

hosted by  **Alibaba** Group
阿里巴巴集团

APACHE
HBASE 

HBase at China Telecom

Chen Ze

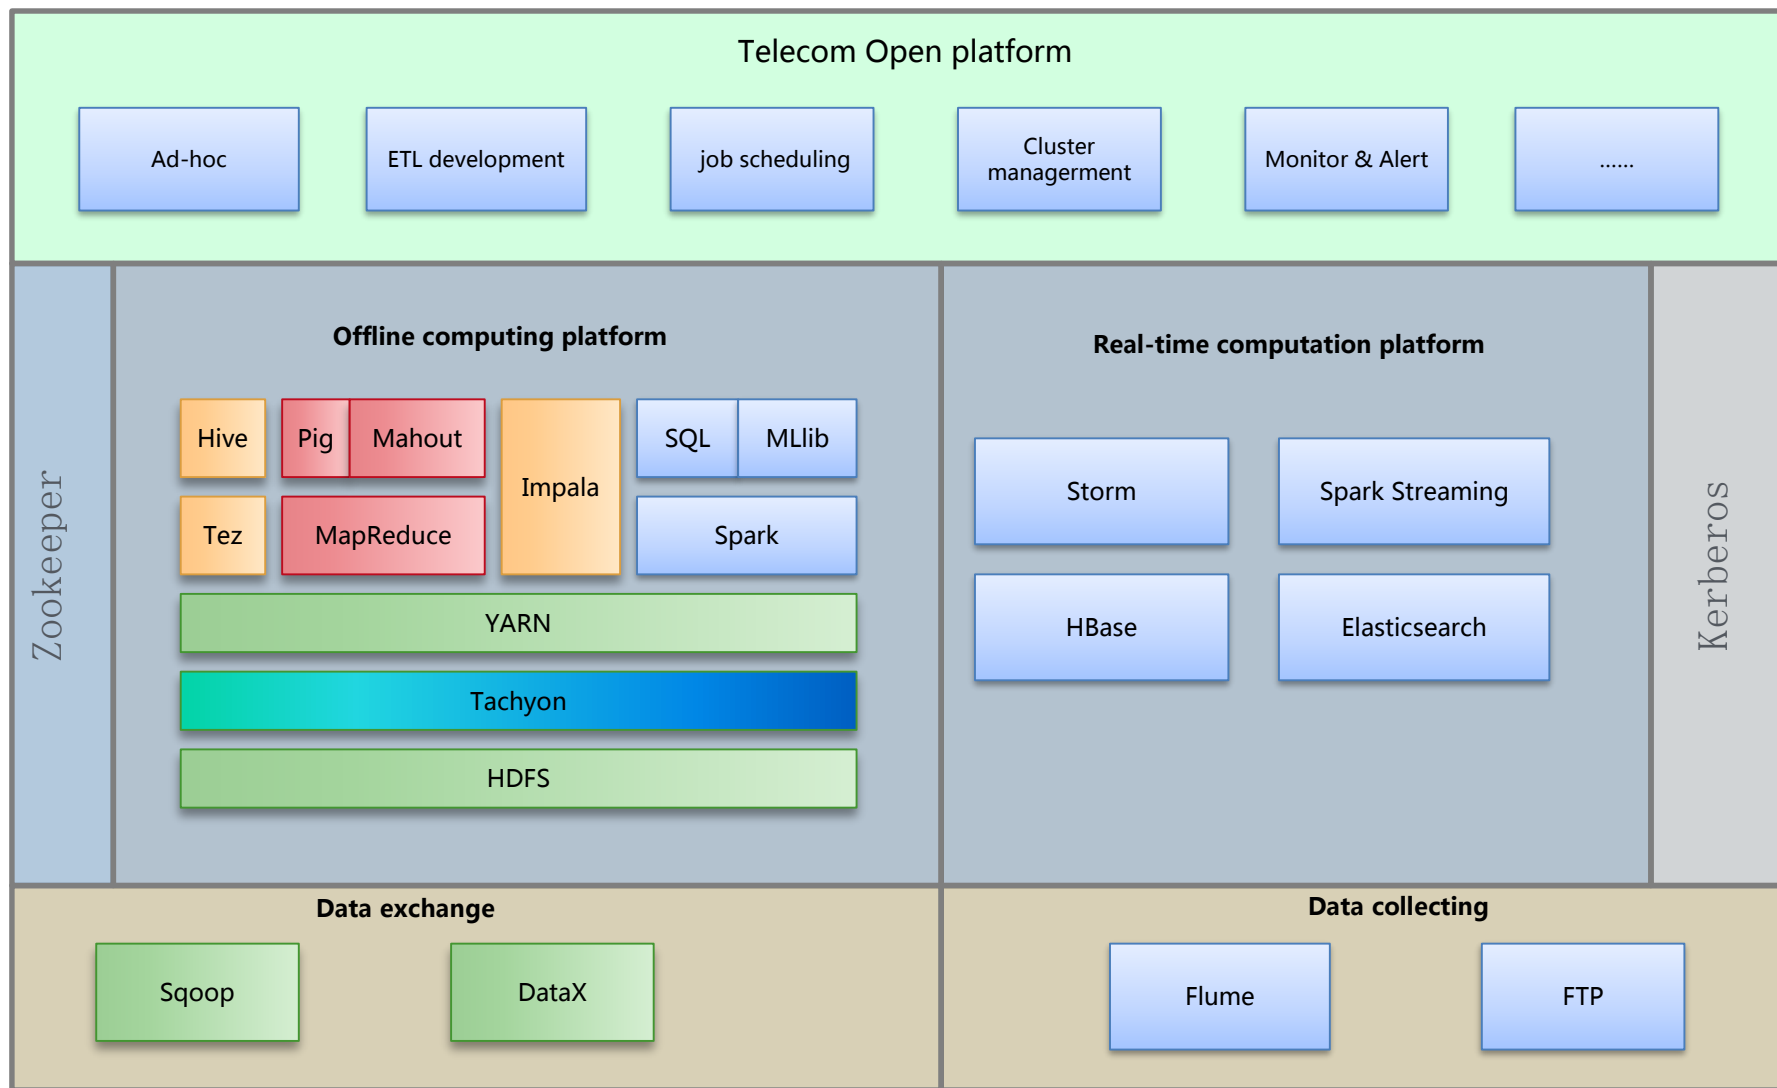
Content

- 01 Infrastructure & Application
HBase Usage at China Telecom
- 02 Monitor & Optimization
How do we monitor HBase and optimization
- 03 Q&A

01 Infrastructure & Application

HBase usage at China Telecom

Infrastructure



China Telecom HBase Platform

- 322 hosts in a independent HDFS cluster ,32 cores, 256GB memory, 3.6T*12 disk
- 6 HBase clusters with different kinds of application
 - Persistence for Streaming Jobs
 - Online writing/reading
 - Kylin support
- 520 TB data , 1 TB/day
- HBase1.2.0--CDH 5.12.1

