Hyland™

Hyland™

Alfresco
Tech Talk Live #156

March 22, 2024

Agenda

- Community news
- Towards a sustainable deployment

Resources

Alfresco

Deploying Alfresco Outlook Transform Engine with Docker Compose @ GitHub

Script Object for RenditionService2 @ GitHub and thanks to @hi-ko

Alfresco GenAI @ GitHub

Resources to come

- Adapting your logging configuration to log4jv2
- How to set up messages for increased reliability
- Upgrading your addons to Jakarta EE 10 and Spring 6
- Upgrading to Apache Tomcat 10
- Using Control Center App with Community Edition
- Deploying Alfresco with Helm in Kubernetes

Collaboration

Blog posts

- Summarization of textual content in Alfresco repository with Amazon Bedrock by @abhinavmishra14

Contributions

- https://github.com/aborroy/alfresco-dockerx-builder/issues/11 by @uvukasinovic
- https://github.com/Alfresco/alfresco-docker-installer/issues/185 by @luca86r
- https://github.com/Alfresco/alfresco-docker-installer/issues/186 by @N00BTellaBrot

Conferences

TTL Speakers wanted!

- Take the opportunity to showcase your work with the community
- About Alfresco, Nuxeo, and associated technologies
- Best practices, integration, scaling, cloud, …
- In your native language

Today’s talk

Towards a sustainable deployment

Alexandre Chapellon and Giovanni Toraldo,
Hyland

Hyland™

Hyland™

TTL #156
Towards a Sustainable deployment

Alexandre Chapellon, Valerio Provaggi, Giovanni Toraldo

March 22, 2024

Digital transformation has the potential to mitigate climate impacts

So basically is just a matter of consuming less resources?

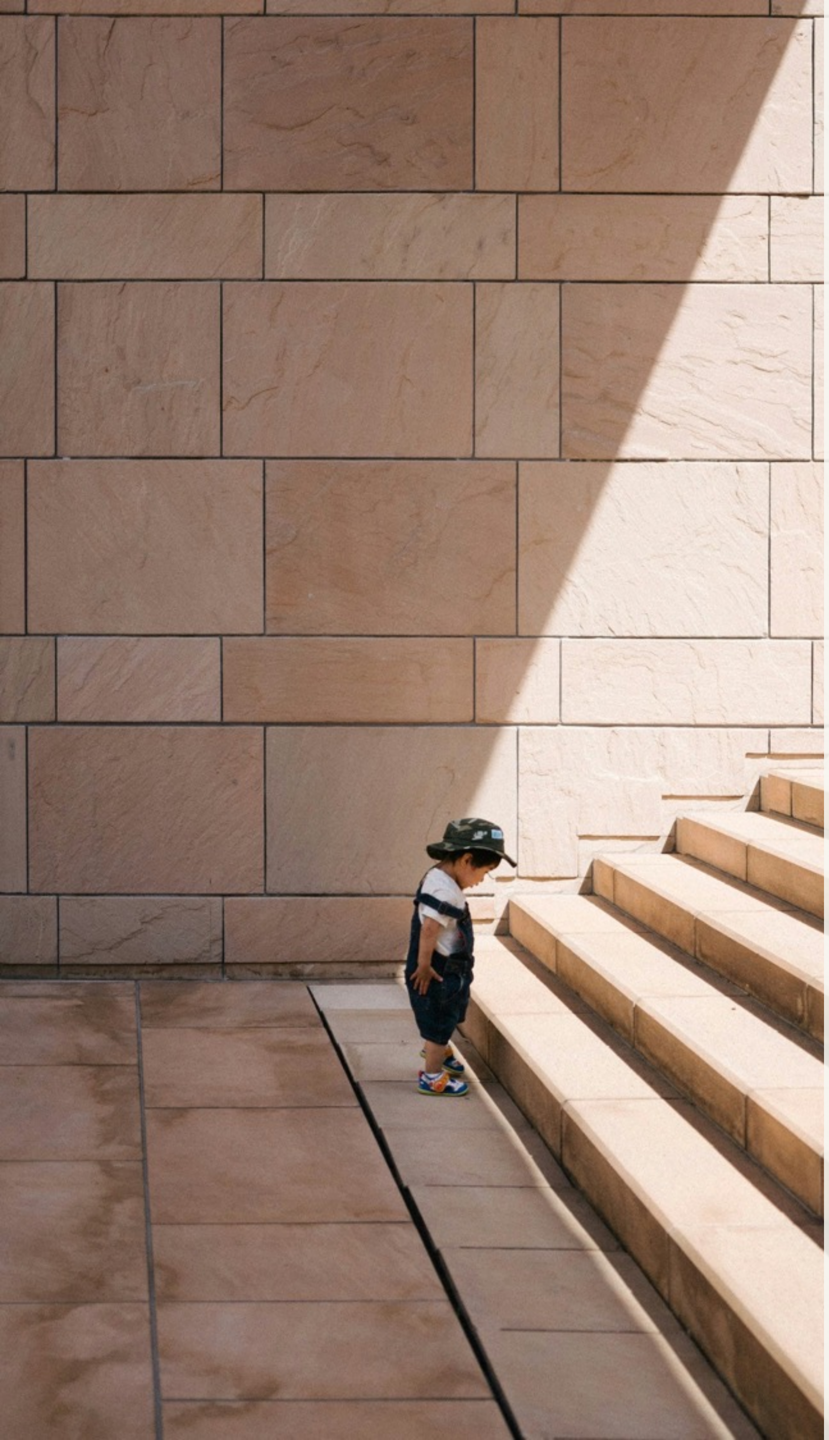

A bit more complex than that

- Efficiency
- Awareness

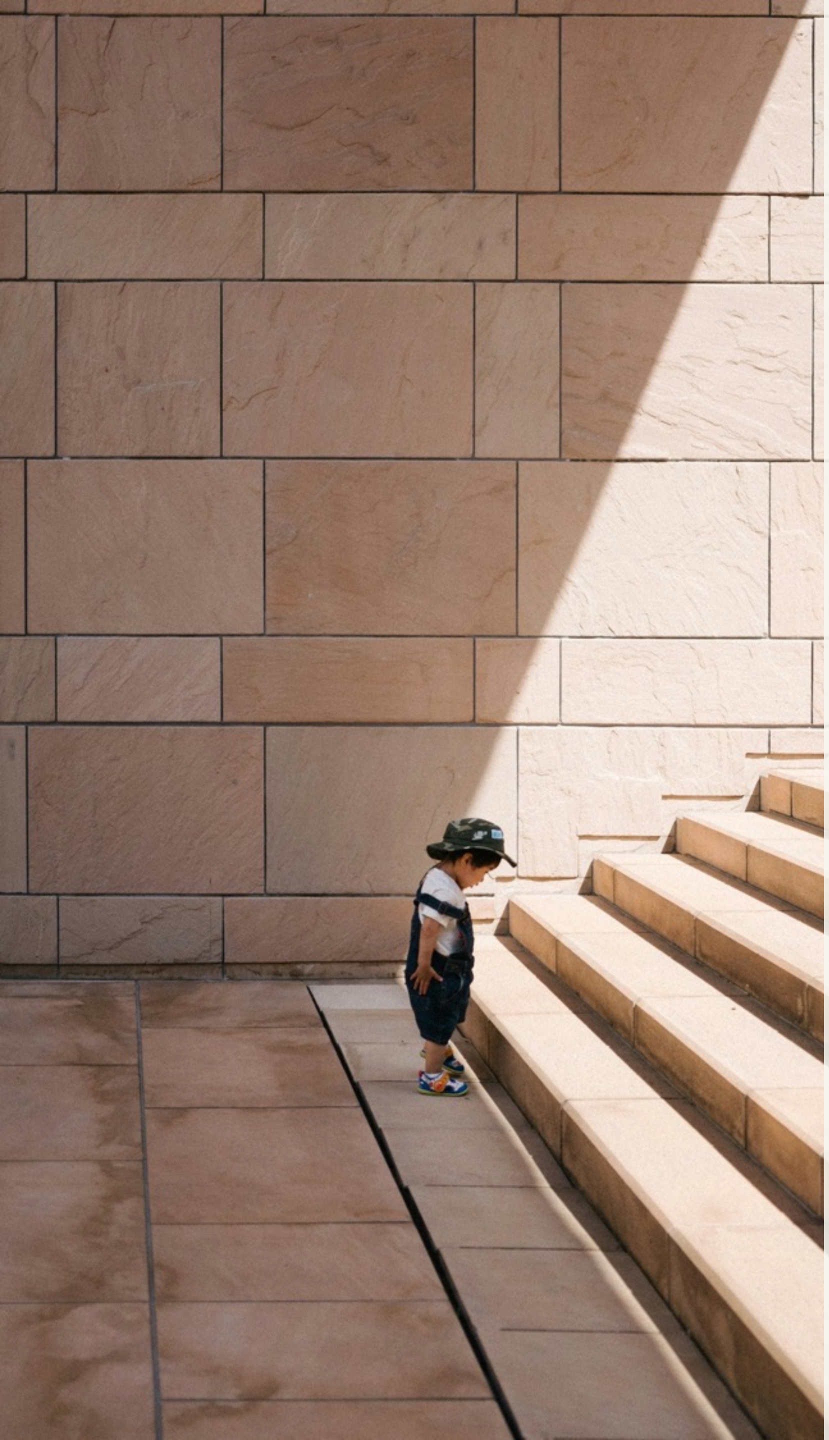

A bit more complext than that

⚡ 🔍

Experience Design

Applicaton Design

Applicaton Deployment

Infrastructure

Agenda

- Helm charts modularity
- Graviton deployment
- Next steps

Helm charts modularity

Why modularity?

- Deploy only what you need and save resources
- Makes Alfresco platform composable
- More generic means of configuration, hence more flexibility
- Gives more freedom in choosing dependencies for 3rd party components

Comparison

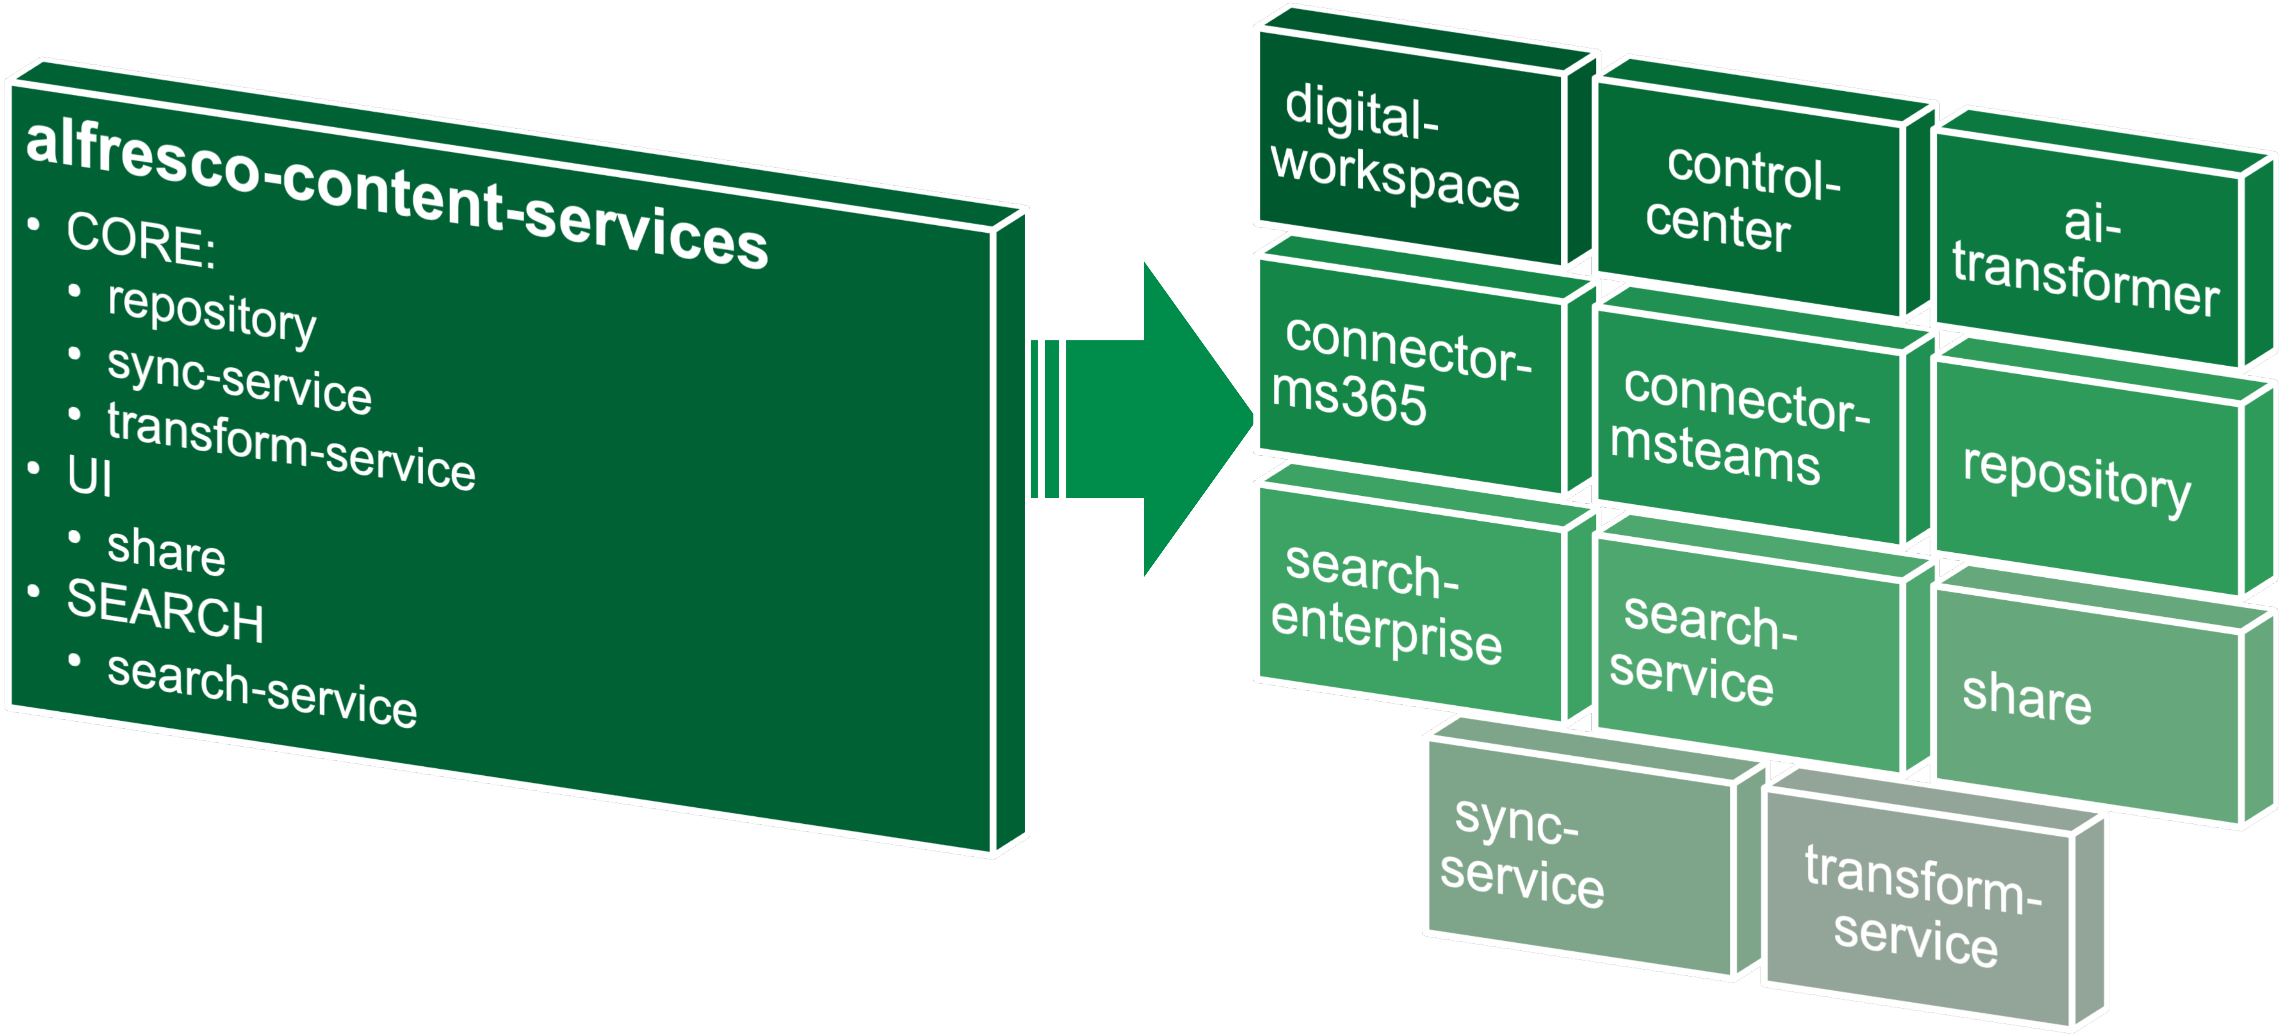

alfresco-content-services

- CORE:
  - repository
  - sync-service
  - transform-service
- UI
  - share
- SEARCH
  - search-service

