# Monitoring the Cloud

Open Science Data Cloud
University of Edinburgh
16-19th July 2012

**Jonathan Stuart Ward**
**Big Data Group,**
**School of Computer Science,**
**University of St Andrews**

**jonathan.stuart.ward@st-andrews.ac.uk**
**@jauntyward**

Introduction

Current Monitoring Systems

Cloud Monitoring

How can we better understand large scale systems which frequently change?

Cloud Computing allows anyone to deploy large scale systems

IaaS cloud systems
are unlike
conventional
systems

conventional { clusters
grids
static servers

# rapid elasticity,

rapid elasticity,
transient, identical VMs,

rapid elasticity,
transient, identical VMs,
metered service,

rapid elasticity,
transient, identical VMs,
metered service,
VM termination

rapid elasticity,
transient, identical VMs,
metered service,
VM termination
Irregular deployment

Despite these differences we still use the same largely unmodified software stacks

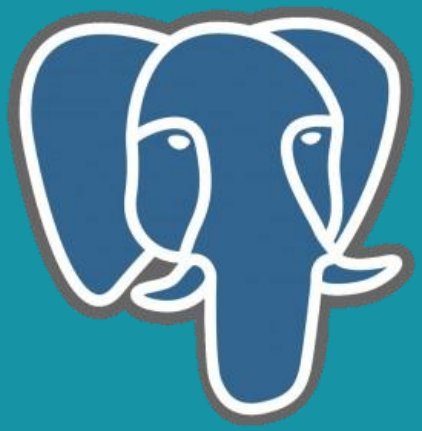

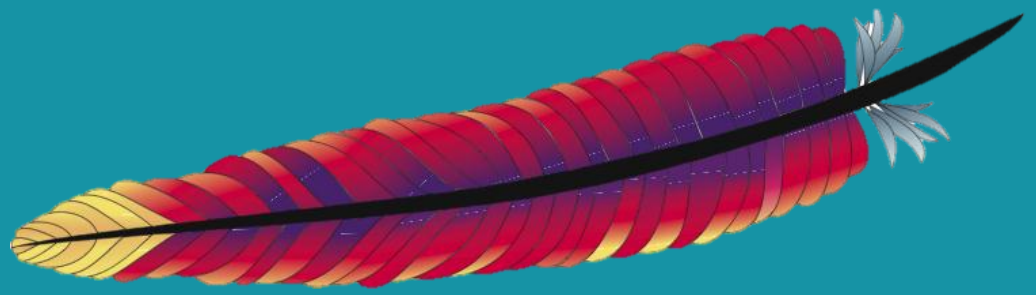


RAILS

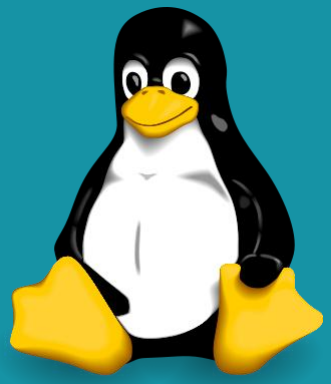


php


Nagios


MySQL™

# Monitoring

# System monitoring

Monitoring and reporting resource usage

# Infrastructure monitoring

Monitoring the current state of a system

Application monitoring, fault detection, fault correction, fault prediction, capacity planning, intrusion detection, reporting, alerts, SLA compliance, business process coordination

The capture of relevant state

Introduction

Current Monitoring Systems

Cloud Monitoring

Current monitoring is slow to adapt, scales poorly and structures information in a way that neither machines or humans can easily understand

