# Understanding Snort Performance

**An Evidence-Based Approach**

1

- Mike Lococo
- Senior Network Security Analyst
- New York University, largest private university in US
- <40k for-credit students +12k non-credit
- >15k employees
- Decentralized IT reporting structure
- 5 security staff

## Survey

Familiar with **Perfmonitor** Preprocessor?

Familiar with **--enable-perfprofiling**?

Understand their sensor **bottlenecks**?

Familiar with **Zabbix** or another monitoring system?



Snort + graphs is better than snort without graphs, that's pretty much all we're going to cover. You'll still make all the same dumb mistakes, they'll just be easier to find.

We'll cover snort instrumentation, combine it with some basic OS and other instrumentation, and then look at graphs. It's not rocket science, but it makes it easier to develop a situational awareness of what is “normal” for your environment.

Click to add title

<Capture Frameworks>

## Libpcap

**Free**, works on **commodity** hardware

Scales to **200-300** Mbits/sec

**Single**-Queue/CPU

This is the default if you use libpcap prior to ~1.0



Tweaks to sysctl.conf may help:

net.core.rmem_max = 33554432
net.core.netdev_max_backlog = 10000
net.core.rmem_default = 33554432

## Mmapped Libpcap

**Free**, works on **commodity** hardware

Scales to **100** Mbits/sec

**Single**-Queue/CPU

Default for libpcap > ~1.0

Small hardcoded buffer-size limits performance, unlike the abandoned Phil Woods patches.



Note that performance may be **worse** with later libpcaps.

## AFPACKET

**Free**, works on **commodity** hardware

Scales to a **200-300** Mbits/sec (if you increase your buffer size)

**Single**-Queue/CPU

Also called af_packet



Snort Manual:
1.5.1 for syntax to configure from snort.conf
1.5.3 for syntax to configure from command-line

Sourcefire Howto:
http://vrt-blog.snort.org/2010/08/snort-29-essentials-daq.html

Kernel Interface Manpage for af_packet:
http://manpages.ubuntu.com/manpages/jaunty/man7/packet.7.html

## PF_RING+TNAPI

**Free-$250**, works on **commodity** hardware

Scales to a **>1G** Gig

**Multi**-Queue/CPU

Highest performance with certain intel cards via proprietary drivers



Explanation of TNAPI at ntop.org:
http://www.ntop.org/TNAPI.html

Buy the high-performance drivers:
http://www.ntop.org/shop/cart.php

Might scale close to 10gig, but no independent reports that I'm aware of confirm the ntop.org numbers.

# Intel X520 and Myricom 10G

**$1k-2k** dedicated **capture cards**

Scales to **10Gig**

**Multi**-Queue/CPU

Most reports are preliminary, these are relatively new cards. Linux drivers are supposed to be in the kernel.

# Endace/Napatech

**$10k-20k** dedicated **capture cards**

Scales to **10Gig**

**Multi**-Queue/CPU

Well-established, well-tested. Expensive and a bit of a pain to manage.

## Suricata

**Free**, runs on **commodity** hardware

Scales according to your CPU-count, but generally **slower than snort**.

Single-Queue/**Multi**-CPU

Evolving quickly



Holisticinfosec Performance Test from August 2010: http://holisticinfosec.org/toolsmith/docs/august2010.html

Summary is that Suricata is 4x less CPU-efficient than snort, but if you compare a large multi-cpu Suricata instance to a single-cpu snort instance it can be faster given sufficient hardware.

## Your Mileage **will** Vary


http://www.flickr.com/photos/usnationalarchives/4272498928/



These are all anecdotes.

There are many site-specific factors, and you'll have to test locally to determine what works for you.

The remainder of this presentation is aimed at giving you the data you need to perform that local assessment.

