

# eFLOW Supervise

## User Guide

Version 5.2

Copyright © TIS, Top Image Systems. All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, translated, transcribed or transmitted in any form, or by any means manual, electric, electronic, electromagnetic, mechanical, chemical, optical, or otherwise, without the prior explicit written permission of TIS.

# Contents

<b>About Supervise</b> .....	<b>5</b>
Terms and definitions .....	5
<b>Installation</b> .....	<b>6</b>
Security settings .....	6
Software requirements .....	6
Localization .....	6
Customization .....	6
<b>Supervise main view</b> .....	<b>7</b>
Stations view .....	7
Customize the Stations view .....	7
Stations view parameters .....	8
Stations view actions .....	9
Configure automatic stations .....	10
Operators view .....	10
Customize the Operators view .....	10
Operators view parameters .....	11
Operators view actions .....	12
SLAs view .....	13
Customize the SLAs view .....	13
SLA card view .....	14
Utilitization view .....	14
Processing CPM view .....	14
Global Performance view .....	14
Alerts view .....	15
Customize the Alerts view .....	15
Alerts view parameters .....	16
Alerts view actions .....	17
Alerts groups .....	17
Reports view .....	17
Customize the reports view .....	17
Reports types .....	18
Print reports .....	18
<b>Supervise Settings</b> .....	<b>19</b>
SLA configuration .....	19
Managing the SLAs .....	19
Basic SLA properties .....	19
SLA types .....	20

Performance properties .....	20
SLA display .....	20
SLA rules .....	21
Checking the SLA rules .....	21
SLA rule editor .....	21
SLA assignment .....	22
View assignments .....	22
Inherited assignment .....	23
Manual assignment .....	23
Change the assignments .....	23
Alerts configuration .....	24
Alert types .....	24
Alerts configuration .....	25
User definitions .....	25
Users and Groups dialog box .....	25
Creating a new team .....	25
Adding a new operator .....	25
<b>Tutorials .....</b>	<b>26</b>
Managing system performance .....	26
Recognizing the performance problem .....	26
Why does the system work slower than expected? .....	27
Automatic stations performance .....	27
Manual stations performance .....	28
SLA usage scenarios .....	28
Multi-language project example .....	29
Advantages of the SLA approach .....	29
Analyzing alerts .....	29
Performance alerts .....	30
IT alerts .....	30
Workload alerts .....	30

# About Supervise

Supervise is the 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.

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

## Terms and definitions

Term	Definition
SLA	A Service Level Agreement 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 till the task will be done.
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 <b>Users</b> table of the eFLOW management database. All operators should also be defined as valid STS users.
Team	The 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	The eFLOW client intended for running the automatic stations. In modern environments, the automatic stations are usually run on powerful multi-CPU server stations, hence the name "Worker server" for this type of client.

# Installation

Supervise can be installed either as one of the eFLOW client installation features or as a separate application using the ClickOnce deployment mechanism. See <http://msdn.microsoft.com/en-us/library/t71a733d.aspx> for information on ClickOnce.

The ClickOnce Supervise server-side installation is deployed as a separate installation package.

**Note:** The ClickOnce Supervise package is installed as a trusted application and requires relevant access permissions.

## Security settings

It is important to keep in mind that the Supervise station provides the most comprehensive information about the system status and allows users to change it in many ways. Only qualified users should have full access to the Supervise functionality.

## Software requirements

- Microsoft Windows 8
- Microsoft Windows 7
- Microsoft Windows Server 2008 R2
- Microsoft Windows Vista
- Microsoft Windows Server 2003

## Localization

The standard Supervise installation supports the following languages:

- English
- German
- Spanish
- French

If an additional language is required, a change request should be sent to TIS. Currently there is no possibility to create custom language files.

The language selection can be changed by editing the *efOrchestrate.exe.config* file. The English language is set by default.

## Customization

The current version does not support customizations.

# Supervise main view

The main view of the eFLOW Supervise module provides access to the following views:

- [Stations](#)
- [Operators](#)
- [SLAs](#)
- [Alerts](#)
- [Reports](#)

It is possible to access the **Settings** dialog box, and to check the status of the system alerts grouped by type.

**Note:** The views are automatically refreshed once per minute.

## Stations view

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

### Customize the Stations view

When opened for the first time, the **Stations** view contains all the station-related information available in the system. In most cases, this view is difficult to handle. You can easily customize the view to show only the relevant SLAs and to group them as required.

