

#FABCONSQLCON2026

FABCON

Microsoft Fabric
COMMUNITY CONFERENCE

SQLCON

Microsoft SQL
COMMUNITY CONFERENCE

ATLANTA MARCH 16 - 20, 2026

Lessons from the Trenches – Fabric Features That Would Have Changed My Data Warehouse Design

From early release POV stages to a robust, centralized platform that enables developers and end-users alike, Microsoft Fabric's accelerated growth enables value at an enterprise scale.

Meet Your Presenters



Ajay Punyapu
RSM
Solutions Architect,
Data Analytics + AI



Nathan Campbell
RSM
Data Architect, Data
Analytics + AI



Find us on LinkedIn!



FABCON | **SQLCON**
Microsoft Fabric | Microsoft SQL
COMMUNITY CONFERENCE | COMMUNITY CONFERENCE

#FABCONSQLCON2026

Session Context and Disclosures



Lessons from the Trenches: Session Overview



Journey of Data Warehouse

Exploring the challenges and growth in building a Microsoft Fabric analytics solution before advanced features existed.



Lessons Learned Insights

Real implementation stories reveal what worked and what didn't, guiding better design decisions.



Business and Technical Impact

Understanding how platform evolution affects both technical engineers and executives' outcomes.



Driving Efficiency and Value

Emerging Fabric features help reduce complexity and support scalable analytics solutions effectively.

Guardrails: Preview Features & Before We Start



Planning to Scale

Plan for the future of your data estate. Features and tools within Fabric are only as strong as the solution architecture behind it. Having a strong workspace management plan will ease adoptability and scale.

NC2



Importance of Staying in the Know

Knowing the cadence of how preview features are released, and that these features can potentially require additional development and testing.



Understanding GA vs Public Preview vs Private Preview

Who is it available for, and how others can access what is in preview can vary. It is important to understand that features in preview are not yet perfected.

NC1



Varying Experience

Know that these features and enhancements are from our practical experiences in this field, and that what has worked in our use-cases may not be the answer for all.

FABCON | **SQLCON**
Microsoft Fabric | Microsoft SQL
COMMUNITY CONFERENCE | COMMUNITY CONFERENCE

#FABCONSQLCON2026

Fabric Platform Transformation



From Then to Now: Fabric Timeline & Major Milestones

Early Platform Limitations



Initial versions required complex workarounds, causing engineering challenges and increasing operational overhead.

Recent updates introduced features enabling more robust, scalable, and maintainable data solutions.



Platform Maturity Gains

Improved Deployment and Security



New features streamline deployment, enhance data integration, and improve platform security for organizations.

Architects today can leverage new capabilities that were previously constrained by platform gaps.



Evolving Architecture Choices

FABCON | **SQLCON**
Microsoft Fabric | Microsoft SQL
COMMUNITY CONFERENCE | COMMUNITY CONFERENCE

#FABCONSQLCON2026

Deployment and Environment Management


Continuous Integration & Continuous Deployment: **Then**



Inheriting Deployment Pipelines

When Microsoft had first released Fabric, CI/CD was available in limited fashion and existing only from what was supported for Power BI.

NC1



Early Limitations

Deployments couldn't be parameterized, git could not be integrated and automated repetition was near-impossible. External workflows were required to handle deployments, introducing risk.



What was done


Risk mitigation was difficult due to the manual nature of deployments. YAML pipelines and sql project artifacts were created to manage continuous integration & deployment.

Continuous Integration & Continuous Deployment: **Now**



Early Deployment Risk for Fabric Enabled Items

Risk became apparent due to unsupported Fabric resources throughout the early days of the deployment pipelines. Most Fabric resources are available for deployment.



Deployment Pipeline Evolution

Fabric Deployment Pipelines introduced controlled promotion and improved release visibility, reducing inefficiencies and risk.



Improved Governance and Configuration

Adoption of resource tracking and environment-specific configurations enhanced governance and deployment reliability.



