Palantir

PALANTIR PLATFORM

SERVICE DEFINITION DOCUMENT

Prepared For: G-Cloud 15 Framework

palantir.com/uk

Foundry AIP Gotham

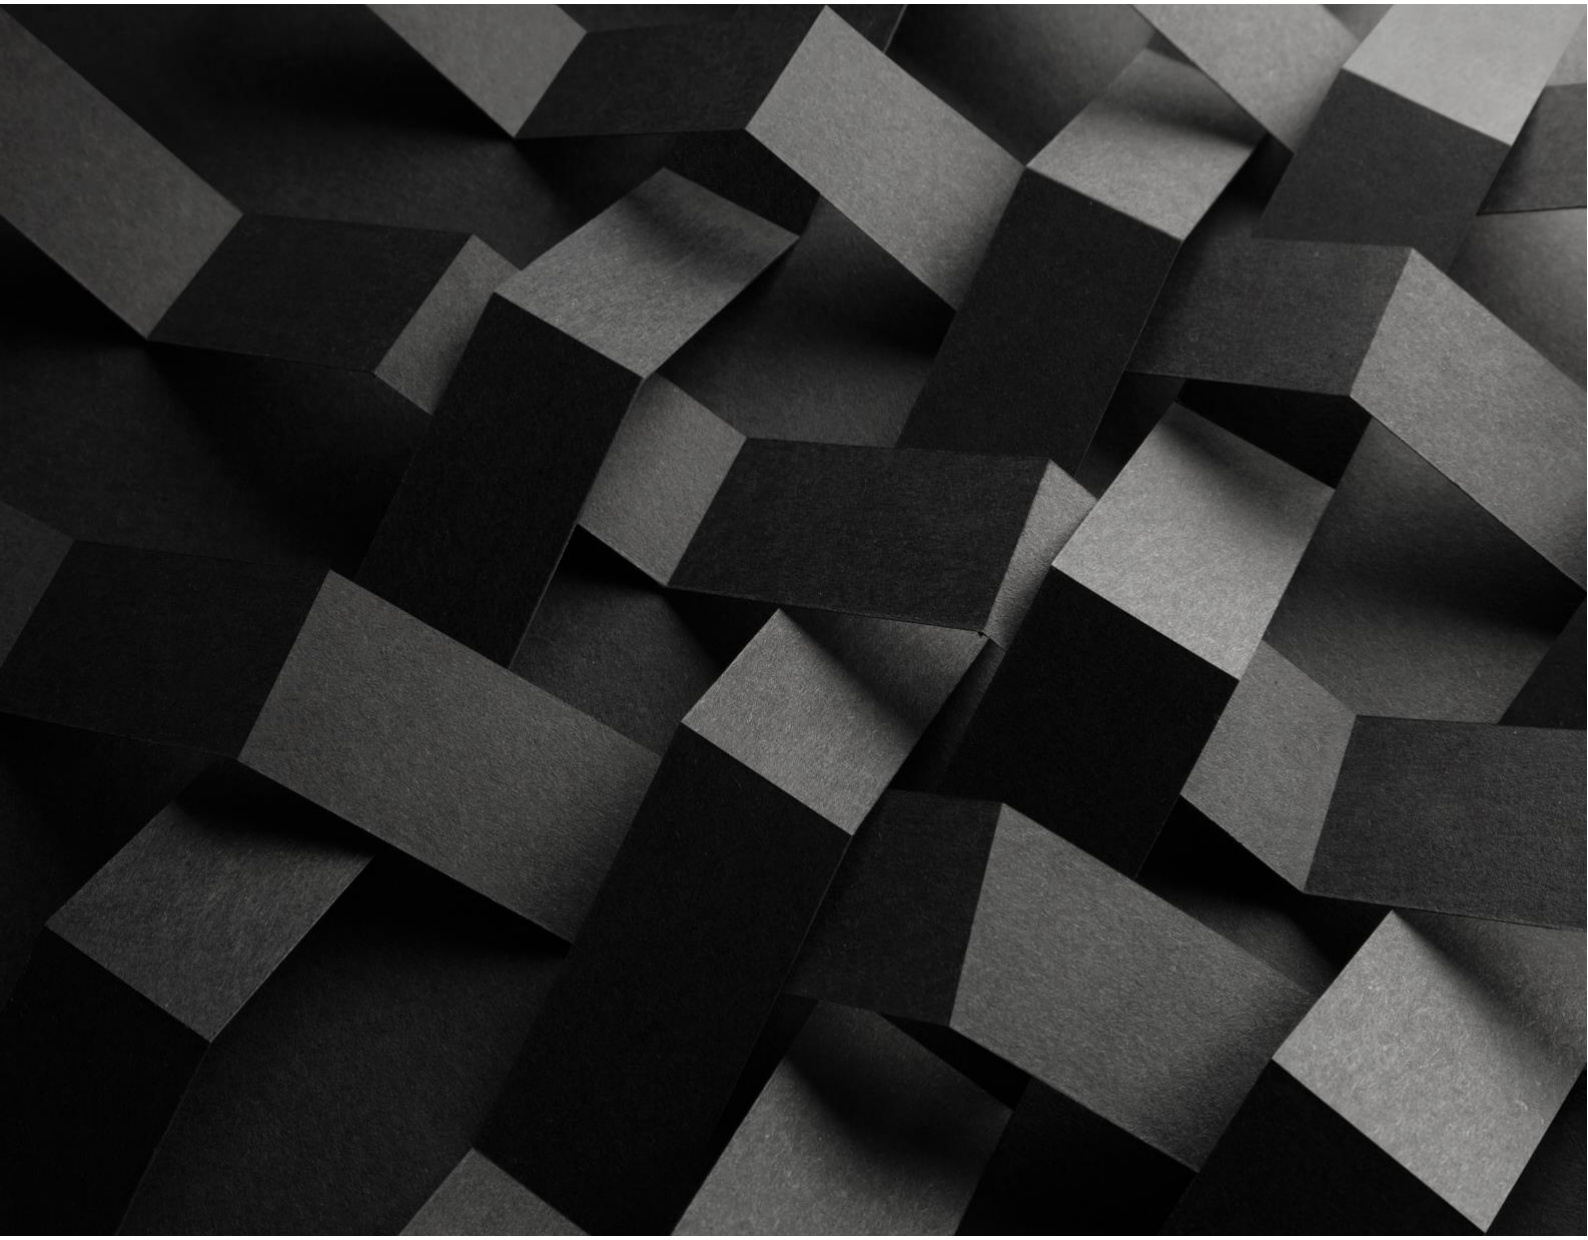

Table of Contents

1.0 PALANTIR PLATFORM: FOUNDRY OVERVIEW

The Palantir Platform offers three core products: Palantir Foundry (Foundry), Palantir Gotham (Gotham) and the Artificial Intelligence Platform (AIP). The following section describes Foundry, which provides foundational data management functionality for the Palantir Platform. Please note that all data shown in screenshots is notional.

Please visit our website (www.palantir.com/uk) and YouTube channel (www.youtube.com/@palantirtech) to access a variety of case studies, technical documentation, product demonstrations and additional information about the Palantir Platform.

1.1 Introduction

Palantir Foundry (“Foundry”) is an enterprise software product for integrating, managing, analysing, modelling and operationalising data at massive scale; empowering organisations to unlock and realise value through data-driven insights, secure automation, and enhanced collaboration across teams.

Foundry is the result of Palantir’s 20+ years of experience enabling data-driven operations in disparate and complex environments. Palantir’s customers include healthcare providers, banks and financial regulators, energy majors, large manufacturing and logistics firms, and law enforcement and defence organisations in multiple countries. In all these sectors, Palantir recognises data as a fundamental business asset and delivers its software to help organisations integrate, secure, govern, and operationalise that asset.

While many enterprise data systems address the challenges of storing and processing large quantities of data, Foundry was built to manage big data complexity. With Foundry, authorised users from across an organisation can intuitively access the data they need, securely collaborate with their peers, and confidently make data-driven decisions. Foundry achieves this by establishing a single source of truth comprised of all available data, brought together from disparate source systems and mapped to a customer-specific data model of familiar business objects. Back-end data pipelines transform and surface data to users via a built-in suite of highly configurable analytics, reporting, and workflow management applications, all backed by automatic data lineage tracking, granular access controls, and platform-wide auditing.

