ADS Tuning Guide

—

Nicolas Peulvast
Performance Architect

Disclaimer Official

This presentation has been prepared by IBM and individual employees and reflect their own views. It is provided for informational purposes only, and is neither intended to, nor shall have the effect of being, legal or other guidance or advice to any readers. While efforts were made to verify the completeness and accuracy of the information contained in this presentation, it is provided AS IS without warranty of any kind, express or implied. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this presentation or any other materials. Nothing contained in this presentation is intended to, nor shall have the effect of, creating any warranties or representations from IBM or its suppliers or licensors, or altering the terms and conditions of the applicable license agreement governing the use of IBM software.

References in this presentation to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates. Product release dates and/or capabilities referenced in this presentation may change at any time at IBM's sole discretion based on market opportunities or other factors, and are not intended to be a commitment to future product or feature availability in any way. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth or other results.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Contents

Executive Summary

Authentication

The Authentication mode chosen has an impact on the throughput of the ADS Runtime Pod and is managed by the Common Services Layer: be sure to have an adapted T-Shirt size of the Common Services Layer to optimize the Authentication performance.

Network

The Network layer may have an important impact of the overall performance and especially on the throughput of the ADS Runtime Pod in case of multiple users in parallel : be sure to have an adapted T-Shirt size of the Zen Layer to optimize the network performance.

BAI

Activating the BAI event emission has an overhead, that is dependent of the additional collected information.

You must increase the resource used for the system to get back to the expected throughput.

Note that you must tune the BAI stack to have a similar average event throughput ingestion from BAI in order to not overload your system.

Tuning

A fine-tuning of the different layers is the key to a good throughput on the ADS Runtime, and especially the network layers.

See the following slides for the tuning guidelines.

General Starter Tuning Guidelines

- In order to propose a Demo/Minimal sizing, the following configuration is suggested for the ADS Runtime service:
  - CPU Request == CPU Limit == 0.5
  - While startup time is longer, it does not reach probe limits
  - With this configuration, we still have a good level of performance at the ADS Runtime level

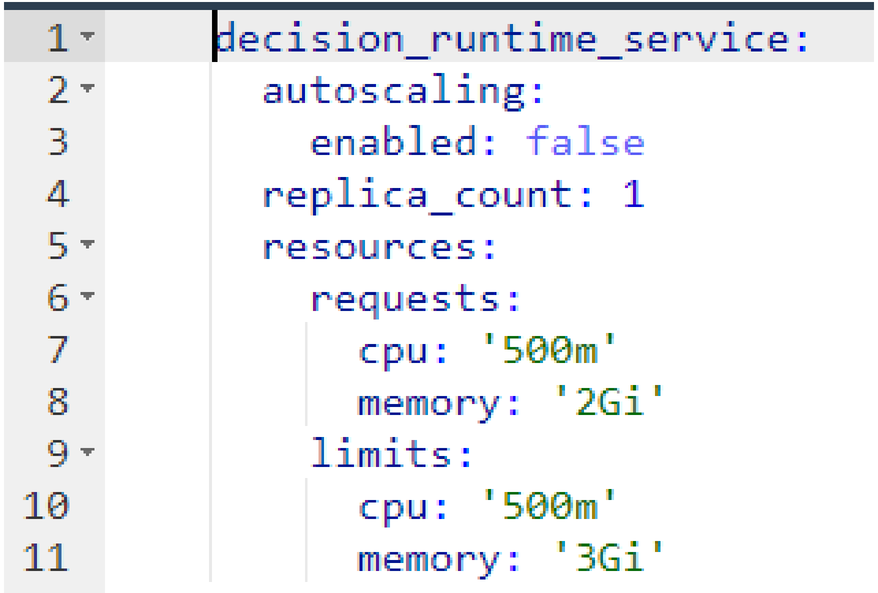

```yaml
decision_runtime_service:
  autoscaling:
    enabled: false
  replica_count: 1
  resources:
    requests:
      cpu: '500m'
      memory: '2Gi'
    limits:
      cpu: '500m'
      memory: '3Gi'
```

