MarkLogic®

Top Data Security Concerns When Integrating Data

MARKLOGIC WHITE PAPER · AUGUST 2018

Data security is a top priority for organizations and there are a plethora of tactical details that DevOps and security experts need to worry about. But, what should CIOs, architects, and business leaders focus on at a strategic level? In this white paper, we discuss the top data security concerns – with a particular focus on data integration – and provide an overview of how MarkLogic® addresses those concerns as a database.

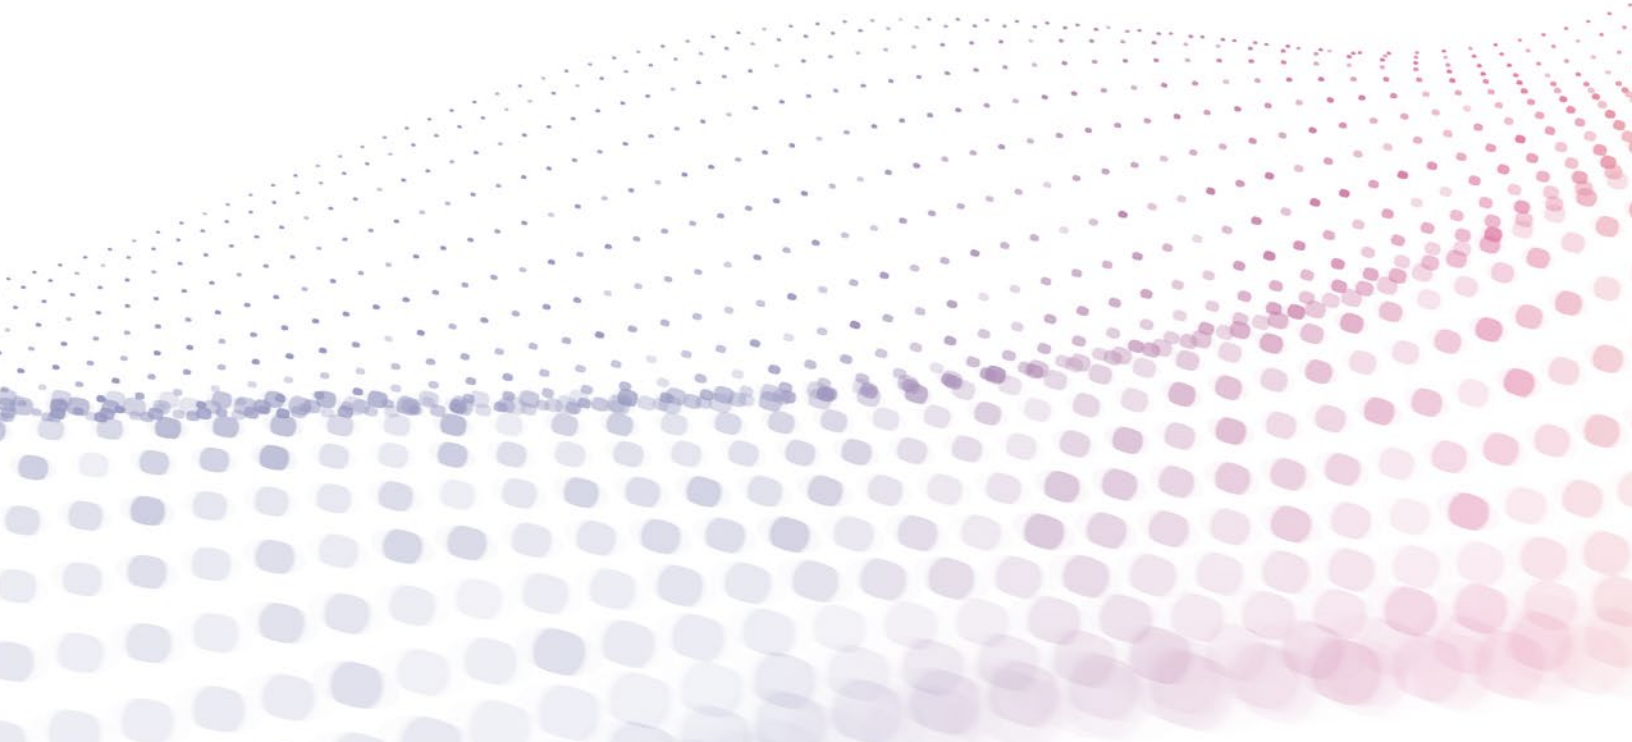

Contents

Introduction 1

Concern #1: Traditional Data Integration Creates Security Vulnerabilities 2

Concern #2: Application Developers Are Burdened With Data Security 3

Concern #3: Unknown, Unmanaged Risks From Insider Threats 6

Introduction to MarkLogic Security 7

Certified, Granular, Government-Grade, & Comprehensive
Security Certifications and Standards

