Content

中国人寿 CHINA LIFE

01

Integration, Processing, Query, Export

Scenarios

中国人寿
CHINA LIFE

Scenarios

Overview
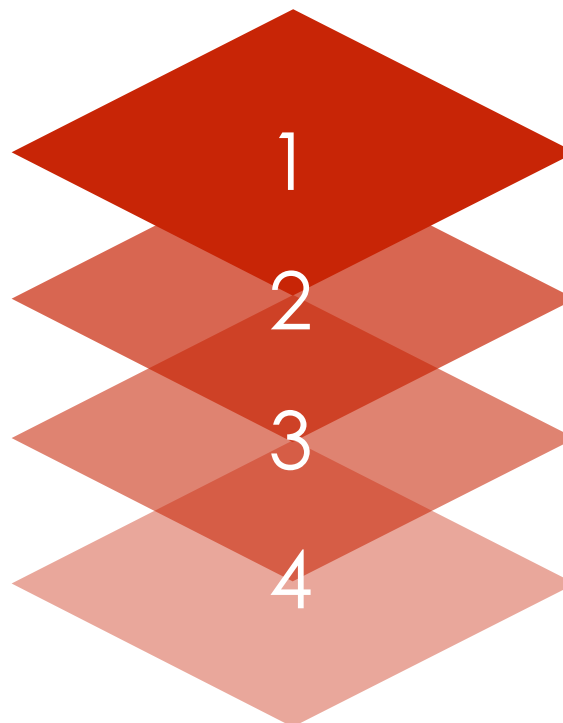
- Storage for integration
- Data source for query and processing

Integration

Easy update, schemaless, millions of columns

1

2

3

4

Processing

Fast batch R/W through snapshot, bulkload...

Query

High concurrency & Low latency

Export

Hive external table support

中国人寿
CHINA LIFE

Cluster

3 Clusters
200+ Nodes

Data

HBase data 300TB+
largest table 30TB+ , 2500 regions

Processing

Hundreds of MR/Hive/Spark jobs per day
50TB+ Incremental data for update&insert

Queries

Tens of millions of queries per day

中国人寿 CHINA LIFE

Integration

- Integrate policy, customer, agent… data from various systems

policy_detail	Agent_detail	Customer_detail	......

Business Systems	Sales management systems	ECIF	CallCenter	Customer contact system	......

Scenarios

Integration

- Data integrated by the key of business entity.
- One row for all information

Contract	Holder	Claim	Modification	income
Contract_no Agent_no Cust_no	Name Birthday	Claim_no Contract_no	Mtn_no Mtn_type	Contract_no Income type amount ......

......

Policy_detail (HBase Wide Table)

Contract_no row_key	Contract JSON	holder JSON	Claim1 JSON	Claim2 JSON	...	Mtn1 JSON	Mtn2 JSON	...	Income1 JSON	Income2 JSON	......

