

API Pricing Recalculate

Dataiku API Service: pricing

Created in API designer menu in dataiku project.

Endpoint: Recalculate

Background

The Pricing backend initially finds similarities between CPCs and from the 10 closest or so, compute a recommended price for the target CPC (median + business rules, overlays, etc.)

However some of the initial neighbors CPCs might not be comparable from business point of view and they remove them. The new recommended price must be recomputed, that is where the API is used.

Goal

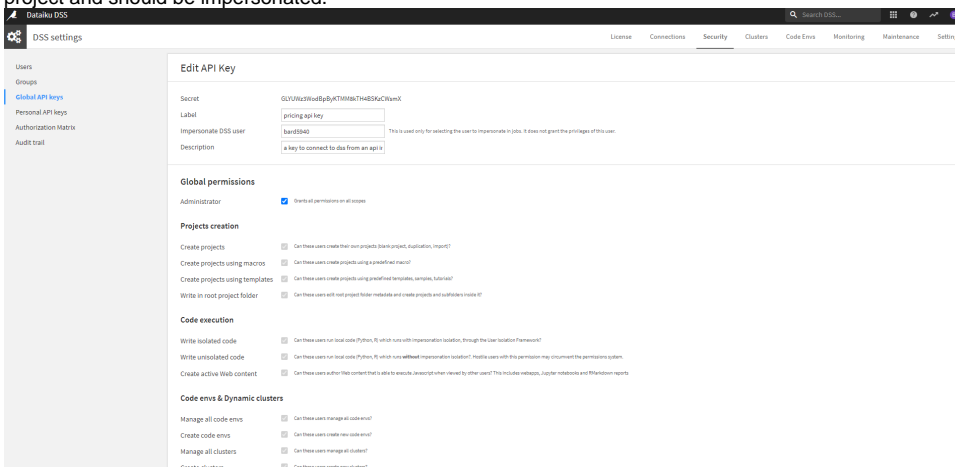
Recalculate the recommended price given the target and a list of neighbor CPCs.

Requirements

- Code shared between API and project, not duplicated
- Access to project's intermediate datasets to re compute only what is necessary
- Dataiku project API Webapp connections in different environment should match

Solution

- Code shared between project initial computation and API is in functions in project's library, not directly in recipes
- Datasets: a mix of PostgreSQL tables and managed folders in project, so we choose to connect to project's and access datasets with dataiku api (alternative would have been direct connection to postgresql if only postgresql tables for example)
- API connects to remote DSS instance using a "Global API key" that should be set at instance level. This key must have appropriate rights for project and should be impersonated.



- Remote DSS is set depending on API host

```
18 # host configuration
19 # hostname is something like: dku-mad-pricing-on-dss-test-standard-1-855fdd5784-kqvxc
20 hostname = socket.gethostname()
21 + if hostname == "ce-bda-dss-design-ew1-prod": # PROD DESIGN
22     dataiku.set_remote_dss("http://ce-bda-dss-design-ew1-prod:10000", "GLYUwz3wodBpBykTM8kTH4BSKzCWsmX")
23     os.environ["DKU_CURRENT_PROJECT_KEY"] = "PBE_MASTER"
24     version = get_config_version()
25 + elif "dss-prod-standard" in hostname: # PROD AUTOMATION
26     dataiku.set_remote_dss("http://ce-bda-dss-automation-ew1-prod:10000", "3rMX9t90tkocF1opB4co5jMrpQRhwsN")
27     os.environ["DKU_CURRENT_PROJECT_KEY"] = "PBE_MASTER"
28     version = get_config_version()
29 + elif hostname == "ce-bda-dss-design-ew1-test": # TEST DESIGN
30     dataiku.set_remote_dss("http://ce-bda-dss-design-ew1-test:10000", "ytYjYL a7j9Nr h1Q66yeKkT9DjGcZqo7")
31     os.environ["DKU_CURRENT_PROJECT_KEY"] = "PBE_MASTER"
32     version = get_config_version()
33 + elif "dss-test-standard" in hostname: # TEST AUTOMATION
34     dataiku.set_remote_dss("http://ce-bda-dss-automation-ew1-test:10000", "1zFckw5vc189j4qN19UPdKxhs1S0d0QP")
35     os.environ["DKU_CURRENT_PROJECT_KEY"] = "PBE_MASTER"
36     version = get_config_version()
37 + else:
38     class ConfigError(Exception):
39         pass
40     raise ConfigError(f"Hostname '{hostname}' is not mapped to a DSS config.")
```

Additional settings

- **Security:** The API service is not public but an API key is created for webapp that should use it when calling the API. If other components need to use the API they should use another API key.

Deployment

Actually done by data ops team.

Since API host (hostname) is important, make sure deployment follows needed naming convention (for prod, test env).

Additionally we deploy on a Kubernetes cluster with a load balancer, some settings are needed.

Steps:

1. In API Designer, Create a Version and Publish it on deployer
2. In Local Deployer, API Services, find the correct service and version, and deploy it: new deployment or update existing deployment.
3. Before the actual deployment, check the settings

The image shows two screenshots of the Local Deployer settings interface for a deployment named 'pricing-on-dss-prod-standard'. The top screenshot shows the 'Replication settings' section, which includes options to override infrastructure settings, set the replica number (currently 1), and enable the Kubernetes Horizontal Pod Autoscaler. The bottom screenshot shows the 'Service exposition' section, which includes options to override infrastructure settings, set the exposition mode (currently Load balancer), port (currently 80), and service annotations (currently networking.gke.io/load-bali: Internal). Both screenshots show a sidebar with navigation options like General settings, Authentication, Endpoints tuning, Audit logging, Pod setup, Sizing and scaling, and Service exposition.