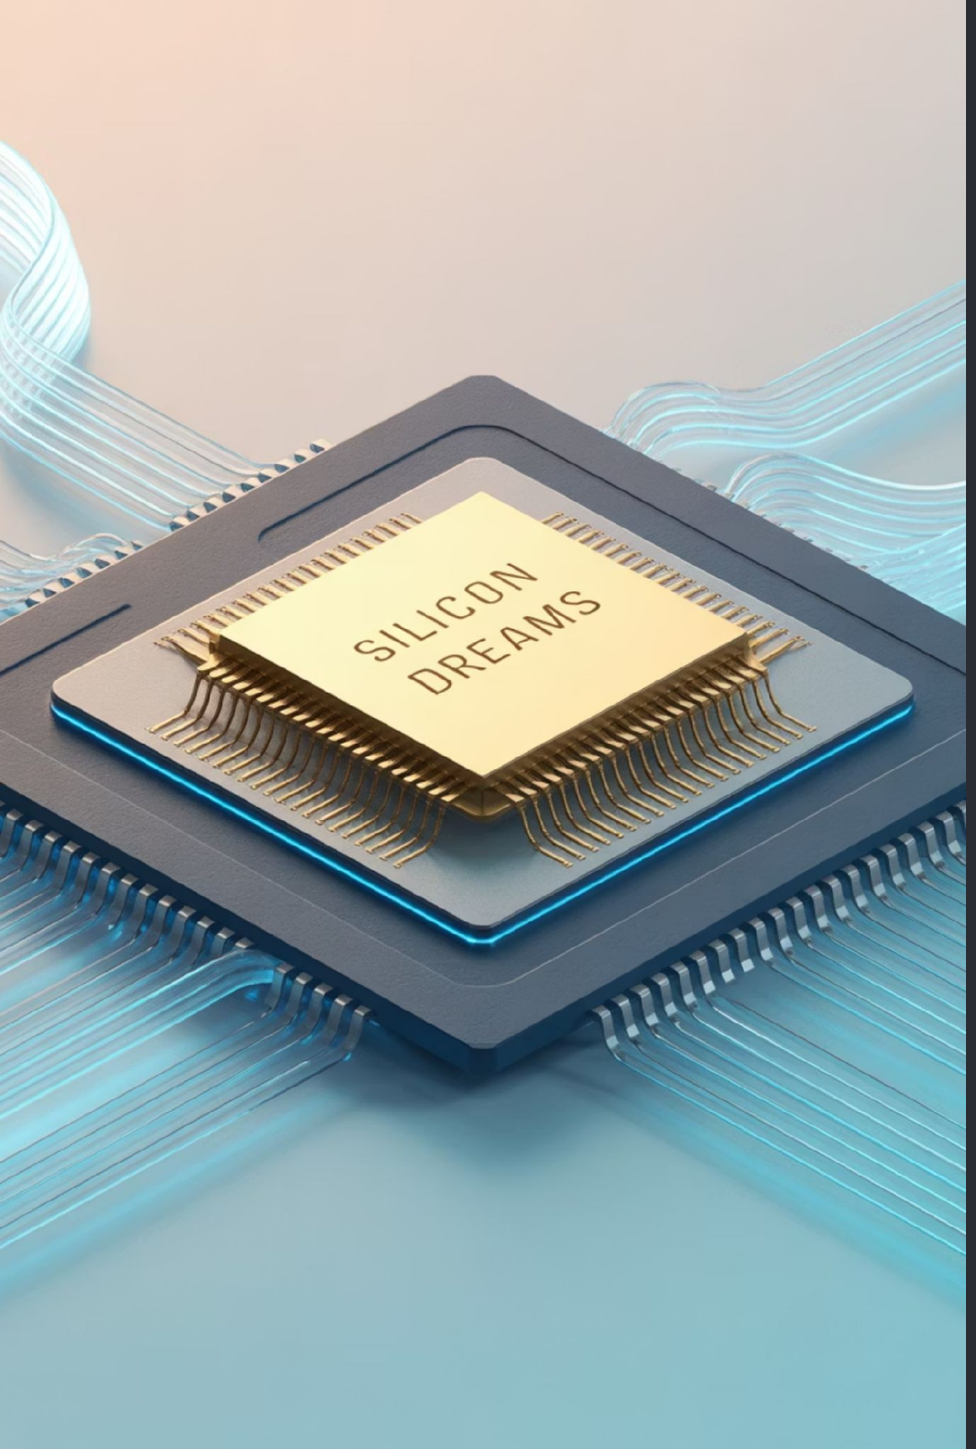

Scaling Analytics with KNIME Business Hub at AMD

Driving Operational Excellence and Data Maturity

Presented by:

Deepshikha Shekhawat

KNIME NYC Summit

June 18th, 2026

Powering the next wave of AI innovation at the edge

Real-Time Processing + Data, Security and Privacy + Personalization Customization

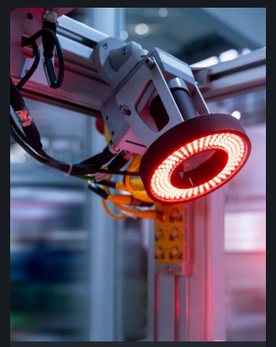

Predictive Maintenance

Professional Robots

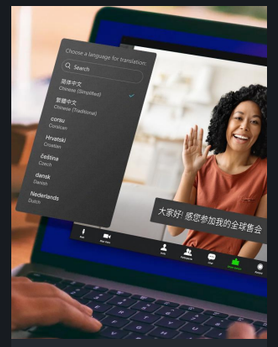

Video Conferencing

Smart Camera

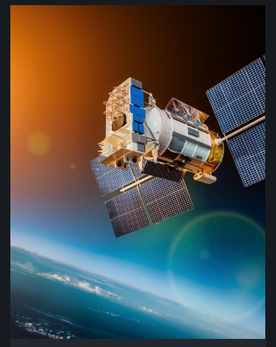

Aerospace and Defense

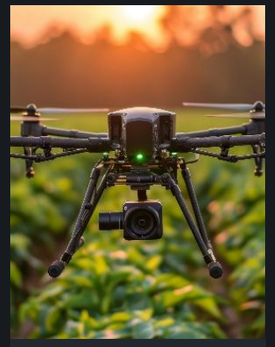

Edge Appliances

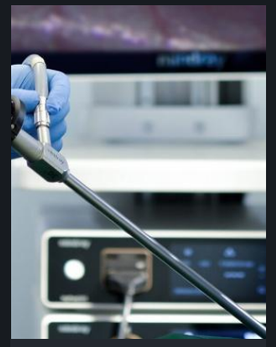

Healthcare

Professional Media Production

ADAS

AMD
together we advance_

57 years

Founded May 1, 1969
Headquartered in Santa Clara, CA

31,000+ employees

Accelerating next-generation computing

$34.6B annual revenue in 2025

Up ~34% year over year from $25.8B in 2024

5x market cap growth in 5 years

Top 100 most valuable companies in the world

100+ locations

Around the world

AMD
together we advance_

Data Maturity: The Cornerstone of Semiconductor Success

Innovation Engine

Data intelligence fuels innovation in operations process enhancements and product development

Operational Excellence

Enables predictive analytics for metric forecasting and KPI reporting

Competitive Advantage

Data maturity ensures agility, scalability, and market responsiveness in a fast-evolving industry

In today's semiconductor landscape, the ability to transform vast quantities of complex data into actionable insights determines market leadership and innovation velocity.

KNIME: Revolutionizing Data Workflows

KNIME Analytics Platform is a powerful open-source tool for accessing, blending, analyzing, and visualizing complex datasets with a visual programming interface.

