Reinforcing your enterprise data security strategy with **Transparent Data Encryption**

White paper

FUJITSU

Reinforcing your enterprise data security strategy with **Transparent Data Encryption**

The reality of today's business is that while IT systems need the flexibility of cloud and DevOps, they also need to satisfy the increasing need for data security and customer data protection.

Stricter legislation requires all types of companies across industries to protect their own data as well as their customers' in order to comply with privacy legislation such as the European Union's GDPR. The need to protect their data is also caused by malicious hackers trying to piece together information so they can impersonate someone's identity, for either financial gain or to hide their own identity.

Every organization, irrespective of its size, relies on data and information exchange to perform certain business operations. Effective use of data also enables organizations to offer new services and engage with their customers better and in a timelier fashion.

The first point of protection is realized by storage devices and file systems that encrypted their data, but for databases the problem is much larger than a single installation. Encryption must be easily applied and must be carried out in such a way that the data is secure and easily recovered in case of failure or disaster recovery scenarios.

Anticipating this, Fujitsu Enterprise Postgres implemented Transparent Data Encryption, which securely stores individual tablespaces and logs encrypted with their own encryption keys, ensuring security of user data while preserving ease of management for database administrators.

Fujitsu's implementation of Transparent Data Encryption hardens your organization's data security with minimum performance overhead, no additional storage required, and without changes to existing applications.

Executive Summary

Organizations are accountable for the safety and confidentiality of its business data, client data, and employee information. There are several regulatory and compliance requirements that must be adhered to. Data breaches can have severe consequences such as downtime, expensive legal fees, and more importantly, reputational damage. It is therefore imperative that companies employ data security mechanisms and procedures to protect their data against threats.

Realizing an effective security policy

An effective data security policy leads to the protection of data from unauthorized access, use, change, disclosure, and destruction, and comprises two essential components.

Data classification

Data classification is an important aspect of the security policy. Classification is based on the level of sensitivity and impact on an organization, should that data be accessed, modified, or deleted without authorization. This classification helps determine what baseline security controls are appropriate for safeguarding that data. These levels vary between organizations, depending on the nature of the business. Classifying data as confidential, private, or public is one such example.

- Confidential data: Leakage of data classified as confidential can cause a significant level of risk to the organization or its partners. The highest level of security controls should be applied to such data.
- Private data: Unauthorized access of data that is classified as private could result in a moderate level of risk to the organization or its partners. A reasonable level of security controls should be applied.
- Public data: Generally, this is public information. While little or no control is required to protect the confidentiality of public data, care should be taken to prevent its unauthorized modification.

This classification is never static and needs to be constantly assessed over the life cycle of the data itself.

Layered security framework

The second component of data security generally follows a layered approach in protecting sensitive data from intruders, often referred to as defense in depth.

The layers include human, physical, network, application, data, and other detection technologies, deployed in such a way that a breach in one layer does not compromise the entire system of data protection.

The level of security and costs is often determined by the value of the data, which in turn determines the classification of data. The tools and techniques deployed are further challenged, as hackers become more sophisticated in their attacks, bypassing security measures.

The implementation of the security policy, can be viewed as securing the data in its three digital states:

- Data at Rest
- Data in Motion
- Data in Use

In this white paper, we will specifically discuss technologies that enables securing data at rest.

Securing data at rest

Data at rest refers to persistent data that is stored in any digital form – *i.e.*, files, spreadsheets, databases, etc. To prevent this data from being accessed, modified, or stolen, organizations use different security protection measures.

There are many ways to protect data, and some of them include strong user authentication, role-based access control, multi-factor authentication, data encryption or a combination of methods. There are also dynamic monitoring tools used to detect and prevent intrusion, based on rules, patterns and policies.

It is crucial for organizations to know where their sensitive data resides and deploy a combination of techniques that best match their needs. Following are some of the key enabling technologies for securing data at rest.

Customer data protection

Data is a valuable business asset that is the target of malicious actors both inside and outside the organization. Breaches incur all types of costs that far outweigh the costs associated with protecting the data in the first place.

Passwords

Passwords are the most widely used method to prevent unauthorized access to systems, applications, files, and data. Having a good password policy is essential to keeping computer systems secure. There are several other methods of authentication as well, including biometric, multi-factor authentication and token-based access control.

Role-based Access Control

Role-Based Access Control (RBAC) is based on the premise that users do not have discretionary access to enterprise objects. Instead, access permissions are associated with roles. Users are made members of roles as determined by their responsibilities, which determines access permissions. The benefit of RBAC is that users can be easily reassigned from one role to another without modifying the underlying access structure.

Encryption

Encryption offers protection by scrambling data, so only the owner of the key or password can read the data. This protects the confidentiality of the data so that if an unauthorized person gains access to the storage device or service, they will not be able to obtain any information. It also protects the integrity of the data so that it cannot be tampered with without the owner knowing it. There are several types of encryption:

- Full-disk encryption (FDE) is the encryption of all data on a disk drive, including the program that encrypts the bootable OS partition. FDE prevents unauthorized drive and data access.
  Some disk encryption solutions have support for a Trusted Platform Module (TPM). These implementations can wrap the decryption key using the TPM, thus tying the hard disk drive (HDD) to a particular device.
- Filesystem-level encryption, often called file-based encryption or file/folder encryption, is a form of disk encryption where individual files or directories are encrypted by the file system itself. Types of filesystem-level encryption include:
  - Cryptographic filesystems on top of the main file system
  - Encrypted general purpose filesystems
- Database Encryption

  Database administrators and database users must understand the sensitivity or classification associated with a database and its contents to ensure that sufficient security controls are applied.

