Scaling to Millions of Simultaneous Connections

Rick Reed
WhatsApp

Erlang Factory SF
March 30, 2012

About ...

- Joined WhatsApp in 2011
- New to Erlang
- Background in performance of C-based systems on FreeBSD and Linux
- Prior work at Yahoo!, SGI

Overview

- The “good problem to have”
- Performance Goals
- Tools and Techniques
- Results
- General Findings
- Specific Scalability Fixes

The Problem

- A good problem, but a problem nonetheless
- Growth, Earthquakes, and Soccer!

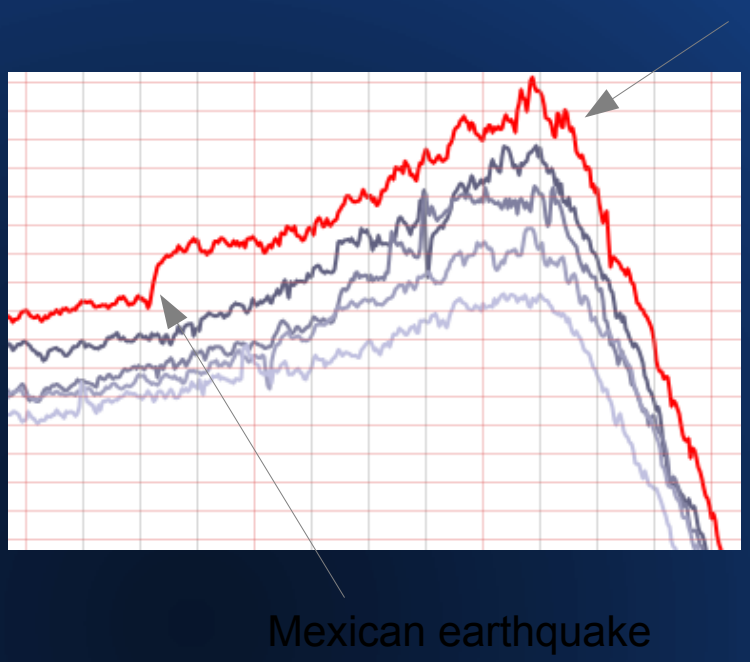

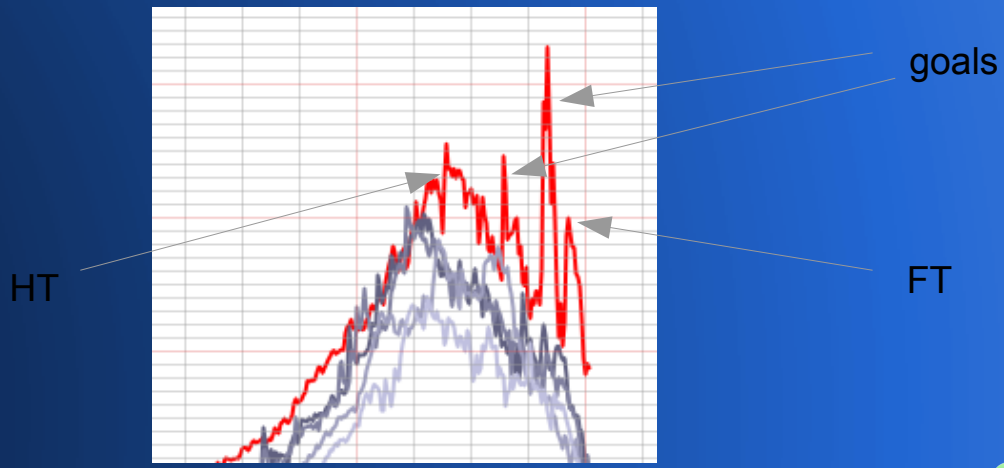

The Problem

- Initial server loading: ~200k connections
- Discouraging prognosis for growth
- Cluster brittle in the face of failures/overloads

Performance Goals

- 1 Million connections per server … !
- Resilience against disruptions under load
  - Software failures
  - Hardware failures (servers, network gear)
  - World events (sports, earthquakes, etc.)

