

Advanced Security

MarkLogic® is the most secure NoSQL database. MarkLogic has focused on having enterprise-grade security from the start, and has fine-grained, certified security that organizations require—providing a shield against today’s cyber threats. Out-of-the-box, MarkLogic provides Document Level Security, Element Level Security, Auditing, Support for External Authentication (LDAP and Kerberos), Compliance Archives, Encryption, and more.

For certain use cases, there is an Advanced Security option, which includes three additional capabilities:

- **External Key Management** – Provides support for external, third-party key management systems
- **Redaction** – Prevents leakage of sensitive information to unauthorized users when importing, exporting, or copying data in and out of MarkLogic
- **Compartment Security** – Additional security control to specify that a user must have all of the right roles to interact with a document rather than just one of the right roles

Overview of MarkLogic Security

MarkLogic has been in the business of protecting and securing data for over a decade. MarkLogic is the only NoSQL database that is Common Criteria certified, and one of only six database vendors with the certification. The Common Criteria for Information Technology Security Evaluation (or “Common Criteria”) is the driving force for the widest available mutual recognition of secure IT products worldwide.

MarkLogic is installed and operational on systems that require databases to meet extremely rigorous requirements. These requirements include stringent measures for access, authentication, management, audits, role separation, and system assurance. It is for this reason that MarkLogic is chosen to run the most demanding, mission-critical applications at the heart of large investment banks, major healthcare organizations, and classified government systems.

By default, MarkLogic uses a role based access control (RBAC) security model in which each user is assigned any number of roles, and these roles are associated with any number of privileges and permissions. Privileges govern the creation of documents and execution of functions (URI and execute privileges) and permissions govern what can be done with a document (read, insert, update, execute).

ROLE-BASED ACCESS CONTROL AT THE DOCUMENT LEVEL

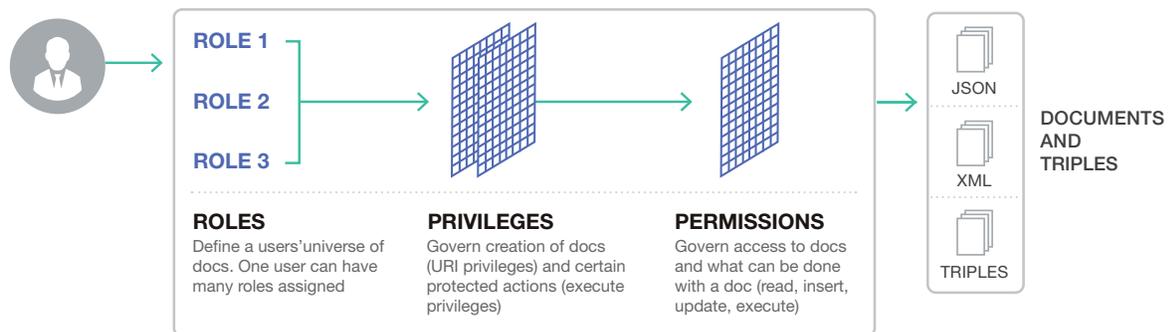


Figure 1: Each specific role contains both privileges (write and execute) and permissions (read and modify)

External Key Management

Encryption at Rest provides transparent encryption of databases, logs, configuration files and backup, with separate and powerful key management through a local Key Management System (KMS) or external KMS. A KMS, or “keystore,” is a secure location that stores the encryption keys. There are two important aspects of how MarkLogic encrypts data:

- **End-to-end encryption** – MarkLogic provides Encryption at Rest, which enables transparent and selective encryption of data residing on disk (locally or in the cloud) to ensure confidentiality and prevent information tampering of data residing on disk
- **Separation of duties** – Encryption at Rest significantly enhances data security controls by enforcing separation of duties. The system administrator with access to the host is not the same person who has control over the encryption keys and the encryption key lifecycle. This reduces the potential threat from insiders as well as hostile entities residing in the network such as Advanced Persistent Threats (APTs)

Another way to increase separation of duties and the overall security profile is to use an external KMS. By default, MarkLogic uses a local KMS. But, the best practice for Encryption at Rest is to employ an external, third-party KMS that is deployed and managed separately from the application servers. The external KMS securely stores authentication or encryption keys entrusted to it and provides them on demand to authorized systems. This provides an additional level of security by storing authentication keys separate from the storage system. Additionally, authentication keys are always handled and stored securely. The keys are never displayed in clear text. An external KMS enabled by the Advanced Security option provides:

- **Additional security** – An external KMS offers additional security for encryption keys, along with key management capabilities like automatic key rotation, key revocation, and key deletion
- **Additional separation of duties** – If an external KMS is used, then neither an unauthorized database admin, system admin, nor the storage admin can access the database files. The external KMS admin controls access to the encryption keys
- **Fast and easy integration** – MarkLogic interoperates with leading third party KMS systems, including Amazon Web Services (AWS) Key Management Service (KMS) and others such as Gemalto’s SafeNet.

