



# 百度Elasticsearch大数据分析 实践

百度大数据部 高攀  
2015年10月17日

# 大纲

- 背景介绍
- 典型应用场景
- 遇到的问题及经验分享
- 对ES的优化与改进
- 后期计划

# ES是啥

open-source

RESTful  
API

JSON  
over HTTP

scales  
massively

full text search

Elasticsearch

Lucene  
based

high  
availability

real time,  
search and  
analytics engine

distributed

schema  
free

multi  
tenancy

# 有谁在用

- 面向搜索 ( for full text search )

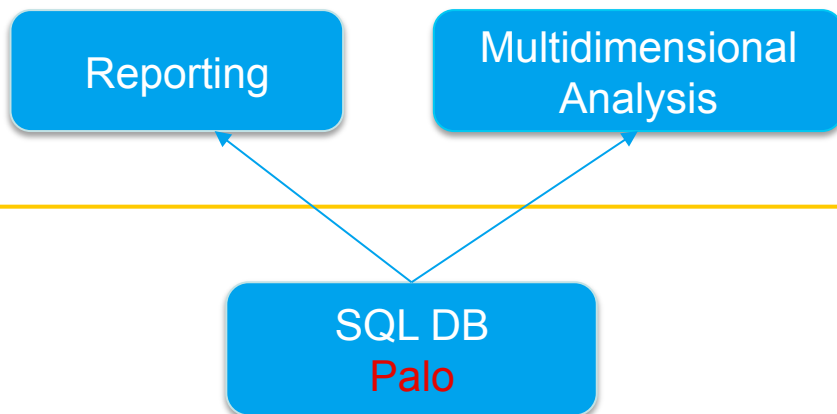
- Wikipedia
- Github
- Quora
- Facebook

- 面向分析 ( for analytics )


