vSphere - the Best Platform for Big Data

Bo Dong

Product Line Manager, VMware

dbo@vmware.com

Agenda

- **Hadoop Market Landscape**
- Hadoop Journey
- Virtualize Hadoop Values
- Summary
- Q & A

Data is exploding & Hadoop is driving growth

Unstructured data driving growth

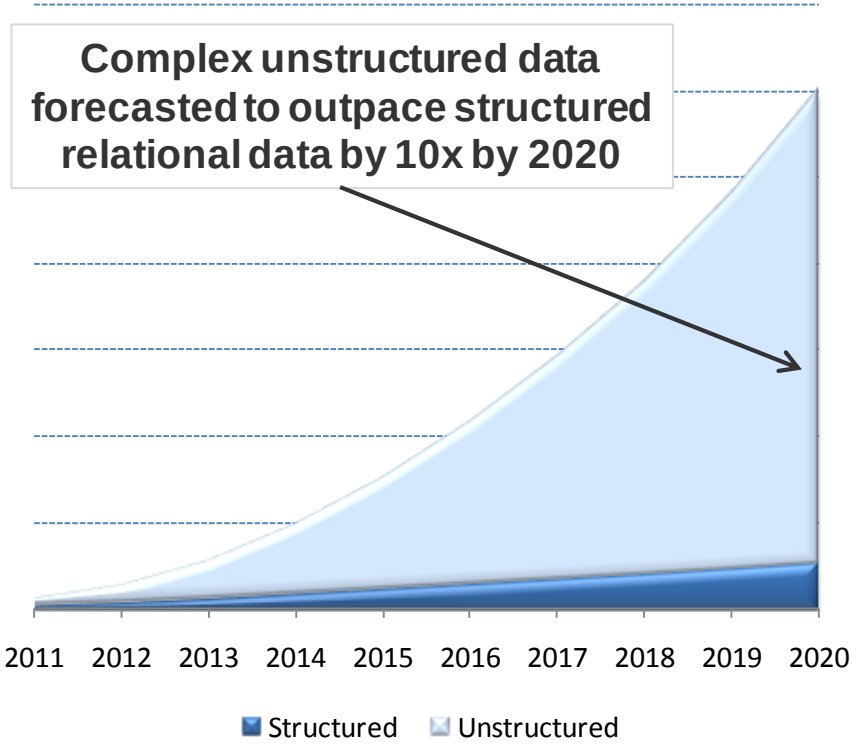

Hadoop adoption is ramping

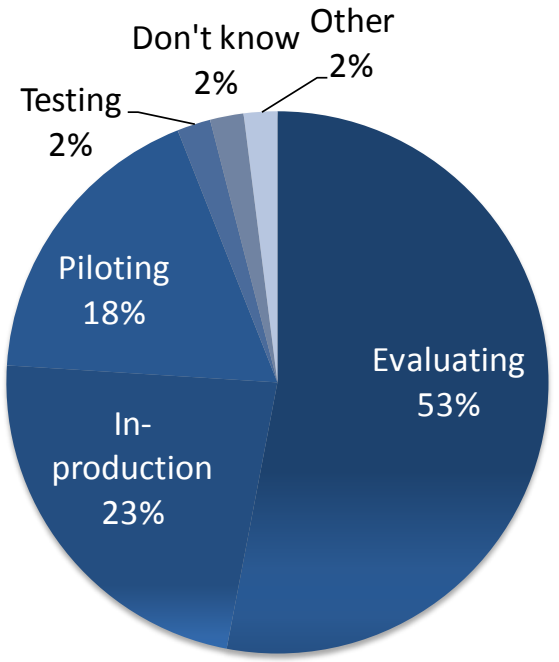

Source: Forrester Survey of 60 CIOs , September 2011

- **Unstructured data explosion and Hadoop capabilities causing CIOs to reconsider Enterprise data strategy**
  - Gartner predicts +800% data growth over next 5 years
  - Hadoop's ability to process raw data at cost presents intriguing value prop for CIOs

Today's Big Data System:

Real Time Streams

Real-Time Processing

(s4, storm)

Analytics

ETL

Real Time Structured Database

Big SQL

Data Parallel Batch Processing

Unstructured Data (HDFS)

Agenda

- Hadoop Market Landscape
- **Hadoop Journey**
- Virtualize Hadoop Values
- Summary
- Q & A

Hadoop Journey in Enterprises

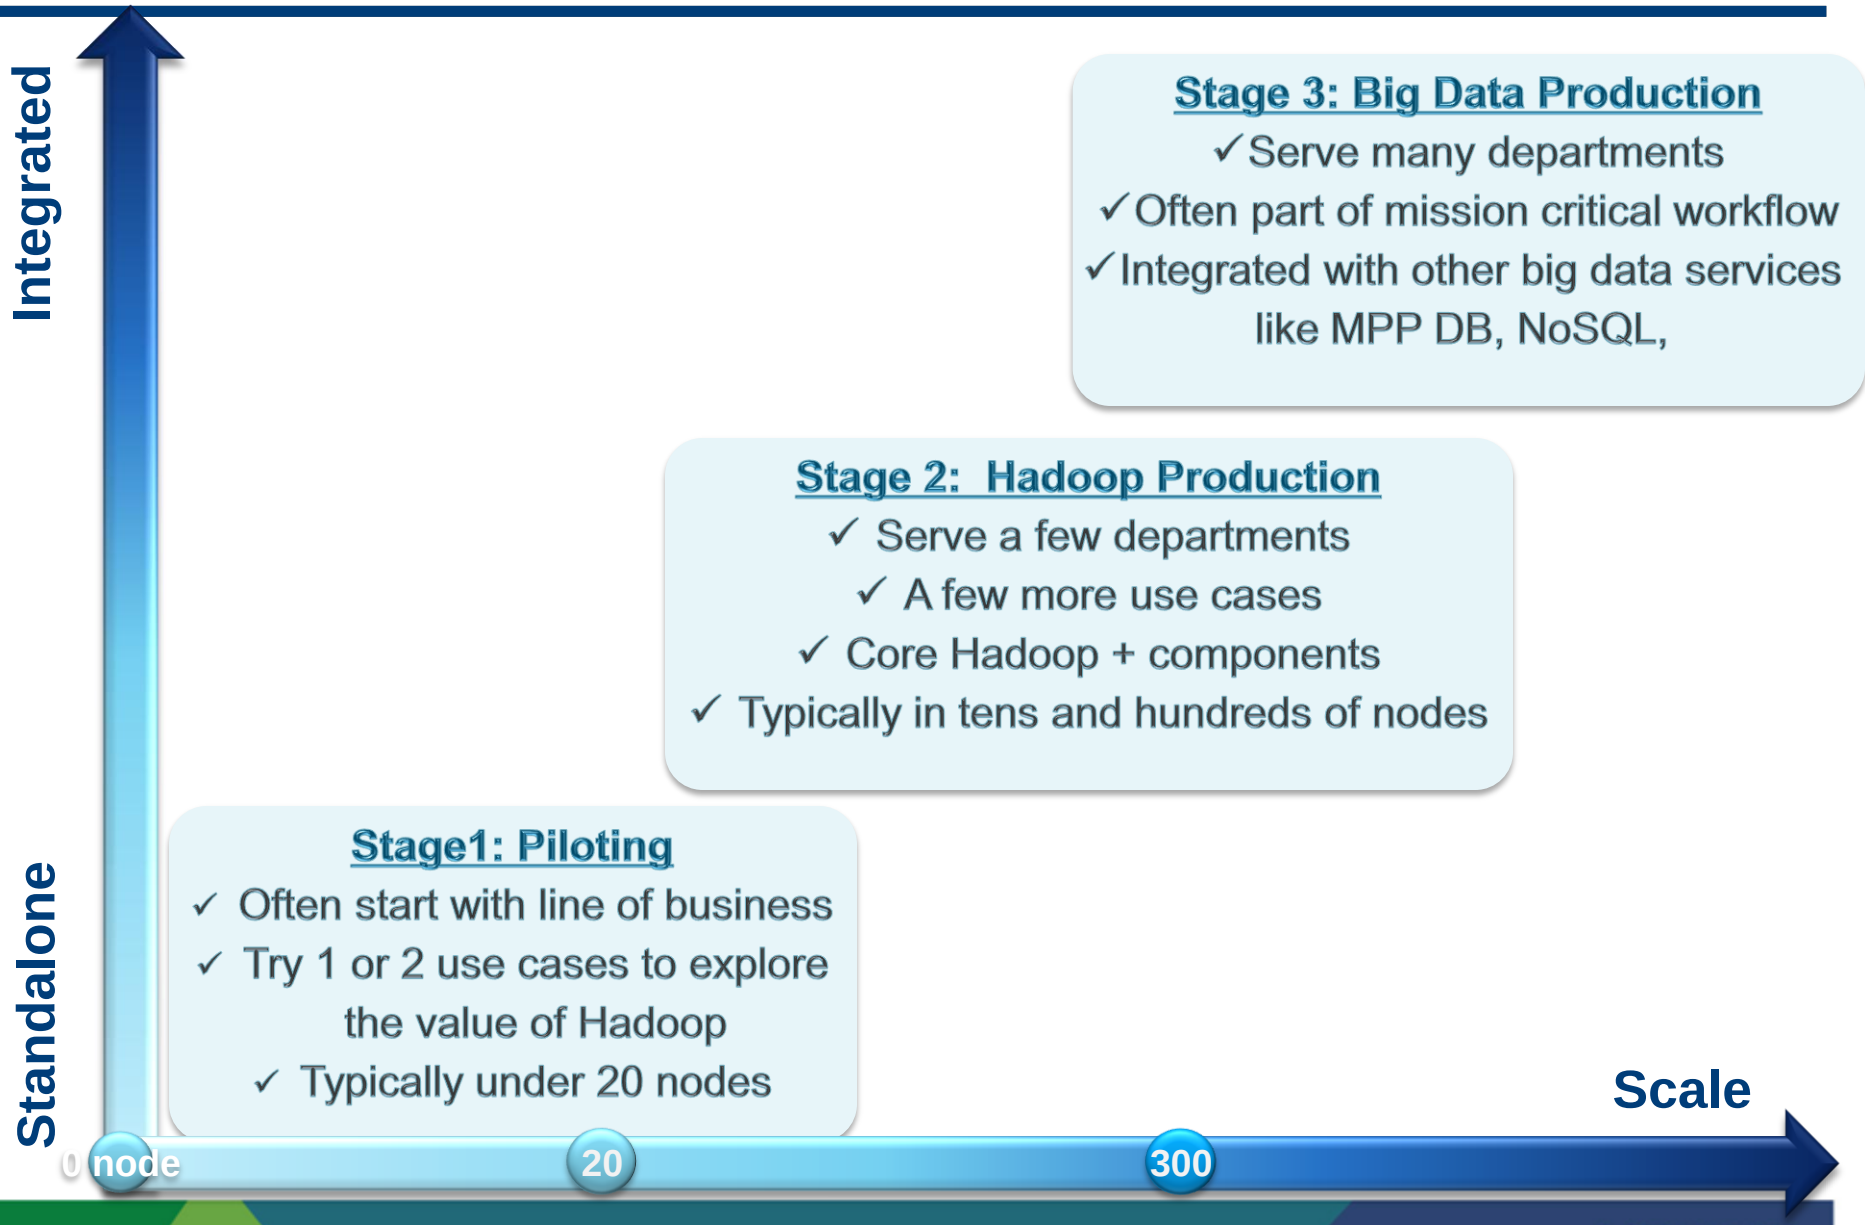

vmware

Stage 1: Piloting

Stage1: Piloting

- ✓ Often start with line of business
- ✓ Try 1 or 2 use cases to explore the value of Hadoop
- ✓ Typically under 20 nodes
- ✓ Led by either data team or infrastructure team

- **Requirements:**

✓ Make it quick

- ✓ I don't want to wait for weeks or months
- ✓ Get me a Hadoop cluster quick

✓ Make it easy

- ✓ Make it easy for me to access the data
- ✓ Make it easy for me to try different algorithms and data sets

Stage 2: Hadoop Production

Stage 2: Hadoop Production

- ✓Serve a few departments
- ✓ A few use cases
- ✓Core Hadoop + some non-core components
- ✓ Dedicated Hadoop administrator

Requirements:

- ✓ High availability
  - ✓ We are in production and need SLA
  - ✓ High availability of the entire Hadoop Stack
- ✓ Agility
  - ✓ We are getting new Hadoop use requests all the time, make it easy for me to scale the cluster
  - ✓ We need to configure and reconfigure the clusters often
- ✓ Differentiated level of services
  - ✓ We have production Hadoop jobs, need to ensure high priority
  - ✓ We also have people trying “ad hoc” Hadoop jobs, need to satisfy their request too

Stage 3: Big Data Production

Stage 3: Big Data Production

- ✓Serve many departments
- ✓Often part of mission critical workflow
- ✓Offer other big data services like MPP DB, NoSQL, more non-core components

Requirements:

