

IBM i2 Analyze Architecture and Services White Paper

Version 4 Release 1

Note

Before using this information and the product it supports, read the information in "Notices" on page 23.

This edition applies to version 4, release 1, modification 1 of IBM i2 Analyze and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 2012, 2016.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Chapter 1. IBM i2 Analyze architecture

IBM® i2® Analyze is designed to help analysts acquire and discover information, and to produce and publish intelligence products by fusing that information. The logical architecture, physical architecture, and security architecture of i2 Analyze support these functions.

The logical architecture defines the components of i2 Analyze and how they interact with each other. The logical architecture of i2 Analyze includes the methods by which data enters and is handled, the services provided, and the clients through which users interact with i2 Analyze.

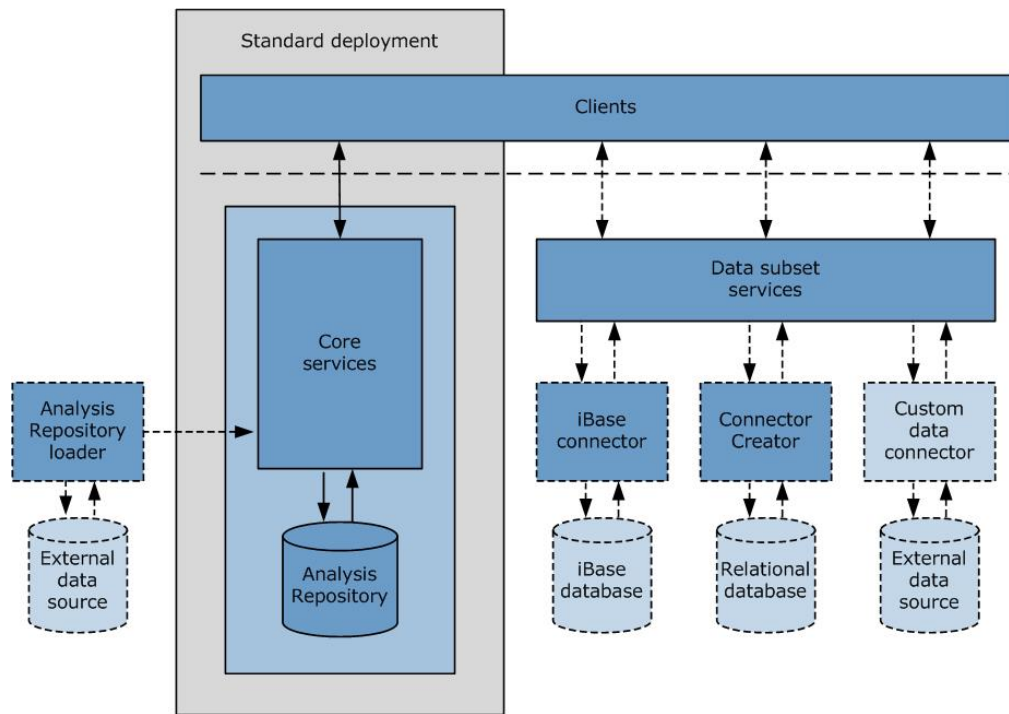
In every deployment of i2 Analyze, the logical architecture maps to a physical architecture, where the components are deployed and connected to each other. The physical architecture of i2 Analyze can vary depending on the purpose of your deployment, your performance and availability requirements, and any custom functionality that you implement.

The information in this document summarizes the physical architecture of an i2 Analyze deployment. For details of the deployment process, see the *IBM i2 Analyze Deployment Guide* for the target operating system.

The security architecture enables i2 Analyze to authenticate users, and to determine their authorization level for every item that it manages. Your implementation of the security architecture can use alternative technologies to those used in a default deployment of i2 Analyze, provided that the alternative technologies fulfill the requirements specified by the security architecture.

IBM i2 Analyze logical architecture

The logical architecture can be viewed in terms of the data that is handled by i2 Analyze, the services that handle the data, and the clients that access the data.



Data

i2 Analyze can accept data from a variety of external sources and make it available to users for analysis. i2 Analyze can store ingested data to provide to clients, or provide the data directly to clients without ingesting it. The characteristics of a data source determine the most suitable location for storing or handling the data:

Analysis Repository

The Analysis Repository is a collaborative repository for verified intelligence data. The Analysis Repository contains high value records that analysts have created from the original source data, then curated into the Analysis Repository. Data can also be placed directly into the Analysis Repository through data load direct, using the Analysis Repository loader. Data in the Analysis Repository can be retrieved for further analysis, then modified with the results of the analysis.

External data sources through data connectors

Data from an external data source can be accessed through a data connector, which uses customized services to read data from an external data source in response to a user's request. When the external data source is IBM i2 iBase, or a relational database with a reasonably simple structure, i2 Analyze has tools that make it easier to create data connectors. An organization can also implement a custom data connector. Data connectors convert the data into an analysis-ready form, and provide it directly to the Intelligence Portal and Analyst's Notebook Premium. The data in the original data source is not changed as a result of the analysis.

Services

i2 Analyze services enable users to perform operations on the data that is stored and provided by i2 Analyze. The following services enable the movement of data into and out of the i2 Analyze server:

- For the Analysis Repository, the core services enable users to discover and retrieve items for further analysis, and also enable users to carry out operations on the data stored in the Analysis Repository.
- For an external data source accessed through a data connector, i2 Analyze provides data subset services that the data connector uses to acquire and provide external data for users.

Other core services enable users to set up and subscribe to alerts for item changes, to audit item changes, to purge items, and to administer the components of i2 Analyze.

Clients

Analysts use the following clients to discover and analyze data that is stored and provided by i2 Analyze:

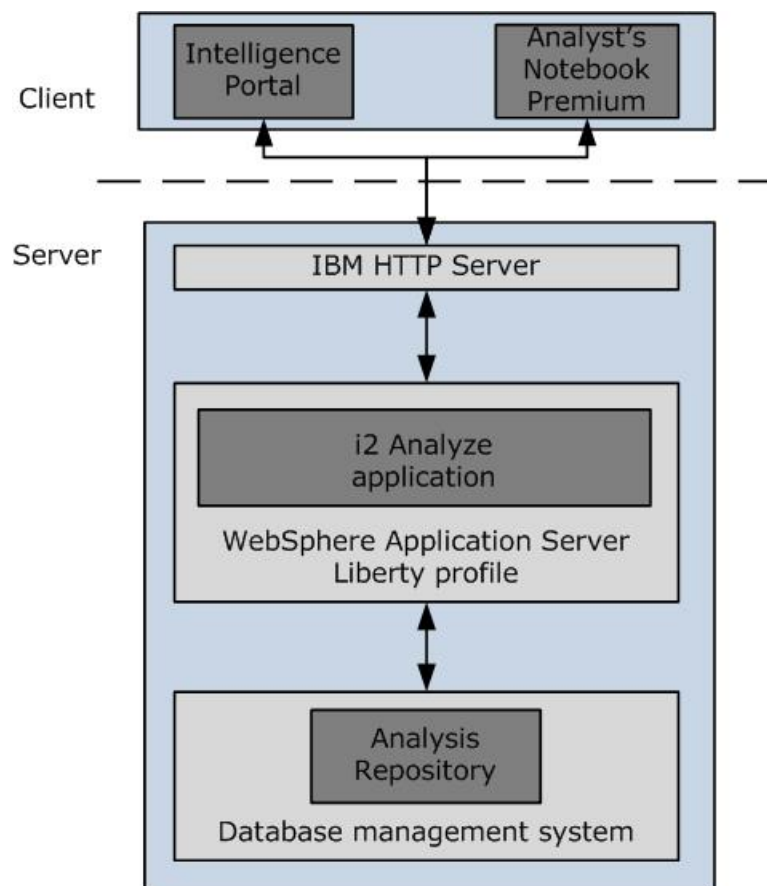
- The Intelligence Portal provides search, discovery, and visualization capabilities through web browsers.
- Analyst's Notebook Premium provides in-depth discovery, visual analysis, and intelligence production capabilities through a rich desktop client.

IBM i2 Analyze physical architecture

The components of IBM i2 Analyze support various physical configuration options. Plan the physical architecture to fit the logical architecture and the performance and availability requirements of the organization.

The logical architecture is designed for flexible deployment of the i2 Analyze data storage, the i2 Analyze services, and data connectors for external data sources. In small deployments, many of the components can be co-located. To improve scalability, or to fit better with existing network architectures, the components can be deployed separately across multiple servers. For information on the deployment process, refer to the *IBM i2 Analyze Deployment Guide*.

The following diagram shows the physical architecture of an i2 Analyze deployment that reflects the instructions in the *IBM i2 Analyze Deployment Guide*:



An i2 Analyze deployment uses the following physical architecture components to provide the elements of the logical architecture:

Components for data

- The Analysis Repository hosted in a compatible database management system, which can be an IBM DB2[®] database, or alternatively, an Oracle or Microsoft SQL Server database.

Components for services

- An IBM WebSphere[®] Application Server Liberty profile server.
- The i2 Analyze application that provides the core i2 Analyze services. This application is hosted in WebSphere Application Server Liberty profile.
- An instance of IBM HTTP Server, which acts as a proxy server for client requests.

Components for clients

- The Intelligence Portal, a web client that is available to users through web browsers. The Intelligence Portal is a Microsoft Silverlight web client that requires a browser plug-in.
- Analyst's Notebook Premium, a rich client that is installed locally on users' workstations.

For information on the software versions required to support i2 Analyze, search for the relevant Software Product Compatibility Report at <http://www.ibm.com/software/reports/compatibility/clarity/softwareReqsForProduct.html>.

Analysis Repository

The Analysis Repository is a collaborative repository for verified intelligence data in IBM i2 Analyze.

The Analysis Repository contains high value records that analysts have created from the original source data, then curated into the Analysis Repository. Data may also be placed directly into the Analysis Repository through data load direct, using the Analysis Repository loader.

The items in the Analysis Repository provide high value source data for discovery by analysts, and may be modified as a result of later analysis or fresh information. Items in the Analysis Repository can be searched in a variety of ways, and a browse function is also available.

i2 Analyze provides the following functions for managing items in the Analysis Repository:

- Individual creation of items
- Full editing of items
- Merging items
- Grouping items into sets
- Deleting items
- Purging items
- Versioning for items
- Auditing and history of individual item changes
- Alert subscriptions for item changes

- Security grading for individual items

Clients use the command service provided by i2 Analyze to create, update, and delete items in the Analysis Repository, and the purge service to remove items completely. The Analysis Repository maintains a version history for every item that it contains:

- For creation, "version 1" of the item is created and stored in the Analysis Repository.
- For modification, a new version of the item is created, and the old version is retained.
- For deletion, a version of the item is created, and the new version is marked as deleted.
- For purging, all versions of the item are removed from the Analysis Repository.

The Analysis Repository stores a permanent record of every command that has been successfully issued using the command service, including the request, the issuer, when they issued it, and where (that is, what IP address) they issued it from.

Documents

Documents that are stored in the Analysis Repository are associated with items. An item can be associated with several documents, but any single document is associated with exactly one item.

The association between items and documents works through properties whose values identify documents. When a client uploads a document to the Analysis Repository, the servlet returns a unique identifier for that document. The client then uses the command service to store that identifier in a suitable item property.

i2 Analyze extracts text from documents and adds it to the search index. A search for items matches against property values and the text content of any associated documents. The search service must therefore track changes to documents in the Analysis Repository. The standard i2 Analyze search service keeps the search index in step with the documents in the Analysis Repository.

Leases

i2 Analyze prevents multiple clients from changing the same item in the Analysis Repository at the same time. Clients must take a lease before they change an item and release it after they change an item. Any attempt to edit an item that is not leased will fail.

- Before a client issues a command to modify an item in the Analysis Repository, it must take a lease on that item. If another client has already taken the lease, then the command to take a lease on the item fails. Otherwise, i2 Analyze records the action.
- When a client issues a command to modify an item, i2 Analyze checks the lease records. If this client has not taken the lease for that item, then the command fails.
- After a client modifies an item, it must release its lease on that item. Alternatively, i2 Analyze releases a lease after a configurable period of inactivity. In both cases, i2 Analyze records the action.

Data load direct for the Analysis Repository

You can load external data directly into the Analysis Repository. Clients can interact with the imported data in the same way as they interact with any other data in the repository.

In data acquisition solutions that use data load direct, an Analysis Repository loader acts as an intermediary between the external data source and the command service provided by i2 Analyze. The loader retrieves data from the source, and calls the command service to add that data to the Analysis Repository.

For more information on data load direct, see the *IBM i2 Analyze Data Acquisition Overview White Paper*.

Data connectors

A data connector can be used to acquire data from an external data source. When a user makes a request to an external data source, the data is converted from its original format into entities, links, and properties, and is then displayed to the user.

When data is provided to i2 Analyze through a data connector, the external data source is not modified and remains complete and unchanged. When the data from the external data source has been used in analysis and vetted and verified, the new or modified entities and links can be added to the Analysis Repository.

An i2 Analyze deployment requires a separate data connector for each external data source to which it provides access. When the external data source is IBM i2 iBase, or a relational database with a reasonably simple structure, i2 Analyze has tools that make it easier to create data connectors.

An iBase connector enables i2 Analyze to work with any deployment of iBase. In the Intelligence Portal, users can conduct the same range of search operations against iBase as they can conduct against the Analysis Repository. Through an iBase connector, organizations can deploy i2 Analyze with a schema that matches an existing iBase database, or deploy iBase as a data store that supports i2 Analyze.

Connector Creator enables organizations to create data connectors for relational databases without writing Java code. In this model, each database requires an XML configuration file that defines the mapping between its data and the i2 Analyze schema, and contains the SQL queries that users can run against it. From the configuration file, Connector Creator generates both the data connector and the user interface that appears in the Intelligence Portal.

To connect to an external data source that is not a relational database, or to allow analysts to specify their own queries, an organization can implement a custom data connector. For information on coding a data connector, see the GitHub repository for IBM i2 Analyze Developer Essentials, and the IBM i2 Analyze Data Acquisition white paper.

Security architecture

The IBM i2 Analyze security architecture supports the behavior that the i2 Analyze security model requires. Every i2 Analyze service can interact with the security architecture to determine what rights the current user has for the operation that they want to perform.

i2 Analyze authenticates users through a choice of technologies, and determines their authorization level for every item that it manages. The i2 Analyze security model bases its behavior on the interaction between the security dimension values that items have, and the security permissions that user groups have.

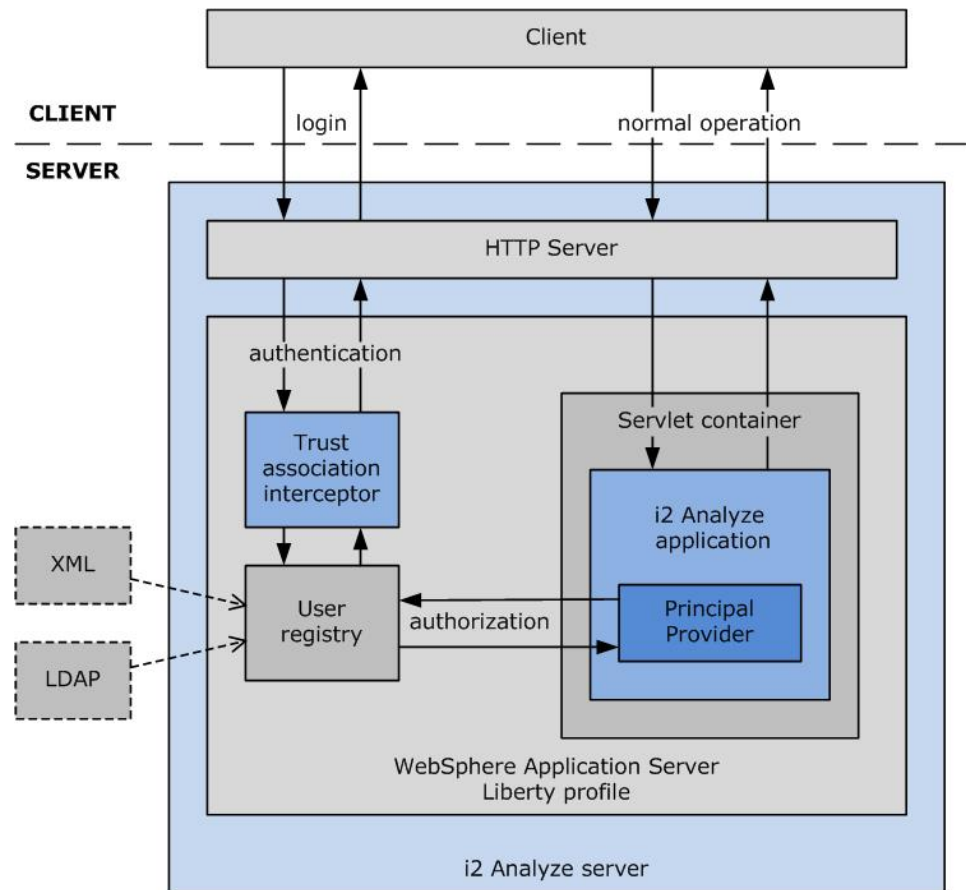
- Users categorize items in the Analysis Repository by assigning them values from security dimensions. The value that an item has within a particular security dimension affects whether users can view or edit that item (access security), or change its dimension values (grant security).
- Security permissions apply to groups of users. On a per-group basis, they associate security access or grant levels with particular dimension values that items can have. Group membership is often determined by the job titles or the security clearance of the users that they contain.

For more details about the security model for i2 Analyze, see the *IBM i2 Analyze Security White Paper*.

The components of an i2 Analyze deployment interact with the security architecture in the following ways:

- At login, WebSphere Application Server Liberty profile requires clients to authenticate before they can interact with i2 Analyze. On successful authentication, the client receives a Lightweight Third-Party Authentication (LTPA) token in a cookie.
- During normal operation, the client passes the cookie back to i2 Analyze, where each i2 Analyze service enforces data access rights in accordance with its particular function.

The following diagram shows how security works in a standard i2 Analyze deployment:



In a standard i2 Analyze deployment that uses the Intelligence Portal client and the Analysis Repository, authentication and authorization take place as follows:

- Authentication between the Intelligence Portal client and WebSphere Application Server Liberty profile is handled through a trust association interceptor (TAI) that IBM provides with i2 Analyze.
- The TAI interacts with the WebSphere Application Server Liberty profile user registry to verify the credentials that the user supplies through the client. The user registry is a service that provides access to user and group information, which can be held in an XML file, one or more LDAP registries, or in any similar store that WebSphere Application Server Liberty profile can use.
- To authorize users of the Intelligence Portal client to access items in the Analysis Repository, the i2 Analyze application communicates with the WebSphere Application Server Liberty profile user registry to retrieve information about the current user's membership of groups. A principal provider then maps the retrieved information to group permissions that are defined in the security permissions section of the i2 Analyze security schema. This mapping is deployment-specific because the security schema is deployment-specific.
- Code in each i2 Analyze service compares the permissions of the current user with the security dimension values for items, to determine what rights the user receives for each item.

The technologies in the diagram are not fixed. It is possible to provide a different TAI, and to use any supported store for the user registry. The requirements are as follows:

- An i2 Analyze service must be able to derive information about a user from the credentials they present.
- A (potentially) deployment-specific module must map user information onto membership of the groups that are named in the security permissions section of the i2 Analyze security schema.

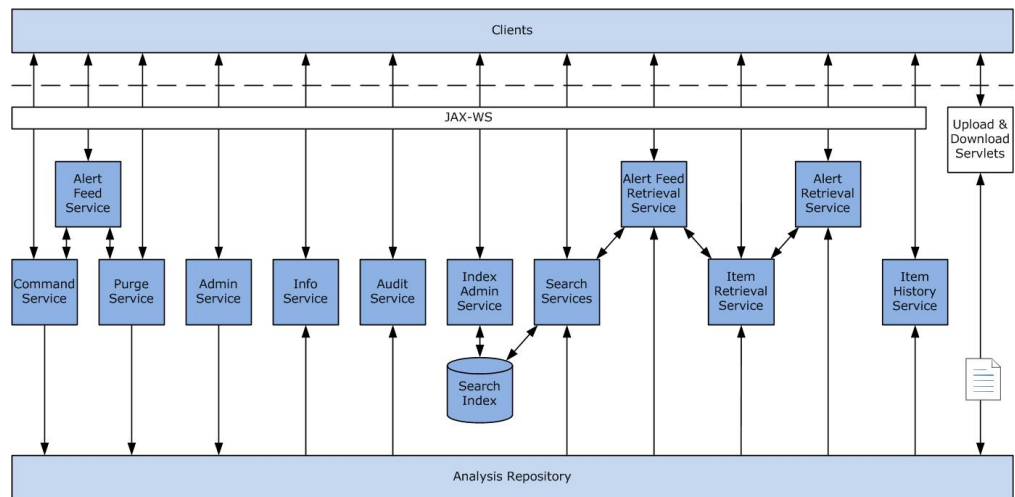
If an implementation of the security architecture fulfills these requirements, then it is suitable for use in an i2 Analyze deployment.

Chapter 2. IBM i2 Analyze services

The IBM i2 Analyze application provides core services and infrastructure that enable clients to create, view, analyze, and update the data to which i2 Analyze has access.

Services for the Analysis Repository

The core services enable interactions between clients and the Analysis Repository, as shown in the following diagram.



Note: JAX-WS is the Java™ API for XML Web Services.

Clients do not use services to upload documents to, or download them from, the Analysis Repository. Instead, document upload and download are handled by a pair of servlets.

The core services provide the following functions to users through clients:

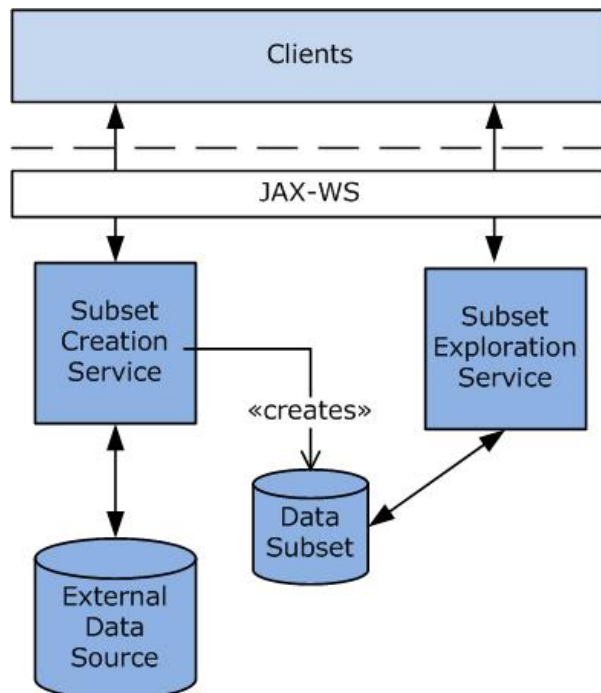
Service	Purpose
Admin	<ul style="list-style-type: none"> Set up Analysis Repository and schema Set up security schema Change schema for Analysis Repository or for security
Alert Feed	Create and subscribe current user to alert feeds for changes to items
Alert Feed Retrieval	Retrieve alert feed information and check subscriptions for current user
Alert Retrieval	Retrieve alerts generated by alert feeds
Audit	Retrieve all changes to items in Analysis Repository

Service	Purpose
Command	<ul style="list-style-type: none"> • Create and merge items in Analysis Repository • Delete items (retrievable) • Amend details of items • Add and remove items from sets • Change item security permissions • Take and release leases on items
Index Admin	Take snapshot of search index for backup
Info	<ul style="list-style-type: none"> • View information about the i2 Analyze instance • View information about the current user • Access logging mechanism
Item History	Retrieve change history for items in Analysis Repository
Item Retrieval	Retrieve specified or latest versions of items in Analysis Repository
Purge	Delete items or sets from Analysis Repository (non-retrievable)
Search	Browse, query, and expand items in Analysis Repository

Services for data connectors

i2 Analyze provides services to support data connectors. Clients use these services to query and retrieve results from external data sources through the Intelligence Portal user interface.

Data connectors create and interrogate relatively short-lived subsets from larger volumes of data. These operations are supported by two services that developers must customize for each data connector, according to the requirements of the external data source in question.



Service	Purpose
Subset Creation	Search external data source and create a subset of data
Subset Exploration	Search within a subset of data and retrieve items from it

Admin service

The i2 Analyze admin service enables clients to initialize and configure i2 Analyze and the Analysis Repository. The admin service also provides control over the security schema.

The admin service contains four pairs of web methods.

Web method	Description
Set/GetChartingSchemes	Sets the charting schemes for this instance of i2 Analyze.
Set/GetRepositoryName	Sets the human-readable name of the Analysis Repository for this instance of i2 Analyze.
Set/GetSchema	Sets the schema for this instance of i2 Analyze. If a schema has previously been set, there are strict rules about how the new schema can differ from the current one.
Set/GetSecuritySchema	Sets the security schema for this instance of i2 Analyze. If a security schema has previously been set, there are strict rules about how the new security schema can differ from the current one.

Operations that modify the i2 Analyze schema or the security schema are tightly constrained after i2 Analyze is deployed. Making significant changes to the schema or the security schema is equivalent to starting a new deployment of i2 Analyze.

Alert feed service

The i2 Analyze alert feed service allows clients to create, manage, and delete alert feeds in the Analysis Repository. The alert feed service acts as an interface to the command and purge services.

In terms of the i2 Analyze data model, an alert feed is a special type of item. When a user subscribes to an alert feed, i2 Analyze creates a link between the user item and the alert feed item.

The alert feed service presents clients with a domain-specific interface for creating and editing alert feed items in the Analysis Repository.

Web method	Description
CreateAlertFeed	Creates an alert feed for a specified item in the Analysis Repository. When a feed exists, i2 Analyze generates alerts in response to any changes on the item (and, optionally, on any items in a set that the item owns).
SubscribeCurrentUserToAlertFeeds	Subscribes the current user to specified alert feeds.
UnsubscribeCurrentUserFromAlertFeeds	Unsubscribes the current user from specified alert feeds.

The pattern through which the alert feed service uses the command and purge services to manipulate the Analysis Repository is a model for custom services. Developers can extend i2 Analyze by writing their own services that use the standard i2 Analyze facilities.

Alert feed retrieval service

The standard i2 Analyze alert feed retrieval service enables clients to retrieve alert feeds, and to determine which feeds they subscribe to.

The alert feed retrieval service has a single method that searches for alert feeds.

Web method	Description
SearchAlertFeeds	Retrieves alert feeds according to criteria such as the feed descriptions, and whether the current user is subscribed to them.

Alert retrieval service

The standard i2 Analyze alert retrieval service enables clients to retrieve the alerts generated by alert feeds.

The alert retrieval service has a single method for retrieving alerts, subject to a number of constraints.

Web method	Description
RetrieveAlerts	Retrieves alerts from i2 Analyze according to different criteria. Valid constraints include the alert feeds that generated the alerts, and the time since alerts from the same feeds were last retrieved.

Audit service

The standard i2 Analyze audit service maintains a permanent record of all the changes that take place to the items in the Analysis Repository. The audit service reports the type of any change, when it took place, who made it, and the IP address from which they made it.

The audit service provides a single web method for clients to call.

Web method	Description
RetrieveAuditEntries	Returns all the audit entries between two time stamps. The results contain entries for all items that the client has permission to see.

When an item is purged from the Analysis Repository, the audit service adds a record of that event, and retains its records of all previous events. Audit entries for purged items do not contain information about the item itself; they store only the types of event that affected the item.

Command service

The i2 Analyze command service enables clients to execute commands that change data in the Analysis Repository. The command service also manages leases that help to determine which commands a client can execute on a particular item at a particular time.

The command service enables clients to create, merge, modify, and delete items, to group items into sets, and to change the security permissions for items. The functions provided by the command service are fully documented in the IBM i2 Analyze SDK API documentation, which is available at <http://ibm-i2.github.io/Analyze/docs/index.html>.

To protect against concurrent modifications, the command service requires clients to take a lease on an item before they are allowed to modify it. Clients take leases on items implicitly when they create them, or explicitly by executing a "take lease" command.

Only one client can hold the lease for a particular item at a time. Leases expire after a configurable period, but clients that require shorter or longer leases can release or renew them.

To use the command service, a client creates commands, places them in a list, and submits them to the service for execution. The order of commands in the list is significant, since the command that takes a lease on an item must come before the command that modifies it. The command service validates the list of submitted commands, which includes checking leases, versions, and access permissions. The command service then executes the commands in the list atomically, so that the whole list completely succeeds, or completely fails.

Web method	Description
ExecuteCommands	Executes named commands in a command list, in the order in which they appear.
RenewLeases	Renews the leases on the items with particular identifiers.

Index admin service

The standard i2 Analyze index admin service provides clients with the ability to take a snapshot of the search index.

The snapshot that the index admin service takes can be backed up without being affected by simultaneous changes to the index.

The index admin service has one method for taking a snapshot of the search index, and another for releasing the snapshot after a backup is made.

Web method	Description
TakeSnapshot	Takes a snapshot of the search index so that a backup of the index can be performed.
ReleaseSnapshot	Releases a snapshot of the search index when it is no longer required.

Info service

The i2 Analyze info service enables clients to discover information about a running instance of i2 Analyze. The info service also provides some simple utility methods.

Through the info service, clients can retrieve:

- The data sources to which this instance of i2 Analyze is connected
- The schema that defines the item types that are available in this instance
- The charting scheme that describes how items from i2 Analyze appear when they are visualized
- The security schema that controls how users can view and interact with items in i2 Analyze
- The principal name of the current user (so that the client can display it in the user interface, for example)

The info service also provides utility methods through which a client can discover more information for the current user:

- The security levels that the user would have on an item with particular permissions
- The security dimensions that the user has permission to know about
- Globalized data about the time zones that this instance of i2 Analyze supports

Finally, the information service provides access to a logging mechanism that other services can use, and a way of checking whether a given date and time value is valid.

Item history service

The standard i2 Analyze item history service provides clients with details of how items in the Analysis Repository changed over time. Items in the Analysis Repository can change as a result of edit and merge operations.

The item history service is similar to the item retrieval service, but it provides a different view of the information in the Analysis Repository.

Clients supply the item history service with the identifiers of items they retrieved at some earlier time. Clients can ask the service for details of how the items evolved, the differences between versions of the same item, or merge information.

Web method	Description
CompareItemVersions	Returns a structure that details the changes to the cards, property groups, properties, and source information between two specific versions of the same item.
RetrieveItemHistory	Returns a list of the kinds of changes that an item underwent between two points in its lifetime, and the users responsible for those changes. The item history service provides information about changes to item data; to the contents of any set the item owns; and to security permissions.
RetrieveMergeHistories	Returns structures that contain summary information about any items that were merged to create specific items.

Web method	Description
RetrieveMergePopulations	Returns structures that contain summary information about any items that were merged to create specific items, and any items into which the specific items were merged.
RetrieveMergeProducts	Returns structures that contain summary information about the most recent items to which specific items contributed content, through merge operations.

Item retrieval service

The standard i2 Analyze item retrieval service enables clients to retrieve specific versions of specific information in the Analysis Repository.

An i2 Analyze client supplies the item retrieval service with one or more item identifiers. The client retrieved the item identifiers from either a search service or the item retrieval service itself. The identifiers might be further qualified with version numbers. The item retrieval service then provides the following functionality.

Web method	Description
GetItems	Returns the specified versions of the items with the specified identifiers.
GetLatestItems	Returns the latest versions of the items with the specified identifiers.

The results that the item retrieval service returns to clients are subject to the security rules of i2 Analyze. Clients never see items whose permissions dictate that the user does not have access to them.

Purge service

The i2 Analyze purge service enables clients to remove all trace of specified items (apart from their record identifiers) from the Analysis Repository. By intention, it is impossible to retrieve or reconstitute a purged item.

Architecturally, the purge service is like the command service. Functionally, however, the purge service is different. Because of the actions it takes, the ability of users to access the purge service is controlled separately from their ability to issue other commands. Users who have access to the purge service are able to purge any item in the Analysis Repository, regardless of their other permissions.

The purge service contains two web methods for clients to call.

Web method	Description
PurgeItems	Removes specified items from the Analysis Repository, making all data non-retrievable. If a specified item is an entity with links to other entities, those links are also purged from the Analysis Repository.
PurgeSet	Removes a specified entity from the Analysis Repository, and also removes the items in any set that is owned by the specified entity.

On receiving an instruction to purge an item or a set from the Analysis Repository, the purge service takes the following actions:

- It takes a lease on the item, by force if necessary.
- It takes leases on any links that are connected to the item, with the same provision.
- It deletes all data about the leased items from the Analysis Repository.

The standard audit service reports that the purge took place, but retains no further information.

Search services

The standard i2 Analyze search and network search services enable clients to browse into, query, and expand against the latest versions of data in the Analysis Repository.

Clients of the search and network search services retrieve results that can be sorted, filtered, or scored according to a range of criteria.

For i2 Analyze clients, running a search against the Analysis Repository is generally a two-part operation:

1. The client specifies some or all of the following information:
 - The text or other property values to be searched for
 - The set that contains the items to be considered
 - The relationships that must exist between any returned items
 - The types of items that the search must return
 - The properties by which to filter any items found

In response, the search services return a results summary that contains the number of matching items and their associated types.

2. The client specifies a token for the executed search, and some or all of the following information:
 - The types of results it wants to retrieve
 - The properties by which to perform further filtering
 - The criteria through which to sort the items

The search services then return actual item results to the client, including (when appropriate) an indication of how well each item matches the original search criteria.

This approach is reflected in the web methods that the search services make available for clients to call. The following methods are for the search service:

Web method	Description
Search	Searches the Analysis Repository for items that meet a range of criteria, and returns a summary of the results.
DumbbellSearch	Searches the Analysis Repository for dumbbells (entity-link-entity structures) that meet a range of criteria, and returns a summary of the results.
Retrieve	Returns results that contain items from an earlier call to Search or DumbbellSearch, allowing further filtering and sorting to take place.

Web method	Description
SearchAndRetrieve	Combines the functionality of Search and Retrieve into a single call, at the cost of the ability to perform multiple filtering and sorting operations on the same set of initial results.
Expand	Searches the Analysis Repository for links that meet a range of criteria, and that are connected to previously discovered entities. Returns the links and their associated entities.
PagedExpand	Searches the Analysis Repository for links that meet a range of criteria, and that are connected to previously discovered entities. Returns the links and their associated entities, in paged form.
FindDuplicateItems	Searches the Analysis Repository for items with similar values for a range of properties, and that might therefore be duplicates of one another.

The network search service provides the following methods:

Web method	Description
Search	Searches the Analysis Repository for items that together meet the assembled criteria for the nodes in a network, and returns a summary of the results.
Retrieve	Returns results that contain items from an earlier call to Search, allowing further filtering and sorting to take place.
SearchAndRetrieve	Combines the functionality of Search and Retrieve into a single call, at the cost of the ability to perform multiple filtering and sorting operations on the same set of initial results.
ValidateSearch	Validates the criteria that are to be used in a search operation without actually performing that operation.

The search services implement the security rules of i2 Analyze. The summary of results contains information only about items that the user who performed the search can access.

Data connector services

IBM i2 Analyze defines services that enable data acquisition through data connectors. Clients use these services to query and retrieve results from external data sources through the Intelligence Portal user interface.

Data connectors create and interrogate relatively short-lived subsets from larger volumes of data. These operations are supported by services that developers can use in each data connector, according to the requirements of the external data source in question.

External data subset creation service

In i2 Analyze, an external data subset creation service is used as the first stage of a data connector. This stage of the connector retrieves data from an external data source. The subset creation service enables users to run searches against an external data source from the standard Intelligence Portal user interface. It then

returns an identifier that represents the subset of data returned from the query, against which subsequent analysis takes place.

External data subset creation services can support the following web methods. In an implementation where the external data source does not support particular query types, you can disable some of the methods for the data connector.

Web method	Description
Search	Searches an external data source for items that meet a range of criteria, and creates an external data subset from the results.
FullTextSearch	Searches an external data source for items whose content includes the specified string, and creates an external data subset from the results.
DumbbellSearch	Searches an external data source for dumbbells (entity-link-entity structures) that meet a range of criteria, and creates an external data subset from the results.
NetworkSearch	Searches an external data source for groups of items that together meet a range of criteria, and creates an external data subset from the results.
Expand	Finds links (and their associated entities) that are connected to specific entities in an external data source, and creates an external data subset from the results.

External data subset exploration service

In i2 Analyze, an external data subset exploration service is used as the second stage of a data connector. This stage of the connector enables Intelligence Portal users to perform analysis on the subset of data that was retrieved from the external data source. The subset exploration service enables clients to browse into, query, and expand against the retrieved data. Typically, the second stage of the data connector is also responsible for converting the data in the subset to a format compatible with the i2 Analyze schema.

External data subset exploration services support the following web methods. After data is retrieved from an external source and a subset is created, all of these methods are available for the data in the subset.

Web method	Description
Initialize	Initializes a specified external data subset so that it is ready to receive further requests.
GetLatestItems	Retrieves items with particular identifiers from a specified external data subset.
Search	Searches a specified external data subset for items that meet a range of criteria.
DumbbellSearch	Searches a specified external data subset for dumbbells (entity-link-entity structures) that meet a range of criteria.
NetworkSearch	Searches a specified external data subset for groups of items that together meet a range of criteria.
Retrieve	Retrieves items from a specified external data subset that were matched by a previous call to Search or DumbbellSearch, allowing sorting and extra filtering to be carried out.

Web method	Description
NetworkRetrieve	Retrieves items from a specified external data subset that were matched by a previous call to NetworkSearch, allowing sorting and extra filtering to be carried out.
Expand	Retrieves links (and their associated entities) that are connected to specific entities within a specified external data subset.

The *IBM i2 Analyze Data Acquisition Overview White Paper* contains more information about data connectors, and the operation and implementation of services for creating and exploring data subsets.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM United Kingdom Limited
Hursley House

Hursley Park
Winchester, Hants, SO21 2JN
UK

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, i2, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Other names may be trademarks of their respective owners. Other company, product, and service names may be trademarks or service marks of others.