




Next Generation Monitoring: Moving Beyond Nagios



Intro - Us

- who are we?
- why do we care about this?



Intro - You

do you like your servers?



Nagios

- "It was here when I got here"
- initially released 1999



What this talk is about

- Nagios pains
- Sensu <3



Sensu <3

- client: runs checks
- server:
 - maintains state
 - executes handlers
 - answers questions about state



Nagios pain - Configuration

- large central configuration
- NRPE configs on server & client
- Perl



Nagios pain - Configuration

- Nagios config in separate git repo
- vast differences b/w dev, stages, & prod
- chicken-and-egg w/ code changes
- complexity grew w/ number of teams



Nagios Pain - Scaling

Weak scaling story

- each check adds load
- scheduled checks may not run


No HA

- no good failover
- can't share state



Nagios pain - Interface

the interface is brutal



General

Home

Documentation

Current Status

Tactical Overview

Map

Hosts

Services

Host Groups

Summary

Grid

Service Groups

Summary

Grid

Problems

Services

(Unhandled)

Hosts (Unhandled)

Network Outages

Quick Search:

Reports

Availability

Trends

Alerts

History

Summary

Histogram

Notifications

Event Log

System

Comments

DownTime

Current Network Status
Last Updated: Fri Apr 10 22:04:31 PDT 2015
Updated every 90 seconds
Nagios® Core™ 3.4.1 - www.nagios.org
Logged in as janyder

View History For all hosts
View Notifications For All Hosts
View Host Status Detail For All Hosts

Host Status Totals
Up Down Unreachable Pending
77 0 0 1
All Problems All Types
0 75

Service Status Totals
Ok Warning Unknown Critical Pending
497 6 6 3 0
All Problems All Types
15 512

Display Filters:
Host Status Types: Pending | Up
Host Properties: Any
Service Status Types: All Problems
Service Properties: Not In Scheduled Downtime
Limit Results: All

Service Status Details For All Hosts

Host	Service	Status	Last Check	Duration	Attempt	Status Information
	hostname	UNKNOWN	2015-04-10 22:02:03	176d 5h 20m 33s	3/3	Invalid host name 'batch'
	load	UNKNOWN	2015-04-10 22:03:45	176d 5h 19m 53s	3/3	Invalid host name 'batch'
	mem_free	UNKNOWN	2015-04-10 22:04:08	176d 5h 17m 52s	3/3	Invalid host name 'batch'
	nrpe	UNKNOWN	2015-04-10 22:03:45	176d 5h 22m 0s	3/3	Invalid host name 'batch'
	certification	WARNING	2015-04-10 12:29:28	35d 21h 35m 29s	3/3	WARNING: Hyperthreading should be disabled on fusion to hosts, but is enabled.
	md	UNKNOWN	2015-04-10 22:04:25	26d 19h 39m 44s	3/3	UNKNOWN
	certification	CRITICAL	2015-04-10 12:32:02	197d 12h 30m 43s	3/3	CRITICAL: /etc/mntab and /etc/fstab do not match: /nail/databases (xfs)
	load	WARNING	2015-04-10 22:03:28	0d 0h 18m 19s	3/3	WARNING - load average: 1.00, 1.04, 1.09
	md	UNKNOWN	2015-04-10 22:03:30	4d 1h 54m 7s	3/3	UNKNOWN
	soa_biz_app	WARNING	2015-01-09 21:36:43	90d 23h 27m 48s	1/3	HTTP WARNING: Status line output matched "200 OK" - 267 bytes in 0.902 second response time
	certification	WARNING	2015-04-10 12:30:43	394d 6h 0m 6s	3/3	WARNING: md0 does not have bitmaps enabled, md1 does not have bitmaps enabled
	certification	CRITICAL	2015-04-10 12:33:40	24d 21h 31m 14s	3/3	CRITICAL: /etc/mntab and /etc/fstab do not match: /storage/xrmbakcup (nfs)
	load	WARNING	2015-04-10 22:03:28	0d 0h 3m 5s	3/3	WARNING - load average: 1.15, 0.88, 0.75
	jenkins_queue_status	WARNING	2015-04-10 22:03:47	0d 3h 29m 48s	3/3	Warning: stuck job(s) in the build queue. check https://jenkins.yelpcorp.com/queue/api/json?pretty=true
	sensu_cluster_health	CRITICAL	2015-04-10 22:03:58	0d 0h 0m 33s	1/20	Sensu API: HTTP CRITICAL: HTTP/1.1 503 Service Unavailable - 155 bytes in 0.074 second response time

Results: 1 - 15 of 15 Matching Services

Nagios pain - Workflow

This one is mostly our fault.



