

**IBM® Workload Scheduler
AI Data Advisor User's Guide
Version 10.2.4**

Note

Before using this information and the product it supports, read the information in [Notices on page lxi](#).

This edition applies to version 10, release 2, modification level 4 of IBM® Workload Scheduler (program number 5698-T09) and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

About this publication.....	v
Accessibility	v
Technical training.....	v
Support information.....	v
Chapter 1. Overview.....	6
Business scenario.....	7
Detecting anomalies in the workload.....	7
Basic concepts.....	9
AIDA Architecture.....	11
Getting started.....	13
Accessing AIDA.....	13
Chapter 2. Administrative tasks.....	15
Adding engines to AIDA.....	15
Configuring email alert settings.....	17
Configuring security.....	18
Managing special days.....	19
Chapter 3. Working with KPIs.....	23
KPIs for IBM Workload Scheduler.....	23
Managing KPIs in AIDA.....	26
Analyzing KPIs data.....	28
Chapter 4. Managing alerts in AIDA.....	36
Receiving alert notifications.....	36
Overview dashboard.....	36
Alert definitions.....	40
Alert details.....	45
Analyzing an alert instance.....	49
Chapter 5. Troubleshooting AIDA	59
Logging and tracing in AIDA	59
Troubleshooting AIDA	60
Notices.....	lxi
Index.....	65

About this publication

This guide provides information about how to use AI Data Advisor (AIDA).

Accessibility

Accessibility features help users with a physical disability, such as restricted mobility or limited vision, to use software products successfully.

With this product, you can use assistive technologies to hear and navigate the interface. You can also use the keyboard instead of the mouse to operate all features of the graphical user interface.

For full information, see the Accessibility Appendix in the *IBM Workload Scheduler User's Guide and Reference*.

Technical training

Cloud & Smarter Infrastructure provides technical training.

For Cloud & Smarter Infrastructure technical training information, see: <http://www.ibm.com/software/tivoli/education>

Support information

IBM provides several ways for you to obtain support when you encounter a problem.

If you have a problem with your IBM software, you want to resolve it quickly. IBM provides the following ways for you to obtain the support you need:

- Searching knowledge bases: You can search across a large collection of known problems and workarounds, Technotes, and other information.
- Obtaining fixes: You can locate the latest fixes that are already available for your product.
- Contacting IBM Software Support: If you still cannot solve your problem, and you need to work with someone from IBM, you can use a variety of ways to contact IBM Software Support.

For more information about these three ways of resolving problems, see the appendix about support information in *IBM Workload Scheduler: Troubleshooting Guide*.

Chapter 1. Overview

Learn about AIDA built-in intelligence and how it can help you detect anomalies in your workload, prevent problems and reach operational excellence.

Modern businesses need to process data faster and more efficiently to make informed, data-driven decisions. To reach this objective, items or events that do not conform to an expected pattern must be detected immediately and prompt responses must be provided.

Manual methods are not effective to handle and analyze complex data since they leave room for human errors, false positives, or missed anomalies. Therefore, a new proactive approach to detect anomalous behaviors and predict issues is needed.

With Artificial Intelligence and Machine Learning techniques, automated anomaly detection is becoming a reality. By analyzing and predicting time series, **AI powered anomaly detection** can be a key to anticipate and prevent issues, saving energy that companies can utilize to grow their business.

Anomaly Detection and Problem Prevention in IBM® Workload Scheduler

A new component is available in IBM® Workload Scheduler – **AI Data Advisor (AIDA)**– based on AI and ML techniques. AIDA enables fast and simplified data-driven decision making, for an intelligent workload management. By analyzing historical data and metrics gathered by IBM® Workload Scheduler and IBM Z Workload Scheduler, and predicting their future patterns, AIDA identifies anomalies in KPI trends (such as number of completed jobs in the current plan, job duration, and job end-time) and sends alerts immediately to anticipate and prevent problems and delays. Alerts show up on the Workload Dashboard and can be notified via email.

Also, when an alert is issued, an event rule can be defined in IBM® Workload Scheduler to open a ticket on the supported service platform.

AIDA frees up Product Administrators, Lines of Business Administrators, and Operators from the burden of managing workload issues, so they can better focus on workload management and optimization. It provides a proactive approach to minimizing operational risk since alerts are not triggered by issues but are sent to prevent issues.

AIDA provides a dedicated User Interface from where you can:

- Obtain an interval estimation of a KPI trend.
- Analyze a KPI trend over time.
- Identify and analyze anomalies in a KPI trend.

AIDA Benefits

- Provides AI-powered automation, ensuring workload runs as expected, smoothly and without delay
- Provides a proactive approach to minimizing operational risk since alerts are sent before problems or delays occur
- Enables fast and simplified data-driven decision making

- Improves root cause analysis
- Provides new monitoring capabilities in cloud native architecture
- Improves stability through risk assessment
- Enables proactive SLA (Service Level Agreement) management
- Increases IBM® Workload Scheduler reliability of both infrastructure and workload

A business scenario

This business scenario shows how an IBM Workload Scheduler operator can benefit from AIDA.

Detecting anomalies in the workload

By predicting KPIs time series, AIDA can identify anomalies in the number of completed jobs in the current plan.

Jason works as an **IBM Workload Scheduler operator** for a large bank. In his job, he needs to quickly find issues with the workload, understand the impacts, and alert the right people.

While some critical batch processing is running, an alert shows up on the Workload Dashboard: **the number of jobs completed in the *Loan and Mortgage* folder is lower than expected**. From the Dashboard, Jason can open AIDA UI and analyze the reported anomaly.

On the Anomaly Data Analysis page, he can see a graph with the number of jobs completed in the folder, compared to the expected range of values. The anomaly is highlighted in the graph where a light-blue area represents the expected range of values, statistically defined based on IBM Workload Scheduler historical data. Jason can also compare the anomalous trend with the trend on a similar day. He can also add more days for comparison.

The screenshot displays the AIDA interface for an alert instance. At the top, the header shows 'AI DATA ADVISOR (AIDA)', 'Time Zone Europe/Rome (GMT + 02:00)', and 'Usernamehere Logout'. The main content area is titled 'Alert Instance Details' and includes the following information:

- Instance ID:** #2020_21_04_1235413316
- Detection Time:** 2020/11/09 - 12:08 PM CEST
- Severity:** High
- Resolution Status:** Resolved
- Special Days:** Christmas
- Related Alert:** Anomaly in the progress of jobs
- Alert Trigger:** >= 10 consecutive anomalies within 1h
- Engine:** 10.14.37.90.311

Below the details, the 'Anomaly Source KPI' section is active, showing two line graphs for 'PROGRESS OF JOBS FOR LOAN_AND_MORTGAGE Folder'. Both graphs show 'Observed Data' (blue line), 'Estimated Data Range' (light blue shaded area), and 'Anomaly' (red circle). The x-axis represents time from 08:00 AM to 01:00 PM on 01/11/2021, and the y-axis represents the number of jobs (0 to 250). The graphs show a steady upward trend in the number of jobs, with an anomaly point marked at approximately 11:00 AM.

An anomaly can have multiple causes with different severities: to find root causes faster, Jason clicks the Add KPI button to add correlated Key Performance Indicators to the anomaly data analysis. From the Add KPI panel, he selects the workload trend by workstation.

Workstation SAP1 shows a certain percentage of anomaly with the number of waiting jobs. Jason selects this KPI. A new graph is added to AIDA UI showing the today trend of jobs waiting on SAP1. Jason adds a comparison graph to analyze the KPI trend on different days. For the selected days, the KPI shows a regular trend, therefore he understands that to manage the unusual number of waiting jobs on the SAP workstation, he needs to contact the SAP administrator to free up some resources.

Jason opens a high severity ticket to have the issue quickly resolved, then marks the alert as resolved.

By using AIDA, and leveraging the correlated KPIs analysis, Jason was able to easily analyze the detected anomaly, quickly identifying and addressing the root causes of problem, without compromising the SLA.

AIDA is much more than anomaly detection and analysis.

A special page can be opened directly from the Monitor Job UI of the Dynamic Workload Console, showing the trend of all the KPIs available for a job.

AIDA can also be used by **IBM Workload Scheduler administrators** to:

- Pause and activate an alert generation.
- Set Special Days for each KPI in a dedicated UI, to include them in the prediction model with a higher tolerance.
- Fine-tune the KPI prediction.

Basic concepts

A few basic concepts are necessary when you use AIDA.

KPIs (Key Performance Indicators)

KPIs for IBM Workload Scheduler processes that are constantly monitored by AIDA. For example, the number of completed jobs in the current plan.

For more information about IBM Workload Scheduler KPIs managed by AIDA, see: [KPIs for IBM Workload Scheduler on page 23](#).

Anomaly Source KPI

The KPI whose anomalous trend has triggered an alert.

For more information about how to analyze an anomaly source KPI, see: [Analyzing an alert instance on page 49](#).

Correlated KPI

KPI correlated with the anomalous KPI. You can add one or more correlated KPIs to the anomaly data analysis.

For more information about how to add correlated KPIs to the anomaly data analysis, see: [Analyzing an alert instance on page 49](#).

Data Point

Each singular observation of a KPI.

Anomaly

Unexpected KPI data point.

AIDA detects an anomaly when a KPI falls outside the expected range of values which is statistically defined based on KPI historical data.

For more information about anomaly data analysis, see [Analyzing an alert instance on page 49](#).

Alert

Sequence of anomalies of a KPI.

An alert is defined by a set of parameters and conditions (see **Alert trigger**). When the conditions are met, the alert triggers, and an alert instance is created.

For example: 10 consecutive KPI data points that fall outside the expected range of values within 1 hour.

For more information, see: [Alert definitions on page 40](#).

Alert Instance

A single occurrence of an alert, given its definition. As AIDA continuously monitors KPIs, when an alert is triggered, a record is created into OpenSearch database with the alert instance information.

For more information about alert instances, see: [Overview dashboard on page 36](#).

Alert Severity

For each detected anomaly, AIDA calculates its percent deviation from the interval estimation. When an alert is generated, given its definition, the alert severity is calculated as average of percent deviations of the anomalies that concur to the alert generation. Alert severity classification by severity is:

- High, when the average of percent deviations is > 30
- Medium, when the average of percent deviations falls in the interval 20-30
- Low, when the average of percent deviations is < 20

AIDA displays the highest severity of all the alerts in an alert instance.

Anomaly Bounds

The upper and lower bounds of the expected range of values for a KPI.

Alert Trigger

Set of conditions that define an alert. For example: 10 consecutive KPI data points that fall outside the expected range of values within 1 hour.

When triggering conditions are satisfied, a new alert instance is created inside OpenSearch database.

There are two types of alert triggers available for selection:

- Continuous: Triggers when anomalous data points fall above or below the predicted range.
- Total: Triggers for anomalous data points that are either above or below the predicted range, as well as those that exceed both thresholds.

Alerts are notified on the Workload Dashboard or via email.

For more information about receiving alert notifications, see: [Receiving alert notifications on page 36](#).

Anomaly %

The percentage of observed KPI data points that fall outside the expected range of values in the reference time interval:

- < 6 : Low
- 6-10: Medium
- >10: High

A KPI trend can show some anomalies, however an alert might not be issued if the trigger condition is not met.

Anomaly Data Analysis

Anomaly Data Analysis is part of AIDA User Interface. When anomalies in a KPI trend generate an alert, you can compare the anomalous trend with trends in one or more different time intervals. You can also add correlated KPIs to the anomaly data analysis to find root causes faster.

For more information about anomaly data analysis, see [Analyzing an alert instance on page 49](#).

Alert Details

Alert Details provides detailed information about an alert, it's status, the current opened instances and its history.

For more information, see: [Alert details on page 45](#).

Alert History

Calendar graph showing previous alert instances and related severity.

For more information, see: [Alert details on page 45](#).

Timerange

How often a KPI is checked to detect anomalies (for example: every day, or every 10 minutes). It is set through the PROPHET_ORCHESTRATOR **schedule_alert** parameter of the common.env configuration file (or in the value.yaml file for Kubernetes deployments).

