TIZOR™ Enterprise Database Monitoring and Protection

WHITEPAPER

Data Discovery and Risk in the Datacenter

The Data Discovery Challenge

One of the biggest challenges facing IT organizations is pinpointing the location of critical data throughout the enterprise. As businesses grow, data and its use grow exponentially. From personally identifiable information and customer information to trade secrets and data governed by regulations, much of this data is sensitive and critical to the business. It must be identified, monitored, audited, and protected from misuse and theft.

The need to comply with regulations, meet auditor demands, and minimize data risk adds to the challenge. To meet these needs, the focus must shift to the data itself. Traditional security solutions that focus on the external threat are not the answer. Data-focused technology is required.

Understanding where data is located is the foundation of a sound framework for assessing governance and compliance risk. This is why data discovery is a critical component of risk mitigation. Without discovery, organizations cannot gain control of data. While much attention has been paid to discovering data on laptops and other end point devices, a bigger problem looms. Data in the data center is not all accounted for. The same principles of data discovery must be considered inside the data center itself. This paper will focus on data center data discovery.

UNDERSTANDING DATA DISCOVERY

Typically, sensitive information is scattered across the enterprise. Customer information (credit card numbers, Social Security numbers), employee information (SSNs, addresses, salary, and medical information), and operational information (financial data, IP) can reside in databases and file shares hidden and unprotected in the data center. Regulations such as the Payment Card Industry Standard (PCI), Gramm-Leach-Bliley (GLBA), Sarbanes Oxley (SOX), and the Personal Data Privacy Act of 2007 require companies to protect data determined to be private. All of this data is typically stored in a database in the data center.

However, over the years, the data center has grown to include potentially hundreds or thousands of database servers that store this sensitive information.   These servers can be scattered across the globe. Add to that the fact that developers and quality assurance testers create their own databases with sensitive information and we know have a situation where most companies simply do not know where all their data is inside the data center.

This lack of visibility into critical data assets leaves companies exposed to significant risks such as data theft, data breaches, and unapproved data access.

The ability to understand where data stores are located, what is in those stores, and who has access to stored data is a critical step in understanding data risk and achieving data governance and compliance. The ability to identify data and determine where it is located, whether it is in-use or at-rest, allows companies to assess the effectiveness of data classification policies and procedures. Additionally, visibility into data used for application development and testing is often uncovered in data discovery initiatives. Data discovery allows an organization to control data migration and replication for both application developers and testing needs.

The Verizon Report

During 2008, the Verizon Business Risk Team published a report called "2008 Data Breach Investigations Report". This report was compiled from information from over 500 forensic engagements handled by the Verizon Business Investigative Response team over a 4 year period. This is the definitive report in the industry on the causes of data breaches.

The findings in this report are eye opening. Consider these statistics:

- 66% of breaches involved a system storing data that the organization did not know existed on that system
- 27% of the breaches involved a system that had unknown network connections
- 10% of the breaches involved a system that had unknown accounts or privileges
- 7% of the breaches involved a system unknown to the organization.

Further, when the study looked at the type of data being breached, they found that over 84% of the time it involved payment card/credit card or personal information. And while most attention is paid to external hackers, the biggest risk for data breaches is from trusted insiders.

The implications of this study are significant:

- Compromises are overwhelmingly happening to core database systems in the data center
- You can't protect what you don't know about
- Insiders are the biggest threat

Clearly, companies need to start by getting a solid inventory in place of all data assets.

BEST PRACTICES FOR DATA DISCOVERY AND RISK MANAGEMENT

Data risk can pose significant business problems and it needs to be controlled. Fortunately, understanding and managing business risk are well-understood disciplines in most corporations. The same disciplines need to be applied to data risk.

Risk management requires knowledge about data assets--where they are, what is happening to them, what bad things might happen to them and, most importantly, the costs associated with the bad things that could happen to them. For companies with stable assets, this is built into standard operations. For companies with changing assets and evolving centers of value, the challenge is to become aware of the shifts and deal with them as quickly as possible. The repercussions of not dealing with changing assets could severely damage a business. Data is one of those changing assets.

To this end, there are three key areas of data risk that organizations must control. These risks fall into the category of "unknown unknowns" --unknown data stores, unknown user access, and unknown location of sensitive data. Simply, organizations must:

1. Discover Sensitive Data

Discovery is the first step in uncovering data risk. It begins with finding and identifying critical data assets in the network, determining what data stores exist, and finding information inside of those data stores. Once this is accomplished, a data risk assessment becomes much simpler.

