Adobe Deploys Hadoop as a Service on VMware vSphere®

A TECHNICAL CASE STUDY

APRIL 2015

vmware®

Table of Contents

A Technical Case Study

This case study paper describes the deployment of virtualized Hadoop clusters in production and preproduction use at the Adobe Digital Marketing Business Unit. It gives a description of the business needs motivating this set of technologies and the use of a portal for providing internal and external users with a Hadoop-as-a-service capability. Adobe has been using tools such as VMware vSphere® Big Data Extensions™ and VMware vRealize™ Automation™ (formerly VMware vCloud® Automation Center™) for managing virtualized Hadoop clusters since early 2013.

Background

Adobe is a well-known international software company whose products are used by millions of people worldwide. Adobe's business is composed of many different parts, one of which is the Adobe Digital Marketing Business Unit, where products and services cover functionality for analytics, campaign management, experience management, media optimization, social networking, and targeting.

Adobe has been using Hadoop internally for a set of different applications for several years and has achieved significant gains in analytics on its key customer data. The Hadoop ecosystem has been deployed successfully on VMware vSphere for several different groups of Adobe users within the Digital Marketing Business Unit.

Why Virtualize Hadoop on vSphere?

Adobe chose to use vSphere as the strategic platform for hosting its Hadoop-based applications for several reasons:

- To provide Hadoop clusters as a service to the end users and development community, reducing time to insight into the data
- To increase the adoption of Hadoop within the homegrown analytics platform across the company by enabling rapid provisioning of custom clusters
- To reduce costs by increasing utilization on existing hardware rather than purchasing dedicated servers
- To avoid the "shadow IT" use of Hadoop by different units in an uncontrolled manner
- To develop an in-house platform as a service for the developer community to use

The Adobe Marketing Cloud and the Digital Marketing Business Unit

In 2009, Adobe acquired Omniture, which became the foundation of the Adobe Digital Marketing Business Unit. This business unit manages the digital assets of its customers, such as email, social media, and Web presence. One key part of that business is the generation of regular reports for customers on the types of users that utilize the customers' Web sites. These regular reports provide details on the end-customer interaction, such as answers to questions on how much of their Web site usage is repeat business and where the time is being spent online. This information is generated in reports but is also available online to customers so they can form their own queries against it and draw new business conclusions. One of the main motivations for adopting big data analytics was to extend this capability to put the power into the customers' hands to interact intelligently with their own data.

The Adobe Virtual Private Cloud: the Platform for Running Hadoop

The Adobe Digital Marketing Business Unit has created a private cloud infrastructure-as-a-service offering called the Adobe Virtual Private Cloud (VPC), which is supported by vSphere technologies. The Adobe VPC is utilized to various degrees by all of the Adobe Marketing Cloud applications. Adobe currently has five production data centers worldwide that support the VPC, with expansion plans for three additional data centers by mid 2015. Each of these data centers has a Hadoop platform capability deployed by the Big Data Extensions technology across hundreds of servers. The VPC is the foundational layer for all applications and products within the business unit. Adobe did not build a separate infrastructure for the Hadoop clusters; rather, it made use of the servers and foundational software within its VPC for this purpose.

Two Hadoop User Communities – Requirements

Adobe separates its use of Hadoop for two diverse groups of people.

1. **The Engineering Environment (Preproduction)**

   Multiple teams in Adobe engineering wanted to work with various configurations and distributions of Hadoop software for their own purposes. This area conducts new application development, functional workload testing, and experimentation with a variety of data stored in Hadoop. Adobe's forward-thinking development of new applications takes place here. Hadoop clusters are provisioned, used for a period of time, and then can be torn down so that new ones can be created for different purposes. These systems are not utilized for performance testing. Performance testing is done in a special environment to help identify the actual workload after the application has incorporated Hadoop.

   Some of the preproduction users had not used Hadoop prior to the Hadoop-as-a-service functionality's becoming available. Their requirement was to be able to bring up a cluster easily for learning purposes.

   The preproduction hardware setup matches the production environment for Hadoop at the company. This is done to guarantee that an application under test is experiencing the same conditions as it would in the production environment and that versions of the various software components are compatible.

