imperva
a Thales company

DATASHEET

Imperva Data Risk Analytics

Detect data threats before they become security incidents or breaches

Protecting sensitive data is hard for enterprise security groups with limited resources and tools. Often it is the group's tools themselves that make data breach detection so difficult. Tools that cannot properly contextualize alerts overwhelm security staff with an avalanche of mostly "false positives," making it very hard to know what to do or even where to begin.

Organizations need advanced data risk analytics to eliminate all the noise and help security staff gain actionable threat insights to accelerate risk mitigation and breach detection.

Reduce false positives through data context

Imperva Data Risk Analytics, a key capability of Imperva Data Security Fabric (DSF), helps identify data breach threats without all the noise. Typical User and Entity Behavior Analytics (UEBA) tools just focus on network or system access anomalies such as login and logout. Data Risk Analytics takes into account what data users access, the user's roles, whether the data is sensitive or not, and what the user does with it. By correlating all of this event information, Data Risk Analytics contextually determines if an activity is simply an anomaly without risk, or an actual serious threat to sensitive data before generating an alert. This filters out false positives and enables teams to act only on higher-risk incidents that should be further investigated.

Actionable insights make staff more effective

Investigating data threats through the information that a typical UEBA solution (such as a SIEM) provides often requires pre-knowledge of the accessed data set, or deep knowledge of data access languages like Structured Query Language (SQL), to know if any sensitive data has been misused or if users are accessing data inappropriately.

Imperva Data Risk Analytics identifies abnormal user behavior that can lead to bad practices, hostile intrusions, and data compromise. Data Risk Analytics translates the aforementioned technical events into plain language that IT operations teams and security staff members can immediately understand. Data Risk Analytics provides an intuitive dashboard that provides a prioritized incident summary of questionable events that anyone can click through, which in turn provides a full description of the threat with actionable intelligence for remediation.

HIGHLIGHTS

Reduce false positives and prioritize what matters most

Gain actionable insights that make staff more efficient and effective

Detect complex or evasive behavior that indicate threats such as privilege abuse or compromised accounts

Achieve Fast Time to Value from features that work right out of the box

Unravel complex threat behavior and prioritize what matters most

Imperva Data Risk Analytics uses unique, purpose-built data threat detection techniques proven to identify data-centric threats many security tools miss. For example, one of the many techniques looks for activity such as a user abusing a service account to access data, a potential sign the account is compromised or someone is trying to conceal their identity for malicious intent.

Data Risk Analytics prioritizes critical incidents by applying grouping and scoring algorithms that factor in variables such as sensitive data type, privileged account, amount of data involved and more. If multiple incidents are related (e.g. they are all associated with the same user account or multiple users are abusing the same service account), they will be grouped into one issue. Security staff are prominently shown the high-risk incidents, and false positive noise is suppressed.

Active Attack Detection

During an attack, minutes count, and that's why Imperva Data Risk Analytics includes two sets of Data Risk Detection Sets:

1. Active Attacks
2. Data Risk Activities

Leveraging an analysis of exploits observed in large numbers of breaches performed by Imperva Labs, Imperva Data Risk Analytics recognizes known attack exploit behaviors and immediately triggers a critical alert to notify the security team. The types of exploits recognized include

Active attacks on data sources:

- Audit Tampering
- Command Execution
- Credentials Extraction
- Data Exfiltration
- Database Weaponization
- Malware Deployment
- OS File Read
- Privilege Escalation
- Ransomware

Data risk activities:

- Database Access at Non-standard Time
- Database Service Account Abuse
- Excessive Database Record Access
- Excessive Failed Logins
- Excessive Failed Logins from Application Server
- Excessive Multiple Database Access
- Machine Takeover
- Suspicious Application Data Access
- Suspicious Database Command Execution
- Suspicious Dynamic SQL Activity
- Suspicious OS Command Execution
- Suspicious Sensitive System Tables Scan

Leverage AI-driven Data Risk Analytics to quickly convert terabytes of raw data into actionable information

Data Risk Analytics (DRA) provides protection against a wide variety of user-related security threats via statistical models created and configured directly in the Imperva platform. DRA models are designed to detect and flag outlier activity within large datasets and can be configured to generate automatic alerts as needed.

