
GalaxyCloudRunner Documentation

Release 0.3.0

Galaxy and GVL projects

May 24, 2019

Contents:

1	How it works	3
2	Getting Started	5
2.1	Configuring Galaxy	5
2.2	Obtaining a CloudLaunch API key	8
2.3	Adding new Pulsar nodes	11
2.4	Job configuration for Galaxy 19.01 or higher	14
2.5	Job configuration for Galaxy versions lower than 19.01	17
2.6	Additional Configuration and Limitations	20
3	Indices and tables	23

GalaxyCloudRunner enables bursting of user jobs to remote compute resources for the [Galaxy application](#). It provides several dynamic job rules that can be plugged into Galaxy, enabling Galaxy to submit jobs to remote compute nodes.

CHAPTER 1

How it works

The GalaxyCloudRunner provides a library of rules that can be plugged into Galaxy through its configuration via `job_conf.xml`. Once configured, Galaxy jobs can be automatically routed to a Galaxy remote job runner, called [Pulsar](#), on nodes running on the cloud. Adding a new node is a simple matter of visiting the [CloudLaunch](#) site and launching a new worker node on your desired cloud. The GalaxyCloudRunner will discover what Pulsar nodes are available by querying the [CloudLaunch](#) API.

CHAPTER 2

Getting Started

Getting started with the GalaxyCloudRunner is a simple process.

1. First, install `galaxycloudrunner` library into your Galaxy's virtual environment
2. Configure Galaxy to use GalaxyCloudRunner job destination rules
3. Launch as many worker nodes as you need through [CloudLaunch](#)
4. Submit jobs as usual

2.1 Configuring Galaxy

2.1.1 Configuring Galaxy 19.01 or higher

1. First install the GalaxyCloudRunner into your Galaxy virtual environment.

```
cd <galaxy_home>
source .venv/bin/activate
pip install galaxycloudrunner
```

2. Edit your `job_conf.xml` in the `<galaxy_home>/config` folder and add the highlighted sections to it.

You will need to add your own value for the `cloudlaunch_api_token` to the file. Instructions on how to obtain your CloudLaunch API key are given below.

```
1 <?xml version="1.0"?>
2 <job_conf>
3   <plugins>
4     <plugin id="local" type="runner" load="galaxy.jobs.runners.
↳ local:LocalJobRunner" workers="4"/>
5     <plugin id="pulsar" type="runner" load="galaxy.jobs.runners.
↳ pulsar:PulsarRESTJobRunner"/>
6   </plugins>
```

(continues on next page)

(continued from previous page)

```

7  <destinations default="galaxycloudrunner">
8      <destination id="local" runner="local"/>
9      <destination id="galaxycloudrunner" runner="dynamic">
10         <param id="type">python</param>
11         <param id="function">cloudlaunch_pulsar_burst</param>
12         <param id="rules_module">galaxycloudrunner.rules</param>
13         <param id="cloudlaunch_api_endpoint">https://launch.usegalaxy.org/
14         ↪ cloudlaunch/api/v1</param>
15         ↪ <!-- Obtain your CloudLaunch token by visiting: https://launch.usegalaxy.
16         ↪ org/profile -->
17         <param id="cloudlaunch_api_token">37c46c89bcbea797bc7cd76fee10932d2c6a2389
18         ↪ </param>
19         ↪ <!-- id of the PulsarRESTJobRunner plugin. Defaults to "pulsar" -->
20         <param id="pulsar_runner_id">pulsar</param>
21         ↪ <!-- Destination to fallback to if no nodes are available -->
22         <param id="fallback_destination_id">local</param>
23         ↪ <!-- Pick next available server and resubmit if an unknown error occurs --
24         ↪ >
25         <resubmit condition="unknown_error and attempt &lt;= 3" destination=
26         ↪ "galaxycloudrunner" />
27     </destination>
</destinations>
<tools>
    <tool id="upload1" destination="local"/>
</tools>
</job_conf>

```

3. Launch as many Pulsar nodes as you need through [CloudLaunch](#). The job rule will periodically query CloudLaunch, discover these new nodes, and route jobs to them. Instructions on how to launch new Pulsar nodes are below.
4. Submit your jobs as usual.

