Palantir

PALANTIR PLATFORM: FOUNDRY & AIP

SERVICE DEFINITION DOCUMENT

Prepared For: G-Cloud 14 Framework

palantir.com/uk

Foundry AIP

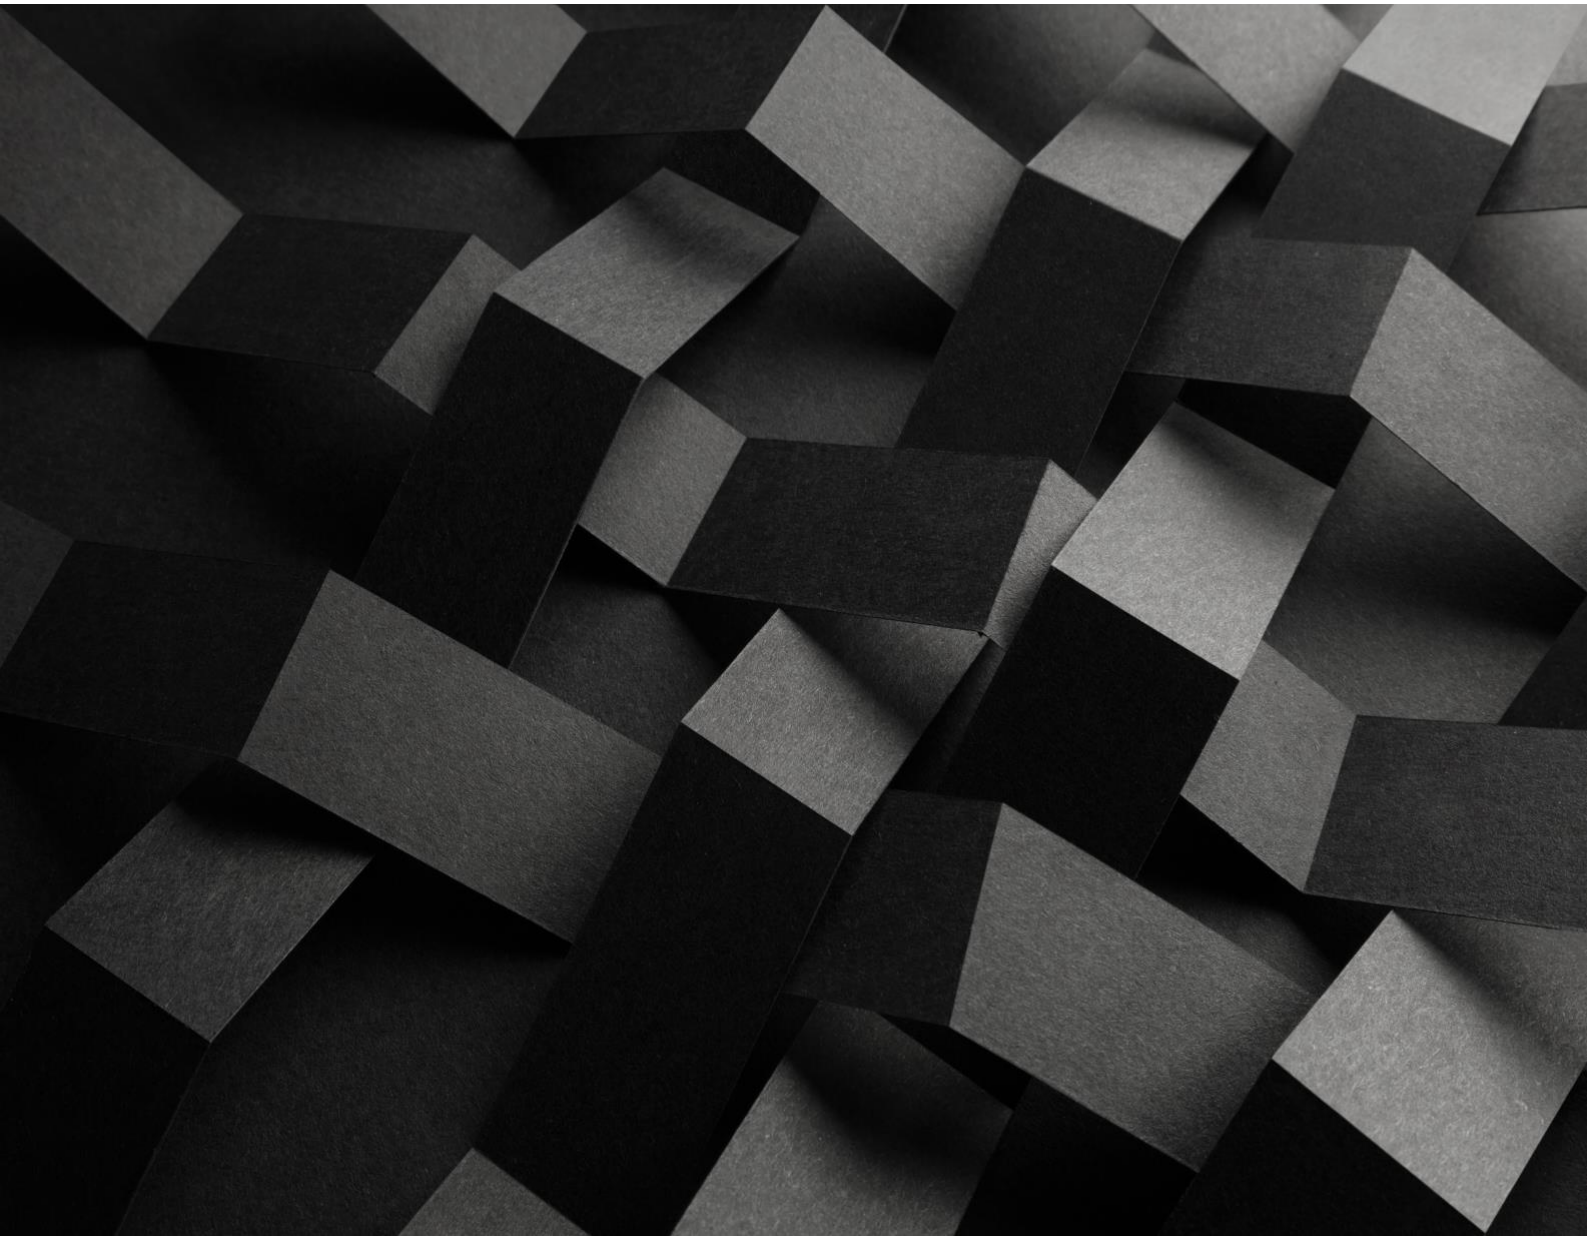

Table of Contents

1.0 PALANTIR FOUNDRY SOFTWARE OVERVIEW

Useful Resources: Please visit our website and YouTube channel 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 fifteen years of experience enabling data-driven operations in the most dispersed and complex environments. Palantir’s customers include, intelligence, and defence organisations in multiple countries; healthcare providers; banks and financial regulators; energy majors; and large manufacturing and logistics firms. In all of 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 the 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.

