
ska-sdp-dataproduct-dashboard

Documentation

Release 0.4.3

Snehal Valame

Aug 08, 2023

CONTENTS:

1	Overview	3
1.1	Usage	3
2	Developer Guide	5
2.1	Development setup	5
2.2	Running scripts	6
3	Deployment Guide	7
3.1	Kubernetes Deployment	7
3.2	Steps to run the system locally in Minikube	8

The *ska-sdp-dataproduct-dashboard* repository contains the code for the dashboard that is used to access Data Products in a storage volume.

OVERVIEW

The data product dashboard is used to list and download data products that are saved on a shared data volume within the environment where it is deployed.

1.1 Usage

The dashboard contains a table of all the data products in the shared data volume, as seen by the [Data Product API](#). This table has built-in functionality for sorting and filtering. When a data product is selected, its metadata is displayed on the right as well as the download options for the selected product.

The screenshot shows the SKAO Data Product Dashboard. At the top, there's a header with the SKAO logo and the title "Data Product Dashboard". Below the header, there are two buttons: "INDEX DATA PRODUCTS" and "RELOAD". The main part of the dashboard is a table with columns: Execution Block ID, Date Created, Command Line, Commit, Image, Processing Block, Processing Script, and Version. The table contains several rows of data products. To the right of the table, there's a sidebar. The top of the sidebar shows the "Selected File" as "eb-m001-20191031-12345" and a "DOWNLOAD" button. Below this, the "Meta Data" section displays various fields: interface, execution_block, context, config, files, date_created, dataproduct_file, and metadata_file.

Execution Block ID	Date Created	Command Line	Commit	Image	Processing Block	Processing Script	Version
eb-test-20200325-00001	2020-03-25			artifact.skao.int/ska-sdp-s...	pb-test-20200425-00000	vis-receive	0.6.0
eb-test-20230401-12345	2023-04-01	-dump /product/eb-test-20...	516fb5a693f9dc9aff5d461...	registry.gitlab.com/ska-tele...	pb-test-20230401-12345	vis-receive	0.8.0
eb-m001-20191031-12345	2019-10-31	-dump /product/eb-m001-2...	516fb5a693f9dc9aff5d461...	artifact.skao.int/ska-dock...	pb-m001-20191031-12345	receive	0.1.3
eb-m003-20191131-12345	2023-06-29	-dump /product/eb-m003-2...	516fb5a693f9dc9aff5d461...	artifact.skao.int/ska-dock...	pb-m003-20191131-12345	receive	0.1.3
eb-test-20230621-12345	2023-06-21	-dump /product/eb-test-20...	516fb5a693f9dc9aff5d461...	registry.gitlab.com/ska-tele...	pb-test-20230401-12345	vis-receive	0.8.0
eb-m002-20221212-12345	2022-12-12	-dump /product/eb-m002-2...	516fb5a693f9dc9aff5d461...	artifact.skao.int/ska-dock...	pb-m002-20221212-12345	receive	0.1.4
eb-test-20230214-07904	2023-02-14			registry.gitlab.com/ska-tele...	pb-test-20230214-05876	vis-receive	0.8.0
eb-m004-20191031-12345	2019-10-31	-dump /product/eb-m004-2...	516fb5a693f9dc9aff5d461...	artifact.skao.int/ska-dock...	pb-m004-20191031-12345	receive	0.1.3

1 row selected

1 - 8 of 8

Selected File: eb-m001-20191031-12345

DOWNLOAD

Meta Data

interface
http://schema.skao.int/ska-data-product-meta/0.1

execution_block
eb-m001-20191031-12345

context
{ "observer": "AIV_person_1", "intent": "Experimental run as part of XYZ-123", "notes": "Running that signal from XX/YY/ZZ through again, things seem a bit flaky" }

config
{ "processing_block": "pb-m001-20191031-12345", "processing_script": "receive", "image": "artifact.skao.int/ska-docker/vis_receive", "version": "0.1.3", "commit": "516fb5a693f9dc9aff5c", "dump": "/product/eb-m001-20191031-12345/ska-sdp/pb-m001-20191031-12345/vis.ms" }

files
[[{"path": "vis.ms", "status": "working", "description": "Raw visibility dump from receive"}]]

date_created
2019-10-31

dataproduct_file
product/eb-m001-20221212-12345

metadata_file
product/eb-m001-20221212-12345/ska-data-product.yaml

Fig. 1: Example Data Product Dashboard

Above the table is a panel which contains functions related to the data store in use. These can include the following:

- **Index Data Products**

This button will re-index the shared data volume from disk into the metadata store. This allows the user to find items that might have been added to the volume without being ingested by the API.

- **Reload**

This function will update the table on the dashboard with the latest list of data products in the metadata store. This allows the user to find items that have been ingested by the API since the dashboard last loaded.

When the data product API has access to an Elasticsearch backend, additional search functionality will become available. This allows the user to use Elasticsearch to search for a key value pair within the metadata.