How A Secure Database Supports the Data Governance Lifecycle 10

Data Quality
Provenance & Lineage
Security & Privacy
Compliance
Lifecycle
Availability

Deeper Dive Into MarkLogic Security 12

Conclusion 12

Additional Resources

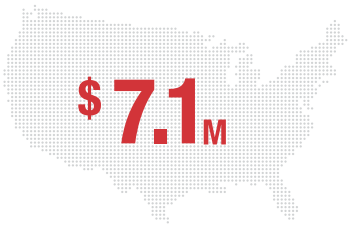

AVERAGE COST OF CYBER INCIDENTS ON U.S. COMPANIES[1]

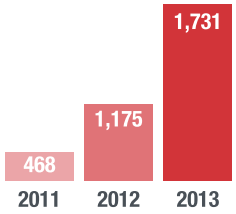

MAJOR BREACHES REPORTED. SEE THE TREND?[2]

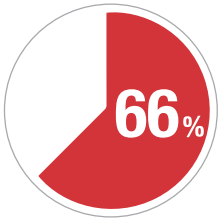

% OF COMPANIES THAT SAY THEY DEPLOY NEW IT WITHOUT APPROPRIATE SECURITY MEASURES IN PLACE[3]

Introduction

Headlines reporting cyberattacks, ransomware, and compromises in data security are increasingly common. It makes sense that data security is now a top priority—the risk of not securing data is simply too high. There is no shortage of splashy numbers that highlight the problem:

- Each cyber incident costs U.S. companies a reported $7.1 million on average, or $221 per record[1]
- In 2011, there were 468 major breaches recorded. In 2012, 1,175. In 2013, 1,731. See the trend?[2]
- Two-thirds (63%) of organizations deploy new IT prior to having appropriate data security measures in place[3]

Despite increasing awareness and spending, the problem with data security is getting worse.

In addressing the problem, it is easy to get buried figuring out how to protect against the latest incident patterns and attack vectors. The tactical details are important, and your DevOps and security experts need to be working together to tackle them. But, there are also more strategic concerns to consider.

At a more strategic level, we see three primary concerns that CIO's, architects, and business leaders should consider. These concerns are commonly shared across industries and have particular relevance to data integration:

- **Concern #1:** How traditional data integration with relational databases creates security vulnerabilities
- **Concern #2:** How application developers are unduly burdened with data security
- **Concern #3:** How insider threats create unknown, unmanaged data security risks within the network perimeter

In this white paper, we take a closer look at these concerns, and discuss how MarkLogic helps organizations address them as the most secure NoSQL database available today. It is one of the reasons that large investment banks, major healthcare organizations, and classified government systems around the world run their most demanding, mission-critical systems on MarkLogic.

1 Survey of 64 U.S. based organizations. IBM and Ponemon Research, 2016 Cost of Data Breach Study: United States, 2016. <http://www-03.ibm.com/security/data-breach/>

2 For the gory details, check out the Veris Community Database at http://vcdb.org/explore.html. This data is summarized in the Verizon Breach Report, also a helpful resource: http://www.verizonenterprise.com/verizon-insights-lab/dbir/.

3 451 Research Data Threat Report, 2016. <http://enterprise-encryption.vormetric.com/rs/480-LWA-970/images/2017-Thales-Data-Threat-Report-Advanced-Technology-Edition.pdf>

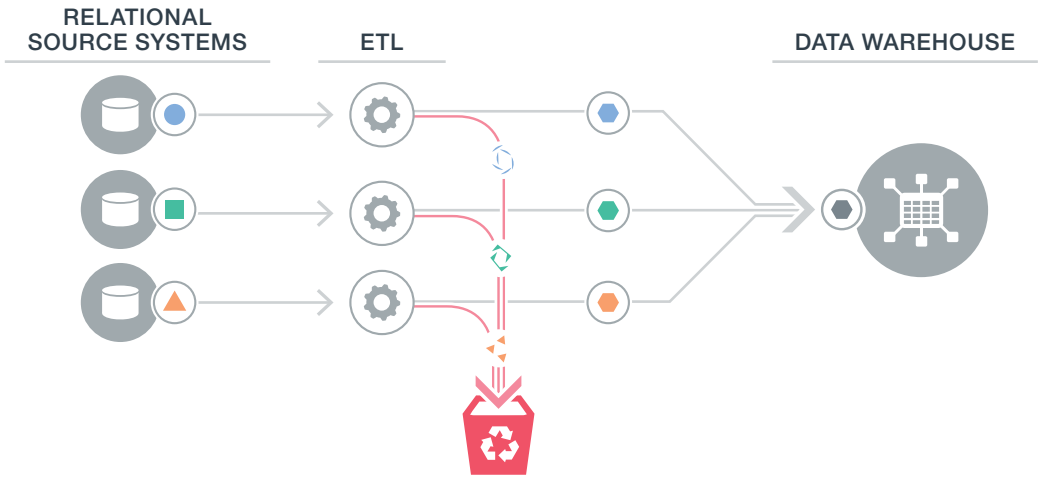

Figure 2: The traditional approach to data integration with relational databases and ETL leads to data loss and goveranance problems.

