# Kofax eFlow
**Supervise User's Guide**

Version: 6.1.0

Date: 2023-08-29

**KOFAX**

# Table of Contents

# Preface

This guide describes about Supervise module. Supervise is eFlow runtime monitoring tool. It helps managers to plan and monitor jobs, enables smart resource allocation according to availability and performance, and provides time-to-completion estimations. See the Tutorials for examples.

Different views provide immediate visual information on the status of all elements of the production environment.

## Product documentation

To access the full Kofax eFlow documentation set online, see the Kofax eFlow Product Documentation page. For a complete set of Kofax eFlow documents, refer to the Kofax eFlow Release Notes.

# About Supervise

This chapter describes how to open the Supervise tool, terms and definitions used, settings, SLA configuration, alert configuration, and user definitions applicable for the Supervise tool.

## Terms and definitions

The following table describes the terms used in Supervise.

| Term | Definition |
|------|-----------|
| SLA | A Service Level Agreement: It is the part of a service contract where the level of service is formally defined. It is a useful method for setting expectations on deliverability. eFlow uses the term SLA to refer to the contracted delivery time. |
| TTC | Time to complete: The estimated time left until the task will be completed. |
| PPM | Pages per minute: The unit of measuring the performance of the automatic stations. |
| CPM | Collections per minute: The unit of measuring station performance. Relevant for all workflow stations. |
| Operator | One of the manual stations operators listed in the Users table of the eFlow management database. All operators should also be defined as valid STS users. |
| Team | A group of operators working under the same configuration. This term is used in the Supervise station for presentation purposes only. |
| Alert | A system notification issued when a potentially problematic situation occurs. |
| Worker server | An eFlow client that runs the automatic stations. In modern environments, automatic stations usually run on powerful multi-CPU server stations, hence the name "Worker server" for this type of client. |

## Supervise window

To open **Supervise**, in the **eFlow Launch Pro**, from the **Reports** list, select **Supervise**.

The **Supervise** window appears.
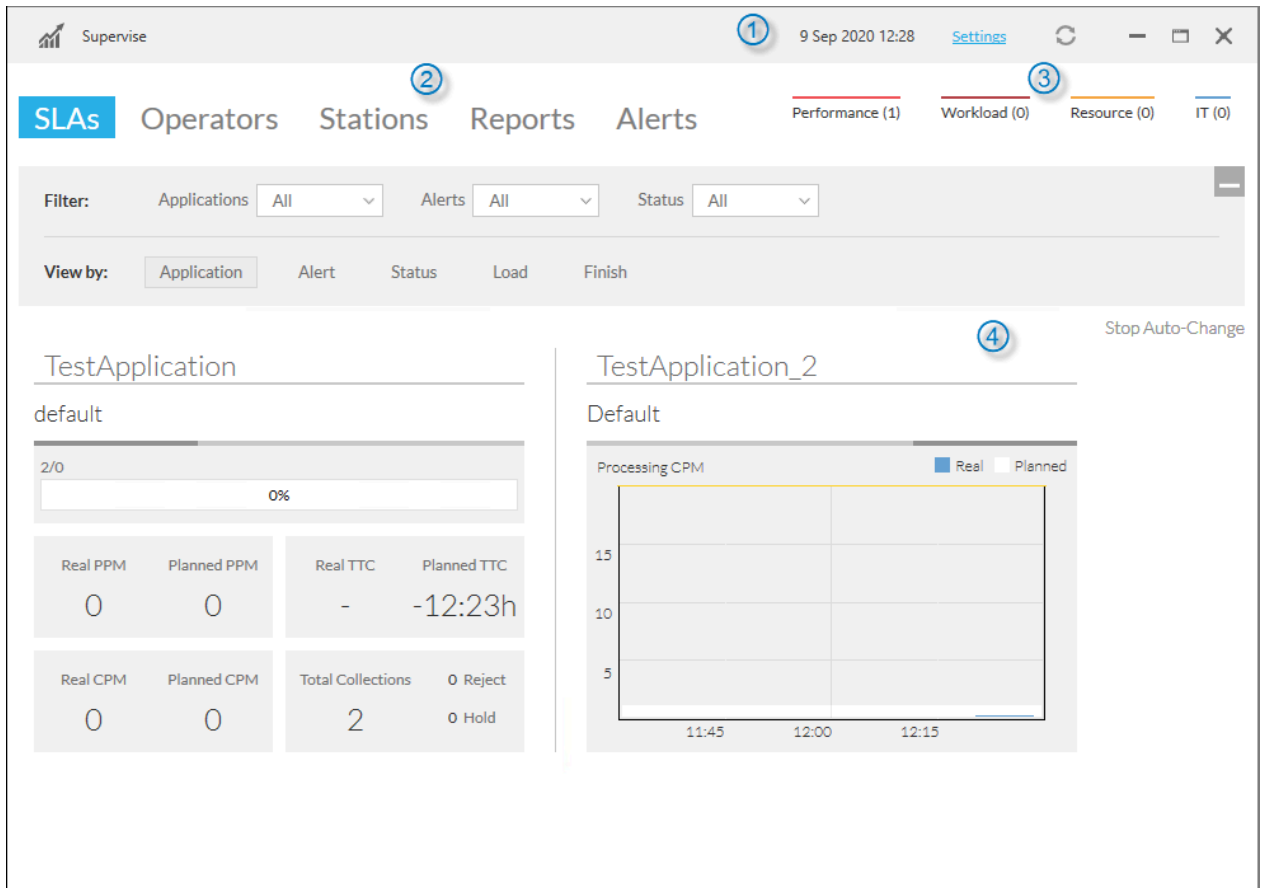
The **Supervise** window includes the following sections.

