



JULY 16–20, 2012 **PORTLAND, OR**

**#oscon**

Apache httpd v2.4:  
***Watch Out Cloud***

***-or-***

***Hello Cloud: Buy you a drink?***

Jim Jagielski

# About me

- Jim Jagielski

- Hacker and developer
- Co-founder of the ASF
- Member, Director and President
- Director: Outercurve and OSI
- Council member: MARSEC-XL
- Consulting Engineer with Red Hat
- @jimjag



# What we will cover

- Overview of Apache httpd 2.4
  - General improvements
  - Reverse proxy improvements
- How the Cloud is a game-changer for web
- Performance Related Enhancements

# Apache httpd 2.4

- Currently at version 2.4.2 (2.4.1 went GA Feb 21, 2012)
- 2.4.3 RSN
- Significant Improvements
  - high-performance
  - cloud suitability

# Apache httpd 2.4 - design drivers

- Support for async I/O w/o dropping support for older systems
- Larger selection of usable MPMs: added Event, Simple, etc...
- Leverage higher-performant versions of APR
- Increase performance
- Reduce memory utilization
- The Cloud

# What's New: Apache httpd 2.4

- Bandwidth control now standard
  - mod\_ratelimit
- Finer control of timeouts, esp. during requests
  - mod\_reqtimeout
  - KeepAliveTimeout down to the millisecond
- Finer control over logging
  - per module/per directory
  - new logging levels (TRACE[1-8])
- <If> supports per-request conditions
- slot-based shared memory capability

# What's New: Apache httpd 2.4

- Controllable buffering of I/O
  - mod\_buffer
- Support for Lua (*still experimental as of 2.4.2*)
- Loadable MPMs
- General purpose Response Body substitution
  - mod\_sed
- Auto-convert Response -> RFC 2397 data URL
  - mod\_data
- Config file variables
- Cache improvements
- Proxy improvements ('natch)

# Why Proxy Matters

- Cloud puts big focus on horizontal scaling
- Apache httpd still the most frequently used front-end
- Proxy capabilities must be cloud friendly

# Proxy Design Drivers

- Becoming a robust but generic proxy implementation
- Support various protocols
  - HTTP, HTTPS, CONNECT, FTP
  - AJP, FastCGI, SCGI, WSGI (soon)
  - Load balancing
- Clustering, failover
- Performance

# What's New: Apache httpd 2.4 proxy

- Reverse Proxy Improvements
  - Supports FastCGI, SCGI in balancer
  - Additional load balancing mechanisms
  - Runtime changing of clusters w/o restarts
  - Support for dynamic configuration
  - mod\_proxy\_express
  - mod\_proxy\_html
  - mod\_fcgid

# Load Balancer

- `mod_proxy_balancer.so`
- `mod_proxy` can do native load balancing
  - weight by actual requests
  - weight by traffic
  - weight by busyness
  - `lb` factors

# Load Balancer

- Backend connection pooling
- Available for named workers:
  - eg: `ProxyPass /foo http://bar.example.com`
- Reusable connection to origin
  - For threaded MPMs, can adjust size of pool (min, max, smax)
  - For prefork: singleton
- Shared data held in shared memory

# Load Balancer

- Sticky session support
  - aka “session affinity”
- Cookie based
  - stickysession=PHPSESSID
  - stickysession=JSESSIONID
- Natively easy with Tomcat
- May require more setup for “simple” HTTP proxying

# Load Balancer

- Cluster set with failover
- Group backend servers as numbered sets
  - balancer will try lower-valued sets first
  - If no workers are available, will try next set
- Hot standby

# Putting it all together

```
<Proxy balancer://foo>
    BalancerMember http://php1:8080/      loadfactor=1
    BalancerMember http://php2:8080/      loadfactor=4
    BalancerMember http://phpbkup:8080/    loadfactor=1 status=+h
    BalancerMember http://phpexp:8080/     lbset=1
    ProxySet lbmethod=bytraffic
</Proxy>

<Proxy balancer://javaapps>
    BalancerMember ajp://tc1:8089/        loadfactor=1
    BalancerMember ajp://tc2:8089/        loadfactor=4
    ProxySet lbmethod=byrequests
</Proxy>

ProxyPass /apps/ balancer://foo/
ProxyPassReverse /apps/ balancer://foo/
ProxyPass /serv/ balancer://javaapps/
ProxyPass /images/ http://images:8080/
```

# Embedded Admin

- Allows for real-time
  - Monitoring of stats for each worker
  - Adjustment of worker params
    - lbset
    - load factor
    - route
    - enabled / disabled
    - ...

# Embedded Admin

- Allows for real-time
  - Addition of new workers/nodes
  - Change of LB methods
  - Can be persistent
  - More RESTful
  - Can be CLI-driven

# Easy setup

```
<Location /balancer-manager>  
    SetHandler balancer-manager  
    Order Deny,Allow  
    Deny from all  
    Allow from 192.168.2.22  
</Location>
```

# Admin

Balancer Manager

+

http://localhost:8880/balancer-manager

Google

## Load Balancer Manager for localhost

Server Version: Apache/2.3.15-dev (Unix) DAV/2  
Server Built: Nov 1 2011 06:19:34

### LoadBalancer Status for [balancer://acna11](#)

MaxMembers	StickySession	DisableFailover	Timeout	FailoverAttempts	Method
8 [3 Used]	(None)	Off	0	2	bytraffic

Worker URL	Route	RouteRedir	Factor	Set	Status	Elected	Busy	Load To	From
<a href="#">http://www1.example.com</a>			1	0	Init Ok	5	0	0	2.1K 110
<a href="#">http://www2.example.com</a>			1	0	Init Ok	5	0	0	2.1K 110
<a href="#">http://www3.example.com/soap/brackletool/</a>			1	0	Init Stby Ok	0	0	0	0

# Admin

Balancer Manager

http://localhost:8880/balancer-manager

## Load Balancer Manager for localhost

Server Version: Apache/2.3.15-dev (Unix) DAV/2  
Server Built: Nov 1 2011 06:19:34

---

LoadBalancer Status for [balancer://acna11](#)

MaxMembers StickySession DisableFailover Timeout FailoverAttempts Method  
8 [3 Used] (None) Off 0 2 bytraffic

Worker URL	Route	RouteRedir	Factor	Set	Status	Elected	Busy	Load	To	From
<a href="http://www1.example.com">http://www1.example.com</a>			1	0	Init Ok	5	0	0	2.1K	110
<a href="http://www2.example.com">http://www2.example.com</a>			1	0	Init Ok	5	0	0	2.1K	110
<a href="http://www3.example.com/snap/crackle/pool/">http://www3.example.com/snap/crackle/pool/</a>			1	0	Init Stby Ok	0	0	0	0	0

Click here

# Admin

Balancer Manager

http://localhost:8880/balancer-manager/

Google

## Load Balancer Manager for localhost

Server Version: Apache/2.3.15-dev (Unix) DAV/2  
Server Built: Nov 1 2011 06:19:34

---

### LoadBalancer Status for [balancer://acna11](#)

MaxMembers	StickySession	DisableFallover	Timeout	FalloverAttempts	Method
8 [3 Used]	(None)	Off	0	2	bytraffic

---

Worker URL	Route	RouteRedir	Factor	Set	Status	Elected	Busy	Load To	From
<a href="#">http://www1.example.com</a>			1	0	Init Ok	5	0	0	2.1K 110
<a href="#">http://www2.example.com</a>			1	0	Init Ok	5	0	0	2.1K 110
<a href="#">http://www3.example.com/snap/crackle/pop/</a>			1	0	Init Sby Ok	0	0	0	0 0

---

### Edit worker settings for [http://www3.example.com/snap/crackle/pop/](#)

Load factor:

LB Set:

Route:

Route Redirect:

Status:

Ign	Drn	Dis	Sby
On <input type="radio"/>	On <input type="radio"/>	On <input type="radio"/>	On <input checked="" type="radio"/>
Off <input checked="" type="radio"/>	Off <input checked="" type="radio"/>	Off <input checked="" type="radio"/>	Off <input type="radio"/>

# Admin

Balancer Manager

http://localhost:8880/balancer-manager/

Google

Load Balancer Manager for localhost

Server Version: Apache/2.3.15-dev (Unix) DAV/2  
Server Built: Nov 1 2011 06:19:34

LoadBalancer Status for [balancer://acna11](#)

MaxMembers

StickySession

DisableFallover

Timeout

FalloverAttempts

Method

8 [3 Used]

(None)

Off

0

2

bytraffic

Worker URL

Route

RouteRedir

Factor

Weight

Status

Elected

Busy

Load To

From

[http://www1.example.com](#)

1

0

Init Ok

5

0

0

2.1K

110

[http://www2.example.com](#)

1

0

Init

5

0

0

2.1K

110

[http://www3.example.com/snap/crackle/pop/](#)

1

0

Init Stby

0

0

0

0

0

Edit worker settings for [http://www3.example.com/snap/crackle/pop/](#)

Load factor:

1

LB Set:

0

Route:

Route Redirect:

Status:

Ign

On ☐

Off ☒

Drn

On ☐

Off ☒

Dis

On ☐

Off ☒

Stby

On ☒

Off ☐

Submit

Click here

Click here

# Admin

Balancer Manager

http://localhost:8880/balancer-manager/

Google

## Load Balancer Manager for localhost

Server Version: Apache/2.3.15-dev (Unix) DAV/2  
Server Built: Nov 1 2011 06:19:34

---

### LoadBalancer Status for [balancer://acna11](#)

MaxMembers	StickySession	DisableFailover	Timeout	FailoverAttempts	Method
8 [3 Used]	(None)	Off	0	2	bytraffic

---

Worker URL	Route	RouteRedir	Factor	Set	Status	Elected	Busy	Load To	From
<a href="#">http://www1.example.com</a>			1	0	Init Ok	5	0	0	2.1K 110
<a href="#">http://www2.example.com</a>			1	0	Init Ok	5	0	0	2.1K 110
<a href="#">http://www3.example.com/snapcracklepop/</a>			1	0	Init Stry Ok	0	0	0	0

---

### Edit balancer settings for [balancer://acna11](#)

LBmethod:

Timeout:

Failover Attempts:

Disable Failover: ☐ On ☒ Off

Sticky Session:  (Use ^ to delete)

Add New Worker:  Are you sure? ☐

O'REILLY

oscon

open source convention

# Admin

**Load Balancer Manager for localhost**

Server Version: Apache/2.3.15-dev (Unix) DAV/2  
Server Built: Nov 1 2011 06:19:34

**LoadBalancer Status for [balancer://acna11](#)**

MaxMembers	StickySession	DisableFailover	Timeout	FalloverAttempts	Method
8 [3 Used]	(None)	Off	0	2	bytraffic

Worker URL	Route	RouteRedir	Factor	Set	Status	Elected	Busy	Load To	From
<a href="#">http://www1.example.com</a>			1	0	Init Ok	5	0	0	2.1K 110
<a href="#">http://www2.example.com</a>			1	0	Init Ok	5	0	0	2.1K 110
<a href="#">http://www3.example.com/snapcracklepop/</a>			1	0	Init Stuck	0	0	0	0

**Edit balancer settings for [balancer://acna11](#)**

LBmethod:  **Changing the LBmethod**

Timeout:

Fallover Attempts:

Disable Failover: ☐ On ☒ Off

Sticky Session:  (Use ^ to del)

Add New Worker:  Are you sure? ☒ **Adding new worker**

# Admin

Balancer Manager

http://localhost:8880/balancer-manager

Google

## Load Balancer Manager for localhost

Server Version: Apache/2.3.15-dev (Unix) DAV/2  
Server Built: Nov 1 2011 06:19:34

### LoadBalancer Status for [balancer://acna11](#)

MaxMembers	StickySession	DisableFailover	Timeout	FailoverAttempts	Method
8 [4 Used]	(None)	Off	0	2	byrequests

Worker URL	Route	RouteRedir	Factor	Set	Status	Elected	Busy	Load To	From
<a href="#">http://www1.example.com</a>			1	0	Init Ok	5	0	0	2.1K 110
<a href="#">http://www2.example.com</a>			1	0	Init Ok	5	0	0	2.1K 110
<a href="#">http://www3.example.com/snapcracklepop/</a>			1	0	Init Stby Ok	0	0	0	0
<a href="#">http://www4.example.com/acna</a>			1	0	Init Dis	0	0	0	0

### Edit balancer settings for [balancer://acna11](#)

LBmethod:

byrequests

Timeout:

0

Failover Attempts:

2

Disable Failover:

On

Off

Sticky Session:

(Use ^ to delete)

Add New Worker:

Are you sure?

Submit

Wow!

Wow!

# Mass Reverse Proxy

- Use the new mod\_proxy\_express module
  - ProxyPass mapping obtained via db file
  - Fast and efficient
  - Still dynamic, with no config changes required

## ProxyExpress map file

```
##
##express-map.txt:
##

www1.example.com      http://192.168.002.2:8080
www2.example.com      http://192.168.002.12:8088
www3.example.com      http://192.168.002.10
...
www6341.example.com   http://192.168.211.26
```

```
mmm034j.example.com  http://192.168.211.26
```

```
...
```

```
www1.example.com      http://192.168.002.2:8080
```

# What's on the horizon?

- Improving AJP
- Adding additional protocols
- More dynamic configuration
  - Adding balancers!

# Cloud and Performance

- The Cloud is a game changer for web servers
  - Horizontal scalability is no longer as painful
  - Concurrency is no longer the sole consideration
  - ... or even the primary one
  - What's important now? Transaction Time!
    - Low latency
    - Fast req/resp turnover
  - Does density still matter? *Of course!*
  - Are there environs where concurrency is the bugaboo? *You betcha! (but the cloud makes these more and more rare)*

# Apache httpd vs nginx

- Why nginx? Everyone asks about it...
- Benchmark: local and reverse proxy transaction times
  - Apache httpd 2.4.1-dev, nginx 1.2.0
  - Fedora 16, Dual Xeon 2.28GHz
  - 4GB memory
  - localhost loopback and external (no firewall)
  - Double checked results: OSX, Ubuntu 10.04

# Setup

# Setup

Setup 1:

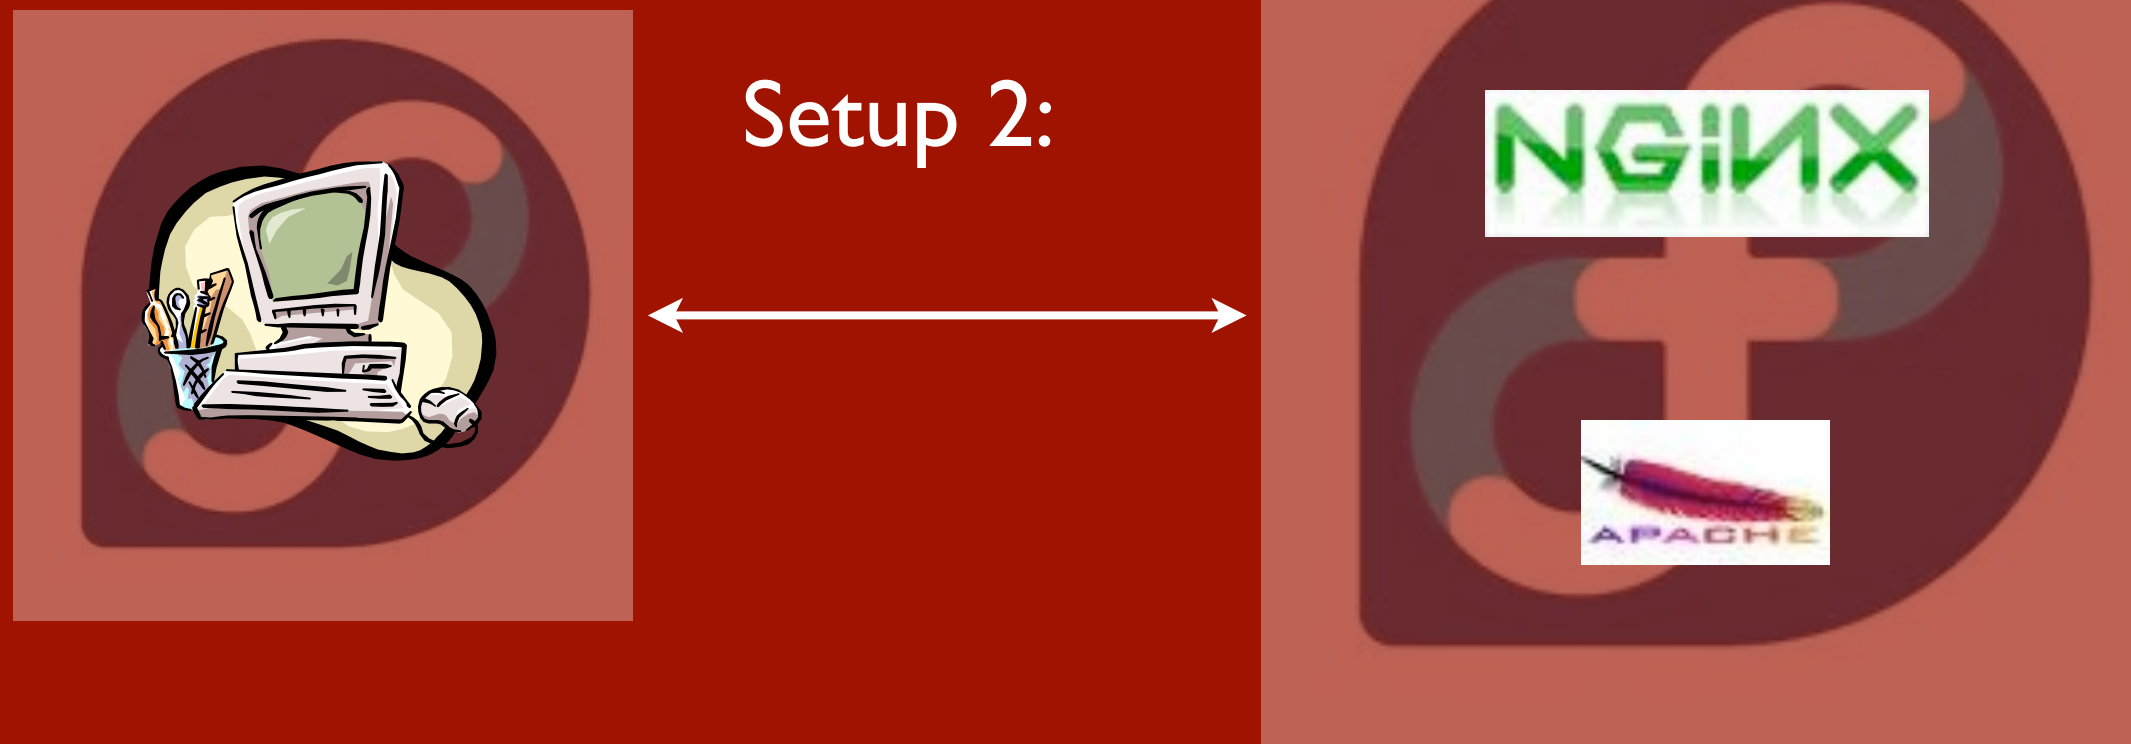


# Setup

Setup 1:



Setup 2:



# Setup

Setup 1:



Setup 2:



Setup 3:

Setup 3:

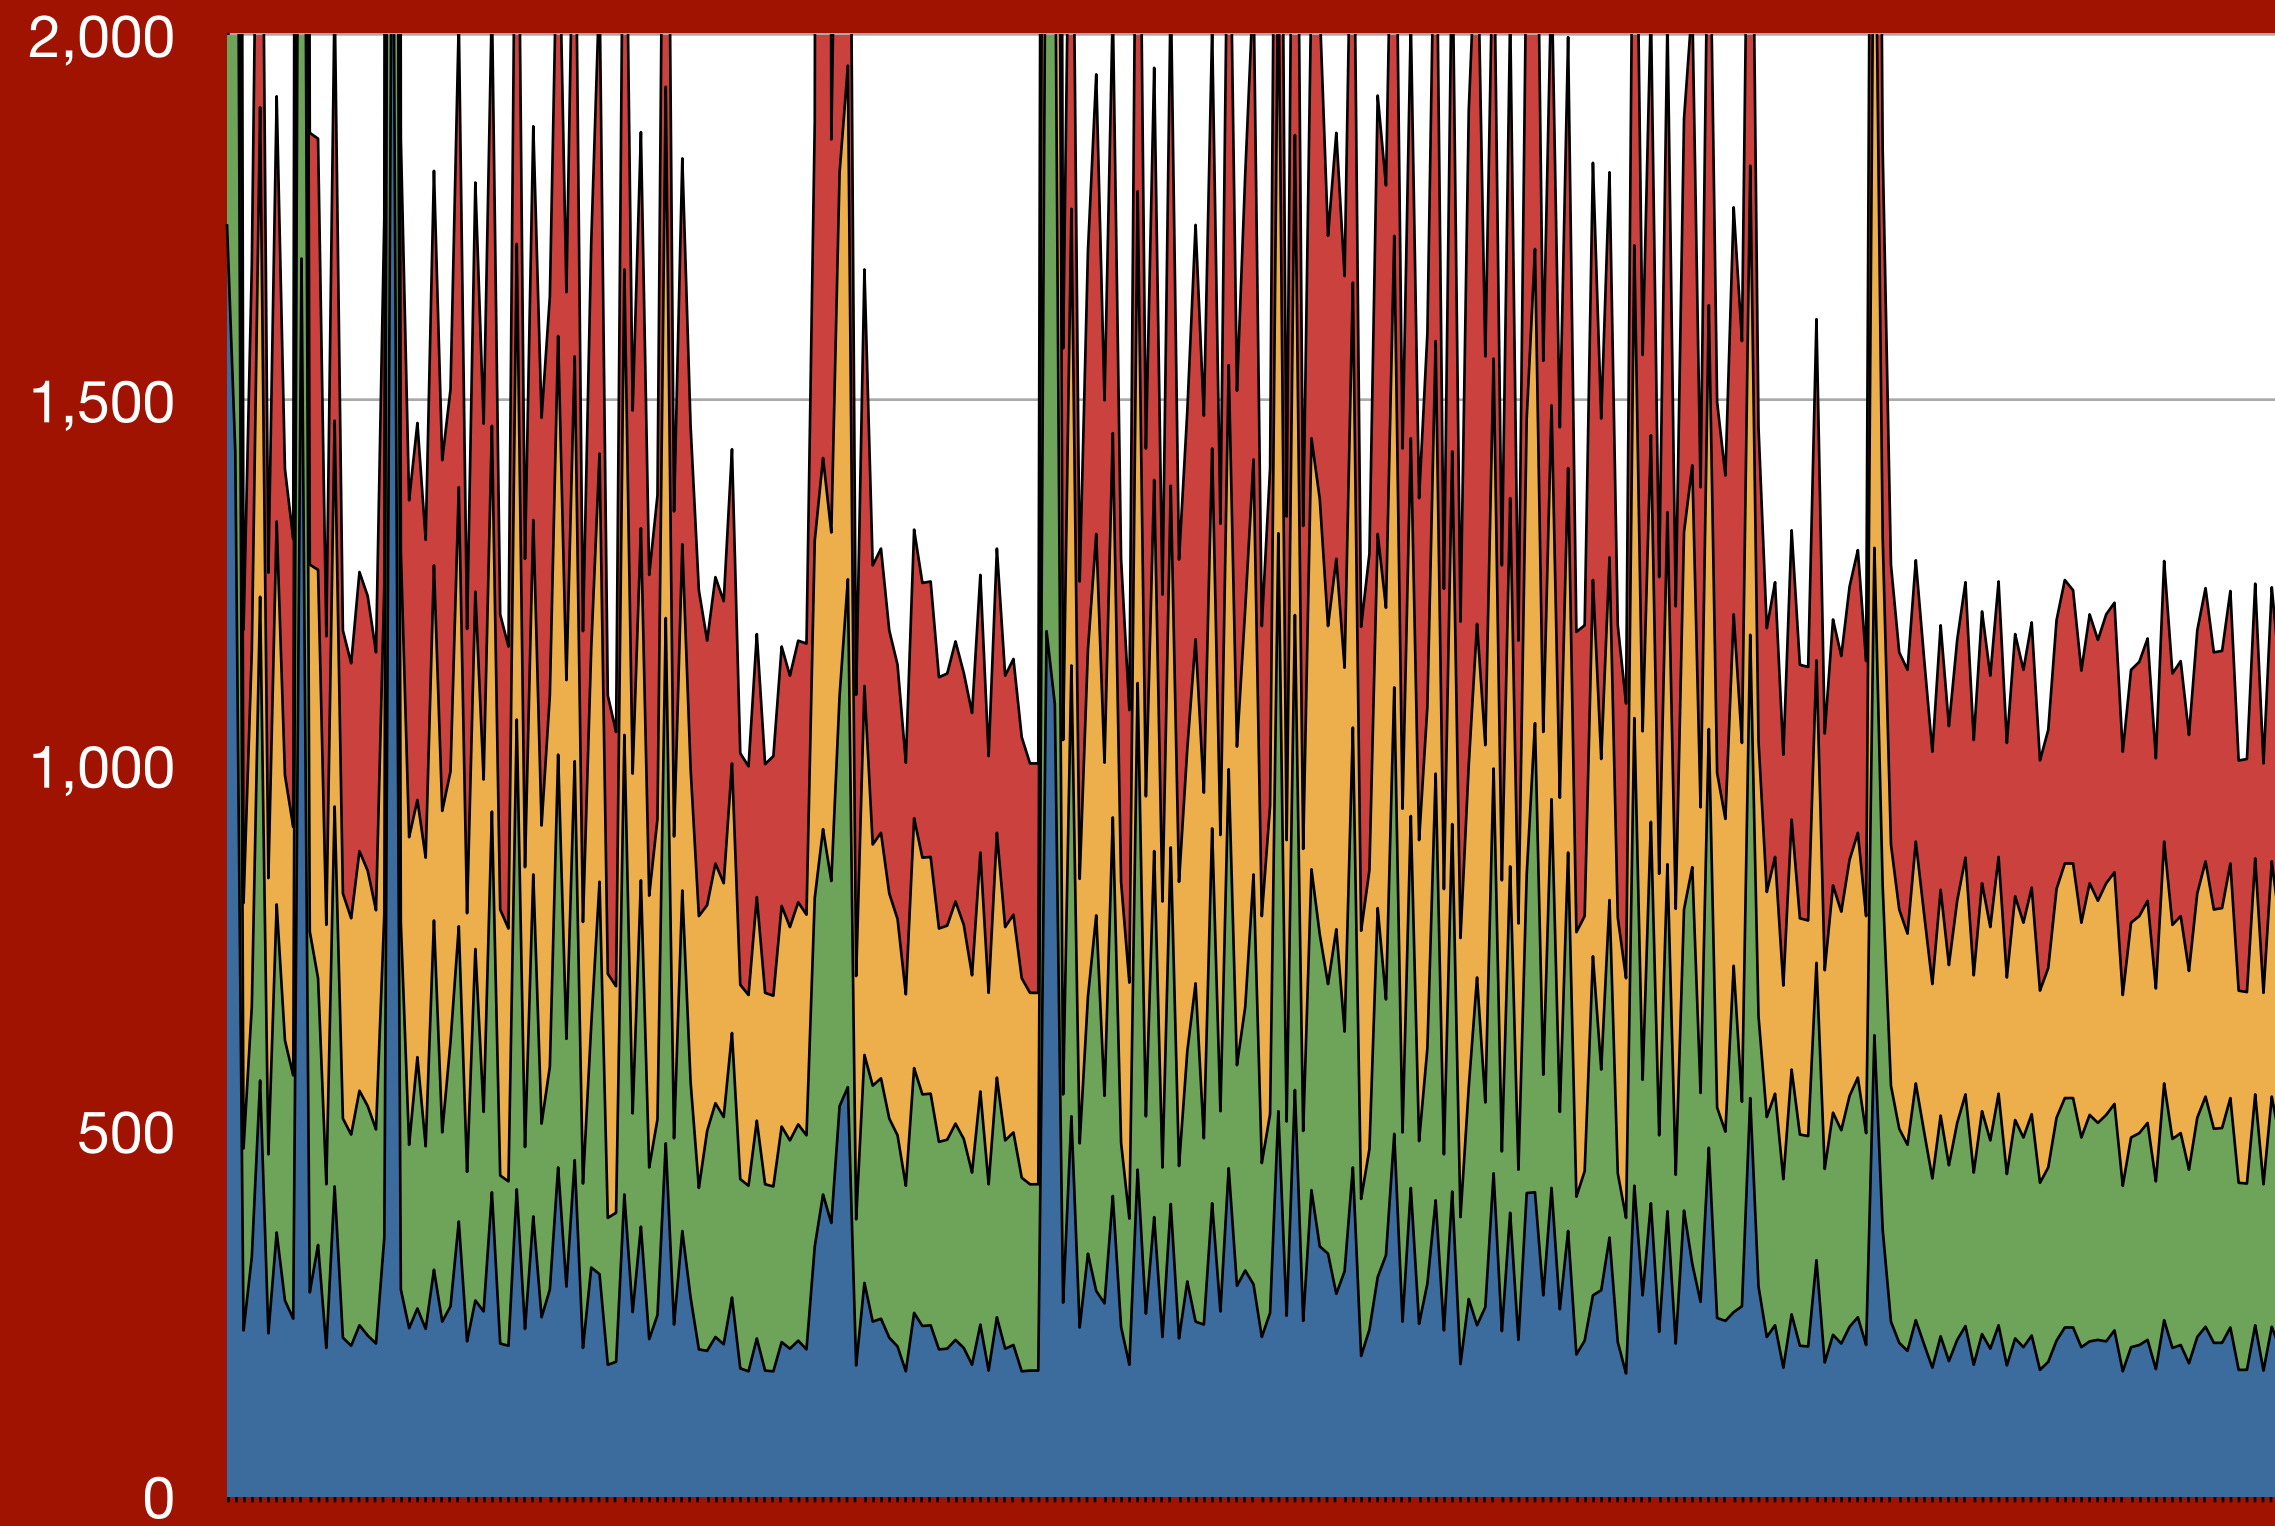


# Considerations

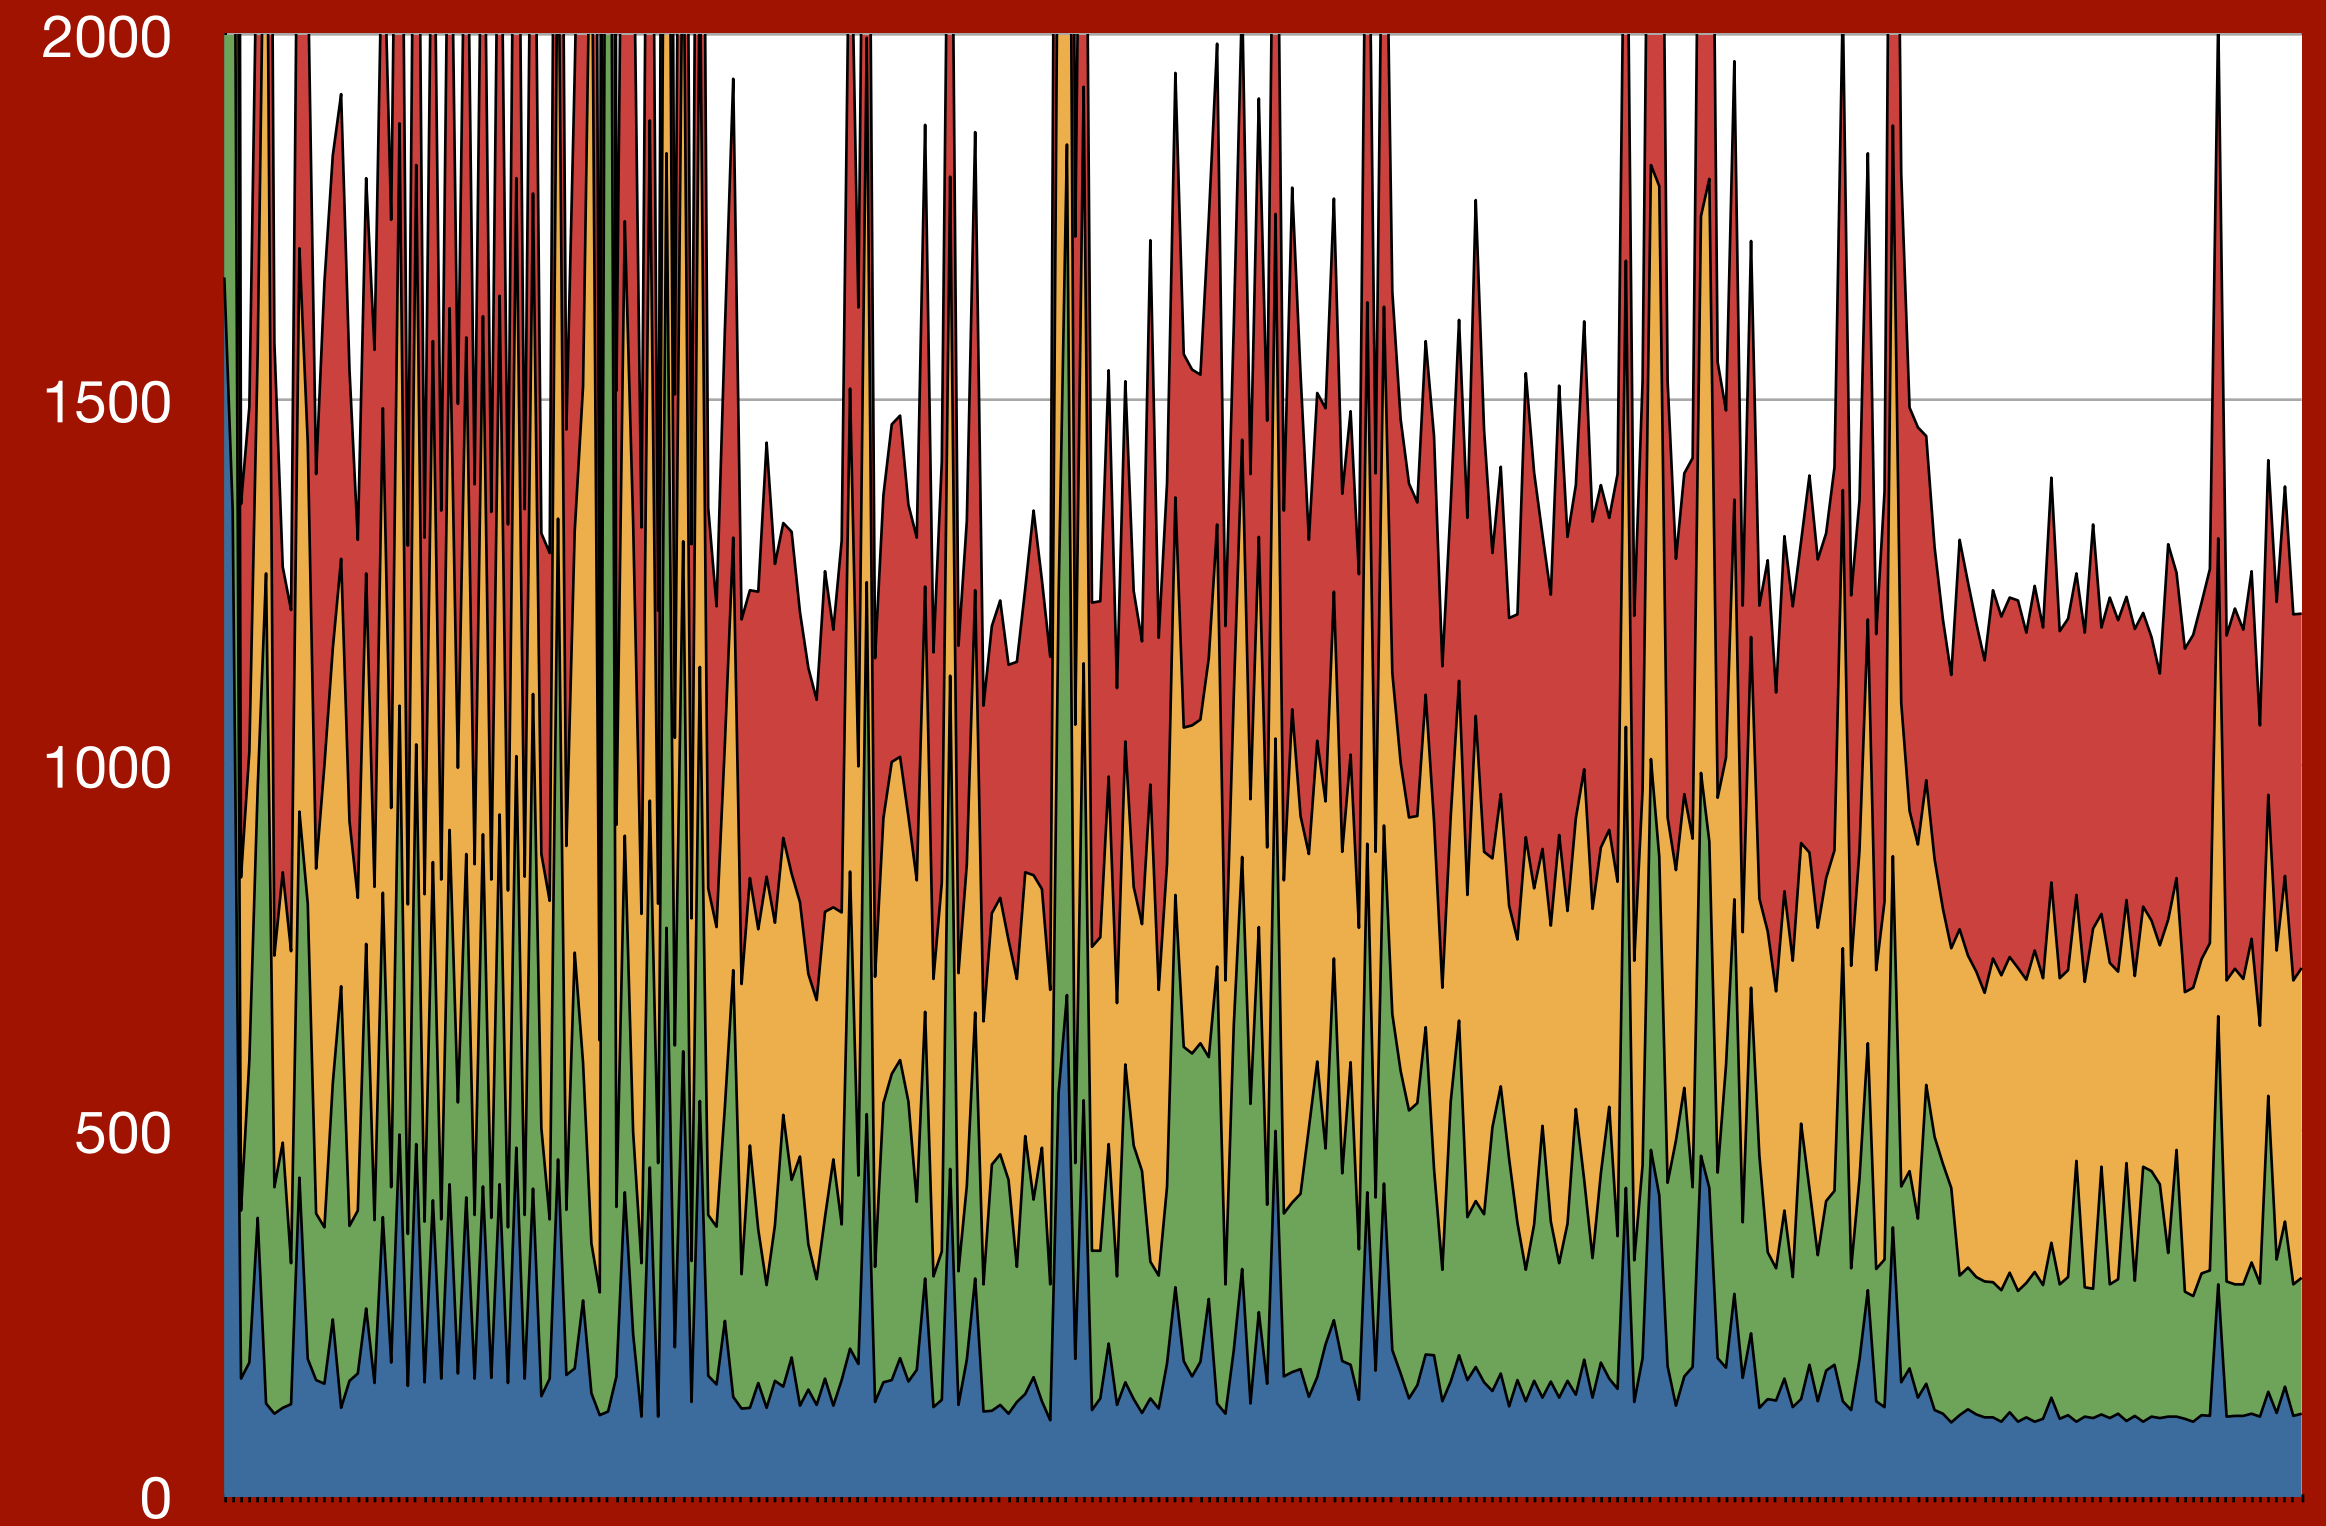
- Multiple benchmarking systems:
  - flood (50/250/5/2, 50/100/5/2, 50/5/5/2)
  - httpperf (num-conns=100->10000, numcalls=3)
- Full URL requests (www.example.com/index.html)
- Static local requests
- Static reverse proxy requests
- All Apache httpd MPMs
- No significant “tuning” efforts (mostly out of the box configs)

