

**IBM® Workload Scheduler
AI Data Advisor User's Guide
Version 10.2.1**

Note

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

This edition applies to version 10, release 2, modification level 1 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.....	16
Adding engines to AIDA.....	16
Configuring email alert settings.....	16
Configuring AIDA for KPIs prediction.....	17
Configuring security.....	21
AIDA settings.....	22
Managing special days.....	23
Chapter 3. Working with KPIs.....	27
KPIs for IBM Workload Scheduler.....	27
Managing KPIs in AIDA.....	29
Analyzing KPIs data.....	32
Chapter 4. Managing alerts in AIDA.....	40
Receiving alert notifications.....	40
Alert overview dashboard.....	40
Alert definitions.....	42
Alert definition details.....	47
Analyzing an alert instance.....	49
Chapter 5. Troubleshooting AIDA	58
Logging and tracing in AIDA	58
Notices.....	lx
Index.....	64

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.

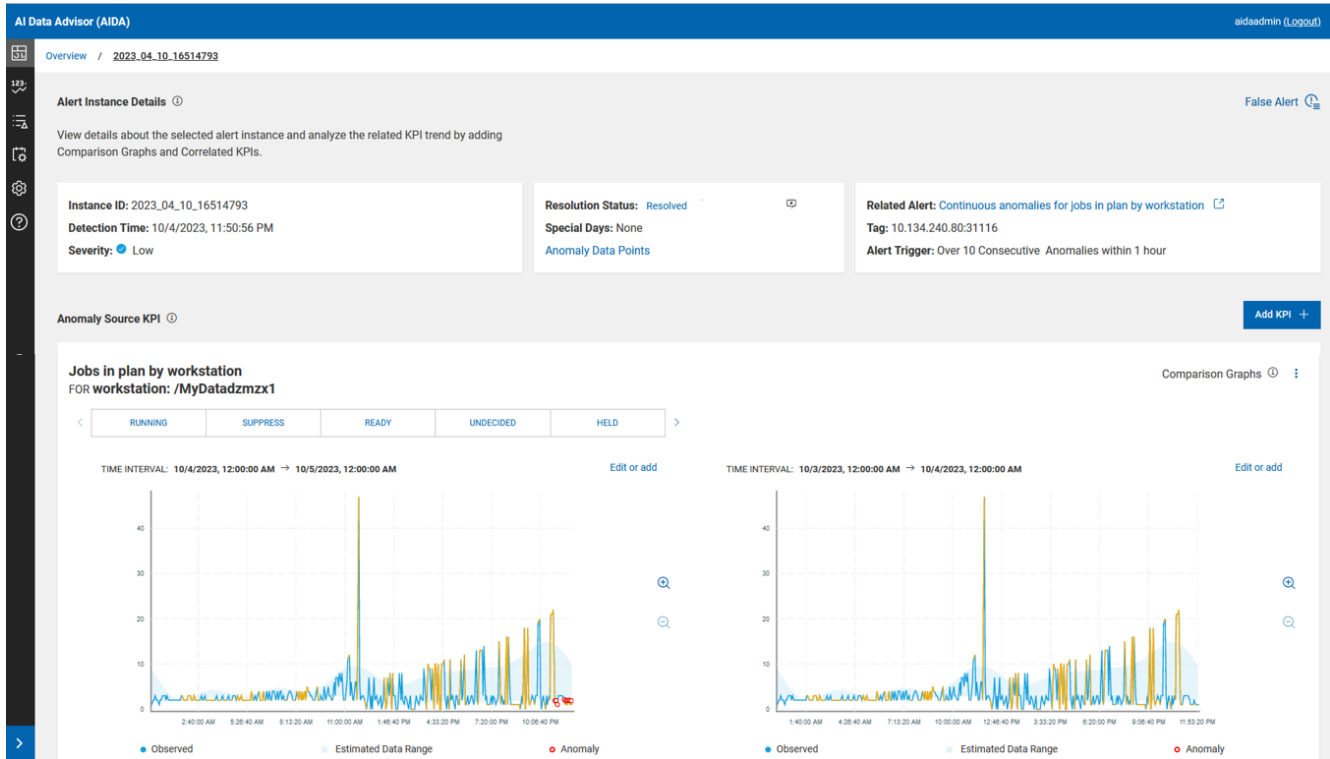
Detecting anomalies in the workload

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.



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.
- Retrain the ML prediction model.

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.

AIDA groups the KPIs in categories and sub-categories.

For more information about IBM Workload Scheduler KPIs managed by AIDA, see: [KPIs for IBM Workload Scheduler on page 27](#).

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](#).

Metrics

The full set of measurements of a KPI.

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

An alert is defined by a set of conditions (see **Alert trigger**). For example: 3 consecutive KPI data points that fall outside the expected range of values.

An alert definition is usually based on one KPI, but multiple KPIs can also be used.

For more information, see: [Alert definitions on page 42](#).

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: [Alert overview dashboard on page 40](#).

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

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: 3 consecutive KPI data points that fall outside the expected range of values.

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

Alerts are notified on the Workload Dashboard or via email.

For more information about receiving alert notifications, see: [Receiving alert notifications on page 40](#).

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 Definition Details

Alert Definition Details provides detailed information about an alert definition, its status, and history.

For more information, see: [Alert definition details on page 47](#).

Alert History

Calendar graph showing previous alert instances and related severity.

For more information, see: [Alert definition details on page 47](#).

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).

Timeframe (or Periodicity)

The time interval whose anomalies must be considered to detect an alert (for example: the last hour, or the last 10 minutes). It is a parameter of each alert definition inside the IBM Workload Scheduler alert definition file.