| 1 | Header bar | Shows the current date and time. |
|---|---|---|
| | | **Settings**: Opens the Settings dialog box where you can configure the settings for the Supervise report. See Supervise Settings for more information. |
| | | ⟳: Refreshes the page. |

| 2 | Views | This section provides access to five different views. <br><br> • **SLAs**: This view displays information about performance against SLAs. SLAs define how fast collections are expected to pass through the workflow. <br><br> • **Operators**: This view displays information about the activity of the eFlow manual stations. <br><br> • **Stations**: This view displays information about the activity of the eFlow automatic stations. <br><br> • **Reports**: This view displays statistical reports as graphs. <br><br> • **Alerts**: This view displays system alerts. Alerts warn you about currently existing or potential problems with the system. <br><br> The views are refreshed automatically once per minute to display the latest information. To manually refresh a view, click ⟳ at the top-right of the window. <br><br> To open a view click on it. <br><br> Each view provides filters and grouping options (View by) to enable you to focus on specific information. The filters and options available depend on the selected view. To hide the filter area, click ▬ at the right of the filter section. See Views for more information. |
|---|---|---|
| 3 | Alerts | Alerts warn you about currently existing or potential problems within the system. See Alert configuration for more information. |
| 4 | Data display area | This section displays the data from different reports. |

# Supervise settings

This chapter describes the configuration of Supervise station settings.

## Manage SLAs

You can create an SLA, edit existing SLAs, change assignments for an SLA, and delete an SLA.

### Create an SLA

You can create an SLA, set its color, and define rules.

1. On the **Settings** page, click **New**.

2. Configure the properties.

| Property | Description |
|---|---|
| SLA Name | The user defined name of the SLA. The SLA name is mandatory and cannot be changed. |
| Application | The name of the application to which the SLA belongs. It cannot be changed after the SLA is created. SLAs are always defined per application. |
| SLA Type | The type of SLA:<br><br>• **Fixed**: A fixed SLA always begins and ends on the same day. You must set the Start and Finish time for the SLA. You can set the Recurrence to Daily or Weekly. If the recurrence is set to weekly, the SLA is executed on selected days. You cannot define the SLA that starts on Sunday and ends on Monday.<br><br>• **Time to complete**: The time that takes one collection to finish the SLA. It is defined in seconds, minutes, hours, or days. |

| Property | Description |
| --- | --- |
| CPM | The planned amount of collections to be processed per minute.<br><br>ⓘ We recommend defining this property though it is not mandatory. |
| PPM | The planned number of pages to be processed per minute. It is relevant for automatic stations.<br><br>ⓘ We recommend defining this property though it is not mandatory. |
| Quantity | Optional. The planned total amount of collections to be processed while SLA is active. |
| Export queue | Defines the last workflow station for the SLA. The list of possible values includes all the workflow stations of the current application and the generic constant Out.<br>Out means that SLA must finish when the workflow ends.<br>If the current SLA is finished and the next one is not yet started, the batch remains connected to its last SLA in the database even though the SLA is no longer working. |
| Color Settings | Optional. Select one of the predefined colors for SLA. Setting colors help in distinguishing between different SLAs. |
| Rules | SLA rules help to determine to which SLA the collection belongs. See SLA rules for more information. |

3. Click **OK**.

# SLA rules

Every eFlow collection must be attached to an SLA. SLA rules help to determine to which SLA the collection should belong.

The SLA rules are defined per flow in the selected application. They can be based on any of the application metatags defined in the Design tool.

## How SLA rules are checked

The SLA rules check is done on the server as follows:

1. The collection looks for a rule that is relevant for it.

2. If no relevant rules are found, the collection looks for an SLA without rules.

3. If an SLA without rules is not found, the collection is sent to the Pending queue.

There should be no overlapping of rules between different SLAs; if this happens, the order of execution of the SLAs cannot be guaranteed.

An SLA can contain more than one rule. To attach the collection to a specific SLA, all the rule checks should be successful. Schematically, the following condition should be met.

(RuleA==true) AND (RuleB==true) AND (RuleC==true)

All SLA rules are checked on every Put collection on every station. If an SLA ends at a specific station and the workflow continues, you should define the rules so that they fail on the next SLAs.

## Define SLA rules

You can add, edit, and remove an SLA rule.

In the **Rules** section, from the **Flow** list, select the flow for which to set the SLA rules.

An SLA rule comprises of:

- **Parameters**

  These parameters are based on the flow metatags defined in the Design tool.
- **Operator**

  The SLA rules support the following comparison operator types:
  - Logical operators:
    - \>
    - <
    - =
    - <>
    - \>=
    - <=
  - String functions:
    - Contained
    - StartWith
    - EndWith
  - RegEx

    > ⓘ All rule comparisons are case sensitive.

