Spotfire®

Spotfire for the enterprise

An overview for administrators

This white paper is intended for those wanting information on the Spotfire® platform administration and deployment capabilities: architecture, data connectivity, security, scalability, administration, and performance.

What is Spotfire?

Spotfire is the visual data science platform for engineers, scientists, and subject matter experts that makes smart people smarter by combining interactive visualizations and advanced analytics to solve complex industry-specific business problems.

The Spotfire visual data science platform uses a powerful visual-first approach to data science, combining data wrangling, visualization, modeling, and predictive analytics in a single, collaborative platform. It breaks down silos between tools and teams—freeing data scientists from time-consuming data prep and giving them more time to focus on meaningful analysis. At the same time, business users can easily access insights with just a few clicks.

Built for data-driven innovation, Spotfire champions this visual-first approach to data science, making it fast and easy to explore, visualize, and operationalize data across the enterprise. It combines market-leading visual analytics, data science, and data wrangling to allow experts to analyze data at-rest, in-motion, and at-scale—solving problems that require human creativity and modern computing.

Whether you're building predictive models, analyzing real-time streaming data, or uncovering trends through interactive dashboards, Spotfire brings data science and business insight together—visually and intuitively. Enjoy a point-and-click user experience where engineers, scientists, and other experts can rapidly build sophisticated, interactive visual data science applications and deploy those to thousands of end users, enhancing business-critical operations and daily decision-making. Trusted by leading organizations worldwide, Spotfire helps teams accelerate discovery, scale impact, and unlock the full value of their data.

Spotfire platform

Spotfire has end-user-focused products that can be used standalone or in combination with Spotfire® Enterprise, the platform for enterprise deployment within your own infrastructure, cloud, or data center.

The end-user Spotfire® products are:

- **Spotfire® Analytics** provides interactive visualization, advanced analytics, and AI-powered insights, helping users to explore data and solve their analysis challenges. Spotfire Analytics is available as an installed Windows application or using a web browser.*
- **Spotfire® Data Science** includes everything in Spotfire Analytics, and adds specific analytics capabilities for Energy, Manufacturing, and other industries, additional built-in data science capabilities, and all statistics and machine learning features of Spotfire Statistica®. Spotfire Data Science is available as an installed Windows application or using a web browser.*

Note: Accessing the application using a web browser requires Spotfire® Enterprise. Not all functionality may be available using a web browser.

For medium and large organizations, **Spotfire Enterprise** enables publishing and sharing analyses and analytic applications with others across the organization with robust security and governance.

Spotfire Enterprise provides the required capabilities for enterprise deployments:

- **Centralized management:** A web-based, centralized interface allows administrators to administer users, groups and licenses, and system configurations.
- **Analysis sharing and collaboration:** Publish, share, and embed analytics across departments using the Spotfire library.
- **Security and governance:** Enforce enterprise-grade security with role-based access control, encryption, and auditability, ensuring compliance with internal and regulatory standards.
- **Scalability and performance:** Support for large volumes of data and thousands of concurrent users without compromising responsiveness.
- **Flexible integration and extensibility:** Data connectivity to a wide range of databases, data lakes, APIs, and third-party systems—on-premises or in the cloud.
- **Automation and scheduling:** Refresh data, generate reports, and trigger analytics workflows automatically based on time or events, improving operational efficiency.
- **Statistics and machine learning services:** Share analytic applications using Python and R scripts natively, supporting predictive and statistical analysis.

In addition, there are two optional add-ons to Spotfire® Enterprise:

- **Spotfire® Enterprise Data Stream** enables visual analytics for streaming data.
- **Spotfire® Enterprise Advanced Data Services** provides a virtual data layer unifying data access for disparate data sources.

You can read more about these add-ons at Spotfire Enterprise Add-ons.

Note: Spotfire® Enterprise consolidates and modernizes analytics infrastructure by unifying capabilities previously spread across multiple offerings—replacing Spotfire® Server, Spotfire® Automation Services, Spotfire® Statistics Services, Spotfire® Consumer, and more—into a single, scalable platform.

Architecture

The main Spotfire components in Spotfire Enterprise are:

- **Spotfire clients:** Spotfire end users connect to the Spotfire Server using either an installed client, a web client, or mobile/tablet app.
- **Spotfire platform:**
  - The Spotfire Server is the administrative center of the Spotfire platform.
  - The Spotfire database stores the Spotfire platform's configuration, analyses, etc.
  - Spotfire has several specialized Spotfire services for enabling web client access, automation tasks, advanced modeling via data functions, etc.

For a quick overview, see the article "Introduction to the Spotfire environment."

