



基于Ambari的PG企业数据架构 运维管理方案



嘉宾：赖伟

公司：神州飞象（北京）数据科技有限公司

邮箱：daniel.lai@postgresdata.com





目录



Ambari 运维管理平台简介



PG高可用集群方案简介



PG高可用集群运维管理平台



01

Ambari 运维管理平台 简介



什么是Ambari ?

- Apache 顶级开源项目
- 创建、管理、监视 Hadoop 的集群
- 支持自定义Stacks and Services
- 内置多种服务组件(如Ambari Metrics)



Ambari功能亮点：向导式安装部署

Add Service Wizard

ADD SERVICE WIZARD

[Choose Services](#)

Assign Masters

[Assign Slaves and Clients](#)

[Customize Services](#)

[Configure Identities](#)

[Review](#)

[Install, Start and Test](#)

[Summary](#)

Assign Masters

Assign master components to hosts you want to run them on.

ZooKeeper Server:

Kafka Broker:

Kafka Broker:

Kafka Broker:

www.hadoop01.zgw (3.7 GB, 2 cores)

ZooKeeper Server Kafka Broker

www.hadoop02.zgw (2.1 GB, 2 cores)

Kafka Broker

www.hadoop03.zgw (2.1 GB, 2 cores)

Kafka Broker

[← Back](#)

[Next →](#)



Ambari功能亮点：集中配置

Add Service Wizard

Assign Slaves and Clients

Customize Services

Configure Identities

Review

Install, Start and Test

Summary

ZooKeeper

PostgreSQL_HA

Misc

Group

Default (3)

Manage Config

Advanced all-env

Advanced etcd-env

Advanced haproxy-env

Advanced keepalived-env

Advanced patroni-env

Advanced pg_install-conf

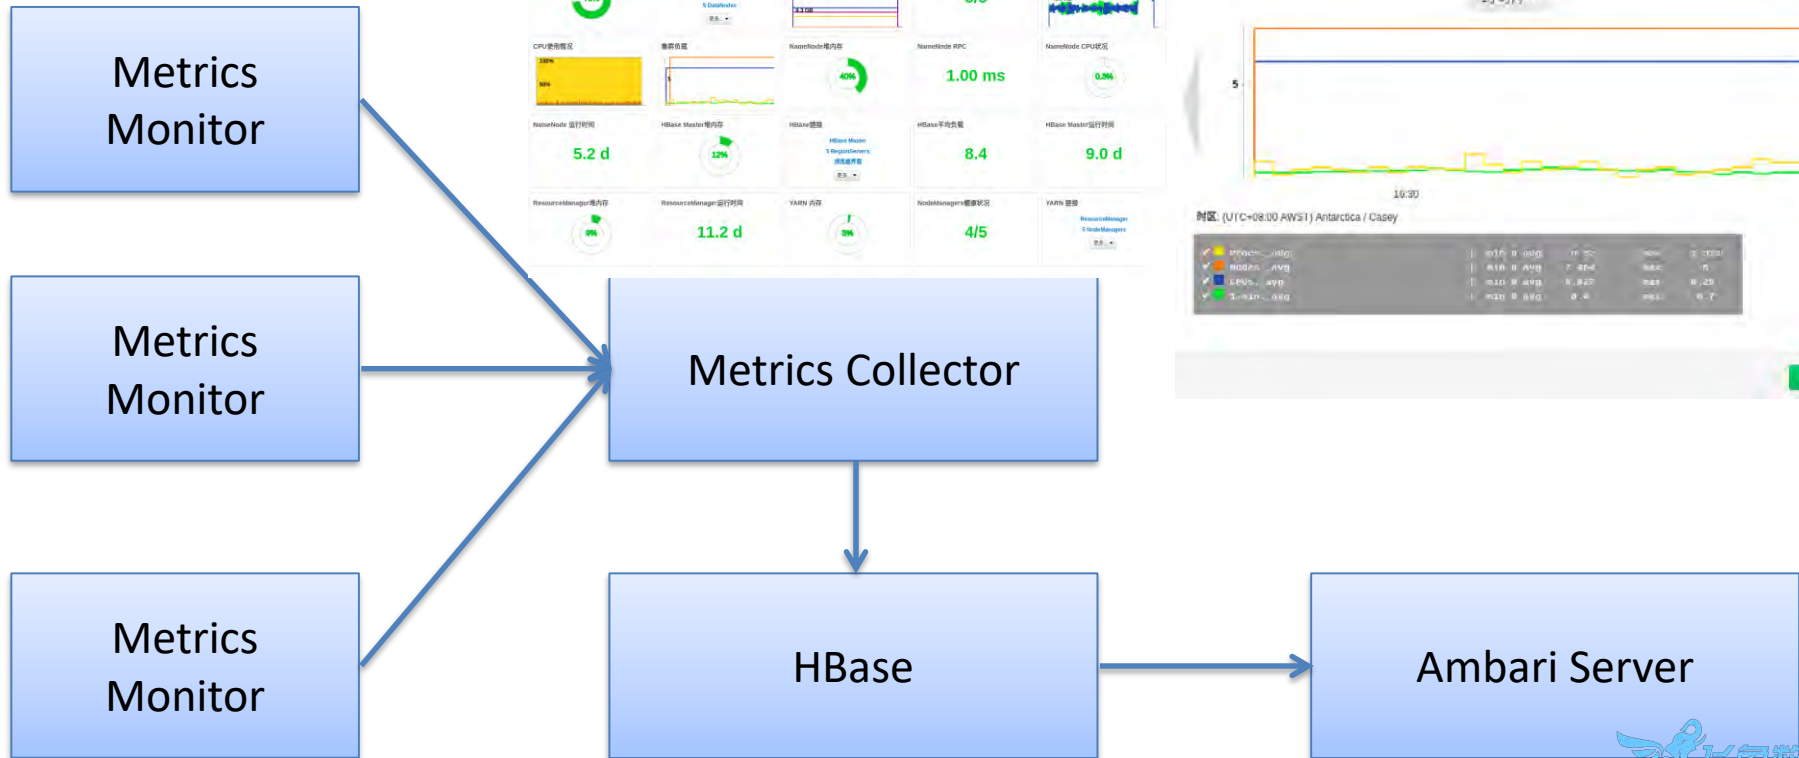
Advanced patroni-env

master_database_port	5432	+	C
master_rest_port	8008	+	C
master_yml_name	postgresq0	+	C
patroni_bin_dir	/home/postgres/patroni/	+	C
patroni_need_package_zip_url	http://192.168.120.130/ambari/postgresql_ha/patroni_need_package.zip	+	C
patroni_zip_url	http://192.168.120.130/ambari/postgresql_ha/patroni-1.2.5.zip	+	C
slave_database_port	5433	+	C
slave_rest_port	8009	+	C
slave_yml_name	postgresq1	+	C

```
scope: batman
#namespace: /service/
name: {yml_name}
```