Special Days

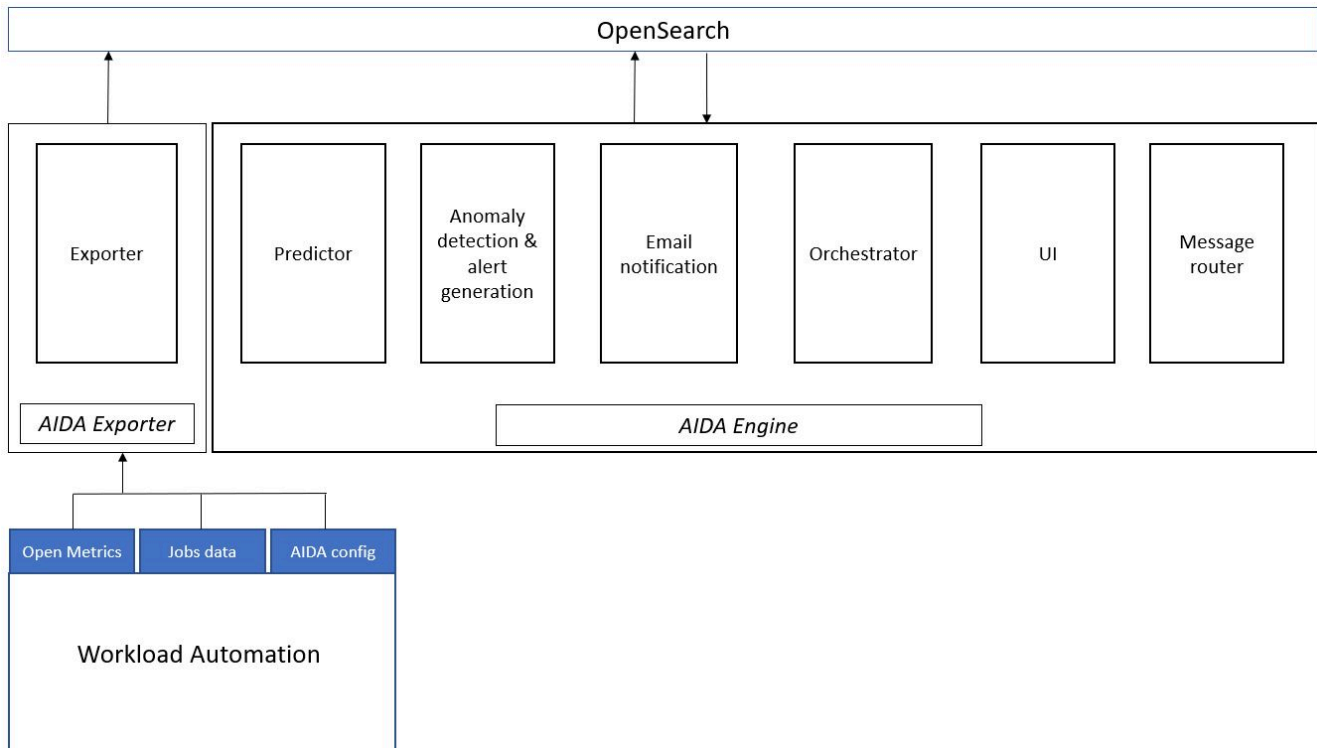
Special days are days on which a KPI trend is affected by seasonality factors such as national holidays, vacation, business cycles, recurring events. To avoid false positive alerts, the special days are included in AIDA prediction model with a higher tolerance level than the standard days.

For more information, see: [Managing special days on page 23](#).

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 (optional)

To manage security and user access in AIDA (Docker deployment only). Keycloak is optional. 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.

Dynamic Workload Console users can access AIDA with a **single sign-on** authentication method: authentication to the Dynamic Workload Console enable users to securely authenticate with AIDA. In AIDA, each user is authorized to the same scheduling objects (jobs, job streams, etc...) he is authorized to 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, retrain KPIs prediction, 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.

From the Dynamic Workload Console, Monitor Job interface

You can also access AIDA UI to analyze the KPIs related to a scheduled job, even without any alert notification:

1. From the Dynamic Workload Console, Monitor Workload > Monitor Jobs interface, right click on the job that you want to analyze.
2. In the drop-down menu, select **Open Data Analysis**. The KPIs Data Analysis page is opened, where you can analyze all the KPIs trends for the job. For details, see [Analyzing KPIs data on page 32](#).

From the Dynamic Workload Console, Workload Designer interface

To analyze the KPIs related to a scheduled job, even without any alert notification:

1. From the Workload Designer of the Dynamic Workload Console, for the job that you want to analyze click **Select an Action**.
2. In the drop-down menu, select **Open Data Analysis**. The KPIs Data Analysis page is opened, where you can analyze all the KPIs trends for the job. For details, see [Analyzing KPIs data on page 32](#).

The same option is available from the Graphical View, right-clicking on a job.



