




Bla Bla Car

How we monitor **1 billion** km
of monthly ride sharing

Jean Baptiste Favre
Ops Lead


@jbfavre

Bla Bla Car

How are we ?

5 million

members in december 2013




A brass telescope is shown diagonally across the frame. A dark blue semi-transparent rectangular overlay is positioned in the center, containing the text "20 million members monitoring" in white. The telescope's barrel is brass, and its lens is visible at the bottom right.

20 million members
monitoring

7 million

members in april 2014






50 million members
monitoring

20 million

members in april 2015





2015



100 million members
monitoring

The background is a blurred image of a road with cars. Overlaid on the image are several digital elements: a large white speech bubble at the top containing the text 'Bla Bla Car'; a semi-transparent dark blue banner across the middle with white text; and various futuristic digital displays and a speedometer in the lower half. The displays show times like '02:12:34' and '01:03', and a speedometer shows '108 km/h'.

Bla Bla Car

How we monitor **1 billion** km
of monthly ride sharing

Zabbix



KEEP CALM
AND
MONITOR
ALL
THE THINGS

How many items ?



How many new VPS ?



Load ? What load ? :)

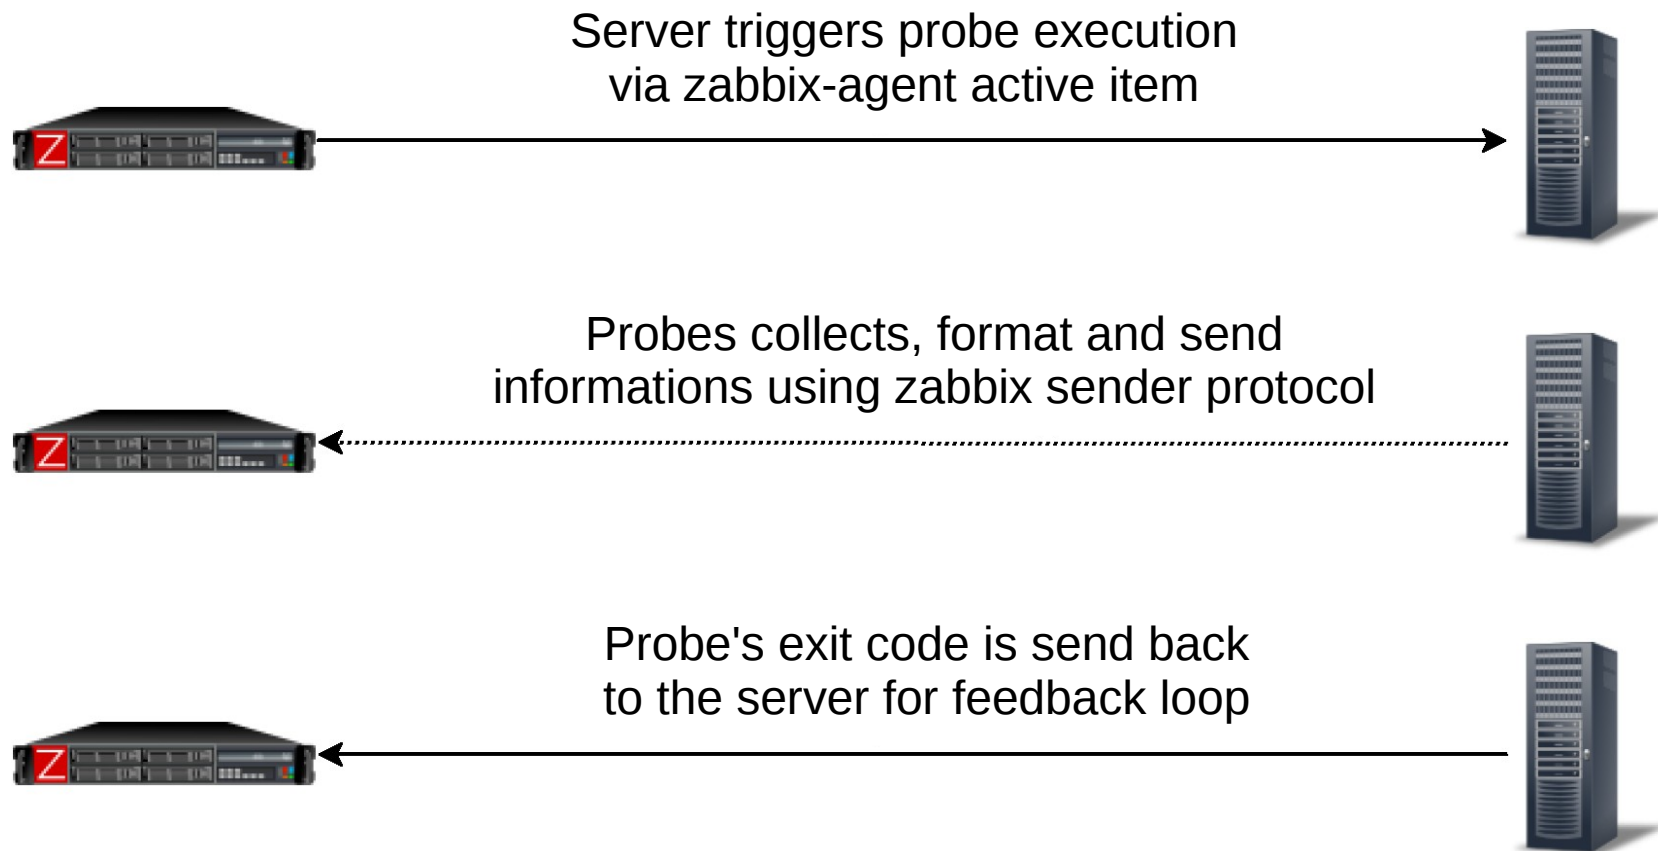


How ?

Standardization



Standardization



Exit codes

Standard :

0 => OK

1 => fail during init

2 => fail while getting informations

3 => fail during Container update

4 => fail during Send phase

Client side probes

Python or Java

LLD wherever possible
trappers always

Only 2 zabbix-agent (active) items per template



python-protobix



Almost



python-protobix



<https://github.com/jbfavre/python-protobix>
(also on pypi.python.org)

LLD example

```
#!/usr/bin/env python