中国人寿 CHINA LIFE
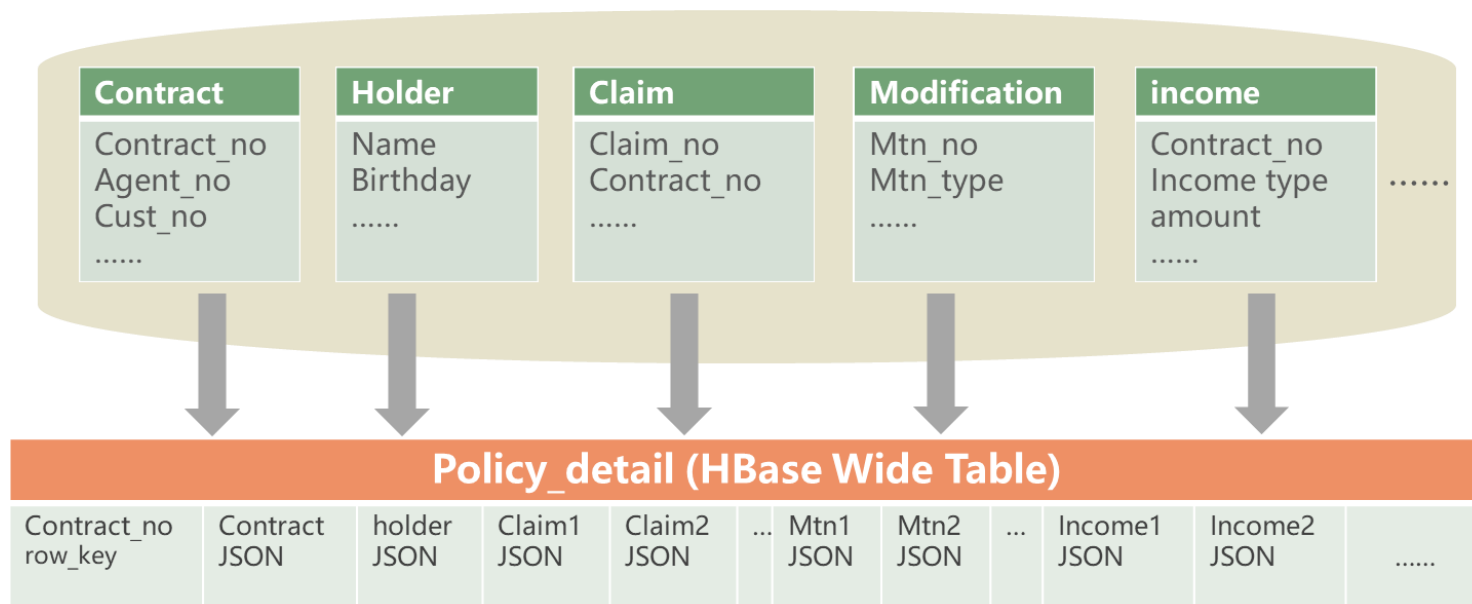
- RDBMS->Textfile->Hfile->BulkLoad to HBase

rowkey	income01	income02	income03	......
001	{sid:01,cntr_no:001,type:A,amount:100}	{sid:02, cntr_no:001, type:B, amount:100}	{sid:03, cntr_no:001, type:C, amount:100}	......

SSIS+MR
Sqoop+MR
SparkStreaming

Hfile

Textfile

001::income01::{sid:01,cntr_no:001,type:A,amount:100}
001::income02::{sid:02, cntr_no:001, type:B, amount:100}
001::income02::{sid:03, cntr_no:001, type:C, amount:100}
…

income			
sid	cntr_no	type	amount
01	001	A	100
02	001	B	100
03	001	C	100

....

中国人寿 CHINA LIFE

Scenarios
Processing

- Analysis of entities (customer,policy …) .
- Build indexes between entities.

| Policy_labels | Customer_labels | Agent_labels | Family Map | …… |

Customer Value Evaluation

Loss Ration | Blacklist label

Customer Risk Warning

Agent Risk Warning

Customer information quality

Labels For Machine Learning

Purchasing Power Prediction

Insurance Coverage Analysis

Family Map Recognition

Index of Customer to Contract

Index of agent to Customer& Contract

……

| policy_detail | Agent_detail | Customer_detail | …… |

Scenarios
Processing

- Processing hundreds of labels within one time I/O

Summary

rowkey	SumAmount	......
001	300	......

500+

MR/Spark

One MR for all labels

Detail

rowkey	income01	income02	income03	......
001	{sid:01,cntr_no:001,type:A,amount:100}	{sid:02, cntr_no:001, type:B, amount:100}	{sid:03, cntr_no:001, type:C, amount:100}	......