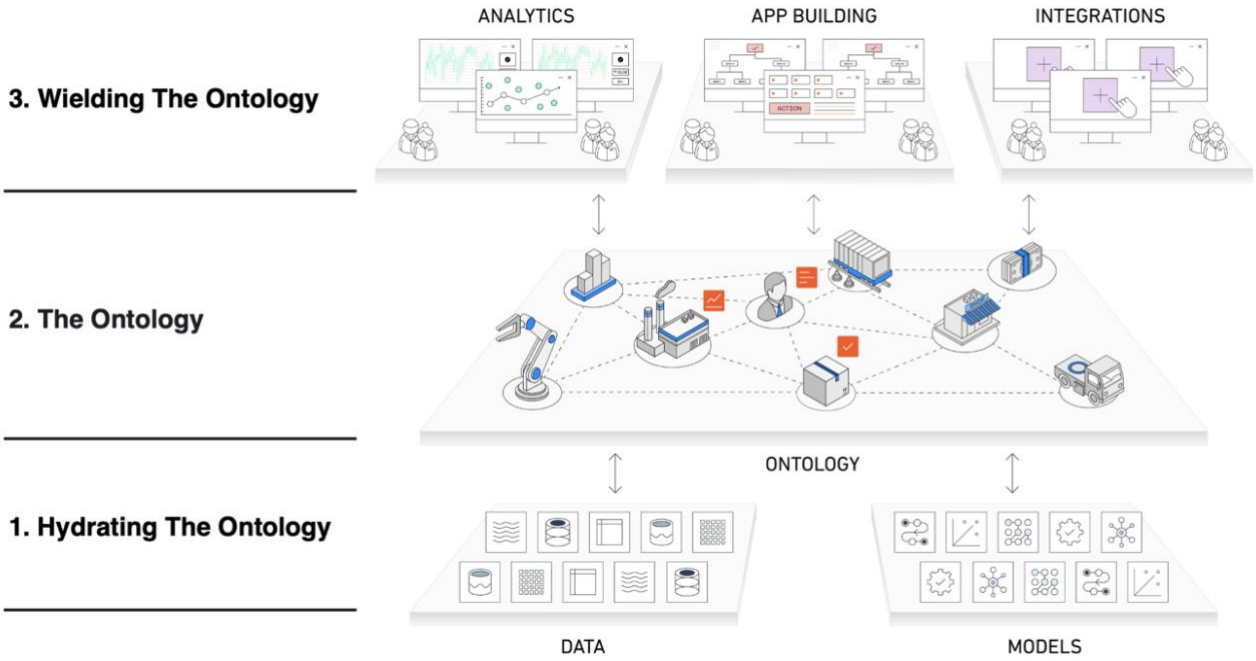

1.2 Value

Foundry allows organisations to:

- Integrate data from and export data to virtually any system, including existing applications and streaming sources. Foundry can process data in any form, including audio and video;
- Handle all data transformations, enrichments, and metadata;
- Manage and govern integrated data, and apply a consistent security, audit, and data provenance framework;
- Conduct queries and business intelligence analytics on integrated data;
- Build and interact with live, data-driven insights in configurable interfaces;
- Allow everyone, from data engineers to non-technical subject matter experts, to use data to make informed decisions;
- Deploy organisation-wide workflows to drive and refine operations over time;
- Respond to changes as they occur and undertake forward-looking “what-if” analysis by eliminating the uncertainty and long lead times of analyses based on periodically-refreshed cuts of static data;
- Unlock secure collaboration between teams, departments and organisations; and
- Optimise and operationalise machine learning and artificial intelligence tools.

1.3 How Does It Work?

Palantir Foundry is a platform that reimagines how people use data by removing the barriers between back-end data management and front-end data analysis. Palantir Foundry enables users with varying technical ability and deep subject matter expertise to work meaningfully with data. With Palantir Foundry, anyone can source, connect, and transform data into any shape they desire, then use it to take action.

Palantir Foundry is backed by a suite of best-in-class capabilities for data integration that run on data and business logic in tandem:

- Versioning semantics to keep data and business logic in sync;
- Dynamic, systemwide security and access controls to replace unreliable one-off policies;
- Branching of code, analyses, and reports to enable safe experimentation;
- Microservice architecture with built-in coordination, security, and upgrades to keep individual components in sync;
- Open APIs and data formats to interoperate with an organisation's entire data ecosystem; and
- Flexible data protection frameworks to keep up with evolving regulations and industry best practices.

Palantir Foundry's front-end capabilities let every user tap into the power of their organisation's data:

- A central data foundation to drive collaboration and discovery across functions;
- A common ontology to turn a complex data landscape into a human-readable representation of the entire organisation;
- Datasets and analyses that feed back into the platform to allow users to build on one another's work, rather than constantly starting from scratch;
- Human-readable data lineage to let users jump from insights to the data and logic that feed them; and
- Diverse analytical tooling to supercharge traditionally non-technical functions and accelerate advanced analytical initiatives.