Special Days

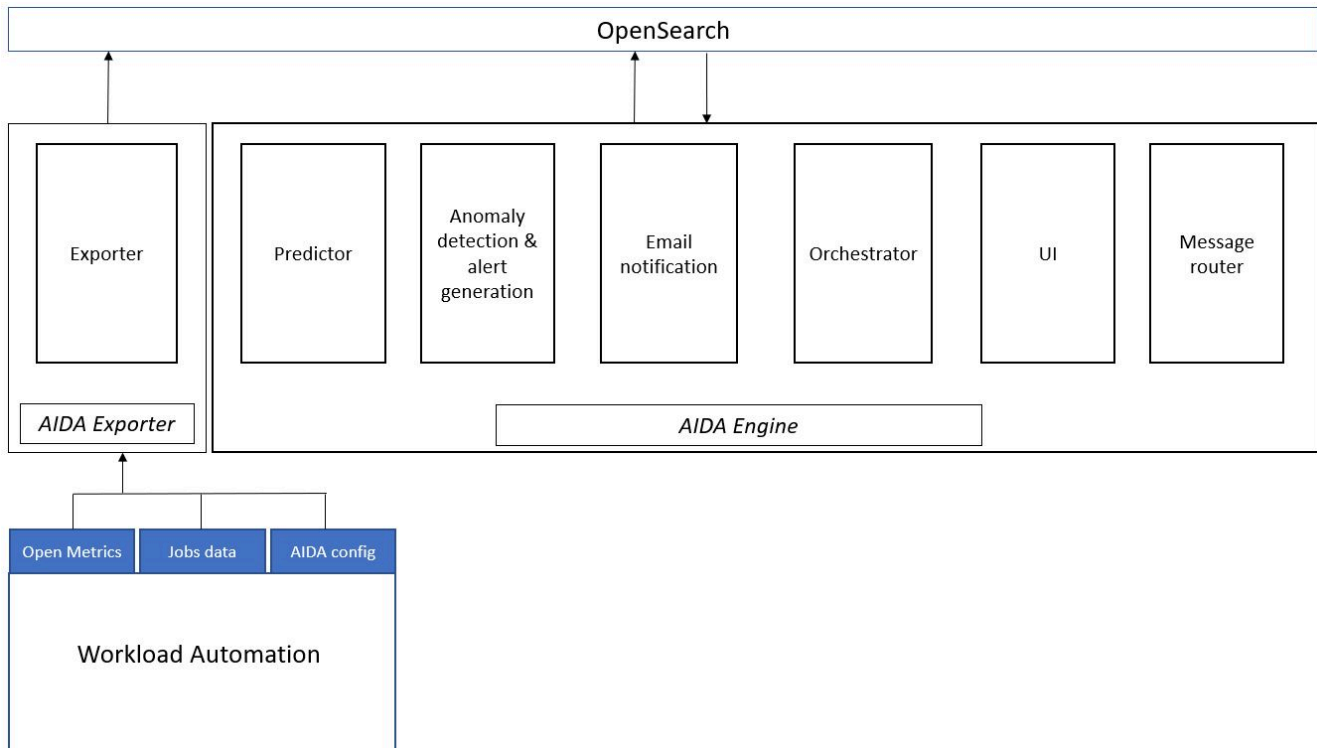
Special days are days on which a KPI trend is affected by seasonality factors such as national holidays, vacation, business cycles, recurring events. To avoid false positive alerts, the special days are included in AIDA prediction model with a higher tolerance level than the standard days.

For more information, see: [Managing special days on page 19](#).

AIDA Architecture

AIDA is built adopting a microservices-based architecture.

Figure 1. AIDA architecture



AIDA is composed of two major components: AIDA Exporter and AIDA Engine. Each component contains a number of services:

AIDA Exporter

Exporter

Through IBM Workload Scheduler APIs, exports KPIs metrics from IBM Workload Scheduler (according to OpenMetrics standard) and stores them into AIDA OpenSearch database.

Also, it exports Alert definitions from IBM Workload Scheduler and imports them into OpenSearch.

AIDA Engine

Predictor

Calculates the expected values of each KPI, also considering special days.

Anomaly detection and alert generation

Detects anomalies in KPIs trend by comparing observed KPI data points with expected values, and generates alerts when trigger conditions are met.

Email notification

Sends email notification when alerts are generated.

Orchestrator

Orchestrates KPI prediction and anomaly detection.

UI

AIDA User Interface.

Internal event manager

Manages communication among AIDA services.

Also, AIDA uses:

OpenSearch (an Elasticsearch technology)

To store and analyze data.

Keycloak

To manage security and user access in AIDA (Docker deployment only). If not deployed, the Dynamic Workload Console user authentication roles will be used.

Nginx

As a reverse proxy for its components.

Getting started

Information about AI Data Advisor (AIDA) installation and configuration.

For information about how to deploy AI Data Advisor (AIDA) to monitor IBM Workload Scheduler and IBM Z Workload Scheduler engines, see *Deploying AI Data Advisor* in the *Planning and Installation Guide*.

Accessing AIDA

You can access AIDA user interface from different entry points.

In AIDA, each user is granted access to the same scheduling objects (such as jobs, job streams, and so on..) as they are in the Dynamic Workload Console. As a result, in AIDA, each user can view only KPIs and alerts related to these scheduling objects. Dynamic Workload Console administrators are also AIDA administrators and can work with all KPIs, manage special days, customize prediction tuning parameters, and pause alerts.

Dynamic Workload Console users can access AIDA from any of the following entry points:

From the Workload Dashboard of the Dynamic Workload Console

When anomalies in a KPI trend generate an alert, the alert is notified by AIDA through the Anomaly Widget on the Workload Dashboard.

For instructions about how to customize the Workload Dashboard with AIDA Anomaly Widget, see the topic *Creating a customized dashboard for monitoring* in the Dynamic Workload Console User's Guide.

The Anomaly Widget on the Workload Dashboard indicates the number of Anomaly Alerts that have been generated in the last 24 hours. To analyze the alerts, run the following steps:

1. Click on the Anomaly Widget. A panel opens containing the list of the latest Anomaly Alerts. For each alert, the following information is displayed:
 - Alert severity
 - Alert description
 - A link to AIDA UI where you can find detailed alert information to quickly identify the root cause of the issue.
2. Follow the link for the alert that you want to analyze. For details, see [Analyzing an alert instance on page 49](#).
3. Click the **View all alerts** button to view the full list of alerts.



Note:

When accessing AIDA UI from the Dynamic Workload Console, the connection authentication is based on the public key of the WebSphere Application Server Liberty SSL certificates, which is the default public key set during AIDA deployment. If you use custom certificates for the Dynamic Workload Console, you must change AIDA default public key accordingly. Properly customize the parameter:

- `DWC_PUBLIC_KEY` in the `common.env` file, during AIDA deployment with Docker.
- `aida-nginx.waConsoleCertSecretName` in the `values.yaml` file, during AIDA deployment with Kubernetes.

If AIDA is already installed, after changing the parameter run the command:

```
docker -compose up -d --build
```

Users defined with Keycloak are typically AIDA administrators and can access AIDA from the following entry points:

From a direct login

AIDA administrators can directly login to AIDA user interface by using a dedicated userid and password. For details about defining users with Keycloak, see [Configuring security on page 18](#).

AIDA user interface can be accessed at the link

```
https://aida-ip:aida-port/
```

where `aida-host` and `aida-port` are the values specified during AIDA deployment.

By following the link provided in the alert email notification

By following the link provided in the notification email, administrators can access the alert instance page in AIDA UI and run an anomaly analysis. For details, see [Analyzing an alert instance on page 49](#).

Chapter 2. Administrative tasks

Administrative tasks are required to work with AIDA

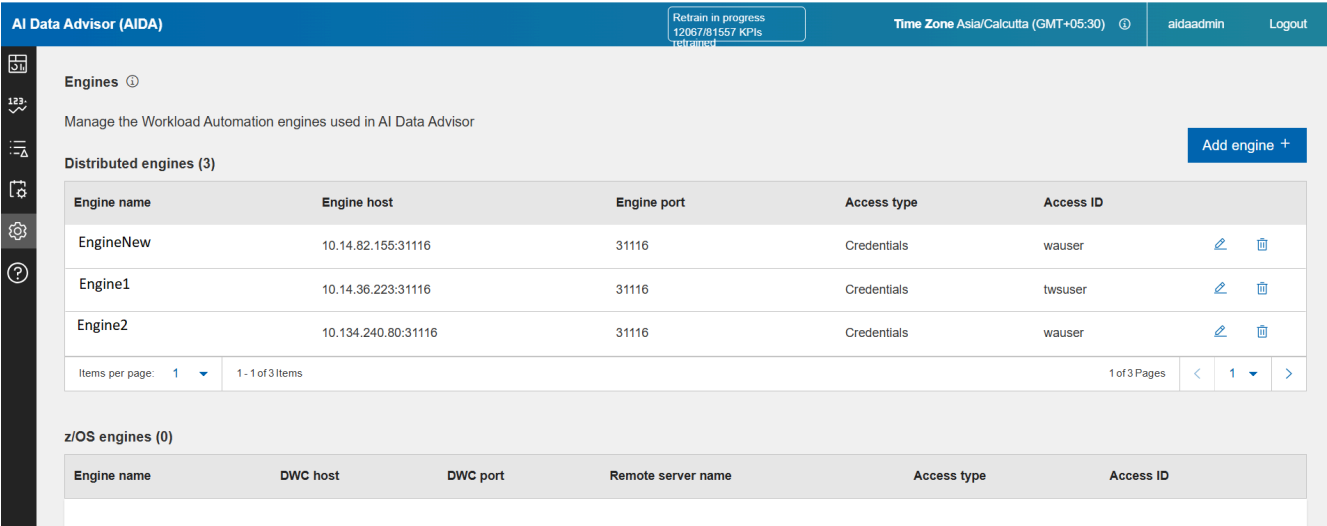
Adding IBM Workload Scheduler engines to AIDA

How to add IBM® Workload Scheduler engines to AIDA for monitoring purposes.

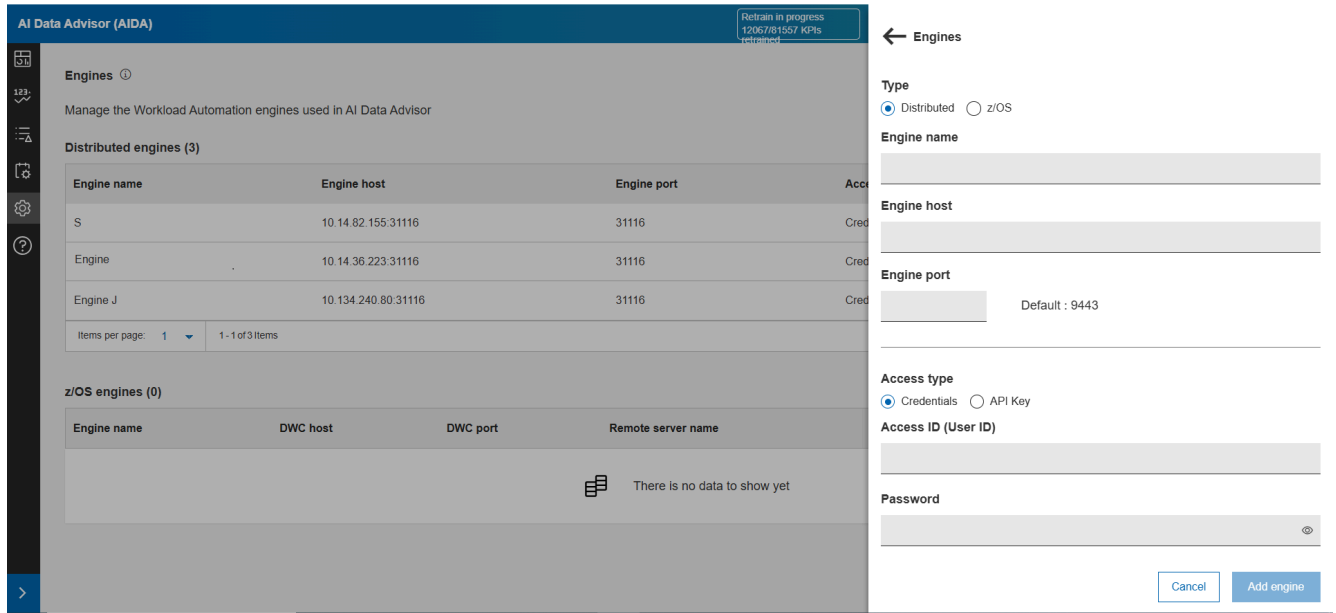
About this task

From AIDA left-hand sidebar, select **Engines** to open the Engines page.

As an administrator you can add and manage engines used in the Workload Automation environments.



This page shows you the Engines managed by AIDA.



You can add engines for distributed and z/OS environment. To add the engines do the following:

1. Click on the **Add Engine** button. A panel opens containing the list of the fields.

Property	Description
Engine name	Enter the name of the engine. This is a unique identifier for the engine being added to the system.
Engine host (only for Distributed)	Enter the host name or the or IP address of the server where the engine is hosted.
Engine port (only for Distributed)	Enter the port number on which the application server is running. The default port number is 9443.
DWC host (only for z/OS)	Enter the DWC host name or the or IP address of the server where the engine is hosted.
DWC port (only for z/OS)	Enter the DWC host name or the or IP address of the server where the engine is hosted.
Remote server name (only for z/OS)	Enter the name or IP address of the remote server.
Access type	Supports the following authentication methods to interact with the system: <ul style="list-style-type: none"> ◦ Credentials: Requires a username and password for user verification. ◦ API Keys: A unique key assigned to each user or application to authenticate API requests.
Access ID (User ID)(only for Credentials)	Enter the user identifier or the username used to access the system.

Property	Description
Password (only for Credentials)	Enter the password used to access the system.
Access ID (API key name) (only for API Key)	Enter the unique identifier assigned to the user for API key management and system access.
API key (only for API Key)	Enter the unique API key assigned to the user for authentication and system access.

2. Enter the required details, and then click **Add engine**.

The added engines appears in the Engines page.

3. Click the **Edit** button to modify an engine.
4. Click the **Delete** icon to delete an engine.

Configuring email alert settings

How to configure email alert settings in AIDA.

About this task

When anomalies in a KPI trend generate an alert, AIDA can notify it via email.

