


Best practices for sensitive data using **Data Masking**

White paper



 **FUJITSU**

The hybrid multi-cloud Postgres you can rely on

Best practices for sensitive data using Data Masking

Data breaches damage a company's reputation, destroy customer confidence, and have major financial implications.

Many organizations unwittingly reveal sensitive data when they provide unrestricted access to data to their own staff or staff of third parties, or copy sensitive or regulated data into nonproduction environments, making data in those environments a target for cybercriminals.

To maintain customers' trust, organizations must take responsibility and put more focus on the management of sensitive data.

This white paper discusses the root causes contributing to organizational noncompliance with regulations governing data management, and explains how to implement best practices for sensitive data.



Executive summary

Are you aware that sensitive data may be at risk in your organization?

- Your staff may have unrestricted access to sensitive data
- Sensitive data is made available to third parties with which you have a commercial arrangement
- Production data has been copied to development and test environments, and that data may be accessible by cybercriminals, and can be lost or stolen

Data breaches not only damage your organization's reputation and brand. They also have legal and financial implications.

Fujitsu has identified a need to help organizations better protect their customers' sensitive data. In response, it has added a Data Masking feature to its flagship enterprise database management system, Fujitsu Enterprise Postgres.

Read on as we discuss issues relating to sensitive data management, from the causes of data breaches, to implementing best practices for managing sensitive data in your organization. Learn about Fujitsu's implementation of Data Masking, and how it can help mitigate data breaches.

Introduction

Hackers have stolen headlines across the world with highly publicized information security breaches over the last few years. Some of these, such as the hacking of the Australian Bureau of Meteorology, increase the public's awareness and general acceptance of data vulnerability in today's society.

While average consumers are mindful of public recommendations on how to protect their personal data, the convenience of online shopping, bill paying and the draw of social media is overpowering. This often leads to increasingly placed trust in those organizations to which they supply personal information.

The problem

Lack of responsibility of organizations

As a result of the perceived convenience to the consumer, many organizations are now accumulating vast amounts of personally identifiable information. This information is at risk of attack if sufficient security measures are not implemented and maintained. Furthermore, unrestricted access by the staff of these organizations, as well as third parties with which they may have a commercial relationship, considerably exacerbates the potential for harm.

Consider for a moment, two of the most popular trends today, one a business trend, the other a technology trend, and their impact on the securing of sensitive data.

Outsourcing

Three of the most commonly outsourced services, "Customer Support", "Accounting", and "Tax Preparation", all involve making personal information available to third parties. What is an adequate level of protection for this data? Nondisclosure agreements certainly don't cut it, even if those countries' laws enforced significant penalties for breaches.

Application security frameworks

On the technology side, many organizations now manage access (authorization) to data at the application layer, using various security frameworks that come embedded within their chosen application development framework. These often see a single generalized database role with an elevated degree of privileges being used by all requests from one or more applications. In some instances, the superuser role is even used.

Organizations must take responsibility and put more focus on the management of sensitive data.

Most organizations store data in one type of database or another, as this is the most efficient way of finding information when it comes to using it. So this is where security of data must begin.



Challenges in protecting data

While just modifying sensitive data sounds very straightforward, there are a number of challenges that include identification and use of sensitive data, auditing change, maintaining data integrity, managing volume, and maintaining flexibility of data usage.

Regulatory noncompliance

The Global Financial Crises of 2007-2008 saw the emanation of regulatory bodies across many industries introducing new and stricter regulations governing data management.

Unsurprisingly, fewer industries fought with more vigor than the financial services sector, which is now one of the most highly regulated sectors today.

According to Steven Norton, The Wall Street Journal, of all the data breaches investigated by Verizon over the last 10 years, not a single company that handled payment card data was found to be compliant with all 12 Payment Card Industry (PCI) requirements at the time the investigated breach of sensitive data occurred. Compliance has since increased in every area except that of testing. Verizon recommends that organizations should limit the number of places where sensitive data is at risk by implementing procedures such as data masking. ^[1]

Similar recommendations can be found in the results of compliance audits of many of today's data compliance regulations.