   The perception among the engineering departments was that Technical Operations was not moving fast enough for them in the big data area. They had already gone out to the public cloud providers to get what they wanted in terms of MapReduce functionality. This sidestepping of the traditional operations function eventually produced unwanted side effects such as virtual machine sprawl and higher costs. This occurs among early adopters of Hadoop after the operations team no longer controls the resources. The essential requirement was to accelerate the software engineering life cycle for the teams engaged in Hadoop work by provisioning clusters for them more rapidly.

   Much of the data they need is shared across different Adobe Digital Marketing Business Unit products. The engineers needed answers to two essential questions:

   - Where is my data?
   - How am I going to access that data?

   There is a separate dedicated environment for performance testing. However, sizing of the environment to suit a particular Hadoop workload can be done in preproduction. This is a common occurrence with the newer types of applications and new frameworks, such as YARN. In the preproduction world, users can provision any vendor's distribution of Hadoop at any version level they want. There are no limitations imposed by the operations team. This situation changes when we get to a production environment.

2. The Production Environment

In this environment, the business need was to enable the Marketing Cloud product teams to take advantage of existing data that is managed by Technical Operations. The analytics processing requirement was to be able to process quantities of up to 30Tb of stored data over long periods of time (several years). The purpose of the production Hadoop clusters is to serve both Adobe's needs and those of its customers who want to be able to ask questions of their own data that is stored in Adobe's VPC. The end goal is to provide users outside the company with enough tools to enable them to express any query they have. There are some rules applied to the production infrastructure, such as the placement of more limits on the types and versions of the Hadoop distributions that are allowed to be provisioned.

Access control to the Hadoop clusters is therefore a prime concern. One Adobe customer should not be able to see or work with another customer's data. This multitenant requirement made security and isolation of individual virtual machines from each other a high priority for the deployment engineers in Adobe Technical Operations. As one security measure, only one distribution of Hadoop of a particular version number is allowed to be provisioned in the production environment. Limited numbers of default virtual machine sizes and configurations were also instituted as an operating principle at the beginning and helped simplify the task.

- Standard – virtual machine (hardware version 10) with four vCPUs and 8GB of RAM. The disk size depends on cluster sizing needs.
- Large – virtual machine (hardware version 10) with eight vCPUs and 16GB of RAM. The disk size depends on cluster sizing needs.

The Adobe Hadoop-as-a-Service Portal

Adobe Technical Operations customized the vRealize Automation user interface to conform to the standard of the organization, as shown in Figure 1. This illustrates a set of catalog entries that make up the standard configurations of Hadoop and other tools that Adobe supplied to its user communities. An end user with suitable entitlements can log in to this portal and provision an entry to get a Hadoop cluster up and running rapidly. The end user is not concerned about the exact hardware configuration that supports the new Hadoop cluster. They know that there is a pool of hardware and storage that supports the cluster and delivers functionality for the application to use.

