

Configuring NiFi CR

Date published: 2024-06-11

Date modified: 2024-06-11



Legal Notice

© Cloudera Inc. 2024. All rights reserved.

The documentation is and contains Cloudera proprietary information protected by copyright and other intellectual property rights. No license under copyright or any other intellectual property right is granted herein.

Unless otherwise noted, scripts and sample code are licensed under the Apache License, Version 2.0.

Copyright information for Cloudera software may be found within the documentation accompanying each component in a particular release.

Cloudera software includes software from various open source or other third party projects, and may be released under the Apache Software License 2.0 (“ASLv2”), the Affero General Public License version 3 (AGPLv3), or other license terms. Other software included may be released under the terms of alternative open source licenses. Please review the license and notice files accompanying the software for additional licensing information.

Please visit the Cloudera software product page for more information on Cloudera software. For more information on Cloudera support services, please visit either the Support or Sales page. Feel free to contact us directly to discuss your specific needs.

Cloudera reserves the right to change any products at any time, and without notice. Cloudera assumes no responsibility nor liability arising from the use of products, except as expressly agreed to in writing by Cloudera.

Cloudera, Cloudera Altus, HUE, Impala, Cloudera Impala, and other Cloudera marks are registered or unregistered trademarks in the United States and other countries. All other trademarks are the property of their respective owners.

Disclaimer: EXCEPT AS EXPRESSLY PROVIDED IN A WRITTEN AGREEMENT WITH CLOUDERA, CLOUDERA DOES NOT MAKE NOR GIVE ANY REPRESENTATION, WARRANTY, NOR COVENANT OF ANY KIND, WHETHER EXPRESS OR IMPLIED, IN CONNECTION WITH CLOUDERA TECHNOLOGY OR RELATED SUPPORT PROVIDED IN CONNECTION THEREWITH. CLOUDERA DOES NOT WARRANT THAT CLOUDERA PRODUCTS NOR SOFTWARE WILL OPERATE UNINTERRUPTED NOR THAT IT WILL BE FREE FROM DEFECTS NOR ERRORS, THAT IT WILL PROTECT YOUR DATA FROM LOSS, CORRUPTION NOR UNAVAILABILITY, NOR THAT IT WILL MEET ALL OF CUSTOMER’S BUSINESS REQUIREMENTS. WITHOUT LIMITING THE FOREGOING, AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, CLOUDERA EXPRESSLY DISCLAIMS ANY AND ALL IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, QUALITY, NON-INFRINGEMENT, TITLE, AND FITNESS FOR A PARTICULAR PURPOSE AND ANY REPRESENTATION, WARRANTY, OR COVENANT BASED ON COURSE OF DEALING OR USAGE IN TRADE.

Contents

Configuring a NiFi instance.....	5
Resource recommendations for NiFi deployments.....	5
Example NiFi cluster size.....	5
Group, version, kind, meta.....	8
Configuring NiFi Image.....	8
Configuring cluster size.....	8
Configuring bootstrap JVM settings.....	9
Configuring persistence.....	9
Configuring assets.....	10
Configuring Kubernetes state management.....	11
Configuring node certificate generation.....	11
Configuring NiFi properties.....	12
Overriding NiFi settings using ConfigMaps and Secrets.....	12
Configuring scaling.....	13
Configuring pod affinity.....	13
Configuring connections to NiFi and NiFi Registry.....	14
Configuring session affinity.....	14
Configuring arbitrary connections.....	14
Configuring NiFi Web UI connection.....	15
Hostname-only ingress example.....	15
Hostname-only route example.....	16
Ingress with context path example.....	16

Configuring authentication for NiFi and NiFi Registry..... 16
 Configuring single user authentication..... 17
 Configuring LDAP authentication.....17

Example CR..... 19

Configuring a NiFi instance

NiFi instances are configured through the CRs used to deploy them.

A custom resource (CR) is a YAML file that describes your desired NiFi deployments. This single file contains all configuration information required for the NiFi instance, no additional configuration is required after deployment.

This documentation provides sample configuration code snippets to help you create a CR.

Resource recommendations for NiFi deployments

Learn about the recommended resource sizes for NiFi deployments. Every NiFi deployment is unique on the basis of the purpose it serves, therefore the values here are just recommendations not requirements. Actual values may substantially differ depending on your use case.

