
Toil Documentation

Release 5.3.0

UCSC Computational Genomics Lab

Mar 26, 2021

1	Installation	3
1.1	Preparing Your Python Runtime Environment	3
1.2	Basic Installation	4
1.3	Installing Toil with Extra Features	4
1.4	Building from Source	5
2	Quickstart Examples	7
2.1	Running a basic workflow	7
2.2	Running a basic CWL workflow	8
2.3	Running a basic WDL workflow	9
2.4	A (more) real-world example	9
2.5	Launching a Toil Workflow in AWS	16
2.6	Running a CWL Workflow on AWS	17
2.7	Running a Workflow with Autoscaling - Cactus	18
3	Introduction	21
3.1	Job Store	21
3.2	Batch System	22
3.3	Provisioner	22
4	Commandline Options	23
4.1	The Job Store	23
4.2	Commandline Options	23
4.3	Restart Option	28
4.4	Running Workflows with Services	28
4.5	Setting Options directly with the Toil Script	28
5	Toil Debugging	31
5.1	Introspecting the Jobstore	31
5.2	Stats and Status	31
5.3	Using a Python debugger	32
6	Running in the Cloud	33
6.1	Managing a Cluster of Virtual Machines (Provisioning)	33
6.2	Storage (Toil jobStore)	33
7	Cloud Platforms	35

7.1	Running on Kubernetes	35
7.2	Running in AWS	43
7.3	Running in Google Compute Engine (GCE)	49
7.4	Cluster Utilities	53
7.5	Stats Command	54
7.6	Status Command	56
7.7	Clean Command	56
7.8	Launch-Cluster Command	57
7.9	Ssh-Cluster Command	58
7.10	Rsync-Cluster Command	59
7.11	Destroy-Cluster Command	59
7.12	Kill Command	59
8	HPC Environments	61
8.1	Standard Output/Error from Batch System Jobs	61
9	CWL in Toil	63
9.1	Running CWL Locally	63
9.2	Detailed Usage Instructions	63
9.3	Running CWL in the Cloud	64
9.4	Running CWL within Toil Scripts	65
9.5	Toil & CWL Tips	66
10	WDL in Toil	71
10.1	How to Run a WDL file in Toil	71
10.2	ENCODE Example from ENCODE-DCC	71
10.3	GATK Examples from the Broad	72
10.4	toilwdl.py Options	73
10.5	Running WDL within Toil Scripts	73
10.6	WDL Specifications	74
11	Developing a Workflow	75
11.1	Scripting Quick Start	75
11.2	Job Basics	76
11.3	Invoking a Workflow	76
11.4	Specifying Commandline Arguments	77
11.5	Resuming a Workflow	78
11.6	Functions and Job Functions	78
11.7	Workflows with Multiple Jobs	79
11.8	Dynamic Job Creation	81
11.9	Promises	82
11.10	Promised Requirements	83
11.11	FileID	84
11.12	Managing files within a workflow	84
11.13	Using Docker Containers in Toil	87
11.14	Services	88
11.15	Checkpoints	90
11.16	Encapsulation	90
11.17	Depending on Toil	91
11.18	Best Practices for Dockerizing Toil Workflows	91
12	Toil Class API	93
13	Job Store API	95

14	Toil Job API	97
14.1	FunctionWrappingJob	97
14.2	JobFunctionWrappingJob	97
14.3	EncapsulatedJob	97
14.4	Promise	97
15	Job Methods API	99
16	Job.Runner API	101
17	job.fileStore API	103
18	Batch System API	105
18.1	Batch System Environmental Variables	105
18.2	Batch System API	106
19	Job.Service API	107
20	Exceptions API	109
21	Running Tests	111
21.1	Running Tests with pytest	112
21.2	Running Integration Tests	112
21.3	Test Environment Variables	112
21.4	Using Docker with Quay	113
21.5	Running Mesos Tests	113
22	Developing with Docker	115
22.1	Making Your Own Toil Docker Image	115
22.2	Running a Cluster Locally	116
23	Maintainer’s Guidelines	119
23.1	Naming Conventions	119
23.2	Pull Requests	120
23.3	Publishing a Release	120
23.4	Adding Retries to a Function	121
24	Pull Request Checklists	125
24.1	Reviewing Pull Requests	125
24.2	Merging Pull Requests	126
25	Toil Architecture	127
25.1	Optimizations	129
25.2	Toil support for Common Workflow Language	130
26	Minimum AWS IAM permissions	133
27	Auto-Deployment	135
27.1	Auto Deployment with Sibling Modules	136
27.2	Auto-Deploying a Package Hierarchy	137
27.3	Relying on Shared Filesystems	138
28	Environment Variables	139
	Index	143

Toil is an open-source pure-Python workflow engine that lets people write better pipelines.

Check out our [website](#) for a comprehensive list of Toil's features and read our [paper](#) to learn what Toil can do in the real world. Please subscribe to our low-volume [announce](#) mailing list and feel free to also join us on [GitHub](#) and [Gitter](#).

If using Toil for your research, please cite

Vivian, J., Rao, A. A., Nothhaft, F. A., Ketchum, C., Armstrong, J., Novak, A., ... Paten, B. (2017). Toil enables reproducible, open source, big biomedical data analyses. *Nature Biotechnology*, 35(4), 314–316. <http://doi.org/10.1038/nbt.3772>

This document describes how to prepare for and install Toil. Note that Toil requires that the user run all commands inside of a Python `virtualenv`. Instructions for installing and creating a Python virtual environment are provided below.

1.1 Preparing Your Python Runtime Environment

Toil currently supports Python 2.7, 3.5, and 3.6, and requires a `virtualenv` to be active to install.

If not already present, please install the latest Python `virtualenv` using `pip`:

```
$ sudo pip install virtualenv
```

And create a virtual environment called `venv` in your home directory:

```
$ virtualenv ~/venv
```

If the user does not have root privileges, there are a few more steps, but one can download a specific `virtualenv` package directly, untar the file, create, and source the `virtualenv` (version 15.1.0 as an example) using

```
$ curl -O https://pypi.python.org/packages/d4/0c/  
→ 9840c08189e030873387a73b90ada981885010dd9aea134d6de30cd24cb8/virtualenv-15.1.0.tar.  
→ gz  
$ tar xvfz virtualenv-15.1.0.tar.gz  
$ cd virtualenv-15.1.0  
$ python virtualenv.py ~/venv
```

Now that you've created your `virtualenv`, activate your virtual environment:

```
$ source ~/venv/bin/activate
```

1.2 Basic Installation

If you need only the basic version of Toil, it can be easily installed using pip:

```
$ pip install toil
```

Now you're ready to run *your first Toil workflow!*

(If you need any of the extra features don't do this yet and instead skip to the next section.)

1.3 Installing Toil with Extra Features

Python headers and static libraries

Needed for the `mesos`, `aws`, `google`, and `encryption` extras.

On Ubuntu:

```
$ sudo apt-get install build-essential python-dev
```

On macOS:

```
$ xcode-select --install
```

Encryption specific headers and library

Needed for the `encryption` extra.

On Ubuntu:

```
$ sudo apt-get install libssl-dev libffi-dev
```

On macOS:

```
$ brew install libssl libffi
```

Or see [Cryptography](#) for other systems.

Some optional features, called *extras*, are not included in the basic installation of Toil. To install Toil with all its bells and whistles, first install any necessary headers and libraries (*python-dev*, *libffi-dev*). Then run

```
$ pip install toil[aws,mesos,google,encryption,cwl]
```

or

```
$ pip install toil[all]
```

Here's what each extra provides:

Extra	Description
all	Installs all extras (though htcondor is linux-only and will be skipped if not on a linux computer).
aws	Provides support for managing a cluster on Amazon Web Service (AWS) using Toil's built in <i>Cluster Utilities</i> . Clusters can scale up and down automatically. It also supports storing workflow state.
google	Experimental. Stores workflow state in <i>Google Cloud Storage</i> .
mesos	<p>Provides support for running Toil on an <i>Apache Mesos</i> cluster. Note that running Toil on other batch systems does not require an extra. The <code>mesos</code> extra requires the following native dependencies:</p> <ul style="list-style-type: none"> • <i>Apache Mesos</i> (Tested with Mesos v1.0.0) • <i>Python headers and static libraries</i> <hr/> <p>Important: If launching toil remotely on a mesos instance, to install Toil with the <code>mesos</code> extra in a virtualenv, be sure to create that virtualenv with the <code>--system-site-packages</code> flag (only use remotely!):</p> <pre>\$ virtualenv ~/venv --system-site-packages</pre> <p>Otherwise, you'll see something like this:</p> <pre>ImportError: No module named mesos.native</pre> <hr/>
htcondor	Support for the htcondor batch system. This currently is a linux only extra.
encryption	<p>Provides client-side encryption for files stored in the AWS job store. This extra requires the following native dependencies:</p> <ul style="list-style-type: none"> • <i>Python headers and static libraries</i> • <i>libffi headers and library</i>
cwl	Provides support for running workflows written using the <i>Common Workflow Language</i> .
wdl	Provides support for running workflows written using the <i>Workflow Description Language</i> . This extra has no native dependencies.

1.4 Building from Source

If developing with Toil, you will need to build from source. This allows changes you make to Toil to be reflected immediately in your runtime environment.

First, clone the source:

```
$ git clone https://github.com/DataBiosphere/toil.git
$ cd toil
```

Then, create and activate a virtualenv:

```
$ virtualenv venv
$ . venv/bin/activate
```

From there, you can list all available Make targets by running `make`. First and foremost, we want to install Toil's build requirements (these are additional packages that Toil needs to be tested and built but not to be run):

```
$ make prepare
```

Now, we can install Toil in development mode (such that changes to the source code will immediately affect the virtualenv):

```
$ make develop
```

Or, to install with support for all optional *Installing Toil with Extra Features*:

```
$ make develop extras=[aws,mesos,google,encryption,cwl]
```

Or:

```
$ make develop extras=[all]
```

To build the docs, run `make develop` with all extras followed by

```
$ make docs
```

To run a quick batch of tests (this should take less than 30 minutes) run

```
$ export TOIL_TEST_QUICK=True; make test
```

For more information on testing see *Running Tests*.

2.1 Running a basic workflow

A Toil workflow can be run with just three steps:

1. Install Toil (see *Installation*)
2. Copy and paste the following code block into a new file called `helloWorld.py`:

```
from toil.common import Toil
from toil.job import Job

def helloWorld(message, memory="1G", cores=1, disk="1G"):
    return "Hello, world!, here's a message: %s" % message

if __name__ == "__main__":
    parser = Job.Runner.getDefaultArgumentParser()
    options = parser.parse_args()
    options.clean = "always"
    with Toil(options) as toil:
        output = toil.start(Job.wrapFn(helloWorld, "You did it!"))
    print(output)
```

3. Specify the name of the *job store* and run the workflow:

```
(venv) $ python helloWorld.py file:my-job-store
```

Note: Don't actually type `(venv) $` in at the beginning of each command. This is intended only to remind the user that they should have their *virtual environment* running.

Congratulations! You've run your first Toil workflow using the default *Batch System*, `singleMachine`, using the `file` job store.

Toil uses batch systems to manage the jobs it creates.

The `singleMachine` batch system is primarily used to prepare and debug workflows on a local machine. Once validated, try running them on a full-fledged batch system (see [Batch System API](#)). Toil supports many different batch systems such as [Apache Mesos](#) and [Grid Engine](#); its versatility makes it easy to run your workflow in all kinds of places.

Toil is totally customizable! Run `python helloWorld.py --help` to see a complete list of available options.

For something beyond a “Hello, world!” example, refer to [A \(more\) real-world example](#).

2.2 Running a basic CWL workflow

The [Common Workflow Language](#) (CWL) is an emerging standard for writing workflows that are portable across multiple workflow engines and platforms. Running CWL workflows using Toil is easy.

1. First ensure that Toil is installed with the `cwl` extra (see [Installing Toil with Extra Features](#)):

```
(venv) $ pip install 'toil[cwl]'
```

This installs the `toil-cwl-runner` executable.

2. Copy and paste the following code block into `example.cwl`:

```
cwlVersion: v1.0
class: CommandLineTool
baseCommand: echo
stdout: output.txt
inputs:
  message:
    type: string
    inputBinding:
      position: 1
outputs:
  output:
    type: stdout
```

and this code into `example-job.yaml`:

```
message: Hello world!
```

3. To run the workflow simply enter

```
(venv) $ toil-cwl-runner example.cwl example-job.yaml
```

Your output will be in `output.txt`:

```
(venv) $ cat output.txt
Hello world!
```

To learn more about CWL, see the [CWL User Guide](#) (from where this example was shamelessly borrowed).

To run this workflow on an AWS cluster have a look at [Running a CWL Workflow on AWS](#).

For information on using CWL with Toil see the section [CWL in Toil](#)

2.3 Running a basic WDL workflow

The [Workflow Description Language](#) (WDL) is another emerging language for writing workflows that are portable across multiple workflow engines and platforms. Running WDL workflows using Toil is still in alpha, and currently experimental. Toil currently supports basic workflow syntax (see [WDL in Toil](#) for more details and examples). Here we go over running a basic WDL helloworld workflow.

1. First ensure that Toil is installed with the wdl extra (see [Installing Toil with Extra Features](#)):

```
(venv) $ pip install 'toil[wdl]'
```

This installs the `toil-wdl-runner` executable.

2. Copy and paste the following code block into `wdl-helloworld.wdl`:

```
workflow write_simple_file {
  call write_file
}
task write_file {
  String message
  command { echo ${message} > wdl-helloworld-output.txt }
  output { File test = "wdl-helloworld-output.txt" }
}

and this code into ``wdl-helloworld.json``::

{
  "write_simple_file.write_file.message": "Hello world!"
}
```

3. To run the workflow simply enter

```
(venv) $ toil-wdl-runner wdl-helloworld.wdl wdl-helloworld.json
```

Your output will be in `wdl-helloworld-output.txt`:

```
(venv) $ cat wdl-helloworld-output.txt
Hello world!
```

To learn more about WDL, see the main [WDL website](#).

2.4 A (more) real-world example

For a more detailed example and explanation, we've developed a sample pipeline that merge-sorts a temporary file. This is not supposed to be an efficient sorting program, rather a more fully worked example of what Toil is capable of.

2.4.1 Running the example

1. Download the example code
2. Run it with the default settings:

```
(venv) $ python sort.py file:jobStore
```

The workflow created a file called `sortedFile.txt` in your current directory. Have a look at it and notice that it contains a whole lot of sorted lines!

This workflow does a smart merge sort on a file it generates, `fileToSort.txt`. The sort is *smart* because each step of the process—splitting the file into separate chunks, sorting these chunks, and merging them back together—is compartmentalized into a **job**. Each job can specify its own resource requirements and will only be run after the jobs it depends upon have run. Jobs without dependencies will be run in parallel.

Note: Delete `fileToSort.txt` before moving on to #3. This example introduces options that specify dimensions for `fileToSort.txt`, if it does not already exist. If it exists, this workflow will use the existing file and the results will be the same as #2.

3. Run with custom options:

```
(venv) $ python sort.py file:jobStore \
      --numLines=5000 \
      --lineLength=10 \
      --overwriteOutput=True \
      --workDir=/tmp/
```

Here we see that we can add our own options to a Toil script. As noted above, the first two options, `--numLines` and `--lineLength`, determine the number of lines and how many characters are in each line. `--overwriteOutput` causes the current contents of `sortedFile.txt` to be overwritten, if it already exists. The last option, `--workDir`, is an option built into Toil to specify where temporary files unique to a job are kept.

2.4.2 Describing the source code

To understand the details of what's going on inside. Let's start with the `main()` function. It looks like a lot of code, but don't worry—we'll break it down piece by piece.

```
def main(options=None):
    if not options:
        # deal with command line arguments
        parser = ArgumentParser()
        Job.Runner.addToilOptions(parser)
        parser.add_argument('--numLines', default=defaultLines, help='Number of lines_
↳in file to sort.', type=int)
        parser.add_argument('--lineLength', default=defaultLineLen, help='Length of_
↳lines in file to sort.', type=int)
        parser.add_argument("--fileToSort", help="The file you wish to sort")
        parser.add_argument("--outputFile", help="Where the sorted output will go")
        parser.add_argument("--overwriteOutput", help="Write over the output file if_
↳it already exists.", default=True)
        parser.add_argument("--N", dest="N",
                            help="The threshold below which a serial sort function is_
↳used to sort file. "
                            "All lines must of length less than or equal to N or_
↳program will fail",
                            default=10000)
        parser.add_argument('--downCheckpoints', action='store_true',
                            help='If this option is set, the workflow will make_
↳checkpoints on its way through'
                            'the recursive "down" part of the sort')
        parser.add_argument("--sortMemory", dest="sortMemory",
```

(continues on next page)

(continued from previous page)

```

        help="Memory for jobs that sort chunks of the file.",
        default=None)

    parser.add_argument("--mergeMemory", dest="mergeMemory",
                        help="Memory for jobs that collate results.",
                        default=None)

    options = parser.parse_args()
    if not hasattr(options, "sortMemory") or not options.sortMemory:
        options.sortMemory = sortMemory
    if not hasattr(options, "mergeMemory") or not options.mergeMemory:
        options.mergeMemory = sortMemory

    # do some input verification
    sortedFileName = options.outputFile or "sortedFile.txt"
    if not options.overwriteOutput and os.path.exists(sortedFileName):
        print(f'Output file {sortedFileName} already exists. '
              f'Delete it to run the sort example again or use --overwriteOutput=True
→')
    exit()

    fileName = options.fileToSort
    if options.fileToSort is None:
        # make the file ourselves
        fileName = 'fileToSort.txt'
        if os.path.exists(fileName):
            print(f'Sorting existing file: {fileName}')
        else:
            print(f'No sort file specified. Generating one automatically called:
→{fileName}.')
            makeFileToSort(fileName=fileName, lines=options.numLines, lineLen=options.
→lineLength)
    else:
        if not os.path.exists(options.fileToSort):
            raise RuntimeError("File to sort does not exist: %s" % options.fileToSort)

    if int(options.N) <= 0:
        raise RuntimeError("Invalid value of N: %s" % options.N)

    # Now we are ready to run
    with Toil(options) as workflow:
        sortedFileURL = 'file://' + os.path.abspath(sortedFileName)
        if not workflow.options.restart:
            sortFileURL = 'file://' + os.path.abspath(fileName)
            sortFileID = workflow.importFile(sortFileURL)
            sortedFileID = workflow.start(Job.wrapJobFn(setup,
                                                         sortFileID,
                                                         int(options.N),
                                                         options.downCheckpoints,
                                                         options=options,
                                                         memory=sortMemory))
        else:
            sortedFileID = workflow.restart()
    workflow.exportFile(sortedFileID, sortedFileURL)

```

First we make a parser to process command line arguments using the `argparse` module. It's important that we add the call to `Job.Runner.addToilOptions()` to initialize our parser with all of Toil's default options. Then we add

the command line arguments unique to this workflow, and parse the input. The help message listed with the arguments should give you a pretty good idea of what they can do.

Next we do a little bit of verification of the input arguments. The option `--fileToSort` allows you to specify a file that needs to be sorted. If this option isn't given, it's here that we make our own file with the call to `makeFileToSort()`.

Finally we come to the context manager that initializes the workflow. We create a path to the input file prepended with `'file:/'` as per the documentation for `toil.common.Toil()` when staging a file that is stored locally. Notice that we have to check whether or not the workflow is restarting so that we don't import the file more than once. Finally we can kick off the workflow by calling `toil.common.Toil.start()` on the job setup. When the workflow ends we capture its output (the sorted file's `fileID`) and use that in `toil.common.Toil.exportFile()` to move the sorted file from the job store back into "userland".

Next let's look at the job that begins the actual workflow, `setup`.

```
def setup(job, inputFile, N, downCheckpoints, options):
    """
    Sets up the sort.
    Returns the FileID of the sorted file
    """
    RealtimeLogger.info("Starting the merge sort")
    return job.addChildJobFn(down,
                             inputFile, N, 'root',
                             downCheckpoints,
                             options = options,
                             preemptable=True,
                             memory=sortMemory).rv()
```

`setup` really only does two things. First it writes to the logs using `Job.log()` and then calls `addChildJobFn()`. Child jobs run directly after the current job. This function turns the 'job function' `down` into an actual job and passes in the inputs including an optional resource requirement, `memory`. The job doesn't actually get run until the call to `Job.rv()`. Once the job `down` finishes, its output is returned here.

Now we can look at what `down` does.

```
def down(job, inputFileStoreID, N, path, downCheckpoints, options, memory=sortMemory):
    """
    Input is a file, a subdivision size N, and a path in the hierarchy of jobs.
    If the range is larger than a threshold N the range is divided recursively and
    a follow on job is then created which merges back the results else
    the file is sorted and placed in the output.
    """

    RealtimeLogger.info("Down job starting: %s" % path)

    # Read the file
    inputFile = job.fileStore.readGlobalFile(inputFileStoreID, cache=False)
    length = os.path.getsize(inputFile)
    if length > N:
        # We will subdivide the file
        RealtimeLogger.critical("Splitting file: %s of size: %s"
                                % (inputFileStoreID, length))
        # Split the file into two copies
        midPoint = getMidPoint(inputFile, 0, length)
        t1 = job.fileStore.getLocalTempFile()
        with open(t1, 'w') as fH:
            fH.write(copySubRangeOfFile(inputFile, 0, midPoint+1))
```

(continues on next page)

(continued from previous page)

```

    t2 = job.fileStore.getLocalTempFile()
    with open(t2, 'w') as fh:
        fh.write(copySubRangeOfFile(inputFile, midPoint+1, length))
        # Call down recursively. By giving the rv() of the two jobs as inputs to the
        ↪follow-on job, up,
        # we communicate the dependency without hindering concurrency.
        result = job.addFollowOnJobFn(up,
                                     job.addChildJobFn(down, job.fileStore.
                                     ↪writeGlobalFile(t1), N, path + '/0',
                                                         downCheckpoints,
                                     ↪checkpoint=downCheckpoints, options=options,
                                                         preemptable=True,
                                     ↪memory=options.sortMemory).rv(),
                                     job.addChildJobFn(down, job.fileStore.
                                     ↪writeGlobalFile(t2), N, path + '/1',
                                                         downCheckpoints,
                                     ↪checkpoint=downCheckpoints, options=options,
                                                         preemptable=True,
                                     ↪memory=options.mergeMemory).rv(),
                                     path + '/up', preemptable=True, options=options,
                                     ↪memory=options.sortMemory).rv())
    else:
        # We can sort this bit of the file
        RealtimeLogger.critical("Sorting file: %s of size: %s"
                               % (inputFileStoreID, length))
        # Sort the copy and write back to the fileStore
        shutil.copyfile(inputFile, inputFile + '.sort')
        sort(inputFile + '.sort')
        result = job.fileStore.writeGlobalFile(inputFile + '.sort')

    RealtimeLogger.info("Down job finished: %s" % path)
    return result

```

Down is the recursive part of the workflow. First we read the file into the local fileStore by calling `job.fileStore.readGlobalFile()`. This puts a copy of the file in the temp directory for this particular job. This storage will disappear once this job ends. For a detailed explanation of the fileStore, job store, and their interfaces have a look at [Managing files within a workflow](#).

Next down checks the base case of the recursion: is the length of the input file less than N (remember N was an option we added to the workflow in main)? In the base case, we just sort the file, and return the file ID of this new sorted file.

If the base case fails, then the file is split into two new tempFiles using `job.fileStore.getLocalTempFile()` and the helper function `copySubRangeOfFile`. Finally we add a follow on Job up with `job.addFollowOnJobFn()`. We've already seen child jobs. A follow-on Job is a job that runs after the current job and *all* of its children (and their children and follow-ons) have completed. Using a follow-on makes sense because up is responsible for merging the files together and we don't want to merge the files together until we *know* they are sorted. Again, the return value of the follow-on job is requested using `Job.rv()`.

Looking at up

```

def up(job, inputFileID1, inputFileID2, path, options, memory=sortMemory):
    """
    Merges the two files and places them in the output.
    """

    RealtimeLogger.info("Up job starting: %s" % path)

```

(continues on next page)

(continued from previous page)

```

with job.fileStore.writeGlobalFileStream() as (fileHandle, outputFileStoreID):
    fileHandle = codecs.getwriter('utf-8')(fileHandle)
    with job.fileStore.readGlobalFileStream(inputFileID1) as inputFileHandle1:
        inputFileHandle1 = codecs.getreader('utf-8')(inputFileHandle1)
        with job.fileStore.readGlobalFileStream(inputFileID2) as inputFileHandle2:
            inputFileHandle2 = codecs.getreader('utf-8')(inputFileHandle2)
            RealtimeLogger.info("Merging %s and %s to %s"
                               % (inputFileID1, inputFileID2, outputFileStoreID))
            merge(inputFileHandle1, inputFileHandle2, fileHandle)
        # Cleanup up the input files - these deletes will occur after the completion_
    is successful.
    job.fileStore.deleteGlobalFile(inputFileID1)
    job.fileStore.deleteGlobalFile(inputFileID2)

    RealtimeLogger.info("Up job finished: %s" % path)

return outputFileStoreID

```

we see that the two input files are merged together and the output is written to a new file using `job.fileStore.writeGlobalFileStream()`. After a little cleanup, the output file is returned.

Once the final up finishes and all of the `rv()` promises are fulfilled, `main` receives the sorted file's ID which it uses in `exportFile` to send it to the user.

There are other things in this example that we didn't go over such as *Checkpoints* and the details of much of the *Toil Class API*.

At the end of the script the lines

```

if __name__ == '__main__':
    main()

```

are included to ensure that the main function is only run once in the `'__main__'` process invoked by you, the user. In Toil terms, by invoking the script you created the *leader process* in which the `main()` function is run. A *worker process* is a separate process whose sole purpose is to host the execution of one or more jobs defined in that script. In any Toil workflow there is always one leader process, and potentially many worker processes.

When using the single-machine batch system (the default), the worker processes will be running on the same machine as the leader process. With full-fledged batch systems like Mesos the worker processes will typically be started on separate machines. The boilerplate ensures that the pipeline is only started once—on the leader—but not when its job functions are imported and executed on the individual workers.

Typing `python sort.py --help` will show the complete list of arguments for the workflow which includes both Toil's and ones defined inside `sort.py`. A complete explanation of Toil's arguments can be found in *Commandline Options*.

2.4.3 Logging

By default, Toil logs a lot of information related to the current environment in addition to messages from the batch system and jobs. This can be configured with the `--logLevel` flag. For example, to only log CRITICAL level messages to the screen:

```

(venv) $ python sort.py file:jobStore \
        --logLevel=critical \
        --overwriteOutput=True

```

This hides most of the information we get from the Toil run. For more detail, we can run the pipeline with `--logLevel=debug` to see a comprehensive output. For more information, see [Commandline Options](#).

2.4.4 Error Handling and Resuming Pipelines

With Toil, you can recover gracefully from a bug in your pipeline without losing any progress from successfully completed jobs. To demonstrate this, let's add a bug to our example code to see how Toil handles a failure and how we can resume a pipeline after that happens. Add a bad assertion at line 52 of the example (the first line of `down()`):

```
def down(job, inputFileStoreID, N, downCheckpoints, memory=sortMemory):
    ...
    assert 1 == 2, "Test error!"
```

When we run the pipeline, Toil will show a detailed failure log with a traceback:

```
(venv) $ python sort.py file:jobStore
...
---TOIL WORKER OUTPUT LOG---
...
m/j/jobonrSMP      Traceback (most recent call last):
m/j/jobonrSMP      File "toil/src/toil/worker.py", line 340, in main
m/j/jobonrSMP      job._runner(jobGraph=jobGraph, jobStore=jobStore,
    ↪fileStore=fileStore)
m/j/jobonrSMP      File "toil/src/toil/job.py", line 1270, in _runner
m/j/jobonrSMP      returnValues = self._run(jobGraph, fileStore)
m/j/jobonrSMP      File "toil/src/toil/job.py", line 1217, in _run
m/j/jobonrSMP      return self.run(fileStore)
m/j/jobonrSMP      File "toil/src/toil/job.py", line 1383, in run
m/j/jobonrSMP      rValue = userFunction(*(self,) + tuple(self._args)), **self._
    ↪kwargs)
m/j/jobonrSMP      File "toil/example.py", line 30, in down
m/j/jobonrSMP      assert 1 == 2, "Test error!"
m/j/jobonrSMP      AssertionError: Test error!
```

If we try and run the pipeline again, Toil will give us an error message saying that a job store of the same name already exists. By default, in the event of a failure, the job store is preserved so that the workflow can be restarted, starting from the previously failed jobs. We can restart the pipeline by running

```
(venv) $ python sort.py file:jobStore \
        --restart \
        --overwriteOutput=True
```

We can also change the number of times Toil will attempt to retry a failed job:

```
(venv) $ python sort.py file:jobStore \
        --retryCount 2 \
        --restart \
        --overwriteOutput=True
```

You'll now see Toil attempt to rerun the failed job until it runs out of tries. `--retryCount` is useful for non-systemic errors, like downloading a file that may experience a sporadic interruption, or some other non-deterministic failure.