Causes of noncompliance

Some root causes contributing to such a high degree of organizational noncompliance with regulations include:

- Sharing of customer data with business partners or market research organizations
- Duplication of production data into development and test environments
- Addition of applications to existing production data where security was moved from the data source layer into the application layer

Publication of noncompliance has resulted in protection of sensitive data in nonproduction environments coming to the forefront of IT tasks in current years.

In order for organizations to meet these challenges and fulfil their compliance obligations, best practices for sensitive data are gaining significance and seeing a greater rate of adoption.

The solution

Sensitive data best practices

Most responsible organizations implement a set of best practices which are generally accepted as being the right way of doing things for a particular industry or business/technology area. Database administrators usually follow a set of DBA best practices, of which comprehensive versions include a subset of practices for managing sensitive data.

Best practices for the securing of sensitive data include:

- Identifying the data owner
- Determining the importance of keeping the data
- Knowing where the data is kept and who has access to it
- Classification of data and the impact of it being stolen or lost
- Removal of data when no longer required
- Only allowing access to data if needed

Fujitsu Enterprise Postgres already implements strong measures for protecting sensitive data. In addition to correct placement (that is, behind an appropriately configured firewall), it also provides:

- Configuration of which servers to allow connections from
- Host based authentication
- Encryption of underlying data with Transparent Data Encryption
- Row-Level Security to restrict access to row data

With the new Data Masking feature in Fujitsu Enterprise Postgres, sensitive data is only made available to people who need it, and even then, that data can be partially obfuscated.



Data Masking policies

Data masking minimizes security risk by enabling user-based confidentiality, altering original data while maintaining its usability. Fujitsu Enterprise Postgres implements data masking using a powerful yet user-friendly policy approach that allows flexibility in the redaction of different types of data to different types of roles.



Data Masking

Data masking provides the ability to obfuscate specific columns or parts of a column while still maintaining the usability of the data.

There are many scenarios where data masking is of benefit to organizations. One of the more common use-cases is partial obfuscation of information such as credit card numbers to staff in outsourced customer service centers, allowing staff to validate the card number with the customer without allowing access to the full personal information.




Figure 1 - Online masking

Another common scenario is allowing testing of new systems with realistic data without exposing sensitive information to testing staff who may not have appropriate security clearance to view such information. And, as realistic data is used, there is no need to change any applications.




Figure 2 - Offline masking

Fujitsu's Data Masking Implementation

The Data Masking feature in Fujitsu Enterprise Postgres has been implemented using a flexible and easy-to-use policy approach. This allows a set of sensitive data policies to be developed for different classifications of data and different classifications of people without getting too entrenched in the complexities of the technology. A policy can then be applied to tables for the different columns that fall under one classification or another.

Once policies have been applied to tables, they can be disabled or enabled as required without having to remove or reapply them.

A number of tables are available for querying data masking policy information to assess the current sensitive data policy state of a database.

Three different types of data masking are available:

- **Full Masking** - A whole column value can be obfuscated with alternate values
- **Partial Masking** - Part of a column value can be obfuscated with alternate values
- **Regular Expression Masking** - The value of a column can be obfuscated via a regular expression statement

Design parameters

The Data Masking feature has been developed with the following design parameters in consideration:

Non-reversible

It should be possible to mask data such that the original sensitive value cannot be derived from it.

Data Masking provides this ability in all three masking types (full, partial and regular expression) by allowing replacement characters to be specified. Where meaningful obfuscated values are required, post-masking processing is easily applied severing any relationship between the original and obfuscated values.



3 types of data redaction available

Our Data Masking increases the safety of your data

- Full masking - Redacts the entire specified column with replacement values
- Partial masking - Masks only portion of the specified column
- Regular expression masking - Applies a regular expression to the column's value

Flexible and easy to use

There are many business reasons for the obfuscation and protection of sensitive data, and each may have its own preferred solution. Fujitsu's goal is to provide architecture and tools that are flexible enough together to provide as wide a selection of quality solutions as possible.

Unfortunately, with flexibility, often comes complexity. Fujitsu's Data Masking feature has managed to avoid this compromise by building on top of a clear architectural design and complementing existing data security features rather than competing with them.

