

IBM SPSS Analytic Server  
Version 2

*Administrator's Guide*



**Note**

Before using this information and the product it supports, read the information in "Notices" on page 11.

**Product Information**

This edition applies to version 2, release 0, modification 0 of IBM SPSS Analytic Server and to all subsequent releases and modifications until otherwise indicated in new editions.

---

## Contents

### Chapter 1. Overview . . . . . 1

Architecture . . . . . 2

What is new for administrators in version 2 . . . . . 2

### Chapter 2. Tenant management. . . . . 3

Naming rules . . . . . 4

### Chapter 3. Getting users started . . . . . 5

### Chapter 4. Analytic Server job names . . . . . 7

### Chapter 5. Troubleshooting . . . . . 9

Logging . . . . . 9

Version information . . . . . 9

Log collector . . . . . 9

Common issues . . . . . 10

### Notices . . . . . 11

Trademarks . . . . . 13



---

## Chapter 1. Overview

IBM® SPSS® Analytic Server is a solution for big data analytics that combines IBM SPSS technology with big data systems and allows you to work with familiar IBM SPSS user interfaces to solve problems on a previously unattainable scale.

### Why big data analytics matters

Data volumes collected by organizations are growing exponentially; for example, financial and retail businesses have all customer transactions for a year (or two years, or ten years), telco providers have call data records (CDR) and device sensor readings, and internet companies have the results of web crawls.

Big data analytics is needed where there exists:

- A large volume of data (terabytes, petabytes, exabytes), especially when it is a mixture of structured & unstructured data
- Rapidly changing/accumulating data

Big data analytics also assists when:

- A large number (thousands) of models are being built
- Models are frequently built/refreshed

### Challenges

The same organizations that collect large volumes of data often have difficulty actually making use of it, for a variety of reasons:

- The architecture of traditional analytic products are not suited to distributed computation, and
- Existing statistical algorithms are not designed to work with big data (these algorithms expect the data to come to them, but big data is too costly to move), thus
- Performing state of the art analytics on big data requires new skills and intimate knowledge of big data systems. Very few analysts have these skills.
- In-memory solutions work for medium-size problems, but do not scale well to truly big data.

### Solution

Analytic Server provides:

- A data-centric architecture that leverages big data systems, such as Hadoop Map/Reduce with data in HDFS.
- A defined interface to incorporate new statistical algorithms designed to go to the data.
- Familiar IBM SPSS user interfaces that hide the details of big data environments so that analysts can focus on analyzing the data.
- A solution that is scalable to any size problem.

---

## Architecture




Figure 1. Architecture

Analytic Server sits between a client application and Hadoop cloud. Assuming that the data resides in the cloud, the general outline for working with Analytic Server is to:

1. Define Analytic Server data sources over the data in the cloud.
2. Define the analysis you want to perform in the client application. For the current release, the client application is IBM SPSS Modeler.
3. When you run the analysis, the client application submits an Analytic Server execution request.
4. Analytic Server orchestrates the job to run in the Hadoop cloud and reports the results to the client application.
5. You can use the results to define further analyses, and the cycle repeats.

---

## What is new for administrators in version 2

### Analytic Server console

#### New layout

The layout has been changed so that the Tenants, Projects, and Data Sources pages are accessed via a home page, rather than accordions.

#### Tenants

Each tenant can be associated with a different authentication registry.

---

## Chapter 2. Tenant management

Tenants provide a high-level division of users, projects, and data sources so that objects cannot be shared between tenants. Each user accesses the system in the context of a tenant to which they are assigned.

You manage tenants, and assign users to tenants, in the Analytic Server console. The view of the Tenants page depends upon the role of the user that is logged in to the console:

- The "super user" administrator that is set up during installation is the tenant manager. Only this user can create new tenants and edit the properties of any tenant.
- Users with the Administrator role can edit the properties of the tenant they are logged in to.
- Users with the User role cannot edit tenant properties. The Tenants page is hidden from them.