Resource Type	Amount
CPU	2+ per Pod
Memory	4Gi+ per Pod
PVC/PV	5 per Pod
Secrets	4 + #Pods
ConfigMaps	13
Services	1 + #Connections
Pods	1 min, 3+ recommended
StatefulSet	1
Deployment	0
Ingress	1 + #IngressConnection

Example NiFi cluster size

The following list of resources represents the whole of a deployed NiFi cluster managed by the Cloudera Flow Management - Kubernetes Operator. The following example is run in kind with cert-manager and ingress-nginx deployed as dependencies.

```
apiVersion: cfm.cloudera.com/v1alpha1
kind: Nifi
metadata:
  name: mynifi
spec:
  replicas: 3
  nifiVersion: 1.0.0
  image:
    repository: container.repository.cloudera.com/cloudera/cfm-nifi-k8s
    tag: 2.8.0-b94-nifi_1.25.0.2.3.13.0-36
    pullSecret: cloudera-container-repository-credentials
    pullPolicy: IfNotPresent
  tiniImage:
    repository: container.repository.cloudera.com/cloudera/cfm-tini
    tag: 2.8.0-b94
    pullSecret: cloudera-container-repository-credentials
    pullPolicy: IfNotPresent
  persistence:
```

```

size: 1Gi
contentRepo:
  size: 1Gi
flowfileRepo:
  size: 1Gi
provenanceRepo:
  size: 2Gi
data: {}
security:
  initialAdminIdentity: anonymous
  nodeCertGen:
    issuerRef:
      name: self-signed-ca-issuer
      kind: ClusterIssuer
  singleUserAuth:
    enabled: true
    credentialsSecretName: creds
hostName: nifi.io
uiConnection:
  type: Ingress
  annotations:
    nginx.ingress.kubernetes.io/affinity: cookie
    nginx.ingress.kubernetes.io/affinity-mode: persistent
    nginx.ingress.kubernetes.io/backend-protocol: HTTPS
    nginx.ingress.kubernetes.io/ssl-passthrough: "true"
    nginx.ingress.kubernetes.io/ssl-redirect: "true"
resources:
  nifi:
    requests:
      cpu: "1"
      memory: 2Gi
    limits:
      cpu: "4"
      memory: 4Gi
  log:
    requests:
      cpu: 50m
      memory: 128Mi

```

StatefulSet

```
$ kubectl get statefulset
```

NAME	READY	AGE
mynifi	3/3	24h

Pods

```
$ kubectl get pod
```

NAME	READY	STATUS	RESTARTS	AGE
mynifi-0	7/7	Running	0	23h
mynifi-1	7/7	Running	0	23h
mynifi-2	7/7	Running	0	23h

