



# Thinking in a Highly Concurrent, Mostly-functional Language

Chicago Erlang

Chicago, September 22<sup>nd</sup> 2014

**Francesco Cesarini**

**Founder & Technical Director**

**@francescoC**

**francesco@erlang-solutions.com**



Erlang Training and Consulting Ltd

# Thinking in a Highly Concurrent, Mostly-functional Language

QCON London, March 12<sup>th</sup>, 2009

**Francesco Cesarini**

**[francesco@erlang-consulting.com](mailto:francesco@erlang-consulting.com)**

```
counter_loop(Count) ->  
  receive  
    increment ->  
      counter_loop(Count + 1);  
  {count, To} ->  
    To ! {count, Count},  
    counter_loop(Count)  
end.
```

# Erlang



**After you've opened the top of your head, reached in and turned your brain inside out, this starts to look like a natural way to count integers. And Erlang does require some fairly serious mental readjustment.**

**However... having spent some time playing with this, I tell you...**

Tim Bray, Director of Web Technologies - Sun Microsystems

... If somebody came to me and wanted to pay me a lot of money to build a large scale message handling system that really had to be up all the time, could never afford to go down for years at the time, I would unhesitatingly choose Erlang to build it in.

Tim Bray, Director of Web Technologies - Sun Microsystems

# Syntax

# Concurrency

# Erlang Highlights: Concurrency

## Creating a new process using spawn

```
-module(ex3).  
-export([activity/3]).  
  
activity(Name,Pos,Size) ->  
.....
```





```
Pid = spawn(ex3,activity,[Joe,75,1024])
```



# Erlang Highlights: Concurrency

*Processes communicate by asynchronous message passing*



```
Pid ! {data,12,13}
```

```
receive  
    {start} -> .....  
    {stop} -> .....  
    {data,X,Y} -> .....  
end
```

# Products: AXD301 Switch - 1996

A Telephony-Class, scalable (10 - 160 GBps) ATM switch

Designed from scratch in less than 3 years

## AXD 301 Success factors:

- Competent organisation and people
- Efficient process
- Excellent technology (e.g. Erlang/OTP)



# Products: AXD301 Switch - 1996

## Erlang: ca 1.5 million lines of code

- Nearly all the complex control logic
- Operation & Maintenance
- Web server and runtime HTML/JavaScript generation

## C/C++: ca 500k lines of code

- Third party software
- Low-level protocol drivers
- Device drivers

## Java: ca 13k lines of code

- Operator GUI applets



# Concurrency Modeling




Model for the natural concurrency in your problem

In the old days, processes were a critical resource

- Rationing processes led to complex and unmanageable code

Nowadays, processes are very cheap: if you need a process - create one!