- **Value**

  The value should correspond to the metatag type. The values of the DateTime type parameters (such as, CreationTime) should use the system date format of the eFlow server machine.
- **Action**

  1. Click **Add**.
  2. Select an operator for the parameter.
  3. Type a value for the parameter.
  4. To save the rule, click **Accept** or to decline rule, click **Cancel**.

## Edit an SLA

You can edit the settings for an existing SLA and change its assignments.

## Edit SLA settings

1. On the **Settings** page, select the SLA to edit and click **Edit**.



2. Make changes to the properties as needed.
3. Click **OK**.

## Change assignments

You can change only one assignment at a time and other assignments are updated accordingly.

1. On the **Settings** page, select the SLA for which to edit assignment and click **Assignment**.

2. Make changes as needed.

3. Click **OK** to save the change.

You can undo the change by clicking **Reset**. Resetting returns the current SLA assignment to its previous state.

## Delete an SLA

1. To delete an SLA, select the SLA and click **Delete**.

2. Click **Ok**.

# Alerts configuration

You can either keep the default alert configuration or configure specific settings.

## Default

The default alert configuration is applied to every alert for which no specific configuration is defined.

Alerts can be configured to send e-mail notifications. To send the e-mail notification to several recipients, separate the addresses with a semicolon.

The default alert configuration uses the system SMPT settings.

Select **Enable SSL** per alert if required.

## Specific settings

For any alert, you can create a configuration that differs from the default one.

Email recipients defined per alert are added to those set in the default configuration.

The alert-specific email SMTP settings overwrite the default configuration.

It is possible to configure alerts without activating them. Alert activation can be done later when needed.

## Alert types

The following table summarizes the alert types supported by eFlow:

| Alert type | Description | Alert resolution |
|---|---|---|
| CPU | CPU usage exceeds 90% on one of the station servers (clients on which the eFlow Autorun Stations Starter service is running). This alert does not include the problematic machine name. It is issued on the system level, that is, for all applications and all stations. If the problem occurs on several servers, still only one alert will appear. | This alert is removed automatically if the CPU falls below 90% on all the station servers. |
| Reject | There are batches in the Reject queue. This alert is issued per application. | This alert is removed automatically if the Reject queue is empty. If you close this alert manually, it appears again only if new batches are sent to the Reject queue. |
| Behind schedule | There are batches in the system that will not be finished in time. | This alert is removed automatically if the situation improves. It makes no sense to close it manually, as it appears again if the problem was not solved. |
| License expired | This alert is issued a week before the license expiration date. | This alert is removed automatically if the new license is installed. It can be closed manually by the project manager. |

| Alert type | Description | Alert resolution |
|---|---|---|
| Queue empty | The specific queue in the application is empty for the last 10 minutes and there is a station that is logged in and was working before. If the queue suddenly becomes empty, this may be a sign that the previous station does not produce enough output. | This alert is removed automatically when batches appear in the queue or the station logged out. It makes no sense to close it manually, as it appears again if the problem was not solved. |
| Queue load | The average input for the specific queue is 10% bigger than output for the last 10 minutes and there is a station in this queue that is currently logged in. This alert can be issued either per SLA or for all SLAs in the application. Depending on the alert information, the administrator can take the relevant measures to resolve the situation. | This alert is removed automatically when the input and output volumes of the station do not differ. It makes no sense to close it manually, as it appears again if the problem was not solved. |

# User definitions

This section describes how to set users and groups definitions in Supervise.

1. On the **Supervise** window, click **Settings**. The **Settings** page appears.
2. In the **General** section, click **Users**. The **Users & Groups** dialog box is displayed.
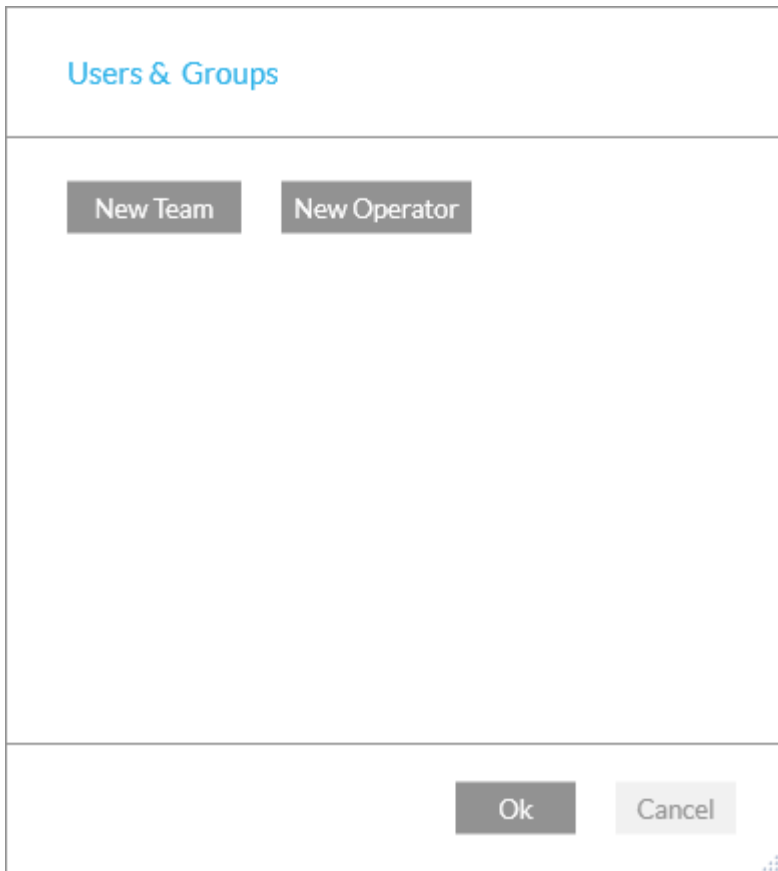
3. Create a new team.