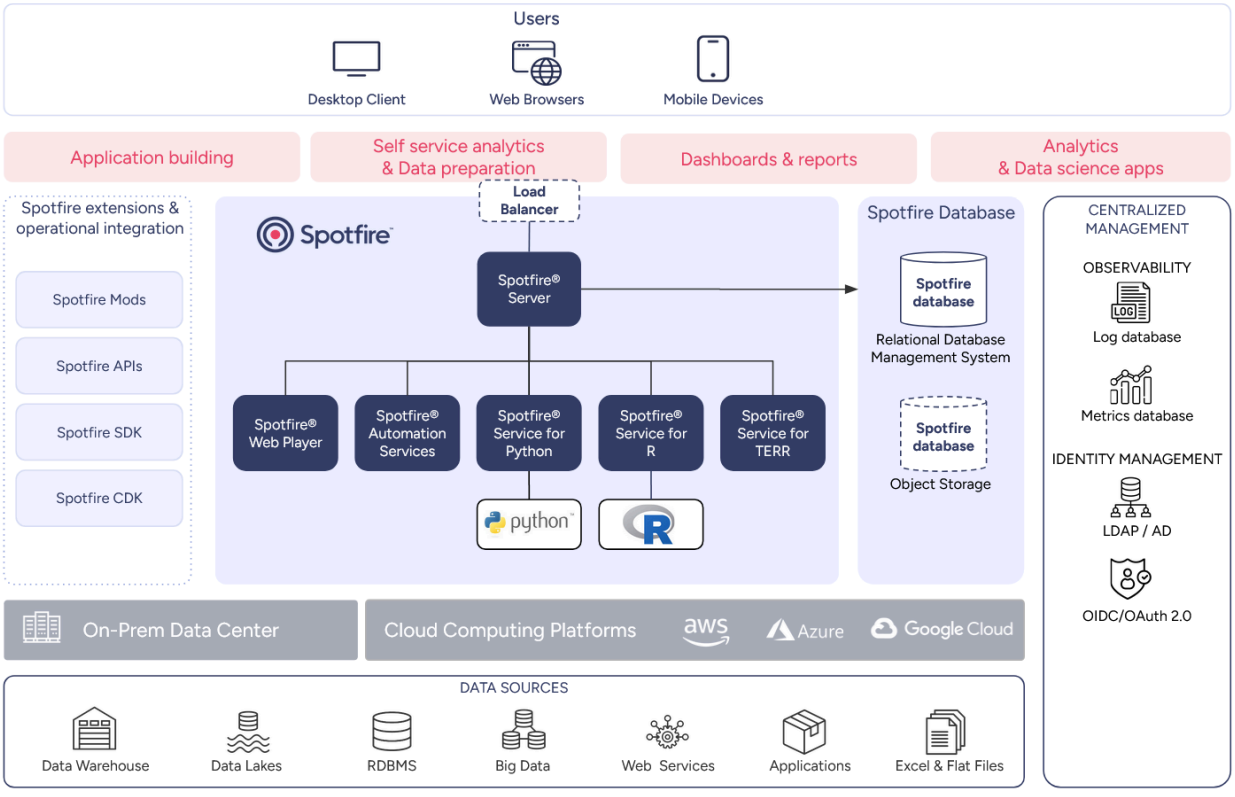

Figure 1: Spotfire can be deployed anywhere—on-premises or in any cloud environment—offering centralized management, the ability to extend through Spotfire extensions, and seamless operational integration. Furthermore,

it connects to a multitude of data sources, enabling organizations to harness the power of visual data science even further.

Note: The Load Balancer is optional and recommended when using multiple Spotfire server instances.

You can extend, customize, and automate Spotfire using:

- **Spotfire® Mods:** A lightweight framework enabling rapid development of reusable custom visualizations and actions within Spotfire.
- **Spotfire APIs:** A set of REST and low-level language APIs that allow interaction with Spotfire features, automation of tasks, and integration of Spotfire into your broader system for identity management, logging, metrics, and more.
- **Spotfire SDK:** A C# / .NET-based framework for building comprehensive Spotfire extensions, including visualizations, panels, tools, and data connectivity.
- **Spotfire CDK (Cloud deployment kit for Spotfire):** A set of recipes for building and customizing Spotfire container images and Helm charts. It provides a reference and starting point for deploying Spotfire using containers on a Kubernetes cluster.

You can manage the Spotfire platform via:

- Spotfire CLI
- Spotfire web user interface
- Spotfire APIs

Spotfire clients

Spotfire end users connect to the Spotfire Server using either an installed client, a web client, or the mobile/tablet app. The Spotfire clients present different capabilities depending on the platform they are running on and the user´s licenses:

- The Spotfire desktop client is a Microsoft Windows application that can be installed on users' local computers or shared servers accessed via remote application servers.
- The Spotfire iOS/Android apps, available from the Google Play Store / Apple App Store, allow users to interact with analyses using mobile or tablet devices.
- Spotfire is also available from any web browser when using the Spotfire platform.

Spotfire Server

The Spotfire® Server is the administrative center of the Spotfire platform. The Spotfire Server:

- Facilitates the configuration and administration of the Spotfire platform.
- Enables sharing of analytic assets and collaboration with the Spotfire Library.
- Provides access to external data sources through Information Links.
- Routes the traffic through the Spotfire platform.
- Distributes software updates to the Spotfire services and clients.