To successfully restart our pipeline, we can edit our script to comment out line 30, or remove it, and then run

```
(venv) $ python sort.py file:jobStore \
        --restart \
        --overwriteOutput=True
```

The pipeline will run successfully, and the job store will be removed on the pipeline's completion.

2.4.5 Collecting Statistics

Please see the *Stats Command* section for more on gathering runtime and resource info on jobs.

2.5 Launching a Toil Workflow in AWS

After having installed the `aws` extra for Toil during the *Installation* and set up AWS (see *Preparing your AWS environment*), the user can run the basic `helloWorld.py` script (*Running a basic workflow*) on a VM in AWS just by modifying the run command.

Note that when running in AWS, users can either run the workflow on a single instance or run it on a cluster (which is running across multiple containers on multiple AWS instances). For more information on running Toil workflows on a cluster, see *Running in AWS*.

Also! Remember to use the *Destroy-Cluster Command* command when finished to destroy the cluster! Otherwise things may not be cleaned up properly.

1. Launch a cluster in AWS using the *Launch-Cluster Command* command:

```
(venv) $ toil launch-cluster <cluster-name> \  
        --keyPairName <AWS-key-pair-name> \  
        --leaderNodeType t2.medium \  
        --zone us-west-2a
```

The arguments `keyPairName`, `leaderNodeType`, and `zone` are required to launch a cluster.

2. Copy `helloWorld.py` to the `/tmp` directory on the leader node using the *Rsync-Cluster Command* command:

```
(venv) $ toil rsync-cluster --zone us-west-2a <cluster-name> helloWorld.py :/tmp
```

Note that the command requires defining the file to copy as well as the target location on the cluster leader node.

3. Login to the cluster leader node using the *Ssh-Cluster Command* command:

```
(venv) $ toil ssh-cluster --zone us-west-2a <cluster-name>
```

Note that this command will log you in as the `root` user.

4. Run the Toil script in the cluster:

```
$ python /tmp/helloWorld.py aws:us-west-2:my-S3-bucket
```

In this particular case, we create an S3 bucket called `my-S3-bucket` in the `us-west-2` availability zone to store intermediate job results.

Along with some other INFO log messages, you should get the following output in your terminal window: `Hello, world!, here's a message: You did it!`.

5. Exit from the SSH connection.

```
$ exit
```

6. Use the *Destroy-Cluster Command* command to destroy the cluster:

```
(venv) $ toil destroy-cluster --zone us-west-2a <cluster-name>
```

Note that this command will destroy the cluster leader node and any resources created to run the job, including the S3 bucket.

2.6 Running a CWL Workflow on AWS

After having installed the `aws` and `cwl` extras for Toil during the *Installation* and set up AWS (see *Preparing your AWS environment*), the user can run a CWL workflow with Toil on AWS.

Also! Remember to use the *Destroy-Cluster Command* command when finished to destroy the cluster! Otherwise things may not be cleaned up properly.

1. First launch a node in AWS using the *Launch-Cluster Command* command:

```
(venv) $ toil launch-cluster <cluster-name> \
      --keyPairName <AWS-key-pair-name> \
      --leaderNodeType t2.medium \
      --zone us-west-2a
```

2. Copy `example.cwl` and `example-job.yaml` from the *CWL example* to the node using the *Rsync-Cluster Command* command:

```
(venv) $ toil rsync-cluster --zone us-west-2a <cluster-name> example.cwl :/tmp
(venv) $ toil rsync-cluster --zone us-west-2a <cluster-name> example-job.yaml :/
↪tmp
```

3. SSH into the cluster's leader node using the *Ssh-Cluster Command* utility:

```
(venv) $ toil ssh-cluster --zone us-west-2a <cluster-name>
```

4. Once on the leader node, it's a good idea to update and install the following:

```
sudo apt-get update
sudo apt-get -y upgrade
sudo apt-get -y dist-upgrade
sudo apt-get -y install git
sudo pip install mesos.cli
```

5. Now create a new virtualenv with the `--system-site-packages` option and activate:

```
virtualenv --system-site-packages venv
source venv/bin/activate
```

6. Now run the CWL workflow:

```
(venv) $ toil-cwl-runner \
      --provisioner aws \
      --jobStore aws:us-west-2a:any-name \
      /tmp/example.cwl /tmp/example-job.yaml
```

Tip: When running a CWL workflow on AWS, input files can be provided either on the local file system or in S3 buckets using `s3://` URI references. Final output files will be copied to the local file system of the leader node.

7. Finally, log out of the leader node and from your local computer, destroy the cluster:

```
(venv) $ toil destroy-cluster --zone us-west-2a <cluster-name>
```

2.7 Running a Workflow with Autoscaling - Cactus

Cactus is a reference-free, whole-genome multiple alignment program that can be run on any of the cloud platforms Toil supports.

Note: Cloud Independence:

This example provides a “cloud agnostic” view of running Cactus with Toil. Most options will not change between cloud providers. However, each provisioner has unique inputs for `--leaderNodeType`, `--nodeType` and `--zone`. We recommend the following:

Option	Used in	AWS	Google
<code>--leaderNodeType</code>	launch-cluster	t2.medium	n1-standard-1
<code>--zone</code>	launch-cluster	us-west-2a	us-west1-a
<code>--zone</code>	cactus	us-west-2	
<code>--nodeType</code>	cactus	c3.4xlarge	n1-standard-8

When executing `toil launch-cluster` with `gce` specified for `--provisioner`, the option `--boto` must be specified and given a path to your `.boto` file. See [Running in Google Compute Engine \(GCE\)](#) for more information about the `--boto` option.

Also! Remember to use the [Destroy-Cluster Command](#) when finished to destroy the cluster! Otherwise things may not be cleaned up properly.

1. Download `pestis.tar.gz`
2. Launch a leader node using the [Launch-Cluster Command](#) command:

```
(venv) $ toil launch-cluster <cluster-name> \  
      --provisioner <aws, gce> \  
      --keyPairName <key-pair-name> \  
      --leaderNodeType <type> \  
      --zone <zone>
```

Note: A Helpful Tip

When using AWS, setting the environment variable eliminates having to specify the `--zone` option for each command. This will be supported for GCE in the future.

```
(venv) $ export TOIL_AWS_ZONE=us-west-2c
```

3. Create appropriate directory for uploading files:

```
(venv) $ toil ssh-cluster --provisioner <aws, gce> <cluster-name>  
$ mkdir /root/cact_ex  
$ exit
```


- Copy the required files, i.e., seqFile.txt (a text file containing the locations of the input sequences as well as their phylogenetic tree, see [here](#)), organisms' genome sequence files in FASTA format, and configuration files (e.g. blockTrim1.xml, if desired), up to the leader node:

```
(venv) $ toil rsync-cluster --provisioner <aws, gce> <cluster-name> pestis-short-
↪aws-seqFile.txt :/root/cact_ex
(venv) $ toil rsync-cluster --provisioner <aws, gce> <cluster-name> GCF_000169655.
↪1_ASM16965v1_genomic.fna :/root/cact_ex
(venv) $ toil rsync-cluster --provisioner <aws, gce> <cluster-name> GCF_000006645.
↪1_ASM664v1_genomic.fna :/root/cact_ex
(venv) $ toil rsync-cluster --provisioner <aws, gce> <cluster-name> GCF_000182485.
↪1_ASM18248v1_genomic.fna :/root/cact_ex
(venv) $ toil rsync-cluster --provisioner <aws, gce> <cluster-name> GCF_000013805.
↪1_ASM1380v1_genomic.fna :/root/cact_ex
(venv) $ toil rsync-cluster --provisioner <aws, gce> <cluster-name> setup_
↪leaderNode.sh :/root/cact_ex
(venv) $ toil rsync-cluster --provisioner <aws, gce> <cluster-name> blockTrim1.
↪xml :/root/cact_ex
(venv) $ toil rsync-cluster --provisioner <aws, gce> <cluster-name> blockTrim3.
↪xml :/root/cact_ex
```

- Log in to the leader node:

```
(venv) $ toil ssh-cluster --provisioner <aws, gce> <cluster-name>
```

- Set up the environment of the leader node to run Cactus:

```
$ bash /root/cact_ex/setup_leaderNode.sh
$ source cact_venv/bin/activate
(cact_venv) $ cd cactus
(cact_venv) $ pip install --upgrade .
```

- Run **Cactus** as an autoscaling workflow:

```
(cact_venv) $ TOIL_APPLIANCE_SELF=quay.io/ucsc_cgl/toil:3.14.0 cactus \
--provisioner <aws, gce> \
--nodeType <type> \
--maxNodes 2 \
--minNodes 0 \
--retry 10 \
--batchSystem mesos \
--logDebug \
--logFile /logFile_pestis3 \
--configFile \
/root/cact_ex/blockTrim3.xml <aws, google>:<zone>:cactus-pestis_
↪\
/root/cact_ex/pestis-short-aws-seqFile.txt \
/root/cact_ex/pestis_output3.hal
```

Note: Pieces of the Puzzle:

TOIL_APPLIANCE_SELF=quay.io/ucsc_cgl/toil:3.14.0 — specifies the version of Toil being used, 3.14.0; if the latest one is desired, please eliminate.

--nodeType — determines the instance type used for worker nodes. The instance type specified here must be on the same cloud provider as the one specified with --leaderNodeType

--maxNodes 2 — creates up to two instances of the type specified with --nodeType and launches Mesos

worker containers inside them.

`--logDebug` — equivalent to `--logLevel DEBUG`.

`--logFile /logFile_pestis3` — writes logs in a file named *logFile_pestis3* under / folder.

`--configFile` — this is not required depending on whether a specific configuration file is intended to run the alignment.

`<aws, google>:<zone>:cactus-pestis` — creates a bucket, named `cactus-pestis`, with the specified cloud provider to store intermediate job files and metadata. **NOTE:** If you want to use a GCE-based jobstore, specify `google` here, not `gce`.

The result file, named `pestis_output3.hal`, is stored under `/root/cact_ex` folder of the leader node.

Use `cactus --help` to see all the Cactus and Toil flags available.

8. Log out of the leader node:

```
(cact_venv) $ exit
```

9. Download the resulted output to local machine:

```
(venv) $ toil rsync-cluster \  
      --provisioner <aws, gce> <cluster-name> \  
      :/root/cact_ex/pestis_output3.hal \  
      <path-of-folder-on-local-machine>
```

10. Destroy the cluster:

```
(venv) $ toil destroy-cluster --provisioner <aws, gce> <cluster-name>
```

Toil runs in various environments, including *locally* and *in the cloud* (Amazon Web Services and Google Compute Engine). Toil also supports two DSLs: *CWL* and (Amazon Web Services and Google Compute Engine). Toil also supports two DSLs: *CWL* and *WDL* (experimental).

Toil is built in a modular way so that it can be used on lots of different systems, and with different configurations. The three configurable pieces are the

- *Job Store API*: A filepath or url that can host and centralize all files for a workflow (e.g. a local folder, or an AWS s3 bucket url).
- *Batch System API*: Specifies either a local single-machine or a currently supported HPC environment (lsf, parasol, mesos, slurm, torque, htcondor, kubernetes, or grid_engine). Mesos is a special case, and is launched for cloud environments.
- *Provisioner*: For running in the cloud only. This specifies which cloud provider provides instances to do the “work” of your workflow.

3.1 Job Store

The job store is a storage abstraction which contains all of the information used in a Toil run. This centralizes all of the files used by jobs in the workflow and also the details of the progress of the run. If a workflow crashes or fails, the job store contains all of the information necessary to resume with minimal repetition of work.

Several different job stores are supported, including the file job store and cloud job stores.

3.1.1 File Job Store

The file job store is for use locally, and keeps the workflow information in a directory on the machine where the workflow is launched. This is the simplest and most convenient job store for testing or for small runs.

For an example that uses the file job store, see *Running a basic workflow*.

3.1.2 Cloud Job Stores

Toil currently supports the following cloud storage systems as job stores:

- *AWS Job Store*: An AWS S3 bucket formatted as “aws:<zone>:<bucketname>” where only numbers, letters, and dashes are allowed in the bucket name. Example: *aws:us-west-2:my-aws-jobstore-name*.
- *Google Job Store*: A Google Cloud Storage bucket formatted as “gce:<zone>:<bucketname>” where only numbers, letters, and dashes are allowed in the bucket name. Example: *gce:us-west2-a:my-google-jobstore-name*.

These use cloud buckets to house all of the files. This is useful if there are several different worker machines all running jobs that need to access the job store.

3.2 Batch System

A Toil batch system is either a local single-machine (one computer) or a currently supported HPC cluster of computers (lsf, parasol, mesos, slurm, torque, htcondor, or grid_engine). Mesos is a special case, and is launched for cloud environments. These environments manage individual worker nodes under a leader node to process the work required in a workflow. The leader and its workers all coordinate their tasks and files through a centralized job store location.

See *Batch System API* for a more detailed description of different batch systems.

3.3 Provisioner

The Toil provisioner provides a tool set for running a Toil workflow on a particular cloud platform.

The *Cluster Utilities* are command line tools used to provision nodes in your desired cloud platform. They allow you to launch nodes, ssh to the leader, and rsync files back and forth.

For detailed instructions for using the provisioner see *Running in AWS* or *Running in Google Compute Engine (GCE)*.

Commandline Options

A quick way to see all of Toil's commandline options is by executing the following on a toil script:

```
$ toil example.py --help
```

For a basic toil workflow, Toil has one mandatory argument, the job store. All other arguments are optional.

4.1 The Job Store

Running toil scripts requires a filepath or url to a centralizing location for all of the files of the workflow. This is Toil's one required positional argument: the job store. To use the *quickstart* example, if you're on a node that has a large **/scratch** volume, you can specify that the jobstore be created there by executing: `python HelloWorld.py /scratch/my-job-store`, or more explicitly, `python HelloWorld.py file:/scratch/my-job-store`.

Syntax for specifying different job stores:

Local: `file:job-store-name`

AWS: `aws:region-here:job-store-name`

Google: `google:projectID-here:job-store-name`

Different types of job store options can be found below.

4.2 Commandline Options

Core Toil Options

--workDir WORKDIR Absolute path to directory where temporary files generated during the Toil run should be placed. Temp files and folders, as well as standard output and error from batch system jobs (unless `--noStdOutErr`), will be placed in a directory `toil-<workflowID>` within

workDir. The workflowID is generated by Toil and will be reported in the workflow logs. Default is determined by the variables (TMPDIR, TEMP, TMP) via mkdtemp. This directory needs to exist on all machines running jobs; if capturing standard output and error from batch system jobs is desired, it will generally need to be on a shared file system.

- noStdOutErr** Do not capture standard output and error from batch system jobs.
- stats** Records statistics about the toil workflow to be used by 'toil stats'.
- clean=STATE** Determines the deletion of the jobStore upon completion of the program. Choices: 'always', 'onError', 'never', or 'onSuccess'. The --stats option requires information from the jobStore upon completion so the jobStore will never be deleted with that flag. If you wish to be able to restart the run, choose 'never' or 'onSuccess'. Default is 'never' if stats is enabled, and 'onSuccess' otherwise
- cleanWorkDir STATE** Determines deletion of temporary worker directory upon completion of a job. Choices: 'always', 'never', 'onSuccess'. Default = always. **WARNING:** This option should be changed for debugging only. Running a full pipeline with this option could fill your disk with intermediate data.
- clusterStats FILEPATH** If enabled, writes out JSON resource usage statistics to a file. The default location for this file is the current working directory, but an absolute path can also be passed to specify where this file should be written. This option only applies when using scalable batch systems.
- restart** If --restart is specified then will attempt to restart existing workflow at the location pointed to by the --jobStore option. Will raise an exception if the workflow does not exist.

Logging Options

Toil hides stdout and stderr by default except in case of job failure. Log levels in toil are based on priority from the logging module:

- logOff** Only CRITICAL log levels are shown. Equivalent to --logLevel=OFF or --logLevel=CRITICAL.
- logCritical** Only CRITICAL log levels are shown. Equivalent to --logLevel=OFF or --logLevel=CRITICAL.
- logError** Only ERROR, and CRITICAL log levels are shown. Equivalent to --logLevel=ERROR.
- logWarning** Only WARN, ERROR, and CRITICAL log levels are shown. Equivalent to --logLevel=WARNING.
- logInfo** All log statements are shown, except DEBUG. Equivalent to --logLevel=INFO.
- logDebug** All log statements are shown. Equivalent to --logLevel=DEBUG.
- logLevel=LOGLEVEL** May be set to: OFF (or CRITICAL), ERROR, WARN (or WARNING), INFO, or DEBUG.
- logFile FILEPATH** Specifies a file path to write the logging output to.

- rotatingLogging** Turn on rotating logging, which prevents log files from getting too big (set using `--maxLogFileSize BYTESIZE`).
- maxLogFileSize BYTESIZE** Sets the maximum log file size in bytes (`--rotatingLogging` must be active).

Batch System Options

- batchSystem BATCHSYSTEM** The type of batch system to run the job(s) with, currently can be one of `lsf`, `Mesos`, `slurm`, `torque`, `htcondor`, `single_machine`, `parasol`, `grid_engine`, `kubernetes`. (default: `single_machine`)
- parasolCommand PARASOLCOMMAND** The name or path of the parasol program. Will be looked up on `PATH` unless it starts with a slash. (default: `parasol`)
- parasolMaxBatches PARASOLMAXBATCHES** Maximum number of job batches the Parasol batch is allowed to create. One batch is created for jobs with a unique set of resource requirements. (default: 1000)
- scale SCALE** A scaling factor to change the value of all submitted tasks' submitted cores. Used in `singleMachine` batch system. (default: 1)
- linkImports** When using Toil's `importFile` function for staging, input files are copied to the job store. Specifying this option saves space by symlinking imported files. As long as caching is enabled Toil will protect the file automatically by changing the permissions to read-only.
- mesosMaster MESOSMASTERADDRESS** The host and port of the Mesos master separated by a colon. (default: `169.233.147.202:5050`)

Autoscaling Options

- provisioner CLOUDPROVIDER** The provisioner for cluster auto-scaling. The currently supported choices are `'aws'` or `'gce'`. The default is `None`.
- nodeTypes NODETYPES** List of node types separated by commas. The syntax for each node type depends on the provisioner used. For the `cgccloud` and `AWS` provisioners this is the name of an EC2 instance type, optionally followed by a colon and the price in dollars to bid for a spot instance of that type, for example `'c3.8xlarge:0.42'`. If no spot bid is specified, nodes of this type will be non-preemptable. It is acceptable to specify an instance as both preemptable and non-preemptable, including it twice in the list. In that case, preemptable nodes of that type will be preferred when creating new nodes once the maximum number of preemptable-nodes have been reached.
- minNodes MINNODES** Minimum number of nodes of each type in the cluster, if using auto-scaling. This should be provided as a comma-separated list of the same length as the list of node types. default=0
- maxNodes MAXNODES** Maximum number of nodes of each type in the cluster, if using autoscaling, provided as a comma-separated list. The first value is used as a default if the list length is less than the number of node-Types. default=10
- preemptableCompensation PREEMPTABLECOMPENSATION** The preference of the autoscaler to replace preemptable nodes with non-preemptable nodes, when preemptable nodes cannot be started for some reason. Defaults to 0.0. This value must be between 0.0 and 1.0, inclusive.

A value of 0.0 disables such compensation, a value of 0.5 compensates two missing preemptable nodes with a non-preemptable one. A value of 1.0 replaces every missing pre-emptable node with a non-preemptable one.

- nodeStorage NODESTORAGE** Specify the size of the root volume of worker nodes when they are launched in gigabytes. You may want to set this if your jobs require a lot of disk space. The default value is 50.
- nodeStorageOverrides NODESTORAGEOVERRIDES** Comma-separated list of nodeType:nodeStorage that are used to override the default value from --nodeStorage for the specified nodeType(s). This is useful for heterogeneous jobs where some tasks require much more disk than others.
- metrics** Enable the prometheus/grafana dashboard for monitoring CPU/RAM usage, queue size, and issued jobs.
- defaultMemory INT** The default amount of memory to request for a job. Only applicable to jobs that do not specify an explicit value for this requirement. Standard suffixes like K, Ki, M, Mi, G or Gi are supported. Default is 2.0G
- defaultCores FLOAT** The default number of CPU cores to dedicate a job. Only applicable to jobs that do not specify an explicit value for this requirement. Fractions of a core (for example 0.1) are supported on some batch systems, namely Mesos and singleMachine. Default is 1.0
- defaultDisk INT** The default amount of disk space to dedicate a job. Only applicable to jobs that do not specify an explicit value for this requirement. Standard suffixes like K, Ki, M, Mi, G or Gi are supported. Default is 2.0G
- defaultPreemptable BOOL** Set if jobs that do not specifically prohibit it should be able to run on preemptable (spot) nodes.
- maxCores INT** The maximum number of CPU cores to request from the batch system at any one time. Standard suffixes like K, Ki, M, Mi, G or Gi are supported.
- maxMemory INT** The maximum amount of memory to request from the batch system at any one time. Standard suffixes like K, Ki, M, Mi, G or Gi are supported.
- maxDisk INT** The maximum amount of disk space to request from the batch system at any one time. Standard suffixes like K, Ki, M, Mi, G or Gi are supported.
- retryCount RETRYCOUNT** Number of times to retry a failing job before giving up and labeling job failed. default=1
- doubleMem** If set, batch jobs which die due to reaching memory limit on batch schedulers will have their memory doubled and they will be retried. The remaining retry count will be reduced by 1. Currently only supported by LSF. default=False.
- maxJobDuration MAXJOBURATION** Maximum runtime of a job (in seconds) before we kill it (this is a lower bound, and the actual time before killing the job may be longer).

- rescueJobsFrequency RESCUEJOBSFREQUENCY** Period of time to wait (in seconds) between checking for missing/overlong jobs, that is jobs which get lost by the batch system.
- maxServiceJobs MAXSERVICEJOBS** The maximum number of service jobs that can be run concurrently, excluding service jobs running on preemptable nodes. default=9223372036854775807
- maxPreemptableServiceJobs MAXPREEMPTABLESERVICEJOBS** The maximum number of service jobs that can run concurrently on preemptable nodes. default=9223372036854775807
- deadlockWait DEADLOCKWAIT** Time, in seconds, to tolerate the workflow running only the same service jobs, with no jobs to use them, before declaring the workflow to be deadlocked and stopping. default=60
- deadlockCheckInterval DEADLOCKCHECKINTERVAL** Time, in seconds, to wait between checks to see if the workflow is stuck running only service jobs, with no jobs to use them. Should be shorter than --deadlockWait. May need to be increased if the batch system cannot enumerate running jobs quickly enough, or if polling for running jobs is placing an unacceptable load on a shared cluster. default=30
- statePollingWait STATEPOLLINGWAIT** Time, in seconds, to wait before doing a scheduler query for job state. Return cached results if within the waiting period.

Miscellaneous Options

- disableCaching** Disables caching in the file store. This flag must be set to use a batch system that does not support cleanup, such as Parasol.
- disableChaining** Disables chaining of jobs (chaining uses one job's resource allocation for its successor job if possible).
- maxLogFileSize MAXLOGFILESIZE** The maximum size of a job log file to keep (in bytes), log files larger than this will be truncated to the last X bytes. Setting this option to zero will prevent any truncation. Setting this option to a negative value will truncate from the beginning. Default=62.5 K
- writeLogs FILEPATH** Write worker logs received by the leader into their own files at the specified path. Any non-empty standard output and error from failed batch system jobs will also be written into files at this path. The current working directory will be used if a path is not specified explicitly. Note: By default only the logs of failed jobs are returned to leader. Set log level to 'debug' to get logs back from successful jobs, and adjust 'maxLogFileSize' to control the truncation limit for worker logs.
- writeLogsGzip FILEPATH** Identical to --writeLogs except the logs files are gzipped on the leader.
- realTimeLogging** Enable real-time logging from workers to masters.
- sseKey SSEKEY** Path to file containing 32 character key to be used for server-side encryption on awsJobStore or googleJobStore. SSE will not be used if this flag is not passed.
- setEnv NAME** NAME=VALUE or NAME, -e NAME=VALUE or NAME are also valid. Set an environment variable early on in the worker. If VALUE

is omitted, it will be looked up in the current environment. Independently of this option, the worker will try to emulate the leader's environment before running a job. Using this option, a variable can be injected into the worker process itself before it is started.

- servicePollingInterval SERVICEPOLLINGINTERVAL** Interval of time service jobs wait between polling for the existence of the keep-alive flag (default=60)
- debugWorker** Experimental no forking mode for local debugging. Specifically, workers are not forked and stderr/stdout are not redirected to the log. (default=False)
- disableProgress** Disables the progress bar shown when standard error is a terminal.

4.3 Restart Option

In the event of failure, Toil can resume the pipeline by adding the argument `--restart` and rerunning the python script. Toil pipelines can even be edited and resumed which is useful for development or troubleshooting.

4.4 Running Workflows with Services

Toil supports jobs, or clusters of jobs, that run as *services* to other *accessor* jobs. Example services include server databases or Apache Spark Clusters. As service jobs exist to provide services to accessor jobs their runtime is dependent on the concurrent running of their accessor jobs. The dependencies between services and their accessor jobs can create potential deadlock scenarios, where the running of the workflow hangs because only service jobs are being run and their accessor jobs can not be scheduled because of too limited resources to run both simultaneously. To cope with this situation Toil attempts to schedule services and accessors intelligently, however to avoid a deadlock with workflows running service jobs it is advisable to use the following parameters:

- `--maxServiceJobs`: The maximum number of service jobs that can be run concurrently, excluding service jobs running on preemptable nodes.
- `--maxPreemptableServiceJobs`: The maximum number of service jobs that can run concurrently on preemptable nodes.

Specifying these parameters so that at a maximum cluster size there will be sufficient resources to run accessors in addition to services will ensure that such a deadlock can not occur.

If too low a limit is specified then a deadlock can occur in which toil can not schedule sufficient service jobs concurrently to complete the workflow. Toil will detect this situation if it occurs and throw a `toil.DeadlockException` exception. Increasing the cluster size and these limits will resolve the issue.

4.5 Setting Options directly with the Toil Script

It's good to remember that commandline options can be overridden in the Toil script itself. For example, `toil.job.Job.Runner.getDefaultOptions()` can be used to run toil with all default options, and in this example, it will override commandline args to run the default options and always run with the `"/toilWorkflow"` directory specified as the jobstore:

```
options = Job.Runner.getDefaultOptions("./toilWorkflow") # Get the options object

with Toil(options) as toil:
    toil.start(Job()) # Run the script
```

However, each option can be explicitly set within the script by supplying arguments (in this example, we are setting `logLevel = "DEBUG"` (all log statements are shown) and `clean="ALWAYS"` (always delete the jobstore) like so:

```
options = Job.Runner.getDefaultOptions("./toilWorkflow") # Get the options object
options.logLevel = "DEBUG" # Set the log level to the debug level.
options.clean = "ALWAYS" # Always delete the jobStore after a run

with Toil(options) as toil:
    toil.start(Job()) # Run the script
```

However, the usual incantation is to accept commandline args from the user with the following:

```
parser = Job.Runner.getDefaultArgumentParser() # Get the parser
options = parser.parse_args() # Parse user args to create the options object

with Toil(options) as toil:
    toil.start(Job()) # Run the script
```

Which can also, of course, then accept script supplied arguments as before (which will overwrite any user supplied args):

```
parser = Job.Runner.getDefaultArgumentParser() # Get the parser
options = parser.parse_args() # Parse user args to create the options object
options.logLevel = "DEBUG" # Set the log level to the debug level.
options.clean = "ALWAYS" # Always delete the jobStore after a run

with Toil(options) as toil:
    toil.start(Job()) # Run the script
```


Toil has a number of tools to assist in debugging. Here we provide help in working through potential problems that a user might encounter in attempting to run a workflow.