Nagios pain - Workflow

- acknowledge and forget
- downtime and forget
- disable notifications and forget
- disable checks and forget



Nagios pain - Workflow

We had a weekly email that would bug us about acknowledged or silenced alerts...




Nagios pain - Workflow

....no one read it.



Let's talk about greener grass



Alternative Nagios strategies

generate Nagios config via puppet

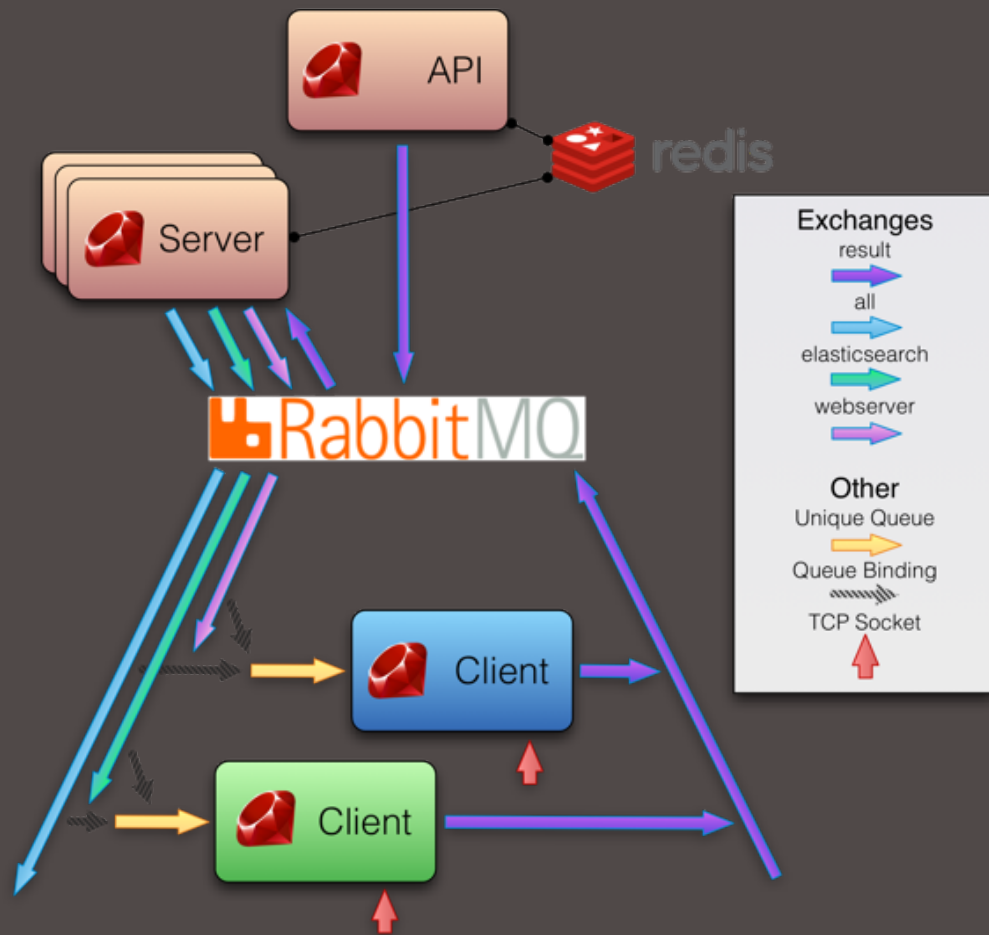
- passive checks?
- centralized checks not associated with nodes?



Sensu is compatible

- uses the Nagios plugin interface
 - existing Nagios checks work in Sensu





Scared?

“Sensu has so many moving parts that I wouldn’t be able to sleep at night unless I set up a Nagios instance to make sure they were all running.”