General Prod Tuning Guidelines

- The network latency between the Cluster and the database has a huge impact on the Designer performance as the Git Service performs a lot of small requests on each repository access.
- Ephemeral Storage & impact on stability: if the Ephemeral storage is configured too small, then the ADS Pods will be evicted.
- A Horizontal Pod Autoscaler (HPA) is available out-of-the-box in the ADS Runtime delivery
  - Using this HPA increase the throughput of the ADS runtime but also increase the number of VPC billed to the customer
  - By default, this HPA is not set for the Small, Medium and Large T-Shirt size but only for X-Large T-Shirt size

Tuning Foundational service

Align the profile (small, medium, large) of your Common service/Foundational Service to the targeted profile of your ADS product.

- In the OpenShift console:
  - Search (ibm-common-services ns) > Resources ‘CommonService’ > common-service
  - In the YAML, change the spec.size value
    - starterset
    - small
    - medium
    - large

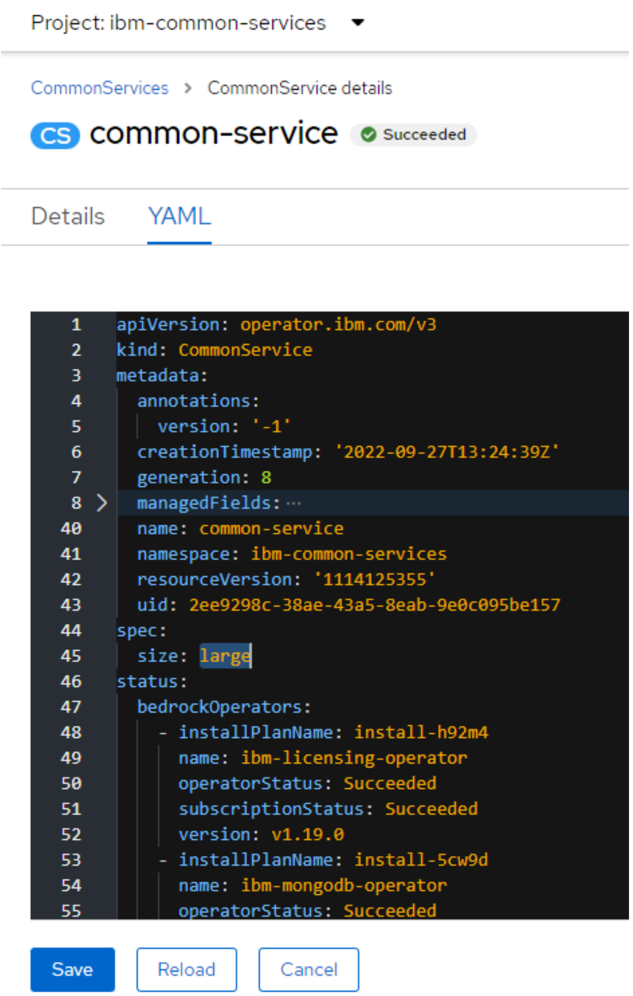

```yaml
apiVersion: operator.ibm.com/v3
kind: CommonService
metadata:
  annotations:
    version: '-1'
  creationTimestamp: '2022-09-27T13:24:39Z'
  generation: 8
  managedFields:···
  name: common-service
  namespace: ibm-common-services
  resourceVersion: '1114125355'
  uid: 2ee9298c-38ae-43a5-8eab-9e0c095be157
spec:
  size: large
status:
  bedrockOperators:
    - installPlanName: install-h92m4
      name: ibm-licensing-operator
      operatorStatus: Succeeded
      subscriptionStatus: Succeeded
      version: v1.19.0
    - installPlanName: install-5cw9d
      name: ibm-mongodb-operator
      operatorStatus: Succeeded
```

Tuning Shared Configuration

Align the profile (small, medium, large) of your Shared configuration to the targeted profile of your ADS product: it's particularly useful when the BAI events are used in ADS as it tunes the InsightEngine.

- Change it using the Shared Configuration
- In the OpenShift console:
  - Search (your ns) > Resources 'ICP4ACluster' > select your CR deployment
  - In the YAML, change the spec.shared_configuration.sc_deployment_profile_size value
    - small
    - medium
    - large

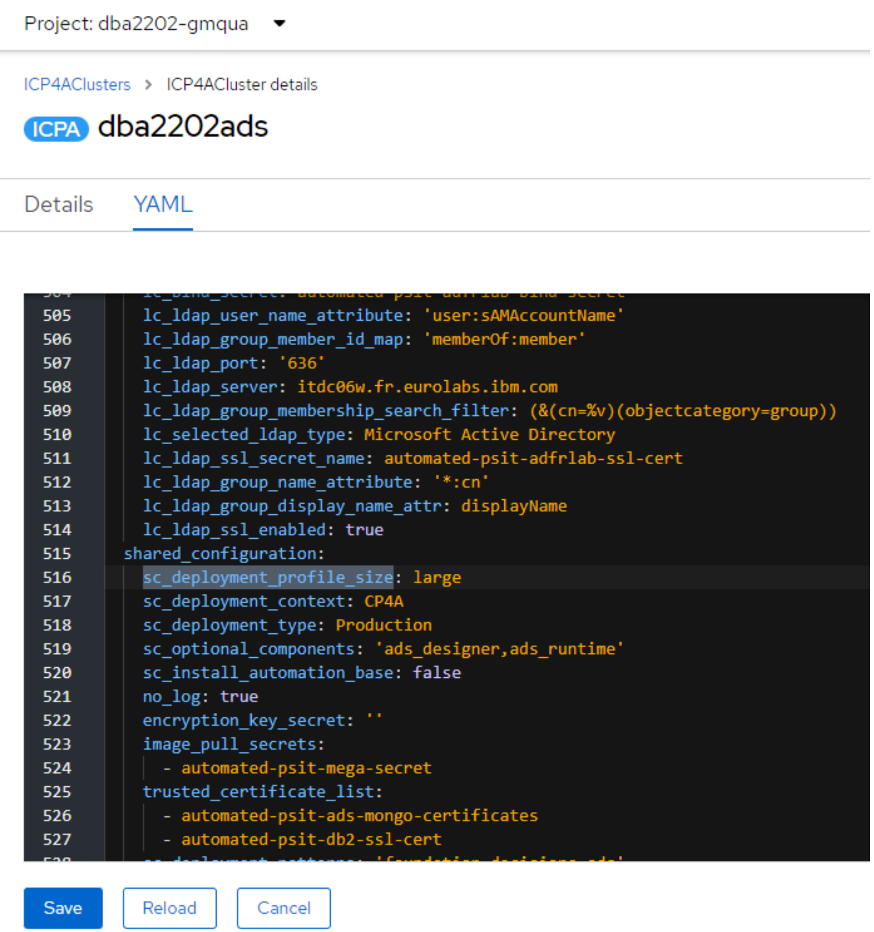

```yaml
    lc_bind_secret: automated-psit-adfrlab-bind-secret
    lc_ldap_user_name_attribute: 'user:sAMAccountName'
    lc_ldap_group_member_id_map: 'memberOf:member'
    lc_ldap_port: '636'
    lc_ldap_server: itdc06w.fr.eurolabs.ibm.com
    lc_ldap_group_membership_search_filter: (&(cn=%v)(objectcategory=group))
    lc_selected_ldap_type: Microsoft Active Directory
    lc_ldap_ssl_secret_name: automated-psit-adfrlab-ssl-cert
    lc_ldap_group_name_attribute: '*:cn'
    lc_ldap_group_display_name_attr: displayName
    lc_ldap_ssl_enabled: true
  shared_configuration:
    sc_deployment_profile_size: large
    sc_deployment_context: CP4A
    sc_deployment_type: Production
    sc_optional_components: 'ads_designer,ads_runtime'
    sc_install_automation_base: false
    no_log: true
    encryption_key_secret: ''
    image_pull_secrets:
      - automated-psit-mega-secret
    trusted_certificate_list:
      - automated-psit-ads-mongo-certificates
      - automated-psit-db2-ssl-cert
    sc_deployment_patterns: 'foundation,decisions_ads'
```