1.4 Platform Principles

Palantir Foundry's capabilities comprise the four (4) core pillars of a flexible and enduring data transformation:

Data Security

Protect data confidently with automatic propagation from source system to final insight

Understand how an insight came to be with lineage and versioning of both data and code

Protect production without disconnecting it from the sandbox environment

Business Ontology

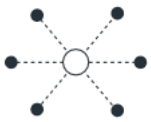

Unify the organisation by capturing every business concept in a common ontology

Compound business intelligence by feeding insights back into the ontology

Improve the quality of ontology data automatically and continuously

Analytical Diversity

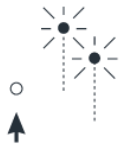

Empower business analysis with point-and-click environments that unlock complex analytics

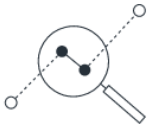

Supercharge advanced analytics for data engineers and data scientists

Accelerate machine learning and artificial intelligence with quality data and seamless deployment to production

Openness & Extensibility

Enhance the value of existing IT investments by centralising data operations

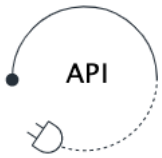

Plug in to in-house and third-party solutions through open data formats and open APIs

Accelerate future projects and reduce their cost with reusable data pipelines and centralised management

1.5 Platform Features

Deliver immediate, compounding business value.
With the whole organisation collaborating on the same data foundation, the cost of new data projects drops, and the value of the data asset increases over time. Instead of putting success at the end of a multi-year roadmap, Palantir Foundry lets organisations achieve critical outcomes from the start.

Unite the organisation around a common ontology.
Collaboration takes off when the whole organisation is speaking the same language. Palantir Foundry lets organisations translate their entire business into an ontology: a set of building blocks that map business concepts to the data that describes them. With one flexible ontology as a starting point for every user, new questions, analyses, and projects enhance organisational knowledge rather than fragment it.

Manage data and business logic in tandem.
Business logic codifies the knowledge that holds an organisation together. Palantir Foundry manages business logic in tandem with the data it runs on so that as logic evolves, insights do too. Users can always trace an insight back to the data and logic that feed it.

Secure the data once; secure the system in perpetuity.
Palantir Foundry lets organisations define granular access control policies at the integration stage, then propagates those policies intelligently across the system. Organisations can promote data access confidently with granular data security and transparent data governance.

Instil trust in data with continuous improvement.
In a living data ecosystem, data integrity is a moving target that requires continuous improvement over time. Palantir Foundry combines automated data quality checks with tools for users to flag issues when they see them, sustaining the integrity of the data asset over the long term.

Make operations analytical and analytics operational.
Successful data transformation calls for collaboration across the entire organisation. Palantir Foundry blurs the lines between functions so that subject matter experts answer mission-critical questions without learning to code, and data scientists operate at the heart of the business.

1.6 Scalability and Interoperability

Palantir 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.

To enable efficient and elastic scaling, Palantir Foundry:

- Scales horizontally across commodity servers, whether using on-premises hardware or commercial cloud infrastructure.
- Uses computing resources efficiently, with Distributed File Systems (e.g., Hadoop, Amazon S3) as a backing store to improve query efficiency and spread operations over multiple systems.
- Maximises performance and storage density.

Avoid vendor lock-in and manage data flexibly with an open system.
Many data management systems lock customers into a solution by storing data in proprietary formats or closing off systems with proprietary APIs. Palantir Foundry stores data in open formats and exposes open APIs to facilitate interoperability. Palantir Foundry exposes data to external services via several interfaces. Customers can easily export structured data from Palantir Foundry to CSVs and databases via direct query or a web API.