# nginx vs Event (typical)

nginx



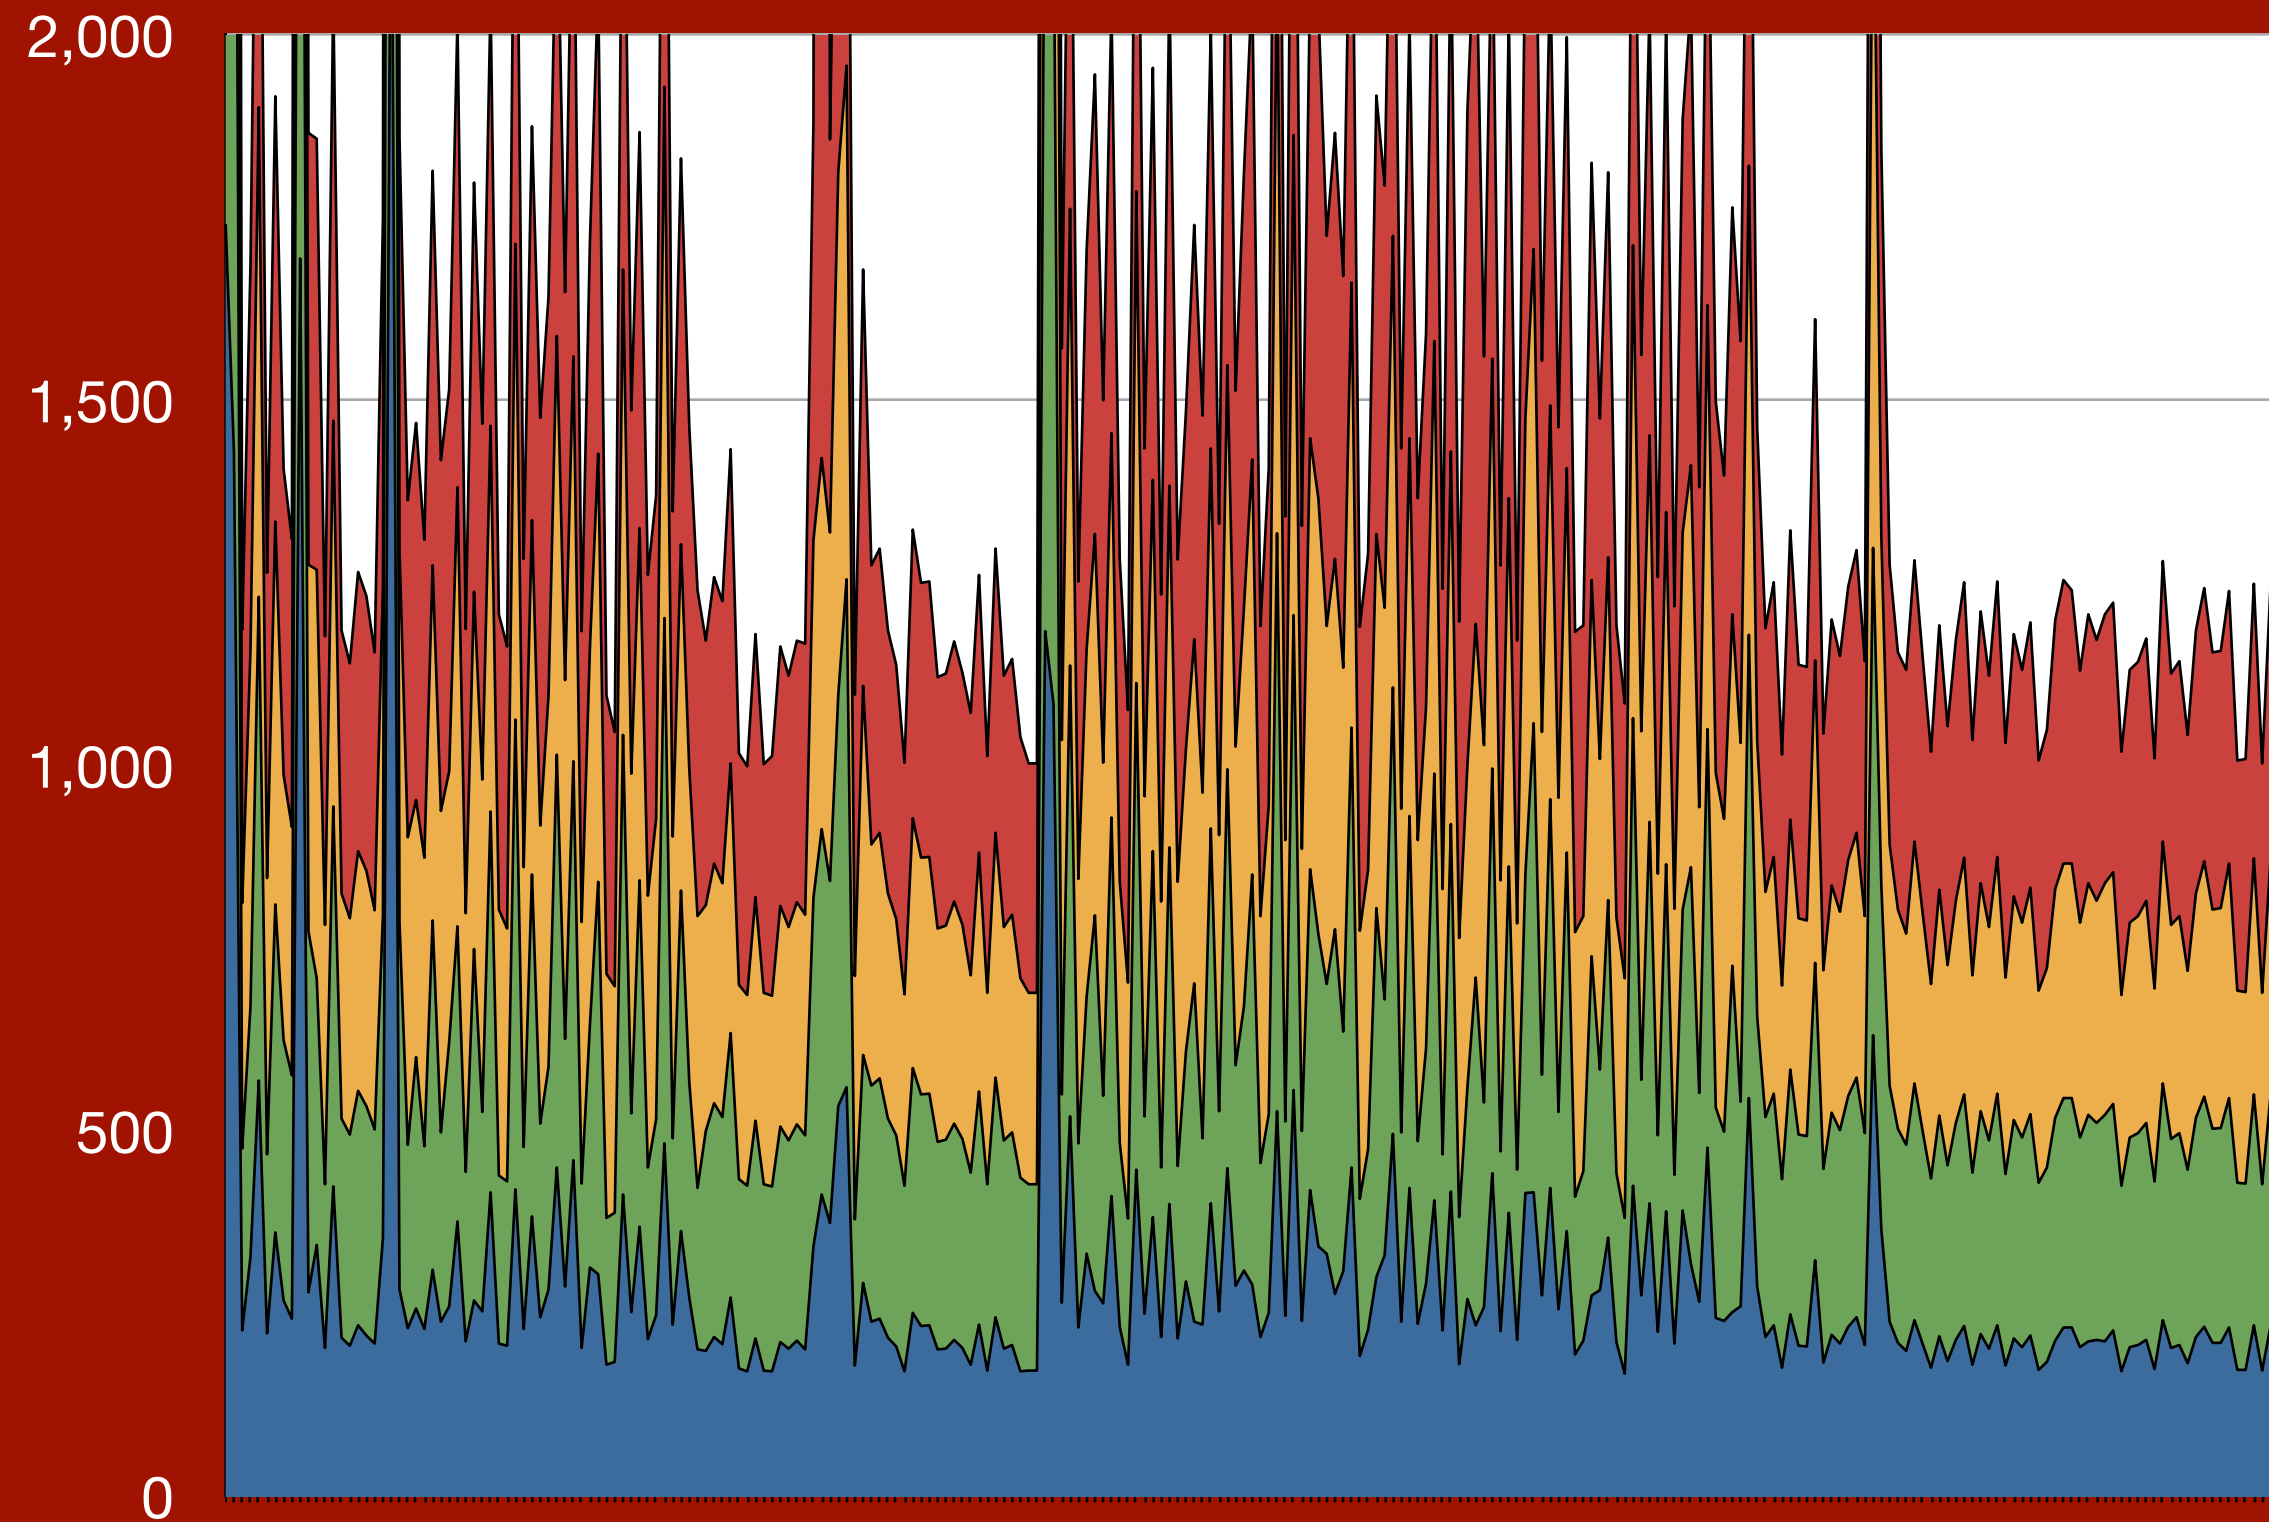
Apache - Event MPM



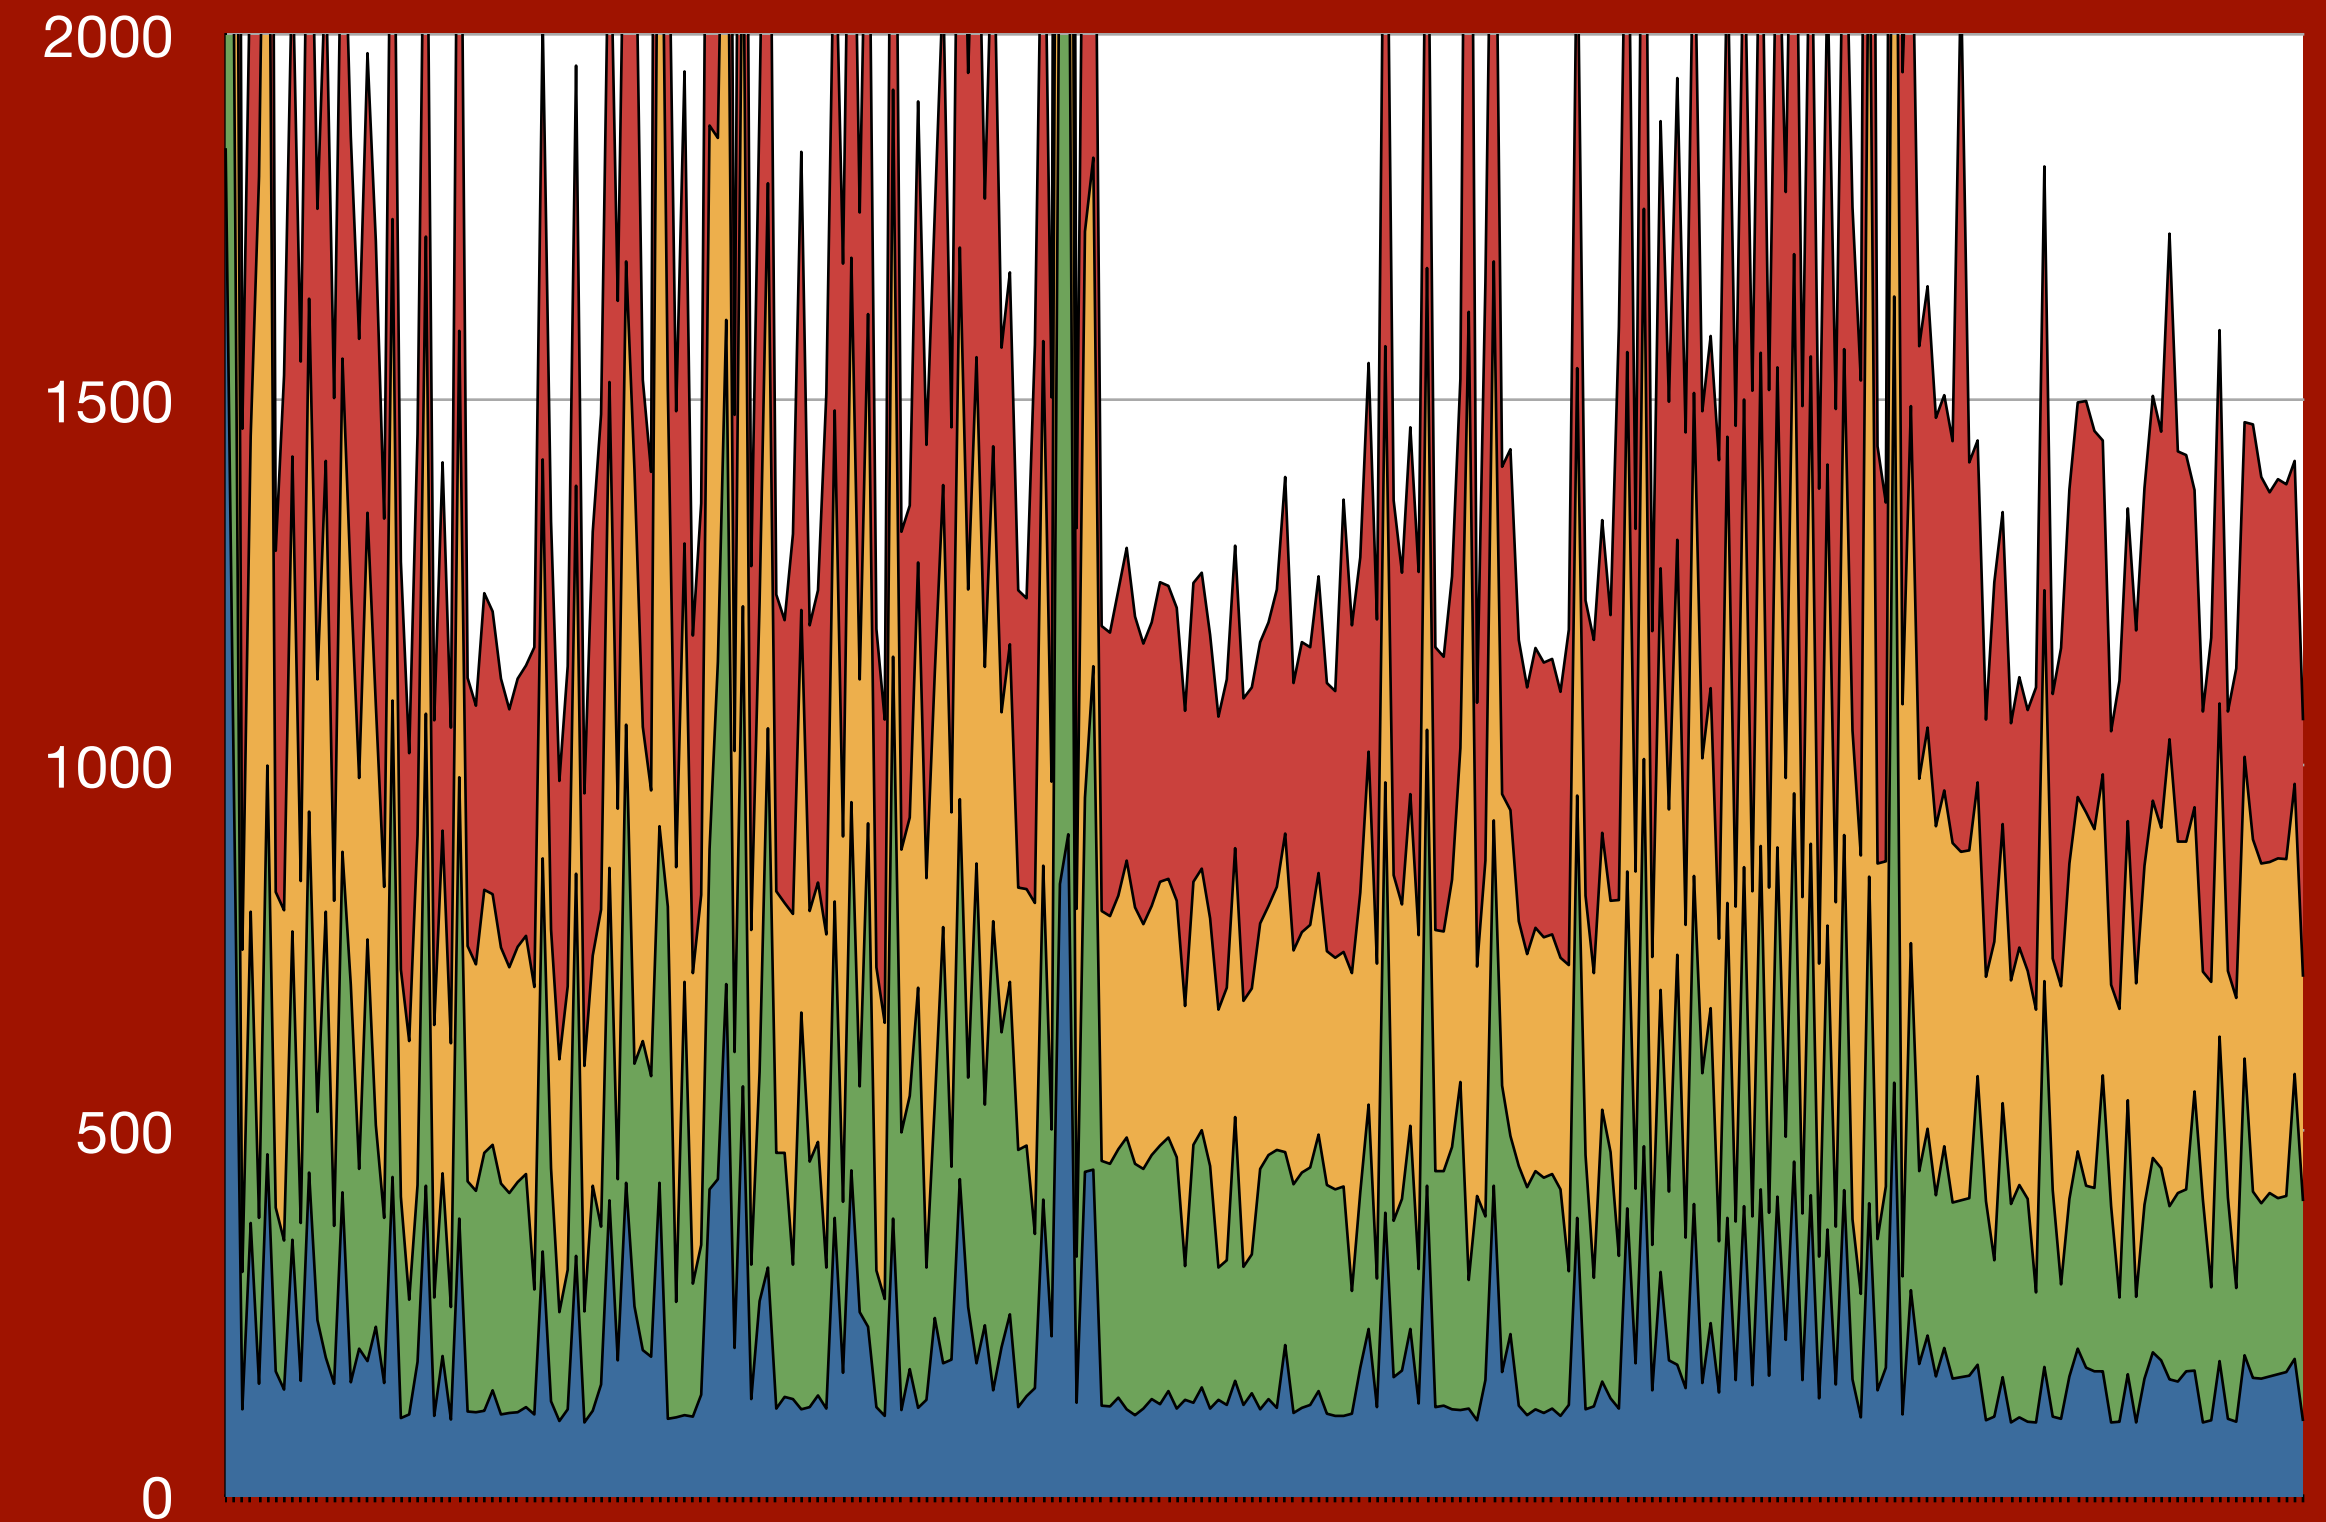
Open Write Read Close