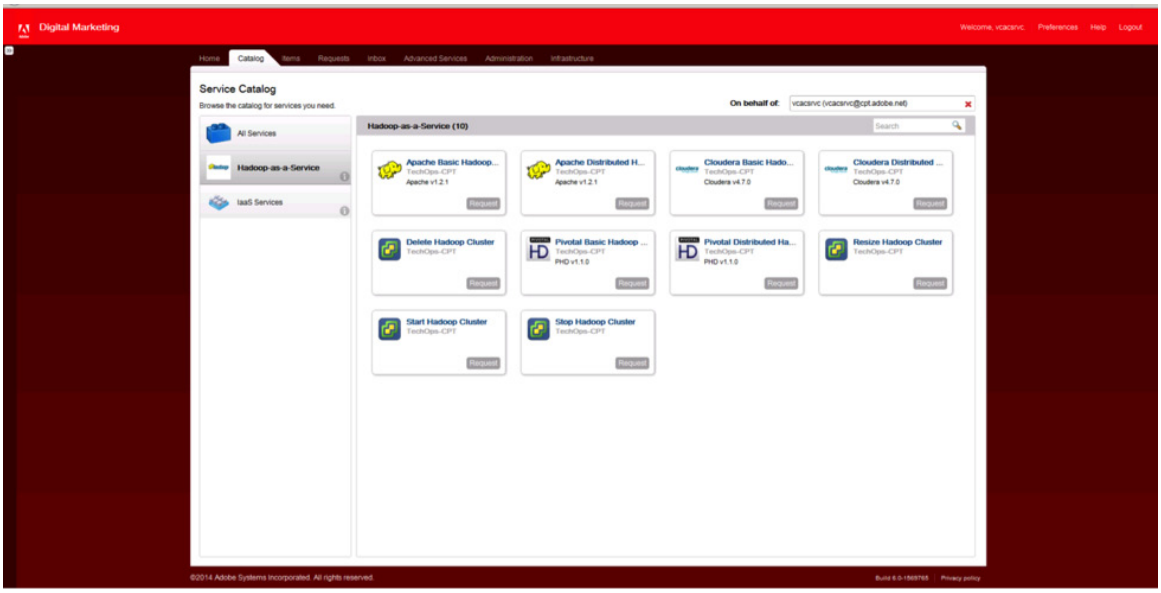

Figure 1. Top-Level View of the Hadoop Provisioning Portal

Figure 2 illustrates the details of the individual catalog items, indicating the nature of each cluster along with its distribution (Apache, Cloudera CDH, or Pivotal PHD) and version number. The end user is not confined to just one type of cluster. They can choose to work with different topologies such as compute-only clusters, resizable clusters, and basic clusters and can decide on the correct Hadoop distribution for their needs.

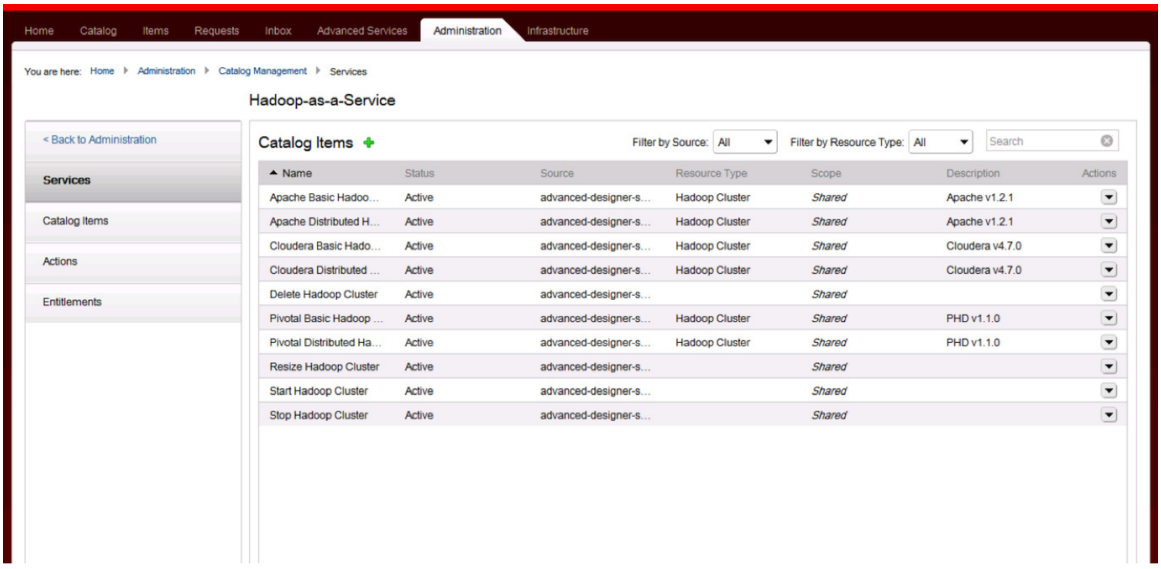

Figure 2. Details on the Catalog Items in the Hadoop Provisioning Portal User Interface

The portal based on vRealize Automation for managing the Hadoop clusters shown in Figure 2 enables those clusters to be started, stopped, and resized by the appropriate users. The "cluster resize" functionality enables virtual machines to be added to the compute tier of a Hadoop cluster to increase its processing capacity. Adobe's engineering staff has customized the workflows that drive these actions, so that one vRealize Automation instance is controlling Hadoop clusters across several VMware vCenter™ domains and even across data centers.