1.2 Foundry Architecture

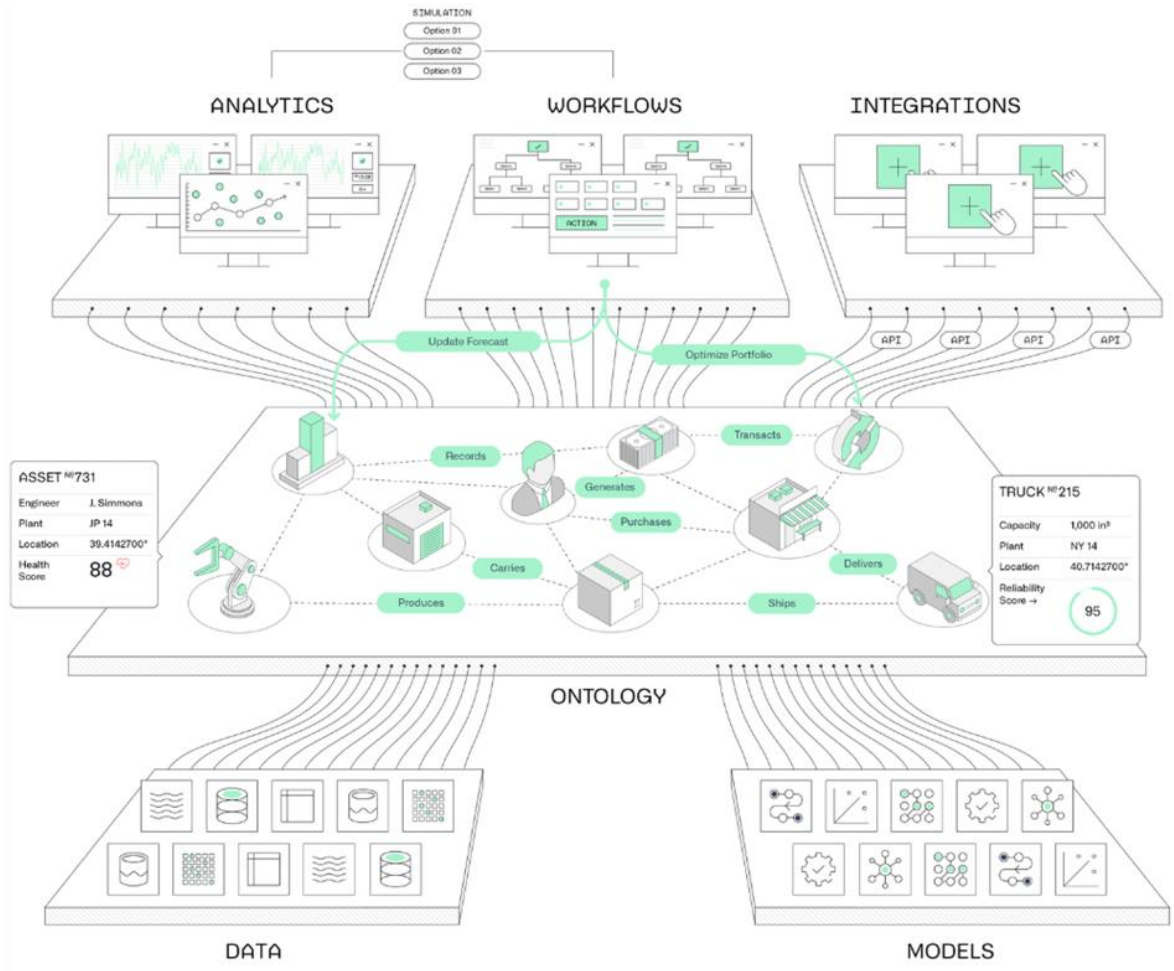

Figure 1: Foundry's conceptual architecture, showing a sample object-based data model for supply chain workflows. The platform unifies data, models, harmonised objects, and flexible user-facing tools.

Data Layer

This layer provides the foundational integrations between Foundry and external data sources. Over 200+ out-of-the-box systems connectors link to databases, data lakes, streaming sources, and file systems, enabling both batch and real-time ingestion from on-premises systems, cloud storage, and SaaS applications.

This layer handles authentication, encryption, and connection management whilst transforming raw data into analysis-ready datasets using Apache Spark and other distributed computing technologies. It offers both code-based and low-code transformation tools through Code Repositories and Pipeline Builder, managing data lineage, versioning, and incremental processing throughout.

Critical for AI/ML workflows, the data layer enables sophisticated feature engineering, automated data preparation, and feature stores ensuring consistent computation across training and inference.

Ontology Layer

The Ontology is a semantic model that represents data as their real-world counterparts—from physical assets like hospitals, equipment, and medical supplies to concepts like operating theatre utilisation levels and patient test results. It effectively builds a digital twin of the organisation, bridging technical data representations and business concepts through both semantic elements (objects, properties, links) and kinetic elements (actions, functions, dynamic security). Machine learning models are registered as first-class objects with versioning and relationships to their input/output data, with model predictions becoming properties of business objects. This enables seamless AI integration into workflows whilst ensuring decisions remain traceable, explainable, and aligned with business logic, supporting use cases of all types across the enterprise.

Application Layer

The Application Layer offers analytical and operational applications accessed through native tools like Quiver, Workshop, and Code Workbooks (all described below), plus custom applications via APIs and the OSDK (Ontology Software Development Kit). Data scientists can develop and deploy models using frameworks like TensorFlow, PyTorch, and scikit-learn through Modelling Objectives and Code Workbooks. For operational users, such as clinicians or police officers, Workshop enables AI-powered operational applications with real-time model-driven decisions. Foundry supports AutoML for data scientists, MLOps monitoring dashboards, and integration of large language models and generative AI (see Palantir AIP below), with both batch and real-time inference endpoints synchronised to the Ontology.

Security and Governance

Enterprise-grade security and compliance permeates all layers through fine-grained, role/classification/attribute/purpose-based access controls, comprehensive audit logging, and data lineage tracking. For AI/ML workflows, it extends governance to model lineage, training data versioning, bias detection, and fairness monitoring, tracking data provenance, deployment approvals, and prediction audit trails. Foundry ensures responsible AI through explainability tools, performance monitoring, and drift detection. It enables privacy-preserving machine learning via federated learning and differential privacy, supporting powerful AI capabilities whilst maintaining strict data protection standards.

1.3 Foundry Key Applications

Below is a non-exhaustive list of Foundry's main applications – more information can be found in our published product documentation at https://www.palantir.com/docs/foundry.

Pipeline Builder offers visual, low-code data transformation through intuitive drag-and-drop interfaces. Users connect pre-built transformation blocks for filtering, joining, and aggregating data, with automatic optimised Spark code generation and dependency management. It provides data quality monitoring, error handling, and testing capabilities. Pipeline Builder enables developers to create enterprise-grade pipelines rapidly whilst maintaining performance and reliability standards, bridging the gap between business users and data engineers for production data transformation workflows.

Object Explorer provides the primary interface for searching, filtering, and exploring Ontology objects through an intuitive point-and-click experience. Users perform queries ranging from simple keyword searches to complex property filters and relationship traversals, visualising results through exploration views, tables, or charts. Object Explorer supports advanced search syntax with logical operators, fuzzy matching, and wildcards, whilst enabling users to save explorations, compare object sets, and execute bulk Actions. The application includes object set management, export capabilities, and integration with other Foundry tools, making it essential for operational users discovering and analysing Ontology data without technical expertise.

Object Views serve as customisable, reusable representations providing central hubs for all information related to individual objects. Available in full and panel form factors, Object Views display biographical data, linked objects, key metrics, and embedded analyses or applications relevant to each object. Built using Workshop modules, they offer flexibility in how objects appear across the platform whilst automatically inheriting permissions from their object types. Object Views integrate property lists, charts, timelines, maps, and embedded Slate or Contour content, delivering comprehensive 360-degree perspectives that support informed decision-making throughout operational workflows.

