

## Executive Summary: Palantir Technologies – ICD 501 Compliant

On January 21<sup>st</sup>, 2009 the Office of the Director of National Intelligence (ODNI) signed Intelligence Community Directive 501 (ICD 501). ICD 501 succinctly outlines the roles, rights and responsibilities of the Intelligence Community (IC) with regards to the ‘Discovery and Dissemination or Retrieval of Information within the Intelligence Community’. A review of these policies reveals that many of the outlined responsibilities will require the adoption of new technologies if they are going to be implemented.

In the attached document, we have highlighted the portions of ICD 501 that are relevant to the development of new technologies and discussed in detail how the Palantir Technologies solution can be brought to bear on these various challenges. What we have shown is that despite the fact that ICD 501 was meant to be a road map for future action, Palantir meets all the requirements of ICD 501 today. Specifically, Palantir meets and exceeds the requirements of the following six sections:

- *Responsibility to Provide (D.2.a)*
  - Palantir provides a rich, distributed data environment where data stewards can make their data available to all analysts while maintain fine-grained access control
- *Responsibility to Discover (D.3)*
  - ‘Discovery’ is the notion that users can be alerted to the existence of information without actually seeing the information itself. Palantir provides a powerful suite of search functions allowing users to discover information. Additionally, Palantir’s collaboration model allows users to discover the work of one another while maintaining a secure multi-level environment.
- *Responsibility to Request (D.4.a)*
  - ‘Request’ is the notion that once information is ‘Discovered’ users should seek to be allowed to see it. Palantir supports this by providing rich auditing and custom discovery messages so that users can quickly identify the data steward.
- *Privacy Rights and Civil Liberties(D.6)*
  - Palantir has a firm and lasting commitment to protecting privacy and civil liberties. Palantir has proven that there does not need to be a trade-off between securing these basic rights and securing the nation itself.
- *User Authorization and Auditing(E.2.a-b)*
  - Palantir provides powerful and extensible user authentication and authorization systems that are tied directly to user access control. Additionally, all user actions are monitored and logged to provide a powerful user-level audit trail. In conjunction with user auditing, all data is audited as well so that users can identify when, by whom and from what source it was created and also when, if ever, the data was modified.
- *Responsibilities of IC Leaders (G.2.i)*
  - IC leaders are required to make arrangements and agreements that are commensurate with ICD 501. Given that Palantir is the only product that meets these requirements today, the path forward is clear.

The IC could take one of two paths: the first is the well-worn path of trying to build a solution in-house costing hundreds of millions of dollars and taking years to complete; the second is to acquire the solution that works today for 10% of the cost and 0% of the risk. Palantir Technologies is the solution.

## Palantir Technologies – ICD 501 Compliant

On January 21<sup>st</sup>, 2009 the Office of The Director of National Intelligence (ODNI) signed Intelligence Community Directive 501 (ICD 501). ICD 501 succinctly outlines the roles, rights and responsibilities of the Intelligence Community (IC) with regards to the ‘Discovery and Dissemination or Retrieval of Information within the Intelligence Community’. A review of these policies reveals that many of the outlined responsibilities will require the adoption of new technologies to maximize their effectiveness. In the complex and dynamic security environment that defines the modern Intelligence Community, it behooves us to identify the best technologies available to meet this challenge.

In this document we will review the portions of ICD 501 that are relevant to the integration of new technologies and then discuss point-by-point how Palantir enables these policies today without necessitating any further custom development. By the end of this document, it should be clear that Palantir is the right solution to meet the requirements of ICD 501.

### Intelligence Community Directive Number 501

*Below find a condensed summary of ICD 501<sup>1</sup> that focuses on the portions of the Directive that relate directly to how technology will be necessary to achieve the Objectives outlined in section B (‘Purpose’).*

## ICD 501 INTELLIGENCE COMMUNITY DIRECTIVE NUMBER 501

### DISCOVERY AND DISSEMINATION OR RETRIEVAL OF INFORMATION WITHIN THE INTELLIGENCE COMMUNITY

(EFFECTIVE: 21 JANUARY 2009 BY ORDER OF THE DIRECTOR OF NATIONAL INTELLIGENCE)

Policy Objectives: (B.1-3)

