Outline of the Talk

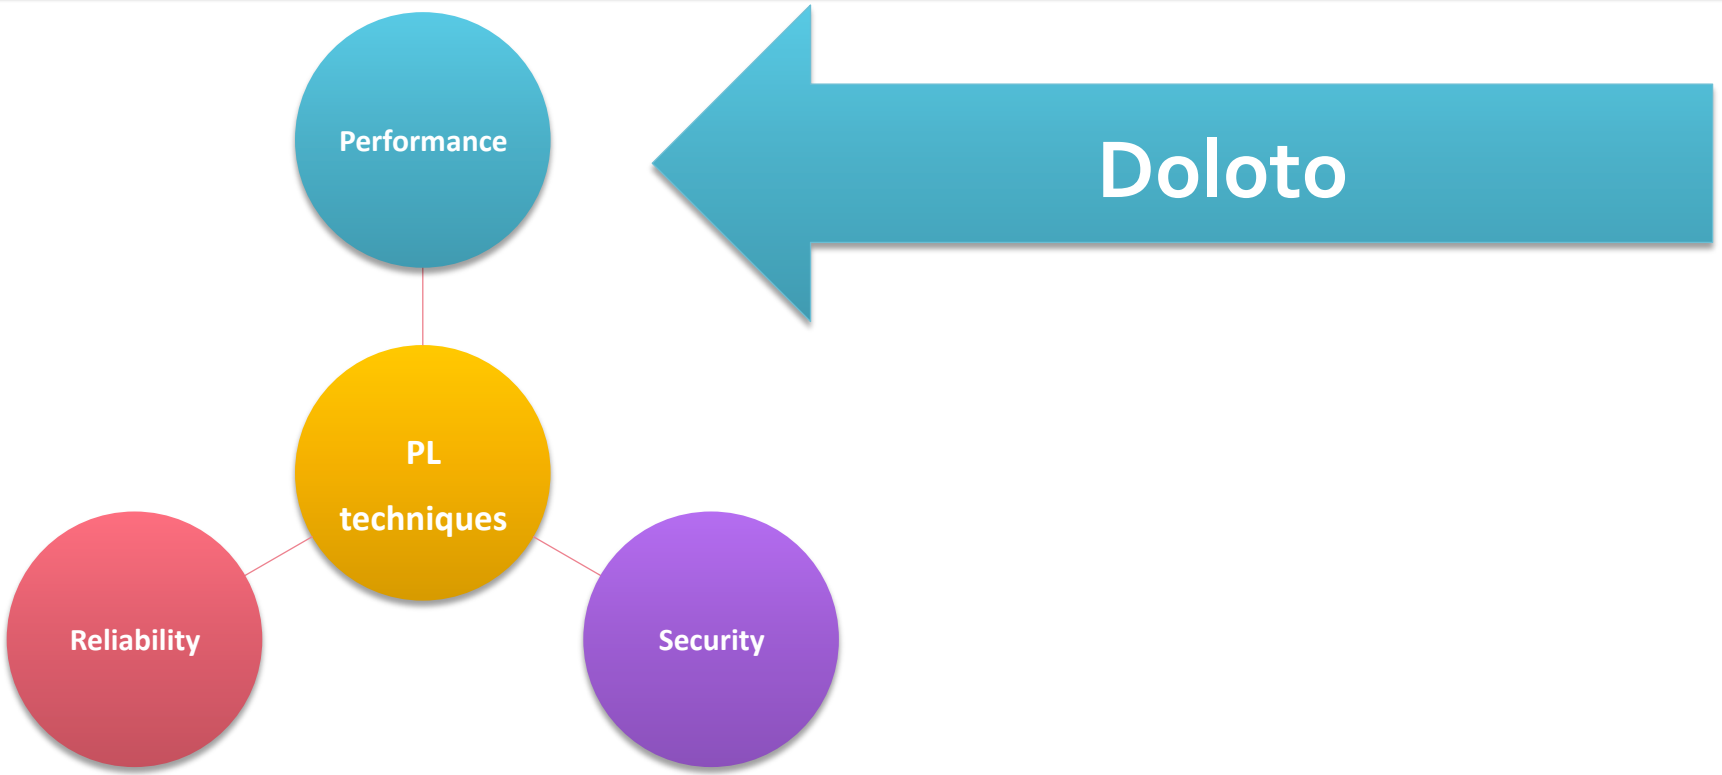

A Web 2.0 Application Disected

Talks to 14 backend services (traffic, images, directions, ads, …)

1+ MB code

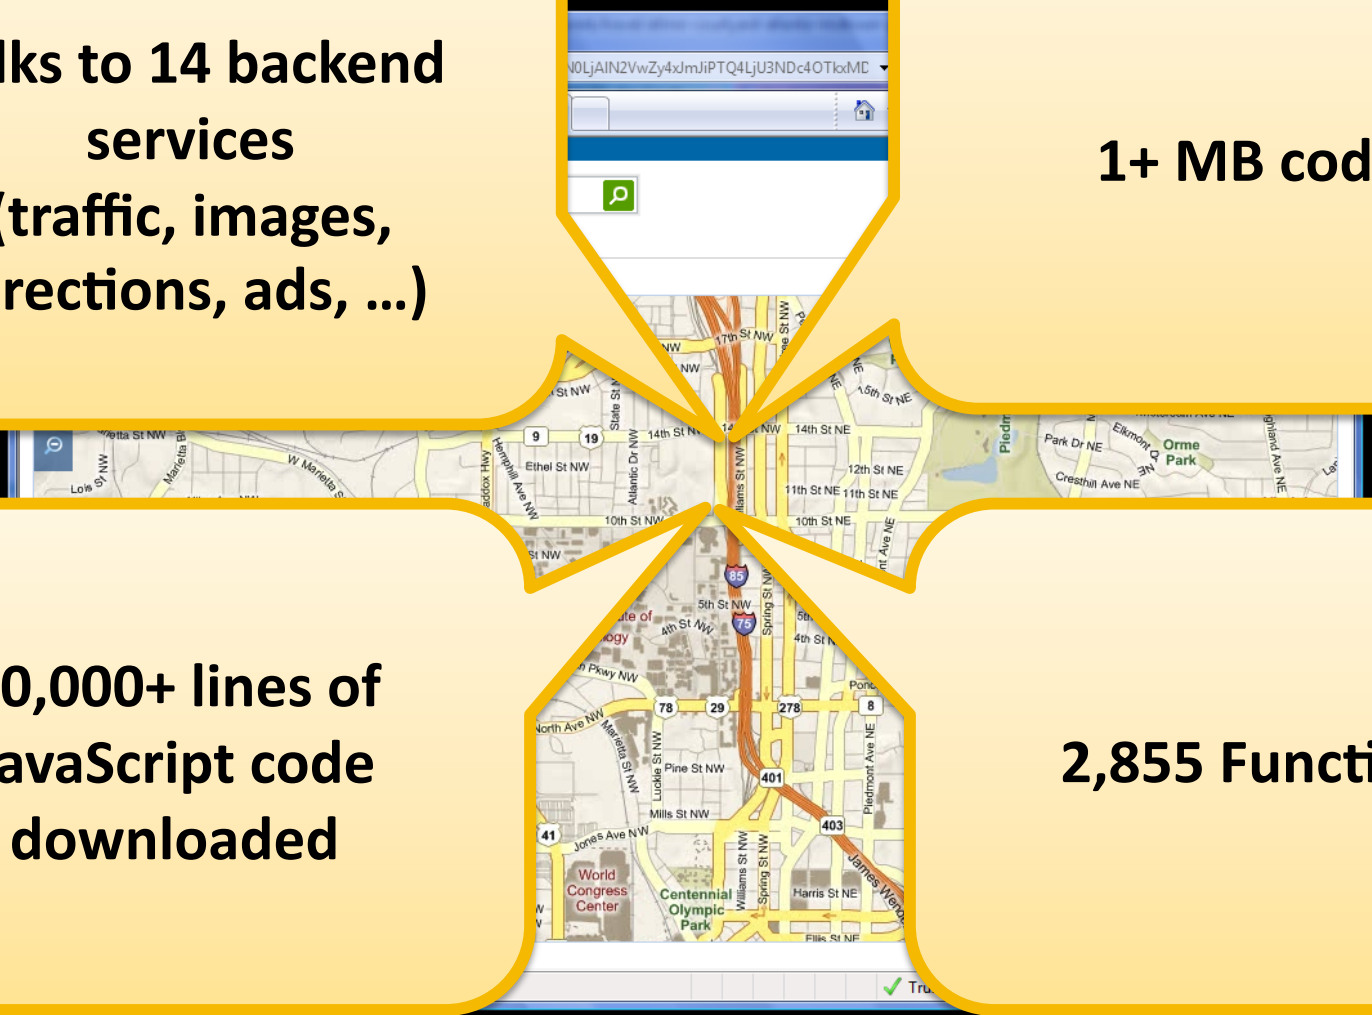

70,000+ lines of JavaScript code downloaded

2,855 Functions

Most of the application download is JavaScript code

Slows down application execution

AJAX Responsiveness: Catch-22

Move code to client for responsiveness

Execution can't start without the code

Motivation for Doloto

- Idea behind Doloto
  - Start with a small piece of code on the client
  - Download required code on demand (pull)
  - Send code when bandwidth available (push)
- Leads to better application responsiveness
  - Interleave code download & execution
  - Faster startup times
  - Rarely executed code is rarely downloaded

