

---

**funcX**

***Release 0.0.1a5***

**May 22, 2020**



---

## Contents:

---

<b>1 Quickstart</b>	<b>3</b>
1.1 Installation . . . . .	3
<b>2 Endpoints</b>	<b>5</b>
2.1 First time setup . . . . .	5
2.2 Configuring funcX . . . . .	5
2.3 Configuring an Endpoint . . . . .	6
2.4 Starting an Endpoint . . . . .	6
2.5 Stopping an Endpoint . . . . .	7
2.6 Listing Endpoints . . . . .	7
<b>3 funcX Tutorial</b>	<b>9</b>
3.1 funcX Client . . . . .	9
3.2 Registering a function . . . . .	10
3.3 Running a function . . . . .	10
<b>4 funcX Client</b>	<b>11</b>
4.1 Registering Functions . . . . .	12
4.2 Running Functions . . . . .	12
4.3 Client Throttling . . . . .	12
4.4 FuncXClient Reference: . . . . .	12
<b>5 Reference</b>	<b>17</b>
5.1 funcx package . . . . .	17
<b>6 Debugging</b>	<b>57</b>
6.1 Setting up the forwarder locally . . . . .	57
<b>7 Indices and tables</b>	<b>59</b>
<b>Python Module Index</b>	<b>61</b>
<b>Index</b>	<b>63</b>



funcX is a High performance function serving system designed to orchestrate scientific workloads across heterogeneous computing resources, from laptops and local servers, through to campus clusters, clouds, and supercomputers.



**funcX** is currently in Alpha and early testing releases are available on [PyPI](#).

The latest version available on PyPI is `v0.0.1a4`.

You can try funcX on [Binder](#)

## 1.1 Installation

**funcX** comes with two components: the **endpoint agent** which is a user-launched services that make computation resources accessible for function executions, and the **funcX client** that enables the registration, execution and tracking of functions across **endpoints**.

Here are some pre-requisites for both the *endpoints* and the *funcX client*

1. Python3.6+
2. The machine must have outbound network access

To check if you have the right Python version, run the following commands:

```
>>> python3 --version
```

This should return the Python version, for eg: `Python 3.6.7`. Please note that that only the first two version numbers need to match.

To check if you have network access, run

```
>>> curl http://dev.funcx.org/api/v1/version
```

This should return a version string, for eg: `"0.0.1"`

### 1.1.1 Installation using Pip

While `pip` and `pip3` can be used to install funcX we suggest the following approach for reliable installation when many Python environments are available.

1. Install funcX:

```
$ python3 -m pip install funcx
```

To update a previously installed funcX to a newer version, use: `python3 -m pip install -U funcx`

---

**Note:** The endpoint and client must use the same Python version. This is due to serialization differences between versions.

---

2. Install Jupyter for Tutorial notebooks:

```
$ python3 -m pip install jupyter
```

---

**Note:** For more detailed info on setting up Jupyter with Python3.5 go [here](#)

---



An endpoint is a persistent service launched by the user on their compute system that serves as a conduit for routing and executing functions to their compute system. This could be their laptop, the login node of a campus cluster, grid, or supercomputing facility.

The endpoint can be configured to connect to the funcX webservice at [funcx.org](https://funcx.org). Once the endpoint is registered you can invoke functions to be executed on it.

### 2.1 First time setup

For general use, running the initialize command via the *funcx-endpoint* CLI tool is the quickest way to setup your installation:

```
$ funcx-endpoint init
```

Initiating funcX will ask you to authenticate with Globus Auth. We require authentication in order to associate endpoints with users and enforce authentication and access control on the endpoint. As part of this step we request access to your identity information (to retrieve your email address) and Globus Groups management. We use Groups information to facilitate sharing of functions and endpoints by checking the Group membership of a group associated with a function.

Once you've run this command, a directory will be created at *\$HOME/funcx* and a set of configuration files will be generated. A default endpoint profile is also created that will be used, whenever you do not explicitly specify a profile to use for endpoint actions.

### 2.2 Configuring funcX

A configuration file will be created at *\$HOME/funcx/config.py*. This contains base information regarding the endpoint, such as a username and email. By default, this includes an address and port for a local broker, which is used if a local broker is deployed. You can also set the address of the endpoint for your workers to connect to. This is necessary if

your workers are not deployed on the same resource as the endpoint (e.g., when using a batch submission system, or cloud workers).

---

**Note:** If your funcX workers are not deployed on the same resource as the endpoint you must set the endpoint address for the workers to find the endpoint. This is done by setting the `endpoint_address`.

---

For example

```
import getpass
from parsl.addresses import address_by_route, address_by_hostname

global_options = {
    'username': getpass.getuser(),
    'email': 'USER@USERDOMAIN.COM',
    'broker_address': '127.0.0.1',
    'broker_port': 8088,
    'endpoint_address': address_by_hostname(),
}
```

## 2.3 Configuring an Endpoint

FuncX endpoints are designed to act as gateways to computational resources such as Clusters, Clouds, Supercomputers, and even your laptop. To make the best use of your resources, the endpoint must be configured to match the resources' capabilities and reflect the needs of the workloads you plan to execute. For example, you may want to limit the number of cores available to your endpoint.

FuncX provides you a rich class based configuration model that allows you to specify the shape of the resources (# of nodes, # of cores per worker, walltime etc) as well as place limits on how funcX may scale the resources in response to changing workload demands.

To generate the appropriate directories and default config template, run the following command:

```
$ funcx-endpoint configure <ENDPOINT_NAME>
```

The above command will create a profile for your endpoint in `$HOME/.funcx/<ENDPOINT_NAME>/` and will instantiate a `config.py` file. This file should be updated with the appropriate configurations for the computational system you are targeting before you start the endpoint. The funcX builds on Parsl and is configured using a `Config` object. For more information, see the `Config` class documentation.

---

**Note:** If the `ENDPOINT_NAME` is not specified, a default endpoint named “default” is configured.

---

## 2.4 Starting an Endpoint

To start a new endpoint run the following command:

```
$ funcx-endpoint start <ENDPOINT_NAME>
```

The above command will create a profile for your endpoint in `$HOME/.funcx/<ENDPOINT_NAME>/config.py`. This file should be updated with the appropriate configurations for the computational system you are targeting before you start the endpoint. To launch the endpoint, simply rerun the above command.

---

**Note:** If the `ENDPOINT_NAME` is not specified, a default endpoint named “default” is started.

---

Starting an endpoint will perform a registration process with the funcX Web Service. The registration process provides funcX with information regarding the endpoint. The Web Service then creates a Forwarder process for the endpoint and returns a UUID and connection information to the Forwarder. The endpoint will use this connection information to connect to the Forwarder. The endpoint establishes three outbound ZeroMQ channels to the forwarder (on the three ports returned during registration) to retrieve tasks, send results, and communicate command information.

Once started, the endpoint uses a daemon process to run in the background.

**Warning:** Only the owner of an endpoint is authorized to start an endpoint. Thus if you register with a different Globus Auth identity and try to start an endpoint owned by another identity, it will fail.

## 2.5 Stopping an Endpoint

To stop an endpoint, run the following command:

```
$ funcx-endpoint stop <ENDPOINT_NAME>
```

---

**Note:** If the `ENDPOINT_NAME` is not specified, the default endpoint is stopped.

---

**Warning:** Run the `funcx-endpoint stop` command **twice** to ensure that the endpoint is shutdown.

## 2.6 Listing Endpoints

To list available endpoints on the current system, run:

```
$ funcx-endpoint list
+-----+-----+-----+
| Endpoint Name | Status | Endpoint ID |
+-----+-----+-----+
| default      | Active | 1e999502-b434-49a2-a2e0-d925383d2dd4 |
+-----+-----+-----+
| KNL_test     | Inactive | 8c01d13c-cfc1-42d9-96d2-52c51784ea16 |
+-----+-----+-----+
| gpu_cluster  | Initialized | None |
+-----+-----+-----+
```

Endpoints can be the following states:

- **Initialized:** This status means that the endpoint has been created, but not started following configuration and not registered with the *funcx service*
- **Active:** This status means that the endpoint is active and available for executing functions
- **Inactive:** This status means that endpoint is not running right now and therefore, cannot service any functions.



funcX is a Function-as-a-Service (FaaS) platform for science that enables you to convert almost any computing resource into a high-performance function serving device. To do this, you deploy a funcX endpoint agent on the resource, which integrates it into the function serving fabric, allowing you to dynamically send, monitor, and receive results from function invocations. funcX is built on top of `Parsl`, enabling a funcX endpoint to use large compute resources via traditional batch queues, where funcX will dynamically provision, use, and release resources on-demand to fulfill function requests. The function service fabric, which is run centrally as a service, is hosted in AWS.

Here we provide an example of using funcX to register a function and run it on a publicly available tutorial endpoint.

### 3.1 funcX Client

We start by instantiating a funcX client as a programmatic means of communicating with the function service fabric. The client allows you to: - Register functions - Register containers and execution environments - Launch registered functions against endpoints - Check the status of launched functions - Retrieve outputs from functions

#### 3.1.1 Authentication

Instantiating a client will force an authentication flow where you will be asked to authenticate with Globus Auth. Every interaction with funcX is authenticated to allow us to enforce access control on both functions and endpoints. As part of the authentication process we request access to your identity information (to retrieve your email address) and Globus Groups management access. We require Groups access in order to facilitate sharing.

```
[ ]: from funcx.sdk.client import FuncXClient  
  
fxc = FuncXClient()
```

Next we define a Python function, which we will later register with funcX. This function simply sums its input.

When defining a function you can specify `*args` and `**kwargs` as inputs.

**Note: any dependencies for a funcX function must be specified inside the function body.**

```
[ ]: def funcx_sum(items):  
      return sum(items)
```

## 3.2 Registering a function

To use a function with funcX, you must first register it with the service, using `register_function`. You can optionally include a description of the function.

The registration process will serialize the function body and transmit it to the funcX function service fabric.

Registering a function returns a UUID for the function, which can then be used to invoke it.

```
[ ]: func_uuid = fxc.register_function(funcx_sum,  
                                    description="A summation function")  
print(func_uuid)
```

## 3.3 Running a function

To invoke (perform) a function, you must provide the function's UUID, returned from the registration process, and an `endpoint_id`. Note: here we use the funcX public tutorial endpoint, which is running on AWS.

The client's `run` function will serialize any `*args` and `**kwargs`, and pass them to the function when invoking it. Therefore, as our example function simply takes an arg input (`items`), we can specify an input arg and it will be used by the function. Here we define a small list of integers for our function to sum.

The Web service will return the UUID for the invocation of the function, which we call a task. This UUID can be used to check the status of the task and retrieve the result.

```
[ ]: endpoint_uuid = '4b116d3c-1703-4f8f-9f6f-39921e5864df' # Public tutorial endpoint  
  
items = [1, 2, 3, 4, 5]  
  
res = fxc.run(items, endpoint_id=endpoint_uuid, function_id=func_uuid)  
print(res)
```

You can now retrieve the result of the invocation using `get_result()` on the UUID of the task.

```
[ ]: fxc.get_result(res)
```

The **funcX Client** is the programmatic interface to FuncX from Python. The client provides a simple and intuitive interface to:

1. Register functions
2. Register containers and execution environments
3. Launch registered function against endpoints
4. Check the status of launched functions and
5. Retrieve outputs from functions

The following shows an example of creating a client.

```
from funcx.sdk.client import FuncXClient
fxc = FuncXClient()
```

Instantiating a client will start an authentication process where you will be asked to authenticate with Globus Auth. We require every interaction with the funcX Web Service to be authenticated using a Bearer token, this allows funcX to enforce access control on both functions and endpoints. As part of the authentication process we request access to your identity information (to retrieve your email address) and Globus Groups management access. We require Groups access in order to facilitate sharing. Users can share functions with others by associating a Globus Group with the function. If a group is set, the funcX Web Service uses Groups Management access to determine whether a user is a member of the specified group.

---

**Note:** The funcX Web Service internally caches function, endpoint, and access control lookups. Caches are based on user authentication tokens. To refresh the caches you can re-authenticate your client with *force\_login=True*.

---

## 4.1 Registering Functions

You can register a Python function with funcX via `register_function()`. Function registration will serialize the function body and transmit it to the funcX Web Service. Once a function is registered with the service it will return a UUID that can be used to invoke the function.

---

**Note:** You must import any dependencies required by the function inside the function body.

---

You can associate a Globus Group with a function to enable sharing. If set, the Web Service will check if the invoking user is a member of the group. This is achieved by setting `group=<globus_group_id>` when registering a function.

You can also set a function to be publicly accessible by setting `public=True` when registering the function.

## 4.2 Running Functions

You can invoke a function using the UUID returned when registering the function. The client's `run()` function requires you to specify the `function_id` and `endpoint_id`. In addition, you can pass `*args` and `**kwargs` to the run function and they will be used when invoking your function. We serialize all inputs and outputs when running a function.

The result of your function's invocation can be retrieved using the client's `get_result()` function. This will either raise an exception if the task is still pending, or give you back the deserialized result of your invocation.

---

**Note:** If your function's execution raises an exception, `get_result()` will reraise it.

---

To minimize the overhead of communicating with the funcX Web Service we provide batch request and status capabilities. Documentation on batch requests can be found below.

## 4.3 Client Throttling

In order to avoid accidentally DoS'ing the funcX Web Service we place soft throttling restrictions on the funcX client. There are two key throttling measures: firstly, we limit the number of requests a client can make to the Web Service to 5 every 5 seconds, and secondly, we limit the size of inputs and outputs transmitted through the service to 2MB.

Batching requests and status can help reduce the number of requests made to the Web Service. In addition, the limit on the number of requests made to the Web Service can be removed by setting `throttling_enabled` to `False`.

```
func = FuncXClient()
func.throttling_enabled = False
```

## 4.4 FuncXClient Reference:

```
class funcx.FuncXClient (http_timeout=None,                                funcx_home='~/funcx',
                        force_login=False,                            fx_authorizer=None,
                        funcx_service_address='https://dev.funcx.org/api/v1', **kwargs)
```

Main class for interacting with the funcX service

Holds helper operations for performing common tasks with the funcX service.



**add\_to\_whitelist** (*endpoint\_id, function\_ids*)

Adds the function to the endpoint's whitelist

**Parameters**

- **endpoint\_id** (*str*) – The uuid of the endpoint
- **function\_ids** (*list*) – A list of function id's to be whitelisted

**Returns** The response of the request

**Return type** json

**batch\_run** (*batch*)

Initiate a batch of tasks to funcX

**Parameters** **batch** (*a Batch object*) –

**Returns** **task\_ids**

**Return type** a list of UUID strings that identify the tasks

**create\_batch** ()

Create a Batch instance to handle batch submission in funcX

**Returns** Status block containing “status” key.

**Return type** Batch instance

**delete\_from\_whitelist** (*endpoint\_id, function\_ids*)

List the endpoint's whitelist

**Parameters**

- **endpoint\_id** (*str*) – The uuid of the endpoint
- **function\_ids** (*list*) – A list of function id's to be whitelisted

**Returns** The response of the request

**Return type** json

**get\_batch\_result** (*task\_id\_list*)

Request results for a batch of task\_ids

**get\_batch\_status** (*task\_id\_list*)

Request status for a batch of task\_ids

**get\_container** (*container\_uuid, container\_type*)

Get the details of a container for staging it locally.

**Parameters**

- **container\_uuid** (*str*) – UUID of the container in question
- **container\_type** (*str*) – The type of containers that will be used (Singularity, Shifter, Docker)

**Returns** The details of the containers to deploy

**Return type** dict

**get\_containers** (*name, description=None*)

Register a DLHub endpoint with the funcX service and get the containers to launch.

**Parameters**

- **name** (*str*) – Name of the endpoint

- **description** (*str*) – Description of the endpoint

**Returns** The port to connect to and a list of containers

**Return type** `int`

**get\_endpoint\_status** (*endpoint\_uuid*)

Get the status reports for an endpoint.

**Parameters** **endpoint\_uuid** (*str*) – UUID of the endpoint in question

**Returns** The details of the endpoint's stats

**Return type** `dict`

**get\_result** (*task\_id*)

Get the result of a funcX task

**Parameters** **task\_id** (*str*) – UUID of the task

**Returns** `Result obj`

**Return type** If task completed

**Raises** Exception obj: Exception due to which the task failed

**get\_task\_status** (*task\_id*)

Get the status of a funcX task.

**Parameters** **task\_id** (*str*) – UUID of the task

**Returns** Status block containing "status" key.

**Return type** `dict`

**get\_whitelist** (*endpoint\_id*)

List the endpoint's whitelist

**Parameters** **endpoint\_id** (*str*) – The uuid of the endpoint

**Returns** The response of the request

**Return type** `json`

**logout** ()

Remove credentials from your local system

**map\_run** (*\*args, endpoint\_id=None, function\_id=None, asynchronous=False, \*\*kwargs*)

Initiate an invocation

**Parameters**

- **\*args** (*Any*) – Args as specified by the function signature
- **endpoint\_id** (*uuid str*) – Endpoint UUID string. Required
- **function\_id** (*uuid str*) – Function UUID string. Required
- **asynchronous** (*bool*) – Whether or not to run the function asynchronously

**Returns**

- **task\_id** (*str*)
- *UUID string that identifies the task*

**register\_container** (*location, container\_type, name=", description="*)

Register a container with the funcX service.

**Parameters**

- **location** (*str*) – The location of the container (e.g., its docker url). Required
- **container\_type** (*str*) – The type of containers that will be used (Singularity, Shifter, Docker). Required
- **name** (*str*) – A name for the container. Default = ''
- **description** (*str*) – A description to associate with the container. Default = ''

**Returns** The id of the container

**Return type** *str*

**register\_endpoint** (*name*, *endpoint\_uuid*, *description=None*)

Register an endpoint with the funcX service.

**Parameters**

- **name** (*str*) – Name of the endpoint
- **endpoint\_uuid** (*str*) – The uuid of the endpoint
- **description** (*str*) – Description of the endpoint

**Returns**

{'endpoint\_id' [⟨>], 'address' : ⟨>, 'client\_ports': ⟨>}

**Return type** A dict

**register\_function** (*function*, *function\_name=None*, *container\_uuid=None*, *description=None*,  
*public=False*, *group=None*)

Register a function code with the funcX service.

**Parameters**

- **function** (*Python Function*) – The function to be registered for remote execution
- **function\_name** (*str*) – The entry point (function name) of the function. Default: None
- **container\_uuid** (*str*) – Container UUID from registration with funcX
- **description** (*str*) – Description of the file
- **public** (*bool*) – Whether or not the function is publicly accessible. Default = False
- **group** (*str*) – A globus group uuid to share this function with

**Returns** **function uuid** – UUID identifier for the registered function

**Return type** *str*

**run** (*\*args*, *endpoint\_id=None*, *function\_id=None*, *\*\*kwargs*)

Initiate an invocation

**Parameters**

- **\*args** (*Any*) – Args as specified by the function signature
- **endpoint\_id** (*uuid str*) – Endpoint UUID string. Required
- **function\_id** (*uuid str*) – Function UUID string. Required
- **asynchronous** (*bool*) – Whether or not to run the function asynchronously

**Returns**

- **task\_id** (*str*)
- *UUID string that identifies the task*

**update\_table** (*return\_msg, task\_id*)

Parses the return message from the service and updates the internal func\_tables

**Parameters**

- **return\_msg** (*str*) – Return message received from the funcx service
- **task\_id** (*str*) – task id string

<i>funcx</i>	funcX : Fast function serving for clouds, clusters and supercomputers.
<i>funcx.sdk</i>	
<i>funcx.sdk.utils</i>	
<i>funcx.serialize</i>	
<i>funcx.endpoint</i>	
<i>funcx.executors</i>	
<i>funcx.providers</i>	
<i>funcx.strategies</i>	
<i>funcx.queues</i>	
<i>funcx.mock_broker</i>	

## 5.1 funcx package

### 5.1.1 Subpackages

**funcx.endpoint package**

**Subpackages**

**funcx.endpoint.utils package**

**Submodules**

**funcx.endpoint.utils.MDP module**

Majordomo Protocol definitions

**funcx.endpoint.utils.zhelpers module**

Helper module for example applications. Mimics ZeroMQ Guide's zhelpers.h.

`funcx.endpoint.utils.zhelpers.dump(msg_or_socket)`

Receives all message parts from socket, printing each frame neatly

`funcx.endpoint.utils.zhelpers.set_id(zsocket)`

Set simple random printable identity on socket

`funcx.endpoint.utils.zhelpers.socket_set_hwm(socket, hwm=-1)`

libzmq 2/3/4 compatible sethwm

`funcx.endpoint.utils.zhelpers.zpipe(ctx)`

build inproc pipe for talking to threads mimic pipe used in czmq zthread\_fork. Returns a pair of PAIRs connected via inproc

**funcx.endpoint.utils.zmq\_worker module**

**class** `funcx.endpoint.utils.zmq_worker.ZMQWorker(broker, service, verbose=False)`

Bases: `object`

Majordomo Protocol Worker API, Python version

Implements the MDP/Worker spec at <http://rfc.zeromq.org/spec:7>.

**HEARTBEAT\_LIVENESS = 3**

**broker = None**

**ctx = None**

**destroy()**

**expect\_reply = False**

**heartbeat = 2500**

**heartbeat\_at = 0**

**liveness = 0**

**reconnect = 2500**

**reconnect\_to\_broker()**

Connect or reconnect to broker

**recv()**

Send reply, if any, to broker and wait for next request.

**reply\_to = None**

**send(reply, reply\_to)**

Send a reply to the broker.

**send\_to\_broker(command, option=None, msg=None)**

Send message to broker.

If no msg is provided, creates one internally

**service = None**

**timeout = 2500**

**verbose = False**

`worker = None`

## Module contents

### Submodules

#### funcx.endpoint.auth module

`funcx.endpoint.auth.load_auth_client()`

Create an AuthClient for the portal

No credentials are used if the server is not production

**Returns** Client used to perform GlobusAuth actions

**Return type** `globus_sdk.ConfidentialAppAuthClient`

#### funcx.endpoint.config module

#### funcx.endpoint.endpoint module

`funcx.endpoint.endpoint.check_pidfile(filepath, match_name, endpoint_name)`

Helper function to identify possible dead endpoints

`funcx.endpoint.endpoint.cli_run()`

Entry point for funcx-endpoint

`funcx.endpoint.endpoint.configure_endpoint(args, config_file=None, global_config=None)`

Configure an endpoint

Drops a `config.py` template into the `funcx configs` directory. The template usually goes to `~/.funcx/<ENDPOINT_NAME>/config.py`

#### Parameters

- **args** (*args object*) – Args object from the arg parsing
- **config\_file** (*str*) – Path to a config file to be used instead of the funcX default config file
- **global\_config** (*dict*) – Global config dict

`funcx.endpoint.endpoint.init_endpoint(args)`

Setup funcx dirs and config files including a default endpoint config

TODO : Every mechanism that will update the config file, must be using a locking mechanism, ideally something like `fcntl` <https://docs.python.org/3/library/fcntl.html> to ensure that multiple endpoint invocations do not mangle the funcx config files or the lockfile module.

`funcx.endpoint.endpoint.init_endpoint_dir(funcx_dir, endpoint_name, config_file=None)`

Initialize a clean endpoint dir

Returns if an `endpoint_dir` already exists

#### Parameters

- **funcx\_dir** (*str*) – Path to the `funcx_dir` on the system

- **endpoint\_name** (*str*) – Name of the endpoint, which will be used to name the dir for the endpoint.
- **config\_file** (*str*) – Path to a config file to be used instead of the funcX default config file

`funcx.endpoint.endpoint.load_endpoint(endpoint_dir)`

**Parameters** `endpoint_dir` (*str*) – endpoint directory path within `funcx_dir`

`funcx.endpoint.endpoint.register_endpoint(funcx_client, endpoint_name, endpoint_uuid, endpoint_dir)`

Register the endpoint and return the registration info.

**Parameters**

- **funcx\_client** (`FuncXClient`) – The auth'd client to communicate with the funcX service
- **endpoint\_name** (*str*) – The name to register the endpoint with
- **endpoint\_uuid** (*str*) – The uuid to register the endpoint with
- **endpoint\_dir** (*str*) – The directory to write endpoint registration info into.

`funcx.endpoint.endpoint.register_with_hub(endpoint_uuid, endpoint_dir, address, redis_host='funcx-redis.wtgh6h.0001.use1.cache.amazonaws.com')`


This currently registers directly with the Forwarder micro service.

Can be used as an example of how to make calls this it, while the main API is updated to do this calling on behalf of the endpoint in the second iteration.

`funcx.endpoint.endpoint.start_endpoint(args, global_config=None)`

Start an endpoint

This function will do: 1. Connect to the broker service, and register itself 2. Get connection info from broker service 3. Start the interchange as a daemon



**Parameters**

- **args** (*args object*) – Args object from the arg parsing
- **global\_config** (*dict*) – Global config dict

`funcx.endpoint.endpoint.stop_endpoint(args, global_config=None)`

Stops an endpoint using the pidfile

**Parameters**

- **args** –



- `global_config`–

## `funcx.endpoint.list_endpoints` module

`funcx.endpoint.list_endpoints.list_endpoints` (*args*)

List all available endpoints

## `funcx.endpoint.version` module

### Module contents

## `funcx.executors` package

### Subpackages

## `funcx.executors.high_throughput` package

### Submodules

## `funcx.executors.high_throughput.container_sched` module

`funcx.executors.high_throughput.container_sched.naive_scheduler` (*task\_qs*,  
*outstanding\_task\_count*,  
*max\_workers*,  
*old\_worker\_map*,  
*to\_die\_list*,  
*logger*)

**Return two items (as one tuple) dict kill\_list :: KILL [(worker\_type, num\_kill), ...] dict create\_list :: CREATE [(worker\_type, num\_create), ...]**

In this scheduler model, there is minimum 1 instance of each nonempty task queue.

## `funcx.executors.high_throughput.default_config` module

## `funcx.executors.high_throughput.executor` module

HighThroughputExecutor builds on the Swift/T EMEWS architecture to use MPI for fast task distribution

There's a slow but sure deviation from Parsl's Executor interface here, that needs to be addressed.

```
class funcx.executors.high_throughput.executor.HighThroughputExecutor (label='HighThroughputExecu
    provider=LocalProvider(
        chan-
        nel=LocalChannel(
            envs={},
            script_dir=None,
            user-
            home='/home/docs/checkouts/
        ),
        cmd_timeout=30,
        init_blocks=4,
        launcher=SingleNodeLaunche
        max_blocks=10,
        min_blocks=0,
        move_files=None,
        nodes_per_block=1,
        par-
        al-
        lelism=1,
        wall-
        time='00:15:00',
        worker_init=""
    ),
    launch_cmd=None,
    ad-
    dress='127.0.0.1',
    worker_ports=None,
    worker_port_range=(54000,
    55000),
    inter-
    change_port_range=(55000,
    56000),
    stor-
    age_access=None,
    work-
    ing_dir=None,
    worker_debug=False,
    cores_per_worker=1.0,
    max_workers=inf,
    heart-
    beat_threshold=120,
    heart-
    beat_period=30,
    poll_period=10,
    con-
    tainer_image=None,
    worker_mode='singularity_re
    sup-
    press_failure=False,
    end-
    point_id=None,
    end-
    point_db=None,
    man-
    aged=True)
```

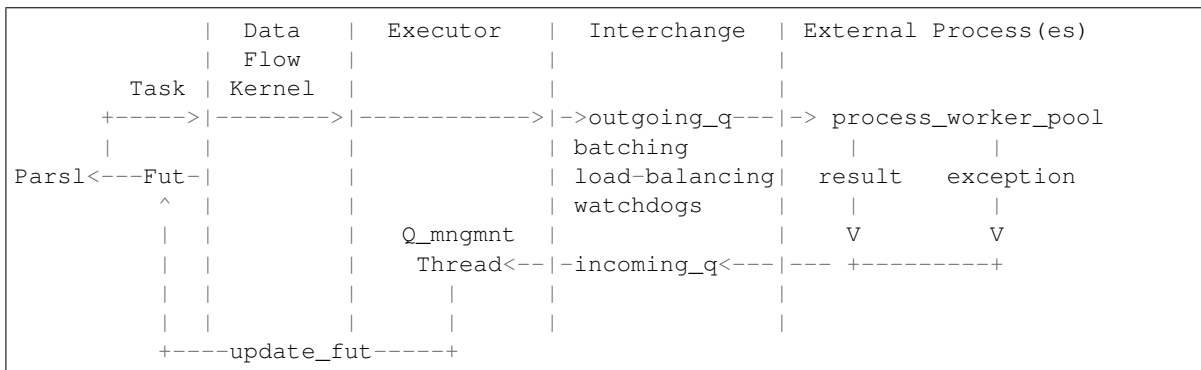
Bases: `parsl.executors.base.ParslExecutor`, `parsl.utils.RepresentationMixin`

Executor designed for cluster-scale

**The HighThroughputExecutor system has the following components:**

1. The HighThroughputExecutor instance which is run as part of the Parsl script.
2. The Interchange which acts as a load-balancing proxy between workers and Parsl
3. The multiprocessing based worker pool which coordinates task execution over several cores on a node.
4. ZeroMQ pipes connect the HighThroughputExecutor, Interchange and the `process_worker_pool`

Here is a diagram



### Parameters

- **provider** (`ExecutionProvider`) – Provider to access computation resources. Can be one of **EC2Provider**, **Cobalt**, **Condor**, **GoogleCloud**, **GridEngine**, **Jetstream**, **Local**, **GridEngine**, **Slurm**, or **Torque**.
- **label** (`str`) – Label for this executor instance.
- **launch\_cmd** (`str`) – Command line string to launch the `process_worker_pool` from the provider. The command line string will be formatted with appropriate values for the following values (`debug`, `task_url`, `result_url`, `cores_per_worker`, `nodes_per_block`, `heartbeat_period`, `heartbeat_threshold`, `logdir`). For eg: `launch_cmd="process_worker_pool.py {debug} -c {cores_per_worker} -task_url={task_url} -result_url={result_url}"`
- **address** (`string`) – An address to connect to the main Parsl process which is reachable from the network in which workers will be running. This can be either a hostname as returned by `hostname` or an IP address. Most login nodes on clusters have several network interfaces available, only some of which can be reached from the compute nodes. Some trial and error might be necessary to identify what addresses are reachable from compute nodes.
- **worker\_ports** (`(int, int)`) – Specify the ports to be used by workers to connect to Parsl. If this option is specified, `worker_port_range` will not be honored.
- **worker\_port\_range** (`(int, int)`) – Worker ports will be chosen between the two integers provided.
- **interchange\_port\_range** (`(int, int)`) – Port range used by Parsl to communicate with the Interchange.
- **working\_dir** (`str`) – Working dir to be used by the executor.

- **worker\_debug** (*Bool*) – Enables worker debug logging.
- **managed** (*Bool*) – If this executor is managed by the DFK or externally handled.
- **cores\_per\_worker** (*float*) – cores to be assigned to each worker. Oversubscription is possible by setting `cores_per_worker < 1.0`. Default=1
- **max\_workers** (*int*) – Caps the number of workers launched by the manager. Default: infinity
- **suppress\_failure** (*Bool*) – If set, the interchange will suppress failures rather than terminate early. Default: False
- **heartbeat\_threshold** (*int*) – Seconds since the last message from the counterpart in the communication pair: (interchange, manager) after which the counterpart is assumed to be un-available. Default:120s
- **heartbeat\_period** (*int*) – Number of seconds after which a heartbeat message indicating liveness is sent to the counterpart (interchange, manager). Default:30s
- **poll\_period** (*int*) – Timeout period to be used by the executor components in milliseconds. Increasing `poll_periods` trades performance for cpu efficiency. Default: 10ms
- **container\_image** (*str*) – Path or identifier to the container image to be used by the workers
- = **None** (*endpoint\_db*) – Endpoint DB object
- **worker\_mode** (*str*) – Select the mode of operation from `no_container`, `singularity_reuse`, `singularity_single_use` Default: `singularity_reuse`

#### **connected\_workers**

#### **connection\_info**

All connection info necessary for the endpoint to connect back

**Returns** Dict with connection info

#### **hold\_worker** (*worker\_id*)

Puts a worker on hold, preventing scheduling of additional tasks to it.

This is called “hold” mostly because this only stops scheduling of tasks, and does not actually kill the worker.

**Parameters** **worker\_id** (*str*) – Worker id to be put on hold

#### **initialize\_scaling** ()

Compose the launch command and call the `scale_out`

This should be implemented in the child classes to take care of executor specific oddities.

#### **outstanding**

#### **request\_status\_info** ()

#### **scale\_in** (*blocks*)

Scale in the number of active blocks by specified amount.

The scale in method here is very rude. It doesn’t give the workers the opportunity to finish current tasks or cleanup. This is tracked in issue #530

**Raises** `NotImplementedError`

#### **scale\_out** (*blocks=1*)

Scales out the number of blocks by “blocks”

**Raises** `NotImplementedError`

**scaling\_enabled**

Specify if scaling is enabled.

The callers of `ParSIExecutors` need to differentiate between `Executors` and `Executors` wrapped in a resource provider

**shutdown** (*hub=True, targets='all', block=False*)

Shutdown the executor, including all workers and controllers.

This is not implemented.

**Kwargs:**

- `hub` (Bool): Whether the hub should be shutdown, Default:True,
- `targets` (list of intsl 'all'): List of block id's to kill, Default:'all'
- `block` (Bool): To block for confirmations or not

**Raises** `NotImplementedError`

**start** ()

Create the Interchange process and connect to it.

**status** ()

Return status of all blocks.

**submit** (*bufs, task\_id=None*)

Submits work to the the outgoing\_q.

The outgoing\_q is an external process listens on this queue for new work. This method behaves like a submit call as described here [Python docs](#):

**Parameters**

- - Pickled buffer with (`b'<Function>', b'<args>', b'<kwargs>'`) (*Bufs*) -
- **Returns** - Future

**wait\_for\_endpoint** ()

**weakref\_cb** (*q=None*)

We do not use this yet.

`funcx.executors.high_throughput.executor.executor_starter` (*htex, logdir, end-point\_id, log-ging\_level=10*)

**funcx.executors.high\_throughput.funcx\_manager module**

```
class funcx.executors.high_throughput.funcx_manager.Manager (task_q_url='tcp://127.0.0.1:50097',
re-
sult_q_url='tcp://127.0.0.1:50098',
max_queue_size=10,
cores_per_worker=1,
max_workers=inf,
uid=None, heart-
beat_threshold=120,
heart-
beat_period=30,
logdir=None,
debug=False,
block_id=None,
inter-
nal_worker_port_range=(50000,
60000),
worker_mode='singularity_reuse',
sched-
uler_mode='hard',
worker_type=None,
worker_max_idletime=60,
poll_period=100)
```

Bases: `object`

Manager manages task execution by the workers

```
0mq | Manager | Worker Processes
||
<---Request N task---+---Count task reqs | Request task<--+
```

**Interchange** | **Receive task batch** | |

```
| Distribute tasks+---> Get(block) & |
|| Execute task |
||||
<-----+ Return results---+--- Post result |
||||
|| +-----+
| IPC-Queues
```

**create\_reg\_message** ()

Creates a registration message to identify the worker to the interchange

**heartbeat** ()

Send heartbeat to the incoming task queue

**pull\_tasks** (kill\_event)

Pull tasks from the incoming tasks 0mq pipe onto the internal pending task queue

**While** : receive results and task requests from the workers receive tasks/heartbeats from the Interchange  
match tasks to workers if task doesn't have appropriate worker type:

launch worker of type.. with LRU or some sort of caching strategy.

**if workers >> tasks:** advertize available capacity

**kill\_event** [threading.Event] Event to let the thread know when it is time to die.

**push\_results** (*kill\_event, max\_result\_batch\_size=1*)

Listens on the pending\_result\_queue and sends out results via 0mq

**kill\_event** [threading.Event] Event to let the thread know when it is time to die.

**remove\_worker\_init** (*worker\_type*)

Kill/Remove a worker of a given worker\_type.

Add a kill message to the task\_type queue.

Assumption : All workers of the same type are uniform, and therefore don't discriminate when killing.

**start** ()

- **while True:** Receive tasks and start appropriate workers Push tasks to available workers Forward results

`funcx.executors.high_throughput.funcx_manager.cli_run()`

### funcx.executors.high\_throughput.funcx\_worker module

```
class funcx.executors.high_throughput.funcx_worker.FuncXWorker (worker_id,
                                                             address, port,
                                                             logdir, de-
                                                             bug=False,
                                                             worker_type='RAW')
```

Bases: `object`

The FuncX worker :param worker\_id: Worker id string :type worker\_id: str :param address: Address at which the manager might be reached. This is usually 127.0.0.1 :type address: str :param port: Port at which the manager can be reached :type port: int :param logdir: Logging directory :type logdir: str :param debug: Enables debug logging :type debug: Bool

**Funcx worker will use the REP sockets to:** task = recv () result = execute(task) send(result)

**execute\_task** (*message*)

Deserialize the buffer and execute the task.

Returns the result or throws exception.

**registration\_message** ()

**start** ()

`funcx.executors.high_throughput.funcx_worker.cli_run()`

### funcx.executors.high\_throughput.global\_config module

### funcx.executors.high\_throughput.interchange module

```
exception funcx.executors.high_throughput.interchange.BadRegistration (worker_id,
                                                                    criti-
                                                                    cal=False)
```

Bases: `Exception`

A new Manager tried to join the executor with a BadRegistration message

```
class funcx.executors.high_throughput.interchange.Interchange (config,
                                                            client_address='127.0.0.1',
                                                            inter-
                                                            change_address='127.0.0.1',
                                                            client_ports=(50055,
                                                            50056, 50057),
                                                            worker_ports=None,
                                                            worker_port_range=(54000,
                                                            55000),
                                                            cores_per_worker=1.0,
                                                            worker_debug=False,
                                                            launch_cmd=None,
                                                            heart-
                                                            beat_threshold=60,
                                                            logdir='.', log-
                                                            ging_level=20,
                                                            poll_period=10,
                                                            end-
                                                            point_id=None,
                                                            sup-
                                                            press_failure=False,
                                                            max_heartbeats_missed=2)
```

Bases: `object`

Interchange is a task orchestrator for distributed systems.

1. Asynchronously queue large volume of tasks (>100K)
2. Allow for workers to join and leave the union
3. Detect workers that have failed using heartbeats
4. Service single and batch requests from workers
5. Be aware of requests worker resource capacity, eg. schedule only jobs that fit into walltime.

TODO: We most likely need a PUB channel to send out global commands, like shutdown

**get\_container** (*container\_uuid*)  
Get the container image location if it is not known to the interchange

**get\_outstanding\_breakdown** ()  
Get outstanding breakdown per manager and in the interchange queues

**Returns**

- *List of status for online elements*
- [*element, tasks\_pending, status*] ... ]

**get\_status\_report** ()  
Get utilization numbers

**get\_tasks** (*count*)  
Obtains a batch of tasks from the internal pending\_task\_queue

**Parameters** *count* (*int*) – Count of tasks to get from the queue

**Returns** eg. [{ 'task\_id':<x>, 'buffer':<buf> } ... ]

**Return type** List of upto count tasks. May return fewer than count down to an empty list



**get\_total\_live\_workers** ()

Get the total active workers

**get\_total\_tasks\_outstanding** ()

Get the outstanding tasks in total

**hold\_manager** (*manager*)

Put manager on hold :param manager: Manager id to be put on hold while being killed :type manager: str

**load\_config** ()

Load the config

**migrate\_tasks\_to\_internal** (*kill\_event*, *status\_request*)

Pull tasks from the incoming tasks 0mq pipe onto the internal pending task queue

**kill\_event** [threading.Event] Event to let the thread know when it is time to die.

**provider\_status** ()

Get status of all blocks from the provider

**scale\_in** (*blocks=None*, *block\_ids=[]*, *task\_type=None*)

Scale in the number of active blocks by specified amount.

#### Parameters

- **blocks** (*int*) – # of blocks to terminate
- **block\_ids** (*[str.. ]*) – List of external block ids to terminate

**scale\_out** (*blocks=1*, *task\_type=None*)

Scales out the number of blocks by “blocks”

**Raises** `NotImplementedError`

**start** (*poll\_period=None*)

Start the Interchange

**poll\_period** [int] poll\_period in milliseconds

**stop** ()

Prepare the interchange for shutdown

**exception** `funcx.executors.high_throughput.interchange.ManagerLost` (*worker\_id*)

Bases: `Exception`

Task lost due to worker loss. Worker is considered lost when multiple heartbeats have been missed.

**exception** `funcx.executors.high_throughput.interchange.ShutdownRequest`

Bases: `Exception`

Exception raised when any async component receives a ShutdownRequest

`funcx.executors.high_throughput.interchange.cli_run` ()

`funcx.executors.high_throughput.interchange.start_file_logger` (*filename*,  
*name='interchange'*,  
*level=10*, *format\_string=None*)

Add a stream log handler.

#### Parameters

- **filename** (*string*) – Name of the file to write logs to. Required.
- **name** (*string*) – Logger name. Default=”parsl.executors.interchange”

- **level** (*logging.LEVEL*) – Set the logging level. Default=logging.DEBUG - format\_string (string): Set the format string
- **format\_string** (*string*) – Format string to use.

**Returns**

**Return type** None.

`funcx.executors.high_throughput.interchange.starter` (*comm\_q, \*args, \*\*kwargs*)

Start the interchange process

The executor is expected to call this function. The args, kwargs match that of the `Interchange.__init__`

**funcx.executors.high\_throughput.interchange\_task\_dispatch module**

`funcx.executors.high_throughput.interchange_task_dispatch.dispatch` (*interesting\_managers, pending\_task\_queue, ready\_manager\_queue, scheduler\_mode='hard', loop='first', task\_dispatch=None, dispatched\_tasks=0*)

This is the core task dispatching algorithm for interchange. The algorithm depends on the scheduler mode and which loop.

`funcx.executors.high_throughput.interchange_task_dispatch.get_tasks_hard` (*pending\_task\_queue, manager\_ads, real\_capacity*)

`funcx.executors.high_throughput.interchange_task_dispatch.get_tasks_soft` (*pending\_task\_queue, manager\_ads, real\_capacity, loop='first'*)

`funcx.executors.high_throughput.interchange_task_dispatch.naive_interchange_task_dispatch` (*i*)

This is an initial task dispatching algorithm for interchange. It returns a dictionary, whose key is manager, and the value is the list of tasks to be sent to manager, and the total number of dispatched tasks.

**funcx.executors.high\_throughput.worker\_map module**

**class** `funcx.executors.high_throughput.worker_map.WorkerMap` (*max\_worker\_count*)

Bases: `object`

WorkerMap keeps track of workers

**add\_worker** (*worker\_id*='0.3343746937276747', *mode*='no\_container', *worker\_type*='RAW', *container\_uri*=None, *walltime*=1, *address*=None, *debug*=None, *worker\_port*=None, *logdir*=None, *uid*=None)

Launch the appropriate worker

#### Parameters

- **worker\_id** (*str*) – Worker identifier string
- **mode** (*str*) – Valid options are no\_container, singularity
- **walltime** (*int*) – Walltime in seconds before we check status

**get\_next\_worker\_q** (*new\_worker\_map*)

**Helper function to generate a queue of next workers to spin up .** From a mapping generated by the scheduler

**Parameters** *new\_worker\_map* (*dict*) – {worker\_type: total\_number\_of\_containers,... }

#### Returns

**Return type** Queue containing the next workers the system should spin-up.

**get\_worker** (*worker\_type*)

Get a task and reduce the # of worker for that type by 1. Raises queue.Empty if empty

**get\_worker\_counts** ()

Returns just the dict of worker\_type and counts

**put\_worker** (*worker*)

Adds worker to the list of waiting workers

**ready\_worker\_count** ()

**register\_worker** (*worker\_id*, *worker\_type*)

Add a new worker

**remove\_worker** (*worker\_id*)

Remove the worker from the WorkerMap

Should already be KILLED by this point.

**spin\_down\_workers** (*new\_worker\_map*, *worker\_max\_idletime*=60, *need\_more*=False, *scheduler\_mode*='hard')

Helper function to call 'remove' for appropriate workers in 'new\_worker\_map'.

**Parameters** *new\_worker\_map* (*dict*) – {worker\_type: total\_number\_of\_containers,... }.

#### Returns

**Return type** List of removed worker types.

**spin\_up\_workers** (*next\_worker\_q*, *address*=None, *debug*=None, *uid*=None, *logdir*=None, *worker\_port*=None)

Helper function to call 'remove' for appropriate workers in 'new\_worker\_map'.

#### Parameters

- **new\_worker\_q** (*queue.Queue* ()) – Queue of worker types to be spun up next.
- **address** (*str*) – Address at which to connect to the workers.
- **debug** (*bool*) – Whether debug logging is activated.
- **uid** (*str*) – Worker ID to be assigned to worker.

- **logdir** (*str*) – Directory in which to write logs
- **worker\_port** (*int*) – Port at which to connect to the workers.

**Returns**

**Return type** Total number of spun-up workers.

**update\_worker\_idle** (*worker\_type*)

Update the workers' last idle time by worker type

**funcx.executors.high\_throughput.zmq\_pipes module**

**class** funcx.executors.high\_throughput.zmq\_pipes.**CommandClient** (*ip\_address*,  
*port\_range*)

Bases: object

**close** ()

**run** (*message*)

This function needs to be fast at the same time aware of the possibility of ZMQ pipes overflowing.

The timeout increases slowly if contention is detected on ZMQ pipes. We could set copy=False and get slightly better latency but this results in ZMQ sockets reaching a broken state once there are ~10k tasks in flight. This issue can be magnified if each the serialized buffer itself is larger.

**class** funcx.executors.high\_throughput.zmq\_pipes.**ResultsIncoming** (*ip\_address*,  
*port\_range*)

Bases: object

Incoming results queue from the Interchange to the executor

**close** ()

**get** (*block=True, timeout=None*)

**request\_close** ()

**class** funcx.executors.high\_throughput.zmq\_pipes.**TasksOutgoing** (*ip\_address*,  
*port\_range*)

Bases: object

Outgoing task queue from the executor to the Interchange

**close** ()

**put** (*message, max\_timeout=1000*)

This function needs to be fast at the same time aware of the possibility of ZMQ pipes overflowing.

The timeout increases slowly if contention is detected on ZMQ pipes. We could set copy=False and get slightly better latency but this results in ZMQ sockets reaching a broken state once there are ~10k tasks in flight. This issue can be magnified if each the serialized buffer itself is larger.

**Parameters**

- **message** (*py object*) – Python object to send
- **max\_timeout** (*int*) – Max timeout in milliseconds that we will wait for before raising an exception

**Raises** zmq.EAGAIN if the send failed.



- **provider** (*ExecutionProvider*) –  
**Provider to access computation resources. Can be one of EC2Provider, Cobalt, Condor, GoogleCloud, GridEngine, Jetstream, Local, GridEngine, Slurm, or Torque.**
- **label** (*str*) – Label for this executor instance.
- **launch\_cmd** (*str*) – Command line string to launch the `process_worker_pool` from the provider. The command line string will be formatted with appropriate values for the following values (`debug`, `task_url`, `result_url`, `cores_per_worker`, `nodes_per_block`, `heartbeat_period`, `heartbeat_threshold`, `logdir`). For eg: `launch_cmd="process_worker_pool.py {debug} -c {cores_per_worker} -task_url={task_url} -result_url={result_url}"`
- **address** (*string*) – An address to connect to the main Parsl process which is reachable from the network in which workers will be running. This can be either a hostname as returned by `hostname` or an IP address. Most login nodes on clusters have several network interfaces available, only some of which can be reached from the compute nodes. Some trial and error might be necessary to indentify what addresses are reachable from compute nodes.
- **worker\_ports** (*(int, int)*) – Specify the ports to be used by workers to connect to Parsl. If this option is specified, `worker_port_range` will not be honored.
- **worker\_port\_range** (*(int, int)*) – Worker ports will be chosen between the two integers provided.
- **interchange\_port\_range** (*(int, int)*) – Port range used by Parsl to communicate with the Interchange.
- **working\_dir** (*str*) – Working dir to be used by the executor.
- **worker\_debug** (*Bool*) – Enables worker debug logging.
- **managed** (*Bool*) – If this executor is managed by the DFK or externally handled.
- **cores\_per\_worker** (*float*) – cores to be assigned to each worker. Oversubscription is possible by setting `cores_per_worker < 1.0`. Default=1
- **max\_workers** (*int*) – Caps the number of workers launched by the manager. Default: infinity
- **suppress\_failure** (*Bool*) – If set, the interchange will suppress failures rather than terminate early. Default: False
- **heartbeat\_threshold** (*int*) – Seconds since the last message from the counterpart in the communication pair: (interchange, manager) after which the counterpart is assumed to be un-available. Default:120s
- **heartbeat\_period** (*int*) – Number of seconds after which a heartbeat message indicating liveness is sent to the counterpart (interchange, manager). Default:30s
- **poll\_period** (*int*) – Timeout period to be used by the executor components in milliseconds. Increasing poll\_periods trades performance for cpu efficiency. Default: 10ms
- **container\_image** (*str*) – Path or identifier to the container image to be used by the workers
- **= None** (*endpoint\_db*) – Endpoint DB object
- **worker\_mode** (*str*) – Select the mode of operation from `no_container`, `singularity_reuse`, `singularity_single_use` Default: `singularity_reuse`

**connected\_workers****connection\_info**

All connection info necessary for the endpoint to connect back

**Returns** Dict with connection info

**hold\_worker** (*worker\_id*)

Puts a worker on hold, preventing scheduling of additional tasks to it.

This is called “hold” mostly because this only stops scheduling of tasks, and does not actually kill the worker.

**Parameters** *worker\_id* (*str*) – Worker id to be put on hold

**initialize\_scaling** ()

Compose the launch command and call the scale\_out

This should be implemented in the child classes to take care of executor specific oddities.

**outstanding****request\_status\_info** ()**scale\_in** (*blocks*)

Scale in the number of active blocks by specified amount.

The scale in method here is very rude. It doesn’t give the workers the opportunity to finish current tasks or cleanup. This is tracked in issue #530

**Raises** `NotImplementedError`

**scale\_out** (*blocks=1*)

Scales out the number of blocks by “blocks”

**Raises** `NotImplementedError`

**scaling\_enabled**

Specify if scaling is enabled.

The callers of ParslExecutors need to differentiate between Executors and Executors wrapped in a resource provider

**shutdown** (*hub=True*, *targets='all'*, *block=False*)

Shutdown the executor, including all workers and controllers.

This is not implemented.

**Kwargs:**

- *hub* (Bool): Whether the hub should be shutdown, Default:True,
- *targets* (list of ints| ‘all’): List of block id’s to kill, Default:’all’
- *block* (Bool): To block for confirmations or not

**Raises** `NotImplementedError`

**start** ()

Create the Interchange process and connect to it.

**status** ()

Return status of all blocks.

**submit** (*bufs, task\_id=None*)

Submits work to the the outgoing\_q.

The outgoing\_q is an external process listens on this queue for new work. This method behaves like a submit call as described here [Python docs](#):

**Parameters**

- - Pickled buffer with (b'<Function>', b'<args>', b'<kwargs>') (*Bufs*) -
- Returns - Future

**wait\_for\_endpoint** ()

**weakref\_cb** (*q=None*)

We do not use this yet.

## funcx.mock\_broker package

### Submodules

#### funcx.mock\_broker.forwarder module

**class** funcx.mock\_broker.forwarder.**Forwarder** (*task\_q, result\_q, executor, endpoint\_id, logdir='forwarder', logging\_level=20*)

Bases: multiprocessing.context.Process

Forwards tasks/results between the executor and the queues

**Tasks\_Q Results\_Q**

^

|

V |

Executors

Todo : We need to clarify what constitutes a task that comes down the task pipe. Does it already have the code fragment? Or does that need to be sorted out from some DB ?

**connection\_info**

Get the client ports to which the interchange must connect to

**handle\_app\_update** (*task\_id, future*)

Triggered when the executor sees a task complete.

This can be further optimized at the executor level, where we trigger this or a similar function when we see a results item inbound from the interchange.

**run** ()

Process entry point.

funcx.mock\_broker.forwarder.**double** (*x*)

funcx.mock\_broker.forwarder.**failer** (*x*)



```
funcx.mock_broker.forwarder.spawn_forwarder (address, executor=None,
                                             task_q=None, result_q=None,
                                             endpoint_id=UUID('b69ddb5a-4c9e-42fa-
                                             8887-5ab463394c57'), logging_level=20)
```

Spawns a forwarder and returns the forwarder process for tracking.

#### Parameters

- **address** (*str*) – IP Address to which the endpoint must connect
- **executor** (*Executor object. Optional*) – Executor object to be instantiated.
- **task\_q** (*Queue object*) – Queue object matching funcx.queues.base.FuncxQueue interface
- **logging\_level** (*int*) – Logging level as defined in the logging module. Default: logging.INFO (20)
- **endpoint\_id** (*uuid string*) – Endpoint id for which the forwarder is being spawned.
- **Returns** – A Forwarder object

### funcx.mock\_broker.mock\_broker module

The broker service

This REST service fields incoming registration requests from endpoints, creates an appropriate forwarder to which the endpoint can connect up.

```
funcx.mock_broker.mock_broker.list_mappings ()
```

```
funcx.mock_broker.mock_broker.register ()
```

Register an endpoint request

1. Start an executor client object corresponding to the endpoint
2. Pass connection info back as a json response.

### funcx.mock\_broker.mock\_tester module

```
funcx.mock_broker.mock_tester.test (address)
```

### funcx.mock\_broker.test module

```
funcx.mock_broker.test.double (x)
```

```
funcx.mock_broker.test.fail (x)
```

```
funcx.mock_broker.test.test_1 ()
```

```
funcx.mock_broker.test.test_2 ()
```

```
funcx.mock_broker.test.test_3 ()
```

## Module contents

### funcx.providers package

#### Subpackages

### funcx.providers.kubernetes package

#### Submodules

### funcx.providers.kubernetes.kube module

```
class funcx.providers.kubernetes.kube.KubernetesProvider (image: str, namespace: str = 'default', nodes_per_block: int = 1, init_blocks: int = 4, min_blocks: int = 0, max_blocks: int = 10, max_cpu: float = 2, max_mem: str = '500Mi', init_cpu: float = 1, init_mem: str = '250Mi', parallelism: float = 1, worker_init: str = "", pod_name: Optional[str] = None, user_id: Optional[str] = None, group_id: Optional[str] = None, run_as_non_root: bool = False, secret: Optional[str] = None, persistent_volumes: List[Tuple[str, str]] = [])
```

Bases: `parsl.providers.provider_base.ExecutionProvider`, `parsl.utils.RepresentationMixin`

Kubernetes execution provider :param namespace: Kubernetes namespace to create deployments. :type namespace: str :param image: Docker image to use in the deployment. :type image: str :param nodes\_per\_block: Nodes to provision per block. :type nodes\_per\_block: int :param init\_blocks: Number of blocks to provision at the start of the run. Default is 1. :type init\_blocks: int :param min\_blocks: Minimum number of blocks to maintain. :type min\_blocks: int :param max\_blocks: Maximum number of blocks to maintain. :type max\_blocks: int :param max\_cpu: CPU limits of the blocks (pods), in cpu units.

This is the cpu “limits” option for resource specification. Check kubernetes docs for more details. Default is 2.

#### Parameters

- **max\_mem** (*str*) – Memory limits of the blocks (pods), in Mi or Gi. This is the memory “limits” option for resource specification on kubernetes. Check kubernetes docs for more details. Default is 500Mi.

- **init\_cpu** (*float*) – CPU limits of the blocks (pods), in cpu units. This is the cpu “requests” option for resource specification. Check kubernetes docs for more details. Default is 1.
- **init\_mem** (*str*) – Memory limits of the blocks (pods), in Mi or Gi. This is the memory “requests” option for resource specification on kubernetes. Check kubernetes docs for more details. Default is 250Mi.
- **parallelism** (*float*) – Ratio of provisioned task slots to active tasks. A parallelism value of 1 represents aggressive scaling where as many resources as possible are used; parallelism close to 0 represents the opposite situation in which as few resources as possible (i.e., min\_blocks) are used.
- **worker\_init** (*str*) – Command to be run first for the workers, such as *python start.py*.
- **secret** (*str*) – Docker secret to use to pull images
- **pod\_name** (*str*) – The name for the pod, will be appended with a timestamp. Default is None, meaning parl automatically names the pod.
- **user\_id** (*str*) – Unix user id to run the container as.
- **group\_id** (*str*) – Unix group id to run the container as.
- **run\_as\_non\_root** (*bool*) – Run as non-root (True) or run as root (False).
- **persistent\_volumes** (*list[(str, str)]*) – List of tuples describing persistent volumes to be mounted in the pod. The tuples consist of (PVC Name, Mount Directory).

**cancel** (*num\_pods*, *task\_type=None*)

Cancels the resources identified by the *job\_ids* provided by the user.

**Parameters** *job\_ids* (-) – A list of job identifiers

**Returns**

- A list of status from cancelling the job which can be True, False

**Raises**

- ExecutionProviderException or its subclasses

**label**

Provides the label for this provider

**status** (*job\_ids*)

Get the status of a list of jobs identified by the job identifiers returned from the submit request. :param - *job\_ids*: A list of job identifiers :type - *job\_ids*: list

**Returns**

- A list of status from ['PENDING', 'RUNNING', 'CANCELLED', 'COMPLETED', 'FAILED', 'TIMEOUT'] corresponding to each *job\_id* in the *job\_ids* list.