2.1.2 Configuring Galaxy versions lower than 19.01

1. First install the GalaxyCloudRunner into your Galaxy virtual environment.

```

cd <galaxy_home>
source .venv/bin/activate
pip install galaxycloudrunner

```

2. For prior prior to Galaxy 19.01, you will need to add a GalaxyCloudRunner job rule to your Galaxy configuration by pasting the following file contents into your Galaxy job rules folder in: `<galaxy_home>/lib/galaxy/jobs/rules/`.

Create a file named `galaxycloudrunner.py` and paste the following contents into the file at the location above.

```

1  from galaxycloudrunner.runners.cl_pulsar_burst import get_destination
2
3
4  def cloudlaunch_pulsar_burst(app, referrer,
5                               cloudlaunch_api_endpoint=None,
6                               cloudlaunch_api_token=None,
7                               pulsar_runner_id="pulsar",
8                               pulsar_file_action_config=None,
9                               fallback_destination_id=None):

```

(continues on next page)

(continued from previous page)

```

10     return get_destination(app, referrer,
11                           cloudlaunch_api_endpoint,
12                           cloudlaunch_api_token,
13                           pulsar_runner_id,
14                           pulsar_file_action_config,
15                           fallback_destination_id)

```

3. Edit your `job_conf.xml` in the `<galaxy_home>/config` folder and add the highlighted sections to it.

You will need to add your own `cloudlaunch_api_token` to the file. Instructions on how to obtain your CloudLaunch API key are given below. If you have a Galaxy version prior to 19.01, the line `<param id="rules_module">galaxycloudrunner.rules</param>` passed to your destination will not work. This is the reason that we need to perform step 2.

```

1  <?xml version="1.0"?>
2  <job_conf>
3    <plugins>
4      <plugin id="local" type="runner" load="galaxy.jobs.runners.
↳ local:LocalJobRunner" workers="4"/>
5      <plugin id="pulsar" type="runner" load="galaxy.jobs.runners.
↳ pulsar:PulsarRESTJobRunner"/>
6    </plugins>
7    <destinations default="galaxycloudrunner">
8      <destination id="local" runner="local"/>
9      <destination id="galaxycloudrunner" runner="dynamic">
10        <param id="type">python</param>
11        <param id="function">cloudlaunch_pulsar_burst_compat</param>
12        <param id="cloudlaunch_api_endpoint">https://launch.usegalaxy.org/
↳ cloudlaunch/api/v1</param>
13        <!-- Obtain your CloudLaunch token by visiting: https://launch.usegalaxy.
↳ org/profile -->
14        <param id="cloudlaunch_api_token">37c46c89bcbea797bc7cd76fee10932d2c6a2389
↳ </param>
15        <!-- id of the PulsarRESTJobRunner plugin. Defaults to "pulsar" -->
16        <param id="pulsar_runner_id">pulsar</param>
17        <!-- Destination to fallback to if no nodes are available -->
18        <param id="pulsar_fallback_destination_id">local</param>
19      </destination>
20    </destinations>
21    <tools>
22      <tool id="upload1" destination="local"/>
23    </tools>
24  </job_conf>

```

4. Launch as many Pulsar nodes as you need through [CloudLaunch](#). The job rule will periodically query CloudLaunch, discover these new nodes, and route jobs to them. Instructions on how to launch new Pulsar nodes are below.
5. Submit your jobs as usual.

2.1.3 Reducing data transfers

If you would like to control the data transfer configurations for Pulsar, an additional option can be specified in the `job_conf` destination for the GalaxyCloudRunner rule. This is particularly useful for Galaxy's reference data because the remote Pulsar nodes have been configured to mount the Galaxy public file system repository with pre-formatted reference data for a number of tools. In turn, this speeds up job execution and reduces data transfers from

your Galaxy instance because the relevant files do not need to be transferred to the remote node with each job.

