



Solving Hadoop® Security

A Holistic Approach to a Secure Data Lake

A Hortonworks White Paper
JULY 2015

Contents

Overview	3
Understanding the security implications of the Data Lake	3
Only as secure as the weakest link	4
Commitment to enterprise-readiness	6
Administration	6
Authentication and perimeter security	7
Authorization	9
Audit	11
Data protection	12
Enabling extensibility through a pluggable framework	13
Comparison across security pillars	13
Summary	14
About Hortonworks	14

Overview

As companies rush to put Big Data to work for their business, new ways of operating can sometimes get ahead of IT's ability to digest their full implications. There's no question that the creation of a Hadoop-powered Data Lake can provide a robust foundation for a new generation of analytics and insight, but it's important to consider security before launching or expanding a Hadoop initiative. By making sure that data protection and governance are built into your Big Data environment, you can leverage the full value of advanced analytics without exposing your business to new risks.

Hortonworks understands the importance of security and governance for every business. To ensure effective protection for our customers, we use a holistic approach based on five pillars:

- Administration
- Authentication and perimeter security
- Authorization
- Audit
- Data protection

In each of these areas, Hortonworks provides differentiated capabilities beyond those of other vendors to help customers achieve the highest possible level of protection. As a result, Big Data doesn't have to incur big risks—and companies can put it to work without sacrificing peace of mind.

Understanding the security implications of the Data Lake

The consensus is strong among leading companies in every industry: data is an essential new driver of competitive advantage. Hadoop plays a critical role in the [modern data architecture](#) by providing low-cost, scale-out data storage and value-add processing. The successful Hadoop journey typically starts with [Data Architecture Optimization](#) or new [Advanced Analytic Applications](#), which leads to the formation of a [Data Lake](#). As existing and new types of data from sensors and machines, server logs, clickstreams, and other sources flow into the Data Lake, it serves as a central repository based on shared Hadoop services that power deep organizational insights across a large, broad and [diverse set of data](#).

The need to protect the Data Lake with comprehensive security is clear. As large and growing volumes of diverse data are stored in the Data Lake, it comes to hold the crown jewels of your company—the vital and often highly sensitive data that has shaped and driven your business over a long history. However, the external ecosystem of data and operational systems feeding the Data Lake is highly dynamic and can introduce new security threats on a regular basis. Users across multiple business units can access the Data Lake freely and refine, explore and enrich its data at will, using methods of their own choosing, thereby increasing risks of exposure to unauthorized users. Any internal or external breach of this enterprise-wide data can be catastrophic, from privacy violations, to regulatory infractions, to damage to corporate image and long-term shareholder value. To prevent damage to the company's business, customers, finances and reputation, IT leaders must ensure that their Data Lake meets the same high standards of security as any legacy data environment.

Only as secure as the weakest link

Piecemeal protections are no more effective for a Data Lake than they would be in a traditional repository. There's no point in securing the primary access path to the data lake when a user can simply access the same data through a different path.

Hortonworks firmly believes that effective Hadoop security depends on a holistic approach. Our framework for comprehensive security revolves around five pillars: administration, authentication/ perimeter security, authorization, audit and data protection.

Five pillars of enterprise security

Administration Central management and consistent security	How do I set policy across the entire cluster?
Authentication Authenticate users and systems	Who am I/ prove it?
Authorization Provision access to data	What can I do?
Audit Maintain a record of data access	What did I do?
Data Protection Protect data at rest and in motion	How can I encrypt data at rest and over the wire?

Figure 1: Requirements for enterprise-grade security

Security administrators must address questions and provide enterprise-grade coverage across each of these pillars as they design the infrastructure to secure data in Hadoop. If any of these pillars remains weak, it introduces threat vectors to the entire data lake. In this light, your Hadoop security strategy must address all five pillars, with a consistent implementation approach to ensure their effectiveness.