Key takeaways and Best Practices

Recent releases in Microsoft Fabric CI/CD features have focused on **risk reduction**. This is to ensure minimal business downtime, repeatable and auditable processes.

Key Feature: Environment Variable Libraries



Challenges of Previous Methods

Earlier environment management relied on altering JSON using pipelines and notebooks causing configuration drift and promotion issues.



Benefits of Variable Libraries

Variable libraries provide a first-class platform feature that reduces errors and operational risks.



Enhanced Promotion Confidence

Standardized environment management supports controlled rollouts and simplifies promotion processes.

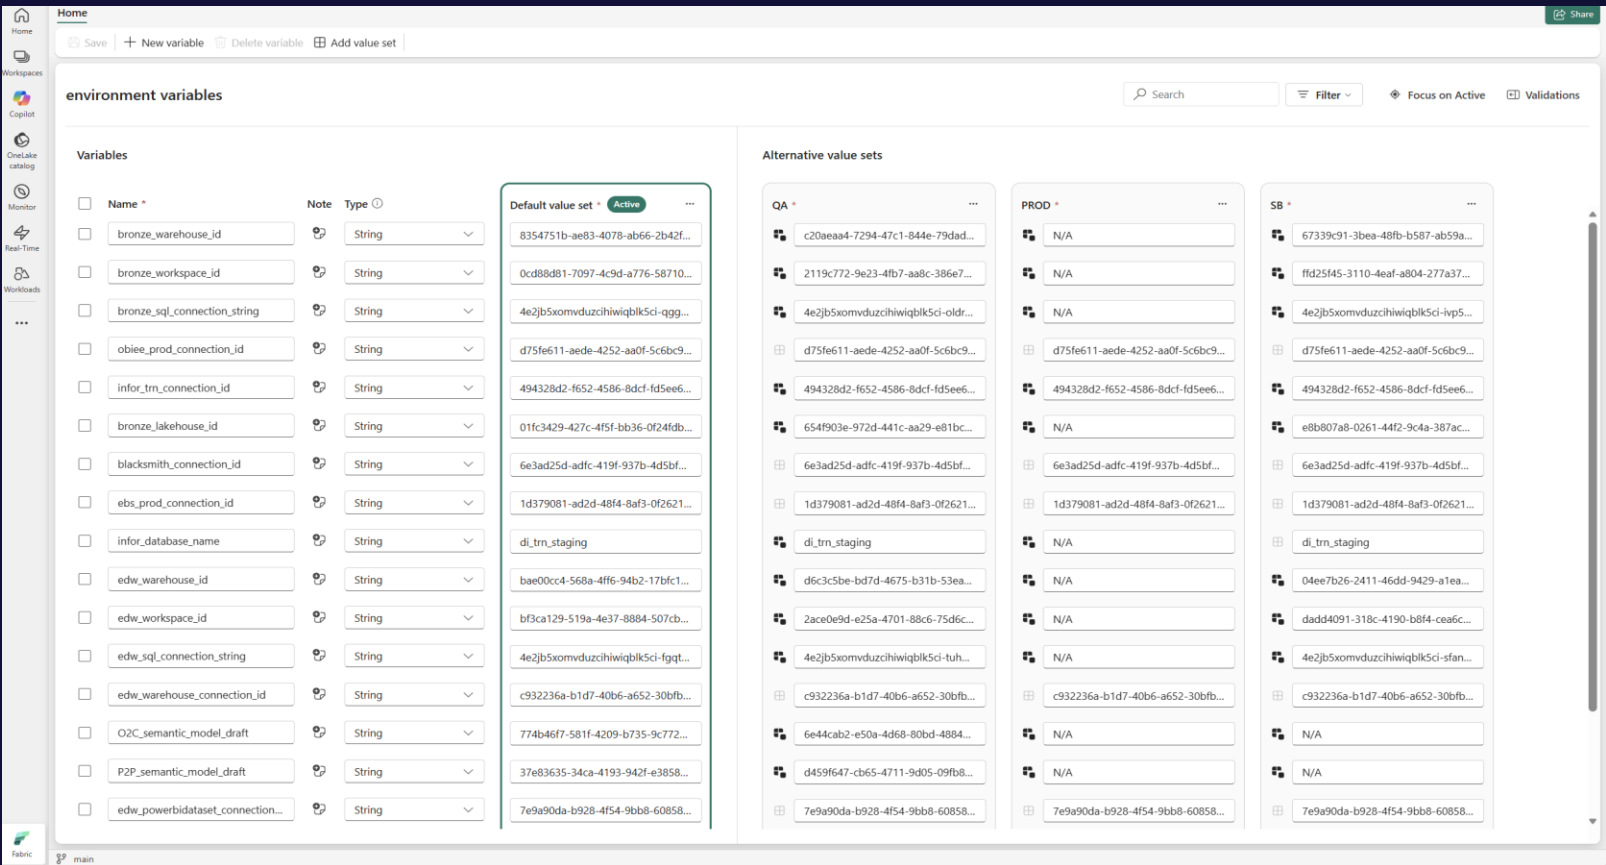





Collaborative Workflow

Variable libraries foster better collaboration between technical and business stakeholders for faster deployments.

Key Feature: Examples Configurations

NC1



FABCON | **SQLCON**
Microsoft Fabric | Microsoft SQL
COMMUNITY CONFERENCE | COMMUNITY CONFERENCE

#FABCONSQLCON2026

Data Engineering and Transformation Patterns

Preparing and Transforming Data: Then

Temporary & Staged Data Support

When initially released, and up until recently temporary tables were not supported. This created a need for persisting data tables mid process, causing increased workloads for moving data from source to target.

Key Management

Without support for IDENTITY columns, dimensional modelling required additional discovery and development. Aggregation and sequencing logic was required to maintain key logic, causing compute-heavy operations.

Merge Capabilities

Action-oriented patterns were relied on in order to identify changed records, causing heavier reads and compute. This pushed insert/update designs back to spark-driven workloads, increasing the barrier to entry.

Implementation Challenges

Remaining caveats include seeding issue, inability to rename object, and perform streamlined schema comparisons with the overall ETL design.

