

ENABLING GLOBAL HADOOP WITH DELL EMC ECS

Hadoop Storage-as-a-Service

ABSTRACT

This White Paper illustrates how Dell EMC® ECS™ can be used to streamline the Hadoop data analytics workflow. It describes the core architectural components, and highlights the features an enterprise can use to leverage Hadoop globally for reliable business insights

March, 2018

The information in this publication is provided “as is.” EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

EMC², EMC, the EMC logo, and ECS are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners. © Copyright 2018 EMC Corporation. All rights reserved. Published in the USA. 04/18, white paper, H14844.1

EMC believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

EMC is now part of the Dell group of companies.

TABLE OF CONTENTS

EXECUTIVE SUMMARY	4
LIMITATIONS OF TRADITIONAL HADOOP DAS	4
ECS ELIMINATES THE LIMITATIONS OF HADOOP DAS	4
ECS OVERVIEW.....	5
ECS Multi-protocol Architecture.....	5
Hadoop on ECS.....	6
ECS HCFS (Hadoop Compatible File System).....	6
ECS Hadoop Supported Versions and Services.....	6
ECS Hadoop Specific Features	7
ECS Hadoop Authentication Modes	7
ECS Hadoop Shared Architecture	7
HOW ECS Protects Global Hadoop Data	8
The Advantages of Geo-distributed Hadoop.....	9
How ECS Streamlines Hadoop Workflow	9
ECS CERTIFICATION WITH HORTONWORKS	10
ECS IS THE DELL EMC DATA LAKE	10
CONCLUSION	10

EXECUTIVE SUMMARY

Apache Hadoop has emerged as the preferred tool for performing powerful data analysis. Modern Cloud Computing has created new opportunities to leverage the power of Hadoop by extending analytics from the traditional core of the organization (the main datacenter) to the edge (or field locations). Traditional Hadoop analytics involves a series of complex workflows and data movement. Dell EMC's ECS is a modern cloud architecture that enables organizations to simplify Hadoop workflow as well as extending the benefits of Hadoop to multiple locations anywhere in the world. This paper will discuss how ECS simplifies and modernizes Hadoop storage by providing a geo-distributed environment with enterprise-class data protection.

Limitations of Traditional Hadoop DAS

A traditional Hadoop storage environment consists of local DAS (Direct-Attach-Storage). This configuration was developed before the advent of high speed networking as large Internet companies had unlimited budgets to deploy and manage silos of Hadoop compute and storage clusters. The limitations to this approach is that an organization end ups with multiple instances of Hadoop DAS for every cluster, increasing the management task and limiting flexibility. In addition, the data to be analyzed needs to be migrated and converted for use with the Hadoop cluster in an ETL process. This data ingestion and conversion can take hours depending on the size of the data set.

Hadoop DAS has the following limitations:

- **Centralized NameNode.** The HDFS file system namespace is managed by a single server and is maintained in memory. The size of namespace that can be managed by the central NameNode is limited by the amount of memory available on the NameNode and performance of the file system is limited by the performance of the NameNode because all NameNode lookups need to be done from a single server.
- **Data Protection.** Traditional Hadoop DAS requires at least 3 copies of the data. First the data must be migrated to the local DAS, and then the copies must be made prior to analyses. This is time and storage consuming process.
- **File system reliability.** Prior to Hadoop 2.x, NameNode was a single-point of failure. Failure of the NameNode resulted in the cluster being unavailable. Recently, High Availability features have been added to HDFS but they have limitations: the hot Standby NameNode cannot actively process requests and new hardware is needed for the Stand-by NameNode.
- **Single-protocol.** HDFS DAS implementation provides for a single protocol for data access and does not expose other storage protocols like Object or File interfaces.
- **Storage overhead.** By default, HDFS DAS performs a 3X replication of all data blocks. This results in 2x overhead which is overkill for certain use cases like archiving.
- **Small-file problem.** HDFS DAS is inefficient at handling a large volume of small files because metadata for each file in the file system needs to be stored in the memory of a single server, the NameNode. For a million files in HDFS, a small number in big data terms, the NameNode consumes about 3GB RAM.
- **Architecture.** HDFS DAS was architected almost a decade ago and was designed for unreliable commodity disks and legacy network infrastructure (1GbE). A distributed system was assumed to be a bottle-neck due to network speeds and not the disk. These assumptions don't hold true in today's infrastructure with the proliferation of high speed networks
- **Lack of enterprise-grade features.** HDFS DAS lacks enterprise-grade features like Geo-Distribution, Disaster Recovery, Consistent Snapshots, Deduplication, Metering, etc.
- **Lack of multi-tenancy.** HDFS DAS lacks robust multi-tenancy features, which can provide data and performance isolation guarantees to tenant users. This results in multiple silo'd clusters in a large enterprise, resulting in low cluster utilization.