**Raises**

- ExecutionProviderExceptions or its subclasses

**submit** (*cmd\_string*, *tasks\_per\_node*, *task\_type*, *job\_name='FuncX'*)

Submit a job :param - *cmd\_string*: (String) - Name of the container to initiate :param - *tasks\_per\_node*: command invocations to be launched per node :type - *tasks\_per\_node*: int

**Kwargs:**

- *job\_name* (String): Name for job, must be unique

**Returns** At capacity, cannot provision more - `job_id`: (string) Identifier for the job

**Return type**

- None

## funcx.providers.kubernetes.template module

### Module contents

### Module contents

```
class funcx.providers.KubernetesProvider (image: str, namespace: str = 'default',
                                         nodes_per_block: int = 1, init_blocks: int =
                                         4, min_blocks: int = 0, max_blocks: int =
                                         10, max_cpu: float = 2, max_mem: str =
                                         '500Mi', init_cpu: float = 1, init_mem: str =
                                         '250Mi', parallelism: float = 1, worker_init: str
                                         = "", pod_name: Optional[str] = None, user_id:
                                         Optional[str] = None, group_id: Optional[str]
                                         = None, run_as_non_root: bool = False, se-
                                         cret: Optional[str] = None, persistent_volumes:
                                         List[Tuple[str, str]] = [])
```

Bases: `parsl.providers.provider_base.ExecutionProvider`, `parsl.utils.RepresentationMixin`

Kubernetes execution provider :param namespace: Kubernetes namespace to create deployments. :type namespace: str :param image: Docker image to use in the deployment. :type image: str :param nodes\_per\_block: Nodes to provision per block. :type nodes\_per\_block: int :param init\_blocks: Number of blocks to provision at the start of the run. Default is 1. :type init\_blocks: int :param min\_blocks: Minimum number of blocks to maintain. :type min\_blocks: int :param max\_blocks: Maximum number of blocks to maintain. :type max\_blocks: int :param max\_cpu: CPU limits of the blocks (pods), in cpu units.

This is the cpu “limits” option for resource specification. Check kubernetes docs for more details. Default is 2.

### Parameters

- **max\_mem** (*str*) – Memory limits of the blocks (pods), in Mi or Gi. This is the memory “limits” option for resource specification on kubernetes. Check kubernetes docs for more details. Default is 500Mi.
- **init\_cpu** (*float*) – CPU limits of the blocks (pods), in cpu units. This is the cpu “requests” option for resource specification. Check kubernetes docs for more details. Default is 1.
- **init\_mem** (*str*) – Memory limits of the blocks (pods), in Mi or Gi. This is the memory “requests” option for resource specification on kubernetes. Check kubernetes docs for more details. Default is 250Mi.
- **parallelism** (*float*) – Ratio of provisioned task slots to active tasks. A parallelism value of 1 represents aggressive scaling where as many resources as possible are used; parallelism close to 0 represents the opposite situation in which as few resources as possible (i.e., `min_blocks`) are used.
- **worker\_init** (*str*) – Command to be run first for the workers, such as `python start.py`.

- **secret** (*str*) – Docker secret to use to pull images
- **pod\_name** (*str*) – The name for the pod, will be appended with a timestamp. Default is None, meaning parl automatically names the pod.
- **user\_id** (*str*) – Unix user id to run the container as.
- **group\_id** (*str*) – Unix group id to run the container as.
- **run\_as\_non\_root** (*bool*) – Run as non-root (True) or run as root (False).
- **persistent\_volumes** (*list[(str, str)]*) – List of tuples describing persistent volumes to be mounted in the pod. The tuples consist of (PVC Name, Mount Directory).

**cancel** (*num\_pods*, *task\_type=None*)

Cancels the resources identified by the *job\_ids* provided by the user.

**Parameters** *job\_ids* (-) – A list of job identifiers

**Returns**

- A list of status from cancelling the job which can be True, False

**Raises**

- ExecutionProviderException or its subclasses

**label**

Provides the label for this provider

**status** (*job\_ids*)

Get the status of a list of jobs identified by the job identifiers returned from the submit request. :param - *job\_ids*: A list of job identifiers :type - *job\_ids*: list

**Returns**

- A list of status from ['PENDING', 'RUNNING', 'CANCELLED', 'COMPLETED', 'FAILED', 'TIMEOUT'] corresponding to each *job\_id* in the *job\_ids* list.

**Raises**

- ExecutionProviderExceptions or its subclasses

**submit** (*cmd\_string*, *tasks\_per\_node*, *task\_type*, *job\_name='FuncX'*)

Submit a job :param - *cmd\_string*: (String) - Name of the container to initiate :param - *tasks\_per\_node*: command invocations to be launched per node :type - *tasks\_per\_node*: int

**Kwargs:**

- *job\_name* (String): Name for job, must be unique

**Returns** At capacity, cannot provision more - *job\_id*: (string) Identifier for the job

**Return type**

- None

## funcx.queues package

### Subpackages

### funcx.queues.redis package

## Submodules

### funcx.queues.redis.redis\_q module

**class** `funcx.queues.redis.redis_q.RedisQueue` (*prefix, hostname, port=6379*)

Bases: `funcx.queues.base.FuncxQueue`

A basic redis queue

The queue only connects when the `connect` method is called to avoid issues with passing an object across processes.

#### Parameters

- **hostname** (*str*) – Hostname of the redis server
- **port** (*int*) – Port at which the redis server can be reached. Default: 6379

**connect** ()

Connects to the Redis server

**get** (*timeout=1*)

Get an item from the redis queue

**Parameters** **timeout** (*int*) – Timeout for the blocking get in seconds

**is\_connected**

Returns the connected status of the queue.

#### Returns

**Return type** Bool

**put** (*key, payload*)

Put's the key:payload into a dict and pushes the key onto a queue :param key: The task\_id to be pushed :type key: str :param payload: Dict of task information to be stored :type payload: dict

`funcx.queues.redis.redis_q.test` ()

## Module contents

### Submodules

#### funcx.queues.base module

**class** `funcx.queues.base.FuncxQueue`

Bases: `object`

Queue interface required by the Forwarder

This is a metaclass that only enforces concrete implementations of functionality by the child classes.

**connect** (*\*args, \*\*kwargs*)

Connects and creates the queue. The queue is not active until this is called

**get** (*\*args, \*\*kwargs*)

Get an item from the Queue

**is\_connected**

Returns the connected status of the queue.

**Returns****Return type** Bool**put** (*\*args, \*\*kwargs*)

Put an item into the Queue

**exception** `funcx.queues.base.NotConnected` (*queue*)Bases: `funcx.errors.FuncxError`

Queue is not connected/active

**Module contents****class** `funcx.queues.RedisQueue` (*prefix, hostname, port=6379*)Bases: `funcx.queues.base.FuncxQueue`

A basic redis queue

The queue only connects when the `connect` method is called to avoid issues with passing an object across processes.**Parameters**

- **hostname** (*str*) – Hostname of the redis server
- **port** (*int*) – Port at which the redis server can be reached. Default: 6379

**connect** ()

Connects to the Redis server

**get** (*timeout=1*)

Get an item from the redis queue

**Parameters** **timeout** (*int*) – Timeout for the blocking get in seconds**is\_connected**

Returns the connected status of the queue.

**Returns****Return type** Bool**put** (*key, payload*)

Put's the key:payload into a dict and pushes the key onto a queue :param key: The task\_id to be pushed :type key: str :param payload: Dict of task information to be stored :type payload: dict

**funcx.sdk package****Subpackages****funcx.sdk.utils package****Submodules****funcx.sdk.utils.futures module**

Tools for dealing with asynchronous execution

Credit: Logan Ward

**class** `funcx.sdk.utils.futures.FuncXFuture` (*client, task\_id: str, ping\_interval: float*)  
Bases: `concurrent.futures._base.Future`

Utility class for simplifying asynchronous execution in funcX

**cancel** ()  
Stop the execution of the function

**running** ()  
Return True if the future is currently executing.

## **funcx.sdk.utils.throttling module**

**exception** `funcx.sdk.utils.throttling.MaxRequestSizeExceeded`  
Bases: `funcx.sdk.utils.throttling.ThrottlingException`

**exception** `funcx.sdk.utils.throttling.MaxRequestsExceeded`  
Bases: `funcx.sdk.utils.throttling.ThrottlingException`

**class** `funcx.sdk.utils.throttling.ThrottledBaseClient` (*\*args, \*\*kwargs*)  
Bases: `globus_sdk.base.BaseClient`

Throttled base client allows for two different types of Throttling. Note, this is ‘polite’ client side throttling, to avoid well intentioned users from mistakenly harming the funcX service, this does not prevent DOS attacks.

Request Flood Throttling: Restricts the number of requests that can be made in a given time period, and raises an exception when that has been exceeded

Request Size Throttling: Restricts the size of the payload that can be sent to the FuncX web service.

Both of these raise exceptions when the values have been exceeded. Both can be caught with `ThrottlingException`. Throttling can also be turned off with:

```
cli.throttling_enabled = False
```

```
DEFAULT_MAX_REQUESTS = 5
```

```
DEFAULT_MAX_REQUEST_SIZE = 2097152
```

```
throttle_max_requests ()
```

```
throttle_request_size (*request_args, **request_kwargs)
```

**exception** `funcx.sdk.utils.throttling.ThrottlingException`  
Bases: `Exception`

## **Module contents**

### **Submodules**

#### **funcx.sdk.client module**

**class** `funcx.sdk.client.FuncXClient` (*http\_timeout=None, funcx\_home='~/funcx',  
force\_login=False, fx\_authorizer=None,  
funcx\_service\_address='https://dev.funcx.org/api/v1',  
\*\*kwargs*)  
Bases: `funcx.sdk.utils.throttling.ThrottledBaseClient`



Main class for interacting with the funcX service

Holds helper operations for performing common tasks with the funcX service.

**CLIENT\_ID** = '4cf29807-cf21-49ec-9443-ff9a3fb9f81c'

**TOKEN\_DIR** = '/home/docs/.funcx/credentials'

**TOKEN\_FILENAME** = 'funcx\_sdk\_tokens.json'

**add\_to\_whitelist** (*endpoint\_id, function\_ids*)

Adds the function to the endpoint's whitelist

**Parameters**

- **endpoint\_id** (*str*) – The uuid of the endpoint
- **function\_ids** (*list*) – A list of function id's to be whitelisted

**Returns** The response of the request

**Return type** json

**batch\_run** (*batch*)

Initiate a batch of tasks to funcX

**Parameters** **batch** (*a Batch object*) –

**Returns** **task\_ids**

**Return type** a list of UUID strings that identify the tasks

**create\_batch** ()

Create a Batch instance to handle batch submission in funcX

**Returns** Status block containing “status” key.

**Return type** Batch instance

**delete\_from\_whitelist** (*endpoint\_id, function\_ids*)

List the endpoint's whitelist

**Parameters**

- **endpoint\_id** (*str*) – The uuid of the endpoint
- **function\_ids** (*list*) – A list of function id's to be whitelisted

**Returns** The response of the request

**Return type** json

**get\_batch\_result** (*task\_id\_list*)

Request results for a batch of task\_ids

**get\_batch\_status** (*task\_id\_list*)

Request status for a batch of task\_ids

**get\_container** (*container\_uuid, container\_type*)

Get the details of a container for staging it locally.

**Parameters**

- **container\_uuid** (*str*) – UUID of the container in question
- **container\_type** (*str*) – The type of containers that will be used (Singularity, Shifter, Docker)

**Returns** The details of the containers to deploy

**Return type** `dict`

**get\_containers** (*name*, *description=None*)

Register a DLHub endpoint with the funcX service and get the containers to launch.

**Parameters**

- **name** (*str*) – Name of the endpoint
- **description** (*str*) – Description of the endpoint

**Returns** The port to connect to and a list of containers

**Return type** `int`

**get\_endpoint\_status** (*endpoint\_uuid*)

Get the status reports for an endpoint.

**Parameters** **endpoint\_uuid** (*str*) – UUID of the endpoint in question

**Returns** The details of the endpoint's stats

**Return type** `dict`

**get\_result** (*task\_id*)

Get the result of a funcX task

**Parameters** **task\_id** (*str*) – UUID of the task

**Returns** **Result obj**

**Return type** If task completed

**Raises** Exception obj: Exception due to which the task failed

**get\_task\_status** (*task\_id*)

Get the status of a funcX task.

**Parameters** **task\_id** (*str*) – UUID of the task

**Returns** Status block containing “status” key.

**Return type** `dict`

**get\_whitelist** (*endpoint\_id*)

List the endpoint's whitelist

**Parameters** **endpoint\_id** (*str*) – The uuid of the endpoint

**Returns** The response of the request

**Return type** `json`

**logout** ()

Remove credentials from your local system

**map\_run** (*\*args*, *endpoint\_id=None*, *function\_id=None*, *asynchronous=False*, *\*\*kwargs*)

Initiate an invocation

**Parameters**

- **\*args** (*Any*) – Args as specified by the function signature
- **endpoint\_id** (*uuid str*) – Endpoint UUID string. Required
- **function\_id** (*uuid str*) – Function UUID string. Required
- **asynchronous** (*bool*) – Whether or not to run the function asynchronously

**Returns**

- **task\_id** (*str*)
- *UUID string that identifies the task*

**register\_container** (*location, container\_type, name=""*, *description=""*)

Register a container with the funcX service.

**Parameters**

- **location** (*str*) – The location of the container (e.g., its docker url). Required
- **container\_type** (*str*) – The type of containers that will be used (Singularity, Shifter, Docker). Required
- **name** (*str*) – A name for the container. Default = ""
- **description** (*str*) – A description to associate with the container. Default = ""

**Returns** The id of the container

**Return type** *str*

**register\_endpoint** (*name, endpoint\_uuid, description=None*)

Register an endpoint with the funcX service.

**Parameters**

- **name** (*str*) – Name of the endpoint
- **endpoint\_uuid** (*str*) – The uuid of the endpoint
- **description** (*str*) – Description of the endpoint

**Returns**

{'endpoint\_id' [⟨>], 'address' : ⟨>, 'client\_ports': ⟨>}

**Return type** A dict

**register\_function** (*function, function\_name=None, container\_uuid=None, description=None, public=False, group=None*)

Register a function code with the funcX service.