Given the explosion of raw data, DRA has become a critical facility for transforming data into actionable insights. Many of today's security and compliance tools provide little to no capability in this area and as a result often simply deliver the data downstream in the hopes that "other" tools may be able to discern behavioral patterns. While there is clear benefit to enterprise-wide UEBA solutions, DRA enables you to bring significant additional value by applying UEBA engines at the individual tool level in order to more effectively isolate anomalies earlier in the inspection process.

Highlights

- Transforms raw activity data into valuable information via unsupervised learning
- Identifies data-centric threats and activates Playbooks that lock users, change security groups, call other playbooks and more
- Optimizes SOC integration by eliminating alert storms and heavy false positives
- Improves security visibility across all on-premise and cloud-based database sources
- Provides out-of-the-box models for more than twenty outliers
- Offers customizable and DIY UEBA options

Sample threats Imperva DRA detects

- Account abuse
- Account compromise
- Code injection
- Insider threat
- Privilege misuse
- Sentiment analysis

How the DRA works in the Imperva platform

DRA provides protection against a wide variety of user-related security threats via statistical models created and configured directly in the platform.

DRA UEBA models in the platform are designed to detect and flag outlier activity within large datasets and can be configured to generate automatic alerts as needed.

DRA includes a set of preconfigured models that can be used as-is, or cloned and customized to meet the specific data security needs of your organization. DRA UEBA models are differentiated by the calculations **(behavior types)** and **comparison**s they perform. The comparison and behavior types defined during the DRA UEBA model setup process are described below.

Preconfigured DRA UEBA models

Imperva DSF includes a set of pre-configured UEBA model templates designed to provide customers with baseline templates that represent various threats.

These detection models can be used to understand how to build automated logic that analyzes audit data originating from all sources across your data estate.

It is possible to clone and customize these pre-configured models to detect user-related security events tailored to your organization.

Each pre-configured DRA UEBA model represents a different threat vector.

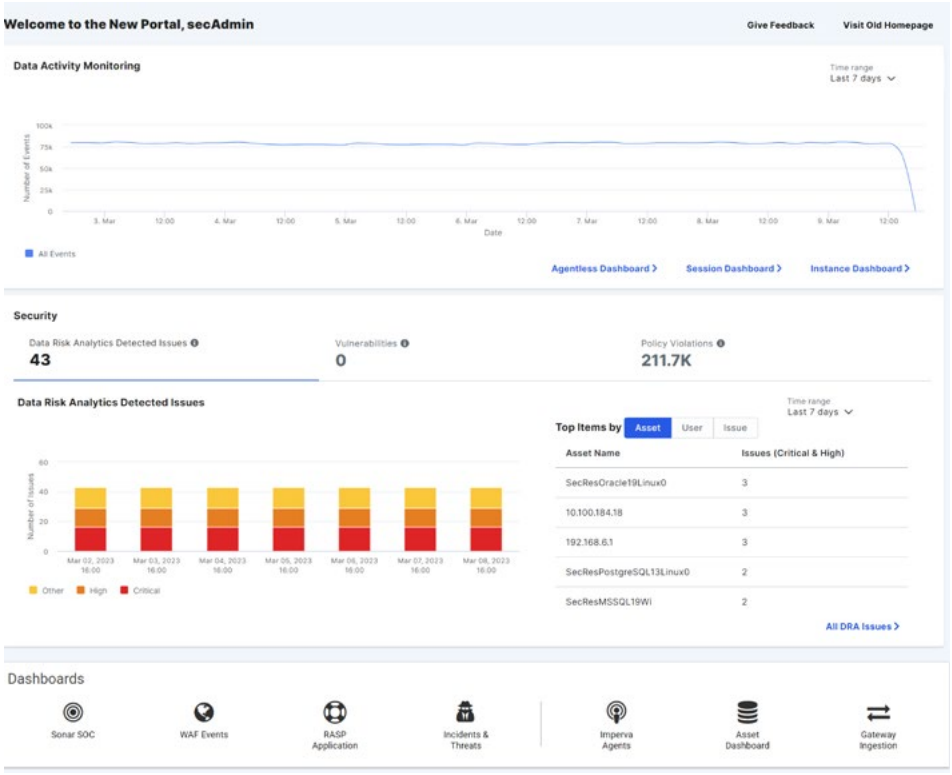

Figure 1: The dashboard provides visibility security teams need to investigate data breach threats.

