Capital One Docker analytic garage

dockercon 15

SF JUNE 22-23

Capital One at a glance

- **A leading diversified bank with $306.2 billion in assets, $204.0 billion in loans and $210.4 billion in deposits**
  - 8th largest bank based on U.S. deposits[1]
  - 5th largest retail depository institution in metro New York[2]
  - Largest consumer and commercial banking institution headquartered in the Washington, DC region
  - 4th largest credit card issuer in the U.S.[3]
  - The 3rd largest issuer of small business Visas and MasterCards in the U.S.[4]
  - The 3rd largest independent auto loan originator[5]
  - Largest US direct bank[6]
- **Major operations in 15 U.S. cities, Canada, U.K.**
- **More than 65 million customer accounts and 46,000 associates**
- **A FORTUNE 500 Company - #124**
- **Numerous recent awards including:**
  - Named to 100 Best Companies to Work For by *FORTUNE* Magazine
  - No. 2 on 2014 *Information Week* Elite 100
  - Received J.D. Power & Associates Call Center Certification
  - Aon Hewitt's Top Companies for Leaders
  - Named to *Working Mother's* 100 Best Companies list & Best Companies for Hourly Workers
  - Ranked #8 on *Military Times*' 2015 "Best for Vets"
  - Recipient of the Secretary of Defense Employer Support Freedom Award
  - Named a "NAFE Top Companies For Executive Women"

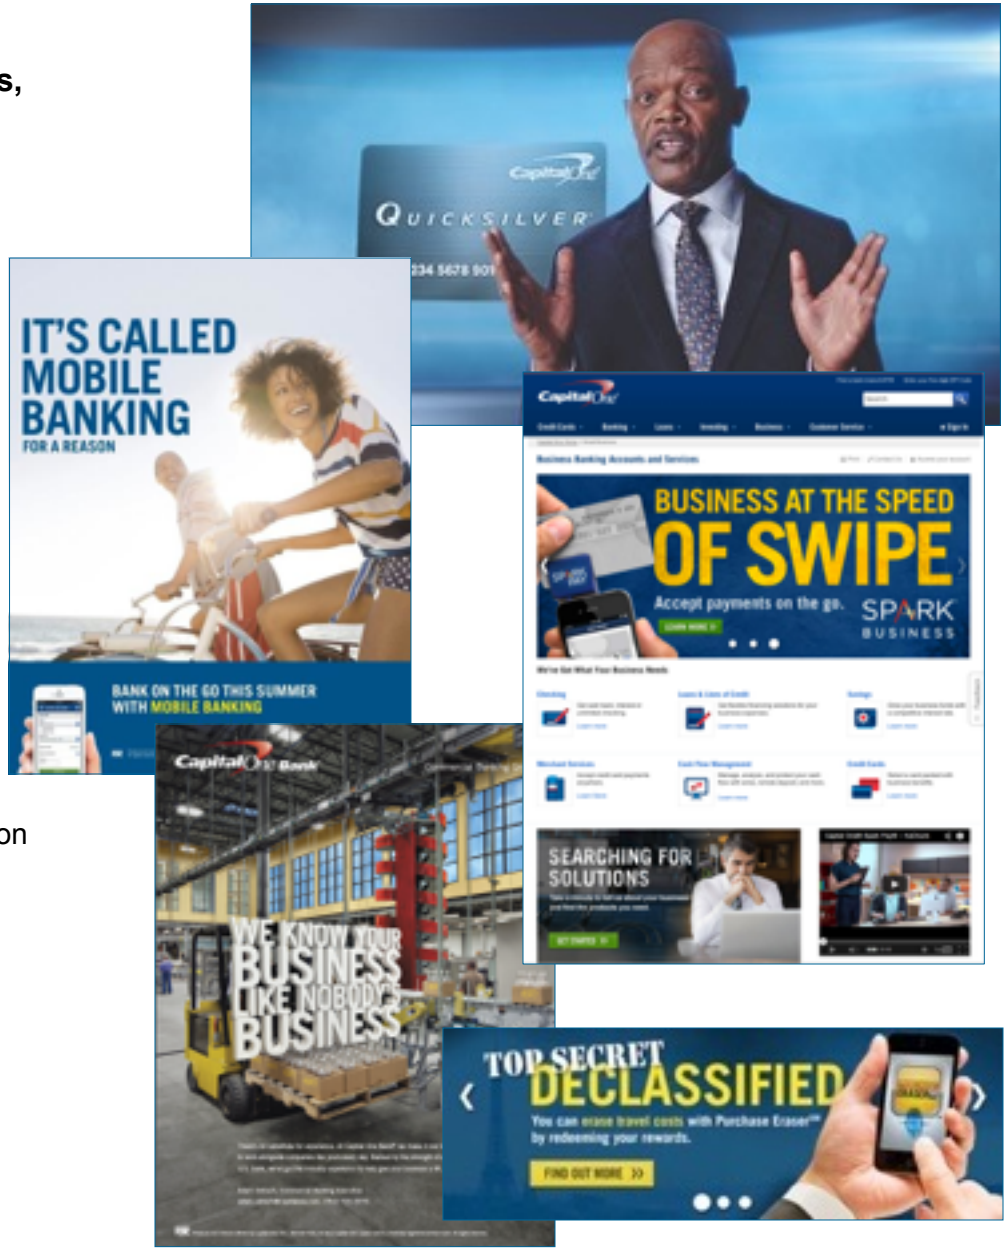

1) Domestic deposits ranking as of Q4'14
2) Source: FDIC, June 2014, deposits capped at $1B per branch
3) Company-reported domestic credit card outstandings, Q1'15, American Express GNS, Discover Cards
4) Source: Nilson Report, Q4'13
5) Source: JD Power, 2014
6) FDIC, company reports as of Q4'14

We are building the technology foundation to ensure our analytics leadership as we move to the world of Big Data

Key objectives

Build an analytics architecture centered on a Hadoop-based Enterprise Data Hub

Provide state-of-the-art analytical tools, unconstrained data storage & processing

Empower our associates to dream and disrupt

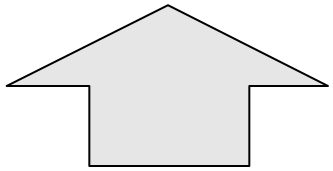

Delivery principles

- **Fast prototyping, scaled agile delivery**
- **Smaller, cross functional teams and integration of new talent**
- **Collaboration and leverage the power of Open Source**