By applying masking policies at the time, the data is accessed, policies can be configured to alter data for specific conditions, such as a particular application role, providing greater flexibility.

A heightened level of security can be applied to the sensitive data stored on disk by utilizing Fujitsu's Transparent Data Encryption.

Policies are easily created and modified. The example below shows how an existing policy can be modified to add a new partial masking column called "cc_num" that stores credit card numbers with only the last four digits to be readable.

```
postgres=# ALTER TABLE emp
postgres=# ADD COLUMN (cc_num text);
ALTER TABLE

postgres=# INSERT INTO emp
postgres=# VALUES (9, 'John Smith', '3453-3454-5343-
3433');
INSERT 0 1

postgres=# select pgx_alter_confidential_policy
postgres=# (table_name      := 'emp'      ,
postgres=# policy_name     := 'policyone' ,
postgres=# action          := 'ADD_COLUMN',
postgres=# column_name    := 'cc_num'   ,
postgres=# function_type   := 'PARTIAL'  ,
postgres=# function_parameters := 'VVVVFVVVVFVVVVFVVV,
postgres=#                VVVV-VVVV-VVVV-VVVV, *, 0, 12');

postgres=# SELECT * from emp;
 id | name      | cc_num
-----
  9 | John Smith | *****-****-3433
```

Maintenance of original representation

Data masking policies are applied at the time the data is accessed. Queried data is modified according to the masking policy before being returned up the query chain.

This allows policies to be disabled and enabled, or even configured to only be applied for specific roles or conditions, so that the original data can still be viewed with appropriate privileges.

Maintenance of referential integrity

As policies are applied at data access time, data integrity is maintained in the source system. When generating the test environments (offline approach), it becomes more of a design issue in the use of data masking policies and post masking processing to maintain referential integrity.

Repeatable representation

It is important when dealing with testing that test cases are reproducible. Therefore, a masking process should be able to reproduce the same masked data each time it is run.

While a generated test database can always be backed up and restored for each test phase, it is often necessary to regenerate the data again from production for a number of reasons. In this situation it is often a requirement to ensure masked values used in previous tests are consistent.

The policy implementation provided by Fujitsu ensures that masking is consistent each time it is applied.

Conclusion

It is imperative that your organization applies best practices for sensitive data, not only to comply with regulations governing data management—in particular, those relating to protecting the privacy of customer information—but to maintain customer trust. Fujitsu's Data Masking feature provides an easy-to-use and flexible method of providing just the right amount of sensitive data to those who need it. That includes safely sharing production data with nonproduction users in development and test environments.

References

1. Norton, Steven. (2015, March 12). "Most Companies Fail Compliance Tests for Payment Data Security: Report" (The Wall Street Journal)

Read more

To learn more about Fujitsu Enterprise Postgres capabilities to realize a thorough data security strategy, we recommend:

- [Transparent Data Encryption white paper](#) 
- [Dedicated Audit Log white paper](#) 
- [High Availability white paper](#) 
- [Fujitsu Enterprise Postgres Resource Center webpage](#) 

Contact us

If you would like to find out more about Fujitsu Enterprise Postgres, contact us at enterprisepostgresql@fujitsu.com.



Our PostgreSQL experts are at hand


Our PostgreSQL experts contribute to the community in a variety of activities, from feature development to code reviews, and test coverage.

If you have any concerns about the security of your data or would like to find out more about Fujitsu Enterprise Postgres, please contact us at enterprisepostgresql@fujitsu.com

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/



Contact

Fujitsu Limited

Email: enterprisepostgresql@fujitsu.com

Website: fast.fujitsu.com

2024-04-02 WW EN

Copyright 2024 Fujitsu Limited. Fujitsu, the Fujitsu logo and Fujitsu brand names are trademarks or registered trademarks of Fujitsu Limited in Japan and other countries. Other company, product and service names may be trademarks or registered trademarks of their respective owners. All rights reserved. No part of this document may be reproduced, stored or transmitted in any form without prior written permission of Fujitsu Australia Software Technology. Fujitsu Australia Software Technology endeavors to ensure the information in this document is correct and fairly stated, but does not accept liability for any errors or omissions