XCrypt Hadoop™

Granular, Next-Generation Data Encryption with Integrity Protection for HDFS data

Fast, Scalable, Affordable

- **Designed to meet the unique data protection requirements of high-volume,** HDFS data stores, in the cloud and on-premises
- **All-software solution simplifies deployment**, easily scales, and delivers exceptional price/performance
- **Highly-granular**, user-defined policy management, with support for a unique key per file
- **Provides ultra-secure authenticated encryption** using associated data (AEAD) to protect HDFS ciphertext and access control lists (ACLs) from unauthorized modification
- **Protects** against stealthy and highly-damaging chosen-ciphertext attacks (CCAs)

XCrypt Hadoop from Zettaset is a true next-generation encryption solution that combines Authenticated Encryption with Associated Data (AEAD) and Galois/Counter mode (GCM) to not only encrypt data, but also provide integrity protection for the encrypted data. XCrypt Hadoop can detect and mitigate the exposure associated with cyberattacks that manipulate or delete encrypted data. This advanced capability ensures the accuracy and consistency of data over its entire life-cycle.

XCrypt Hadoop is designed for selective HDFS encryption down to the file-level. Encrypting and decrypting data at the file-level opens up the possibilities of unauthorized access, and calls for greater levels of data protection. For that reason, XCrypt Hadoop provides additional protection in two unique ways: (1) Authenticated encryption and protection for cyphertext using associated data (AEAD), and (2) Cryptographic protection for access control lists (ACLs).

Why Protection for Data That's Already Encrypted?

There is a common misperception that encrypted data is fully protected, but even data which has been encrypted is exposed to malicious attacks and unauthorized modification. File-based encryption protects data on nodes against attackers reading files, but still is still vulnerable to write attacks on those same encrypted files (ciphertext).

If attacker can write to the ciphertext, he/she can either (1) erase data without detection or (2) mount a chosen ciphertext attack (CCA) to try to obtain the data key. Data-in-motion is an even easier target since an attacker can simply modify ciphertext by performing a man-in-the-middle attack.

XCrypt Hadoop addresses these vulnerabilities and includes unique features not found in any other HDFS data encryption solution.

Next-Generation Encryption for Added Ciphertext Integrity Protection

XCrypt Hadoop is a true, next-generation encryption solution that takes an extra step, providing protection for encrypted data (ciphertext). XCrypt Hadoop uses AEAD architecture (vs. Encryption + MAC) to protect encrypted data from ciphertext modification. This approach ensures that encrypted data is verifiably secure. AEAD enables encryption and authentication to happen concurrently, making it easier to use and optimize than older, commonly-used modes such as CCM.

Authenticated encryption using associated data has additional advantages as part of the XCrypt Hadoop encryption solution.

- **AEAD is verifiably secure**, while Encrypt + MAC is not
- **AEAD using Galois/Counter mode (GCM)** protects against an intruder writing to files at-rest by way of a chosen-ciphertext attack (CCA)
- **AEAD improves efficiencies** because of its ability to concurrently perform encryption and authentication

Some otherwise secure encryption schemes, including non-authenticated encryption modes, can allow attackers to discover the encryption keys using a CCA. Non-authenticated encryption only prevents an attacker from reading the plaintext. It does not prevent an attacker from modifying the ciphertext.

The authenticated encryption mode used by XCrypt Hadoop is able to prove that the ciphertext was made by someone who was authorized to possess the encryption. Existing non-authenticated encryption products depend on the user to detect any data modification by noticing plaintext that appears to be wrong, an unreliable approach that exposes organizations to unnecessary risk.

High Performance Encryption with Galois/ Counter mode (GCM)

Authenticated encryption will become mandatory due to its more stringent security properties. However, it must not compromise performance and scalability. To ensure optimal performance levels, Zettaset XCrypt Hadoop uses the Galois/ Counter mode (GCM) for authenticated encryption. GCM has been identified as the only encryption mode that addresses requirements for high data-rate authenticated encryption. GCM mode encryption can efficiently achieve speeds of up to 10+ gigabits per second, and can be pipelined and parallelized.

Cryptographic Protection for Access Control Lists (ACLs)

Zettaset XCrypt Hadoop cryptographically protects access control lists and prevents an attacker from modifying ACLs and using those changes to gain access to data. Application of ACLs at every layer of access for data is critical to secure a system. An ACL is a list of permissions attached to an object that specifies which users or system processes are granted access to objects and are typically applied to data to restrict access to data to approved entities.

XCrypt Hadoop enforces file access according to policies established by the administrator. Policies can enforce access control as well as encryption. The solution has been designed for interoperability in existing IT infrastructure, and works with Active Directory, LDAP, Kerberos, and UNIX authentication mechanisms. XCrypt Hadoop co-exists with and enhances the Apache Ranger offering by adding ACL integrity to HDFS.

Protection for HDFS Extended Attributes (Xattrs)

Extended attributes (Xattrs) provide a storage place for tags and IVs. Zettaset XCrypt Hadoop cryptographically ties Xattrs to associated files for additional protection against malicious attacks. XCrypt Hadoop uses Xattrs to (1) keep track of authentication tags, and (2) generate unique initialization vectors (IVs) for each block

Automated, Granular Key Management

XCrypt Hadoop's key management is system is infinitely scalable and highly granular, and can support a unique key per file. A distributed policy server enables user-defined policy enforcement on a granular level, ensuring that a "lost" key has the potential for only minor impact. The automated policy server also ensures that your keys never go to clients.

Supports Strategic IT Compliance Initiatives

- **Provides a proven defense** for sensitive data in regulated industries such healthcare, financial services, and retail from the accelerating frequency and scope of data breaches
- **Supports corporate and regulatory compliance initiatives** including PCI/DSS, HIPAA, FISMA, and GDPR.

Simple to Deploy, Low TCO

- **No proprietary appliances required,** making XCrypt Hadoop non-disruptive and much more cost-effective compared to legacy approaches. This is especially valuable in highly elastic cloud environments, and provides power users with greater operational efficiencies.
- **Supports Java API and REST API** for ease of integration
- **Simplified installation** using CLI tools (Ansible, Puppet, Chef)

Fits into Existing Security Infrastructure

- **Adheres to OASIS encryption open standards**, and is compatible with KMIP-compliant key management systems and PKCS#11-compliant hardware security modules (HSM).
- **Included with the XCrypt Hadoop data encryption solution** is a software-based automated Virtual Enterprise Key Manager (V-EKM) and Virtual Hardware Security Module (V-HSM) to facilitate initial deployment.

Interoperability Certifications

- **Certified interoperability with key manager solutions** from MicroFocus, Fornetix, Gemalto, HyTrust, Thales and others
- **Certified interoperability with HSMs** from Utimaco, Gemalto, and others.

XCrypt Hadoop is just one of the advanced, industry-leading encryption solutions built on Zettaset's XCrypt Data Encryption Platform. Additional solutions include XCrypt Full Disk and XCrypt Object. For more information, please contact us at sales@zettaset.com

About Zettaset

Zettaset is a software-defined encryption solution that protects against data theft and can be transparently deployed across all physical and virtual environments. Its products are designed for medium to large enterprises that deal with sensitive information that would expose them, financially and reputationally, in the event of a breach. Unlike traditional solutions that are appliance-based, Zettaset is a cost-effective, software-only solution that is easy to deploy, does not impact performance, and scales with your business from on premise to the cloud.

Zettaset
465 Fairchild Drive, Suite 207, Mountain View, CA 94043 // USA: +1.650.314.7920 // Fax: +1.650.314.7950
sales@zettaset.com // www.zettaset.com