# Nagios

De facto standard
open source
monitoring system

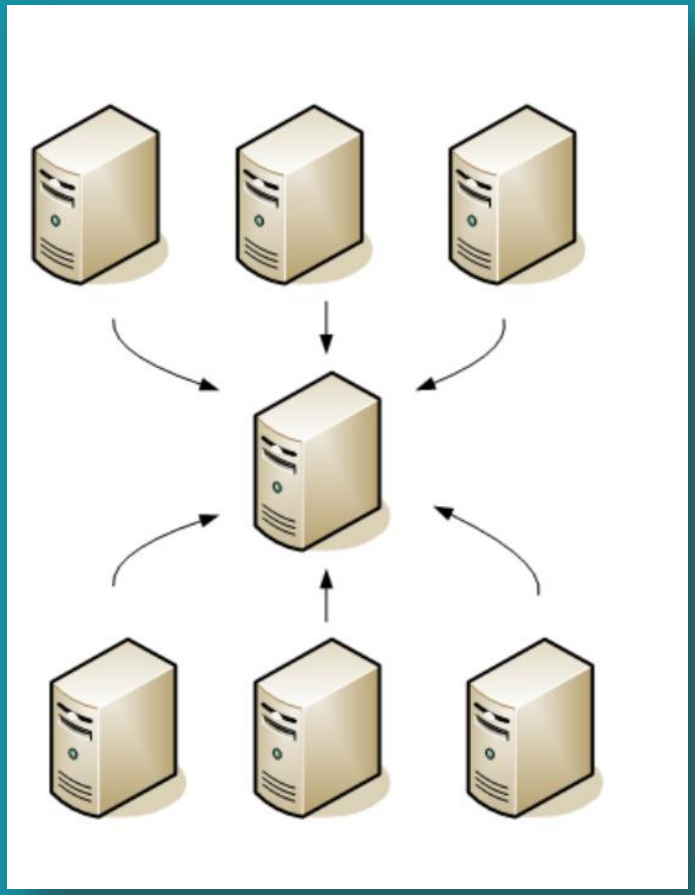

nagios.com

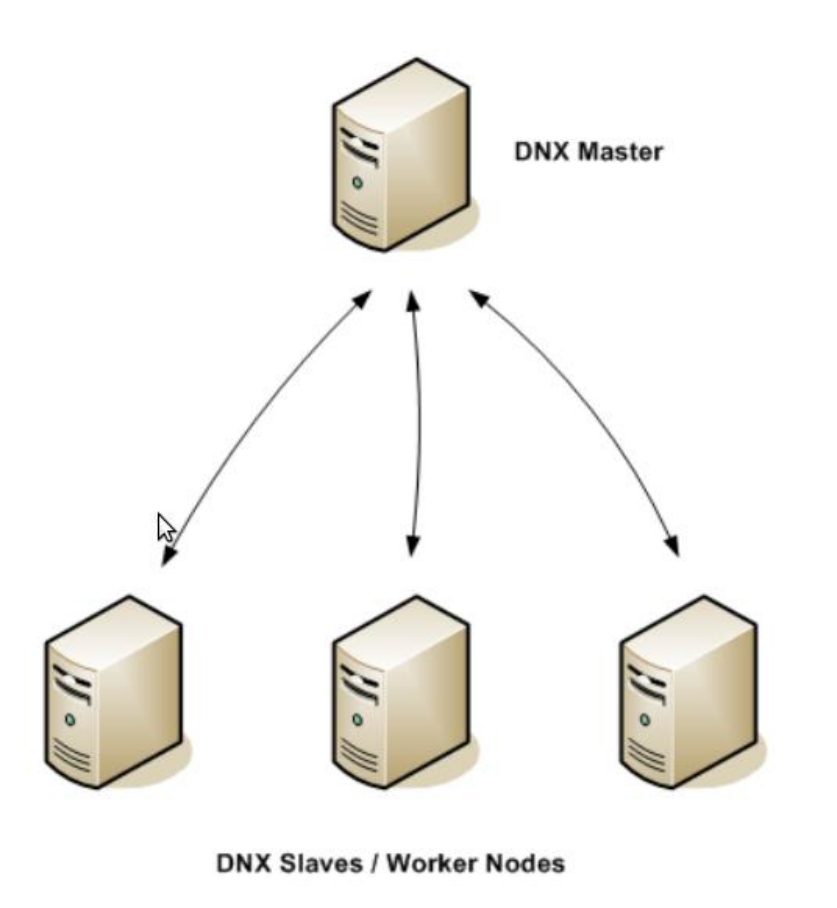


nagios.com

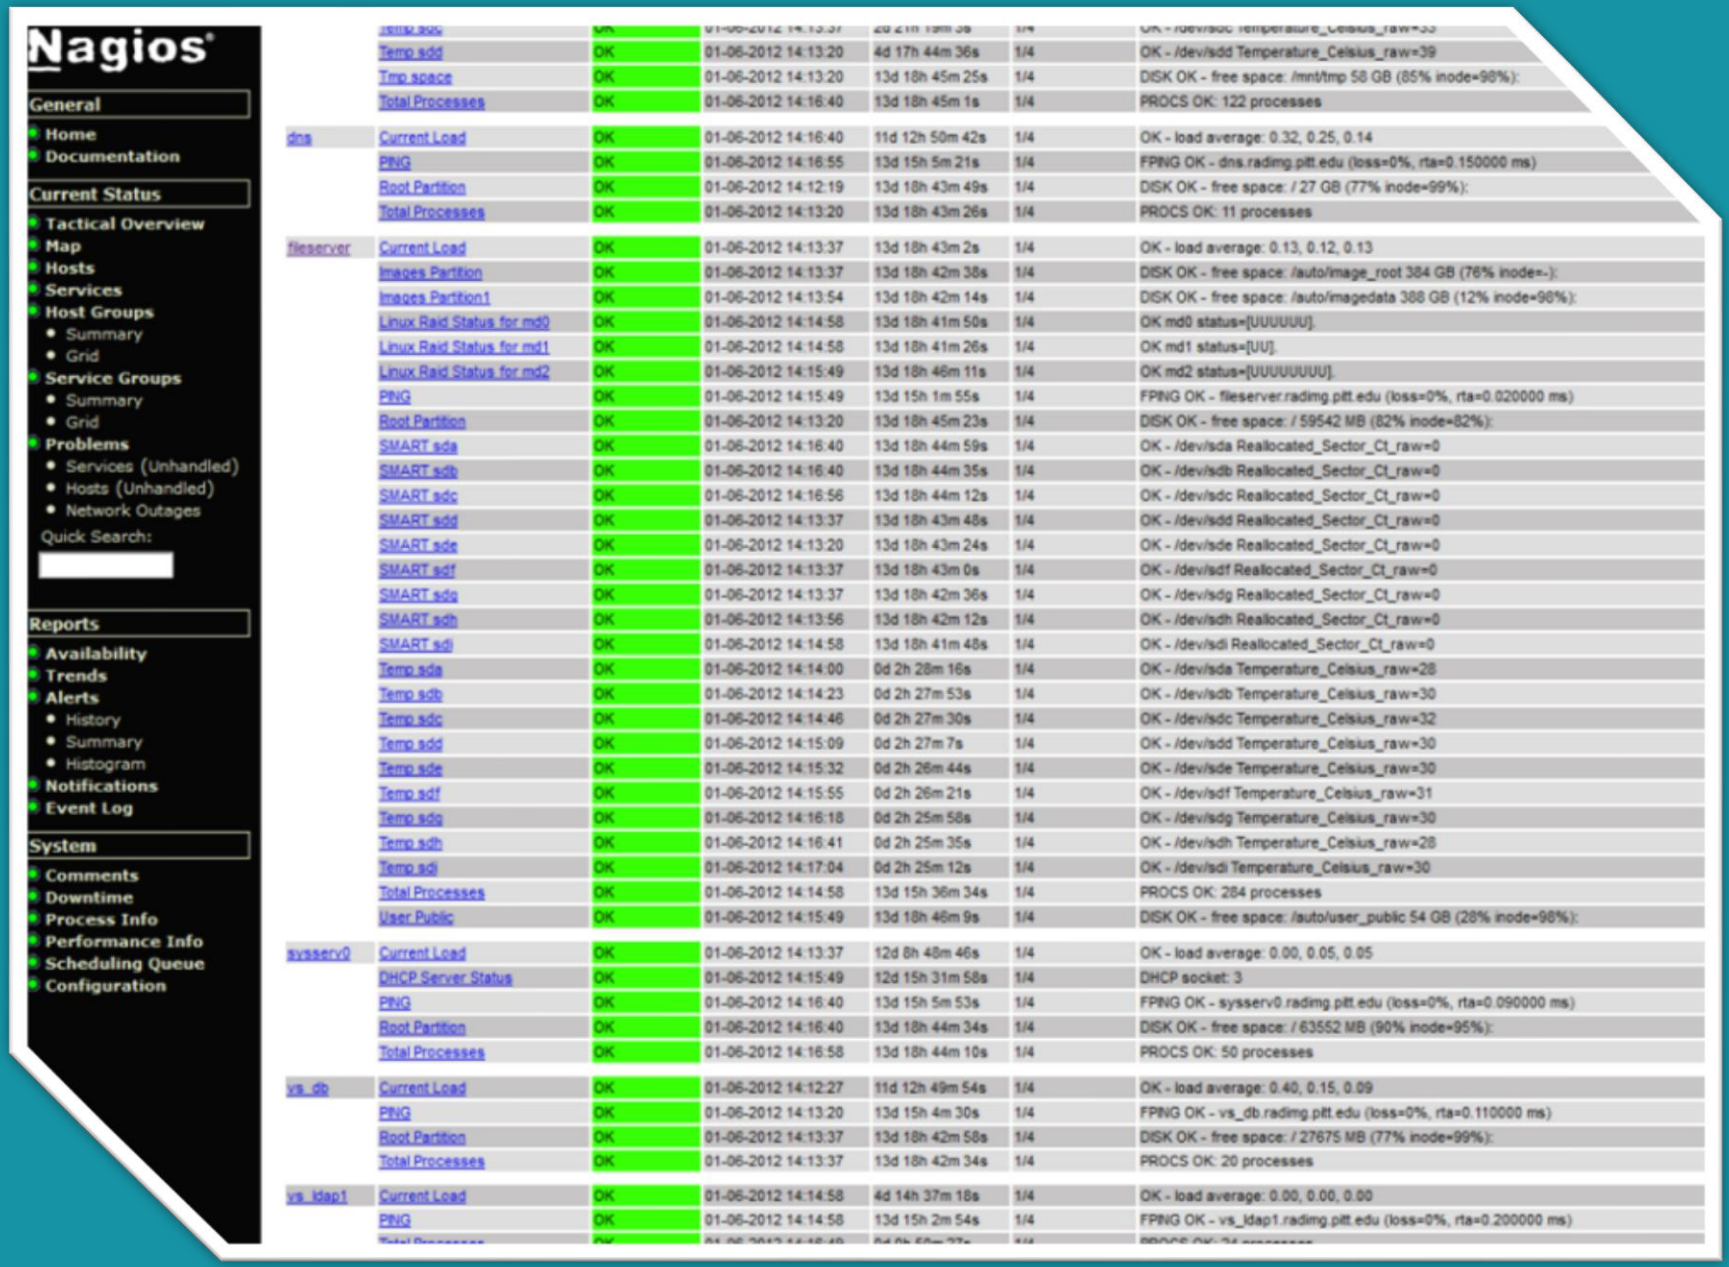


| | | | | | | |
|---|---|---|---|---|---|---|
| | Temp sdd | OK | 01-06-2012 14:13:20 | 4d 17h 44m 36s | 1/4 | OK - /dev/sdd Temperature_Celsius_raw=39 |
| | Tmp space | OK | 01-06-2012 14:13:20 | 13d 18h 45m 25s | 1/4 | DISK OK - free space: /mnt/tmp 58 GB (85% inode=98%): |
| | Total Processes | OK | 01-06-2012 14:16:40 | 13d 18h 45m 1s | 1/4 | PROCS OK: 122 processes |
| dns | Current Load | OK | 01-06-2012 14:16:40 | 11d 12h 50m 42s | 1/4 | OK - load average: 0.32, 0.25, 0.14 |
| | PING | OK | 01-06-2012 14:16:55 | 13d 15h 5m 21s | 1/4 | FPING OK - dns.radimg.pitt.edu (loss=0%, rta=0.150000 ms) |
| | Root Partition | OK | 01-06-2012 14:12:19 | 13d 18h 43m 49s | 1/4 | DISK OK - free space: / 27 GB (77% inode=99%): |
| | Total Processes | OK | 01-06-2012 14:13:20 | 13d 18h 43m 26s | 1/4 | PROCS OK: 11 processes |
| fileserver | Current Load | OK | 01-06-2012 14:13:37 | 13d 18h 43m 2s | 1/4 | OK - load average: 0.13, 0.12, 0.13 |
| | Images Partition | OK | 01-06-2012 14:13:37 | 13d 18h 42m 38s | 1/4 | DISK OK - free space: /auto/image_root 384 GB (76% inode=-): |
| | Images Partition1 | OK | 01-06-2012 14:13:54 | 13d 18h 42m 14s | 1/4 | DISK OK - free space: /auto/imagedata 388 GB (12% inode=98%): |
| | Linux Raid Status for md0 | OK | 01-06-2012 14:14:58 | 13d 18h 41m 50s | 1/4 | OK md0 status=[UUUUUU]. |
| | Linux Raid Status for md1 | OK | 01-06-2012 14:14:58 | 13d 18h 41m 26s | 1/4 | OK md1 status=[UU]. |
| | Linux Raid Status for md2 | OK | 01-06-2012 14:15:49 | 13d 18h 46m 11s | 1/4 | OK md2 status=[UUUUUUUU]. |
| | PING | OK | 01-06-2012 14:15:49 | 13d 15h 1m 55s | 1/4 | FPING OK - fileserver.radimg.pitt.edu (loss=0%, rta=0.020000 ms) |
| | Root Partition | OK | 01-06-2012 14:13:20 | 13d 18h 45m 23s | 1/4 | DISK OK - free space: / 59542 MB (82% inode=82%): |
| | SMART sda | OK | 01-06-2012 14:16:40 | 13d 18h 44m 59s | 1/4 | OK - /dev/sda Reallocated_Sector_Ct_raw=0 |
| | SMART sdb | OK | 01-06-2012 14:16:40 | 13d 18h 44m 35s | 1/4 | OK - /dev/sdb Reallocated_Sector_Ct_raw=0 |
| | SMART sdc | OK | 01-06-2012 14:16:56 | 13d 18h 44m 12s | 1/4 | OK - /dev/sdc Reallocated_Sector_Ct_raw=0 |
| | SMART sdd | OK | 01-06-2012 14:13:37 | 13d 18h 43m 48s | 1/4 | OK - /dev/sdd Reallocated_Sector_Ct_raw=0 |
| | SMART sde | OK | 01-06-2012 14:13:20 | 13d 18h 43m 24s | 1/4 | OK - /dev/sde Reallocated_Sector_Ct_raw=0 |
| | SMART sdf | OK | 01-06-2012 14:13:37 | 13d 18h 43m 0s | 1/4 | OK - /dev/sdf Reallocated_Sector_Ct_raw=0 |
| | SMART sdg | OK | 01-06-2012 14:13:37 | 13d 18h 42m 36s | 1/4 | OK - /dev/sdg Reallocated_Sector_Ct_raw=0 |