A key objective of our Big data journey is to enable access to the best tools for all Associates

- Self-evaluation
- Local testing
- On-demand evaluation environment
- Rapid progression from evaluation to production
- Workload isolation
- Tool governance

We had to engineer an effective architecture to evaluate and integrate the variety and volume of tools / SW packages

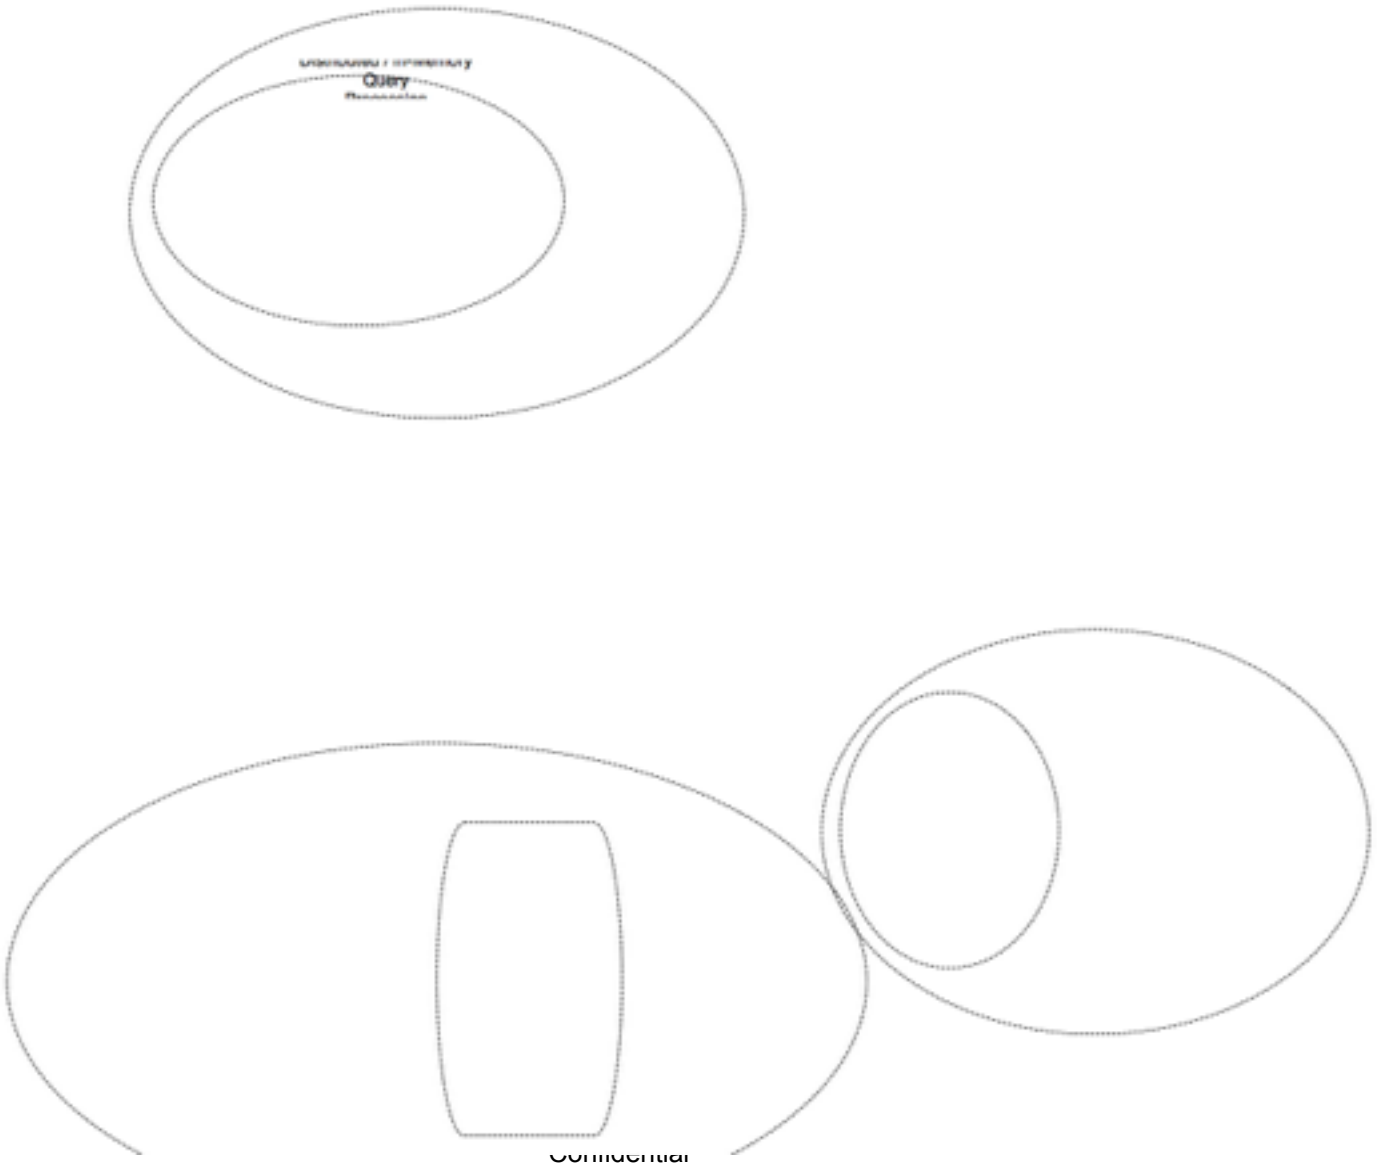

The analytic garage was stood up to create a separate environment for users to fast prototype new tools / insights with integrated resource management

Version 0.1

Bare metal
RHEL 6.x

- Instability
- Very few users / tools per server

Version 0.8

Bare metal
RHEL 6.x
cgroups
LXC

- Some stability
- Limited users / tools per server
- Challenges with speed of tool evaluations

Version 1.0

Mesos
Marathon
Docker
cgroups
RHEL 6.x
Gluster FS

- Improved stability
- Improved resource utilization, more users, more tools
- Self-evaluation
- Isolated workload

The analytic garage was integrated with the rest of Capital One's Big Data ecosystem to enable agile progression of insights to deployment