ECS Eliminates the Limitations of Hadoop DAS

Dell ECS is a single, shared, global architecture that modernizes Hadoop with a new paradigm. Global Hadoop eliminates the need for Hadoop DAS while accelerating time to results by reducing complexity:

- **Global Namespace.** ECS can be configured as one global namespace. It is designed for massive scale so the Hadoop file system is not burden by limitations with NameNode memory.

- **“In Place” Analytics.** ECS modernized and accelerates Hadoop workflow because the Hadoop cluster analyzes data on ECS without the need for copies or data movement.
- **File system reliability.** ECS is “enterprise grade” and delivers seven 9’s of availability for Hadoop data. In addition Hadoop data is mirrored and erasure coded for high availability.
- **Multi-protocol.** ECS supports multiple protocols (Object, File, and HDFS) enabling a central shared storage architecture for Hadoop data
- **Storage overhead.** ECS offers a much lower storage overhead than Hadoop DAS and does not require you to make 3 copies of the data for data protection. ECS also offers configurable “cold data” schemes to further increase storage efficiency
- **Small-file problem.** ECS integrates a patented feature called box-carting which increases “write” performance by aggregating multiple small data objects in memory and then writing them in a single write command. The result is much faster small file performance for Hadoop. This is ideal for IoT (Internet of Things) applications where millions of tiny data points need to be captured and analyzed. On the other hand, ECS has no limitation regarding the size of large objects.
- **High Performance Architecture.** ECS uses high speed 10gb Ethernet switches. A non-blocking 10 Gbps switch (up to 2500 MB/sec full duplex) can provide more bandwidth than a typical DAS subsystem with 8 disks (600 – 1200 MB/sec).
- **Enterprise-Grade features.** ECS delivers “enterprise-class” availability for Hadoop by eliminating single points of failure. ECS geo-distribution uses erasure coding and replication to ensure the availability of Hadoop. Hadoop applications would continue to run during a temporary site outage as ECS provides a highly available, self-healing data repository.
- **Robust Multi-tenancy and Tenant isolation.** ECS delivers a single, shared, global storage system for Hadoop which eliminates the complexity of maintaining and managing multiple DAS silos. Cloud Service Providers leverage ECS’s HDFS and robust multitenant features such as metering, chargeback and show back to deliver Hadoop Storage-As-A-Service similar to public cloud services. ECS can be divided into multiple namespaces providing an isolated and secure multitenant environment.

ECS Overview

ECS is a software-defined, scale-out object storage platform that deploys entirely in software. ECS features a unique unstructured storage engine that supports Object, File (NFS v.3), and HDFS storage. ECS enables HDFS access on the same unstructured storage engine. In addition ECS supports popular API’s like Amazon S3 and OpenStack Swift.

Customers have the choice of deploying ECS in two ways:

- **ECS Software on Commodity:** Customer can deploy the ECS software on their choice of commodity server and operating systems and writing data to certified 3rd party commodity JBOD .The customer is primarily responsible for management and operations, though ECS can monitor the health of the nodes.
- **ECS Appliance:** Customers can deploy ECS as an integrated appliance that includes all software and hardware. Dell EMC provides management and operational support.
- **Dedicated Cloud:** ECS Dedicated Cloud is a fully managed ECS service that gives customers the flexibility of choosing between ECS in their own datacenters, ECS in Dell EMC datacenters, or a hybrid option. It has rich capabilities of ECS and private cloud “control,” coupled with hands-off operations and the flexibility of the Public Cloud.

ECS provides a complete cloud-scale storage infrastructure designed to meet the requirements of cloud, mobile, Big Data and social applications. Dell EMC ECS Appliance brings the cost-profile, simplicity, and scale of a public cloud to anyone - with the trust, reliability, and support you expect from Dell EMC.

ECS MULTI-PROTOCOL ARCHITECTURE

ECS supports the storage, manipulation, and analysis of unstructured data on a massive scale using commodity infrastructure. ECS features an unstructured storage engine that supports object and HDFS data services:

Object

The Object data service enables the storage, access, and manipulation of unstructured data as objects on commodity-based systems, such as the HP SL4540, and the ECS Appliance. The Object data service is compatible with Amazon S3, OpenStack Swift, Dell EMC Atmos and Centera CAS object storage APIs.