Tuning IAF Layer

- This performance tuning is only applicable if you select the BAI event emitter
- You can add additional JVM parameter in the AutomationBase configuration
- If you do that, you must be sure to set shared_configuration.sc_install_automation_base to false

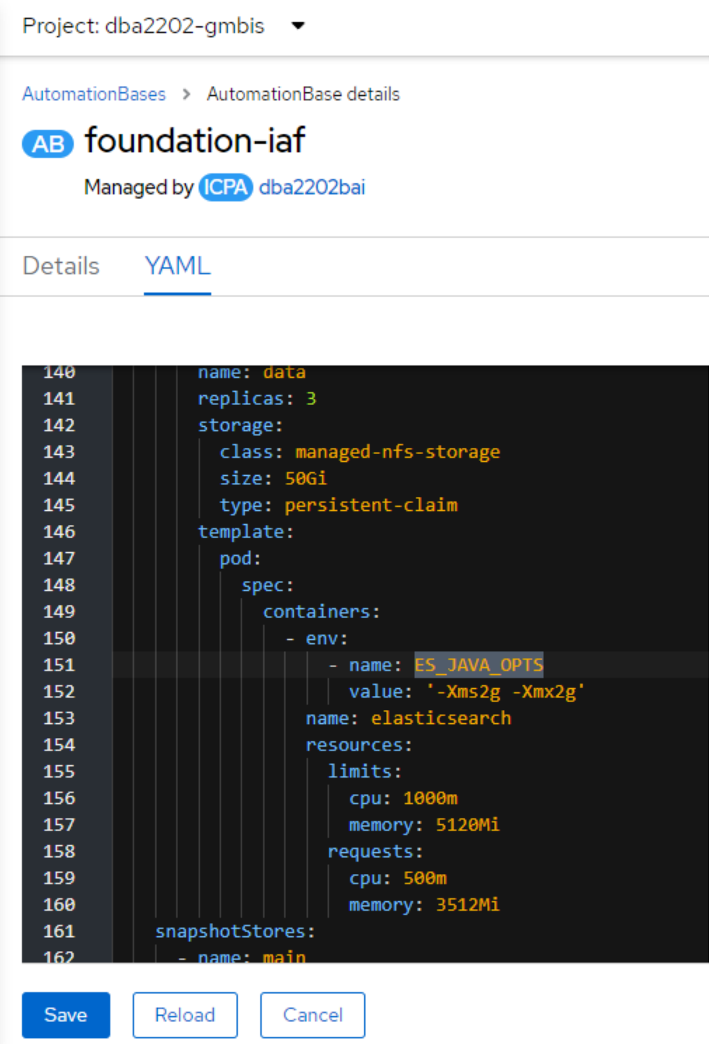

Project: dba2202-gmbis
AutomationBases > AutomationBase details
AB foundation-iaf
Managed by ICPA dba2202bai
Details
YAML
Save
Reload
Cancel

```
    name: data
    replicas: 3
    storage:
      class: managed-nfs-storage
      size: 50Gi
      type: persistent-claim
    template:
      pod:
        spec:
          containers:
            - env:
                - name: ES_JAVA_OPTS
                  value: '-Xms2g -Xmx2g'
              name: elasticsearch
              resources:
                limits:
                  cpu: 1000m
                  memory: 5120Mi
                requests:
                  cpu: 500m
                  memory: 3512Mi
snapshotStores:
  - name: main
```

