




Multi-tenant Apache Hadoop Clusters

Dániel Schöberle | Cloudera


Designated Support Engineer for Bank of America



What is multi-tenancy?



Single tenant



Free-for-all



Multi-tenancy

Why do we need it?

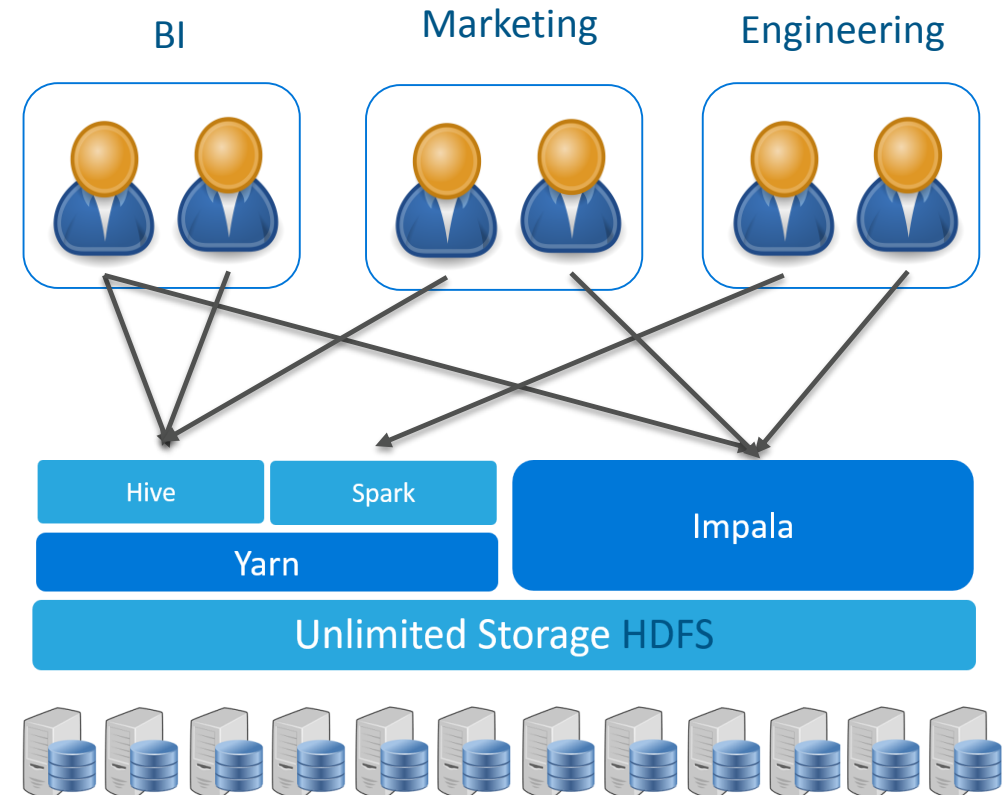
- **Optimize** resource usage
- **Share** infrastructure
- Allow different groups **access** to storage/data
- Support wide audience (developers, analysts, data scientists from different organizational units)
- Allow the little guy access to big resources

What should multi-tenancy solve?

- Resource Management / Sharing
- Access control / Security
- Reporting / Operations / Management considerations


What is a multi-tenant Hadoop cluster?

- Single General Purpose Hadoop Cluster
- Multiple distinct user groups with code & data that need to be separated
- Sharing storage (HDFS) & processing resources (cores & RAM)
- Mixed work loads storage only, batch & interactive processing
- Typically run by an in-house data center team on-premise or in the cloud



