MarkLogic®

How to Deploy MarkLogic Securely

MARKLOGIC WHITE PAPER · AUGUST 2018

Maintaining security throughout the deployment process is critical—all the way from planning to going live into production. In this white paper, we provide an introduction on how to deploy MarkLogic securely, looking at the ways in which MarkLogic enables secure deployments from both organizational and product perspectives..

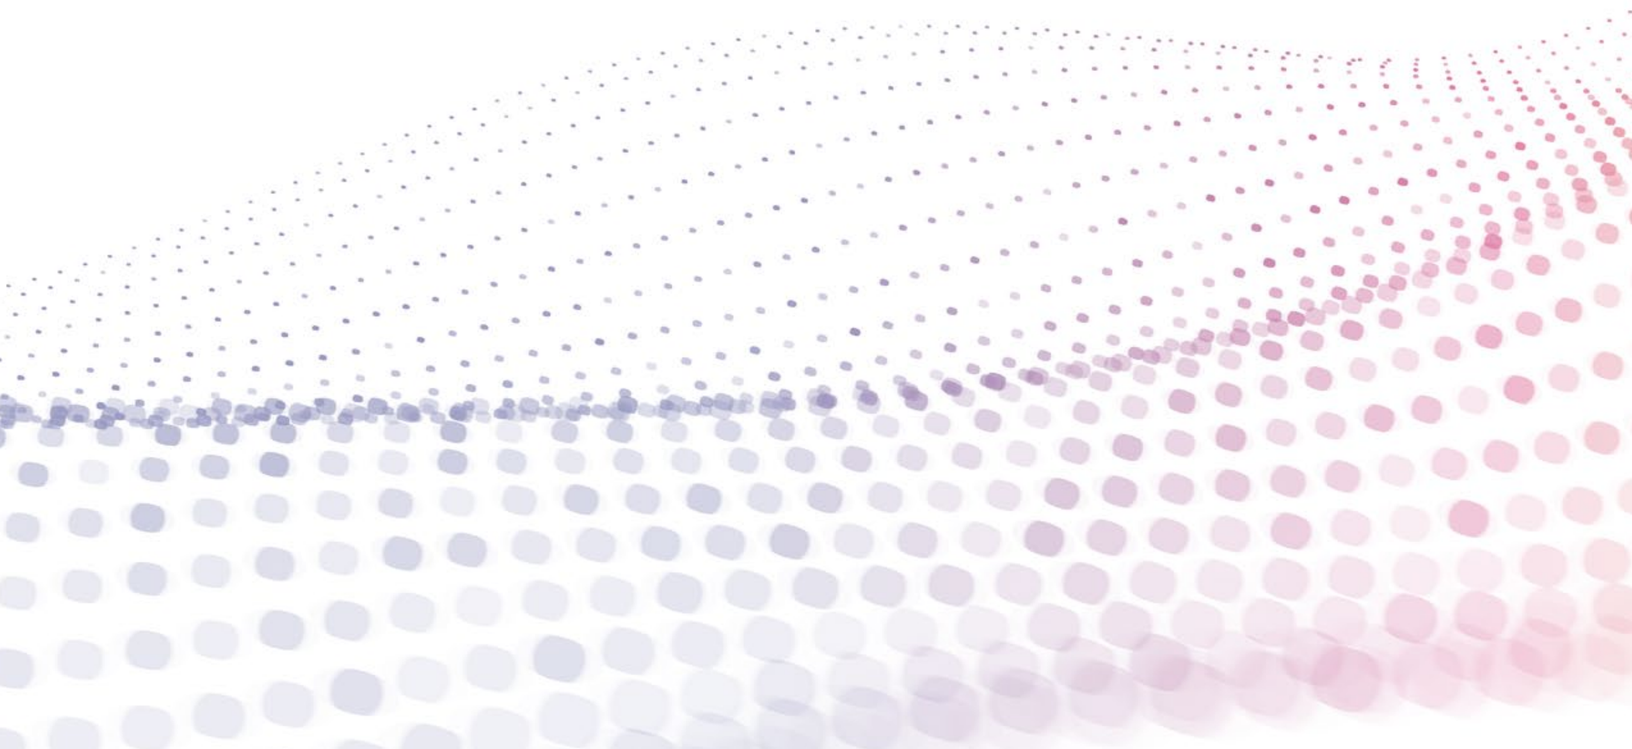

Contents

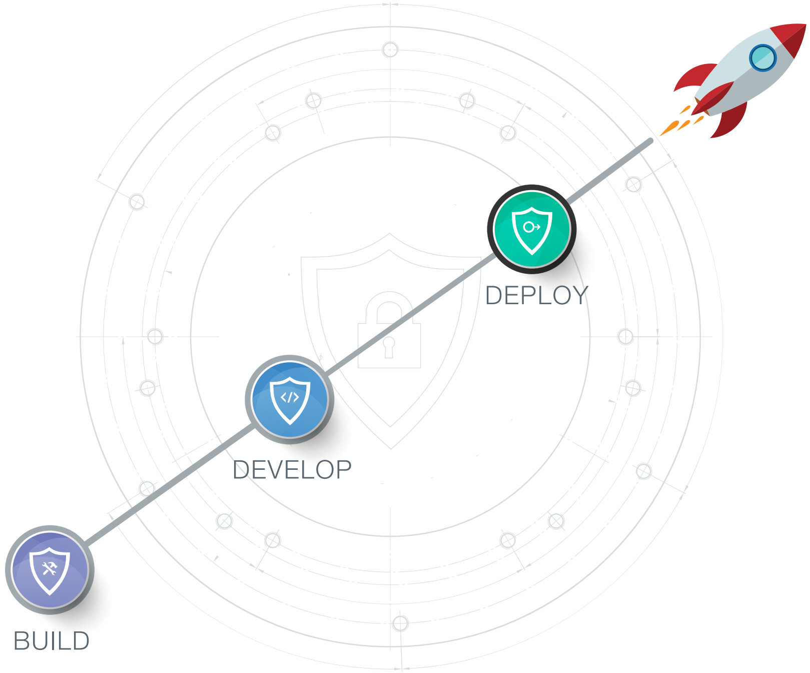

Introduction

Maintaining security through the deployment process is critical—all the way from planning to going live into production. Our overall approach to security involves looking at an integrated security ecosystem so that organizations can deploy MarkLogic securely while maintaining complete peace of mind. The MarkLogic security ecosystem is framed by three main components:

- **How We Build a Secure Product**
  This area focuses on how our company's engineering team applies best practices, tools, and techniques to build the most secure product possible.

- **How to Develop Secure Applications on MarkLogic**
  This area focuses on the use of integrated security services and capabilities built into the MarkLogic platform that are available for use by application developers during the development lifecycle.

- **How to Deploy MarkLogic Securely**
  This area focuses on ensuring that MarkLogic is deployed into a secure environment. It includes the ability to work with industry-standard security technologies (e.g., LDAP, Kerberos, SSL/TLS, and KMIP) and also organizational support such as education and consulting.

In this white paper, we focus on that third aspect—how to deploy MarkLogic securely. We provide an in-depth look at the ways in which MarkLogic enables secure deployments from both organizational and product perspectives. We also provide a checklist for DBAs that helps guide the process.

Key Ways MarkLogic Helps

