# Federal & Sovereign-Cloud Reality Check

Microsoft Fabric and Azure Databricks in Azure Government and GCC High

Version 4.0 · 19 September 2026 · Frank Garofalo · Chief Data Officer, Federal Civilian
Companion to *The Skills Gap Myth*. Written in a personal capacity. Every claim below is dated and sourced, because all of it moves — treat this as a snapshot to check against, not a state to plan from.

**BOTTOM LINE UP FRONT**

1. **The sovereign tiers are further behind than the commercial deck implies, and in different places on each platform.** Azure Government has no Databricks SQL and no Unity Catalog. Fabric in GCC High is in public preview, with GA announced to start **1 October 2026**.
2. **The multi-engine architecture most reference diagrams draw does not exist in GCC High today.** Every mirroring source except Azure SQL Database is unavailable there, so "mirror Databricks into Fabric" is not an option you can select. There is a shortcut-based pattern whose parts are each supported, with one network-isolation constraint to resolve first — §7.
3. **Treat a sovereign requirement as a reason to raise the bar on multi-platform, not as a point in its favor.**

## 1 Getting the tiers right

These are routinely drawn as nested rings in architecture decks. They are not rings, they are different services with different assessments, and the differences decide what you can actually build.

| Tier | What it is | What it is assessed against |
|---|---|---|
| **Microsoft 365 GCC** | Government Community Cloud for US public sector. | FedRAMP High, DFARS, and the requirements for criminal-justice (CJI) and federal tax (FTI) data types. |
| **Microsoft 365 GCC High** | For DoD contractors and organizations handling CUI or ITAR-controlled data. Eligibility is validated before a tenant exists. | NIST SP 800-53 at a FIPS 199 High categorisation. Microsoft states it **can demonstrate equivalency to IL4**, or the inheritance needed for CMMC. **Not IL5.** |
| **Microsoft 365 DoD** | Exclusive to the Department of Defense. Non-DoD entities cannot buy it. | DoD Cloud Computing SRG, **up to Impact Level 5**. |
| **Azure Government** | Microsoft's physically and logically isolated Azure cloud, operated by screened US persons — separate from commercial Azure. This is where your Databricks and Fabric compute actually runs. | FedRAMP High P-ATO; DoD SRG IL2, IL4 and IL5 workload-isolation provisional authorizations, per service. |

### Three terms that get conflated in nearly every requirements document

- **ITAR is not an authorization.** It is a set of contractual and access controls. Microsoft will only agree to ITAR contract language for GCC High — so if a requirements document says "ITAR" and the proposed tenant is GCC, that is a contradiction to resolve before design, not during.
- **HIPAA and CJIS do not require a sovereign cloud.** Both can be met on commercial Azure with the right configuration. A sovereign tier chosen for these reasons is usually an unexamined assumption carrying real cost and real feature loss.
- **HYOK is an information-protection term, not a data-platform capability.** It is a different kind of control from network isolation, and the two are routinely traded against each other as if they were substitutes.

## 2 Azure Databricks in Azure Government

The common instinct is to make Databricks the primary engine in Gov and lean on Unity Catalog and Genie throughout, because that is what the commercial reference architectures do. It does not survive contact with the regional feature documentation.

> **The Azure Government regions `usgovaz` and `usgovva` do not support Databricks SQL or Unity Catalog.** Stated plainly in Microsoft's own Azure Databricks regional-feature documentation, which was last revised 16 September 2026. Every capability that depends on Unity Catalog — Genie, lineage, RLS and column masks, Delta Sharing — depends on something that is not there.

Two things follow, and the second is the one Microsoft-authored documents usually leave out.

- **It is on the roadmap.** Databricks and Microsoft have announced that the full Azure Databricks Data Intelligence Platform comes to Azure Government in 2026. Confirm the current state and the date with both account teams — not with this document, and not with a

slide.