Scenarios
Processing

> • Encapsulated, Configurable, Sql-like development framework
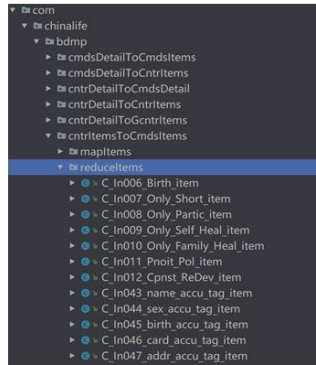
Tools: select | Where | aggregate | order | …

rowkey	income01	income02	income03	……
001	{sid:01,cntr_no:001, type:A,amount:100}	{sid:02, cntr_no:001, type:B, amount:100}	{sid:03, cntr_no:001, type:C, amount:100}	……

Development:
• Sql-like&Native API
• Off-line Unit Test
……

Runtime Configuration:
• Name of Source/target table
• Read Table/Snapshot
• Write Hive/Hbase
• Rowkey blacklist
• Init/Incr switch
• BulkLoad switch
• ColumnPrefixFilter switch
• Single Row Processing
• Single Label Processing
……

Scenarios
Query

> • Unified query services centering on entities such as customer

| Unified Customer View | Unified Agent View | Family View | …… |

| Policy_labels | Customer_labels | Agent_labels | Family Map | …… |

| policy_detail | Agent_detail | Customer_detail | …… |

Scenarios

Query

- Sql-like input
- One Service for All labels

Goal

Select
Name,......
From xx
Where age>=40
......

Input parameters

GetList

{" IDList":"1,......"
,"itemsStruct":{"name":"姓名"......}
,"Condition":{"age>=40"......}
......
}

Filter

Output parameters

{"resultCode":"1"
,"result":{
 itemsStruct":{"name":"姓名"......}
 ,itemResult":[
 {"name":"张三",......}
 ,......
 ]
 }
}

Export

hosted by **Alibaba** Group APACHE HBASE

- Exporting to Hive to support batch query and analysis

Hive

Table to table
Choose required Columns

Policy_labels   Customer_labels   Agent_labels   ......
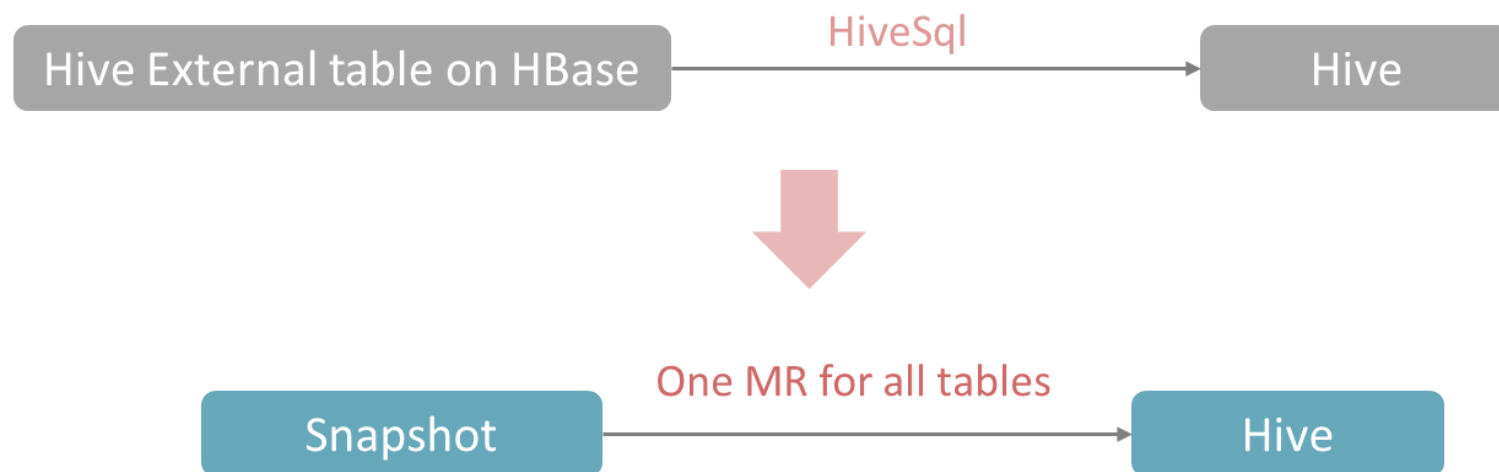
中国人寿
CHINA LIFE