   a. Click **New Team**. A default team gets created.

   b. To replace the default team's name with a meaningful one, select the default team, click **Rename** and type the name.

   > ℹ️ Teams defined in the Users & Groups dialog box are valid for Supervise only.

4. Add some user definitions to the team.

   a. Select the team and click **New Operator**.

      > ℹ️ Every new manual station operator should belong to one of the Supervise teams. A user cannot be included in more than one team.

   b. Provide a name to the operator.

      The operator names should match those defined in the current STS system. For example, if Windows authentication is selected, the names should include the domain name. A user that is not defined in the STS will not be able to log in.

5.  Click **OK**. The operators and groups are saved in the User table in the eFlow_Management SQL database.

> ⛔ Currently there is no automatic utility that synchronizes between the system security settings (STS) and the Users table. The eFlow system administrator should make sure that users and teams are properly defined.

# SLA assignment

You can configure each application to keep the required balance between the amount of collections bound to different SLAs.

## Inherited assignment

By default, all SLA assignments are inherited from the upper level assignments. The assignments hierarchy allows to achieve the maximal flexibility in routing the batches. You can override the inherited definitions on every level.

The assignments are set using the following hierarchy (starting from the top level).

1.  Instance
2.  Operator teams
3.  Queues
4.  Application

For example, assignments set manually per operator team, will override those set by queue (the lower hierarchy level).

The default application level assignments are all set to zero. It means that collections will be taken using the simple "FIFO" (first-in-first-out) algorithm.

## Manual assignment

You can set SLA assignments manually. The manual assignment always overwrites the inherited assignments.

Changing the assignment for one SLA in the application influences all the others. The sum of all the assignments is always equal to 100%. The assignment status is calculated every time when the Get collection operation is performed.

For example, if each of the two SLAs are set to 50%, but all the new batches belong to the first SLA, they will be taken by the system one after another as long as there are no new batches belonging to the second SLA. When a batch bound to the second SLA appears, it is taken immediately, as the system tries to keep the required proportion of 50-50.

1.  On the **Supervise** window, click **Settings**. The **Settings** page appears.

2. Under **General**, click **SLA Assignments**. The **Global SLA Assignments** dialog box is displayed. It displays all the system SLA assignments ordered by application name.

3. Select one of the applications in the tree to see its SLA assignments.



4. Select either option:
   - **Manual Assignment**: Use the slider control to set SLA assignments manually.
   - **Inherited Assignment** (default): Allow all SLA assignments to be inherited from the upper level assignments.

5. Click **OK**.

# Chapter 3

# Views

The Supervise main window provides access to five different views. Click on a view name to open that view.

- SLAs: This view displays information about performance against SLAs. SLAs define how fast collections are expected to pass through the workflow.
- Operators: This view displays information about the activity of the eFlow manual stations.
- Stations: This view displays information about the activity of the eFlow automatic stations.
- Reports: This view displays statistical reports as graphs.
- Alerts: This view displays system alerts. Alerts warn you about currently existing or potential problems with the system.

The views are refreshed automatically once per minute to display the latest information. To manually refresh a view, click ⟳ at the top right of the window.

Each view provides filters and grouping options (View by) to enable you to focus on specific information.

Which filters and options are available depends on the selected view.

## SLAs view

The **SLAs** view is the central part of the Supervise station.

The eFlow SLA is a contract that defines how fast should the collection go through the workflow. This contract is always defined per specific application.

Usually an application has multiple SLAs attached to different types of collections. Each collection is bound to some SLA in accordance with the SLA rules.

### View filters

The following filters are available.

| Filter | Description |
|---|---|
| Applications | Shows only the SLAs that belong to the selected application. If **All** is selected, the applications filter is not applied. |
| Alerts | You can select to view only those SLAs for which active **Performance** alerts exist. If **All** is selected, the alerts filter is not applied. |

| Filter | Description |
|---|---|
| Status | You can select to view only those SLAs that are behind schedule. If **All** is selected, the status filter is not applied. |

The following grouping options are available.

| View by | Description |
|---|---|
| Application | Groups the SLA cards by application name. |
| Alert | Groups the SLA cards by active alerts. |
| Status | Groups the SLA cards by their status. |
| Load | Separately shows the **Empty** SLAs that have no collections and the **Loaded** SLAs that have some collections attached. |
| Finish | Separately shows SLAs that are **On time** and **Late**. SLAs that have no **Finish** time defined are shown as **Unknown**. |

ⓘ A filter set for one view will be shared by all views.