2. Assess Data Activity Risk

In order to understand data risk an organization must know who has access to which data. Visibility into data usage and the associated risks is essential for developing the appropriate compliance and security strategy. This includes identifying how data is being used and which users and applications are accessing data from where and when. This assessment must encompass all applications, users, and processes relating to the access of sensitive data.

3. Ensure Data Compliance

To comply with a variety of regulations such PCI, SOX, HIPAA, GLBA and others, dozens of data protection requirements must be addressed. Compliance with these regulations includes the ability to provide regular and detailed reports that address the information requirements of outside assessors and internal stakeholders.

THE DATA DISCOVERY MARKET LANDSCAPE

There are many solutions in the market today that claim to help identify data assets. Lets take a closer look at some of these tools:

- Data Classification tools – these tools will help "categorize" data, primarily for the purpose of tiered storage. Typically these tools are focused on finding unstructured data on a variety of file shares. This data can be categorized by content, file type, useage and many other variables. Once categorized the tools can also help identify the most appropriate storage solution.
- Data Leak Prevention Tools – these tools are designed to prevent sensitive data from leaving the enterprise over email, instant message, or illegal copying of data to removable devices. The discovery component of these tools scans file shares, identifies different types of unstructured data, and then classifies it. Polices are then written to monitor the flow of this data and stop unapproved activity.
- Fileshare crawlers – there are a number of free tools on the market that will simply crawl a file share looking for different file types and create an inventory list of what they find.

While these tools do server a useful purpose, they are all missing a critical area of the enterprise – the data center. None of these tools help identify database servers, database systems or discover sensitive content in the databases. And as we know from the Verizon report, this creates tremendous risk in the enterprise.

THE SOLUTION – DATABASE DISCOVERY IN THE DATA CENTER

Database activity monitoring (DAM) solutions mitigate data risk by discovering critical data in the data center; intelligently monitoring and analyzing the activity that affects it; providing detailed auditing trails; and reporting on all user access to data stored in open systems such as databases, fileservers, and legacy applications such as mainframe systems.

Database auditing and monitoring helps assure core data by addressing four critically important data issues:

- Data Discovery – Where is sensitive data stored in the data center
- Data Activity Monitoring – How, where, what, when, and by whom is data being accessed?
- Data Risk Assessment – If data risks are detected, can this risk be managed in an automated way to assure data for compliance and governance?
- Data Risk Management - Can stored data be protected from data theft, including data theft by authorized users?

HOW DATA ACTIVITY MONITORING WORKS

An ideal data auditing and monitoring solution cost-effectively answers key questions about stored data such as: where is the data, who is accessing it, can data governance/compliance be achieved, and can data breaches be prevented? In order to answer these questions, a DAM solution must possess four key capabilities:

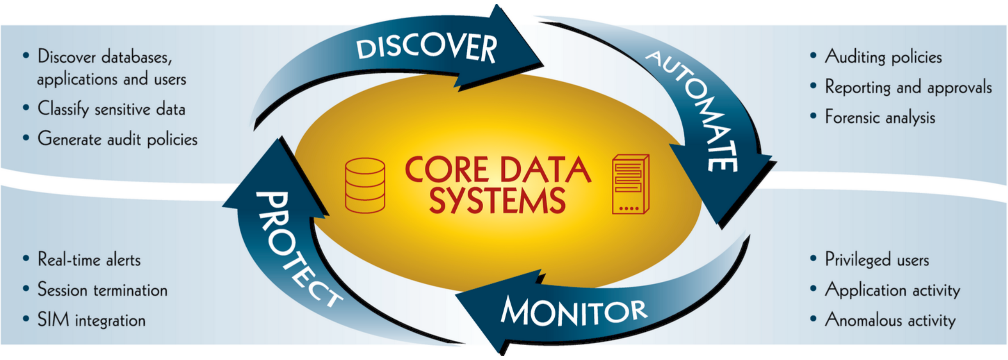

- Discover databases, applications and users
- Classify sensitive data
- Generate audit policies

- Auditing policies
- Reporting and approvals
- Forensic analysis

- Real-time alerts
- Session termination
- SIM integration

- Privileged users
- Application activity
- Anomalous activity

1. Discover

Customer information (credit card numbers, Social Security numbers), employee information (SSNs, addresses, salary, and medical information), and operational information (financial data, IP) must be located in order to protect it. A data auditing and monitoring solution must pro-actively find sensitive data-at-rest and data-in-motion, then classify it. This capability is critical in order to mitigate risk and identify gaps in compliance initiatives. Discovery and classification of data also provides insight into what policies need to be implementing for a sound data activitiy monitoring project to be successful.