ConfigMaps

```
$ kubectl get configmap
```

NAME	DATA	AGE
mynifi-authorizers	1	24h
mynifi-authorizers-empty	1	24h
mynifi-bootstrap	1	24h
mynifi-certificate-setup-script	1	24h
mynifi-decommission-script	1	24h
mynifi-identities-config	1	24h
mynifi-logback	1	24h
mynifi-login-identity-providers	1	24h
mynifi-nifi-cli-properties	1	24h
mynifi-nifi-properties	1	24h
mynifi-start-script	1	24h
mynifi-state-management	1	24h

```

mynifi-stop-script          1          24h
Secrets
$ kubectl get secret
NAME                        TYPE            DATA    AGE
creds                      Opaque          2        27h
mynifi-0-node-cert         kubernetes.io/tls 5        24h
mynifi-1-node-cert         kubernetes.io/tls 5        23h
mynifi-2-node-cert         kubernetes.io/tls 5        23h
mynifi-keystorepassword    Opaque          1        24h
mynifi-proxy-cert          kubernetes.io/tls 3        27h
mynifi-sensitive-props-key Opaque          1        24h
Certificates
$ kubectl get certificates
NAME            READY    SECRET    AGE
mynifi-0-node-cert    True    mynifi-0-node-cert    24h
mynifi-1-node-cert    True    mynifi-1-node-cert    23h
mynifi-2-node-cert    True    mynifi-2-node-cert    23h
mynifi-proxy-cert     True    mynifi-proxy-cert     24h
Services
$ kubectl get service
NAME            TYPE            CLUSTER-IP    EXTERNAL-IP    PORT(S)
AGE
mynifi          ClusterIP    None          <none>          6007/TCP,5000/TCP    24h
mynifi-web      ClusterIP    10.96.28.159  <none>          8443/TCP
24h
Ingresses
$ kubectl get ingresses
NAME            CLASS    HOSTS    ADDRESS    PORTS    AGE
mynifi-web      <none>   nifi.io  localhost  80       24h

PersistentVolumeClaims
$ kubectl get persistentvolumeclaim
NAME            CAPACITY    ACCESS MODES    STATUS    VOLUME    STORAGECLASS    AGE
content-repository-mynifi-0  61fa4b  1Gi    RWO    Bound    pvc-d5b00d05-d8ee-4b5c-abe4-2cae61  standard    24h
content-repository-mynifi-1  480d4b31  1Gi    RWO    Bound    pvc-3a510ebf-2f63-409b-992b-5a08  standard    23h
content-repository-mynifi-2  ac994b  1Gi    RWO    Bound    pvc-flbaf8e-b8b2-485c-a0c8-ddb583  standard    23h
data-mynifi-0  9417a441  1Gi    RWO    Bound    pvc-67072ae8-b1ae-445c-b81a-0771  standard    24h
data-mynifi-1  506686  1Gi    RWO    Bound    pvc-4d20e11b-0d93-4b1b-95ff-a19189  standard    23h
data-mynifi-2  79896f59  1Gi    RWO    Bound    pvc-71b92ed3-3a70-4c4d-a848-44f8  standard    23h
flowfile-repository-mynifi-0  62c059  1Gi    RWO    Bound    pvc-c9c0ea11-0f59-4eeb-b9ac-a795b5  standard    24h
flowfile-repository-mynifi-1  f7d39a75  1Gi    RWO    Bound    pvc-58200919-b8b2-43be-8d2f-16b1  standard    23h
flowfile-repository-mynifi-2  709692  1Gi    RWO    Bound    pvc-869c0159-51c7-4857-aa52-8c775c  standard    23h
provenance-repository-mynifi-0  756fb344  2Gi    RWO    Bound    pvc-97dc6f41-b8ea-4682-9f60-74c8  standard    24h
provenance-repository-mynifi-1  53277b  2Gi    RWO    Bound    pvc-bd529e22-7024-4bef-a951-0a3fb7  standard    23h
provenance-repository-mynifi-2  e2bfdb13  2Gi    RWO    Bound    pvc-cbcc2b50-3404-4fea-8c3c-6a2f  standard    23h
state-mynifi-0  2599a4  1Gi    RWO    Bound    pvc-d0e72637-2b7d-40b7-a147-724094  standard    24h
state-mynifi-1  e794c671  1Gi    RWO    Bound    pvc-cl7f3fe2-b93d-4774-b49d-5564  standard    23h

```

state-mynifi-2	Bound	pvc-19a010f0-0397-496b-a77d-89944b
94a9b0 1Gi RWO	standard	23h

Related Information

[kind documentation](#)

Group, version, kind, meta

This is the initial section of your YAML file that you need to specify in all cases.

You need to add the following section to the top of each NiFi CR you write. It defines the group “cfm.cloudera.com”, the version “v1alpha1”, the kind “NiFi”, and the name of your cluster and the NiFi nodes. It can also specify the namespace in which resources will be deployed. It is expected that a single NiFi cluster is deployed in a given namespace. You can also specify namespace during deployment, if that is what you want, omit namespace from the CR

```
apiVersion: cfm.cloudera.com/v1alpha
kind: Nifi
metadata:
  name: [***NIFI CLUSTER NAME***]
  namespace: [***NIFI CLUSTER NAMESPACE***]
```

Replace [***NIFI CLUSTER NAME***] and [***NIFI CLUSTER NAMESPACE***] with the desired cluster name and cluster namespace respectively.

Configuring NiFi Image

Specify location of the image used for deployment.

This is how you specify the NiFi image repository and image version to be used for deployment. This describes the images used for running NiFi. This also provides a way of manually upgrading the NiFi version in an existing cluster or very quickly rolling out NiFi clusters with new versions.

```
spec:
  image:
    repository: container.repository.cloudera.com/cdp-private/cfm-nifi-k8s
    tag: []
```



Note:

container.repository.cloudera.com/cdp-private/cfm-nifi-k8s is the default repository for Cloudera Kubernetes images. If your Kubernetes cluster has no internet access or you want to use a self-hosted repository, replace it with the relevant path.