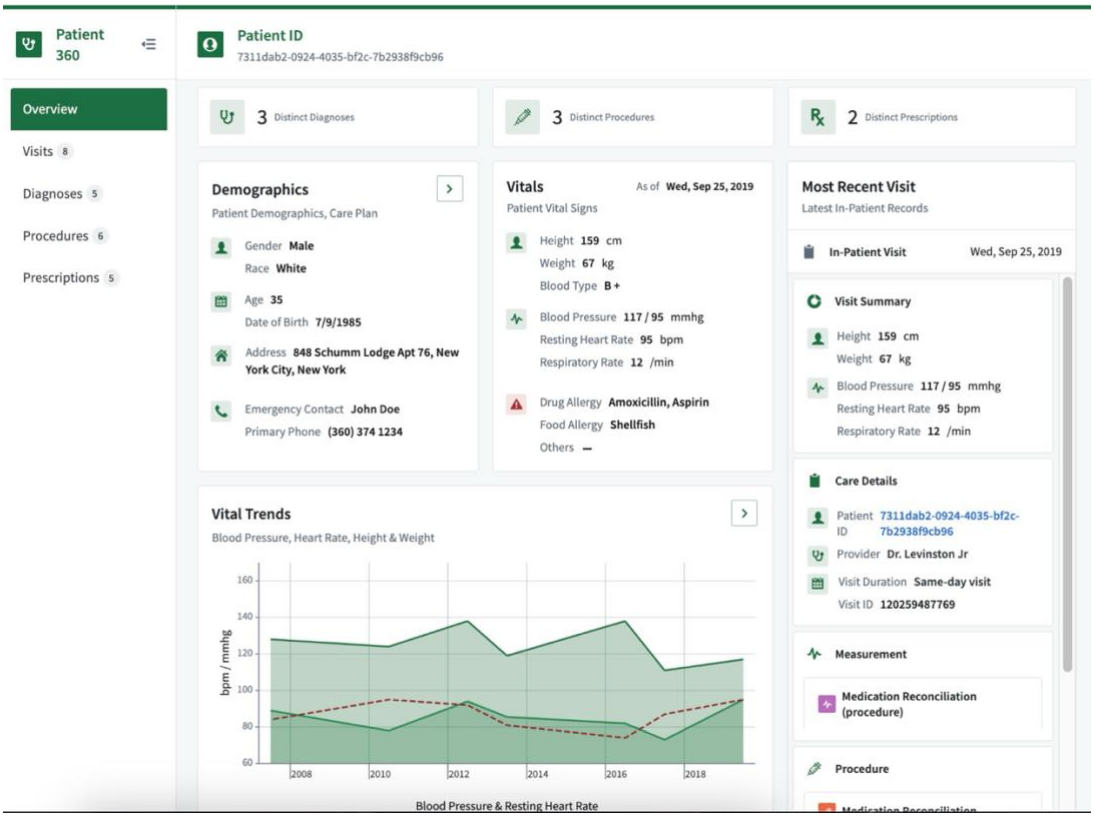

Figure 2: This full sample Object View could serve as an exhaustive resource for all relevant information about a patient, facilitating better-informed healthcare decisions.

Workshop is Foundry's no-code application builder for creating interactive operational applications through drag-and-drop interfaces. Users build custom UIs with widgets like maps, charts, and forms connected to the Ontology, configuring object explorers and orchestrating workflows without code. Workshop applications integrate AI model outputs, enable scenario modelling, and support collaborative decision-making. The platform handles responsive design, permissions, and real-time synchronisation automatically, making it ideal for operational dashboards, case management, and command centres where moderate complexity and maintainability are priorities.

Slate provides a flexible application development environment combining drag-and-drop interfaces with full HTML, CSS, and JavaScript customisation capabilities. Builders create sophisticated operational applications and interactive dashboards with complete design freedom, enabling pixel-perfect implementations and complex business logic. Slate supports both integrated applications leveraging the Foundry ecosystem and public applications accessible to external users without Foundry accounts. With capabilities for custom styling, event-driven workflows, and database integration through Postgres, Slate accommodates highly customised use cases requiring maximum flexibility, though with greater complexity than Workshop.

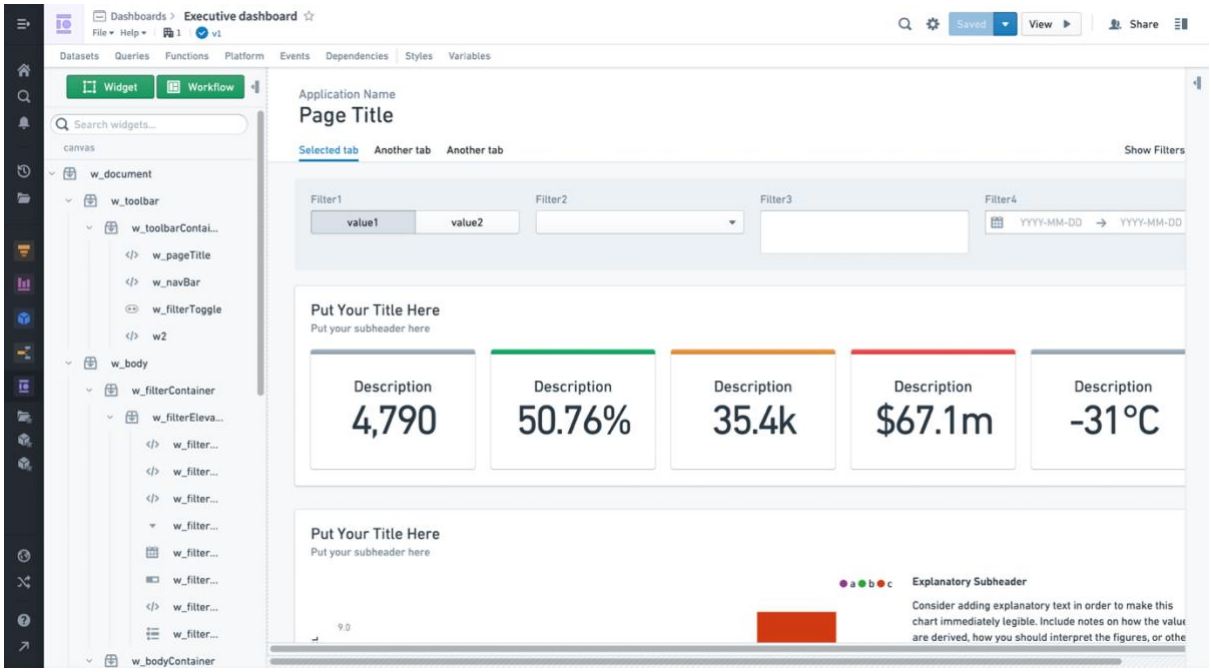

Figure 3: Sample executive dashboard created in Slate

Contour delivers point-and-click analytics for tabular dataset exploration and transformation through a board-based, notebook-style interface. Users apply pre-built transformation blocks for filtering, joining, pivoting, and aggregating data without writing code, whilst creating interactive dashboards with parameterised analyses. Contour generates optimised Spark code automatically and enables conversion of analyses into production pipelines, bridging exploratory work and operational deployment. Supporting scheduled reports, collaborative editing, and extensive visualisation options, Contour serves analysts performing dataset-centric work, data quality validation, and rapid prototyping before productionisation.

Quiver provides advanced time-series analytics and visualisation capabilities for creating sophisticated dashboards and exploratory analysis. It offers extensive chart types from business graphics to geospatial and network visualisations with interactive filtering and drill-down. Quiver blends Ontology objects with raw datasets through SQL and point-and-click interfaces, supporting real-time collaboration.

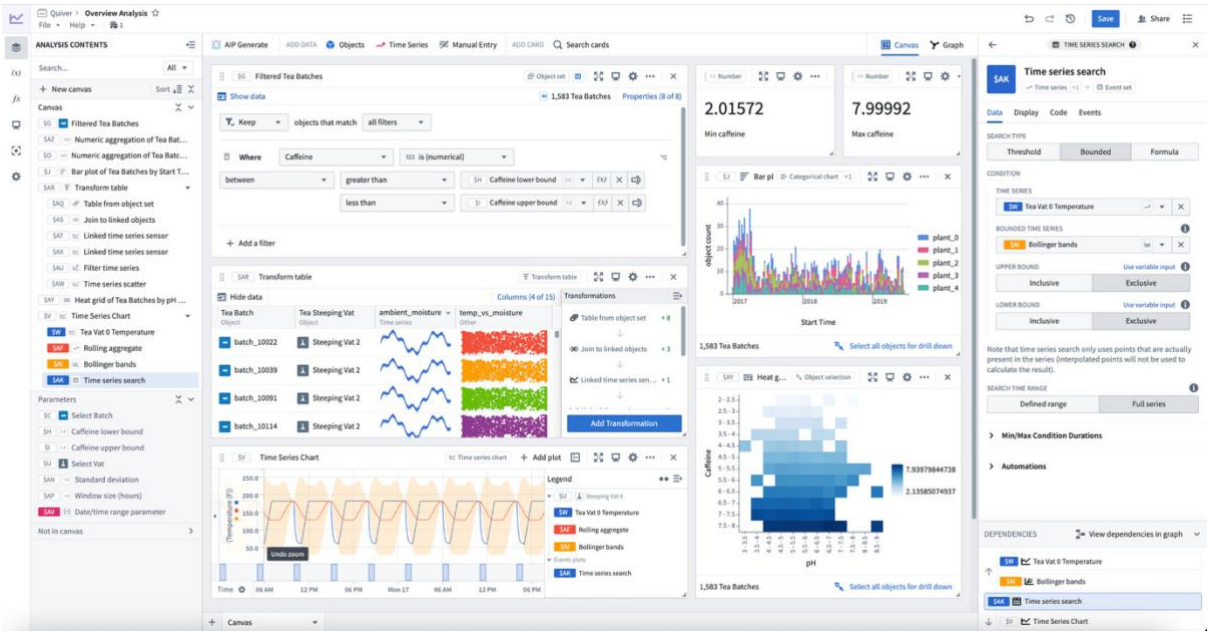

Figure 4: *Sample Quiver analysis, showing how it can be used to create interactive dashboards that allow others to explore and investigate the data in operational workflows.*

Code Workspaces delivers integrated development environments bringing JupyterLab, RStudio Workbench, and VS Code to Foundry with native platform integration. Users access Ontology objects and datasets through familiar IDEs whilst benefiting from Foundry's data security, branching, build scheduling, and resource management. The environment supports Python and R for exploratory analysis, model development, and transform prototyping, with capabilities for creating Dash, Streamlit, and Shiny applications. Code Workspaces enables Git workflow support, customisable compute profiles, and seamless publishing of transforms to production pipelines, serving data scientists performing cell-by-cell analysis, experimentation, and collaborative development on single-node workloads.

