


[DISTRIBUTIONAL]

# Distributional's Product



# Intro

**Distributional has updated and continues to update our core functionality.**

**See [docs.dbnl.com](https://docs.dbnl.com) for the latest on our data pipeline, analytics workflow, and enterprise platform.**

Production AI behavior is a black box, which leaves AI teams struggling to improve, fix, and scale their products. Distributional's adaptive analytics platform empowers these teams to have confidence in AI behavior—the interplay and correlations between users, context, tools, models, and metrics. The product (DBNL) provides a detailed snapshot of aggregate AI product behavior, regularly surfaces behavioral signals, links evidence to facilitate deep root cause analysis, and enables teams to track relevant signals over time. This empowers AI teams to better understand the behavior of their users and AI products so that they can fix and improve those products with confidence.

Distributional's adaptive analytics product includes an analytics workflow that is supported by a data pipeline and powered by an enterprise platform. The focus of this paper is to give you a high level introduction into each of these three components that comprise the product: the platform, pipeline, and workflow.



# Table Of Contents

Page 4:      Product

Page 5:      Enterprise Platform

Page 7:      Data Pipeline

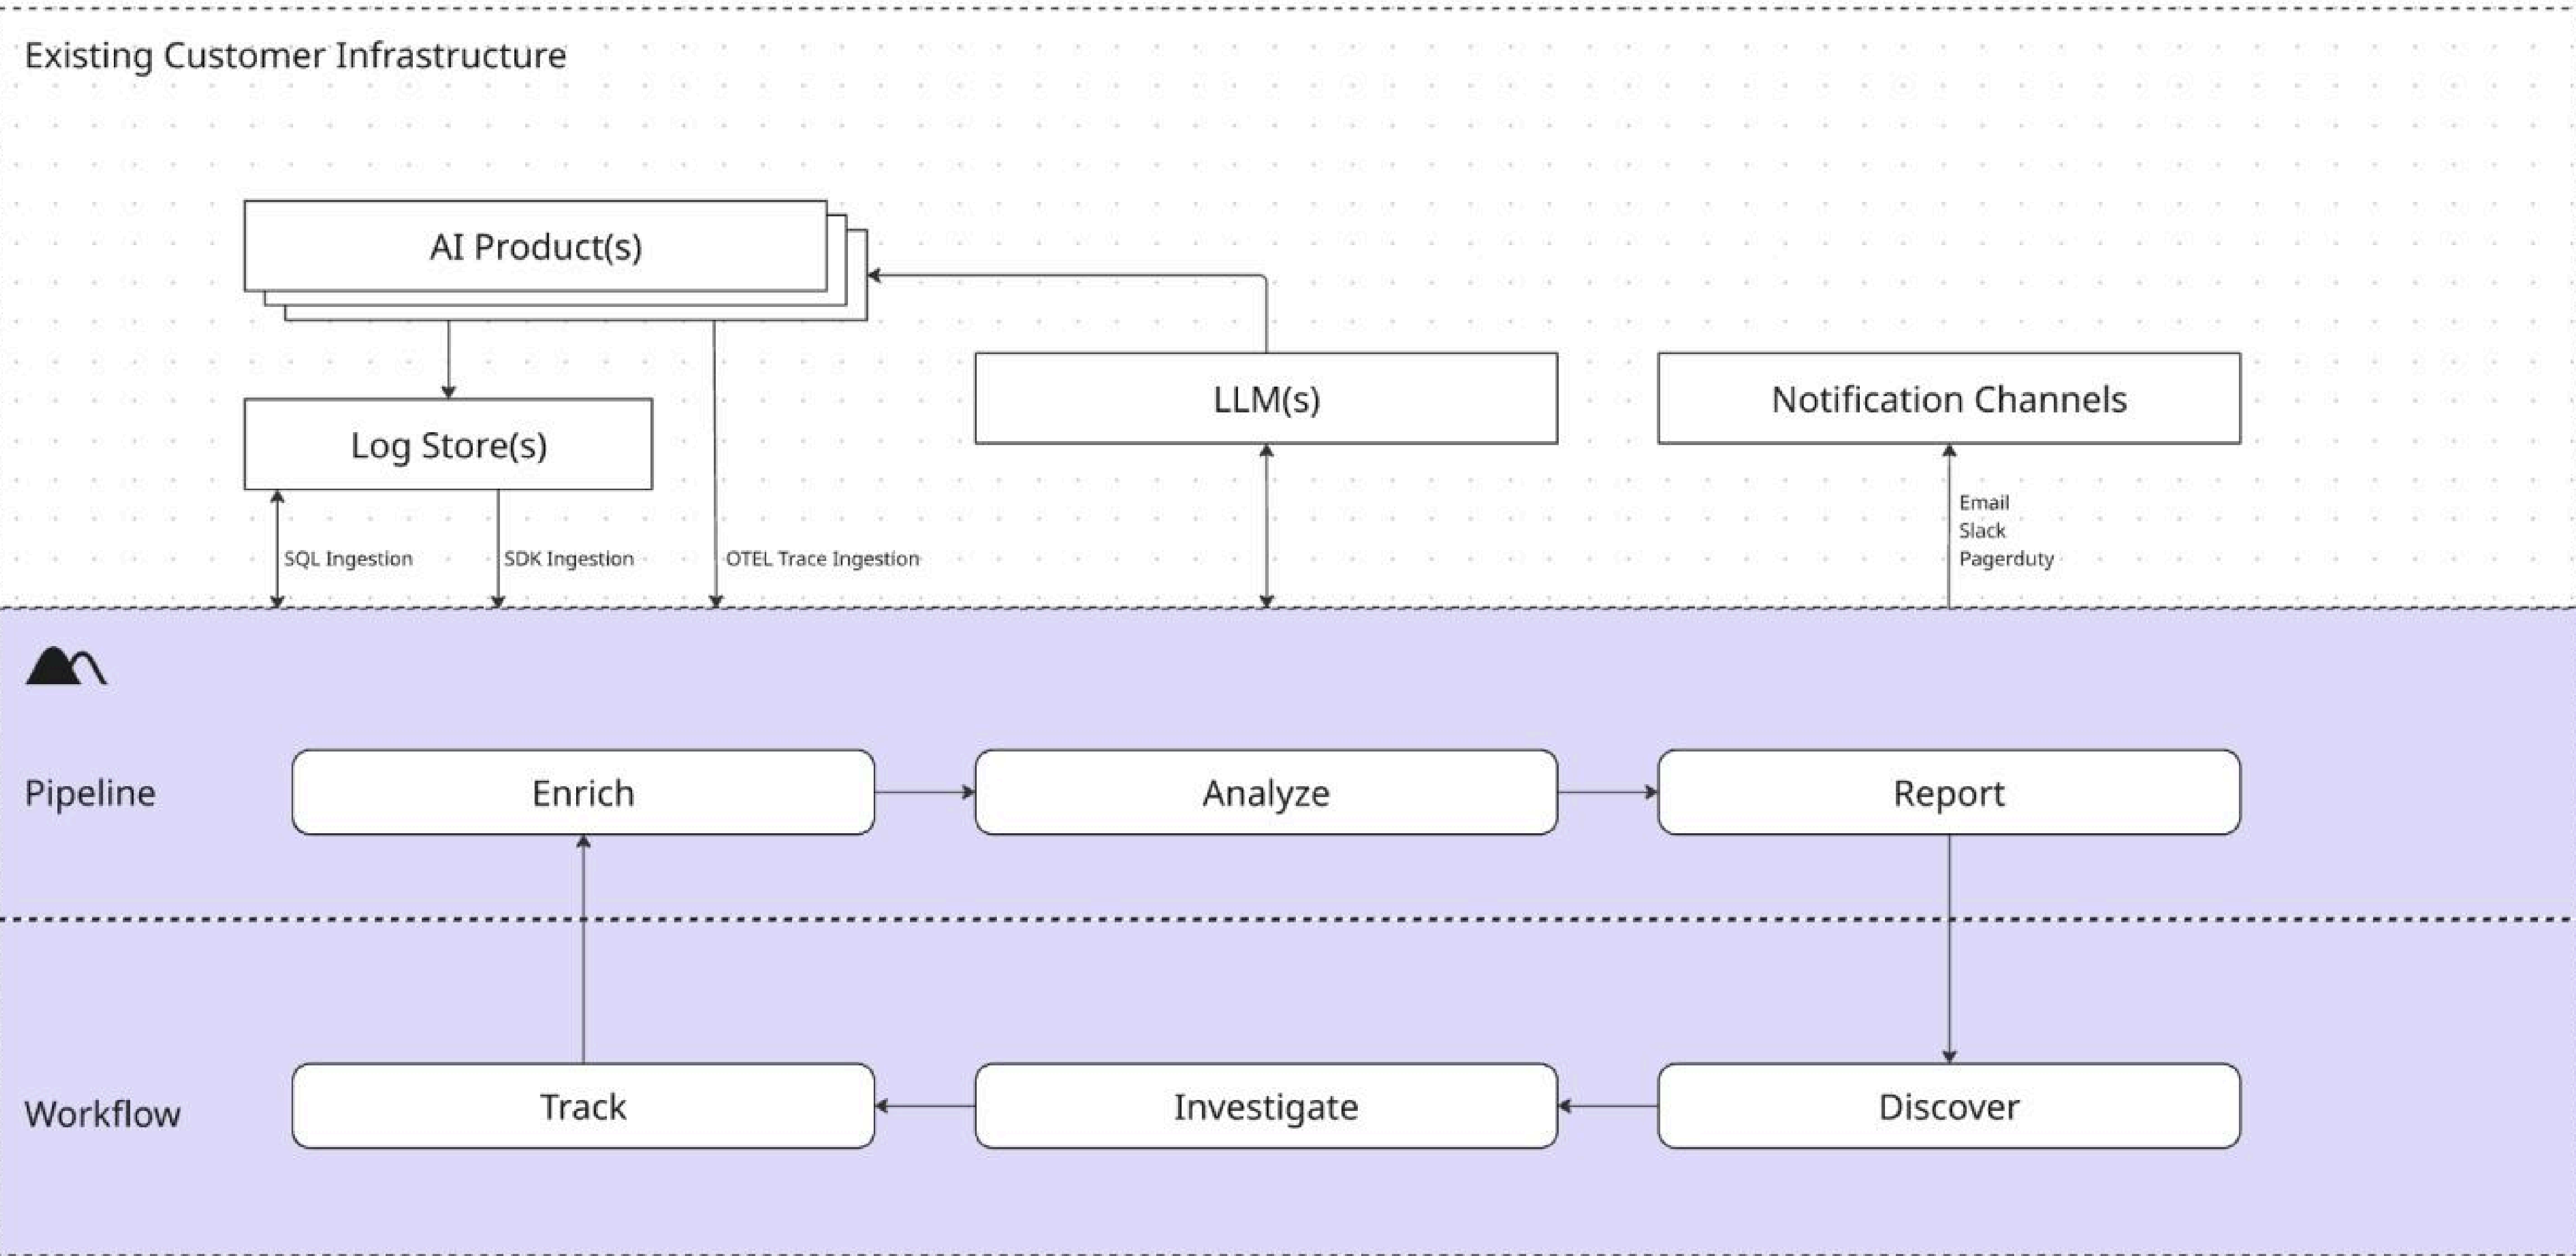
Page 8:      Analytics Workflow

Page 9:      Complete The AI Feedback Cycle



# Product

Distributional’s adaptive analytics product has three components. The enterprise platform enables free deployment of the full service in your environment of choice, seamlessly integrates with your existing stack, and includes enterprise grade administration features to keep data fully secure and managed. The data pipeline automatically enriches production AI logs, analyzes these enriched Logs to discover behavioral signals, and publishes insights from these signals to your channels of choice.





DBNL is designed with compatibility, insights, automation, scale, and collaboration in mind.