Configuring cluster size

Specify the number of pods in your deployment.

This section configures the number and capacity of your pods in the cluster.

```
spec:
  replicas: [***NUMBER OF REPLICAS***]
  resources:
    nifi:
      requests:
```



```

cpu: "[***CPU IN CORES***]"
memory: [***MEMORY IN BITES***]
limits:
  cpu: "[***CPU IN CORES***]"
  memory: [***MEMORY IN BITES***]
log:
  requests:
    cpu: [***CPU IN CORES***]
    memory: [***MEMORY IN BITES***]

```

Configuring bootstrap JVM settings

Learn about Java memory calculations, defaults and how to override them with custom values.

Cloudera Flow Management - Kubernetes Operator calculates JVM memory setting Max Direct Memory Size, Min Heap Size (xms), and Max Heap Size (xmx) based on container memory limits or requests.

In bootstrap settings, java.arg.10 is DirectMem, java.arg.2 is Min Heap, and java.arg.3 is Max Heap.

Java memory is calculated and set in the following order:

1. Based on memory of the NiFi resource

- The minimum DirectMem allowed for NiFi is 512MB. DirectMem is set to the maximum value between 512MB and 10% of the memory limit. If you do not provide a memory limit, the same calculation is made on the memory request in the specifications.
- Min Heap Size (xms) and Max Heap Size (xmx) is set to 75% of the memory limit subtracting the calculated DirectMem.

2. Defaults

If you do not specify memory for the NiFi resource, the following default values are automatically set:

- java.arg.2: -Xms2g
- java.arg.3: -Xmx2g
- java.arg.10: -XX:MaxDirectMemorySize=512m

Advanced configuration: Custom values to override inbuilt memory calculations

You can set each java argument for memory as part of NiFi specifications, under the configOverride key.

```

spec:
  configOverride:
    bootstrapConf:
      upsert:
        java.arg.2: -Xms2g
        java.arg.3: -Xmx2g
        java.arg.10: -XX:MaxDirectMemorySize=512m

```

Configuring persistence

Specify storage size and class globally, or for individual repositories.

This section specifies the storage to be used for the NiFi repositories. You can define storage globally, or have overrides for specific repositories. In case of OpenShift, the storage classes have to be specified at the OpenShift level to match the IOPS expectations for your NiFi workloads.

The Cloudera Flow Management - Kubernetes Operator can configure persistent disk storage for the following directories:

- state
- data
- FlowFile Repository
- Content Repository
- Provenance Repository

In the persistence spec, you can define a default size and StorageClass which applies to each of the directories. The spec can be further configured to define specific sizes and StorageClasses for each directory, if necessary.

```
spec:
  persistence:
    size: [***SIZE IN GIGABITES***]
    storageClass: [***STORAGE CLASS***]
    contentRepo:
      size: [***SIZE IN GIGABITES***]
      storageClass: [***STORAGE CLASS***]
    flowfileRepo:
      size: [***SIZE IN GIGABITES***]
    provenanceRepo:
      size: [***SIZE IN GIGABITES***]
```

Configuring assets

Learn about configuring access to NiFi assets.

You can make NiFi assets, such as custom processor JARs, or configuration files available to your NiFi cluster using the assets field. This field allows you to specify a mount path within the NiFi Pods to which the provided pre-existing Persistent Volume Claim (PVC) is mounted. Cloudera Flow Management - Kubernetes Operator does not provide a method of loading assets into this volume. Using the example below, all files located in the volume associated with my-nifi-assets-volume-claim are accessible at the path /opt/nifi/nifi-assets/ for use within your flow.

Before you can start using assets with NiFi deployed through Cloudera Flow Management - Kubernetes Operator, you have to provide the following:

- A volume provisioner which supports creating volumes that are ReadWriteMany (RWX), for example nfs-provisioner.
- A pod running some kind of software or process (for example, an FTP server) attached to the RWX volume for loading files onto the volume. Alternatively, you can use the `kubectl cp` command to directly copy files into most containers, such as Ubuntu.