cloudera®

Resource Isolation & Management

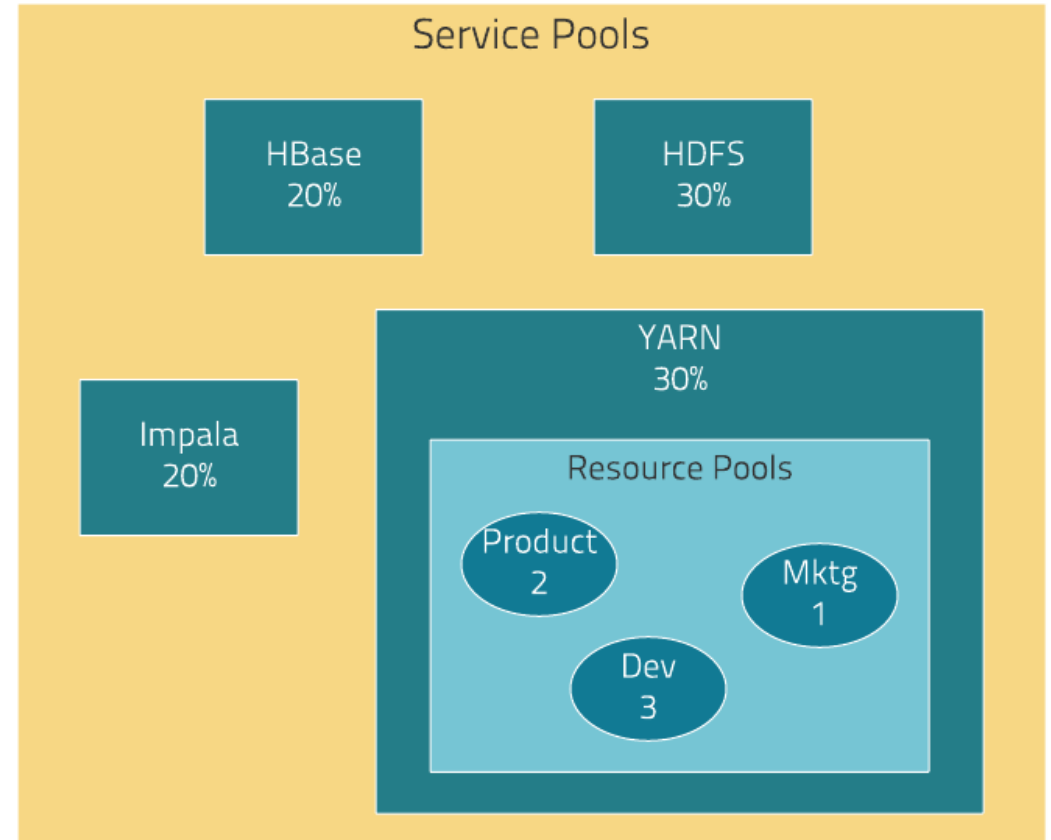


YARN scheduler (Dynamic resource allocation)

- Capacity or fair scheduler (Cloudera recommended) should be used to control to cluster resource by YARN applications
- Allocation is dynamic, based on queues
- Resources are divided between queues. If a queue is not allocating any resources, they can be distributed to other queues
- Access to queues can be restricted based on user/group executing the YARN job
- Works with: [MapReduce](#), [Spark](#), [Hive](#), [Oozie](#) , ...

Cgroups (Static resource allocation)

- Applications outside of YARN need to be tamed
- Linux Control Groups (cgroups) allows for per-resource isolation between services and roles
- Services are allocated a static percentage of total resources:
 - CPU shares
 - I/O weight
 - Memory usage