5.1 Introspecting the Jobstore

Note: Currently these features are only implemented for use locally (single machine) with the fileJobStore.

To view what files currently reside in the jobstore, run the following command:

```
$ toil debug-file file:path-to-jobstore-directory \  
  --listFilesInJobStore
```

When run from the commandline, this should generate a file containing the contents of the job store (in addition to displaying a series of log messages to the terminal). This file is named “jobstore_files.txt” by default and will be generated in the current working directory.

If one wishes to copy any of these files to a local directory, one can run for example:

```
$ toil debug-file file:path-to-jobstore \  
  --fetch overview.txt *.bam *.fastq \  
  --localFilePath=/home/user/localpath
```

To fetch `overview.txt`, and all `.bam` and `.fastq` files. This can be used to recover previously used input and output files for debugging or reuse in other workflows, or use in general debugging to ensure that certain outputs were imported into the jobStore.

5.2 Stats and Status

See *Stats Command* for more about gathering statistics about job success, runtime, and resource usage from workflows.

5.3 Using a Python debugger

If you execute a workflow using the `--debugWorker` flag, Toil will not fork in order to run jobs, which means you can either use `pdb`, or an IDE that supports debugging Python as you would normally. Note that the `--debugWorker` flag will only work with the `singleMachine` batch system (the default), and not any of the custom job schedulers.

Running in the Cloud

Toil supports Amazon Web Services (AWS) and Google Compute Engine (GCE) in the cloud and has autoscaling capabilities that can adapt to the size of your workflow, whether your workflow requires 10 instances or 20,000.

Toil does this by creating a virtual cluster with [Apache Mesos](#). [Apache Mesos](#) requires a leader node to coordinate the workflow, and worker nodes to execute the various tasks within the workflow. As the workflow runs, Toil will “autoscale”, creating and terminating workers as needed to meet the demands of the workflow.

Once a user is familiar with the basics of running toil locally (specifying a *jobStore*, and how to write a toil script), they can move on to the guides below to learn how to translate these workflows into cloud ready workflows.

6.1 Managing a Cluster of Virtual Machines (Provisioning)

Toil can launch and manage a cluster of virtual machines to run using the *provisioner* to run a workflow distributed over several nodes. The provisioner also has the ability to automatically scale up or down the size of the cluster to handle dynamic changes in computational demand (autoscaling). Currently we have working provisioners with AWS and GCE (Azure support has been deprecated).

Toil uses [Apache Mesos](#) as the *Batch System*.

See here for instructions for *Running in AWS*.

See here for instructions for *Running in Google Compute Engine (GCE)*.

6.2 Storage (Toil jobStore)

Toil can make use of cloud storage such as AWS or Google buckets to take care of storage needs.

This is useful when running Toil in single machine mode on any cloud platform since it allows you to make use of their integrated storage systems.

For an overview of the job store see *Job Store*.

For instructions configuring a particular job store see:

- *AWS Job Store*
- *Google Job Store*

7.1 Running on Kubernetes

Kubernetes is a very popular container orchestration tool that has become a *de facto* cross-cloud-provider API for accessing cloud resources. Major cloud providers like Amazon, Microsoft, Kubernetes owner Google, and DigitalOcean have invested heavily in making Kubernetes work well on their platforms, by writing their own deployment documentation and developing provider-managed Kubernetes-based products. Using `minikube`, Kubernetes can even be run on a single machine.

Toil supports running Toil workflows against a Kubernetes cluster, either in the cloud or deployed on user-owned hardware.

7.1.1 Preparing your Kubernetes environment

1. Get a Kubernetes cluster

To run Toil workflows on Kubernetes, you need to have a Kubernetes cluster set up. This will not be covered here, but there are many options available, and which one you choose will depend on which cloud ecosystem if any you use already, and on pricing. If you are just following along with the documentation, use `minikube` on your local machine.

Note that currently the only way to run a Toil workflow on Kubernetes is to use the AWS Job Store, so your Kubernetes workflow will currently have to store its data in Amazon's cloud regardless of where you run it. This can result in significant egress charges from Amazon if you run it outside of Amazon.

Kubernetes Cluster Providers:

- Your own institution
- Amazon EKS
- Microsoft Azure AKS
- Google GKE
- DigitalOcean Kubernetes

- [minikube](#)

2. Get a Kubernetes context on your local machine

There are two main ways to run Toil workflows on Kubernetes. You can either run the Toil leader on a machine outside the cluster, with jobs submitted to and run on the cluster, or you can submit the Toil leader itself as a job and have it run inside the cluster. Either way, you will need to configure your own machine to be able to submit jobs to the Kubernetes cluster. Generally, this involves creating and populating a file named `.kube/config` in your user's home directory, and specifying the cluster to connect to, the certificate and token information needed for mutual authentication, and the Kubernetes namespace within which to work. However, Kubernetes configuration can also be picked up from other files in the `.kube` directory, environment variables, and the enclosing host when running inside a Kubernetes-managed container.

You will have to do different things here depending on where you got your Kubernetes cluster:

- [Configuring for Amazon EKS](#)
- [Configuring for Microsoft Azure AKS](#)
- [Configuring for Google GKE](#)
- [Configuring for DigitalOcean Kubernetes Clusters](#)
- [Configuring for minikube](#)

Toil's internal Kubernetes configuration logic mirrors that of the `kubectl` command. Toil workflows will use the current `kubectl` context to launch their Kubernetes jobs.

3. If running the Toil leader in the cluster, get a service account

If you are going to run your workflow's leader within the Kubernetes cluster (see [Option 1: Running the Leader Inside Kubernetes](#)), you will need a service account in your chosen Kubernetes namespace. Most namespaces should have a service account named `default` which should work fine. If your cluster requires you to use a different service account, you will need to obtain its name and use it when launching the Kubernetes job containing the Toil leader.

4. Set up appropriate permissions

Your local Kubernetes context and/or the service account you are using to run the leader in the cluster will need to have certain permissions in order to run the workflow. Toil needs to be able to interact with jobs and pods in the cluster, and to retrieve pod logs. You as a user may need permission to set up an AWS credentials secret, if one is not already available. Additionally, it is very useful for you as a user to have permission to interact with nodes, and to shell into pods.

The appropriate permissions may already be available to you and your service account by default, especially in managed or ease-of-use-optimized setups such as EKS or minikube.

However, if the appropriate permissions are not already available, you or your cluster administrator will have to grant them manually. The following Role (`toil-user`) and ClusterRole (`node-reader`), to be applied with `kubectl apply -f filename.yaml`, should grant sufficient permissions to run Toil workflows when bound to your account and the service account used by Toil workflows. Be sure to replace `YOUR_NAMESPACE_HERE` with the namespace you are running your workflows in

```
apiVersion: rbac.authorization.k8s.io/v1
kind: Role
metadata:
  namespace: YOUR_NAMESPACE_HERE
  name: toil-user
rules:
- apiGroups: ["*"]
  resources: ["*"]
  verbs: ["explain", "get", "watch", "list", "describe", "logs", "attach", "exec",
→ "port-forward", "proxy", "cp", "auth"]
```

(continues on next page)

(continued from previous page)

```

- apiGroups: ["batch"]
  resources: ["*"]
  verbs: ["get", "watch", "list", "create", "run", "set", "delete"]
- apiGroups: [""]
  resources: ["secrets", "pods", "pods/attach", "podtemplates", "configmaps",
↪ "events", "services"]
  verbs: ["patch", "get", "update", "watch", "list", "create", "run", "set",
↪ "delete", "exec"]
- apiGroups: [""]
  resources: ["pods", "pods/log"]
  verbs: ["get", "list"]
- apiGroups: [""]
  resources: ["pods/exec"]
  verbs: ["create"]

```

```

apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
  name: node-reader
rules:
- apiGroups: [""]
  resources: ["nodes"]
  verbs: ["get", "list", "describe"]
- apiGroups: [""]
  resources: ["namespaces"]
  verbs: ["get", "list", "describe"]
- apiGroups: ["metrics.k8s.io"]
  resources: ["*"]
  verbs: ["*"]

```

To bind a user or service account to the Role or ClusterRole and actually grant the permissions, you will need a RoleBinding and a ClusterRoleBinding, respectively. Make sure to fill in the namespace, username, and service account name, and add more user stanzas if your cluster is to support multiple Toil users.

```

apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: toil-developer-member
  namespace: toil
subjects:
- kind: User
  name: YOUR_KUBERNETES_USERNAME_HERE
  apiGroup: rbac.authorization.k8s.io
- kind: ServiceAccount
  name: YOUR_SERVICE_ACCOUNT_NAME_HERE
  namespace: YOUR_NAMESPACE_HERE
roleRef:
  kind: Role
  name: toil-user
  apiGroup: rbac.authorization.k8s.io

```

```

apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: read-nodes
subjects:

```

(continues on next page)

(continued from previous page)

```

- kind: User
  name: YOUR_KUBERNETES_USERNAME_HERE
  apiGroup: rbac.authorization.k8s.io
- kind: ServiceAccount
  name: YOUR_SERVICE_ACCOUNT_NAME_HERE
  namespace: YOUR_NAMESPACE_HERE
roleRef:
  kind: ClusterRole
  name: node-reader
  apiGroup: rbac.authorization.k8s.io

```

7.1.2 AWS Job Store for Kubernetes

Currently, the only job store, which is what Toil uses to exchange data between jobs, that works with jobs running on Kubernetes is the AWS Job Store. This requires that the Toil leader and Kubernetes jobs be able to connect to and use Amazon S3 and Amazon SimpleDB. It also requires that you have an Amazon Web Services account.

1. Get access to AWS S3 and SimpleDB

In your AWS account, you need to create an AWS access key. First go to the IAM dashboard; for “us-west1”, the link would be:

```
https://console.aws.amazon.com/iam/home?region=us-west-1#/home
```

Then create an access key, and save the Access Key ID and the Secret Key. As documented in [the AWS documentation](#):

1. On the IAM Dashboard page, choose your account name in the navigation bar, and then choose My Security Credentials.
2. Expand the Access keys (access key ID and secret access key) section.
3. Choose Create New Access Key. Then choose Download Key File to save the access key ID and secret access key to a file on your computer. After you close the dialog box, you can’t retrieve this secret access key again.

Make sure that, if your AWS infrastructure requires your user to authenticate with a multi-factor authentication (MFA) token, you obtain a second secret key and access key that don’t have this requirement. The secret key and access key used to populate the Kubernetes secret that allows the jobs to contact the job store need to be usable without human intervention.

2. Configure AWS access from the local machine

This only really needs to happen if you run the leader on the local machine. But we need the files in place to fill in the secret in the next step. Run:

```
$ aws configure
```

Then when prompted, enter your secret key and access key. This should create a file `~/.aws/credentials` that looks like this:

```

[default]
aws_access_key_id = BLAH
aws_secret_access_key = blahblahblah

```

3. Create a Kubernetes secret to give jobs access to AWS

Go into the directory where the `credentials` file is:

```
$ cd ~/.aws
```

Then, create a Kubernetes secret that contains it. We'll call it `aws-credentials`:

```
$ kubectl create secret generic aws-credentials --from-file credentials
```

7.1.3 Configuring Toil for your Kubernetes environment

To configure your workflow to run on Kubernetes, you will have to configure several environment variables, in addition to passing the `--batchSystem kubernetes` option. Doing the research to figure out what values to give these variables may require talking to your cluster provider.

1. `TOIL_AWS_SECRET_NAME` is the most important, and **must** be set to the secret that contains your AWS credentials file, **if** your cluster nodes don't otherwise have access to S3 and SimpleDB (such as through IAM roles). This is required for the AWS job store to work, which is currently the only job store that can be used on Kubernetes. In this example we are using `aws-credentials`.
2. `TOIL_KUBERNETES_HOST_PATH` **can** be set to allow Toil jobs on the same physical host to share a cache. It should be set to a path on the host where the shared cache should be stored. It will be mounted as `/var/lib/toil`, or at `TOIL_WORKDIR` if specified, inside the container. This path must already exist on the host, and must have as much free space as your Kubernetes node offers to jobs. In this example, we are using `/data/scratch`. To actually make use of caching, make sure to also pass `--disableCaching false` to your Toil workflow.
3. `TOIL_KUBERNETES_OWNER` **should** be set to the username of the user running the Toil workflow. The jobs that Toil creates will include this username, so they can be more easily recognized, and cleaned up by the user if anything happens to the Toil leader. In this example we are using `demo-user`.

Note that Docker containers cannot be run inside of unprivileged Kubernetes pods (which are themselves containers). The Docker daemon does not (yet) support this. Other tools, such as Singularity in its user-namespace mode, are able to run containers from within containers. If using Singularity to run containerized tools, and you want downloaded container images to persist between Toil jobs, you will also want to set `TOIL_KUBERNETES_HOST_PATH` and make sure that Singularity is downloading its containers under the Toil work directory (`/var/lib/toil` by default) by setting `SINGULARITY_CACHEDIR`. However, you will need to make sure that no two jobs try to download the same container at the same time; Singularity has no synchronization or locking around its cache, but the cache is also not safe for simultaneous access by multiple Singularity invocations. Some Toil workflows use their own custom workaround logic for this problem; this work is likely to be made part of Toil in a future release.

7.1.4 Running workflows

To run the workflow, you will need to run the Toil leader process somewhere. It can either be run inside Kubernetes as a Kubernetes job, or outside Kubernetes as a normal command.

Option 1: Running the Leader Inside Kubernetes

Once you have determined a set of environment variable values for your workflow run, write a YAML file that defines a Kubernetes job to run your workflow with that configuration. Some configuration items (such as your username, and the name of your AWS credentials secret) need to be written into the YAML so that they can be used from the leader as well.

Note that the leader pod will need your workflow script, its other dependencies, and Toil all installed. An easy way to get Toil installed is to start with the Toil appliance image for the version of Toil you want to use. In this example, we use `quay.io/ucsc_cgl/toil:4.1.0`.

Here's an example YAML file to run a test workflow:

```
apiVersion: batch/v1
kind: Job
metadata:
  # It is good practice to include your username in your job name.
  # Also specify it in TOIL_KUBERNETES_OWNER
  name: demo-user-toil-test
  # Do not try and rerun the leader job if it fails
spec:
  backoffLimit: 0
  template:
    spec:
      # Do not restart the pod when the job fails, but keep it around so the
      # log can be retrieved
      restartPolicy: Never
      volumes:
      - name: aws-credentials-vol
        secret:
          # Make sure the AWS credentials are available as a volume.
          # This should match TOIL_AWS_SECRET_NAME
          secretName: aws-credentials
      # You may need to replace this with a different service account name as
      # appropriate for your cluster.
      serviceAccountName: default
      containers:
      - name: main
        image: quay.io/ucsc_cgl/toil:4.1.0
        env:
          # Specify your username for inclusion in job names
          - name: TOIL_KUBERNETES_OWNER
            value: demo-user
          # Specify where to find the AWS credentials to access the job store with
          - name: TOIL_AWS_SECRET_NAME
            value: aws-credentials
          # Specify where per-host caches should be stored, on the Kubernetes hosts.
          # Needs to be set for Toil's caching to be efficient.
          - name: TOIL_KUBERNETES_HOST_PATH
            value: /data/scratch
        volumeMounts:
          # Mount the AWS credentials volume
          - mountPath: /root/.aws
            name: aws-credentials-vol
      resources:
        # Make sure to set these resource limits to values large enough
        # to accomodate the work your workflow does in the leader
        # process, but small enough to fit on your cluster.
        #
        # Since no request values are specified, the limits are also used
        # for the requests.
        limits:
          cpu: 2
          memory: "4Gi"
          ephemeral-storage: "10Gi"
      command:
      - /bin/bash
      - -c
```

(continues on next page)

(continued from previous page)

```

- |
  # This Bash script will set up Toil and the workflow to run, and run them.
  set -e
  # We make sure to create a work directory; Toil can't hot-deploy a
  # script from the root of the filesystem, which is where we start.
  mkdir /tmp/work
  cd /tmp/work
  # We make a virtual environment to allow workflow dependencies to be
  # hot-deployed.
  #
  # We don't really make use of it in this example, but for workflows
  # that depend on PyPI packages we will need this.
  #
  # We use --system-site-packages so that the Toil installed in the
  # appliance image is still available.
  virtualenv --python python3 --system-site-packages venv
  . venv/bin/activate
  # Now we install the workflow. Here we're using a demo workflow
  # script from Toil itself.
  wget https://raw.githubusercontent.com/DataBiosphere/toil/releases/4.1.0/src/
  ↪toil/test/docs/scripts/tutorial_helloworld.py
  # Now we run the workflow. We make sure to use the Kubernetes batch
  # system and an AWS job store, and we set some generally useful
  # logging options. We also make sure to enable caching.
  python3 tutorial_helloworld.py \
    aws:us-west-2:demouser-toil-test-jobstore \
    --batchSystem kubernetes \
    --realTimeLogging \
    --logInfo \
    --disableCaching false

```

You can save this YAML as `leader.yaml`, and then run it on your Kubernetes installation with:

```
$ kubectl apply -f leader.yaml
```

To monitor the progress of the leader job, you will want to read its logs. If you are using a Kubernetes dashboard such as [k9s](#), you can simply find the pod created for the job in the dashboard, and view its logs there. If not, you will need to locate the pod by hand.

Monitoring and Debugging Kubernetes Jobs and Pods

The following techniques are most useful for looking at the pod which holds the Toil leader, but they can also be applied to individual Toil jobs on Kubernetes, even when the leader is outside the cluster.

Kubernetes names pods for jobs by appending a short random string to the name of the job. You can find the name of the pod for your job by doing:

```

$ kubectl get pods | grep demo-user-toil-test
demo-user-toil-test-g5496                1/1      Running   ↪
↪0                                     2m

```

Assuming you have set `TOIL_KUBERNETES_OWNER` correctly, you should be able to find all of your workflow's pods by searching for your username:

```
$ kubectl get pods | grep demo-user
```

If the status of a pod is anything other than `Pending`, you will be able to view its logs with:

```
$ kubectl logs demo-user-toil-test-g5496
```

This will dump the pod's logs from the beginning to now and terminate. To follow along with the logs from a running pod, add the `-f` option:

```
$ kubectl logs -f demo-user-toil-test-g5496
```

A status of `ImagePullBackoff` suggests that you have requested to use an image that is not available. Check the image section of your YAML if you are looking at a leader, or the value of `TOIL_APPLIANCE_SELF` if you are delaying with a worker job. You also might want to check your Kubernetes node's Internet connectivity and DNS function; in Kubernetes, DNS depends on system-level pods which can be terminated or evicted in cases of resource oversubscription, just like user workloads.

If your pod seems to be stuck `Pending`, `ContainerCreating`, you can get information on what is wrong with it by using `kubectl describe pod`:

```
$ kubectl describe pod demo-user-toil-test-g5496
```

Pay particular attention to the `Events:` section at the end of the output. An indication that a job is too big for the available nodes on your cluster, or that your cluster is too busy for your jobs, is `FailedScheduling` events:

Type	Reason	Age	From	Message
----	-----	----	----	-----
Warning	FailedScheduling	13s (x79 over 100m)	default-scheduler	0/4 nodes are_
				↪available: 1 Insufficient cpu, 1 Insufficient ephemeral-storage, 4 Insufficient_
				↪memory.

If a pod is running but seems to be behaving erratically, or seems stuck, you can shell into it and look around:

```
$ kubectl exec -ti demo-user-toil-test-g5496 /bin/bash
```

One common cause of stuck pods is attempting to use more memory than allowed by Kubernetes (or by the Toil job's memory resource requirement), but in a way that does not trigger the Linux OOM killer to terminate the pod's processes. In these cases, the pod can remain stuck at nearly 100% memory usage more or less indefinitely, and attempting to shell into the pod (which needs to start a process within the pod, using some of its memory) will fail. In these cases, the recommended solution is to kill the offending pod and increase its (or its Toil job's) memory requirement, or reduce its memory needs by adapting user code.

When Things Go Wrong

The Toil Kubernetes batch system includes cleanup code to terminate worker jobs when the leader shuts down. However, if the leader pod is removed by Kubernetes, is forcibly killed or otherwise suffers a sudden existence failure, it can go away while its worker jobs live on. It is not recommended to restart a workflow in this state, as jobs from the previous invocation will remain running and will be trying to modify the job store concurrently with jobs from the new invocation.

To clean up dangling jobs, you can use the following snippet:

```
$ kubectl get jobs | grep demo-user | cut -f1 -d' ' | xargs -n10 kubectl delete job
```

This will delete all jobs with `demo-user`'s username in their names, in batches of 10. You can also use the UUID that Toil assigns to a particular workflow invocation in the filter, to clean up only the jobs pertaining to that workflow invocation.

Option 2: Running the Leader Outside Kubernetes

If you don't want to run your Toil leader inside Kubernetes, you can run it locally instead. This can be useful when developing a workflow; files can be hot-deployed from your local machine directly to Kubernetes. However, your local machine will have to have (ideally role-assumption- and MFA-free) access to AWS, and access to Kubernetes. Real time logging will not work unless your local machine is able to listen for incoming UDP packets on arbitrary ports on the address it uses to contact the IPv4 Internet; Toil does no NAT traversal or detection.

Note that if you set `TOIL_WORKDIR` when running your workflow like this, it will need to be a directory that exists both on the host and in the Toil appliance.

Here is an example of running our test workflow leader locally, outside of Kubernetes:

```
$ export TOIL_KUBERNETES_OWNER=demo-user # This defaults to your local username if
↪not set
$ export TOIL_AWS_SECRET_NAME=aws-credentials
$ export TOIL_KUBERNETES_HOST_PATH=/data/scratch
$ virtualenv --python python3 --system-site-packages venv
$ . venv/bin/activate
$ wget https://raw.githubusercontent.com/DataBiosphere/toil/releases/4.1.0/src/toil/
↪test/docs/scripts/tutorial_helloworld.py
$ python3 tutorial_helloworld.py \
    aws:us-west-2:demouser-toil-test-jobstore \
    --batchSystem kubernetes \
    --realTimeLogging \
    --logInfo \
    --disableCaching false
```

7.2 Running in AWS

Toil jobs can be run on a variety of cloud platforms. Of these, Amazon Web Services (AWS) is currently the best-supported solution. Toil provides the *Cluster Utilities* to conveniently create AWS clusters, connect to the leader of the cluster, and then launch a workflow. The leader handles distributing the jobs over the worker nodes and autoscaling to optimize costs.

The *Running a Workflow with Autoscaling* section details how to create a cluster and run a workflow that will dynamically scale depending on the workflow's needs.

The *Static Provisioning* section explains how a static cluster (one that won't automatically change in size) can be created and provisioned (grown, shrunk, destroyed, etc.).

7.2.1 Preparing your AWS environment

To use Amazon Web Services (AWS) to run Toil or to just use S3 to host the files during the computation of a workflow, first set up and configure an account with AWS:

1. If necessary, create and activate an [AWS account](#)
2. Next, generate a key pair for AWS with the command (do NOT generate your key pair with the Amazon browser):

```
$ ssh-keygen -t rsa
```

3. This should prompt you to save your key. Please save it in

```
~/ .ssh/id_rsa
```

4. Now move this to where your OS can see it as an authorized key:

```
$ cat ~/.ssh/id_rsa.pub >> ~/.ssh/authorized_keys
$ eval `ssh-agent -s`
$ ssh-add
```

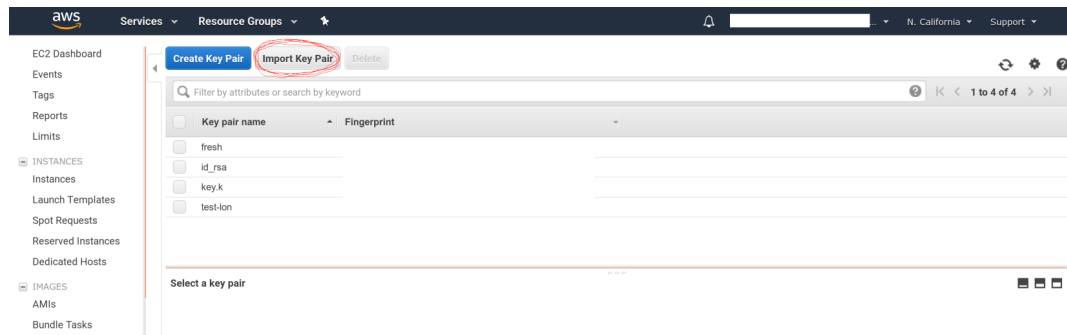
5. You'll also need to chmod your private key (good practice but also enforced by AWS):

```
$ chmod 400 id_rsa
```

6. Now you'll need to add the key to AWS via the browser. For example, on us-west1, this address would be accessible at:

```
https://us-west-1.console.aws.amazon.com/ec2/v2/home?region=us-west-1
↪ #KeyPairs:sort=keyName
```

7. Now click on the “Import Key Pair” button to add your key:



8. Next, you need to create an AWS access key. First go to the IAM dashboard, again; for “us-west1”, the example link would be here:

```
https://console.aws.amazon.com/iam/home?region=us-west-1#/home
```

9. The directions (transcribed from: <https://docs.aws.amazon.com/general/latest/gr/managing-aws-access-keys.html>) are now:

1. On the IAM Dashboard page, choose your account name in the navigation bar, and then choose My Security Credentials.
2. Expand the Access keys (access key ID and secret access key) section.
3. Choose Create New Access Key. Then choose Download Key File to save the access key ID and secret access key to a file on your computer. After you close the dialog box, you can't retrieve this secret access key again.

10. Now you should have a newly generated “AWS Access Key ID” and “AWS Secret Access Key”. We can now install the AWS CLI and make sure that it has the proper credentials:

```
$ pip install awscli --upgrade --user
```

11. Now configure your AWS credentials with:

```
$ aws configure
```

12. Add your “AWS Access Key ID” and “AWS Secret Access Key” from earlier and your region and output format:

```
" AWS Access Key ID [*****Q65Q]: "
" AWS Secret Access Key [*****G0ys]: "
" Default region name [us-west-1]: "
" Default output format [json]: "
```

This will create the files `~/.aws/config` and `~/.aws/credentials`.

13. If not done already, install toil (example uses version 3.12.0, but we recommend the latest release):

```
$ virtualenv venv
$ source venv/bin/activate
$ pip install toil[all]==3.12.0
```

14. Now that toil is installed and you are running a virtualenv, an example of launching a toil leader node would be the following (again, note that we set `TOIL_APPLIANCE_SELF` to toil version 3.12.0 in this example, but please set the version to the installed version that you are using if you're using a different version):

```
$ TOIL_APPLIANCE_SELF=quay.io/ucsc_cgl/toil:3.12.0 \
  toil launch-cluster clustername \
  --leaderNodeType t2.medium \
  --zone us-west-1a \
  --keyPairName id_rsa
```

To further break down each of these commands:

TOIL_APPLIANCE_SELF=quay.io/ucsc_cgl/toil:latest — This is optional. It specifies a mesos docker image that we maintain with the latest version of toil installed on it. If you want to use a different version of toil, please specify the image tag you need from https://quay.io/repository/ucsc_cgl/toil?tag=latest&tab=tags.

toil launch-cluster — Base command in toil to launch a cluster.

clustername — Just choose a name for your cluster.

--leaderNodeType t2.medium — Specify the leader node type. Make a t2.medium (2CPU; 4Gb RAM; \$0.0464/Hour). List of available AWS instances: <https://aws.amazon.com/ec2/pricing/on-demand/>

--zone us-west-1a — Specify the AWS zone you want to launch the instance in. Must have the same prefix as the zone in your awscli credentials (which, in the example of this tutorial is: “us-west-1”).

--keyPairName id_rsa — The name of your key pair, which should be “id_rsa” if you’ve followed this tutorial.

7.2.2 AWS Job Store

Using the AWS job store is straightforward after you’ve finished *Preparing your AWS environment*; all you need to do is specify the prefix for the job store name.

To run the sort example *sort example* with the AWS job store you would type

```
$ python sort.py aws:us-west-2:my-aws-sort-jobstore
```

7.2.3 Toil Provisioner

The Toil provisioner is included in Toil alongside the `[aws]` extra and allows us to spin up a cluster.

Getting started with the provisioner is simple:

1. Make sure you have Toil installed with the AWS extras. For detailed instructions see [Installing Toil with Extra Features](#).
2. You will need an AWS account and you will need to save your AWS credentials on your local machine. For help setting up an AWS account see [here](#). For setting up your AWS credentials follow instructions [here](#).