**Parameters**

- **function** (*Python Function*) – The function to be registered for remote execution
- **function\_name** (*str*) – The entry point (function name) of the function. Default: None
- **container\_uuid** (*str*) – Container UUID from registration with funcX
- **description** (*str*) – Description of the file
- **public** (*bool*) – Whether or not the function is publicly accessible. Default = False
- **group** (*str*) – A globus group uuid to share this function with

**Returns** **function uuid** – UUID identifier for the registered function

**Return type** *str*

**run** (*\*args, endpoint\_id=None, function\_id=None, \*\*kwargs*)

Initiate an invocation

**Parameters**

- **\*args** (*Any*) – Args as specified by the function signature

- **endpoint\_id** (*uuid str*) – Endpoint UUID string. Required
- **function\_id** (*uuid str*) – Function UUID string. Required
- **asynchronous** (*bool*) – Whether or not to run the function asynchronously

**Returns**

- **task\_id** (*str*)
- *UUID string that identifies the task*

**update\_table** (*return\_msg, task\_id*)

Parses the return message from the service and updates the internal func\_tables

**Parameters**

- **return\_msg** (*str*) – Return message received from the funcx service
- **task\_id** (*str*) – task id string

## funcx.sdk.config module

`funcx.sdk.config.write_option` (*option, value*)

Write an option to disk – doesn't handle config reloading

`funcx.sdk.config.lookup_option` (*option*)

`funcx.sdk.config.remove_option` (*option*)

`funcx.sdk.config.internal_auth_client` ()

Get the globus native app client.

**Returns**

`funcx.sdk.config.check_logged_in` ()

Check if the user is already logged in.

**Returns**

`funcx.sdk.config.safeprint` (*s*)

Catch print errors.

**Parameters** *s* –

**Returns**

`funcx.sdk.config.format_output` (*dataobject*)

Use safe print to make sure jobs are correctly printed.

**Parameters** *dataobject* –

**Returns**

## funcx.sdk.version module

### Module contents

### funcx.serialize package

### Submodules

**funcx.serialize.base module**

**exception** funcx.serialize.base.**DeserializationError** (*reason*)

Bases: `Exception`

Base class for all deserialization errors

**class** funcx.serialize.base.**RemoteExceptionWrapper** (*e\_type, e\_value, traceback*)

Bases: `object`

**reraise** ()

**class** funcx.serialize.base.**SerializerError** (*reason*)

Bases: `object`

**class** funcx.serialize.base.**fxPicker\_enforcer**

Bases: `object`

Ensure that any concrete class will have the serialize and deserialize methods

**deserialize** (*payload*)

**serialize** (*data*)

**class** funcx.serialize.base.**fxPicker\_shared**

Bases: `object`

Adds shared functionality for all serializer implementations

**check** (*payload*)

**chomp** (*payload*)

If the payload starts with the identifier, return the remaining block

**Parameters** *payload* (*str*) – Payload blob

**identifier**

Get the identifier of the serialization method

**Returns** *identifier*

**Return type** *str*

**funcx.serialize.concretes module**

funcx.serialize.concretes.**bar** (*x, y={‘a’: 3}*)

**class** funcx.serialize.concretes.**code\_pickle**

Bases: `funcx.serialize.base.fxPicker_shared`

**deserialize** (*payload*)

**serialize** (*data*)

**class** funcx.serialize.concretes.**code\_text\_dill**

Bases: `funcx.serialize.base.fxPicker_shared`

We use dill to get the source code out of the function object and then exec the function body to load it in. The function object is then returned by name.

**deserialize** (*payload*)

**serialize** (*data*)

```
class funcx.serialize.concretes.code_text_inspect
```

```
    Bases: funcx.serialize.base.fxPicker_shared
```

We use dill to get the source code out of the function object and then exec the function body to load it in. The function object is then returned by name.

```
    deserialize (payload)
```

```
    serialize (data)
```

```
class funcx.serialize.concretes.json_base64
```

```
    Bases: funcx.serialize.base.fxPicker_shared
```

```
    deserialize (payload)
```

```
    serialize (data)
```

```
class funcx.serialize.concretes.pickle_base64
```

```
    Bases: funcx.serialize.base.fxPicker_shared
```

```
    deserialize (payload)
```

```
    serialize (data)
```

## **funcx.serialize.facade module**

```
class funcx.serialize.facade.FuncXSerializer
```

```
    Bases: object
```

Information that we want to be able to ship around:

- Function run information ->
- Function id <—> Container id ? (is there a 1-1 mapping here?)
- Container id
- Endpoint id
- Function body
- Potentially in a byte compiled form ?
- All parameters ->
- Args + Kwargs

From the client side. At function invocation we need to capture

```
    deserialize (payload)
```

```
        Parameters payload (str) – Payload object to be deserialized
```

```
    pack_buffers (buffers)
```

```
        buffers : list of
```

```
        terminated strings
```

```
    serialize (data)
```

```
    unpack_and_deserialize (packed_buffer)
```

```
        Unpacks a packed buffer and returns the deserialized contents :param packed_buffers: :type  
        packed_buffers: packed buffer as string
```

```
    unpack_buffers (packed_buffer)
```

Parameters **packed\_buffers** (*packed buffer as string*)–

## Module contents

### funcx.strategies package

#### Submodules

#### funcx.strategies.base module

**class** funcx.strategies.base.**BaseStrategy** (\*args, threshold=20, interval=5)

Bases: `object`

Implements threshold-interval based flow control.

The overall goal is to trap the flow of apps from the workflow, measure it and redirect it the appropriate executors for processing.

This is based on the following logic:

```
BEGIN (INTERVAL, THRESHOLD, callback) :
    start = current_time()

    while (current_time()-start < INTERVAL) :
        count = get_events_since(start)
        if count >= THRESHOLD :
            break

    callback()
```

This logic ensures that the callbacks are activated with a maximum delay of *interval* for systems with infrequent events as well as systems which would generate large bursts of events.

Once a callback is triggered, the callback generally runs a strategy method on the sites available as well as queue

TODO: When the debug logs are enabled this module emits duplicate messages. This issue needs more debugging. What I've learnt so far is that the duplicate messages are present only when the timer thread is started, so this could be from a duplicate logger being added by the thread.

**close** ()

Merge the threads and terminate.

**make\_callback** (*kind=None*)

Makes the callback and resets the timer.

**KWargs:**

- *kind* (str): Default=None, used to pass information on what triggered the callback

**notify** (*event\_id*)

Let the FlowControl system know that there is an event.

This method is to be called from the Interchange to notify the flowcontrol

**start** (*interchange*)

Actually start the strategy :param interchange: Interchange to bind the strategy to :type interchange: funcx.executors.high\_throughput.interchange.Interchange

**strategize** (\*args, \*\*kwargs)

Strategize is called everytime the threshold or the interval is hit

**class** funcx.strategies.base.**Timer** (*callback, \*args, interval=5*)

Bases: `object`

This timer is a simplified version of the FlowControl timer. This timer does not employ notify events.

This is based on the following logic :

```
BEGIN (INTERVAL, THRESHOLD, callback) :
    start = current_time()

    while (current_time()-start < INTERVAL) :
        wait()
        break

    callback()
```

**close** ()

Merge the threads and terminate.

**make\_callback** (*kind=None*)

Makes the callback and resets the timer.

### **funcx.strategies.kube\_simple module**

**class** funcx.strategies.kube\_simple.**KubeSimpleStrategy** (*\*args, threshold=20, interval=1, max\_idletime=60*)

Bases: `funcx.strategies.base.BaseStrategy`

Implements the simple strategy for Kubernetes

**strategize** (*\*args, \*\*kwargs*)

Strategize is called everytime the threshold or the interval is hit

### **funcx.strategies.simple module**

**class** funcx.strategies.simple.**SimpleStrategy** (*\*args, threshold=20, interval=1, max\_idletime=60*)

Bases: `funcx.strategies.base.BaseStrategy`

Implements the simple strategy

**strategize** (*\*args, \*\*kwargs*)

Strategize is called everytime the threshold or the interval is hit

### **funcx.strategies.test module**

#### **Module contents**

**class** funcx.strategies.**BaseStrategy** (*\*args, threshold=20, interval=5*)

Bases: `object`

Implements threshold-interval based flow control.

The overall goal is to trap the flow of apps from the workflow, measure it and redirect it the appropriate executors for processing.

This is based on the following logic:



```

BEGIN (INTERVAL, THRESHOLD, callback) :
    start = current_time()

    while (current_time()-start < INTERVAL) :
        count = get_events_since(start)
        if count >= THRESHOLD :
            break

    callback()

```

This logic ensures that the callbacks are activated with a maximum delay of *interval* for systems with infrequent events as well as systems which would generate large bursts of events.

Once a callback is triggered, the callback generally runs a strategy method on the sites available as well as queue

TODO: When the debug logs are enabled this module emits duplicate messages. This issue needs more debugging. What I've learnt so far is that the duplicate messages are present only when the timer thread is started, so this could be from a duplicate logger being added by the thread.

**close()**

Merge the threads and terminate.

**make\_callback** (*kind=None*)

Makes the callback and resets the timer.

**KWargs:**

- *kind* (str): Default=None, used to pass information on what triggered the callback

**notify** (*event\_id*)

Let the FlowControl system know that there is an event.

This method is to be called from the Interchange to notify the flowcontrol

**start** (*interchange*)

Actually start the strategy ;param interchange: Interchange to bind the strategy to :type interchange: funcx.executors.high\_throughput.interchange.Interchange

**strategize** (*\*args, \*\*kwargs*)

Strategize is called everytime the threshold or the interval is hit

**class** funcx.strategies.**SimpleStrategy** (*\*args, threshold=20, interval=1, max\_idletime=60*)

Bases: *funcx.strategies.base.BaseStrategy*

Implements the simple strategy

**strategize** (*\*args, \*\*kwargs*)

Strategize is called everytime the threshold or the interval is hit

**class** funcx.strategies.**KubeSimpleStrategy** (*\*args, threshold=20, interval=1, max\_idletime=60*)

Bases: *funcx.strategies.base.BaseStrategy*

Implements the simple strategy for Kubernetes

**strategize** (*\*args, \*\*kwargs*)

Strategize is called everytime the threshold or the interval is hit

## 5.1.2 Submodules

### 5.1.3 funcx.config module

```
class funcx.config.Config (provider=LocalProvider( channel=LocalChannel( envs={},
script_dir=None, userhome='/home/docs/checkouts/readthedocs.org/user_builds/funcx/checkouts/funcx'),
cmd_timeout=30, init_blocks=4, launcher=SingleNodeLauncher(),
max_blocks=10, min_blocks=0, move_files=None, nodes_per_block=1,
parallelism=1, walltime='00:15:00', worker_init=" "), scaling_enabled=True,
worker_ports=None, worker_port_range=(54000, 55000),
strategy=<funcx.strategies.simple.SimpleStrategy object>,
max_workers_per_node=inf, cores_per_worker=1.0,
mem_per_worker=None, launch_cmd=None,
worker_mode='no_container', scheduler_mode='hard',
container_type=None, prefetch_capacity=10, heartbeat_period=2,
heartbeat_threshold=10, poll_period=10, working_dir=None,
worker_debug=False)
```

Bases: `parsl.utils.RepresentationMixin`

Specification of FuncX configuration options.

#### Parameters

- **max\_workers\_per\_node** (*int*) – Maximum # of worker per node. Default: `inf`
- **cores\_per\_worker** (*float*) – cores to be assigned to each worker. Oversubscription is possible by setting `cores_per_worker < 1.0`. Default=1
- **mem\_per\_worker** (*float*) – GB of memory required per worker. If this option is specified, the node manager will check the available memory at startup and limit the number of workers such that there's sufficient memory for each worker. Default: `None`
- **working\_dir** (*str*) – Working dir to be used by the executor. Default to the endpoint directory if not specified
- **worker\_debug** (*Bool*) – Enables worker debug logging.
- **worker\_mode** (*str*) – Select the mode of operation from `no_container`, `singularity_reuse`, `singularity_single_use` Default: `no_container`
- **scheduler\_mode** (*str*) – Select the mode of how the container is managed from `hard`, `soft` Default: `hard`
- **container\_type** (*str*) – Select the type of container from `Docker`, `Singularity`, `Shifter` Default: `None`
- **scaling\_enabled** (*Bool*) – Allow Interchange to manage resource provisioning. If set to `False`, interchange will not do any scaling. Default: `True`

### 5.1.4 funcx.errors module

```
exception funcx.errors.FuncXUnreachable (address)
```

Bases: `funcx.errors.FuncxError`

FuncX remote service is unreachable

```
exception funcx.errors.FuncxError
```

Bases: `Exception`

Base class for all funcx exceptions

**exception** `funcx.errors.MalformedResponse` (*response*)

Bases: `funcx.errors.FuncxError`

FuncX remote service responded with a Malformed Response

**exception** `funcx.errors.RegistrationError` (*reason*)

Bases: `funcx.errors.FuncxError`

Registering the endpoint has failed

## 5.1.5 funcx.version module

Set module version.

<Major>.<Minor>.<maintenance>[alpha/beta/..] Alphas will be numbered like this -> 0.4.0a0

## 5.1.6 Module contents

funcX : Fast function serving for clouds, clusters and supercomputers.

`funcx.set_file_logger` (*filename*, *name='funcx'*, *level=10*, *format\_string=None*)

Add a stream log handler.

### Parameters

- **filename** (-) – Name of the file to write logs to
- **name** (-) – Logger name
- **level** (-) – Set the logging level.
- **format\_string** (-) – Set the format string

### Returns

- None

`funcx.set_stream_logger` (*name='funcx'*, *level=10*, *format\_string=None*)

Add a stream log handler.

### Parameters

- **name** (-) – Set the logger name.
- **level** (-) – Set to logging.DEBUG by default.
- **format\_string** (-) – Set to None by default.

### Returns

- None



To sidestep the FuncX web-service the endpoint can be forced to connect directly to a debug-forwarder. To enable this behavior update `~/funcx/config`

```
* 'broker_address' : Address at which the broker is listening, eg: "http://127.0.0.1:  
↪50005"  
* 'broker_test' : True  
* 'redis_host' : Point redis host to a locally running redis server. Eg: "127.0.0.1"
```

## 6.1 Setting up the forwarder locally

You can run the forwarder-service in debug mode on your local system and skip the web-service entirely. For this, make sure you have the redis package installed and running. You can check this by running:

```
>>> redis-cli
```

This should output a prompt that says : 127.0.0.1:6379. This string needs to match.

Now, you can start the forwarder service for testing:

```
>>> forwarder-service --address 127.0.0.1 --port 50005 --debug
```

Once you have this running, we can update the endpoint configs to point to this local service.