digital-workspace

control-center

ai-transformer

connector-ms365

connector-msteams

repository

search-enterprise

search-service

share

sync-service

transform-service

Must deploy at least all of

Can deploy any of

alfresco-content-services as an umbrella chart

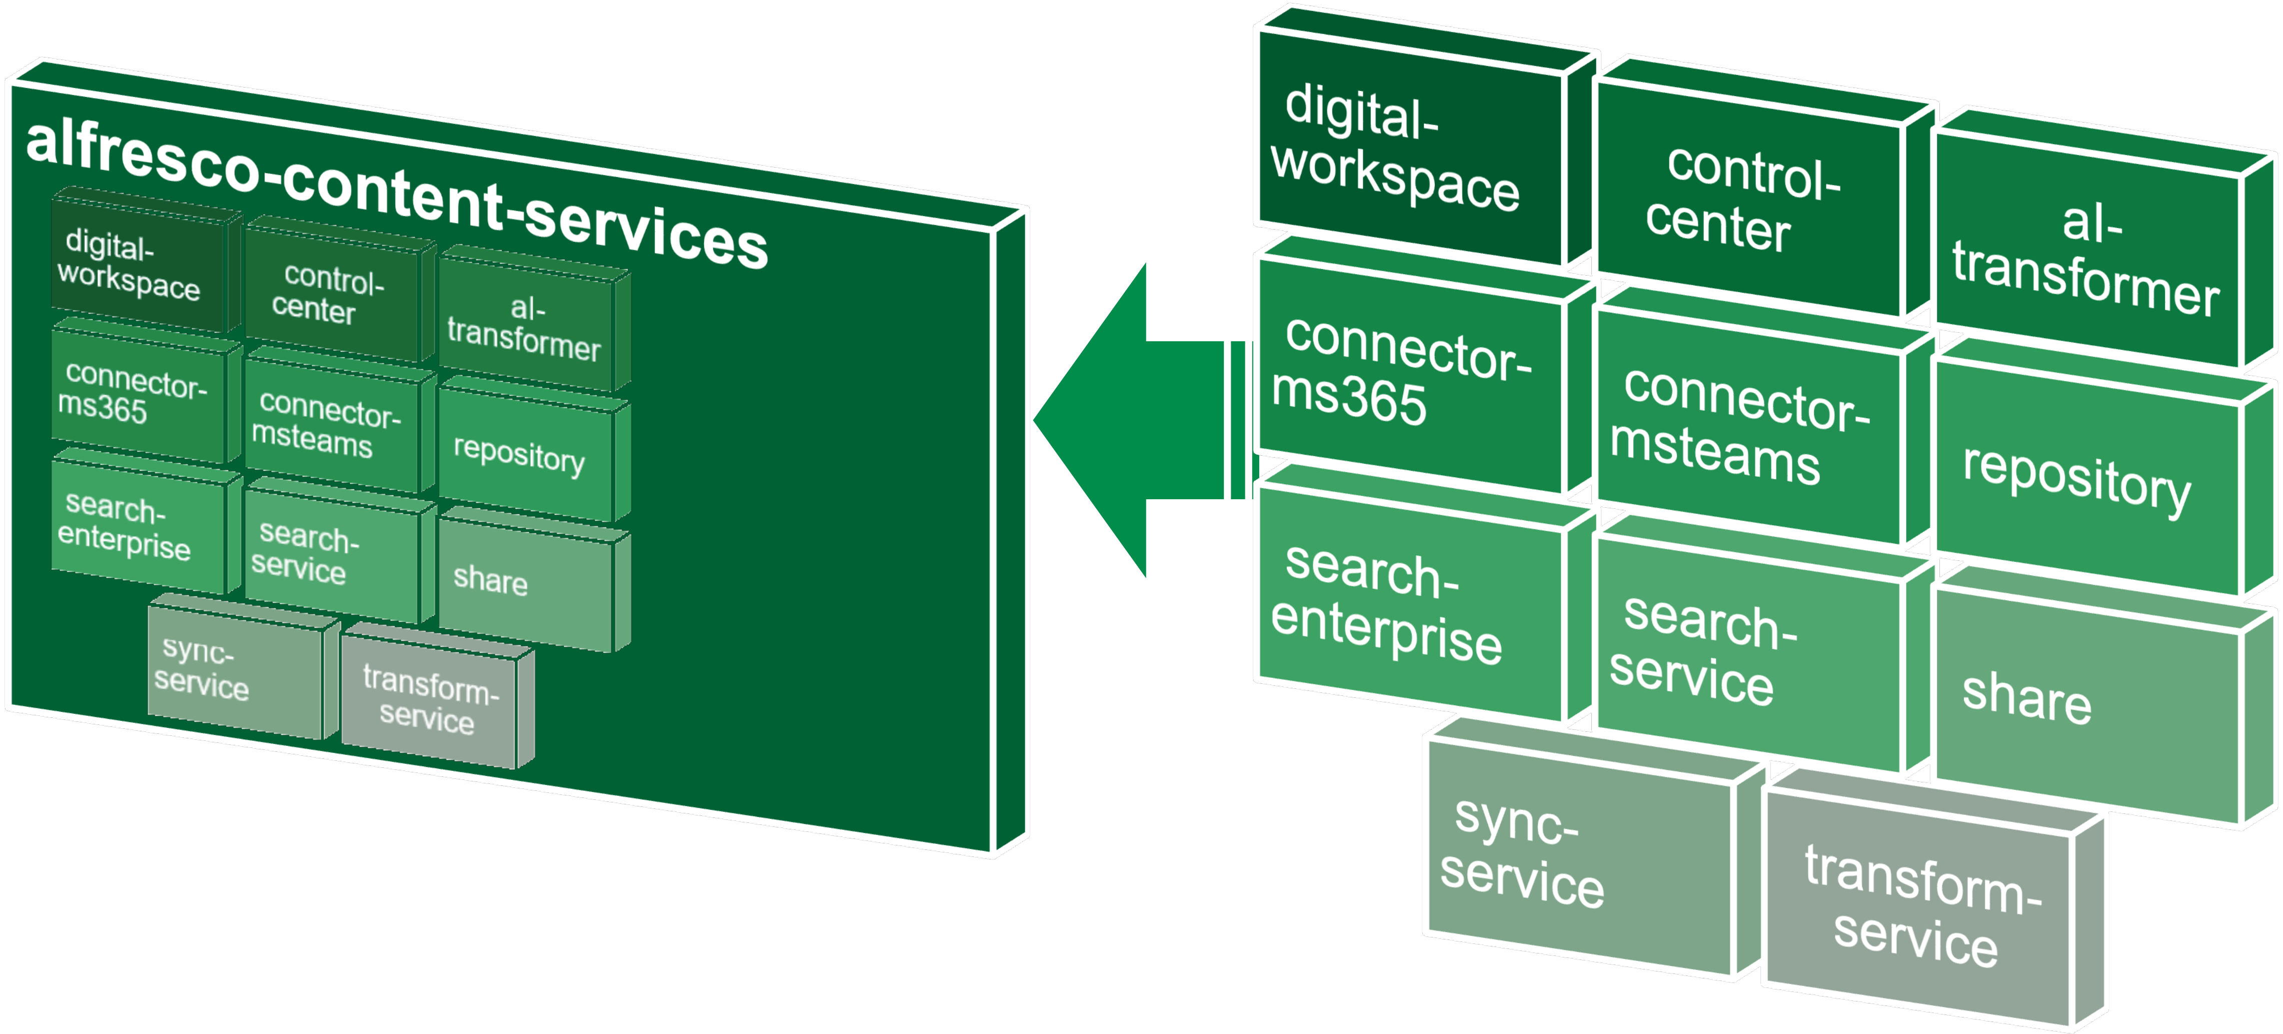

Components as building blocks for alfresco chart or your own chart

Creating your own chart is just plumbing

- ConfigMaps to share configuration parameters across charts (alfresco-global.properties is back)

- Secrets to share credentials

- alfresco-common library

chart brings helpers to ease creation of ConfigMap entries

- No third-party dependencies with component charts

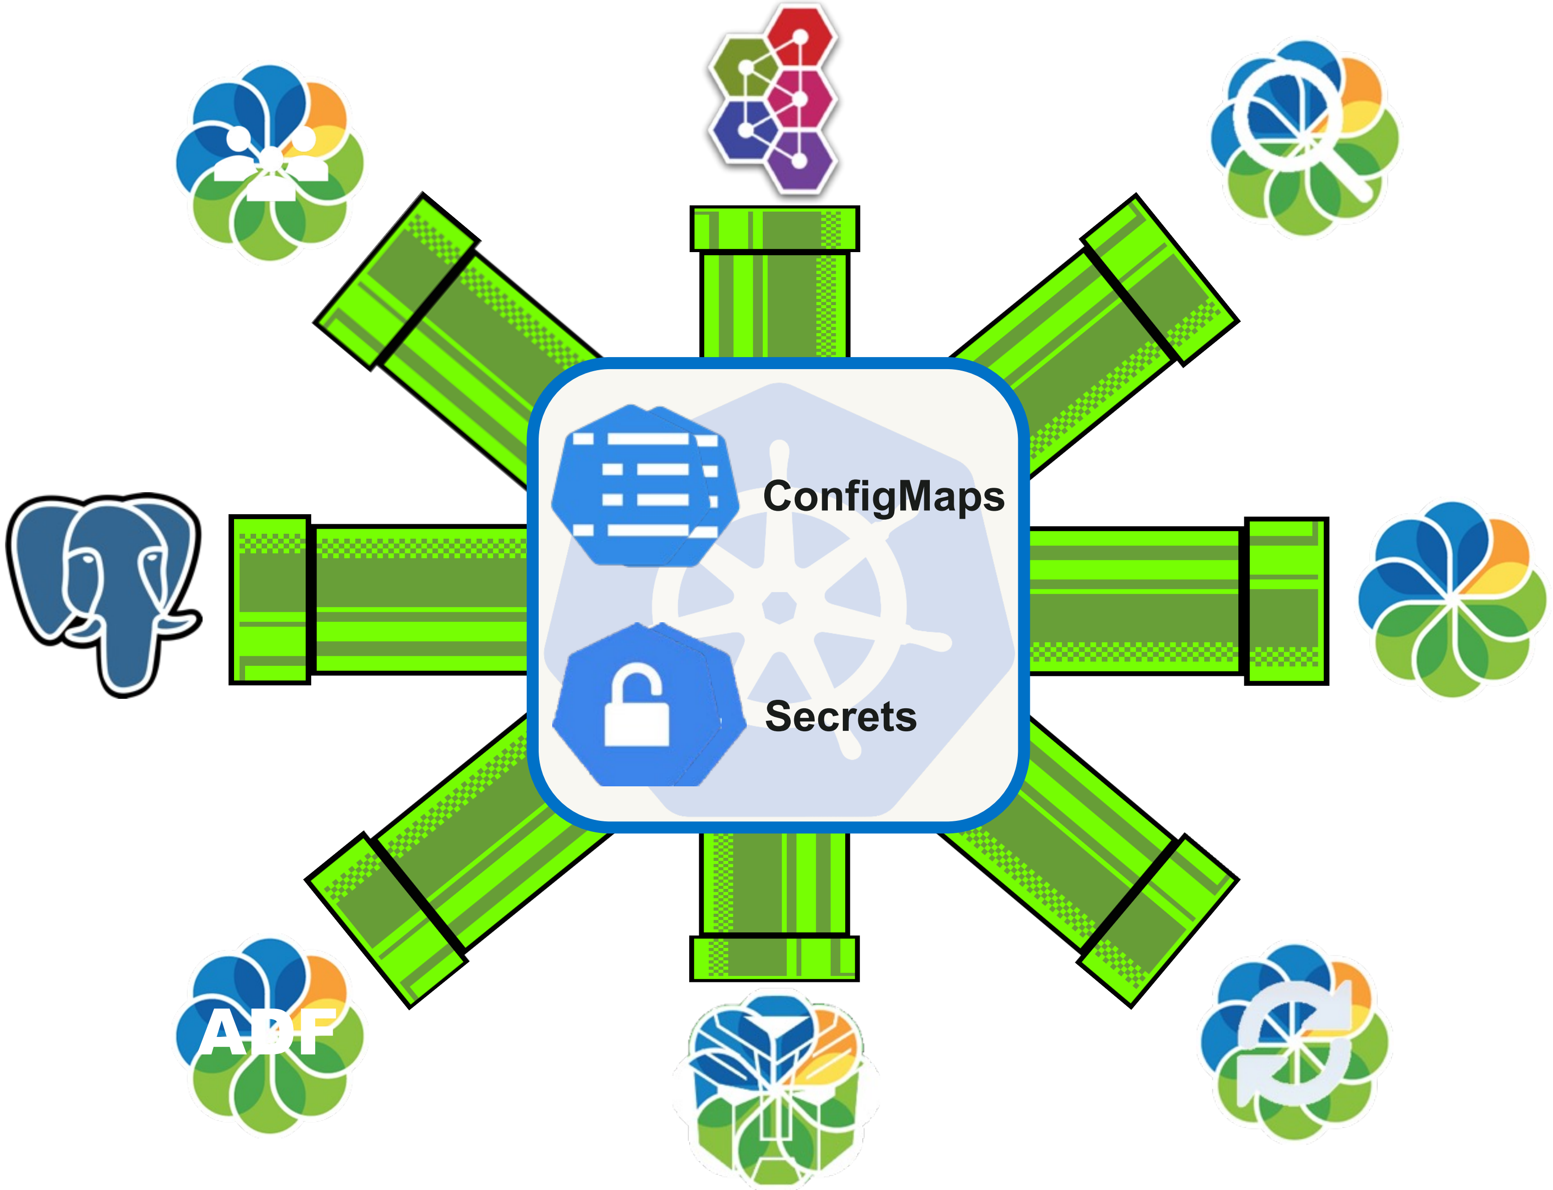

Example dependencies (Chart.yaml)

```
…
dependencies:
  - name: postgresql
    repository: oci://registry-1.docker.io/bitnamicharts
    version: 13.4.0
  - name: activemq
    repository: https://alfresco.github.io/alfresco-helm-charts/
    version: 3.4.1
  - name: alfresco-repository
    repository: https://alfresco.github.io/alfresco-helm-charts/
    version: 0.1.3
  - name: alfresco-transform-service
    version: 1.2.0
    repository: https://alfresco.github.io/alfresco-helm-charts/
  - name: my-custom-search-service
    repository: https://charts.partner.com/alfresco/
    version: 1.0.0
```

