A framework for success

Tackling the challenges of multicloud security

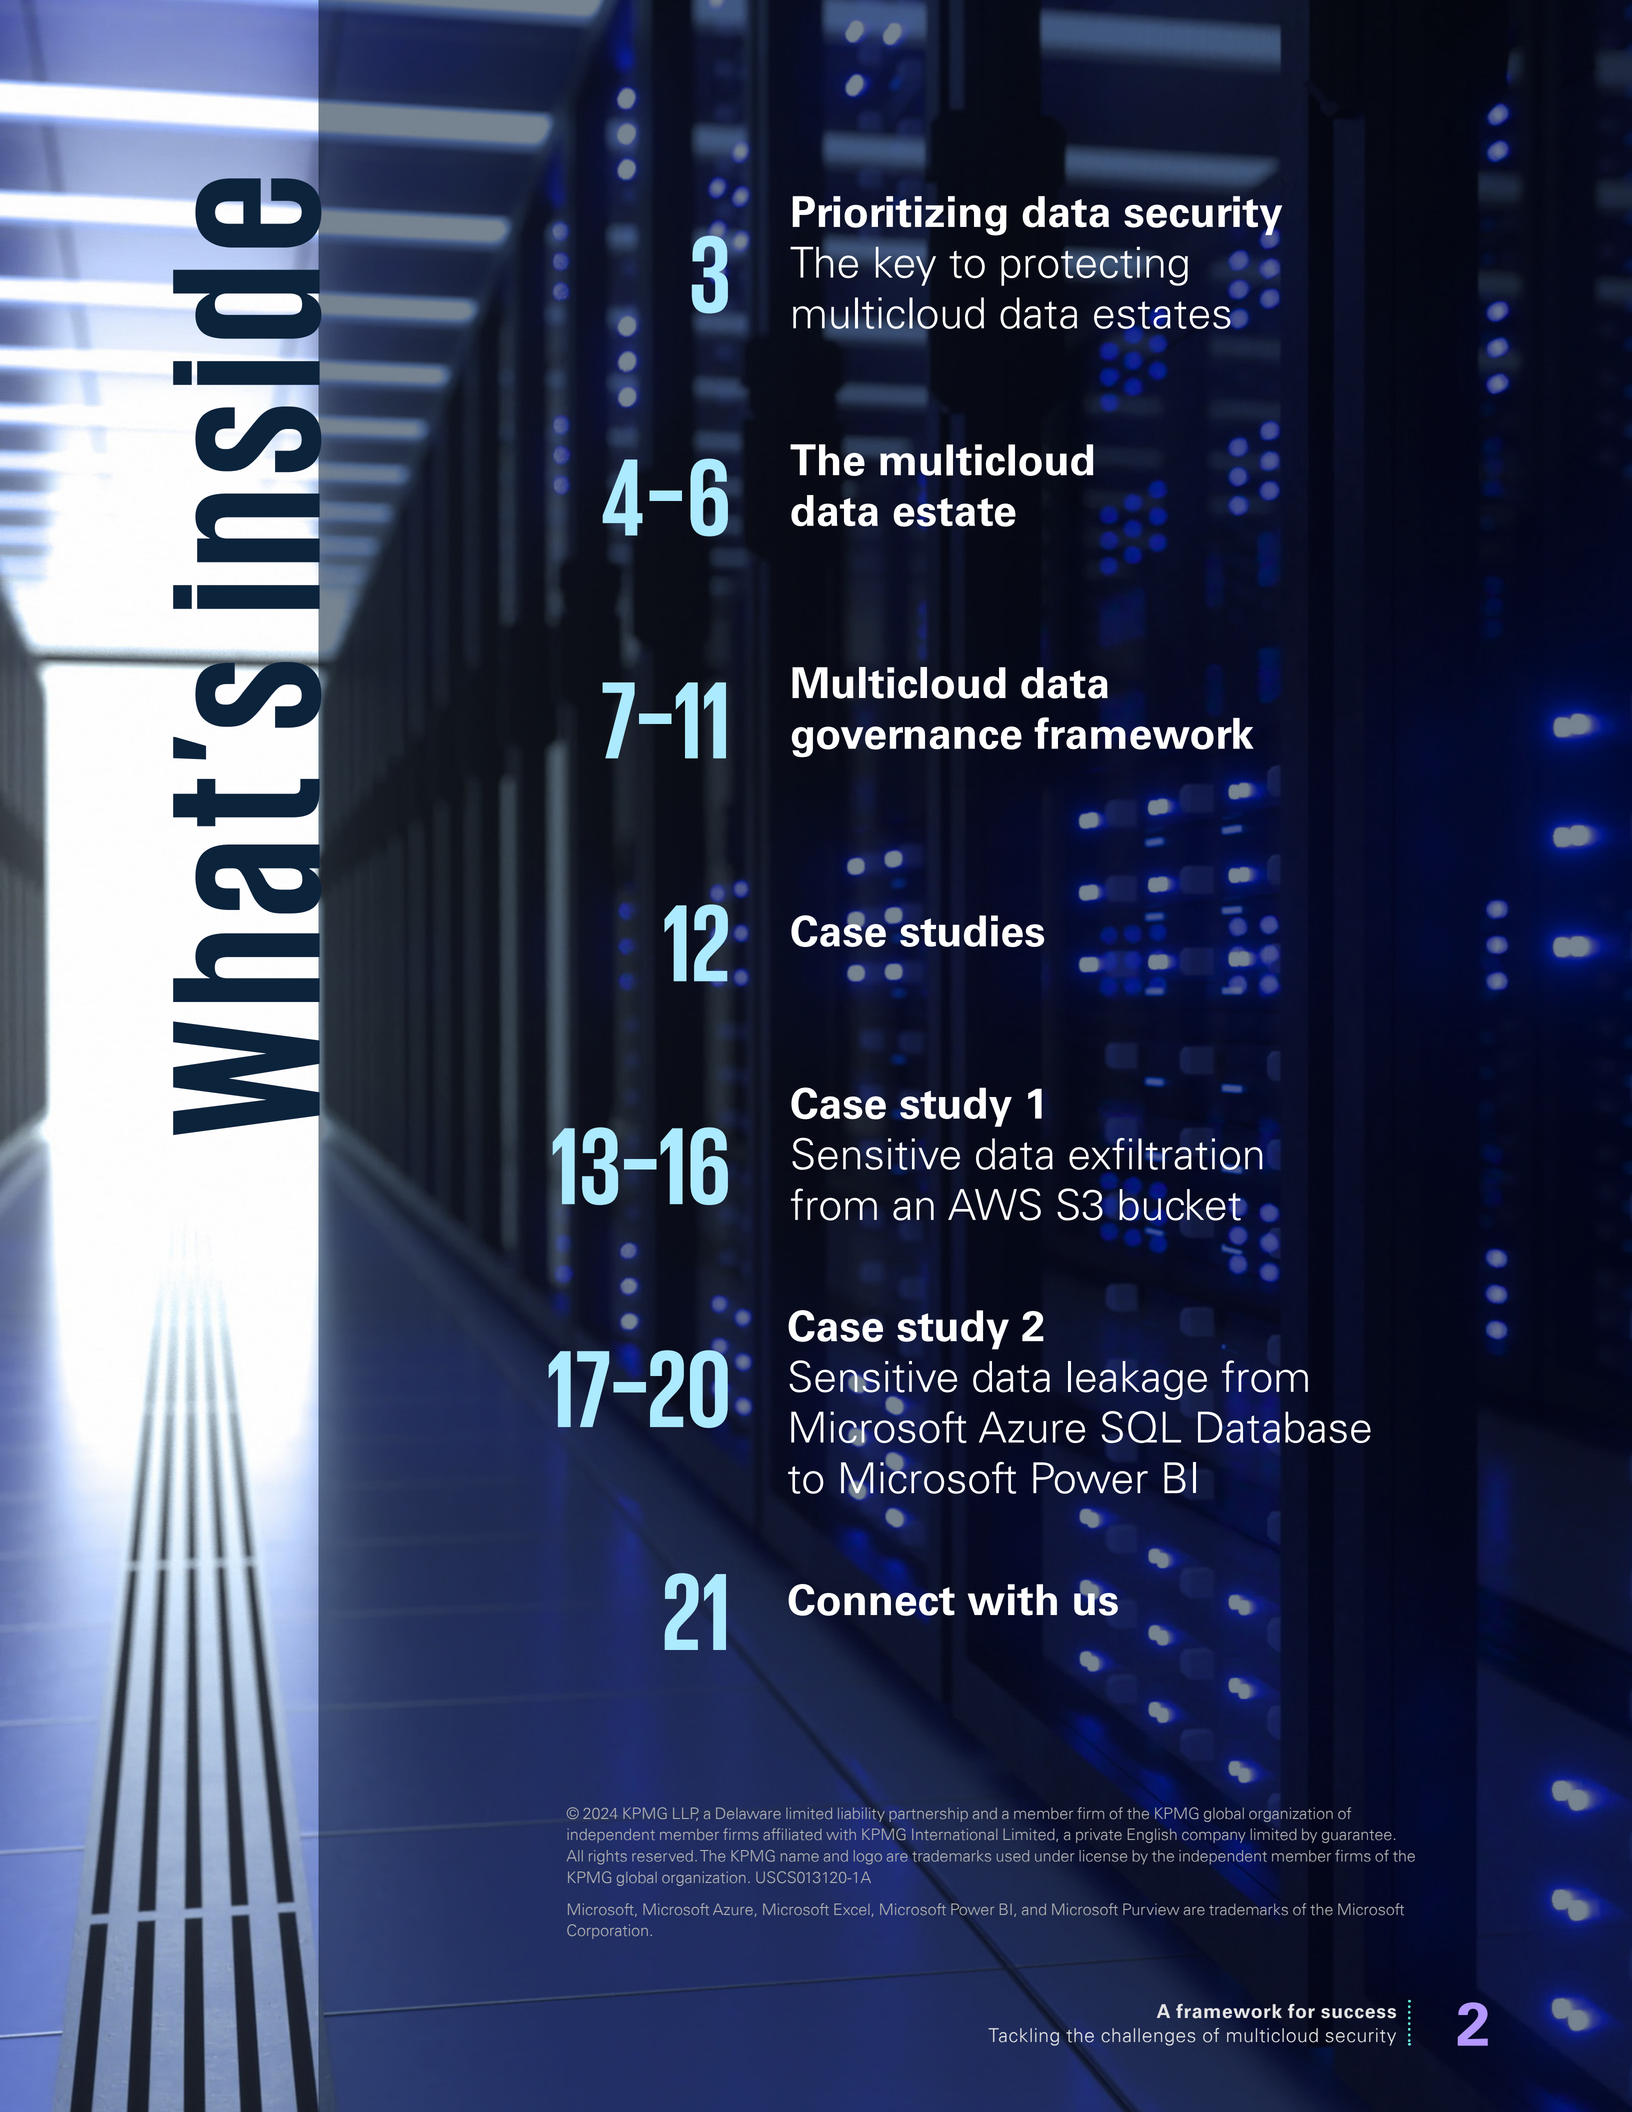

What's inside

Prioritizing data security

The key to protecting multicloud data estates

As organizations move their workloads to the cloud, an increasing number utilize multiple cloud providers to optimize costs and leverage specific features that support their needs. Results from a recent survey of 750 information technology professionals and executive leaders by Flexera[1] indicate that 87 percent of organizations have implemented a multicloud strategy, signaling the trend's continuation.

While the move to multicloud shows clear benefits, security teams experience increasing pressure to address the new risks accompanying this transition. The annual Cost of a Data Breach[2] report published by IBM reported an average data breach cost of $9.8 million in the US. Moreover, the results showed that 82 percent of these breaches involved cloud-stored data, with data spanning multiple cloud environments being the most common occurrence.

Organizations often struggle to address this concern. Identifying sensitive data poses a significant challenge when it exists in a high volume distributed across various cloud environments. Blind spots complicate data security, making it challenging to know where to begin. Security teams that manage to identify sensitive data also face challenges applying security controls effectively. Controls implemented in one environment will not typically transfer when the data moves to another environment, especially after the transformation of structured to unstructured data.