## CHAPTER 7

---

### Indices and tables

---

- `genindex`
- `search`





### f

- funcx, 55
- funcx.config, 54
- funcx.endpoint, 21
- funcx.endpoint.auth, 19
- funcx.endpoint.config, 19
- funcx.endpoint.endpoint, 19
- funcx.endpoint.list\_endpoints, 21
- funcx.endpoint.utils, 19
- funcx.endpoint.utils.MDP, 17
- funcx.endpoint.utils.zhelpers, 18
- funcx.endpoint.utils.zmq\_worker, 18
- funcx.endpoint.version, 21
- funcx.errors, 54
- funcx.executors, 33
- funcx.executors.high\_throughput, 33
- funcx.executors.high\_throughput.container\_sched,  
21
- funcx.executors.high\_throughput.default\_config,  
21
- funcx.executors.high\_throughput.executor,  
21
- funcx.executors.high\_throughput.funcx\_manager,  
26
- funcx.executors.high\_throughput.funcx\_worker,  
27
- funcx.executors.high\_throughput.global\_config,  
27
- funcx.executors.high\_throughput.interchange,  
27
- funcx.executors.high\_throughput.interchange\_task\_dispatch,  
30
- funcx.executors.high\_throughput.worker\_map,  
30
- funcx.executors.high\_throughput.zmq\_pipes,  
32
- funcx.mock\_broker, 38
- funcx.mock\_broker.forwarder, 36
- funcx.mock\_broker.mock\_broker, 37
- funcx.mock\_broker.mock\_tester, 37
- funcx.mock\_broker.test, 37
- funcx.providers, 40
- funcx.providers.kubernetes, 40
- funcx.providers.kubernetes.kube, 38
- funcx.providers.kubernetes.template, 40
- funcx.queues, 43
- funcx.queues.base, 42
- funcx.queues.redis, 42
- funcx.queues.redis.redis\_q, 42
- funcx.sdk, 48
- funcx.sdk.client, 44
- funcx.sdk.config, 48
- funcx.sdk.utils, 44
- funcx.sdk.utils.futures, 43
- funcx.sdk.utils.throttling, 44
- funcx.sdk.version, 48
- funcx.serialize, 51
- funcx.serialize.base, 49
- funcx.serialize.concretes, 49
- funcx.serialize.facade, 50
- funcx.strategies, 52
- funcx.strategies.base, 51
- funcx.strategies.kube\_simple, 52
- funcx.strategies.simple, 52
- funcx.version, 55



## A

`add_to_whitelist()` (*funcx.FuncXClient* method), 12

`add_to_whitelist()` (*funcx.sdk.client.FuncXClient* method), 45

`add_worker()` (*funcx.executors.high\_throughput.worker\_map.WorkerMap* method), 30

## B

`BadRegistration`, 27

`bar()` (in module *funcx.serialize.concretes*), 49

`BaseStrategy` (class in *funcx.strategies*), 52

`BaseStrategy` (class in *funcx.strategies.base*), 51

`batch_run()` (*funcx.FuncXClient* method), 13

`batch_run()` (*funcx.sdk.client.FuncXClient* method), 45

`broker` (*funcx.endpoint.utils.zmq\_worker.ZMQWorker* attribute), 18

## C

`cancel()` (*funcx.providers.kubernetes.kube.KubernetesProvider* method), 39

`cancel()` (*funcx.providers.KubernetesProvider* method), 41

`cancel()` (*funcx.sdk.utils.futures.FuncXFuture* method), 44

`check()` (*funcx.serialize.base.fxPicker\_shared* method), 49

`check_logged_in()` (in module *funcx.sdk.config*), 48

`check_pidfile()` (in module *funcx.endpoint.endpoint*), 19

`chomp()` (*funcx.serialize.base.fxPicker\_shared* method), 49

`cli_run()` (in module *funcx.endpoint.endpoint*), 19

`cli_run()` (in module *funcx.executors.high\_throughput.funcx\_manager*), 27

`cli_run()` (in module *funcx.executors.high\_throughput.funcx\_worker*), 27

`cli_run()` (in module *funcx.executors.high\_throughput.interchange*), 29

`CLIENT_ID` (*funcx.sdk.client.FuncXClient* attribute), 45

`close()` (*funcx.executors.high\_throughput.zmq\_pipes.CommandClient* method), 32

`close()` (*funcx.executors.high\_throughput.zmq\_pipes.ResultsIncoming* method), 32

`close()` (*funcx.executors.high\_throughput.zmq\_pipes.TasksOutgoing* method), 32

`close()` (*funcx.strategies.base.BaseStrategy* method), 51

`close()` (*funcx.strategies.base.Timer* method), 52

`close()` (*funcx.strategies.BaseStrategy* method), 53

`code_pickle` (class in *funcx.serialize.concretes*), 49

`code_text_dill` (class in *funcx.serialize.concretes*), 49

`code_text_inspect` (class in *funcx.serialize.concretes*), 49

`CommandClient` (class in *funcx.executors.high\_throughput.zmq\_pipes*), 32

`Config` (class in *funcx.config*), 54

`configure_endpoint()` (in module *funcx.endpoint.endpoint*), 19

`connect()` (*funcx.queues.base.FuncxQueue* method), 42

`connect()` (*funcx.queues.redis.redis\_q.RedisQueue* method), 42

`connect()` (*funcx.queues.RedisQueue* method), 43

`connected_workers` (*funcx.executors.high\_throughput.executor.HighThroughputExecutor* attribute), 24

`connected_workers` (*funcx.executors.HighThroughputExecutor* attribute), 34

`connection_info` (*funcx.executors.high\_throughput.executor.HighThroughputExecutor* attribute), 24

- attribute*), 24
  - connection\_info (*funcx.executors.HighThroughputExecutor* *attribute*), 35
  - connection\_info (*funcx.mock\_broker.forwarder.Forwarder* *attribute*), 36
  - create\_batch() (*funcx.FuncXClient* method), 13
  - create\_batch() (*funcx.sdk.client.FuncXClient* method), 45
  - create\_reg\_message() (*funcx.executors.high\_throughput.funcx\_manager.Manager* method), 26
  - ctx (*funcx.endpoint.utils.zmq\_worker.ZMQWorker* *attribute*), 18
- D**
- DEFAULT\_MAX\_REQUEST\_SIZE (*funcx.sdk.utils.throttling.ThrottledBaseClient* *attribute*), 44
  - DEFAULT\_MAX\_REQUESTS (*funcx.sdk.utils.throttling.ThrottledBaseClient* *attribute*), 44
  - delete\_from\_whitelist() (*funcx.FuncXClient* method), 13
  - delete\_from\_whitelist() (*funcx.sdk.client.FuncXClient* method), 45
  - DeserializationError, 49
  - deserialize() (*funcx.serialize.base.fxPicker\_enforcer* method), 49
  - deserialize() (*funcx.serialize.concretes.code\_pickle* method), 49
  - deserialize() (*funcx.serialize.concretes.code\_text\_dill* method), 49
  - deserialize() (*funcx.serialize.concretes.code\_text\_inspect* method), 50
  - deserialize() (*funcx.serialize.concretes.json\_base64* method), 50
  - deserialize() (*funcx.serialize.concretes.pickle\_base64* method), 50
  - deserialize() (*funcx.serialize.facade.FuncXSerializer* method), 50
  - destroy() (*funcx.endpoint.utils.zmq\_worker.ZMQWorker* method), 18
  - dispatch() (*in module funcx.executors.high\_throughput.interchange\_task\_dispatcher*), 30
  - double() (*in module funcx.mock\_broker.forwarder*), 36
  - double() (*in module funcx.mock\_broker.test*), 37
  - dump() (*in module funcx.endpoint.utils.zhelpers*), 18
- E**
- execute\_task() (*funcx.executors.high\_throughput.funcx\_worker.FuncXWorker* method), 27
  - executor\_starter() (*in module funcx.executors.high\_throughput.executor*), 25
  - expect\_reply (*funcx.endpoint.utils.zmq\_worker.ZMQWorker* *attribute*), 18
- F**
- fail() (*in module funcx.mock\_broker.test*), 37
  - failer() (*in module funcx.mock\_broker.forwarder*), 36
  - format\_output() (*in module funcx.sdk.config*), 48
  - Forwarder (*class in funcx.mock\_broker.forwarder*), 36
  - funcx (module), 55
  - funcx.config (module), 54
  - funcx.endpoint (module), 21
  - funcx.endpoint.auth (module), 19
  - funcx.endpoint.config (module), 19
  - funcx.endpoint.endpoint (module), 19
  - funcx.endpoint.list\_endpoints (module), 21
  - funcx.endpoint.utils (module), 19
  - funcx.endpoint.utils.MDP (module), 17
  - funcx.endpoint.utils.zhelpers (module), 18
  - funcx.endpoint.utils.zmq\_worker (module), 18
  - funcx.endpoint.version (module), 21
  - funcx.errors (module), 54
  - funcx.executors (module), 33
  - funcx.executors.high\_throughput (module), 33
  - funcx.executors.high\_throughput.container\_sched (module), 21
  - funcx.executors.high\_throughput.default\_config (module), 21
  - funcx.executors.high\_throughput.executor (module), 21
  - funcx.executors.high\_throughput.funcx\_manager (module), 26
  - funcx.executors.high\_throughput.funcx\_worker (module), 27
  - funcx.executors.high\_throughput.global\_config (module), 27
  - funcx.executors.high\_throughput.interchange (module), 27
  - funcx.executors.high\_throughput.interchange\_task\_dispatcher (module), 30
  - funcx.executors.high\_throughput.worker\_map (module), 30
  - funcx.executors.high\_throughput.zmq\_pipes (module), 32
  - funcx.mock\_broker (module), 38
  - funcx.mock\_broker.forwarder (module), 36
  - funcx.mock\_broker.mock\_broker (module), 37
  - funcx.mock\_broker.mock\_tester (module), 37
  - funcx.mock\_broker.test (module), 37

- funcx.providers (module), 40  
 funcx.providers.kubernetes (module), 40  
 funcx.providers.kubernetes.kube (module), 38  
 funcx.providers.kubernetes.template (module), 40  
 funcx.queues (module), 43  
 funcx.queues.base (module), 42  
 funcx.queues.redis (module), 42  
 funcx.queues.redis.redis\_q (module), 42  
 funcx.sdk (module), 48  
 funcx.sdk.client (module), 44  
 funcx.sdk.config (module), 48  
 funcx.sdk.utils (module), 44  
 funcx.sdk.utils.futures (module), 43  
 funcx.sdk.utils.throttling (module), 44  
 funcx.sdk.version (module), 48  
 funcx.serialize (module), 51  
 funcx.serialize.base (module), 49  
 funcx.serialize.concretes (module), 49  
 funcx.serialize.facade (module), 50  
 funcx.strategies (module), 52  
 funcx.strategies.base (module), 51  
 funcx.strategies.kube\_simple (module), 52  
 funcx.strategies.simple (module), 52  
 funcx.version (module), 55  
 FuncXClient (class in funcx), 12  
 FuncXClient (class in funcx.sdk.client), 44  
 FuncXError, 54  
 FuncXFuture (class in funcx.sdk.utils.futures), 44  
 FuncXQueue (class in funcx.queues.base), 42  
 FuncXSerializer (class in funcx.serialize.facade), 50  
 FuncXUnreachable, 54  
 FuncXWorker (class in funcx.executors.high\_throughput.funcx\_worker), 27  
 fxPicker\_enforcer (class in funcx.serialize.base), 49  
 fxPicker\_shared (class in funcx.serialize.base), 49
- ## G
- get () (funcx.executors.high\_throughput.zmq\_pipes.ResultsIncoming method), 32  
 get () (funcx.queues.base.FuncXQueue method), 42  
 get () (funcx.queues.redis.redis\_q.RedisQueue method), 42  
 get () (funcx.queues.RedisQueue method), 43  
 get\_batch\_result () (funcx.FuncXClient method), 13  
 get\_batch\_result () (funcx.sdk.client.FuncXClient method), 45  
 get\_batch\_status () (funcx.FuncXClient method), 13  
 get\_batch\_status () (funcx.sdk.client.FuncXClient method), 45  
 get\_container () (funcx.executors.high\_throughput.interchange.Interchange method), 28  
 get\_container () (funcx.FuncXClient method), 13  
 get\_container () (funcx.sdk.client.FuncXClient method), 45  
 get\_containers () (funcx.FuncXClient method), 13  
 get\_containers () (funcx.sdk.client.FuncXClient method), 46  
 get\_endpoint\_status () (funcx.FuncXClient method), 14  
 get\_endpoint\_status () (funcx.sdk.client.FuncXClient method), 46  
 get\_next\_worker\_q () (funcx.executors.high\_throughput.worker\_map.WorkerMap method), 31  
 get\_outstanding\_breakdown () (funcx.executors.high\_throughput.interchange.Interchange method), 28  
 get\_result () (funcx.FuncXClient method), 14  
 get\_result () (funcx.sdk.client.FuncXClient method), 46  
 get\_status\_report () (funcx.executors.high\_throughput.interchange.Interchange method), 28  
 get\_task\_status () (funcx.FuncXClient method), 14  
 get\_task\_status () (funcx.sdk.client.FuncXClient method), 46  
 get\_tasks () (funcx.executors.high\_throughput.interchange.Interchange method), 28  
 get\_tasks\_hard () (in module funcx.executors.high\_throughput.interchange\_task\_dispatch), 30  
 get\_tasks\_soft () (in module funcx.executors.high\_throughput.interchange\_task\_dispatch), 30  
 get\_total\_live\_workers () (funcx.executors.high\_throughput.interchange.Interchange method), 28  
 get\_total\_tasks\_outstanding () (funcx.executors.high\_throughput.interchange.Interchange method), 29  
 get\_whitelist () (funcx.FuncXClient method), 14  
 get\_whitelist () (funcx.sdk.client.FuncXClient method), 46  
 get\_worker () (funcx.executors.high\_throughput.worker\_map.WorkerMap method), 31  
 get\_worker\_counts () (funcx.executors.high\_throughput.worker\_map.WorkerMap method), 31

## H

handle\_app\_update ()  
 (*funcx.mock\_broker.forwarder.Forwarder*  
*method*), 36

heartbeat (*funcx.endpoint.utils.zmq\_worker.ZMQWorker*  
*attribute*), 18

heartbeat () (*funcx.executors.high\_throughput.funcx\_manager.Manager*  
*method*), 26

heartbeat\_at (*funcx.endpoint.utils.zmq\_worker.ZMQWorker*  
*attribute*), 18

HEARTBEAT\_LIVENESS  
 (*funcx.endpoint.utils.zmq\_worker.ZMQWorker*  
*attribute*), 18