| | SMART sdh | OK | 01-06-2012 14:13:56 | 13d 18h 42m 12s | 1/4 | OK - /dev/sdh Reallocated_Sector_Ct_raw=0 |
| | SMART sdi | OK | 01-06-2012 14:14:58 | 13d 18h 41m 48s | 1/4 | OK - /dev/sdi Reallocated_Sector_Ct_raw=0 |
| | Temp sda | OK | 01-06-2012 14:14:00 | 0d 2h 28m 16s | 1/4 | OK - /dev/sda Temperature_Celsius_raw=28 |
| | Temp sdb | OK | 01-06-2012 14:14:23 | 0d 2h 27m 53s | 1/4 | OK - /dev/sdb Temperature_Celsius_raw=30 |
| | Temp sdc | OK | 01-06-2012 14:14:46 | 0d 2h 27m 30s | 1/4 | OK - /dev/sdc Temperature_Celsius_raw=32 |
| | Temp sdd | OK | 01-06-2012 14:15:09 | 0d 2h 27m 7s | 1/4 | OK - /dev/sdd Temperature_Celsius_raw=30 |
| | Temp sde | OK | 01-06-2012 14:15:32 | 0d 2h 26m 44s | 1/4 | OK - /dev/sde Temperature_Celsius_raw=30 |
| | Temp sdf | OK | 01-06-2012 14:15:55 | 0d 2h 26m 21s | 1/4 | OK - /dev/sdf Temperature_Celsius_raw=31 |
| | Temp sdg | OK | 01-06-2012 14:16:18 | 0d 2h 25m 58s | 1/4 | OK - /dev/sdg Temperature_Celsius_raw=30 |
| | Temp sdh | OK | 01-06-2012 14:16:41 | 0d 2h 25m 35s | 1/4 | OK - /dev/sdh Temperature_Celsius_raw=28 |
| | Temp sdi | OK | 01-06-2012 14:17:04 | 0d 2h 25m 12s | 1/4 | OK - /dev/sdi Temperature_Celsius_raw=30 |
| | Total Processes | OK | 01-06-2012 14:14:58 | 13d 15h 36m 34s | 1/4 | PROCS OK: 284 processes |
| | User Public | OK | 01-06-2012 14:15:49 | 13d 18h 46m 9s | 1/4 | DISK OK - free space: /auto/user_public 54 GB (28% inode=98%): |
| sysserv0 | Current Load | OK | 01-06-2012 14:13:37 | 12d 8h 48m 46s | 1/4 | OK - load average: 0.00, 0.05, 0.05 |
| | DHCP Server Status | OK | 01-06-2012 14:15:49 | 12d 15h 31m 58s | 1/4 | DHCP socket: 3 |
| | PING | OK | 01-06-2012 14:16:40 | 13d 15h 5m 53s | 1/4 | FPING OK - sysserv0.radimg.pitt.edu (loss=0%, rta=0.090000 ms) |
| | Root Partition | OK | 01-06-2012 14:16:40 | 13d 18h 44m 34s | 1/4 | DISK OK - free space: / 63552 MB (90% inode=95%): |
| | Total Processes | OK | 01-06-2012 14:16:58 | 13d 18h 44m 10s | 1/4 | PROCS OK: 50 processes |
| vs_db | Current Load | OK | 01-06-2012 14:12:27 | 11d 12h 49m 54s | 1/4 | OK - load average: 0.40, 0.15, 0.09 |
| | PING | OK | 01-06-2012 14:13:20 | 13d 15h 4m 30s | 1/4 | FPING OK - vs_db.radimg.pitt.edu (loss=0%, rta=0.110000 ms) |
| | Root Partition | OK | 01-06-2012 14:13:37 | 13d 18h 42m 58s | 1/4 | DISK OK - free space: / 27675 MB (77% inode=99%): |
| | Total Processes | OK | 01-06-2012 14:13:37 | 13d 18h 42m 34s | 1/4 | PROCS OK: 20 processes |
| vs_ldap1 | Current Load | OK | 01-06-2012 14:14:58 | 4d 14h 37m 18s | 1/4 | OK - load average: 0.00, 0.00, 0.00 |
| | PING | OK | 01-06-2012 14:14:58 | 13d 15h 2m 54s | 1/4 | FPING OK - vs_ldap1.radimg.pitt.edu (loss=0%, rta=0.200000 ms) |