The Spotfire Server supports automatic clustering. All the Spotfire server instances are always active in a Spotfire clustered server. Typically, a load balancer with session affinity is added in front of the clustered Spotfire servers to help distribute the workload among the Spotfire servers.

Clustered Spotfire servers provide high availability and scalability for enterprise applications.

Spotfire database

The Spotfire server requires a database for centralized persistent storage of the Spotfire platform's configuration information.

The Spotfire database supports the following Relational Database Management Systems (RDBMS): PostgreSQL, Oracle Database server, and Microsoft SQL Server. You can use them on-premises or via the corresponding database services in the major cloud service providers.

Optionally, you can also use supported object storage for some of the Spotfire assets (the Library items) and keep only their metadata in the RDBMS. Using an object storage provides better storage performance and scalability compared to using a database for large items like analysis with embedded data, and it is also more cost-effective.

See the article "Spotfire system requirements" for more details on which vendors and versions are supported.

Spotfire services

The Spotfire platform has a services-oriented architecture, with several specialized services of two different types native to Spotfire Enterprise:

- Visualization services:
    - **Spotfire® Web Player:** A remote application server that enables analysis consumption from web browsers and mobile devices.
        - The Spotfire Web Player generates visualizations displayed in the Spotfire web clients and mobile apps.

        - To optimize resource allocation, the Spotfire Server handles client session routing towards Spotfire Web Player instances. Resource Pools and Routing Rules can be defined for fine-grained QoS control for different users, groups, or analyses.
    - **Spotfire® Automation Services** is a service for automatically executing scheduled jobs within your Spotfire environment. The most typical usages are automatic report creation, distribution, and caching of report data.
- Data function services:
    - **Spotfire® Service for Python (Python service):** Provides remote execution of Python data functions for users from web clients and mobile devices.
    - **Spotfire® Service for R (R service):** Provides remote execution of R data functions, R predictive analytics, or R custom expressions for users from web clients and mobile devices.
    - **Spotfire® Enterprise Runtime for R-Server Edition (aka TERR service):** Provides remote execution of TERR data functions, TERR predictive analytics, or TERR custom expressions for users from web clients and mobile devices.

No load balancer is required between the Spotfire Server and the service nodes because the Spotfire Server's routing capability features built-in load balancing. This enables analyses to be routed to the proper service instances based on resource utilization, session affinity, and custom resource pool routing rules.

The Spotfire platform's services-oriented architecture allows you to deploy the services you need and the number of instances of each Spotfire service you need. Each Spotfire service can scale independently based on default or custom metrics.

Deploy anywhere

Any major cloud provider

Spotfire can be deployed anywhere: On-premises or in any private, public, hybrid, or multi-cloud environment. Additionally:

- For very restricted and sensitive deployments, you can deploy the Spotfire platform on-premises on your private data center.
- For higher elasticity, you can also deploy Spotfire on any major cloud service provider: AWS, Azure, and Google Cloud, and integrate Spotfire natively with their most common cloud services.
- You can deploy Spotfire using a hybrid or multi-cloud setup for widespread availability and cost flexibility.