- **The same capability already exists on AWS GovCloud.** Databricks on AWS GovCloud carries FedRAMP High and DoD IL5 authorization *with* Databricks SQL, Unity Catalog and Delta Sharing. If your programme's requirement is genuinely Unity-Catalog-governed analytics at IL5 today, Azure is not currently the platform that does it, and you should hear that from me rather than discover it in month four.

## 3 Microsoft Fabric in GCC High

Fabric for GCC High is in **public preview**, and Microsoft's US-government cloud blog announced on **2 September 2026** that **general availability starts 1 October 2026**. Note where that date does and does not appear: it is in the announcement, not on the Learn page, which still carries the preview banner and the limitations below. Plan against the limitations until the GA feature list is published — GA is a date, not a guarantee that everything in this section cleared.

Available in **US Gov Virginia** and **US Gov Texas** only. Free licences and trials do not exist in government clouds — you need Power BI Pro.

### What is not supported during public preview

Quoted from the limitations list, because paraphrasing this one costs people quarters:

- **Private Link** is not supported.
- **Customer-managed keys (CMK)** are not supported.
- **Outbound access protection** is not supported.
- **OneLake security:** Spark support for OneLake security is not available.
- **Workspace identity** is not supported.
- **OneLake disaster recovery** is not supported.
- **Fabric monitoring** is not supported.
- **Shortcuts** have partial support — Azure Blob Storage and ADLS Gen2 only.
- **Activator** has partial support — no create, edit or delete of alert rules for Power BI reports.

> **The one that reshapes the architecture: every mirroring source except Mirrored Azure SQL Database is unavailable in GCC High.** So is the whole Fabric IQ family (graph model, graph queryset, operations agent, ontology). If your design mirrors a Databricks Unity Catalog or a Snowflake estate into Fabric, that design has no GCC High implementation today — not a degraded one, none.

> **And on the AI layer:** Copilot and data agents do not appear anywhere in the GCC High feature table. The "AI at every tier" architecture does not extend here yet. Plan for its absence and treat its arrival as upside.

## 4 The co-location constraint

Fabric for GCC High runs in US Gov Virginia and US Gov Texas. Azure Databricks in Azure Government runs in `usgovva` and `usgovaz`. **US Gov Virginia is the only region where both can sit together.** That is not a preference, it is the constraint — and it collapses your DR options, because the second region you would fail over to does not have both halves of the architecture.

## 5 Mirroring is three mechanisms, not one word

The failure modes differ by mechanism, so "we mirror it" is not a design statement. Know which one you are running before reasoning about staleness, cost, or what happens when it stops.

| Mechanism | Sources | Data duplicated? | What to watch |
|---|---|---|---|
| **Database mirroring** | Azure SQL DB and MI, Cosmos DB, PostgreSQL, Snowflake, Oracle, SAP, BigQuery, SQL Server | **Yes** — continuously replicated into OneLake as Delta, publishing as fast as every 15 seconds. Not a snapshot. | A paused or deleted capacity stops replication entirely. Storage is free to 1 TB per capacity unit; you pay past that, and while paused. |
| **Metadata mirroring** | Azure Databricks Unity Catalog; Dremio catalog (preview) | **No** — only the catalog structure is mirrored; data is reached in place through OneLake shortcuts. | Latency is source access time, not replication lag. Materialized views, streaming tables and non-Delta external tables do not appear at all. |
| **Open mirroring** | Anything you write an adapter for | **Yes** — you land change data in a OneLake zone; Fabric merges it to Delta. | The correctness of the change feed is yours. No vendor connector to blame for a missed delete. |

> **Which is why "no data duplication" is not a true statement about mirroring in general.** It is true of the middle row and false of the other two — and the middle row is the one a Databricks estate uses. In GCC High, as above, only the top row exists, and only for Azure SQL Database.

## 6 What actually happens to access control at the mirror seam

This is the finding most often called a talent problem. It is not one, and it does not work the way it is usually described. Three different things happen depending on the object and the path.