Ambari功能亮点：分布式监控





Ambari功能亮点：告警通知

创建警告通知

组 全部 自定义

- YARN Default
- ATLAS Default
- SPARK2 Default
- KAFKA Default

选择所有 | 清除所有

严重程度

- OK
- WARNING
- CRITICAL
- UNKNOWN

选择所有 | 清除所有

描述

方法

电子邮件

SMTP 服务器

SMTP 端口

电子邮件来自于

使用身份验证

取消 保存

配置

Description This host-level alert is triggered if the amount of disk space used goes above specific thresholds. The default threshold values for WARNING and 80% for CRITICAL.

Check Interval 1 Minute

Minimum Free Space **WARNING** 5000000000 bytes

Critical **CRITICAL** 80 %

Warning **WARNING** 50 %

实例

服务	主机	状态	24小时
Ambari	bigdata-cn-01.cars.com	CRIT for 4 days	0
Ambari	bigdata-dn-01.cars.com	WARN for 4 days	0





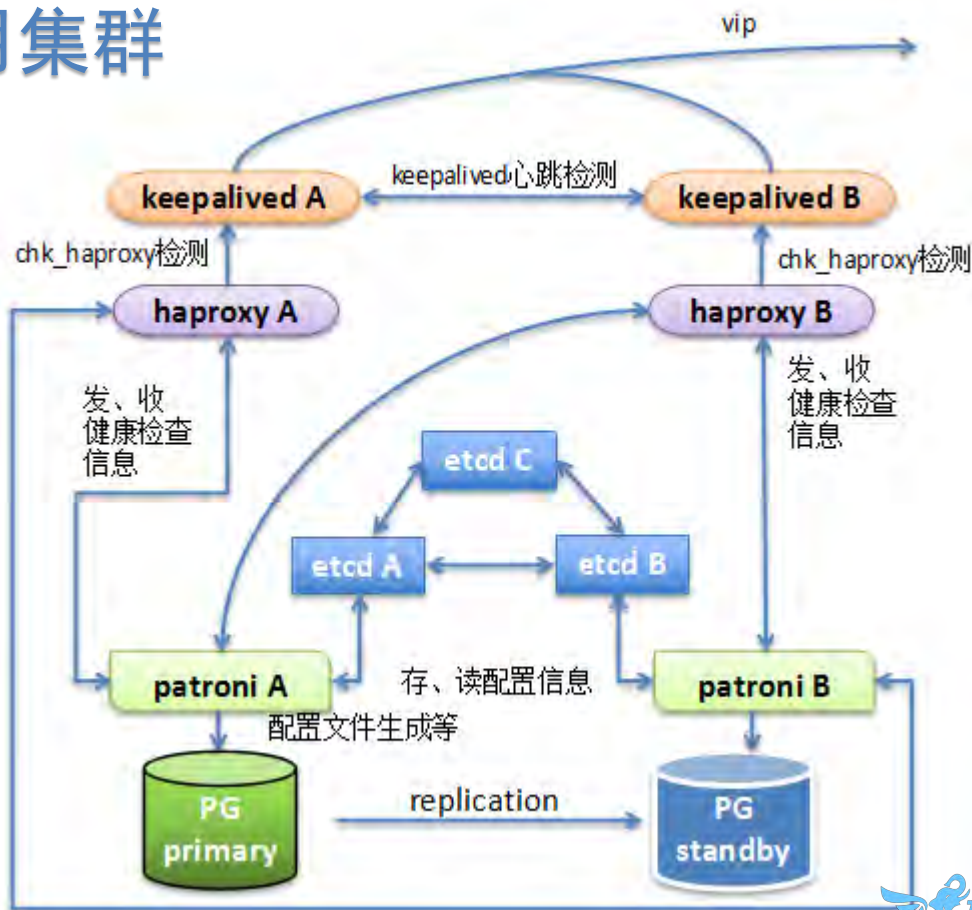
02

PG高可用集群 方案简介



基于Patroni的PG高可用集群

patroni是一款运用etcd集群及其主备状态与配置信息,通过patroni来检测并且实现自动切换的软件。并且可运用haproxy及keepalived保持主备切换或者节点故障后,访问地址、端口对上层不变。





基于Patroni的PG高可用集群优势

- 自动检测主备状态进行切换
- 统一模板配置
- 利用etcd等DCS系统，有效防止脑裂
- 在线添加etcd、patroni节点以及数据节点



基于Patroni的PG高可用集群优势

- 支持同步异步流复制，级联流复制
- 异步流复制可设置最小丢失数据量
- 可用pg_rewind进行恢复，缩短恢复时间
- haproxy+keepalived可以保持在主备切换或者节点故障后，实现ip漂移。对外的ip+端口不变。