Needless to say, you can't achieve comprehensive protection across the Hadoop stack through an ad-hoc approach. Security must be an integral part of the platform on which your Data Lake is built with a combination of bottom-up and top down approach. This makes it possible to enforce and manage security across the stack through a central point of administration and prevent gaps and inconsistencies. This approach is especially important for Hadoop implementations where new applications or data engines are always on the horizon in the form of new Open Source projects, a dynamic scenario that can quickly exacerbate any vulnerabilities.

Hortonworks helps customers maintain the high levels of protection their enterprise data demands by building centralized security administration and management into the DNA of the Hortonworks Data Platform (HDP). HDP provides an enterprise-ready data platform with rich capabilities spanning security, governance and operations. By implementing security at the platform level, Hortonworks ensures that security is consistently administered to any application built on top of the data platform, and makes it easier to build or retire data application without impacting security.

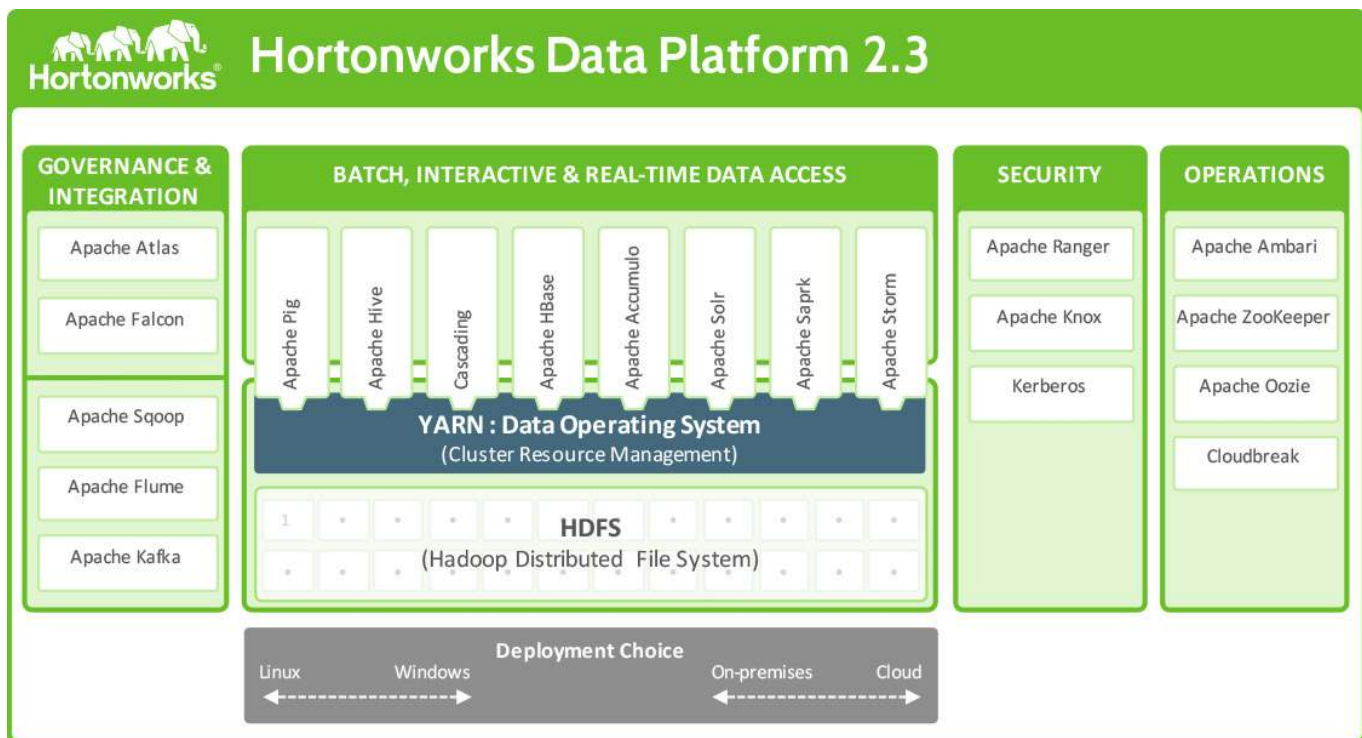


Figure 2: Hortonworks Data Platform 2.3

Commitment to enterprise-readiness

Hortonworks was founded with the objective to make Hadoop ready for the enterprise and has a strong legacy of significant contributions in this area. This goal of enterprise-readiness led the original Hadoop team at Yahoo! to develop Kerberos as the basis for strong authentication