

#FABCONSQLCON2026

FABCON

Microsoft Fabric
COMMUNITY CONFERENCE

SQLCON

Microsoft SQL
COMMUNITY CONFERENCE

ATLANTA MARCH 16 - 20, 2026



Engineering Fabric at scale: Accelerating time to value





The experts with you today




Shishir Tejpal

Esri | Principal Data Architect
stejpal@esri.com



Brandy Weatherly

PwC US | Senior Manager,
Data Engineering & Analytics,
brandy.l.weatherly@pwc.com



Eugene Grib

PwC US | Director, Data
Engineering & Analytics
eugene.grib@pwc.com



Esri—powering location intelligence at a global scale

What Esri does

Core business

- Global leader in Geographic Information Systems (GIS) software
- Powers location intelligence, mapping, and spatial analytics
- Enables governments, utilities, enterprises, and public sector organizations to make data-driven decisions

Market position and scale

Market leadership

- Recognized leader in GIS and geospatial analytics
- Strategic technology partner to Microsoft
- ArcGIS for MS Fabric was part of MS Ignite 2025 / FabCon 2025

Platform reach

- ArcGIS platform deployed worldwide
- Used across the defense, utilities, transportation, telecommunications, environmental, and public safety sectors

Enterprise scale

- 6,000+ employees
- 90 offices worldwide
- Large enterprise data estate supporting global operations

The real problem—scaling innovation on a legacy foundation

01

Business drivers

Growing demand for faster analytics delivery

Executive focus on near-real-time insights

Growth across product, sales, and operational data domains

Need for governed, self-service analytics

02

Transition triggers

Microsoft Fabric release aligned with Esri's unified data strategy (one data platform used across all analytics and operational decisions)

Strategic opportunity to modernize intentionally—not 'lift-and-shift'

03

Operational friction

Cross-domain dependencies slowing change cycles

High effort required for validation and impact analysis

Difficulty scaling new subject areas efficiently

04

Technical reality

Large number of tightly-coupled warehouse objects

Heavy SQL dependency and accumulated technical debt

Limited automation in lineage and orchestration

Translating business ambition into technical reality

What the business requires

Business objectives



Technical implications

Near-real-time insight



reduced latency across ingestion and transformation

Self-service analytics



stable, governed semantic models

Domain expansion



repeatable onboarding patterns

Faster analytics delivery



automated testing and promotion workflows


Enterprise complexity

Enterprise data estate characteristics

- Multiple domains and sub-domains (e.g. sales, products, HR, finance, etc.)
- Thousands of tables across operational data stores
- High attribute counts with cross-domain dependencies
- Changing and evolving source systems (e.g. SAP, Salesforce, etc.)
- Legacy modeling patterns embedded in downstream reports
- Over 4,000 tables, 3,000 views, 1,200 stored procedures, 100s of SSIS packages
- Over 1,200 curated objects
- Multiple ODS migrations in parallel

Architecting for scale—Microsoft Fabric and OneLake


Overview of Esri's high-level architecture




Illustrative and non-exhaustive

Metadata-driven ingestion—eliminating manual pipeline build


Framework



Deploy



Audit



- Source-to-target mapping is defined in metadata (not hard-coded pipelines).
- Parameterized PySpark ingestion templates are reused across domains.
- Dynamic schema handling and column-mapping are driven by configuration.
- Automated data quality and validation rules are applied consistently.
- A standardized ingestion framework enables the rapid onboarding of new tables.

Rebuilding the foundation—a governed Common Data Model (CDM)


CDM discipline determines whether Fabric would simplify - or amplify - complexity

What we changed

- Anchored CDM to two production-critical reports
- Untangled tightly-coupled dimensions
- Re-established star schema discipline
- Standardized business keys and date modeling

What it changed

- Greatly reduced snowflake sprawl
- Enabled repeatable domain onboarding
- Stabilized semantic models
- Positioned platform for AI readiness




Migration without CDM discipline is just technical debt in a new platform.

Single-source ingestion—zero duplication architecture

All reporting workspaces leverage governed CDM data using repeatable ingestion- and shortcut-based access

- Data is **sourced once** into the CDM from EFS.
- All Fabric workspaces consume governed CDM data (**not independent copies**).
- Where data is not yet modeled in CDM, it is accessed from EFS using the **same metadata-driven ingestion framework**.
- Data moves across workspaces using **Fabric shortcuts**, with no duplication or shadow datasets.
- Consistent ingestion patterns ensure **auditability, lineage, and repeatability** across domains.