Example: AXD301 process model




1<sup>st</sup> prototype:  
6 processes/call




2 processes/call



1 process/all calls



2 processes/  
call transaction

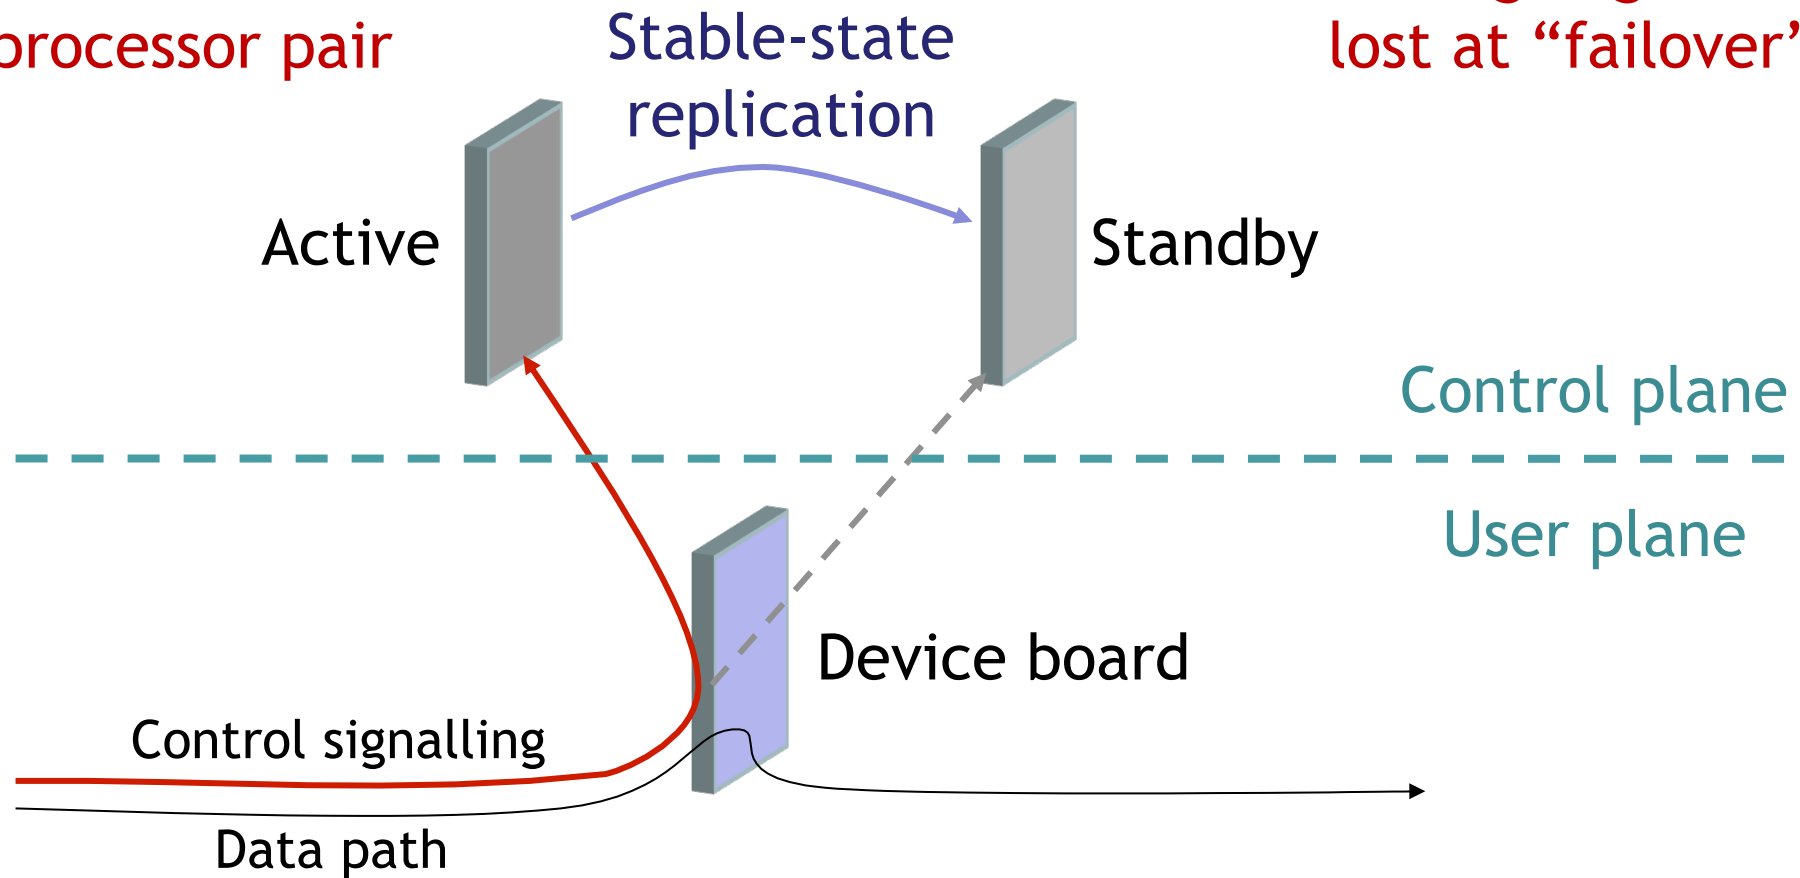


4-5 processes/  
call transaction

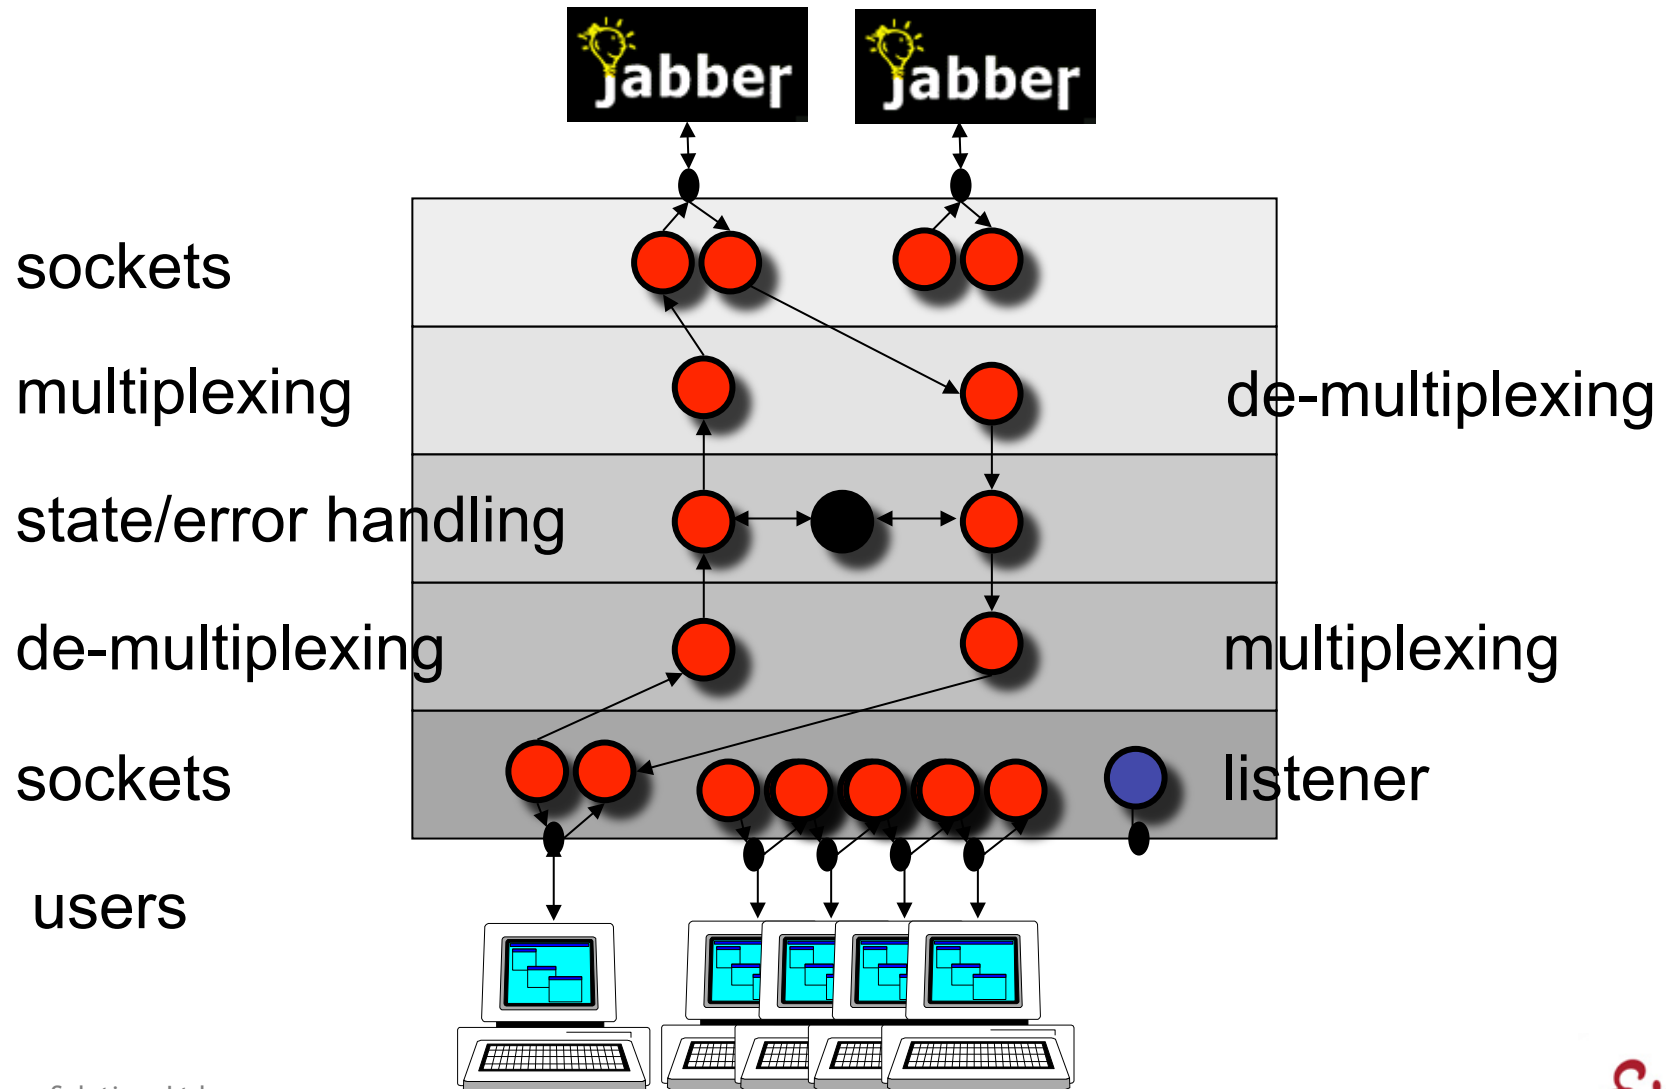
# 1+1 Redundancy - Good ol' Telecoms

~ 35 000 calls  
per processor pair

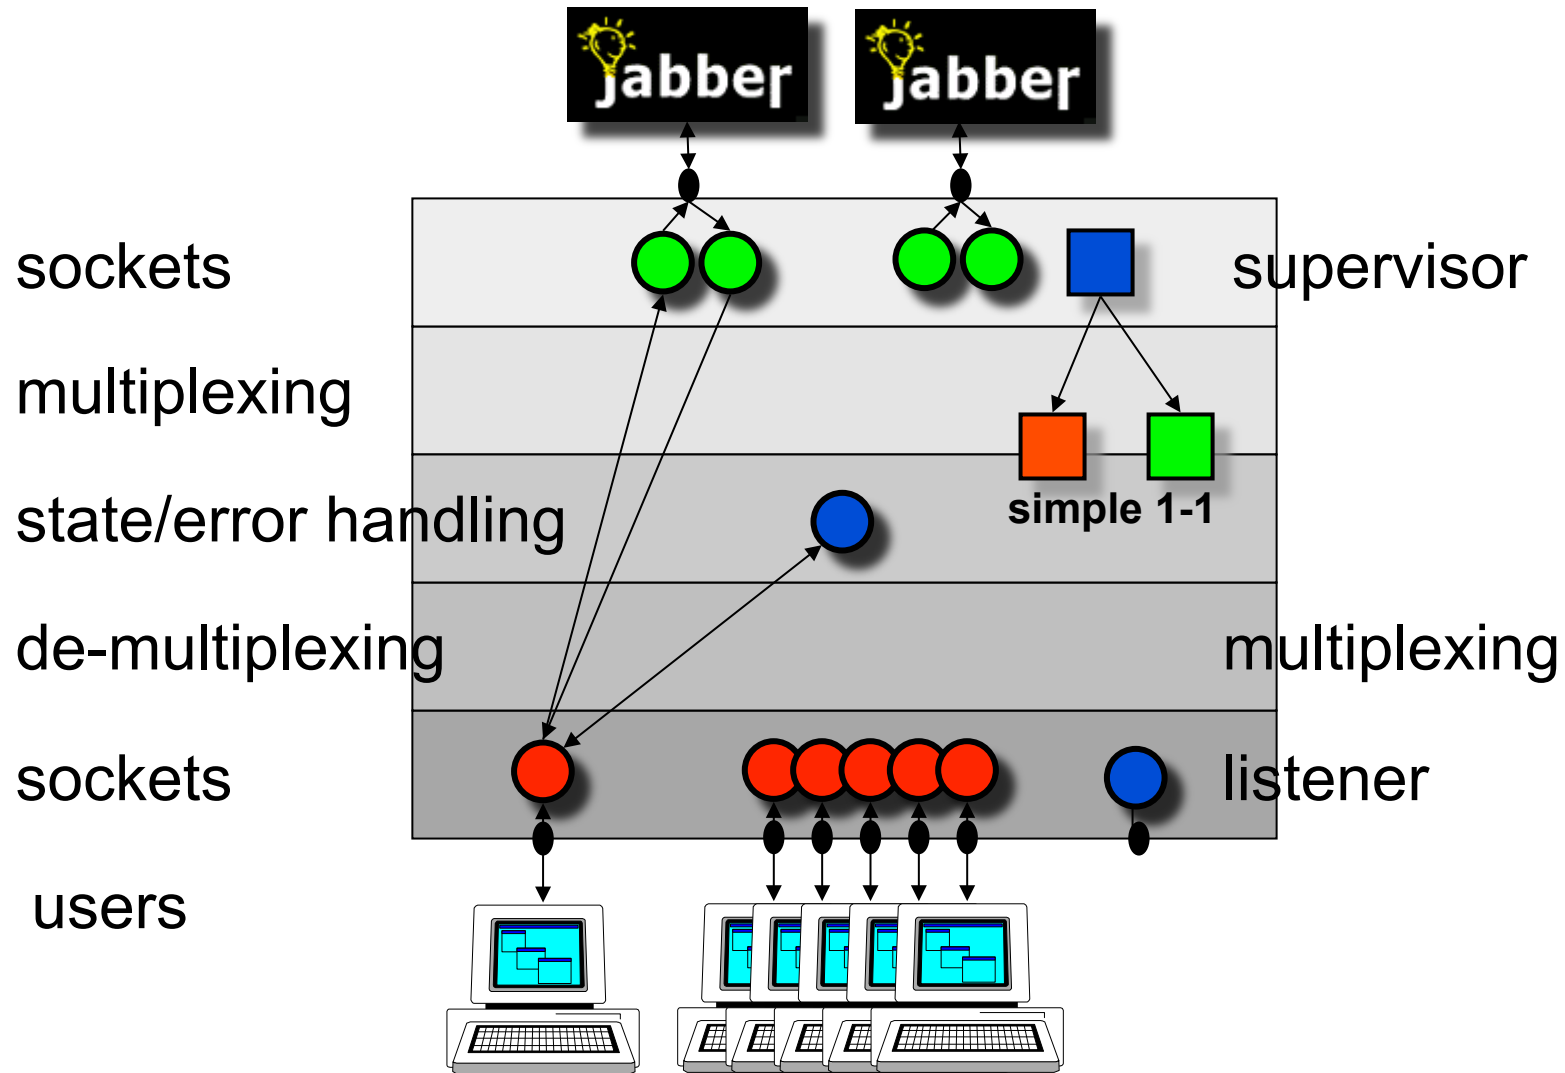
No ongoing sessions  
lost at "failover"



# First IM Proxy Prototype - 2000



# First IM Proxy Prototype - 2000



# Products: EjabberD IM Server - 2002

A distributed XMPP server

Started as an Open Source Project by *Alexey Shchepin*

Commercially Supported by Process-One (Paris)

- 40% of the XMPP IM market
- Used as a transport layer
- Managed 30,000 users / node





# Products: EjabberD IM Server - 2002

A distributed XMPP server

Started as an Open Source Project by *Alexey Shchepin*

Commercially Supported by Process-One (Paris)

- 40% of the XMPP IM market
- Used as a transport layer
- 2008, Managed 30,000 users / node