To address these challenges, organizations must prioritize data security and take a multicloud approach. This whitepaper highlights the risks modern multicloud data estates face and provides a framework to address those risks. Simulated case studies then describe how Microsoft Purview's innovative enhancements can implement technical controls that support the framework.

By adopting this framework and utilizing Purview's new technological advancements, organizations can start reducing data security risks within their multicloud data estate.

[1] Flexera.com, Flexera, "2023 State of the Cloud Report" (2023)

[2] IBM.com, IBM Security, "Cost of a Data Breach Report 2023" (2023)

The multicloud data estate

Discover the power of a multicloud data estate, managing company data across different cloud environments while mitigating risks and ensuring compliance.

Harnessing data, **embracing clouds.**

What is a multicloud data estate?

A data estate refers to the underlying infrastructure required to manage data, regardless of the data's location within the organization.[3] This means a multicloud data estate involves managing company data across different cloud environments. A security team should work with key stakeholders to align on their scope for managing the multicloud data estate. This usually involves alignment with data governance, legal, compliance, privacy, and business teams. Many companies will establish or use an existing data governance council consisting of these stakeholders to streamline discussions and decision-making.

[3] Microsoft.com, Microsoft, "The data estate and data modernization in Dynamics 365 - Dynamics 365" (2024)

Platforms in the multicloud data estate

When modernizing a data estate for a multicloud strategy, consider which cloud platforms and services the organization uses, particularly those handling sensitive information. Aim to minimize the number of technologies used to manage data across these platforms and consider both structured and unstructured data types.

Given the widespread usage of Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform, organizations likely use a combination of these providers for various data-centric services, such as data lakes, cloud-based relational databases, and data warehouses. Consider SaaS-based products as well because they increasingly store sensitive data. Ideally, a solution designed for managing a data estate should cover the most critical platforms at the organization.

Strategic intelligence

After addressing scope and relevant platforms, an organization should identify key risks to the data within their multicloud data estate. According to a recent survey by Microsoft,[4] 83 percent of business leaders see managing multicloud environments as their biggest pain point. This pain point becomes apparent given the likelihood of many risks significantly increasing when adopting a multicloud strategy.

The examples provided in the graphic below offer additional context around the key risks facing organizations with a multicloud data estate.

Lack of visibility
Difficult to know where sensitive data resides and to track sharing of the data

Misconfigurations
Incorrect security settings in cloud services or resources

Data leakage
Unintended exposure of sensitive information

Insider threats
High-risk employees with access to cloud resources pose a risk of data theft

Unauthorized access
Attackers or unauthorized users trying to access sensitive data or critical systems

Compliance
Concerns of regulatory violations when storing data in different geographical regions

Reducing risk

The risks outlined above can be overwhelming, but organizations can begin their journey toward a multicloud data estate by establishing a repeatable framework. A framework can set the stage for an efficient, secure, and well-governed data estate. Once this foundation is in place, organizations can better identify and manage risk, streamline processes, and assist with regulatory compliance.

[4] Microsoft.com, Microsoft.com, "Multicloud Security News" (2022)

Multicloud data governance framework

This section provides guidance for key considerations within a multicloud data governance framework. This framework can be used to help reduce the common risks found within most organizations.

1 KNOW YOUR DATA

Understand your data landscape and identify important data across your hybrid environment.

2 PROTECT YOUR DATA

Apply flexible protection actions, including encryption, access restrictions and visual markings.

3 PREVENT DATA LOSS

Detect risky behavior and prevent accidental oversharing of sensitive information.

4 GOVERN YOUR DATA

Automatically retain, delete, and store data and records in a compliant manner.

Powered by an intelligent platform

Unified approach to automatic data classification, policy management, analytics, and APIs

1 Know your data

Understand your data landscape and identify important data across your hybrid environment

Data discovery

The first step should be to conduct scans to understand what data exists, the data's location, and how the data traverses across the environment. Data spread across numerous cloud environments makes this difficult.

Microsoft Purview's Data Map can help accomplish this because out-of-the-box connectors can quickly scan data sources across many cloud providers, services, and resources. Data lineage features also provide visibility into the movement of data.

Data classification

Next, the organization should classify discovered data into sensitive information types. Common examples of sensitive information types include bank account numbers, credit card numbers, and Social Security numbers.

Purview provides the opportunity to classify discovered data using both traditional content-based matching as well as artificial intelligence-driven trainable classifiers.

2 Protect your data

Apply flexible protection actions, including encryption, access restrictions and visual markings

Sensitivity labels

Once classified, the organization should apply sensitivity labels to data based on classification. A label should be persistent and remain applied to data even after transmission from one cloud environment to another. Most organizations will align labels to a four-tier schema, with labels such as Public, General, Confidential, and Highly Confidential.

Purview can apply persistent data labels using a data catalog that will remain with data, even as it traverses across different cloud environments. These labels can also be customized to align with the schema defined in an organization's Information Classification and Handling Standards.

Protection controls

Different protection controls should be applied to data based on the associated sensitivity label. Typical controls include encryption, access controls, and content markings (e.g., headers, footers, watermarks). Similar to the sensitivity label, these protection controls should be persistent. For example, access restrictions should remain enforced even after data moves to a different location.

Purview's newest features provide these types of protection controls, and the latest release provides persistent protection as data moves across key resources in a multicloud environment, such as Microsoft Azure SQL, Microsoft Azure Data Lake Storage, and Amazon S3 buckets.

3 Prevent data loss

Detect risky behavior and prevent accidental oversharing of sensitive information

Data loss prevention (DLP)

Organizations benefit greatly from integrations between their classification, labeling, and DLP tooling. If a DLP can recognize the data classification and label or data, then it can be used within a DLP policy.

Purview's DLP policies provide a built-in ability to use both sensitive information types from data classification and sensitivity labels within DLP policies, which can significantly help organizations prevent unintentional sharing of sensitive information.

Insider risk management (IRM)

High-risk users may conduct suspicious activities that pose significant risk to sensitive data. Organizations should monitor user activity for high-risk users, especially for events related to cloud misconfigurations.

Purview now addresses this through enhanced risk indicators included for Google Drive, Box, Dropbox, GitHub, Microsoft Azure, and AWS. The indicators align to key MITRE ATT&CK tactics, such as defense evasion and exfiltration, and can be used to enhance data security.