# Click to add title

</Capture Frameworks>

Click to add title

<Snort Instrumentation>

## Perfmon Preprocessor

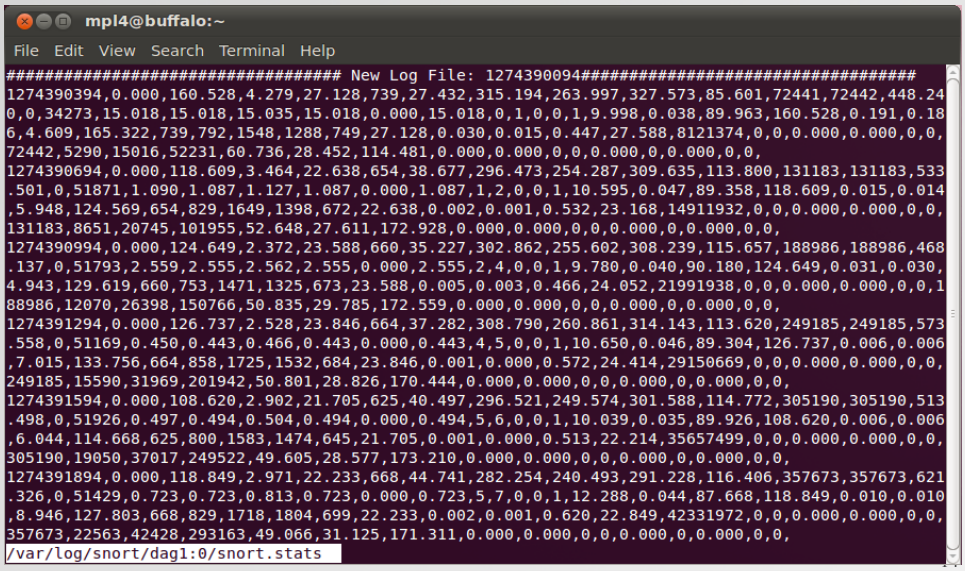


CSV output.

Simple to configure.

Little/no cpu overhead to collect.

Rotate with logrotate, or just let it grow to a few hundred meg over the lifetime of your sensor.

Not much fun to interpret unless you have some way to process it beyond reading the CSV file.

# Configuring Perfmon

In Snort.conf:

```
preprocessor perfmonitor: time 300
file /var/log/snort/bogus.stats
pktcnt 10000
```

## Perfmon Data Fields

The fields are “documented” in section 2.2.5 of the manual... poorly.

Drop rate, Mbits/sec, Packets/sec, Alerts/sec, Packets received, Packets dropped, many more



Be aware that the drop-rate is averaged over the lifetime of the snort-process, not over the data-collection period. A 5% drop rate may be 5% over 24 hours, or 60% for 2 hours.

Other fields are generally averaged over the data-collection period, but one needs to be careful.

# Perfprofiling Preprocessors

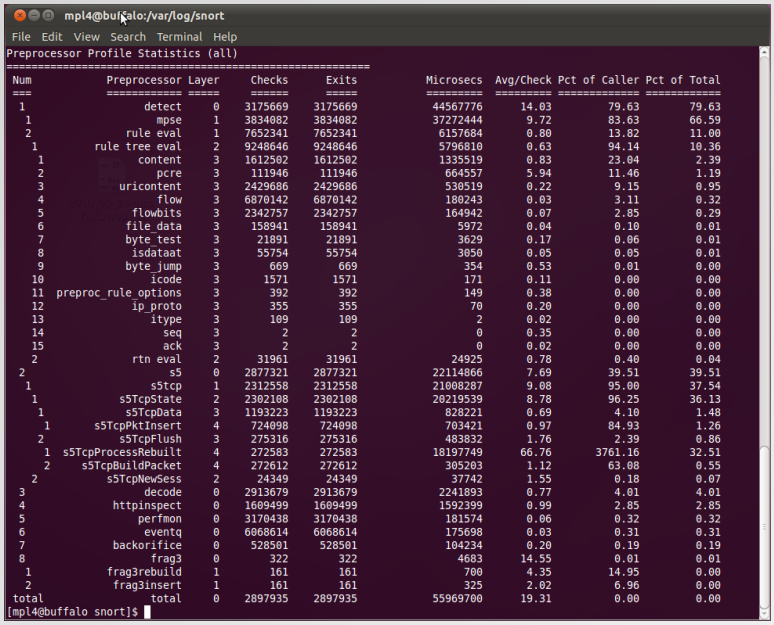




“Num” column shows component nesting within the snort process. For example, “mpse” and “rule eval” are components of the “detect” module.