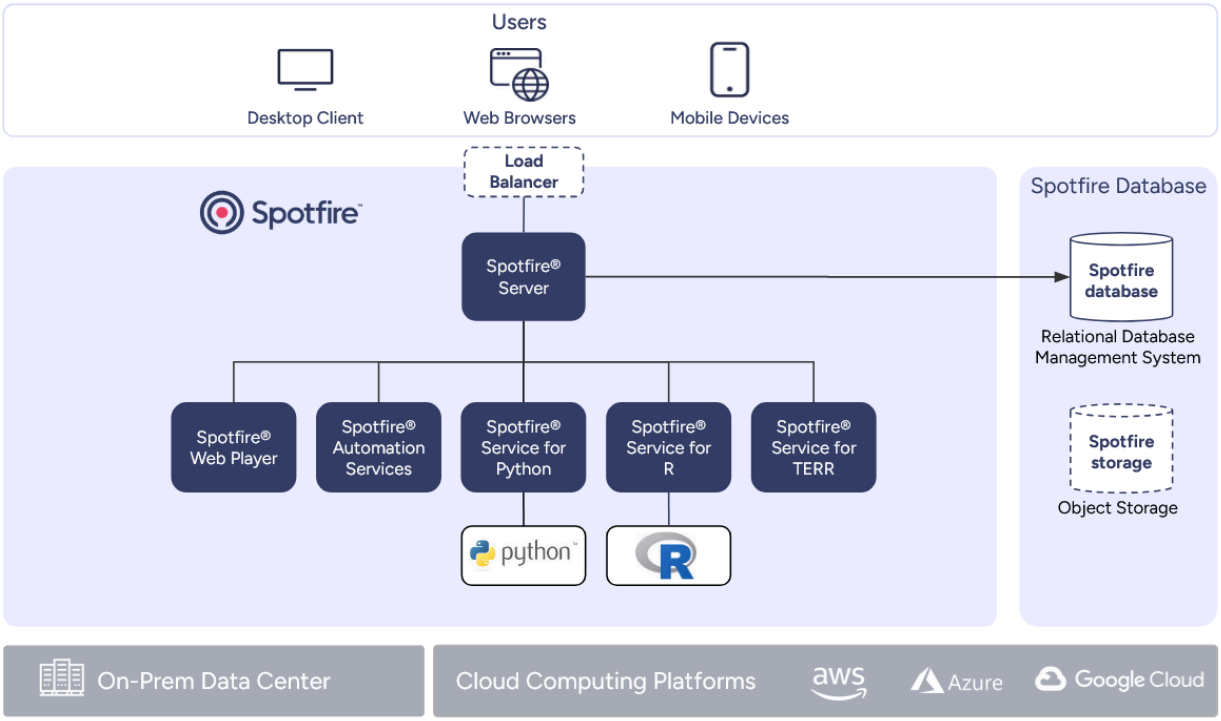

Figure 2: Spotfire can be deployed anywhere: On-premises or in any private, public, hybrid, or multi-cloud environment.

Any major operating system

You can deploy the Spotfire platform services on Microsoft Windows Server or GNU/Linux distributions, as Debian, Ubuntu, and Red Hat Enterprise Linux.

You can consolidate your deployments and use a single operating system for all Spotfire services, making managing licensing and security easier.

When deploying GNU/Linux distributions, enterprises can achieve significant cost savings, better resource utilization, and increased security. On the negative side, there are a small number of known differences and limitations when running web player or automation services on Linux, which might not be relevant for most of the deployments.

See the article "System requirements for Spotfire products" for the complete list of supported operating systems and versions.

Virtualization and containerization

The Spotfire platform is based on modern distributed architecture and is composed of several independent services.

You can deploy the Spotfire platform using bare-bone machines, Virtual Machines, or containers.

Spotfire supports deployment using containers and Helm charts on Kubernetes, offering a more efficient alternative to traditional Virtual Machines.

Spotfire on Kubernetes

You can use Spotfire on Kubernetes®, the set of official pre-built container images and Helm charts that enable you to deploy the Spotfire platform on any certified Kubernetes distribution.

The recipes and examples for creating these official Spotfire containers and Helm charts are available as open-source on the Cloud Deployment Kit for Spotfire (Spotfire CDK). This further allows organizations to extend and customize them to meet their specific requirements.

Simpler installation and upgrades

With Spotfire on Kubernetes, installations and upgrades are easier than ever.

With the Spotfire platform Helm chart, you can deploy a complete Spotfire platform environment with just one simple one-line command. All preparation steps are handled via Helm chart configuration values and predefined templates. For example, when installing or upgrading to a new version of Spotfire, the Spotfire database is automatically upgraded to the new schema.

Product openness

Spotfire supports standard interfaces that integrate with the most commonly used cloud-native components.

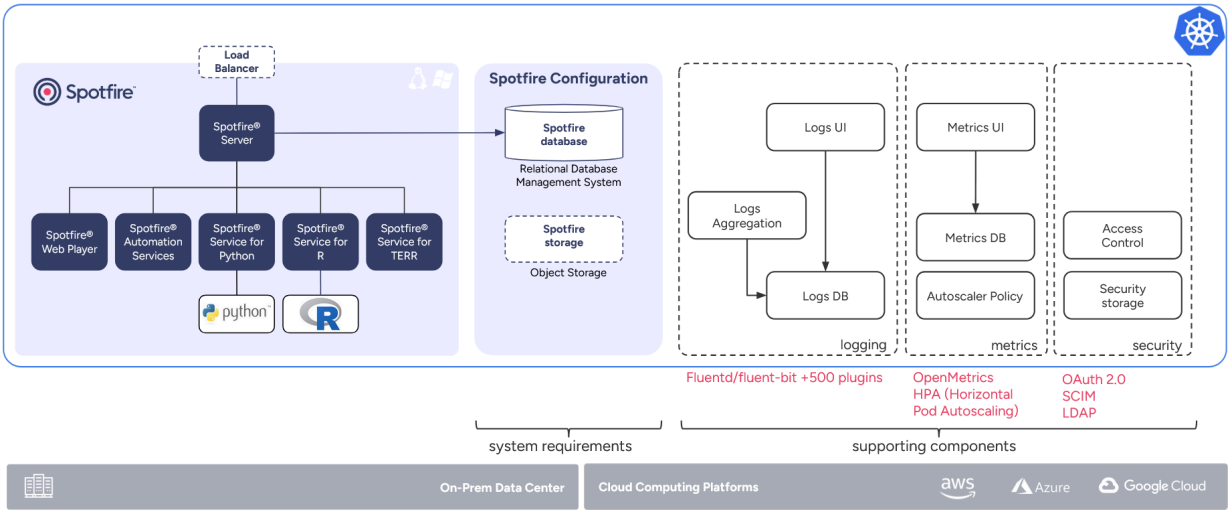

Figure 3: Spotfire integrates with key ecosystem components like logging, monitoring, and security—supporting standard interfaces that work across both traditional and cloud-native deployments.