To receive alert email, you must set up an SMTP server and, during the aida-email container deployment with Docker or Kubernetes, you must configure the following settings in the configuration file (.env or .yaml):

SMTP_SERVER

Fully qualified hostname of the SMTP Server that will be used by AIDA to send alert email [Example: smtp.gmail.com]

SMTP_PORT

The port of the SMTP mail server [Must be a TLS Port. Example for Gmail: 587]

SENDER_MAILID

The SMTP email account [Example: smtp@server.com]

SENDER_MAILPWD

The password associated with the SMTP email account

RECIPIENT_MAILIDS

Comma separated list of recipient emails [Example: test@mail.com,test2@mail.com]

HOST_IP

AIDA host IP address and port [Example: 10.10.10.10:1111]

For example, to configure an SMTP server for Google Gmail, run the following steps:

1. Sign in to your Gmail account.
2. In the top right corner of Gmail window, click **Settings**, and then **All settings**.
3. Select the **Forwarding and POP/IMAP** tab and click the **Enable IMAP** radio button.
4. From **Manage your Google Account**, select **Security**.
5. Turn on **Less secure app access** .

Configuring security

To manage access to AIDA, use **Keycloak**.

Before you begin

During AIDA deployment with Docker, deploy a Keycloak container. In this case, you can manage access to AIDA through Keycloak.

In Keycloak, each application has its own Realm with different users and authorization settings. AIDA authorization settings are stored in a Realm named AIDA.

With Keycloak deployment, the following users are automatically generated :

1. **userid:** aidaadmin, **password:** admin, **role:** aida-admin.

With this role, a user can directly login to AIDA UI, from where, besides analyzing anomalies and alerts, he can work with all KPIs, manage special days, customize prediction tuning parameters, and pause alerts. This user is typically an AIDA Administrator.

2. **userid:** admin, **password:** admin, **role:** keycloak-admin.

With this role, a user can access Keycloak admin console to define additional users, or change default passwords.

Defining users from Keycloak admin console

About this task

Use Keycloak admin console to define new users, new roles, or change user passwords.

For example, to create a new AIDA user with administrator role, run the following steps:

1. Access Keycloak admin console **https://<IP:PORT>/keycloak/auth/admin** by using the following credentials:
 - **userid=admin**
 - **password=password**
2. If you want, you can change Keycloak default password:
 - a. From Keycloak admin console, in the upper right corner, click **Admin**.
 - b. Select **Manage account -> password**
3. Under **Clients -> nginx -> roles tab**, click the **Add role** button.
4. Provide the role name **admin** and click **save**.
5. Under **users**, click the **add user** button.
6. Provide a user name and click **save**.

7. Under **Credentials**, provide a password for the user, turn the **temporary** field to **off**, click the **Reset Password** button and confirm.
8. Under **Role Mappings**, in the **Client Roles** dropdown, select **nginx**. Some boxes appear on the right.
9. Under **Available Roles**, select **admin** and click the **Add Selected** button. The **admin role** appears in the **Assigned Roles** box.
10. On the left navigation bar, select the **Realm Settings** page and go to the **Themes** tab.
11. In the **Login Theme** parameter, select the **Keycloak** theme, then click save.

For details about Keycloak, see [Keycloak documentation](#).

Managing special days

From the Special Days UI, AIDA administrator can define special days in the prediction model.

Before you begin

Special days are days on which a KPI trend is affected by seasonality factors such as national holidays, vacation, business cycles, recurring events. To avoid false positive alerts, the special days are included in AIDA prediction model with a higher tolerance level than the standard days.

Adding special days

About this task

Use the Special Days UI to define special days in the KPIs trend. From AIDA left hand sidebar, select **Special Days**.

The screenshot shows the 'Special Days' interface in AIDA. At the top, there's a header with 'AI DATA ADVISOR (AIDA)', 'Time Zone Europe/Rome (GMT + 02:00)', 'Usernamehere', and 'Logout'. Below the header, the 'Special Days' title is followed by a sub-header: 'In AIDA prediction model, special days are included with a higher tolerance level than standard days to avoid false positive alerts.' A search bar with the placeholder 'Search an item...' and an 'Add Special Day +' button are present. The main content is a table with the following data:

Name	Date	Repeat	End Repeat	Description	Status
New Year	01-Jan-2022 Sat	every year	never	The first day of the new year	Active
USA Holidays	-	none	none	All Country Holidays	Active
My custom range	01-Jan-2022 Sat	none	none	Custom description	Draft

At the bottom of the table, there's a pagination control: 'Items per page: 10 | 1 - 10 of 40 items' and 'Page 2 of 4' with navigation arrows.

The Special Days UI landing page contains a table with all the special days defined in AIDA. For each special day, the following information is displayed:

Name

The name of the special day (or interval between two dates).

Date

A single day (or interval).

Repeat

For recurring special days, define how often you want the special day (or interval) to repeat. It can be:

- None
- Every day
- Every week
- Every month

- Every year
- Custom

End Repeat

Define when you want to stop repeating

Description

A description of the special day (or interval).

Status

The status of the special day (or interval). Can be Active or Draft.

The table shows the row action icons available for each special day:

Edit

Click this action icon to edit the special day.

Delete

Click this action icon to delete the special day. By deleting a special day, the day will be considered as a standard day.

To define a new special day in AIDA, click the **Add Special Day** button.

In the **Type** drop-down menu you can select one of the following types of special days.

- Country specific holidays
- Custom date

AIDA prediction model is automatically retrained every 24 hours.

Adding country specific holidays

About this task

You can add your country specific holidays. Run the following steps:

1. On the **Add Special Day** panel, in the **Type** drop-down menu, select **Country specific holidays**.
2. In the **Country** drop-down menu, select your country. The list of all your country specific holidays is displayed.
3. Click the **Add Special Day** button.

Adding a custom date

About this task

You can add a custom special day (or interval between two dates) like, for example, the monthly financial closing dates. Run the following steps:

1. On the **Add Special Day** panel, in the **Type** drop-down menu, select **Custom date**.
2. In the **Name** field, provide the name of the special day or interval.
3. Select a day in the **Start Date** calendar.
4. If you are defining a special interval, check the **Add End Date** check-box and select a day in the **End Date** calendar .
5. Select the frequency in the **Repeat** drop-down list. To customize the frequency, select **Custom**. In the **Frequency** drop-down menu, select one of the following options:

Daily

Specify **Every n days**

Weekly

Specify **Every n weeks**, and select the days of the week

Monthly

Specify **Every n months**, and select the days of the month.

Yearly

Specify **Every n years**, and select the months of the year.

6. Specify the **End Repeat** date (default value is *Never*).
7. Add a description in the **Description** field.
8. Set **Status** toggle to Active.
9. Click the **Add Special Day** button.

Chapter 3. Working with Key Performance Indicators (KPIs)

Learn how to manage and analyze KPIs in AIDA.

KPIs for IBM Workload Scheduler

Find out the IBM Workload Scheduler KPIs managed by AIDA.

IBM Workload Scheduler and IBM Z Workload Scheduler expose metrics and KPIs definitions according to the OpenMetrics standard.

KPIs definitions and data retrieval frequency are defined into a json file inside IBM Workload Scheduler. This file is retrieved by AIDA Exporter component once a day.

According to the frequency of data retrieval defined in the json file, AIDA's Exporter component retrieves the metrics through ad-hoc APIs and stores them into AIDA OpenSearch database.

KPIs definitions and KPIs metrics cannot be modified by AIDA users.

For details about IBM Workload Scheduler exposed metrics, see *Exposing metrics to monitor your workload* in the IBM Workload Scheduler *User's Guide and Reference*.

For details about IBM Z Workload Scheduler exposed metrics, see *Exposing metrics to monitor your workload* in the IBM Z Workload Scheduler *Managing the Workload*.

AIDA also collects a special KPI named **Job history**, containing the duration for each job that has been defined in IBM Workload Scheduler with the advanced analytics option enabled and for all its predecessor jobs. Every day, this KPI generates one data point for each job execution (KPI frequency = 86400 seconds) .

On a daily basis, starting from the KPIs time series, AIDA uses Machine Learning algorithms to predict the KPIs trends.

According to the **Timerange** parameter in the common.env configuration file (or values.yaml file for Kubernetes deployment), KPIs current values are compared with their predicted values. Alerts can be generated, based on alerts definition rules. For details, see [Alert Definitions on page 40](#).

IBM Workload Scheduler KPIs are grouped in the following categories:

Category	KPI name	HCL Workload Automation KPI metric string	Metric for instances division	Anomaly type	Object monitored	Data frequency
Jobs	Number of jobs in plan by folder (by status)	application_wa_JobsByFolder_jobs	Job status (10)	Higher/lower jobs number	Folder	1 data point every 4 minutes (240 seconds)

	Number of jobs in plan by workstation (by status)	application_wa_JobsByWorkstation_jobs	Job status (10)	Higher/lower jobs number	Workstation	1 data point every 4 minutes (240 seconds)
	Number of jobs in plan by status	application_wa_JobsInPlan_Count_job	Job status (10)	Higher/lower jobs number	All jobs in plan	1 data point every 4 minutes (240 seconds)
	Number of total jobs in plan	application_wa_JobsInPlan_Count_job_total	/	Higher/lower jobs number	All jobs in plan	1 data point every 4 minutes (240 seconds)
	Job history (start time & duration)	job_history	Start time Duration	Earlier/later start time Longer/shorter duration	Job	1 data point per each daily job executions (86400 seconds)
Queue	Available space for WA message files	application_wa_msgFileFill_percent	Queues (12)	Finishing space for queue	All queues	1 data point every 4 minutes (240 seconds)



Note:

- **Job status:** WAITING, READY, RUNNING, SUCCESSFUL, ERROR, CANCELED, HELD, UNDECIDED, BLOCKED, SUPPRESS.
- **Queues:** Appserverbox.msg, Courier.msg, mirrorbox.msg, Mailbox.msg, Monbox.msgn, Moncmd.msg, auditbox.msg, clbox.msg, planbox.msg, Intercom.msg, pobox messages, server.ms

IBM Workload Scheduler KPIs json file

In the KPIs json file inside IBM Workload Scheduler, each entry defines a KPI. The frequency parameter represents the frequency of the KPI data retrieval, expressed in seconds. This file cannot be modified by users.

```
[
{
```

```

    "name": "Job history KPI",
    "metric_name": "job_history",
    "frequency": 86400,
    "category": "Jobs",
    "subcategory": "history",
    "labels": [
      "job"
    ],
    "keyprop": "attributes",
    "keyPropValues": ["duration"],
    "type": "records"
  },
  {
    "name": "Total jobs in plan",
    "metric_name": "application_wa_JobsInPlanCount_job",
    "frequency": 240,
    "category": "Jobs",
    "subcategory": "Trend",
    "type": "total"
  },
  {
    "name": "Jobs in plan by status",
    "metric_name": "application_wa_JobsInPlanCount_job",
    "frequency": 240,
    "category": "Jobs",
    "subcategory": "Trend",
    "keyprop": "jobstatus"
  },
  {
    "name": "Jobs in plan by workstation",
    "metric_name": "application_wa_JobsByWorkstation_jobs",
    "frequency": 240,
    "category": "Jobs",
    "subcategory": "Trend_by_wks",
    "keyprop": "jobstatus",
    "labels": [
      "workstation"
    ]
  },
  {
    "name": "Jobs in plan by folder",
    "metric_name": "application_wa_JobsByFolder_jobs",
    "frequency": 240,
    "category": "Jobs",
    "subcategory": "Trend_by_folder",
    "keyprop": "jobstatus",
    "labels": [
      "folder"
    ]
  },
  {
    "name": "WA Message files fill percentile",
    "metric_name": "application_wa_msgFileFill_percent",
    "frequency": 240,
    "category": "Queue",
    "subcategory": "Msg file fill",
    "keyprop": "msgfile"
  }

```

```
}  
]
```

Managing KPIs in AIDA

With AIDA you can fine-tune KPIs prediction and analyze the KPIs trend over time.

About this task

From AIDA left-hand sidebar, select **KPIs** to open the KPIs page.

The screenshot shows the 'KPIs' management interface in AIDA. At the top, there's a header with 'AI DATA ADVISOR (AIDA)', 'Time Zone Europe/Rome (GMT + 02:00)', and 'Usernamehere Logout'. The main content area is titled 'KPIs' and includes a sub-header 'Select one or more KPIs for tuning, retraining, or data analysis.' Below this are three filter dropdowns: 'KPI name' (set to 'All'), 'Engine' (set to '10.10.10.10'), and 'Object name' (set to 'All'). To the right of these filters, it says 'Automatic retrain every 24h' and 'Scheduled retrain time: 13.45.00', along with a 'Global tuning' button. The main part of the page is a table with the following data:

<input type="checkbox"/>	KPI name	Engine	Monitored object	Last 24h Alerts Instances	
<input type="checkbox"/>	Job history (start time & duration)	10.10.10.10	Job ABCD123	250	
<input type="checkbox"/>	Total number jobs in plan	10.10.10.10	All Jobs in plan	67	
<input type="checkbox"/>	Number Jobs in plan by status	10.10.10.10	All Jobs in plan	30	
<input type="checkbox"/>	Number Jobs in plan by Workstation (by status)	10.10.10.10	Workstation WKST123	4	
<input type="checkbox"/>	Number Jobs in plan by status	10.10.10.10	All Jobs in plan	12	
<input type="checkbox"/>	Number jobs in plan by folder (by status)	10.10.10.10	Folder 12345	0	
<input type="checkbox"/>	Number jobs in plan by folder (by status)	10.10.10.10	Folder 6789	0	
<input type="checkbox"/>	Available space for WA message files	10.10.10.10	All queues	0	

At the bottom of the table, there's a pagination control: 'Items per page: 10 | 1 - 10 of 40 items' and 'Page 1 of 4' with navigation arrows.

