


CF.Cumulus™

An open-source cloud data platform product to accelerate time to insights.

Cloud Formations

What is CF.Cumulus

Cloud Formations






CF.Cumulus Framework



- **Open Source*:**
Zero cost, unlimited access to code
- **Scalable:**
Horizontal & vertical scaling for Big Data executions
- **Cost-Efficient:**
Faster time-to-insight reduces delivery costs
- **Extensible:**
Supports complex & unique data handling
- **Collaborative:**
Brings development & operations together
- **Rapid Development:**
Accelerates data solution implementation
- **Proven Framework:**
Adopted and used with 100+ repository forks
- **Transparent & Trustworthy:**
Tried & tested patterns using mature technologies



Control

- Metadata-driven workflow orchestration
- Managed dependency chains
- Concurrent batch processing
- Multi-threaded execution
- Resilient exception handling
- Enterprise grade reliability

Ingest

- Metadata-driven data ingestion
- Adaptable to all scenarios
- Extensible data sources & formats
- Flexible operating modes:
 - Deploy-time code gen
 - Run-time code gen[†]

Transform

- Metadata-driven data transformation
- Rapid engineering of common transformations
- Scalable data processing
- Baked in support for dimensional models
- Supports AI & Machine Learning model inference
- Complete flexibility through custom notebooks

* In accordance with MIT licencing.

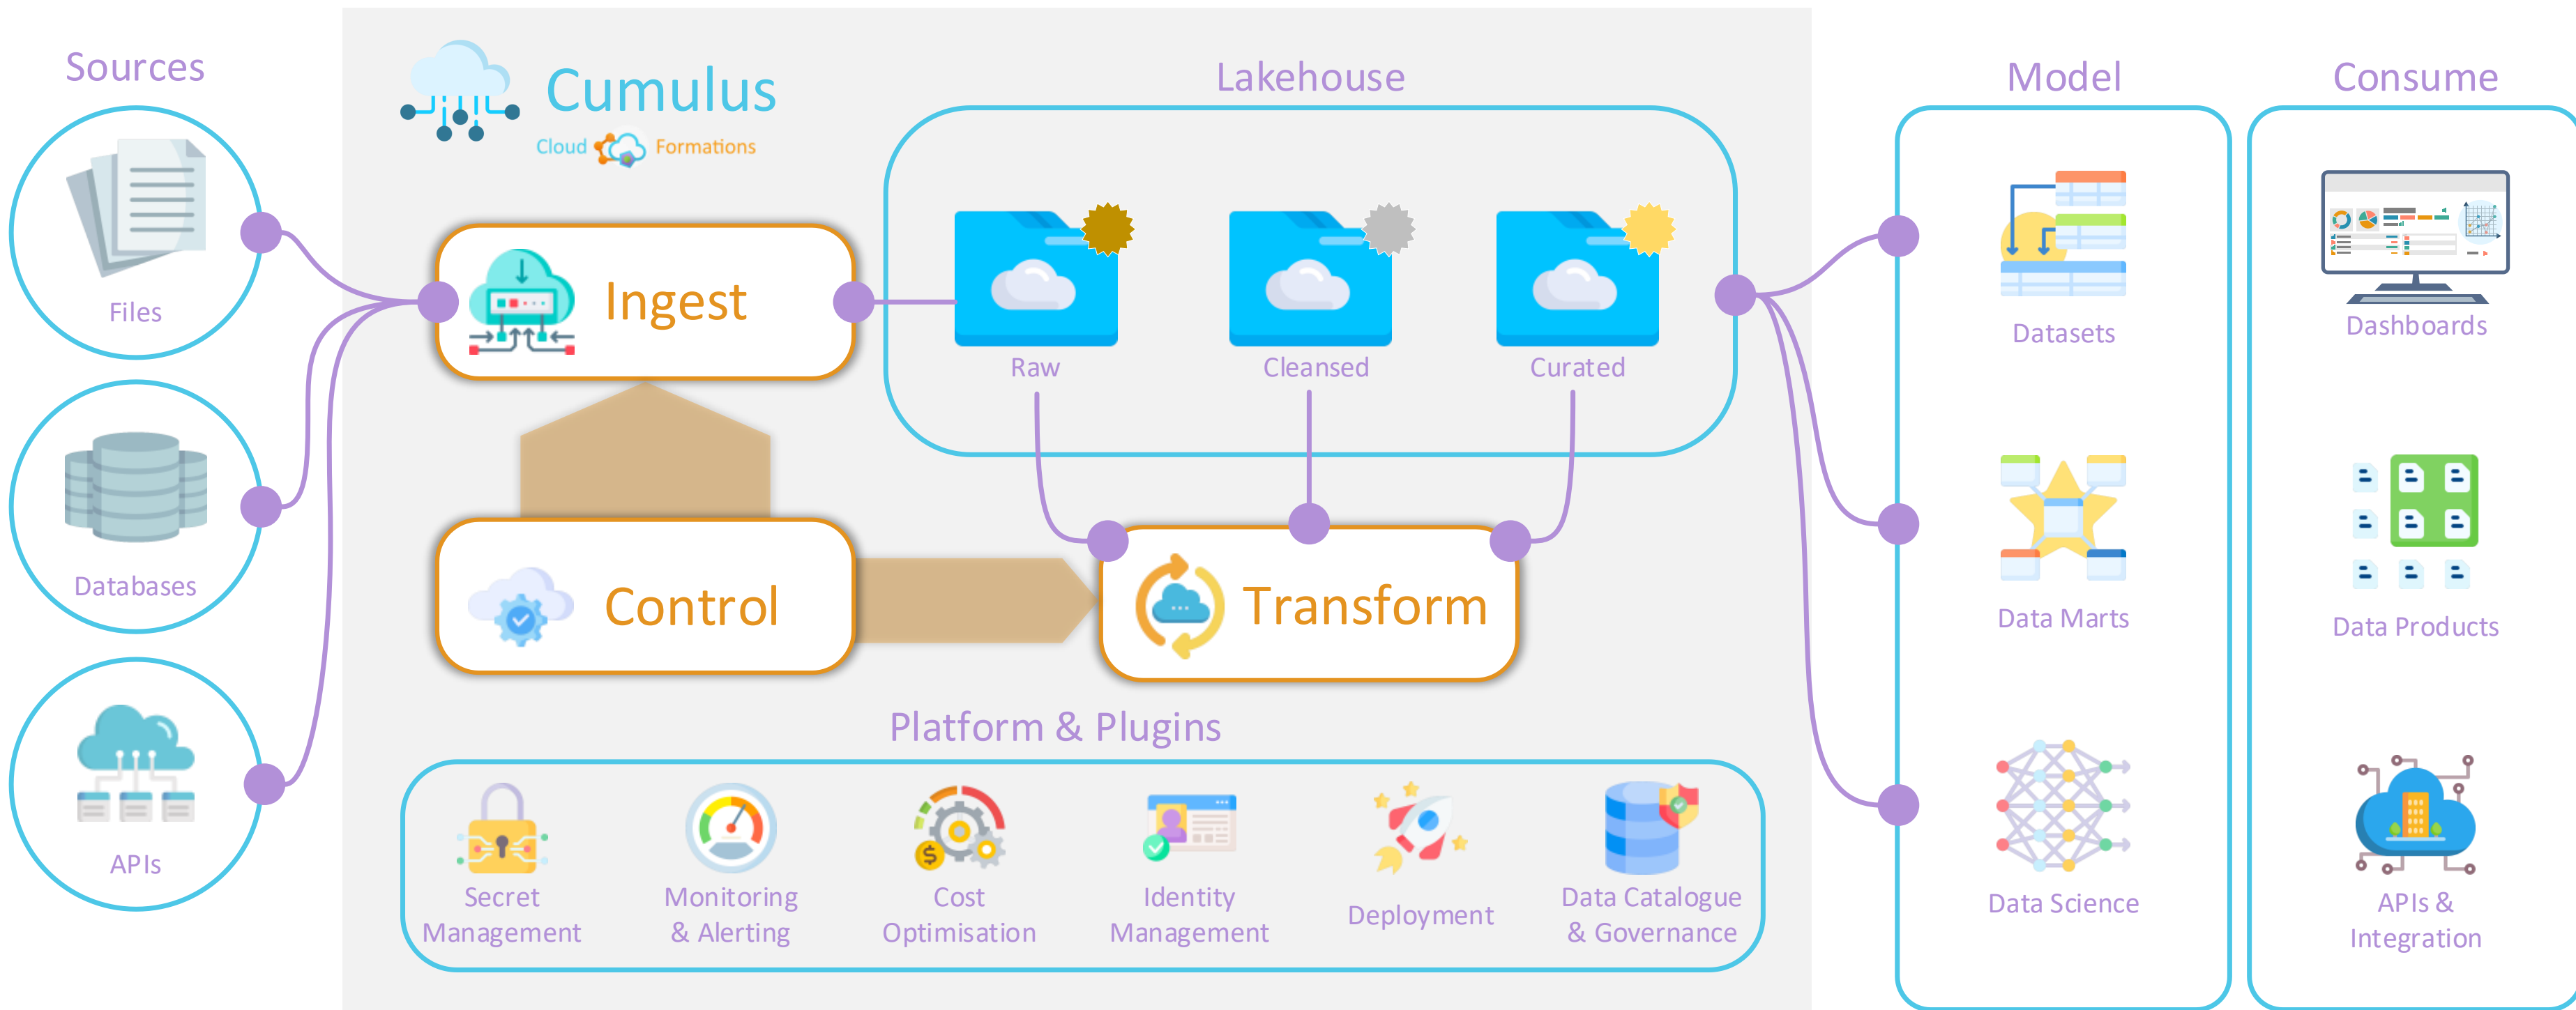
† Coming soon.



CF.Cumulus Framework



Cloud Formations - Logical Architecture



Underlying Technology: Spark DELTA LAKE C# JSON

Works With:



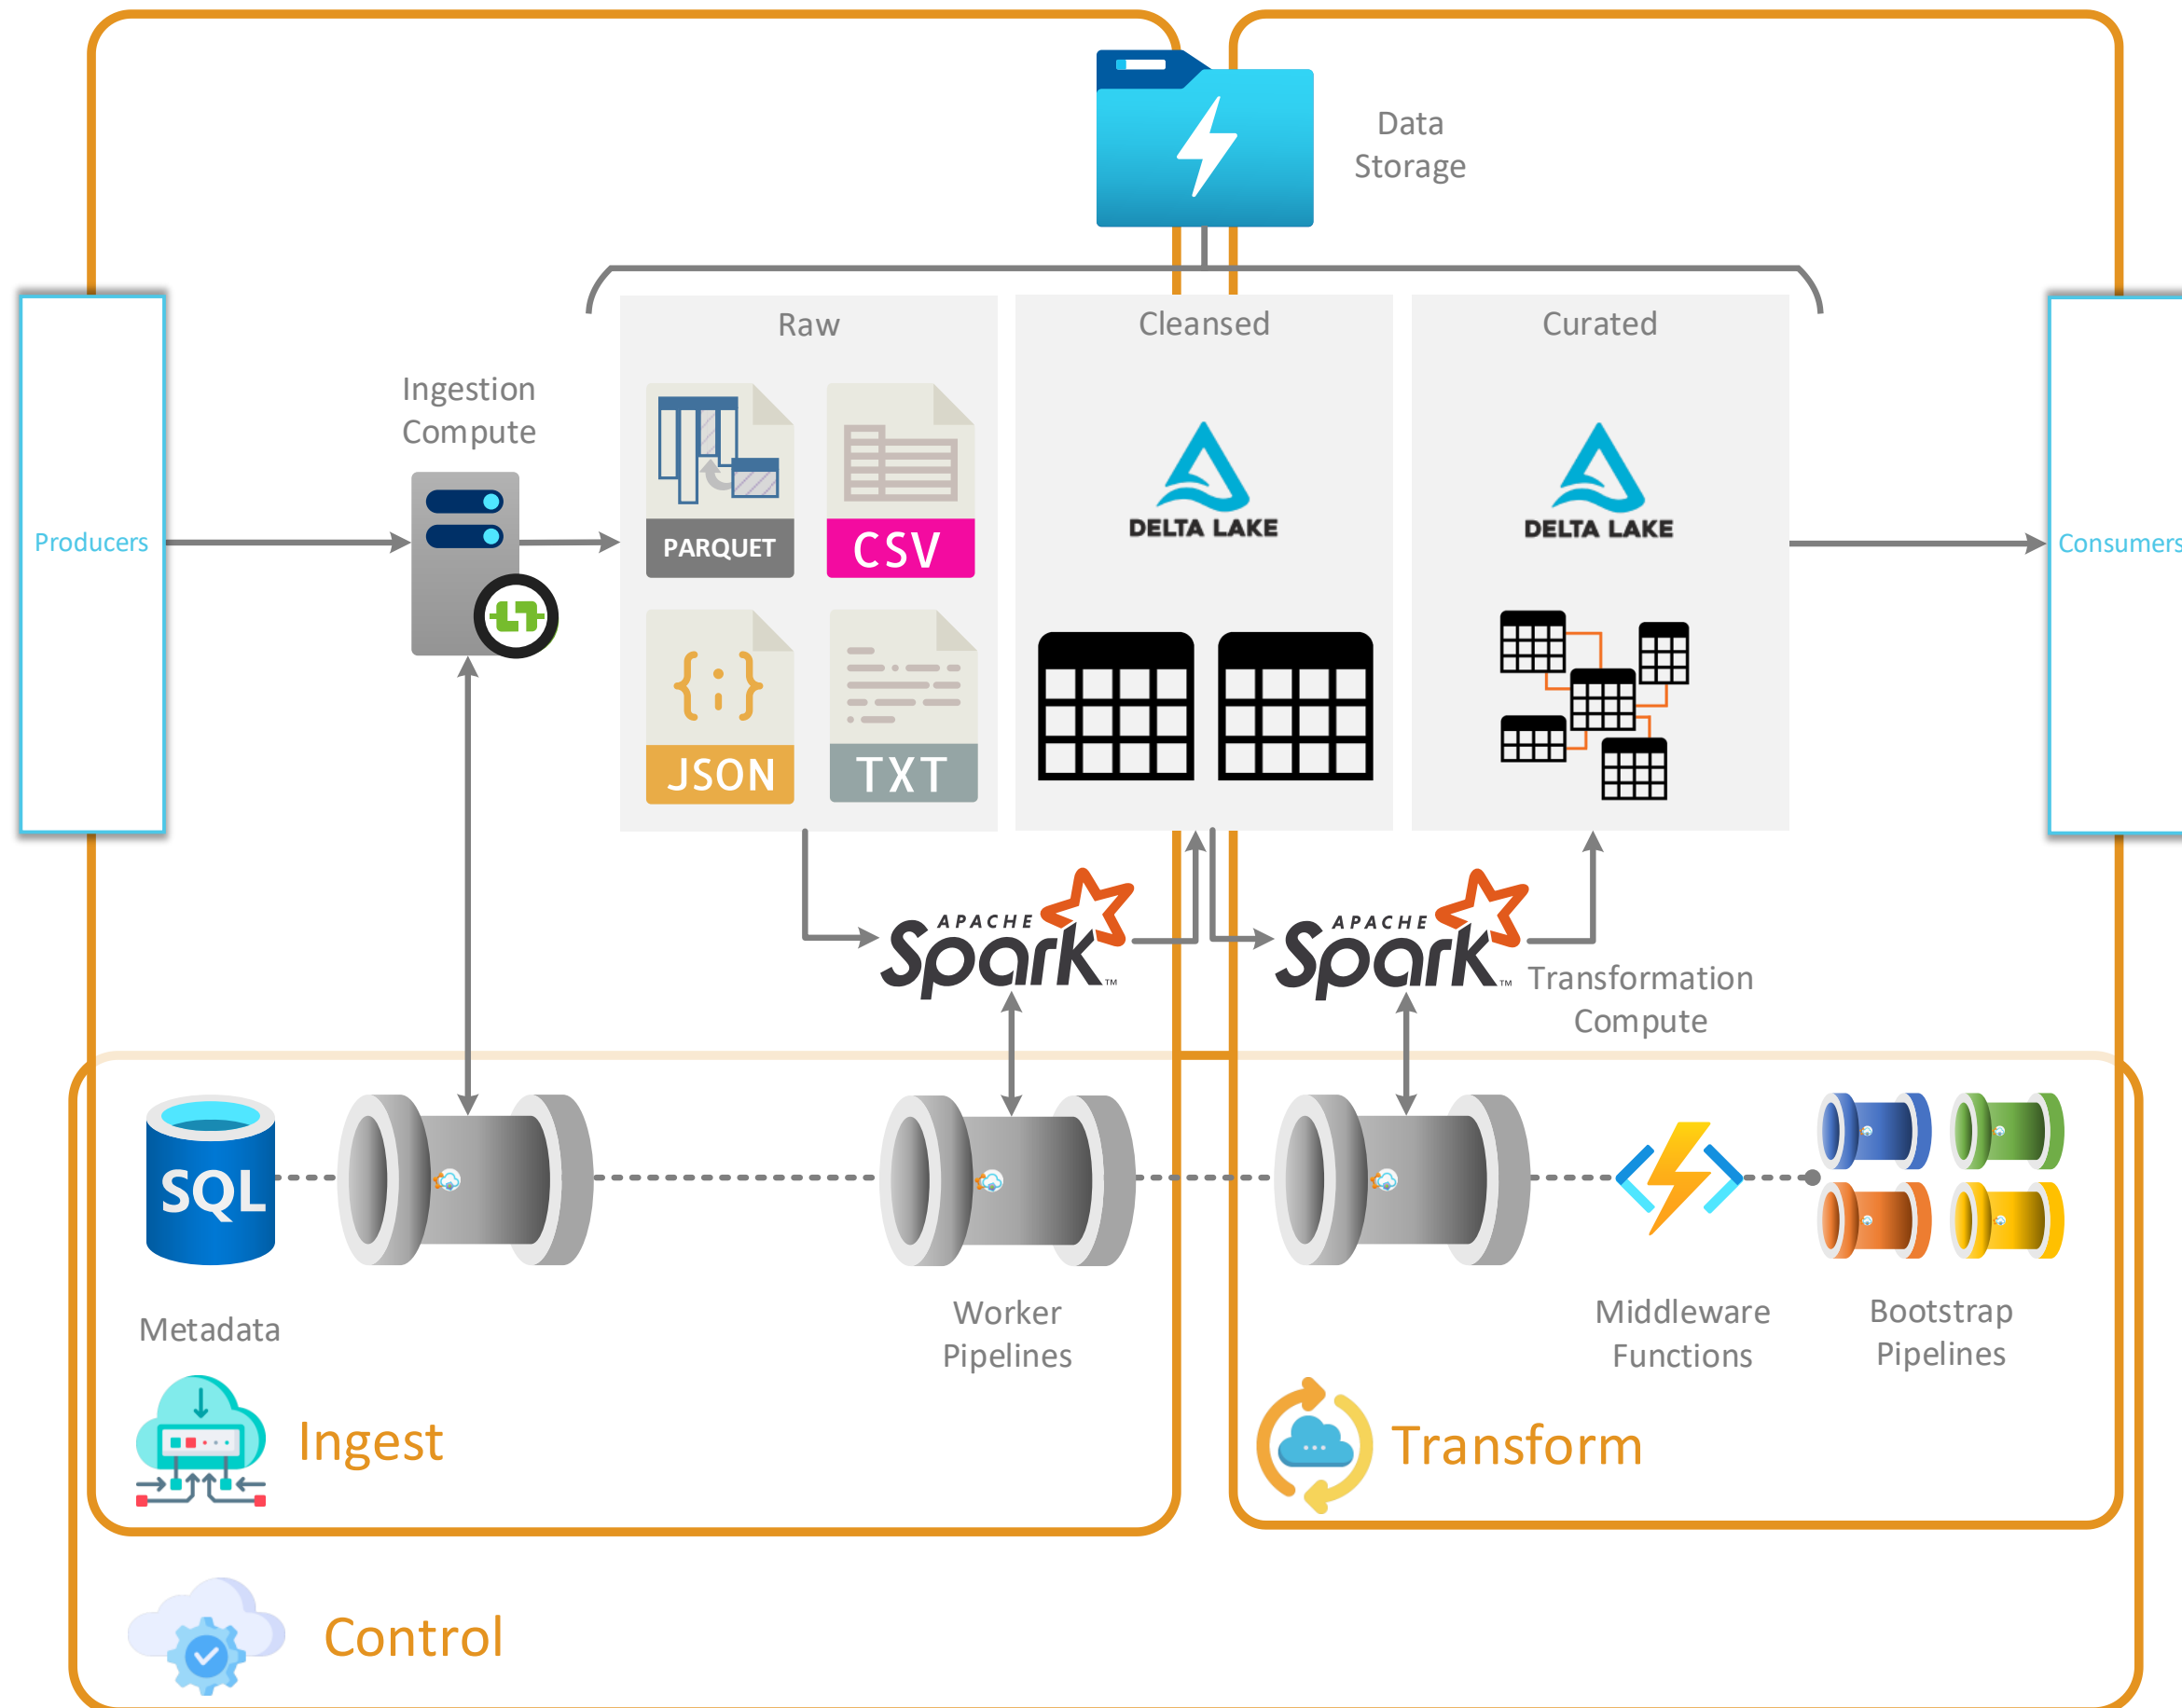
CF.Cumulus Framework



An Azure product agnostic view of compute, storage and orchestration resources as used by CF.Cumulus.

A primary design focus for CF.Cumulus is to align with open-standards and open technologies. Offering a flexible, plug and play approach to the core components.

In this context, using Apache Spark compute and Delta Lake structures for the storage layer.

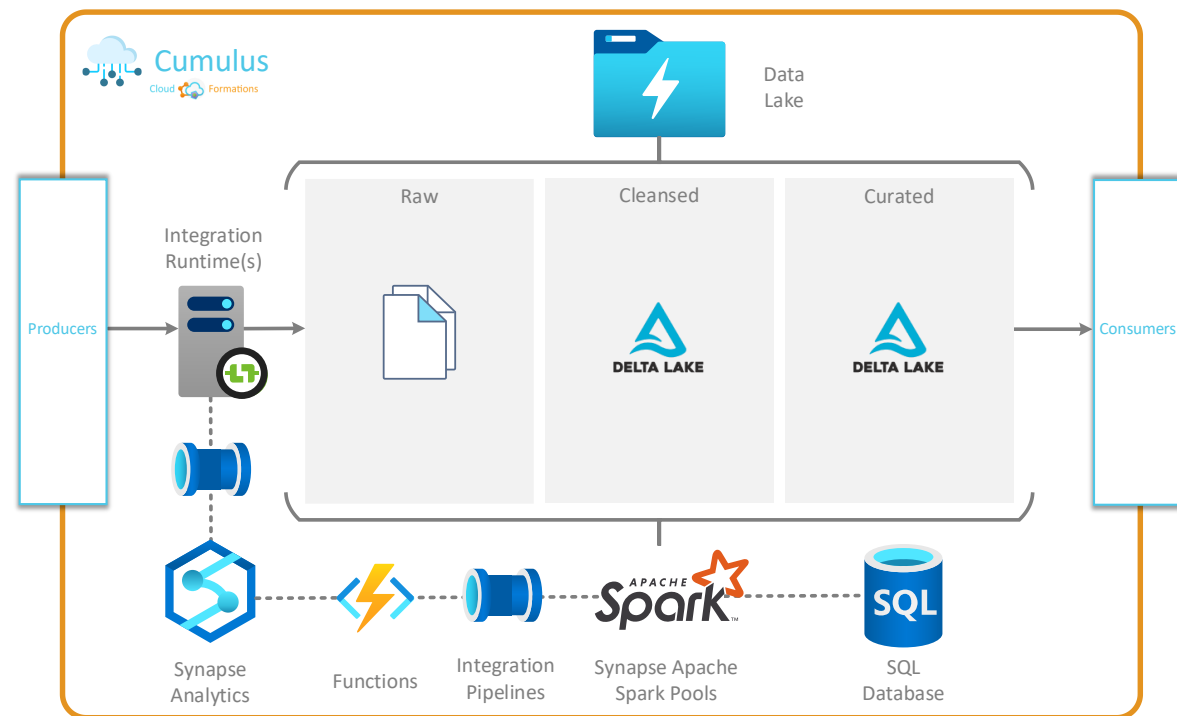




Flexible Technology Choices



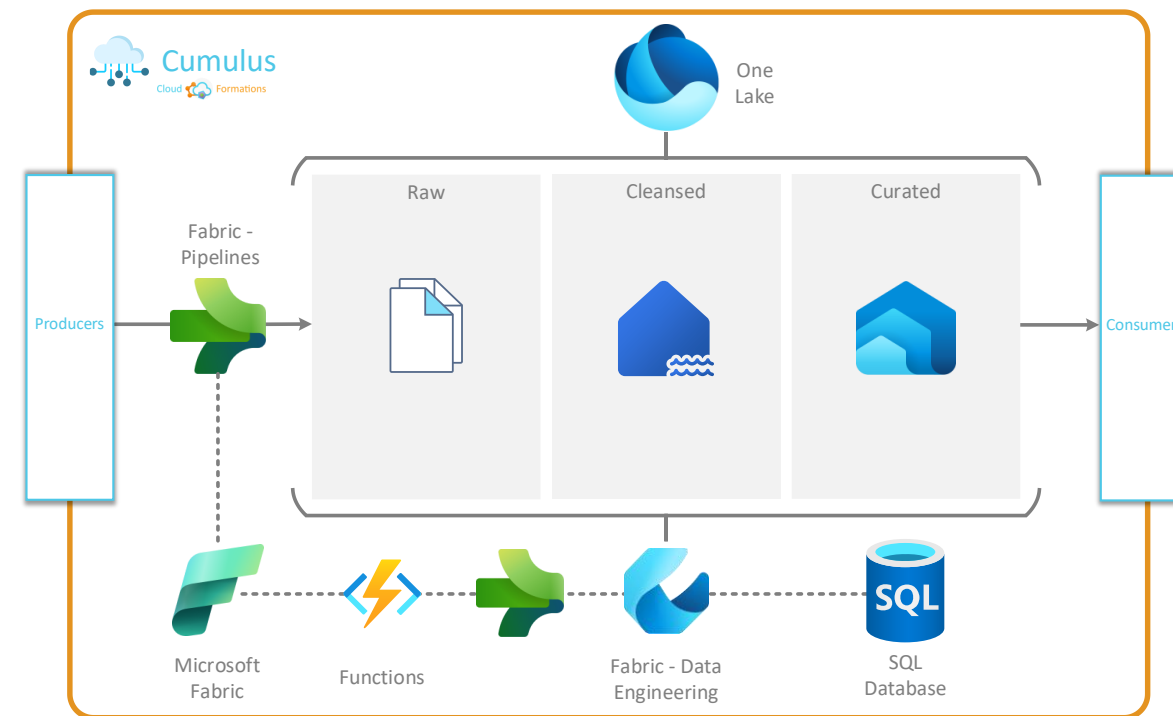
Building on the generic visuals above, but now applying a pure Microsoft product perspective to the logical delivery of Cumulus. The following diagrams could apply to a solution depending on the choice of tooling for Orchestration, Compute and Storage with each offering different advantages and coupling of resources.