Administrators can access the Projects and Data sources pages and manage any project or data source for cleanup and administration. See the *IBM SPSS Analytic Server User's Guide* for more information.

### Tenant listing

The main Tenants page displays the existing tenants in a table. Only the "super user" administrator can make edits on this page.

- Click a tenant's name to display its details and edit its properties.
- Click a tenant's URL to open the console in the context of that tenant.

**Note:** You will be logged out of the console and will need to log in with valid credentials for the tenant.

- Type in the search area to filter the listing to display only tenants with the search string in their name.
- Click **New** to create a new tenant with the name you specify in the **Add new tenant** dialog. See "Naming rules" on page 4 for restrictions on the names you can give to tenants.
- Click **Delete** to remove the selected tenant(s).
- Click **Refresh** to update the listing.

### Individual tenant details

The content area is divided into several collapsible sections.

#### Details

**Name** An editable text field that displays the name of the tenant.

#### Description

An editable text field that allows you to provide explanatory text about the tenant.

**URL** This is the URL to give to users to log in to the tenant through the Analytic Server console, and to use to configure SPSS Modeler server. See *IBM SPSS Analytic Server Installation and Configuration Guide* for details on configuring SPSS Modeler.

**Status** **Active** tenants are currently in use. Making a tenant **Inactive** prevents users from logging in to that tenant, but does not delete any of the underlying information.

#### Principals

Principals are users and groups that are drawn from the security provider that is set up during installation. You can add principals to a tenant as Administrators or Users.

- Typing in the text box filters on users and groups with the search string in their name. Select **Administrator** or **User** from the drop-down list to assign their role within the tenant. Click **Add participant** to add them to the list of authors.
- To remove a participant, select a user or group in the member list and click **Remove participant**.

### Metrics

Allows you to configure resource limits for a tenant. Reports the disk space currently used by the tenant.

- You can set a maximum disk space quota for the tenant; when this limit is reached, no more data can be written to disk on this tenant until enough disk space is cleared to bring the tenant disk space usage below the quota.
- You can set a disk space warning level for the tenant; when the quota is exceeded, no analytic jobs can be submitted by principals on this tenant until enough disk space is cleared to bring the tenant disk space usage below the quota.
- You can set a maximum number of parallel jobs that can be run at a single time on this tenant; when the quota is exceeded, no analytic jobs can be submitted by principals on this tenant until a currently running job completes.
- You can set the maximum number of fields a data source can have. The limit is checked whenever a data source is created or updated.
- You can set the maximum file size in megabytes. The limit is checked when a file is uploaded.

### Security provider configuration

Allows you to specify the user authentication provider. **Default** uses the default tenant's provider, which was set up during installation and configuration. **LDAP** allows you to authenticate users with an external LDAP server such as Active Directory or OpenLDAP. Specify the settings for the provider and optionally specify filter settings to control the users and groups available in the Principals section.

---

## Naming rules

For anything that can be given a unique name in Analytic Server, such as data sources and projects, the following rules are applied to those names.

- Within a single tenant, names must be unique within objects of the same type. For example, two data sources cannot both be named `insuranceClaims`, but a data source and a project could each be named `insuranceClaims`.
- Names are case-sensitive. For example, `insuranceClaims` and `InsuranceClaims` are considered unique names.
- Names ignore leading and trailing white space.
- The following characters are invalid in names.  
`~ , # , % , & , * , { , } , \ , : , < , > , ? , / , | , " , \t , \r , \n`



---

## Chapter 3. Getting users started

Tell users to navigate to `http://<host>:<port>/<context-root>/admin/<tenant>` and enter their username and password to log on to the Analytic Server console.

**<host>**

The address of the Analytic Server host.

**<port>**

The port that Analytic Server is listening on. By default this is 8080.

**<context-root>**

The context root of the Analytic Server. By default this is `analyticserver`.

**<tenant>**

In a multi-tenant environment, the tenant you belong to. In a single-tenant environment, the default tenant is **ibm**.