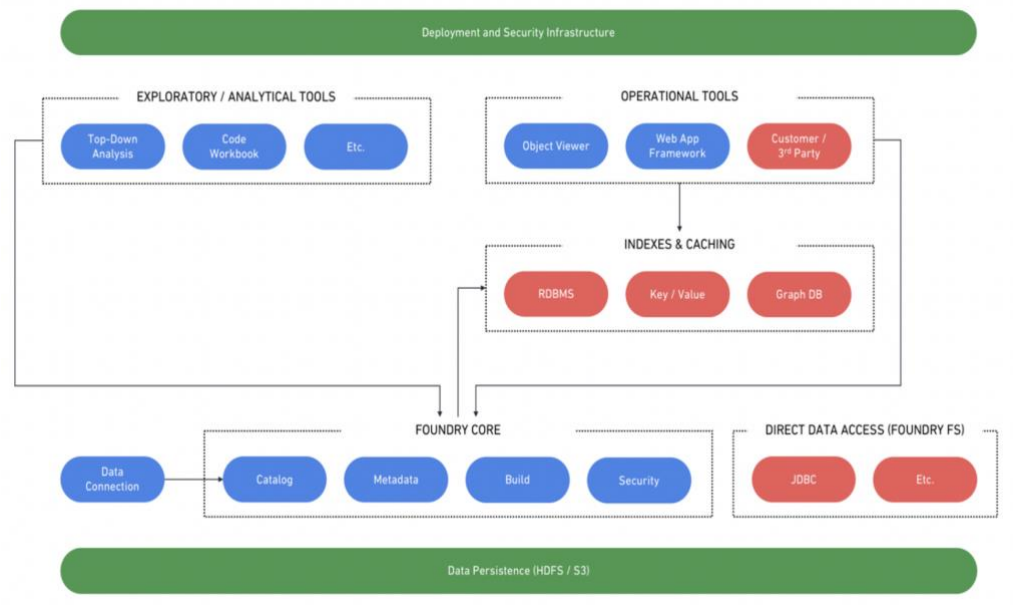

Palantir Foundry provides several “push” mechanisms, including file transfer and common standard connections (e.g., JDBC/ODBC drivers that enable Java applications to interact with databases). External systems can also “pull” from Palantir Foundry via methods like common drivers and RESTful Web Service APIs and by connecting directly to the Distributed File System where data lives. As an integrated ecosystem, Palantir Foundry includes the base data management layer, an authoring environment for data transformations, a suite of user-facing analytical applications, and developer frameworks & open APIs for building operational applications.

1.7 Security and Collaboration

Palantir Foundry’s granular access control framework lets customers secure information at the dataset level and assign specific degrees of access for different user groups. For each individual dataset, the customer can define the users who are permitted to discover, read, modify, and delete the data.

Palantir Foundry maintains an audit trail that captures all user activity within the platform. For every user action—read, write, delete—Palantir Foundry captures what data was accessed, where, when, and by whom. Palantir Foundry also captures a detailed history of integration, including time of connection, source, and revision history. This metadata is used to track data provenance and manage compliance with data auditing and retention policies.

