



IBM ECM System Monitor Field Guides

Monitoring Listener Metrics

August 29th, 2023
CENIT AG
Michael Wohland

Content

- Introduction 3
- Overview 3
- Disclaimer 3
- Purpose of this guide..... 4
- Requirements..... 4
- Step by Step setup 5
 - 1. Extract System Dashboard and run it 5
 - 2. Preparing ESM for using Listeners 8
 - 3. How to read information from the system dashboard..... 9
 - 4. Setting up a Listener probe in ESM 11
 - 5. Aggregation of Listener information in reports and showing the report in the monitoring dashboard..... 16
- Contact Information 19

Introduction

Overview

This guide describes in detail, how ESM can be used to monitor performance and capacity metrics provided by IBM Listeners in real-time. The system dashboard is used to gather the Listener objects which will then be used in ESM for monitoring and alerting.

Disclaimer

The content of this document is based on ESM Version 5.5.5.1. The descriptions and guidelines in this document are for informational purposes only. Up-to-dateness, content completeness, appropriateness and validity for all possible scenarios cannot be guaranteed. All information is provided on an as-is basis. The author is not liable for any errors or omissions in this document or any losses, injuries and damages arising from its use.

If you are planning to setup or configure ESM or to adjust an existing installation, it is absolutely necessary to take into account current security whitepapers, release notes and announcements from the official IBM ECM System Monitor product documentation website.

Purpose of this guide

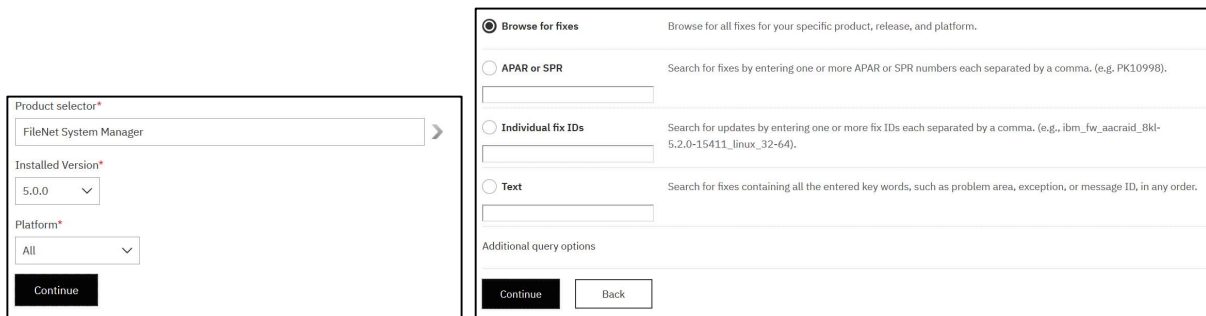
Many of the IBM Automation engines that are monitored by ESM provide so called Listener / PCH information with performance and capacity metrics. ESM can access this API and read the information. Afterwards it can be displayed in the ESM console. This guide will support you in setting up such probes.

Requirements

Port 32775 must be accessible on the Servers that run the Software like CPE

The System Dashboard is required. The latest Version can be download from Fix Central.

Search for the following:



The image shows two screenshots of the Fix Central interface. The left screenshot is the 'Product selector' for 'FileNet System Manager'. It has a dropdown for 'Installed Version' set to '5.0.0' and a dropdown for 'Platform' set to 'All'. A 'Continue' button is at the bottom. The right screenshot is the 'Browse for fixes' search interface. It has three radio buttons for search criteria: 'APAR or SPR', 'Individual fix IDs', and 'Text'. Each has a corresponding input field. There are 'Continue' and 'Back' buttons at the bottom.

Download this file:



The image shows a file download list. It has a checkbox, a '1' icon, and the text 'fix pack: → 5.0.0.4-P8SD-FP004'. Below it are two links: 'More Information' and '5.0.0.4-P8SD-FP004 Readme'.

You can also try this link:

<https://www.ibm.com/support/fixcentral/swg/selectFixes?parent=FileNet%20Product%20Family&product=ibm/Information+Management/FileNet+System+Manager&release=5.0.0&platform=All&function=all>

Step by Step setup

1. Extract System Dashboard and run it

The downloaded file is a zip file. It must be extracted to any location. There is no need to do an installation. After the file is extracted you'll find two files named Dashboard.bat and Dashboard.sh. In this example I am using it on Windows, so I need to adjust the .bat file. Open the Dashboard.bat in an editor.