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
```

For AIDA Docker installation only, to manage access to AIDA you can optionally use **Keycloak**.

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 21](#).

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

For Docker deployment, at any time you can add IBM® Workload Scheduler or IBM® Z Workload Scheduler engines to AIDA by using *AIDA.sh* installation script.

Run the script with the *add-credentials* option: a guided configuration procedure will start. Follow the procedure and provide the credentials for the engine that you want to monitor with AIDA.

To update already existing credentials, run the script with the *update-credentials* option.

For details about *AIDA.sh* script, see AIDA readme file for Docker:

- [HCL AI Data Advisor for IBM Workload Scheduler](#)

In AIDA Kubernetes deployment, with every AIDA instance you can monitor just one IBM® Workload Scheduler engine defined in AIDA helm chart. Monitoring of IBM® Z Workload Scheduler engines is not supported. For details, see AIDA readme file for Kubernetes:

- [HCL AI Data Advisor for IBM Workload Scheduler](#)

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 AIDA for KPIs prediction

To optimize KPIs predictions, AIDA administrator can configure tuning parameters.

Before you begin

By default, AIDA uses the **Neural Prophet** Machine Learning (ML) engine. ML algorithms used by AIDA to predict KPIs trend are **additive regression models** based on four main components:

- A **linear** or **logistic** growth trend model. AIDA automatically detects changes in KPIs trend by selecting changepoints from time series.
- A **yearly seasonal** component.
- A **weekly and daily seasonal** component.
- A user-provided list of **holidays and special days**.

Tuning parameters can be set the same way for all KPIs, and AIDA provides default values. However, since tuning parameters adjust the impact of seasonality and special days on prediction, it might be convenient to set them differently for each KPI. After you set tuning parameters, you must retrain the prediction model to recalculate the predictions.



Note: When you start using AIDA, you need to train the prediction model with some data for at least one week before you can get a reliable prediction of the KPIs trend.

About this task

From AIDA left-hand sidebar, select **KPIs** to open the KPIs page. In the **KPIs** page, you, as an AIDA administrator, can select the **Tuning** option for one or more KPIs. The same option is also available:

- From each KPI graph in the **KPIs Data Analysis** page. For detail, see: [Analyzing KPIs data on page 32](#).
- From each KPI graph in the **Alert Instance Details** page. For detail, see: [Analyzing an alert instance on page 49](#).
- From AIDA **Settings** page, where you can change the global default values for all KPIs. For details, see: [AIDA global settings on page 22](#).

By selecting the **Tuning** option, a side panel opens where you can customize some **hyper parameters** to fine-tune the Machine Learning model for KPIs prediction.

The screenshot displays the 'KPI prediction tuning' interface in the AI Data Advisor (AIDA). On the left, a list of KPIs is shown with checkboxes for selection. Two KPIs are selected: 'application_wa_Jobs ByWorkstation_jobs' (Workstation, SAP 1) and 'application_wa_Jobs ByWorkstation_jobs' (Workstation, FTA 1). The right panel shows the tuning options for these selected KPIs. It includes a 'General Tuning' section with a 'Tolerance interval' parameter. The 'Tolerance interval' section explains that increasing or decreasing the tolerance interval for anomaly detection is done by using the 'tolerance_interval_width' parameter, where a smaller value identifies more anomalies. A slider for 'interval_width' is shown with 'Default' and 'Current' values both at 0.8, and an 'Override Default Value' checkbox. Below this is an 'Advanced Tuning' section which is currently collapsed. At the bottom of the tuning panel are four buttons: 'Cancel', 'Restore Defaults Values', 'Save', and 'Apply and retrain'.

Check the option **Override default value** if you want to change default values for the selected KPIs. You can apply your changes immediately by clicking the **Apply and retrain** button: if a retraining process is already in progress, it will be stopped. Alternatively, you can save changes by clicking the **Save** button: in this case, changes will be applied automatically on the next planned retrain. To restore the default settings, click the **Restore default values** button.

Tuning parameters

General tuning

Tolerance interval

Increase or decrease the tolerance interval for anomaly detection by using the parameter: **tolerance_interval_width**.

The smaller it is, the more anomalies can be identified.

Valid range is 0 - 1. Default value is 0.8.

Advanced tuning

Growth trend model

Set the parameter **growth_trend_model** to establish if KPIs prediction should follow a linear or logistic growth trend model.

Valid values: [linear|logistic]. Default value is linear.

Trend flexibility

Adjust the flexibility of trend changes (how the trend changes are being fit) by using the parameter: **trend_flexibility**.

Increasing it will make the trend more flexible.

Valid range is 0 - 1. Default value is 0.05.

Seasonality effect

Adjust the extent to which the seasonality model will fit the data by using the parameter: **seasonality_effect**. A large value allows the seasonality model to fit large fluctuations, a small value shrinks the magnitude of the seasonality effect.

Valid range is 0 - 10. Default value is 10 (which provides very little regularization).

Special days effect

Adjust the extent to which the special days model will fit the data by using the parameter: **special_days_effect**.

A large value allows the special days model to fit large fluctuations, a small value shrinks the magnitude of the special days effect.

Valid range is 0 - 10. Default value is 10 (which provides very little regularization).

Seasonality mode

To get the prediction, the effect of seasonality can be added or not to a KPI trend . Set the parameter:

- **seasonality_mode** = multiplicative

when the seasonality is not a constant additive factor, rather it grows with the trend, so it is not convenient to add its effect to the trend.

- **seasonality_mode** = additive

when the seasonality is a constant additive factor so it is convenient to add its effect to the trend.

Valid values:[multiplicative|additive]. Default value is additive.

About scaling predictions

In this section, you can find some considerations about scaling predictions in AIDA .

About this task

AIDA deployment on Kubernetes enables automatic Pod scalability for KPIs prediction. A new Pod is deployed when the percentage of RAM used for prediction exceeds the 80% of RAM limit.

Special days have little to no impact on performance.

Every prediction uses about 200 MB of RAM. When the prediction is completed, the RAM is released. AIDA optimizes the number of predictions it can handle concurrently.

The time consumed by prediction has a linear growth curve over the number of predictions.

About OpenSearch configuration

