ZABBIX 2012 Conference

Zabbix Trending:

Investigating Production Problems

By Erik Skytthe, DBC, Denmark

DBC: http://www.dbc.dk/english/about_dbc

Erik: http://www.linkedin.com/pub/erik-skytthe/20/a77/444

Zabbix IRC/Forum: eskytthe

DBC’s main task is to develop and maintain the bibliographic and IT-infrastructure in **Danish libraries** by:

- Producing the **Danish National Bibliography** and user oriented cataloguing
- Developing and maintaining **Dan*Bib** as Union Catalog for public, educational and academic libraries as part of an automated ILL
- Developing and maintaining **library.dk** as the common access for citizens to Danish publications and Danish library holdings.

- DBC’s IT systems are based on **open source** and service oriented architecture
- … E.g https://github.com/ding2/ding2

A specialized world

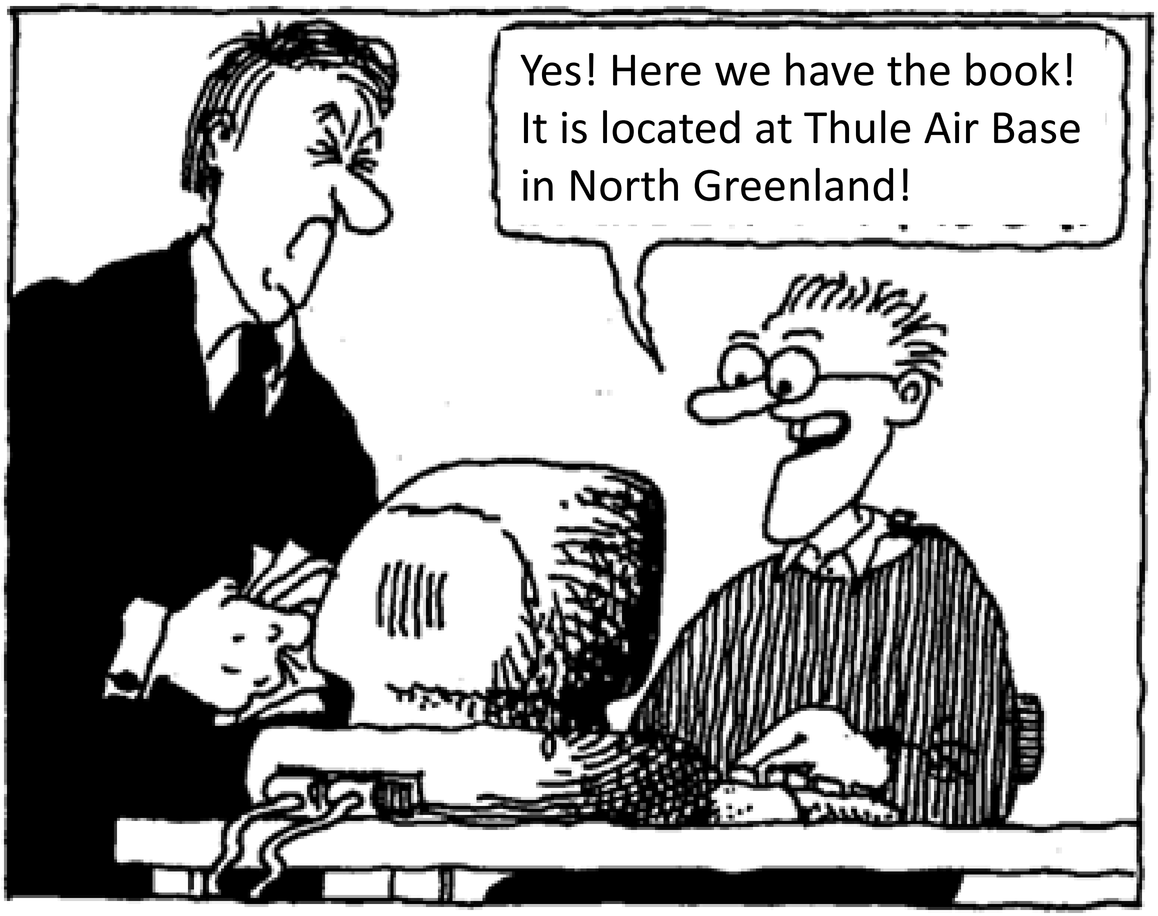

And something completely different:

”If you have 50 search fields, a librarian will use them all!”

The Ghost case

- Victims: Dead or unconscious services
- Detail: Spontaneous “ghost” outages in multiple services
- … and in non related services and systems !
- Suspect: Network?
- Evidence: Trading graphs, monitoring alarms

THE DAME SAID SHE HAD A CASE. SHE SOUNDED LIKE A CASE HERSELF, BUT I CAN'T CHOOSE MY CLIENTS.
SHE WAS THE PUSHY TYPE. THE KIND WHO'D BREAK YOUR HEART, OR MAYBE YOUR ARMS. I HURRIED OVER.

© Universal Press Syndicate

The Ghost case

- **Central Firewall was found guilty!**
- **Root cause:** Setup of max connections was to low!
- Background: New Ajax based web services have over time increased the demand of connections allowed (and a bug)

- Tools: Central syslog, Hobbit and Zabbix

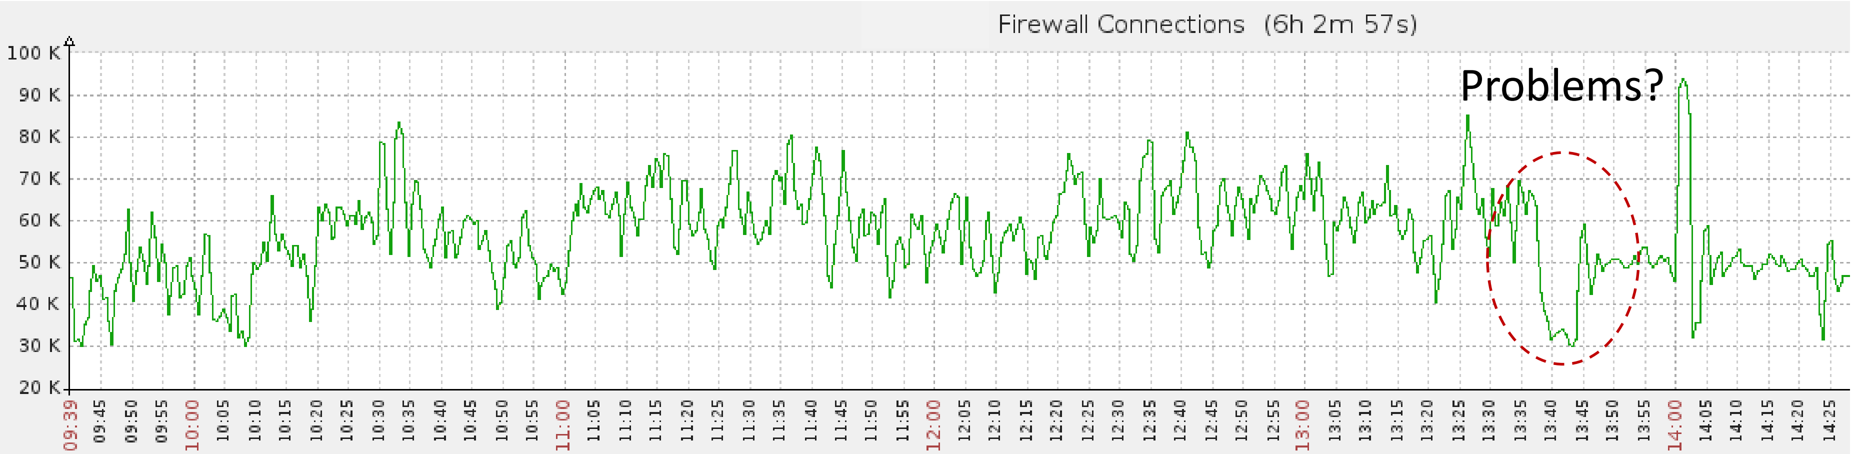

Problems?

The Ghost case

- Made a trigger on max connections
- Had to increase the setting several times
- Periodic peeks because of testing from development team

```
$ snmpwalk -v1 -c community aaa-fw-1 1.3.6.1.4.1.2620.1.1.25.3

SNMPv2-SMI::enterprises.2620.1.1.25.3.0 = INTEGER: 49776
```

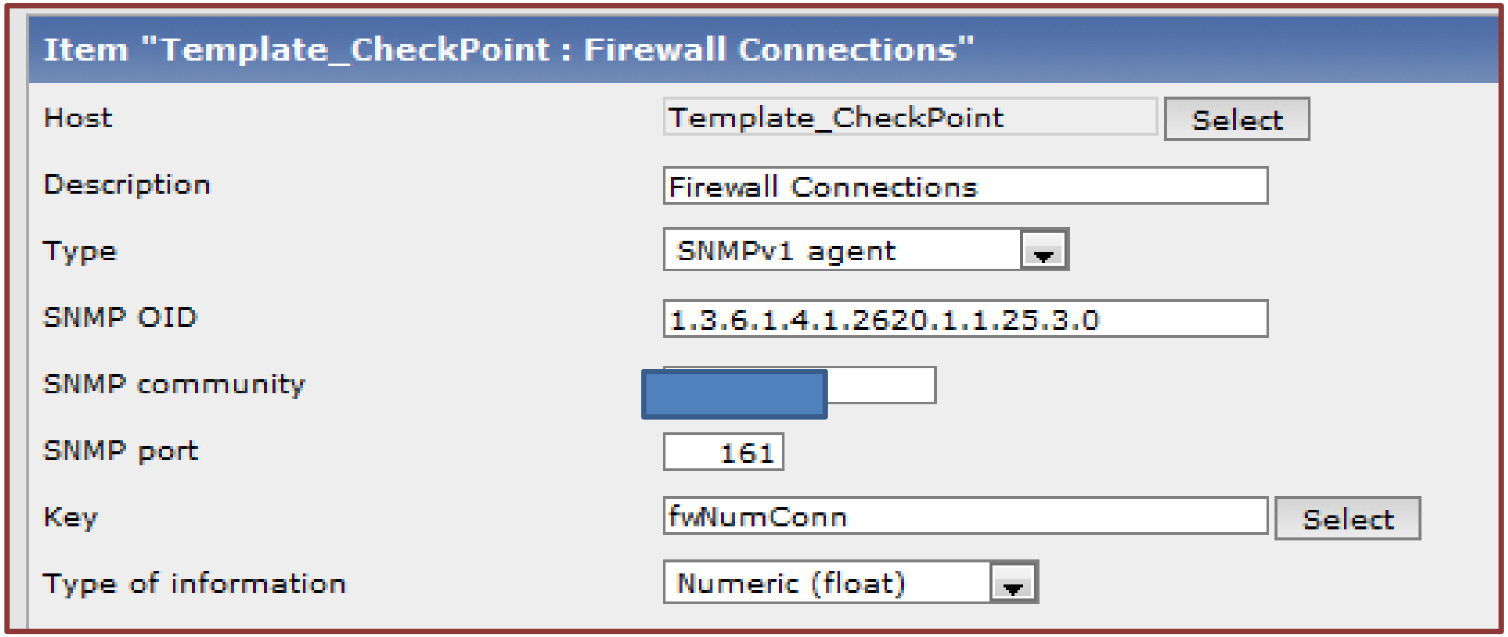

Item "Template_CheckPoint : Firewall Connections"

Host	Template_CheckPoint
Description	Firewall Connections
Type	SNMPv1 agent
SNMP OID	1.3.6.1.4.1.2620.1.1.25.3.0
SNMP community	[redacted]
SNMP port	161
Key	fwNumConn
Type of information	Numeric (float)

The Ghost case … continued

- Started to setup central syslog …
- This message dropped in from local iptable firewall on host1
- Needed to increase the iptable connection tracking table on the host

```
Aug  2 13:09:47 host1 kernel: [8453153.139095] nf_conntrack: table full, dropping packet.

Aug  2 13:09:47 host1 kernel: [8453153.141819] nf_conntrack: table full, dropping packet.
```

EITHER SHE HAD A PSYCHOTIC DECORATOR, OR HER PLACE HAD BEEN RANSACKED BY SOMEONE IN A BIG HURRY.
WELL?! HOW DO YOU EXPLAIN THIS?
THE DAME WAS HYSTERICAL. DAMES USUALLY ARE.

© Universal Press Syndicate

The Ghost case … continued

- Have “integrated” syslog files in Zabbix screens
- URL’s to files on central syslog server
- syslog-today, syslog-yesterday, syslog-weekly
- Will filter syslog on central syslog server and send critical lines to Zabbix server
- Rsyslog has some nice filter options

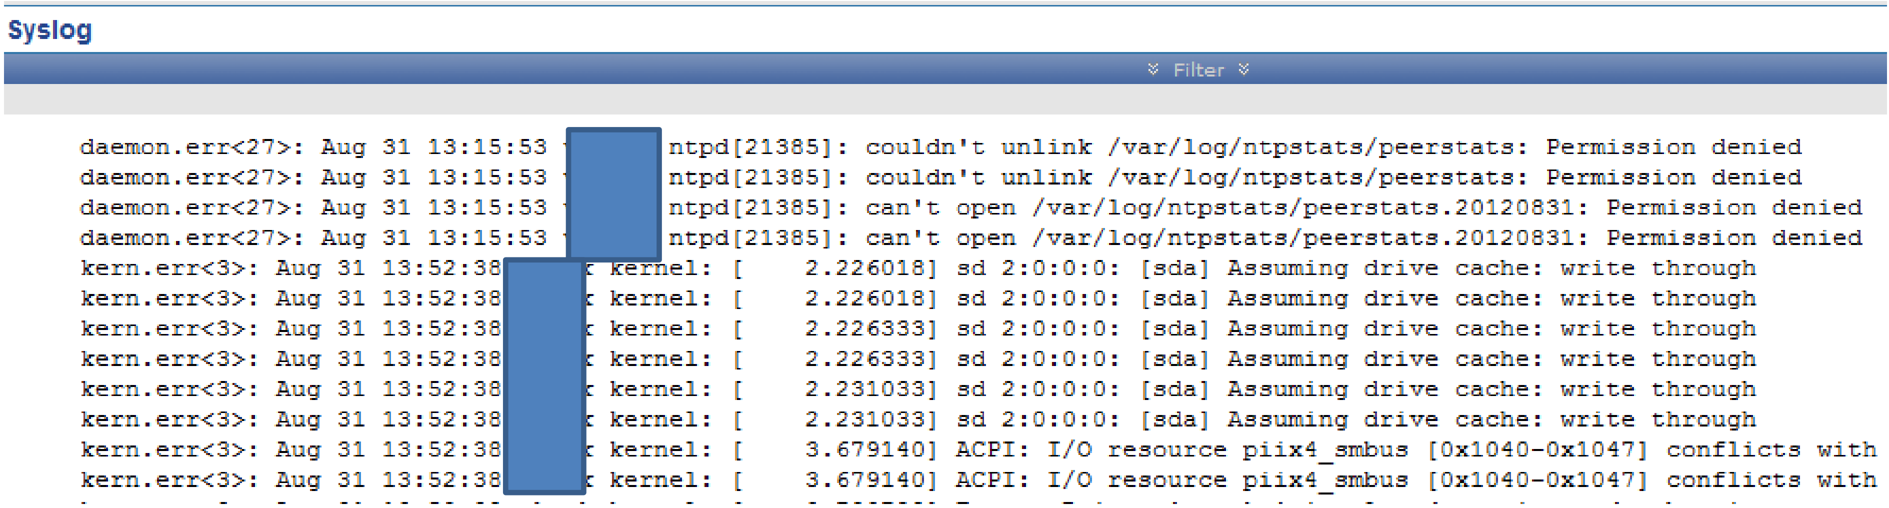

Syslog

```
daemon.err<27>: Aug 31 13:15:53 [redacted] ntpd[21385]: couldn't unlink /var/log/ntpstats/peerstats: Permission denied
daemon.err<27>: Aug 31 13:15:53            ntpd[21385]: couldn't unlink /var/log/ntpstats/peerstats: Permission denied
daemon.err<27>: Aug 31 13:15:53            ntpd[21385]: can't open /var/log/ntpstats/peerstats.20120831: Permission denied
daemon.err<27>: Aug 31 13:15:53            ntpd[21385]: can't open /var/log/ntpstats/peerstats.20120831: Permission denied
kern.err<3>: Aug 31 13:52:38 [redacted] kernel: [    2.226018] sd 2:0:0:0: [sda] Assuming drive cache: write through
kern.err<3>: Aug 31 13:52:38            kernel: [    2.226018] sd 2:0:0:0: [sda] Assuming drive cache: write through
kern.err<3>: Aug 31 13:52:38            kernel: [    2.226333] sd 2:0:0:0: [sda] Assuming drive cache: write through
kern.err<3>: Aug 31 13:52:38            kernel: [    2.226333] sd 2:0:0:0: [sda] Assuming drive cache: write through
kern.err<3>: Aug 31 13:52:38            kernel: [    2.231033] sd 2:0:0:0: [sda] Assuming drive cache: write through
kern.err<3>: Aug 31 13:52:38            kernel: [    2.231033] sd 2:0:0:0: [sda] Assuming drive cache: write through
kern.err<3>: Aug 31 13:52:38            kernel: [    3.679140] ACPI: I/O resource piix4_smbus [0x1040-0x1047] conflicts with
kern.err<3>: Aug 31 13:52:38            kernel: [    3.679140] ACPI: I/O resource piix4_smbus [0x1040-0x1047] conflicts with
```

The Ghost case … continued … cont.

