



INTER IIT
TECH MEET 14.0

HIGH PREP



.pathway

Financial Agents that work with
Real-time Insights

About the Company

Pathway is shaking the foundations of AI by introducing the world's first post-transformer model that adapts and thinks just like humans.

BDH Architecture paper on Hugging Face: <https://huggingface.co/papers/2509.26507>

The breakthrough architecture outperforms Transformer and provides the enterprise with full visibility into how the model works. Combining the foundational model with the **fastest data processing engine on the market** (**the framework at your disposal for this year's inter IIT**), Pathway enables enterprises to move beyond incremental optimization and toward truly contextualized, experience-driven intelligence. We are trusted by organizations such as NATO, La Poste, and Formula 1 racing teams.

Pathway is led by co-founder & CEO Zuzanna Stamirowska, a complexity scientist who created a team consisting of AI pioneers, including CTO Jan Chorowski who was the first person to apply Attention to speech and worked with Nobel laureate Geoff Hinton at Google Brain, as well as CSO Adrian Kosowski, a leading computer scientist and quantum physicist who obtained his PhD at the age of 20 & co-founded SPOJ, one of the earliest popular competitive programming platforms. The company is backed by leading investors and advisors, including Lukasz Kaiser, co-author of the Transformer ("the T" in ChatGPT) and a key researcher behind OpenAI's reasoning models.

Pathway is headquartered in Palo Alto, California, has offices in Paris and Wrocław, and hires exceptional talent remotely as well.

To stay updated and join Pathway's OS community, star the GitHub repositories below:

- github.com/pathwaycom/llm-app: Pathway's LLM xPack and tooling for RAG/agents
- github.com/pathwaycom/pathway: The core Pathway engine
- github.com/pathwaycom/bdh: The post transformer architecture

The Pathway Framework

The Pathway framework is a Python data processing framework designed for analytics and AI pipelines over data streams. It is the ideal solution for real-time processing use cases such as streaming ETL or Retrieval Augmented Generation (RAG) pipelines for unstructured, changing data.

Key Components of this definition:

- 1. Python Framework:** Written in Rust for speed and efficiency, Pathway is usable via Python, making it powerful yet simple to use with just Python know-how.
- 2. Data Processing:** The Pathway framework excels at processing large-scale, real-time data and is recognized as the world's fastest data processing engine. As a developer, you can use it for tasks like performing JOINS on incoming data streams (real-time data flow) or updating vector/hybrid indexes in real time. These are just simple examples, its potential goes much further.
- 3. AI Pipelines Over Data Streams:** The Pathway framework helps AI systems learn from real-time data streams, enabling applications like sentiment analysis, anomaly detection, and RAG pipelines that automatically adapt to incoming data.

Problem Statement Description

Financial institutions run on continuous signals across investing, risk, fraud, and customer engagement. Decisions lose value with every second of delay. Yet most stacks are batch oriented, siloed, and brittle.

We know you have been learning. This year we push you to go pro mode within clear rules. The difficulty is in applied GenAI. Expect a higher bar that mirrors real production and delivers real impact

This is your mission, should you choose to accept it: build smarter agentic applications that add value to real systems in the finance world.

Your application must run on the Pathway framework as the streaming data engine. It must use what streaming unlocks in practice:

- Always up to date metrics
- Incrementally updated deterministic models (Streaming ML)
- Ambient agents that react without prompts
- Live indexes that refresh as documents and events change

Build something a real team could deploy and a real company could grow from, not a classroom demo.

When companies deploy GenAI, they often start by redesigning processes. As engineers, use a clear decision lens to select the right mix of agents, classical machine learning, generative models, and RPA.



INTER IIT TECH MEET 14.0

Exploit the Pathway framework as much as possible. Go beyond connectors and live document indexing. Bring in stream processing, stateful windows, incremental joins, the MCP server for real time analytics, and incrementally trained classical models.

LLMs powered by Pathway BDH architecture learn on the go. Those models are not available for this challenge. Use current (static) open or closed source LLMs in combination with deterministic models that support incremental training/ Streaming ML.