- Goldman Sachs
- Ubnt
- Foursquare (LBS)
- Linkedin
- Netflix

# ES在百度的使用定位 — OLAP

结构化数据分析



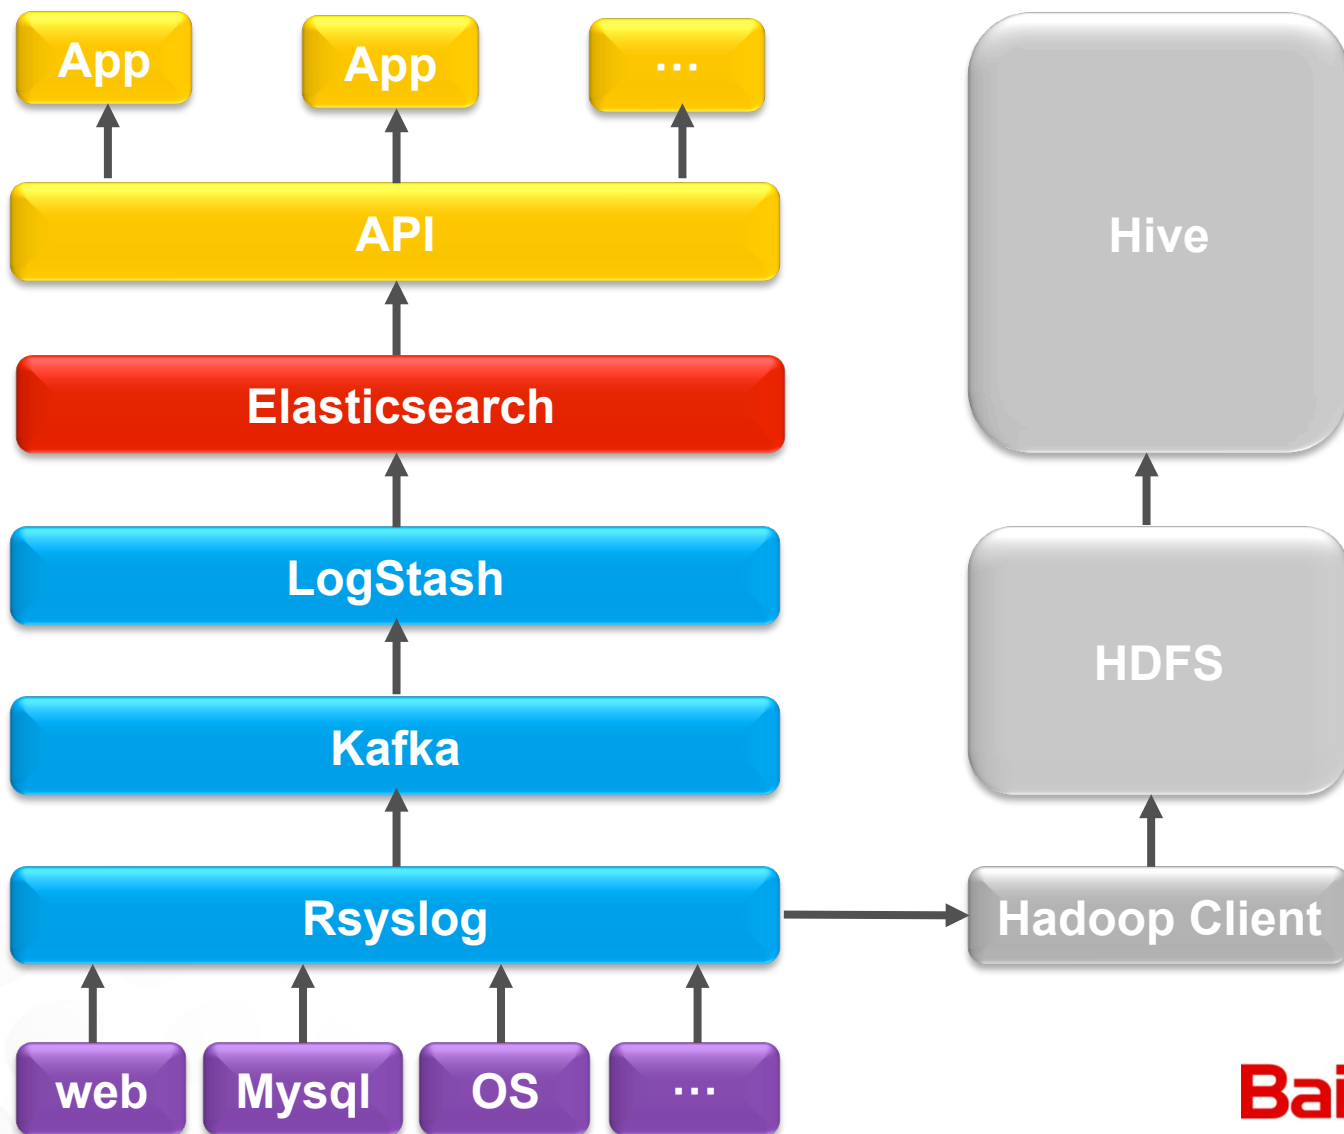
文本数据分析



# ES在百度现状

- 2013年10月开始使用
- 目前覆盖百度内部20多个业务线
  - 包括casio、云分析、网盟、预测、文库、直达号、钱包、风控等
- 单集群每天导入30TB+数据，总共每天60TB+
- 单集群最大100台机器，200个ES节点
- 共使用近300台机器，启动500+ES节点