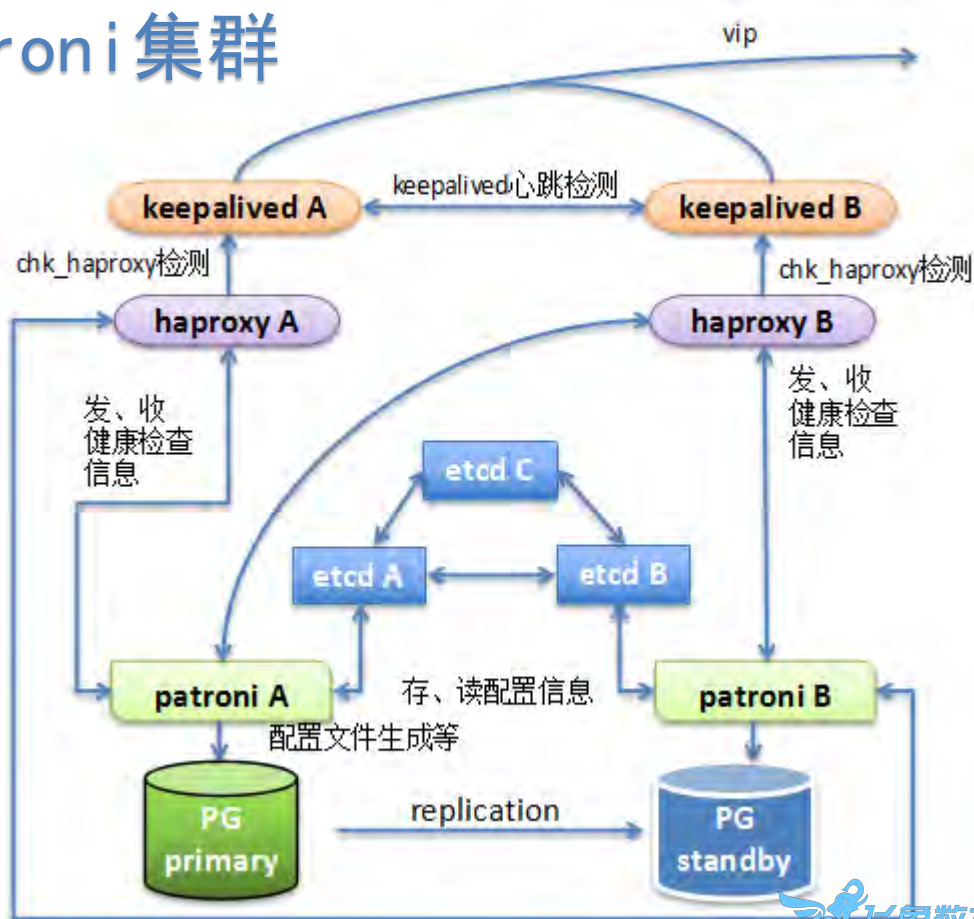
03

PG高可用集群运维 管理平台




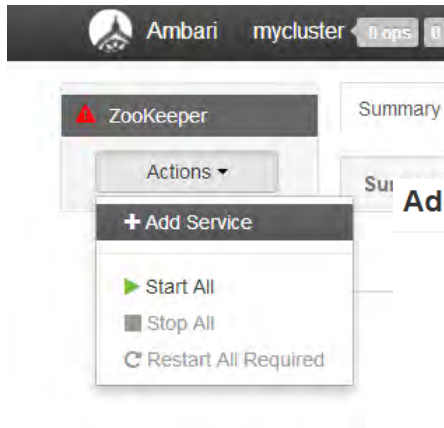
为什么要用Ambar i管理patroni集群

- 集群搭建较繁琐
- 状态信息欠直观
- 组件启停不方便
- 操作审计难追踪






搭建PG高可用集群： 第一步（选择服务）



点击主页面上的**Actions**中的**Add Service**进入服务添加页面

Add Service Wizard

勾选**PostgreSQL_HA**
并点击**next**



<input type="checkbox"/>	SmartSense	1.3.1.0-136	SmartSense - Hortonworks SmartSense Tool (HST) helps quickly gather configuration, metrics, logs from common HDP services that aids to quickly troubleshoot support cases and receive cluster-specific recommendations.
<input type="checkbox"/>	Spark	1.6.2	Apache Spark is a fast and general engine for large-scale data processing.
<input type="checkbox"/>	Spark2	2.0.0	Apache Spark 2.0 is a fast and general engine for large-scale data processing. This service is Technical Preview .
<input type="checkbox"/>	Zeppelin Notebook	0.6.0	A web-based notebook that enables interactive data analytics. It enables you to make beautiful data-driven, interactive and collaborative documents with SQL, Scala and more.
<input type="checkbox"/>	Mahout	0.9.0	Project of the Apache Software Foundation to produce free implementations of distributed or otherwise scalable machine learning algorithms focused primarily in the areas of collaborative filtering, clustering and classification
<input checked="" type="checkbox"/>	PostgreSQL_HA	1.0.0	PostgreSQL_HA_PATRONI
<input type="checkbox"/>	Slider	0.91.0	A framework for deploying, managing and monitoring existing distributed applications on YARN.

Next →



搭建PG高可用集群：第二步（规划Master服务安装）

根据服务架构合理选择组件安装

如：

haproxy需和**keepalived**安装在同一服务器上。

Add Service Wizard

ADD SERVICE WIZARD

Choose Services

Assign Masters

Assign Slaves and Clients

Customize Services

Configure Identities

Review

Install, Start and Test

Summary

Assign Masters

Assign master components to hosts you want to run them on.

ZooKeeper Server:

HAprouxy_Server:

Patroni_Server_Master:

keepalived_Server:

etcd_Server:

www.hadoop01.zgw (3.7 GB, 2 c

- ZooKeeper Server
- HAprouxy_Server
- Patroni_Server_Master
- keepalived_Server
- etcd_Server

2 hosts not running master services

← Back

Next →

www.hadoop01.zgw (3.7 GB, 2 cores)	ZooKeeper Server	etcd_Server		
www.hadoop02.zgw (2.1 GB, 2 cores)	HAprouxy_Server	Patroni_Server_Master	keepalived_Server	etcd_Server
www.hadoop03.zgw (2.1 GB, 2 cores)	HAprouxy_Server	keepalived_Server	etcd_Server	



搭建PG高可用集群：第三步（规划Slave及client服务安装）

Client为类似于仅安装服务文件，不存在启动、停止及状态检测
该处的**Client**被定义为PG_HOME。

Add Service Wizard

x

ADD SERVICE WIZARD

[Choose Services](#)[Assign Masters](#)[Assign Slaves and Clients](#)[Customize Services](#)[Configure Identities](#)[Review](#)[Install, Start and Test](#)[Summary](#)

Assign Slaves and Clients

Assign slave and client components to hosts you want to run them on.

Hosts that are assigned master components are shown with *.