- Later having major problems on host after linux upgrade
- Have to optimize some TCP kernel parameters

- Bug: TCP sockets were not reused in applications (old bug)
- Effect of fixed bug can be seen in graph

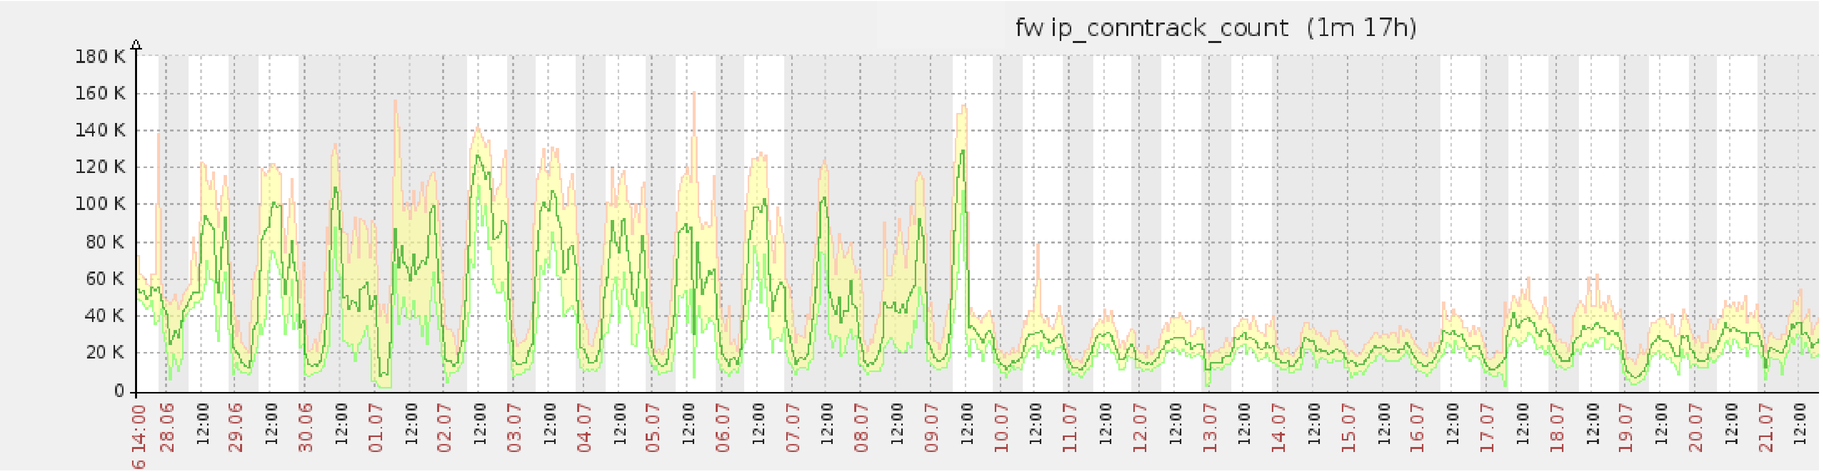

The Java case

- Victim: Dead java application
- Suspect: ??
- Evidence: BIG BIG java log files
- Starting monitoring of java (Tomcat etc)
- Tool: Zapcat JMX Zabbix Bridge

I INTRODUCED THE DAME TO A FRIEND WHO'S VERY CLOSE TO MY HEART. JUST A LITTLE DOWN AND LEFT, TO BE SPECIFIC.
MY FRIEND IS AN ELOQUENT SPEAKER. HE MADE THREE PROFOUND ARGUMENTS WHILE I EXCUSED MYSELF FROM THE ROOM. I ALWAYS LEAVE WHEN THE TALK GETS PHILOSOPHICAL.
!!!
YOU'RE IN REAL TROUBLE NOW, YOUNG MAN!!

© Universal Press Syndicate

The Java case

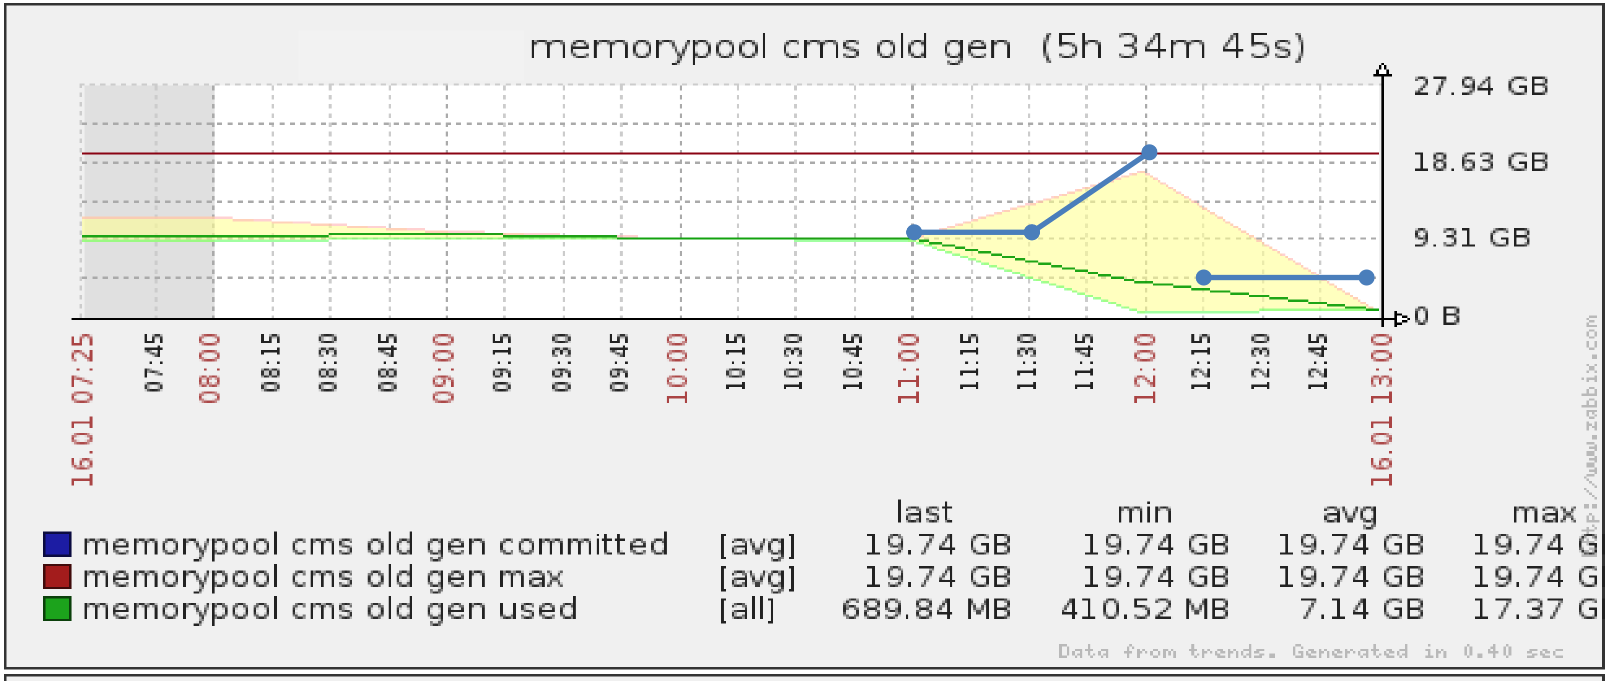

		last	min	avg	max
🟦 memorypool cms old gen committed	[avg]	19.74 GB	19.74 GB	19.74 GB	19.74 G
🟥 memorypool cms old gen max	[avg]	19.74 GB	19.74 GB	19.74 GB	19.74 G
🟩 memorypool cms old gen used	[all]	689.84 MB	410.52 MB	7.14 GB	17.37 G

Data from trends. Generated in 0.40 sec

Observations and actions:

- Short before crime: “cms old gen” went unusually high
- Making trigger on 70% of cms old gen used - alert/early warning
- Using graph to point out time for problem start – problem reporting to development team

The Java case

- **Internal search update routine was found guilty!**
- **Route Cause:** Did not handle timeout from data store system ok
- – start accumulated search over and over again!
- Still using graphs to point out start time for problems

- **Extension of java monitoring and templates**
- Solr and other production data
- Own developed attributes used in new templates

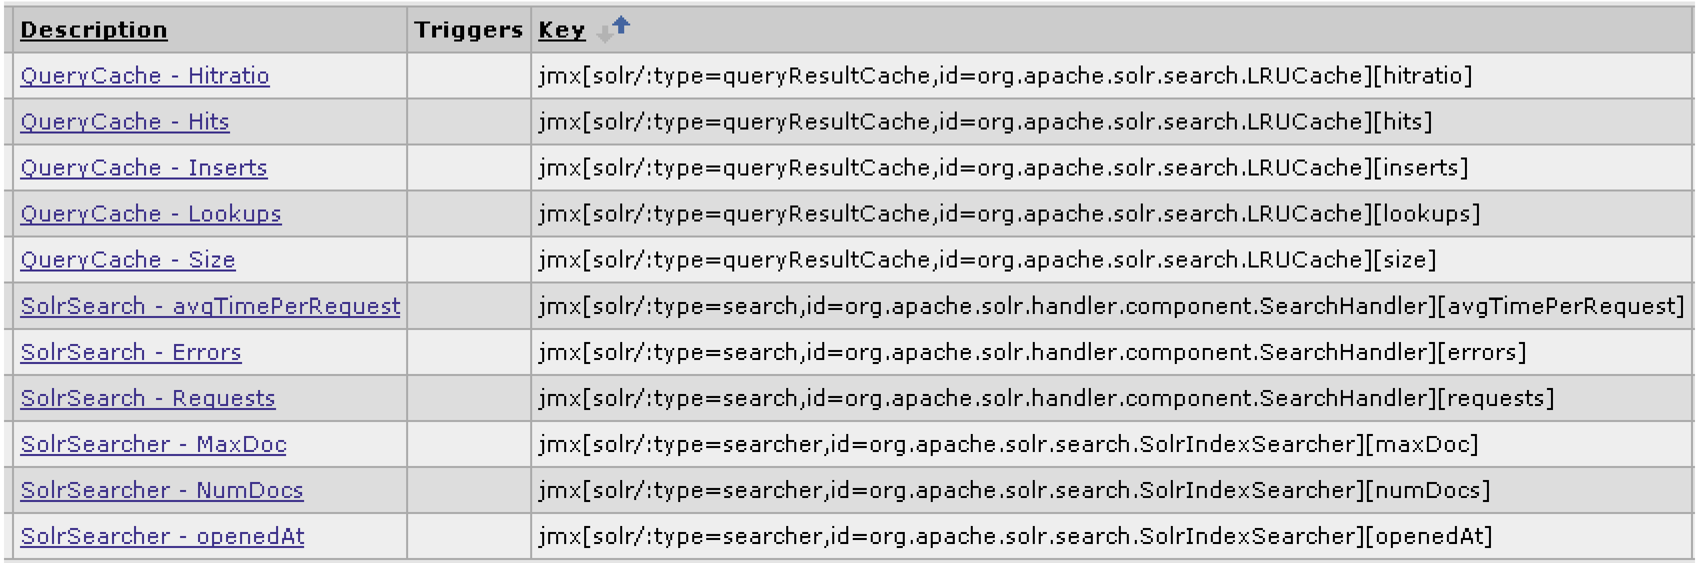

Description	Triggers	Key
QueryCache - Hitratio		jmx[solr/:type=queryResultCache,id=org.apache.solr.search.LRUCache][hitratio]
QueryCache - Hits		jmx[solr/:type=queryResultCache,id=org.apache.solr.search.LRUCache][hits]
QueryCache - Inserts		jmx[solr/:type=queryResultCache,id=org.apache.solr.search.LRUCache][inserts]
QueryCache - Lookups		jmx[solr/:type=queryResultCache,id=org.apache.solr.search.LRUCache][lookups]
QueryCache - Size		jmx[solr/:type=queryResultCache,id=org.apache.solr.search.LRUCache][size]
SolrSearch - avgTimePerRequest		jmx[solr/:type=search,id=org.apache.solr.handler.component.SearchHandler][avgTimePerRequest]
SolrSearch - Errors		jmx[solr/:type=search,id=org.apache.solr.handler.component.SearchHandler][errors]
SolrSearch - Requests		jmx[solr/:type=search,id=org.apache.solr.handler.component.SearchHandler][requests]
SolrSearcher - MaxDoc		jmx[solr/:type=searcher,id=org.apache.solr.search.SolrIndexSearcher][maxDoc]
SolrSearcher - NumDocs		jmx[solr/:type=searcher,id=org.apache.solr.search.SolrIndexSearcher][numDocs]
SolrSearcher - openedAt		jmx[solr/:type=searcher,id=org.apache.solr.search.SolrIndexSearcher][openedAt]

Getting Production data

(Enter the crime scene)

- Discovered that many DBC services were using nearly the same log format
- Python script to parse the logfiles
- Running from cron locally now.
- Moving to central solution, with syslog etc.
- Parse logfile from EOF, because of BIG logfiles (back to *–B sec*.)
- Counting transactions, errors and avg. transactions time
- Sending items to Zabbix server with zabbix_sender

- *parsetracelog.py -r --tag=Sog --sysname=systemX -B 300 -z file.log*

Getting Production data

- Big moment when first real life production data was showed up in the Zabbix GUI !!

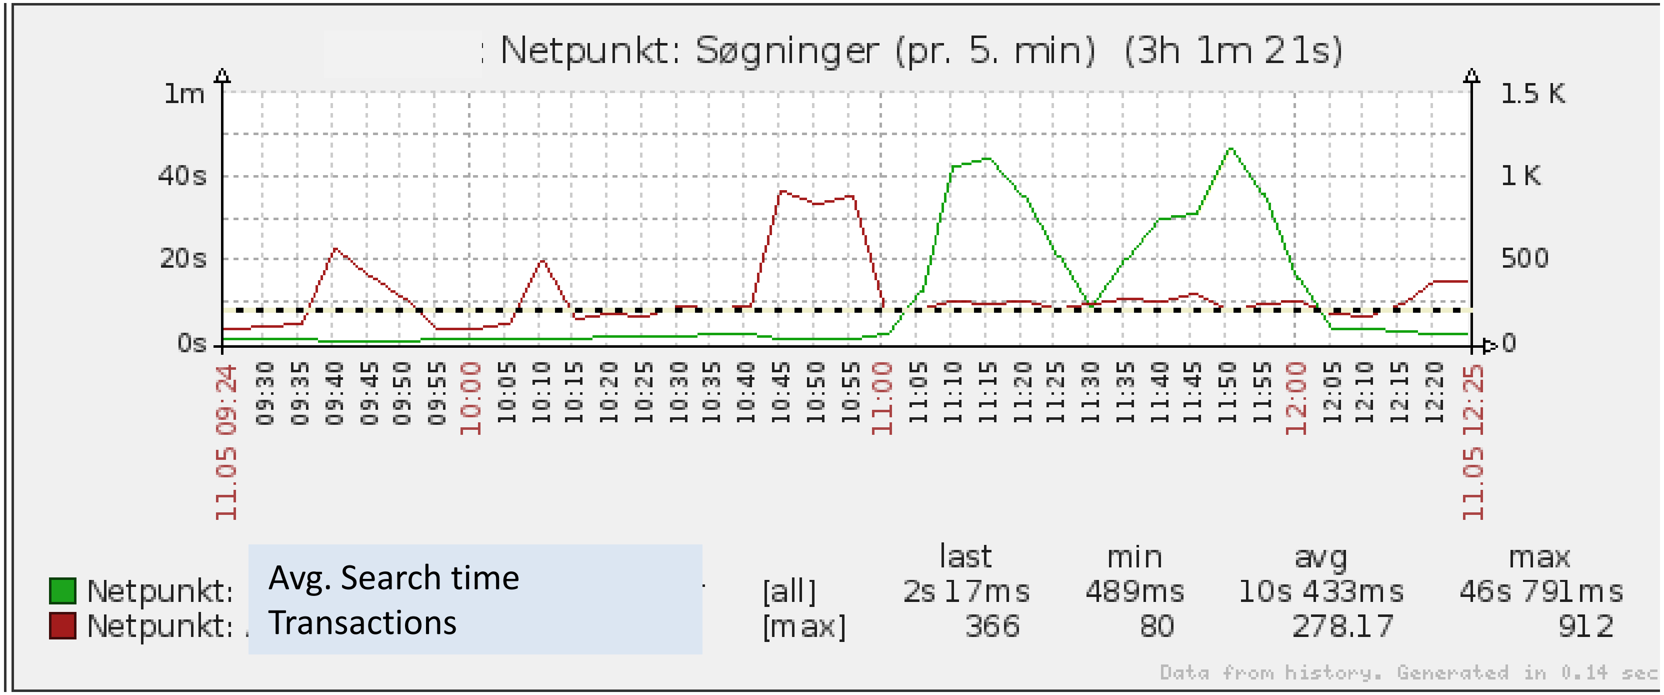

		last	min	avg	max
🟩 Netpunkt: Avg. Search time	[all]	2s 17ms	489ms	10s 433ms	46s 791ms
🟥 Netpunkt: Transactions	[max]	366	80	278.17	912

Data from history. Generated in 0.14 sec

Getting Production data

- Discovered test utility at a workshop
- Also running from cron – every 1 min.
- Test all main backend subsystems and return times for steps
- *yaz-client –f inputfile | grep Elapsed ….*

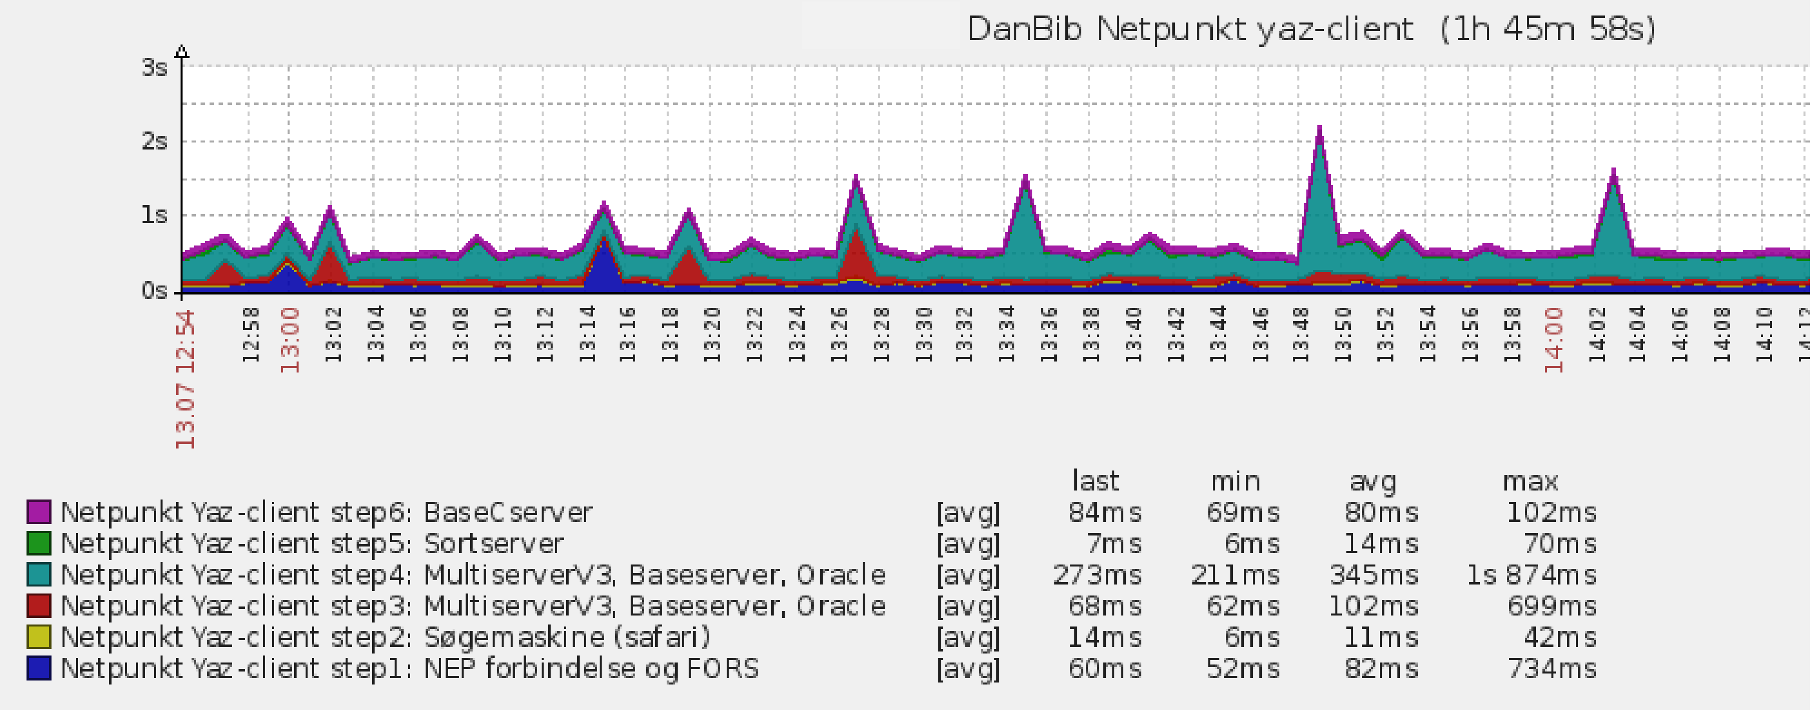

		last	min	avg	max
🟪 Netpunkt Yaz-client step6: BaseCserver	[avg]	84ms	69ms	80ms	102ms
🟩 Netpunkt Yaz-client step5: Sortserver	[avg]	7ms	6ms	14ms	70ms
🟦 Netpunkt Yaz-client step4: MultiserverV3, Baseserver, Oracle	[avg]	273ms	211ms	345ms	1s 874ms
🟥 Netpunkt Yaz-client step3: MultiserverV3, Baseserver, Oracle	[avg]	68ms	62ms	102ms	699ms
🟨 Netpunkt Yaz-client step2: Søgemaskine (safari)	[avg]	14ms	6ms	11ms	42ms
🟦 Netpunkt Yaz-client step1: NEP forbindelse og FORS	[avg]	60ms	52ms	82ms	734ms

Looking for traces

- **Talk with people!**
- Get different people's views and opinions
- … the programmer, the customer, the super user, the application people, the database admins, the end user support team …
- Let them show you how they work with the systems
- Get overview drawings etc.
- Make logs / notes about what you see and find

- **Get data!**
- From logs
- From scripts
- From polling
- From Zabbix database – SQL , event etc.