Spotfire is built with standard observability by default:

- Spotfire metrics are available using OpenMetrics (the Prometheus metrics format), the industry standard for metrics, backed by the CNCF.
- Spotfire logs are pushed with Fluent-bit. This log forwarder decouples data sources from backend systems by providing a unified logging layer. Fluent-bit/Fluentd has 500+ plugins to connect it to most data sources and outputs.

With the Spotfire Helm charts, you get off-the-shelf integration for Spotfire with common CNCF observability applications, using, for example, the EFK stack (Elastic, Fluent-bit/Fluentd, Kibana), Prometheus, and Grafana.

For identity management, you can integrate with any application using common standards, such as OpenID/OAuth2.0, LDAP, or SCIM, among many others.

Auto-scaling

All Spotfire services can scale independently using predefined or custom metrics. On top of that, Spotfire services support static and dynamic smart routing for better resource utilization in distributed deployments.

Spotfire Helm charts enable autoscaling using the Kubernetes HPA (Horizontal Pod Autoscaler) component via direct integration with KEDA (Kubernetes Event-driven Autoscaling).

Instead of investing resources in development, writing scripts, and integrations to manage autoscaling, customers simply need to configure which metrics each Spotfire service will use to auto-scale. This allows for automatic "spinning-up / spinning-down," which adds Spotfire pods as loads increase and removes pods as loads decrease for better cost control, adjusting resource utilization to your business needs.

Enterprise features

System administration

The Spotfire platform provides diverse capabilities for easier management of enterprise environments with thousands of users and analyses:

- **Identity management:** For user authentication and authorization, group, roles, and preferences management.
- **Library:** To enable collaboration and reuse, the library is a centralized storage for all Spotfire Analysis files, assets, and metadata.
- **Scheduled updates:** This feature enables scheduled preloading and updating of analyses, making analyses with large data and heavy data functions ready for exploration when needed, without loading time.
- **Automation services (job scheduling):** Allows scheduling of custom automated tasks.
- **Data access:** Provides a centralized data access point and metadata management for relational data sources.
- **Monitoring:** Monitors the health and activities of the Spotfire environment and provides diagnostic information.
- **Audit:** Provides centralized collection of action logs for accountability, tuning, and troubleshooting purposes.

Identity management

Spotfire supports the most common user authentication protocols:

- Users can be checked against the internal Spotfire user directory, an external LDAP directory, or a custom Java Authentication and Authorization Service (JAAS) module for username and password authentication.
- For single sign-on, Spotfire supports web authentication (such as OpenID/OAuth and SAML), X.509 Certificates, and Kerberos.
- A preconfigured Spotfire user identity is used to authenticate with the Spotfire Server for anonymous authentication.

The Spotfire server checks the Spotfire user directory to determine which licenses, preferences, and permissions have been set for the user.

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

You can configure Spotfire Server to synchronize users, groups, or group memberships from one provider to another, by using LDAP or the System for Cross-domain Identity Management (SCIM) standard APIs.

Users, groups, licenses, and preferences

Users and groups

All Spotfire users are registered in the Spotfire database, where they are organized in groups. Administrators can set access to library folders for both groups and individual users. Groups can be added as subgroups to other groups, and administrators can build a hierarchy of groups to meet company requirements.

Product licenses and features

Administrators assign product licenses, such as Spotfire Analytics and Spotfire Data Science, to user groups. These product licenses assign the right to use a certain default set of capabilities and features in Spotfire, rather than the use of a specific application.

Administrators can further restrict the available capabilities and features for a user's group by enabling or disabling the specific license features available for that group.

Since a user can belong to several groups, when a user logs into Spotfire, the user can only access the license features enabled for the product license and the groups to which the user belongs.

Spotfire provides several hundred license features to allow precise, fine-grained control over what a user can do within the Spotfire platform, allowing administrators to create custom roles. Since groups can have subgroups and users can also be members of different groups, administrators can configure roles that take advantage of group mapping and inheritance.

Examples of license features:

- Opening an analysis file from disk or 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.

Preferences

Preferences set the default behavior of the Spotfire clients. Spotfire administrators can set various preferences per group, such as a default color scheme for analyses or data optimization options.

Spotfire Library

Why?

Analysts, data scientists, and other experts demand tight collaboration on complex analyses in enterprise environments. They also need to publish their results to share them with other users in a trusted and secure environment.

What?

The Spotfire Library provides centralized storage of Spotfire Analysis files, assets, and metadata.

> The Library stores Spotfire analysis files, Spotfire data files, Spotfire data functions, Spotfire data sources, scripts, color schemes, visualization Mods, action Mods, and many more item types.

Use cases

- The Library enables user collaboration. It allows users to share and reuse their work efficiently.
- The Library is organized into hierarchical folders, which are also used to control access to folder content by organizations and teams.
- The Library supports version history, which facilitates reverting changes, troubleshooting, and provides change traceability.
- The Library enables quick access to recent items, promoted assets, favorites, etc.

How does it work?