Persistence for Streaming Jobs

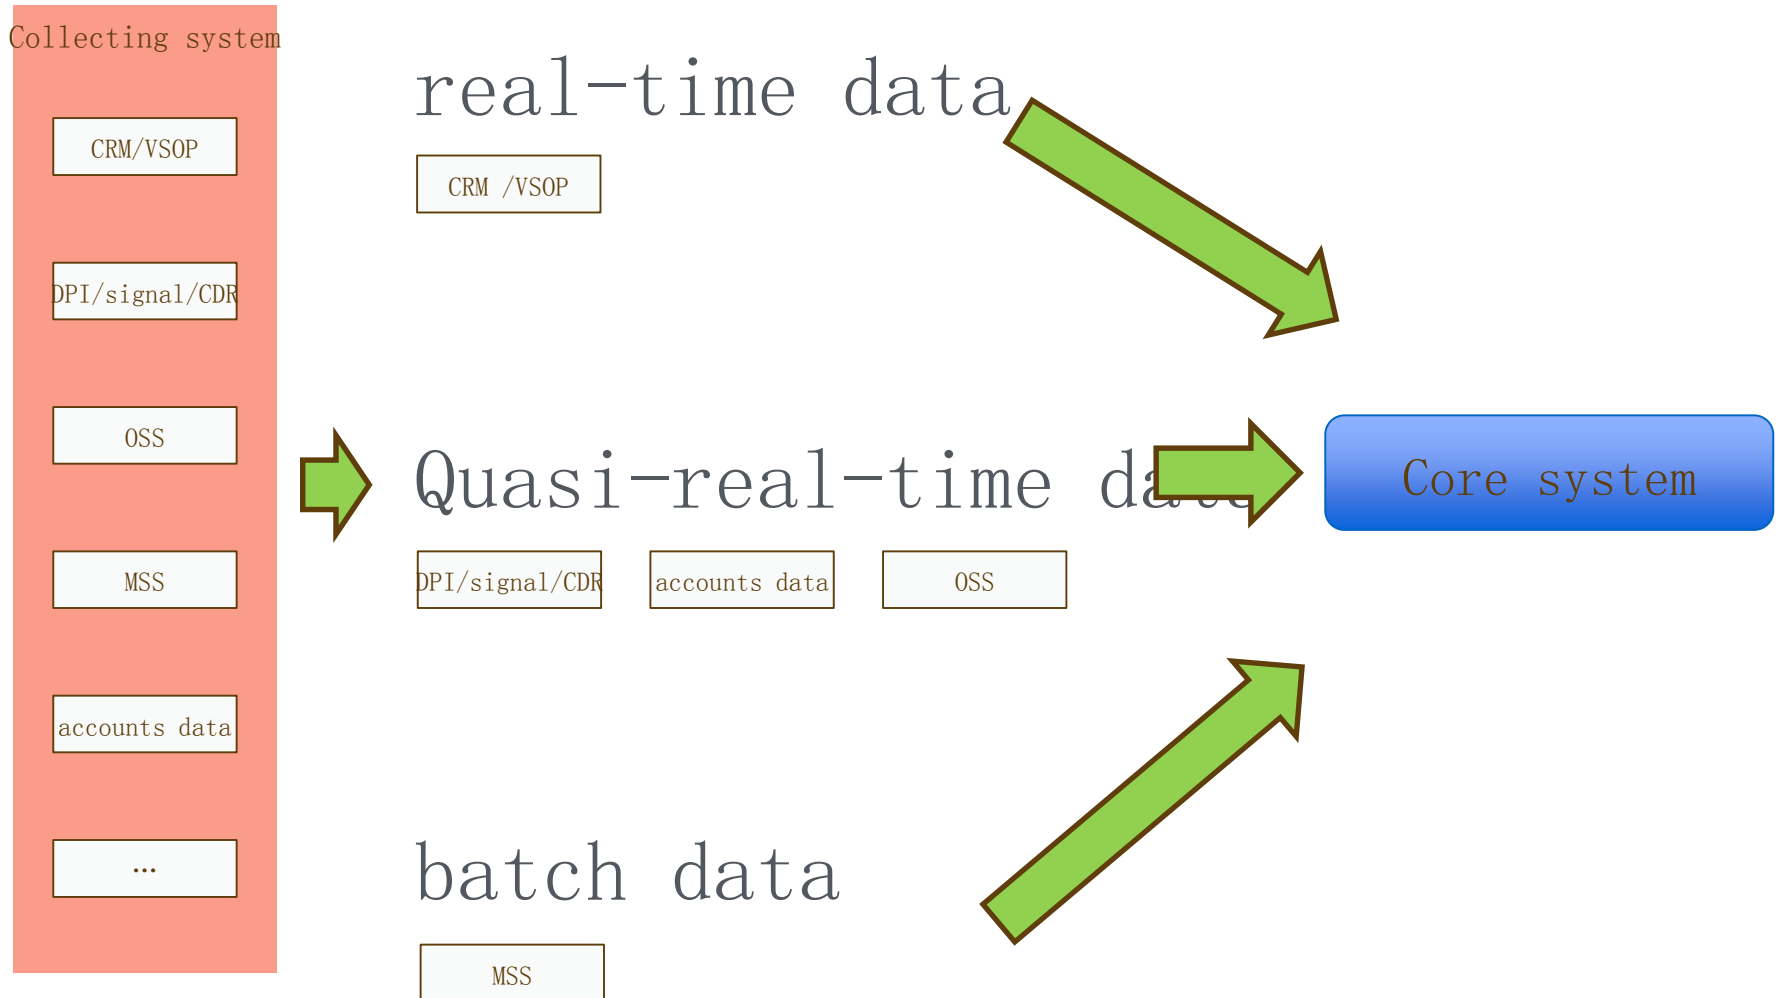
hosted by



- core system
- data collecting system

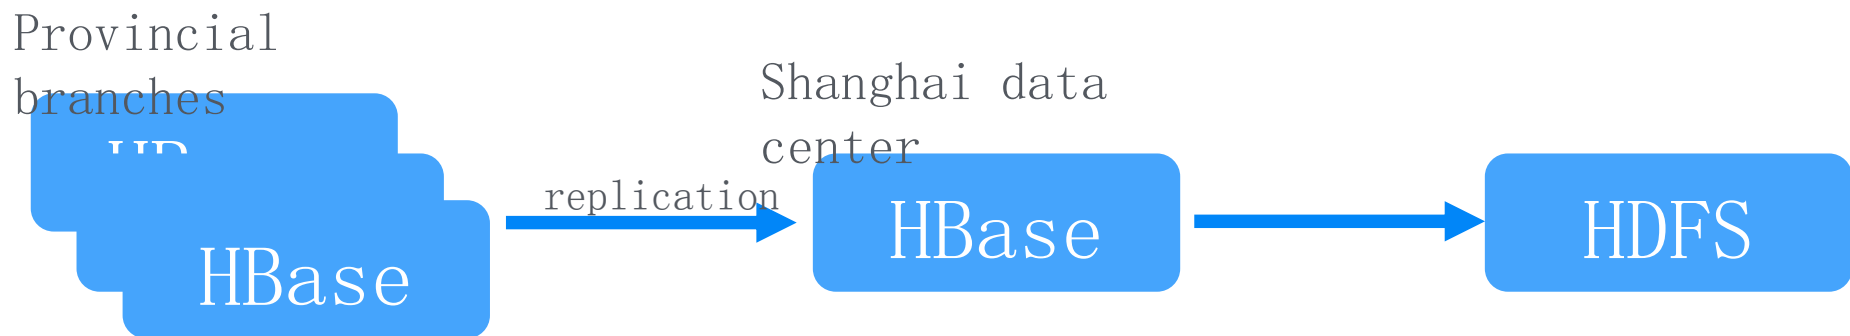
Data collecting system

Collect different kinds of data with different kinds of collecting method



Data collecting system


OSS data collecting use HBase replication to receive data from 31 Provincial branches



Core System

- core system use HBase to store middle layer data and Aggregation layer data through Spark Streaming

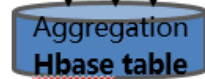
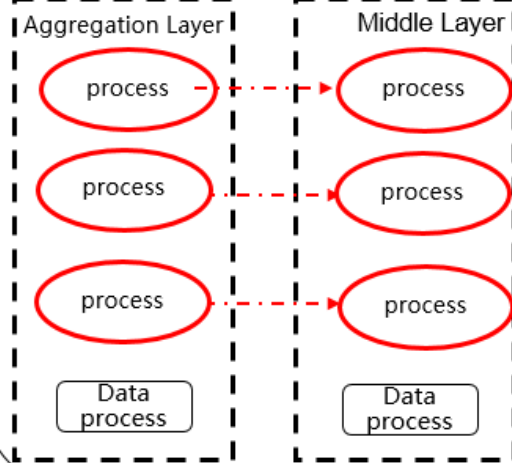
Collecting system



Kafka





Spark Streaming



Schedule Kafka



dispatch



Core system

Online writing/reading application

hosted by



- DPI data service
- Location-based data service

DPI data service