“Percent of total” column gives a rough idea of how much cpu time that component takes up. It's not precise (note the top-level components do not add up to 100) but relative comparisons are useful.

In this sample, the pattern-matcher and the stream-preprocessor take up the vast-majority of time. This is relatively healthy for a snort instance with lots (~7k) rules enabled.

This snort can't be effectively tuned by stripping individual rules, we'd need to cut down the number to reduce the pattern matcher (mpse) CPU usage.

# Perfprofiling Rules

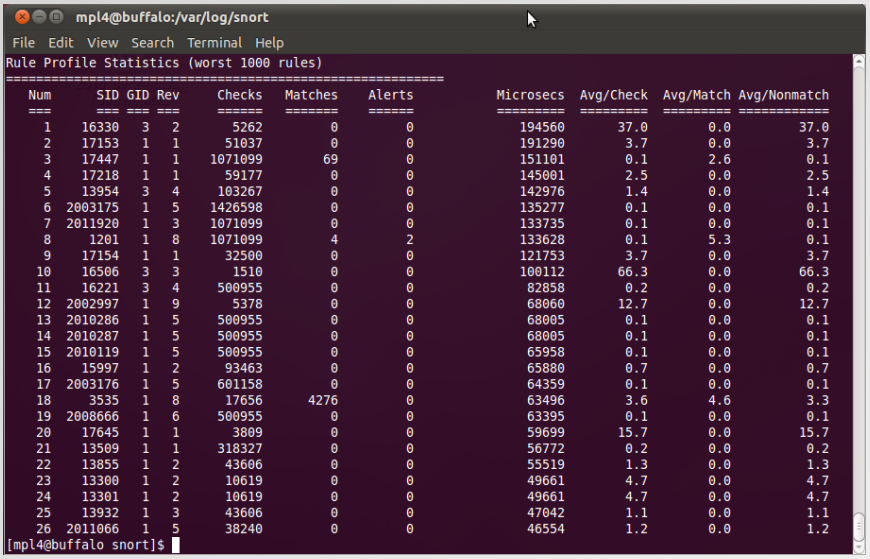


| Num | SID | GID | Rev | Checks | Matches | Alerts | Microsecs | Avg/Check | Avg/Match | Avg/Nonmatch |
|---|---|---|---|---|---|---|---|---|---|---|
| 1 | 16330 | 3 | 2 | 5262 | 0 | 0 | 194560 | 37.0 | 0.0 | 37.0 |
| 2 | 17153 | 1 | 1 | 51037 | 0 | 0 | 191290 | 3.7 | 0.0 | 3.7 |
| 3 | 17447 | 1 | 1 | 1071099 | 69 | 0 | 151101 | 0.1 | 2.6 | 0.1 |
| 4 | 17218 | 1 | 1 | 59177 | 0 | 0 | 145001 | 2.5 | 0.0 | 2.5 |
| 5 | 13954 | 3 | 4 | 103267 | 0 | 0 | 142976 | 1.4 | 0.0 | 1.4 |
| 6 | 2003175 | 1 | 5 | 1426598 | 0 | 0 | 135277 | 0.1 | 0.0 | 0.1 |
| 7 | 2011920 | 1 | 3 | 1071099 | 0 | 0 | 133735 | 0.1 | 0.0 | 0.1 |
| 8 | 1201 | 1 | 8 | 1071099 | 4 | 2 | 133628 | 0.1 | 5.3 | 0.1 |
| 9 | 17154 | 1 | 1 | 32500 | 0 | 0 | 121753 | 3.7 | 0.0 | 3.7 |
| 10 | 16506 | 3 | 3 | 1510 | 0 | 0 | 100112 | 66.3 | 0.0 | 66.3 |
| 11 | 16221 | 3 | 4 | 500955 | 0 | 0 | 82858 | 0.2 | 0.0 | 0.2 |
| 12 | 2002997 | 1 | 9 | 5378 | 0 | 0 | 68060 | 12.7 | 0.0 | 12.7 |
| 13 | 2010286 | 1 | 5 | 500955 | 0 | 0 | 68005 | 0.1 | 0.0 | 0.1 |
| 14 | 2010287 | 1 | 5 | 500955 | 0 | 0 | 68005 | 0.1 | 0.0 | 0.1 |
| 15 | 2010119 | 1 | 5 | 500955 | 0 | 0 | 65958 | 0.1 | 0.0 | 0.1 |
| 16 | 15997 | 1 | 2 | 93463 | 0 | 0 | 65880 | 0.7 | 0.0 | 0.7 |
| 17 | 2003176 | 1 | 5 | 601158 | 0 | 0 | 64359 | 0.1 | 0.0 | 0.1 |
| 18 | 3535 | 1 | 8 | 17656 | 4276 | 0 | 63496 | 3.6 | 4.6 | 3.3 |
| 19 | 2008666 | 1 | 6 | 500955 | 0 | 0 | 63395 | 0.1 | 0.0 | 0.1 |
| 20 | 17645 | 1 | 1 | 3809 | 0 | 0 | 59699 | 15.7 | 0.0 | 15.7 |
| 21 | 13509 | 1 | 1 | 318327 | 0 | 0 | 56772 | 0.2 | 0.0 | 0.2 |
| 22 | 13855 | 1 | 2 | 43606 | 0 | 0 | 55519 | 1.3 | 0.0 | 1.3 |
| 23 | 13300 | 1 | 2 | 10619 | 0 | 0 | 49661 | 4.7 | 0.0 | 4.7 |
| 24 | 13301 | 1 | 2 | 10619 | 0 | 0 | 49661 | 4.7 | 0.0 | 4.7 |
| 25 | 13932 | 1 | 3 | 43606 | 0 | 0 | 47042 | 1.1 | 0.0 | 1.1 |
| 26 | 2011066 | 1 | 5 | 38240 | 0 | 0 | 46554 | 1.2 | 0.0 | 1.2 |