Concern #1: Traditional Data Integration Creates Security Vulnerabilities

Role-and policy-based access controls are essential to govern, preserve, and audit data and associated entitlements. If these controls are not managed, you introduce unnecessary complexity and risk.

Unfortunately, most organizations have a proliferation of relational database silos. Each one has separate security access controls that make it virtually impossible to adequately track and protect all of the data. Additionally, there are multiple ETL tools with obfuscated code and integration points, not to mention their own access controls that need to be managed. With an increasing number of data silos, there are more opportunities for exploits.

Often, what happens is a team builds a complex ETL process from multiple databases to a centralized analytical data warehouse—all using relational databases. The ETL is done to (a) simply make the system be able to function, and (b) to "cleanse" the data because there is a business process that requires standardization so that the system can be used to count things, do math, or disambiguate the data. But, far from ensuring quality, this cleansing process may actually be reducing quality by removing important data.

To a data analyst, some metadata may seem like "data lint" that needs to be laundered, but to a compliance analyst or data modeler, that same "data lint" may be required for critical business reasons (say, to prove to a regulatory agency that your trades were legal in order to avoid a hefty fine).

Over time, it becomes more and more difficult to maintain data governance. Data governance, which we define simply as the application of policy to data, includes many aspects: Data quality, lineage and provenance, security and privacy, compliance requirements, availability. If an organization fails to pay close attention to each aspect of data governance across the entire lifecycle of data, they open themselves up to additional cyber risk.

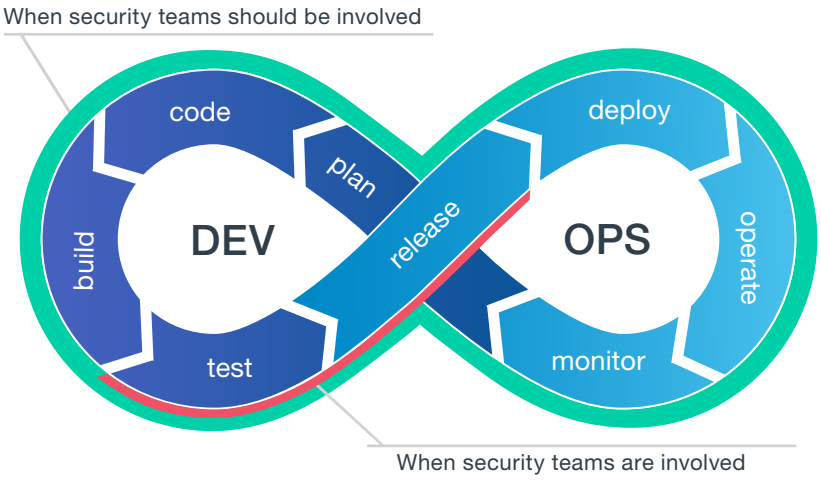

Figure 3: There is a disconnect between DevOps and security teams. Security is often only worked on during testing and release rather than through the whole lifecycle.

How MarkLogic Helps

MarkLogic makes data integration a good thing for security and data governance.

First, MarkLogic reduces the burden of traditional ETL. By handling the process of ingesting source data *as is* and transforming and harmonizing the data inside MarkLogic, the whole process of integrating data becomes faster and more seamless. No data gets discarded during the process.

Second, MarkLogic's multi-model approach using documents and triples is better for governing data over time. You can manage high level business concepts from multiple silos, materializing them as entities and relationships. Data and metadata stay together and you can track the details across the lifecycle—its provenance, who can see it, how it changed—all in a single system. By taking a more comprehensive approach, MarkLogic reduces opportunities for exploits and provides a more agile platform to handle new and changing regulations.

Concern #2: Application Developers Are Burdened With Data Security

It's really hard to secure your data up and down the stack and across multiple data silos. Unfortunately, security is often not tested and maintained in a single data layer. Instead, the burden is put on developers to secure data at the application layer for every new application. With regulation around data privacy and security that organizations now have to account for (HIPAA, SEC17a-4, FINRA, GDPR, etc.), the stakes are higher and the burden is growing.

This is problematic because development and security teams are often disconnected. A disconnect has grown because of the move towards DevOps and agile development. Both are positive improvements to software development that enable shorter release cycles. With these approaches, it may be common to commit small batches of code every few days—if not hours.

