

Build with Palantir Foundry

PLATFORM MODULES

01 – Data Integration

02 – Digital Twin

03 – BI & Analytics

04 – Model Integration


05 – Application Building

OVERVIEW


Palantir Foundry collapses the trade-off between the “build from scratch” approach of hyperscalers and the “buy pre-packaged” approach of typical SaaS.

Originally developed for Palantir engineers to deliver mission-critical outcomes on rapid timelines, Foundry offers a suite of software primitives to help power your architecture — from data integration to application building:

- ✔ **Connect data, analytics, and operations** to drive better decision-making across the enterprise.
- ✔ **Empower every user archetype** to transform data, perform analysis, build models, and build applications — freeing your technical resources to focus on building the most valuable things.
- ✔ **Choose what you need** — leveraging modular, interoperable building blocks — to accelerate your architecture roadmap and maximize value from existing investments.



Software Modules



01

[Data Integration](#)

Accelerate your pipelining with built-in automation, while safeguarding data integrity across collaboration.

Accelerate pipeline development speed by eliminating redundant, time-consuming aspects of data engineering with integrated automations (CI/CD, dependency management, autoscaling compute, etc.).

Unlock resources by enabling everyone to build data pipelines without using code.

Decrease infrastructure costs by reducing data duplication and compute usage. Pipeline versioning and row/column level permissions eliminate the need to manage multiple environments or fork pipelines.

02

[Digital Twin](#)

Beyond a data model, Palantir Foundry maps actions and relationships to form an interactive model of the enterprise — complete with write-back, scenario-modeling, and integrated security.

Eliminate data fragmentation and siloes. The digital twin is a common operating layer with interchangeable backends, multimodal storage, and various consumption methods.

Get next-gen master data management. More than just a data model, the highly dynamic digital twin reflects the evolution of relationships, actions, and dependencies between objects.

Synchronize data changes seamlessly. Direct, auditable write-back to the underlying data asset maintains the digital twin as an authoritative, actionable source of truth.

Run what-if analyses and compound simulations across your entire business by leveraging Foundry's native “scenario-aware” functionality.

Software Modules

03

[BI & Analytics](#)

Empower users of all technical abilities to manipulate, analyze and act on data via intuitive, interactive applications.

Engage all users on heterogeneous data types with both point-and-click and code-based tools for table-based, top-down visual, geospatial, and time-series analysis.

Build and deploy custom dashboards in hours, not months. Transform ad-hoc analysis into production-ready dashboards with templating and parameterization.

Go beyond conventional read-only paradigms. Connect third-party BI, visualization and analytical apps to Foundry via standard APIs and interfaces to capture decisions made on all operational fronts and write them back to the digital twin.

04

[Data Science & Modeling](#)

Enable full-stack data science with an integrated environment to develop, test and operationalize models.

Accelerate time to production. Write production-ready model code from the first line. Foundry's modeling framework eliminates notebook refactors, irreproducibility, and orphan ML models.

Operationalize models faster. Users can deploy new models directly into applications with a few clicks.

Keep models fresh. Live connectivity with production data pipelines means model inputs are always up-to-date and models are trained (and re-trained) on the freshest available data. Foundry exposes model predictions and features and factors influencing outputs to build confidence.

Iterate rapidly. Foundry's integrated feedback loop creates an automated process for model drift detection, tuning, and re-training. Foundry also captures operator feedback on deployed models so data science teams can rapidly improve them.

05

[Application Building](#)

Give teams the ability to create read-write applications and run simulations as needed.

Empower technical and non-technical users to create sophisticated, custom applications via low-code or pro-code tooling.

Go from zero to production application in hours. Deploy use-case templates on your data to accelerate application development.

Build complex workflows without complex integrations. Embed business rules, consume data, and enable writeback without needing engineering resources to translate logic into SQL queries.

Simplify model consumption. Natively integrate with models deployed via Foundry, powering sophisticated simulation and 'what-if' analysis in a few clicks.

Key Advantages

ADVANTAGE	DESCRIPTION
<p>✔ Accelerate time to value</p>	<p>Foundry's tooling empowers any size organization to harmonize data in hours and build applications in days.</p>
<p>✔ Reduce total cost of ownership</p>	<p>Gain transparency and control over compute costs. Empower your teams to accelerate your speed of development & value creation.</p>
<p>✔ Minimize integration time and effort</p>	<p>Foundry applications are built to interoperate in complex architectures, helping to reduce integration time and cost.</p>
<p>✔ Scale every dimension of your business</p>	<p>Leverage dynamic autoscaling and deliver future use cases on your existing data foundation, solving problems quickly as your business evolves.</p>
<p>✔ Focus on building</p>	<p>Leverage your engineering resources to focus on building what brings the most value to your company – not managing infrastructure, integrations and data access requests.</p>