cloudera®

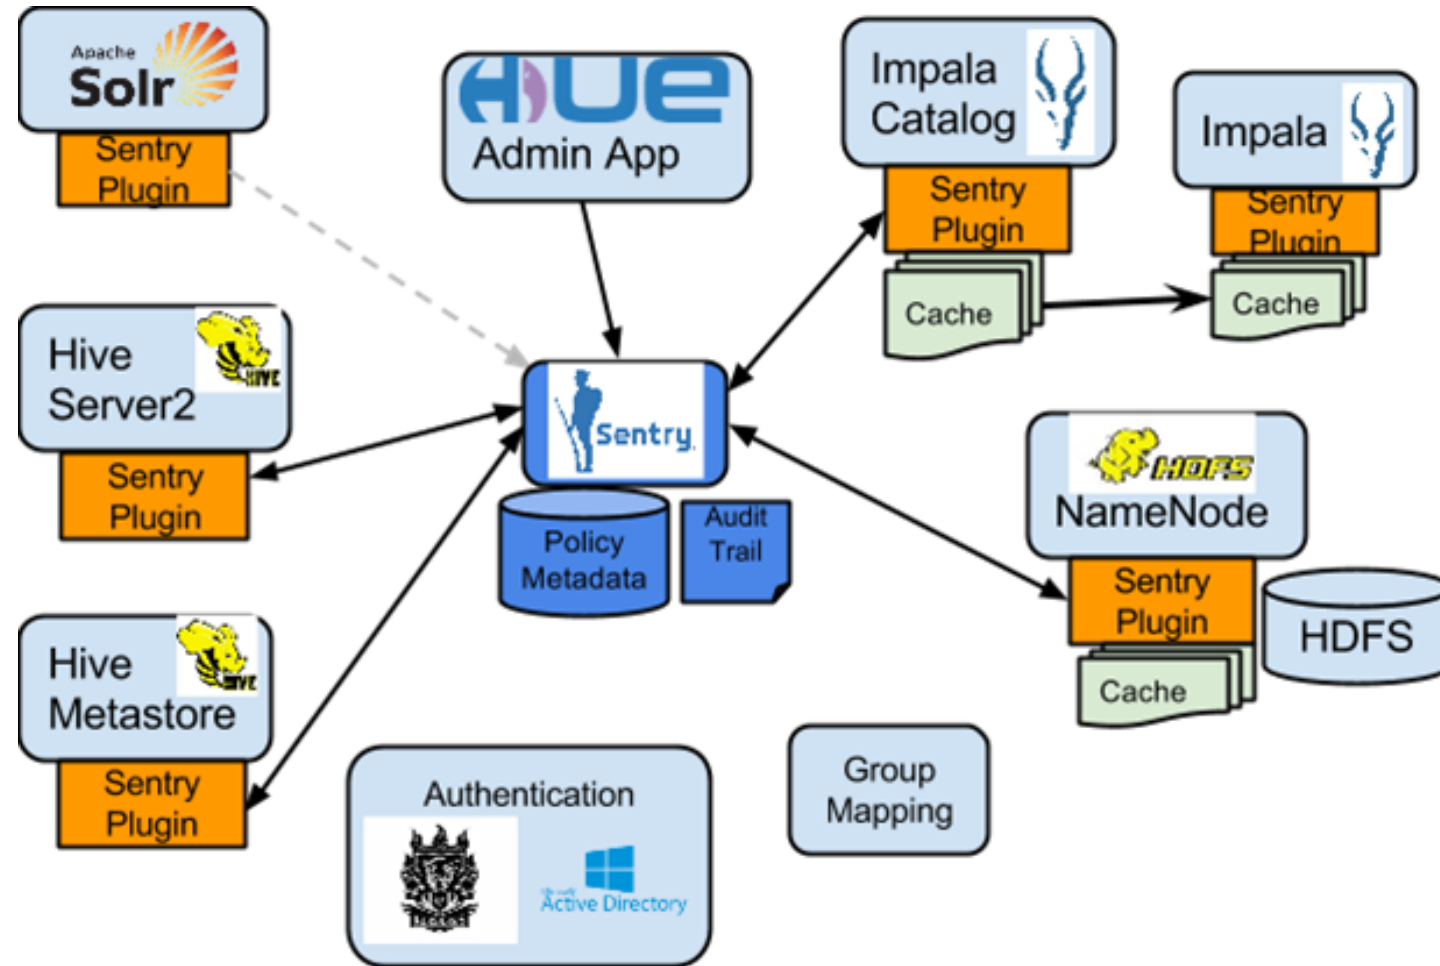
Security



Authentication / Authorization

- We need to know who the users are
 - We need to know which groups they belong to
 - We need to know what can they access
 - And what level of permissions they have
- Kerberos is the only authentication method supported by most components
 - LDAP can be used for some components (HiveServer 2 / Impala)
 - LDAP allows group management by integrating with Identity Management solutions (AD, Centrify, SSSD)