4 Govern your data

Automatically retain, delete, and store data and records in a compliant manner

Centralized data management

Organizations will need a streamlined process to manage data across multiple cloud platforms. This process will become more efficient if the organization can quickly discover, classify, and apply security controls—all from a single location.

Purview's release of a unified platform addresses this, combining its previously separate data governance and compliance portals. The visibility provided from scans, such as a data's location, lineage, and classification across various platforms can be used to reduce data risk through the application of appropriate controls or remediation activities.

Records management

Organizations should identify data for archiving and deletion. Removal of data from certain environments where possible reduces the overall attack surface that a malicious attacker can use to obtain sensitive data.

Purview provides capabilities to apply retention labels and can automatically create alerts for data ready to be achieved or deleted.

Case studies

The case studies illustrate how leveraging the described framework can be used to enhance the security of data throughout the lifecycle. The organizations within the case studies will leverage Microsoft Purview to enable the framework.

Case study 1

Sensitive data exfiltration from an AWS S3 bucket

Challenge

A data engineer flagged as a high-risk user creates a new AWS S3 bucket. While attempting to share the data in the bucket with another application, the engineer encounters issues and modifies the S3 bucket policy during troubleshooting. Unfortunately, the modification unintentionally makes the S3 bucket publicly accessible.

The data stored in the bucket lacks classification and labeling, and the S3 bucket ends up storing personally identifiable information (PII) that should have been encrypted and protected by access restrictions. As a result, unauthorized users find the open S3 bucket, access its contents, and download unencrypted PII data to their personal devices, leading to a sensitive data breach that causes severe financial and reputational damage to the organization.

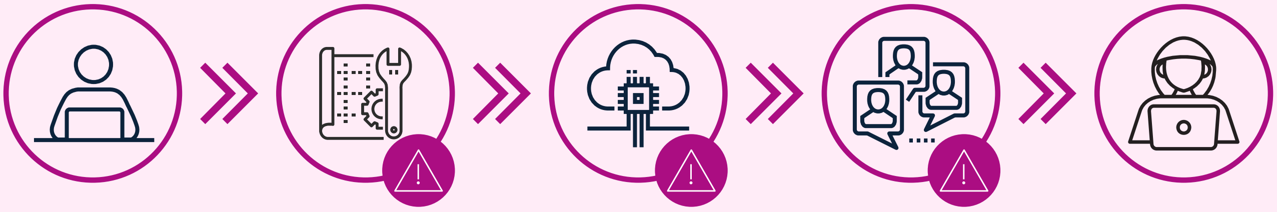

Data Engineer

Configuration Settings

A data engineer flagged as a high-risk user accidentally makes an S3 bucket public

S3 Bucket

PII is unknowingly stored in the S3 bucket

PII Data

Malicious Actors

Malicious actors access, download, and view PII data

How the framework can address this challenge

Four core concepts drive our strategies to address issues:

1 KNOW YOUR DATA

Understand your data landscape and identify important data across your hybrid environment

The organization can use Purview to perform discovery scans on their S3 buckets. Then, the data can be classified into sensitive information types, which would provide the ability to identify PII data elements, including but not limited to full names, addresses, and Social Security numbers.

Microsoft Purview's new unified portal allows both discovery and classification from a single location. Additionally, protecting data using the controls outlined in the next step can be simplified since the controls can be applied within the same portal.

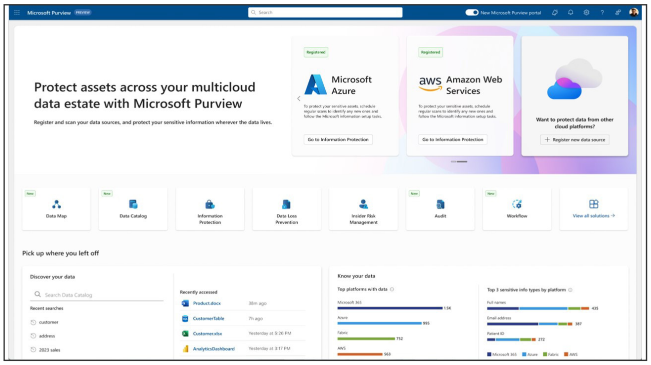

Protect assets across your multicloud data estate with Microsoft Purview

Register and scan your data sources, and protect your sensitive information wherever the data lives.

To protect your sensitive assets, schedule regular scans to identify any new ones and follow the Microsoft information setup tasks.

To protect your sensitive assets, schedule regular scans to identify any new ones and follow the Microsoft information setup tasks.

Want to protect data from other cloud platforms?

Unified platform

Image provided by Microsoft, 2024

2 PROTECT YOUR DATA

Apply flexible protection actions. including encryption, access restrictions and visual markings

The organization can apply "Highly Confidential" sensitivity labels to PII data in the S3 bucket. The "Highly Confidential" label can then be configured to apply access restrictions that only permit internally authenticated and authorized users to access the data. These restrictions will persist, even after exporting the data from the S3 environment. This means that even if unauthorized users managed to download the PII data, they would not be able to access the information.

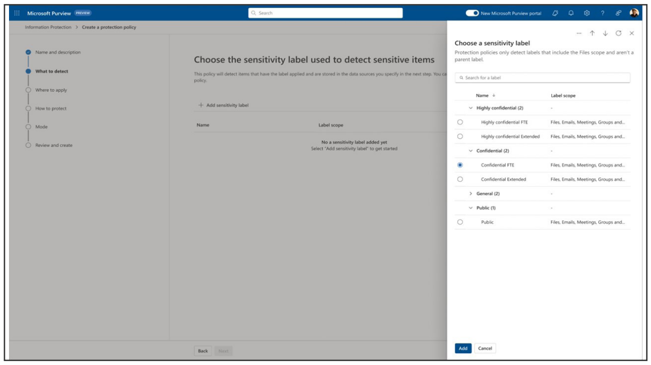

Choose the sensitivity label used to detect sensitive items

This policy will detect items that have the label applied and are stored in the data sources you specify in the next step. You ca policy.