In AIDA Docker deployment, OpenSearch is used with a single node. In Kubernetes deployment it can be configured with multiple nodes. The following table can help you configuring OpenSearch with multiple nodes.

The table represents the space required to store daily KPIs data in OpenSearch for the time period defined by the `MAXIMUM_DAYS_OF_OLDER_DATA` parameter configured for AIDA Exporter component.

Number of KPIs	Time period	Indexes
100	1 month (30days)	466 MB
	6 months (180 days)	3051 MB
	12 months (365 days)	6150 MB
500	1 month (30days)	3 GB
	6 months (180 days)	18 GB
	12 months (365 days)	36 GB
1000	1 month (30days)	6 GB

	6 months (180 days)	36 GB
	12 months (365 days)	72 GB

Configuring security

To manage access to AIDA, you can optionally use **Keycloak**.

Before you begin

During AIDA deployment with Docker, you can optionally 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, retrain KPIs prediction, 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](#).

AIDA global settings

AIDA administrator can configure global settings in AIDA.

About this task

From AIDA left-hand sidebar, select **Settings** to configure AIDA global settings.

The screenshot displays the 'Settings' page in the AIDA interface. The page title is 'Settings' with a help icon. Below the title, it says 'Configure global settings in AIDA.' The settings are organized into three main sections:

- KPIs prediction and alert generation:** This section contains two toggle switches. The first is for 'KPIs prediction and alert generation', which is currently active. The second is for 'alert generation only', which is also active. Both sections note that data collection from Workload Automation will not be affected.
- KPIs prediction tuning:** This section includes a button labeled 'KPIs prediction tuning' to open a side panel for adjusting hyperparameters.
- Time Zone:** This section shows the current time zone is 'Europe/Rome - Ora dell'Europa centrale (GMT+1:00)'. It instructs users to change the time zone via their browser settings.

The interface includes a top navigation bar with 'AI DATA ADVISOR (AIDA)' and a user profile 'Tim (Logout)'. A left-hand sidebar contains navigation icons for home, dashboard, settings, and help.

AIDA administrator can configure the following global settings:

KPIs prediction and alert generation

- Click the toggle button to enable or disable KPIs prediction and alert generation. KPIs data collection from IBM Workload Scheduler will not be affected.
- Click the toggle button to enable or disable alert generation only. KPIs data collection and prediction will not be affected.

KPIs prediction tuning

Click the **Tuning** button to open the side panel where you can change default hyperparameters settings for KPIs prediction tuning. For details, see: [Configuring AIDA for KPIs prediction on page 17](#).

Time Zone

To change the current Time Zone, go to your browser settings and select your preferred Time Zone.

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**.

AI DATA ADVISOR (AIDA) Tim (Logout)

Special Days ⓘ

In AIDA prediction model, special days are included with a higher tolerance level than standard days to avoid false positive alerts.

Search an item... 3 Results found Retrain all KPIs Add Special Day +

Name	Date	Repeat	End Repeat	Description	Status	Actions
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	✎ 🗑️

Items per page: 10 | 1 - 10 of 40 items Page 2 of 4 < 2 >

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. However, when you change or add special days, to immediately activate their effect on prediction, you can retrain the prediction model upon request. Run the following steps:

1. Click the **Retrain all KPIs** button and confirm the request. A notification appears on the page header, informing that the retraining process is in progress. The number of KPIs involved is displayed.
2. Click **See details** to view details about the retraining of each single KPI.
3. A page notification informs you about the retraining process completion.

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 datababase.

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 algoritms 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 42](#).

IBM Workload Scheduler KPIs are grouped in the following categories:

Category	KPI name	Metric name	Description	Data frequency
Jobs	Job history	job_history	Duration of each job with the advanced analytics option enabled and all its predecessor jobs.	1 data point per each daily job executions (86400 seconds)
	Total jobs in plan	application_wa_JobsInPlanCount_job_total	The total number of jobs in the current plan.	1 data point every 4 minutes (240 seconds)

	Jobs in plan by status	application_wa_JobsInPlanCount_job	Jobs in the current plan with a specific status. The status can be: WAITING, READY, RUNNING, SUCCESSFUL, ERROR, CANCELED, HELD, UNDECIDED, BLOCKED, and SUPPRESS.	1 data point every 4 minutes (240 seconds)
	Jobs in plan by workstation	application_wa_JobsByWorkstation_jobs	Jobs in the current plan, with a specific status, running on a specific workstation.	1 data point every 4 minutes (240 seconds)
	Jobs in plan by folder	application_wa_JobsByFolder_jobs	Jobs in the current plan, with a specific status, in a specific folder.	1 data point every 4 minutes (240 seconds)
Queue	WA message files fill percentage	application_wa_msgFileFill_percent	Internal message queue usage for Appserverbox.msg, Courier.msg, mirrorbox.msg, Mailbox.msg, Monbox.msgn, Moncmd.msg, auditbox.msg, clbox.msg, planbox.msg, Intercom.msg, pobox messages, and server.msg	1 data point every 4 minutes (240 seconds)

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 retrain the prediction model after changes in the configuration. You can also access the KPIs Data Analysis page to analyze the KPIs trend over time.

About this task

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

AI DATA ADVISOR (AIDA) Retrain in progress
5/57 KPIs retrained [See details](#) Tim (Logout)

KPIs ⓘ

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

Latest planned retrain: 2020/04/14 - 00:00:00

Quick search by name or access advanced filters

<input type="checkbox"/>	KPI Name	Object Type	Object Name	Tag	Anomaly % ⓘ	Last completed retrain	Retrain status		
<input type="checkbox"/>	application_wa_Jobs ByWorkstation_jobs	Workstation	SAP 1	Engine 1	30%	2020/04/15 09:00:03	● Complete	↺	⋮
<input type="checkbox"/>	application_wa_Jobs ByWorkstation_jobs	Workstation	FTA 1	Engine 1	5%	2020/04/15 09:00:03	● Complete	↺	⋮
<input type="checkbox"/>	application_wa_JobsIn PlanCount_job_total	Job	EU-HWS-LNX104-XA#FINAL POSTEPORTCHECKSYNC	Engine 1	20%	2020/04/15 09:00:03	● Complete	↺	⋮
<input type="checkbox"/>	application_wa_JobsIn PlanCount_job_total	Job	EU-HWS-LNX104-XA#FINAL POSTEPORTCHECKSYNC	Engine 1	30%	2020/04/15 09:00:03	● Complete	↺	⋮
<input type="checkbox"/>	application_wa_JobsIn PlanCount_job	Job	EU-HWS-LNX104-XA#FINAL POSTEPORTCHECKSYNC	Engine 2	30%	2020/04/15 09:00:03	○ In progress	↺	⋮
<input type="checkbox"/>	application_wa_JobsIn PlanCount_job	Job	EU-HWS-LNX104-XA#FINAL POSTEPORTCHECKSYNC	Engine 2	30%	2020/04/15 09:00:03	○ In progress	↺	⋮
<input type="checkbox"/>	job_history	Job	EU-HWS-LNX104-XA#FINAL POSTEPORTCHECKSYNC	Engine 2	30%	2020/04/15 09:00:03	○ In progress	↺	⋮
<input type="checkbox"/>	job_history	Job	EU-HWS-LNX104-XA#FINAL POSTEPORTCHECKSYNC	Engine 3	30%	2020/04/15 09:00:03	● Failed	↺	⋮