```

```
import protobix
```

```
''' create DataContainer, providing data_type, zabbix server and port '''  
zbx_container = protobix.DataContainer('lld', 'localhost', 10051)
```

```
hostname='myhost'  
item='hardware.power_supply'  
value=[  
    { '#SLOT': 0, '#PLUGGED' : 1 },  
    { '#SLOT': 1, '#PLUGGED' : 0 },  
]  
zbx_container.add_item( hostname, item, value)  
  
try:  
    zbx_response = zbx_container.send()  
except protobix.SenderException:  
    print 'Oups...'
```

PUT YOUR OWN LOGIC HERE :)

item example

```
#!/usr/bin/env python
```

```
import protobix
```

```
''' create DataContainer, providing data_type, zabbix server and port '''  
zbx_container = protobix.DataContainer('items', 'localhost', 10051)
```

```
hostname='myhost'  
item='hardware.power_supply[0,status]'  
value=1  
zbx_container.add_item( hostname, item, value)
```

```
try:  
    zbx_response = zbx_container.send()  
except protobix.SenderException:  
    print 'Oups...'
```

PUT YOUR OWN LOGIC HERE :)

RabbitMQ example

Low Level Discovery

vhosts & queues

thresholds

Update values

message number

in/out ratio

Who is master of this queue

☐	Name
☐	DL Message count for {#RMQQUEUEUENAME} in RabbitMQ vhost {#RMQVHOSTNAME}
☐	Message count for {#RMQQUEUEUENAME} in RabbitMQ vhost {#RMQVHOSTNAME}
☐	Master node for {#RMQQUEUEUENAME} in RabbitMQ vhost {#RMQVHOSTNAME}
☐	ACK message rate for {#RMQQUEUEUENAME} in RabbitMQ vhost {#RMQVHOSTNAME}
☐	DELIVER_GET message rate for {#RMQQUEUEUENAME} in RabbitMQ vhost {#RMQVHOSTNAME}
☐	GET message rate for {#RMQQUEUEUENAME} in RabbitMQ vhost {#RMQVHOSTNAME}
☐	PUBLISH message rate for {#RMQQUEUEUENAME} in RabbitMQ vhost {#RMQVHOSTNAME}
☐	REDELIVER message rate for {#RMQQUEUEUENAME} in RabbitMQ vhost {#RMQVHOSTNAME}

