Palantir

PALANTIR DATA PLATFORM FOR POPULATION HEALTH MANAGEMENT

SERVICE DEFINITION DOCUMENT

Prepared For:
G Cloud 12 Framework

Prepared By:
Palantir Technologies UK, Ltd.
www.palantir.com

Palantir

Table of Contents

1. SERVICE OVERVIEW & SOLUTION 3

1.1 BASIC SUMMARY 3
1.2 HOW DOES IT WORK? 3
1.2.1 Integrating and Analysing Data 4
1.2.2 Managing Data Flexibly with an Open System 5
1.2.3 Access Control and Data Governance 5
1.2.4 Population Health Management 5

2. INFORMATION ASSURANCE 6

3. PALANTIR SECURITY 7

4. ON-BOARDING PROCESSES 7

4.1 ON-BOARDING PROCESS 7
4.2 AGILE METHODOLOGY 7

5. TRAINING 8

6. SERVICE MANAGEMENT AND SUPPORT 8

6.1 MAINTENANCE AND SUPPORT APPROACH 8
6.2 SUPPORT SERVICES, INCLUDING SERVICE DESK SUPPORT 8
6.3 INCIDENT MANAGEMENT 9
6.4 CHANGE MANAGEMENT 9
6.5 RELEASE MANAGEMENT AND PREVENTIVE MAINTENANCE 9
6.6. FEEDBACK 9

7. SERVICE CONSTRAINTS AND LEVELS 10

8. SERVICE TERMS 10

9. BACKUP/RESTORE & DISASTER RECOVERY PROCESSES 10

9.1 BUSINESS CONTINUITY 10
9.2 DISASTER RECOVERY 11

10. TECHNICAL REQUIREMENTS 11

11. PRICING 11

12. OFF-BOARDING PROCESSES 11

13. DATA REMOVAL AND EXTRACTION 12

13.1 DATA REMOVAL 12
13.2 DATA EXTRACTION 12

Palantir

1. SERVICE OVERVIEW & SOLUTION

1.1 Basic Summary

The Palantir Data Platform for Population Health Management (the "Platform") is a configuration of Palantir's enterprise data management platform that enables care providers and health authorities to assess care quality and outcomes, to optimise care pathways and to identify, pilot and monitor care interventions for population cohorts.

Electronic health records are typically optimised for delivering care for individual patients at the level of an individual provider. It is often difficult for providers and health authorities to access clean and structured patient data, especially across different providers within an area, region or country. Efforts to pull together patient data to support population health efforts are often manual and time-consuming and technical challenges in ensuring appropriate access controls and security end up limiting the extent to which users can access data and collaborate. In addition, it is really difficult to complete the feedback loop between a population health "study" and population health interventions or changes to care pathways being implemented and then monitored for their effectiveness over time.

The Platform enables providers and health authorities to stratify patients into cohorts based on needs and risk factors, analyse care pathways in terms of outcomes and costs, identify interventions that could be beneficial to certain cohorts, pilot care interventions or changes to care pathways, monitor the results and operationalise successful interventions. It does this by:

- Integrating all relevant clinical and service data and metadata from source systems into a single secure environment;
- Linking previously distinct, siloed data and mapping it to a customer-defined ontology, making it easy for users to access and explore the information by connecting the data to real world entities and concepts such as patients, diagnoses, procedures and care pathways;
- Providing a range of analytical tools such as a cohort explorer, a data science coding environment and a machine learning model management framework, and operational tools for defining, piloting, monitoring and redefining care pathway interventions; and
- Applying rigorous security and access controls to ensure that only authorised users are able to see individual data points that are appropriate based on their role.

1.2 How Does It Work?

The Platform ingests and stores data, offers basic and advanced data transformation tools, and removes the barriers between back-end data management and front-end data consumption. Underling these tools are best-in-class security features and granular access controls.

The Platform includes a suite of capabilities for robust data integration and data governance, which include:

- Versioning semantics to keep data and business logic in sync;
- Granular, systemwide and provenance-aware access controls that are applied to all resources in the Platform including data, analyses, reports and models, replacing

unreliable one-off policies and ensure data cannot be mistakenly shared with users not authorised to see it;
- Git-like branching of both code and data to ensure that changes to data transformation pipelines can be made quickly and robustly;
- Microservice architecture with built-in coordination, security and upgrades to keep individual components in sync;
- Open APIs and data formats to interoperate with a cross-organisation data ecosystem; and
- Flexible data protection frameworks to keep up with evolving regulations and industry best practices.

