NUCYPHER

NuCypher Hadoop: Delegated Access Control and Encryption Management System

HIGH-PERFORMANCE ACCESS MANAGEMENT AND DATA PROTECTION FOR HADOOP DATA LAKES

NuCypher Hadoop leads the advancement of big data network security with an innovative re-encryption-based data access control system for Hadoop. Purpose-built to enable the secure sharing of sensitive data in compliance with governance and residency requirements, NuCypher Hadoop arms security executives with new capabilities to protect sensitive, globally distributed corporate data lakes from ever-evolving global security threats.

NuCypher Hadoop is a cost-effective data encryption solution that allows organizations to deliver scalable encryption and fine-grained access control without sacrificing speed, security, or data usability. NuCypher Hadoop enables enterprises to leverage Hadoop across globally distributed infrastructure and clouds to deliver better analysis faster and to maximize the value of their data – without increasing risk or compromising security, privacy, and regulatory compliance mandates

Removes Major Barriers To Data Sharing in Hadoop and Cloud Deployments

Companies are adopting Hadoop as a lower cost, higher performance alternative to proprietary data warehousing and are now using Hadoop data lakes for storing both structured and unstructured data, which quite often contains sensitive information. Data ingested and residing in the data lake needs to be accessed by different teams, including business analysts, data scientists, developers, and others in the enterprise and also needs to be shared with partners and external customers. It's critical that organizations maintain a separation of duties between administrative and authorized user access; and prevent unauthorized access to meet privacy, security, and compliance requirements. Enforcing existing enterprise security models in this new infrastructure has proven difficult because Hadoop was not designed with security in mind and the current set of data encryption and tokenization products require unacceptable trade-offs between security, performance, and cost.

NuCypher Hadoop enforces role-based access control for fine-grained permissions management, which restricts users and services to only the data they are allowed to access, enabling precise control of datasets down to the field level. Sensitive data in the cluster is secured at rest as well as in motion using encryption, which protects data in the cluster from unauthorized visibility. Unlike masking, tokenization, or less sophisticated encryption products, which create operational bottlenecks and limit performance, data availability, and usability, or fail to protect all data elements, NuCypher Hadoop allows organizations to encrypt entire data lakes without latency, downtime, or meaningful overhead.

NuCypher Hadoop uses a proxy re-encryption scheme and premise-based key management system, which provides military strength encryption and fine-grained access control to protect data on-premise, in the cloud, or in hybrid environments - in compliance with the most demanding governance standards, laws, and regulations. NuCypher Hadoop enables shared data lake deployments by ensuring that data can never be accessed by cloud providers or attackers without permission.

NUCYPHER

Key Advantages:

- Quick time-to-value. Hadoop distribution agnostic, easy installation, drop-in replacement for transparent data encryption (TDE)
- Zero downtime rekeying your Hadoop platform. No more SLA disruptions or late night downtime!
- Inherits existing roles and policies - Easily and consistently manage policies.
- Easy to use key and encryption management – no more complex, confusing management tools and processes.
- Encryption policy that follows the workload. Easy migration from on-premise to cloud and from cloud to cloud.
- Enables data de-identification in test, development, and production – Keeps business agile and flexible.

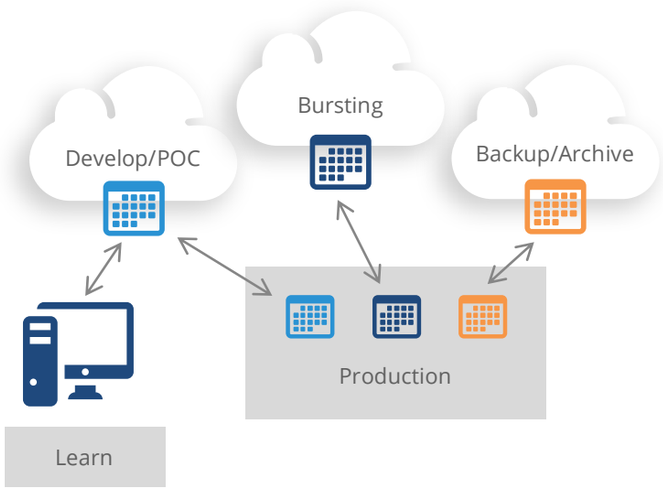

Extend Hadoop to more users. Securely store sensitive data with non-sensitive data in Hadoop and share it with users across all departments and organizations through multi-tenant authorization and fine-grained access controls.

How It Works

NuCypher Hadoop uses a breakthrough encryption scheme developed to secure distributed storage among many users. In a single-user cryptographic file system, access control is straightforward. The user creates all the keys protecting content and there is no key distribution problem. With group sharing in cryptographic storage, group members must communicate with content owners to obtain decryption keys for accessing files. This creates operational barriers and security risks that render this method unusable in Hadoop environments which need to share access among many users. NuCypher Hadoop provides an encryption layer with access controls embedded directly into the encryption layer. As a result, decryption permissions can be managed without access to the secret keys, allowing you to protect data in a safe and compliant manner. Privacy, confidentiality and data integrity is maintained by encrypting the contents of stored files and allows these encrypted files to be stored on outsourced servers.

NuCypher Hadoop supports granular encryption, allowing fine-grained decryption permissions, combined with delegated access both internally (e.g. employees, roles) and externally (e.g. partners, customers) in multi-tenant cloud environments and centralized data lakes. And it works seamlessly across all your environments—on-premise, hybrid, and public cloud—while eliminating single points of trust and enabling enterprises to leverage the elasticity of the public cloud while maintaining on-premise key management.

NuCypher Hadoop is a cryptographically enforced access control system based on proxy re-encryption where data remains encrypted in the filesystem until it is needed for processing. Data is stored encrypted in the cluster and the owner can delegate access rights to the computing cluster for processing. This proxy re-encryption scheme uses a public-key encryption scheme that permits a proxy to transform ciphertexts. The proxy needs a re-encryption key generated by the delegating entity and proxy re-encryption then enables delegation of decryption rights. In this scheme there are three entities: The data owner, with public and private keys, the delegatee, with public and private keys, and the proxy entity, which permits access through re-encryption. The result is that stored data is always encrypted and encryption keys do not need to be shared between different data sources. And the use of proxy re-encryption allows for the delegation of access to the stored data to the computing cluster.

