April, 2025

Microsoft Fabric, Unraveled

Closing the Gaps in Capacity Monitoring

DATALOGZ

Gaps in Microsoft Fabric

The 2023 launch of Microsoft Fabric presented a vision of a more unified data and analytics environment, where increasingly advanced AI tools were more easily integrated into workflows. This has major implications for BI, and many enterprises are making the transition to Microsoft Fabric.

'Break Out' Moment

By early 2025, Microsoft Fabric reached its "break out" moment at the data layer. Microsoft CEO Satya Nadella said it became "the fastest-growing analytics product in our history," with corporate adoption being particularly explosive. Migration to Microsoft Fabric was identified as a top priority for 2025 in Datalogz' recent State of BI report, which was based on interviews with 55 data and analytics leaders. Microsoft Fabric also serves as a tailwind to the growth of Power BI, as the decade-old platform continued to experience 40% year-over-year growth in Q2 of 2025.

But as Microsoft Fabric is implemented by enterprises at scale, there is a growing gulf between the promises of tomorrow and realities of today.

More than a year since the launch of Microsoft Fabric, there is continued potential that AI will reshape workflows and deliver powerful performance tomorrow. But, today, the implementation of AI at scale is still being worked out.

Today, the administrators who are operating Microsoft Fabric need to deliver successful data products and a seamless experience with the tools, workflows, and environments that are being used to drive business forward inside the companies that power the U.S. economy.

Complexity at the Consumption Layer

There is a particular need for optimization at the data consumption layer, where the most business users interact with data. The popularity of Microsoft Power BI, now 10 years old, has brought enterprises closer than ever to making self-serve analytics a reality. But administrators need accurate and responsive capacity monitoring to keep costs under control, deliver high performance, and ensure that resources are distributed proportionally between different departments that are using the product, and other parts of the analytics stack.

With Microsoft Fabric, all Microsoft data and analytics tools are delivered and metered under one umbrella. Previously, Power BI was available only as a standalone platform, with Power BI-specific capacities purchased to allocate compute. Now, Microsoft Fabric houses a complete data stack, including business intelligence (Power BI) data lakes (OneLake), data integration (Azure Data Factory), and data engineering (Azure Synapse Analytics). Compute is allocated through the purchase of a unified capacity, called the F-SKU, which covers usage across the entire stack. Power BI-specific SKUs have been retired, meaning data engineering compute and business intelligence compute come from the same pool.

While the tools and metrics are new, the needs of administrators have not changed. They need to monitor environments carefully, and respond quickly to overages, throttling, and other capacity challenges that can drive up costs. In this report, we identify three gaps that admins want to close today. These concern centralized reporting, real-time data, and chargebacks. By identifying areas for improvement, we hope to plant the seed for collaboration that can lead to meaningful solutions.

Today, Capacity Monitoring in Microsoft Fabric is:

Fragmented

Administrators lack a high-level view of overall capacity health, reducing visibility across departments and geographies.

Reactive

Data on capacity anomalies only becomes available after the fact, preventing action in real-time.

Chargebacks

Compute is not apportioned to individual business units and users, making it difficult to distribute internal spend based on consumption.

FRAGMENTED

Administrators lack a single pane to monitor the health of capacities, across geographies and departments.

When operating at scale, BI administrators must closely monitor the health of all capacities.

A typical BI admin could be overseeing more than a dozen capacities, across multiple departments, located in different geographies. Problems, such as delays, rejections, and throttling, are often just as complex. An issue that is slowing down load times for everyone could be traced to a single refresh, or a spike in compute usage may be caused by a series of interdependent tasks.

Administrators must correctly apportion compute across time and workflows, analyze the root cause of issues, and act quickly to solve them. Today, administrators using Microsoft Fabric must dig through records that are tied to specific capacities, geographies, workflows or storage to effectively monitor compute.

The need

Right now, administrators are looking at the equivalent of a series of individual X-Rays. They need a picture of overall health of capacities, delivered in a consumable format. With a source of truth, they can pinpoint the problem, and administer the remedy.

REACTIVE

Data on capacity anomalies only becomes available after the fact, preventing corrective action in real-time

When there's a capacity issue, every second counts. Capacity issues compound. A spike in usage, if left unaddressed, could grow, leading to throttling or downtime that affects cost and performance. If admins can identify anomalies in usage when they occur, they can move quickly to correct the issue.

The time between an event and the time an admin learns about it and takes action is too large. Currently, Microsoft Fabric capacity monitoring data is delivered after the fact. This means they find out about a problem only when a user reports an issue, or when conducting an analysis after the issue occurs. This leaves admins unable to take immediate action when an issue happens, and prevents proactive monitoring that could root out problems before they drive up cost and compute. Instead, they are stuck responding after the fact.

The need

Right now, admins are stuck looking at events that already happened, and piecing together what went wrong. With real-time data and alerting, they would have a fighting chance to take action immediately, and contain issues before they impact the entire environment.