Framework vs. delivery—a multi-workspace strategy

Why this matters


- Enables scale without chaos
- Isolates domain changes
- Supports parallel development

Framework workspace

- Metadata-driven ingestion templates
- Reusable orchestration logic
- Core CDM structures
- Central governance and RBAC controls

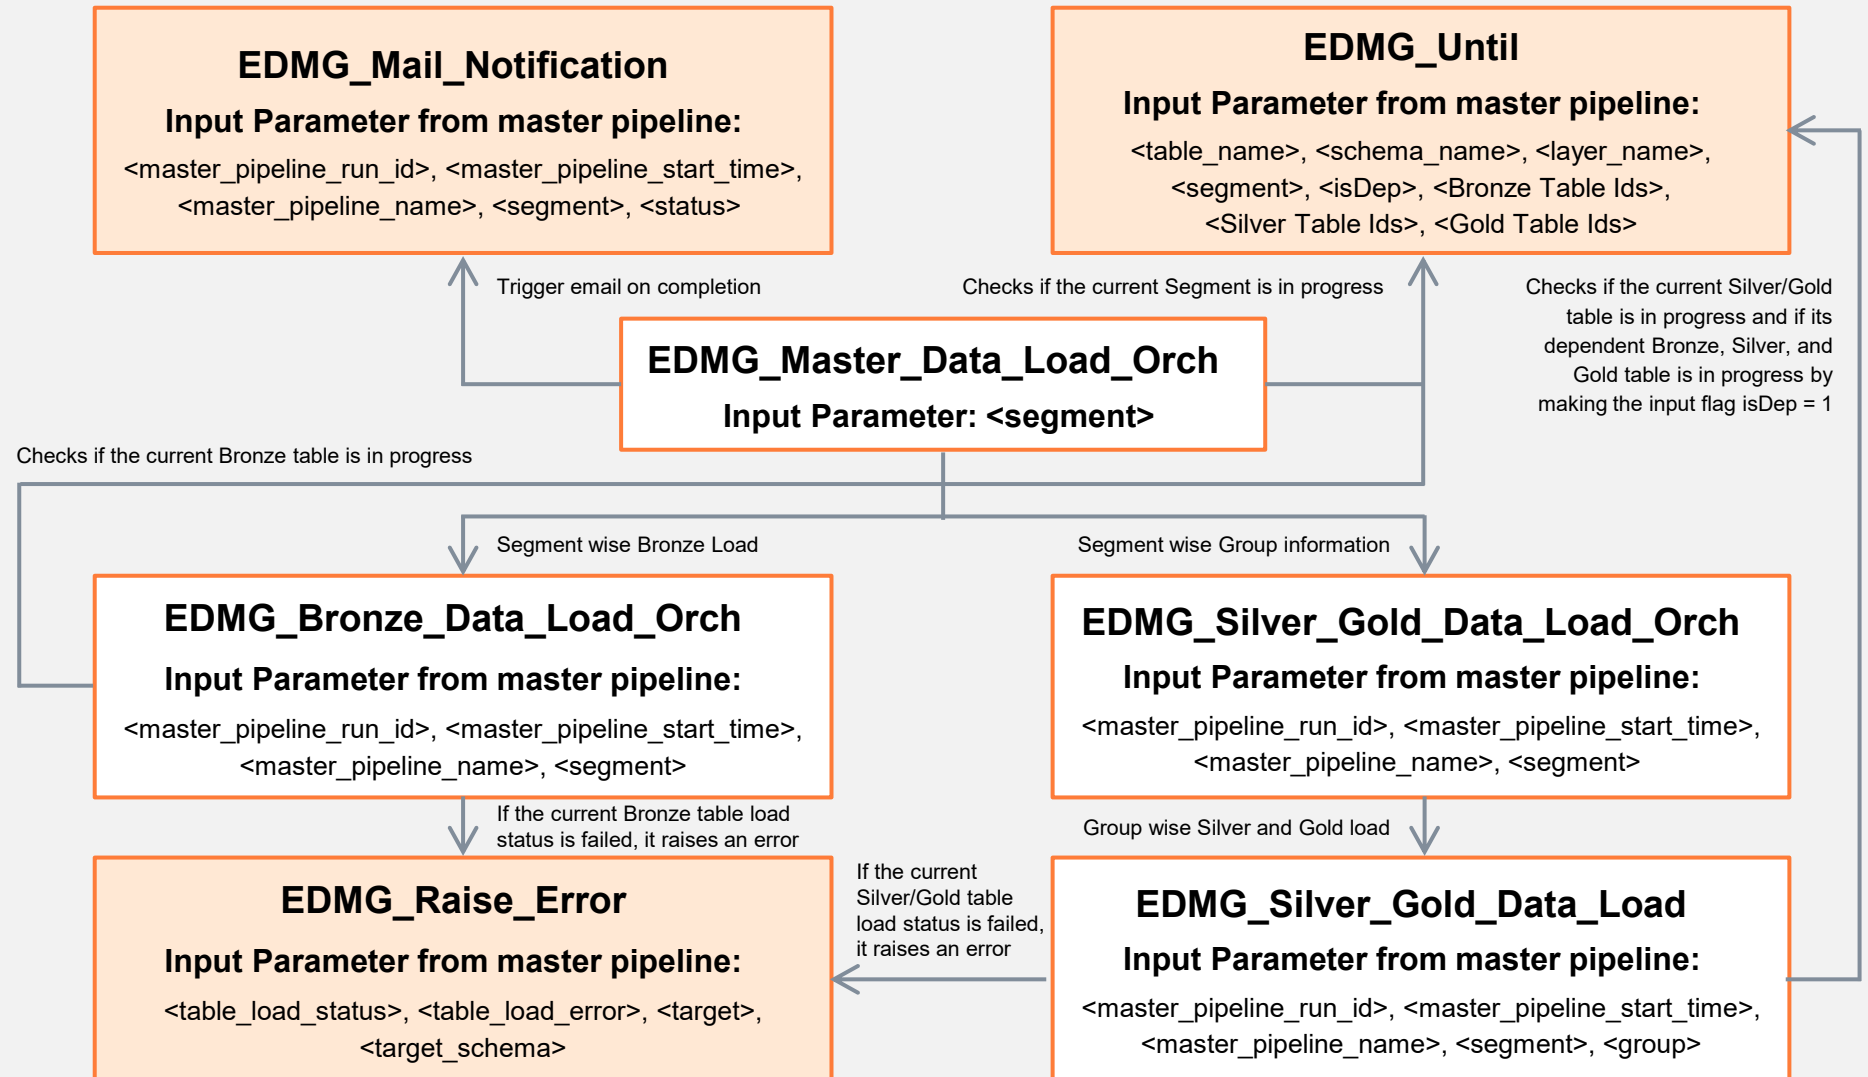
Delivery workspaces

- Subject-area aligned (e.g. sales, finance, HR, etc.)
- Independent lifecycle management
- Controlled semantic model deployment
- Reduced cross-domain risk
- No data duplication for shared datasets (shortcuts)



Orchestration at scale—managing thousands of objects

- Domain-based orchestration strategy
- Dependency-aware pipeline* execution
- Automated promotion across development and QA
- Integrated testing and reconciliation
- Lineage visibility through the Framework Metadata
- Segmented loads (ERP vs. non-ERP)
- EFS integration strategy through shortcuts



*Pipelines used for orchestration

Advice for enterprise Fabric modernization

Key considerations for organizations beginning a Fabric modernization journey

Governance before migration

Architecture with intent

Platform alignment

Do this

- Align business, data, and platform stakeholders early
- Define and govern your common data model before migration
- Clarify business keys, dimensional boundaries, and ownership

- Invest in metadata-driven ingestion and transformation
- Define workspace strategy (framework vs. domain delivery) up front