Performance Goals

- Our standard configuration
  - Dual Westmere Hex-core (24 logical CPUs)
  - 100GB RAM, SSD
  - Dual NIC (user-facing, back-end/distribution)
  - FreeBSD 8.3
  - OTP R14B03

Tools and Techniques

- System activity monitoring (wsar)
  - OS-level

```
c114          | CPU-------------------------------| TCP     send rexmt recv-------- listn     pcb conn-| VM---------------------------------
time          | %util %user %nice  %sys %intr %idle| spkt/s  kb/s     % rpkt/s  kb/s ovflw   count  /sec| %util  %act %inac %cach %wire %free
03/19 16:00:21   40.4  25.9   0.0   9.7   4.7  59.6   70211  9280  5.58  60351  6659     0 1629907  1547   77.9  57.1  19.1   0.1  20.7   3.0
03/19 16:30:18   34.6  22.0   0.0   8.4   4.1  65.4   60492  7846  6.09  52173  5598     0 1567237  1455   77.8  57.0  19.4   0.1  20.7   2.8
03/19 17:00:15   29.9  18.9   0.0   7.5   3.6  70.1   52937  6678  6.50  45334  4701     0 1515203  1442   79.2  56.8  18.0   1.7  20.7   2.8
03/19 17:30:13   26.9  16.9   0.0   6.8   3.2  73.1   47918  5900  6.85  40618  4088     0 1478899  1462   79.9  56.7  17.4   2.5  20.7   2.7
03/19 18:00:13   24.6  15.3   0.0   6.3   3.0  75.4   43464  5256  7.29  36865  3628     0 1450185  1388   79.7  56.5  17.6   2.5  20.7   2.7
03/19 18:30:13   23.6  14.7   0.0   6.1   2.8  76.4   41605  4965  7.36  35177  3410     0 1435880  1375   79.6  56.4  17.7   2.5  20.7   2.7
```

  - BEAM

```
c114          | ERL   msgqlen-------------- dist---------------------------------------- inio outio------ sched---------------------------------- gc------------------------- mem----------------    nproc
time          | nodes     tot    max   nonz  nonzq   busy    kbq kbqmax   msgin  msgout  kb/s  kb/s  busy  load %util   csw/s wait/s sleep/s kred/s   /sec kwrd/s   long longms tot Mb sys Mb procMb     tot
03/20 07:30:25 [illegible]  924    789      8      0      0      0      0   45111   67136  7235 13040     0  0.58  59.3  181438  96261   92338  36955 100185 104625      1  ovflw  55075  25665  29410 1944015
03/20 08:00:29 [illegible] 1192   1054     10      0      0      0      0   46097   68425  7413 13329     0  0.58  61.4  184409  92777   89016  37808 101619 107098      2  ovflw  55585  25731  29855 1950798
03/20 08:30:31 [illegible] 1522   1388     10      0      0      0      0   46698   69192  7510 13482     0  0.58  61.9  185666  92035   88313  38133 102111 107971      2  ovflw  55807  25797  30010 1952132
03/20 09:00:33 [illegible] 1522   1388      6      0      0      0      0   47089   69739  7535 13510     0  0.61  62.2  185892  91312   87624  38114 102548 107787      2  ovflw  55976  25855  30122 1955793
```

Tools and Techniques

- Processor hardware perf counters (pmcstat)

```
PMC: [CPU_CLK_UNHALTED_CORE] Samples: 2563984 (100.0%) , 556 unr

%SAMP IMAGE      FUNCTION             CALLERS
 16.1 beam.smp.W process_main         process_main
  9.2 kernel     AcpiOsReadPort       AcpiHwReadPort
  3.8 kernel     cpu_search_highest   cpu_search_highest
  2.6 pmcstat    _init
  1.6 libthr.so. pthread_mutex_lock
  1.6 beam.smp.W schedule
  1.3 libthr.so. _pthread_mutex_unloc
  1.2 beam.smp.W cmp
  1.2 kernel     ipfw_chk             ipfw_check_hook
  1.1 kernel     sched_switch         mi_switch
  1.0 kernel     amd64_syscall
  1.0 beam.smp.W sweep_one_area
  1.0 kernel     _thread_lock_flags
  1.0 kernel     cpu_switch           mi_switch
  0.9 libthr.so. pthread_mutex_tryloc
  0.9 beam.smp.W minor_collection
  0.9 beam.smp.W handle_error
  0.8 beam.smp.W get_free_block
  0.7 kernel     tcp_output
  0.7 kernel     acpi_cpu_c1          acpi_cpu_idle
  0.6 libc.so.7  bsearch
  0.6 kernel     Xfast_syscall
  0.5 kernel     _rw_rlock
  0.5 beam.smp.W copy_struct
```

