Palantir

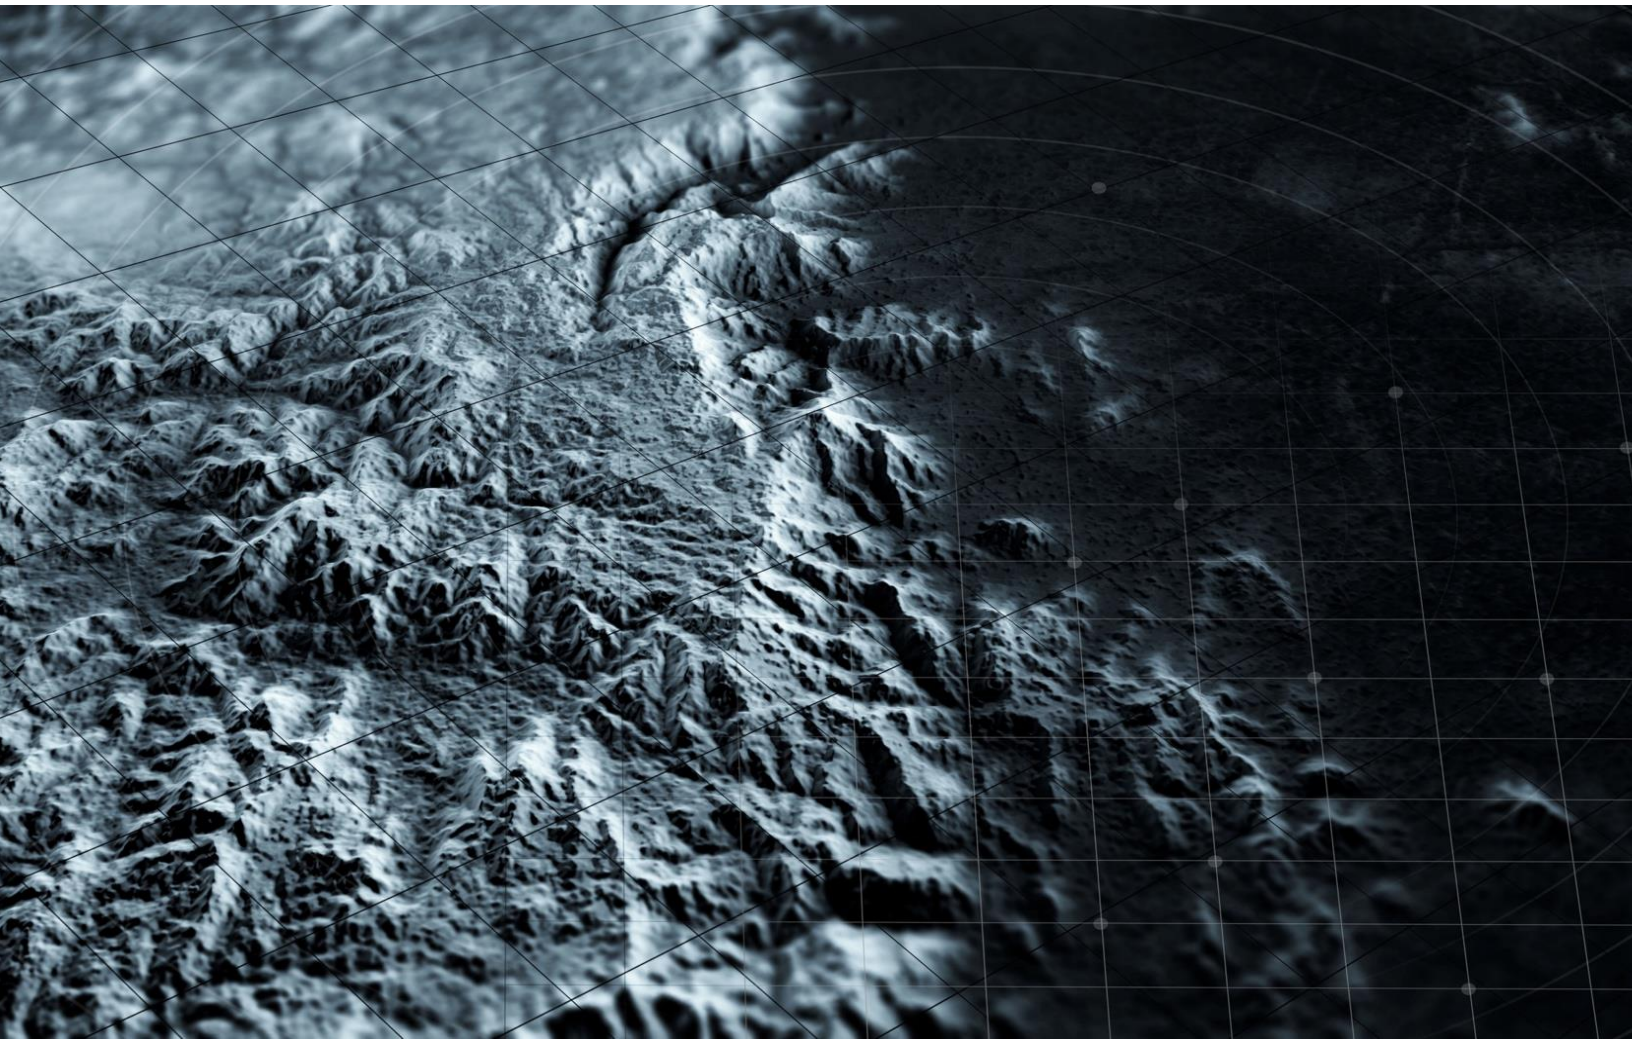
Palantir Civilian Harm Mitigation Technical Whitepaper

November 2026

PALANTIR.COM

COPYRIGHT © 2026
PALANTIR TECHNOLOGIES INC.

ALL RIGHTS RESERVED



Actual results and experiences may vary. © Palantir Technologies Inc. Any data contained herein is notional

Table of Contents

1 Introduction 3

2 Civilian Considerations in the Ontology 5

3 Modeling the Civilian Environment in Gaia 9

4 Facilitating Process and Doctrine in the Platform 17

5 Legal Review Manager and Target Auditability 24

6 AI for Civilian Harm Mitigation 27

7 Conclusion 29

Actual results and experiences may vary. © Palantir Technologies Inc. Any data contained herein is notional

1 Introduction

Software is an essential tool of modern warfare. Militaries must operate in and through digital technology systems to competently and responsibly perform their roles. Technologies that facilitate the collection, organization, management, and analysis of information and intelligence increasingly serve as critical infrastructure to support operational decision making on everything from resource management to force readiness to kinetic strikes. While this has been the case for decades, recent years have seen significant developments in the sophistication of these systems and a growing role of artificial intelligence (AI) in analysis and decision making.

Palantir builds decision support system (DSS) software that is used in many industries and contexts. DSS technologies used in military targeting processes must be designed and configured with particular care to support operators as they navigate difficult, time sensitive decisions, balancing mission objectives and doctrinal, legal, and ethical obligations to protect non-combatants and civilian infrastructure. Ultimately, the decisions to act are made by the users of our software, but the tools Palantir provides must facilitate those decisions clearly to account for nuances in military doctrine and process, while still holding a high bar for presenting relevant information about the civilian environment, using all the data to which users feasibly have access.

What is a Military Decision Support System

A DSS implemented for defense is a system that is configured specifically to aid military decision makers as they conduct mission planning and targeting operations, among other highly sensitive and consequential actions. Compared to DSS offerings Palantir and others provide to non-defense customers, military DSSs must meet and account for additional standards, contextual requirements, and use cases:

- **Technical requirements:** Militaries possess tremendous amounts of data, generally spread amongst a number of databases and systems, often in different formats, and operating in disconnected environments. Just managing data at the scale and variety of military information sources is a challenge, and that's before considering that data must often be held in compliance with strict classification requirements.
- **Decision types:** Militaries bear the grave responsibility of making decisions that other organizations do not, and should not, ever encounter. The processes a military uses when deciding when and how to use lethal force are unique.
- **Doctrinal and procedural measures:** Militaries have protocols and standards that articulate requirements for conducting warfighting activities and that also must be rigorously accounted for in the course of operations supported by DSS capabilities.

Actual results and experiences may vary. © Palantir Technologies Inc. Any data contained herein is notional

Palantir's Investment in Civilian Protection Capabilities and Workflows

Before walking through the product features themselves, it is important to ground this document in the philosophy behind Palantir's investment in civilian protection capabilities.

As the prevalence and sophistication of technological systems used in defense grow, we believe there is a serious need to modernize capabilities that detect and protect the civilian environment, in line with International Humanitarian Law (IHL) and existing military doctrine. Builders of the software and other systems that contribute to defense missions, such as Palantir, have a responsibility to consider the ways in which the use and misuse of their offerings can both enable and mitigate civilian harm. While respecting the fact that states hold the ultimate responsibility for operational decision making, including the decision to use any given technology in a given situation, vendors operating in this space should proactively develop features and capabilities explicitly intended to help decision makers identify and mitigate risk to civilians. Existing doctrinal practices should be well accommodated but should be seen as a floor rather than a ceiling for design and development considerations; vendors should use their technical expertise to explore and propose novel capabilities that militaries can use to adapt and supplement existing procedure.

We see this as a continuing and perpetual investment. One that began years ago and will continue as our products, our partner organizations, military doctrine, and the state of conflict evolve. We are committed to constantly pushing the art of the possible with respect to civilian protection capabilities in DSSs, as well as transparency and dialogue around the existence of those capabilities to not only our customers but also external parties in civil society and policy-making communities. Each participant in this lifecycle - militaries, governments, industry, civil society, and more - have an important role to play in the continual development and strengthening of civilian harm mitigation in modern warfare.

For more, see our written publications on how we view the compatibility of working in defense and upholding fundamental rights, as well as the role of ethical AI in defense.

Goal of this Whitepaper

This whitepaper intends to detail many of the capabilities Palantir builds into its military DSS that support civilian harm mitigation, noting the ways in which those features promote compliance with principles from IHL. As these capabilities will change and grow over time, we will seek to update this resource with future versions, in line with the transparency goals we outlined above.

With this whitepaper, we hope readers can grasp tangible product capabilities that seek to support civilian harm mitigation for the military users who ultimately use these systems.

Actual results and experiences may vary. © Palantir Technologies Inc. Any data contained herein is notional

All data and images throughout this whitepaper are notional. Locations are not selected with any relevance to ongoing military operations, but rather purely to demonstrate product capabilities in a notional context.

2 Civilian Considerations in the Ontology

The Palantir Ontology System

The Ontology system - "Ontology" in the rest of this whitepaper - sits at the heart of all of Palantir's platform architecture. An Ontology models the complex, interconnected data that an organization possesses and our software platform helps organizations design an Ontology and use it for operational decision-making.

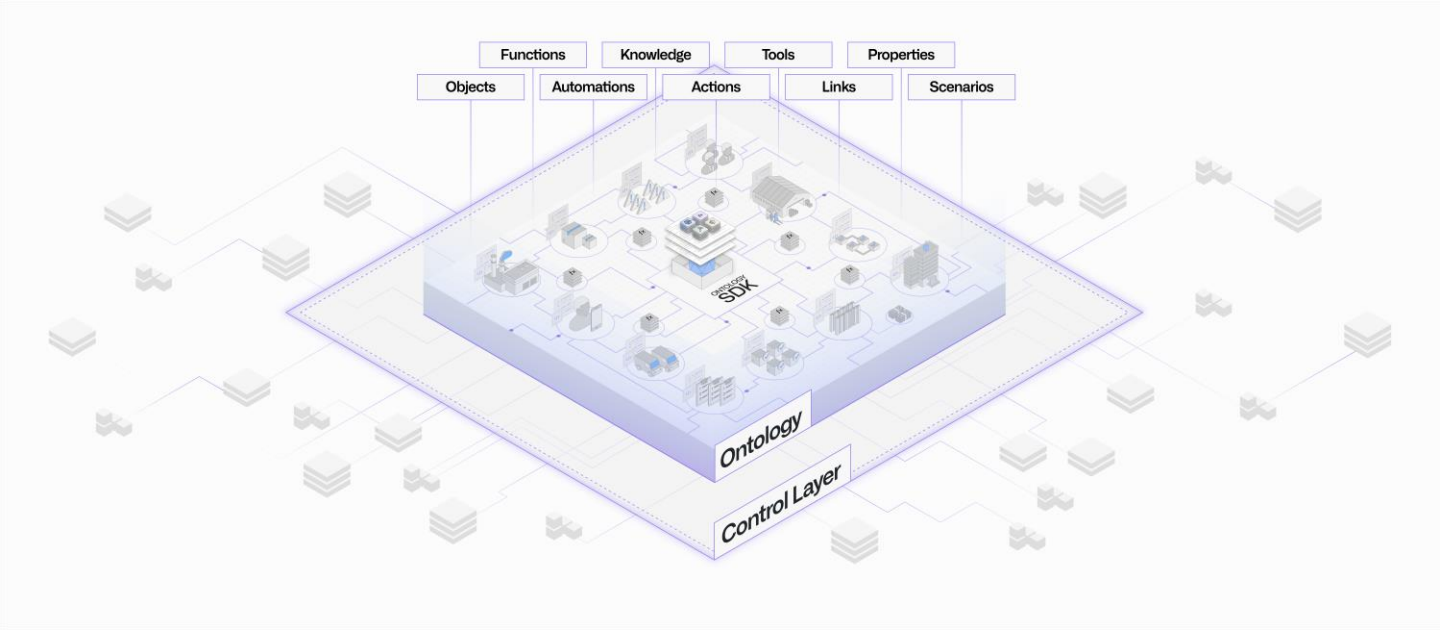
The Ontology models decisions through the four-fold integration of data, action, logic, and security.