HDFS

The HDFS data service provides support for the Hadoop Distributed File System (HDFS). The HDFS data service enables organizations to build a big data repository at scale. With the HDFS data service, organizations can use the ECS storage environment as a big data repository against which they can run Hadoop analytic applications.

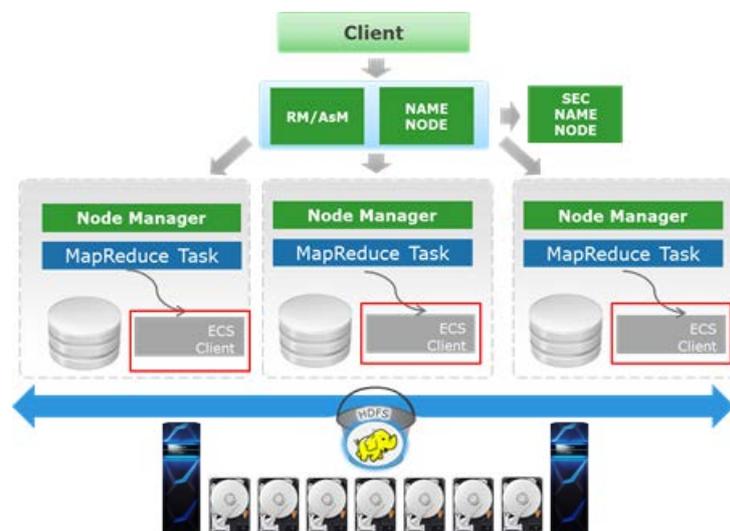
The ECS unstructured storage engine is the primary component that ensures data availability and protection against data corruption, hardware failures, and data center disasters. The engine exposes multi-head access across object and HDFS so that an application can write object and read through HDFS or vice versa, and allows access to the same data concurrently through multiple protocols. Today, ECS supports object and HDFS interfaces.

File (NFS)

With the introduction of ECS v2.2 a file head feature has been added to enable global namespace file capability for ECS. Key features include native NFS v.3 support, rich ACL's, global namespace, and global locking. ECS NFS supports file sizes up to 4 TB and offers the ability to scale NFS data across 8 ECS sites. Because of ECS cloud scale design, there are no limits to the number of files or directories. ECS does not support CIFS at this time, but this capability can be added via various third party gateway solutions.

HADOOP ON ECS

The ECS HDFS is a Hadoop Compatible File System (HCFS) that enables you to run Apache Hadoop 2.7 applications on ECS. Customers can configure any existing Hadoop cluster running any Hadoop distribution to run on ECS HDFS.



ECS HDFS Architecture

When customers set up the Hadoop client to use ECS HDFS instead of traditional HDFS, its configuration points to the ECS HDFS to do all the HDFS activity. On each Hadoop HDFS client node, any traditional Hadoop component would use the ECS HDFS client (JAR) to perform the HDFS activity. Starting with ECS 2.2, customers can use Apache Ambari to configure ECS HDFS as a primary file system.

ECS HDFS (HADOOP DISTRIBUTED FILE SYSTEM)

The ECS unstructured storage engine is the primary component that ensures data availability and protection against data corruption, hardware failures, and data center disasters. The engine exposes multi-head access across object, file and HDFS so that an application can write object and read through HDFS or vice versa, and allows access to the same data concurrently through multiple protocols. ECS HDFS (Hadoop Distributed File System) enables you to use your portal storage infrastructure as a Big Data repository. Consequently all data stored within ECS is accessible by Hadoop. An important benefit of ECS is that Hadoop analytical applications can be run in-place without moving data to local storage. ECS acts as a Hadoop data node in a Hadoop cluster.

ECS HADOOP SUPPORTED VERSIONS AND SERVICES

ECS HDFS is compatible with Apache Hadoop 2.7 and support fine grained ACL's (Access Control Lists) and extended file system attributes. ECS is certified with and has deep integration with Hortonworks Ambari. ECS HDFS can be administered and

configured quickly and easily with Ambari's wizard driven GUI. In addition ECS has been tested with Cludera 5.4 and Pivotal 3.0. Other ECS feature enhancements include:

- Proxy User Authentication – impersonation for Hive, HBase, and Oozie
- Server side ACL enforcement with and the addition of a Hadoop “Superuser” and “SuperUser” groups.