### Set the filter

You can apply the following filters:

Filter	Description
Machine name	
Applications	
Queue	
Usage	
SLA name	Used to exclude from the view the automatic stations that are not working with the specific SLA.

## Group the filtered views

After filtering the view, you can decide how you want to view the controls. The following viewing 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 would like to add some work load to the stations that are less busy.

## Stations view parameters

The stations grid view contains the following columns:

Column name	Description
Name	Automatic station name
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.
PPM	Number of pages 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.
Usage	The ratio between the net working time and total working time.
Queue	The name of the queue (the logged in station name).
Load	<p>The number of collections that are waiting in the queue for this station instance. This value is only an estimation, since it does not take into account the stations performance.</p> <p>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.</p> <p>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.</p>



Column name	Description
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 worker server name where this station instance is running. Note:It is not the eFLOW server!
CPU	The meaning depends on the <b>View by</b> selection: <ul style="list-style-type: none"> <li>■ Station: CPU usage of the station instance</li> <li>■ Server : CPU usage of the whole server</li> </ul>
Memory	The meaning depends on the <b>View by</b> selection: <ul style="list-style-type: none"> <li>■ Station: Memory consumed by the station instance (MB)</li> <li>■ Server : Percent of the memory usage of the worker server.</li> </ul>
Hard disk	The meaning depends on the <b>View by</b> selection: <ul style="list-style-type: none"> <li>■ Station: Amount of read/write bytes per second.</li> <li>■ Server : Percentage of the free disk space.</li> </ul>

### Stations view actions

Every command will be applied to selected rows only.

Action	Description
Logout	Sends the command to perform the 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 <a href="#">Run</a> dialog.

## Configure automatic stations

To configure the eFLOW automatic stations:

1. Open the **Stations** view in the Supervise station.
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 <b>Autorun</b> and <b>StandBy</b> properties set to <b>True</b> in the Visual Designer.
Add instances	The number of the 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 the specific machine for running the stations of the selected type.
Action	<b>Apply to all</b> action appears after changing the station properties. Select it to apply the changes to all the stations.

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

## Operators view

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

### Customize the Operators view

When opened for the first time, the **Operators** view contains information about all the manual stations. In most cases, this view is difficult to handle. You can easily customize the view to show only the relevant items and to group them as required.

## Set the filters

You can apply the following filters:

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.

## Group the filtered views

After filtering the view, you can decide how you want to group the controls. 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 would like to add some work load to the users that are less busy.

## Operators 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 have to be defined in the dialog box.
Application	Name of the application currently logged in.

Column name	Description
CPM	Number of collections done by the operator per minute during the working time. By "working time" we mean the total time when the station was logged in, including the idle time ( operator breaks, etc .)
PPM	Number of pages done by the operator per minute during the working time. By "working time" we mean the total time when the station was logged in, including the idle time ( operator breaks, etc .)
Usage	The ratio between the net working time and total working time.
Queue	The name of the queue ( the logged in station name ).
Load	<p>The number of collections that are waiting in the queue for this station instance. This value is only an estimation, since it does not take into account 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.</p> <p>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.</p>
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	Machine name where the station instance is running.
Team	User 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 <b>View By</b> selection. If the SLA assignment is FIFO (all assignment values are equal to zero), this column will be empty.

## Operators view actions

Every command will be applied to selected rows only.

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

## 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.

### Customize the SLAs view

When opened for the first time, the **SLAs** view includes all the SLAs available in the system. In most cases this view contains too much information and is difficult to handle. The view can be easily customized to show only the relevant SLAs and to group them as required.

### Set the SLA filters

The following SLA filters can be applied:

Filter	Description
Applications	Show only the SLAs that belong to selected application. If <b>All</b> is selected, the applications filter will not be applied.
Alerts	User can decide to see only those SLAs where active <b>Performance</b> alerts exist. If <b>All</b> is selected, the alerts filter will not be applied.
Status	User can decide to see only those SLAs that are behind schedule. If <b>All</b> is selected, the status filter will not be applied.

### Group the filtered SLAs

After filtering the SLAs, the user can decide how does he want to group the SLA cards. The following grouping options are available:

View by	Description
Application	Group the SLA cards by the application name
Alert	Group the SLA cards by the active alerts
Status	Group the SLA cards by their status
Load	Show separately the <b>Empty</b> SLAs that have no collections and the <b>Loaded</b> SLAs that have some collections attached.
Finish	Show separately SLAs that are <b>On time</b> and <b>Late</b> . SLAs that have no <b>Finish</b> time defined will be shown as <b>Unknown</b> .

