

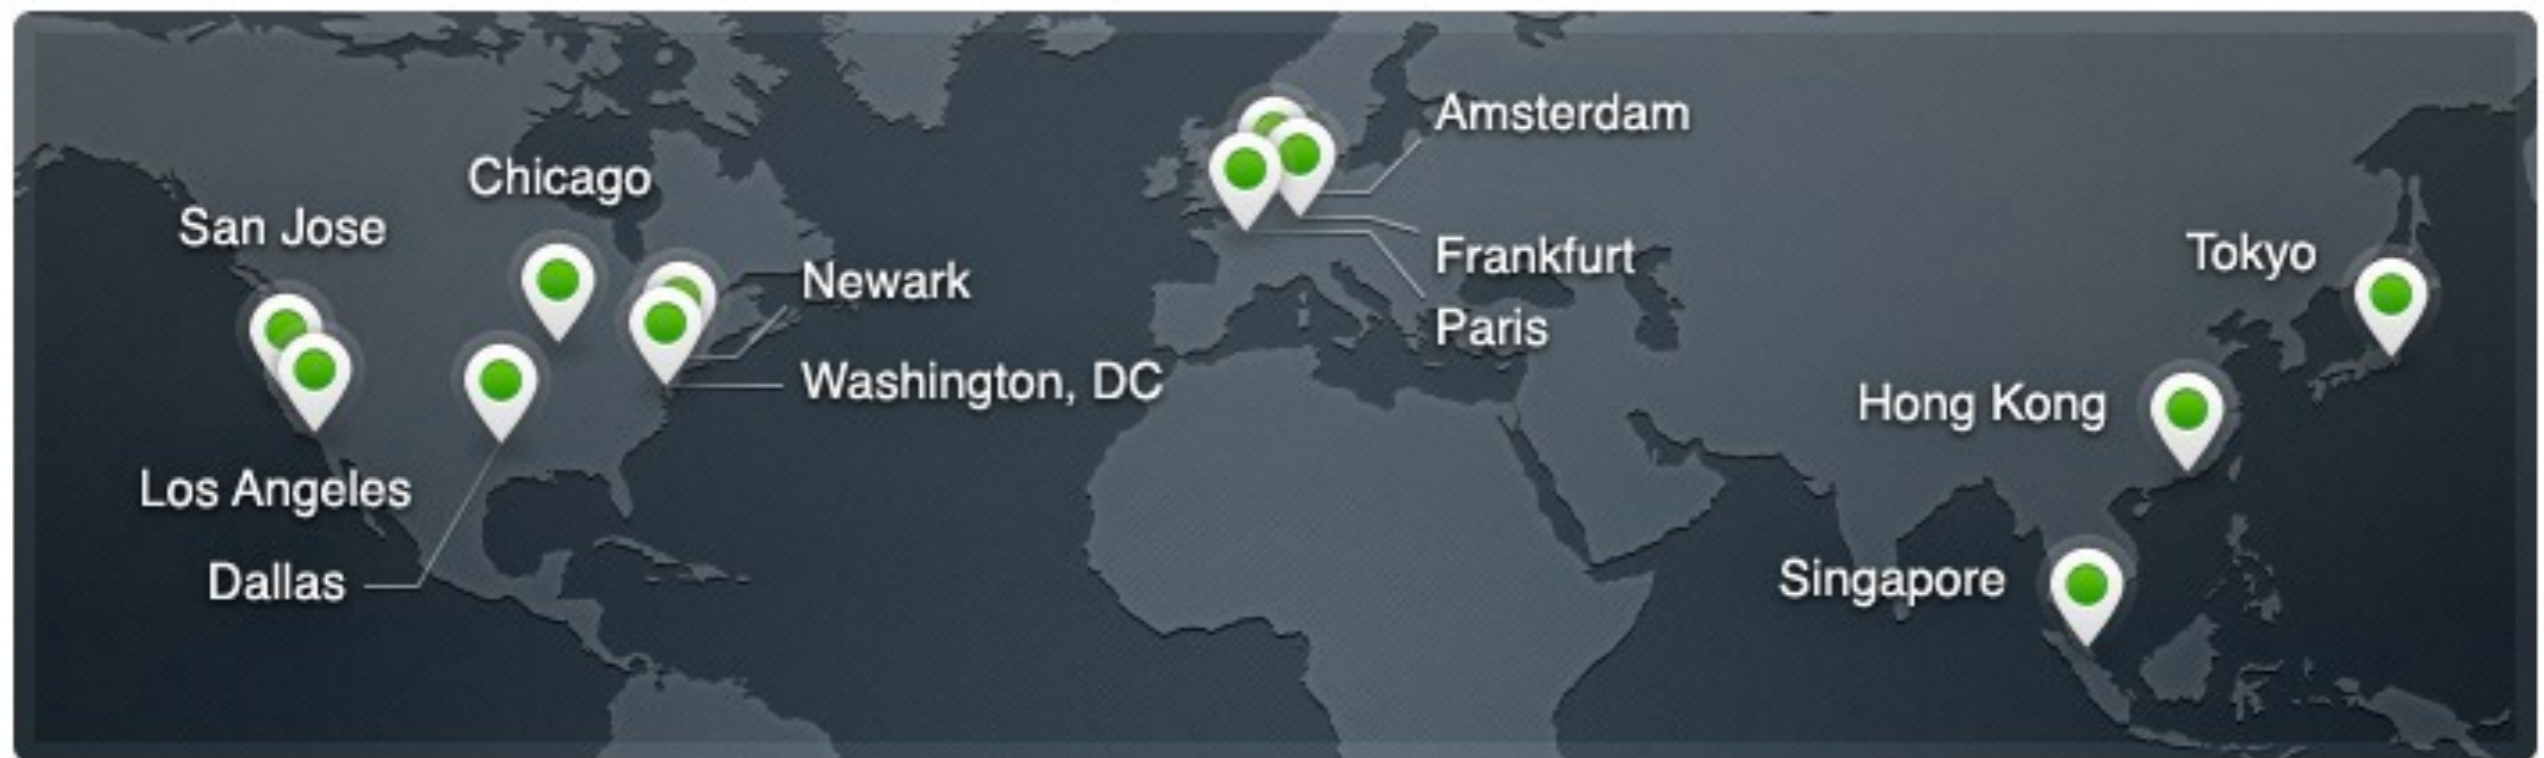
Locks, Deadlocks and Abstractions: Experiences with Multi-Threaded Programming at CloudFlare, Inc.