This page shows you the KPIs managed by AIDA.

KPIs ⓘ

Select one or more KPIs for tuning, retraining, or data analysis.

KPI name: All | Engine: 10.10.10.10 | Object name: All | Automatic retrain every 24h | Scheduled retrain time: 13.45.00 | Global tuning

<input type="checkbox"/> KPI name	Engine	Monitored object	Last 24h Alerts	
<input type="checkbox"/> Job history (start time & duration)	10.10.10.10	Job ABCD123	250	
<input type="checkbox"/> Total number jobs in plan	10.10.10.10	All Jobs in plan	67	
<input type="checkbox"/> Number Jobs in plan by status	10.10.10.10	All Jobs in plan	30	
<input type="checkbox"/> Number Jobs in plan by Workstation (by status)	10.10.10.10	Workstation WKST123	4	
<input type="checkbox"/> Number Jobs in plan by status	10.10.10.10	All Jobs in plan	12	
<input type="checkbox"/> Number jobs in plan by folder (by status)	10.10.10.10	Folder 12345	0	
<input type="checkbox"/> Number jobs in plan by folder (by status)	10.10.10.10	Folder 6789	0	
<input type="checkbox"/> Available space for WA message files	10.10.10.10	All queues	0	

Items per page: 10 | 1 - 10 of 40 items | Page 1 of 4

KPIs prediction global tuning ✕

Sensibility for anomaly detection
The higher it is, the more anomalies can be identified. Default is "medium".

Low — Medium — High

Modifications will override the single KPI settings. Changes will be applied after next retrain.

From this page, depending on your permissions, you can run the following tasks:

- Tune KPIs prediction parameters, the sensibility for anomaly detection.



Note: Every 24 hours AIDA runs an automatic retrain of all KPIs. The prediction area is not visible while retraining is in progress.

- Access the KPIs Data Analysis page where you can:
 - Obtain an interval estimation of the KPIs trend.
 - Analyze the KPIs trend over time.
 - Identify and analyze anomalies in the KPIs trend.
- Delete the KPI.

The table displays the following information:

KPI Name

The name of the KPI.

Engine

The engine where KPI runs.

Monitored Object

The name of the object measured by the KPI.

Last 24h alert instances

The number of alert instances opened in the last 24 hours, allowing to identify any problematic KPIs.

Select one or more KPIs to run the following actions:

Tuning

You can tune a KPI prediction parameter, the anomaly detection sensitivity in the KPIs prediction tuning popover. This action allows you to increase or decrease sensibility, with higher sensibility identifying more anomalies.

You can also modify all KPIs using the **Global tuning** button, which overrides individual KPI settings. These changes will be applied after the next retraining.

Open

To open the **KPI Data Analysis** page where you can:

- Obtain an interval estimation of the KPI trend.
- Analyze the KPI trend over time.
- Identify anomalies in the KPI trend.

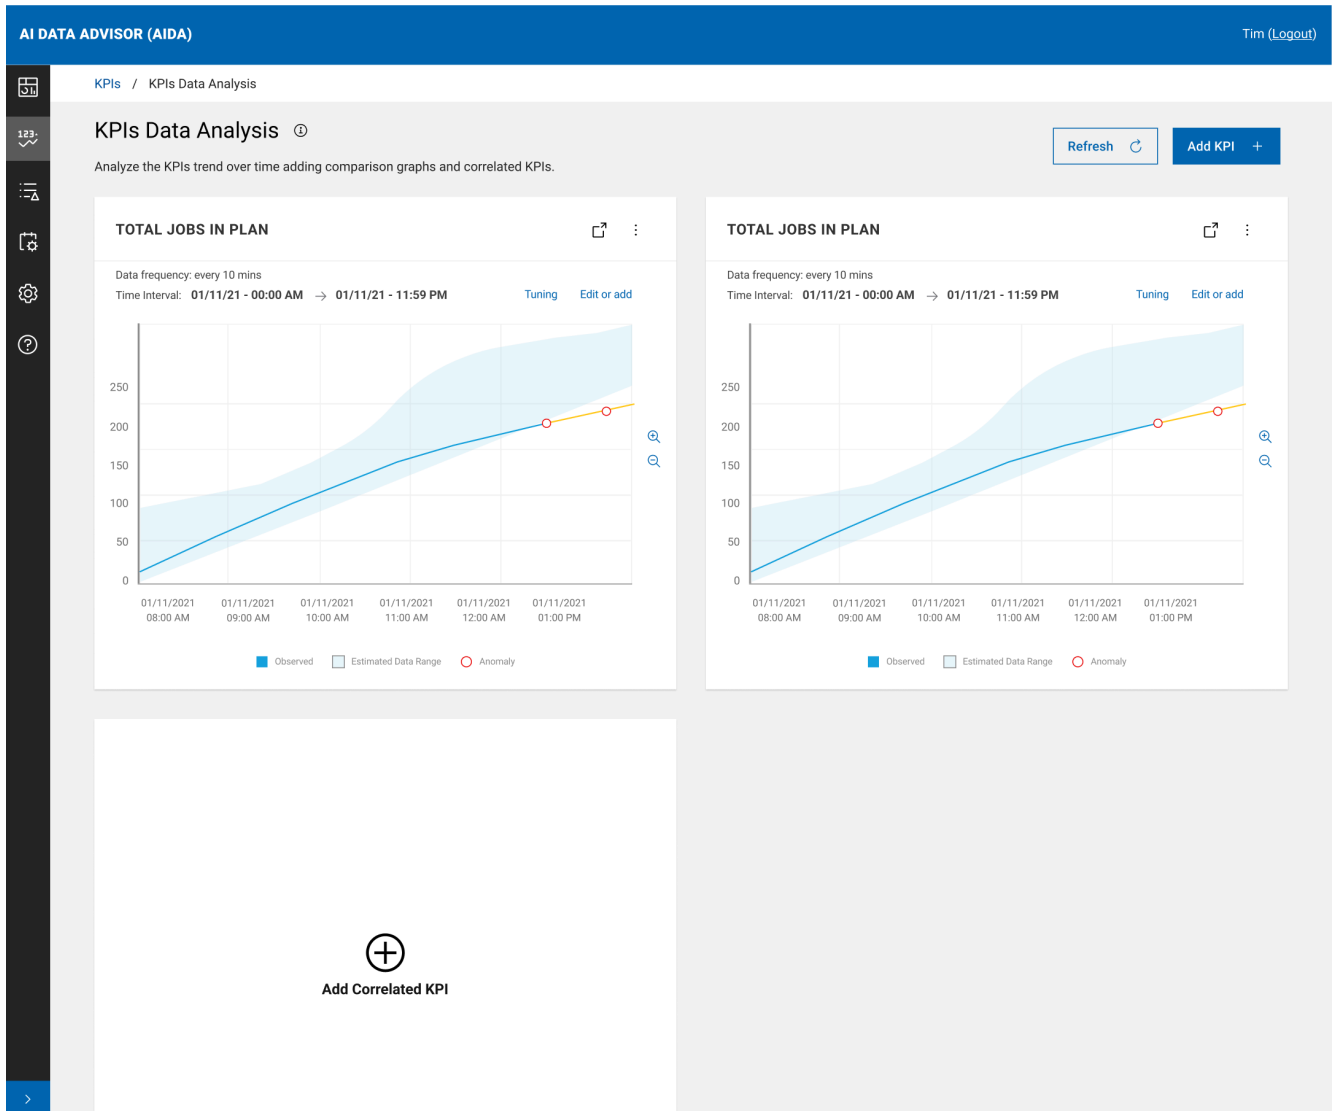
For details, see: [Analyzing KPIs data on page 28](#).

Analyzing KPIs data

See how to analyze and compare KPIs trend over time.

About this task

From AIDA left-hand sidebar, select **KPIs** to open the KPIs page. Select one or more KPIs and then select the **Open** action to open the **KPIs Data Analysis** page.



From this page you can:

- Obtain an interval estimation of the KPIs trend.
- Analyze the KPIs trend over time.
- Identify and analyze anomalies in the KPIs trend.

You can deepen your analysis by comparing the KPIs trend over different time intervals or adding correlated KPIs to your investigation.

In each KPI graph, data is displayed in time buckets. The KPI data frequency and the reference time interval are indicated in the graph header. To view all data points, click on the **Edit or add** link and reduce the time interval. A **light-blue area** represents the expected range of values for the KPI in the reference time interval, statistically defined by AIDA based on

historical data. The **blue line** represents the observed data that falls within the expected range of values. The **orange line** represents data that falls outside the expected range of values. You can zoom in or zoom out on the graph. On hovering over the KPI trend, data points appear. For each data point, a popover window displays the following information:

- Date and time of the observation.
- Current value: the KPI observed value.
- Estimated: the KPI interval estimation.

The anomalies in the KPI trend are represented by **red circles**. Among these, the anomalies that contribute to generating an alert are represented by **red dots**. On hovering over an anomaly, the following information is displayed:

- Date and time of the observation.
- Current value: the KPI observed value.
- Estimated: the KPI interval estimation.
- Deviation: the minimum distance (with - or + sign) of the KPI observed value from its interval estimation.

In each KPI graph, you can run a number of actions:

- Click on **Edit or add** to edit the graph time interval, or add time intervals to the graph for comparison purposes. For details, see the task **Setting time intervals with the Datpicker** below.
- The menu icon in the upper right corner of the graph contains the following additional actions:
 - **Compare graph**, to create a comparison graph with single or multiple time intervals for comparison purposes.
 - **Tuning**, to tune a KPI prediction parameter, the anomaly detection sensitivity in the KPIs prediction tuning popover. This action allows you to increase or decrease sensibility, with higher sensibility identifying more anomalies. You can also modify all KPIs using the **Global tuning** button, which overrides individual KPI settings. These changes will be applied after the next retraining. Tuning is available to AIDA administrators only.
 - **Refresh**, to refresh the graph after you run some tuning adjustments.
 - **Delete**, to delete the graph.
- For KPIs belonging to the **Jobs** category, an action icon is also present to open the workstation or job properties panel directly in IBM Workload Scheduler.

To deepen your KPIs data analysis, you can add additional graphs to this page:

- Comparison graphs with KPIs trend over different time intervals
- Graphs for additional KPIs.

Adding comparison graphs

About this task

You can edit the time interval, or add multiple time intervals to any KPI graph for comparison:

- In the graph, click **Edit or add** to open the Datepicker panel that allows you to:
 - edit the time interval
 - add multiple time intervals for comparison

To enhance the analysis, you can generate an additional graph.

- From the menu icon in the upper right corner of the graph, select **Compare graph**. The Datepicker panel opens where you can create a comparison graph with single or multiple time intervals.

For details about how to use the Datepicker widget, see [Setting time intervals with the Datepicker on page 32](#).



Note: In the graphs showing KPI trends in multiple time intervals, the gray area representing the expected KPI values in each time interval is not displayed.

Adding correlated KPIs

About this task

You can add one or more KPIs to your data analysis by clicking the **Add KPI** button in the upper right corner of the KPIs Data Analysis page.

On the left-hand side of the **Add KPI** panel, select a KPIs category.

For each KPI of the selected category, the following information is displayed:

KPI Name

Name of the KPI

Object Name

The name of the object measured by the KPI.

Tag

A search tag for the KPI. Usually, the tag is used to identify the engine to which the KPI refers.

Anomaly %

The percentage of observed KPI data points that fall outside the expected range of values in the reference time interval:

- < 6 : Low
- 6-10: Medium
- >10: High

To select additional KPIs, run the following steps:

1. Use the search bar to refine your search.
2. Select one or more KPIs.
3. Click the **Add KPI** button.

Results

A new graph for each selected KPI is added to the KPI Data Analysis page, representing the KPI trend in the reference time interval.

Setting time intervals with the Datepicker

Before you begin

Use the Datepicker to set single or multiple time intervals in a KPI graph.

In the Datepicker panel, select the type of interval:

Single Interval

- To edit a time interval in a KPI graph
- To add a KPI comparison graph with a single time interval

Multiple Intervals

- To edit multiple time intervals in a KPI graph
- To add a KPI comparison graph with multiple time intervals

Setting a single time interval

About this task

The **Single interval** section contains the following fields:

- **Start Date**
- **Start Time**
- **End Date**
- **End Time**

When you first open the Datepicker panel, these fields are set to the current time interval values in the KPI graph.

<
Edit or add intervals for comparison

Jobs in plan by workstation FOR **workstation: /WA-SERVER**

TIME INTERVAL: 18/1/2022, 00:00:00 → 19/1/2022, 00:00:00

Select the type of interval and set the corresponding information

Calendar indications: Anomaly % ● 0 - 5 ● 6 - 10 ● > 10 ■ Special day Selected day Current day

Single Interval
Displays observed KPI data, estimated data range and anomalies for a single time interval.