Using Azure Synapse Analytics

This implementation represents a partly decoupled architecture, in terms of cloud resources used to delivery Cumulus.

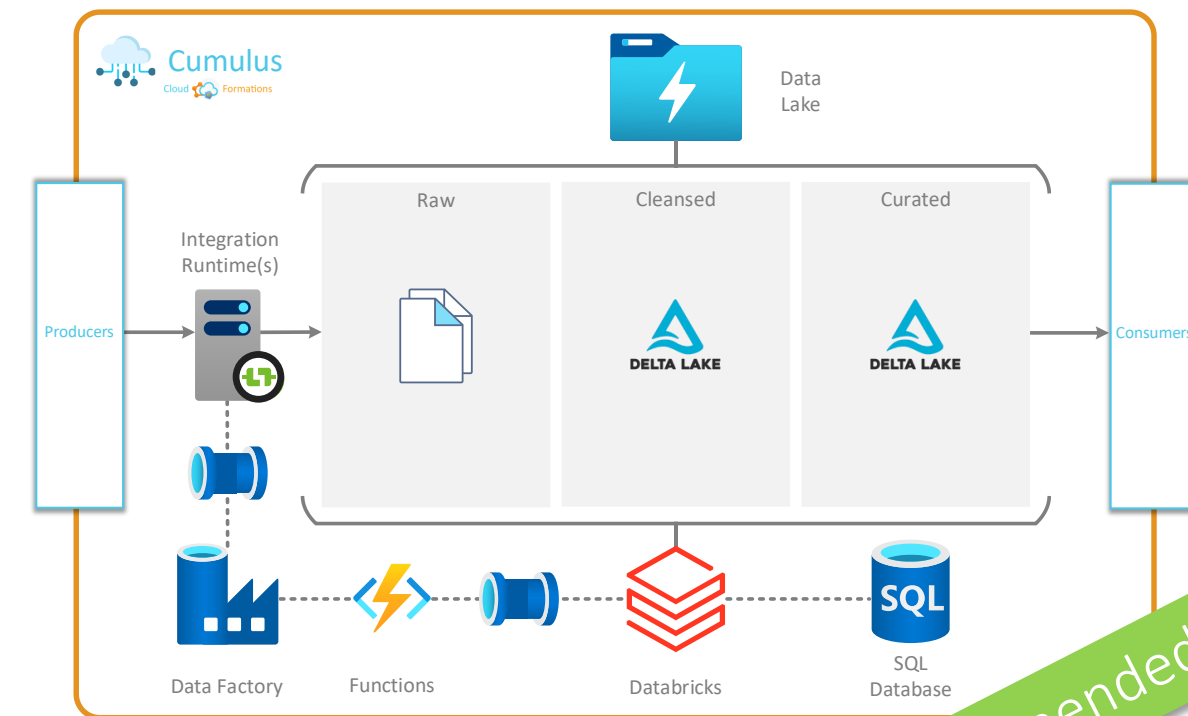
- **Orchestration** and **Compute** delivered using Azure Synapse Analytics.
- **Storage** delivered using Azure Data Lake.



Using Microsoft Fabric

This implementation represents a coupled architecture, in terms of cloud resources used to delivery Cumulus.

- **Orchestration, Compute** and **Storage** delivered using Microsoft Fabric experiences.



Using Azure Data Factory & Azure Databricks


This implementation represents a fully decoupled architecture, in terms of cloud resources used to delivery Cumulus.

- **Orchestration** delivered using Azure Data Factory.
- **Compute** delivered using Azure Databricks.
- **Storage** delivered using Azure Data Lake.

Recommended

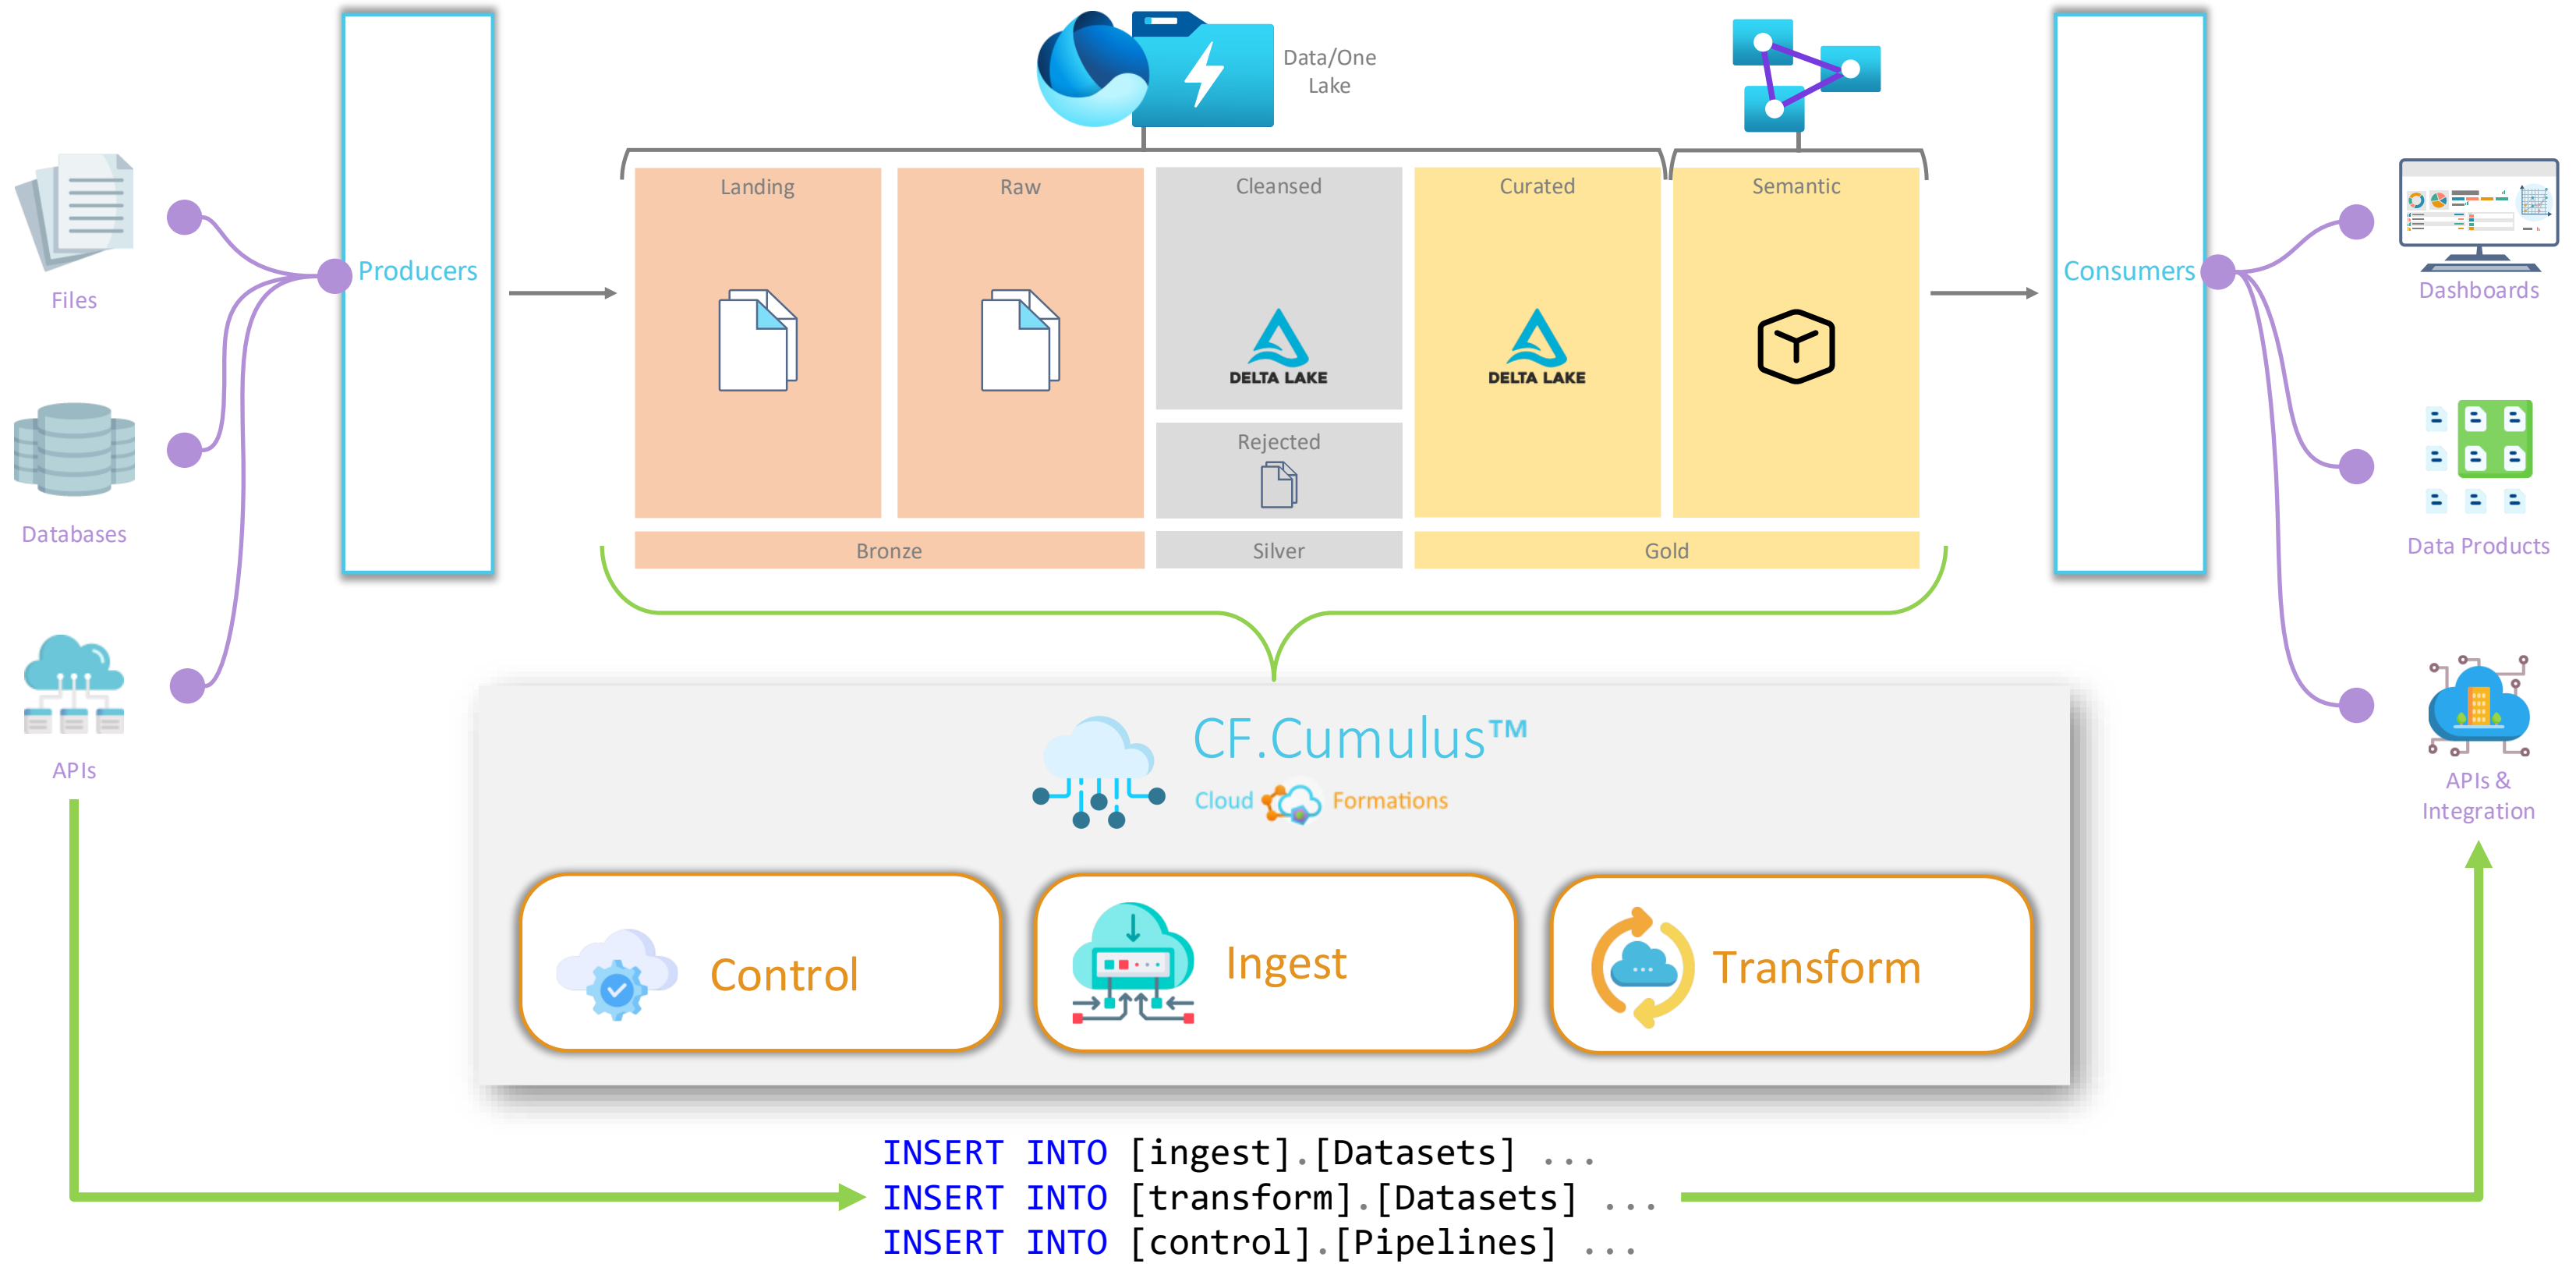
Why CF.Cumulus

Cloud Formations



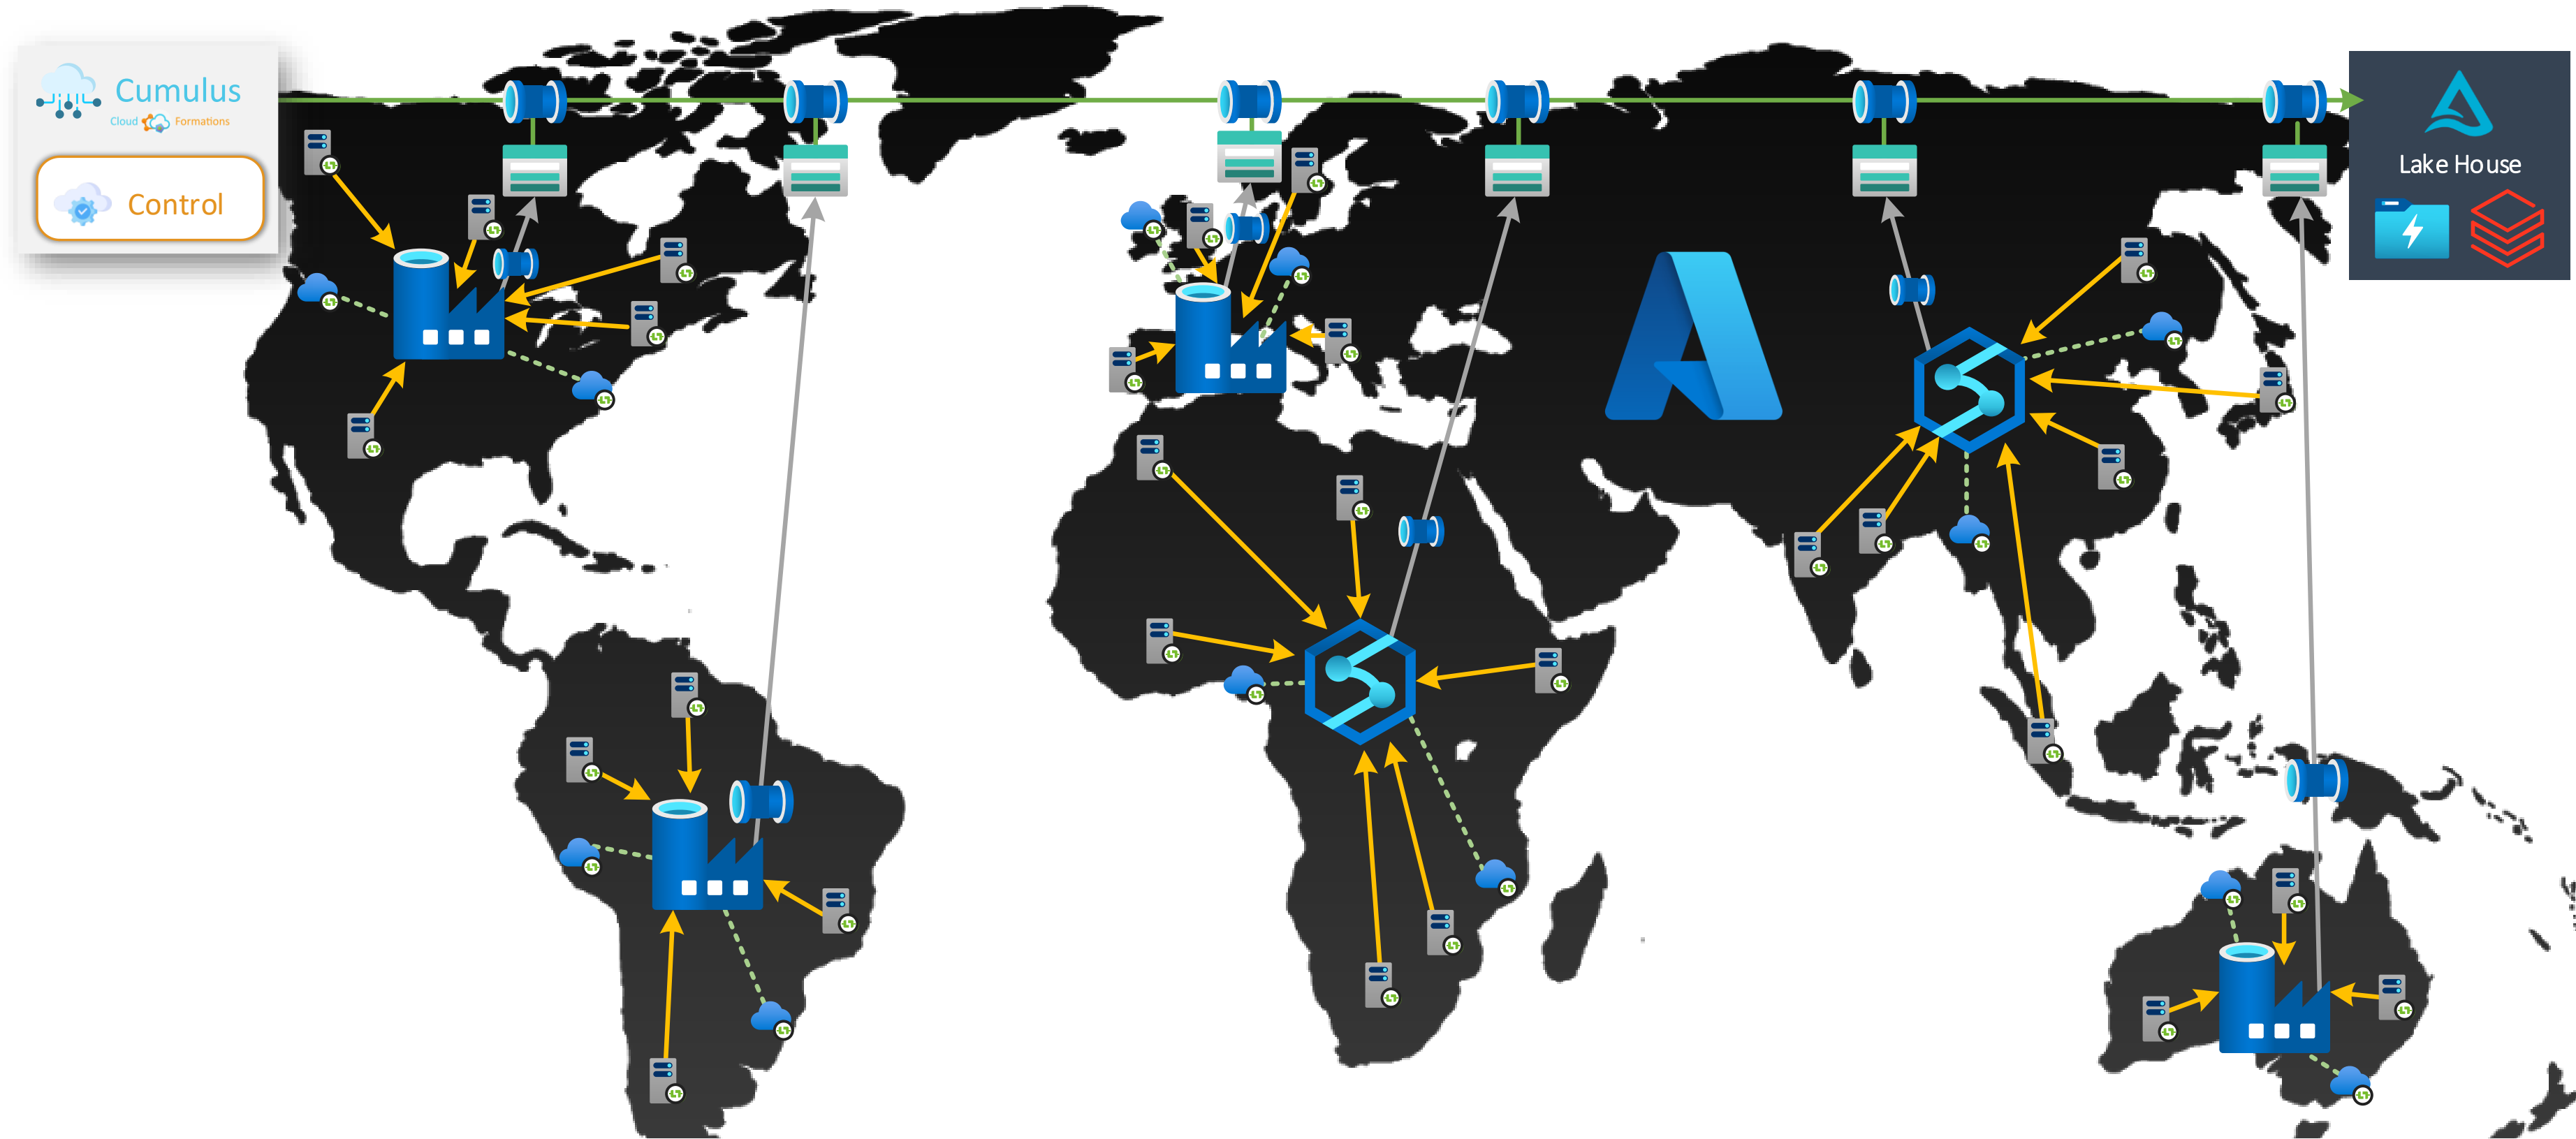
Lakehouse Implementations - Simplified Through Metadata

New source datasets delivered from producers to consumers with ease using only a few lines of metadata and open cloud native tooling.