<input type="checkbox"/>	application_wa_msgFile Fill_percent	Queue	/	Engine 3	30%	2020/04/15 09:00:03	● Failed	↺	⋮
<input type="checkbox"/>	application_wa_msgFile Fill_percent	Queue	/	Engine 3	30%	2020/04/15 09:00:03	● Failed	↺	⋮

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

This page shows you the KPIs managed by AIDA.

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

- Retrain the Machine Learning prediction model for one or more KPIs after changes in the configuration (for example, after adding a special day definition or changing a prediction tuning parameter).



Note: Every 24 hours AIDA runs an automatic retrain of all KPIs. If you want the configuration changes to take immediate effect, you can request a retrain at any time.

- Tune KPIs prediction parameters.
- 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.

The page header shows information about the **latest planned retrain** for all the KPIs (by default, automatic retrains run every 24 hours):

The search bar allows you to search for one or more KPIs. The table displays the following information:

KPI Name

The name of the KPI.

Object Type

The type of the object measured by 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

Last Completed Retrain

Date and time of the last completed retraining process for the KPI.

Retrain Status

Status of the retraining process for the KPI. It can be: **Completed, In Progress, Failed, Not enough data.**

The **Not enough data** status means that there is not enough data available to generate a specific KPI prediction. Depending on the KPI, data needs to be collected from 1 to 7 days before the model can generate a prediction for the KPI.

Select one or more KPIs to run the following actions:

Retrain

To start the prediction model retraining process for the KPIs. KPIs Retrain cannot be scheduled by AIDA users. By default, the retraining process runs automatically on a daily basis (its start time is the start time of the Orchestrator container), but you can start it upon needs (for example, after modifying a special day or a tuning parameter). This action is available to AIDA administrators only.

Tuning

To configure the KPI prediction parameters in the **Tuning** side panel. For example, you might run this action to increase or decrease the tolerance interval for anomaly detection. This action is available to AIDA administrators only. For details, see: [Configuring AIDA for prediction on page 17](#).

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 32](#).

Notes:


KPIs cannot be delete.

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 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 configure the KPI prediction parameters in the Tuning side panel. For example, you might run this action to increase or decrease the tolerance interval for anomaly detection. After tuning, a retrain is needed to immediately apply your changes. Refresh the graph to see the tuning result on the KPI prediction. Tuning is available to AIDA administrators only. For details, see [Configuring AIDA for prediction on page 17](#).
 - **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 36](#).



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.

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 12:00 → Interval duration 1 Days 0:0 Hours End time 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 16](#).
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 .

Alert overview dashboard

In the alert overview dashboard you can view the full list of alert instances and change their status to resolved.

About this task

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

The screenshot shows the AIDA dashboard with the following data:

Instance ID	Alert Description	Severity	Detection Time	Anomaly Source KPI	Tag	Resolution Status
2023_04_10_16514793	Continuous Anomalies 5Sbc-YoBlaDbYk4dLd4n	Low	10/4/2023, 11:50:56 PM		10.134.240.80:31116	Open
2023_04_10_30800633	Continuous Anomalies 6ffc-YoBlaDbYk4dQ2t6	Low	10/4/2023, 11:34:55 PM		10.134.240.80:31116	Open
2023_04_10_17346777	Continuous Anomalies 6ffc-YoBlaDbYk4dQ2t6	Low	10/4/2023, 11:18:56 PM		10.134.240.80:31116	Open
2023_04_10_82673283	Continuous Anomalies 5Sbc-YoBlaDbYk4dLd4n	Low	10/4/2023, 11:18:55 PM		10.134.240.80:31116	Open
2023_04_10_24515603	Continuous Anomalies 7yfc-YoBlaDbYk4dVv11	Low	10/4/2023, 11:02:55 PM		10.134.240.80:31116	Open
2023_04_10_17047241	Continuous Anomalies 5Sbc-YoBlaDbYk4dLd4n	Low	10/4/2023, 11:02:55 PM		10.134.240.80:31116	Open

A summary section contains the following information:

- The number of alerts defined in AIDA.
- The number of alert instances in open status.
- Date and time of the latest retraining process for all KPIs.

You can click the **Retrain all KPIs** button to retrain AIDA prediction model for all KPIs immediately. KPIs Retrain cannot be scheduled by AIDA users.



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 prediction tuning parameters, if you want the configuration changes to take immediate effect, you can request a retraining at any time.

The **Alert Instances** section contains a table with all the alert instances that have been generated. The search bar allows you to search for one or more alert instances. The table displays the following information:

Instance ID

The ID of the alert instance.

Alert Description

The description of the related alert.

Severity

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

Detection Time

Date and time when the alert instance was generated.

Anomaly Source KPI

The KPI that generated the alert instance.

Resolution Status

The alert instance status. Can be: **Open** or **Resolved**. If you want to dismiss an alert instance because you already resolved it, select the status **Resolved** in the instance row. If you want to resume an alert instance, select the status **Open**.

To run a detailed analysis on an alert instance, click on the **Instance ID** in the alert instance row. For details, see [Analyzing an alert instance on page 49](#).

An action icon is available for each alert instance to **Add a Comment** about the alert instance and its resolution. Click **Save** to save your comments.

To view only the resolved instances, click the **Resolved** button on top of the Alert Instances table. To view only the open instances, click the **Open** button.



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 the following parameters :

Trigger

A custom set of conditions, for example: 10 anomalous data points falling outside the expected range of values.

Timeframe (or Periodicity)

The time interval whose anomalies must be considered to detect an alert (for example: the last hour, or the last 10 minutes).

By default, two alert definitions are provided for each KPI:

- ALERT_TOTAL, triggered by 3 (or more) anomalous data points **above and below** the predicted range, in 1 hour (for the job_history KPI, timeframe is 2 days).
- ALERT CONTINUOUS, triggered by 10 (or more) anomalous data points **all above or below** the predicted range, in 1 hour (for the job_history KPI, timeframe is 2 days).

By default, the ALERT_TOTAL type alerts are paused, while the ALERT_CONTINUOUS type alerts are active.

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"
  }
]

