



Migration to Zabbix 3.2 & Galera Cluster in large Banking sector environment

Xavier SCHNEIDER
September 16th 2017



SUMMARY

1. NATIXIS AT A GLANCE

2. ZABBIX IN NATIXIS

**3. MIGRATION METHODOLOGY TO ZABBIX 3.2 IN
LARGE ENVIRONMENT**

4. GALERA ARCHITECTURE IN NATIXIS

5. GALERA RUNTIME EXPERIENCE

6. QUESTIONS



NATIXIS AT A GLANCE



1 | NATIXIS AT A GLANCE

Natixis is part of BPCE Group, the second largest banking Group in France



1

NATIXIS AT A GLANCE

Natixis is the international corporate, investment, insurance and financial services arm of Groupe



Natixis is listed on Euronext Paris
and is included in CAC NEXT 20 index

- +16 000 employees
- + 36 countries
- € 8,7bn NET REVENUES

Our distinctive features



IT : Our expertise at the service
of our clients



1

NATIXIS IT AT A GLANCE



Natixis IT

3 300 people

24/7 support



Technology



Scope



Value to Dev

Cloud

From an infrastructure provider to a cloud broker

Devops

Be an infrastructure resources broker rather than a technical provider.

Attractive projects



+ 7 200 physical hosts (+100 000 cores)
+ 10 600 Vm hosts (+34000 cores)

Service Quality

Provide data powered proactive service improvement

API enablement

Enable API development, deployment and management





Creation



Management platform



Culture & Organization







ZABBIX IN NATIXIS



2

ZABBIX IN NATIXIS


Number of hosts	20 324
Number of items	2 089 841
Number of triggers	1 538 461
Number of users	1 244
NVPS	7 500



Zabbix Server :

2 Blade

- 2 - INTEL(R) XEON(R) CPU E5-2620 V3 @ 2.40GHZ:X86_64(2394Mhz)
- 128 Go RAM / 2 x 300Go HD
- Linux : Redhat 7
- Cluster VCS




Frontend

On 2 VMs

- 2 Apache instances
 - Frontend Users
 - API

Pairs dedicated

- DMZ
 - Tests web
 - Servers
- Hexadata Database
- Tests web



14 Proxy servers :

7 pairs

- 2 - INTEL(R) XEON(R) CPU L5640 @ 2.27GHZ:X86_64(2266Mhz)
- 20 Go RAM / 2 x 300Go HD
- Linux : Redhat 6 & 7
- Active/passive loadbalancing with F5

2 | Zabbix in Natixis : Industrial processes

All tasks of affectation are automatically set :

Users

- Admin in Zabbix on their perimeter
- Are responsibled about their own supervision
- Declared automatically according to the outlook Email list associated with the production team



Hosts

- Automatically attached to Hostgroup associated with the production Team
- System Template automatically link to the Host.
- Linked with System Team



API Frontend

About 100 jobs in Perl or Php



2

Zabbix in Natixis

ZABBIX **ToolBox** **NATIXIS**

DBA Host Item Maintenance MAP NAS Pilotage Reports Template

Summary
Change host status : enable or disable

Description
Allow to change the status of a particularly host
Please see parameters section

Parameters

Name	Located In	Description	Required	Type
physicalHostName	query	name of the physical host in Zabbix	Yes	String
HostGroupname	query	name of the hostgroup in Zabbix where the physicalHostName belongs	Yes	String
option	query	enable or disable	yes	String

Team
---Browse Usergroups---

HostGroup
---Browse Hostgroups---

Host
---Browse Hosts---

Status

Enable Server **Disable Server**

Copy API Pattern URL

http://toolbox.zabbix.intranatixis.com/cgi-bin/host/changeStatusHost.cgi?option=[enable|disable]&hostGroupName=xxxx&physicalHostName=yyyy


Toolbox Frontend :

- Repetitive users actions
- API calls





3

MIGRATION METHODOLOGY TO ZABBIX 3.2 IN LARGE ENVIRONMENT



3 Migration Methodology to Zabbix 3.2 in large environment



Target architecture

For Galera and partitioning, add columns and primary key

3

Migration Methodology to Zabbix 3.2 in large environment

- Where are the limits ?