Administrators manage the Spotfire Library's folder permissions for the users and groups. Analysts can work with local files or with files stored in the Spotfire Library.

The Spotfire Library content is stored in the Spotfire database (a supported RDBMS, on-premises, or in a major cloud service provider). Administrators can also choose to store the Library artifacts on external persistent storage (e.g., AWS S3, Azure Blob Storage, or Google Cloud Storage) and use the Spotfire database only for the metadata.

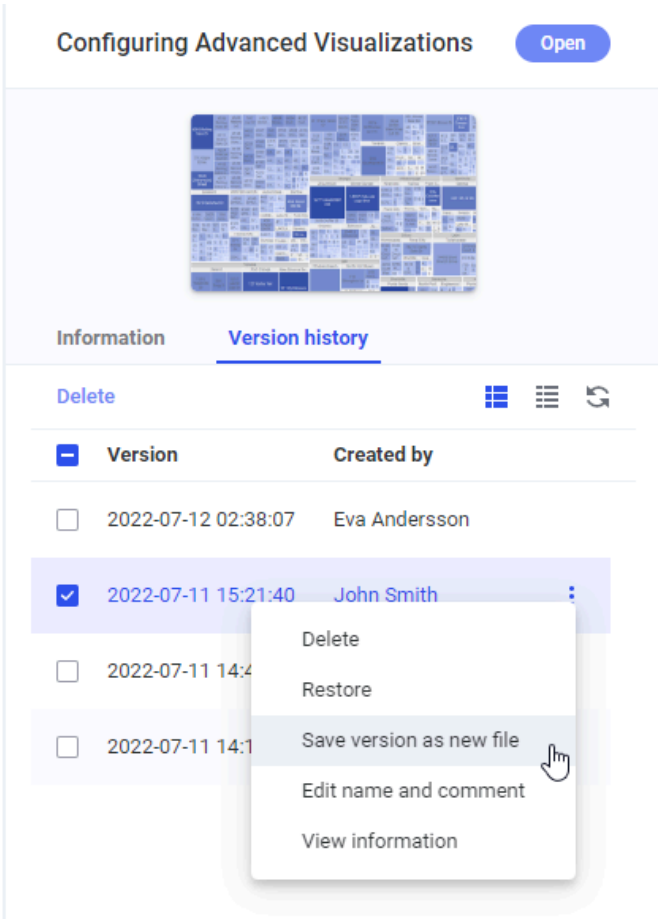

Figure 4: Easily save different versions of a file in the Spotfire Library and manage permissions centrally.

Scheduled updates

Why?

Some analyses require a long time for data preparation and loading. Administrators need a tool to schedule the loading so that the data is ready for users to consume at the expected time. Scheduled updates are often used for applications and dashboards that are used by large numbers of users and need to be available to the user very quickly.

What?

Scheduled updates enable scheduled preloading and analysis updates to a configured resource pool.

A scheduled updates job is a special case of routing rules where time events trigger the routing. You can specify how often the schedule update job is executed, repetition patterns, such as which day of the week, hourly, daily, weekly, etc., and when to begin and end the schedule repetition pattern.

Use cases

- Preload an analysis in a resource pool to handle an expected large number of users at a certain time.
- Preload a large analysis so it is ready at the expected times (e.g., before the morning meeting)
- Refresh the analysis continuously at regular intervals.
- Trigger an analysis refresh externally via Spotfire APIs, e.g., right after the data is updated in the data source, so that Spotfire always gets the latest data.

How does it work?

Users with the appropriate privileges can define schedules and routing rules (see below) for analyses in the Spotfire Library. The analyses and their data are then loaded into an available Web Player based on the resource pool or default pool. Users can then consume the analyses in their web browser without waiting for the data to load.

Automation services (job scheduling)

Automation Services allows you to schedule custom automation jobs.

An automation services job is an automated procedure that carries out multi-step tasks. Jobs can be created with the job builder in the Spotfire desktop and are stored in the Spotfire Library. A schedule specifies how often a job is executed. You can specify complex repetition patterns like which day of the week, hourly, daily, weekly, etc. Also, from when to begin and end the schedule repetition pattern.

Use cases

- Automatic update and distribution of analyses and reports on a time schedule to organizations, teams, or users (for example, a morning mail report with specific sections of an analysis with refreshed data to specific managers in PDF format every Monday and Wednesday at 07:00 AM).
- Automatic caching of data in analyses or the Spotfire Library for faster loading.

Besides triggering the automation service jobs with a custom schedule, administrators can also trigger the automation services job on demand, using the Spotfire API. This is useful, for example, for starting a job when certain conditions are met, like a data table has been updated with new data.

Automation Services ships with several built-in tasks like:

- Opening and saving analysis files in the Spotfire Library.
- Updating visualization content.
- Importing and exporting data tables, images, or PDFs.
- Updating data connection information.
- Executing IronPython scripts.
- *ForEach* to allow looping over report tasks.