Doloto Workflow

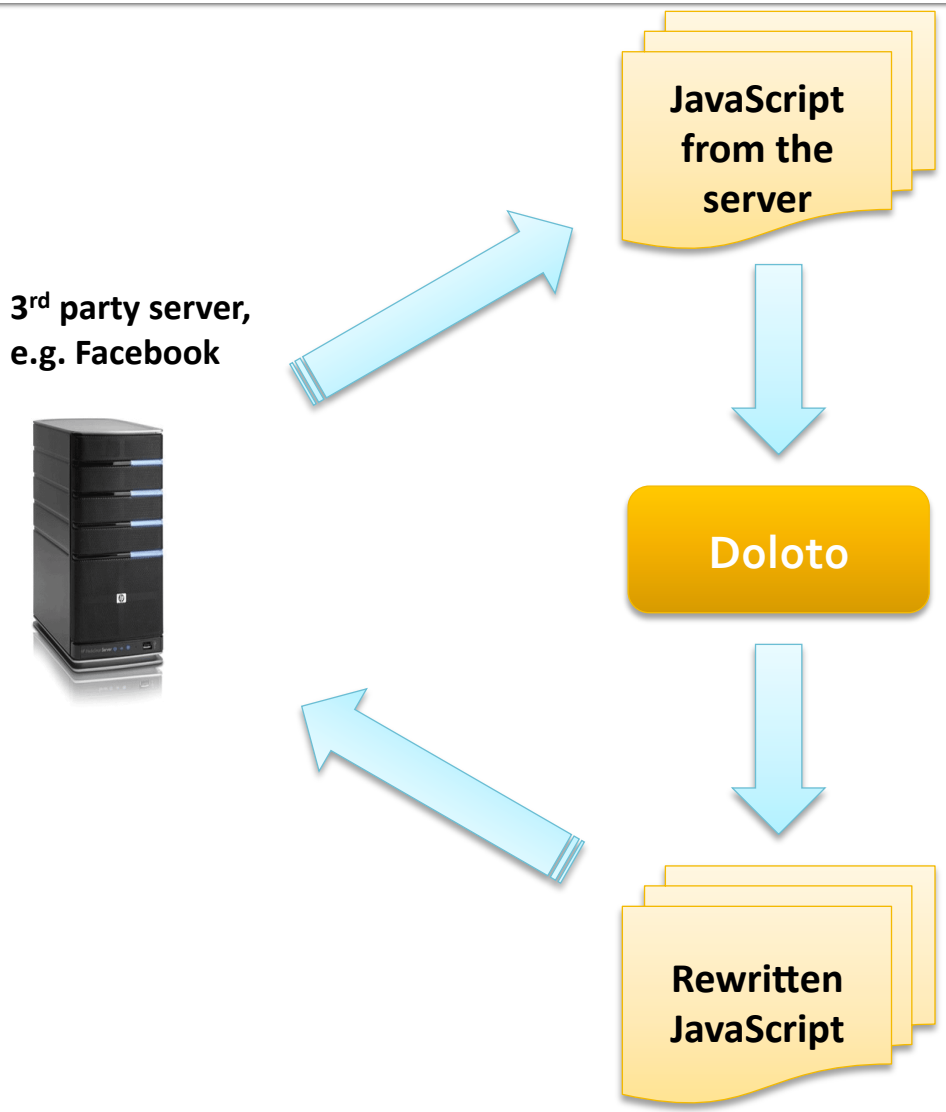

- Doloto intercepts JavaScript from the server using a proxy
- Instruments and rewrites it on the client
- Deploy it back to the server as the last step

Doloto: the Steps

1. [training] Runtime training to collect

- Instrument every function
- Record the first-execute timestamp
- Look for gaps to find clusters

stubbing" for on-demand code loading

Doloto Training Tool

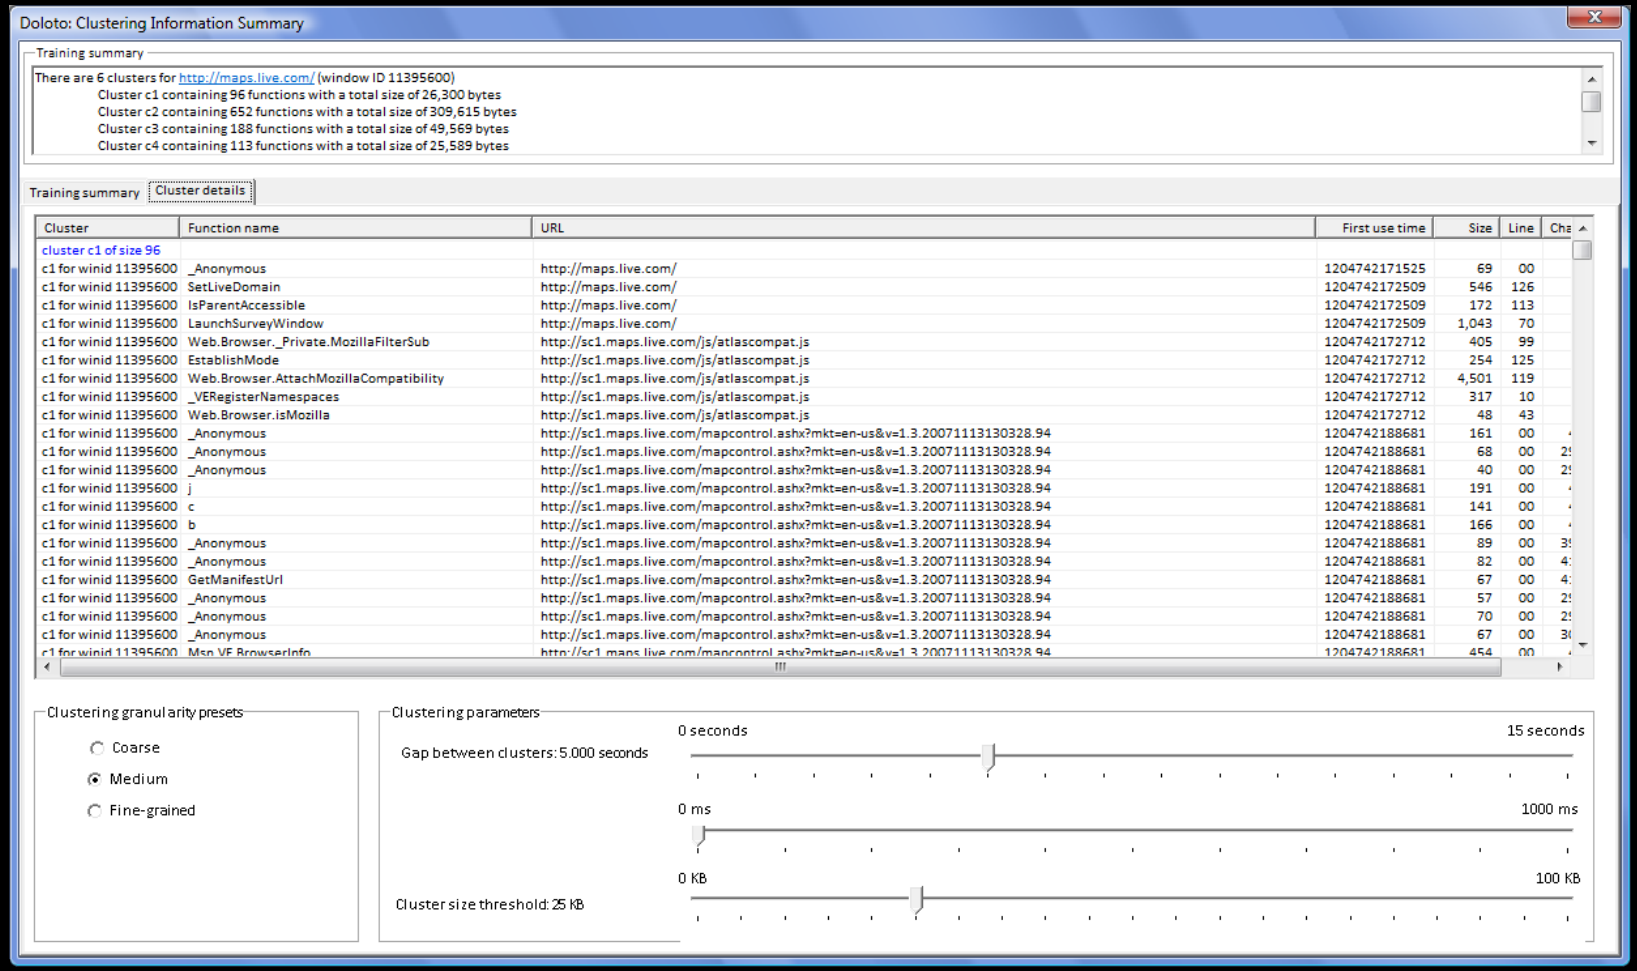

Architecture of Doloto

1. [**training**] Runtime training to collect access profile

3. [**clustering**] Grouping related functions together

4. [**rewriting**] Function rewriting or "stubbing" for on-demand code loading

Inserting Function Stubs

```
var g = 10;
function f1(y){
    var x=g+1;
    …
    …
    …
    …
    …
    return …;
}
```

```
var g = 10;

var real_f1;
function f1(y) {
    if(!real_f1){
        var code = load("f1");
        real_f1 = eval(code);
        f1 = real_f1;
    }
    return real_f1.apply(this,
                         arguments);
}
```

```
eval($exp("f1"), "y");        // 22 bytes
```