- dtrace, kernel lock-counting, gprof

Tools and Techniques

- fprof (w/ and w/o cpu_timestamp)

lists,do_flatten,2	234817	750.760	4.24%	4.24%	561.161
wap_1_1,decode_attrs,3	59473	425.213	2.40%	6.64%	994.159
wlib,get_option_value,3	84821	325.764	1.84%	8.48%	450.080
wap_1_1,encode_keyword_or_binary,1	73166	313.390	1.77%	10.25%	624.413
lists,map,2	72372	312.713	1.77%	12.01%	2369.478
envelope,parse_xml,1	38449	271.991	1.54%	13.55%	620.758
wlib,hex,2	51302	255.198	1.44%	14.99%	276.617
chatd_router,'-route_to_node/4-lc$^0/1-0-',2	75077	232.190	1.31%	16.30%	162.922
envelope,from_xml,2	23798	230.361	1.30%	17.60%	1416.922
xml,get_attr_s,3	55258	201.735	1.14%	18.74%	278.533
lists,keysearch,3	96324	201.082	1.14%	19.88%	201.082
wap_1_1,encode_attr,1	33905	199.083	1.12%	21.00%	785.417
binary_stream,parse,2	26403	198.421	1.12%	22.12%	2513.605
xml,normalize_attr,4	38453	194.789	1.10%	23.22%	314.607
wap_1_1,decode_thing,1	48607	185.607	1.05%	24.27%	243.223
envelope,fill_element,1	25059	184.037	1.04%	25.31%	456.130

- BEAM lock-counting (invaluable!!!)

Tools and Techniques

- Synthetic workload
  - Good for subsystems with simple interfaces
  - Limited value for user-facing systems

Tools and Techniques

- Tee'd workload
  - Where side-effects can be contained
  - Extremely useful for tuning

Tools and Techniques

- Diverted workload
  - Add additional production load to server
  - DNS via extra IP aliases
    - TTL issues
  - IPFW forwarding
    - Ran into a few kernel panics at high conn counts

Results

- Initial bottlenecks appeared around 425k
- First round of fixes got us to 1M conns
- Fruit was hanging pretty low

```
09/22 13:43:06   72.4  47.4   0.0  19.8   5.3  27.6   82183  8064  3.34  50116  4941     0 1086672  1414   75.1  62.6   0.0   0.0  12.5  24.9 996713
09/22 13:44:07   73.2  47.8   0.0  20.1   5.3  26.8   81902  8032  3.31  49977  4915     0 1088954  1403   75.4  62.9   0.0   0.0  12.5  24.6 998765
c114          | CPU-----------------------------| TCP     send rexmt recv-------- listn     pcb conn-| VM-------------------------------     open
time          | %util %user %nice  %sys %intr %idle| spkt/s  kb/s     % rpkt/s  kb/s ovflw   count  /sec| %util  %act %inac %cach %wire %free   files
09/22 13:45:07   72.8  47.5   0.0  20.0   5.3  27.2   82368  8066  3.33  50323  4938     0 1091250  1412   75.9  63.4   0.0   0.0  12.5  24.0 1000879
09/22 13:46:08   73.7  48.1   0.0  20.2   5.3  26.3   81200  7940  3.40  49433  4853     0 1094420  1415   76.0  63.5   0.0   0.0  12.5  24.0 1002723
09/22 13:47:07   72.7  47.3   0.0  20.0   5.4  27.3   83310  8115  3.47  50744  4948     0 1097600  1451   76.3  63.7   0.0   0.0  12.5  23.7 1004706
09/22 13:48:08   72.6  47.3   0.0  19.9   5.3  27.4   81302  7929  3.40  49421  4831     0 1099899  1418   76.3  63.8   0.0   0.0  12.5  23.7 1006536
09/22 13:49:08   73.4  47.9   0.0  20.1   5.4  26.6   82539  8060  3.39  50493  4927     0 1101326  1428   76.4  63.8   0.0   0.0  12.6  23.6 1007721
```

```
-------------  inio outio------ sched------------------------------------------------------------ gc------------------------ mem-----------------   nproc r
 msgin msgout  kb/s  kb/s  busy  load  runq %util %proc  %sys %port  csw/s wait/s sleep/s kred/s  /sec kwrd/s   long longms tot Mb sys Mb procMb     tot
 25358  57872  3496  1830     0     -     1  45.5  34.8   5.3   5.5 184614 110129   37843  26415 93910  72296      0      0  25365   3731  21635 1994028
 25263  57597  3472  1820     0     -     1  45.2  34.4   5.3   5.5 184409 111285   38255  26355 93691  72178      0      0  25398   3733  21665 1998180
 25447  57960  3493  1823     0     -     0  46.1  35.2   5.4   5.6 185335 108567   37285  26442 94237  72364      0      0  25475   3737  21738 2002343
 25026  57185  3434  1792     0     -    15  45.6  34.6   5.5   5.5 182235 109577   37308  25942 92693  70805      0      0  25526   3743  21784 2005993
 25696  58732  3492  1812     0     -     0  45.4  34.6   5.3   5.5 187135 111210   38124  26637 95587  72641      0      0  25524   3742  21782 2010065
 25000  57264  3405  1779     0     -     2  45.5  34.7   5.3   5.5 182249 108503   37347  25962 93047  70835      0      0  25485   3744  21741 2013718
 25473  58031  3482  1814     0     -    20  45.3  34.4   5.4   5.5 186396 112905   38706  26474 94828  72314      0      0  25528   3745  21783 2016089
```