```

About this task

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

Alert ID	Alert Description	Anomaly Source KPI	Alert Trigger	Status
<input type="checkbox"/> CONTINUOUS_JOB_FOLDER	Continuous Anomalies Jobs in plan by Folder	application_wa_jobsByFolder_jobs	Over 10 Consecutive Anomalies within 1 hour	● Active 00
<input type="checkbox"/> CONTINUOUS_JOB_HISTORY	Continuous Anomalies for job history	job_history	Over 10 Consecutive Anomalies within 2 days	● Active 00
<input type="checkbox"/> CONTINUOUS_JOB_STATUS	Continuous Anomalies for Jobs in plan by status	application_wa_jobsInPlanCount_job	Over 10 Consecutive Anomalies within 1 hour	● Active 00
<input type="checkbox"/> CONTINUOUS_JOB_TOTAL	Continuous Anomalies for total jobs	application_wa_jobsInPlanCount_job_total	Over 10 Consecutive Anomalies within 1 hour	● Active 00
<input type="checkbox"/> CONTINUOUS_JOB_WKS	Continuous anomalies for jobs in plan by workstation	application_wa_jobsByWorkstation_jobs	Over 10 Consecutive Anomalies within 1 hour	● Active 00
<input type="checkbox"/> CONTINUOUS_MESSAGE	Continuous Anomalies for message files fill percentile	application_wa_msgFileFill_percent	Over 10 Consecutive Anomalies within 1 hour	● Active 00
<input type="checkbox"/> TOTAL_JOB_FOLDER	Total Anomalies for Jobs in plan by Folder	application_wa_jobsByFolder_jobs	Over 3 Anomalies within 1 hour	○ Paused ▷
<input type="checkbox"/> TOTAL_JOB_HISTORY	Total Anomalies for job history	job_history	Over 3 Anomalies within 2 days	○ Paused ▷
<input type="checkbox"/> TOTAL_JOB_STATUS	Total Anomalies for jobs in plan by status	application_wa_jobsInPlanCount_job	Over 3 Anomalies within 1 hour	○ Paused ▷
<input type="checkbox"/> TOTAL_JOB_TOTAL	Total Anomalies for total jobs	application_wa_jobsInPlanCount_job_total	Over 3 Anomalies within 1 hour	○ Paused ▷
<input type="checkbox"/> TOTAL_JOB_WKS	Total Anomalies for Jobs in plan by Workstation	application_wa_jobsByWorkstation_jobs	Over 3 Anomalies within 1 hour	○ Paused ▷
<input type="checkbox"/> TOTAL_MESSAGE	Total Anomalies for message files fill percentile	application_wa_msgFileFill_percent	Over 3 Anomalies within 1 hour	○ Paused ▷

In this page you can view the full list of alert definitions in table format. The search bar allows you to search for one or more alert definitions.

The table displays the following information:

Alert ID

The ID of the alert.

Alert Description

The description of the alert.

Anomaly Source KPI

The KPIs that contribute to generate the alert (can be more than one KPI).

Alert Trigger

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

Status

The status of the alert. It can be *Active* or *Paused*.

Click on an alert definition to open the **Alert Definition Details** where you can get detailed information about the alert, its definition and history. For details, see [Alert definition details on page 47](#).