Ian Pye
ian@cloudflare.com

Roadmap

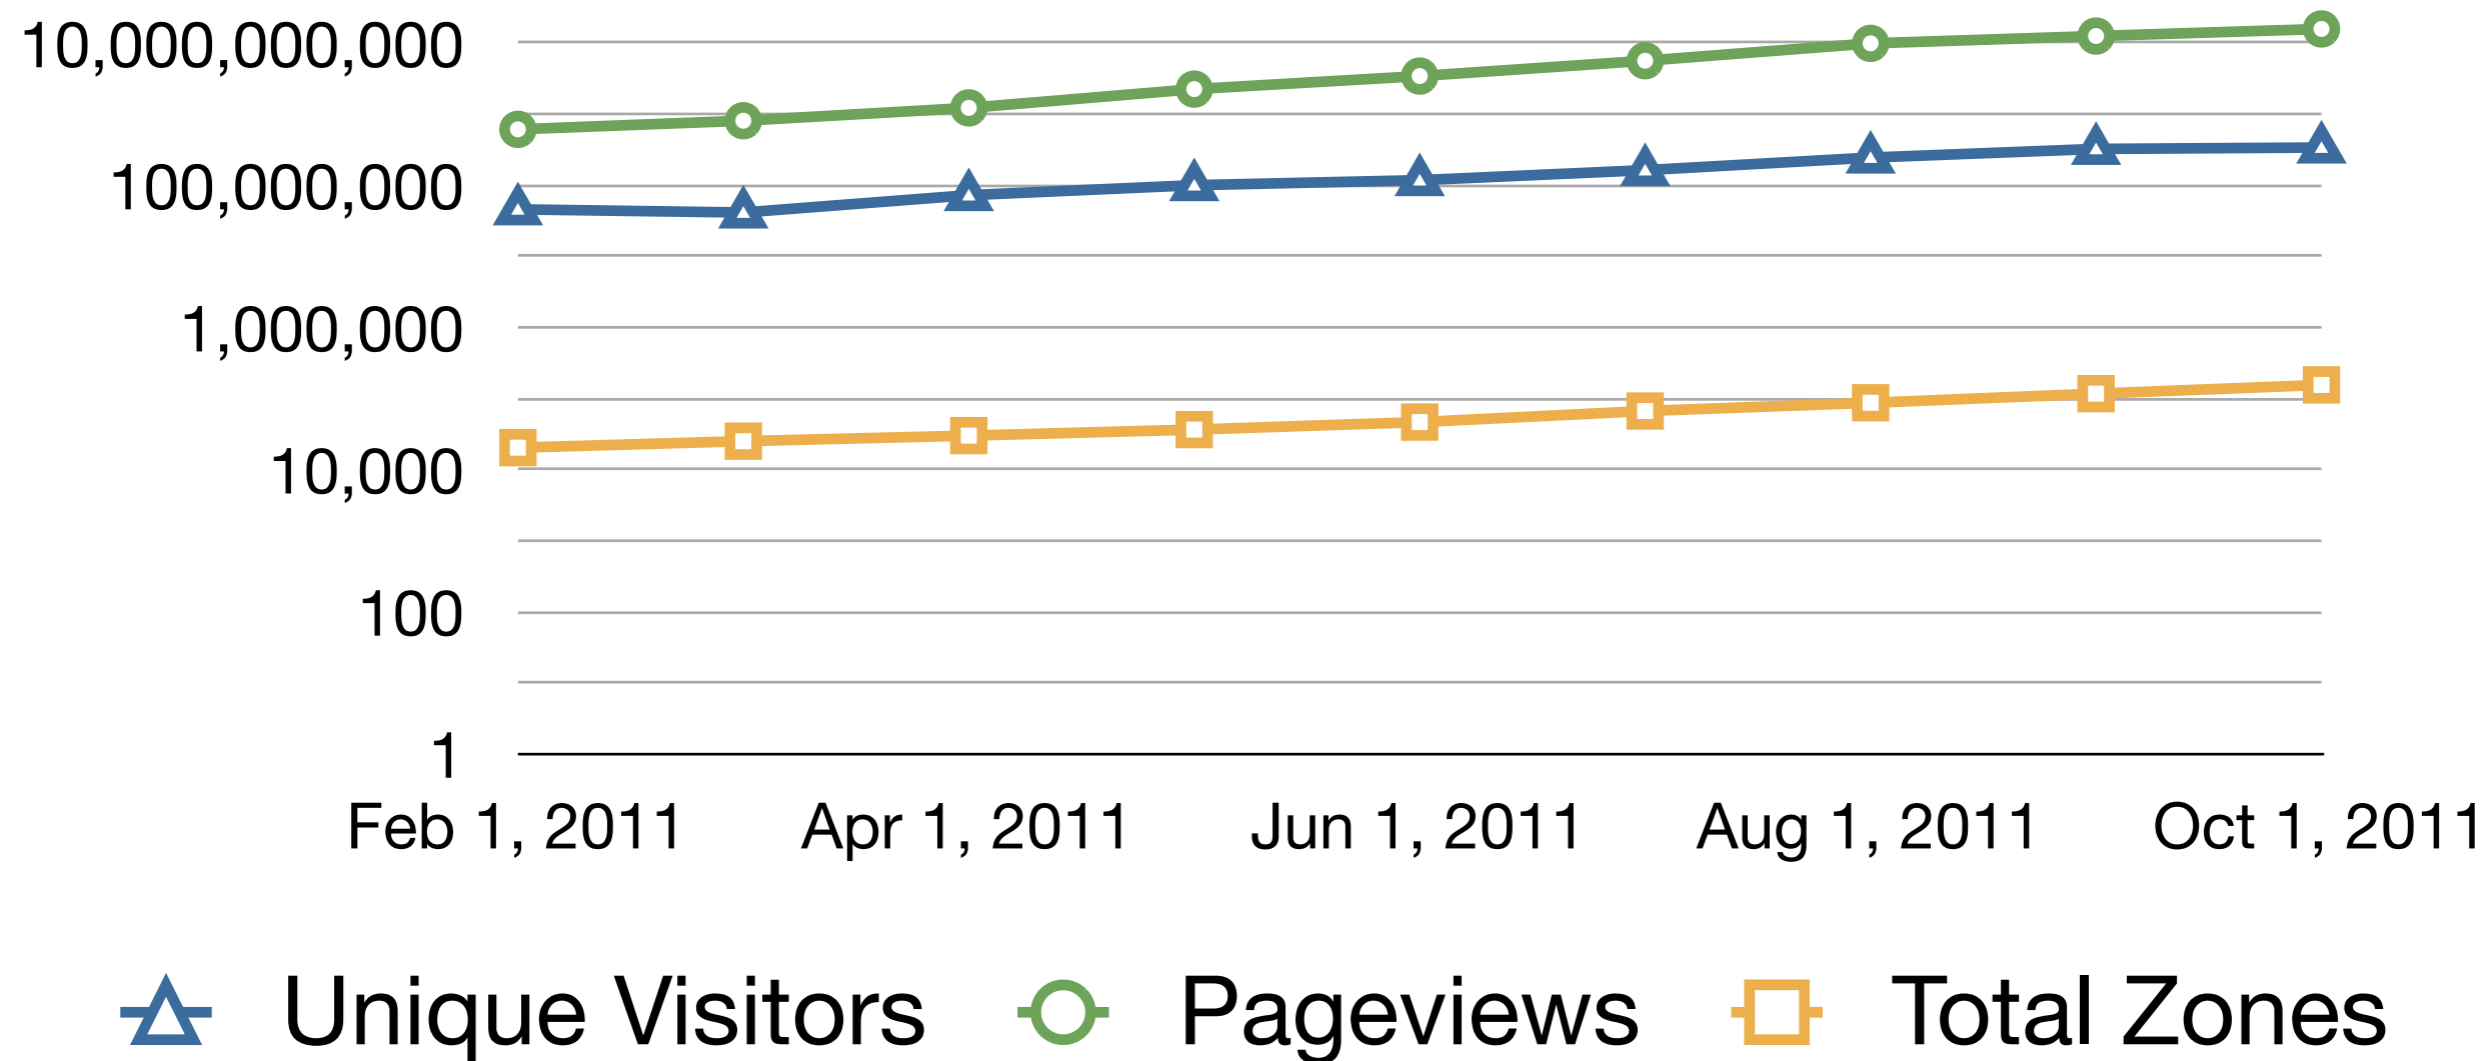
- What is CloudFlare
- How we use multi-threading
- Concurrency bugs in two example programs
- Concurrency analysis tools we use