- ✓Multi-tenancy
  - ✓We have many tenants on the cluster now, and need ensure resource isolation, configuration isolation between different tenants
- ✓Elasticity
  - ✓With more and more users and jobs on the system, we need to make sure the Hadoop cluster is elastic and adjust to changing demands
- ✓Integrated big data production
  - ✓It's not just about Hadoop anymore, Hadoop is now critical part of overall big data analytics workflow

Agenda

- Hadoop Market Landscape
- Hadoop Journey
- **Virtualize Hadoop Values**
- Summary
- Q & A

VMWare brings Agility, Efficiency, and Elasticity to Big Data

Agility

- Deploy, configure and monitor Hadoop clusters on the fly
- Dynamic reconfiguring of Hadoop to meet changing business demands
- One click HA set up

Efficiency

- Consolidate Hadoop to achieve higher utilization
- Pool resources to allow for increased performance and priority job processing

Elasticity

- Enable full elasticity through separation of Data and Compute
- Scale In/Out Hadoop with Resource Constrain

VMWare brings Agility, Efficiency, and Elasticity to Big Data

Agility

- Deploy, configure and monitor Hadoop clusters on the fly
- Dynamic reconfiguring of Hadoop to meet changing business demands
- One click HA set up

Efficiency

- Consolidate Hadoop to achieve higher utilization
- Pool resources to allow for increased performance and priority job processing

Elasticity

- Enable full elasticity through separation of Data and Compute
- Scale In/Out Hadoop with Resource Constrain

Agility: Automation of Hadoop cluster management

Deploy

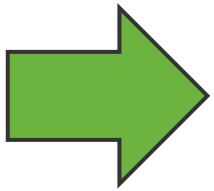

Customize

Incorporate best practices

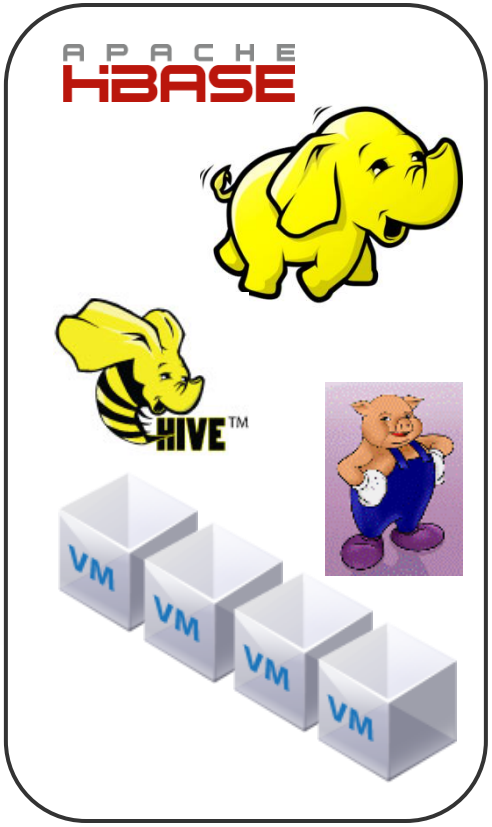

Manage

Tune configuration

Resize

Elastic scaling

Run

Execute jobs

Access HDFS

1/1000 human efforts.
You don't need to be
a Hadoop expert.

SERENGETI
Run Apache Hadoop Virtually

vmware

Rapid Deployment of a Hadoop/HBase Cluster with Serengeti

Step 1: Deploy Serengeti virtual appliance on vSphere.

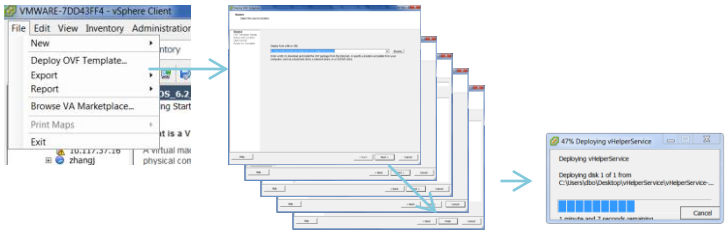

Step 2: A few clicks to stand up Hadoop Cluster.

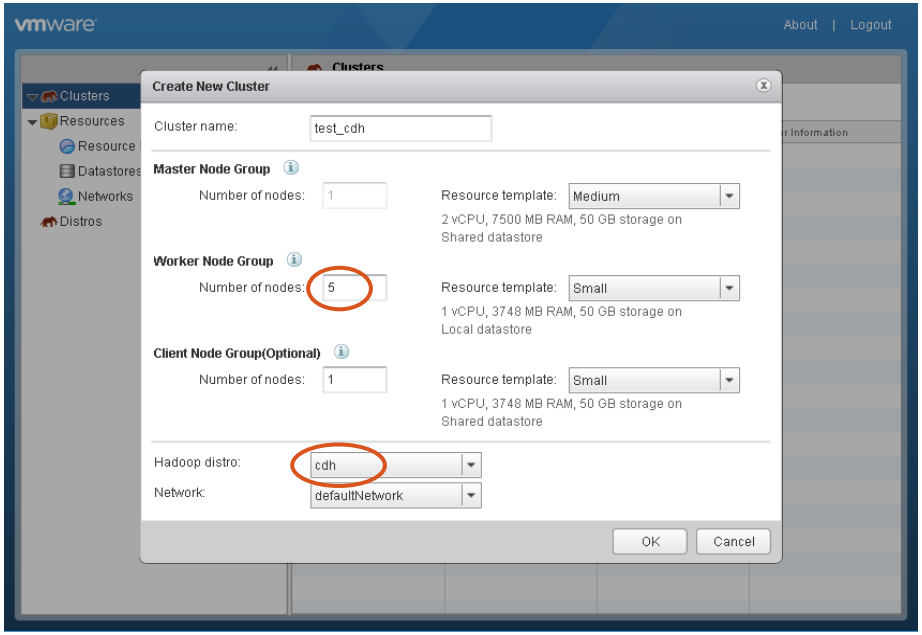

Done

vmware

Customizing your Hadoop/HBase cluster with Serengeti

```
…
 "distro":"apache",
 "groups":[
  { "name":"master",
    "roles":[
      "hadoop_namenode",
      "hadoop_jobtracker”],
"storage": {
       "type": "SHARED",
       "sizeGB": 20},
    "instance_type":MEDIUM,
    "instance_num":1,
    "ha":true},
  {"name":"worker",
    "roles":[
      "hadoop_datanode",
      "hadoop_tasktracker"
    ],
    "instance_type":SMALL,
    "instance_num":5,
    "ha":false
  …
```

- **Choice of distros**
- **Storage configuration**
  - Choice of shared storage or Local disk
- **Resource configuration**
- **High availability option**
- **# of nodes**

One click to scale out your cluster with Serengeti

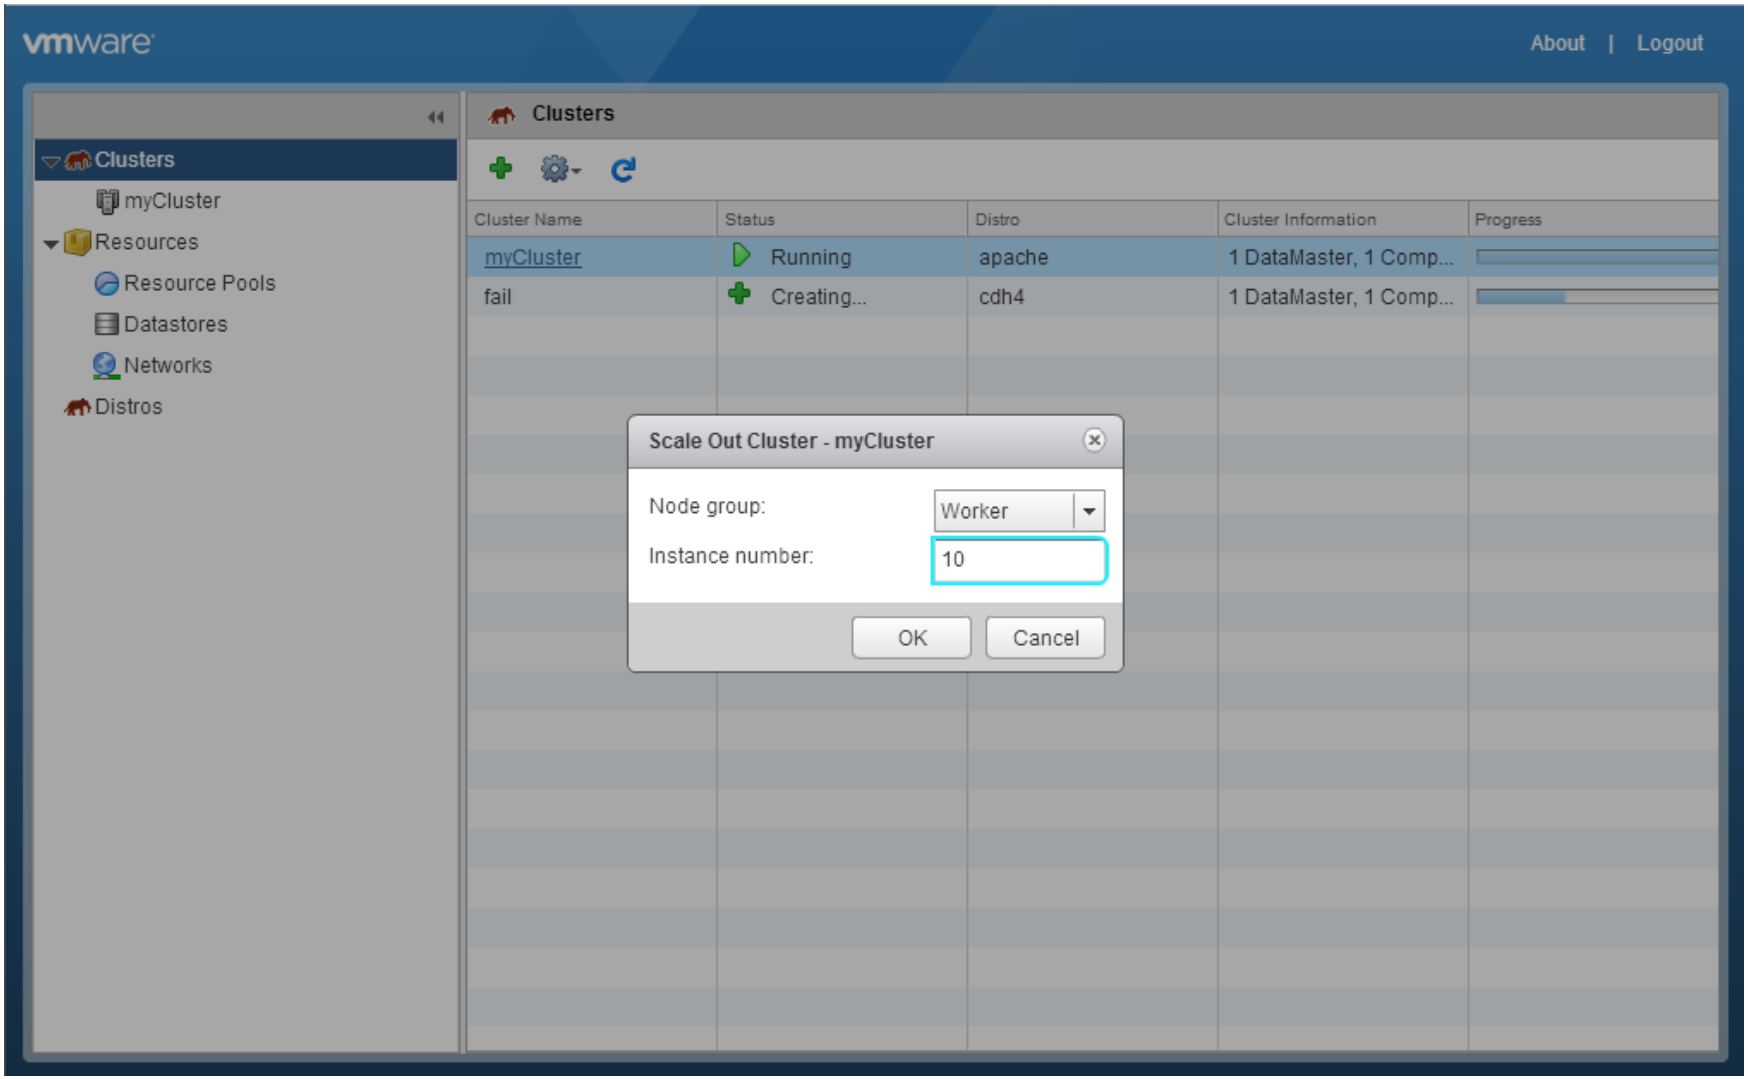

vmware

Configure/reconfigure Hadoop with ease by Serengeti

- **Modify Hadoop cluster configuration from Serengeti**
  - Use the “configuration” section of the json spec file
  - Specify Hadoop attributes in core-site.xml, hdfs-site.xml, mapred-site.xml, hadoop-env.sh, log4j.properties

```
"configuration": {
  "hadoop": {
    "core-site.xml": {
      // check for all settings at http://hadoop.apache.org/common/docs/r1.0.0/core-default.html
    },
    "hdfs-site.xml": {
      // check for all settings at http://hadoop.apache.org/common/docs/r1.0.0/hdfs-default.html
    },
    "mapred-site.xml": {
      // check for all settings at http://hadoop.apache.org/common/docs/r1.0.0/mapred-default.html
      "io.sort.mb": "300"
    } ,
    "hadoop-env.sh": {
      // "HADOOP_HEAPSIZE": "",
      // "HADOOP_NAMENODE_OPTS": "",
      // "HADOOP_DATANODE_OPTS": "",
  …
```