## SLA card view

The SLA card consists of three tabs representing different views. To switch between the tabs, either click on the upper border of the card or enable the Auto-change option.
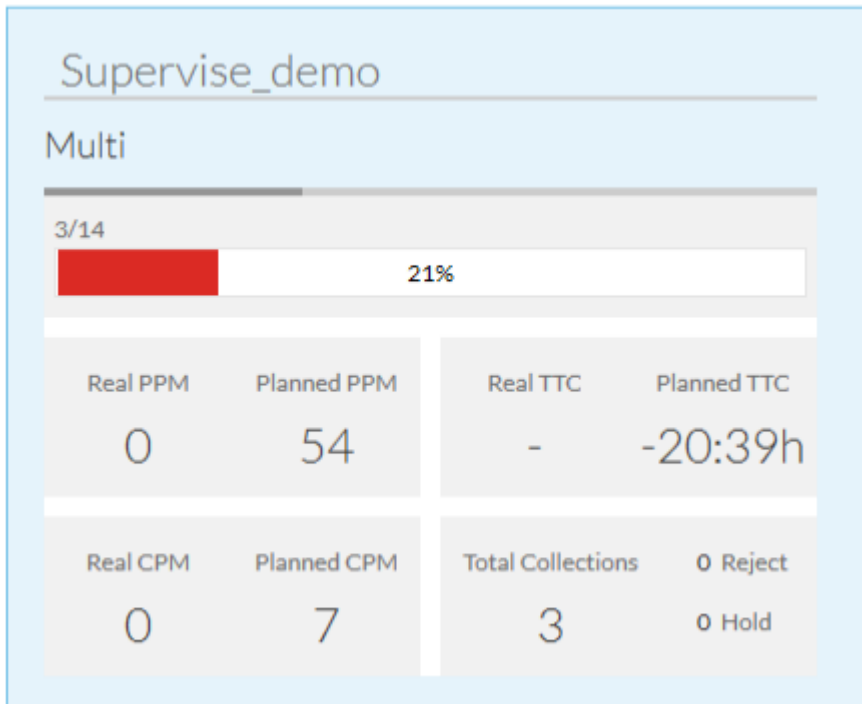


The line between the SLA name and the card serves as the alert indicator. If any performance alerts are issued, the line's color changes.

> ℹ️ It is not possible to stop or enable Auto-change for one SLA only. It is the global setting for the whole view.

## Global Performance view

The Global Performance view shows current and planned values of all the performance parameters.

The numbers above the progress bar in the upper part of this view inform about the real and planned amount of collections. The real quantity includes both collections that have already left the system and those still in the workflow.

The progress bar shows the percentage of the work done compared to the planned quantity. All collections that are on time or have already finished, are shown in green. Red color indicates that collections are behind schedule.

If the planned quantity is not defined, the progress bar is empty.

The controls in the view include the following information.

| Property | Description | Values | Notes |
|---|---|---|---|
| CPM (Collections Per Minute) | The unit of measuring the station performance | Real Planned | Relevant for all workflow stations. |
| PPM (Pages Per Minute) | The unit of measuring the station performance | Real Planned | Relevant for automatic stations only. |
| TTC (Time to complete) | The estimated time left until the task will be done | Real Planned | TTC can be negative for batches that are in the system for long time. |
| Collections | Number of collections | Total number In Rejected queue On hold | |

## Utilitization view

The Utilization view is a burndown graph showing how much work is left now. It always shows the work progress during the last hour.

If **Start** and **Finished** times are defined, the graph always starts at the maximal value and goes down until there are no collections left in this SLA. The comparison between planned and real graphs shows immediately if there are any problems.

For the SLA of the **TTC** type, the planned performance is represented by the horizontal line, meaning that every collection executes this SLA for the same amount time. If problems occur, the real line appears to be above the planned, meaning that these collections are moving slowly.

## Processing CPM view

The Processing CPM view shows the CPM values for the specific SLA during the last hour.

# Operators view

The Operators view provides information about the activity of eFlow manual stations. Manual stations resource management is never a simple task; this view makes it easier by presenting many different parameters.

## View parameters

The **Operators** tree view can be organized in different ways depending on the **View by** selection. The selected **View by** element determines the root of the tree view, and every row represents one of the logged in stations. By default, the view is grouped by applications.