# 典型应用场景一：云分析业务



# 典型应用场景一：云分析业务

## ● 主要挑战

- 大量文本数据的分词及建立索引
- 数万用户，索引分片过多，元数据过大
- fielddata占用大量内存，容易OOM
- 任意多维度关键词聚合查询秒级返回



# 典型应用场景一：云分析业务

## ● 解决方案

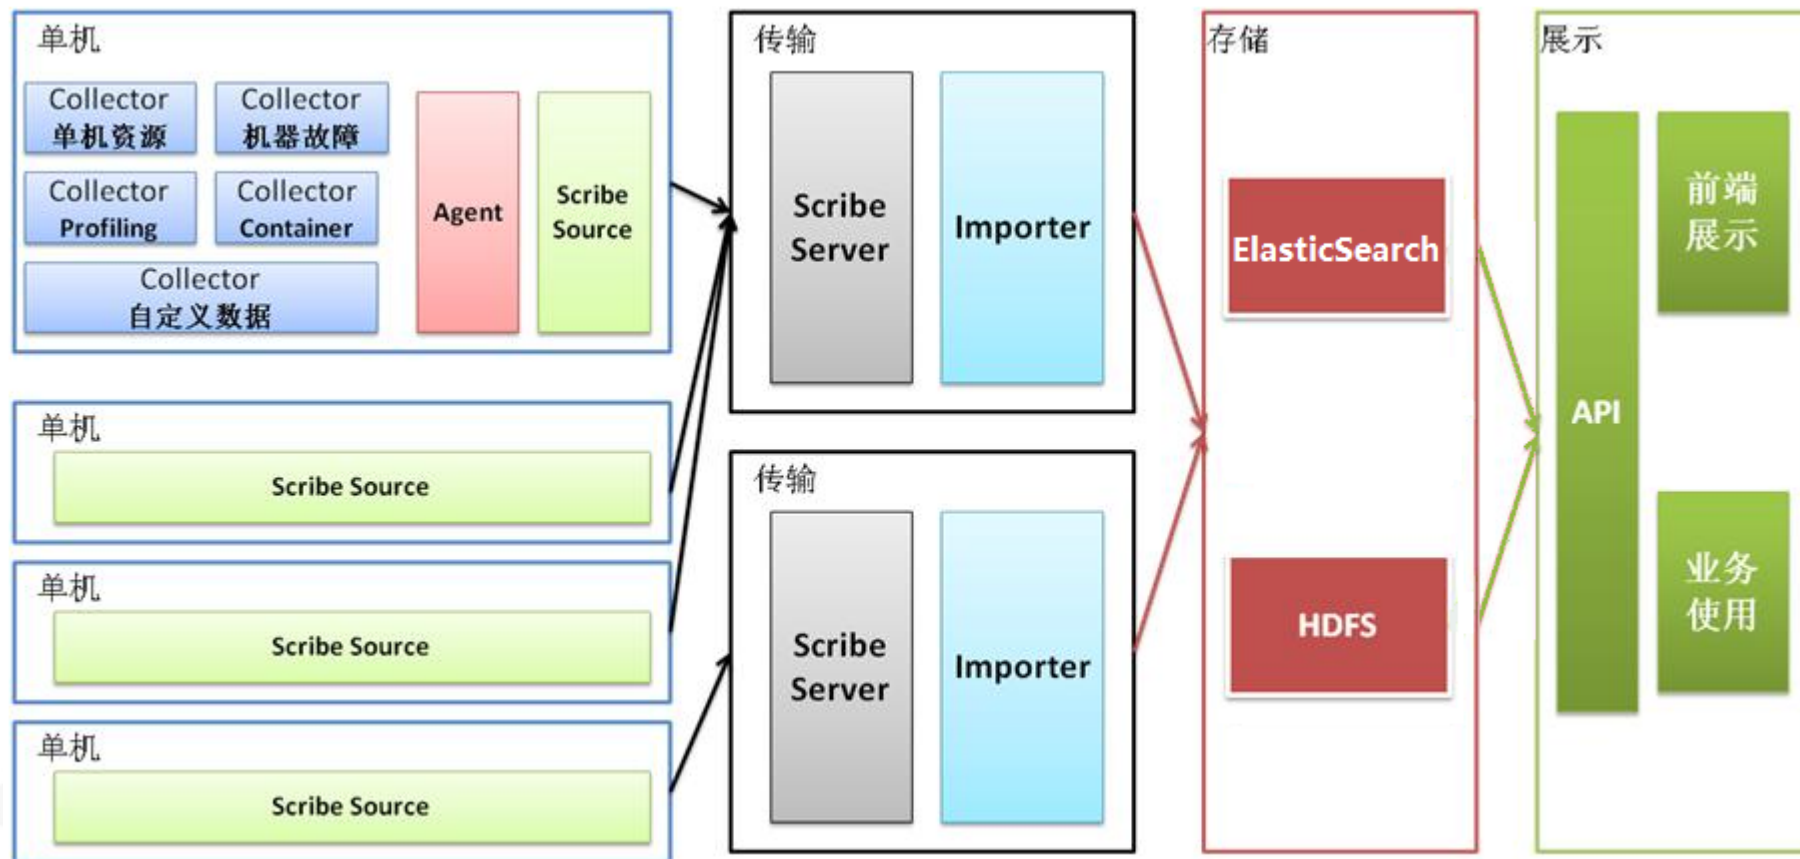
- 根据索引大小分别设置分片数，充分利用type合并索引
- 除分词字段外，其他字段全部存储为doc value
- master node、data node、client node 分离部署
- 保守设置fielddata内存占用软硬限，及其他内存占用限制
- 设置fielddata有效期