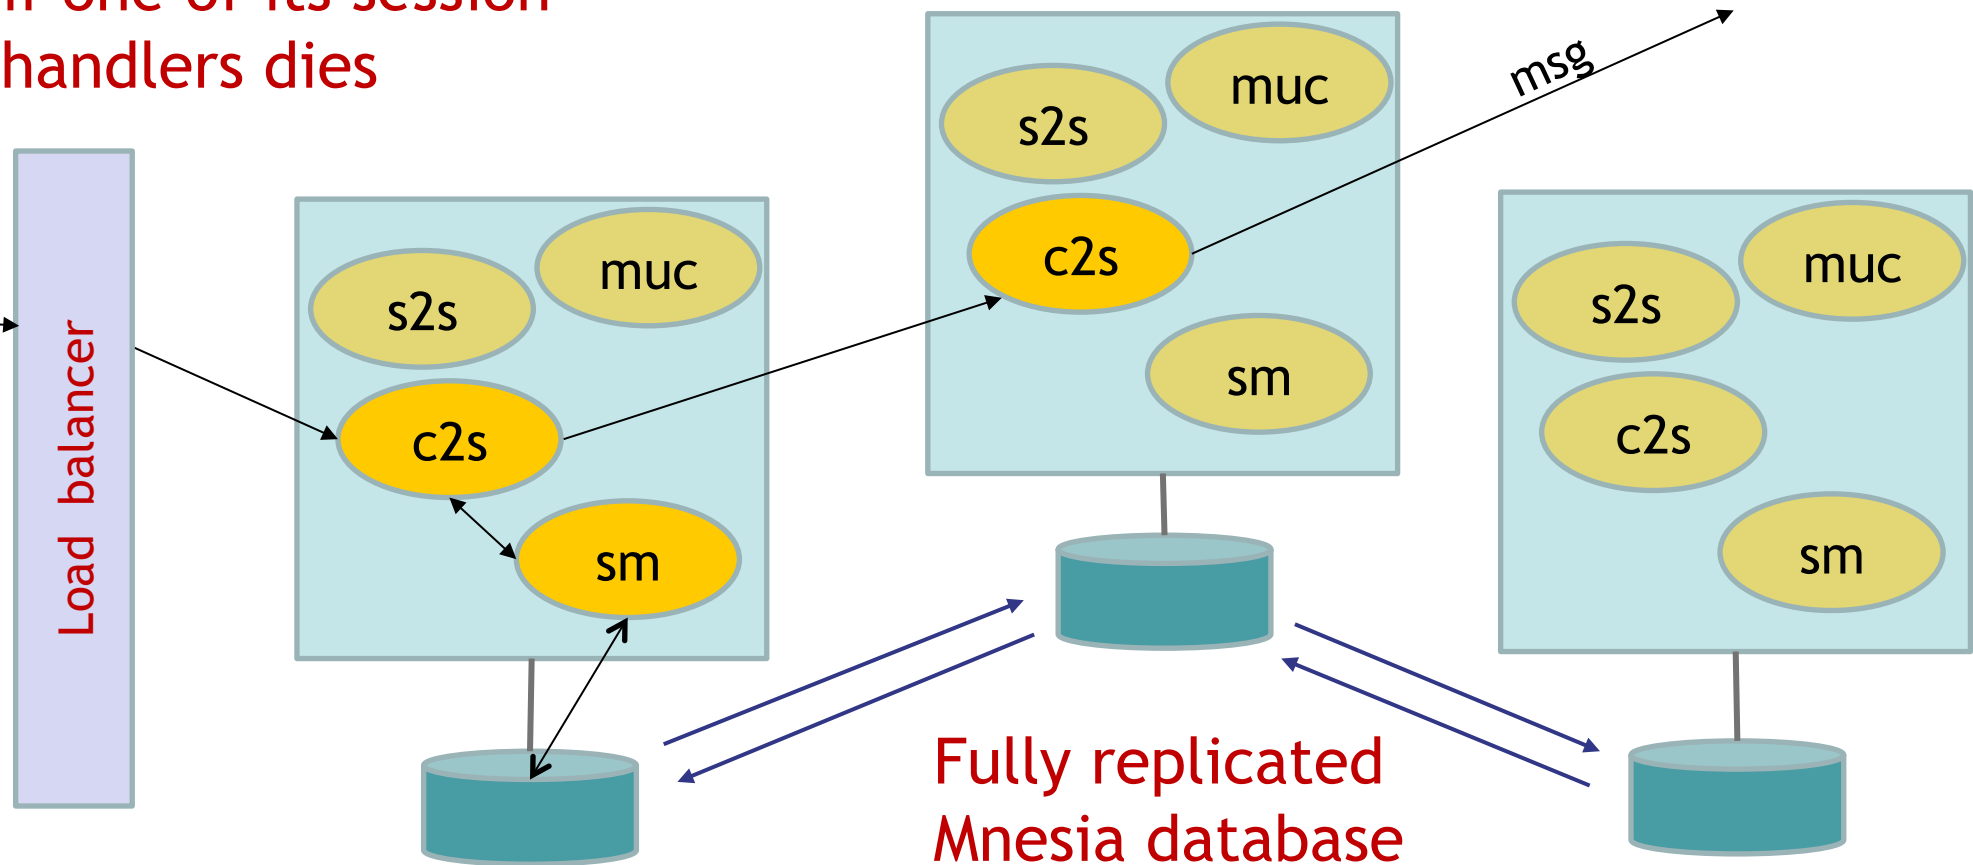
MongooseIM is a fork and rewrite

- Open Source, supported by Erlang Solutions
- Used for Messaging and Device Management
- 2014, managed 1 million users / node