# nginx vs Worker (typical)

nginx



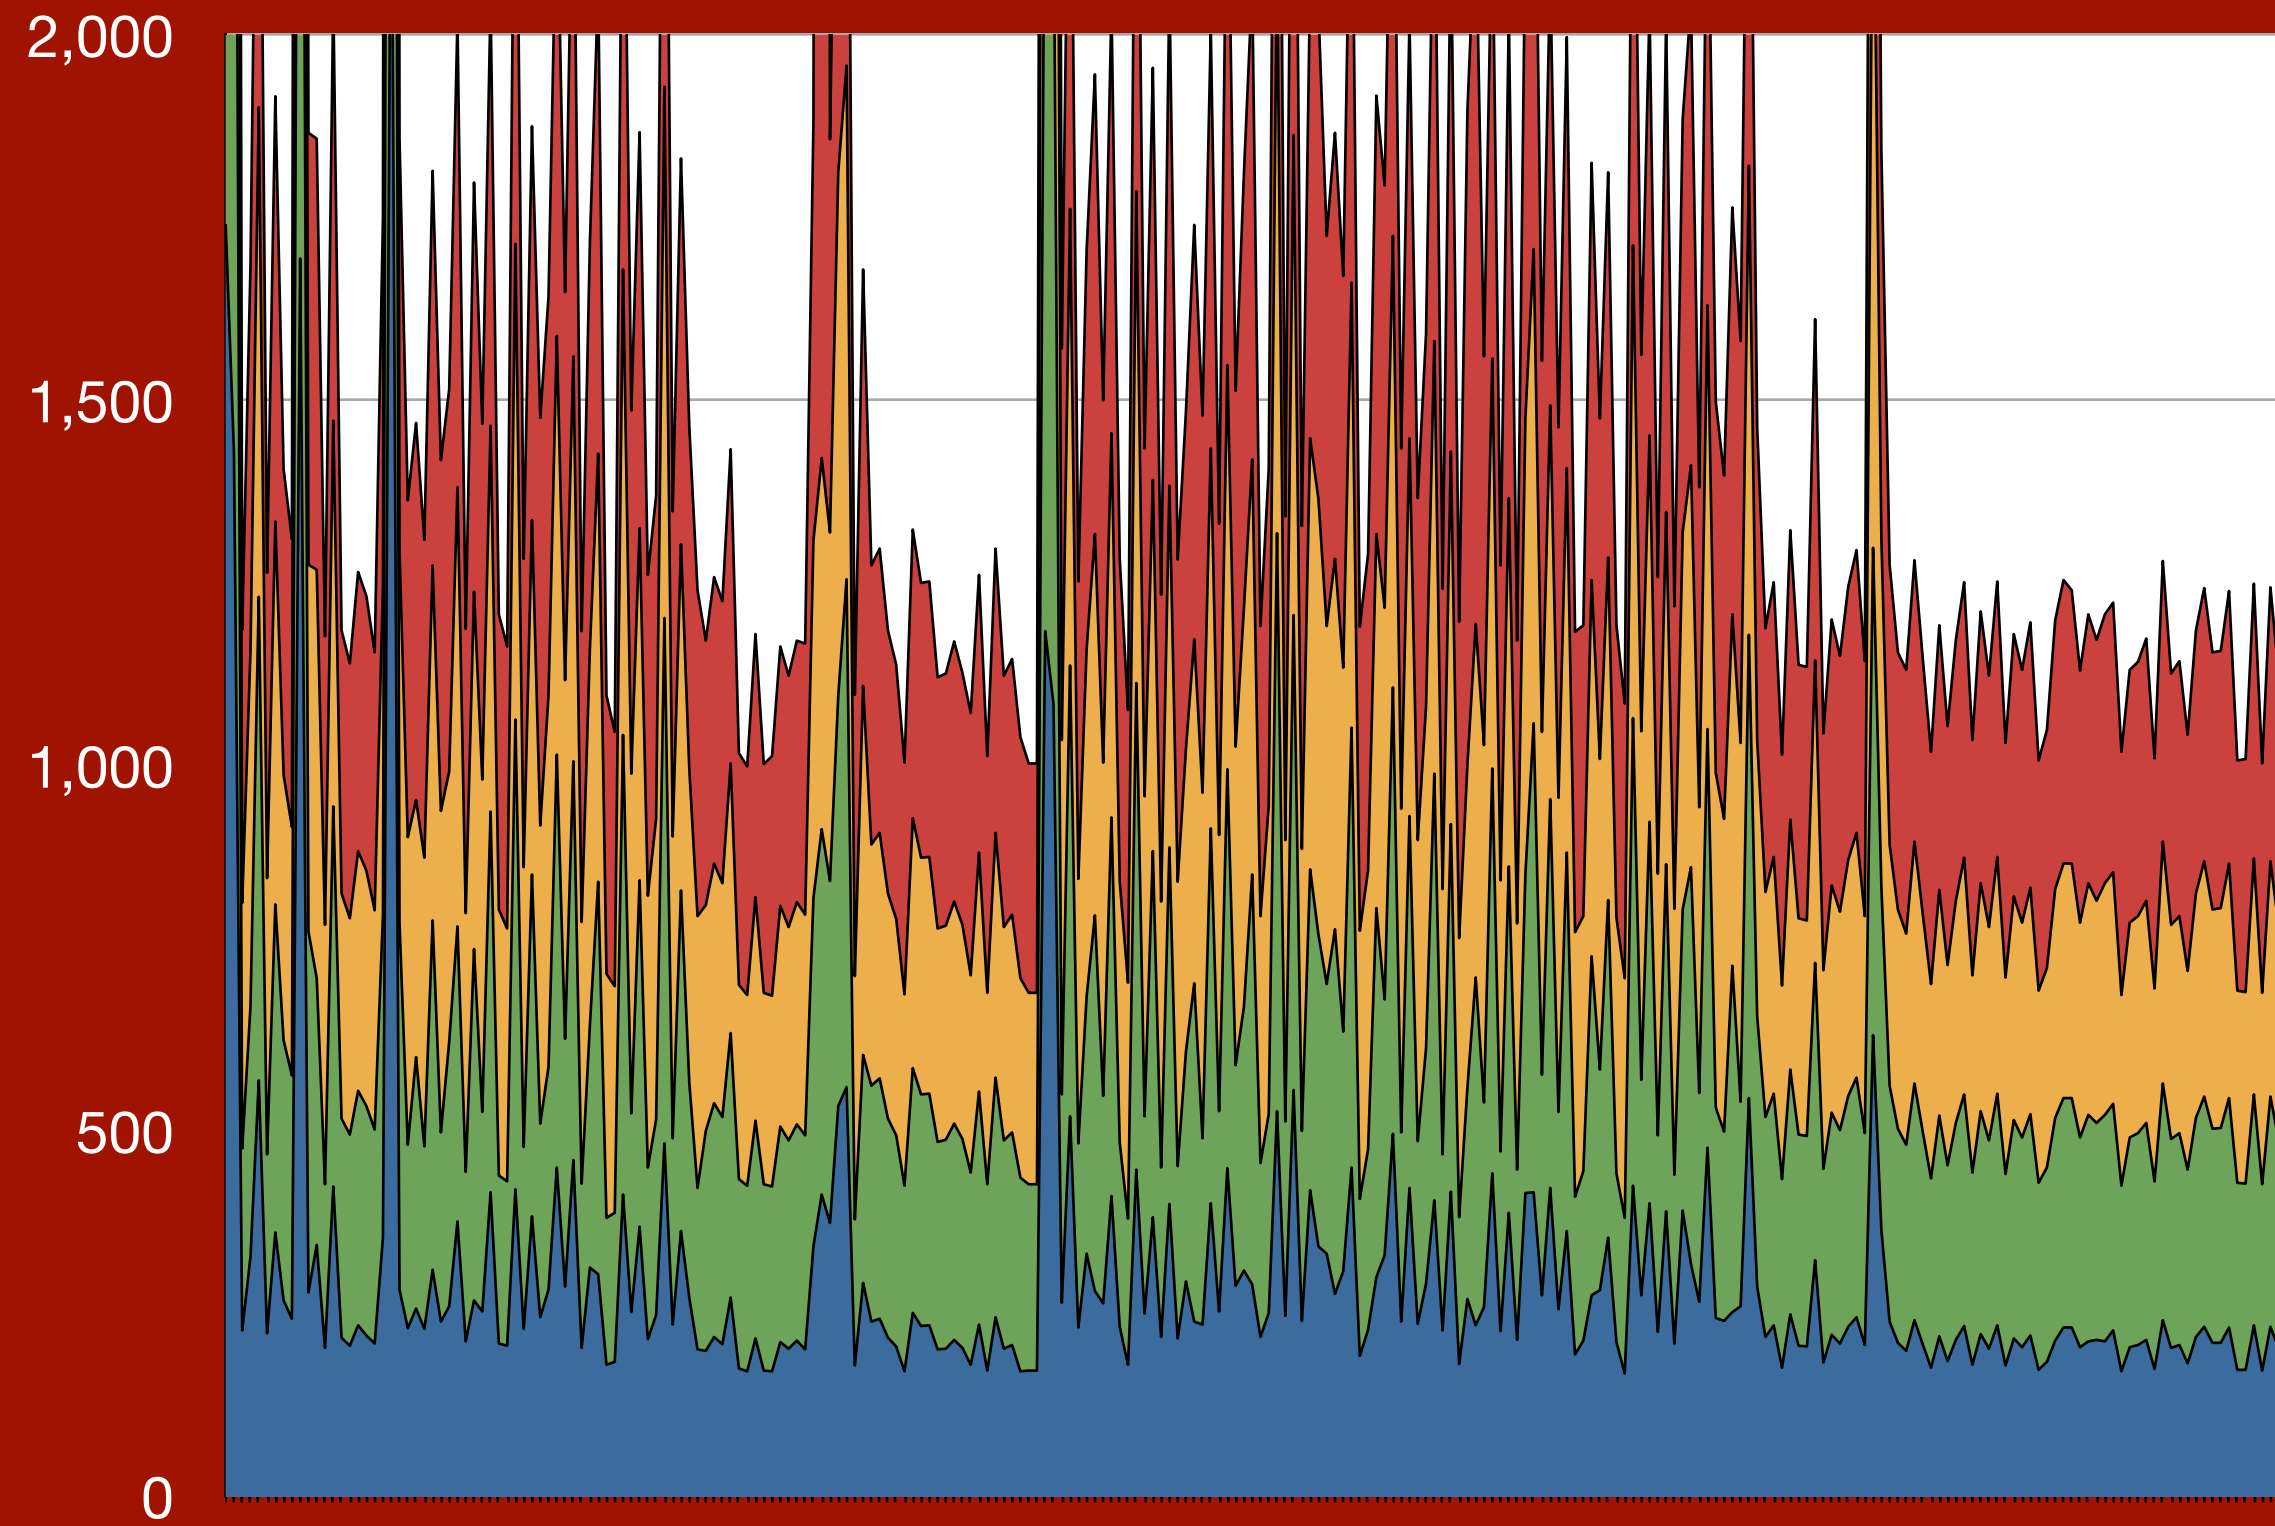
Apache - Worker MPM



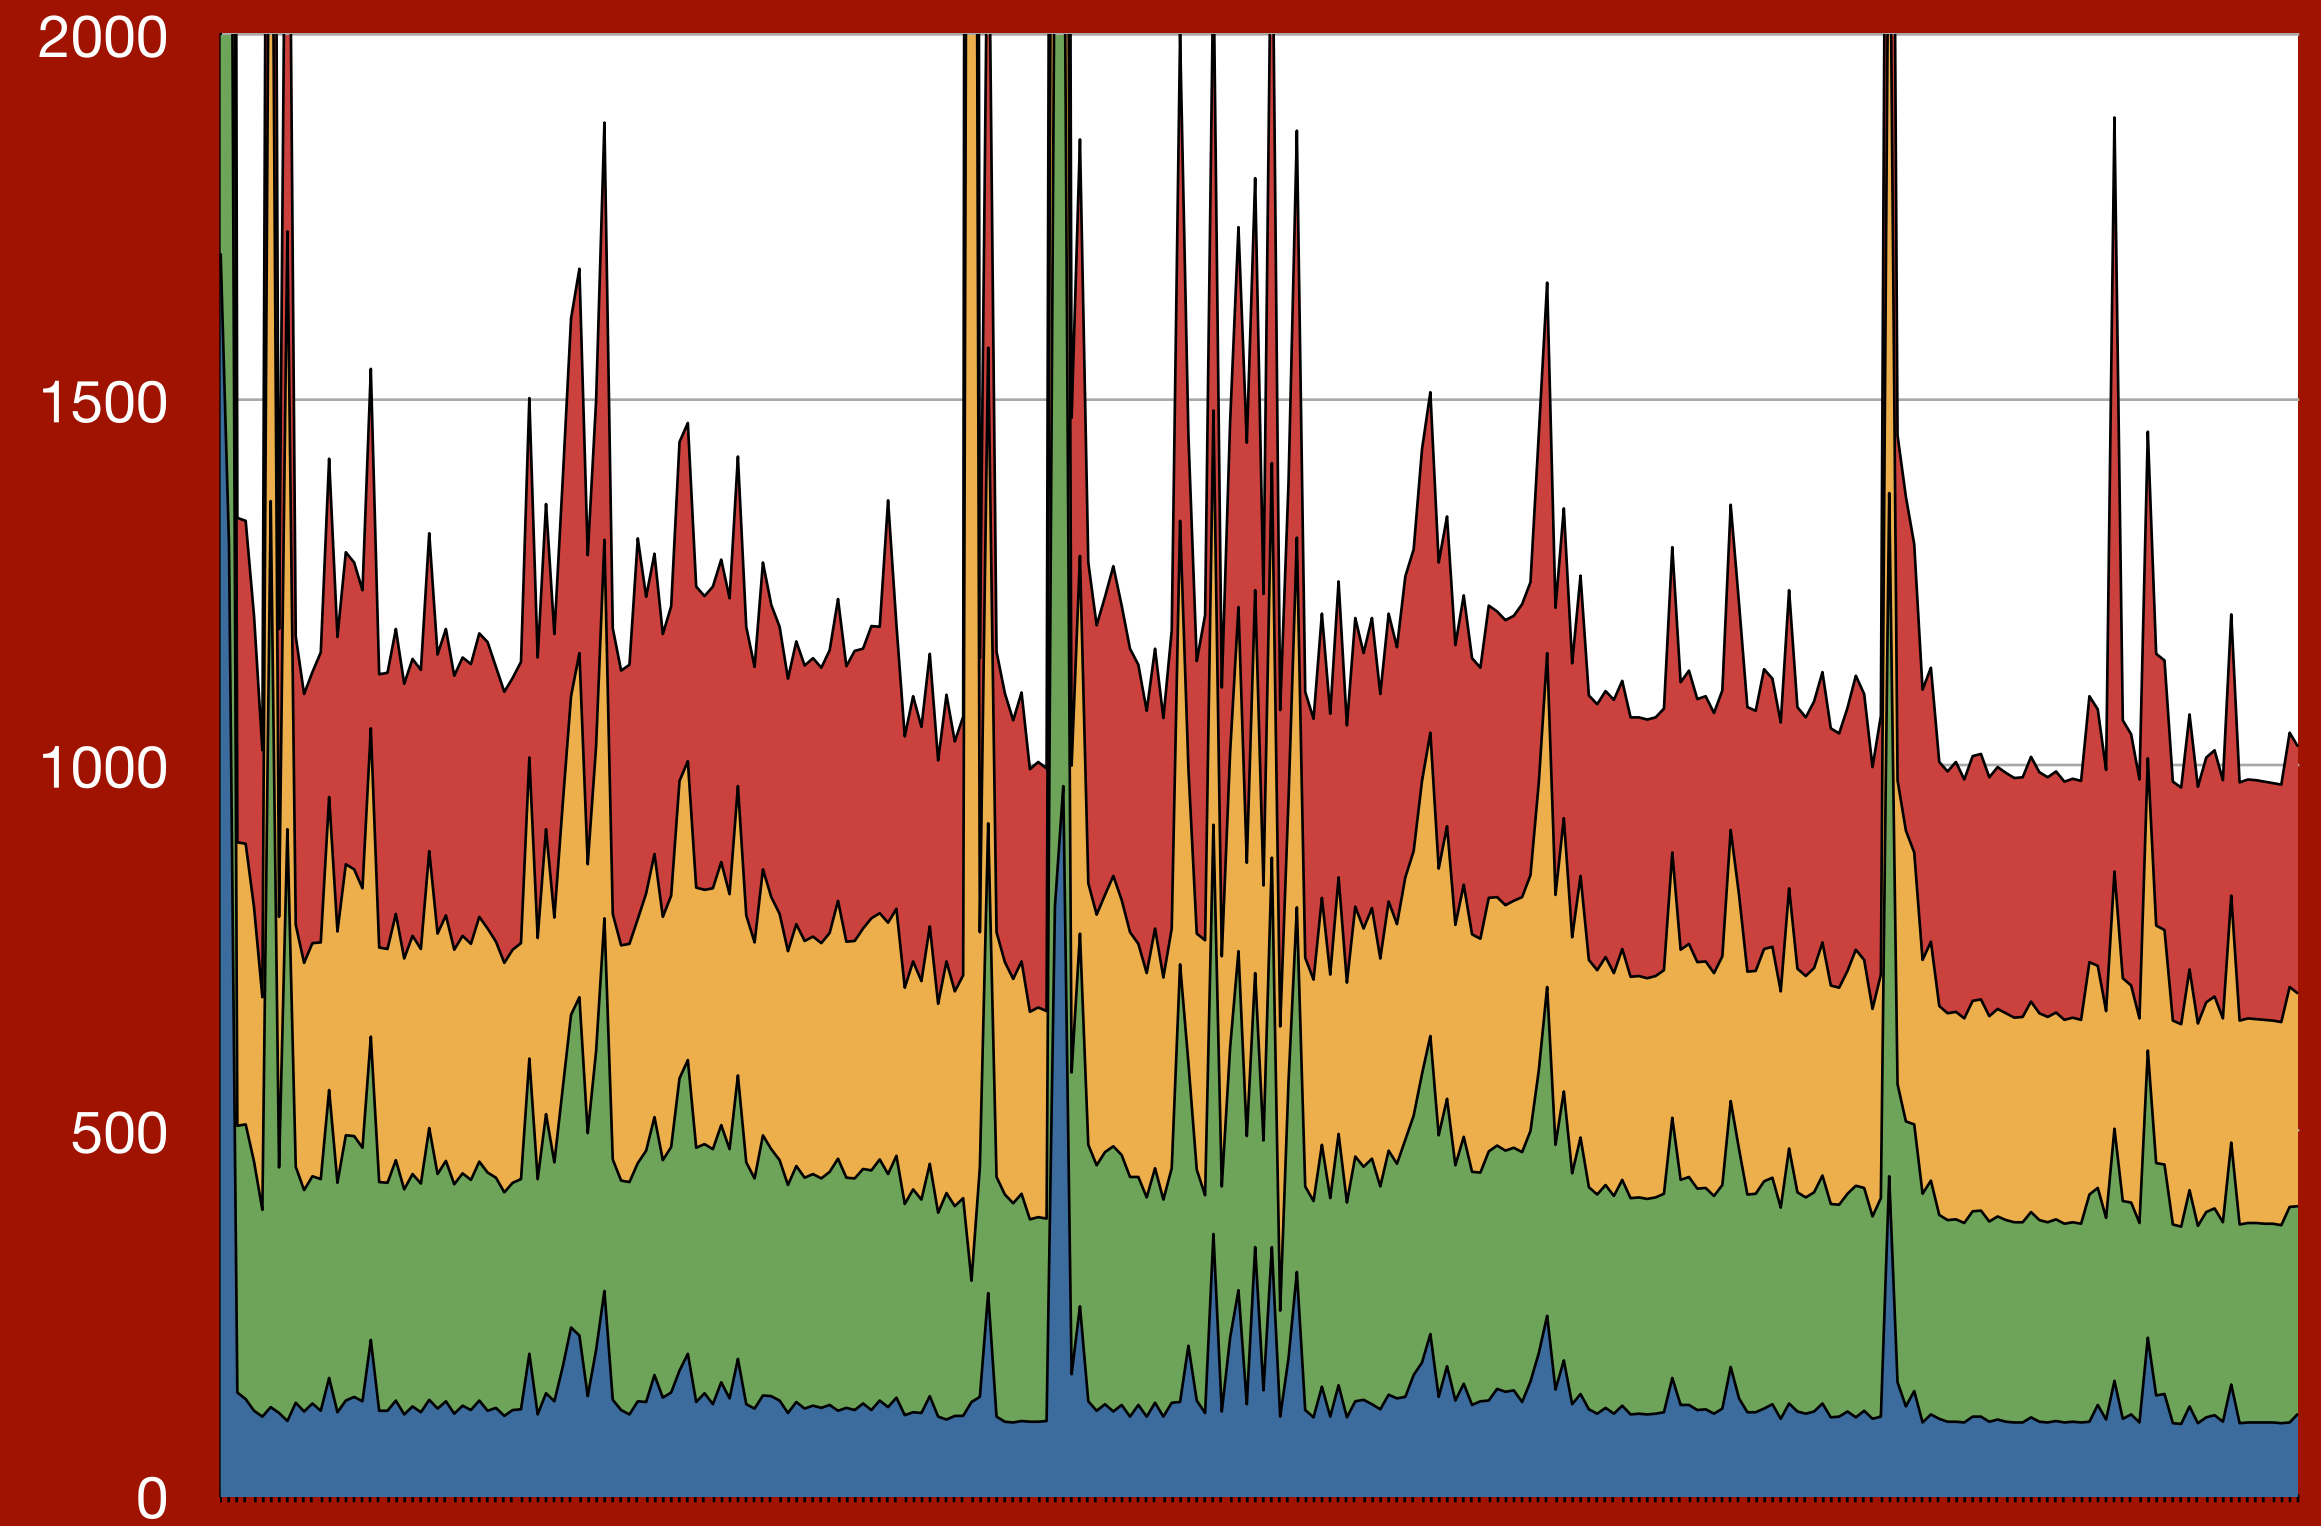
Open Write Read Close

# nginx vs Prefork (typical)

nginx