Start day	Start time	AM	→	End day	End time	AM
📅 01/18/2022	12:00	AM ▼		📅 01/19/2022	12:00	AM ▼

Multiple intervals
Displays observed KPI data for up to 5 different time intervals.

Cancel

Reset to default

Apply

Two calendar widgets are provided to assist you in setting a new interval: the left calendar assists you in setting the start date, while the right calendar assists you in setting the end date.

To further assist you in setting a new interval, both calendars highlight:

Anomaly %

The percentage of observed KPI data points that fall outside the expected range of values in the reference time interval:

- < 6 : Low
- 6-10: Medium
- >10: High

Special days

Days on which a KPI trend is affected by seasonality factors such as holidays, vacation, business cycles, recurring events.

To set a time interval, run the following steps:

1. Modify the **Start Date** and **End Date** current values, or select the new start date and end date directly on the calendars. To set an interval within a single day, select the same day on both calendars.
2. Modify the **Start Time** and **End Time** current values.
3. Click **Apply**.

Results

A graph with the KPI trend in the new time interval is displayed.

Setting multiple time intervals

About this task

The **Multiple interval** section contains the following fields:

- **Start Time**
- **Interval duration** (days + hours)
- **End Time**

When you first open the Datepicker panel, these fields are set to the current time interval values in the KPI graph.

You can customize up to five intervals for comparison.

<
Edit or add intervals for comparison

Jobs in plan by workstation FOR **workstation: /WA-SERVER**

TIME INTERVAL: 18/1/2022, 00:00:00 → 19/1/2022, 00:00:00

Select the type of interval and set the corresponding information

Calendar indications: Anomaly % ● 0-5 ● 6-10 ● >10 ■ Special day Selected day Current day

Single Interval
Displays observed KPI data, estimated data range and anomalies for a single time interval.

Multiple intervals
Displays observed KPI data for up to 5 different time intervals.

1 - Define the common duration of all the time intervals.

Start time		Interval duration		End time
12:00 A...	→	1 Days	0:0 Hours	00:00 AM

2 - Add the time intervals that you want to compare (up to 5), and set the starting date for each of them.

Interval 1

🗓
01/18/2022
Add new interval +

Cancel

Reset to default

Apply

A calendar widget is provided to assist you in setting intervals. The calendar highlights:

Anomaly %

The percentage of observed KPI data points that fall outside the expected range of values in the reference time interval:

- < 6 : Low
- 6-10: Medium
- >10: High

Special days

Days on which a KPI trend is affected by seasonality factors such as holidays, vacation, business cycles, recurring events.

To set multiple time intervals (up to five), run the following steps:

1. Modify the **Start Time** and **Interval duration** values. The **End Time** value updates automatically.
2. For each time interval that you want to set, fill in the **Starting Date** field or use the calendar to set it.
3. Click **Add new interval** to set a new time interval.
4. When you have set all the desired time intervals, click **Apply**.
5. Select **Reset to default** to return to the original time interval, or **Close** to close the Datepicker panel.

Results

A graph with the KPI trend in the multiple time intervals is displayed.



Note: In the graphs showing KPI trends in multiple time intervals, the gray area representing the expected KPI values in each time interval is not displayed. On hovering over the KPI trends, a popover window displays the following information:

- Observation time
- KPI observed value for each time interval

Chapter 4. Managing alerts in AIDA

Learn how to manage and analyze alerts in AIDA.

Receiving alert notifications

Alerts can be notified through the Anomaly Widget on the Workload Dashboard or via email.

Receiving notifications through the Anomaly Widget

About this task

The Anomaly Widget on the Workload Dashboard indicates the number of Anomaly Alerts that have been generated in the last 24 hours. To analyze the alerts, run the following steps:

1. Click on the Anomaly Widget. A panel opens containing the list of the latest Anomaly Alerts. For each alert, the following information is displayed:
 - Alert severity
 - Alert description
 - A link to AIDA UI where you can find detailed alert information to quickly identify the root cause of the issue.
2. Follow the link for the alert that you want to analyze. For details, see [Analyzing an alert instance on page 49](#).
3. Click the **View all alerts** button to view the full list of alerts.

Receiving notifications via email

About this task

1. Alerts can be notified via email. To setup alert notifications via email, some configuration steps must be executed. For details, see [Configuring email alert settings on page 17](#).
2. By following the link provided in the notification email, administrators can access the alert instance page in AIDA UI and run an anomaly analysis. For details, see [Analyzing an alert instance on page 49](#).

Advanced notification

About this task

When an alert is found in AIDA, you can define an event rule in IBM Workload Scheduler to create a ticket on the supported service platform .

Overview dashboard

In the overview dashboard you can view the full list of alert instances.

About this task

From AIDA left-hand sidebar, select **Overview** to open the overview dashboard.

The screenshot shows the AIDA Overview dashboard. At the top, there's a header with 'AI DATA ADVISOR (AIDA)', 'Time Zone Europe/Rome (GMT + 02:00)', and 'Usernamehere Logout'. Below the header, the 'Overview' section displays five summary cards: 'Open Alerts instances' (28), 'Alerts Definitions' (6), 'Monitored Engines' (5), 'KPIs' (200), and 'Special days' (1). Each card has a 'See all' link. To the right, it says 'Latest KPI retrain for prediction : 29/03/2021 03:54 P.M.'. Below this is the 'Open alert instances' section, which includes filters for 'Anomaly KPI Source', 'Engine', and 'Object name', all set to 'All'. A table lists the instances with columns for 'Instances', 'Severity', 'Alert ID', 'Anomaly KPI source', 'Engine', and 'Object(s)'. The table contains 8 rows of data. At the bottom, there's a pagination control showing 'Items per page: 10' and 'Page 1 of 3'.

Instances	Severity	Alert ID	Anomaly KPI source	Engine	Object(s)
> 8	High	JOBWKS	Number jobs in plan by workstation (by status)	10.10.10.10	Workstation WKST123
> 2	High	JOBHISTORY	Job history (start time & duration)	00.00.00.00	Job ABCD123
> 1	Medium	JOBHISTORY	Job history (start time & duration)	00.00.00.00	Job JOB12345
> 14	Medium	MESSAGE	Available space for WA message files	10.10.10.10	All queues in plan
> 1	Medium	JOBFOLDER	Number jobs in plan by folder (by status)	00.00.00.00	Folder 12345
> 1	Medium	JOBTOTAL	Total number jobs in plan	10.10.10.10	All jobs in plan
> 1	Medium	JOBHISTORY	Job history (start time & duration)	10.10.10.10	Job 458JOB

A summary section contains the following information:

- The number of alert instances in open status.
- The number of alerts defined in AIDA.
- The number of the monitored engines.
- The number of the KPIs.
- The number of Special days set.



Note: Every 24 hours AIDA runs an automatic retrain of all KPIs. The start time of the automatic retrain is the start time of the Orchestrator container. After adding a special day to AIDA prediction model or making changes to the KPI prediction tuning, you will have to wait for the next retrain for the configuration to change.

The **Open Alert Instances** section contains a table with all the alert instances that have been generated. The instances have been grouped by alerts, an icon on the first column of each row allows to see the details of all the instances.

The data refreshes automatically every 15 minutes, you can also manually update it using the **Refresh Data** button .

The table displays the following information:

Instances

The number of open instances for each alert.

Severity

The severity of the alert instance. Displays the highest severity for the alert.

Alert ID

The name of the Alert definition corresponding to the alert.

Anomaly Source KPI

The KPI that generated the alert instance.

Engine

The engine where the KPI runs.

Objects

The name of the object measured by the KPI.



Note: To access the [Alert detail on page 36](#)page, click the icon in the last column.

To view the instance details for each alert, click the icon in the first column. This will expand the row and display the instances in a table format. For details, see [Alert detail on page 36](#)page.

AI DATA ADVISOR (AIDA) Time Zone Europe/Rome (GMT + 02:00) John Doe Logout

Overview Latest KPI retrain for prediction : 29/03/2021 03:54 P.M.

Open Alerts instances

28

Alerts Definitions

6

[See all](#)

Monitored Engines

5

[See all](#)

KPIs

200

[See all](#)

Special days

1

[See all](#)

Open alert instances Data update: every 15 min [Refresh data](#)

Anomaly KPI Source: All Engine: All Object name: All

Instances	Severity	Alert ID	Anomaly KPI source	Engine	Object(s)																																																
8	High	JOBWKS	Number jobs in plan by workstation (by status)	10.10.10.10	Workstation WKST123																																																
<table border="1"> <thead> <tr> <th>Job status</th> <th>Alert highest severity</th> <th>Trigger type</th> <th>First alert detection time</th> <th>Latest alert detection time</th> <th>Open alerts</th> </tr> </thead> <tbody> <tr> <td>Waiting</td> <td>High</td> <td>Continuous</td> <td>10/11/2024 - 04:45 P.M.</td> <td>10/11/2024 - 06:10 P.M.</td> <td>28</td> </tr> <tr> <td>Ready</td> <td>High</td> <td>Continuous</td> <td>10/11/2024 - 05:10 P.M.</td> <td>10/11/2024 - 04:45 P.M.</td> <td>24</td> </tr> <tr> <td>Running</td> <td>High</td> <td>Continuous</td> <td>10/11/2024 - 05:12 P.M.</td> <td>10/11/2024 - 04:45 P.M.</td> <td>21</td> </tr> <tr> <td>Successful</td> <td>High</td> <td>Continuous</td> <td>10/11/2024 - 04:45 P.M.</td> <td>10/01/2021 - 04:45 P.M.</td> <td>16</td> </tr> <tr> <td>Cancelled</td> <td>High</td> <td>Continuous</td> <td>10/11/2024 - 04:45 P.M.</td> <td>10/01/2021 - 04:45 P.M.</td> <td>13</td> </tr> <tr> <td>Hold</td> <td>High</td> <td>Continuous</td> <td>10/11/2024 - 04:45 P.M.</td> <td>10/01/2021 - 04:45 P.M.</td> <td>8</td> </tr> <tr> <td>Blocked</td> <td>Medium</td> <td>Continuous</td> <td>10/11/2024 - 04:45 P.M.M.</td> <td>10/01/2021 - 04:45 P.M.</td> <td>3</td> </tr> </tbody> </table>						Job status	Alert highest severity	Trigger type	First alert detection time	Latest alert detection time	Open alerts	Waiting	High	Continuous	10/11/2024 - 04:45 P.M.	10/11/2024 - 06:10 P.M.	28	Ready	High	Continuous	10/11/2024 - 05:10 P.M.	10/11/2024 - 04:45 P.M.	24	Running	High	Continuous	10/11/2024 - 05:12 P.M.	10/11/2024 - 04:45 P.M.	21	Successful	High	Continuous	10/11/2024 - 04:45 P.M.	10/01/2021 - 04:45 P.M.	16	Cancelled	High	Continuous	10/11/2024 - 04:45 P.M.	10/01/2021 - 04:45 P.M.	13	Hold	High	Continuous	10/11/2024 - 04:45 P.M.	10/01/2021 - 04:45 P.M.	8	Blocked	Medium	Continuous	10/11/2024 - 04:45 P.M.M.	10/01/2021 - 04:45 P.M.	3
Job status	Alert highest severity	Trigger type	First alert detection time	Latest alert detection time	Open alerts																																																
Waiting	High	Continuous	10/11/2024 - 04:45 P.M.	10/11/2024 - 06:10 P.M.	28																																																
Ready	High	Continuous	10/11/2024 - 05:10 P.M.	10/11/2024 - 04:45 P.M.	24																																																
Running	High	Continuous	10/11/2024 - 05:12 P.M.	10/11/2024 - 04:45 P.M.	21																																																
Successful	High	Continuous	10/11/2024 - 04:45 P.M.	10/01/2021 - 04:45 P.M.	16																																																
Cancelled	High	Continuous	10/11/2024 - 04:45 P.M.	10/01/2021 - 04:45 P.M.	13																																																
Hold	High	Continuous	10/11/2024 - 04:45 P.M.	10/01/2021 - 04:45 P.M.	8																																																
Blocked	Medium	Continuous	10/11/2024 - 04:45 P.M.M.	10/01/2021 - 04:45 P.M.	3																																																
2	High	JOBHISTORY	Job history (start time & duration)	00.00.00.00	Job ABCD123																																																
1	Medium	JOBHISTORY	Job history (start time & duration)	00.00.00.00	Job JOB12345																																																
14	Medium	MESSAGE	Available space for WA message files	10.10.10.10	All queues in plan																																																
1	Medium	JOBFOLDER	Number jobs in plan by folder (by status)	10.10.10.10	Folder 12345																																																

Items per page: 10 | 1 - 8 of 8 items Page 1 of 1

The table displays the following information:

Metric

The subdivision criteria of the alert instances. Depending on the KPI, it can be job status, star time/duration or queue.

Trigger type

Continuous or Total.

Highest Severity

The highest severity of all detected alerts of the alert instance. For details, see [Basic concepts on page 9](#).

First alert detection time

Date and time of the initial alert generation and the creation of the corresponding instance.

Latest alert detection time

Date and time of the most recent alert generated to date.

Total alerts

Total number of alerts generated since the initial detection.