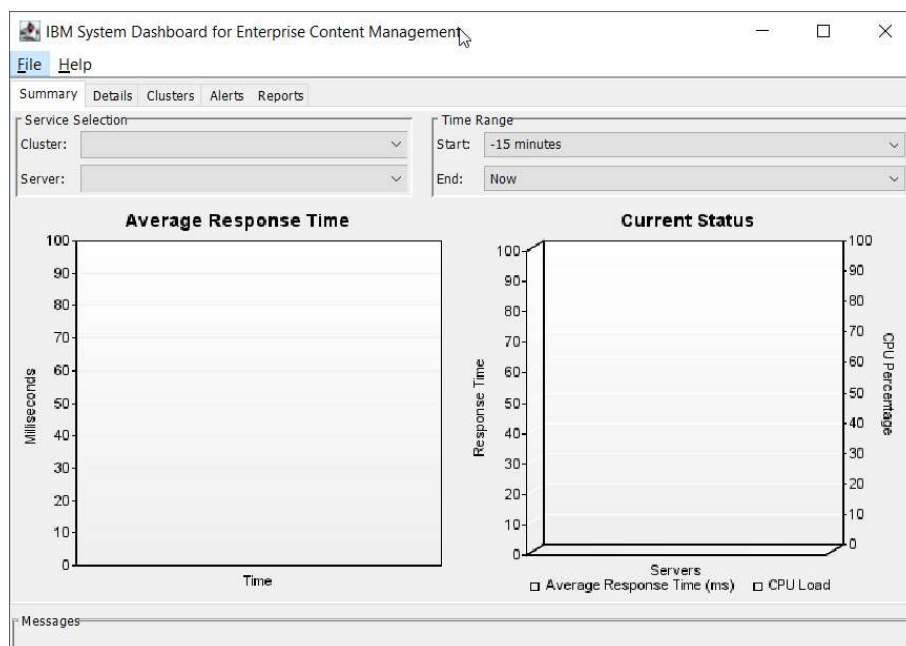
```
[@ECHO OFF

start "" "javaw" -Xms128m -Xmx1100m -cp
.;libs\forms_rt.jar;libs\chart.jar;libs\chart.ext.jar;libs\log4j-
1.2.17.jar;mgrlib.jar;mgr.jar com/filenet/PchMgr/Manager AcceptLicense

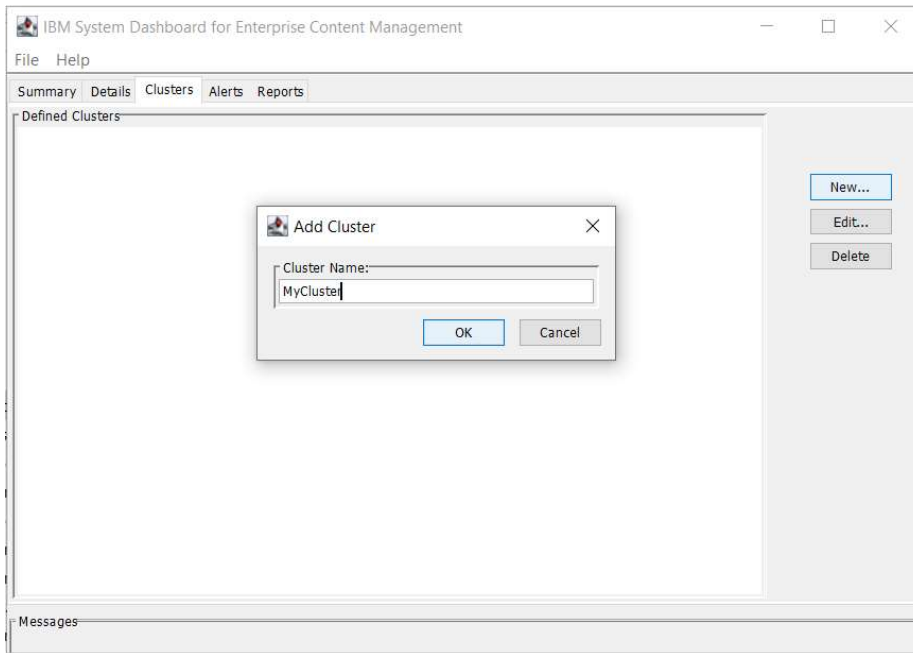
@echo In order to run the System Dashboard, you must indicate that you
accept the license. Please update the line above to indicate that you
accept the license
@echo by adding AcceptLicense to the end of the line. The line below
provides an example of the command updated with the AcceptLicense parameter
@echo start "" "javaw" -Xms128m -Xmx1100m -cp
.;libs\forms_rt.jar;libs\chart.jar;libs\chart.ext.jar;libs\log4j-
1.2.17.jar;mgrlib.jar;mgr.jar com/filenet/PchMgr/Manager AcceptLicense
```

In the “start ...” line the red highlighted AcceptLicense must be added at the end. Otherwise the tool will not start. Save the file afterwards.