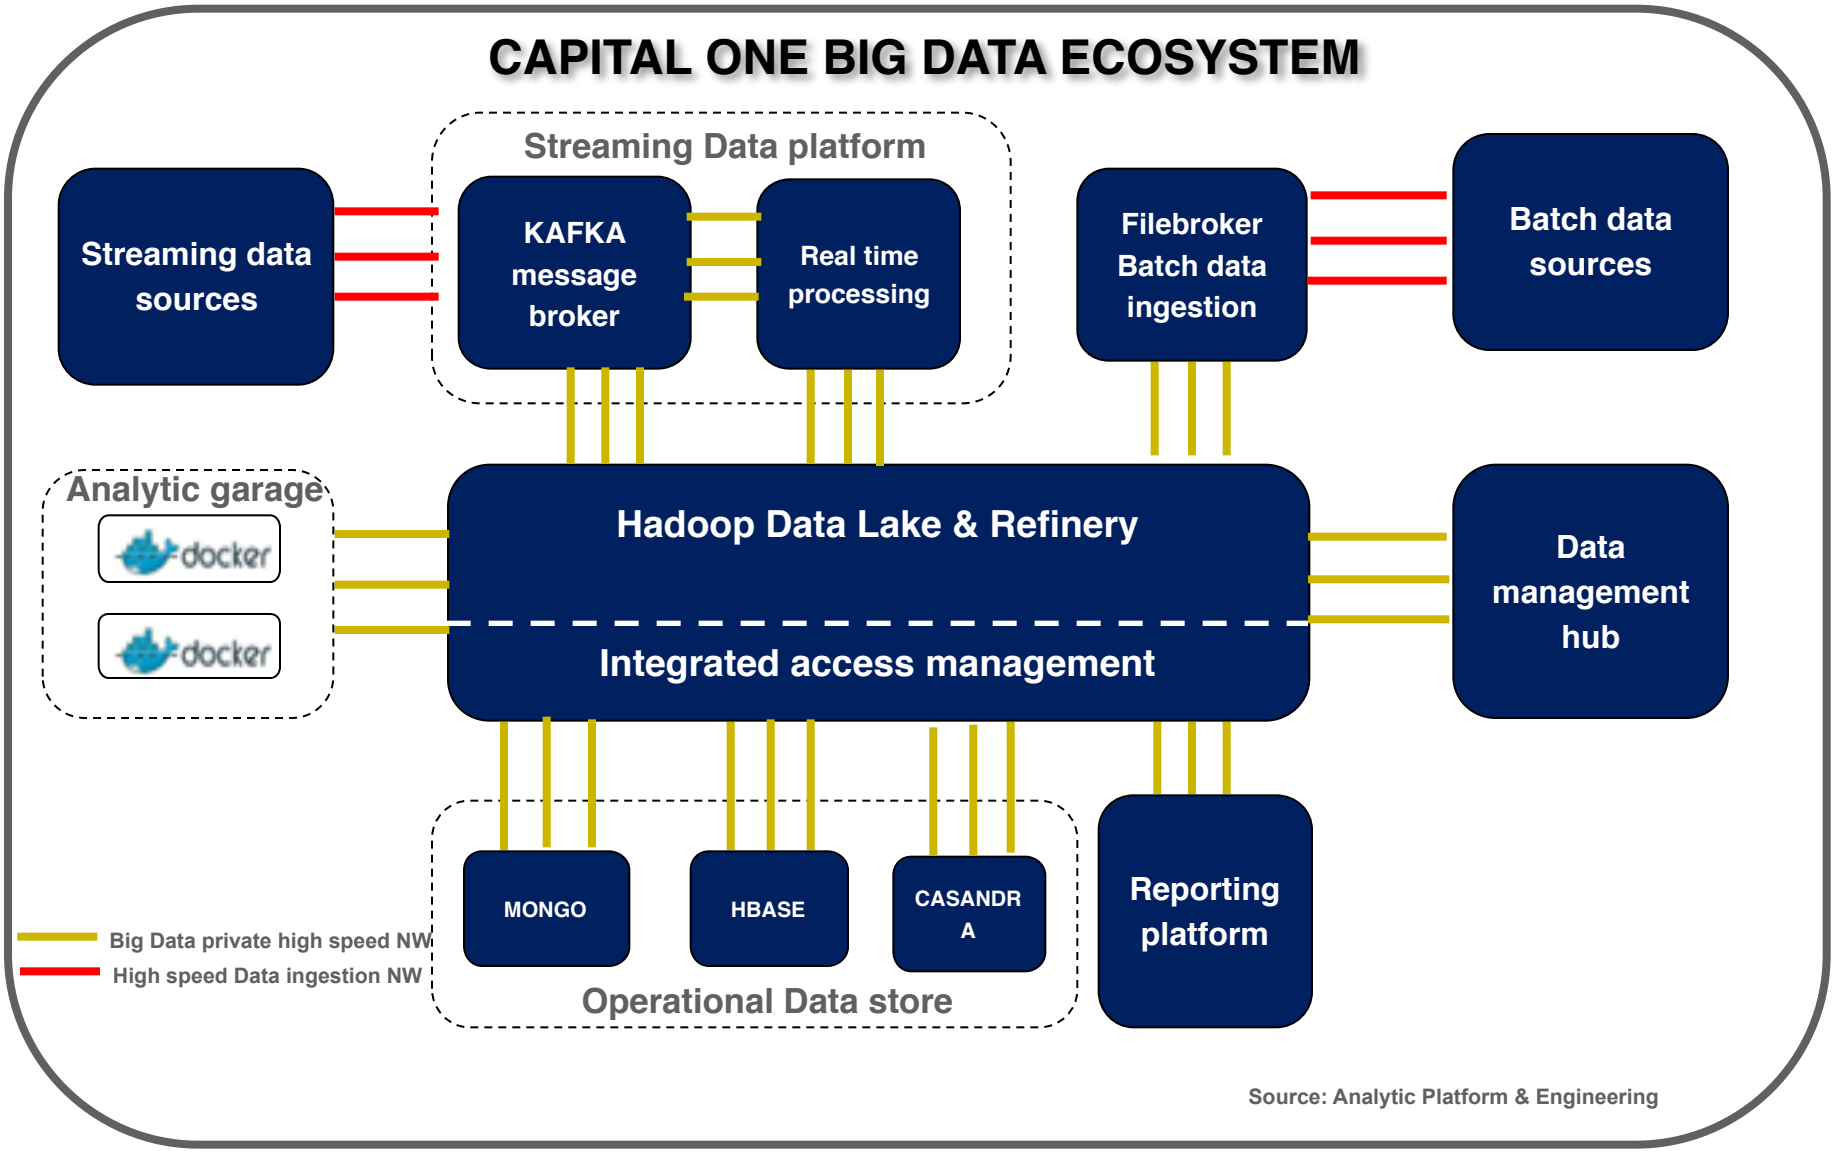

To increase adoption of the Docker analytic garage across the analyst community we developed a self-service UI