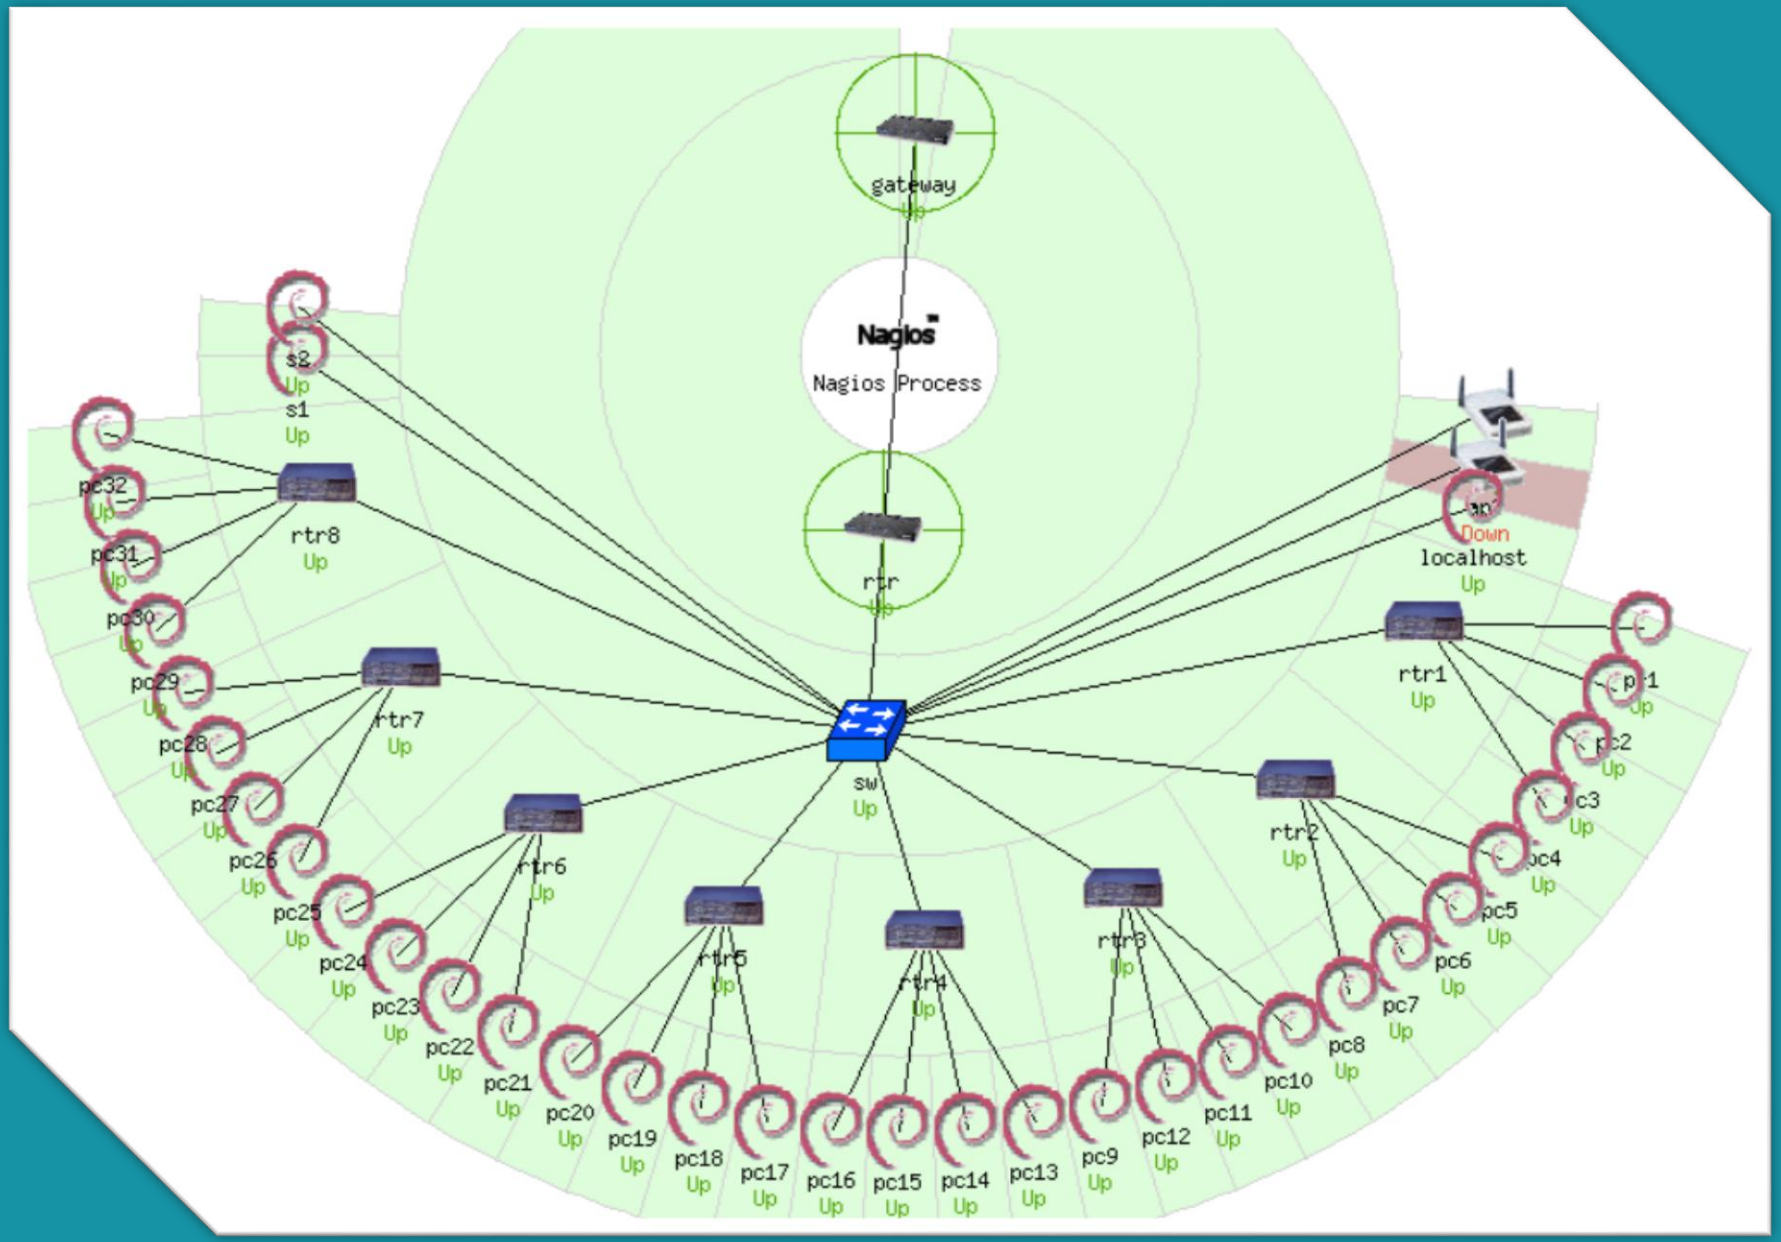
gateway
Up
Nagios
Nagios Process
rtr
Up
sw
Up
s2
Up
s1
Up
rtr8
Up
rtr7
Up
rtr6
Up
rtr5
Up
rtr4
Up
rtr3
Up
rtr2
Up
rtr1
Up
localhost
Up
pc32
pc31
pc30
pc29
pc28
pc27
pc26
pc25
pc24
pc23
pc22
pc21
pc20
pc19
pc18
pc17
pc16
pc15
pc14
pc13
pc9
pc12
pc11
pc10
pc8
pc7
pc6
pc5
pc4
pc3
pc2
Down

Nagios

General
- Home
- Documentation

Monitoring
- Tactical Overview
- Service Detail
- Host Detail
- Hostgroup Overview
- Hostgroup Summary
- Hostgroup Grid
- Servicegroup Overview
- Servicegroup Summary
- Servicegroup Grid
- Status Map
- 3-D Status Map
- Service Problems
- Host Problems
- Network Outages

Show Host:

- Comments
- Downtime
- Process Info
- Performance Info
- Scheduling Queue

Reporting
- Trends
- Availability
- Alert Histogram
- Alert History
- Alert Summary
- Notifications
- Event Log

Configuration
- View Config

**Network Map For All Hosts**
Last Updated: Tue Sep 12 18:43:45 MDT 2006
Updated every 90 seconds
Nagios® - www.nagios.org
Logged in as *a9k*

View Status Detail For All Hosts
View Status Overview For All Hosts

Layout Method: Circular
Scaling factor: 0.0
Drawing Layers: DNS, HotSpots, MAIL, Routers
Layer mode: Include / Exclude
Suppress popups:
Update

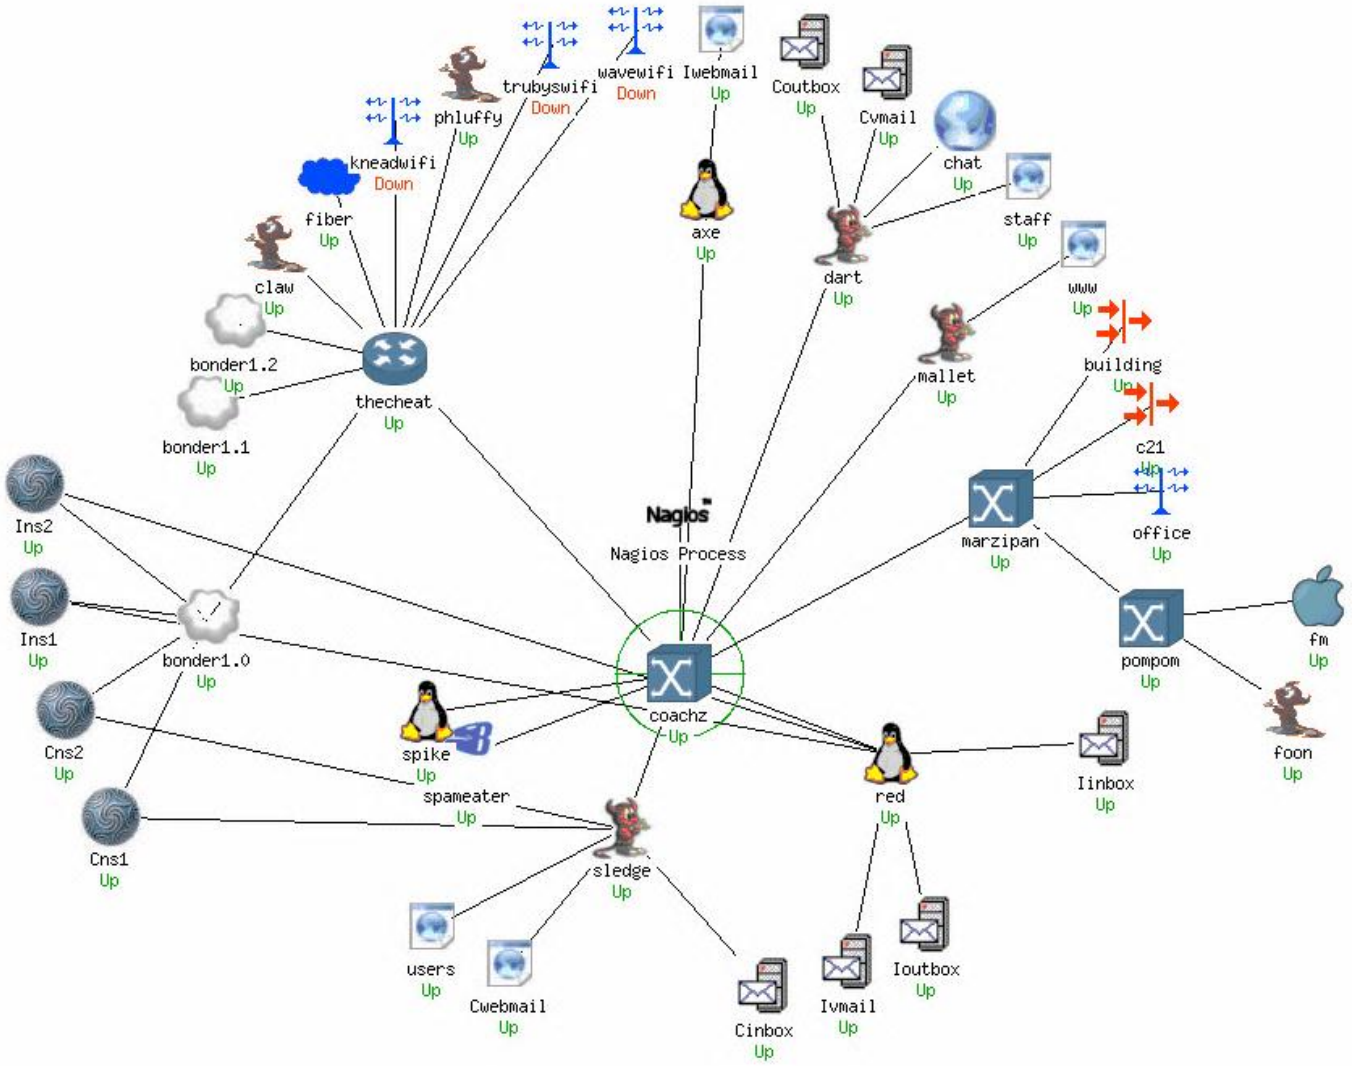

```
define service{
use                 generic-http
hostgroup_name      firewall-htt
contact_groups      firewall-admins
}

define hostgroup {
hostgroup_name      linux-host
alias               Debian  Servers
members             localhost, web-server
}

define host{
    use             linux-host
    host_name       web-server.domain
    alias           web-server.domain
    address         10.0.0.1
    hostgroups      http-servers, file-servers
}
```