Design for interoperability and orchestration so your system fits a world where agents are common. Add guardrails. You can consider supporting Model Context Protocol for tool access and agent to agent communication, including A2A patterns when they make sense. Make agents observable, traceable, and explainable. Log inputs, tools, parameters, and reasoning. Add scenario tests, real time error detection, and sensible fallbacks.

This should lead to: one production grade, streaming native, agentic system on the Pathway framework that ingests live or simulated streams, transforms data incrementally, maintains live indexes for retrieval, reasons over a continuously updated state, and takes the next best action with operator grade transparency and controls.

While you design, you can try to envision how your architecture could evolve/ simplify when post-transformer LLMs like Pathway Baby Dragon arrive with stronger reasoning, continual learning capability, and built-in interpretability.

Use-cases for your solution

This year's problem statement is use-case agnostic. However, engineering solutions for particular niche challenges is often a good idea from a value creation perspective.



INTER IIT TECH MEET 14.0

Ahead in this document, we've listed a few use cases (with hints) where demand exists in the market.

Pick one or combine several if you can keep the scope coherent. You can go beyond these, provided you validate the use case with substantiated research or via conversations & testimonials with industry practitioners.

Also note that the same use cases might seem overwhelming. They're doable if you have team members with complimentary skills and it's okay if you're still unable to implement all the specs. This use case list is just for you to get hints about various components that matter.

Sample use cases/ideas

1. Real Time AI Investment Decision

Imagine a production grade system on Pathway that behaves like an AI investment team. It ingests live signals, scores and ranks them, backtests on the same code path, proposes trades, and adapts as outcomes arrive.

You could include:

- Multimodal ingestion via Pathway connectors from market APIs, news and RSS, social, filings as PDFs, and earnings transcripts, etc with parsing and time alignment
- Streaming signal detection with sentiment, event classification, anomaly and risk indicators stored in a stateful knowledge base
- Hypothesis generation and ranking using rules or simple ML/LLMs guided by mission objectives
- Batch backtests inside the same pipeline with returns, volatility, drawdown, Sharpe, and constraint checks for limits and compliance
- Trade proposal artifacts plus real time alerts when breaking signals land
- Incremental updates using Pathway state persistence and online learning patterns
- Any sector or asset class, with alternative data when it supports the thesis

2. Live Credit Risk Evaluation and Next Best Action

Picture a real time credit flow on Pathway that keeps each borrower profile current and suggests the next best action with clear operator visibility.

You could include:

- Continuous ingestion from transactions, repayment behavior, bureau updates, and macro or news signals
- Real time recomputation of credit score and probability of default with incremental state per borrower
- Threshold driven actions such as review, tighten terms, request collateral, adjust limits, or cross sell when risk improves
- Human in the loop approvals where policy requires it LLM explanations grounded in live indexed policy documents with simple Q and A for reviewers
- Reliable, scalable outputs to dashboards, databases, portals, or APIs

3. Continuous Fraud Monitoring and Instant Response

Model a live fraud shield on Pathway that inspects every transaction stream and acts in sub second.

You could include:

- Streaming ingestion for cards, transfers, or crypto with in memory enrichment from profiles, geo, and device data
- Rules, statistical anomalies, and graph style links for rings, with velocity and sequence features using windows, aggregations, and joins
- Incremental per event computations for high throughput
- Instant alerts or actions such as decline, lock, or analyst notification via queues or APIs
- Joins with live sanctions or blacklist feeds and optional consortium pattern feeds
- LLM assisted SAR drafting grounded in current pipeline facts

4. Real Time Customer 360 and Personalized Engagement

Compose a streaming Customer 360 on Pathway for financial services that keeps a live profile per customer and triggers timely, personalized engagement.

You could include:

- Continuous ingestion from card payments, instant payment rails, ACH and wire transfers, core banking ledger updates, trading and brokerage events, loan and KYC updates, web and app events, support logs, and optional consented social signals into a table keyed by customer or account ID. This is a common streaming use case akin to the [Kafka timestamps example](#) on Pathway's website.
- Triggers over the streaming profile using filters, joins, and windowed aggregations such as churn risk, inferred life events, preapproved offer eligibility, and fraud or AML anomalies. Focus on triggers that must respond during authorization and instant payment flow.
- LLM messages and 3 point summaries grounded in live indexed facts from core systems, CRM, trading platforms, and service tickets. Keep each message explainable with links to source records
- A small scenario that measures event to action latency and demonstrates relevance, for example a card decline, a large deposit, or a risky transfer that triggers a personalized message, an agent script, or an extra verification step. Report event time, decision time, and customer or agent action time