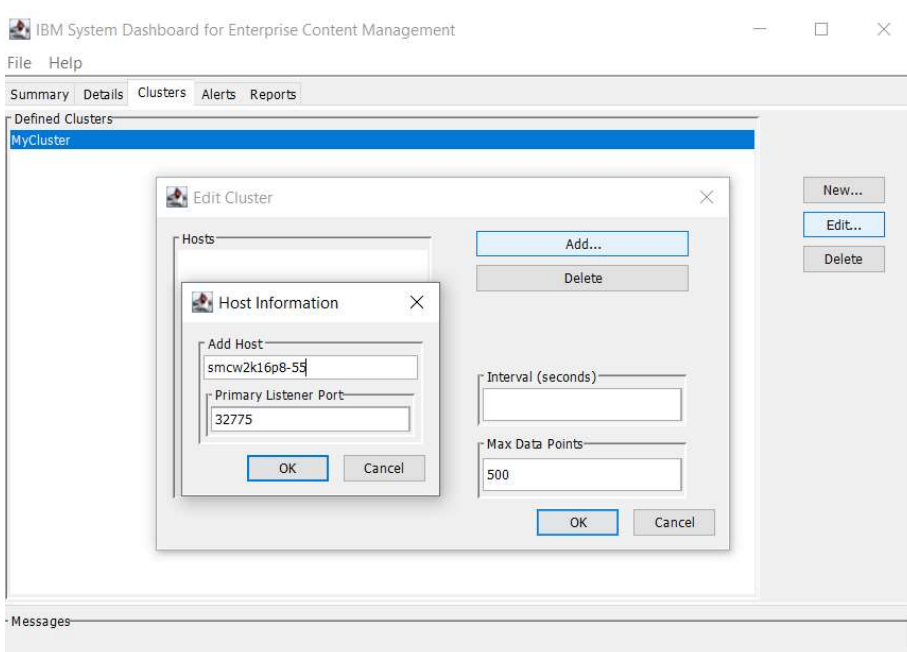
Start the System Dashboard by executing the Dashboard.bat



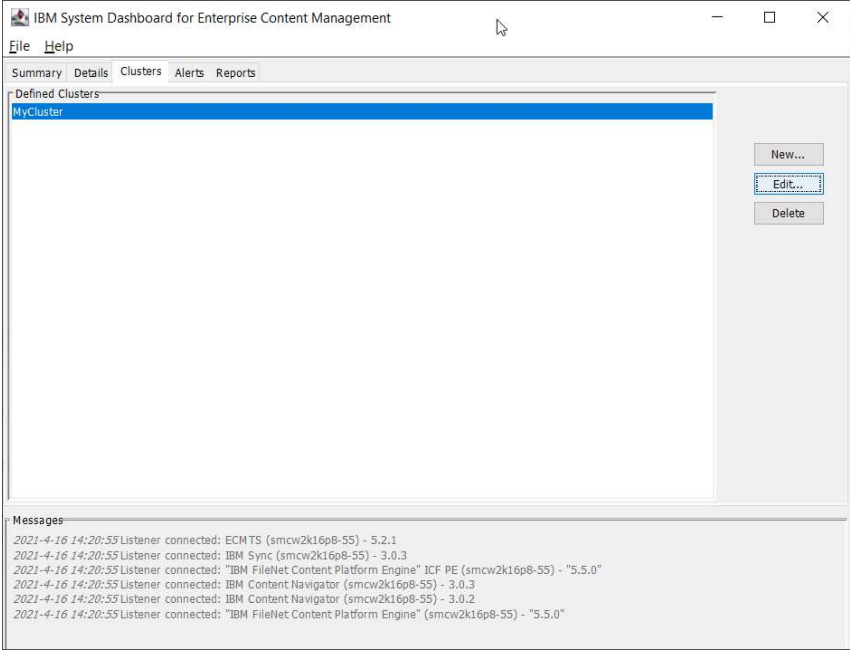
Once the Software is started switch to the Clusters tab and click on the “New ...” button on the right. Specify a name for the cluster:



Select the new Cluster from the Defined Clusters Window and click on the “Edit ...” button on the right. This will open the Cluster Editor in which you can add hosts. The port is usually 32775. For the interval do not go lower than 120 seconds. The interval defines the push interval (update) of the object. The default is every 15 min.



Once you have completed the list of hosts, click on “OK” in the Cluster Editor. The System Dashboard will connect to the Listener and will display all connected Listeners at the bottom.

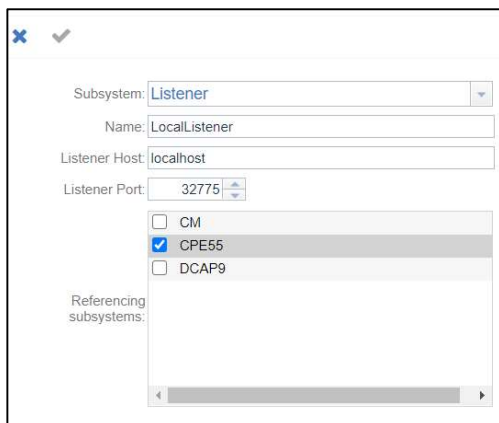


Via “File” in the menu you can save the cluster configuration if you like.

2. Preparing ESM for using Listeners

Before an ESM Agent can read the Listener information a subsystem from type Listener must be defined, that configures the connection to the corresponding server and port. If the ESM agents are running on the same server like the software monitored I usually prefer to setup just one Listener subsystem which I call LocalListener, as the agent can connect to localhost on port 32775. If that is not possible, the connection can also be established as a remote connection. In the subsystem the server name must be used in that case and of course the port must be accessible from “outside” of the system.

