
Rally Documentation

Release 1.0.1.dev0

Daniel Mitterdorfer

Jul 11, 2018

- 1 Getting Help or Contributing to Rally 3**

- 2 Source Code 5**
 - 2.1 Quickstart 5
 - 2.2 Installation 7
 - 2.3 Configuration 10
 - 2.4 Run a Benchmark: Races 13
 - 2.5 Compare Results: Tournaments 16
 - 2.6 Tips and Tricks 19
 - 2.7 Define Custom Workloads: Tracks 24
 - 2.8 Developing Rally 43
 - 2.9 Command Line Reference 45
 - 2.10 Offline Usage 57
 - 2.11 Track Reference 58
 - 2.12 Configure Elasticsearch: Cars 82
 - 2.13 Using Elasticsearch Plugins 86
 - 2.14 Telemetry Devices 92
 - 2.15 Rally Daemon 95
 - 2.16 Pipelines 96
 - 2.17 Metrics 98
 - 2.18 Summary Report 102
 - 2.19 Migration Guide 106
 - 2.20 Frequently Asked Questions (FAQ) 115
 - 2.21 Glossary 117
 - 2.22 Community Resources 117

- 3 License 119**

Warning: This documentation is for the version of Rally currently under development. Were you looking for the [documentation of the latest stable version](#)?

You want to benchmark Elasticsearch? Then Rally is for you. It can help you with the following tasks:

- Setup and teardown of an Elasticsearch cluster for benchmarking
- Management of benchmark data and specifications even across Elasticsearch versions
- Running benchmarks and recording results
- Finding performance problems by attaching so-called telemetry devices
- Comparing performance results

We have also put considerable effort in Rally to ensure that benchmarking data are reproducible.

In general, Rally works with all versions of Elasticsearch starting from 1.x. *Benchmarking with plugins* and *benchmarking source builds* will only work from Elasticsearch 5.0 onwards.

CHAPTER 1

Getting Help or Contributing to Rally

- Use our [Discuss forum](#) to provide feedback or ask questions about Rally.
- See our [contribution guide](#) on guidelines for contributors.

Rally's source code is available on [Github](#). You can also check the [changelog](#) and the [roadmap](#) there.

2.1 Quickstart

Rally is developed for Unix and is actively tested on Linux and MacOS. Rally supports [benchmarking Elasticsearch clusters running on Windows](#) but Rally itself needs to be installed on machines running Unix.

2.1.1 Install

Install Python 3.4+ including `pip3`, `git` 1.9+ and an [appropriate JDK to run Elasticsearch](#) Be sure that `JAVA_HOME` points to that JDK. Then run the following command, optionally prefixed by `sudo` if necessary:

```
pip3 install esrally
```

If you have any trouble or need more detailed instructions, look in the [detailed installation guide](#).

2.1.2 Configure

Just invoke `esrally configure`.

For more detailed instructions and a detailed walkthrough see the [configuration guide](#).

2.1.3 Run your first race

Now we're ready to run our first *race*:

```
esrally --distribution-version=6.0.0
```

This will download Elasticsearch 6.0.0 and run Rally's default *track* - the *geonames track* - against it. After the race, a *summary report* is written to the command line::

```

-----
/ _ _ _ _ _ \ / _ _ _ _ \ / _ _ _ _ \
/ / _ _ \ / / _ _ \ / / _ _ \ / / _ _ \
/ / _ _ \ / / _ _ \ / / _ _ \ / / _ _ \
/ / _ _ \ / / _ _ \ / / _ _ \ / / _ _ \
-----

| Metric | Task | Value | Unit |
|-----:|-----:|-----:|-----:|
| Total indexing time | | 28.0997 | min |
| Total merge time | | 6.84378 | min |
| Total refresh time | | 3.06045 | min |
| Total flush time | | 0.106517 | min |
| Total merge throttle time | | 1.28193 | min |
| Median CPU usage | | 471.6 | % |
| Total Young Gen GC | | 16.237 | s |
| Total Old Gen GC | | 1.796 | s |
| Index size | | 2.60124 | GB |
| Totally written | | 11.8144 | GB |
| Heap used for segments | | 14.7326 | MB |
| Heap used for doc values | | 0.115917 | MB |
| Heap used for terms | | 13.3203 | MB |
| Heap used for norms | | 0.0734253 | MB |
| Heap used for points | | 0.5793 | MB |
| Heap used for stored fields | | 0.643608 | MB |
| Segment count | | 97 | |
| Min Throughput | index-append | 31925.2 | docs/s |
| Median Throughput | index-append | 39137.5 | docs/s |
| Max Throughput | index-append | 39633.6 | docs/s |
| 50.0th percentile latency | index-append | 872.513 | ms |
| 90.0th percentile latency | index-append | 1457.13 | ms |
| 99.0th percentile latency | index-append | 1874.89 | ms |
| 100th percentile latency | index-append | 2711.71 | ms |
| 50.0th percentile service time | index-append | 872.513 | ms |
| 90.0th percentile service time | index-append | 1457.13 | ms |
| 99.0th percentile service time | index-append | 1874.89 | ms |
| 100th percentile service time | index-append | 2711.71 | ms |
| ... | ... | ... | ... |
| ... | ... | ... | ... |
| Min Throughput | painless_dynamic | 2.53292 | ops/s |
| Median Throughput | painless_dynamic | 2.53813 | ops/s |
| Max Throughput | painless_dynamic | 2.54401 | ops/s |
| 50.0th percentile latency | painless_dynamic | 172208 | ms |
| 90.0th percentile latency | painless_dynamic | 310401 | ms |
| 99.0th percentile latency | painless_dynamic | 341341 | ms |
| 99.9th percentile latency | painless_dynamic | 344404 | ms |
| 100th percentile latency | painless_dynamic | 344754 | ms |
| 50.0th percentile service time | painless_dynamic | 393.02 | ms |
| 90.0th percentile service time | painless_dynamic | 407.579 | ms |
| 99.0th percentile service time | painless_dynamic | 430.806 | ms |
| 99.9th percentile service time | painless_dynamic | 457.352 | ms |
| 100th percentile service time | painless_dynamic | 459.474 | ms |
-----

```

(continues on next page)

(continued from previous page)

```
[INFO] SUCCESS (took 2634 seconds)
-----
```

2.1.4 Next steps

Now you can check *how to run benchmarks*, get a better understanding how to interpret the numbers in the *summary report* or start to *create your own tracks*. Be sure to check also some *tips and tricks* to help you understand how to solve specific problems in Rally.

Also run `esrally --help` to see what options are available and keep the *command line reference* handy for more detailed explanations of each option.

2.2 Installation

This is the detailed installation guide for Rally. If you are in a hurry you can check the *quickstart guide*.

2.2.1 Hardware Requirements

Use an SSD on the load generator machine. If you run bulk-indexing benchmarks, Rally will read one or more data files from disk. Usually, you will configure multiple clients and each client reads a portion of the data file. To the disk this appears as a random access pattern where spinning disks perform poorly. To avoid an accidental bottleneck on client-side you should therefore use an SSD on each load generator machine.

2.2.2 Prerequisites

Rally does not support Windows and is only actively tested on MacOS and Linux. Install the following packages first.

Python

- Python 3.4 or better available as `python3` on the path. Verify with: `python3 --version`.
- Python3 header files (included in the Python3 development package).
- `pip3` available on the path. Verify with `pip3 --version`.

Debian / Ubuntu

```
sudo apt-get install gcc python3-pip python3-dev
```

RHEL 6/ CentOS 6

Tested on CentOS release 6.9 (Final).

Note: You will need to enable `EPEL` before.

```
sudo yum install -y gcc python34.x86_64 python34-devel.x86_64 python34-setuptools.
↪noarch
# installs pip as it is not available as an OS package
sudo python3 /usr/lib/python3.4/site-packages/easy_install.py pip
```

RHEL 7 / CentOS 7

Note: You will need to enable [EPEL](#) before.

Tested on CentOS Linux release 7.4.1708 (Core).

```
sudo yum install -y gcc python34.x86_64 python34-devel.x86_64 python34-pip.noarch
```

Amazon Linux

```
sudo yum install -y gcc python35-pip.noarch python35-devel.x86_64
```

MacOS

We recommend that you use [Homebrew](#):

```
brew install python3
```

git

Git is not required if **all** of the following conditions are met:

- You are using Rally only as a load generator (`--pipeline=benchmark-only`) or you are referring to Elasticsearch configurations with `--team-path`.
- You create your own tracks and refer to them with `--track-path`.

In all other cases, Rally requires `git 1.9` or better. Verify with `git --version`.

Debian / Ubuntu

```
sudo apt-get install git
```

Red Hat / CentOS / Amazon Linux

```
sudo yum install git
```

Note: If you use RHEL, install a recent version of git via the [Red Hat Software Collections](#).

MacOS

git is already installed on MacOS.

JDK

A JDK is required on all machines where you want to launch Elasticsearch. If you use Rally just as a load generator to *benchmark remote clusters*, no JDK is required.

We recommend to use Oracle JDK but you are free to use OpenJDK as well. For details on how to install a JDK check your operating system's documentation pages.

To find the JDK, Rally expects the environment variable `JAVA_HOME` to be set on all targeted machines. To have more specific control, for example when you want to benchmark across a wide range of Elasticsearch releases, you can also set `JAVAx_HOME` where `x` is the major version of a JDK (e.g. `JAVA8_HOME` would point to a JDK 8 installation). Rally will then choose the highest supported JDK per version of Elasticsearch that is available.

Note: If you have Rally download, install and benchmark a local copy of Elasticsearch (i.e., the [default Rally behavior](#)) be sure to configure the Operating System (OS) of your Rally server with the [recommended kernel settings](#)

2.2.3 Installing Rally

Simply install Rally with pip: `pip3 install esrally`

Note: Depending on your system setup you may need to prepend this command with `sudo`.

If you get errors during installation, it is probably due to the installation of `psutil` which we use to gather system metrics like CPU utilization. Ensure that you have installed the Python development package as documented in the prerequisites section above.

2.2.4 Non-sudo Install

If you don't want to use `sudo` when installing Rally, installation is still possible but a little more involved:

1. Specify the `--user` option when installing Rally (step 2 above), so the command to be issued is: `python3 setup.py develop --user`.
2. Check the output of the install script or lookup the [Python documentation on the variable site.USER_BASE](#) to find out where the script is located. On Linux, this is typically `~/.local/bin`.

You can now either add `~/.local/bin` to your path or invoke Rally via `~/.local/bin/esrally` instead of just `esrally`.

2.2.5 VirtualEnv Install

You can also use Virtualenv to install Rally into an isolated Python environment without `sudo`.

1. Set up a new virtualenv environment in a directory with `virtualenv --python=python3 .`
2. Activate the environment with `source /path/to/virtualenv/dir/bin/activate`
3. Install Rally with `pip install esrally`

Whenever you want to use Rally, run the activation script (step 2 above) first. When you are done, simply execute `deactivate` in the shell to exit the virtual environment.

2.2.6 Offline Install

Warning: This documentation is for the version of Rally currently under development. We do not provide offline installation packages for development versions. Were you looking for the [documentation of the latest stable version](#)?

If you are in a corporate environment where your servers do not have any access to the Internet, you can use Rally's offline installation package. Follow these steps to install Rally:

1. Install all prerequisites as documented above.

2. Download the offline installation package for the [latest release](#) and copy it to the target machine(s).
3. Decompress the installation package with `tar -xzf esrally-dist-*.tar.gz`.
4. Run the install script with `sudo ./esrally-dist-*/install.sh`.

2.2.7 Next Steps

After you have installed Rally, you need to configure it. Just run `esrally configure` or follow the [configuration help page](#) for more guidance.

2.3 Configuration

Rally has to be configured once after installation. If you just run `esrally` after installing Rally, it will detect that the configuration file is missing and asks you a few questions.

If you want to reconfigure Rally at any later time, just run `esrally configure` again.

2.3.1 Simple Configuration

By default, Rally will run a simpler configuration routine and autodetect as much settings as possible or choose defaults for you. If you need more control you can run Rally with `esrally configure --advanced-config`.

Rally can build Elasticsearch either from sources or use an [official binary distribution](#). If you have Rally build Elasticsearch from sources, it can only be used to benchmark Elasticsearch 5.0 and above. The reason is that with Elasticsearch 5.0 the build tool switched from Maven to Gradle. As Rally utilizes the Gradle Wrapper, it is limited to Elasticsearch 5.0 and above.

Let's go through an example step by step: First run `esrally`:

```
dm@io:~ $ esrally

  _____
 /  _  \  _  /  _  /  _  /
/_  _  /  _  /  _  /  _  /
/_  _  /  _  /  _  /  _  /
/_  _  /  _  /  _  /  _  /
  _____

Running simple configuration. Run the advanced configuration with:

    esrally configure --advanced-config

* Setting up benchmark data directory in /Users/dm/.rally/benchmarks
* Setting up benchmark source directory in /Users/dm/.rally/benchmarks/src/
→elasticsearch

Configuration successfully written to /Users/dm/.rally/rally.ini. Happy benchmarking!

More info about Rally:

* Type esrally --help
* Read the documentation at https://esrally.readthedocs.io/en/latest/
* Ask a question on the forum at https://discuss.elastic.co/c/elasticsearch/rally
```

Congratulations! Time to [run your first benchmark](#).

2.3.2 Advanced Configuration

If you need more control over a few variables or want to store your metrics in a dedicated Elasticsearch metrics store, then you should run the advanced configuration routine. You can invoke it at any time with `esrally configure --advanced-config`.

Prerequisites

When using the advanced configuration, you can choose that Rally stores its metrics not in-memory but in a dedicated Elasticsearch instance. Therefore, you will also need the following software installed:

- **Elasticsearch:** a dedicated Elasticsearch instance which acts as the metrics store for Rally. If you don't want to set it up yourself you can also use [Elastic Cloud](#).
- **Optional:** Kibana (also included in [Elastic Cloud](#)).

Preparation

First [install Elasticsearch 5.0](#) or higher. A simple out-of-the-box installation with a single node will suffice. Rally uses this instance to store metrics data. It will setup the necessary indices by itself. The configuration procedure of Rally will you ask for host and port of this cluster.

Note: Rally will choose the port range 39200-39300 (HTTP) and 39300-39400 (transport) for the benchmark cluster, so do not use this port range for the metrics store.

Optional but recommended is to install also [Kibana](#). However, note that Kibana will not be auto-configured by Rally.

Configuration Options

Rally will ask you a few more things in the advanced setup:

- **Benchmark data directory:** Rally stores all benchmark related data in this directory which can take up to several tens of GB. If you want to use a dedicated partition, you can specify a different data directory here.
- **Elasticsearch project directory:** This is the directory where the Elasticsearch sources are located. If you don't actively develop on Elasticsearch you can just leave the default but if you want to benchmark local changes you should point Rally to your project directory. Note that Rally will run builds with the Gradle Wrapper in this directory (it runs `./gradlew clean` and `./gradlew :distribution:tar:assemble`).
- **Metrics store type:** You can choose between `in-memory` which requires no additional setup or `elasticsearch` which requires that you start a dedicated Elasticsearch instance to store metrics but gives you much more flexibility to analyse results.
- **Metrics store settings** (only for metrics store type `elasticsearch`): Provide the connection details to the Elasticsearch metrics store. This should be an instance that you use just for Rally but it can be a rather small one. A single node cluster with default setting should do it. When using self-signed certificates on the Elasticsearch metrics store, certificate verification can be turned off by setting the `datastore.ssl.verification_mode` setting to `none`. Alternatively you can enter the path to the certificate authority's signing certificate in `datastore.ssl.certificate_authorities`. Both settings are optional.
- **Name for this benchmark environment** (only for metrics store type `elasticsearch`): You can use the same metrics store for multiple environments (e.g. local, continuous integration etc.) so you can separate metrics from different environments by choosing a different name.

- whether or not Rally should keep the Elasticsearch benchmark candidate installation including all data by default. This will use lots of disk space so you should wipe `~/ .rally/benchmarks/races` regularly.

2.3.3 Proxy Configuration

Rally downloads all necessary data automatically for you:

- Elasticsearch distributions from elastic.co if you specify `--distribution-version=SOME_VERSION_NUMBER`
- Elasticsearch source code from Github if you specify a revision number e.g. `--revision=952097b`
- Track meta-data from Github
- Track data from an S3 bucket

Hence, it needs to connect via http(s) to the outside world. If you are behind a corporate proxy you need to configure Rally and git. As many other Unix programs, Rally relies that the HTTP proxy URL is available in the environment variable `http_proxy` (note that this is in lower-case). Hence, you should add this line to your shell profile, e.g. `~/ .bash_profile`:

```
export http_proxy=http://proxy.acme.org:8888/
```

Afterwards, source the shell profile with `source ~/ .bash_profile` and verify that the proxy URL is correctly set with `echo $http_proxy`.

Finally, you can set up git (see also the [Git config documentation](#)):

```
git config --global http.proxy $http_proxy
```

Verify that the proxy setup for git works correctly by cloning any repository, e.g. the `rally-tracks` repository:

```
git clone https://github.com/elastic/rally-tracks.git
```

If the configuration is correct, git will clone this repository. You can delete the folder `rally-tracks` after this verification step.

To verify that Rally will connect via the proxy server you can check the log file. If the proxy server is configured successfully, Rally will log the following line on startup:

```
Rally connects via proxy URL [http://proxy.acme.org:3128/] to the Internet (picked up ↵  
↵ from the environment variable [http_proxy]).
```

Note: Rally will use this proxy server only for downloading benchmark-related data. It will not use this proxy for the actual benchmark.

2.3.4 Logging

Logging in Rally is configured in `~/ .rally/logging.json`. For more information about the log file format please refer to the following documents:

- [Python logging cookbook](#) provides general tips and tricks.
- The Python reference documentation on the [logging configuration schema](#) explains the file format.
- The [logging handler documentation](#) describes how to customize where log output is written to.

By default, Rally will log all output to `~/ .rally/logs/rally.log` and rotate this file daily.