Supported services include: HDFS, YARN, MapReduce, Pig, Hive/Hiveserver 2, HBase, Zookeeper, Flume, Spark, Sqoop.

ECS HADOOP SPECIFIC FEATURES

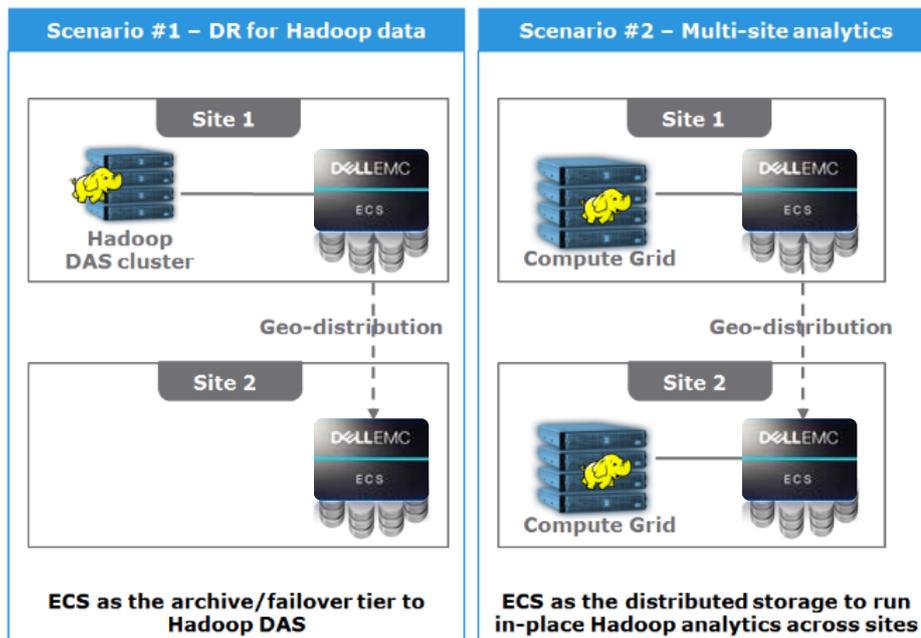
ECS incorporates specific features that benefit Hadoop users, these include:

- Support for Apache Hadoop 2.7 including fine grained ACL's and extended file system attributes
- Native Ambari integration now greatly simplifies registering ECS with the Hadoop cluster
- Proxy User Authentication for Hive, HBase, and Oozie
- Security features including server-side ACL enforcement, the ability to create ad Hadoop superuser and superusergroup
- Hadoop performance enhancements with memory and pool-size tuning

ECS HADOOP AUTHENTICATION MODES

ECS HDFS integrates with Hadoop clusters configured to use either simple or Kerberos authentication modes. Hadoop applications access data stored in ECS buckets, so the Hadoop user must have permissions to read and write to those objects. Hadoop services such as MapReduce, hive, hbase must have the appropriate permissions to write system files.

ECS HADOOP SHARED ARCHITECTURE



Scenario #1 Disaster Recovery for Hadoop Data

ECS's geo-distribution is leveraged to provide a highly available environment for Hadoop. The benefits include:

- Provides high-availability for Hadoop by eliminating single points of failure
- Low cost data archive
- 3 copies of the data are not required
- “Data in Place” analytics
- Higher scalability and lower cost than DAS

Scenario #2 Global Hadoop Storage

Multiple ECS's are leveraged as part of a data lake scenario to provide a global platform for Hadoop

- Huge cost savings by eliminating the need for multiple Hadoop DAS environments
- Multiple Hadoop users can analyze the same data concurrently regardless of geographic location
- Multiple users can analyze data with their choice of Hadoop distribution
- Provides high-availability for Hadoop by eliminating single points of failure
- Low cost data archive
- 3 copies of the data are not required
- "Data in Place" analytics
- Higher scalability and lower cost than DAS

HOW ECS PROTECTS GLOBAL HADOOP DATA

The ECS HDFS brings critical enterprise-grade features to Hadoop storage. The ECS HDFS increases the efficiency, speed and reliability of Hadoop and provides:

- **Massive scalability** - The ECS Appliance has been proven to effortlessly scale to Petabyte and Exabyte storage demands. The ECS architecture allows compute and storage to scale independently.
- **Multi-Protocol Access** - ECS offers unparalleled accessibility in a single platform with support for multiple Object APIs, NFS File and HDFS access. Developers are free to focus on their application and not operations.
- **Geo-protection** - ECS geo-protection provides full protection against a site failure should a disaster or other calamity force an entire site offline. A geo-protection layer in ECS protects data across geo- distributed sites and ensures that applications seamlessly function in the event of a site failure.
- **Multi-Site Access** - ECS provides strong, consistent views of data regardless of where the data resides. With geo-protection enabled, applications can access data immediately through any ECS site, regardless of where the data was last written.
- **Efficiency** - Erasure coding provides storage efficiency without compromising data protection or access. The ECS storage engine implements the Reed Solomon 12/4 erasure-coding scheme in which a chunk is broken into 12 data fragments and 4 coding fragments. The resulting 16 fragments are dispersed across nodes at the local site. The storage engine can reconstruct a chunk from a minimum of 12 fragments.
- **Flexibility** - The ECS HDFS allows IT departments to choose multiple Hadoop vendors and connect them all to use the ECS HDFS as the storage substrate to allow in-place analytics. Snapshots, Versioning and Journaling are also natively built into the platform.
- **Storage Efficiency for Large and Small Files** - The ECS storage engine is adept at handling both a high volume of small files as well as very large. Using a technique called box-carting, ECS can execute a large number of user transactions concurrently with very little latency. This enables ECS to support workloads with high transaction rates. ECS is also very efficient when handling very large files. All ECS nodes can process write requests for the same object simultaneously and each node can write to a set of three disks.
- **Security and Compliance** - ECS is Section 508 compliant. Beginning with ECS 2.2, ECS supports SEC Rule 17a-4(f) which regulates the storage of electronic records. SEC compliance is enabled at the bucket level as users can select "Compliance" when creating a bucket. Existing buckets cannot be upgraded to compliance mode and "Compliance" and only be applied to buckets with S3 or CAS data. ECS SEC compliance enforces the following policy; retention periods can be extended but not reduced or deleted. Records with "compliance" status cannot be deleted (this includes privileged delete); a bucket cannot be deleted if it has record objects under retention.

THE ADVANTAGES OF GEO-DISTRIBUTED HADOOP

ECS's geo-distributed architecture provides unique benefits that modernize and expand Hadoop benefits. Because ECS is a single, shared, global architecture it enables unique advantages for global organizations.



These include:

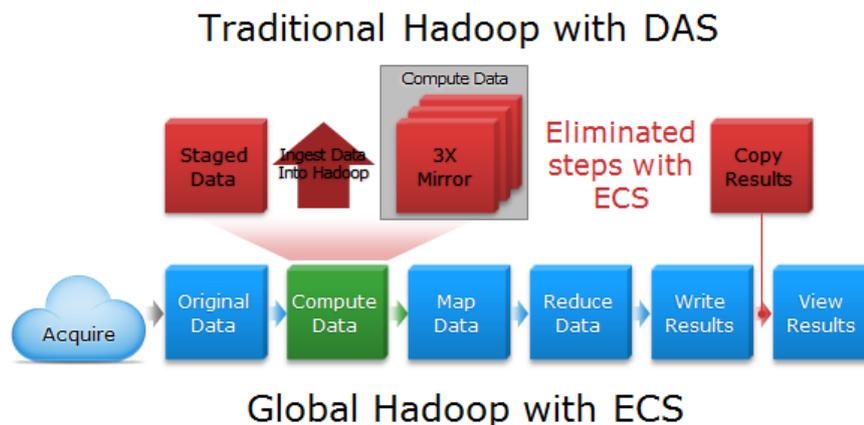
- The ability to reduce complexity by replacing multiple Hadoop DAS installations with one easily managed, ECS system
- Data "In Place" analytics mean that data does not have to be moved or migrated, it can be analyzed right on ECS from anywhere in the world
- ECS geo-distribution enables companies to expand the benefits of Hadoop analytics to the edge by consolidating all data in one repository.

Consequently multiple Hadoop clusters anywhere in the world can easily access data in ECS

- ECS supports multiple Hadoop distributions so the field can choose their preferred Hadoop application
- ECS multitenant architecture enables the delivery of Hadoop storage-as-a –service to multiple lines of business or customers
- The field can analyze data store in ECS concurrently without bottlenecks accelerating time to results

HOW ECS STREAMLINES HADOOP WORKFLOW

ECS provides significant benefits in the reduction in time and complexity for Hadoop workflow. Because of ECS's "in place" analytics feature, data does not have to be migrated to a local Hadoop DAS. Multiple Hadoop compute clusters anywhere in the world can access and analyze the ECS data store without the need for migration or the requirement to make a minimal three copies of the data. The graphic below illustrates the steps that ECS eliminates from the workflow.

