Tools for the Hadoop Stack

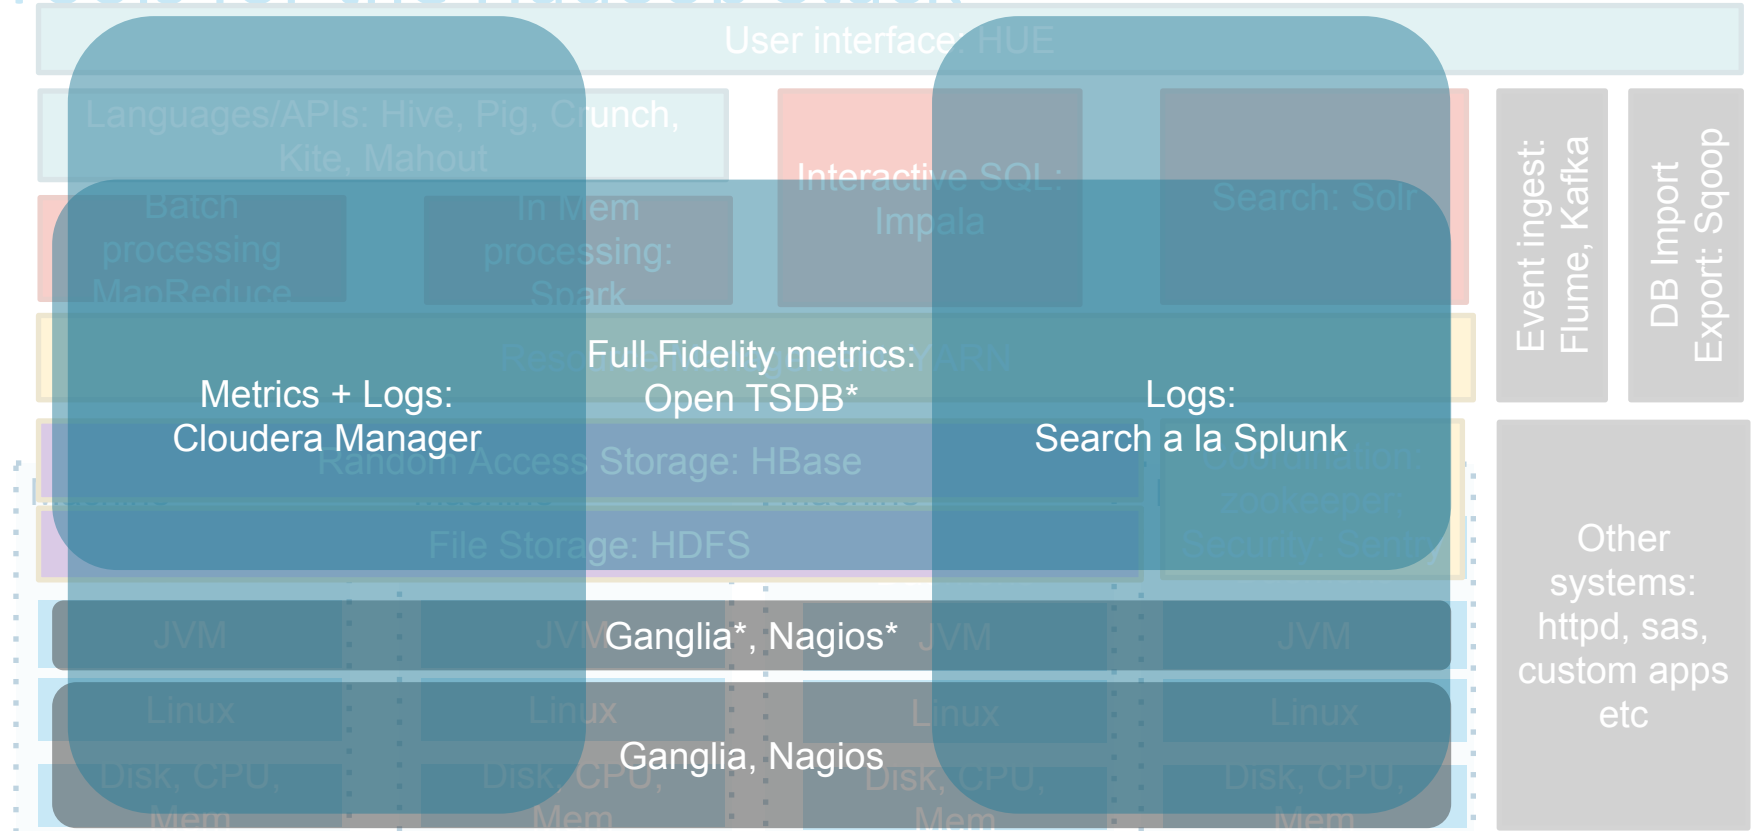

cloudera

Activity

How many file descriptors do the datanodes have open?

What is the current latency of the HDFS canary?

Troubleshooting

Managing a Hadoop Clusters

Troubleshooting Hadoop Systems

Debugging Hadoop Applications

cloudera

The Law of Cluster Inertia

A cluster in a good state stays in a good state,

and

a cluster in a bad state stays in a bad state,

unless

acted upon by an external force.