Every organization has unique constraints and opportunities based on their current IT environment and capabilities. MarkLogic is designed to fit into existing security infrastructure, following industry best practices and standards, in addition to providing enterprise documentation, training, support, and consulting services. Below is a list of the key ways MarkLogic helps you deploy securely:

1. **Integration with your environment** – MarkLogic is designed to seamlessly integrate with your existing environment, leveraging security infrastructure and best practices to ensure that the database is itself a secure product, and that it has all the features you need to manage your data throughout its lifecycle.
2. **Support for standards** – MarkLogic supports proven industry recognized and accepted security standards that are used across the database industry, including everything from authentication protocols and PKI standards for hash algorithms to encryption standards used to ensure data integrity.
3. **Timely updates** – MarkLogic responds quickly to newly discovered threats with updates and patches delivered in a timely manner. For example, MarkLogic made patches to address the Heartbleed OpenSSL vulnerability in only five days.
4. **Documentation** – MarkLogic provides thorough documentation, including specific guidance that focuses on deployments, including a whole guide on Securing Your Product Deployment.
5. **Training** – MarkLogic provides free training through MarkLogic University, with an entire track and certifications just for DBAs.
6. **Support** – If your organization has a support contract, you have a direct line to the full engineering resources of our customer support. The Customer Support Handbook has more information.
7. **Consulting Services** – You can leverage our best practices and extensive experience in deploying MarkLogic into mission-critical environments through MarkLogic Consulting Services.