The subsystem that I have created in my ESM Console looks like this:



The screenshot shows a configuration window for a Listener subsystem. The 'Subsystem' dropdown is set to 'Listener'. The 'Name' field contains 'LocalListener'. The 'Listener Host' field contains 'localhost'. The 'Listener Port' field contains '32775'. Below these fields is a list of checkboxes for 'Referencing subsystems': 'CM' (unchecked), 'CPE55' (checked), and 'DCAP9' (unchecked). The 'Referencing subsystems' label is positioned to the left of the list.

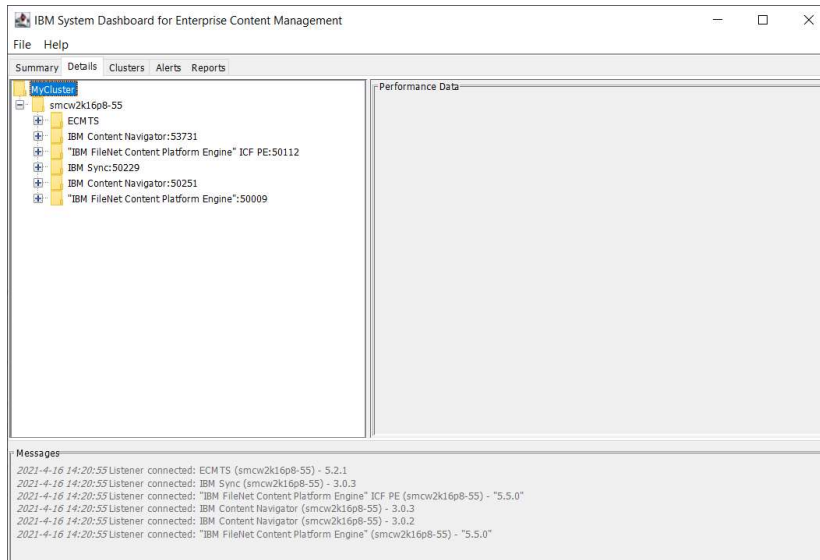
I am just referencing my CPE Subsystem here since I don't have more to reference. If you have several CPEs or additional Servers like ICN, Case Manager etc. you can also reference those.

At some subsystems you might get asked for Listener Application Names or Instances, e.g. on the CPE Subsystem. This information can be gathered from the System Dashboard – see next step.

3. How to read information from the system dashboard

Open the system dashboard and establish the connection to your server(s) as described in Step 1. If you have saved the cluster configuration as an xml file, you can also load it if needed.

Switch to the “Details” tab and double click on the name of the cluster configuration – in my case “MyCluster”, the window will look like this afterwards:



For each server that is part of you configuration you will get one entry with the server name. Below server name available Listener applications are given. In the screen these are

- ECMTS → Master
- IBM Content Navigator: 53731 → Content Navigator
- “IBM FileNet Content Platform Engine” ICF PE: 50112 → PE Part
- IBM Sync → IBM Sync Tool
- IBM Content Navigator: 50251 → acce Console
- “IBM FileNet Content Platform Engine”: 50009 → CE Part

ECMTS is the first (Master) Listener in this case, this runs on 32775. All others are having their internal port number given after the “.”. Nevertheless ESM just needs the master port and still can read from the others.

I do have two Instances of IBM Content Navigator, one reflects the acce page of the CPE in my case.

It is possible to run two or more instances of the same type on one server. This is where the Listener instance parameter in the ESM steps determines which of the instances it should read from. Unfortunately not all of the managed engines do offer the instance.

The instance is shown in the environment object as `app_instance` and usually reflects the installation path of the software or deployed application. Also the actual application name (`app_name`) and application version (`app_version`) are given here. You can use these three for the configuration of the subsystem or probe in ESM.

The screenshot shows the IBM System Dashboard for Enterprise Content Management. The left sidebar displays a tree view of the system hierarchy, with the 'Environment' object selected. The main area shows the 'Performance Data' for this object, with the following properties highlighted in yellow:

Property	Value
<code>app_guid</code>	"3f8e6ad03cc94a259465fce2bae4a47e"
<code>app_instance</code>	E:\IBM\FileNet\ContentEngine
<code>app_name</code>	"IBM FileNet Content Platform Engine"
<code>app_version</code>	"5.5.0"
<code>hostname</code>	smcw2k16p8-55
<code>java.runtime.version</code>	pwa6470_27sr3fp30-20160112_01 (SR3 FP30)
<code>java.version</code>	1.7.0
<code>jvm_max_memory</code>	4294967296
<code>listener_version</code>	5.0.0.146
<code>os.arch</code>	amd64
<code>os.name</code>	Windows 10
<code>os.version</code>	10.0
<code>processors</code>	4
<code>protocol</code>	1.3
<code>time</code>	1618808613878
<code>user.language</code>	de

Below the performance data, there is a 'Messages' section showing several log entries related to listener connections for various components like ECMTS, IBM Sync, and IBM FileNet Content Platform Engine.

For the configuration of the probe you will also need the so called Listener path. The Listener path reflects the "drill down" of the object in system dashboard and the requested information. In the example below this is `/USER/BASEOS/Document/Loads/Total Count` where `/USER/BASEOS/Document/Load` is the object and `/Total Count` is the information.