KNIME's Integration Power

- Accelerates data-driven decision-making across R&D and business operations
- Enables rapid integration of disparate systems, streamlining data pipelines
- Reduces time-to-insight, allowing teams to focus on innovation rather than manual processes

Cross-functional teams can harness data without extensive technical expertise

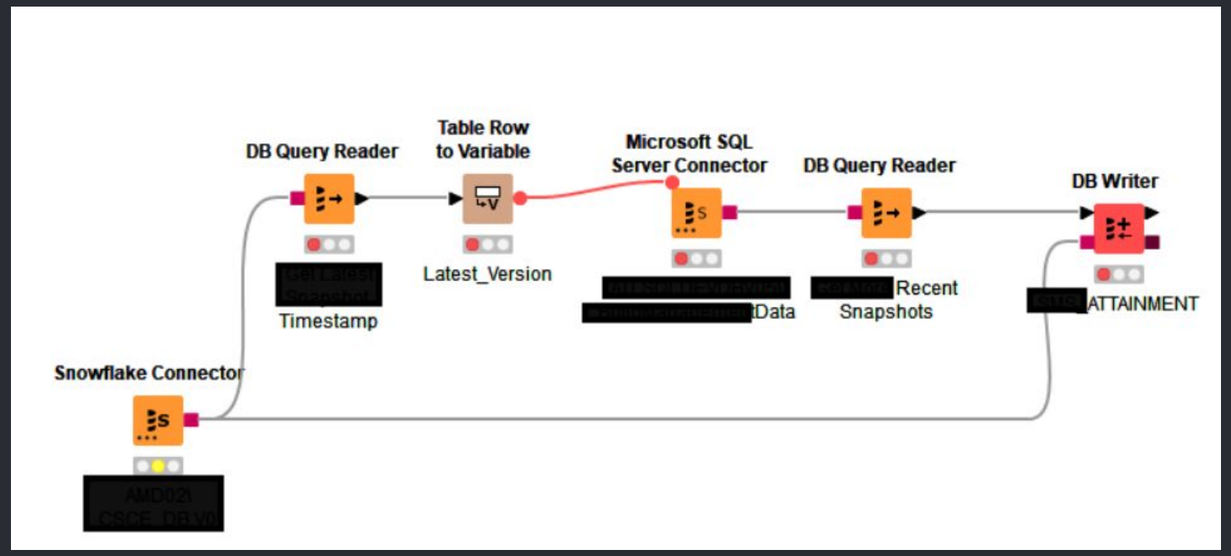

User-Friendly Learning Curve

01

Comprehensive ‘Getting Started’ documentation and YouTube tutorials simplify workflow creation.

02

Empowers AMD’s teams to build custom data workflows tailored to team’s analytics/data needs

03

Reduces onboarding time for non-technical users, fostering data literacy across the organization