Example

With the following configuration Rally will log all output to standard error:

```

{
  "version": 1,
  "formatters": {
    "normal": {
      "format": "%(asctime)s,%(msecs)d %(actorAddress)s/PID:%(process)d %(name)s
↪ %(levelname)s %(message)s",
      "datefmt": "%Y-%m-%d %H:%M:%S",
      "():": "esrally.log.configure_utc_formatter"
    }
  },
  "filters": {
    "isActorLog": {
      "():": "thespian.director.ActorAddressLogFilter"
    }
  },
  "handlers": {
    "console_log_handler": {
      "class": "logging.StreamHandler",
      "formatter": "normal",
      "filters": ["isActorLog"]
    }
  },
  "root": {
    "handlers": ["console_log_handler"],
    "level": "INFO"
  },
  "loggers": {
    "elasticsearch": {
      "handlers": ["console_log_handler"],
      "level": "WARNING",
      "propagate": false
    }
  }
}

```

2.4 Run a Benchmark: Races

2.4.1 Definition

A “race” in Rally is the execution of a benchmarking experiment. You can choose different benchmarking scenarios (called *tracks*) for your benchmarks.

2.4.2 List Tracks

Start by finding out which tracks are available:

```
esrally list tracks
```

This will show the following list:

Name	Description	Documents	Documents
↳ Compressed Size	Uncompressed Size	↳ All Challenges	
geonames	POIs from Geonames	11396505	252.4 MB
↳	3.3 GB	append-no-conflicts	append-no-conflicts, appe...
geopoint	Point coordinates from PlanetOSM	60844404	481.9 MB
↳	2.3 GB	append-no-conflicts	append-no-conflicts, appe...
http_logs	HTTP server log data	247249096	1.2 GB
↳	31.1 GB	append-no-conflicts	append-no-conflicts, appe...
nested	StackOverflow Q&A stored as nested docs	11203029	663.1 MB
↳	3.4 GB	nested-search-challenge	nested-search-challenge, ...
noaa	Global daily weather measurements from NOAA	33659481	947.3 MB
↳	9.0 GB	append-no-conflicts	append-no-conflicts, appe...
nyc_taxis	Taxi rides in New York in 2015	165346692	4.5 GB
↳	74.3 GB	append-no-conflicts	append-no-conflicts, appe...
percolator	Percolator benchmark based on AOL queries	2000000	102.7 kB
↳	104.9 MB	append-no-conflicts	append-no-conflicts, appe...
pmc	Full text benchmark with academic papers from PMC	574199	5.5 GB
↳	21.7 GB	append-no-conflicts	append-no-conflicts, appe...

The first two columns show the name and a description of each track. A track also specifies one or more challenges which describe the workload to run.

2.4.3 Starting a Race

Note: Do not run Rally as root as Elasticsearch will refuse to start with root privileges.

To start a race you have to define the track and challenge to run. For example:

```
esrally --distribution-version=6.0.0 --track=geopoint --challenge=append-fast-with-
↳ conflicts
```

Rally will then start racing on this track. If you have never started Rally before, it should look similar to the following output:

```
dm@io:~ $ esrally --distribution-version=6.0.0 --track=geopoint --challenge=append-
↳ fast-with-conflicts

  / _ \ _ \ _ \ / / / _ \
 / / _ \ / _ \ \ / / / / /
 / _ \ / _ \ / / / / / /
 / _ \ | _ \ \ _ \ / / _ \ /
 / _ \ / _ \

[INFO] Racing on track [geopoint], challenge [append-fast-with-conflicts] and car [
↳ 'defaults'] with version [6.0.0].
[INFO] Downloading Elasticsearch 6.0.0 ... [OK]
[INFO] Rally will delete the benchmark candidate after the benchmark
[INFO] Downloading data from [http://benchmarks.elasticsearch.org.s3.amazonaws.com/
↳ corpora/geopoint/documents.json.bz2] (482 MB) to [/Users/dm/.rally/benchmarks/data/
↳ geopoint/documents.json.bz2] ... [OK]
[INFO] Decompressing track data from [/Users/dm/.rally/benchmarks/data/geopoint/
↳ documents.json.bz2] to [/Users/dm/.rally/benchmarks/data/geopoint/documents.json]
↳ (resulting size: 2.28 GB) ... [OK]
```

(continues on next page)

(continued from previous page)

```
[INFO] Preparing file offset table for [/Users/dm/.rally/benchmarks/data/geopoint/
↳documents.json] ... [OK]
Running index-update [ 0%↳
↳done]
```

The benchmark will take a while to run, so be patient.

When the race has finished, Rally will show a summary on the command line:

Metric	Task	Value	Unit
-----:-----:-----:-----:			
Total indexing time		124.712	min
Total merge time		21.8604	min
Total refresh time		4.49527	min
Total merge throttle time		0.120433	min
Median CPU usage		546.5	%
Total Young Gen GC		72.078	s
Total Old Gen GC		3.426	s
Index size		2.26661	GB
Totally written		30.083	GB
Heap used for segments		10.7148	MB
Heap used for doc values		0.0135536	MB
Heap used for terms		9.22965	MB
Heap used for points		0.78789	MB
Heap used for stored fields		0.683708	MB
Segment count		115	
Min Throughput	index-update	59210.4	docs/s
Median Throughput	index-update	65276.2	docs/s
Max Throughput	index-update	76516.6	docs/s
50.0th percentile latency	index-update	556.269	ms
90.0th percentile latency	index-update	852.779	ms
99.0th percentile latency	index-update	1854.31	ms
99.9th percentile latency	index-update	2972.96	ms
99.99th percentile latency	index-update	4106.91	ms
100th percentile latency	index-update	4542.84	ms
50.0th percentile service time	index-update	556.269	ms
90.0th percentile service time	index-update	852.779	ms
99.0th percentile service time	index-update	1854.31	ms
99.9th percentile service time	index-update	2972.96	ms
99.99th percentile service time	index-update	4106.91	ms
100th percentile service time	index-update	4542.84	ms
Min Throughput	force-merge	0.221067	ops/s
Median Throughput	force-merge	0.221067	ops/s
Max Throughput	force-merge	0.221067	ops/s
100th percentile latency	force-merge	4523.52	ms
100th percentile service time	force-merge	4523.52	ms

[INFO] SUCCESS (took 1624 seconds)			

Note: You can save this report also to a file by using `--report-file=/path/to/your/report.md` and save it as CSV with `--report-format=csv`.

What did Rally just do?

- It downloaded and started Elasticsearch 6.0.0
- It downloaded the relevant data for the geopoint track
- It ran the actual benchmark
- And finally it reported the results

If you are curious about the operations that Rally has run, inspect the [geopoint track specification](#) or start to *write your own tracks*. You can also configure Rally to *store all data samples in Elasticsearch* so you can analyze the results with Kibana. Finally, you may want to *change the Elasticsearch configuration*.

2.5 Compare Results: Tournaments

Suppose, we want to analyze the impact of a performance improvement.

First, we need a baseline measurement. For example:

```
esrally --track=pmc --revision=latest --user-tag="intention:baseline_github_1234"
```


Above we run the baseline measurement based on the latest source code revision of Elasticsearch. We can use the command line parameter `--user-tag` to provide a key-value pair to document the intent of a race.

Then we implement our changes and finally we want to run another benchmark to see the performance impact of the change. In that case, we do not want Rally to change our source tree and thus specify the pseudo-revision `current`:

```
esrally --track=pmc --revision=current --user-tag="intention:reduce_alloc_1234"
```

After we've run both races, we want to know about the performance impact. With Rally we can analyze differences of two given races easily. First of all, we need to find two races to compare by issuing `esrally list races`:

```
dm@io:~ $ esrally list races
```



```
Recent races:
```

Race Timestamp	Track	Track Parameters	Challenge	Car	User Tag
20160518T122341Z	pmc		append-no-conflicts	defaults	↵
↵intention:reduce_alloc_1234					
20160518T112057Z	pmc		append-no-conflicts	defaults	↵
↵intention:baseline_github_1234					
20160518T101957Z	pmc		append-no-conflicts	defaults	

We can see that the user tag helps us to recognize races. We want to compare the two most recent races and have to provide the two race timestamps in the next step:

```
dm@io:~ $ esrally compare --baseline=20160518T112057Z --contender=20160518T122341Z
```



(continues on next page)

(continued from previous page)

```

 / /_ / /_ \ / / / / / /
 / _ / _ / / / / / / /
 /_ / |_ | \_ /_ /_ /_ /
              /_ /

```

Comparing baseline
Race timestamp: 2016-05-18 11:20:57
Challenge: append-no-conflicts
Car: defaults

with contender
Race timestamp: 2016-05-18 12:23:41
Challenge: append-no-conflicts
Car: defaults

```

  _____
 / _____ \
 / /_ / /_ \ \ \ \ \ \ \ \ \
 / /_ / / / / \ \ \ \ \ \ \ \ \
 /_ / / / / / / / \ \ \ \ \ \ \ \ \
 /_ / / / / / / \ \ \ \ \ \ \ \ \

```

	Metric	Baseline	Contender	
↔	Diff			↔
↔				
↔383.00000	Min Indexing Throughput [docs/s]	19501	19118	-
↔304.45833	Median Indexing Throughput [docs/s]	20232	19927.5	-
↔323.00000	Max Indexing Throughput [docs/s]	21172	20849	-
↔+0.53603	Total indexing time [min]	55.7989	56.335	↔
↔+0.33495	Total merge time [min]	12.9766	13.3115	↔
↔+0.00030	Total refresh time [min]	5.20067	5.20097	↔
↔+0.00332	Total flush time [min]	0.0648667	0.0681833	↔
↔+0.08285	Total merge throttle time [min]	0.796417	0.879267	↔
↔+0.05372	Query latency term (50.0 percentile) [ms]	2.10049	2.15421	↔
↔+0.06630	Query latency term (90.0 percentile) [ms]	2.77537	2.84168	↔
↔+0.63287	Query latency term (100.0 percentile) [ms]	4.52081	5.15368	↔
↔1.66392	Query latency country_agg (50.0 percentile) [ms]	112.049	110.385	-
↔4.42138	Query latency country_agg (90.0 percentile) [ms]	128.426	124.005	-
↔22.19185	Query latency country_agg (100.0 percentile) [ms]	155.989	133.797	-
↔1.62519	Query latency scroll (50.0 percentile) [ms]	16.1226	14.4974	-
↔1.83043	Query latency scroll (90.0 percentile) [ms]	17.2383	15.4079	-

(continues on next page)

(continued from previous page)

	Query latency scroll (100.0 percentile) [ms]	18.8419	18.4241	-
↔0.41784				
	Query latency country_agg_cached (50.0 percentile) [ms]	1.70223	1.64502	-
↔0.05721				
	Query latency country_agg_cached (90.0 percentile) [ms]	2.34819	2.04318	-
↔0.30500				
	Query latency country_agg_cached (100.0 percentile) [ms]	3.42547	2.86814	-
↔0.55732				
	Query latency default (50.0 percentile) [ms]	5.89058	5.83409	-
↔0.05648				
	Query latency default (90.0 percentile) [ms]	6.71282	6.64662	-
↔0.06620				
	Query latency default (100.0 percentile) [ms]	7.65307	7.3701	-
↔0.28297				
	Query latency phrase (50.0 percentile) [ms]	1.82687	1.83193	⌋
↔+0.00506				
	Query latency phrase (90.0 percentile) [ms]	2.63714	2.46286	-
↔0.17428				
	Query latency phrase (100.0 percentile) [ms]	5.39892	4.22367	-
↔1.17525				
	Median CPU usage (index) [%]	668.025	679.15	⌋
↔+11.12499				
	Median CPU usage (stats) [%]	143.75	162.4	⌋
↔+18.64999				
	Median CPU usage (search) [%]	223.1	229.2	⌋
↔+6.10000				
	Total Young Gen GC [s]	39.447	40.456	⌋
↔+1.00900				
	Total Old Gen GC [s]	7.108	7.703	⌋
↔+0.59500				
	Index size [GB]	3.25475	3.25098	-
↔0.00377				
	Totally written [GB]	17.8434	18.3143	⌋
↔+0.47083				
	Heap used for segments [MB]	21.7504	21.5901	-
↔0.16037				
	Heap used for doc values [MB]	0.16436	0.13905	-
↔0.02531				
	Heap used for terms [MB]	20.0293	19.9159	-
↔0.11345				
	Heap used for norms [MB]	0.105469	0.0935669	-
↔0.01190				
	Heap used for points [MB]	0.773487	0.772155	-
↔0.00133				
	Heap used for points [MB]	0.677795	0.669426	-
↔0.00837				
	Segment count	136	121	-
↔15.00000				
	Indices Stats(90.0 percentile) [ms]	3.16053	3.21023	⌋
↔+0.04969				
	Indices Stats(99.0 percentile) [ms]	5.29526	3.94132	-
↔1.35393				
	Indices Stats(100.0 percentile) [ms]	5.64971	7.02374	⌋
↔+1.37403				
	Nodes Stats(90.0 percentile) [ms]	3.19611	3.15251	-
↔0.04360				
	Nodes Stats(99.0 percentile) [ms]	4.44111	4.87003	⌋
↔+0.42892				

(continues on next page)

(continued from previous page)

↔+0.44450	Nodes Stats(100.0 percentile) [ms]	5.22527	5.66977	↵
-----------	------------------------------------	---------	---------	---

2.6 Tips and Tricks


This section covers various tips and tricks in a recipe-style fashion.

2.6.1 Benchmarking an existing cluster

Warning: If you are just getting started with Rally and don't understand how it works, do NOT run it against any production or production-like cluster. Besides, benchmarks should be executed in a dedicated environment anyway where no additional traffic skews results.

Note: We assume in this recipe, that Rally is already properly *configured*.

Consider the following configuration: You have an existing benchmarking cluster, that consists of three Elasticsearch nodes running on 10.5.5.10, 10.5.5.11, 10.5.5.12. You've setup the cluster yourself and want to benchmark it with Rally. Rally is installed on 10.5.5.5.



First of all, we need to decide on a track. So, we run `esrally list tracks`:

Name	Description	Documents	↵
↔Compressed Size	Uncompressed Size	Default Challenge	All Challenges
↔	-----	-----	-----
geonames	POIs from Geonames	11396505	252.4 MB ↵
↔	3.3 GB	append-no-conflicts	append-no-conflicts, appe...
geopoint	Point coordinates from PlanetOSM	60844404	481.9 MB ↵
↔	2.3 GB	append-no-conflicts	append-no-conflicts, appe...
http_logs	HTTP server log data	247249096	1.2 GB ↵
↔	31.1 GB	append-no-conflicts	append-no-conflicts, appe...
nested	StackOverflow Q&A stored as nested docs	11203029	663.1 MB ↵
↔	3.4 GB	nested-search-challenge	nested-search-challenge, ...
noaa	Global daily weather measurements from NOAA	33659481	947.3 MB ↵
↔	9.0 GB	append-no-conflicts	append-no-conflicts, appe...

(continues on next page)

(continued from previous page)

nyc_taxi	Taxi rides in New York in 2015	165346692	4.5 GB	↳
↳	74.3 GB	append-no-conflicts	append-no-conflicts, appe...	
percolator	Percolator benchmark based on AOL queries	2000000	102.7 kB	↳
↳	104.9 MB	append-no-conflicts	append-no-conflicts, appe...	
pmc	Full text benchmark with academic papers from PMC	574199	5.5 GB	↳
↳	21.7 GB	append-no-conflicts	append-no-conflicts, appe...	

We're interested in a full text benchmark, so we'll choose to run `pmc`. If you have your own data that you want to use for benchmarks *create your own track* instead; the metrics you'll gather will be more representative and useful than some default track.

Next, we need to know which machines to target which is easy as we can see that from the diagram above.

Finally we need to check which *pipeline* to use. For this case, the `benchmark-only` pipeline is suitable as we don't want Rally to provision the cluster for us.

Now we can invoke Rally:

```
esrally --track=pmc --target-hosts=10.5.5.10:9200,10.5.5.11:9200,10.5.5.12:9200 --
↳pipeline=benchmark-only
```

If you have `X-Pack Security` enabled, then you'll also need to specify another parameter to use https and to pass credentials:

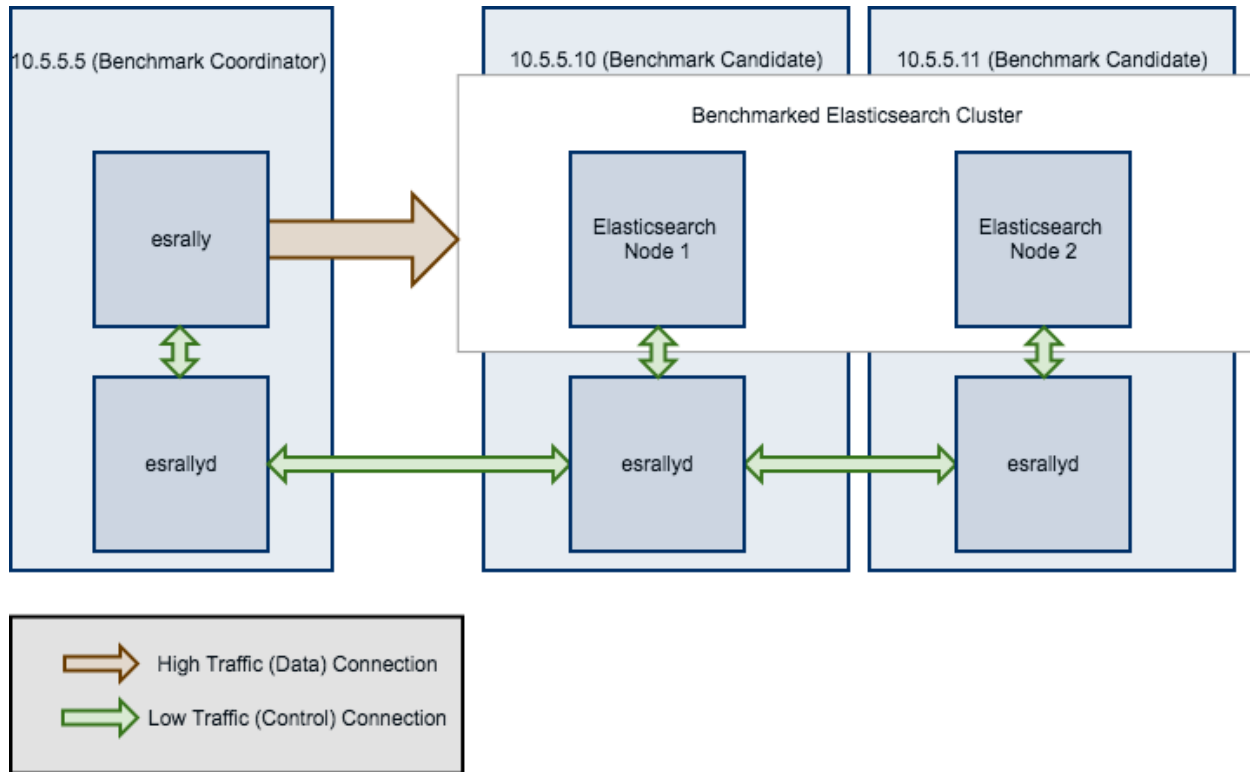
```
esrally --track=pmc --target-hosts=10.5.5.10:9243,10.5.5.11:9243,10.5.5.12:9243 --
↳pipeline=benchmark-only --client-options="use_ssl:true,verify_certs:true,basic_auth_
↳user:'elastic',basic_auth_password:'changeme'"
```

2.6.2 Benchmarking a remote cluster

Contrary to the previous recipe, you want Rally to provision all cluster nodes.

