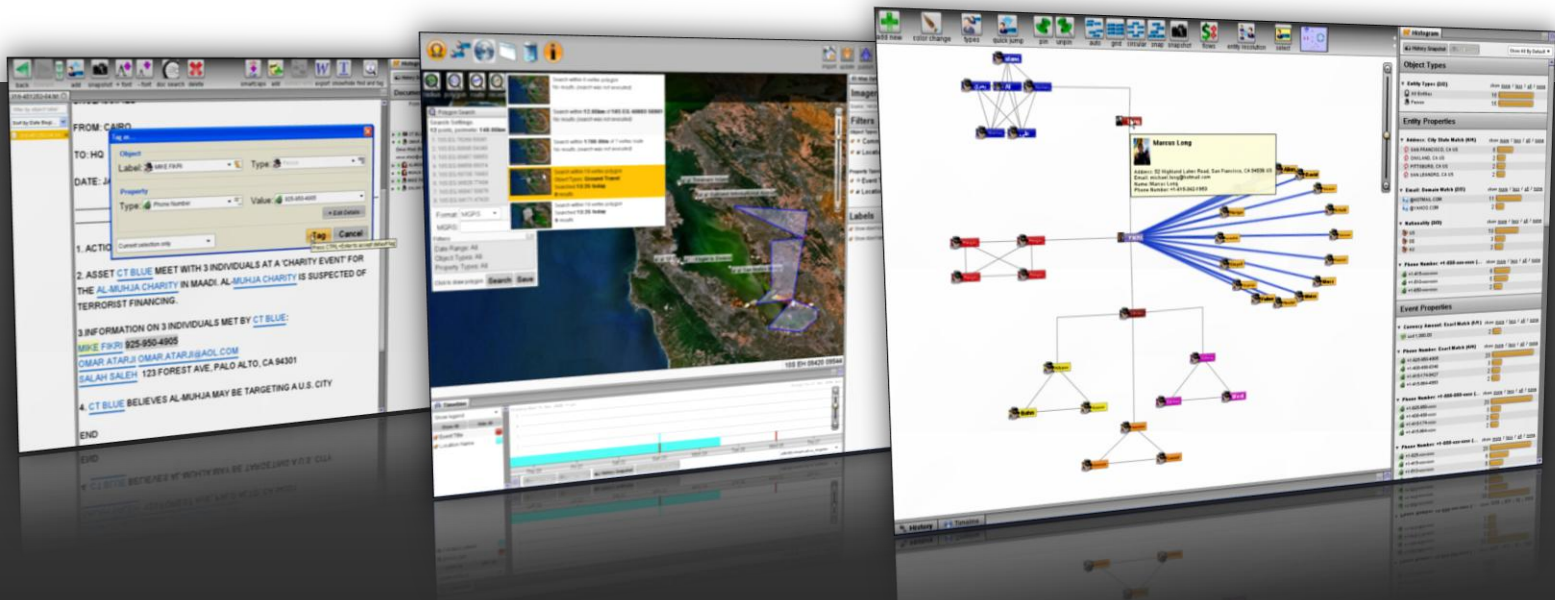# The Palantir Platform
# The Platform for Information Analysis

Palantir Technologies was founded in 2004 based on a simple idea – that the human mind is the most effective way to identify patterns in information while computers are the most effective way to manage enormous amounts of data. By identifying the sweet spot where humans are able to fully leverage their intellect against the computational power of an analytical software platform, it becomes possible to revolutionize our ability to perform analysis.

From this simple idea was born the Palantir Platform, an end–to–end solution for the performance of analysis.

The Palantir Platform

From its inception, the Palantir Platform was designed as an enterprise solution to the analytical needs of the intelligence community.

To that end, Palantir employs a Client–Server model knitting together a practically limitless number of users into a single enterprise–grade dataspace. Here analysts and operators can work together to explore, discover and create the intelligence necessary to excel at their tactical and strategic tasks. Only an enterprise solution is capable of delivering this truly collaborative and secure functionality.
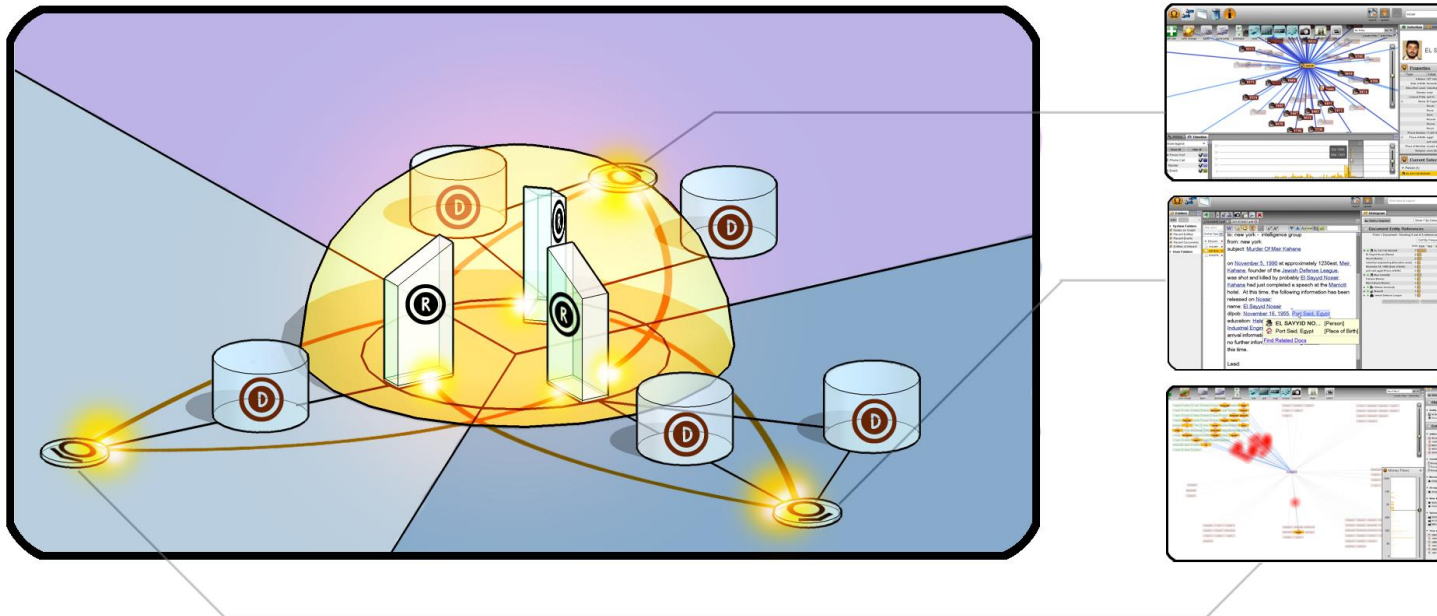
Openness and Extensibility

Before a single line of code was written, the Palantir Platform was designed to be **open**, exposing numerous points for integration and extensibility including an entirely open API (Application Programming Interface)

While being open is considered Standard Operating Procedure in Silicon Valley, it is a revolutionary idea in the field of intelligence. Ultimately, as has been proven time–and–time–again, openness leads to better solutions by combining the vast expertise and interests of broader communities to solve complex problems.

An example of this approach is the Palantir Workspace which is divided into two primary functional units: **Applications and Helpers**. Applications are used to perform large analytical tasks such as link analysis, persistent searching or geospatial analysis. Helpers operate against the data in a given Application to assist in specific analytical processes such as statistical analysis, filtered searching or temporal analysis.

Most importantly, the Application and Helper framework is extensible so that any organization or integrator can build what they need for their workflows.
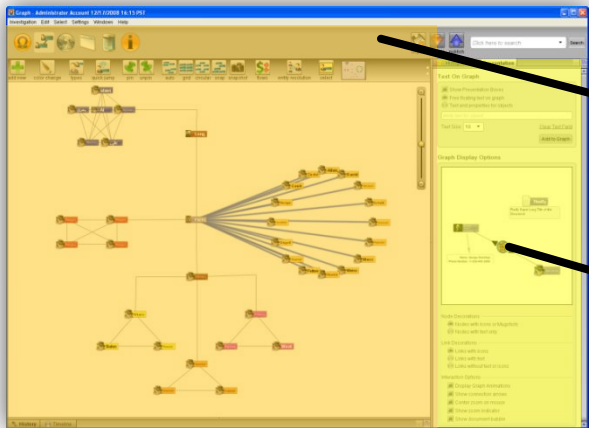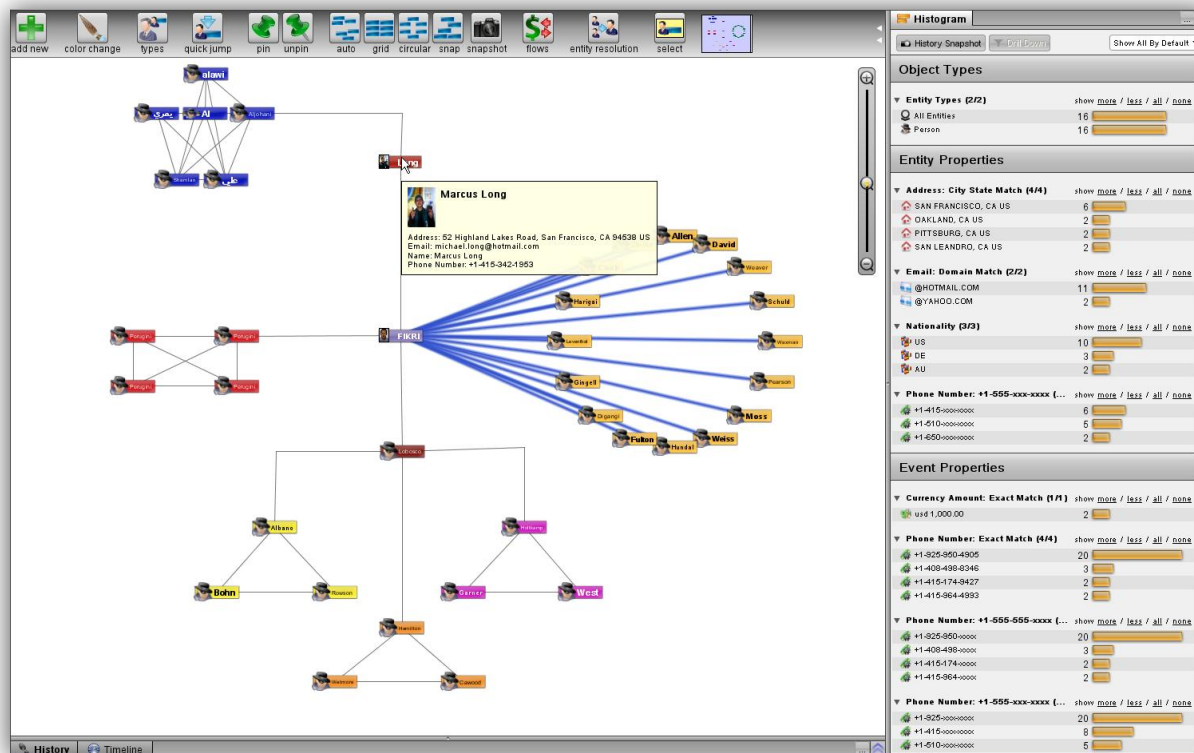

The Application Bar

The Helper Tabs

The Graph Application

The Graph is one of the native applications in the Palantir Platform. It provides a space for performing link analysis and is one of the most critical applications for communicating analytical conclusions.

The graph is fully interactive and can be used both to explore and create connections. The graph is equipped with a host of automated layout and presentation tools including a dedicated Presentation Helper.
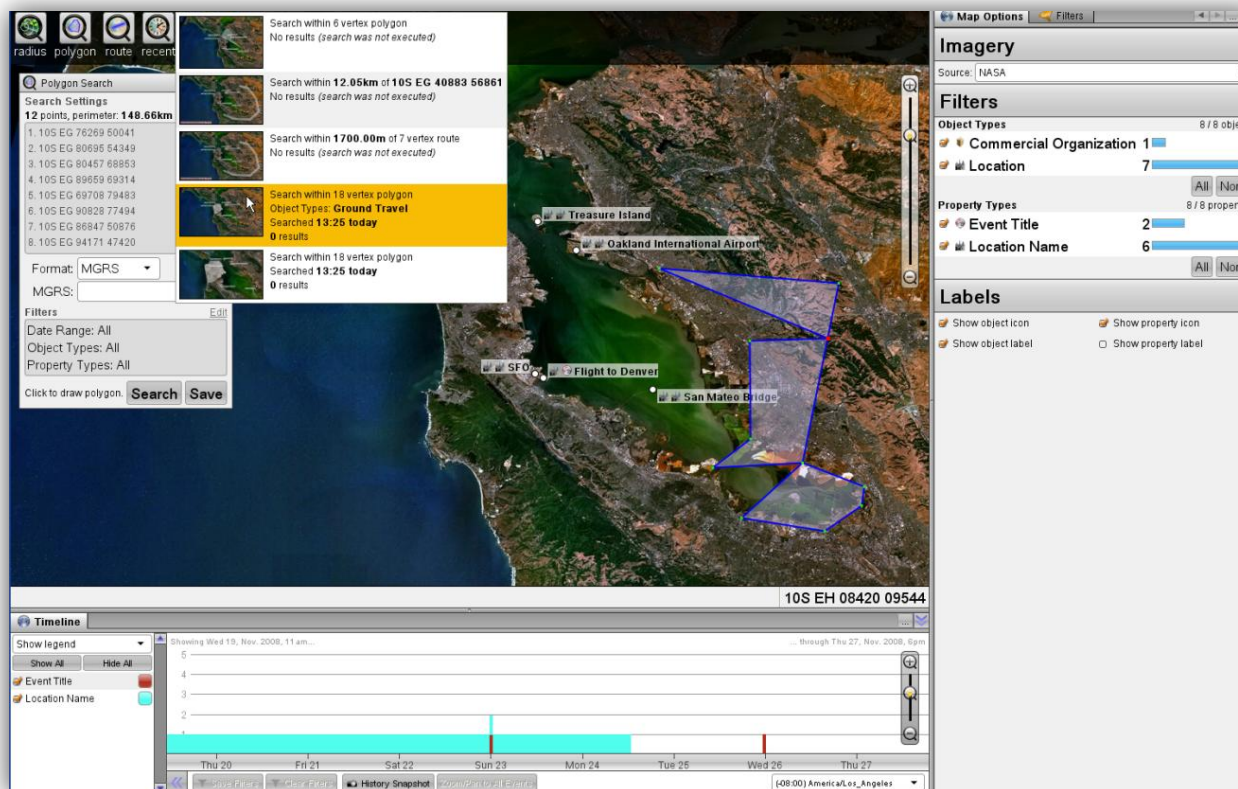

actual screen shot

The Map Application

In version 2.1, a Map Application was added as a native application in the Palantir Platform. This application allows for geospatial analysis, geo–searching and geo–referencing. The Map can operate using any enterprise tile source.

Once constructed, a geospatial search can be added to any filter search and then persisted allowing for the constant monitoring of a n Area of Interest.

actual screen shot

The Browser Application

The Browser is a native application designed to interact with both structured and unstructured data. Any object can be 'opened' in a Browser tab so that all its properties, notes, media, relationships and data sources can be instantly viewed. The Browser also provides a space for reading unstructured documents and assigning structure to them using a process called 'tagging'.

*actual screen shot

The Dashboard, Data Sources and Summary

The other native Applications in the Palantir Platform are the Dashboard, Data Sources and Summary.

The **Dashboard** operates like an analyst's homepage and is the first place an analyst sees when they log into the client. The Dashboard provides a space for the management of investigations and the construction of persistent searches. Persistent searches are queries that run constantly against the enterprise data store so that whenever new data enters the system it is included in the updated query results.

Data Sources provides a high level review of all the data ingested by the Palantir Enterprise. All data sources can be explored here and through the use of collections, data sources can be organized.

Summary provides functionality for rapidly producing reports based on the work accomplished during an investigation. By leveraging the investigation history, important graph states can be quickly collated into a PowerPoint presentation that can be instantly exported to rapidly produce a brief.
This automated report generation displays not just analytic conclusions, but also the analytic path that led to them.

actual screen shot

Palantir
Technologies
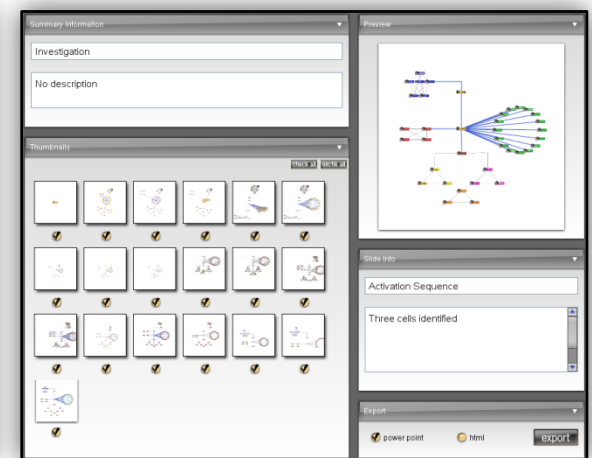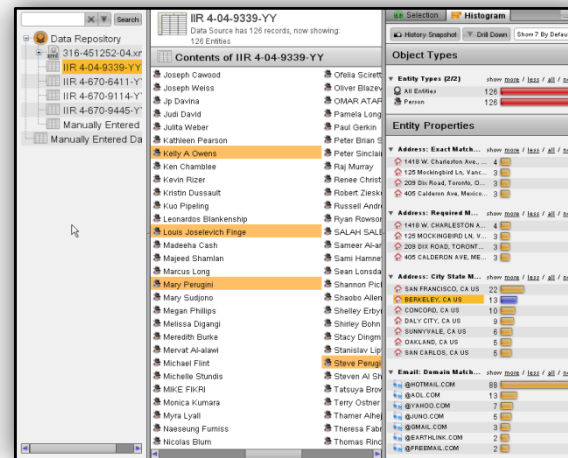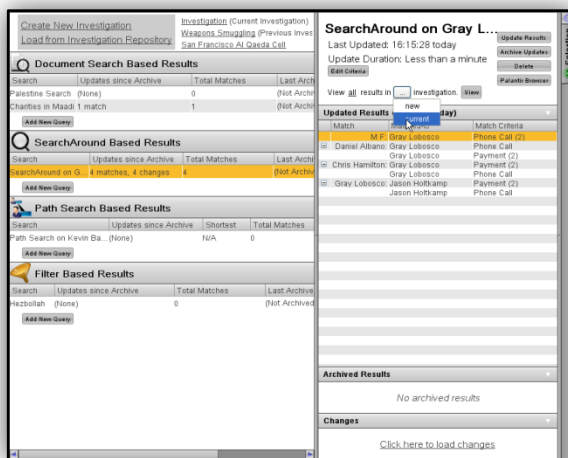
Data Integration	In order to explore lots of data, it is first necessary to bring that data into a single environment where it can be accessed. Palantir can import data from structured sources (DBs, .xls, csv, .anb, .axp...) unstructured sources (.doc, .pdf, .txt, .xml...) as well as federating against external data stores where direct ingest is not an option. In addition, Media (images, video, bio-metrics...) can be associated with any object.
Search & Discovery	Once data has been integrated into the Palantir Platform, it is only useful in so far as it can be explored. To that end, the platform provides a host of approaches for searching into the data in an intuitive and human readable fashion that requires no knowledge of query languages. All data in the system is searched simultaneously which naturally leads to the discovery of links and connections that would have otherwise been missed.
Collaboration	As new information is discovered, it is important that this information is available to other users as though it were any other data source. Palantir utilizes a revolutionary 'version control' model to simultaneously enable broad collaboration as well as individual control of data. Analysts are presented with a virtual private space in which to explore their conjectures which they can then 'publish' to the broader enterprise when they are ready.
Knowledge Management	One of the core technologies of the Palantir Platform is the Revisioning Database which tracks all information in the Palantir environment. This tracking occurs at a very fine-grained level where we track the source of all information, when it was created and by whom. This enables collaboration, search & discovery and data integration by allowing analysts to see who is doing what in their enterprise as well as allowing data to be explored along many different paths.

The Palantir Object Model

Palantir employs a highly flexible **Object Model** which is capable of modeling almost any organizational data model.

The Palantir Object Model consists of Objects, Properties and Relationships. Objects are containers which hold Properties. Relationships are the connections between Objects. An example of an Object, Property and Relationship would be a Person, a Person's Name and Friendship respectively. Each Property and Relationship can be sourced back to original document sources, either structured or unstructured.
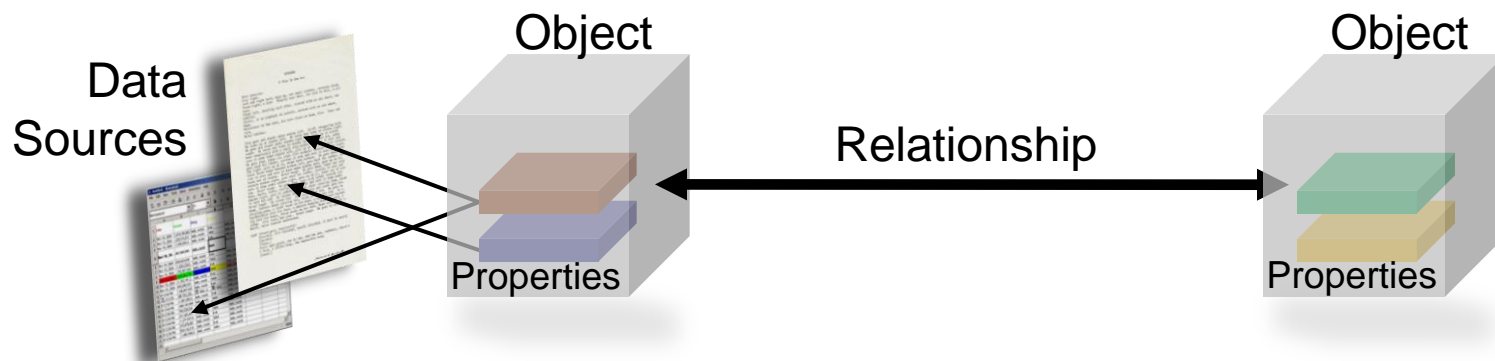
The Dynamic Ontology

Due to the flexible structure of the Object Model, it is possible to use the Palantir **Dynamic Ontology** to create almost any organizational data model.

The Dynamic Ontology is the way to describe what is important to an organization. Palantir does not assume to know how any organization would like to model its information. Rather, Palantir provides an interface to create whatever semantics are useful to the analytic task. The Ontology is entirely customizable and can actually be changed after it has been deployed and data has been integrated.

The Revisioning Database

Palantir employs a revolutionary database structure referred to as the **Revisioning Database** for storing all information introduced to the enterprise.

In this structure, every object in Palantir can be thought of as a stack of cards describing all the attributes of those objects. Each card describes the addition, modification or deletion of a single attribute along with when it happened, who did it, what the security level is and where the information is sourced to. This structure allows for any object to be 're–shuffled' along almost any analytical axis enabling secure collaboration.

Collaboration with Security

It may seem paradoxical at first, but the more that information can be effectively secured, the more it can be effectively shared.

The **Security Model** enabled in Palantir is the most sophisticated and flexible approach to information security in the Intelligence Community. Through rich knowledge management and automated information sourcing, Palantir is able to track the security/classifications associated with every shred of information. By leveraging this extremely fine–grained model, it becomes possible to share as much information as could possibly be shared without revealing the most sensitive information except to those with the permission to see it.

PalantirWeb and Palantir Forward

In addition to the standard client, Palantir can also be operated via a read-only web client (PalantirWeb) and via a re-syncable stand-alone client (Palantir Forward).
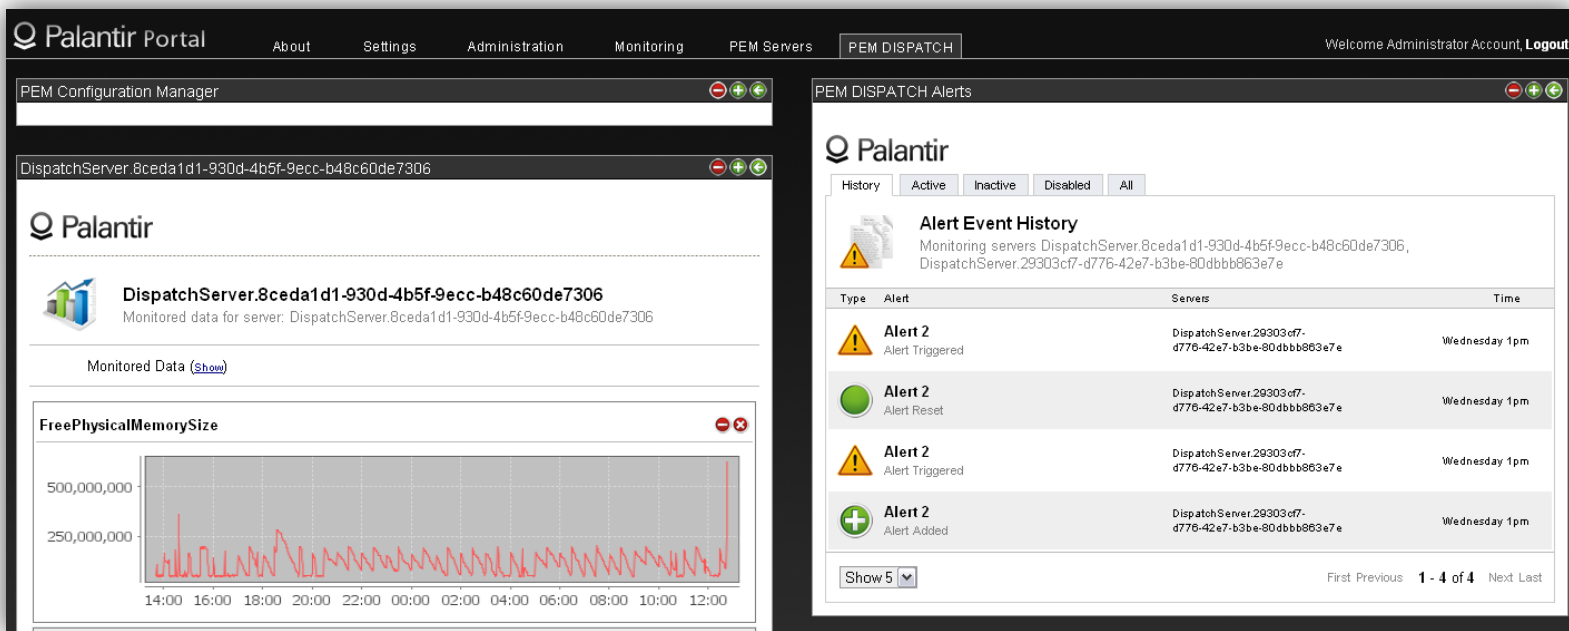
In addition to the Client-Server architecture described previously, it is also possible to interact with a Palantir deployment via **Palantir Web**. Palantir Web is a light weight, high performance web client that allows for the secure, read-only exploration of all data in an Enterprise with both simple keyword and advanced filter searching.

The Palantir client can also operate in a stand-alone mode referred to as **Palantir Forward**. As the name suggests Palantir Forward allows analysts/operators to take their investigations *forward*, copying the portion of the database relevant to that investigation to their local machine. When the analyst/operator returns home, they re-sync their investigation with the enterprise so no information is lost.
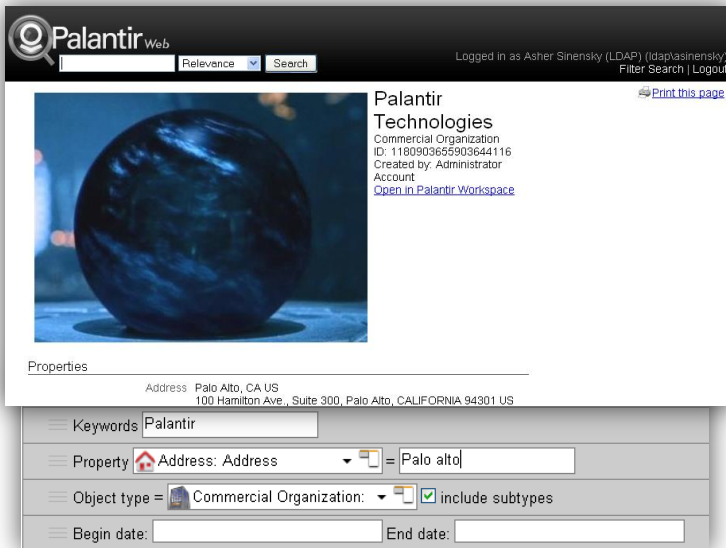
The Palantir Portal

The Palantir Portal is an administrative interface which allows an organization to monitor their entire Palantir Enterprise. The Portal monitors system and usage analytics and provides an extensible 'portlet' framework so that new metrics can be easily added to the already existing ones.

The Portal has an extensible system of alerts and warnings which will instantly contact system administrators if any aspect of the enterprise crosses a threshold including server and usage metrics.

- Sample Workflows:
  - http://www.palantirtech.com/videos
- Technical White Videos
  - http://www.palantirtech.com/whitevideos
- Operation Trade Stop
  - http://www.optradestop.com
- Project GreyGoose
  - http://www.palantirtech.com/greygoose/

inquiries@palantirtech.com

National Security Archive,

Suite 701, Gelman Library, The George Washington University,

2130 H Street, NW, Washington, D.C., 20037,

Phone: 202/994-7000, Fax: 202/994-7005, nsarchiv@gwu.edu