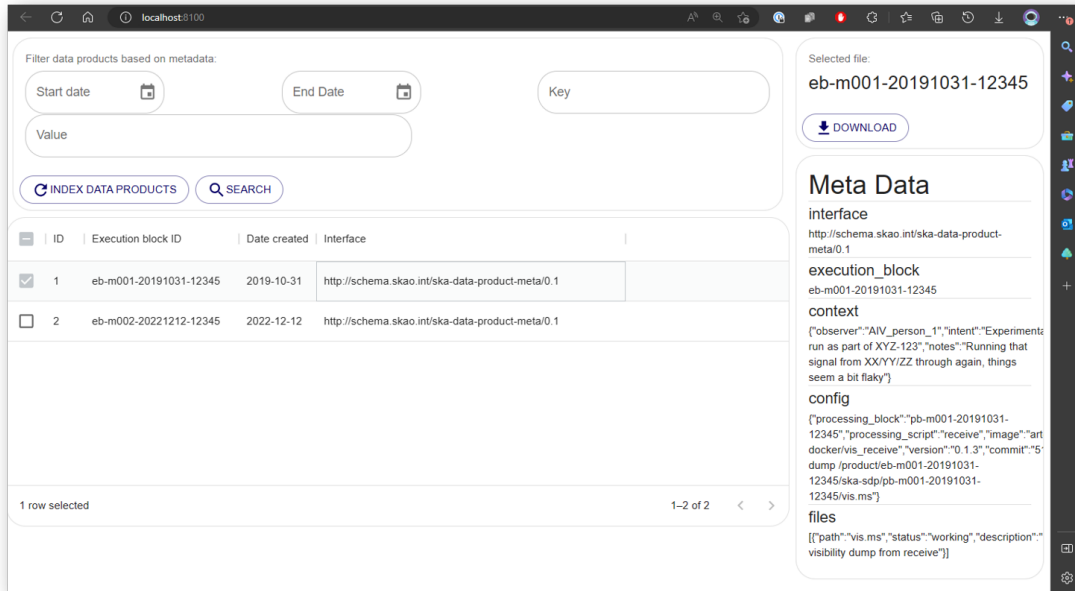


Fig. 2: Example Data Product Dashboard with Elasticsearch enabled.

The application can be run as a standalone front-end application or used as a remote (Webpack 5 Module) within the SKA Portal [SKA Landing Page](#).

DEVELOPER GUIDE

This document complements the guidelines set out in the [SKA telescope developer portal](#)

Tooling Pre-requisites

This project requires **Node** and **YARN** to install and run. To install please follow the instructions for your operating system at [nodejs downloads](#).

Alternatively, the official Node docker image can be used. Instructions can be found on the [official Node docker image site](#).

2.1 Development setup

To run the application directly on your host machine

All the following notes assume you are at the command prompt for your chosen environment.

Confirm Node and YARN are installed and configured correctly, both the following commands should return the relevant version number.

```
> node --version  
> yarn --version
```

Clone the repository and its submodules:

```
git clone git@gitlab.com:ska-telescope/sdp/ska-sdp-dataproduct-dashboard.git  
git submodule update --init --recursive
```

Scripts for running, testing, and building the application are defined in the scripts section of the package.json file. These are run using YARN

To run the application locally on your host machine, install all the necessary project dependencies with the YARN package manager:

```
> yarn install
```

2.2 Running scripts

You should now be able to run the scripts defined in the `package.json` within the project directory.

Running the application in development mode

The app can be run with the node environment set to `NODE_ENV=test` (allowing Istanbul to instrument the code) and webpack serve with `--mode development`. You can access your application at <http://localhost:8100>. The app will recompile and restart if you make any edits to the source files.

```
> yarn start
```

Running the application tests using Cypress

Cypress has been set up to provide component and end to end testing. For information on the use of Cypress, see [Cypress component-testing](#).

Code coverage is implemented with [Istanbul](#) and [NYC](#) for instrumenting the code, and [cobertura reporter](#) as it is used for reporting for the Gitlab CI of coverage statistics.

Cypress can be opened in a browser by running:

```
> yarn cypress:open
```

Or alternatively unit and end to end tests can be run headless by:

```
> yarn test:component:headless
> yarn test:e2e:headless
```

Code coverage can be viewed by opening the `build/coverage/index.html` in a browser after running:

```
> yarn test:coverage:report
```

Running the production code

The build script builds the app for production to the `dist` folder. The build is minified and any JSX is transpiled to JavaScript. Your app is ready to be deployed!

```
> yarn build
```

Running the application inside a docker container on your host machine

To run the application using docker, build the docker file in the root directory and run the container exposing port 8100.

```
docker build -t ska-sdp-dataproduct-dashboard .
docker run -p 8100:8100 ska-sdp-dataproduct-dashboard
```

The project will then be accessible at the URL <http://localhost:8100/>

DEPLOYMENT GUIDE

3.1 Kubernetes Deployment

This is the documentation for the Data Product Dashboard Helm Chart.