5. Smart KYC Screener

KYC teams juggle identity docs, sanctions hits, PEP checks, and adverse media. Workloads spike. False positives waste hours. You can build an interactive screening copilot that automates level one checks, keeps a human in the loop for edge cases, and explains every decision with evidence. Things you can consider prioritizing:

- Clean UI and understanding of how it either rivals or integrates into KYC startups.
- Interactive level one screening. Upload ID and company docs, extract names and details, screen across OFAC SDN, UN, EU and UK consolidated lists, local and vendor lists, and summarize hits with confidence, reasons, and evidence links
- Adverse media sweep. Pull recent articles and filings, cluster by topic, highlight potential risk, and show the exact lines that triggered a flag
- Deterministic plus LLM plus RPA. Use a rule engine for exact and fuzzy matches, PEP and RCA classification, ownership checks like the fifty percent rule, and freshness checks on lists. Use LLMs for disambiguation and plain language summaries. Use RPA to open or update cases in your workflow and core systems
- Automate data capture from IDs and forms with grounded evidence via advanced multimodal parsers. Handle passports, licenses, proof of address, company registers, and beneficial ownership. Keep cropped image snippets and field level citations for traceability
- Track and display quality metrics such as hit rate, true positive rate, false positive reduction, time to decision, and analyst time saved

Conclusion: How to use these ideas

The items above are ideas to spark direction, not tracks or requirements. Use them as building blocks that you can mix and adapt to your domain while keeping one coherent, production minded system on Pathway. You may pursue a single end to end flow or compose parts across ideas, as long as the streaming core, observability, grounding, and evaluation standards remain clear. Feel free to substitute data sources, models, and actions where justified. You can also go beyond these examples if you validate the use case with a practitioner.



INTER IIT TECH MEET 14.0

Whether you focus on risk management, fraud prevention, investment analysis, or customer engagement, the common essence is automation with immediacy – the system reacts to data as soon as it arrives, enabling intelligent next steps without delay.

By combining Pathway's core abstractions (i.e. connectors, tables, transformations, and the powerful Rust-based incremental engine) with modern AI elements such as deterministic models, incremental training, and agent orchestration, you will demonstrate what real-time AI looks like in practice.

Production Bar Checklist

- Pathway pipeline for streaming ingest. Handle streaming tables and changing documents. See Pathway app templates in the ETL section.
- Live indexing for retrieval. Watch folders and endpoints. Auto refresh embeddings and metadata.
- Orchestrator agent that plans tool use and handoff across skills. Allow agent-to-agent calls.
- Human in the loop or fully autonomous is acceptable, but be explicit. Provide citations, review points, approvals, and an easy override with a kill switch.
- Guardrails including role and permission checks, safe tool invocation, PII protection, and redaction.
- Observability, including traces, logs, metrics, prompt and tool call history, and a dashboard with alerts.
- Evaluation including offline retrieval and answer quality, scenario tests, and failure drills.
- Ethics and risk, including misuse cases, bias checks, consent flow, and an audit trail.
- Product fit, including a crisp problem statement, KPIs, SLOs, and a business impact model validated with a practitioner.
- Runbooks including install steps, environment and secrets, seed data, and operations.
- Retrieval quality, including precision and grounded answers with citations. Accuracy must be demonstrated.
- No unnecessary vendor usage. If Pathway can supply a component, use it to keep the stack practical



INTER IIT TECH MEET 14.0

- Modularity that plugs into common toolsets and developer environments.
- **Do not use conflicting components.** Avoid unverified frameworks that reduce reliability in critical settings. Avoid tools that cannot support streaming or real-time processing.
- Forward-looking note: Document where a future BabyDragon model would fit for complex reasoning and on-the-go learning. Keep the current build production grounded.

Deliverables at the end

Codebase

A well-structured, production-ready repository using Pathway as the data engine. It should demonstrate both streaming and batch capabilities where relevant, and may integrate additional frameworks for ML, LLM, or automation components.