- Compatibility: Fits your environment, stack, security, privacy, and data needs with minimal ongoing overhead and maintenance.
- Insights: Built to generate regularly compelling insights so you can fix, improve, and scale your AI products with confidence.
- Automation: Does the repetitive data analysis tasks for you, so that instead you can spend your time applying these insights to improve, fix, or scale your AI product.
- Scale: Runs daily batch analysis on all production AI logs as you scale your product so you don't miss any relevant usage patterns.
- Collaboration: Designed to organize behavioral signals on your product in a single place so you can make collaborative decisions to move the AI product forward.

Let's dive deeper on each of the components of the product and how they adhere to these principles to produce a robust analytics experience.

# Enterprise Platform

Distributional's enterprise platform enables you to deploy, manage, scale, and integrate our product seamlessly with your stack.

## DEPLOYMENT

The full DBNL service is available for free with an open distribution model to facilitate adoption by any team that needs AI product analytics. It is installable rather than cloud SaaS so data stays in your environment. This approach ensures your data stays private and secure.

DBNL has three options for deployment. The [sandbox](#) bundles all components and dependencies into a single docker container that can run on a laptop. It is a scaled down version of our service that is quick to run for a proof of concept or pilot. The full production grade DBNL service can be deployed as a [Helm chart](#) that uses existing infrastructure that you've provisioned in your environment. The full production grade Distributional service can also be deployed as a [Terraform module](#) on infrastructure provisioned by the module alongside the service in AWS, GCP, or Azure.





## INTEGRATION

DBNL has few dependencies, runs as a fully contained service, connects with your existing stack, and is designed to produce analysis on your data wherever it is rather than force you to adopt an entirely new data workflow.