  - Apply new Hadoop configuration using the edited spec file

```
> cluster config --name myHadoop --specFile /home/serengeti/myHadoop.json
```

Proactive monitoring and tuning with VCOPs

- Proactively monitoring through VCOPs
- Gain comprehensive visibility
- Eliminate manual processes with intelligent automation
- Proactively manage operations

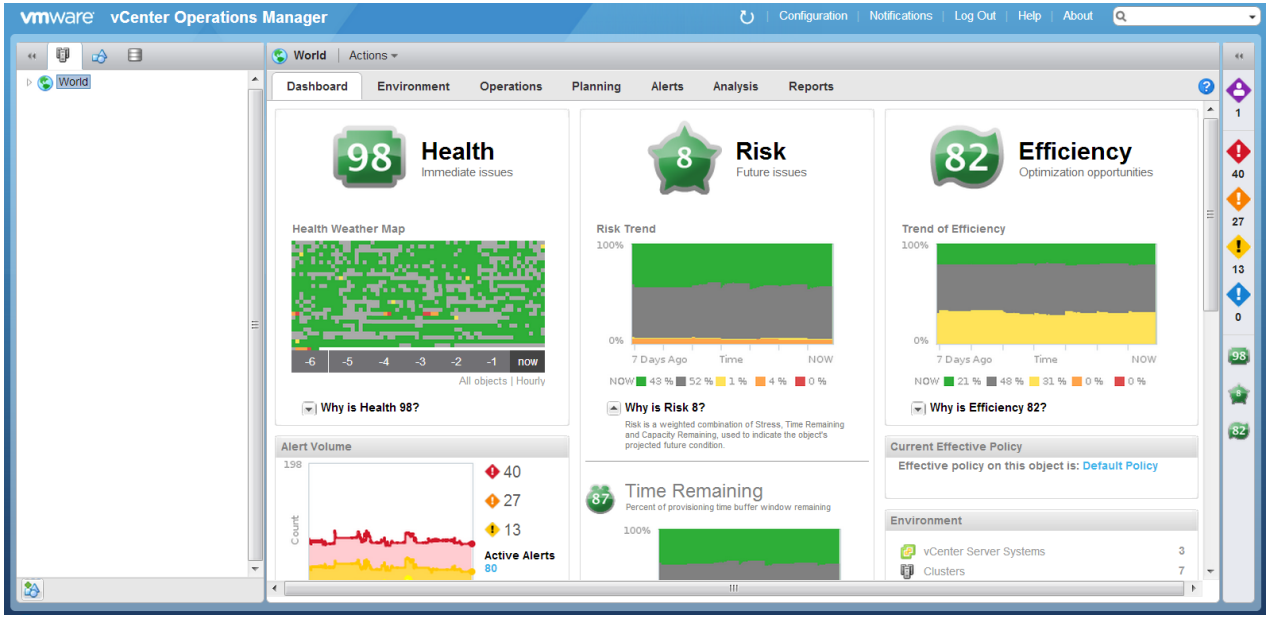

High availability for the Hadoop stack

- Increase availability of whole Hadoop stack
- Battle tested solution
- Use vMotion to eliminate planned downtime
- Use vSphere HA to decrease unplanned downtime with automatic fail over
- Use vSphere FT to provide zero downtime zero data lost protection

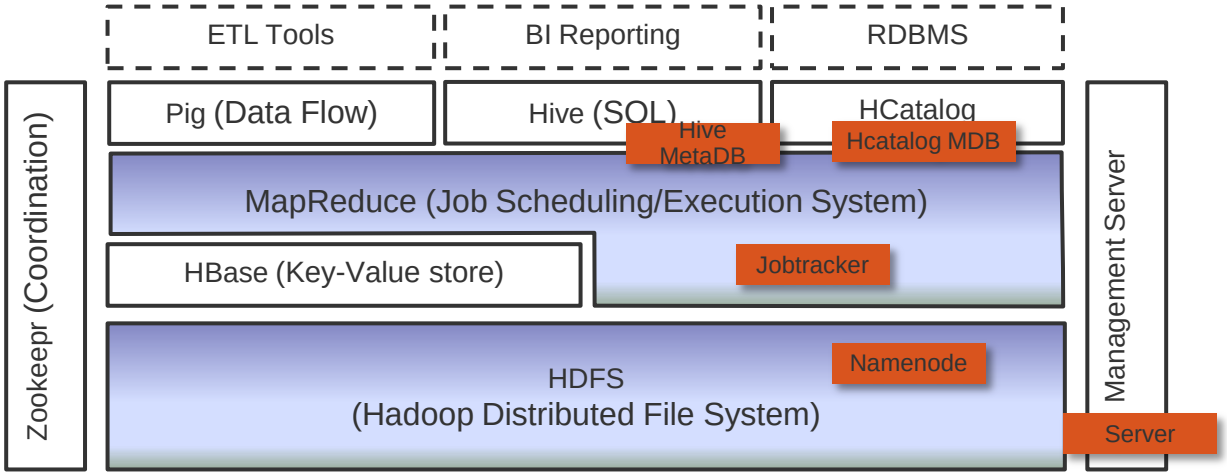

Support all popular Hadoop distributions and tools

Distributions

hadoop

cloudera

Greenplum

Hortonworks

MAPR TECHNOLOGIES

Community

APACHE HBASE

- Flexibility to choose and try out major distributions
- Support for multiple projects
- Open architecture to welcome industry participation
- Contributing Hadoop Virtualization Extensions (HVE) to open source community

Virtualization is much more agile than physical

	Dedicated Physical Clusters	Virtual Clusters
Cluster Construction	•Server procurement •Data Center considerations •All kinds of manual steps	•Centralized IT management •No case by case consideration •Fully end to end automation
Cluster Operation	•Need immediate reaction when failure happens	•Higher tolerance to failure with large distributed resource pool •Automatic fail over
Capacity Planning	•Plan for future, requires unutilized capacity	•Plan for now, get only used capacity.
Enlarge Capacity	Requires server procurement and setup when compute or storage capacity is not enough	Carve from large sharing pool, apply and get

VMWare brings Agility, Efficiency, and Elasticity to Big Data

Agility

- Deploy, configure and monitor Hadoop clusters on the fly
- Dynamic reconfiguring of Hadoop to meet changing business demands
- One click HA set up

Efficiency

- Consolidate Hadoop to achieve higher utilization
- Pool resources to allow for increased performance and priority job processing

Elasticity

- Enable full elasticity through separation of Data and Compute
- Scale In/Out Hadoop with Resource Constrain

Customer Example: Enterprise Adoption of Hadoop

Dept A: recommendation engine

Dept B: ad targeting

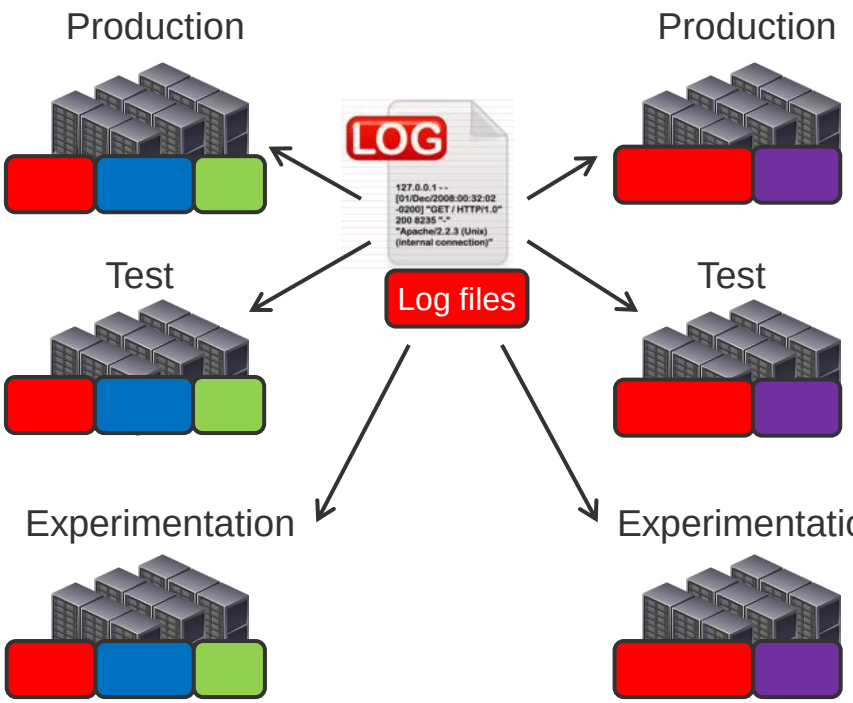

SLA: Jobs complete in 15 minutes

Bandwidth limited to 30 nodes at peak

Issues:

1. Multiple clusters to manage
2. Redundant common data in separate clusters
3. Peak compute and I/O resource is limited to number of nodes in each independent cluster

Transaction data

Social data

Historical cust behavior

vmware

Consolidate Big Data clusters into a unified virtual infrastructure

- **Multiple Big Data clusters co-exist in hosts**
  - Share hardware resource to gain high utilization
  - Data co-exist avoids cross network movements
  - Single infrastructure to maintain
- **vSphere ensures strong isolation between clusters.**
  - Resource isolation.
  - Failure isolation
  - Configure isolation
  - Security isolation