Example ConfigMap (ActiveMQ)

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: message-broker
data:
  {{- with .Values.activemq }}
  {{- $mqCtx := dict "Values" . "Chart" $.Chart "Release" $.Release }}
  MQ_URL: {{ include "alfresco-common.activemq.url.withFailover" (printf "nio://%s-broker:%v" (include "activemq.fullname" $mqCtx) 61616) | quote }}
  {{- end }}
```

- **Principles**:
  - Build the subchart context ($mqCtx)
  - Create a ConfigMap item to pass the expected information the target chart expects (Check individual README.md)
  - Leverage common helpers (alfresco-common.activemq.url.withFailover) & source charts' named templates (activemq.fullname) to reliably build the ActiveMQ URL

Example ConfigMap (ActiveMQ)

```
alfresco-repository:
  configuration:
    messageBroker:
      existingConfigMap:
        name: message-broker
        keys:
          url: MQ_URL
       existingSecret:
         name: message-broker
         keys:
           username: MQ_USER
           password: MQ_PASS
```

- **Principles**:
  - Leverage configuration values as expected by the target chart (Check individual README.md) to point to the right key in the right ConfigMap.

There's more than one way to do it

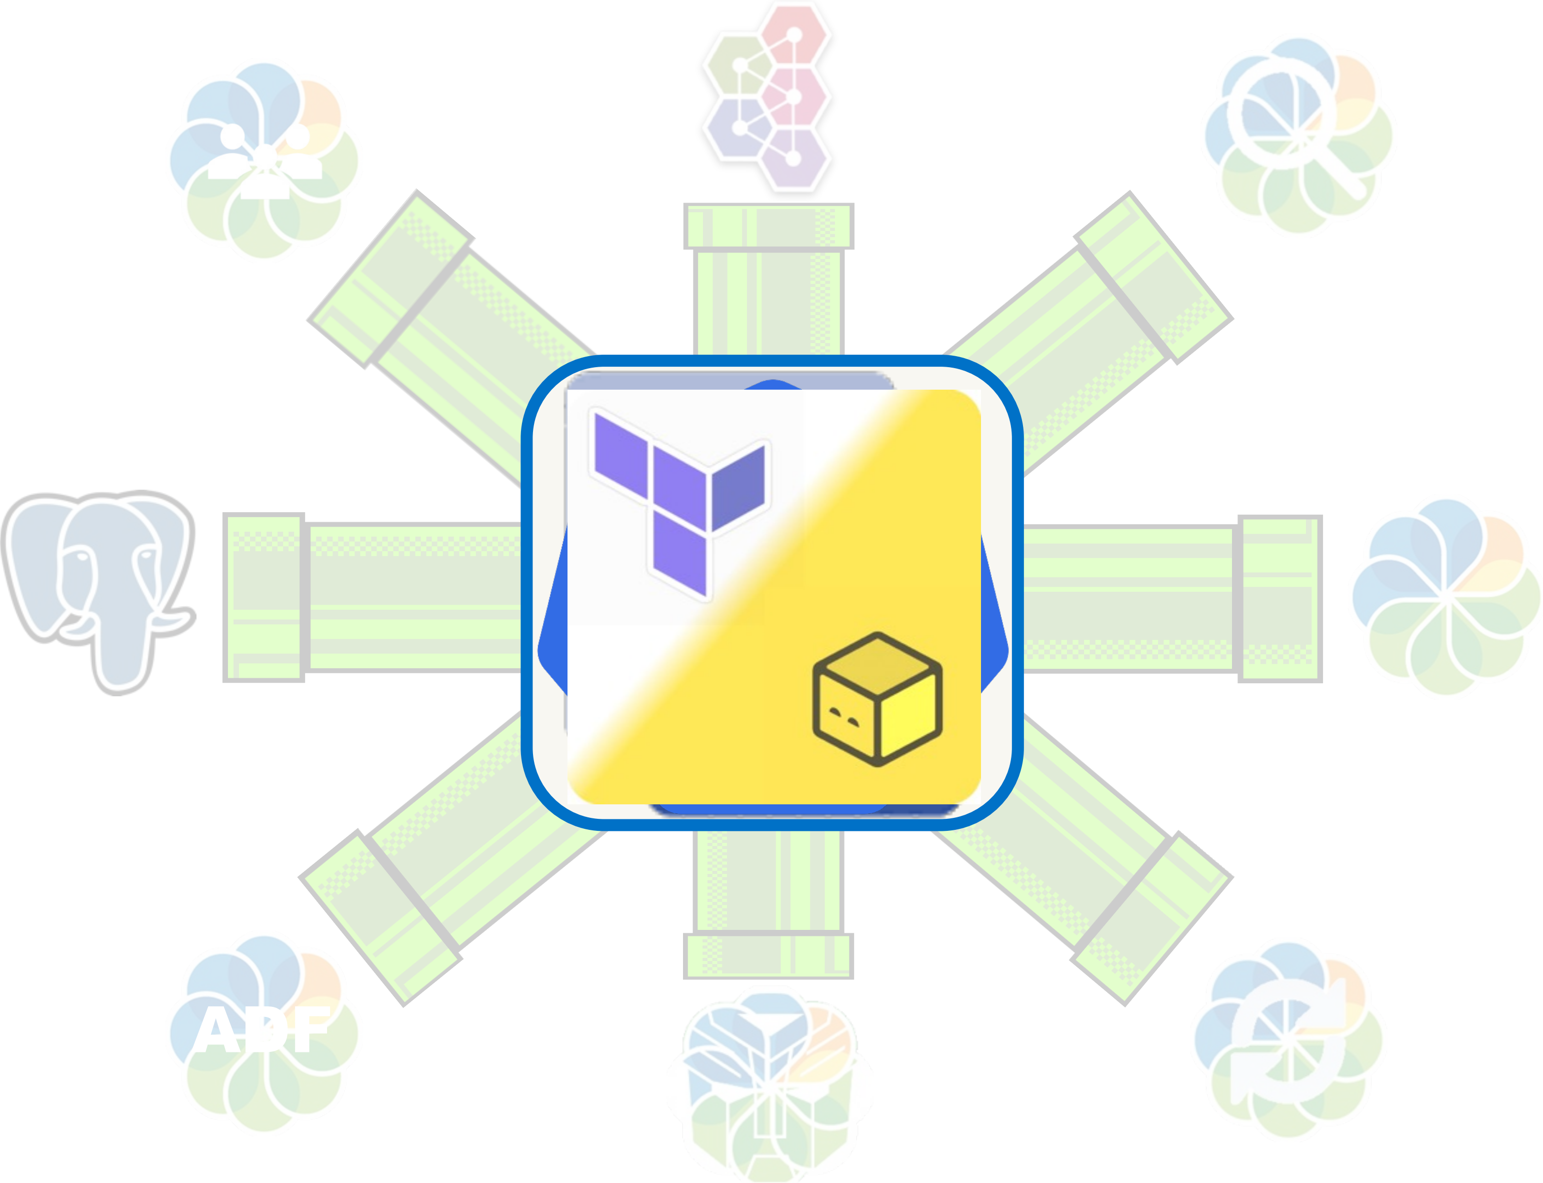

More details:

acs-sso-example available in acs-deployment with step-by-step documentation

Standardized means of configuration

- extraVolumes & extraVolumeMounts
  - Can be used to mount additional files from ConfigMaps (e.g. config files, keystores)
- extraInitContainers & extraSideContainers
  - Preliminary actions
  - System integration
- extraEnv
  - Additional config options through environment variables
- nodeSelector
- tolerations
- affinity
- podLabels & podAnnotations

Compositions examples

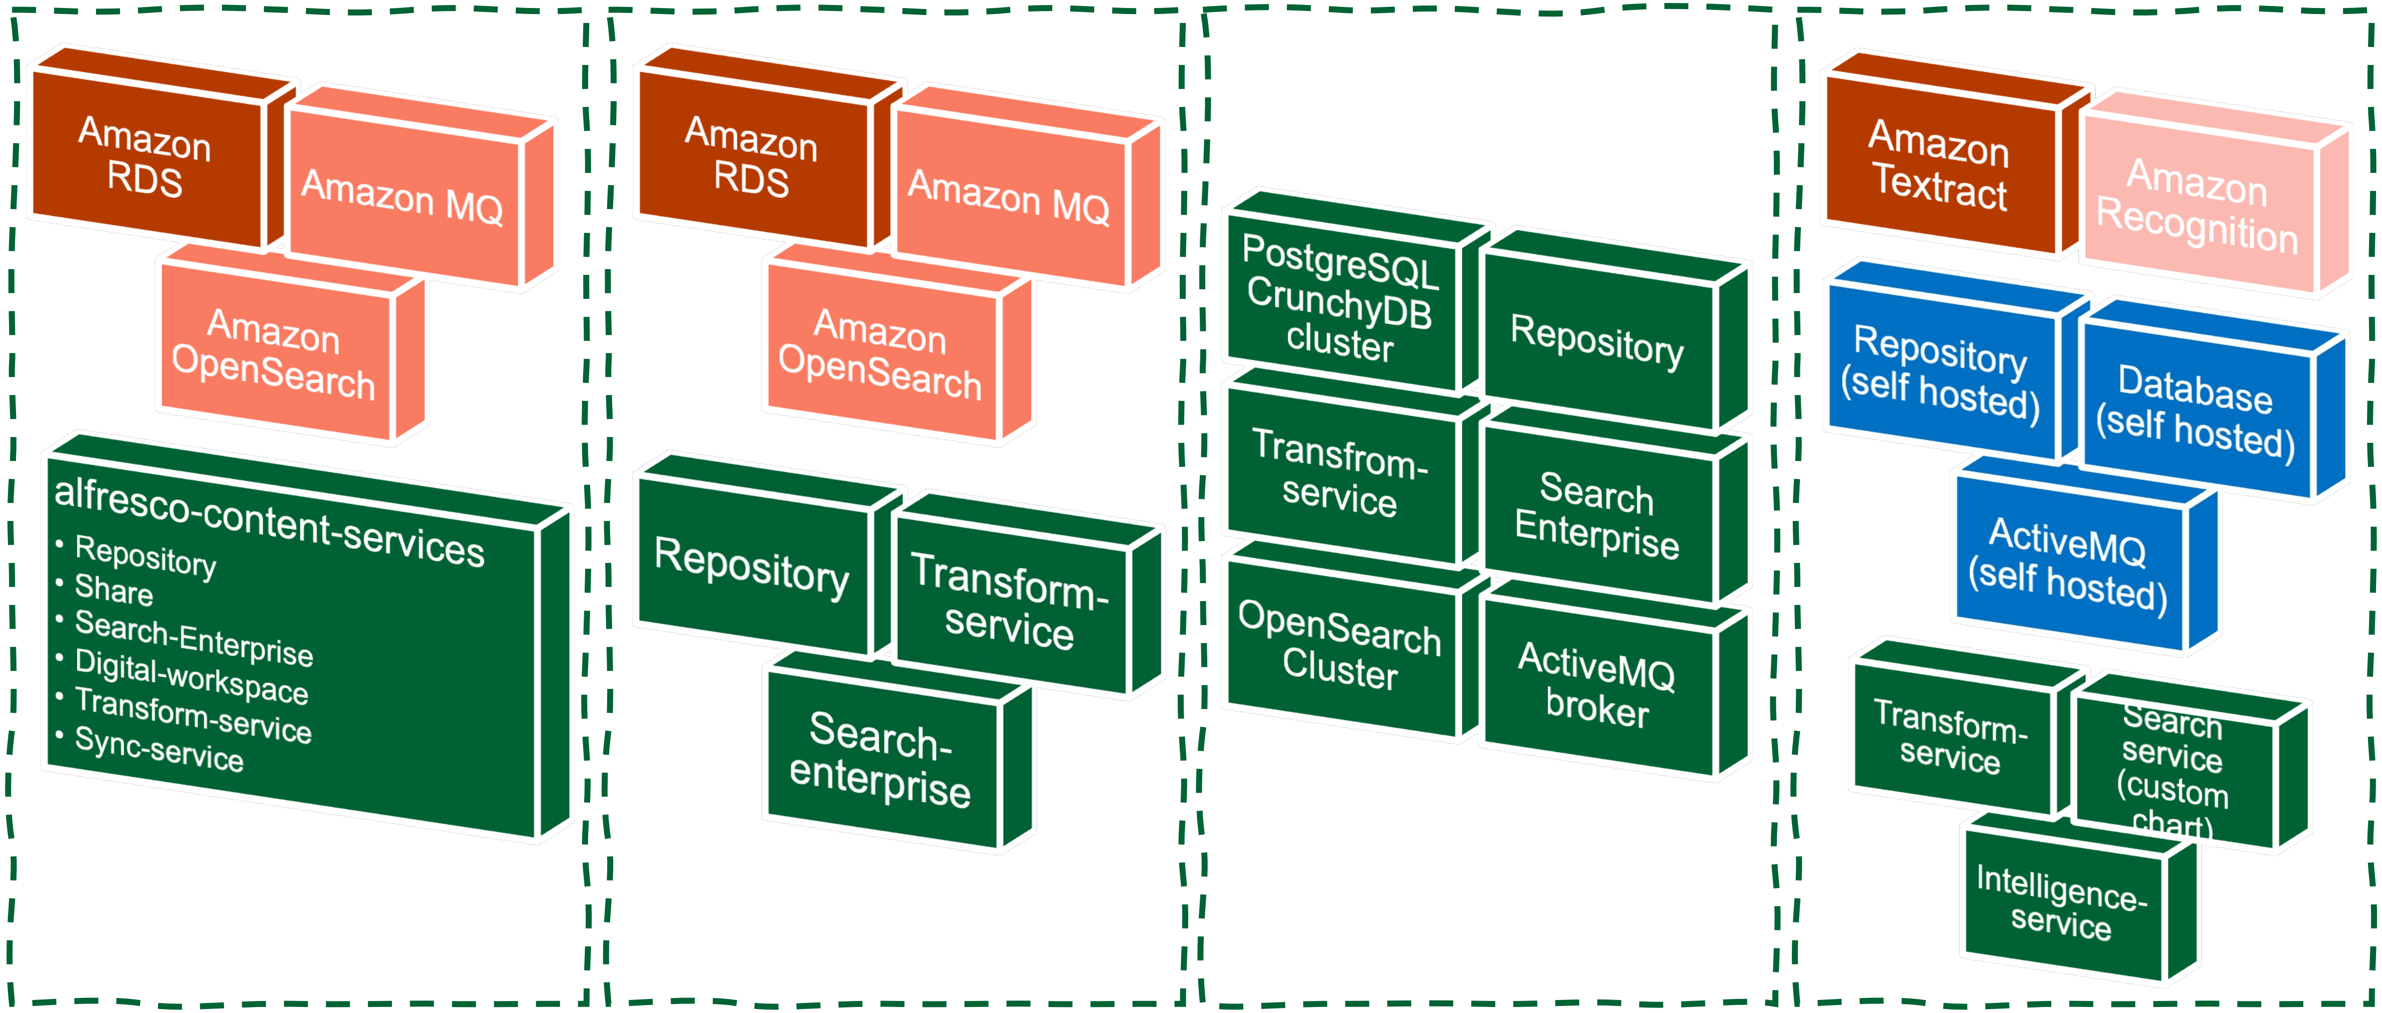

Amazon RDS

Amazon MQ

Amazon OpenSearch

alfresco-content-services

- Repository
- Share
- Search-Enterprise
- Digital-workspace
- Transform-service
- Sync-service

Amazon RDS

Amazon MQ

Amazon OpenSearch

Repository

Transform-service

Search-enterprise

PostgreSQL CrunchyDB cluster

Repository

Transfrom-service

Search Enterprise

OpenSearch Cluster

ActiveMQ broker

Amazon Textract

Amazon Recognition

Repository (self hosted)

Database (self hosted)

ActiveMQ (self hosted)

Transform-service

Search service (custom chart)

Intelligence-service

Charts Repositories

- New git repository: https://github.com/Alfresco/alfresco-helm-charts
- New chart registry: https://alfresco.github.io/alfresco-helm-charts/
- Each chart now has its own documentation
  - Main README.md for values
  - /docs/* for specific use cases
- Charts are tested & released using standard Helm tooling on KinD
  - Chart Tester (ct) & Chart Releaser (cr)
- Semver:
  - Major bump: expect breaking changes in architecture or configuration interface
  - Minor bump: new features or non-breaking changes
  - Patch bump: bugfix only

Beyond deployment modularity

Elasticity

- Deployments tend to be sized to cope with the peak workload
- Autoscaling can help lowering the resources impact & costs
- Alfresco repository and Alfresco Transform Service good first candidates
- CPU based scaling & investigate other metrics

Graviton deployment

Docker multi-arch support requirement

- Alfresco ACS 23.2 includes multi-arch docker images (almost)
- Check if a specific image is multi-arch:

```
$ docker buildx imagetools inspect quay.io/alfresco/alfresco-content-repository:23.2.1
Name:      quay.io/alfresco/alfresco-content-repository:23.2.1
MediaType: application/vnd.docker.distribution.manifest.list.v2+json
Digest:    sha256:8af8a3c1590b3318a0d292a20d7aeed4b253e897d8d4b4fc9352949812a0fabb

Manifests:
  Name:      quay.io/alfresco/alfresco-content-repository:23.2.1@sha256:eed0f4833bd1dfbce1828cdae535623ec59ba32d367a1e0e03dadcd60c196ffd
  MediaType: application/vnd.docker.distribution.manifest.v2+json
  Platform:  linux/amd64

  Name:      quay.io/alfresco/alfresco-content-repository:23.2.1@sha256:15f2bddae502fd2c59ed89a7eba18d0de9ab634841d1f76c5a106671e448cdae
  MediaType: application/vnd.docker.distribution.manifest.v2+json
  Platform:  linux/arm64
```

- Multi-arch docker image can be built easily with:

```
$ docker buildx build --platform linux/amd64,linux/arm64 -t my-namespace/my-image:latest .
```

Graviton testbed architecture

- Single EKS cluster
  - Graviton (arm64) node group with t4g.xlarge
    - 4 burstable vCPUs, 16.0 GiB of memory and up to 5 Gibps of bandwidth, starting at $0.1344 per hour - $96.76 per month per node
  - Classic (amd64) node group with t3a.xlarge
    - 4 burstable vCPUs, 16.0 GiB of memory and up to 5 Gibps of bandwidth, starting at $0.1504 per hour - $108.28 per month per node

- Helm install of Alfresco against each node group
  - (Deployment would fit 2x nodes each with room for scaling up)

Provision EKS cluster with graviton nodes

- Deploy with Terraform EKS module

```
[..]
 eks_managed_node_groups = {
    graviton = {
      min_size     = 0
      max_size     = 2
      desired_size = 1

      instance_types = ["t4g.xlarge"]
      ami_type       = "AL2_ARM_64"
    }
  }
[..]
```

Cluster Autoscaler

Automatically adjust the desired size of node groups so that all pods have a place to run and there are no unneeded nodes.

- Runs as a deployment within the cluster
- Should have permissions to examine and modify EC2 Auto Scaling Groups
- Do not require additional configuration except for fine tuning

Links:

- https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/cloudprovider/aws/README.md
- https://github.com/lablabs/terraform-aws-eks-cluster-autoscaler

nodeSelector to target nodeGroup

- nodeSelector is a Pod spec to tell K8s scheduler to run only on nodes that match that constraint
- Helm charts allow specifying a nodeSelector via values for each component:

```
[..]
alfresco-repository:
  nodeSelector:
    "kubernetes.io/arch": arm64
share:
  nodeSelector:
    "kubernetes.io/arch": arm64
[..]
```

Install alfresco on EKS

- Values oriented for testing purposes only:
  - No persistency
  - 1 replicas

```
$ helm upgrade --install acs alfresco/alfresco-content-services --version ~8.0.0
  --namespace alfresco
  --set global.search.sharedSecret=$(openssl rand -hex 24)
  --values helm-install/acs_values.yml
  --values helm-install/arm64_selector.yml
```

acs_values.yaml excerpt

```
alfresco-repository:
  ingress:
    annotations:
      cert-manager.io/cluster-issuer: letsencrypt-prod
  persistence:
    enabled: false
  configuration:
    search:
      flavor: elasticsearch
[..]
```

arm64_selector.yml excerpt

```
alfresco-repository:
  nodeSelector:
    "kubernetes.io/arch": amd64
  ingress:
    hosts:
      - host: FQDN
        paths:
          - path: /
            pathType: Prefix
    tls:
      - hosts:
          - FQDN
        secretName: letsencrypt-prod-hydra
[..]
```

Sources at https://github.com/gionn/alfresco-eks-graviton-example

Cluster dashboard

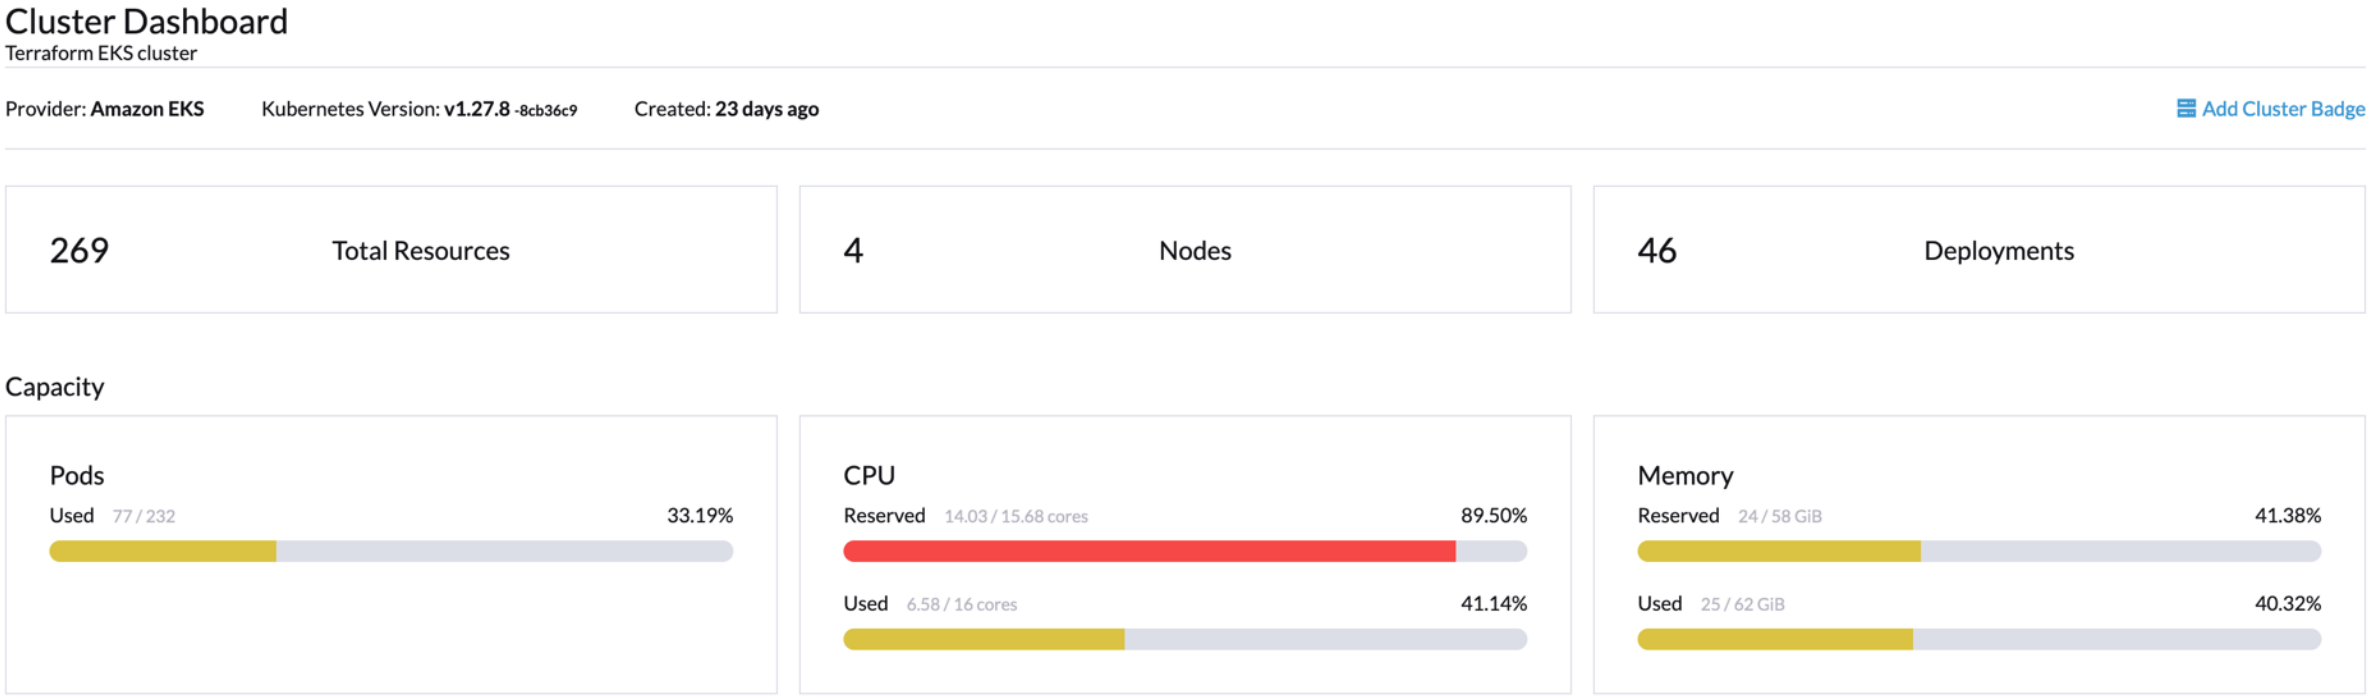

Cluster Dashboard

Terraform EKS cluster

Provider: **Amazon EKS** Kubernetes Version: **v1.27.8** -8cb36c9 Created: **23 days ago**

269 Total Resources

4 Nodes

46 Deployments

Capacity

Pods

Used 77 / 232 33.19%

CPU

Reserved 14.03 / 15.68 cores 89.50%

Used 6.58 / 16 cores 41.14%

Memory

Reserved 24 / 58 GiB 41.38%

Used 25 / 62 GiB 40.32%

Nodes overview

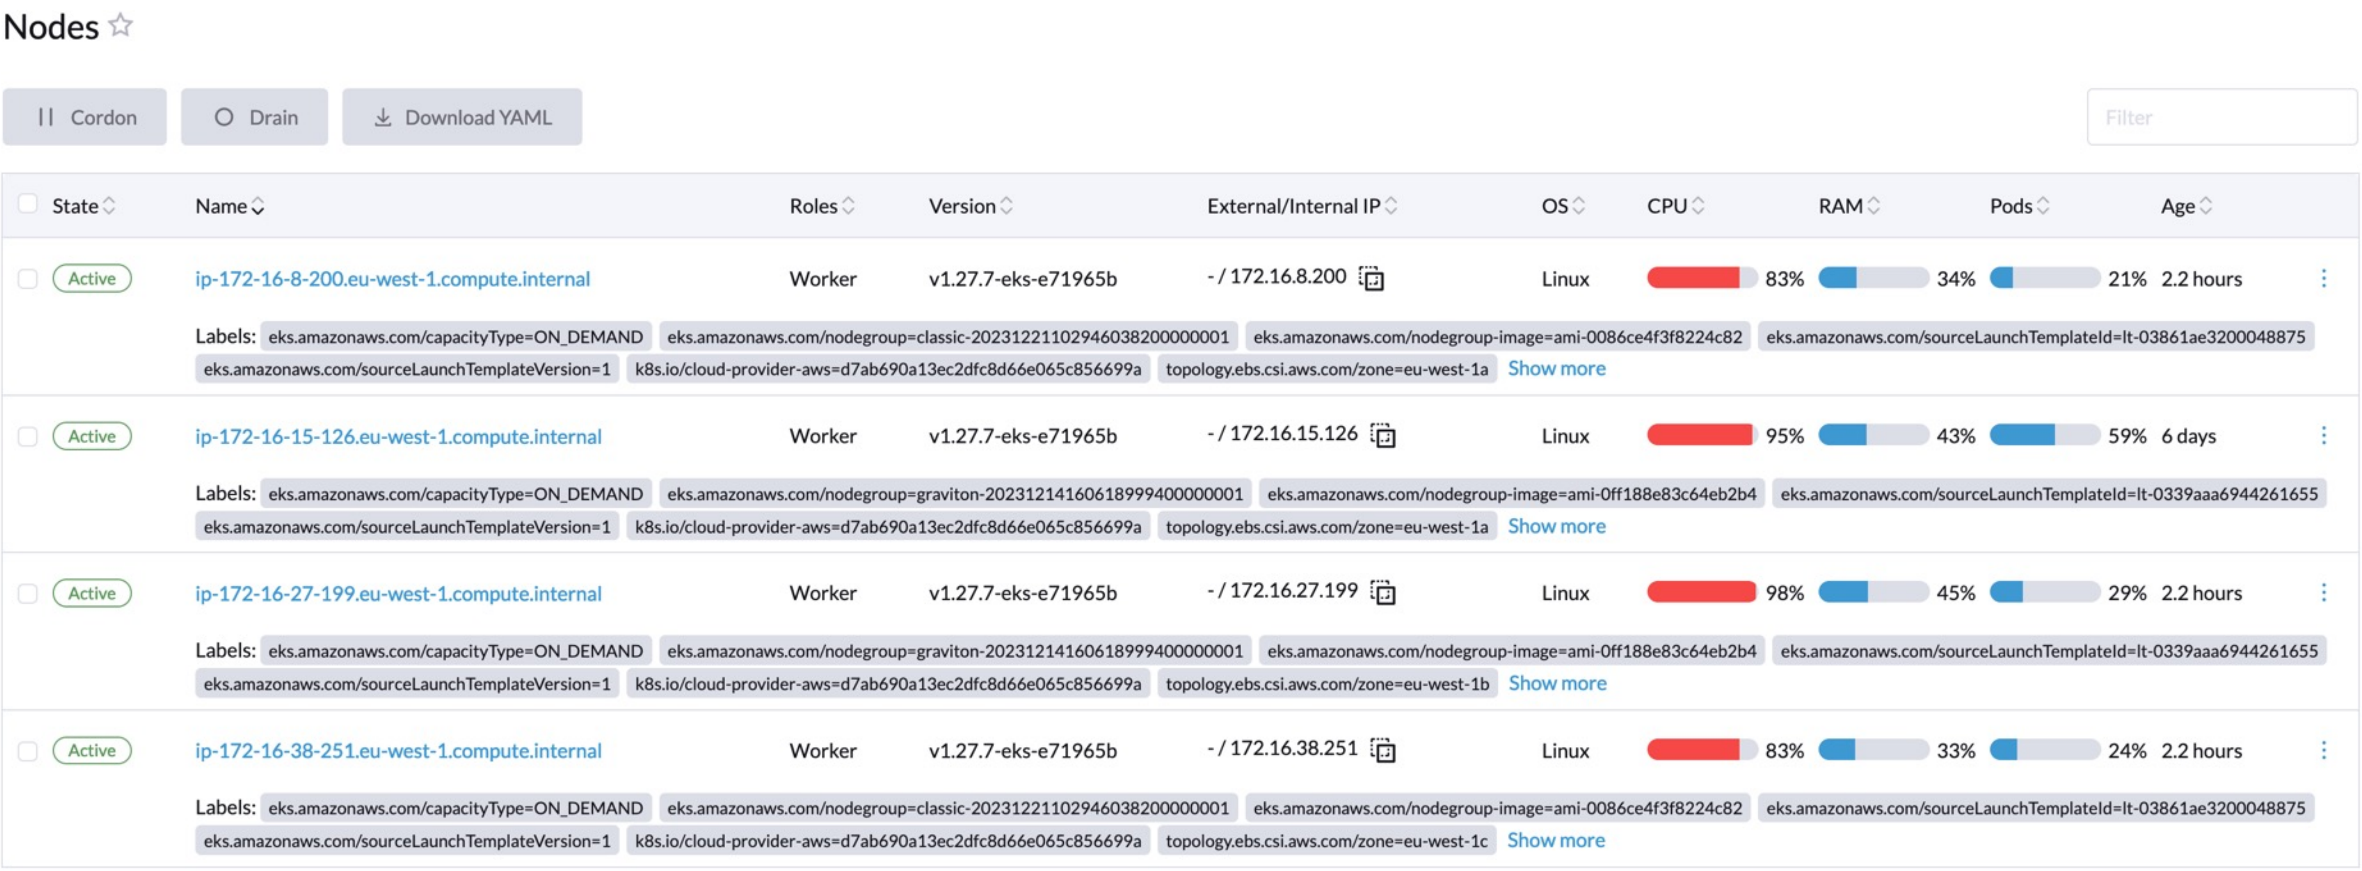

Nodes ☆

|| Cordon ○ Drain ⤓ Download YAML

Filter

□	State	Name	Roles	Version	External/Internal IP	OS	CPU	RAM	Pods	Age
□	Active	ip-172-16-8-200.eu-west-1.compute.internal	Worker	v1.27.7-eks-e71965b	- / 172.16.8.200	Linux	83%	34%	21%	2.2 hours
		Labels: eks.amazonaws.com/capacityType=ON_DEMAND eks.amazonaws.com/nodegroup=classic-20231221102946038200000001 eks.amazonaws.com/nodegroup-image=ami-0086ce4f3f8224c82 eks.amazonaws.com/sourceLaunchTemplateId=lt-03861ae3200048875 eks.amazonaws.com/sourceLaunchTemplateVersion=1 k8s.io/cloud-provider-aws=d7ab690a13ec2dfc8d66e065c856699a topology.ebs.csi.aws.com/zone=eu-west-1a Show more								
□	Active	ip-172-16-15-126.eu-west-1.compute.internal	Worker	v1.27.7-eks-e71965b	- / 172.16.15.126	Linux	95%	43%	59%	6 days
		Labels: eks.amazonaws.com/capacityType=ON_DEMAND eks.amazonaws.com/nodegroup=graviton-20231214160618999400000001 eks.amazonaws.com/nodegroup-image=ami-0ff188e83c64eb2b4 eks.amazonaws.com/sourceLaunchTemplateId=lt-0339aaa6944261655 eks.amazonaws.com/sourceLaunchTemplateVersion=1 k8s.io/cloud-provider-aws=d7ab690a13ec2dfc8d66e065c856699a topology.ebs.csi.aws.com/zone=eu-west-1a Show more								
□	Active	ip-172-16-27-199.eu-west-1.compute.internal	Worker	v1.27.7-eks-e71965b	- / 172.16.27.199	Linux	98%	45%	29%	2.2 hours
		Labels: eks.amazonaws.com/capacityType=ON_DEMAND eks.amazonaws.com/nodegroup=graviton-20231214160618999400000001 eks.amazonaws.com/nodegroup-image=ami-0ff188e83c64eb2b4 eks.amazonaws.com/sourceLaunchTemplateId=lt-0339aaa6944261655 eks.amazonaws.com/sourceLaunchTemplateVersion=1 k8s.io/cloud-provider-aws=d7ab690a13ec2dfc8d66e065c856699a topology.ebs.csi.aws.com/zone=eu-west-1b Show more								
□	Active	ip-172-16-38-251.eu-west-1.compute.internal	Worker	v1.27.7-eks-e71965b	- / 172.16.38.251	Linux	83%	33%	24%	2.2 hours
		Labels: eks.amazonaws.com/capacityType=ON_DEMAND eks.amazonaws.com/nodegroup=classic-20231221102946038200000001 eks.amazonaws.com/nodegroup-image=ami-0086ce4f3f8224c82 eks.amazonaws.com/sourceLaunchTemplateId=lt-03861ae3200048875 eks.amazonaws.com/sourceLaunchTemplateVersion=1 k8s.io/cloud-provider-aws=d7ab690a13ec2dfc8d66e065c856699a topology.ebs.csi.aws.com/zone=eu-west-1c Show more								

Deployment/statefulset overview