Integration with Your Environment

MarkLogic has numerous interfaces to ensure interoperability with existing enterprise security environments. Existing infrastructure investments do not have to change. In this section, we take a closer look at MarkLogic's capabilities that make it easy to seamlessly integrate MarkLogic into your environment.

Fine-Grained Access Controls

One of the important capabilities to consider when deploying a new system is how access controls are managed and the role granularity.

MarkLogic is designed with Role Based Access Control (RBAC) at the document level to govern who can access what data based on their privileges and permissions. These privileges and permissions are very fine-grained, guarding access at the document level, though you can also get even finer-grained control by setting element- and property-level permissions on data inside documents.

Another important aspect of managing database security is looking at who actually has access to the operating system and file system that the database runs on. With MarkLogic, the administrator does not need access to the operating system or file system. Most MarkLogic administrative tasks can be scripted and then executed without ever needing to share credentials with administrators.

When needed, administrative privileges can be "amped" to provide increased privileges necessary for a short duration tasks. Amps allow some administrative duties to be delegated without needing to give users full permission to become database administrators (DBAs). For example, an amp may be helpful when a DBA needs to delegate administrative duties for managing user passwords for a defined set of users in a Platform as a Service (PAAS) cloud environment. The DBA can amp a "change-password" function to run as "admin," allowing users to call that function without actually being a database admin.

Authentication

In order to authenticate users, MarkLogic leverages existing enterprise authentication databases and the existing enterprise namespace for user and resource names. MarkLogic provides Single Sign-On (SSO) capabilities and consistently enforces security policies using LDAP and Kerberos. Mutual authentication is supported through use of industry standards such as SSL/TLS, PKI, and X.509.

MarkLogic supports external authentication by means of LDAP and Kerberos and certificate based authentication. When a user attempts to access a MarkLogic application server configured for external authentication, the requested application server sends a hash of the username and password to the LDAP server for authentication.

Communications Encryption

Secure Sockets Layer (SSL) and Transport Layer Security (TLS) are communications security standards for providing encrypted communication for data-in-motion, usually HTTPS (Note: TLS is newer and replaced SSL). Typically, a handshake procedure authenticates the server so that the client can trust the server but the client remains unauthenticated to the server. MarkLogic supports mutual authentication where the client also holds a digital certificate, which it sends to the server.

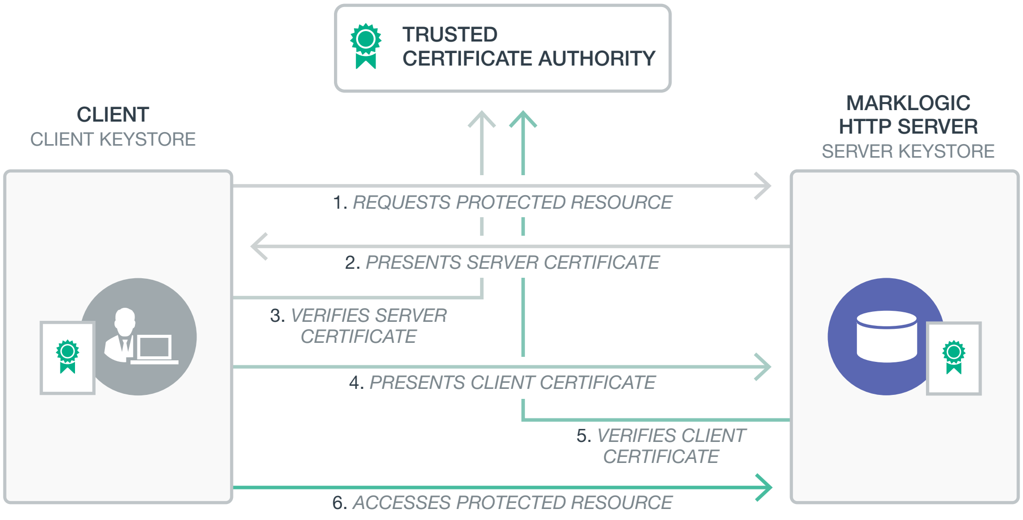
TRUSTED
CERTIFICATE AUTHORITY
CLIENT
CLIENT KEYSTORE
MARKLOGIC
HTTP SERVER
SERVER KEYSTORE
1. REQUESTS PROTECTED RESOURCE
2. PRESENTS SERVER CERTIFICATE
3. VERIFIES SERVER CERTIFICATE
4. PRESENTS CLIENT CERTIFICATE
5. VERIFIES CLIENT CERTIFICATE
6. ACCESSES PROTECTED RESOURCE