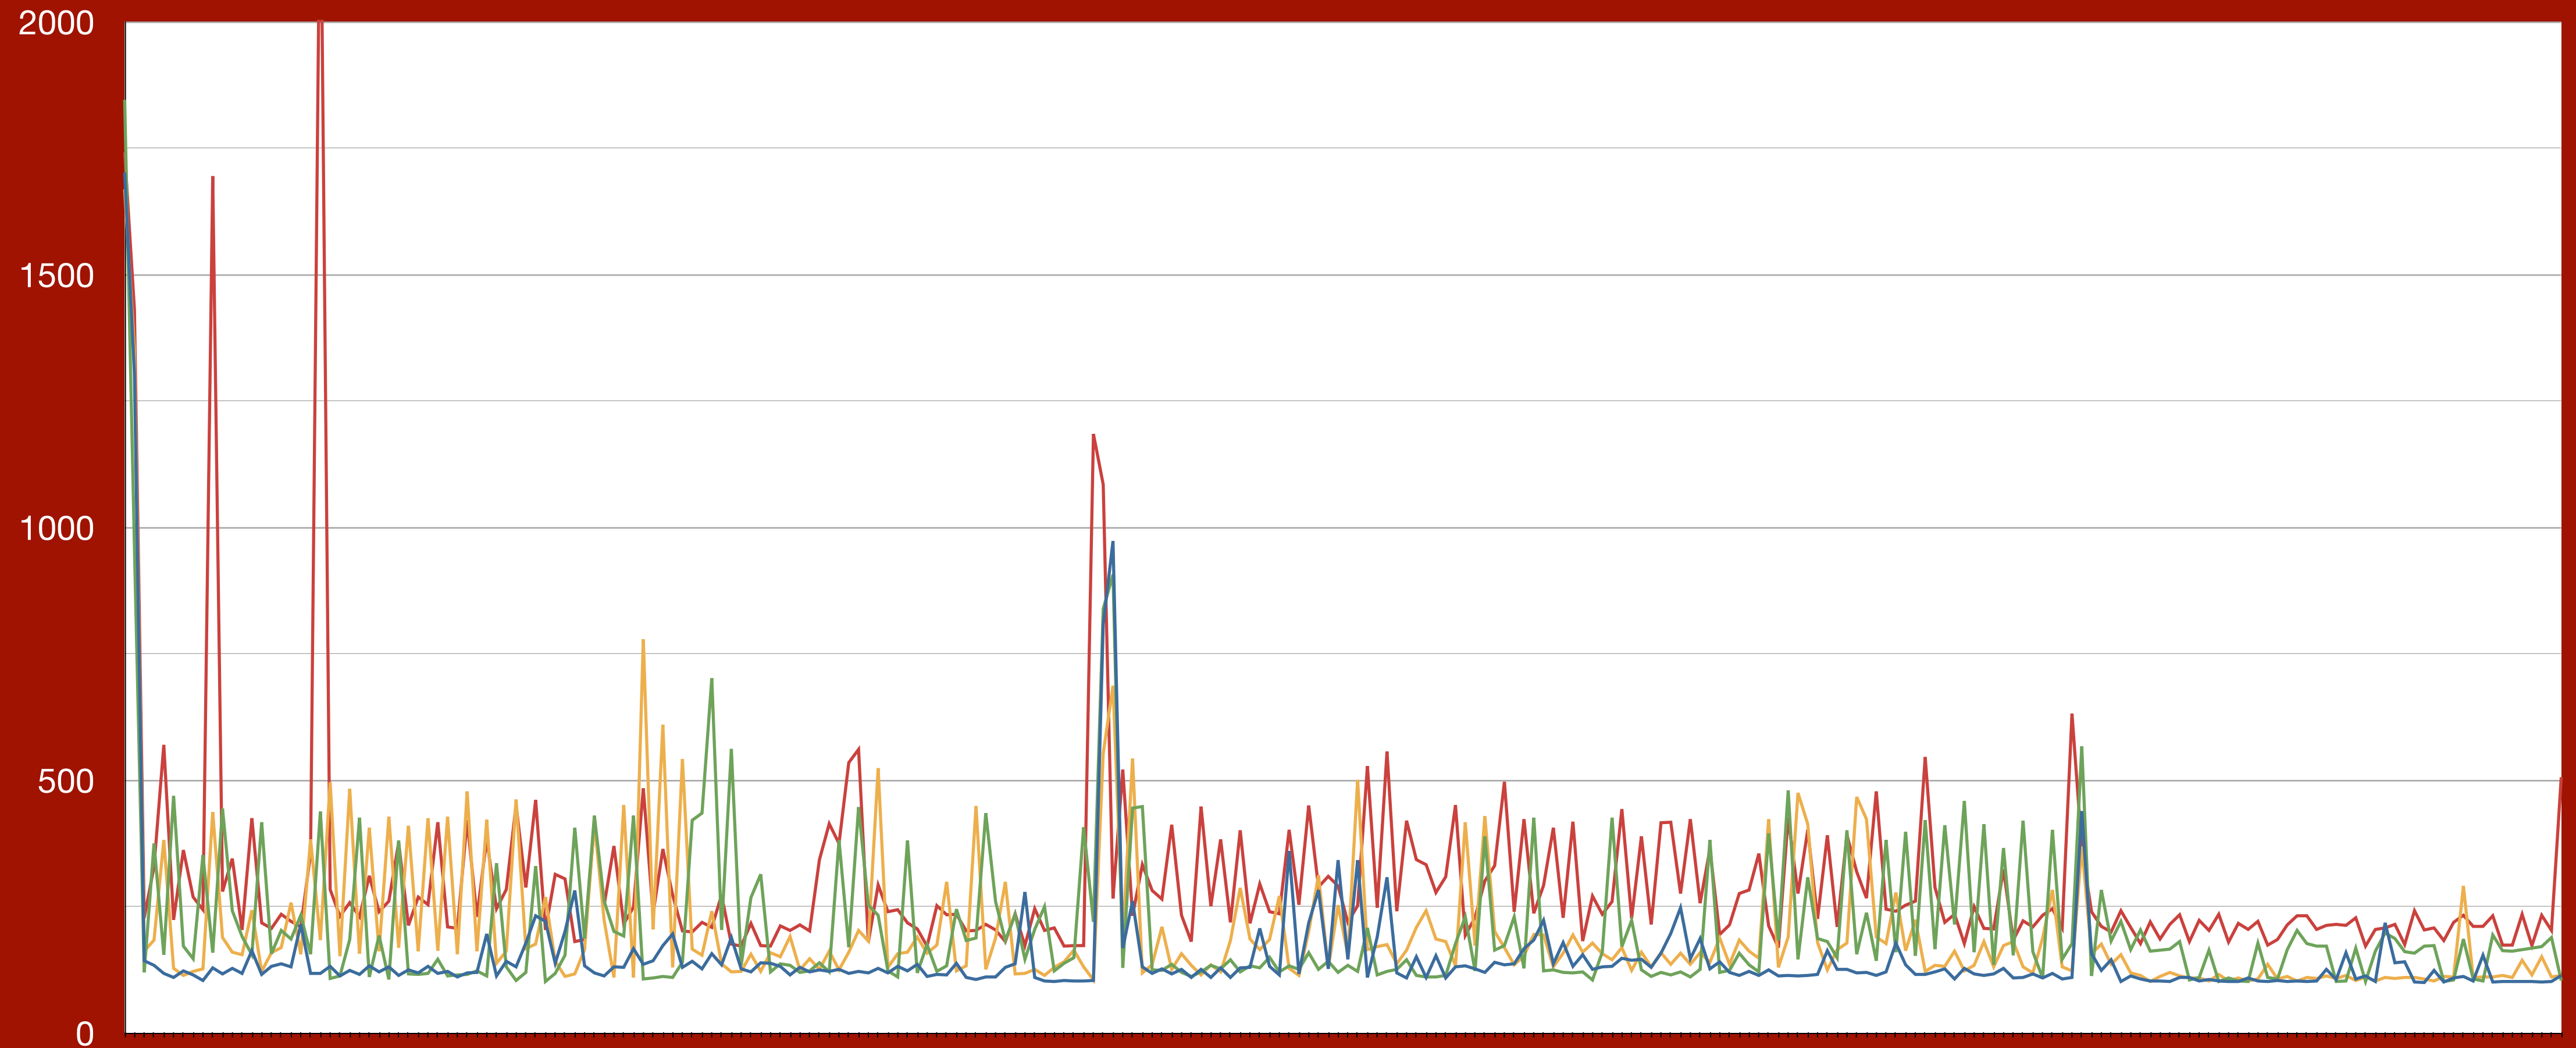
Apache - Prefork MPM



Open Write Read Close

# Focus on open()

Comparison - opens



— Prefork

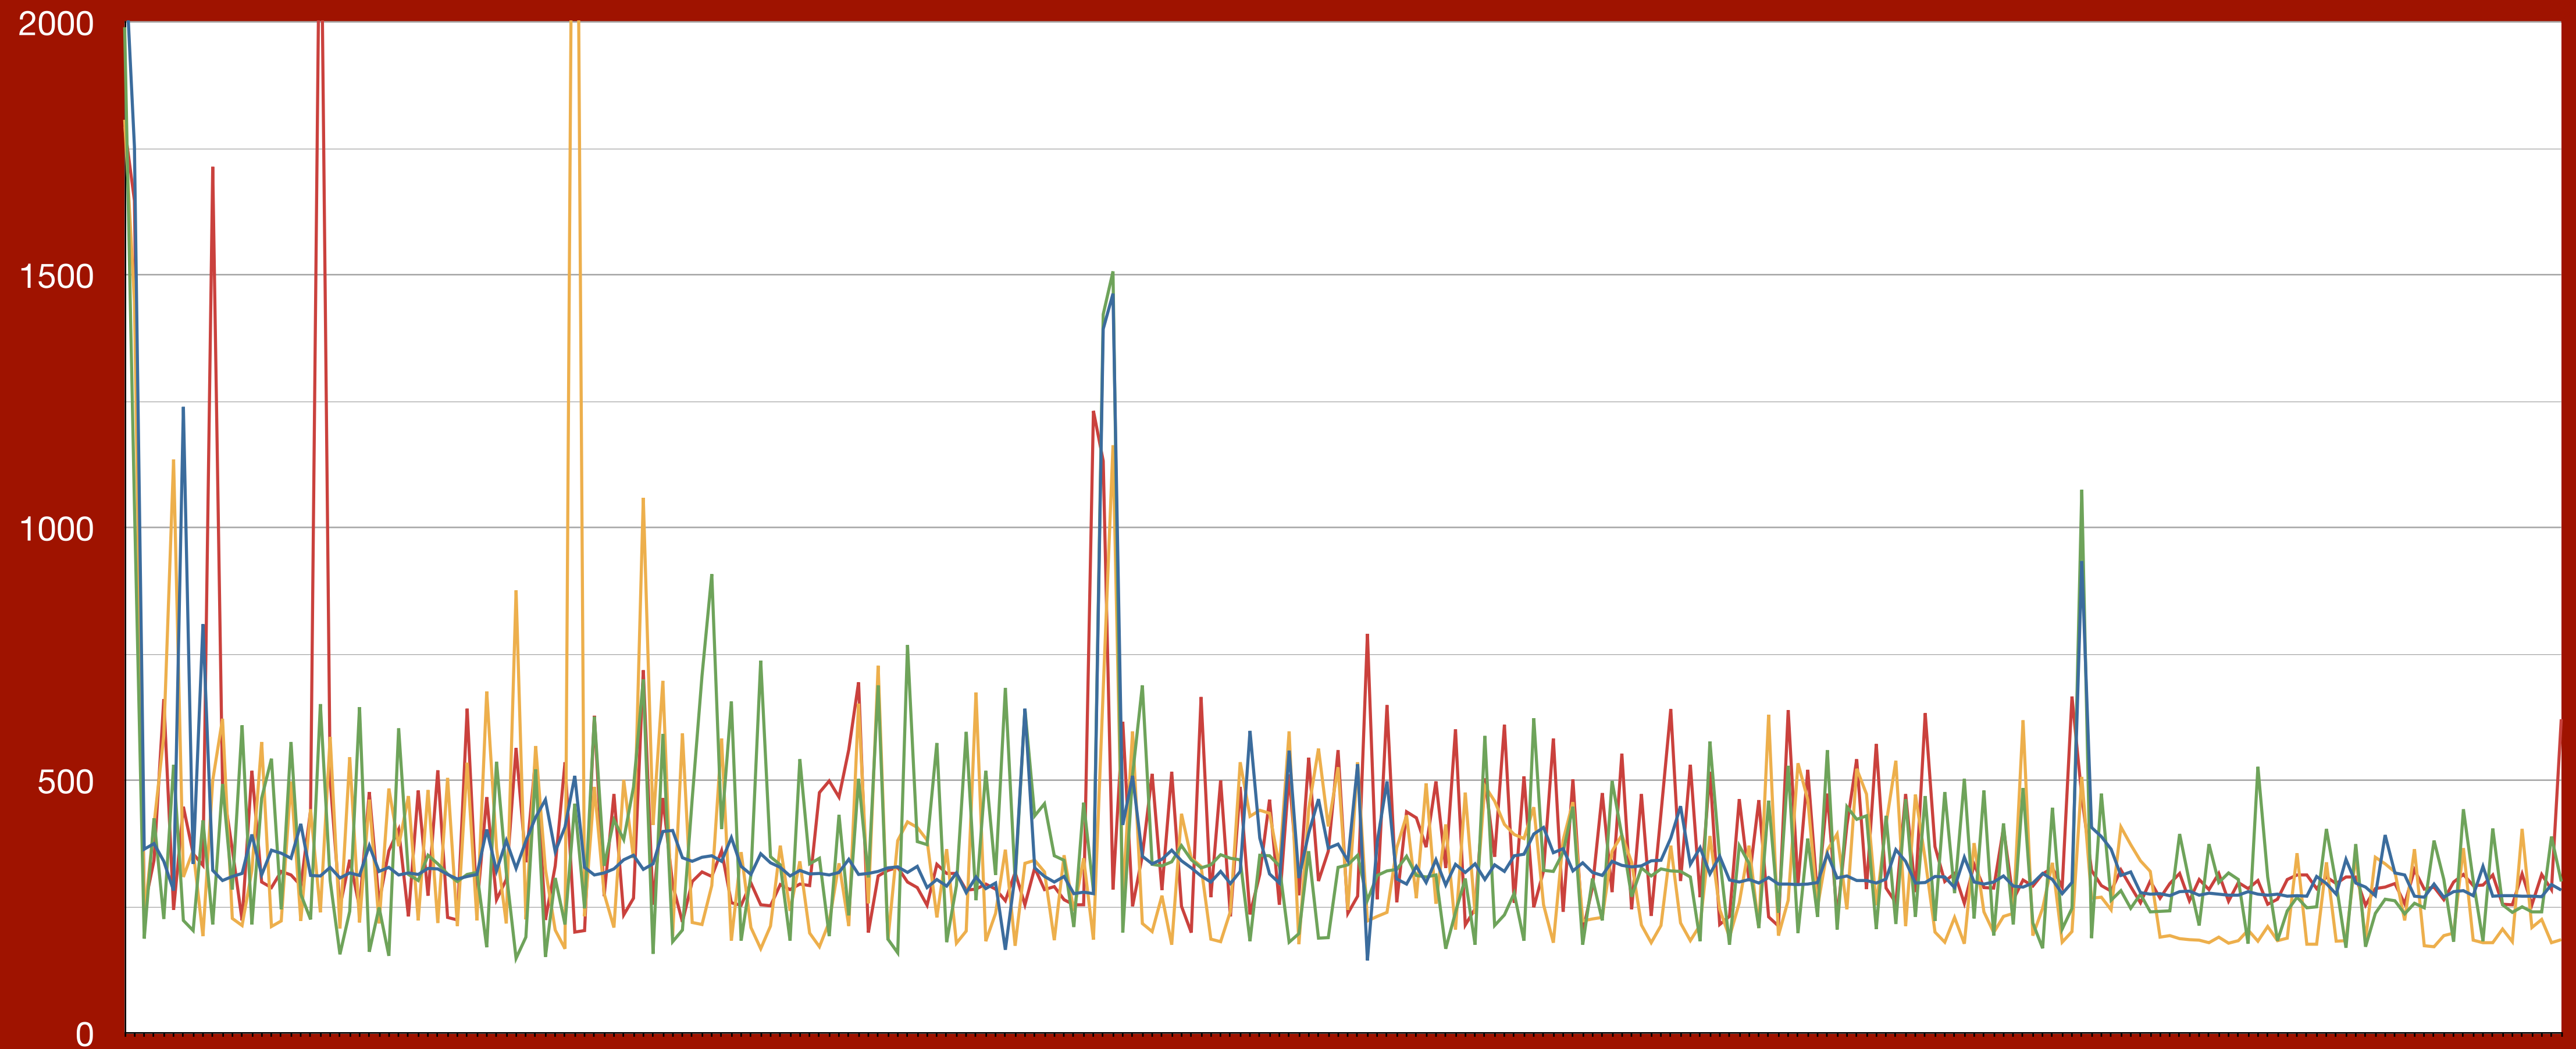