Export

- Universal MR, read snapshot

| Hive External table on HBase | → HiveSql → | Hive |

⬇

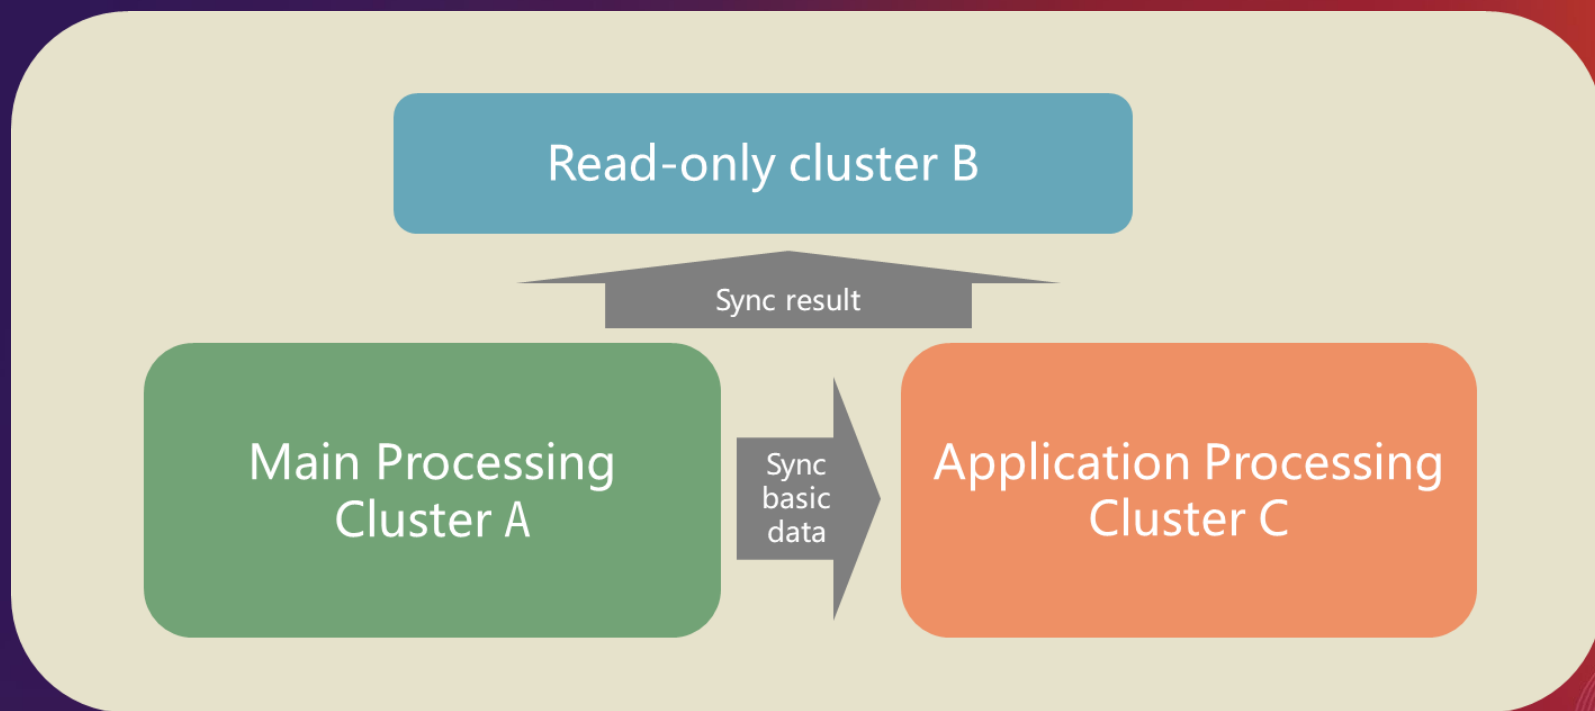| Snapshot | → One MR for all tables → | Hive |

02

Cluster, Configuration, Writing, Reading

Optimizations

Optimizations
Multiple Cluster

- Read/Write Splitting
- Platform/Application Splitting

Optimizations
Configuration

- Minimal impact on RS

- Region Balanced By Table

- G1GC

- Daily/Weekly MajorCompaction

- Pre-split Regions

- ......