Take care !!
- With avg. in graphs
- Your statistics

Looking for traces

- My main work board / Tool is **Zabbix Screens**!
- I combine different graphs
- Move around with them, play with them
- Zoom in and out
- Try to make overviews
- Make several different screens, looking at different aspects
- Go down in details at other times

- **In this way I:**
- Try to analyse data and looking for trends
- Try to find patterns and correlate graphs
- Try to isolate the problems

Looking for traces

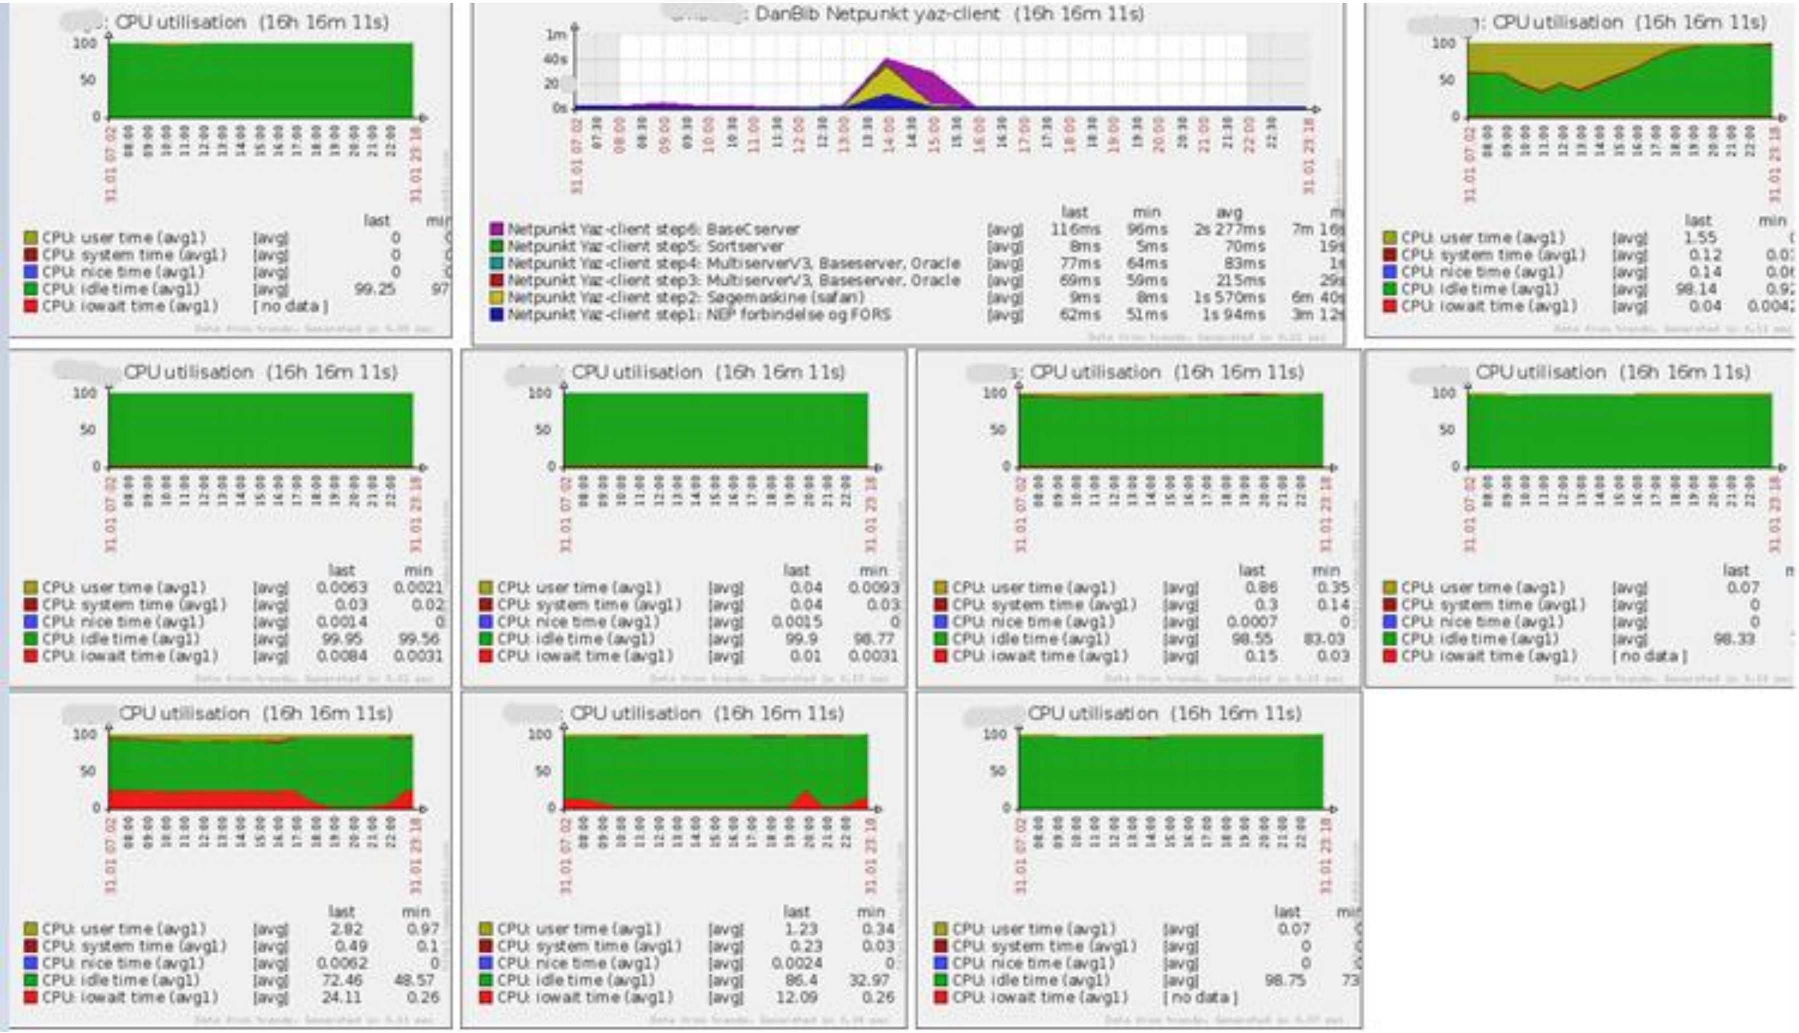
[redacted]: CPU utilisation (16h 16m 11s)
CPU: user time (avg1) [avg] last 0
CPU: system time (avg1) [avg] 0
CPU: nice time (avg1) [avg] 0
CPU: idle time (avg1) [avg] 99.25 97
CPU: iowait time (avg1) [no data]
[redacted]: DanBib Netpunkt yaz-client (16h 16m 11s)
1m
40s
20
0s
last min avg
Netpunkt Yaz-client step6: BaseCserver [avg] 116ms 96ms 2s 277ms 7m 16
Netpunkt Yaz-client step5: Sortserver [avg] 8ms 5ms 70ms 19
Netpunkt Yaz-client step4: MultiserverV3, Baseserver, Oracle [avg] 77ms 64ms 83ms 1
Netpunkt Yaz-client step3: MultiserverV3, Baseserver, Oracle [avg] 69ms 59ms 215ms 29
Netpunkt Yaz-client step2: Søgemaskine (safan) [avg] 9ms 8ms 1s 570ms 6m 40
Netpunkt Yaz-client step1: NEP forbindelse og FORS [avg] 62ms 51ms 1s 94ms 3m 12
[redacted]: CPU utilisation (16h 16m 11s)
last min
CPU: user time (avg1) [avg] 1.55
CPU: system time (avg1) [avg] 0.12 0.0
CPU: nice time (avg1) [avg] 0.14 0.0
CPU: idle time (avg1) [avg] 98.14 0.9
CPU: iowait time (avg1) [avg] 0.04 0.004
[redacted] CPU utilisation (16h 16m 11s)
last min
CPU: user time (avg1) [avg] 0.0063 0.0021
CPU: system time (avg1) [avg] 0.03 0.02
CPU: nice time (avg1) [avg] 0.0014 0
CPU: idle time (avg1) [avg] 99.95 99.56
CPU: iowait time (avg1) [avg] 0.0084 0.0031
[redacted] CPU utilisation (16h 16m 11s)
last min
CPU: user time (avg1) [avg] 0.04 0.0093
CPU: system time (avg1) [avg] 0.04 0.03
CPU: nice time (avg1) [avg] 0.0015 0
CPU: idle time (avg1) [avg] 99.9 98.77
CPU: iowait time (avg1) [avg] 0.01 0.0031
[redacted]: CPU utilisation (16h 16m 11s)
last min
CPU: user time (avg1) [avg] 0.86 0.35
CPU: system time (avg1) [avg] 0.3 0.14
CPU: nice time (avg1) [avg] 0.0007 0
CPU: idle time (avg1) [avg] 98.55 83.03
CPU: iowait time (avg1) [avg] 0.15 0.03
[redacted] CPU utilisation (16h 16m 11s)
last
CPU: user time (avg1) [avg] 0.07
CPU: system time (avg1) [avg] 0
CPU: nice time (avg1) [avg] 0
CPU: idle time (avg1) [avg] 98.33
CPU: iowait time (avg1) [no data]
[redacted] CPU utilisation (16h 16m 11s)
last min
CPU: user time (avg1) [avg] 2.82 0.97
CPU: system time (avg1) [avg] 0.49 0.1
CPU: nice time (avg1) [avg] 0.0062 0
CPU: idle time (avg1) [avg] 72.46 48.57
CPU: iowait time (avg1) [avg] 24.11 0.26
[redacted] CPU utilisation (16h 16m 11s)
last min
CPU: user time (avg1) [avg] 1.23 0.34
CPU: system time (avg1) [avg] 0.23 0.03
CPU: nice time (avg1) [avg] 0.0024 0
CPU: idle time (avg1) [avg] 86.4 32.97
CPU: iowait time (avg1) [avg] 12.09 0.26
[redacted] CPU utilisation (16h 16m 11s)
last
CPU: user time (avg1) [avg] 0.07
CPU: system time (avg1) [avg] 0
CPU: nice time (avg1) [avg] 0
CPU: idle time (avg1) [avg] 98.75 73
CPU: iowait time (avg1) [no data]
100
50
0
31.01 07:02
31.01 23:18

Looking for traces

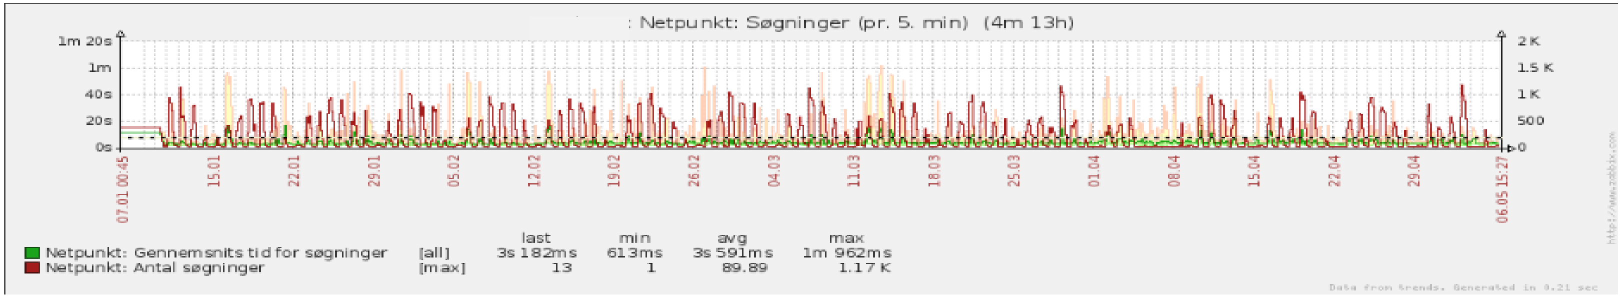

: Netpunkt: Søgninger (pr. 5. min) (4m 13h)

		last	min	avg	max
🟩 Netpunkt: Gennemsnits tid for søgninger	[all]	3s 182ms	613ms	3s 591ms	1m 962ms
🟥 Netpunkt: Antal søgninger	[max]	13	1	89.89	1.17 K

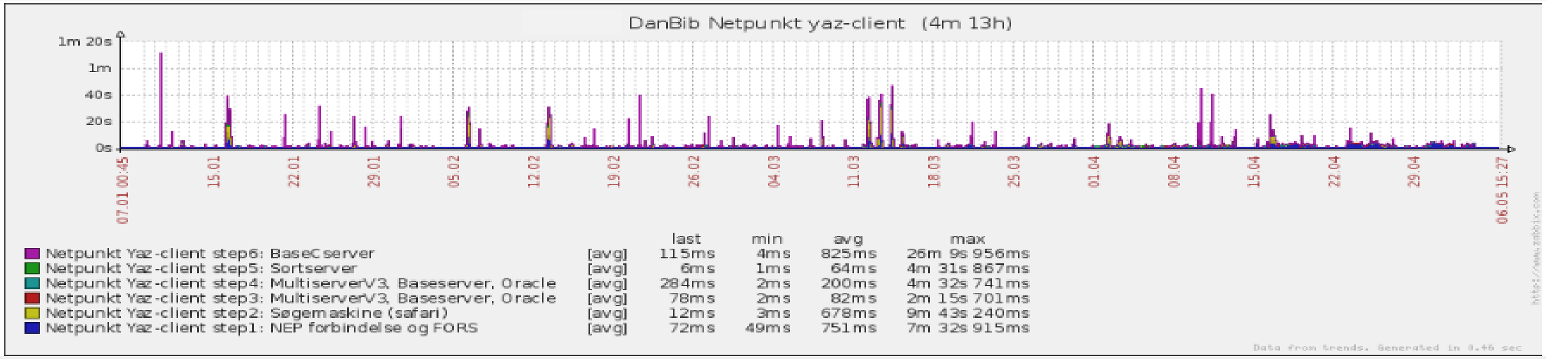

DanBib Netpunkt yaz-client (4m 13h)

		last	min	avg	max
🟪 Netpunkt Yaz-client step6: BaseCserver	[avg]	115ms	4ms	825ms	26m 9s 956ms
🟩 Netpunkt Yaz-client step5: Sortserver	[avg]	6ms	1ms	64ms	4m 31s 867ms
🟦 Netpunkt Yaz-client step4: MultiserverV3, Baseserver, Oracle	[avg]	284ms	2ms	200ms	4m 32s 741ms
🟥 Netpunkt Yaz-client step3: MultiserverV3, Baseserver, Oracle	[avg]	78ms	2ms	82ms	2m 15s 701ms
🟨 Netpunkt Yaz-client step2: Søgemaskine (safari)	[avg]	12ms	3ms	678ms	9m 43s 240ms
🟦 Netpunkt Yaz-client step1: NEP forbindelse og FORS	[avg]	72ms	49ms	751ms	7m 32s 915ms

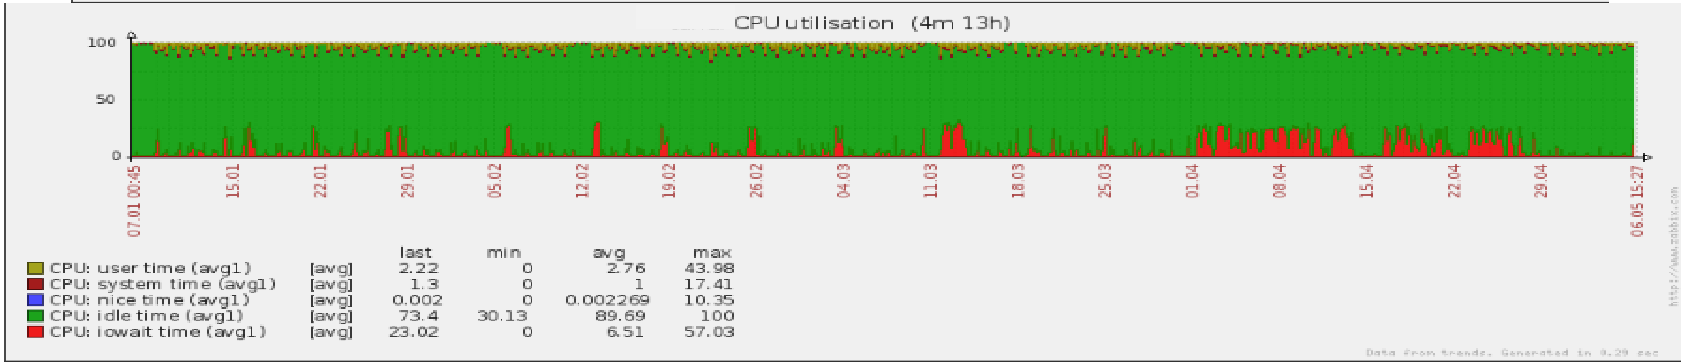

CPU utilisation (4m 13h)

		last	min	avg	max
🟨 CPU: user time (avg1)	[avg]	2.22	0	2.76	43.98
🟥 CPU: system time (avg1)	[avg]	1.3	0	1	17.41
🟦 CPU: nice time (avg1)	[avg]	0.002	0	0.002269	10.35
🟩 CPU: idle time (avg1)	[avg]	73.4	30.13	89.69	100
🟥 CPU: iowait time (avg1)	[avg]	23.02	0	6.51	57.03