— Worker

— Event

— nginx

# Focus on write()

Comparison - writes



— Prefork

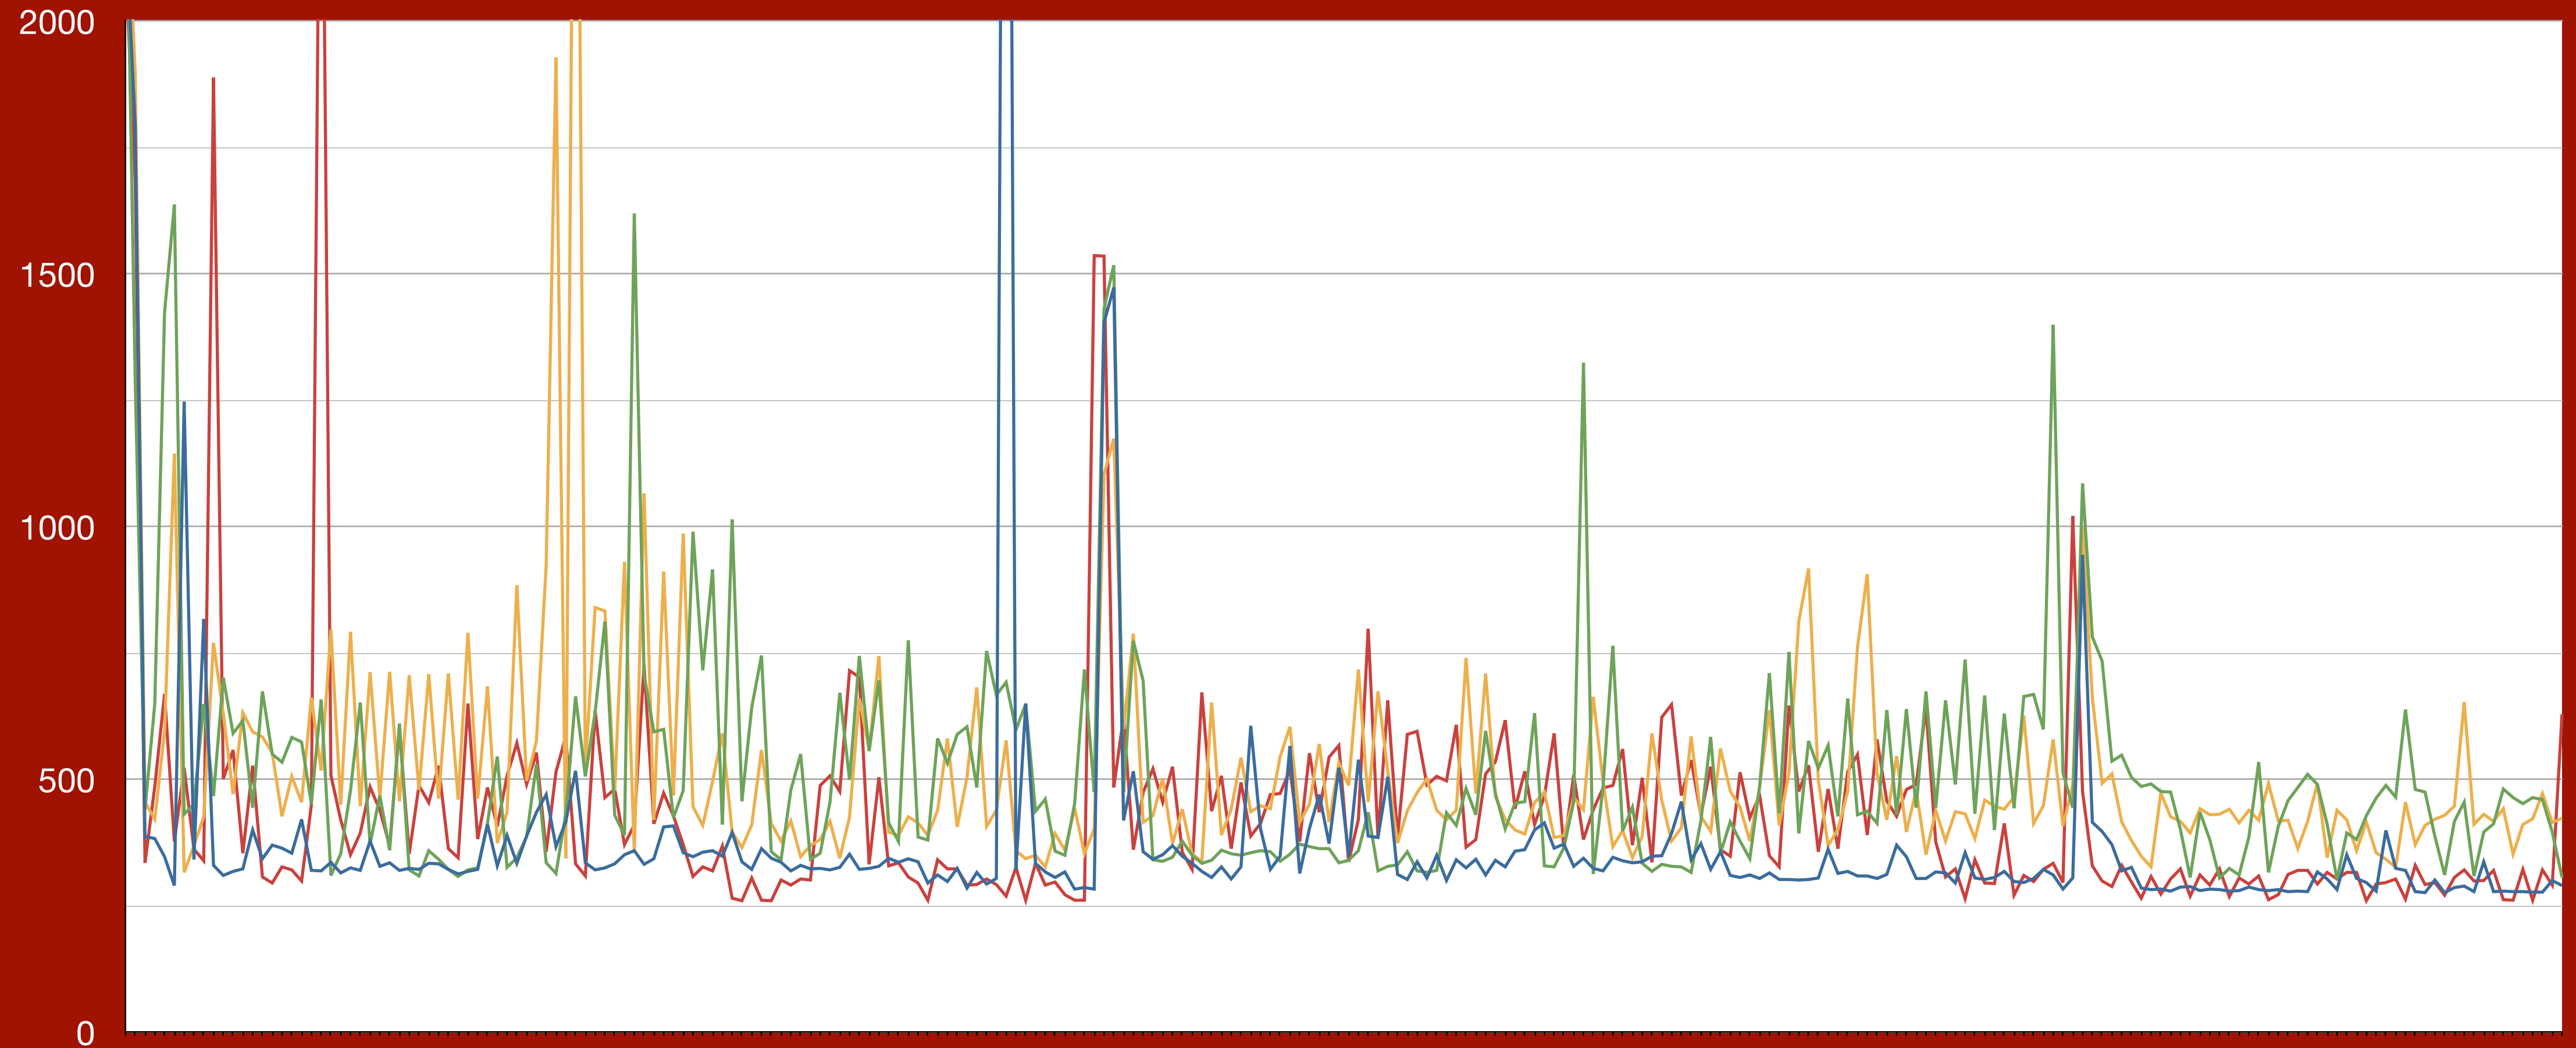
— Worker

— Event

— nginx