Shows most “expensive” rules in terms of cpu consumption, sorted by microsecs.

A high “check” count indicates a short/common content match causing the rule to evaluate often.

A high “Avg/check” value usually indicates an expensive regex.

Microsecs is checks * avg/check and represents the relative cpu-cost of running the rule.

Rules that “match” a lot but don't “alert” are probably setting flowbits, check to see if you care about that flowbit and if not disable all the rules that set/check it.

## Configuring Perfprofiling

Add `--enable-perfprofiling` to `configure` prior to compiling.

In snort.conf:

```
config profile_rules: print 1000, sort
total_ticks, filename /var/log/snort/snort-
perftest_rule.log

config profile_preprocs: print all, sort
total_ticks, filename /var/log/snort/snort-
perftest_preproc.log
```



Documented in the perfprofiling readme package with Snort's source:

http://cvs.snort.org/viewcvs.cgi/snort/doc/README.PerfProfiling?rev=HEAD&content-type=text/vnd.viewcvs-markup

There is allegedly a performance overhead for doing this, but it's clearly not all that high. Perhaps 10%-20% at the most, I suspect it's negligible.

Note that this is not trivially collectible or trendable, but it's a good exercise to go through once a year or so during a major upgrade.


You need data from many sources to evaluate snort performance. Use a tool that can accept data from many sources.

Zabbix has a relatively gentle learning curve and is relatively featureful. Especially dynamically generated graphs are great for exploring data.

I've heard very good things about Zenoss, but haven't used it.

Nagios is well established, but has more legacy baggage and a higher-learning curve. Plus there are no graphing capabilities built in.  Go for it if you have expertise, but I wouldn't stand up a new instance today.

## Instrumenting Beyond Snort

**Tap/Span**: SNMP

**Packet Transport**: SNMP or NOC

**Capture**: Varies

**Snort Analysis**: PerfMonitor

**Alerts to Disk**: OS Stats

**Database**: OS or DB stats

**Frontend**: OS stats or App Stats



Snort is an integrated system that depends on many components, you want a system that can instrument all (or many) of those components, not just snort.

# Click to add title

</Snort Instrumentation>

Click to add title

<Zabbix>

## Zabbix Concepts

**Item**: Data element to be collected

**Graph**: Visual trend for numeric data

**Screen**: A collection of graphs

**Trigger**: Nominal ranges for items



Items can be collected from an agent that supports a fairly wide variety of OS items natively. SNMP is also supported, as well as trivial checks like pings or http-requests from the server.

