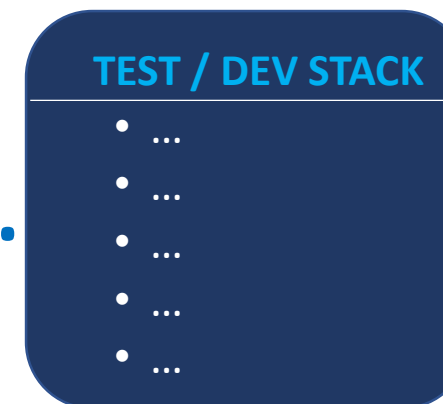





Thin Fedora O/S + Docker Engine + LXC Engine





Thin **Fedora O/S** + **Docker Engine** + **LXC Engine**




- TEST / DEV STACK**
- Storm Cluster
 - Redis Cluster
 - Cassandra Cluster
 - Kafka Brokers

- TEST / DEV STACK**
- Cloudera Cluster
 - CDH-5 / CDH-6
 - CDH Mgt Console


LXC/Linux Container Example Use-Case

Circa 2012, while architecting a **Real-Time Trade-Stream Anomaly Analytics** platform for BA/ML, I used **Apache Storm** for Complex Event Processing. During the design I discovered a show-stopping deficiency with Storm and needed to reverse-engineer its behavior under failure scenarios; to see if/how I could work around them. At the time Storm had only a small GitHub page discussing it, and no books. And since my query to the Storm community about this issue went unanswered (<http://bit.ly/1bsBooT>), I used these LXC containers to instrument behaviors myself. I devised a solution based on observations and published my work and methods here: <http://bit.ly/prisma003>





Thin **Fedora O/S** + Docker Engine + **LXC Engine**



- TEST / DEV STACK**
- Storm Cluster
 - Redis Cluster
 - Cassandra Cluster
 - Kafka Brokers

- TEST / DEV STACK**
- Cludera Cluster
 - CDH-5 / CDH-6
 - CDH Mgt Console











LXC/Linux Container Example Use-Case
Like **Apache Storm**, stacks requiring a proper O/S are deployed in **LXC Containers**. Another example of this is a **Cludera Hadoop** cluster, where Docker containers are neither provided nor supported.

Six (qty. 6) LXC Containers serve this CDH use-case:

- vps00: **Master** node for big data services
- vps01: **Worker** node for shards / executors
- vps02: **Worker** node for shards / executors
- vps03: **Worker** node for shards / executors
- vps04: **Worker** node for shards / executors
- vps10: CDH SCM **Management** node

Issuing `lxc-start(1)` against these six (qty. 6) containers erect this CDH cluster on-demand.

✔ **PRISMALYTICS** (CDH 6.1.0, Parcels) ▾




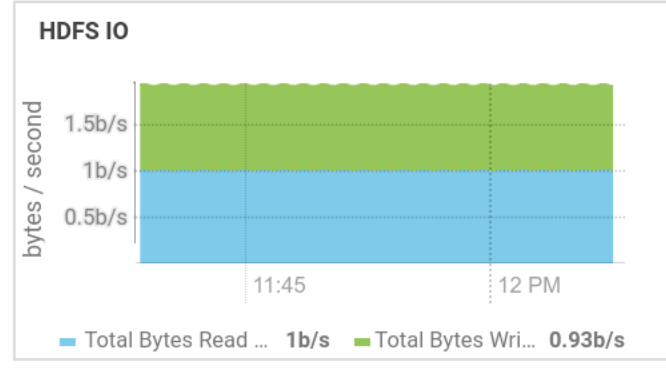
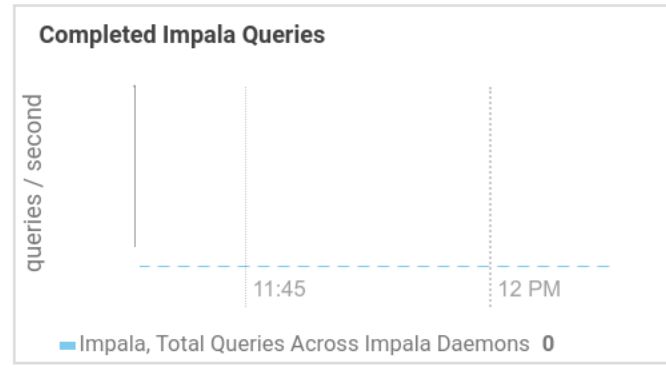
- ✔  6 Hosts
- ✔  HBase ▾
- ✔  HDFS ▾
- ✔  Hive ▾
- ✔  Hue ▾
- ✔  Impala ▾
- ✔  Oozie ▾
- ✔  Spark ▾
- ✔  YARN (MR2 In... ▾
- ✔  ZooKeeper ▾


Cloudera Management Service

- ✔  Cloudera Man... ▾

Charts

30m 1h 2h 6h 12h 1d 7d 30d 






Thin Fedora O/S + Docker Engine + LXC Engine

DOCKER Container Example Use-Case

Conversely, stack components with lightweight O/S dependencies are deployed in **Docker Containers** and run as coordinated services via **docker-compose.yml** specs. For the stack shown, Docker provides **Kafka**, **ZooKeeper** and **MongoDB** service endpoints that **Spark 2.x** connects to; thus erecting a streaming-data analytics pipeline on-demand. The service definition specifies three Kafka brokers, which is needed to simulate distribution of messages across topic-partitions by partition key. **Tearing down a stack and replacing it with another is measured in minutes.** Some results and Jupyter Notebooks derived from using these stacks are shared here: <https://jupyter.ai>

Docker Containers (Hub + Docker Files)



TEST / DEV STACK

- Spark 2.x PySpark
- Kafka Brokers x 3
- ZooKeeper
- MongoDB
- Python-3 VENV

TEST / DEV STACK

- ...
- ...
- ...
- ...
- ...