"Client" will install PostgreSQL_install

Host	all none	all none
www.hadoop01.zgw*	<input type="checkbox"/> Patroni_Server_Slave	<input type="checkbox"/> Client
www.hadoop02.zgw*	<input type="checkbox"/> Patroni_Server_Slave	<input checked="" type="checkbox"/> Client
www.hadoop03.zgw*	<input checked="" type="checkbox"/> Patroni_Server_Slave	<input checked="" type="checkbox"/> Client

Show: 25 1 - 3 of 3

[← Back](#)[Next →](#)



搭建PG高可用集群：第四步（确定配置文件）

一般为默认参数，也可根据需求修改配置文件参数

Add Service Wizard

Assign Slaves and Clients

Customize Services

Configure Identities

Review

Install, Start and Test

Summary

ZooKeeper PostgreSQL_HA Misc

Group Default (3) Manage Config G

Advanced all-env

Advanced etcd-env

Advanced haproxy-env

Advanced keepalived-env

Advanced patroni-env

Advanced pg_install-conf

master_database_port	5432	+	C
master_rest_port	8008	+	C
master_yml_name	postgresql0	+	C
patroni_bin_dir	/home/postgres/patroni/	+	C
patroni_need_package_zip_url	http://192.168.120.130/ambari/postgresql_ha/patroni_need_package.zip	+	C
patroni_zip_url	http://192.168.120.130/ambari/postgresql_ha/patroni-1.2.5.zip	+	C
slave_database_port	5433	+	C
slave_rest_port	8009	+	C
slave_yml_name	postgresql1	+	C

```
scope: batman
#namespace: /service/
name: {yml_name}
```



搭建PG高可用集群：第五步（核验信息）

确认无误后可点击**Deploy**开始执行安装

Add Service Wizard

[Assign Masters](#)

[Assign Slaves and Clients](#)

[Customize Services](#)

[Configure Identities](#)

[Review](#)

[Install, Start and Test](#)

[Summary](#)

Please review the configuration before installation

Admin Name : admin

Cluster Name : mycluster

Total Hosts : 3 (0 new)

Repositories:

redhat7 (HDP-2.5):
http://192.168.120.130/ambari/HDP/centos7/

redhat7 (HDP-UTILS-1.1.0.21):
http://192.168.120.130/ambari/HDP-UTILS/

Services:

PostgreSQL_HA

HProxy_Server : 2 hosts
Patroni_Server_Master : www.hadoop02.zgw
Patroni_Server_Slave : 1 host
etcd_Server : 3 hosts
keepalived_Server : 2 hosts

[← Back](#)

[Print](#)

[Deploy →](#)



搭建PG高可用集群：第六步（等待安装完成）

待进度条100%并且全部显示为Success即为所有服务安装完成并启动成功。



Add Service Wizard

ADD SERVICE WIZARD

- Choose Services
- Assign Masters
- Assign Slaves and Clients
- Customize Services
- Configure Identities
- Review
- Install, Start and Test**
- Summary

Install, Start and Test

Please wait while the selected services are installed and started.



Show: **All (3)** | [In Progress \(3\)](#) | [Warning \(0\)](#) | [Success \(0\)](#) | [Fail \(0\)](#)

Host	Status	Message
www.hadoop01.zgw	33%	Install complete (Waiting to start)
www.hadoop02.zgw	23%	Installing etcd_Server
www.hadoop03.zgw	23%	Installing etcd_Server

3 of 3 hosts showing - [Show All](#)

Show: 25 | 1 - 3 of 3

Next →

Install, Start and Test

Please wait while the selected services are installed and started.



Show: **All (3)** | [In Progress \(0\)](#) | [Warning \(0\)](#) | [Success \(3\)](#) | [Fail \(0\)](#)

Host	Status	Message
www.hadoop01.zgw	100%	Success
www.hadoop02.zgw	100%	Success
www.hadoop03.zgw	100%	Success

3 of 3 hosts showing - [Show All](#)

Show: 25 | 1 - 3 of 3

所有服务均为先安装完成再启动。



搭建PG高可用集群：第七步（生成摘要信息）

生成完成摘要信息，点击**Complete**即可

Add Service Wizard

ADD SERVICE WIZARD

Choose Services

Assign Masters

Assign Slaves and Clients

Customize Services


Configure Identities

Review

Install, Start and Test

Summary

Summary

Important: You may also need to restart other services for the newly added services to function properly (for example, HDFS and YARN/MapReduce need to be restarted after adding Oozie). After closing this wizard, please restart all services that have the restart indicator  next to the service name.

Here is the summary of the install process.