HighThroughputExecutor (class in  
*funcx.executors*), 33

HighThroughputExecutor (class in  
*funcx.executors.high\_throughput.executor*),  
 21

hold\_manager () (*funcx.executors.high\_throughput.interchange.Interchange*  
*method*), 29

hold\_worker () (*funcx.executors.high\_throughput.executor.HighThroughputExecutor*  
*method*), 24

hold\_worker () (*funcx.executors.HighThroughputExecutor*  
*method*), 35

## I

identifier (*funcx.serialize.base.fxPicker\_shared* at-  
*tribute*), 49

init\_endpoint () (in module  
*funcx.endpoint.endpoint*), 19

init\_endpoint\_dir () (in module  
*funcx.endpoint.endpoint*), 19

initialize\_scaling ()  
 (*funcx.executors.high\_throughput.executor.HighThroughputExecutor*  
*method*), 24

initialize\_scaling ()  
 (*funcx.executors.HighThroughputExecutor*  
*method*), 35

Interchange (class in  
*funcx.executors.high\_throughput.interchange*),  
 28

internal\_auth\_client () (in module  
*funcx.sdk.config*), 48

is\_connected (*funcx.queues.base.FuncxQueue* at-  
*tribute*), 42

is\_connected (*funcx.queues.redis.redis\_q.RedisQueue*  
*attribute*), 42

is\_connected (*funcx.queues.RedisQueue* attribute),  
 43

## J

json\_base64 (class in *funcx.serialize.concretes*), 50

## K

KubernetesProvider (class in *funcx.providers*), 40

KubernetesProvider (class in  
*funcx.providers.kubernetes.kube*), 38

KubeSimpleStrategy (class in *funcx.strategies*), 53

KubeSimpleStrategy (class in  
*funcx.strategies.kube\_simple*), 52

label (*funcx.providers.kubernetes.kube.KubernetesProvider*  
*attribute*), 39

label (*funcx.providers.KubernetesProvider* attribute),  
 41

list\_endpoints () (in module  
*funcx.endpoint.list\_endpoints*), 21

list\_mappings () (in module  
*funcx.mock\_broker.mock\_broker*), 37

liveness (*funcx.endpoint.utils.zmq\_worker.ZMQWorker*  
*attribute*), 18

load\_auth\_client () (in module  
*funcx.endpoint.endpoint*), 19

load\_config () (*funcx.executors.high\_throughput.interchange.Interchange*  
*method*), 29

load\_endpoint () (in module  
*funcx.endpoint.endpoint*), 20

logout () (*funcx.FuncXClient* method), 14

logout () (*funcx.sdk.client.FuncXClient* method), 46

lookup\_option () (in module *funcx.sdk.config*), 48

## M

make\_callback () (*funcx.strategies.base.BaseStrategy*  
*method*), 51

make\_callback () (*funcx.strategies.base.Timer*  
*method*), 52

make\_callback () (*funcx.strategies.BaseStrategy*  
*method*), 53

MalformedResponse, 54

Manager (class in *funcx.executors.high\_throughput.funcx\_manager*),  
 26

ManagerLost, 29

map\_run () (*funcx.FuncXClient* method), 14

map\_run () (*funcx.sdk.client.FuncXClient* method), 46

MaxRequestsExceeded, 44

MaxRequestSizeExceeded, 44

migrate\_tasks\_to\_internal ()  
 (*funcx.executors.high\_throughput.interchange.Interchange*  
*method*), 29

## N

naive\_interchange\_task\_dispatch () (in  
 module *funcx.executors.high\_throughput.interchange\_task\_dispatch*)  
 30

naive\_scheduler() (in module `funcx.executors.high_throughput.container_sched`), 21  
 NotConnected, 43  
 notify() (*funcx.strategies.base.BaseStrategy* method), 51  
 notify() (*funcx.strategies.BaseStrategy* method), 53  
**O**  
 outstanding (*funcx.executors.high\_throughput.executor.HighThroughputExecutor* attribute), 24  
 outstanding (*funcx.executors.HighThroughputExecutor* attribute), 35  
**P**  
 pack\_buffers() (*funcx.serialize.facade.FuncXSerializer* method), 50  
 pickle\_base64 (class in *funcx.serialize.concretes*), 50  
 provider\_status() (*funcx.executors.high\_throughput.interchange.Interchange* method), 29  
 pull\_tasks() (*funcx.executors.high\_throughput.funcx\_manager.Manager* method), 26  
 push\_results() (*funcx.executors.high\_throughput.funcx\_manager.Manager* method), 27  
 put() (*funcx.executors.high\_throughput.zmq\_pipes.TasksOutgoing* method), 32  
 put() (*funcx.queues.base.FuncxQueue* method), 43  
 put() (*funcx.queues.redis.redis\_q.RedisQueue* method), 42  
 put() (*funcx.queues.RedisQueue* method), 43  
 put\_worker() (*funcx.executors.high\_throughput.worker\_map.WorkerMap* method), 31  
**R**  
 ready\_worker\_count() (*funcx.executors.high\_throughput.worker\_map.WorkerMap* method), 31  
 reconnect (*funcx.endpoint.utils.zmq\_worker.ZMQWorker* attribute), 18  
 reconnect\_to\_broker() (*funcx.endpoint.utils.zmq\_worker.ZMQWorker* method), 18  
 recv() (*funcx.endpoint.utils.zmq\_worker.ZMQWorker* method), 18  
 RedisQueue (class in *funcx.queues*), 43  
 RedisQueue (class in *funcx.queues.redis.redis\_q*), 42  
 register() (in module *funcx.mock\_broker.mock\_broker*), 37  
 register\_container() (*funcx.FuncXClient* method), 14  
 register\_container() (*funcx.sdk.client.FuncXClient* method), 47  
 register\_endpoint() (*funcx.FuncXClient* method), 15  
 register\_endpoint() (*funcx.sdk.client.FuncXClient* method), 47  
 register\_endpoint() (in module *funcx.endpoint.endpoint*), 20  
 register\_function() (*funcx.FuncXClient* method), 15  
 register\_function() (*funcx.sdk.client.FuncXClient* method), 47  
 register\_with\_hub() (in module *funcx.endpoint.endpoint*), 20  
 register\_worker() (*funcx.executors.high\_throughput.worker\_map.WorkerMap* method), 31  
 registration\_message() (*funcx.executors.high\_throughput.funcx\_worker.FuncXWorker* method), 27  
 RegistrationError, 55  
 RemoteExceptionWrapper (class in *funcx.serialize.base*), 49  
 remove\_option() (in module *funcx.sdk.config*), 48  
 remove\_worker() (*funcx.executors.high\_throughput.worker\_map.WorkerMap* method), 31  
 remove\_worker\_init() (*funcx.executors.high\_throughput.funcx\_manager.Manager* method), 27  
 reply\_to (*funcx.endpoint.utils.zmq\_worker.ZMQWorker* attribute), 18  
 request\_close() (*funcx.executors.high\_throughput.zmq\_pipes.ResultsIncoming* method), 32  
 request\_status\_info() (*funcx.executors.high\_throughput.executor.HighThroughputExecutor* method), 24  
 request\_status\_info() (*funcx.executors.HighThroughputExecutor* method), 35  
 reraise() (*funcx.serialize.base.RemoteExceptionWrapper* method), 49  
 ResultsIncoming (class in *funcx.executors.high\_throughput.zmq\_pipes*), 32  
 run() (*funcx.executors.high\_throughput.zmq\_pipes.CommandClient* method), 32  
 run() (*funcx.FuncXClient* method), 15  
 run() (*funcx.mock\_broker.forwarder.Forwarder* method), 36  
 run() (*funcx.sdk.client.FuncXClient* method), 47  
 running() (*funcx.sdk.utils.futures.FuncXFuture* method), 44  
**S**  
 safeprint() (in module *funcx.sdk.config*), 48

*scale\_in()* (*funcx.executors.high\_throughput.executor.HighThroughputExecutor* method), 24  
*scale\_in()* (*funcx.executors.high\_throughput.worker\_map.WorkerMap* method), 24  
*scale\_in()* (*funcx.executors.high\_throughput.interchange.Interchange* method), 29  
*scale\_in()* (*funcx.executors.HighThroughputExecutor* method), 35  
*scale\_out()* (*funcx.executors.high\_throughput.executor.HighThroughputExecutor* method), 24  
*scale\_out()* (*funcx.executors.high\_throughput.interchange.Interchange* method), 27  
*scale\_out()* (*funcx.executors.HighThroughputExecutor* method), 35  
*scaling\_enabled()* (*funcx.executors.high\_throughput.executor.HighThroughputExecutor* attribute), 25  
*scaling\_enabled()* (*funcx.executors.HighThroughputExecutor* attribute), 35  
*send()* (*funcx.endpoint.utils.zmq\_worker.ZMQWorker* method), 18  
*send\_to\_broker()* (*funcx.endpoint.utils.zmq\_worker.ZMQWorker* method), 18  
*serialize()* (*funcx.serialize.base.fxPicker\_enforcer* method), 49  
*serialize()* (*funcx.serialize.concretes.code\_pickle* method), 49  
*serialize()* (*funcx.serialize.concretes.code\_text\_dill* method), 49  
*serialize()* (*funcx.serialize.concretes.code\_text\_inspect* method), 50  
*serialize()* (*funcx.serialize.concretes.json\_base64* method), 50  
*serialize()* (*funcx.serialize.concretes.pickle\_base64* method), 50  
*serialize()* (*funcx.serialize.facade.FuncXSerializer* method), 50  
*SerializerError* (class in *funcx.serialize.base*), 49  
*service* (*funcx.endpoint.utils.zmq\_worker.ZMQWorker* attribute), 18  
*set\_file\_logger()* (in module *funcx*), 55  
*set\_id()* (in module *funcx.endpoint.utils.zhelpers*), 18  
*set\_stream\_logger()* (in module *funcx*), 55  
*shutdown()* (*funcx.executors.high\_throughput.executor.HighThroughputExecutor* method), 25  
*shutdown()* (*funcx.executors.HighThroughputExecutor* method), 35  
*ShutdownRequest*, 29  
*SimpleStrategy* (class in *funcx.strategies*), 53  
*SimpleStrategy* (class in *funcx.strategies.simple*), 52  
*socket\_set\_hwm()* (in module *funcx.endpoint.utils.zhelpers*), 18  
*spawn\_forwarder()* (in module *funcx.mock\_broker.forwarder*), 36  
*spin\_down\_workers()* (*funcx.executors.high\_throughput.worker\_map.WorkerMap* method), 31  
*start()* (*funcx.executors.high\_throughput.executor.HighThroughputExecutor* method), 25  
*start()* (*funcx.executors.high\_throughput.funcx\_manager.Manager* method), 27  
*start()* (*funcx.executors.high\_throughput.funcx\_worker.FuncXWorker* method), 27  
*start()* (*funcx.executors.high\_throughput.interchange.Interchange* method), 29  
*start()* (*funcx.executors.HighThroughputExecutor* method), 51  
*start()* (*funcx.strategies.base.BaseStrategy* method), 53  
*start()* (*funcx.strategies.BaseStrategy* method), 53  
*start\_endpoint()* (in module *funcx.endpoint.endpoint*), 20  
*start\_file\_logger()* (in module *funcx.executors.high\_throughput.interchange*), 29  
*starter()* (in module *funcx.executors.high\_throughput.interchange*), 30  
*status()* (*funcx.executors.high\_throughput.executor.HighThroughputExecutor* method), 25  
*status()* (*funcx.executors.HighThroughputExecutor* method), 35  
*status()* (*funcx.providers.kubernetes.kube.KubernetesProvider* method), 39  
*status()* (*funcx.providers.KubernetesProvider* method), 41  
*stop()* (*funcx.executors.high\_throughput.interchange.Interchange* method), 29  
*stop\_endpoint()* (in module *funcx.endpoint.endpoint*), 20  
*strategize()* (*funcx.strategies.base.BaseStrategy* method), 51  
*strategize()* (*funcx.strategies.BaseStrategy* method), 53  
*strategize()* (*funcx.strategies.kube\_simple.KubeSimpleStrategy* method), 52  
*strategize()* (*funcx.strategies.KubeSimpleStrategy* method), 53  
*strategize()* (*funcx.strategies.simple.SimpleStrategy* method), 52  
*strategize()* (*funcx.strategies.SimpleStrategy* method), 53  
*submit()* (*funcx.executors.high\_throughput.executor.HighThroughputExecutor* method), 25  
*submit()* (*funcx.executors.HighThroughputExecutor* method), 35  
*submit()* (*funcx.providers.kubernetes.kube.KubernetesProvider* method), 39



submit() (*funcx.providers.KubernetesProvider method*), 41

## T

TasksOutgoing (class in *funcx.executors.high\_throughput.zmq\_pipes*), 32

test() (in module *funcx.mock\_broker.mock\_tester*), 37

test() (in module *funcx.queues.redis.redis\_g*), 42

test\_1() (in module *funcx.mock\_broker.test*), 37

test\_2() (in module *funcx.mock\_broker.test*), 37

test\_3() (in module *funcx.mock\_broker.test*), 37

throttle\_max\_requests() (*funcx.sdk.utils.throttling.ThrottledBaseClient method*), 44

throttle\_request\_size() (*funcx.sdk.utils.throttling.ThrottledBaseClient method*), 44

ThrottledBaseClient (class in *funcx.sdk.utils.throttling*), 44

ThrottlingException, 44

timeout (*funcx.endpoint.utils.zmq\_worker.ZMQWorker attribute*), 18

Timer (class in *funcx.strategies.base*), 52

TOKEN\_DIR (*funcx.sdk.client.FuncXClient attribute*), 45

TOKEN\_FILENAME (*funcx.sdk.client.FuncXClient attribute*), 45

## U

unpack\_and\_deserialize() (*funcx.serialize.facade.FuncXSerializer method*), 50

unpack\_buffers() (*funcx.serialize.facade.FuncXSerializer method*), 50

update\_table() (*funcx.FuncXClient method*), 16

update\_table() (*funcx.sdk.client.FuncXClient method*), 48

update\_worker\_idle() (*funcx.executors.high\_throughput.worker\_map.WorkerMap method*), 32

## V

verbose (*funcx.endpoint.utils.zmq\_worker.ZMQWorker attribute*), 18

## W

wait\_for\_endpoint() (*funcx.executors.high\_throughput.executor.HighThroughputExecutor method*), 25

wait\_for\_endpoint() (*funcx.executors.HighThroughputExecutor method*), 36

weakref\_cb() (*funcx.executors.high\_throughput.executor.HighThroughputExecutor method*), 25

weakref\_cb() (*funcx.executors.HighThroughputExecutor method*), 36

worker (*funcx.endpoint.utils.zmq\_worker.ZMQWorker attribute*), 18

WorkerMap (class in *funcx.executors.high\_throughput.worker\_map*), 30

write\_option() (in module *funcx.sdk.config*), 48

## Z

ZMQWorker (class in *funcx.endpoint.utils.zmq\_worker*), 18

zpipe() (in module *funcx.endpoint.utils.zhelpers*), 18