cloudera

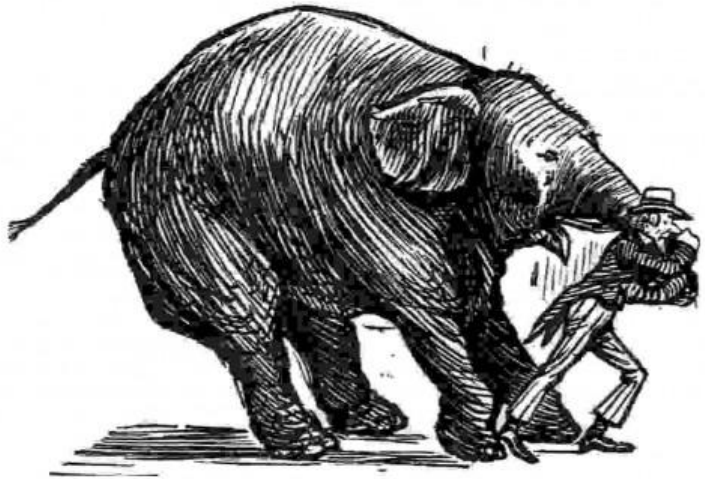

External Forces

- Failures
- Acts of God
- Users
- Admins

cloudera

Failures as an External Force

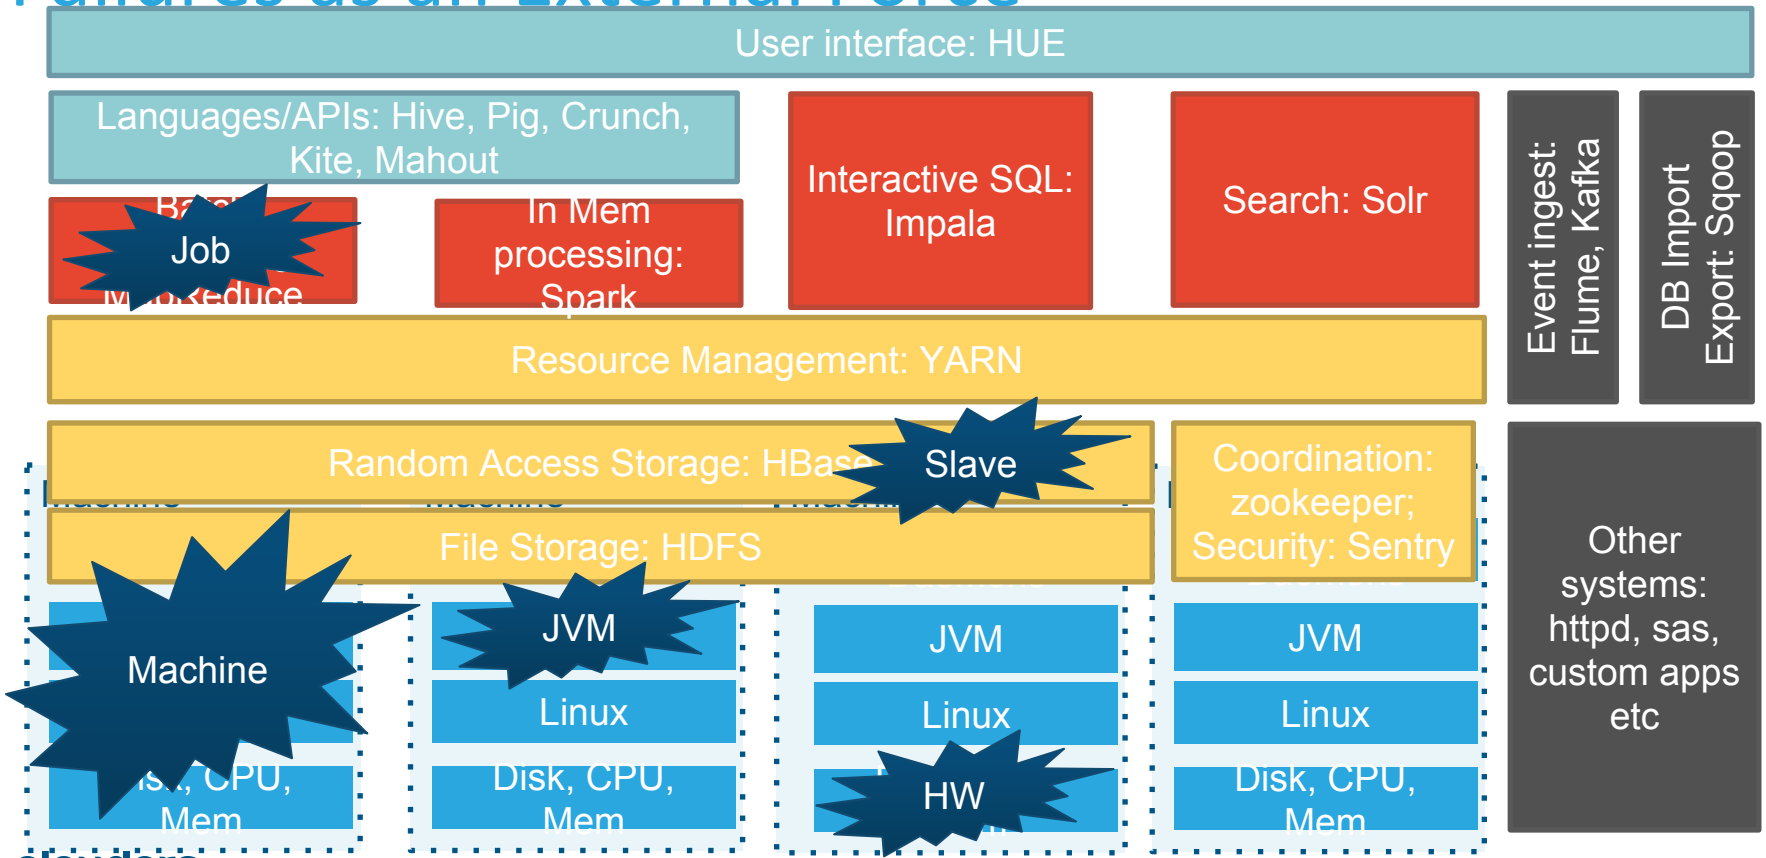

Acts of God as an External Force

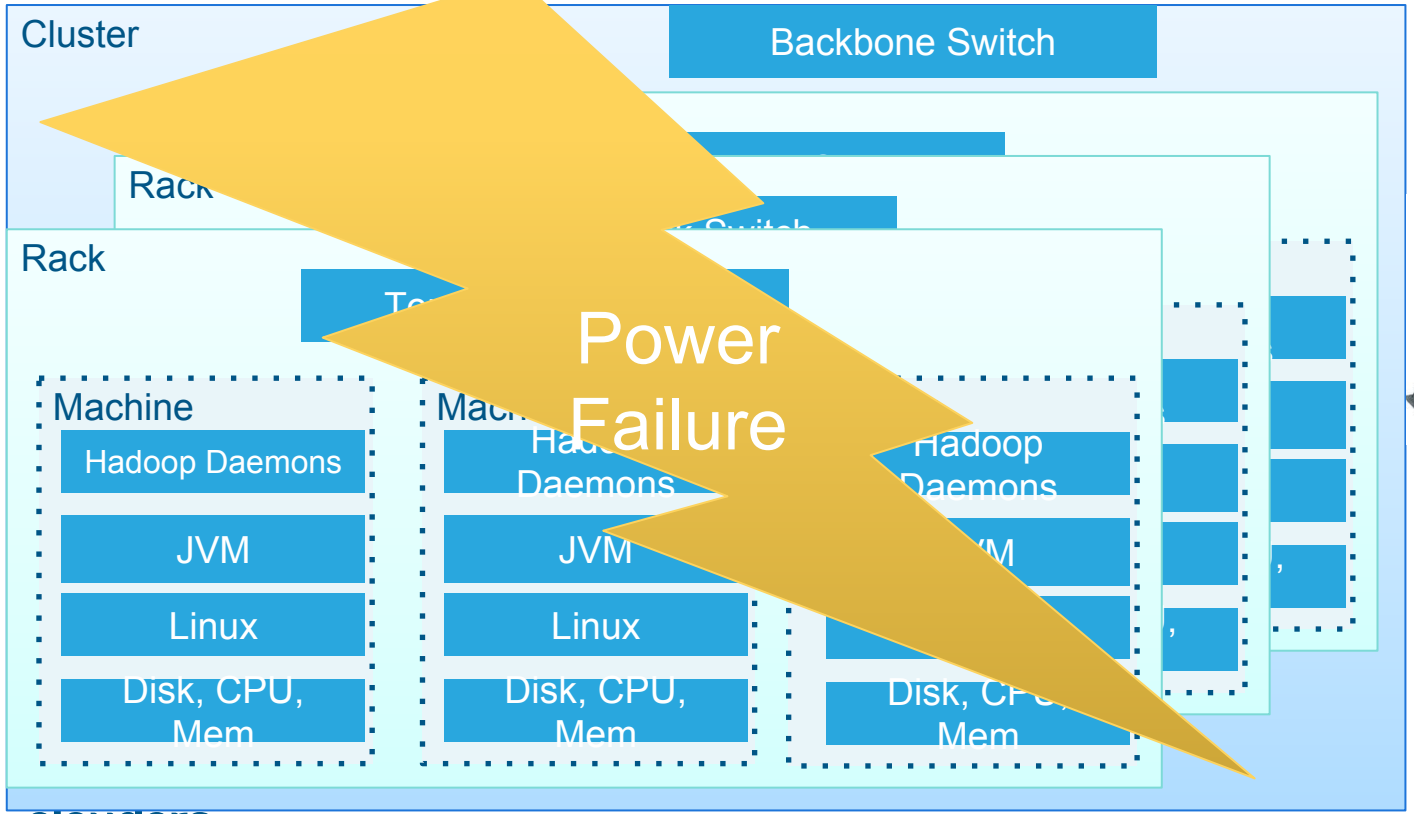

Users as an external force

- Use Security to protect systems from users
  - Prevent and track

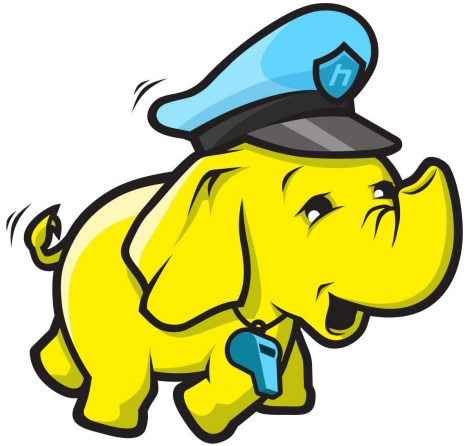

- Authentication – proving who you are
  - LDAP, Kerberos
- Authorization – deciding what you are allowed to do
  - Apache Sentry (incubating), Hadoop security, HBase security
- Audit – who and when was something done?
  - Cloudera Navigator

Admins as an external force

Upgrades

- Linux
- Hadoop
- Java

Misconfiguration

- Memory Mismanagement
  - TT OOME
  - JT OOME
  - Native Threads
- Thread Mismanagement
  - Fetch Failures
  - Replicas
- Disk Mismanagement
  - No File
  - Too Many Files

