

Job Execution Service (JES) Overview

For Cora SeQuence and Cora Orchestration

April 2022

Notice

All rights reserved. No part of this document (including the text, images, graphics or the selection or arrangement of content, which forms an original compilation), may be reproduced or transmitted in any material form or by any means, electronic or mechanical, including photocopying or recording in any medium (whether or not transiently) without the written permission of the copyright holder. Such written permission must also be obtained before any part of this publication is stored in a retrieval system of any nature. The same shall apply to the export of this publication from India, and a violation of this condition will lead to civil and criminal prosecution.

Table of contents

General document information.....	4
Purpose.....	4
Prerequisites.....	4
Introduction	4
JES deployment topologies	5
Managing JES jobs.....	5
Job status	6
JES jobs	6
Archiving jobs (V9.4) (Not supported in Orchestration Services).....	6
Elasticsearch jobs (V9.2) (Not supported in Orchestration Services)	7
Email Sender (V9.4) (Not supported in Orchestration Services).....	7
File Migration (V9.2) (Not supported in Orchestration Services).....	7
Personal Data Purge (V9.7) (Not supported in Orchestration Services).....	7
Purge Files (V9.7) (Not supported in Orchestration Services).....	8
Start Workflow.....	8
Workflow Redirect (Preview) (Not supported in Orchestration Services)	8
Email Listener (V7.9).....	8
Kafka Subscriber (V9.2.2).....	8
Azure Service Bus (V8.8)	8
File Listener	8
Message Bus Listener	9
Multiple JES deployment for Cora SeQUENCE	9
Troubleshooting.....	9
SVC log files	9
JES performance counters (Not supported in Cora Orchestration).....	10
JES stops.....	10
JES won't start	10
Security exception in reminder thread	10
Email sending issues (Not supported in Cora Orchestration)	10

General document information

Purpose

This document provides an overview of the Job Execution Service (JES) and lists the jobs it executes.

Prerequisites

You should have access to the server on which JES is deployed.

Introduction

The Job Execution Service (JES) is an integral part of the Cora SeSequence and Cora Orchestration functionality. JES is deployed as part of the installation and continues to run in the background at all times.

JES is a Microsoft Windows service that executes several jobs, each of which performs one of the following actions:

Action	Supported in Cora SeSequence and Cora Orchestration	Supported in Orchestration Services
Collect and archive data	Yes	No
Synchronize the database with Elasticsearch database	Yes	No
Send conversation emails	Yes	No
Migrate files	Yes	No
Start a workflow	Yes	Yes
Listen to emails	Yes	Yes
Listen to a Kafka Subscriber	Yes	Yes
Listen to Azure Service Bus queues	Yes	Yes
File Listener	Yes	Yes

Each of the above jobs includes configuration options, such as how often and when to perform the action.

Unlike the BRS threads that are configured via the BRS config file, all the above jobs are configured from the **Cora SeSequence and the Cora Orchestration Admin site**.

The configuration options are described in the following sections.

JES deployment topologies

There can be several instances of the JES. They can all be deployed on a dedicated server or divided into separate servers.

A new JES instance is registered automatically. The system is updated automatically when services are shut down or go offline, without any need for manual action.

Multiple JES deployments provide the following configuration options:

- Job redundancy:** ensures that jobs run even when the hosting Job Execution Service (JES) is down.

When the JES that hosts the job is down, the redundant job is reassigned to another JES instance. Consequently, the job keeps its availability and ability to run.

A redundant job is always assigned to a JES host with the least amount of assigned jobs.
- You can define job redundancy and load balancing in systems that include more than one server running a JES.

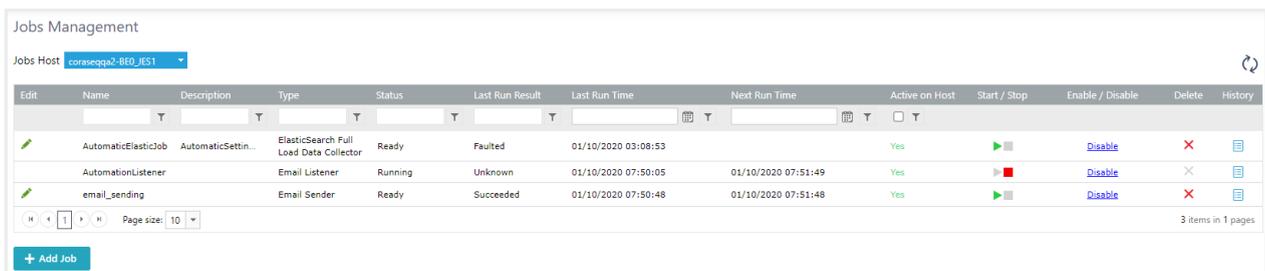
All job types support redundancy.
- Job scaling:** enables higher job throughput and improves system performance. You can set up scaling for *Service Bus Listener* and *Kafka Subscriber* jobs. With scaling you can have multiple JES instances running the same job simultaneously.

