

The challenge of external data

Martin Simons
6 februari 2024

Whoami?

Born a Millenium ago

Unofficial husband

Official father of three daughters

Love to drink beer

Debian

Get it right kind of guy

Self employed since 1998

Hit me on LinkedIn, Webhuis



Old Configuration Management

It all started in the nineties...

With bare metal machines

Having many combined roles on one machine

An ntp server was stuffed to a central machine:

- The name server

Bigger data centres to meet the Internet needs

The config management problem was born

New Configuration Management

Virtual Machines emerged

Containers

IOT

Small units having a simple role

There are **behavioral** aspects to actors in the field

Throwing everything away

So we do config management for 30+ years

- Tons of scripts have emerged
- Numerous tools come and go
- Where has all the information gone?
- Why has all the information gone?

There is high value in the information

Data / Information is more **persistent** than tools

A fresh start every day

It is natural and very human

There never is time to do it right,
But there is always time to do it again

Orchestration, The Holy Grail



Aspects of Orchestration

- There is a **plan**
- There is a **Desired State**

Orchestration

The conductor

- Has a Desired State in mind
- Conveys the Desired State he desires
- The different actors focus on their own role
- Actors do not mind other roles too much
- Works highly **imparative**
- Might even **enforce** synchronisation of moves
- Actors are highly dependend
- Is continuously sending messages
- Does he arrive at the **Desired State**

Orchestration, Span of Control



Limits of Orchestration

Trying to go beyond the limits is looking **funny**

- When do we reach the Desired State?
- To what extent do we have a desired state?
- Is it possible to arrive at the desired state?
- Is the Desired State persistent?
- No autonomy in the above examples
- Highly centralised
- Hardly any interaction between the actors

Orchestration?



Autonomous Actors – 1

Real World

- Actors and the Desired State(s)
- Can will still speak of a Desired State?
- How do we arrive at the desired state?
- Is the Desired State persistent?
- Enforcing becomes very difficult
- Autonomy vs dependency

Autonomous Actors – 2

Real World


- It looks messy, but it works if....
- Actors use a common infrastructure
- Actors go from a(n) to b(m)
- Actors are aware of other Actors
- Actors use standardised messages to interact
- **Intentions** are being **advertised**

Messaging

All sorts of messages involved

- Interaction
- Behavior
- Types of messages
- Context of messages
- Ways of managing messages

Classic broadcast message



Broadcasting aspects

- One way delivery
- TV Radio
- Flyering
- Message in a bottle
- Verified delivery
- Message brokering

Formal messaging system



MATT



'Actually, this is the back of the queue for Westminster Hall, in London'

Postal services aspects

It was the only way

- Upstream
- Processing
- Downstream

- The Postal service acts like a broker
- Single point of failure, strike

There are ways of signaling

Most of them broadcasting

- Lighthouse
- Telephone tree
- Smoke signals

Messaging in an IT landscape

Why would there be messages?

- The IT World gets more and more **dynamic**
- Numerous **different** types of devices
- Devices become more and more **intelligent**
- Devices use actual information in their **role**
- Devices need to adapt to change
- Many Devices need **dynamic configuration**

Dealing with information

Every Actor in the field provides information

- Processing information
- Generated **Role** based information
- Valuable for collection and distribution
- Actual and historic information
- Distributing sensitive data out of a database

Enter **Data**

How does Data work?

Data is an App and a PostgreSQL database.

The database has three schema's:

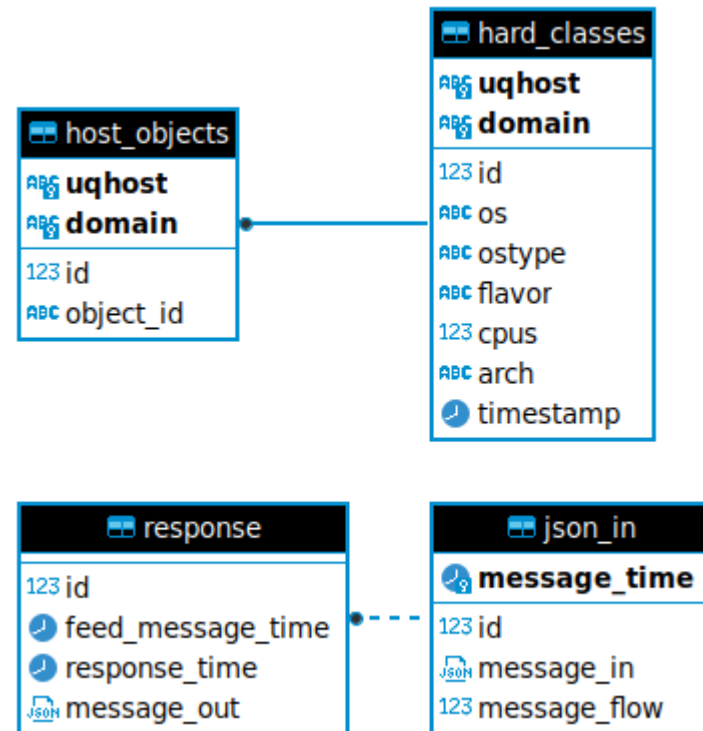
- Feeds
- Context
- Knowledge

The App is an object oriented Python3 program

- Three tier (I hope)
- Multi threading
- <https://github.com/Webhuis/Data/tree/master>