Graphs are generated dynamically at view time, which makes exploration a cinch.

Screens, unlike most other resources can't be templated, which can be frustrating.

Triggers support many conditions including string tests.

## Zabbix Concepts

**Action**: Performed on trigger condition(s)

**Host/Tempate**: Collections of (most of) the above with inheritance

**Reports**: Weak. Hardcoded red-light/green-light grids, plus uptime reports for triggers.



Notifications are supported via actions, as are IPMI commands (I haven't tested), and remote agent commands (I also haven't tested, but know they require sudo/nopasswd access for the zabbix-agent user). Actions are another resource that can't be templated for no reason that I can fathom.

## Collecting Perfmon Data

Create a userparameter in zabbix_agentd.conf:

```
UserParameter=tss.snort.perfmon[*],tail
-n1 '/var/log/snort/$1/snort.stats' |
awk -F ',' '{print $ $2}'
```



This userparameter accepts an interface-name ($1) and an awk column-number ($2). It pulled the most recent line from snort.stats for the appropriate interface, and prints out the column requested (note the extra $, as awk uses them to denote variables just like zabbix does.

# Create an Item

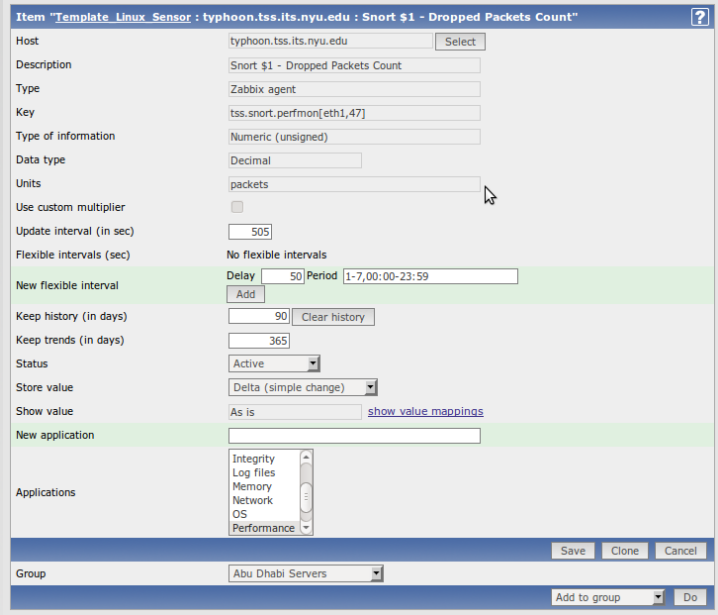




This is my dropped packets item.

Note “Store Value” → Delta setting. The snort dropped packet stats are poor in my opinion due to averaging over the lifetime of the process, and this stat is given as a simple counter. Zabbix converts the counter into a number of dropped packets per collection-period and gracefully handles overflow-rollovers.

