



Microsoft Purview Azure Pricing Model

Rolf Tesmer
Senior Cloud Architect
23 Sep 2022

Today's data realities

What **data** do I have?

Is it **trustworthy**?

Can people access the **data** needed to make the right decisions?

How can I **enable faster business insights**?

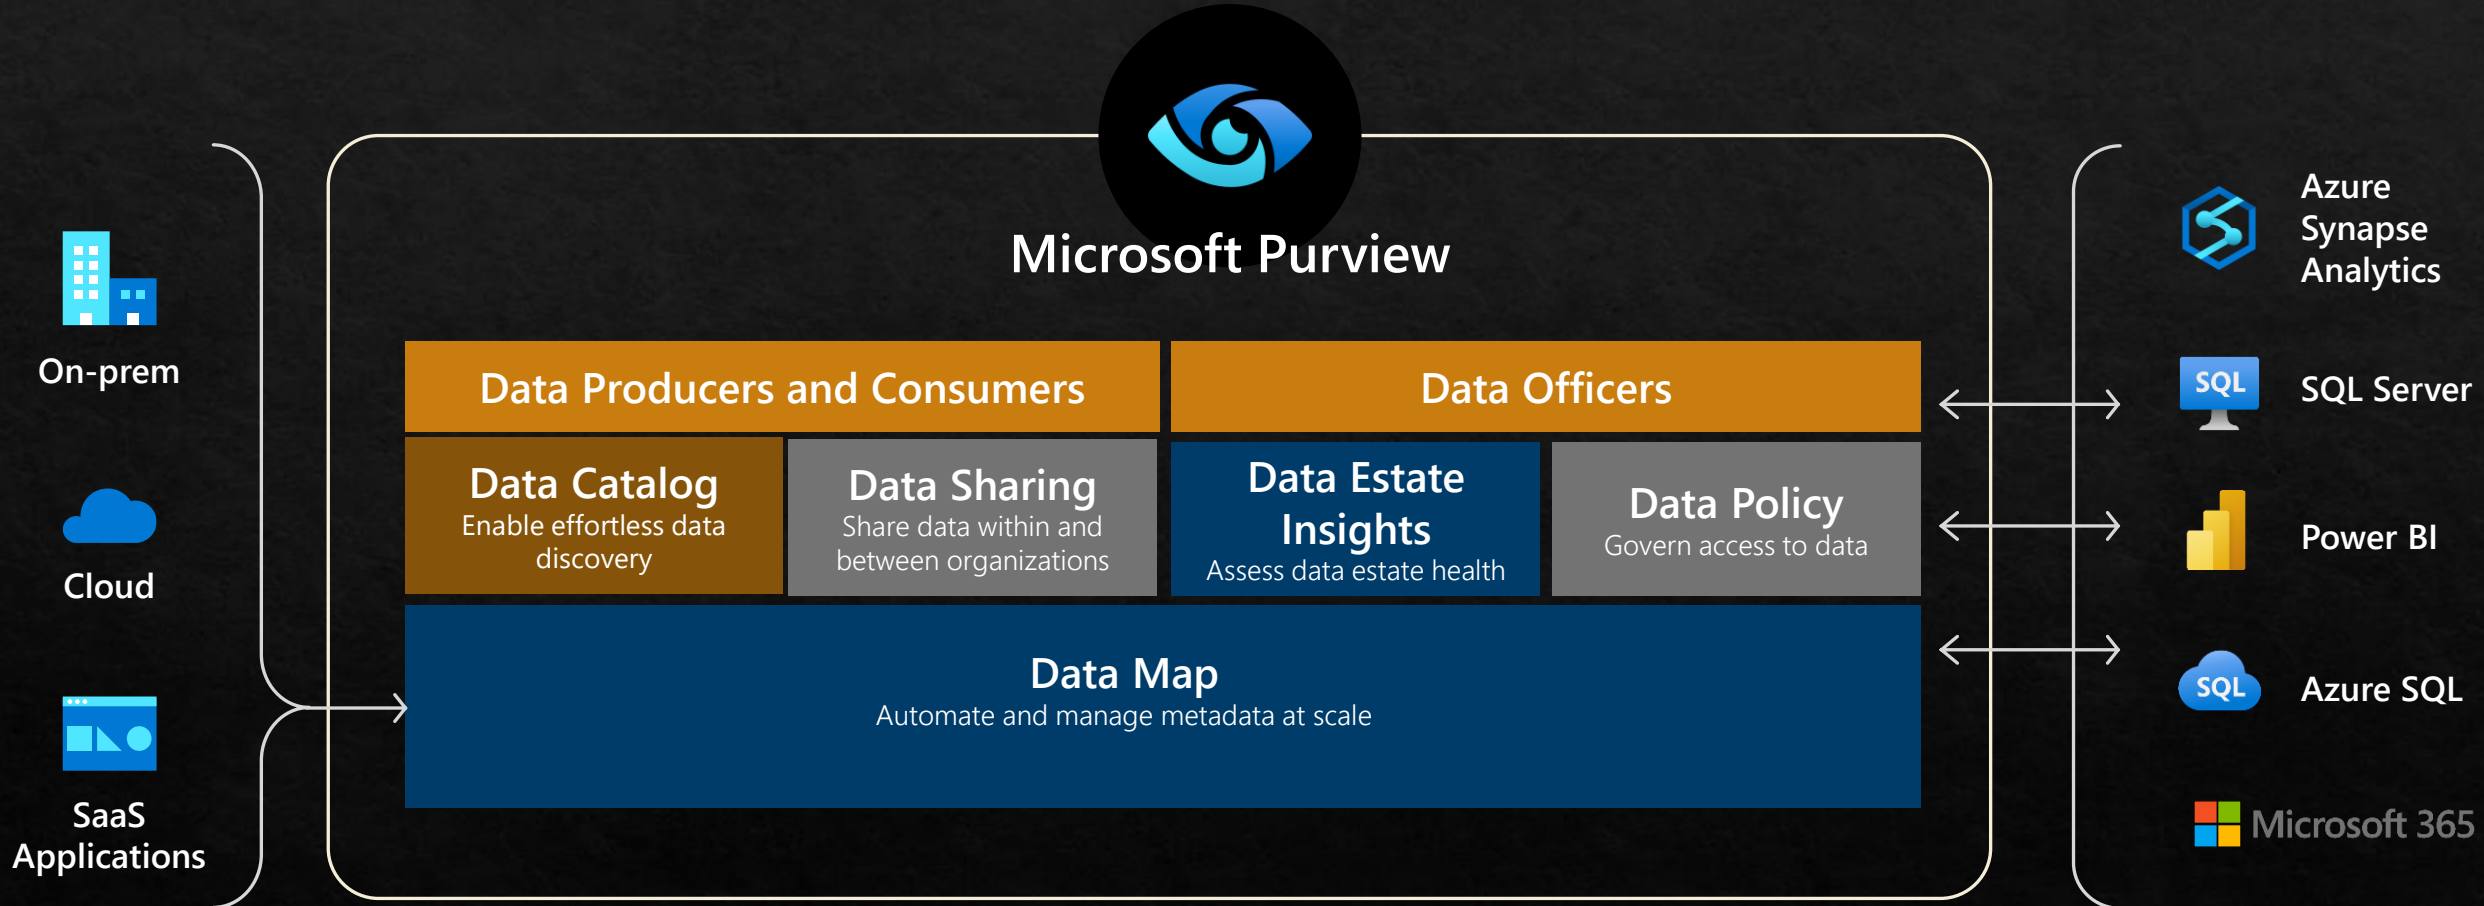
What's my **compliance exposure**?