The Toil provisioner is built around the Toil Appliance, a Docker image that bundles Toil and all its requirements (e.g. Mesos). This makes deployment simple across platforms, and you can even simulate a cluster locally (see [Developing with Docker](#) for details).

Choosing Toil Appliance Image

When using the Toil provisioner, the appliance image will be automatically chosen based on the pip-installed version of Toil on your system. That choice can be overridden by setting the environment variables `TOIL_DOCKER_REGISTRY` and `TOIL_DOCKER_NAME` or `TOIL_APPLIANCE_SELF`. See [Environment Variables](#) for more information on these variables. If you are developing with autoscaling and want to test and build your own appliance have a look at [Developing with Docker](#).

For information on using the Toil Provisioner have a look at [Running a Workflow with Autoscaling](#).

7.2.4 Details about Launching a Cluster in AWS

Using the provisioner to launch a Toil leader instance is simple using the `launch-cluster` command. For example, to launch a cluster named “my-cluster” with a `t2.medium` leader in the `us-west-2a` zone, run

```
(venv) $ toil launch-cluster my-cluster \
      --leaderNodeType t2.medium \
      --zone us-west-2a \
      --keyPairName <your-AWS-key-pair-name>
```

The cluster name is used to uniquely identify your cluster and will be used to populate the instance’s Name tag. Also, the Toil provisioner will automatically tag your cluster with an Owner tag that corresponds to your keypair name to facilitate cost tracking. In addition, the `ToilNodeType` tag can be used to filter “leader” vs. “worker” nodes in your cluster.

The `leaderNodeType` is an [EC2 instance type](#). This only affects the leader node.

The `--zone` parameter specifies which EC2 availability zone to launch the cluster in. Alternatively, you can specify this option via the `TOIL_AWS_ZONE` environment variable. Note: the zone is different from an EC2 region. A region corresponds to a geographical area like `us-west-2` (Oregon), and availability zones are partitions of this area like `us-west-2a`.

By default, Toil creates an IAM role for each cluster with sufficient permissions to perform cluster operations (e.g. full S3, EC2, and SDB access). If the default permissions are not sufficient for your use case (e.g. if you need access to ECR), you may create a custom IAM role with all necessary permissions and set the `--awsEc2ProfileArn` parameter when launching the cluster. Note that your custom role must at least have [these permissions](#) in order for the Toil cluster to function properly.

In addition, Toil creates a new security group with the same name as the cluster name with default rules (e.g. opens port 22 for SSH access). If you require additional security groups, you may use the `--awsEc2ExtraSecurityGroupId` parameter when launching the cluster. **Note:** Do not use the same name as the cluster name for the extra security groups as any security group matching the cluster name will be deleted once the cluster is destroyed.

For more information on options try:

```
(venv) $ toil launch-cluster --help
```

Static Provisioning

Toil can be used to manage a cluster in the cloud by using the *Cluster Utilities*. The cluster utilities also make it easy to run a toil workflow directly on this cluster. We call this static provisioning because the size of the cluster does not change. This is in contrast with *Running a Workflow with Autoscaling*.

To launch worker nodes alongside the leader we use the `-w` option:

```
(venv) $ toil launch-cluster my-cluster \
    --leaderNodeType t2.small -z us-west-2a \
    --keyPairName your-AWS-key-pair-name \
    --nodeTypes m3.large,t2.micro -w 1,4
```

This will spin up a leader node of type `t2.small` with five additional workers — one `m3.large` instance and four `t2.micro`.

Currently static provisioning is only possible during the cluster’s creation. The ability to add new nodes and remove existing nodes via the native provisioner is in development. Of course the cluster can always be deleted with the *Destroy-Cluster Command* utility.

Uploading Workflows

Now that our cluster is launched, we use the *Rsync-Cluster Command* utility to copy the workflow to the leader. For a simple workflow in a single file this might look like

```
(venv) $ toil rsync-cluster -z us-west-2a my-cluster toil-workflow.py :/
```

Note: If your toil workflow has dependencies have a look at the *Auto-Deployment* section for a detailed explanation on how to include them.

Running a Workflow with Autoscaling

Autoscaling is a feature of running Toil in a cloud whereby additional cloud instances are launched to run the workflow. Autoscaling leverages Mesos containers to provide an execution environment for these workflows.

Note: Make sure you’ve done the AWS setup in *Preparing your AWS environment*.

1. Download `sort.py`
2. Launch the leader node in AWS using the *Launch-Cluster Command* command:

```
(venv) $ toil launch-cluster <cluster-name> \
    --keyPairName <AWS-key-pair-name> \
    --leaderNodeType t2.medium \
    --zone us-west-2a
```

3. Copy the `sort.py` script up to the leader node:

```
(venv) $ toil rsync-cluster -z us-west-2a <cluster-name> sort.py :/root
```

4. Login to the leader node:

```
(venv) $ toil ssh-cluster -z us-west-2a <cluster-name>
```

5. Run the script as an autoscaling workflow:

```
$ python /root/sort.py aws:us-west-2:<my-jobstore-name> \
    --provisioner aws \
    --nodeTypes c3.large \
    --maxNodes 2 \
    --batchSystem mesos
```

Note: In this example, the autoscaling Toil code creates up to two instances of type *c3.large* and launches Mesos slave containers inside them. The containers are then available to run jobs defined by the *sort.py* script. Toil also creates a bucket in S3 called *aws:us-west-2:autoscaling-sort-jobstore* to store intermediate job results. The Toil autoscaler can also provision multiple different node types, which is useful for workflows that have jobs with varying resource requirements. For example, one could execute the script with `--nodeTypes c3.large, r3.xlarge --maxNodes 5, 1`, which would allow the provisioner to create up to five *c3.large* nodes and one *r3.xlarge* node for memory-intensive jobs. In this situation, the autoscaler would avoid creating the more expensive *r3.xlarge* node until needed, running most jobs on the *c3.large* nodes.

1. View the generated file to sort:

```
$ head fileToSort.txt
```

2. View the sorted file:

```
$ head sortedFile.txt
```

For more information on other autoscaling (and other) options have a look at [Commandline Options](#) and/or run

```
$ python my-toil-script.py --help
```

Important: Some important caveats about starting a toil run through an ssh session are explained in the [Ssh-Cluster Command](#) section.

Preemptability

Toil can run on a heterogeneous cluster of both preemptable and non-preemptable nodes. Being preemptable node simply means that the node may be shut down at any time, while jobs are running. These jobs can then be restarted later somewhere else.

A node type can be specified as preemptable by adding a [spot bid](#) to its entry in the list of node types provided with the `--nodeTypes` flag. If spot instance prices rise above your bid, the preemptable node will be shut down.

While individual jobs can each explicitly specify whether or not they should be run on preemptable nodes via the boolean `preemptable` resource requirement, the `--defaultPreemptable` flag will allow jobs without a `preemptable` requirement to run on preemptable machines.

Specify Preemptability Carefully

Ensure that your choices for `--nodeTypes` and `--maxNodes` <> make sense for your workflow and won't cause it to hang. You should make sure the provisioner is able to create nodes large enough to run the largest job in the workflow, and that non-preemptable node types are allowed if there are non-preemptable jobs in the workflow.

Finally, the `--preemptableCompensation` flag can be used to handle cases where preemptable nodes may not be available but are required for your workflow. With this flag enabled, the autoscaler will attempt to compensate for a shortage of preemptable nodes of a certain type by creating non-preemptable nodes of that type, if non-preemptable nodes of that type were specified in `--nodeTypes`.

Using MinIO and S3-Compatible object stores

Toil can be configured to access files stored in an [S3-compatible object store](#) such as [MinIO](#). The following environment variables can be used to configure the S3 connection used:

- `TOIL_S3_HOST`: the IP address or hostname to use for connecting to S3
- `TOIL_S3_PORT`: the port number to use for connecting to S3, if needed
- `TOIL_S3_USE_SSL`: enable or disable the usage of SSL for connecting to S3 (True by default)

Examples:

```
TOIL_S3_HOST=127.0.0.1
TOIL_S3_PORT=9010
TOIL_S3_USE_SSL=False
```

7.2.5 Dashboard

Toil provides a dashboard for viewing the RAM and CPU usage of each node, the number of issued jobs of each type, the number of failed jobs, and the size of the jobs queue. To launch this dashboard for a toil workflow, include the `--metrics` flag in the toil script command. The dashboard can then be viewed in your browser at `localhost:3000` while connected to the leader node through `toil ssh-cluster`. On AWS, the dashboard keeps track of every node in the cluster to monitor CPU and RAM usage, but it can also be used while running a workflow on a single machine. The dashboard uses Grafana as the front end for displaying real-time plots, and Prometheus for tracking metrics exported by toil. In order to use the dashboard for a non-released toil version, you will have to build the containers locally with `make docker`, since the `prometheus`, `grafana`, and `mtail` containers used in the dashboard are tied to a specific toil version.

7.3 Running in Google Compute Engine (GCE)

Toil supports a provisioner with Google, and a *Google Job Store*. To get started, follow instructions for *Preparing your Google environment*.

7.3.1 Preparing your Google environment

Toil supports using the [Google Cloud Platform](#). Setting this up is easy!

1. Make sure that the `google` extra (*Installing Toil with Extra Features*) is installed
2. Follow [Google's Instructions](#) to download credentials and set the `GOOGLE_APPLICATION_CREDENTIALS` environment variable
3. Create a new ssh key with the proper format. To create a new ssh key run the command

```
$ ssh-keygen -t rsa -f ~/.ssh/id_rsa -C [USERNAME]
```

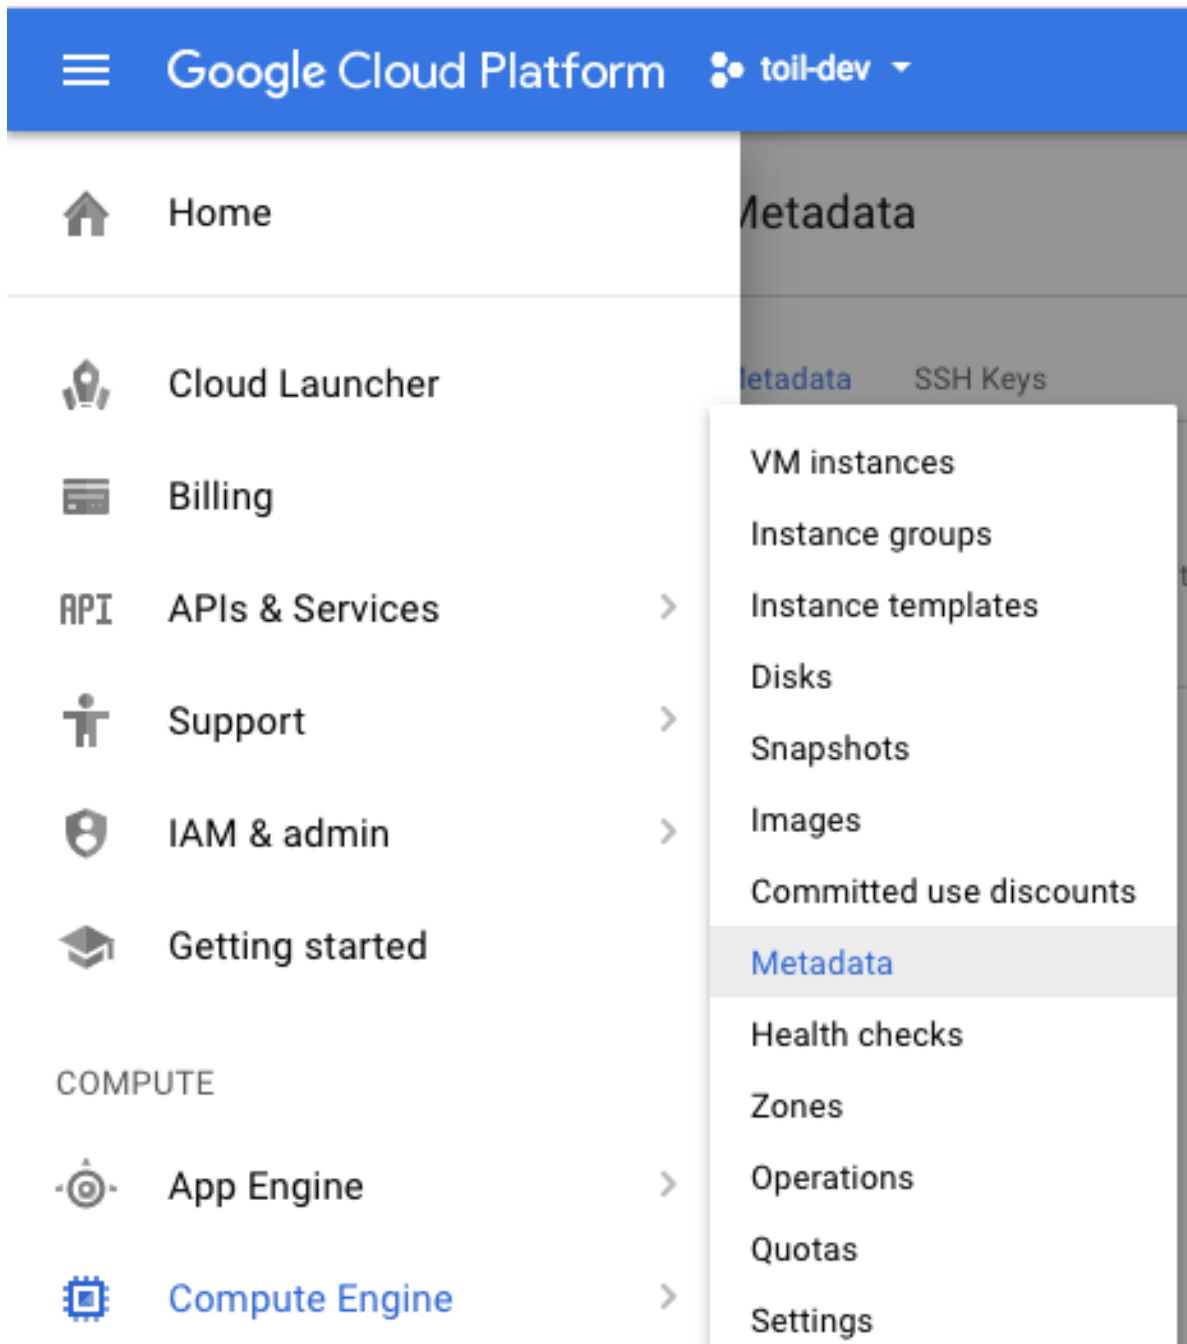
where [USERNAME] is something like `jane@example.com`. Make sure to leave your password blank.

Warning: This command could overwrite an old ssh key you may be using. If you have an existing ssh key you would like to use, it will need to be called `id_rsa` and it needs to have no password set.

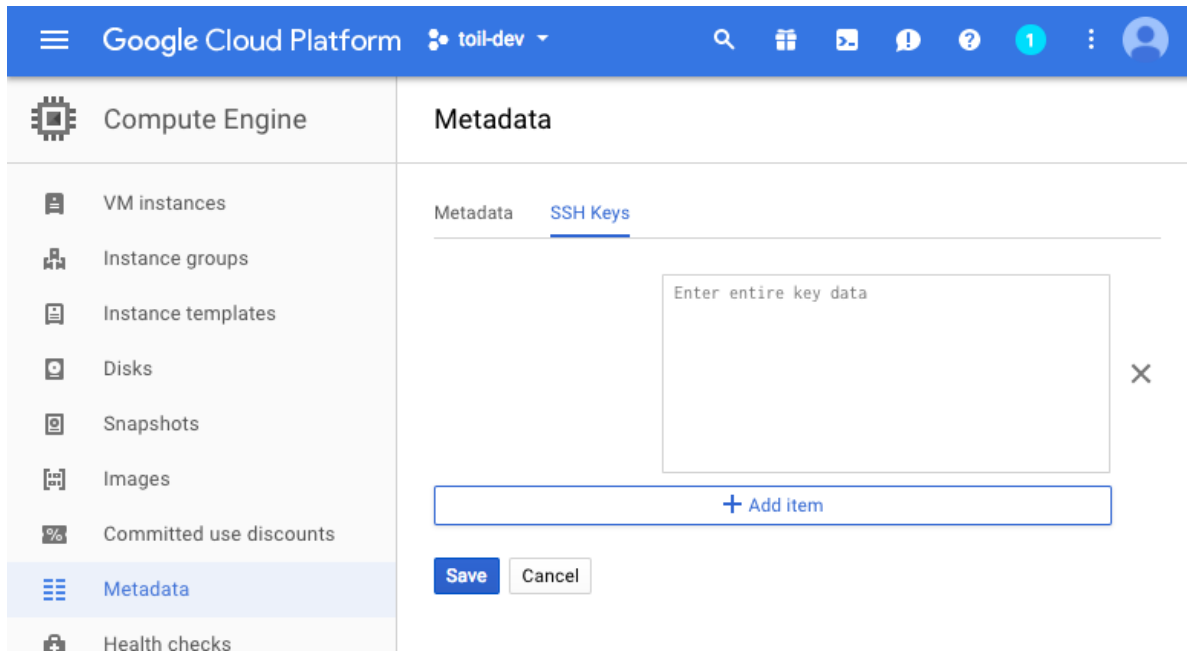
Make sure only you can read the SSH keys:

```
$ chmod 400 ~/.ssh/id_rsa ~/.ssh/id_rsa.pub
```

4. Add your newly formatted public key to Google. To do this, log into your Google Cloud account and go to [metadata](#) section under the Compute tab.



Near the top of the screen click on 'SSH Keys', then edit, add item, and paste the key. Then save:



For more details look at Google’s instructions for [adding SSH keys](#).

7.3.2 Google Job Store

To use the Google Job Store you will need to set the `GOOGLE_APPLICATION_CREDENTIALS` environment variable by following [Google’s instructions](#).

Then to run the sort example with the Google job store you would type

```
$ python sort.py google:my-project-id:my-google-sort-jobstore
```

7.3.3 Running a Workflow with Autoscaling

Warning: Google Autoscaling is in beta!

The steps to run a GCE workflow are similar to those of AWS (*Running a Workflow with Autoscaling*), except you will need to explicitly specify the `--provisioner gce` option which otherwise defaults to `aws`.

1. Download `sort.py`
2. Launch the leader node in GCE using the *Launch-Cluster Command* command:

```
(venv) $ toil launch-cluster <CLUSTER-NAME> \
        --provisioner gce \
        --leaderNodeType n1-standard-1 \
        --keyPairName <SSH-KEYNAME> \
        --zone us-west1-a
```

Where `<SSH-KEYNAME>` is the first part of `[USERNAME]` used when setting up your ssh key. For example if `[USERNAME]` was `jane@example.com`, `<SSH-KEYNAME>` should be `jane`.

The `--keyPairName` option is for an SSH key that was added to the Google account. If your ssh key [USERNAME] was `jane@example.com`, then your key pair name will be just `jane`.

3. Upload the sort example and ssh into the leader:

```
(venv) $ toil rsync-cluster --provisioner gce <CLUSTER-NAME> sort.py :/root
(venv) $ toil ssh-cluster --provisioner gce <CLUSTER-NAME>
```

4. Run the workflow:

```
$ python /root/sort.py google:<PROJECT-ID>:<JOBSTORE-NAME> \
  --provisioner gce \
  --batchSystem mesos \
  --nodeTypes n1-standard-2 \
  --maxNodes 2
```

5. Clean up:

```
$ exit # this exits the ssh from the leader node
(venv) $ toil destroy-cluster --provisioner gce <CLUSTER-NAME>
```

7.4 Cluster Utilities

There are several utilities used for starting and managing a Toil cluster using the AWS provisioner. They are installed via the `[aws]` or `[google]` extra. For installation details see *Toil Provisioner*. The cluster utilities are used for *Running in AWS* and are comprised of `toil launch-cluster`, `toil rsync-cluster`, `toil ssh-cluster`, and `toil destroy-cluster` entry points.

Cluster commands specific to `toil` are:

`status` — Reports runtime and resource usage for all jobs in a specified jobstore (workflow must have originally been run using the `--stats` option).

`stats` — Inspects a job store to see which jobs have failed, run successfully, etc.

`destroy-cluster` — For autoscaling. Terminates the specified cluster and associated resources.

`launch-cluster` — For autoscaling. This is used to launch a toil leader instance with the specified provisioner.

`rsync-cluster` — For autoscaling. Used to transfer files to a cluster launched with `toil launch-cluster`.

`ssh-cluster` — SSHs into the toil appliance container running on the leader of the cluster.

`clean` — Delete the job store used by a previous Toil workflow invocation.

`kill` — Kills any running jobs in a rogue toil.

For information on a specific utility run:

```
toil launch-cluster --help
```

for a full list of its options and functionality.

The cluster utilities can be used for *Running in Google Compute Engine (GCE)* and *Running in AWS*.

Tip: By default, all of the cluster utilities expect to be running on AWS. To run with Google you will need to specify the `--provisioner gce` option for each utility.

Note: Boto must be [configured](#) with AWS credentials before using cluster utilities.

Running in Google Compute Engine (GCE) contains instructions for

7.5 Stats Command

To use the stats command, a workflow must first be run using the `--stats` option. Using this command makes certain that toil does not delete the job store, no matter what other options are specified (i.e. normally the option `--clean=always` would delete the job, but `--stats` will override this).

An example of this would be running the following:

```
python discoverfiles.py file:my-jobstore --stats
```

Where `discoverfiles.py` is the following:

```
import os
import subprocess

from toil.common import Toil
from toil.job import Job

class discoverFiles(Job):
    """Views files at a specified path using ls."""
    def __init__(self, path, *args, **kwargs):
        self.path = path
        super(discoverFiles, self).__init__(*args, **kwargs)

    def run(self, fileStore):
        if os.path.exists(self.path):
            subprocess.check_call(["ls", self.path])

def main():
    options = Job.Runner.getDefaultArgumentParser().parse_args()
    options.clean = "always"

    job1 = discoverFiles(path="/sys/", displayName='sysFiles')
    job2 = discoverFiles(path=os.path.expanduser("~"), displayName='userFiles')
    job3 = discoverFiles(path="/tmp/")

    job1.addChild(job2)
    job2.addChild(job3)

    with Toil(options) as toil:
        if not toil.options.restart:
            toil.start(job1)
        else:
            toil.restart()
```

(continues on next page)

(continued from previous page)

```
if __name__ == '__main__':
    main()
```

Notice the `displayName` key, which can rename a job, giving it an alias when it is finally displayed in stats. Running this workflow file should record three job names: `sysFiles (job1)`, `userFiles (job2)`, and `discoverFiles (job3)`. To see the runtime and resources used for each job when it was run, type

```
toil stats file:my-jobstore
```

This should output the following:

```
Batch System: singleMachine
Default Cores: 1 Default Memory: 2097152K
Max Cores: 9.22337e+18
Total Clock: 0.56 Total Runtime: 1.01
Worker
  Count |
  ↪ Clock |
  ↪ Memory
    n | min med* ave max total | min med ave
  ↪max total | min med ave max total | min med ave
  ↪ max total
    1 | 0.14 0.14 0.14 0.14 0.14 | 0.13 0.13 0.13 0.
  ↪13 0.13 | 0.01 0.01 0.01 0.01 0.01 | 76K 76K 76K
  ↪ 76K 76K
Job
Worker Jobs | min med ave max
          | 3 3 3 3
  Count |
  ↪ Clock |
  ↪ Memory
    n | min med* ave max total | min med ave
  ↪max total | min med ave max total | min med ave
  ↪ max total
    3 | 0.01 0.06 0.05 0.07 0.14 | 0.00 0.06 0.04 0.
  ↪07 0.12 | 0.00 0.01 0.00 0.01 0.01 | 76K 76K 76K
  ↪ 76K 229K
sysFiles
  Count |
  ↪ Clock |
  ↪ Memory
    n | min med* ave max total | min med ave
  ↪max total | min med ave max total | min med ave
  ↪ max total
    1 | 0.01 0.01 0.01 0.01 0.01 | 0.00 0.00 0.00 0.
  ↪00 0.00 | 0.01 0.01 0.01 0.01 0.01 | 76K 76K 76K
  ↪ 76K 76K
userFiles
  Count |
  ↪ Clock |
  ↪ Memory
    n | min med* ave max total | min med ave
  ↪max total | min med ave max total | min med ave
  ↪ max total
    1 | 0.06 0.06 0.06 0.06 0.06 | 0.06 0.06 0.06 0.
  ↪06 0.06 | 0.01 0.01 0.01 0.01 0.01 | 76K 76K 76K
  ↪ 76K 76K
(continues on next page)
```

(continued from previous page)

[illegible]

Once we're done, we can clean up the job store by running

```
toil clean file:my-jobstore
```

7.6 Status Command

Continuing the example from the stats section above, if we ran our workflow with the command

```
python discoverfiles.py file:my-jobstore --stats
```

We could interrogate our jobstore with the status command, for example:

```
toil status file:my-jobstore
```

If the run was successful, this would not return much valuable information, something like

```
2018-01-11 19:31:29,739 - toil.lib.bioio - INFO - Root logger is at level 'INFO',
↳ 'toil' logger at level 'INFO'.
2018-01-11 19:31:29,740 - toil.utils.toilStatus - INFO - Parsed arguments
2018-01-11 19:31:29,740 - toil.utils.toilStatus - INFO - Checking if we have files_
↳ for Toil
The root job of the job store is absent, the workflow completed successfully.
```

Otherwise, the `status` command should return the following:

There are x unfinished jobs, y parent jobs with children, z jobs with services, a services, and b totally failed jobs currently in c .

7.7 Clean Command

If a Toil pipeline didn't finish successfully, or was run using `--clean=always` or `--stats`, the job store will exist until it is deleted. `toil clean <jobStore>` ensures that all artifacts associated with a job store are removed. This is particularly useful for deleting AWS job stores, which reserves an SDB domain as well as an S3 bucket.

The deletion of the job store can be modified by the `--clean` argument, and may be set to `always`, `onError`, `never`, or `onSuccess` (default).

Temporary directories where jobs are running can also be saved from deletion using the `--cleanWorkDir`, which has the same options as `--clean`. This option should only be run when debugging, as intermediate jobs will fill up disk space.

7.8 Launch-Cluster Command

Running `toil launch-cluster` starts up a leader for a cluster. Workers can be added to the initial cluster by specifying the `-w` option. An example would be

```
$ toil launch-cluster my-cluster \
  --leaderNodeType t2.small -z us-west-2a \
  --keyPairName your-AWS-key-pair-name \
  --nodeTypes m3.large,t2.micro -w 1,4
```

Options are listed below. These can also be displayed by running

```
$ toil launch-cluster --help
```

`launch-cluster`'s main positional argument is the `clusterName`. This is simply the name of your cluster. If it does not exist yet, Toil will create it for you.

Launch-Cluster Options

- help** -h also accepted. Displays this help menu.
- tempDirRoot TEMPDIRROOT** Path to the temporary directory where all temp files are created, by default uses the current working directory as the base.
- version** Display version.
- provisioner CLOUDPROVIDER** -p CLOUDPROVIDER also accepted. The provisioner for cluster auto-scaling. Both AWS and GCE are currently supported.
- zone ZONE** -z ZONE also accepted. The availability zone of the leader. This parameter can also be set via the `TOIL_AWS_ZONE` or `TOIL_GCE_ZONE` environment variables, or by the `ec2_region_name` parameter in your `.boto` file if using AWS, or derived from the instance metadata if using this utility on an existing EC2 instance.
- leaderNodeType LEADERNODETYPE** Non-preemptable node type to use for the cluster leader.
- keyPairName KEYPAIRNAME** The name of the AWS or ssh key pair to include on the instance.
- boto BOTOPATH** The path to the boto credentials directory. This is transferred to all nodes in order to access the AWS jobStore from non-AWS instances.
- tag KEYVALUE** KEYVALUE is specified as `KEY=VALUE`. `-t KEY=VALUE` also accepted. Tags are added to the AWS cluster for this node and all of its children. Tags are of the form: `-t key1=value1 --tag key2=value2`. Multiple tags are allowed and each tag needs its own flag. By default the cluster is tagged with: { "Name": clusterName, "Owner": IAM username }.
- vpcSubnet VPCTSUBNET** VPC subnet ID to launch cluster in. Uses default subnet if not specified. This subnet needs to have auto assign IPs turned on.
- nodeTypes NODETYPES** Comma-separated list of node types to create while launching the leader. The syntax for each node type depends on the provisioner used. For the AWS provisioner this is the name of an EC2 instance type followed by a colon and the price in dollars to bid for a

spot instance, for example 'c3.8xlarge:0.42'. Must also provide the `--workers` argument to specify how many workers of each node type to create.