Code must be modular, documented, and runnable with clear setup instructions for a clean environment. Include configuration files, environment variables, and dependency management.

Documentation

A concise but complete description of the architecture, design decisions, and workflow.

Include:

- A high-level architecture diagram and brief component explanation
- How Pathway is used for streaming, state handling, and incremental computation
- How your system achieves real-time behavior and integrates external tools or APIs
- Notes on scalability, observability, and extension to other use cases or domains

Example Scenario

Demonstrate your solution in one realistic scenario. This could simulate a continuous stream (financial, operational, or customer data) and show how the system reacts, updates, and outputs decisions or insights.

You may script it through a Jupyter notebook, CLI replay, or lightweight dashboard.



INTER IIT TECH MEET 14.0

Demonstration Pipeline

A running instance or recorded demo illustrating live updates, reasoning flow, and outputs. Deployment to cloud is optional, but the demo should clearly show streaming ingestion, transformation, and action.

Short video demonstrations are encouraged.

Testing & Evaluation

Basic unit or component tests validating the correctness of core functions.

Optionally, include load or latency tests that illustrate system responsiveness and reliability.

Team Presentation

A short presentation summarizing your concept, architecture, challenges, and key learnings. Highlight how you leveraged Pathway's streaming engine and why your design matters in a real-world context.

Submission Format

All deliverables should be shared via a zip file containing:

- README with installation steps, architecture diagram, and demo links
- Documentation (PDF or markdown)
- Example data or scripts to replay the scenario
- Recorded video or live demo link

Final presentations will showcase how teams applied real-time reasoning and streaming intelligence in their chosen domain or mix of directions.

Evaluation Criteria

Midterm report (research + experiments) - 10%

- Problem validated and user journey defined.
- Architecture sketch.
- Pathway usage clarity
- Data plan (lawful sources, grounding) and backtest/live parity plan.
- Scenario test plan (late events, API timeout, failure drill)

Novelty of use-case - 20%

- Clear pain point for the finance/financial services industry with evidence.
- Streaming-first technique used meaningfully.
- Novelty in techniques used for higher accuracy, continual learning, explainability, etc.
- Differentiation from existing solutions discussed in conferences/offered by startups

Technical Implementation - 35%

- Pathway as core for streaming, pre/post processing, and live documents indexing.
- Enterprise readiness via grounded answers with citations, modularity, guardrails, exhaustivity, etc.
- Agent planner, retries/timeouts, idempotency; MCP/A2A/Agentic Web where apt.
- Observability and tests: traces/logs/metrics, scenario/CI; latency/throughput and cost reported.

Solution utility & simplicity of architecture - 20%

- Minimal, modular design with clear boundaries.
- One-command run, container image, env-based config, seed data.
- Explainable outputs, citations, approvals, useful errors, and easy to navigate UI.
- Stable APIs/outputs to DB/queues/dashboards; basic runbooks

End-term report - 15%

- Results vs KPIs with charts; limits noted.
- Repro steps, configs, infra notes, license.
- Benchmarks & Evaluation details: retrieval quality, latency, cost, parity evidence, etc.
- Zip file with clear README and demo.
- Future work and where BDH could slot in (less weightage given limited context available at participants' end).

Developer Resources

To ensure your project demonstrates strong real-time and production-readiness capabilities, all teams **must adhere** to the following requirements and leverage the official Pathway developer ecosystem. Exploit pathway to the max

1. Live Data Ingestion with Pathway Connectors

Your system must utilize **Pathway's real-time connectors** to ingest streaming data relevant to your chosen use case.

Pathway provides built-in connectors for **files, databases, message queues, APIs, and web sources**, all operating in streaming mode, ensuring results update in real time as data changes.

If your desired data source is not directly supported, you must extend Pathway's ingestion layer by implementing your own connector using the Custom Python Connector. This allows you to adapt Pathway to new sources and contributes to enhancing the engine's capabilities.

Documentation: <https://pathway.com/developers/user-guide/connect/connectors-in-pathway>

Create a custom Python connector: <https://pathway.com/developers/user-guide/connect/connectors/custom-python-connectors>

Python web scraper example: <https://pathway.com/developers/user-guide/connect/python-web-scraping>

Artificial Data Streams with the demo Module (in case you find it difficult to access free streaming data APIs): <https://pathway.com/developers/user-guide/connect/artificial-streams>

At least one live or simulated data feed must be integrated.