Note that this configuration is necessary only if your file system paths differ from those on the remote Pulsar nodes. Specifically for the reference data, Pulsar nodes mount Galaxy Project's CVMFS repository, which is available under `/cvmfs/data.galaxyproject.org/` directory. The layout of that directory can be inspected here: <https://gist.github.com/afgane/b527eb857244f43a680c9654b30deb1f>

To enable this feature for the GalaxyCloudRunner, it is necessary to add the following param to the existing job destination in `job_conf.xml`:


```
<!-- Path for the Pulsar destination config file for path rewrites. -->
<param id="pulsar_file_action_config">config/pulsar_actions.yml</param>
```

In addition, *transfer actions* need to be defined that specify how paths should be translated between the systems. This is done in a dedicated file pointed to in the above param tag, in above example `config/pulsar_actions.yml`. A basic example of the file is available below while complete details about the available transfer action options are available as part of the [Pulsar documentation](#).


```
paths:
- path: /galaxyx/server/tool-data/sacCer2/bwa_mem_index/sacCer2/
  path_types: unstructured
  action: rewrite
  source_directory: /galaxy/server/sacCer2/bwa_mem_index/sacCer2/
  destination_directory: /cvmfs/data.galaxyproject.org/managed/bwa_mem_index/
  ↪ sacCer2/
```

2.2 Obtaining a CloudLaunch API key


1. Visit the CloudLaunch site: <https://launch.usegalaxy.org/>
2. Select *Login* on the top menu bar and sign in through a 3rd party provider.
3. Once logged in, select the 'My Profile' option from the menu bar as shown.



4. Get a new API token for CloudLaunch by expanding the collapsed *API Tokens* panel. You can give the API key any name you like (we have given `galaxycloudrunner`) and click the *Add New Token* button.




5. Copy the token value and paste it into your `job_conf.xml`.



2.3 Adding new Pulsar nodes

1. To launch a new Pulsar node, go to <https://launch.usegalaxy.org/>
2. Select the Galaxy Cloud Bursting appliance.



3. You may be asked to login through a social network provider.
4. Once logged in, fill in the following fields:
 - a. The target cloud you want to launch in
 - b. Provide or choose your credentials for the selected cloud
 - c. Click the 'Test and use these Credentials button' to validate them
 - d. Click next

CloudLaunch v0.2

https://launch.usegalaxy.org/catalog/appliance/pulsar-standalone

CloudLaunch

Q Catalog Public Appliances My Appliances User (enis)

Launching Galaxy Cloud Bursting appliance

Fill out the form below to launch the selected appliance.

1 Select Deployment Target 2 Select Appliance Settings 3 Launch!

Which version of this appliance would you like to launch?

0.1

On which target would you like to launch your appliance?

cloud: Google Cloud Platform, region: us-central1, zone: us-central1-b

What type of credentials do you want to use?

☒ Temporary Credentials ☐ Saved Credentials

You can manually enter the required credentials in the form below or by uploading a text file from your computer.

Load credentials from file

GCP Credentials as JSON

```
{
  "type": "service_account",
  "project_id": "cloudlaunch",
  "private_key_id": "ra16d696da65taa787evcx2d9cc8vxc53aVdxbv",
  "private_key": "-----BEGIN PRIVATE KEY-----
\nMIIEEdgIBADbbNBgqhkiG9w0BAQEFAASCBKggwggSkAgEAAoIBAQD1CmQP/hf1Zmv\nTWi49CXmMoSPy0ZcDWcEOUY9mfSMFmjI2RNVtVgsSu5EwURt
sexXgOPF+PrUsooN5mFZ0d6w\nZnA95b+M9xSLix4Tinn8IWVGeIqnrX85cmcrXUli32H6Mho4/HH5dFYg86RZKabi\nnvNKXGJXJ/J/G2KbRo2qgaanK81EToP
S5KRz/vcWCPkGSahvOmCXGgE5fRfIEhqtH\nnwtLi3jCxAgMBAAECggEAE8XIH7s84y4sHDl2IRAMW9ghNrosHglt2Fsnxhl9d61a\n\nPLS5eKdpkT9lnRlSS56z
B7eMiZ3G3bFbCsdGRSoEE7/Mo+7zQuXzA8uY14W5xu5n\nP6rPMZm5UUIW08HEXzK4jj2kGAm+ua/aE6KK+G9pVZtEoj9u7jwtnvdqS40e2F5\n\nnxM0zfnS
2dzoo238VDUF0Jlt7ccsdcbzXf8kiGP1iCE5b18j8Dn1MjJaEuqOcldgzV2+h\n\nTjfhPlI4N02If/PBk7XYxP/Qwk04SeRtcnvco71bQ9RnZVyywO/fOisdr0dUwXec
u\nnCwGrXBSSyZX5oB3N1xHm/hotHo9ojxGyKpbLWaSo0QKBgQD9AsjUG42/UWDpRFBq\n\nnbd2koTQneeSA0VDMge4x4ZX9uqkv2MWme+gzOKpBdWa/V9L
etHVYw4fCwX7DHomO\n\nndJQqi6H+H9euwEy56GniuVhE87JslhYsvXedsYk4SxDzWG0DcslcjWpsnHRZlp6r\n\nnOnkpvj1BMFUCup6BuJwFRrAIQKBgQD374A
```

