eBook

5 Steps to Secure Cloud Data Governance

Introduction

Today's business environment is flooded with data. From critical product development intellectual property to sensitive customer PII to logistics and sales information, data is coming at us from all directions. And that data is moving through the business in ways that it never could before. In the past, your customer and prospect data may have stayed securely behind a firewall in a company-owned datacenter. But from the moment Salesforce introduced the Software-as-a-Service CRM, that data has been moving into the cloud. And much more has followed.

Now, cloud data platforms like Snowflake or Amazon Redshift offer anyone the ability to host and analyze data on the cloud with just a credit card and a spreadsheet. This has opened a pandora's box of data analysis possibilities that comes with the attendant risks. Luckily, the very technologies that make it easier than ever to extract value from data also allow for improved and simpler avenues to execute data governance and security. We recommend you take these five steps to building your secure cloud data governance ecosystem.

5 Steps:

Identify "sensitive" data across your network

Get granular on who has access and why

Prioritize visibility and understand consumption

Implement data consumption controls

Mitigate risk with easy-to-apply yet powerful data security

Identify “sensitive” data across your network

Almost all the data a company stores is valuable – no one wants plans for a new product or this quarter’s sales data leaked. Even email communications among executives can lead to embarrassments or reputation damage – think the 2014 Sony hack. However, at this point only certain data requires protection by law. Europe’s GDPR and California’s CCPA specify the kinds of individual “personally identifiable information” that must have higher levels of protection and can result in massive fines if leaked. CCPA, for example, can result in civil penalties of $2,500 for each violation or $7,500 for each intentional violation.

That makes it essential that this data be discovered and classified across your company’s ecosystem. It could reside in e-commerce platforms, CRMs like Salesforce, or in ERPs like NetSuite. If possible, this data should be identified before being loaded onto cloud data storage. And, discovery and classification should be automated, include on-premises and cloud locations, within structured and unstructured data formats.

Solutions

- **OneTrust** automates data discovery across your entire IT infrastructure with deep scans of the actual data. It leverages advanced ML-based classification to label and tag data with both out-of-the-box and custom classifiers. Its unique data discovery architecture allows for discovery and classification of all data types across cloud, on-premises, and legacy systems, in unstructured file shares, structured databases, big data storage, SaaS applications and other cloud solutions.
- With **Collibra Data Catalog** users can easily register data sources using out-of-the-box integrations for commonly-used source systems to ingest data and get up and running faster. It also automates the process of adding context to new data with a proprietary algorithm that learns to identify data classes through machine learning.
- **Snowflake’s** cloud data platform can host and analyze data from multiple sources, and its recently announced classification capability automatically detects personally identifiable information (PII) in a given table and annotates it using Snowflake’s tagging framework. ALTR’s integration with Snowflake can provide consumption intelligence, automated governance policy, and protection based on the Snowflake classification.

Get granular on who has access and why

While sensitive data comes with regulatory strings attached, it often has the most potential value. It can provide insights into customer demographics and sales opportunities and is ripe for analysis. That means it needs to be available to those who need it – but ONLY those who need it to do their jobs.

Create your policies not just around "roles," i.e. RBAC (role-based access control), but instead around "purpose." This is a more focused approach to granting data access based on why a user needs the data, rather than simply their spot on the org chart. When you know why they need the data, you can build access control policies to fit just that purpose.

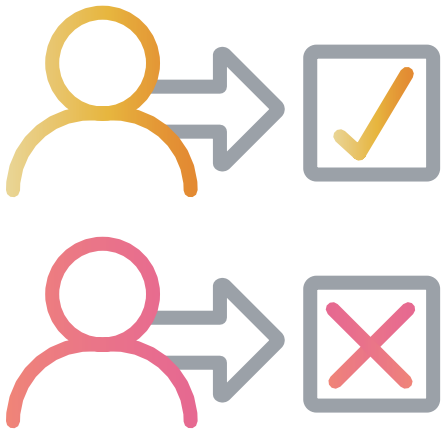

Solutions

- **OneTrust** helps users understand the privacy impact of business use of data and create governance frameworks to minimize risk. It can sync data usage requests with privacy awareness training to ensure compliant access, and OneTrust DataGuidance allows users to research and build applicable governance policies that comply with regulatory obligations.
- Using **ALTR** you can create policies that limit access based on which data is being requested, who is requesting it, the access rate, time of day, day of week, and IP address. ALTR's cloud-based policy engine and management console allow you to control data consumption across multiple cloud and on-premises applications from one central location.
- **Amazon Redshift** supports granular column level security controls to ensure users see only the data they should have access to. Column level access control for local tables means you can control access to individual columns of a table or view by granting / revoking column level privileges to a user or a user-group.

Prioritize visibility and understand consumption

Once you've built your purpose-based access controls, it's important to keep watching: visibility into actual usage will let you know if you need to adjust your policies. Maybe someone initially needed 10,000 records per week to complete a task, but the number has dropped to 2,000. You can adjust your policy to match and reduce your risk. In fact, you could actually start watching consumption before Step 2 so that you have insight into real data usage when you create your purpose-based access controls.

The ability to document data access and usage is also critical to meeting the audit requirements of many privacy regulations. In addition, enterprises that focus only on security, but lack observability into how data is consumed put a heavy burden on their security teams by forcing them to react instead of proactively address threats.

Solutions

- **ALTR Data Intelligence** enables auditable records, including granular details around each request for data so you can better understand where it is, how it's used, and who your top users are.
- **ALTR's Snowflake** integration allows sensitive data to be shared and utilized easily across the organization while usage is monitored and visualized with BI tools.
- Cloud-based BI tools like **Domo, Looker and ThoughtSpot** let you visualize that data to identify typical consumption patterns and quickly spot out-of-normal spikes. This can help you see if your roles and policies are working well, but also understand your overall data requirements and risk.

