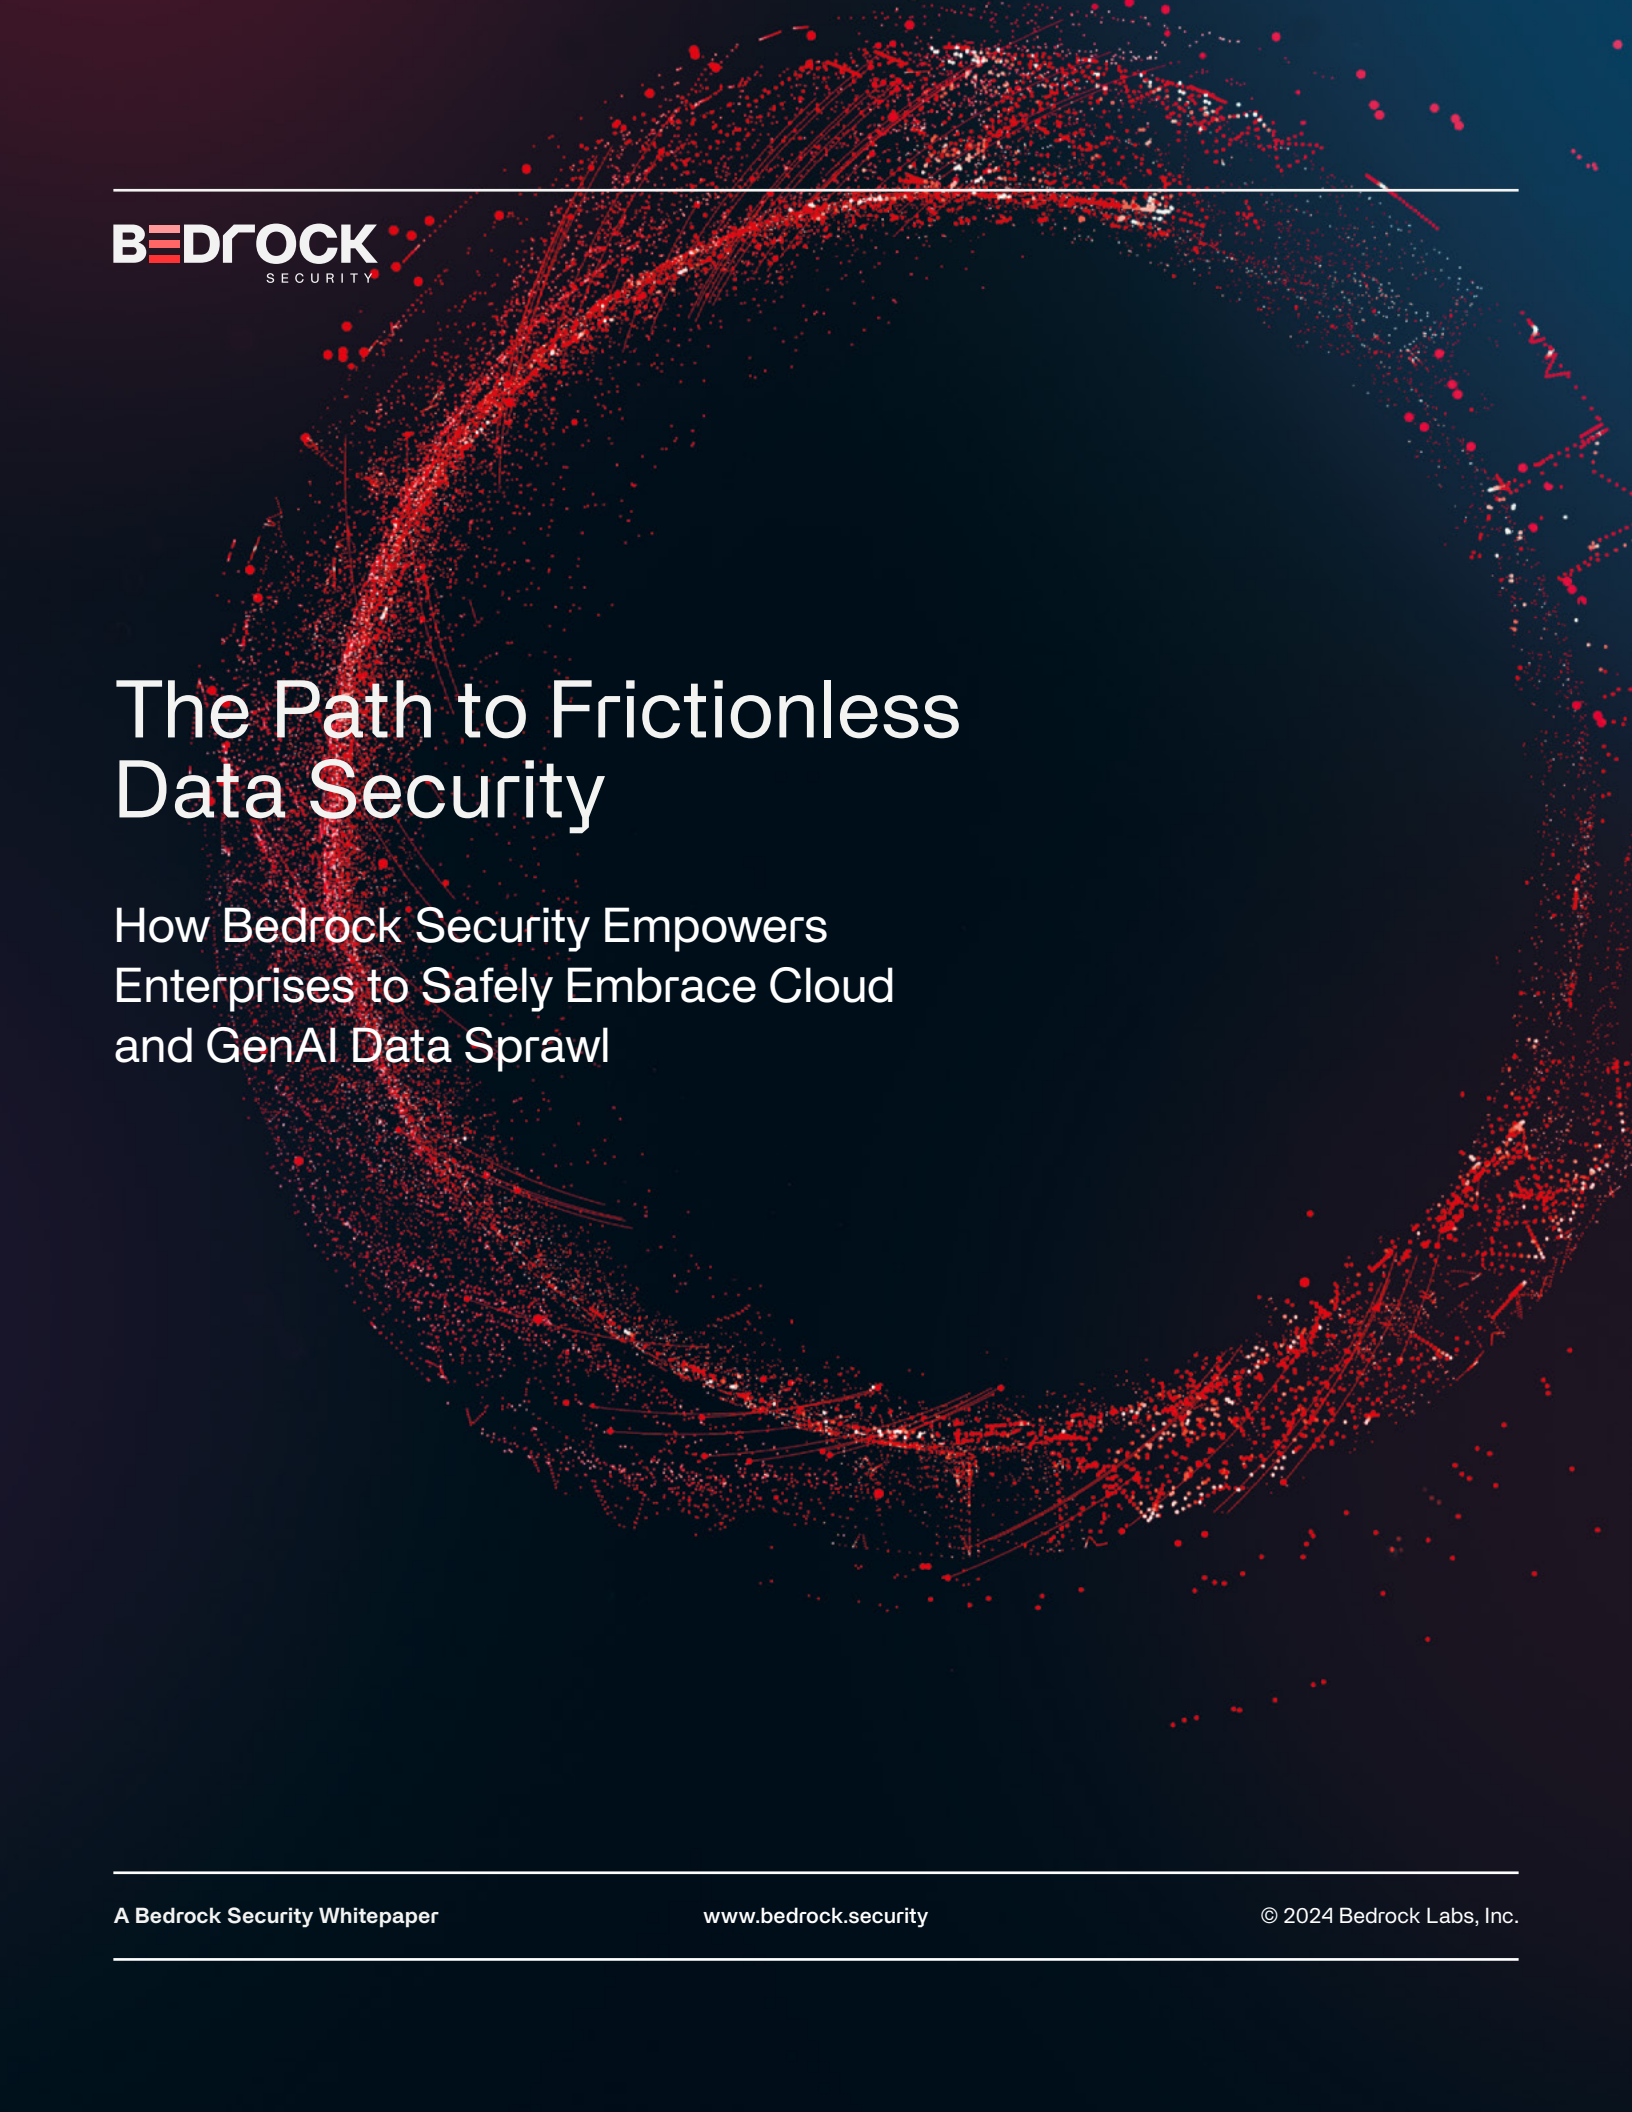

The Path to Frictionless Data Security

How Bedrock Security Empowers Enterprises to Safely Embrace Cloud and GenAI Data Sprawl

Introduction

Data is the lifeblood of modern business and society. But protecting data is becoming increasingly challenging. Even the most sophisticated security teams are struggling to keep up with the rapidly increasing volume, variety, and velocity of enterprise data. The cloud and now generative AI are creating even more complexities.

Critically, conventional data security posture management (DSPM) technologies were not built to address today's evolving demands. As a result, organizations are struggling to ensure that they maintain robust data security and compliance practices without stifling their business operations.

To address the challenges of safeguarding modern data, Bedrock Security has developed a new approach to how organizations protect their business information, providing full visibility and control of their most important data ... automatically, cost-effectively, and at scale.

In this white paper, we will detail how the Bedrock data security platform helps support the development of fundamental and advanced data security practices for meeting today's compliance, risk, and security challenges.

> "98% of global tech executives say that the increasing complexity of managing data across clouds has been a challenge for their businesses."
>
> NetApp 2023 report

The Path to Better Data Security

Bedrock Security's AI Reasoning (AIR) Engine is the foundation to a new approach for modern data security.