To run a detailed analysis on an alert instance, click the open icon in the last column. For details, see [Analyzing an alert instance. on page 49](#)



Note: Alert instances in **Open** status are automatically marked as **Resolved** after a time period defined by the RESOLVE_ALERTS_AFTER_DAYS parameter configured for AIDA Exporter component (default value = 1 day).

Alert definitions

See all the alerts defined in AIDA. Pause one or more alerts immediately.

Before you begin

Alert definitions are contained in a json file inside IBM Workload Scheduler. It is retrieved by AIDA Exporter component and stored into the OpenSearch database. Alert definitions cannot be changed by users.

Each alert definition is based on a trigger, a specific set of conditions: the amount of consecutive anomalies during a specific time interval.

For example, 10 anomalous data points falling outside the expected range of values within 1 hour.

There are two types of alert triggers available for selection:

- **Continuous:** Triggers when anomalous data points fall above or below the predicted range.
- **Total:** Triggers for anomalous data points that are either above or below the predicted range, as well as those that exceed both thresholds.

See the content of the alert definition json file retrieved by IBM Workload Scheduler:

```
[
  {
    "definitionID": "CONTINUOUS_JOBWKS",
    "name": "Continuous anomalies for jobs in plan by workstation",
    "kpi": "application_wa_JobsByWorkstation_jobs",
    "trigger":{
      "type": "continuous",
      "value": 10,
      "timeFrame": 60,
      "description": "Over 10 Consecutive Anomalies within 1 hour"
    },
    "periodicity": "1 hour",
    "isActive": "true"
  },
  {
    "definitionID": "TOTAL_JOBWKS",
    "name": "Total Anomalies for Jobs in plan by Workstation",
    "kpi": "application_wa_JobsByWorkstation_jobs",
    "trigger":{
```

```

        "type": "total",
        "value": 3,
        "timeFrame": 60,
        "description": "Over 3 Anomalies within 1 hour"
    },
    "periodicity": "1 hour",
    "isActive": "false"
},
{
    "definitionID": "CONTINUOUS_JOBFOLDER",
    "name": "Continuous Anomalies Jobs in plan by Folder",
    "kpi": "application_wa_JobsByFolder_jobs",
    "trigger":{
        "type": "continuous",
        "value": 10,
        "timeFrame": 60,
        "description": "Over 10 Consecutive Anomalies within 1 hour"
    },
    "periodicity": "1 hour",
    "isActive": "true"
},

{
    "definitionID": "TOTAL_JOBFOLDER",
    "name": "Total Anomalies for Jobs in plan by Folder",
    "kpi": "application_wa_JobsByFolder_jobs",
    "trigger":{
        "type": "total",
        "value": 3,
        "timeFrame": 60,
        "description": "Over 3 Anomalies within 1 hour"
    },
    "periodicity": "1 hour",
    "isActive": "false"
},
{
    "definitionID": "CONTINUOUS_JOBSTATUS",
    "name": "Continuous Anomalies for Jobs in plan by status",
    "kpi": "application_wa_JobsInPlanCount_job",
    "trigger":{
        "type": "continuous",
        "value": 10,
        "timeFrame": 60,
        "description": "Over 10 Consecutive Anomalies within 1 hour"
    },
    "periodicity": "1 hour",
    "isActive": "true"
},
{
    "definitionID": "TOTAL_JOBSTATUS",
    "name": "Total Anomalies for jobs in plan by status",
    "kpi": "application_wa_JobsInPlanCount_job",
    "trigger":{
        "type": "total",
        "value": 3,
        "timeFrame": 60,
        "description": "Over 3 Anomalies within 1 hour"
    },

```

```

"periodicity": "1 hour",
"isActive": "false"
},
{
"definitionID": "CONTINUOUS_JOBTOTAL",
"name": "Continuous Anomalies for total jobs",
"kpi": "application_wa_JobsInPlanCount_job_total",
"trigger":{
  "type": "continuous",
  "value": 10,
  "timeFrame": 60,
  "description": "Over 10 Consecutive Anomalies within 1 hour"
},
"periodicity": "1 hour",
"isActive": "true"
},
{
"definitionID": "TOTAL_JOBTOTAL",
"name": "Total Anomalies for total jobs",
"kpi": "application_wa_JobsInPlanCount_job_total",
"trigger":{
  "type": "total",
  "value": 3,
  "timeFrame": 60,
  "description": "Over 3 Anomalies within 1 hour"
},
"periodicity": "1 hour",
"isActive": "false"
},
{
"definitionID": "CONTINUOUS_JOBHISTORY",
"name": "Continuous Anomalies for job history",
"kpi": "job_history",
"trigger":{
  "type": "continuous",
  "value": 2,
  "timeFrame": 2880,
  "description": "Over 10 Consecutive Anomalies within 2 days"
},
"periodicity": "1 hour",
"isActive": "true"
},
{
"definitionID": "TOTAL_JOBHISTORY",
"name": "Total Anomalies for job history",
"kpi": "job_history",
"trigger":{
  "type": "total",
  "value": 2,
  "timeFrame": 2880,
  "description": "Over 3 Anomalies within 2 days"
},
"periodicity": "1 hour",
"isActive": "false"
},
{
"definitionID": "CONTINUOUS_MESSAGE",
"name": "Continuous Anomalies for message files fill percentile",

```

```

"kpi": "application_wa_msgFileFill_percent",
"trigger":{
    "type": "continuous",
    "value": 10,
    "timeFrame": 60,
    "description": "Over 10 Consecutive Anomalies within 1 hour"
},
"periodicity": "1 hour",
"isActive": "true"
},
{
"definitionID": "TOTAL_MESSAGE",
"name": "Total Anomalies for message files fill percentile",
"kpi": "application_wa_msgFileFill_percent",
"trigger":{
    "type": "total",
    "value": 3,
    "timeFrame": 60,
    "description": "Over 3 Anomalies within 1 hour"
},
"periodicity": "1 hour",
"isActive": "false"
},
{
"definitionID": "CONTINUOUS_INCOMPLETEPREDECESSOR_CRITICAL",
"name": "WA critical job incomplete predecessor",
"kpi": "application_wa_criticalJob_incompletePredecessor_jobs",
"trigger": {
    "type": "continuous",
    "value": 10,
    "timeFrame": 60,
    "description": "Over 10 Consecutive Anomalies within 1 hour"
},
"periodicity": "1 hour",
"isActive": "true",
>alert-definition": "CONTINUOUS"
},
{
"definitionID": "TOTAL_INCOMPLETEPREDECESSOR_CRITICAL",
"name": "WA critical job incomplete predecessor",
"kpi": "application_wa_criticalJob_incompletePredecessor_jobs",
"trigger": {
    "type": "total",
    "value": 10,
    "timeFrame": 60,
    "description": "Over 10 Anomalies within 1 hour"
},
"periodicity": "1 hour",
"isActive": "true",
>alert-definition": "TOTAL"
}
]

```

About this task

From AIDA left-hand sidebar, select **Alert Definitions**.

Alert ID	Related anomaly source KPI	Anomaly type	Metric (division by)	Alert triggers (both type continuous & total)	KPI
JOBWKS	Number of jobs in plan by workstation (by status)	Higher/lower jobs number	Job Status (10)	>= 10 consecutive anomalies within 1h	63 >
JOBFOLDER	Number of jobs in plan by folder (by status)	Higher/lower jobs number	Job Status (10)	>= 10 consecutive anomalies within 1h	48 >
JOBSTATUS	Number of jobs in plan by status	Higher/lower jobs number	Job Status (10)	>= 10 consecutive anomalies within 1h	6 >
JOBTOTAL	Number total of jobs in plan	Higher/lower jobs number	Job Status (10)	>= 10 consecutive anomalies within 1h	6 >
MESSAGE	Available space for WA message files	Finishing space for queue	Queues	>= 10 consecutive anomalies within 1h	6 >
PREDECESSOR	Number of incomplete predecessors for critical job	Higher/Lower incomplete predecessor number	Critical job	>= 5 consecutive anomalies within 30 min	28 >

In this page you can view the full list of alert definitions in table format.

The table displays the following information:

Alert ID

The ID of the alert definition.

Anomaly Source KPI

The KPIs that contribute to generate the alert.

Anamoly type

The type of anomaly that generates the alert (for example: a higher or lower amount of job than expected).

Metric

The sub-division criteria of the alert instances. Depending on the KPI, it can be job status, star time/duration or queue.

Alert Trigger

Set of conditions defining the alert. For example: Over 3 anomalies within 1 hour.

KPI

The number of KPIs defined by this alert definition.

Click on an alert definition row (or the icon on the right) to open a side panel displaying detailed information and a tabular list of all alerts organized by the KPIs specified in the alert definition.

The screenshot displays the AIDA interface. On the left, the 'Alert Definitions' section lists various alert types with their related KPIs and anomaly types. On the right, the 'Alert ID: JOBWKS' detail page is shown, including alert definition information, a table of active alerts by KPIs, and a 'Deactivate all alerts' button.

Alert Definitions

All available alerts definitions for each anomaly source KPI.
Alert trigger types: CONTINUOUS (above OR below expected data) or TOTAL (above OR below + above AND below)

Alert ID	Related anomaly source KPI	Anomaly type
JOBWKS	Number of jobs in plan by workstation (by status)	Higher/lower jobs number
JOBFOLDER	Number of jobs in plan by folder (by status)	Higher/lower jobs number
JOBSTATUS	Number of jobs in plan by status	Higher/lower jobs number
JOBTOTAL	Number total of jobs in plan	Higher/lower jobs number
MESSAGE	Available space for WA message files	Finishing space for queue
PREDECESSOR	Number of incomplete predecessors for critical job	Higher/Lower incomplete p

Alert ID: JOBWKS

Alert definition information

Anomaly Source KPI: Number of jobs in plan by workstation (by status)
Metric: Job Status (10)
Anomalies: Higher/lower jobs number

Alert trigger: ≥ 10 consecutive anomalies within 1h

- CONTINUOUS (above OR below expected data)
- TOTAL (above OR below + above AND below)

Alerts by KPIs (63) Deactivate all alerts

Engines	Workstation	Trigger option	Alert status
10.14.82.155.31116	/P20_200	<input type="radio"/> Continuous <input checked="" type="radio"/> Total	<input checked="" type="checkbox"/> Active 🔗
10.14.82.155.31116	/TEST_FTA_1	<input type="radio"/> Continuous <input checked="" type="radio"/> Total	<input checked="" type="checkbox"/> Active 🔗
10.14.82.155.31116	/RMMYCLDDL73...	<input type="radio"/> Continuous <input checked="" type="radio"/> Total	<input checked="" type="checkbox"/> Active 🔗
10.14.82.155.31116	/RMMYCLDDL73611	<input type="radio"/> Continuous <input checked="" type="radio"/> Total	<input checked="" type="checkbox"/> Active 🔗
10.14.82.155.31116	/DAUNIX	<input type="radio"/> Continuous <input checked="" type="radio"/> Total	<input checked="" type="checkbox"/> Active 🔗
10.14.82.155.31116	/ADCAPP364	<input type="radio"/> Continuous <input checked="" type="radio"/> Total	<input checked="" type="checkbox"/> Active 🔗
10.14.82.155.31116	/ADCAPP241	<input type="radio"/> Continuous <input checked="" type="radio"/> Total	<input checked="" type="checkbox"/> Active 🔗
10.14.82.155.31116	/DAUNIX2	<input type="radio"/> Continuous <input checked="" type="radio"/> Total	<input checked="" type="checkbox"/> Active 🔗
10.14.82.155.31116	/MASTER_DA	<input type="radio"/> Continuous <input checked="" type="radio"/> Total	<input checked="" type="checkbox"/> Active 🔗

For each alert, you can modify the trigger type and activate or deactivate alert generation. A global option, **Deactivate All Alerts**, is also available to disable alert generation for all alerts in the list. Click the action icon on the right side of each row to access the [Alert detail page on page 45](#), where you can view alert information, open alert instances, and the history for the past 12 months.

Alert details

See detailed information about an alert definition, its status, and history.

Before you begin

The **Alert Details** page provides you with information about an alert, its status, the current opened instances and its history.