Figure 1: This graphic depicts MarkLogic's mutual authentication between the database and client applications using FIPS enabled TLS/SSL. A similar process is used to secure communication between database nodes, cluster to cluster, and Admin GUI to database.

MarkLogic uses OpenSSL to implement SSL/TLS, thus enabling mutual authentication, and maintains the current version of OpenSSL. Any fixes available will always be incorporated into MarkLogic. MarkLogic also supports a FIPS 140-enabled version of OpenSSL and supports all the operations necessary to enabling full SSL/TLS encryption, including the automated certificate setup.

Key Management & Encryption

Every database that has data encryption uses keys and has to securely manage those keys. MarkLogic goes above and beyond what most databases offer in two ways: (1) key rotation; and, (2) the separation of data and keys.

Data keys are encrypted with an intermediary key to enable fast rotation. Continually updating keys provides an additional level of security just like updating your password on a frequent basis. And, even if a key was cracked before it was rotated, new data keys get generated regularly and frequently and each key only protects a very small subset of data. So, the risk of data compromise is very low. Dictionary, known plaintext, or brute force attacks would compromise just one file.

In addition to fast key rotation, MarkLogic stores and manages the encryption keys separate from the encrypted data. Many other databases do not do this—they store keys and data in the database together. MarkLogic's approach is more secure and is an industry recommended best practice per OWASP security recommendations.

MarkLogic uses enveloped encryption keys with keys managed by separate entities other than the database administrator, ensuring separation of duties. MarkLogic sends the 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.

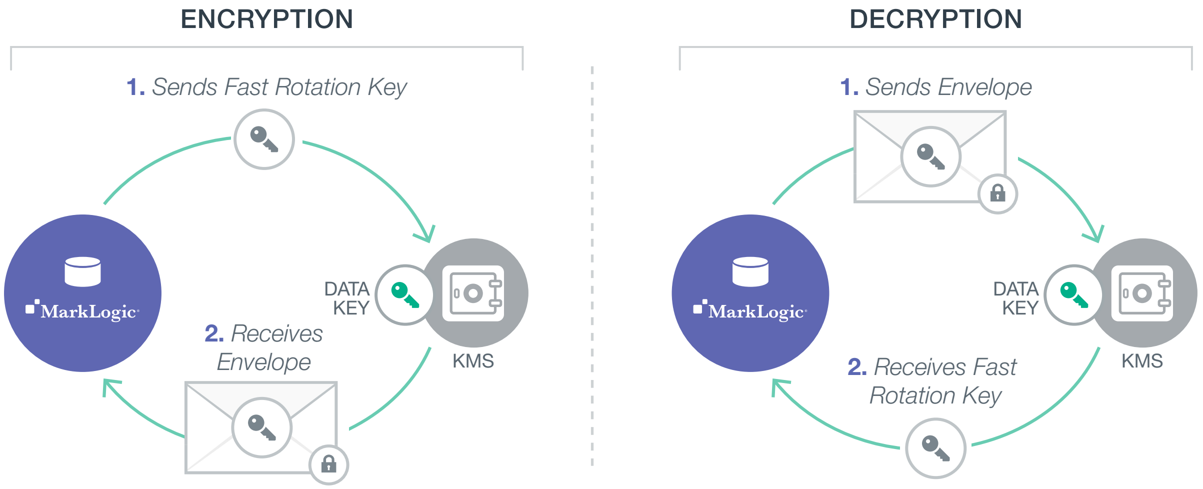

Figure 2: With Encryption at Rest, to access low level keys and read files, MarkLogic sends an envelope to the KMS, which then sends back the unencrypted key. If using an external KMS, MarkLogic has no access to enveloped keys, which means no access to files.