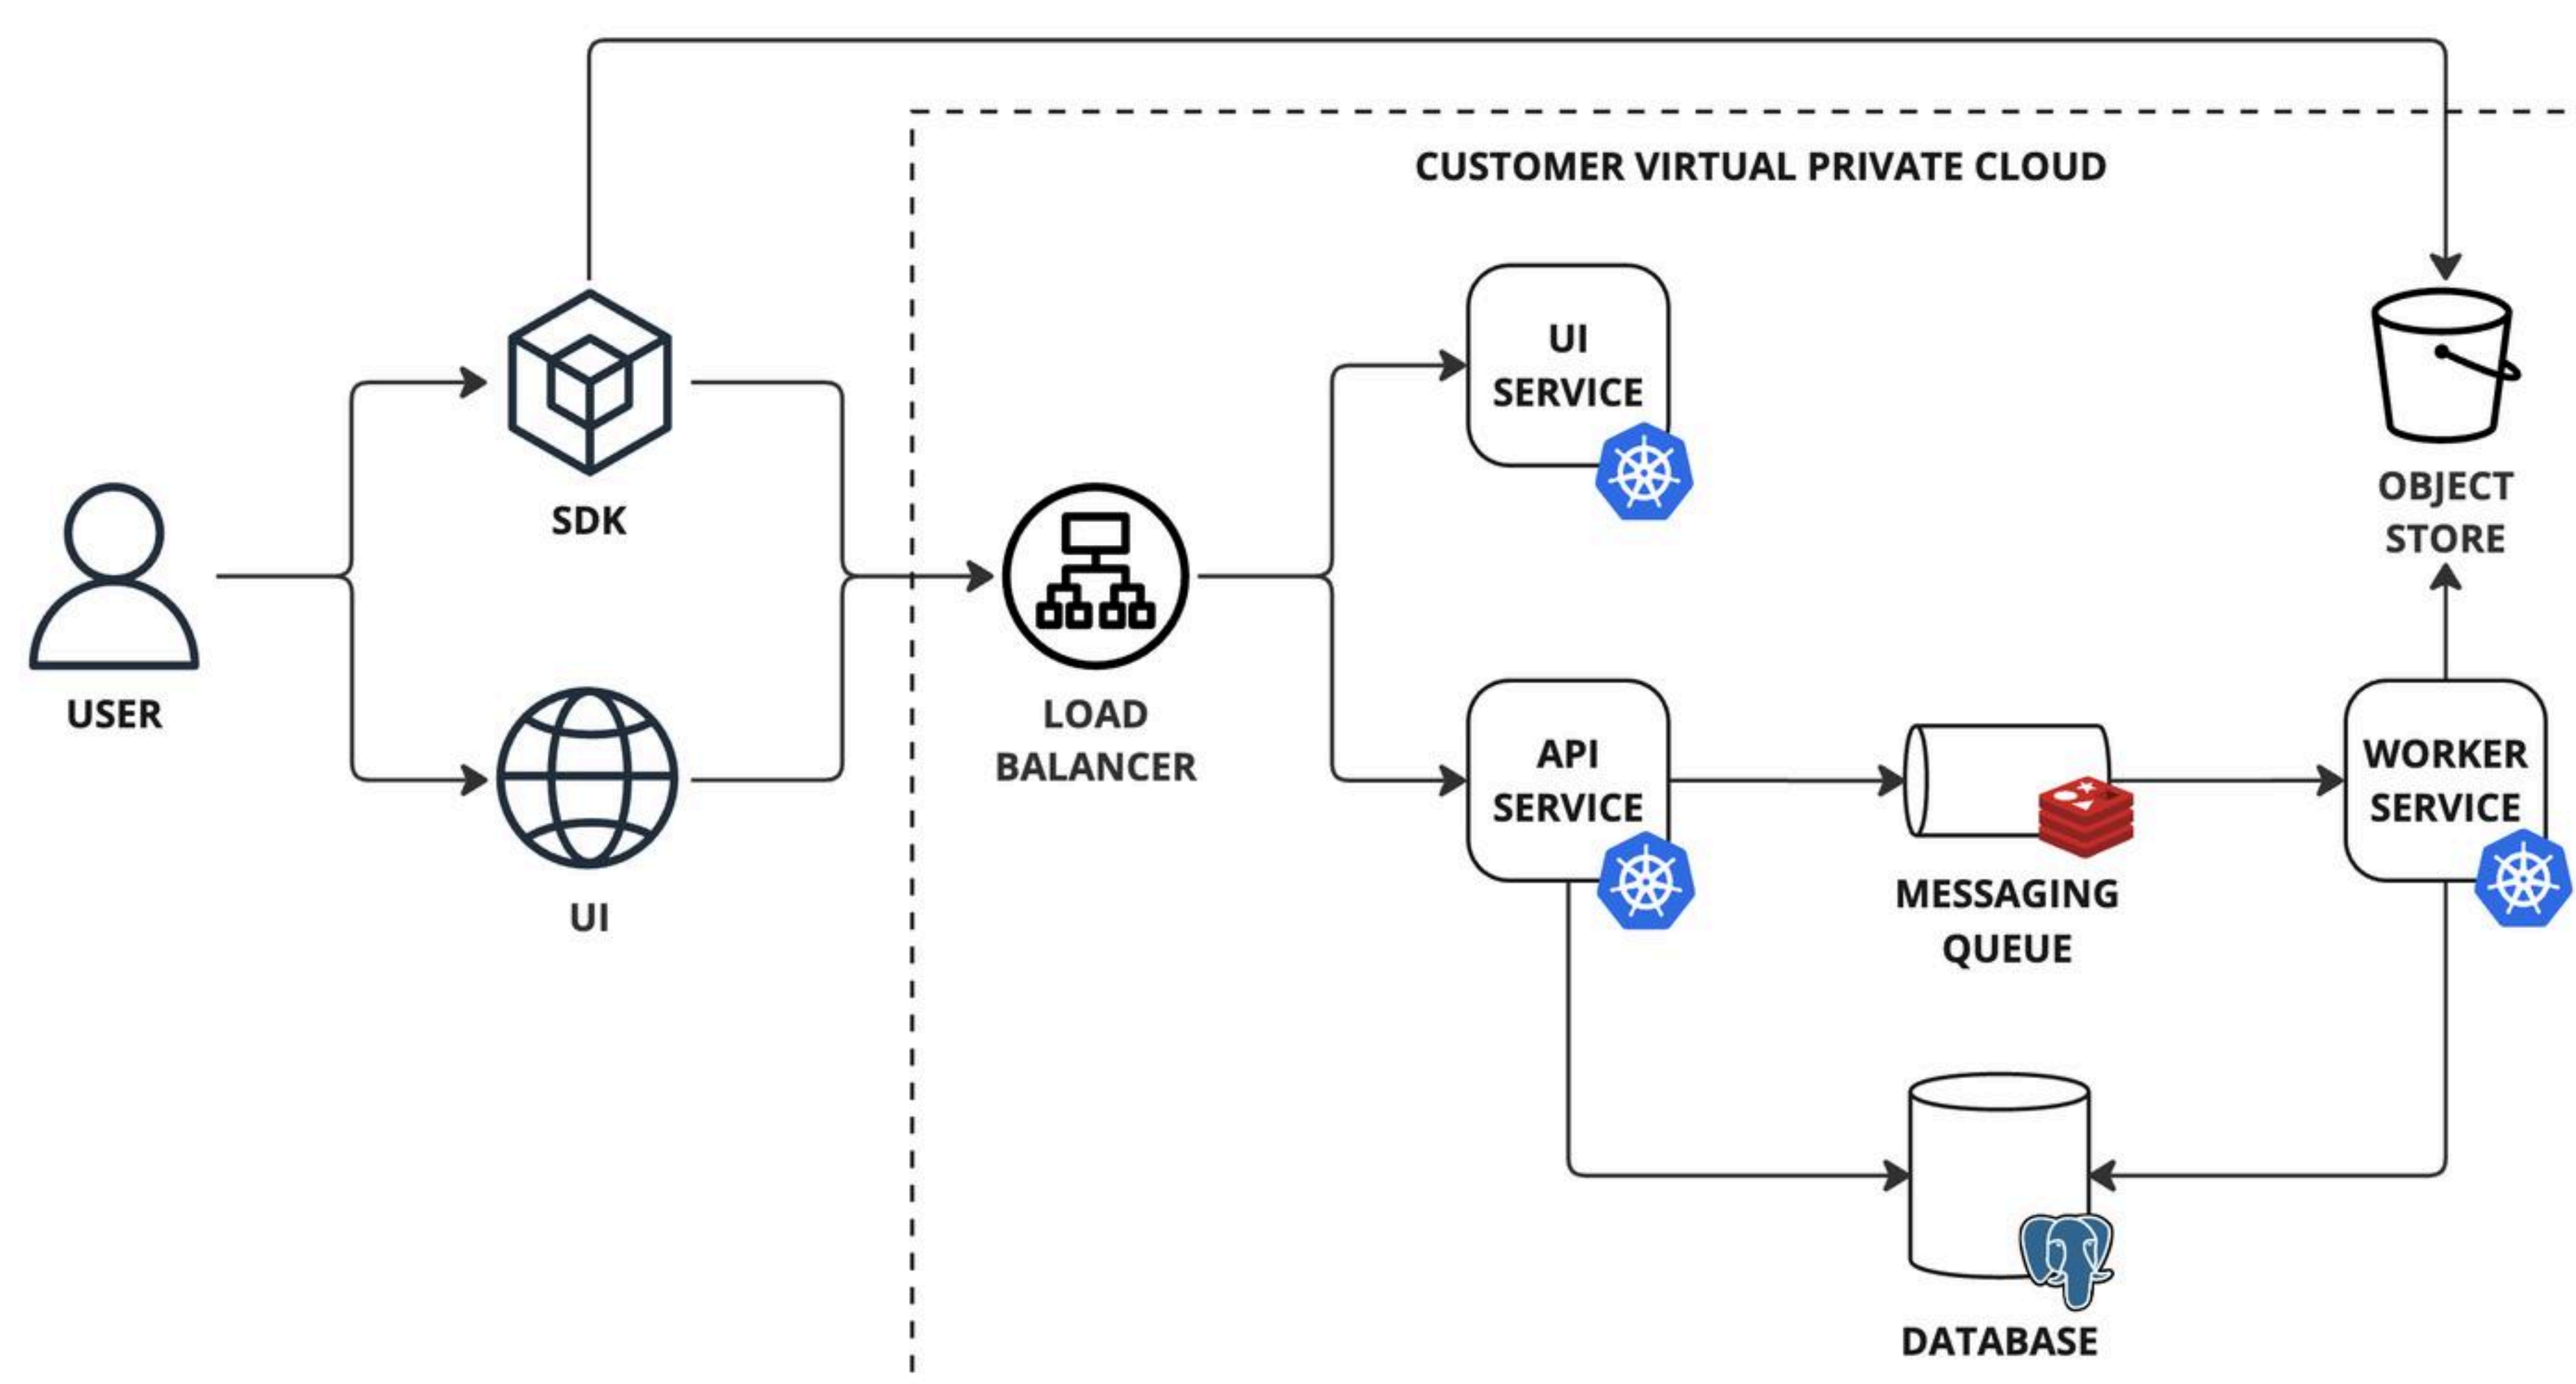
Distributional's [Data Connections](#) includes support for three approaches for DBNL to ingest production AI logs. [OTEL Trace Ingestion](#) publishes Open Telemetry traces directly to DBNL as the product runs. [SDK Log Ingestion](#) pushes data manually or as part of a daily orchestration job using the Python SDK. [SQL Integration Ingestion](#) pulls data from a SQL table into DBNL on a schedule.