```
$ kubectl get deploy -n alfresco
NAME                                     READY   UP-TO-DATE   AVAILABLE
acs-activemq                             1/1     1            1
acs-alfresco-cc                          1/1     1            1
acs-alfresco-dw                          1/1     1            1
acs-alfresco-repository                  1/1     1            1
acs-alfresco-search-enterprise-content   1/1     1            1
acs-alfresco-search-enterprise-metadata  1/1     1            1
acs-alfresco-search-enterprise-path      1/1     1            1
acs-filestore                            1/1     1            1
acs-imagemagick                          1/1     1            1
acs-libreoffice                          1/1     1            1
acs-pdfrenderer                          1/1     1            1
acs-share                                1/1     1            1
acs-tika                                 1/1     1            1
acs-transform-misc                       1/1     1            1
acs-transform-router                     1/1     1            1
```

```
$ kubectl -n alfresco get statefulset
NAME                                   READY
acs-postgresql-acs                     1/1
alfresco-search-enterprise-mediation   1/1
elasticsearch-master                   1/1
```

Verify deployment with DTAS

- Deployment Test Automation Scripts (DTAS) is an internal (not yet opensource) pytest suite to verify the correctness of an Alfresco install
- API testing
  - content lifecycle: create user, create a new site, create a folder structure, uploading documents
  - Alfresco Transformation Service (ATS): requesting renditions for various document types such as docx, xlsx, pdf, odt, jpeg, and png
- Basic benchmarking using pytest-benchmark library
  - pytest suites executed multiple times within a configurable time frame or number of rounds

Run DTAS via helm charts on Alfresco Enterprise

- Helm install as usual plus `--set dtas.enabled=true`
- Run the suite via `helm test RELEASE`