Code Repositories provide a complete development environment for building production-grade pipelines and applications. Supporting Python, Java, SQL, and TypeScript with Git integration, it includes intelligent code completion, debugging, and automated testing. Developers create reusable libraries and sophisticated orchestration logic using Foundry's compute infrastructure. The platform enables CI/CD workflows, branch-based development, and code review processes, ensuring enterprise software standards throughout the data pipeline lifecycle for teams building scalable, maintainable production systems.

Modelling Objectives delivers an integrated ML platform managing the complete machine learning lifecycle. It provides AutoML alongside support for TensorFlow, PyTorch, and scikit-learn frameworks, handling experiment tracking, hyperparameter tuning, and model versioning. Modelling Objectives includes model monitoring, performance metrics, and drift detection for production accuracy. Models seamlessly consume Ontology data and publish predictions as object properties, embedding AI capabilities directly within operational workflows whilst maintaining governance and explainability standards for enterprise ML operations.

Actions and Functions provide foundational capabilities enabling data writeback and business logic execution throughout the Ontology. Actions define rule-based workflows for creating, modifying, deleting, and linking objects through automatically generated forms, ensuring secure and governed user interactions with data. Functions implement complex business logic using code, enabling sophisticated data transformations, validations, and calculations that extend beyond standard Ontology operations. These capabilities integrate across Workshop, Object Explorer, and other applications, forming the kinetic layer that transforms Foundry from a passive data platform into an operational system supporting decision execution and process automation.

1.3 Scalability and Interoperability

Foundry was built to grow with the customer's enterprise by targeting the issues that make systems difficult to resize. When capacity must be defined up front, systems either quickly become too small or incur costs for unused capacity.

Scalability on Demand: Foundry dynamically scales infrastructure to accommodate a practically unlimited level of data storage and compute requirements using secure tooling developed on top of open-source software. Rubix is a next-generation cloud infrastructure layer that builds on top of Kubernetes for improved compute performance, cost control, and security. Rubix provides dynamic infrastructure scaling to ensure that Foundry runs the precise amount of compute needed to support actual demand. New hosts spin up in near-real-time and are run as part of auto-scaling groups. Rubix reduces the need for job queueing, leading to an improved user experience in the platform, while ensuring that costs to the customer are associated with increased performance in the platform.

Open Standards and Avoiding Vendor Lock-in: Foundry maintains a firm commitment to open data formats and interoperability across the full spectrum of enterprise systems. All data is stored in its original format and remains accessible through standard interfaces including REST, JDBC, and secure filesystem access. Transformed data is available by default in open formats such as Parquet, ensuring customers maintain complete control and flexibility over their data assets.

Comprehensive Integration Patterns: Foundry provides multiple integration mechanisms across six key domains:

- **Data Interoperability**: Beyond traditional integration, Virtual Tables connect to data at rest in existing systems, enabling flexible, blended approaches. The platform supports both "push" mechanisms (file transfer, JDBC/ODBC drivers) and "pull" methods (RESTful APIs, direct Distributed File System access).
- **Metadata & Semantic Interoperability**: Rich integration with data catalogues, metadata management tools, and master data systems. The Ontology exposes all elements through REST APIs and JSON-driven authoring, enabling bidirectional synchronisation with existing semantic modelling tools.
- **Code & Analytics Interoperability**: Open language support (Python, Java, R, SparkSQL) with bindings for open runtimes (Spark, Flink). Out-of-the-box connectors for Power BI, Tableau, Jupyter, and RStudio enable users to leverage existing

analytical investments whilst benefiting from Foundry's governance and data management.

- **Security Interoperability**: Leverages existing authentication systems via SAML and authorisation systems like Active Directory, supporting role-based, classification-based, and purpose-based permission regimes.

As an integrated ecosystem, Foundry provides the base data management layer, authoring environments for data transformations, a suite of analytical applications, and comprehensive developer frameworks with open APIs for building operational applications, all whilst maintaining complete interoperability with existing and future technology investments.

2.0 PALANTIR PLATFORM: ARTIFICIAL INTELLIGENCE PLATFORM (AIP) OVERVIEW

2.1 Introduction

Palantir's Artificial Intelligence Platform (AIP) powers real-time, AI-driven decision-making in critical commercial and government settings around the world.
From healthcare to manufacturing, organisations depend on AIP to leverage AI safely, securely, and effectively across their enterprises to drive operational results, efficiencies and cost savings.

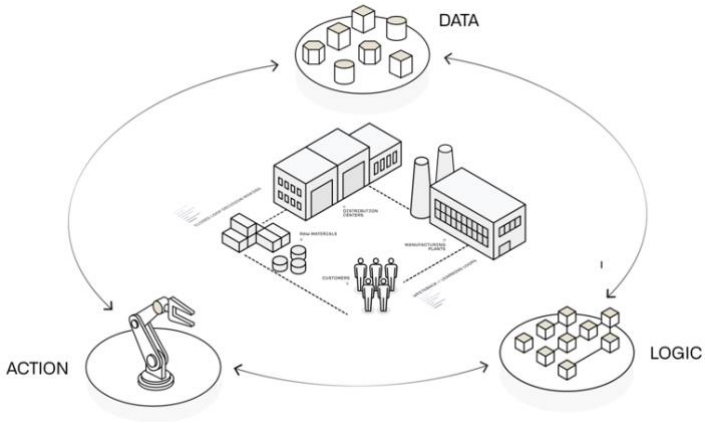

Full AIP documentation is available online (https://www.palantir.com/docs/foundry/aip/overview/) and we offer AIP Bootcamps where customers get hands-on-keyboard with AIP to build tools and achieve real-world outcomes in a matter of hours.

2.2 Palantir AIP Overview

AIP connects AI with your data and operations. AIP is designed to drive automation across operational processes, providing a comprehensive suite of tools that can be used by everyone in an organisation, from developers to frontline users.

AIP's builder tools, like AIP Logic, AIP Agent Studio, and AIP Evals, enable the development of production-ready AI-powered workflows, agents, and functions on top of the Ontology and developer toolchain. Additionally, AIP transforms the application environment by allowing sandboxed, autoscaling applications to integrate generative AI seamlessly within existing security, audit, and resource management frameworks.

Together with Foundry (Palantir's data operations platform) and Apollo (Palantir's mission control for autonomous software deployment), AIP forms an operating system that can deliver a full range of AI-driven products, from large language model (LLM) powered web applications to mobile applications using vision-language models to edge applications that embed localised AI.

AIP offers the following capabilities:

- **Securely access LLMs of your choice.** AIP supports a wide range of LLMs and text embedding models from leading providers, including xAI, OpenAI, Anthropic, Meta, and Google (some subject to certain georestrictions). An end-to-end framework for deploying and managing LLMs with an integrated LLMOps framework is available for capturing feedback to train and fine-tune custom models.

- **Bring your own model (BYOM).** BYOM, also known as "registered models" in the Palantir Platform, is a capability that provides first-class support for customers who want to connect their own LLMs or accounts to use in AIP with all Palantir developer products.

- **Integrate data and build pipelines with natural language.** The platform's extensible, multimodal data connection and integration framework works with enterprise data systems out-of-the-box. Pipeline Builder enables using LLMs to power pipeline-based transforms such as classification, sentiment analysis, summarisation, entity extraction, or translation through point-and-click interfaces.

