

Best Data Platform for Data Engineering: Ilum Earns 49 Top-3 Rankings in G2 Winter 2026

Ilum appears in 103 G2 reports and earns 116 badges, highlighting a Kubernetes-native data lakehouse approach built on open standards, without vendor lock-in.

ALBUQUERQUE, NM, UNITED STATES,
December 17, 2025 /
EINPresswire.com/ -- Ilum, a
Kubernetes-native, open-source-first
data lakehouse and data platform built
around Apache Spark, today
announced its recognition in G2's

Winter 2026 Reports. Ilum earned 116 badges, appeared in 103 G2 reports, and achieved Top-3 placement in 49 reports (including 18 #1 rankings, 19 #2 rankings, and 12 #3 rankings).



G2 Winter 2026 Highlights for Ilum (Data Engineering, Analytics, MLOps)

- Broad recognition: Top-3 placement in 49 G2 Winter 2026 reports, reflecting consistent performance across multiple data categories.
- #1 rankings in key enterprise indexes: Ilum ranked #1 in the Enterprise Usability Index and Results Index for Data Governance and Big Data Integration Platforms.
- Enterprise implementation momentum: Ilum ranked #2 in the Big Data Analytics and ETL Enterprise Implementation Index, and #2 in the Enterprise Usability Index for Data Warehouse.

Ilum is a modern alternative for organizations evaluating platforms like Databricks, Snowflake, or Cloudera but who want a Kubernetes-first, infrastructure-agnostic architecture and an ecosystem anchored in open-source tools. Built for teams who want "Databricks-class" outcomes with open-source control.

At its core, Ilum is a Kubernetes-native control plane for the open-source data stack, built to run Apache Spark jobs, serve SQL and notebook workflows, and standardize governance and cost controls across cloud, on-prem, and hybrid environments.

Why users award Ilum so many badges?

Across data engineering and analytics organizations, Ilum is adopted for one repeatable outcome: delivering a "managed-platform" experience without abandoning open-source tools or infrastructure choice.

Users consistently choose Ilum because it:

- Accelerates time-to-value by unifying SQL, notebooks, orchestration, governance, and operations in one platform
- Reduces Kubernetes complexity with a control plane that makes day-2 operations predictable
- Keeps data portable by aligning to open standards and formats instead of proprietary lock-in
- Improves trust and compliance through automated lineage and versioning
- Makes cost actionable with pipeline/workload-aware cost attribution and optimization opportunities

From [Spark on Kubernetes](#) (2019) to a full data lakehouse platform

Ilum began in 2019 with a sharp focus on Apache Spark on Kubernetes, helping teams run Spark reliably on cloud-native infrastructure. Since then, Ilum has evolved into a full data lakehouse platform: still Spark-centric at the compute layer, but expanded to cover the operational and governance layers enterprises need to run production data workloads.

Data Lakehouse architecture built around the [Medallion](#) pattern

Ilum supports structured data lakehouse approaches like the Medallion architecture (Bronze/Silver/Gold) so teams can standardize ingestion, refinement, and serving layers with consistent metadata, access patterns, and operational controls. This turns “best practice diagrams” into something teams can actually implement and sustain in production. especially across multiple environments and teams.

[Spark Connect](#): a modern way to build and connect to Spark

Ilum’s Spark Connect capability adds a modern, developer-friendly way to connect clients to Spark while maintaining the platform’s operational model (multi-tenancy, resource controls, and governance). This improves the developer experience for interactive work while keeping production discipline for shared infrastructure.

Open standards: catalogs and table formats that keep data portable

Ilum’s lakehouse model is designed around interoperability. It works with a data catalog layer (commonly Hive Metastore / Unity) and supports open table formats such as Apache Iceberg, enabling reliable schema evolution, incremental processing patterns, and long-term portability.

For organizations operating in ecosystems that use Unity Catalog, Ilum is positioned to co-exist with existing catalog strategies while keeping the compute and operations layer open and Kubernetes-native.

Unified SQL, notebooks, and fully integrated Jupyter

Ilum includes its own SQL editor and notebook experience, and also supports fully integrated Jupyter Notebook inside the platform. The goal is simple: let teams build, test, and iterate using the interfaces they prefer, without breaking governance, operations, or reproducibility.

Automated lineage, ERDs, column lineage, and table versioning

Ilum automatically tracks what happens to data across the platform, table creation, changes, reads, and downstream usage then visualizes it via data lineage diagrams, ERD views, and column-level lineage. Ilum also tracks table versions and shows differences over time, helping teams audit changes, troubleshoot issues faster, and reduce the risk of silent data breakage.

Cost intelligence: pipeline-aware FinOps for the data stack

Ilum connects cost to reality: pipelines, jobs, clusters, and workloads. Teams can see the cost of specific data pipelines, identify inefficient resource usage, and find optimization opportunities, supporting a "FinOps-by-design" operating model instead of post-hoc billing reviews.

Integrated open-source modules for the modern data stack

Ilum provides an open-source-first platform experience with integrated modules such as Apache Airflow, MageAI, n8n, Kestra, Apache Superset (as the main BI tool), MLflow, and Gidea so teams can assemble a modern stack with fewer integration projects and a more consistent operating model.

The Winter 2026 results, 116 badges and inclusion in 103 G2 reports underscore why teams choose Ilum: it started with Spark on Kubernetes and matured into a complete open-source lakehouse platform that makes the medallion architecture practical in production, adds modern connectivity through Spark Connect, and unifies catalogs, open table formats, lineage, and cost controls under one Kubernetes-native operating model.

Learn more at illum.cloud, view documentation, or contact Ilum to discuss an evaluation.

Stan Ross
Ilum Labs LLC
stan.ross@illum.cloud

EIN Presswire's priority is source transparency. We do not allow opaque clients, and our editors try to be careful about weeding out false and misleading content. As a user, if you see something we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire, Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable in today's world. Please see our Editorial Guidelines for more information.

© 1995-2025 Newsmatics Inc. All Right Reserved.