The screenshot shows the IBM System Dashboard for Enterprise Content Management. The left sidebar displays a tree view of the system hierarchy, with the 'USER' object selected. The main area shows the 'Performance Data' for this object, with the following columns highlighted in yellow:

Time	Count	Rate per second	Total Count
2021-4-18 14:28:55	57	0.47498020915795175	1254058
2021-4-18 14:30:55	82	0.6833219446342561	1254140
2021-4-18 14:32:55	139	1.15826576783021	1254279
2021-4-18 14:34:55	57	0.4750593824228029	1254336
2021-4-18 14:36:55	82	0.6832877808146124	1254418
2021-4-18 14:38:55	139	1.1583236806359947	1254557
2021-4-18 14:40:55	57	0.47498020915795175	1254614
2021-4-18 14:42:55	82	0.6833333333333333	1254696
2021-4-18 14:44:55	139	1.1583526392106536	1254835
2021-4-18 14:46:55	57	0.4750079167986133	1254892
2021-4-18 14:48:55	82	0.6833731967698116	1254974
2021-4-18 14:50:55	139	1.1583526392106536	1255113
2021-4-18 14:52:55	57	0.47497625118744063	1255170
2021-4-18 14:54:55	82	0.6832991683749146	1255252
2021-4-18 14:56:55	139	1.1584491782511586	1255391
2021-4-18 14:58:55	57	0.47496437767167446	1255448
2021-4-18 15:00:55	82	0.6833675017084188	1255530
2021-4-18 15:02:55	139	1.1583236806359947	1255669
2021-4-18 15:04:55	57	0.4749960416996525	1255726
2021-4-18 15:06:55	82	0.6833105563147894	1255808
2021-4-18 15:08:55	139	1.1583236806359947	1255947
2021-4-18 15:10:55	57	0.47500395836631976	1256004

Below the performance data, there is a 'Messages' section showing several log entries related to listener connections for various components like ECMTS, IBM Sync, and IBM FileNet Content Platform Engine.

In this example you can also see the `Count`. The Count is calculated by the System Dashboard out of the current and previous `Total Count`. It cannot be used as parameter info in the Listener path in ESM. The information available is:

- min
- max
- standard deviation
- average
- value
- total count

Check the System Dashboard which of those is available for the object.

Hint:

The Listener object in most cases just contains raw data information, the counter is reset to 0 once the software is restarted in this case. In the example above the Total Count will be 0 after the CPE is started. This also leads to the fact that some of the objects are not available at the start of the software and will only become available once such an object was counted. E.g. /USER/BASEOS/Document/Delete instead of Load

4. Setting up a Listener probe in ESM

As already mentioned some of the information given above is either needed in the subsystem already or during the probe setup.

CPE Listener probes using predefined information:

There are some CPE specific Listener probes. Therefore some CPE Listener information needs to be defined in the subsystem already:

Subsystem

Specify the type of subsystem you want to configure and fill in the parameters. If applicable you can configure related subsystems on the following pages

Note: when switching panels in this editor changes will be persisted automatically.

Subsystem:	ContentPlatformEngine
Name:	CPE55
Listener Application:	"IBM FileNet Content Platform Engine"
Listener Instance:	*
Listener Application Version:	
URL:	http://smcw2k16p8-55:9080/wsi/FNCEWS40MTOM
User:	p8admin
Password:	unchanged

As you can see, * is given in this example for the "Listener Instance" and "Listener Application Version". This is possible in this case as only one CPE Instance is running on my CPE Server. If you have several instances running, specify the instance and or version here as well. Otherwise the information is requested from all of the instances.

Setting up the probes:

There are some predefined probes for the CPE Subsystem where the full object path is automatically set. In some cases additional probe parameters like the Objectstore name needs to be given, as this is a part that can differ in the Listener object. The predefined probes are:

- CeObjectstoreChangesFailed
- CeObjectstoreClbInterestBaseObjectsFailed
- CeObjectstoreClbNotificationItemObjectsFailed
- CeObjectstoreClbNotificationUserObjectsFailed
- CeObjectstoreClbSeedlistEntryObjectsFailed
- CeObjectstoreClbShareObjectsFailed
- CeObjectstoreIndexingExceptions
- CeObjectstorePublishRequestsFailed
- CeObjectstoreQueueItemsFailed
- CeObjectstoreSweepActivityQueueFailed
- CeObjectstoreSweepAnnotationsFailed
- CeObjectstoreSweepBackgroundSearchesFailed
- CeObjectstoreSweepDocVersionsFailed
- CeObjectstoreSweepEmailQueueFailed
- CeObjectstoreSweepFrameworkExceptions
- CeObjectstoreSweepNotificationQueueFailed
- CeObjectstoreSweepThumbnailRequestFailed
- CeObjectstoreSweepTranscriptionRequestFailed
- CeRpcRequestsFailed

Example (CeObjectstoreChangesFailed):

Probe Configuration

Select the subsystem you want to monitor and the type of the probe that should be used.

Subsystem:

Probe:

Name:

active:

Objectstore:

Agent Assignment

Select the location where the Probe should run.

Schedule

This probe is connected permanently and listens for external events.