Name	Label scope
No a sensitivity label added yet Select "Add sensitivity label" to get started	

Choose a sensitivity label

Protection policies only detect labels that include the Files scope and aren't a parent label.

Search for a label

	Name ↓	Label scope
	Highly confidential (2)	-
○	Highly confidential FTE	Files, Emails, Meetings, Groups and...
○	Highly confidential Extended	Files, Emails, Meetings, Groups and...
	Confidential (2)	-
●	Confidential FTE	Files, Emails, Meetings, Groups and...
○	Confidential Extended	Files, Emails, Meetings, Groups and...
	General (2)	-
	Public (1)	-
○	Public	Files, Emails, Meetings, Groups and...

Configure sensitivity label

Image provided by Microsoft, 2024

3 PREVENT DATA LOSS

Detect risky behavior and prevent accidental oversharing of sensitive information

The organization properly identified the data engineer as a high-risk user; however, it did not use appropriate alerts to identify risky activities around misconfigurations within its multicloud environment. These misconfigurations ended up exposing its data.

Purview now has enhanced risk indicators within IRM that can detect and send out alerts on high-risk user activities, such as in this scenario when a high-risk user made the Amazon S3 bucket publicly accessible.

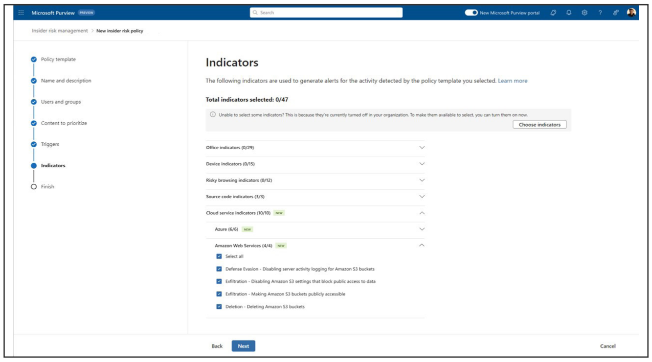

Indicators

The following indicators are used to generate alerts for the activity detected by the policy template you selected. Learn more

Total indicators selected: 0/47

Unable to select some indicators? This is because they're currently turned off in your organization. To make them available to select, you can turn them on now.

Office indicators (0/29)

Device indicators (0/15)

Risky browsing indicators (0/12)

Source code indicators (3/3)

Cloud service indicators (10/10) NEW

Azure (6/6) NEW

Amazon Web Services (4/4) NEW

- ☑ Select all
- ☑ Defense Evasion - Disabling server activity logging for Amazon S3 buckets
- ☑ Exfiltration - Disabling Amazon S3 settings that block public access to data
- ☑ Exfiltration - Making Amazon S3 buckets publicly accessible
- ☑ Deletion - Deleting Amazon S3 buckets

Amazon S3 indicators

Image provided by Microsoft, 2024

4 GOVERN YOUR DATA

Automatically retain, delete, and store data and records in a compliant manner

To maintain ongoing governance, the organization can utilize the results obtained from discovery and classification scans to identify the presence of PII stored in unapproved locations, such as specific S3 buckets not approved for managing data with this level of sensitivity.

With Purview's data catalog providing precise resource information, the security teams can quickly relay the needed details to the information owners to take corrective actions, such as deleting the data.

Resulting in success

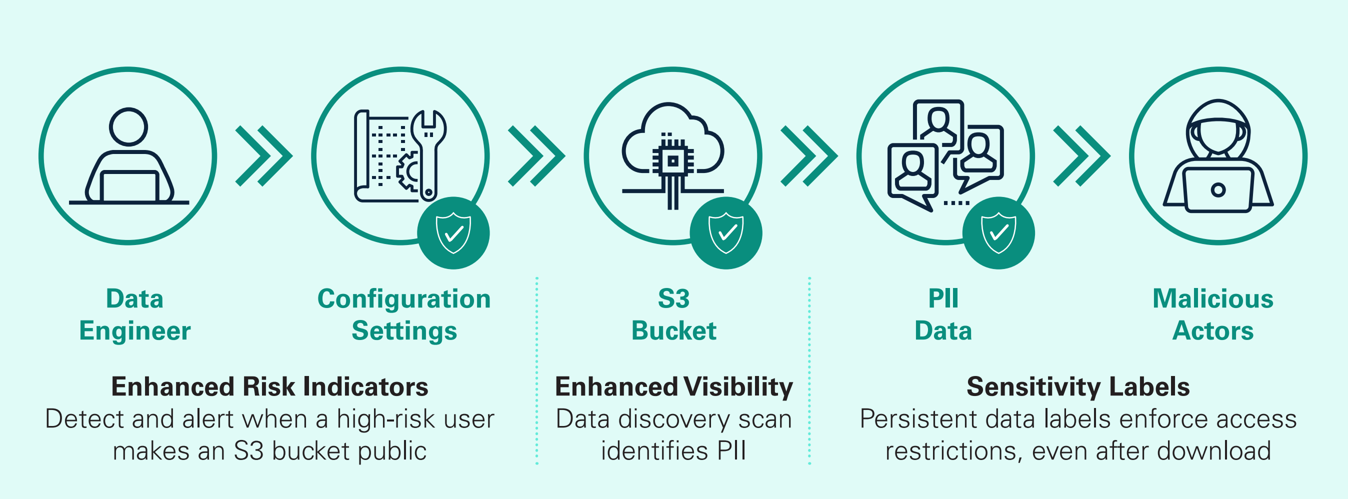

Enhanced Risk Indicators
Detect and alert when a high-risk user makes an S3 bucket public

Enhanced Visibility
Data discovery scan identifies PII

Sensitivity Labels
Persistent data labels enforce access restrictions, even after download

Enhanced risk indicators

By implementing IRM with Purview, high-risk user activity can detect and alert security teams of unintentional misconfigurations in the S3 bucket that expose data, such as making the S3 bucket public.

Enhanced visibility

Purview's unified platform facilitates data discovery and classification within a multicloud data estate, including the identification of PII within S3 buckets, as demonstrated within this example.

Sensitivity labels

In this scenario, even if misconfiguration of the S3 bucket provided public access, the access restrictions applied by the "Highly Confidential" label would prevent unauthorized users from accessing the PII data.