Examples include:

- Subscribing to live market APIs (e.g., Alpha Vantage, Polygon.io)
- Reading transaction streams from Kafka or sockets
- Simulating live events via Pathway's demo utilities

If live data is unavailable, teams may simulate streaming input by replaying static datasets with realistic time intervals.

2. Core Concepts

Before starting development, familiarize yourself with Pathway's foundational ideas and architecture. These concepts will help you understand incremental computation, table semantics, and event-driven design principles that power every Pathway pipeline.

Pathway Core Concepts: <https://pathway.com/developers/user-guide/introduction/concepts/#core-concepts>

3. Streaming Transformations and Feature Engineering

All data transformations must be performed in streaming mode using Pathway's transformation APIs.

Your pipeline should support:

- Incremental joins, filters, and aggregations
- Stateful window computations
- Real-time feature engineering for signals and indicators

Documentation: <https://pathway.com/developers/user-guide/data-transformation/table-operations>

Temporal Data Windows: <https://pathway.com/developers/user-guide/temporal-data/windows-manual>

Ensure computations are low-latency and modular, with clear separation between ingestion, transformation, and output modules

4. LLM Integration for Real-Time Insights

To make your system interactive and human-centric, integrate **Pathway's LLM xPack**, enabling smooth orchestration of retrieval, summarization, and reasoning over live data.

You may use it for:

- Live retrieval-augmented generation (RAG)
- Automated report generation
- Explainable insights (e.g., credit decision rationale, fraud summary reports)

Documentation: <https://pathway.com/developers/user-guide/llm-xpack/overview>

Pathway MCP Server: https://pathway.com/developers/user-guide/llm-xpack/pathway_mcp_server

5. Mandatory Learning Resources and Templates

These resources will accelerate development and ensure alignment with Pathway's architecture.

Templates

RAG App Templates (YAML): <https://pathway.com/developers/templates/>

ETL and ML Time-Series Pipelines (Live Data Framework): <https://pathway.com/developers/templates/?tab=live-data-framework>

Hands-On Tutorials

Vanilla RAG with OpenAI (Python):

<https://pathway.com/bootcamps/rag-and-llms/coursework/module-5-hands-on-development/1-first-rag-pipeline/building-with-open-ai>

RAG with Gemini:

<https://pathway.com/bootcamps/rag-and-llms/coursework/module-5-hands-on-development/1-first-rag-pipeline/rag-with-gemini-and-other-open-ai-alternatives>

Real-Time RAG using LlamaIndex:

<https://pathway.com/bootcamps/rag-and-llms/coursework/module-5-hands-on-development/3-realtime-rag-with-llamaindex-langchain-and-pathway/implementation-with-llamaindex>

Real-Time RAG using LangChain:

<https://pathway.com/bootcamps/rag-and-llms/coursework/module-5-hands-on-development/3-realtime-rag-with-llamaindex-langchain-and-pathway/implementation-with-langchain>

Evaluation & Benchmarks

Evaluating RAG Applications with RAGAS: <https://pathway.com/blog/evaluating-rag>

Deployment

Docker Deployment Guide: <https://pathway.com/developers/user-guide/deployment/docker-deployment>

Persistence and Fault Tolerance: <https://pathway.com/developers/user-guide/deployment/persistence>

Licensing Guide (For unlocking Advanced Features): <https://pathway.com/developers/templates/licensing-guide>

Reference Implementations

Option Greeks Computation with Databento: <https://pathway.com/developers/templates/etl/option-greeks/>

Real-Time Multimodal Data Processing (Docling): <https://pathway.com/blog/multimodal-data-processing>

La Poste ETA Microservices Case Study: <https://pathway.com/blog/pathway-laposte-microservices/>

Pathway Community Spotlights: <https://pathway.com/blog/?tag=community>, <https://pathway.com/blog/ai-agents-finance-due-diligence/>, <https://pathway.com/blog/live-ai-multi-agentic-rag>, <https://pathway.com/blog/financial-intelligence-with-event-based-state-machine>