cloudera

Troubleshooting

Managing Hadoop Clusters

Troubleshooting Hadoop Systems

Debugging Hadoop Applications

Example application pipeline with strict SLAs

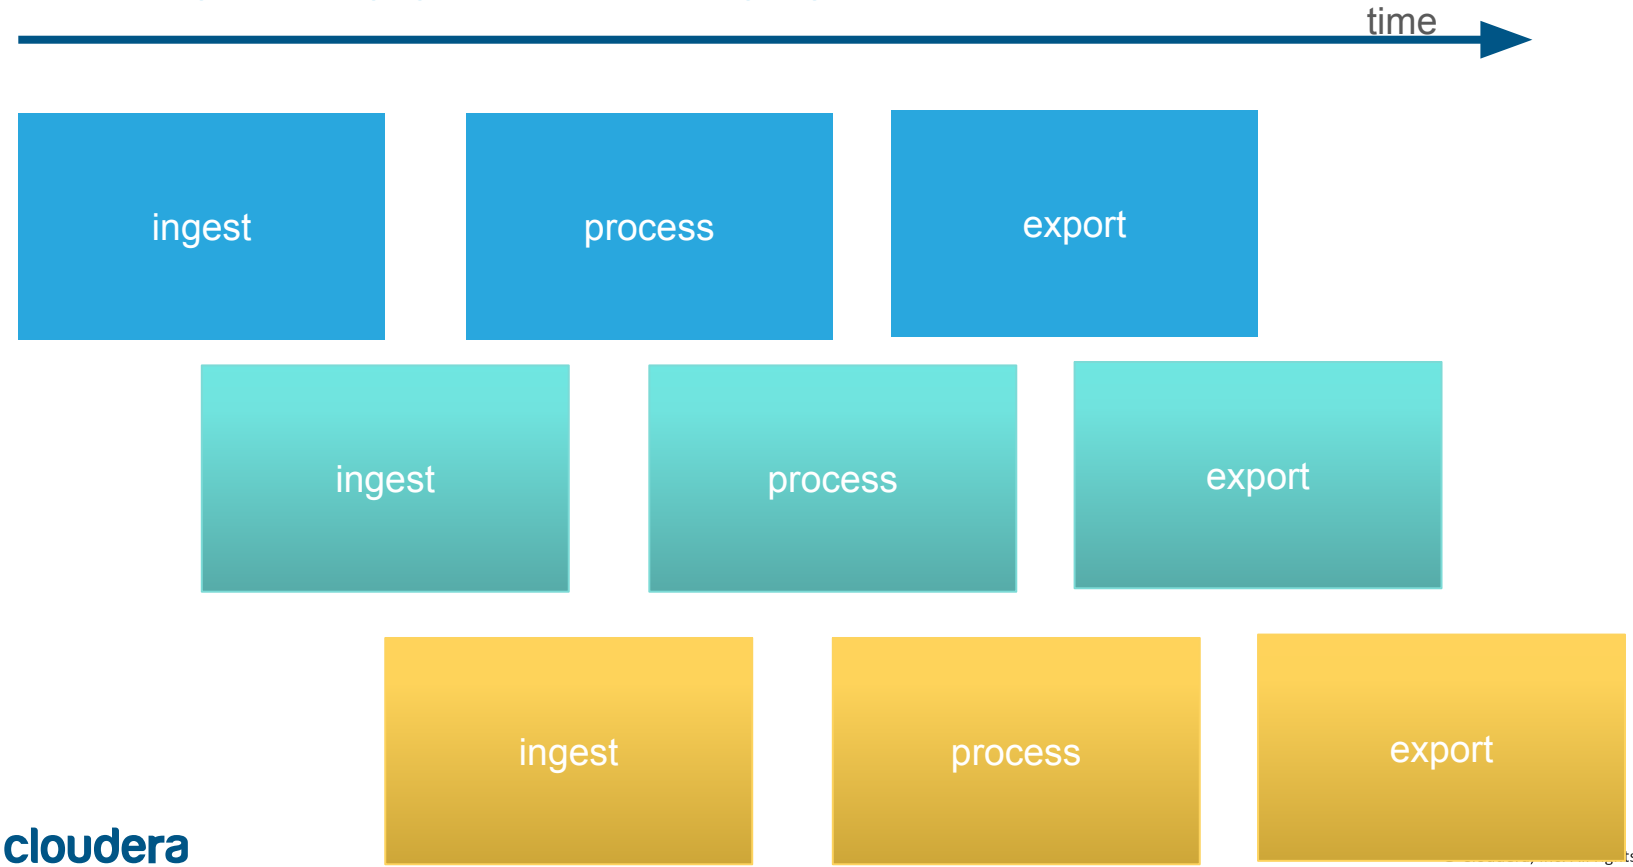

Example Application Pipeline - Ingest

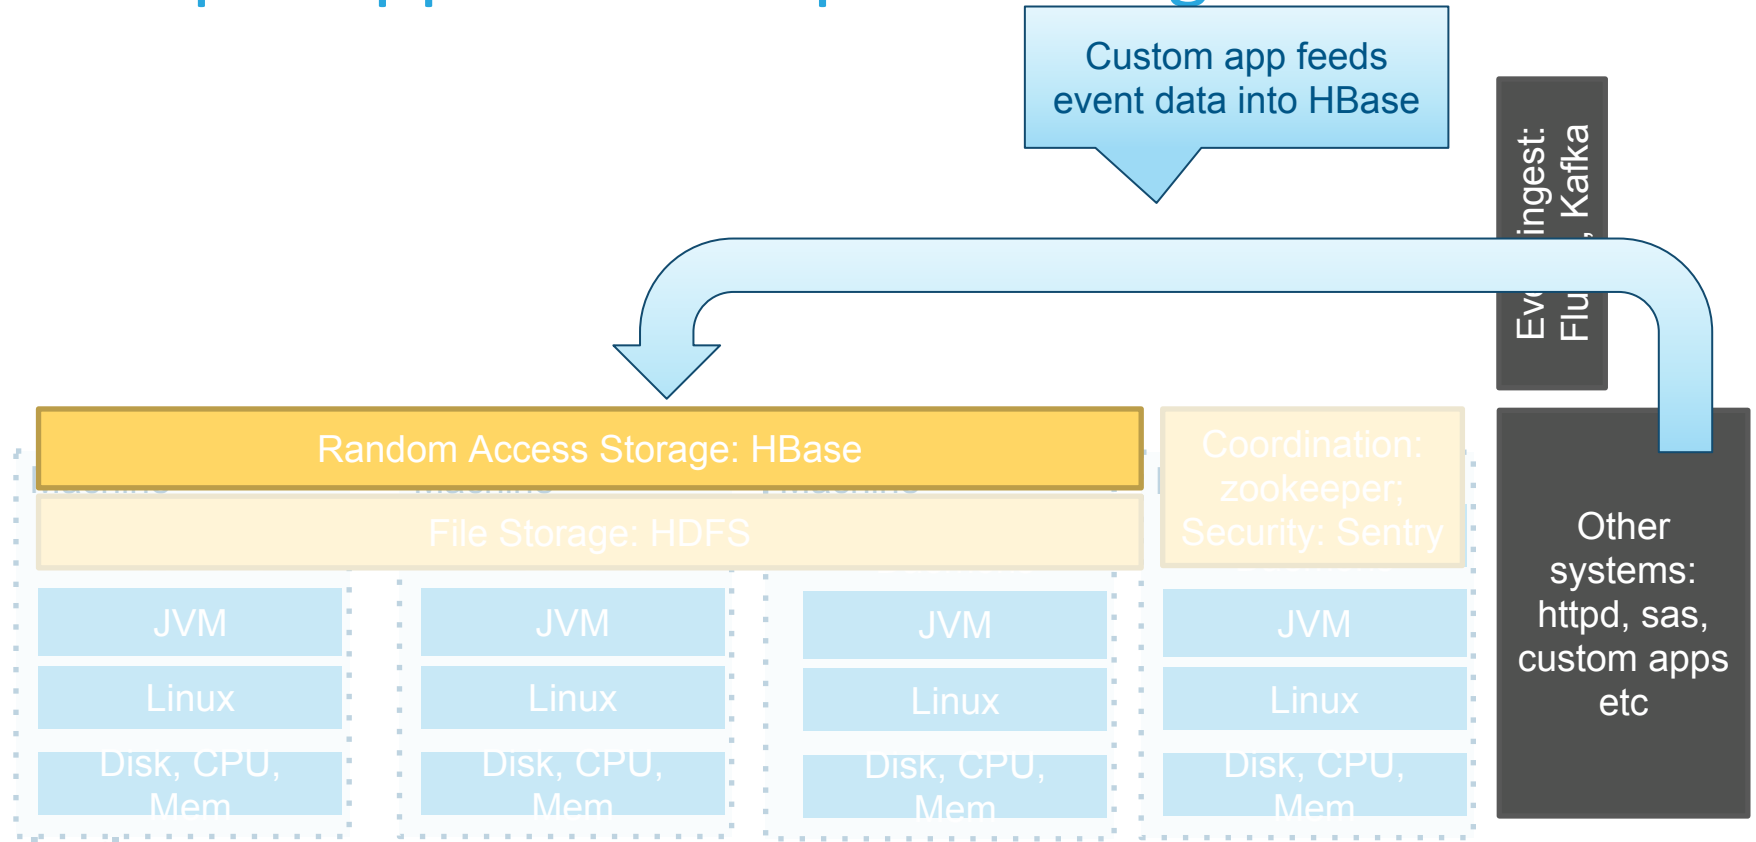

cloudera

Example Application Pipeline - processing

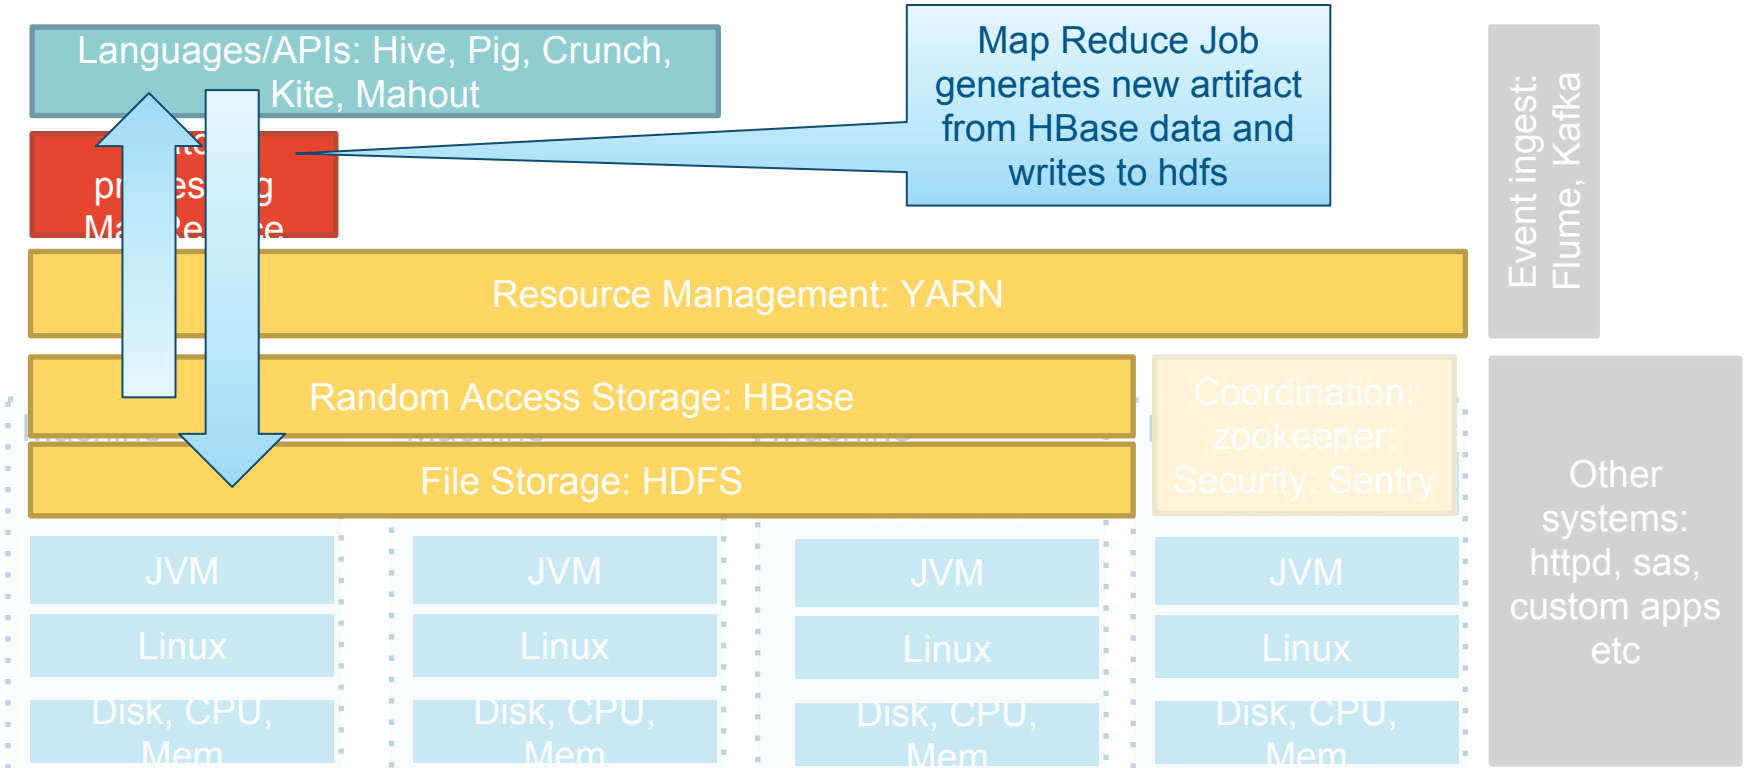

Example Application Pipeline - export

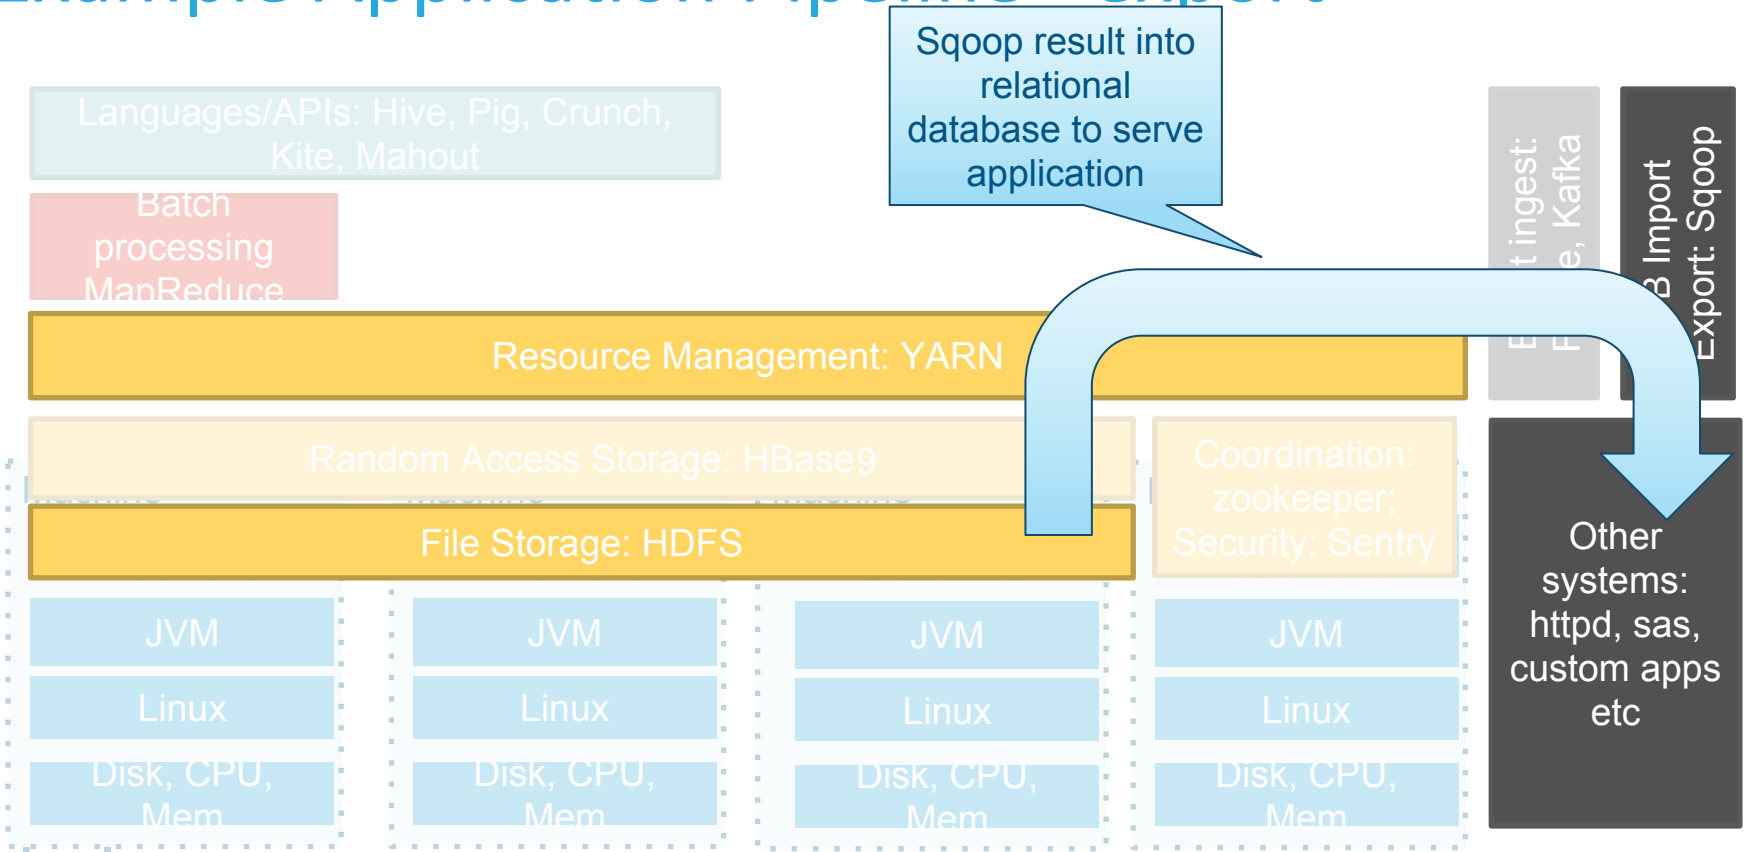

Case study 1: slow jobs after Hadoop upgrade

Symptom:

After an upgrade, activity on the cluster eventually began to slow down and the job queue overflowed.

Finding the right part of the stack

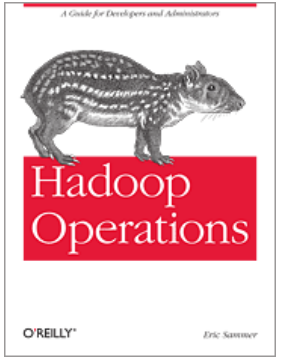

E-SPORE (from Eric Sammer's *Hadoop Operations)*

- Environment
  - What is different about the environment now from the last time everything worked?
- Stack