Results

- Continued attacking similar bottlenecks
- Achieved 2M conns about a month later

c114	TCP		send	rexmt	recv--------		listn	pcb	conn-	CPU----------						VM----------						open
time	spkt/s	kb/s	%	rpkt/s	kb/s	ovflw	count	/sec	%util	%user	%nice	%sys	%intr	%idle	%util	%act	%inac	%cach	%wire	%free	files
11/03 03:00:21	110690	10686	4.57	67959	6505	0	2150655	1982	71.3	48.6	0.0	15.7	7.0	28.7	79.2	57.6	18.8	3.1	18.5	2.0	1982322
11/03 03:10:13	117273	11316	4.64	71752	6891	0	2174200	2109	73.5	50.1	0.0	16.1	7.3	26.5	79.9	58.7	18.9	2.7	18.5	1.2	1999202
11/03 03:20:14	115909	11227	4.65	71277	6871	0	2183105	2057	73.8	50.4	0.0	16.1	7.3	26.2	79.9	59.1	19.0	2.2	18.6	1.0	2007357
11/03 03:30:14	116855	11333	4.60	71920	6944	1	2191993	2063	74.6	51.0	0.0	16.2	7.4	25.4	80.8	59.5	18.1	2.8	18.6	1.0	2016065
11/03 03:40:14	118619	11513	4.64	73021	7065	0	2204591	2094	75.6	51.6	0.0	16.4	7.6	24.4	82.3	60.1	16.7	3.6	18.6	1.0	2025539
11/03 03:50:14	120747	11721	4.55	74150	7181	0	2213468	2141	76.7	52.5	0.0	16.5	7.7	23.3	82.2	60.5	16.8	3.1	18.6	1.0	2035142

		inio	outio------	sched----------											gc----------		mem----------				nproc
msgin	msgout	kb/s	kb/s	busy	load	runq	%util	%proc	%sys	%port	csw/s	wait/s	sleep/s	kred/s	/sec	kwrd/s	long	longms	tot Mb	sys Mb	procMb	tot
35866	84022	4876	2617	0	0.79	22	76.5	58.7	9.6	8.3	266873	63415	61206	55107	100241	126079	0	0	54213	5359	48854	4033240
36443	85278	4964	2671	0	0.75	30	77.9	59.9	9.6	8.4	270788	60657	58584	56006	101723	128218	0	0	54760	5377	49384	4052189
37116	86897	5047	2708	0	0.74	22	79.3	61.1	9.7	8.5	274414	57424	55521	56862	103163	130241	0	0	55187	5391	49797	4071473
37448	87484	5120	2744	0	0.73	31	79.8	61.6	9.7	8.6	276563	56566	54701	57503	104058	131868	0	0	55545	5402	50143	4086118
38617	90429	5266	2820	0	1.42	81	82.7	64.4	9.5	8.8	282287	48701	47185	58975	106369	135123	0	0	56734	5439	51296	4129050
38588	89964	5286	2834	0	0.79	31	82.9	64.4	9.6	8.8	284019	49194	47672	59228	106991	135949	0	0	56913	5450	51464	4146068