After meeting the above prerequisites, you need to create a Persistent Volume Claim (PVC) which creates the required RWX volume. For example:

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: nifi-assets
spec:
  storageClassName: [***STORAGE CLASS***]
  accessModes:
    - ReadWriteMany
  resources:
    requests:
      storage: 1Gi
```

where `[***STORAGE CLASS***]` refers to the Storage Class associated with the RWX volume provisioner.

With the PVC created and the PV provisioned, attach the pod you have created, expose the pod via Ingress or Service if required, and load the asset files.

```
spec:
  assets:
    mountPath: [***ASSETS_PATH***]
    persistentVolumeClaim:
      name: nifi-assets
```

where [***ASSETS_PATH***] is the filesystem path within the NiFi container where the assets are located, for example, /opt/nifi/assets.

Related Information

[ReadWriteMany \(RWX\)](#)

[nfs-provisioner](#)

[Persistent Volume Claim \(PVC\)](#)

Configuring Kubernetes state management

Specify Kubernetes native state management provider as the state management provider of your cluster.

Cloudera's distribution of NiFi comes with a Kubernetes native state management provider. This is the recommended state management for use with Cloudera Flow Management - Kubernetes Operator. However, as it is not the default state management provider set by Cloudera Flow Management - Kubernetes Operator, you need to add this section to the configuration. Without this configuration, a ZooKeeper cluster is expected.

To configure the Kubernetes state management provider, use the below YAML.

```
spec:
  stateManagement:
    clusterProvider:
      id: kubernetes-provider
      class: org.apache.nifi.kubernetes.state.provider.KubernetesConfigMapsStateProvider
    configOverride:
      nifiProperties:
        upsert:
          nifi.cluster.leader.election.implementation: "KubernetesLeaderElectionManager"
```

Configuring node certificate generation

Learn about certificate generation options.

Cloudera Flow Management - Kubernetes Operator provides automatic certificate generation for each NiFi node in a given cluster by way of cert-manager certificates to secure intra-cluster communication between NiFis. To configure nodeCertGen, a cert-manager Issuer or ClusterIssuer is required. A self-signed Issuer setup is sufficient for development environments. In production environments use a third-party authority, or internal signing CAs.

```
spec:
  security:
    nodeCertGen:
      issuerRef:
        name: self-signed-ca-issuer
        kind: ClusterIssuer
```

Related Information[Issuers and ClusterIssuers](#)

Configuring NiFi properties

Learn how to override default NiFi configuration settings provided by Cloudera Flow Management - Kubernetes Operator from the CR file.

NiFi settings are available as part of the specification, under the `configOverride` key. They can be provided in one of the following ways:

- inline,
- as a ConfigMap
- as a Secret

```
spec:
  configOverride:
    nifiProperties:
      upsert:
        nifi.cluster.load.balance.connections.per.node: "1"
        nifi.cluster.load.balance.max.thread.count: "4"
        nifi.cluster.node.connection.timeout: "60 secs"
        nifi.cluster.node.read.timeout: "60 secs"
    bootstrapConf:
      upsert:
        java.arg.2: -Xms2g
        java.arg.3: -Xmx2g
        java.arg.13: -XX:+UseConcMarkSweepGC
```

Overriding NiFi settings using ConfigMaps and Secrets

Learn about overriding default NiFi settings using ConfigMaps and Secrets.

The ConfigMap or Secret values are available to inject into the environment for the following files:

- authorizers.xml
- bootstrap.conf
- logback.xml
- login-identity-providers.xml
- nifi.properties
- state-management.xml

Each of these config overrides must be in an individual ConfigMap with the key being the filename to be replaced. Using this ConfigMap or Secret reference method entirely overrides the defaults provided by the Cloudera Flow Management - Kubernetes Operator, which may impact cluster operation.

```
NiFiSpec
spec:
  configOverride:
    authorizersObjectReference:
      kind: "ConfigMap"
      name: "custom-authorizers"

ConfigMapSpec
data:
  authorizers.xml: |
    <authorizers>
```

```

    <authorizer>
      <identifier>single-user-authorizer</identifier>
      <class>org.apache.nifi.authorization.single.user.SingleUserAuth
orizer</class>
    </authorizer>
  </authorizers>

```

Configuring scaling

Learn about scaling NiFi clusters either manually or automatically, using HPA.

It is possible to manually scale up and down the NiFi cluster size by editing the replicas value in the deployment file and applying the changes. It is also possible to specify an HPA to automatically scale the NiFi cluster (replica count) based on the Kubernetes resources (CPU/memory).

To manually scale the cluster, simply edit the replicas field to your desired replica count.

For autoscaling, apply a Horizontal Pod Autoscaling (HPA) resource targeting the NiFi CR, as follows:

```

apiVersion: autoscaling/v2
kind: HorizontalPodAutoscaler
metadata:
  name: nifi-hpa
spec:
  maxReplicas: 3
  minReplicas: 1
  metrics:
    - type: Resource
      resource:
        name: cpu
        target:
          type: Utilization
          averageUtilization: 75
  scaleTargetRef:
    apiVersion: cfm.cloudera.com/v1alpha1
    kind: Nifi
    name: [***NIFI CLUSTER NAME***]

```

Configuring pod affinity

Pod affinity controls where pods are deployed based on node configuration and placement of other pods.

To learn more about Pod Affinity, read *Assigning Pods to Nodes* in the Kubernetes documentation.

You can configure the affinity settings of the NiFi pod in the NiFi Custom Resource under spec.statefulset. The following example represents the default configuration which will be added to the Custom Resource in the defaulting webhook.



Note:

If any affinity is provided in spec.statefulset, the default in the example will not be applied.

```

spec:
  statefulset:
    affinity:
      podAntiAffinity:
        preferredDuringSchedulingIgnoredDuringExecution:
          - podAffinityTerm:
              labelSelector:

```

```
matchExpressions:
- key: app.kubernetes.io/instance
  operator: In
  values:
  - mynifi
topologyKey: kubernetes.io/hostname
weight: 1
```

This default configuration attempts to spread NiFi cluster pods to different nodes of the Kubernetes cluster. If there are more NiFi pods than available Kubernetes nodes, then some pods will coexist on the same node.

Related Information

[Assigning Pods to Nodes | Kubernetes documentation](#)

Configuring connections to NiFi and NiFi Registry

Learn about configuring connections for your NiFi cluster.

Cloudera Flow Management - Kubernetes Operator provides a flexible method of configuring connections to NiFi called Connections. Using Connections, a Service, Ingress, or Route can be configured to route to a specific port on NiFi. For defining Connections targeting an arbitrary port on NiFi, use the `spec.connections` array. For configuring connection to the NiFi Web UI, use the `spec.uiConnection` field. This documentation provides a full reference for Connections.

Configuring session affinity

Learn about configuring session affinity. It makes possible to keep connection to the web UI alive in clusters with several nodes.

Regardless of your connection type, a NiFi cluster with more than one node requires session affinity of some type for the Web UI to operate. This is because each NiFi node can supply its own web UI and if a LoadBalancer shifts you to another instance, your authentication tokens become invalid. The best method of applying session affinity varies greatly depending on the Kubernetes cluster provider. In the simplest case, defining session affinity on the web Service resource itself is sufficient:

```
spec:
  uiConnection:
    serviceConfig:
      sessionAffinity: ClientIP
```

In certain clouds, for example AWS, the backing LoadBalancer resources do not support session affinity, and cause provisioning to break.

Configuring arbitrary connections

Learn about configuring a connections array.

You can use the `connections` array to flexibly define routing to ports on NiFi. The below example configures an Ingress resource with some annotations and labels provided. The Ingress will expose a URL `https://nifi.io/listenTCP` which routes to port 9432 on NiFi. Additionally, the backing Service is configured to have two extra ports, 8496 and 8495.

```
spec:
  connections:
  - type: Ingress
    name: someConnection
```

```

annotations:
  someanno: myanno
labels:
  somelabel: mylabel
ingressConfig:
  hostname: nifi.io
  paths:
  - port: 9432
    path: /listenTCP
    name: listentcp
serviceConfig:
  ports:
  - port: 8496
    protocol: TCP
    name: porta
  - port: 8495
    protocol: UDP
    name: portb

```

Configuring NiFi Web UI connection

Learn about configuring a connection to the NiFi web UI.

You can configure a connection to the NiFi Web UI using the `spec.uiConnection` field. It is a standard connection field with special validation and handling. The name of this connection is always ignored and set to `[**CR NAME**]-web`. For Ingress type connections, a maximum of one path may be specified. When you configure a `uiConnection`, the `spec.hostname` field is required.

The `uiConnection` can support hostname routing with and without an additional context path. It is not recommended to use a context path for routing as NiFi does not support it well, but it is possible. For more information, see NiFi documentation on proxy configuration. An example using `ingress-nginx` is included in this section.

Related Information

[NiFi proxy configuration](#)

Hostname-only ingress example

Learn about configuring an Ingress resource using TLS files generated by Cloudera Flow Management - Kubernetes Operator.

This YAML snippet configures an Ingress resource for accessing the NiFi Web UI. It uses the TLS files generated by a Cloudera Flow Management - Kubernetes Operator created Certificate as defined in `spec.security.ingressCertGen`. The supplied annotations are for the `ingress-nginx` Ingress controller. The affinity settings enable a persistent session so that UI interactions go to the same NiFi node in the cluster. The `backend-protocol` setting is needed for when NiFi is configured to be secure, as it will reject any non-HTTPS connection attempts.