The nosql feed schema

Upstream schema feeds

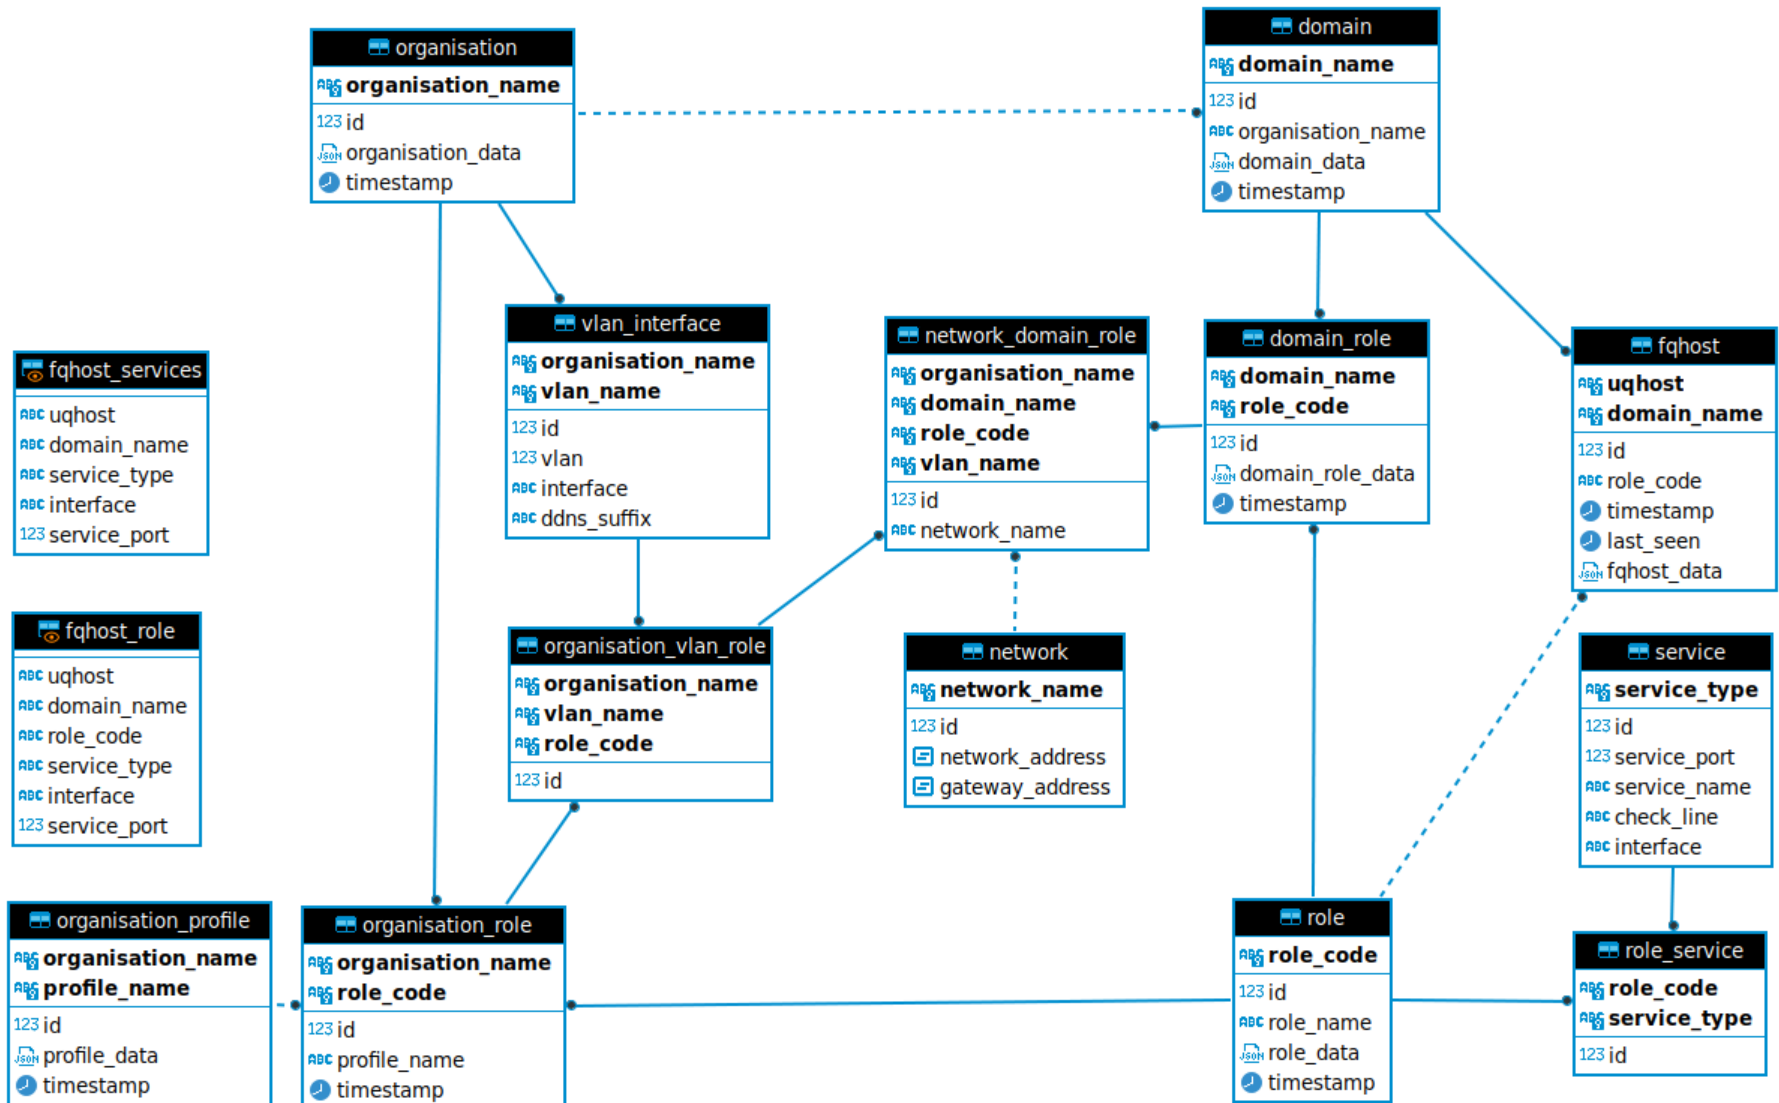


Making data available

So how becomes data made manageable?

- Make individual nodes **share** their data
- Every message is a (json) container
- Nodes send **feeds**
- Data processes feeds
- Nodes send queries
- Data gives taylorred **responses** to queries

The enriched data ER model



How does the complex work?

Data and the Actors exchange messages

- Using 0mq, zeromq or ZMQ
- Periodically
- Complex messages, having containers
- Using the JSON format
- Works for **every** 0mq and JSON **enabled** Actor
- May contain sensitive information

The communication is **not** yet encrypted

TCP/IP supports all

The protocol allows

- Broadcasting
- Verified delivery
- Brokered messaging
- **Brokerless** point to point messaging

Every Actor feeds Data

This is the hard_classes feed:

```
{"message": "feed",  
"query": "hard_classes",  
"content": {  
  "feeding_host": "mnmutl0001",  
  "domain": "webhuis.nl",  
  "os": "linux",  
  "ostype": "linux_x86_64",  
  "flavor": "debian_8",  
  "cpus": "1",  
  "arch": "x86_64" } }
```

It is fast!