Case study 2

Sensitive data leakage from Microsoft Azure SQL Database to Microsoft PowerBI

Challenge

A banking company stored sensitive client financial data, such as routing and bank account numbers, within an Azure SQL Database. Although the company applied encryption and access restrictions to the sensitive financial data stored within the database, the data lacked persistent controls. As a result, export of the data from the database meant that the access restrictions and encryption no longer applied.

When a business analyst requested the creation of a Power BI dashboard that required data from the Azure SQL Database, the data engineer accidentally included routing and bank account numbers in the data set. As a result, sensitive financial data was exposed to a wide range of stakeholders due to the broad sharing of the dashboard. This put the privacy of clients' financial records at risk and made the financial services company vulnerable to regulatory sanctions, damage to its reputation, and harm to its business relationships.

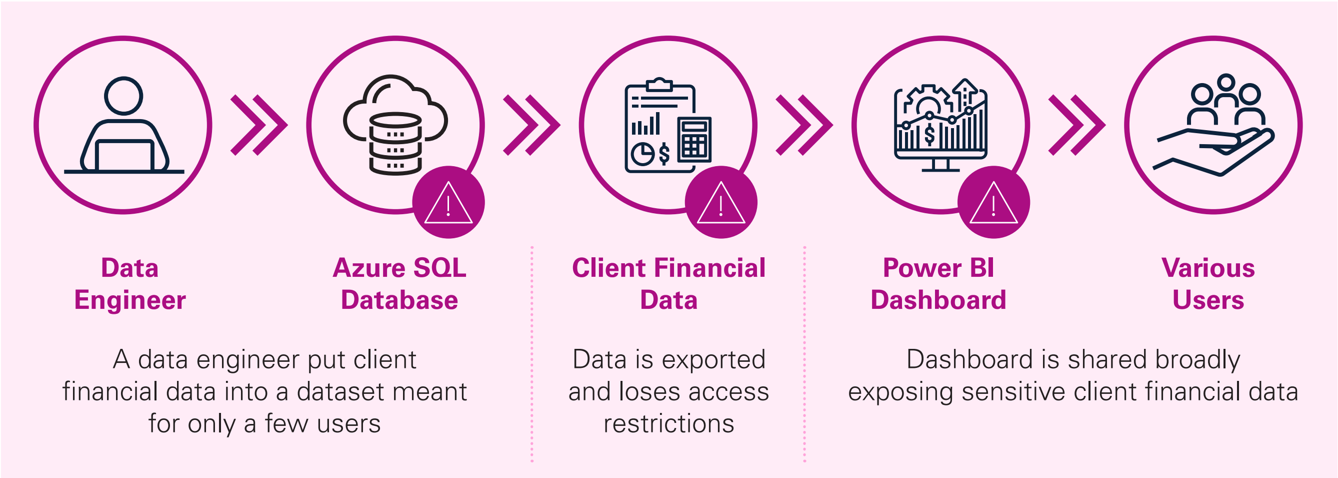

Data Engineer

Azure SQL Database

Client Financial Data

Power BI Dashboard

Various Users

A data engineer put client financial data into a dataset meant for only a few users

Data is exported and loses access restrictions

Dashboard is shared broadly exposing sensitive client financial data

How the framework can address this challenge

Four core concepts drive our strategies to address issues:

1 KNOW YOUR DATA

Understand your data landscape and identify important data across your hybrid environment

Purview can discover and classify data in both Azure SQL Databases and Power BI dashboards. This helps identify sensitive financial data, such as routing and account numbers. Additionally, it shows data lineage, providing details on which downstream Power BI dashboards consume the sensitive data sets from Azure SQL Databases.

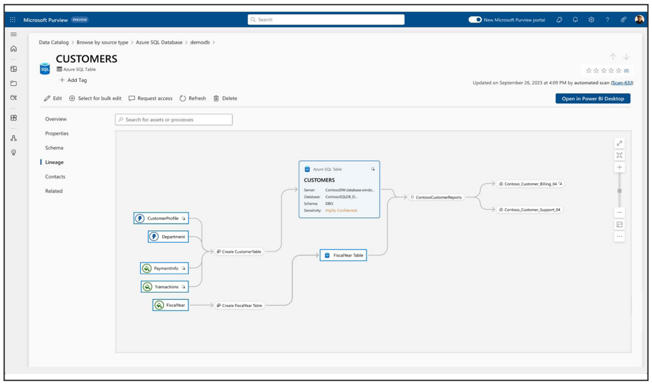

Updated on September 26, 2023 at 4:09 PM by automated scan (Scan-63J)

Azure SQL Table	
CUSTOMERS	
Server:	ContosoDW.database.windo...
Database:	ContosoSQLDB_D...
Schema:	DBO
Sensitivity:	Highly Confidential

Data lineage

Image provided by Microsoft, 2024

2 PROTECT YOUR DATA

Apply flexible protection actions, including encryption, access restrictions and visual markings

The organization can apply sensitivity labels to data fields within the Azure SQL table, such as routing numbers and other data types such as Social Security numbers, credit card numbers, and passport numbers. These labels can enforce encryption and access controls, helping restrict usage of the sensitive data in other locations such as Power BI dashboards.

The original access controls applied to the source data will also persist across platforms, even after export of data from Power BI as a report in an unstructured data format, such as an Microsoft Excel file.