- Web portal to instantiate containers & analytic services
- Kerberos integration with Hadoop, Hive
- Integrated monitoring and metrics
- Lifecycle management (container expiration)
- Highly available cluster with Mesos-Marathon
- Shared storage using GlusterFS

Docker analytic garage workflow

New analytic container workflow

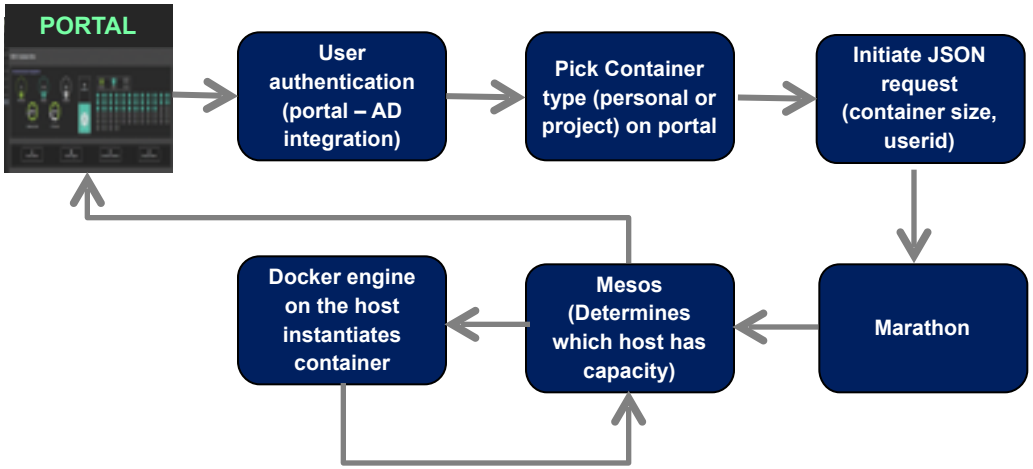

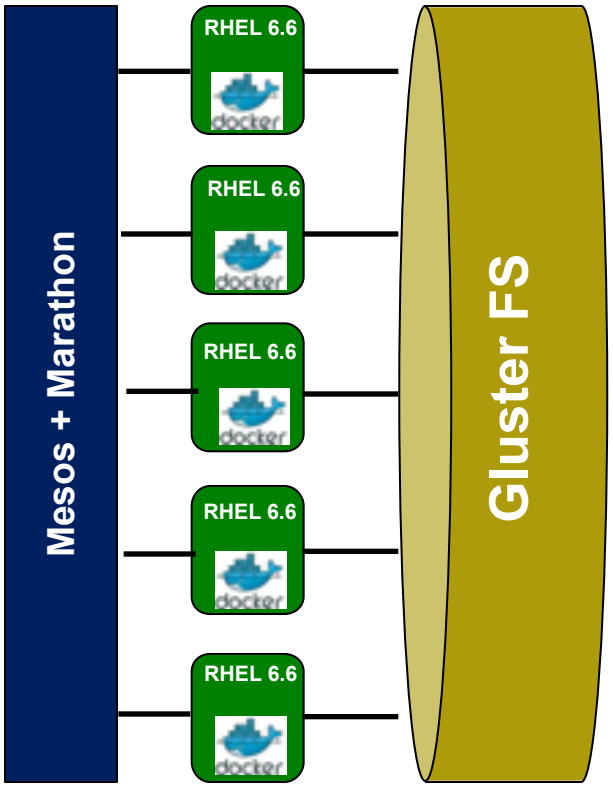

To minimize complexity of adoption, we created a virtual private server by integrating multiple analytic services, applications & tools into one Docker image