- Put further optimizations on back burner

Results

- Began optimizing app code after New Years
- Unintentional record attempt in Feb
- Peaked at 2.8M conns before we intervened

c114	CPU						soft	hard	sysc	trap	csw	TCP	send	rexmt	recv		listn	pcb	conn-	IGB 0		pkts	IGB 1		pkts	VM						open
time	%util	%user	%nice	%sys	%intr	%idle	int/s	int/s	/sec	/sec	/sec	spkt/s	kb/s	%	rpkt/s	kb/s	ovflw	count	/sec	rxd	txd	rxd	txd	%util	%act	%inac	%cach	%wire	%free	files
02/10 08:11:55	86.7	57.0	0.0	18.4	11.3	13.3	1108	264980	2368976	238120	1705338	164399	19800	5.42	135200	14671	0	3117995	3730	114883	131997	131917	126831	70.1	48.9	17.0	0.0	21.2	12.8	2840133
02/10 08:12:57	86.6	57.0	0.0	18.4	11.2	13.4	1003	239825	2160137	215401	1557662	148493	17931	5.39	122262	13329	0	3118270	3356	104334	119783	119266	114352	70.2	48.9	17.1	0.0	21.2	12.8	2841741
02/10 08:13:58	84.8	55.7	0.0	18.0	11.1	15.2	1047	246350	2267368	226498	1585579	153385	18430	5.45	126761	13741	9	3117883	3469	107812	123619	122287	119682	70.2	49.0	17.1	0.0	21.2	12.7	2841534
02/10 08:14:59	86.7	57.0	0.0	18.4	11.2	13.3	1011	240120	2127992	215794	1526388	150646	18220	5.37	123375	13348	0	3116209	3356	104857	120618	119800	116072	70.2	49.0	17.1	0.0	21.2	12.7	2841173
02/10 08:17:11	86.8	57.2	0.0	18.3	11.3	13.2	952	228098	2042882	204583	1474679	141511	17071	5.47	116695	12652	0	3120540	3200	99536	114129	113795	109519	70.3	49.0	17.1	0.0	21.2	12.6	2843567
02/10 08:18:52	85.7	56.3	0.0	18.0	11.4	14.3	1233	290008	2649440	258062	1903005	185917	22205	5.47	153040	16687	124	3131984	4604	132186	151907	145156	141908	70.7	49.4	17.2	0.0	21.2	12.1	2848794

msgqlen		dist							inio	outio		sched									gc			mem			nproc
tot	max	nonz	nonzq	busy	kbq	kbqmax	msgin	msgout	kb/s	kb/s	busy	load	runq	%util	%proc	%sys	%port	csw/s	wait/s	sleep/s	kred/s	/sec	kwrd/s	long	longms	tot Mb	sys Mb	procMb	tot
17354	17330	20	0	0	0	0	71117	107474	10898	19784	0	7.84	3299	94.7	76.4	6.9	11.4	265025	15066	14724	57314	185798	154423	0	0	45344	7627	37718	2843819
17129	16954	20	0	0	0	0	64576	97327	9930	17991	0	5.87	2283	94.8	76.6	6.8	11.3	240476	12948	12674	52063	168468	140181	0	0	45439	7638	37801	2845517
75604	75482	16	0	0	0	0	66997	100943	10294	18483	0	28.02	5079	96.8	78.7	6.7	11.4	241920	7360	7205	53435	173541	143651	1	284	45679	7654	38026	2846089
75604	75482	16	0	0	0	0	64813	98360	9900	18176	0	11.02	981	96.4	78.6	6.3	11.5	238699	8530	8344	52375	168353	141143	2	342	45450	7644	37806	2844836
22392	22372	14	0	0	0	0	61307	92417	9419	17069	0	8.13	2686	96.0	78.2	6.5	11.3	226986	9523	9322	49457	160327	133160	0	0	45630	7666	37965	2847343
571109	570833	16	0	0	0	0	82111	123968	12444	22192	0	27.38	8585	96.8	79.0	6.3	11.4	288432	10024	9801	63429	207925	169649	1	537	45889	7877	38013	2852986