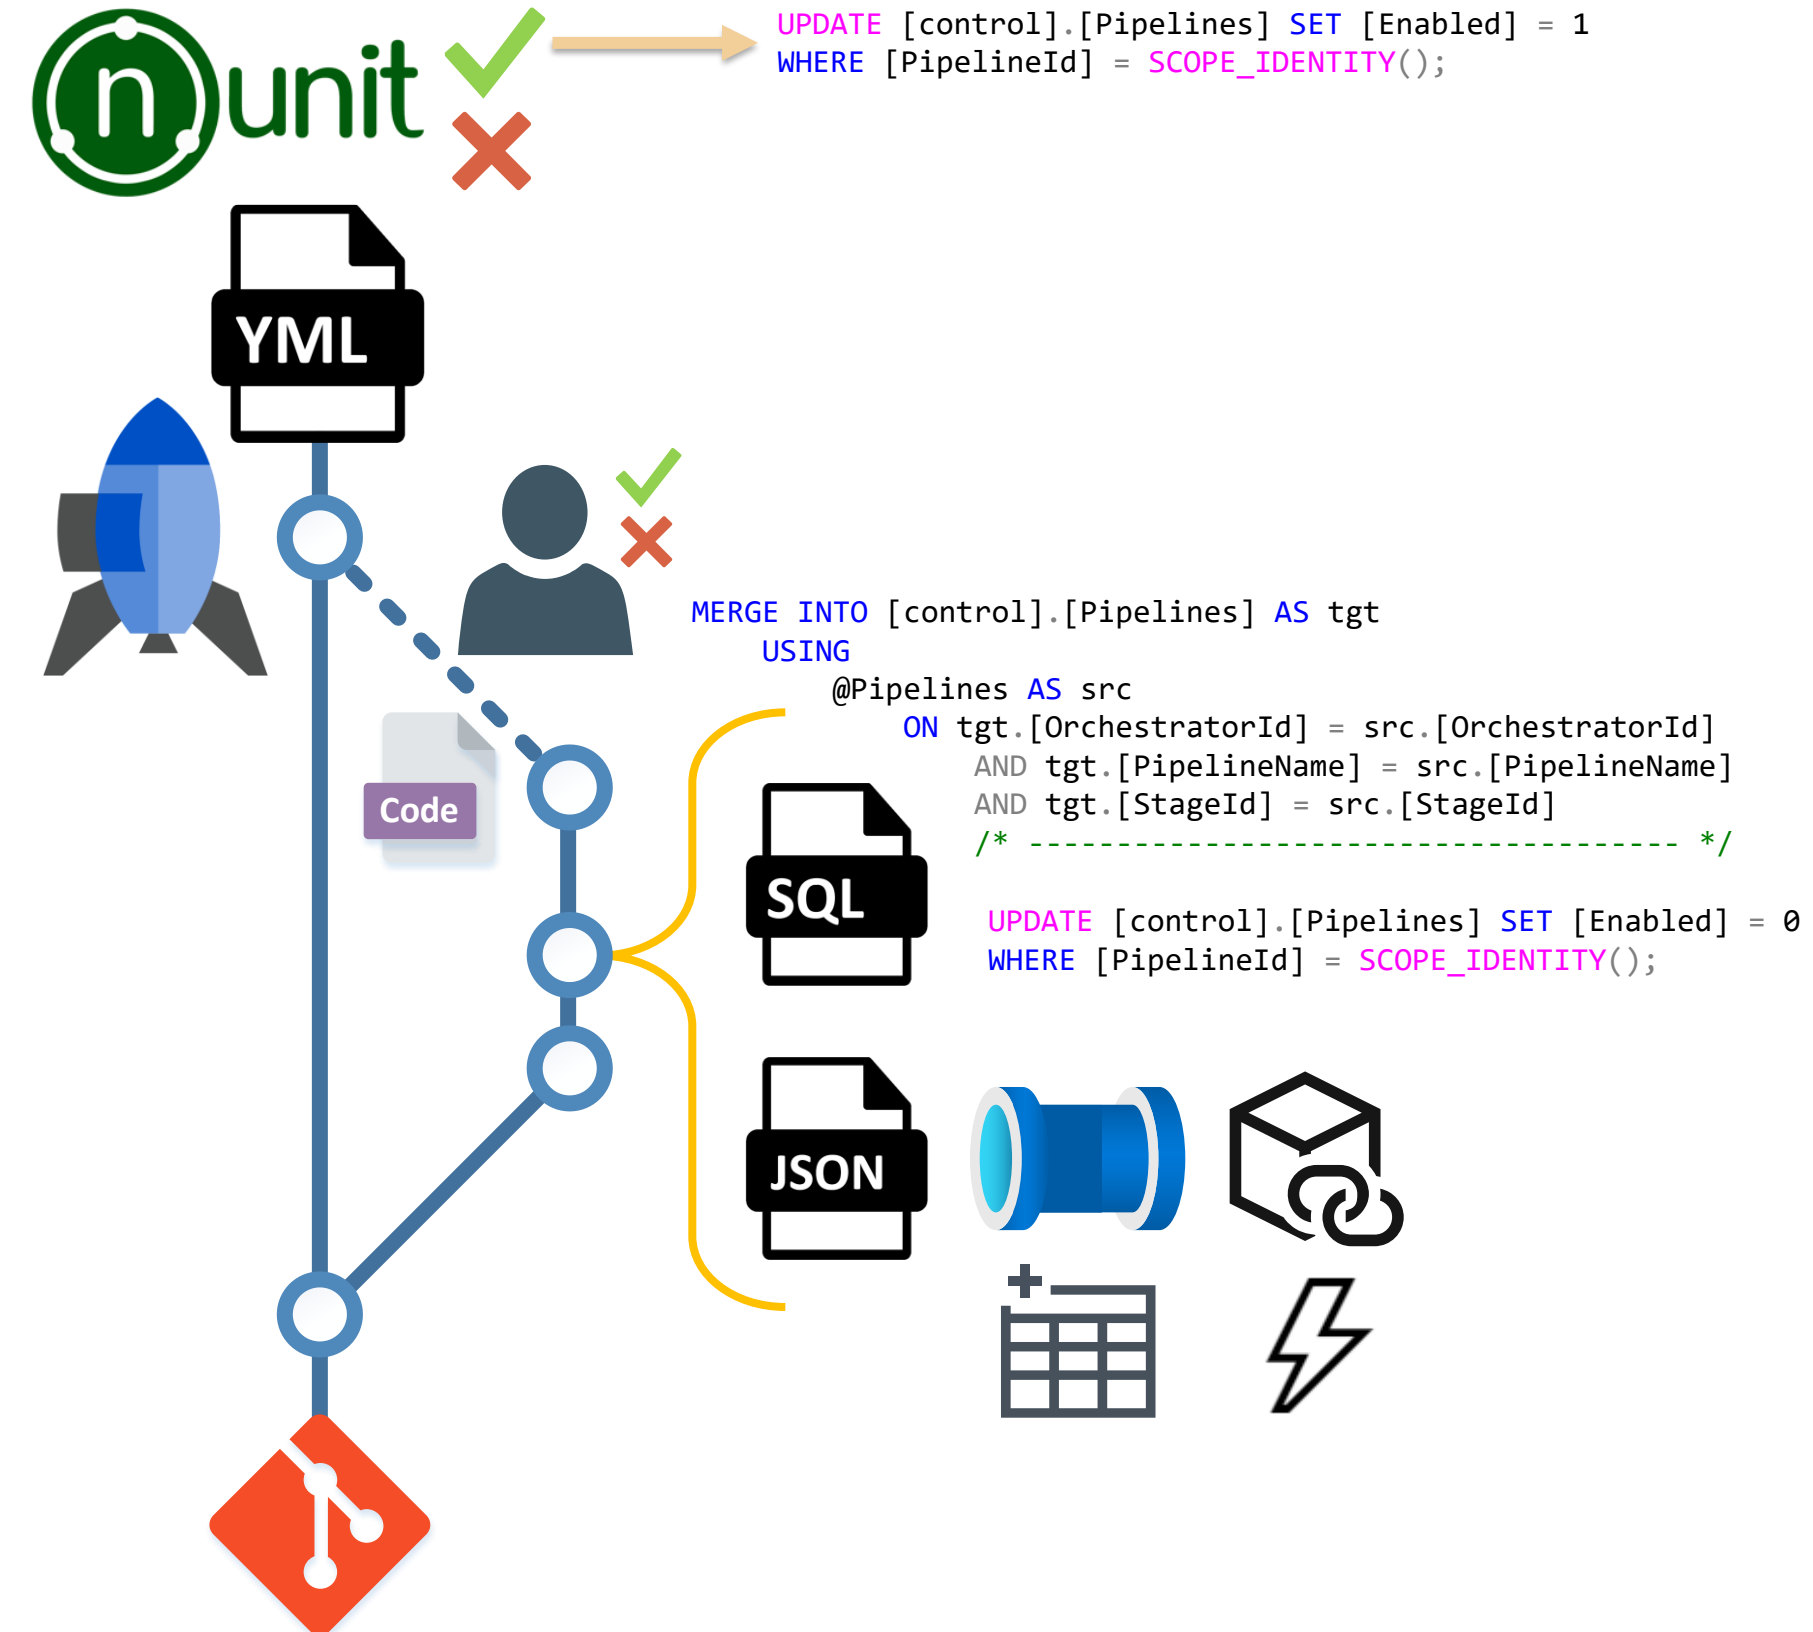
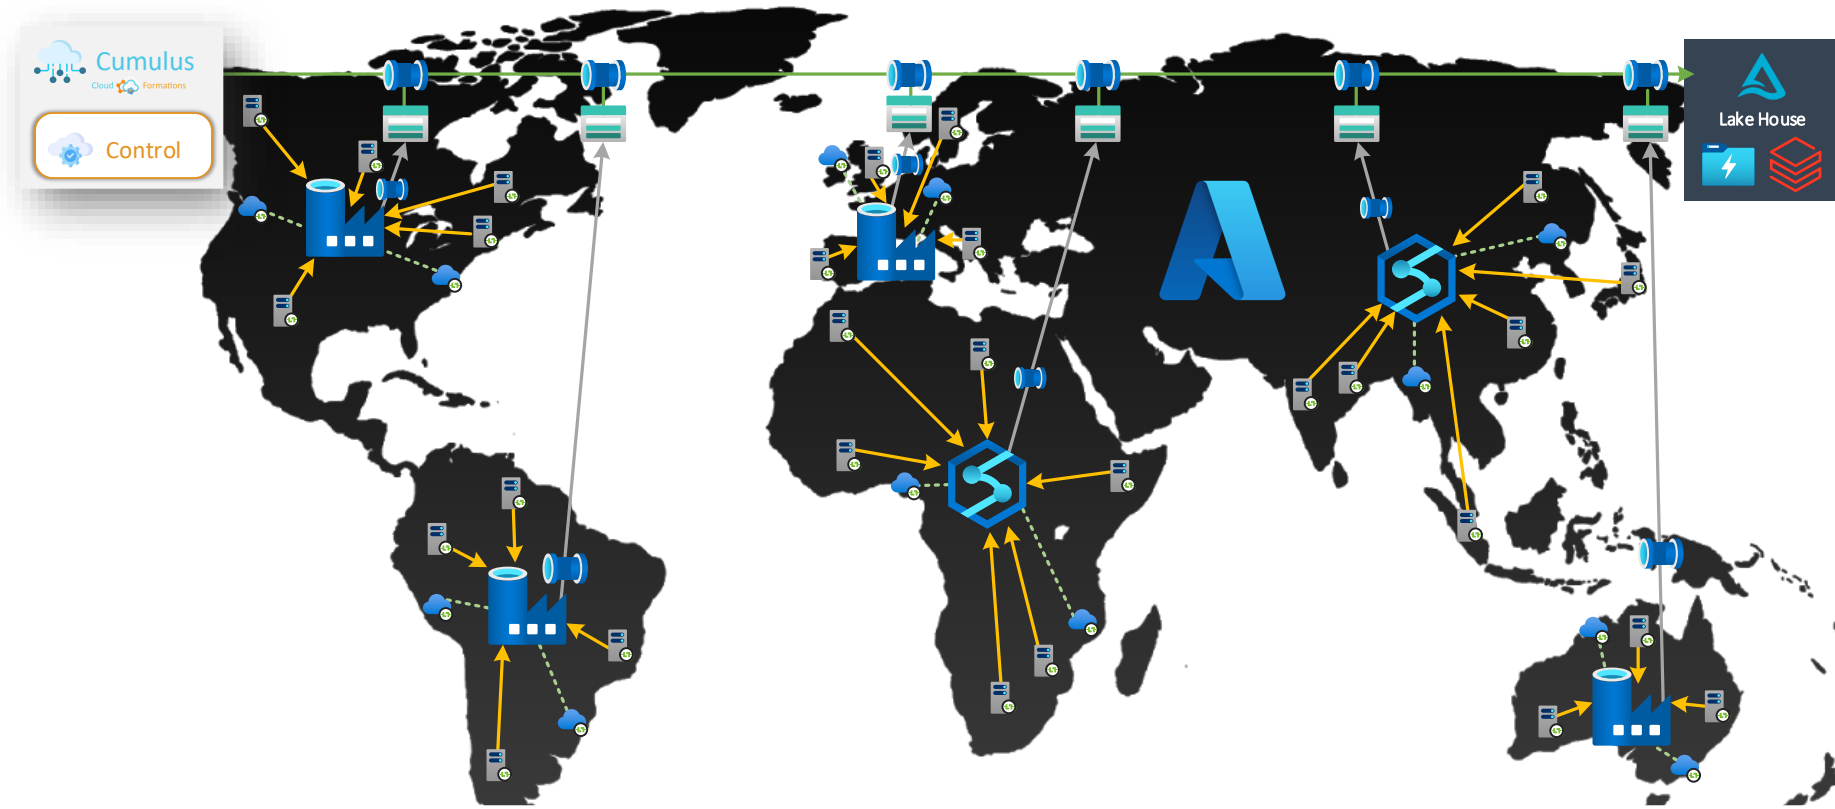
Lakehouse Deployments - Simplified Through Metadata

Multinational companies used CF.Cumulus to deliver cloud native orchestration at scale across 6 Azure Regions and 25 data source locations.

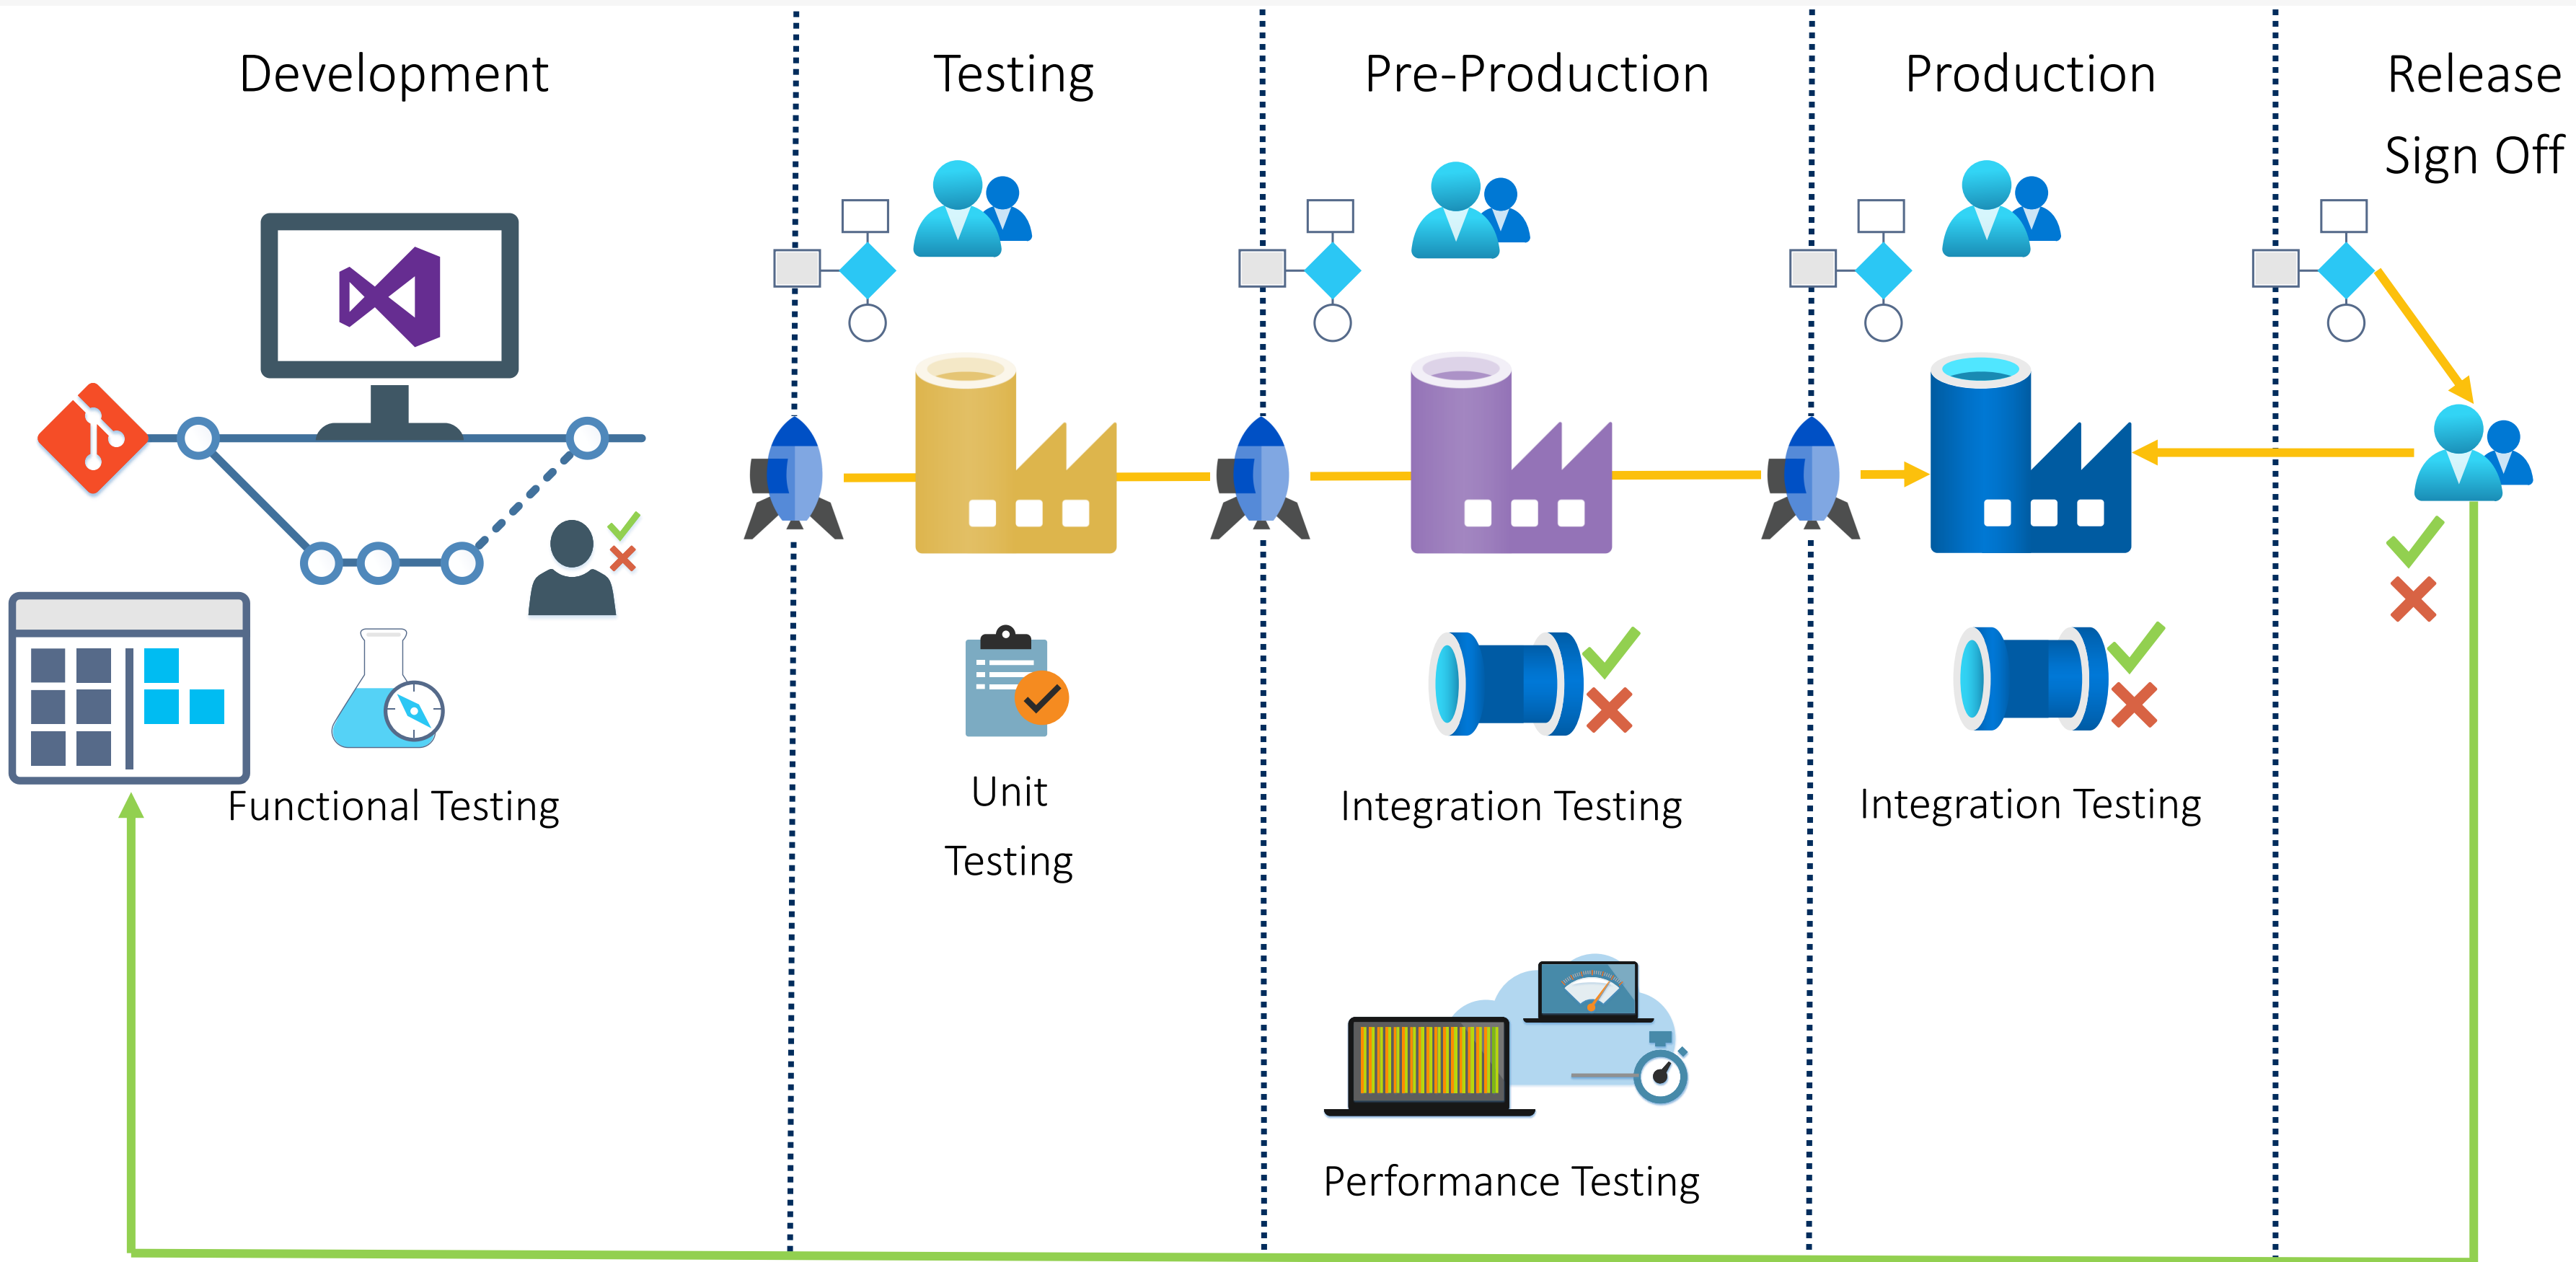


Lakehouse Deployments - Simplified Through Metadata

Multinational companies used CF.Cumulus to deliver cloud native orchestration at scale across 6 Azure Regions and 25 data source locations.