--workers WORKERS `-w WORKERS` also accepted. Comma-separated list of the number of workers of each node type to launch alongside the leader when the cluster is created. This can be useful if running toil without auto-scaling but with need of more hardware support.

--leaderStorage LEADERSTORAGE Specify the size (in gigabytes) of the root volume for the leader instance. This is an EBS volume.

--nodeStorage NODESTORAGE Specify the size (in gigabytes) of the root volume for any worker instances created when using the `-w` flag. This is an EBS volume.

--nodeStorageOverrides NODESTORAGEOVERRIDES Comma-separated list of `nodeType:nodeStorage` that are used to override the default value from `--nodeStorage` for the specified `nodeType(s)`. This is useful for heterogeneous jobs where some tasks require much more disk than others.

Logging Options

--logOff Same as `--logCritical`.

--logCritical Turn on logging at level `CRITICAL` and above. (default is `INFO`)

--logError Turn on logging at level `ERROR` and above. (default is `INFO`)

--logWarning Turn on logging at level `WARNING` and above. (default is `INFO`)

--logInfo Turn on logging at level `INFO` and above. (default is `INFO`)

--logDebug Turn on logging at level `DEBUG` and above. (default is `INFO`)

--logLevel LOGLEVEL Log at given level (may be either `OFF` (or `CRITICAL`), `ERROR`, `WARN` (or `WARNING`), `INFO` or `DEBUG`). (default is `INFO`)

--logFile LOGFILE File to log in.

--rotatingLogging Turn on rotating logging, which prevents log files getting too big.

7.9 Ssh-Cluster Command

Toil provides the ability to ssh into the leader of the cluster. This can be done as follows:

```
$ toil ssh-cluster CLUSTER-NAME-HERE
```

This will open a shell on the Toil leader and is used to start an *Running a Workflow with Autoscaling* run. Issues with docker prevent using `screen` and `tmux` when sshing the cluster (The shell doesn't know that it is a TTY which prevents it from allocating a new screen session). This can be worked around via

```
$ script
$ screen
```

Simply running `screen` within `script` will get things working properly again.

Finally, you can execute remote commands with the following syntax:


```
$ toil ssh-cluster CLUSTER-NAME-HERE remoteCommand
```

It is not advised that you run your Toil workflow using remote execution like this unless a tool like `nohup` is used to ensure the process does not die if the SSH connection is interrupted.

For an example usage, see [Running a Workflow with Autoscaling](#).

7.10 Rsync-Cluster Command

The most frequent use case for the `rsync-cluster` utility is deploying your Toil script to the Toil leader. Note that the syntax is the same as traditional `rsync` with the exception of the hostname before the colon. This is not needed in `toil rsync-cluster` since the hostname is automatically determined by Toil.

Here is an example of its usage:

```
$ toil rsync-cluster CLUSTER-NAME-HERE \  
  ~/localFile :/remoteDestination
```

7.11 Destroy-Cluster Command

The `destroy-cluster` command is the advised way to get rid of any Toil cluster launched using the [Launch-Cluster Command](#). It ensures that all attached nodes, volumes, security groups, etc. are deleted. If a node or cluster is shut down using Amazon's online portal residual resources may still be in use in the background. To delete a cluster run

```
$ toil destroy-cluster CLUSTER-NAME-HERE
```

7.12 Kill Command

To kill all currently running jobs for a given jobstore, use the command

```
toil kill file:my-jobstore
```

HPC Environments

Toil is a flexible framework that can be leveraged in a variety of environments, including high-performance computing (HPC) environments. Toil provides support for a number of batch systems, including [Grid Engine](#), [Slurm](#), [Torque](#) and [LSF](#), which are popular schedulers used in these environments. Toil also supports [HTCondor](#), which is a popular scheduler for high-throughput computing (HTC). To use one of these batch systems specify the “`--batchSystem`” argument to the toil script.

Due to the cost and complexity of maintaining support for these schedulers we currently consider them to be “community supported”, that is the core development team does not regularly test or develop support for these systems. However, there are members of the Toil community currently deploying Toil in HPC environments and we welcome external contributions.

Developing the support of a new or existing batch system involves extending the abstract batch system class `toil.batchSystems.abstractBatchSystem.AbstractBatchSystem`.

8.1 Standard Output/Error from Batch System Jobs

Standard output and error from batch system jobs (except for the Parasol and Mesos batch systems) are redirected to files in the `toil-<workflowID>` directory created within the temporary directory specified by the `--workDir` option; see [Commandline Options](#). Each file is named as follows: `toil_job_<Toil job ID>_batch_<name of batch system>_<job ID from batch system>_<file description>.log`, where `<file description>` is `std_output` for standard output, and `std_error` for standard error. HTCondor will also write job event log files with `<file description> = job_events`.

If capturing standard output and error is desired, `--workDir` will generally need to be on a shared file system; otherwise if these are written to local temporary directories on each node (e.g. `/tmp`) Toil will not be able to retrieve them. Alternatively, the `--noStdOutErr` option forces Toil to discard all standard output and error from batch system jobs.

The Common Workflow Language (CWL) is an emerging standard for writing workflows that are portable across multiple workflow engines and platforms. Toil has full support for the CWL v1.0.1 specification.

9.1 Running CWL Locally

The `toil-cwl-runner` command provides cwl-parsing functionality using cwltool, and leverages the job-scheduling and batch system support of Toil.

To run in local batch mode, provide the CWL file and the input object file:

```
$ toil-cwl-runner example.cwl example-job.yml
```

For a simple example of CWL with Toil see *Running a basic CWL workflow*.

9.1.1 Note for macOS + Docker + Toil

When invoking CWL documents that make use of Docker containers if you see errors that look like

```
docker: Error response from daemon: Mounts denied:
The paths /var/...tmp are not shared from OS X and are not known to Docker.
```

you may need to add

```
export TMPDIR=/tmp/docker_tmp
```

either in your startup file (`.bashrc`) or add it manually in your shell before invoking toil.

9.2 Detailed Usage Instructions

Help information can be found by using this toil command:

```
$ toil-cwl-runner -h
```

A more detailed example shows how we can specify both Toil and cwltool arguments for our workflow:

```
$ toil-cwl-runner \  
  --singularity \  
  --jobStore my_jobStore \  
  --batchSystem lsf \  
  --workDir `pwd` \  
  --outdir `pwd` \  
  --logFile cwltoil.log \  
  --writeLogs `pwd` \  
  --logLevel DEBUG \  
  --retryCount 2 \  
  --maxLogFileSize 20000000000 \  
  --stats \  
  standard_bam_processing.cwl \  
  inputs.yaml
```

In this example, we set the following options, which are all passed to Toil:

`--singularity`: Specifies that all jobs with Docker format containers specified should be run using the Singularity container engine instead of the Docker container engine.

`--jobStore`: Path to a folder that already exists, which will contain the Toil jobstore and all related job-tracking information.

`--batchSystem`: Use the specified HPC or Cloud-based cluster platform.

`--workDir`: The directory where all temporary files will be created for the workflow. A subdirectory of this will be set as the `$TMPDIR` environment variable and this subdirectory can be referenced using the CWL parameter reference `$(runtime.tmpdir)` in CWL tools and workflows.

`--outdir`: Directory where final File and Directory outputs will be written. References to these and other output types will be in the JSON object printed to the stdout stream after workflow execution.

`--logFile`: Path to the main logfile with logs from all jobs.

`--writeLogs`: Directory where all job logs will be stored.

`--retryCount`: How many times to retry each Toil job.

`--maxLogFileSize`: Logs that get larger than this value will be truncated.

`--stats`: Save resources usages in json files that can be collected with the `toil stats` command after the workflow is done.

9.3 Running CWL in the Cloud

To run in cloud and HPC configurations, you may need to provide additional command line parameters to select and configure the batch system to use.

To run a CWL workflow in AWS with toil see [Running a CWL Workflow on AWS](#).

9.4 Running CWL within Toil Scripts

A CWL workflow can be run indirectly in a native Toil script. However, this is not the *standard* way to run CWL workflows with Toil and doing so comes at the cost of job efficiency. For some use cases, such as running one process on multiple files, it may be useful. For example, if you want to run a CWL workflow with 3 YML files specifying different samples inputs, it could look something like:

```
import os
import subprocess

from toil.common import Toil
from toil.job import Job

def initialize_jobs(job):
    job.fileStore.logToMaster('initialize_jobs')

def runQC(job, cwl_file, cwl_filename, yml_file, yml_filename, outputs_dir, output_
↳ num):
    job.fileStore.logToMaster("runQC")
    tempDir = job.fileStore.getLocalTempDir()

    cwl = job.fileStore.readGlobalFile(cwl_file, userPath=os.path.join(tempDir, cwl_
↳ filename))
    yml = job.fileStore.readGlobalFile(yml_file, userPath=os.path.join(tempDir, yml_
↳ filename))

    subprocess.check_call(["toil-cwl-runner", cwl, yml])

    output_filename = "output.txt"
    output_file = job.fileStore.writeGlobalFile(output_filename)
    job.fileStore.readGlobalFile(output_file, userPath=os.path.join(outputs_dir,
↳ "sample_" + output_num + "_" + output_filename))
    return output_file

if __name__ == "__main__":
    options = Job.Runner.getDefaultOptions("./toilWorkflowRun")
    options.logLevel = "INFO"
    options.clean = "always"
    with Toil(options) as toil:

        # specify the folder where the cwl and yml files live
        inputs_dir = os.path.join(os.path.dirname(os.path.abspath(__file__)),
↳ "cwlExampleFiles")
        # specify where you wish the outputs to be written
        outputs_dir = os.path.join(os.path.dirname(os.path.abspath(__file__)),
↳ "cwlExampleFiles")

        job0 = Job.wrapJobFn(initialize_jobs)

        cwl_filename = "hello.cwl"
        cwl_file = toil.importFile("file://" + os.path.abspath(os.path.join(inputs_
↳ dir, cwl_filename)))

        # add list of yml config inputs here or import and construct from file
```

(continues on next page)

(continued from previous page)

```

    yml_files = ["hello1.yml", "hello2.yml", "hello3.yml"]
    i = 0
    for yml in yml_files:
        i = i + 1
        yml_file = toil.importFile("file://" + os.path.abspath(os.path.
→ join(inputs_dir, yml)))
        yml_filename = yml
        job = Job.wrapJobFn(runQC, cwl_file, cwl_filename, yml_file, yml_filename,
→ outputs_dir, output_num=str(i))
        job0.addChild(job)

    toil.start(job0)

```

9.5 Toil & CWL Tips

See logs for just one job by using the full log file

This requires knowing the job's toil-generated ID, which can be found in the log files.

```
cat cwltoil.log | grep jobVM1fIs
```

Grep for full tool commands from toil logs

This gives you a more concise view of the commands being run (note that this information is only available from Toil when running with `-logDebug`).

```

pcregrep -M "[job .*\..cwl.*$\n(.*)\n.*$\n)*" cwltoil.log
#      ^allows for multiline matching

```

Find Bams that have been generated for specific step while pipeline is running:

```
find . | grep -P '^./out_tmpdir.*_MD\.bam$'
```

See what jobs have been run

```
cat log/cwltoil.log | grep -oP "[job .*\..cwl\]" | sort | uniq
```

or:

```
cat log/cwltoil.log | grep -i "issued job"
```

Get status of a workflow

```

$ toil status /home/johnsoni/TEST_RUNS_3/TEST_run/tmp/jobstore-09ae0acc-c800-11e8-
→ 9d09-70106fb1697e
<hostname> 2018-10-04 15:01:44,184 MainThread INFO toil.lib.bioio: Root logger is at
→ level 'INFO', 'toil' logger at level 'INFO'.
<hostname> 2018-10-04 15:01:44,185 MainThread INFO toil.utils.toilStatus: Parsed
→ arguments
<hostname> 2018-10-04 15:01:47,081 MainThread INFO toil.utils.toilStatus: Traversing
→ the job graph gathering jobs. This may take a couple of minutes.

Of the 286 jobs considered, there are 179 jobs with children, 107 jobs ready to run,
→ 0 zombie jobs, 0 jobs with services, 0 services, and 0 jobs with log files
→ currently in file:/home/user/jobstore-09ae0acc-c800-11e8-9d09-70106fb1697e.

```


Toil Stats

You can get run statistics broken down by CWL file. This only works once the workflow is finished:

```
$ toil stats /path/to/jobstore
```

The output will contain CPU, memory, and walltime information for all CWL job types:

```
<hostname> 2018-10-15 12:06:19,003 MainThread INFO toil.lib.bioio: Root logger is at
↳level 'INFO', 'toil' logger at level 'INFO'.
<hostname> 2018-10-15 12:06:19,004 MainThread INFO toil.utils.toilStats: Parsed
↳arguments
<hostname> 2018-10-15 12:06:19,004 MainThread INFO toil.utils.toilStats: Checking if
↳we have files for toil
<hostname> 2018-10-15 12:06:19,004 MainThread INFO toil.utils.toilStats: Checked
↳arguments
Batch System: lsf
Default Cores: 1 Default Memory: 10485760K
Max Cores: 9.22337e+18
Total Clock: 106608.01 Total Runtime: 86634.11
Worker
  Count |
  ↳      Clock |
  ↳
  ↳      n |      min      Memory
  ↳      max total |      med* ave      max      total |      min      med      ave
  ↳      med ave      max      total
  ↳      1659 |      0.00      0.80 264.87 12595.59 439424.40 |      0.00      0.46      449.05
  ↳42240.74 744968.80 | -35336.69      0.16 -184.17 4230.65 -305544.39 |      48K
  ↳ 223K 1020K 40235K 1692300K
Job
  Worker Jobs |      min      med      ave      max
  ↳           |      1077      1077      1077      1077
  Count |
  ↳      Clock |
  ↳
  ↳      n |      min      Memory
  ↳      max total |      med* ave      max      total |      min      med      ave
  ↳      med ave      max      total
  ↳      1077 |      0.04      1.18 407.06 12593.43 438404.73 |      0.01      0.28      691.17
  ↳42240.35 744394.14 | -35336.83      0.27 -284.11 4230.49 -305989.41 |      135K
  ↳ 268K 1633K 40235K 1759734K
ResolveIndirect
  Count |
  ↳      Clock |
  ↳
  ↳      n |      min      Memory
  ↳      max total |      med* ave      max      total |      min      med      ave
  ↳      med ave      max      total
  ↳      205 |      0.04      0.07 0.16      2.29      31.95 |      0.01      0.02      0.02
  ↳ 0.14      3.60 |      0.02 0.05      0.14      2.28      28.35 |      190K
  ↳266K 256K 314K 52487K
CWLgather
  Count |
  ↳      Clock |
  ↳
  ↳      n |      min      Memory
  ↳      max total |      med* ave      max      total |      min      med      ave
  ↳      med ave      max      total
  ↳      205 |      0.04      0.07 0.16      2.29      31.95 |      0.01      0.02      0.02
  ↳ 0.14      3.60 |      0.02 0.05      0.14      2.28      28.35 |      190K
  ↳266K 256K 314K 52487K
```

(continues on next page)

(continued from previous page)

(continued from previous page)											
40		0.05	0.17	0.29	1.90	11.62		0.01	0.02	0.02	↳
↳ 0.05		0.80		0.03	0.14	0.27	1.88		10.82		188K
↳ 265K	250K	316K	10039K								
CWLWorkflow											
Count		Time*				Wait					
↳		Clock						Wait			
↳		Memory									
n		min	med*	ave	max	total		min	med	ave	
↳ max		total		min	med	ave	max		total		min
↳ med		ave	max	total							
205		0.09	0.40	0.98	13.70	200.82		0.04	0.15	0.16	
↳ 1.08	31.78		0.04	0.26	0.82	12.62		169.04		190K	
↳ 270K	257K	316K	52826K								
file:///home/johnsoni/pipeline_0.0.39/ACCESS-Pipeline/cwl_tools/expression_tools/											
↳ group_waltz_files.cwl											
Count		Time*				Wait					
↳		Clock						Wait			
↳		Memory									
n		min	med*	ave	max	total		min	med	ave	
↳ max		total		min	med	ave	max		total		min
↳ med		ave	max	total							
99		0.29	0.49	0.59	2.50	58.11		0.14	0.26	0.29	
↳ 1.04	28.95		0.14	0.22	0.29	1.48		29.16		135K	
↳ 135K	135K	136K	13459K								
file:///home/johnsoni/pipeline_0.0.39/ACCESS-Pipeline/cwl_tools/expression_tools/											
↳ make_sample_output_dirs.cwl											
Count		Time*				Wait					
↳		Clock						Wait			
↳		Memory									
n		min	med*	ave	max	total		min	med	ave	
↳ max		total		min	med	ave	max		total		min
↳ med		ave	max	total							
11		0.34	0.52	0.74	2.63	8.18		0.20	0.30	0.41	
↳ 1.17	4.54		0.14	0.20	0.33	1.45		3.65		136K	
↳ 136K	136K	136K	1496K								
file:///home/johnsoni/pipeline_0.0.39/ACCESS-Pipeline/cwl_tools/expression_tools/											
↳ consolidate_files.cwl											
Count		Time*				Wait					
↳		Clock						Wait			
↳		Memory									
n		min	med*	ave	max	total		min	med	ave	
↳ max		total		min	med	ave	max		total		min
↳ med		ave	max	total							
8		0.31	0.59	0.71	1.80	5.69		0.18	0.35	0.37	
↳ 0.63	2.94		0.13	0.27	0.34	1.17		2.75		136K	
↳ 136K	136K	136K	1091K								
file:///home/johnsoni/pipeline_0.0.39/ACCESS-Pipeline/cwl_tools/bwa-mem/bwa-mem.cwl											
Count		Time*				Wait					
↳		Clock						Wait			
↳		Memory									
n		min	med*	ave	max	total		min	med	ave	
↳ max		total		min	med	ave	max		total		min
↳ med		ave	max	total							
22		895.76	3098.13	3587.34	12593.43	78921.51		2127.02	7910.31	8123.06	
↳ 16959.13	178707.34		-11049.84	-3827.96	-4535.72	19.49		-99785.83		5659K	
↳ 5950K	5854K	6128K	128807K								

Understanding toil log files

There is a *worker_log.txt* file for each job, this file is written to while the job is running, and deleted after the job finishes. The contents are printed to the main log file and transferred to a log file in the *-logDir* folder once the job is completed successfully.

The new log file will be named something like:

```
file:<path to cwl tool>.cwl_<job ID>.log  
  
file:---home-johnsoni-pipeline_1.1.14-ACCESS--Pipeline-cwl_tools-marianas-  
↪ProcessLoopUMIFastq.cwl_I-O-jobfGsQQw000.log
```

This is the toil job command with spaces replaced by dashes.

Support is still in the alpha phase and should be able to handle basic wdl files. See the specification below for more details.

10.1 How to Run a WDL file in Toil

Recommended best practice when running wdl files is to first use the Broad's wdltool for syntax validation and generating the needed json input file. Full documentation can be found on the [repository](#), and a precompiled jar binary can be downloaded here: [wdltool](#) (this requires [java7](#)).

That means two steps. First, make sure your wdl file is valid and devoid of syntax errors by running

```
java -jar wdltool.jar validate example_wdlfile.wdl
```

Second, generate a complementary json file if your wdl file needs one. This json will contain keys for every necessary input that your wdl file needs to run:

```
java -jar wdltool.jar inputs example_wdlfile.wdl
```

When this json template is generated, open the file, and fill in values as necessary by hand. WDL files all require json files to accompany them. If no variable inputs are needed, a json file containing only '{ }' may be required.

Once a wdl file is validated and has an appropriate json file, workflows can be run in toil using:

```
toil-wdl-runner example_wdlfile.wdl example_jsonfile.json
```

See options below for more parameters.

10.2 ENCODE Example from ENCODE-DCC

To follow this example, you will need docker installed. The original workflow can be found here: <https://github.com/ENCODE-DCC/pipeline-container>

We’ve included the wdl file and data files in the toil repository needed to run this example. First, download the example `code` and unzip. The file needed is “testENCODE/encode_mapping_workflow.wdl”.

Next, use `wdltool` (this requires `java7`) to validate this file:

```
java -jar wdltool.jar validate encode_mapping_workflow.wdl
```

Next, use `wdltool` to generate a json file for this wdl file:

```
java -jar wdltool.jar inputs encode_mapping_workflow.wdl
```

This json file once opened should look like this:

```
{
"encode_mapping_workflow.fastqs": "Array[File]",
"encode_mapping_workflow.trimming_parameter": "String",
"encode_mapping_workflow.reference": "File"
}
```

The `trimming_parameter` should be set to ‘native’. Download the example `code` and unzip. Inside are two data files required for the run

```
ENCODE_data/reference/GRCh38_chr21_bwa.tar.gz    ENCODE_data/ENCFF000VOL_chr21.fq.gz
```

Editing the json to include these as inputs, the json should now look something like this:

```
{
"encode_mapping_workflow.fastqs": ["/path/to/unzipped/ENCODE_data/ENCFF000VOL_chr21.fq.gz"],
"encode_mapping_workflow.trimming_parameter": "native",
"encode_mapping_workflow.reference": "/path/to/unzipped/ENCODE_data/reference/GRCh38_chr21_bwa.tar.gz"
}
```

The wdl and json files can now be run using the command

```
toil-wdl-runner encode_mapping_workflow.wdl encode_mapping_workflow.json
```

This should deposit the output files in the user’s current working directory (to change this, specify a new directory with the ‘-o’ option).

10.3 GATK Examples from the Broad

Simple examples of WDL can be found on the Broad’s website as tutorials: <https://software.broadinstitute.org/wdl/documentation/topic?name=wdl-tutorials>.

One can follow along with these tutorials, write their own wdl files following the directions and run them using either `cromwell` or `toil`. For example, in tutorial 1, if you’ve followed along and named your wdl file ‘helloHaplotype-Call.wdl’, then once you’ve validated your wdl file using `wdltool` (this requires `java7`) using

```
java -jar wdltool.jar validate helloHaplotypeCaller.wdl
```

and generated a json file (and subsequently typed in appropriate filepaths* and variables) using

```
java -jar wdltool.jar inputs helloHaplotypeCaller.wdl
```

- Absolute filepath inputs are recommended for local testing.

then the wdl script can be run using

```
toil-wdl-runner helloHaplotypeCaller.wdl helloHaplotypeCaller_inputs.json
```

10.4 toilwdl.py Options

‘-o’ or ‘-outdir’: Specifies the output folder, and defaults to the current working directory if not specified by the user.

‘-dev_mode’: Creates “AST.out”, which holds a printed AST of the wdl file and “mappings.out”, which holds the printed task, workflow, csv, and tsv dictionaries generated by the parser. Also saves the compiled toil python workflow file for debugging.

Any number of arbitrary options may also be specified. These options will not be parsed immediately, but passed down as toil options once the wdl/json files are processed. For valid toil options, see the documentation: <http://toil.readthedocs.io/en/latest/running/cliOptions.html>

10.5 Running WDL within Toil Scripts

Note: A cromwell.jar file is needed in order to run a WDL workflow.

A WDL workflow can be run indirectly in a native Toil script. However, this is not the *standard* way to run WDL workflows with Toil and doing so comes at the cost of job efficiency. For some use cases, such as running one process on multiple files, it may be useful. For example, if you want to run a WDL workflow with 3 JSON files specifying different samples inputs, it could look something like:

```
import os
import subprocess

from toil.common import Toil
from toil.job import Job

def initialize_jobs(job):
    job.fileStore.logToMaster("initialize_jobs")

def runQC(job, wdl_file, wdl_filename, json_file, json_filename, outputs_dir, jar_loc,
    ↪output_num):
    job.fileStore.logToMaster("runQC")
    tempDir = job.fileStore.getLocalTempDir()

    wdl = job.fileStore.readGlobalFile(wdl_file, userPath=os.path.join(tempDir, wdl_
    ↪filename))
    json = job.fileStore.readGlobalFile(json_file, userPath=os.path.join(tempDir,
    ↪json_filename))

    subprocess.check_call(["java", "-jar", jar_loc, "run", wdl, "--inputs", json])

    output_filename = "output.txt"
    output_file = job.fileStore.writeGlobalFile(outputs_dir + output_filename)
    job.fileStore.readGlobalFile(output_file, userPath=os.path.join(outputs_dir,
    ↪"sample_" + output_num + "_" + output_filename))
    return output_file

if __name__ == "__main__":
    options = Job.Runner.getDefaultOptions("./toilWorkflowRun")
```

(continues on next page)

(continued from previous page)

```

options.logLevel = "INFO"
options.clean = "always"

with Toil(options) as toil:

    # specify the folder where the wdl and json files live
    inputs_dir = "wdlExampleFiles/"
    # specify where you wish the outputs to be written
    outputs_dir = "wdlExampleFiles/"
    # specify the location of your cromwell jar
    jar_loc = os.path.abspath("wdlExampleFiles/cromwell-35.jar")

    job0 = Job.wrapJobFn(initialize_jobs)

    wdl_filename = "hello.wdl"
    wdl_file = toil.importFile("file://" + os.path.abspath(os.path.join(inputs_
→dir, wdl_filename)))

    # add list of yml config inputs here or import and construct from file
    json_files = ["hello1.json", "hello2.json", "hello3.json"]
    i = 0
    for json in json_files:
        i = i + 1
        json_file = toil.importFile("file://" + os.path.join(inputs_dir, json))
        json_filename = json
        job = Job.wrapJobFn(runQC, wdl_file, wdl_filename, json_file, json_
→filename, outputs_dir, jar_loc, output_num=str(i))
        job0.addChild(job)

    toil.start(job0)

```

10.6 WDL Specifications

WDL language specifications can be found here: <https://github.com/broadinstitute/wdl/blob/develop/SPEC.md>

Implementing support for more features is currently underway, but a basic roadmap so far is:

CURRENTLY IMPLEMENTED:

- Scatter
- Many Built-In Functions
- Docker Calls
- Handles Priority, and Output File Wrangling
- Currently Handles Primitives and Arrays

TO BE IMPLEMENTED:

- Integrate Cloud Autoscaling Capacity More Robustly
- WDL Files That “Import” Other WDL Files (Including URI Handling for ‘<http://>’ and ‘<https://>’)

Developing a Workflow

This tutorial walks through the features of Toil necessary for developing a workflow using the Toil Python API.

Note: “script” and “workflow” will be used interchangeably

11.1 Scripting Quick Start

To begin, consider this short toil script which illustrates defining a workflow:

```
from toil.common import Toil
from toil.job import Job

def helloWorld(message, memory="2G", cores=2, disk="3G"):
    return "Hello, world!, here's a message: %s" % message

if __name__ == "__main__":
    options = Job.Runner.getDefaultOptions("./toilWorkflowRun")
    options.logLevel = "OFF"
    options.clean = "always"

    hello_job = Job.wrapFn(helloWorld, "Woot")

    with Toil(options) as toil:
        print(toil.start(hello_job)) #Prints Hello, world!, ...
```

The workflow consists of a single job. The resource requirements for that job are (optionally) specified by keyword arguments (memory, cores, disk). The script is run using `toil.job.Job.Runner.getDefaultOptions()`. Below we explain the components of this code in detail.

11.2 Job Basics

The atomic unit of work in a Toil workflow is a `Job`. User scripts inherit from this base class to define units of work. For example, here is a more long-winded class-based version of the job in the quick start example:

```
from toil.job import Job

class HelloWorld(Job):
    def __init__(self, message):
        Job.__init__(self, memory="2G", cores=2, disk="3G")
        self.message = message

    def run(self, fileStore):
        return "Hello, world! Here's a message: %s" % self.message
```

In the example a class, `HelloWorld`, is defined. The constructor requests 2 gigabytes of memory, 2 cores and 3 gigabytes of local disk to complete the work.

The `toil.job.Job.run()` method is the function the user overrides to get work done. Here it just logs a message using `toil.job.Job.log()`, which will be registered in the log output of the leader process of the workflow.