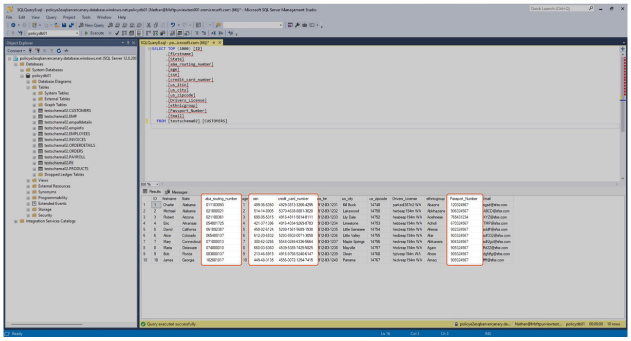

```
SELECT TOP (1000) [ID]
      ,[firstname]
      ,[State]
      ,[aba_routing_number]
      ,[age]
      ,[ssn]
      ,[credit_card_number]
      ,[us_itin]
      ,[us_city]
      ,[us_zipcode]
      ,[Drivers_License]
      ,[ethnicgroup]
      ,[Passport_Number]
      ,[Email]
  FROM [testschema02].[CUSTOMERS]
```

	ID	firstname	State	aba_routing_number	age	ssn	credit_card_number	us_itin	us_city	us_zipcode	Drivers_License	ethnicgroup	Passport_Number	Email
1	1	Charlie	Alabama	011103093	1	489-36-8350	4929-3813-3266-4295	912-83-1231	Kill Buck	14748	parkedl367n2 WA	Abazins	120324567	sgsd@sfss.com
2	2	Michael	Alabama	021000021	2	514-14-8905	5370-4638-8881-3020	912-83-1232	Lakewood	14750	herbesp194m WA	Abkhazians	906324567	ABCD@sfss.com
3	3	Robert	Arizona	021100361	3	690-05-5315	4916-4811-5814-8111	912-83-1233	Lily Dale	14752	heabesp194m WA	Acehnese	765431234	XYZ@sfss.com
4	4	Eric	Arkansas	054001725	4	421-37-1396	4916-4034-9269-8783	912-83-1234	Limestone	14753	hebbesp194m WA	Acholi	670324567	TMP@sfss.com
5	5	David	California	061092387	5	458-02-6124	5299-1561-5689-1938	912-83-1235	Little Genesee	14754	hecbesp194m WA	Afemai	902324567	sddff@sfss.com
6	6	Alice	Colorado	065400137	6	612-20-6832	5293-8502-0071-3058	912-83-1236	Little Valley	14755	hedbesp194m WA	Afar	903324567	sdf332@sfss.com
7	7	Mary	Connecticut	071000013	7	300-62-3266	5548-0246-6336-5664	912-83-1237	Maple Springs	14756	hedvesp194m WA	Afrikaners	904324567	sdf2gd@sfss.com
8	8	Maria	Delaware	074000010	8	660-03-8360	4539-5385-7425-5825	912-83-1238	Mayville	14757	hhdvesp194m WA	Agaw	905324567	ffd22@sfss.com
9	9	Bob	Florida	083000137	9	213-46-8915	4916-9766-5240-6147	912-83-1239	Olean	14760	hjdvesp194m WA	Ahom	909324567	dgfdfg@sfss.com
10	10	James	Georgia	102001017	10	449-48-3135	4556-0072-1294-7415	912-83-1240	Panama	14767	hkdvesp194m WA	Aimaq	909324567	ffff@sfss.com

Query executed successfully. | policye2esqlservercanary.da... | Nathan@Msftpurviewtest... | policydb01 | 00:00:00 | 10 rows

Sensitive data in SQL Database

Image provided by Microsoft, 2024

3 PREVENT DATA LOSS

Detect risky behavior and prevent accidental oversharing of sensitive information

Microsoft Purview supports DLP policies for Power BI, which can alert security teams and end users in real time upon the detection of sensitive data within a dashboard.

For this scenario, Purview can detect when routing and bank account numbers make it into the dashboard, promptly notifying security to contain potential incidents. Additionally, policy tips provide awareness to users that their dashboards contain sensitive information that they should avoid sharing broadly across the organization.

4 GOVERN YOUR DATA

Automatically retain, delete, and store data and records in a compliant manner

Full visibility into both Azure SQL Databases and Power BI dashboards enhances governance of data. Organizations can use Purview to conduct periodic reviews and identify which Power BI dashboards consume sensitive data elements from Azure SQL. Information owners can be notified, and a coordinated review can be conducted to make sure the data should be consumed by Power BI dashboards.

Resulting in success

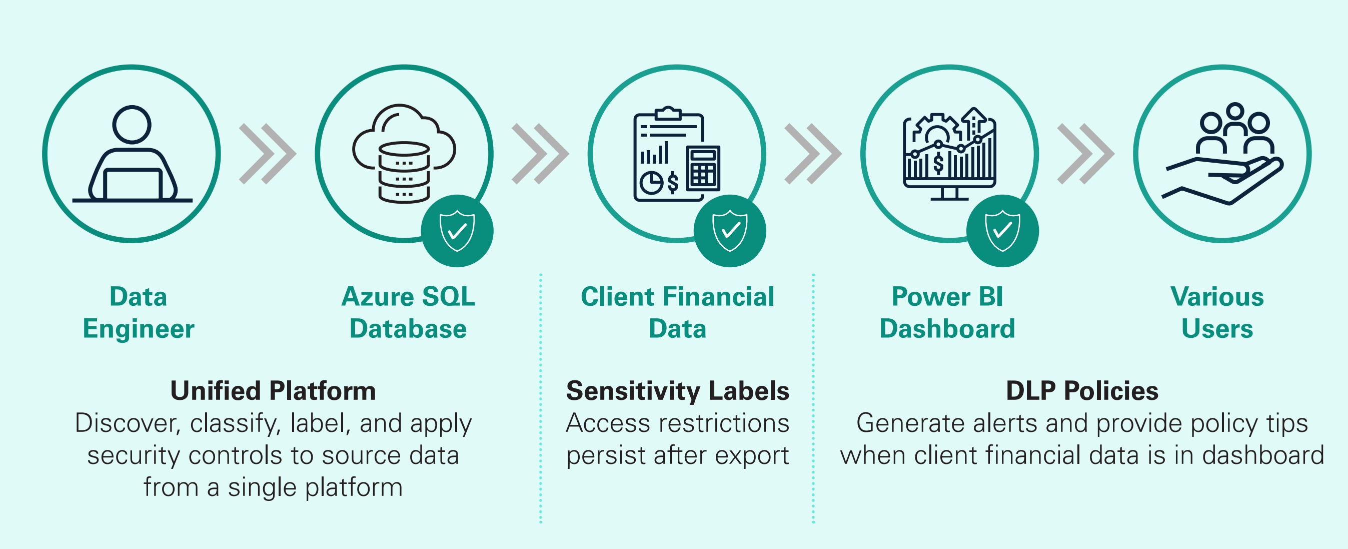

Unified Platform
Discover, classify, label, and apply security controls to source data from a single platform

Sensitivity Labels
Access restrictions persist after export

DLP Policies
Generate alerts and provide policy tips when client financial data is in dashboard

Unified platform

Purview's unified platform allows organizations to discover, classify, and label sensitive client financial data from a single location, including the Azure SQL Database within this example.