Comparison types

Historical

Each entity's behavior is compared to its historical behaviors. An entity is only compared to its historical behaviors within the scope defined in the setup process. Behaviors that are older than the scope parameters are not included. When a historical model is first turned on, it assumes (looking back as per the scope) that all behaviors are normal and creates a baseline according to that assumption. New behaviors for existing entities will be compared to that baseline.

Peer

Each entity's behavior is compared to the behavior of its peer group. A peer group is defined by a classifier that has been configured in the "Account Classifier" application. If, for example, a classifier has been configured to classify accounts into "Service account" and "DB admin," DRA will determine whether each entity belongs to the "Service account" or the "DB admin" group, and compare its behavior to all other entities in this group. If DRA is to determine which entity belongs to which group, fields must be mapped from DRA to a classifier during the setup process for peer comparison models.

None

Entities are not compared at all.

Behavior types

Numeric

Input behavior is expected to be numeric. Each entity's z score is calculated where the average and standard deviation are obtained from the entity's reference behavior, as determined by the comparison type. If the absolute value of an entity's z score is above the threshold, it is flagged as an outlier. Numeric calculations consider the following parameters and are defined during setup of the model:

- **Threshold** – An entity will be considered an outlier if the absolute value of its z score is above this number (assuming the Allow Lower / Greater Than Threshold options are unchecked).
- **Allow Lower Than Threshold** – If this option is checked, entities with z scores below the threshold will not be considered outliers.
- **Allow Greater Than Threshold** – If this option is checked, entities with z scores above (threshold) will not be considered outliers.
- **Minimum / Maximum Outlier Value** – If an entity's behavior is below/above this value (respectively), it will not be considered an outlier.

Discrete

The entity's behaviors are collected into a set; when the entity's current behavior set is abnormal as compared to the reference behavior set (determined by the comparison type), it is flagged as an outlier. Discrete calculations consider the following parameters and are defined during setup of the model:

- **Score Type**
  - **Count** – The score is the number of new behaviors.
  - **Percentage** – The score is the number of new behaviors divided by the number of reference behaviors.
- **Whitelist** – A list of behaviors that are never considered abnormal.
- **Threshold** – An entity will be considered an outlier if its score is above this number.

Conditional

The behavior analysis is done by the user in the pipeline builder page. Conditional calculations consider the following parameters and are defined during setup of the model:

- **Score Input Field / Expression** – The field or expression in which a score has been calculated (expected to be numeric).
- **Outlier Input Field / Expression** – The field or expression containing a boolean indicating whether or not the entity is an outlier.

General configurations and definitions

Attack Vector
A name under which different models that analyze the same category of behavior or threat can be grouped.

Behavior Input Field / Behavior Input Expression
A field or expression that contains the behavior being analyzed.

Score
A number representing the relative abnormality of an entity's behavior. Entities who are not outliers will have a score of 0.

Preconfigured DRA UEBA models

Imperva platform includes a set of preconfigured DRA model templates designed to provide customers with baseline templates that represent various threats.

These detection models can be used to understand how to build automated logic that analyzes audit data originating from all sources across your data estate.

It is possible to clone and customize these pre-configured models to detect user-related security events tailored to your organization.

Each preconfigured DRA model represents a different threat vector.

Account abuse

This threat category refers to a broad spectrum of unexpected or suspicious activities by users within an organization, e.g. unusual login activity and unexpected data movement. The preconfigured models for this category are described below.

Excessive Connections
Detects excessive rates of user connections to the database.

Excessive Data Extrusion
Detects unauthorized or otherwise unusual movement of data by users.

Excessive Data Modifications
Detects unusually large amounts of modification to data.

Excessive SQL Error against Sensitive Objects
Detects unusually large amounts of SQL errors against the Sensitive Objects group.

Nonstandard Access Times and Nonstandard Login Times
These models analyze user access and login times respectively; access/login times are flagged as outliers if they occur outside of a user's usual/expected login times.

Users and Verb Categories
Detects anomalous amounts of commands by a user against a particular database.

Account compromise

This threat category refers to suspicious activity wherein a third party (inside or outside your company) attempts to gain control of machines within your organization using existing account credentials, e.g. brute force login attempts. The preconfigured models for this category are described below.

Brute Force Login Attacks
Detects attempts to access accounts via brute force attacks.

Machine Takeover Breakout Attempt
Detects attempts to use a compromised machine within your organization to launch attacks against other machines outside your network.