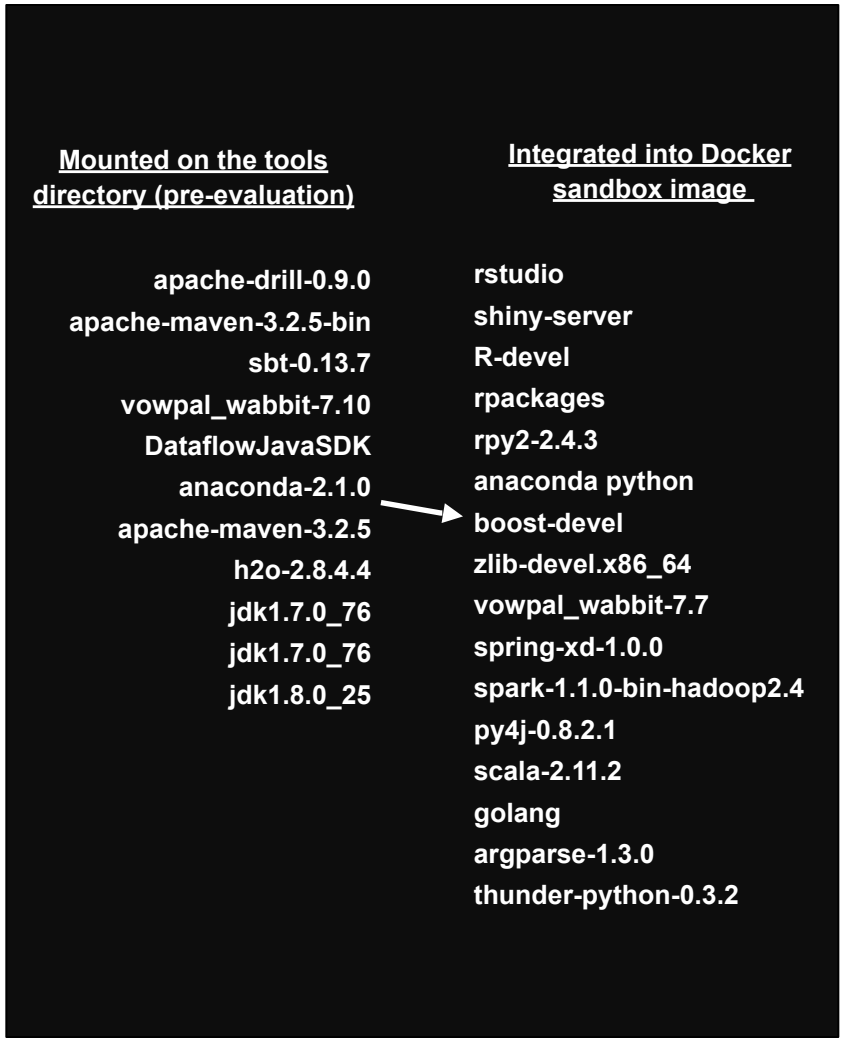

- Familiar data centric sandbox image
- Maximize portability & performance
- Allow for hybrid tool use
  - Rpy2 (Python + R)
  - rPython (R + Python)
  - Sparkling Water (Spark+H2O)
  - R + H2O
- Reproducibility and Auditability with a Versioned Environment
- Volume mounted Tool directory to screen new tools before integrated into the sandbox image
- In spite of the size of the image, we are able to instantiate containers in seconds

Some of the challenges we encountered were around the overall ecosystem surrounding Docker

- **Challenges Coordinating the Initialization order of the Services**
  - Gluster, Docker, Mesos-Master, Mesos-Slave, Marathon
- **Gluster:**
  - Open Source Gluster Resiliency is Fragile
  - RHEL Gluster is Much more Reliable
  - Quotas are not effective
- **Docker is not fully supported or stable on 2.x Linux Kernels**
  - Cgroups Bug, Random Reboots
- **Devicemapper is much too complex for use with Docker**
  - Too many moving parts
  - The “thin” tools
- **Networking**
  - IPtables configuration is fragile in complex mutli-home network topologies
  - HA Proxy updates via Marathon scripting are sluggish for VPS deployments

Docker garage has significantly reduced the time taken to evaluate and onboard new tools & solutions and has helped accelerate the evolution of our Data Technology Strategy (DTS)

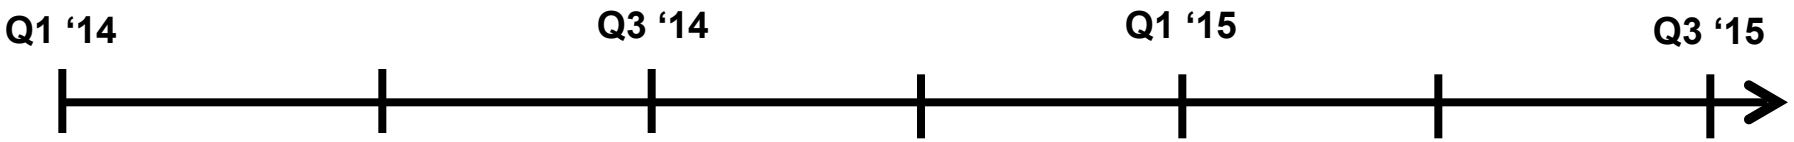

DTS 1.0

Start of our Big Data journey with a central Hadoop Lake

DTS 2.0

Flood of tools and operationalize Big Data use cases

DTS 3.0

Fast Data, Real-time, cross-data center resiliency

The analytic garage has enabled us to build, test and iterate complete application prototypes using a ‘Lego block’ approach