  - The entire cluster also has shared dependency on data center infrastructure such as the network, DNS, and other services.
- Patterns
  - Are the tasks from the same job? Are they all assigned to the same tasktracker? Do they all use a shared library that was changed recently?
- Output
  - Always check log output for exceptions but don't assume the symptom correlates to the root cause.
- Resources
  - Do local disks have enough? Is the machine swapping? Does the network utilization look normal? Does the CPU utilization look normal?
- Event correlation
  - It's important to know the order in which the events led to the failure.

Example Application Pipeline - processing

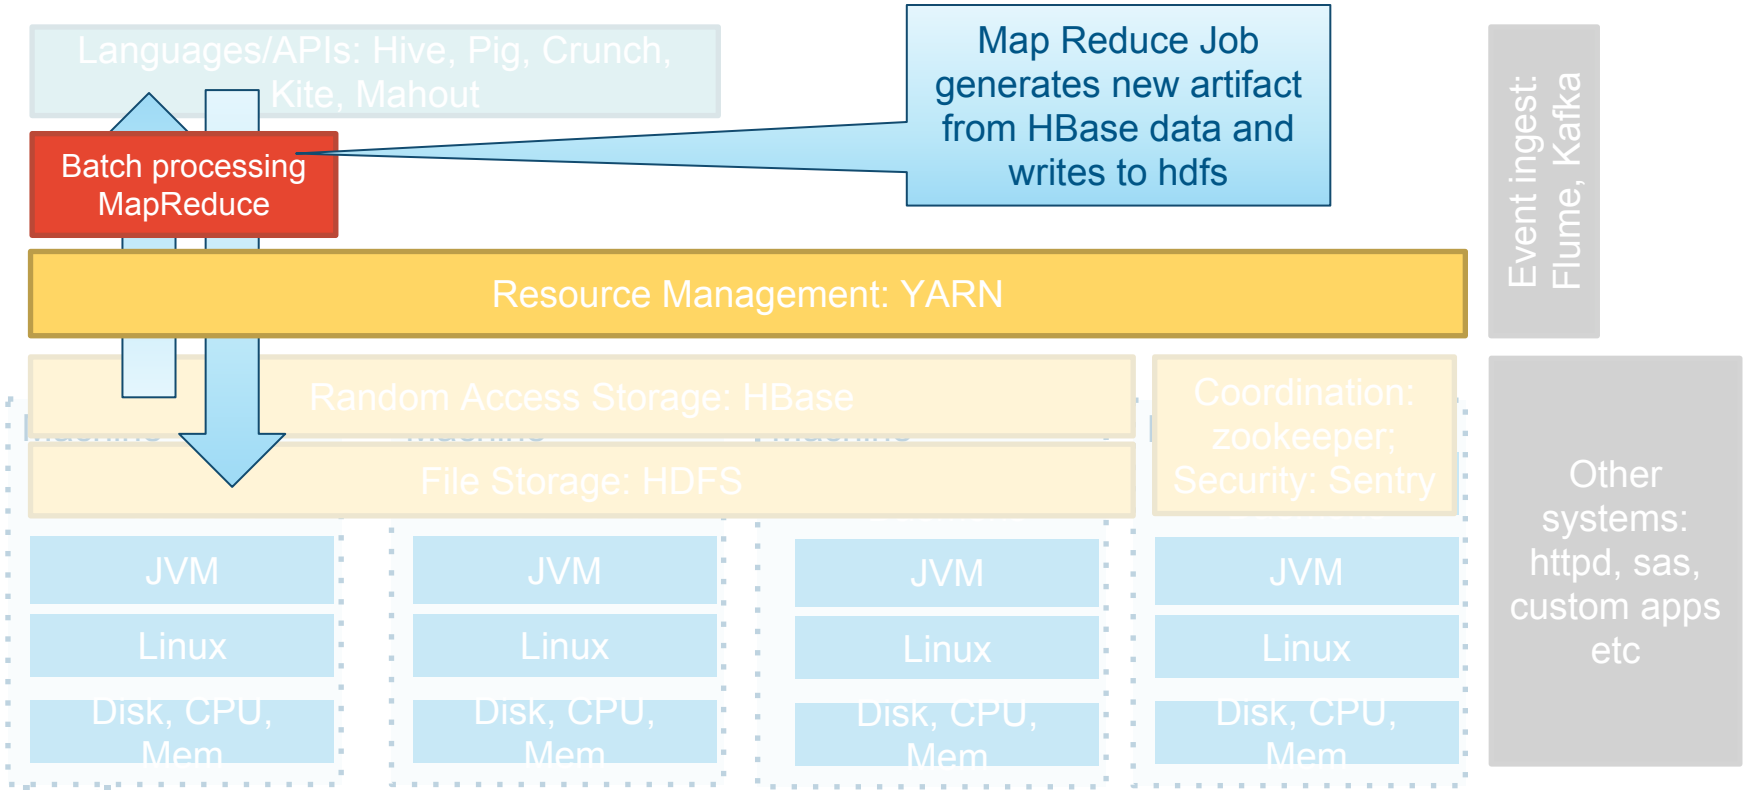

cloudera

Case study 1: slow jobs after Hadoop upgrade

```
INFO
org.apache.hadoop.
mapred.JobInProgre
ss: Too many
fetch-failures for
output of task
```

Evidence:

Isolated to Processing phase (MR).