in Hadoop. Since that time, Hortonworks has continued to make significant investments in security. In May 2014, Hortonworks acquired XA Secure, a leading data security company, to accelerate the delivery of a comprehensive approach to Hadoop security. To be consistent with its mission to develop, distribute and support 100% open source Apache Hadoop data platform, Hortonworks immediately incorporated the XA Secure technology into the Hortonworks Data Platform (HDP), while also converting the commercial solution into an open Apache community project called Apache Ranger.

As part of HDP, Hortonworks features comprehensive security that spans across the five security pillars. With this platform approach, HDP enables IT to meet the requirements of Hadoop security better than any other vendor.

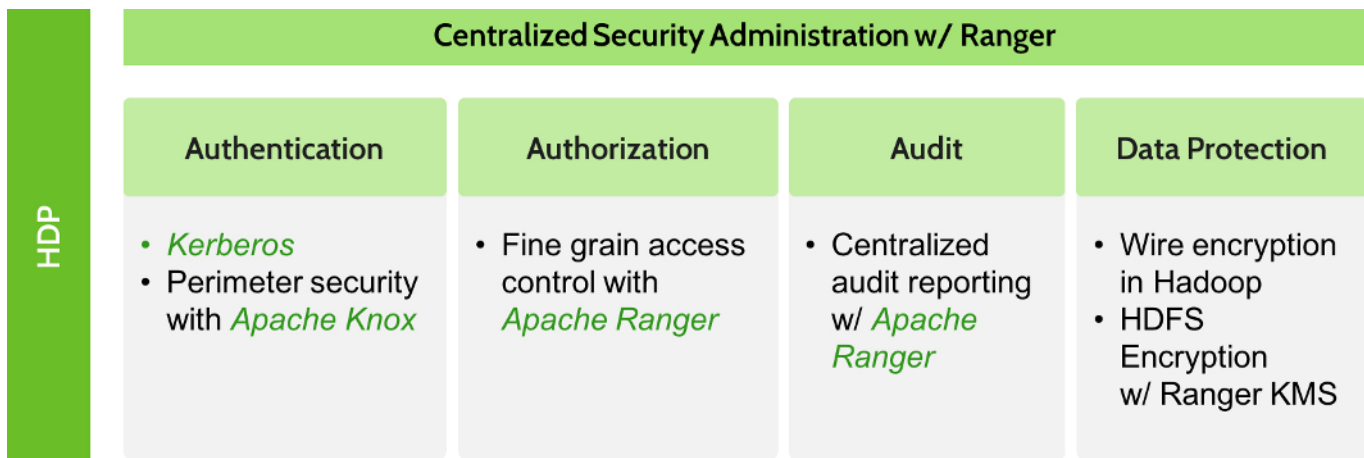


Figure 3: Comprehensive security in HDP

Administration

In order to deliver consistent security administration and management, Hadoop administrators require a centralized user interface—a single pane of glass that can be used to define, administer and manage security policies consistently across all the components of the Hadoop stack. Hortonworks addressed this requirement through Apache Ranger, an integral part of HDP, which provides a central point of administration for the other four functional pillars of Hadoop security.

Ranger enhances the productivity of security administrators and reduces potential errors by empowering them to define security policy once and apply it to all the applicable components across the Hadoop stack from a central location.