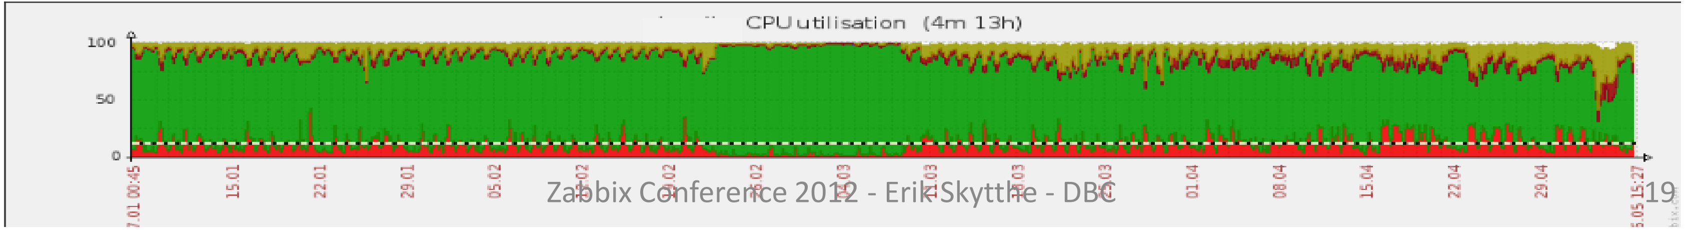

CPU utilisation (4m 13h)

Looking for traces

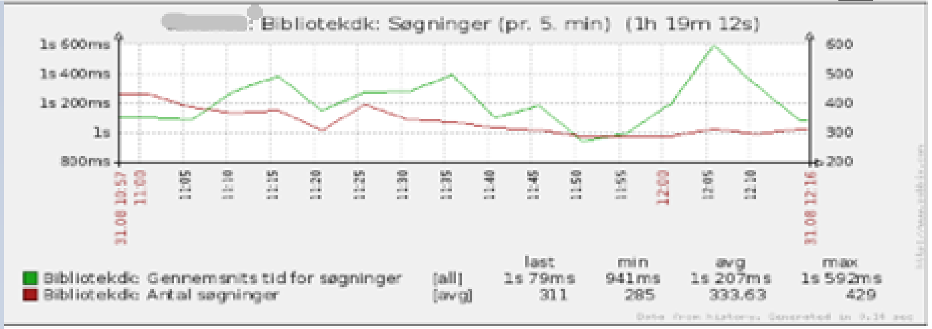

[redacted]: Bibliotekdk: Søgninger (pr. 5. min) (1h 19m 12s)

		last	min	avg	max
🟩 Bibliotekdk: Gennemsnits tid for søgninger	[all]	1s 79ms	941ms	1s 207ms	1s 592ms
🟥 Bibliotekdk: Antal søgninger	[avg]	311	285	333.63	429

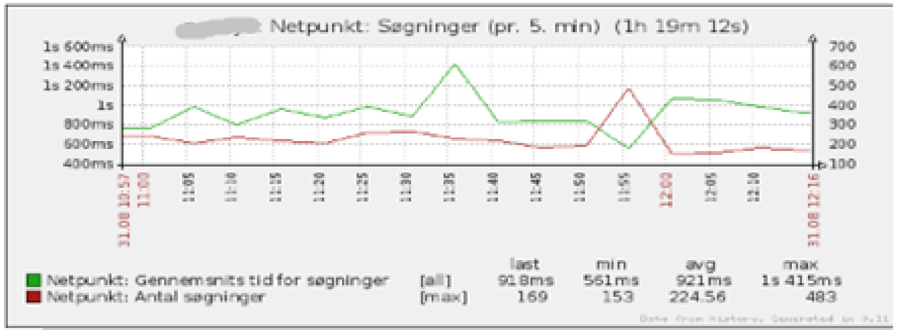

[redacted] Netpunkt: Søgninger (pr. 5. min) (1h 19m 12s)

		last	min	avg	max
🟩 Netpunkt: Gennemsnits tid for søgninger	[all]	918ms	561ms	921ms	1s 415ms
🟥 Netpunkt: Antal søgninger	[max]	169	153	224.56	483

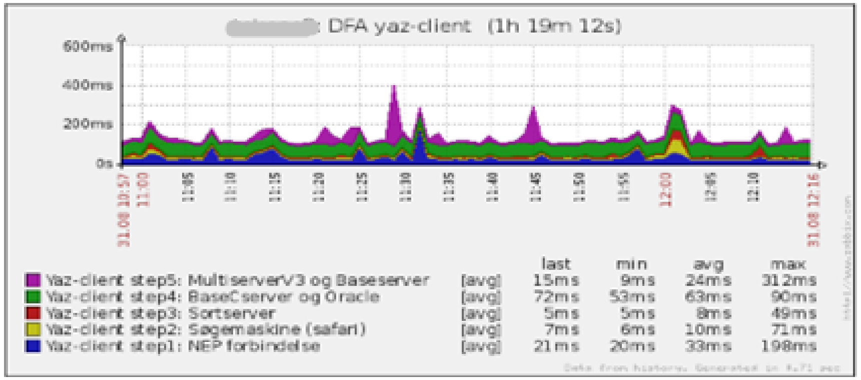

[redacted]: DFA yaz-client (1h 19m 12s)

		last	min	avg	max
🟪 Yaz-client step5: MultiserverV3 og Baseserver	[avg]	15ms	9ms	24ms	312ms
🟩 Yaz-client step4: BaseCserver og Oracle	[avg]	72ms	53ms	63ms	90ms
🟥 Yaz-client step3: Sortserver	[avg]	5ms	5ms	8ms	49ms
🟨 Yaz-client step2: Søgemaskine (safari)	[avg]	7ms	6ms	10ms	71ms
🟦 Yaz-client step1: NEP forbindelse	[avg]	21ms	20ms	33ms	198ms

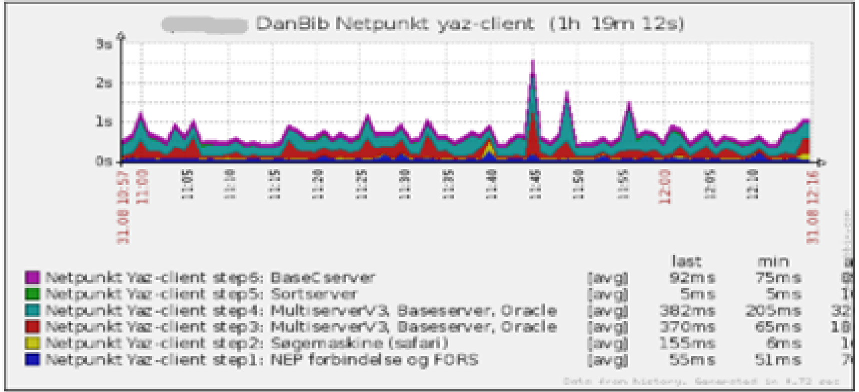

[redacted] DanBib Netpunkt yaz-client (1h 19m 12s)

		last	min
🟪 Netpunkt Yaz-client step6: BaseCserver	[avg]	92ms	75ms
🟩 Netpunkt Yaz-client step5: Sortserver	[avg]	5ms	5ms
🟦 Netpunkt Yaz-client step4: MultiserverV3, Baseserver, Oracle	[avg]	382ms	205ms
🟥 Netpunkt Yaz-client step3: MultiserverV3, Baseserver, Oracle	[avg]	370ms	65ms
🟨 Netpunkt Yaz-client step2: Søgemaskine (safari)	[avg]	155ms	6ms
🟦 Netpunkt Yaz-client step1: NEP forbindelse og FORS	[avg]	55ms	51ms

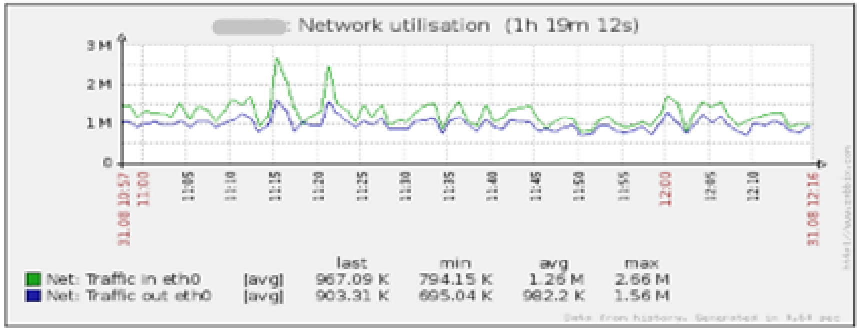

[redacted]: Network utilisation (1h 19m 12s)

		last	min	avg	max
🟩 Net: Traffic in eth0	[avg]	967.09 K	794.15 K	1.26 M	2.66 M
🟦 Net: Traffic out eth0	[avg]	903.31 K	695.04 K	982.2 K	1.56 M

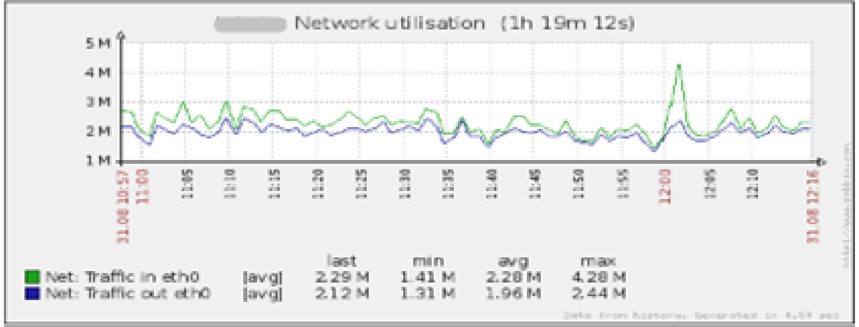

[redacted] Network utilisation (1h 19m 12s)

		last	min	avg	max
🟩 Net: Traffic in eth0	[avg]	2.29 M	1.41 M	2.28 M	4.28 M
🟦 Net: Traffic out eth0	[avg]	2.12 M	1.31 M	1.96 M	2.44 M

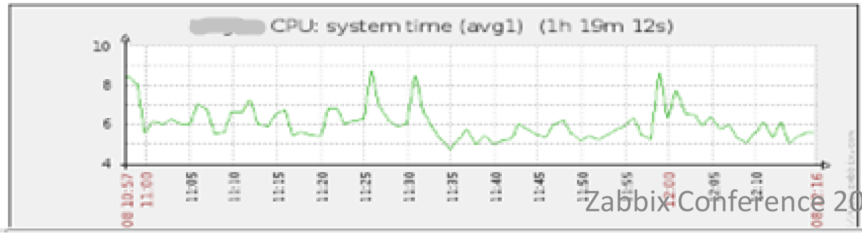

[redacted] CPU: system time (avg1) (1h 19m 12s)

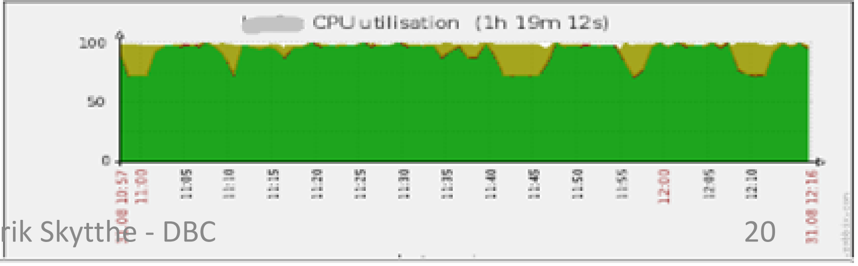

[redacted] CPU utilisation (1h 19m 12s)

The “Wednesday monster” case

- Victims: Users of application.
- **Details:**
- On most Wednesdays, major performance problems, or even service breakdown are seen. Sometimes it is Thursday instead.

- Suspects: Backend, frontend systems?
- Evidence: Trading graphs, monitoring alarms

CALVIN, YOU'RE NOT PAYING ATTENTION AGAIN!
7 +2
SPACEMAN SPIFF, CONQUEROR OF THE COSMOS, IS TRAPPED BY A HIDEOUS ZONDARG!

© Universal Press Syndicate

The “Wednesday monster” case

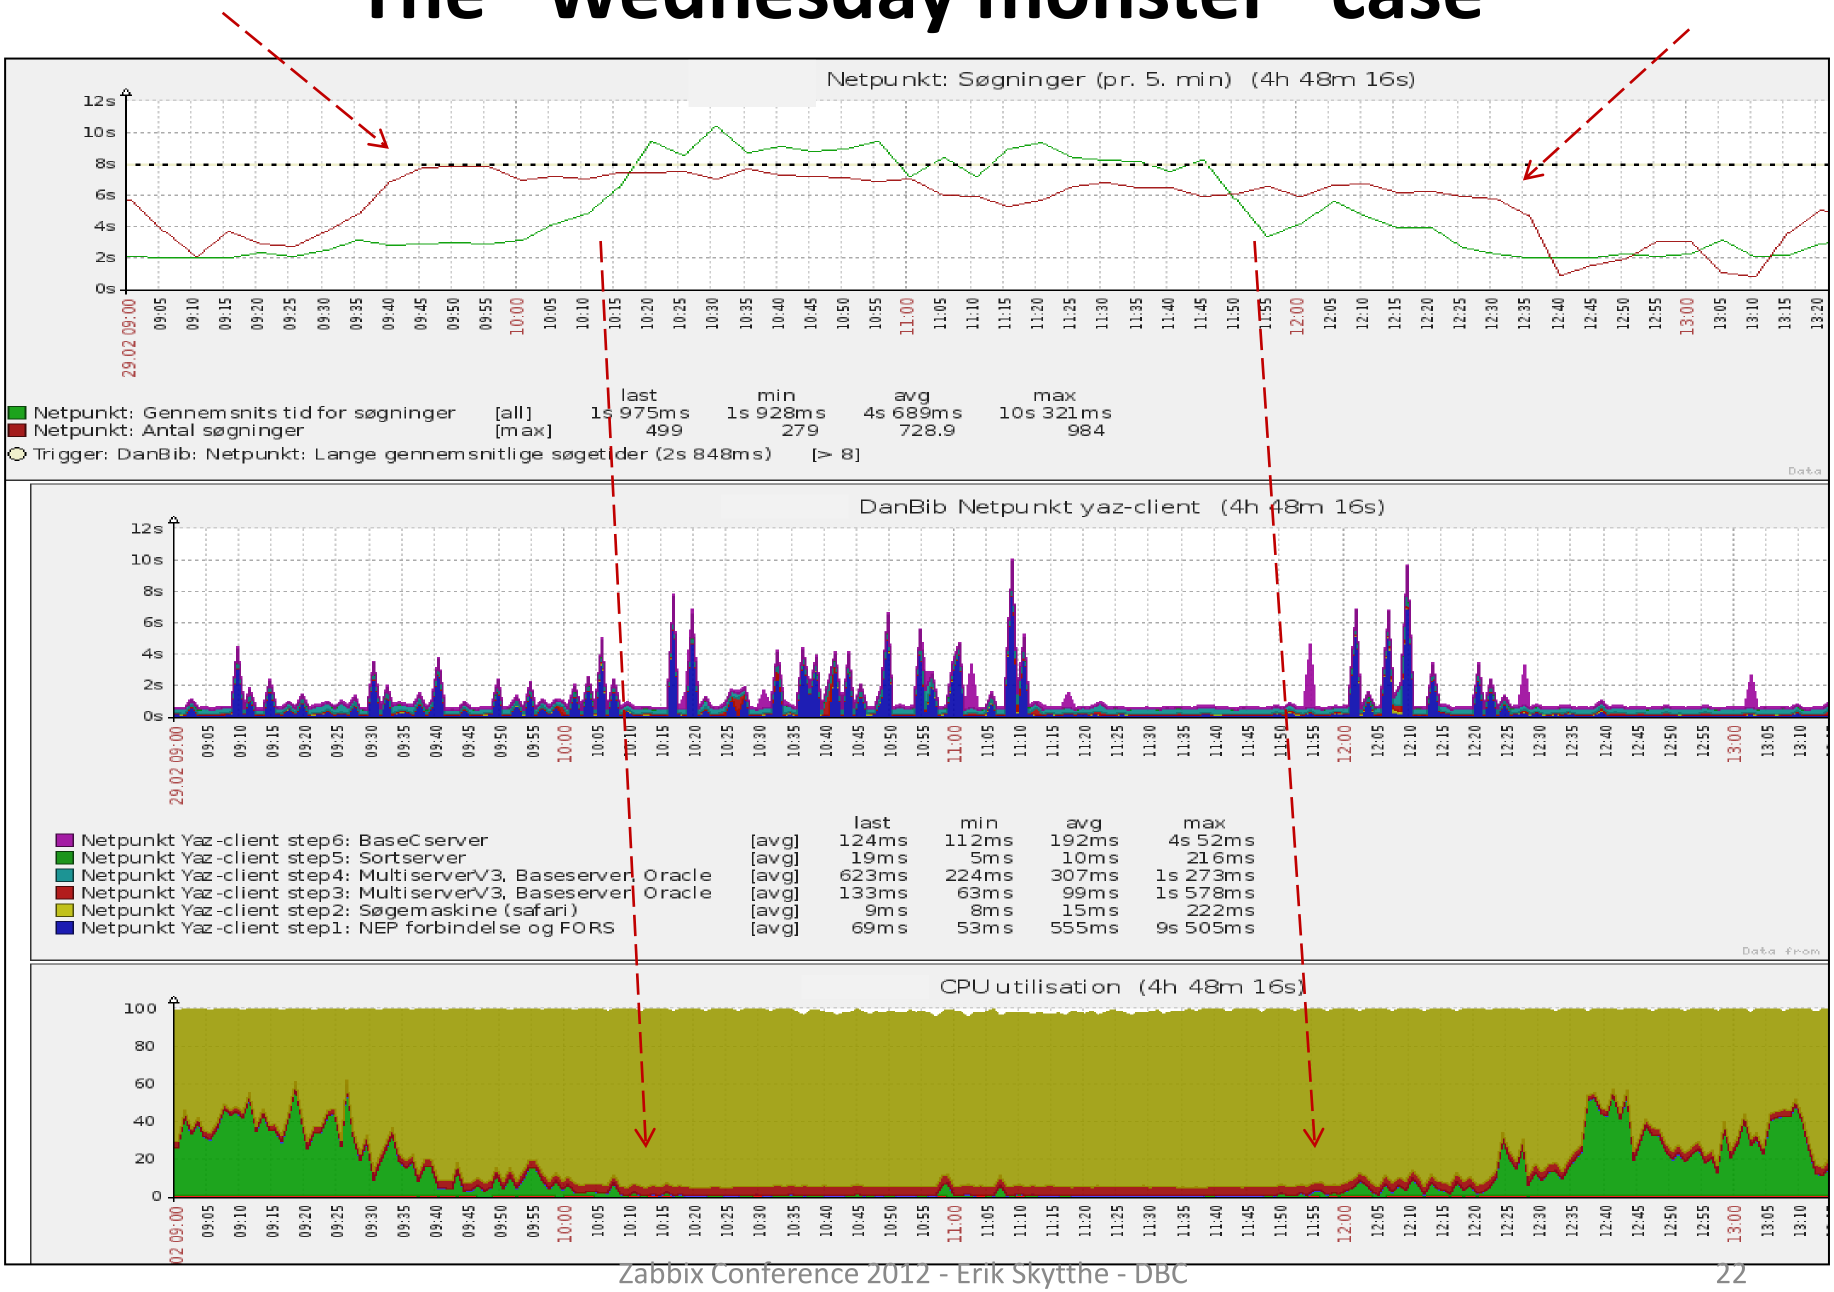

		last	min	avg	max
🟩 Netpunkt: Gennemsnits tid for søgninger	[all]	1s 975ms	1s 928ms	4s 689ms	10s 321ms
🟥 Netpunkt: Antal søgninger	[max]	499	279	728.9	984