The cluster consists of 3 hosts

Installed and started services successfully on 3 new hosts

Install and start completed in 36 seconds

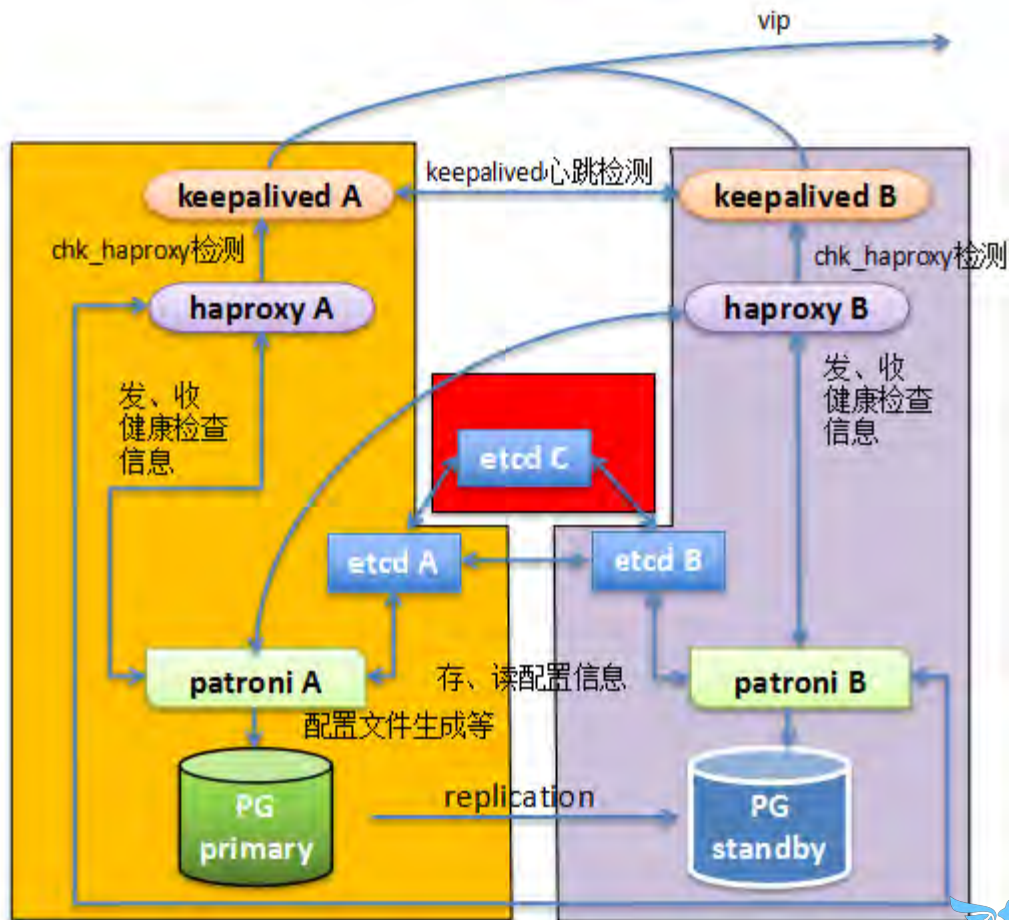
Complete →





完成集群搭建

- www.hadoop01.zgw
- www.hadoop02.zgw
- www.hadoop03.zgw





管理PG高可用集群：组件的启、停

Ambari mycluster 0 ops 0 alerts Dashboard Services

Ambari mycluster 0 ops 0 alerts

PostgreSQL_H A

Actions ▾

Summary Configs

Summary

- [HAproxy_Server](#) ✓ Started **No alerts**
- [HAproxy_Server](#) ✓ Started **No alerts**
- [Patroni_Server_Master](#) ✓ Started **No alerts**
- [etcd_Server](#) ✓ Started **No alerts**
- [etcd_Server](#) ✓ Started **No alerts**
- [etcd_Server](#) ✓ Started **No alerts**
- [keepalived_Server](#) ✓ Started **No alerts**
- [keepalived_Server](#) ✓ Started **No alerts**
- [Patroni_Server_Slave](#) 1/1 Patroni_Server_Slave Live
- [PostgreSQL_installs](#) 2 PostgreSQL_installs Installed

www.hadoop02.zgw

← Back

Summary Configs Alerts 0 Versions

Components + Add

- ✓ HAproxy_Server / PostgreSQL_HA Started ▾
- ✓ Patroni_Server_M... / PostgreSQL_HA Started ▾
- ✓ etcd_Server / PostgreSQL_HA Started ▾
- ✓ keepalived_Server / PostgreSQL_HA Clients / PostgreSQL_install Started ▾
- Installed ▾

Started ▾

- Restart
- Stop
- Turn On Maintenance Mode
- Delete



管理PG高可用集群：组件状态

kill掉hadoop02上的etcd进程，GUI上会显示“stopped”

```
[root@hadoop02 ~]# ps -ef |grep etcd
postgres  8461      1  1 09:21 ?        00:14:14 /home/postgres/etcd//etcd --name etcd2 --data-dir=/home/postgres/etcd/data/ --initial-advertise-peer-urls http://www.hadoop02.zgw:2380 --listen-peer-urls http://www.hadoop02.zgw:2380 --listen-client-urls http://www.hadoop02.zgw:2379,http://127.0.0.1:2379 --advertise-client-urls http://www.hadoop02.zgw:2379 --initial-cluster-token etcd-cluster-1 --initial-cluster ,etcd1=http://www.hadoop03.zgw:2380,etcd2=http://www.hadoop02.zgw:2380,etcd3=http://www.hadoop01.zgw:2380 --initial-cluster-state new
root      25119    25010  0  22:22 pts/1    00:00:00 grep --color=auto etcd
[root@hadoop02 ~]# kill -15 8461
[root@hadoop02 ~]# ps -ef |grep etcd
root      25207    25010  0  22:23 pts/1    00:00:00 grep --color=auto etcd
[root@hadoop02 ~]#
```

Summary

HAproxy_Server	✔ Started	No alerts
HAproxy_Server	✔ Started	No alerts
Patroni_Server_Master	✔ Started	No alerts
etcd_Server	✔ Started	No alerts
etcd_Server	⚠ Stopped	No alerts
etcd_Server	✔ Started	No alerts
keepalived_Server	✔ Started	No alerts
keepalived_Server	✔ Started	No alerts
Patroni_Server_Slave	1/1 Patroni_Server_Slave Live	
PostgreSQL_installs	2 PostgreSQL_installs Installed	