```

spec:
  uiConnection:
    type: Ingress
    ingressConfig:
      ingressClassName: myIngressClass
      ingressTLS:
        - hosts:
            - nifi.localhost
          secretName: mynifi-ingress-cert
    annotations:
      nginx.ingress.kubernetes.io/affinity: cookie
      nginx.ingress.kubernetes.io/affinity-mode: persistent
      nginx.ingress.kubernetes.io/backend-protocol: HTTPS

```

Hostname-only route example

Learn about configuring a Route resource to access the NiFi web UI.

This YAML snippet configures a Route resource for accessing the NiFi web UI.

```
spec:
  uiConnection:
    type: Route
    routeConfig:
      tls:
        termination: passthrough
```

Ingress with context path example

Learn about configuring an Ingress resource that rewrites the connection path in incoming requests and does a reverse-rewrite on UI calls going to the backend.

This YAML code snippet configures an ingress UI Connection with a path. The annotations here are for the ingress-nginx ingress controller and all are required for NiFi to correctly understand the incoming requests.

In the example the path includes some regex at the end: `(/|$)(.*)`. This regex informs the rewrite directives in the configuration-snippet and rewrite-target annotations. NiFi does not handle proxy paths well, it does not understand that `https://nifi.localhost/some/path/to/nifi` coming through the defined Ingress is intended to call the `/nifi` API to load the UI. The rewrite-target annotation addresses this by capturing the `/nifi` and anything that comes after and sends that as the path to the NiFi pod. It translates `/some/path/to/nifi/` to `/nifi/`. Similarly, the NiFi web UI does not correctly form API calls going to the backend, attempting to call `/nifi/` instead of `/some/path/to/nifi/`. This is addressed by the configuration-snippet rewrite instruction. It does the reverse of the rewrite-target, reapplying the removed context path `/some/path/to`. The remaining configuration-snippet lines are headers required by a NiFi behind a proxy. For more information, see the *NiFi System Administrator's Guide*.

```
spec:
  uiConnection:
    type: Ingress
    ingressConfig:
      ingressClassName: myIngressClass
      ingressTLS:
        - hosts:
            - nifi.localhost
          secretName: mynifi-ingress-cert
      paths:
        - port: 8443
          path: "/some/path/to(/|$)(.*)"
      annotations:
        nginx.ingress.kubernetes.io/affinity: cookie
        nginx.ingress.kubernetes.io/affinity-mode: persistent
        nginx.ingress.kubernetes.io/backend-protocol: HTTPS
        nginx.ingress.kubernetes.io/configuration-snippet: |-
          proxy_set_header X-ProxyScheme $scheme;
          proxy_set_header X-ProxyHost $host;
          proxy_set_header X-ProxyPort $server_port;
          proxy_set_header X-ProxyContextPath /some/path/to;
          rewrite (.*\nifi)$ $1/ redirect;
          proxy_ssl_name mynifi.default.svc.cluster.local;
        nginx.ingress.kubernetes.io/rewrite-target: /$2
```

Configuring authentication for NiFi and NiFi Registry

Learn about configuring single user authentication for development purposes and TLS authentication more suited for production environments.

Configuring single user authentication

Single user authentication is NiFi's most basic authentication option, sufficient for individual development clusters and also production clusters where flows are deployed in a controlled manner, such as continuous integration (CI) or site reliability engineering (SRE). A single user is granted all permissions on the NiFi cluster, no other users can be configured.

- Configuration snippet for letting NiFi generate the password.