In TT Logs, found an innocuous but anomalous log entry about “fetch failures.”

Many users had run in to this MR problem using different versions of MR.

Workaround provided: remove the problem node from the cluster.

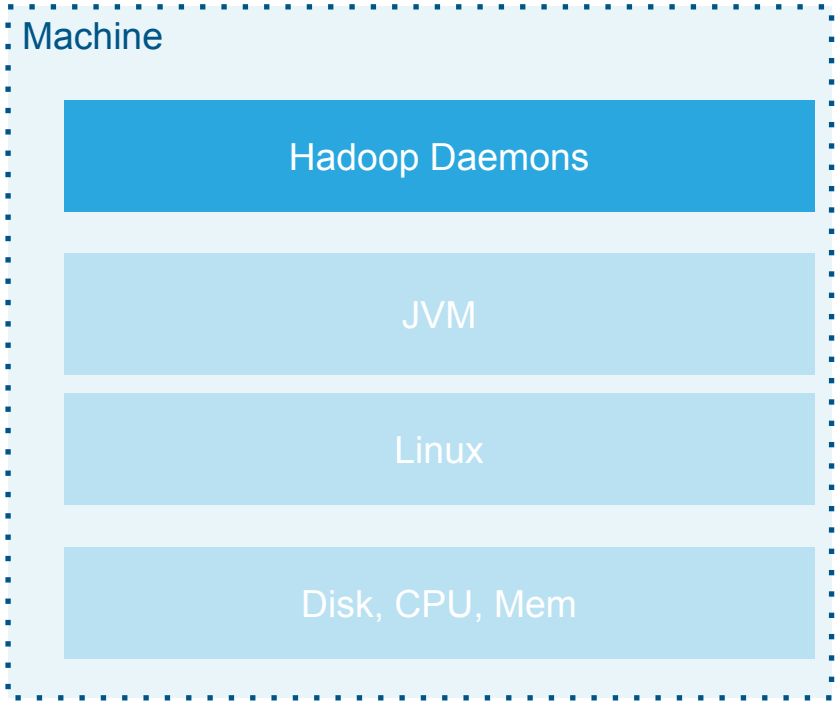

Case study 1: slow jobs after Hadoop upgrade

Root cause:

All MR versions had a common dependency on a particular version of Jetty (Jetty 6.1.26) .

Dev was able to reproduce and fix the bug in Jetty.

Case study 2: slow jobs after Linux upgrade

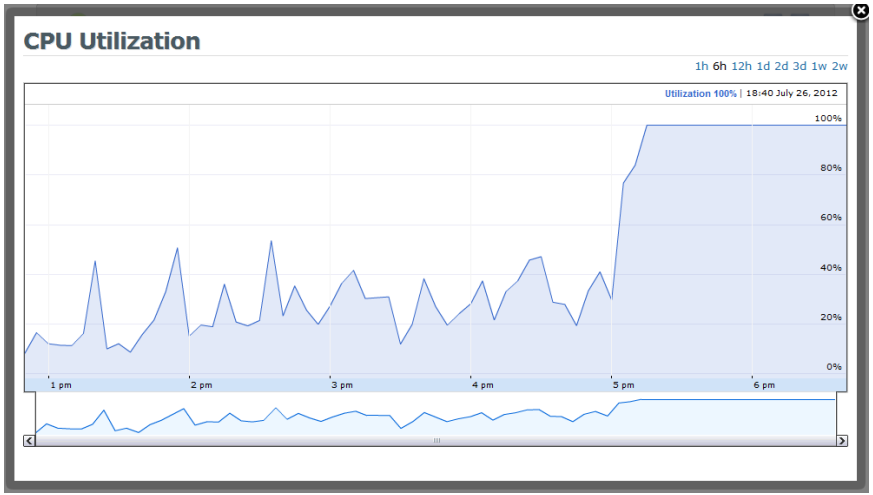

Symptom:

After an upgrade, system CPU usage peaked at 30% or more of the total CPU usage.

cloudera

Case study 2: slow jobs after Linux upgrade

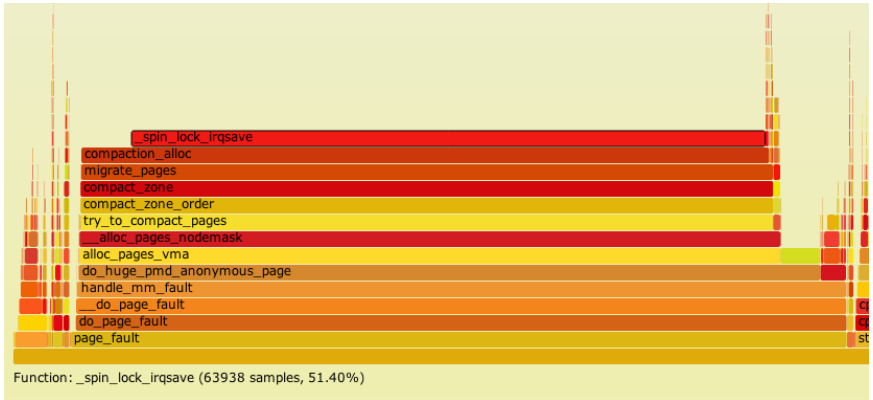

Evidence:

Used tracing tools to isolate a majority of time was inexplicably spent in virtual memory calls.

http://structureddata.org/2012/06/18/linux-6-transparent-huge-pages-and-hadoop-workloads/

cloudera

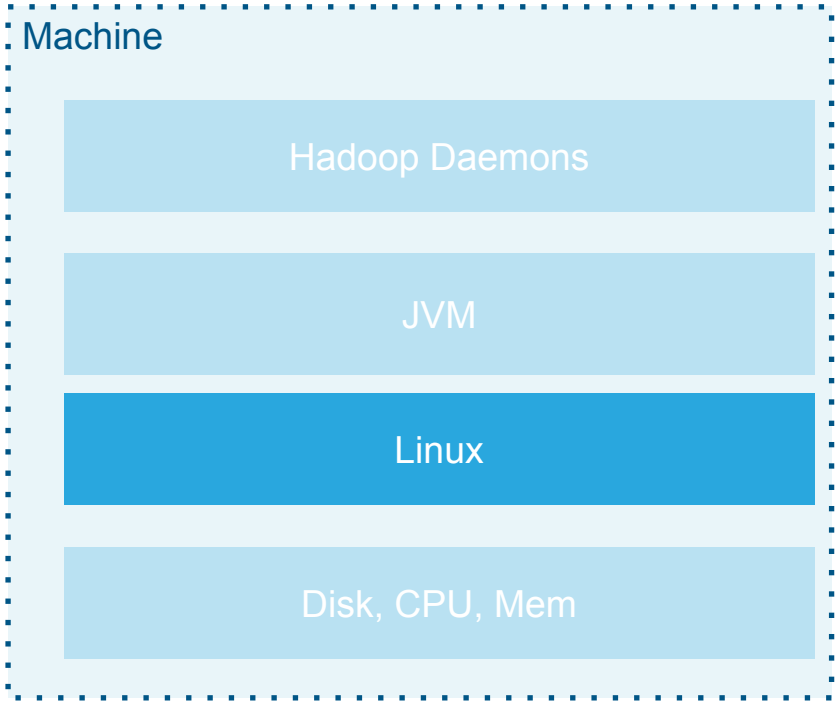

Case study 2: slow jobs after Linux upgrade

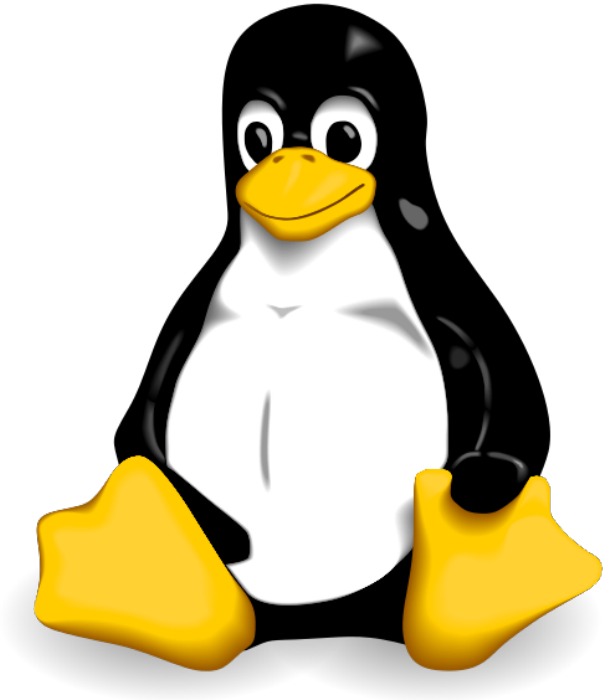

Root cause:

Running RHEL or CentOS versions 6.2, 6.3, and 6.4 or SLES 11 SP2 has a feature called "Transparent Huge Page (THP)" compaction which interacts poorly with Hadoop workloads.

cloudera

Case study 3: slow jobs at a precise moment

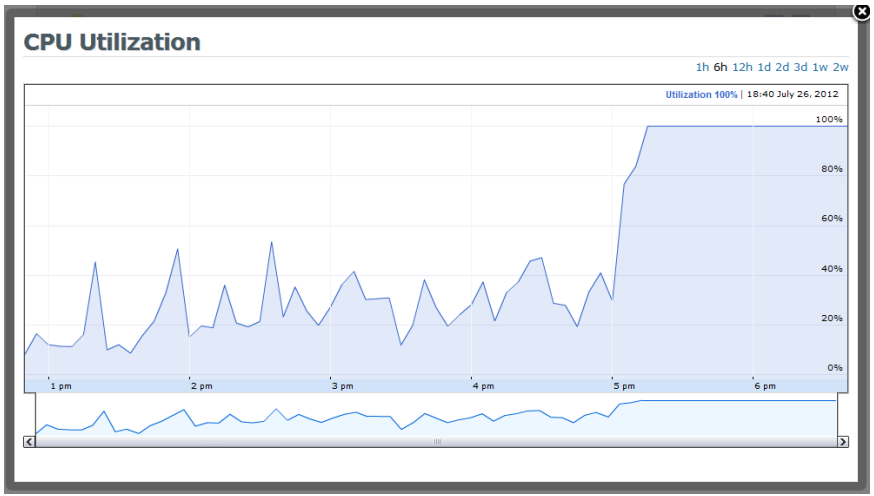

Symptom:

High CPU usage and responsive but sluggish cluster - even non-Hadoop apps e.g. MySQL.