It harnesses advanced AI and machine learning technologies that provide a multi-dimensional understanding of your data, going far beyond the common approaches of traditional DSPM vendors using pattern-based rules (RegEx).

Bedrocks's AIR Engine provides rich business context by classifying data based on its contents including such factors as:

- IP awareness
- Similarity to known sensitive data
- Identities that have access
- Compliance and regulatory categories

With Bedrock's AIR Engine, you can quickly and easily understand what data you have, where it lives, where it could travel, its importance to your business, and who has access to which data.

Bedrock goes beyond simply providing visibility to your data. It helps organizations understand the risks to all their data while giving them the tools to quickly remediate any exposures.

Critically, Bedrock not only provides essential visibility for modern data security, it offers unprecedented efficiencies through a highly scalable lightweight architecture, innovative sampling and classification capabilities, context-aware issue prioritization, and an intuitive interface for non-technical users.

> "
>
> **56% of technology executives note that data security processes slow down access to data."**
>
> Immuta 2023 report

As a result, the Bedrock platform delivers the quickest deployment, highest accuracy, fastest data discovery, and the lowest OpEx (compute costs) in the industry for finding, understanding, monitoring, and protecting your data.

In the following chapters, we'll detail how Bedrock's AIR Engine and other innovations bring unprecedented support to organizations looking to master today's growing compliance and security challenges.

The Bedrock Platform: Data Security Architecture

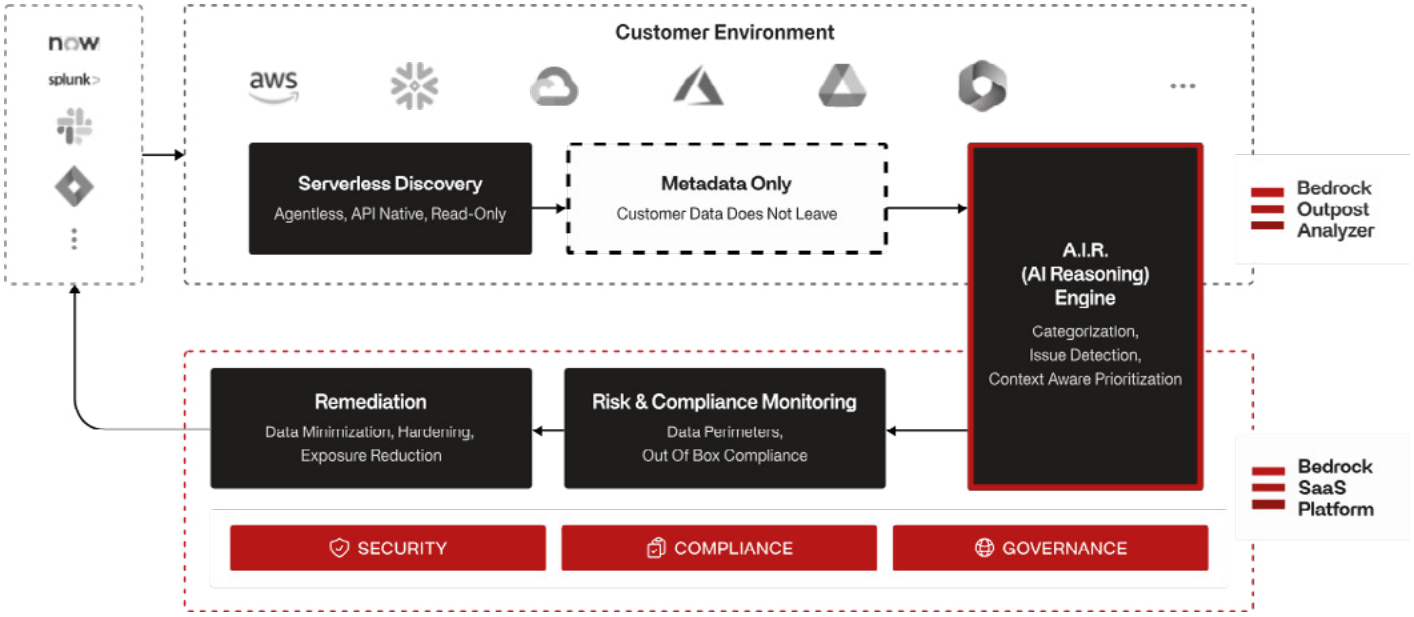

STAGE 1: Understand Your Data

Primary Need

If you are a chief information security officer (CISO), understanding the cyber risks to your organization's data is paramount. At the same time, growing regulatory requirements demand that you have full visibility and control of your data regardless of where it lives.