Technical Architecture

The high-level technical architecture for the Hadoop cluster provisioning portal solution built by Adobe is shown in Figure 3. Here we see that the vSphere virtualization layer is the base platform for the whole system. When Hadoop is provisioned by the end user, the different roles or daemons in Hadoop become processes that are running in various virtual machines that vSphere manages. These are the ResourceManager, NodeManager, and ApplicationMaster parts of the Hadoop infrastructure that are mapped into one or more virtual machines by the Big Data Extensions component in the middle (green) layer.

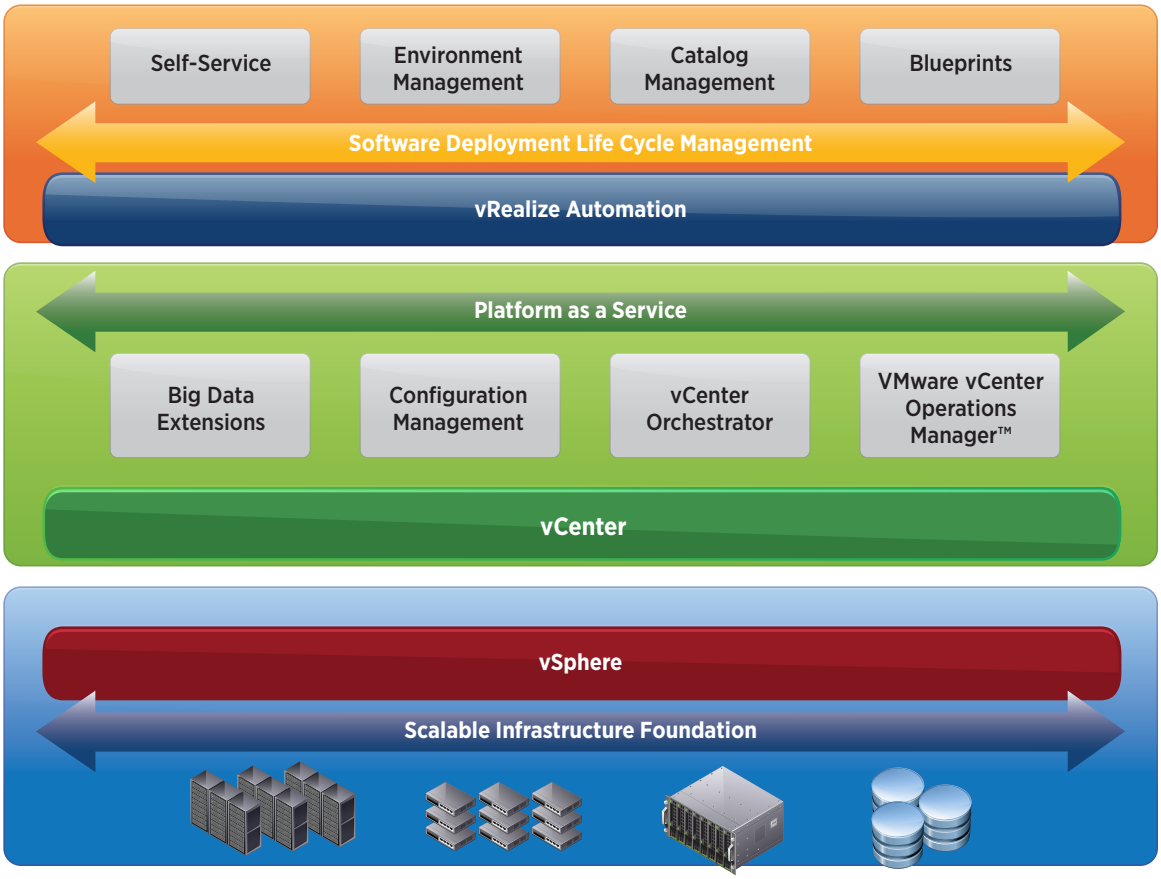

Figure 3. Outline – Technical Architecture for the Adobe Hadoop Provisioning Portal