30 customers all hit this at the exact same time: 6/30/12 at 5pm PDT.

Case study 3: slow jobs at a precise moment

```
Clock: inserting leap second 23:59:60 UTC
```

Evidence:

Checked the kernel message buffer (run dmesg) and look for output confirming the leap second injection.

Other systems had same problem.

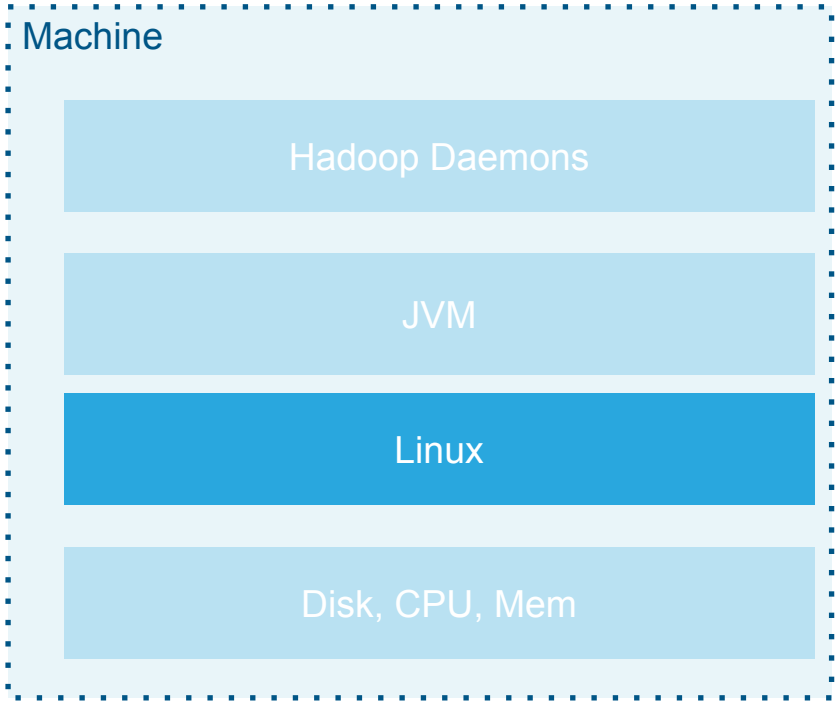

Case study 3: slow jobs at a precise moment

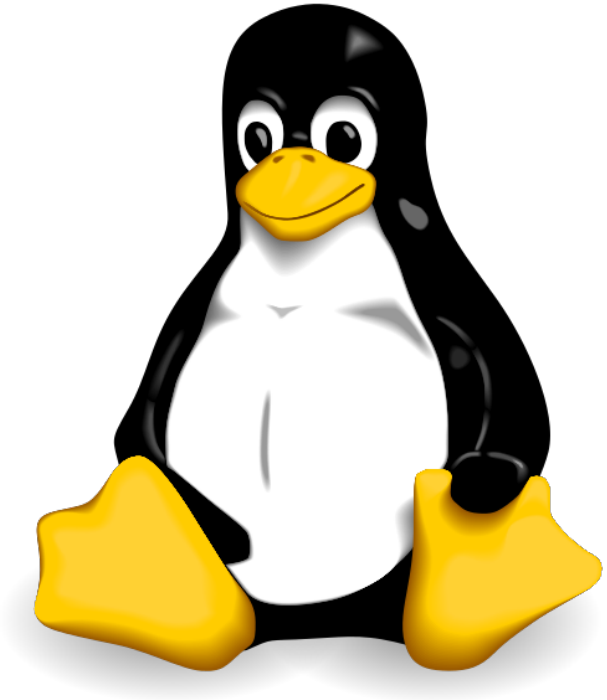

Root cause:

Linux OS kernel mishandled a leap second added.

Similar symptoms, different problem

cloudera

Case study 1: slow jobs after Hadoop upgrade

```
INFO
org.apache.hadoop.
mapred.JobInProgre
ss: Too many
fetch-failures for
output of task
```

Symptom

After an upgrade, activity on the cluster eventually began to slow down and the job queue overflowed.

Evidence

In TT Logs, found an innocuous but anomalous log entry about “fetch failures.”

Many users had run in to this MR problem using different versions of MR.

Workaround provided: remove the problem node from the cluster.

Root Cause

All MR versions had a common dependency on a particular version of Jetty (Jetty 6.1.26) .

Dev was able to reproduce and fix the bug in Jetty.

cloudera

Case study 2: slow jobs after Linux upgrade

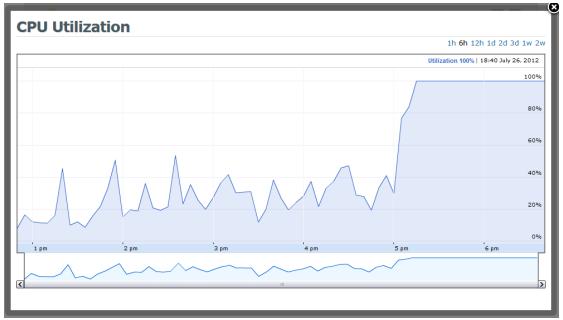

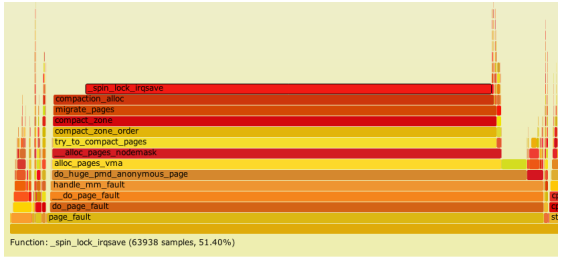

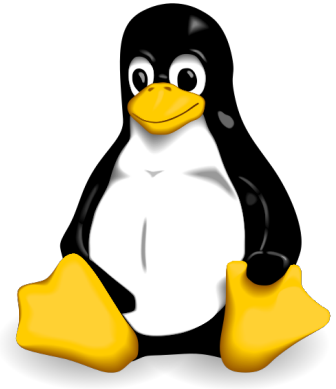

Symptom

After an upgrade, system CPU usage peaked at 30% or more of the total CPU usage.

Evidence

Perf tool which proved that majority of time was inexplicably spent in virtual memory calls.

Root Cause

Running RHEL or CentOS versions 6.2, 6.3, and 6.4 or SLES 11 SP2 has a feature called "Transparent Huge Page (THP)" compaction which interacts poorly with Hadoop workloads.

cloudera

Case study 3: slow jobs at a precise moment

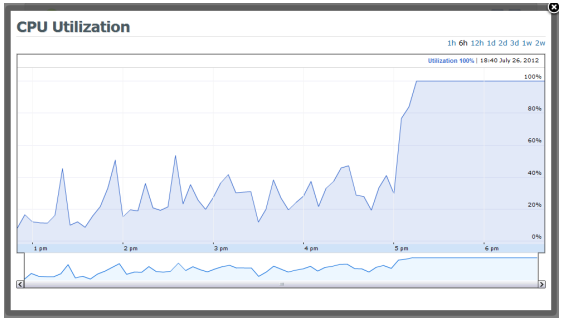


```
Clock: inserting leap second 23:59:60 UTC
```

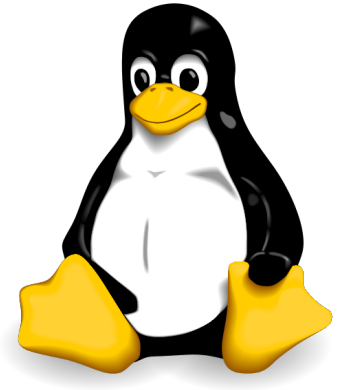

Symptom

High CPU usage and responsive but sluggish cluster.

30 customers all hit this at the exact same time.

Evidence

Checked the kernel message buffer (run dmesg) and look for output confirming the leap second injection.

Root Cause

Linux OS kernel mishandled a leap second added on 6/30/12 at 5pm PDT.

cloudera

Lessons learned

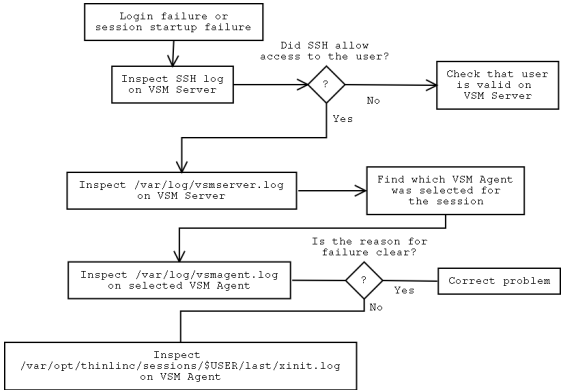

Methodology

More crucial than the specific troubleshooting methodology used is to use one.

Tools

More crucial than the specific tool used is the type of data analyzed and how it's analyzed.