Message in: 2024-02-05 15:07:47.493

Response out: 2024-02-05 15:07:47.507

- A grand total of **14 msec** (on the Data box)
- Carried out seven selects
- Did at least two inserts
- The response size is around 1.5 Kb

Example Nagios Monitor

The screenshot displays the Nagios Core web interface. The browser window title is "Nagios Core - Iceweasel" and the address bar shows "nagios.webhuis.nl/nagios3/". The interface includes a sidebar with navigation links for General, Current Status, Reports, and System. The main content area is divided into several sections:

- Current Network Status:** Last Updated: Sun Jan 31 23:27:34 GMT 2016. Updated every 90 seconds. Nagios® Core™ 3.5.1 - www.nagios.org. Logged in as CFEngine_Roadshow.
- Host Status Totals:** Up: 1, Down: 0, Unreachable: 0, Pending: 0. All Problems: 0, All Types: 1.
- Service Status Totals:** Ok: 6, Warning: 0, Unknown: 0, Critical: 0, Pending: 0. All Problems: 0, All Types: 6.
- Service Status Details For All Hosts:** A table showing details for localhost services.

Limit Results: 100

Host	Service	Status	Last Check	Duration	Attempt	Status Information
localhost	Current Load	OK	2016-01-31 23:23:21	35d 7h 48m 14s	1/4	OK - load average: 0.00, 0.04, 0.05
	Current Users	OK	2016-01-31 23:27:02	35d 7h 47m 24s	1/4	USERS OK - 2 users currently logged in
	Disk Space	OK	2016-01-31 23:23:21	35d 7h 46m 34s	1/4	DISK OK
	HTTP	OK	2016-01-31 23:26:51	15d 11h 31m 52s	1/4	HTTP OK: HTTP/1.1 200 OK - 11378 bytes in 0.003 second response time
	SSH	OK	2016-01-31 23:23:46	35d 7h 44m 54s	1/4	SSH OK - OpenSSH_6.7p1 Debian-5 (protocol 2.0)
	Total Processes	OK	2016-01-31 23:26:51	35d 7h 44m 4s	1/4	PROCS OK: 91 processes

Results 1 - 6 of 6 Matching Services

Actors feed their details

It is a dynamic setup

- Every time an actor comes by
- Data registers details
- The monitor gets the relevant details
- The Actors **autonomously** process messages

Hosts are added

The screenshot shows the Nagios Core web interface in a browser window titled "Nagios Core - Iceweasel". The address bar shows "nagios.webhuis.nl/nagios3/". The interface includes a sidebar with navigation menus for General, Current Status, Problems, Reports, and System. The main content area displays "Current Network Status" (last updated Sun Jan 31 23:36:50 GMT 2016), "Host Status Totals" (2 Up, 0 Down, 0 Unreachable, 5 Pending), and "Service Status Totals" (6 Ok, 0 Warning, 0 Unknown, 0 Critical, 7 Pending). Below these is a table titled "Host Status Details For All Host Groups" with columns for Host, Status, Last Check, Duration, and Status Information. The table lists 7 hosts, with 1 UP and 6 PENDING.

Current Network Status
Last Updated: Sun Jan 31 23:36:50 GMT 2016
Updated every 90 seconds
Nagios® Core™ 3.5.1 - www.nagios.org
Logged in as CFEngine_Roadshow

Host Status Totals

Up	Down	Unreachable	Pending
2	0	0	5

All Problems All Types

0	7
---	---

Service Status Totals

Ok	Warning	Unknown	Critical	Pending
6	0	0	0	7

All Problems All Types

0	13
---	----

Host Status Details For All Host Groups

Limit Results: 100

Host	Status	Last Check	Duration	Status Information
apjums0001.webhuis.nl	UP	2016-01-31 23:36:33	0d 0h 0m 17s+	PING OK - Packet loss = 0%, RTA = 3.61 ms
cfeutl0001.webhuis.nl	PENDING	N/A	0d 0h 0m 17s+	Host check scheduled for Sun Jan 31 23:37:15 GMT 2016
dtautl0001.webhuis.nl	PENDING	N/A	0d 0h 0m 17s+	Host check scheduled for Sun Jan 31 23:37:58 GMT 2016
localhost	UP	2016-01-31 23:34:12	35d 7h 57m 30s	PING OK - Packet loss = 0%, RTA = 0.06 ms
mnmultl0001.webhuis.nl	PENDING	N/A	0d 0h 0m 17s+	Host check scheduled for Sun Jan 31 23:38:41 GMT 2016
webapp0180.webhuis.nl	PENDING	N/A	0d 0h 0m 17s+	Host check scheduled for Sun Jan 31 23:39:24 GMT 2016
webjen0001.webhuis.nl	PENDING	N/A	0d 0h 0m 17s+	Host check scheduled for Sun Jan 31 23:40:07 GMT 2016