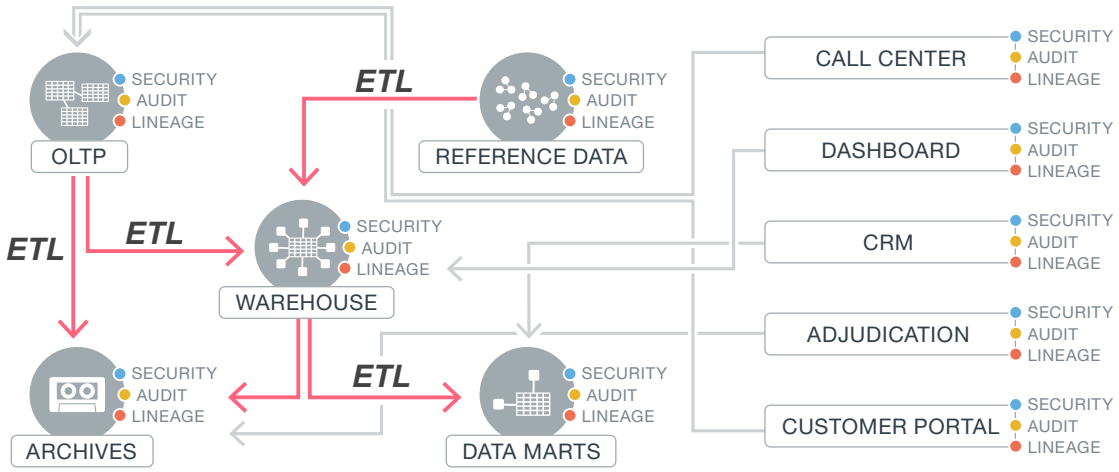

Figure 4: Unless security is handled in a more centralized database, what results is a spaghetti architecture that leads to more vulnerabilities. This graphic does not even depict the systems for backup and recovery, development, and testing that also require security monitoring maintenance.

Unfortunately, security teams cannot keep up. Security review cycles are designed to take weeks or months, and security certification and accreditations are bound to waterfall methods, not continuous improvement. Most developers know the OSWAP Top Ten, but the real security experts are only brought into the development process to do a final check before go-live.

According to Gartner, 90 percent of companies using DevOps consider security an afterthought.[4] It is no surprise then, that according to the Department of Homeland Security, 90 percent of exploits are due to defective software.[5]

One example showing the disconnect between teams is at Intuit, which adopted an agile, DevOps approach for their 3,000-person team. Shannon Lietz, senior manager for cloud security engineering at Intuit, said in an interview, "We realized that the DevOps teams were throwing [responsibility] over the wall to security, and [security] had all the information; they knew all the attacks that were coming in, and the DevOps people did not have the information to make the decisions."[6]

While most organizations are not the size of Intuit, the challenge is often similar. A development team is tasked with stitching together multiple technologies with different, usually quite limited security capabilities. The security team is out of sync and cannot keep up.

To solve this problem, organizations should implement many tactical recommendations: Develop closer integration between security and DevOps teams to close the feedback loop, make security checks more automated by performing dynamic code analysis (and perform such checks earlier and more frequently in the sprint lifecycle), design for security, improve Identity and Access Management (IAM) systems, enforce segregation of duties, and conduct risk and threat modeling for applications.

Additionally, it is important to take a broader, more strategic look at how data is managed at the lowest possible level—*in the database*.

4 Gartner, DevSecOps: How to Seamlessly Integrate Security Into DevOps. September, 2016. <https://www.gartner.com/doc/reprints?id=1-3KWUQXV&ct=161028&st=sg>

5 Department of Homeland Security Infosheet, reporting on research done by the Security Engineering Insitute at Carnegie Mellon. <https://www.us-cert.gov/sites/default/files/publications/infosheet_SoftwareAssurance.pdf>

6 TechTarget. Robert Lemos, <http://searchsecurity.techtarget.com/feature/DevOps-security-requires-new-mindset-and-tools-for-visibility-automation>

How MarkLogic Helps

The goal is to keep data governance governable across the stack. If you move to using a centralized database to govern and secure the data, securing applications becomes easier and faster. The work of data governance happens in one place so that one change in data policy at the database level can be automatically applied to a hundred applications rather than having developers make a hundred manual changes to application code.

MarkLogic has extensive capabilities to govern and secure data in the database, which in turn helps with many of the aspects of application security.

The SANS Institute, a well-known cybersecurity training organization, provides a SWAT checklist to help development teams.[7]

SWAT Checklist
(Securing Web Application Technologies)

1. Error handling and logging
2. Data protection
3. Configuration and operations
4. Authentication
5. Session management
6. Input and output
7. Access control

Of this list, MarkLogic fully addresses numbers 1, 2, and 7 – error handling and logging, data protection, and access control – and also helps address the rest (3, 4, 5, and 6). By addressing many of these concerns in the database, the attack surface is decreased significantly.

One of MarkLogic's key underlying capabilities that makes data security stronger and easier to implement is **Role Based Access Control (RBAC)**.

RBAC governs who can access what data based on their privileges and permissions. These privileges and permissions work to secure data at the document level. MarkLogic also has **Element Level Security**, which makes it possible to secure pieces of data inside documents (more on this later). Working together, these features make life easier on developers by managing the access controls in the database.