MarkLogic provides powerful key management either through a local KMS or external Key Management System (KMS)[1]. A KMS, or "keystore," is a secure location where the enveloped encryption keys used to encrypt data are stored and managed. Both options for key management, local and external, provide high performance and are transparent to users:

1. **Local KMS** (enabled by default) – Default rapid key rotation is provided with the product and is enabled if the local KMS option is selected and encryption is enabled.

2. **External KMS** (add on) – If you have purchased the Advanced Security option, then you can enable *Advanced Encryption*, which makes it possible to use an external third party KMS that conforms to the KMIP 1.2 standard.

When Encryption at Rest has a Key Management System (KMS) deployed and managed externally, separately from the application servers in the MarkLogic cluster, it provides additional separation of duties between the Security Administrator, Application Administrators, and other System Administrators.

MarkLogic interoperates with third-party external KMS systems that are KMIP 1.2 compliant. Key Management Interoperability Protocol (KMIP) is a communication protocol standard that defines message formats for the manipulation of cryptographic keys on an industry-standard key management server.

Auditable Events

Auditing is the monitoring and recording of selected operational actions of application users and administrative users. Auditing provides many security benefits. It deters users or potential intruders from inappropriate actions and provides the ability to investigate suspicious activity and notify an auditor of inappropriate actions from an unauthorized user.

Auditing is used to verify other security controls as well. For example, if an audit policy is designed such that no audit events of a certain type should be detected during normal operation, then any detected audit event of that type will point to an abnormal or unexpected condition which could be the result of an inadequate security control. As a result, the security control can be evaluated and improved. Auditing is an essential part of almost every regulatory compliance standard.

MarkLogic supports a comprehensive auditing system which logs information about data accesses, security reconfigurations and administrative changes. Some notable examples include logs about authentication failures, configuration changes, disabling of FIPS mode, security access successes and failures, and more. These logs can be exposed to log file management tools to make monitoring and escalation easier and more automated. MarkLogic can also encrypt these logs to reduce leakage.

[1] To use an external Key Management System (KMS), you must purchase the Advanced Security option separately. The Advanced Security option includes the ability to use an external KMS, Redaction, and Compartment Security.

Support for Standards

MarkLogic supports proven security standards in order to integrate seamlessly into enterprise environments and support the use of best practices. Below is a partial list of standards supported by MarkLogic.

Category	Standard	Security Attribute
Authentication	Kerberos	Authentication standard used to authentication entities on a network
	Lightweight Directory Access Protocol (LDAP)	Authentication, authorization
	Secure Sockets Layer (SSL)	Mutual authentication used for data in motion encryption
	Transport Layer Security (TLS)	Mutual authentication used for data in motion encryption
	Public Key Infrastructure (PKI)	Mutual authentication
	RSA	Asymmetric encryption used for data in motion confidentiality
	X.509 v3	Certificate standard used for authentication
	Public Key Certificate Standards (PKCS)	Certificate standard used for authentication
Authorization	Role Based Access Control (RBAC)	Fine-grained, role-based authorization
	Lightweight Directory Access Protocol (LDAP)	Authentication, authorization
	Compartment Security	Authorization for highly secure environments
Confidentiality	American Encryption Standard (AES)	Symmetric encryption used for data in motion confidentiality
	Key Management System/Key Management Interoperability Protocol (KMS/KMIP) 1.2	Key management protocol standard used for data at rest encryption
	FIPS enabled communications	Data in motion confidentiality
Integrity	Secure Hash Algorithm (SHA-1, 2 or 3)	Hashing/Digest algorithms to ensure data integrity

Deployment Checklist

In this section, we provide a closer look at a typical MarkLogic deployment environment from the perspective of security. The typical environment, depicted in Figure 3, consists of a 3-tier architecture: Tier 1 – End user clients; Tier 2 – Application servers; and, Tier 3 – Database running on a cluster.