CloudFlare: Next-Gen Content Delivery Network



CloudFlare: exponential traffic, linear hardware

(note log scale)



Supporting:

- 155,000 Sites
- 350,000 Requests / Second (peak)
- 17 Billion Page Views (last 30 days)
- 350 Million Unique Visitors (last 30 days)

Typical Production Server:

- Intel 24 Core CPU (with Hyper-threading)
- 48 MB RAM
- SSD storage

Ways We Write Multi-Threaded Code

- Callback Framework
(someone else does the hard work)
- Custom Threading
(up to the implementor)

Callback Frameworks

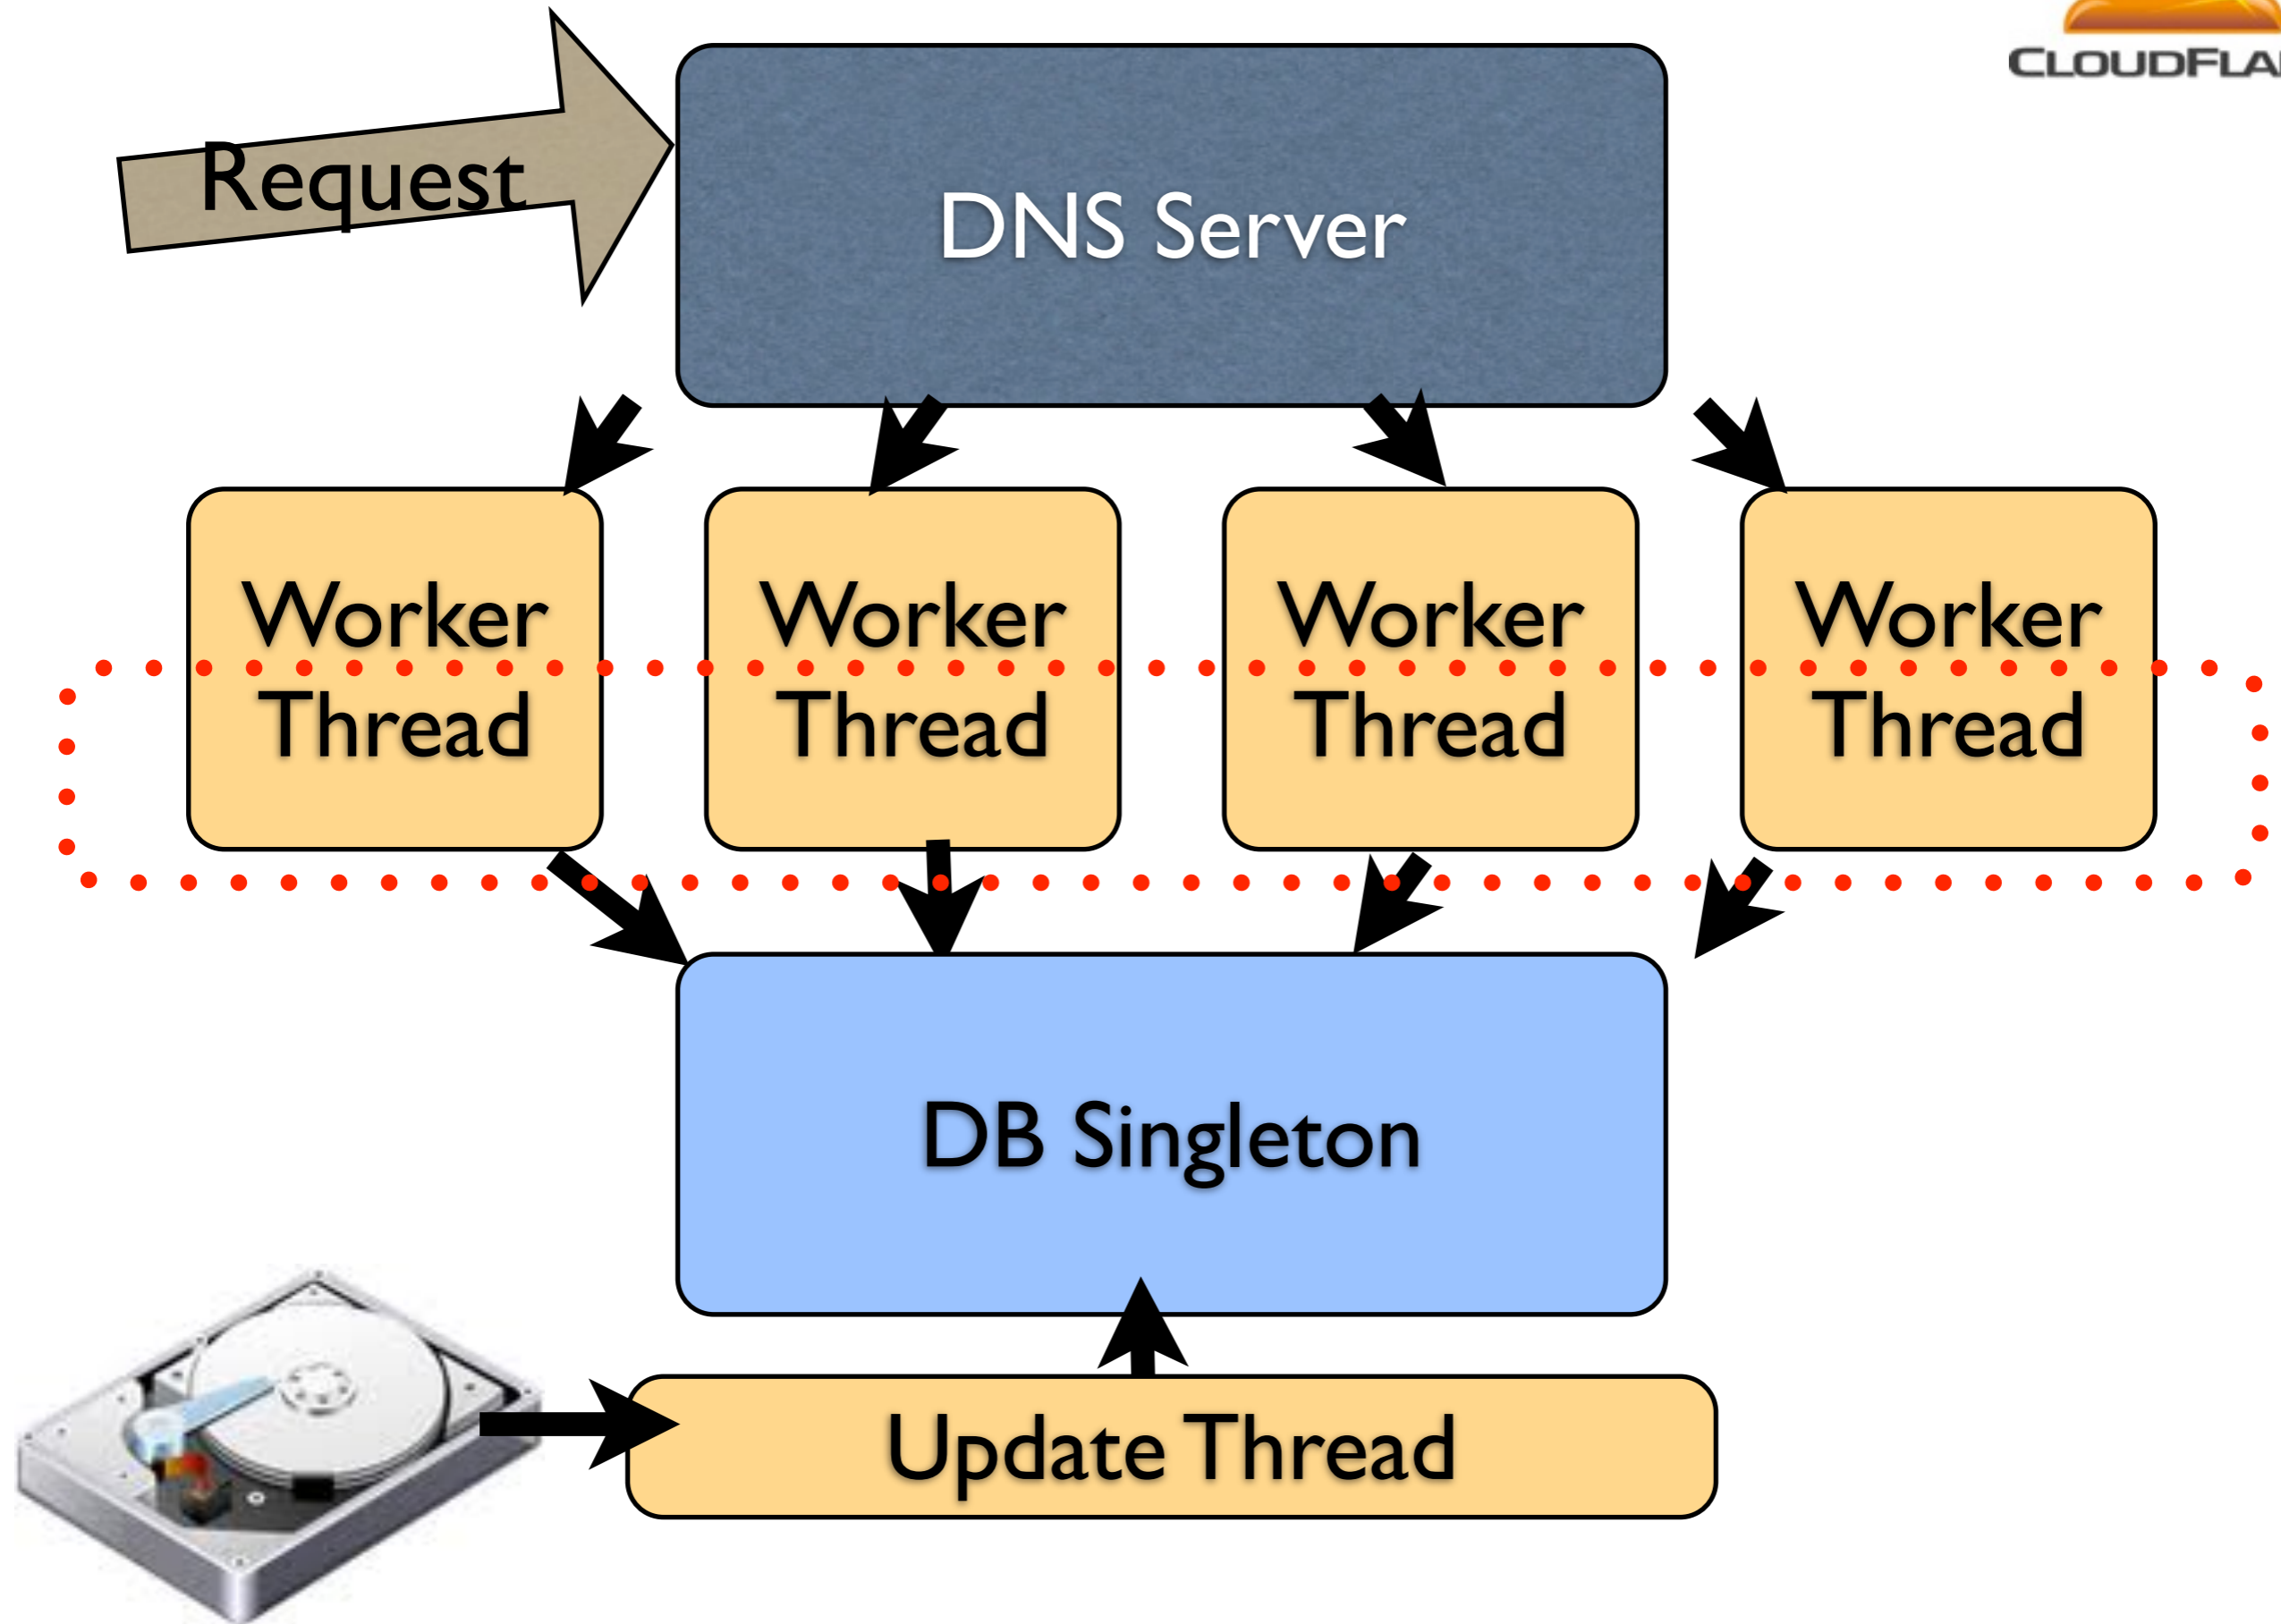
- Request handling consists of multiple sequential calls
- Multiple requests handled in parallel
- Implement callback functions
- Framework persists data across calls
- Examples:
 - FUSE
 - NGINX
 - DNS