Furthermore there are some CPE Listener Probes where the first part of the object path is already given. The partially predefined probes are:

Probe	Related Object
ListenerCEUser	/USER/<ObjectStore>/
ListenerCpu	/CPU/
ListenerDisk	/Disk/
ListenerNetwork	/Network/
ListenerUser	/USER/
ListenerUserServer	/USER/Server Based Counters/
ListenerUserStatic	/USER/Static

Additional parameter for the probes are:

- <ObjectStore> in case of the ListenerCEUser
- Param Prefix
- Parameter To Check

This combined reflects the following Example (ListenerCEUser):

`/USER/<object_store>/<param_prefix>/<param_to_check>/Total Count`

Probe Configuration

Select the subsystem you want to monitor and the type of the probe that should be used.

Subsystem:

Probe:

Name:

active:

Objectstore:

Param Prefix:

Param To Check:

Agent Assignment

Select the location where the Probe should run.

Schedule

This probe is connected permanently and listens for external events.

Hint:

Unfortunately the requested information for the partially predefined probes is always “Total Count” which does not fit all of the possible objects and therefore these probes are not usable in such a case. Please check the “System Dashboard” first if it fits.

Listener Probes without predefinition:

There are two probes (Listener and ListenerForLicences) for the subsystem type Listener. For both the “Application Name”, “Application Instance” and “Application Version” can be given in the probe setup.

The “ListenerForLicences” probe requests a specific object and therefore in addition only the “Request Type” must be given. No “Listener Path” is needed.

Example (ListenerForLicences):

The screenshot shows the 'Probe Configuration' window. It includes a title bar with a refresh icon and the text 'Probe Configuration'. Below the title, there is a subtitle: 'Select the subsystem you want to monitor and the type of the probe that should be used.' The configuration fields are: 'Subsystem' (dropdown menu with 'LocalListener' selected), 'Probe' (dropdown menu with 'ListenerForLicences' selected), 'Name' (text field with 'ListenerForLicences @ LocalListener'), 'active' (checkbox checked), 'Request Type' (dropdown menu with 'LICENCE' selected), 'Application Name' (text field with 'IBM FileNet Content Platform Engine'), 'Application Instance' (text field with '*'), and 'Application Version' (text field with '*'). Below these fields is the 'Agent Assignment' section with the subtitle 'Select the location where the Probe should run.' and a dropdown menu with 'CPE55_Agent' selected. At the bottom, there is a 'Schedule' section with a calendar icon and the text 'This probe is connected permanently and listens for external events.'

The Listener probe is the generic implementation of all of the other “Listener” probes. The full “Listener Path” must be given here.

Example (Listener):

The screenshot shows the 'Probe Configuration' window. It includes a title bar with a refresh icon and the text 'Probe Configuration'. Below the title, there is a subtitle: 'Select the subsystem you want to monitor and the type of the probe that should be used.' The configuration fields are: 'Subsystem' (dropdown menu with 'LocalListener' selected), 'Probe' (dropdown menu with 'Listener' selected), 'Name' (text field with 'Listener @ LocalListener'), 'active' (checkbox checked), 'Application Name' (text field with 'IBM FileNet Content Platform Engine'), 'Application Instance' (text field with '*'), 'Application Version' (text field with '*'), and 'Listener Path' (text field with '/CPU/jvm_free_memory/total count'). Below these fields is the 'Agent Assignment' section with the subtitle 'Select the location where the Probe should run.' and a dropdown menu with 'CPE55_Agent' selected. At the bottom, there is a 'Schedule' section with a calendar icon and the text 'This probe is connected permanently and listens for external events.'

Setup of the “Evaluation” (Page 2 of the probe config) for all Listener probes:

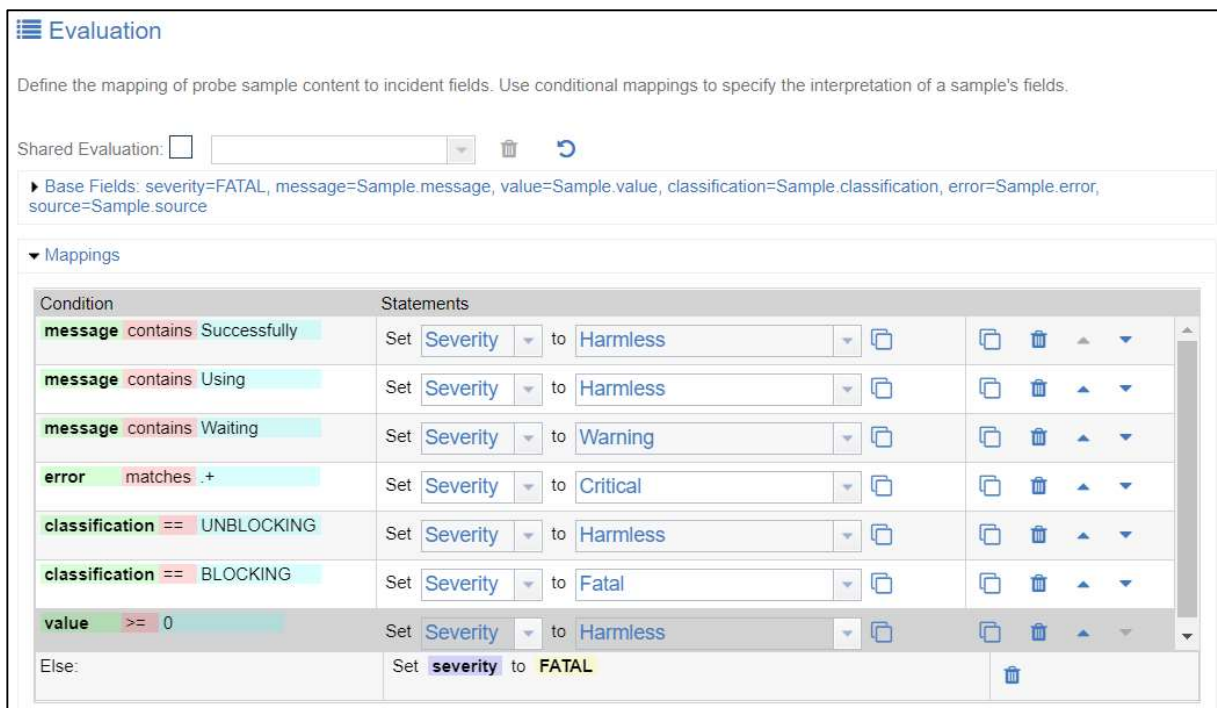
One thing all of the Listener probes have in common is the setup of the evaluation. No matter which kind of value mapping you want to use, at the beginning of the table you should add three additional entries for message mapping.

message contains Successfully → Set Severity to Harmless

message contains Using → Set Severity to Harmless

message contains Waiting → Set Severity to Warning

Example:



The screenshot shows the 'Evaluation' configuration page. It includes a 'Shared Evaluation' checkbox and a dropdown menu. Below this, there is a list of 'Base Fields' and a 'Mappings' section. The 'Mappings' section contains a table with the following entries:

Condition	Statements
message contains Successfully	Set Severity to Harmless
message contains Using	Set Severity to Harmless
message contains Waiting	Set Severity to Warning
error matches .+	Set Severity to Critical
classification == UNBLOCKING	Set Severity to Harmless
classification == BLOCKING	Set Severity to Fatal
value >= 0	Set Severity to Harmless
Else:	Set severity to FATAL

The three entries reflect the three possible status information of the Listener probes that will be send at the startup of the agent. The agent connects to the Listener port and then sends a status information:

Successfully connected ... → The agent could connect to the Listener ...

Using existing ... → The agent uses an already established connection ...

Waiting for ... → The agent was not able to establish the connection so far ...

Select from the drop down between “First” or “Always”.

The actual incident for the probe can be reviewed as usual in the “Monitoring Dashboard”.

5. Aggregation of Listener information in reports and showing the report in the monitoring dashboard

As mentioned above, most of the Listener object will be reset to 0 at software startup. Due to that the generated report might look odd at the time the software was restarted.

Since the information is usually only a counted value after the restart of the software, ESM offers the possibility of aggregation, so far delta aggregation is possible. ESM uses the report functionality to create a report which then can be shown in the monitoring dashboard.

The “Monitoring Dashboard” always shows the latest report that was created. Meaning if the report is created e.g. every 2 minutes, the portlet in the dashboard is also updated also every 2 minutes.

Example setup of the Report:

Report Configuration

Enter a name, select one or more situations and the type of the chart that should be used.

*Report Name:	Object Request Delta
Type:	Aggregation
*Situations:	FindObjectCalls @ IS42 FullTextSearchServerStatus @ CMoD101 FullTextSearchServerStatus @ CMoD105 FullTextSearchServerStatus @ CMoD105 FullTextSearchStatistics @ CMoD101 FullTextSearchStatistics @ CMoD105 GeneralDiskSpace @ CMoD105_Agent GeneralDiskSpace @ CSM_Server Get Object Requests -Total Count @ CPE
Smoothing:	None
Aggregation Type:	Delta
Aggregation Period:	Hour(s)

Schedule

Server time is used but browser displays local time.

active:	<input checked="" type="checkbox"/>
*Start Date:	01.02.2021 00:00
E-mail Recipients:	Enter E-mail Recipients
SMTP Server:	
Period:	12 Hour(s)
Interval:	2 Minute(s)

Description for this kind of report:

Type: Must be set to “Aggregation”.

Situation: Select only one Situation → Only one will be shown in the dashboard portlet

Aggregation Type: Currently only “Delta” is possible

Aggregation Period: Select the aggregation period from the drop down. You can choose between "Auto", "Hours(s)", "Day(s)" and "Week(s)". This defines the period for which the aggregation is created. It is always the full hour, day and week. Auto will automatically define the period, it depends on the period that is given for the schedule of the report.

Active: This report must be scheduled

Start Date: Defines the date for the first generation of the report – could be in future or past.

Period: The period defines the timeframe for the reported graph. Once the scheduler starts it is going back the interval and reports over this period. Possible selections are.