# Ganglia

Scalable monitoring system for HPC systems

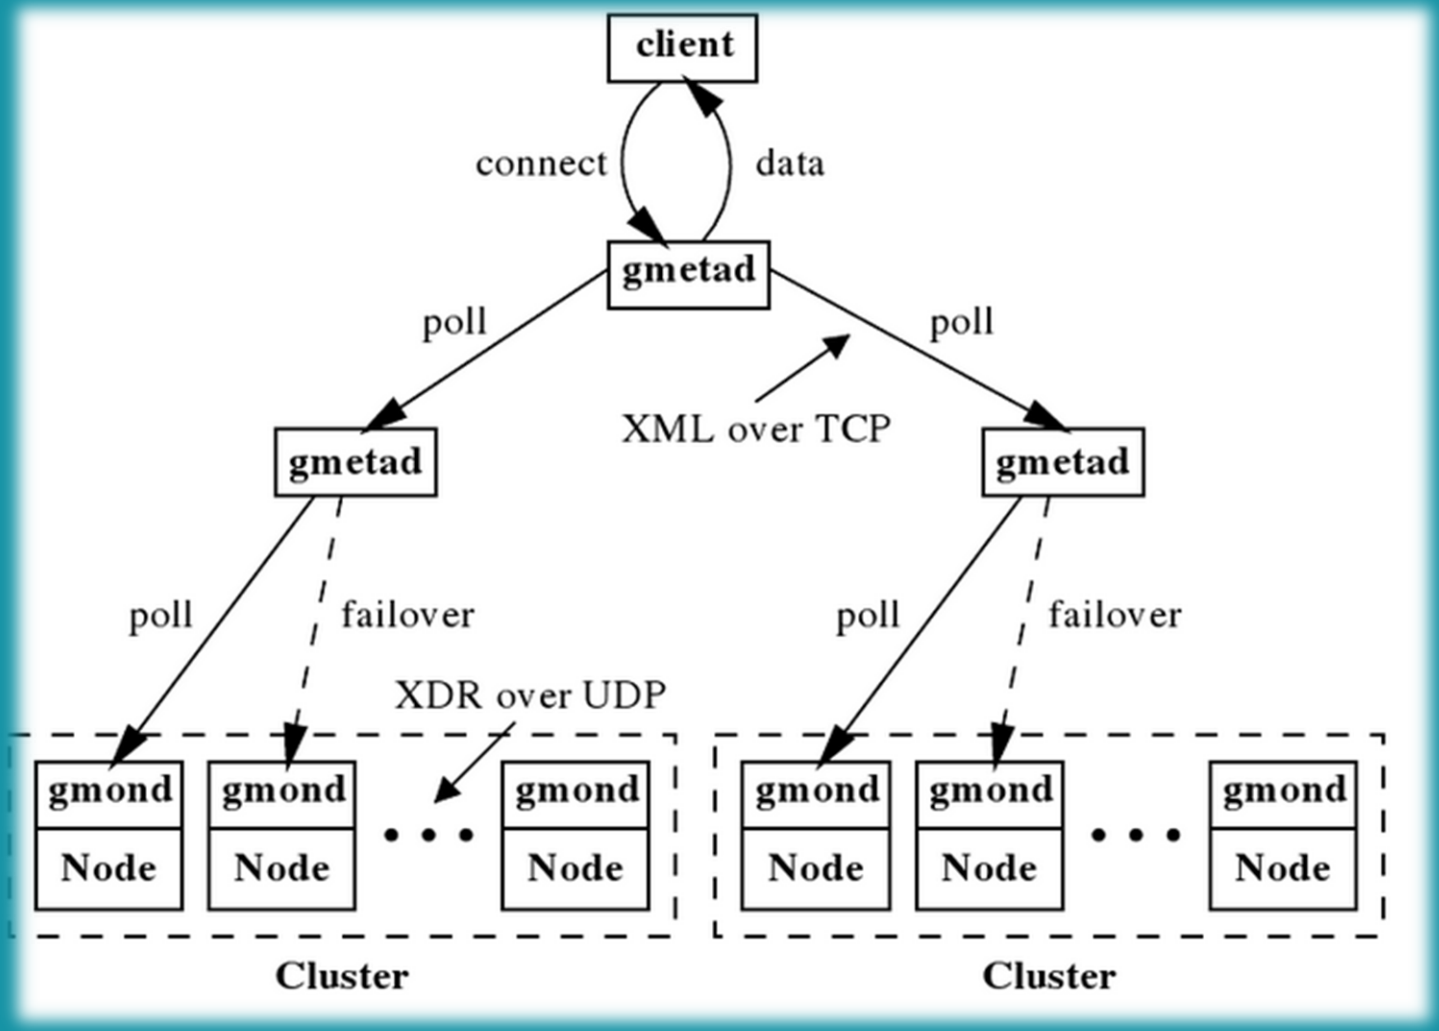


Massie et al, 2004

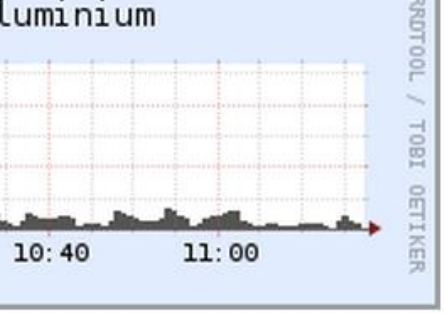


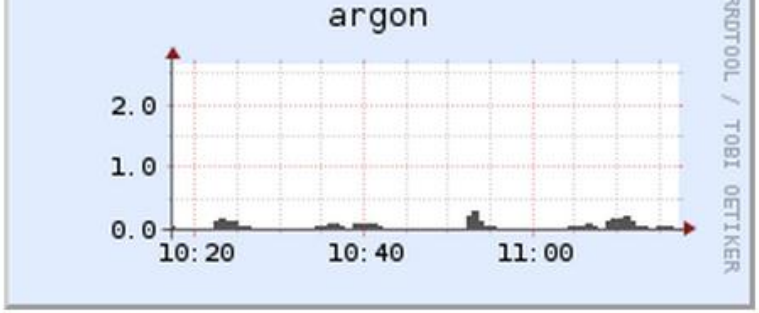


dataset1001.wikimedia.org

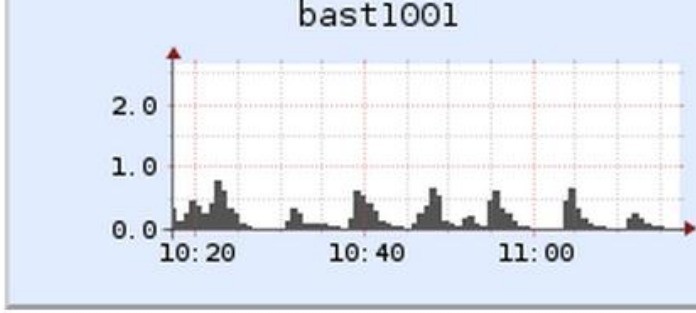


gallium.wikimedia.org

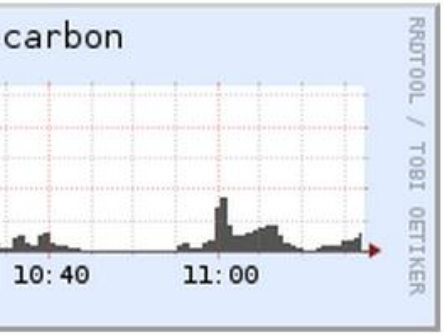


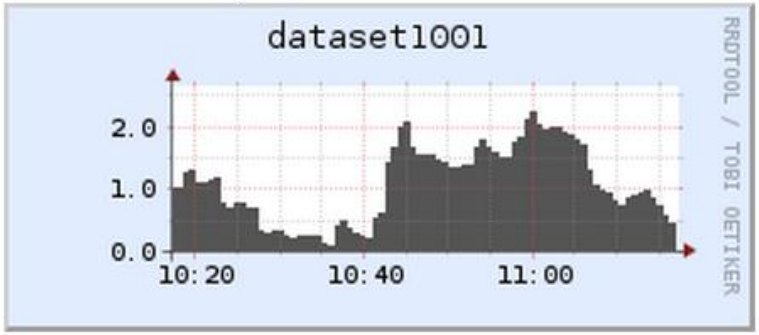


lvs1002.wikimedia.org

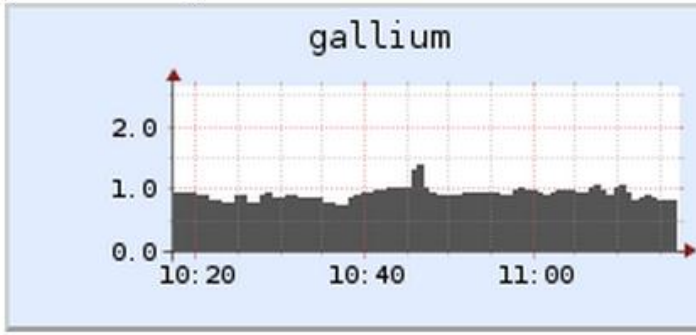


lvs1003.wikimedia.org

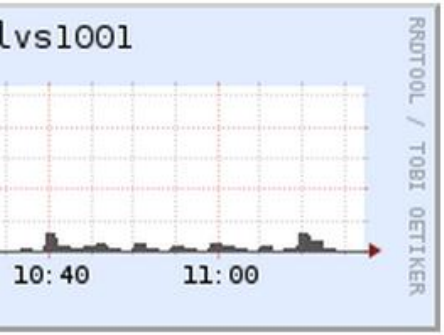


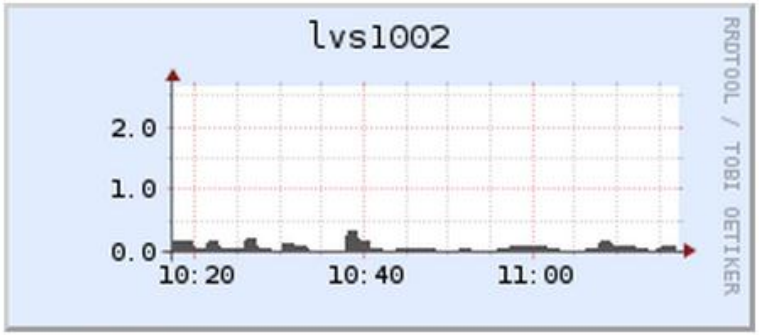


lvs1006.wikimedia.org

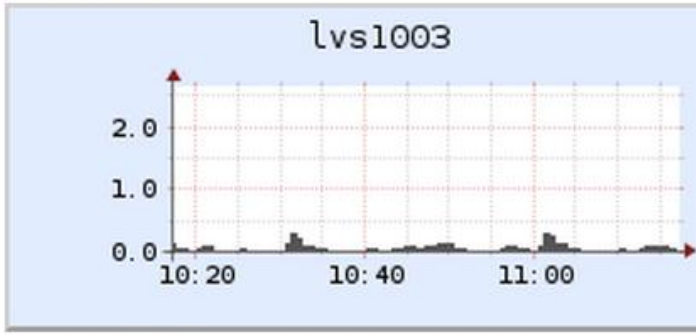


manganese.wikimedia.org

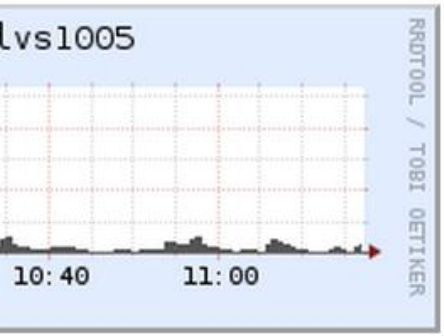


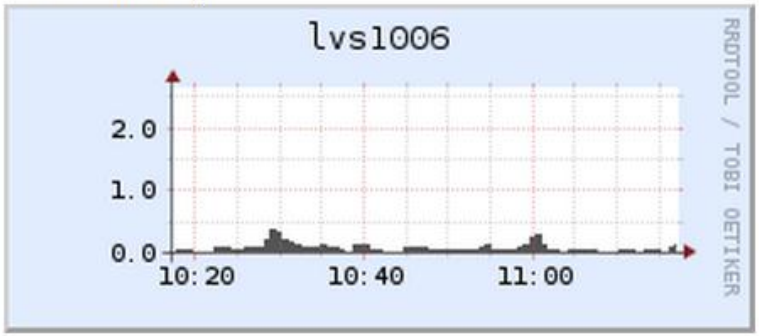


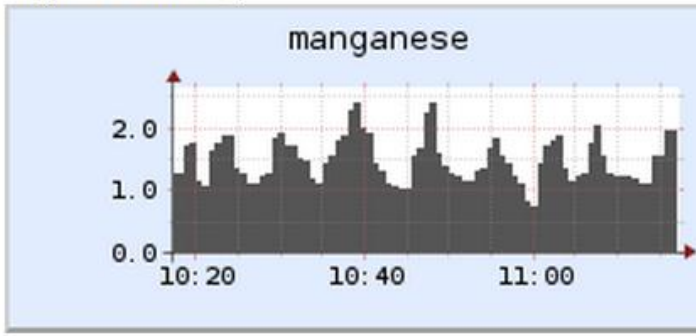

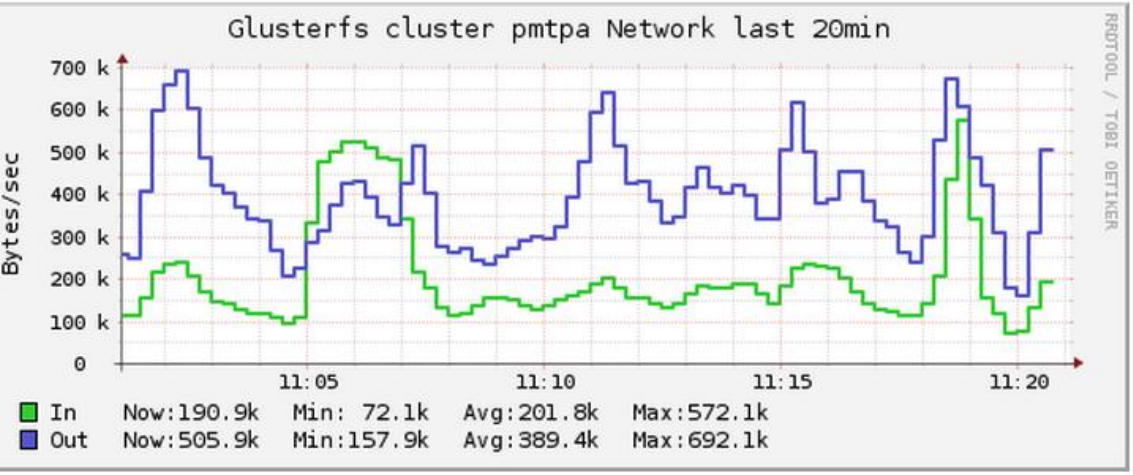


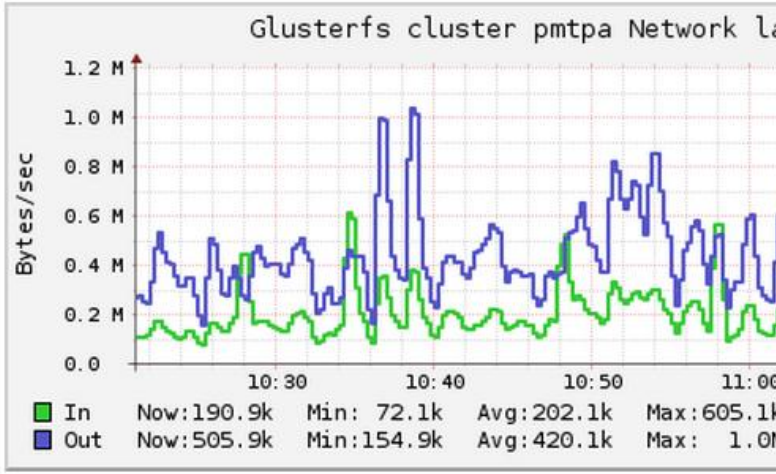


2hr CSV JSON Inspect

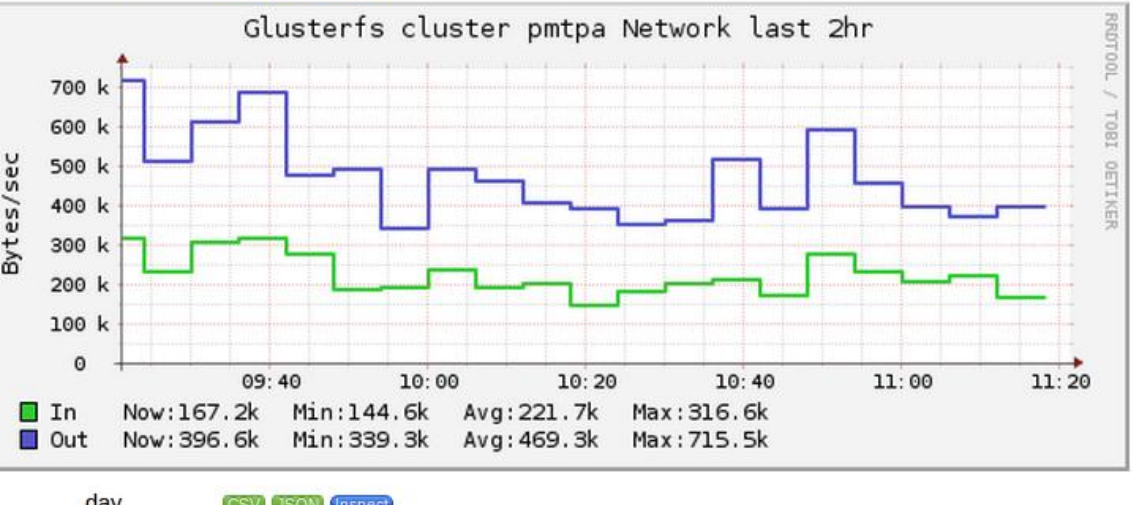


4hr CSV JSON Inspect

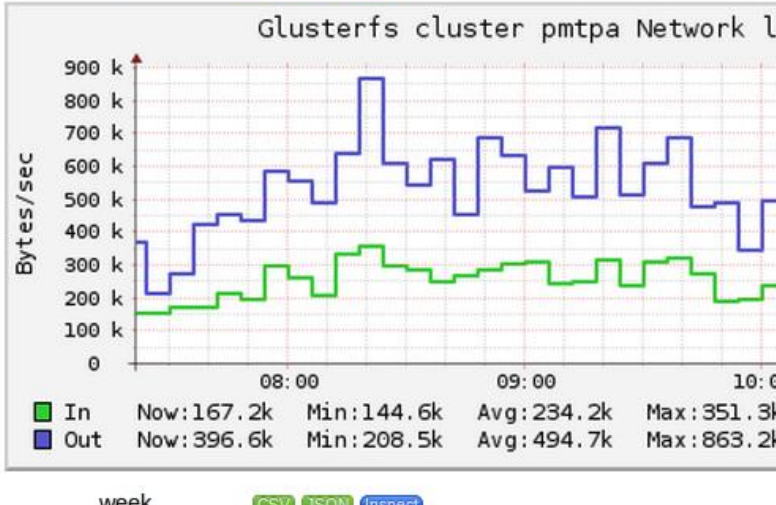