An action icon is available for each alert definition to **Pause** or **Resume** the alert generation. For details about pausing an alert, see [Pause alert generation on page 49](#).

By selecting multiple alert definitions, the following bulk action icon appears in the table header:

Pause now/Resume

To pause or resume the generation of the selected alerts. Only the **Pause now** action is available for multiple alerts.

Alert definition details

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

Before you begin

The **Alert Definition Details** page provides you with information about an alert definition, its status, and history. You can also **Pause** and **Resume** an alert generation. This page is composed of three sections:

- Alert Definition
- Alert Instances
- Last 12 months history

The screenshot displays the 'Alert Definition Details' page in the AIDA interface. The top navigation bar shows 'AI Data Advisor (AIDA)' and the user 'aldadmin (Logout)'. The breadcrumb trail indicates the current page is 'Alert definition / Continuous Anomalies for jobs in plan by workstation'.

The main content area is divided into three sections:

- Alert Definition Details:** This section provides instructions: 'View the selected alert definition. Pause the alert. Analyze the alert instances.' It shows the 'Alert Description' as 'Continuous anomalies for jobs in plan by workstation' with a status of 'Active'. A 'Pause' button is visible. The 'Alert Trigger' is 'Over 10 Consecutive Anomalies within 1 hour' and the 'Anomaly Source KPI' is 'application_wa_jobsByWorkstation_jobs'.
- Alert Instances:** This section prompts the user to 'Select the alert instance that you want to analyze.' It includes a search bar with the placeholder 'Search for instance id or open advanced search' and 'Open' and 'Resolved' buttons.
- Table of Alert Instances:** A table with columns for Instance ID, Severity, Detection Time, Tag, and Resolution Status. It lists five instances, all with a severity of 'Low' and a resolution status of 'Open'.

Instance ID	Severity	Detection Time	Tag	Resolution Status
2023_04_10_16514793	Low	10/4/2023, 11:50:56 PM	10.134.240.80:31116	Open
2023_04_10_30800633	Low	10/4/2023, 11:34:55 PM	10.134.240.80:31116	Open
2023_04_10_17346777	Low	10/4/2023, 11:18:56 PM	10.134.240.80:31116	Open
2023_04_10_82673283	Low	10/4/2023, 11:18:55 PM	10.134.240.80:31116	Open
2023_04_10_24515603	Low	10/4/2023, 11:02:55 PM	10.134.240.80:31116	Open

Alert Definition

About this task

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

Alert Description

The description of the alert.

Anomaly Source KPI

The KPIs that generated the alert instance (can be more than one KPI).

Alert Trigger

Set of conditions defining the alert.

Status

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

You can pause or resume an alert generation, see [Pause alert generation on page 49](#).

Alert Instances

About this task

The **Alert Instances** section contains a table with all the alert instances that have been generated. The search bar allows you to search for one or more alert instances. The table displays the following information:

Instance ID

The ID of the alert instance.

Severity

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

Detection Time

Date and time when the alert instance was generated.

Tag

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

Resolution Status

The alert instance status. Can be: **Open** or **Resolved**. If you want to dismiss an alert instance because you already resolved it, select the status **Resolved** in the instance row. If you want to resume an alert instance, select the status **Open**.

To run a detailed analysis on an alert instance, click on the **Instance ID** in the alert instance row. For details, see [Analyzing an alert instance on page 49](#).

An action icon is available for each alert instance to **Add a Comment** about the alert instance and its resolution. Click **Save** to save your comments.

To view only the resolved instances, click the **Resolved** button on top of the Alert Instances table. To view only the open instances, click the **Open** button.



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).

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.

To stop all alerts generation, from AIDA left hand sidebar, select **Settings** and switch the alert generation toggle on off.