# 典型应用场景一：云分析业务

```
"persistent": {  
  "indices": {  
    "breaker": {  
      "total": {  
        "limit": "30%"  
      },  
      "fielddata": {  
        "limit": "20%"  
      },  
      "request": {  
        "limit": "10%"  
      }  
    }  
  },  
}
```

```
indices.fielddata.cache.size: 20%  
indices.fielddata.cache.expire : "48h"  
indices.cache.filter.size : 10%
```

## 典型应用场景二：casio业务



## 典型应用场景二：casio业务

### ● 主要挑战

- 字段不确定
- 数据量较大，每天30TB+数据，24小时不间断导入
- 小时级任意维度聚合分析毫秒级返回，天级秒级返回
- 集群规模较大（100台机器），机器较旧，宕机为常态

# 典型应用场景二：casio业务

## ● 解决方案

- 使用动态mapping自动匹配未知字段
- 数据分发到所有节点批量导入
- 全部使用doc value存储，减少内存消耗
- 使用模板，分天级、小时级自动创建索引
- SSD与SATA分组，冷数据定期自动迁移

## 典型应用场景二：casio业务

```
"casio_machines": {
  "order": 0,
  "template": "casio-machines-*",
  "settings": {
    "index.translog.flush_threshold_size": "1g",
    "index.refresh_interval": "10s",
    "index.routing.allocation.include.group": "ssd",
    "index.number_of_replicas": "1",
    "index.number_of_shards": "93",
    "index.translog.flush_threshold_period": "20m",
    "index.routing.allocation.total_shards_per_node": "3"
  },
  "mappings": {
    "_default_": {
      "_source": {
        "enabled": false
      },
      "dynamic_templates": [
        {
          "kvs": {
            "mapping": {
              "index": "no",
              "store": false,
              "doc_values": true,
              "type": "{dynamic_type}"
            },
            "path_match": "*_kvs.*"
          }
        ]
      }
    }
  }
},
```

POST /casio-machines-20151017/all  
{  
 "instance":1001,  
 "double\_kvs":{"cpu":35.6},  
 "long\_kvs":{"mem":12}  
}

```
"properties": {
  "float_kvs": {
    "properties": {
      "cpu": {
        "type": "double",
        "index": "no",
        "doc_values": true
      }
    }
  },
  "instance": {
    "type": "long"
  },
  "long_kvs": {
    "properties": {
      "mem": {
        "type": "long",
        "index": "no",
        "doc_values": true
      }
    }
  }
}
```

# 典型应用场景三：网盟DMP业务

## ● 业务简介

- 把百度内部数据（用户特征、历史浏览以及搜索数据等）开放给广告主，让广告主更好的定义精准受众，协助广告主制定投放策略，提升营销效果。

## ● 主要挑战

- 用户标签千万级别，相当于每张表数千万字段
- 每次请求涉及数千维度的组合过滤，数亿条数据的分组聚合
- 高并发，QPS 100以上
- 要求平均响应时间为秒级