- Foster an enduring culture of responsible sharing and collaboration within an integrated IC
- Provide an improved capacity to warn of and disrupt threats to the United States (U.S.) homeland, and U.S. persons and interests
- Provide more accurate, timely, and insightful analysis to inform decision making by the President, senior military commanders, national security advisers, and other executive branch officials

<sup>1</sup> Summarized from [http://www.dni.gov/electronic\\_reading\\_room/ICD\\_501.pdf](http://www.dni.gov/electronic_reading_room/ICD_501.pdf)

The Policy:

- *Responsibility to Provide (D.2.a)*
  - IC elements shall make all intelligence and intelligence-related information discoverable<sup>2</sup> by authorized IC personnel
  - Stewards shall make all information collected and analysis produced by an IC element available for discovery by making as much information as possible available for automated retrieval upon discovery; and by presuming that authorized IC personnel who request information discovered possess a “need to know”
- *Responsibility to Discover (D.3)*
  - Authorized IC personnel have a “responsibility to discover” information believed to have the potential to contribute to their assigned mission need.
- *Responsibility to Request (D.4.a)*
  - Authorized IC personnel have a corresponding “responsibility to request” relevant information they have discovered that has the potential to contribute to an analytic judgment, to optimize collection, to inform collection strategies and priorities, or to otherwise advance the intelligence mission
- *Privacy Rights and Civil Liberties(D.6)*
  - All IC personnel shall carry out their responsibilities under this Directive, including the discovery, dissemination, retrieval, and use of information collected or analysis produced, consistent with applicable law and in a manner that protects fully the privacy rights and civil liberties of all U.S. persons
- *User Authorization and Auditing(E.2.a-b)*
  - Until such time as an attribute-based identity management capability that enables automated user authorization, discovery, retrieval, and auditing services for IC personnel is approved by the DNI and implemented throughout the IC, IC element heads shall identify authorized IC personnel within their IC element who have discovery rights to information collected and analysis produced by other IC elements
  - The total number of authorized IC personnel approved by stewards to retrieve information collected or analysis produced may be subject to the IC’s ability to adequately audit such activities.
- *Responsibilities of IC Leaders (G.2.i)*
  - [IC Leaders shall] negotiate arrangements, agreements, understandings, or commercial contracts that shall, to the greatest extent possible, seek to obtain terms that permit discovery, and dissemination or retrieval by authorized IC personnel

In the remainder of this document, we will discuss these policies and indicate specifically how Palantir is compliant. Furthermore, it should become clear that Palantir is in fact the only product available that meets and exceeds the compliance threshold.

---

<sup>2</sup> Discovery is the act of obtaining knowledge of the existence, but not necessarily the content, of information collected or analysis produced by any IC element. The act of discovery does not itself constitute a request for receipt of the information collected or analysis produced.

## ICD 501: Policy and Palantir

### Responsibility to Provide (D.2.a)

The Responsibility to Provide is one of the most important requirements outlined in ICD 501 and is, in essence, a call for greater collaboration. One of the paradoxes of collaboration is that in order to share data we must be able to secure data. The more fine-grained our ability to secure data, the more data we can share. As such, a responsibility to provide is also a responsibility to secure.

ICD 501 outlines a notion referred to as 'Discovery' which is defined as:

*the act of obtaining knowledge of the existence, but not necessarily the content, of information collected or analysis produced by any IC element.*

This notion of discovery allows users who are searching for information to be notified if their search matches an existing record without revealing the record itself. If the data owner ('Steward') determines that the user performing the search is authorized to see the record, it can be released. Under the Responsibility to Provide, the Steward is asked to generally assume that the user should be allowed to see the information unless there is a compelling reason to keep it secure.

Palantir is ideally suited to enable this responsibility. The Palantir security model enables the most expansive notion of collaboration possible because Palantir secures every individual scrap of information separately. This means that the maximum possible amount of information can be shared with authorized users. Practically speaking, if one part of a record is classified but the rest is not, the record can still be shared and Palantir will automatically filter the classified parts for unauthorized users. The security model can operate with an unlimited number of user groups, classification levels, security compartments, caveats and dissemination controls.

Palantir not only controls access but also the degree of access allowed for any user to any particular piece of information. Palantir natively enables five degrees of access including the notion of Discovery through a security concept that we also refer to as *Discovery*.