Preparing and Transforming Data: Now

Optimized Data Transformation

Temporary tables now available, streamlining the consolidation of data from various sources with additional transformation in a singular operation.

Simplified Key Identification

Recent introduction to IDENTITY columns systematically ease the management of historical data from an operational standpoint.



Insert & Update Enhancements

Row change tracking and merge capabilities allow for optimized upsert functionalities within the warehouse. Upsert* functionality is also available when utilizing a copy data activity.

Refining Solution Architecture

Additional enhancements such as truncating capabilities, renaming data assets and alter functionality enable developers to enhance and refine their ETL/ELT processes.

Key Feature(s): Enhanced ELT/ETL Capabilities

Nested CTE*

Complex, multi-layer transformations become maintainable, readable and structured with the ability to allow for nested CTEs. This will decrease development time, reduce maintenance, and generate faster insights.

Metadata Sync REST API

Fabric metadata can be programmatically discovered and synchronized through a REST API, enabling automated inventory, governance, and validation across environments. This reduces manual tracking, improves visibility into data assets, and strengthens platform governance at scale.

Case insensitive Collation Support

Case-insensitive collation aligns Fabric with real-world enterprise data behavior, eliminating common join and comparison issues caused by casing mismatches. This improves data quality, reduces transformation overhead, and lowers the risk of subtle reporting defects.

Materialized Lake Views & Direct Lake

Materialized lake views simplify your design by materializing your SQL transformation as refreshed Delta tables on top of OneLake . This ability to persist pre-computed logic greatly decreased the amount of compute necessary to perform more complex transformations, as they are cached in memory.

