dagster +

A data orchestrator developers love

Eliminate distractions and speed up your development work with data engineering best practices

Dagster is a data orchestrator built with developers in mind. It has been crafted to provide the best possible developer experience across the entire development lifecycle. It makes your pipelines easy to design, run, test, observe, and debug.

The engine is designed to be fault-tolerant, eliminating those late-night on-call alerts.

The Dagster UI is built with the practitioner in mind. It gives you:

- A bird's-eye view of the entire data platform,
- The surgical precision of each asset status and run log
- And everything in between.

Dagster allows other teams to work autonomously to observe the state of pipelines, explore the data catalog, and locate the assets critical to their work. This, in turn, frees you up from responding to requests for help so you can focus on the most value-added work.

Dagster core concepts

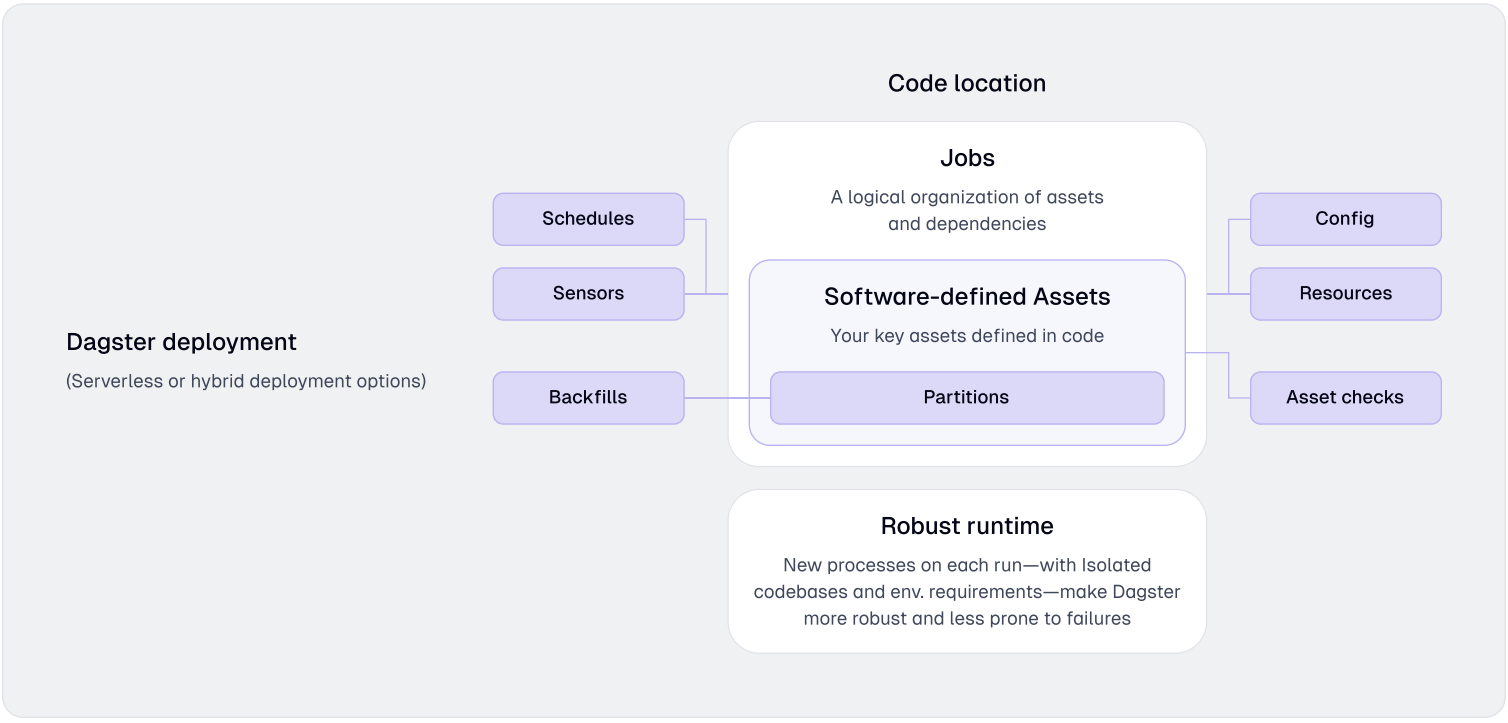

Dagster provides a logical framework for building data pipelines which encourages data engineering best practices, a DRY approach, and observability of all key components of the system.

dagster +

A data orchestrator developers love

What makes Dagster different

An Asset-oriented framework

By building with Software-defined Assets, you can manage the complexity in your data environment, write reusable, low-maintenance code, and deliver a logical data platform with rich metadata.

Built for performance

New processes on each run make Dagster more robust and less prone to failures. Isolated codebases and environmental requirements make Dagster more flexible and stable.

Full SDLC support

With carefully crafted Pythonic APIs, in-built testability, CI/CD support. . With Dagster, you can adopt software development best practices in data engineering.

Robust, flexible scheduling

Thanks to it's asset-oriented approach, Dagster provides both schedule- and sensor-based triggers for runs. We also offer rules-based auto-materialization.

Built-in testing and checks

Dagster is designed to support code tests, unit tests, as well as asset checks for data quality.

A true collaboration platform

Dagster fosters collaboration in the building of pipelines, the support for heterogeneous technologies, and the sharing of data products with stakeholders.

A unified control plane

Once your pipeline is built, you will come to love Dagster's UI. With detailed logging, run timelines, partition tracking, data cataloging, operational and cost observability features... Dagster becomes your single source of truth for your data operations and data assets.

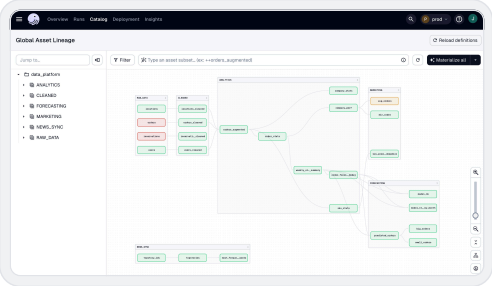

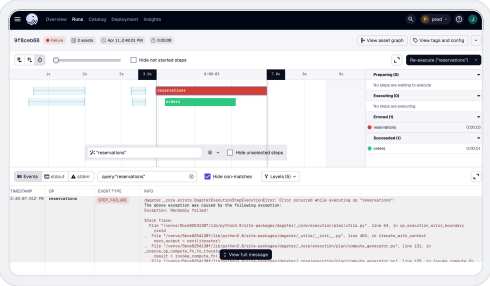

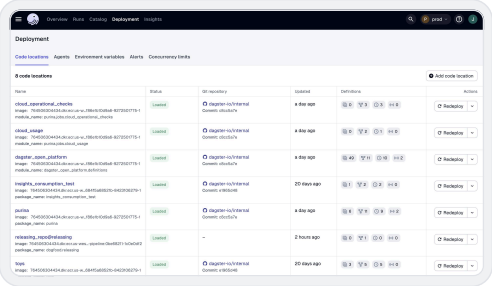

Try Dagster+ today

Go open source

You can
`pip install dagster`
today, and start building

Go cloud

Start a free trial at
dagster.cloud/signup
and get up and running in minutes

Learn Dagster

Learn the fundamentals with Dagster University

courses.dagster.io

Join the community

Join us on Slack, Github discussions, and get support from Dagster team members