2. Automate

Automation of the data activity monitoring system provides significant reduction of cost to compliance and governance as opposed to following a manual process. A sound framework of auditing policies provide the ability to only audit what is required in an automated way across data stores. Implementing an automated workflow process to receive, review and approve automatically generated reports ensures compliance, and minimizes efforts to manually perform these tasks. Lastly providing real-time alerting of critical events and automating forensic analysis of that event can minimize risk in real-time.

3. Monitor

Effective monitoring across many data types, users and applications is required to ensure a successful data activity monitoring project to be a success. There are many types of users accessing data, this includes privileged users who have unrestricted access to data, users who access the data through an application and business partners accessing data. All of these users must be monitored to minimize risk and be compliant. In order to effectively monitor users and applications a solution must be able to identify suspicious or anomalous user activity or application behavior in real time.

4. Protect

The ability to take action based on data activity provides the remediation needed to minimize data risk against non-compliance, misuse or theft of data. Key capabilities in protecting critical data assets includes alerting in real time based on policy violation or suspicious activity, terminating a user or application session to the data store when a violation occurs and notifying an enterprise security incident management (SIM) system for broad security event correlation.

TIZOR MANTRA FOR DATA DISCOVERY AND DATA RISK MITIGATION

Mantra provides all of this functionality in a policy-driven, intelligent, scalable, and cost effective appliance. Mantra is transparent and high performance. It continuously monitors and audits all data access traffic to and from database servers and file systems. Mantra reduces business risk and lowers IT costs by enabling the highest level of compliance assurance, data protection, and data privacy.

Mantra helps meet data compliance and governance requirements in addition to database monitoring and protection requirements. Mantra gives you the ability to:

- Automatically discover where sensitive data resides in your databases
- Determine exactly who is doing what with this data
- Determine who has access to the data and their associated entitlements to data
- Audit database and file server traffic from a single appliance without impacting production systems
- Audit all privileged user activity, including DBAs and system administrators, in real time
- Employ real-time analytics to identify anomalous user behavior in time to mitigate cardholder data risk
- Generate a broad range of reports for auditors, managers, executives, and other stakeholders

Mantra monitors, alerts, and reports on all critical data, structured or unstructured, wherever it resides, with no negative impact on systems or processes.

Mantra provides the ability to discovery data-at-rest and data-in-motion, this includes both activity discovery and content discovery:

Activity Discovery

Mantra's activity discovery mechanisms begin operating as soon as the appliance is connected to a network segment. During Activity Discovery, Mantra automatically scans traffic across databases, file shares, and the activity taking place on these systems. In just hours, Mantra has collected an inventory of databases, users, including their activity and presents this information in summary graphs. These graphs are "active" and can be clicked on to automatically generate monitoring policies without writing a single rule.

Database Discovery

Mantra's Discovery capability uses Tizor's patent-pending technology to automatically locate databases and determine precisely where critical data resides. Mantra identifies the location of all databases, tables, columns, and specific types or classes of data. This enables companies to build an inventory of data assets and locations as well as identify potential risks.

Content Discovery

Mantra's discovery can quickly locate and classify data including specific kinds of regulated data such as Social Security numbers and credit card numbers. Mantra helps companies classify data and apply the most appropriate polices based on the type of data. For example, a class of "confidential" can be defined as any PCI related data that includes a credit card number.

User and Role Discovery

Mantra provides the ability to answer the question, "Who has access to what?" This lack of understanding who has access to the data stores creates a considerable amount of risk, such as finding dormant users and priveleges. Mantra's user discovery can not only tell you who can access the database but it can provide detailed reports displaying what roles each user is assigned. Lastly, Mantra can also provide detailed role reports displaying the exact priveleges entitled to a role.

Mantra advantages include:

Deploys Faster : Mantra makes it easy to deploy an auditing and monitoring solution. Every aspect of the product has been designed for ease of use, productivity, and speed of deployment. Mantra deploys faster and requires fewer resources than other solutions.