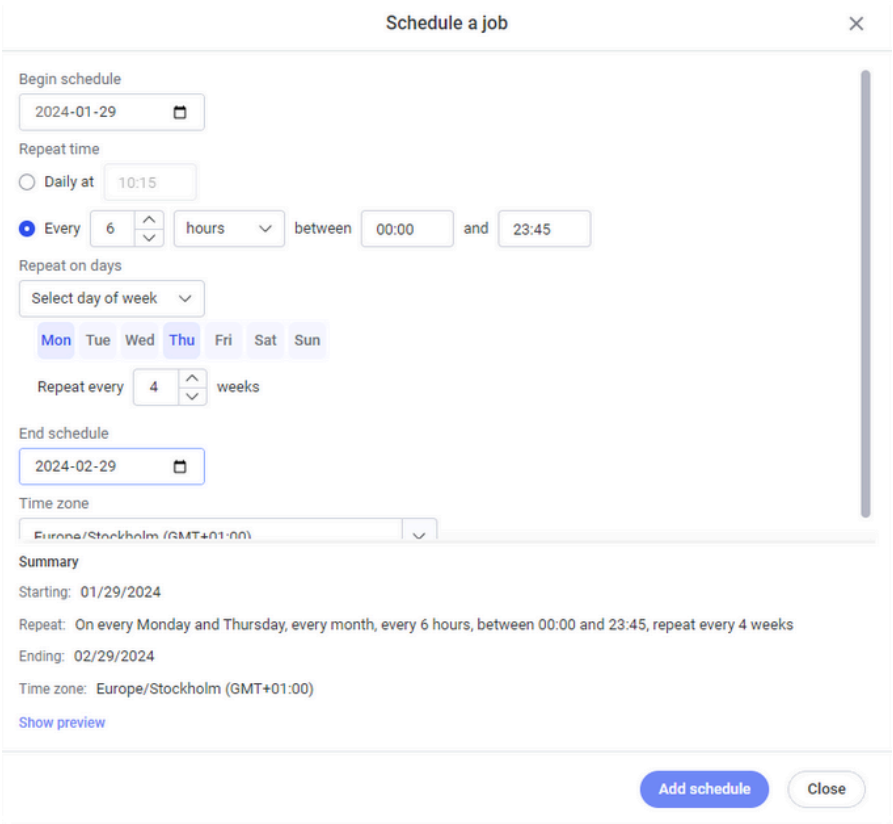

Figure 5: Using the Spotfire SDK, you can also develop and add custom tasks to fit your specific needs.

Services grouping

We mentioned that the Spotfire platform comprises several independent, specialized services that can scale independently. In addition, you can create different types of logical service groups for the Spotfire services in your environment for different purposes, like for distributing different software packages, for managing multiple geographical locations, for separating by quality of service, or for applying different authentication methods.

Using these different logical groupings, Spofire provides many solutions for common enterprise needs:

- **Deployment areas:** For easier distribution of software updates to services and clients.
- **Resource pools and routing rules:** To control resource utilization and guarantee service quality for the critical analyses or the most important consumers.
- **Sites:** To reduce latency for multi-geographic deployments and enable the use of different authentication methods within the same deployment.

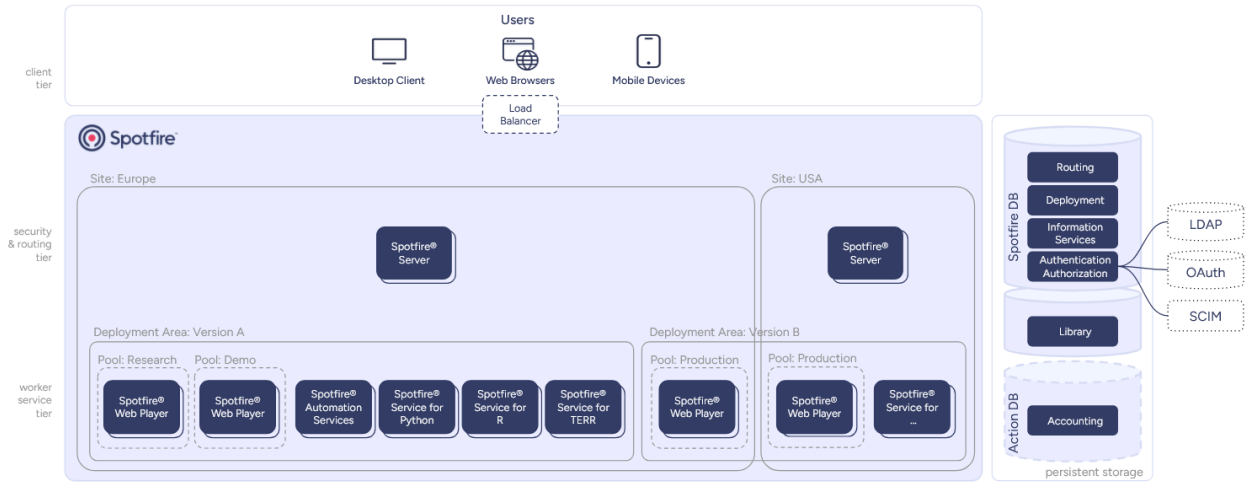

Figure 6: A resource pool is a logical group of Spotfire Web Player instances. Resource pools and routing rules can be used together to control the quality of service for different users, groups, or analyses.