Additionally, MarkLogic has programming APIs so developers can create and execute policies utilizing all of the security and data protection capabilities in MarkLogic (e.g., backup, retention, data access, data lifecycle, and authentication). Policies can be associated with data, metadata, and data attributes so that policies such as those for privacy or compliance can be easily executed. And, the security controls and checks are transparent to developers.

Beyond these features, MarkLogic also has additional out-of-the-box features designed to help organizations with compliance. **Bitemporal** data management ensures that historical data remains unchanged and that you have a full audit trail of data. Also, **Compliance Archive** provides a mechanism to protect data from changes, and save the data to WORM (Write Once, Read Many) storage.

All of these features means smarter data management in the database, less work for developers to do at the application level, reduced time and complexity around security testing, and better security resilience.

7 Note: This checklist includes references to the common weakness enumerators that map very closely to those referenced by the OWASP Top Ten, which many people are more familiar with.

Concern #3:
Unknown, Unmanaged Risks From Insider Threats

Typically, most organizations put an immense focus on implementing endpoint, application, perimeter, and network security—and for good reason. Preventing intrusion into your network is a critical part of securing your infrastructure. Some companies see hundreds of thousands of intrusion attempts against their network—*every single day*.

But focusing only on *network security* is like creating a hard shell around a soft, squishy middle. If you can get in, you're in. The truth is, no network perimeter will ever be impenetrable. There are likely bad actors already in the network.

Some of the biggest data breaches have occurred because an insider got the keys to the kingdom. And, the number of incidents involving internal actors is increasing.

The numbers vary, but in general, internal actors are involved in 25 percent of all breaches.[8] In the healthcare industry, insiders are responsible for 68 percent of breaches.[9] Unfortunately, many systems are vulnerable to such attacks because they only have all-or-none data access rather than fine-grained security controls.

Complicating the insider threat problem is the fact that modern enterprises have staff, contractors, sub-contractors, trading partners, consultants, auditors, and other people involved. It is very difficult to discern just who is 'inside' and who is 'outside.'

Sometimes, it is relatively innocuous data management decisions that can create the biggest insider threats. For example, many organizations have data lakes that are virtual treasure troves

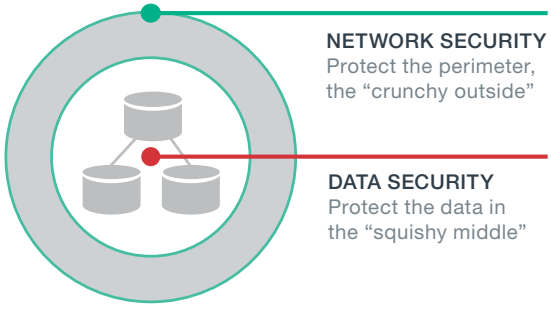

Figure 5: Focusing only on network security, the perimeter might be somewhat secure, but the data lives in the "squishy middle" that becomes extremely vulnerable.

of data with broad access to users. One global bank we work with spent years building a data lake using another technology. But, they shut it down for security and compliance reasons when they realized the new system did not have proper controls and that were potentially violating certain rules and regulations regarding customer data.

Organizations today need better data security. It is not an option, however, to just lock everything down. While the most secure database in the world might be one that is locked in a safe and dropped in the bottom of the ocean, that data would not be very shareable.

In the quest for data security, it is important to still maintain data sharing. Organizations must have proper security controls to ensure that the right portions of data are accessible and shareable with those in and outside the company who are granted proper access. And, there must be a separation of duties so that administrators granting access do not themselves have access to the crown jewels.

8 Verizon. 2017 Data Breach Investigations Report: 10th Edition. < http://www.verizonenterprise.com/verizon-insights-lab/dbir/.>

9 IBM. Security trends in the healthcare industry. February, 2017. https://www-01.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=SEL03123USEN

How MarkLogic Helps

As discussed in the previous section, MarkLogic has fine-grained access controls designed to provide optimal data security even when sharing data. One additional feature that directly addresses the problem of insider threats is **Advanced Encryption**.

Without encryption, or even with file system encryption, the system administrator, cloud operator, or hacker could access or modify files—including the files that comprise the database.

MarkLogic's Advanced Encryption allows data, configuration, and logs to be encrypted on disk (i.e., encrypted at rest). This feature requires no modification to applications developed on MarkLogic. And, the optional use of an **External Key Management System** (**KMS**) further ensures separation of duties and integration into existing security infrastructure.

Introduction to MarkLogic Security

With the context of these concerns and a bit about how MarkLogic addresses them, let us now provide a full overview of MarkLogic's security capabilities.

As a company, we focused on security from the start. Without strong data security, you cannot safely manage enterprise data, and that is what MarkLogic was originally designed for—integrating, storing, searching, and managing enterprise data. Some database vendors forget about the "M" in DB**M**S, but "management" is central to how MarkLogic is designed.