The Platform's front-end capabilities enable users across care providers and health authorities to drill-down into cohorts of patients, examine care pathways, and plan and manage interventions. Its core features include:

- A central data foundation to ensure a common operating picture across providers;
- A common ontology to turn a complex data landscape into a human-readable representation of the data with rapidly configurable front-end views of ontological concepts;
- Platform-wide data lineage that allows users to move from front end views into the data and logic that feed them;
- Diverse analytical tooling to supercharge traditionally non-technical functions and accelerate advanced analytical initiatives;
- First class model management capabilities for tracking assumptions, objectives, versions and validations of models from development through to testing and deployment; and
- Extensive tooling for constructing custom user interfaces which support high value workflows.

1.2.1 Integrating and Analysing Data

The Platform has been built using industry standard technologies for efficient storage, processing and transformation of massive-scale data. It integrates common data sources (e.g. HDFS, JDBC and SQL databases, flat files) out of the box and can be configured to support any other source system or legacy technology. Using the Platform, data integrations with source systems can be built, tested and managed from development environments integrated seamlessly into the Platform.

Once data has been ingested it is typically subject to cleaning, normalisation and other kinds of processing. Crucially, incoming data is mapped to a customer defined ontology which corresponds to tangible concepts representing real-world entities, relations and events. By mapping data in this way, downstream users are able to navigate and analyse the data asset without a need for in depth knowledge of the various source system's conventions and characteristics.

As data is stored and transformed incrementally on the Platform, users responsible for maintaining integrations are able to quickly examine, test and adjust the logic responsible for each stage of processing. This renders the provenance of downstream data highly transparent and enables the rapid diagnosis and correction of any problems arising through the data integration and transformation process.

1.2.2 Managing Data Flexibly with an Open System

Many data management systems lock customers in by storing data in proprietary formats, or by closing off access with proprietary APIs. In contrast, the Platform stores data in open formats and exposes open APIs to facilitate interoperability.

The Platform provides various “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” data from the Platform via methods like common drivers, RESTful APIs, and by connecting directly to the Platform’s storage layer.

As an integrated ecosystem, the Platform includes the base data management layer, an authoring environment for data transformations, a suite of user-facing analytical applications, and developer frameworks and open APIs for building operational applications. All data and business logic stored on the Platform is stored in open, industry standard formats and can be exported as required in conformity with the customer’s business requirements and access control policies.

1.2.3 Access Control and Data Governance

The Platform features a granular access control framework that allows authorised users to secure information at the object level and assign specific degrees of access for different user groups. The Platform provides intuitive interfaces for customer administrators to define the users that are permitted to discover, read, modify and delete each object. Permissions are inherited for all derived intelligence products, ensuring that access controls persist as data is transformed and analysed.

The Platform integrates seamlessly with common enterprise authorisation systems (e.g. Active Directory, LDAP, SAML, Kerberos, etc.) without need for extensive manual configuration. To protect data at rest, data is stored in an encrypted format. In transit, data is encrypted via SSL/TLS, during both client-to-server and server-to-server communications.

To further safeguard integrated data, the Platform maintains an audit trail that captures all user activity within the Platform. For every user action (i.e. read, write, delete, etc.) the Platform captures what resource was accessed, where, when and by whom. The Platform also captures a detailed history of integrations, 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.

As such, the Platform’s security model and versioning capabilities contribute to a collaboration framework that breaks down the barriers that prevent inter- and cross-organisational information sharing. When sharing data, Platform tenants can be confident that data is only exposed to users with the right permissions.

1.2.4 Population Health Management

Users can leverage highly configurable user interfaces which present the most pertinent information about care pathways in an intuitive, comprehensible format, for example, volumes of patients on a given pathway over time, associated procedures, therapies and costs, typical waiting times and patient outcomes. Users can drill-down to better

understand care pathways to prioritise for improvement based on any combination of metrics, for example, based on number of impacted patients, longest waiting time, highest cost per patient, worst outcomes, etc. Alternatively, users can start with patients and drill-down to identify sub-cohorts that may benefit from additional interventions using the cohort explorer tool, for example, patients with multiple morbidities or patients with late stage diagnoses.

The Platform also provides easily configurable scenario planning tooling to enable users to plan interventions and operational improvements and understand the expected impact in terms of patient outcomes and costs. The model management capabilities within the Platform enable more technical users to develop and test model to feed into the scenario planning tooling. Model assumptions, objectives, versions and validations are tracked from development through to testing and deployment ensuring that operational users have full transparency into any tools that feed into their decision-making process.