For all the configuration details of job scaling and redundancy, click [here](#).

Managing JES jobs

You can control all the JES jobs from the Jobs Management screen.

In the Administration site navigate to **Administration>Global Settings>Jobs Management**.



Edit	Name	Description	Type	Status	Last Run Result	Last Run Time	Next Run Time	Active on Host	Start / Stop	Enable / Disable	Delete	History
	AutomaticElasticJob	AutomaticSettin...	ElasticSearch Full Load Data Collector	Ready	Failed	01/10/2020 03:08:53		Yes		Disable		
	AutomationListener	Email Listener	Email Listener	Running	Unknown	01/10/2020 07:50:05	01/10/2020 07:51:49	Yes		Disable		
	email_sending	Email Sender	Email Sender	Ready	Succeeded	01/10/2020 07:50:48	01/10/2020 07:51:48	Yes		Disable		

Page size: 10 | 3 items in 1 pages

[+ Add Job](#)

Here you can manage (start/stop, enable/disable) existing jobs running on a specific host and monitor their status. You can also add new jobs.

Job status

The table below lists all possible job statuses:

Status	Description
Disabled	When a job is disabled. This is the default status of a new JES job when the “Job is enabled” check box is not selected.
Queued	When a scheduler has triggered a job execution and the job is waiting for a free thread to run on.
Ready	When a job is waiting for a scheduler to trigger its execution.
RecyclePending	When a job definition is changed, but it can’t be refreshed because the job is still running.
Running	When a job is running.
Suspended	When a job is suspended on the current host due to redundancy.
Unknown	When Cora SeSequence is not able to determine a job’s status.

JES jobs

The JES jobs listed below appear in the same order of the “Add Job” window. The last three jobs (Email Listener, Kafka Subscriber and Azure Service Bus) are activities that you configure in the App studio.

Archiving jobs (V9.4) (Not supported in Orchestration Services)

Jobs required when you set up the archiving configuration.

Archiving Data Collector

The long-term archiving feature extracts the closed workflow instances' data from the Cora SeSequence and Cora Orchestration database and moves it to the archive database. This job collects data into an archiving queue according to configuration filters options.

Archiving Data Worker

This job executes archiving on collected data queue (created by Archiving Data Collector Job) to a separate database defined in connection string.

Archiving Files Worker

This job executes files archiving on collected data queue (created by Archiving Data Worker Job) to a separate database or connection defined in connection string.

For all the configuration details of the archiving jobs, click [here](#).

Elasticsearch jobs (V9.2) (Not supported in Orchestration Services)

Jobs required when you integrate with the Elasticsearch server.

Elasticsearch Full Load Data Collector

The search feature provides Google-like search capabilities. When search is enabled, users can retrieve any data related to processes and cases.

The search feature uses the Elasticsearch engine.

This job synchronizes between database and Elasticsearch.

Elasticsearch Incremental Load Data Collector

This job synchronizes between the Cora SeQUENCE and Cora Orchestration database and the Elasticsearch database.

For all the configuration details of the Elasticsearch jobs, click [here](#).

Email Sender (V9.4) (Not supported in Orchestration Services)

This job sends conversation mails from Cora SeQUENCE and Cora Orchestration, as part of the Conversations View. It enables users to compose and send ad-hoc emails from within forms, and the Conversations View component helps display the incoming and outgoing emails.

For all the configuration details of the email sender job, click [here](#).

File Migration (V9.2) (Not supported in Orchestration Services)

This job migrates files to a chosen file storage.

After you set up a new storage location for attachment files, only new files uploaded to the system are stored in the new location. Existing files remain in the previous storage location.

Migrating the existing files to the new storage location is a recommendation for all systems that upgrade from previous product versions to make sure that all attachment files are managed from a central location.

During migration, the system transfers all binary records from the attachments tables (old and new) to the new storage location.

For all the configuration details of the file migration job, click [here](#).