### 1. Table, schema and catalog GRANTs do not carry over. This one is unambiguous.

Microsoft states it plainly: *"Unity Catalog policies and permission aren't mirrored in Fabric… Permissions set on catalogs, schemas, and tables inside Azure Databricks doesn't carry over to Fabric workspaces."* You re-author them in Fabric's permission model, or they are not there.

> **And every Fabric reader is the same identity.** *"The credential used to create the connection to Unity Catalog of this catalog mirroring is used for all data queries."* Not the calling user — the connection. So the blast radius of the mirror is whatever that one principal can see, for everyone in the Fabric workspace. Scope that credential deliberately; it is the actual security boundary.

### 2. Tables with row filters or column masks fail CLOSED, not open.

A table carrying a row filter or column mask is not readable through the external-access path at all. Databricks lists both routes as unsupported: *"You cannot use Iceberg REST catalog or Unity REST APIs to access tables with row filters or column masks"* and *"Path-based access to files in tables with policies is not supported."* The design intent is stated for the runtime case in the same list — those clients *"fail securely, meaning that if you try to access tables from these runtimes, no data is returned."*

> **So the risk is not a silent leak. It is a silent gap, and then a workaround.** The table arrives unreadable or missing, someone is blocked, and the fastest fix in front of them is to drop the row filter or land an unfiltered copy for Fabric to read. That is how the control actually gets lost — by hand, under delivery pressure, usually without a change record. Watch for it in review; it does not show up as an error.

### 3. The one path where data really does arrive unfiltered.

If the Fabric side reads the ADLS storage *directly* rather than through Unity Catalog — the trusted-workspace-access pattern used to reach firewall-enabled ADLS Gen2 — Microsoft states the consequence explicitly: *"Unity Catalog policies such as RLS/CLM or ABAC are not enforced at the storage layer and will not be applied if a connection is used to directly access storage."* That configuration leans on Fabric workspace identity governance instead.

> **This is the one to put in front of your security team,** because it is a configuration choice that looks like a networking decision. Reaching firewalled storage and bypassing fine-grained policy are the same setting.

### Moving, and worth a re-check before you design around it

> **Cross-engine ABAC is in Beta** (Databricks, as of 11 September 2026). It enforces ABAC policies, row filters and column masks for external engines by routing the read through a Databricks serverless layer that returns sanitized data. It requires external data access enabled on the metastore, the `EXTERNAL USE SCHEMA` privilege, a **managed** table with catalog commits, and OAuth M2M or PAT authentication — reads only, and the serverless compute is billable. It is Beta, so do not put it on a critical path yet, but "fine-grained policy cannot cross the boundary" is a statement with an expiry date on it now. Re-check before you architect around the limitation.

### And the Fabric side has its own three properties to design around

- **OneLake security is a grant model, and workspace roles sit above it.** You cannot deny access granted elsewhere, and **Admins, Members and Contributors are not restricted by RLS or CLS at all**. If your sensitive-data control depends on RLS, every Contributor on that workspace is outside it. That is a workspace-membership review, not a rules review.
- **Enforcement differs by engine.** On a `SELECT *` where CLS hides a column: Spark notebooks succeed and return the allowed columns; the SQL analytics endpoint and semantic models return an error instead. Same policy, two behaviors — test both.
- **In GCC High, Spark support for OneLake security is not available** (§3). So the one engine that degrades gracefully is the one you do not have there.

> **One more, further down the same path.** Through **Direct Lake over SQL**, or T-SQL in **delegated identity mode**, the calling user's identity is not passed to the shortcut target — the calling item's *owner* is. Any per-user rule you expected downstream evaluates against the wrong principal. The documented fix is Direct Lake over OneLake, or T-SQL in user identity mode.