www.hadoop02.zgw

Summary | Configs | Alerts 0 | Versions

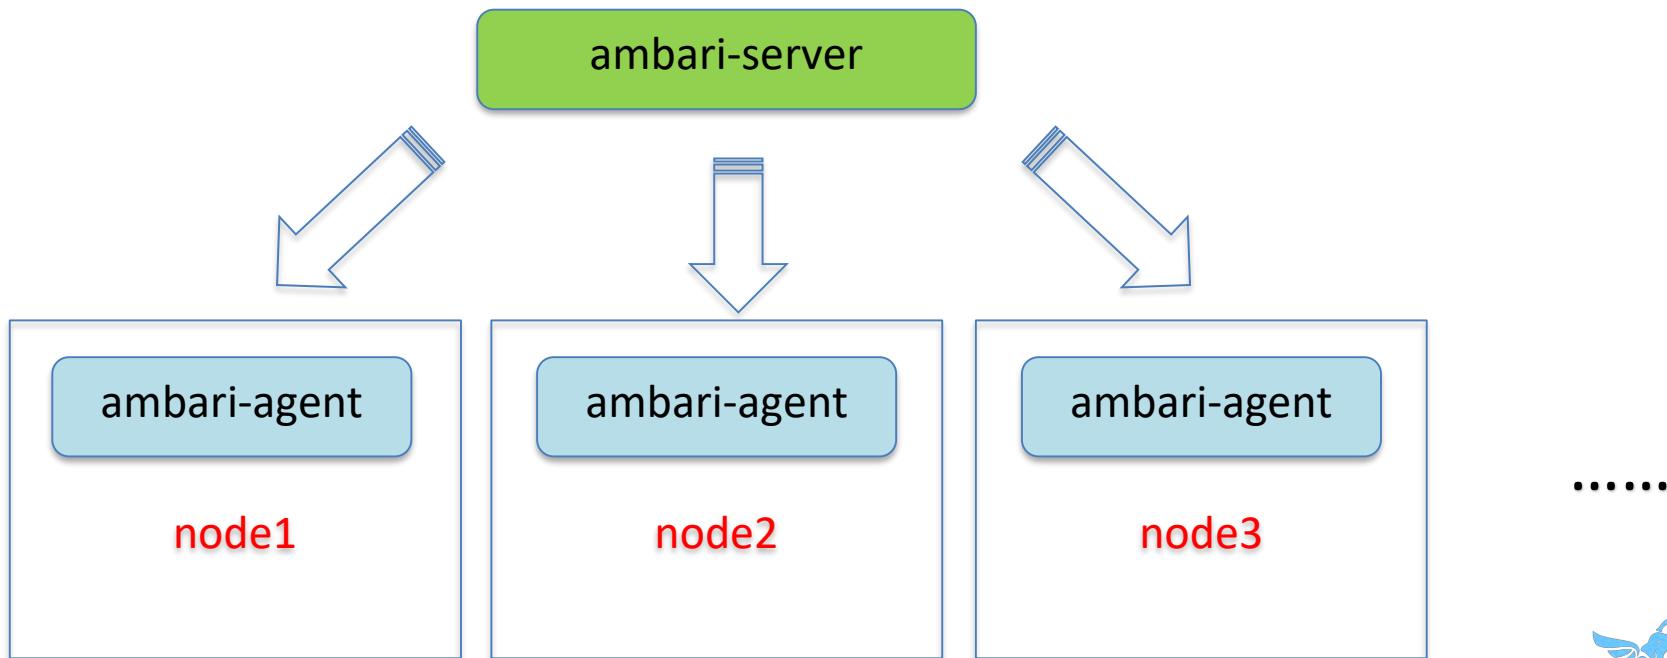
Components + Add

✔ HAproxy_Server / PostgreSQL_HA	Started
✔ Patroni_Server_M... / PostgreSQL_HA	Started
⚠ etcd_Server / PostgreSQL_HA	Stopped
✔ keepalived_Server / PostgreSQL_HA	Started
Clients / PostgreSQL_install	Installed



管理PG高可用集群：安装、启停、状态检查解析

Ambari为分布式软件，在执行脚本时会将ambari-server上的python脚本分发到各个指定的ambari-agent节点上执行。





管理PG高可用集群：安装、启停、状态检查解析

```
def install(self, env):
```

```
def start(self, env):
```

```
def stop(self, env):
```

```
def status(self, env):
```

Master/Slave类型脚本

```
def install(self, env):
```

Client类型脚本

添加服务时执行“Deploy”被调用

对应GUI中的“start”按键，添加服务时会被调用一次

对应GUI中的“stop”按键

Ambari循环检测调用，官方采用 `check_process_status(pid_file_path)` 函数进行检测，每次循环检测后来刷新组件GUI上的状态（“started”或“stopped”）



管理PG高可用集群：安装、启停、状态检查解析

check_process_status()函数解析：

```
def check_process_status(pid_file):
    from resource_management.core import sudo
    if not pid_file or not os.path.isfile(pid_file):
        Logger.info("Pid file {0} is empty or does not exist".format(str(pid_file)))
        raise ComponentIsNotRunning()
    try:
        pid = int(sudo.read_file(pid_file))
    except:
        Logger.info("Pid file {0} does not exist or does not contain a process id
number".format(pid_file))
        raise ComponentIsNotRunning()
    try:
        sudo.kill(pid, 0)
    except OSError:
        Logger.info("Process with pid {0} is not running. Stale pid file"
            " at {1}".format(pid, pid_file))
        raise ComponentIsNotRunning()
```



管理PG高可用集群： 组件参数修改

Save Configuration

Notes: test update bind_port 5001

Cancel Discard Save

PostgreSQL_HA

Summary **Configs**

Group: Default (3) Manage Config Groups

V1 admin 6 minutes ago HDP-2.5

V1 admin authored on Tue, Oct 17, 2017 16:20

- Advanced all-env
- Advanced etcd-env
- Advanced haproxy-env**
- Advanced keepalived-env
- Advanced patroni-env

V1 admin authored on Tue, Oct 17, 2017 16:20

Advanced haproxy-env

haproxy_config_file_dir /home/(linux_user)/haproxy

haproxy_rpm_url http://192.168.120.130/ambari/postgresql_ha/haproxy-1.5.18-6.el7.x86_64.rpm

```
global
  maxconn 100

defaults
  log global
  mode tcp
  retries 2
  timeout client 30m
  timeout connect 4s
  timeout server 30m
  timeout check 5s

frontend ft_postgresql
  bind *:5001
  default_backend bk_db
```



管理PG高可用集群：组件参数修改

Metrics Heatmaps **Config History**

Service	Config Group	Created	Author	Notes
All	All	Any	Any	Any
V2 PostgreSQL_HA	Default Current	Tue, Oct 17, 2017 16:47	admin	test update bind_port 5001
V1 PostgreSQL_HA	Default	Tue, Oct 17, 2017 16:20	admin	Initial configurations for PostgreSQL_HA

www.hadoop02.zgw

← Back

Summary **Configs** Alerts 0 Versions

Components

+ Add

- HAProxy_Server / PostgreSQL_HA
 - Patroni_Server_M... / PostgreSQL_HA
 - etcd_Server / PostgreSQL_HA
 - keepalived_Server / PostgreSQL_HA
- Clients / PostgreSQL_install

Started

Restart

Stop

Turn On Maintenance Mode

Delete

Installed

旧的配置文件在Ambari中不会被删除，将会生成一个新版本，并可查看到修改记录

重启组件新的参数将会覆盖原有的配置文件

```
[root@hadoop02 haproxy]# netstat -apn|grep 5001
tcp        0      0 0.0.0.0:5001          0.0.0.0:*              LISTEN     88426/haproxy
```



管理PG高可用集群：告警

告警可以用来显示非组件服务故障，如patroni中的PGdb是否在运行。

PGdb Alert OK

[← Back](#)

Configuration [Edit](#)

Description: This is PGdb alert .