# Create a Graph

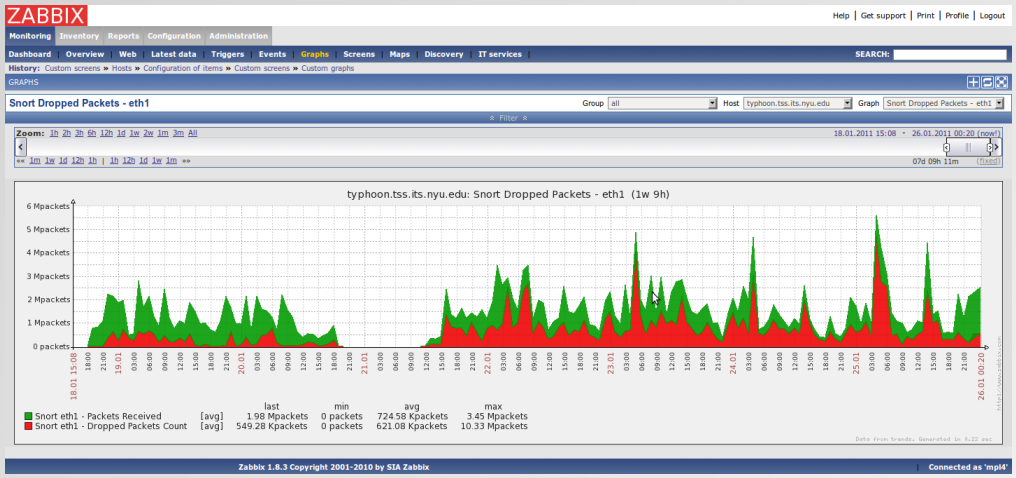




Line graphs or pie graphs.

Line graphs can be stacked, as above, to show cumulative values over time. This snort sensor is angry, red is dropped packets, green is processed packets.

## Create a Screen

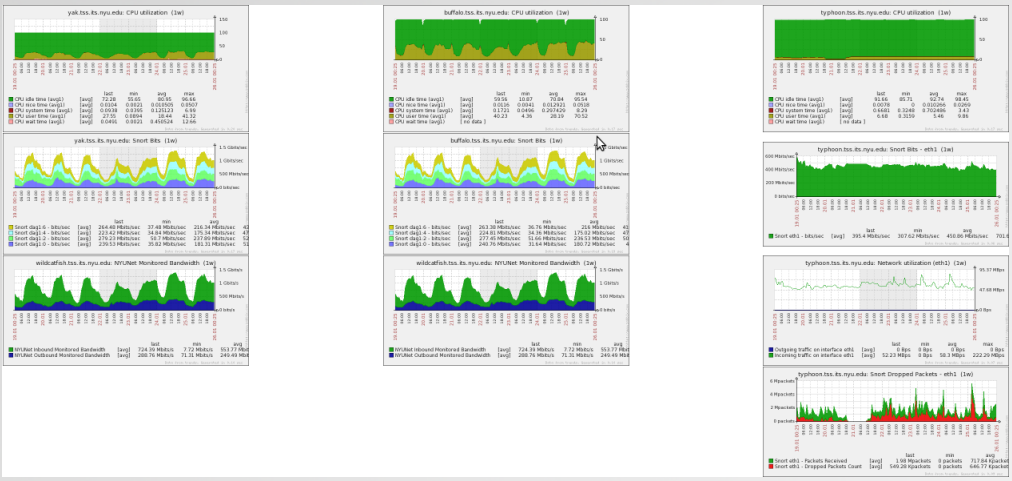

29

Three sensors, in columns.

In rows from top to bottom are: cpu-utilization (normalized against the number of cpu's present), bandwidth reported by snort (stacked where there are multiple snort instances), bandwidth reported by the gigamon (or kernel on the right), packet processed vs dropped packets where available.

This kind of at-a-glance analysis really changed the way we operate.

# Find Problems – DB Disk Outage

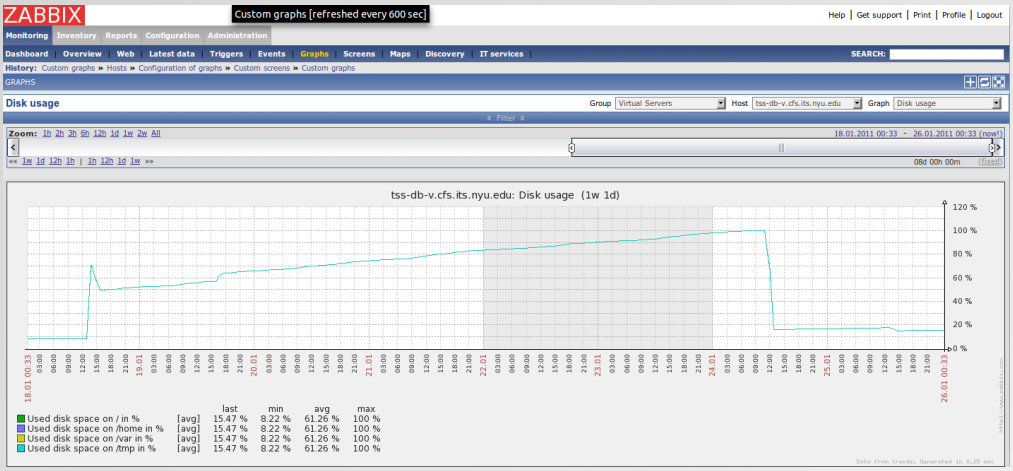




This was a real snort outage, and this is the real graph that I found the problem with.