Learn from failure

Capture for posterity in a knowledge base article, blog post, or conference presentation.

cloudera

cloudera

Questions?

cloudera

Apache Hadoop Operations for Production Systems: Enterprise Considerations

Miklos Christine

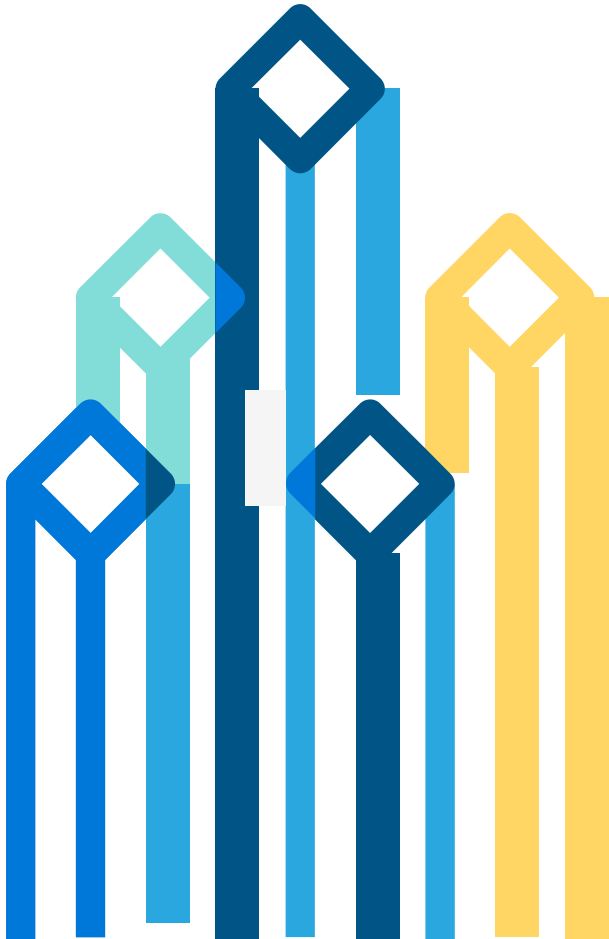

Scale Considerations

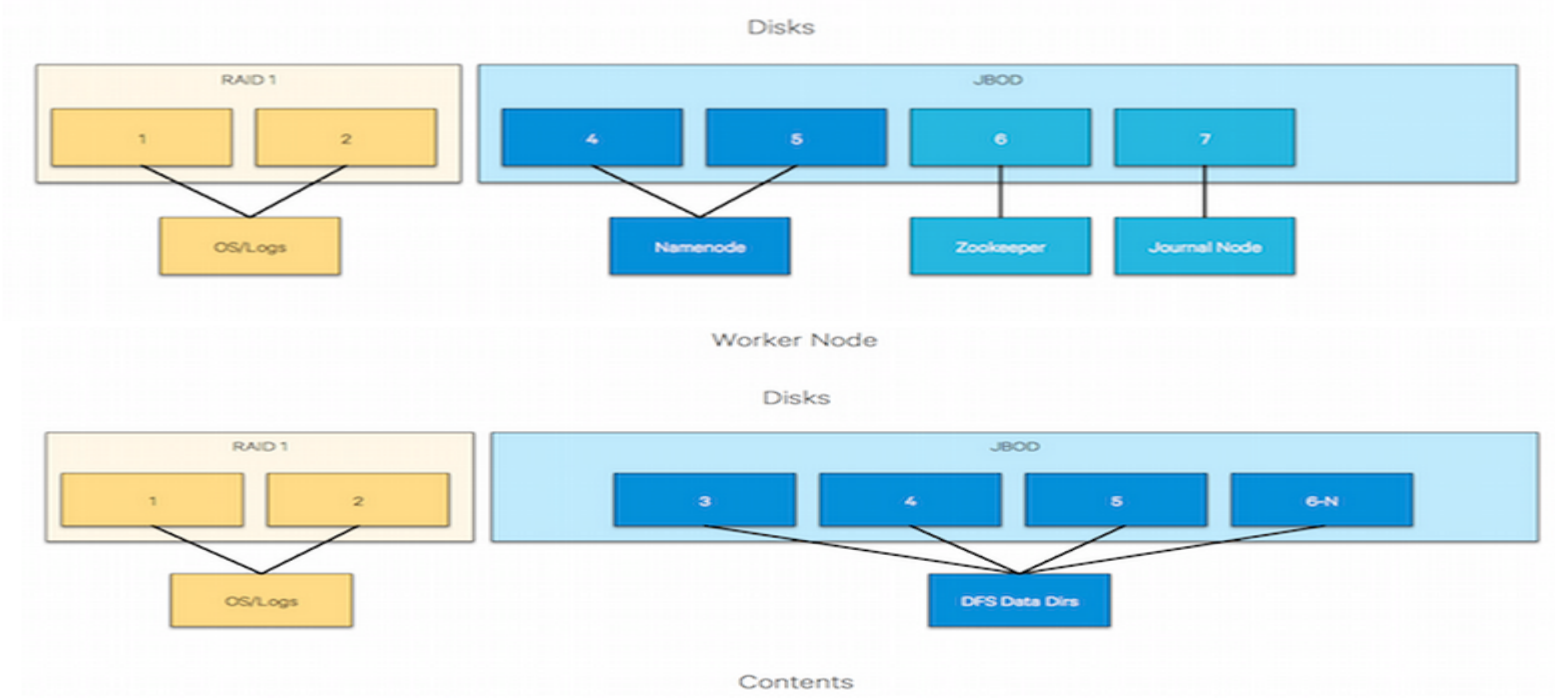

Ref: http://blog.cloudera.com/blog/2015/01/how-to-deploy-apache-hadoop-clusters-like-a-boss/

Scale Considerations

- HDFS
  - Namenode Heap Settings
  - Namenode RPC Configurations

Property Name	Default	Recommended
dfs.namenode.servicerpc-address	N/A	8022
dfs.namenode.handler.count	10	ln(# of DNs) * 20
dfs.namenode.service.handler.count	10	ln(# of DNs) * 20

Scale Considerations

- YARN
  - ResourceManager High Availability
    - yarn.resourcemanager.zk-address
  - Application recovery
    - yarn.resourcemanager.work-preserving-recovery.enabled
    - yarn.nodemanager.recovery.dir
  - User cache disk space
    - yarn.nodemanager.local-dirs

Metrics: HDFS

- HDFS is the core of the platform. What's important?
  - Is the Standby NN checkpointing?
  - Are the NNs garbage collecting?
  - Percentage of heap used at steady state?


```
"name" : "Hadoop:service=NameNode,name=FSNamesystem",
"modelerType" : "FSNamesystem",
"tag.Context" : "dfs",
"tag.HAState" : "active",
"tag.Hostname" : "mwc-2.ent.cloudera.com",
"MissingBlocks" : 0,
"MissingReplOneBlocks" : 0,
"ExpiredHeartbeats" : 0,
"TransactionsSinceLastCheckpoint" : 115,
"TransactionsSinceLastLogRoll" : 115,
"LastWrittenTransactionId" : 80716,
"LastCheckpointTime" : 1423453166789,
"CapacityTotal" : 262529737116,
"CapacityTotalGB" : 244.0,
"CapacityUsed" : 675131392,
"CapacityUsedGB" : 1.0,
"CapacityRemaining" : 185955512320,
"CapacityRemainingGB" : 173.0,
"CapacityUsedNonDFS" : 75899093404,
"TotalLoad" : 6,
"SnapshottableDirectories" : 0,
"Snapshots" : 0,
"BlocksTotal" : 297,
"FilesTotal" : 2354,
"PendingReplicationBlocks" : 0,
"UnderReplicatedBlocks" : 1,
"CorruptBlocks" : 1,
"ScheduledReplicationBlocks" : 0,
"PendingDeletionBlocks" : 0,
"ExcessBlocks" : 0,
"PostponedMisreplicatedBlocks" : 0,
"PendingDataNodeMessageCount" : 0,
"MillisSinceLastLoadedEdits" : 0,
"BlockCapacity" : 1048576,
"StaleDataNodes" : 0,
"TotalFiles" : 2354
```


Logs are your friend

- Logs are verbose but necessary
  - Namenode Logs:
    - 10 * 200MB log files = 2GB
    - 2GB of logs span 3 hours
    - 3 days of logs = ~48GB
- Retain enough logs for debugging. Plan for the worst case
  - Adjust log retention as the cluster grows

Logs are your friend

- Just reduce the log level to save space?
  - NO!
  - INFO logging is important!
- Yarn containers write logs locally, then migrate to HDFS.
  - Ensure application log space is sufficient

Logs are your friend

- GC Logging
  - `-verbose:gc -XX:+PrintGCDetails`
  - `-XX:+PrintGCTimeStamps -XX:+PrintGCDateStamps`
  - `-Xloggc:/var/log/hdfs/nn-hdfs.log`
  - `-XX:+UseGCLogFileRotation -XX:NumberOfGCLogFiles=5 -XX:GCLogFileSize=20M`
- Great resource:
  - https://stackoverflow.com/questions/895444/java-garbage-collection-log-messages

Debugging Techniques : Hung Process

- jstack
  - Use the same JDK
  - Must be run as user of the process
- kill -3 <PID>
  - Dumps jstack to stdout

```
2015-02-14 15:47:30
Full thread dump Java HotSpot(TM) 64-Bit Server VM (24.65-b04 mixed mode):

"process reaper" daemon prio=10 tid=0x00007f8844031000 nid=0x3191 waiting on condition [0x00007f88310e4000]
   java.lang.Thread.State: TIMED_WAITING (parking)
        at sun.misc.Unsafe.park(Native Method)
        - parking to wait for  <0x00000000ef7728c0> (a java.util.concurrent.SynchronousQueue$TransferStack)
        at java.util.concurrent.locks.LockSupport.parkNanos(LockSupport.java:226)
        at java.util.concurrent.SynchronousQueue$TransferStack.awaitFulfill(SynchronousQueue.java:460)
        at java.util.concurrent.SynchronousQueue$TransferStack.transfer(SynchronousQueue.java:359)
        at java.util.concurrent.SynchronousQueue.poll(SynchronousQueue.java:942)
        at java.util.concurrent.ThreadPoolExecutor.getTask(ThreadPoolExecutor.java:1068)
        at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1130)
        at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:615)
        at java.lang.Thread.run(Thread.java:745)

"48403724@qtp-644610485-9" daemon prio=10 tid=0x00000000007df000 nid=0x3110 in Object.wait() [0x00007f8841754000]
   java.lang.Thread.State: TIMED_WAITING (on object monitor)
        at java.lang.Object.wait(Native Method)
        - waiting on <0x00000000e9d43598> (a org.mortbay.thread.QueuedThreadPool$PoolThread)
        at org.mortbay.thread.QueuedThreadPool$PoolThread.run(QueuedThreadPool.java:626)
        - locked <0x00000000e9d43598> (a org.mortbay.thread.QueuedThreadPool$PoolThread)

"Trash Emptier" daemon prio=10 tid=0x00007f884ca10000 nid=0x2839 waiting on condition [0x00007f883131f000]
   java.lang.Thread.State: TIMED_WAITING (sleeping)
        at java.lang.Thread.sleep(Native Method)
        at org.apache.hadoop.fs.TrashPolicyDefault$Emptier.run(TrashPolicyDefault.java:265)
        at java.lang.Thread.run(Thread.java:745)
```