Lakehouse Deployments - Simplified Through Metadata













CF.Cumulus Indicative Run Costs



Assuming the use of an Azure UK Region priced for GBP, with all CF.Cumulus compute processes running for 4 hours every day of the year, with locally redundant storage, the following Azure run costs **per month** can be assumed:

SQL Database 250GB of Metadata £12.08 	Data Factory Azure IR Usage £25.84 	Functions 2 Node Dedicated App Service Plan £44.30 	Databricks 3 Node Cluster with 4 vCPUs & 16GB RAM £152.92 
Blob Storage 1TB of Log Data £18.45 	Data Lake 1TB Analytics Data £38.82 	Key Vault 50,000 Requests £0.14 	Monitoring 500MB Log & App Insights Data per Day £45.61 

£338.16*

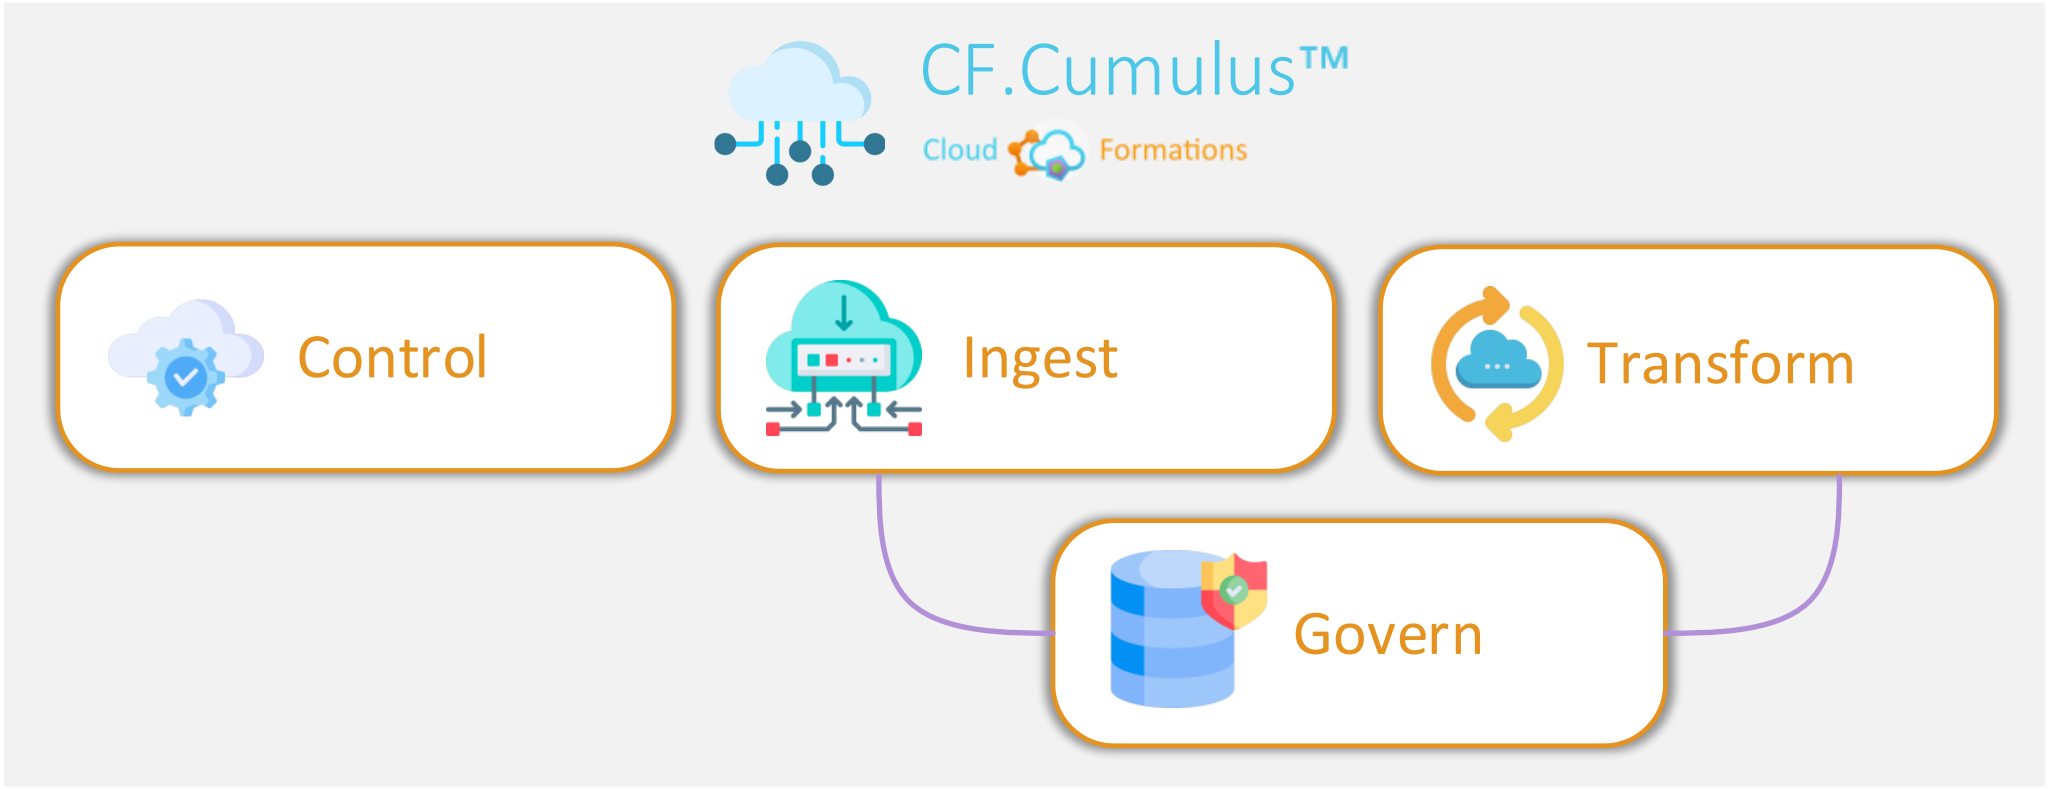
*With every customer implementation data volumes, table complexity and serving queries (to dashboards) may differ. All affecting the above run costs.

Azure Price Calculator workings available for review and adjustment: <https://azure.com/e/30f8aeedfd564ba189d33f278845bb89>

CF.Cumulus – Govern Plugin



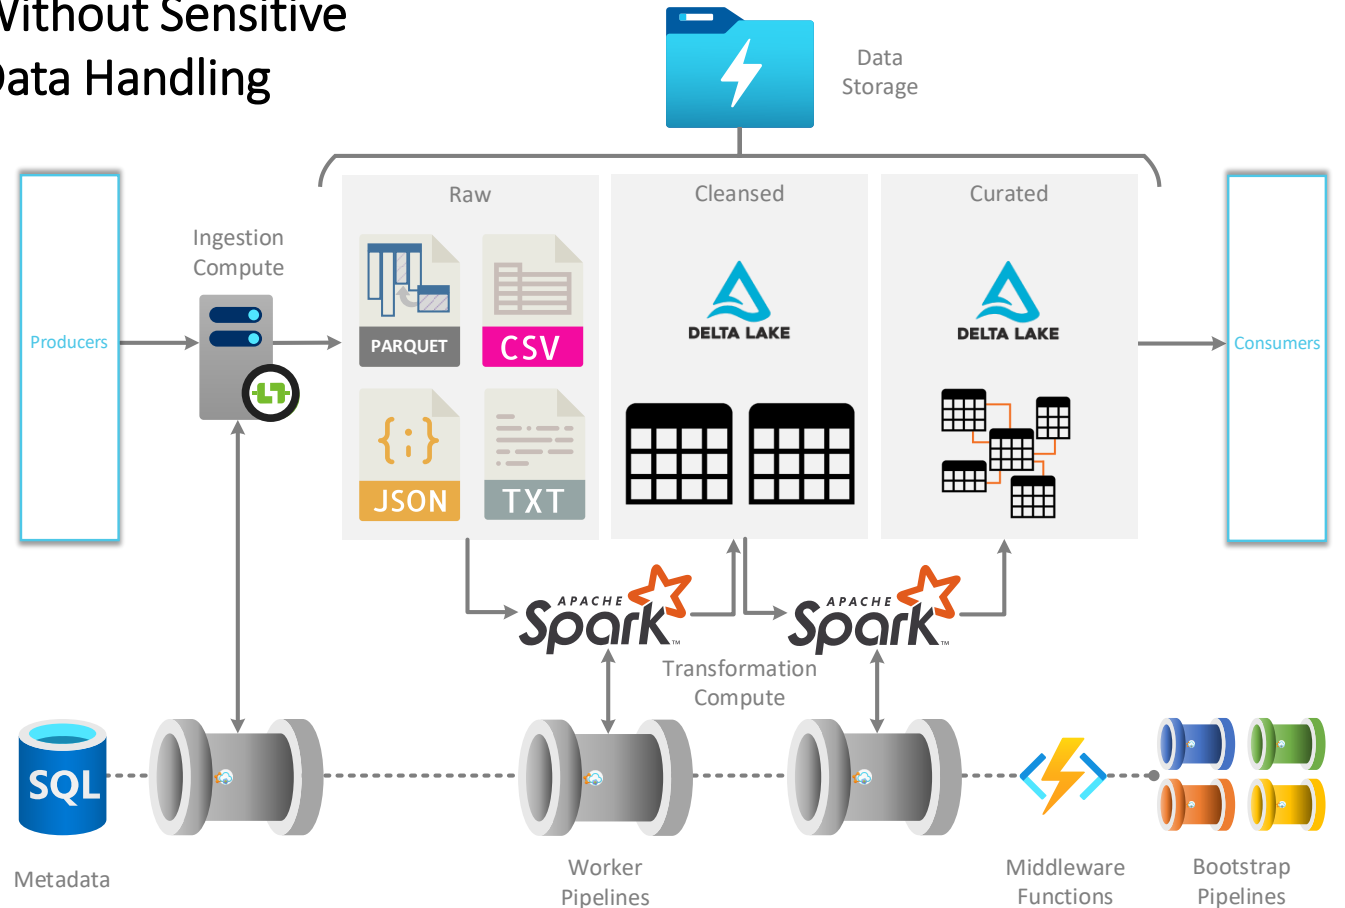
Cloud Formations - Logical Architecture Overview



Govern

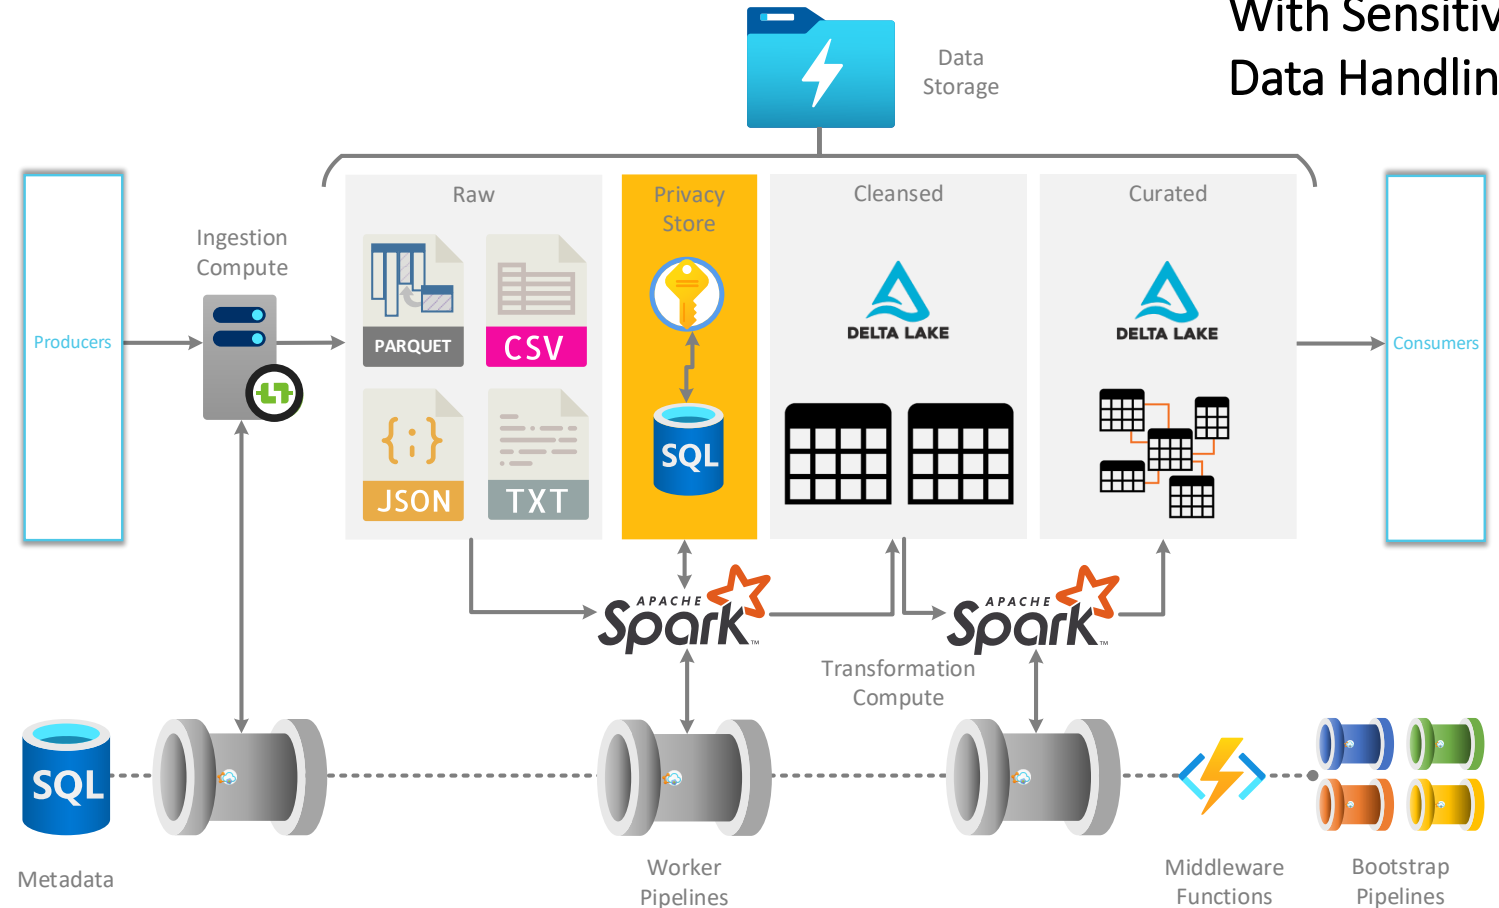
- Metadata-driven sensitive data handling.
- Apply ingestion data policies to support GDPR use cases.
- Custom value encryption and decryption.
- Custom value type handling for all sensitive values.
- Downstream reporting content/aggregations maintained.
- Isolation of sensitive data into a dedicated Privacy Store.
- Implemented data retention policies across Raw data.
- Synchronous and asynchronous, performant value handling.
- Integration with CF.Cumulus.Ingest and Control.
- API value retrieval for trusted Workers.

Without Sensitive Data Handling



A logical view of the CF.Cumulus product architect with and without sensitive data handling implemented as a plugin for the Ingest and Control components. Relational values diverted to a dedicated Privacy Store with custom encryption and full GDPR support.