- **Run on your structured and unstructured data.** Customer data can be made usable by the LLM through Vector Stores and the Ontology, which transform data into LLM-understandable objects, actions and processes into tools for humans and LLM-driven agents.

- **Build human-in-the-loop applications.** Operational use of AI and LLM(s) is powered with interfaces for decision making, feedback, and safe hand-off among AI agents and human operators. A unified orchestration engine coordinates among LLM execution, operational tools, user-driven action, security controls, resource management, auditing, and more.

- **Validation and oversight.** Full-spectrum security and audit controls allow for granular authority over model usage. Integrated human review checkpoints throughout workflows, with guardrails as a first-class concept, ensure that LLMs never execute outside of boundaries. AIP Observability provides tools for monitoring and optimising the performance of agents and applications.

2.3 AIP Applications

AIP has a rapidly growing suite of dedicated tools to enable business and technical users to embed LLMs and other AI/ML models into their work. These include:

- **AIP Agent Studio:** AIP Agent Studio enables you to create interactive assistants that can leverage enterprise-specific data in the Ontology and a range of tools to complete tasks and achieve goals.
- **AIP Assist.** An LLM-powered support tool designed to help users navigate, understand, and generate value with the Palantir platform. Users can ask AIP Assist questions in natural language and receive real-time, context-aware assistance and troubleshooting.
- **AIP Evals:** AIP Evals are the foundation of stable, reliable AIP workflows in production environments; by using AIP Evals to test and evaluate LLM-based functions and prompts, you can build confidence in LLM-backed workflows. By setting up test cases and evaluation criteria in AIP Evals, you can systematically debug, iterate, and improve your implementations, compare different models, and examine variance across runs.
- **AIP Logic:** AIP Logic is a no-code development environment for creating, testing, and deploying AI-powered functions - giving you a point-and-click way to use the power of LLMs backed by the data in your Ontology. AIP Logic provides an intuitive interface for engineering prompts, setting up automation, and integrating structured or unstructured Ontology data.
- **AIP Model Catalogue:** AIP Model Catalogue helps with discovery and orientation of AIP-provided models, enabling users to view and select the right models for their use case, get started with building workflows and test models in sandbox settings.
- **AIP Observability:** AIP Observability enables you to gain comprehensive insights into your AIP and Ontology workflow executions, helping you understand the performance of your agents, functions, language models, automations, Actions, and Ontology.
- **Palantir MCP (Model Context Protocol):** Palantir MCP enables external AI IDEs and agents to connect to the Palantir platform and gain context on your Ontology and Foundry tools. Use Palantir MCP to let external AI systems query data, access documentation, and build applications more efficiently.

2.4 AIP Bootcamps

Palantir created AIP Bootcamps to help existing and prospective clients get hands on with our software, tackle real problems most pertinent to them, and enable swift identification of tools and workflows that can enhance time-critical missions. AIP Bootcamps, whether single or multi-day, provide an opportunity to experience generative AI through use cases relevant to your organisation. You'll build live alongside Palantir engineers and AI specialists, all working toward the common goal of deploying operational AI to solve your most difficult problems.

3.0 PALANTIR PLATFORM: GOTHAM OVERVIEW

3.1 Introduction

Palantir Gotham (Gotham) is a proven, end-to-end, COTS platform for AI-enabled defence, intelligence, and law enforcement operations. It enables secure collaboration, intelligence analysis, and investigation tools to support front-line operations, specialised investigations, multi-jurisdictional strategic intelligence operations and strategic initiatives. It is the result of more than a decade and a half of partnerships with military, intelligence, and law enforcement agencies. Gotham is operationally deployed for international and national security across all Five Eyes ("FVEY") nations and coalition partners, NATO members.

3.2 How Does It Work?

Gotham solves the problem of increasing data and threat complexity by bringing together data, users, and analysis into a single, cohesive platform. Rather than working in isolation with data siloed across disparate systems, authorised users from across an organisation can search, access, and analyse the same data foundation and share data with their peers under the platform's fine-grained security model. All platform application and features are backed by stringent, platform-wide security. In addition, Gotham incorporates and leverages the broader Palantir Platform’s capabilities, offering a suite of analytical and operational applications which can accommodate a wide range of workflows within a unified, interoperable solution.

Gotham enables operational advantage and proactive deterrence for military, law enforcement, and intelligence agencies by providing a common operating picture ("COP") and common intelligence picture ("CIP"), and by supporting intelligence collection and production, investigative network analysis at scale, mission planning and delivery, and after-action analysis. This lets commanders and staff visualise and analyse information from multiple systems in real time, and to work collaboratively and securely across the operating environment to achieve successful mission outcomes.

Key capabilities:

- Creates an integrated data asset serving as a single source of truth;
- Imports and exports data to any system, processing all data formats, including audio and video;
- Conducts investigative workflows including link analysis, geospatial analysis, and object analysis;
- Enables secure team collaboration with granular access controls;
- Provides customisable interfaces tailored to specific workflows;
- Features open APIs for interoperability and extensibility;
- Incorporates AI capabilities for enhanced analysis and decision-making.

3.3 Platform Principles

Data Integration

Gotham transforms structured and unstructured data into objects representing real concepts (people, organisations, places, events) and their relationships. This adaptable data model – the Ontology (described Section 1.2 above) – changes based on organisational needs.

Unified Search and Discovery

The platform offers single-point search across all data sources. Users can explore data in its original format or enriched views, with relevant data automatically surfaced for review.

Secure Collaboration

Every piece of data maintains access restrictions down to individual attributes. All interactions are audit-logged, enabling secure insight sharing across organisational boundaries while maintaining data integrity.

Openness and Extensibility

Gotham interoperates with common programming languages and external systems through REST APIs. All data can be exported in open formats, and the platform evolves with technology advances. Overall platform interoperability is described in Section 1.3 above.

3.4 Platform Features

Please note, there are applications which are used for workflows such as ISR and ISINST analysis which are not listed below. For more information on these capabilities please contact Palantir directly.

Users access Gotham's analytical functions through the Palantir Workspace - a modern interface connecting applications and services, accessible from any browser or desktop application.

Browser enables viewing and editing information about objects within the platform's data model. Users can view relationships, edit properties, add notes, and track changes. Features comprehensive search capabilities across all integrated data sources.

Custom Object Views (COVs) configurable dashboards in Browser that allow users to customise how information is presented. COVs can display different views for different teams and support custom widgets answering specific questions.

Object Explorer provides intuitive analysis of vast data sets, enabling users to find entities with similar characteristics, run analyses on millions of records, and perform drill-down analyses with embedded tools like timelines and aggregation functions.

Chat enables secure communication application for messages, files, and data in classification-controlled environments. Features include multi-channel views, automatic language translation, and integration with third-party messaging systems.

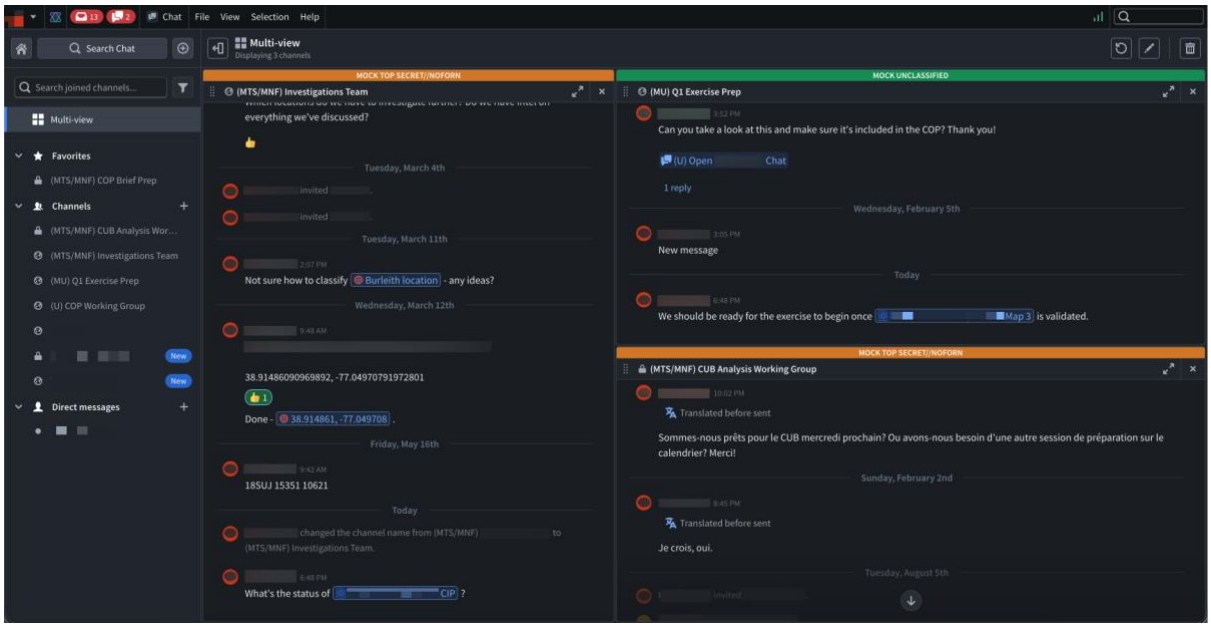

Figure 5: *Notional view of secure* collaboration between teams in 'Chat'

Inbox centralises notifications and alerts in an inbox-style interface. Users can subscribe to various alert types including search feeds, object watch feeds and geofence alerts.