Sensitivity labels

Purview's sensitivity labels enforce encryption and access controls that persist from the source data in Azure SQL Database through the Power BI dashboards.

DLP policies

By implementing a DLP policy with Purview, security teams and end users can be notified in real time of sensitive information types found in Power BI dashboards.

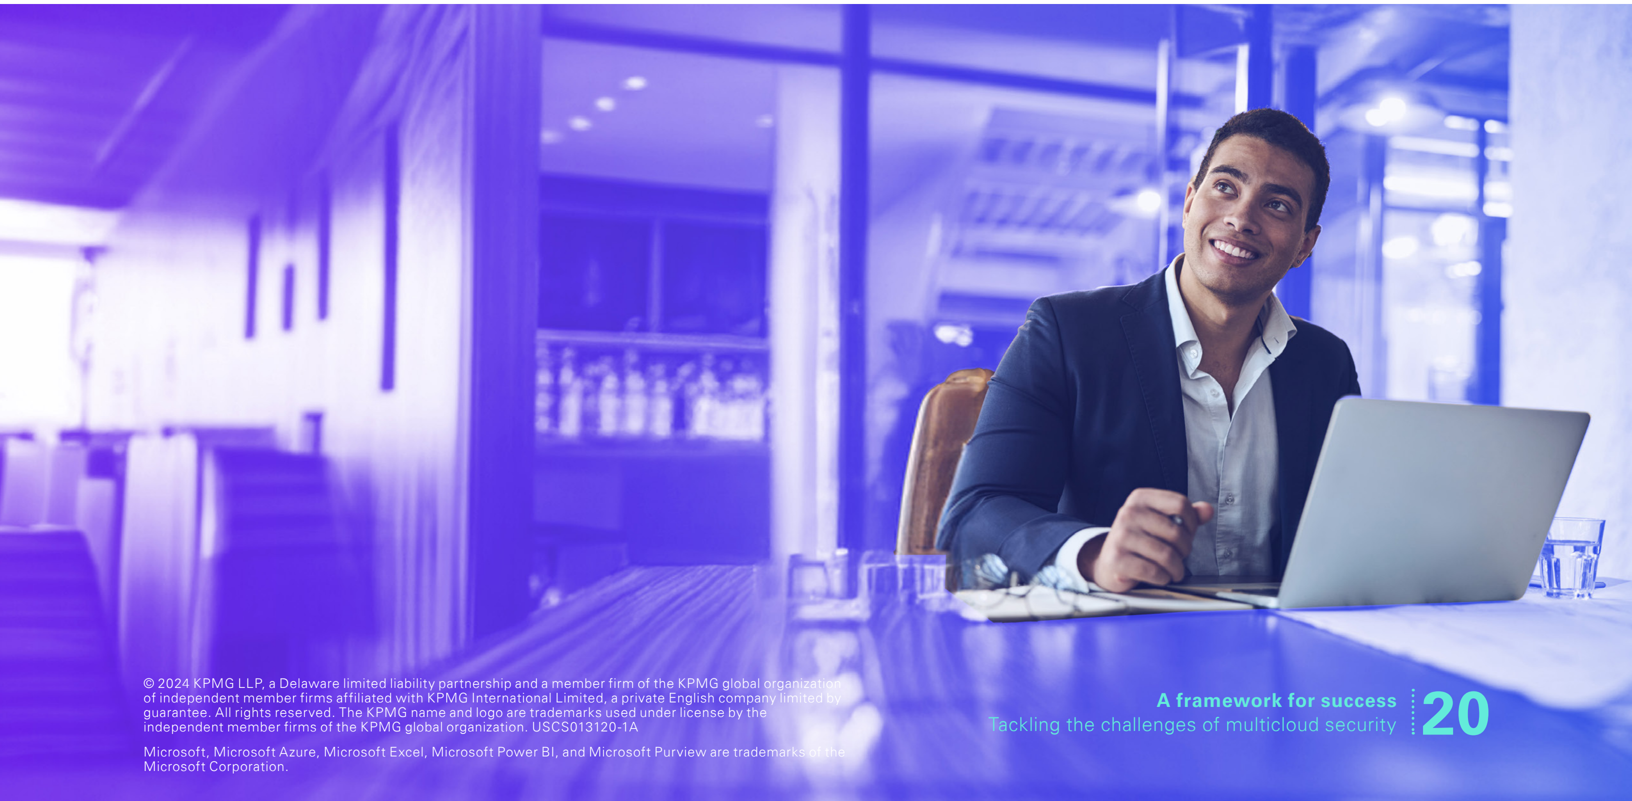

Connect with us

KPMG LLP
345 Park Avenue
New York, NY 10154

Raman Kalyan
E: raman.kalyan@microsoft.com

Daniel Hidalgo
E: daniel.hidalgo@microsoft.com

Ravi Kiran Poluri
E: ravikip@microsoft.com

Annapurna Saripalli
E: annapurna.saripalli@microsoft.com

Pankaj Parikh
E: pankaj.parikh@microsoft.com

Michael D Gomez
E: michaelgomez@kpmg.com

Jim Wilhelm
E: jameswilhelm@kpmg.com

Venoth Lal
E: venothlal@kpmg.com

Ryan McGurgan
E: rmcgurgan@kpmg.com

Some or all of the services described herein may not be permissible for KPMG audit clients and their affiliates or related entities.

Learn about us: in | **kpmg.com**

The information contained herein is of a general nature and is not intended to address the circumstances of any particular individual or entity. Although we endeavor to provide accurate and timely information, there can be no guarantee that such information is accurate as of the date it is received or that it will continue to be accurate in the future. No one should act upon such information without appropriate professional advice after a thorough examination of the particular situation.

Microsoft Intelligent Security Association

Microsoft Security

Microsoft Verified Managed XDR Solution