- The first level is *Owner*, which would be the level held by a Steward and allows modification of the permissions for the information in question.
- The second is *Write*, which would allow an authorized user to modify the information but not change the permissions associated with it.
- The third is *Read*, which would allow an authorized user to view the information but not alter it in any way.
- The fourth is *None*, which would completely inhibit the ability of a user to view or discover a piece of information (this is the default permission state).
- The fifth is *Discovery*, which would allow a user to become aware of the existence of a record, but not reveal the record itself. The user is given a configurable *Discovery Message* which alerts them as to whom they must contact in order to view the record in question. Palantir is the only product with this ability today.

## Responsibility to Discover (D.3)

The Responsibility to Discover is a mandate both for Authorized Users as well as for the technology they will use to discover. Essentially, users are required to do whatever is in their power to discover any information that is relevant to their mission. Technologically, this mandates the implementation of a system that enables the user to make these discoveries. Palantir has tackled this problem by solving a series of very difficult problems that individually are very powerful and together create an ideal platform for ICD 501 compliance:

- The first is *Data Integration*: in order to discover information, it is important that the data be made accessible through as few interfaces as possible so that discovery can be made as efficient as possible while limiting duplication. Palantir's ability to integrate data – whether it be structured (Databases, spreadsheets, XML) or unstructured (documents, cables, media) – is unparalleled in the IC.
- The second is *Search*: just having all the data accessible through a common interface is not enough. It must also be possible to search across all the data to identify the records of interest. Palantir has the most robust, flexible and scalable search framework available in the IC, allowing everything from the simplest keyword search, to Boolean logic, to complex filters and graph expansions. Most importantly, all these searches are simple and intuitive requiring no specialized knowledge of structured query languages.
- The third is *Knowledge Management*: as users search their data, make discoveries and derive new connections it is extremely valuable to add those connections to the growing collection of discoveries that other users can make while maintaining security. In Palantir, when a user uncovers a new piece of information or a new connection, that new datum is automatically sourced back to the record from which it was derived. In so doing, Palantir maintains a rich audit trail for all information while also inheriting the security of the source record. This newly created and secured piece of information will now be available for other users in their searches.
- The last problem is *Collaboration*: As discussed in the previous section (Responsibility to Provide), true collaboration is not possible without powerful security. In the previous section we discussed the Palantir security model and how it relates to Collaboration; in this section we will give more detail on the rest of the collaboration model. A critical part of discovery is being able to follow one's own line of thought without being intruded upon by others. Much like security, this also seems to be incompatible with Collaboration. Palantir uses a 'virtual private sandbox' model to allow each analyst to follow their own lines of inquiry and then 'publish' those findings when they need to be shared. Similarly, analysts can accept the updates of other analysts as needed. This model allows analysts to simultaneously explore their own ideas and conjectures while also allowing the creation of an enterprise view of the world.

By combining these four pillars – *Data Integration, Search, Knowledge Management* and *Collaboration* – Palantir enables the analyst to discover new information and develop new insights in ways that are simply not possible with current systems. This ability is critical to demonstrating ICD 501 Compliance and is central to all Palantir deployments. No other product has the capability to meet these four challenges in a single integrated package.

## Responsibility to Request (D.4.a)

The Responsibility to Request places a burden on analysts to dig deeper when they discover information that they need to learn more about. As discussed above, Palantir already has a built in discovery model that enables data stewards and analysts to uncover information in a secure and collaborative way. The configurable discovery message for any data record allows data stewards to ensure that analysts have the information they need to meet their obligation to request more information.

However, Palantir also enables the Responsibility to Request through the *Knowledge Management* model described in the preceding section. As described there, whenever any information is derived from a data source, the derived information is linked to that data source and inherits any security permissions associated with that data source. During this sourcing process, Palantir makes a record of when the information was derived and which analyst derived it. The result is that whenever analysts look at a record in Palantir, they can see the history of that record including the analyst who derived it. Consequently, all analysts become data stewards in a sense, and the community of analysts can now work together to request information when more is needed to understand the results of their searches.

## Privacy Rights and Civil Liberties (D.6)

Section D.6 of ICD 501 states:

All IC personnel shall carry out their responsibilities under this Directive, including the discovery, dissemination, retrieval, and use of information collected or analysis produced, consistent with applicable law and in a manner that protects fully the privacy rights and civil liberties of all U.S. persons

For too long it has been presumed that there is always a trade-off between information awareness and privacy, between security and civil liberties. This manifests itself in many forms including the policy surrounding the sharing of information between intelligence organizations. When our foes no longer respect borders, we find our intelligence apparatus left in the lurch when it has structural limitations imposed by those same borders.