Copy your GCP credentials json file into this field, or use the upload button above.

Which username would you like to use as the default VM system user?

ubuntu

When launching a new virtual machine, the keypair you specify at launch will be associated with this user's account and you must use this username to ssh in. You can leave the defaults as is if you're unsure.

TEST AND USE THESE CREDENTIALS

SAVE TO PROFILE

NEXT >

5. Finally, select the size of the Virtual Machine you want, and click Launch.

CloudLaunch v0.2

https://launch.usegalaxy.org/catalog/appliance/pulsar-standalone

CloudLaunch

Q Catalog Public Appliances My Appliances User (enis)

Launching Galaxy Cloud Bursting appliance

Fill out the form below to launch the selected appliance. ☒

1 Select Deployment Target 2 Select Appliance Settings 3 Launch!

Provide a name for your deployment
enis-pulsar-standalone-19-05-21t14-25

A deployment name helps you identify your appliance. The name must be at most 63 characters long and can consist of lowercase letters, numbers, and dashes. It must start with a letter and not end with a dash.

What type of virtual hardware would you like to use?
n1-standard-2

2 VCPUs 7.5 GB RAM 55536 GB Disk

☐ Advanced cloud launch options

[< PREVIOUS](#)

[LAUNCH](#)

ABOUT
About
Contact Us

DEVELOPER
Developer API (beta)
Source Code

f t G+ +

- Simply launching the node is enough, the GalaxyCloudRunner will now pick up your new nodes by querying the CloudLaunch API.

2.4 Job configuration for Galaxy 19.01 or higher

2.4.1 Simple configuration

The following is a simple job configuration sample that you can use to get started.

```

1 <?xml version="1.0"?>
2 <job_conf>
3   <plugins>
4     <plugin id="local" type="runner" load="galaxy.jobs.runners.
↳ local:LocalJobRunner" workers="4"/>
5     <plugin id="pulsar" type="runner" load="galaxy.jobs.runners.
↳ pulsar:PulsarRESTJobRunner"/>
6   </plugins>
7   <destinations default="galaxycloudrunner">
8     <destination id="local" runner="local"/>
9     <destination id="galaxycloudrunner" runner="dynamic">
10       <param id="type">python</param>
11       <param id="function">cloudlaunch_pulsar_burst</param>
12       <param id="rules_module">galaxycloudrunner.rules</param>
13       <param id="cloudlaunch_api_endpoint">https://launch.usegalaxy.org/
↳ cloudlaunch/api/v1</param>

```

(continues on next page)

(continued from previous page)

```

14      <!-- Obtain your CloudLaunch token by visiting: https://launch.usegalaxy.
↳org/profile -->
15      <param id="cloudlaunch_api_token">37c46c89bcbea797bc7cd76fee10932d2c6a2389
↳</param>
16      <!-- id of the PulsarRESTJobRunner plugin. Defaults to "pulsar" -->
17      <param id="pulsar_runner_id">pulsar</param>
18      <!-- Destination to fallback to if no nodes are available -->
19      <param id="fallback_destination_id">local</param>
20      <!-- Pick next available server and resubmit if an unknown error occurs --
↳>
21      <resubmit condition="unknown_error and attempt &lt;= 3" destination=
↳"galaxycloudrunner" />
22      </destination>
23  </destinations>
24  <tools>
25      <tool id="upload1" destination="local"/>
26  </tools>
27 </job_conf>