Profile applications using a proxy

Deploy rewritten code + cluster file to the server

Doloto Savings

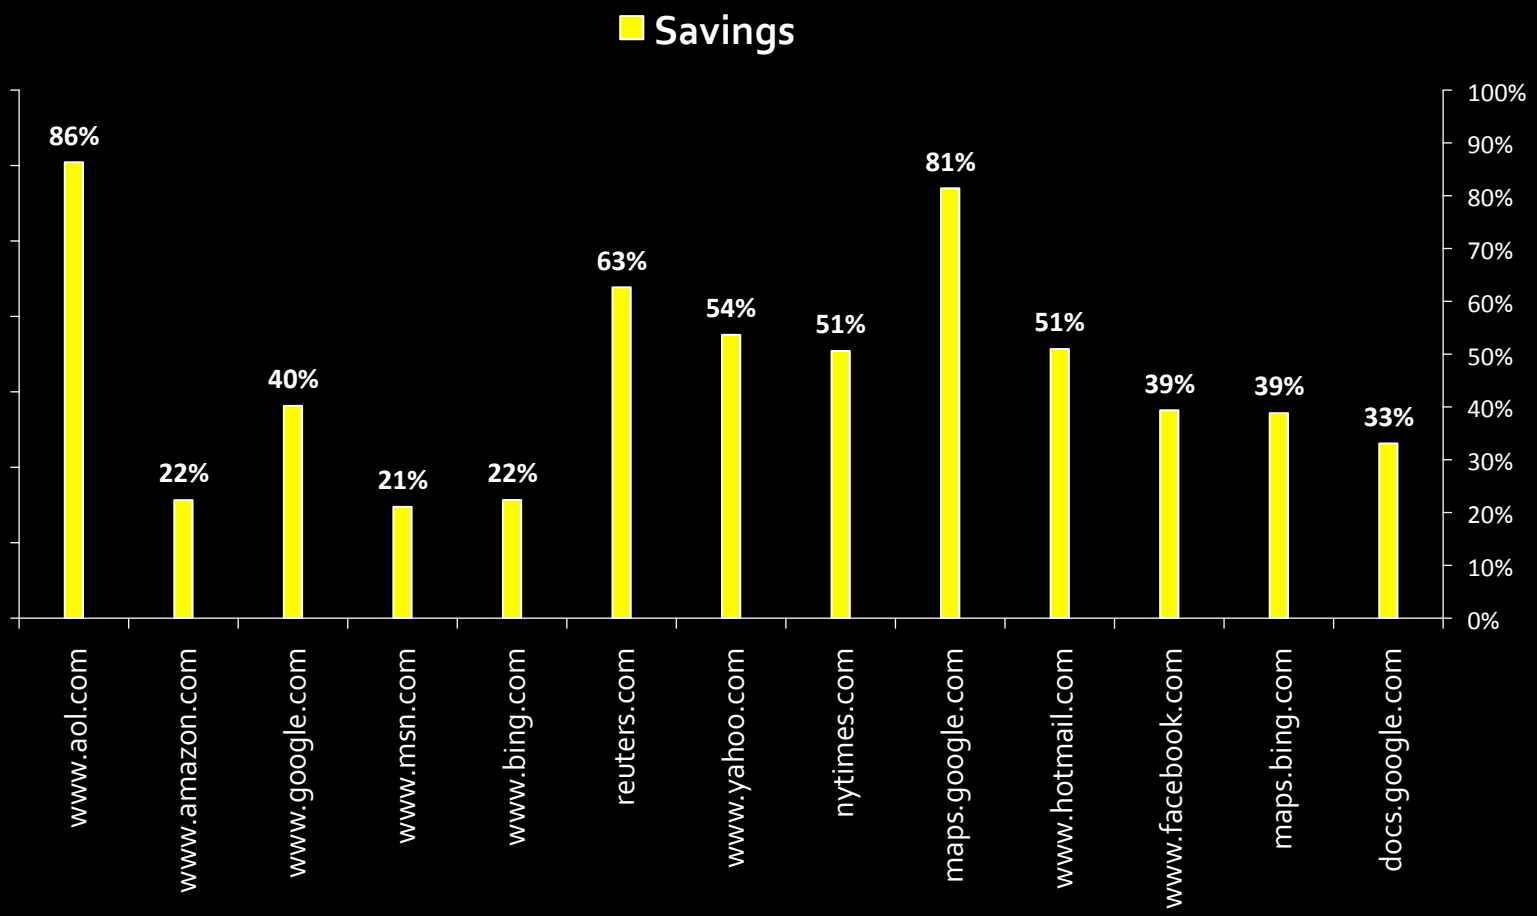

Doloto Savings

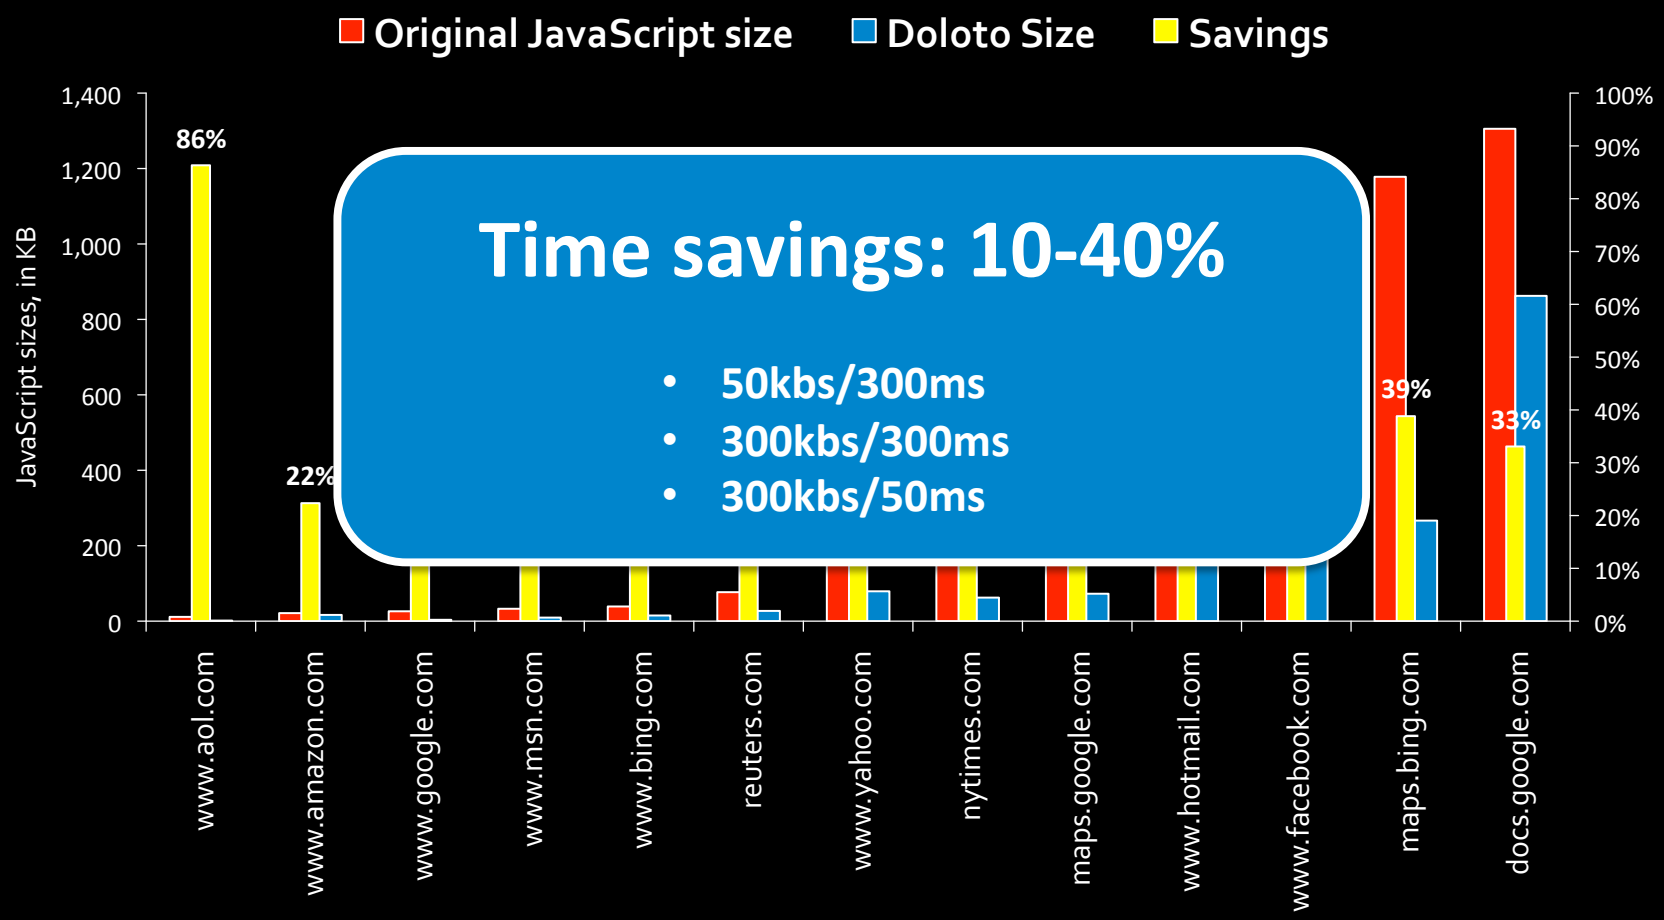

Runtime Savings with Doloto

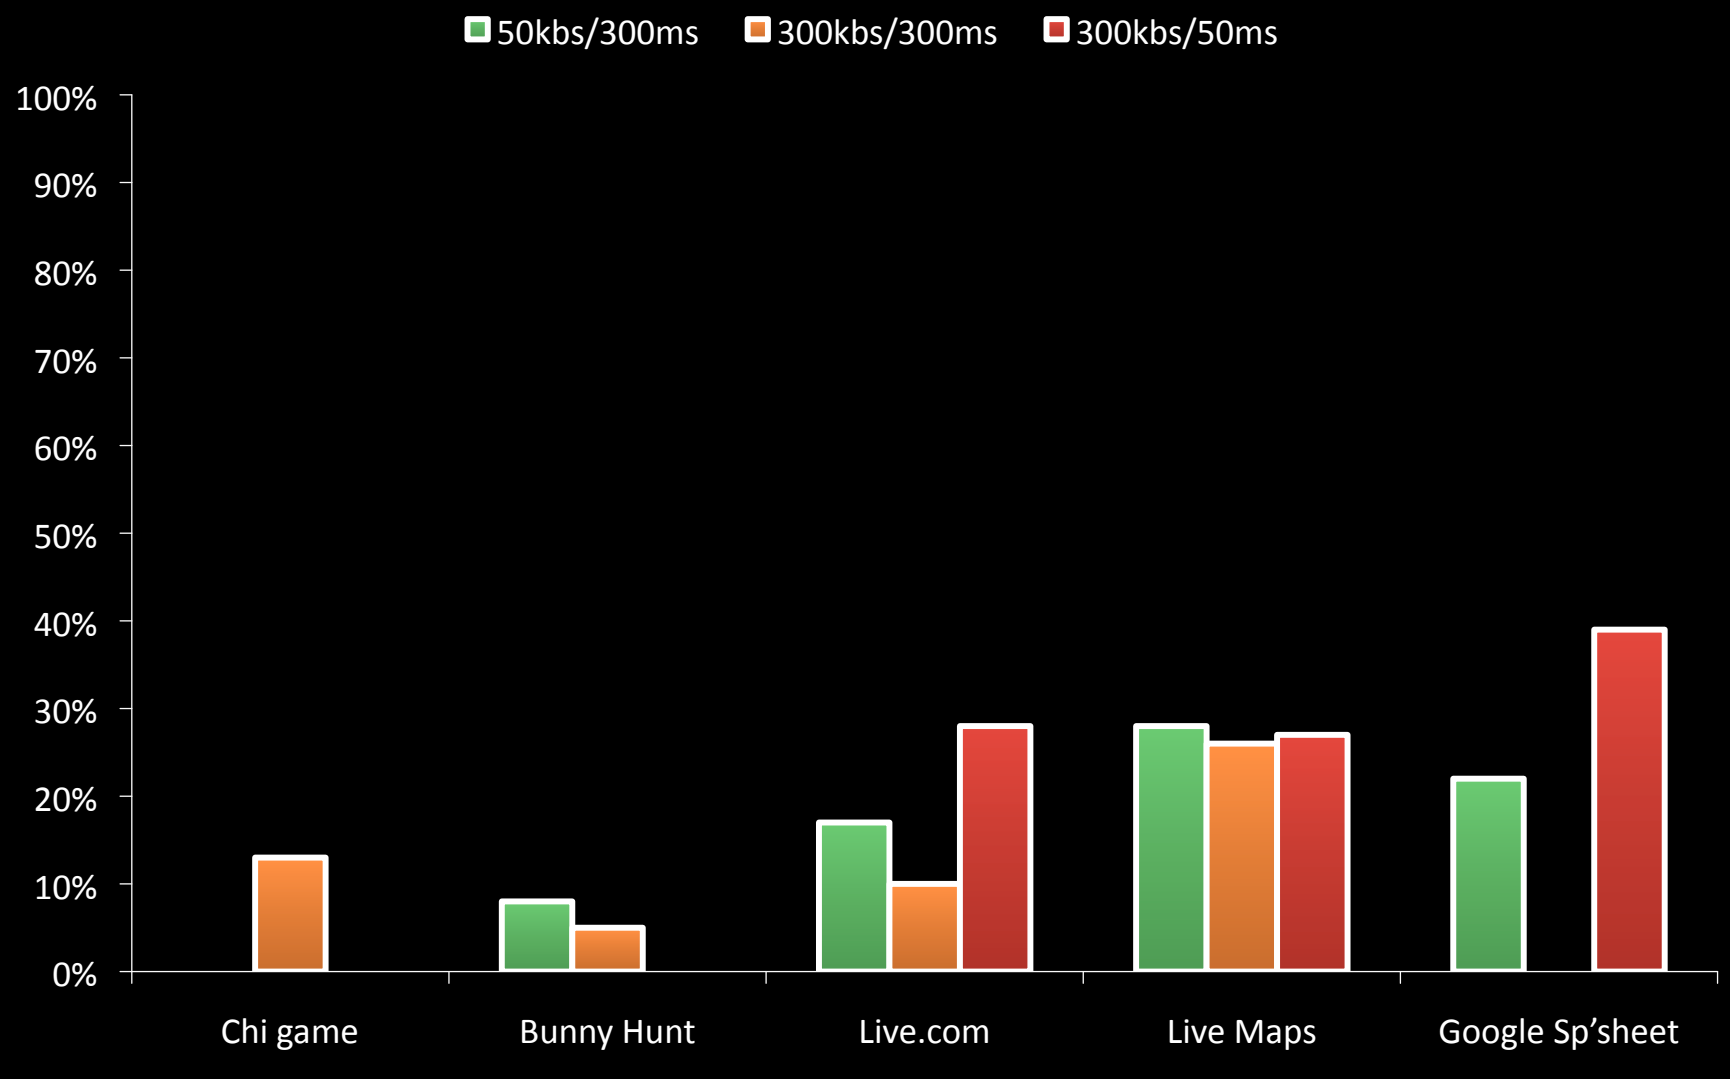

Doloto: Conclusions

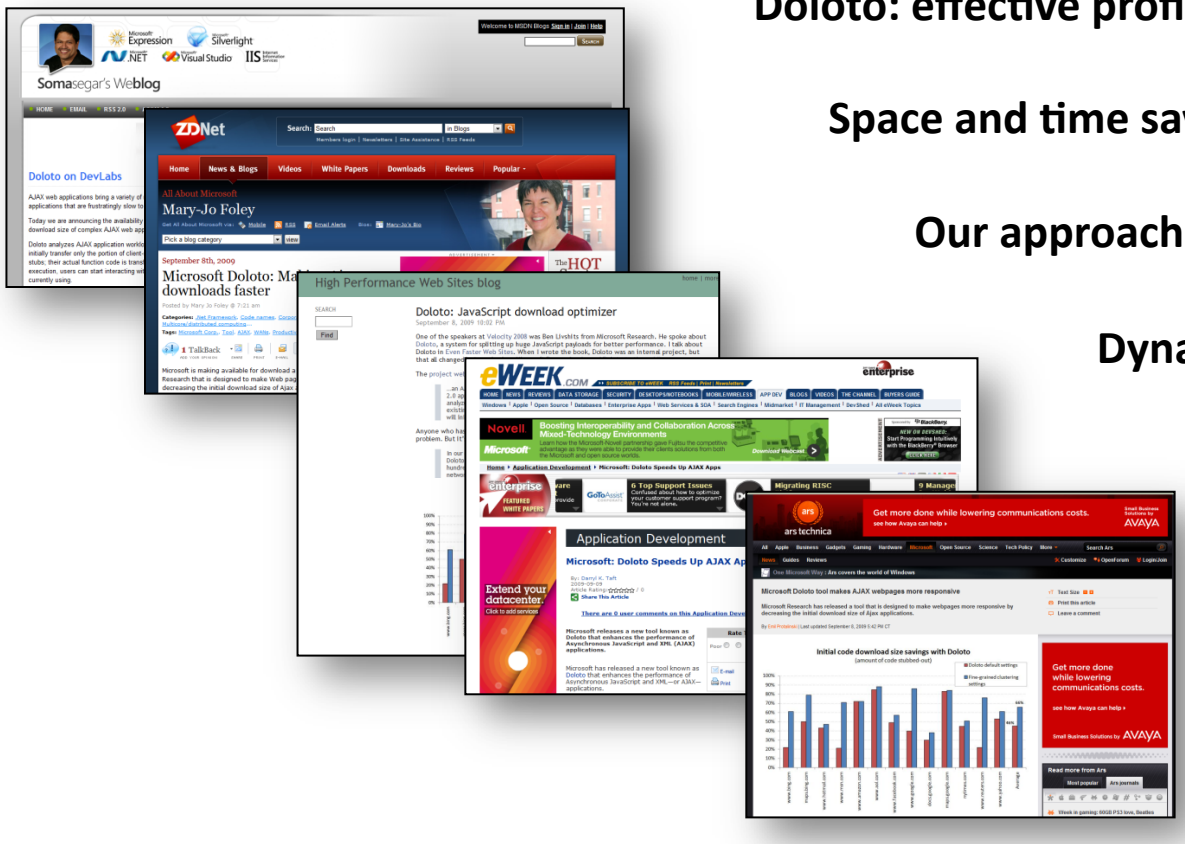