Distributional's [Model Connections](#) are how DBNL interfaces with the LLMs required for the DBNL [Data Pipeline](#) to compute LLM-as-judge metrics, perform unsupervised analytics, and translate behavioral signals into human readable insights. DBNL supports externally managed LLM endpoints, cloud managed LLM services that are part of your VPC, or locally managed LLMs that are deployed on a cluster. This includes AWS, Azure, Google, and NVIDIA options, as well as any LLM provider compatible with the OpenAI API convention.

Distributional's [Notification Connections](#) empower you to connect your preferred notification tools so you're informed when there are new behavioral signals or meaningful changes to expected usage patterns. Distributional supports Slack and email notifications with more channels coming soon.

## ARCHITECTURE

Distributional's [architecture](#) consists of a set of services packaged as Docker images and a set of standard infrastructure components that are deployed into your infrastructure (e.g., a VPC in AWS or on premises on your own infrastructure). The architecture is designed to be scalable, modular, and self contained. It does not require an external connection to hosted Distributional services to operate.





## MANAGEMENT

Distributional's full service includes administration features to keep data private, secure, and organized to meet strict enterprise standards.

DBNL [uses](#) personal access tokens for API authentication, OpenID connect or OIDC for user authentication, and configuration options depending on OIDC provider. DBNL also includes [namespaces and user permissions](#) that can be centrally managed by organization admins from an admin dashboard. And [networking](#) requirements are designed to fit best practices for running these services in your environment.

These features ensure data stays secure and private as you scale the number of AI products and AI product teams that use Distributional.

# Data Pipeline

Distributional's Data Pipeline automatically enriches production AI logs, analyzes these enriched Logs to uncover behavioral signals, and publishes signals to various channels to empower efficient, effective product decision making.

## ENRICH PRODUCTION AI LOGS

The first step in the pipeline is to [augment](#) raw, unstructured production AI logs into a variety of [Metrics](#) that can be parsed, clustered, analyzed, correlated, and judged. The pipeline automatically computes these properties off of your production AI logs. If you have metrics you already prefer from development, these can easily be added to give you a richer picture of behavior against performance expectations. The goal is to create as rich, deep, and broad a picture of behavior as possible so you don't miss potential interesting signals.

## ANALYZE ENRICHED LOGS TO UNCOVER BEHAVIORAL SIGNALS

This pipeline automatically runs daily analysis of these enriched AI product Logs to tease out interesting [signals](#) on AI product behavior. This pipeline applies a variety of techniques for this analysis. The pipeline classifies each trace according to topics that our pipeline has learned and defined for your team. It clusters distributions of these properties according to a variety of conventions, and correlates these clusters to





identify any interesting, unexpected, or anomalous connection between properties. It runs an efficient LLM-as-judge on inputs and responses, and provides insight to any that violate a previously defined threshold or have drifted from prior behavior. And the pipeline applies statistical tests to detect any change in the distributions of these metrics or properties.