# Focus on read()

Comparison - reads



— Prefork

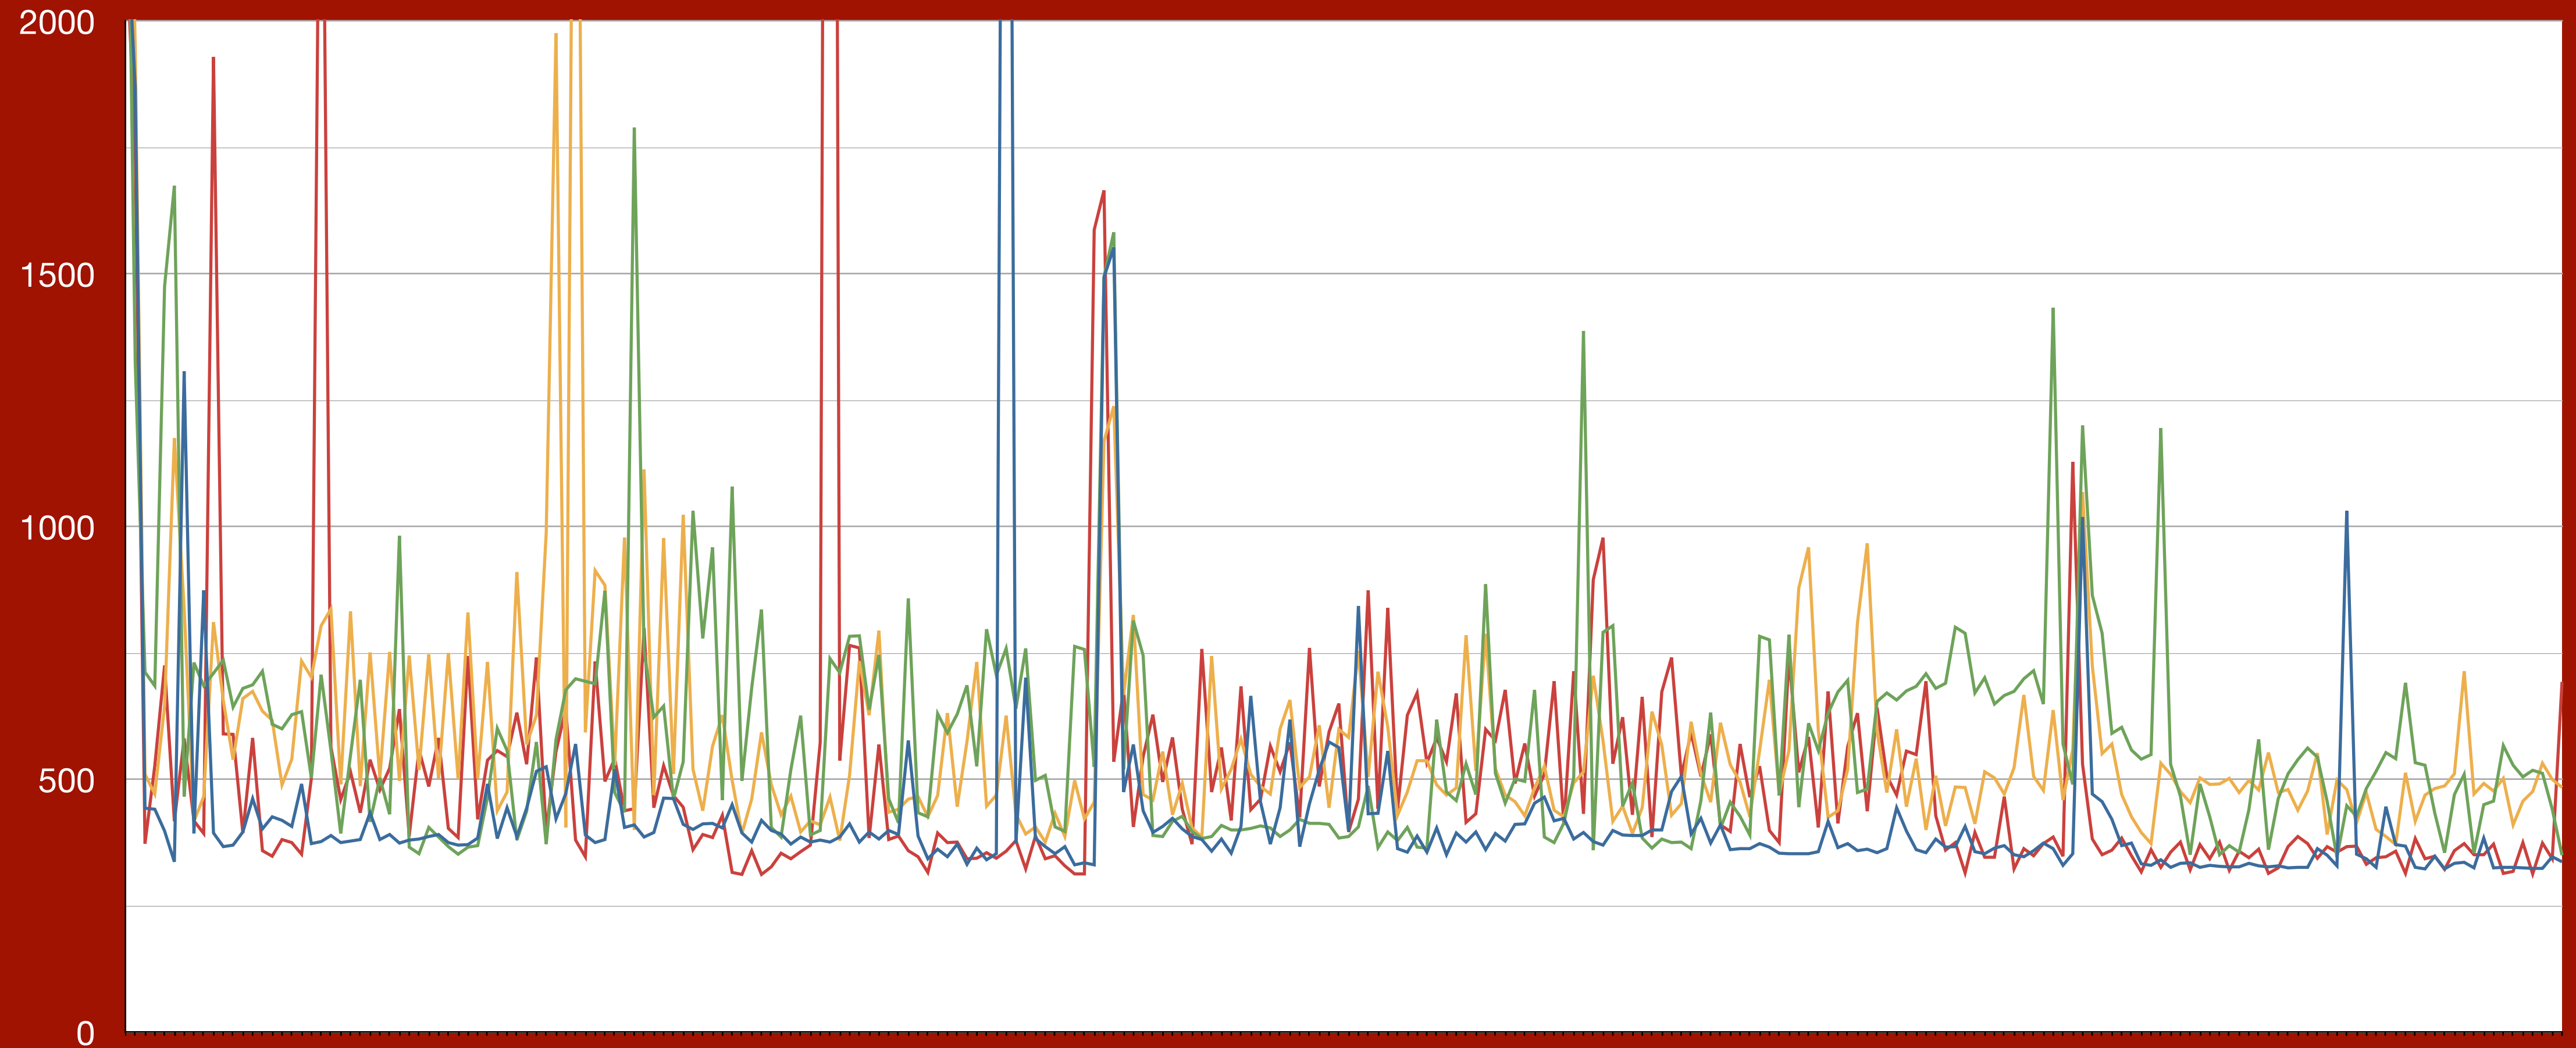
— Worker

— Event

— nginx

# Total req/resp time

Comparison - total transaction (close)