- How to use DPI data ?

月盘点 专题分析 经营管控 运营支撑 营销支撑 报表中心 服务支撑 客户视图 渠道协同 针对性营销

用户管理 用户管理 客户视图

客户名片-详细+

用户名: 李凌志	证件类型: 身份证	证件号码: 220182198204020212
时间: 20150527	地市: 长春	县区: 二道区
客户群:	联系人: 李凌志	证件地址:
联系地址: 自由大路5899号	业务类型: 移动后付	(宽带)端口速率:
手机号码: 18904300663	城市描述:	用户状态:
渠道类别: 10	渠道状态:	渠道生效时间: 2013-06-09
渠道:	入网时间: 2009-09-18	在网时长: 68.3
第一次通话时间:	入网方式: 购机入网	套餐基本费: 59
宽带包月时长:	集团名称: 5002891380	最近一次停机时间: 2011-07-26
套餐名称: (87000146)乐享		

客户特征+

- 爱沟通
- 爱近球
- 爱旅游

综合评估+

客户营销度
此客户营销度为66%

活跃 -
价值 -
稳定 -



66%

客户健康度
此客户健康度为8%

流失风险 -
价值提升 -



8%

消费变化-表格+

账期	用户消	本地	长途	漫游	短信	月租	赠款	当月	手机	彩信	其他	
1	201510	80	56	29	22	24	6	31	52	78	3	51
2	201509	81	48	79	88	92	99	78	29	82	13	3
3	201508	14	57	98	54	6	28	60	5	73	62	11
4	201507	46	7	25	16	3	2	45	9	3	73	35
5	201506	79	57	62	29	35	41	69	21	21	79	59