11.3 Invoking a Workflow

We can add to the previous example to turn it into a complete workflow by adding the necessary function calls to create an instance of `HelloWorld` and to run this as a workflow containing a single job. This uses the `toil.job.Job.Runner` class, which is used to start and resume Toil workflows. For example:

```
from toil.common import Toil
from toil.job import Job

class HelloWorld(Job):
    def __init__(self, message):
        Job.__init__(self, memory="2G", cores=2, disk="3G")
        self.message = message

    def run(self, fileStore):
        return "Hello, world!, here's a message: %s" % self.message

if __name__ == "__main__":
    options = Job.Runner.getDefaultOptions("./toilWorkflowRun")
    options.logLevel = "OFF"
    options.clean = "always"

    hello_job = HelloWorld("Woot")

    with Toil(options) as toil:
        print(toil.start(hello_job))
```

Note: Do not include a `.` in the name of your python script (besides `.py` at the end). This is to allow toil to import the types and functions defined in your file while starting a new process.

Alternatively, the more powerful `toil.common.Toil` class can be used to run and resume workflows. It is used as a context manager and allows for preliminary setup, such as staging of files into the job store on the leader node. An

instance of the class is initialized by specifying an options object. The actual workflow is then invoked by calling the `toil.common.Toil.start()` method, passing the root job of the workflow, or, if a workflow is being restarted, `toil.common.Toil.restart()` should be used. Note that the context manager should have explicit if else branches addressing restart and non restart cases. The boolean value for these if else blocks is `toil.options.restart`.

For example:

```
from toil.common import Toil
from toil.job import Job

class HelloWorld(Job):
    def __init__(self, message):
        Job.__init__(self, memory="2G", cores=2, disk="3G")
        self.message = message

    def run(self, fileStore):
        self.log("Hello, world!, I have a message: {}".format(self.message))

if __name__ == "__main__":
    options = Job.Runner.getDefaultOptions("./toilWorkflowRun")
    options.logLevel = "INFO"
    options.clean = "always"

    with Toil(options) as toil:
        if not toil.options.restart:
            job = HelloWorld("Woot!")
            toil.start(job)
        else:
            toil.restart()
```

The call to `toil.job.Job.Runner.getDefaultOptions()` creates a set of default options for the workflow. The only argument is a description of how to store the workflow's state in what we call a *job-store*. Here the job-store is contained in a directory within the current working directory called "toilWorkflowRun". Alternatively this string can encode other ways to store the necessary state, e.g. an S3 bucket object store location. By default the job-store is deleted if the workflow completes successfully.

The workflow is executed in the final line, which creates an instance of `HelloWorld` and runs it as a workflow. Note all Toil workflows start from a single starting job, referred to as the *root* job. The return value of the root job is returned as the result of the completed workflow (see promises below to see how this is a useful feature!).

11.4 Specifying Commandline Arguments

To allow command line control of the options we can use the `toil.job.Job.Runner.getDefaultArgumentParser()` method to create a `argparse.ArgumentParser` object which can be used to parse command line options for a Toil script. For example:

```
from toil.common import Toil
from toil.job import Job

class HelloWorld(Job):
    def __init__(self, message):
        Job.__init__(self, memory="2G", cores=2, disk="3G")
        self.message = message
```

(continues on next page)

(continued from previous page)

```

def run(self, fileStore):
    return "Hello, world!, here's a message: %s" % self.message

if __name__ == "__main__":
    parser = Job.Runner.getDefaultArgumentParser()
    options = parser.parse_args()
    options.logLevel = "OFF"
    options.clean = "always"

    hello_job = HelloWorld("Woot")

    with Toil(options) as toil:
        print(toil.start(hello_job))

```

Creates a fully fledged script with all the options Toil exposed as command line arguments. Running this script with “-help” will print the full list of options.

Alternatively an existing `argparse.ArgumentParser` or `optparse.OptionParser` object can have Toil script command line options added to it with the `toil.job.Job.Runner.addToilOptions()` method.

11.5 Resuming a Workflow

In the event that a workflow fails, either because of programmatic error within the jobs being run, or because of node failure, the workflow can be resumed. Workflows can only not be reliably resumed if the job-store itself becomes corrupt.

Critical to resumption is that jobs can be rerun, even if they have apparently completed successfully. Put succinctly, a user defined job should not corrupt its input arguments. That way, regardless of node, network or leader failure the job can be restarted and the workflow resumed.

To resume a workflow specify the “restart” option in the options object passed to `toil.common.Toil.start()`. If node failures are expected it can also be useful to use the integer “retryCount” option, which will attempt to rerun a job retryCount number of times before marking it fully failed.

In the common scenario that a small subset of jobs fail (including retry attempts) within a workflow Toil will continue to run other jobs until it can do no more, at which point `toil.common.Toil.start()` will raise a `toil.leader.FailedJobsException` exception. Typically at this point the user can decide to fix the script and resume the workflow or delete the job-store manually and rerun the complete workflow.

11.6 Functions and Job Functions

Defining jobs by creating class definitions generally involves the boilerplate of creating a constructor. To avoid this the classes `toil.job.FunctionWrappingJob` and `toil.job.JobFunctionWrappingTarget` allow functions to be directly converted to jobs. For example, the quick start example (repeated here):

```

from toil.common import Toil
from toil.job import Job

def helloWorld(message, memory="2G", cores=2, disk="3G"):
    return "Hello, world!, here's a message: %s" % message

```

(continues on next page)

(continued from previous page)

```

if __name__=="__main__":
    options = Job.Runner.getDefaultOptions("./toilWorkflowRun")
    options.logLevel = "OFF"
    options.clean = "always"

    hello_job = Job.wrapFn(helloWorld, "Woot")

    with Toil(options) as toil:
        print(toil.start(hello_job)) #Prints Hello, world!, ...

```

Is equivalent to the previous example, but using a function to define the job.

The function call:

```
Job.wrapFn(helloWorld, "Woot")
```

Creates the instance of the `toil.job.FunctionWrappingTarget` that wraps the function.

The keyword arguments *memory*, *cores* and *disk* allow resource requirements to be specified as before. Even if they are not included as keyword arguments within a function header they can be passed as arguments when wrapping a function as a job and will be used to specify resource requirements.

We can also use the function wrapping syntax to a *job function*, a function whose first argument is a reference to the wrapping job. Just like a *self* argument in a class, this allows access to the methods of the wrapping job, see `toil.job.JobFunctionWrappingTarget`. For example:

```

from toil.common import Toil
from toil.job import Job

def helloWorld(job, message):
    job.log("Hello world, I have a message: {}".format(message))

if __name__=="__main__":
    options = Job.Runner.getDefaultOptions("./toilWorkflowRun")
    options.logLevel = "INFO"
    options.clean = "always"

    hello_job = Job.wrapJobFn(helloWorld, "Woot!")

    with Toil(options) as toil:
        toil.start(hello_job)

```

Here `helloWorld()` is a job function. It uses the `toil.job.Job.log()` to log a message that will be printed to the output console. Here the only subtle difference to note is the line:

```
hello_job = Job.wrapJobFn(helloWorld, "Woot")
```

Which uses the function `toil.job.Job.wrapJobFn()` to wrap the job function instead of `toil.job.Job.wrapFn()` which wraps a vanilla function.

11.7 Workflows with Multiple Jobs

A *parent* job can have *child* jobs and *follow-on* jobs. These relationships are specified by methods of the job class, e.g. `toil.job.Job.addChild()` and `toil.job.Job.addFollowOn()`.

Considering a set of jobs the nodes in a job graph and the child and follow-on relationships the directed edges of the graph, we say that a job B that is on a directed path of child/follow-on edges from a job A in the job graph is a *successor* of A, similarly A is a *predecessor* of B.

A parent job's child jobs are run directly after the parent job has completed, and in parallel. The follow-on jobs of a job are run after its child jobs and their successors have completed. They are also run in parallel. Follow-ons allow the easy specification of cleanup tasks that happen after a set of parallel child tasks. The following shows a simple example that uses the earlier `helloWorld()` job function:

```
from toil.common import Toil
from toil.job import Job

def helloWorld(job, message, memory="2G", cores=2, disk="3G"):
    job.log("Hello world, I have a message: {}".format(message))

if __name__=="__main__":
    options = Job.Runner.getDefaultOptions("./toilWorkflowRun")
    options.logLevel = "INFO"
    options.clean = "always"

    j1 = Job.wrapJobFn(helloWorld, "first")
    j2 = Job.wrapJobFn(helloWorld, "second or third")
    j3 = Job.wrapJobFn(helloWorld, "second or third")
    j4 = Job.wrapJobFn(helloWorld, "last")
    j1.addChild(j2)
    j1.addChild(j3)
    j1.addFollowOn(j4)

    with Toil(options) as toil:
        toil.start(j1)
```

In the example four jobs are created, first `j1` is run, then `j2` and `j3` are run in parallel as children of `j1`, finally `j4` is run as a follow-on of `j1`.

There are multiple short hand functions to achieve the same workflow, for example:

```
from toil.common import Toil
from toil.job import Job

def helloWorld(job, message, memory="2G", cores=2, disk="3G"):
    job.log("Hello world, I have a message: {}".format(message))

if __name__=="__main__":
    options = Job.Runner.getDefaultOptions("./toilWorkflowRun")
    options.logLevel = "INFO"
    options.clean = "always"

    j1 = Job.wrapJobFn(helloWorld, "first")
    j2 = j1.addChildJobFn(helloWorld, "second or third")
    j3 = j1.addChildJobFn(helloWorld, "second or third")
    j4 = j1.addFollowOnJobFn(helloWorld, "last")

    with Toil(options) as toil:
        toil.start(j1)
```

Equivalently defines the workflow, where the functions `toil.job.Job.addChildJobFn()` and `toil.job.Job.addFollowOnJobFn()` are used to create job functions as children or follow-ons of an earlier job.

Jobs graphs are not limited to trees, and can express arbitrary directed acyclic graphs. For a precise definition of legal graphs see `toil.job.Job.checkJobGraphForDeadlocks()`. The previous example could be specified as a DAG as follows:

```
from toil.common import Toil
from toil.job import Job

def helloWorld(job, message, memory="2G", cores=2, disk="3G"):
    job.log("Hello world, I have a message: {}".format(message))

if __name__=="__main__":
    options = Job.Runner.getDefaultOptions("./toilWorkflowRun")
    options.logLevel = "INFO"
    options.clean = "always"

    j1 = Job.wrapJobFn(helloWorld, "first")
    j2 = j1.addChildJobFn(helloWorld, "second or third")
    j3 = j1.addChildJobFn(helloWorld, "second or third")
    j4 = j2.addChildJobFn(helloWorld, "last")
    j3.addChild(j4)

    with Toil(options) as toil:
        toil.start(j1)
```

Note the use of an extra child edge to make `j4` a child of both `j2` and `j3`.

11.8 Dynamic Job Creation

The previous examples show a workflow being defined outside of a job. However, Toil also allows jobs to be created dynamically within jobs. For example:

```
from toil.common import Toil
from toil.job import Job

def binaryStringFn(job, depth, message=""):
    if depth > 0:
        job.addChildJobFn(binaryStringFn, depth-1, message + "0")
        job.addChildJobFn(binaryStringFn, depth-1, message + "1")
    else:
        job.log("Binary string: {}".format(message))

if __name__=="__main__":
    options = Job.Runner.getDefaultOptions("./toilWorkflowRun")
    options.logLevel = "INFO"
    options.clean = "always"

    with Toil(options) as toil:
        toil.start(Job.wrapJobFn(binaryStringFn, depth=5))
```

The job function `binaryStringFn` logs all possible binary strings of length `n` (here `n=5`), creating a total of $2^{(n+2)} - 1$ jobs dynamically and recursively. Static and dynamic creation of jobs can be mixed in a Toil workflow, with jobs defined within a job or job function being created at run time.

11.9 Promises

The previous example of dynamic job creation shows variables from a parent job being passed to a child job. Such forward variable passing is naturally specified by recursive invocation of successor jobs within parent jobs. This can also be achieved statically by passing around references to the return variables of jobs. In Toil this is achieved with promises, as illustrated in the following example:

```
from toil.common import Toil
from toil.job import Job

def fn(job, i):
    job.log("i is: %s" % i, level=100)
    return i+1

if __name__ == "__main__":
    options = Job.Runner.getDefaultOptions("./toilWorkflowRun")
    options.logLevel = "INFO"
    options.clean = "always"

    j1 = Job.wrapJobFn(fn, 1)
    j2 = j1.addChildJobFn(fn, j1.rv())
    j3 = j1.addFollowOnJobFn(fn, j2.rv())

    with Toil(options) as toil:
        toil.start(j1)
```

Running this workflow results in three log messages from the jobs: `i is 1` from `j1`, `i is 2` from `j2` and `i is 3` from `j3`.

The return value from the first job is *promised* to the second job by the call to `toil.job.Job.rv()` in the following line:

```
j2 = j1.addChildFn(fn, j1.rv())
```

The value of `j1.rv()` is a *promise*, rather than the actual return value of the function, because `j1` for the given input has at that point not been evaluated. A promise (`toil.job.Promise`) is essentially a pointer to for the return value that is replaced by the actual return value once it has been evaluated. Therefore, when `j2` is run the promise becomes 2.

Promises also support indexing of return values:

```
def parent(job):
    indexable = Job.wrapJobFn(fn)
    job.addChild(indexable)
    job.addFollowOnFn(raiseWrap, indexable.rv(2))

def raiseWrap(arg):
    raise RuntimeError(arg) # raises "2"

def fn(job):
    return (0, 1, 2, 3)
```

Promises can be quite useful. For example, we can combine dynamic job creation with promises to achieve a job creation process that mimics the functional patterns possible in many programming languages:


```

from toil.common import Toil
from toil.job import Job

def binaryStrings(job, depth, message=""):
    if depth > 0:
        s = [ job.addChildJobFn(binaryStrings, depth-1, message + "0").rv(),
              job.addChildJobFn(binaryStrings, depth-1, message + "1").rv() ]
        return job.addFollowOnFn(merge, s).rv()
    return [message]

def merge(strings):
    return strings[0] + strings[1]

if __name__=="__main__":
    options = Job.Runner.getDefaultOptions("./toilWorkflowRun")
    options.loglevel = "OFF"
    options.clean = "always"

    with Toil(options) as toil:
        print(toil.start(Job.wrapJobFn(binaryStrings, depth=5)))

```

The return value `l` of the workflow is a list of all binary strings of length 10, computed recursively. Although a toy example, it demonstrates how closely Toil workflows can mimic typical programming patterns.

11.10 Promised Requirements

Promised requirements are a special case of *Promises* that allow a job's return value to be used as another job's resource requirements.

This is useful when, for example, a job's storage requirement is determined by a file staged to the job store by an earlier job:

```

import os

from toil.common import Toil
from toil.job import Job, PromisedRequirement

def parentJob(job):
    downloadJob = Job.wrapJobFn(stageFn, "File://" + os.path.realpath(__file__),
    ↪cores=0.1, memory='32M', disk='1M')
    job.addChild(downloadJob)

    analysis = Job.wrapJobFn(analysisJob, fileStoreID=downloadJob.rv(0),
                             disk=PromisedRequirement(downloadJob.rv(1)))
    job.addFollowOn(analysis)

def stageFn(job, url, cores=1):
    importedFile = job.fileStore.importFile(url)
    return importedFile, importedFile.size

def analysisJob(job, fileStoreID, cores=2):
    # now do some analysis on the file
    pass

```

(continues on next page)

(continued from previous page)

```
if __name__ == "__main__":
    options = Job.Runner.getDefaultOptions("./toilWorkflowRun")
    options.logLevel = "INFO"
    options.clean = "always"

    with Toil(options) as toil:
        toil.start(Job.wrapJobFn(parentJob))
```

Note that this also makes use of the `size` attribute of the *FileID* object. This promised requirements mechanism can also be used in combination with an aggregator for multiple jobs' output values:

```
def parentJob(job):
    aggregator = []
    for fileNum in range(0,10):
        downloadJob = Job.wrapJobFn(stageFn, "File://" + os.path.realpath(__file__),
    ↪ cores=0.1, memory='32M', disk='1M')
        job.addChild(downloadJob)
        aggregator.append(downloadJob)

    analysis = Job.wrapJobFn(analysisJob, fileStoreID=downloadJob.rv(0),
        disk=PromisedRequirement(lambda xs: sum(xs), [j.rv(1)
    ↪ for j in aggregator]))
    job.addFollowOn(analysis)
```

Limitations

Just like regular promises, the return value must be determined prior to scheduling any job that depends on the return value. In our example above, notice how the dependent jobs were follow ons to the parent while promising jobs are children of the parent. This ordering ensures that all promises are properly fulfilled.

11.11 FileID

The `toil.fileStore.FileID` class is a small wrapper around Python's builtin string class. It is used to represent a file's ID in the file store, and has a `size` attribute that is the file's size in bytes. This object is returned by `importFile` and `writeGlobalFile`.

11.12 Managing files within a workflow

It is frequently the case that a workflow will want to create files, both persistent and temporary, during its run. The `toil.fileStores.abstractFileStore.AbstractFileStore` class is used by jobs to manage these files in a manner that guarantees cleanup and resumption on failure.

The `toil.job.Job.run()` method has a file store instance as an argument. The following example shows how this can be used to create temporary files that persist for the length of the job, be placed in a specified local disk of the node and that will be cleaned up, regardless of failure, when the job finishes:

```
from toil.common import Toil
from toil.job import Job
```

(continues on next page)

(continued from previous page)

```

class LocalFileStoreJob(Job):
    def run(self, fileStore):
        # self.TempDir will always contain the name of a directory within the
        ↪ allocated disk space reserved for the job
        scratchDir = self.tempDir

        # Similarly create a temporary file.
        scratchFile = fileStore.getLocalTempFile()

if __name__ == "__main__":
    options = Job.Runner.getDefaultOptions("./toilWorkflowRun")
    options.logLevel = "INFO"
    options.clean = "always"

    # Create an instance of FooJob which will have at least 2 gigabytes of storage
    ↪ space.
    j = LocalFileStoreJob(disk="2G")

    #Run the workflow
    with Toil(options) as toil:
        toil.start(j)

```

Job functions can also access the file store for the job. The equivalent of the `LocalFileStoreJob` class is

```

def localFileStoreJobFn(job):
    scratchDir = job.tempDir
    scratchFile = job.fileStore.getLocalTempFile()

```

Note that the `fileStore` attribute is accessed as an attribute of the `job` argument.

In addition to temporary files that exist for the duration of a job, the file store allows the creation of files in a *global* store, which persists during the workflow and are globally accessible (hence the name) between jobs. For example:

```

import os

from toil.common import Toil
from toil.job import Job

def globalFileStoreJobFn(job):
    job.log("The following example exercises all the methods provided"
           " by the toil.fileStores.abstractFileStore.AbstractFileStore class")

    # Create a local temporary file.
    scratchFile = job.fileStore.getLocalTempFile()

    # Write something in the scratch file.
    with open(scratchFile, 'w') as fh:
        fh.write("What a tangled web we weave")

    # Write a copy of the file into the file-store; fileID is the key that can be
    ↪ used to retrieve the file.
    # This write is asynchronous by default
    fileID = job.fileStore.writeGlobalFile(scratchFile)

    # Write another file using a stream; fileID2 is the

```

(continues on next page)

(continued from previous page)

```

# key for this second file.
with job.fileStore.writeGlobalFileStream(cleanup=True) as fH, fileID2:
    fH.write(b"Out brief candle")

# Now read the first file; scratchFile2 is a local copy of the file that is read-
→only by default.
scratchFile2 = job.fileStore.readGlobalFile(fileID)

# Read the second file to a desired location: scratchFile3.
scratchFile3 = os.path.join(job.tempDir, "foo.txt")
job.fileStore.readGlobalFile(fileID2, userPath=scratchFile3)

# Read the second file again using a stream.
with job.fileStore.readGlobalFileStream(fileID2) as fH:
    print(fH.read()) #This prints "Out brief candle"

# Delete the first file from the global file-store.
job.fileStore.deleteGlobalFile(fileID)

# It is unnecessary to delete the file keyed by fileID2 because we used the_
→cleanup flag,
# which removes the file after this job and all its successors have run (if the_
→file still exists)

if __name__ == "__main__":
    options = Job.Runner.getDefaultOptions("./toilWorkflowRun")
    options.logLevel = "INFO"
    options.clean = "always"

    with Toil(options) as toil:
        toil.start(Job.wrapJobFn(globalFileStoreJobFn))

```

The example demonstrates the global read, write and delete functionality of the file-store, using both local copies of the files and streams to read and write the files. It covers all the methods provided by the file store interface.

What is obvious is that the file-store provides no functionality to update an existing “global” file, meaning that files are, barring deletion, immutable. Also worth noting is that there is no file system hierarchy for files in the global file store. These limitations allow us to fairly easily support different object stores and to use caching to limit the amount of network file transfer between jobs.

11.12.1 Staging of Files into the Job Store

External files can be imported into or exported out of the job store prior to running a workflow when the `toil.common.Toil` context manager is used on the leader. The context manager provides methods `toil.common.Toil.importFile()`, and `toil.common.Toil.exportFile()` for this purpose. The destination and source locations of such files are described with URLs passed to the two methods. A list of the currently supported URLs can be found at `toil.jobStores.abstractJobStore.AbstractJobStore.importFile()`. To import an external file into the job store as a shared file, pass the optional `sharedFileName` parameter to that method.

If a workflow fails for any reason an imported file acts as any other file in the job store. If the workflow was configured such that it not be cleaned up on a failed run, the file will persist in the job store and needs not be staged again when the workflow is resumed.

Example:

```

import os

from toil.common import Toil
from toil.job import Job

class HelloWorld(Job):
    def __init__(self, id):
        Job.__init__(self, memory="2G", cores=2, disk="3G")
        self.inputFileID = id

    def run(self, fileStore):
        with self.fileStore.readGlobalFileStream(self.inputFileID) as fi:
            with self.fileStore.writeGlobalFileStream() as (fo, outputFileID):
                fo.write(fi.read() + 'World!')
        return outputFileID

if __name__ == "__main__":
    options = Job.Runner.getDefaultOptions("./toilWorkflowRun")
    options.logLevel = "INFO"
    options.clean = "always"

    with Toil(options) as toil:
        if not toil.options.restart:
            ioFileDirectory = os.path.join(os.path.dirname(os.path.abspath(__file__)),
↪ "stagingExampleFiles")
            inputFileID = toil.importFile("file://" + os.path.abspath(os.path.
↪ join(ioFileDirectory, "in.txt")))
            outputFileID = toil.start(HelloWorld(inputFileID))
        else:
            outputFileID = toil.restart()

        toil.exportFile(outputFileID, "file://" + os.path.abspath(os.path.
↪ join(ioFileDirectory, "out.txt")))

```

11.13 Using Docker Containers in Toil

Docker containers are commonly used with Toil. The combination of Toil and Docker allows for pipelines to be fully portable between any platform that has both Toil and Docker installed. Docker eliminates the need for the user to do any other tool installation or environment setup.

In order to use Docker containers with Toil, Docker must be installed on all workers of the cluster. Instructions for installing Docker can be found on the [Docker](#) website.

When using Toil-based autoscaling, Docker will be automatically set up on the cluster's worker nodes, so no additional installation steps are necessary. Further information on using Toil-based autoscaling can be found in the [Running a Workflow with Autoscaling](#) documentation.

In order to use docker containers in a Toil workflow, the container can be built locally or downloaded in real time from an online docker repository like [Quay](#). If the container is not in a repository, the container's layers must be accessible on each node of the cluster.

When invoking docker containers from within a Toil workflow, it is strongly recommended that you use `dockerCall()`, a toil job function provided in `toil.lib.docker`. `dockerCall` leverages docker's own

python API, and provides container cleanup on job failure. When docker containers are run without this feature, failed jobs can result in resource leaks. Docker’s API can be found at [docker-py](#).

In order to use `dockerCall`, your installation of Docker must be set up to run without `sudo`. Instructions for setting this up can be found [here](#).

An example of a basic `dockerCall` is below:

```
dockerCall(job=job,
           tool='quay.io/ucsc_cgl/bwa',
           workDir=job.tempDir,
           parameters=['index', '/data/reference.fa'])
```

Note the assumption that *reference.fa* file is located in */data*. This is Toil’s standard convention as a mount location to reduce boilerplate when calling `dockerCall`. Users can choose their own mount locations by supplying a *volumes* kwarg to `dockerCall`, such as: `volumes={working_dir: {'bind': '/data', 'mode': 'rw'}}`, where *working_dir* is an absolute path on the user’s filesystem.

`dockerCall` can also be added to workflows like any other job function:

```
import os

from toil.common import Toil
from toil.job import Job
from toil.lib.docker import apiDockerCall

align = Job.wrapJobFn(apiDockerCall,
                      image='ubuntu',
                      working_dir=os.getcwd(),
                      parameters=['ls', '-lha'])

if __name__ == "__main__":
    options = Job.Runner.getDefaultOptions("./toilWorkflowRun")
    options.logLevel = "INFO"
    options.clean = "always"

    with Toil(options) as toil:
        toil.start(align)
```

`cgl-docker-lib` contains `dockerCall`-compatible Dockerized tools that are commonly used in bioinformatics analysis.

The documentation provides guidelines for developing your own Docker containers that can be used with Toil and `dockerCall`. In order for a container to be compatible with `dockerCall`, it must have an `ENTRYPOINT` set to a wrapper script, as described in `cgl-docker-lib` containerization standards. This can be set by passing in the optional keyword argument, ‘`entrypoint`’. Example:

```
entrypoint=["/bin/bash","-c"]
```

`dockerCall` supports currently the 75 keyword arguments found in the python [Docker API](#), under the ‘`run`’ command.

11.14 Services

It is sometimes desirable to run *services*, such as a database or server, concurrently with a workflow. The `toil.job.Job.Service` class provides a simple mechanism for spawning such a service within a Toil workflow, allowing precise specification of the start and end time of the service, and providing start and end methods to use for initialization and cleanup. The following simple, conceptual example illustrates how services work:

```

from toil.common import Toil
from toil.job import Job

class DemoService(Job.Service):

    def start(self, fileStore):
        # Start up a database/service here
        # Return a value that enables another process to connect to the database
        return "loginCredentials"

    def check(self):
        # A function that if it returns False causes the service to quit
        # If it raises an exception the service is killed and an error is reported
        return True

    def stop(self, fileStore):
        # Cleanup the database here
        pass

j = Job()
s = DemoService()
loginCredentialsPromise = j.addService(s)

def dbFn(loginCredentials):
    # Use the login credentials returned from the service's start method to connect_
    ↪to the service
    pass

j.addChildFn(dbFn, loginCredentialsPromise)

if __name__=="__main__":
    options = Job.Runner.getDefaultOptions("./toilWorkflowRun")
    options.logLevel = "INFO"
    options.clean = "always"

    with Toil(options) as toil:
        toil.start(j)

```

In this example the `DemoService` starts a database in the `start` method, returning an object from the `start` method indicating how a client job would access the database. The service's `stop` method cleans up the database, while the service's `check` method is polled periodically to check the service is alive.

A `DemoService` instance is added as a service of the root job `j`, with resource requirements specified. The return value from `toil.job.Job.addService()` is a promise to the return value of the service's `start` method. When the promise is fulfilled it will represent how to connect to the database. The promise is passed to a child job of `j`, which uses it to make a database connection. The services of a job are started before any of its successors have been run and stopped after all the successors of the job have completed successfully.

Multiple services can be created per job, all run in parallel. Additionally, services can define sub-services using `toil.job.Job.Service.addChild()`. This allows complex networks of services to be created, e.g. Apache Spark clusters, within a workflow.

11.15 Checkpoints

Services complicate resuming a workflow after failure, because they can create complex dependencies between jobs. For example, consider a service that provides a database that multiple jobs update. If the database service fails and loses state, it is not clear that just restarting the service will allow the workflow to be resumed, because jobs that created that state may have already finished. To get around this problem Toil supports *checkpoint* jobs, specified as the boolean keyword argument `checkpoint` to a job or wrapped function, e.g.:

```
j = Job(checkpoint=True)
```

A checkpoint job is rerun if one or more of its successors fails its retry attempts, until it itself has exhausted its retry attempts. Upon restarting a checkpoint job all its existing successors are first deleted, and then the job is rerun to define new successors. By checkpointing a job that defines a service, upon failure of the service the database and the jobs that access the service can be redefined and rerun.

To make the implementation of checkpoint jobs simple, a job can only be a checkpoint if when first defined it has no successors, i.e. it can only define successors within its run method.