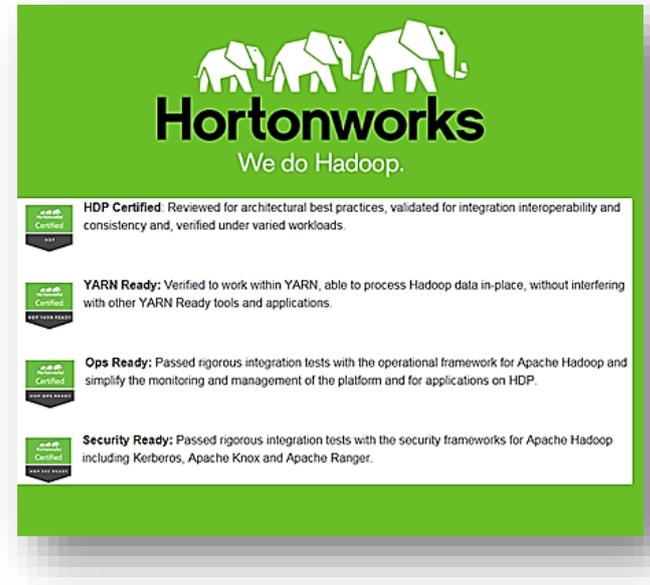


ECS Certification with HortonWorks

Dell EMC has entered into an agreement with Hortonworks to license their proprietary test suite to qualify HDP on ECS. Dell EMC will also be committing code changes to enable Ambari to deploy ECS backed Hadoop Clusters. As part of the agreement, Dell EMC will handle level 1 support with Hortonworks handling level 2 and 3 escalations. ECS has been certified for the Hortonworks services in the adjacent graphic.

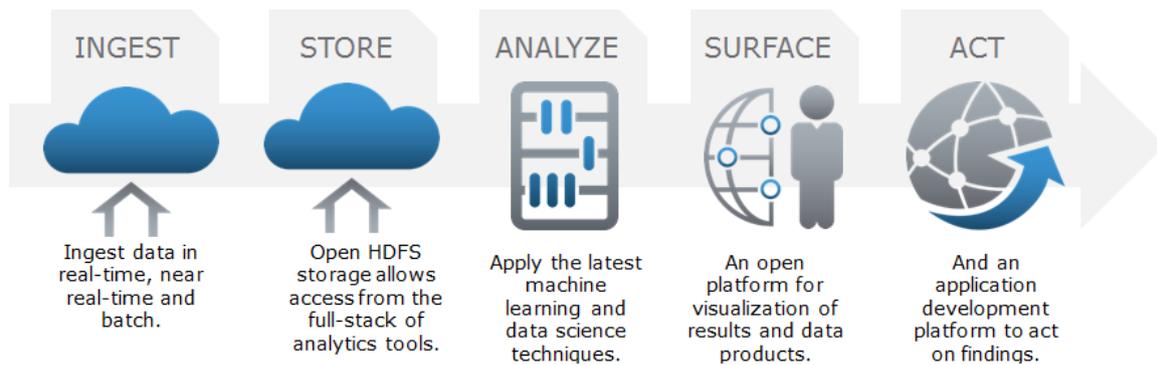
Benefits for customers include:

1. Certified Integration for faster time to market
2. Easy deployment with Ambari configuration wizard
3. Integrated support delivers a better user experience



ECS is the Dell EMC Data Lake

ECS is an important component of the Dell EMC Data Lake which enables the cloud scale storage for modern applications and the Internet of Things. ECS is the cloud between data ingestion or collection and Hadoop analyses. Because ECS is object storage, data can be captured and stored directly via HTTP without the need for authentication servers or cloud gateways. And because all ECS data is accessible via HDFS, data does not need to be transformed because it is already stored in a format that Hadoop understands. These benefits result in lower costs and a faster time to results for analytic workflows.



Conclusion

Hadoop is a low-cost, highly scalable, distributed analytics engine that can significantly reduce the time and resources needed by an enterprise to derive valuable insight from their Big Data assets. Hadoop and the ECS HDFS data service integration allows organizations to utilize a software-defined storage infrastructure as a native part of their Hadoop architecture while also providing simplicity, flexibility and rapid deployment. In addition ECS geo-distribution and inherent data protection modernizes Hadoop workflow and simplifies management while accelerating time to results.