Doloto: effective profile-driven optimization

Space and time savings

Our approach is general: Silverlight

Dynamic code loading for future web apps

Press Coverage

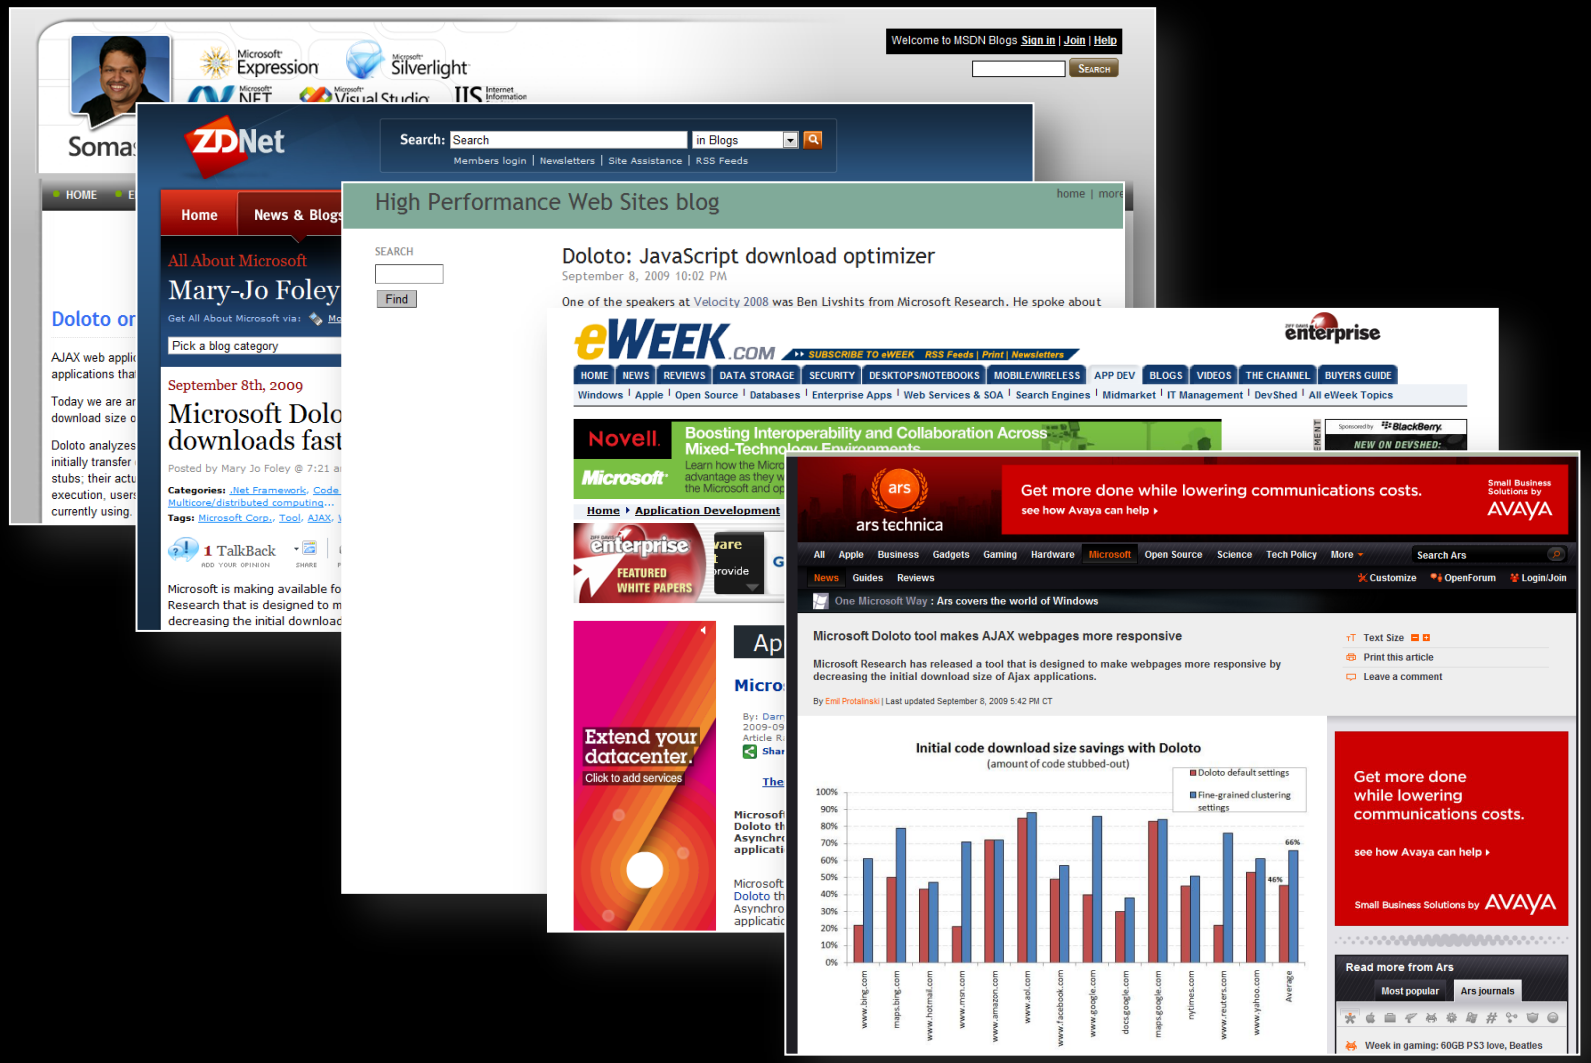

Outline of the Talk

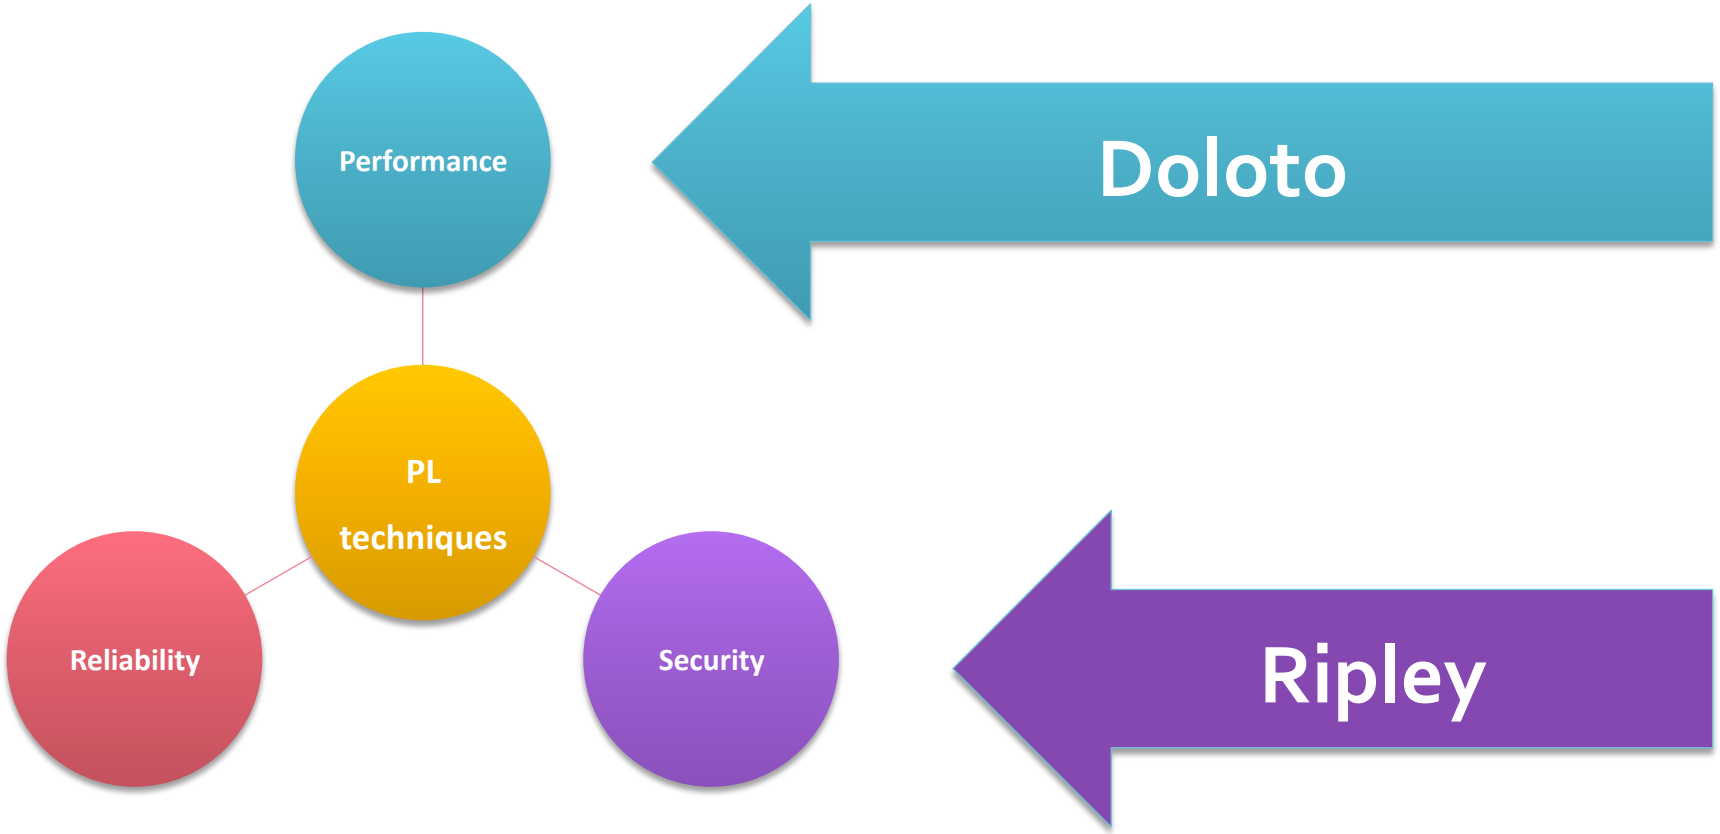

AUTOMATICALLY SECURING WEB 2.0 APPLICATIONS