Discovers Data : Mantra can determine where sensitive cardholder data resides and how it may be vulnerable to theft and misuse. Mantra identifies the location of all databases, tables, columns, and specific types or classes of data.

Most Intelligent : Event capture, analysis, and storage are rule-driven with pioneering real-time filtering, forensics and analytics—including patent-pending Behavioral Fingerprinting® technology.

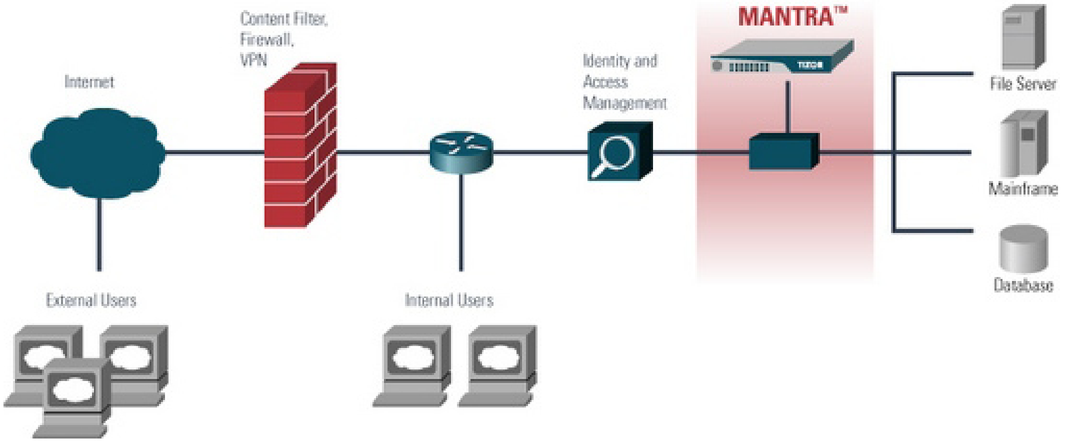

Most Scalable : Mantra was architected specifically to meet the high-performance requirements of the largest enterprise data centers. Auditing over 50,000 transactions per second with no dropped packets, Mantra captures all critical data activity with no impact on networks, databases, or file systems.

Broadest Coverage : Mantra provides the most comprehensive data monitoring coverage today with support for relational databases, file servers and mainframe applications in a wide variety of current and legacy systems – coverage across the largest, most diverse data centers.

Content Scanning : Mantra can scan database and file server traffic for specific data patterns that may represent sensitive data, such as credit card numbers, SSNs, or other site-specific data items. Content scanning can be combined with Tizor's Behavioral Fingerprinting® technology to detect suspicious user activity in real time.

Three-way auditing : Mantra offers real-time, policy-based, agent-less auditing of network traffic plus a choice of agent-less or agent-based local auditing—depending on what best suits your company's local auditing and privileged user monitoring needs.

English-like Policy Language : Pre-defined PCI policies come with Mantra, but if they are needed, custom policies for PCI and other compliance regulations are easy to create and deploy—without DBA or programming skills.

Workflow and Reporting : Pre-defined PCI reports are built in. Custom reports are simple to create and schedule. Automated review and approval functionality makes it easy to generate and manage reports for a wide variety of stakeholders including PCI auditors.

CONCLUSION

Whether it takes the form of intellectual property, health care, financial, credit card or customer information, data is the lifeblood of the enterprise. Data auditing and monitoring is a critical technology for data protection, security and compliance.

Using Tizor's Mantra, organizations can gain unprecedented control over the sensitive data in their care. From locating sensitive data and identifying data risk to obtaining ongoing intelligence about how data is being used, Mantra data auditing and monitoring is a powerful and highly cost-effective tool for mitigating the most costly and damaging forms of data risk..

About Tizor

Tizor provides the world's largest companies with the only enterprise Data Monitoring and Protection solutions capable of auditing and reporting on all critical data activity across the enterprise data center—databases, file servers, and mainframe applications—for compliance assurance, data protection, and theft detection. Tizor's Mantra solutions enable the highest level of compliance assurance, data security, and privacy by providing a complete life cycle of intelligent data auditing capabilities including data discovery, audit reporting, theft detection, real-time alerting, and data protection. Founded by former Bell Labs researchers and headquartered in Maynard, Massachusetts, Tizor is led by industry veterans in networking, security, and software and systems management. More information on Tizor can be found at www.tizor.com.