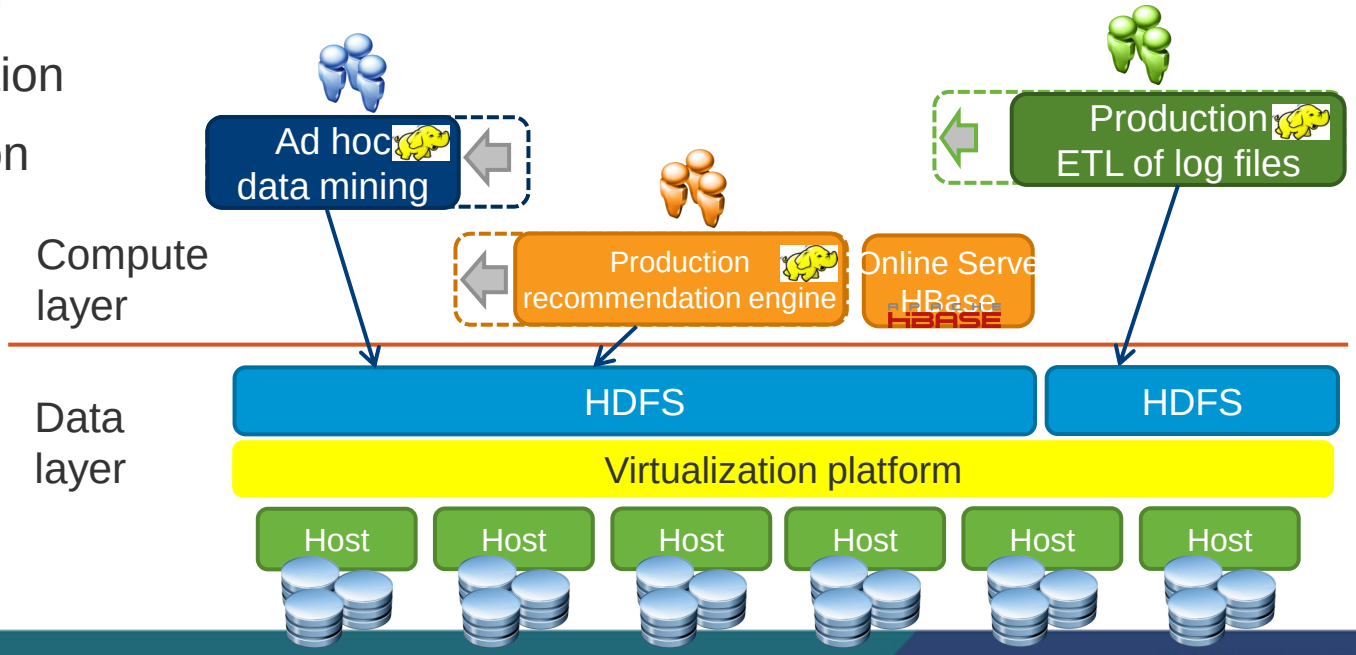

Lower CAPEX with sharing compute resource with consolidation

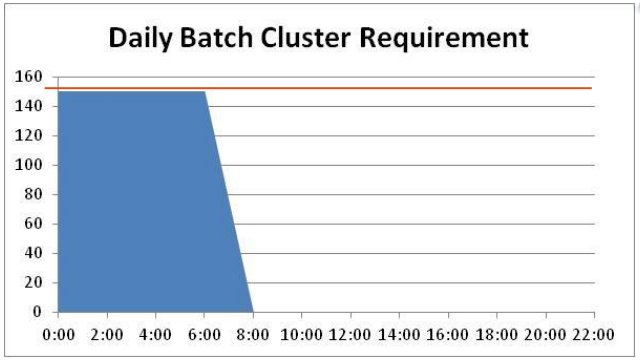

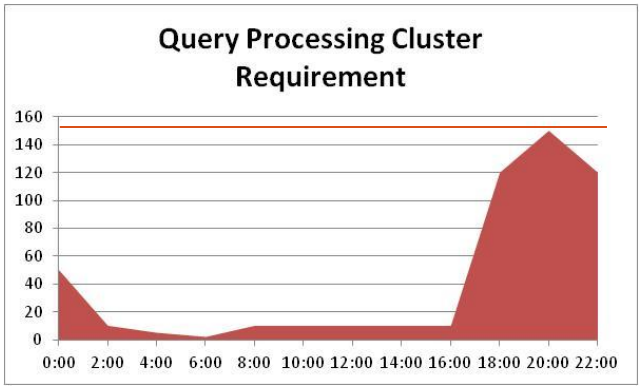

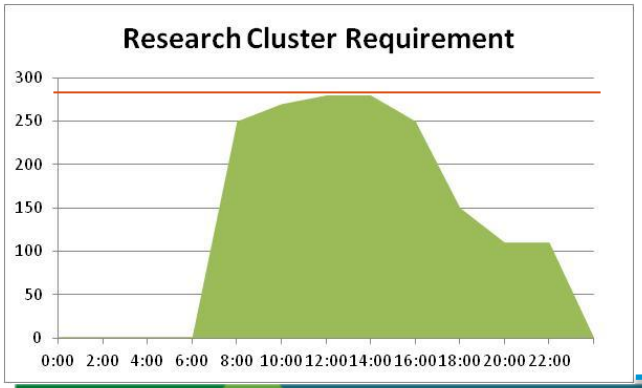

Σ(Max) → Max(Σ)

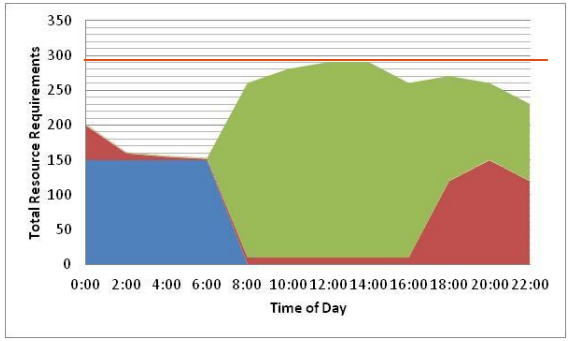

- **Without virtualization, CAPEX is calculated by the sum of the maximum workload of each cluster.**
- **Virtualized**
  - Clusters share a big pool of resource
  - CAPEX is calculated by maximum of overall workload
  - 2~4 to 1 consolidation rate

vmware

Lower CAPEX with sharing of common data

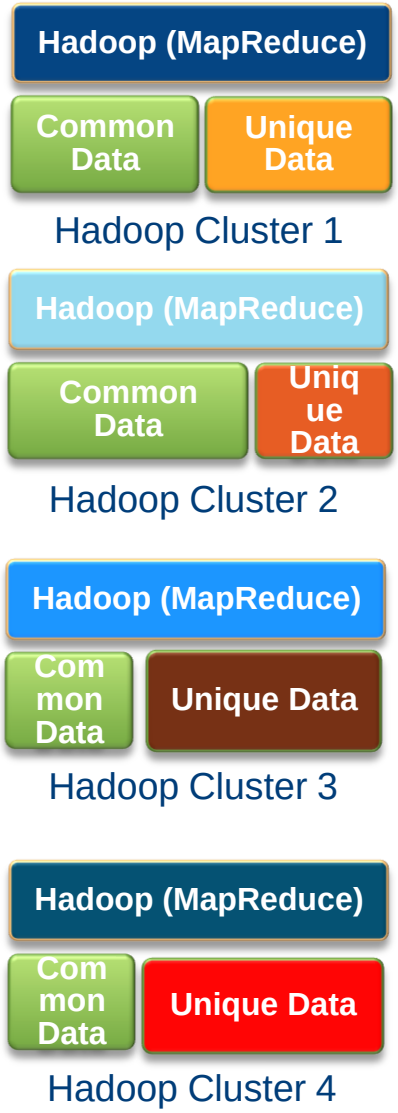

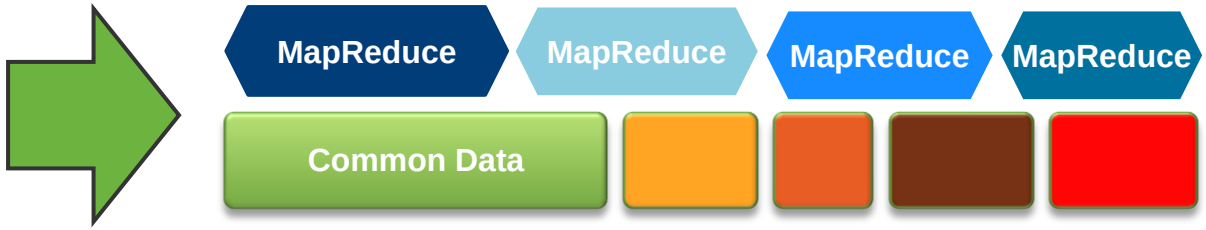

Without Virtualization, multiple copies of common data in separate Hadoop clusters

Virtualized

- Single HDFS to serve multiple MR clusters without losing data locality
- Single copy of common data results in less storage required and lower hardware requirement
- 3 to 2 consolidation rate

Utilize all your resources to solve the priority problem

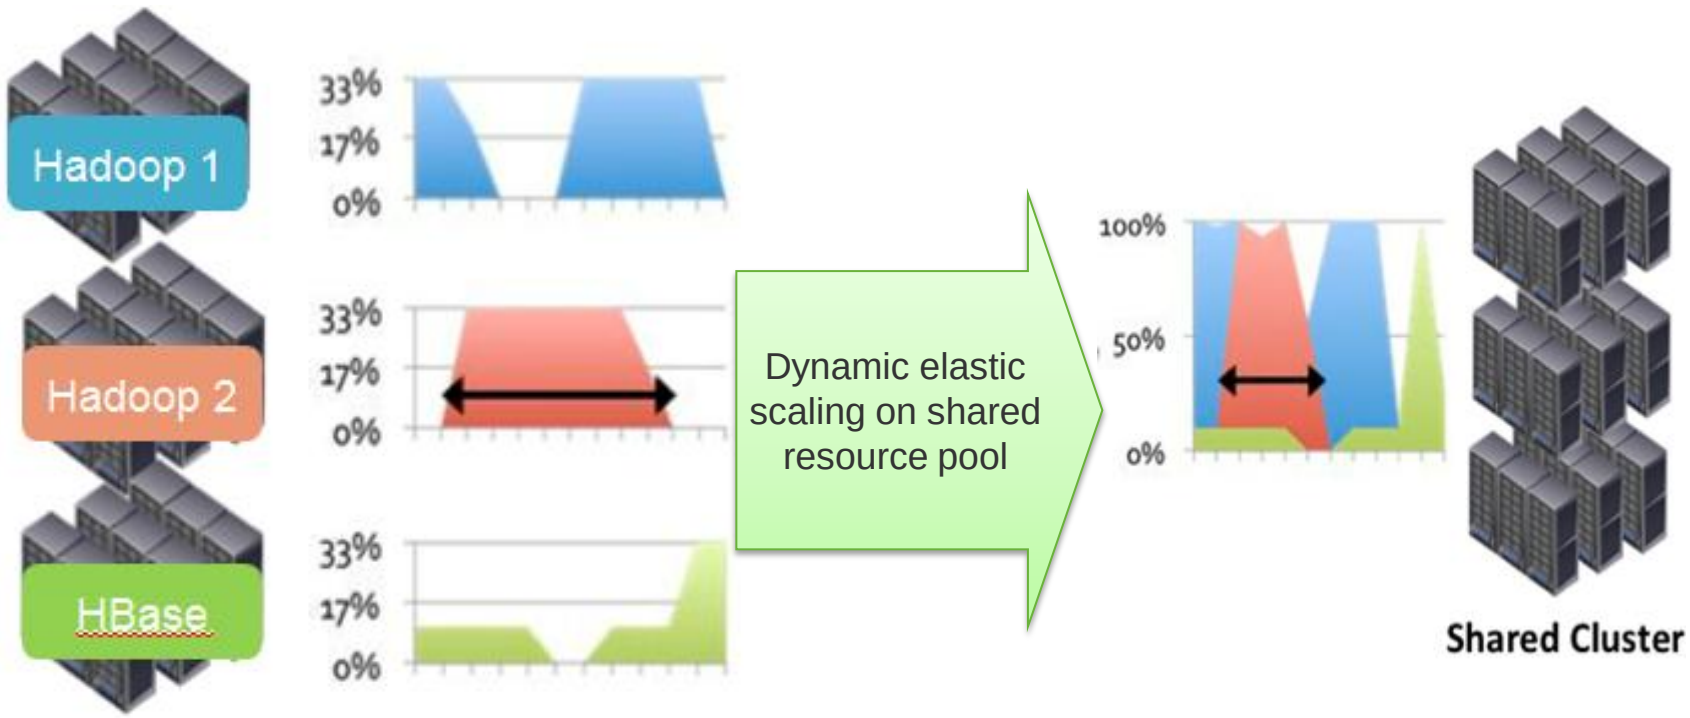

50%+ resources are sitting idle while high priority job is burning up its cluster.

Utilize all resources from pool on demand.

There're ways to consolidate, but virtualization is the best

	Physical	Virtual
Resource Sharing	Yes, Users share a common Hadoop cluster	Yes, Users share common physical servers in different Hadoop clusters
Data Sharing	Yes, Users share a common Hadoop cluster	Yes, Different compute clusters share a common HDFS cluster
Performance Isolation	Weak, by slot number	Strong, by CPU, RAM, Disk IO
Failure Isolation	No, Bad job fails entire cluster	Strong, Failure impact only one cluster
Configuration Isolation	No, Same configuration, same distro, same version	Yes, Free to use different distro, version, configuration
Security Isolation	Weak, Enforced by Hadoop authentication and authorization	Strong, Cluster level isolation.
Scalability	Single master node capacity will become a bottle neck	As many Namenode and Jobtracker as needed

VMWare brings Agility, Efficiency, and Elasticity to Big Data

Agility

- Deploy, configure and monitor Hadoop clusters on the fly
- Dynamic reconfiguring of Hadoop to meet changing business demands

Efficiency

- Consolidate Hadoop to achieve higher utilization
- Pool resources to allow for increased performance and priority job processing

Elasticity

- Enable full elasticity through separation of Data and Compute
- Scale In/Out Hadoop with Resource Constrain

Business requires elasticity but Hadoop don’t Work That way

- **Business requires elasticity**
  - It’s never easy to predict future workload. How to change quickly?
  - Every workload shows peak and valley. How to avoid waste in off-peak time?
  - You always want to use everything for urgent jobs. How to boost up?
- **Hadoop is not as elastic as expected**
  - You can add more servers to scale out, but data don’t allow you to scale in.
  - Scale in/out requires careful configuration
  - Scale in/out requires huge data movement