Check Interval: 1 Minute

Thresholds:

- OK** TCP OK - {0..3f}s response on port (1)
- WARNING** 1.5 Seconds TCP OK - {0..3f}s response on port {1}
- CRITICAL** 5 Seconds Connection failed: {0} to {1};{2}

State: ✔ Enabled

Service: PostgreSQL_HA

Component: Patroni_Server_Master

Type: PORT

Groups: PostgreSQL_HA Default

Last Changed: Wed, Oct 18, 2017 22:53

Check Count: 1 (default) [✎](#)

es Hosts Alerts Admin admin

Instances

Service	Host	Status	24-Hour	Response
PostgreSQL_HA	www.hadoop02.zgw	✔ OK	for 3 hours	19 TCP OK - 0.000s response on port 5433
Ambari Agent Heartbeat		✔ OK (3)		Ambari
Ambari Server Performance		✔ OK		Ambari
Ambari Server Alerts		✔ OK		Ambari
PGdb Alert		✔ OK		PostgreSQL_HA

Last Status Changed	State
about a day ago	✔ Enabled
11 hours ago	✔ Enabled
about a day ago	✔ Enabled
6 hours ago	✔ Enabled
3 hours ago	✔ Enabled



管理PG高可用集群：告警

kill 掉后台的PG主节点进程，
GUI上会出现一条CRITICAL
级别的告警信息



● PostgreSQL_ HA 1

Actions ▾

Summary
Configs

⚙️ Restart Required: 9 Components on 3 Hosts

Summary

- [HAproxy Server](#) ✔ Started No alerts
- [HAproxy Server](#) ✔ Started No alerts
- [Patroni Server Master](#) ✔ Started 1 alert
- [etcd Server](#) ✔ Started No alerts
- [etcd Server](#) ✔ Started No alerts
- [etcd Server](#) ✔ Started No alerts
- [keepalived Server](#) ✔ Started No alerts
- [keepalived Server](#) ✔ Started No alerts
- [Patroni Server Slaves](#) 2/2 Patroni_Server_Slaves Live
- [PostgreSQL installs](#) 3 PostgreSQL_installs Installed

```
[root@hadoop02 ~]# ps -ef |grep postgres
postgres 16746      1  0 10:52 ?                00:00:32 python /home/postgres/pat
stgresql0.yml
postgres 16760      1  0 10:52 ?                00:00:00 postgres -D data/postgres
cesses=8 --max_locks_per_transaction=64 --wal_keep_segments=8 --wal_level=r
ck_commit_timestamp=off --max_prepared_transactions=0 --port=5433 --max_rep
postgres 16762    16760  0 10:52 ?                00:00:00 postgres: batman: startup
postgres 16765    16760  0 10:52 ?                00:00:00 postgres: batman: checkpo
postgres 16766    16760  0 10:52 ?                00:00:00 postgres: batman: writer
postgres 16767    16760  0 10:52 ?                00:00:00 postgres: batman: stats collector process
postgres 16771    16760  0 10:52 ?                00:00:15 postgres: batman: postgres postgres 192.168.120.131(59558) idle
postgres 16809    16760  0 10:52 ?                00:00:09 postgres: batman: wal receiver process streaming 0/50C2180
root      46397    46314  0 13:57 pts/1          00:00:00 grep --color=auto postgres
postgres 109175      1  0 06:43 ?                00:00:21 haproxy -f /home/postgres/haproxy/haproxy.cfg
postgres 109640      1  2 Oct17 ?                00:19:07 /home/postgres/etcd//etcd --name etcd2 --data-dir=/home/postgres/etcd/data/ --initial-advertise-pe
er-urls http://www.hadoop02.zgw:2380 --listen-peer-urls http://www.hadoop02.zgw:2380 --listen-client-urls http://www.hadoop02.zgw:2379,http://127.0.
w:2380,etcd2=http://www.hadoop02.zgw:2380,etcd3=http://www.hadoop01.zgw:2380 --initial-cluster-token etcd-cluster-1 --initial-cluster ,etcd1=http://www.hadoop02.zgw:2380,etcd2=http://www.hadoop02.zgw:2380,etcd3=http://www.hadoop01.zgw:2380 --initial-cluster-state new
[root@hadoop02 ~]# kill -9 16760
```





管理PG高可用集群：添加节点

向导方式添加节点，删除节点也只需在相应节点页面点击 **delete host**，需要先停止并删除节点上所有服务。

Add Host Wizard

ADD HOST WIZARD

Install Options

Confirm Hosts

Assign Slaves and Clients

Configurations

Review

Install, Start and Test

Summary

Install Options

Enter the list of hosts to be included in the cluster and provide yo

Target Hosts

Enter a list of hosts using the Fully Qualified Domain Name (FQD

host names

Host Registration Information

Provide your [SSH Private Key](#) to automatically register hosts

未选择任何文件

ssh private key



Ambari

mycluster

0 ops

0 alerts

Actions

+ Add New Hosts

Selected Hosts (0)

Filtered Hosts (3)

All Hosts (3)

www.hadoop02.zgw

192.168.120.130/default-rack

www.hadoop03.zgw

192.168.120.131/default-rack

2



管理PG高可用集群：添加组件

Ambari mycluster 0 ops 0 alerts

www.hadoop01.zgw

Back

Summary Configs Alerts 0 Versions

Components

+ Add

etcd_Server / PostgreSQL_HA

Clients / PostgreSQL_install

HProxy_Server

Patroni_Server_Slave

keepalived_Server

Summary

Hostname: www.hadoop01.zgw
IP Address: 192.168.120.130

```
postgres=# select * from pg_replication_slots ;
 slot_name | plugin | slot_type | datoid | database | active | active_pid | xmin | catalog_xmin | restart_lsn | confirmed_flush_lsn
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
 postgresql1_0 |      | physical |      |      | t      | 109819 |      |      | 0/3000000 |
(1 row)
```

```
postgres=# select * from pg_replication_slots ;
 slot_name | plugin | slot_type | datoid | database | active | active_pid | xmin | catalog_xmin | restart_lsn | confirmed_flush_lsn
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----
 postgresql1_0 |      | physical |      |      | t      | 109819 |      |      | 0/5000060 |
 postgresql1_1 |      | physical |      |      | t      | 110079 |      |      | 0/5000060 |
(2 rows)
```