We will use the following configuration for the example:

- You will start Rally on `10.5.5.5`. We will call this machine the “benchmark coordinator”.
- Your Elasticsearch cluster will consist of two nodes which run on `10.5.5.10` and `10.5.5.11`. We will call these machines the “benchmark candidate”s.



Note: All `esrallyd` nodes form a cluster that communicates via the “benchmark coordinator”. For aesthetic reasons we do not show a direct connection between the “benchmark coordinator” and all nodes.

To run a benchmark for this scenario follow these steps:

1. *Install* and *configure* Rally on all machines. Be sure that the same version is installed on all of them and fully *configured*.
2. Start the *Rally daemon* on each machine. The Rally daemon allows Rally to communicate with all remote machines. On the benchmark coordinator run `esrallyd start --node-ip=10.5.5.5 --coordinator-ip=10.5.5.5` and on the benchmark candidate machines run `esrallyd start --node-ip=10.5.5.10 --coordinator-ip=10.5.5.5` and `esrallyd start --node-ip=10.5.5.11 --coordinator-ip=10.5.5.5` respectively. The `--node-ip` parameter tells Rally the IP of the machine on which it is running. As some machines have more than one network interface, Rally will not attempt to auto-detect the machine IP. The `--coordinator-ip` parameter tells Rally the IP of the benchmark coordinator node.
3. Start the benchmark by invoking Rally as usual on the benchmark coordinator, for example: `esrally --distribution-version=5.0.0 --target-hosts=10.5.5.10:9200,10.5.5.11:9200`. Rally will derive from the `--target-hosts` parameter that it should provision the nodes `10.5.5.10` and `10.5.5.11`.
4. After the benchmark has finished you can stop the Rally daemon again. On the benchmark coordinator and on the benchmark candidates run `esrallyd stop`.

Note: Logs are managed per machine, so all relevant log files and also telemetry output is stored on the benchmark candidates but not on the benchmark coordinator.

Now you might ask yourself what the differences to benchmarks of existing clusters are. In general you should aim to

give Rally as much control as possible as benchmark are easier reproducible and you get more metrics. The following table provides some guidance on when to choose which option:


Your requirement	Recommendation
You want to use Rally's telemetry devices	Use Rally daemon, as it can provision the remote node for you
You want to benchmark a source build of Elasticsearch	Use Rally daemon, as it can build Elasticsearch for you
You want to tweak the cluster configuration yourself	Use Rally daemon with a <i>custom configuration</i> or set up the cluster by yourself and use <code>--pipeline=benchmark-only</code>
You need to run a benchmark with plugins	Use Rally daemon if the <i>plugins</i> are supported or set up the cluster by yourself and use <code>--pipeline=benchmark-only</code>
You need to run a benchmark against multiple nodes	Use Rally daemon if all nodes can be configured identically. For more complex cases, set up the cluster by yourself and use <code>--pipeline=benchmark-only</code>

Rally daemon will be able to cover most of the cases described above in the future so there should be almost no case where you need to use the `benchmark-only` pipeline.

2.6.3 Distributing the load test driver

By default, Rally will generate load on the same machine where you start a benchmark. However, when you are benchmarking larger clusters, a single load test driver machine may not be able to generate sufficient load. In these cases, you should use multiple load driver machines. We will use the following configuration for the example:

- You will start Rally on `10.5.5.5`. We will call this machine the “benchmark coordinator”.
- You will start two load drivers on `10.5.5.6` and `10.5.5.7`. Note that one load driver will simulate multiple clients. Rally will simply assign clients to load driver machines in a round-robin fashion.
- Your Elasticsearch cluster will consist of three nodes which run on `10.5.5.11`, `10.5.5.12` and `10.5.5.13`. We will call these machines the “benchmark candidate”. For simplicity, we will assume an externally provisioned cluster but you can also use Rally to setup the cluster for you (see above).



1. *Install* and *configure* Rally on all machines. Be sure that the same version is installed on all of them and fully *configured*.
2. Start the *Rally daemon* on each machine. The Rally daemon allows Rally to communicate with all remote machines. On the benchmark coordinator run `esrallyd start --node-ip=10.5.5.5 --coordinator-ip=10.5.5.5` and on the load driver machines run `esrallyd start --node-ip=10.5.5.6 --coordinator-ip=10.5.5.5` and `esrallyd start --node-ip=10.5.5.7 --coordinator-ip=10.5.5.5` respectively. The `--node-ip` parameter tells Rally the IP of the machine on which it is running. As some machines have more than one network interface, Rally will not attempt to auto-detect the machine IP. The `--coordinator-ip` parameter tells Rally the IP of the benchmark coordinator node.
3. Start the benchmark by invoking Rally on the benchmark coordinator, for example: `esrally --pipeline=benchmark-only --load-driver-hosts=10.5.5.6,10.5.5.7 --target-hosts=10.5.5.11:9200,10.5.5.12:9200,10.5.5.13:9200`.

4. After the benchmark has finished you can stop the Rally daemon again. On the benchmark coordinator and on the load driver machines run `esrallyd stop`.

Note: Rally neither distributes code (i.e. *custom runners* or *parameter sources*) nor data automatically. You should place all tracks and their data on all machines in the same directory before starting the benchmark. Alternatively, you can store your track in a custom track repository.

Note: As indicated in the diagram, track data will be downloaded by each load driver machine separately. If you want to avoid that, you can run a benchmark once without distributing the load test driver (i.e. do not specify `--load-driver-hosts`) and then copy the contents of `~/.rally/benchmarks/data` to all load driver machines.

2.6.4 Changing the default track repository

Rally supports multiple track repositories. This allows you for example to have a separate company-internal repository for your own tracks that is separate from Rally's [default track repository](#). However, you always need to define `--track-repository=my-custom-repository` which can be cumbersome. If you want to avoid that and want Rally to use your own track repository by default you can just replace the default track repository definition in `~/.rally/rally.ini`. Consider this example:

```
...
[tracks]
default.url = git@github.com:elastic/rally-tracks.git
teamtrackrepo.url = git@example.org/myteam/my-tracks.git
```

If `teamtrackrepo` should be the default track repository, just define it as `default.url`. E.g.:

```
...
[tracks]
default.url = git@example.org/myteam/my-tracks.git
old-rally-default.url=git@github.com:elastic/rally-tracks.git
```

Also don't forget to rename the folder of your local working copy as Rally will search for a track repository with the name `default`:

```
cd ~/.rally/benchmarks/tracks/
mv default old-rally-default
mv teamtrackrepo default
```

From now on, Rally will treat your repository as default and you need to run Rally with `--track-repository=old-rally-default` if you want to use the out-of-the-box Rally tracks.

2.7 Define Custom Workloads: Tracks

2.7.1 Definition

A track describes one or more benchmarking scenarios. Its structure is described in detail in the [track reference](#).

2.7.2 Example track

We will create the track “tutorial” step by step. We store everything in the directory `~/rally-tracks/tutorial` but you can choose any other location.

First, get some data. [Geonames](#) provides geo data under a [creative commons license](#). Download `allCountries.zip` (around 300MB), extract it and inspect `allCountries.txt`.

The file is tab-delimited but to bulk-index data with Elasticsearch we need JSON. Convert the data with the following script:

```
import json

cols = (("geonameid", "int", True),
        ("name", "string", True),
        ("asciiname", "string", False),
        ("alternatenames", "string", False),
        ("latitude", "double", True),
        ("longitude", "double", True),
        ("feature_class", "string", False),
        ("feature_code", "string", False),
        ("country_code", "string", True),
        ("cc2", "string", False),
        ("admin1_code", "string", False),
        ("admin2_code", "string", False),
        ("admin3_code", "string", False),
        ("admin4_code", "string", False),
        ("population", "long", True),
        ("elevation", "int", False),
        ("dem", "string", False),
        ("timezone", "string", False))

def main():
    with open("allCountries.txt", "rt", encoding="UTF-8") as f:
        for line in f:
            tup = line.strip().split("\t")
            record = {}
            for i in range(len(cols)):
                name, type, include = cols[i]
                if tup[i] != "" and include:
                    if type in ("int", "long"):
                        record[name] = int(tup[i])
                    elif type == "double":
                        record[name] = float(tup[i])
                    elif type == "string":
                        record[name] = tup[i]
            print(json.dumps(record, ensure_ascii=False))

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

Store the script as `toJSON.py` in the tutorial directory (`~/rally-tracks/tutorial`). Invoke it with `python3 toJSON.py > documents.json`.

Then store the following mapping file as `index.json` in the tutorial directory:

```
{
  "settings": {
    "index.number_of_replicas": 0
  },
  "mappings": {
    "docs": {
      "dynamic": "strict",
      "properties": {
        "geonameid": {
          "type": "long"
        },
        "name": {
          "type": "text"
        },
        "latitude": {
          "type": "double"
        },
        "longitude": {
          "type": "double"
        },
        "country_code": {
          "type": "text"
        },
        "population": {
          "type": "long"
        }
      }
    }
  }
}
```

For details on the allowed syntax, see the Elasticsearch documentation on [mappings](#) and the [create index API](#).

Finally, store the track as `track.json` in the tutorial directory:

```
{
  "version": 2,
  "description": "Tutorial benchmark for Rally",
  "indices": [
    {
      "name": "geonames",
      "body": "index.json",
      "types": [ "docs" ]
    }
  ],
  "corpora": [
    {
      "name": "rally-tutorial",
      "documents": [
        {
          "source-file": "documents.json",
          "document-count": 11658903,
          "uncompressed-bytes": 1544799789
        }
      ]
    }
  ],
  "schedule": [
```

(continues on next page)

(continued from previous page)

```

{
  "operation": {
    "operation-type": "delete-index"
  }
},
{
  "operation": {
    "operation-type": "create-index"
  }
},
{
  "operation": {
    "operation-type": "cluster-health",
    "request-params": {
      "wait_for_status": "green"
    }
  }
},
{
  "operation": {
    "operation-type": "bulk",
    "bulk-size": 5000
  },
  "warmup-time-period": 120,
  "clients": 8
},
{
  "operation": {
    "operation-type": "force-merge"
  }
},
{
  "operation": {
    "name": "query-match-all",
    "operation-type": "search",
    "body": {
      "query": {
        "match_all": {}
      }
    }
  },
  "clients": 8,
  "warmup-iterations": 1000,
  "iterations": 1000,
  "target-throughput": 100
}
]
}

```

The numbers under the `documents` property are needed to verify integrity and provide progress reports. Determine the correct document count with `wc -l documents.json` and the size in bytes with `stat -f "%z" documents.json`.

Note: You can store any supporting scripts along with your track. However, you need to place them in a directory starting with “_”, e.g. “_support”. Rally loads track plugins (see below) from any directory but will ignore directories

starting with “_”.

Note: We have defined a [JSON schema for tracks](#) which you can use to check how to define your track. You should also check the tracks provided by Rally for inspiration.

The new track appears when you run `esrally list tracks --track-path=~ /rally-tracks/tutorial`:

```
dm@io:~ $ esrally list tracks --track-path=~ /rally-tracks/tutorial

  _____  _____
 /  _  \  _  \  /  _  \  /  _  \
/_  _  \/_  _  \/_  _  \/_  _  \
/_  _  \/_  _  \/_  _  \/_  _  \
/_  _  \/_  _  \/_  _  \/_  _  \
  _____  _____

Available tracks:

Name          Description          Documents    Compressed Size  Uncompressed_  
↪Size
-----
↪-----
tutorial      Tutorial benchmark for Rally    11658903    N/A              1.4 GB
```

Congratulations, you have created your first track! You can test it with `esrally --distribution-version=6.0.0 --track-path=~ /rally-tracks/tutorial`.

2.7.3 Adding support for test mode

You can check your track very quickly for syntax errors when you invoke Rally with `--test-mode`. Rally postprocesses its internal track representation as follows:

- Iteration-based tasks run at most one warmup iteration and one measurement iteration.
- Time-period-based tasks run at most for 10 seconds without warmup.

Rally also postprocesses all data file names. Instead of `documents.json`, Rally expects `documents-1k.json` and assumes the file contains 1.000 documents. You need to prepare these data files though. Pick 1.000 documents for every data file in your track and store them in a file with the suffix `-1k`. We choose the first 1.000 with `head -n 1000 documents.json > documents-1k.json`.

2.7.4 Challenges

To specify different workloads in the same track you can use so-called challenges. Instead of specifying the `schedule` property on top-level you specify a `challenges` array:

```
{
  "version": 2,
  "description": "Tutorial benchmark for Rally",
  "indices": [
    {
      "name": "geonames",
      "body": "index.json",

```

(continues on next page)

(continued from previous page)

```

    "types": [ "docs" ]
  }
],
"corpora": [
  {
    "name": "rally-tutorial",
    "documents": [
      {
        "source-file": "documents.json",
        "document-count": 11658903,
        "uncompressed-bytes": 1544799789
      }
    ]
  }
],
"challenges": [
  {
    "name": "index-and-query",
    "default": true,
    "schedule": [
      {
        "operation": {
          "operation-type": "delete-index"
        }
      },
      {
        "operation": {
          "operation-type": "create-index"
        }
      },
      {
        "operation": {
          "operation-type": "cluster-health",
          "request-params": {
            "wait_for_status": "green"
          }
        }
      },
      {
        "operation": {
          "operation-type": "bulk",
          "bulk-size": 5000
        },
        "warmup-time-period": 120,
        "clients": 8
      },
      {
        "operation": {
          "operation-type": "force-merge"
        }
      },
      {
        "operation": {
          "name": "query-match-all",
          "operation-type": "search",
          "body": {
            "query": {

```

(continues on next page)

(continued from previous page)

```

        "match_all": {}
    }
},
"clients": 8,
"warmup-iterations": 1000,
"iterations": 1000,
"target-throughput": 100
}
]
}
]
}

```

Note: If you define multiple challenges, Rally runs the challenge where `default` is set to `true`. If you want to run a different challenge, provide the command line option `--challenge=YOUR_CHALLENGE_NAME`.

When should you use challenges? Challenges are useful when you want to run completely different workloads based on the same track but for the majority of cases you should get away without using challenges:

- To run only a subset of the tasks, you can use *task filtering*, e.g. `--include-tasks="create-index, bulk"` will only run these two tasks in the track above.
- To vary parameters, e.g. the number of clients, you can use *track parameters*

2.7.5 Structuring your track

`track.json` is the entry point to a track but you can split your track as you see fit. Suppose you want to add more challenges to the track but keep them in separate files. Create a `challenges` directory and store the following in `challenges/index-and-query.json`:

```

{
  "name": "index-and-query",
  "default": true,
  "schedule": [
    {
      "operation": {
        "operation-type": "delete-index"
      }
    },
    {
      "operation": {
        "operation-type": "create-index"
      }
    },
    {
      "operation": {
        "operation-type": "cluster-health",
        "request-params": {
          "wait_for_status": "green"
        }
      }
    }
  ],
  {

```

(continues on next page)

(continued from previous page)

```

    "operation": {
      "operation-type": "bulk",
      "bulk-size": 5000
    },
    "warmup-time-period": 120,
    "clients": 8
  },
  {
    "operation": {
      "operation-type": "force-merge"
    }
  },
  {
    "operation": {
      "name": "query-match-all",
      "operation-type": "search",
      "body": {
        "query": {
          "match_all": {}
        }
      }
    },
    "clients": 8,
    "warmup-iterations": 1000,
    "iterations": 1000,
    "target-throughput": 100
  }
]
}

```

Include the new file in `track.json`:

```

{
  "version": 2,
  "description": "Tutorial benchmark for Rally",
  "indices": [
    {
      "name": "geonames",
      "body": "index.json",
      "types": [ "docs" ]
    }
  ],
  "corpora": [
    {
      "name": "rally-tutorial",
      "documents": [
        {
          "source-file": "documents.json",
          "document-count": 11658903,
          "uncompressed-bytes": 1544799789
        }
      ]
    }
  ],
  "challenges": [
    {% include "challenges/index-and-query.json" %}
  ]
}

```

(continues on next page)

(continued from previous page)

}

We replaced the challenge content with `{% include "challenges/index-and-query.json" %}` which tells Rally to include the challenge from the provided file. You can use `include` on arbitrary parts of your track.

To reuse operation definitions across challenges, you can define them in a separate `operations` block and refer to them by name in the corresponding challenge:

```
{
  "version": 2,
  "description": "Tutorial benchmark for Rally",
  "indices": [
    {
      "name": "geonames",
      "body": "index.json",
      "types": [ "docs" ]
    }
  ],
  "corpora": [
    {
      "name": "rally-tutorial",
      "documents": [
        {
          "source-file": "documents.json",
          "document-count": 11658903,
          "uncompressed-bytes": 1544799789
        }
      ]
    }
  ],
  "operations": [
    {
      "name": "delete",
      "operation-type": "delete-index"
    },
    {
      "name": "create",
      "operation-type": "create-index"
    },
    {
      "name": "wait-for-green",
      "operation-type": "cluster-health",
      "request-params": {
        "wait_for_status": "green"
      }
    },
    {
      "name": "bulk-index",
      "operation-type": "bulk",
      "bulk-size": 5000
    },
    {
      "name": "force-merge",
      "operation-type": "force-merge"
    },
    {
      "name": "query-match-all",
```

(continues on next page)

(continued from previous page)