```

In this simple configuration, all jobs are routed to GalaxyCloudRunner by default. This works as follows:

1. If a Pulsar node is available, it will return that node.
2. If multiple Pulsar nodes are available, they will be returned in a round-robin loop.
3. You can add or remove Pulsar nodes at any time. However, there's a caching period (currently 5 minutes) to avoid repeatedly querying the server, which will result in a short period of time before the change is detected by the GalaxyCloudRunner. This has implications for node addition and in particular removal. When adding a node, there could be a delay of a few minutes before the node is picked up. If a Pulsar node is removed, your jobs may be routed to a dead node for the duration of the caching period. Therefore, we recommend attempting a job resubmission through the resubmit tag as shown in the example. See [Additional Configuration and Limitations](#) on how to change this cache period.
4. If no node is available, it will return the `fallback_destination_id`, if specified, in which case the job will be routed there. If no `fallback_destination_id` is specified, the job will be re-queued till a node becomes available.

2.4.2 To burst or not to burst?

In the above example, all jobs are routed to the GalaxyCloudRunner by default. However, it is often the case that jobs should be routed to the remote cloud nodes only if the local queue is full. To support this scenario, we recommend a configuration like the following.

```

1 <?xml version="1.0"?>
2 <job_conf>
3   <plugins>
4     <plugin id="local" type="runner" load="galaxy.jobs.runners.
↳local:LocalJobRunner" workers="4"/>
5     <plugin id="drmaa" type="runner" load="galaxy.jobs.runners.
↳drmaa:DRMAAJobRunner">
6     <plugin id="pulsar" type="runner" load="galaxy.jobs.runners.
↳pulsar:PulsarRESTJobRunner"/>
7   </plugins>
8   <destinations default="burst_if_queued">
9     <destination id="local" runner="local"/>
10    <destination id="burst_if_queued" runner="dynamic">

```

(continues on next page)

(continued from previous page)

```

11     <param id="type">burst</param>
12     <param id="from_destination_ids">local,drmaa</param>
13     <param id="to_destination_id">galaxycloudrunner</param>
14     <param id="num_jobs">2</param>
15     <param id="job_states">queued</param>
16 </destination>
17 <destination id="galaxycloudrunner" runner="dynamic">
18     <param id="type">python</param>
19     <param id="function">cloudlaunch_pulsar_burst</param>
20     <param id="rules_module">galaxycloudrunner.rules</param>
21     <param id="cloudlaunch_api_endpoint">https://launch.usegalaxy.org/
    ↪ cloudlaunch/api/v1</param>
22     <!-- Obtain your CloudLaunch token by visiting: https://launch.usegalaxy.
    ↪ org/profile -->
23     <param id="cloudlaunch_api_token">37c46c89bcbea797bc7cd76fee10932d2c6a2389
    ↪ </param>
24     <!-- id of the PulsarRESTJobRunner plugin. Defaults to "pulsar" -->
25     <param id="pulsar_runner_id">pulsar</param>
26     <!-- Destination to fallback to if no nodes are available -->
27     <param id="fallback_destination_id">local</param>
28     <!-- Pick next available server and resubmit if an unknown error occurs --
    ↪ >
29     <resubmit condition="unknown_error and attempt <= 3" destination=
    ↪ "galaxycloudrunner" />
30 </destination>
31 </destinations>
32 <tools>
33     <tool id="upload1" destination="local"/>
34 </tools>
35 </job_conf>

```

Note the emphasized lines. In this example, we route to the built-in rule `burst_if_queued` first, which determines whether or not the cloud bursting should occur. It examines how many jobs in the `from_destination_ids` are in the given state (queued in this case), and if there are above `num_jobs`, routes to the `to_destination_id` destination (`galaxycloudrunner` in this case). If bursting should not occur, it routes to the first destination in the `from_destination_ids` list. This provides a simple method to scale to Pulsar nodes only if a desired queue has a backlog of jobs. You may need to experiment with these values to find ones that work best for your requirements.