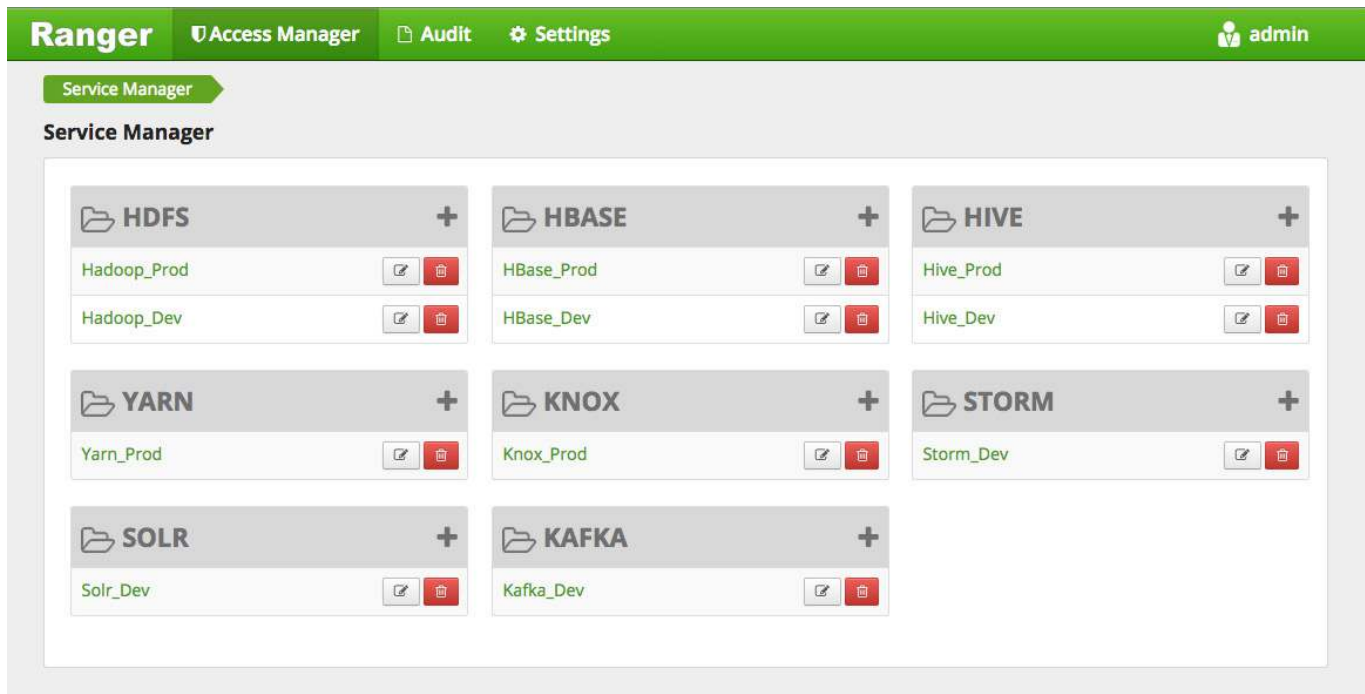


Figure4: Apache Ranger provides a "single pane of glass" for the security administrator

Other solutions for Hadoop enterprise security only offer partial administration across authentication, authorization, auditing and data protection/encryption—and lack the centralized administration and management needed for efficient and comprehensive security.

	APACHE RANGER
Centralized security administration	<ul style="list-style-type: none"> ● Apache Ranger provides a centralized platform for security policy administration

Authentication and perimeter security

Establishing user identity with strong authentication is the basis for secure access in Hadoop. Users need to reliably identify themselves and then have that identity propagated throughout the Hadoop cluster to access resources such as files and directories, and to perform tasks such as running MapReduce jobs. Hortonworks uses Kerberos, an industry standard, to authenticate users and resources within Hadoop cluster. Hortonworks has also simplified Kerberos setup, configuration and maintenance through Ambari 2.0.

Apache Knox Gateway ensures perimeter security for Hortonworks customers. With Knox, enterprises can confidently extend the Hadoop REST API to new users without Kerberos complexities, while also maintaining compliance with enterprise security policies. Knox provides a central gateway for Hadoop REST APIs that have varying degrees of authorization, authentication, SSL and SSO capabilities to enable a single access point for Hadoop.