```
NAME: acs
LAST DEPLOYED: Mon Mar 11 17:28:56 2024
NAMESPACE: alfresco
STATUS: deployed
REVISION: 1
TEST SUITE:     dtas-hbusb
Last Started:   Mon Mar 11 17:36:58 2024
Last Completed: Mon Mar 11 17:49:34 2024
Phase:          Succeeded
```

CPU/Memory overview (arm64 vs amd64)

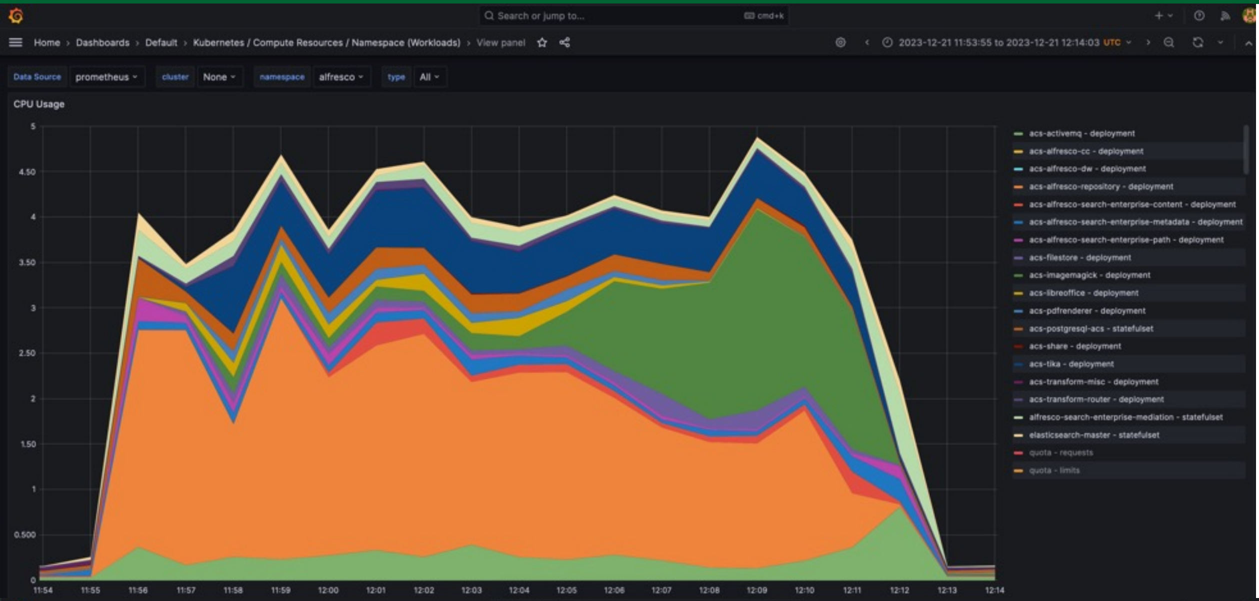

Home › Dashboards › Default › Kubernetes / Compute Resources / Namespace (Workloads) › View panel

2023-12-21 11:53:55 to 2023-12-21 12:14:03 UTC

Data Source prometheus · cluster None · namespace alfresco · type All

CPU Usage

Search or jump to...
cmd+k
5
4.50
4
3.50
3
2.50
2
1.50
1
0.500
0
11:54
11:55
11:56
11:57
11:58
11:59
12:00
12:01
12:02
12:03
12:04
12:05
12:06
12:07
12:08
12:09
12:10
12:11
12:12
12:13
12:14
acs-activemq - deployment
acs-alfresco-cc - deployment
acs-alfresco-dw - deployment
acs-alfresco-repository - deployment
acs-alfresco-search-enterprise-content - deployment
acs-alfresco-search-enterprise-metadata - deployment
acs-alfresco-search-enterprise-path - deployment
acs-filestore - deployment
acs-imagemagick - deployment
acs-libreoffice - deployment
acs-pdfrenderer - deployment
acs-postgresql-acs - statefulset
acs-share - deployment
acs-tika - deployment
acs-transform-misc - deployment
acs-transform-router - deployment
alfresco-search-enterprise-mediation - statefulset
elasticsearch-master - statefulset
quota - requests
quota - limits

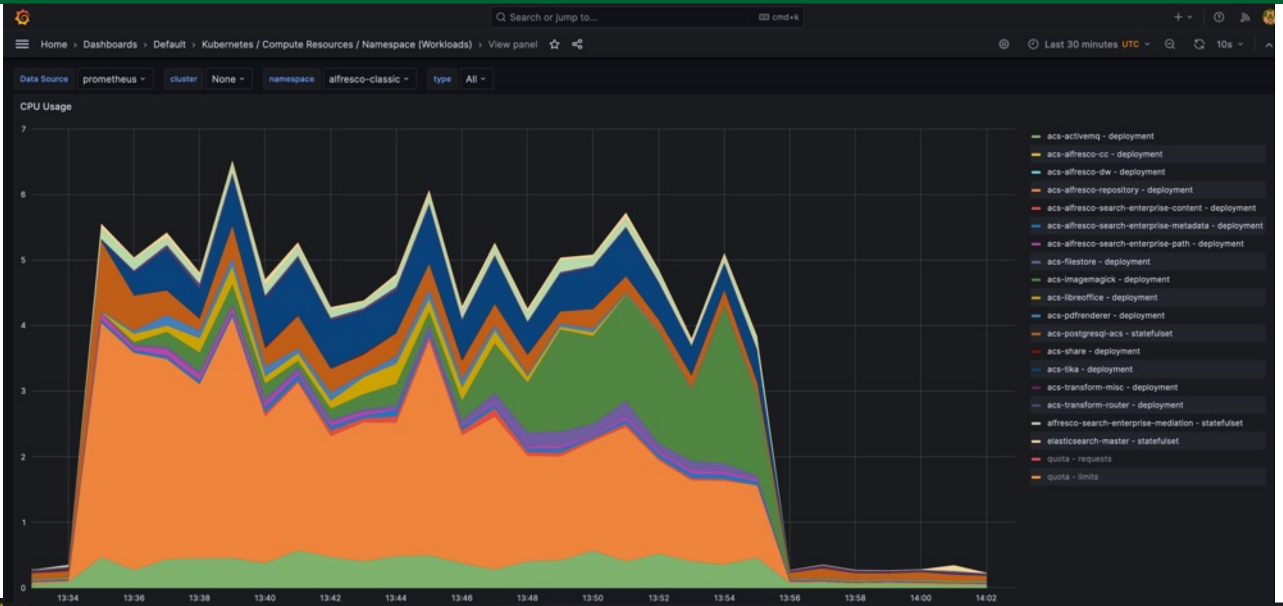

Home › Dashboards › Default › Kubernetes / Compute Resources / Namespace (Workloads) › View panel

Last 30 minutes UTC · 10s

Data Source prometheus · cluster None · namespace alfresco-classic · type All

CPU Usage

Search or jump to...
cmd+k
7
6
5
4
3
2
1
0
13:34
13:36
13:38
13:40
13:42
13:44
13:46
13:48
13:50
13:52
13:54
13:56
13:58
14:00
14:02
acs-activemq - deployment
acs-alfresco-cc - deployment
acs-alfresco-dw - deployment
acs-alfresco-repository - deployment
acs-alfresco-search-enterprise-content - deployment
acs-alfresco-search-enterprise-metadata - deployment
acs-alfresco-search-enterprise-path - deployment
acs-filestore - deployment
acs-imagemagick - deployment
acs-libreoffice - deployment
acs-pdfrenderer - deployment
acs-postgresql-acs - statefulset
acs-share - deployment
acs-tika - deployment
acs-transform-misc - deployment
acs-transform-router - deployment
alfresco-search-enterprise-mediation - statefulset
elasticsearch-master - statefulset
quota - requests
quota - limits

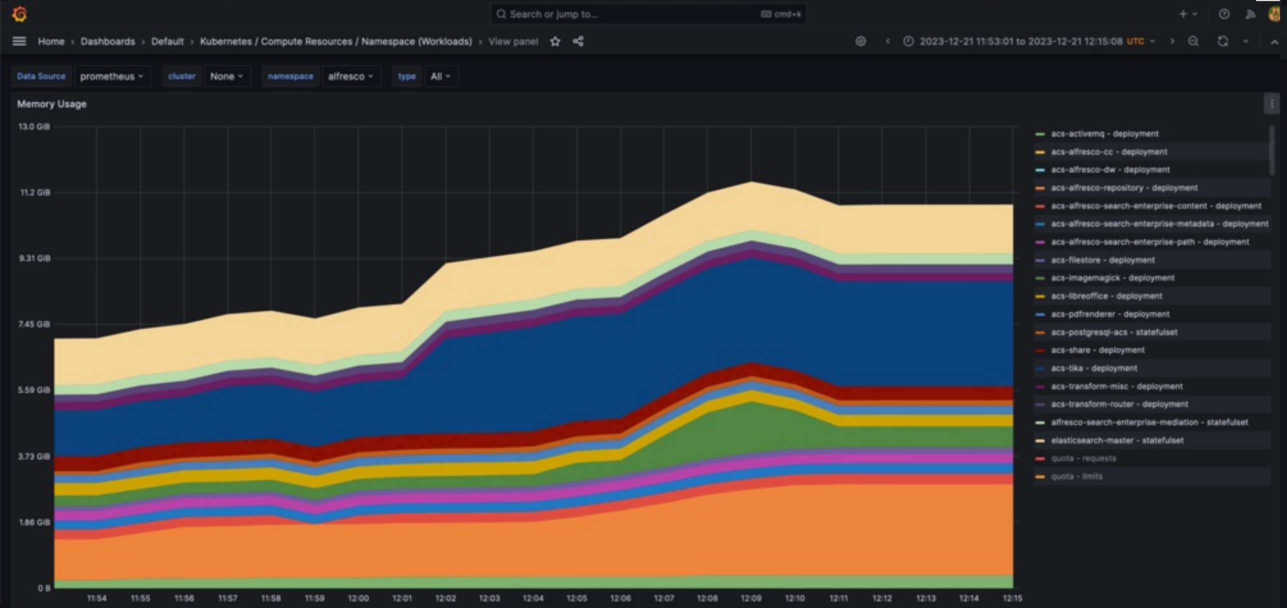

Home › Dashboards › Default › Kubernetes / Compute Resources / Namespace (Workloads) › View panel

2023-12-21 11:53:01 to 2023-12-21 12:15:08 UTC

Data Source prometheus · cluster None · namespace alfresco · type All

Memory Usage

Search or jump to...
cmd+k
13.0 GiB
11.2 GiB
9.31 GiB
7.45 GiB
5.59 GiB
3.73 GiB
1.86 GiB
0 B
11:54
11:55
11:56
11:57
11:58
11:59
12:00
12:01
12:02
12:03
12:04
12:05
12:06
12:07
12:08
12:09
12:10
12:11
12:12
12:13
12:14
12:15
acs-activemq - deployment
acs-alfresco-cc - deployment
acs-alfresco-dw - deployment
acs-alfresco-repository - deployment
acs-alfresco-search-enterprise-content - deployment
acs-alfresco-search-enterprise-metadata - deployment
acs-alfresco-search-enterprise-path - deployment
acs-filestore - deployment
acs-imagemagick - deployment
acs-libreoffice - deployment
acs-pdfrenderer - deployment
acs-postgresql-acs - statefulset
acs-share - deployment
acs-tika - deployment
acs-transform-misc - deployment
acs-transform-router - deployment
alfresco-search-enterprise-mediation - statefulset
elasticsearch-master - statefulset
quota - requests
quota - limits

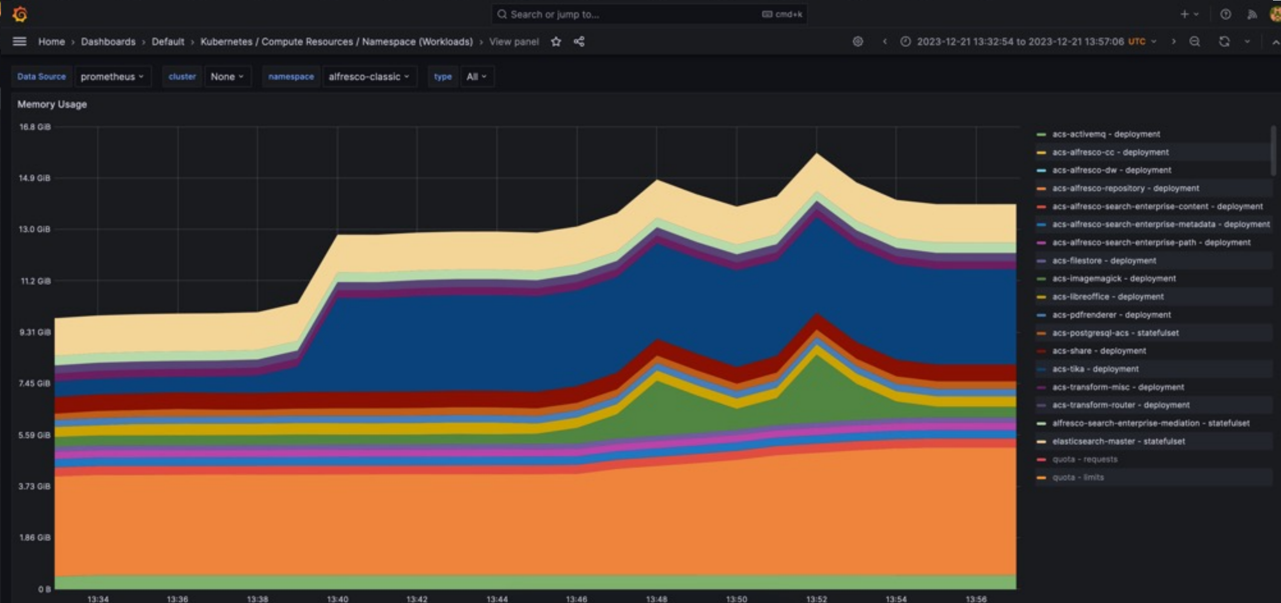

Home › Dashboards › Default › Kubernetes / Compute Resources / Namespace (Workloads) › View panel

2023-12-21 13:32:54 to 2023-12-21 13:57:06 UTC

Data Source prometheus · cluster None · namespace alfresco-classic · type All

Memory Usage

Search or jump to...
cmd+k
16.8 GiB
14.9 GiB
13.0 GiB
11.2 GiB
9.31 GiB
7.45 GiB
5.59 GiB
3.73 GiB
1.86 GiB
0 B
13:34
13:36
13:38
13:40
13:42
13:44
13:46
13:48
13:50
13:52
13:54
13:56
acs-activemq - deployment
acs-alfresco-cc - deployment
acs-alfresco-dw - deployment
acs-alfresco-repository - deployment
acs-alfresco-search-enterprise-content - deployment
acs-alfresco-search-enterprise-metadata - deployment
acs-alfresco-search-enterprise-path - deployment
acs-filestore - deployment
acs-imagemagick - deployment
acs-libreoffice - deployment
acs-pdfrenderer - deployment
acs-postgresql-acs - statefulset
acs-share - deployment
acs-tika - deployment
acs-transform-misc - deployment
acs-transform-router - deployment
alfresco-search-enterprise-mediation - statefulset
elasticsearch-master - statefulset
quota - requests
quota - limits

Network bandwidth overview (arm64 vs amd64)

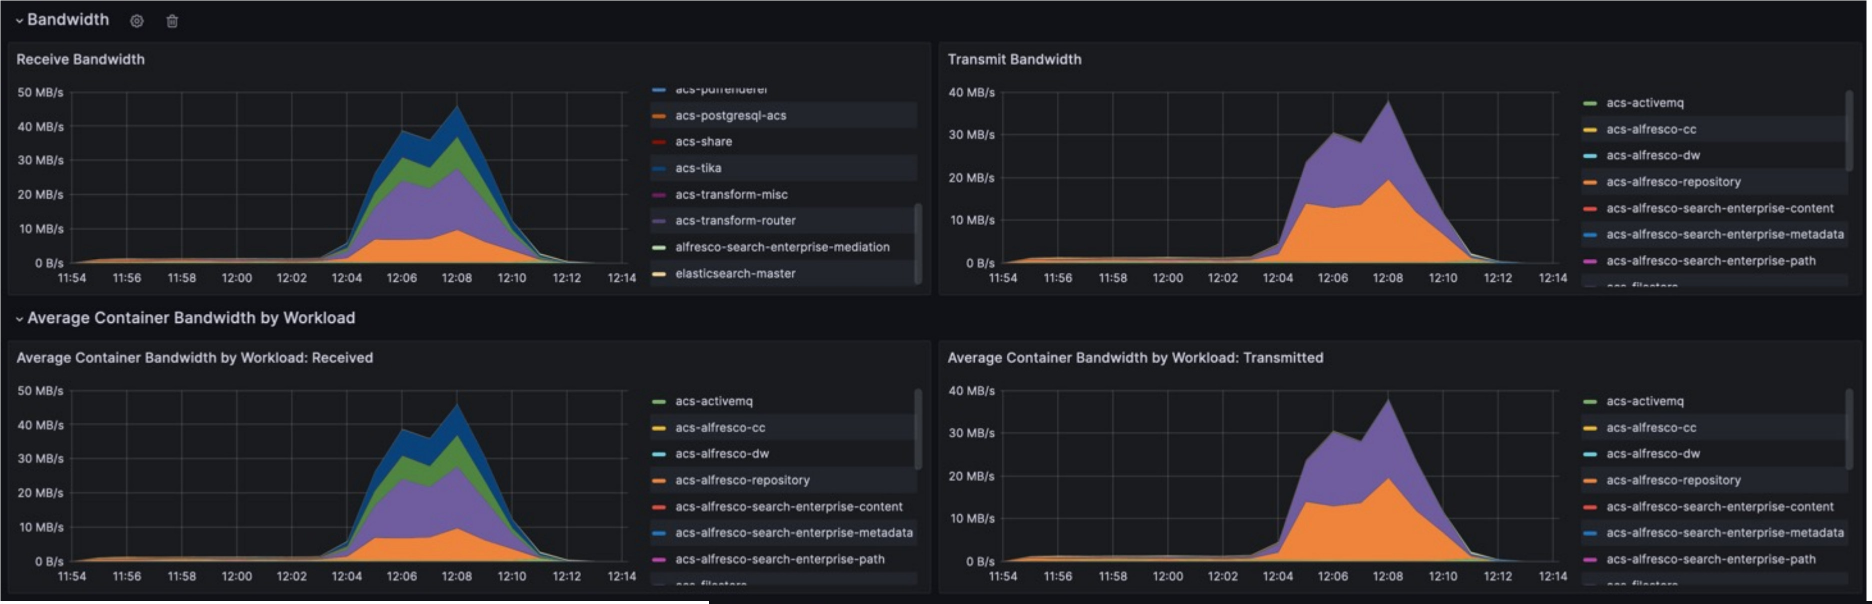

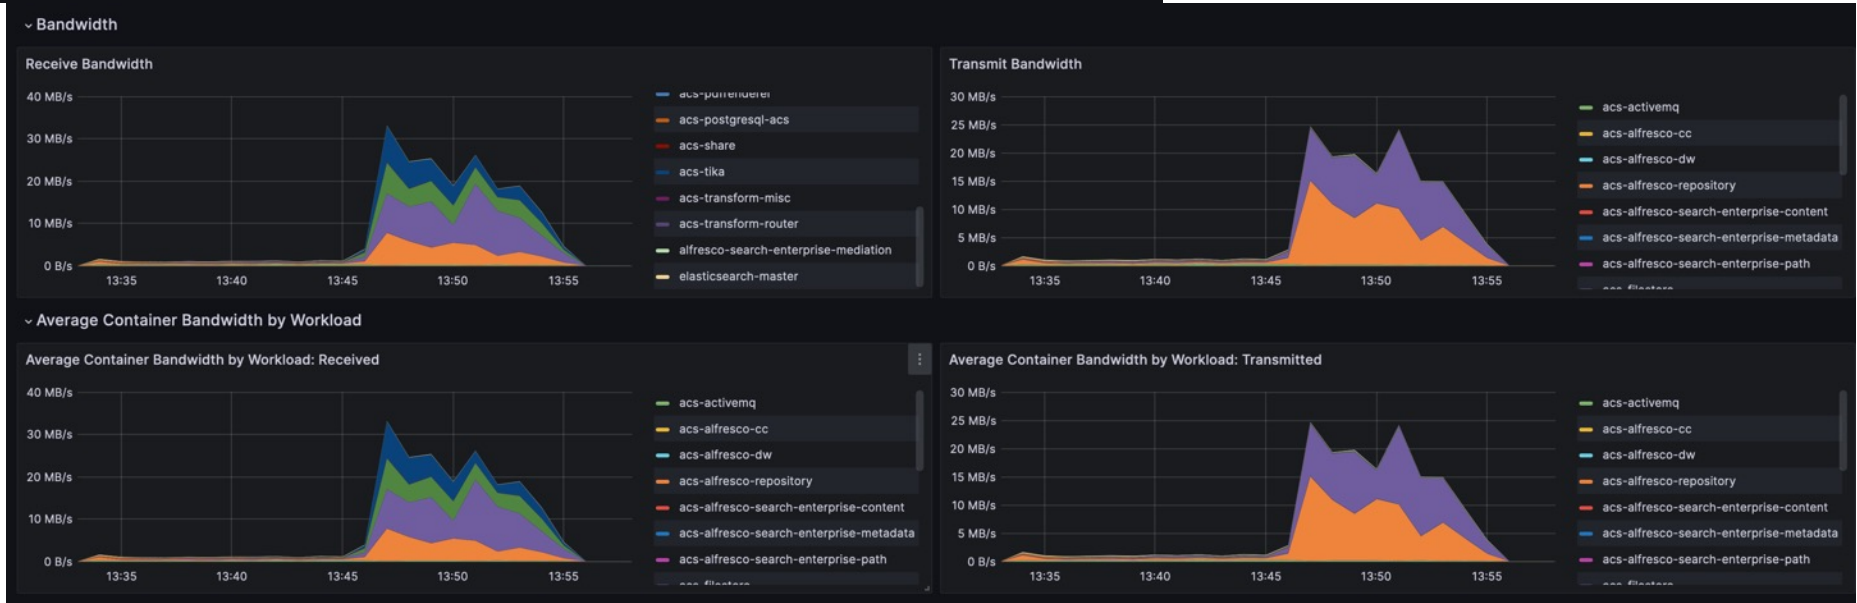

Benchmark results (arm64 vs amd64)

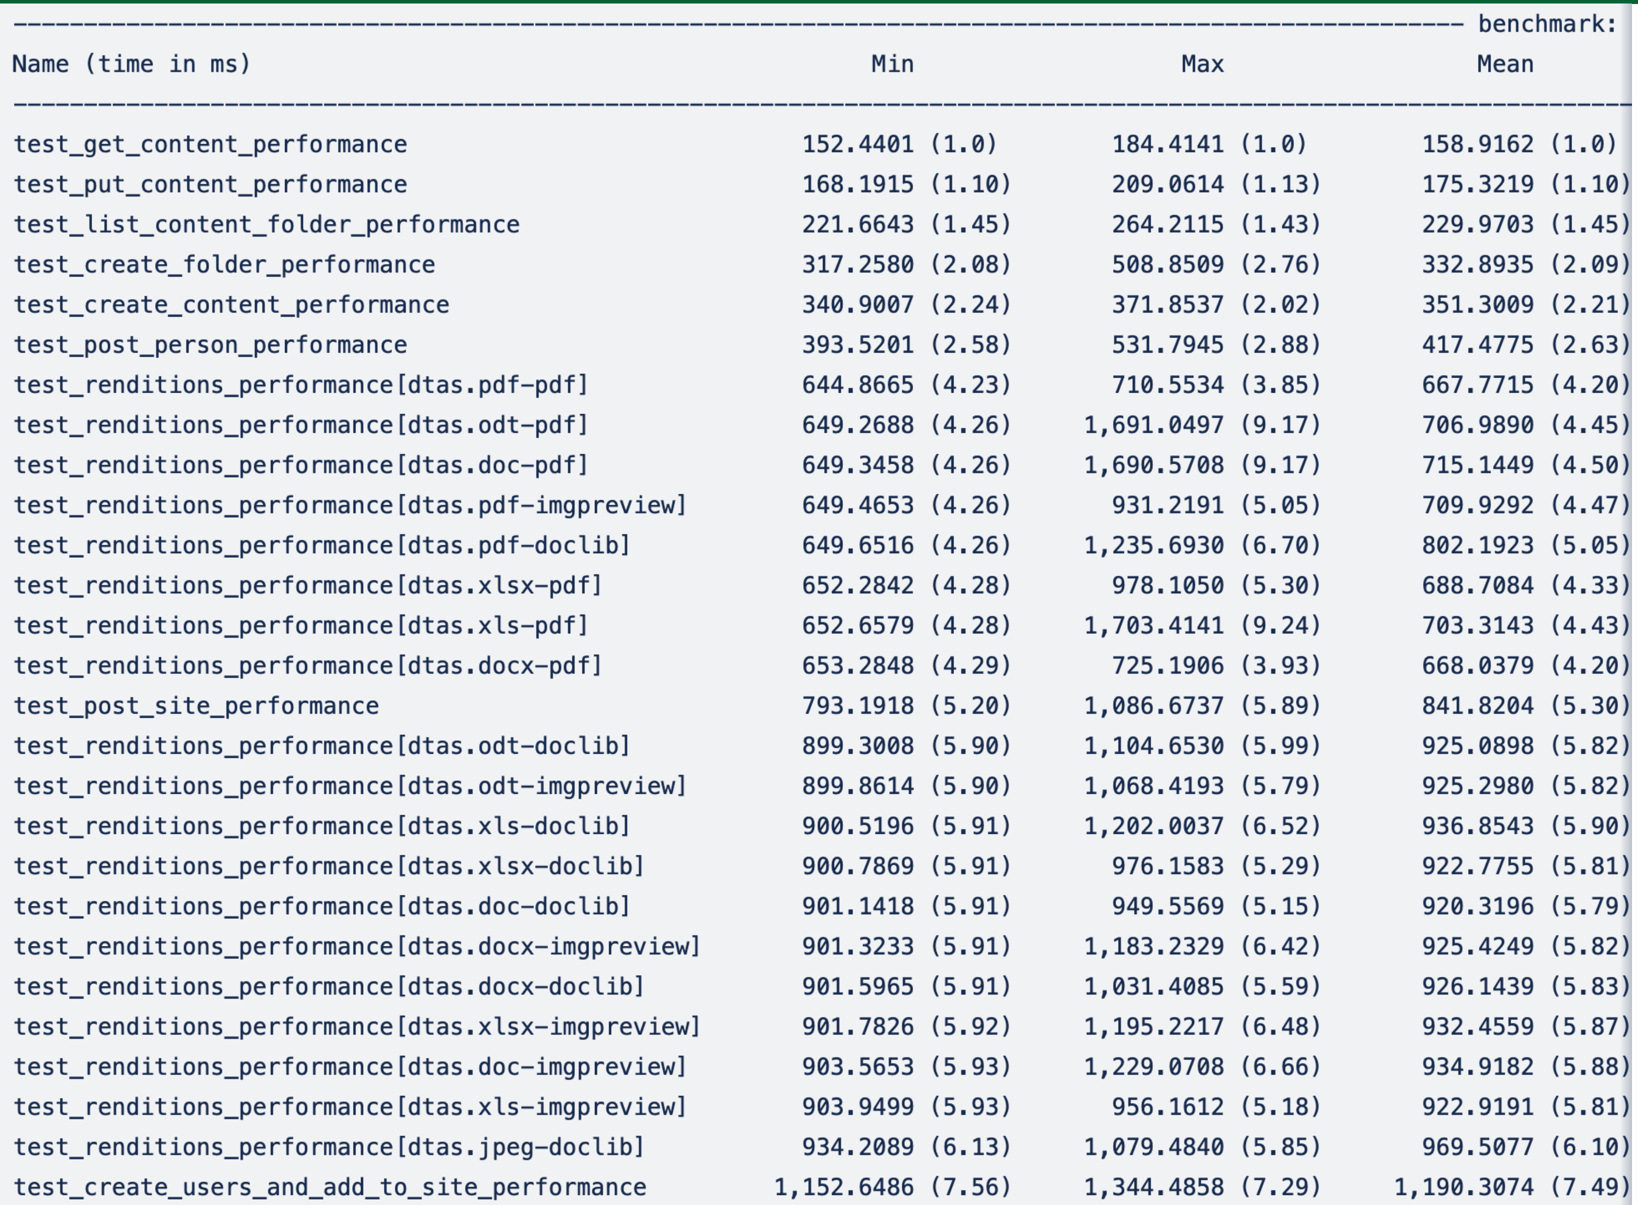