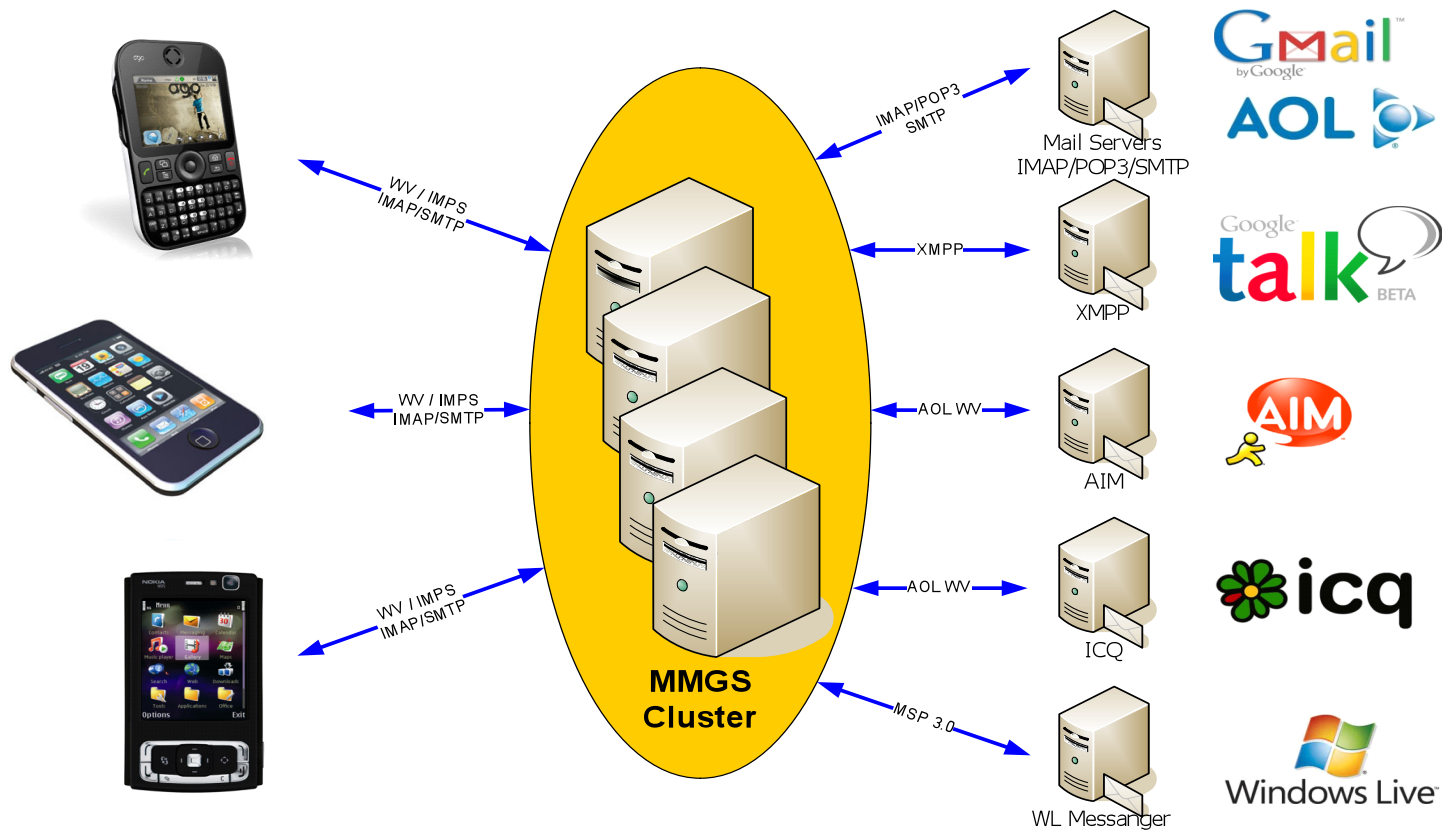


# Fully Replicated Cluster - Ejabberd 2002

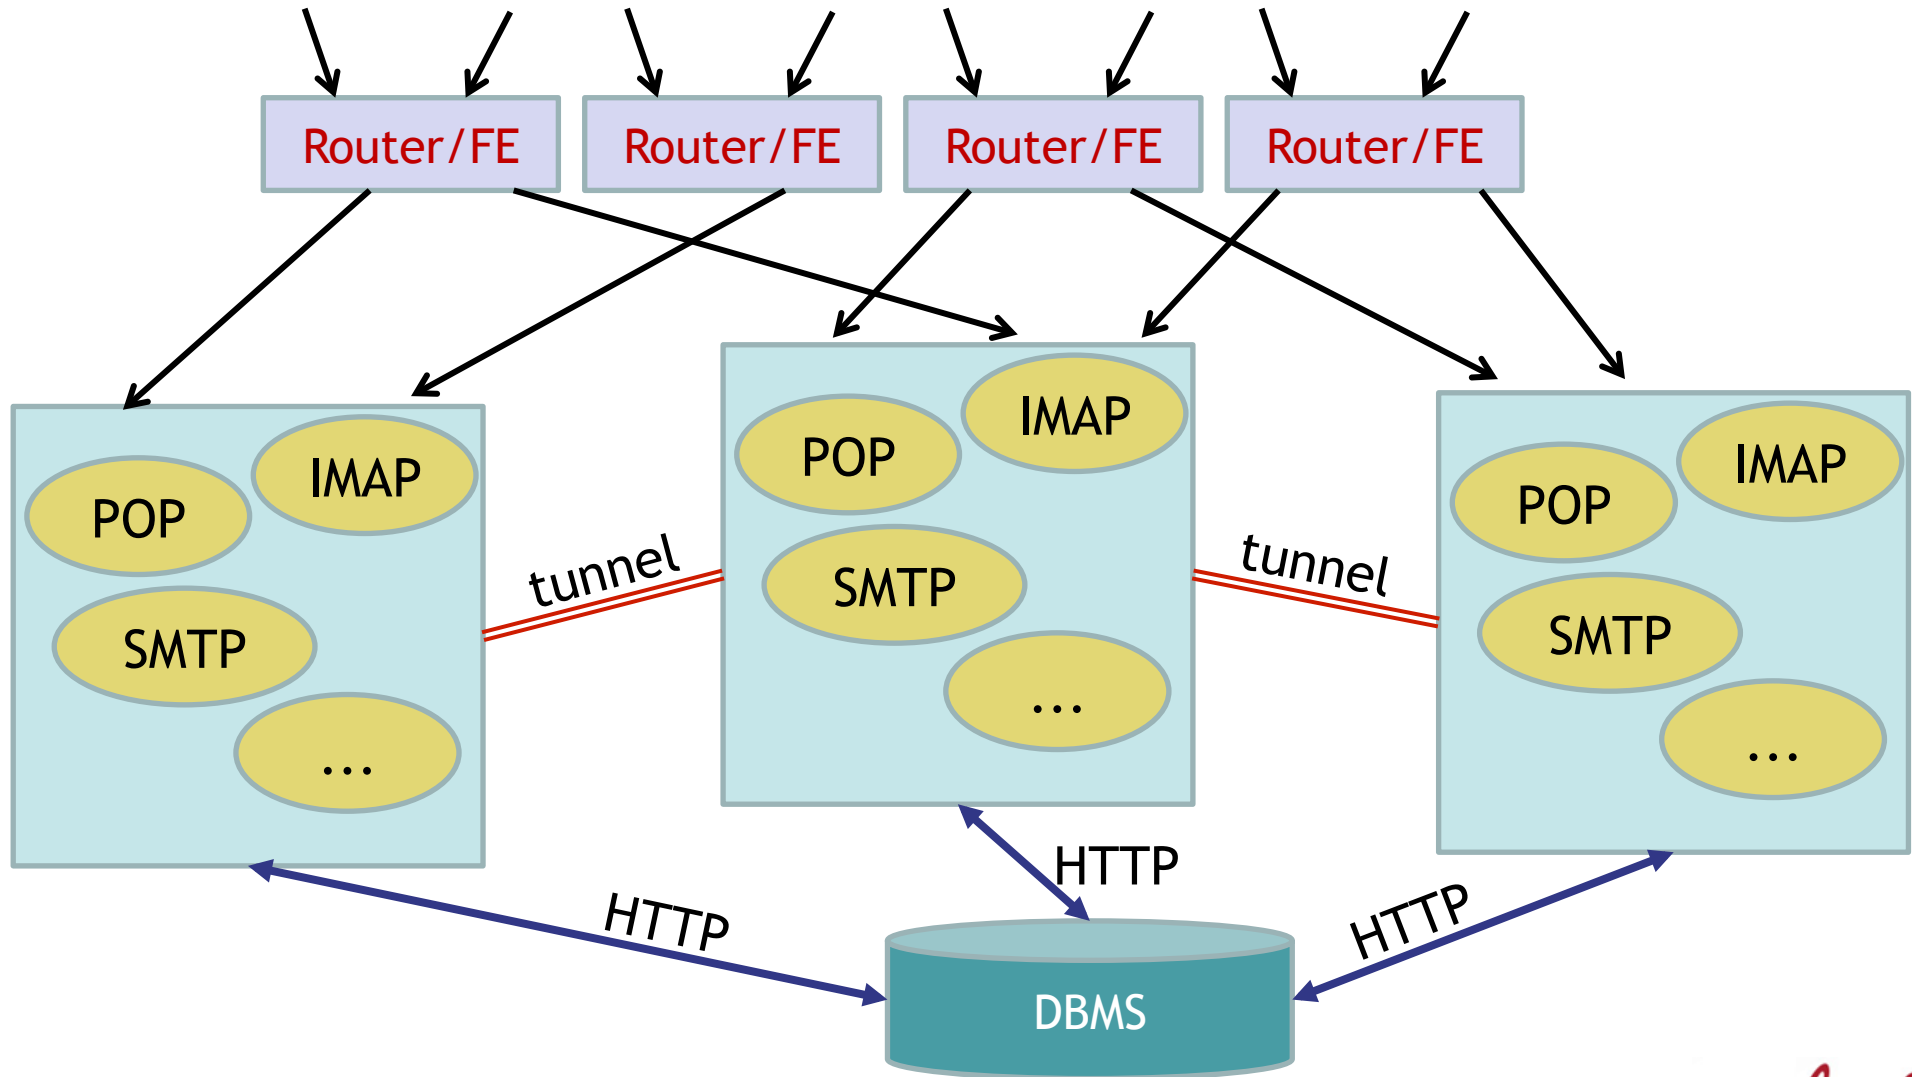
Client must re-connect  
if one of its session  
handlers dies



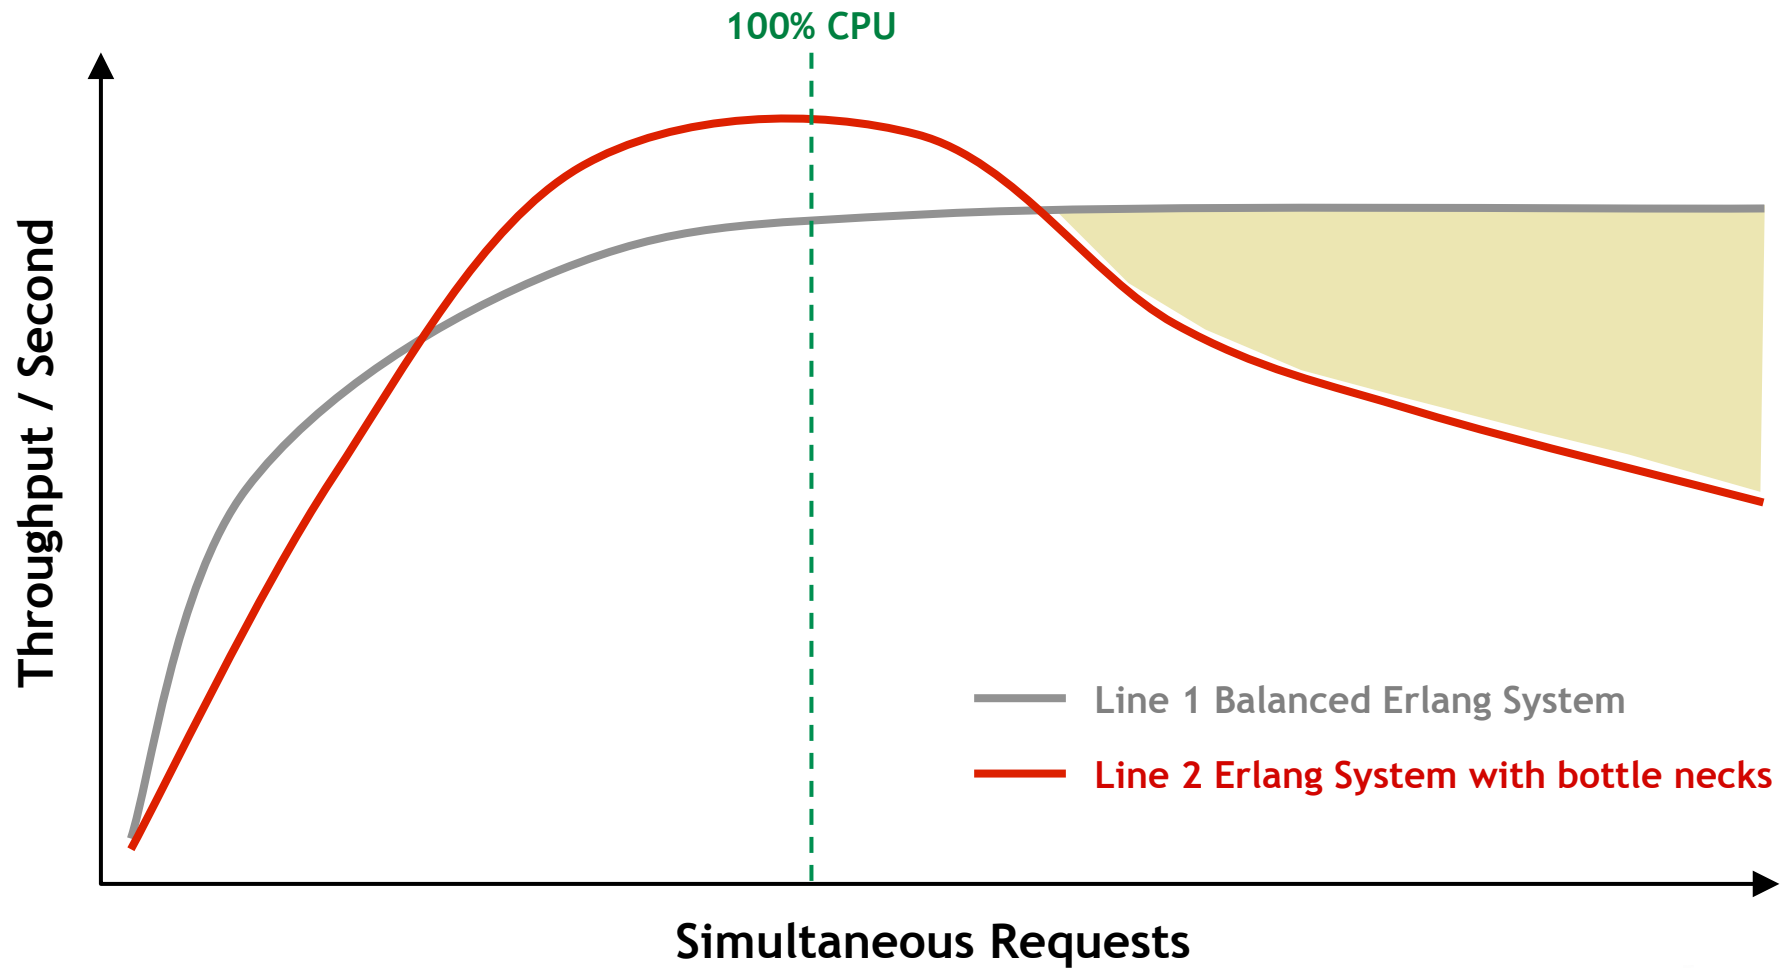
# Share-nothing Architecture - Messaging Gateway