‘IT analytic portal’ tested and prototype deployed on the analytic garage

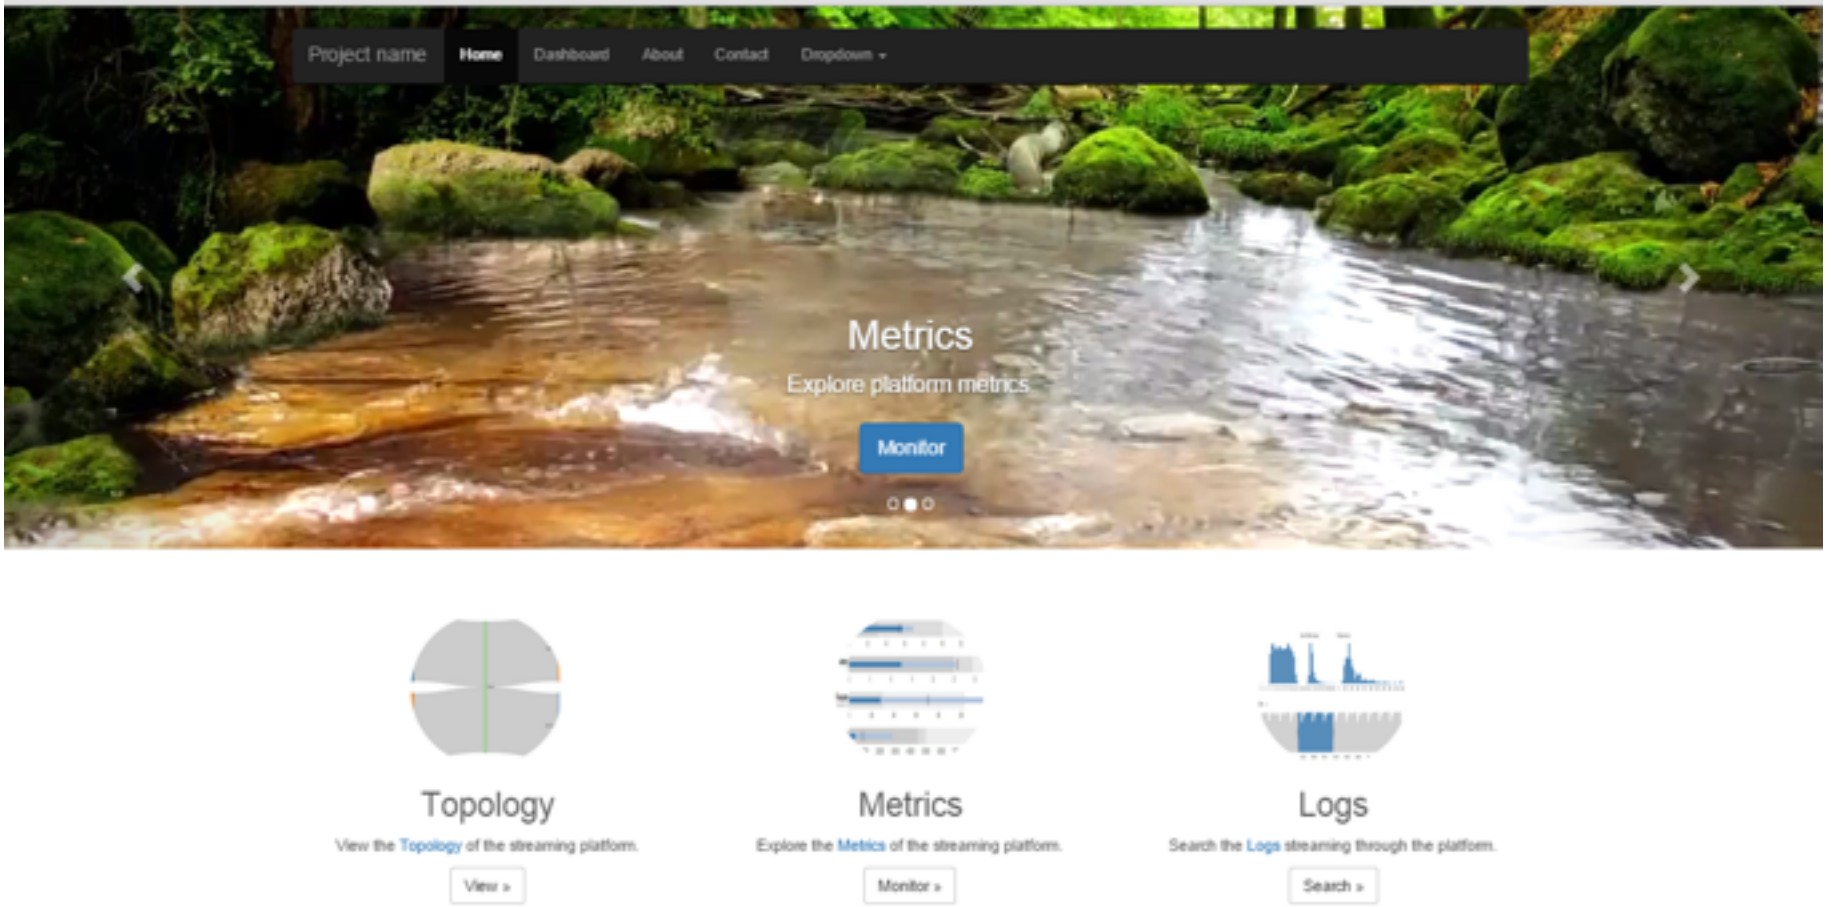

What's in your wallet ? A personal analytic garage powered by Boot2docker

Performance within Docker containers is comparable to bare metal , enabling our analysts to run complex models