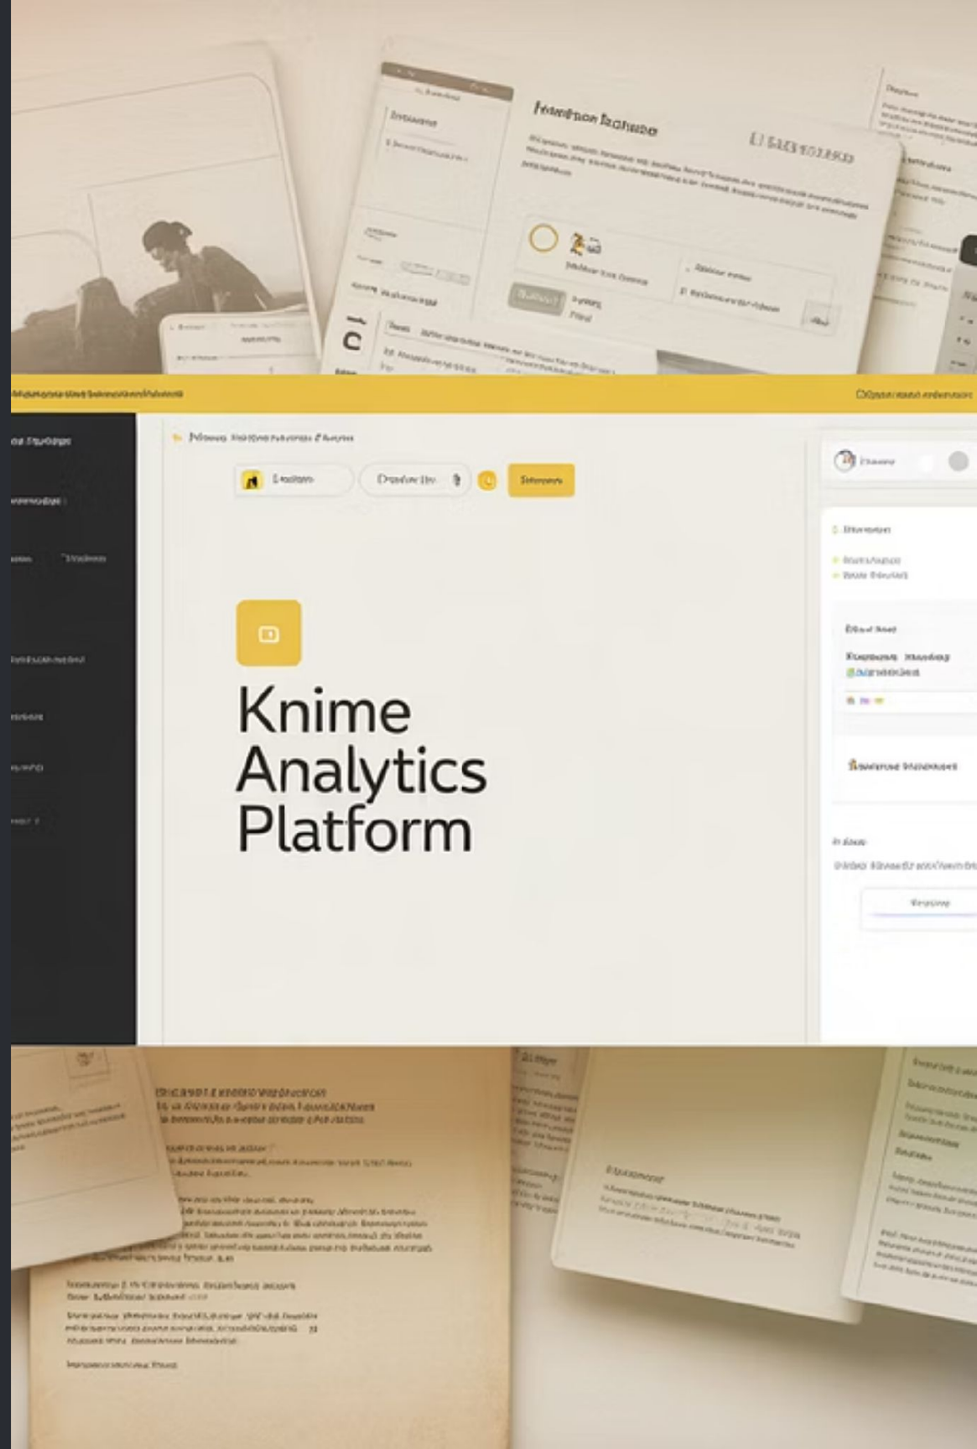

Automation for Operational Resilience

- Continuity: workflows run reliably even when team members are out, on leave, or change roles
- Reliability: built-in validation steps and schema checks reduce manual errors in data pipelines
- Consistency: predictable refresh cadences our dashboards and downstream consumers can trust