Personal Data Purge (V9.7) (Not supported in Orchestration Services)

This job purges the marked personal data from the tables.

1. Go to **Administration > Global Settings > Jobs Management**.
2. Click **+Add Job**.

3. In the Create New Job window, select the **Personal Data Purge** job.
4. Click **Create**.
5. Configure the job like any other job and click **Finish**.

Purge Files (V9.7) (Not supported in Orchestration Services)

This job purges files from the database.

Start Workflow

This job starts a new workflow according to a schedule.

You need to provide the workflow that will be started and the schedule of the job.

Workflow Redirect (Preview) (Not supported in Orchestration Services)

Email Listener (V7.9)

This job listens to an email inbox and starts a dynamic workflow when a new email arrives to the inbox. You can then use the returned email object in other workflow stages.

For all the configuration details of the email listener job, click [here](#).

Kafka Subscriber (V9.2.2)

This job connects to Kafka and retrieves the message from a topic. Each subscriber connects to a specific topic. The Kafka Subscriber is a JES job.

Apache Kafka[®] is a distributed streaming platform that is designed to be fast, scalable, and durable. Kafka is generally used to move data between systems or applications and to enable applications to consume data required to perform specific actions.

Kafka activities are configured to set up a messaging mechanism within Cora SeQUENCE and Cora Orchestration or between system and other applications.

For all the configuration details of the Kafka subscriber job, click [here](#).

Azure Service Bus (V8.8)

The Service Bus Listener job is a messaging mechanism to integrate your workflow with other applications or services. This activity integrates with Microsoft Azure Service Bus, which is a cloud-based messaging service used to connect and send information between applications, services and devices.

For all the configuration details of the service bus listener job, click [here](#).

File Listener

File Listener is a JES job to listen to an external file storage system to determine whether a certain file exists.

For all the configuration details of the File listener job, click [here](#).

Message Bus Listener

Message Bus Listener is a JES job to get messages from the internal queue from within workflows.

For all the configuration details of the Message bus listener job, click [here](#).

Multiple JES deployment for Cora SeQUENCE

It's possible to deploy multiple JES instances on different servers (or on the same server) that will provide redundancy and scaling to the relevant jobs, depending on the job type and configuration.

Multiple JES topology reduces the redirection and progress time of workflow instances, which is important for large capacity installations.

To deploy multiple JES, follow the steps below:

1. Deploy a JES app on the first server using the PowerShell function.
2. For additional JES on the same server:
 - a. Deploy a second JES app using the PowerShell function. Provide an instance number using the ApplicationID parameter.
 - b. Repeat step a as many times you'll need.
3. For additional JES on a different server, repeat step 1.
4. Run the Set-CoraSeQUENCEConnectionString function (on each server) to configure all apps connection string together.

Troubleshooting

Here are some tools and techniques to troubleshoot common JES issues.

SVC log files

A generic tool that can help you with the troubleshooting is the workflow execution service trace file, also called the **SVC log file**.

This log is generated during runtime for every workflow execution and can be found in the Cora SeQUENCE server. It could be very helpful when investigating issues either on the runtime level or on the job host level.

Note that you must enable the logs in the JES config file and specify their path. You locate the relevant SVC log according to the workflow execution instance number and open it with the Microsoft Service Trace Viewer.

IMPORTANT

Remember to stop the log file execution after you complete the troubleshooting process, since creating the SVCLOG files consumes server resources.

For all the configuration details of SVC log file, click [here](#).

JES performance counters (Not supported in Cora Orchestration)

Use performance counters to investigate issues related to the execution of workflow instances, and JES performance.

For configuration details, click [here](#).

JES stops

If the JES stops unexpectedly, it might indicate a host environment issue. We recommend looking at the server logs to identify the root cause.

JES won't start

If you try to start the JES service and it doesn't start, there may be several reasons:

- The user that is defined to run the JES during installation doesn't have permissions to run windows services.
- The JES can't connect to the Cora SeQUENCE Database using the connection string you provided during deployment Either the connection string is wrong, or there is a problem with the Database server.
- The password of the user that is defined to run the JES has changed.

Security exception in reminder thread

If you have received a security exception in the redirect or reminder thread, check that the user is a valid Cora SeQUENCE user.

Email sending issues (Not supported in Cora Orchestration)

To troubleshoot email sending failures, check that the SMTP port (25) is open on the server and check the SMTP server configuration.

Another common mistake is that emails are not defined for the activity.