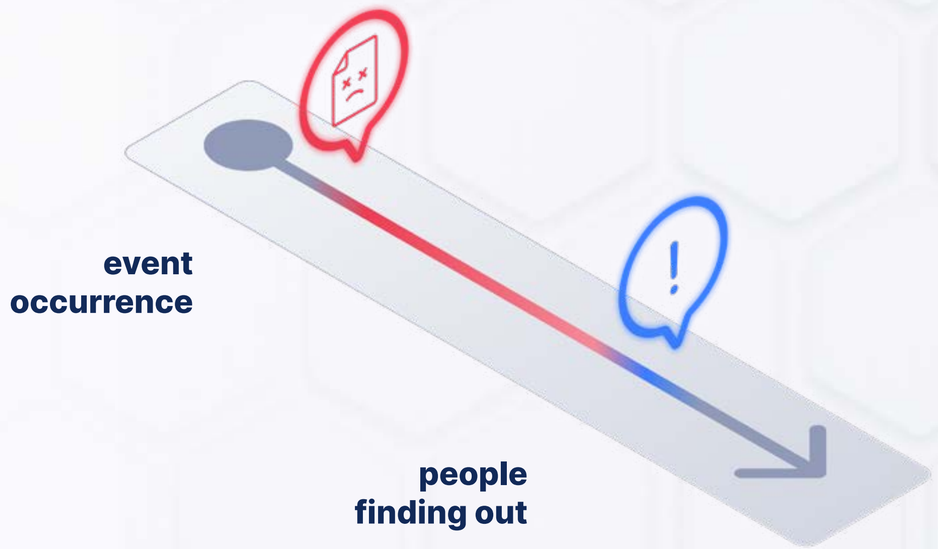

CHARGEBACKS

Capacity monitoring does not tie consumption to spend on the level of the individual business unit or user, making it difficult to allocate cost to departments based on usage.

When managing BI environments, cost and compute are intertwined. BI usage is measured and billed in compute capacity. When overages create issues such as throttling, it not only degrades performance, but could also drive up spend. When managing multiple capacities, admins already face a balancing act to achieve the right balance of compute to optimize costs.

Microsoft Fabric adds even more complexity. All of the tools in the data stack are measured in the same capacity. Admins aren't just measuring BI compute, they are weighing how BI compute factors into overall compute used by data engineering, data lakes, and more.

Take this scenario: A finance department wants to move to Azure, and will use Power BI for benchmarking. This means not only that additional capacity will be needed, but also requires BI admins to apportion compute to specific workloads, so they can in turn calculate the cost associated with those costs, and percentage of cost the finance department should be responsible to cover. Then, the BI admin can compare that percentage to the overall Fabric capacity expenditure for the organization, and issue chargebacks to account for compute consumed by the finance department.

The need

Today, BI admins are left to piece together the above scenario on their own. They need a chargeback mechanism to understand the percentage of compute used by each business unit and user, and easily distribute costs appropriately across departments.

The Datalogz Control Tower

Capacity Monitoring for the Consumption Layer

Datalogz is on a mission to end BI and reporting sprawl. Through its BI ops platform, Control Tower, it provides automated BI administration that enables organizations tocontinually and instantly cut costs, reduce risks, set up BI governance and increase BI efficiency.

Capacity Monitoring with Microsoft Fabric

The Datalogz Control Tower optimizes compute in Microsoft Fabric to enable **high-performing analytics** and **reduced infrastructure spend.**

Control Tower includes the following:

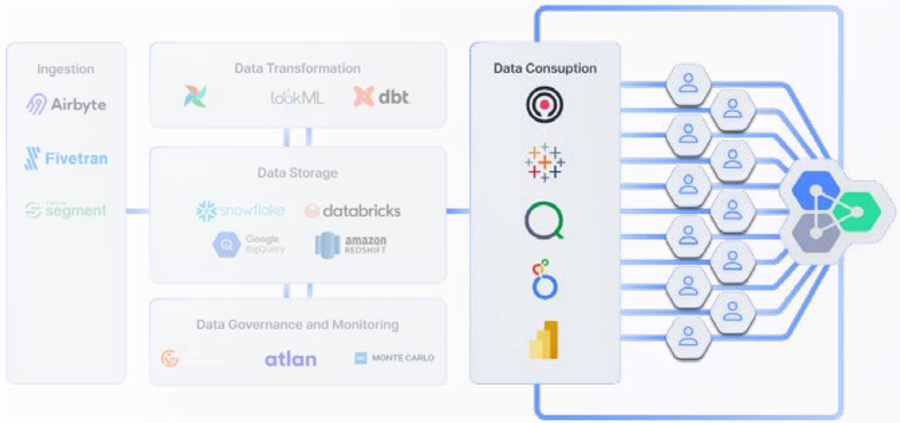

Understand Your Existing BI Environment: The Datalogz Control Tower ingests metadata from across BI tools, including Microsoft Fabric, bringing all capacities together in one place.

Deploy Real-Time Monitoring and Alerts across multiple capacities to root out issues, when they happen, allowing you to take action at the source.

Manage Performance and Cost: Synchronize compute with spend, across capacities, so you can correctly apportion budget, and save money.

About 70% of companies have three or more business intelligence platforms, such as Power BI, Looker, Qlik, or Tableau. The BI sprawl that inevitably comes with this report proliferation is overtaking data-mature organizations. What if an organization has the same KPI being reported in twenty different dashboards with slightly different calculations? Unfortunately, this is the reality for most data-mature organizations. This leads to wasted compute, but more importantly could result in major business consequences, as costs, risk, and labor efforts rise exponentially. BI and reporting make up 5-10% of total IT spend, yet at most orgs it's the wild west in terms of governance, monitoring & controls.

Datalogz provides the guardrails that make enterprise-wide analytics initiatives successful by preventing sprawl and safely enabling self-service analytics.

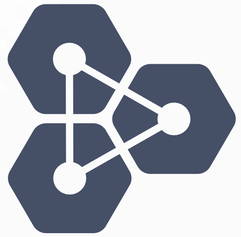

DATALOGZ

On a mission to End BI Sprawl

Book a demo!