Slides is a data-centric presentation application with real-time information updates. It

supports collaborative editing, template-based creation, and broadcasting capabilities for presentations. Data and decks from Slides can be published to standard document formats (PPTX and PDF) for broader dissemination.

Dossier is a collaborative text editor for creating dynamic intelligence products. Users can embed objects, graphs, maps, and document snippets that maintain links to source data and update automatically.

Graph is a network analysis application for creating visual representations of networked data. Users can view property statistics, organise and annotate graphs, and utilise analytical tools like histograms, timelines, and search functions.

Gaia is a collaborative mapping application with drawing tools and geospatial analysis capabilities. Enables users to work on shared maps with unified data assets, perform geographic searches, and create heatmap analyses. Integrates with third-party platforms and supports various map formats.

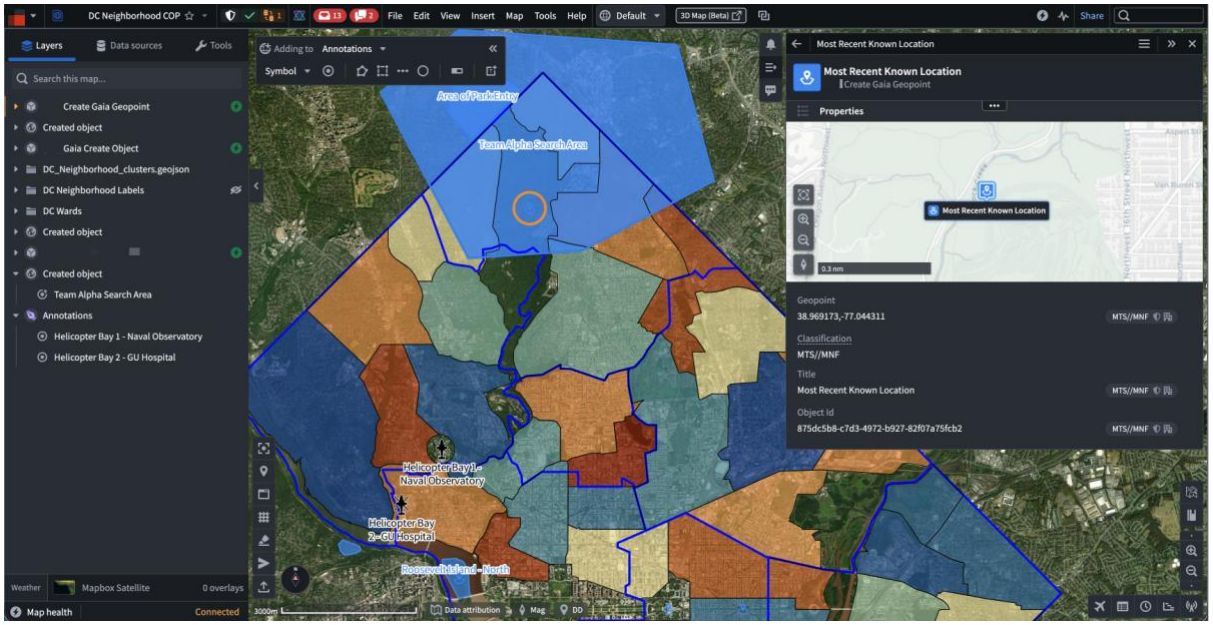

Figure 6: Notional geospatial analysis and map overlays in Gaia

Video provides analysis and enrichment of full motion video within Gotham. Features include real-time AR overlays, AI/ML model integration for detection, data tagging, and export capabilities with redaction options.

Kairos creates interactive Gantt charts for operational planning, entity tracking, and activity synchronisation. Kairos integrates with Gaia for combined temporal and geospatial visualisation.

Workbench is a visual workflow management system that models complex processes through customisable states. It enables tracking of items through different stages with varying levels of automation and human oversight. It leverages the Ontology to ensure workflows remain consistent with organisational data models and business rules.

4.0 PALANTIR PLATFORM SUPPORT OVERVIEW

4.1 Summary

Palantir can provide support for the Palantir Platform, encompassing technical support, infrastructure management (per the Platform's standard SaaS model), release management, and training. To help customers implement and onboard to the platform, Palantir can provide assistance with planning, set up and migration, security services, quality assurance and performance testing, training, and ongoing support.

4.2 Support Features

- Cloud software support; including public, private, community and hybrid clouds
- Migration and set up assistance

- Business analysis and cloud solution guidance
- User and service management
- End-to-end managed support
- Intuitive, easy to use interface
- Training on how to navigate the platform
- Helpdesk support
- AI-assisted troubleshooting

4.3 Support Benefits

- Ensure ease of use for technical and non-technical employees
- Enhanced strategic planning, reducing business risk and costs
- Rapid migration to chosen cloud service
- Fully managed migration process
- Optimised use of cloud software
- Flexible training delivery for varying user adoption rates and locations
- Multiple training modes; instructor-led, internet webinars and self-guided learning
- Bespoke, tailored training materials based on specific contract/project requirements

4.4 Platform Updates

As part of our SaaS-based licensing model, Palantir continuously updates the platform through Palantir Apollo, our end-to-end continuous delivery and maintenance system. Updates follow Palantir's product roadmaps, which in turn are informed by close feedback loops between customer users and the Palantir implementation team assigned to each customer. Excluding any sensitive information or customer data, feedback is shared with our Product Development Teams and aggregated anonymously across all customers to inform future development. In this way, customers benefit from the lessons learned at other organisations facing similar challenges.

Palantir uses a Continuous Integration (CI) / Continuous Delivery (CD) method of upgrading our software products. This ensures that every code change is restricted to the changes required to address a bug or implement a feature, traceable to a set of tests, documented in mandatory change logs (fix/feature/break), and reviewed by other developers and our Information Security Team where relevant. By continuously delivering incremental changes, Palantir enables our customers to benefit from ongoing improvements to product-level capabilities, and a reduced change management burden due to the relatively small footprint of each CI/CD-based release.

4.5 Customer Success Services

In addition to the support, training and documentation provided as part of our standard software licensing, Palantir can also provide enhanced support, solution and use case

design consultancy and training services for platform users (subject to additional fees) via our Customer Success programme. Please refer to our Lot 3: 'Palantir Cloud Support Services' offering for a description of supplementary support and implementation services.

5.0 SECURITY

Palantir is firmly committed to protecting data security, privacy and civil liberties. As such, Palantir has made security its highest priority at every point in the development of the Palantir Platform. The platform adheres to NIST 800-53 and NIST 800-171 standards. Additionally, it undergoes annual audits, including SOC 2 Type 2 (Security, Confidentiality, and Availability) and SOC 1 Type 2, and holds certifications for FedRAMP High and FedRAMP Moderate, U.S. DoD Impact Level 5 and 6, CMMC LvL 2, ISO 27001/27017/27018, ISO 9001, Cyber Essentials and Cyber Essentials Plus, ENS, TISAX, and IRAP. Palantir also performs an annual NHS Data Security Protection Toolkit assessment for the platform

As a data processor, Palantir has significant experience in assisting clients in meeting specific regulatory and industry requirements, such as HIPAA, GxP, GDPR, CCPA, CJIS, and FISMA High.

Palantir operates in a broad variety of security environments with extremely diverse authentication requirements. In summary, the Palantir Platform's Security Model provides:

Platform Security:

The Palantir platform has security as a core development philosophy. The Palantir security model enables strict enforcement of granular access controls with transparency and usability to build a collaborative and trusted ecosystem:

- **Strict enforcement:** Ensures users only have access to data that they have been authorised to interact with.
- **Granular controls:** Powerful enough to achieve flexible levels of access control granularity.
- **Transparency:** Enables users to reason about *who* has access to *what* resource and *why*.
- **Usability:** Empowers users to reason about and manage access controls with confidence.