3.1.1 Usage

The data product dashboard is intended to be deployed as a standalone deployment, running as a service accessible to other deployments through its API or to users through the dashboard URL. Typical deployments are done from within the GitLab pipelines, deploying into pre-configured environments to one of three namespaces (ci-dev, integration or staging)

During development, developers can deploy the development branches into the ci-dev namespace from the Gitlab pipeline:

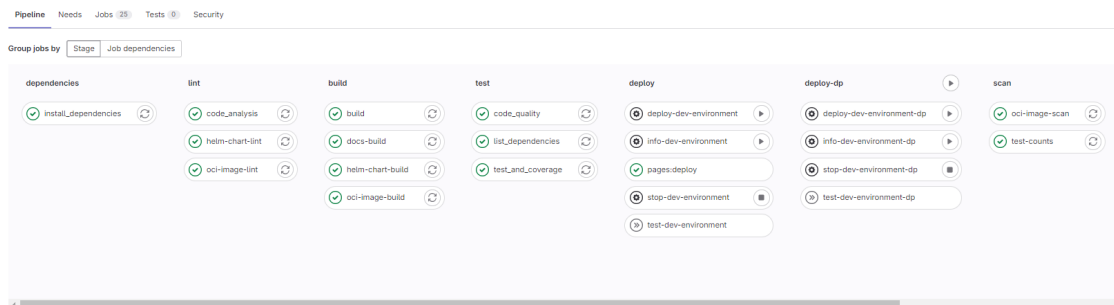


Fig. 1: Deployment from pipeline on dev branch

From the master branch, the application can be deployed into the integration or staging namespace of each environment.

The deployed Data Product Dashboard should then be accessible at: “https://sdhp.stfc.skao.int/\protect\T1\textdollarKUBE_NAMESPACE/dashboard/”, and the backend should be accessible at: “https://sdhp.stfc.skao.int/\protect\T1\textdollarKUBE_NAMESPACE/api/”

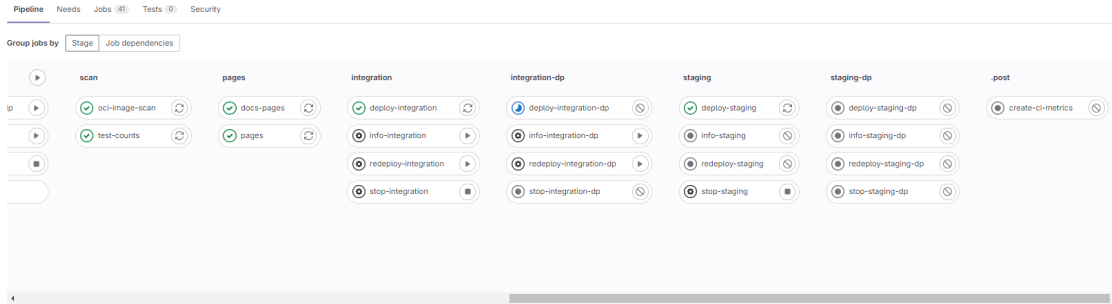


Fig. 2: Deployment from pipeline on master branch

3.2 Steps to run the system locally in Minikube

The following steps will assume that you have the repo checked out, or have the chart locally.

1. Start Minikube if it is not already running:

```
minikube start
minikube status
```

If needed, build images, tag and load them to Minikube.

```
docker build -t ska-sdp-dataproduct-dashboard .
docker images
docker tag [Image ID] ska-sdp-dataproduct-dashboard:[Tag]
minikube image load ska-sdp-dataproduct-dashboard:[Tag]
minikube image ls
```

2. Change to the chart directory in the repository: `cd charts/ska-sdp-dataproduct-dashboard/`. Make the needed changes to image versions and enable the deployments as required in the values files. Then update the Helm dependencies.

```
helm dependency update .
helm dependency build
```

3. Create a new namespace (optional): `kubectl create namespace [namespace]`

4. Install the helm chart with the following values:

```
helm install [namespace] . -n [namespace] --set helmdeploy.namespace=[namespace] --values values_local_deployment.yaml
```

On a system with limited resources / slow connection, run with the following additional flags:

```
helm install [namespace] . -n [namespace] --set helmdeploy.namespace=[namespace] --values values_local_deployment.yaml --set diagnosticMode.enabled=true --timeout=60m
```

Once the above is complete you will have the following running:

- The Data Product API
- The Data Product Dashboard

To be able to access the API and the dashboard run the following:

```
kubectl -n [namespace] port-forward service/ska-sdp-dataproduct-api 8000:8000
kubectl -n [namespace] port-forward service/ska-sdp-dataproduct-dashboard 8100:8100
```

You should now be able to access the API and the Dashboard on the following URL's:

- <http://localhost:8000/filelist>
- <http://localhost:8100/>

To get data onto the PV:

```
kubectl get pod -n [namespace]
kubectl cp [host path]/ska-sdp-dataproduct-api/tests/test_files/product [ska-sdp-
↳ dataproduct-api pod]:/usr/data -n [namespace]
```