```

    "operation-type": "search",
    "body": {
      "query": {
        "match_all": {}
      }
    }
  ],
  "challenges": [
    {% include "challenges/index-and-query.json" %}
  ]
}

```

challenges/index-and-query.json then becomes:

```

{
  "name": "index-and-query",
  "default": true,
  "schedule": [
    {
      "operation": "delete"
    },
    {
      "operation": "create"
    },
    {
      "operation": "wait-for-green"
    },
    {
      "operation": "bulk-index",
      "warmup-time-period": 120,
      "clients": 8
    },
    {
      "operation": "force-merge"
    },
    {
      "operation": "query-match-all",
      "clients": 8,
      "warmup-iterations": 1000,
      "iterations": 1000,
      "target-throughput": 100
    }
  ]
}

```

Note how we reference to the operations by their name (e.g. create, bulk-index, force-merge or query-match-all).

You can also use Rally's collect helper to simplify including multiple challenges:

```

{% import "rally.helpers" as rally %}
{
  "version": 2,
  "description": "Tutorial benchmark for Rally",
  "indices": [
    {

```

(continues on next page)

```
    "name": "geonames",
    "body": "index.json",
    "types": [ "docs" ]
  }
],
"corpora": [
  {
    "name": "rally-tutorial",
    "documents": [
      {
        "source-file": "documents.json",
        "document-count": 11658903,
        "uncompressed-bytes": 1544799789
      }
    ]
  }
],
"operations": [
  {
    "name": "delete",
    "operation-type": "delete-index"
  },
  {
    "name": "create",
    "operation-type": "create-index"
  },
  {
    "name": "wait-for-green",
    "operation-type": "cluster-health",
    "request-params": {
      "wait_for_status": "green"
    }
  },
  {
    "name": "bulk-index",
    "operation-type": "bulk",
    "bulk-size": 5000
  },
  {
    "name": "force-merge",
    "operation-type": "force-merge"
  },
  {
    "name": "query-match-all",
    "operation-type": "search",
    "body": {
      "query": {
        "match_all": {}
      }
    }
  }
],
"challenges": [
  { rally.collect(parts="challenges/*.json") }
]
}
```

The changes are:

1. We import helper functions from Rally by adding `{% import "rally.helpers" as rally %}` in line 1.
2. We use Rally's `collect` helper to find and include all JSON files in the `challenges` subdirectory with the statement `{{ rally.collect(parts="challenges/*.json") }}`.

Note: Rally's log file contains the fully rendered track after it has loaded it successfully.

You can even use [Jinja2 variables](#) but then you need to import the Rally helpers a bit differently. You also need to declare all variables before the `import` statement:

```
{% set clients = 16 %}
{% import "rally.helpers" as rally with context %}
```

If you use this idiom you can refer to the `clients` variable inside your snippets with `{{ clients }}`.

2.7.6 Sharing your track with others

So far the track is only available on your local machine. To share your track you could check it into version control. To avoid committing the potentially huge data file you can expose it via http (e.g. via an S3 bucket) and reference it in your track with the property `base-url`. Rally expects that the URL points to the parent path and appends the document file name automatically.

You should also compress your document corpus to save network bandwidth; `bzip2` works well. You can create a compressed archive with the following command:

```
bzip2 -9 -c documents.json > documents.json.bz2
```

If you want to support Rally's test mode, also compress your test mode corpus with:

```
bzip2 -9 -c documents-1k.json > documents-1k.json.bz2
```

Then upload `documents.json.bz2` and `documents-1k.json.bz2` to the remote location.

Finally, specify the compressed file name in the `source-file` property and also add the `base-url` property:

```
{
  "version": 2,
  "description": "Tutorial benchmark for Rally",
  "corpora": [
    {
      "name": "rally-tutorial",
      "documents": [
        {
          "base-url": "http://benchmarks.elasticsearch.org.s3.amazonaws.com/corpora/
↪geonames",
          "source-file": "documents.json.bz2",
          "document-count": 11658903,
          "compressed-bytes": 197857614,
          "uncompressed-bytes": 1544799789
        }
      ]
    }
  ]
},
```

(continues on next page)

(continued from previous page)

```
...
}
```

Specifying `compressed-bytes` (file size of `documents.json.bz2`) and `uncompressed-bytes` (file size of `documents.json`) is optional but helps Rally to provide progress indicators and also verify integrity.

You've now mastered the basics of track development for Rally. It's time to pat yourself on the back before you dive into the advanced topics!

2.7.7 Advanced topics

Template Language

Rally uses [Jinja2](#) as a template language so you can use Jinja2 expressions in track files.

Extensions

Rally also provides a few extensions to Jinja2:

- `now`: a global variable that represents the current date and time when the template is evaluated by Rally.
- `days_ago()`: a [filter](#) that you can use for date calculations.

You can find an example in the `http_logs` track:

```
{
  "name": "range",
  "index": "logs-*",
  "type": "type",
  "body": {
    "query": {
      "range": {
        "@timestamp": {
          "gte": "now-{{ '15-05-1998' | days_ago(now) }}d/d",
          "lt": "now/d"
        }
      }
    }
  }
}
```

The data set that is used in the `http_logs` track starts on 26-04-1998 but we want to ignore the first few days for this query, so we start on 15-05-1998. The expression `{{ '15-05-1998' | days_ago(now) }}` yields the difference in days between now and the fixed start date and allows us to benchmark time range queries relative to now with a predetermined data set.

Custom parameter sources

Warning: Your parameter source is on a performance-critical code-path. Double-check with *Rally's profiling support* that you did not introduce any bottlenecks.

Consider the following operation definition:

```
{
  "name": "term",
  "operation-type": "search",
  "body": {
    "query": {
      "term": {
        "body": "physician"
      }
    }
  }
}
```

This query is defined statically but if you want to vary parameters, for example to search also for “mechanic” or “nurse”, you can write your own “parameter source” in Python.

First, define the name of your parameter source in the operation definition:

```
{
  "name": "term",
  "operation-type": "search",
  "param-source": "my-custom-term-param-source"
  "professions": ["mechanic", "physician", "nurse"]
}
```

Rally recognizes the parameter source and looks for a file `track.py` next to `track.json`. This file contains the implementation of the parameter source:

```
import random

def random_profession(track, params, **kwargs):
    # choose a suitable index: if there is only one defined for this track
    # choose that one, but let the user always override index and type.
    if len(track.indices) == 1:
        default_index = track.indices[0].name
        if len(track.indices[0].types) == 1:
            default_type = track.indices[0].types[0].name
        else:
            default_type = None
    else:
        default_index = "_all"
        default_type = None

    index_name = params.get("index", default_index)
    type_name = params.get("type", default_type)

    # you must provide all parameters that the runner expects
    return {
        "body": {
            "query": {
                "term": {
                    "body": "%s" % random.choice(params["professions"])
                }
            }
        },
        "index": index_name,
```

(continues on next page)

(continued from previous page)

```

        "type": type_name,
        "cache": params.get("cache", False)
    }

def register(registry):
    registry.register_param_source("my-custom-term-param-source", random_profession)

```

The example above shows a simple case that is sufficient if the operation to which your parameter source is applied is idempotent and it does not matter whether two clients execute the same operation.

The function `random_profession` is the actual parameter source. Rally will bind the name “my-custom-term-param-source” to this function by calling `register`. `register` is called by Rally before the track is executed.

The parameter source function needs to declare the parameters `track`, `params` and `**kwargs`. `track` contains a structured representation of the current track and `params` contains all parameters that have been defined in the operation definition in `track.json`. We use it in the example to read the professions to choose. The third parameter is there to ensure a more stable API as Rally evolves.

We also derive an appropriate index and document type from the track’s index definitions but allow the user to override this choice with the `index` or `type` parameters:

```

{
  "name": "term",
  "operation-type": "search",
  "param-source": "my-custom-term-param-source",
  "professions": ["mechanic", "physician", "nurse"],
  "index": "employee*",
  "type": "docs"
}

```

If you need more control, you need to implement a class. Below is the implementation of the same parameter source as a class:

```

import random

class TermParamSource:
    def __init__(self, track, params, **kwargs):
        # choose a suitable index: if there is only one defined for this track
        # choose that one, but let the user always override index and type.
        if len(track.indices) == 1:
            default_index = track.indices[0].name
            if len(track.indices[0].types) == 1:
                default_type = track.indices[0].types[0].name
            else:
                default_type = None
        else:
            default_index = "_all"
            default_type = None

        # we can eagerly resolve these parameters already in the constructor...
        self._index_name = params.get("index", default_index)
        self._type_name = params.get("type", default_type)
        self._cache = params.get("cache", False)
        # ... but we need to resolve "profession" lazily on each invocation later
        self._params = params

```

(continues on next page)

(continued from previous page)

```

def partition(self, partition_index, total_partitions):
    return self

def size(self):
    return 1

def params(self):
    # you must provide all parameters that the runner expects
    return {
        "body": {
            "query": {
                "term": {
                    "body": "%s" % random.choice(self._params["professions"])
                }
            }
        },
        "index": self._index_name,
        "type": self._type_name,
        "cache": self._cache
    }

def register(registry):
    registry.register_param_source("my-custom-term-param-source", TermParamSource)

```

In `register` you bind the name in the track specification to your parameter source implementation class similar to the previous example. `TermParamSource` is the actual parameter source and needs to fulfill a few requirements:

- The constructor needs to have the signature `__init__(self, track, params, **kwargs)`.
- `partition(self, partition_index, total_partitions)` is called by Rally to “assign” the parameter source across multiple clients. Typically you can just return `self`. If each client needs to act differently then you can provide different parameter source instances here as well.
- `size(self)`: This method helps Rally to provide a proper progress indication to users if you use a warmup time period. For bulk indexing, return the number of bulks (for a given client). As searches are typically executed with a pre-determined amount of iterations, just return 1 in this case.
- `params(self)`: This method returns a dictionary with all parameters that the corresponding “runner” expects. This method will be invoked once for every iteration during the race. In the example, we parameterize the query by randomly selecting a profession from a list.

For cases, where you want to provide a progress indication but cannot calculate `size` up-front (e.g. when you generate bulk requests on-the fly up to a certain total size), you can implement a property `percent_completed` which returns a floating point value between 0.0 and 1.0. Rally will query this value before each call to `params()` and uses it to indicate progress. However:

- Rally will not check `percent_completed` if it can derive progress in any other way.
- The value of `percent_completed` is purely informational and does not influence when Rally considers an operation to be completed.

Note: The method `params(self)` is called on a performance-critical path. Don’t do anything in this method that takes a lot of time (avoid any I/O). For searches, you should usually throttle throughput anyway and there it does not matter that much but if the corresponding operation is run without throughput throttling, double-check that your custom parameter source does not introduce a bottleneck.

Custom parameter sources can use the Python standard API but using any additional libraries is not supported.

You can also implement your parameter sources and runners in multiple Python files but the main entry point is always `track.py`. The root package name of your plugin is the name of your track.

Custom runners

Warning: Your runner is on a performance-critical code-path. Double-check with *Rally's profiling support* that you did not introduce any bottlenecks.

Runners execute an operation against Elasticsearch. Rally supports many operations out of the box already, see the *track reference* for a complete list. If you want to call any other Elasticsearch API, define a custom runner.

Consider we want to use the percolate API with an older version of Elasticsearch which is not supported by Rally. To achieve this, we implement a custom runner in the following steps.

In `track.json` set the `operation-type` to "percolate" (you can choose this name freely):

```
{
  "name": "percolator_with_content_google",
  "operation-type": "percolate",
  "body": {
    "doc": {
      "body": "google"
    },
    "track_scores": true
  }
}
```

Then create a file `track.py` next to `track.json` and implement the following two functions:

```
def percolate(es, params):
    es.percolate(
        index="queries",
        doc_type="content",
        body=params["body"]
    )

def register(registry):
    registry.register_runner("percolate", percolate)
```

The function `percolate` is the actual runner and takes the following parameters:

- `es`, is an instance of the Elasticsearch Python client
- `params` is a dict of parameters provided by its corresponding parameter source. Treat this parameter as read-only.

This function can return:

- Nothing at all. Then Rally will assume by default 1 and "ops" (see below).
- A tuple of `weight` and a `unit`, which is usually 1 and "ops". If you run a bulk operation you might return the bulk size here, for example in number of documents or in MB. Then you'd return for example (5000, "docs") Rally will use these values to store throughput metrics.

- A dict with arbitrary keys. If the dict contains the key `weight` it is assumed to be numeric and chosen as weight as defined above. The key `unit` is treated similarly. All other keys are added to the `meta` section of the corresponding service time and latency metrics records.

Similar to a parameter source you also need to bind the name of your operation type to the function within `register`.

If you need more control, you can also implement a runner class. The example above, implemented as a class looks as follows:

```
class PercolateRunner:
    def __call__(self, es, params):
        es.percolate(
            index="queries",
            doc_type="content",
            body=params["body"]
        )

    def __repr__(self, *args, **kwargs):
        return "percolate"

def register(registry):
    registry.register_runner("percolate", PercolateRunner())
```

The actual runner is implemented in the method `__call__` and the same return value conventions apply as for functions. For debugging purposes you should also implement `__repr__` and provide a human-readable name for your runner. Finally, you need to register your runner in the `register` function. Runners also support Python's `context manager` interface. Rally uses a new context for each request. Implementing the context manager interface can be handy for cleanup of resources after executing an operation. Rally uses it, for example, to clear open scrolls.

If you have specified multiple Elasticsearch clusters using `target-hosts` you can make Rally pass a dictionary of client connections instead of one for the default cluster in the `es` parameter.

To achieve this you need to:

- Use a runner class
- Specify `multi_cluster = True` as a class attribute
- Use any of the cluster names specified in `target-hosts` as a key for the `es` dict

Example (assuming Rally has been invoked specifying default and remote in `target-hosts`):

```
class CreateIndexInRemoteCluster:
    multi_cluster = True

    def __call__(self, es, params):
        es['remote'].indices.create(index='remote-index')

    def __repr__(self, *args, **kwargs):
        return "create-index-in-remote-cluster"

def register(registry):
    registry.register_runner("create-index-in-remote-cluster",
↵CreateIndexInRemoteCluster())
```

Note: You need to implement `register` just once and register all parameter sources and runners there.

Custom schedulers

Warning: Your scheduler is on a performance-critical code-path. Double-check with *Rally's profiling support* that you did not introduce any bottlenecks.

If you want to rate-limit execution of tasks, you can specify a `target-throughput` (in operations per second). For example, Rally attempts to run this term query 20 times per second:

```
{
  "operation": "term",
  "target-throughput": 20
}
```

By default, Rally uses a [deterministic distribution](#) to determine when to schedule the next operation. Hence it executes the term query at 0, 50ms, 100ms, 150ms and so on. The scheduler is also aware of the number of clients. Consider this example:

```
{
  "operation": "term",
  "target-throughput": 20,
  "clients": 4
}
```

If Rally would not take the number of clients into account and would still issue requests (from each of the four clients) at the same points in time (i.e. 0, 50ms, 100ms, 150ms, ...), it would run at a target throughput of $4 * 20 = 80$ operations per second. Hence, Rally will automatically reduce the rate at which each client will execute requests. Each client will issue requests at 0, 200ms, 400ms, 600ms, 800ms, 1000ms and so on. Each client issues five requests per second but as there are four of them, we still have a target throughput of 20 operations per second. You should keep this in mind, when writing your own custom schedules.

To create a custom scheduler, create a file `track.py` next to `track.json` and implement the following two functions:

```
import random

def random_schedule(current):
    return current + random.randint(10, 900) / 1000.0

def register(registry):
    registry.register_scheduler("my_random", random_schedule)
```

You can then use your custom scheduler as follows:

```
{
  "operation": "term",
  "schedule": "my_random"
}
```

The function `random_schedule` returns a floating point number which represents the next point in time when Rally should execute the given operation. This point in time is measured in seconds relative to the beginning of the execution of this task. The parameter `current` is the last return value of your function and is `0.0` for the first invocation. So, for example, this scheduler could return the following series: 0, 0.119, 0.622, 1.29, 1.343, 1.984, 2.233.

This implementation is usually not sufficient as it does not take into account the number of clients. Therefore, you typically want to implement a full-blown scheduler which can also take parameters. Below is an example for our

random scheduler:

```
import random

class RandomScheduler:
    def __init__(self, params):
        # assume one client by default
        clients = self.params.get("clients", 1)
        # scale accordingly with the number of clients!
        self.lower_bound = clients * self.params.get("lower-bound-millis", 10)
        self.upper_bound = clients * self.params.get("upper-bound-millis", 900)

    def next(self, current):
        return current + random.randint(self.lower_bound, self.upper_bound) / 1000.0

def register(registry):
    registry.register_scheduler("my_random", RandomScheduler)
```

This implementation achieves the same rate independent of the number of clients. Additionally, we can pass the lower and upper bound for the random function from the track:

```
{
  "operation": "term",
  "schedule": "my_random",
  "clients": 4,
  "lower-bound-millis": 50,
  "upper-bound-millis": 250
}
```

2.8 Developing Rally

2.8.1 Prerequisites

Install the following software packages:

- Python 3.4 or better available as `python3` on the path (verify with: `python3 --version`)
- `pip3` available on the path (verify with `pip3 --version`)
- JDK 8 or 9
- git 1.9 or better

Check the *installation guide* for detailed installation instructions for these packages.

Rally does not support Windows and is only actively tested on MacOS and Linux.

2.8.2 Installation Instructions for Development

```
git clone https://github.com/elastic/rally.git
cd rally
virtualenv -p python3 .venv
source .venv/bin/activate
./rally
```

If you get errors during installation, it is probably due to the installation of `psutil` which we use to gather system metrics like CPU utilization. Check the [installation instructions of psutil](#) in this case. Keep in mind that Rally is based on Python 3 and you need to install the Python 3 header files instead of the Python 2 header files on Linux.

Automatic Updates

Rally has a built-in auto-update feature when you install it from sources. By default, it will update from the remote named `origin`. If you want to auto-update from a different remote, provide `--update-from-remote=YOUR_REMOTE_NAME` as first parameter.

To work conveniently with Rally, we suggest that you add the Rally project directory to your `PATH`. In case you use a different remote, you should also define aliases in your shell's config file, e.g.:

```
alias rally='rally --update-from-remote=elastic '
alias rallyd='rallyd --update-from-remote=elastic '
```

Then you can invoke Rally or the *Rally daemon* as usual and have auto-update still work.

Also note that automatic updates are disabled in the following cases:

- There are local (uncommitted) changes in the Rally project directory
- A different branch than `master` is checked out
- You have specified `--skip-update` as the first command line parameter
- You have specified `--offline` as a command line parameter for Rally

Configuring Rally

Before we can run our first benchmark, we have to configure Rally. Just invoke `./rally configure` and Rally will automatically detect that its configuration file is missing and prompt you for some values and write them to `~/.rally/rally.ini`. After you've configured Rally, it will exit.

For more information see [configuration help page](#).

2.8.3 Key Components of Rally

To get a rough understanding of Rally, it makes sense to get to know its key components:

- *Race Control*: is responsible for proper execution of the race. It sets up all components and acts as a high-level controller.
- *Mechanic*: can build and prepare a benchmark candidate for the race. It checks out the source, builds Elasticsearch, provisions and starts the cluster.
- *Track*: is a concrete benchmarking scenario, e.g. the `http_logs` benchmark. It defines the data set to use.
- *Challenge*: is the specification on what benchmarks should be run and its configuration (e.g. `index`, then run a search benchmark with 1000 iterations)
- *Car*: is a concrete system configuration for a benchmark, e.g. an Elasticsearch single-node cluster with default settings.
- *Driver*: drives the race, i.e. it is executing the benchmark according to the track specification.
- *Reporter*: A reporter tells us how the race went (currently only after the fact).

There is a dedicated [tutorial on how to add new tracks to Rally](#).

2.8.4 How to contribute code

See the [contributors guide](#). We strive to be PEP-8 compliant but don't follow it to the letter.

2.9 Command Line Reference

You can control Rally with subcommands and command line flags:

- Subcommands determine which task Rally performs.
- Command line flags are used to change Rally's behavior but not all command line flags can be used for each subcommand. To find out which command line flags are supported by a specific subcommand, just run `esrally <<subcommand>> --help`.

2.9.1 Subcommands

race

The `race` subcommand is used to actually run a benchmark. It is the default one and chosen implicitly if none is given.

list

The `list` subcommand is used to list different configuration options:

- `telemetry`: Will show all *telemetry devices* that are supported by Rally.
- `tracks`: Will show all tracks that are supported by Rally. As this *may* depend on the Elasticsearch version that you want to benchmark, you can specify `--distribution-version` and also `--distribution-repository` as additional options.
- `pipelines`: Will show all *pipelines* that are supported by Rally.
- `racess`: Will show a list of the most recent races. This is needed for the *tournament mode*.
- `cars`: Will show all cars that are supported by Rally (i.e. Elasticsearch configurations).
- `elasticsearch-plugins`: Will show all Elasticsearch plugins and their configurations that are supported by Rally.

To list a specific configuration option, place it after the `list` subcommand. For example, `esrally list pipelines` will list all pipelines known to Rally.

compare

This subcommand is needed for *tournament mode* and its usage is described there.

configure

This subcommand is needed to *configure* Rally. It is implicitly chosen if you start Rally for the first time but you can rerun this command at any time.

2.9.2 Command Line Flags

track-path

Can be either a directory that contains a `track.json` file or a `.json` file with an arbitrary name that contains a track specification. `--track-path` and `--track-repository` as well as `--track` are mutually exclusive. See the [track reference](#) to decide whether you should use `--track-path` or `--track-repository` / `--track`.

Examples:

```
# provide a directory - Rally searches for a track.json file in this directory
# Track name is "app-logs"
esrally --track-path=~/.Projects/tracks/app-logs
# provide a file name - Rally uses this file directly
# Track name is "syslog"
esrally --track-path=~/.Projects/tracks/syslog.json
```

track-repository

Selects the track repository that Rally should use to resolve tracks. By default the default track repository is used, which is available in the Github project [rally-tracks](#). See the [track reference](#) on how to add your own track repositories. `--track-path` and `--track-repository` as well as `--track` are mutually exclusive.

track

Selects the track that Rally should run. By default the `geonames` track is run. For more details on how tracks work, see [adding tracks](#) or the [track reference](#). `--track-path` and `--track-repository` as well as `--track` are mutually exclusive.

track-params

With this parameter you can inject variables into tracks. The supported variables depend on the track and you should check the track JSON file to see which variables can be provided.

It accepts a list of comma-separated key-value pairs or a JSON file name. The key-value pairs have to be delimited by a colon.

Examples:

Consider the following track snippet showing a single challenge:

```
{
  "name": "index-only",
  "schedule": [
    {
      "operation": {
        "operation-type": "bulk",
        "bulk-size": {{ bulk_size|default(5000) }}
      },
      "warmup-time-period": 120,
      "clients": {{ clients|default(8) }}
    }
  ]
}
```

Rally tracks can use the Jinja templating language and the construct `{{ some_variable|default(0) }}` that you can see above is a [feature of Jinja](#) to define default values for variables.

We can see that it defines two variables:

- `bulk_size` with a default value of 5000
- `clients` with a default value of 8

When we run this track, we can override these defaults:

- `--track-params="bulk_size:2000,clients:16"` will set the bulk size to 2000 and the number of clients for bulk indexing to 16.
- `--track-params="bulk_size:8000"` will just set the bulk size to 8000 and keep the default value of 8 clients.
- `--track-params="params.json"` will read the track parameters from a JSON file (defined below)

Example JSON file:

```
{
  "bulk_size": 2000,
  "clients": 16
}
```

All track parameters are recorded for each metrics record in the metrics store. Also, when you run `esrally list races`, it will show all track parameters:

Race Timestamp	Track	Track Parameters	Challenge	Car	
<code>↔User Tag</code>					<code>↔</code>
<code>↔-----</code>					
20160518T122341Z	pmc	bulk_size=8000	append-no-conflicts	defaults	
20160518T112341Z	pmc	bulk_size=2000,clients=16	append-no-conflicts	defaults	

Note that the default values are not recorded or shown (Rally does not know about them).

challenge

A track consists of one or more challenges. With this flag you can specify which challenge should be run. If you don't specify a challenge, Rally derives the default challenge itself. To see the default challenge of a track, run `esrally list tracks`.

include-tasks

Each challenge consists of one or more tasks but sometimes you are only interested to run a subset of all tasks. For example, you might have prepared an index already and want only to repeatedly run search benchmarks. Or you want to run only the indexing task but nothing else.

You can use `--include-tasks` to specify a comma-separated list of tasks that you want to run. Each item in the list defines either the name of a task or the operation type of a task. Only the tasks that match will be executed. Currently there is also no command that lists the tasks of a challenge so you need to look at the track source.

Examples:

- Execute only the tasks with the name `index` and term: `--include-tasks="index,term"`
- Execute only tasks of type `search`: `--include-tasks="type:search"`

- You can also mix and match: `--include-tasks="index,type:search"`

team-repository

Selects the team repository that Rally should use to resolve cars. By default the `default` team repository is used, which is available in the Github project `rally-teams`. See the documentation about *cars* on how to add your own team repositories.

team-path

A directory that contains a team configuration. `--team-path` and `--team-repository` are mutually exclusive. See the *car reference* for the required directory structure.

Example:

```
esrally --team-path=~ /Projects/es-teams
```

car

A *car* defines the Elasticsearch configuration that will be used for the benchmark. To see a list of possible cars, issue `esrally list cars`. You can specify one or multiple comma-separated values.

Example

```
esrally --car="4gheap,ea"
```

Rally will configure Elasticsearch with 4GB of heap (`4gheap`) and enable Java assertions (`ea`).

car-params

Allows to override config variables of Elasticsearch. It accepts a list of comma-separated key-value pairs or a JSON file name. The key-value pairs have to be delimited by a colon.

Example

```
esrally --car="4gheap" --car-params="data_paths:'/opt/elasticsearch'"
```

The variables that are exposed depend on the *car's configuration*. In addition, Rally implements special handling for the variable `data_paths` (by default the value for this variable is determined by Rally).

elasticsearch-plugins

A comma-separated list of Elasticsearch plugins to install for the benchmark. If a plugin supports multiple configurations you need to specify the configuration after the plugin name. To see a list of possible plugins and configurations, issue `esrally list elasticsearch-plugins`.

Example:

```
esrally --elasticsearch-plugins="analysis-icu,xpack:security"
```

In this example, Rally will install the `analysis-icu` plugin and the `x-pack` plugin with the `security` configuration. See the reference documentation about *Elasticsearch plugins* for more details.

plugin-params

Allows to override variables of Elasticsearch plugins. It accepts a list of comma-separated key-value pairs or a JSON file name. The key-value pairs have to be delimited by a colon.

Example:

```
esrally --distribution-version=6.1.1. --elasticsearch-plugins="x-pack:monitoring-http
↪" --plugin-params="monitoring_type:'https',monitoring_host:'some_remote_host',
↪monitoring_port:10200,monitoring_user:'rally',monitoring_password:'m0nlt0rlng'"
```

This enables the HTTP exporter of [X-Pack Monitoring](#) and exports the data to the configured monitoring host.

pipeline

Selects the *pipeline* that Rally should run.

Rally can autodetect the pipeline in most cases. If you specify `--distribution-version` it will auto-select the pipeline `from-distribution` otherwise it will use `from-sources-complete`.

laps

Allows to run the benchmark for multiple laps (defaults to 1 lap). Each lap corresponds to one full execution of a track but note that the benchmark candidate is not restarted in between.

enable-driver-profiling

This option enables a profiler on all tasks that the load test driver performs. It is intended to help track authors spot accidental bottlenecks, especially if they implement their own runners or parameter sources. When this mode is enabled, Rally will enable a profiler in the load driver module. After each task and for each client, Rally will add the profile information to a dedicated profile log file. For example:

```
2017-02-09 08:23:24,35 rally.profile INFO
=== Profile START for client [0] and task [index-append-1000] ===
  16052402 function calls (15794402 primitive calls) in 180.221 seconds

  Ordered by: cumulative time

  ncalls  tottime  percall  cumtime  percall  filename:lineno(function)
  130     0.001    0.000    168.089    1.293    /Users/dm/Projects/rally/esrally/driver/
↪driver.py:908(time_period_based)
  129     0.260    0.002    168.088    1.303    /Users/dm/.rally/benchmarks/tracks/
↪develop/bottleneck/parameter_sources/bulk_source.py:79(params)
  129000  0.750    0.000    167.791    0.001    /Users/dm/.rally/benchmarks/tracks/
↪develop/bottleneck/parameter_sources/randomevent.py:142(generate_event)
  516000  0.387    0.000    160.485    0.000    /Users/dm/.rally/benchmarks/tracks/
↪develop/bottleneck/parameter_sources/weightedarray.py:20(get_random)
  516000  6.199    0.000    160.098    0.000    /Users/dm/.rally/benchmarks/tracks/
↪develop/bottleneck/parameter_sources/weightedarray.py:23(__random_index)
  516000  1.292    0.000    152.289    0.000    /usr/local/Cellar/python3/3.6.0/
↪Frameworks/Python.framework/Versions/3.6/lib/python3.6/random.py:96(seed)
  516000  150.783  0.000    150.783    0.000    {function Random.seed at 0x10b7fa2f0}
  129000  0.363    0.000    45.686    0.000    /Users/dm/.rally/benchmarks/tracks/
↪develop/bottleneck/parameter_sources/randomevent.py:48(add_fields)
  129000  0.181    0.000    41.742    0.000    /Users/dm/.rally/benchmarks/tracks/
↪develop/bottleneck/parameter_sources/randomevent.py:79(add_fields) (continues on next page)
```

(continued from previous page)

```
....  
=== Profile END for client [0] and task [index-append-1000] ===
```

In this example we can spot quickly that `Random.seed` is called excessively, causing an accidental bottleneck in the load test driver.

test-mode

Allows you to test a track without running it for the whole duration. This mode is only intended for quick sanity checks when creating a track. Don't rely on these numbers at all (they are meaningless).

If you write your own track you need to *prepare your track to support this mode*.

telemetry

Activates the provided *telemetry devices* for this race.

Example

```
esrally --telemetry=jfr,jit
```

This activates Java flight recorder and the JIT compiler telemetry devices.

telemetry-params

Allows to set parameters for telemetry devices. It accepts a list of comma-separated key-value pairs or a JSON file name. The key-value pairs have to be delimited by a colon. See the *telemetry devices* documentation for a list of supported parameters.

Example:

```
esrally --telemetry=jfr --telemetry-params="recording-template:'profile'"
```

This enables the Java flight recorder telemetry device and sets the `recording-template` parameter to "profile".

runtime-jdk

By default, Rally will derive the appropriate runtime JDK versions automatically per version of Elasticsearch. For example, it will choose JDK 8 or 7 for Elasticsearch 2.x but only JDK 8 for Elasticsearch 5.0.0. It will choose the highest available version.

This command line parameter sets the major version of the JDK that Rally should use to run Elasticsearch. It is required that either `JAVA_HOME` or `JAVAx_HOME` (where `x` is the major version, e.g. `JAVA8_HOME` for a JDK 8) points to the appropriate JDK.

Example:

```
# Run a benchmark with defaults (i.e. JDK 8)  
esrally --distribution-version=2.4.0  
# Force to run with JDK 7  
esrally --distribution-version=2.4.0 --runtime-jdk=7
```

revision

If you actively develop Elasticsearch and want to benchmark a source build of Elasticsearch (which Rally will create for you), you can specify the git revision of Elasticsearch that you want to benchmark. But note that Rally uses and expects the Gradle Wrapper in the Elasticsearch repository (`./gradlew`) which effectively means that it will only support this for Elasticsearch 5.0 or better. The default value is `current`.

You can specify the revision in different formats:

- `--revision=latest`: Use the HEAD revision from origin/master.
- `--revision=current`: Use the current revision (i.e. don't alter the local source tree).
- `--revision=abc123`: Where `abc123` is some git revision hash.
- `--revision=@2013-07-27T10:37:00Z`: Determines the revision that is closest to the provided date. Rally logs to which git revision hash the date has been resolved and if you use Elasticsearch as metrics store (instead of the default in-memory one), *each metric record will contain the git revision hash also in the meta-data section.*

Supported date format: If you specify a date, it has to be ISO-8601 conformant and must start with an @ sign to make it easier for Rally to determine that you actually mean a date.

If you want to create source builds of Elasticsearch plugins, you need to specify the revision for Elasticsearch and all relevant plugins separately. Revisions for Elasticsearch and each plugin need to be comma-separated (`,`). Each revision is prefixed either by `elasticsearch` or by the plugin name and separated by a colon (`:`). As core plugins are contained in the Elasticsearch repo, there is no need to specify a revision for them (the revision would even be ignored in fact).

Examples:

- Build latest Elasticsearch and plugin “my-plugin”: `--revision="elasticsearch:latest,my-plugin:latest"`
- Build Elasticsearch tag `v5.6.1` and revision `abc123` of plugin “my-plugin”: `--revision="elasticsearch:v5.6.1,my-plugin:abc123"`

Note that it is still required to provide the parameter `--elasticsearch-plugins`. Specifying a plugin with `--revision` just tells Rally which revision to use for building the artifact. See the documentation on [Elasticsearch plugins](#) for more details.

distribution-version

If you want to benchmark a binary distribution, you can specify the version here.

Example

```
esrally --distribution-version=2.3.3
```

Rally will then benchmark the official Elasticsearch 2.3.3 distribution.

Rally works with all releases of Elasticsearch that are [supported by Elastic](#).

The following versions are already end-of-life:

- `0.x`: Rally is not tested, and not expected to work for this version; we will make no effort to make Rally work.
- `1.x`: Rally works on a best-effort basis with this version but support may be removed at any time.

Additionally, Rally will always work with the current development version of Elasticsearch (by using either a snapshot repository or by building Elasticsearch from sources).

distribution-repository

Rally does not only support benchmarking official distributions but can also benchmark snapshot builds. This option is really just intended for [our benchmarks that are run in continuous integration](#) but if you want to, you can use it too. The only supported value out of the box is `release` (default) but you can define arbitrary repositories in `~/.rally/rally.ini`.

Example

Say, you have an in-house repository where Elasticsearch snapshot builds get published. Then you can add the following in the `distributions` section of your Rally config file:

```
in_house_snapshot.url = https://www.example.org/snapshots/elasticsearch/elasticsearch-
↳ {{VERSION}}.tar.gz
in_house_snapshot.cache = false
```

The `url` property defines the URL pattern for this repository. The `cache` property defines whether Rally should always download a new archive (`cache=false`) or just reuse a previously downloaded version (`cache=true`). Rally will replace the `{{VERSION}}` placeholder of in the `url` property with the value of `distribution-version` provided by the user on the command line.

You can use this distribution repository with the name “`in_house_snapshot`” as follows:

```
esrally --distribution-repository=in_house_snapshot --distribution-version=7.0.0-
↳ SNAPSHOT
```

This will benchmark the latest 7.0.0 snapshot build of Elasticsearch.

report-format

The command line reporter in Rally displays a table with key metrics after a race. With this option you can specify whether this table should be in `markdown` format (default) or `csv`.

show-in-report

By default, the command line reporter will only show values that are available (`available`). With `all` you can force it to show a line for every value, even undefined ones, and with `all-percentiles` it will show only available values but force output of all possible percentile values.

This command line parameter is not available for comparing races.

report-file

By default, the command line reporter will print the results only on standard output, but can also write it to a file.

Example