day CSV JSON Inspect

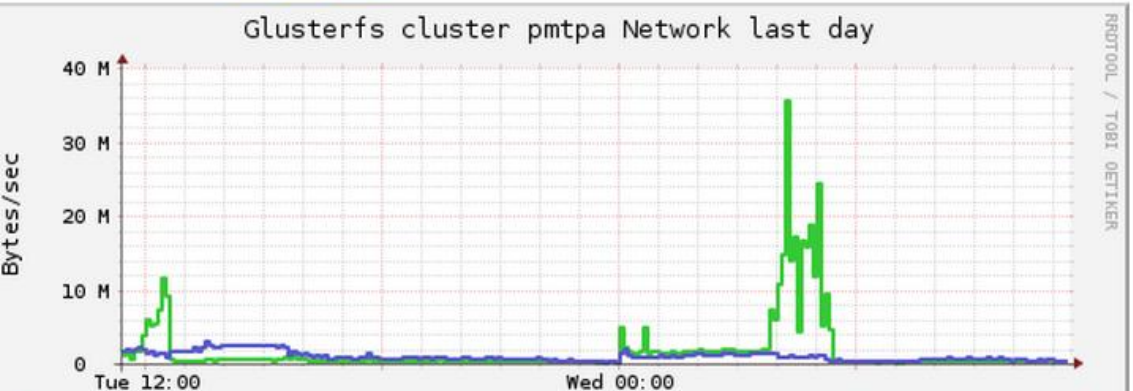


week CSV JSON Inspect

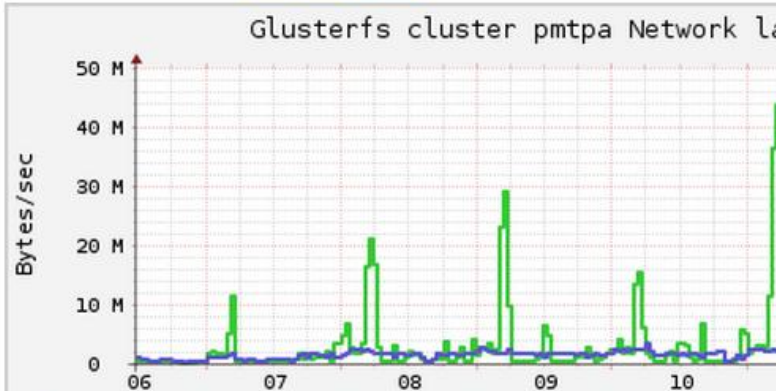

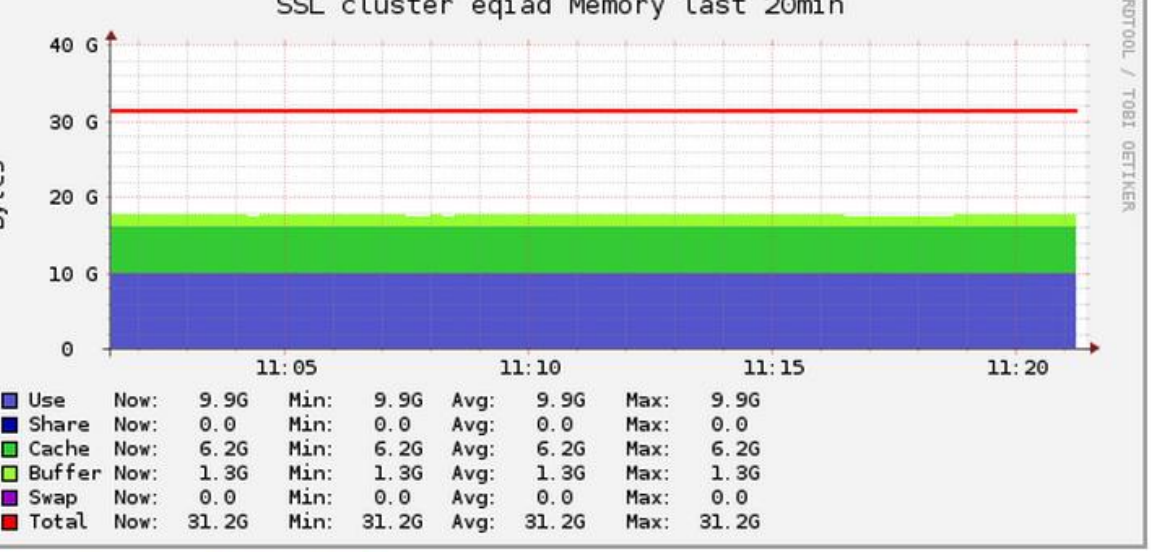


2hr CSV JSON Inspect


day CSV JSON Inspect

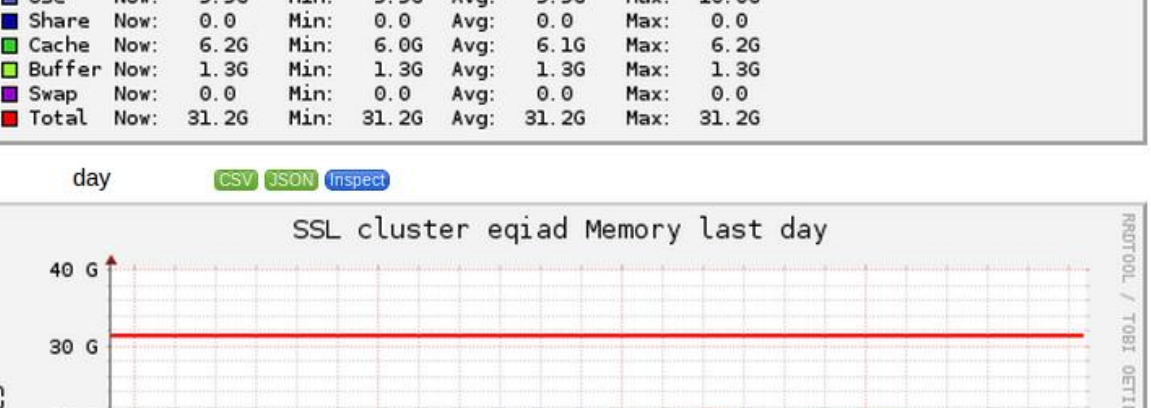


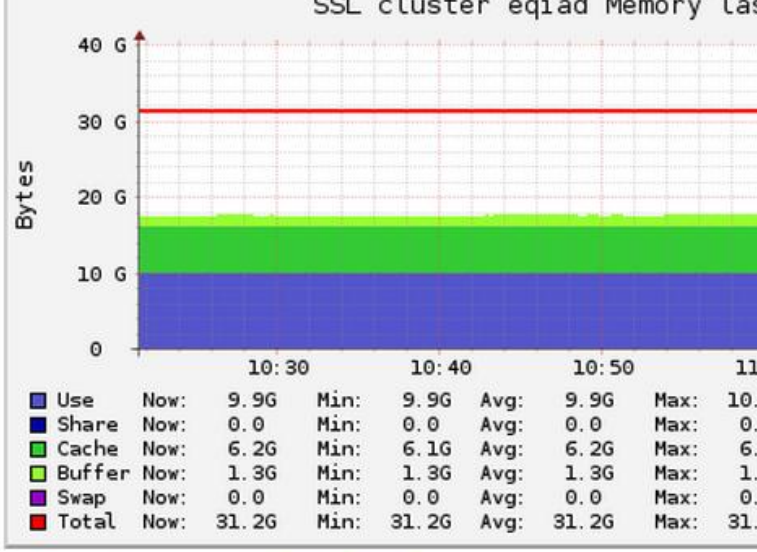


4hr CSV JSON Inspect

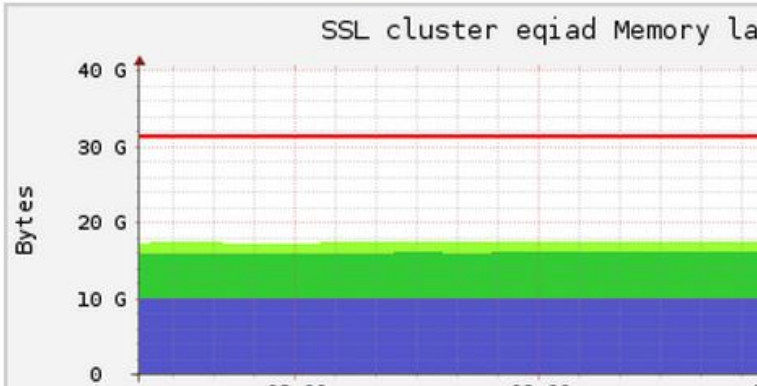


week CSV JSON Inspect

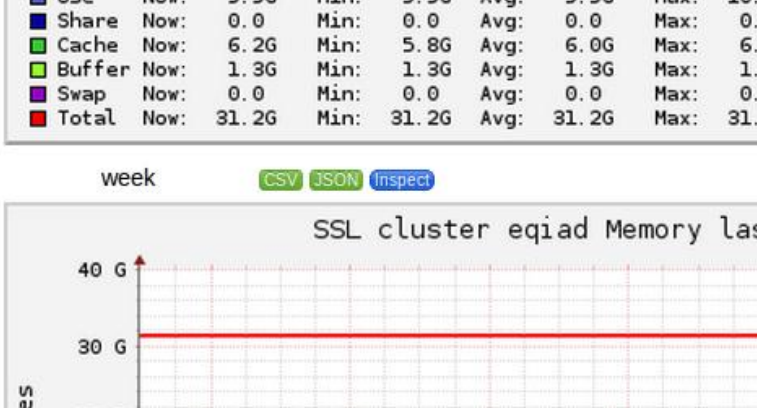


http://ganglia.wikimedia.org

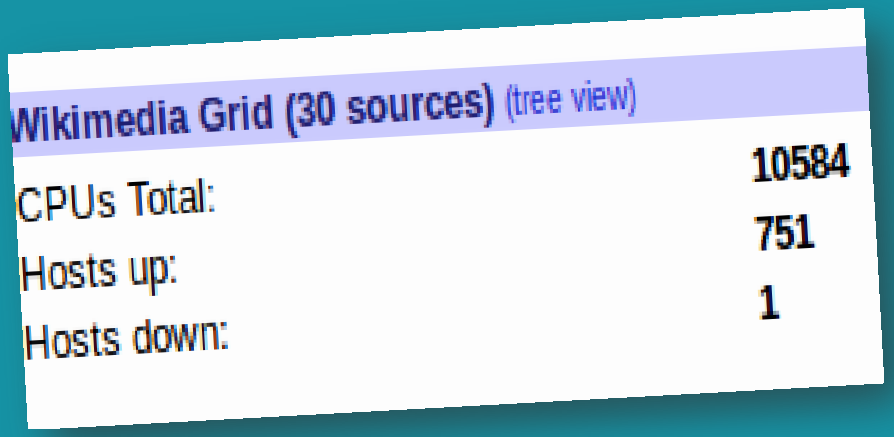


CPU, memory, load,
network,
disk, processes

751 hosts x 23 metrics =

# 17,273 graphs

Unstructured,
specific
&
Non Machine
Readable

Introduction

Current Monitoring Systems

Cloud Monitoring

How can we better understand large scale systems which frequently change?

# Decentralization & fault tolerance

Lessen failure,
Improve scale

# Location awareness

Knowledge of logical location within the cloud minimize costs, maximize performance