Unified Data Governance with Microsoft Purview

Generally Available

Preview



Microsoft Purview Pricing Overview

- Microsoft Purview pricing include both “*Direct*” and “*Indirect*” costs
- “**Indirect**” Costs
 - **Storage Account** (*required*) – cater for secured scanning. < \$0-\$100/mth (~2% scan charge)
 - **Event Hub** (*optional*) – enable monitoring with Atlas Kafka topics events. < \$30/mth
 - **Private Endpoint** (*optional*) – enable secure access over a private link / vnet. < \$20/mth/PE
 - **SHIR VM** (*optional*) – enable secure DataMap scanning within a vnet. < \$100-\$800/mth/VM
 - **MS 365** (*optional*) – enable Information Protection (IP) Labels requires [Microsoft 365 license](#)
 - **Data Egress** (*varies*) – egress charges apply if data [exits Azure](#), or 3rd party clouds (ie [AWS](#))
- “**Direct**” Costs – *see next page for overview*
 - [Purview Data Map](#) (*required*) – includes 3 activities; Population, Enrichment, Consumption
 - [Data Estate Insights](#) (*optional*) - understand health and governance posture of the data estate

Microsoft Purview Pricing – Data Map

- Microsoft Purview pricing “*Direct*” costs → [Purview Data Map](#) (*required*)
- Pricing for **Data Map** is based on 3 core dimensions / functions
 - [Elastic data map](#)
 - [Automated scanning & classification](#)
 - [Advanced resource sets](#) (*optional*)

<p><u>Elastic Data Map</u> Charged in Capacity Unit (CU). \$0.594 / CU / hour (rolled to hour)</p> <p>Purview provisions min 1 CU (by default)</p> <p>Each CU provisioned in the Data Map;</p> <ul style="list-style-type: none"> • Provides 10 GB metadata storage • Serves 25 data map API operations/sec • Ops = Create/Read/Write/Update/Delete on any stored metadata (avg/hour) • CU is “<i>auto-elastic</i>” to max 10x CU limit • Can increase CU max limit to 100 (via SR) <p>Cost = % duration a CU was used/month</p>	<p><u>Scanning & Classification</u> Charged in scan compute used/hour. \$0.910 per 1 vCore/Hour (rolled to minute)</p> <p>Metadata is ingested via Native Connector Source Scan or ADF Pipeline Scan. Includes data scan, classification, metadata creation</p> <p>Native Connector Source Scan Cost</p> <ul style="list-style-type: none"> • Scans can be Full or Incremental. • Rules set what/how source is scanned <p>ADF Pipeline Scan Cost</p> <ul style="list-style-type: none"> • Metadata/Lineage ingested each time a pipeline run is completed 	<p><u>Advanced Resource Sets</u> (<i>optional</i>) Charged in scan compute used/hour. \$0.304 per 1 vCore/Hour (rolled to minute)</p> <p>Resource Sets map many data assets into a single logical resource, ie scan all Data Lake files to find patterns, group as single asset</p> <p>Scanning enables enriched computed info such as Total Size, Partition Count, etc – and cost dependent on the size/number of resource sets registered in your catalog.</p> <p>By default scanning runs every 12 hours for all systems with “resource set” enabled</p>
--	---	--

Microsoft Purview Pricing – Data Estate Insights

- Microsoft Purview pricing “*Direct*” costs → [Data Estate Insights](#) (*optional*)
- Track, monitor and understand health and governance posture of the Data Estate
- Pricing for the **Data Estate Insights** is based on 2 core dimensions / functions
 - Report Generation
 - Report Consumption

Report Generation

Charged in compute used/hour.

\$1.185 per 1 vCore/Hour (rolled to minute)

Aggregates metadata and classifications into enriched, executive-ready reports that can be visualized in Data Estate Insights.

Enabled at provisioning (default). Can be disabled (optional) – however existing reports can still be consumed.