```
---------------------------------------------------------------------------------------------------------- benchmark:
Name (time in ms)                                             Min                    Max                   Mean
---------------------------------------------------------------------------------------------------------------------
test_get_content_performance                           152.4401 (1.0)        184.4141 (1.0)        158.9162 (1.0)
test_put_content_performance                           168.1915 (1.10)       209.0614 (1.13)       175.3219 (1.10)
test_list_content_folder_performance                   221.6643 (1.45)       264.2115 (1.43)       229.9703 (1.45)
test_create_folder_performance                         317.2580 (2.08)       508.8509 (2.76)       332.8935 (2.09)
test_create_content_performance                        340.9007 (2.24)       371.8537 (2.02)       351.3009 (2.21)
test_post_person_performance                           393.5201 (2.58)       531.7945 (2.88)       417.4775 (2.63)
test_renditions_performance[dtas.pdf-pdf]              644.8665 (4.23)       710.5534 (3.85)       667.7715 (4.20)
test_renditions_performance[dtas.odt-pdf]              649.2688 (4.26)     1,691.0497 (9.17)       706.9890 (4.45)
test_renditions_performance[dtas.doc-pdf]              649.3458 (4.26)     1,690.5708 (9.17)       715.1449 (4.50)
test_renditions_performance[dtas.pdf-imgpreview]       649.4653 (4.26)       931.2191 (5.05)       709.9292 (4.47)
test_renditions_performance[dtas.pdf-doclib]           649.6516 (4.26)     1,235.6930 (6.70)       802.1923 (5.05)
test_renditions_performance[dtas.xlsx-pdf]             652.2842 (4.28)       978.1050 (5.30)       688.7084 (4.33)
test_renditions_performance[dtas.xls-pdf]              652.6579 (4.28)     1,703.4141 (9.24)       703.3143 (4.43)
test_renditions_performance[dtas.docx-pdf]             653.2848 (4.29)       725.1906 (3.93)       668.0379 (4.20)
test_post_site_performance                             793.1918 (5.20)     1,086.6737 (5.89)       841.8204 (5.30)
test_renditions_performance[dtas.odt-doclib]           899.3008 (5.90)     1,104.6530 (5.99)       925.0898 (5.82)
test_renditions_performance[dtas.odt-imgpreview]       899.8614 (5.90)     1,068.4193 (5.79)       925.2980 (5.82)
test_renditions_performance[dtas.xls-doclib]           900.5196 (5.91)     1,202.0037 (6.52)       936.8543 (5.90)
test_renditions_performance[dtas.xlsx-doclib]          900.7869 (5.91)       976.1583 (5.29)       922.7755 (5.81)
test_renditions_performance[dtas.doc-doclib]           901.1418 (5.91)       949.5569 (5.15)       920.3196 (5.79)
test_renditions_performance[dtas.docx-imgpreview]      901.3233 (5.91)     1,183.2329 (6.42)       925.4249 (5.82)
test_renditions_performance[dtas.docx-doclib]          901.5965 (5.91)     1,031.4085 (5.59)       926.1439 (5.83)
test_renditions_performance[dtas.xlsx-imgpreview]      901.7826 (5.92)     1,195.2217 (6.48)       932.4559 (5.87)
test_renditions_performance[dtas.doc-imgpreview]       903.5653 (5.93)     1,229.0708 (6.66)       934.9182 (5.88)
test_renditions_performance[dtas.xls-imgpreview]       903.9499 (5.93)       956.1612 (5.18)       922.9191 (5.81)
test_renditions_performance[dtas.jpeg-doclib]          934.2089 (6.13)     1,079.4840 (5.85)       969.5077 (6.10)
test_create_users_and_add_to_site_performance        1,152.6486 (7.56)     1,344.4858 (7.29)     1,190.3074 (7.49)
```

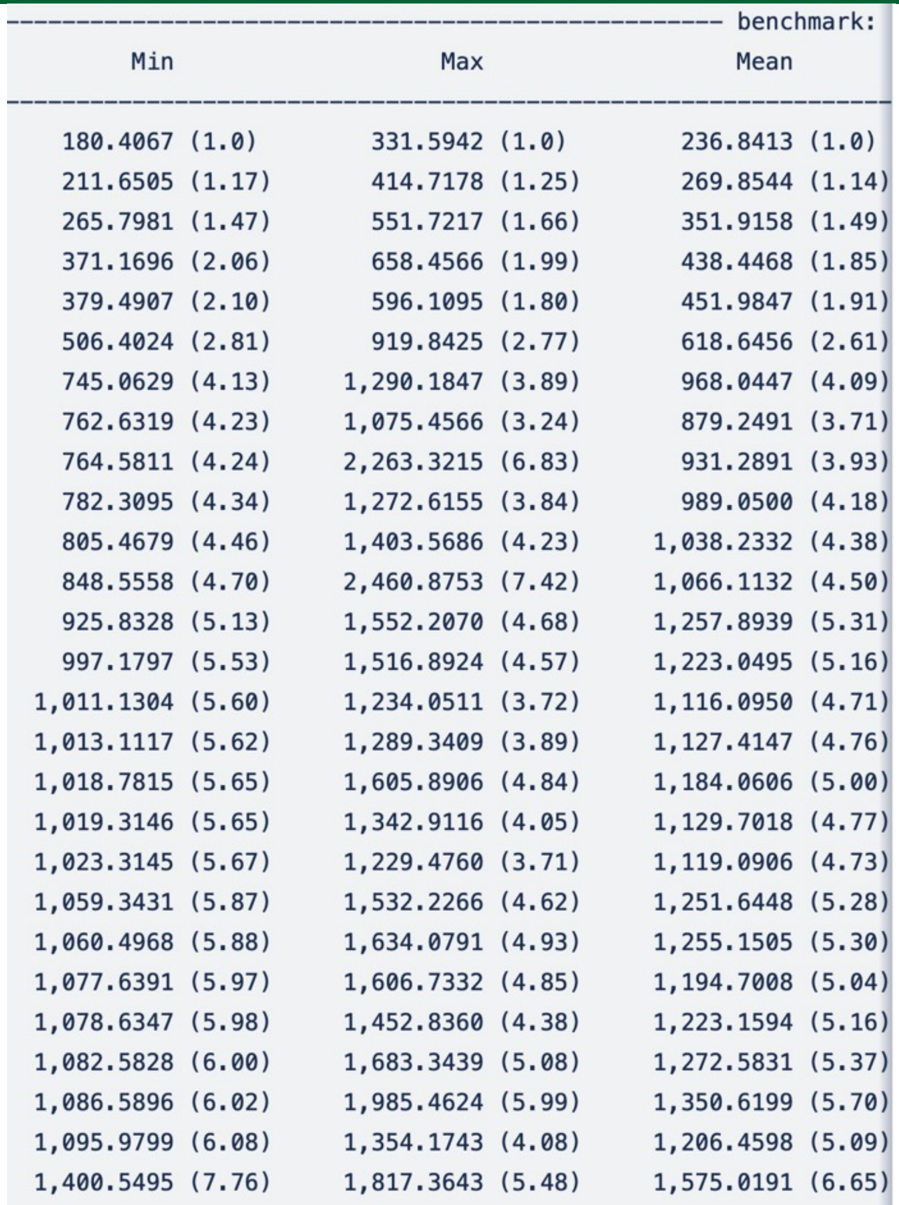