○ Trigger: DanBib: Netpunkt: Lange gennemsnitlige søgetider (2s 848ms) [> 8]

		last	min	avg	max
🟪 Netpunkt Yaz-client step6: BaseCserver	[avg]	124ms	112ms	192ms	4s 52ms
🟩 Netpunkt Yaz-client step5: Sortserver	[avg]	19ms	5ms	10ms	216ms
▨ Netpunkt Yaz-client step4: MultiserverV3, Baseserver, Oracle	[avg]	623ms	224ms	307ms	1s 273ms
🟥 Netpunkt Yaz-client step3: MultiserverV3, Baseserver, Oracle	[avg]	133ms	63ms	99ms	1s 578ms
🟨 Netpunkt Yaz-client step2: Søgemaskine (safari)	[avg]	9ms	8ms	15ms	222ms
🟦 Netpunkt Yaz-client step1: NEP forbindelse og FORS	[avg]	69ms	53ms	555ms	9s 505ms

The “Wednesday monster” case

- **Graph conclusions:**
- Major increase in amount of transactions for period
- Correlation between “bad” avg. search time and user CPU time at frontend host
- User CPU time consumption is HEAVY …
- … But is it because of slow backend systems or ???
- … does not seem so – the middle graph is showing backend subsystems. No clear correlation there.
- (Well, middle graph is polled test values, but should be more clear if correlations – Upper graph have higher weight)
- **Other conclusions:**
- Twin system does not show any problems … so no general infrastructure problems … (problem isolation)

The “Wednesday monster” case

- Zoomed out 4 weeks – only looking at transactions
- Seems like Wednesday hang out

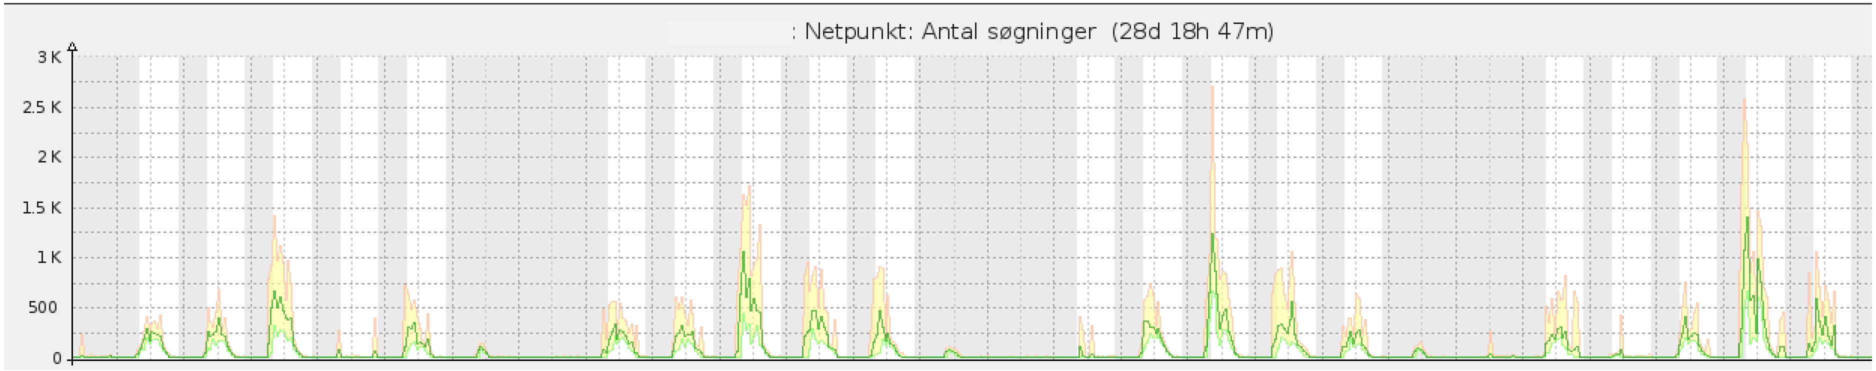

The “Wednesday monster” case

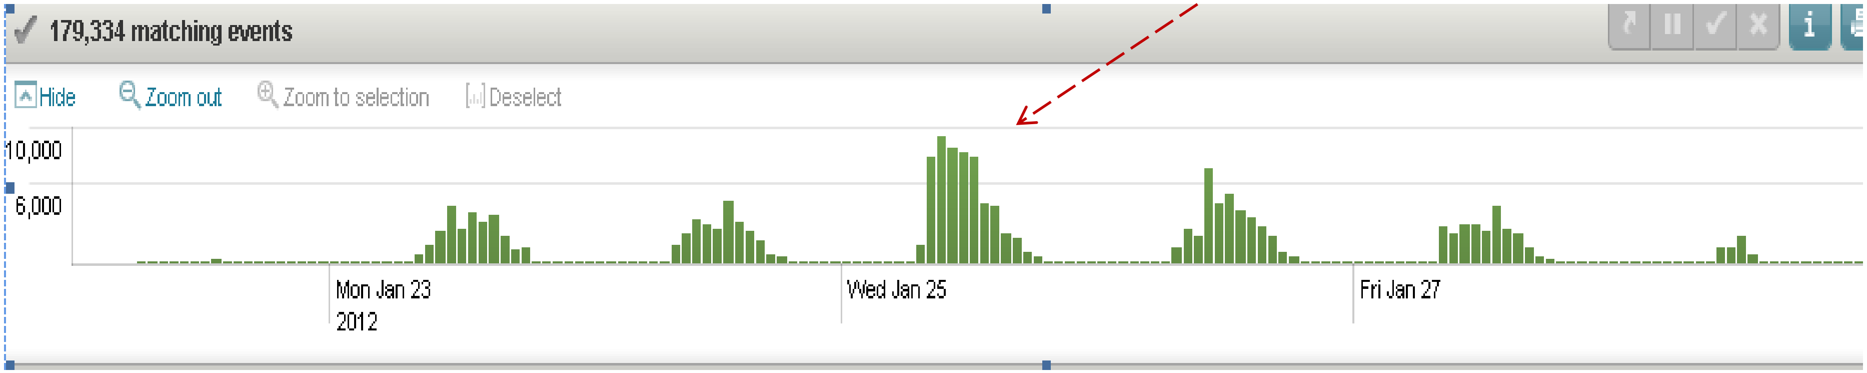

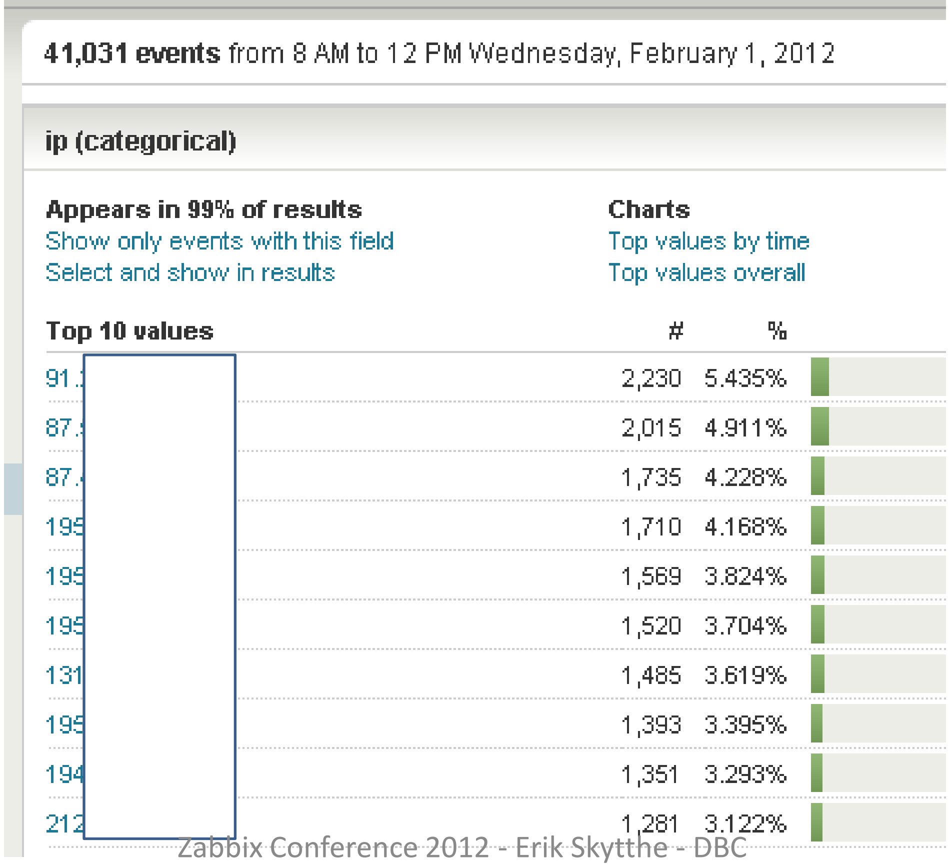
41,031 events from 8 AM to 12 PM Wednesday, February 1, 2012

ip (categorical)

Appears in 99% of results
Show only events with this field
Select and show in results

Charts
Top values by time
Top values overall

Top 10 values	#	%
[redacted]	2,230	5.435%
[redacted]	2,015	4.911%
[redacted]	1,735	4.228%
[redacted]	1,710	4.168%
[redacted]	1,569	3.824%
[redacted]	1,520	3.704%
[redacted]	1,485	3.619%
[redacted]	1,393	3.395%
[redacted]	1,351	3.293%
[redacted]	1,281	3.122%

The “Wednesday monster” case

- Did a deeper analysis of the transaction file
- Using “Splunk” - a not so free log analyzer tool (other can do)

- Looking into Wednesday …
- … Top10 source ip ~ more then 40 % of transactions
- Looking at one Top10 source ip …
- … Most transactions on Wednesday!

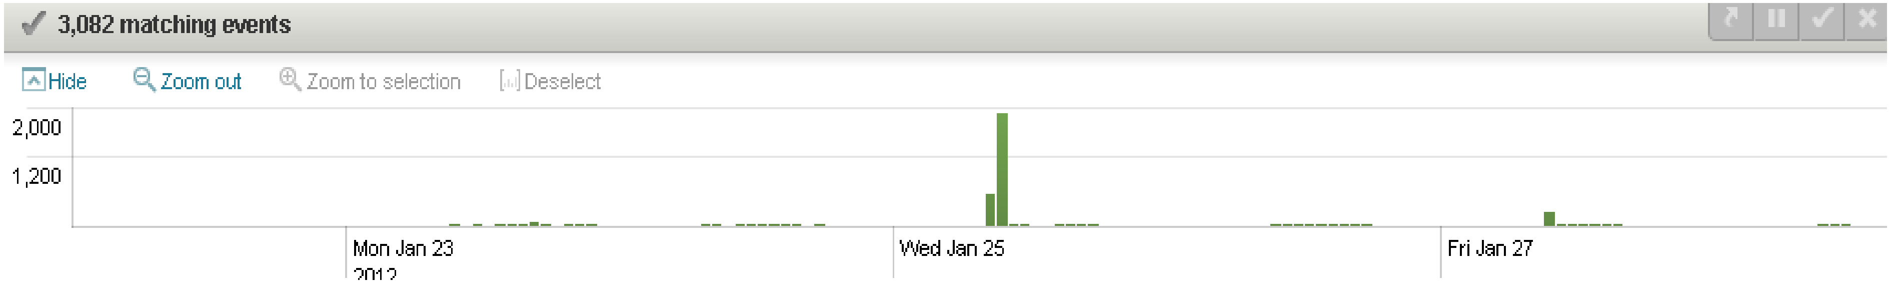

3,082 matching events

The “Wednesday monster” case

- **After some talk with application people:**
- Every Tuesday night there is a update of the central database
- Some local library systems do a data harvest the day after, via DBC web frontend system (to update their local databases)
- Yes Wednesday !!!
- … and sometime Thursday instead
- … if DBC central update have been delayed !

- **Case closed:**
- Customers were found guilty ☺ or ?
- No **root cause was old slow server hardware** at frontend host
- … problem solved after hardware upgrade

The “Search monster” case

- Victims: Users of application.
- **Details:**
- General performance problems in intervals of 5 – 30 min
- Some days or even weeks there is no problems at all.

- Suspects: Backend or frontend systems or?
- Evidence: Trading graphs, monitoring alarms

VERY AMUSING, YOU TWISTED SPACE FROG. WHAT'S *THIS* FIENDISH DEVICE CALLED?
A CHIN-UP BAR. GET ON IT.
SPIFF READIES HIS DARING ESCAPE...
GYM INSTRUCTOR
WATTERSON
12-2
©1988 Universal Press Syndicate

The “Search monster” case

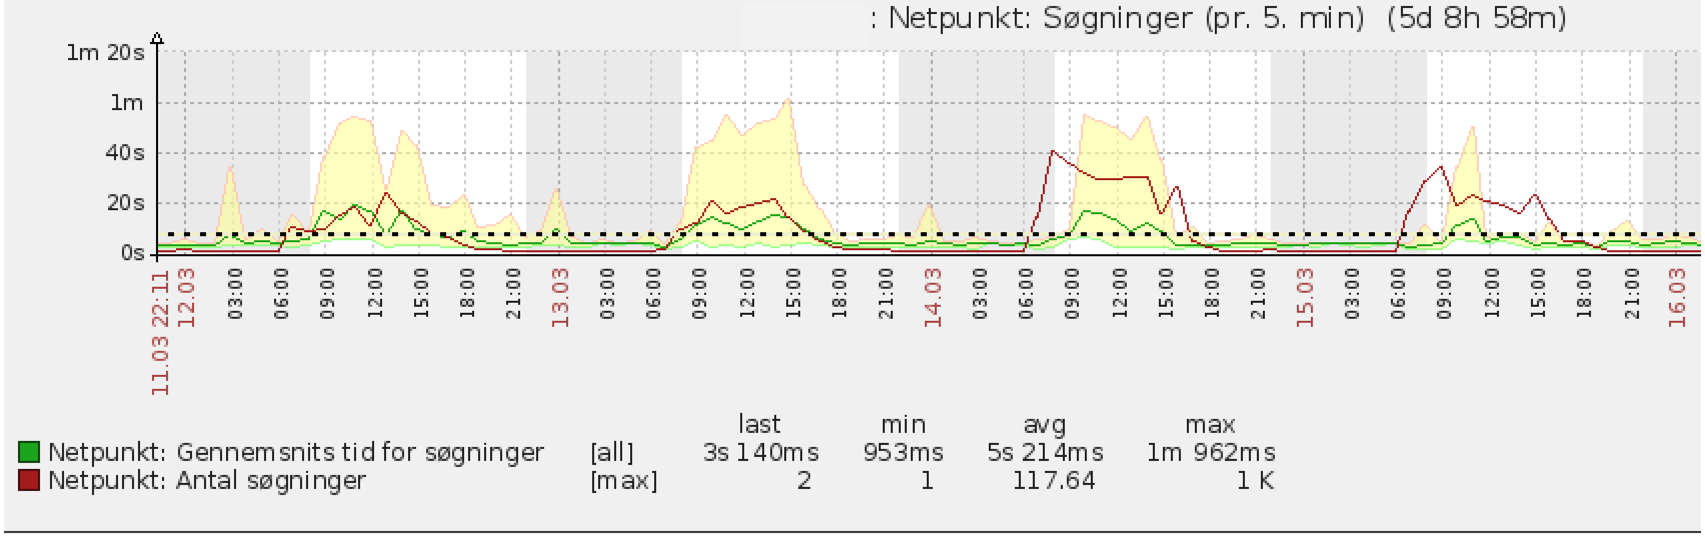

			last	min	avg	max
🟩	Netpunkt: Gennemsnits tid for søgninger	[all]	3s 140ms	953ms	5s 214ms	1m 962ms
🟥	Netpunkt: Antal søgninger	[max]	2	1	117.64	1 K