2.4.3 Advanced bursting

In this final example, we show how a complex chain of rules can be used to exert fine-grained control over the job routing process.

```

1 <?xml version="1.0"?>
2 <job_conf>
3     <plugins>
4         <plugin id="local" type="runner" load="galaxy.jobs.runners.
    ↪ local:LocalJobRunner" workers="4"/>
5         <plugin id="drmaa" type="runner" load="galaxy.jobs.runners.
    ↪ drmaa:DRMAAJobRunner">
6         <plugin id="pulsar" type="runner" load="galaxy.jobs.runners.
    ↪ pulsar:PulsarRESTJobRunner"/>
7     </plugins>
8     <destinations default="burst_if_queued">
9         <destination id="local" runner="local"/>

```

(continues on next page)

(continued from previous page)

```

10     <destination id="burst_if_queued" runner="dynamic">
11         <param id="type">burst</param>
12         <param id="from_destination_ids">local,drmaa</param>
13         <param id="to_destination_id">burst_if_size</param>
14         <param id="num_jobs">2</param>
15         <param id="job_states">queued</param>
16     </destination>
17     <destination id="burst_if_size" runner="dynamic">
18         <param id="type">python</param>
19         <param id="function">to_destination_if_size</param>
20         <param id="rules_module">galaxycloudrunner.rules</param>
21         <param id="max_size">1g</param>
22         <param id="to_destination_id">galaxycloudrunner</param>
23         <param id="fallback_destination_id">local</param>
24     </destination>
25     <destination id="galaxycloudrunner" runner="dynamic">
26         <param id="type">python</param>
27         <param id="function">cloudlaunch_pulsar_burst</param>
28         <param id="rules_module">galaxycloudrunner.rules</param>
29         <param id="cloudlaunch_api_endpoint">https://launch.usegalaxy.org/
30         <!-- Obtain your CloudLaunch token by visiting: https://launch.usegalaxy.
31         org/profile -->
32         <param id="cloudlaunch_api_token">37c46c89bcbea797bc7cd76fee10932d2c6a2389
33         </param>
34         <!-- id of the PulsarRESTJobRunner plugin. Defaults to "pulsar" -->
35         <param id="pulsar_runner_id">pulsar</param>
36         <!-- Destination to fallback to if no nodes are available -->
37         <param id="fallback_destination_id">local</param>
38         <!-- Pick next available server and resubmit if an unknown error occurs --
39         <resubmit condition="unknown_error and attempt <= 3" destination=
40         "galaxycloudrunner" />
41     </destination>
42 </destinations>
43 <tools>
44     <tool id="upload1" destination="local"/>
45 </tools>
46 </job_conf>

```

Jobs are first routed to the built-in `burst_if_queued` rule, which determines whether the bursting should occur. If it should, it is then routed to the `burst_if_size` destination, which will check the total size of the input files. If they are less than 1GB, they are routed to the `galaxycloudrunner` destination. If not, they are routed to a local queue.

2.5 Job configuration for Galaxy versions lower than 19.01

2.5.1 Simple configuration

The following is a simple job configuration sample that you can use to get started.

```

1 <?xml version="1.0"?>
2 <job_conf>
3     <plugins>

```

(continues on next page)

(continued from previous page)

```

4     <plugin id="local" type="runner" load="galaxy.jobs.runners.
↪local:LocalJobRunner" workers="4"/>
5     <plugin id="pulsar" type="runner" load="galaxy.jobs.runners.
↪pulsar:PulsarRESTJobRunner"/>
6     </plugins>
7     <destinations default="galaxycloudrunner">
8         <destination id="local" runner="local"/>
9         <destination id="galaxycloudrunner" runner="dynamic">
10             <param id="type">python</param>
11             <param id="function">cloudlaunch_pulsar_burst_compat</param>
12             <param id="cloudlaunch_api_endpoint">https://launch.usegalaxy.org/
↪cloudlaunch/api/v1</param>
13             <!-- Obtain your CloudLaunch token by visiting: https://launch.usegalaxy.
↪org/profile -->
14             <param id="cloudlaunch_api_token">37c46c89bcbea797bc7cd76fee10932d2c6a2389
↪</param>
15             <!-- id of the PulsarRESTJobRunner plugin. Defaults to "pulsar" -->
16             <param id="pulsar_runner_id">pulsar</param>
17             <!-- Destination to fallback to if no nodes are available -->
18             <param id="pulsar_fallback_destination_id">local</param>
19         </destination>
20     </destinations>
21     <tools>
22         <tool id="upload1" destination="local"/>
23     </tools>
24 </job_conf>