```
----------------------------------------------------- benchmark:
       Min                    Max                   Mean
-----------------------------------------------------------------
   180.4067 (1.0)        331.5942 (1.0)        236.8413 (1.0)
   211.6505 (1.17)       414.7178 (1.25)       269.8544 (1.14)
   265.7981 (1.47)       551.7217 (1.66)       351.9158 (1.49)
   371.1696 (2.06)       658.4566 (1.99)       438.4468 (1.85)
   379.4907 (2.10)       596.1095 (1.80)       451.9847 (1.91)
   506.4024 (2.81)       919.8425 (2.77)       618.6456 (2.61)
   745.0629 (4.13)     1,290.1847 (3.89)       968.0447 (4.09)
   762.6319 (4.23)     1,075.4566 (3.24)       879.2491 (3.71)
   764.5811 (4.24)     2,263.3215 (6.83)       931.2891 (3.93)
   782.3095 (4.34)     1,272.6155 (3.84)       989.0500 (4.18)
   805.4679 (4.46)     1,403.5686 (4.23)     1,038.2332 (4.38)
   848.5558 (4.70)     2,460.8753 (7.42)     1,066.1132 (4.50)
   925.8328 (5.13)     1,552.2070 (4.68)     1,257.8939 (5.31)
   997.1797 (5.53)     1,516.8924 (4.57)     1,223.0495 (5.16)
 1,011.1304 (5.60)     1,234.0511 (3.72)     1,116.0950 (4.71)
 1,013.1117 (5.62)     1,289.3409 (3.89)     1,127.4147 (4.76)
 1,018.7815 (5.65)     1,605.8906 (4.84)     1,184.0606 (5.00)
 1,019.3146 (5.65)     1,342.9116 (4.05)     1,129.7018 (4.77)
 1,023.3145 (5.67)     1,229.4760 (3.71)     1,119.0906 (4.73)
 1,059.3431 (5.87)     1,532.2266 (4.62)     1,251.6448 (5.28)
 1,060.4968 (5.88)     1,634.0791 (4.93)     1,255.1505 (5.30)
 1,077.6391 (5.97)     1,606.7332 (4.85)     1,194.7008 (5.04)
 1,078.6347 (5.98)     1,452.8360 (4.38)     1,223.1594 (5.16)
 1,082.5828 (6.00)     1,683.3439 (5.08)     1,272.5831 (5.37)
 1,086.5896 (6.02)     1,985.4624 (5.99)     1,350.6199 (5.70)
 1,095.9799 (6.08)     1,354.1743 (4.08)     1,206.4598 (5.09)
 1,400.5495 (7.76)     1,817.3643 (5.48)     1,575.0191 (6.65)
```

Key takeaway

- Alfresco can run smoothly on arm64
- EKS cluster can be configured easily to run arm64 nodes
- Performances/latency is slightly better on AWS Graviton2 instances vs cost-near t3a instances
  - Increased cost efficiency
- Migrating to arm64 reduce workloads energy consumption
  - AWS claim up to 60 percent for the same performance of comparable EC2

Next Steps

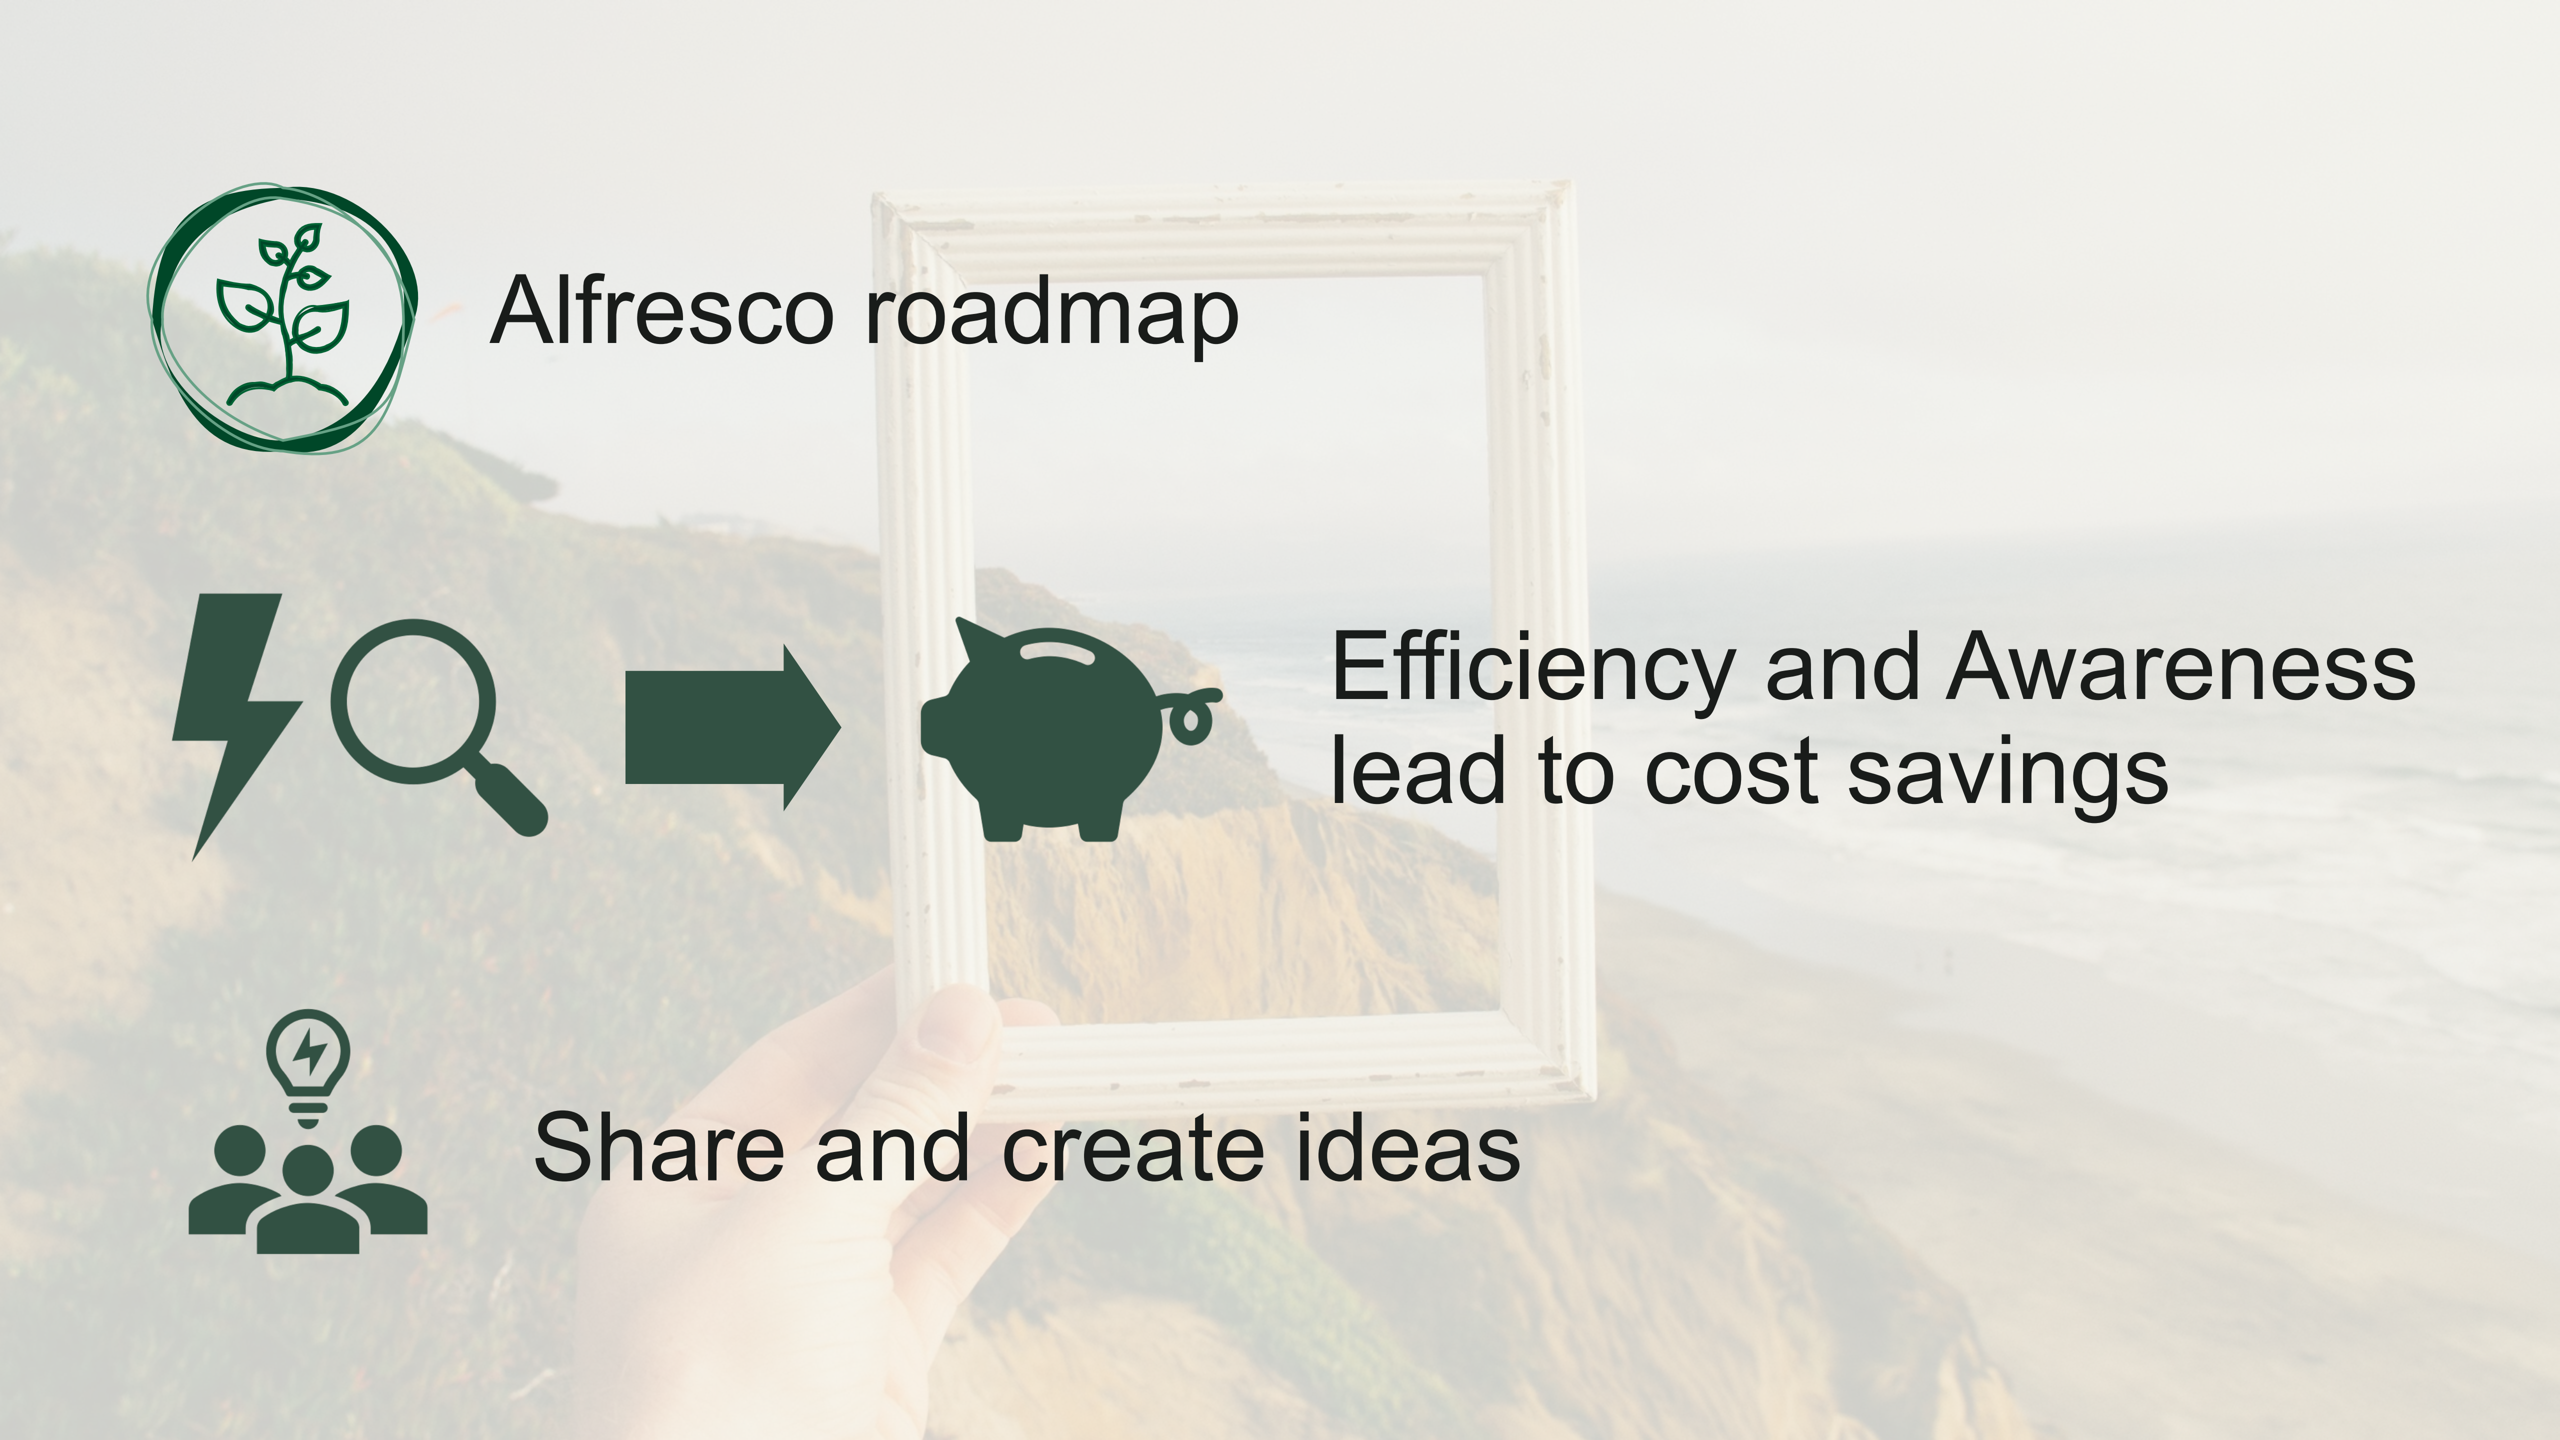

Alfresco roadmap

Efficiency and Awareness
lead to cost savings

Share and create ideas

Thank you

Hyland™

Questions?