- 571k pkts/sec, >200k dist msgs/sec

Results

- Still trying to obtain elusive 3M conns
- St. Patrick's Day wasn't as lucky as hoped

c114	CPU						soft	hard	sysc	trap	csw	TCP	send	rexmt	recv		listn	pcb	conn-	IGB 0	pkts	IGB 1	pkts	VM						open
time	%util	%user	%nice	%sys	%intr	%idle	int/s	int/s	/sec	/sec	/sec	spkt/s	kb/s	%	rpkt/s	kb/s	ovflw	count	/sec	rxd	txd	rxd	txd	%util	%act	%inac	%cach	%wire	%free	files
03/17 07:03:33	78.2	50.3	0.0	18.0	9.8	21.8	1119	271204	2009381	253310	1349209	143449	17938	6.39	122083	13284	0	2982894	3768	105617	120758	111423	106799	80.6	59.8	15.4	0.0	20.8	4.0	2695134
03/17 07:05:40	77.7	50.2	0.0	17.9	9.7	22.3	967	234706	1733664	218113	1174989	122308	15336	6.39	104352	11403	0	2989099	3178	90105	103053	95520	91461	80.7	59.9	15.4	0.0	20.8	3.9	2700078

msgqlen	dist								inio	outio		sched										gc				mem			nproc
tot	max	nonz	nonzq	busy	kbq	kbqmax	msgin	msgout	kb/s	kb/s	busy	load	runq	%util	%proc	%sys	%port	csw/s	wait/s	sleep/s	kred/s	/sec	kwrd/s	long	longms	tot Mb	sys Mb	procMb	tot
110	49	7	0	0	0	0	64709	95333	9692	17416	0	0.79	60	79.8	58.9	9.8	11.1	238552	58460	56629	47403	135010	133058	1	814	54583	19865	34719	2695014
110	49	6	0	0	0	0	61744	90988	9343	16742	0	0.82	6	80.7	59.7	9.8	11.2	227256	53812	52156	45582	128369	128019	0	0	54896	19886	35011	2695736
133	49	6	0	0	0	0	52958	78100	7950	14313	0	0.78	69	81.6	60.6	9.8	11.2	194380	45056	43680	38990	109996	109469	0	0	55015	19897	35119	2697622
133	49	6	0	0	0	0	66777	98413	9969	17930	0	0.85	3	82.2	60.8	9.9	11.5	243570	53391	51820	48803	138782	136769	0	0	55355	19915	35440	2705748

20120317070502										
PID	Group	InitialCall	CurrentCall	Reductions	Messages	Heap	DeqRate	QDelay	QFillRate	ETDrain
<0.287.0>	chatd_local(F)	proc_lib:init_p/5	gen_factory:factory_method/3	235395876870	606770	177508148	40273	15	-14644/-283	41/2144
20120317070513										
PID	Group	InitialCall	CurrentCall	Reductions	Messages	Heap	DeqRate	QDelay	QFillRate	ETDrain
<0.287.0>	chatd_local(F)	proc_lib:init_p/5	gen_server:handle_common_reply/6	235411174533	367643	181346591	46443	7	-20744/-7557	17/48
20120317070523										
PID	Group	InitialCall	CurrentCall	Reductions	Messages	Heap	DeqRate	QDelay	QFillRate	ETDrain
<0.287.0>	chatd_local(F)	proc_lib:init_p/5	gen_server:handle_common_reply/6	235428828436	159169	90777778	48105	3	-22190	7
<0.55.0>	wa_log(F)	proc_lib:init_p/5	gen_server:loop/6	14404935976	49	3571	1574	-	-	49/9

General Findings

- Erlang has awesome SMP scalability
  - >85% cpu utilization across 24 logical cpus
  - FreeBSD shines as well