```

In this simple configuration, all jobs are routed to GalaxyCloudRunner by default. This works as follows:

1. If a Pulsar node is available, it will return that node.
2. If multiple Pulsar nodes are available, they will be returned in a round-robin loop.
3. You can add or remove Pulsar nodes at any time. However, there's a caching period (currently 5 minutes) to avoid repeatedly querying the server, that will result in a short period of time before the change is detected by the GalaxyCloudRunner. This has implications for node addition and in particular removal. When adding a node, there could be a delay of a few minutes before the node is picked up. If a Pulsar node is removed, your jobs may be routed to a dead node for the duration of the caching period. Therefore, we recommend a job resubmission through a resubmit tag. However, Galaxy versions prior to 19.01 do not support resubmissions for Pulsar, and you may need to change the cache period to zero to handle this scenario. See [Additional Configuration and Limitations](#) on how to change this cache period.
4. If no node is available, it will return the `fallback_destination_id`, if specified, in which case the job will be routed there. If no `fallback_destination_id` is specified, the job will be re-queued till a node becomes available.

Note that you must manually add the galaxy rule as described here: [Configuring Galaxy versions lower than 19.01](#)

2.5.2 To burst or not to burst?

In the above example, all jobs are routed to the GalaxyCloudRunner by default. However, it is often the case that jobs should be routed to the remote cloud nodes only if the local queue is full. To support this scenario, we recommend a configuration like the following.

```

1 <?xml version="1.0"?>
2 <job_conf>

```

(continues on next page)

(continued from previous page)

```

3     <plugins>
4         <plugin id="local" type="runner" load="galaxy.jobs.runners.
↳ local:LocalJobRunner" workers="4"/>
5         <plugin id="drmaa" type="runner" load="galaxy.jobs.runners.
↳ drmaa:DRMAAJobRunner">
6         <plugin id="pulsar" type="runner" load="galaxy.jobs.runners.
↳ pulsar:PulsarRESTJobRunner"/>
7     </plugins>
8     <destinations default="galaxycloudrunner">
9         <destination id="local" runner="local"/>
10        <destination id="galaxycloudrunner" runner="dynamic">
11            <param id="type">python</param>
12            <param id="function">cloudlaunch_pulsar_burst_compat</param>
13            <param id="cloudlaunch_api_endpoint">https://launch.usegalaxy.org/
↳ cloudlaunch/api/v1</param>
14            <!-- Obtain your CloudLaunch token by visiting: https://launch.usegalaxy.
↳ org/profile -->
15            <param id="cloudlaunch_api_token">37c46c89bcbea797bc7cd76fee10932d2c6a2389
↳ </param>
16            <!-- id of the PulsarRESTJobRunner plugin. Defaults to "pulsar" -->
17            <param id="pulsar_runner_id">pulsar</param>
18            <!-- Destination to fallback to if no nodes are available -->
19            <param id="pulsar_fallback_destination_id">local</param>
20            <param id="burst_enabled">true</param>
21            <param id="burst_from_destination_ids">local,drmaa</param>
22            <param id="burst_num_jobs">2</param>
23            <param id="burst_job_states">queued</param>
24        </destination>
25    </destinations>
26    <tools>
27        <tool id="upload1" destination="local"/>
28    </tools>
29 </job_conf>

```

Galaxy versions prior to 19.01 do not support chaining dynamic rules, and therefore, we have provided a single monolithic rule that can handle both scenarios.