With fine-grained access controls, separation of duties, data segmentation, advanced encryption, and more, MarkLogic has the features you need to deliver the triad of Confidentiality, Integrity, and Availability (CIA). Whether you're an IT executive or security manager, performing security audits and reviewing controls, responsible for deploying applications, or for ensuring software supply chain safety—MarkLogic provides the necessary data protection to exceed modern enterprise requirements.

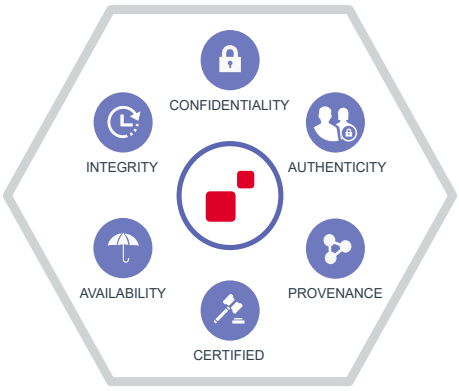

Certified, Granular, Government-Grade, & Comprehensive

Highly *Certified* & Compliant with Major Systems Security Standards

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. It is not easy to meet the requirements to be Common Criteria certified, and the list of vendors is short.

MarkLogic is one of only six vendors that offers a database that is Common Criteria certified, and MarkLogic is the only NoSQL database with the certification.

MarkLogic is also installed and operational on government systems with demanding security policies. These policies include stringent measures for access, authentication, management, audits, role separation, and system assurance. For example:

- **NIACAP** (National Information Assurance Certification and Accreditation Process) – Developed by the U.S. intelligence community for certification of computer and telecommunications systems that handle U.S. national security information
- **NIST Special Publication** 800-37 – Guide for applying risk management to federal information systems. It supports the six-step Risk Management Framework (RMF)

Security Certifications and Standards

MarkLogic's Security Certifications

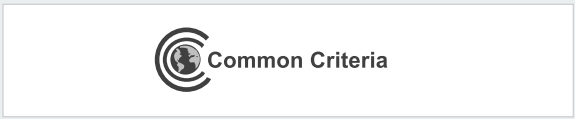

Common Criteria Certification – The driving force for the widest available mutual recognition of secure IT products worldwide. It is not easy to meet the requirements to be Common Criteria certified, and the list of vendors is short. MarkLogic is one of only six vendors that offers a database that is Common Criteria certified, and MarkLogic is the only NoSQL database with the certification.

Additional Security Standards

- NIACAP
- NIST Special Publication 800-37
- NIST 800-53
- ICD 503
- FIPS 140-2
- HIPAA
- SOX 302/404
- FedRAMP
- SSAE 18
- EU 95/46/EC

Customers have also received Authority to Operate (ATO) for information systems utilizing MarkLogic that involve almost all of the major systems security standards. These standards continue to evolve and MarkLogic stays up to date on the latest changes (for example, SSAE 18 replaced SSAE 16).

The system security standards currently in place on systems running MarkLogic include the following:

- NIST 800-53
- ICD 503
- FIPS 140-2
- HIPAA
- SOX 302/404
- FedRAMP
- SSAE 18
- EU 95/46/EC

Granular Security at the Lowest Levels

As mentioned, MarkLogic uses a Role Based Access Control (RBAC) security model by default, 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). Security checks verify the necessary credentials before granting the requested action, and security information is stored in a specific security database in MarkLogic.

MarkLogic closely monitors database activity and makes it possible to audit document access and updates, configuration changes, administrative actions, code execution, and changes to access controls.

MarkLogic supports **external authentication** using Lightweight Directory Access Protocol (LDAP) or Kerberos. MarkLogic also supports **strong certificate-based authentication** with Public Key Infrastructure (PKI) and Certificate Authorities (CAs).

Additionally, beyond RBAC, MarkLogic supports **Attribute Based Access Control (ABAC)** and **Policy Based Access Control (PBAC)**. These models further restrict access based on attributes, (i.e., metadata about the data such as provenance, geo-location, time of day, etc.), policy information stored in document metadata, or simple labels representing “high” or “low” levels of trust.

Beyond just securing data at the level of individual documents, MarkLogic also has even more fine-grained security. **Element Level Security** provides access control at the level of JSON properties or XML elements within documents, regardless of schema. Specific information inside a document may be hidden from users based on their role, while still providing access to other information in the document. Element Level Security is akin to “cell-level” security in relational databases. But, it is a step above “cell-level,” as it is not restricted to protecting a certain set of cells in a relational database schema.

Government-Grade & Trusted for Mission Critical Use Cases

MarkLogic has been in the business of protecting and securing data for over a decade, and is installed and operational on sensitive government systems that require databases to meet extremely rigorous requirements.

MarkLogic exceeds the security requirements to serve as the trusted platform to run the most demanding, mission-critical applications at the heart of large investment banks, major healthcare organizations, and classified government systems.