Machine Takeover Propagation
Detects successful uses of a compromised machine within your organization to launch attacks against other machines outside your network.

Code injection

This threat category refers to activity related to the injection and execution of malicious code into an application. The preconfigured models for this category are described below.

SQL Injection
Detects attempts to access data via SQL injection.

Suspicious Character Usage
Detects suspicious attempts to place meta character(s) into data input.

Suspicious Command Usage / Injection
Detects attempts to access data via vulnerabilities in an application's user input functionality.

Suspicious Data Unload
Detects attempts to access data via data unload.

Insider threat

Suspicious Account Creation
Detects attempts to create accounts with data access privileges to non-existent or unauthorized users.

Suspicious Grants
Detects attempts to grant data access to malicious applications.

Privilege misuse

This threat category refers to the misuse or abuse of a user account's particular privileges. The preconfigured models for this category are described below.

Service Account Abuse
Detects outlier behavior by non-human accounts.

Super Users and Verb Categories
Detects anomalous amounts of commands by a super user against a particular database.

Suspicious Connections
Detects connections from unusual, unexpected or otherwise suspicious endpoints.

Suspicious Data Access
Detects access to specific datasets by unexpected or unauthorized users.

Suspicious Data Change
Detects unusual changes to data.

Integration with Imperva Data Security

The Threats and Outliers dashboard includes various visualizations designed to assist in threat assessment.

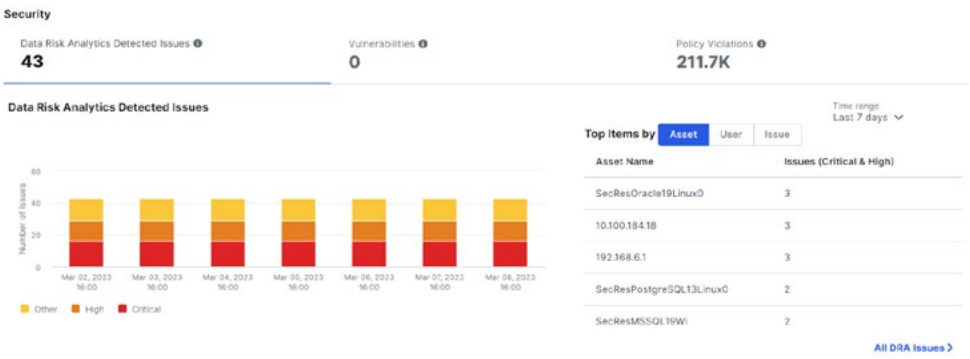

Imperva is an analyst-recognized, **cybersecurity leader** championing the fight to **secure data and applications** wherever they reside.

Enterprise coverage, speed, and scale

To mitigate the risk of data breaches enterprise-wide, you need to be able to detect threats across all your sensitive data repositories on-premises, in the cloud, or across multiple clouds. Human beings just can't do it at the speed and scale required. Imperva Data Risk Analytics seamlessly leverages the reach of Imperva Data Security Fabric (DSF) to access data everywhere. Through automation and machine learning, Imperva Data Risk Analytics uncovers suspicious data access and risky behavior across millions and even billions of data access events that happen across potentially thousands of databases every day in a large, data-driven organization. Over time, the analytics engine continuously learns the details of who the users are, what they typically access, and how they typically use the data, using this contextual behavior baseline to constantly fine tune its accuracy.

Fast Time to Value

Data Risk Analytics is a key component of Imperva Data Security Fabric. It helps security teams detect and pinpoint critical threats to data, prioritizes what matters most, and provides actionable insights allowing you to accelerate threat investigation and response - even if you don't know much about the data - and don't know database languages.

Imperva Data Risk Analytics does not require you to create policies before it can recognize non-compliant or risky behavior. Purpose built threat recognition intelligence comes right out of the box, so you can start seeing the benefits and changes in days, not months. Then it continuously tunes and adapts to changing circumstances. Imperva Data Risk Analytics helps you spot and mitigate data breach risks before they become damaging incidents.

Imperva Data Risk Analytics is part of a holistic Imperva Data Security Fabric solution for all enterprise data assets across on-premises, cloud, hybrid and multi-cloud enterprise environments. To learn more, visit Imperva.com and read about how Imperva Data Security Fabric can help your organization.