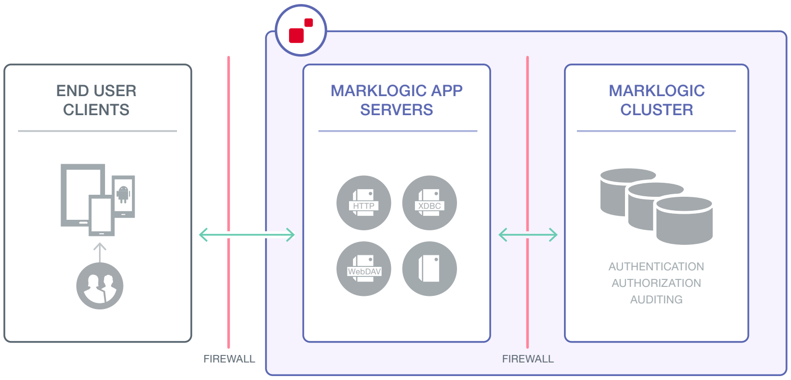

Figure 3: The typical deployment environment using MarkLogic consists of end user clients, application servers, and a database.

Not every deployment is the same. However, there are some common elements and some frequently used steps to take in order to ensure a good outcome. Below are some common steps used to secure the environment and the application.

Harden the Enterprise Environment

Devices, servers, database machines. Secure these at the physical level. Infrastructure is all registered, certifications, patches, etc.

End User Endpoints

- Authenticate device endpoints (all devices are untrusted to begin with)
- Provision credentials using a well-known trusted Certificate Authority (CA)
- Enforce a password policy for end users and devices
- *Optional:* Implement MFA (Multi-Factor Authentication)

App Server Machines

- Implement OS and security patches
- Remove unsecure services (e.g., FTP, Telnet, etc.)
- Disable un-used TCP/UDP ports (e.g., port 80)
- *Optional:* Implement M2M (Machine-to-Machine) authentication
- Ensure load balancers are configured correctly
- Always use mutual authentication between endpoints
- Monitor SysLog and EventLog for unusual activity
- Close all ports except 8000, 8001, 8002 (and except for 7997 – 7999 used for XDQP communication)
- IdAM Services (for LDAP, disable 389 to ensure communication is only on 443)

Database Cluster Machines

- Implement OS and security patches
- *Optional:* Implement M2M (Machine-to-Machine) authentication
- Implement separation of roles (e.g., SysAdmin, DBA)
- Create network fencing
- Change all default passwords
- Monitor SysLog and EventLog for unusual activity

Secure the Application Environment

After hardening the environment, the next step is to secure the application environment. The high level steps are described below.

Secure End User Client Authentication

- Strong and centralized password/credential management policy
- Strong mutual authentication for end users (certificate-based)
- External Key Management Systems for Certificate management

Implement Communications Security (i.e. "Secure the pipes")

- Mutual strong authentication between services and hosts
- Security policy for key management lifecycle
- Encrypt communications security per NIST guidelines

Secure the App Server and Database Connections

- Minimal secure configuration for app services and Databases.
- Strong load balancer and network security policy
- Encryption at rest for data, logs, and configuration files

Other General Tips

It is important to stay up to date on security issues. Two key resources include:

- Information Assurance Support Environment – Website that contains Security Technical Implementation Guides (STIGs) to 'lock down' information systems and software
- CERT Program – The Computer Emergency Response Team is part of the Software Engineering Institute, a federally funded research and development center that is operated by Carnegie Mellon University and provides best practices for cybersecurity

Timely Updates

No environment is immune to issues and it is often impossible to predict what issues may arise. The most important thing is that when issues are uncovered, they are addressed as quickly as possible.

MarkLogic responds to issues extremely quickly, in line with the severity of the issue. For example, when the industry-wide Heartbleed vulnerability discovered in OpenSSL became known, MarkLogic took swift action to release patches within only five days. In addition, MarkLogic actively engaged with CERT and major OS vendors to ensure that MarkLogic customers were protected.