Deployment areas

Why?

A common problem for administrators managing enterprise platforms is making sure all the servers are correctly updated with the right software versions and custom extensions. The solution for this is the Spotfire Deployment Areas.

What?

A deployment area is a software package repository to facilitate software deployment to a specific group of Spotfire services and desktop clients.

Deployment areas facilitate new software deployments and software upgrades, allowing administrators to manage different deployment areas with software packages for different user groups or to test new deployments before rolling them out to the entire client base.

A Spotfire deployment area stores and distributes Spotfire packages (.SPK files). A group of Spotfire packages (.SPK files) can be bundled into a Spotfire distribution (.SDN file).

Use cases

- Deploy a new Spotfire Platform environment.
- Deploy a Spotfire software upgrade or hotfix.
- Deploy custom extensions.
- Distribute software updates to Spotfire desktop clients.

How does it work?

An administrator can create different deployment areas and assign specific services (e.g., Spotfire Server, Web Player, Automation Services) to each. The administrator places software packages in a deployment area to deploy Spotfire software. This way:

- Services belonging to that deployment area can be updated from the web-based admin GUI.
- The client can be automatically updated if new software is available when a user logs in from an installed Spotfire desktop client.

Administrators can create multiple named deployment areas, such as “Production,” “Staging,” and so on. 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.

Administrators can copy all the Spotfire packages from a deployment area into a new deployment area in the same Spotfire Server or export them as a Spotfire distribution (.SDN) to another Spotfire Server. This allows for easy replication of environments.

Resource pools and routing rules

Why?

In environments with thousands of analyses and thousands of users, it is important to control resource utilization and guarantee the quality of service for the critical analyses or the most important consumers.

What?

A resource pool is a logical group of Spotfire Web Player instances. IT administrators can modify the number and the size of the Web Player services, and Spotfire administrators can modify the number of Web Player instances on each Web Player service.

A routing rule specifies the resource pool on which an analysis opens, the decision object (user, group, or file), and the priority.

Resource pools and routing rules can be used together to control the quality of service for different users, groups, or analyses.

Use cases

- Guarantee reserved resource allocation for critical users, groups, or analyses.
- Isolate non-production analyses or heavy analyses within their own sandboxed pool to avoid any noisy neighbour problems.
- Define a new resource pool for validating a new custom extension, data driver, or version of Spotfire Web Player.
- Distribute the load of a highly requested analysis among different pools. Load an analysis on several Spotfire Web Player instances to handle a large number of users.

Spotfire resource-based routing

When a request is made to open a specific analysis file, the Spotfire routing mechanism goes through the configured routing rules in priority order. The applicable rule could be a routing rule, a scheduled update (SU) rule, or the default routing rule, resulting in one or more Spotfire Web Player service instances to choose between.

When service instances are available for selection, the Spotfire Server routing mechanism focuses on optimizing resource utilization (using as many of the instance resources as possible and not distributing user sessions evenly among different instances) by default. Administrators can tune several resource utilization thresholds to change the balance between performance (response time) and efficiency (resource utilization).

Please see the articles "Understanding routing" and "Tuning the load distribution between Spotfire Web Player instances" for more details.

Sites

Why?

Multinational enterprises might benefit from using Spotfire services near their region, not only for quicker analysis, interaction, and discovery, but also to comply with different data regulations and security policies.

Also, within a multinational enterprise, business units, subsidiaries, companies, and affiliates require different authentication methods.

What?

A Site is a logical group of Spotfire Servers and Nodes co-located within the same country or region to reduce latency for multi-geographic deployments.

Sites also enable the use of different authentication methods within the same deployment.

> Spotfire sites help to reduce latency for multi-geographic deployments, comply with specific country specific regulations and allow different authentication methods for each region.

Use cases

- Route user requests from a particular office to the services closest to that office. This reduces the impact of network latency between servers located in different geographic regions.
- Enable different authentication methods for various user groups within the same Spotfire deployment. For example, internal users may use LDAP authentication while external users, such as customers and partners, may use OAuth.

How does it work?

Each site includes one or more Spotfire Server clusters along with their connected Spotfire services. Those nodes can only communicate within their own site.

If required, the administrator can assign different authentication methods for each site.

All Sites share the same Spotfire database (Library, routing, deployment areas, etc.) and action database (user traceability).

Learn more about Spotfire

Spotfire Enterprise offers unparalleled value by combining advanced analytics, interactive visualizations, automation, and governance in one powerful, scalable platform.

Ready to get smarter with Spotfire visual data science? Talk to an expert today at spotfire.com/contact-us.

Cloud Software Group
Headquarters
851 W Cypress Creek Rd.
Fort Lauderdale, FL 33090
www.spotfire.com

Spotfire® is a visual data science platform that makes smart people smarter by combining interactive visualizations and advanced analytics to solve complex, industry-specific business problems.

8July2025