消费变化-折线+

消费变化——折线图



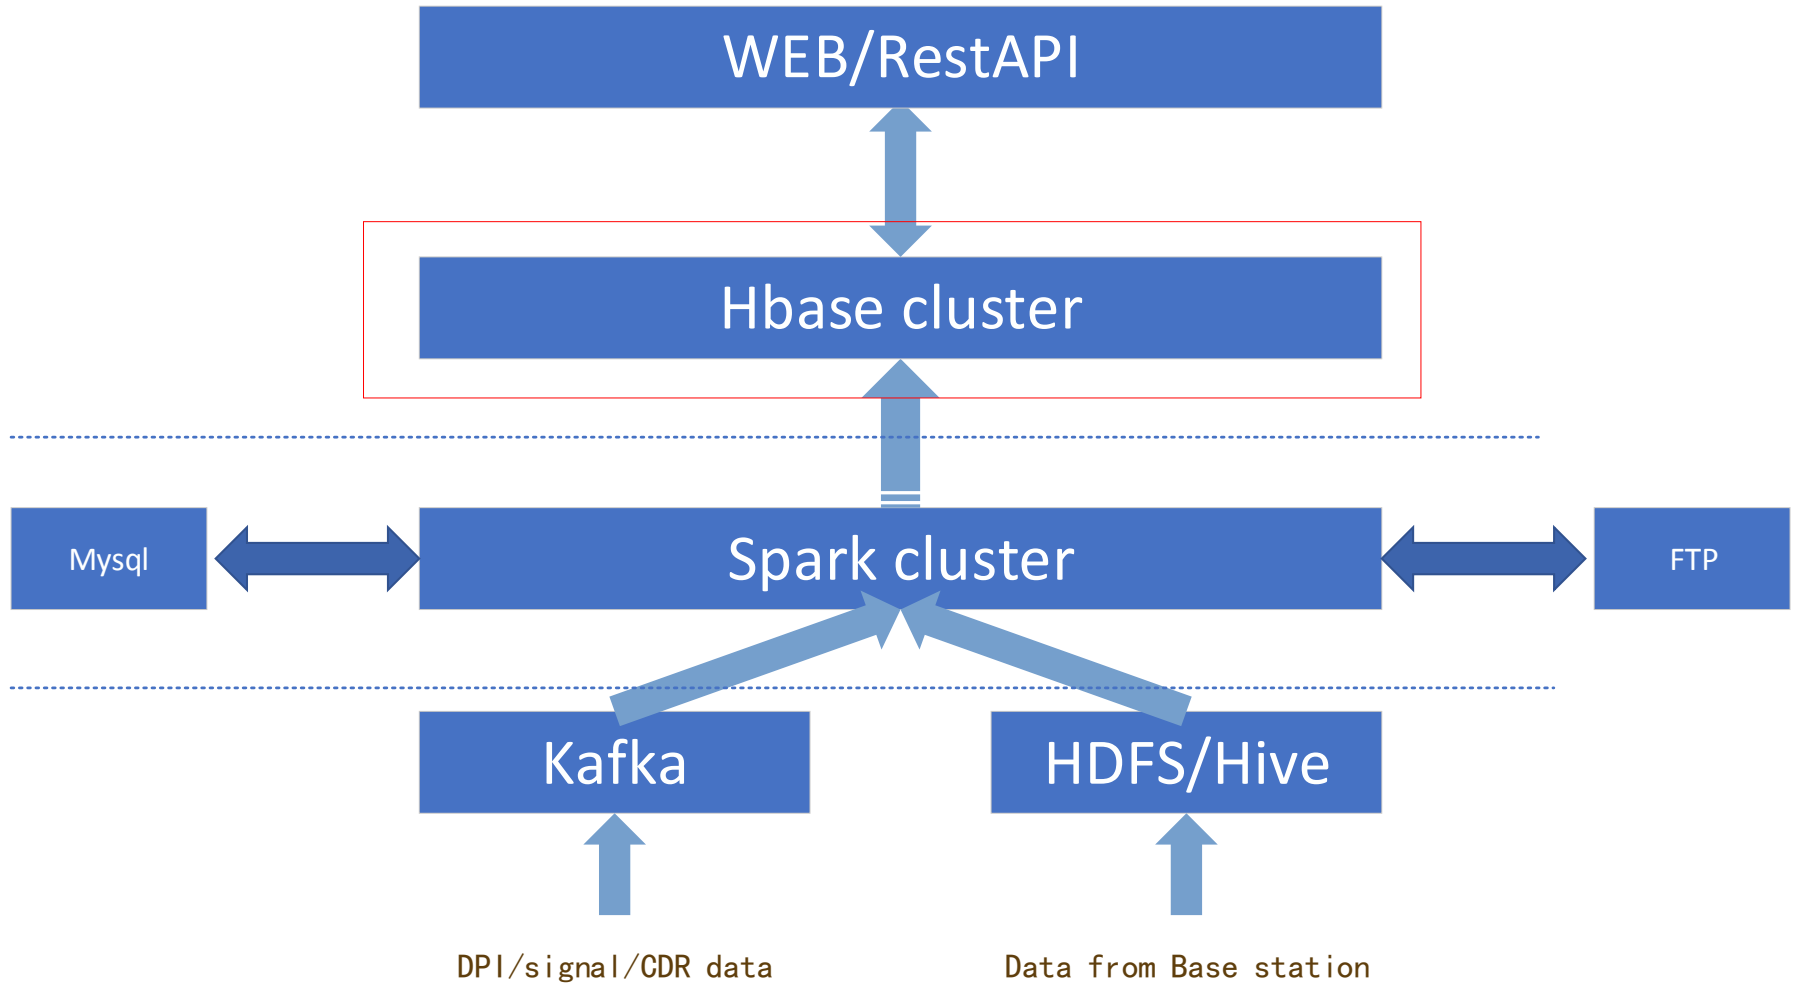
消费变化-饼图+



76% ↑ 2.9K/s +

↓ 4.3K/s -

Location-based data service



• Location-based data service

wzfwcontroller : 位置服务相关api

Show/Hide | List Operations | Expand Operations

GET	/add	test
GET	/current-city	个人用户当前城市查询功能
GET	/current-location	个人用户当前位置查询功能
GET	/history-location	个人用户历史位置查询功能
GET	/history-track	个人用户历史轨迹查询功能
GET	/residence-location	居住地分析
GET	/roaming-destination	漫游地查询功能