**Note:** 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 color changes.

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

### Utilitization view

This is actually a burndown graph showing how much work is left at the moment. It always shows the work progress during the last hour.

If **Start** and **Finished** times are defined, the graph will always start at the maximal value and go down till 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 will be 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 there collections that are moving slowly.

### Processing CPM view

This tab shows the CPM values for the specific SLA during the last hour.

### Global Performance view

This 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 include 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 will be empty.

The controls inside 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.

Property	Description	Values	Notes
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 till the task will be done.	Real Planned	TTC can be negative for the batches that are in the system for long time.
Collections	Number of collections.	Total number  In <b>Rejected</b> queue  On hold	

## Alerts view

The eFLOW alert system is used to warn about the 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 to see all the system alerts and helps 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 alert was updated ( new batches added to reject queue ), the view will be changed but new notification will not be sent.

All alerts are removed automatically when the error condition is not met any more. They can also be closed manually by the project administrator, though this action is not always recommended.

### Customize the Alerts view

When opened for the first time, the **Alerts** view contains all the alerts information available in the system. In most cases this view is difficult to handle. The view can be easily customized to show only the relevant alerts and group them as required.

### Set the filters

Filter	Description
Applications	The application name
SLA name	The SLA name
Alert type	One if the predefined alert types.

Filter	Description
Date	<p>Allows to see the alerts that were recently issued or those that remain for a long time in the system. One can select one of the following predefined values:</p> <ul style="list-style-type: none"> <li>■ Recent</li> <li>■ 1 hour</li> <li>■ 1 day</li> <li>■ Old</li> </ul>

### Group the filtered views

After filtering the view the user can decide how does he want to view the controls. The following viewing options are available:

View by	Description
Alert name	Alert name
Date	Alert creation date

### 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 some information about the problem. It is updated every minute (for example, the time behind schedule can grow), but the e-mail notification is sent only for the first time.
Published	Indicates whether notification was already sent.
Application	An optional column. It will be filled if only if it 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 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.



## Alerts view actions

Action	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 will give the full information about the problem and a possibility to fix it. For example, the navigation from the <b>Queue Load</b> alert will bring you to the <b>Stations</b> or <b>Operators</b> view (for automatic and manual stations accordingly) 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

## Reports view

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

### Customize the reports view

When opened for the first time, the **Reports** view contains all reports available in the system. The view can be easily customized to show only the relevant reports and to group them as required.

The following filters can be selected:

Filter	Description
Application	Application name

Filter	Description
Dates	<p>The period of time used to create the report. One of the following predefined values can be selected:</p> <ul style="list-style-type: none"> <li>■ Last 1 hour</li> <li>■ Last 24 hours</li> <li>■ Last 7 days</li> </ul>

## Reports types

Report type	Description
Burn down	The burn down report the amount of work left vs the time. By comparing a real graph to the planned one can say whether the system is working slower or faster than expected.
Collection load	The total amount of collections that entered the system during the specified period of time.
Performance history	<p>Distribution of collections per queue during the specified period of time.</p> <p>The workflow manager should pay special attention to stations that differ significantly from the others.</p>
Form types total	The total number of forms that were processed during the specified period of time.
Form types history	Form types distribution for the specified period of time.
Target vs processed collection	Comparison between 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	Changes in the selected queue load for the specific period of time.
Station license usage	The system license usage during the specified period of time.

## Print reports

Click the **Print** button to print the selected report.

# Supervise Settings

## SLA configuration

This topic describes how to configure SLAs in the Supervise station.

### Managing the SLAs

The SLA list shows all the SLAs of the current eFLOW system.

Action name	Description
Create	Select <b>New</b> to create a new SLA. See SLA properties for details.
Edit	Select one of the existing SLAs and click <b>Edit</b> to change the current settings.
Assign	Select one of the existing SLAs and click <b>Assign</b> to <a href="#">change the assignments</a> of the application.
Delete	Select an existing SLA and double-click <b>Delete</b> to delete it.

### Basic SLA properties

Basic properties are the minimal set of parameters that are required to define a valid SLA.