# Holistic

Provide a high level view of entire systems

# Autonomic

Self managing,
less human interaction,
faster adaptation


Decentralised
Semantic Data
Collection

# Structured p2p overlay

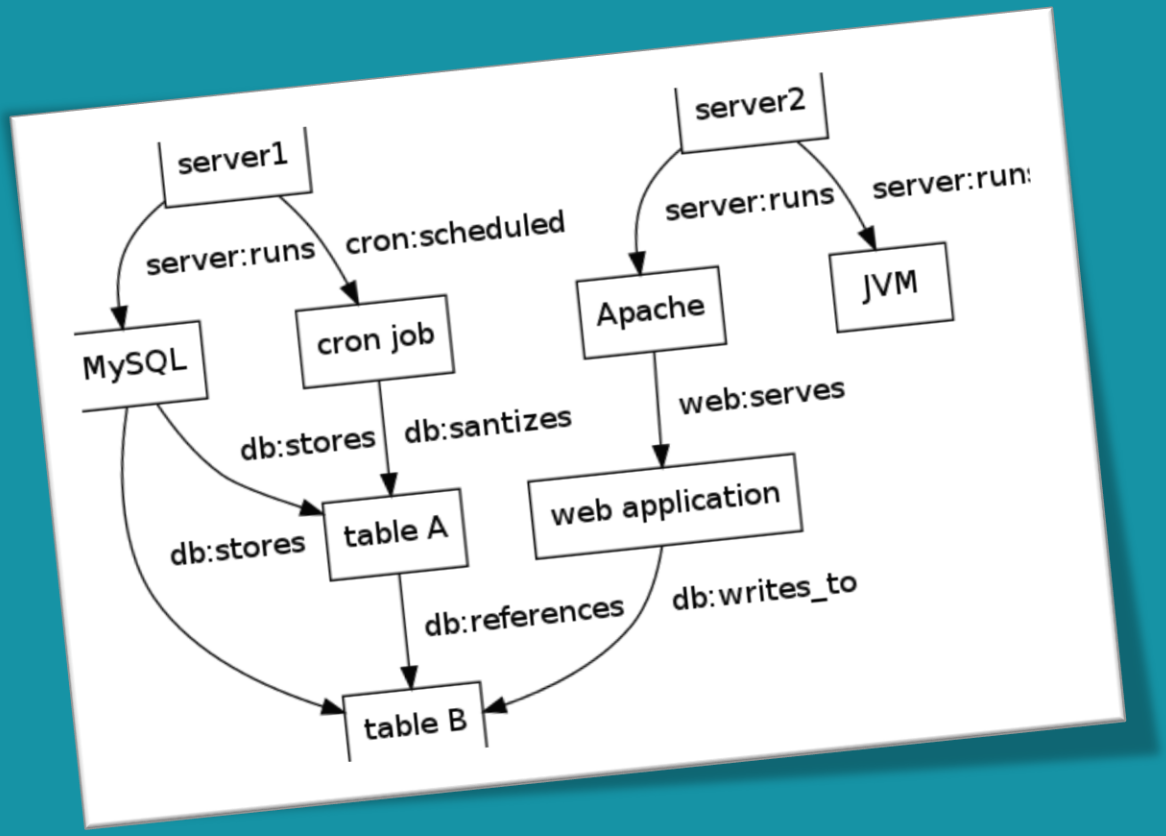


# Linked data model

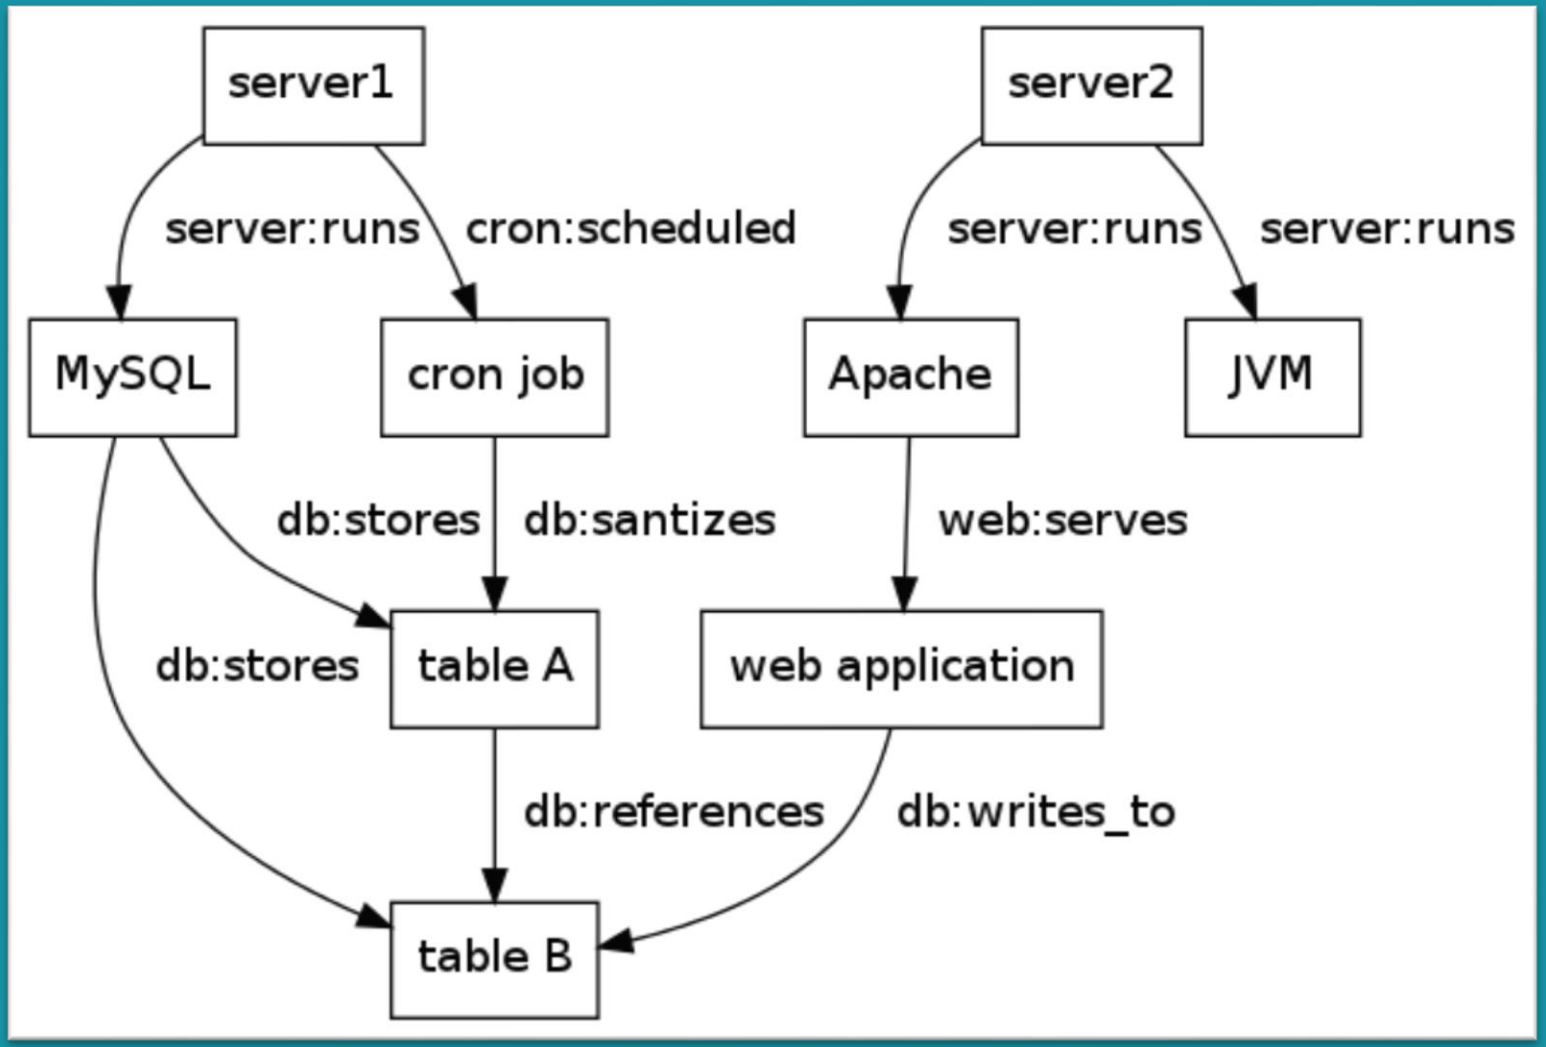
server1
server2
server:runs
cron:scheduled
server:runs
server:runs
MySQL
cron job
Apache
JVM
db:stores
db:santizes
web:serves
db:stores
table A
web application
db:references
db:writes_to
table B

Use the semantic model to restructure the monitoring system in order to adapt to change

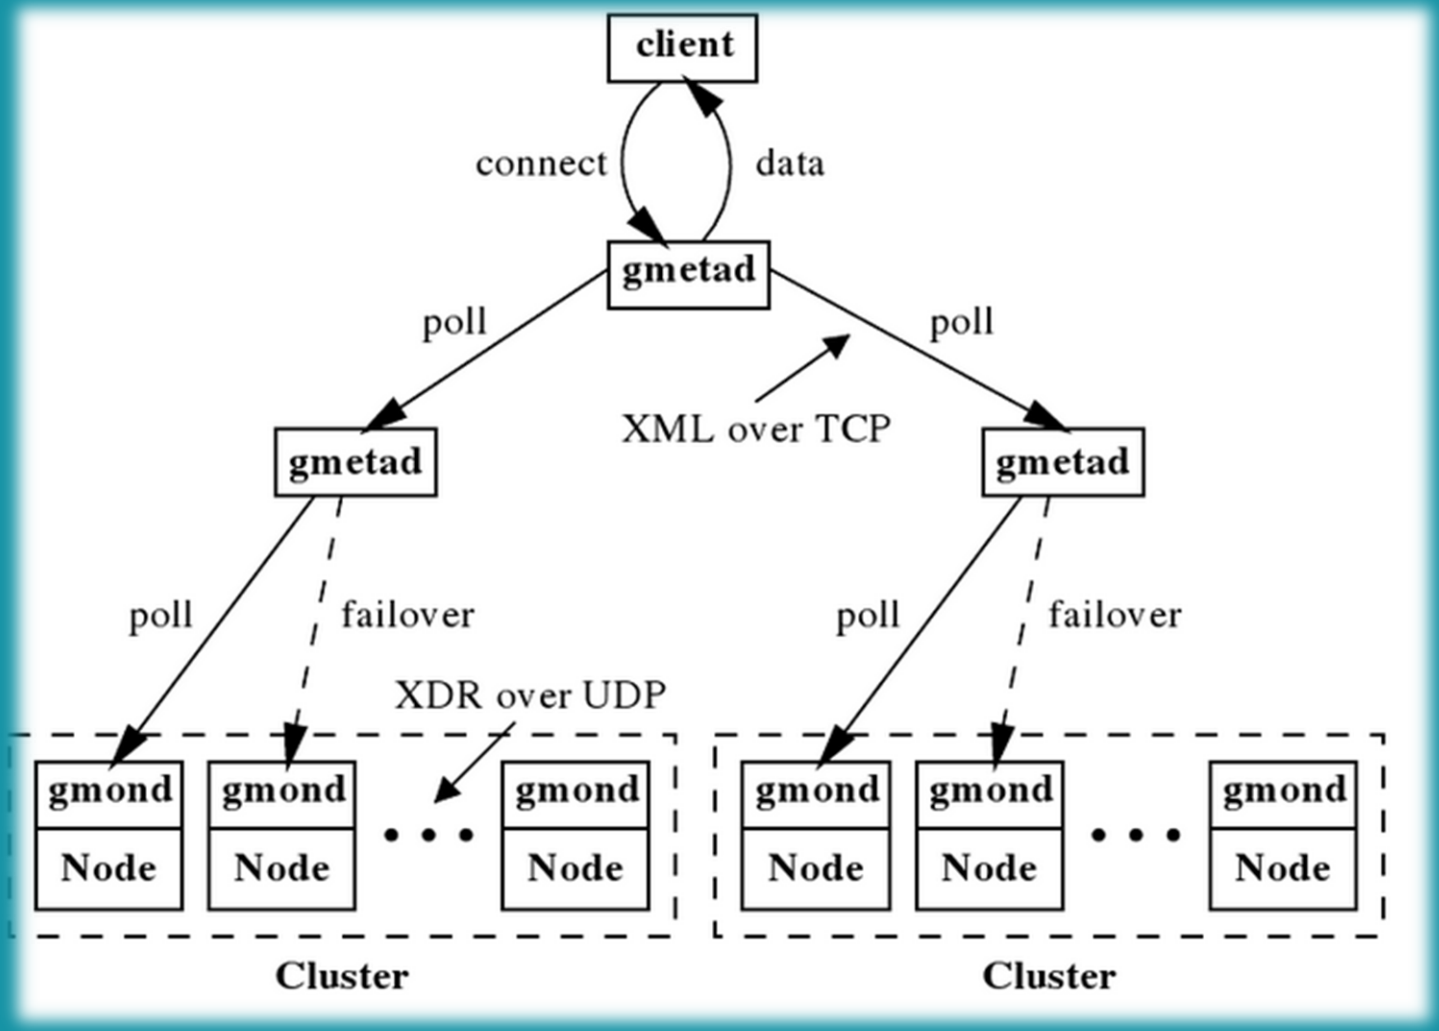


Massie et al, 2004

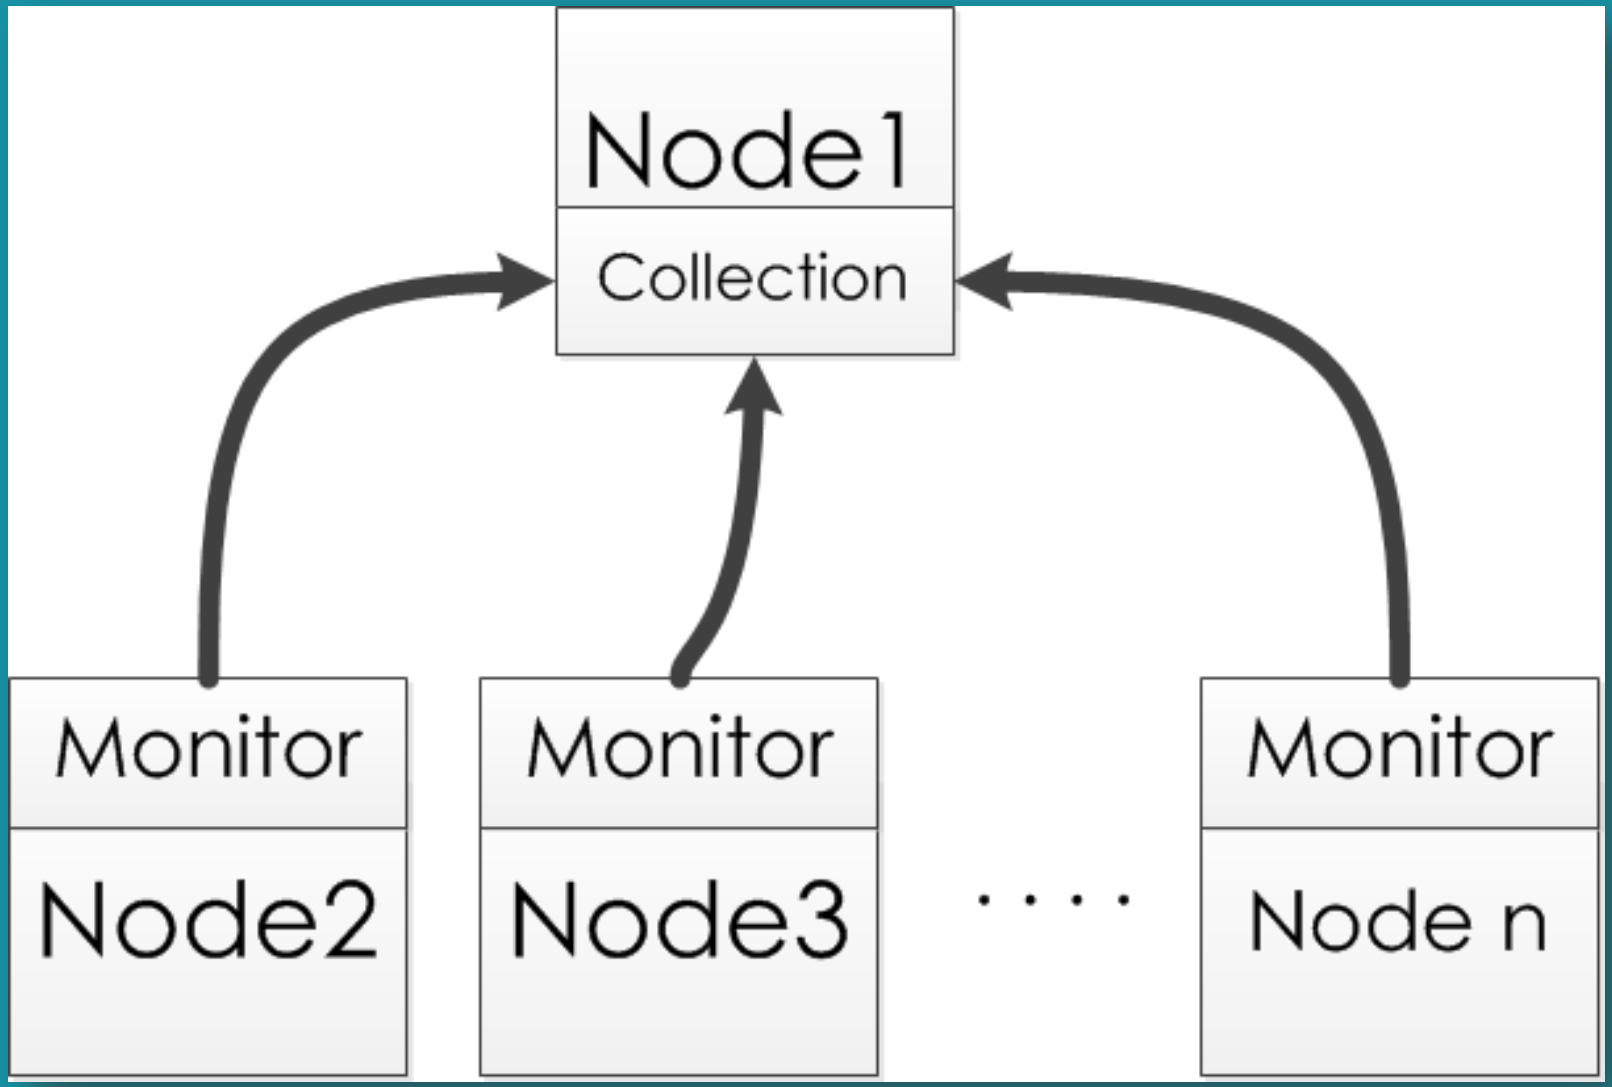
Node1
Collection
Monitor
Node2
Monitor
Node3
. . . .
Monitor
Node n