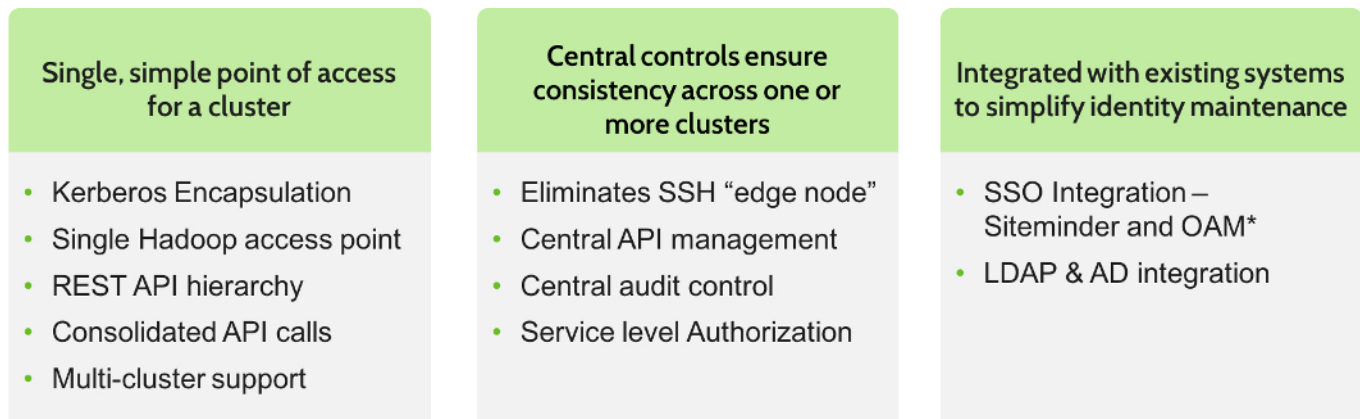


Figure 5: Perimeter security with Apache Knox

Other vendors fail to provide a comprehensive solution in this area, instead positioning Kerberos for perimeter security. Kerberos is an essential step for user authentication, but it is not sufficient in itself as it lacks the ability to hide cluster entry points and block access at the perimeter. By comparison, Apache Knox was built as a secure API gateway for Hadoop, with the ability to block services at the perimeter of the cluster. When using Apache Knox for REST APIs, cluster’s multiple access points are hidden from end users, adding another layer of protection for perimeter security.

Apache Knox is a pluggable framework, and a new REST API service can be added easily using a configurable services definition (Knox Stacks).

	KERBEROS
Kerberos-based authentication	<ul style="list-style-type: none"> ● Ambari simplifies the setup, configuration and maintenance of Kerberos ● Ambari includes support for Apache Ranger installation and configuration
	APACHE KNOX
Perimeter security	<ul style="list-style-type: none"> ● Provide security to all of Hadoop’s REST and HTTP services

Authorization

Ranger manages fine-grained access control through a rich user interface that ensures consistent policy administration across Hadoop data access components. Security administrators have the flexibility to define security policies for a database, table and column or a file, and administer permissions for specific LDAP based groups or individual users. Rules based on dynamic conditions such as time or geography can also be added to an existing policy rule. The Ranger authorization model is highly pluggable and can be easily extended to any data source using a service-based definition.

Administrators can use Ranger to define centralized security policy for the following components:

- Apache Hadoop HDFS
- Apache Hadoop YARN
- Apache Hive
- Apache HBase
- Apache Storm
- Apache Knox
- Apache Solr
- Apache Kafka

Ranger works with standard authorization APIs in each Hadoop component and is able to enforce centrally administered policies for any method of accessing the data lake.

The screenshot displays the Apache Ranger web interface for editing a policy. The browser address bar shows the URL: 127.0.0.1:8080/index.html#/service/2/policies/18/edit. The interface includes a navigation bar with 'Ranger', 'Access Manager', 'Audit', and 'Settings', and a user profile for 'admin'. The main content area is titled 'Edit Policy' and contains the following sections:

- Policy Details:**
 - Policy ID: 18
 - Policy Name: Call_Details_Table (Status: enabled)
 - Hive Database: xademo (Status: include)
 - table: call_detail_records (Status: include)
 - Hive Column: phone_number (Status: exclude)
 - Description: (Empty text area)
 - Audit Logging: YES
- User and Group Permissions:**

Permissions	Select Group	Select User	Permissions	Delegate Admin
	x developer	Select User	select	<input type="checkbox"/>
	+			

Figure 6: Fine-grained security policy definition with Apache Ranger

Solutions from other vendors lack the flexibility and rich user interface to enable administrators configure security policy for specific groups and individual users.