# Share-nothing Architecture - Messaging Gateway




# Erlang Concurrency Under Stress - Pre-SMP



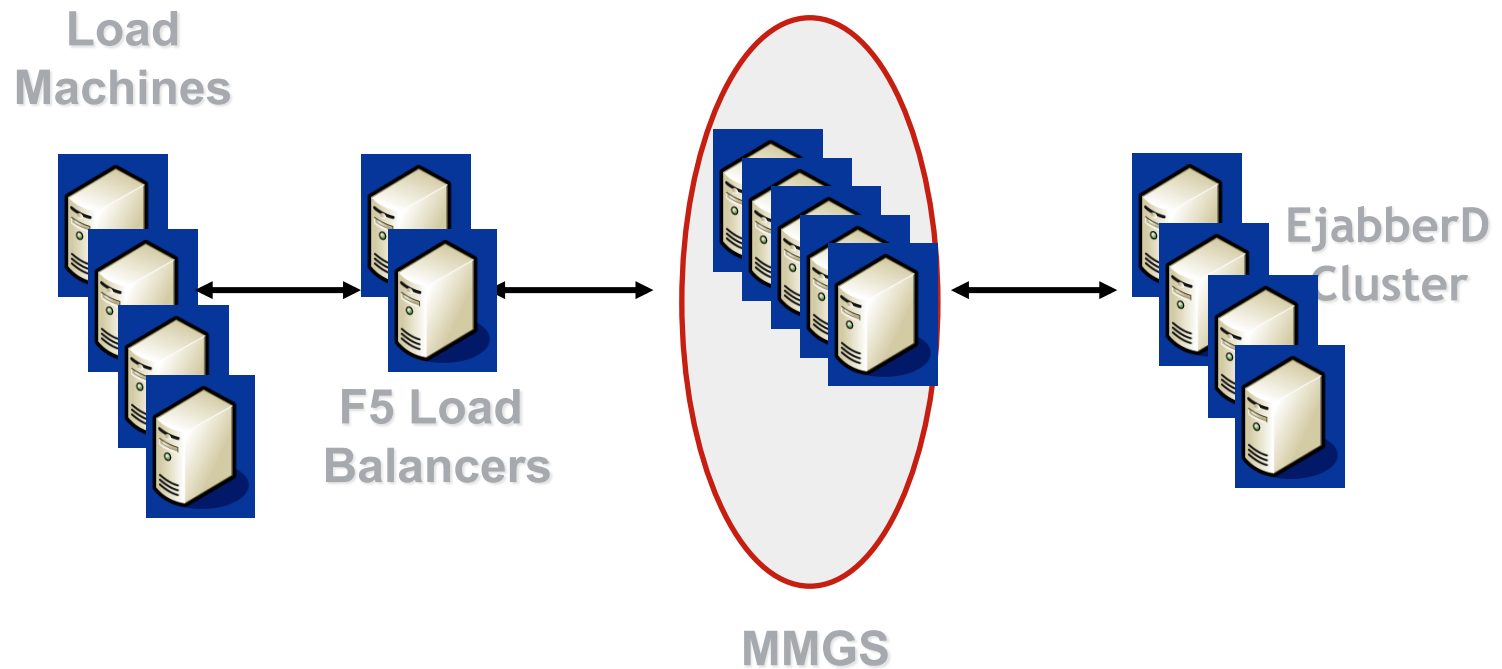
# Erlang Concurrency Under Stress - Pre-SMP

YAWS Throughput  
(KBytes/second)