  In cases where all of a database's contents are of the same sensitivity or classification, an organization may choose to classify the entire database at this level. Alternatively, in cases where a database's contents are of varying sensitivity or classification levels, and database users have differing levels of access, an organization may choose to apply classification at a more granular level within the database.

  Limiting a database user's ability to access, insert, modify, or remove content based on their responsibilities ensures that the need-to-know principle is applied, and the likelihood of unauthorized modifications is reduced.

  Since its inception, encryption has been held as one of the top data protection techniques available. This security approach enables the user to scramble the content of protected systems using keys and utilize a decryption key to decipher it.

Transparent Data Encryption

This is a database encryption technology that solves the problem of encrypting data at rest. It is an integral element in the data security continuum. The term *transparent* denotes the fact that the encryption method is transparent to authorized users of the database, as no change is required in the applications or existing access policies.

At a high level, the encryption method protects data in the database by encrypting the underlying files. No meaningful information can be obtained if data is accessed through unauthorized access to the disk or database. To access the data, the original encryption certificate and key are required.

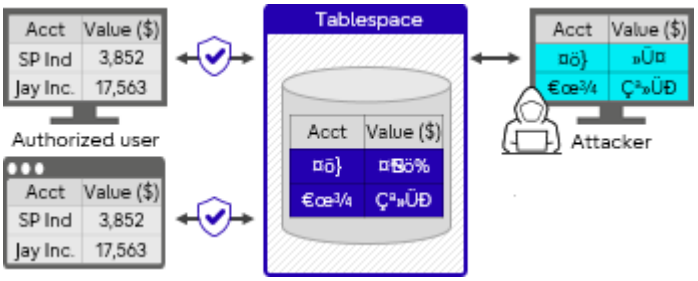

Figure 1 - Even if a data breach occurs, data is not compromised

Transparent Data Encryption key attributes

- Compliant with Payment Card Industry Data Security Standard (PCI DSS)
- 256-bit encryption
- No need to modify existing business applications for migration
- Encryption key can be changed without having to re-encrypt data
- Storage data security using Advanced Encryption Standard (AES)
- Support for streaming replication, as objects encrypted on the primary server are transferred in its encrypted format to the standby server

Minimum overhead allows you to protect more data without impacting performance.

Encryption/decryption is performed by manipulating entire blocks, instead of one bit at a time, which increases its performance and results in minimum overhead to the process. Overhead can be further minimized by using AES-NI built into processors that provide this feature, as is the case with several Intel and AMD processors.

As a result, you no longer need to reduce the scope of encryption to ensure application performance, and you can encrypt all data of an application with minimum impact.

Fujitsu Enterprise Postgres and Transparent Data Encryption

Fujitsu Enterprise Postgres is bundled with Transparent Data Encryption out of the box. Some of the key features are:

- Existing applications require no change, as data is transparently encrypted when it is written to the disk and decrypted when it is read from the disk
- Fast encryption/decryption, with minimum overhead
- Unlike other commercial databases, additional licenses are not required to use this functionality
- No overhead in storage areas, as the encryption algorithm does not alter the size of the object being encrypted.
- It is possible to encrypt a subset of the data as per the organization's data classification policy, so that stringent rules can be applied to that portion of the data.
- Multiple encryption keys can be deployed, which in turn are encrypted by the master encryption key. The master encryption key is also encrypted based on a passphrase.
- Encryption is extended to logs, backups, temporary tables, and temporary indexes, providing comprehensive security
- Support for streaming replication, as objects encrypted on the primary server are transferred in its encrypted format to the standby server.

Tablespace granularity

Encryption is applied at the tablespace level –all tables, indexes, temporary tables, and temporary indexes created in the specified tablespace will be automatically encrypted.

This allows you to encrypt important data, but also maintain metadata and other reference data in a non-encrypted tablespace, to avoid encrypting/decrypting overhead.

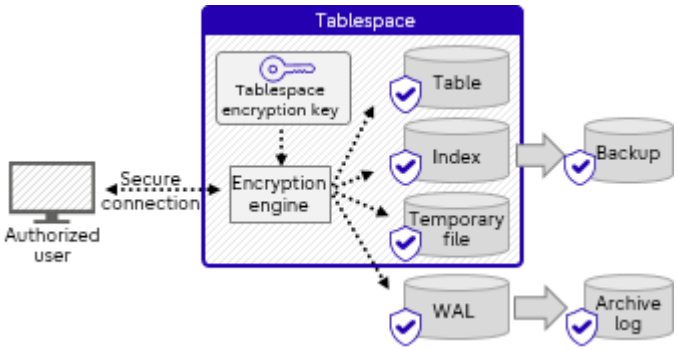

Figure 2 - Indexes on row-oriented data (reading data)