Custom Parallelism

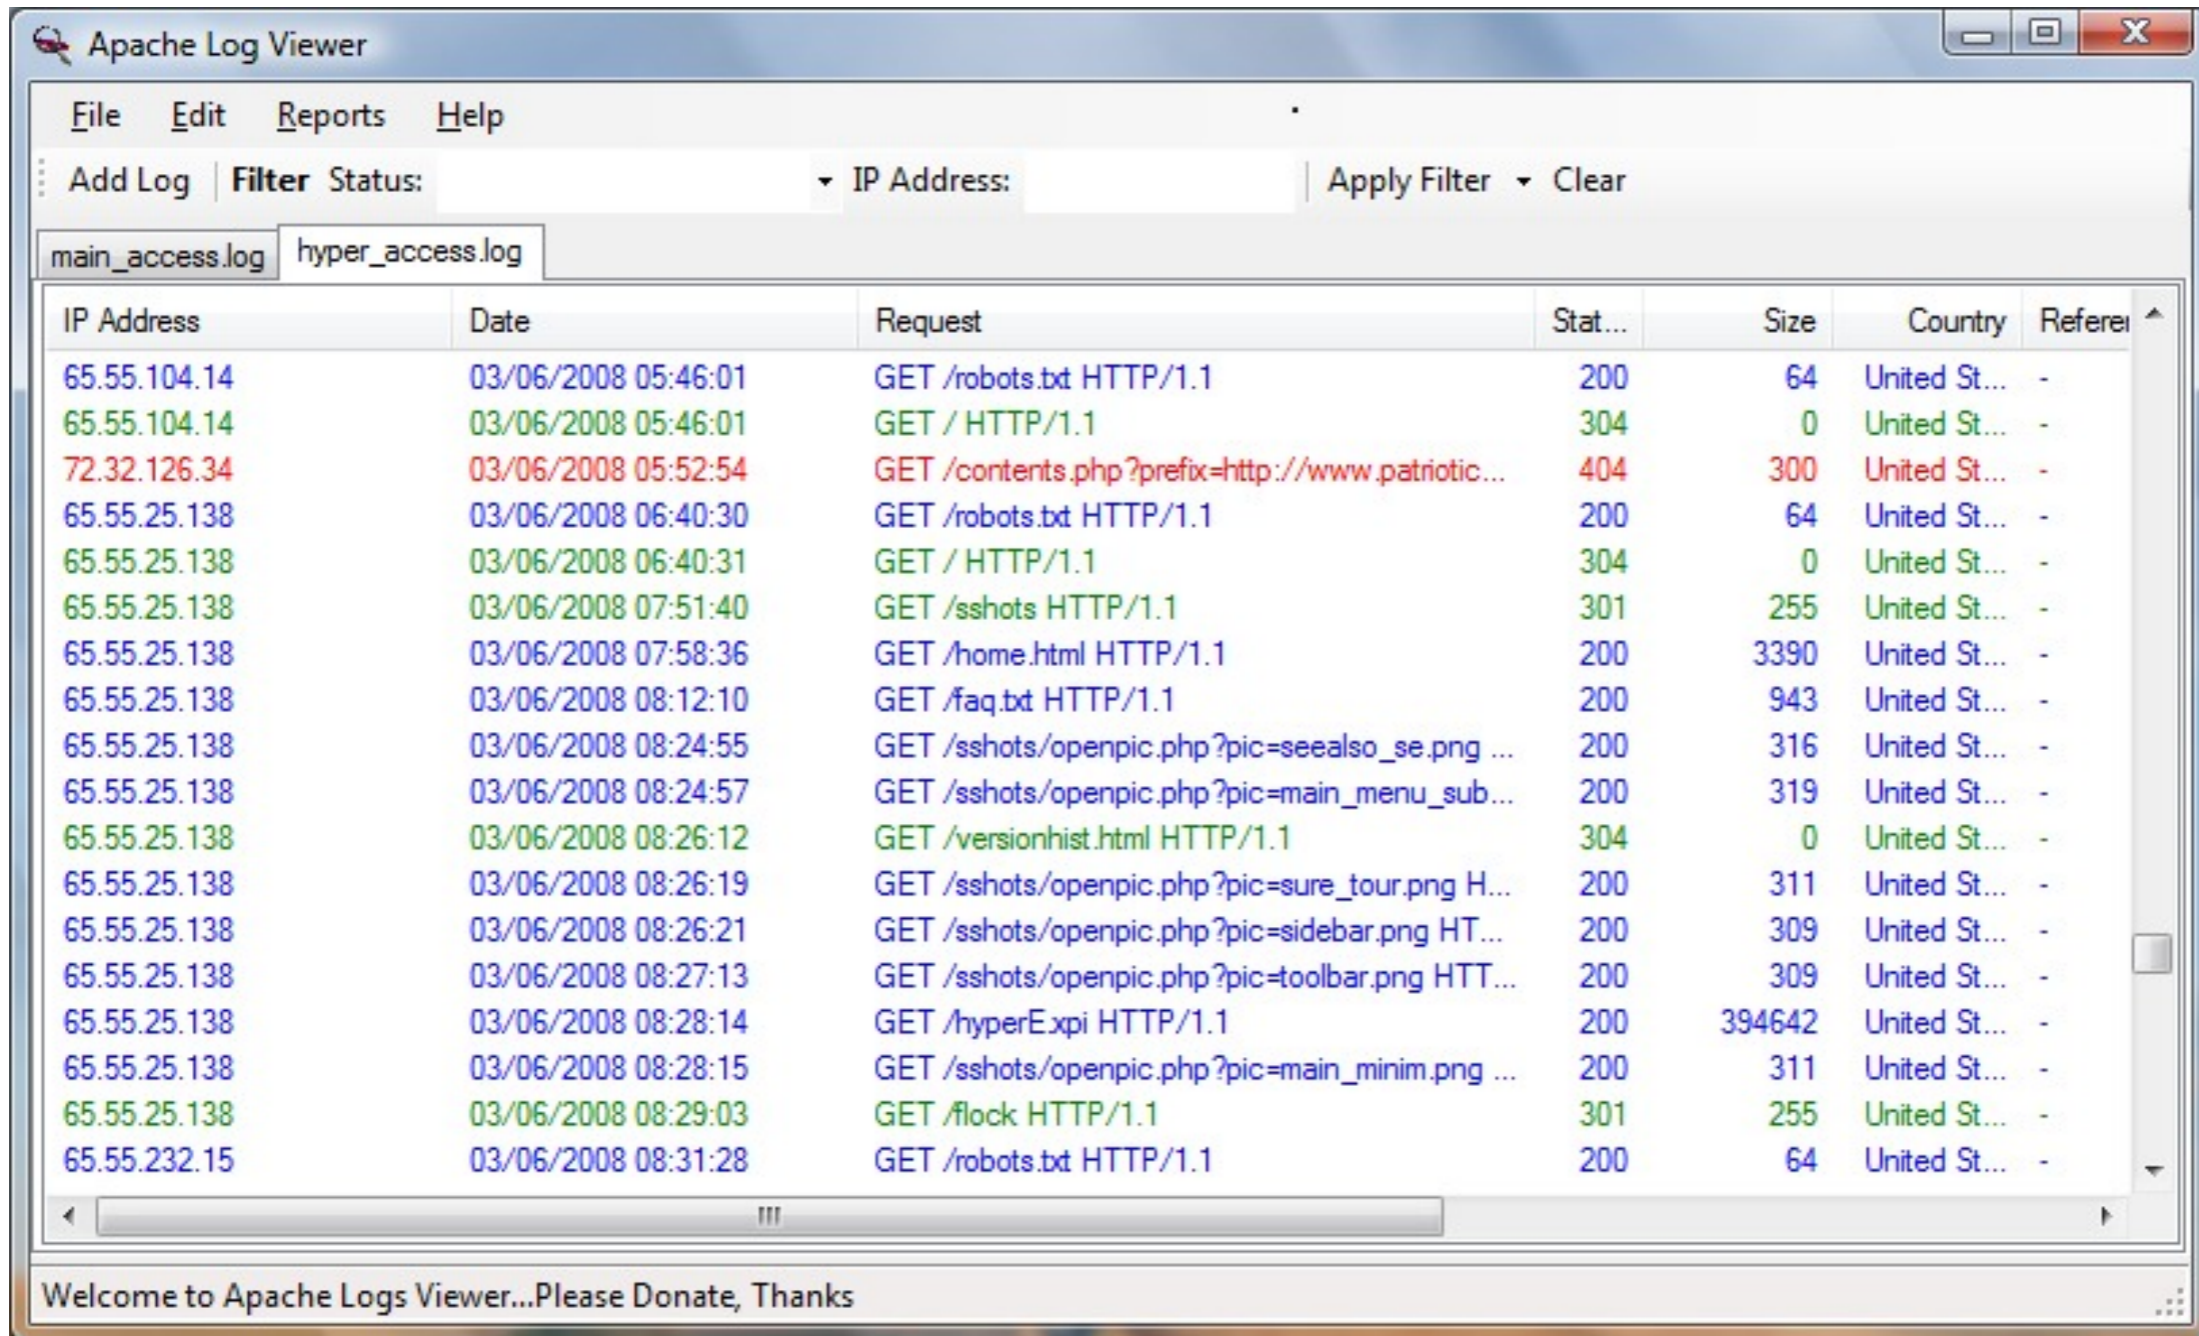
- Boost Library
- Synchronization with locks, db transactions
- Examples:
 - **Log Analysis**

