Digital Channels Private Edition Guide

Configure AI Connector

5/12/2026

Contents

Complete the steps on this page to configure your AI Connector deployment.

Related documentation:

-
-
-
-

RSS:

- For private edition

Override Helm chart values

You can specify parameters for the deployment by overriding Helm chart values in the **values.yaml** file. See the **Parameters** table for a full list of overridable values.

For more information about how to override Helm chart values, see Overriding Helm chart values.

Parameters

Parameter	Description	Valid values	Default
version	The AI Connector version.	A valid release version	nil
nameOverride	A string to partially override the athena.fullname template. This string is prepended to the release name.	String	nil
fullnameOverride	A string to fully override the athena.fullname template.	String	nil
image.registry	The AI Connector image registry.	A valid registry URL	nil
image.repository	The AI Connector image name.	A valid image name	nexus/athena
image.pullPolicy	Specifies when Kubernetes pulls images from the registry on start up.	IfNotPresent or Always	IfNotPresent

Parameter	Description	Valid values	Default
image.pullSecrets	An array of docker-registry secret names.	An array of secret names	[] (does not add image pull secrets to deployed pods)
serviceAccount.create	Specifies whether a service account must be created.	false or true	false
serviceAccount.name	The name of the service account to use. If this is not set and create is true, a name is generated using the fullname template.	String	""
serviceAccount.annotations	Annotations to add to the service account.	A valid set of labels as "name: value"	{}
podAnnotations	Custom annotations for each pod.	A valid set of labels as "name: value"	{}
podSecurityContext.runAsNonRoot	Specifies whether the container must run as a non-root user.	true or false	true
podSecurityContext.runAsUser	The user ID to run the entry point of the container process.	A valid user ID	500
podSecurityContext.runAsGroup	The group ID to run the entry point of the container process.	A valid group ID	500
podSecurityContext.fsGroup	A supplemental group ID that applies to all containers in a pod.	A valid group ID	500
configChecksum	Adds SHA-256 checksum of the ConfigMap to the deployment annotations.	true of false	true
secretChecksum	Adds SHA-256 checksum of the Secret to the deployment annotations.	true or false	true
containerPort	TCP port the service is listening on.	A valid port	4084
service.enabled	Enables the Kubernetes service.	true or false	true
service.type	The Kubernetes service type.	See the Kubernetes documentation for details.	ClusterIP
service.annotations	The service annotations.	A valid set of annotations as "name: value"	{}
service.port	The Kubernetes service	A valid port	80

Parameter	Description	Valid values	Default
	HTTP port.		
ingress.enabled	Enables the ingress controller resource.	true or false	false
ingress.annotations	The ingress annotations.	A valid set of annotations as "name: value"	[]
ingress.hosts[0].host	The hostname of your AI Connector installation.	A valid hostname	athena.local
ingress.hosts[0].paths	The paths (within the URL structure) to your AI Connector.	A valid list of paths	[]
ingress.tls[0].secretName	Kubernetes secret name with server.crt certificate and server.key private key file (Only required if you want to configure TLS for ingress resources).	A valide name	nil
ingress.tls[0].hosts	The hostname used to generate cert that matches the TLS certificate (required if you want to configure TLS for ingress resources).	A valid hostname	nil
resources	The requests and limits for CPU and memory usage in Kubernetes. See the Kubernetes documentation for details.	Object	{}
nodeSelector	The labels Kubernetes uses to assign pods to nodes. See the Kubernetes documentation for details.	Object	{}
tolerations	The tolerations Kubernetes uses for advanced pod scheduling. See the Kubernetes documentation for details.	Object	{}
affinity	Specifies the affinity and anti-affinity for Digital Channels pods. See the Kubernetes documentation for details.	Object	{}

Parameter	Description	Valid values	Default
priorityClassName	The class name Kubernetes uses to determine the priority of a pod relative to other pods. See the Kubernetes documentation for details.	A valid priority class name	""
monitoring.enabled	Specifies whether to deploy Custom Resource Definitions (CRD) for ServiceMonitors to determine which services should be monitored.	true or false	false
athena.server.apiPrefix	The prefix for all API endpoints exposed by the service.	String or comma-separated list of strings	/athena/v1,/nexus/v3
athena.nexus.url	The Nexus service URL.	A valid URL	http://nexus-production.nexus.svc.cluster.local
athena.nexus.apiPrefix	The prefix of Nexus API endpoints.	String	/nexus/v3
athena.nexus.apiKey	A Nexus API key used by AI Connector to fetch tenant settings and services from Nexus. See the Digital Channels documentation for details.	A valid cluster Nexus API key with permissions "nexus:cluster:*" assigned	""
athena.nexus.timeout	Timeout for requests to Nexus (access tokens validation, configuration retrieval).	Integer	10000
athena.db.host	The Postgres service URL.	A valid URL	postgres
athena.db.port	The Postgres service port.	A valid port	5432
athena.db.user	The user assigned for the AI Connector application to access Postgres.	A valid user	nexus
athena.db.password	The password assigned for the AI Connector application to access Postgres.	A valid password	""
athena.db.database	The database for the AI Connector application to use in Postgres.	A valid database	nexus

Parameter	Description	Valid values	Default
athena.db.ssl	Use secured connection to Postgres.	true or false	false
athena.redis.nodes	A comma-separated list of Redis nodes to connect.	A valid URL	redis://redis:6379
athena.redis.password	The password for Redis authentication.	A valid password	""
athena.redis.cluster	Specifies whether to deploy Redis as a cluster.	true or false	false
athena.redis.tls	Specifies whether to use TLS on the Redis connection.	true or false	false
athena.google.speechApiKey	API key for Google speech recognition API used to transcribe speech to text for bot providers that do not natively support voice input (LUIS, DialogEngine), not needed for Dialogflow ES, CX and Lex.	A valid Google speech recognition API key	""

Configure security

To learn more about how to configure security for private edition, be sure to read Permissions and OpenShift security settings.

The security context settings define the privilege and access control settings for pods and containers.

By default, the user and group IDs are set in the **values.yaml** file as `500:500:500`, meaning the **genesys** user.

```
podSecurityContext:
    runAsUser: 500
    runAsGroup: 500
    fsGroup: 500
    runAsNonRoot: true
```