NUCYPHER

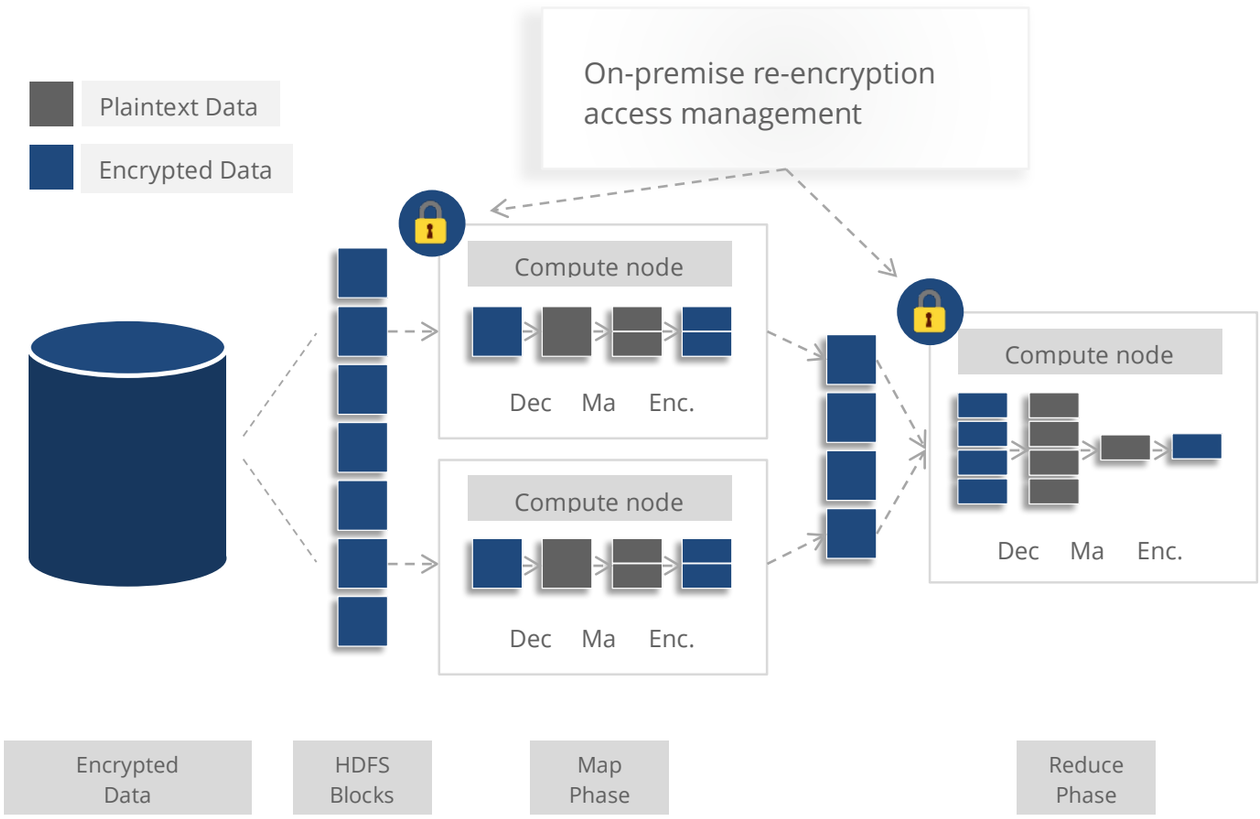

NuCypher Hadoop uniformly enforces fine-grained authorization policies for users based on their role, regardless of how they're accessing data. You create the policy once and then rest assured that your users access only the data for which they have permissions, without unnecessary repetition or policy mirroring. Set granular permission with the flexibility to control access from block-level down to rows and fields. NuCypher Hadoop covers everything and can protect broadly with coarse-grained block, file or volume encryption as well as provide fine-grained access to share documents, rows, columns and fields in semi-structured documents. Access policies are inherited from the Hadoop management platform and tied to LDAP and Kerberos to maintain real-time user provisioning and de-provisioning synchronization.

Enterprise Scale Big Data Protection

Securing big data is different. This is because the massive attack surface of big data stores makes them highly vulnerable to unauthorized intrusion. At the same time, much of the encryption technology that exists today was not designed for deployment in distributed computing architectures such as Hadoop, which consist of multiple servers that are networked together into server clusters. Distributed systems also require automated mechanisms for secure node removal when a server is removed from a cluster, encryption for both data-at-rest and in-motion, as well as rapid and secure encryption key rotation. All of these functions must perform efficiently without requiring the re-encryption of any files or downtime during normal operation.

Key Features

Breakthrough Multi-Layer Security

NuCypher Hadoop supports file and block encryption to secure data at-rest and in-transit. And to ensure that data is protected wherever it goes yet remains flexible enough to be used, fine-grained field/column protection can be applied. Fine-grained protection continues while data is in-use and during analytical processing. This enables you to block access to sensitive information while allowing access to the necessary business intelligence, empowering users and processes that require only part of the original data to function.