The **Operators** grid view contains the following columns.

| Column name | Description |
| --- | --- |
| Name | Name of the operator of the currently logged in station. There can be several lines containing the same user name and the view can change depending on the filter. The users shown in the view must be defined in the dialog box. |
| Application | Name of the application currently logged in. |
| CPM | Number of collections completed by the operator per minute during the working time. The working time is the total time that the station was logged in, including idle time, such as operator breaks. |
| PPM | Number of pages completed by the operator per minute during the working time. The working time is the total time that the station was logged in, including idle time, such as operator breaks. |
| Usage | The ratio between the net working time and total working time. |
| Queue | The name of the queue (the logged in station name). |
| Load | The number of collections that are waiting in the queue for this station instance. This value is only an estimation since it does not consider the operators' performance. For example, if there are 1000 collections in the queue and there are 2 instances of the same station, the Load column value for each instance will be 500.<br><br>The calculation of the number of collections in the queue takes the SLA assignments into consideration. Thus, the numbers in the example above will not be the same if SLA assignments for the station instances differ. |
| Load Rate | Indicates how fast the queue grows or becomes smaller. It becomes red if the input exceeds the output by 25% during the last ten minutes, green if output exceeds the input, and remains neutral if the input is approximately equal to the output. |
| Machine | Name of the machine on which station instance is running. |
| Team | User's team name. |
| SLA assignment | Shows graphically the actual SLA assignment, either inherited or set manually. The SLA assignment level (Instance, Application, Queue, Team) depends on the View By selection. If the SLA assignment is FIFO (all assignment values are equal to zero), this column will be empty. |

## View filters

The following filters are available.

| Filter | Description |
| --- | --- |
| Applications | Logged in application name. |
| Team | Team name. |
| SLA name | You will not see the operators that have no collections belonging to this SLA in their real assignments. |

The following grouping options are available.

| View by | Description |
| --- | --- |
| Application | Logged in application name. |
| Team | Operator team name. |
| Queue | Logged in station queue. |
| Usage | Shows the users who work relatively much compared to others. You can add some work load to the users that are less busy. |

## View actions

Every command will be applied to selected rows only.

| Action | Description |
| --- | --- |
| Logout | Send this command to perform an orderly logout of the selected station instance. |
| Assign | Change the SLA assignment for the specific station instance. |

# Stations view

The **Stations** view provides information about the eFlow automatic stations activity.

## Stations view parameters

The stations grid view contains the following columns.

| Column name | Description |
| --- | --- |
| Name | Automatic station name. |
| PPM | Number of pages processed by the station per minute during the working time. "Working time" means the total time when the station was logged in, including the idle time. |
| CPM | Number of collections processed by the station per minute during the working time. By "Working time" we mean the total time when the station was logged in, including the idle time. |

| Column name | Description |
|---|---|
| Usage | The ratio between the net working time and total working time. |
| Queue | The name of the queue (the logged in station name). |
| Load | The number of collections that are waiting in the queue for this station instance. This value is only an estimation since it does not consider the station's performance.<br><br>For example, if there are 1000 collections in the queue and there are 2 instances of the same station, the **Load** column value for each instance will be 500.<br><br>The calculation of the number of collections in the queue takes the SLA assignments into consideration. Thus, the numbers in the example above will not be the same if the SLA assignments for the station instances differ. |
| Load rate | Indicates how fast the queue grows or shrinks. It becomes red if the input exceeds the output by 25% during the last ten minutes, green if output exceeds the input, and remains neutral if the input is approximately equal to the output. |
| Application name | The logged in application name. |
| Server | The name of the worker server on which this station instance is running.<br><br>ⓘ This is not the eFlow server! |
| CPU | The meaning depends on the **View by** selection:<br>• Station: CPU usage of the station instance<br>• Server: CPU usage of the whole server |
| Memory | The meaning depends on the **View by** selection:<br>• Station: Memory consumed by the station instance (MB)<br>• Server: Percent of the memory usage of the worker server |
| Hard disk | The meaning depends on the **View by** selection:<br>• Station: Number of read/write bytes per second<br>• Server: Percentage of the free disk space |

## View filters

The following filters are available.

| Filter | Description |
|---|---|
| Machine name | |
| Applications | Logged in application name. |
| Queue | List of autorun stations. |
| Usage | |
| SLA name | Used to exclude from the view the automatic stations that are not working with a specific SLA. |

The following grouping options are available.

| View by | Description |
|---------|-------------|
| Server | Station server name where several automatic stations are running. |
| Queue | Specific station queue. |
| Usage | Shows what stations work relatively much compared to others. You can add some workload to the stations that are less busy. |

## View actions

Every command will be applied to selected rows only.

| Action | Description |
|--------|-------------|
| Logout | Sends the command to perform an orderly logout of the selected station instance. |
| Assign | Changes the SLA assignment for the specific station instance. |
| Duplicate | Duplicates the station instance. |
| Run | Opens the Run dialog box. |

## Configure automatic stations

To configure the eFlow automatic stations:

1. On the **Supervise** window, click **Stations**.

2. Click **Run** in the upper right part of the view.
   The **Automatic Stations Assignment** dialog box opens.

3. Start the automatic stations (called also **Autorun** stations) by selecting the application.
   After the application was selected, the list of the available automatic stations appears.
   For each station you need set the following parameters.

| Parameter | Description |
|-----------|-------------|
| Station | The automatic station names list. The station will appear in this list if it has the **Autorun** and **StandBy** properties set to **True** in the Design module. |
| Add instances | The number of station instances that will be added to the currently running ones. |
| Server | The worker server where the station will be running. By default, the **Auto-select** option is set, meaning that the least busy instance of the server is selected. It is also possible to select a specific machine for running stations of the selected type. |
| Action | The **Apply to all** action appears after changing the station properties. Select this option to apply the changes to all the stations. |

4. Click **Run** to start all the Autorun stations or **Cancel** to ignore the configuration changes.

# Reports view

The **Reports** view allows you to analyze the system performance providing the data for the specified period of time.