AI DATA ADVISOR (AIDA)
Time Zone Europe/Rome (GMT + 02:00) ⓘ
Usernamehere Logout

[Alert rules](#) / Alert

Alert : MESSAGE / Engine 00.00.00.00

KPI information

Name: Available space for WA Message files
Engine: 00.00.00.00

Alert Generation

Active

Anomaly Type

Finishing space for WA Message files (by Queue)

Trigger

Type: >= 10 consecutive anomalies within 1h
Options: Continuous Total

Opened Alert Instances ⓘ

Data update: every 15 min Refresh data ↻ Resolve all ✓ See all Resolved Alerts

Queue	Highest Severity ⓘ	Trigger type	First Alert Detection Time	Latest Alert Detection Time	Total Alerts	
Courier.msg	High	Continuous	10/11/2024 - 01:45 P.M.	10/11/2024 - 01:45 P.M.	203	✔ 📄 🔗
mirrorbox.msg	High	Continuous	10/11/2024 - 01:45 P.M.	10/11/2024 - 01:45 P.M.	150	✔ 📄 🔗
Mailbox.msg	High	Continuous	10/11/2024 - 01:45 P.M.	10/11/2024 - 01:45 P.M.	110	✔ 📄 🔗
Monbox.msgn	High	Continuous	10/11/2024 - 01:45 P.M.	10/11/2024 - 01:45 P.M.	98	✔ 📄 🔗
auditbox.msg	High	Continuous	10/11/2024 - 01:45 P.M.	10/11/2024 - 01:45 P.M.	67	✔ 📄 🔗
clbox.msg	Medium	Continuous	10/11/2024 - 01:45 P.M.	10/11/2024 - 01:45 P.M.	50	✔ 📄 🔗
planbox.msg	Medium	Continuous	10/11/2024 - 01:45 P.M.	10/11/2024 - 01:45 P.M.	34	✔ 📄 🔗
ntercom.msg	Medium	Continuous	10/11/2024 - 01:45 P.M.	10/11/2024 - 01:45 P.M.	25	✔ 📄 🔗
pobox messages	Medium	Continuous	10/11/2024 - 01:45 P.M.	10/11/2024 - 01:45 P.M.	12	✔ 📄 🔗

Items per page: 10 | 1 - 9 of 9 items Page 1 of 1 ◀ 1 ▶

Last 12 months history ⓘ

April 2024

Mon	Tue	Wed	Thu	Fri	Sat	Sun
29	30	31	1	2	3	4
5	6	7	8	9	10	11
12	13	14	15	16	17	18
19	20	21	22	23	24	25
26	27	28	29	30		

May 2024

Mon	Tue	Wed	Thu	Fri	Sat	Sun
29	30	31	1	2	3	4
5	6	7	8	9	10	11
12	13	14	15	16	17	18
19	20	21	22	23	24	25
26	27	28	29	30		

June 2024

Mon	Tue	Wed	Thu	Fri	Sat	Sun
29	30	31	1	2	3	4
5	6	7	8	9	10	11
12	13	14	15	16	17	18
19	20	21	22	23	24	25
26	27	28	29	30		

You can also **Activate** and **Deactivate** an alert generation, and change the trigger type. This page is composed of three sections:

- Alert Definition
- Alert Instances
- Last 12 months history

Alert Information

About this task

This section provides the following information about an alert, its definition and status.

Anomaly Source KPI

The KPIs that generated the alert instance .

Alert generation

The status of the alert. It can be Active or Deactive.

Anamoly type

The type of data that is monitored by the KPI that may trigger the alert.

Alert Trigger

Set of conditions defining the alert.

Opened Alert Instances

About this task

The **Alert Instances** section contains a table with all the alert instances that have been generated. The table displays the following information:

Metric

The subdivision criteria of the alert instances. Depending on the KPI, it can be job status, star time/duration or queue.

Trigger type

Continuous or Total.

Highest Severity

The highest severity of all detected alerts of the alert instance. For details, see [Basic concepts on page 9](#).

First alert detection time

Date and time of the initial alert generation and the creation of the corresponding instance.

Latest alert detection time

Date and time of the most recent alert generated to date.

Total alerts

Total number of alerts generated since the initial detection.

Three action icons are available for each alert instance:

- **Resolve:** to resolve the instance.

You can resolve all instances by clicking the button at the top of the table. To view only the resolved instances, click the **See All Resolved Alerts** button at the top of the table. This action will open a side panel displaying the resolved instances from the past 14 days in a tabular format. Additionally, you can filter the view by KPI metrics such as job status, queue, duration, or start time.

Note: Alert instances in **Open** status are automatically marked as **Resolved** after a time period defined by the RESOLVE_ALERTS_AFTER_DAYS parameter configured for AIDA Exporter component (default value = 1 d) .

The screenshot displays the AI Data Advisor (AIDA) interface. The top navigation bar includes the title 'AI DATA ADVISOR (AIDA)', the time zone 'Europe/Rome (GMT + 02:00)', and user information 'Usernamehere' and 'Logout'. The main content area is divided into several sections:

- Alert rules / Alert:** Shows an alert rule for 'MESSAGE / Engine 00.00.00.00' with KPI information: 'Name: Available space for WA Message files' and 'Engine: 00.00.00.00'.
- Opened Alert Instances:** A table listing various message queues and their highest severity levels. For example, 'Courier.msg' has a 'High' severity, while 'clbox.msg' has a 'Medium' severity.
- Last 12 months history:** A calendar view for April 2024, with the 15th highlighted in yellow.
- Resolved Instances:** A table showing resolved alert instances. The table has columns for Severity, Trigger type, Latest Alert Detection, and Total Alerts. The first row shows a 'High' severity instance for 'Appserverbox.msg' with a 'Continuous' trigger type and a total of 150 alerts.

- **Comment:** you can add a comment about the alert instance and its resolution. Click **Save** to save your comments.
- **Open instance page:** to access the Instance detail page and run a detailed analysis. For details, see [Analyzing an alert instance on page 49](#).

Last 12 months history

About this task

This section shows a calendar representation of the alert history for the last 12 months, highlighting the days affected by the alert.

For each day affected by the alert, a colored circle represents the daily highest severity for the alert. On hovering over each day, you can also see the daily number of alert instances generated and, if present, the Special Day label.

Click **Show more** to see the following months.

Pause alert generation

About this task

You can pause and resume an alert generation.

When an alert status is **Active**, you can pause the alert generation with immediate effect by clicking the **Pause** button.

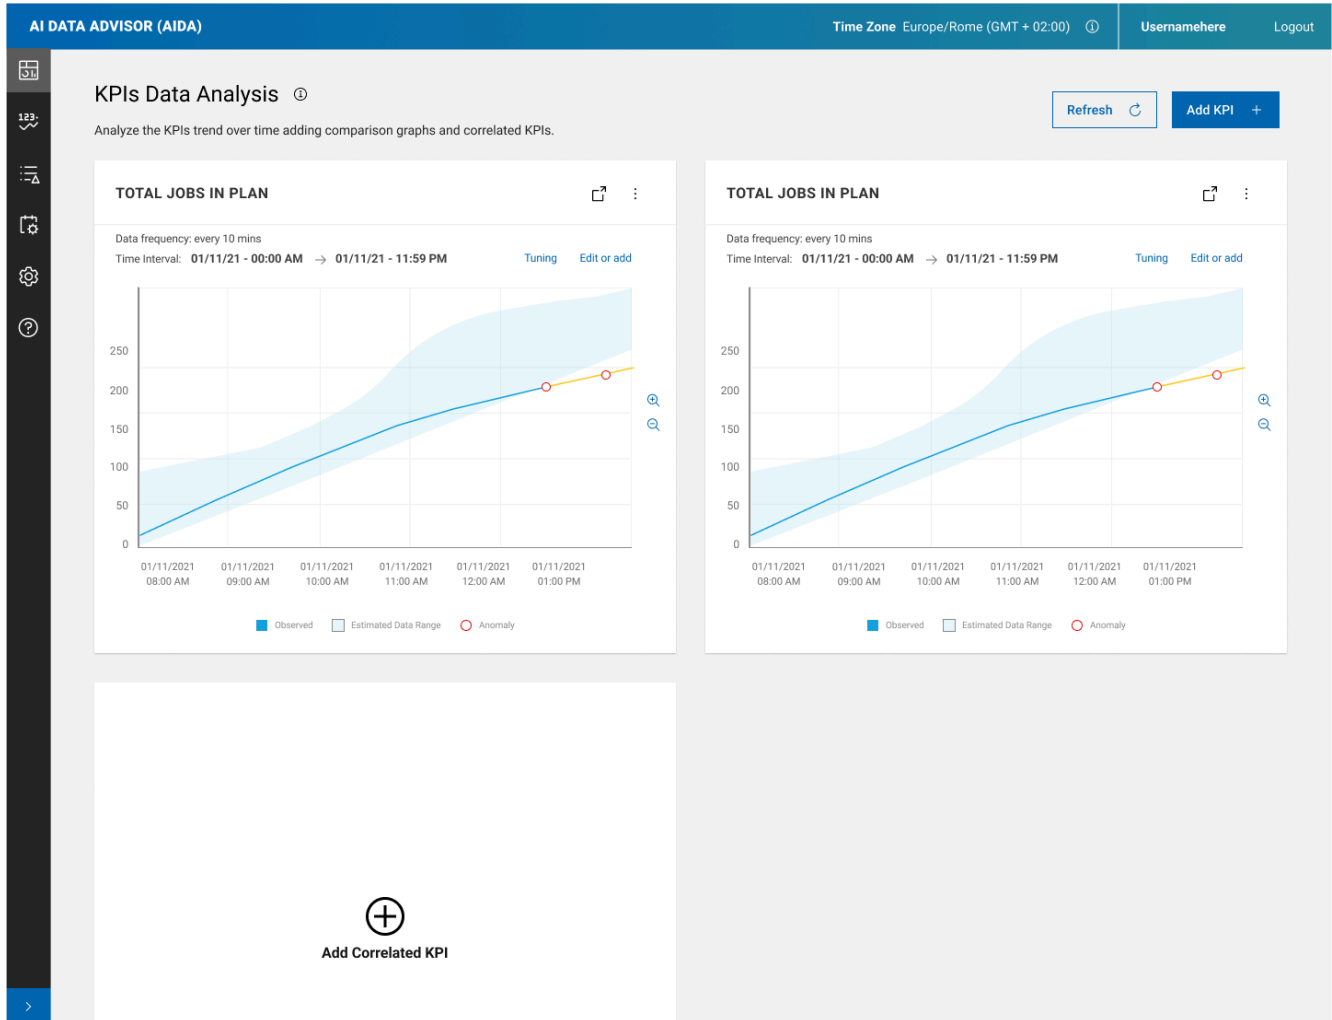
When the alert status is **Paused**, you can resume the alert with immediate effect by clicking the **Resume** button.

Analyzing an alert instance

By using Machine Learning techniques to predict KPIs time series, AIDA can detect anomalies in a KPI trend and help you quickly identify the root cause of problems.

Before you begin

When anomalies in a KPI trend generate an alert, from the **Alert Instance Details** page you can analyze the anomalous trend and compare it with the trend over different time intervals. You can also add correlated KPIs to the data analysis to find root causes faster.



You can reach the Alert Instance Details page in different ways:

- From the Anomaly Widget on the Workload Dashboard.
- From AIDA menu on the left-hand sidebar, by clicking on **Overview** and selecting an alert instance.
- From AIDA menu on the left-hand sidebar, by clicking **Alert Definitions**, selecting an alert and then an alert instance.
- From the link provided in the alert email notification (for AIDA administrators only).

When an alert is detected by AIDA, if you think it is a false alert, consider the following observations:

- **AIDA's prediction model might not have enough data yet**

Aida uses a machine learning model based on historical data. Maybe the model still has little data available to make accurate predictions. In this case, wait for the model to get more data.

- **The Machine Learning algorithm might need some tuning**

KPIs prediction is based on a number of tuning parameters, such as the tolerance interval width, that must be properly customized. Try to better configure the tuning parameters and run a retrain process to recalculate the prediction interval with the new parameters.

- **You might want to pause the alert**

If you think this alert is a false alert, or you don't want to be bothered by this alert for the next few hours or days, you can pause the alert generation.

About this task

In the **Alert Instance Details** page, you can find a summary section with the following alert instance information:

Instance ID

The alert instance identifier.

Detection Time

Date and time when the alert instance was generated.

Severity

The severity of the alert instance. For details, see [Basic concepts on page 9](#).

Resolution Status

The alert instance status: can be: **Open** or **Resolved**. Here you can change the status of the alert instance that you are analyzing. Alert instances in **Open** status are automatically marked as **Resolved** after a time period defined by the RESOLVE_ALERTS_AFTER_DAYS parameter configured for AIDA Exporter component (default value = 1 day)..

Special Days

Specifies if the alert instance was detected on a special day or not.

Related Alert

The name of the related alert.

Tag

The tag identifies the engine to which the KPI refers.

Alert Trigger

Set of conditions defining the alert.

The **Anomaly Source KPI** section shows graphs related to the KPI anomalous trend that you can compare with the trend over different time intervals. A comparison graph is shown, by default, to the right of the KPI graph, representing the KPI trend on the previous day. You can also add correlated KPIs to the data analysis to find the root cause of problem.

In each KPI graph, data is displayed in time buckets. The KPI data frequency and the reference time interval are indicated in the graph header. To view all data points, click on the **Edit or add** link and reduce the time interval. A **light-blue area** represents the expected range of values for the KPI in the reference time interval, statistically defined by AIDA based on historical data. The **blue line** represents the observed data that falls within the expected range of values. The **orange line** represents data that falls outside the expected range of values. You can zoom in or zoom out on the graph. On hovering over the KPI trend, data points appear. For each data point, a popover window displays the following information:

- Date and time of the observation.
- Current value: the KPI observed value.
- Estimated: the KPI interval estimation.

The anomalies in the KPI trend are represented by **red circles**. Among these, the anomalies that contribute to generating an alert are represented by **red dots**. On hovering over an anomaly, the following information is displayed:

- Date and time of the observation.
- Current value: the KPI observed value.
- Estimated: the KPI interval estimation.
- Deviation: the minimum distance (with - or + sign) of the KPI observed value from its interval estimation.