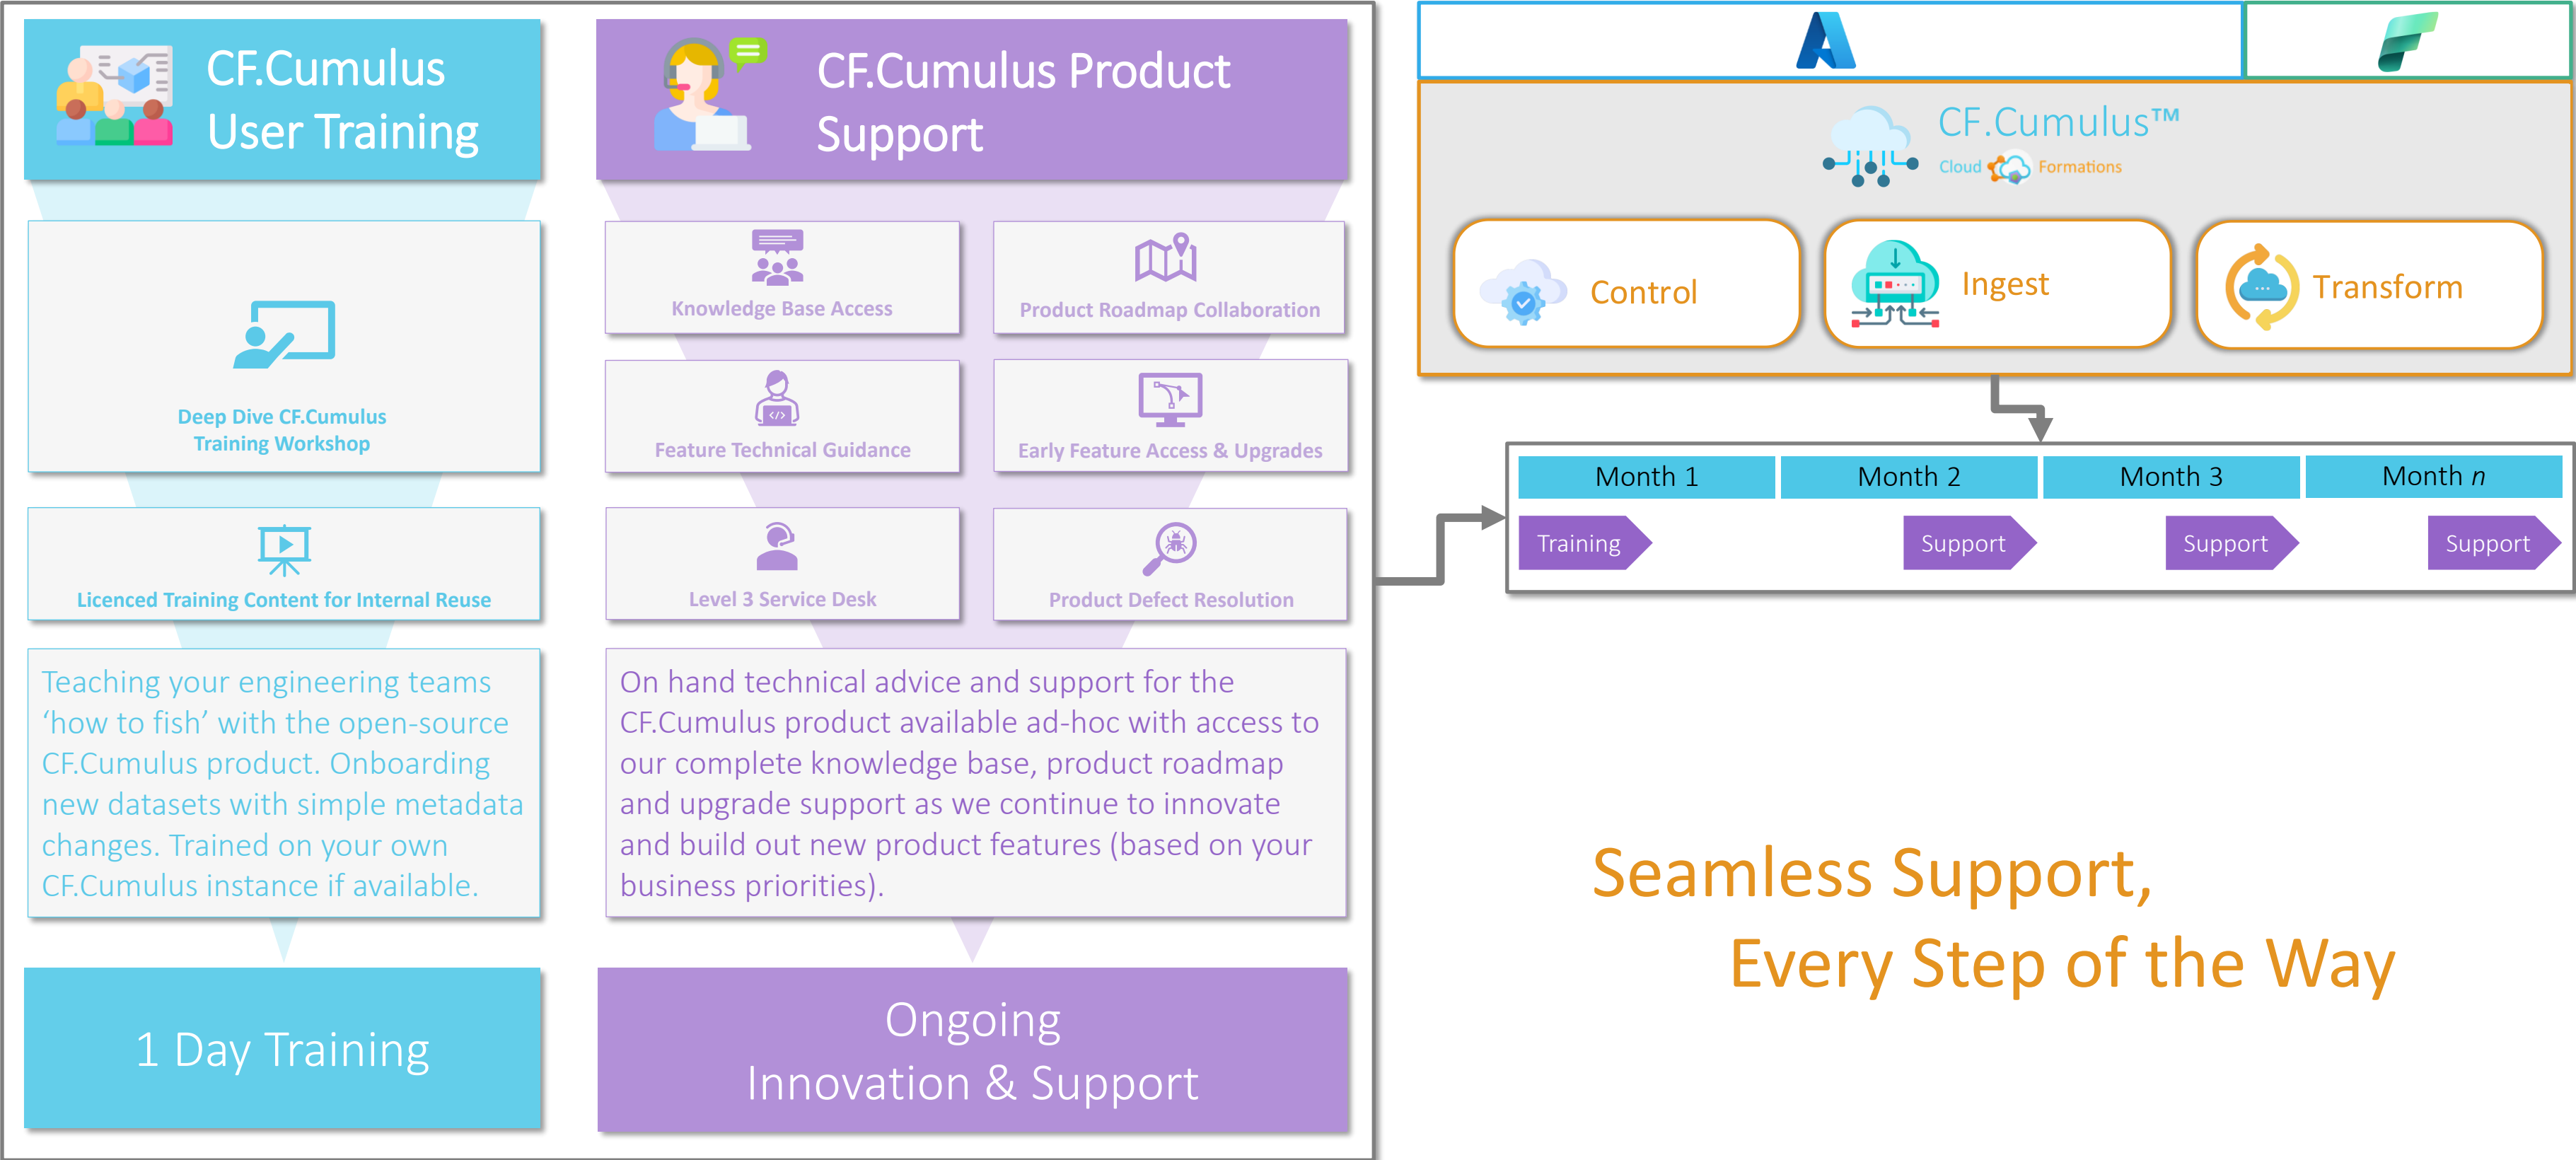
With Sensitive Data Handling



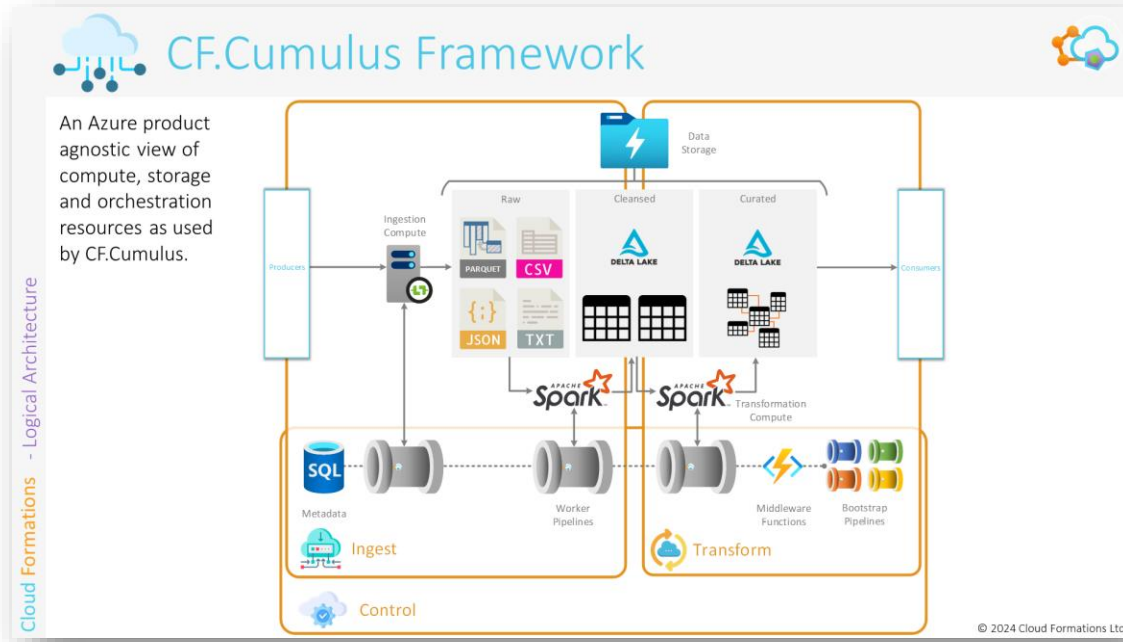
CF.Cumulus.Support



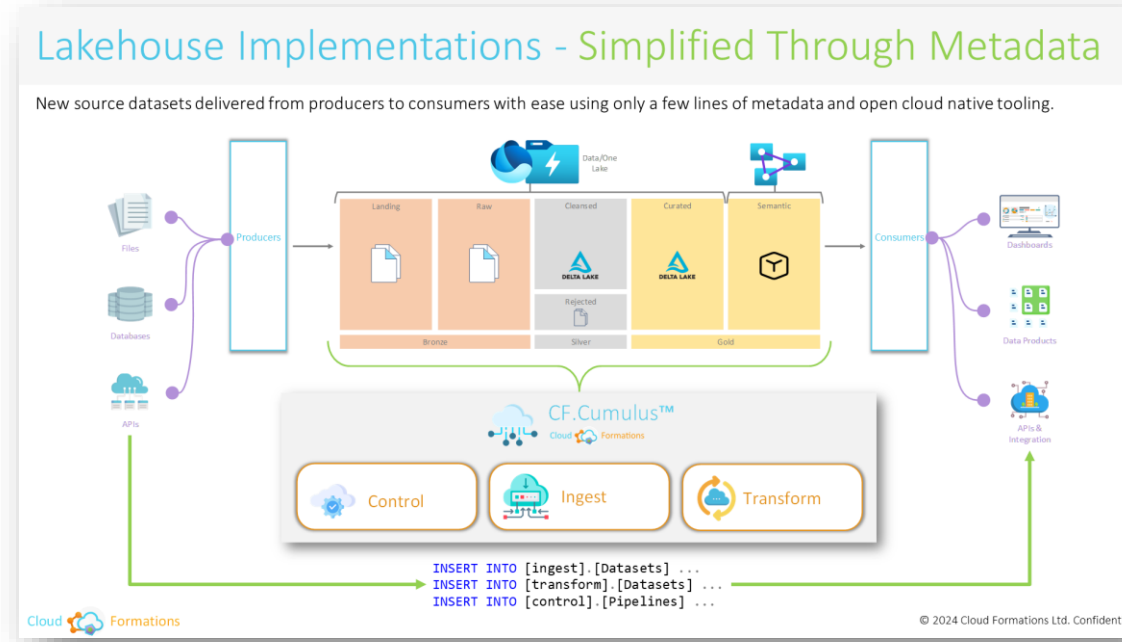
Our enduring support gives you delivery confidence when using CF.Cumulus from the product creators, with level 3 support and guidance.



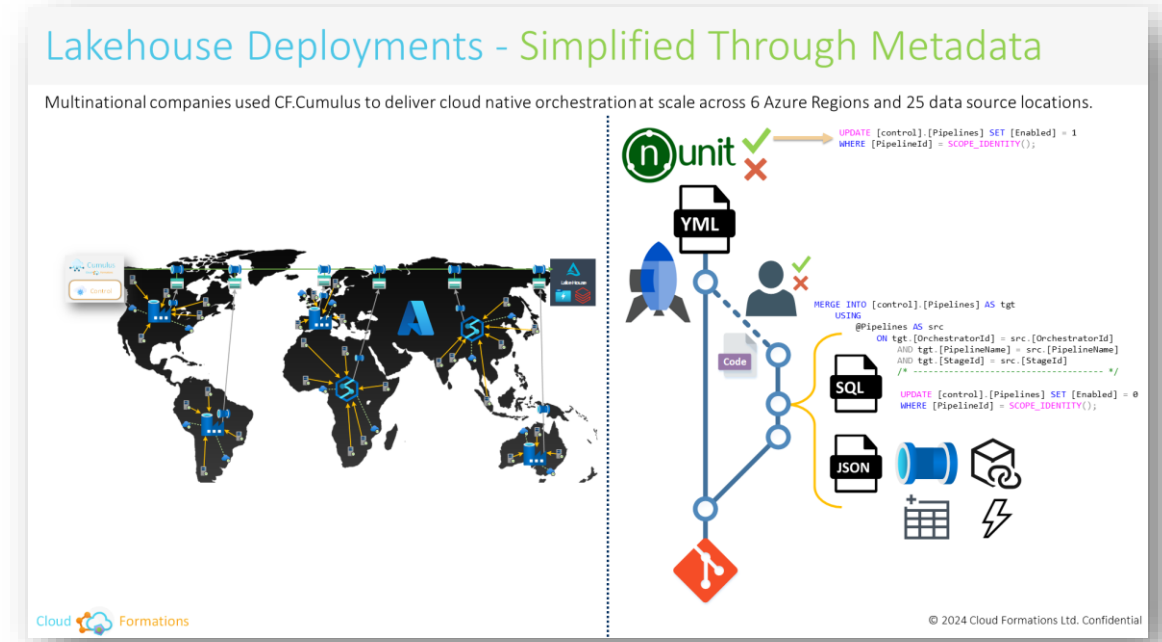
In Summary



1. Open-source and open-standards.



2. Rapid dataset delivery.



3. Deployment control.

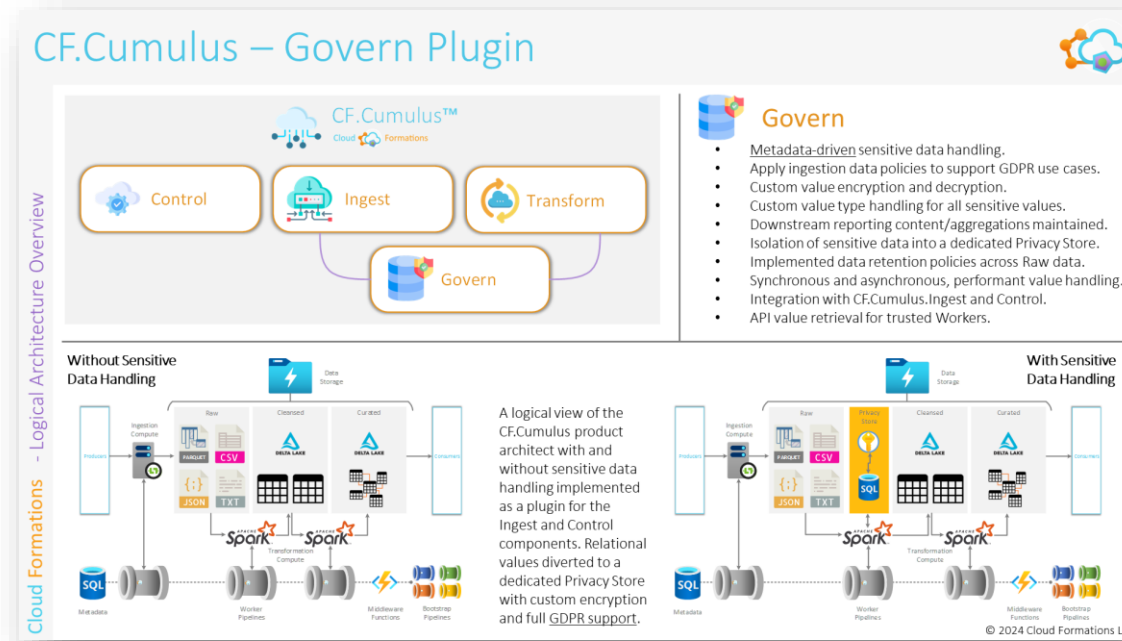
CF.Cumulus Indicative Run Costs

Assuming the use of an Azure UK Region priced for GBP, with all CF.Cumulus compute processes running for 4 hours every day of the year, with locally redundant storage, the following Azure run costs per month can be assumed:

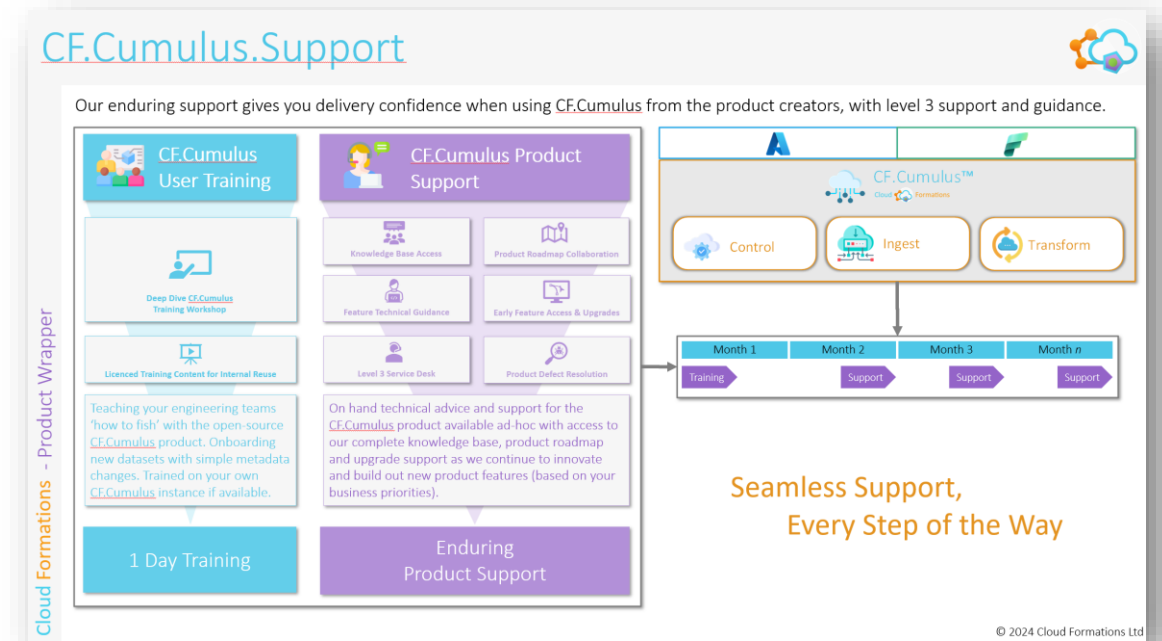
SQL Database 250GB of Metadata £12.08	Data Factory Azure IR Usage £25.84	Functions 2 Node Dedicated App Service Plan £44.30	Databricks 3 Node Cluster with 4 vCPUs & 16GB RAM £152.92
Blob Storage 1TB of Log Data £18.45	Data Lake 1TB Analytics Data £38.82	Key Vault 50,000 Requests £0.14	Monitoring 500MB Log & App Insights Data per Day £45.61
£338.16*			

*With every customer implementation data volumes, table complexity and serving queries (to dashboards) may differ. All affecting the above run costs.
Azure Price Calculator workings available for review and adjustment: <https://azure.com/e/30f8aeedfd564ba189d33f278845bb89>

4. Low running costs.




5. Plugins and extensibility.



6. Comprehensive support.

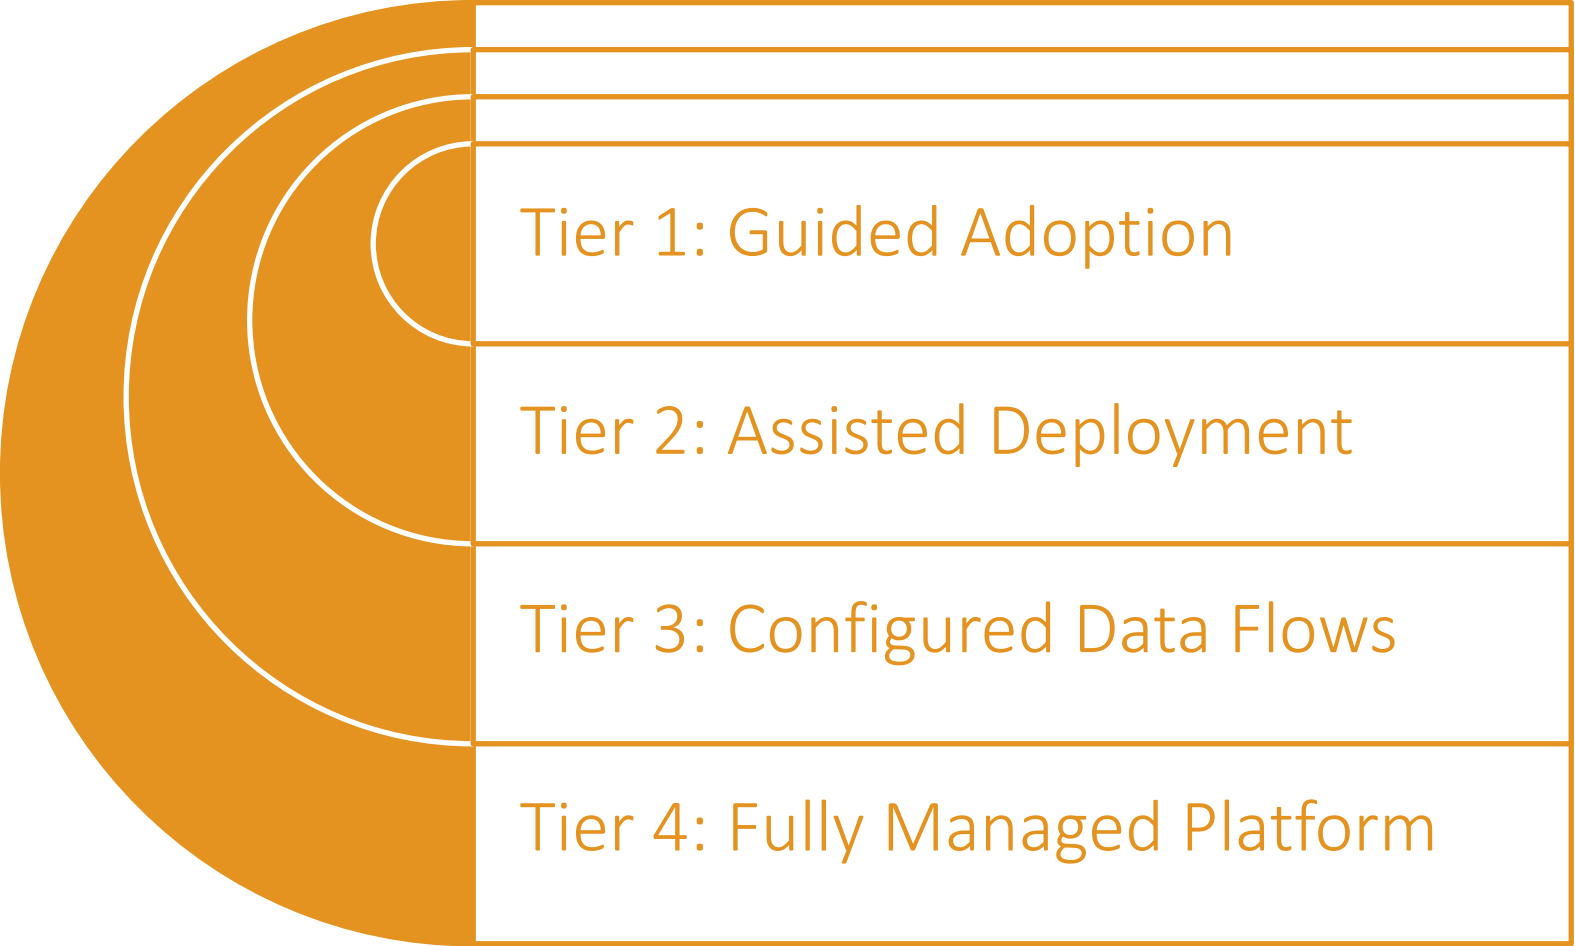
Engagement Options

Cloud Formations



CF.Cumulus: Flexible Tiers of Engagement

Unlock the full potential of the **CF.Cumulus** at your own pace. Whether you're looking to deploy and manage it yourself, or prefer a fully managed service, we have the right level of support for you.




Cumulus is now available as an open-source framework, allowing businesses to leverage cutting-edge data orchestration technology without license fees. However, the real value comes from partnering with us to ensure a successful deployment. Our decades of experience in building scalable, resilient data platforms will ensure that you extract the maximum value from your data.

Open-Source Flexibility, Professional Expertise

CF.Cumulus: Flexible Tiers of Engagement


Unlock the full potential of the **CF.Cumulus** at your own pace. Whether you're looking to deploy and manage it yourself, or prefer a fully managed service, we have the right level of support for you.



Open-Source Flexibility, Professional Expertise

CF.Cumulus: Flexible Tiers of Engagement

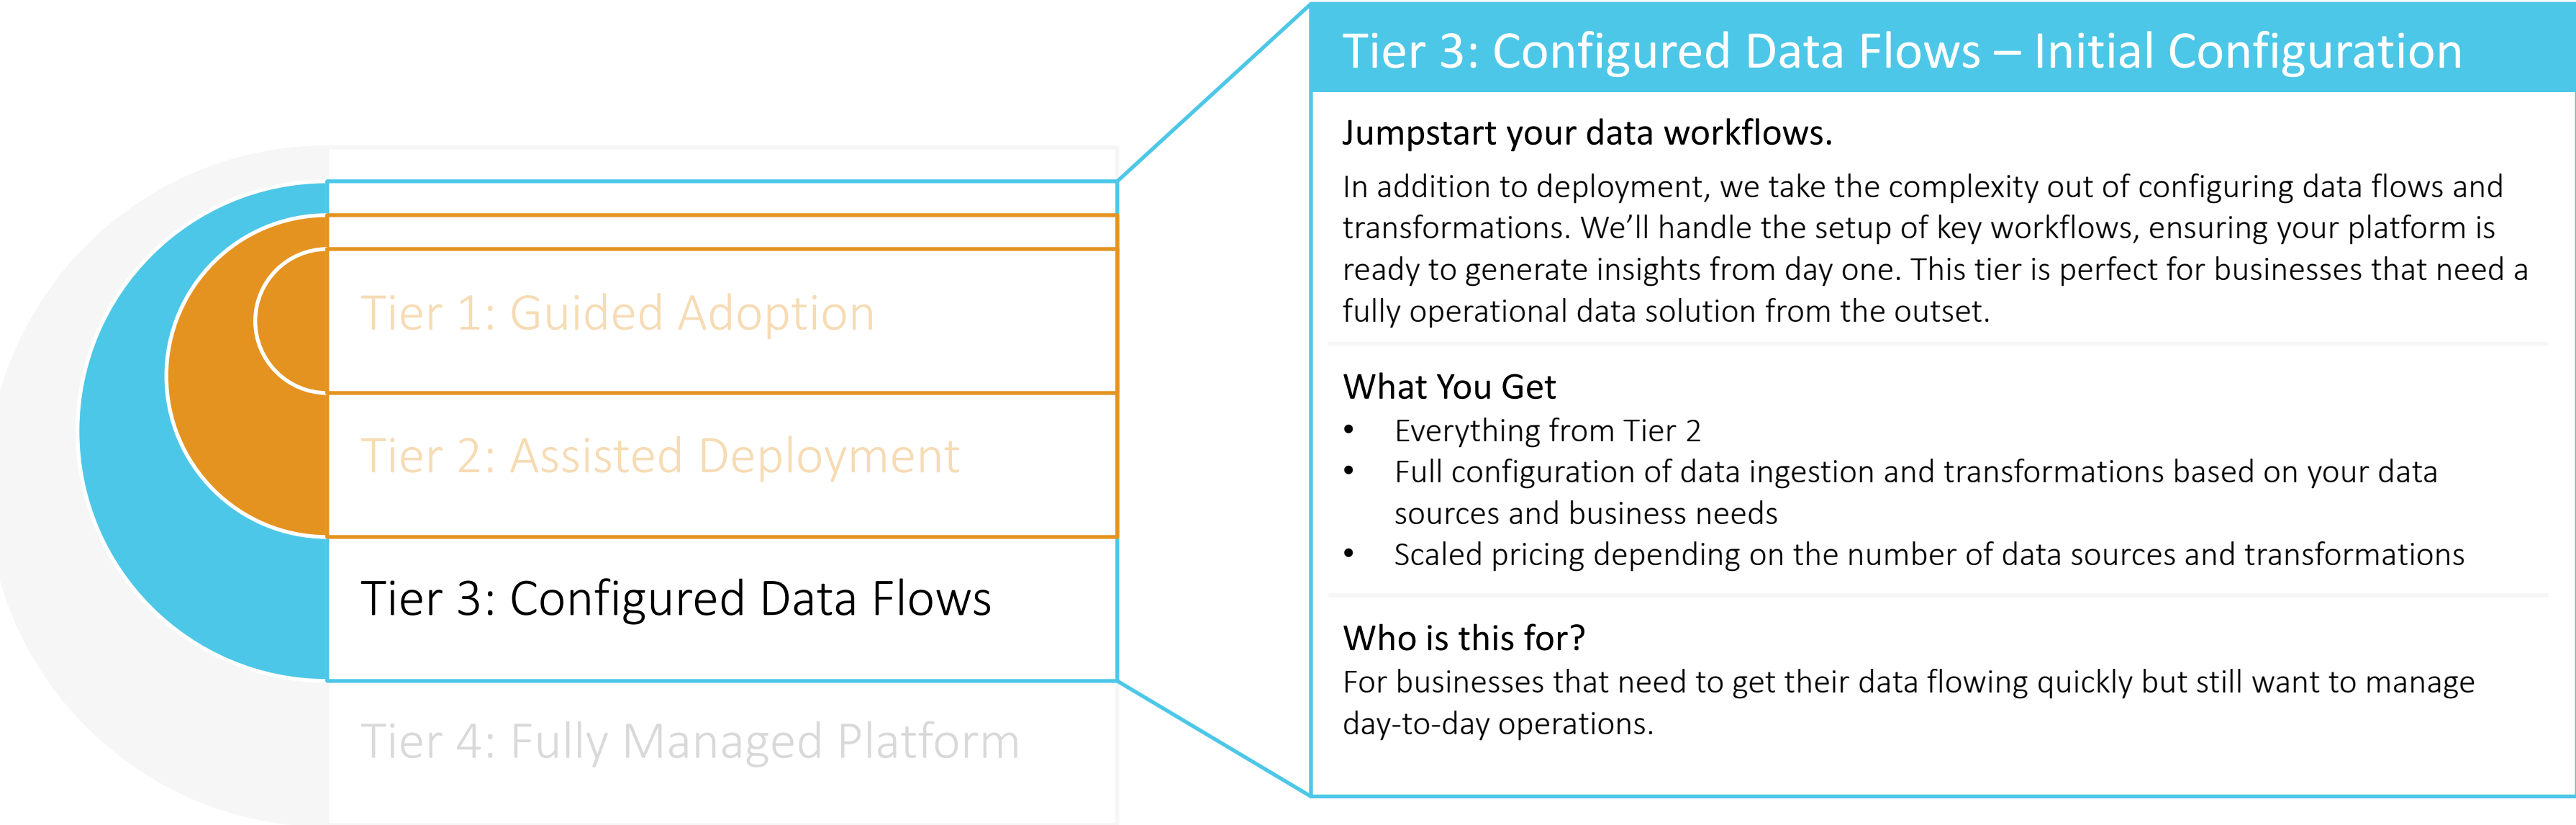
Unlock the full potential of the **CF.Cumulus** at your own pace. Whether you're looking to deploy and manage it yourself, or prefer a fully managed service, we have the right level of support for you.