Property	Description
SLA name	The user defined name of the SLA. This is a mandatory property that cannot be changed.
Application	The SLAs are always defined per application. This mandatory property contains the name of the application to which this SLA belongs. It cannot be changed after the SLA was created.
Export queue	<p>This property defines the last workflow station for the SLA.</p> <p>The list of possible values includes all the workflow stations of the current application and the generic constant <b>Out</b>.</p> <p><b>Out</b> means that SLA should be finished with the end of the workflow. If current SLA has finished and the next one has not yet started, the batch will be still connected in the database to its last SLA, though it is not really working any more.</p>

## SLA types

Type	Description
Fixed	This type requires <b>Start</b> and <b>Finish</b> time to be set. Its recurrence can be set to <b>Daily</b> or <b>Weekly</b> . The <b>Fixed</b> SLA always begins and ends on the same day. If its recurrence is set to <b>Weekly</b> , it will be executed on selected days.  There is no possibility to define the SLA that starts on Sunday and ends on Monday.
TTC (Time to complete )	The time that takes one collection to finish the SLA. It can be defined in seconds, minutes, hours or days.

## Performance properties

Property	Description	Notes
PPM (Pages per minute)	The planned amount of pages to be processed per minute. It is relevant for automatic stations.	We recommended to define this property though it is not mandatory.
CPM (Collections per minute)	The planned amount of collections to be processed per minute.	We recommended to define this property though it is not mandatory
Quantity	The planned total amount of collections to be processed while SLA is active.	This is an optional property.

## SLA display

Currently, **Color** is the only display setting available for the SLA display. It is not a mandatory setting, but can often help to distinguish quickly between different SLAs. You can select one of the predefined colors to be used with this SLA.

Any SLA property besides the SLA name and application can be changed at any moment.

**Note:** Though it is possible to change the SLA rules while there are still collections in the workflow, we do not recommend it.

The safe way to change the SLA rules is to make sure there are no collections and change the rules afterwards.

## SLA rules

Every eFLOW collection should be attached to some SLA. The SLA rules help to determine to which SLA the collection should belong.

Rules are defined per flow in the selected application.

The SLA rule can be based on any of the application metatags defined in Visual Designer.

### Checking the SLA rules

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

1. The collection is looking for a rule that is relevant for it.
2. If no relevant rules are found, it is looking for the SLA without rules.
3. If the 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, we cannot guarantee the order of the execution of SLAs.

The 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 the SLA rules are checked on every **Put collection** on every station. If SLA ends on the specific station and the workflow continues, you one should define the rules so that they will fail on the next SLAs.

### SLA rule editor

The SLA rule editor allows to add, remove and edit the rules.

#### Parameters

The list of available rule parameters is based on the flow metatags defined in the Visual Designer. The metatags list cannot be changed in the Supervise station SLA rule editor.

#### Operators

The rule editor supports the following comparison operator types :

- Logical operators:
  - >
  - <
  - =
  - <>

- >=
- <=
- String functions:
  - *Contains*
  - *StartsWith*
  - *EndsWith*.
- Regular expressions.

**Important:** All rule comparisons are case-sensitive.

## Values

The values should correspond to the metatag type.

The values of the **DateTime** type parameters (for example, **CreationTime** ) should use the system date format of the eFLOW server machine.

## Actions

The table below describes the rule editor actions.

Action	Description	Usage
Add	Add a new rule in the rule editor	Always visible
Accept	Save a new rule in the database	Visible after adding a new rule in the editor
Cancel	Do not save a new rule in the database	Visible after adding a new rule
Edit	Edit the existing rule	Visible when selecting one of the existing rules
Delete	Delete the existing rule	Visible when selecting one of the existing rules

## SLA assignment

Application can be configured to keep the required balance between the amount of collections bound to different SLAs. These configuration parameters in the Supervise station are called SLA assignments.

### View assignments

To see the existing SLA assignments per application:

1. Click the **SLA Assignments** button in the **General** section of the **Settings** view to open the **Global SLA Assignments** dialog.

The **Global SLA Assignments** dialog contains all the system eFLOW SLA assignments ordered by

application name.

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

## 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. It is possible to 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 ), etc.

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

SLA assignments can be set manually using the slider control. 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.

## Change the assignments

You can change only one assignment at a time, since other assignments should be updated accordingly.

After setting the new values it is necessary either to save this change or undo it by clicking **Reset**. The **Reset** button returns the current SLA assignments to its previous state.

## Alerts configuration

This topic describes the alerts used in the Supervise station.

### Alert types

The table below 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 TIS 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 <b>Reject</b> queue. This alert is issued per application.	This alert is removed automatically if the <b>Reject</b> queue is empty. If you close this alert manually, it appears again only if new batches are sent to the <b>Reject</b> 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.
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.



Alert type	Description	Alert resolution
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.

## Alerts configuration

### 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 SMTP settings.

Check **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.

## User definitions

This topic describes how to set users and groups definitions in the Supervise station.

### Users and Groups dialog box

To open the **Users & Groups dialog** click the **Users** button in the **General** section of the **Settings** view.