Essential Capabilities

To address these fundamental requirements for risk assessment and compliance, organizations must be be able to:

- **Gain full data visibility:** Know what data you have (regulatory, IP, GenAI, etc.)
- **Understand data movement:** Know where data is going (within the organization and across clouds, systems, and environments)
- **Understand data access:** Know who has access to which data

How Bedrock Helps

Bedrock provides a fast, accurate, and highly scalable platform for finding, classifying, understanding, and, most importantly, protecting your data.

First, Bedrock is quick and easy to set up. Bedrock deploys in one click, using a lightweight serverless architecture while harnessing an organization's existing APIs to enumerate datastores and classify that data. The Bedrock Outpost data analyzer conducts discovery and classification on-site, sending only metadata back to the Bedrock Cloud, ensuring that Bedrock can never see your organization's critical data.

Once deployed, Bedrock's AIR Engine learns and finds ALL your data. It provides a full inventory of regulated data types, like personal identifiable information (PII), but it also can classify any individual data type, making it easy to define custom data types and even identifying data you don't know about.

Most importantly, the AIR Engine can determine the business context of your data, leveraging technology such as large language models (LLMs). Bedrock is able to model both the type of document (e.g. a credit application) and the likely business owner (e.g. finance department). A Social Security number, for example, could exist in many different kinds of documents such as a W-2 U.S. tax form or within a customer record. These two types of documents require different kinds of access policies with separate business owners.

As an additional benefit, Bedrock's AIR Engine can find and understand data without the help of data owners or lines of businesses — greatly simplifying the data classification process.

Such capabilities are particularly helpful for addressing the challenges of generative AI (GenAI). Sensitive data can easily be identified, classified, and tracked to ensure it never enters into an AI training model or other GenAI environment.

The intelligence of the Bedrock platform also greatly reduces the number of false positives and false negatives by understanding the context of the data (e.g. other data types stored in the same table) while easily recognizing new data or changes in data usage using Bedrock's Fingerprinting capabilities to track data lineage.

The AIR Engine uses Bedrock's patented Adaptive Sampling technology to dramatically lower the operational costs of classifying and controlling access to data. Adaptive Sampling intelligently groups similar data together and focuses attention on datastores containing sensitive data, greatly reducing the computational work required to gain an accurate and comprehensive view of an organization's data landscape.

For example, if an Amazon S3 storage bucket with minimal internal access contains hundreds of values (CSVs) with the same column names and overall structure, Bedrock's Adaptive Sampling will inspect less of the total corpus of files, recognizing that the files are all of the same kind. But for unstructured data found in proximity to highly sensitive information, Bedrock will more thoroughly inspect this data to ensure it is correctly classified.

Bedrock's Benefits

Bedrock's AIR Engine and other data classification technologies provide substantial improvements to traditional data security tools to help organizations better identify security and compliance risks.

- **99 percent accuracy** in classifying data to understand business context (such as GenAI usage)
- **Exceptional scalability** for classifying and securing all your data across infrastructures and platforms
- **Dramatically lower costs**, up to 10X less OpEx (compute costs) than traditional DSPM platforms
- **No trade-offs** between greater data security and budgets

Using Bedrock, a financial services SaaS vendor was able to analyze 2 petabytes of data in 24 hours.

STAGE 2: Defend Your Data

Primary Need

As security teams gain full visibility and understanding of their data ecosystem, the next step is to implement more consistent and comprehensive controls to better protect their organization's data.

Security teams must be able to easily create and enforce data usage policies and data perimeters across today's labyrinth of platforms, applications, and personnel, providing consistent enforcement spanning dynamic hybrid and multi-cloud environments. At the same time, they need to be able to consistently monitor their data and its usage, quickly identify changes or new risks.

Essential Capabilities

To build better protection for their data, organizations need to focus on establishing a robust data detection and response (DDR) practice.