Optimizations
Reading

- Minimal impact on RS

- Snapshot read

- ColumnPrefixFilter

- Data Blacklist

- Incremental read

- ......

Optimizations
writing

- Minimal impact on RS

- Write HFiles ->Do bulkLoad
- Skip WAL
- ......

03

Copy failed, Compact never end

Problems

Problems

Failed to copy tables

Processes

1. Create snapshot
2. Export snapshot
3. Disable table
4. Restore snapshot
5. Enable table
6. Delete snapshot

Problems

1. Failed to create snapshot
   (Wasn't complet in expectedTime)
2. Failed to export snapshot
   (File not found)

Problems

Failed to copy tables

Problems

1. Failed to create snapshot
   (Timeout/Wasn't complet in expectedTime)

Reasons

1. compacting/splitting/transiting while creating snapshot
2. Spending to much time on flushing

Solutions

1. Put required column into a small table.
2. Wait and retry.
3. Disable table.
4. Disable region balance temporarily.
5. skip_flush
6. Increase the timeout limit

Problems

Failed to copy tables

Problems

1. Failed to export snapshot
   (File not found)

```
Caused by: org.apache.hadoop.ipc.RemoteException(java.io.FileNotFoundException): File /datafs/.tmp/data/default/mytable/c48642fecae3913e0d09ba236b014667/info/7358baf78dda4c11973a7a7fef943fcc not found.
        at org.apache.hadoop.hdfs.web.JsonUtil.toRemoteException(JsonUtil.java:119)
        at org.apache.hadoop.hdfs.web.WebHdfsFileSystem.validateResponse(WebHdfsFileSystem.java:419)
        at org.apache.hadoop.hdfs.web.WebHdfsFileSystem.access$200(WebHdfsFileSystem.java:107)
        at org.apache.hadoop.hdfs.web.WebHdfsFileSystem$AbstractRunner.connect(WebHdfsFileSystem.java:595)
        at org.apache.hadoop.hdfs.web.WebHdfsFileSystem$ReadRunner.connect(WebHdfsFileSystem.java:1855)
        at org.apache.hadoop.hdfs.web.WebHdfsFileSystem$AbstractRunner.runWithRetry(WebHdfsFileSystem.java:673)
        ... 23 more
```

Reasons

1. compaction/split happend after snapshot creation
2. Bug in ExportSnapshot

Solutions

1. Disable table before copy.
2. Fix bug to search hfile correctly

Problems

Compact never ends

Problems

```
hbase(main):004:0> desc '...'
DESCRIPTION
'...', {NAME => 'ci', DATA BLOCK ENCODING => 'PREFIX TREE',
ERSIONS => '1', COMPRESSION => 'SNAPPY', MIN_VERSIONS => '0', TTL =>
 => '65536', IN_MEMORY => 'false', BLOCKCACHE => 'true', METADATA =>
1 row(s) in 1.2900 seconds
```

HBase / HBASE-12959

Compact never end when table's dataBlockEncoding using PREFIX_TREE

Details

				People	
Type:	Bug	Status:	RESOLVED	Assignee:	
Priority:	↑ Critical	Resolution:	Won't Fix	Reporter:	
Affects Version/s:	0.98.7	Fix Version/s:	None	Votes:	

∨ 🗨 stack added a comment - 06/Nov/17 06:17

   hbase-prefix-tree was removed in hbase2. Resolving as won't fix.