Here is the summarized timeline of MarkLogic's updating the product to protect against Heartbleed:

- Advisory published by CERT April 7, 2014
- Testing complete and security advisory published to customers April 10, 2014
- Windows and Linux Updates for all supported versions released April 11, 2014
- All other platforms (Solaris, Macintosh, etc.) patched and released April 12, 2014

The MarkLogic Support Handbook provides more information about response times to support cases for those customers who have a support contract.

Documentation

MarkLogic provides guidance and support to enable organizations to effectively leverage existing customer infrastructure and security best practices in order to deploy a secure data environment based on MarkLogic.

Access the Security Guide at http://docs.marklogic.com/guide/security. As part of the documentation there are Sample Security Recipes, which are step-by-step instructions on how to properly implement security.

Training

Free training is available for application developers, database administrators, and system administrators through classes delivered online (at globally convenient times) from MarkLogic University.

Go to http://www.marklogic.com/training/ to register for on-demand and instructor-led training.

Support

Customers bet their business on MarkLogic, so MarkLogic Support is available for customers 24x7 with timely help available from knowledgeable support engineers.

Enterprise customers supporting mission-critical applications should not waste time in multi-tier support queues, so MarkLogic does not use them. Instead, MarkLogic connects the customer directly with an engineer who has real experience building, deploying and supporting production MarkLogic applications. Security issues encountered by customers can be reported directly to MarkLogic support.

The MarkLogic global support organization can follow the sun to work urgent issues until they are resolved. And, with the backing of MarkLogic product development and professional services organizations, MarkLogic provides holistic solutions to ensure customer long-term success.

See http://www.marklogic.com/services/support/ for more information.

How to Contact Support

Once registered as a support contact, you can contact MarkLogic Technical Support via:

- Email – support@marklogic.com
- Web – https://help.marklogic.com
- Phone – 1-855-882-8323

We recommend that all support requests be submitted via either email or web, to enhance the process of reporting, tracking and resolving issues. Support requests for urgent issues (as defined in Understanding Case Priority and Response Time Targets) should be submitted at any time via email to urgent@marklogic.com.

Consulting Services

MarkLogic Consulting Services provides recommendations for best practices based on a decade of NoSQL database deployments of the most demanding and mission-critical applications in the world. MarkLogic consulting is very familiar with the *MarkLogic Security Model* and will work with customer application developers to design security into every application.

See http://www.marklogic.com/services/consulting-services/ for more information.

Conclusion

In this white paper, we provided an overview of how to deploy MarkLogic securely. We discussed MarkLogic's ability to work with industry-standard security technologies (e.g., LDAP, Kerberos, SSL/TLS, and KMIP) and also the organizational support we provide such as education and consulting.

Additional white papers, linked to below, are intended to answer additional questions for developers building applications and for security and business professionals interested in learning how we build security into the product itself.

Lastly, it is worth noting again that data security is a constantly evolving topic. This white paper is only an introduction to how to deploy MarkLogic securely, and you should always refer to the documentation and MarkLogic support team for the most up to date information.

Key Resources

MarkLogic Concepts Guide on Security
https://docs.marklogic.com/guide/concepts/security

Understanding and Using Security Guide
https://docs.marklogic.com/guide/security

White Paper – Building Security Into MarkLogic
http://www.marklogic.com/resources/building-security-marklogic/

White Paper – Developing Secure Applications on MarkLogic
http://www.marklogic.com/resources/developing-secure-apps-marklogic/

White Paper – Secure Architectures When Deploying MarkLogic On-Premises
https://www.marklogic.com/resources/secure-architectures-deploying-marklogic-premises/

MARKLOGIC CORPORATION
999 Skyway Road, Suite 200 San Carlos, CA 94070
+1 650 655 2300 | +1 877 992 8885 | www.marklogic.com | sales@marklogic.com

MarkLogic®

999 Skyway Road, Suite 200 San Carlos, CA 94070
+1 650 655 2300 | +1 877 992 8885
www.marklogic.com | sales@marklogic.com