To consistently monitor and enforce data security policies, organizations need to:

- **Understand the state of their data** at all times, across all infrastructures
- **Create effective policies** for protecting the data and keeping it compliant while controlling GenAI data usage
- **Proactively alert security teams** to policy violations and unnecessary data exposure
- **Quickly take remediation steps** to reduce data risks

How Bedrock Helps

Bedrock can not only help organizations know what data they have, it can also tell them what they should do to remediate potential security exposures.

Importantly, Bedrock's AIR Engine understands the nature of an organization's data, helping determine the role and importance of the data to better guide how security teams prioritize

and remediate any issues. With such contextual understanding of data, Bedrock delivers more accurate forensic impact analysis than alternative DSPMs.

And because of Bedrock's highly efficient architecture and Adaptive Sampling technology, security teams can rescan data quickly and cost-effectively without limitations because of computational costs. Freed from such operational and budgetary constraints, organizations can continuously assess usage patterns for new security or compliance risks, always being up-to-date on what data they have and how it is being used.

With a contextual understanding of the data, Bedrock provides accurate analysis of data security threats, ranking issues by the sensitivity of the data involved to help organizations quickly prioritize the most serious security issues.

Bedrock also provides pre-built, best-practices policies for controlling access to data, reducing the impact of a breach, and complying with regulatory requirements. Organizations can also create their own custom policies using plain-language descriptions to support their unique data-handling policies and data structure.

With Bedrock's intuitive enforcement policies and data maps, security and risk teams can easily build Bedrock's patented Trust Boundaries, which put protective perimeters around key datasets by delineating where sensitive data should live and who should have access to it. For example, you can specify that sensitive data should not exist outside of production environments, or that only human resources team members should be able to access HR-related data.

Trust Boundaries are particularly powerful for ensuring protected information, such as regulated data and intellectual property (IP), is not accessed and used in GenAI models.

Bedrock can generate alerts for policy-based violations, as well as monitor for risky or suspicious data usage based on a wide set of behaviors. It looks for anomalies such as unusually high downloads of sensitive data, unusual roles assumed by a user, or unusual locations used to access

data — behaviors that could indicate a breach, attempted breach, or risky behavior that could lead to a breach. Policy alerts can also be easily funneled into a SIEM or ticketing system.

Key Benefits

With the accuracy and highly efficient scaling of Bedrock's data security platform, organizations can greatly increase their cybersecurity posture based on best practices rather than budget limitations.

- **Reduce the mean time to detection** (MTTD) and mean time to respond (MTTR)
- **Eliminate gaps** in visibility and risk changes
- **Maintain continuous compliance** with data regulations
- **Easily create perimeters** around GenAI data

By using Bedrock, a healthcare provider saved 60+ hours a week managing its data compliance.

Stage 3: Reduce Your Exposure

Primary Need

The ultimate goal for any data security team is to progressively reduce the cybersecurity exposure for their organization, especially to protect core IP. At the same time, risk teams must consistently enforce data management policies to ensure compliance with data regulations and laws.

Critically, these efforts need to be as seamless and non-disruptive to business operations as possible.

Essential Capabilities

In order for organizations to systematically reduce their risk exposure, they need to be able to easily implement key remediation tactics to:

- **Minimize the total amount of data** maintained by the organization (retire “stale” data)
- **Limit data access** to “least privilege”
- **Harden sensitive data** via encryption, tokenization, or other protective measures
- **Protect intellectual property** (IP) with the utmost certainty

How Bedrock Helps

The Bedrock AIR Engine takes data classification a step further by providing Dynamic Fingerprinting, which can track data lineage across diverse data stores, including unstructured data. This capability allows security teams to further increase their understanding of how data flows and changes within their organizations.

Combined with Bedrock’s ability to integrate with your infrastructure’s existing controls and APIs, the AIR Engine and Dynamic Fingerprinting makes it possible for organizations to systematically reduce their data security attack surface through risk impact analysis.

Bedrock's remediation playbooks guide security teams in how they can find and remove unnecessary data, reduce the number of people who have access to various datasets, and identify data that should have additional protections. For example, Bedrock can reduce an organization's data exposure by identifying which employees are not actively using a data source. It then can provide a suggested workflow to remove their access.