# 典型应用场景三：网盟DMP业务

## ● 解决方案

- 通过ES的嵌套文档类型，将用户标签属性由key转换为value

qt: "1001":{"date":20150805,"freq":1},"1002":{"date":20150806,"freq":2}

hct: "1002":{"date":20150806,"ag":25}

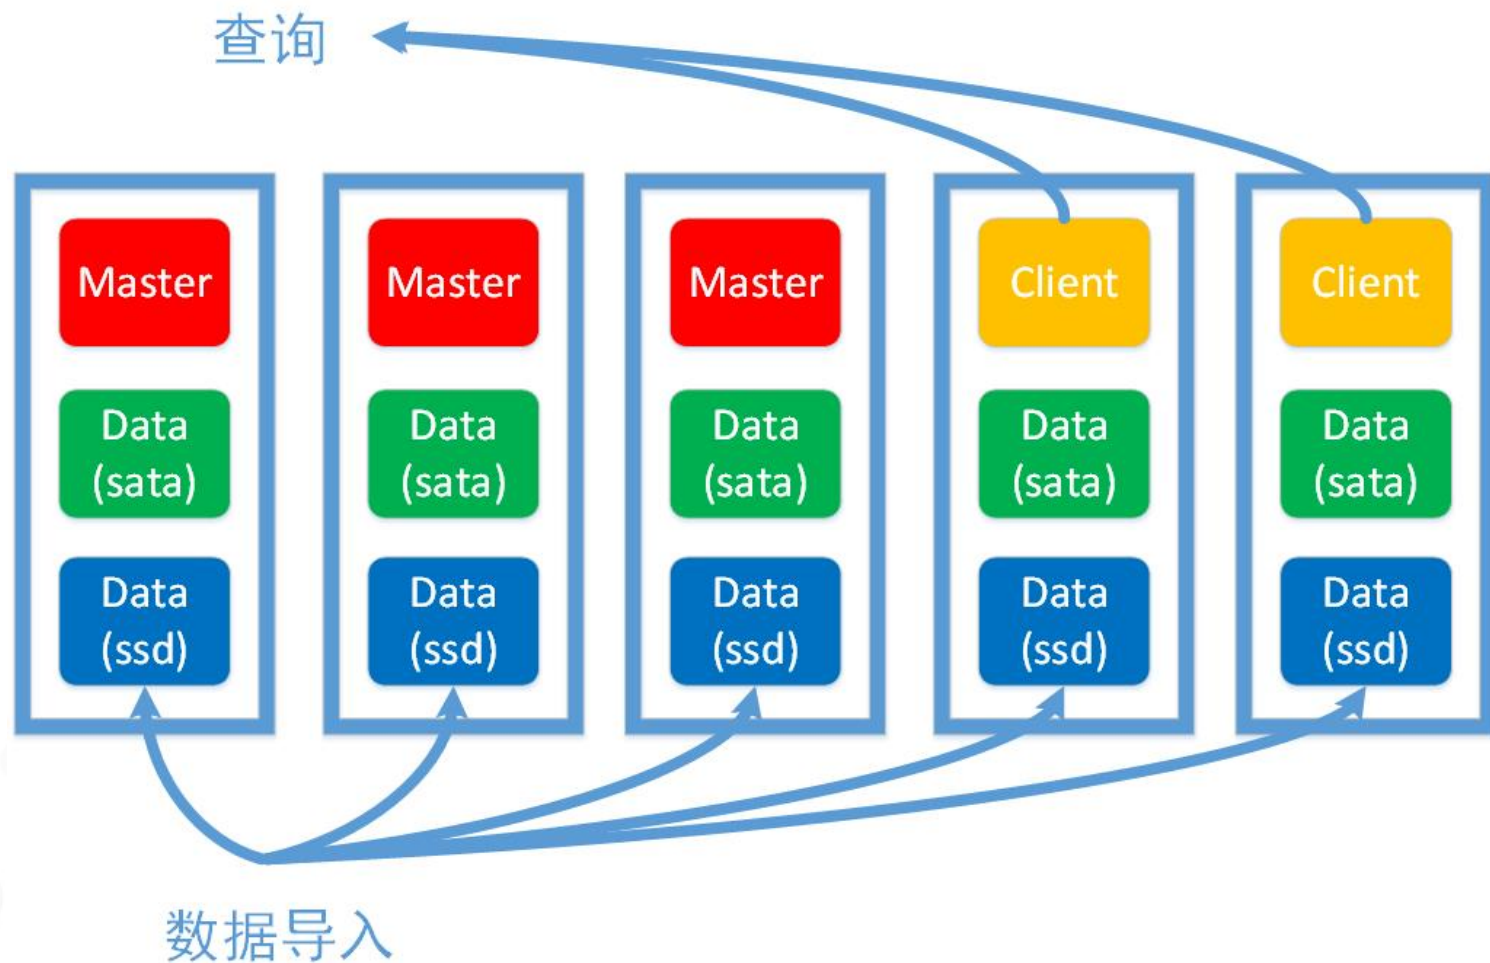
dmp: {"qt":[{"key":"1001", "date":20150805,"freq":1},{"key":"1002",  
"date":20150806,"freq":2}],"hct":[{"key":"1002","date":20150806,"ag":25}]}

- 配置total\_shards\_per\_node参数，最大化均衡分片分布
- 查询请求负载均衡到集群所有节点



# 遇到的问题及经验分享

## ● 集群规划



# 遇到的问题及经验分享

## 索引规划

- 根据机器数，磁盘数，索引大小等设置分片数，单个分片最好不超过10GB
- 配置total\_shards\_per\_node参数，限制每个index每个节点最多分配多少个分片