### PUBLISH SIGNALS TO YOUR DASHBOARDS AND CHANNELS

The DBNL Data Pipeline translates behavioral signals into human readable [Insights](#) that are published in relevant [Dashboards](#) and [channels](#) with supporting evidence that is linked in the [Explorer](#) and [Logs](#) tabs. This approach is designed to facilitate rapid root cause analysis and data-driven AI product decision making.

## Analytics Workflow

DBNL automates analysis of enriched production AI Logs to surface interesting signals of AI product behavior—the interrelation of inputs, prompts, context, tools, model, and response. These signals drive a DBNL-powered analytics [workflow](#) that drives improvements, fixes, or changes to AI products. This workflow enables AI product teams to regularly discover new insights on AI product usage, investigate these insights to quickly make decisions on their products, and track any changes to confirm the result. This workflow gives you a better understanding of your AI product as it evolves over time.

### DISCOVER BEHAVIORAL SIGNALS IN THE CONTEXT OF STANDARD USAGE PATTERNS

The starting point for your interaction with DBNL is to discover new signals on AI product behavior. With the first week of data, DBNL learns a standard set of topics and canonical pathways a user takes through your AI product—including all of the data sources, LLM requests, and tool calls that define intermediate steps in this flow. Going forward, the platform classifies future daily usage according to this rubric, so you always know how yesterday's behavior compares to this standard.

Alongside this daily assessment of AI product status, DBNL is uncovering behavioral [signals](#) hidden in your AI production logs. These signals include increased propensity for a given topic, anomalous response quality as assessed by LLM as judge, changes in the distributions of Metrics that define AI product performance, or interesting correlations between clusters of multiple properties. DBNL provides these signals—as well as any alerts on previously defined threshold violations—to the user in a daily report that serves as the starting point for deeper understanding of AI product behavior.





## INVESTIGATE SIGNALS WITH CURATED EVIDENCE

DBNL links each signal to a curated set of evidence to the claim. Although the signals themselves are distilled sentences explaining the discovery, the evidence includes detailed historical [analysis](#) and a specific subset of [Logs](#) from which the signal was derived. DBNL also places these signals in the context of topics and flows that define canonical usage patterns.

This approach is particularly important as AI product usage scales. Without unsupervised analysis that points you in the right direction and surfaces the relevant traces for deeper analysis, your team is stuck randomly sampling or reviewing only human feedback. The goal is to take this evidence and take action to improve or debug your product, so these signals drive product improvement over time.

## TRACK SIGNALS TO VALIDATE PRODUCT BEHAVIOR

As you investigate these signals and take action to improve your product, DBNL makes it easy to identify or apply thresholds to these signals as new tracked [Metrics](#) and [Segments](#) that define your AI product behavior. With a click of a button, you can add these Metrics and Segments to DBNL, and they'll appear as alerts via notification channels and as tracked Metrics or Segments on the dashboard going forward. Each time you indicate your preference for tracking one of these signals, DBNL reinforces these preferences by tailoring future analysis.

# Complete The AI Feedback Cycle

When you build your AI product, you can look at most traces to gauge performance. When you productionalize, monitoring helps you feel comfortable about aggregated, basic performance metrics tied to speed, cost, and quality. But all of this breaks when you scale your AI product. It becomes impossible to look at each trace and random sampling usually misses something. It also becomes hard to understand usage of your product if you are just looking at real-time monitoring signals.

As you use the Distributional platform, you develop a richer, more complete definition of AI product behavior that evolves with shifting usage patterns over time. And, most importantly, it aligns your team on their shared definition of performance. This completes the AI product feedback cycle, which empowers more efficient, effective, and scalable continuous improvement of these products.

Learn more:

- [docs.dbnl.com](https://docs.dbnl.com)
- [distributional.com/blog](https://distributional.com/blog)



[DISTRIBUTIONAL]

# About Distributional

Distributional's platform is designed to give customers complete control and optionality over their data and integrations. The platform can adapt to specific security, privacy, and regulatory needs, while still providing the quality and insights necessary to test AI applications.

If you'd like to learn more, reach out to the Distributional team. We're here to help!

Learn more at [distributional.com](https://distributional.com)

**Follow us on:**

- [LinkedIn](#)
- [YouTube](#)