GET	/time-interval	时间段活动地分析
GET	/weekend-location	周末活动地分析
GET	/work-location	工作地分析

02 Monitor & Optimization

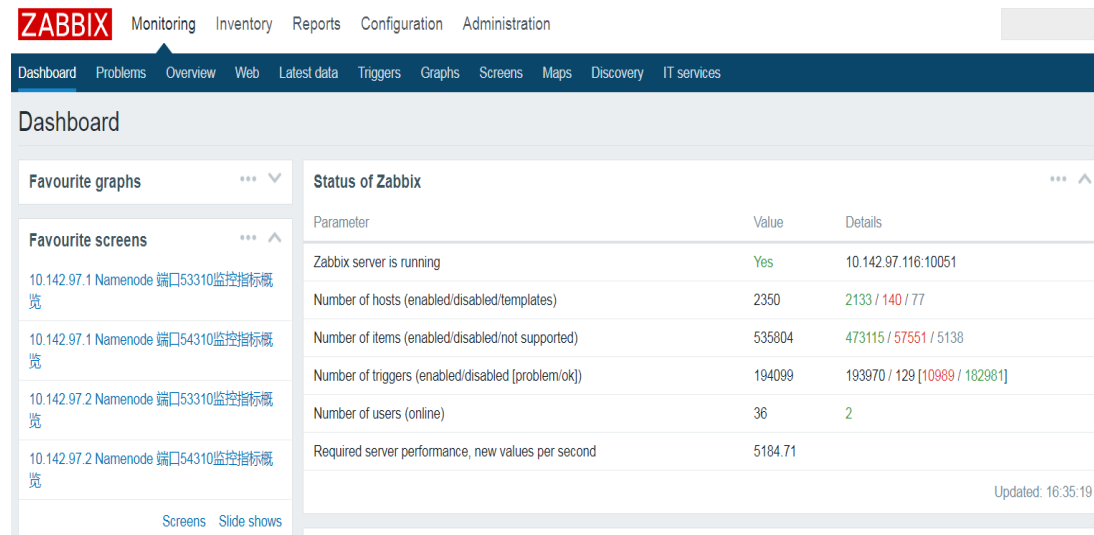
How do we monitor HBase and optimization

HBase Monitor

Tools

ganglia-3.6.0
Zabbix-3.2.4

- Ganglia for HBase basic metrics
- Zabbix for important items to alert



ZABBIX Monitoring Inventory Reports Configuration Administration

Dashboard Problems Overview Web Latest data Triggers Graphs Screens Maps Discovery IT services