The end user does not interact with the Big Data Extensions user interface directly, however. Instead, a separate software layer, shown in orange in the top layer of Figure 3, provides that interaction. The control flow and user interface or portal is built using the vRealize Automation software. The vRealize Automation runtime engine, augmented with certain workflows that are customized by the Adobe Technical Operations architects, invokes the Big Data Extensions APIs on the user's behalf to implement the Hadoop cluster. Big Data Extensions works with vCenter and enables a vSphere administrator to provision Hadoop clusters onto virtual machines in a user-friendly way. More information on this feature set can be found at http://www.vmware.com/bde.

Taking this outline of the architecture down one level to look at some of the details, we see the components in Figure 4. Clients use functionality provided by the vRealize Automation environment, which has a set of workflows and blueprints preinstalled in it for interacting with the Big Data Extensions plug-in to vCenter. The Big Data Extensions management server invokes the vCenter APIs to clone a template virtual machine, which is provided inside a vSphere virtual appliance. This virtual appliance is created when Big Data Extensions is first installed.

The vRealize Automation component was customized by the Adobe engineers to have the correct look and feel for the company as well as unique workflows to satisfy the particular Adobe Digital Marketing Business Unit requirements. This involved some changes to VMware vCenter Orchestrator™ workflows that are built into the vRealize Automation solution. Likewise, the Big Data Extensions environment was also customized by the addition of YUM repositories that contain the correct components for installation of a particular distribution of Hadoop. This is shown in the top-right corner of Figure 4, shaded in light blue. After this setup is in place, users can utilize a Web portal that hides some of the details in Big Data Extensions from them, making the provisioning process simpler for them to use.

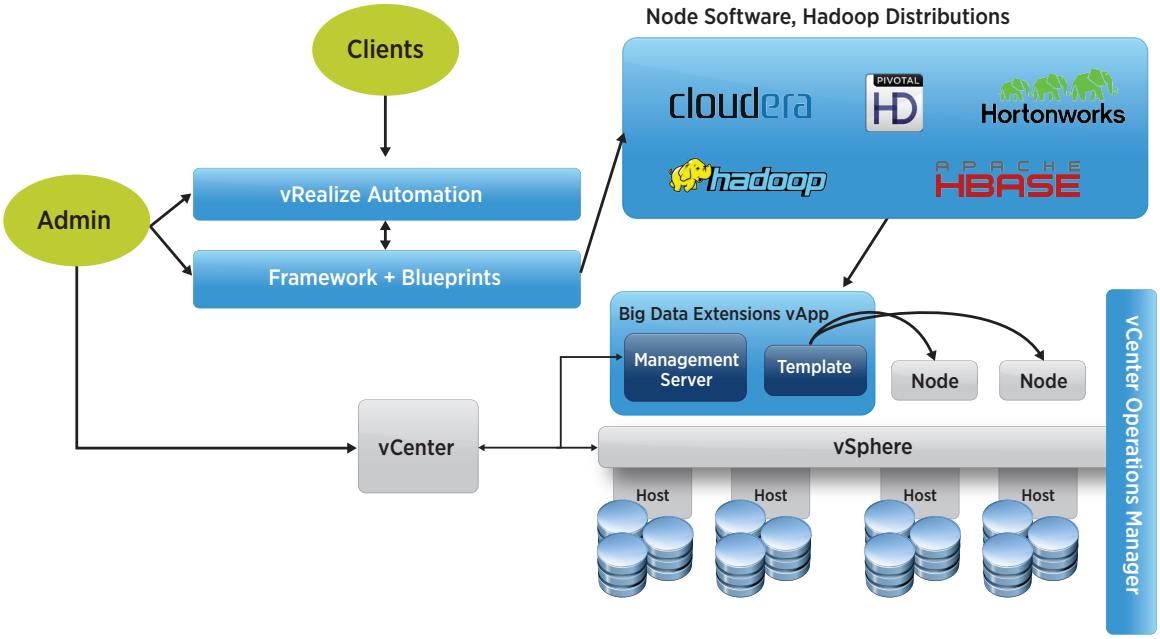

Figure 4. Details – Technical Architecture for the Adobe Hadoop Provisioning Portal