Simultaneous Requests

# Erlang Concurrency Under Stress - Post-SMP





# Stress Tests With SMP

I/O Starvation

TCP/IP Congestion

Memory Spikes

Timeout Fine-tuning

OS Limitations

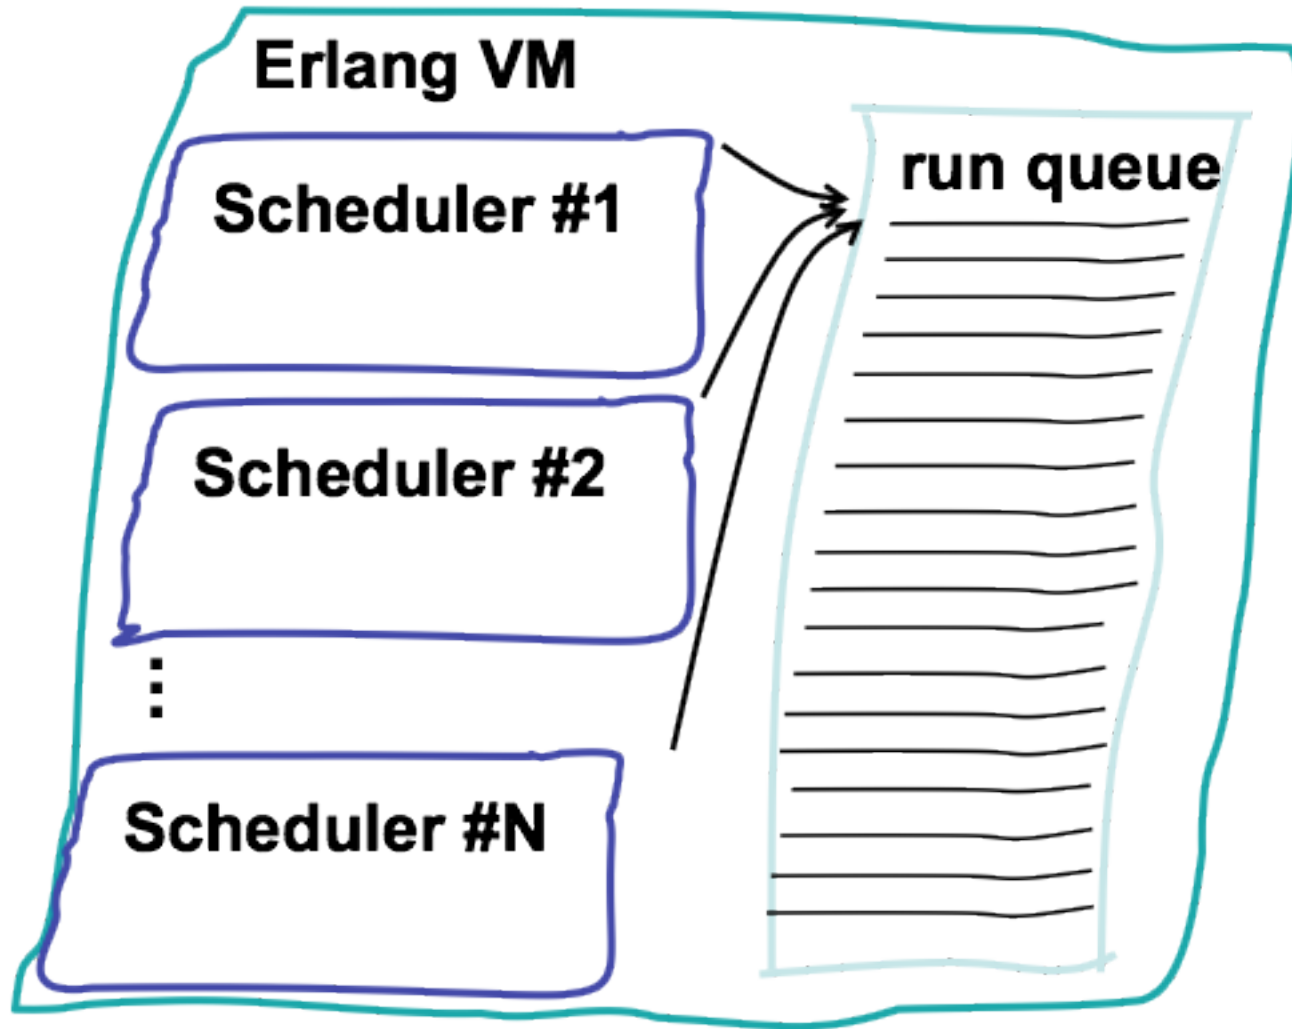
ERTS Configuration Flags

Shut down Audit Logs

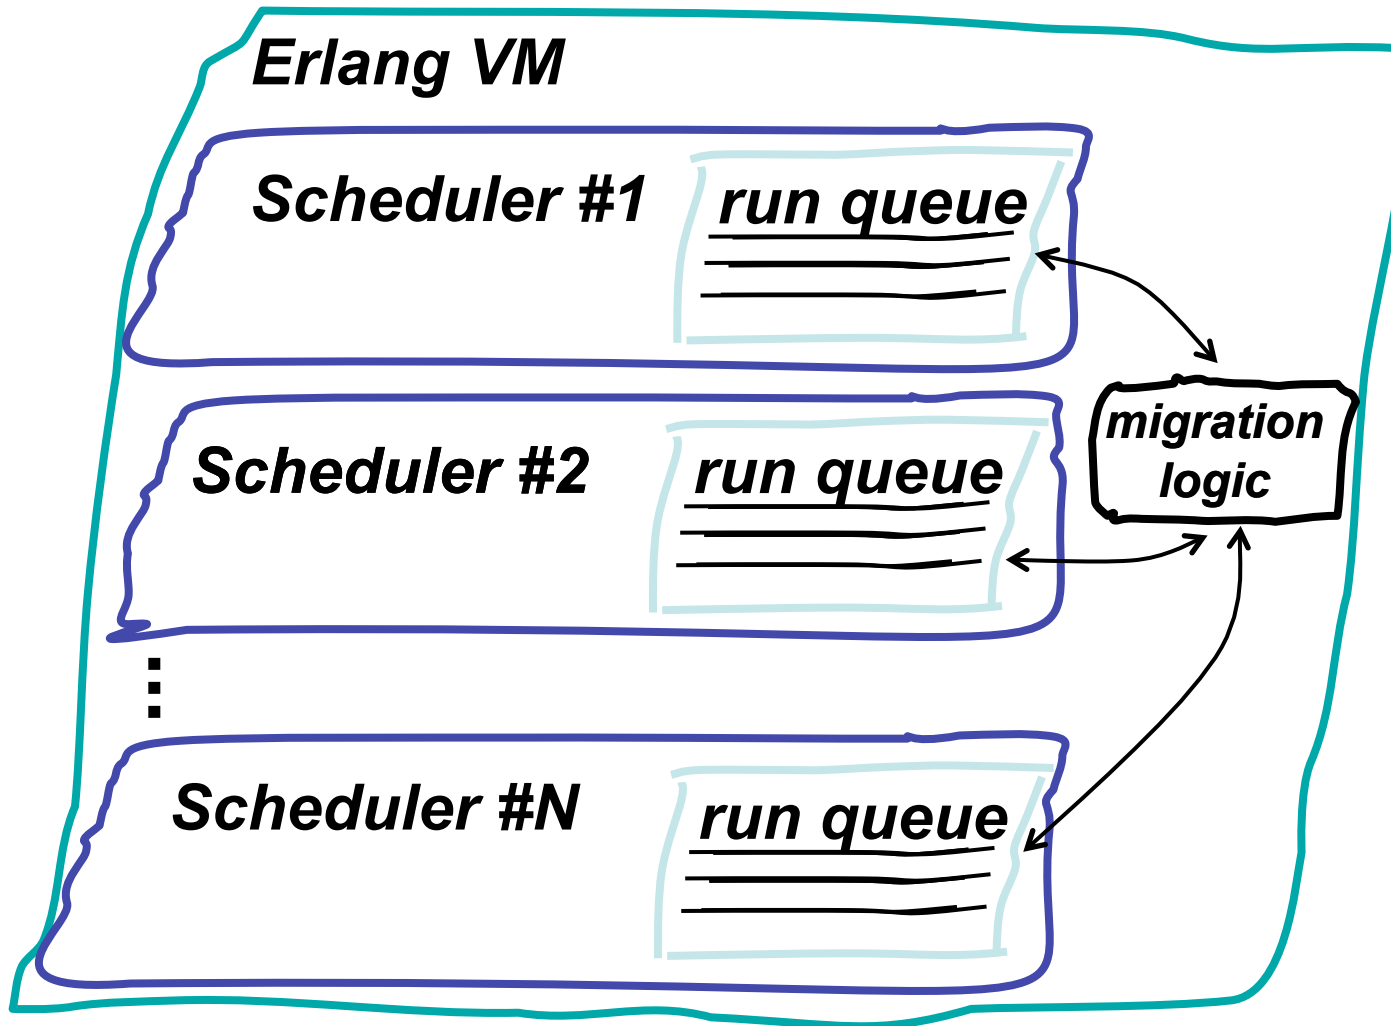




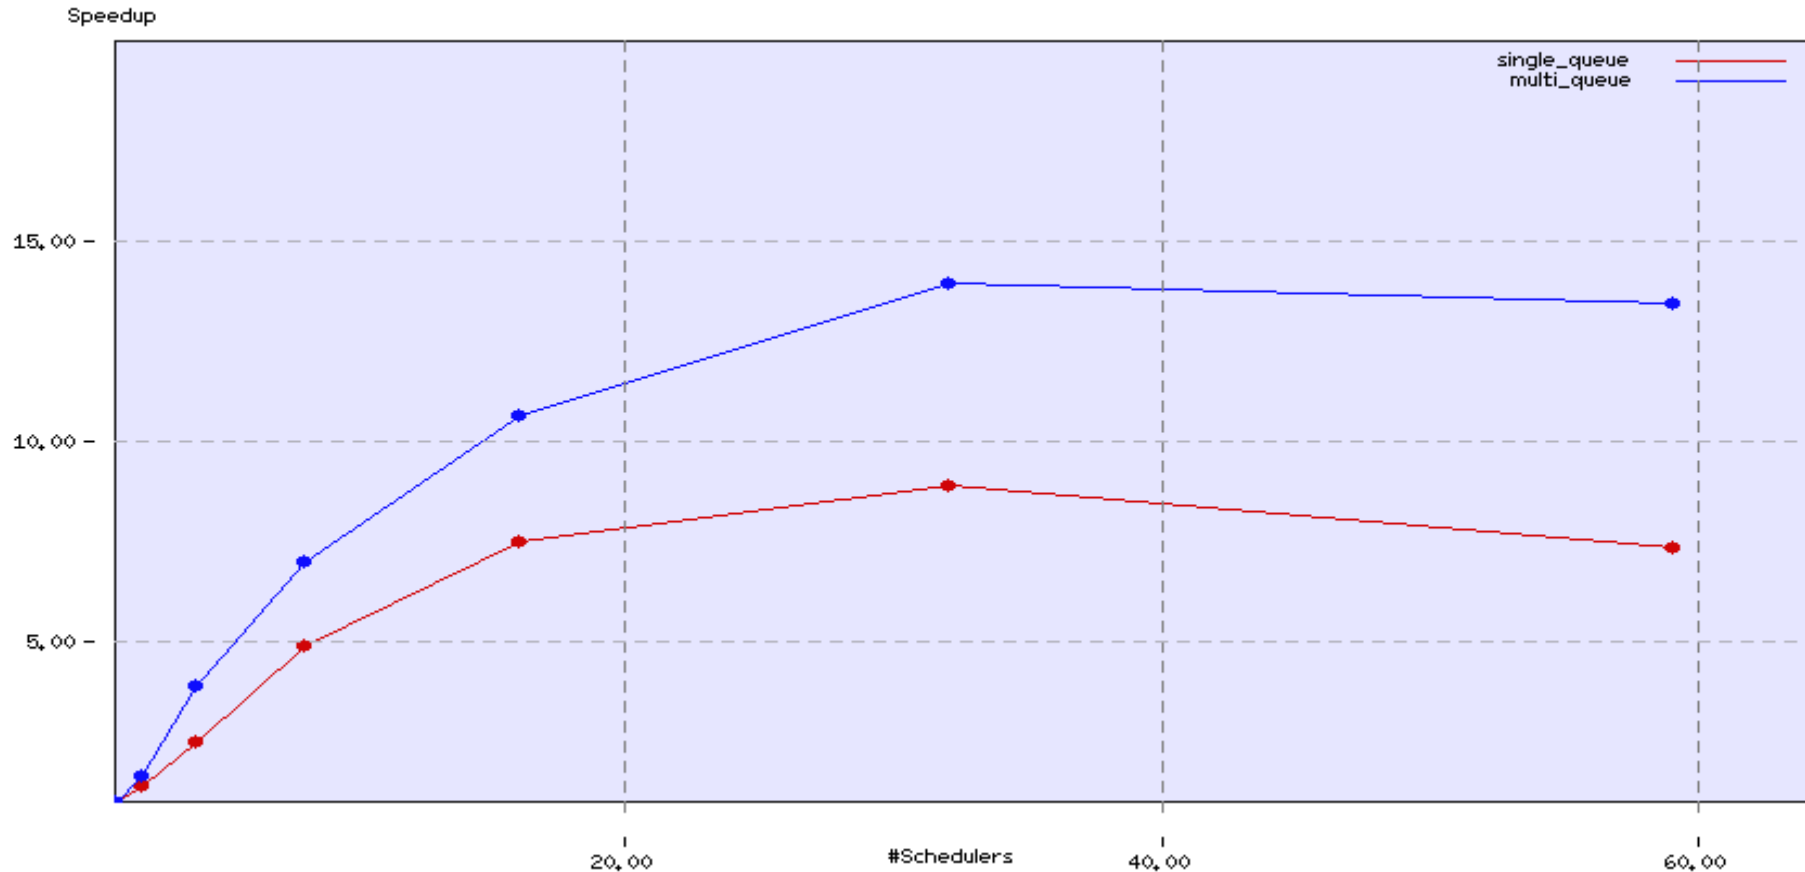
# SMP bottlenecks - pre 2008



# SMP bottlenecks - post 2008



# Big Bang Benchmark - post 2008




**Red:** Single Queue, **Blue:** Multiple Run Queue on a **Tileria TilePro64 (64 cores)**

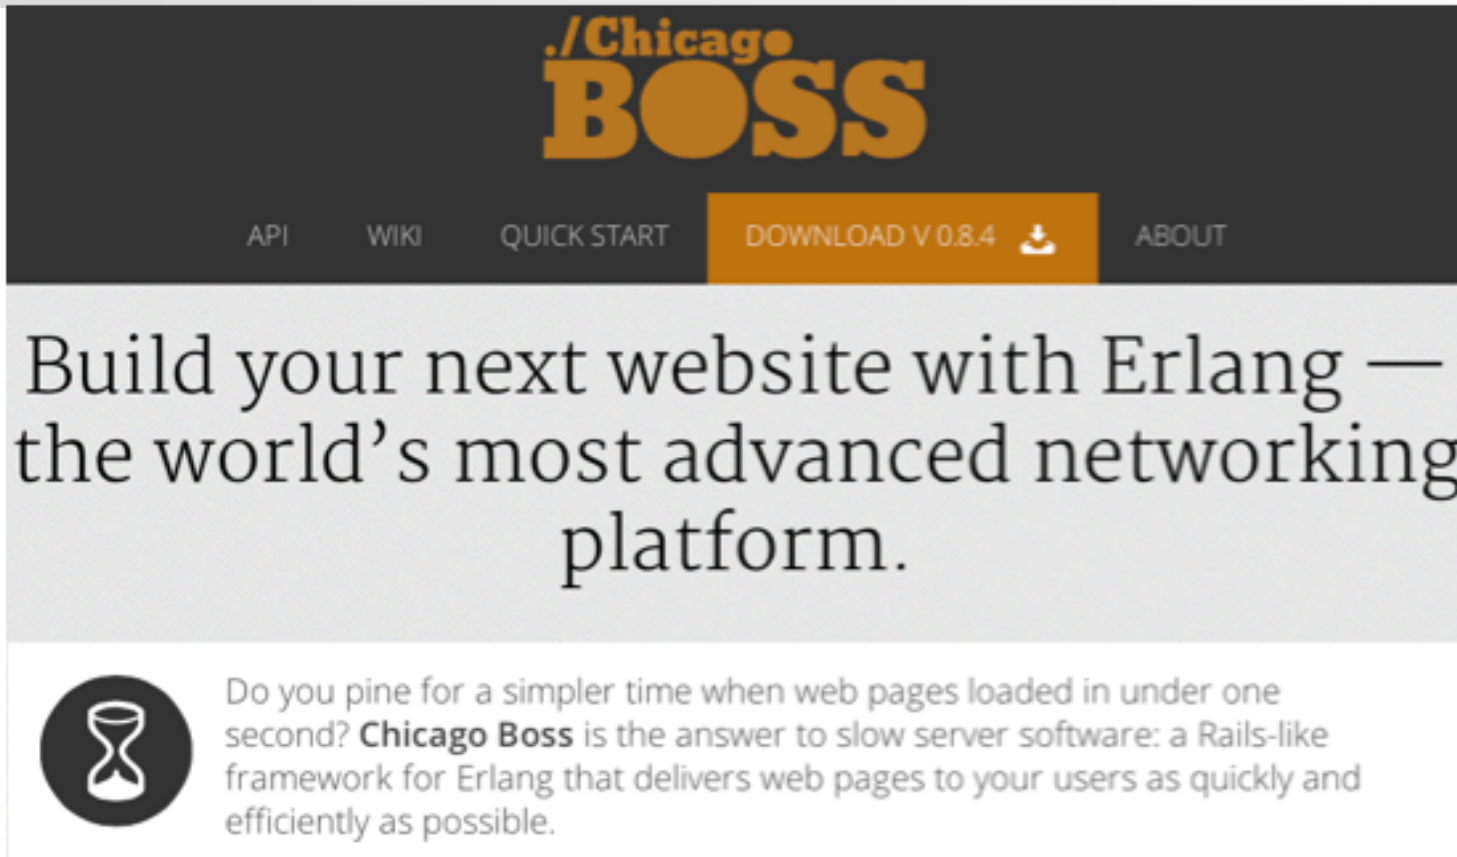
# Mandelbrot- 2013



**Mandelbrot throughput**



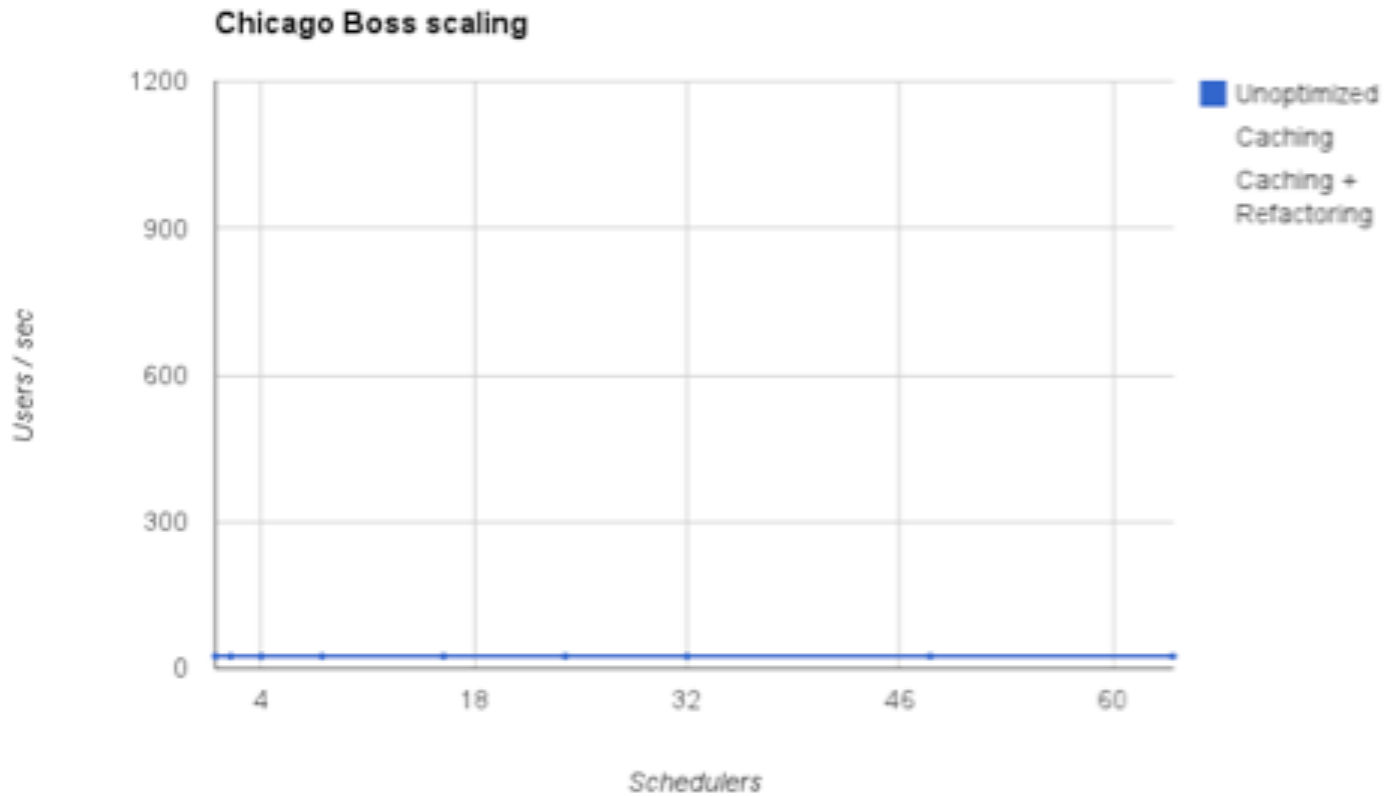