NUCYPHER

Policy-Based Encryption for Total Control

NuCypher Hadoop's comprehensive encryption capabilities provide access to raw data safely. Masking and tokenization require you to know where your data is and often misses sensitive data such as PII and credit card numbers. NuCypher Hadoop covers everything and can protect broadly with coarse-grained file or block encryption as well as provide fine-grained access to share documents, rows, columns, and fields in semi-structured documents. It allows you to set access based on policies: who can use the data, where they can use it, and when they can use it. By sharing data with granular permissions, third parties or other departments within an organization can only decrypt the specific slices of data (columns, fields), which they are allowed to decrypt.

Block Cloud Providers Access to Data

NuCypher Hadoop helps defend and protect your most important assets in the cloud – your data lakes. For most organizations moving to the cloud is complicated by data governance and control issues. In addition to the need to protect data from being hacked, organizations need assurance the even their cloud provider can't access data that is stored on their infrastructure or be forced to turn over data to external parties under subpoena or the threat of legal action. NuCypher Hadoop ensures that data remains encrypted at-rest and in-transit, where it can't be accessed by anyone under any conditions that aren't authorized. On-premise key management means that nobody can access your data without your consent – nobody!

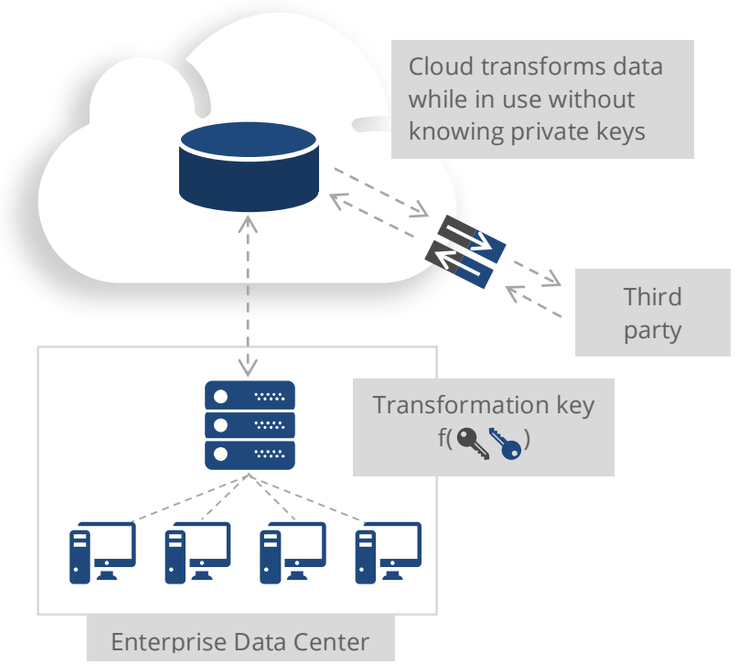

Maintain Data Sovereignty

NuCypher Hadoop provides boundary control that gives you the ability control where your data resides, who can access it, and from where. You can define access controls based on geographic area, sets of nodes meeting certain compliance or regulatory requirements or a number of other factors. You control where your data resides to easily comply with data sovereignty requirements protect against insider threats and external breaches.

Broad and Deep Data Encryption for Shared Data Lakes

NuCypher Hadoop enables state of the art proxy re-encryption with both coarse and fine-grained data encryption options and key management. Control access within a pooled resource of data and compute either internally between departments/groups of one organization or externally among multiple organizations. Deployment is fast and easy with intuitive GUI-based installation and management options. Encryption key rotation can be performed with zero downtime to ensure that users always have access to the data they need.

Automated and Simplified Key Management That's Blazingly Fast

When it comes to securing big data in the enterprise, existing encryption solutions are insufficient in many ways. Hadoop's native open-source transparent data encryption (TDE) lacks basic key rotation capabilities and forces manual, time-consuming workarounds, making it difficult to fit into existing enterprise security policies and extremely difficult to comply with standards that require key rotation such as PCI. Attempts to perform key rotation often result in leaving encryption keys exposed, which is a big problem. Most legacy

encryption products were designed for static relational databases and lack the ability to scale and perform well in highly dynamic, petabyte-scale, distributed computing architectures. Very few are optimized for big data platforms and distributed file systems like HDFS, especially in cloud environments.