— Prefork

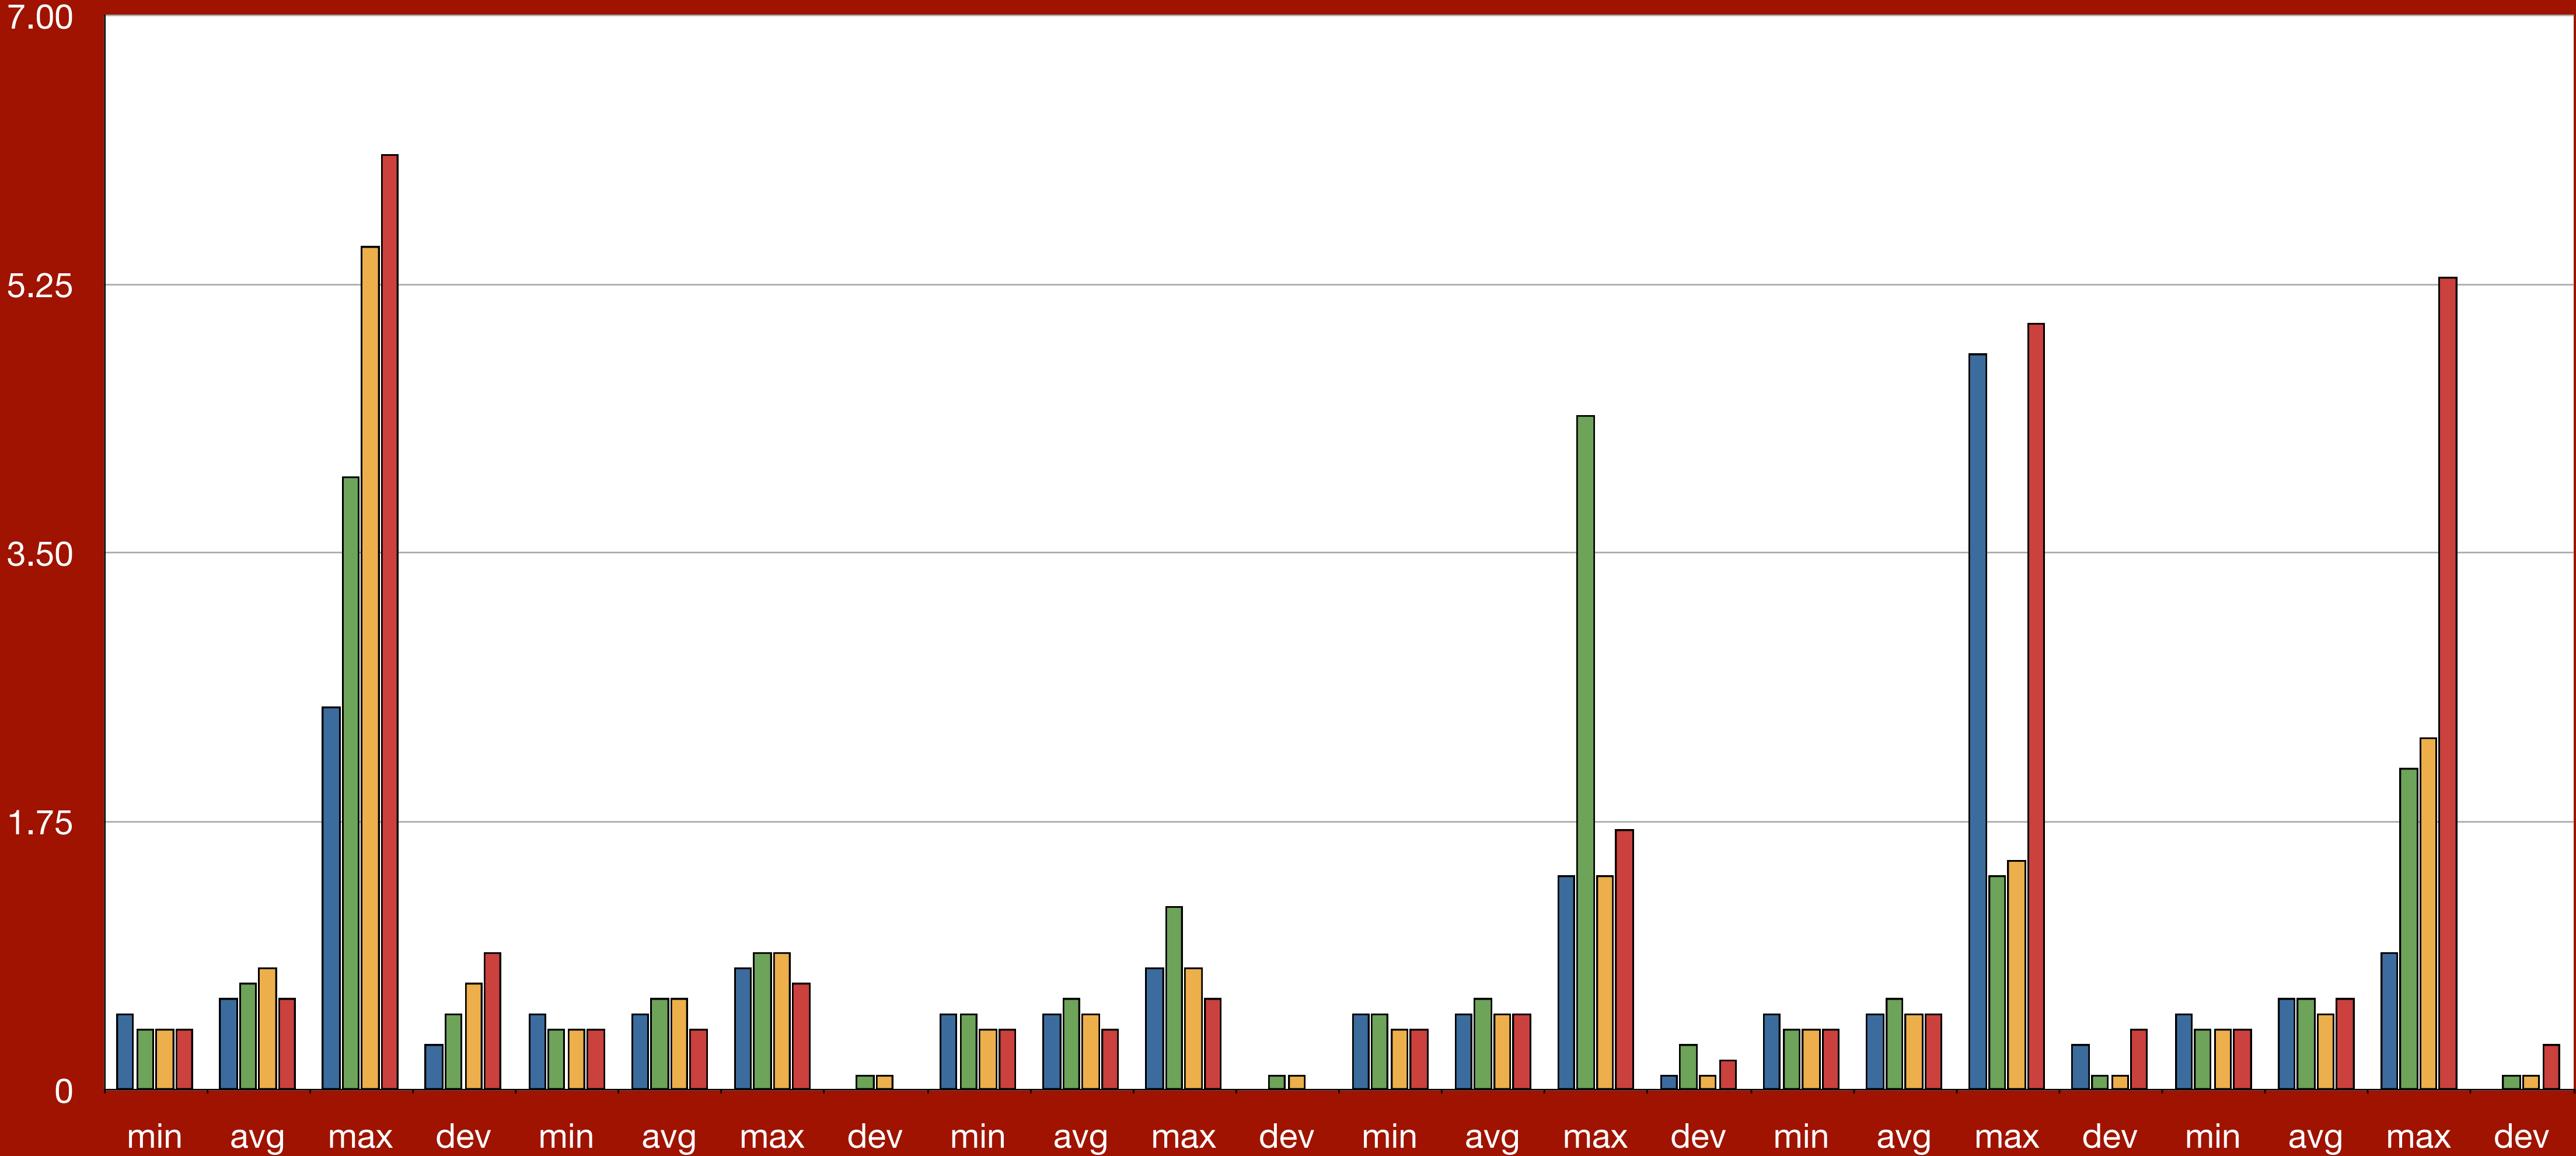
— Worker

— Event

— nginx

# Resp to Req. Bursts - httpref

100 ---> 10000



# Independent benchmark

表 1 テスト環境 (クライアント)	
Table 1 Experimentation Environment for Client	
PREV	クライアント
	Intel Core2Duo E8400 3.00GHz
Memory	4GB
NIC	Realtek RTL8111/8168B 1Gbps
OS	CentOS 5.6

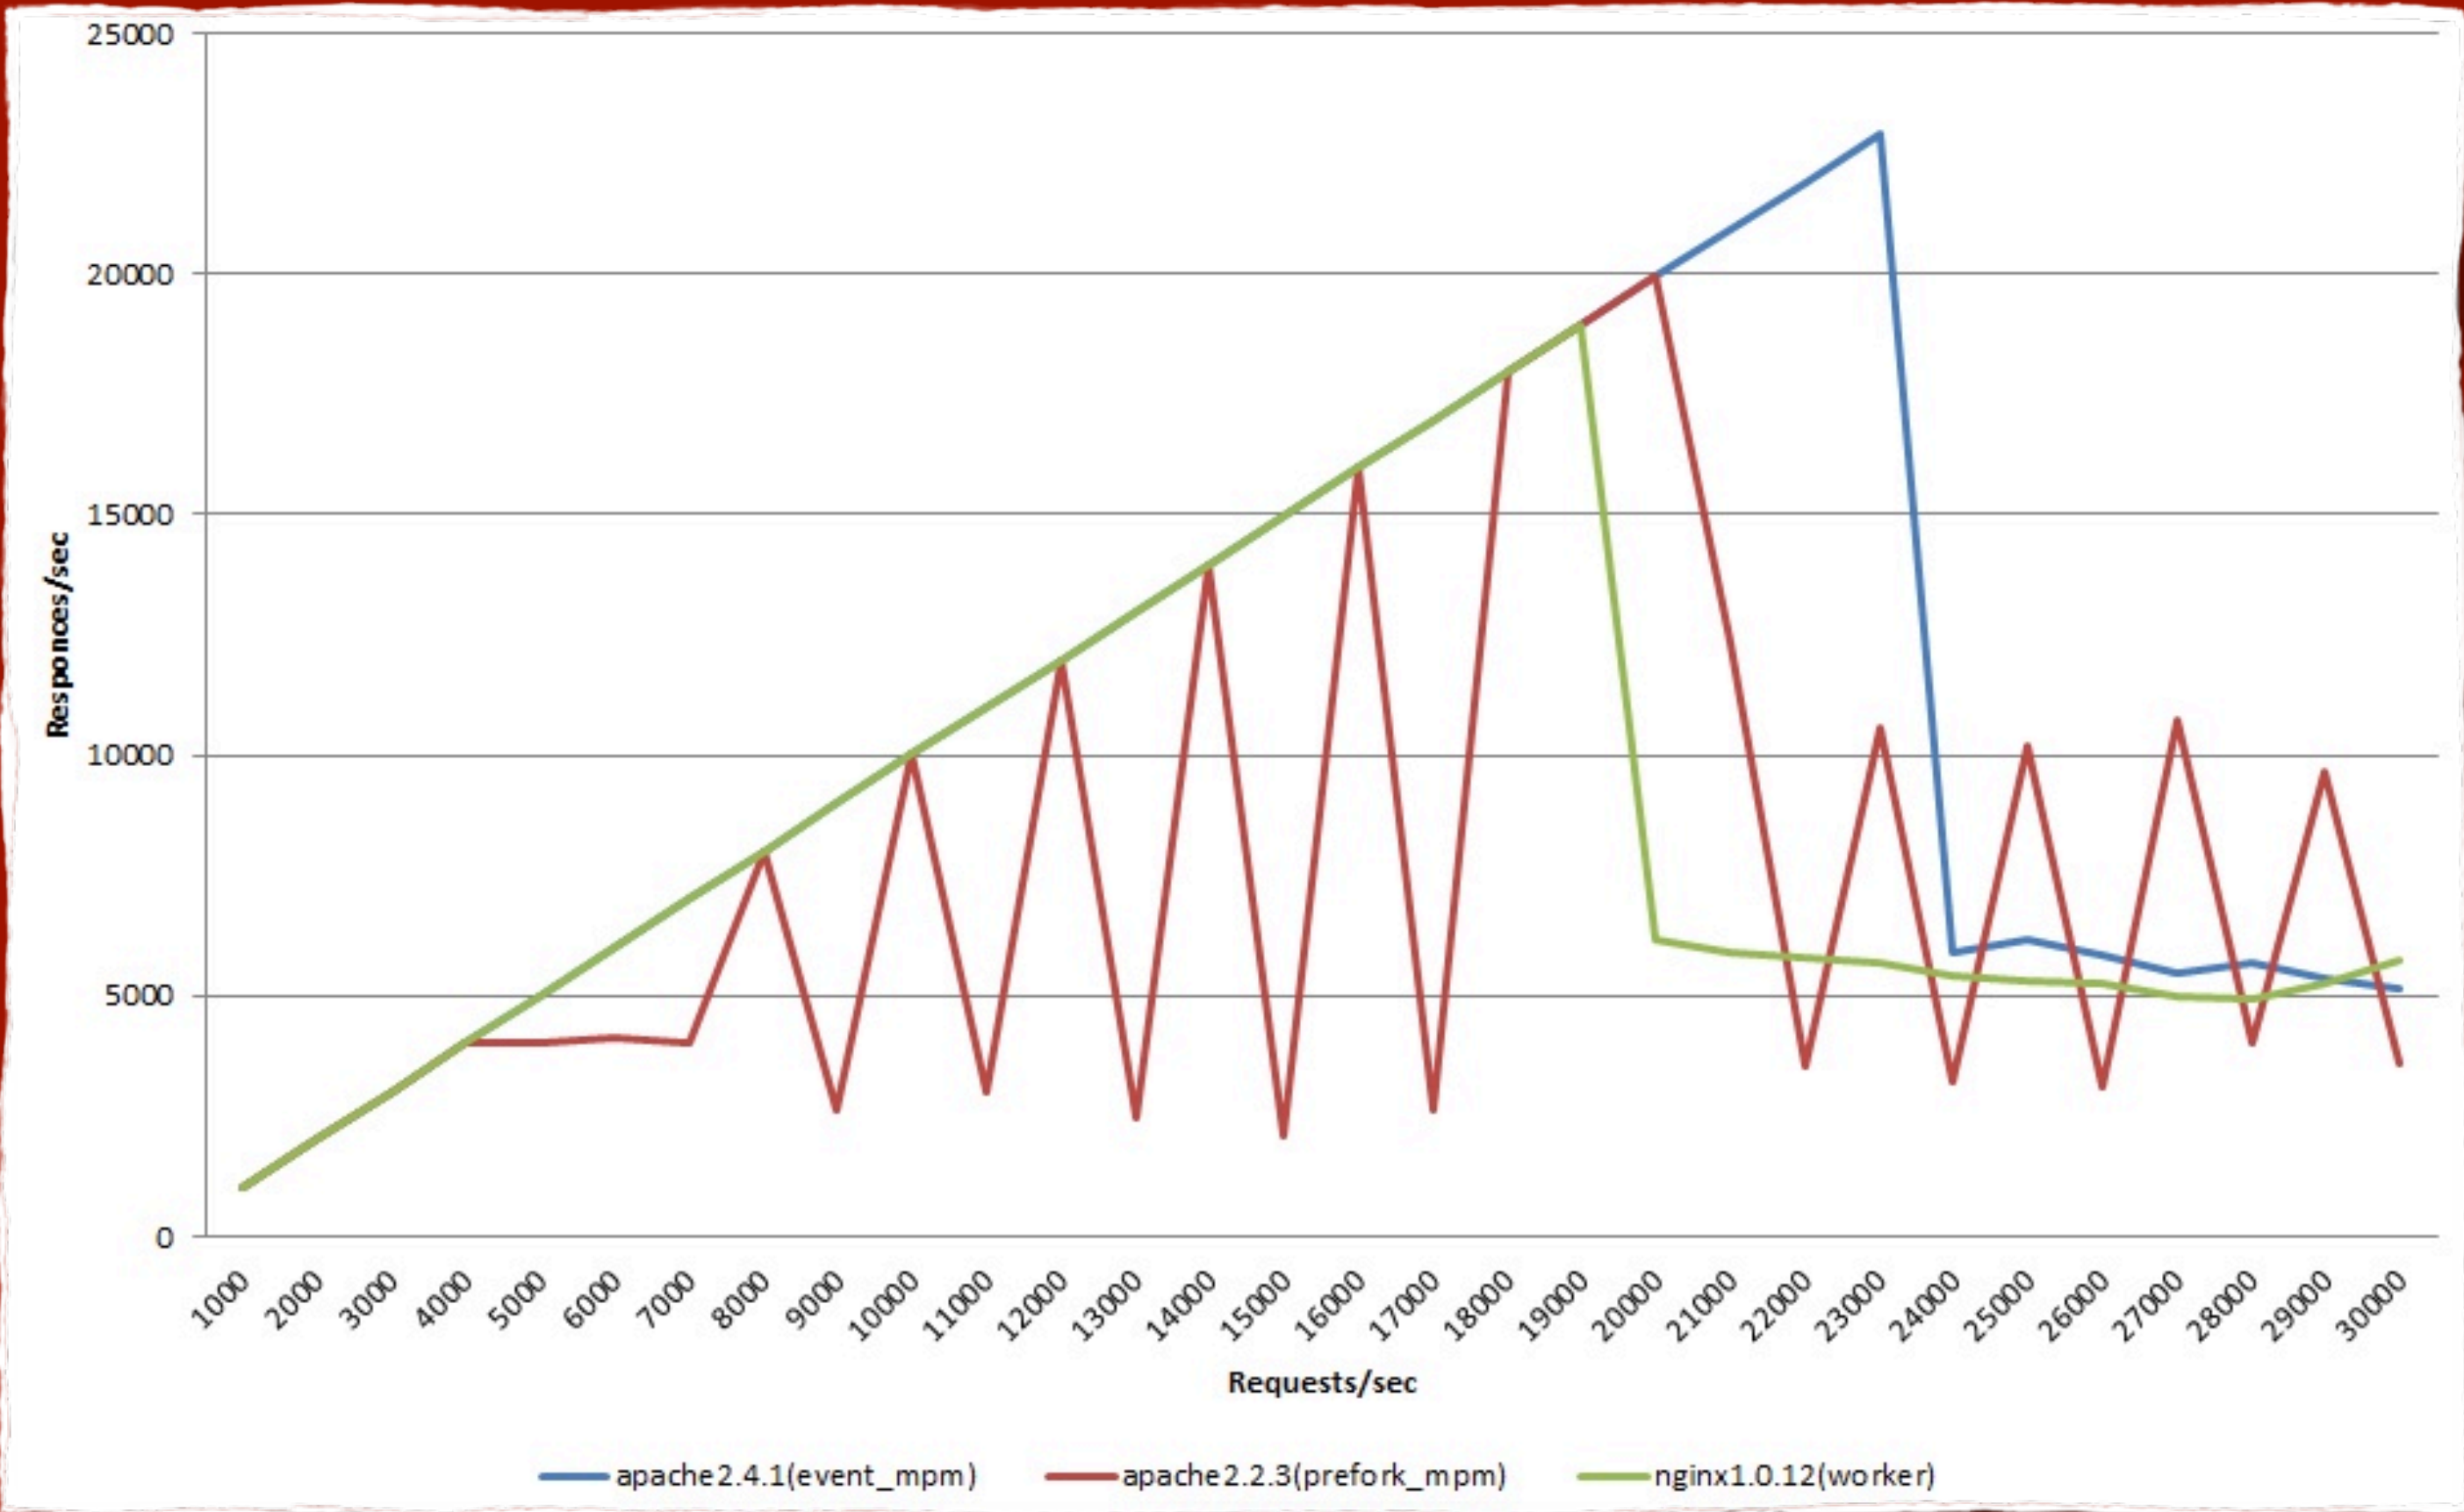
表 2 テスト環境 (サーバ)	
Table 2 Experimentation Environment for Server	
	サーバ
CPU	Intel Xeon X5355 2.66GHz
Memory	8GB
NIC	Broadcom BCM5708 1Gbps
OS	CentOS 5.6

Image 2 of 4

CLOSE X

```
#!/bin/sh
RESULT='./result.txt'

for port in 80 8080 8888
do
    #for count in 1000 2000 3000 4000 5000 6000 7000 8000 9000 10000
    #for count in 11000 12000 13000 14000 15000 16000 17000 18000 19000 20000
    for count in 21000 22000 23000 24000 25000 26000 27000 28000 29000 30000
    do
        echo -n "$port $count " >> $RESULT
        httpperf --rate $count --num-conns 25000 --server ipaddr --port $port \
            --uri=/test.html | grep "Request rate:" >> $RESULT.$port
        sleep 60
    done
done
```



Source: Ryosuke Matsumoto : <http://blog.matsumoto-r.jp/?p=1812>

# Benchmark Conclusions

- Events, polling and fork/spawn creates overhead: good for “more bang for buck” system, bad for performance for that request
- For concurrency, Event & Worker on par with nginx\*
- For transaction speed, prefork shines
- Let's reboot “Simple” mpm (*currently being done*)
- \*Main Caveats:
  - Apache is never resource starved
  - If memory is a scarce resource, nginx still better (*for now ;)* )
  - More work can (and should) be done

# In conclusion...

- Performance of Apache httpd 2.4 still in the big leagues (and on par with the “big boys” and the fanboi webserver du jure)
- For cloud environs, the performance and dynamic control of Apache httpd 2.4 in reverse proxies is just what the Dr. ordered (and flexibility remains a big strength)
- Architecture of Apache httpd 2.4 allows a lot of room for growth and additional functionality (both for the cloud and not)
- There’s still a category of “edge cases” that require nginx, lighttpd, G-WAN, Apache Traffic Server, etc... If that’s you, don’t try to use Apache httpd (but if you do, provide patches!)
- lies, damned lies and benchmarks (sorry, statistics).

# Thanks!

- Contact Info:

- Jim Jagielski
- [jim@jaguNET.com](mailto:jim@jaguNET.com)      [jimjag@redhat.com](mailto:jimjag@redhat.com)
- [@jimjag](#)      [www.jimjag.com](http://www.jimjag.com)
- [people.apache.org/~jim/presos/](http://people.apache.org/~jim/presos/)