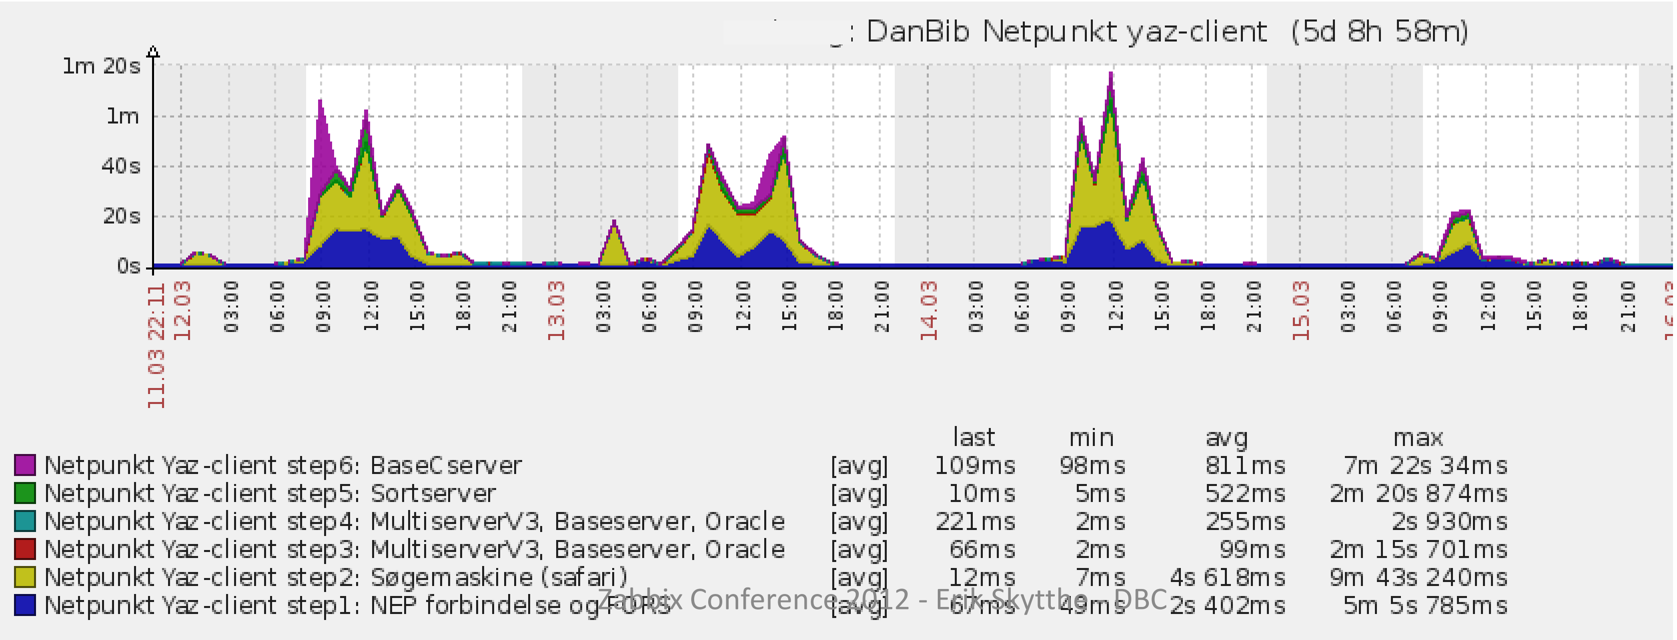

			last	min	avg	max
🟪	Netpunkt Yaz-client step6: BaseCserver	[avg]	109ms	98ms	811ms	7m 22s 34ms
🟩	Netpunkt Yaz-client step5: Sortserver	[avg]	10ms	5ms	522ms	2m 20s 874ms
🟦	Netpunkt Yaz-client step4: MultiserverV3, Baseserver, Oracle	[avg]	221ms	2ms	255ms	2s 930ms
🟥	Netpunkt Yaz-client step3: MultiserverV3, Baseserver, Oracle	[avg]	66ms	2ms	99ms	2m 15s 701ms
🟨	Netpunkt Yaz-client step2: Søgemaskine (safari)	[avg]	12ms	7ms	4s 618ms	9m 43s 240ms
🟦	Netpunkt Yaz-client step1: NEP forbindelse og FORS	[avg]	67ms	49ms	2s 402ms	5m 5s 785ms

The “Search monster” case

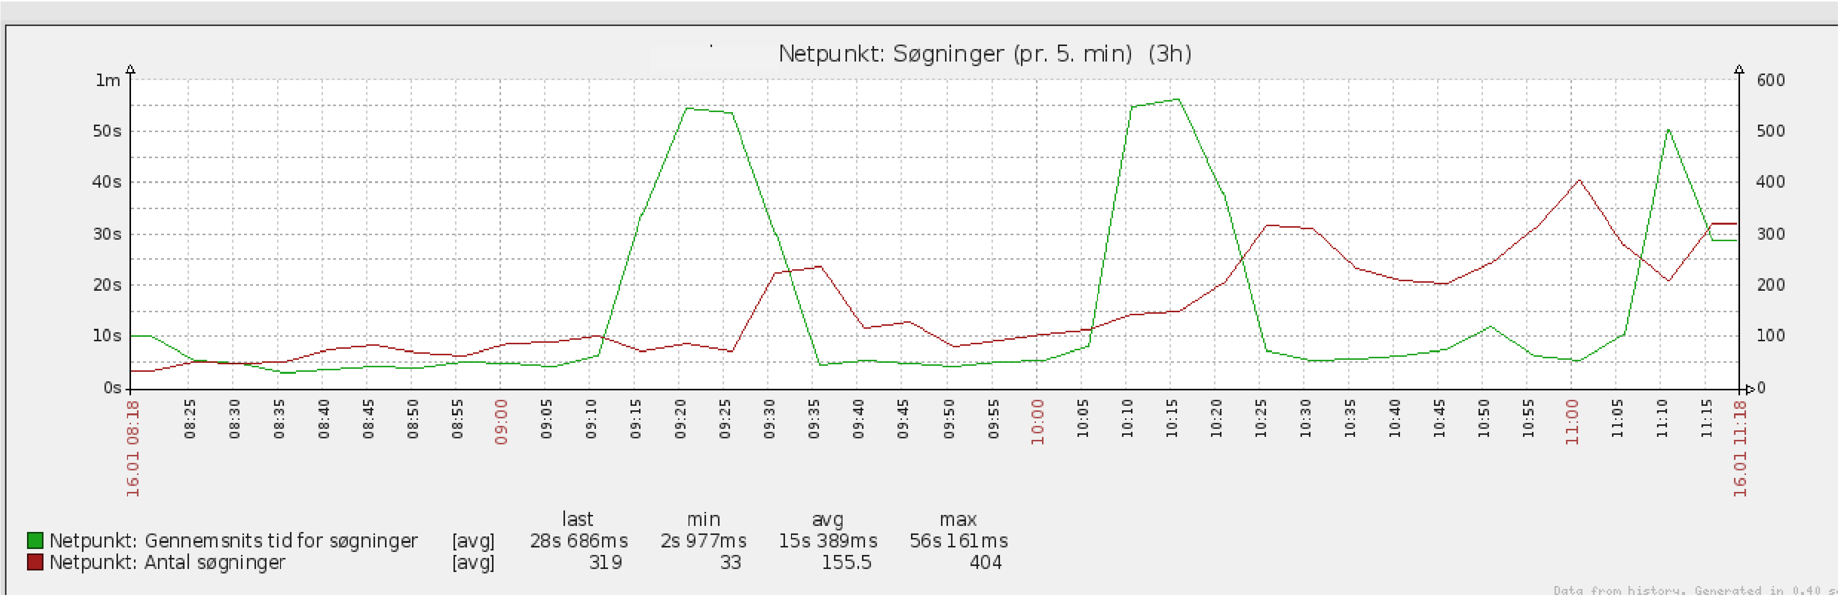

		last	min	avg	max
🟩 Netpunkt: Gennemsnits tid for søgninger	[avg]	28s 686ms	2s 977ms	15s 389ms	56s 161ms
🟥 Netpunkt: Antal søgninger	[avg]	319	33	155.5	404

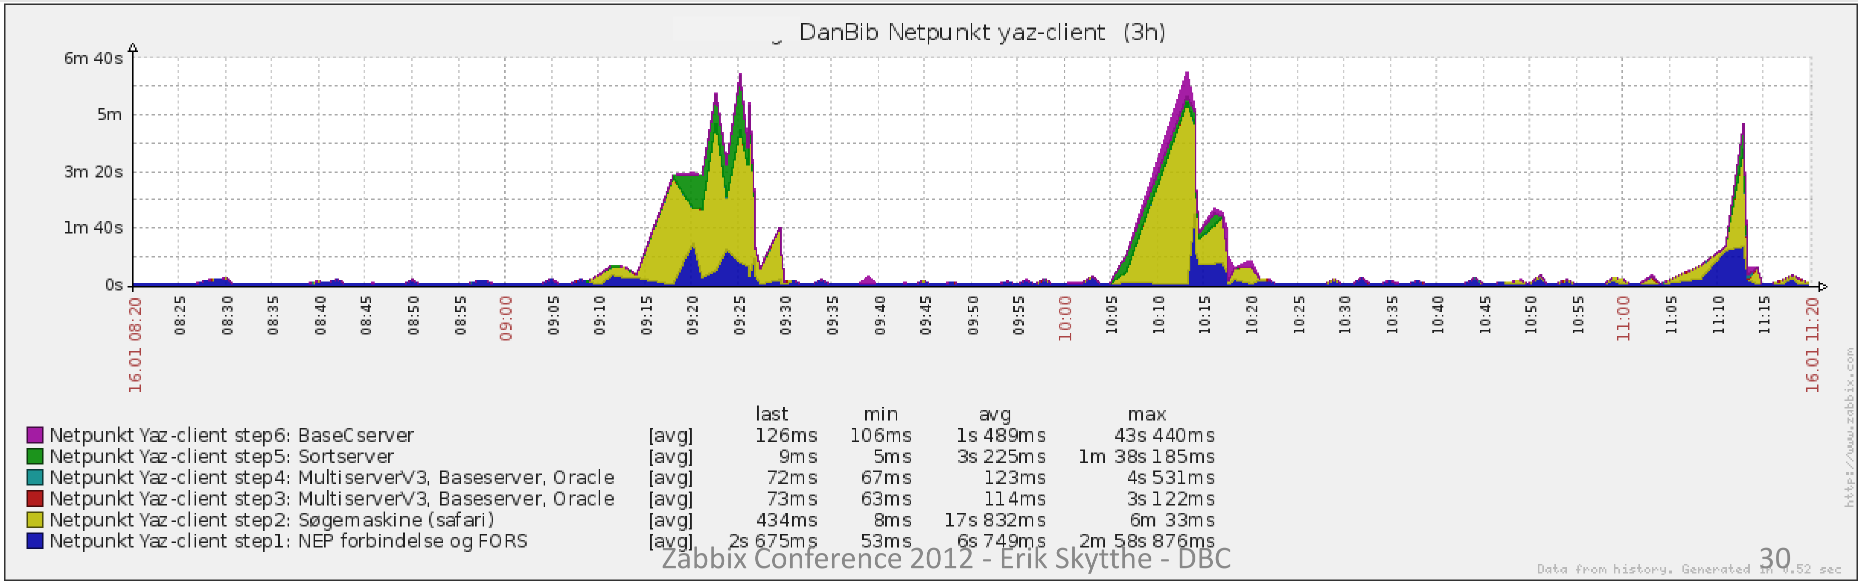

		last	min	avg	max
🟪 Netpunkt Yaz-client step6: BaseCserver	[avg]	126ms	106ms	1s 489ms	43s 440ms
🟩 Netpunkt Yaz-client step5: Sortserver	[avg]	9ms	5ms	3s 225ms	1m 38s 185ms
🟦 Netpunkt Yaz-client step4: MultiserverV3, Baseserver, Oracle	[avg]	72ms	67ms	123ms	4s 531ms
🟥 Netpunkt Yaz-client step3: MultiserverV3, Baseserver, Oracle	[avg]	73ms	63ms	114ms	3s 122ms
🟨 Netpunkt Yaz-client step2: Søgemaskine (safari)	[avg]	434ms	8ms	17s 832ms	6m 33ms
🟦 Netpunkt Yaz-client step1: NEP forbindelse og FORS	[avg]	2s 675ms	53ms	6s 749ms	2m 58s 876ms

The “Search monster” case

- **Graph week overview:**
- Clear correlation
- Search subsystem hangout
- Small BaseCserver subsystem peak Monday
- Not as bad as it looked like – not whole day
  … do a zoom in

- **Graph hour:**
- Search subsystem peeks of 10 – 15 min
- Clear correlation with bad performance (avg. transaction)

The “Search monster” case

- **Update:**
- **Commit:**

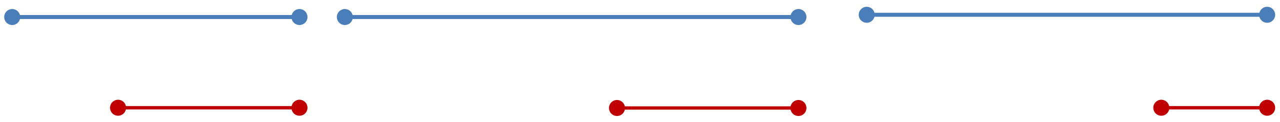

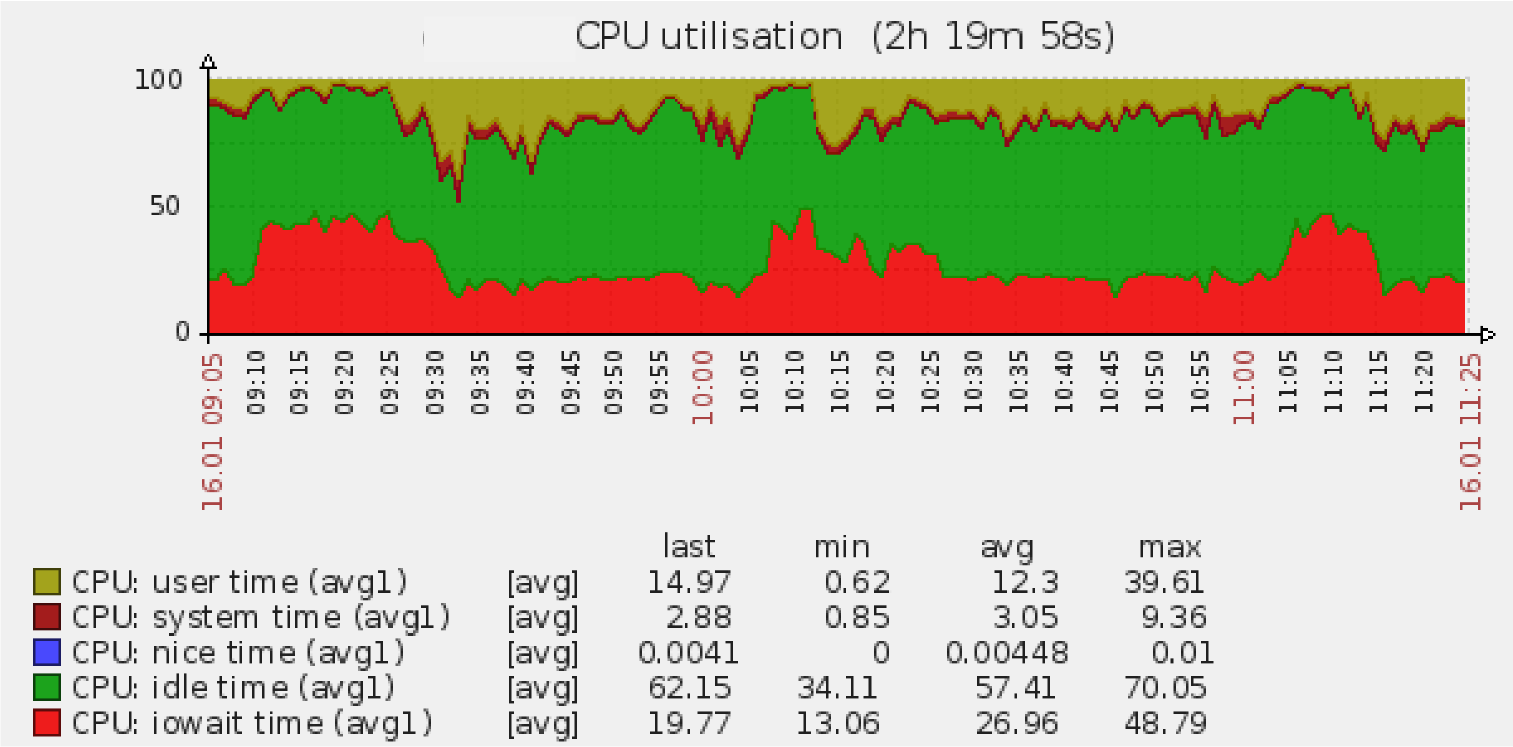

		last	min	avg	max
🟨 CPU: user time (avg1)	[avg]	14.97	0.62	12.3	39.61
🟫 CPU: system time (avg1)	[avg]	2.88	0.85	3.05	9.36
🟦 CPU: nice time (avg1)	[avg]	0.0041	0	0.00448	0.01
🟩 CPU: idle time (avg1)	[avg]	62.15	34.11	57.41	70.05
🟥 CPU: iowait time (avg1)	[avg]	19.77	13.06	26.96	48.79

The “Search monster” case

- **Added CPU graph from a search server**
- I/O wait peeks correlate with application performance problems
- Application colleague saw the correlation immediately:

- **Explanation / Root cause:**
- Blue lines is updates of search registers in search subsystem
- Red lines is commits of these updates!
- … and these commits are HEAVY … at disk I/O
- Updates are running in periods (sometimes with weeks in between)
- New search engine did also make more load on servers

- **Case closed:**
- **Search server hardware was found guilty!**
- … problem solved after hardware upgrade
- … AND heavy Disk system upgrade (read: faaaast disks)

The “Monday monster” case

- Victims: Users of application.