THROUGH REPLICATED EXECUTION

Web 1.0 → Web 2.0

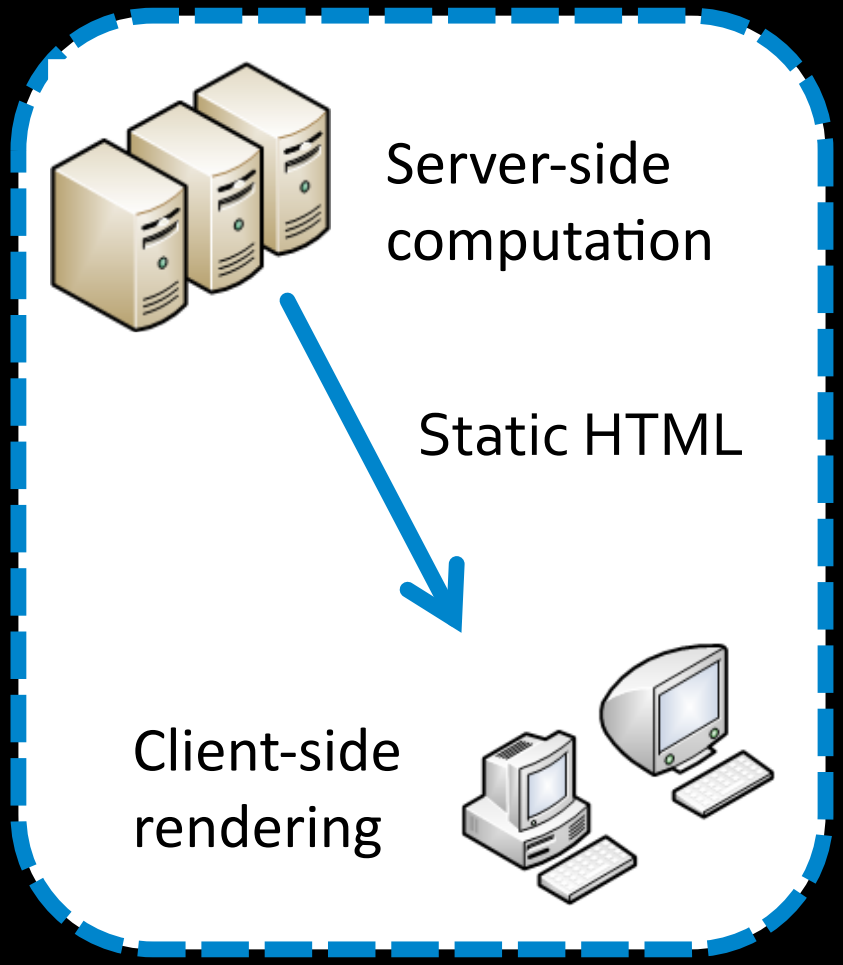
Server-side
computation
Static HTML
Client-side
rendering

How Do We Program This Mess?

- Currently:
  - J2EE + JavaScript?
  - PHP + Flash?
  - ASP.NET + Silverlight?
- One alternative:
  - Distributing compilers
  - Volta, GWT, Hops, Links