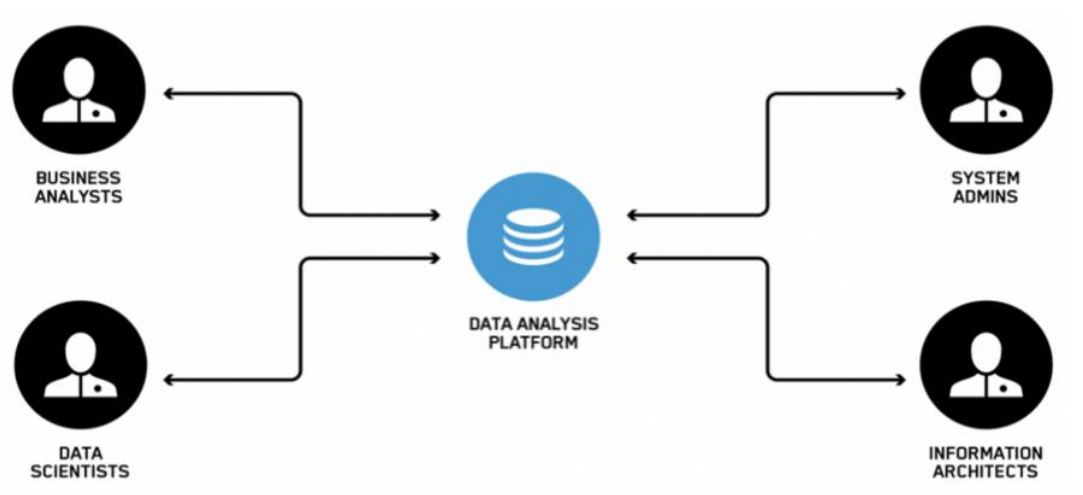

Palantir Foundry’s role-based security model enables seamless cross-functional collaboration among different user groups.

All volumes and buckets associated with the platform is encrypted by default using best in their class key management systems such as Amazon Web Services Key Management Service, Microsoft Azure Key Vault, or Google Cloud Platform Key Management Service, which generate encryption keys via FIPS 140-2 validated hardware security modules to provide centralised management of cryptographic keys and operations in the cloud. All calls, both management and cryptographic requests, to the hardware security modules are centrally logged (as Security Logs) and monitored by Palantir. Ephemeral instance storage, where integration with hardware security modules is not offered by cloud providers, is encrypted at the instance hardware level or with OS-level LUKS encryption using AES-256-XTS. Palantir Foundry offers an additional layer of application-level encryption, whereby each file is encrypted with a distinct symmetric key (AES-256) before being stored in the appropriate bucket at rest. AES keys are then envelope encrypted with an asymmetric key pair (RSA-2048) and stored alongside the data at rest.

All data transferred to, from, and within the platform is encrypted without exception. This includes not only boundary-traversing traffic and traffic containing Customer Data, but also internal data flows and all types of intra-service communication. The platform only supports modern, strong, and industry-accepted TLS protocol versions (1.2+) and encryption ciphers. HTTP Strict Transport Security (HSTS) is enforced for platform domains to ensure network traffic is encrypted from end user web browsers to the platform. Emails sent from the platform to customer users for notification workflows are similarly encrypted. Palantir specifies a DNS Certification Authority Authorisation (CAA) record for production domains. The CAA record designates which Certificate Authorities may issue a TLS certificate. This is a defence-in-depth control to ensure that certificates may not be issued for Palantir domains without violating this record. Certificate Transparency (CT) logs are monitored to detect unusual or malicious certificate issuance in violation of the CAA record.

Palantir Foundry's security model and versioning capabilities contribute to a collaboration framework that breaks down the barriers that prevent cross-organisational information sharing. Palantir Foundry lets the customer define types or groups of users who can collaborate (e.g., data scientists, system administrators, business analysts, etc.). The security model ensures that data is only exposed to users with the right permissions.

2.0 PALANTIR ARTIFICIAL INTELLIGENCE PLATFORM (AIP)

2.1 Introduction

Palantir AIP (Artificial Intelligence Platform) powers real-time, AI-driven decision-making in the most critical commercial and government contexts around the world. From public health to battery production, organisations depend on Palantir to leverage AI safely, securely, and effectively in their enterprises to drive operational results. You can learn more about the platform in this documentation, or from an AIP Bootcamp where customers are hands-on-keyboard and achieving outcomes with AI in a matter of hours.