-- @murphy_slaw



How we use Sensu

- set up in HA clusters
- individual components are HA
- entirely configured in Puppet



Checks are *just* Puppet

- all of our checks run standalone
- Puppet configures the checks that run on a node
- the Sensu server requires no reconfiguration to add new checks



Checks sit next to functionality

```
if $manage_mount {  
  mount { '/nail/databases':  
    ensure => 'present',  
    device => '/dev/vg1/databases',  
    fstype  => 'xfs',  
    options => 'nobootwait,noatime,nobarrier,noquota',  
    before  => Exec['service mysql start'],  
  } ->  
  monitoring_check { "disk_nail_databases":  
    runbook => 'y/disk-nail_databases',  
    team    => 'dba',  
    command => '/usr/lib/nagios/plugins/check_disk --warning=10% --critical=5%  
--iwarning=10% --icritical=10%'  
  }  
}
```

```
{
  "checks": {
    "disk_nail_databases": {
      "command": "/usr/lib/nagios/plugins/check_disk --warning=10% --critical=5% --iwarning=10%
--icritical=10%",
      "interval": 60,
      "handlers": [
        "default"
      ],
      "standalone": true,
      "timeout": 60,
      "dependencies": [],
      "alert_after": 0,
      "realert_every": "-1",
      "runbook": "y/disk-nail_databases",
      "annotation": "modules/profile_services_mysql/manifests/bootstrap.pp:40",
      "sla": "No SLA defined.",
      "team": "dba",
      "irc_channels": "dba-notifications",
      "notification_email": "undef",
      "ticket": false,
      "project": false,
      "page": false,
      "tip": false
    }
  }
}
```

Sensu + Puppet == <3

```
db40-r11-sfo2 ~ >>> ls --color /etc/sensu/conf.d/checks/
check_disk_databases.json      disk_ro_mounts.json          mysql_puppet_config_differs.json
check_kernel_bugs.json         em1_link.json               mysql_read_only_state.json
check_ntp_timesync.json        fusion_io_driver_version.json mysql_slave_status.json
check_puppet_fast_feedback.json fusion_io_free_lebs.json     mysql_status.json
check_puppet_staleness.json    fusion_io_reserve_space.json mysql_temp_tables_on_disk.json
check_slash.json               hachek_process_running.json mysql_temp_tables_size_on_disk.json
cron_haproxy_synapse_reaper_staleness.json hardware_raid.json          nerve_config_generation_staleness.json
cron_innodb-log_staleness.json ipmi_certification.json     nerve_process_running.json
cron_iproute2_git_commit_staleness.json lldpd.json                  proc_xinetd.json
cron.json                      mcelog.json                 smart.json
cron_long-transactions_staleness.json mcollective_check_connectivity.json sshd.json
cron_manage_grants_staleness.json  mysql_auto_inc_capacity.json swap-rate.json
cron_ps-log_staleness.json         mysql_config_diff.json      synapse_config_generation_staleness.json
cron_set_ipmi_sol_baud_staleness.json mysql_config_diff_long.json synapse_haproxy_process_running.json
cron_update_hardware_data_staleness.json mysql_has_role.json         synapse_process_running.json
cron_yelp_ipmi_set_credentials_staleness.json mysql_heartbeat.json        syslog-ng.json
default-routes.json             mysql_history_length.json    tmp-file-count.json
dhclient_not_running.json        mysql_long_running_queries.json yelpsoa-configs-staleness.json
diamond_running.json            mysql_port.json
```



How do MySQL and Sensu interact?

- host-based condition monitoring
 - not graphing
- monitor health of MySQL and its hardware



What do we check?

```
db40-r11-sfo2 ~ >>> ls --color /etc/sensu/conf.d/checks/
```

check_disk_databases.json	disk_ro_mounts.json	mysql_puppet_config_differs.json
check_kernel_bugs.json	em1_link.json	mysql_read_only_state.json
check_ntp_timesync.json	fusion_io_driver_version.json	mysql_slave_status.json
check_puppet_fast_feedback.json	fusion_io_free_lebs.json	mysql_status.json
check_puppet_staleness.json	fusion_io_reserve_space.json	mysql_temp_tables_on_disk.json
check_slash.json	hacheck_process_running.json	mysql_temp_tables_size_on_disk.json
cron_haproxy_synapse_reaper_staleness.json	hardware_raid.json	nerve_config_generation_staleness.json
cron_innodb-log_staleness.json	ipmi_certification.json	nerve_process_running.json
cron_iproute2_git_commit_staleness.json	lldpd.json	proc_xinetd.json
cron.json	mcelog.json	smart.json
cron_long-transactions_staleness.json	mcollective_check_connectivity.json	sshd.json
cron_manage_grants_staleness.json	mysql_auto_inc_capacity.json	swap-rate.json
cron_ps-log_staleness.json	mysql_config_diff.json	synapse_config_generation_staleness.json
cron_set_ipmi_sol_baud_staleness.json	mysql_config_diff_long.json	synapse_haproxy_process_running.json
cron_update_hardware_data_staleness.json	mysql_has_role.json	synapse_process_running.json
cron_yelp_ipmi_set_credentials_staleness.json	mysql_heartbeat.json	syslog-ng.json
default-routes.json	mysql_history_length.json	tmp-file-count.json
dhclient_not_running.json	mysql_long_running_queries.json	yelpsoa-configs-staleness.json
diamond_running.json	mysql_port.json	