The Palantir security model encompasses both authentication and authorisation. Authentication verifies the identity of a user, while authorisation grants access based on a user's attributes and permissions. Data security is guaranteed through a combination of mandatory and discretionary controls. Mandatory controls propagate along with each unit of data or resource type, via Palantir's sophisticated provenance and lineage capabilities. Discretionary permissions are granted to users on individual resources, in the form of roles with different operations (for example, view or edit). In addition, granular row or column-level controls based on a user's attributes can be put in place on resources too.

For highly sensitive data, markings are another form of mandatory controls that can be applied to data or resources that require special protection (for example, PII or financially sensitive data). Users must have special permission to discover or access such data, in addition to Organisation membership.

Enterprise Security

- Mandatory **encryption** of all data, both in transit and at rest, that uses robust, modern cryptography standards.
- Strong **authentication** and identity protection controls, including single sign-on and multi-factor authentication.
- Strong **authorisation** controls, including mandatory and discretionary access controls.
- Robust **security audit logging** for detecting and investigating potential abuse.
- Highly extensible **information governance, management, and privacy controls** to meet the needs of any use case.

Infrastructure Security

Palantir's hosted infrastructure has additional layers of security controls to help protect your data:

- Robust **security architecture** built around principles of zero trust, least privilege, and defence-in-depth.
- Enforced **security baseline configurations** with rigorous change management and security monitoring processes.
- Strong **network security** hardening and segmentation.
- Host-based and network-based **intrusion detection systems** to detect and defeat anomalous activity.
- Aggressive **infrastructure and application vulnerability management and patching**, with industry-leading SLAs.
- **Web application firewall** (WAF) inspection of incoming web requests to detect and block attacks.
- **Security monitoring** at every layer of the environment, including users, hosts, networks, and applications.

6.0 INFORMATION ASSURANCE

Our approach to information assurance takes all appropriate protective measures to establish administrative, technical and physical safeguards to assure information in all its forms.

The Palantir Platform complies with NIST 800-53 and NIST 800-171 standards and undergo annual SOC 2 Type 2 and SOC 1 Type 2 audits. The platforms maintain certifications including FedRAMP Moderate and High, U.S. DoD Impact Level 5 and 6, CMMC LvL 2, ISO 27001/27017/27018, ISO 9001, Cyber Essentials and Cyber Essentials Plus, ENS, TISAX, and IRAP, and complete annual NHS Data Security Protection Toolkit assessments. Palantir services encompass threat mitigation, information security management, data

traceability, access controls, encryption, systems development and maintenance, infrastructure security, and disaster recovery. Palantir platforms include built-in technical measures that protect and defend the system by ensuring its availability, integrity, authentication, confidentiality and non-repudiation. These measures were factored into the original software development process and continue to be priorities for current and future configuration.

Availability: The Palantir Platform is designed for high availability and have been successfully deployed in environments with stringent uptime requirements. The applications, hosted in the cloud on behalf of customers, operate in an active-active configuration within a Region, with multiple duplicate nodes processing system requests in parallel and bidirectional data replication. This setup ensures that if one node goes down, another node will seamlessly pick up the load with zero failover time. To achieve this high availability, Palantir deploys redundant instances of each core service across isolated Availability Zones within a single Region, along with redundant storage and a fault-tolerant architecture. Palantir provides robust system patching on a frequent schedule, and all team members are experienced in regularly updating systems to minimise vulnerabilities. The platform conducts regular, incremental backups of all data and provides the option to restore remotely.

Integrity: Palantir provides a high degree of system integrity by encrypting all data at rest and in transit and by regularly patching the operating system and all applications. Palantir platforms also include an audit log to ensure that all system activity and data usage is aligned with any governing rules or policies.

Authentication: The Palantir Platform provides internal authentication and authorisation services that can store user groups and permissions and authenticate user credentials for the system. These services support Single Sign On (SSO), which allows for a single point of entry and access control. Palantir platforms can also integrate with third-party integration services such as Public Key Infrastructure (PKI) and Active Directory, and supports multi-factor authentication.

Confidentiality: Palantir utilises appropriate controls such as firewalls, password protection, encryption and digital certificates at all times to protect confidential information that is processed by, stored in or transmitted from the system. Palantir platforms include robust, granular access controls so that each individual piece of information can be marked with the appropriate classification. Palantir also ensures confidentiality by encrypting all data in transit and at rest.

Non-repudiation: Palantir offers non-repudiation by authenticating, monitoring and auditing all user activity in the system in protected access and activity logs.

7.0 ONBOARDING PROCESS

Palantir uses a multi-phase approach to the on-boarding process, balancing robust governance with an agile and rapid approach to capability provision that has been tested,

refined and proven across over hundreds of deployments. This multi-phase approach typically includes:

- Scoping and Clarification / Preparation Phase
- Infrastructure Setup Phase
- Implementation Phase
- Support & Maintenance Phase

8.0 AGILE METHODOLOGY

Palantir uses an Agile methodology to implement and configure our platforms. Structuring the Implementation Phase into Agile Sprints (time-bound periods of work, generally 1-2 weeks) gives customer stakeholders at all levels the opportunity to provide timely feedback, design input and updates on task prioritisation. Palantir then uses this information to inform Sprint planning, allowing them to quickly address changes in requirements with minimal disruption to the project schedule.

Unlike the Waterfall methodology, which can be linear or difficult to alter when project circumstances shift, Agile project plans are designed to incorporate the unexpected without disrupting the delivery of higher-level capabilities. In this way, an Agile approach will be better placed to build upon customer feedback on newly delivered capabilities, as well as to conduct any future unforeseen configuration work that may be required beyond the initial goals of this project. Further discussion and detail on Palantir's Agile methodology can be provided upon request.

9.0 TRAINING

Palantir has extensive training material available via https://learn.palantir.com/. This includes dozens of separate courses in different languages, as well as guided training tracks for Data Engineers, Application Developers, AI Engineers, Data Scientists, and Frontend Developers. The training resources include documentation, YouTube content and a community for support in learning and troubleshooting.

For customer projects, Palantir designs and executes customer-specific training plans based on the user profiles, scale, workflows, and production timelines for each deployment of our software. Our approach to user training is flexible and aims to accommodate the different user groups of the relevant platform and their technical competencies to ensure all users become proficient at using the platform for their specific roles. Training can be delivered on site by Palantir, with expert support from Palantir's global engineering resources as needed, or via a variety of other self-guided methods.

The general types of training provided are:

- **In-person, instructor-led training**: Palantir can hold specific training sessions at customer locations, tailored according to user profile, specific contract requirements and

project stage.

- **Webinars**: Webinars are available on a variety of topics, based on ongoing assessments of end user needs. Webinars allow flexibility scheduling, varying user adoption rates and location.
- **Self-guided learning**: Palantir also provides for self-paced training through our web-based video training application that includes features such as videos and documentation.

10.0 SERVICE MANAGEMENT AND SUPPORT

10.1 Maintenance and Support Approach

Palantir's maintenance and support approach is designed to minimise system downtime and ensure that users have the access they require to perform mission-critical work. Support is accessible by phone or email. Our quality assurance measures ensure that support meets or exceeds industry best practices and is tailored to each individual organisation.

10.2 Support Services, including Service Desk Support

To ensure that users and system administrators are fully supported throughout the duration of the contract, Palantir can provide a combination of in-person and remote support services as needed. Help Desk support is handled by the Palantir's embedded engineering teams who provide prompt triage, response, and resolution of issues.

Palantir will implement, install, monitor, and test all parts of the platform for correct operation and can fully support the needs of the customer on an ongoing basis by providing maintenance services, including troubleshooting during critical periods and ongoing configuration work.

Within the Palantir Platform, a support portal called Issues provides ticket logging, AI-assisted help, and immediate access to system information and help documents, including user guides, training manuals, frequently asked questions and troubleshooting documentation.

10.3 Security Incident Management

Our support model utilises Palantir's best practices for quality assurance and quality assurance surveillance. Palantir has a standard incident management process, annually updated, and approved by management, for security events and suspected security incidents that may affect the integrity, availability, or confidentiality of Palantir systems, such as data breaches. This process specifies courses of action, procedures for notification, escalation, mitigation, and documentation.

Palantir maintains both a process for manually reporting security incidents and an automatic alerting system of security events. Palantir employs activity and integrity monitoring tools with automated alerting and uses industry standard tools for triaging, notifying relevant parties, and tracking alerts and incidents. When the response teams receive a security alert, they tag the alert with a priority level and gather all relevant and precipitating details. Alerts are triaged on the basis of highest-priority alerts first. High-severity alerts receive human action promptly. Other alerts receive a response based on predetermined SLAs. Alerts are categorised and tagged with relevant information for later review and analysis. If relevant, alerts are escalated to a Security Incident.