Open-Source Flexibility, Professional Expertise

CF.Cumulus: Flexible Tiers of Engagement

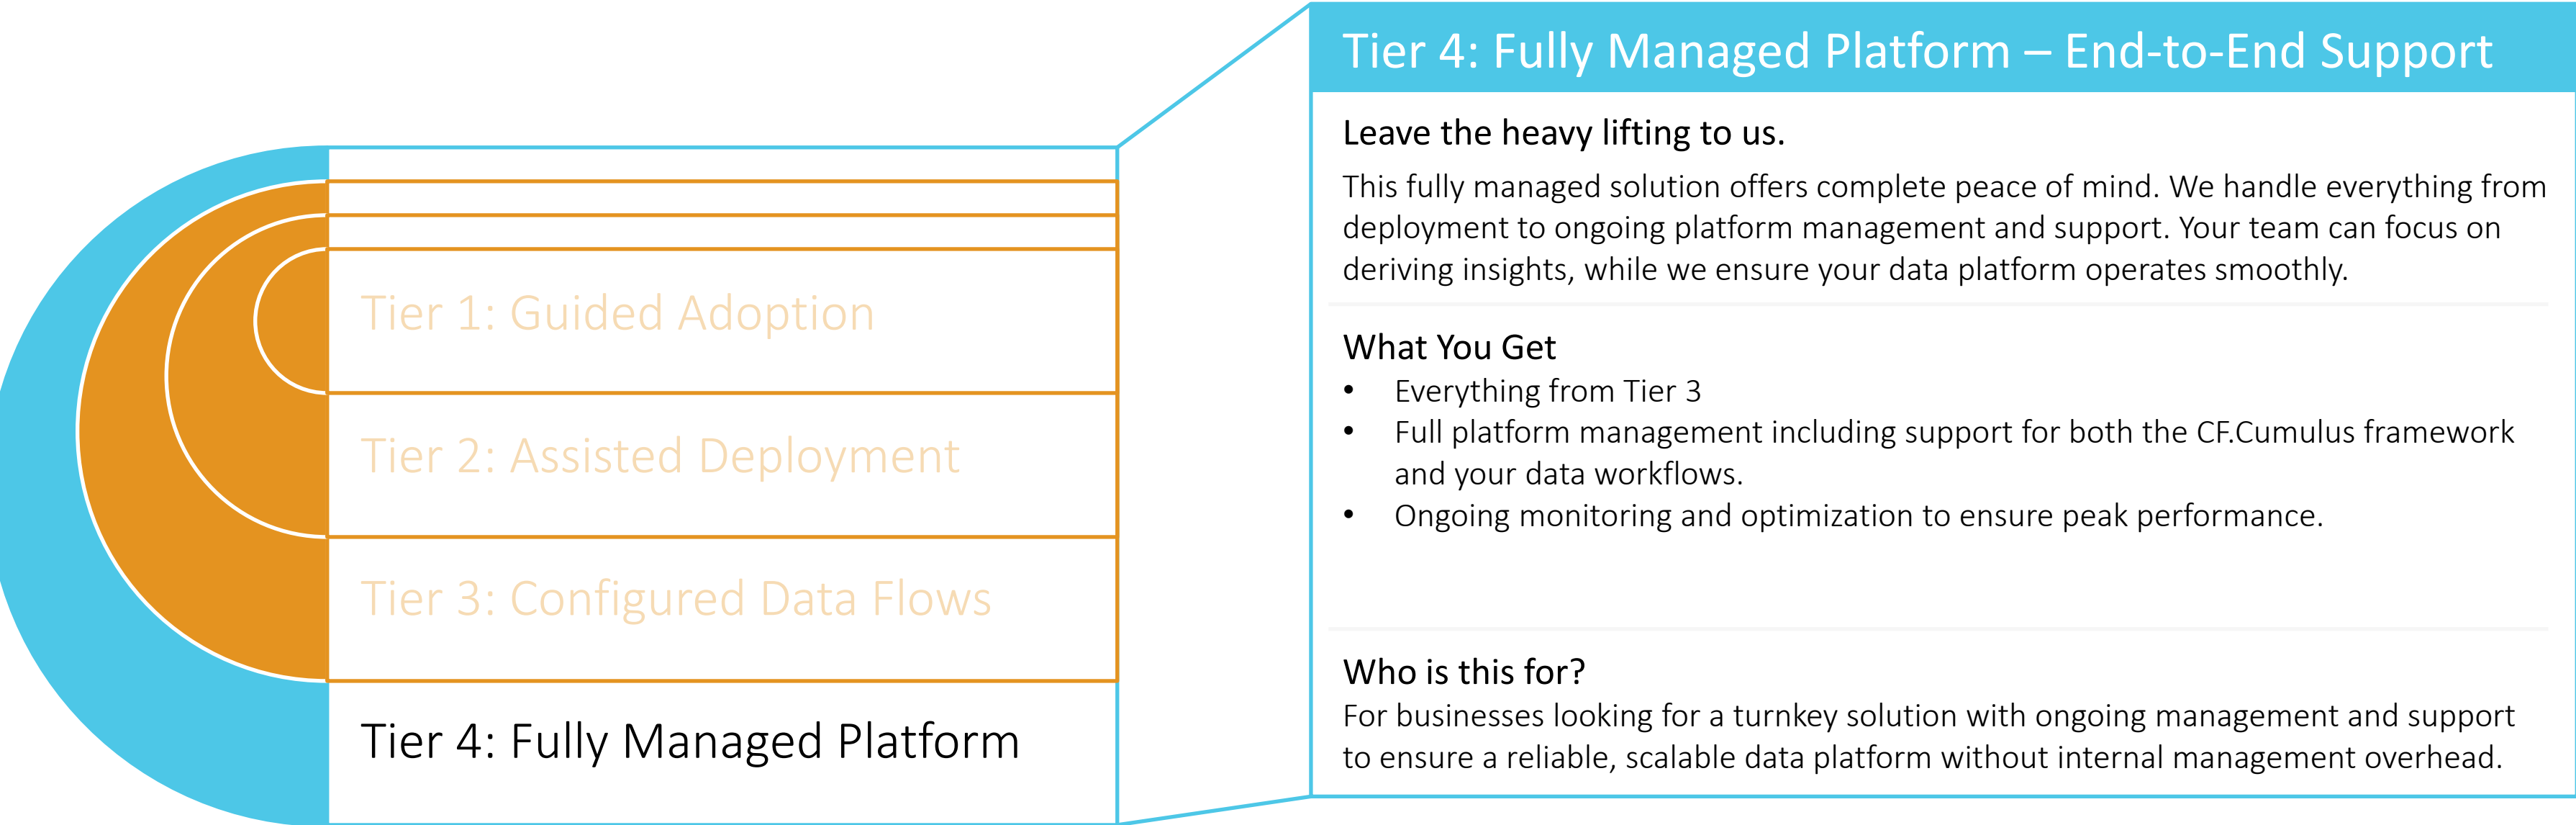
Unlock the full potential of the **CF.Cumulus** at your own pace. Whether you're looking to deploy and manage it yourself, or prefer a fully managed service, we have the right level of support for you.



Open-Source Flexibility, Professional Expertise

CF.Cumulus: Flexible Tiers of Engagement

Unlock the full potential of the **CF.Cumulus** at your own pace. Whether you're looking to deploy and manage it yourself, or prefer a fully managed service, we have the right level of support for you.



Open-Source Flexibility, Professional Expertise

CF.Cumulus: Flexible Tiers of Engagement

Tier 1: Guided Adoption – Training & Product Support

Take full control with expert guidance

We equip your team with the knowledge and tools to confidently deploy, configure, and use the framework. After training, our team remains available to support you through your initial implementation phase, ensuring a smooth transition to self-sufficiency.

What You Get

- Hands-on training covering deployment, configuration, and usage
- On-demand support during early implementation stages
- Easy transition into ongoing product support with continued guidance and access to our L3/L4 service desk for Cumulus-specific issues

Who is this for?

For teams that want to take ownership of the Cumulus platform but need expert training and a support safety net during the initial stages.

Tier 2: Assisted Deployment – Design & Deployment

Handle complex environments with confidence.

For organizations with more intricate security, networking, or infrastructure requirements, we go beyond training by helping you design and deploy the Cumulus framework in your environment. You'll get the best of both worlds—our expertise to handle complexity, plus the flexibility to manage the platform post-deployment.

What You Get

- Everything from Tier 1
- Custom architecture design to accommodate security and networking needs
- Initial deployment in your environment with hands-on assistance

Who is this for?

For businesses with complex infrastructure that want expert help with deployment but still wish to manage the platform internally after setup.

Tier 3: Configured Data Flows – Initial Configuration

Jumpstart your data workflows.

In addition to deployment, we take the complexity out of configuring data flows and transformations. We'll handle the setup of key workflows, ensuring your platform is ready to generate insights from day one. This tier is perfect for businesses that need a fully operational data solution from the outset.

What You Get

- Everything from Tier 2
- Full configuration of data ingestion and transformations based on your data sources and business needs
- Scaled pricing depending on the number of data sources and transformations

Who is this for?

For businesses that need to get their data flowing quickly but still want to manage day-to-day operations.

Tier 4: Fully Managed Platform – End-to-End Support

Leave the heavy lifting to us.

This fully managed solution offers complete peace of mind. We handle everything from deployment to ongoing platform management and support. Your team can focus on deriving insights, while we ensure your data platform operates smoothly.

What You Get

- Everything from Tier 3
- Full platform management including support for both the CF.Cumulus framework and your data workflows.
- Ongoing monitoring and optimization to ensure peak performance.


Who is this for?

For businesses looking for a turnkey solution with ongoing management and support to ensure a reliable, scalable data platform without internal management overhead.

Open-Source Flexibility, Professional Expertise

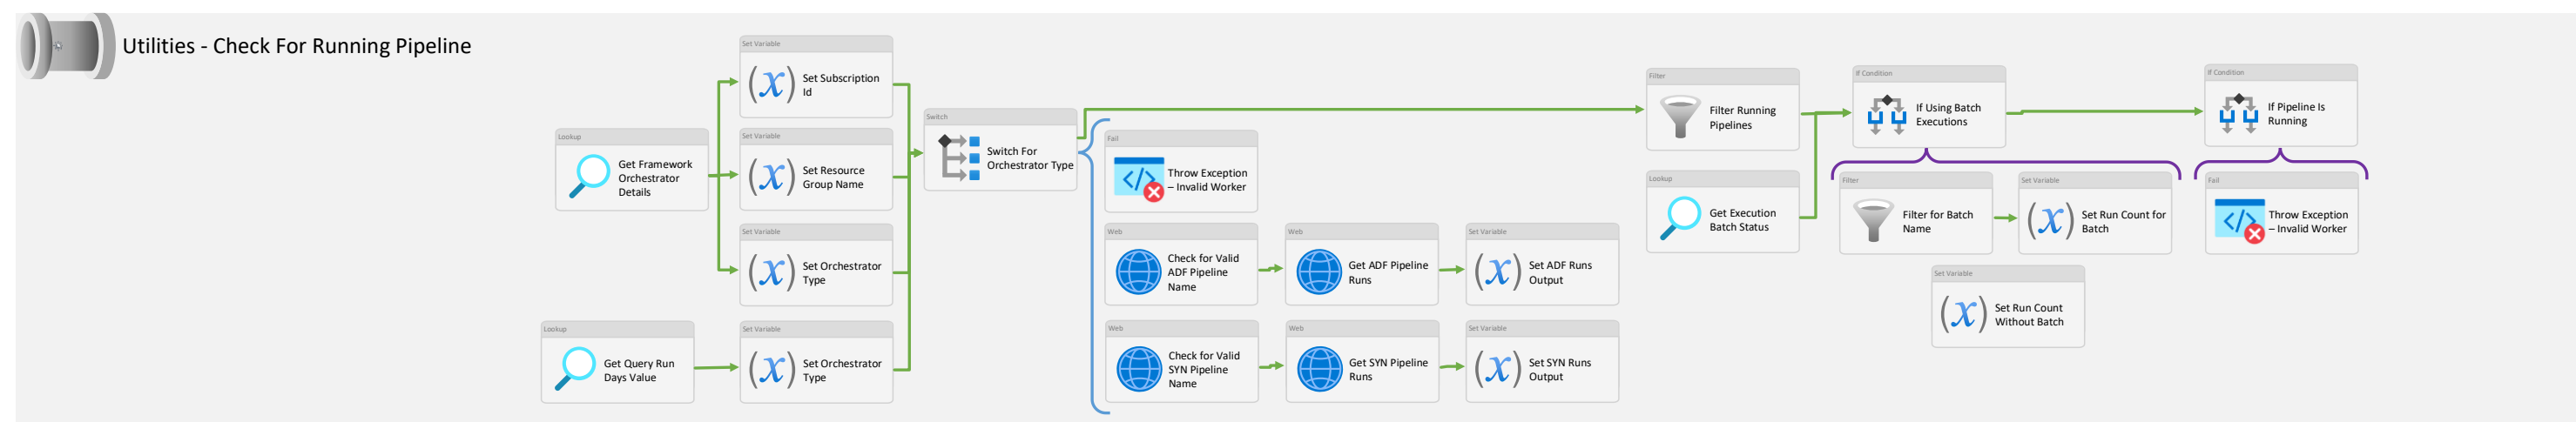
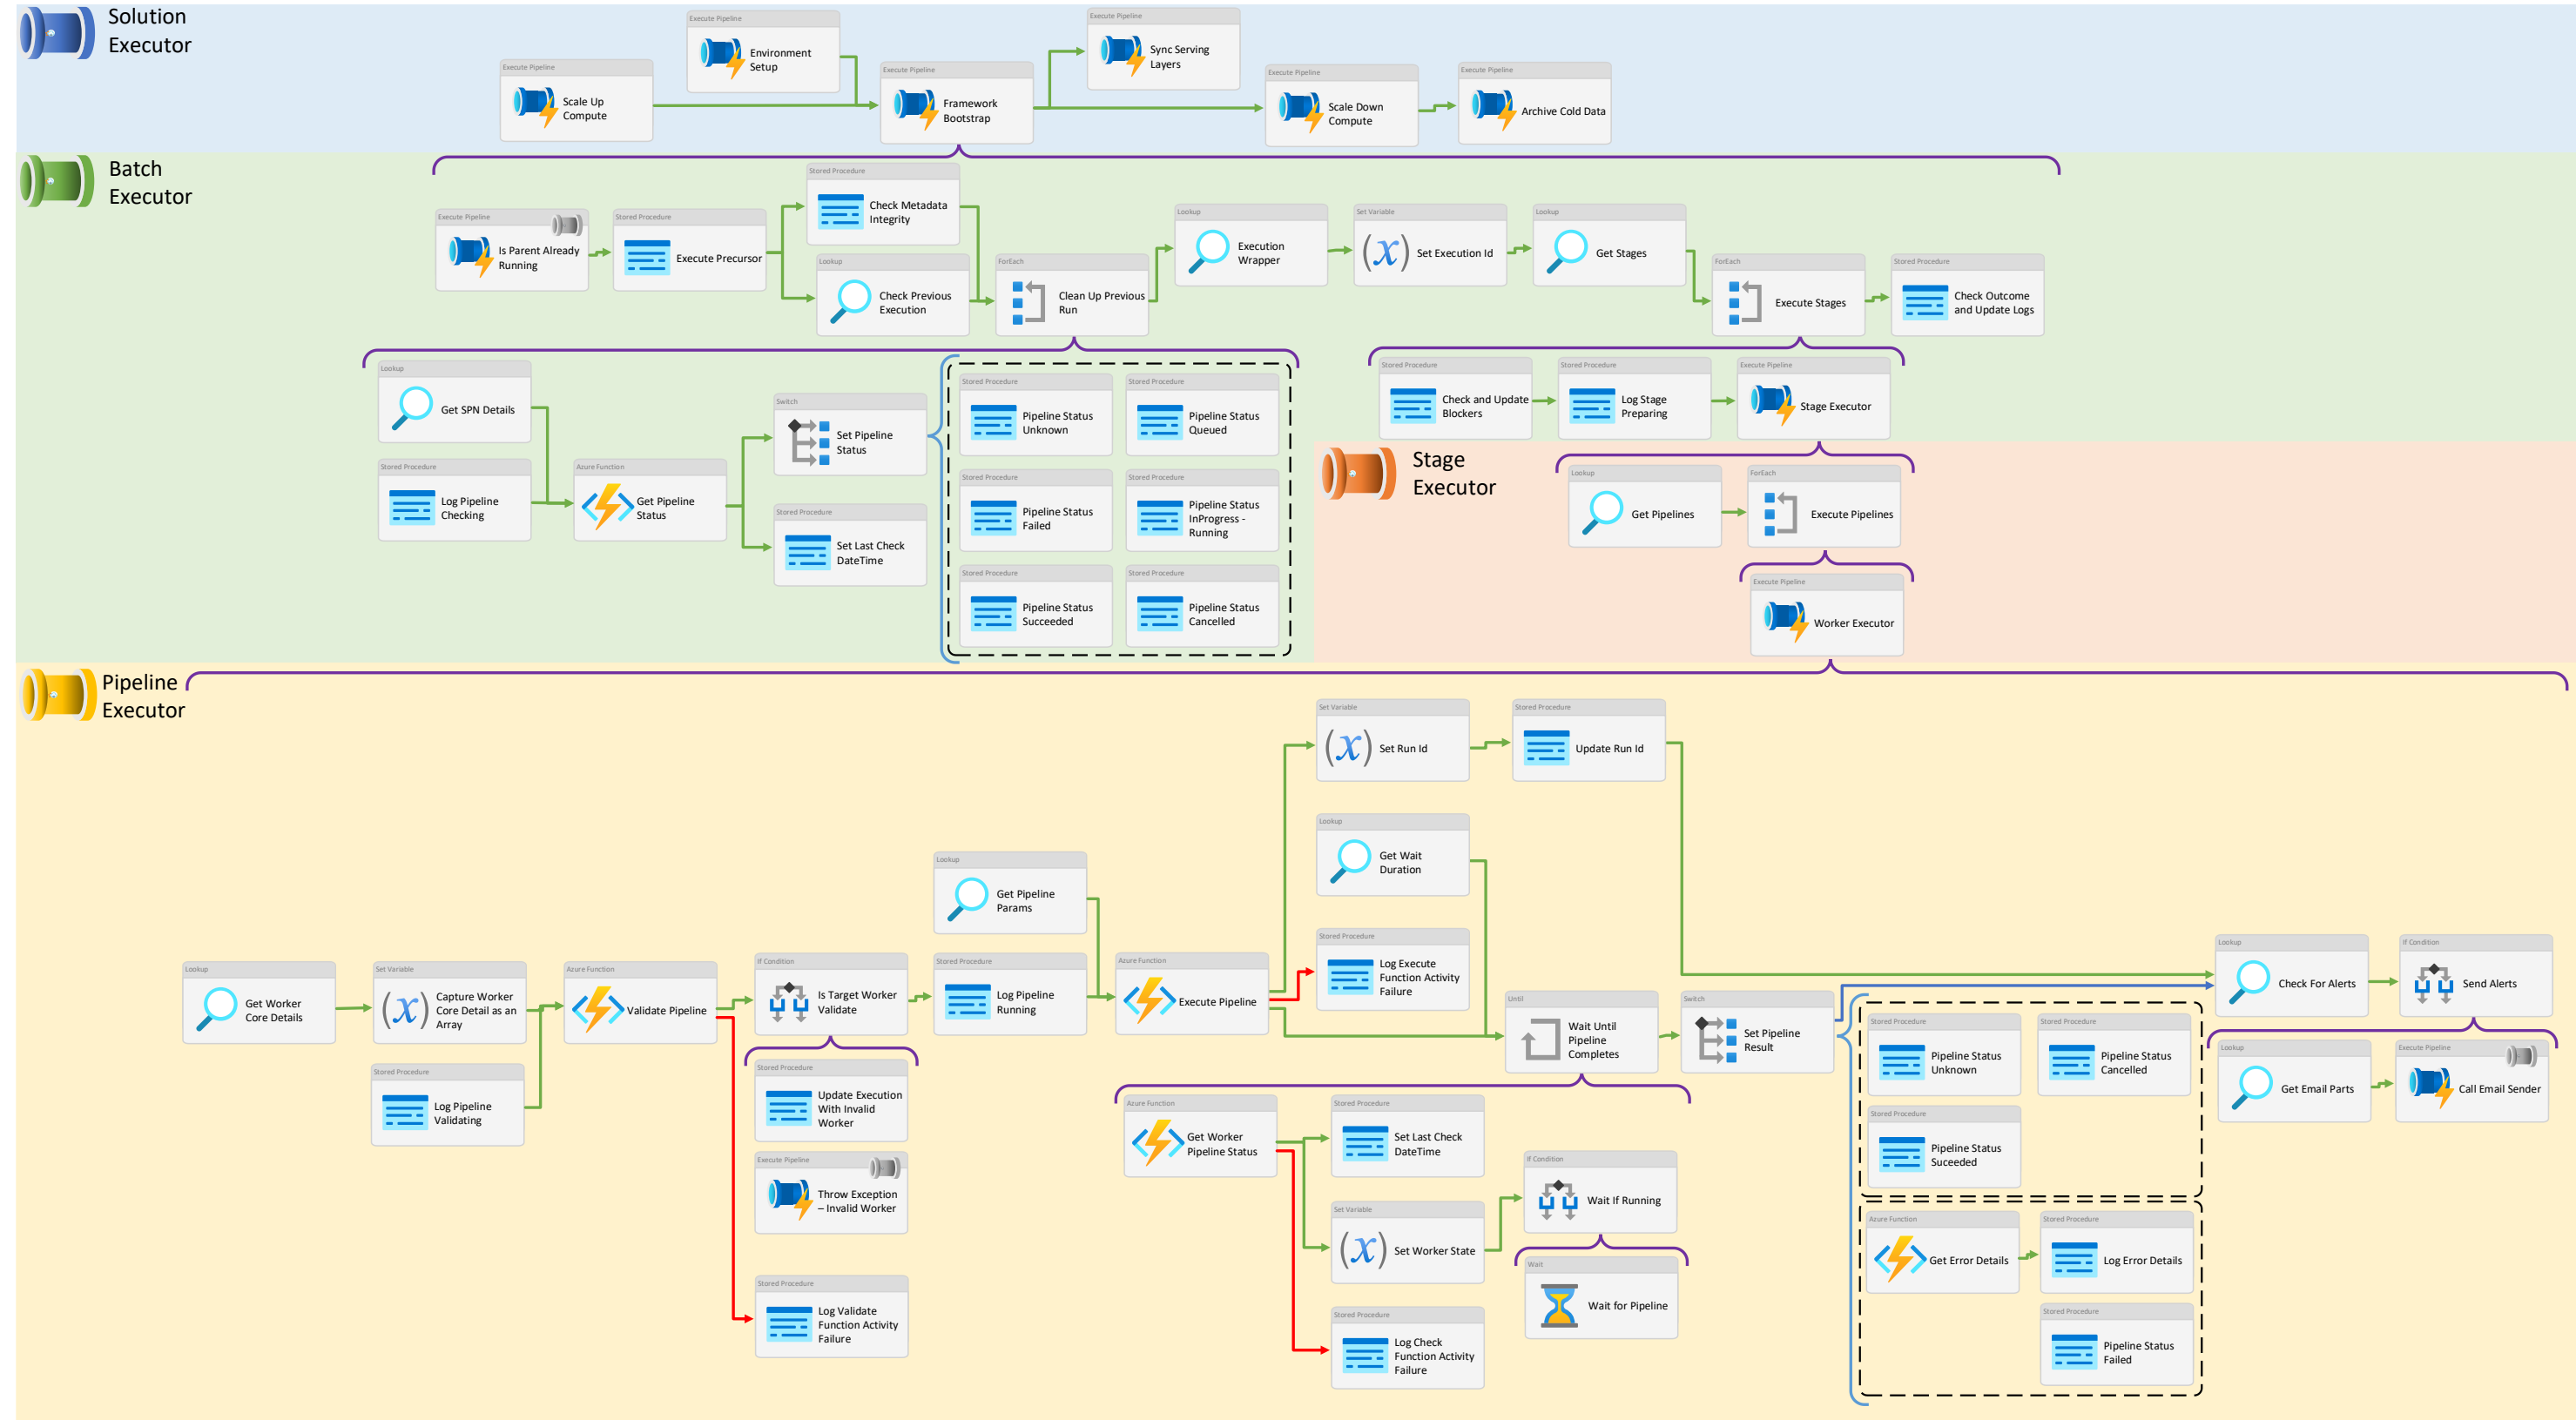
Product Internals