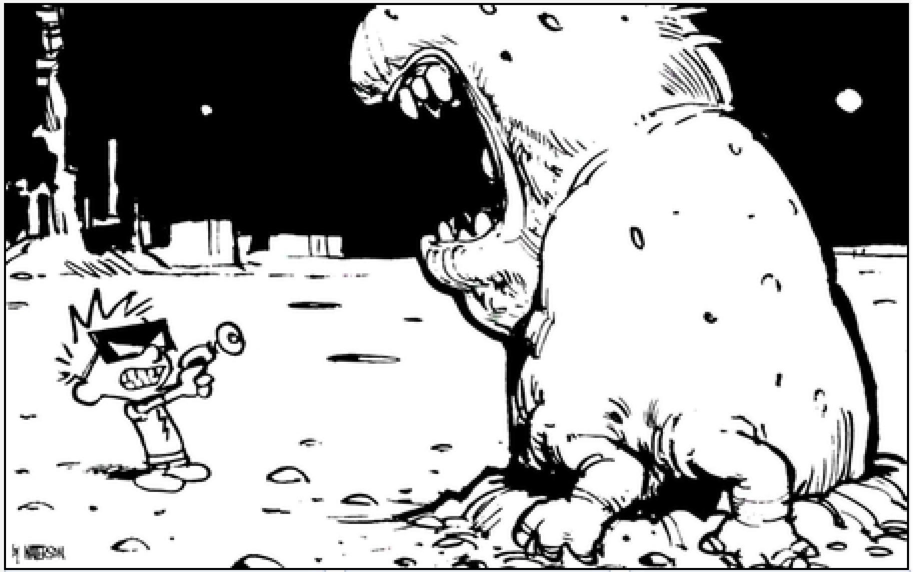

© Universal Press Syndicate

- **Details:**
- A performance problem. Even that it is called the Monday problem, it do not have special connections with Mondays. It is seen at all days and all times of days. Even at weekends. It is not seen that often … 1 – 3 times a week … but can block activity for some time

- Suspects: Backend systems, database system, heavy batch jobs?
- Evidence: Historical data, trading graphs, monitoring alarms

The "Monday monster" case

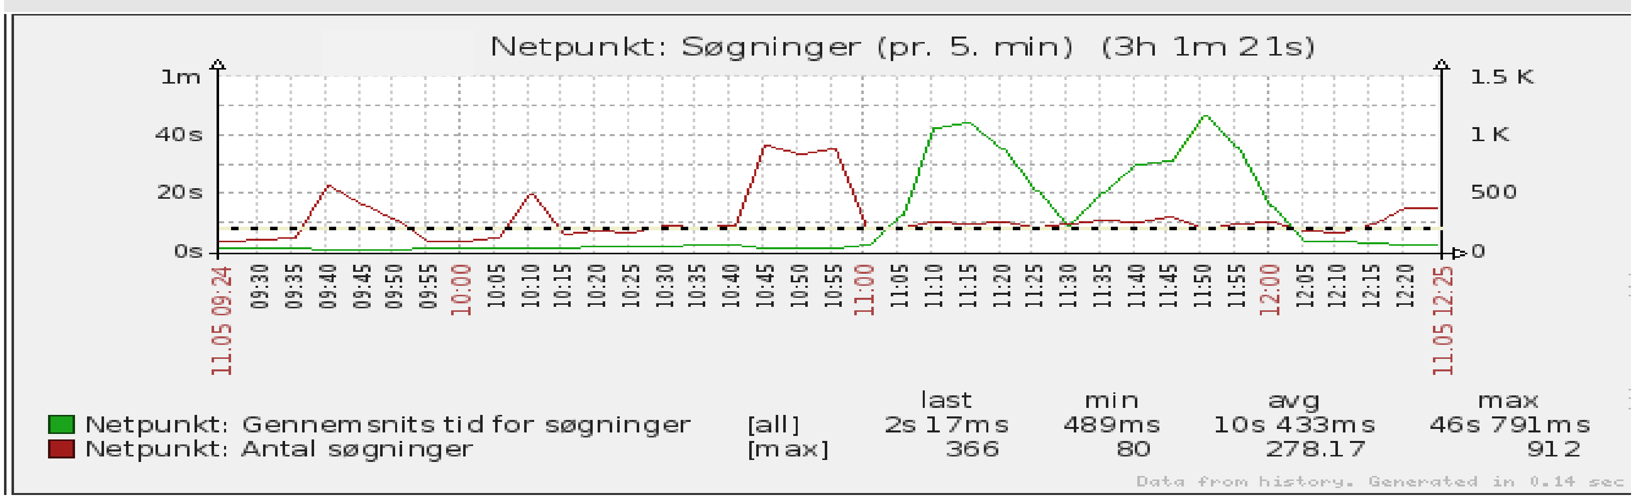

		last	min	avg	max
🟩 Netpunkt: Gennemsnits tid for søgninger	[all]	2s 17ms	489ms	10s 433ms	46s 791ms
🟥 Netpunkt: Antal søgninger	[max]	366	80	278.17	912

Data from history. Generated in 0.14 sec

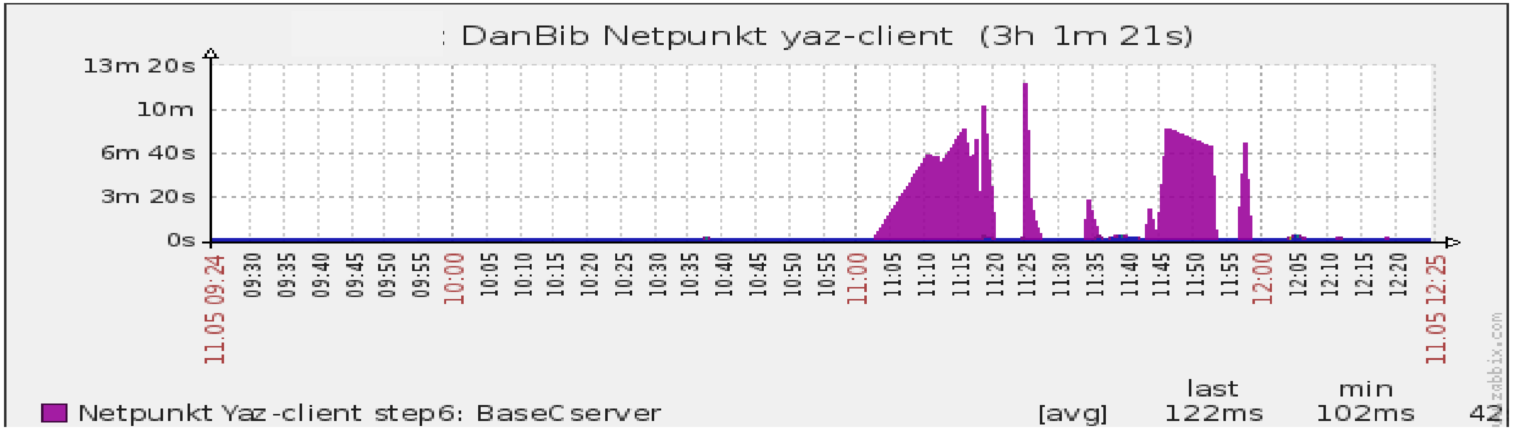

		last	min	
🟪 Netpunkt Yaz-client step6: BaseCserver	[avg]	122ms	102ms	42

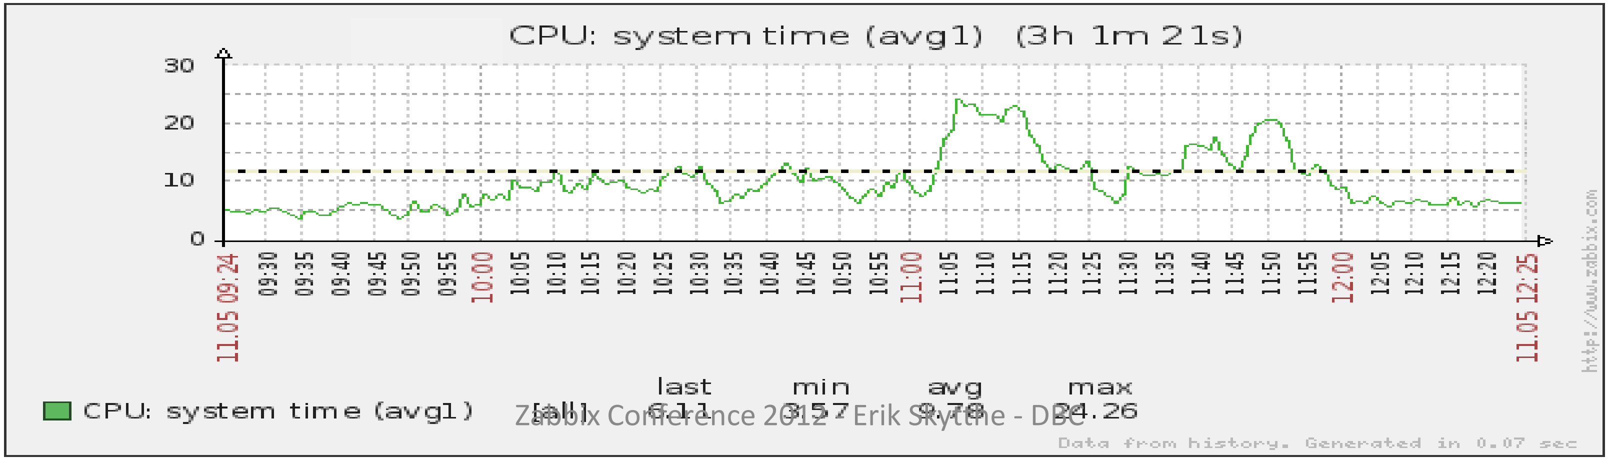

		last	min	avg	max
🟩 CPU: system time (avg1)	[all]	6.11	3.57	9.78	24.26

Data from history. Generated in 0.07 sec

The “Monday monster” case

- **BaseCserver (backend subsystem)**
- Did early see correlation with baseCserver tests
- However Logs and statistics for baseCserver did not indicate problems
- Did not have logs and statistics for systems in between frontend and baseCserver

- **Analyses of frontend log files:**
- Only transaction times from 0 – 20 sec and > 1 min (timeout) ??
- 25 % of transactions was timeouts ?

The “Monday monster” case

- **Database server**
- Did later find correlation with systemtime on DB server
- But response time statistics etc. from DB server was fine
- Even specialized DB monitor tool did not show problems

- Further search for correlations did not succeed
- Eyes was turned onto batch jobs and other systems doing heavy database searches etc.

CALVIN! YOU'RE NOT PAYING ATTENTION!
..WE JOIN SPACEMAN SPIFF ON THE DISTANT PLANET ZORG...
GRONK! ARGH!
ZOUNDS!

© Universal Press Syndicate

The “Monday monster” case

- **Other systems – the long search**
- Did get list of “heavy” systems
- Did go for the logs again
- Different log formats, but a matter of scripting

- Did not want to wait for things (problems) to happen this time
- So pulled out historical data of system logs and transferred it into Zabbix

TRAPPED BY A HIDEOUS GRAKNIL, SPIFF DRAWS HIS TRUSTY ATOMIC NAPALM NEUTRALIZER!
CHEW ELECTRIC DEATH, SNARLING CUR!
BUT THE WEAPON IS USELESS! SPIFF IS DOOMED!!

© Universal Press Syndicate

The “Monday monster” case

- **Other systems – the long search … continued**
- Struggle with how to visualize data from log files
  - Time span
  - Records / transactions processed
- Miss bar graphs
- Tried to “simulate” bar graphs

- **Did not find any correlation with heavy systems**
- … at least not the systems I did investigate

- Did however find jobs that should not run in day time …

The ”Monday monster” case

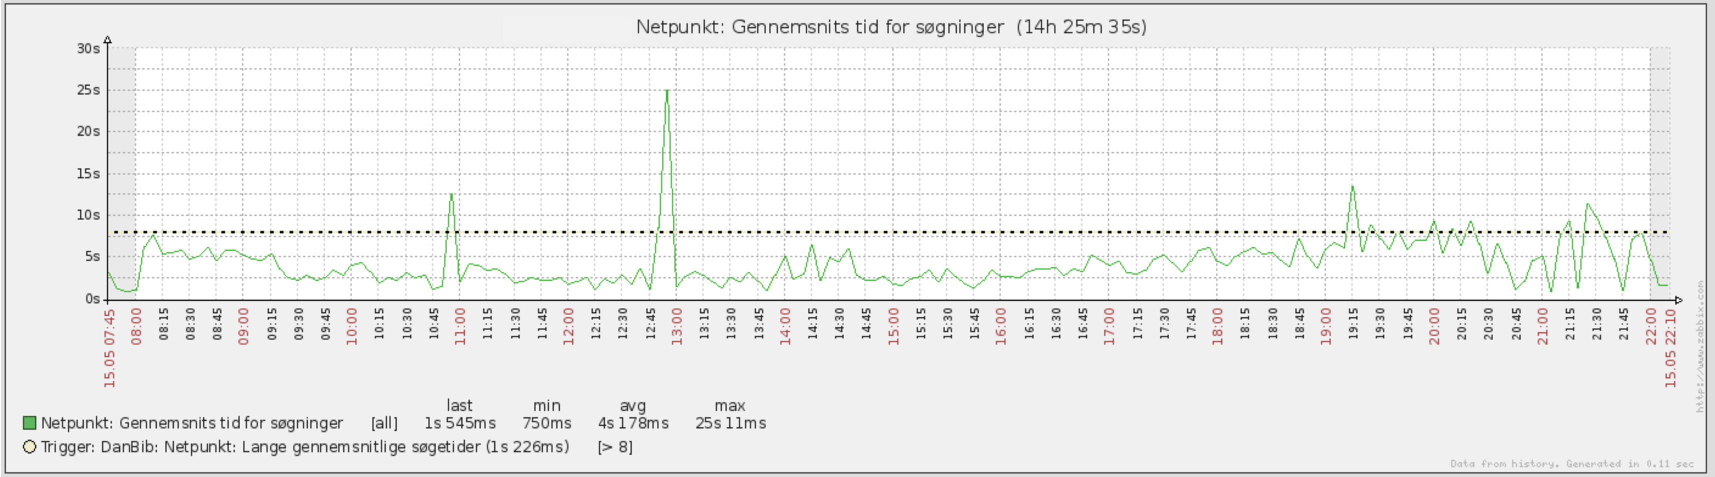

		last	min	avg	max
🟩 Netpunkt: Gennemsnits tid for søgninger	[all]	1s 545ms	750ms	4s 178ms	25s 11ms
○ Trigger: DanBib: Netpunkt: Lange gennemsnitlige søgetider (1s 226ms)	[> 8]				

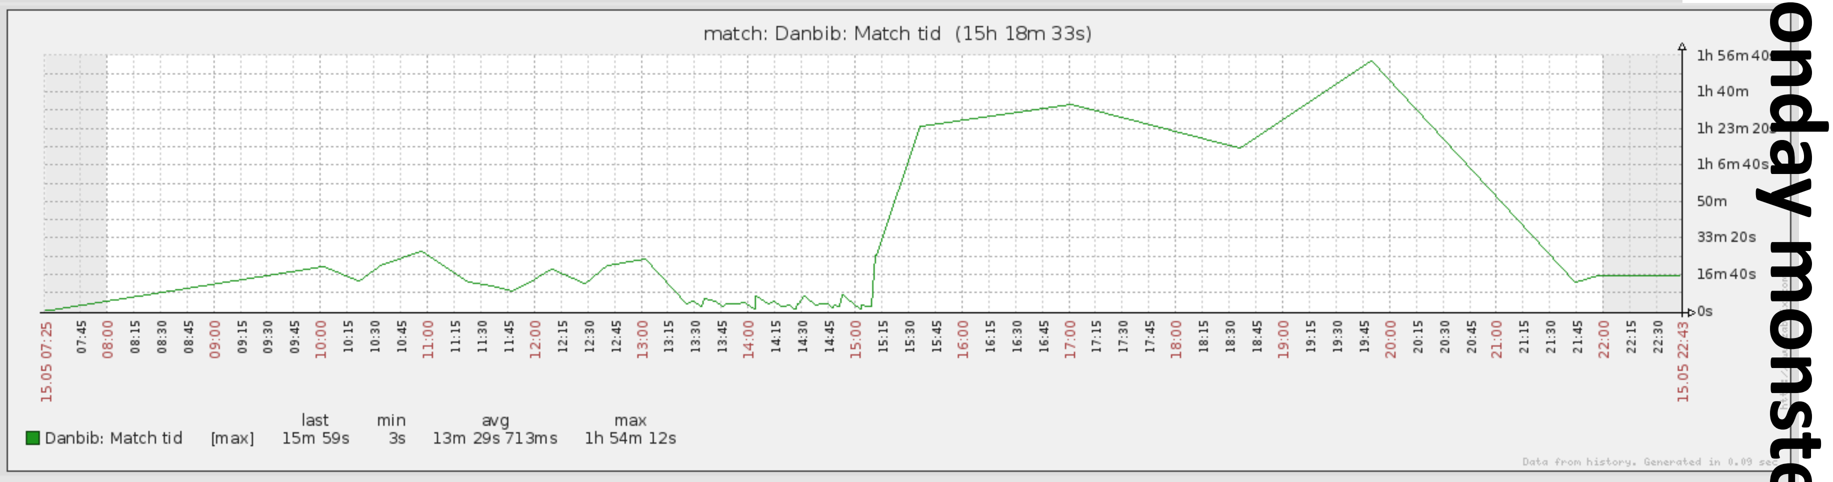

		last	min	avg	max
🟩 Danbib: Match tid	[max]	15m 59s	3s	13m 29s 713ms	1h 54m 12s

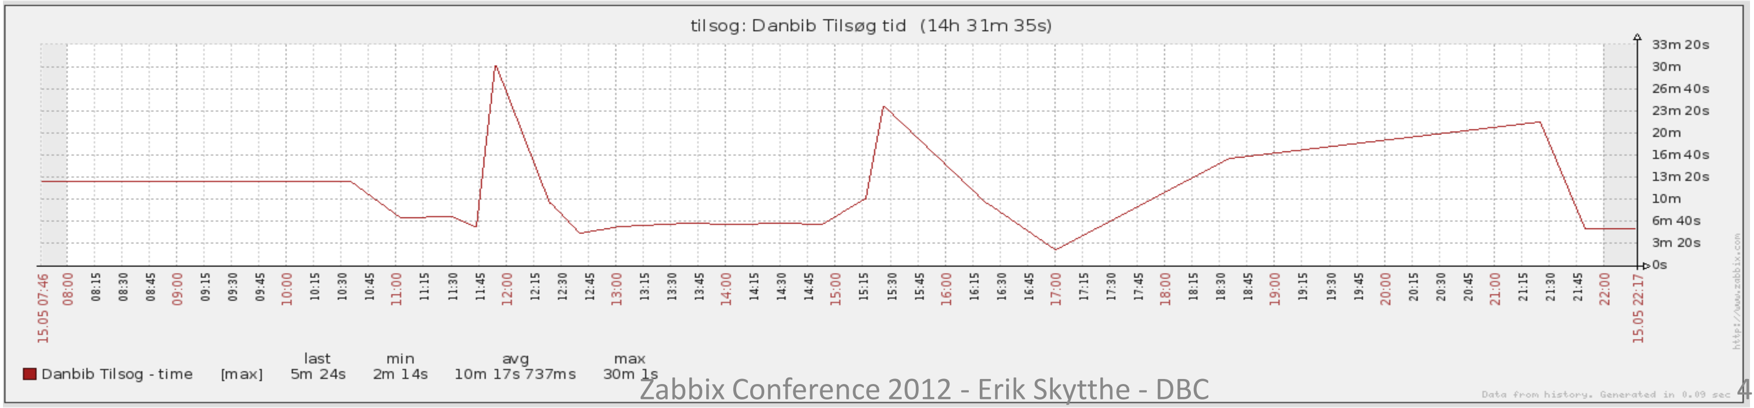

		last	min	avg	max
🟥 Danbib Tilsog - time	[max]	5m 24s	2m 14s	10m 17s 737ms	30m 1s

The “Monday monster” case

- **Zabbix_sender**
- **$ awk -f adhl.awk adhl-log > outputfile**
- **Outputfile: (*simulation of bar graphs)***
- adhl danbib.adhlimport.time 1343180920 0 *(edge start)*
- adhl danbib.adhlimport.time 1343180930 3783
- adhl danbib.adhlimport.time 1343184704 3783
- adhl danbib.adhlimport.time 1343184714 0 *(edge end)*
- adhl danbib.adhlimport.time 1343184725 0
- adhl danbib.adhlimport.time 1343184735 877
- adhl danbib.adhlimport.time 1343185603 877
- adhl danbib.adhlimport.time 1343185613 0
-
- **$ Zabbix_sender -z Zabbix -i outputfile -T**
- Info from server: "Processed 250 Failed 0 Total 250 Seconds spent 0.001821"
- …
- Info from server: "Processed 250 Failed 0 Total 250 Seconds spent 0.001739"
- Info from server: "Processed 80 Failed 0 Total 80 Seconds spent 0.000577"
- sent: 1580; skipped: 0; total: 1580