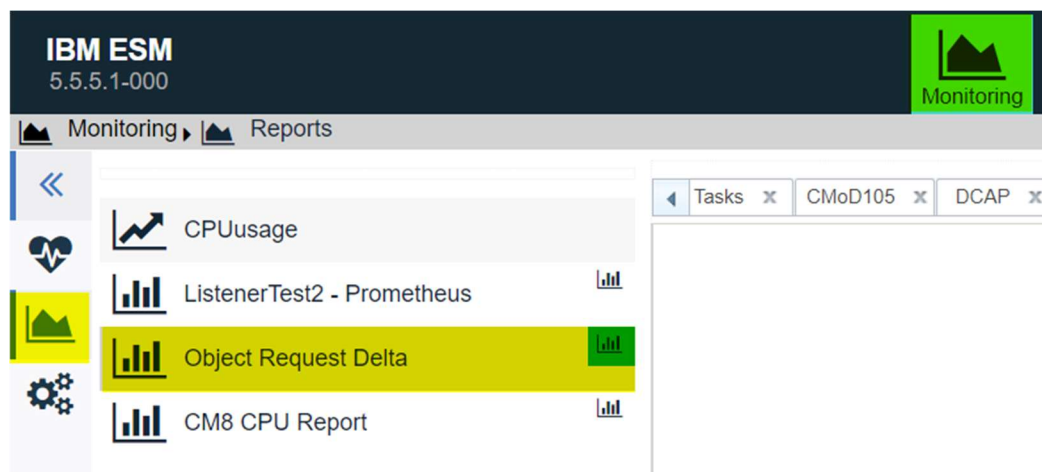
- Hours(s)
- Day(s)
- Week(s)
- Month(s)
- Year(s)

Interval: Defines the timeframe for the next run of the report. Possible selections are.

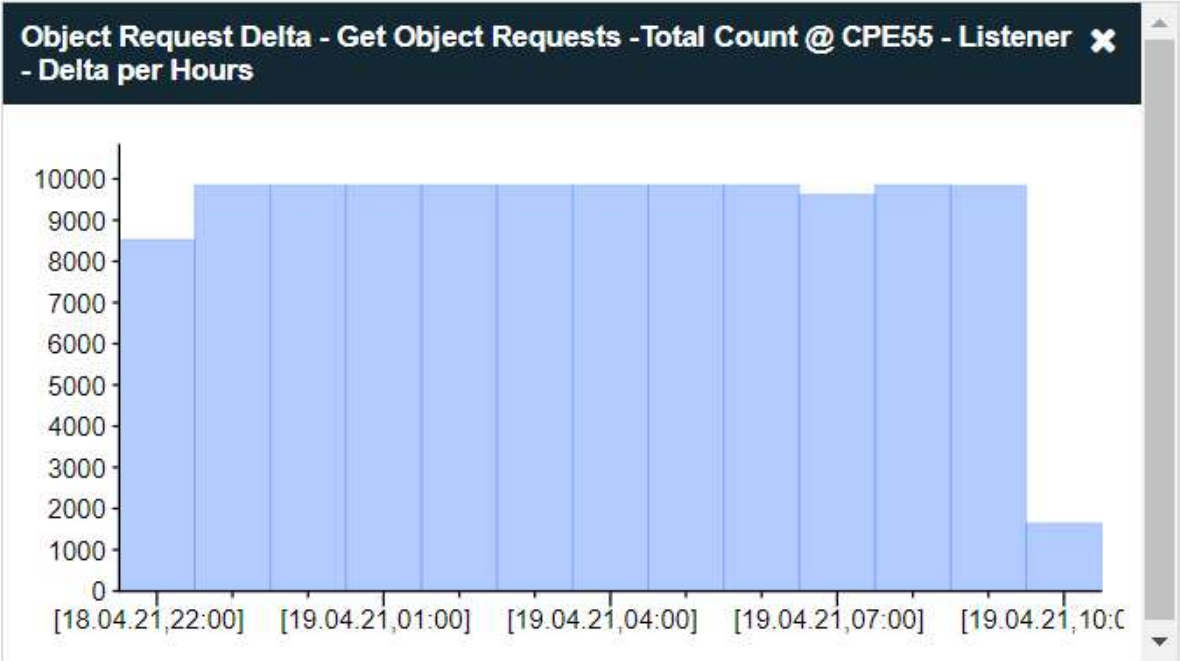
- Second(s)
- Minute(s)
- Hours(s)
- Day(s)
- Week(s)
- Month(s)
- Year(s)

These settings will generate a new report every 2 minutes (Interval). The report will contain 12 hours of data (Period). For each hour (Aggregation Period) on bar is created. For the current "Aggregation Period" the delta is created out of the current information and last information from the previous "Aggregation Period" (Hour in this case). For the previous (completed) "Aggregation period(s)" the delta is always created out of the last information from those.

After the report is generated it can be added to the monitoring dashboard. Select the "Monitoring Dashboard" and the "Report" from the left tree. All reports that are from type aggregation have the add button on the right.



The created report portlet looks like this:



Contact Information

If you have any questions, please contact us at ECM.SystemMonitor@cenit.com.

CENIT AG
Phone: +49 711 7825 30
Email: ECM.SystemMonitor@cenit.com