In contrast, Ranger provides administrators with deep visibility into the security administration process that is required for auditing purposes. The combination of Ranger's rich user interface with deep audit visibility makes it highly intuitive to use, enhancing productivity for security administrators.

The screenshot shows the Apache Ranger web interface. The browser address bar displays the URL `127.0.0.1:6080/index.html#/service/2/policies`. The interface has a green header with the 'Ranger' logo and navigation tabs for 'Access Manager', 'Audit', and 'Settings'. The user is logged in as 'admin'. The main content area shows the 'List of Policies : sandbox_hive' with a search bar and an 'Add New Policy' button. Below is a table of policies:

Policy ID	Policy Name	Status	Audit Logging	Groups	Users	Action
3	sandbox_hive-1-20150529142947	Enabled	Enabled	--	xapolicymgr	[Edit] [Delete]
4	Hive Global Tables Allow	Disabled	Enabled	public	--	[Edit] [Delete]
5	Hive Global UDF Allow	Disabled	Enabled	public	--	[Edit] [Delete]
18	Call_Details_Table	Enabled	Enabled	developer	--	[Edit] [Delete]
19	Customer_Details_Table	Disabled	Enabled	Marketing	--	[Edit] [Delete]
20	Hive Demo Table Loader	Enabled	Enabled	--	hive	[Edit] [Delete]
21	Hive Demo UDF Loader	Enabled	Enabled	--	hive	[Edit] [Delete]
29	admin policy	Enabled	Enabled	--	admin	[Edit] [Delete]

Figure 7: With Apache Ranger, administrators have complete visibility into the security administration process

	APACHE RANGER
Platform-wide coverage across Hadoop stack	<ul style="list-style-type: none"> Coverage across HDFS, YARN, Hive, HBase, Storm, Knox, Solr and Kafka
Fine grain authorization	<ul style="list-style-type: none"> Authorize security policies for a database, table and column or a file as well as LDAP based groups or individual user
Provide hooks for dynamic policy-based authorization	<ul style="list-style-type: none"> Specify dynamic conditions in service definitions Flexibility to define unique conditions by service (HDFS, Hive etc.)
Built on pluggable service-based model	<ul style="list-style-type: none"> Custom plugins can be created for any data store

Audit

As customers deploy Hadoop into corporate data and processing environments, metadata and data governance must be vital parts of any enterprise-ready data lake. For these reasons, Hortonworks established the [Data Governance Initiative \(DGI\)](#) with [Aetna](#), [Merck](#), [Target](#) and [SAS](#) to introduce a common approach to Hadoop data governance into the open source community. This initiative has since evolved into a new open source project called Apache Atlas. Apache Atlas is a set of core foundational governance services that enables enterprises to effectively and efficiently meet their compliance requirements within Hadoop and allows integration with the complete enterprise data ecosystem. These services include:

- [Search and lineage for datasets](#)
- [Metadata-driven data access control](#)
- [Indexed and searchable centralized auditing operational events](#)
- [Data lifecycle management from ingestion to disposition](#)
- [Metadata interchange with other tools](#)

Ranger also provides a centralized framework for collecting access audit history and easily reporting on this data, including the ability to filter data based on various parameters. Together with Apache Atlas, this makes it possible for users to gain a comprehensive view of data lineage and access audit, with an ability to query and filter audit based on data classification, users or groups, and other filters.