The Volta Distributing Compiler Illustrated

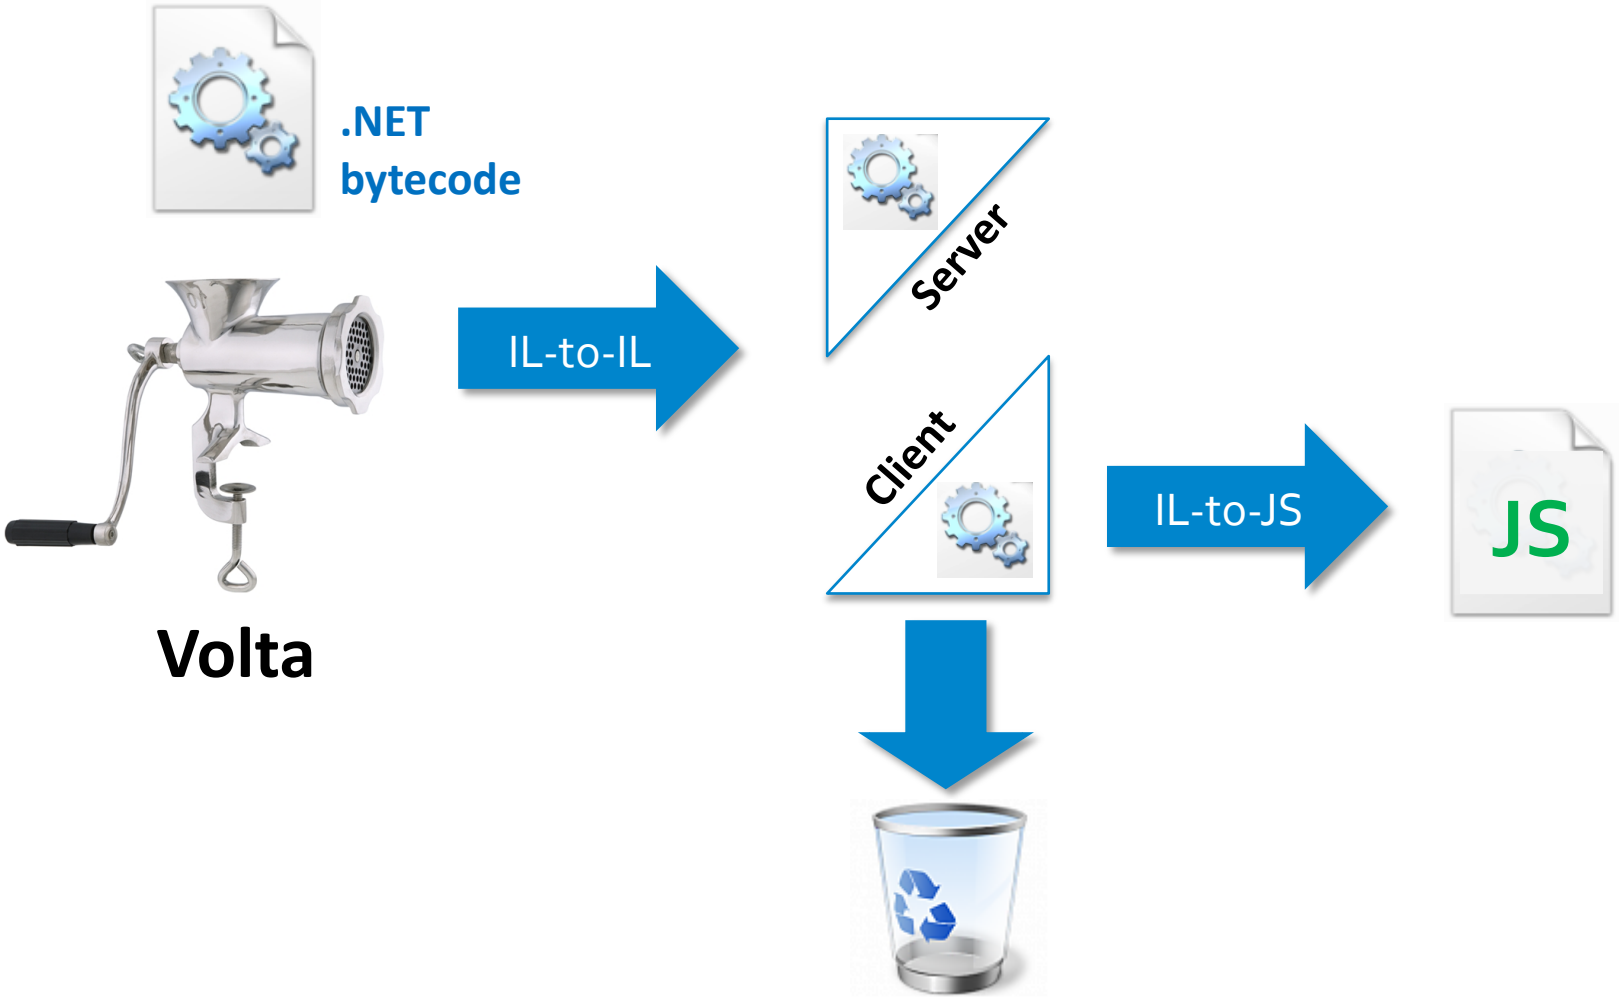

The Volta Distributing Compiler: Deployment

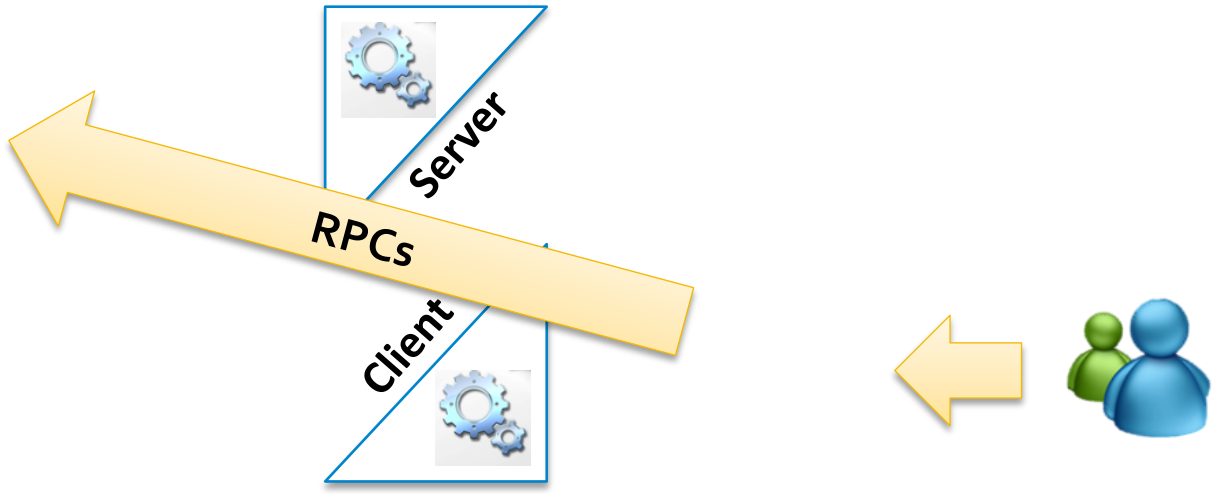

Server

Client

Web Developer's Mantra

Thou shall not trust the client

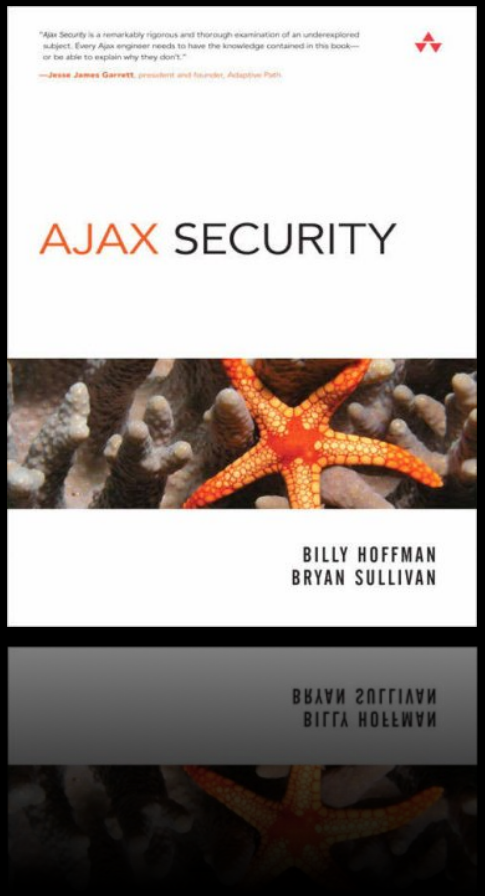

- No data integrity
- No code integrity

AJAX-based Shopping Cart

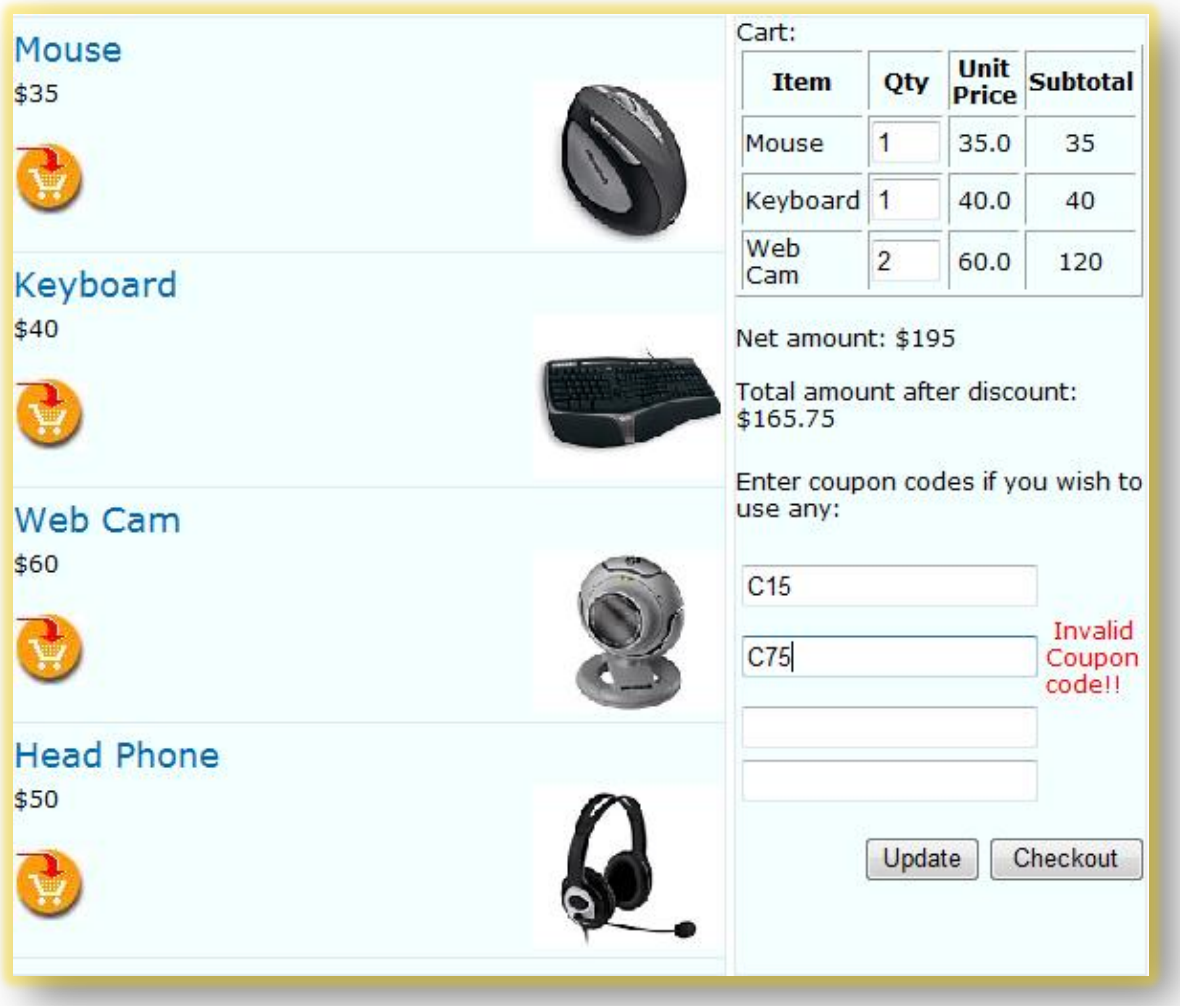

Security vs. Performance

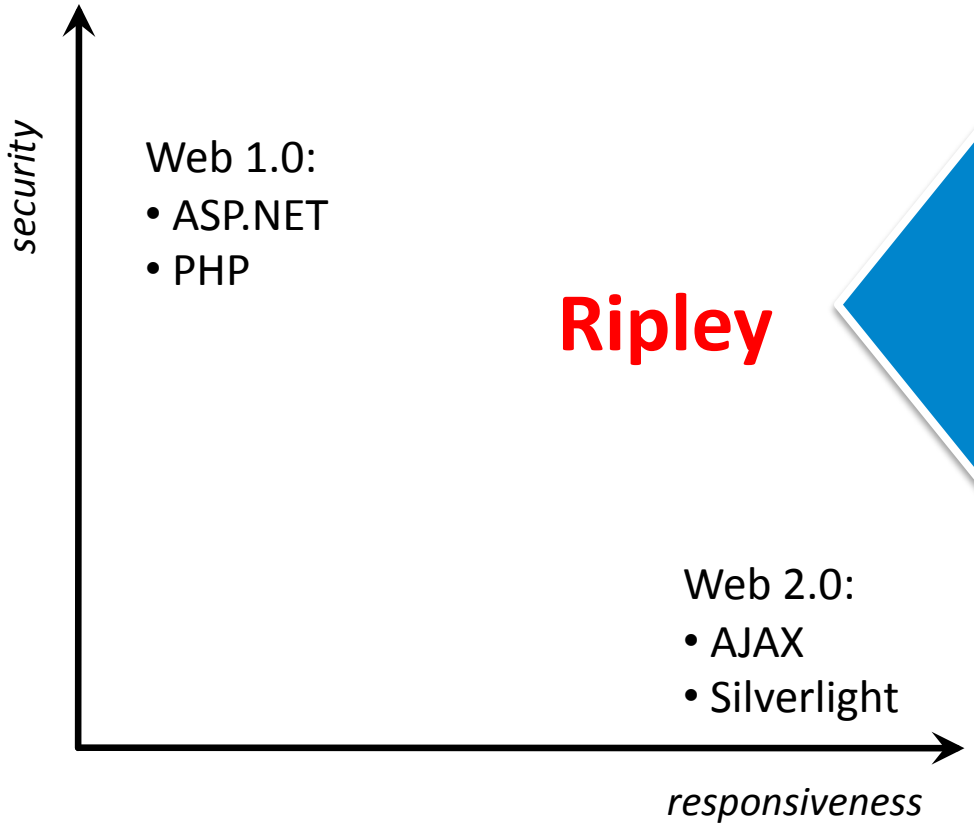

With Ripley, placing computation on the client *does not reduce* computational integrity

Ripley Architecture

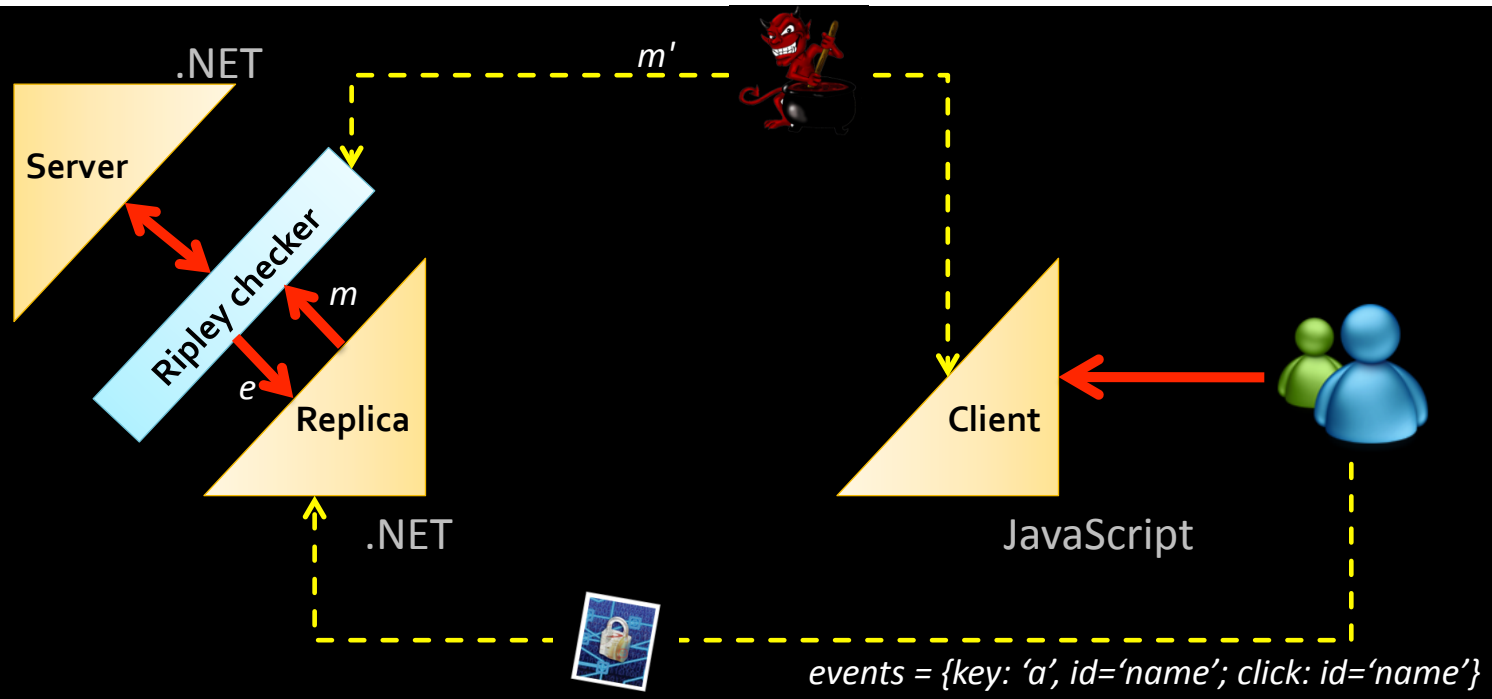

1. Keep a replica of the client code
2. Capture user events & transmit to server for replay
3. Compare server and client results

Ripley Architecture

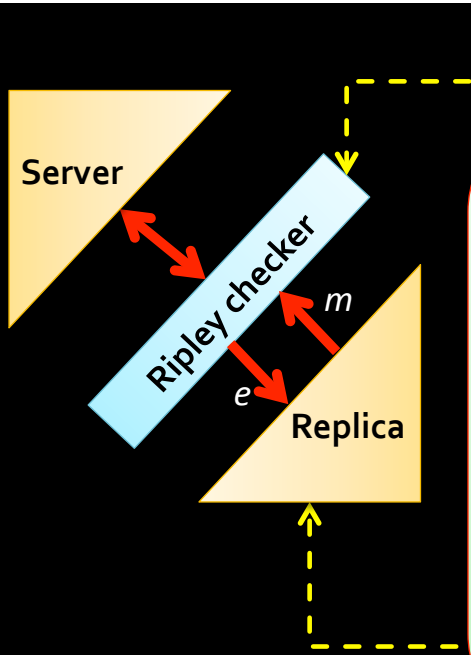

Client-side code instrumented
- Rewrite event handlers
- Capture "default" events
- Buffer events for performance