## 内存溢出

- 尽量使用doc value存储
- 保守配置内存限制参数
- 查询时限制size、from参数

# 遇到的问题及经验分享

- 提升导入性能

- 增大refresh及flush刷新间隔



# 遇到的问题及经验分享

## ● 提升导入性能

- 使用SSD盘时增大索引限制：

`"indices.store.throttle.max_bytes_per_sec" : "200mb"`

- 适当提高bulk队列：

`"threadpool.bulk.queue_size" : 1024`

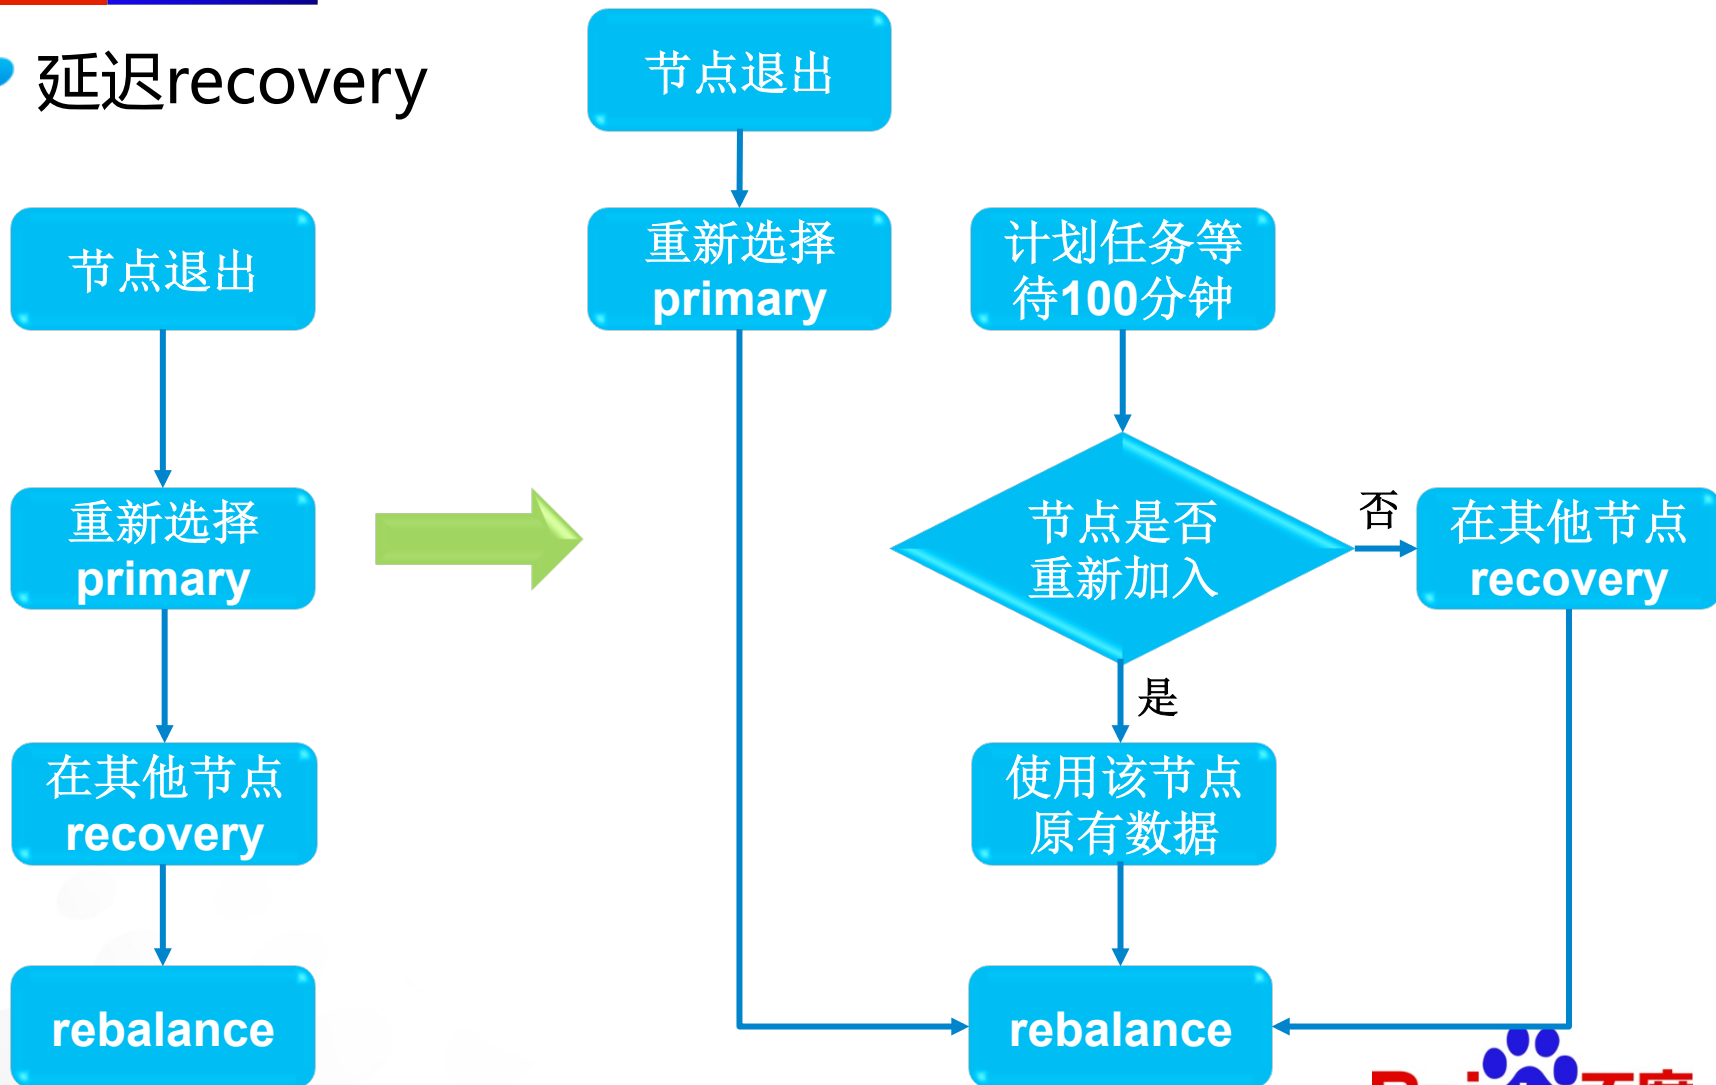
# 遇到的问题及经验分享

## ● 集群异常恢复过慢

- 关闭relocations size及磁盘使用率检查，重启完成后再打开
- 增大每个节点同时允许恢复的分片数
- 增大rebalance最小平衡阈值
- 调整 Index、shard、primary平衡优先级
- 增大集群同时允许rebalance的分片数
- 增大Recovery时回放数据块大小