DNS

```
ian@carrot:/home/q/import-log$ dig linuxdating.com
; <<>> DiG 9.7.1-P2 <<>> linuxdating.com
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 46055
;; flags: qr rd ra; QUERY: 1, ANSWER: 2, AUTHORITY: 0, ADDITIONAL: 0
;; QUESTION SECTION:
;linuxdating.com.                IN      A
;; ANSWER SECTION:
linuxdating.com.                 5       IN      A      173.245.60.112
linuxdating.com.                 5       IN      A      199.27.134.39
;; Query time: 212 msec
;; SERVER: 192.168.222.2#53(192.168.222.2)
;; WHEN: Thu Oct 20 20:59:22 2011
;; MSG SIZE rcvd: 65
```



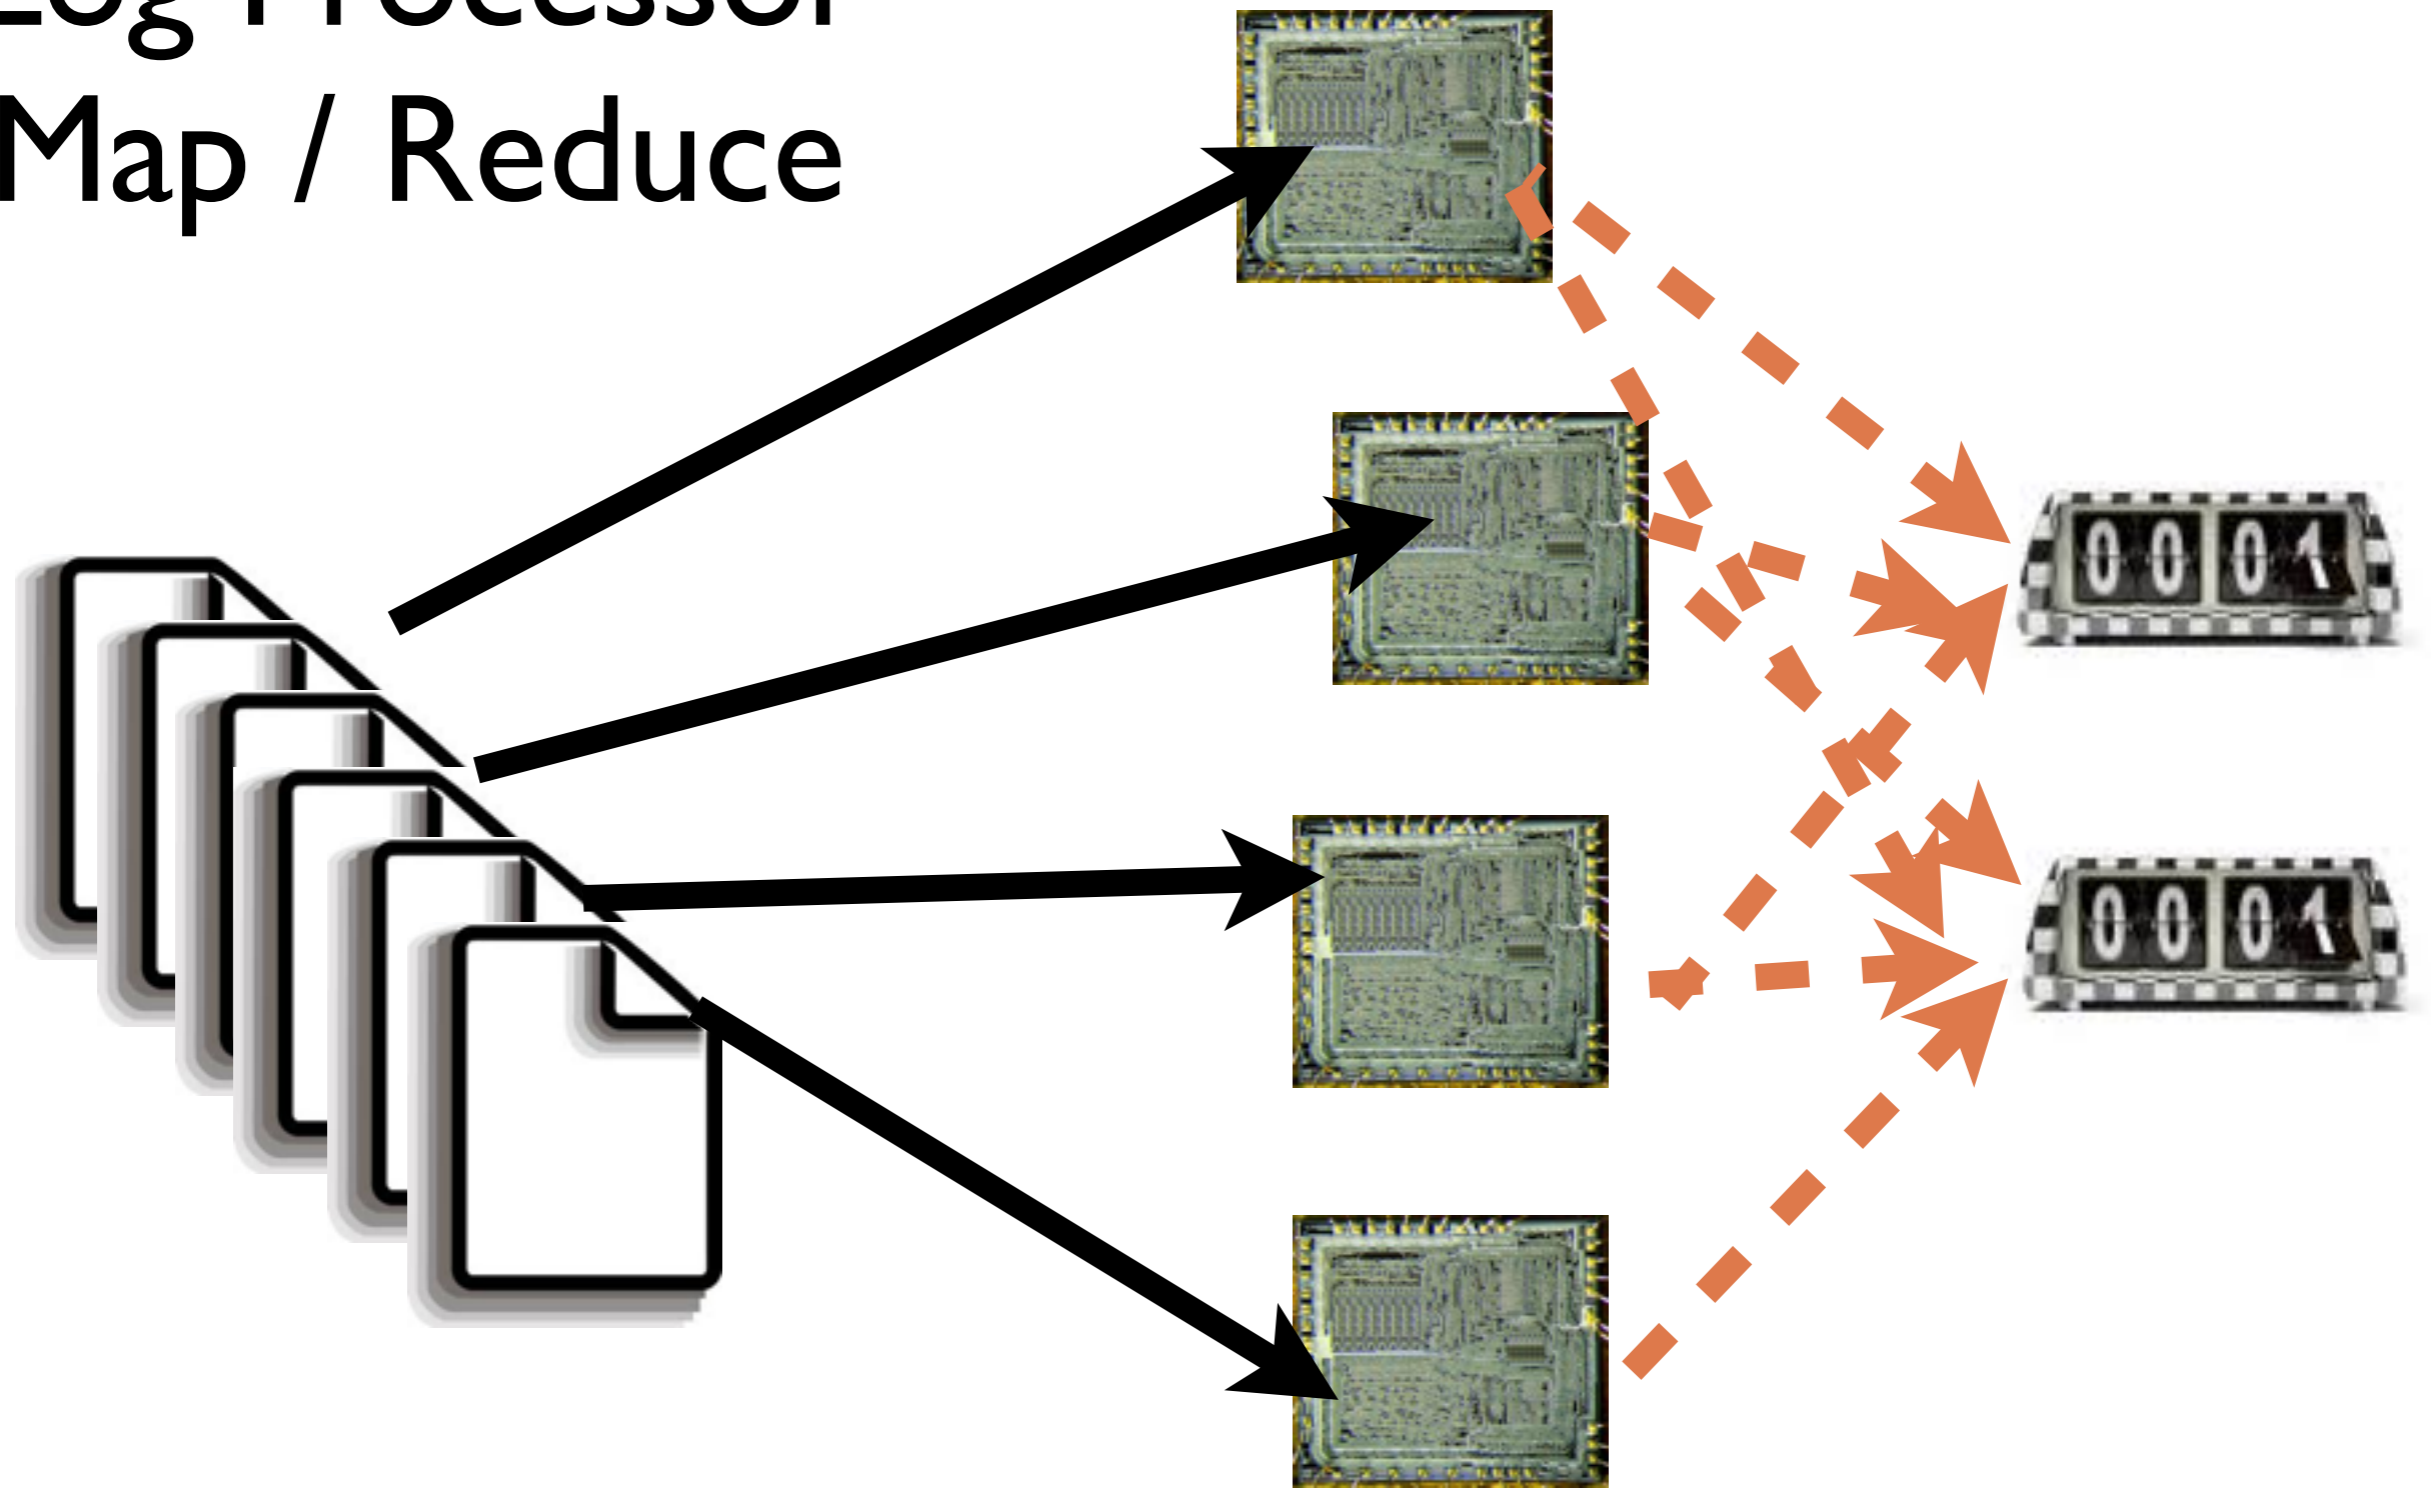
Log Processor: Logs Into Analysis



The screenshot shows the Apache Log Viewer application window. The title bar reads "Apache Log Viewer". The menu bar includes "File", "Edit", "Reports", and "Help". Below the menu bar, there are controls for "Add Log", "Filter Status:", "IP Address:", "Apply Filter", and "Clear". Two log files are selected: "main_access.log" and "hyper_access.log". The main area displays a table of log entries with columns for IP Address, Date, Request, Status, Size, Country, and Referer. The status codes are color-coded: 200 (blue), 304 (green), 404 (red), and 301 (green). The bottom status bar says "Welcome to Apache Logs Viewer...Please Donate, Thanks".

IP Address	Date	Request	Stat...	Size	Country	Referer
65.55.104.14	03/06/2008 05:46:01	GET /robots.txt HTTP/1.1	200	64	United St...	-
65.55.104.14	03/06/2008 05:46:01	GET / HTTP/1.1	304	0	United St...	-
72.32.126.34	03/06/2008 05:52:54	GET /contents.php?prefix=http://www.patriotic...	404	300	United St...	-