General Findings

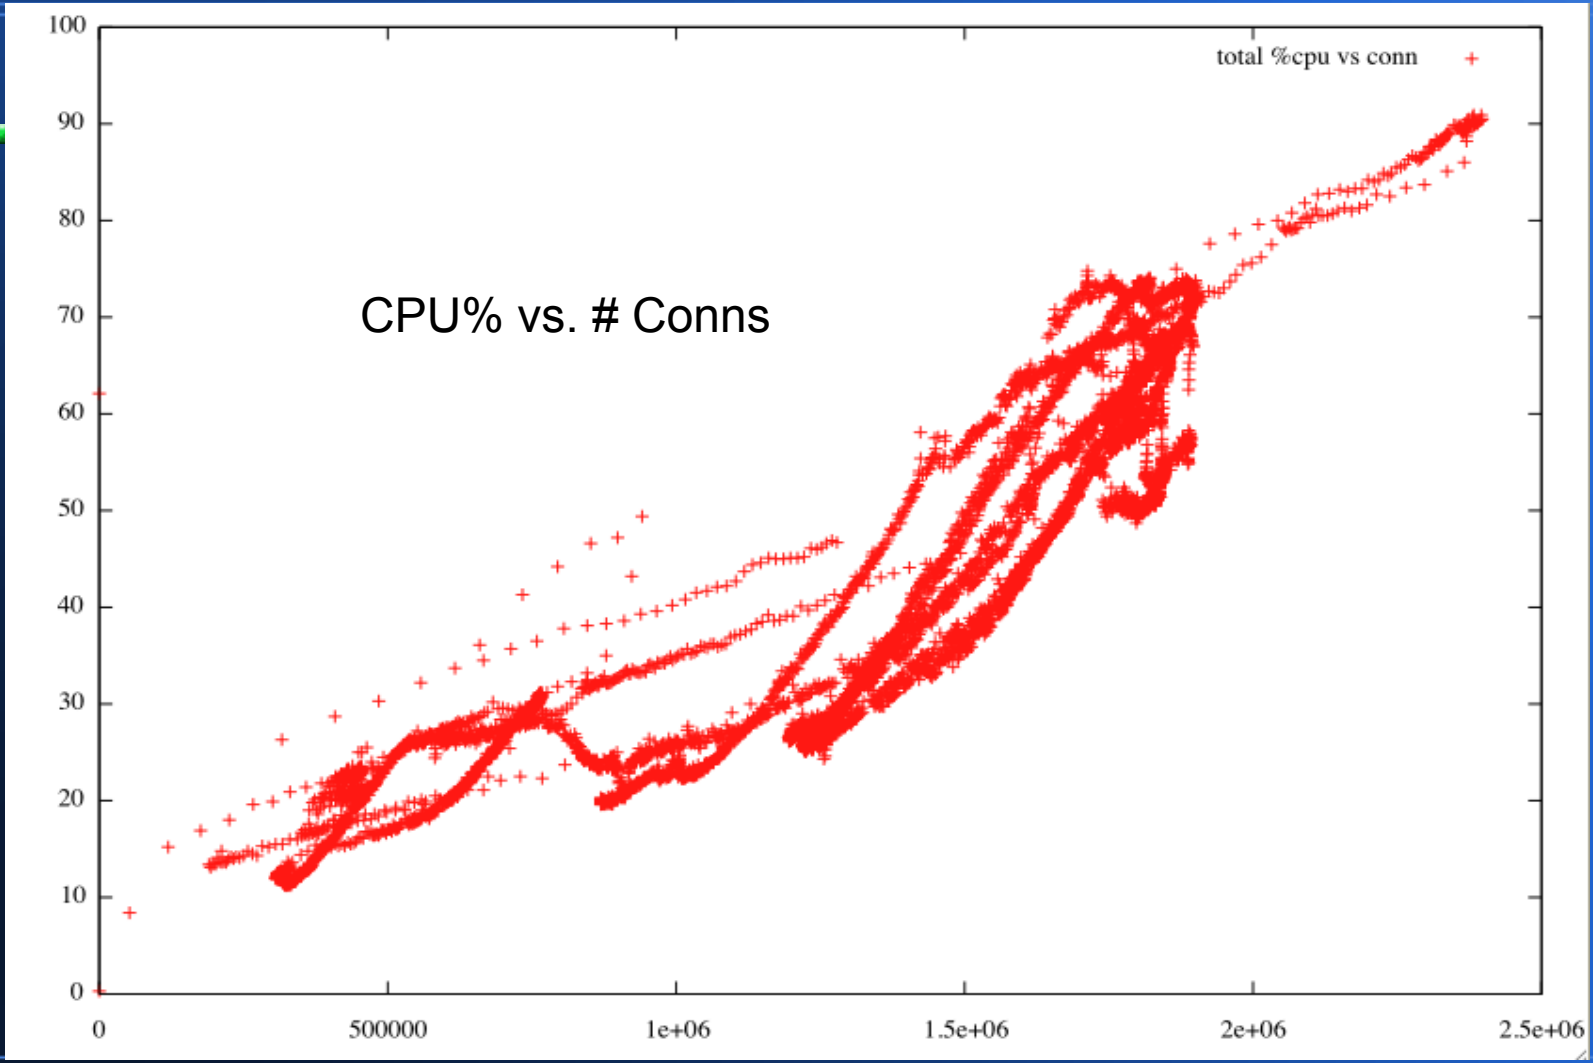

General Findings

- Contention, contention, contention
  - From 200k to 2M were all contention fixes
  - Some issues are internal to BEAM
    - Some addressable with app changes
    - Most required BEAM patches
  - Some required app changes
    - Especially: partitioning workload correctly
    - Some common Erlang idioms come at a price

Specific Scalability Fixes

- FreeBSD
  - Backported TSC-based kernel timecounter
    - gettimeofday(2) calls much less expensive
  - Backported igb network driver
    - Had issues with MSI-X queue stalls
  - sysctl tuning
    - Obvious limits (e.g., kern.ipc.maxsockets)
    - net.inet.tcp.tcphashsize=524288

Specific Scalability Fixes

- BEAM metrics
  - Scheduler (%util, csw, waits, sleeps, …)
  - statistics(message_queues)
    - Msgs queued, #non-empty queues, longest queue
  - process_info(message_queue_stats)
    - Enq/deq/send count & rates (1s, 10s, 100s)
  - statistics(message_counts)
    - Aggregation of message_queue_stats
  - Enable fprof cpu_timestamp for FreeBSD

Specific Scalability Fixes

- BEAM metrics (cont.)
  - Make lock-counting work for larger async thread counts (e.g., +A 1024)
  - Add suspend, location, and port_locks options to erts_debug:lock_counters
  - Enable/disable process/port lock counting at runtime
  - Fix missing accounting for outbound dist bytes

Specific Scalability Fixes

- BEAM tuning
  - +swt low
    - Avoid scheduler perma-sleep
  - +Mummc/mmmbc/mmsbc 99999
    - Prefer mseg over malloc
  - +Mut 24
    - Want allocator instance per scheduler