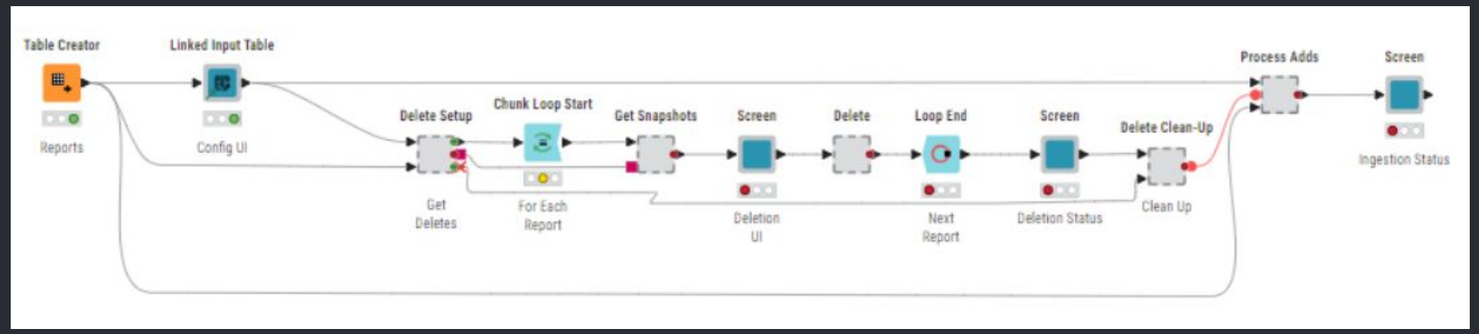

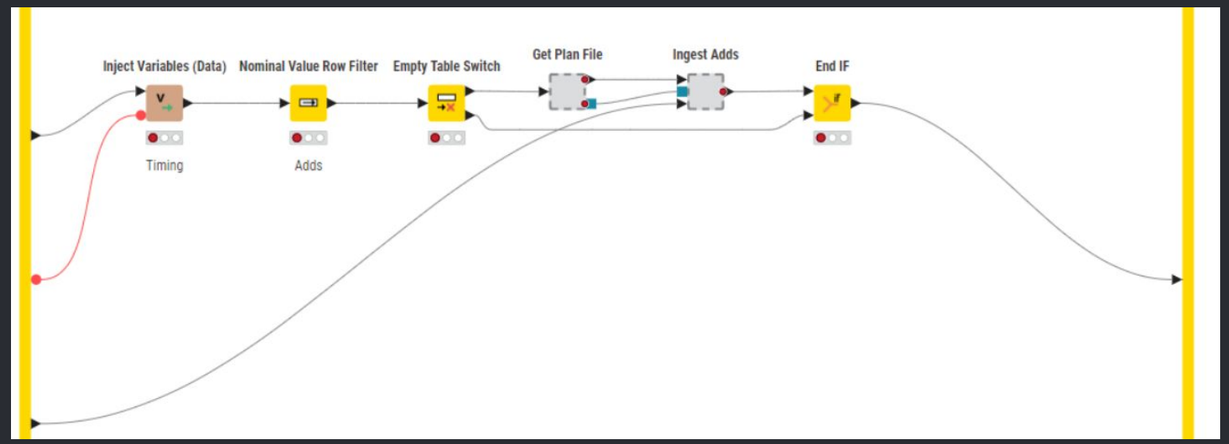

Mitigating Risks with Automated Data Flows

Security and Privacy Compliance

- KNIME’s transparent tools facilitate approvals for secure data handling and privacy compliance
- Supports obfuscation and encryption to protect sensitive semiconductor and customer data
- Simplifies audits, ensuring low-risk, secure data transfers across AMD’s ecosystem

Bridging Data Pipeline Gaps

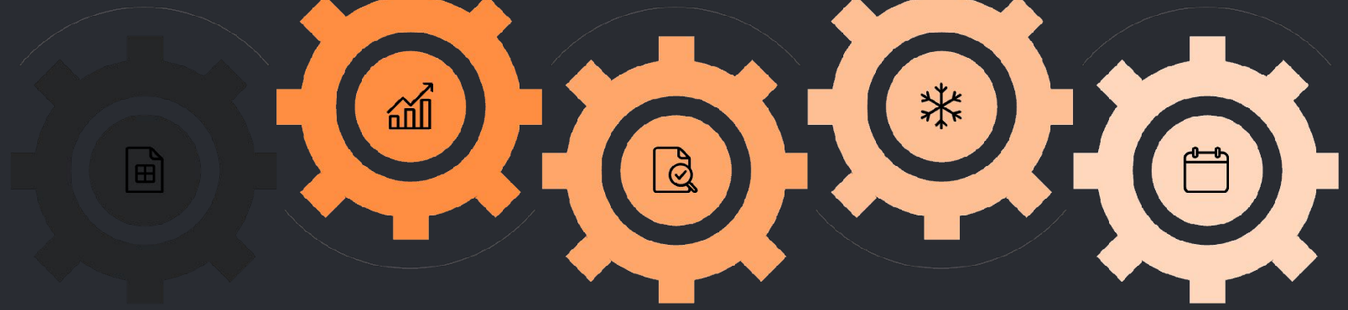

Scheduled Updates

Automates data refreshes to maintain currency in semiconductor operations dashboards

On-Demand Access

Pre-defined workflows enable non-technical users to extract detailed data as needed

Data Consolidation

Integrates Excel worksheets from multiple teams, cleaning and standardizing for database storage