65.55.25.138	03/06/2008 06:40:30	GET /robots.txt HTTP/1.1	200	64	United St...	-
65.55.25.138	03/06/2008 06:40:31	GET / HTTP/1.1	304	0	United St...	-
65.55.25.138	03/06/2008 07:51:40	GET /sshots HTTP/1.1	301	255	United St...	-
65.55.25.138	03/06/2008 07:58:36	GET /home.html HTTP/1.1	200	3390	United St...	-
65.55.25.138	03/06/2008 08:12:10	GET /faq.txt HTTP/1.1	200	943	United St...	-
65.55.25.138	03/06/2008 08:24:55	GET /sshots/openpic.php?pic=seealso_se.png ...	200	316	United St...	-
65.55.25.138	03/06/2008 08:24:57	GET /sshots/openpic.php?pic=main_menu_sub...	200	319	United St...	-
65.55.25.138	03/06/2008 08:26:12	GET /versionhist.html HTTP/1.1	304	0	United St...	-
65.55.25.138	03/06/2008 08:26:19	GET /sshots/openpic.php?pic=sure_tour.png H...	200	311	United St...	-
65.55.25.138	03/06/2008 08:26:21	GET /sshots/openpic.php?pic=sidebar.png HT...	200	309	United St...	-
65.55.25.138	03/06/2008 08:27:13	GET /sshots/openpic.php?pic=toolbar.png HTT...	200	309	United St...	-
65.55.25.138	03/06/2008 08:28:14	GET /hyperE.xpi HTTP/1.1	200	394642	United St...	-
65.55.25.138	03/06/2008 08:28:15	GET /sshots/openpic.php?pic=main_minim.png ...	200	311	United St...	-
65.55.25.138	03/06/2008 08:29:03	GET /flock HTTP/1.1	301	255	United St...	-
65.55.232.15	03/06/2008 08:31:28	GET /robots.txt HTTP/1.1	200	64	United St...	-


Log Processor Map / Reduce



Errors Across Projects:

	<i>Logger</i>	<i>DNS</i>
Model	Native	Callback
LoC (Excluding Frameworks)	2,833	1,276
Source Commits	85	79
Transactions	5	6
Locks	1	0
Maximum Concurrent Threads	16	6
Race Conditions	8	3
Deadlocks	0	2