11.16 Encapsulation

Let A be a root job potentially with children and follow-ons. Without an encapsulated job the simplest way to specify a job B which runs after A and all its successors is to create a parent of A, call it A_p , and then make B a follow-on of A_p . e.g.:

```
from toil.common import Toil
from toil.job import Job

if __name__ == "__main__":
    # A is a job with children and follow-ons, for example:
    A = Job()
    A.addChild(Job())
    A.addFollowOn(Job())

    # B is a job which needs to run after A and its successors
    B = Job()

    # The way to do this without encapsulation is to make a parent of A, Ap, and make_
    ↪ B a follow-on of Ap.
    Ap = Job()
    Ap.addChild(A)
    Ap.addFollowOn(B)

    options = Job.Runner.getDefaultOptions("./toilWorkflowRun")
    options.logLevel = "INFO"
    options.clean = "always"

    with Toil(options) as toil:
        print(toil.start(Ap))
```

An *encapsulated job* E (A) of A saves making A_p , instead we can write:

```
from toil.common import Toil
from toil.job import Job
```

(continues on next page)

(continued from previous page)

```

if __name__ == "__main__":
    # A
    A = Job()
    A.addChild(Job())
    A.addFollowOn(Job())

    # Encapsulate A
    A = A.encapsulate()

    # B is a job which needs to run after A and its successors
    B = Job()

    # With encapsulation A and its successor subgraph appear to be a single job,
    ↪ hence:
    A.addChild(B)

    options = Job.Runner.getDefaultOptions("./toilWorkflowRun")
    options.logLevel = "INFO"
    options.clean = "always"

    with Toil(options) as toil:
        print(toil.start(A))

```

Note the call to `toil.job.Job.encapsulate()` creates the `toil.job.Job.EncapsulatedJob`.

11.17 Depending on Toil

If you are packing your workflow(s) as a pip-installable distribution on PyPI, you might be tempted to declare Toil as a dependency in your `setup.py`, via the `install_requires` keyword argument to `setup()`. Unfortunately, this does not work, for two reasons: For one, Toil uses Setuptools’ *extra* mechanism to manage its own optional dependencies. If you explicitly declared a dependency on Toil, you would have to hard-code a particular combination of extras (or no extras at all), robbing the user of the choice what Toil extras to install. Secondly, and more importantly, declaring a dependency on Toil would only lead to Toil being installed on the leader node of a cluster, but not the worker nodes. Auto-deployment does not work here because Toil cannot auto-deploy itself, the classic “Which came first, chicken or egg?” problem.

In other words, you shouldn’t explicitly depend on Toil. Document the dependency instead (as in “This workflow needs Toil version X.Y.Z to be installed”) and optionally add a version check to your `setup.py`. Refer to the `check_version()` function in the `toil-lib` project’s `setup.py` for an example. Alternatively, you can also just depend on `toil-lib` and you’ll get that check for free.

If your workflow depends on a dependency of Toil, consider not making that dependency explicit either. If you do, you risk a version conflict between your project and Toil. The `pip` utility may silently ignore that conflict, breaking either Toil or your workflow. It is safest to simply assume that Toil installs that dependency for you. The only downside is that you are locked into the exact version of that dependency that Toil declares. But such is life with Python, which, unlike Java, has no means of dependencies belonging to different software components within the same process, and whose favored software distribution utility is *incapable* of properly resolving overlapping dependencies and detecting conflicts.

11.18 Best Practices for Dockerizing Toil Workflows

Computational Genomics Lab’s [Dockstore](#) based production system provides workflow authors a way to run Dock-

erized versions of their pipeline in an automated, scalable fashion. To be compatible with this system of a workflow should meet the following requirements. In addition to the Docker container, a common workflow language [descriptor file](#) is needed. For inputs:

- Only command line arguments should be used for configuring the workflow. If the workflow relies on a configuration file, like [Toil-RNAseq](#) or [ProTECT](#), a wrapper script inside the Docker container can be used to parse the CLI and generate the necessary configuration file.
- All inputs to the pipeline should be explicitly enumerated rather than implicit. For example, don't rely on one FASTQ read's path to discover the location of its pair. This is necessary since all inputs are mapped to their own isolated directories when the Docker is called via Dockstore.
- All inputs must be documented in the CWL descriptor file. Examples of this file can be seen in both [Toil-RNAseq](#) and [ProTECT](#).

For outputs:

- All outputs should be written to a local path rather than S3.
- Take care to package outputs in a local and user-friendly way. For example, don't tar up all output if there are specific files that will care to see individually.
- All output file names should be deterministic and predictable. For example, don't prepend the name of an output file with PASS/FAIL depending on the outcome of the pipeline.
- All outputs must be documented in the CWL descriptor file. Examples of this file can be seen in both [Toil-RNAseq](#) and [ProTECT](#).

CHAPTER 12

Toil Class API

The Toil class configures and starts a Toil run.

CHAPTER 13

Job Store API

The job store interface is an abstraction layer that that hides the specific details of file storage, for example standard file systems, S3, etc. The `AbstractJobStore` API is implemented to support a give file store, e.g. S3. Implement this API to support a new file store.

Functions to wrap jobs and return values (promises).

14.1 FunctionWrappingJob

The subclass of Job for wrapping user functions.

14.2 JobFunctionWrappingJob

The subclass of FunctionWrappingJob for wrapping user job functions.

14.3 EncapsulatedJob

The subclass of Job for *encapsulating* a job, allowing a subgraph of jobs to be treated as a single job.

14.4 Promise

The class used to reference return values of jobs/services not yet run/started.

CHAPTER 15

Job Methods API

Jobs are the units of work in Toil which are composed into workflows.

CHAPTER 16

Job.Runner API

The Runner contains the methods needed to configure and start a Toil run.

The `AbstractFileStore` is an abstraction of a Toil run's shared storage.

class `toil.fileStores.FileID` (*fileStoreID, size, executable=False*)

A small wrapper around Python's builtin string class. It is used to represent a file's ID in the file store, and has a `size` attribute that is the file's size in bytes. This object is returned by `importFile` and `writeGlobalFile`.

Calls into the file store can use bare strings; size will be queried from the job store if unavailable in the ID.

__init__ (*fileStoreID, size, executable=False*)

Initialize self. See `help(type(self))` for accurate signature.

pack ()

Pack the `FileID` into a string so it can be passed through external code.

classmethod unpack (*packedFileStoreID*)

Unpack the result of `pack()` into a `FileID` object.

Batch System API

The batch system interface is used by Toil to abstract over different ways of running batches of jobs, for example Slurm, GridEngine, Mesos, Parasol and a single node. The `toil.batchSystems.abstractBatchSystem.AbstractBatchSystem` API is implemented to run jobs using a given job management system, e.g. Mesos.

18.1 Batch System Environmental Variables

Environmental variables allow passing of scheduler specific parameters.

For SLURM:

```
export TOIL_SLURM_ARGS="-t 1:00:00 -q fatq"
```

For TORQUE there are two environment variables - one for everything but the resource requirements, and another - for resources requirements (without the `-l` prefix):

```
export TOIL_TORQUE_ARGS="-q fatq"
export TOIL_TORQUE_REQS="walltime=1:00:00"
```

For GridEngine (SGE, UGE), there is an additional environmental variable to define the [parallel environment](#) for running multicore jobs:

```
export TOIL_GRIDENGINE_PE='smp'
export TOIL_GRIDENGINE_ARGS='-q batch.q'
```

For HTCondor, additional parameters can be included in the submit file passed to `condor_submit`:

```
export TOIL_HTCONDOR_PARAMS='requirements = TARGET.has_sse4_2 == true; accounting_
↪group = test'
```

The environment variable is parsed as a semicolon-separated string of `parameter = value` pairs.

18.2 Batch System API

CHAPTER 19

Job.Service API

The Service class allows databases and servers to be spawned within a Toil workflow.

CHAPTER 20

Exceptions API

Toil specific exceptions.

CHAPTER 21

Running Tests

Test make targets, invoked as `$ make <target>`, subject to which environment variables are set (see *Running Integration Tests*).

TARGET	DESCRIPTION
test	Invokes all tests.
integration_test	Invokes only the integration tests.
test_offline	Skips building the Docker appliance and only invokes tests that have no docker dependencies.
integration_test_local	Makes integration tests easier to debug locally by running the integration tests serially and doesn't redirect output. This makes it appears on the terminal as expected.

Before running tests for the first time, initialize your virtual environment following the steps in *Building from Source*.

Run all tests (including slow tests):

```
$ make test
```

Run only quick tests (as of Jul 25, 2018, this was ~ 20 minutes):

```
$ export TOIL_TEST_QUICK=True; make test
```

Run an individual test with:

```
$ make test tests=src/toil/test/sort/sortTest.py::SortTest::testSort
```

The default value for `tests` is `"src"` which includes all tests in the `src/` subdirectory of the project root. Tests that require a particular feature will be skipped implicitly. If you want to explicitly skip tests that depend on a currently installed *feature*, use

```
$ make test tests="-m 'not aws' src"
```

This will run only the tests that don't depend on the `aws` extra, even if that extra is currently installed. Note the distinction between the terms *feature* and *extra*. Every extra is a feature but there are features that are not extras, such

as the `gridengine` and `parasol` features. To skip tests involving both the `parasol` feature and the `aws` extra, use the following:

```
$ make test tests="-m 'not aws and not parasol' src"
```

21.1 Running Tests with pytest

Often it is simpler to use `pytest` directly, instead of calling the `make` wrapper. This usually works as expected, but some tests need some manual preparation. To run a specific test with `pytest`, use the following:

```
python -m pytest src/toil/test/sort/sortTest.py::SortTest::testSort
```

For more information, see the [pytest documentation](#).

21.2 Running Integration Tests

These tests are generally only run using in our CI workflow due to their resource requirements and cost. However, they can be made available for local testing:

- Running tests that make use of Docker (e.g. autoscaling tests and Docker tests) require an appliance image to be hosted. First, make sure you have gone through the set up found in [Using Docker with Quay](#). Then to build and host the appliance image run the `make target push_docker`.

```
$ make push_docker
```

- Running integration tests require activation via an environment variable as well as exporting information relevant to the desired tests. Enable the integration tests:

```
$ export TOIL_TEST_INTEGRATIVE=True
```

- Finally, set the environment variables for keyname and desired zone:

```
$ export TOIL_X_KEYNAME=[Your Keyname]
$ export TOIL_X_ZONE=[Desired Zone]
```

Where `X` is one of our currently supported cloud providers (GCE, AWS).

- See the above sections for guidance on running tests.

21.3 Test Environment Variables

TOIL_TEST_TEMP	An absolute path to a directory where Toil tests will write their temporary files. Defaults to the system's standard temporary directory .
TOIL_TEST_INTEGRATIVE	GRATIVE, this allows the integration tests to run. Only valid when running the tests from the source directory via <code>make test</code> or <code>make test_parallel</code> .
TOIL_AWS_KEYNAME	AWES keyname (see Preparing your AWS environment), which is required to run the AWS tests.
TOIL_GOOGLE_PROJECTID	GOOGLE Cloud account projectID (see Running in Google Compute Engine (GCE)), which is required to to run the Google Cloud tests.
TOIL_TEST_QUICK	If <code>True</code> , long running tests are skipped.

Partial install and failing tests

Some tests may fail with an `ImportError` if the required extras are not installed. Install Toil with all of the extras to prevent such errors.

21.4 Using Docker with Quay

[Docker](#) is needed for some of the tests. Follow the appropriate installation instructions for your system on their website to get started.

When running `make test` you might still get the following error:

```
$ make test
Please set TOIL_DOCKER_REGISTRY, e.g. to quay.io/USER.
```

To solve, make an account with [Quay](#) and specify it like so:

```
$ TOIL_DOCKER_REGISTRY=quay.io/USER make test
```

where `USER` is your Quay username.

For convenience you may want to add this variable to your `bashrc` by running

```
$ echo 'export TOIL_DOCKER_REGISTRY=quay.io/USER' >> $HOME/.bashrc
```

21.5 Running Mesos Tests

If you're running Toil's Mesos tests, be sure to create the `virtualenv` with `--system-site-packages` to include the Mesos Python bindings. Verify this by activating the `virtualenv` and running `pip list | grep mesos`. On macOS, this may come up empty. To fix it, run the following:

```
for i in /usr/local/lib/python2.7/site-packages/*mesos*; do ln -snf $i venv/lib/
↳python2.7/site-packages/; done
```

Developing with Docker

To develop on features reliant on the Toil Appliance (the docker image toil uses for AWS autoscaling), you should consider setting up a personal registry on [Quay](#) or [Docker Hub](#). Because the Toil Appliance images are tagged with the Git commit they are based on and because only commits on our master branch trigger an appliance build on Quay, as soon as a developer makes a commit or dirties the working copy they will no longer be able to rely on Toil to automatically detect the proper Toil Appliance image. Instead, developers wishing to test any appliance changes in autoscaling should build and push their own appliance image to a personal Docker registry. This is described in the next section.

22.1 Making Your Own Toil Docker Image

Note! Toil checks if the docker image specified by `TOIL_APPLIANCE_SELF` exists prior to launching by using the docker v2 schema. This should be valid for any major docker repository, but there is an option to override this if desired using the option: `-forceDockerAppliance`.

Here is a general workflow (similar instructions apply when using Docker Hub):

1. Make some changes to the provisioner of your local version of Toil
2. Go to the location where you installed the Toil source code and run

```
$ make docker
```

to automatically build a docker image that can now be uploaded to your personal [Quay](#) account. If you have not installed Toil source code yet see [Building from Source](#).

3. If it's not already you will need Docker installed and need to [log into Quay](#). Also you will want to make sure that your Quay account is public.
4. Set the environment variable `TOIL_DOCKER_REGISTRY` to your Quay account. If you find yourself doing this often you may want to add

```
export TOIL_DOCKER_REGISTRY=quay.io/<MY_QUAY_USERNAME>
```

to your `.bashrc` or equivalent.

5. Now you can run

```
$ make push_docker
```

which will upload the docker image to your Quay account. Take note of the image's tag for the next step.

6. Finally you will need to tell Toil from where to pull the Appliance image you've created (it uses the Toil release you have installed by default). To do this set the environment variable `TOIL_APPLIANCE_SELF` to the url of your image. For more info see [Environment Variables](#).
7. Now you can launch your cluster! For more information see [Running a Workflow with Autoscaling](#).

22.2 Running a Cluster Locally

The Toil Appliance container can also be useful as a test environment since it can simulate a Toil cluster locally. An important caveat for this is autoscaling, since autoscaling will only work on an EC2 instance and cannot (at this time) be run on a local machine.

To spin up a local cluster, start by using the following Docker run command to launch a Toil leader container:

```
docker run \
  --entrypoint=mesos-master \
  --net=host \
  -d \
  --name=leader \
  --volume=/home/jobStoreParentDir:/jobStoreParentDir \
  quay.io/ucsc_cgl/toil:3.6.0 \
  --registry=in_memory \
  --ip=127.0.0.1 \
  --port=5050 \
  --allocation_interval=500ms
```

A couple notes on this command: the `-d` flag tells Docker to run in daemon mode so the container will run in the background. To verify that the container is running you can run `docker ps` to see all containers. If you want to run your own container rather than the official UCSC container you can simply replace the `quay.io/ucsc_cgl/toil:3.6.0` parameter with your own container name.

Also note that we are not mounting the job store directory itself, but rather the location where the job store will be written. Due to complications with running Docker on MacOS, I recommend only mounting directories within your home directory. The next command will launch the Toil worker container with similar parameters:

```
docker run \
  --entrypoint=mesos-slave \
  --net=host \
  -d \
  --name=worker \
  --volume=/home/jobStoreParentDir:/jobStoreParentDir \
  quay.io/ucsc_cgl/toil:3.6.0 \
  --work_dir=/var/lib/mesos \
  --master=127.0.0.1:5050 \
  --ip=127.0.0.1 \
  ---attributes=preemptable:False \
  --resources=cpus:2
```

Note here that we are specifying 2 CPUs and a non-preemptable worker. We can easily change either or both of these in a logical way. To change the number of cores we can change the 2 to whatever number you like, and to change the worker to be preemptable we change `preemptable:False` to `preemptable:True`. Also note that the same

volume is mounted into the worker. This is needed since both the leader and worker write and read from the job store. Now that your cluster is running, you can run

```
docker exec -it leader bash
```

to get a shell in your leader 'node'. You can also replace the `leader` parameter with `worker` to get shell access in your worker.

Docker-in-Docker issues

If you want to run Docker inside this Docker cluster (Dockerized tools, perhaps), you should also mount in the Docker socket via `-v /var/run/docker.sock:/var/run/docker.sock`. This will give the Docker client inside the Toil Appliance access to the Docker engine on the host. Client/engine version mismatches have been known to cause issues, so we recommend using Docker version 1.12.3 on the host to be compatible with the Docker client installed in the Appliance. Finally, be careful where you write files inside the Toil Appliance - 'child' Docker containers launched in the Appliance will actually be siblings to the Appliance since the Docker engine is located on the host. This means that the 'child' container can only mount in files from the Appliance if the files are located in a directory that was originally mounted into the Appliance from the host - that way the files are accessible to the sibling container. Note: if Docker can't find the file/directory on the host it will silently fail and mount in an empty directory.

Maintainer's Guidelines

In general, as developers and maintainers of the code, we adhere to the following guidelines:

- We strive to never break the build on master. All development should be done on branches, in either the main Toil repository or in developers' forks.
- Pull requests should be used for any and all changes (except truly trivial ones).
- Pull requests should be in response to issues. If you find yourself making a pull request without an issue, you should create the issue first.

23.1 Naming Conventions

- **Commit messages** *should* be *great*. Most importantly, they *must*:
 - Have a short subject line. If in need of more space, drop down **two** lines and write a body to explain what is changing and why it has to change.
 - Write the subject line as a command: *Destroy all humans*, not *All humans destroyed*.
 - Reference the issue being fixed in a Github-parseable format, such as *(resolves #1234)* at the end of the subject line, or *This will fix #1234*. somewhere in the body. If no single commit on its own fixes the issue, the cross-reference must appear in the pull request title or body instead.
- **Branches** in the main Toil repository *must* start with `issues/`, followed by the issue number (or numbers, separated by a dash), followed by a short, lowercase, hyphenated description of the change. (There can be many open pull requests with their associated branches at any given point in time and this convention ensures that we can easily identify branches.)

Say there is an issue numbered #123 titled *Foo does not work*. The branch name would be `issues/123-fix-foo` and the title of the commit would be *Fix foo in case of bar (resolves #123)*.

23.2 Pull Requests

- All pull requests must be reviewed by a person other than the request’s author. Review the PR by following the [Reviewing Pull Requests](#) checklist.
- Modified pull requests must be re-reviewed before merging. **Note that Github does not enforce this!**
- Merge pull requests by following the [Merging Pull Requests](#) checklist.
- When merging a pull request, make sure to update the [Draft Changelog](#) on the Github wiki, which we will use to produce the changelog for the next release. The PR template tells you to do this, so don’t forget. New entries should go at the bottom.
- Pull requests will not be merged unless Travis and Gitlab CI tests pass. Gitlab tests are only run on code in the main Toil repository on some branch, so it is the responsibility of the approving reviewer to make sure that pull requests from outside repositories are copied to branches in the main repository. This can be accomplished with (from a Toil clone):

```
./contrib/admin/test-pr theirusername their-branch issues/123-fix-description-here
```

This must be repeated every time the PR submitter updates their PR, after checking to see that the update is not malicious.

If there is no issue corresponding to the PR, after which the branch can be named, the reviewer of the PR should first create the issue.

Developers who have push access to the main Toil repository are encouraged to make their pull requests from within the repository, to avoid this step.

- Prefer using “Squash and marge” when merging pull requests to master especially when the PR contains a “single unit” of work (i.e. if one were to rewrite the PR from scratch with all the fixes included, they would have one commit for the entire PR). This makes the commit history on master more readable and easier to debug in case of a breakage.

When squashing a PR from multiple authors, please add [Co-authored-by](#) to give credit to all contributing authors.

See [Issue #2816](#) for more details.

23.3 Publishing a Release

These are the steps to take to publish a Toil release:

- Determine the release version **X.Y.Z**. This should follow [semantic versioning](#); if user-workflow-breaking changes are made, **X** should be incremented, and **Y** and **Z** should be zero. If non-breaking changes are made but new functionality is added, **X** should remain the same as the last release, **Y** should be incremented, and **Z** should be zero. If only patches are released, **X** and **Y** should be the same as the last release and **Z** should be incremented.
- If it does not exist already, create a release branch in the Toil repo named **X.Y.x**, where **x** is a literal lower-case “x”. For patch releases, find the existing branch and make sure it is up to date with the patch commits that are to be released. They may be [cherry-picked over](#) from master.
- On the release branch, edit `version_template.py` in the root of the repository. Find the line that looks like this (slightly different for patch releases):

```
baseVersion = 'X.Y.0a1'
```

Make it look like this instead:

```
baseVersion = 'X.Y.Z'
```

Commit your change to the branch.

- Tag the current state of the release branch as `releases/X.Y.Z`.
- Make the Github release [here](#), referencing that tag. For a non-patch release, fill in the description with the changelog from [the wiki page](#), which you should clear. For a patch release, just describe the patch.
- For a non-patch release, set up the main branch so that development builds will declare themselves to be alpha versions of what the next release will probably be. Edit `version_template.py` in the root of the repository on the main branch to set `baseVersion` like this:

```
baseVersion = 'X.Y+1.0a1'
```

Make sure to replace `X` and `Y+1` with actual numbers.

23.4 Adding Retries to a Function

See `toil.lib.retry`.

`retry()` can be used to decorate any function based on the list of errors one wishes to retry on.

This list of errors can contain normal `Exception` objects, and/or `RetryCondition` objects wrapping `Exceptions` to include additional conditions.

For example, retrying on a one `Exception` (`HTTPError`):

```
from requests import get
from requests.exceptions import HTTPError

@retry(errors=[HTTPError])
def update_my_wallpaper():
    return get('https://www.deviantart.com/')
```

Or:

```
from requests import get
from requests.exceptions import HTTPError

@retry(errors=[HTTPError, ValueError])
def update_my_wallpaper():
    return get('https://www.deviantart.com/')
```

The examples above will retry for the default interval on any errors specified the “errors=” arg list.

To retry on specifically 500/502/503/504 errors, you could specify an `ErrorCondition` object instead, for example:

```
from requests import get
from requests.exceptions import HTTPError

@retry(errors=[
    ErrorCondition(
        error=HTTPError,
        error_codes=[500, 502, 503, 504]
    )])
def update_my_wallpaper():
    return requests.get('https://www.deviantart.com/')
```

To retry on specifically errors containing the phrase “NotFound”:

```
from requests import get
from requests.exceptions import HTTPError

@retry(errors=[
    ErrorCondition(
        error=HTTPError,
        error_message_must_include="NotFound"
    )])
def update_my_wallpaper():
    return requests.get('https://www.deviantart.com/')
```

To retry on all HTTPError errors EXCEPT an HTTPError containing the phrase “NotFound”:

```
from requests import get
from requests.exceptions import HTTPError

@retry(errors=[
    HTTPError,
    ErrorCondition(
        error=HTTPError,
        error_message_must_include="NotFound",
        retry_on_this_condition=False
    )])
def update_my_wallpaper():
    return requests.get('https://www.deviantart.com/')
```

To retry on boto3’s specific status errors, an example of the implementation is:

```
import boto3
from botocore.exceptions import ClientError

@retry(errors=[
    ErrorCondition(
        error=ClientError,
        boto_error_codes=["BucketNotFound"]
    )])
def boto_bucket(bucket_name):
    boto_session = boto3.session.Session()
    s3_resource = boto_session.resource('s3')
    return s3_resource.Bucket(bucket_name)
```

Any combination of these will also work, provided the codes are matched to the correct exceptions. A ValueError will not return a 404, for example.

The retry function as a decorator should make retrying functions easier and clearer. It also encourages smaller independent functions, as opposed to lumping many different things that may need to be retried on different conditions in the same function.

The ErrorCondition object tries to take some of the heavy lifting of writing specific retry conditions and boil it down to an API that covers all common use-cases without the user having to write any new bespoke functions.

Use-cases covered currently:

1. Retrying on a normal error, like a KeyError.
2. Retrying on HTTP error codes (use ErrorCondition).
3. Retrying on boto’s specific status errors, like “BucketNotFound” (use ErrorCondition).

4. Retrying when an error message contains a certain phrase (use `ErrorCondition`).
5. Explicitly NOT retrying on a condition (use `ErrorCondition`).

If new functionality is needed, it's currently best practice in Toil to add functionality to the `ErrorCondition` itself rather than making a new custom retry method.

Pull Request Checklists

This document contains checklists for dealing with PRs. More general PR information is available at [Pull Requests](#).

24.1 Reviewing Pull Requests

This checklist is to be kept in sync with the checklist in the pull request template.

When reviewing a PR, do the following:

- **Make sure it is coming from `issues/XXXX-fix-the-thing` in the Toil repo, or from an external repo.**
 - If it is coming from an external repo, make sure to pull it in for CI with:

```
contrib/admin/test-pr otheruser theirbranchname issues/XXXX-fix-the-thing
```
 - If there is no associated issue, [create one](#).
- **Read through the code changes. Make sure that it doesn't have:**
 - Addition of trailing whitespace.
 - New variable or member names in `camelCase` that want to be in `snake_case`.
 - New functions without [type hints](#).
 - New functions or classes without informative docstrings.
 - Changes to semantics not reflected in the relevant docstrings.
 - New or changed command line options for Toil workflows that are not reflected in `docs/running/cliOptions.rst`
 - New features without tests.
- Comment on the lines of code where problems exist with a review comment. You can shift-click the line numbers in the diff to select multiple lines.

- Finish the review with an overall description of your opinion.

24.2 Merging Pull Requests

This checklist is to be kept in sync with the checklist in the pull request template.

When merging a PR, do the following:

- Make sure the PR passes tests.
- Make sure the PR has been reviewed **since its last modification**. If not, review it.
- **Merge with the Github “Squash and merge” feature.**
 - If there are multiple authors’ commits, add **Co-authored-by** to give credit to all contributing authors.
- Copy its recommended changelog entry to the [Draft Changelog](#).
- Append the issue number in parentheses to the changelog entry.

Toil Architecture

The following diagram layouts out the software architecture of Toil.

These components are described below:

- **the leader:** The leader is responsible for deciding which jobs should be run. To do this it traverses the job graph. Currently this is a single threaded process, but we make aggressive steps to prevent it becoming a bottleneck (see *Read-only Leader* described below).
- **the job-store:** Handles all files shared between the components. Files in the job-store are the means by which the state of the workflow is maintained. Each job is backed by a file in the job store, and atomic updates to this state are used to ensure the workflow can always be resumed upon failure. The job-store can also store all user files, allowing them to be shared between jobs. The job-store is defined by the `AbstractJobStore` class. Multiple implementations of this class allow Toil to support different back-end file stores, e.g.: S3, network file systems, Google file store, etc.
- **workers:** The workers are temporary processes responsible for running jobs, one at a time per worker. Each worker process is invoked with a job argument that it is responsible for running. The worker monitors this job and reports back success or failure to the leader by editing the job's state in the file-store. If the job defines successor jobs the worker may choose to immediately run them (see *Job Chaining* below).
- **the batch-system:** Responsible for scheduling the jobs given to it by the leader, creating a worker command for each job. The batch-system is defined by the `AbstractBatchSystem` class. Toil uses multiple existing batch systems to schedule jobs, including Apache Mesos, GridEngine and a multi-process single node implementation that allows workflows to be run without any of these frameworks. Toil can therefore fairly easily be made to run a workflow using an existing cluster.
- **the node provisioner:** Creates worker nodes in which the batch system schedules workers. It is defined by the `AbstractProvisioner` class.
- **the statistics and logging monitor:** Monitors logging and statistics produced by the workers and reports them. Uses the job-store to gather this information.