2. 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. Our security strategy is based on International Organisation for Standardisation (ISO) 27001, ISO 27002, and other industry best practices, and encompasses threat mitigation, information security management, data traceability, access controls, encryption, systems development and maintenance, infrastructure security, and disaster recovery. Palantir includes 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 is designed for high availability and has been successfully deployed in environments with stringent uptime requirements. Palantir provides 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 also includes an audit log to ensure that all system activity and data usage is aligned with any governing rules or policies.

Authentication: Palantir 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 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 includes 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.

3. 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 is GDPR and SOC 2 compliant, and holds Cyber Essentials Plus certification.

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.

4. ON-BOARDING PROCESSES

4.1 On-Boarding 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

4.2 Agile Methodology

Palantir uses an Agile methodology to implement and configure the Platform. Structuring the implementation phase into Agile Sprints (time-bound periods of work, generally one or two 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.

5. 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 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.

6. SERVICE MANAGEMENT AND SUPPORT

6.1 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 24 hours a day, 7 days a week, 365 days a year. Our quality assurance measures ensure that support meets or exceeds industry best practices and is tailored to each individual organisation.

6.2 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, the Platform also includes 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.

6.3 Incident Management

Our support model utilises Palantir's best practices for quality assurance and quality assurance surveillance. This includes the use of a ticket tracking and task management system to systematically collect and address incidents and monitor compliance with service targets. The customer can report incidents directly to Palantir, via agreed upon means. All reported incidents are tracked and triaged for efficient coordination of the issue through to resolution.

6.4 Change Management

Palantir's change management process establishes an orderly and effective procedure for tracking the submission, prioritisation, implementation and approval for release of all change requests. Our approach considers various factors when prioritising changes, including complexity and scope.

The process typically includes:

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

6.5 Release Management and Preventive Maintenance

Palantir will perform release management and preventive maintenance on the Platform to ensure that it is 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.

6.6. 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 in 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, and 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).

Palantir

7. 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 'Availability and resilience' section of the Service Offering webpage for further information on Palantir's Service Level Agreements (SLAs).

8. SERVICE TERMS

Refer to the 'Terms and Conditions document' on the Service Offering webpage.

9. BACKUP/RESTORE & DISASTER RECOVERY PROCESSES

9.1 Business Continuity

Business continuity is an essential part of making Palantir a resilient organisation. Palantir's policy ensures that business continuity arrangements support organisational objectives and is designed to:

- Provide continuity of service to critical business activities;
- Enable Palantir to recover and return to normal as quickly as possible following a disruptive event;
- Enable compliance with regulatory and customer requirements;
- Support effective risk management; and
- Ensure the safety of staff and visitors.

Palantir works with its customers to ensure that business continuity arrangements are up to date and effective, and that enough people are trained in maintaining and using their business continuity plans when responding to incidents that might cause disruption to their operations. Business continuity plans typically include:

- Business Impact Assessment, listing services provided and their order of importance;
- Backup and restoration plan for services and systems supporting critical services;
- Succession plan;
- Communication plan;
- Plan for annual testing; and
- Requirement for regular review.

In implementing business continuity, Palantir follows the guidance provided by ISO 22313 and the Business Continuity Institute (BCI) in the 2018 version of its Good Practice Guidelines. Both are consistent with ISO 22301, the management systems standard for Business Continuity Management.

9.2 Disaster Recovery

Disaster recovery depends on the hosting infrastructure. With AWS-backed cloud infrastructures, Palantir encrypts and stores all raw data feeds in AWS's S3 data stores, which feature exceptionally high durability. These stores can be used to restore services in the event of an outage. Additionally, all operational databases, including application platforms and Palantir metadata, are backed up by default, with encrypted snapshots stored in AWS S3. The backup and restore of these services are available via Palantir management tooling and should be done automatically. It is typical practice to schedule regular fire drill testing to ensure that the full environment can be duplicated from backups under the required SLA.

Using the global infrastructure of the underlying cloud service providers, Palantir can continuously stream all critical data to multiple live hosts, and this is the Platform's primary backup strategy. In addition, the Platform leverages a secondary backup system, called Rescue, that leverages S3 and EBS snapshots - configurable but by default incrementally taken every 2 hours - which are backed by S3.

10. TECHNICAL REQUIREMENTS

Refer to the 'How users work with your service' section of the Service Offering webpage.

11. PRICING

Refer to the 'Pricing document' on the Service Offering webpage.

12. 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, 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

13. DATA REMOVAL AND EXTRACTION

13.1 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).

13.2 Data Extraction

If the customer wishes to extract their data when the contract ends, Palantir can export all existing data in the Platform 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 Platform 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.