Bedrock offers powerful data mapping visualizations that show locations, access, and data flows, clearly defining the "blast radius" of a potential breach within a data ecosystem and identifying which accounts and services would be highest risk in the case of a compromise. Because Bedrock's AIR Engine understands the context for data, it is able to provide specific risk analysis and remediation recommendations.

Bedrock also provides customized data types and special controls to help organizations harness the AIR Engine and Dynamic Fingerprinting to protect their unique IP, something that other DSPM platforms can't do.

For example, a life sciences company is using Bedrock's customized classification capability to protect its core IP of 10,000+ gene sequences. The Bedrock platform dynamically tracks file copies and derivative data to ensure any information about the gene sequences stays within the organization.

Such data security capabilities are now essential to helping organizations address the data risks of GenAI.

With Bedrock's scale and visibility, it is easy to track all your GenAI data. And Bedrock's customized data classification makes it easy to recognize and protect critical and sensitive data from use in GenAI environments. Dynamic Fingerprinting also makes it possible to track usage of strategic information and set up alerts, ensuring sensitive data wasn't copied, duplicated or modified in GenAI environments.

Key Benefits

With Bedrock's remediation tools and custom IP identification and protection capabilities, organizations can steadily reduce their risks and make their entire data ecosystem more secure.

- **Reduce the risks of GenAI** causing leakage of your most sensitive data
- **Systematically reduce data access** by identifying non-essential users
- **Identify "stale" data** to reduce total data exposure
- **Implement custom classifications and protections** to ensure utmost security for an organization's most important IP
- **Reduce potential "blast radius"** of a breach by steadily eliminating data exposures and risks

> "
>
> **On average, the data in a large organization comprises 35% dark data, 50% redundant, obsolete, or trivial (ROT) data, and only 16% business critical data.** [1]
>
> Veritas Research

The Path to Frictionless Data Security

STAGE	FUNCTIONAL NEED	HOW BEDROCK HELPS
STAGE 1 Understand Your Data	• Gain full data visibility • Understand data movement • Understand data access	• Provides 99 percent data classification accuracy • Offers exceptional scalability • Reduces OpEx 10X
STAGE 2 Defend Your Data	• Monitor data, at all times • Create effective policies • Proactively alert security teams • Quickly remediate risks	• Reduces MTTD/MTTR • Eliminates visibility and risk gaps • Helps maintain continuous compliance
STAGE 3 Reduce Your Exposure	• Retire “stale” data • Limit access to least privilege • Harden sensitive data • Secure core IP	• Protects against GenAI risks • Guides remediation priorities and tasks • Reduces “blast radius” of potential breaches

Conclusion

Bedrock provides a new approach to today's data security challenges. With Bedrock, organizations no longer have to choose between data security and organizational growth.

With Bedrock, organizations can now create Frictionless Data Security to manage and protect their data to meet modern business needs.

> "
>
> **88% of data leaders believe that data security will become an even higher priority in the next 12 months."** [2]
>
> **Report: State of Data Security, Immuta**

Learn more about how Bedrock Security can help your organization address today's most pressing data security challenges.

www.bedrocksecuity.com

About Bedrock Security

In the data-driven era of cloud and AI, modern businesses demand seamlessly integrated data security measures for effective risk management. Bedrock Security utilizes advanced AI reasoning to ensure accurate risk assessment and response. Overcoming time and OpEx constraints, Bedrock enables continuous data security without compromise. With dynamic categorization, fingerprinting, and policy enforcement, security teams can establish adaptable trust boundaries, seamlessly enabling operations while effectively managing risks to brand, revenue, and reputation. Embrace a new era of cybersecurity with Bedrock: Frictionless Data Security.

Sources

1 Veritas Research (2022), “Managing the Top Pain Points of the Cloud,” https://www.veritas.com/content/dam/www/en_us/documents/infographics/ig_managing-top-pain-pointscloud-us_V1609-1pg.pdf

2 Immuta (2023), “Survey Finds Data Governance and Security Are Top Priorities for 2024, Ahead of AI,” https://www.veritas.com/content/dam/www/en_us/documents/infographics/ig_managingtop-pain-points-cloud-us_V1609-1pg.pdf