However, Palantir has shown that this trade-off is actually a false choice if the right technology is in place. Because of the fine-grained nature of the Palantir security model it is now possible to share information between individuals with widely varying security clearances and widely divergent missions without fear of leaking any data that cannot be shared either by policy or statute.

For a more in-depth discussion of the Palantir commitment to privacy and civil liberties, please visit our online resources:

- **Privacy and Civil Liberties:** <http://www.palantirtech.com/privacy-and-civil-liberties>
- **Whitepaper:** [Privacy and Civil Liberties Are In Palantir's DNA](#) (pdf)
- **Whitepaper:** [Hard Technical Problems in Civil Liberties Protection](#) (pdf)
- **Video:** [Protecting Privacy and Civil Liberties](#)
- **Press:** [A Tech Fix For Illegal Government Snooping?](#) (NPR)

## User Authorization and Auditing (E.2.a-b)

Section E.2.a-b of ICD 501 lists possible limitations to the full implementation of the directive in the form of user management, authorization and auditing. These are difficult problems and should not be taken lightly. As the directive states, “[until] automated user authorization, discovery, retrieval, and auditing services” can be implemented, the reach of ICD 501 will necessarily be limited. According to the directive, until such systems can be put in place, it will be left to IC element leaders to identify authorized personnel. Fortunately, none of this needs to be considered a limitation as Palantir meets all these requirements today.

Palantir is capable of directly addressing user authorization and how it bears on discovery and retrieval. The Palantir authentication server works natively with Active Directory, LDAP and PKI authorization systems individually or together. In addition, the authentication server is extensible so that customized attribute-based systems for user authentication can be implemented. Once a user has been authenticated using one of the systems described, the user’s access permissions will be in force for the entirety of the user’s session in Palantir and as such all data discovery and retrieval restrictions will be in force for all the user’s searches. User permissions can be changed by data stewards or system administrators on an as needed basis or they can be synced to LDAP user groups so that they change automatically. The Palantir security model is the only true multi-level security system available today.

In addition to providing powerful authentication tied directly to data security, Palantir also has rich auditing of all user actions. User auditing occurs at many levels, from the most miniscule action to the aggregate of all user actions in an enterprise. At the most fine-grained level, every piece of data in Palantir is traced back to the originating user, allowing an exceptionally detailed understanding of the data in the system. At the next level, all records that are composed of these data have an extensive history indicating everything that has happened to that record and the user or users that are responsible for each change. At the next level, the study and analysis of these records is tracked through what is referred to as an Investigation History which tracks the linking and visualization of these records in aggregate along a branching investigative path. At the highest level is Palantir Usage Analytics which tracks the behavior and tendencies of individual users as well as all users in aggregate so that a high level understanding of the usage of the system can be ascertained. No other product available offers this range of auditing capability.

## Responsibilities of IC Leaders (G.2.i)

[IC Leaders shall] negotiate arrangements, agreements, understandings, or commercial contracts that shall, to the greatest extent possible, seek to obtain terms that permit discovery, and dissemination or retrieval by authorized IC personnel

The Palantir Platform meets all the requirements needed to be compliant with ICD 501 and is in fact the only system that does. Additionally Palantir meets the desired user authorization and auditing functions that are not yet part of the official requirements. The final responsibility outlined in ICD 501 is that IC Leaders have an obligation to obtain the technologies that will enable these requirements. Given that Palantir does this today and is the only product that does this today, the path forward is clear.

**Company Information:**

*Palantir Technologies, Inc*

<http://www.palantirtech.com>

GSA contract #: GS-35F-0086U

650-494-1574

100 Hamilton Ave, Suite 300

Palo Alto CA, 94301

**Palantir POC:**

*Shyam Sankar*

[ssankar@palantirtech.com](mailto:ssankar@palantirtech.com)

415-244-5545 (cell)

650-821-0534 (fax)

100 Hamilton Ave, Suite 300

Palo Alto, CA 94301

**Additional Palantir Resources:**

<http://www.palantirtech.com/videos>

A series of video demonstrations

<http://www.palantirtech.com/whitevideos>

A series of technical talks on the Palantir Platform