MariaDB example

Low Level Discovery

Galera

storage engines

Multi-replication

Update values

Pretty much everything:)

☐	<u>Name</u> 
☐	MySQL Aria discovery
☐	MySQL Cassandra discovery
☐	MySQL Galera discovery
☐	MySQL InnoDB discovery
☒	MySQL replication sources discovery
☐	MySQL Server Audit discovery
☐	MySQL Sphinx discovery
☐	MySQL Spider discovery
☐	MySQL TokuDB discovery

Protobix probes

16 probes available

And more to come

`redis/dynomite`

`zookeeper`

...



<https://github.com/jbfavre/python-zabbix>

jmx-zabbix



KEEP CALM
AND
MONITOR
ALL THE
JAVA
THINGS

jmx-zabbix

Because python is not (always) enough :)



<https://github.com/n0rad/jmx-zabbix>

jmx-zabbix

Embedded inside a Java process

- Internal Java daemons

Aside any Java process (separate service)

- Cassandra
- Elasticsearch
- ...

configuration

```
serverName: <hostname in Zabbix>
```

```
pushIntervalSecond: 60
```

```
inMemoryMaxQueueSize: 10
```

```
zabbix:
```

```
  host: <Zabbix server hostname or IP>
```

```
  port: 10051
```

```
jmx:
```

```
  url: service:jmx:rmi:///jndi/rmi://localhost:7199/jmxrmi
```

```
  username: zabbix
```

```
  password: zabbix
```

```
  timeoutSecond: 30
```

```
[...]
```

JMX to ZBX mapping

[...]

```
metrics:
```

```
cassandra.status.failure:org.apache.cassandra.net:type=FailureDetector
cassandra.status.timeouts:org.apache.cassandra.net:type=MessagingService
cassandra.db.storage: org.apache.cassandra.db:type=StorageProxy
```

valuesCaptured:

[illegible]

Zabbix visualization




KEEP CALM
AND
LOOK
AT THE
GRAPHS

Grafana

Grafana

Grafana
+
Zabbix datasource
=
10 dashboards
in 2 days



<https://github.com/grafana/grafana>
<https://github.com/alexanderzobnin/grafana-zabbix>



MySQL - [dropdown]




Zoom Out


an hour ago to a few seconds ago refreshed every 30s

\$group: MariaDB Servers \$host: [dropdown]


Replication Lag




Connected/Running Threads



Queries




Bytes Sent/Received




- Replication Lag




- Connected/Running Threads



- Queries



- Bytes Sent/Received





RabbitMQ Messages count



Zoom Out

an hour ago to a few seconds ago refreshed every 1m

group: RabbitMQ Servers


host:

queue:


Message count

GLOBAL VIEW


RabbitMQ Message count in vhost




RabbitMQ Message count in vhost




RabbitMQ Message count in vhost



RabbitMQ Message count in vhost




RabbitMQ Message count in vhost



QUEUE DETAILED VIEW

RabbitMQ activity for (+) and DL (-)



RabbitMQ activity for



2015-08-20 15:05:57

ACK message rate for	4
DELIVER_GET message rate	4
GET message rate	4
PUBLISH message rate	0
REDELIVER message rate	0

Dashing

warning

22

Last updated at 17:10

average

3

Last updated at 17:10

high

1

Last updated at 17:10

disaster

0

Last updated at 17:10

<https://gist.github.com/chojayr/7401426>
<https://github.com/tolleiv/dashing-zabbix>

Caveats



KEEP CALM
AND
FIX THINGS
BEFORE
CTO
NOTICES

Beware of

Plugins & templates synchronization

Zabbix configuration automatization

Use same hostname everywhere

What next ?



KEEP CALM
AND
WAIT
FOR
ZABBIX 3.0

What I miss in Zabbix

Announced

- Trends predictions
- More scalable backend
- SSL communications

Not announced (As far as I know)

- Trends from
- Implicit dependency against proxy
- Detailed web scenario
- Per item maintenance
- Anomaly detection

3 Take aways

Now you can wake up :)

1. Define & use standards

2. Use LLD & Trappers

3. Visualization is critical

Let's discuss all that !