```
button.onClick =
  function buttonHandler(e) {
    ripleyEnqueue(e);          // inserted by rewriting
    var obj = eventTrigger(e);
    var notify = document.getElementById &&
                 document.getElementById('notify');
    notify.value = 'You clicked on ' + obj.value;
    return true;
  };
```

1. Keep a replica of the client code
2. Capture user events & transmit to server for replay
3. Compare server and client results

Ripley Architecture

- Run replica in a Ripley emulator
- Run in .NET, not in JavaScript, 100x speed increase

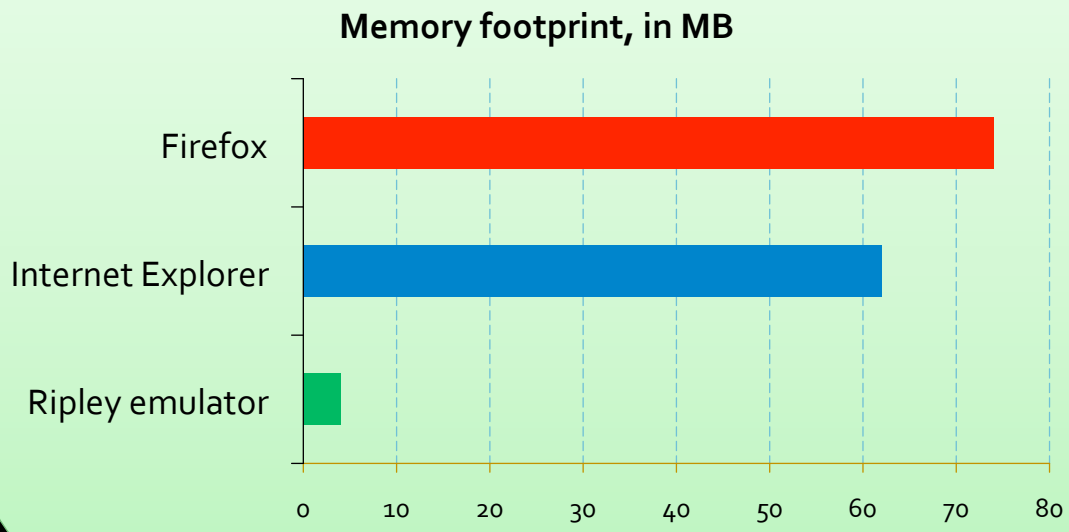

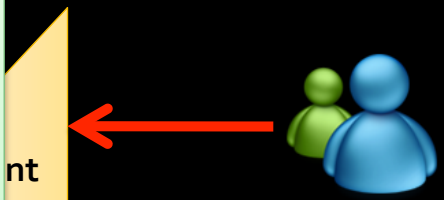

…: 'a', id='name'; click: id='name'}

1. Keep a replica of the client code
2. Capture user events & transmit to server for replay
3. Compare server and client results

Ripley Applications

- ✓ Shopping cart
- ✓ Sudoku
- ✓ Blog
- ✓ Speed typing
- ✓ Online Quiz
- ✓ Distributed online ga