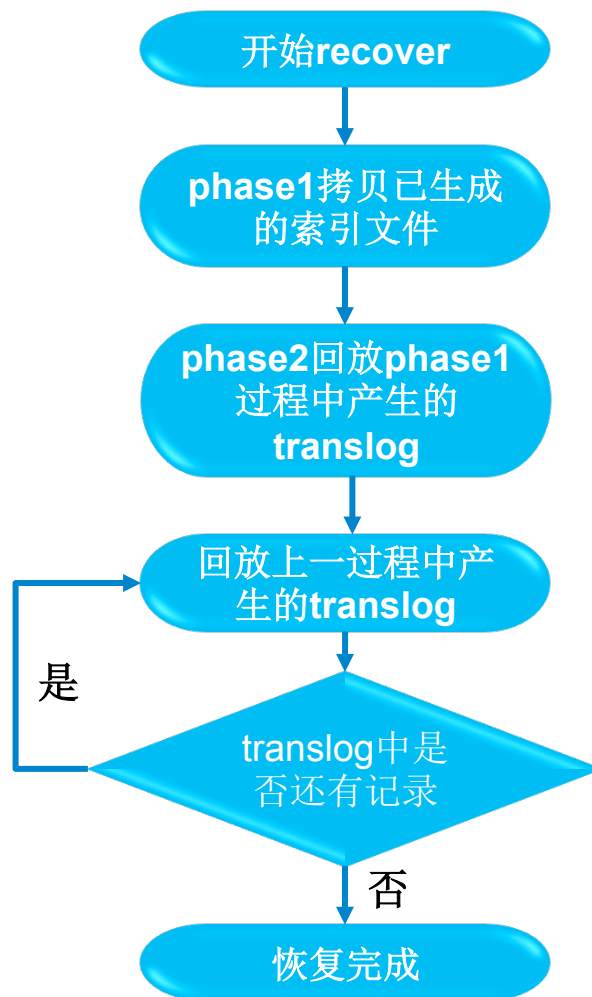

# 对ES的优化与改进

## ● 延迟recovery



# 对ES的优化与改进

- 去除回放translog时的写入锁



# 对ES的优化与改进

- 集成中文分词模块与权限管理模块
- 增加查询时对size和form的限制
- 生成translog时随机选择磁盘



# 后期计划

- 添加SQL解析层

- 支持基本SQL语法，方便用户使用，同时保留原有接口

- 提供ES云化服务




- 一键创建集群
- 权限管理
- 动态伸缩
- 资源隔离
- 监控与报警
- 自动升级与备份

# AWS

## 2015.10.1 AWS在Analytics下面加入了Elasticsearch服务

https://aws.amazon.com/products/?nc2=h\_ql\_ny\_livestream\_blu

Accelerate your Cloud Success with Rich Platform Services

 <b>Analytics</b>	 <b>Enterprise Applications</b>	 <b>Mobile Services</b>	<b>Inter</b>
<ul style="list-style-type: none"><li>Hadoop</li><li>Data Pipelines</li><li><b>Elasticsearch</b></li><li>Streaming Data</li><li>Machine Learning</li><li>Business Intelligence</li><li>Data Warehouse</li></ul>	<ul style="list-style-type: none"><li>Desktop Virtualization</li><li>Email &amp; Calendaring</li><li>Document Sharing &amp; Feedback</li></ul>	<ul style="list-style-type: none"><li>Mobile Development</li><li>API Management</li><li>Identity</li><li>App Testing</li><li>Mobile Analytics</li><li>Notifications</li><li>Development</li></ul>	

Increase Developer Productivity and Operational Efficiency

## PRODUCTS & SERVICES

[Amazon Elasticsearch Service](#) >

[Product Details](#) >

[Pricing](#) >

[Getting Started](#) >

[FAQs](#) >

[What's New](#) >

## RELATED LINKS

[Documentation](#)

[Management Console](#)

[Release Notes](#)

[Discussion Forum](#)

# Amazon Elasticsearch Service

Amazon Elasticsearch Service is a managed service that makes it easy to deploy, operate, and scale Elasticsearch in the AWS Cloud.

Elasticsearch is a popular open-source search and analytics engine for use cases such as log analytics, real-time application monitoring, and click stream analytics. You can set up and configure your Amazon Elasticsearch cluster in minutes from the AWS Management Console. Amazon Elasticsearch Service provisions all the resources for your cluster and launches it. The service automatically detects and replaces failed Elasticsearch nodes, reducing the overhead associated with self-managed infrastructure and Elasticsearch software. Amazon Elasticsearch Service allows you to easily scale your cluster via a single API call or a few clicks in the AWS Management Console. With Amazon Elasticsearch Service, you get direct access to the Elasticsearch open-source API so that code and applications you're already using with your existing Elasticsearch environments will work seamlessly.

To learn more about Elasticsearch and its uses, read the [Getting Started](#) section of [Elasticsearch: The Definitive Guide](#).

# GCE



Google Cloud Platform

Elasticsearch on the ... X Search this site



Sign in

Why Google Products Solutions Launcher Pricing Customers Documentation Support Partners

免费试用

与销售人员联系

产品 > Elasticsearch on the Google Cloud Platform

发送反馈

## Elasticsearch on the Google Cloud Platform

Free Trial Guide

- The Basics
- Start by Programming Language
- Compute
- Storage
- Networking
- Auth and Security
- Big Data
- App Development
- Logging and Monitoring
- Developer Workflow
- Resource Management
- ▼ Click-to-Deploy
  - Aerospike
  - Apache Cassandra
  - Drupal
  - ▼ **Elasticsearch**

## Elasticsearch on Google Compute Engine

**Elasticsearch** is an open-source search and analytics platform built on top of Apache Lucene. Elasticsearch uses a distributed architecture making it a good fit for cloud environments.

### Get Started

#### Use Elasticsearch

Deploy Elasticsearch on Compute Engine in minutes from the Developers Console. [Learn more](#)

Deploy Elasticsearch now

#### Documentation

- [Overview and Features](#)
- [Documentation](#)
- [Getting Started](#)

### Get in Touch

#### Contact sales

Getting started with click-to-deploy Let our sales team help you determine the best way to begin

#### Get support for Elasticsearch

Elasticsearch is supported by a [community](#) of developers and



# 计划

- 2015.10 完成开发
- 2015.11 开始公测
- 2016.02 正式发布
- 如有兴趣，请联系 [palo-rd@baidu.com](mailto:palo-rd@baidu.com)



**Thanks!**