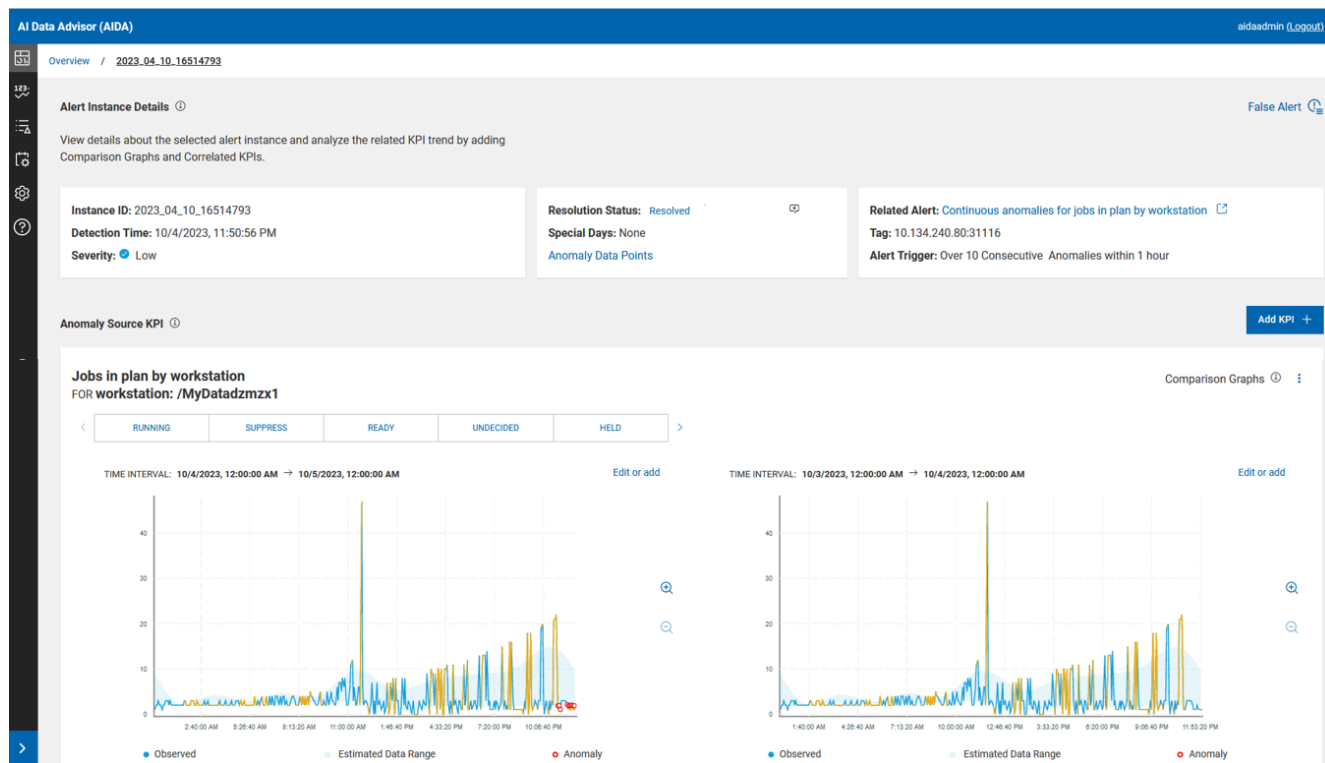
Analyzing an alert instance

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 configure the KPI prediction parameters in the Tuning side panel. For example, you might run this action to increase or decrease the tolerance interval for anomaly detection. After tuning, a retrain is needed to immediately apply your changes. Refresh the graph to see the tuning result on the KPI prediction. Tuning is available to AIDA administrators only. For details, see [Configuring AIDA for prediction on page 17](#).
 - **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 52](#).

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 5. Troubleshooting AIDA

See how to troubleshoot problems in AIDA.

This section describes:

- How to collect logs and activate traces 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

Notices

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

© Copyright IBM Corporation 1993, 2016 / © Copyright HCL Technologies Limited 2016, 2023

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) (2023).

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

© Copyright IBM Corp. 2023

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 52
- Add a special day 23
- Add correlated KPIs for comparison 53
- Add country specific holidays 25
- Adding a country specific holiday 25
- adding a custom date 26
- adding a custom special day 26
- Adding comparison graphs 52
- Adding correlated KPIs 53
- Adding country specific holidays 25
- adding custom special days 26
- adding engines to AIDA 16
- Adding special days 23
- Administrator
 - AIDA
 - 16
 - AI 6, 7, 7
 - AI Data Advisor
 - 47, 49
 - troubleshooting 58
 - AIDA
 - 6, 7, 7, 40, 47, 49
 - administration 16
 - configuration 16
 - troubleshooting 58
 - AIDA architecture 11
 - AIDA
 - basic concepts
 - 9
 - AIDA
 - entry points
 - 13
 - AIDA global settings 22
 - AIDA
 - installation and configuration
 - 13
 - AIDA security 21
 - AIDA security settings 21
 - AIDA settings 22
 - AIDA
 - terminology
 - 9
 - alert 47, 49
 - definition 9
 - Alert definition 47
 - alert definitions 42
 - Alert history 48, 49
 - alert information 47
 - Alert Information 47
 - Alert instances 49
 - alert notification 40
 - alert notifications 40
 - alert overview dashboard 40
 - alerts 40
 - all alert instances 40
 - Analyzing an alert instance 49
 - analyzing KPIs 27, 27
 - Analyzing KPIs 32
 - anomalies 40
 - anomaly 40
 - definition 9
 - anomaly alert definitions 42

- anomaly alerts 42
- anomaly detection 6, 7, 7, 40
- Artificial Intelligence 6, 7, 7
- automatic prioritization 6, 7, 7

B

- business scenario 7, 7

C

- Cloud & Smarter Infrastructure technical training v
- Configuring
 - AIDA
 - 16
 - configuring AIDA for KPIs prediction 17
 - configuring email settings 16
 - configuring KPIs prediction in AIDA 17
- Correlated KPI graphs 53
- current and historical alert information 47

D

- data-driven decisions 6, 7, 7
- datapicker panel 36, 54
- defining anomaly alerts 42
- Dynamic Workload Console
 - accessibility v

E

- education v

H

- how to access
 - AIDA
 - 13
 - how to add engines to AIDA 16
 - How to analyze KPIs in AIDA
 - 32
 - how to deploy
 - AIDA
 - 13
 - How to manage KPIs in AIDA
 - 29
 - how to set time intervals 36, 54

K

- Key Performance Indicators 27, 27
- KPI prediction 17

L

- Logging and tracing in AIDA
 - 58
- Logs 58

M

- Machine Learning 6, 7, 7
- mail notifications 40
- managing alerts 40
- managing KPIs 27, 27
- Managing KPIs 29
- microservices 11
- microservices-based architecture 11

P

- pause alert 49
- pause alert generation 49
- prediction configuration 17
- proactive SLA management 6, 7, 7

- problem prevention 40

R

- receiving alert notifications 40
- retraining 23

S

- seasonality 23
- Setting time intervals 36, 54
- Special Days 23
- special days in KPI trend 23

T

- technical training v
- time intervals 36, 54
- traces 58
- training
 - technical v
- trigger
 - definition 9
- troubleshooting 58
- tuning KPIs prediction 17
- tuning parameters 17

V

- view all alert instances 40
- view all alerts 40

W

- Working with alerts 40
- Working with KPIs 27