Cloud Formations

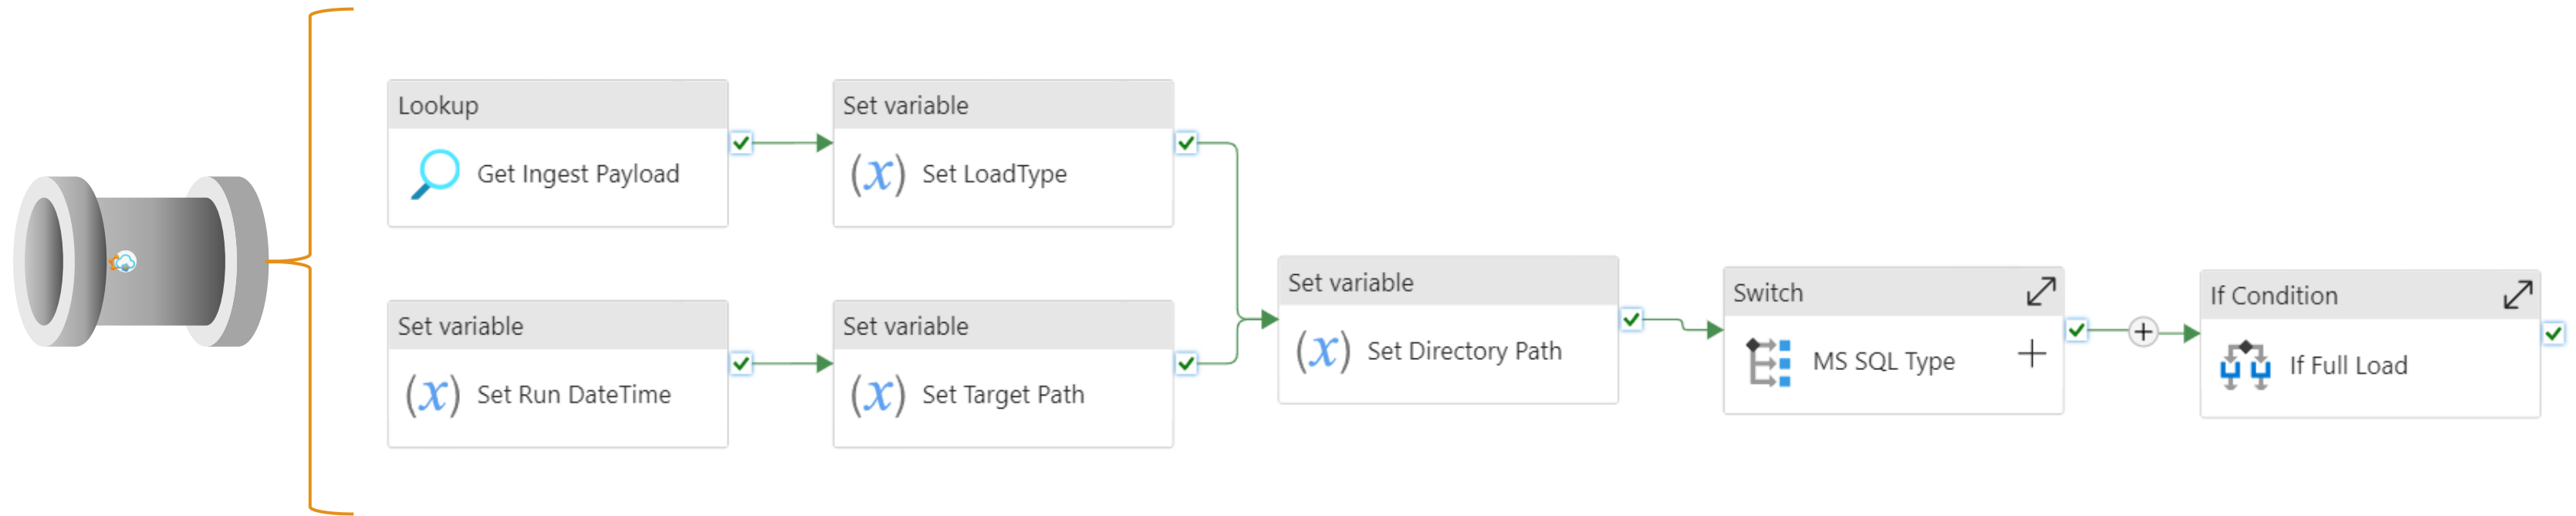


CF.Control Pipeline Internals

A view on the activity structure within the CF.Cumulus control pipelines, executing each of the Ingest and Transform workers.

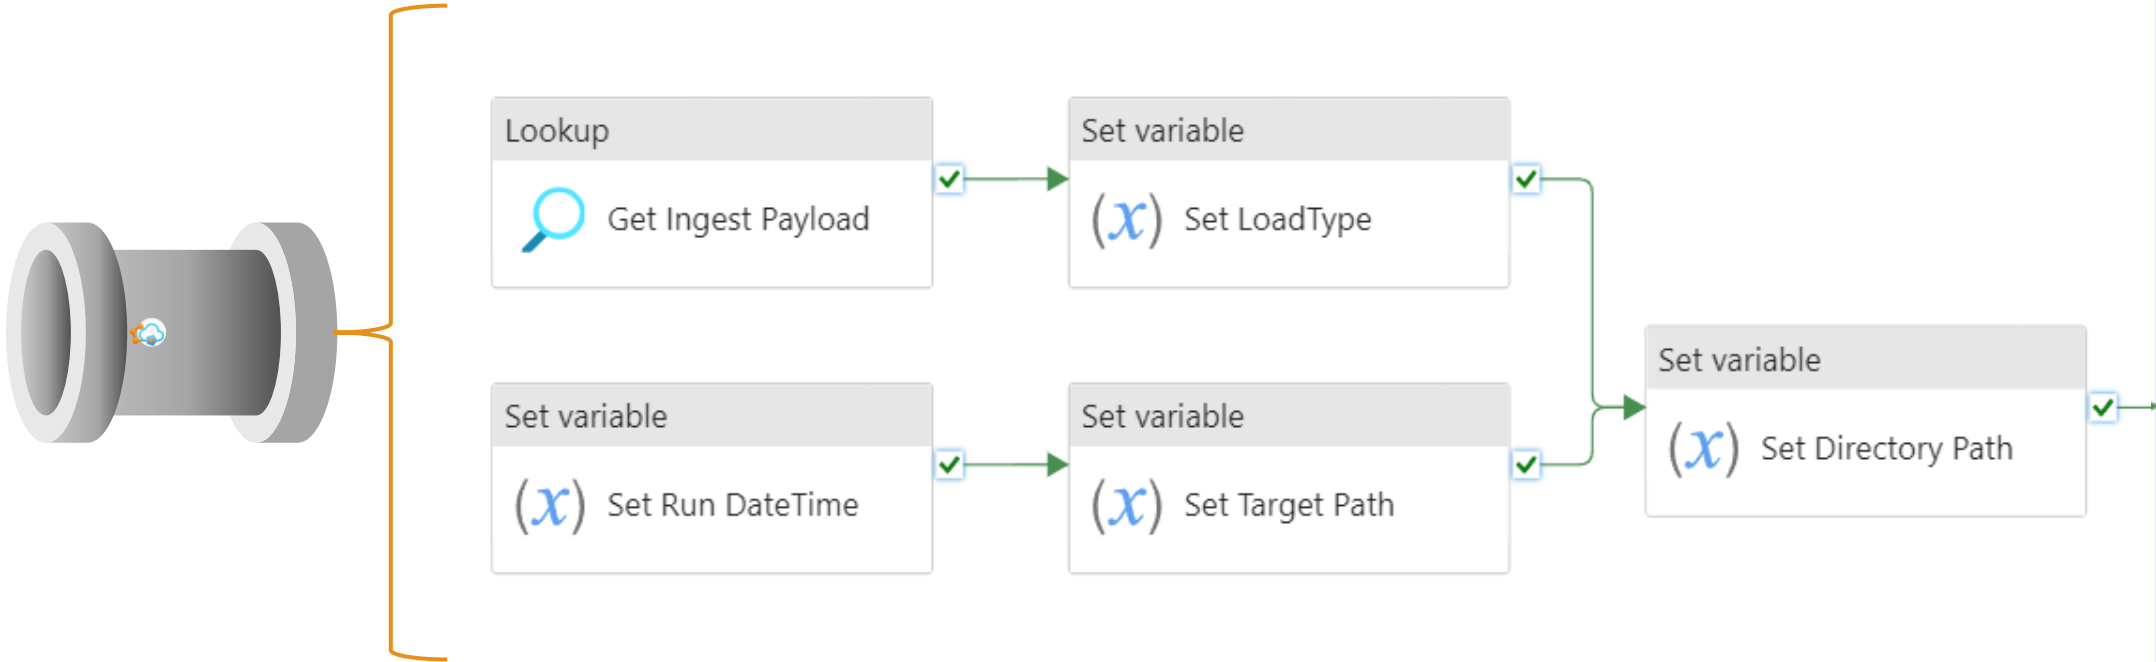


CF.Ingest Pipeline Internals



A view on the activity structure within a typical CF.Cumulus ingest pipeline, handling the flow of data from a Microsoft SQL based data source.

CF.Ingest Pipeline Internals

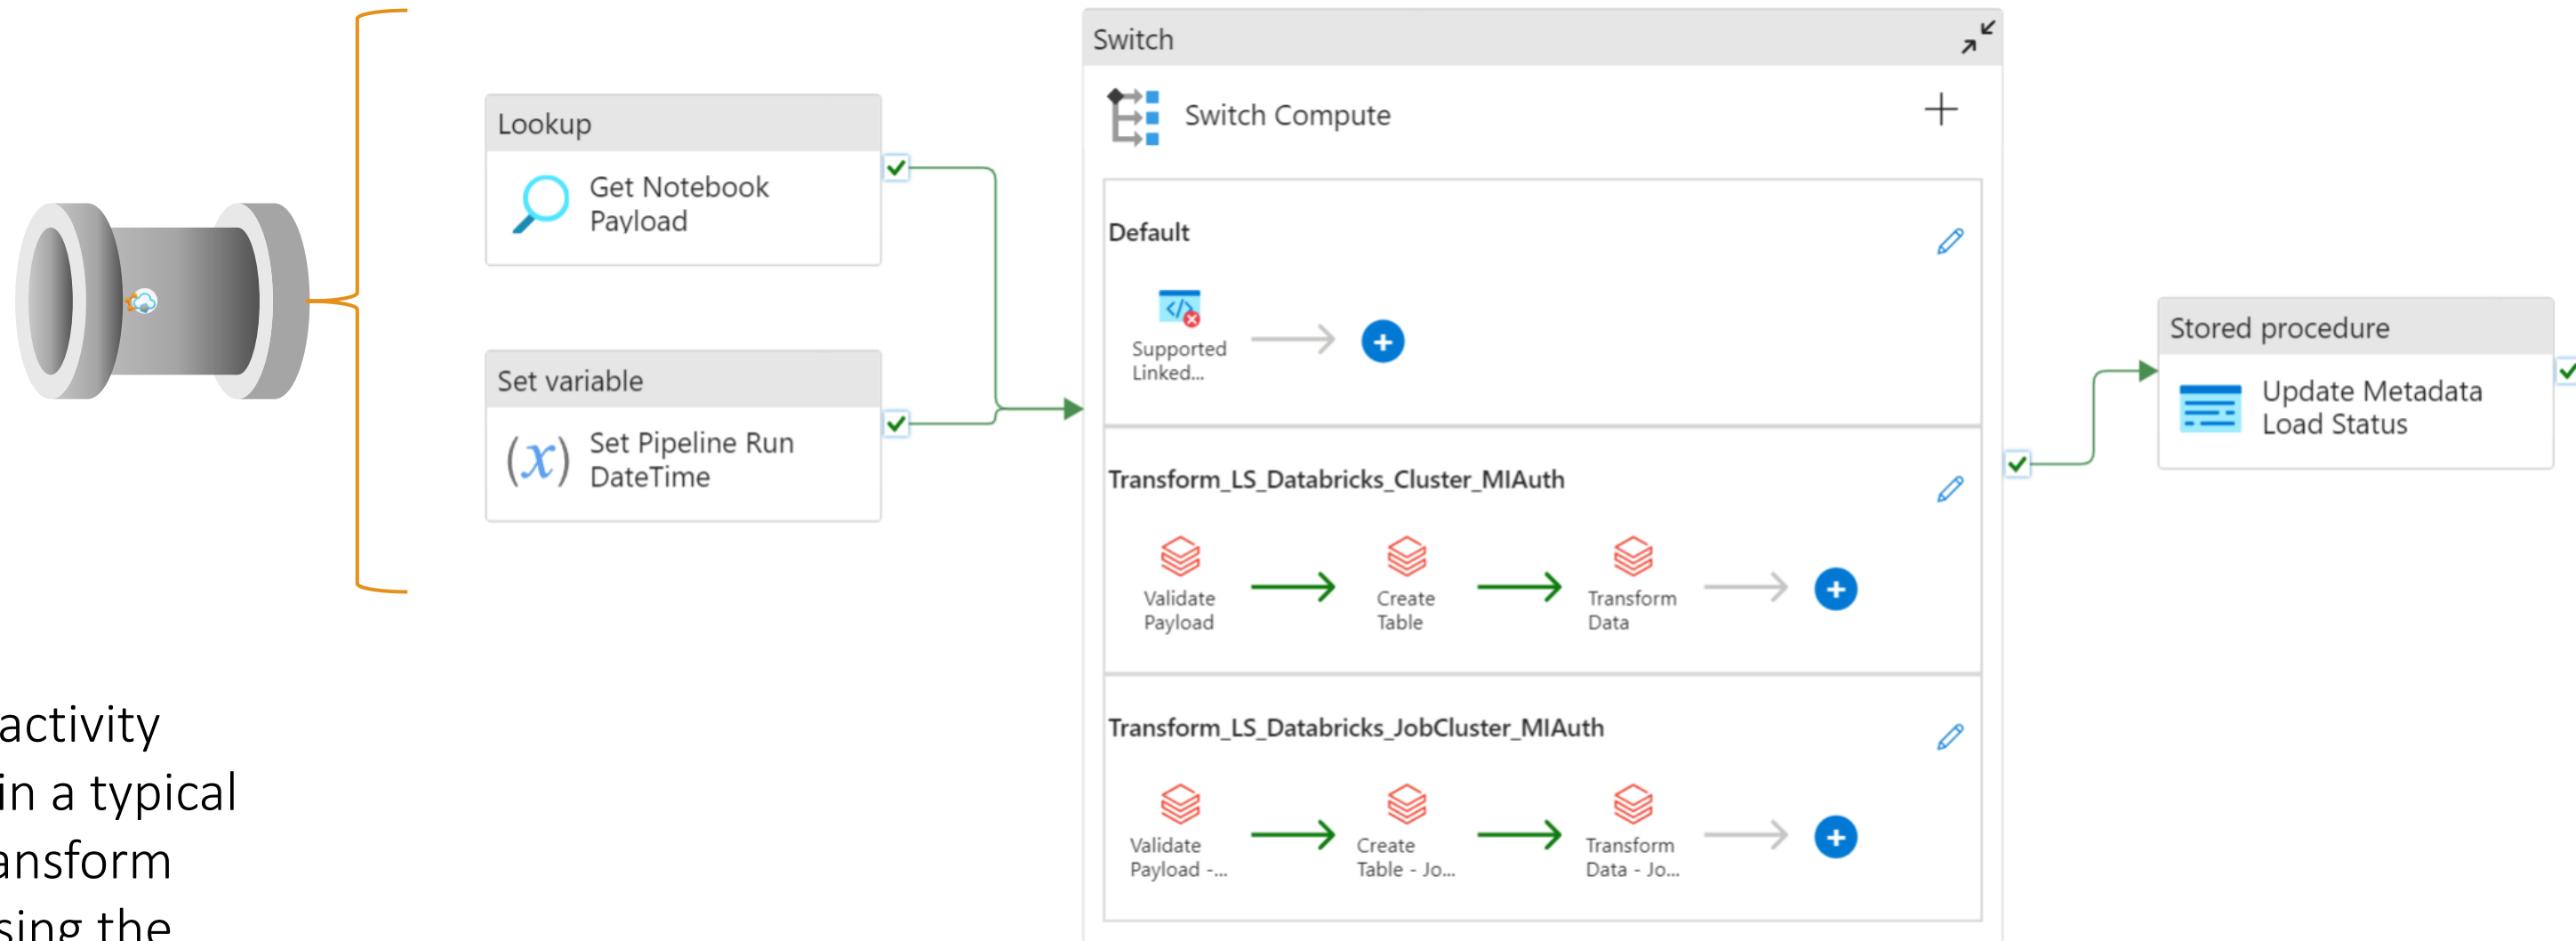


A view on the activity structure within a typical CF.Cumulus ingest pipeline, handling the flow of data from a Microsoft SQL based data source. Expanded to show the variations with the pipeline Switch activity for different version and connections of the SQL instance/resource.

The screenshot shows a 'Switch' activity configuration with five cases:

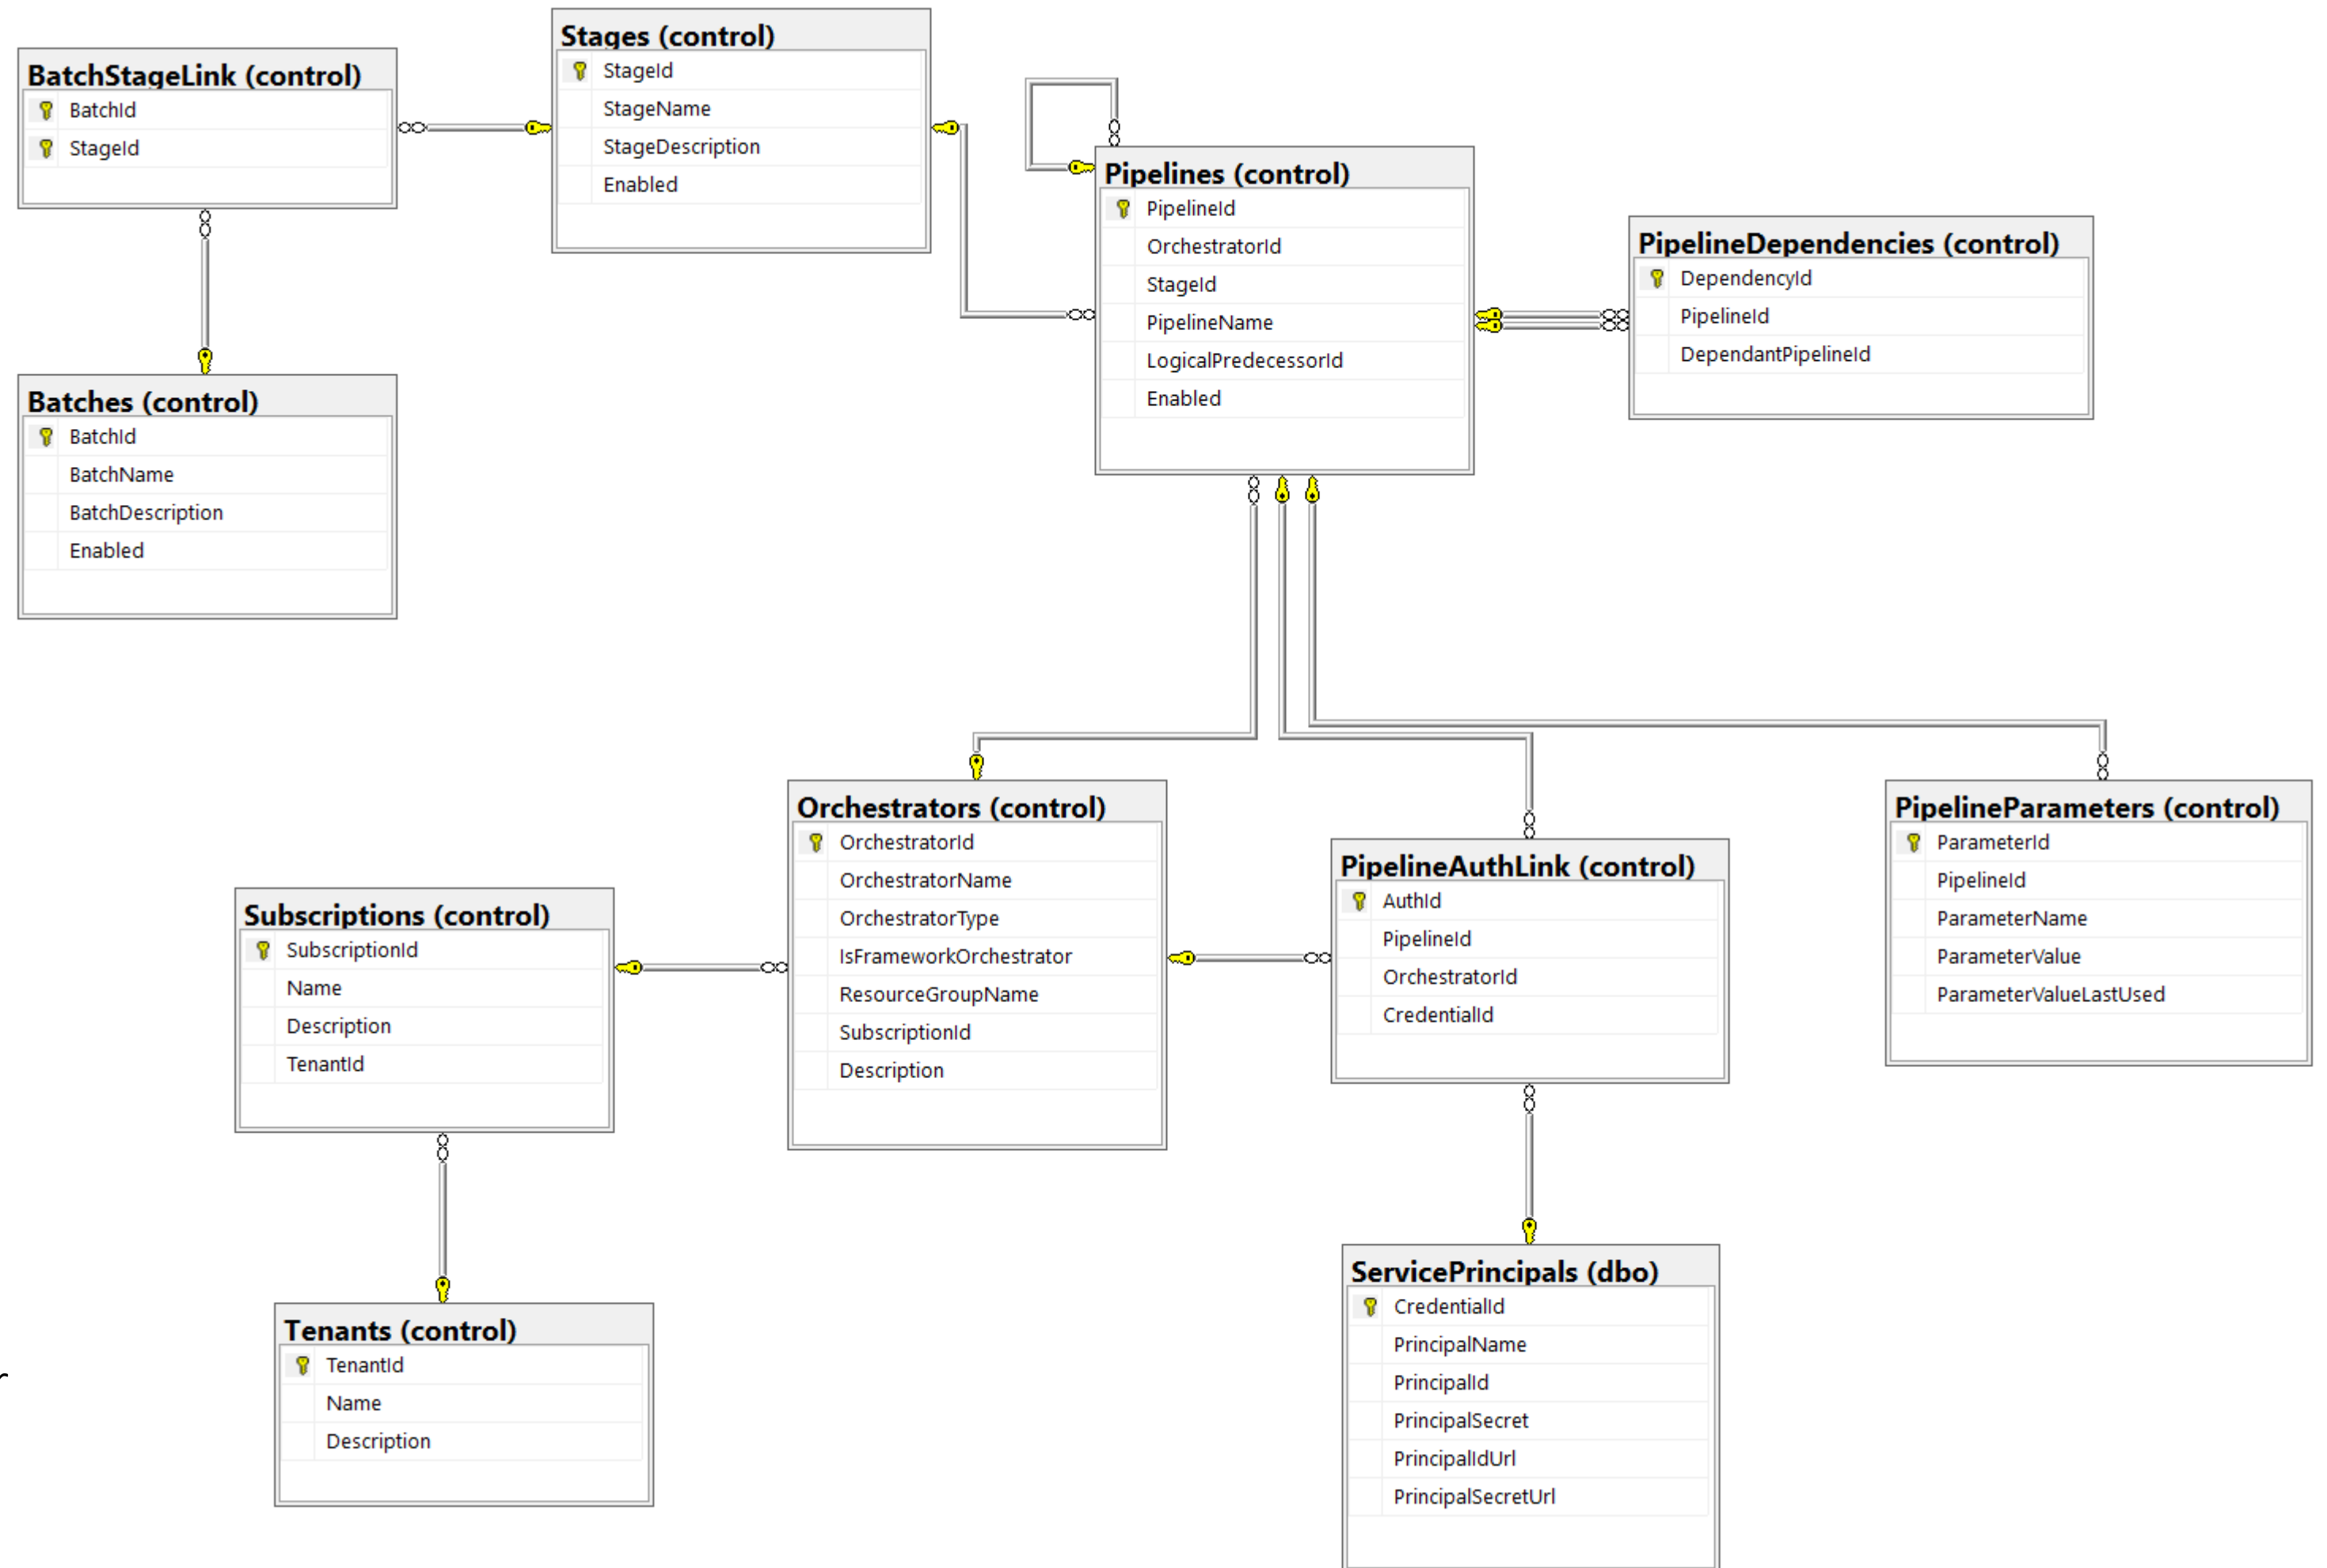
- Default**: Supported Linked... → +
- SQLDB_MIAuth**: SQLDB MIAuth... → +
- SQLDB_SQLAuth**: Fetch SQL Auth... → SQLDB SQLAuth... → +
- SQLServer_WinAuth**: Fetch Windows... → SQLServer WinAuth... → +
- SQLServer_SQLAuth**: Fetch SQL Login... → SQLServer SQLAuth... → +