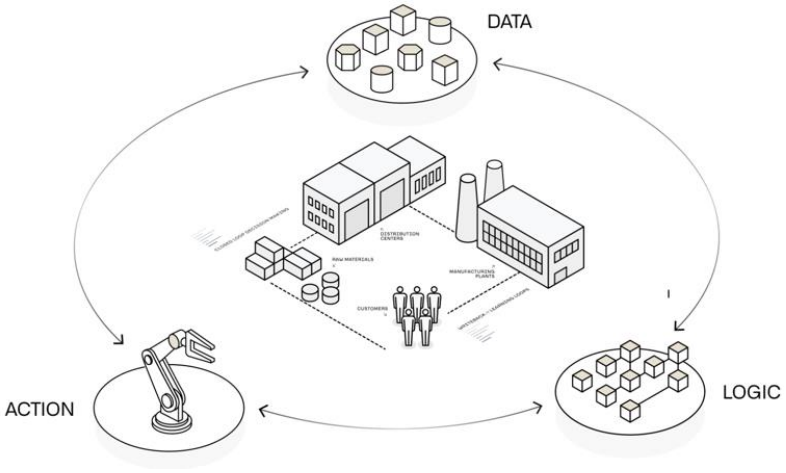

2.2 Palantir AIP Overview

AIP is a comprehensive solution for securely accessing, integrating, and managing Large Language Models (LLMs) and AI technologies. AIP enables customers to leverage the latest generation of LLMs (e.g., OpenAI's GPT-4) for their operations, within secure and private environments hosted across Amazon Web Services (AWS), Microsoft Azure (Azure), and Google Cloud (Google) infrastructure. AIP offers 200+ system connectors and enables users to create custom tools for querying Ontology objects and executing models. The platform supports operational AI use with decision-making interfaces, feedback loops, and seamless collaboration between AI agents and human operators. AIP also provides full-spectrum security, audit controls, and integrated human review checkpoints. In addition to secure mediation of external models, AIP provides organisations the ability to leverage privately trained and managed LLMs seamlessly throughout workflows. AIP offers the following specific capabilities:

- **Securely Access LLMs of Your Choice.** Unified access to a full range of open-source, self-hosted, and commercial LLMs is provided. 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.

- **Integrate Data and Build Pipelines with Natural Language.** Out-of-the-box connectors for 200+ systems are available, leveraging an extensible connection framework. LLMs can be used to produce the data pipelines needed for all enterprise consumers, including LLMs themselves.

- **Run on Your Structured and Unstructured Data.** Private data can be made usable by the LLM through Vector Stores and the Ontology, which transform data into LLM-understandable objects and actions and processes into tools for humans and LLM-driven agents.

- **Create Tool Driven Plugins.** Dozens of out-of-the-box tools are available for querying Ontology objects, executing models, running graph computations, and much more. A Workbench for creating custom tools that can be securely wielded by humans and AI is also provided.

- **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.

- **Automate Execution.** A step-by-step process for converting LLM Logic flows into secure, governed automations is provided. Flexible, real-time monitoring services allow for trigger-based and schedule-based automations. End-to-end traceability ensures that all automation execution and downstream effects can be rigorously audited.

- **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.

2.3 AIP Bootcamps

Palantir built 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 business. 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. These intensive, hands-on-keyboard sessions allow organisations to go from zero to use case in one to five days, with initial bootcamps typically offered free-of-charge.

3.0 PALANTIR FOUNDRY & AIP SUPPORT OVERVIEW

3.1 Basic Summary

Palantir can provide support for Palantir Foundry & AIP. To help users learn to use Palantir Foundry and other Palantir platforms which utilise Palantir Foundry, Palantir can provide assistance with planning, set up and migration, security services, quality assurance and performance testing, training, and ongoing support.

3.2 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 Palantir Foundry
- Helpdesk support

3.3 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

3.4 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 / use case design consultancy and training services for users of our Foundry & AIP platform (subject to additional fees) via our Customer Success program. Please refer to our Lot 3: ‘Palantir Cloud Support Services’ offering for a description of supplementary support and implementation services.

4.0 PALANTIR SECURITY

Palantir believes data analysis software becomes a liability when it lacks robust, built-in access controls, and 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 Palantir, and it is why organisations in national security and global finance trust Palantir to safeguard their most important data assets. Palantir platforms are aligned with NIST 800-53 and NIST 800-171. Additionally, Palantir platforms annually complete a SOC 2, Type 2 audit (Security, Confidentiality, and Availability); are certified compliant with FedRAMP Moderate, U.S. DoD Impact Level 5, ISO 27001/27017/27018, and ISO 9001; hold cyber essentials and cyber essentials plus certificates; and perform an annual NHS Data Security Protection Toolkit assessment.