# Now for the Bottlenecks



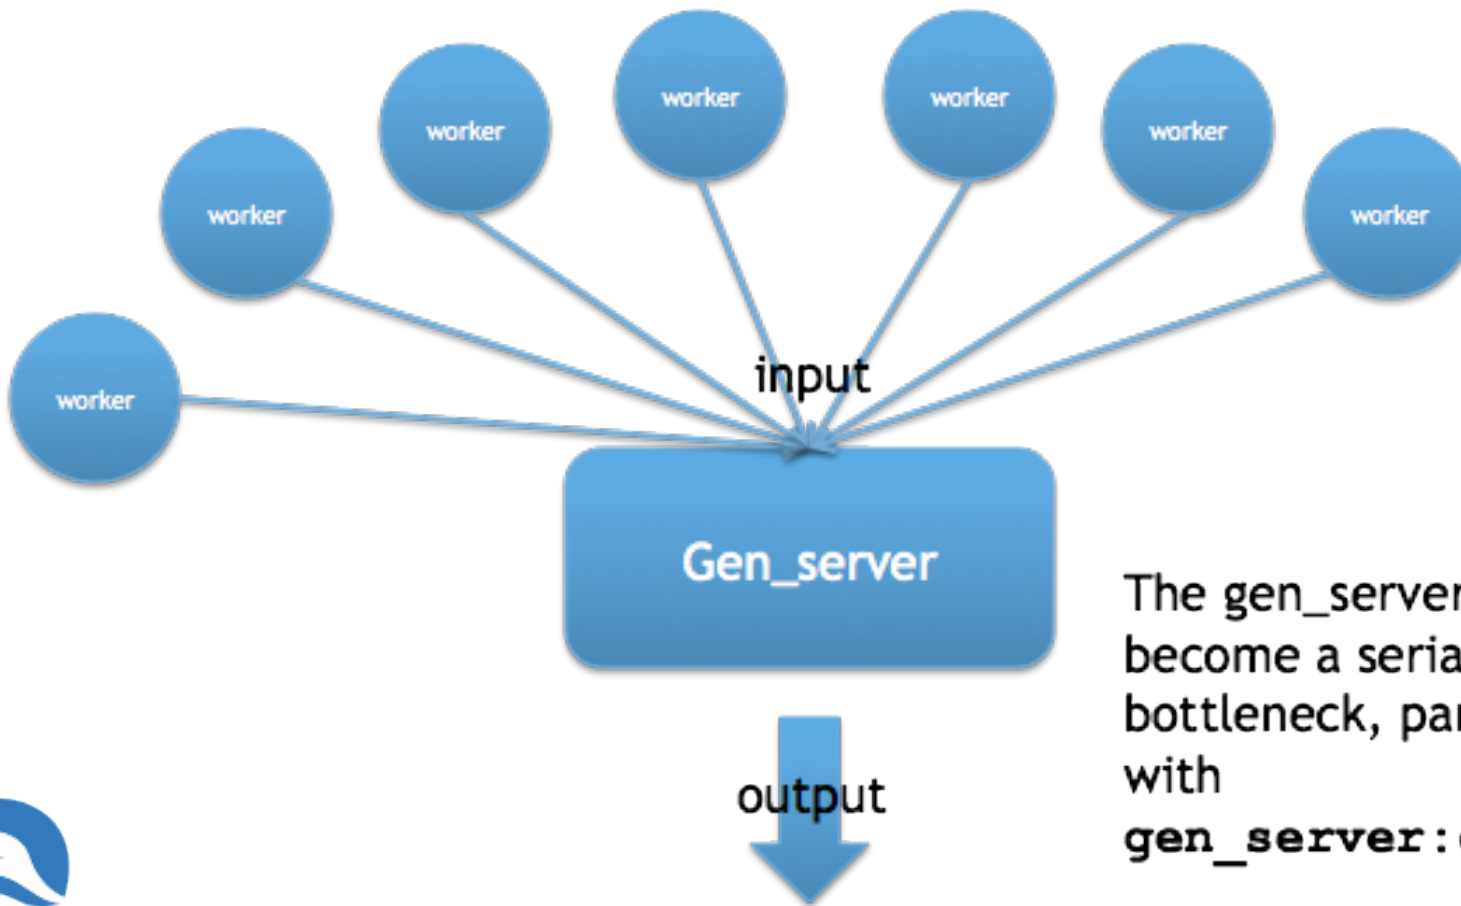
The screenshot shows the Chicago BOSS website. At the top, the logo reads **./Chicago BOSS**. Below the logo is a navigation bar with links for [API](#), [WIKI](#), [QUICK START](#), [DOWNLOAD V 0.8.4](#) (with a download icon), and [ABOUT](#). The main content area features the text: "Build your next website with Erlang — the world's most advanced networking platform." Below this is a section with an hourglass icon and the text: "Do you pine for a simpler time when web pages loaded in under one second? **Chicago Boss** is the answer to slow server software: a Rails-like framework for Erlang that delivers web pages to your users as quickly and efficiently as possible."