> **Observed in a live tenant, 18 September 2026.** The OneLake security properties above are observed rather than cited. There is no preview marker anywhere in that surface, in a UI that labels other features "(preview)" a few pixels away; the new-role dialog offers exactly one role type (*"apply Read and ReadWrite permissions to your data using Grant roles"*) with no Deny option; and `DefaultReader` is present by default as Grant/Read. Note that a tooltip in the same pane describes Grant and Deny roles — the dialog is authoritative, and it offers only Grant.

> **Seam behavior here is documented, not observed.** Confirming it needs a mirrored Unity Catalog carrying a row filter, which is an environment somebody has to build rather than one you find lying around. Everything in this section about what crosses the seam is read from the vendor documentation cited in §10, and that interaction is exactly where documented behavior and observed behavior diverge. Validate it in your own tenant — and if what you see differs, I would genuinely like to know.

## 7 The interim pattern, and the constraint to resolve

Given §3, mirroring is not the route into Fabric in GCC High. Shortcuts are.

**Databricks writes Delta to ADLS Gen2 → Fabric OneLake shortcut → Direct Lake to Power BI.** Shortcuts to Blob and ADLS Gen2 are supported in the preview, and both Direct Lake on SQL and Direct Lake on OneLake are in the GCC High feature table.

> **The catch is network isolation, not authentication.** The shortcut itself authenticates without difficulty: ADLS Gen2 shortcuts support organizational account, service principal, SAS and account key, and none of those needs a workspace identity. What you cannot do in the preview is firewall that storage account and still reach it. **Trusted workspace access** is the documented route to a firewall-enabled ADLS Gen2 account, and it requires a workspace identity — which GCC High preview does not support. Private Link is unsupported there too, and OneLake shortcuts do not support storage accounts using managed private endpoints in any cloud. So the question for your security team is whether a storage account reachable from the public network is acceptable in an environment you chose for isolation. Settle that before you design around the pattern.

## 8 What is genuinely differentiated on Azure

- **Microsoft Entra ID as one identity plane** across Databricks, Fabric, Purview and Power BI. A real structural advantage over hand-bridging AWS IAM with a separate BI auth model, and it compounds in a federal environment where PIV and CAC federation is already solved.
- **Azure Databricks' first-party integration** into Azure networking, Monitor, Policy and cost management — genuine depth that Databricks-on-AWS or GCP does not have.

Neither of those claims implies sovereign-cloud feature completeness, and I would not let them be used that way. On feature completeness in Gov today, §2 and §3 are the honest picture.

## 9 Sequencing for a Gov or GCC High programme

1. **Verify current Databricks SQL and Unity Catalog availability in your target Azure Government region** directly with the Microsoft and Databricks account teams. Do not plan from this document.
2. **Decide the region first, not last.** If you need both platforms co-located, US Gov Virginia is the only answer, and your DR design has to account for that.
3. **If you are pursuing Fabric in GCC High before GA, scope to the shortcut pattern** (§7), not to mirroring — and get storage authentication signed off early.
4. **Do not plan for Copilot or data-agent capability in GCC High** until it is confirmed present in the feature table.
5. **Scope the mirroring connection credential deliberately.** Every Fabric reader of the mirrored catalog queries as that one principal, so it is the real security boundary — not the Unity Catalog grants behind it.
6. **Inventory which tables carry row filters or column masks before you mirror,** and decide what happens to each one. They will not come through. Agree the answer in design, so nobody solves it later by dropping the filter.
7. **Re-author and independently validate access control on the Fabric side** for every consumer of mirrored data. Review workspace membership at the same time — Contributors are outside OneLake RLS and CLS entirely.
8. **Re-check every date in this document before you commit budget.** The sovereign clouds move on their own schedule and the feature tables lag the commercial ones by quarters.

## 10 Sources

All read 17 September 2026. Document revision dates given where Microsoft publishes them, because that is what tells you whether this page has moved since I read it.