Simple to implement

Transparent Data Encryption is set up as follows:

1. Create a master encryption key.
   A. In `postgresql.conf`, set the location of encryption key.
   ```
   keystore_location='/disk1/keystoreloc'
   ```
   B. Create the master key using a pass phrase.
   ```
   SELECT pgx_set_master_key('mysecretkey');
   ```
   The file `keystore.ks` will be created in `keystore_location`.
2. Create an encrypted tablespace.
   A. Specify the encryption type desired– 128-bit or 256-bit Advanced Encryption Standard.
   ```
   SET tablespace_encryption_algorithm = 'AES256';
   ```
   B. Restart the database server, opening the keystore.
   ```
   pg_ctl --keystore-passphrase restart -D /home/db
   ```
   C. Create a tablespace, specifying its physical location.
   ```
   CREATE TABLESPACE tbspc1 LOCATION '/data/encTbs1';
   ```
3. Create data files in the encrypted tablespace.

   Data files for tables and indexes in this tablespace will be encrypted. Even existing unencrypted tables and indexes will be automatically encrypted when moved to this tablespace.

Transparent Data Encryption scope

- Tablespaces (tables, indexes, temporary tables, temporary indexes)
- Backup
- Temporary files
- WAL

Underlying data files will be safe at this point

You can check how data at rest is safe in encrypted tablespaces.

In the command line, gather the relevant database information.

```
SELECT oid FROM pg_tablespace WHERE spcname = tblspc;
SELECT oid FROM pg_database WHERE datname = d;
SELECT relfilenode FROM pg_class WHERE relname = table;
```

B

C

Display the content of the data file.

```
cat PGDATA/pg_tblspc/tblspcOid/PGvers/dbOid/tblNnode;
```

A B C

Unlike with unencrypted data files, the content of the file above is garbled and will not expose any meaningful data.

Conclusion

Transparent Data Encryption should be an integral part of your organization's data security policy, as it protects all the data at rest. It provides the ability to comply with many laws, regulations, and guidelines as required in different industries. It also enables developers to secure their data using secure encryption algorithms without changing their applications.

Contact us

If you have any questions about the Transparent Data Encryption and other enterprise security features of Fujitsu Enterprise Postgres, feel free to contact us at **enterprisepostgresql@fujitsu.com**.

About Fujitsu

Fujitsu is the 5th largest IT service provider in the world, offering a full range of technology products, solutions, and services. Around 126,000 Fujitsu employees support customers in over 100 countries.

Fujitsu Enterprise Postgres provides powerful tools to protect your organization from database security threats – **Transparent Data Encryption** to protect data at rest, **Data Masking** to redact data in-flight, and the **Dedicated Audit Log** to record database access and monitor suspicious/unauthorized activity.

Fujitsu Enterprise Postgres can help your journey

Fujitsu Enterprise Postgres is the enhanced version of PostgreSQL, for enterprises seeking a more robust, secure, and fully supported edition for business-critical applications.

It is fully compatible with PostgreSQL and shares the same operation method, interface for application development, and inherent functionality. Designed to deliver the Quality of Service (QoS) that enterprises demand of their databases in the digital world, while supporting the openness and extensibility expected of open source platforms, all at a lower cost than traditional enterprise databases.

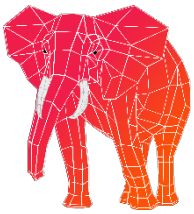

Fujitsu Enterprise Postgres

Combine the strengths of open-source PostgreSQL with the enterprise features developed by Fujitsu.

Enhanced speed, security, and support — without the costs associated with most proprietary systems.

Fujitsu Enterprise Postgres for Kubernetes

Utilize operator capabilities for provisioning and managing operations on the OpenShift Container Platform.

Business-ready database that integrates container operation technology for rapid development-to-production deployments.

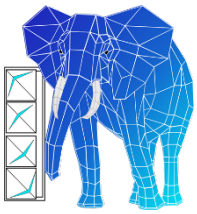

Fujitsu Enterprise Postgres on IBM LinuxONE™

World-class platform that embraces open source and improves data security, performance, and business continuity.

The best of open source flexibility with the peace of mind that comes from knowing it is backed by Fujitsu and IBM.

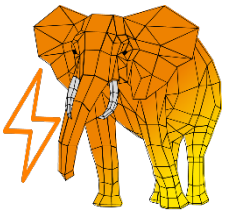

Fujitsu Enterprise Postgres on IBM Power®

Experience frictionless hybrid cloud that can help you modernize to respond faster to business demands.

Fujitsu database designed for security, performance, and reliability, combined with IBM server built for agility in the hybrid cloud.

Discover how Fujitsu Enterprise Postgres' unique security features take PostgreSQL to the next level to protect you from data breaches and ensure you comply with data protection regulations such as GDPR at fast.fujitsu.com/enhanced-security-for-enterprises/

FUJITSU

Contact
Fujitsu Limited
Email: enterprisepostgresql@fujitsu.com
Website: fast.fujitsu.com

2022-09-08 WW EN