The Hardware Layout

Servers and Data Centers

At the time of writing, Adobe has more than 8,000 virtual machines supporting various applications. This number is growing constantly as more applications are virtualized. Within the big data area alone, the following data centers and machines have already been deployed:

- One preproduction data center
  - There are 200 host servers deployed currently.
  - There are dozens of small Hadoop clusters running in preproduction at any given time.
- Five production data centers worldwide
  - There are 1,000 host servers deployed currently.
  - The number of virtual machines containing Hadoop functionality is continually growing and will be in the thousands as the system is adopted widely over time.
- Adobe staff are conducting tests with an additional 256 virtual machines at the time of writing.

All the VPC data centers are backed by a combination of Hitachi SAN and EMC Isilon storage.

- Hitachi Unified Storage VM – Utilizes tiered storage with Flash for acceleration of performance and transaction rates across the array
- EMC Isilon X400 nodes – 4U nodes that provide flexible balance between high-performance and large-capacity clusters

The VPC supports 10GbE networking based on Cisco technology. The server hardware comprises the Dell PowerEdge M1000e blade enclosure fully populated with Dell PowerEdge M620 blade servers. Each blade server is configured with 256GB of RAM and two Intel Xeon Processor E5-2650 CPUs.

Operating system volumes for the various virtual machines are stored on the Hitachi SAN. There is an option also for HDFS data storage on an EMC Isilon NAS system. The Isilon storage holds more than 10Pb of data that it can present to the Hadoop compute tiers using the HDFS protocol.