For example, if the host machine has IP address 9.86.44.232, you have created a "mycompany" tenant and added users to it, and the other settings have been left to their defaults, then users should navigate to `http://9.30.86.232:8080/analyticserver/admin/mycompany` to access the Analytic Server console.



---

## Chapter 4. Analytic Server job names

Analytic Server produces map-reduce jobs, which can be monitored through your Hadoop cluster's jobtracker.

The map-reduce job name has the following structure.

AS/{tenant name}/{user name}/{algorithm name}

**{tenant name}**

This is the name of the tenant under which the job is run.

**{user name}**

This is the user who requested the job.

**{algorithm name}**

This is the primary algorithm in the job. Note that a single stream may generate multiple map-reduce jobs; likewise, several operations within a stream can be contained within a single map-reduce job.



---

## Chapter 5. Troubleshooting

Analytic Server provides several helpful tools for problem determination.

---

### Logging

Analytic Server creates customer log files and trace files that are helpful for diagnosing problems. With the default Liberty installation, you can find the log files in the `{AS_ROOT}/ae_wlpserver/usr/servers/aeserver/logs` directory.

The default logging configuration produces two log files that roll over on a daily basis.

**as.log** This file contains the high-level summary of informational warning and error messages. Check this file first when server errors occur that cannot be resolved by using the error message that is displayed in the User interface.

**as\_trace.log**

This file contains all the entries from `ae.log`, but adds more information that is primarily targeted to IBM support and development for debugging purposes.

Analytic Server uses Apache LOG4J as its underlying logging facility. Using LOG4J, the logging can be dynamically adjusted by editing the `{AS_SERVER_ROOT}/configuration/log4j.xml` configuration file. You may be asked to do this by Support to help diagnose problems, or you may want to modify this to limit the number of log files kept around. Changes to the file are detected automatically within a few seconds so the Analytic Server does not need to be restarted.

For more information about log4j and the configuration file, see documentation at the official Apache website at <http://logging.apache.org/log4j/>.

---

### Version information

You can determine what version of Analytic Server is installed by checking the `{AS_ROOT}/properties/version` folder. The following files contain version information.

**IBM\_SPSS\_Analytic\_Server-\*.swtag**

Contains detailed product information.

**version.txt**

Version and build number for the installed product.

---

### Log collector

When problems cannot be resolved by directly reviewing the log files, you can bundle all the logs and send them to IBM support. There is a utility that is provided to make collecting all the necessary data simpler.

Using a command shell, run the following commands:

```
cd {AS_ROOT}/bin
run >sh ./logcollector.sh
```

These commands create a compressed file under `{AS_ROOT}/bin`. The compressed file contains all the log files and product version information.

---

## Common issues

This section describes some common administration issues and how you can fix them.

### Kerberos authentication fails when trying to access HCatalog data source

If you see errors in the log like the following:

```
cause:javax.security.sasl.SaslException: Failure to initialize security context
```

```
com.spss.analyticframework.api.exceptions.ComponentException: Cannot access HCatalog
```

You must ensure that the HDFS user's Kerberos TGT is cached and available on the Analytic Server server host. To do this:

1. Stop the Analytic Server process.
2. Run `kinit -f $hdfs.user` from the Analytic Server host, where **\$hdfs.user** is as defined in the `config.properties` file and has write permission to analytic root
3. Start Analytic Server.

---

## 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.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing  
Legal and Intellectual Property Law  
IBM Japan Ltd.  
1623-14, Shimotsuruma, Yamato-shi  
Kanagawa 242-8502 Japan

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 want 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 Software Group  
ATTN: Licensing  
200 W. Madison St.  
Chicago, IL; 60606  
U.S.A.

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.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

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

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs.



© Copyright IBM Corp. \_enter the year or years\_. All rights reserved.

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

---

## Trademarks

IBM, the IBM logo, and [ibm.com](http://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](http://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.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

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

ITIL is a registered trademark, and a registered community trademark of The Minister for the Cabinet Office, and is registered in the U.S. Patent and Trademark Office.

UNIX is a registered trademark of The Open Group in the United States and other countries.

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

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.







Printed in USA