```
spec:
  security:
    singleUserAuth:
      enabled: true
```

You find the generated username and password in the app-log container logs.

- Configuration snippet for setting NiFi username and password using a Secret:

```
spec:
  security:
    singleUserAuth:
      enabled: true
      credentialsSecretName: [***YOUR CREDENTIALS SECRET***]
```

Replace:

[*YOUR CREDENTIALS SECRET***]**

Create your credentials secret with the following command:

```
kubectl create secret generic [***YOUR CREDENTIALS SECRET***] --
from-literal=username=[***YOUR USER NAME***] --from-literal=password=[***YOUR PASSWORD***]
```

Replace:

[*YOUR CREDENTIALS SECRET***]**

with the desired credentials secret name

[*YOUR USER NAME***]**

with the generated username in the app-log container logs

[*YOUR PASSWORD***]**

with the generated password in the app-log container logs

Configuring LDAP authentication

Learn how to configure an LDAP server for user authentication in your NiFi or NiFi Registry cluster.

Cloudera Flow Management - Kubernetes Operator can configure NiFi to connect to an LDAP server for user authentication.

Prerequisites:

- Full LDAP URL, i.e. `ldap://[***LDAP SERVER URL***]:[***LDAP PORT***]`
- Desired authentication strategy
- Authentication credentials and key/trust stores if using LDAPS.
- User search filters

For LDAP servers protected with any authentication, a Secret must be created containing the correct authentication credentials and TLS resources (if applicable). The Secret must contain the following data fields:

- managerPassword
- keystore (if TLS is configured)
- keystorePassword (if TLS is configured)
- truststore (if TLS is configured)
- truststorePassword (if TLS is configured)

Create the secret using the kubectl CLI utility:

```
kubectl create secret generic my-ldap-creds \
  --from-literal=managerPassword=myManagerPassword \
  --from-file=keystore=/path/to/keystore \
  --from-literal=keystorePassword=myKeystorePassword \
  --from-file=truststore=/path/to/truststore \
  --from-literal=truststorePassword=myTruststorePassword
```

The following example shows a connection to an LDAP server protected with basic authentication with TLS.

```
spec:
  security:
    initialAdminIdentity: mynifiadmin
    ldap:
      authenticationStrategy: SIMPLE
      managerDN: "cn=admin,dc=example,dc=org"
      secretName: my-openldap-creds
      referralStrategy: FOLLOW
      connectTimeout: 3 secs
      readTimeout: 10 secs
      url: ldap://my-ldap-url:389
      userSearchBase: "dc=example,dc=org"
      userSearchFilter: "(uid={0})"
      identityStrategy: USE_USERNAME
      authenticationExpiration: 12 hours
  tls:
    keystoreType: jks
    truststoreType: jks
    clientAuth: NONE
    protocol: TLSv1.2
```

By default, Cloudera Flow Management - Kubernetes Operator does not deploy a UserGroupProvider using the LDAP target. This means NiFi does not pull down any users, only queries the LDAP server for authentication. This impedes configuring user access, requiring the NiFi administrator to create each user manually.

The following example shows configuring user synchronization with the LDAP server:

```
spec:
  security:
    ldap:
      sync:
        interval: 30 min
        userObjectClass: inetOrgPerson
        userSearchScope: SUBTREE
        userIdentityAttribute: cn
        userGroupNameAttribute: ou
        userGroupNameReferencedGroupAttribute: ou
        groupSearchBase: "dc=example,dc=org"
        groupObjectClass: organizationalUnit
        groupSearchScope: OBJECT
        groupNameAttribute: ou
```

Example CR

The following example NiFi CR deploys a 3 node cluster with Kubernetes-based state management and leader election, and a Route to access the NiFi UI.

```
apiVersion: cfm.cloudera.com/v1alpha1
kind: Nifi
metadata:
  name: mynifi
spec:
  replicas: 3
  image:
    repository: container.repository.cloudera.com/cloudera/cfm-nifi-k8s
    tag: [***NIFI TAG***]
    pullSecret: docker-pull-secret
  tiniImage:
    repository: container.repository.cloudera.com/cloudera/cfm-tini
    tag: [***CFM TINI TAG***]
    pullSecret: docker-pull-secret
  hostName: mynifi.[***OPENSIFT ROUTER DOMAIN***]
  uiConnection:
    type: Route
    serviceConfig:
      sessionAffinity: ClientIP
```