Results 1 - 7 of 7 Matching Hosts

Response from Data

```
{
  "pdnrpcsw001.sw.webhuis.nl": [
    {
      "fqhost_data": "",
      "Webhuis": [
        {
          "organisation_data": [
            {
              "dhcp_ldap": "mihrwpcsw1.sw.webhuis.nl",
              "administrator": "martin@webhuis.nl",
              "ldap_frontend": "ldap.sw.webhuis.nl",
              "ldap_m": [
                [
                  "100",
                  "ldappdcsw001.sw.webhuis.nl",
                  "101",
                  "ldappdcsw002.sw.webhuis.nl"
                ]
              ],
              "primary_dns": [
                {
                  "mdns": "powerdns.webhuis.nl"
                }
              ],
              "mdns_key": "/etc/dhcp/webhuis.nl.update"
            }
          ],
          "local_networks": [
            "10.0.0.0/8",
            "192.168.0.0/16",
            "213.127.130.120/29"
          ],
          "ttl": "300",
          "ntp_server": "ntp.time.nl"
        }
      ]
    },
    {
      "domain_data": [
        {
          "ldap_base": "dc=sw,dc=webhuis,dc=nl",
          "local_resolvers": [
            "10.68.171.24",
            "10.10.117.24",
            "10.10.217.24"
          ],
          "searchlist": "sw.webhuis.nl"
        }
      ],
      "role_data": "",
      "domain_role_data": "",
      "domain_role_network": [
        [
          {
            "vlan_name": "admin",
            "vlan": 22,
            "interface": "eth2",
            "ddns_suffix": "-ssh",
            "network_name": "admin_sw",
            "network_address": "10.168.71.0/24",
            "gateway_address": "10.168.71.0"
          },
          {
            "vlan_name": "support",
            "vlan": 123,
            "interface": "eth1",
            "ddns_suffix": "",
            "network_name": "support_sw",
            "network_address": "10.68.171.0/24",
            "gateway_address": "10.68.171.0"
          }
        ]
      ],
      "organisation_profile": [
        [
          "Webhuis",
          "pdnr",
          "std_sup",
          {
            "organisation_profile": [
              {
                "etc_hosts_nic": "support",
                "admin_nic": "admin",
                "nic_trick": [
                  "0",
                  "1"
                ]
              }
            ]
          }
        ]
      ],
      "pdnr": [
        [
          {
            "service_port": 22,
            "service_name": "sshd",
            "check_line": "check_tcp\\!22",
            "interface": "eth2"
          },
          {
            "service_port": 53,
            "service_name": "Powerdns server",
            "check_line": "check_tcp\\!53",
            "interface": "eth1"
          }
        ]
      ]
    }
  ]
}
```

Services of hosts join in

The screenshot shows the Nagios Core web interface. The browser address bar displays 'nagios.webhuis.nl/nagios3/'. The interface includes a sidebar with navigation links for General, Current Status, Reports, and System. The main content area displays 'Current Network Status' with a last update of Sun Jan 31 23:38:21 GMT 2016. It also shows 'Host Status Totals' (4 Up, 0 Down, 0 Unreachable, 3 Pending) and 'Service Status Totals' (6 Ok, 0 Warning, 0 Unknown, 0 Critical, 7 Pending). Below this is a table titled 'Host Status Details For All Host Groups' with columns for Host, Status, Last Check, Duration, and Status Information. The table lists several hosts, including 'localhost' which is UP, and three hosts in a PENDING state.