NC1

FABCON | **SQLCON**
Microsoft Fabric | Microsoft SQL
COMMUNITY CONFERENCE | COMMUNITY CONFERENCE

#FABCONSQLCON2026

**Lakehouse, Integration,
and Security**

Lakehouse Shortcuts & Source Control



Data Management with Lakehouse Shortcuts

Shortcuts simplify managing raw (bronze) data and workspace separation, leading to more organized and efficient workflows.



Source Control Tracking

Tracking shortcuts in source control improves collaboration and supports agile development practices within teams.



Aligning Tech with Business

Aligning technical capabilities with business processes ensures flexible, secure, and effective data warehouse solutions.

Ingestion & Integration Reality




Reduced Custom Integration

GA connectors minimize custom integration code, improving reliability and scalability of data ingestion.



Mirroring* vs Ingestion Tradeoffs

Understanding tradeoffs between data mirroring and ingestion is crucial for pragmatic integration decisions.



Hybrid Integration Patterns

Hybrid integration approaches are common, requiring tailored solutions to complex data landscapes.



Balancing Solutions & Features

Effective data architectures balance custom solutions with platform capabilities for adaptability.

Fast Fabric & D365 Acceleration



Accelerated Data Ingestion

Fast Fabric enables data refreshes under 15 minutes, vastly improving ingestion and processing speeds over CDC pipelines.



Architectural Benefits

Faster ingestion simplifies architecture, reduces operational complexity, and supports agile business decision-making.



Business Impact

Automated high-performance solutions enable rapid insights, enhancing responsiveness to business needs and improving outcomes.

Notebooks, Security & Key Vault



Secured Version Control

No custom connections and the ability to shortcut data allows for streamline data movement across workspaces, allowing for more enhanced version control and data sharing.



Secret Management

Native integration with Azure Key Vault allows secrets, tokens, and connection details to be centrally managed and securely referenced at runtime. This reduces hard-coded credentials, improves compliance posture, and standardizes secret rotation practices.



Governance and Monitoring Controls

Automated logging of all Fabric activities and the ability to monitor security violation attempts allows for a more proactive approach to identifying breaches.



Integrated Security Practices

Embedding security controls directly into notebooks and data workflows enforces least-privilege access and protects sensitive data by design. This ensures authorized access while aligning day-to-day development with enterprise security standards.

FABCON | **SQLCON**
Microsoft Fabric | Microsoft SQL
COMMUNITY CONFERENCE | COMMUNITY CONFERENCE

#FABCONSQLCON2026

Notable Mentions & Key Takeaways

Notable Mentions



Managed Private Endpoints

Securely connect to customer-managed data sources through managed network isolation and routing. Reducing infrastructure overhead while enabling governed ingestion from private data platforms.



Private Network Access to Fabric

For external services and users that want to access data inside Fabric over private endpoints, rather than public IPs. This supports enterprise security models by keeping analytics workloads and consumers within trusted network boundaries.



Fabric IQ & Ontology*

Fabric Introduce an understanding of data assets, relationships, and business meaning across the platform. Data governance, lineage, and discoverability become seamless through aligning metadata with business context and AI-assisted insights.



Autoscale Features*

Autoscale capabilities enable consumption smoothing, optimizing resource use and cost efficiency in data warehouses.

Key Takeaways and Community Impact



Avoiding Past Mistakes

Encourage revisiting early implementations and leveraging new capabilities to prevent repeating previous errors.



Platform Enhancements

New platform features eliminate many (but not all) prior workarounds, streamlining development and optimizing performance.



Community Collaboration

Continuous learning and knowledge sharing strengthen the Fabric ecosystem and drive innovation.



Driving Business Value

Applying insights boosts efficiency, scalability, and business value in data warehouse projects.

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
Uhova for your chance to WIN
PRIZES!