Debugging Techniques : Hung Process

```
Thread 10175: (state = BLOCKED)
 - java.lang.Object.wait(long) @bci=0 (Interpreted frame)
 - java.lang.Object.wait() @bci=2, line=503 (Interpreted frame)
 - org.apache.hadoop.ipc.Server.join() @bci=8, line=2421 (Interpreted frame)
 - org.apache.hadoop.hdfs.server.namenode.NameNodeRpcServer.join() @bci=4, line=412 (Interpreted frame)
 - org.apache.hadoop.hdfs.server.namenode.NameNode.join() @bci=4, line=790 (Interpreted frame)
 - org.apache.hadoop.hdfs.server.namenode.NameNode.main(java.lang.String[]) @bci=39, line=1495 (Interpreted frame)
```

```
"main" prio=10 tid=0x00007f884c017800 nid=0x27bf in Object.wait() [0x00007f8854507000]
   java.lang.Thread.State: WAITING (on object monitor)
        at java.lang.Object.wait(Native Method)
        - waiting on <0x00000000f01aa070> (a org.apache.hadoop.ipc.ProtobufRpcEngine$Server)
        at java.lang.Object.wait(Object.java:503)
        at org.apache.hadoop.ipc.Server.join(Server.java:2421)
        - locked <0x00000000f01aa070> (a org.apache.hadoop.ipc.ProtobufRpcEngine$Server)
        at org.apache.hadoop.hdfs.server.namenode.NameNodeRpcServer.join(NameNodeRpcServer.java:412)
        at org.apache.hadoop.hdfs.server.namenode.NameNode.join(NameNode.java:790)
        at org.apache.hadoop.hdfs.server.namenode.NameNode.main(NameNode.java:1495)
```

cloudera

Debugging Techniques : LogLevel

- Set the log level without process restarts

```
http://namenode.cloudera.com:50070/logLevel
```

- Scriptable

```
http://namenode.cloudera.com:50070/logLevel?log=org&level=DEBUG
```

Debugging Techniques : Heap Analysis

- jstat -gcutil <PID> 1s 120
  - Checks for current GC activity
- jmap -histo:live <PID>
  - Get a histogram of the current objects within the heap

Controlled Usage

- How to prevent bad behavior from bringing down the cluster?
  - HDFS Quotas
  - Yarn FairScheduler Pools
  - Hive / Impala Access Control with Sentry

Failure Testing

- If the NN fails, how long does it take to recover given the average # of edits?
- If RM HA failover were to occur, would jobs continue?

Security Considerations

- Securing communication channels within the cluster
  - Kerberos
    - Allows secure communication between hosts on an untrusted network.
    - Secures traffic between hosts in the cluster
    - Provides authentication for users to services
  - TLS
    - Used to secure http interfaces
    - Kerberos can be used to authenticate to these interfaces with SPNEGO

Kerberos, Authentication and Authorization

- While often conflated, these are distinct concepts
- They are usually configured together, and we would recommend this, but it's not an absolute requirement
- **Authentication**: Having a user provide and prove their identity
- **Authorization**: Controlling what a user can access or do
- Similarly, there are authentication and authorization mechanisms you can use which don't depend on Kerberos
  - But which are not considered very defensible

Security and Authentication (cont)

- Setting up Kerberos is an exercise that's beyond the scope of this tutorial
  - Main implementations: MIT Kerberos, Active Directory
  - Typically LDAP (or AD) is used for user management
- Cloudera Manager can help you configure Kerberos for your services

Authentication

- Without Kerberos, users are typically identified as whatever Linux system user their client application runs as.
- With Kerberos, the user will obtain a kerberos ticket (typically at login time) that will be used to identify them to the cluster services

Authorization

- Even if you're using an authentication mechanism to limit who can connect to the various services, you probably want to control what they can do. Without authorization, anyone can do anything
- Each service provides different authorization mechanisms. eg:
  - YARN queues can be restricted to certain users
  - HBase tables can be restricted to certain users
- The nature of cluster users will affect authorization requirements
  - Are there different groups with different SLAs?

Office Hours tomorrow (Thursday)

10:40-11:20 am at O’Reilly booth

cloudera

Takeaway

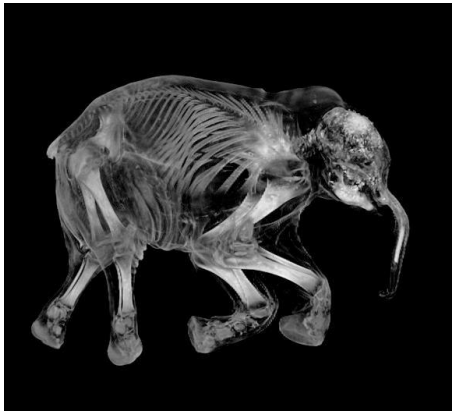

Anatomy of a Hadoop System

A cluster in a good state stays in a good state, and a cluster in a bad state stays in a bad state, unless acted upon by an external force..

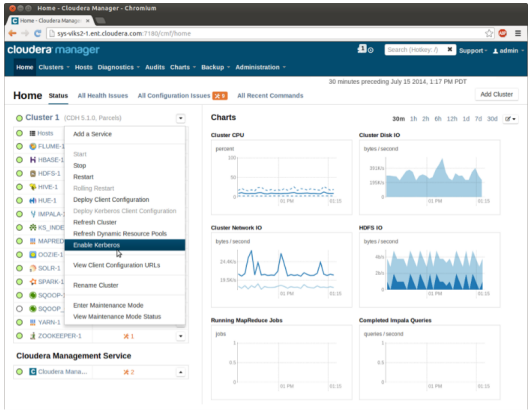

Managing Hadoop Clusters

Cloudera has seen a lot of diverse clusters and used that experience to build tools to help diagnose and understand how Hadoop operates.

Troubleshooting Hadoop Applications

Similar symptoms can lead to different root causes. Use tools to assist with event correlation and pattern determination.

Join the Discussion

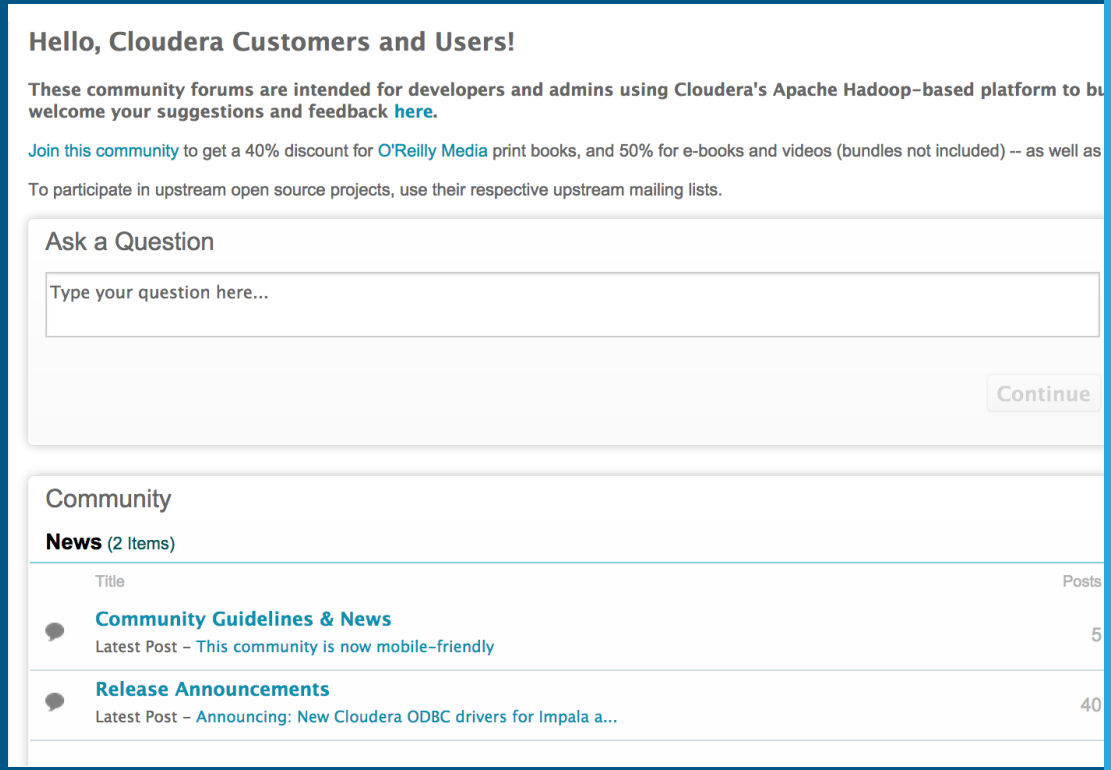

Get community help or provide feedback

cloudera.com/community

cloudera


cloudera LIVE
The fastest way
to get started
with Hadoop

Visit us at Booth #809

HIGHLIGHTS:

Apache Kafka is now fully supported with Cloudera

Learn why Cloudera is the leader for security in Hadoop

cloudera

BOOK SIGNINGS

THEATER SESSIONS

TECHNICAL DEMOS

GIVEAWAYS

cloudera

Questions?