Tuning Zen Layer

- You can tune the NGNIX layer
  - Edit the `ConfigMap` `<crname>-ads-designer-zen-configuration & <crname>-ads-runtime-zen-configuration`
  - Tune the `nginx.conf` part of the `ConfigMap`
- You can verify the zen resources via the `scaleConfig` parameter in Zen Service where the default profile is `small`, and we also support `medium` / `large` / `xlarge`
  - The Zen layer is adapted to the T-Shirt size that you selected in your Custom Resource
  - T-shirt size `medium` - 3 replicas , cpu limit 800M – Throughput of ADS Runtime Pod up to 11000 TPS
  - T-shirt size `large/xlarge` - 5 replicas, cpu limit 2 – Throughput of ADS Runtime Pod up to 11000 TPS
  - We recommend using the large configuration to avoid bottleneck in the Zen Layer
  - We can manually force the size of your Zen layer using `oc patch AutomationUIConfig iaf-system --type=merge -p '{"spec":{"zenService":{"scaleConfig":"large"}}}’`
    But note the operator will change it back to your Custom Resource value after one roundtrip of the Operator (usually 20 minutes)
- Put the `roundrobin` annotation in you `CPD` route
  - `haproxy.router.openshift.io/balance=roundrobin`
  - Since 21.0.3, This annotation is always override by the source value as some element behind the reverse-proxy have bug in using the roundrobin algorithm.
  - If you want to switch to the `roundrobin` algorithm, you have to de-activate the CP4A operator.

Tuning ADS configuration

Adapt profile (small, medium, large) of your ADS configuration to the targeted profile of your ADS product.

- Change it using the ADS Configuration
- In the OpenShift console:
  - Search (your ns) > Resources ‘ICP4ACluster’ > select your CR deployment
  - In the YAML, change the spec.ads_configuration.deployment_profile_size value
    - small
    - medium
    - large
    - extra-large

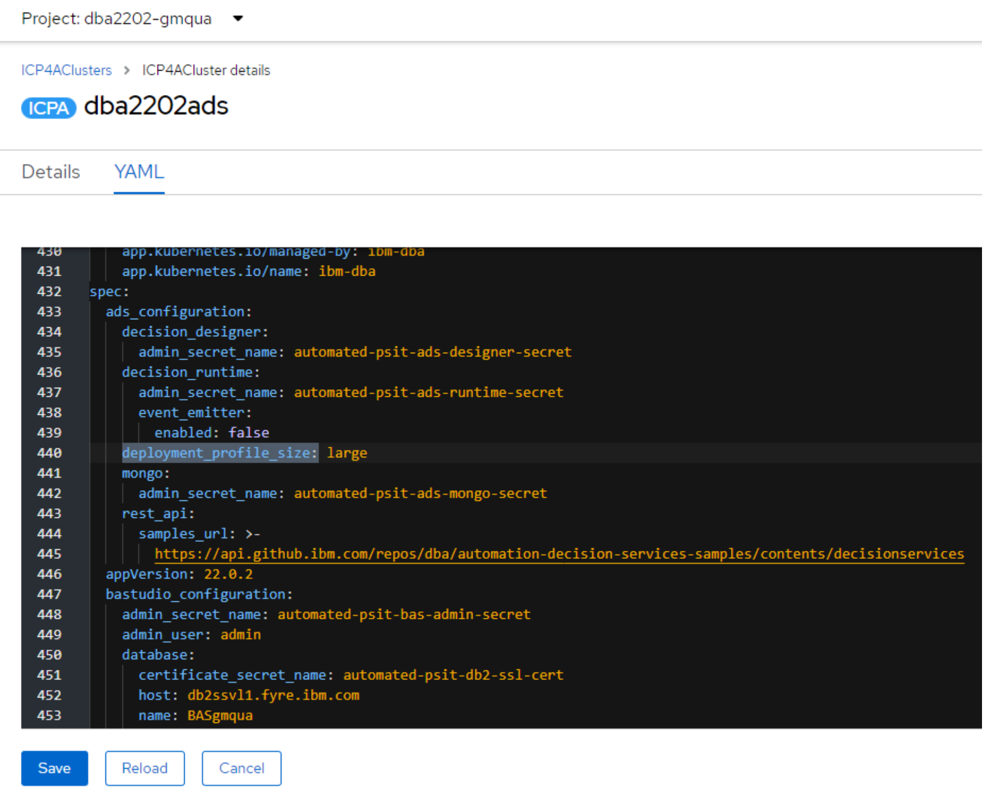

```
    app.kubernetes.io/managed-by: ibm-dba
    app.kubernetes.io/name: ibm-dba
spec:
  ads_configuration:
    decision_designer:
      admin_secret_name: automated-psit-ads-designer-secret
    decision_runtime:
      admin_secret_name: automated-psit-ads-runtime-secret
      event_emitter:
        enabled: false
    deployment_profile_size: large
    mongo:
      admin_secret_name: automated-psit-ads-mongo-secret
    rest_api:
      samples_url: >-
        https://api.github.ibm.com/repos/dba/automation-decision-services-samples/contents/decisionservices
  appVersion: 22.0.2
  bastudio_configuration:
    admin_secret_name: automated-psit-bas-admin-secret
    admin_user: admin
    database:
      certificate_secret_name: automated-psit-db2-ssl-cert
      host: db2ssvl1.fyre.ibm.com
      name: BASgmqua
```