The “Monday monster” case

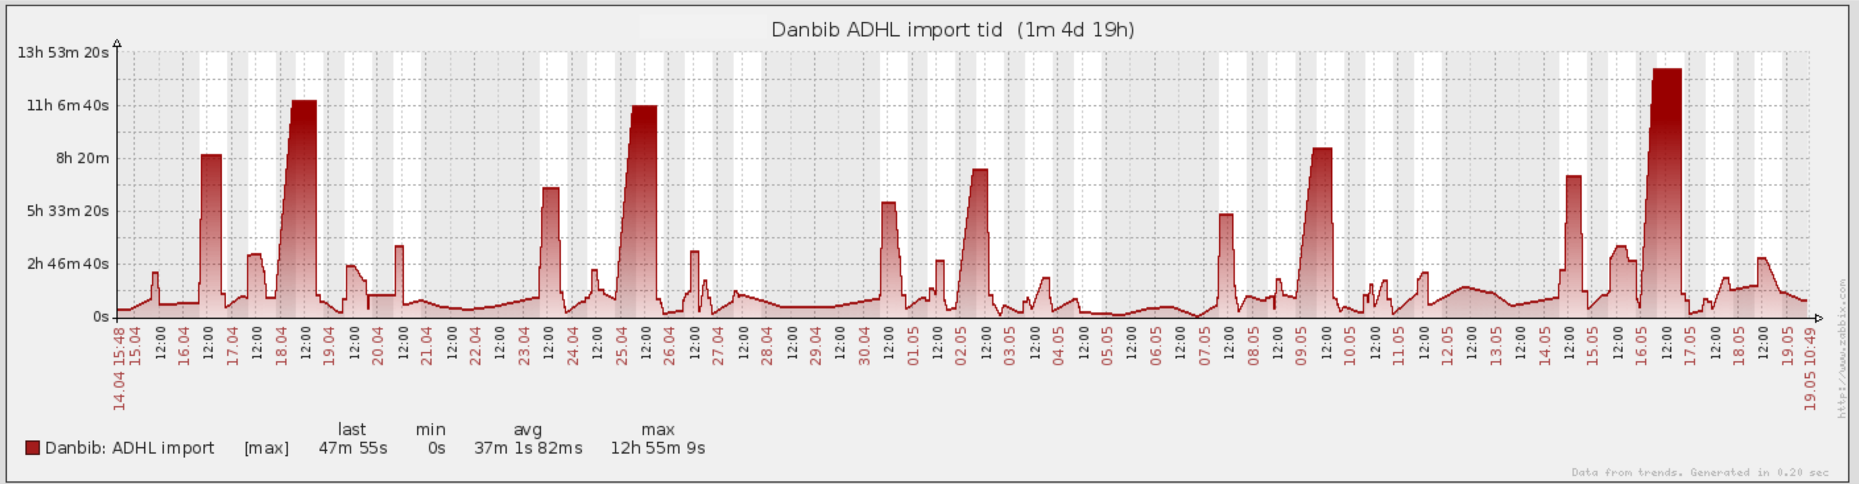

		last	min	avg	max
🟥 Danbib: ADHL import	[max]	47m 55s	0s	37m 1s 82ms	12h 55m 9s

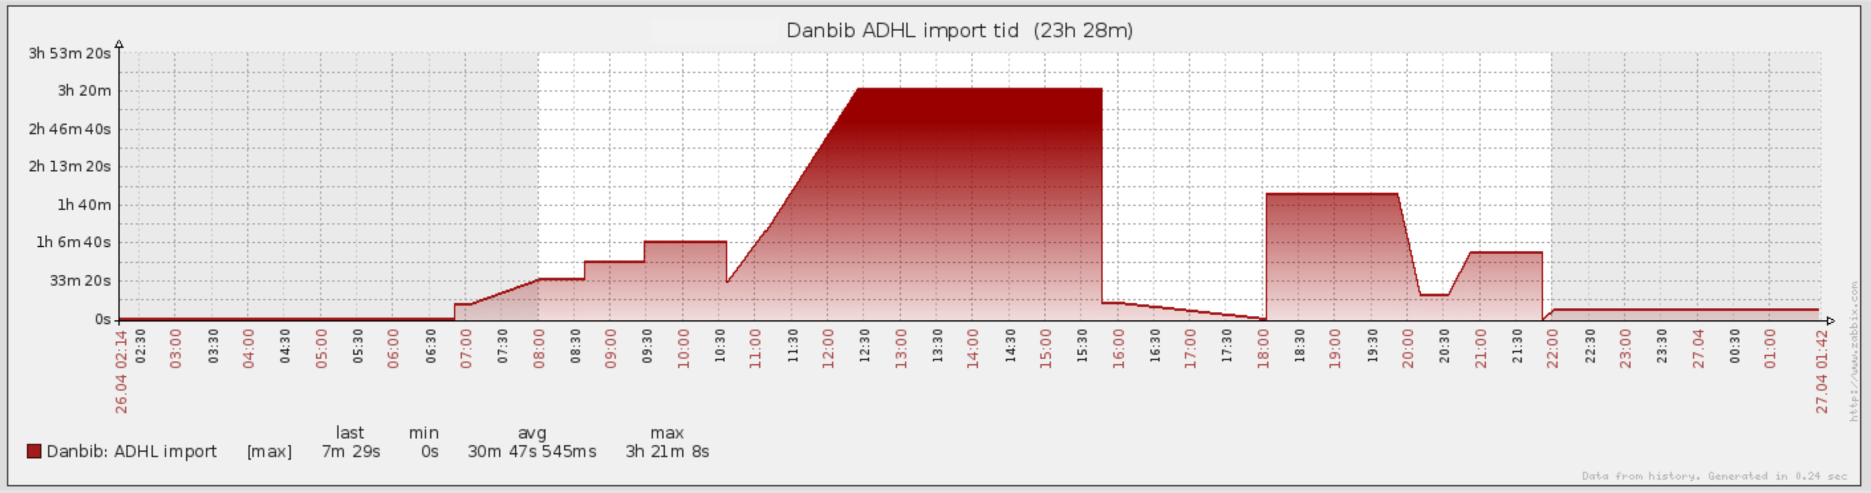

		last	min	avg	max
🟥 Danbib: ADHL import	[max]	7m 29s	0s	30m 47s 545ms	3h 21m 8s

The “Monday monster” case

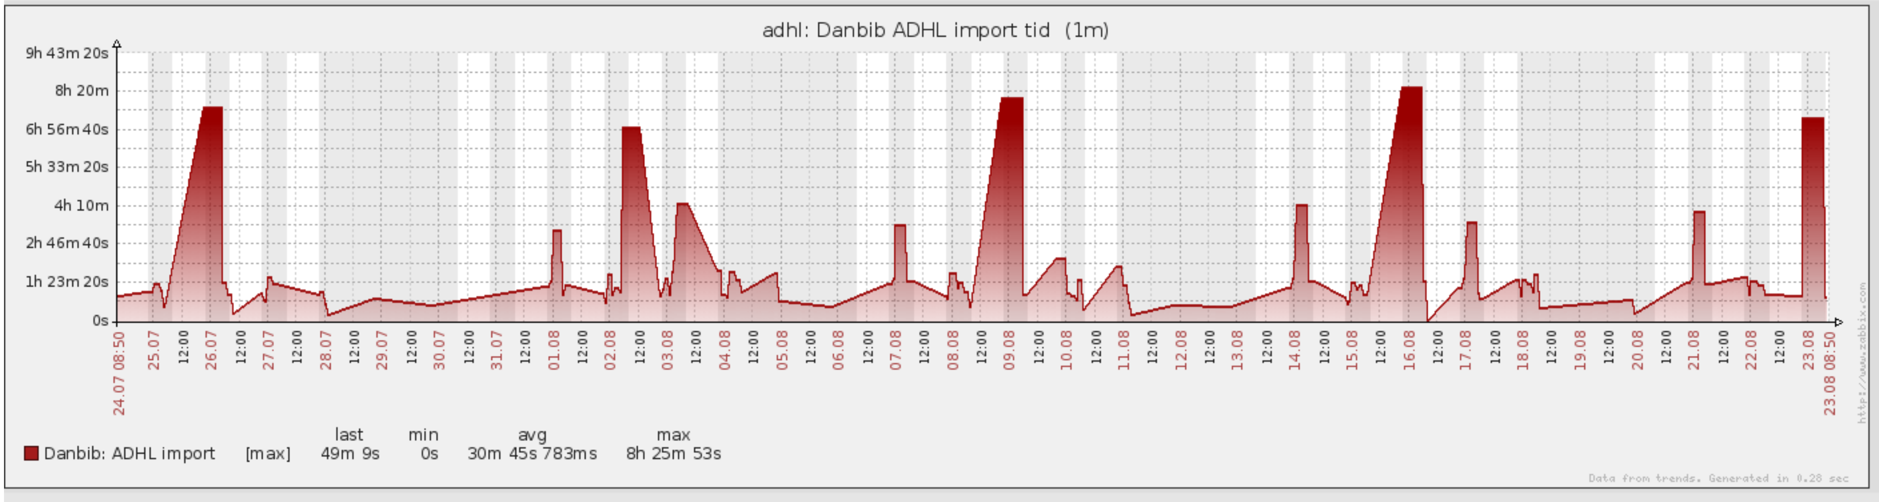

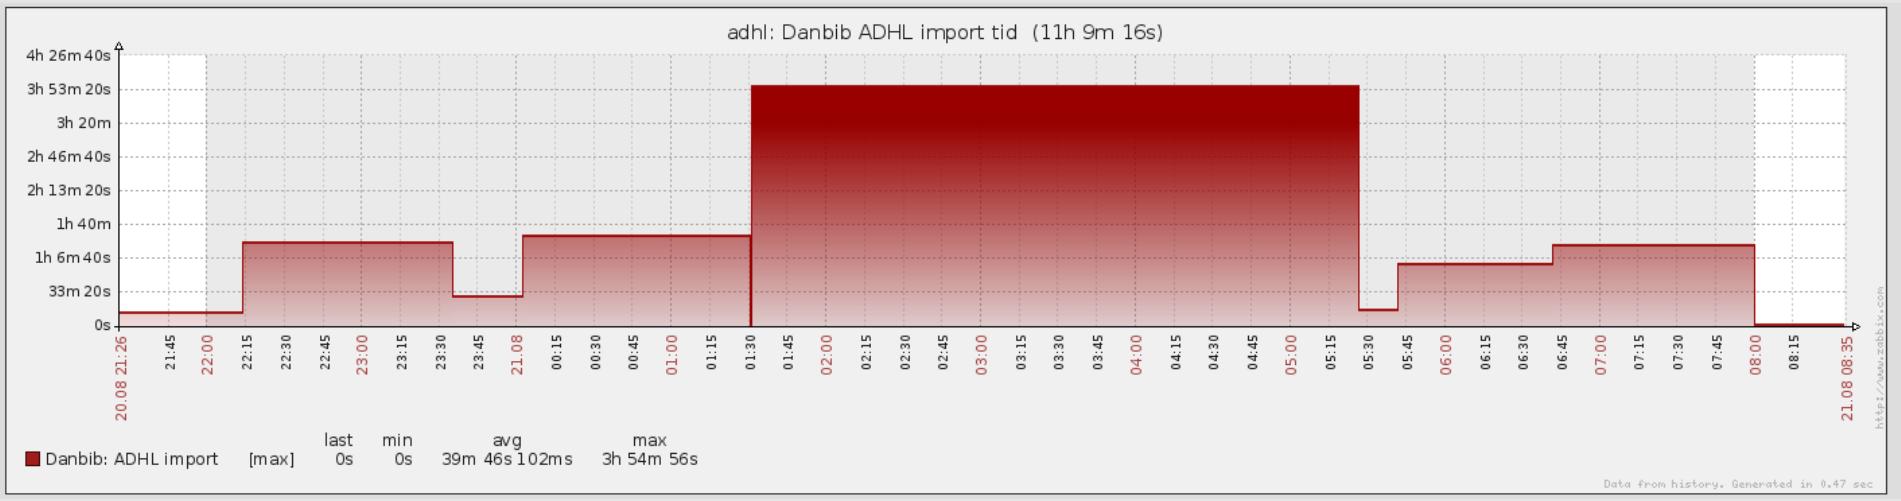

The “Monday monster” case

Frontend

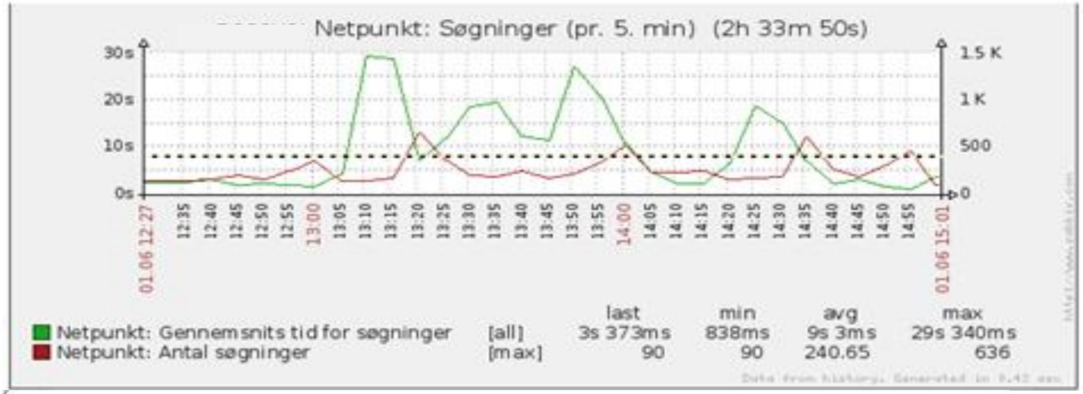

DB server

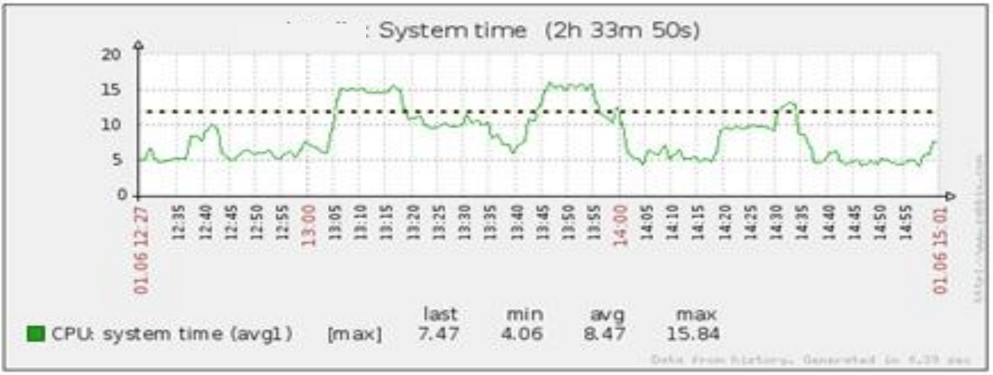

BaseCserver host

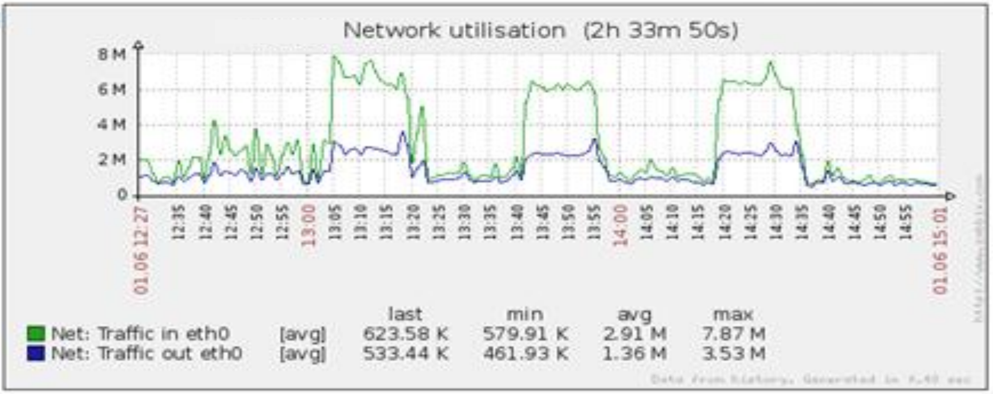

The “Monday monster” case

- **Finally a breakthrough !**
- One morning … looking at other graphs … I went to check the baseCserver hosts OS statistic screen ...
- Peaks with high Network I/O caught my attention
- Some zoom in and out and …
- … **correlation with DB server system time found !!**

- From time period provided by the correlation, my database admin colleague found this in the DB log:
- **<Timing>Total time for request : 399** (sec)
- Development colleague found the request in a baseCserver log …
- … and could explain the whole story …

The “Monday monster” case

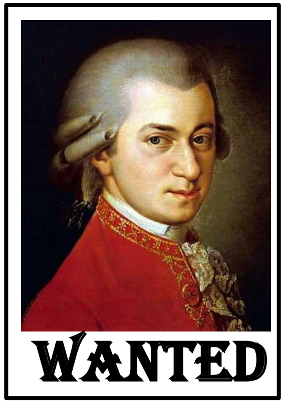

- **The guilty was: Wolfgang Amadeus Mozart ☺**
- **Explanation / Root cause:**
- A subtask in a search, is to do a dynamic matching of all references
- The problem appeared when someone did a search on media with very many references
- **BaseCserver (was the real guilty)**
- The baseCserver does this matching
- The baseCserver has to match on a big amount of data (records) …
- … that’s why we see network peaks on baseCserver hosts
- … and higher systemtime on DB server
- 4 instances of the baseCserver were running
- 1 will be blocked – explaining 25 % transactions timeouts!
- **Solution**
- The “simple” solution was to increase the number of instances of the baseCserver !!!

The Start …

© Universal Press Syndicate

Calvin and Hobbes

- Zabbix 1.8 Tutorials:
- http://www.zabbix.com/forum/showthread.php?t=22211