Refresh occurs weekly (default, cannot be changed).

Report Consumption

Charged in total number of report API calls

\$0.304 per API Call (1x call returns 10K rows)

Enabled at provisioning (default). Can be disabled (optional) – and no existing reports can be consumed, and generation is stopped.

Data Estate Insights API calls serve aggregated and detailed data to users across asset, glossary, classification, sensitive labels, etc.

1x API call returns up to 10,000 rows of tabular result. The data populates the main dashboard or the detail pages as the users consume Insights via the Portal.

Microsoft Purview **Pricing Example**

<u>Microsoft Purview Deployment and Options Enabled</u>	<u>Est Mth Cost (\$AUD, RRP, PAYG, AU Region, No AHUB)</u>
Account = Single Production deployment provisioned in AU East No Event Hub, default Storage Account Private Link / Vnet = Enabled. 1TB data moved through network (at month end) SHIR = Single deployment on Azure VM (D4s v3) with 1x Premium P30 disk Data Map 29GB of Metadata Stored (at month end) Operations usage pattern of average 50 Ops/Sec (at month end) Data Source Scans = run daily (default connectors only) Scans take 2 hours (1 vcore) for all Data Sources Advanced Resource Set = configured for Data Lake Scans take 1 hours (1 vcore) for all assets (default scanning runs every 12 hours) Data Estate Insights = ENABLED Report Generation takes 30 minutes (default generation runs weekly) Report Consumption uses 375 API calls (at month end)	< \$1.00 \$10.95 \$14.85 \$587.16 \$1,325.80 (Required CU = 3) \$56.42 \$18.85 \$2.37 \$114.00
All prices are RRP Azure pricing in \$AUD (excl. GST). Pricing current as at 9 Aug 2022. Does not reflect EA pricing. This is a <u>Pricing Example ONLY</u>	Total: AUD \$ 2131.40 / Month

(assuming 31 days/month)
(assuming 4 weeks/month)

Microsoft Purview Pricing FAQ (as at 9 Aug 2022)

What data source scanning is free (today)	For a limited time, on-premises SQL Server and Power BI online tenants will have free scanning and classification
What classifies as Metadata Storage	Storage for the business and technical metadata and lineage associated with data assets in a searchable graph format
Is any amount of Metadata Storage free	Yes. The first 1MB of metadata storage is free as part of your Microsoft Purview account.
What is the CU max elastic window size	You can provision a Data Map with CU of 1 (or above). Data Map metadata storage scales linearly in 10 GB increments per provisioned CU. The capacity of the Data Map cannot be changed after provisioning. As a cost control measure, a Data Map is configured by default to elastically scale up within the “elasticity window”. Currently the maximum CU within the “elasticity window” is 10 CU. This can be increased to 100 CU via a Support Call.
What is the max number of total vCores available for scanning, per account	160 maximum. This cannot be increased
What is the max amount of time a scan can run for	7 days (or 168 hours). This cannot be increased
Can I use the same self-hosted integration runtime (SHIR) for both Microsoft Purview and Azure Data Factory?	No, these must be physically separate SHIR installs on separate VM's. For more information about the SHIR, see our self-hosted integration runtime documentation .

References

- ◇ [Microsoft Purview \(formerly Azure Purview\) accounts architecture and best practices - Microsoft Purview | Microsoft Docs](#)
- ◇ [Best practices for scanning data sources in Microsoft Purview - Microsoft Purview | Microsoft Docs](#)
- ◇ [Microsoft Purview collections architecture and best practices - Microsoft Purview | Microsoft Docs](#)
- ◇ [Classification best practices for the Microsoft Purview governance portal - Microsoft Purview | Microsoft Docs](#)
- ◇ [Pricing guidelines for Microsoft Purview \(formerly Azure Purview\) - Microsoft Purview | Microsoft Docs](#)
- ◇ [Pricing - Microsoft Purview | Microsoft Azure](#)