### Creating a new team

Click the **New Team** button to create a new team. The default name should be replaced by a meaningful one. After creating a team you need to add some user definitions to it.

**Note:** Teams defined in the **Users & Groups** dialog are valid for the Supervise station only.

### Adding a new operator

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

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. Obviously, the user that is not defined in the STS will not be able to log in.

Operators and groups defined in this dialog are saved in the **User** table in the **eFLOW\_Management** SQL database.

**Important:** Currently there is no automatic utility that synchronizes between the system security settings ( STS ) and the User table. The eFLOW system administrator should make sure that users and teams are properly set.

# Tutorials

## Managing system performance

This topic provides advice on managing system performance.

### Recognizing the performance problem

How do I 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 table below shows what indications may appear in 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 while 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 <a href="#">Alerts</a> view requires the immediate attention.

## Why does the system work slower than expected?

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

The most common reasons for the performance problems are listed below:

Problem	Alerts	How to fix
Some stations are overloaded	Queue load Behind schedule	Automatic stations : add relevant station instances  Manual stations: add operators
Some collections never leave the system due to the wrong SLA assignments	Behind schedule	Change the SLA assignments to ensure the balanced system behavior.
The reject queue is not empty (relevant for all the workflow stations except Input)	Rejects Behind schedule	Make sure that there are no collections in the <b>Rejects</b> queue.
Hardware problems ( machines / network )	CPU Queue load Behind schedule	If the <b>CPU</b> alert appears, check the worker servers.  If there is no <b>CPU</b> 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 an immediate indication.

Symptom	Solution	Action
The <b>Usage</b> column shows that none of the stations is idle, but some stations are still behind schedule.	Add the resources	Click <b>Run</b> or <b>Duplicate</b> to add more stations.
The <b>Queue Load</b> column shows overloaded queues.		

Symptom	Solution	Action
The <b>Usage</b> column shows that some stations are working less than others or not working at all.	Rearrange the existing resources	Check the problematic station queue: <ul style="list-style-type: none"> <li>■ If there are collections in the queue, and the <b>Usage</b> is low, then you need to consider the hardware problem.</li> <li>■ If the queue is empty, it is necessary to change the <b>SLA assignment</b> for this station.</li> </ul>

## Manual stations performance

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

Symptom	Solution	Action
The <b>Queue Load</b> column shows overloaded queues.	Add the resources	Add manual station operators.
The <b>Usage</b> column shows that all operators are busy, but some manual stations are still behind schedule.		
A closer look at the <b>Usage</b> 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: <ul style="list-style-type: none"> <li>■ If there are collections in the queue, and <b>Usage</b> is low, then you need to consider the human factor (lunch break, poor professional skills, etc.).</li> <li>■ If the queue is empty, it is necessary to change the SLA assignment for this station.</li> </ul>

## SLA usage scenarios

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 the 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 till the splitting occurs, and two or more other ones that will be active after the new condition is met.

### Multi-language 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 metatags 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 re-arrange 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 the detailed information about the alerts present in the system.

## Performance alerts

Alert	Explanation
<p><b>Queue Empty</b> is issued for all the queues. No alerts regarding the low load rate are issued.</p>	<p>Input stations do not inject collections to the system. It can be due to one of the following problems:</p> <ul style="list-style-type: none"> <li>■ File Portal: the search path is either empty or contains invalid files (for example, wrong resolution).</li> <li>■ Scan portal: the scanner fails to scan the pages or its feeder is empty.</li> <li>■ Collections are created but are not sent to the server. Check the SQL connection.</li> </ul>
<p><b>Queue load</b> and <b>Behind schedule</b> alerts are issued for the specific queue</p>	<p>The station is blocking the workflow. It can be due to one of the following problems:</p> <ul style="list-style-type: none"> <li>■ The station is working normally, but it cannot process all the images fast enough. Check the station performance.</li> <li>■ Collections are not sent to the server. Check the SQL connection.</li> </ul>

**Important:** 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 (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 <b>Stations</b> view and group the results by server to see the problematic server details.

## Workload alerts

Alert	Explanation
Rejects	<p>There are collections in the <b>Rejects</b> queue. It can happen for the following reasons:</p> <ul style="list-style-type: none"> <li>■ Collections contains invalid images or data and should be deleted.</li> <li>■ Collections were sent to the <b>Rejects</b> queue due to some problem in the workflow ( for example, wrong routing rules ). You need to fix the workflow problem and move the collections to another queue manually.</li> </ul>