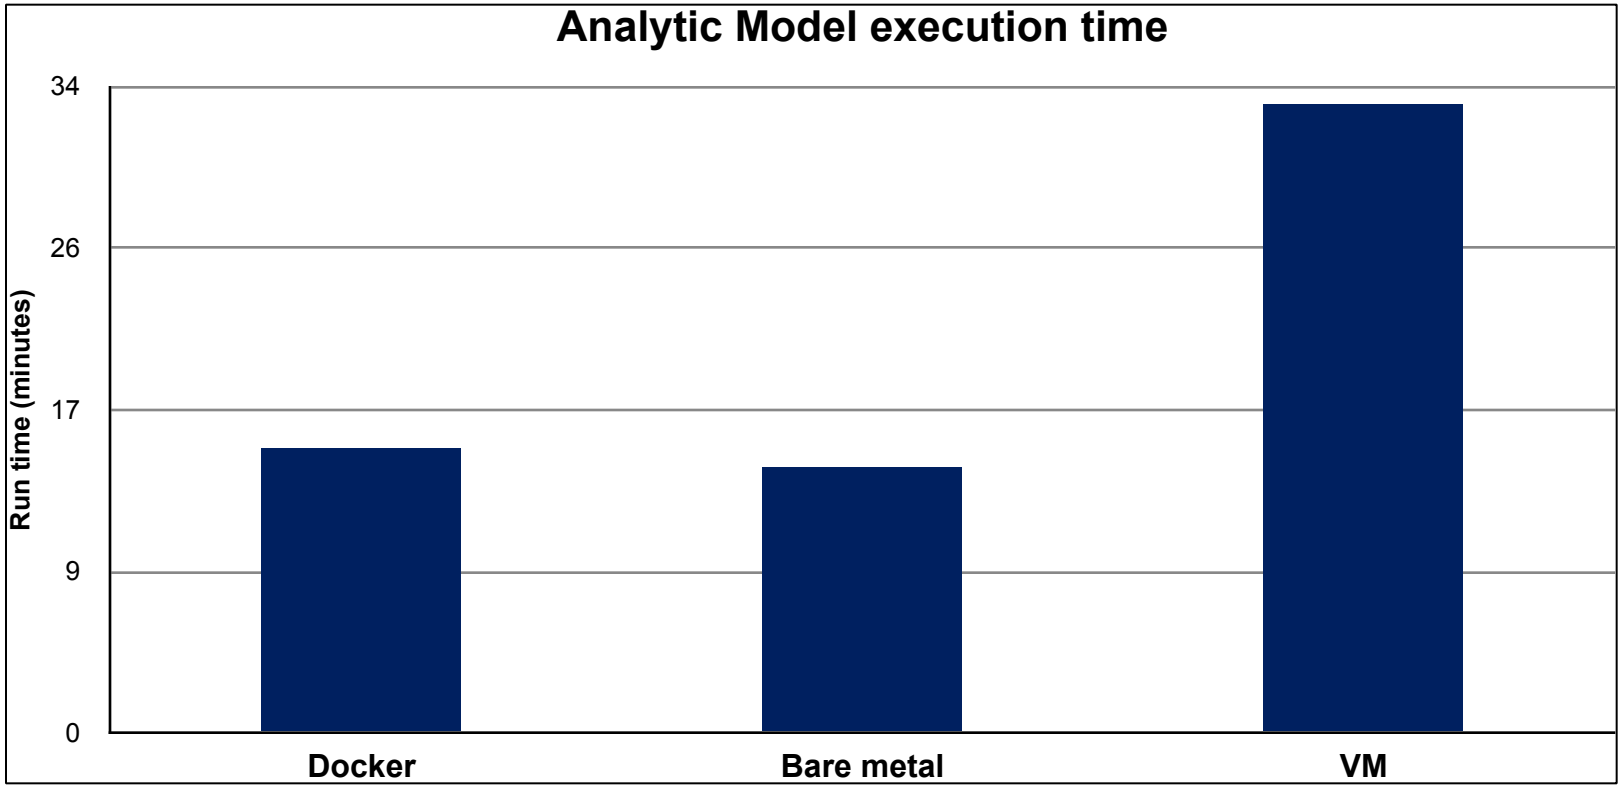

Docker Security and monitoring

- **Audit, Audit, Audit**
  - What goes into the image and changes to the image
  - User level auditing
  - No direct user access to bare metal host, access only to containers
- **Kerberos Ticket Required for Hadoop/HDFS Access**
- **Nested Firewalling, ACLs:**
  - Only Specific Users Allowed Access
- **Single User per Container facilitates accountability and auditability**
- **All container processes visible to host OS**
  - nsenter allows auditing of internal container activity

We have been partnering with Docker to evolve our use cases on a Docker only stack prior to running production workloads

Docker only stack

- Docker
- Swarm
- Compose
- Socketplane/ Openvswitch
- Docker Machine
- Cluster Local Private Registries
- Global DHE

Advantages

- Uniform provisioning methodology for build, dev, test, production, persistent and ephemeral containers using compose

Gaps

- No Native HA
- Without VXLan integration to network, HA Proxy/Portal is still required
- Limited ability to migrate containers across clusters

We are currently testing different approaches to persist databases on Docker as an enhancement to the analytic ecosystem

Integrate Db containers with streaming data platform

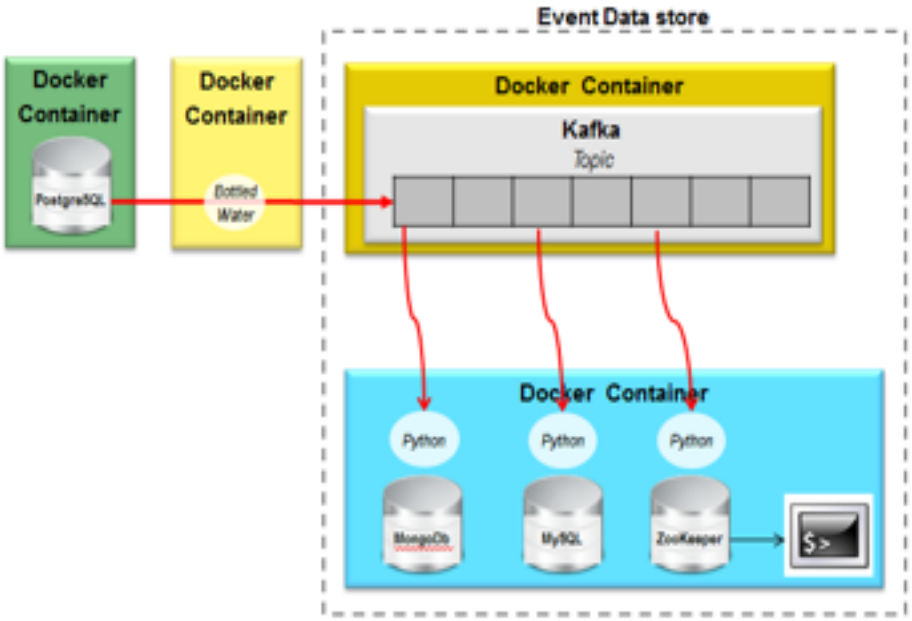

- Stream 'Write Ahead log' data from PostgreSQL into Kafka message broker as events
- Create an event data store (Kafka + Mongo/ Cassandra) for data persistence with SQL search capabilities (such as SOLR)

Integrate Db containers in a micro services architecture

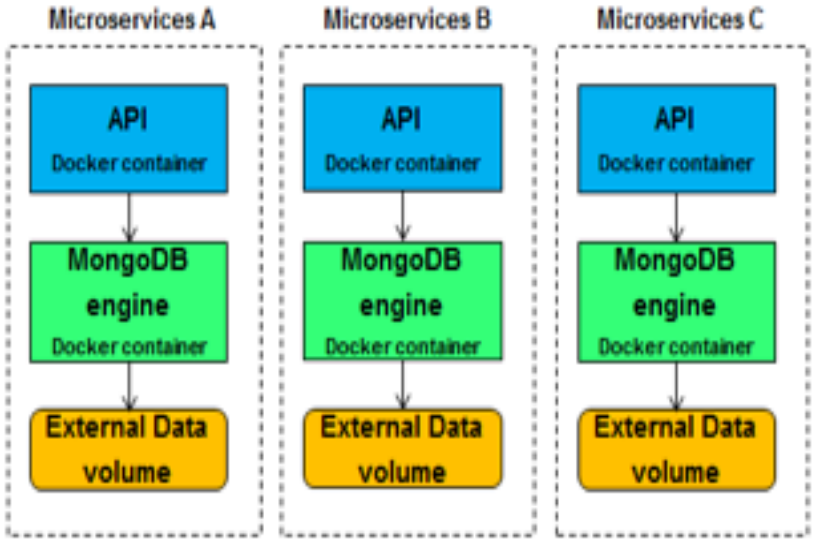

- Each micro service has a local API and Database engine container
- Persist data on external mounted volumes (entire Database/, schema/ partitions)
- Database engine container is ephemeral and can be swapped with newer SW versions
- Challenges – Scaling out Database engine / data stores

We have instituted a Docker working group across the company to collaborate and share learnings