	APACHE ATLAS AND APACHE RANGER	
Data lineage	●	Reporting by entity type or instance
Consolidated audit	●	Ranger provides security audit which can be combined with data lineage in Atlas to provide a comprehensive view
Metadata services	●	Open extensible system with a policy rules engine
Third party support	●	HDP fosters a rich ecosystem of 3rd party vendors

Data protection

Data protection adds a robust layer of security by making data unreadable in transit over the network or at rest on a disk. HDP supports encrypting network traffic as data into and through the Hadoop cluster over RPC, HTTP, Data Transfer Protocol (DTP) and JDBC. Network traffic over each of these protocols can be encrypted to provide privacy for data movement.

HDP also satisfies enterprise requirements for security and compliance through encryption for data at rest. HDP supports the ability to encrypt files stored in Hadoop, including a Ranger-embedded open source Hadoop key management store (KMS). Ranger provides security administrators with the ability to manage keys and authorization policies for KMS. Hortonworks is also working extensively with its encryption partners to integrate HDFS encryption with enterprise grade key management frameworks. With HDP, our customers have the flexibility to leverage open source key management store (KMS) or use enterprise wide KMS solutions provided by the partner ecosystem.

Encryption in HDFS, combined with KMS access policies maintained by Ranger, prevents rogue Linux or Hadoop administrators from accessing data and supports segregation of duties for both data access and encryption.

	HDP TDE (HDFS) + APACHE RANGER (KMS)	
Open source	●	100% open source and open community
Partner ecosystem	●	Flexibility to use open sources KMS or existing enterprise solutions
	●	Hortonworks is currently working with select encryption partners to integrate HDFS encryption with enterprise grade KMS solutions

Enabling extensibility through a pluggable framework

Apache Ranger and Knox provide a service-based architecture for policy administration and enforcement that enables customers to easily support new components or data engines. Applications that integrate with Ranger can simply use the pluggable architecture to leverage the existing security policies in Ranger without the need to redefine the security policy from scratch. Users can also create custom services as plug-ins to any data store and build and manage services centrally for their big data BI applications. Similarly, REST APIs can easily be integrated with Apache Knox using a service-based configuration. The configurable framework also enables partner solutions to easily work with Ranger and Knox to extend a new service to support a new component or data engine.

Comparison across security pillars

Point by point, pillar by pillar, HDP provides rich capabilities to protect sensitive data within your Hadoop Data Lake. In each aspect, Hortonworks meets or exceeds what's possible with competing solutions. The ground-up platform design of the Hortonworks security architecture ensures that each pillar complements the others to ensure that no gaps remain. The centralization of administration and management enables IT to take a holistic approach to security, as well as increasing IT productivity and efficiency. By taking a approach based on a pluggable and completely 100% open source design, Hortonworks enables ecosystem partners to easily embed additional security capabilities. In this way, companies can leverage comprehensive, easy to use security to their data and apply the same level of rigor to Hadoop as they would for traditional data environments—as it should be.

	HDP	OTHER DISTRIBUTIONS
Administration	●	◐
Authentication	●	◐
Perimeter Security	●	○
Authorization	●	◐
Audit	◐	◐
Data Protection	◐	◐

Figure 8: Competitive analysis across security pillars

Summary

No business can afford to have Big Data insight come at the expense of enterprise security. As you plan your Hadoop strategy, make sure that the platform you choose provides a comprehensive and holistic approach to protect your Data Lake and the valuable information it contains. With Hortonworks, companies can implement a platform with all five pillars of Hadoop security woven into its architecture—administration, authentication/perimeter security, authorization, audit and data protection. Unlike other Hadoop platforms that offer only partial security functionality, HDP is all IT needs to support a complete strategy for secure Big Data—making its competitive benefits available to business users throughout the organization without exposing the company to risk.

About Hortonworks

Hortonworks develops, distributes and supports the only 100% open source Apache Hadoop data platform. Our team comprises the largest contingent of builders and architects within the Hadoop ecosystem who represent and lead the broader enterprise requirements within these communities. Hortonworks Data Platform deeply integrates with existing IT investments upon which enterprises can build and deploy Hadoop-based applications. Hortonworks has deep relationships with the key strategic data center partners that enable our customers to unlock the broadest opportunities from Hadoop. For more information, visit www.hortonworks.com.