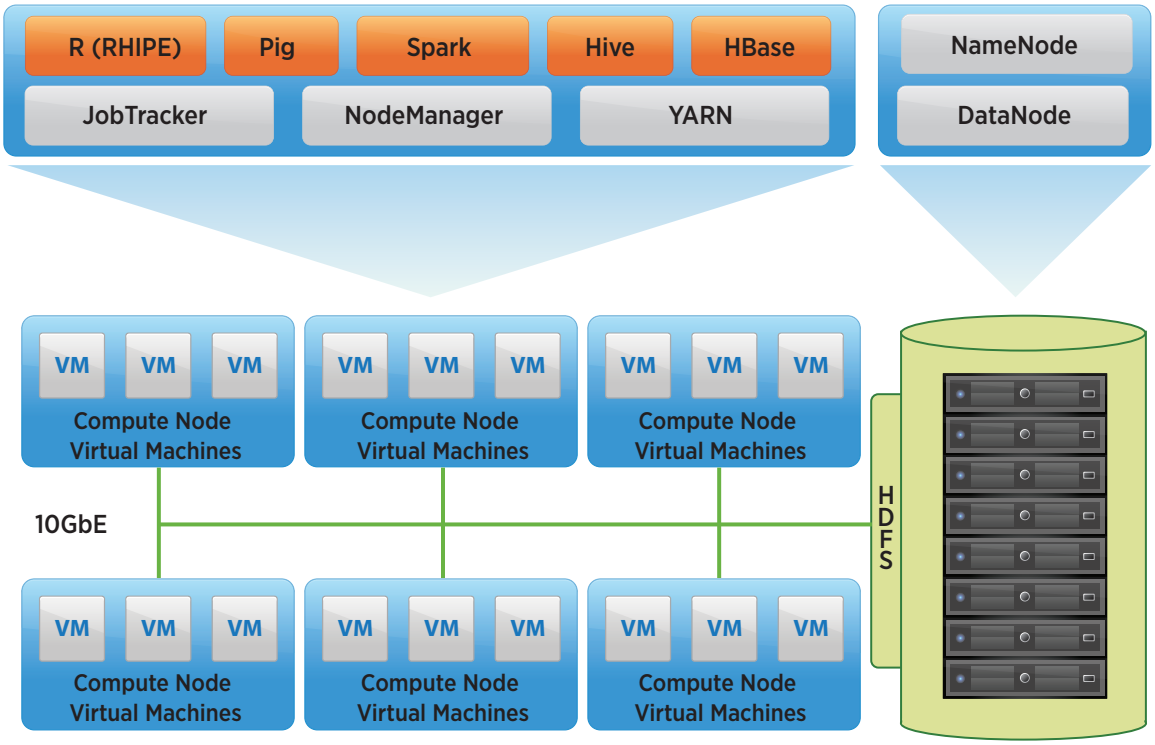

Figure 5. Technical Architecture for the Hadoop Deployment

Configuration of Hadoop Clusters on vSphere

Big Data Extensions, used by the vRealize Automation tool, places virtual machines that it provisions into resource pools by default. Two separate resource pools in vSphere can have different shares, limits, and reservations on CPU and RAM, enabling them to operate independently of one another from a performance perspective while sharing hardware.

The innermost resource pools are created along **Node Group** lines as specified by the user when doing the initial design of a virtualized Hadoop cluster. The **Master** and **Worker** node groups shown in Figure 6 are located in separate resource pools, for example. In vSphere, resource pools can be nested within other resource pools.

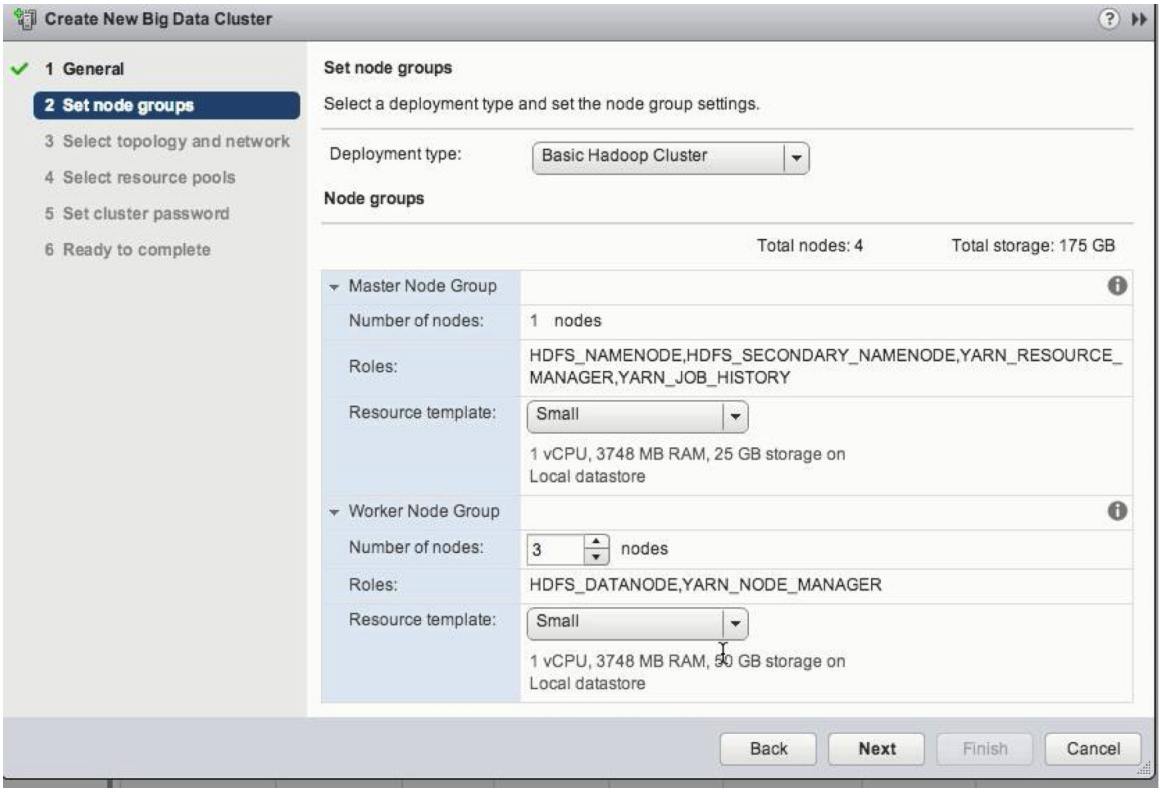

Figure 6. Configuring the Hadoop Roles Within Virtual Machines for a New Cluster

The outermost resource pool, into which the new Hadoop cluster is placed, is set up by the user before Big Data Extensions provisioning starts. There are eight server hosts in an outermost resource pool by default, but this can be changed without affecting the overall architecture. Different resource pools are used in vSphere to isolate workloads from each other from a visibility and performance perspective.