postgres=# █



www.hadoop01.zgw

Back

Summary Configs Alerts 0 Versions

Components

+ Add

Host needs 2 components restarted

Restart

etcd_Server / PostgreSQL_HA



Started

Patroni_Server_S... / PostgreSQL_HA

Clients / PostgreSQL_install



Stopped

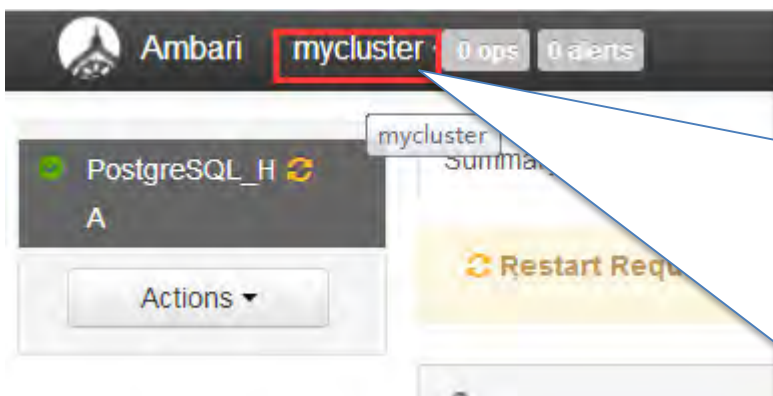
Start

Turn On Maintenance Mode

Delete



管理PG高可用集群： Ambari平台操作审计



0 Background Operations Running

Operations	Start Time	Duration	Show: All (10)
✓ Start Patroni_Server_Slave	Today 00:02	1.62 secs	100%
✓ Install Patroni_Server_Slave	Today 00:02	7.58 secs	100%
✓ Start Added Services	Today 00:00	9.57 secs	100%
✓ Install Services	Today 00:00	20.13 secs	100%
! Install Services	Tue Oct 17 2017 23:52	4.07 secs	100%
✓ Stop ZooKeeper Server	Tue Oct 17 2017 23:47	994 ms	100%
✓ Start Added Services	Tue Oct 17 2017 23:47	12.34 secs	100%
✓ Install Services	Tue Oct 17 2017 23:47	5.92 secs	100%

Do not show this dialog again when starting a background operation

所有启、停操作记录信息都能得以保存



管理PG高可用集群： Ambari平台用户权限

mycluster ROLES



可建立用户\组及配置集群访问控制权限

Roles ? Assign roles to these users Assign roles to these groups

Cluster Administrator

Cluster Operator

Service Administrator

Service Operator

Cluster User

Add Group

Ambari

- Clusters
 - mycluster
 - Roles
 - Go to Dashboard
 - Versions
 - Remote Clusters

Views

- Views
- View URLs

User + Group Management

- Users
- Groups

Users / Create Local User

Username

Type **Local**

Status **Active**

Ambari Admin No

Password

Cancel Save



管理PG高可用集群： 监控

集群监控信息

mycluster 0 ops 0 alerts

Dashboard Services Hosts Alerts Admin

Metrics Heatmaps Config History

Metric Actions Last 1 hour

Memory Usage

1.8 GB

Network Usage

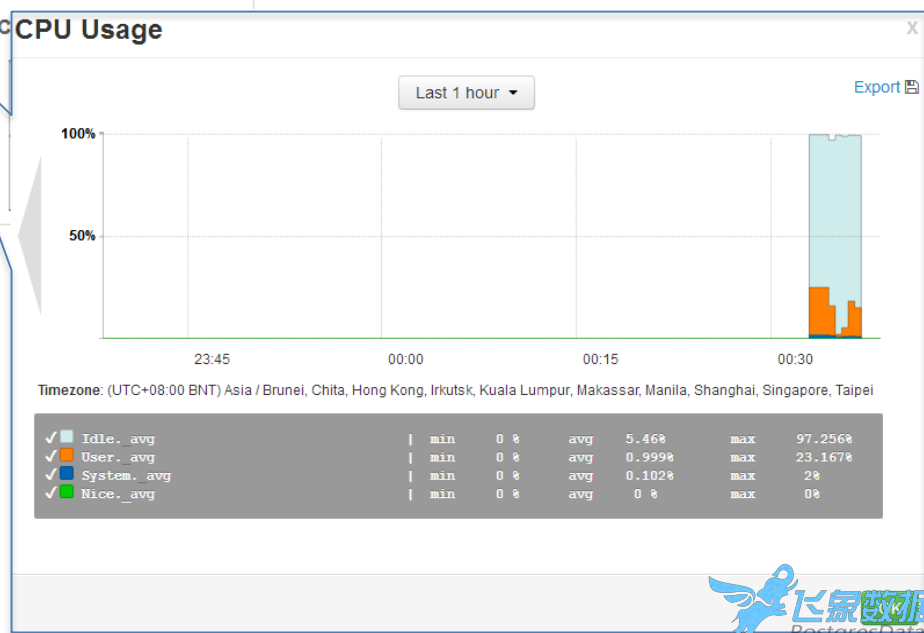
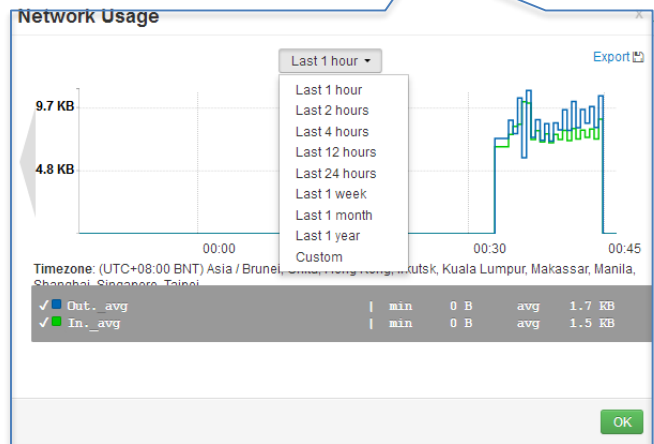
4.8 KB

CPU Usage

100%

CPU Usage

100%





管理PG高可用集群： 监控

节点监控

www.hadoop02.zgw
← Back

Summary Configs Alerts 0 Versions Host Actions

Components + Add

Host needs 5 components restarted Restart

- HAproxy_Server / PostgreSQL_HA Started
- Metrics Collector / Ambari Metrics Started
- Patroni_Server_M... / PostgreSQL_HA Started
- etcd_Server / PostgreSQL_HA Started
- keepalived_Server / PostgreSQL_HA Started
- Metrics Monitor / Ambari Metrics Started
- Clients / PostgreSQL_install Installed

Host Metrics Last 1 hour

CPU Usage 27.9 GB
18.6 GB
9.3 GB

Load 1.8 GB
953.6 MB

Network Usage 14.6 KB
9.7 KB
4.8 KB

Memory Usage 200
100

Summary

Hostname: www.hadoop02.zgw
IP Address: 192.168.120.131
Rack: /default-rack
OS: centos7 (x86_64)
Cores (CPU): 2 (2)
Disk: 7.21GB/32.97GB (21.87% used)
Memory: 2.07GB



更多功能，敬请期待



Thanks!