Virtualization eliminates the elastic boundary of Hadoop

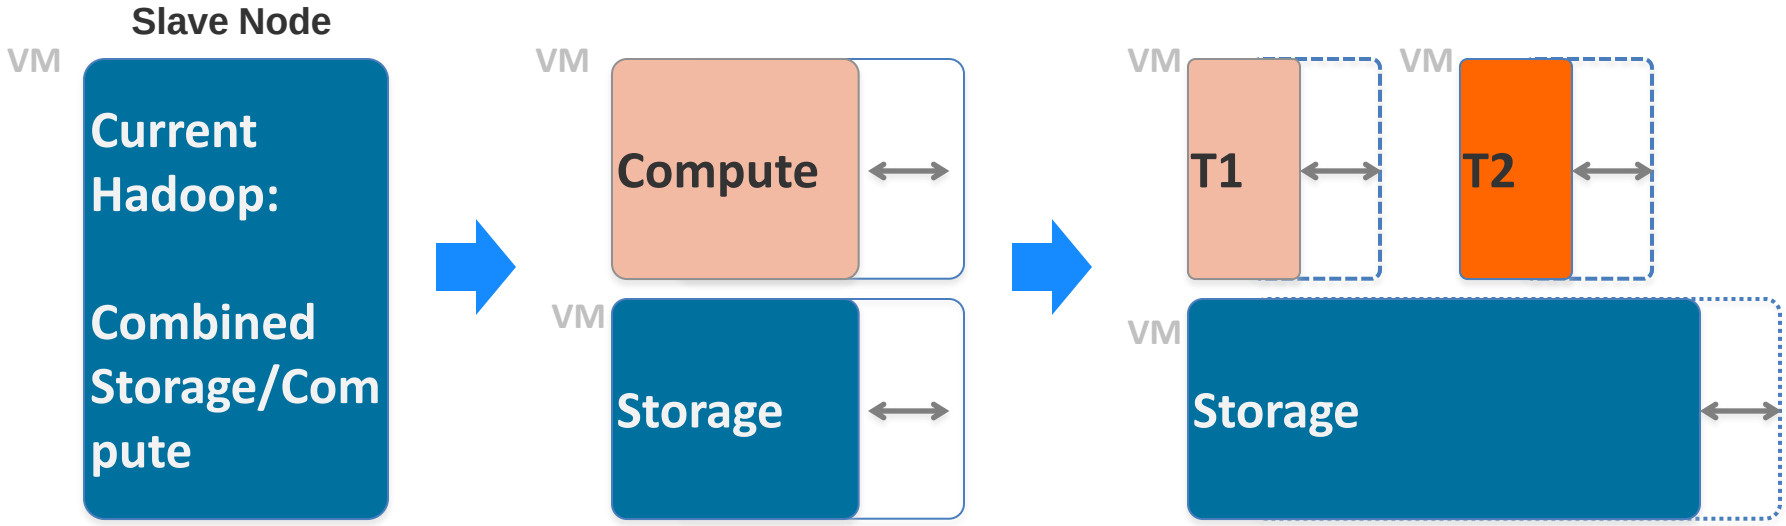

Hadoop in VM

- * VM lifecycle determined by Datanode
- * Limited elasticity

Separate Storage

- * Separate compute from data
- * Remove elastic constrain by Datanode
- * Elastic compute
- * Raise utilization

Separate Compute Clusters

- * Separate virtual compute
- * Compute cluster per tenant
- * Stronger VM-grade security and resource isolation

Locality of data-compute separation in virtualization

Combined mode

1 Combined Compute/Datanode VM per Host

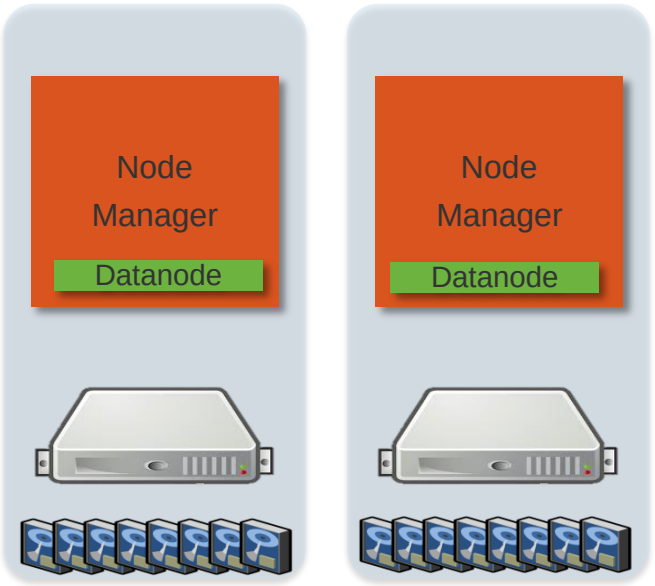

Split Mode

1 Datanode VM, 1 Compute node VM per Host

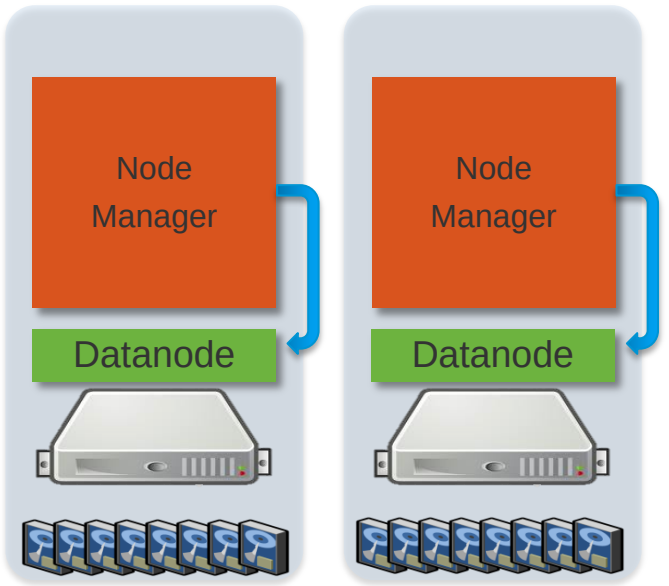

Workload: Teragen, Terasort, Teravalidate
HW Configuration: 8 cores, 96GB RAM, 16 disks per host x 2 nodes

Performance analysis of separation

Minimum performance impact with separation of compute and data

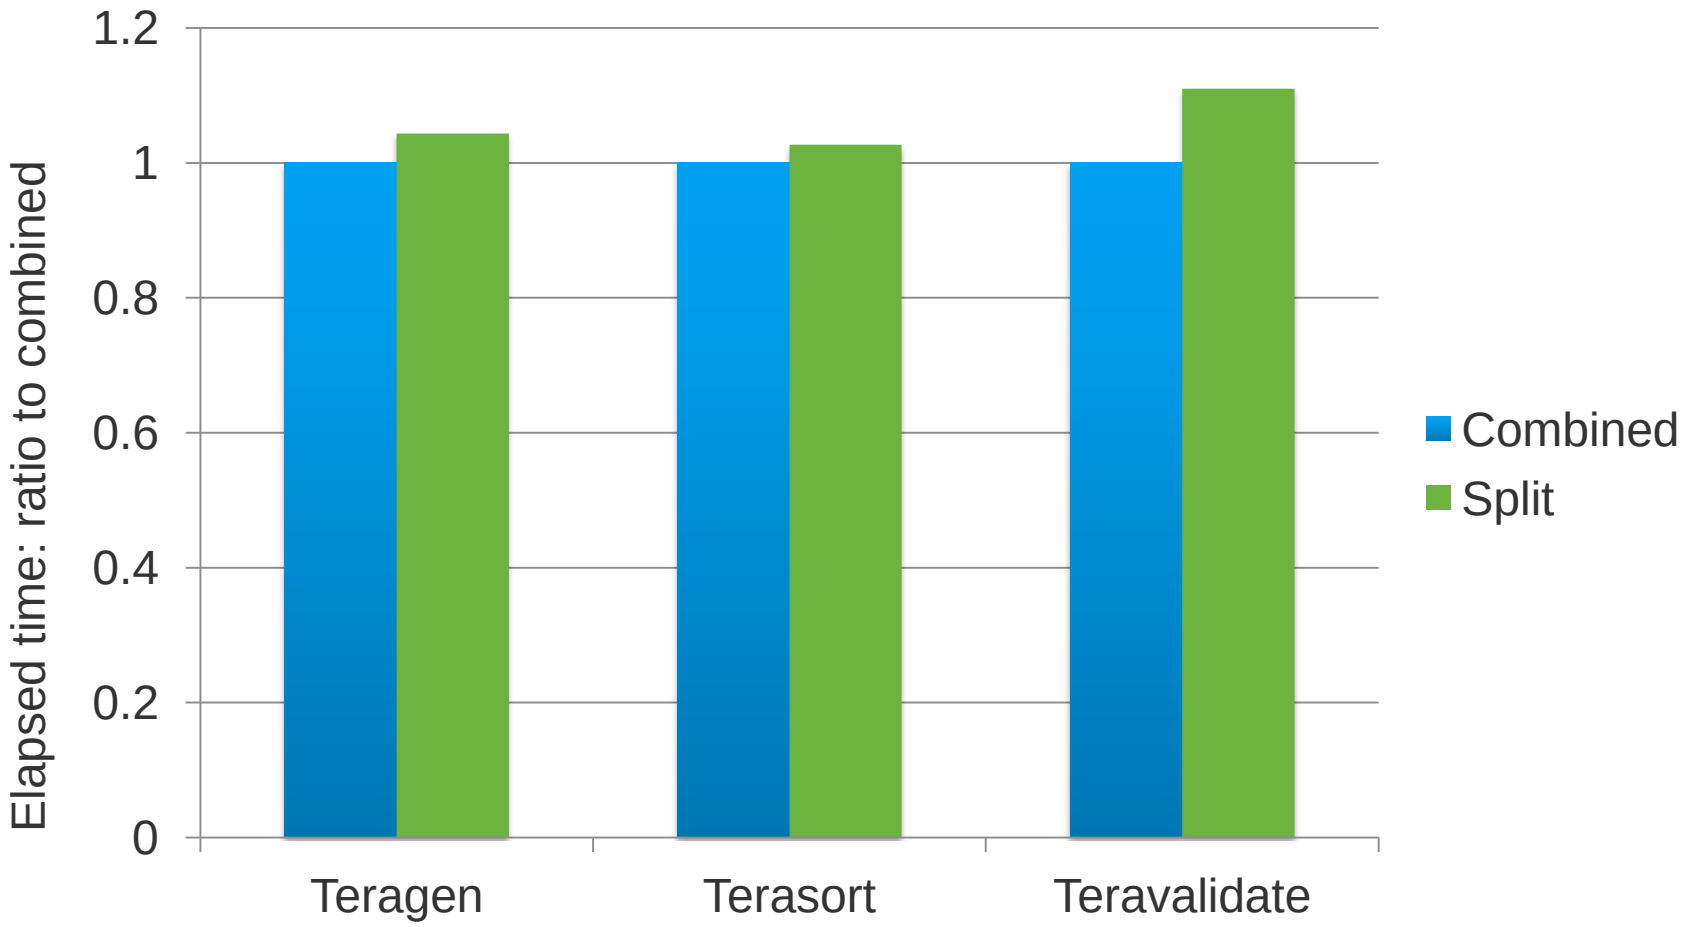

Scale in/out Hadoop dynamically

- **Deploy separate compute clusters for different tenants sharing HDFS.**
- **Commission/decommission task trackers according to priority and available resources**