At the VMware ESXi™ level, there is one Hadoop worker node—that is, one virtual machine containing the NodeManager in Hadoop 2.0—per hardware socket on the server hosts. This is a general best practice from VMware and supports locality of reference in a NUMA node as well as being a good fit of virtual CPUs onto physical cores in a socket.

Monitoring and Management of Hadoop Clusters

vRealize Operations Management

The VMware vRealize Operations™ tool provides Adobe Technical Operations and engineering organization with insights into the behavior patterns and performance trends that the Hadoop clusters are exhibiting. The operations management tool gathers data to interpret the "normal" behavior of the system over time and can then identify growing anomalies from that normal behavior to its users. Garbage collection times within the Java virtual machines supporting many of the Hadoop roles are an important measure that Technical Operations pays a lot of attention to. If these times grow at a pace over a period, vRealize Operations points out these patterns to the end user, helping to create a more timely solution to an impending problem. The tool is also used in the preproduction environment to assess the size of a suitable virtual Hadoop cluster for supporting new applications.

Important Insights from Virtualizing Hadoop

The engineering teams come to Technical Operations for deployment of a new application workload and have a service-level agreement that they must fulfill. Whereas in the past, the engineering team would have had the added burden of carrying out capacity management and infrastructure design themselves for the new application, now Technical Operations can shorten the time to complete these tasks by providing a temporary virtualized environment in which the application can be configured and tested. Ultimately, the business decision on where best to deploy the application can be made by teams who are not the original application developers.

The engineering team might not have had the opportunity to fully test out the applications under load. This is now done in the preproduction Hadoop cluster previously described. With the help of the vRealize Operations management tool, the teams can determine the best deployment design for the application as well as which virtualized data center it would best fit into.

Future Work

1. Cloudera Manager – Complementing the vRealize Operations tool, Technical Operations also uses Cloudera Manager extensively to view the Hadoop-specific measures that indicate the health of any cluster. Indicators such as HDFS and MapReduce health states enable the operator to see whether or not the system needs attention. These measurements are correlated with data from the infrastructure level such as the consumption of resources by the virtual machines. These complementary views enable one to make decisions about optimizing workloads on different hosts in a data center.
2. The Adobe Digital Marketing Business Unit is investigating using Big Data Extensions to provide data-only clusters as described in the joint *Scaling the Deployment of Hadoop in a Multitenant Virtualized Infrastructure* white paper written by Intel, Dell, and VMware.

Conclusion

Providing Hadoop cluster creation and management on VMware vSphere has optimized the way that big data applications are delivered in the Digital Marketing Business Unit at Adobe. The costs of dedicated hardware as well as management and administration for individual clusters have been greatly reduced by sharing pools of hardware among different user communities. Developers and preproduction staff can now obtain a Hadoop cluster for testing or development without concerning themselves with the underlying hardware. Virtualization has also improved the production Hadoop cluster management environment, enabling more operations control over resource consumption and better management of the trade-offs in performance analysis. Virtualizing Hadoop workloads has become an important part of the Adobe Digital Marketing Business Unit, empowering Technical Operations to become a service provider to the business.

Further Reading

1. Adobe Digital Marketing Business Unit Web Site
   http://www.adobe.com/#marketing-cloud
2. VMware vSphere Big Data Extensions
   http://www.vmware.com/bde
3. Virtualized Hadoop Performance with VMware vSphere 5.1
   http://www.vmware.com/resources/techresources/10360
4. Virtualized Hadoop Performance with VMware vSphere 6 on High-Performance Servers
   http://www.vmware.com/resources/techresources/10452
5. EMC Hadoop Starter Kits Using Big Data Extensions and Isilon Storage
   https://community.emc.com/docs/DOC-26892

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com