Click on a report type to view the report.

| Report type | Description |
|---|---|
| Burn down | Shows the amount of work left versus the time. By comparing a real graph to the planned one, you can determine whether the system is working slower or faster than expected. |
| Collection load | Shows the total amount of collections that entered the system during the specified period of time. |
| Performance history | Shows the distribution of collections per queue during the specified time period. You should pay special attention to stations that differ significantly from the others. |
| Form types total | Shows the total number of forms that were processed during the specified time period. |
| Form types history | Shows the form types distribution for the specified period of time. |
| Target vs processed collection | Shows a comparison of the number of collections done and the planned values for the specified period of time. The chart shows the distribution of the workflow data per SLA. |
| Queue performance | Shows changes in the selected queue load for the specified period of time. |
| Station license usage | Shows the system license usage during the specified period of time. |

To print the selected report, click **Print**.

## View filters

The following filters are available.

| Filter | Description |
|---|---|
| Application | Application name. |
| SLA name | The SLA name. |
| Dates | The period used to create the report. You can select one of the following predefined values: <br> • Last 1 hour <br> • Last 24 hours <br> • Last 7 days |

# Alerts view

The eFlow alert system warns you about currently existing or foreseeable system problems. There are several alert types that are issued in problematic situations. Some alerts are relevant for the whole system, while others are connected to the SLA for the specific application and/or queue.

The Alerts view allows you to see all the system alerts and helps you to navigate to other views to get more information about the problem and fix it if possible.

Every alert issued by the system appears in the view and remains there until it is closed. If an alert was updated (new batches added to reject queue), the view will change, but a new notification will not be sent.

All alerts are removed automatically when the error condition is no longer met. They can also be closed manually by the project administrator, though this action is not always recommended.

## Alerts view parameters

The **Alerts** grid view contains the following columns.

| Column name | Description |
|---|---|
| Name | One of the predefined alert names. |
| Creation date | Alert creation date. |
| Message | Provides information about the problem. It is updated every minute (for example, the time behind schedule can grow), but an e-mail notification is sent only for the first time. |
| Published | Indicates whether a notification was already sent. |
| Application Name | An optional column. It will be filled if only if it is relevant for the specific alert. * |
| SLA | An optional column. It will be filled if only if it is relevant for the specific alert. * |
| Queue | An optional column. It will be filled if only if it is relevant for the specific alert. * |

* For example, it may contain the application name for the Behind schedule alert but be empty for the CPU alert.

## View filters

The following filters are available.

| Filter | Description |
|---|---|
| Application | Application name |
| Sla name | The SLA name |
| Alert type | One of the predefined alert types |
| Dates | Allows you see alerts that were recently issued or alerts that remain for a long time in the system. You can select one of the following predefined values:<br>• Recent<br>• 1 hour<br>• 1 day<br>• Old |

The following grouping options are available.

| View by | Description |
|---|---|
| Alert Name | Alert name |
| Date View | Alert creation date |

## View actions

Every command will be applied to selected rows only.

| Filter | Description |
|---|---|
| Remove | Closes the alert. It will disappear from the alerts list but will not be removed from the database immediately. This will be done later internally. |
| Navigate | Navigates to the relevant view that provides full information about the problem and a possibility to fix it.<br><br>For example, the navigation from the Queue Load alert will bring you to the Stations or Operators view (for automatic and manual stations respectively) grouped by the queue and filtered by application. |

## Alerts groups

All alerts are divided into four groups.

| Alert Group | Members |
|---|---|
| Performance | SLA behind schedule |
| | Queue load |
| | Empty queue |
| Workload | Reject |
| IT | CPU |
| | LicenseExpired |

# Tutorials

This chapter details information on how to manage system performance, using SLA, and analyzing alerts.

## Managing system performance

This section provides advice on managing system performance.

### Recognizing the performance problem

How do you identify low system performance?

The first sign is that the **Performance** alerts list is not empty. You can either navigate to the relevant view directly from the list or open the Alerts view to check the details.

Alerts are issued when something goes wrong, or might possibly go wrong, in the system, but they do not provide much information about the problem. The following table shows which indications may appear in the different views:

| View | Low performance signs |
|---|---|
| SLAs | Real CPM/PPM differs significantly from the planned value. The **Utilization** (burn down) graph shows that there are still collections in the system, although according to the plan, the work should be already done. |
| Stations | The load rate column is colored in red indicating that the station is working too slow. There is a problem with the automatic station performance. |
| Operators | The Load Rate column is colored in red indicating that the station input exceeds the queue output. There is a problem with the manual station performance. |
| Alerts | Ideally, this view should always stay empty. Any record that appears in the Alerts view requires immediate attention. |

### Why does the system work slower than expected

There can be many different reasons for low system performance. The Supervise station tools allow you to get to the root of the problem as quickly as possible.