# Now for the Bottlenecks

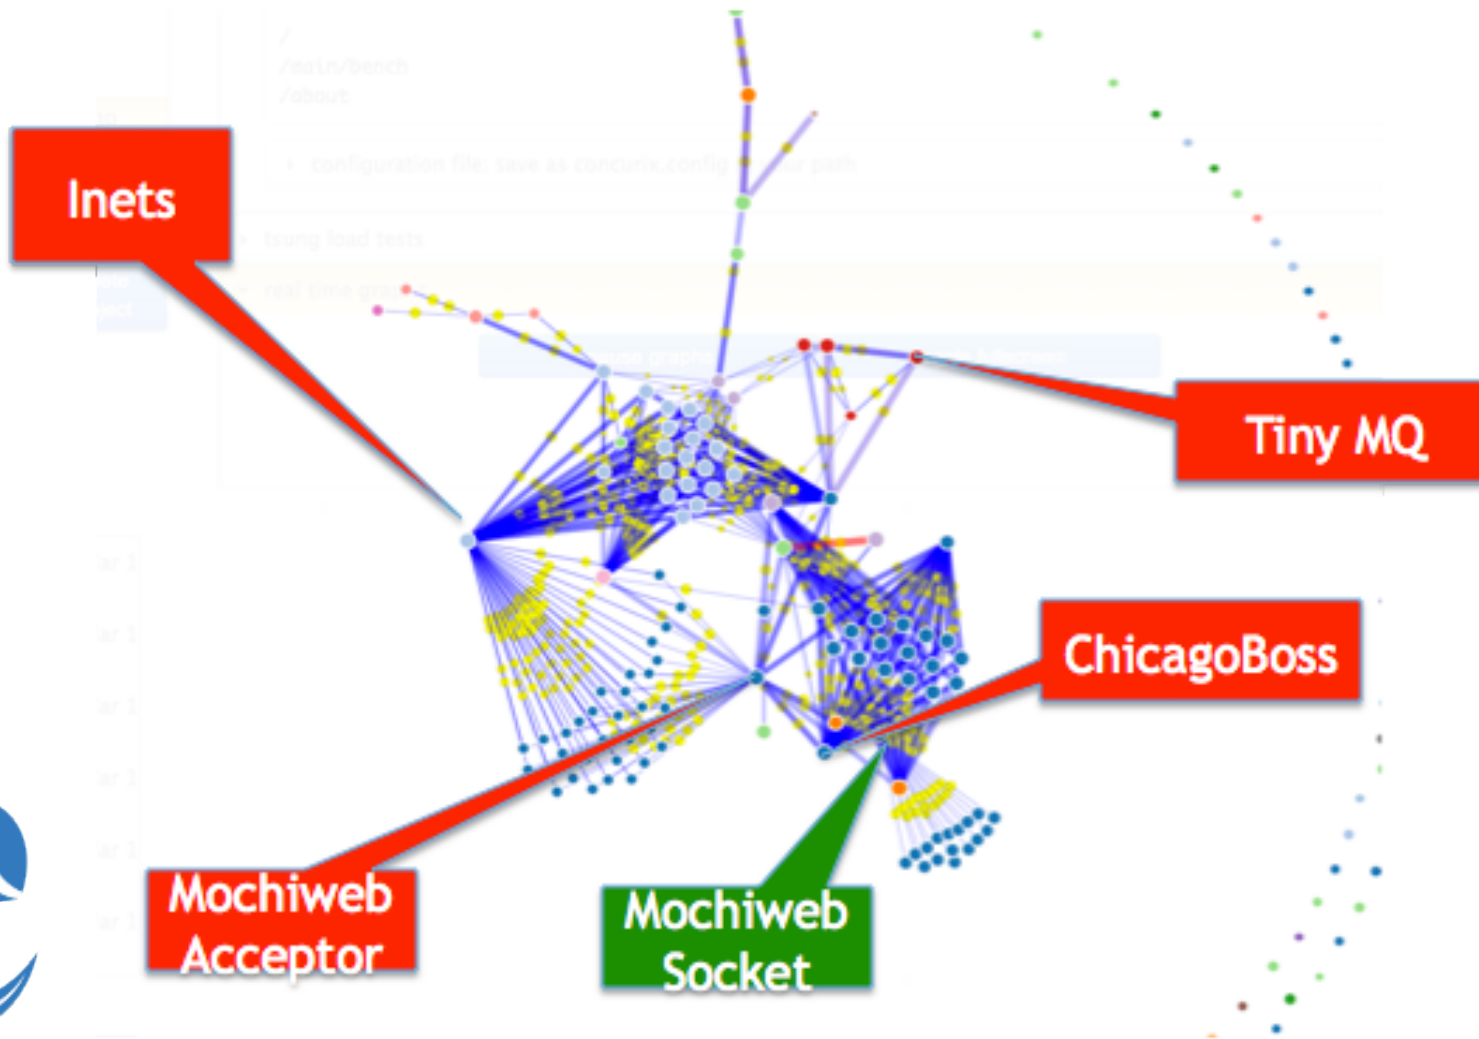


# Now for the Bottlenecks



# Now for the Bottlenecks

[www.concurix.com](http://www.concurix.com)




© 2014 - Erlang Solutions Ltd

*Erlang*  
SOLUTIONS




# Now for the Bottlenecks


Tons of headroom still!




# Riak and other scalable architectures



## N/R/W Values





## N/R/W Values



# Clusters and SD Erlang

- **TWO MAJOR ISSUES**
  - **FULLY CONNECTED** CLUSTERS
  - **EXPLICIT** PROCESS PLACEMENT
- **SCALABLE DISTRIBUTED (SD) ERLANG**
  - **NODES GROUPING**
  - **NON-TRANSITIVE** CONNECTIONS
  - **IMPLICIT** PROCESS PLACEMENT
  - PART OF THE **STANDARD** ERLANG/OTP PACKAGE
- **NEW CONCEPTS** INTRODUCED
  - **LOCALITY, AFFINITY AND DISTANCE**



# Release Statement of Aims



“To scale the radical *concurrency-oriented programming* paradigm to build *reliable* general-purpose software, such as server-based systems, on *massively parallel* machines ( $10^5$  cores).”

 **RELEASE**

University of  
**Kent**



UPPSALA  
UNIVERSITET

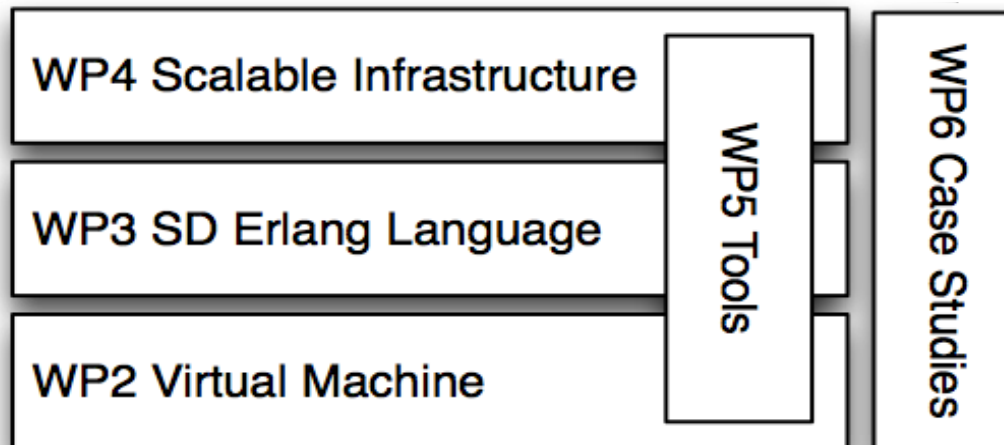


# Release



*“Limitations exist on all levels. You would not want an Erlang VM to run with  $10^5$  schedulers.”*

 **RELEASE**



# Release



Push the responsibility for scalability from the programmer to the VM

Analyze performance and scalability

Identify bottlenecks and prioritize changes and extensions

Tackle well-known scalability issues

Ets tables (shared global data structure)

Message passing, copying and frequently communicating processes



# Thank You!

@francescoc

francesco@erlang-solutions.com