Dashboard

Favourite graphs ... v

Favourite screens ... ^

- 10.142.97.1 Namenode 端口53310监控指标概览
- 10.142.97.1 Namenode 端口54310监控指标概览
- 10.142.97.2 Namenode 端口53310监控指标概览
- 10.142.97.2 Namenode 端口54310监控指标概览

Screens Slide shows

Status of Zabbix ... ^






Parameter	Value	Details
Zabbix server is running	Yes	10.142.97.116:10051
Number of hosts (enabled/disabled/templates)	2350	2133 / 140 / 77
Number of items (enabled/disabled/not supported)	535804	473115 / 57551 / 5138
Number of triggers (enabled/disabled [problem/ok])	194099	193970 / 129 [10989 / 182981]
Number of users (online)	36	2
Required server performance, new values per second	5184.71	

Updated: 16:35:19

CPU's Total: **8672**
Hosts up: **282**
Hosts down: **0**

Current Load Avg (15, 5, 1m):
1%, **1%**, **1%**
Avg Utilization (last hour):
1%

Utilization heatmap



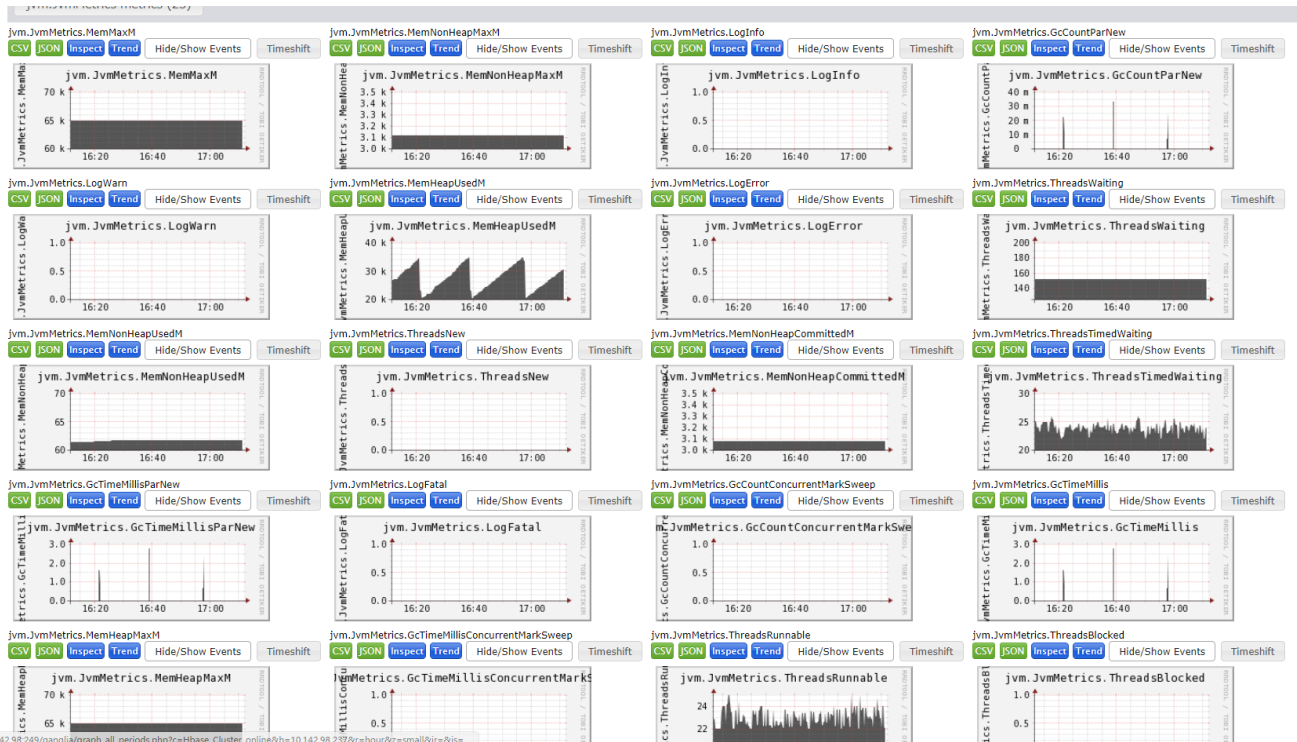
Stacked Graph - load_one

HBase ganglia configuration

Edit HBase config file `hadoop-metrics2-hbase.properties`

```
*.source.filter.class=org.apache.hadoop.metrics2.filter.RegexFilter
*.record.filter.class=${*.source.filter.class}
*.metric.filter.class=${*.source.filter.class}
hbase.sink.ganglia.metric.filter.exclude=^(.*table.*)|(\w+metric)|(\w+region\w+)|(. *Replication.*)|(Balancer\w+)|(\w+Assign\w+)|(\w+percentile)|(\w+max)|(\w+median)|(\w+min)|(MetaLog\w+)|(\w+WAL\w+)|(. *bigdata.chinatelecom.cn.*)$
hbase.sink.ganglia.class=org.apache.hadoop.metrics2.sink.ganglia.GangliaSink31
hbase.sink.ganglia.tagsForPrefix.jvm=ProcessName
*.sink.ganglia.period=20
```

Then we can see many metrics items



• Important metrics

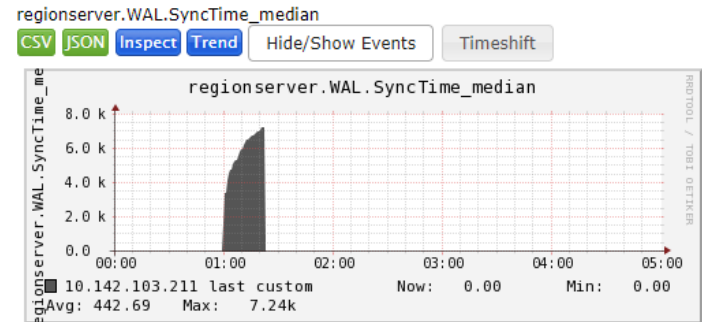
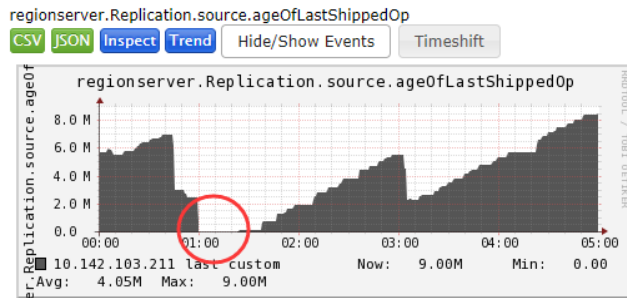