The most common reasons for performance problems are listed below:

| Problem | Alerts | How to fix |
|---------|--------|------------|
| Some stations are overloaded | Queue load<br>Behind schedule | Automatic stations: Add relevant station instances<br>Manual stations: Add operators |
| Some collections never leave the system due to wrong SLA assignments | Behind schedule | Change the SLA assignments to ensure balanced system behavior. |
| The reject queue is not empty (relevant for all workflow stations except input stations) | Rejects<br>Behind schedule | Make sure that there are no collections in the Rejects queue. |
| Hardware problems (machines or network) | CPU<br>Queue load<br>Behind schedule | If the CPU alert appears, check the worker servers.<br>If there is no CPU alert, but the overall system performance is getting low, check the network speed and stability. |

## Automatic stations performance

To see the automatic stations runtime information, open the **Stations** view. If there are any problems, this view will provide you with an immediate indication.

| Symptom | Solution | Action |
|---------|----------|--------|
| The Usage column shows that none of the stations is idle, but some stations are still behind schedule.<br><br>The Queue Load column shows overloaded queues. | Add resources. | Click **Run** or **Duplicate** to add more stations. |
| The Usage column shows that some stations are working less than others, or not working at all. | Rearrange the existing resources. | Check the problematic station queue:<br>• If there are collections in the queue, and the Usage is low, then you need to consider the hardware problem.<br>• If the queue is empty, it is necessary to change the SLA assignment for this station. |

## Manual stations performance

To see the manual stations runtime information, open the **Operators** view. If there are any problems, this view will provide you with an immediate indication.

| Symptom | Solution | Action |
|---------|----------|--------|
| The Queue Load column shows overloaded queues. | Add resources. | Add manual station operators. |

| Symptom | Solution | Action |
|---|---|---|
| The Usage column shows that all operators are busy, but some manual stations are still behind schedule. | | |
| A closer look at the Usage column shows that some operators are working less than others or not working at all. | Rearrange the existing resources. | Check the queue where the less busy users are logged in:<br><br>• If there are collections in the queue, and Usage is low, then you need to consider the human factor (lunch break, poor professional skills, etc.).<br><br>• If the queue is empty, it is necessary to change the SLA assignment for this station. |

# Using SLA

eFlow SLAs should always reflect the project business logic. There are two most common scenarios that require several SLAs in one application:

• Some documents should move through the workflow faster than others.

   For example, urgent insurance claims compared to standard insurance payments.

• Different types of documents need to be processed by different operator teams.

   For example, multi-language or multi-site projects.

If the splitting of SLAs occurs after a certain point on the workflow, then it might be necessary to define an SLA that is active until the splitting occurs, and two or more other ones that will be active after the new condition is met.

## Multilanguage project example

The application should process forms that come from the different countries. All forms are scanned on the same site and are sent to recognition stations where the relevant language metatag is set. Data entry is performed by teams speaking different languages.

In this case the SLA schema should look as follows:

| SLA | |
|---|---|
| SLA 1 | Input Recognition |
| SLA 2 | Data Entry (English speaking team only) Export |
| SLA 3 | Data Entry (German speaking team only) Export |
| SLA 4 | Data Entry (Spanish speaking team only) Export |

The assignments for SLAs containing the Data Entry stations will be defined so that each instance of this station will get only the batches that have the specified metatag set.

## Advantages of the SLA approach

The SLAs simplify the workflow configuration: instead of setting three different Data Entry stations, we're using the different instances of the same station.

Besides, it is easy to notice the process bottlenecks and rearrange the resources if needed.

# Analyzing alerts

Supervise alerts enable you to foresee potentially dangerous conditions, or issue immediate warnings about existing problems.

Active alerts indicators are found in the upper part of the screen and are visible from all the Supervise views. The Alerts view presents detailed information about the alerts present in the system.

## Performance alerts

| Alert | Explanation |
|---|---|
| **Queue Empty** is issued for all the queues. No alerts regarding the low load rate are issued. | Input stations do not inject collections to the system. This can be due to one of the following problems:<br>• Collect: the search path is either empty or contains invalid files (for example, files with the wrong resolution).<br>• Scan: the scanner fails to scan the pages or its feeder is empty.<br>• Collections are created but are not sent to the server. Check the SQL connection. |
| **Queue load** and **Behind schedule** alerts are issued for the specific queue. | The station is blocking the workflow. This can be due to one of the following problems:<br>• The station is working normally, but it cannot process all the images fast enough. Check the station performance.<br>• Collections are not sent to the server. Check the SQL connection. |

⚠ The **Queue Empty** alert will be issued only if the stations were working normally and stopped working while the SLA is still running. If the SLA was not started (that is, the input station did not produce any collections at all), there will be no alerts. In the **Stations** view you will see that the input stations are running, but CPM=0 and PPM=0.

## IT alerts

| Alert | Explanation |
|---|---|
| CPU | One of the worker servers shows high CPU usage. Navigate from the CPU alert to the Stations view and group the results by server to see the problematic server details. |

## Workload alerts

| Alert | Explanation |
|---|---|
| Rejects | There are collections in the **Rejects** queue. It can happen for the following reasons:<br><br>• Collections contain invalid images or data and should be deleted.<br><br>• Collections were sent to the **Rejects** queue due to some problem in the workflow (for example, incorrect routing rules). You need to fix the workflow problem and move the collections to another queue manually. |