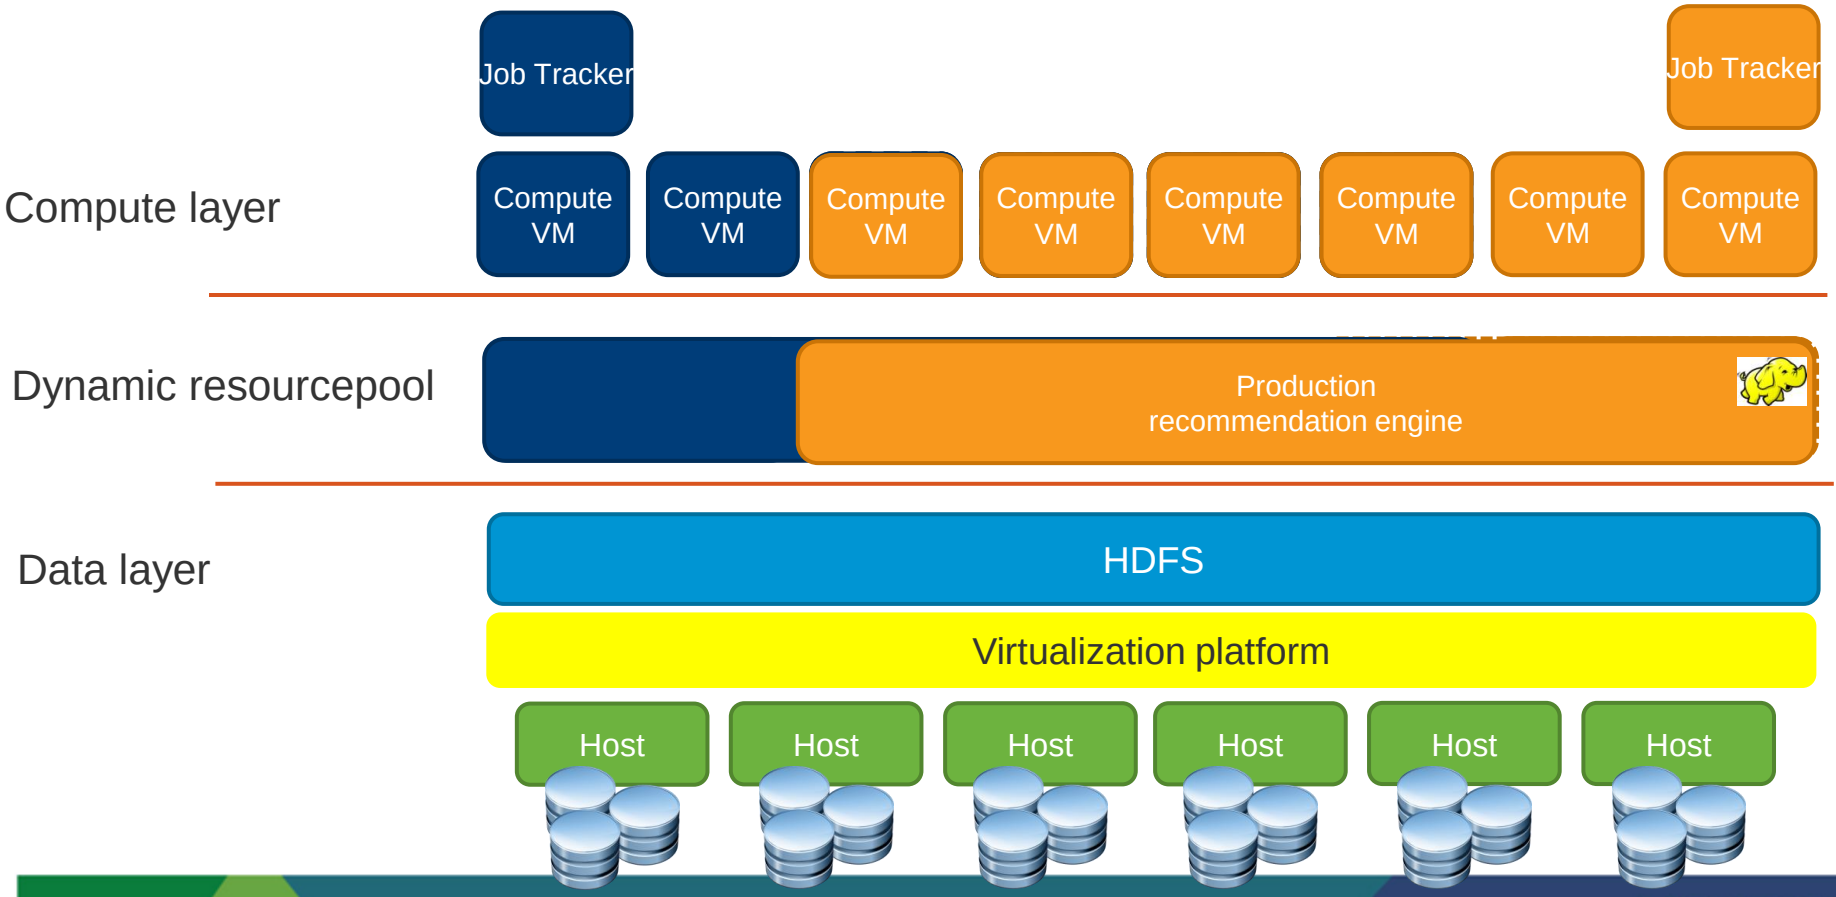

Control resource consumption to satisfy SLA

- >cluster limit --name <clustername> --activeComputeNodeNum <#>
- 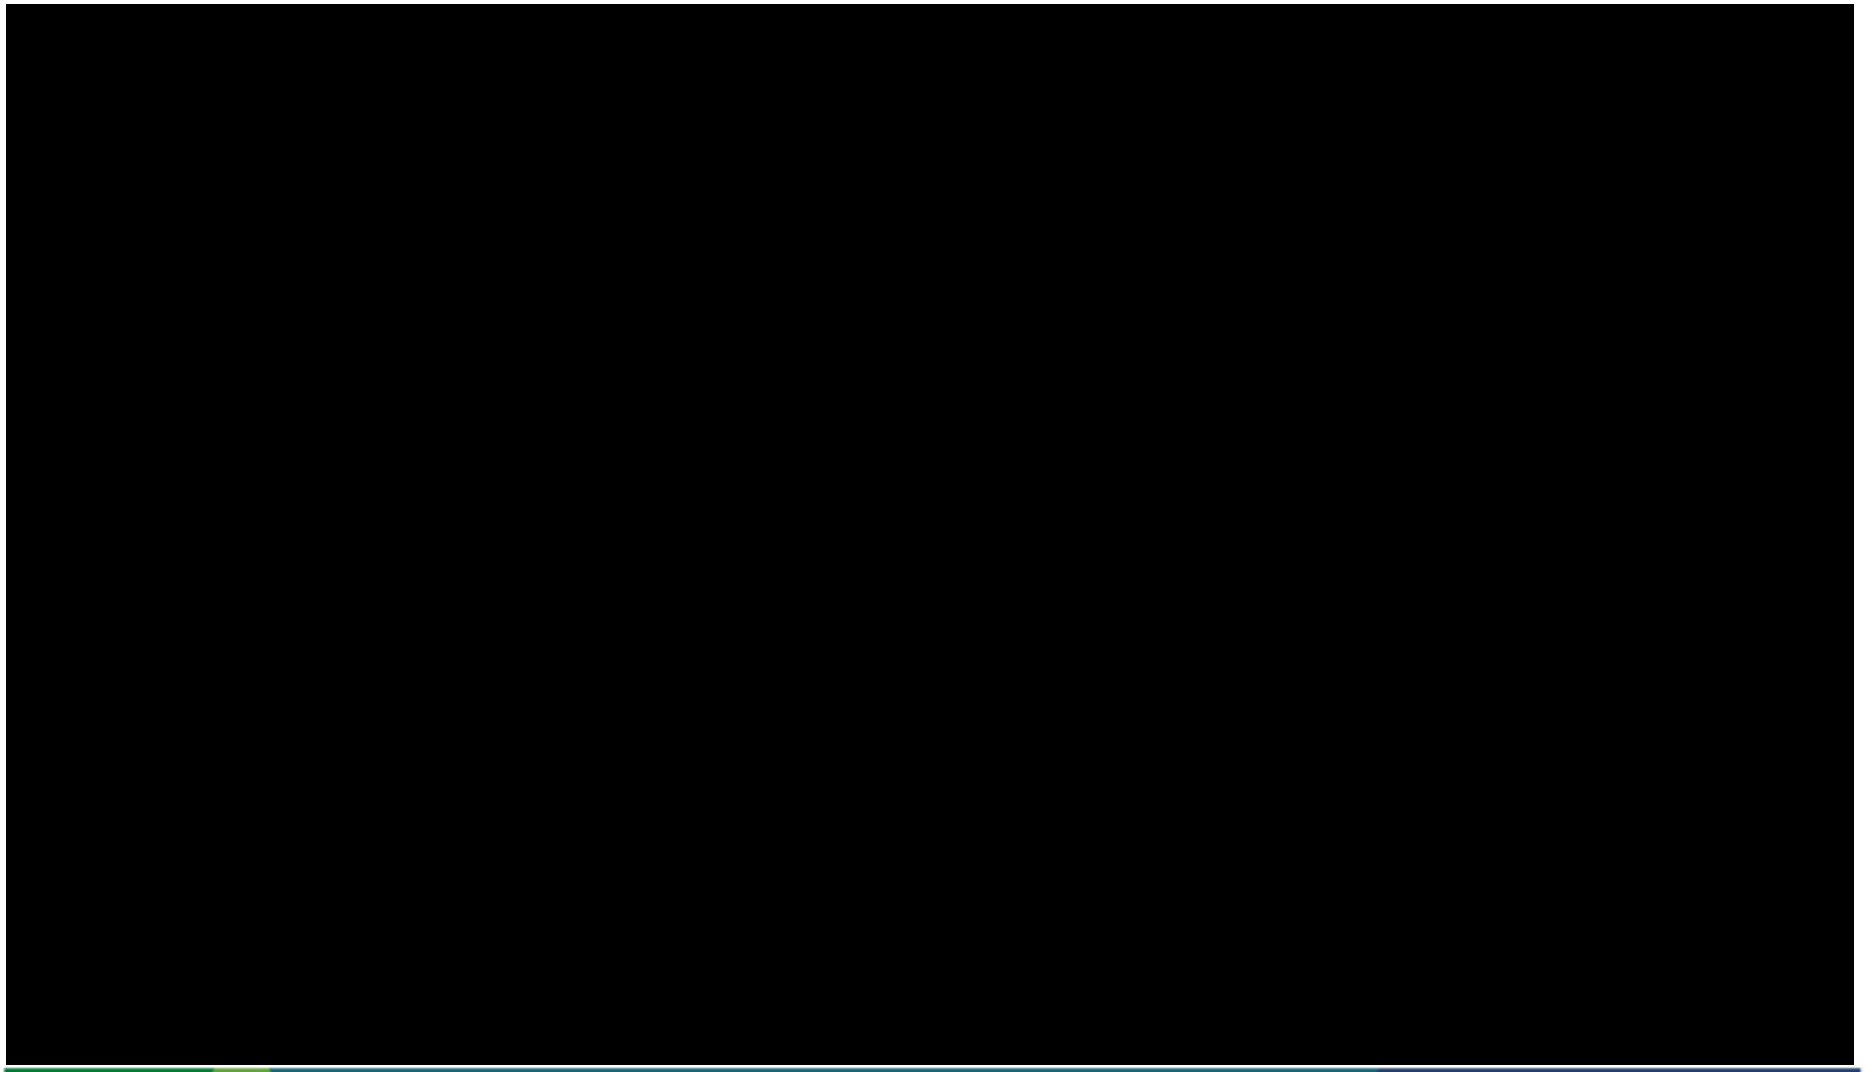

Automatic and adaptive elastic

- Control with knowledge of both physical and virtual
- Control instantly and gracefully

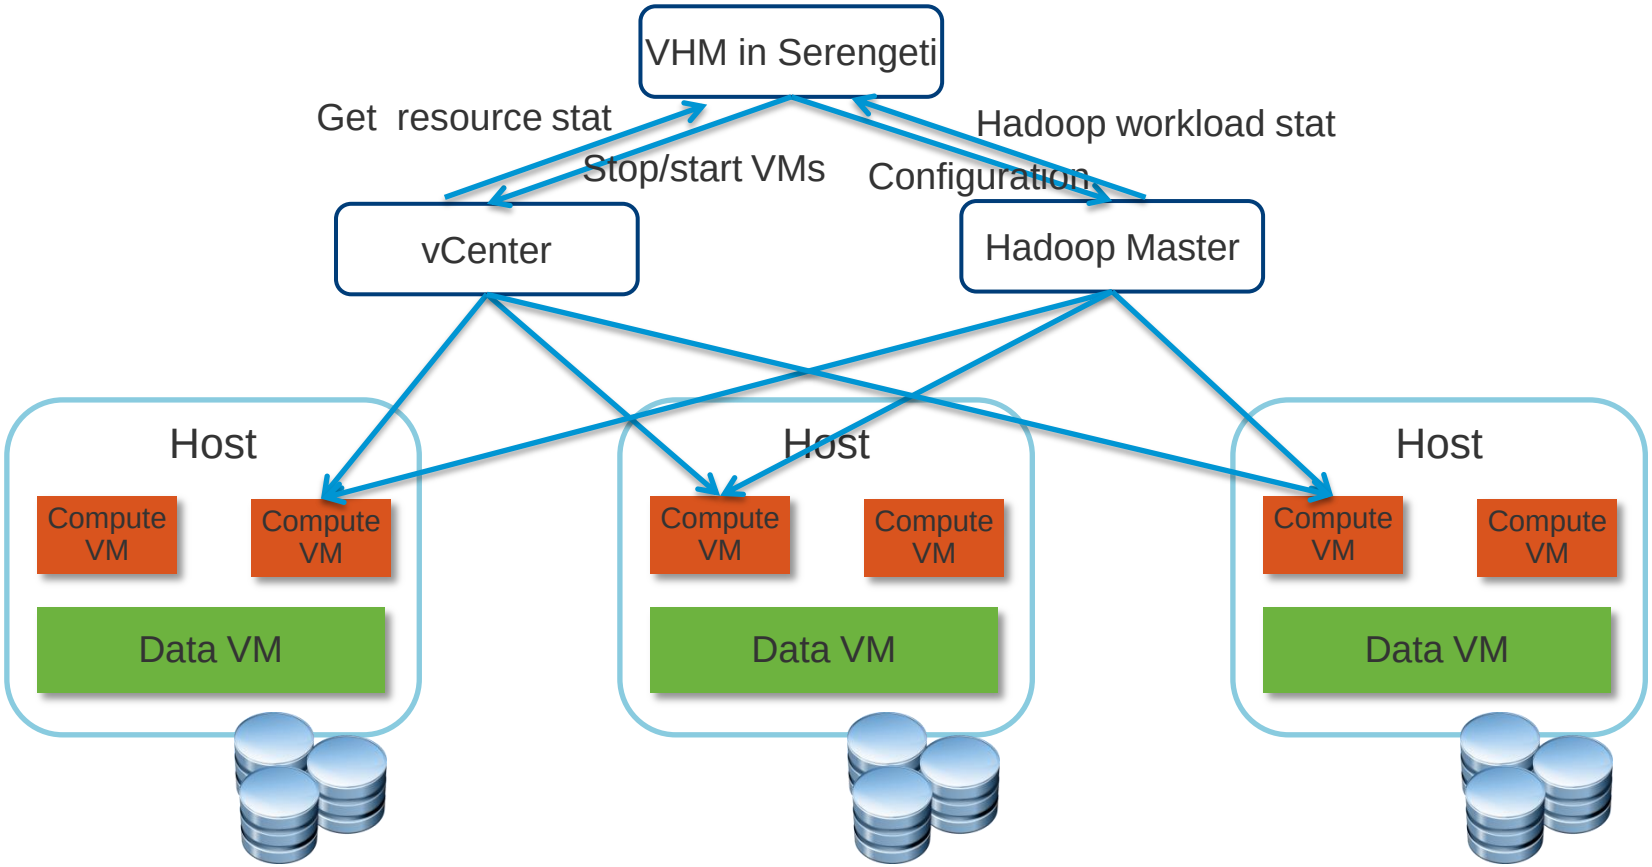

In-house Hadoop as a Service – (Hadoop + Hadoop)