Type	Items
System status	ping Available memory in percent Disk usage
HDFS status	NameNode port DataNode port JournalNode port HDFS available space
HBase status	HBase Master Port HBase Region Server Port
Zookeeper status	ZooKeeper Port ZKFC Port
HBase RPC	regionserver.TotalCallTime regionserver.ProcessCallTime regionserver.QueueCallTime regionserver.numActiveHandler regionserver.ipc.numCallsInGeneralQueue regionserver.ipc.numOpenConnections regionserver.RegionServer.numCallsInWriteQueue regionserver.RegionServer.numCallsInReadQueue
Hbase IO	regionserver.Server.Mutate_99th_percentile regionserver.wal.SyncTime_99th_percentile regionserver.server.Get_99th_percentile regionserver.server.ScanTime_99th_percentile

• Important metrics

Type	Items
HBase region	Region num, size BlockCache hit radio
JVM GC	jvm. JvmMetrics. GcTimeMillis jvm. JvmMetrics. GcCount GC log

HBase Debug case

- One day, our core data system hanged when trying to connection Hbase cluster...



HBase Optimization

- CMS vs G1
- Read/Write Splitting and some Optimization

CMS vs G1

```
"-Xms64g -Xmx64g -Xmn16g -XX:+UseParNewGC -XX:+UseConcMarkSweepGC \  
-XX:CMSInitiatingOccupancyFraction=65 -XX:+CMSParallelRemarkEnabled \  
-XX:SurvivorRatio=28 -XX:PermSize=3g -XX:MaxPermSize=3g \  
|
```

CMS  to G1

```
-server -Xms96g -Xmx96g -XX:PermSize=512m -XX:MaxPermSize=512m \  
-XX:+UseG1GC -XX:InitiatingHeapOccupancyPercent=30 -XX:G1MixedGCCountTarget=3 \  
-XX:G1ReservePercent=15 -XX:G1HeapWastePercent=5 \  
|
```

Why?

- Use replication
- Use Kerberos
- Optimization `hbase.ipc`

hosted by  **Alibaba** Group
阿里巴巴集团

A P A C H E
H B A S E 

Q&A

hosted by  **Alibaba** Group
阿里巴巴集团

APACHE
HBASE 

Thanks