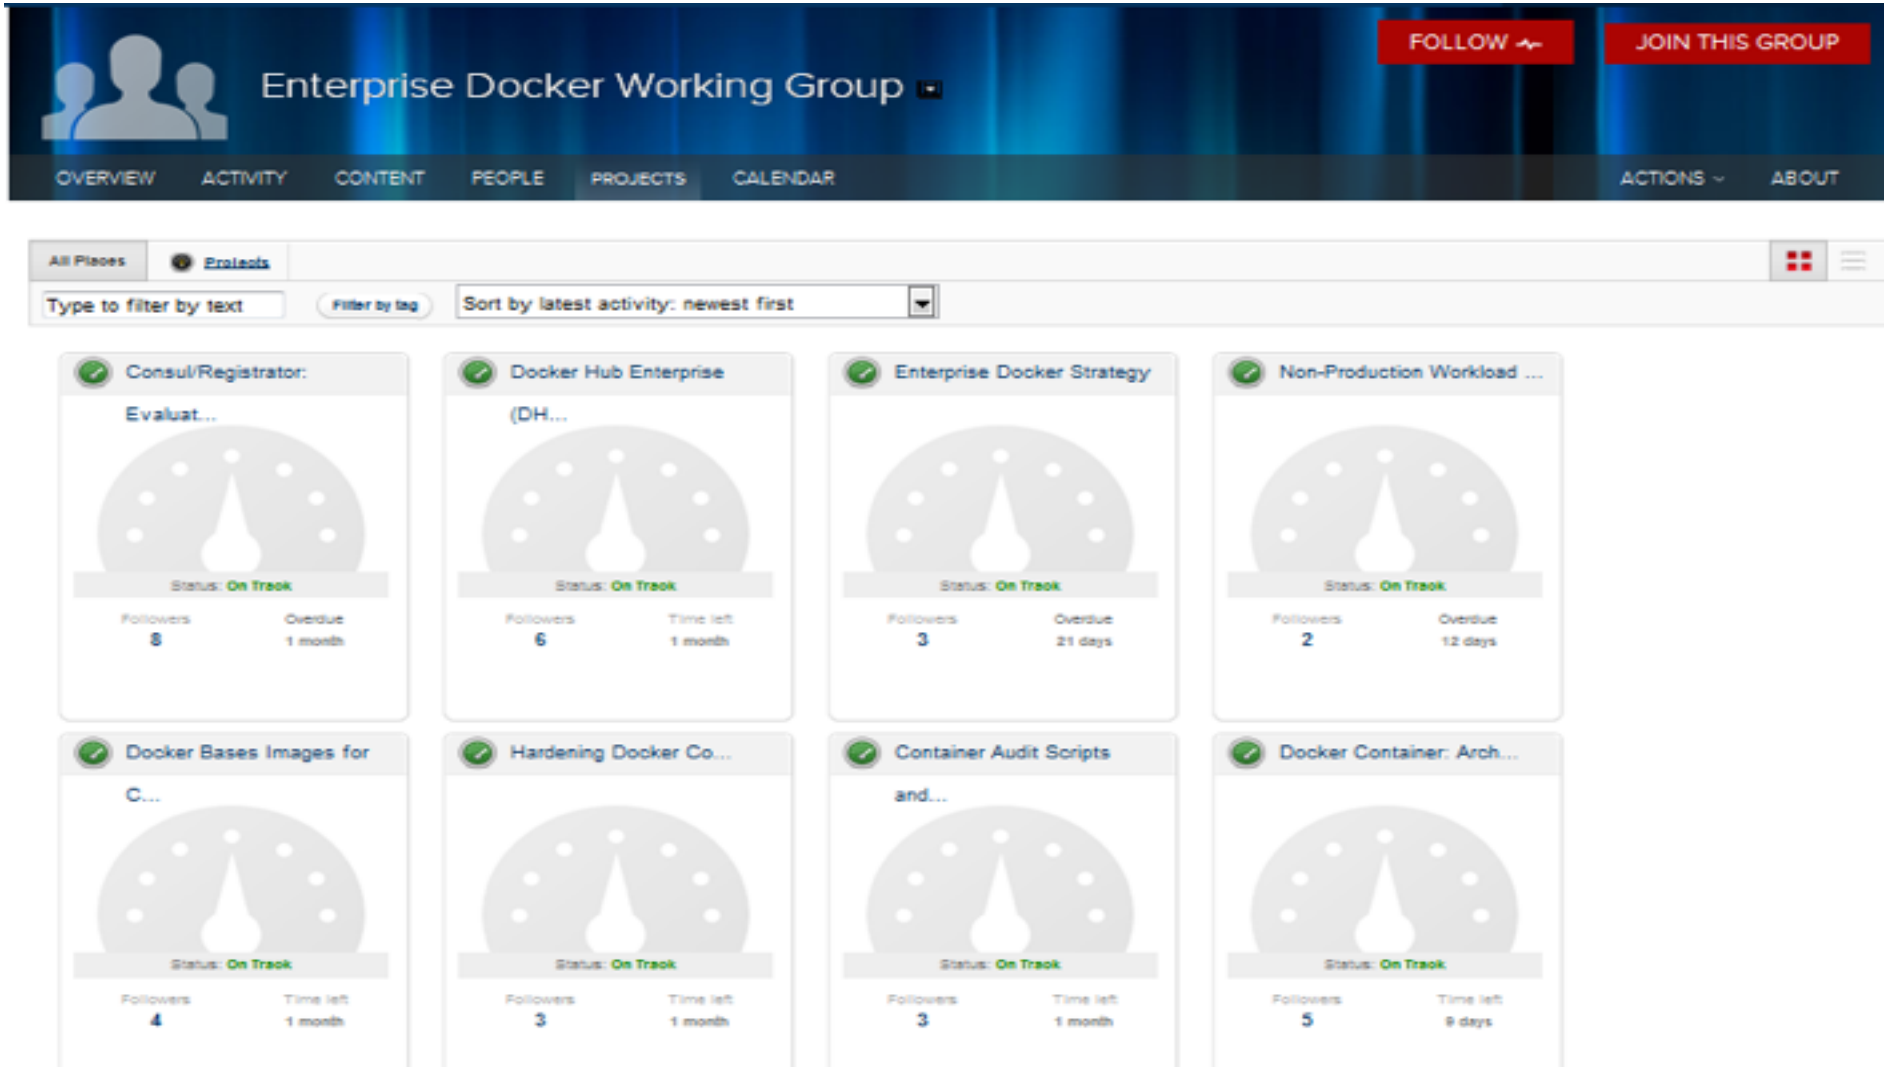

Thank you

Santosh Bardwaj

Senior Director, Engineering

@santoshkb15

https://www.linkedin.com/in/sbardwaj

William Scott Cochran

Master Software Engineer