- Microsoft Cloud Blog (US Government) — *Microsoft Fabric in GCC High: building the data foundation for AI*, 2 Sep 2026 — the public-preview and 1 Oct 2026 GA announcement: microsoft.com/en-us/microsoft-cloud/blog/us-government/2026/09/02/microsoft-fabric-in-gcc-high-building-the-data-foundation-for-ai
- Azure Databricks — *Features with limited regional availability* (rev. 16 Sep 2026): learn.microsoft.com/en-us/azure/databricks/resources/feature-region-support
- Microsoft Fabric — *Fabric for US Government GCC High customers (preview)* (rev. 3 Sep 2026): learn.microsoft.com/en-us/fabric/enterprise/us-government-community-cloud-high
- Microsoft Fabric — *Mirrored databases from Azure Databricks: security* (rev. 1 May 2026) — the "permissions don't carry over" and connection-credential statements: learn.microsoft.com/en-us/fabric/mirroring/azure-databricks-security

- Azure Databricks — *Row filters and column masks* (rev. 11 Sep 2026) — the limitations list, including Unity REST API and path-based access: learn.microsoft.com/en-us/azure/databricks/data-governance/unity-catalog/filters-and-masks
- Azure Databricks — *Cross-engine attribute-based access controls (ABAC)*, Beta (rev. 11 Sep 2026): learn.microsoft.com/en-us/azure/databricks/external-access/cross-engine-abac
- Microsoft Fabric — *Mirroring in Fabric* (rev. 28 Aug 2026): learn.microsoft.com/en-us/fabric/mirroring/overview
- Microsoft Fabric — *Mirrored catalog from Azure Databricks*: learn.microsoft.com/en-us/fabric/mirroring/azure-databricks
- Microsoft Fabric — *OneLake table, column, and row-level security* (rev. 19 Aug 2026): learn.microsoft.com/en-us/fabric/onelake/security/table-column-row-security
- Microsoft Fabric — *OneLake shortcuts* (rev. 11 Aug 2026): learn.microsoft.com/en-us/fabric/onelake/onelake-shortcuts
- Microsoft Fabric — *Create an ADLS Gen2 shortcut* (rev. 1 Jul 2026) — the five supported authorization types, and the managed-private-endpoint limitation: learn.microsoft.com/en-us/fabric/onelake/create-adls-shortcut
- Microsoft Fabric — *Trusted workspace access* (rev. 14 Aug 2026) — the workspace-identity prerequisite for reaching a firewall-enabled ADLS Gen2 account: learn.microsoft.com/en-us/fabric/security/security-trusted-workspace-access
- Service description — *Office 365 GCC High and DoD* (rev. 3 Aug 2026): learn.microsoft.com/.../office-365-us-government/gcc-high-and-dod
- Service description — *Office 365 US Government*: learn.microsoft.com/.../office-365-us-government/office-365-us-government
- Azure Government — *Azure services FedRAMP and DoD SRG audit scope*: learn.microsoft.com/en-us/azure/azure-government/compliance/azure-services-in-fedramp-auditscope
- Databricks — *Databricks partners with Microsoft to bring our Data Intelligence Platform to Azure Government*: databricks.com/blog/databricks-partners-microsoft-bring-our-data-intelligence-platform-azure-government
- Databricks — *GA of AWS GovCloud with FedRAMP High and DoD IL5 authorization*: databricks.com/blog/announcing-general-availability-aws-govcloud-fedramp-high-agency-ato-and-department-defense


---

Version 4.0 · 19 September 2026 · Frank Garofalo · Chief Data Officer, Federal Civilian. Personal capacity; not an official Microsoft publication and not a commitment on behalf of Microsoft or Databricks. The OneLake security properties in §6 were observed in a commercial tenant on 18 September 2026 and are marked as such; nothing here has been tested in a live GCC High tenant, and the rest is derived from the vendor documentation cited above, read 17 September 2026. Verify against current documentation and your account teams before architecting a live workload.