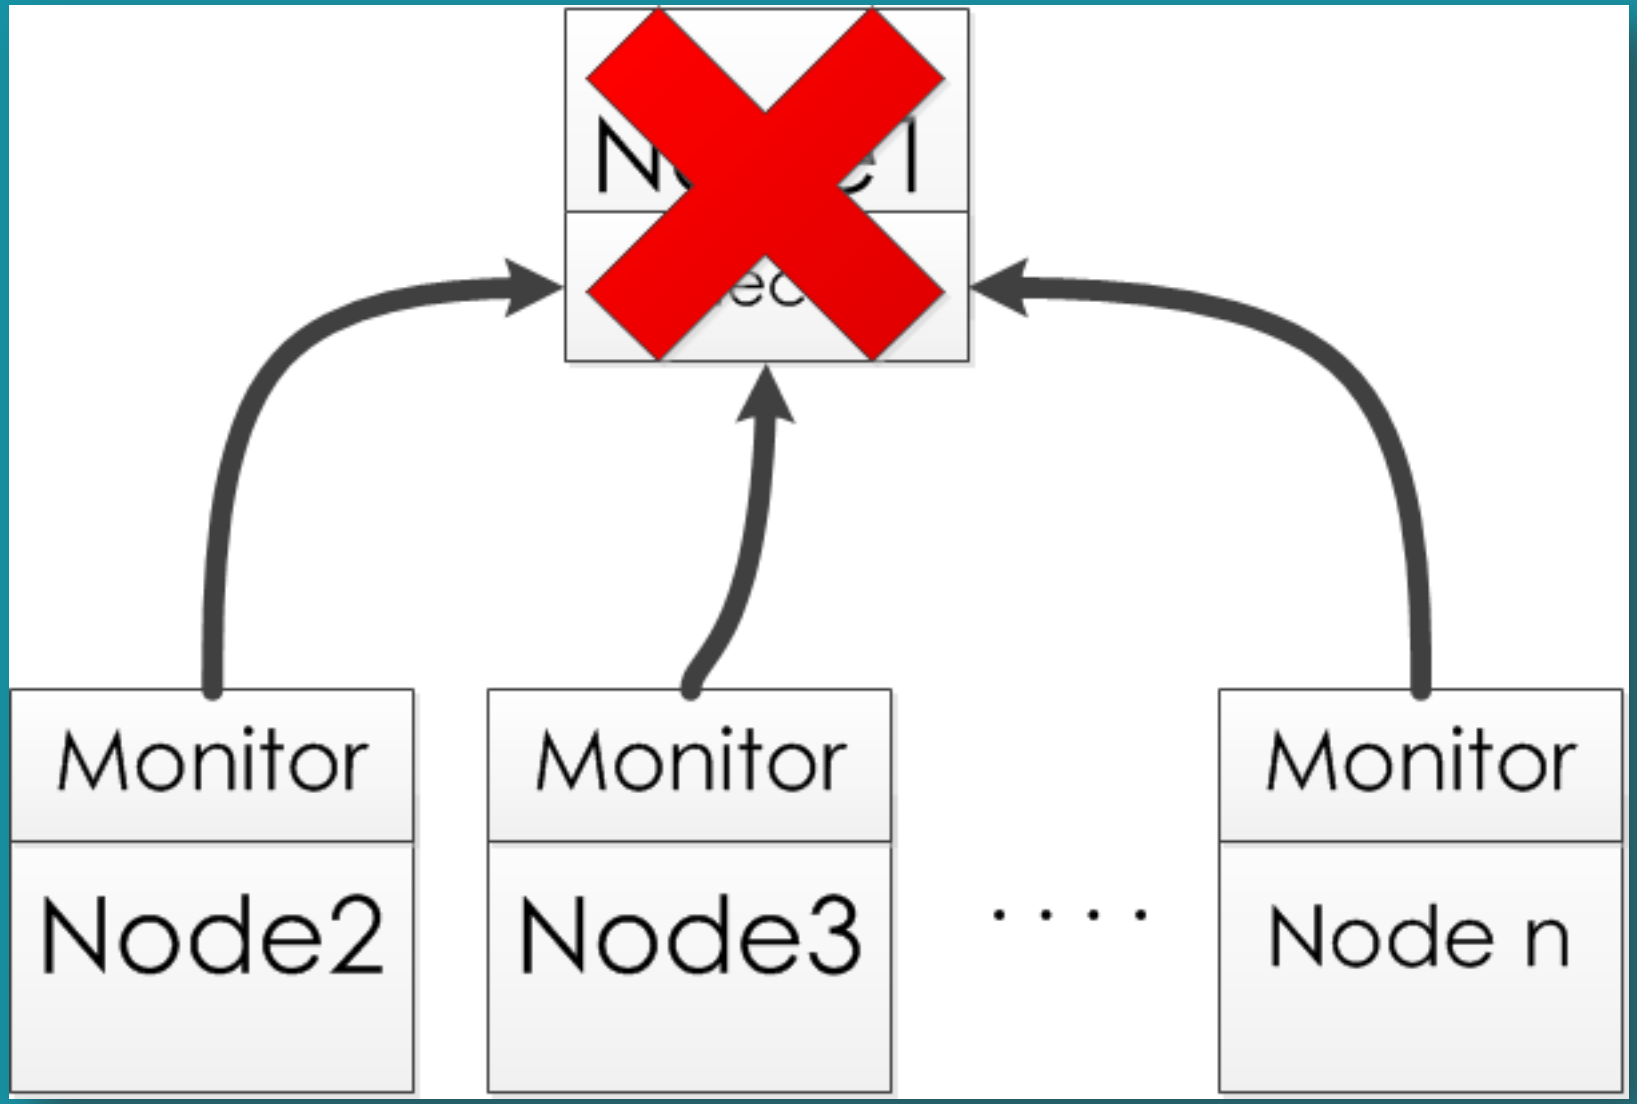
Monitor
Node2
Monitor
Node3
. . . .
Monitor
Node n

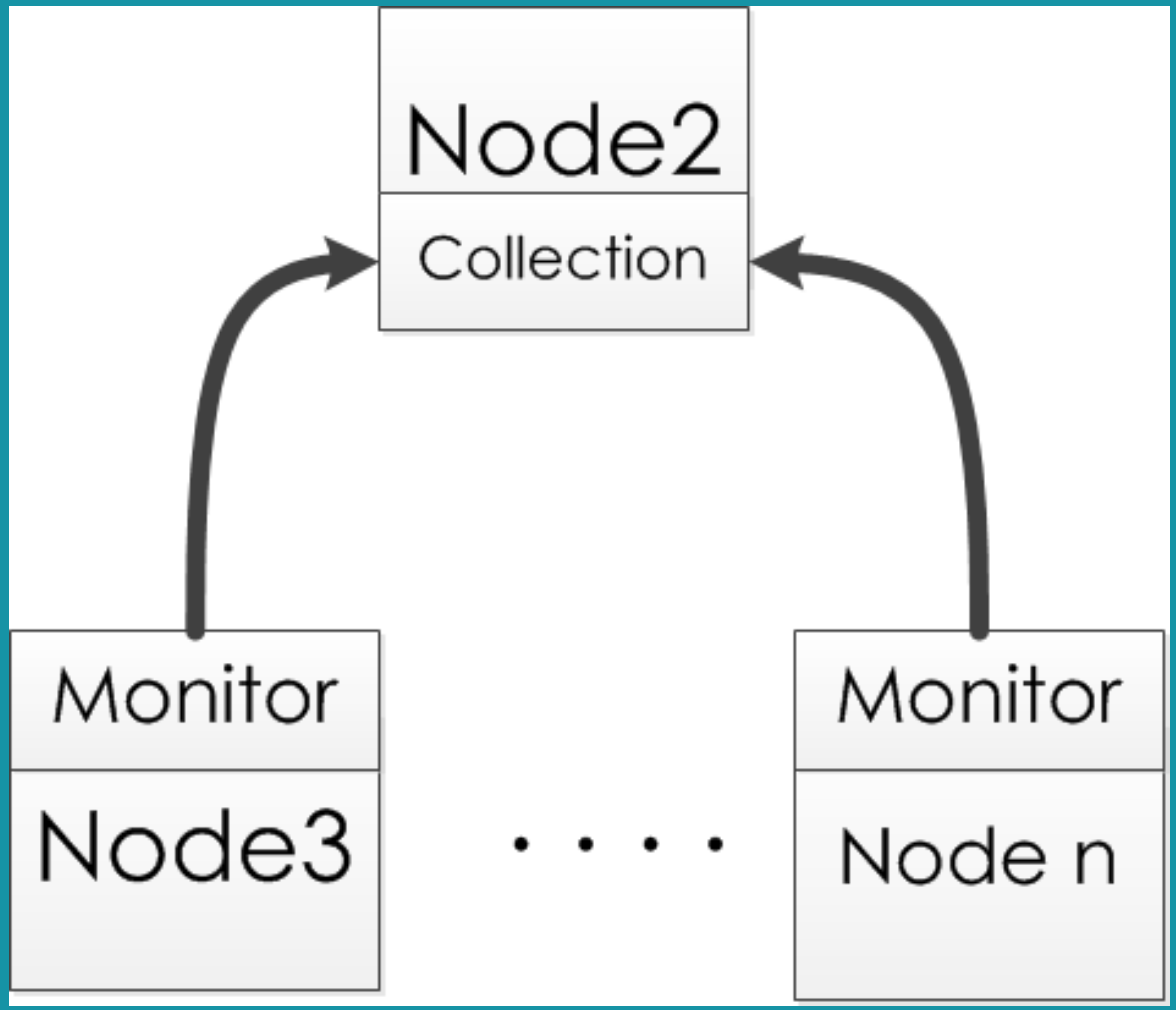
Node2
Collection
Monitor
Node3
Monitor
Node n

A basis for understanding and optimising large scale systems

Current monitoring is slow to adapt, scales poorly and structures information in a way that neither machines or humans can easily understand

Current monitoring is slow to adapt, scales poorly and struggles to make information in a way that neither machines nor humans can easily understand

**decentralised semantic data collection can overcome these difficulties**

**(hopefully)**

# Thank you, Questions?


**Jonathan Stuart Ward**
**School of Computer Science**
**Big Data Group,**
**University of St Andrews**

**jonathan.stuart.ward@st-andrews.ac.uk**
**@jauntyward**