Comprehensive Security, Built-in From the Start

Security is an end-to-end feature in MarkLogic, where data, data security, and data-driven policies are all tied together. In other words, security travels with the data.

Since the first version of MarkLogic was released, we have continued to improve security in each subsequent release. For example, MarkLogic 1 (originally called Cerisent XQE Server) included RBAC. And, MarkLogic has continued to maintain the Common Criteria Certification since MarkLogic 4.

Additionally, the security features are designed to scale. Many MarkLogic customers are running extremely large product systems, and the security checks and data encryption processes do not slow down data access.

Here are a just a few of the many examples of major organizations using MarkLogic:

KPMG

KPMG built a MarkLogic-powered application to support client onboarding primarily for the purposes of compliance with regulation, tax, and reporting. The application uses intelligent automation of complex manual processes and maintains a fully traceable, auditable data workflow.

Deutsche Bank

MarkLogic replaced Oracle as the global trade store for the bank's operational trade data. The first production deployment integrated dozens of trading systems and launched in just six months—all while maintaining secure and consistent transactions.

U.S. Combatant Command

MarkLogic replaced Oracle to serve as the data layer for a command-wide knowledge- and information-sharing system for an increasingly diverse dataset consumed by a wide variety of programs and people in the U.S. Department of Defense.

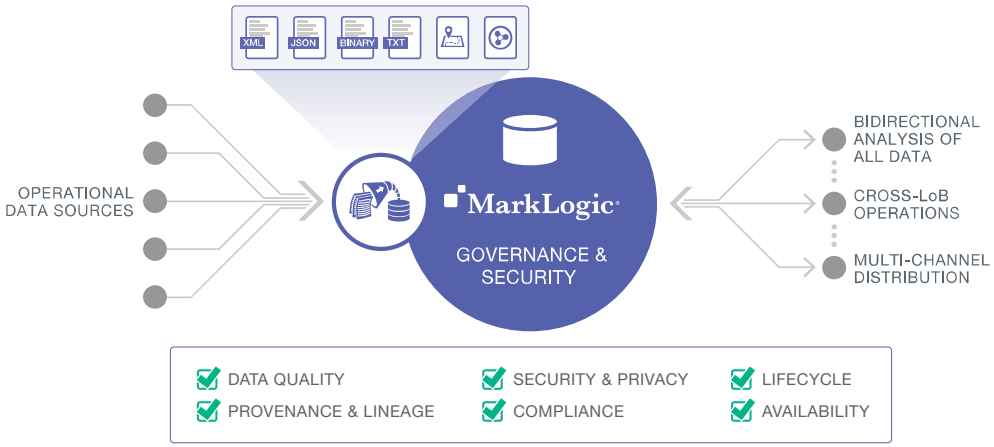

Figure 7: One common architectural pattern is the Operational Data Hub, shown here to illustrate the lifecycle of integrating data from silos with MarkLogic. Throughout the process, MarkLogic maintains the highest standards of data governance and security.

Advanced Security Option for Certain Security Use Cases

In addition to the key features that come out-of-the-box with MarkLogic, some customers need additional capabilities for certain use cases:

- **External KMS Support** OPTION – This option makes it possible to use an external Key Management System, or KMS[10], to help with Advanced Encryption, which is often done for the additional separation of concerns and ease of management
- **Compartment Security** OPTION – With Compartment Security, more complex rules can be applied to documents so that a user must have *all* of the right roles to access or create a document rather than just *one* of the rights roles. This is often useful when handling classified material
- **Redaction** OPTION – Similar to Element Level Security, but focused on securing data on *export* rather than real-time protection when querying data. Redaction eliminates the exposure of sensitive information by making it possible to remove existing information or replace it with other values when exporting data or sharing. The process is simple, flexible, and is designed to work with large volumes of data

How A Secure Database Supports the Data Governance Lifecycle

With a better understanding of MarkLogic's security capabilities, you may now be asking, "How does this help me with data governance when integrating data?"

In this section, we summarize the key components of data governance, the questions that define each component, and how MarkLogic checks each of the boxes. We already discussed some of the features mentioned, and so will skip some of the detail here.

If you are interested in learning more about MarkLogic and data governance, you can watch a recording of a keynote presentation by our SVP of Engineering, David Gorbet, which walks through each of the following components.

Data Quality

Is the data fit for purpose? Is it accurate, timely, consistent, etc.?

MarkLogic's flexible data model makes it easier to track lineage and provenance, and not discard any raw data. With MarkLogic, data and metadata can

10 For example, SafeNet, Vormetric, or other vendors that are KMIP-compliant.

stay together and you can transform data within the database. You can manage multiple schemas and avoid the problems of data loss that you get with a traditional approach.

You can also define flexible and fluid validation rules that execute in MarkLogic. As data is ingested, it can be rejected as invalid or accepted. If accepted, you can flag it so it is not used until reviewed, used for some use cases but not others, or used only for some people and made invisible to others. The database is designed to be flexible and support a variety of policies. It is not restrictive.