This was a disk-space outage on our DB server (the scale goes to 120%, but my disks do not), but I actually had to investigate my whole failure chain to find it. Doing so in zabbix took only a few minutes.

Proper use of triggers and actions/alerts could have prevented this from happening at all.

# Intermittent Packet Loss

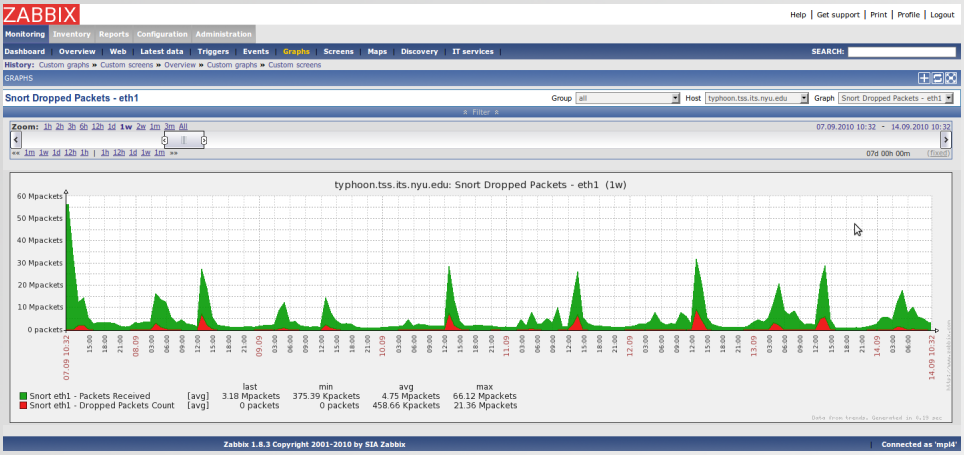




This sensor can achieve acceptable loss stats by filtering one port on one host that does large nightly transfers.

# Traffic Growth

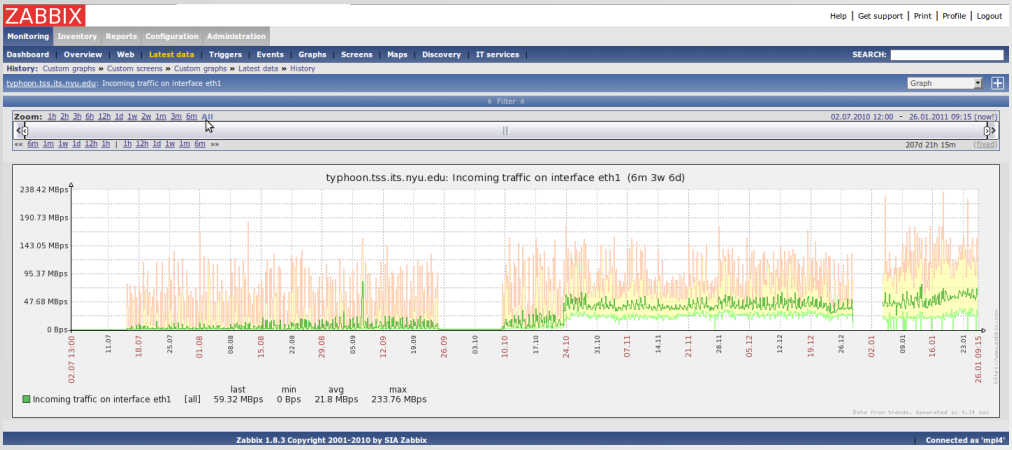




This sensor was spec'ed for 50-100mbits of traffic. The two previous dropped-packets graphs are for this host at different times during it's traffic growth to now 500mbit/sec.

Take note of the gaps for comparison on the next slide.