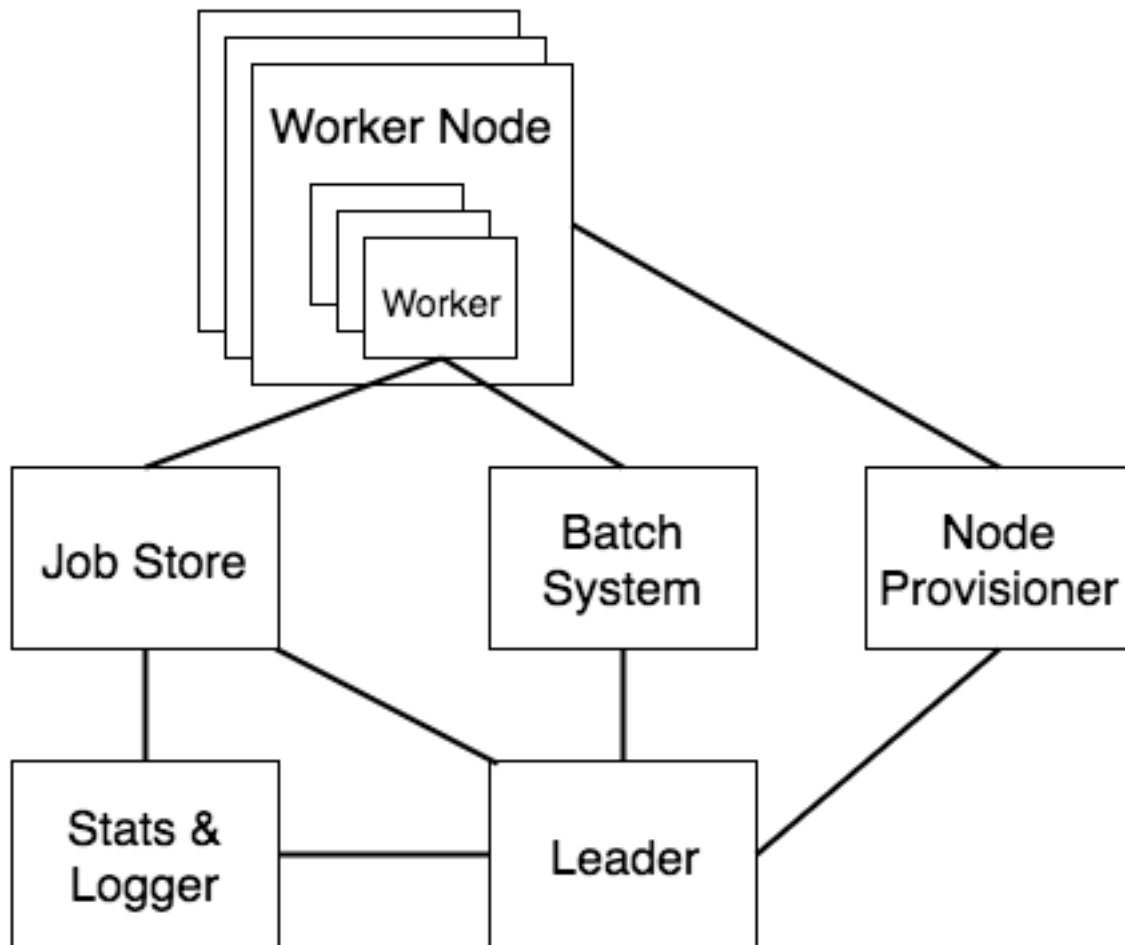


Fig. 1: Figure 1: The basic components of Toil's architecture.

25.1 Optimizations

Toil implements lots of optimizations designed for scalability. Here we detail some of the key optimizations.

25.1.1 Read-only leader

The leader process is currently implemented as a single thread. Most of the leader's tasks revolve around processing the state of jobs, each stored as a file within the job-store. To minimise the load on this thread, each worker does as much work as possible to manage the state of the job it is running. As a result, with a couple of minor exceptions, the leader process never needs to write or update the state of a job within the job-store. For example, when a job is complete and has no further successors the responsible worker deletes the job from the job-store, marking it complete. The leader then only has to check for the existence of the file when it receives a signal from the batch-system to know that the job is complete. This off-loading of state management is orthogonal to future parallelization of the leader.

25.1.2 Job chaining

The scheduling of successor jobs is partially managed by the worker, reducing the number of individual jobs the leader needs to process. Currently this is very simple: if there is a single next successor job to run and its resources fit within the resources of the current job and closely match the resources of the current job then the job is run immediately on the worker without returning to the leader. Further extensions of this strategy are possible, but for many workflows which define a series of serial successors (e.g. map sequencing reads, post-process mapped reads, etc.) this pattern is very effective at reducing leader workload.

25.1.3 Preemptable node support

Critical to running at large-scale is dealing with intermittent node failures. Toil is therefore designed to always be resumable providing the job-store does not become corrupt. This robustness allows Toil to run on preemptible nodes, which are only available when others are not willing to pay more to use them. Designing workflows that divide into many short individual jobs that can use preemptable nodes allows for workflows to be efficiently scheduled and executed.

25.1.4 Caching

Running bioinformatic pipelines often require the passing of large datasets between jobs. Toil caches the results from jobs such that child jobs running on the same node can directly use the same file objects, thereby eliminating the need for an intermediary transfer to the job store. Caching also reduces the burden on the local disks, because multiple jobs can share a single file. The resulting drop in I/O allows pipelines to run faster, and, by the sharing of files, allows users to run more jobs in parallel by reducing overall disk requirements.

To demonstrate the efficiency of caching, we ran an experimental internal pipeline on 3 samples from the TCGA Lung Squamous Carcinoma (LUSC) dataset. The pipeline takes the tumor and normal exome fastqs, and the tumor rna fastq and input, and predicts MHC presented neopeptides in the patient that are potential targets for T-cell based immunotherapies. The pipeline was run individually on the samples on c3.8xlarge machines on AWS (60GB RAM, 600GB SSD storage, 32 cores). The pipeline aligns the data to hg19-based references, predicts MHC haplotypes using PHLAT, calls mutations using 2 callers (MuTect and RADIA) and annotates them using SnpEff, then predicts MHC:peptide binding using the IEDB suite of tools before running an in-house rank boosting algorithm on the final calls.

To optimize time taken, The pipeline is written such that mutations are called on a per-chromosome basis from the whole-exome bams and are merged into a complete vcf. Running mutect in parallel on whole exome bams requires each mutect job to download the complete Tumor and Normal Bams to their working directories – An operation that

quickly fills the disk and limits the parallelizability of jobs. The script was run in Toil, with and without caching, and Figure 2 shows that the workflow finishes faster in the cached case while using less disk on average than the uncached run. We believe that benefits of caching arising from file transfers will be much higher on magnetic disk-based storage systems as compared to the SSD systems we tested this on.

25.2 Toil support for Common Workflow Language

The CWL document and input document are loaded using the `cwltool.load_tool` module. This performs normalization and URI expansion (for example, relative file references are turned into absolute file URIs), validates the document against the CWL schema, initializes Python objects corresponding to major document elements (command line tools, workflows, workflow steps), and performs static type checking that sources and sinks have compatible types.

Input files referenced by the CWL document and input document are imported into the Toil file store. CWL documents may use any URI scheme supported by Toil file store, including local files and object storage.

The `'location'` field of File references are updated to reflect the import token returned by the Toil file store.

For directory inputs, the directory listing is stored in Directory object. Each individual files is imported into Toil file store.

An initial workflow Job is created from the toplevel CWL document. Then, control passes to the Toil engine which schedules the initial workflow job to run.

When the toplevel workflow job runs, it traverses the CWL workflow and creates a toil job for each step. The dependency graph is expressed by making downstream jobs children of upstream jobs, and initializing the child jobs with an input object containing the promises of output from upstream jobs.

Because Toil jobs have a single output, but CWL permits steps to have multiple output parameters that may feed into multiple other steps, the input to a CWLJob is expressed with an “indirect dictionary”. This is a dictionary of input parameters, where each entry value is a tuple of a promise and a promise key. When the job runs, the indirect dictionary is turned into a concrete input object by resolving each promise into its actual value (which is always a dict), and then looking up the promise key to get the actual value for the the input parameter.

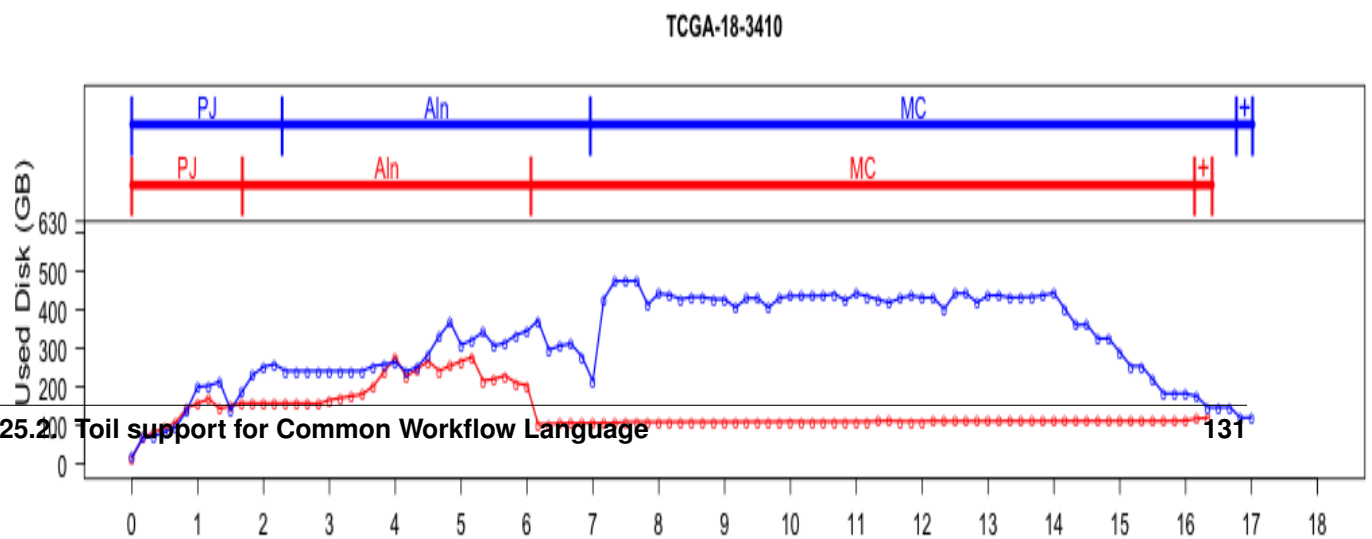
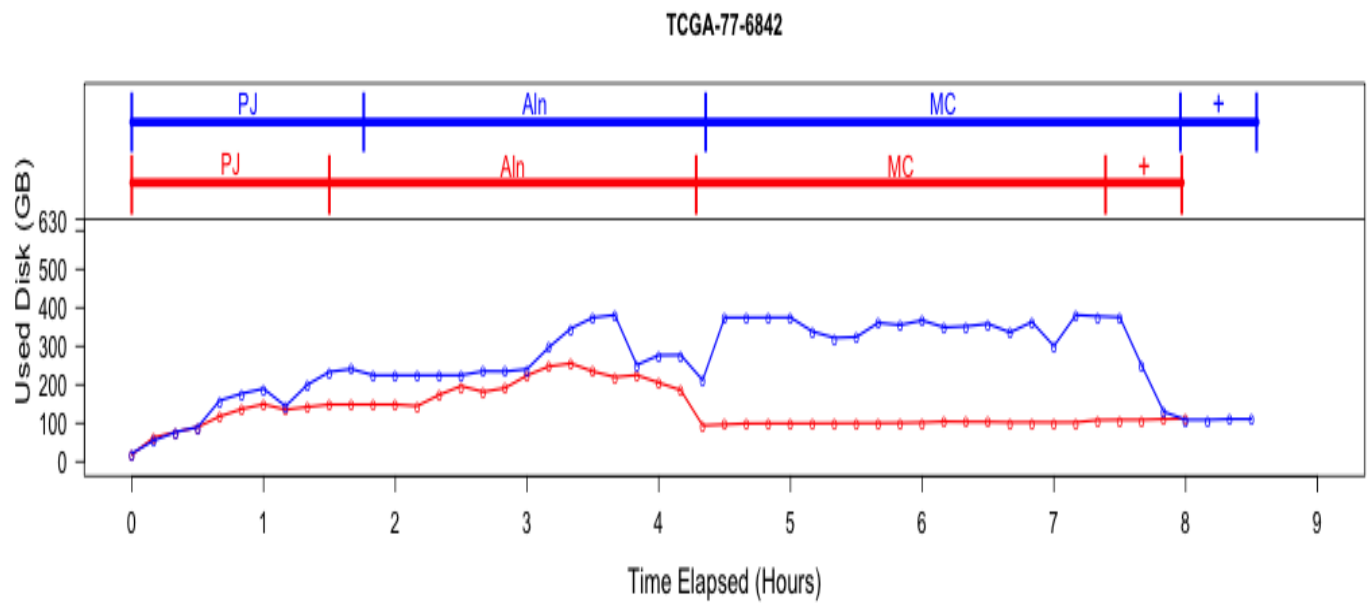
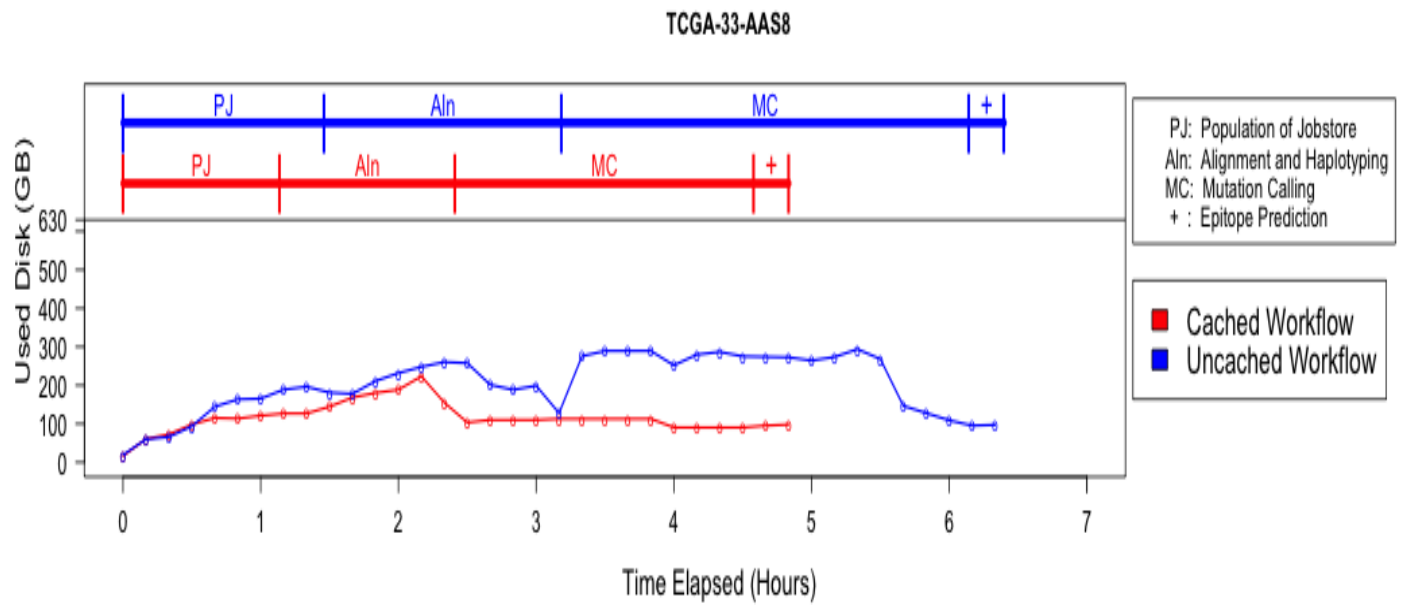
If a workflow step specifies a scatter, then a scatter job is created and connected into the workflow graph as described above. When the scatter step runs, it creates child jobs for each parameterizations of the scatter. A gather job is added as a follow-on to gather the outputs into arrays.

When running a command line tool, it first creates output and temporary directories under the Toil local temp dir. It runs the command line tool using the `single_job_executor` from `CWLTool`, providing a Toil-specific constructor for filesystem access, and overriding the default `PathMapper` to use `ToilPathMapper`.

The `ToilPathMapper` keeps track of a file’s symbolic identifier (the `Toil FileID`), its local path on the host (the value returned by `readGlobalFile`) and the the location of the file inside the Docker container.

After executing `single_job_executor` from `CWLTool`, it gets back the output object and status. If the underlying job failed, raise an exception. Files from the output object are added to the file store using `writeGlobalFile` and the `'location'` field of File references are updated to reflect the token returned by the Toil file store.

When the workflow completes, it returns an indirect dictionary linking to the outputs of the job steps that contribute to the final output. This is the value returned by `toil.start()` or `toil.restart()`. This is resolved to get the final output object. The files in this object are exported from the file store to `'outdir'` on the host file system, and the `'location'` field of File references are updated to reflect the final exported location of the output files.



Minimum AWS IAM permissions

Toil requires at least the following permissions in an IAM role to operate on a cluster. These are added by default when launching a cluster. However, ensure that they are present if creating a custom IAM role when *launching a cluster* with the `--awsEc2ProfileArn` parameter.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:*",
        "s3:*",
        "sdb:*",
        "iam:PassRole"
      ],
      "Resource": "*"
    }
  ]
}
```

Auto-Deployment

If you want to run your workflow in a distributed environment, on multiple worker machines, either in the cloud or on a bare-metal cluster, your script needs to be made available to those other machines. If your script imports other modules, those modules also need to be made available on the workers. Toil can automatically do that for you, with a little help on your part. We call this feature *auto-deployment* of a workflow.

Let's first examine various scenarios of auto-deploying a workflow, which, as we'll see shortly cannot be auto-deployed. Lastly, we'll deal with the issue of declaring *Toil as a dependency* of a workflow that is packaged as a *setuptools* distribution.

Toil can be easily deployed to a remote host. First, assuming you've followed our *Preparing your AWS environment* section to install Toil and use it to create a remote leader node on (in this example) AWS, you can now log into this into using *Ssh-Cluster Command* and once on the remote host, create and activate a *virtualenv* (noting to make sure to use the `--system-site-packages` option!):

```
$ virtualenv --system-site-packages venv
$ . venv/bin/activate
```

Note the `--system-site-packages` option, which ensures that globally-installed packages are accessible inside the *virtualenv*. Do not (re)install Toil after this! The `--system-site-packages` option has already transferred Toil and the dependencies from your local installation of Toil for you.

From here, you can install a project and its dependencies:

```
$ tree
.
├── util
│   ├── __init__.py
│   └── sort
│       ├── __init__.py
│       └── quick.py
└── workflow
    ├── __init__.py
    └── main.py
```

(continues on next page)

(continued from previous page)

```
3 directories, 5 files
$ pip install matplotlib
$ cp -R workflow util venv/lib/python2.7/site-packages
```

Ideally, your project would have a `setup.py` file (see [setuptools](#)) which streamlines the installation process:

```
$ tree
.
├── util
│   ├── __init__.py
│   └── sort
│       ├── __init__.py
│       └── quick.py
├── workflow
│   ├── __init__.py
│   └── main.py
└── setup.py

3 directories, 6 files
$ pip install .
```

Or, if your project has been published to PyPI:

```
$ pip install my-project
```

In each case, we have created a virtualenv with the `--system-site-packages` flag in the `venv` subdirectory then installed the `matplotlib` distribution from PyPI along with the two packages that our project consists of. (Again, both Python and Toil are assumed to be present on the leader and all worker nodes.)

We can now run our workflow:

```
$ python main.py --batchSystem=mesos ...
```

Important: If workflow's external dependencies contain native code (i.e. are not pure Python) then they must be manually installed on each worker.

Warning: Neither `python setup.py develop` nor `pip install -e .` can be used in this process as, instead of copying the source files, they create `.egg-link` files that Toil can't auto-deploy. Similarly, `python setup.py install` doesn't work either as it installs the project as a Python `.egg` which is also not currently supported by Toil (though it [could be](#) in the future).

Also note that using the `--single-version-externally-managed` flag with `setup.py` will prevent the installation of your package as an `.egg`. It will also disable the automatic installation of your project's dependencies.

27.1 Auto Deployment with Sibling Modules

This scenario applies if the user script imports modules that are its siblings:

```
$ cd my_project
$ ls
userScript.py utilities.py
$ ./userScript.py --batchSystem=mesos ...
```

Here `userScript.py` imports additional functionality from `utilities.py`. Toil detects that `userScript.py` has sibling modules and copies them to the workers, alongside the user script. Note that sibling modules will be auto-deployed regardless of whether they are actually imported by the user script—all `.py` files residing in the same directory as the user script will automatically be auto-deployed.

Sibling modules are a suitable method of organizing the source code of reasonably complicated workflows.

27.2 Auto-Deploying a Package Hierarchy

Recall that in Python, a **package** is a directory containing one or more `.py` files—one of which must be called `__init__.py`—and optionally other packages. For more involved workflows that contain a significant amount of code, this is the recommended way of organizing the source code. Because we use a package hierarchy, we can't really refer to the user script as such, we call it the user *module* instead. It is merely one of the modules in the package hierarchy. We need to inform Toil that we want to use a package hierarchy by invoking Python's `-m` option. That enables Toil to identify the entire set of modules belonging to the workflow and copy all of them to each worker. Note that while using the `-m` option is optional in the scenarios above, it is mandatory in this one.

The following shell session illustrates this:

```
$ cd my_project
$ tree
.
├── utils
│   ├── __init__.py
│   └── sort
│       ├── __init__.py
│       └── quick.py
└── workflow
    ├── __init__.py
    └── main.py

3 directories, 5 files
$ python -m workflow.main --batchSystem=mesos ...
```

Here the user module `main.py` does not reside in the current directory, but is part of a package called `util`, in a subdirectory of the current directory. Additional functionality is in a separate module called `util.sort.quick` which corresponds to `util/sort/quick.py`. Because we invoke the user module via `python -m workflow.main`, Toil can determine the root directory of the hierarchy—`my_project` in this case—and copy all Python modules underneath it to each worker. The `-m` option is documented [here](#)

When `-m` is passed, Python adds the current working directory to `sys.path`, the list of root directories to be considered when resolving a module name like `workflow.main`. Without that added convenience we'd have to run the workflow as `PYTHONPATH="$PWD" python -m workflow.main`. This also means that Toil can detect the root directory of the user module's package hierarchy even if it isn't the current working directory. In other words we could do this:

```
$ cd my_project
$ export PYTHONPATH="$PWD"
$ cd /some/other/dir
$ python -m workflow.main --batchSystem=mesos ...
```

Also note that the root directory itself must not be package, i.e. must not contain an `__init__.py`.

27.3 Relying on Shared Filesystems

Bare-metal clusters typically mount a shared file system like NFS on each node. If every node has that file system mounted at the same path, you can place your project on that shared filesystem and run your user script from there. Additionally, you can clone the Toil source tree into a directory on that shared file system and you won't even need to install Toil on every worker. Be sure to add both your project directory and the Toil clone to `PYTHONPATH`. Toil replicates `PYTHONPATH` from the leader to every worker.

Using a shared filesystem

Toil currently only supports a `tempdir` set to a local, non-shared directory.

27.3.1 Toil Appliance

The term Toil Appliance refers to the Mesos Docker image that Toil uses to simulate the machines in the virtual mesos cluster. It's easily deployed, only needs Docker, and allows for workflows to be run in single-machine mode and for clusters of VMs to be provisioned. To specify a different image, see the Toil [Environment Variables](#) section. For more information on the Toil Appliance, see the [Running in AWS](#) section.

CHAPTER 28

Environment Variables

There are several environment variables that affect the way Toil runs.

TOIL_CHECK_ENV	<code>True</code> determines whether Toil will try to refer back to a Python virtual environment in which it is installed when composing commands that may be run on other hosts. If set to <code>True</code> , if Toil is installed in the current virtual environment, it will use absolute paths to its own executables (and the virtual environment must thus be available on at the same path on all nodes). Otherwise, Toil internal commands such as <code>_toil_worker</code> will be resolved according to the <code>PATH</code> on the node where they are executed. This setting can be useful in a shared HPC environment, where users may have their own Toil installations in virtual environments.
TOIL_WORKDIR	Absolute path to a directory where Toil will write its temporary files. This directory must exist on each worker node and may be set to a different value on each worker. The <code>--workDir</code> command line option overrides this. When using the Toil docker container, such as on Kubernetes, this defaults to <code>/var/lib/toil</code> . When using Toil autoscaling with Mesos, this is somewhere inside the Mesos sandbox. In all other cases, the system's standard temporary directory is used.
TOIL_KUBERNETES_HOST_PATH	A path that will be mounted as <code>/tmp</code> in the workers, to allow for shared caching.
TOIL_KUBERNETES_JOB_NAME	Job identification of Kubernetes jobs. If not set, Toil will use the current user name.
KUBE_WATCH_ENABLED	A boolean that allows for users to utilize kubernetes watch stream feature instead of polling for running jobs. Default value is set to <code>False</code> .
TOIL_APPLIANCE_SELF	A self-referencing reference for the Toil Appliance you wish to use, in the form <code>REPO/IMAGE:TAG</code> . <code>quay.io/ucsc_cgl/toil:3.6.0</code> and <code>cket/toil:3.5.0</code> are both examples of valid options. Note that since Docker defaults to Dockerhub repos, only quay.io repos need to specify their registry.
TOIL_DOCKER_REGISTRY	Docker registry of the Toil Appliance image you wish to use. Docker will use Dockerhub by default, but the quay.io registry is also very popular and easily specifiable by setting this option to <code>quay.io</code> .
TOIL_DOCKER_NAME	The Toil Appliance image you wish to use. Generally this is simply <code>toil</code> but this option is provided to override this, since the image can be built with arbitrary names.
TOIL_AWS_SECRET_NAME	For the batch system, the name of a Kubernetes secret which contains a <code>credentials</code> file granting access to AWS resources. Will be mounted as <code>~/aws</code> inside Kubernetes-managed Toil containers. Enables the <code>AWSJobStore</code> to be used with the Kubernetes batch system, if the credentials allow access to S3 and SimpleDB.
TOIL_AWS_ZONE	The AWS zone to provision nodes in if using Toil's provisioner.
TOIL_AWS_AMI	The AMI to use in node provisioning. If in doubt, don't set this variable.
TOIL_AWS_NODE_DEBUG	A boolean to preserve nodes that have failed health checks. If set to <code>True</code> , nodes that fail EC2 health checks won't immediately be terminated so they can be examined and the cause of failure determined. If any EC2 nodes are left behind in this manner, the security group will also be left behind by necessity as it cannot be deleted until all associated nodes have been terminated.
TOIL_GOOGLE_PROJECT_ID	The Google Project ID to use when generating Google job store names for tests or CWL workflows.
TOIL_SLURM_ARGS	Arguments for slurm for the slurm batch system. Do not pass CPU or memory specifications here. Instead, define resource requirements for the job. There is no default value for this variable.
TOIL_GRIDENGINE_ARGS	Arguments for the gridengine batch system. Do not pass CPU or memory specifications here. Instead, define resource requirements for the job. There is no default value for this variable.
TOIL_GRIDENGINE_IP	IP argument arguments for qsub and for the gridengine batch system. There is no default value for this variable.
TOIL_TORQUE_ARGS	Arguments for qsub for the Torque batch system. Do not pass CPU or memory specifications here. Instead, define extra parameters for the job such as queue. Example: <code>-q medium</code> Use <code>TOIL_TORQUE_REQS</code> to pass extra values for the <code>-l</code> resource requirements parameter. There is no default value for this variable.
TOIL_TORQUE_REQS	Resource requirements for Torque batch system. Do not pass CPU or memory specifications here. Instead, define extra resource requirements as a string that goes after the <code>-l</code> argument to qsub. Example: <code>walltime=2:00:00,file=50gb</code> There is no default value for this variable.
TOIL_LSF_ARGS	Arguments for the LSF's bsub command. Instead, define extra parameters for the job such as queue. Example: <code>-q medium</code> . There is no default value for this variable.
TOIL_HTCONDOR_PARAMS	Parameters to include in the HTCondor submit file passed to <code>condor_submit</code> . Do not pass CPU or memory specifications here. Instead define extra parameters which may be required by HTCondor. This variable is parsed as a semicolon-separated string of <code>parameter = value</code> pairs. Example: <code>requirements = TARGET.has_sse4_2 == true; accounting_group = test</code> . There is no default value for this variable.

TOIL_CUSTOM_DOCKER_INIT_COMMAND A custom command to run in the Toil docker container prior to running the Toil services. Can be used for any custom initialization in the worker and/or primary nodes such as private docker authentication. Example for AWS ECR: `pip install awscli&& eval $(aws ecr get-login --no-include-email --region us-east-1)`.

- `genindex`
- `search`

Symbols

`__init__()` (*toil.fileStores.FileID method*), [103](#)

F

`FileID` (*class in toil.fileStores*), [103](#)

P

`pack()` (*toil.fileStores.FileID method*), [103](#)

U

`unpack()` (*toil.fileStores.FileID class method*), [103](#)