Apache Sentry I



Apache Sentry II

- It's an authorization service usable by many components
- Familiar SQL syntax, manages permissions, stores them in private database
- Role-based access control /GRANT SELECT ON TABLE data TO Analyst/
- **Objects** (Hive/Impala) are: server, database, table, column, HDFS URI
- Objects are mapped to HDFS directories for jobs outside of Hive/Impala
- **Roles** are mapped to groups /GRANT ROLE Analyst TO GROUP finance-managers/
- **Permissions** SELECT(rx), INSERT(wx) and ALL(rwx) are mapped to POSIX file permissions outside of Hive/Impala

Encryption

Data in-transit

- SSL/TLS needs to be enabled to encrypt data between clients and services' endpoints (Hive, Hue, ...)
- Certificates and key management tasks are usually outside of scope of Hadoop cluster
- Keys and certificates are configured per service/role

Data at-rest (Key Trustee)


- Multiple encryption zones on HDFS allow only authorized users to access the data.
- Data is transmitted in encrypted form as the encryption is on HDFS block level.
- Keys can be stored in Java keystore or HSM

HDFS considerations

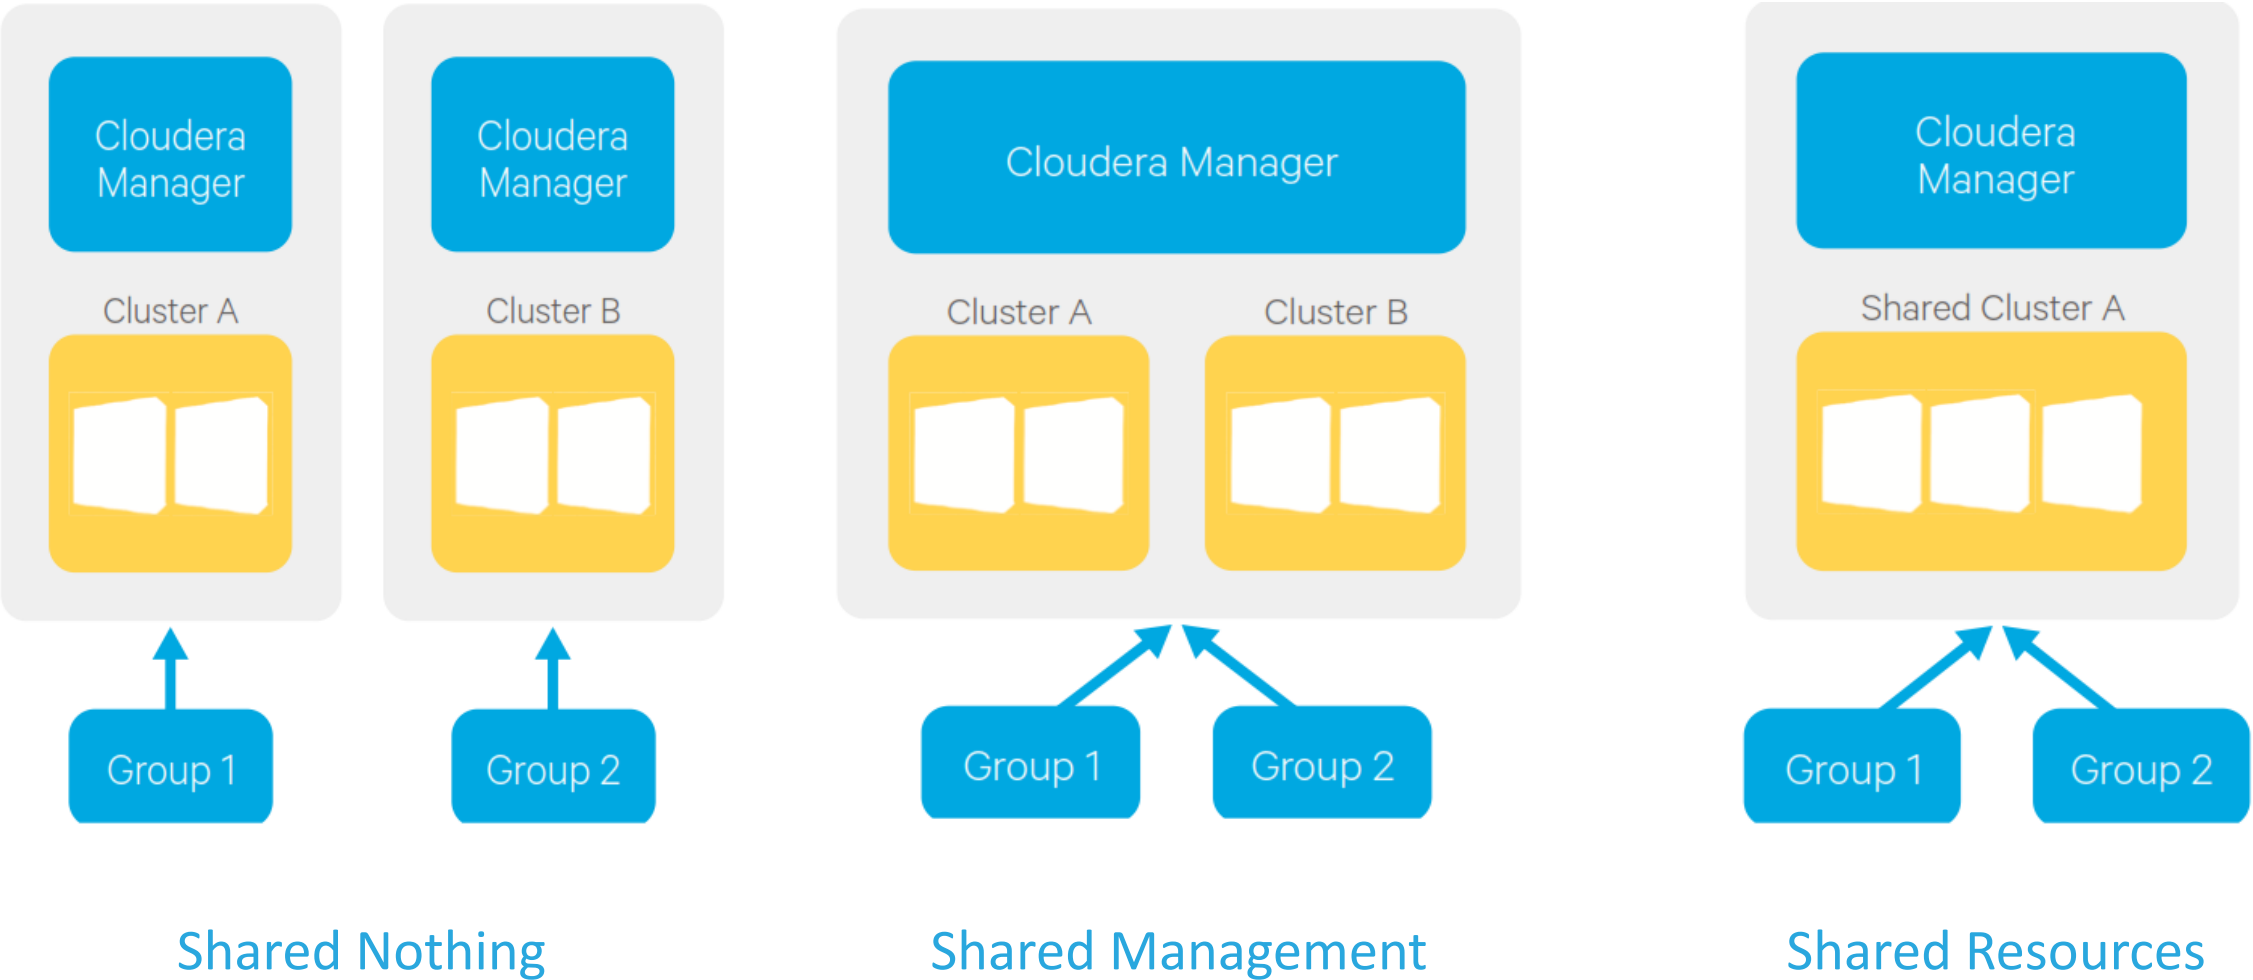
- Organize your data, think namespaces (directory structure and name conventions)
- Make sure nobody uses too much space, enable HDFS quotas
- Unix file permissions are not enough, enable ACLs
- If using Sentry, enable Sentry HDFS sync plug-in