- **Data:** Data can flow into an Ontology from every conceivable source, such as fragmented databases, organization-specific systems of records, geospatial repositories, real-time sensors, document stores, and essentially any other digital alcove. The Ontology unifies and maps an organization's disparate data sources into understandable objects, properties, and links. For example, in the military context, Object Types might represent concepts like Units, Munitions, and Facilities.
- **Action:** While the data models the "nouns" in an organization, actions model the "verbs" or the decisions that a user can take in the platform. Taking an Action in the platform can not only update the Ontology to reflect the outcome of the decision, but also interoperate with external systems to carry out the physical effect of the action as well. For example, in the military context, Action Types might describe decisions like tasking a satellite, nominating a target, or dispatching a unit to a location.
- **Logic:** The logic that powers each action can be modular and evolve over time, reflecting the diversity of calculation and reasoning that drives decision-making. The logic underlying a given action (or enhancing a particular object) could be a simple business rule, a conventional machine learning model, an LLM-driven function, or a complex multi-step orchestration that involves several compute engines. For example, in the military context, Logic of an Ontology could include a CV model for Imagery Analysis, an LLM-backed Agent for classifying OSINT, or a rules-based procedure for finding nearby No Strike List (NSL) entities from a verified data source.
- **Security:** Security is woven directly into the data, actions, and logic of the Ontology. Individual Objects can hold role-, marking-, and classification-based access controls, enabling the granular use and sharing of data. Users can configure Logic with scopes to restrict what data a function can take as an input and return as outputs. Actions can be configured with specific submission criteria to restrict who is able to execute certain decisions under which

Actual results and experiences may vary. © Palantir Technologies Inc. Any data contained herein is notional

circumstances. Security controls are built directly into the Ontology, offering security and privacy by design directly in the Ontology itself.

The Ontology is the foundation for many of the capabilities in our Defense platform, including those described in the following sections of the whitepaper.



Ontology for Civilian Protection

The Ontology itself is uniquely valuable for civilian protection workflows, where information about the civilian environment is typically complex and fragmented. Military-managed no-strike lists, military-collected intelligence, partner forces and humanitarian organizations and OSINT – all storing data in different ways – are among the potential sources of information relevant to the discharge of a military's objectives and IHL obligations. The organization of the data contained in these sources directly impacts a military's ability to use said data both to accomplish its operational objectives and to take meaningful precautions to mitigate civilian harm.

Integrated Data Foundation

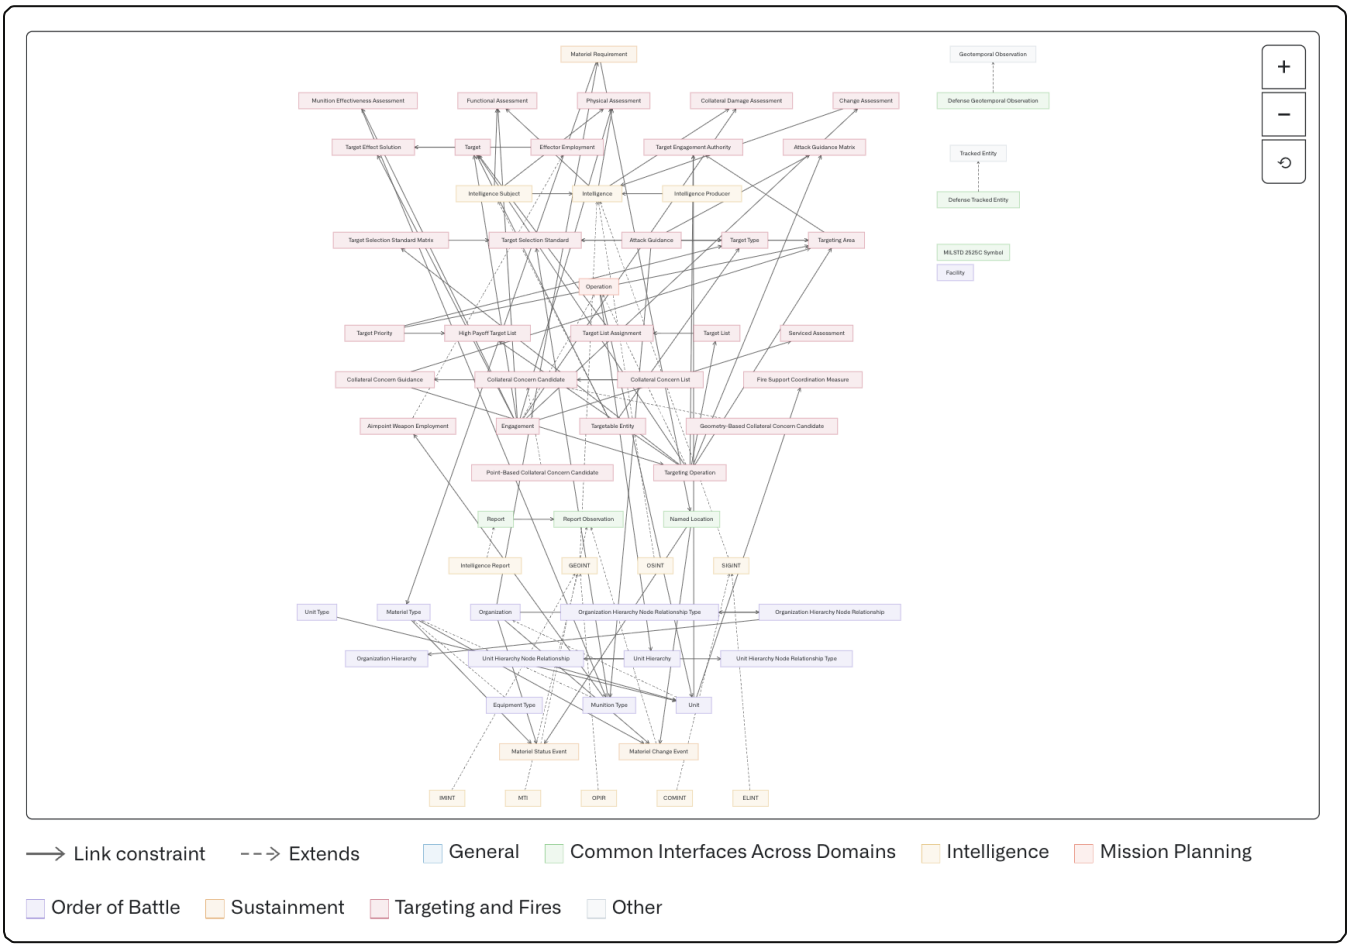
The Ontology provides a governed data foundation for integrated information about the friendly, adversary, and civilian environment. The platform supports integrating data of varying modalities (tabular, streaming, unstructured, multi-media, etc.) from a variety of different data sources. Once the data is ingested into the platform, it can be cleaned, transformed, resolved, and structured in a way that facilitates better exploration, analysis, and operational use for decision-making. As we describe in Section 3 below, data about the "civilian environment" can come from fragmented data sources and

data modalities, and the data integration capabilities of the platform can help Intelligence Officers, Civil Affairs cells, and others in putting together a more comprehensive picture of civilian activity. The data integration capabilities in the platform also come with tools for data health, freshness, release management, and governance to ensure that users can validate the provenance, freshness, and quality of data synced into the Ontology.

Once the data is brought into the Ontology, users can more readily leverage this data for their designated analysis and operational workflows, promoting the re-use of existing, validated sources of data.

Enabling Builders through the Defense Ontology SDK

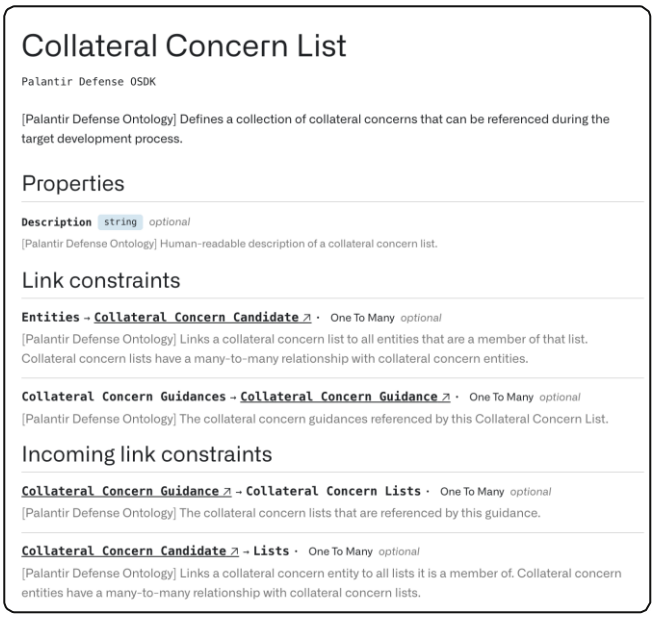
An Ontology serves not only as a foundation for analytical and operational workflows in the Palantir platform, but also as a set of building blocks that users can leverage to develop their own, custom software applications. The Defense Ontology is Palantir's standardized data model for concepts across warfighting functions. Because this model follows a "closed for modification, but open for extension" maxim, Palantir can develop applications that draw on data from this model while allowing customers flexibility to model additional, context-dependent information.



Actual results and experiences may vary. © Palantir Technologies Inc. Any data contained herein is notional

Because this standard Defense Ontology is just a consistent set of definitions, customers can use any of the application building tools in the Palantir platform to build custom workflows and capabilities around this common data layer. Additionally, with the Ontology Software Development Kits (OSDK), developers can build custom applications and workflows using other tools.

The Defense Ontology has native support for key Civilian Harm Mitigation concepts. These object types not only power capabilities in platform applications, but are also available for users to adopt in their own custom workflows. For example, the Collateral Concern Candidate interface in the Defense Ontology models entities that are on a No Strike List (NSL), Restricted Target List (RTL), or any other custom list that targets may need to be deconflicted against in the targeting cycle. The Defense Ontology defines a standard set of fields that may exist on such a NSL or RTL entity as well as the links it has to other Ontology object types.



This offers developers two benefits: First, application developers can reason about entities from the Defense Ontology without needing to contend with their specific implementations (similar to the concept of interfaces in Object Oriented Programming). For example, NSLs, RTLs, and other custom collateral concern lists may have different schemas and formats, but an application developer can reliably know that these lists will have certain properties and linked object type because they all extend the Collateral Concern List interface from the Defense Ontology. Second, by applying this interface on object types in an Ontology, users can mark custom object types as collateral concerns, which means platform-level features (described in later sections) can benefit this custom and heterogeneous data without additional configuration work needed from the customer or Palantir. What this means is that the Ontology enables not only information and doctrinal consistency, but also the flexibility to extend critical civilian harm reduction actions across a dynamic decision-making environment.

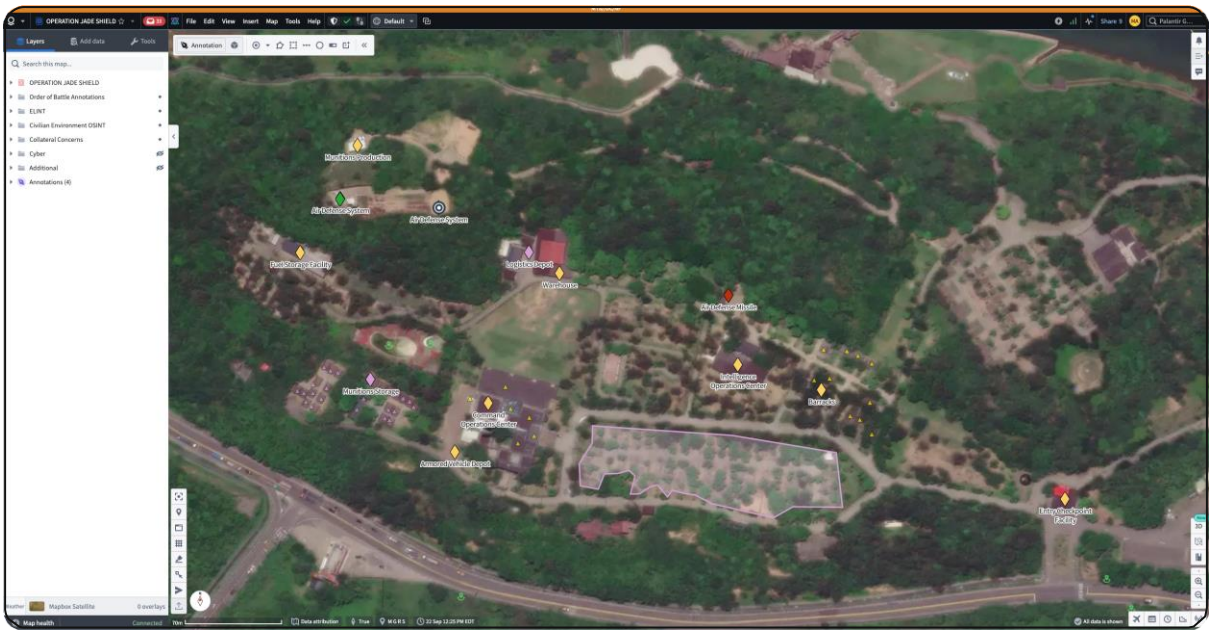
The Defense Ontology allows builders to interoperate with the platform in a standardized, structured manner – and one that reinforces best practices for Civilian Protection. The full Defense Ontology documentation can be found online here.

Actual results and experiences may vary. © Palantir Technologies Inc. Any data contained herein is notional

3 Modeling the Civilian Environment in Gaia

Supporting Intelligence Workflows

A core part of the Palantir defense product suite is Gaia, a geospatial application for intelligence analysis and mission planning. Gaia sits over a foundation of integrated data, allowing quick querying of historical and strategic data sources in a single interface. The civilian considerations from the Ontology, detailed in the previous section, can now be thoroughly represented in a geospatial tool, even if their data structure varies. Intelligence Officers use this product for modeling intelligence and planning various operations, seen in the example below.



NOTIONAL - Mission planning around an adversary base

In this notional example, Intelligence Officers can build out a targeting environment for potential operations, represented by the diamonds (standard military symbology), ranging from targets of interest to potential aim points for strikes. Users can configure their own relevant military symbology accordingly. These potential targets are carried through the targeting lifecycle before any action is taken as determined by the relevant doctrine and process of a combatant command (or non-US equivalent). The targeting lifecycle will be covered later in this whitepaper.

Geospatial tools are helpful starting points for visually understanding the civilian environment. However, the data reflecting the civilian environment is typically housed across different systems, if it exists in a structured format at all. Analysts may be using PPT slides, written assessments, graphics, human intelligence, or other analog means during their intelligence gathering and targeting validation processes. Logging into separate systems can lead to missing essential data in a common operating picture (COP).

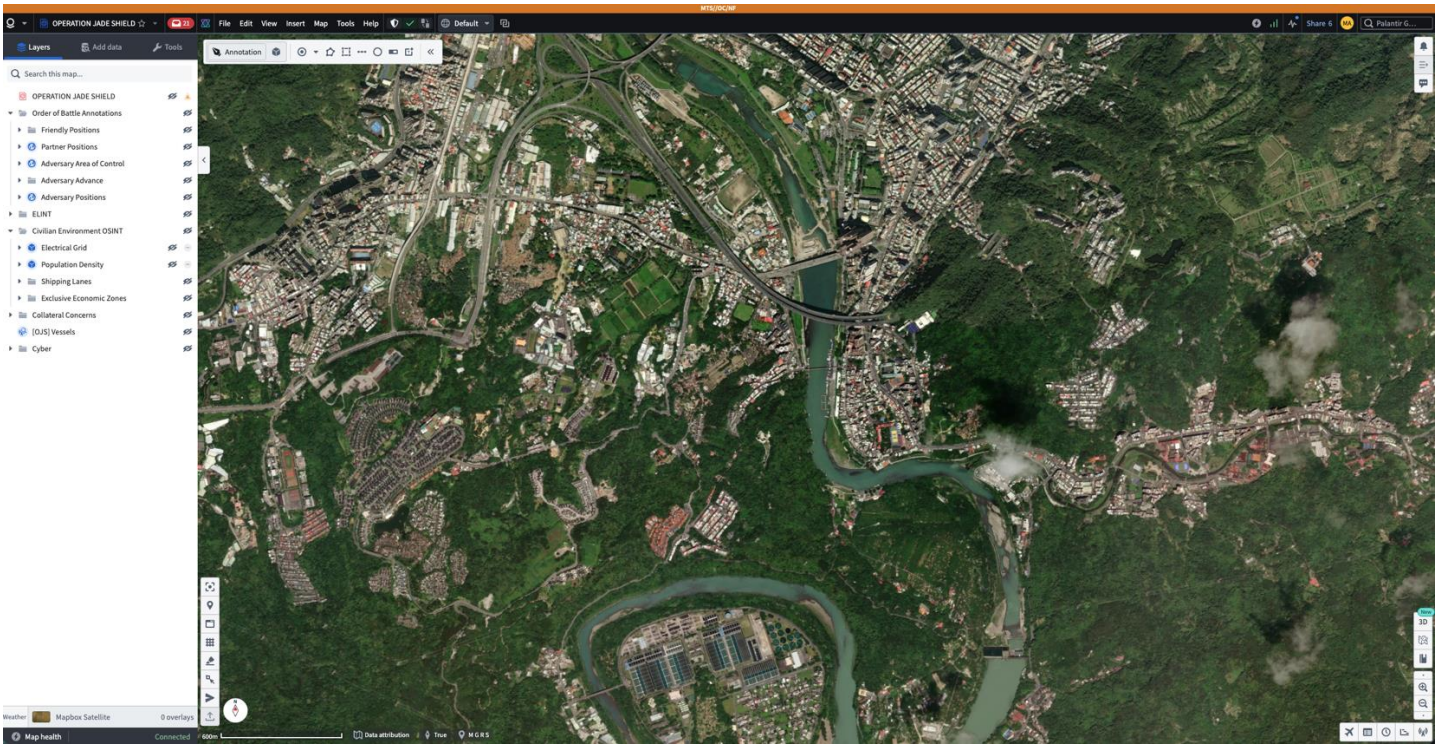
Actual results and experiences may vary. © Palantir Technologies Inc. Any data contained herein is notional

By assembling and layering this data into a single interface, there is a meaningful step change in the capacity to understand the reality on the ground. The composite view helps equip human operators with more comprehensive access to their information as they work to fulfill their own obligations under existing doctrine and IHL.

Layering Civilian Infrastructure in a COP

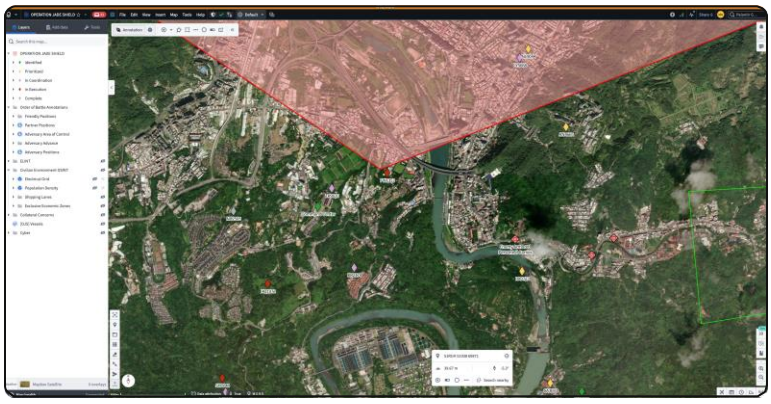
Gaia is configurable to allow for the integration and visualization of any geospatial data a customer owns and models in their Ontology.

Analysts start with a selected map layer of the relevant geography, represented in the screenshot (right). Their primary view is the environment in question, with the various data sources to which they have access represented as layers on the left. Users would want to begin populating this environment - whether with potential targets of interest, civilian infrastructure, or any other data relevant to their current operations.



NOTIONAL - Gaia blank map view

Actual results and experiences may vary. © Palantir Technologies Inc. Any data contained herein is notional



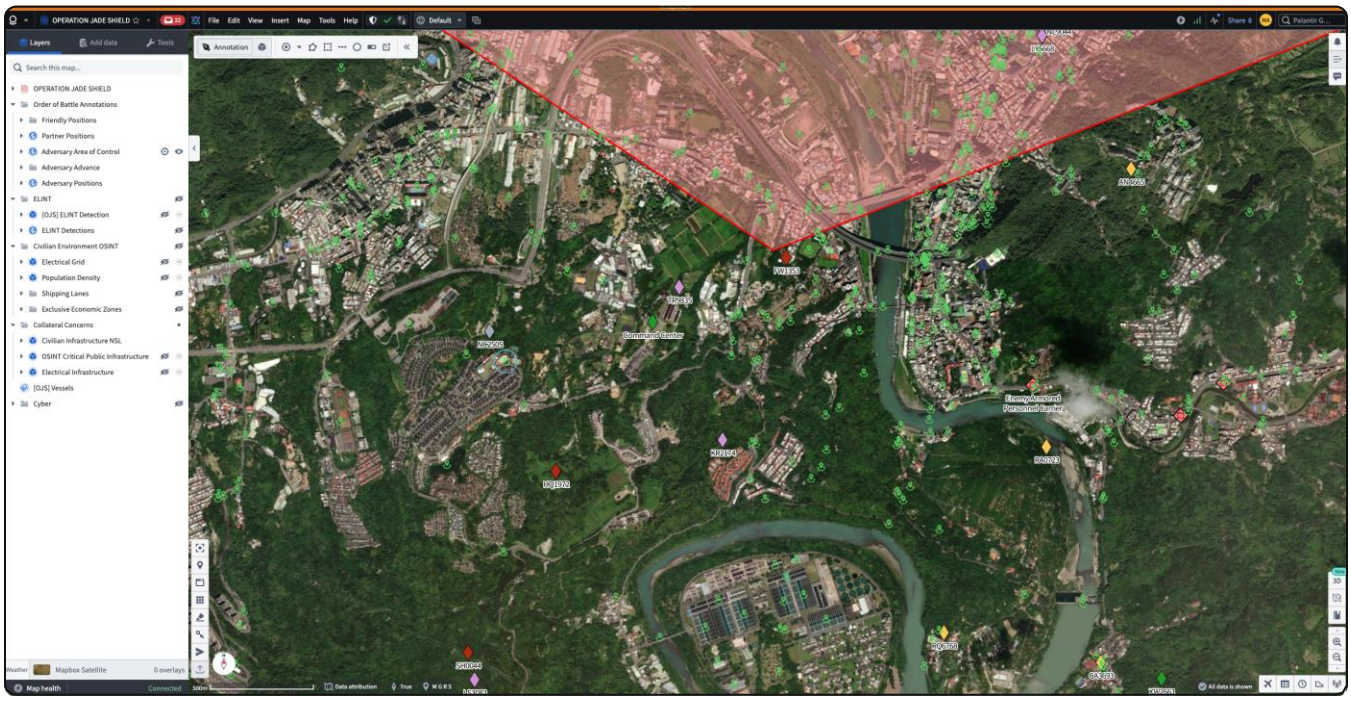
NOTIONAL - Gaia with conflict environment layers

This initial map view reflects the data layers that represent the conflict occurring on the ground, including areas of advance and control, and notional targets, reflected by the diamonds, that would exist across different stages of the targeting lifecycle. This overall view is helpful for understanding potential targets and the associated environment. However, users need to be able to understand the data representing civilian environment, alongside the conflict environment, all in the same COP.

Because information about the civilian environment is just another form of data, users can readily layer this context alongside the targeting environment using familiar Gaia tools and processes. Intelligence Officers and other users can then benefit from this added layer of information as they conduct assessments and identify potential targets of interest. For example, if a No Strike List (NSL) has already been added to the map, the analyst can toggle on that layer, seen in the green points in the screenshot below. A NSL is generally a list of known civilian or noncombatant entities maintained by a military, and typically includes important entities in the civilian environment, such as hospitals, critical infrastructure, and cultural sites that are protected from the effects of military operations under international law and/or rules of engagement. For example, US doctrine mandates that combatant commands maintain NSLs that are consulted during targeting. There are equivalent processes in other military law jurisdictions.

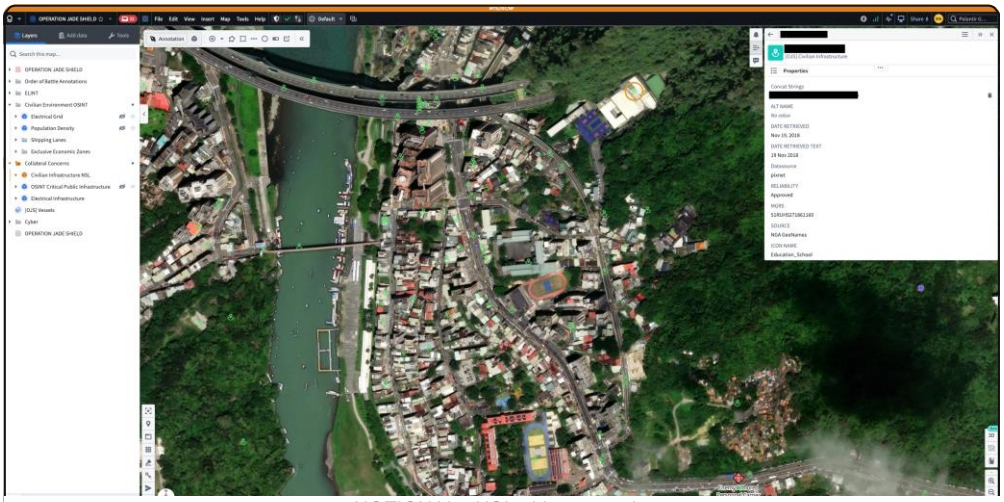
It is important to note the role of NSLs in this process. NSL entities do not represent the entirety of the civilian environment that commanders must take account of. Instead, NSLs are explicitly identified, often high priority features of the civilian environment. There may be other features that — especially in urban conflicts — can be inferred through additional information layers, such as population density maps. The ability to visualize population density, pattern of life, and other dimensions of the civilian environment (where the data is available) further enriches understanding of the civilian environment. Because NSL-shaped data and workflows are somewhat standardized and commonly used, there is tooling to directly support their use. These features are not intended to discourage, and in no way prevent, other procedures and tools from being used to mitigate civilian harm.

Actual results and experiences may vary. © Palantir Technologies Inc. Any data contained herein is notional



NOTIONAL - Civilian infrastructure geopoints

While this specific data example is for a notional No Strike List, this is similar to how any NSL data source, or other data source reflecting the civilian environment, can be integrated into COPs.



NOTIONAL - NSL object metadata

Users can select the NSL entity itself to see associated metadata and better understand data attributes, provenance, and potential staleness. The properties presented here, while served from the Ontology data later, are sourced from additional systems the NSL owner already controls.

Users can further enrich this view by reflecting additional datasets. In this case, they might have access to additional open-source intelligence, or OSINT, that reflects the civilian infrastructure as shapes with a fuller outline, instead of individual points on a map. This may be especially valuable in delineating boundaries in urban environments.

Actual results and experiences may vary. © Palantir Technologies Inc. Any data contained herein is notional



NOTIONAL – OSINT geoshapes



NOTIONAL – OSINT geoshapes

Users now benefit from multiple sources that not only show a fuller volume of entities, but also a more accurate representation of the boundaries of individual entities. They can verify the NSL point data with additional intelligence. These lists might also represent different data sources - one from an official classified system, the other from OSINT sources that have been directly integrated from the public internet.

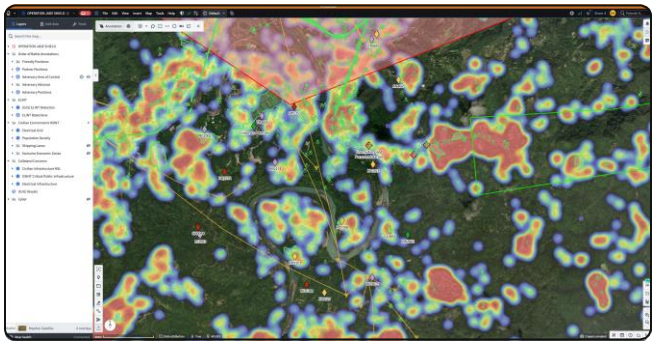
In these notional examples, there are several no strike entities that are captured on both lists. However, each includes a number of auxiliary buildings that may have been missed in the single geopoint view from the NSL. On the other hand, the OSINT data source list also lacks a number of the entities on the NSL. There will never be a single data source that encompasses the entirety of the current civilian environment. The reality on the ground is far too dynamic, so the reflection of that reality has to be as dynamic as possible as well. Without this kind of data integration and visualization capability, analysts would need to log in to different systems or manually compare these distinct lists. The flexibility of Gaia in accommodating a range of data sources into one view, despite the sources themselves coming from different systems, is a major advancement in improving operators' understanding of the ground truth and scope of the civilian environment.

Flexibility of Data Rendering

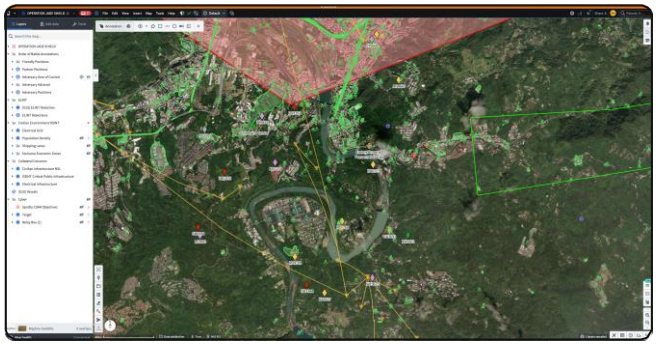
The data modeled in the Ontology will take many different structures and forms. The platform is built to be maximally accommodating to allow for the integration and visualization of any data source, regardless of data structure.

Gaia can display live data feeds to provide users real-time situational awareness of rapidly evolving operations, or even one-off uploads of new data that becomes relevant to the operational environment. In this Gaia Map, users can also layer on some additional sources of intelligence to better understand this environment, including the civilian environment. Some notional examples of platform flexibility include:

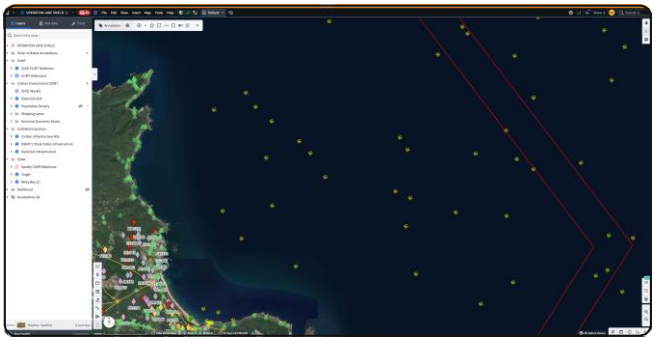
Actual results and experiences may vary. © Palantir Technologies Inc. Any data contained herein is notional



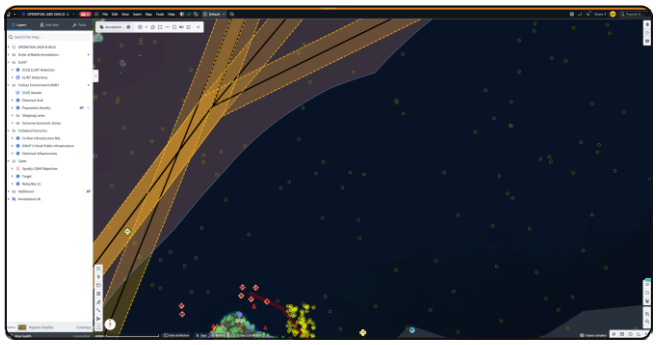
NOTIONAL - Population density



NOTIONAL - Electrical grid data



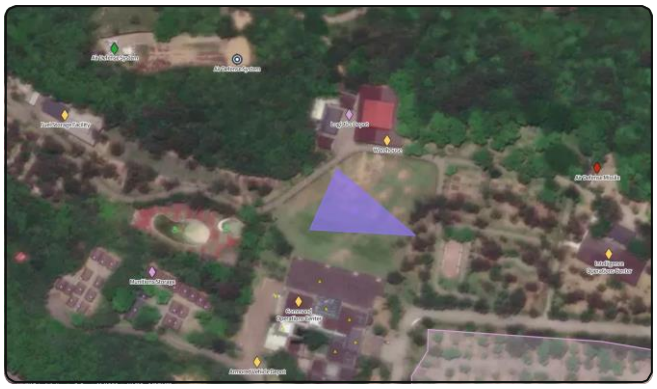
NOTIONAL - ELINT Live data



NOTIONAL - Shipping lanes



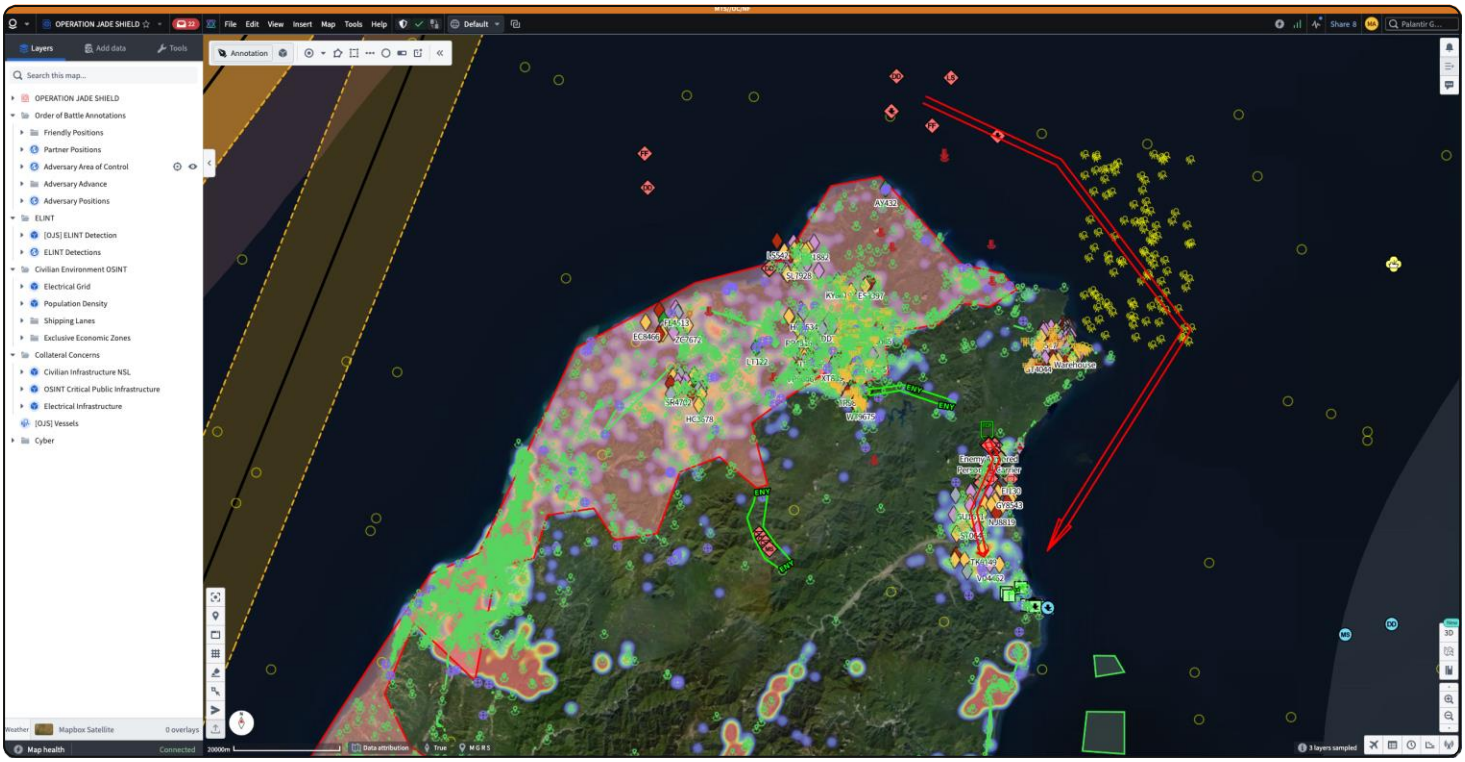
NOTIONAL - Cyber data



NOTIONAL - Manually-drawn shapes

Actual results and experiences may vary. © Palantir Technologies Inc. Any data contained herein is notional

In the final screenshot below, users can see the COP with all the existing layers of data in view.



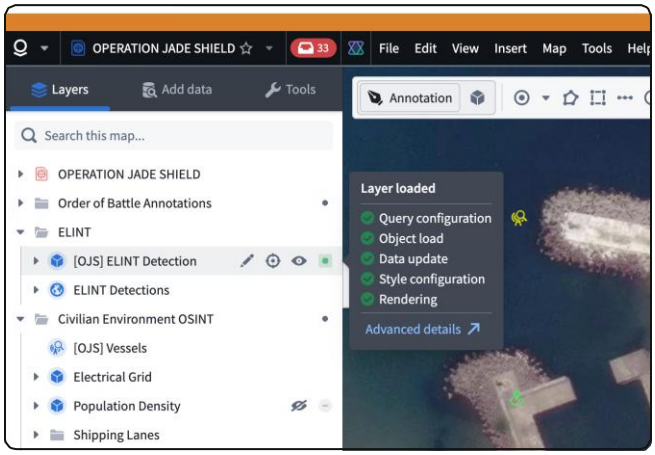
NOTIONAL - COP with all data sources untoggled

Data Refresh and Health Checks

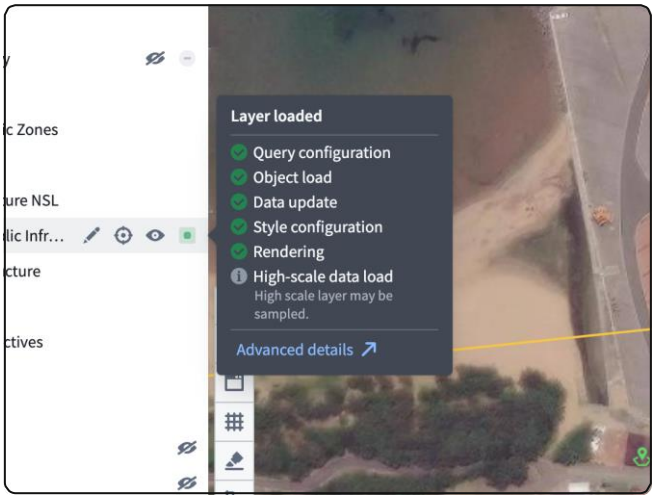
A DSS such as the Palantir Platform must balance capability with institutional control. For militaries, underlying integrated data is theirs as are the procedural requirements for managing and validating that data. The focus and responsibility of DSS platform design is to make available and readily usable the tools and capabilities that support data health, latency, and other technical capabilities to monitor the health of data ingested into the platform. Many of these features exist as core backend capabilities covered in the section above on Ontology modeling, or in features to support data health checks and platform stability. For example, the platform can page users or send alerts if it detects failures in a connection from an upstream data system that could degrade data quality. For more on health checks, see Palantir's public documentation. These capabilities are especially important in settings where builds, transforms, and other data checks inform sensitive decision making, such as in conflict environments.

Actual results and experiences may vary. © Palantir Technologies Inc. Any data contained herein is notional

In addition, there are front-end tools in Gaia that flag if data is not rendering correctly.



NOTIONAL - Data health visuals



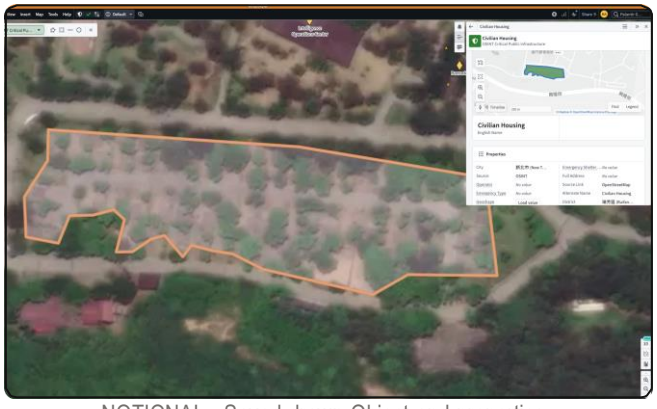
NOTIONAL - Data health flags

Collectively, users can leverage these tools to ensure they are viewing the most recent versions of the underlying data in question.

Drawing from Ontology

Gaia also supports the creation of new geoshape entities directly in the application itself, which can then be written back to the underlying Ontology objects and source datasets. As Gaia and the other applications detailed in this whitepaper leverage the same shared data foundation, these additions will be reflected in any place where those backing Ontology objects are rendered.


NOTIONAL - Drawing new Objects in Gaia


NOTIONAL - Saved drawn Object and properties

In these screenshots, we can see a notional example of an Intelligence Analyst manually drawing a block of civilian housing on a military base, supposing this is new intelligence they recently received and want

Actual results and experiences may vary. © Palantir Technologies Inc. Any data contained herein is notional

to reflect in Gaia and other applications leveraging the notional OSINT data source in question. These tools provide flexibility for users who are often in dynamic situations and need to be able to mark new intelligence in whichever tool they're currently using, knowing it will cascade to other applications that use the same Objects.

Once saved, users can readily interact with the Object and understand associated metadata that is relevant for their operations. As outlined above, metadata is valuable for digging into specific attributes about the entity itself, including properties that indicate recency, source quality, ownership, and more. In this example, users can leverage tools to facilitate the custom creation of new civilian infrastructure based on recently received intelligence.

In conclusion, civilian protection cannot be reduced to a single point solution or data source in a DSS. Gaia's ability to provide a flexible and customizable COP, and users' ability to integrate intelligence on civilian entities into that COP, make it a valuable tool for constituting a fuller understanding of the civilian environment.

4 Facilitating Process and Doctrine in the Platform

Militaries have legal and doctrinal requirements for civilian protection. Further, targeting decisions, and the broader targeting lifecycle, are dictated by existing military law, doctrine, and process. While the structure and flow of targeting lifecycles are appropriately determined in accordance with these military defined procedures and doctrines, a DSS can play an important role in translating them into operational practices configurable to reinforce IHL principles.

A targeting lifecycle incorporates multiple stages that vary by military and combatant commands. For example, in US doctrine, potential targets must be nominated, assessed, checked against multiple checklists and operational requirements including collateral damage assessments, reviewed by legal advisors, assigned weaponeering, and triaged to fires teams before any kinetic action may be approved to take place, among other steps and requirements. Each of these stages, requirements within each stage, users involved in reviewing and approving decisions, and the conduct of reactive governance and oversight must be shaped to align with the demands and expectations of the relevant military and government authorities.

Palantir's platforms are designed to be configurable to allow users to encode their process into the applications they use for their individual roles. This section details examples of how this configuration can be applied in practice.

Actual results and experiences may vary. © Palantir Technologies Inc. Any data contained herein is notional

Collateral Concern Checks in Target Nomination

Analysts can use Gaia to nominate potential targets of interest to initiate review in the targeting lifecycle. The software provides tools to manually and automatically calculate distances to potential areas of civilian infrastructure.

For example, users can manually use ruler tools that exist into the platform to gauge distances between targets and infrastructure.



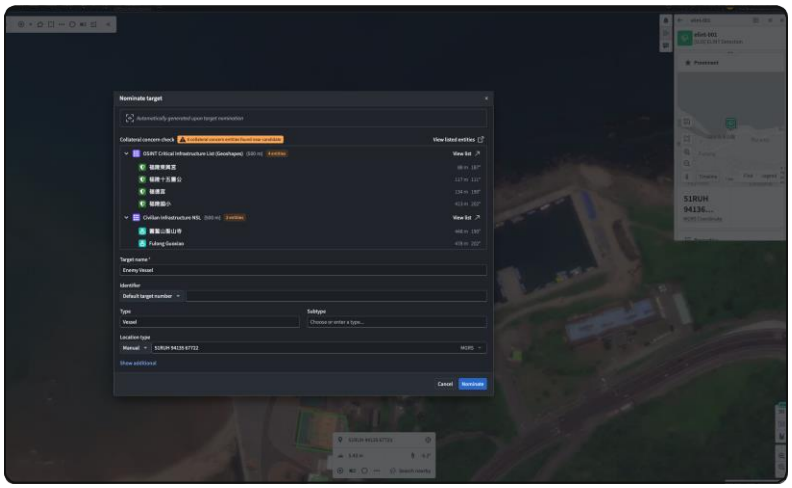
NOTIONAL - Manually measuring distance



NOTIONAL - Manually measuring distance

When nominating a target, users must first populate a nomination window, which prompts for more metadata about the target. This interface will automate the distance calculations to the nearby entities on any relevant collateral concern entities, including the Objects detailed in the previous section.

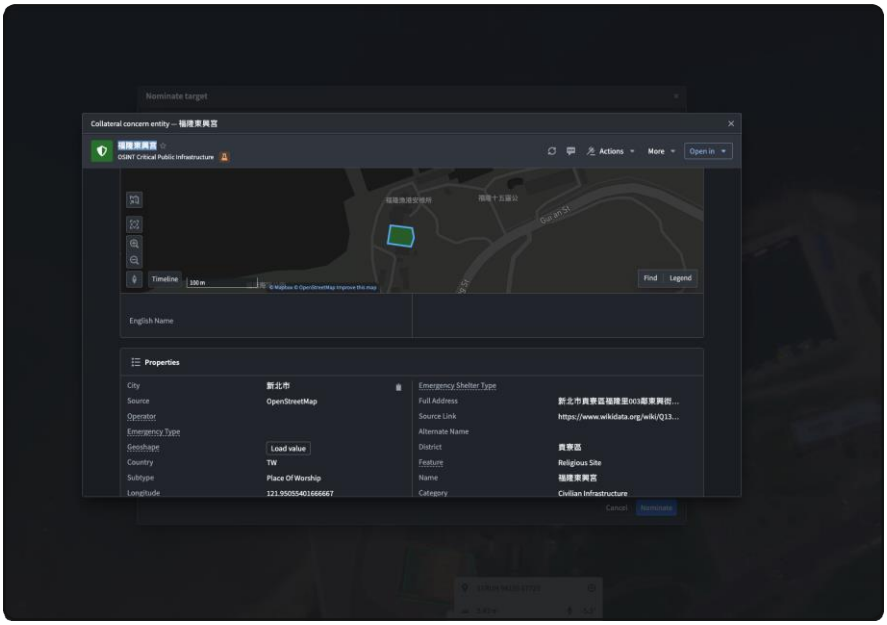
Actual results and experiences may vary. © Palantir Technologies Inc. Any data contained herein is notional



NOTIONAL - CCC in target nomination

Gaia will display both the name and source of the nearby entity, and calculate distance and bearing from the proposed target location. This is called the collateral concern check (CCC).

Users can also dig in further on specific collateral concern entities to understand properties, provenance, potential data staleness, and other data relevant for analysis.



NOTIONAL - Object deep dive in CCC

Collateral concern checks can be configured to run against NSLs, RTLs, or any custom list of collateral concerns – regardless of whether the layer was displayed on the COP in Gaia. The entities that are surfaced, and the radius around the target that will be searched, is configurable. Configurability here is key.

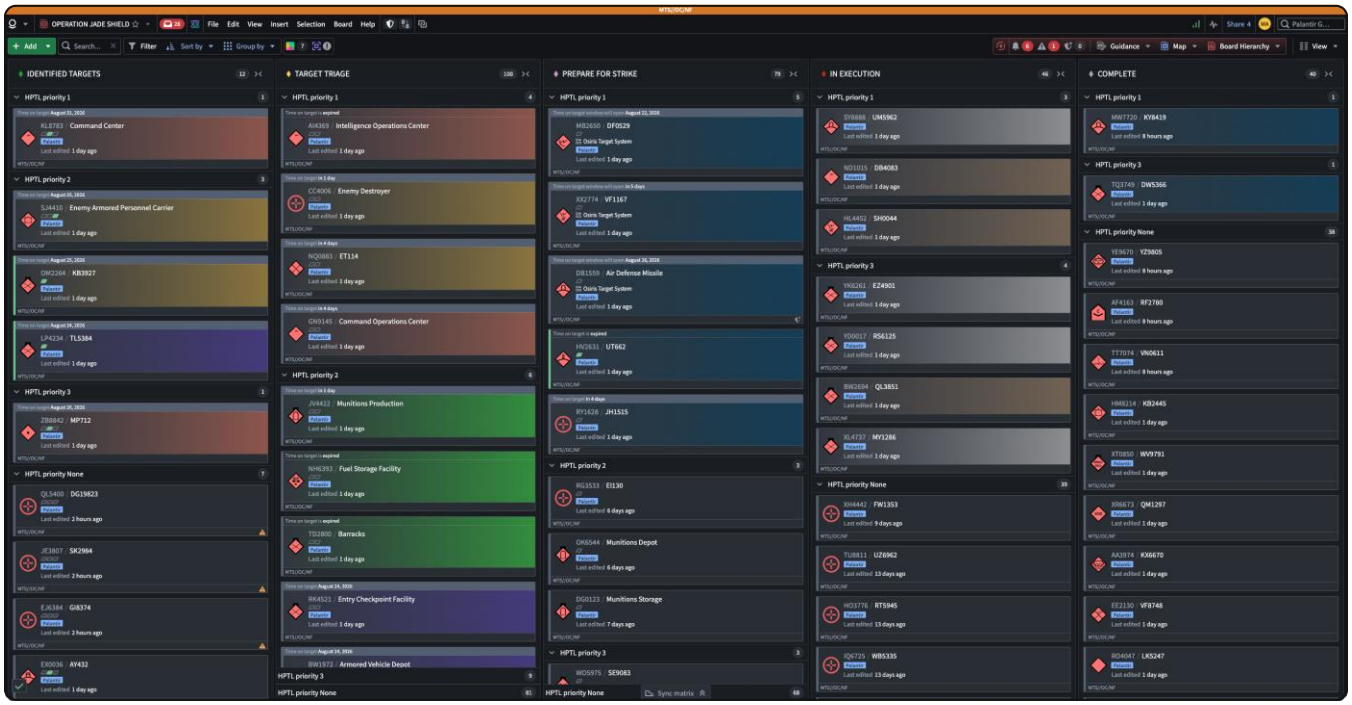
While ensuring that users are considering the civilian environment as part of the COP is vital, this collateral check can help to foreground relevant civilian considerations in this target nomination process. This saves the analyst the time it would take to manually measure out to each NSL entity and enables more consistent application of NSLs each time a target is nominated.

When a target is nominated, Gaia will surface it to Target Workbench, another application specifically

designed to process targets through the Targeting Lifecycle.

Collateral Concern Checks in Target Workbench

Target Workbench, another application in the Palantir platform, is designed to support the management of targets throughout the targeting cycle. Each target is represented as a card, which is gradually advanced across columns to reflect the progressive stages of a military's targeting process. This tool is designed to maximize user configurability - each column can represent a state of the target in the targeting cycle, which contain various activities and verifications that must be performed at these states. For example, US Doctrine specifies that a target is checked against a No Strike List during the targeting process, and that a Judge Advocate (military lawyer) reviews and provides a recommendation on each targeting decision.

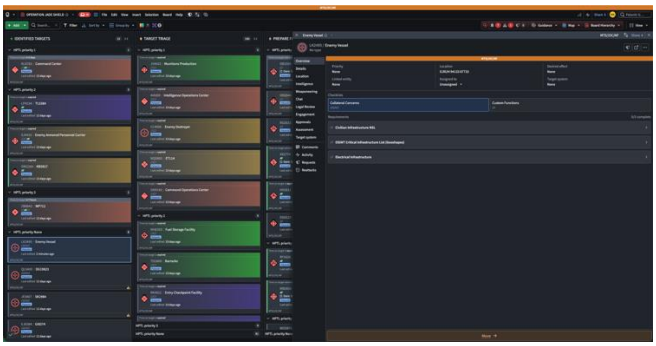


NOTIONAL - Target Workbench

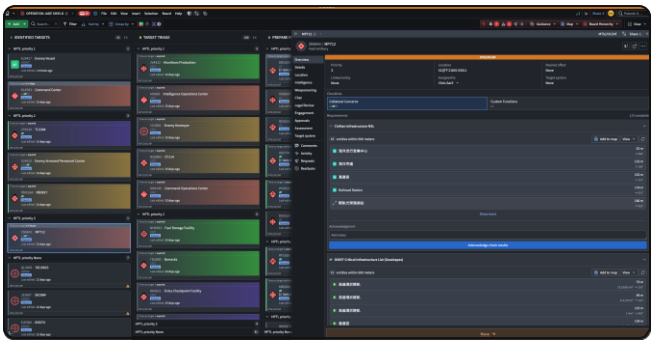
There are a variety of features here - targets are assigned to different users based on their specific roles and purview. Targets are ranked by priority that map to strategic aims of combatant commands. Targets can be assigned time periods before they expire and must be reassessed. Target Workbench is designed to enable users to encode their relevant checks to each stage of the targeting lifecycle, as we'll walk through below:

In the first stage of this notional targeting lifecycle, relevant decision makers have stated that each target must be reviewed against collateral concern checklists. The checklists themselves are populated by the objects seen in the previous section, if they fall within a predetermined radius to the target:

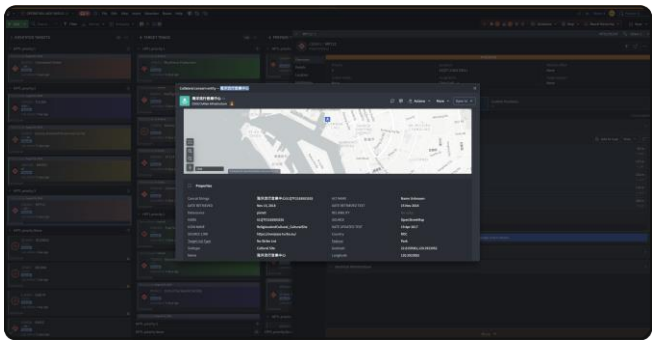
Actual results and experiences may vary. © Palantir Technologies Inc. Any data contained herein is notional



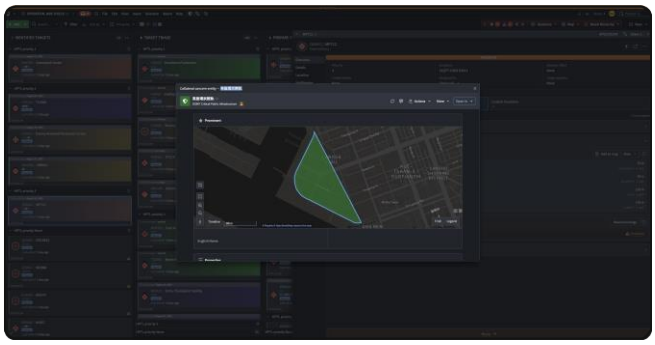
NOTIONAL – Collateral concern checklists



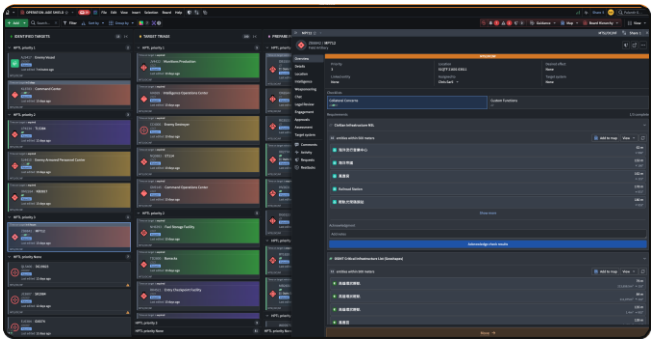
NOTIONAL – Different datasources and proximities



NOTIONAL - Metadata of collateral concern (1)



NOTIONAL - Metadata of collateral concern (2)



NOTIONAL - User acknowledge and logging

Targeteers can review every potential collateral concern within the relevant proximity to the target, better assess the metadata surrounding the entity itself, including source and potential staleness. In this example, acknowledgements must be logged before the target can be moved to the next stage of the lifecycle.

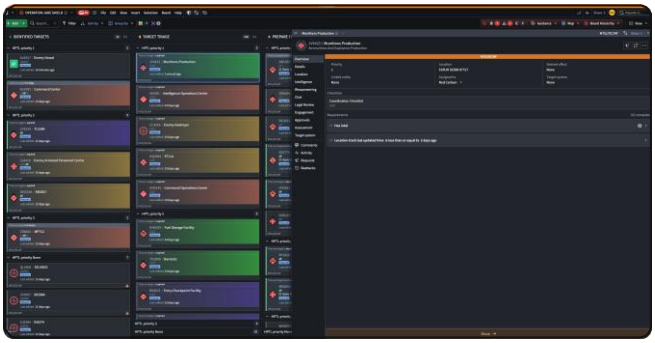
The Collateral Concern Lists themselves, and objects therein, are entirely configurable around the relevant data for this operation, but they can include any data source the entity has access to, even if that underlying structure varies. It should be noted that the utility of this feature is dependent on the quality and availability of the data sources backing the military's representation of the civilian environment.

Actual results and experiences may vary. © Palantir Technologies Inc. Any data contained herein is notional

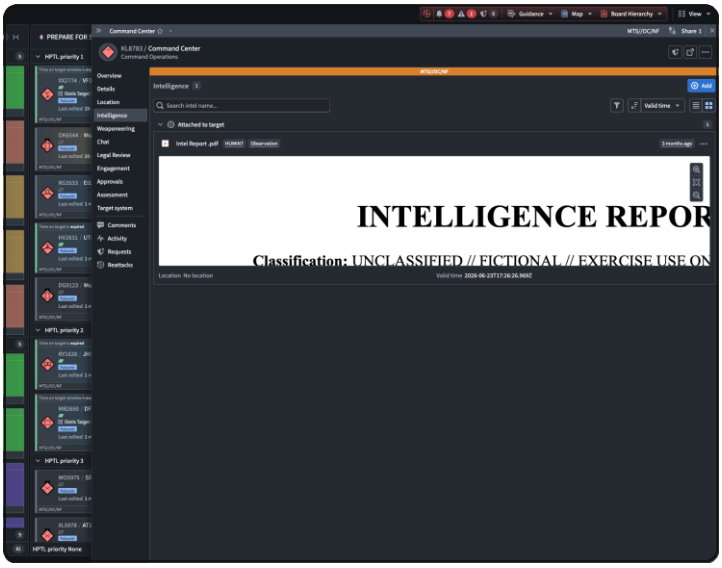
Additional Procedures

Other checks can vary based on the needs of the combatant command (or analogous non-US operational group). The platform is built to accommodate this needed flexibility, as the operational contexts, doctrine, and processes may vary substantially. Some examples of other checks that could be built into the platform include:

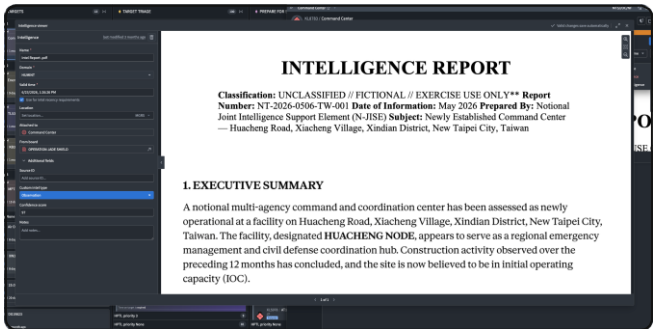
Requirements that intelligence is attached and that the location has been updated within a set time range.



NOTIONAL - Intelligence requirement



NOTIONAL - Intelligence report



NOTIONAL - Intelligence report

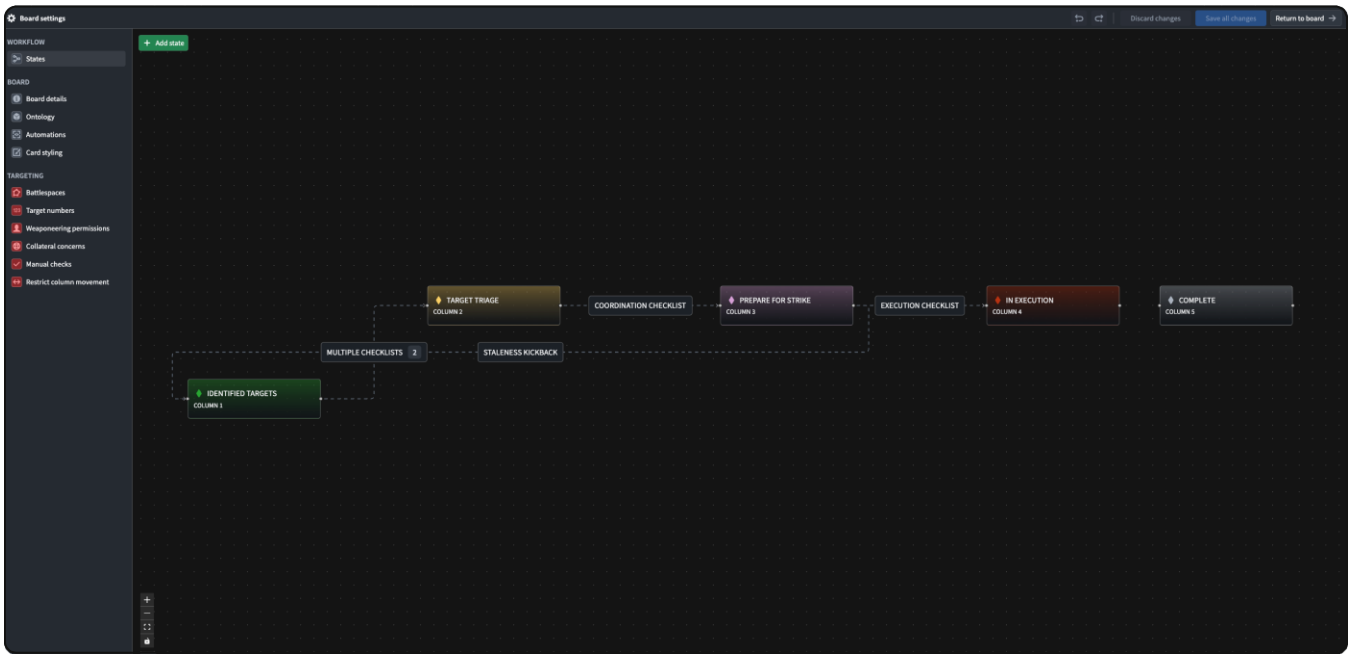
Intelligence reports, often taking the form of long-form written assessments, can be uploaded and attached to relevant targets. Metadata like locations can be derived or extracted from these reports, and expiration dates can also be assigned to automatically flag when intelligence becomes stale.

This reflects the reality that these systems are only as good as the underlying data, and checks should be configured by users in accordance with their procedural standards.

Board Structure & Configurability

The board itself can be easily configured in settings, allowing the relevant CCMDs to structure according to their processes.

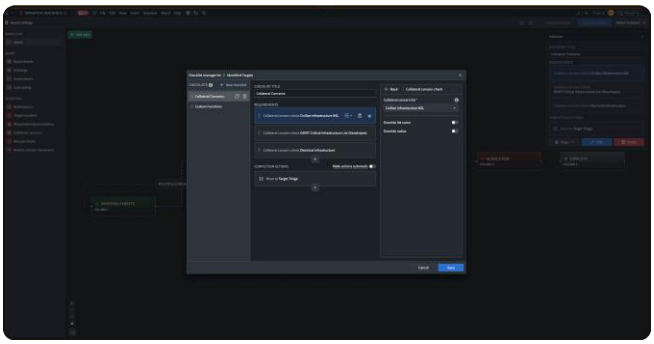
Actual results and experiences may vary. © Palantir Technologies Inc. Any data contained herein is notional



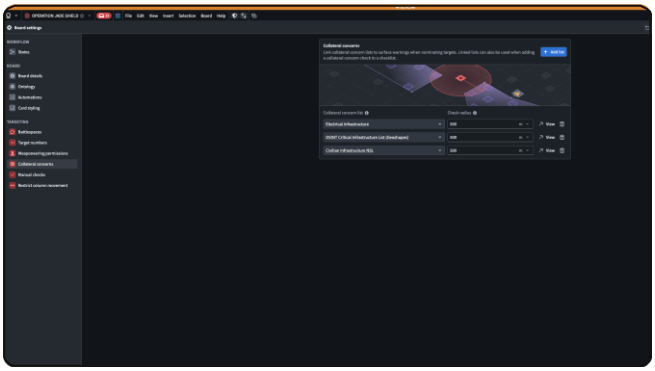
NOTIONAL - Board configuration

This no-code interface structures the columns and stages seen in Target Workbench. Each stage, the corresponding data, views, checklists, and requirements, can be set by relevant users. Additionally, staleness “kickbacks” can be set to automatically move targets backwards in the lifecycle.

As they become available, new data sources can be added as Collateral Concern Lists (left), alongside the relevant radius for consideration.



NOTIONAL - Adding checks into stages



NOTIONAL - Adding new collateral concerns

Those lists, and other checks, can then be readily configured at the relevant states of the lifecycle (right).

Conclusion

There is no single solution for civilian protection in a military DSS. Instead, it is the combination of many engineering, workflow design, and operational choices to build features that military users can leverage

Actual results and experiences may vary. © Palantir Technologies Inc. Any data contained herein is notional

to more reliably and confidently uphold their responsibilities under codes of conduct and applicable laws and regulations. It is the government's role to establish and review its own laws, doctrine, and process surrounding defense capabilities and to determine whether and how to use these tools. Palantir's platform is built to accommodate maximal configurability, so those differing checks can be feasibly encoded into the platform and surfaced to relevant users.

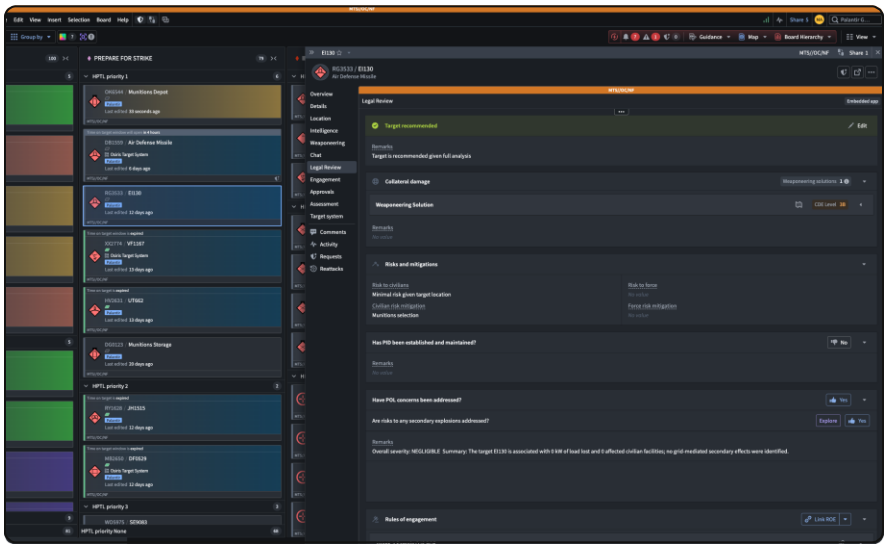
5 Legal Review Manager and Target Auditability

Legal review manager

The platform can be extended by user groups to configure new workflows that operationalize specific processes and user needs, including around Civilian Harm Mitigation, over that same shared source of data (the same Ontology). These can be new standalone applications or custom helpers embedded into applications like Target Workbench.

For example, Judge Advocates (JAGs) - military lawyers - built a custom workflow for capturing a legal review of targeting decisions. JAGs review and sign off on targets during the targeting lifecycle. Without such DSS capabilities, JAGs may be operating in a separate system, or over analog mechanisms, trying to understand the potential target, relevant infrastructure, and then provide and track their approvals and relay those back to relevant operators.

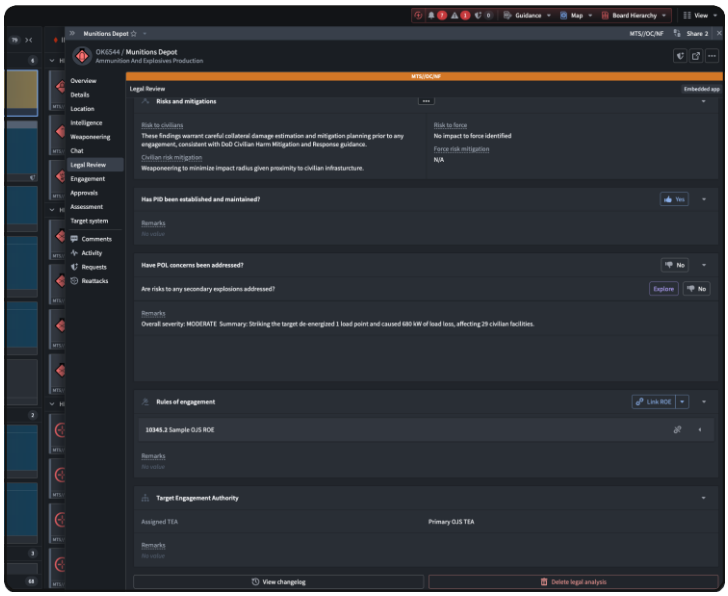
Actual results and experiences may vary. © Palantir Technologies Inc. Any data contained herein is notional



NOTIONAL: Legal review example

Rather than having JAG recommendations stored in a separate system, JAGs can build on the Ontology and create custom workflows that integrate directly into Target Workbench and allow JAGs to review the data from the same pane of glass as operators. This centralizes the key information that each operator needs and writes it back to the same shared foundational data asset in the platform.

The fields of these reviews and recommendations are entirely configurable and dictated by the JAGs themselves to mirror the existing reviews they conduct and to provide flexibility for discretionary variations afforded by the law and demanded by specific combat scenarios.



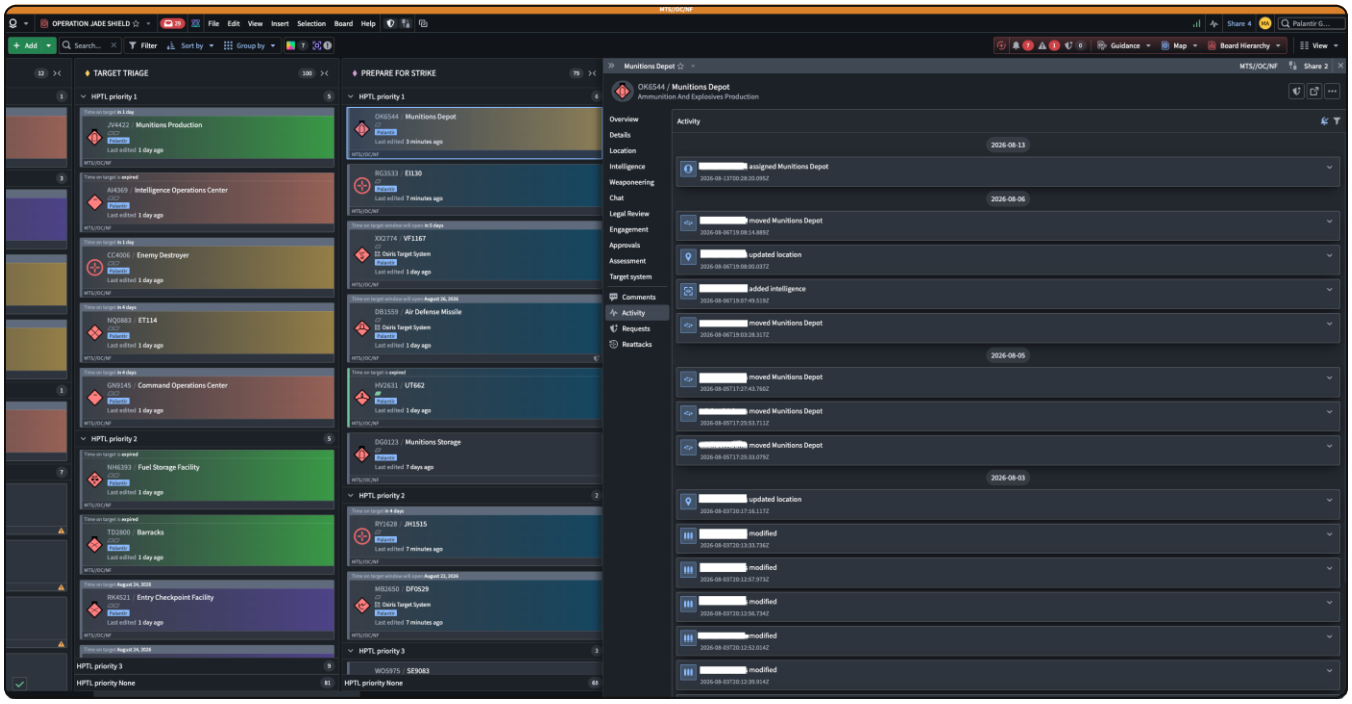
NOTIONAL: Legal review example

Auditability

Palantir's products are built to include an immutable, detailed audit log of every action taken in the platform, enabling relevant security, governance, or legal teams to detect threats, investigate incidents, ensure compliance, and maintain accountability across the platform. Audit logs can be thought of as a distilled record of all actions taken by users in the platform. For more detail on audit logging within Palantir's platforms, read our public documentation.

Actual results and experiences may vary. © Palantir Technologies Inc. Any data contained herein is notional

This system fosters wider, front-end auditability of every change or step the target underwent up to its current point in the targeting lifecycle. This facilitates highly important governance and review in a single tool, which aligns with the existing doctrine and law around after-action reports, and also supports learning and adaptation from past combat scenarios.



NOTIONAL - Auditability of decisions on a given target

Part of the reason why these applications work together so readily is because under the hood they are all reading from and writing to the Ontology.

Conclusion

Thus far, this whitepaper has detailed several capabilities across Gaia and Target Workbench, and highlighted how the COP is built upon a common data foundation – the Ontology. It also outlined how users of the platform can build custom applications with the Ontology, which provides a common data foundation, expanding the platform's utility for an even broader array of roles and workflows, specific to their organization or mission.

Throughout all these features, human oversight, doctrine, and legal process remain integral components of the configuration and use of the platform. By moving away from a world where users leverage multiple, disconnected systems for different parts of their role, this kind of DSS enables military professionals to benefit from the full picture of their data to inform decisions, including a robust understanding of the civilian environment.

Actual results and experiences may vary. © Palantir Technologies Inc. Any data contained herein is notional

6 AI for Civilian Harm Mitigation

In the prior sections, we have discussed the role of software capabilities in DSSs for civilian protection. As militaries increasingly adopt Artificial Intelligence (AI) and AI DSSs, AI capabilities can also be used to improve civilian protection outcomes. This section outlines some of the ways AI can reinforce civilian protection. While not exhaustive, the following case studies demonstrate how different forms of AI may address particular challenges that militaries face when considering civilian harm mitigation practices.

Intelligence and Imagery Analysis

AI can support the integration, cleaning, and structuring of the heterogeneous information sources described in Section 3.

- Large Language Models (LLMs) can be used to analyze text-based Open Source Intelligence (OSINT) about the civilian environment, often performing several steps at once including translation, entity extraction, and Optical Character Recognition (OCR). Purpose-built Natural Language Processing (NLP) models for discrete text-based tasks can also be used in sequence to clean high-scale OSINT data and quickly identify pieces of intelligence that should be rapidly integrated into the civilian environment.
- Visual Language Models (VLMs) and Computer Vision (CV) Models can be used to identify civilian presence in Imagery Intelligence (IMINT). For example, CV models could detect potential civilian presence in satellite imagery, and an analyst could use this finding to potentially supplement a canonical No Strike List or No Fire Area.

The effective application of these AI techniques also requires a rigorous approach to data governance, testing & evaluation, and monitoring of the AI system. The platform offers a breadth of capabilities for Responsible AI practices, including tools for assessing data quality, evaluating traditional AI/ML models, writing "evals" for Generative AI systems, and staged rollouts of AI systems in production settings. Responsible AI practices are crucial for the effective use of AI systems in the military domain, including workflows supporting the protection of civilians. For more, see our public documentation.

Simulation, Scenarios, and Modeling

Operational simulations & Modeling workflows can assist targeteers and military lawyers in understanding the secondary effects of a targeting decision. These secondary or "reverberating" effects are critical considerations under International Humanitarian Law, and militaries need to take feasible precautions to avoid and minimize injury to civilians and damage to civilian objects.

Physics-based models, statistical models, or rules-based simulation models can be integrated into the Ontology to help military users forecast the expected outcomes of their planned decisions. This enables

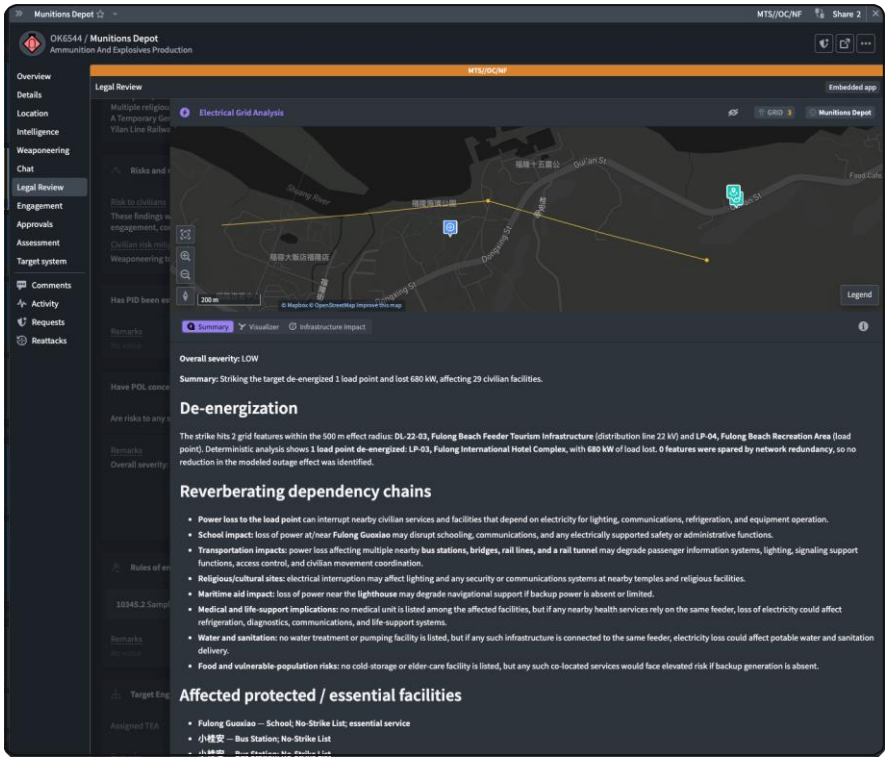
Actual results and experiences may vary. © Palantir Technologies Inc. Any data contained herein is notional

users to proactively evaluate various courses of action.

Moreover, the Ontology natively supports robust capabilities for running these “what-if” scenarios, by creating a sandboxed environment where simulated actions can be run over the Ontology. In this notional example, military lawyers built a simulation tool into their legal review to model the potential impacts of a strike on electrical grids powering civilian infrastructure.

As this data exists in the Ontology backing this board, modeling can pull information from actual power lines and infrastructure in the area.

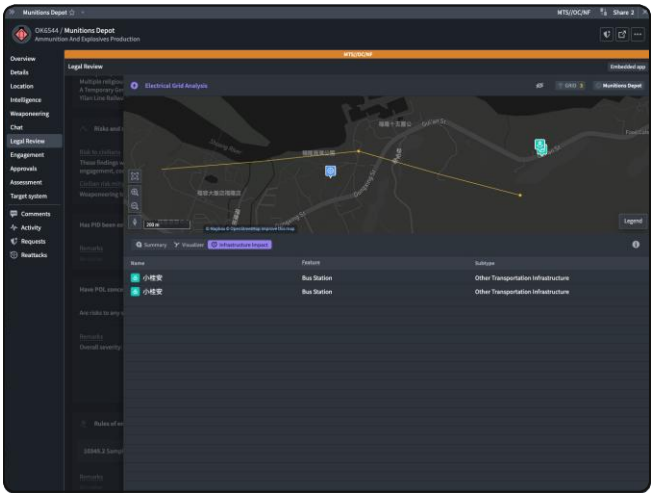
Natural language summaries can be generated around the configurable needs of the users in question.



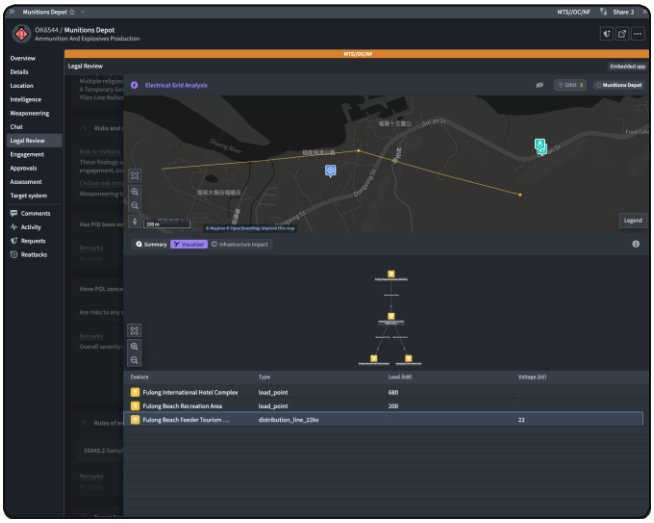
NOTIONAL - Electrical grid impact analysis

Simulations will pull specific objects from the Ontology that would be impacted by a potential strike, allowing for more granular analysis to feed into legal review.

Actual results and experiences may vary. © Palantir Technologies Inc. Any data contained herein is notional



NOTIONAL - Modeling effects and Ontology objects



NOTIONAL - Modeling effects and Ontology objects

Whatever data exists that is relevant can be included here. This is particularly helpful in understanding the often increasing importance of particular infrastructure as a conflict develops, for example if specific pieces of power infrastructure become more integral to an area after the erosion or destruction of other facilities in the same geography.

7 Conclusion

In this whitepaper, we walked through a myriad of capabilities for Civilian Protection in Palantir's DSS platform. There are several core themes underpinning these features:

First, the essential and layered role of humans throughout the targeting lifecycle, especially in surfacing and acting on intelligence that represents the civilian environment. Militaries have doctrine, law, and process - these tools help operationalize those important governance and oversight mechanisms, and reinforce human review across the targeting lifecycle.

Second, the importance of data integration tools for Civilian Protection. All of the capabilities in this whitepaper – from understanding the civilian environment, to collateral concern management, to providing additional tools to augment the Collateral Damage Estimation (CDE) process – require a strong data foundation.

Third, the potential of AI for Civilian Protection. Understanding the civilian environment is no small feat, and there are so many sources of potentially helpful information. But militaries need tools to sift through the noise. There is opportunity to leverage AI in different ways, on top of that common data foundation, to enable Civilian Protection workflows.

This whitepaper also helps to demonstrate how DSSs provide the basis for compounded feedback loops that can carry forward insights and learnings beyond immediate operational contexts. The outcomes and

assessments of targeting decisions should inform future decisions, and this is especially important for Civilian Protection. New intelligence can be dynamically added to a COP which can then inform downstream collateral concern checks and even munitions selection. These kind of connections between workflows are critical for closing the loop between all of the separate workflows and personas who have CHMR responsibilities.

This is a space of continuous product investment for Palantir. As new capabilities become available, we will continue to bring these to the public domain, to the degree possible, to further the important dialogue surrounding technology and defense.

Actual results and experiences may vary. © Palantir Technologies Inc. Any data contained herein is notional