Solutions

Recreate table without this 'PREFIX_TREE'

中国人寿 CHINA LIFE

hosted by

- Table unavailable at the end of copy

   (disable->restore->enable)

- Region's locality drop to 0

- ......

Future Works

Pheonix

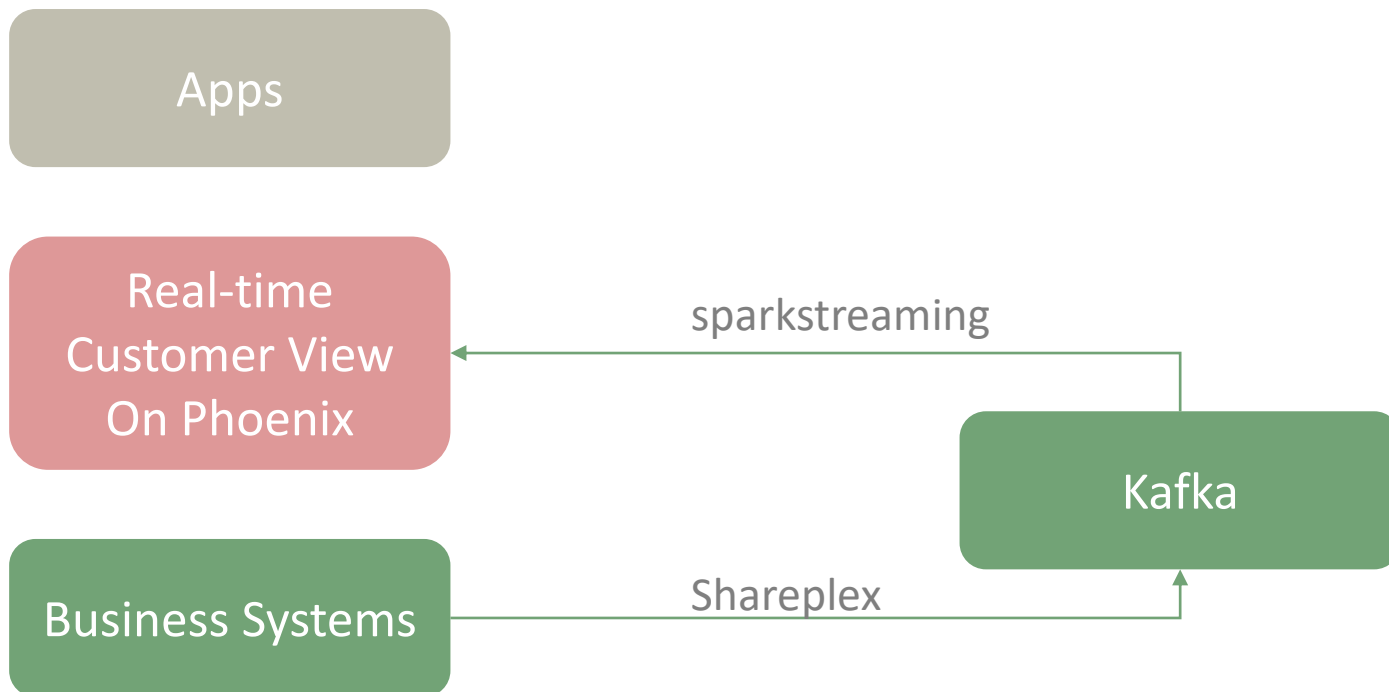- Flexible, real-time, precise query scenario

Apps

Real-time
Customer View
On Phoenix

Business Systems

Kafka

sparkstreaming

Shareplex

Future Works
Real-time

- Real-time detail tables and label tables

| Kudu+Impala | Policy_labels | Customer_labels |

| HBase | Policy_detail | Policy_labels | Customer_labels |

SparkStreaming

policy detail process | Policy_label process | Customer_label process

Kafka

Customer labels topic

Policy labels topic

Policy detail topic

Business data topic

中国人寿
CHINA LIFE