11.0 Change Management

Palantir's Change Management Policy sets standards for upgrading capabilities, responding to threats, adhering to laws/regulations, and complying with contract obligations, while limiting impact and ensuring adequate messaging. All changes to systems must be submitted as a "Change Request". The relevant teams review the request, prioritise, and develop a plan for implementation. Authorised personnel approve changes under Palantir policies and procedures, and customers are notified of any major changes per agreed upon processes.

The process typically includes:

- Change Tracking
- Testing
- Approval
- Communicating Changes
- Scheduled Maintenance
- Monitoring Changes
- Emergency Changes

11.1 Release Management and Preventive Maintenance

Palantir will perform release management and preventive maintenance on the Palantir Platform to ensure that they are kept in proper and reliable working order. Our maintenance structure is adaptive, providing regular and high-priority releases (as needed) and continually reprioritising work based on feedback and trends.

11.2 Feedback

Palantir constantly collects useful feedback from our users, monitors trends in incidents and new feature requests, and prioritises development work against what is observed in the field. Notable incidents are consolidated into internal tracking systems, where they are continually reviewed and analysed against other incidents received from across Palantir. Issues of wider concern are flagged to other engagement teams or to appropriate product developers for

resolution in future product releases. Due to our rapid response times, Palantir have provided releases to customers in as little as one hour for emergency fixes (e.g., for emergency security vulnerabilities).

12.0 SERVICE CONSTRAINTS AND LEVELS

Palantir recognises that each customer has different maintenance schedules and different service continuity needs. The Palantir Platform is configurable to meet the specific needs of a customer's environment and Palantir teams can be flexible to meet the needs of a customer with respect to planned maintenance, service disruption and outages, and will ensure that any such events are appropriately communicated.

Wherever possible, planned maintenance will be carried out without affecting the service. This will generally be achieved by carrying out planned maintenance during periods of anticipated low traffic and by carrying out planned maintenance on part, not all, of the network at any one time. Where emergency maintenance is necessary and is likely to affect the service, Palantir will endeavour to inform the affected parties as soon as possible within the start of the emergency maintenance.

Refer to the 'Ongoing support' and 'User support' sections of the 'Palantir Platform' catalogue entry webpage for further information on Palantir's Service Level Agreements (SLAs).

13.0 SERVICE TERMS

Please refer to the 'Terms and Conditions document' on the 'Palantir Platform' catalogue entry webpage.

14.0 BUSINESS CONTINUITY & DISASTER RECOVERY

14.1 Business Continuity

Palantir conducts a regular, and at least annual, risk assessment and business impact analysis to understand and mitigate risk of business disruption. Results of these assessments are used to update Palantir's contingency policies and procedures. To ensure the reliability of plans in the event of an incident, contingency exercises are performed on representative environments at least annually and results are documented. Results of contingency exercises are then used to update the Continuity Plan.

Palantir has configured its platforms to be highly available and to take regular backups as part of its contingency planning. Palantir has minimised the need for downtime related to

hardware failure or for maintenance to hardware or software products. Additionally, Palantir takes backups of the system to reduce the risk of data loss from corruption, accidental deletion, or to resolve disputes.

High Availability Configuration
The Palantir Platform is designed to be highly available. Customers are hosted in an active-active configuration; meaning that within a Region the fail-over status is Hot.

Data Centres
Palantir uses best in class cloud hosting providers such as Amazon Web Services (AWS), Microsoft Azure (Azure), Oracle, and Google Cloud Platform (GCP) as the underlying cloud service providers for our platforms. The global infrastructure of these services is built around Regions and Availability Zones. A Region consists of multiple Availability Zones, physically separated and isolated, where offered by the underlying cloud service provider, which are connected with low-latency, high-throughput, and highly redundant networking. Cloud Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data centre infrastructures. Using this infrastructure, Palantir platforms are designed and operated such that failover between Availability Zones within a Region is executed automatically and without disruption to users (RPO/RTO: 0/0).

Data Stores
Palantir Foundry uses the underlying cloud providers' storage solutions, such as AWS S3, Azure Blob Storage, or Google Cloud Storage, as its primary customer data store. Palantir Foundry additionally uses Cassandra installed on the cloud providers' compute platform, such as AWS EC2 instances, Azure Virtual Machine hosts, or GCP Compute Engine, as its primary metadata and configuration data store.

Palantir Gotham uses Postgres or Oracle on the underlying cloud providers' storage solutions, such as AWS EC2 or RDS, Azure Virtual Machine hosts, or GCP Compute Engine, as its primary customer data store and metadata and configuration data store.

Additional computations may run on datasets within the platform to produce data transformations and may generate and store results or views in Elasticsearch or Postgres. Any computational index metadata, such as these, are considered secondary data stores and can be automatically rebuilt or rehydrated based on the Palantir platform’s primary data stores.

Palantir platforms leverage the underlying cloud service providers’ infrastructure to ensure all data, in primary or secondary stores, is replicated to multiple active hosts simultaneously, so that the platform is resilient to single points of failure and automatically restores redundancy after hardware faults are resolved. If an Availability Zone were to go down, no data will be lost permanently, and an RPO/RTO of 0/0 is achieved.

The SLAs of the underlying cloud service providers services are determined by the underlying cloud service provider and can be found on their respective websites.

14.2 Disaster Recovery

Backups
Palantir platform backups are archived sets of data used to restore the original after a data loss event. While a highly available configuration, as described above, mitigates risk resulting from a loss of a physical datacentre or other hardware outage, it does not mitigate the risk of data loss due to such things as accidental data deletion or corruption. To mitigate this risk, Palantir also takes regular backups of the system.

Upstream Customer Data Sources
Customers are responsible for the integrity and backups of data in customer managed source systems. The resilience of such data integrated into Palantir platforms is dependent on the customer's business continuity practices.

Platform Configuration Backups
Palantir leverages a system called Rescue to back up and restore the services that make up Palantir platforms. The backup cadence defaults to taking a backup snapshot every one(1) to two (2) hours. Rescue stores these snapshots in an encrypted CSP storage bucket, such as AWS S3, Azure Blob Storage, or Google Cloud Storage. These snapshots can be used to restore the entire state of Palantir platform to a point in time. The default retention period for the Rescue backup is seven (7) days. Palantir continuously monitors Rescue to ensure the health of backups.

15.0 TECHNICAL REQUIREMENTS

Please refer to the 'Using the service' section of the 'Palantir Platform' catalogue entry webpage.

16.0 PRICING AND LICENSING

Palantir licenses its software through:

- **Platform Licences** which provide the set of infrastructure engineering services required to set-up, maintain, monitor and assure enterprise-scale deployments of the Palantir Platform. These licenses allow customer organisations to create user accounts for their personnel to access their platform instance and a range of out-of-the-box tooling to integrate, manage and exploit their data.
- **Capability Licences** which provide Organisations with specialised workflow configurations to enable a specific outcome or set of outcomes.
- **Sustain Licenses** refer to Capabilities which were previously Active Capabilities (above) that are no longer being actively iterated on. No ongoing configuration is being applied to the Capability, but it needs to be sustained in its current operational state. Sustain Capabilities receive engineering services focused on support, monitoring and, where necessary, bug fixes of an in-production Capability to ensure that the customer organisations can continue to use that Capability as it stands.

- **Third Party Licenses** Customers may choose to have third-parties develop additional Capabilities on the Platform. Third-parties shall be defined as any organisation or party (e.g. the customer or a contractor) developing Capabilities on the Platform, where Palantir is not contracted to do so. The development of Capabilities by third-parties shall be subject to similar scoping for the purpose of understanding the Capability size and appropriate corresponding licence.

For pricing information please refer to the 'Pricing document' on the 'Palantir Platform' catalogue entry webpage.

17.0 OFF-BOARDING PROCESSES

Palantir platforms lay an integration layer on top of an organisation's disparate IT landscape, thereby serving as a single point of access to all data within an enterprise without requiring data duplication, data cleansing, or data warehousing. In this way, the cost, time and risk associated with implementation or cessation of service are minimised.

Palantir is an open platform with open, non-proprietary data and file formats, public APIs, a plug-in architecture, and numerous extensibility points. Once data is integrated into Palantir platforms, it is easy to export it in many different formats or otherwise make it accessible to other software systems. Such flexibility allows data to be exported in formats that are easily digestible by other tools as needed. Palantir can assist the customer with any data retention, archiving, transportation or destruction requirements.

Palantir commits to purge and destroy customer data from any computers, storage devices and storage media that are to be retained by Palantir after the end of the contract period, and the subsequent extraction of customer data (if requested by the customer).


Palantir