Specific Scalability Fixes

- BEAM tuning
  - +Mulmbcs 32767 +Mumbcgs 1 +Musmbcs 2047
    - Want large 2M-aligned mseg allocations to maximize superpage promotions
  - Run with real-time scheduling priority
  - +ssct 1 (via patch; scheduler spin count)

Specific Scalability Fixes

- BEAM contention
  - timeofday lock (esp., timeofday delivery)
  - Reduced slot traversals on timer wheel
  - Widened bif timer hash table
    - Ended up moving bif timers to receive timeouts
  - Improved check_io allocation scalability
  - Added prim_file:write_file/3 & /4 (port reuse)
  - Disable mseg max check

Specific Scalability Fixes

- BEAM contention (cont.)
  - Reduce setopts calls in prim_inet:accept and in inet:tcp_controlling_process

Specific Scalability Fixes

- OTP throughput
  - Add gc throttling when message queue is long
  - Increase default dist receive buffer from 4k to 256k (and make configurable)
  - Patch mnesia_tm to dispatch async_dirty txns to separate per-table procs for concurrency
  - Add pg2 denormalized group member lists to improve lookup throughput
  - Increase max configurable mseg cache size

Specific Scalability Fixes

- Erlang usage
  - Prefer os:timestamp to erlang:now
  - Implement cross-node gen_server calls without using monitors (reduces dist traffic and proc link lock contention)
  - Partition ets and mnesia tables and localize access to smaller number of processes
  - Small mnesia clusters

Specific Scalability Fixes

- Operability fixes
  - Added [prepend] option to erlang:send
  - Added process_flag(flush_message_queue)

Questions? Comments?

- rr@whatsapp.com