Tuning Frontend Layer

- Check the number of thread allocated to your frontend/HAProxy and verify that the configuration is using the roundrobin algorithm
- As an example, edit the `/etc/haproxy/haproxy.cfg` file and change
  - the `ingress-https` backend from "`balance source`" to "`balance roundrobin`“
  - the `nbproc` should be at least set to 5

Tuning Network Layer

- During our testing, we reached the network bandwidth capacity (1Gbps) hence the response curve flattens so switch to 10Gbps network should be considered.
- You can also tune the network layer as follow
  - Search [in ns `openshift-ingress-operator`] > Resources ‘`IngressController`’ > select default
    - Edit the yaml and change the `spec.replicas` to at least 5
  - Search [in your cp4ba ns] > Deployment ‘`ibm-nginx`’ > scale up the number of pod in order to reach 100% of CPU usage in your ODM Runtime

Tuning HPAs

- HPAs are automatically created with the extra-large sizing configuration
- You have 3 additional HPA created in your namespace in that case

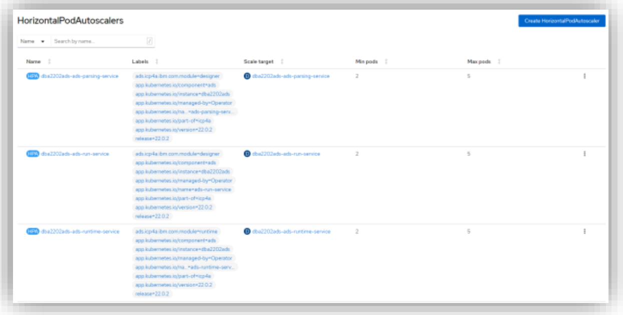

Name	Labels	Scale target	Min pods	Max pods
HPA dba2202ads-ads-parsing-service	ads.icp4a.ibm.com.module=designer app.kubernetes.io/component=ads app.kubernetes.io/instance=dba2202ads app.kubernetes.io/managed-by=Operator app.kubernetes.io/na...=ads-parsing-serv... app.kubernetes.io/part-of=icp4a app.kubernetes.io/version=22.0.2 release=22.0.2	D dba2202ads-ads-parsing-service	2	5
HPA dba2202ads-ads-run-service	ads.icp4a.ibm.com.module=designer app.kubernetes.io/component=ads app.kubernetes.io/instance=dba2202ads app.kubernetes.io/managed-by=Operator app.kubernetes.io/name=ads-run-service app.kubernetes.io/part-of=icp4a app.kubernetes.io/version=22.0.2 release=22.0.2	D dba2202ads-ads-run-service	2	5
HPA dba2202ads-ads-runtime-service	ads.icp4a.ibm.com.module=runtime app.kubernetes.io/component=ads app.kubernetes.io/instance=dba2202ads app.kubernetes.io/managed-by=Operator app.kubernetes.io/na...=ads-runtime-serv... app.kubernetes.io/part-of=icp4a app.kubernetes.io/version=22.0.2 release=22.0.2	D dba2202ads-ads-runtime-service	2	5

- You can change your HPAs using the following Custom Resource customization:

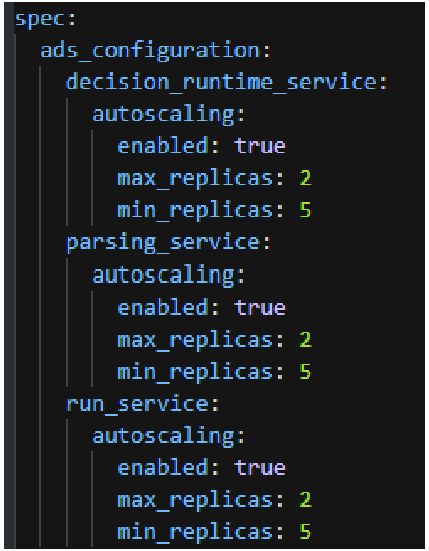

```
spec:
  ads_configuration:
    decision_runtime_service:
      autoscaling:
        enabled: true
        max_replicas: 2
        min_replicas: 5
    parsing_service:
      autoscaling:
        enabled: true
        max_replicas: 2
        min_replicas: 5
    run_service:
      autoscaling:
        enabled: true
        max_replicas: 2
        min_replicas: 5
```

Storage consideration

Official CloudPak documentation

https://www.ibm.com/docs/en/cloud-paks/cp-biz-automation/22.0.2?topic=deployment-storage-considerations

Align the profile (small, medium, large) of your Foundational service to the targeted profile of you ADS product.

Tip for CR updating

- In order to be sure that a CR change is considered as soon as possible, you can delete the pod `ibm-cp4a-operator-XX-XX` in order to shutdown the current reconciliation loop (running on the old CR).
- It will result with a new resolution loop that will consider your new values.

Additional resources

- https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/pdf/scalability_and_performance/openshift_container_platform-4.10-scalability_and_performance-en-us.pdf

Abbreviations

	Definition
ADS	IBM Automation Decision Services – Tested product from the CloudPak that provides a comprehensive environment for authoring, managing, and running decision services
BAI	IBM Business Automation Insights – Product from the CloudPak that processes event data so that you can derive insights into the performance of your business. You can use this data to drive automations and visualize the state of the indicators that matter most to you in near real-time
CPU	Central Processing Unit - In Kubernetes, 1 CPU unit is equivalent to 1 physical CPU core, or 1 virtual core, depending on whether the node is a physical host or a virtual machine running inside a physical machine. See: https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/#meaning-of-cpu
CS	IBM Cloud Pak Common Services – a.k.a IBM Cloud Pak foundational services, this Cloud Pak provides key foundational services See: https://www.ibm.com/docs/en/cpfs?topic=about
CR	Custom Resource: Description YAML file of object ICP4ACluster
PSIT	Performance and System Integration Testing team
PVC	A PersistentVolumeClaim (PVC) is a request for storage by a user See: https://kubernetes.io/docs/concepts/storage/persistent-volumes/