Modernizing AMD's KNIME Workflows on Business Hub

Why we migrated, and how we approached it

THE STARTING POINT

Legacy 4.x workflows on individual desktops

- Person-at-machine bottleneck — workflows ran only when someone clicked Execute
- Credentials management challenges
- Inherited deprecated nodes and circular logic from previous developers
- Version compatibility breaking on newer platform releases

WHAT WE DID

Modernized legacy workflows on KNIME 5.4.4, deployed to Business Hub

01 Audit

Inspect workflows for deprecated nodes and legacy patterns

02 Remediate Locally

Replace deprecated nodes, validate on desktop

03 Publish to Hub

Upload cleaned workflow, verify no legacy nodes remain

04 Deploy

Configure scheduled or WebPortal deployment on shared executor

KNIME Hub: Orchestrating ~80% of Our Data Pipelines

BENEFITS REALIZED

Secure Credentials
Key-pair auth to Snowflake via Hub-managed secrets — no plaintext in workflow files

Hands-Off Scheduling
Workflows now run unattended on the General Shared Executor

Self-Service WebPortal
Business users trigger workflows via browser forms — no KNIME install needed

Reduced Maintenance Debt
Deprecated nodes and inherited anti-patterns systematically cleaned up

Dual-Mode Debugging
Desktop still resolves Hub secrets — engineers diagnose without modifying workflows

ORCHESTRATION LAYER

80%

of our data workflows and pipelines are managed end-to-end by KNIME Business Hub

WHAT IT ORCHESTRATES

- Scheduled ingestion into Snowflake
- Web Portal-triggered interactive runs
- Snowflake stored procedure calls
- Multi-step end-to-end pipelines

ALERTS & VISIBILITY

- In-workflow Send Email node — context-specific alerts (run summaries, validation failures)
- Hub-native deployment alerts — scheduled-run failures surfaced automatically

Three Workflows, Three Capabilities

From self-service ingestion to complex production orchestration

01

SELF-SERVICE INGESTION

Forecast Ingestion

- Business user uploads forecast Excel via WebPortal form
- Authenticates to Snowflake using Hub-managed secret
- Schema robustness: explicit data area locked, fail on schema change

SHOWCASES

Secure auth + Self-service UI

02

AUTOMATED SCHEDULED RUNS

Scheduled Loader

- Runs unattended on cron schedule via General Shared Executor
- Modernized: deprecated Snowflake driver, Row Filter, and Column Expressions replaced
- Now governed, scheduled, and visible to the team

SHOWCASES

Unattended automation + Legacy modernization

03

COMPLEX PRODUCTION WORKFLOW

Open POs

- Joins PO transactional data with Fiscal Dates dimension (year/quarter/month/week)
- Java Snippet preserves real business logic for conditional Quarter derivation
- Inherited circular rejoins and deprecated nodes cleaned up — logic preserved

SHOWCASES

Complex joins + Embedded business logic

Driving Accountability with KNIME

Proactive Monitoring & Alerts

Accountability without the babysitting — clear ownership and runtime visibility:

- SLAs monitored for runtime, row counts, and freshness thresholds
  — deviations surface automatically
- Every workflow has a named owner; alerts route to that owner with context (last success, row counts, rerun link)

- Result: rapid resolution, minimal downtime, and a clean audit trail for every run — no “where’s the latest file?” threads

KNIME: A Catalyst for AMD's Data-Driven Future

10-15

Daily/Weekly Workflows

Automated processes running consistently without manual intervention

5x

Faster Data Availability

Significant reduction in time-to-insight for critical operations data requiring frequent manual input

50%

Non-Technical Users

Half of team members now actively utilizing KNIME without coding expertise

KNIME helps bridge data gaps, creating seamless pipelines across different platforms and enabling the analytics and reporting needed to meet operational goals.

Strategic Benefits

Enhanced Productivity through automation of complex data workflows, significantly reducing manual processing time

Resolved Data Latency by empowering non-technical users to leverage data independently, democratizing analytics capabilities


AMD
High Performance and Adaptive Computing
Creating solutions to the world's most important challenges
1
AMD
together we advance_

Back-up


AMD

AMD: Pioneering Semiconductor Excellence

Advanced Micro Devices (AMD) is a global leader in high-performance computing, delivering cutting-edge CPUs, GPUs, and adaptive SoCs (system-on-chips) since 1969.