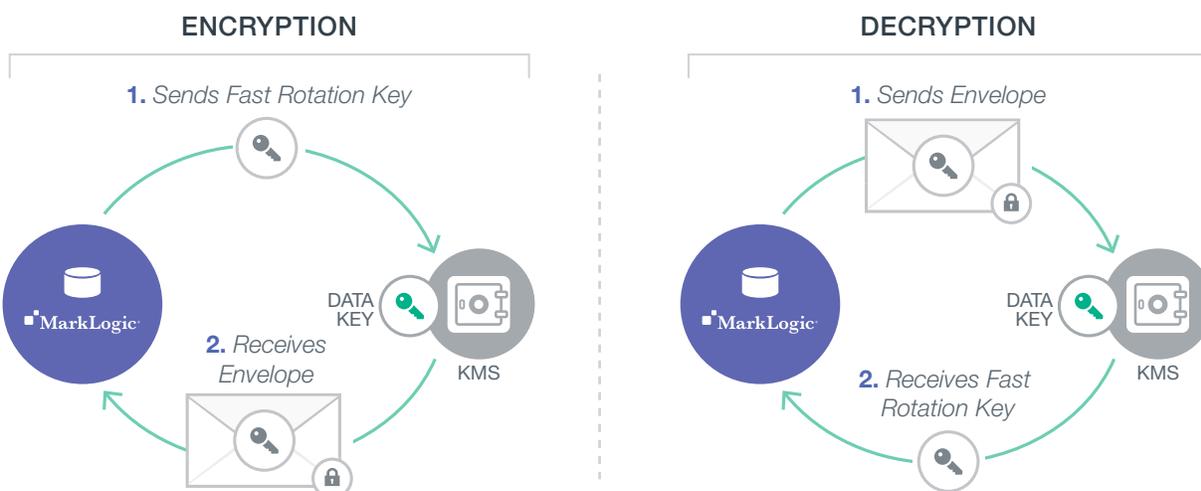


Figure 2: With Encryption at Rest, to access low level keys and read files, MarkLogic sends an envelope to the KMS, which then sends back the unencrypted key. If using an external KMS, MarkLogic has no access to envelope keys, which means no access to files, no ingestion, and no compromises.

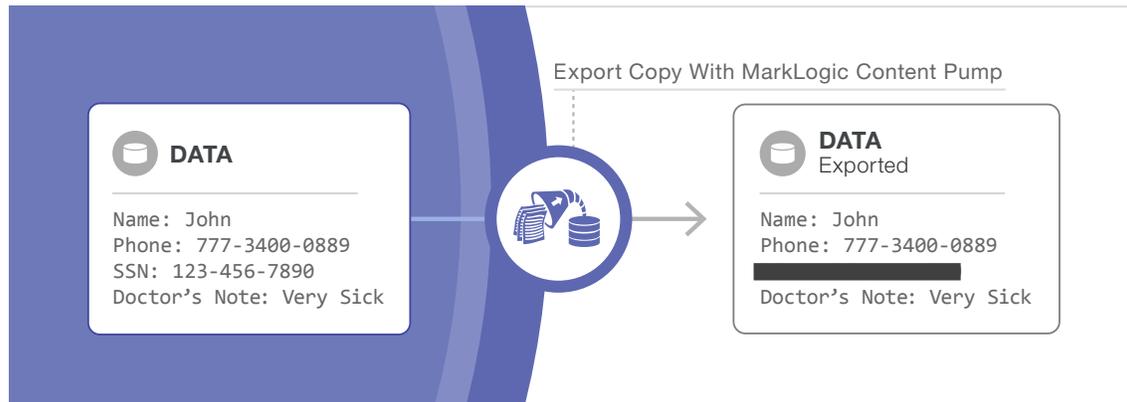


Figure 3: With Redaction, key information can be removed or masked

Redaction

Redaction helps prevent leakage of sensitive information to unauthorized users when importing, exporting, or copying data into and outside of MarkLogic. For example, Redaction is often required when providing data for analysis by data scientists, or when a developer needs production data but should not have access to real credit card data or personally identifiable information. Here are some key characteristics of Redaction:

- **Based on Rules and Policies** – To implement, a MarkLogic security administrator creates redaction policies that contain rules defining which sensitive information should be redacted, and then chooses which policy to apply when running an export. Administrators can combine built-in or custom rules into policies to match different target needs
- **Utilizes Built-in Functions** – Includes built-in functions for different types of redaction:
 - Concealing: Hide elements and/or their values (or properties and/or their values in the case of JSON)
 - Masking: Change the data using random masking (the value varies with each instance), deterministic masking (the same value is applied every time), or dictionary masking (the value is applied from a specified dictionary)
 - Patterns: Change the data using a pattern such as Social Security Number, U.S. phone number, email, IPv4, or Regex
 - Custom: Use server-side JavaScript or XQuery functions to apply unique rules (e.g., redact the name if the person is less than 18 years old)
- **Fully Auditable** – All rules and actions taken by users are logged, ensuring all export activity can be audited later on
- **Performs In Batch at Scale** – Redaction is designed to be used when running large bulk exports. And, by utilizing the MarkLogic Content Pump (mlcp), it is faster and more secure than solutions implemented at the application layer

Compartment Security

With Compartment Security it is possible to apply more complex role-based security rules for data access and updates. It is possible to specify that a user must have all of the right roles to access or create a document rather than just one of the right roles. When a role is compartmented, all privileges associated with a resource must be valid at the same time (AND semantics). However, when roles are not compartmented, then satisfying any privilege authorization condition will be sufficient (OR semantics).

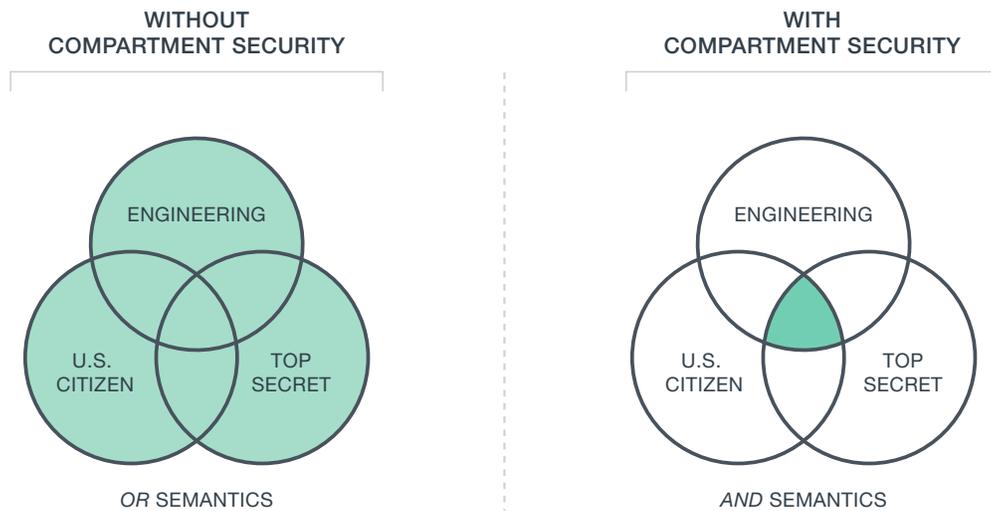


Figure 4: With Compartment Security it can be specified that only those with the combined roles of “Top Secret”, “US Citizen” and “Engineering” can access the data. Without compartment Security, it is only possible to choose that anyone with at least one of the roles can access the data.

For example, with Compartment Security, a government document classified at the “Top Secret” level with an additional security marking of “NOFORN” (“no foreign nationals”) cannot be read unless the user has both a “Top Secret” role and a role that describes the individual as citizen of that particular country.

When to Use

External KMS – Use an external KMS for additional separation of concerns and ease of management for storing encryption keys. This option is also helpful when you want to leverage an external KMS that is already in use at your organization

Redaction – Use Redaction when certain pieces of your data needs to be removed or obscured when exporting data for sharing. This feature may be helpful in meeting compliance guidelines (e.g., HIPAA, SEC17a-4, FINRA, GDPR, etc.)

Compartment Security – Use Compartment Security anytime AND semantics is required to further restrict data access. It is often employed to protect classified material in government systems

About MarkLogic

MarkLogic is the world’s best database for integrating data from silos, providing an operational and transactional Enterprise NoSQL database platform that integrates data better, faster, with less cost. Visit <http://www.marklogic.com> for more information.