Provenance & Lineage

Where did the data originate, and how did it change?

With MarkLogic, you can validate the data and metadata together, without worrying about ETL transformations that may have been done to "cleanse" the data to make it fit a certain schema. MarkLogic is designed to handle messy, changing data—including data from different sources and schemas.

Security & Privacy

Is the data protected with fine-grained access controls? Is it encrypted? Is your database certified?

This is where MarkLogic's robust security features apply. As mentioned, MarkLogic has a huge number of security features designed to secure data at a fine-grained level. These features are all designed to work at scale. And, the database is certified by a third party.

Compliance

Are you in compliance with regulations, and can you demonstrate it?

MarkLogic can handle frequent rule changes because of the flexible data model and through fast data access using built-in search. If a regulator asks a question that was not planned for, that is okay because all data is indexed with the "Ask Anything" Universal Index. Other Compliance features such as Bitemporal and Compliance Archive further help address regulatory concerns. For these reasons, customers use MarkLogic to help comply with specific regulations such as GDPR and MiFID II.

Lifecycle

How is the data changed, stored, and accessed as it ages?

In addition to having a flexible data model that supports messy, changing, complex data, MarkLogic's Tiered Storage feature makes it possible to define a policy-based tiering strategy for data storage—including age-based policies. With this feature, you can automatically manage the movement of data through its lifecycle, from fast storage to slower storage to a queryable archive.

Availability

Does your system meet your Service Level Agreements (SLAs) for durability, consistency, high availability, and disaster recovery?

As emphasized, enterprise capabilities have been built into MarkLogic from the start and they have been proven through thousands of enterprise deployments. MarkLogic has the features required – most notably **HA/DR** and **ACID transactions** – to provide five nines of availability (i.e. available 99.999% of the time, or always available except for 5.26 minutes per year).

Other new features such as **Rolling Upgrades** (eliminates downtime during upgrades), **Ops Director** (single view to monitor and manage clusters), and **Telemetry** (opt-in support line) make management of MarkLogic faster and easier.

Deeper Dive Into MarkLogic Security

Listed below are some additional resources If you are interested in going deeper into MarkLogic security from a more technical perspective.

WHITE PAPER

Building Security Into MarkLogic

Our company's engineering team applies best practices, tools, and techniques to build the most secure product possible, using the ***MarkLogic Security Framework*** to guide the process.

WHITE PAPER

Developing Secure Applications on MarkLogic

We provide integrated security services and capabilities built into the MarkLogic platform that are available for use by developers and DBAs. The ***MarkLogic Security Model*** provides a conceptual, multi-layered view of how MarkLogic implements security.

WHITE PAPER

Deploying MarkLogic Securely

To deploy MarkLogic into a secure environment, we provide guidance on best practices through education and consulting, in addition to ensuring that MarkLogic is compatible with industry-standards (e.g., LDAP, Kerberos, SSL/TLS, KMIP, etc.).

Conclusion

In this white paper, we looked at three top cybersecurity concerns, provided an introduction to MarkLogic security, and highlighted key aspects of data governance. Our view is that by focusing more on security at the level of the database, it is possible to prevent a lot of the common security and data governance issues from happening in the first place.

Tying back to the top data security concerns discussed earlier in this white paper, MarkLogic's main value proposition is the unique data and indexing approach that contributes to the robust security and data governance capabilities. Based on that foundation, we developed a better database to integrate, store, manage, and search your data:

1. **Improved data integration** – Using MarkLogic to integrate data helps prevent a spaghetti architecture of data silos and ETL from consuming time and resources, and opening up unnecessary security vulnerabilities.
2. **Centralized data governance** – Bringing your data together in MarkLogic aids developers and the security experts by centralizing a lot of the security policies and reducing the time they would have ordinarily spent duplicating security and data governance across applications.
3. **Fine-grained security capabilities** – MarkLogic's fine-grained security is on by default, and provides the necessary access controls that modern enterprises need to manage data access while not limiting the ability to share their data securely.

By choosing a database like MarkLogic that is built with the necessary controls for modern data security, you can get the agility you need while still having a system you can trust with your most critical data. In other words, you can have both agility and security.

Additional Resources

The following are additional resources that provide more information about security and data governance.

PRESENTATION

Security Keynote: SVP of Engineering

David Gorbet, SVP of Engineering, MarkLogic

Provides an overview of MarkLogic's approach to security and data governance.

PRESENTATION

Data Security In Practice

Caio Milani, Director of Product Management, MarkLogic

Provides an overview of security features in MarkLogic, including the new features in MarkLogic 9: Encryption, Element Level Security, and Redaction.

PRESENTATION

Data Governance in an Unpredictable World

Damon Feldman, Ph.D., Solutions Director, MarkLogic

Provides examples of how MarkLogic improves data governance in regulated industries.

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

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