Current Network Status
Last Updated: Sun Jan 31 23:38:21 GMT 2016
Updated every 90 seconds
Nagios® Core™ 3.5.1 - www.nagios.org
Logged in as *CFEngine_Roadshow*

Host Status Totals

Up	Down	Unreachable	Pending
4	0	0	3

All Problems All Types

All Problems	All Types
0	7

Service Status Totals

Ok	Warning	Unknown	Critical	Pending
6	0	0	0	7

All Problems All Types

All Problems	All Types
0	13

Host Status Details For All Host Groups

Limit Results:

Host	Status	Last Check	Duration	Status Information
apjums0001.webhuis.nl	UP	2016-01-31 23:36:33	0d 0h 1m 48s+	PING OK - Packet loss = 0%, RTA = 3.61 ms
cfeutl0001.webhuis.nl	UP	2016-01-31 23:37:15	0d 0h 1m 48s+	PING OK - Packet loss = 0%, RTA = 0.52 ms
dtautl0001.webhuis.nl	UP	2016-01-31 23:37:58	0d 0h 1m 48s+	PING OK - Packet loss = 0%, RTA = 0.33 ms
localhost	UP	2016-01-31 23:34:12	35d 7h 59m 1s	PING OK - Packet loss = 0%, RTA = 0.06 ms
mnmutl0001.webhuis.nl	PENDING	N/A	0d 0h 1m 48s+	Host check scheduled for Sun Jan 31 23:38:41 GMT 2016
webapp0180.webhuis.nl	PENDING	N/A	0d 0h 1m 48s+	Host check scheduled for Sun Jan 31 23:39:24 GMT 2016
webjen0001.webhuis.nl	PENDING	N/A	0d 0h 1m 48s+	Host check scheduled for Sun Jan 31 23:40:07 GMT 2016

Results 1 - 7 of 7 Matching Hosts

Do I smell a desired state somewhere?

The screenshot shows the Nagios Core web interface. The browser address bar displays 'nagios.webhuis.nl/nagios3/'. The interface includes a sidebar with navigation menus for General, Current Status, Reports, and System. The main content area features summary statistics for Host Status Totals and Service Status Totals, and a detailed table for Host Status Details For All Host Groups. The table lists 7 hosts, with one host, 'webapp0180.webhuis.nl', in a CRITICAL state (DOWN) due to being unreachable.

Current Network Status
Last Updated: Sun Jan 31 23:44:25 GMT 2016
Updated every 90 seconds
Nagios® Core™ 3.5.1 - www.nagios.org
Logged in as CFEngine_Roadshow

Host Status Totals

Up	Down	Unreachable	Pending
6	1	0	0

Service Status Totals

Ok	Warning	Unknown	Critical	Pending
11	0	0	2	0

Host Status Details For All Host Groups

Limit Results: 100

Host	Status	Last Check	Duration	Status Information
apjums0001.webhuis.nl	UP	2016-01-31 23:41:43	0d 0h 5m 11s	PING OK - Packet loss = 0%, RTA = 4.39 ms
cfeutl0001.webhuis.nl	UP	2016-01-31 23:42:23	0d 0h 2m 29s	PING OK - Packet loss = 0%, RTA = 0.45 ms
dtautl0001.webhuis.nl	UP	2016-01-31 23:43:18	0d 0h 4m 48s	PING OK - Packet loss = 0%, RTA = 2.00 ms
localhost	UP	2016-01-31 23:39:17	35d 8h 5m 5s	PING OK - Packet loss = 0%, RTA = 0.06 ms
mnmutl0001.webhuis.nl	UP	2016-01-31 23:43:43	0d 0h 4m 25s	PING OK - Packet loss = 0%, RTA = 0.05 ms
webapp0180.webhuis.nl	DOWN	2016-01-31 23:44:03	0d 0h 0m 17s	CRITICAL - Host Unreachable (webapp0180.webhuis.nl)
webjen0001.webhuis.nl	UP	2016-01-31 23:40:07	0d 0h 4m 2s	PING OK - Packet loss = 0%, RTA = 0.61 ms

Results 1 - 7 of 7 Matching Hosts

Convergence

Actors

- Will exchange information periodically
- **Convergently** work towards the Desired state

Message timeouts

The duration of a message exchange

- The response times are normally distributed
- Timeout .05 seconds, on a local connection
- Important to configure to prevent overload

Use cases

When an Actor needs actual information

- Monitoring servers in a dynamic landscape
- Certificate distribution
- About other Actors that are around

Use cases

When you need information

- Performance analysis
- Capacity analysis
- Historic information
- Big Data analysis

Conclusion?

Keep the gold to yourself,
Stop throwing it away

Discussion