NuCypher Hadoop is designed to enforce premise-based key management, so that keys never leave your control. NuCypher Hadoop is able to support encryption for unlimited clusters across multiple physical locations and private, public, or hybrid clouds. Cloud providers receive only encrypted data, while encryption keys always stay on-premise, greatly reducing the risk of a data breach and providing enterprises access to the scalability and elasticity of the cloud.

NuCypher Hadoop provides policy-based key management so administrators do not need to be security experts. Key management is easily integrated into the Hadoop platform dashboard through APIs, and keys are managed through policy settings, making the system easy to use. NuCypher Hadoop offers ultra-efficient key rotation where administrators can set policies to re-key in accordance with industry regulations like PCI or with company guidelines. This can be done automatically on a fixed schedule or on an ad-hoc basis as needed. This is especially significant for organizations with lots of data and many encryption keys, and for companies that must meet SLAs for application uptime.

NuCypher Hadoop executes secure key rotation in seconds instead of hours and, unlike TDE, it doesn't require shutting down encryption zones and suffering through long downtimes.

Adaptable to Fit Existing Hadoop Environments

NuCypher Hadoop offers broad platform support, is Hadoop distribution agnostic and can be used with Apache Hadoop, Cloudera and Hortonworks distributions, as well as Amazon's EMR. It's a drop-in replacement for Transparent Data Encryption (TDE) with no change to user experience. Installation is easy to perform and modifies the KMS using the same cluster hardware in your (pre-existing) Hadoop deployment. Implementation is rapid and transparent because it doesn't require changes to applications, the underlying file system or hardware infrastructure. Because NuCypher Hadoop performs encryption and decryption at the HDFS layer it's transparent to users, applications, databases, and storage subsystems. NuCypher Hadoop doesn't require coding or modification to applications or databases. You can easily customize the solution to meet your exact requirements and set policy for fine-grained protection of structured and unstructured data with complete end-to-end protection within a common management interface.

> "Businesses have traditionally managed data within structured and unstructured silos, driven by inherent requirements to deploy relational database management systems, file storage systems and unstructured file shares. However, **the advent of big data and cloud storage environments is transforming the way in which data is stored, accessed, and processed.** Access to public cloud services and infrastructure further complicates this process due to the potential access by cloud service providers and security vendors - **CISOs need to develop a data-centric security approach."**
>
> Brian Lowans, Principal Research Analyst at Gartner

Protection that Addresses Compliance Requirements

NuCypher Hadoop provides a proven defense in regulated industries such as healthcare, financial services, and retail from the accelerating frequency and scope of data breaches. NuCypher Hadoop easily integrates into your security and compliance policies, and can bring Hadoop data stores into compliance with corporate and regulatory data protection initiatives, including GDPR, HIPAA, HITECH, and PCI.

NUCYPHER

Unparalleled Performance and Scalability With Zero Impact On Analytics

When Hadoop uses TDE, the key management server (KMS) is a bottleneck that adds a significant drag on performance because thousands of nodes talk to your KMS to decrypt your data. Additionally, the KMS must be close to the cluster, so that you don't suffer from latency problems. NuCypher Hadoop avoids this issue because operations are delegated to the cluster and performed locally. With NuCypher Hadoop, processing doesn't require communication between compute nodes and the KMS. This means the KMS doesn't have to stay online during a job, removing latency bottlenecks and slow requests over the network. This capability provides scalability for large and complex environments.

Strategic Partnerships

NuCypher Hadoop technology provides proven interoperability between enterprise key managers, cryptographic devices and a range of storage, security and cloud products. It has been tested for interoperability with our strategic partners, including:

- Hortonworks
- Cloudera
- MapR
- Amazon EMR
- Gemalto Safenet
- Thales