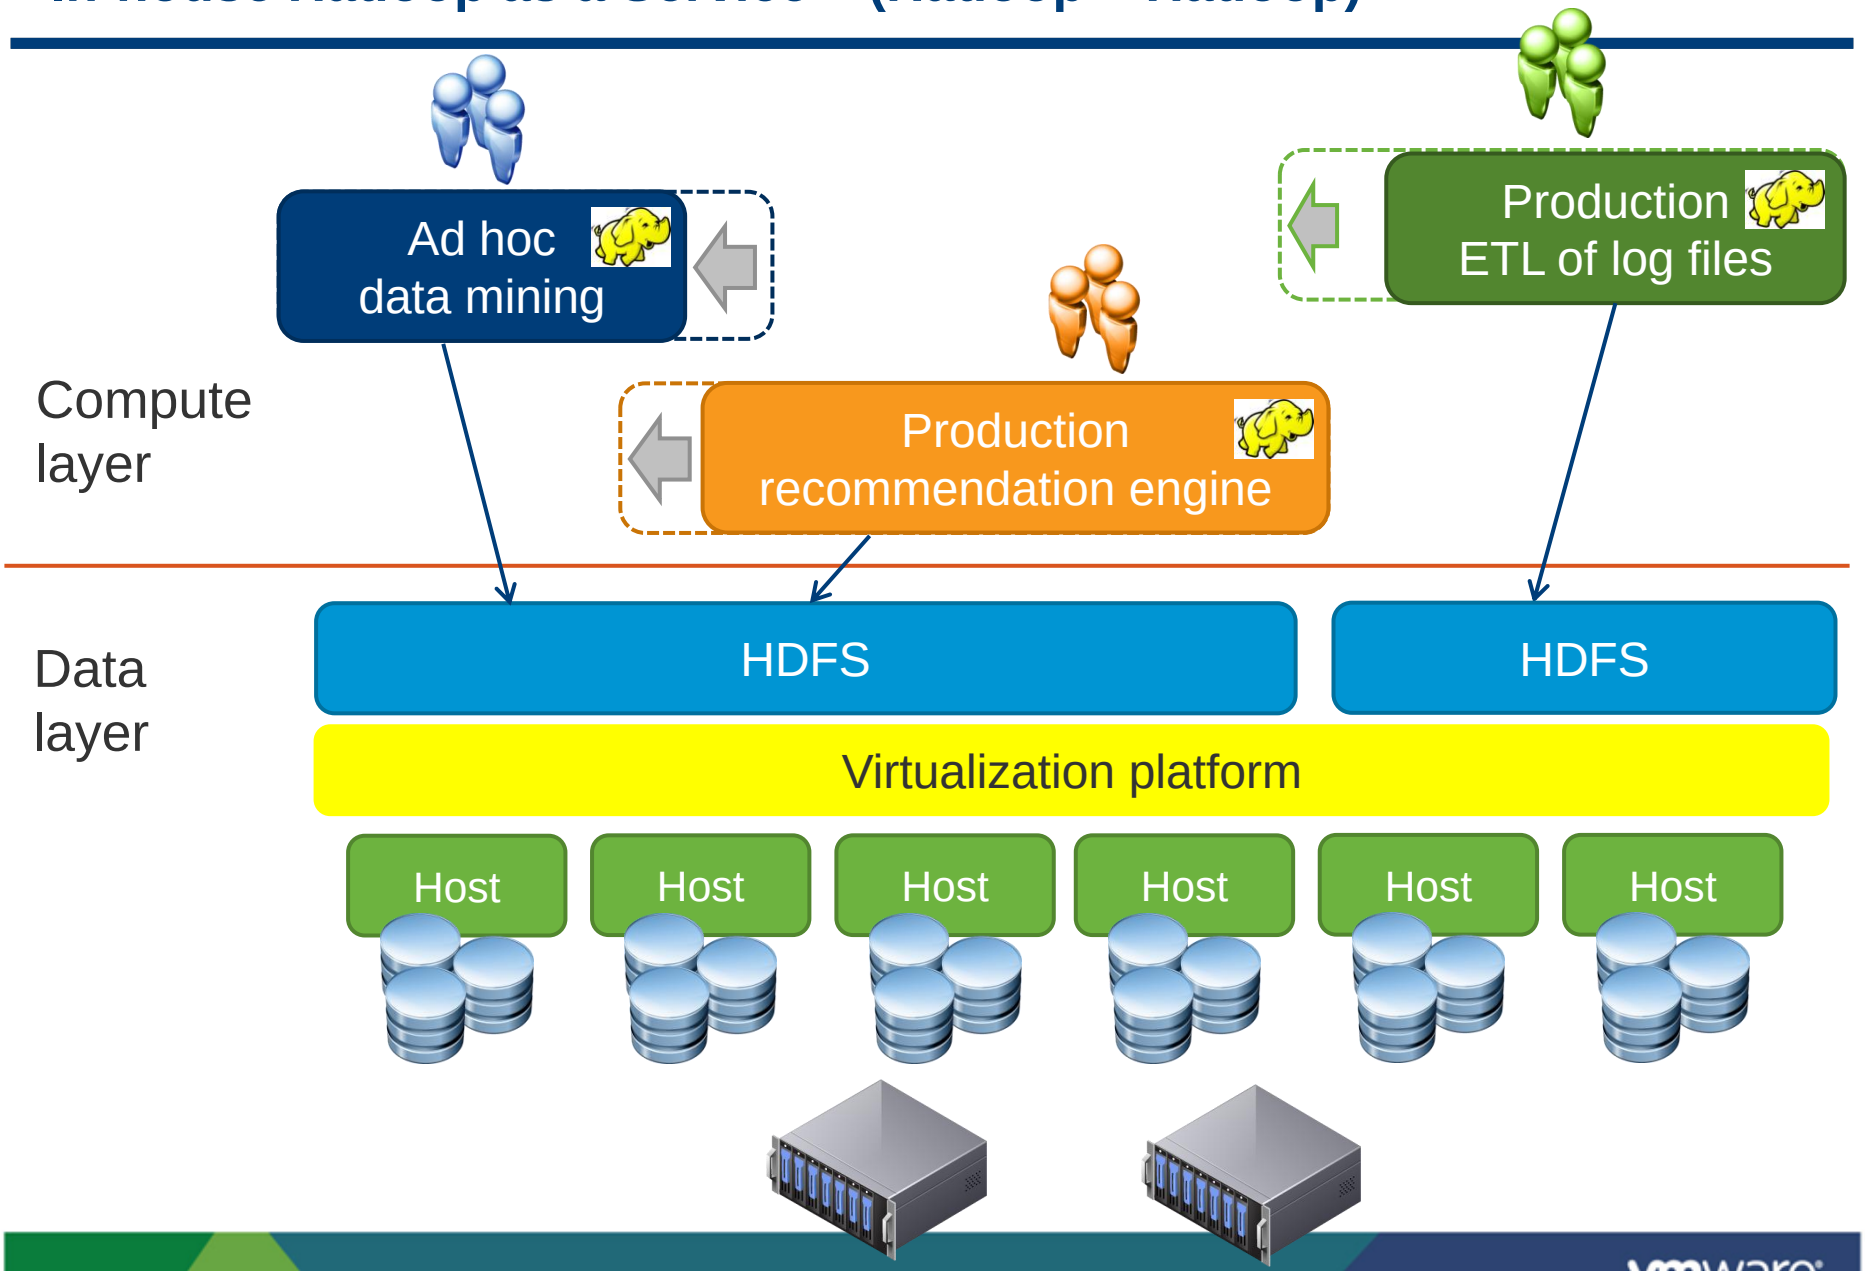

Integrated Hadoop and Webapps – (Hadoop + Other Workloads)