More Puppet <3

```
cron::d { 'long-transactions':  
  command      => "/nail/sys/bin/record_long_transactions ${arguments}",  
  user          => 'mysql',  
  minute        => '*/5',  
  staleness_threshold => '15m',  
  staleness_check_params => {  
    team        => 'dba',  
    page        => false,  
    runbook     => 'y/cron_staleness',  
  },  
}
```



More Puppet <3

```
"checks": {  
  "cron_long-transactions_staleness": {  
    "command": "/nail/sys/bin/cron_staleness_check cron_long-transactions 900",  
    "interval": 180,  
    "handlers": [  
      "default"  
    ],  
    "standalone": true,  
    "timeout": 180,  
    "dependencies": [],  
    "alert_after": 0,  
    "realert_every": "-1",  
    "runbook": "y/cron_staleness",  
  },  
}
```



How do we work with Sensu?



Alerts vs. Events: Nagios

Nagios is normally used for everything:

- scheduling
- emailing
- raising alerts
- notification
- escalation
- acknowledgement



Alerts vs. Events: Sensus

Sensus does these things

- runs checks
- generates events
- calls handlers



Disabling/Stashing checks

- stash checks/host
- sensu-cli tools
 - silence around maintenance/reboots
 - can query API for stashed alerts
- our stashed alerts come back



Stash example

```
jsnyder@dev5-devc:~$ sensu-cli stash list
```

```
-----  
path:  silence/netenglab1-devc.dev.yelpcorp.com/check_puppet_fast_feedback  
content: {"reason"=>"Network labe disabled b/c of cooling issues", "source"=>"uchiwa", "timestamp"=>1428702219.0}  
expire: -1
```

```
-----  
path:  silence/netenglab1-devc.dev.yelpcorp.com/ipmi_certification  
content: {"reason"=>"Network labe disabled b/c of cooling issues", "source"=>"uchiwa", "timestamp"=>1428702219.0}  
expire: -1
```

```
-----  
path:  silence/netenglab1-devc.dev.yelpcorp.com/check_puppet_staleness  
content: {"reason"=>"Network labe disabled b/c of cooling issues", "source"=>"uchiwa", "timestamp"=>1428702219.0}  
expire: -1
```

```
-----  
path:  silence/db8-r1-devc.dev.yelpcorp.com  
content: {"timestamp"=>1428700867}  
expire: -1
```



Sensu notification handlers

our Sensu handlers can

- announce in IRC
- file a (JIRA) ticket
- page via PagerDuty
- send email



Sensu notification handlers

```
from sensu/conf.d/handlers/mailer.json:
{
  "mailer": {
    "teams": {
      "dba": {
        "pagerduty_api_key": "sdlkfjasdklfjasdklfjaskdjflaksjf",
        "pages_irc_channel": "dba-pages",
        "notifications_irc_channel": "dba-notifications",
        "project": "DBA"
      },
      "test": {
        "pagerduty_api_key": "qepoiruqewroiugweriqwperweiruwer",
        "pages_irc_channel": "test-pages",
        "notifications_irc_channel": "test-notifications"
      }
    }
  }
}
```

Less used notification handlers

We also have:

- Graphite
- aws_prune
- OpsGenie



Users and events

Users interact with their notifications via:

- PagerDuty
 - IRC bot, mobile app, SMS, web UI

Can see overview via:

- Uchiwa



Uchiwa

uchiwa



EVENTS >

ALL DATACENTERS ▾

ACTIONS ▾

HIDE ▾

50 OF 187 ▾

Search

Source ↕	Check ↕	Output ↕	⚙️ ↕	☁️ ↕	🕒 ↕
<input type="checkbox"/> 🔊 db8-r1-devc.dev.yelpcorp.com	🔊 mysql_heartbeat	cat: /nail/etc/mysql/roles/current/heartbeat_thresho... 142591		localhost	2015-04-10 08:50:59
<input type="checkbox"/> 🔊 db8-r1-devc.dev.yelpcorp.com	🔊 mysql_slave_status	No replication information for 127.0.0.1 35539		localhost	2015-04-10 08:50:35
<input type="checkbox"/> 🔊 dev34-devc.dev.yelpcorp.com	🔊 check_puppet_staleness	CRITICAL: Run has errors (finished 0h 2m 59s ago) E... 1		localhost	2015-04-10 08:25:00
<input type="checkbox"/> 🔊 dev5-devc.dev.yelpcorp.com	🔊 check_ntp_timesync	NTP CRITICAL: Server not synchronized, Offset unkn... 21		localhost	2015-04-10 08:50:04
<input type="checkbox"/> 🔊 db6-r6-devc.dev.yelpcorp.com	🔊 check_ntp_timesync	NTP CRITICAL: Server not synchronized, Offset unkn... 21		localhost	2015-04-10 08:50:26
<input type="checkbox"/> 🔊 db6-r6-devc.dev.yelpcorp.com	🔊 check_puppet_staleness	CRITICAL: Run has errors (finished 0h 1m 42s ago) E... 1		localhost	2015-04-10 08:09:43
<input type="checkbox"/> 🔊 db6-r6-devc.dev.yelpcorp.com	🔊 check_puppet_fast_feedback	CRITICAL: Run has errors (finished 0h 41m 34s ago) ... 5		localhost	2015-04-10 08:49:35
<input type="checkbox"/> 🔊 irc1-devc.dev.yelpcorp.com	🔊 check_ntp_timesync	NTP CRITICAL: Server not synchronized, Offset unkn... 5		localhost	2015-04-10 08:50:46
<input type="checkbox"/> 🔊 irc1-devc.dev.yelpcorp.com	🔊 check_puppet_staleness	CRITICAL: Run has errors (finished 0h 1m 36s ago) E... 1		localhost	2015-04-10 08:25:37

Event resolution


Sensu

- announces in IRC
- can close a ticket
- PagerDuty
- sends email



motd

```
jsnyder@batch23-r1-sfo2:~$ ssh db23-r5-sfo2
```



Don't Panic! Ops already knows! Failed Sensu checks on this host:

Crit: mysql_puppet_config_differs: 1 config difference Variable ...

Crit: check_puppet_fast_feedback: CRITICAL: Run has errors (finished 0h 1...

cluster dw; role(s) replicas x1

```
jsnyder@db23-r5-sfo2:~$
```


Are we totally dependent on Puppet?



Are we totally dependent on Puppet?

- Puppet runs hourly
- We may want to change checks more frequently
- Problem? Certainly not



(1) Smart checks

```
cursor.execute("select @@read_only")  
(read_only,) = cursor.fetchone()  
if read_only == 1:  
    finish("OK", ["this check only cares about masters"])
```



(2) Dynamic thresholds

```
{  
  "checks": {  
    "mysql_heartbeat": {  
      "command": "/usr/lib/nagios/plugins/yelp/check_heartbeat.py -n 127.0.0.1 -p 3306  
$(cat /nail/etc/mysql/roles/current/heartbeat_threshold)",  
      "interval": 15,  
      "handlers": [  
        "default"  
      ],  
      "standalone": true,  
      "timeout": 15,  
    }  
  }  
}
```



Symlinks to the rescue!

```
db40-r11-sfo2 /n/e/m/roles >>> ls -l
total 12
dr-xr-xr-x 3 root root 4096 Mar 24 17:09 master
dr-xr-xr-x 3 root root 4096 Mar 19 23:09 read
dr-xr-xr-x 3 root root 4096 Mar 24 17:09 reporting
lrwxrwxrwx 1 root root    4 Apr  2 12:38 current -> read
```



(3) Symlinked checks

- not in use at Yelp
- allows tweaking of:
 - intervals
 - handlers



Making the switch

- compatibility helps
- there are still differences
- make the best use of customized check routing
 - start non-alerting
 - gain confidence
 - add in alerting




Conclusion

- it's an adjustment
- Sensus lets us do more with less
- we still need to migrate some checks



Questions?



Follow us! [@hashbrowncipher](#) & [@jcsuperstar](#)

Related GitHub Projects:

github.com/yelp/sensu_handlers

github.com/yelp/pysensu-yelp

Kyle's ([@solarkennedy](#)) talk at SFDevOps:

[Part 1](#)

[Part 2](#)

Tom's ([@bobtfish](#)) from Puppetconf 2014
[Sensu and Sensibility](#)



YelpEngineers



@YelpEngineering



engineeringblog.yelp.com



github.com/yelp

Come see us at booth 402!