- Headquartered in Santa Clara, California, serving markets from gaming and data centers to AI and professional computing
- AMD posted record Q1 2026 revenue of **$10.3 billion**, up 38% year-over-year, with gross margin of **53%** (non-GAAP: 55%). Net income reached **$1.4 billion**, with Data Center revenue up 57% Y/Y to $5.8 billion
- Portfolio includes market-leading Ryzen processors, RDNA GPUs and the MI355X DLC rack packing **128 GPUs**, **36 TB of HBM3e memory**, and delivers up to **2.6 exaflops at FP4 precision**) marking a leap in AI compute density
- AMD sustained its lead in the supercomputing space, with its systems—**including El Capitan and Frontier**—retaining the top two spots on the latest Top500 list
- Recently partnered with IBM to develop hybrid systems integrating quantum computing with classical HPC
- At Financial Analyst Day 2025 AMD outlined long-term targets of >35% revenue CAGR and >60% Data Center CAGR; “Advancing AI 2026” (July 22–23, 2026) will reveal the next product roadmap

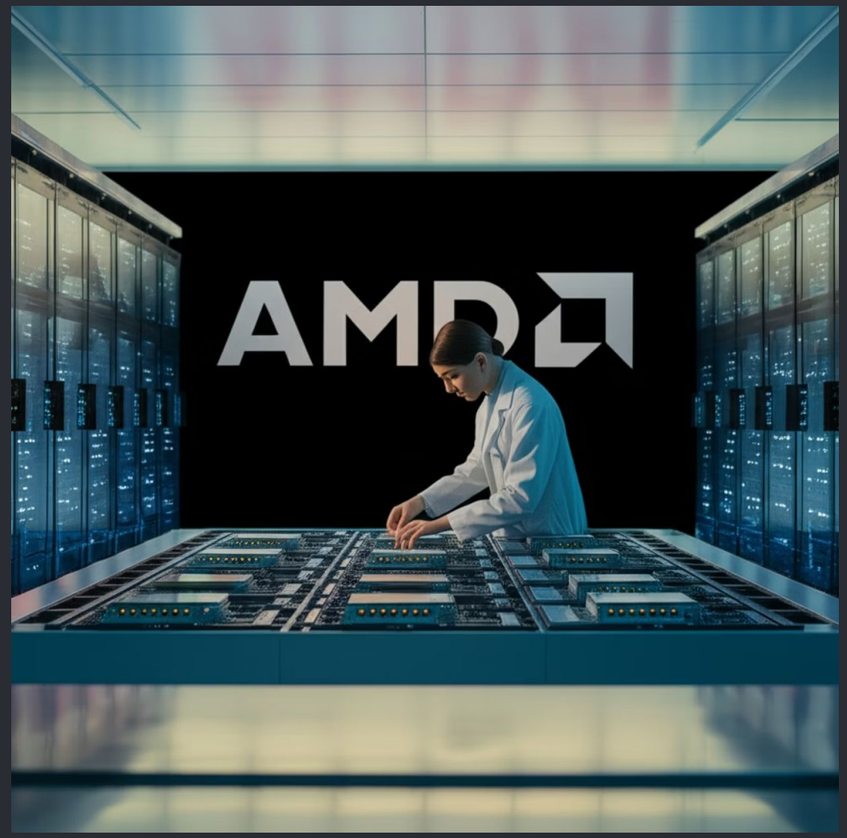

KNIME’s Most Used Features

DB Query Reader

Efficiently extracts specific datasets from relational databases, allowing precise data selection without requiring advanced SQL knowledge

DB SQL Executor

Executes complex SQL commands directly within workflows, enabling data transformation and manipulation operations at the database level

DB Connectors

Seamlessly integrates with various database systems including Snowflake, Oracle, and SQL Server to create unified data pipelines

Excel Writer

Exports processed data to formatted Excel files, facilitating data sharing with stakeholders who rely on spreadsheet-based reporting

These core components form the foundation of AMD's most frequently used KNIME workflows, enabling seamless data integration across platforms.