- Design for repeatable domain onboarding

- Align with Microsoft on Fabric capabilities and roadmap early
- Design with semantic models and downstream consumption in mind
- Plan for AI and advanced analytics, even if not immediate

Avoid this

- Migrating legacy dimensional chaos into fabric
- Treating workspace cleanup as a 'later' activity
- Assuming governance will emerge organically

- Building manual pipelines that don't scale
- Allowing cross-domain workspace sprawl
- Designing only for the initial migration instead of long-term evolution

- Treating fabric as just another storage layer
- Over-optimizing for current use cases only
- Ignoring evolving platform capabilities

Thank you

Sound off.
The mic is all yours.
Influence the product roadmap.

Join the Fabric user panel



Share your feedback directly with
our Fabric product group and
researchers.

<https://aka.ms/JoinFabricUserPanel>

Join the SQL user panel




Influence our SQL roadmap and
ensure it meets your real-life
needs.

<https://aka.ms/JoinSQLUserPanel>

How was the session?



Complete Session Surveys in
Whova for your chance to WIN
PRIZES!



Get Two Fabric Certifications for FREE

Attendees of FABCON can take the Fabric Analytics Engineer or Fabric Data Engineer exam for free. Be part of the 2 fastest growing role-based certifications in Microsoft history.

Request your voucher by March 23, 2026.

<https://aka.ms/fabcon/cert100>