CF.Transform Pipeline Internals



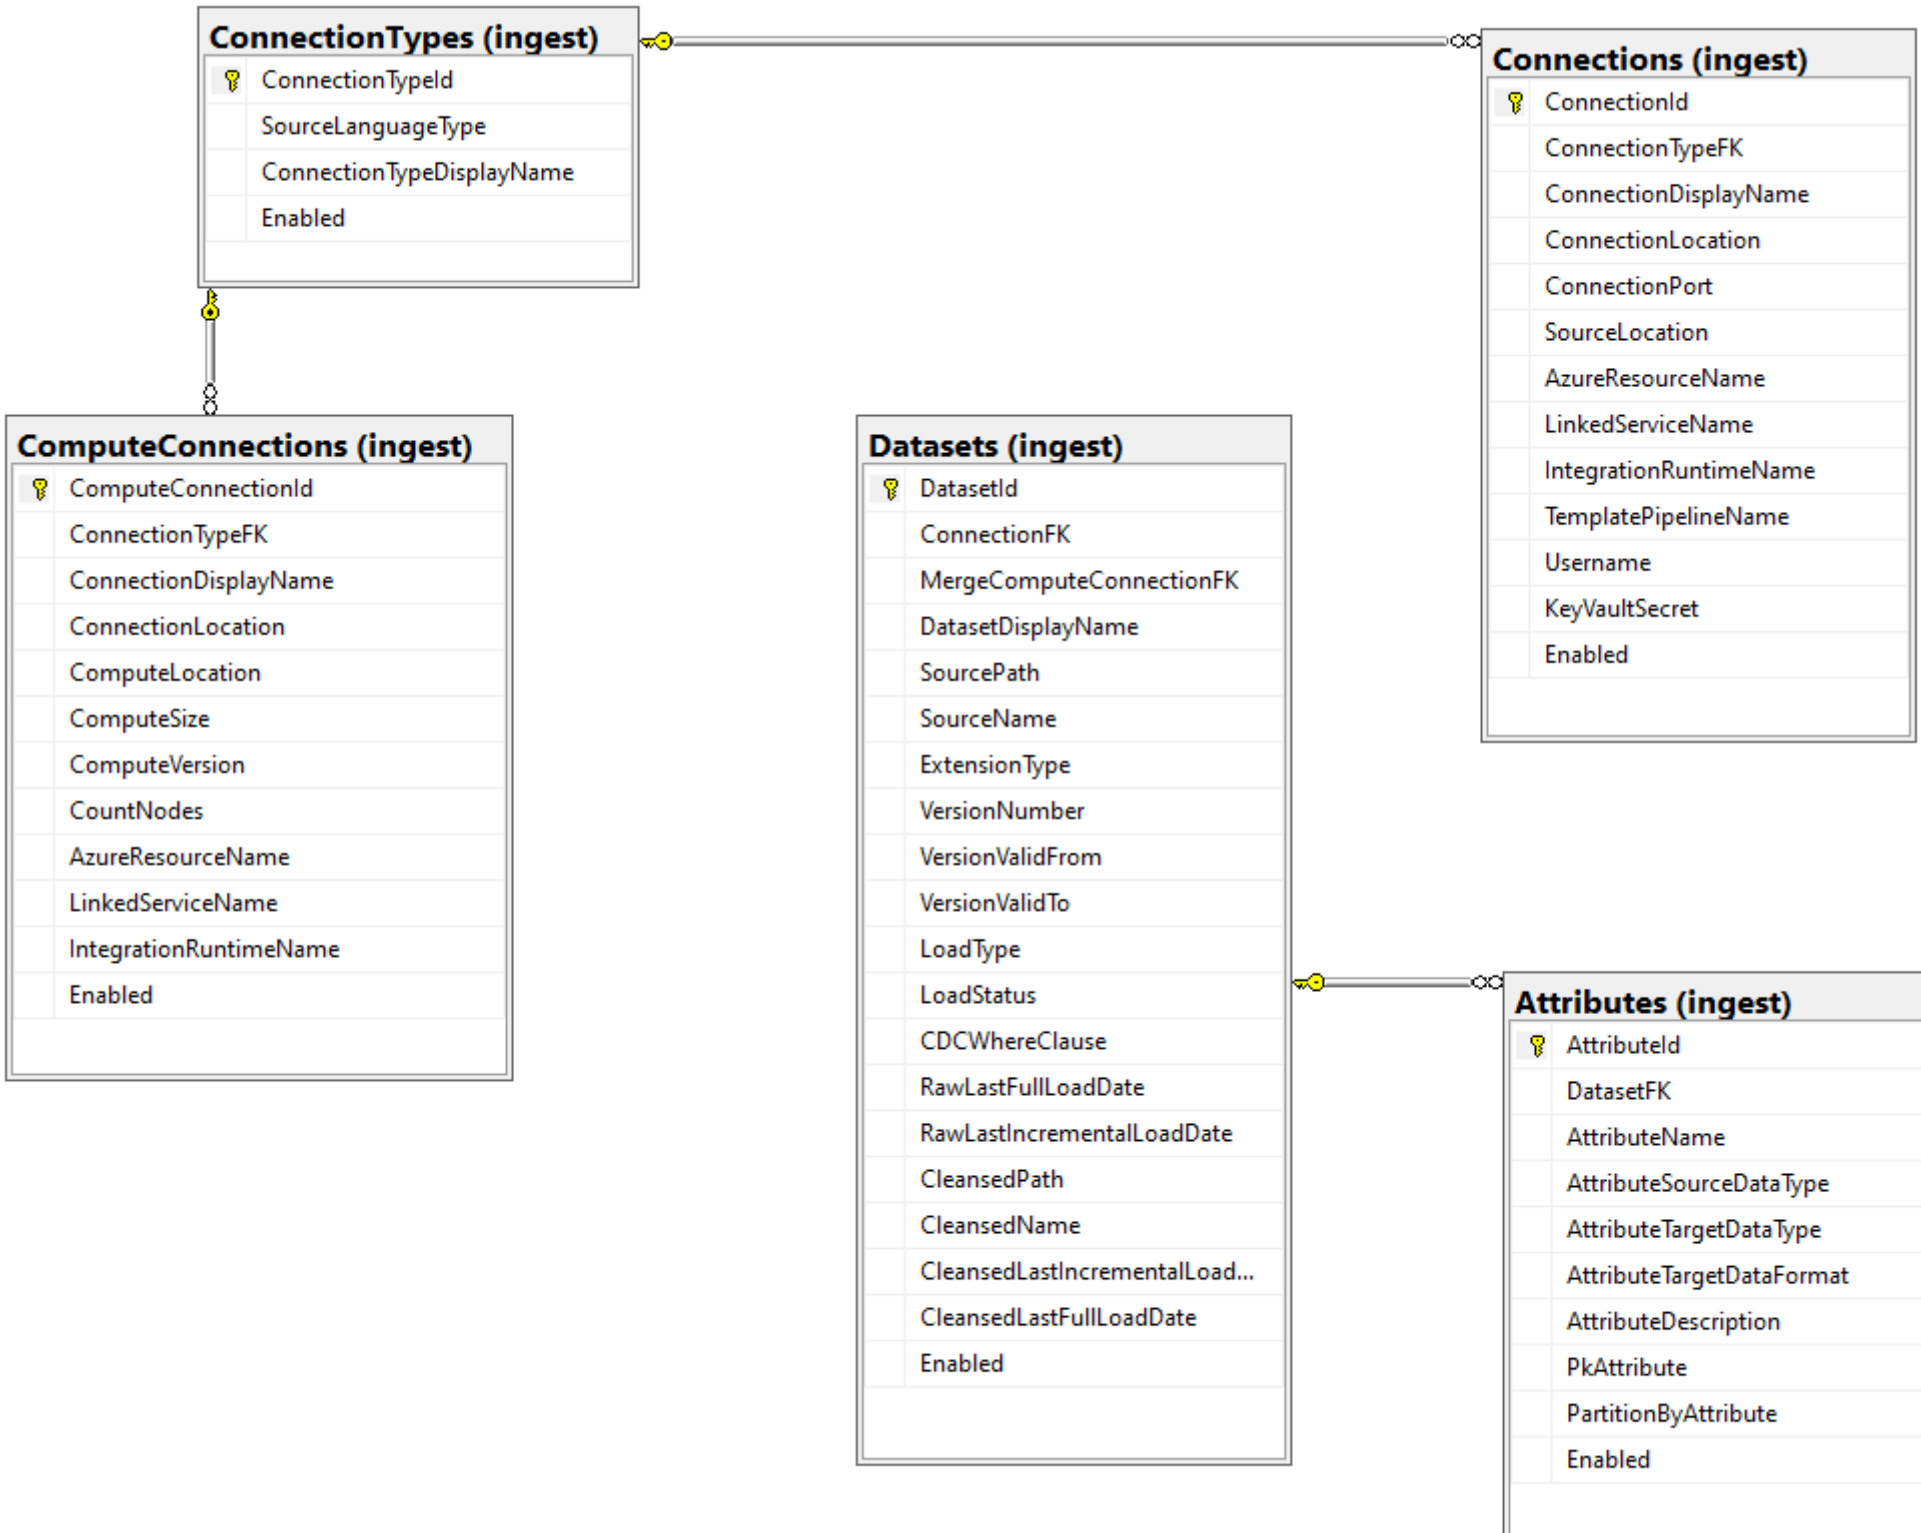
A view on the activity structure within a typical CF.Cumulus transform pipeline, exposing the Switch activity calling different variations of Notebook, depending on cluster sizing and authentication requirements.

CF.Control Database Internals



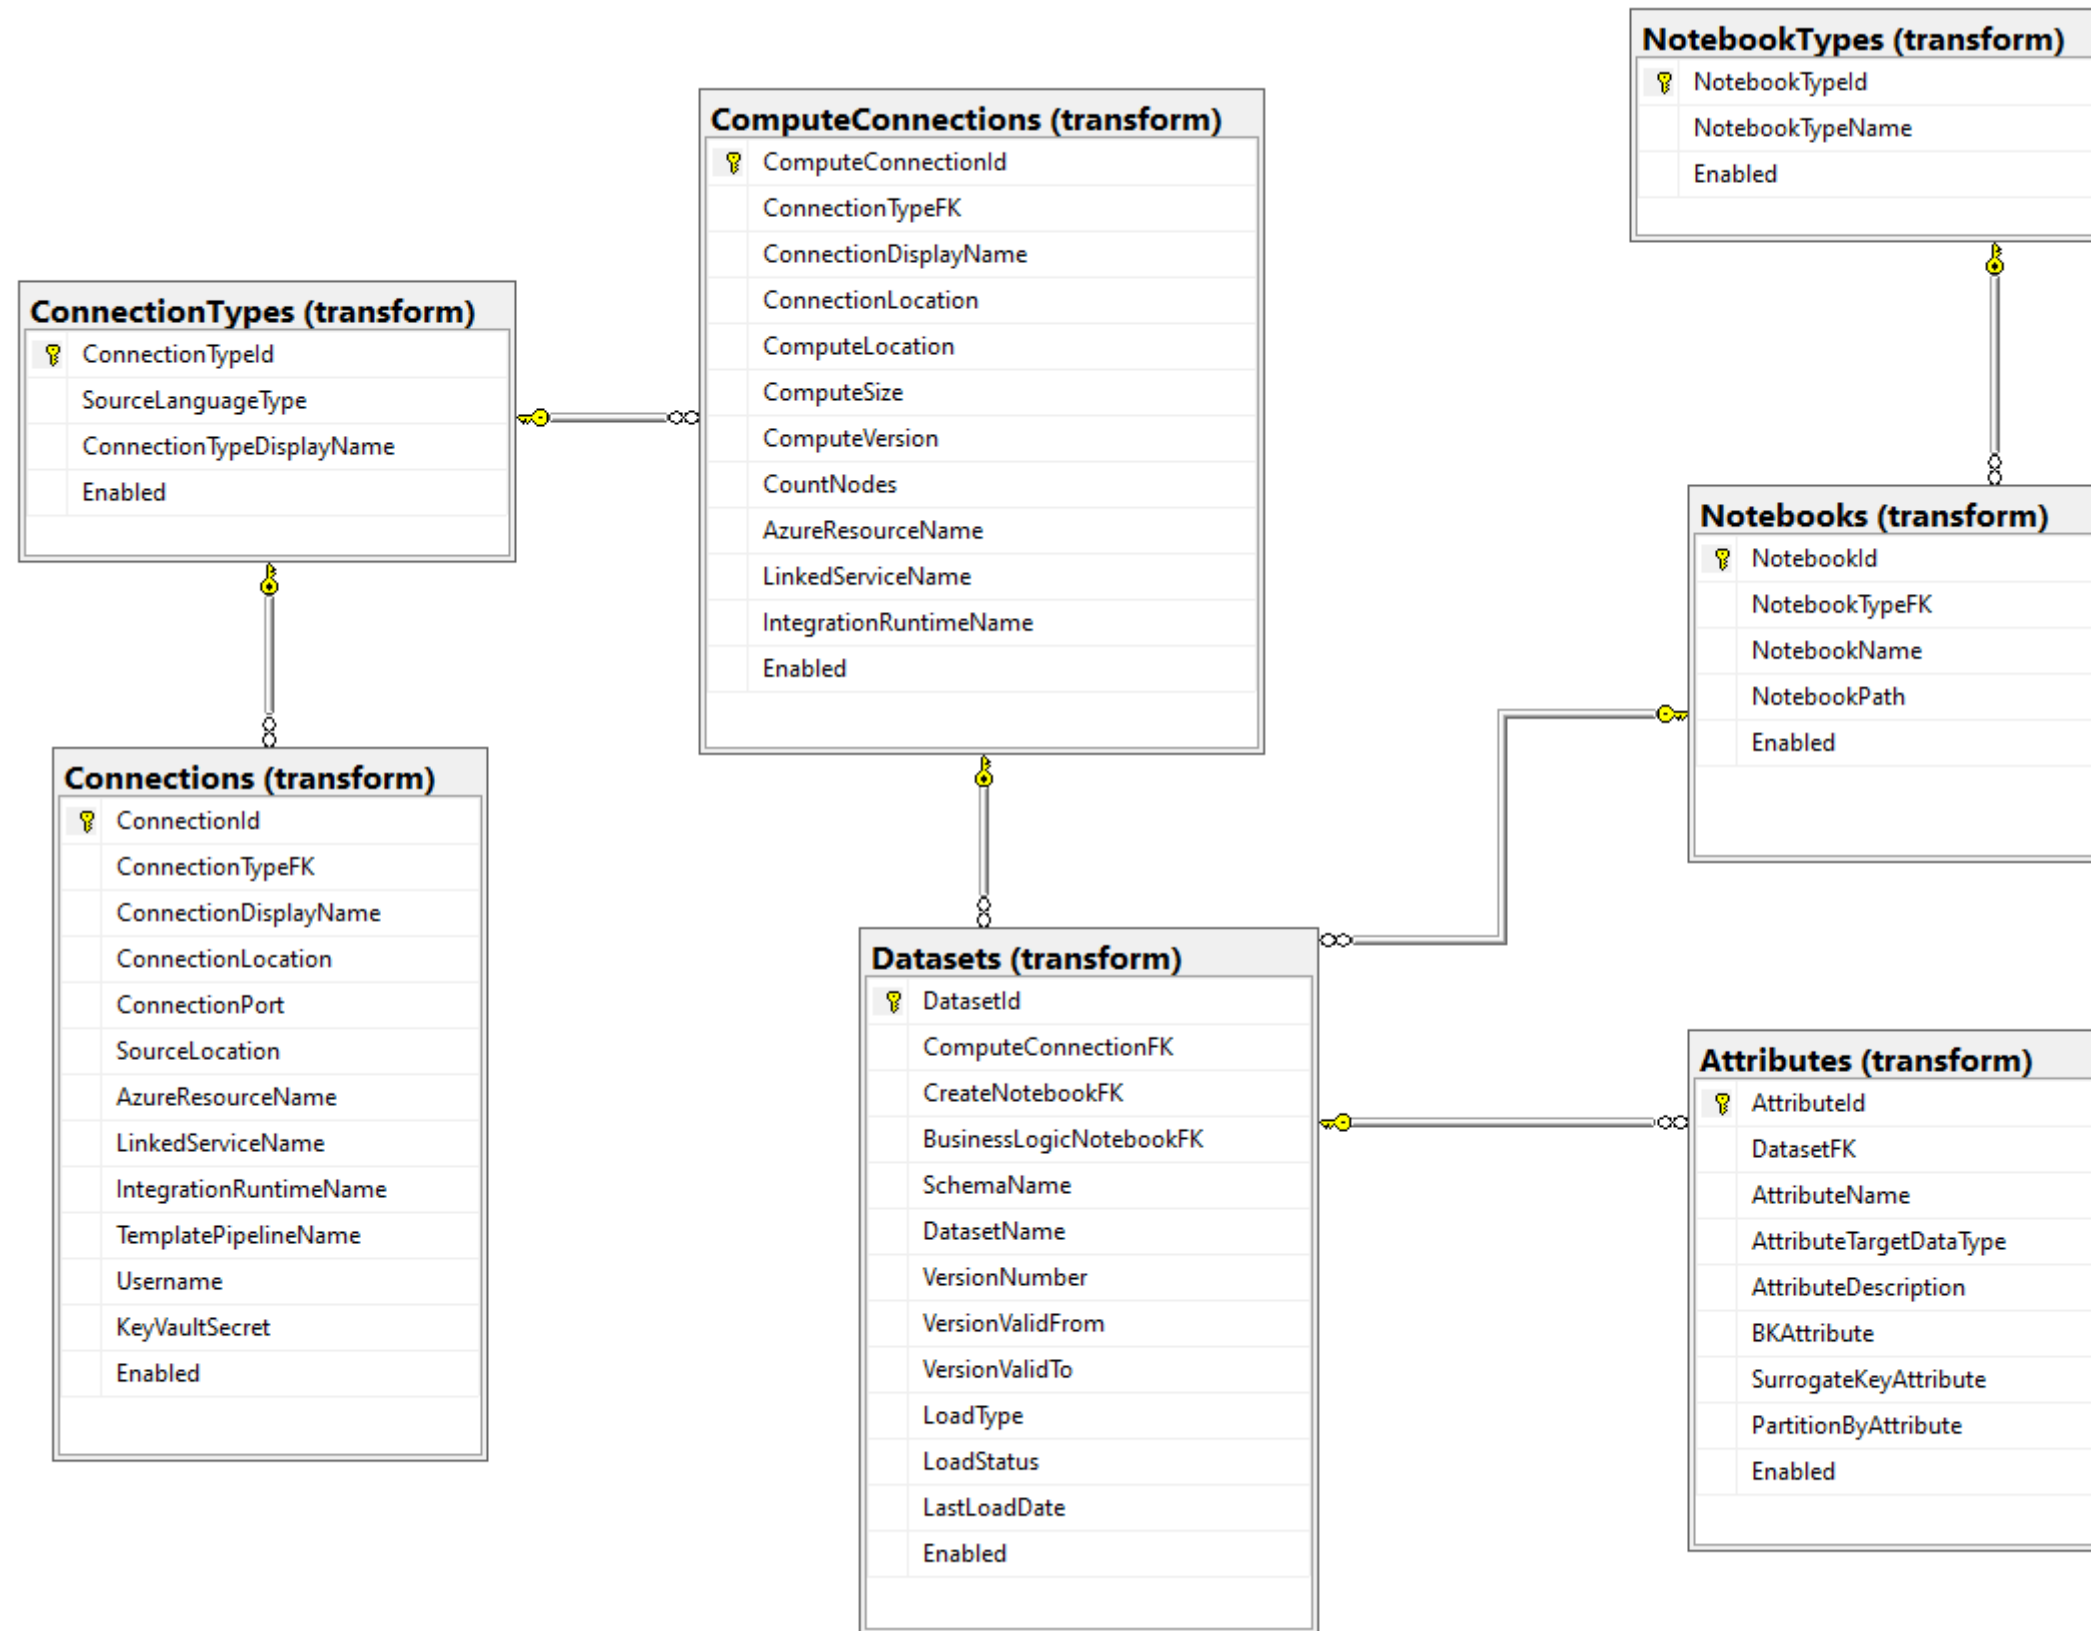
A view on the entities and relationships with CF.Cumulus metadata database. As presented for use by the Control pipelines.

CF.Ingest Database Internals



A view on the entities and relationships with CF.Cumulus metadata database. As presented for use by the Ingest pipelines.

CF.Transform Database Internals



A view on the entities and relationships with CF.Cumulus metadata database. As presented for use by the Transform pipelines.

Cloud Formations



Business

Contact Us

- <https://cloudformations.org>
- contactus@cloudformations.org
- [In/CloudFormations](https://www.linkedin.com/company/cloudformations)
- [@CloudFormsLtd](https://twitter.com/CloudFormsLtd)
- [CloudFormationsLtd](https://www.facebook.com/CloudFormationsLtd)

bit.ly/cf-meet