In addition, as a data processor, Palantir has extensive experience helping customers meet specific regulatory and industry requirements, including HIPAA, GxP, GDPR, CCPA, CJIS, and FISMA High.

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

- Fine-grained access controls that secure every piece of data individually;
- Specific degrees of access including ownership, write, read, discovery and no access permissions;
- Collaboration with security;
- Secure data integration; and
- Full and immutable audit trail.

5.0 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.

Palantir platforms are aligned with NIST 800-53 and NIST 800-171. Additionally, Palantir platforms annually complete a SOC 2, Type 2 audit (Security, Confidentiality, and Availability); are certified compliant with FedRAMP Moderate, U.S. DoD Impact Level 5, ISO 27001/27017/27018, and ISO 9001; hold cyber essentials and cyber essentials plus certificates; and perform an annual NHS Data Security Protection Toolkit assessment.

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: Palantir platforms are designed for high availability and have been successfully deployed in environments with stringent uptime requirements. Palantir platforms provide several features that are designed to increase system availability, including redundant storage and a fault-tolerant architecture. Palantir will provide robust system patching on a frequent schedule, and all team members are experienced in regularly updating systems to minimise vulnerabilities. Palantir can conduct regular, incremental backups of all data in the system 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: Palantir platforms provide 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.

6.0 ONBOARDING PROCESS

Palantir uses a multi-phase approach to the on-boarding process, which provides robust, agile and rapid approach to capability provision that has been tested, refined and proven across over hundreds of deployments with reliably excellent results. This multi-phase approach typically includes:

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

6.1 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 one (1) or two (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.

7.0 TRAINING

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.
- **Internet 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. Palantir have successfully used our web-based video training method at many engagements with diverse user bases.

8.0 SERVICE MANAGEMENT AND SUPPORT

8.1 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.

8.2 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.

As a commercial product, our platforms also include an online support portal that provides authorised staff with immediate access to system information and help documents, including user guides, training manuals, frequently asked questions and troubleshooting documentation.

8.3 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 tools such as Jira and PagerDuty 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.

8.4 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 users 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

8.5 Release Management and Preventive Maintenance

Palantir will perform release management and preventive maintenance on our platforms 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.

8.6 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 (1) hour for emergency fixes (e.g., for emergency security vulnerabilities).

9.0 SERVICE CONSTRAINTS AND LEVELS

Palantir recognises that each customer has different maintenance schedules and different service continuity needs. Palantir platforms are 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. Level of availability varies depending on the specific project.

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

10.0 SERVICE TERMS

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

11.0 BUSINESS CONTINUITY & DISASTER RECOVERY

11.1 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.

Description of the Platform's Business Continuity Configuration
Palantir has configured our 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.

Highly Available Configuration
Palantir platforms are 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), 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 Palantir Foundry or Palantir Gotham 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.

11.2 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 are 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 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.

12.0 TECHNICAL REQUIREMENTS

Please refer to the 'Using the service' section of the 'Palantir Platform: Foundry & AIP' catalogue entry webpage.

13.0 PRICING

Please refer to the 'Pricing document' on the 'Palantir Platform: Foundry & AIP' catalogue entry webpage .

14.0 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.

15.0 DATA REMOVAL AND EXTRACTION

15.1 Data Removal

Palantir provides various capabilities for restricting and removing data from its platforms, including a hard delete capability, which involves a full purging of the specified data object and amounts to complete, irrecoverable removal of the underlying information and destruction of all traces thereof from the hardware. 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).

15.2 Data Extraction

If the customer wishes to extract their data when the contract ends, Palantir can export all existing data in the platforms into raw formats. Palantir's software platforms have been purposefully designed to prevent vendor lock-in. As such, they have an open, pluggable architecture with publicly documented APIs at every tier of the software. All data in the platforms can be securely exported in non-proprietary formats for use in other databases or systems. Palantir will work with the customer to determine the best export format(s) for customer datasets and their destination systems.


Palantir