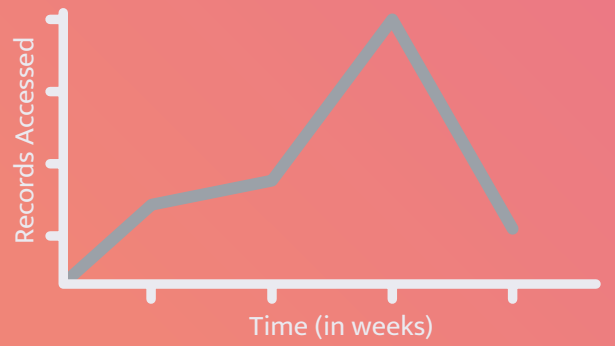

Implement data consumption controls

Banks don't let customers take as much money from ATMs as they like. There are daily limits despite the total amount of cash in the account. They have the original Zero Trust approach to identity – they assume an ATM card could be stolen at any moment. We must have this same attitude toward data consumption. With credentialed access theft a constant threat, we have to assume any "authorized" user could be compromised and act to limit the potential damage. This means enforcing policies via automated data consumption limits.

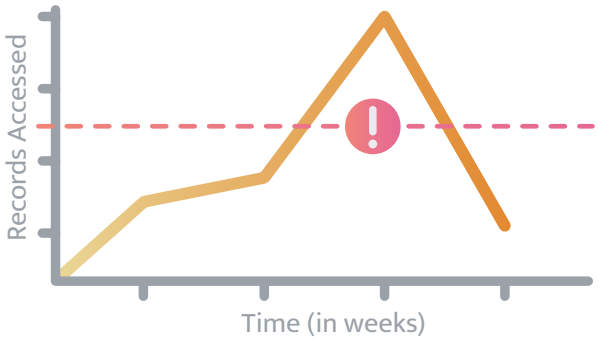

Solutions

- **ALTR's** cloud-based policy engine continually analyzes observed requests for data to determine whether those requests have violated policy. If so, it responds back to connected applications where enforcement is then locally applied. Once policy is violated, an event is triggered within ALTR and the business can be alerted in multiple channels, from email or messaging all the way up to enterprise SIEM/SOAR tools, for response. While the security team investigates, responses can be slowed down or automatically masked to proactively stop data loss before it happens.
- **IBM QRadar Cloud** and **Splunk Cloud** are cloud-based SIEMs that take alerts from ALTR's data security, analyze and respond in real time. These provide an additional layer of defense.

Mitigate risk with easy-to-apply yet powerful data security

Access controls only go so far in protecting sensitive data. The next critical step is to implement real data security. While it may initially seem like data encryption is the “strongest” approach, it can provide a false sense of security while creating roadblocks to the business. Its complexity can slow down authorized data usage and create the risk that encrypted data and key could be stolen at the same time.

The right solution will integrate easily with your data platform and protect data in a way that is transparent to your organization – that can’t affect performance and does not have a key that can be stolen. Data security based on tokenization meets these requirements.

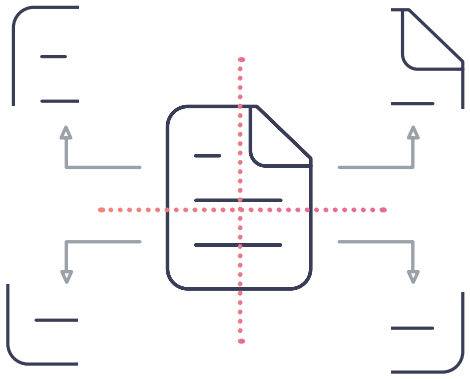

Solutions

- Built for the cloud, **Snowflake** leverages the most sophisticated cloud security technologies available. Snowflake’s government deployments have achieved Federal Risk & Authorization Management Program (FedRAMP) Authorization to Operate (ATO) at the Moderate level. In addition, SOC 2 Type 2, PCI DSS compliance, and support for HIPAA compliance all validate the level of Snowflake security required by industries, and state and federal governments.
- **Amazon Redshift** supports industry-leading security with built-in AWS IAM integration, identity federation for single-sign on (SSO), multi-factor authentication, column-level access control, Amazon Virtual Private Cloud (Amazon VPC), and provides built-in AWS KMS integration to protect your data in transit and at rest.
- **ALTR Data Protection Service** is unique, allowing you to safeguard sensitive datasets without complex and lengthy software installation. You can tokenize structured and unstructured data, there are no keys to maintain and no maps to reduce the security of the data, and using ALTR’s cloud platform, tokenized data can be accessed from anywhere you allow.

Conclusion

ALTR continues to develop relationships with cloud data leaders across the industry. Our goal is to help our customers to get the most from their data by building a secure cloud data ecosystem that allows users to safely share and analyze sensitive data. Our scalable cloud platform uniquely delivers both data governance and security and seamlessly integrates with a wide variety of enterprise tools used to ingest, transform, store, govern, secure, and analyze data. Taking these five steps will set you on the path to secure cloud data governance and data-driven leadership.

5 Steps:

Identify “sensitive” data across your network

Get granular on who has access and why

Prioritize visibility and understand consumption

Implement data consumption controls

Mitigate risk with easy-to-apply yet powerful data security

Complete data control and protection

ALTR simplifies and unifies data governance and security, allowing anyone the ability to confidently store, share, analyze, and use their data. With ALTR, customers gain unparalleled visibility into how sensitive data is used across their organization. This intelligence can be used to create advanced policies to control data access. Through ALTR's cloud platform, customers can implement data governance and security in minutes instead of months.

Get started for free at get.altr.com/free