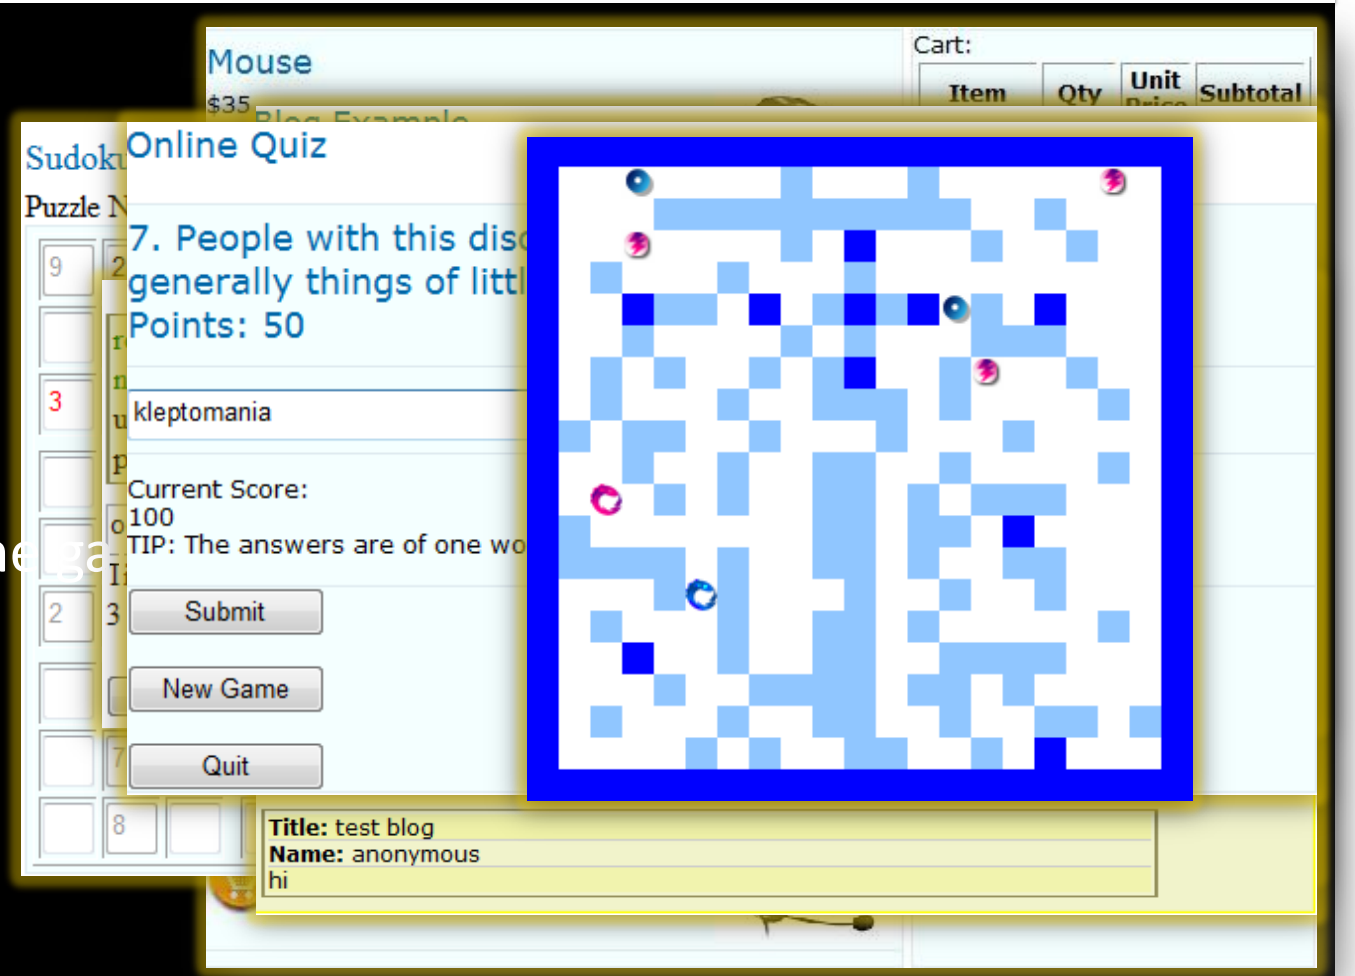

Performance Overhead Summary

Network:

- 2-3 bytes per user event (key press, mouse, etc.)
- Event stream compresses extremely well

CPU:

- Client: Several *ms* of overhead added for event capture
- Server: Several *ms* for server-side checking

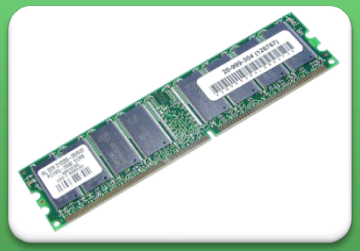

Memory:

- About 1 MB per connected client
- Can scale to 1,000's of clients per server

Ripley: Vision for the Future

- Security by construction

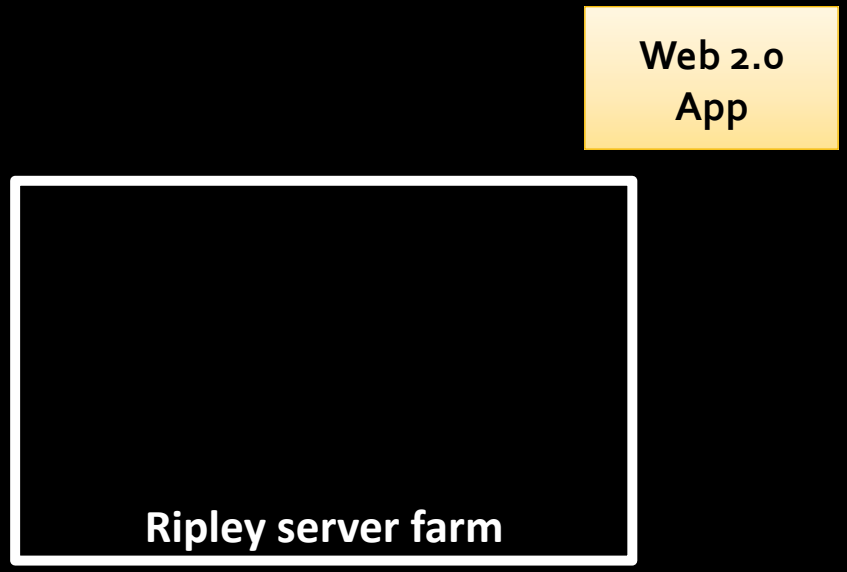

For More…

Doloto|Ripley MSR _

Summary

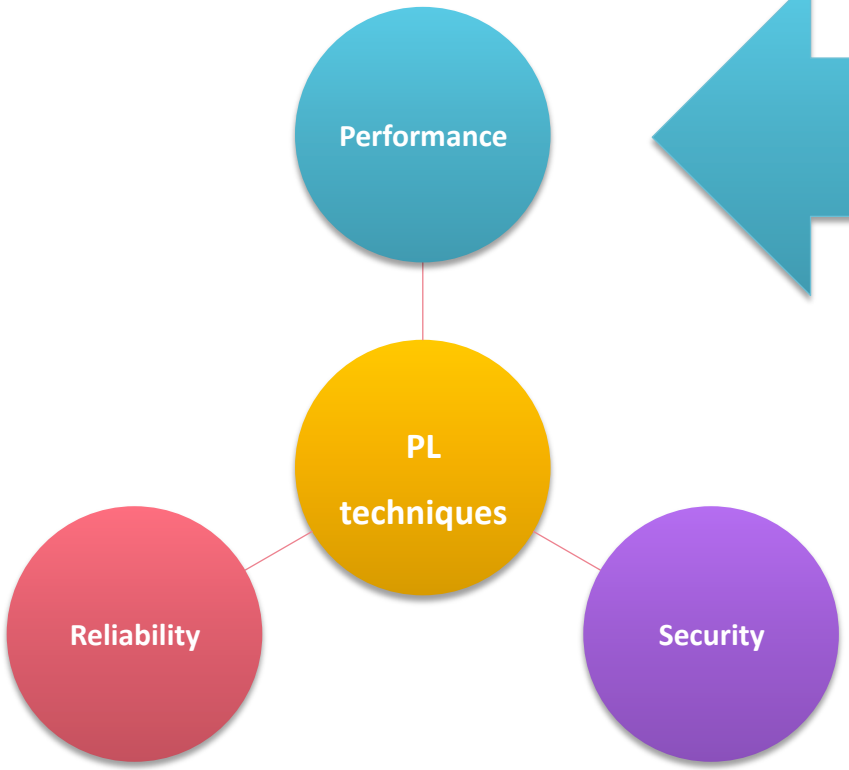

Doloto

[FSE 2008]

Ripley

[CCS 2009]

Call to Arms

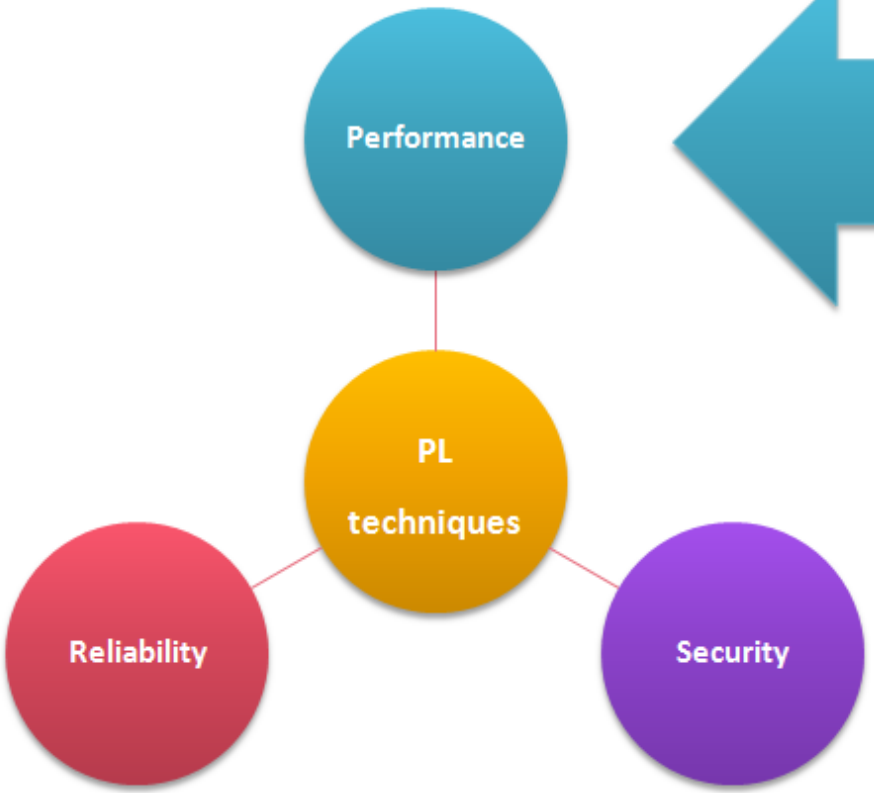

Doloto

- Web applications are here to stay

[FSE 2008]

- Exciting new opportunities for research

Ripley

[CCS 2009]