cloudera®

Operations / Managing the Cluster



Managing the cluster



Reports on user activity

Monitor, monitor, monitor!

- “How much CPU & memory did each tenant use?”
- “I set up fair scheduler. Did each of my tenants get their fair share?”
- “Which tenants had to wait the longest for their applications to get resources?”
- “Which tenants asked for the most memory but used the least?”
- “When do I need to add nodes to my cluster?”

Cloudera Manager reports

Cluster Utilization Report (Cluster 1)

Configuration: Default 02/22/2016 - 03/22/2016

Overview YARN Impala


CPU Utilization

Overview YARN Impala


Utilization Capacity Planning Preemption Tuning

Tenant Capacities

Overall Cluster Utilization



YARN + Impala Utilization



Utilization By Tenant



vc0102.halg.cloudera.com:7180/cm/

Tenant Name
root.cmjobuser
root.hospool
root.systest
root.pigpool
root.hivepool
root.swimpool
root.yarnpool
root.hdfs
hospool
root.testyarnpool1
root.impalatpdcspool
default
root.default


Cluster Utilization Report (Cluster 1)

Configuration: Default 04/10/2016 - 05/09/2016

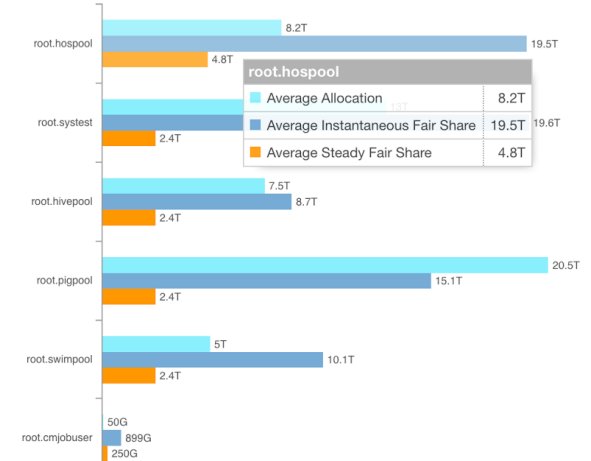
Overview YARN Impala

Utilization Capacity Planning Preemption Tuning

CPU



Memory



How to start?

Start small

- 2-3 tenants

Plan ahead!

- user management
- data governance

Configure Kerberos

- http://www.cloudera.com/documentation/enterprise/latest/topics/cm_sg_authentication.html

Enable HDFS ACLs:

```
<property>
  <name>dfs.namenode.acls.enabled</name>
  <value>>true</value>
</property>
```

Enable fair scheduler:

- http://www.cloudera.com/documentation/enterprise/latest/topics/admin_fair_scheduler.html

Look into Sentry:

- http://www.cloudera.com/documentation/enterprise/latest/topics/cm_sg_sentry_service.html



cloudera
Thank you

Dániel Schöberle
dschoberle@cloudera.com