Stress module

Jan Garaj Project (<https://github.com/monitoringartist/zabbix-server-stress-test>)

- **10 Zabbix Agents**
 - Simulated loads of 4 000 hosts

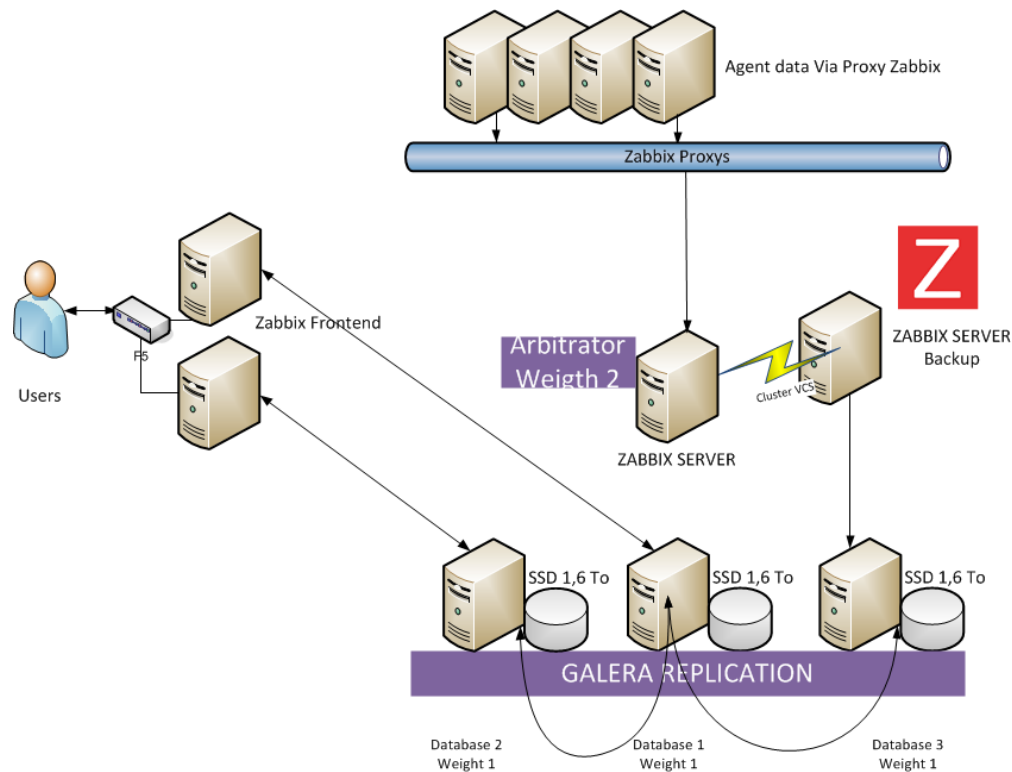


Environment

- **Prod Database**
 - restored on 2 nodes 43/44 (garbd on Zabbix Server),
 - All hosts are disabled except agents for Bench located in host group GALERA_BENCH.
 - All actions are disabled. All proxies are configured as Active.
- **Galera**
 - mysql-wsrep-5.6.33-25.17-linux-x86_64
 - sst xtrabackup plugin.
 - Process GARBD on Zabbix server
- **Zabbix 3.2.3**

3 Migration Methodology to Zabbix 3.2 in large environment

- Where are the limits ?
 - Architecture limits



Zabbix Flow loadbalanced
between Database


➤ Dead locks

3


Migration Methodology to Zabbix 3.2 in large environment

• Where are the limits ?

- Stress limits
 - Tested till 30k Nvps (only numeric) without problem.
 - This test generates only Mysql inserts.



- Tested at 20k Nvps
 - Trigger with 1 condition created in TPL System (4300 hosts)
 - Condition 8 added in same Template
 - Row lock waits appear



3

Migration Methodology to Zabbix 3.2 in large environment

- Which methodology to only have 5 minutes downtime ?

BEFORE D DAY - PREPARATIONS

- Archiving
- Prepare all 3.2.6 binaries
- Prepare new configs for new/obsolete parameters


DRY RUNS

D DAY

- Dump production (11 minutes)
- Start Migration (while Zabbix 2.4 keeps running)
 - Restore Dump on new Database
 - Clean all orphan triggers inherited from 2.4.
 - Update to 3.2.6 on other Zabbix Server linked to new Database (3 minutes)
 - STOP other Zabbix server
 - Reset all triggers
 - Basic check from frontends
- Stop production 2.4 Zabbix Platform
 - Replace all binaries with appropriate symbolic links
 - Replace new Configuration files
- Start Zabbix 3.2.6
 - Check alerts
 - Check Zabbix Internal