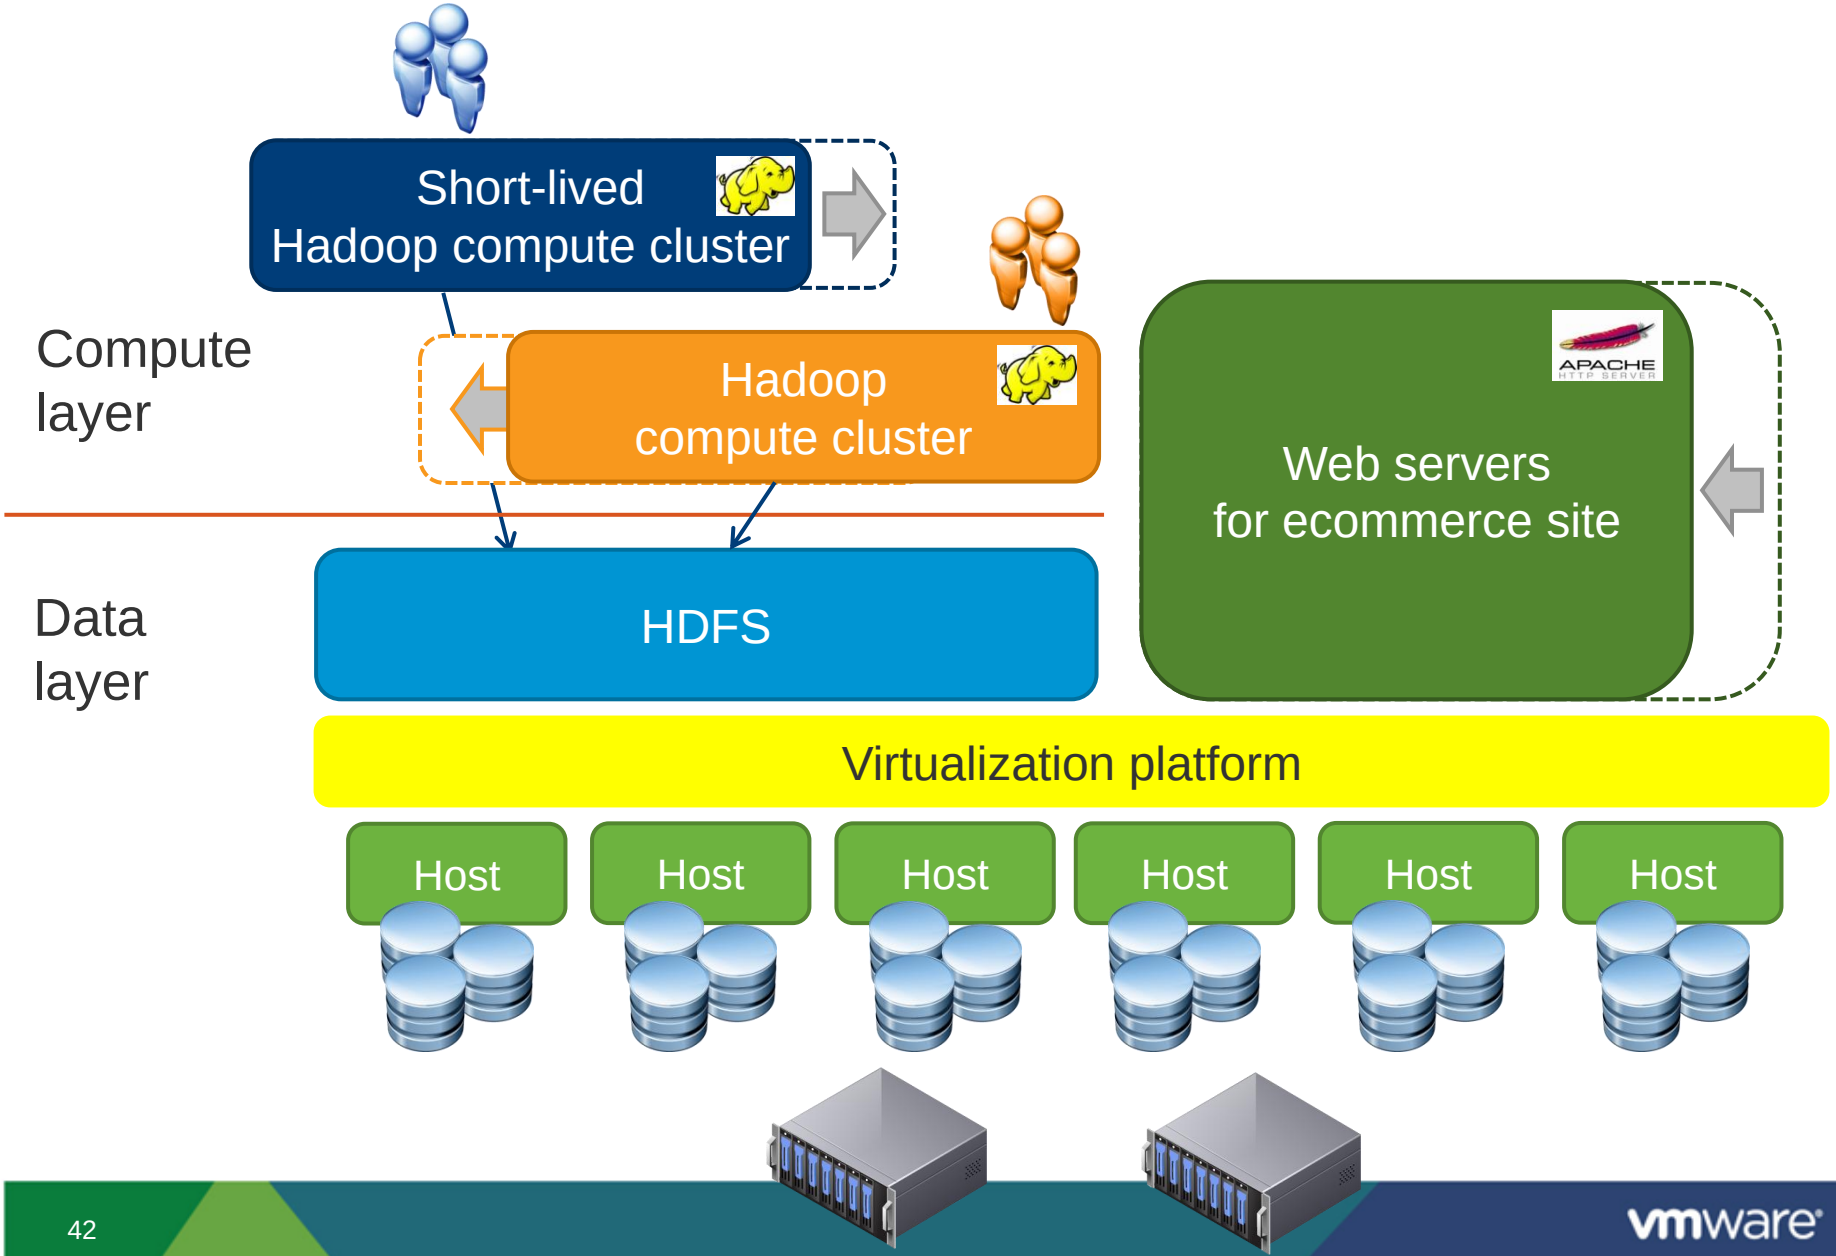

Integrated Big Data Production – (Hadoop + other big data)

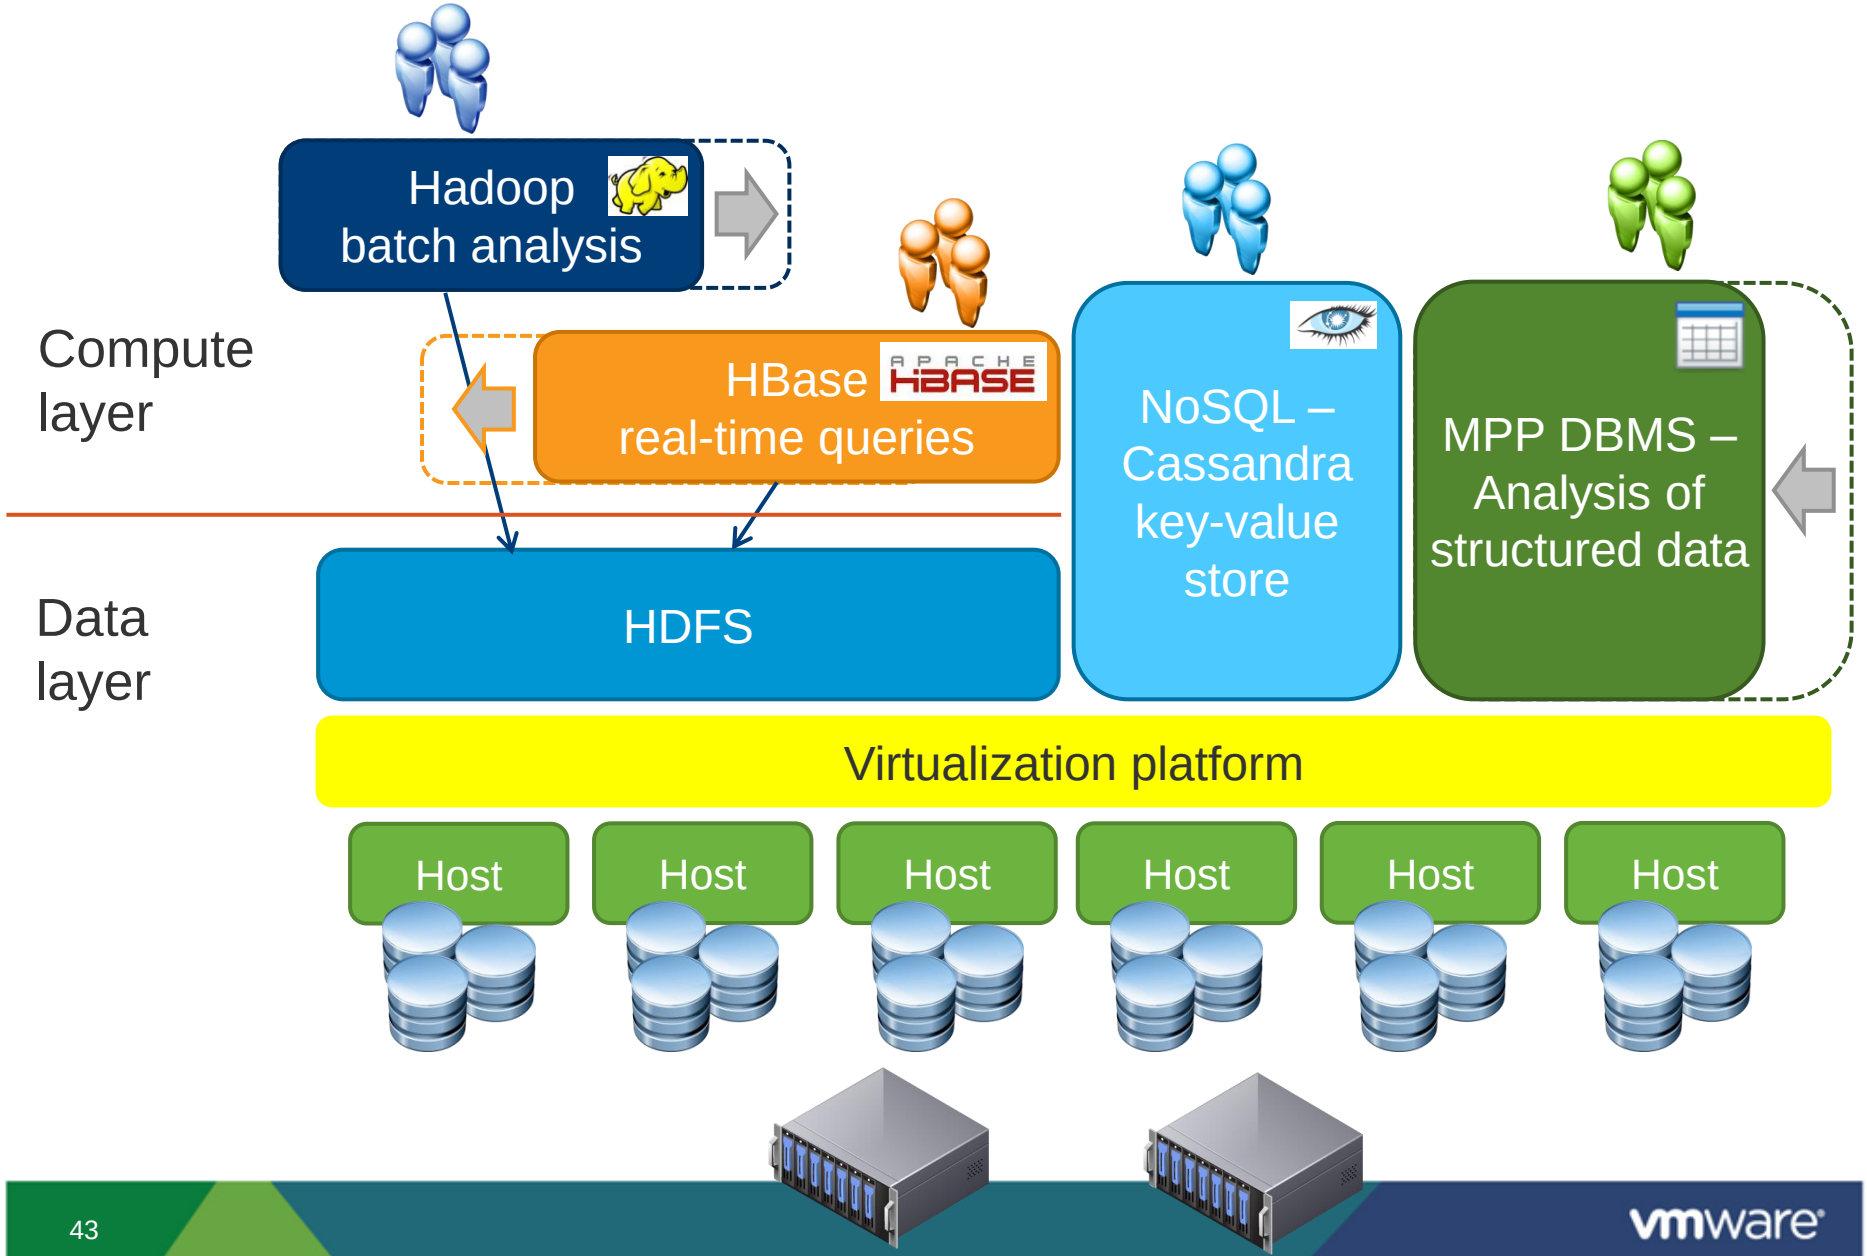

Agenda

- Hadoop Market Landscape
- Hadoop Journey
- Virtualize Hadoop Values
- **Summary**
- Q & A

VMWare brings Agility, Efficiency, and Elasticity to Big Data

Agility

- Deploy, configure and monitor Hadoop clusters on the fly
- Dynamic reconfiguring of Hadoop to meet changing business demands
- One click HA set up

Efficiency

- Consolidate Hadoop to achieve higher utilization
- Pool resources to allow for increased performance and priority job processing

Elasticity

- Enable full elasticity through separation of Data and Compute
- Scale In/Out Hadoop with Resource Constrain

Serengeti Resources

- **Download and try Serengeti**
  - projectserengeti.org

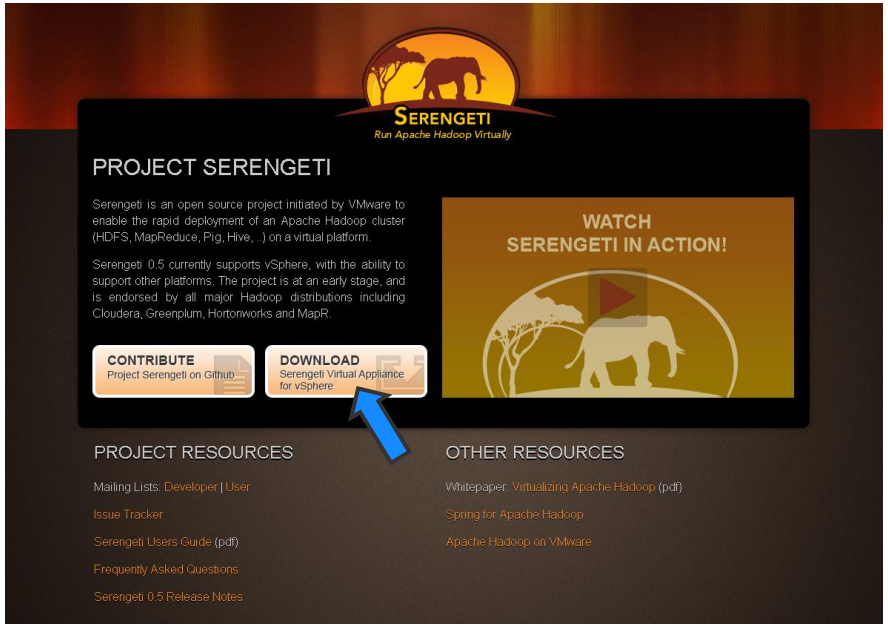

- **VMware Hadoop site**
  - vmware.com/hadoop

- **Hadoop performance on vSphere**
  - vmware.com/files/pdf/VMW-Hadoop-Performance-vSphere5.pdf
- **Hadoop High Availability solution**
  - vmware.com/files/pdf/Apache-Hadoop-VMware-HA-solution.pdf