# Span Outage

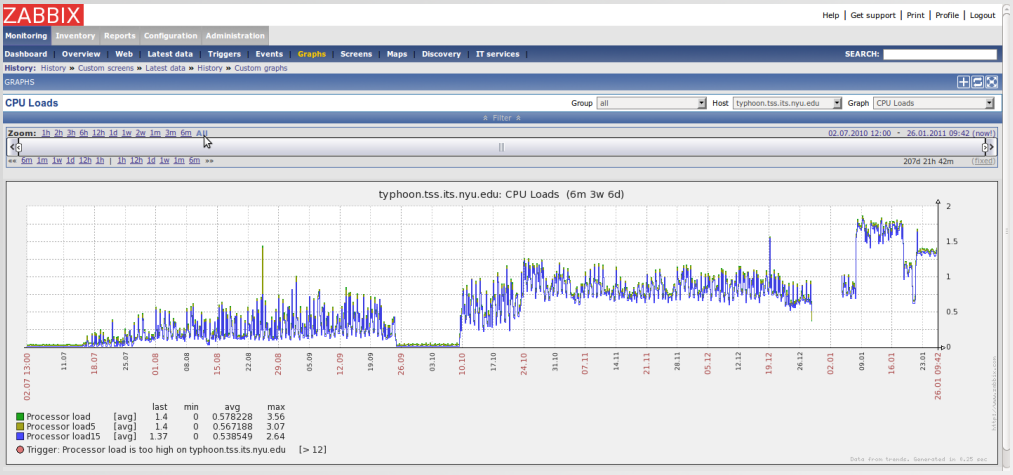




The second gap also has a gap in the system load graph, this was a planned outage.

The first gap has low system load stats. Our span was misconfigured by the NOC and we stopped getting packets. Zabbix told us before our incident monitoring folks realized that alerts had stopped.

## Mean Nice

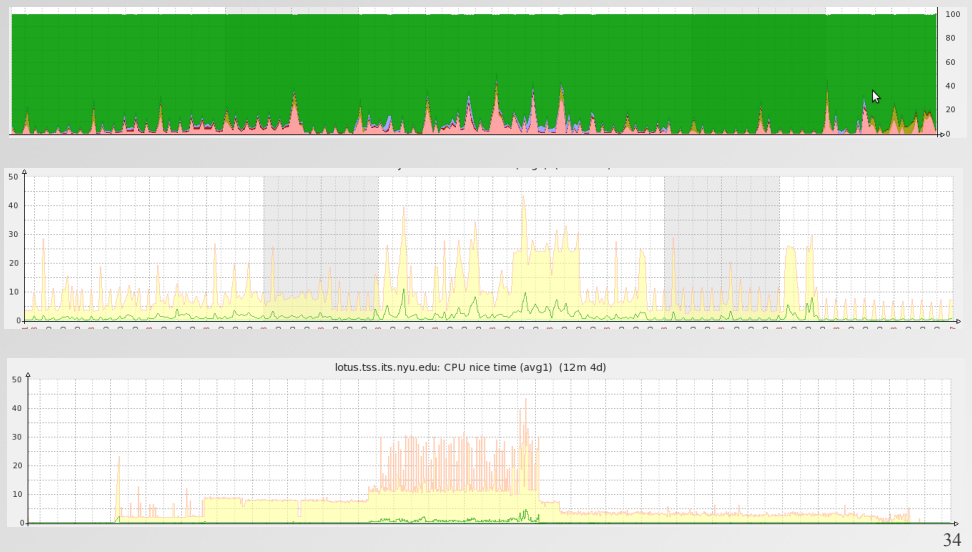


34

This was a bear to troubleshoot. Processes were getting reniced and we couldn't figure out how.

Turns out it was updated processes inheriting the nice-value of yum-updatesd when being restarted after upgrades. We finally made this determination by correlating log timestamps with initial nice-cpu spikes. We couldn't have done it based on logs alone, we needed the CPU trend data.

This is an outstanding redhat bug by the way, watch out for it.

# Click to add title

</Zabbix>

# Click to add title

<Demo />

# Conclusion

- **Instrument** your systems
- **Visualize** your data
- **Troubleshoot** faster
- **Understand** what's normal for you

# Click to add title

<Questions />