Note the `burst_enabled` flag, which will activate the bursting rule. This rule will determine whether or not the cloud bursting should occur. It examines how many jobs in the `burst_from_destinations` are in the given state (queued in this case), and bursts to pulsar only if they are above `burst_num_jobs`. If bursting should not occur, it routes to the first destination in the `from_destinations` list. This provides a simple method to scale to Pulsar nodes only if a desired queue has a backlog of jobs. You may need to experiment with these values to find ones that work best for your requirements.

2.5.3 Advanced bursting

In this final example, we expand this compound rule to also filter jobs by size.

```

1 <?xml version="1.0"?>
2 <job_conf>
3     <plugins>
4         <plugin id="local" type="runner" load="galaxy.jobs.runners.
↳ local:LocalJobRunner" workers="4"/>
5         <plugin id="drmaa" type="runner" load="galaxy.jobs.runners.
↳ drmaa:DRMAAJobRunner">

```

(continues on next page)

(continued from previous page)

```

6      <plugin id="pulsar" type="runner" load="galaxy.jobs.runners.
↳ pulsar:PulsarRESTJobRunner"/>
7      </plugins>
8      <destinations default="galaxycloudrunner">
9          <destination id="local" runner="local"/>
10         <destination id="galaxycloudrunner" runner="dynamic">
11             <param id="type">python</param>
12             <param id="function">cloudlaunch_pulsar_burst_compat</param>
13             <param id="cloudlaunch_api_endpoint">https://launch.usegalaxy.org/
↳ cloudlaunch/api/v1</param>
14             <!-- Obtain your CloudLaunch token by visiting: https://launch.usegalaxy.
↳ org/profile -->
15             <param id="cloudlaunch_api_token">37c46c89bcbea797bc7cd76fee10932d2c6a2389
↳ </param>
16             <!-- id of the PulsarRESTJobRunner plugin. Defaults to "pulsar" -->
17             <param id="pulsar_runner_id">pulsar</param>
18             <!-- Destination to fallback to if no nodes are available -->
19             <param id="pulsar_fallback_destination_id">local</param>
20             <param id="burst_enabled">true</param>
21             <param id="burst_from_destination_ids">local,drmaa</param>
22             <param id="burst_num_jobs">2</param>
23             <param id="burst_job_states">queued</param>
24             <param id="dest_if_size_enabled">true</param>
25             <param id="dest_if_size_max_size">1g</param>
26             <param id="dest_if_size_fallback_destination_id">local</param>
27         </destination>
28     </destinations>
29     <tools>
30         <tool id="upload1" destination="local"/>
31     </tools>
32 </job_conf>

```

Enable the `dest_if_size_enabled` flag as highlighted to filter by size. This will make sure that the job is routed to Pulsar only if the total size of the input files are less than 1GB. If not, they are routed to `dest_if_size_fallback_destination_id`, which in this case, is a local queue.

2.6 Additional Configuration and Limitations

1. Configuring the query timeout

You can set the environment variable `CLOUDLAUNCH_QUERY_CACHE_PERIOD` before starting Galaxy to control the caching period (in seconds). Setting this to 0 will allow you to get around the node removal issue where, if a Pulsar node is removed, jobs may be routed to a dead node for the duration of the caching period. However, we recommend setting a value greater than 0 to avoid repeatedly querying a remote server during each job submission.

2. Incompatible tools

Due to the nature of how Galaxy collects metadata on datasets, certain tools are not compatible with job execution in the bursting mode. Some of these issues will be resolved once Pulsar is upgraded to collect metadata itself but for the time being the following is an (incomplete) list of tools and tool classes that will not operate when executed via the GalaxyCloudRunner: upload tool, data managers, tools that use metadata input, and tools that use custom data discovery.

3. Auto-scaling

Currently, the GalaxyCloudRunner does not support automatic scaling, you must manually add and remove individual nodes but you can add as many as you would like. We will be adding autoscaling features as part of CloudMan v2.0 in future.

4. Galaxy versions prior to 19.01

Galaxy versions prior to 19.01 do not support certain features required by GalaxyCloudRunner and therefore, need more complex configuration steps.

CHAPTER 3

Indices and tables

- `genindex`
- `modindex`
- `search`