In each KPI graph, you can run a number of actions:

- Click on **Edit or add** to edit the graph time interval, or add time intervals to the graph for comparison purposes. For details, see the task **Setting time intervals with the Datepicker** below.
- The menu icon in the upper right corner of the graph contains the following additional actions:
 - **Compare graph**, to create a comparison graph with single or multiple time intervals for comparison purposes.
 - **Tuning**, to tune a KPI prediction parameter, the anomaly detection sensitivity in the KPIs prediction tuning popover. This action allows you to increase or decrease sensibility, with higher sensibility identifying more anomalies. You can also modify all KPIs using the **Global tuning** button, which overrides individual KPI settings. These changes will be applied after the next retraining. Tuning is available to AIDA administrators only.
 - **Refresh**, to refresh the graph after you run some tuning adjustments.
 - **Delete**, to delete the graph.
- For KPIs belonging to the **Jobs** category, an action icon is also present to open the workstation or job properties panel directly in IBM Workload Scheduler.

To deepen your analysis, you can add additional graphs to your anomaly source KPI graph:

- Comparison graphs with the KPI trend over different time intervals
- Correlated KPI graphs

Adding comparison graphs

About this task

You can add comparison graphs, both to anomaly source KPIs and correlated KPIs.

For the KPI: **Jobs in plan by status**

For the KPI **Jobs in plan by status**, a comparison graph is shown, by default, representing the KPI trend during the previous day.

For both graphs, you can edit the time interval, or add time intervals for comparison:

- In the graph that you want to modify, click **Edit or add** to open the Datepicker panel where you can:
 - edit the time interval
 - add time intervals for comparison

For the remaining KPIs

You can edit the time interval, or add multiple time intervals to any KPI graph for comparison:

- In the graph, click **Edit or add** to open the Datepicker panel where you can:
 - edit the time interval
 - add time intervals for comparison

To enhance the analysis, you can generate an additional graph.

- From the menu icon in the upper right corner of the graph, select **Compare graph**. The Datepicker panel opens where you can create a comparison graph with single or multiple time intervals.

For details about Datepicker, see [Setting time intervals with the Datepicker on page 54](#).



Note: In the graphs showing KPI trends in multiple time intervals, the gray area representing the expected KPI values in each time interval is not displayed.

Adding correlated KPIs

About this task

You can add one or more correlated KPIs to the anomaly data analysis from the **Add KPI** panel that you can open in either of the following ways:

- In the Correlated KPI area, click **Add Correlated KPI**.
- In the upper right corner of the Anomaly Data Analysis UI, click the **Add KPI** button.

On the left-hand side of the **Add KPI** panel, select a KPI category.

For each KPI of the selected category, the following information is displayed:

KPI Name

Name of the KPI

Object Name

The name of the object measured by the KPI.

Tag

A search tag for the KPI. Usually, the tag is used to identify the engine to which the KPI refers.

Anomaly %

The percentage of observed KPI data points that fall outside the expected range of values in the reference time interval:

- < 6 : Low
- 6-10: Medium
- >10: High

To select correlated KPIs, run the following steps:

1. Use the search bar to refine your search.
2. Select one or more KPIs.
3. Click the **Add KPI** button.

Results

A new graph for each selected KPI is added to the Correlated KPI area, representing the KPI trend in the reference time interval.

As for the anomaly source KPIs, you can add comparison graphs to the correlated KPIs. For details, see [Adding comparison graphs on page 53](#).

Setting time intervals with the Datepicker

Before you begin

Use the Datepicker to set single or multiple time intervals in a KPI graph.

In the Datepicker panel, select the type of interval:

Single Interval

- To edit a time interval in a KPI graph
- To add a KPI comparison graph with a single time interval

Multiple Intervals

- To edit multiple time intervals in a KPI graph
- To add a KPI comparison graph with multiple time intervals

Setting a single time interval

About this task

The **Single interval** section contains the following fields:

- **Start Date**
- **Start Time**
- **End Date**
- **End Time**

When you first open the Datepicker panel, these fields are set to the current time interval values in the KPI graph.

< Edit or add intervals for comparison

Jobs in plan by workstation FOR workstation: /WA-SERVER

TIME INTERVAL: 18/1/2022, 00:00:00 → 19/1/2022, 00:00:00

Select the type of interval and set the corresponding information

Calendar indications: Anomaly % ● 0-5 ● 6-10 ● >10 ■ Special day Selected day Current day

Single Interval
 Displays observed KPI data, estimated data range and anomalies for a single time interval.

Start day: Start time: AM → End day: End time: AM

Multiple intervals
 Displays observed KPI data for up to 5 different time intervals.

Cancel

Reset to default

Apply

Two calendar widgets are provided to assist you in setting a new interval: the left calendar assists you in setting the start date, while the right calendar assists you in setting the end date.

To further assist you in setting a new interval, both calendars highlight:

Anomaly %

The percentage of observed KPI data points that fall outside the expected range of values in the reference time interval:

- < 6 : Low
- 6-10: Medium
- >10: High

Special days

Days on which a KPI trend is affected by seasonality factors such as holidays, vacation, business cycles, recurring events.

To set a time interval, run the following steps:

1. Modify the **Start Date** and **End Date** current values, or select the new start date and end date directly on the calendars. To set an interval within a single day, select the same day on both calendars.
2. Modify the **Start Time** and **End Time** current values.
3. Click **Apply**.

Results

A graph with the KPI trend in the new time interval is displayed.

Setting multiple time intervals

About this task

The **Multiple interval** section contains the following fields:

- **Start Time**
- **Interval duration** (days + hours)
- **End Time**

When you first open the Datepicker panel, these fields are set to the current time interval values in the KPI graph.

You can customize up to five intervals for comparison.

<
Edit or add intervals for comparison

Jobs in plan by workstation FOR **workstation: /WA-SERVER**

TIME INTERVAL: 18/1/2022, 00:00:00 → 19/1/2022, 00:00:00

Select the type of interval and set the corresponding information

Calendar indications: Anomaly% ● 0-5 ● 6-10 ● >10 ■ Special day Selected day Current day

Single Interval
Displays observed KPI data, estimated data range and anomalies for a single time interval.

Multiple intervals
Displays observed KPI data for up to 5 different time intervals.

1 - Define the common duration of all the time intervals.

Start time Interval duration End time

12:00 A... → 1 Days 0:0 Hours 00:00 AM

2 - Add the time intervals that you want to compare (up to 5), and set the starting date for each of them.

Interval 1

01/18/2022 Add new interval +

Cancel

Reset to default

Apply

A calendar widget is provided to assist you in setting intervals. The calendar highlights:

Anomaly %

The percentage of observed KPI data points that fall outside the expected range of values in the reference time interval:

- < 6 : Low
- 6-10: Medium
- >10: High

Special days

Days on which a KPI trend is affected by seasonality factors such as holidays, vacation, business cycles, recurring events.

To set multiple time intervals (up to five), run the following steps:

1. Modify the **Start Time** and **Interval duration** values. The **End Time** value updates automatically.
2. For each time interval that you want to set, fill in the **Starting Date** field or use the calendar to set it.
3. Click **Add new interval** to set a new time interval.

4. When you have set all the desired time intervals, click **Apply**.
5. Select **Reset to default** to return to the original time interval, or **Close** to close the Datepicker panel.

Results

A graph with the KPI trend in the multiple time intervals is displayed.



Note: In the graphs showing KPI trends in multiple time intervals, the gray area representing the expected KPI values in each time interval is not displayed. On hovering over the KPI trends, a popover window displays the following information:

- Observation time
- KPI observed value for each time interval

Chapter 5. Troubleshooting AIDA

See how to troubleshoot problems in AIDA.

This section describes:

- How to [collect logs and activate traces on page 59](#) in AIDA:
- How to troubleshoot problems

Logging and tracing in AIDA

How to configure logging and tracing in AIDA

Log files for each AIDA component are located inside the respective container.

Each container, except for UI, supports five logging levels: DEBUG, INFO, ERROR, WARNING, and CRITICAL.

UI container supports three logging levels: ERROR, INFO, and TRACE.

By default, after the installation only informational messages are logged. If you want to change log level, run the following commands.

Docker installation

All containers except for UI

Run `docker run --env LOG_LEVEL=log_level -it [container_name]`

where `log_level` can be DEBUG, INFO, ERROR, WARNING, CRITICAL

UI

1. In the `configuration.sh` file available in the UI installation package, locate the script `./configuration.sh {option}`
2. Run the script with option:

`--error-log-level` to update logging level to ERROR

`--info-log-level` to update logging level to INFO

`--trace-log-level` to update logging level to TRACE

Kubernetes installation

All containers except for UI

1. Edit the `values.yaml` file and set `log_level` to any of the available options: DEBUG, INFO, ERROR, WARNING, CRITICAL
2. To update the container configuration, run the command:

`helm upgrade [container_name] [path_of_values.yaml_file]`

UI

1. In the `configuration.sh` file available in the UI installation package, locate the script `./configuration.sh {option}`

2. Run the script with option:

`--error-log-level` to update logging level to ERROR

`--info-log-level` to update logging level to INFO

`--trace-log-level` to update logging level to TRACE

Troubleshooting AIDA

See how to troubleshoot problems in AIDA.

This section describes:

- How to resolve **DNS** error

If you are facing problem to resolve the public name of the AIDA host machine (to connect to AIDA you must use only public address) and execute these steps:

- Open the command prompt and ping the IP address(hostname)



Note: If the ping does not work, the Docker also may not work.

- If the ping on the hostname works, then try restart the Docker.
 - (systemctl stop docker -- systemctl start docker)

Notices

This document provides information about copyright, trademarks, terms and conditions for product documentation.

© Copyright IBM Corporation 1993, 2016 / © Copyright HCL Technologies Limited 2016, 2025

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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, MD-NC119
Armonk, NY 10504-1785
US

For license inquiries regarding double-byte character set (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.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

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 jurisdictions 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide 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 Director of Licensing

IBM Corporation

North Castle Drive, MD-NC119

Armonk, NY 10504-1785

US

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.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

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.

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 actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

© (HCL Technologies Limited) (2025).

Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp. 2016

Trademarks

IBM, the IBM logo, 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.

IT Infrastructure Library™ is a Registered Trade Mark of AXELOS Limited.

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.

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.



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.

ITIL™ is a Registered Trade Mark of AXELOS Limited.

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

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Index

A

- accessibility v
- accessing
- AIDA
 - 13
 - Add a graph for comparison 53
 - Add a special day 19
 - Add correlated KPIs for comparison 53
 - Add country specific holidays 21
 - Adding a country specific holiday 21
 - adding a custom date 21
 - adding a custom special day 21
 - Adding comparison graphs 53
 - Adding correlated KPIs 53
 - Adding country specific holidays 21
 - adding custom special days 21
 - adding engines to AIDA 15
 - Adding special days 19
- Administrator
 - AIDA
 - 15
 - AI 6, 7, 7
 - AI Data Advisor
 - 45, 49
 - troubleshooting 59, 60
- AIDA
 - 6, 7, 7, 36, 45, 49
 - administration 15
 - configuration 15
 - troubleshooting 59, 60
- AIDA architecture 11
- AIDA
 - basic concepts
 - 9
 - entry points
 - 13
 - installation and configuration
 - 13
 - AIDA security 18
 - AIDA security settings 18
 - terminology
 - 9
 - alert 45, 49
 - definition 9
 - Alert definition 46
 - alert definitions 40
 - Alert history 47, 48
 - alert information 46
 - Alert Information 45
 - Alert instances 49
 - alert notification 36
 - alert notifications 36
 - alerts 36
 - all alert instances 36
 - Analyzing an alert instance 49
 - analyzing KPIs 23, 23
 - Analyzing KPIs 28
 - anomalies 36
 - anomaly 36
 - definition 9
 - anomaly alert definitions 40
 - anomaly alerts 40
 - anomaly detection 6, 7, 7, 36
 - Artificial Intelligence 6, 7, 7

- automatic prioritization 6, 7, 7

B

- business scenario 7, 7

C

- Cloud & Smarter Infrastructure technical training v
- Configuring
 - AIDA
 - 15
 - configuring email settings 17
- Correlated KPI graphs 53
- current and historical alert information 45

D

- data-driven decisions 6, 7, 7
- datapicker panel 32, 54
- defining anomaly alerts 40
- Dynamic Workload Console
 - accessibility v

E

- education v

H

- how to access
 - AIDA
 - 13
 - how to add engines to AIDA 15
 - How to analyze KPIs in
 - AIDA
 - 28
 - how to deploy
 - AIDA
 - 13
 - How to manage KPIs in
 - AIDA
 - 26
 - how to set time intervals 32, 54

K

- Key Performance Indicators 23, 23

L

- Logging and tracing in
 - AIDA
 - 59
- Logs 59

M

- Machine Learning 6, 7, 7
- mail notifications 36
- managing alerts 36
- managing KPIs 23, 23
- Managing KPIs 26
- microservices 11
- microservices-based architecture 11

O

- overview dashboard 36

P

- pause alert 49
- pause alert generation 49
- proactive SLA management 6, 7, 7
- problem prevention 36

R

- receiving alert notifications 36
- retraining 19

S

- seasonality 19
- Setting time intervals 32, 54
- Special Days 19
- special days in KPI trend 19

T

- technical training v
- time intervals 32, 54
- traces 59
- training
 - technical v
- trigger
 - definition 9
- troubleshooting 59, 60

V

- view all alert instances 36
- view all alerts 36

W

- Working with alerts 36
- Working with KPIs 23