Without
history and
trends



Only 5
minutes
downtime

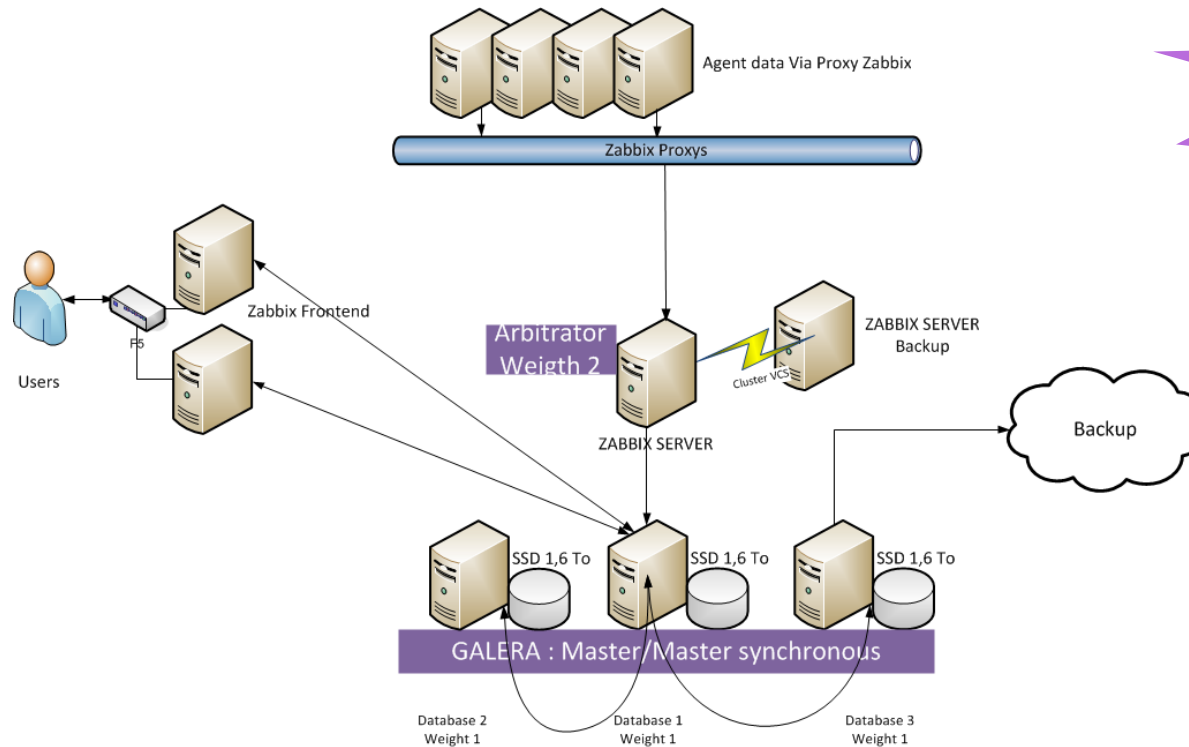


GALERA ARCHITECTURE IN NATIXIS



4

Galera architecture in NATIXIS



Final
architecture

Dedicated server for backup

- Extra Backup plugin
- Stop database on this node easily



Zabbix recommendations

- All access to the same database

Simple architecture

- No Database proxy
- Switch config from Zabbix




GALERA RUNTIME EXPERIENCE



5 Galera Runtime experience

With database in full production (700 Go)




SST Test on backup node

- *Node 1 and 2 Synced*
- *Stop node instance 3*
- *Delete all datas*
- *Start*

Total Time to execute : 1H42


- *SST : 1h09*
- *Prepare : 0h22*
- *Move : 0h06*
- *Start + resynchro : 0h05*





6

QUESTIONS





ANNEX



Add primary keys for Galera

Change keys for Galera:

```
alter table dbversion ADD id4galera INT PRIMARY KEY AUTO_INCREMENT;  
alter table history add id4galera int key auto_increment;  
alter table history drop primary key , add primary key (id4galera,clock);  
alter table history_uint add id4galera int key auto_increment;  
alter table history_uint drop primary key , add primary key (id4galera,clock);  
alter table history_log add id4galera int key auto_increment;  
alter table history_log drop primary key , add primary key (id4galera,clock);  
alter table history_text add id4galera int key auto_increment;  
alter table history_text drop primary key , add primary key (id4galera,clock);  
alter table history_str add id4galera int key auto_increment;  
alter table history_str drop primary key , add primary key (id4galera,clock)
```

History* & trends* Tables Partitionning

Create procedure from http://zabbix.org/wiki/Docs/howto/mysql_partition

Use this for “partition_maintenance_all”: