



Spotfire for the Enterprise: An Overview for IT Administrators

This whitepaper is intended for those wanting information on TIBCO Spotfire^{*} administration and deployment capabilities: its architecture, data connection, security, scalability, administration, and performance.

THE SPOTFIRE ENVIRONMENT

The TIBCO Spotfire platform uses a client/server architecture in which users analyze their data using Spotfire clients (Figure 1).



Figure 1: Spotfire Platform - Manager and Workers

TIBCO Spotfire' Server is the central component that all clients connect to. Multiple nodes are installed and connected to Spotfire Server. TIBCO Spotfire' Web Player and the TIBCO Spotfire' Automation Services are installed on nodes to enable Spotfire web clients and Automation Services jobs. The Server is connected to a Spotfire database that contains a user directory and stores analyses and configuration files. From a Spotfire Server start page entities in the Spotfire environment can be configured and monitored.

SPOTFIRE SERVER

Spotfire Server, a Tomcat web application that runs on Windows and Linux operating systems, is the administrative center of any Spotfire environment. In addition to providing the tools for configuring and administering the Spotfire environment, the Server, through Spotfire clients, enables users to access their data, create visualizations, and share them with their co-workers or with the world.

Spotfire Server performs the following main functions:

- Authenticates and authorizes Spotfire users.
- Provides access to analyses and data stored in the Spotfire library.
- Provides access to external data sources, including Oracle and SQL Server databases, or most JDBC sources, through information links.
- Makes sure that analyses are loaded with updated data according to schedules that are defined by the administrator.
- Provides storage (in the Spotfire database) for configurations, preferences, analyses, and so on.
- Manages the traffic through the Spotfire environment to optimize performance, and in accordance with rules that are defined by the administrator.
- Distributes software updates throughout the implementation.
- Monitors the health and activities of the Spotfire environment and provides diagnostic information both in the server interface and through downloadable logs.

SPOTFIRE DATABASE

Spotfire Server requires access to a Spotfire database.

The Spotfire database stores the information that Spotfire Server needs to control the Spotfire environment, including users, groups, licenses, preferences, shared analyses, and system configuration data.

You must have a database server up and running, preferably on a dedicated computer, before installing Spotfire Server. The Spotfire database can be installed on an Oracle Database server or a Microsoft SQL Server.

NODES AND SERVICES

To enable Spotfire web clients and Automation Services, install nodes in the environment.

With Spotfire Server installed, the installed Spotfire client, called Spotfire Analyst, can be used. To enable the use of Spotfire web clients and Spotfire Automation Services, one or more nodes must also be configured, preferably on dedicated computers.

For each node, the administrator enables Spotfire Web Player services, Automation Services, or both. Spotfire Web Player service is used to analyze in a web browser. Automation Services can be used to automate creation of analysis files, for example, with new data. The enabled services determine the functionality that the node provides to Spotfire end users, through the Spotfire Server. For failover and performance purposes, multiple service instances can be added on each node.

SCALABILITY

You can scale your Spotfire environment by adding or removing nodes and service instances almost without limits. Multiple instances of multiple client requests can be routed through multiple Spotfire servers, which in turn route traffic to backend worker nodes, each providing services such as Spotfire Web Player, Spotfire Automation Services, or Spotfire Statistics Services. This is achieved through smart routing and node managers.

Node managers manage one or more services running on the same server (a node). This enables:

- Lifecycle management starting, stopping
- Centralized deployment and configuration management
- Centralized logging
- Centralized monitoring

These features allow a node manager to provide the following benefits:

- Limit or optimize hardware resource utilization by routing users to an already open analysis file
- Isolate specific analysis files from the effects of other analysis files
- Isolate specific analysis files based on geographical, business unit, or security requirements
- Load an analysis on multiple service instances to handle a large number of concurrent users
- Support rolling upgrades by marking instances as scheduled for maintenance to divert traffic away from them

SPOTFIRE CLIENTS

Spotfire end users connect to the Spotfire Server using either an installed client or a web client. Spotfire Analyst is an installed client while Spotfire Consumer and Spotfire Business Author are web clients. To facilitate interactive analysis in a web browser, a Spotfire Web Player service generates visualizations that are displayed in the web browser. Depending on which of two licenses a user has, the web client will have different capabilities.

TIBCO SPOTFIRE ANALYST

TIBCO Spotfire Analyst is a Microsoft Windows 32/64 bit executable that is installed on users' local computers. It is a fully-featured client for working with data sources and creating complex analyses. This gives a user full authoring capabilities.

TIBCO SPOTFIRE CONSUMER

TIBCO Spotfire Consumer is a zero-install browser-based client that gives a user the ability to interact with analyses without any authoring capabilities.

TIBCO SPOTFIRE BUSINESS AUTHOR

TIBCO Spotfire Business Author is a zero-install browser-based client that gives a user the ability to interact with or create a new analysis project with limited authoring capabilities.

TIBCO SPOTFIRE IPAD/IPHONE APPLICATION

In addition, the native iOS apps available from the Apple App Store allow a user to have Spotfire Consumer like capabilities using an iPad and iPhone.

TIBCO SPOTFIRE AUTOMATION SERVICES

Within most organizations there are usually reports or analyses that need to be updated and distributed on a time schedule. Spotfire Automation Services is the tool to accomplish this batch processing by running automated jobs using TIBCO Spotfire Analyst.

It is easiest to think of Spotfire Automation Services as a copy of Spotfire Analyst that runs on your behalf without a display and executes a series of activities called tasks. Out of the box Spotfire Automation Services ships with 14 built-in tasks with the following general themes:

- Opening and saving analysis files in the Spotfire Library
- Updating visualization content
- Importing and exporting data, images, or PDFs
- Updating data connection information

Some examples of the business use cases that Automation Services fulfills are:

- Opening a quarterly sales dashboard, updating the data, saving a snapshot copy to the Spotfire Library, and emailing a DXP file as an attachment to each sales rep so they can view the data offline.
- Repeating the above use case but using parameters within Spotfire Automation Services job to create a separate DXP file for each region and sales rep.
- Performing the migration of an analysis file between the Test and Prod Spotfire environments, automatically updating the database connection information.

Custom tasks can also be developed using the Spotfire SDK.

TIBCO SPOTFIRE STATISTICS SERVICES

TIBCO Spotfire^{*} Statistics Services is a lightweight, flexible server that provides a communication layer, a service layer, and a TIBCO^{*} Enterprise Runtime for R (TERR), S- PLUS, or open-source R engine pool, among other features. Spotfire Statistics Services does not include user interface features (such as the TIBCO SpotfireS+^{*} Workbench or the S+ GUI). You can also use external engines, such as SAS and MATLAB if you have access to the corresponding software. TIBCO Spotfire Statistics Services can be deployed as a stand-alone server against which you can run analyses in one of its engines, or it can be deployed as part of a Spotfire stack.

ENVIRONMENT COMMUNICATION

All back-end communication in a Spotfire environment is secured by HTTPS/TLS, complying with current security standards and industry best practices (Figure 2).

Spotfire Servers listen to incoming traffic from installed clients and web clients on one HTTP or HTTPS port, the front-end communication port.

Spotfire Servers listen to traffic from services on the nodes on another HTTPS port, the back-end communication port.



Figure 2: Environment Communication in Spotfire Platform

The secured back-end communication is based on certificates. After an administrator has approved the new server or node, the certificates are issued automatically. Without a certificate, a server or a service on a node cannot make requests to, or receive requests from, other entities, except for when requesting a certificate.

After being installed, a node performs a join request to a specific, unencrypted HTTP Spotfire Server port that only handles registration requests. The node remains untrusted until the administrator approves the request by trusting the node. The Spotfire Server start page provides the tools to add nodes to the environment by explicitly trusting them, thereby issuing the certificates. When the node receives its certificate, it can send encrypted communication over the HTTPS/TLS ports, and with this it can start to send more than registration requests.

AUTHENTICATION AND USER DIRECTORY

Installed clients, as well as web clients, connect to the Spotfire Server. When users of either client log in to a Spotfire Server, two things happen before they get access: authentication and authorization (Figure 3).

Authentication is the process of validating the identity of a user. Once the identity is validated, the user is authorized in the user directory. Authorizing users determines what their access rights are within the Spotfire environment; in other words, what they are allowed to do.



Figure 3: Authentication and Authorization in Spotfire

If the username and password is used for authentication, they can be checked against the internal Spotfire user directory, a custom Java Authentication and Authorization Service module, or, the most common option, against an external LDAP directory. Spotfire has built-in support for Microsoft Active Directory and the Directory Server product family, which includes Oracle Directory Server, Sun Java Directory Server, and Sun ONE Directory Server. Other LDAP servers can also be used. The custom JAAS options allow integration with third-party security systems such as SiteMinder or customers' own custom security systems.

For single sign-on, Spotfire supports NTLM, Kerberos, and X.509 certificates. For anonymous, a preconfigured Spotfire user identity is used to authenticate with the Spotfire Server.

Regardless of how the user was authenticated, the process of authorization is the same. The Spotfire Server checks the Spotfire user directory to determine a user's privileges that control which functions and analyses they can access with the Spotfire clients.

Optionally, the user and group accounts in the Spotfire user directory can be configured to be synchronized from an external LDAP directory. Spotfire supports the same LDAP servers for directory synchronization as it does for authentication.

In the user directory, it is possible to organize the users in groups. The user and group information can later be used to assign permissions, licenses, preferences, and other elements to the different resources available within the Spotfire environment.

USERS AND GROUPS

All Spotfire users are registered in the Spotfire database where they are organized in groups.

The authentication method of your Spotfire environment determines how users are added to the database and where they are administered:

- If your Spotfire implementation is configured for authentication towards the Spotfire database, the administrator adds and administers user accounts directly in the database by using Spotfire Server and the Administration Manager tool. Administration Manager is accessed from Spotfire Analyst. Note that users, groups and group membership can also be managed from the Spotfire Server web-based Admin Console (Figure 5).
- If your implementation uses an external user directory such as LDAP, user accounts are added and administered in that context rather than in the server, and changes are automatically copied to the Spotfire database during synchronization.

Spotfire settings, including access to Spotfire features, which are controlled by licenses, are set at the group level, so all users necessarily belong to at least one group. Any user who is entered into the system automatically becomes a member of the Everyone group; this group cannot be deleted and will always contain all registered users.

In addition to the Everyone group, a user can belong to any number of groups and has access to all of the features that are enabled for those groups. Groups can be created and managed locally in the Spotfire database or synchronized from an external source such as an LDAP directory.

LICENSES AND PREFERENCES

Licenses determine which features a group of users should have access to, and preferences set the default behavior of the Spotfire clients.

Licenses also determine which features and functionality are available to Spotfire users. License data is stored in the Spotfire database. When a user logs into Spotfire, the user can only access the features that are enabled for the groups to which the user belongs, such as:

- Opening an analysis file from disk or from the Spotfire Library
- · Creating a table visualization to view the raw data within an analysis file
- Exporting data contained within an analysis file to disk or to the Spotfire Library
- Running statistical functions

There are many hundreds of licenses that allow very fine-grained control over what a user or group of users can do within the Spotfire platform. Licenses are managed in the Administration Manager UI of TIBCO Spotfire client.

Spotfire administrators can set a wide variety of preferences for the members of a group, such as a default color scheme for analyses or data optimization options.

Licenses and preferences are set in the Administration Manager in Spotfire Analyst. See the Administration Manager documentation for details on license and preference administration.

DEPLOYMENT

To deploy Spotfire software, the administrator places software packages in a deployment area on Spotfire Server, and assigns the deployment area to particular groups. Deployments are managed in the Spotfire Server web-based Admin Console (Figure 5).

If a new deployment is available when a user logs in to a Spotfire client, the software packages are downloaded from the Spotfire Server to the client. Deployments are used to:

- Set up a new Spotfire environment.
- Install a product upgrade, extension, or hotfix provided by Spotfire.
- Install a custom tool or extension. Administrators can create multiple deployment areas, such as "Production" and "Staging." This allows administrators to test new deployments before rolling them out to the entire client base, or to maintain different deployments for different groups of users.

SPOTFIRE LIBRARY

The Spotfire database contains the Spotfire Library. The library is accessible to Spotfire Analyst and web clients through the Spotfire Server, allowing users to easily share and reuse their work. The library stores Spotfire analyses, Spotfire data files, custom Spotfire data functions, information links, shared connections created with Spotfire connectors, and visualization color schemes. The library is organized into hierarchical folders, which are also used to control access to folder content. The administrator creates the folder structure and assigns groups with the appropriate read and write permissions to the folders.

LOAD BALANCING AND FAILOVER

A cluster of Spotfire Servers in an environment can be fronted by a load balancer to distribute the traffic to the servers (Figure 4). The load balancer should be able to detect if an instance becomes available or unavailable. Any load balancing technology supporting session affinity may be used. No load balancer is required between Spotfire Server and the nodes because the routing capability of Spotfire Server features built-in load balancing, enabling non-opened analyses to be loaded by the least used Spotfire Web Player service instance.



Figure 4: Load Balancing and Failover in Spotfire

SMART ROUTING

Spotfire provides routing capabilities within the environment. Smart routing provides several improvements over simple round-robin allocation:

- Capacity based optimization, route throttling, and monitoring of resource utilization
- Built-in (default) load-balancing that can be refined by static routing rules enabling fine-grained traffic control based on user, user groups, and analysis files
- Pre-loading and refreshing of analysis files on one or more instances of the Spotfire Web Player (formerly known as Scheduled Updates)
- Run-time clustering of routing data amongst all Spotfire Server nodes in the cluster using a distributed data-grid.

DEFAULT ROUTING

The descriptions below are for Spotfire Web Player, but the principles apply generally to all service types.

By default, any Spotfire Server in a cluster can send requests from clients to any Spotfire Web Player instance. Likewise, any Spotfire Web Player instance can access any Spotfire Server for library data or to execute information links.

After an analysis has been opened in a client, all subsequent requests for the session are forwarded to the instance that was used for the initialization; thus Spotfire Server routing maintains analysis session affinity.

Default routing improves capacity utilization by forwarding requests for a specific analysis file to the instance or instances of the Spotfire Web Player instance where it is already opened, thereby serving multiple users with the same service instance. Analysis data is also shared between users, so additional users accessing the analysis file will have low impact on performance. Furthermore, the router features built-in load balancing capabilities, enabling non-opened analyses to be loaded by the least utilized Spotfire Web Player instance.

The router also performs route throttling as a mechanism for proactive resource utilization and as a capacity saturation watchdog. Continuously evaluating memory and CPU utilization data from all Spotfire Web Player instances, it applies two pre-configured thresholds:

- STRAINED: When a Spotfire Web Player instance reaches this threshold, the router tries to find alternative instances for new requests. It may load an analysis on another strained instance, or itself, if no less-utilized instance is available.
- EXHAUSTED: An exhausted instance is unavailable for new requests, and strictly keeps serving only the current analysis and authoring sessions.

STATIC ROUTING

In addition to default routing, administrators can create named resource pools and assign any Spotfire Web Player service instances to them. The resource pools abstraction enables default routing to be altered by specific static routing rules. Rules can be specified for users, groups or specific analysis files and are defined and applied in priority order, similar to mail sorting rules. Rules can be sorted, enabled, disabled, and re-mapped to a different resource pool.

Also, administrators can attach schedules to routing rules that apply to analysis files, effectively turning a routing rule into a scheduled update. Thereby, the administrator can have the analysis pre-loaded on selected instances in a resource pool and have it refreshed at specified intervals.

Static routing rules are applicable only to initialization requests: Analyses Open or Create requests. In all other cases dynamic routing logic and routing attributes affinity will prevail.

Even if, by default, all Spotfire Web Player instances belong to one default unnamed resource pool, administrators can create named resource pools and assign any Spotfire Web Player instances to them.

SCHEDULING: A SPECIAL CASE OF STATIC ROUTING

A schedule can be attached to a static routing rule defined for an analysis file. By attaching a schedule and specifying additional scheduling attributes, administrators can pre-load analysis files on selected service instances in a resource pool. In addition, scheduling enables the administrator to specify a refresh interval.

LOGGING

In addition to the configurable logs for Spotfire Server, the nodes, and service instances, the action logs and system monitoring feature helps administrators keep an eye on the health of their Spotfire environment.

Action logs collect information about system events that are sent through a web service from Spotfire Analyst, Spotfire Automation Services, and Spotfire Web Player to the Spotfire Server. These event logs, along with those from the Spotfire Server itself, can be saved either to files or in a database.

System monitoring takes periodic snapshots of key metrics on Spotfire Server and Spotfire Web Player services, and stores this information in the same location as the action logs. The logs can then be analyzed in a Spotfire client.

Administrators have many options for how to configure this feature, including which events and system statistics should be logged, from which hosts logging information will be collected, and how the logs are pruned or archived.

WEB-BASED SPOTFIRE SERVER ADMIN CONSOLE

The web-based Spotfire Server Admin Console provides access to most administrative tasks and diagnostic information on your Spotfire environment (Figure 5).



Figure 5: Web-based Spotfire Server Administration Console

- Users and Groups: You can create or add users or groups (including the predefined administrator ones), assign deployment areas to groups, change names, passwords, and emails.
- Scheduling and Routing: You can schedule updates and monitor their status, date, and time, and create routing rules applicable to groups, users, or specific analysis files.
- In Nodes and Services: You can review the servers and services setup, add new nodes and services, upgrade or rollback existing ones, and create resource pools for routing rules (Figure 6).
- In Deployments and Packages: You can manage products, upgrades, extensions, and hotfixes by creating or altering deployment areas, adding distributions and packages, and so forth.
- In Monitoring and Diagnostics: You can monitor the system status, set logging levels, review logs, troubleshoot and download troubleshooting bundle, create memory dumps, and more.
- In Server Tools: You can download the configuration tool for Spotfire Server.

Library administration, licenses, and preferences are configured in the Administration Manager in the installed Spotfire Analyst client.

WHITEPAPER | 11

TIBC Spotfire Nodes & Servi	ces 🗸				•	? (Q) spotfir
				Your network U	ntrusted nodes Activity	Notifications
Network components	Status	From	ID		Time	
Select a view:	0	sf-naw2k1264-12.na.tibco.com	f25579ae-2fb4-4	3c3-a74d-3f717183a67f	2016-04-20 1	4:18:17
Nodes V	Details: No	ode can communicate successfully.				
Search network Q						[-]
sl-naw2k1284-12.na.tiboo.com j Automation Service 1	matalieu				Ir	istall new service
AUTOMATION_SERVICES Service	- Sta	tus Name	Capability	Deployment area	Version	
Web Player 2	 O 	Automation Service 1	8	Production	0	
Web Player 3	► 0	AUTOMATION_SERVICES .		Production	0	
[+] WEB_PLAYER Service	 o 	Web Player 2	-	Production	0	
	 o 	Web Player 3		Production	0	
	• •	WEB_PLAYER Service	WEB_PLAYER	Production	7.5.0 R133306	

Figure 6: Managing Nodes and Services in the Spotfire Server Admin Console.

EXAMPLE SCENARIO - USER OPENING AN ANALYSIS IN A WEB CLIENT

This is an example scenario of what happens in the Spotfire environment when a user opens an analysis in a web client.

- 1 The Spotfire web client user receives an email with a link to an analysis that contains interesting information.
- 2 When the link is opened, an ordinary http (or https) connection is set up from the browser to Spotfire Server. Because the environment is configured for username and password authentication, a login dialog appears.
- 3 If the username and password are correct, the user also needs to be listed in the user directory. Spotfire Server compares the credentials in the Spotfire database for verification.
- 4 A check is made to see that the user has the license privileges to see the analysis, which is stored in the library.
- 5 The analysis is not already loaded on any Spotfire Web Player service instance, so the routing logic of Spotfire Server selects the least used instance to load the analysis. The request is forwarded to this instance.
- 6 The Spotfire Web Player instance loads the analysis from the library.
- 7 Data in an analysis can be linked or embedded. In this example, the analysis contains linked data, loaded through information services. A request for the data goes back from the Spotfire Web Player service instance to a Spotfire Server.
- 8 After the analysis and its data are loaded, Spotfire Server acts as a proxy between the web browser and the Spotfire Web Player instance.
- 9 The user finds the analysis interesting and wants to add an extra visualization. Because the user has the business author license, the menu options to do so are visible.
- 10 After the user has updated the analysis, the user can send a link to interested parties.

CONNECTING TO DATA

The basic Spotfire environment provides three ways for clients to connect to data: by opening a local file, using a native Spotfire connector, or connecting through the Information Services function of the Spotfire Server. Analysts can combine data from multiple sources in a single Spotfire analysis.

Cloud enterprise customers can use all the same data sources and connection methods as we support in on-premises installations, although TIBCO Professional Services Group may need to be involved to set up secure connections.

LOCAL FILES

Spotfire Analyst users can open any file that can be accessed from their local machine or network for analysis. Business author users can upload files to Spotfire Web Player server to use in their analyses.

These are some of the file types Spotfire supports: Microsoft Excel workbooks, text files with comma-separated values, Microsoft Access databases, and SAS data files. For the full list, see the Spotfire Data Sources page.

NATIVE DATA CONNECTORS

Spotfire native connectors provide a mechanism for Spotfire clients to make a direct connection with enterprise data. By default Spotfire connectors will generate queries as analysts work with visualizations, filters, etc. These queries are executed in the database (often referred to as in-database) and the resulting data set is returned to and visualized in Spotfire. Analysts can also choose to load the entire or part of the data set (on-demand) into the Spotfire client's in-memory data engine.

Native data connectors save their meta data in two objects, the data source object and the data connection object. By storing these objects and sharing them through the Spotfire Library, IT can govern both connectivity details (stored in the data source) and data modeling details (stored in the data connection). These objects are by default not stored as separate objects and instead embedded in the analysis file. This allows quick access to data for ad hoc visual data discovery.

If all of the data needed is in a single database, analysts can use existing joins in a database or define their own by creating a virtual join view in Spotfire. SQL skilled analysts can also write (parameterized) custom SQL queries to optimize query performance and use source-specific functions.

When using in-database connections, the rate that Spotfire's visualizations are refreshed is directly dependent on how complex expressions are used in the visualization and the time it takes for the underlying database to process those queries. In certain cases it is preferred to load data into Spotfire and let the Spotfire data engine do the aggregations.

Spotfire has a long list of native connectors, with more being added with every release. For the latest information on system requirements for the connectors, please visit the Spotfire Data Connectors System Requirements page.

Here is the current list of TIBCO Spotfire Connectors:

- Amazon Redshift
- Apache^{*} Spark^{*} SQL (Apache Hadoop^{*}, Cloudera, Hortonworks, BigInsights, MapR)
- Cloudera Hive
- Cloudera Impala
- Cisco ("Composite") Information Server
- Hortonworks (Apache Hive', BigInsights Hive, MapR Hive)

- HP Vertica
- IBM DB2
- IBM Netezza
- Microsoft SQL Server
- Microsoft SQL Server Analysis Services
- OData
- Oracle
- Oracle Essbase
- Oracle MySQL
- Pivotal Greenplum
- Pivotal HAWQ
- PostgreSQL
- Salesforce.com
- SAP HANA
- SAP BW (NetWeaver Business Warehouse)
- Teradata
- Teradata Aster

TIBCO SPOTFIRE INFORMATION SERVICES

Using Spotfire Server Information Services is another option for connecting to enterprise data. In this case, the Spotfire Server makes connections to data sources on the client's behalf using information links saved in the Spotfire Library. The data is streamed through the Spotfire Server to the clients and is not loaded in the memory of the server but the memory of the clients (Analyst, Web Player, Automation Services).

The data sources available out of the box are Oracle, Microsoft SQL Server, Teradata, Sybase, SAS/Share, MySQL, and IBM DB2. On-premises AWS and cloud enterprise customers can also add custom JDBC source types.

For the list of data sources and more details on how to configure them, see the Spotfire Server Installation and Configuration Manual and have a look at the Data Source Templates section of the Advanced Procedures chapter.

TIBCO SPOTFIRE ADVANCED DATA SERVICES

Along with the three methods for accessing data that Spotfire provides out of the box, organizations can also implement an add-on product called Spotfire^{*} Advanced Data Services.

Spotfire Advanced Data Services offers market leading data virtualization and caching across web services and virtually any JDBC compliant data sources. Five add-ons are available for connectivity to SAP BW, SAP ERP, Salesforce. com, Siebel eBusiness Applications, and Oracle E-Business Suite. For the full list of supported data sources, please refer to the Spotfire Advanced Data Services system requirements on the Spotfire support website.

In an environment that includes Spotfire Advanced Data Services, clients can use a native in-database connector or Information Services to connect to a Spotfire Advanced Data Services server. Spotfire Advanced Data Services then connects to the data source and returns the required data to Spotfire.

TIBCO SPOTFIRE ATTIVIO AIE

The add-on product TIBCO Spotfire" Attivio" AIE" extends the scope of Spotfire analytics to include human-created information and unstructured text. For example, full text search can be built into Spotfire applications and dashboards against content loaded into AIE. The product is seamlessly integrated with Spotfire Information Services. Structured and unstructured data is blended to provide answers from sources such as PDF documents, e-mails, XML files, RSS feeds, and many database types. Text analytics tools such as sentiment analysis, key phrase extraction, and entity extraction and classification functionalities are provided out of the box.

LINKING, EMBEDDING, AND SAVING DATA

Once data has been brought into Spotfire, there are a number of options for how it is handled. The default option is for data tables to be linked to the original source. The data will be reloaded automatically when the analysis is opened, which requires all viewers to have access to the data source.

Alternatively, data that was loaded in the memory of the Spotfire client can be embedded in the analysis. In this case, the data will not be reloaded when the analysis is opened. Viewers can choose to reload the data manually if they have access to the data source.

Lastly, all or part of the data set can be saved to the Spotfire Library or exported as a file for use in other analyses. (You can also save the entire analysis, of course!) Analysts can select different options for the various data tables in an analysis.

Finally, analyses can be scheduled for pre-load in memory on the server. This allows thousands of consumers to instantly open analysis files without waiting for data to be loaded from the data source. Data is shared in a smart way to keep memory consumption to a minimum.

SPOTFIRE CLOUD OFFERINGS

SPOTFIRE CLOUD

TIBCO Spotfire^{*} Cloud is a software-as-a-service offering intended for individual business users and small teams who need Spotfire but want TIBCO to manage infrastructure. It consists of Spotfire Cloud downloadable client and web client for authors, editors, and consumers.

SPOTFIRE CLOUD ENTERPRISE

TIBCO Spotfire^{*} Cloud Enterprise is a platform-as-a-service offering intended for companies with tens, hundreds, and thousands of users who need the full Spotfire platform, but want TIBCO to manage their environment and upgrades. It is a full Spotfire platform on the latest software version hosted on AWS and managed by TIBCO. It can include automation services, statistics services, Spotfire Automation Services, etc.

SPOTFIRE FOR AWS

Spotfire for AWS is an infrastructure-as-a-service offering intended for IT administrators and developers who want to rent TIBCO infrastructure and manage their own Spotfire deployment and upgrades. It is the Spotfire platform hosted on AWS infrastructure and available for rental by the hour.



Global Headquarters 3307 Hillview Avenue Palo Alto, CA 94304 +1 650-846-1000 TEL +1 800-420-8450 +1 650-846-1005 FAX www.tibco.com TIBCO Software empowers executives, developers, and business users with Fast Data solutions that make the right data available in real time for faster answers, better decisions, and smarter action. Over the past 15 years, thousands of businesses across the globe have relied on TIBCO technology to integrate their applications and ecosystems, analyze their data, and create real-time solutions. Learn how TIBCO turns data—big or small—into differentiation at www.tibco.com.

©2016, TIBCO Software Inc. All rights reserved. TIBCO, the TIBCO logo, and Spotfire are trademarks or registered trademarks of TIBCO Software Inc. or its subsidiaries in the United States and/or other countries. Apache Hadoop, Hive, and Spark are trademarks of The Apache Software Foundation in the United States and/or other countries. Attiviou and AlE are trademarks of Attivio, inc. All other product and company names and marks in this document are the property of their respective owners and mentioned for identification purposes only. 06/07/16