DNS Errors:



```
// accept a visitor to a record
bool accept(const char* kbuf,
            size_t ksiz,
            Visitor* visitor,
            bool writable = true) {
    _assert_(kbuf && ksiz <= kc::MEMMAXSIZ && visitor);
    kc::ScopedRWLock lock(&mlock_, false);
    size_t rsiz;
    const char* rbuf;
```

Log Errors: Write Synchronization

```
boost::mutex::scoped_lock lock(g_apache_feed_mutex);  
(*apache_file_out) << linebuf << endl;  
lock.unlock();
```

Tools Used:

- gdb
- Valgrind
- Helgrind
- aptitude install

What we need from a tool

- Have heard about it
- Easy to install
- Good documentation



GCC: `g++ -O2`

ICC: `icpc -Ofast`

ICC ~30% Faster for log code than gcc

atop: gamification of performance profiling

```
ATOP - 16log1 2011/08/14 21:00:37
```

PRC		sys	1.29s		user	16.38s		#proc	424
CPU		sys	38%		user	547%		irq	9%
cpu		sys	1%		user	99%		irq	0%
cpu		sys	0%		user	100%		irq	0%
cpu		sys	1%		user	99%		irq	0%
cpu		sys	1%		user	99%		irq	0%
cpu		sys	1%		user	99%		irq	0%
cpu		sys	11%		user	30%		irq	4%
cpu		sys	5%		user	0%		irq	0%
cpu		sys	1%		user	0%		irq	0%
cpu		sys	9%		user	15%		irq	1%
cpu		sys	4%		user	0%		irq	2%
cpu		sys	0%		user	0%		irq	0%
cpu		sys	0%		user	0%		irq	0%
cpu		sys	1%		user	0%		irq	0%
cpu		sys	1%		user	0%		irq	0%
cpu		sys	3%		user	0%		irq	0%
cpu		sys	1%		user	0%		irq	0%
cpu		sys	0%		user	1%		irq	1%
CPL		avg1	1.60		avg5	0.78		avg15	1.03
MEM		tot	47.3G		free	132.5M		cache	37.7G
SWP		tot	20.0G		free	19.9G			
PAG		scan	48171		stall	0			
DSK		cciss/c0d0			busy	58%		read	0
DSK		cciss/c0d1			busy	52%		read	0

**T-Shirts?
T-Shirts!**



CLOUDFLARE™