```
esrally --report-format=csv --report-file=~/.benchmarks/result.csv
```

client-options

With this option you can customize Rally’s internal Elasticsearch client.

It accepts a list of comma-separated key-value pairs. The key-value pairs have to be delimited by a colon. These options are passed directly to the Elasticsearch Python client API. See [their documentation on a list of supported options](#).

We support the following data types:

- Strings: Have to be enclosed in single quotes. Example: `ca_certs: '/path/to/CA_certs'`
- Numbers: There is nothing special about numbers. Example: `sniffer_timeout: 60`
- Booleans: Specify either `true` or `false`. Example: `use_ssl: true`

In addition to the options, supported by the Elasticsearch client, it is also possible to enable HTTP compression by specifying `compressed: true`

Default value: `timeout: 60`

Warning: If you provide your own client options, the default value will not be magically merged. You have to specify all client options explicitly. The only exceptions to this rule is `ca_cert` (see below).

Examples

Here are a few common examples:

- Enable HTTP compression: `--client-options="compressed:true"`
- Enable SSL (e.g. if you have X-Pack Security installed): `--client-options="use_ssl:true, verify_certs:true"`. Note that you don't need to set `ca_cert` (which defines the path to the root certificates). Rally does this automatically for you.
- Enable SSL with a client key and certificate: `--client-options="use_ssl:true, verify_certs:true, ca_certs: '/path/to/cacert.pem', client_cert: '/path/to/client_cert.pem', client_key='/path/to/client_key.pem'"` (see also the [Elasticsearch Python client docs](#))
- Enable basic authentication: `--client-options="basic_auth_user:'user', basic_auth_password:'password'"`. Avoid the characters `,` `,` and `:` in user name and password as Rally's parsing of these options is currently really simple and there is no possibility to escape characters.

on-error

This option controls whether Rally will `continue` or `abort` when a request error occurs. By default, Rally will just record errors and report the error rate at the end of a race. With `--on-error=abort`, Rally will immediately abort the race on the first error and print a detailed error message.

load-driver-hosts

By default, Rally will run its load driver on the same machine where you start the benchmark. However, if you benchmark larger clusters, one machine may not be enough to generate sufficient load. Hence, you can specify a comma-separated list of hosts which should be used to generate load with `--load-driver-hosts`.

Example

```
esrally --load-driver-hosts=10.17.20.5,10.17.20.6
```

In the example, above Rally will generate load from the hosts `10.17.20.5` and `10.17.20.6`. For this to work, you need to start a Rally daemon on these machines, see [distributing the load test driver](#) for a complete example.

target-hosts

If you run the `benchmark-only` pipeline or you want Rally to *benchmark a remote cluster*, then you can specify a comma-delimited list of `hosts:port` pairs to which Rally should connect. The default value is `127.0.0.1:9200`.

Example

```
esrally --pipeline=benchmark-only --target-hosts=10.17.0.5:9200,10.17.0.6:9200
```

This will run the benchmark against the hosts 10.17.0.5 and 10.17.0.6 on port 9200. See `client-options` if you use X-Pack Security and need to authenticate or Rally should use https.

You can also target multiple clusters with `--target-hosts` for specific use cases. This is described in the *Advanced topics section*.

quiet

Suppresses some output on the command line.

offline

Tells Rally that it should assume it has no connection to the Internet when checking for track data. The default value is `false`. Note that Rally will only assume this for tracks but not for anything else, e.g. it will still try to download Elasticsearch distributions that are not locally cached or fetch the Elasticsearch source tree.

preserve-install

Rally usually installs and launches an Elasticsearch cluster internally and wipes the entire directory after the benchmark is done. Sometimes you want to keep this cluster including all data after the benchmark has finished and that's what you can do with this flag. Note that depending on the track that has been run, the cluster can eat up a very significant amount of disk space (at least dozens of GB). The default value is configurable in the advanced configuration but usually `false`.

Note: This option does only affect clusters that are provisioned by Rally. More specifically, if you use the pipeline `benchmark-only`, this option is ineffective as Rally does not provision a cluster in this case.

advanced-config

This flag determines whether Rally should present additional (advanced) configuration options. The default value is `false`.

Example

```
esrally configure --advanced-config
```

assume-defaults

This flag determines whether Rally should automatically accept all values for configuration options that provide a default. This is mainly intended to configure Rally automatically in CI runs. The default value is `false`.

Example

```
esrally configure --assume-defaults=true
```

user-tag

This is only relevant when you want to run *tournaments*. You can use this flag to attach an arbitrary text to the meta-data of each metric record and also the corresponding race. This will help you to recognize a race when you run `esrally list races` as you don't need to remember the concrete timestamp on which a race has been run but can instead use your own descriptive names.

The required format is key ":" value. You can choose key and value freely.

Example

```
esrally --user-tag="intention:github-issue-1234-baseline,gc:cms"
```

You can also specify multiple tags. They need to be separated by a comma.

Example

```
esrally --user-tag="disk:SSD,data_node_count:4"
```

When you run `esrally list races`, this will show up again:

Race	Timestamp	Track	Track Parameters	Challenge	Car	User Tag
↳	-----					
20160518T122341Z		pmc		append-no-conflicts	defaults	↳
↳			intention:github-issue-1234-baseline			
20160518T112341Z		pmc		append-no-conflicts	defaults	disk:SSD,
↳			data_node_count:4			

This will help you recognize a specific race when running `esrally compare`.

2.9.3 Advanced topics

target-hosts

Rally can also create client connections for multiple Elasticsearch clusters. This is only useful if you want to create *custom runners* that execute operations against multiple clusters, for example for [cross cluster search](#) or cross cluster replication.

To define the host:port pairs for additional clusters with `target-hosts` you can specify either a JSON filename (ending in `.json`) or an inline JSON string. The JSON object should be a collection of name:value pairs. name is string for the cluster name and there **must be one** default.

Examples:

- json file: `--target-hosts="target_hosts1.json"`:

```
{ "default": ["127.0.0.1:9200", "10.127.0.3:19200"] }
```

- json file defining two clusters: `--target-hosts="target_hosts2.json"`:

```
{
  "default": [
    {"host": "127.0.0.1", "port": 9200},
    {"host": "127.0.0.1", "port": 19200}
  ],
  "remote": [
    {"host": "10.127.0.3", "port": 9200},
    {"host": "10.127.0.8", "port": 9201}
  ]
}
```

- json inline string defining two clusters:

```
--target-hosts="{\"default\": [\"127.0.0.1:9200\"], \"remote\": [\"127.0.0.1:19200\",  
↪ \"127.0.0.1:19201\"]}"
```

Note: All *built-in operations* will use the connection to the default cluster. However, you can utilize the client connections to the additional clusters in your *custom runners*.

client-options

client-options can optionally specify options for the Elasticsearch clients when multiple clusters have been defined with target-hosts. If omitted, the default is timeout : 60 for all cluster connections.

The format is similar to target-hosts, supporting both filenames ending in .json or inline JSON, however, the parameters are a collection of name:value pairs, as opposed to arrays.

Examples, assuming that two clusters have been specified with --target-hosts:

- json file: --client-options="client_options1.json":

```
{
  "default": {
    "timeout": 60
  },
  "remote": {
    "use_ssl": true,
    "verify_certs": false,
    "ca_certs": "/path/to/cacert.pem"
  }
}
```

- json inline string defining two clusters:

```
--client-options="{\"default\":{\"timeout\": 60}, \"remote\": {\"use_ssl\":true,  
↪ \"verify_certs\":false, \"ca_certs\": \"\"/path/to/cacert.pem\"}}"
```

Warning: If you use client-options you must specify options for **every** cluster name defined with target-hosts. Rally will raise an error if there is a mismatch.

2.10 Offline Usage

In some corporate environments servers do not have Internet access. You can still use Rally in such environments and this page summarizes all information that you need to get started.

2.10.1 Installation and Configuration

We provide a special offline installation package. Follow the *offline installation guide* and *configure Rally as usual* afterwards.

2.10.2 Command Line Usage

Rally will automatically detect upon startup that no Internet connection is available and print the following warning:

```
[WARNING] No Internet connection detected. Automatic download of track data sets etc.
↳ is disabled.
```

It detects this by trying to connect to `github.com`. If you want to disable this probing you can explicitly specify `--offline`.

2.10.3 Using tracks

A Rally track describes a benchmarking scenario. You can either write your own tracks or use the tracks that Rally provides out of the box. In the former case, Rally will work just fine in an offline environment. In the latter case, Rally would normally download the track and its associated data from the Internet. If you want to use one of Rally's standard tracks in offline mode, you need to download all relevant files first on a machine that has Internet access and copy it to the target machine(s).

Use the [download script](#) to download all data for a track on a machine that has access to the Internet. Example:

```
# downloads the script from Github
curl -O https://raw.githubusercontent.com/elastic/rally-tracks/master/download.sh
chmod u+x download.sh
# download all data for the geonames track
./download.sh geonames
```

This will download all data for the `geonames` track and create a tar file `rally-track-data-geonames.tar` in the current directory. Copy this file to the home directory of the user which will execute Rally on the target machine (e.g. `/home/rally-user`).

On the target machine, run:

```
cd ~
tar -xf rally-track-data-geonames.tar
```

The download script does not require a Rally installation on the machine with Internet access but assumes that `git` and `curl` are available.

After you've copied the data, you can list the available tracks with `esrally list tracks`. If a track shows up in this list, it just means that the track description is available locally but not necessarily all data files.

2.10.4 Using cars

Note: You can skip this section if you use Rally only as a load generator.

If you have Rally configure and start Elasticsearch then you also need the out-of-the-box configurations available. Run the following command on a machine with Internet access:

```
git clone https://github.com/elastic/rally-teams.git ~/.rally/benchmarks/teams/default
tar -C ~ -czf rally-teams.tar.gz ~/.rally/benchmarks/teams/default
```

Copy that file to the target machine(s) and run on the target machine:

```
cd ~
tar -xzf rally-teams.tar.gz
```

After you've copied the data, you can list the available tracks with `esrally list cars`.

2.11 Track Reference

2.11.1 Definition

A track is a specification of one or more benchmarking scenarios with a specific document corpus. It defines for example the involved indices, data files and the operations that are invoked. Its most important attributes are:

- One or more indices, each with one or more types
- The queries to issue
- Source URL of the benchmark data
- A list of steps to run, which we'll call "challenge", for example indexing data with a specific number of documents per bulk request or running searches for a defined number of iterations.

2.11.2 Track File Format and Storage

A track is specified in a JSON file.

Ad-hoc use

For ad-hoc use you can store a track definition anywhere on the file system and reference it with `--track-path`, e.g:

```
# provide a directory - Rally searches for a track.json file in this directory
# Track name is "app-logs"
esrally --track-path=~/.Projects/tracks/app-logs
# provide a file name - Rally uses this file directly
# Track name is "syslog"
esrally --track-path=~/.Projects/tracks/syslog.json
```

Rally will also search for additional files like mappings or data files in the provided directory. If you use advanced features like *custom runners* or *parameter sources* we recommend that you create a separate directory per track.

Custom Track Repositories

Alternatively, you can store Rally tracks also in a dedicated git repository which we call a “track repository”. Rally provides a default track repository that is hosted on [Github](#). You can also add your own track repositories although this requires a bit of additional work. First of all, track repositories need to be managed by git. The reason is that Rally can benchmark multiple versions of Elasticsearch and we use git branches in the track repository to determine the best match for each track (based on the command line parameter `--distribution-version`). The versioning scheme is as follows:

- The *master* branch needs to work with the latest *master* branch of Elasticsearch.
- All other branches need to match the version scheme of Elasticsearch, i.e. `MAJOR.MINOR.PATCH-SUFFIX` where all parts except `MAJOR` are optional.

Rally implements a fallback logic so you don’t need to define a branch for each patch release of Elasticsearch. For example:

- The branch `6.0.0-alpha1` will be chosen for the version `6.0.0-alpha1` of Elasticsearch.
- The branch `5` will be chosen for all versions for Elasticsearch with the major version 5, e.g. `5.0.0`, `5.1.3` (provided there is no specific branch).

Rally tries to use the branch with the best match to the benchmarked version of Elasticsearch.

Rally will also search for related files like mappings or custom runners or parameter sources in the track repository. However, Rally will use a separate directory to look for data files (`~/.rally/benchmarks/data/$TRACK_NAME/`). The reason is simply that we do not want to check multi-GB data files into git.

Creating a new track repository

All track repositories are located in `~/.rally/benchmarks/tracks`. If you want to add a dedicated track repository, called `private` follow these steps:

```
cd ~/.rally/benchmarks/tracks
mkdir private
cd private
git init
# add your track now
git add .
git commit -m "Initial commit"
```

If you want to share your tracks with others you need to add a remote and push it:

```
git remote add origin git@git-repos.acme.com:acme/rally-tracks.git
git push -u origin master
```

If you have added a remote you should also add it in `~/.rally/rally.ini`, otherwise you can skip this step. Open the file in your editor of choice and add the following line in the section `tracks`:

```
private.url = <<URL_TO_YOUR_ORIGIN>>
```

If you specify `--track-repository=private`, Rally will check whether there is a directory `~/.rally/benchmarks/tracks/private`. If there is none, it will use the provided URL to clone the repo. However, if the directory already exists, the property gets ignored and Rally will just update the local tracking branches before the benchmark starts.

You can now verify that everything works by listing all tracks in this track repository:

```
esrally list tracks --track-repository=private
```

This shows all tracks that are available on the `master` branch of this repository. Suppose you only created tracks on the branch `2` because you're interested in the performance of Elasticsearch 2.x, then you can specify also the distribution version:

```
esrally list tracks --track-repository=private --distribution-version=2.0.0
```

Rally will follow the same branch fallback logic as described above.

Adding an already existing track repository

If you want to add a track repository that already exists, just open `~/.rally/rally.ini` in your editor of choice and add the following line in the section `tracks`:

```
your_repo_name.url = <<URL_TO_YOUR_ORIGIN>>
```

After you have added this line, have Rally list the tracks in this repository:

```
esrally list tracks --track-repository=your_repo_name
```

When to use what?

We recommend the following path:

- Start with a simple json file. The file name can be arbitrary.
- If you need *custom runners* or *parameter sources*, create one directory per track. Then you can keep everything that is related to one track in one place. Remember that the track JSON file needs to be named `track.json`.
- If you want to version your tracks so they can work with multiple versions of Elasticsearch (e.g. you are running benchmarks before an upgrade), use a track repository.

2.11.3 Anatomy of a track

A track JSON file consists of the following sections:

- `indices`
- `templates`
- `corpora`
- `operations`
- `schedule`

In the `indices` and `templates` sections you define the relevant indices and index templates. These sections are optional but recommended if you want to create indices and index templates with the help of Rally.

In the `corpora` section you define all document corpora (i.e. data files) that Rally should use for this track.

In the `operations` section you describe which operations are available for this track and how they are parametrized. This section is optional and you can also define any operations directly per challenge. You can use it, if you want to share operation definitions between challenges.

In the `schedule` section you describe the workload for the benchmark, for example index with two clients at maximum throughput while searching with another two clients with ten operations per second. The schedule either uses the operations defined in the `operations` block or defines the operations to execute inline.

2.11.4 Track elements

The track elements that are described here are defined in [Rally's JSON schema for tracks](#). Rally uses this track schema to validate your tracks when it is loading them.

Each track defines the following info attributes:

- `version` (optional): An integer describing the track specification version in use. Rally uses it to detect incompatible future track specification versions and raise an error. See the table below for a reference of valid versions.
- `description` (optional): A human-readable description of the track. Although it is optional, we recommend providing it.

Track Specification Version	Rally version
1	<code>>=0.7.3, <0.10.0</code>
2	<code>>=0.9.0</code>

The `version` property has been introduced with Rally 0.7.3. Rally versions before 0.7.3 do not recognize this property and thus cannot detect incompatible track specification versions.

Example:

```
{
  "version": 2,
  "description": "POIs from Geonames"
}
```

meta

For each track, an optional structure, called `meta` can be defined. You are free which properties this element should contain.

This element can also be defined on the following elements:

- `challenge`
- `operation`
- `task`

If the `meta` structure contains the same key on different elements, more specific ones will override the same key of more generic elements. The order from generic to most specific is:

1. `track`
2. `challenge`
3. `operation`
4. `task`

E.g. a key defined on a task, will override the same key defined on a challenge. All properties defined within the merged `meta` structure, will get copied into each metrics record.

indices

The `indices` section contains a list of all indices that are used by this track.

Each index in this list consists of the following properties:

- `name` (mandatory): The name of the index.
- `body` (optional): File name of the corresponding index definition that will be used as body in the create index API call.
- `types` (optional): A list of type names in this index.

Example:

```
"indices": [  
  {  
    "name": "geonames",  
    "body": "geonames-index.json",  
    "types": ["docs"]  
  }  
]
```

templates

The `templates` section contains a list of all index templates that Rally should create.

- `name` (mandatory): Index template name
- `index-pattern` (mandatory): Index pattern that matches the index template. This must match the definition in the index template file.
- `delete-matching-indices` (optional, defaults to `true`): Delete all indices that match the provided index pattern before start of the benchmark.
- `template` (mandatory): Index template file name

Example:

```
"templates": [  
  {  
    "name": "my-default-index-template",  
    "index-pattern": "my-index-*",  
    "delete-matching-indices": true,  
    "template": "default-template.json"  
  }  
]
```

corpora

The `corpora` section contains all document corpora that are used by this track. Note that you can reuse document corpora across tracks; just copy & paste the respective corpora definitions. It consists of the following properties:

- `name` (mandatory): Name of this document corpus. As this name is also used by Rally in directory names, it is recommended to only use lower-case names without whitespaces for maximum compatibility across file systems.
- `documents` (mandatory): A list of documents files.

Each entry in the `documents` list consists of the following properties:

- `base-url` (optional): A `http` or `https` URL that points to the root path where Rally can obtain the corresponding source file.
- `source-format` (optional, default: `bulk`): Defines in which format Rally should interpret the data file specified by `source-file`. Currently, only `bulk` is supported.
- `source-file` (mandatory): File name of the corresponding documents. For local use, this file can be a `.json` file. If you provide a `base-url` we recommend that you provide a compressed file here. The following extensions are supported: `.zip`, `.bz2`, `.gz`, `.tar`, `.tar.gz`, `.tgz` or `.tar.bz2`. It must contain exactly one JSON file with the same name. The preferred file extension for our official tracks is `.bz2`.
- `includes-action-and-meta-data` (optional, defaults to `false`): Defines whether the documents file contains already an action and meta-data line (`true`) or only documents (`false`).
- `document-count` (mandatory): Number of documents in the source file. This number is used by Rally to determine which client indexes which part of the document corpus (each of the `N` clients gets one `N`-th of the document corpus). If you are using parent-child, specify the number of parent documents.
- `compressed-bytes` (optional but recommended): The size in bytes of the compressed source file. This number is used to show users how much data will be downloaded by Rally and also to check whether the download is complete.
- `uncompressed-bytes` (optional but recommended): The size in bytes of the source file after decompression. This number is used by Rally to show users how much disk space the decompressed file will need and to check that the whole file could be decompressed successfully.
- `target-index`: Defines the name of the index which should be targeted for bulk operations. Rally will automatically derive this value if you have defined exactly one index in the `indices` section. Ignored if `includes-action-and-meta-data` is `true`.
- `target-type`: Defines the name of the document type which should be targeted for bulk operations. Rally will automatically derive this value if you have defined exactly one index in the `indices` section and this index has exactly one type. Ignored if `includes-action-and-meta-data` is `true`.

To avoid repetition, you can specify default values on document corpus level for the following properties:

- `base-url`
- `source-format`
- `includes-action-and-meta-data`
- `target-index`
- `target-type`

Examples

Here we define a single document corpus with one set of documents:

```
"corpora": [
  {
    "name": "geonames",
    "documents": [
      {
        "base-url": "http://benchmarks.elasticsearch.org.s3.amazonaws.com/corpora/
↪geonames",
        "source-file": "documents.json.bz2",
        "document-count": 11396505,
        "compressed-bytes": 264698741,
        "uncompressed-bytes": 3547614383,
        "target-index": "geonames",
```

(continues on next page)

(continued from previous page)

```

    "target-type": "docs"
  }
]
}
]

```

We can also define default values on document corpus level but override some of them (`base-url` for the last entry):

```

"corpora": [
  {
    "name": "http_logs",
    "base-url": "http://benchmarks.elasticsearch.org.s3.amazonaws.com/corpora/http_
↪logs",
    "target-type": "docs",
    "documents": [
      {
        "source-file": "documents-181998.json.bz2",
        "document-count": 2708746,
        "target-index": "logs-181998"
      },
      {
        "source-file": "documents-191998.json.bz2",
        "document-count": 9697882,
        "target-index": "logs-191998"
      },
      {
        "base-url": "http://example.org/corpora/http_logs",
        "source-file": "documents-201998.json.bz2",
        "document-count": 13053463,
        "target-index": "logs-201998"
      }
    ]
  }
]

```

operations

The `operations` section contains a list of all operations that are available later when specifying a schedule. Operations define the static properties of a request against Elasticsearch whereas the `schedule` element defines the dynamic properties (such as the target throughput).

Each operation consists of the following properties:

- `name` (mandatory): The name of this operation. You can choose this name freely. It is only needed to reference the operation when defining schedules.
- `operation-type` (mandatory): Type of this operation. See below for the operation types that are supported out of the box in Rally. You can also add arbitrary operations by defining *custom runners*.
- `include-in-reporting` (optional, defaults to `true` for normal operations and to `false` for administrative operations): Whether or not this operation should be included in the command line report. For example you might want Rally to create an index for you but you are not interested in detailed metrics about it. Note that Rally will still record all metrics in the metrics store.


Depending on the operation type a couple of further parameters can be specified.

bulk

With the operation type `bulk` you can execute [bulk requests](#). It supports the following properties:

- `bulk-size` (mandatory): Defines the bulk size in number of documents.
- `ingest-percentage` (optional, defaults to 100): A number between (0, 100] that defines how much of the document corpus will be bulk-indexed.
- `corpora` (optional): A list of document corpus names that should be targeted by this bulk-index operation. Only needed if the `corpora` section contains more than one document corpus and you don't want to index all of them with this operation.
- `indices` (optional): A list of index names that defines which indices should be used by this bulk-index operation. Rally will then only select the documents files that have a matching `target-index` specified.
- `batch-size` (optional): Defines how many documents Rally will read at once. This is an expert setting and only meant to avoid accidental bottlenecks for very small bulk sizes (e.g. if you want to benchmark with a `bulk-size` of 1, you should set `batch-size` higher).
- `pipeline` (optional): Defines the name of an (existing) ingest pipeline that should be used (only supported from Elasticsearch 5.0).
- `conflicts` (optional): Type of index conflicts to simulate. If not specified, no conflicts will be simulated (also read below on how to use external index ids with no conflicts). Valid values are: 'sequential' (A document id is replaced with a document id with a sequentially increasing id), 'random' (A document id is replaced with a document id with a random other id).
- `conflict-probability` (optional, defaults to 25 percent): A number between [0, 100] that defines how many of the documents will get replaced. Combining `conflicts=sequential` and `conflict-probability=0` makes Rally generate index ids by itself, instead of relying on Elasticsearch's [automatic id generation](#).
- `on-conflict` (optional, defaults to `index`): Determines whether Rally should use the action `index` or `update` on id conflicts.
- `recency` (optional, defaults to 0): A number between [0,1] indicating whether to bias conflicting ids towards more recent ids (`recency` towards 1) or whether to consider all ids for id conflicts (`recency` towards 0). See the diagram below for details.
- `detailed-results` (optional, defaults to `false`): Records more detailed meta-data for bulk requests. As it analyzes the corresponding bulk response in more detail, this might incur additional overhead which can skew measurement results.

The image below shows how Rally behaves with a `recency` set to 0.5. Internally, Rally uses the blue function for its calculations but to understand the behavior we will focus on red function (which is just the inverse). Suppose we have already generated ids from 1 to 100 and we are about to simulate an id conflict. Rally will randomly choose a value on the y-axis, e.g. 0.8 which is mapped to 0.1 on the x-axis. This means that in 80% of all cases, Rally will choose an id within the most recent 10%, i.e. between 90 and 100. With 20% probability the id will be between 1 and 89. The closer `recency` gets to zero, the "flatter" the red curve gets and the more likely Rally will choose less recent ids.



You can also explore the recency calculation interactively.

Example:

```
{
  "name": "index-append",
  "operation-type": "bulk",
  "bulk-size": 5000
}
```

Throughput will be reported as number of indexed documents per second.

force-merge

With the operation type `force-merge` you can call the [force merge API](#). On older versions of Elasticsearch (prior to 2.1), Rally will use the `optimize` API instead. It supports the following parameter:

- `max-num-segments` (optional) The number of segments the index should be merged into. Defaults to simply checking if a merge needs to execute, and if so, executes it.

This is an administrative operation. Metrics are not reported by default. If reporting is forced by setting `include-in-reporting` to `true`, then throughput is reported as the number of completed force-merge operations per second.

index-stats

With the operation type `index-stats` you can call the [indices stats API](#). It does not support any parameters.

Throughput will be reported as number of completed *index-stats* operations per second.

node-stats

With the operation type `nodes-stats` you can execute `nodes stats API`. It does not support any parameters.

Throughput will be reported as number of completed *node-stats* operations per second.

search

With the operation type `search` you can execute `request body searches`. It supports the following properties:

- `index` (optional): An `index pattern` that defines which indices should be targeted by this query. Only needed if the `index` section contains more than one index. Otherwise, Rally will automatically derive the index to use. If you have defined multiple indices and want to query all of them, just specify `"index": "_all"`.
- `type` (optional): Defines the type within the specified index for this query. By default, no `type` will be used and the query will be performed across all types in the provided index.
- `cache` (optional): Whether to use the query request cache. By default, Rally will define no value thus the default depends on the benchmark candidate settings and Elasticsearch version.
- `request-params` (optional): A structure containing arbitrary request parameters. The supported parameters names are documented in the [Python ES client API docs](#). Parameters that are implicitly set by Rally (e.g. `body` or `request_cache`) are not supported (i.e. you should not try to set them and if so expect unspecified behavior).
- `body` (mandatory): The query body.
- `pages` (optional): Number of pages to retrieve. If this parameter is present, a scroll query will be executed. If you want to retrieve all result pages, use the value `"all"`.
- `results-per-page` (optional): Number of documents to retrieve per page for scroll queries.

Example:

```
{
  "name": "default",
  "operation-type": "search",
  "body": {
    "query": {
      "match_all": {}
    }
  },
  "request-params": {
    "_source_include": "some_field",
    "analyze_wildcard": false
  }
}
```

For scroll queries, throughput will be reported as number of retrieved scroll pages per second. The unit is `ops/s`, where one `op(eration)` is one page that has been retrieved. The rationale is that each HTTP request corresponds to one operation and we need to issue one HTTP request per result page. Note that if you use a dedicated Elasticsearch metrics store, you can also use other request-level meta-data such as the number of hits for your own analyses.

For other queries, throughput will be reported as number of search requests per second, also measured as `ops/s`.

put-pipeline

With the operation-type `put-pipeline` you can execute the `put pipeline API`. Note that this API is only available from Elasticsearch 5.0 onwards. It supports the following properties:

- *id* (mandatory): Pipeline id
- *body* (mandatory): Pipeline definition

Example:

```
{
  "name": "define-ip-geocoder",
  "operation-type": "put-pipeline",
  "id": "ip-geocoder",
  "body": {
    "description": "Extracts location information from the client IP.",
    "processors": [
      {
        "geoip": {
          "field": "clientip",
          "properties": [
            "city_name",
            "country_iso_code",
            "country_name",
            "location"
          ]
        }
      }
    ]
  }
}
```

This example requires that the `ingest-geoip` Elasticsearch plugin is installed.

This is an administrative operation. Metrics are not reported by default. Reporting can be forced by setting `include-in-reporting` to `true`.

cluster-health

With the operation `cluster-health` you can execute the `cluster health API`. It supports the following properties:

- `request-params` (optional): A structure containing any request parameters that are allowed by the cluster health API.
- `index` (optional): The name of the index that should be used to check.

The `cluster-health` operation will check whether the expected cluster health and will report a failure if this is not the case. Use `--on-error` on the command line to control Rally's behavior in case of such failures.

Example:

```
{
  "name": "check-cluster-green",
  "operation-type": "cluster-health",
  "index": "logs-*",
  "request-params": {
    "wait_for_status": "green",
  }
}
```

(continues on next page)

(continued from previous page)

```

    "wait_for_no_relocating_shards": "true"
  }
}

```

This is an administrative operation. Metrics are not reported by default. Reporting can be forced by setting `include-in-reporting` to `true`.

refresh

With the operation `refresh` you can execute the [refresh API](#). It supports the following properties:

- `index` (optional, defaults to `_all`): The name of the index that should be refreshed.

This is an administrative operation. Metrics are not reported by default. Reporting can be forced by setting `include-in-reporting` to `true`.

create-index

With the operation `create-index` you can execute the [create index API](#). It supports two modes: it creates either all indices that are specified in the track's `indices` section or it creates one specific index defined by this operation.

If you want it to create all indices that have been declared in the `indices` section you can specify the following properties:

- `settings` (optional): Allows to specify additional index settings that will be merged with the index settings specified in the body of the index in the `indices` section.
- `request-params` (optional): A structure containing any [request parameters](#) that are allowed by the create index API.

If you want it to create one specific index instead, you can specify the following properties:

- `index` (mandatory): One or more names of the indices that should be created. If only one index should be created, you can use a string otherwise this needs to be a list of strings.
- `body` (optional): The body for the create index API call.
- `request-params` (optional): A structure containing any [request parameters](#) that are allowed by the create index API.

Examples

The following snippet will create all indices that have been defined in the `indices` section. It will reuse all settings defined but override the number of shards:

```

{
  "name": "create-all-indices",
  "operation-type": "create-index",
  "settings": {
    "index.number_of_shards": 1
  },
  "request-params": {
    "wait_for_active_shards": true
  }
}

```

With the following snippet we will create a new index that is not defined in the `indices` section. Note that we specify the index settings directly in the body:

```
{
  "name": "create-an-index",
  "operation-type": "create-index",
  "index": "people",
  "body": {
    "settings": {
      "index.number_of_shards": 0
    },
    "mappings": {
      "docs": {
        "properties": {
          "name": {
            "type": "text"
          }
        }
      }
    }
  }
}
```

This is an administrative operation. Metrics are not reported by default. Reporting can be forced by setting `include-in-reporting` to `true`.

delete-index

With the operation `delete-index` you can execute the [delete index API](#). It supports two modes: it deletes either all indices that are specified in the track's `indices` section or it deletes one specific index (pattern) defined by this operation.

If you want it to delete all indices that have been declared in the `indices` section, you can specify the following properties:

- `only-if-exists` (optional, defaults to `true`): Defines whether an index should only be deleted if it exists.
- `request-params` (optional): A structure containing any [request parameters](#) that are allowed by the delete index API.

If you want it to delete one specific index (pattern) instead, you can specify the following properties:

- `index` (mandatory): One or more names of the indices that should be deleted. If only one index should be deleted, you can use a string otherwise this needs to be a list of strings.
- `only-if-exists` (optional, defaults to `true`): Defines whether an index should only be deleted if it exists.
- `request-params` (optional): A structure containing any [request parameters](#) that are allowed by the delete index API.

Examples

With the following snippet we will delete all indices that are declared in the `indices` section but only if they existed previously (implicit default):

```
{
  "name": "delete-all-indices",
  "operation-type": "delete-index"
}
```

With the following snippet we will delete all `logs-*` indices:

```
{
  "name": "delete-logs",
  "operation-type": "delete-index",
  "index": "logs-*",
  "only-if-exists": false,
  "request-params": {
    "expand_wildcards": "all",
    "allow_no_indices": true,
    "ignore_unavailable": true
  }
}
```

This is an administrative operation. Metrics are not reported by default. Reporting can be forced by setting `include-in-reporting` to `true`.

create-index-template

With the operation `create-index-template` you can execute the [create index template API](#). It supports two modes: it creates either all index templates that are specified in the track's `templates` section or it creates one specific index template defined by this operation.

If you want it to create index templates that have been declared in the `templates` section you can specify the following properties:

- `template` (optional): If you specify a template name, only the template with this name will be created.
- `settings` (optional): Allows to specify additional settings that will be merged with the settings specified in the body of the index template in the `templates` section.
- `request-params` (optional): A structure containing any [request parameters](#) that are allowed by the create index template API.

If you want it to create one specific index instead, you can specify the following properties:

- `template` (mandatory): The name of the index template that should be created.
- `body` (mandatory): The body for the create index template API call.
- `request-params` (optional): A structure containing any [request parameters](#) that are allowed by the create index template API.

Examples

The following snippet will create all index templates that have been defined in the `templates` section:

```
{
  "name": "create-all-templates",
  "operation-type": "create-index-template",
  "request-params": {
    "create": true
  }
}
```

With the following snippet we will create a new index template that is not defined in the `templates` section. Note that we specify the index template settings directly in the body:

```
{
  "name": "create-a-template",
  "operation-type": "create-index-template",
```

(continues on next page)

(continued from previous page)

```

"template": "defaults",
"body": {
  "index_patterns": ["*"],
  "settings": {
    "number_of_shards": 3
  },
  "mappings": {
    "docs": {
      "_source": {
        "enabled": false
      }
    }
  }
}
}

```

This is an administrative operation. Metrics are not reported by default. Reporting can be forced by setting `include-in-reporting` to `true`.

delete-index-template

With the operation `delete-index-template` you can execute the [delete index template API](#). It supports two modes: it deletes either all index templates that are specified in the track's `templates` section or it deletes one specific index template defined by this operation.

If you want it to delete all index templates that have been declared in the `templates` section, you can specify the following properties:

- `only-if-exists` (optional, defaults to `true`): Defines whether an index template should only be deleted if it exists.
- `request-params` (optional): A structure containing any [request parameters](#) that are allowed by the delete index template API.

If you want it to delete one specific index template instead, you can specify the following properties:

- `template` (mandatory): The name of the index that should be deleted.
- `only-if-exists` (optional, defaults to `true`): Defines whether the index template should only be deleted if it exists.
- `delete-matching-indices` (optional, defaults to `false`): Whether to delete indices that match the index template's index pattern.
- `index-pattern` (mandatory iff `delete-matching-indices` is `true`): Specifies the index pattern to delete.
- `request-params` (optional): A structure containing any [request parameters](#) that are allowed by the delete index template API.

Examples

With the following snippet we will delete all index templates that are declared in the `templates` section but only if they existed previously (implicit default):

```

{
  "name": "delete-all-index-templates",
  "operation-type": "delete-index-template"
}

```

With the following snippet we will delete the *default* index template:

```
{
  "name": "delete-default-template",
  "operation-type": "delete-index-template",
  "template": "default",
  "only-if-exists": false,
  "delete-matching-indices": true,
  "index-pattern": "*"
}
```

Note: If `delete-matching-indices` is set to `true`, indices with the provided `index-pattern` are deleted regardless whether the index template has previously existed.

This is an administrative operation. Metrics are not reported by default. Reporting can be forced by setting `include-in-reporting` to `true`.

raw-request

With the operation `raw-request` you can execute arbitrary HTTP requests against Elasticsearch. This is a low-level operation that should only be used if no high-level operation is available. Note that it is always possible to write a *custom runner*. The `raw-request` operation supports the following parameters:

- `method` (optional, defaults to `GET`): The HTTP request method to use
- `path` (mandatory): Path for the API call (excluding host and port)
- `header` (optional): A structure containing any request headers as key-value pairs.
- `body` (optional): The document body.
- `request-params` (optional): A structure containing HTTP request parameters.
- `ignore` (optional): An array of HTTP response status codes to ignore (i.e. consider as successful).

schedule

The `schedule` element contains a list of tasks that are executed by Rally, i.e. it describes the workload. Each task consists of the following properties:

- `name` (optional): This property defines an explicit name for the given task. By default the operation's name is implicitly used as the task name but if the same operation is run multiple times, a unique task name must be specified using this property.
- `operation` (mandatory): This property refers either to the name of an operation that has been defined in the `operations` section or directly defines an operation inline.
- `clients` (optional, defaults to 1): The number of clients that should execute a task concurrently.
- `warmup-iterations` (optional, defaults to 0): Number of iterations that each client should execute to warmup the benchmark candidate. Warmup iterations will not show up in the measurement results.
- `iterations` (optional, defaults to 1): Number of measurement iterations that each client executes. The command line report will automatically adjust the percentile numbers based on this number (i.e. if you just run 5 iterations you will not get a 99.9th percentile because we need at least 1000 iterations to determine this value precisely).

- `warmup-time-period` (optional, defaults to 0): A time period in seconds that Rally considers for warmup of the benchmark candidate. All response data captured during warmup will not show up in the measurement results.
- `time-period` (optional): A time period in seconds that Rally considers for measurement. Note that for bulk indexing you should usually not define this time period. Rally will just bulk index all documents and consider every sample after the warmup time period as measurement sample.
- `schedule` (optional, defaults to `deterministic`): Defines the schedule for this task, i.e. it defines at which point in time during the benchmark an operation should be executed. For example, if you specify a `deterministic` schedule and a `target-interval` of 5 (seconds), Rally will attempt to execute the corresponding operation at second 0, 5, 10, 15 Out of the box, Rally supports `deterministic` and `poisson` but you can define your own *custom schedules*.
- `target-throughput` (optional): Defines the benchmark mode. If it is not defined, Rally assumes this is a throughput benchmark and will run the task as fast as it can. This is mostly needed for batch-style operations where it is more important to achieve the best throughput instead of an acceptable latency. If it is defined, it specifies the number of requests per second over all clients. E.g. if you specify `target-throughput: 1000` with 8 clients, it means that each client will issue 125 (= 1000 / 8) requests per second. In total, all clients will issue 1000 requests each second. If Rally reports less than the specified throughput then Elasticsearch simply cannot reach it.
- `target-interval` (optional): This is just $1 / \text{target-throughput}$ (in seconds) and may be more convenient for cases where the throughput is less than one operation per second. Define either `target-throughput` or `target-interval` but not both (otherwise Rally will raise an error).

Defining operations

In the following snippet we define two operations `force-merge` and a `match-all` query separately in an operations block:

```
{
  "operations": [
    {
      "name": "force-merge",
      "operation-type": "force-merge"
    },
    {
      "name": "match-all-query",
      "operation-type": "search",
      "body": {
        "query": {
          "match_all": {}
        }
      }
    }
  ],
  "schedule": [
    {
      "operation": "force-merge",
      "clients": 1
    },
    {
      "operation": "match-all-query",
      "clients": 4,
      "warmup-iterations": 1000,
      "iterations": 1000,

```

(continues on next page)

(continued from previous page)

```

    "target-throughput": 100
  }
]
}

```

If we do not want to reuse these operations, we can also define them inline. Note that the `operations` section is gone:

```

{
  "schedule": [
    {
      "operation": {
        "name": "force-merge",
        "operation-type": "force-merge"
      },
      "clients": 1
    },
    {
      "operation": {
        "name": "match-all-query",
        "operation-type": "search",
        "body": {
          "query": {
            "match_all": {}
          }
        }
      },
      "clients": 4,
      "warmup-iterations": 1000,
      "iterations": 1000,
      "target-throughput": 100
    }
  ]
}

```

Contrary to the query, the `force-merge` operation does not take any parameters, so Rally allows us to just specify the `operation-type` for this operation. It's name will be the same as the operation's type:

```

{
  "schedule": [
    {
      "operation": "force-merge",
      "clients": 1
    },
    {
      "operation": {
        "name": "match-all-query",
        "operation-type": "search",
        "body": {
          "query": {
            "match_all": {}
          }
        }
      },
      "clients": 4,
      "warmup-iterations": 1000,

```

(continues on next page)

(continued from previous page)


```
"iterations": 1000,  
  "target-throughput": 100  
}  
]  
}
```

Choosing a schedule

Rally allows you to choose between the following schedules to simulate traffic:

- [deterministically distributed](#)
- [Poisson distributed](#)

The diagram below shows how different schedules in Rally behave during the first ten seconds of a benchmark. Each schedule is configured for a (mean) target throughput of one operation per second.



If you want as much reproducibility as possible you can choose the *deterministic* schedule. A Poisson distribution models random independent arrivals of clients which on average match the expected arrival rate which makes it suitable for modelling the behaviour of multiple clients that decide independently when to issue a request. For this reason, Poisson processes play an important role in [queueing theory](#).

If you have more complex needs on how to model traffic, you can also implement a *custom schedule*.

Time-based vs. iteration-based

You should usually use time periods for batch style operations and iterations for the rest. However, you can also choose to run a query for a certain time period.

All tasks in the `schedule` list are executed sequentially in the order in which they have been defined. However, it is also possible to execute multiple tasks concurrently, by wrapping them in a `parallel` element. The `parallel` element defines of the following properties:

- `clients` (optional): The number of clients that should execute the provided tasks. If you specify this property, Rally will only use as many clients as you have defined on the `parallel` element (see examples)!
- `warmup-time-period` (optional, defaults to 0): Allows to define a default value for all tasks of the `parallel` element.
- `time-period` (optional, no default value if not specified): Allows to define a default value for all tasks of the `parallel` element.
- `warmup-iterations` (optional, defaults to 0): Allows to define a default value for all tasks of the `parallel` element.
- `iterations` (optional, defaults to 1): Allows to define a default value for all tasks of the `parallel` element.
- `completed-by` (optional): Allows to define the name of one task in the `tasks` list. As soon as this task has completed, the whole `parallel` task structure is considered completed. If this property is not explicitly defined, the `parallel` task structure is considered completed as soon as all its subtasks have completed. A task is completed if and only if all associated clients have completed execution.
- `tasks` (mandatory): Defines a list of tasks that should be executed concurrently. Each task in the list can define the following properties that have been defined above: `clients`, `warmup-time-period`, `time-period`, `warmup-iterations` and `iterations`.

Note: `parallel` elements cannot be nested.

Warning: Specify the number of clients on each task separately. If you specify this number on the `parallel` element instead, Rally will only use that many clients in total and you will only want to use this behavior in very rare cases (see examples)!

challenge

If your track defines only one benchmarking scenario specify the `schedule` on top-level. Use the `challenge` element if you want to specify additional properties like a name or a description. You can think of a challenge as a benchmarking scenario. If you have multiple challenges, you can define an array of `challenges`.

This section contains one or more challenges which describe the benchmark scenarios for this data set. A challenge can reference all operations that are defined in the `operations` section.

Each challenge consists of the following properties:

- `name` (mandatory): A descriptive name of the challenge. Should not contain spaces in order to simplify handling on the command line for users.
- `description` (optional): A human readable description of the challenge.
- `default` (optional): If true, Rally selects this challenge by default if the user did not specify a challenge on the command line. If your track only defines one challenge, it is implicitly selected as default, otherwise you need to define `"default": true` on exactly one challenge.

- `schedule` (mandatory): Defines the workload. It is described in more detail above.

Note: You should strive to minimize the number of challenges. If you just want to run a subset of the tasks in a challenge, use *task filtering*.

2.11.5 Examples

A track with a single task

To get started with custom tracks, you can benchmark a single task, e.g. a `match_all` query:

```
{
  "schedule": [
    {
      "operation": {
        "operation-type": "search",
        "index": "_all",
        "body": {
          "query": {
            "match_all": {}
          }
        }
      },
      "warmup-iterations": 100,
      "iterations": 100,
      "target-throughput": 10
    }
  ]
}
```

This track assumes that you have an existing cluster with pre-populated data. It will run the provided `match_all` query at 10 operations per second with one client and use 100 iterations as warmup and the next 100 iterations to measure.

For the examples below, note that we do not show the operation definition but you should be able to infer from the operation name what it is doing.

Running unthrottled

In this example Rally will run a bulk index operation unthrottled for one hour:

```
"schedule": [
  {
    "operation": "bulk",
    "warmup-time-period": 120,
    "time-period": 3600,
    "clients": 8
  }
]
```

Running tasks in parallel

Note: You cannot nest parallel tasks.

If we want to run tasks in parallel, we can use the `parallel` element. In the simplest case, you let Rally decide the number of clients needed to run the parallel tasks (note how we can define default values on the `parallel` element):

```
{
  "parallel": {
    "warmup-iterations": 50,
    "iterations": 100,
    "tasks": [
      {
        "operation": "default",
        "target-throughput": 50
      },
      {
        "operation": "term",
        "target-throughput": 200
      },
      {
        "operation": "phrase",
        "target-throughput": 200
      }
    ]
  }
}
```

Rally will determine that three clients are needed to run each task in a dedicated client. You can also see that each task can have different settings.

However, you can also explicitly define the number of clients:

```
"schedule": [
  {
    "parallel": {
      "warmup-iterations": 50,
      "iterations": 100,
      "tasks": [
        {
          "operation": "match-all",
          "clients": 4,
          "target-throughput": 50
        },
        {
          "operation": "term",
          "clients": 2,
          "target-throughput": 200
        },
        {
          "operation": "phrase",
          "clients": 2,
          "target-throughput": 200
        }
      ]
    }
  }
]
```

(continues on next page)

(continued from previous page)

```
}  
]
```

This schedule will run a match all query, a term query and a phrase query concurrently. It will run with eight clients in total (four for the match all query and two each for the term and phrase query).

In this scenario, we run indexing and a few queries in parallel with a total of 14 clients:

```
"schedule": [  
  {  
    "parallel": {  
      "tasks": [  
        {  
          "operation": "bulk",  
          "warmup-time-period": 120,  
          "time-period": 3600,  
          "clients": 8,  
          "target-throughput": 50  
        },  
        {  
          "operation": "default",  
          "clients": 2,  
          "warmup-iterations": 50,  
          "iterations": 100,  
          "target-throughput": 50  
        },  
        {  
          "operation": "term",  
          "clients": 2,  
          "warmup-iterations": 50,  
          "iterations": 100,  
          "target-throughput": 200  
        },  
        {  
          "operation": "phrase",  
          "clients": 2,  
          "warmup-iterations": 50,  
          "iterations": 100,  
          "target-throughput": 200  
        }  
      ]  
    }  
  }  
]
```

We can use `completed-by` to stop querying as soon as bulk-indexing has completed:

```
"schedule": [  
  {  
    "parallel": {  
      "completed-by": "bulk",  
      "tasks": [  
        {  
          "operation": "bulk",  
          "warmup-time-period": 120,  
          "time-period": 3600,  
          "clients": 8,  
          "target-throughput": 50  
        }  
      ]  
    }  
  }  
]
```

(continues on next page)

(continued from previous page)

```

    "target-throughput": 50
  },
  {
    "operation": "default",
    "clients": 2,
    "warmup-time-period": 480,
    "time-period": 7200,
    "target-throughput": 50
  }
]
}
]

```

We can also mix sequential tasks with the `parallel` element. In this scenario we are indexing with 8 clients and continue querying with 6 clients after indexing has finished:

```

"schedule": [
  {
    "operation": "bulk",
    "warmup-time-period": 120,
    "time-period": 3600,
    "clients": 8,
    "target-throughput": 50
  },
  {
    "parallel": {
      "warmup-iterations": 50,
      "iterations": 100,
      "tasks": [
        {
          "operation": "default",
          "clients": 2,
          "target-throughput": 50
        },
        {
          "operation": "term",
          "clients": 2,
          "target-throughput": 200
        },
        {
          "operation": "phrase",
          "clients": 2,
          "target-throughput": 200
        }
      ]
    }
  }
]

```

Be aware of the following case where we explicitly define that we want to run only with two clients *in total*:

```

"schedule": [
  {
    "parallel": {
      "warmup-iterations": 50,

```

(continues on next page)

(continued from previous page)

```
"iterations": 100,
"clients": 2,
"tasks": [
  {
    "operation": "match-all",
    "target-throughput": 50
  },
  {
    "operation": "term",
    "target-throughput": 200
  },
  {
    "operation": "phrase",
    "target-throughput": 200
  }
]
}
```

Rally will *not* run all three tasks in parallel because you specified that you want only two clients in total. Hence, Rally will first run “match-all” and “term” concurrently (with one client each). After they have finished, Rally will run “phrase” with one client. You could also specify more clients than there are tasks but these will then just idle.

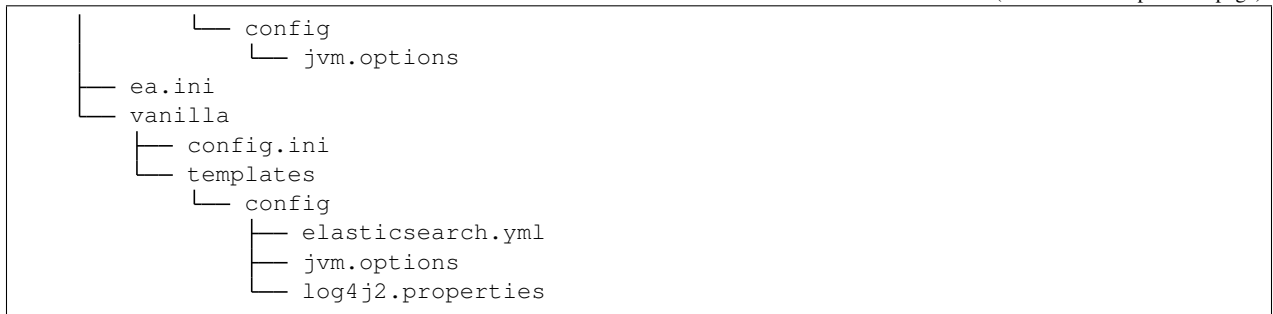
You can also specify a number of clients on sub tasks explicitly (by default, one client is assumed per subtask). This allows to define a weight for each client operation. Note that you need to define the number of clients also on the `parallel` parent element, otherwise Rally would determine the number of totally needed clients again on its own:

```
{
  "parallel": {
    "clients": 3,
    "warmup-iterations": 50,
    "iterations": 100,
    "tasks": [
      {
        "operation": "default",
        "target-throughput": 50
      },
      {
        "operation": "term",
        "target-throughput": 200
      },
      {
        "operation": "phrase",
        "target-throughput": 200,
        "clients": 2
      }
    ]
  }
}
```

This will ensure that the phrase query will be executed by two clients. All other ones are executed by one client.

2.12 Configure Elasticsearch: Cars

(continued from previous page)



The top-level directory “v1” denotes the configuration format in version 1. Below that directory, each `.ini` file defines a car. Each directory (`ea` or `vanilla`) contains templates for the config files. Rally will only copy the files in the `templates` subdirectory. The top-level directory is reserved for a special file, `config.ini` which you can use to define default variables that apply to all cars that are based on this configuration. Below is an example `config.ini` file:

```
[variables]
clean_command=./gradlew clean
```

This defines the variable `clean_command` for all cars that reference this configuration. Rally will treat the following variable names specially:

- `clean_command`: The command to clean the Elasticsearch project directory.
- `build_command`: The command to build an Elasticsearch source distribution.
- `artifact_path_pattern`: A glob pattern to find a previously built source distribution within the project directory.
- `release_url`: A download URL for Elasticsearch distributions. The placeholder `{{VERSION}}` is replaced by Rally with the actual Elasticsearch version.

Let’s have a look at the `1gheap` car by inspecting `1gheap.ini`:

```
[meta]
description=Sets the Java heap to 1GB
type=car

[config]
base=vanilla

[variables]
heap_size=1g
```

The name of the car is derived from the `.ini` file name. In the `meta` section we can provide a `description` and the `type`. Use `car` if a configuration can be used standalone and `mixin` if it needs to be combined with other configurations. In the `config` section we define that this definition is based on the `vanilla` configuration. We also define a variable `heap_size` and set it to `1g`. Note that variables defined here take precedence over variables defined in the `config.ini` file of any of the referenced configurations.

Let’s open `vanilla/config/templates/jvm.options` to see how this variable is used (we’ll only show the relevant part here):

```
# Xms represents the initial size of total heap space
# Xmx represents the maximum size of total heap space

-Xms{{heap_size}}
-Xmx{{heap_size}}
```


So Rally reads all variables and the template files and replaces the variables in the final configuration. Note that Rally does not know anything about `jvm.options` or `elasticsearch.yml`. For Rally, these are just plain text templates that need to be copied to the Elasticsearch directory before running a benchmark. Under the hood, Rally uses Jinja2 as template language. This allows you to use Jinja2 expressions in your car configuration files.

If you open `vanilla/templates/config/elasticsearch.yml` you will see a few variables that are not defined in the `.ini` file:

- `network_host`
- `http_port`
- `node_count_per_host`

These values are derived by Rally internally based on command line flags and you cannot override them in your car definition. You also cannot use these names as names for variables because Rally would simply override them.

If you specify multiple configurations, e.g. `--car="4gheap,ea"`, Rally will apply them in order. It will first read all variables in `4gheap.ini`, then in `ea.ini`. Afterwards, it will copy all configuration files from the corresponding config base of `4gheap` and *append* all configuration files from `ea`. This also shows when to define a separate “car” and when to define a “mixin”: If you need to amend configuration files, use a mixin, if you need to have a specific configuration, define a car.

Simple customizations

For simple customizations you can create the directory hierarchy as outlined above and use the `--team-path` command line parameter to refer to this configuration. For more complex use cases and distributed multi-node benchmarks, we recommend to use custom team repositories.

Custom Team Repositories

Rally provides a default team repository that is hosted on [Github](#). You can also add your own team repositories although this requires a bit of additional work. First of all, team repositories need to be managed by git. The reason is that Rally can benchmark multiple versions of Elasticsearch and we use git branches in the track repository to determine the best match. The versioning scheme is as follows:

- The *master* branch needs to work with the latest *master* branch of Elasticsearch.
- All other branches need to match the version scheme of Elasticsearch, i.e. `MAJOR.MINOR.PATCH-SUFFIX` where all parts except `MAJOR` are optional.

Rally implements a fallback logic so you don’t need to define a branch for each patch release of Elasticsearch. For example:

- The branch `6.0.0-alpha1` will be chosen for the version `6.0.0-alpha1` of Elasticsearch.
- The branch `5` will be chosen for all versions for Elasticsearch with the major version 5, e.g. `5.0.0`, `5.1.3` (provided there is no specific branch).

Rally tries to use the branch with the best match to the benchmarked version of Elasticsearch.

Creating a new team repository

All team repositories are located in `~/rally/benchmarks/teams`. If you want to add a dedicated team repository, called `private` follow these steps:

```
cd ~/.rally/benchmarks/teams
mkdir private
cd private
git init
# add your team now (don't forget to add the subdirectory "cars").
git add .
git commit -m "Initial commit"
```

If you want to share your teams with others (or you want to run remote benchmarks) you need to add a remote and push it:

```
git remote add origin git@git-repos.acme.com:acme/rally-teams.git
git push -u origin master
```

If you have added a remote you should also add it in `~/.rally/rally.ini`, otherwise you can skip this step. Open the file in your editor of choice and add the following line in the section `teams`:

```
private.url = <<URL_TO_YOUR_ORIGIN>>
```

Rally will then automatically update the local tracking branches before the benchmark starts.

Warning: If you run benchmarks against a remote machine that is under the control of Rally then you need to add the custom team configuration on every node!

You can now verify that everything works by listing all teams in this team repository:

```
esrally list cars --team-repository=private
```

This shows all teams that are available on the `master` branch of this repository. Suppose you only created tracks on the branch `2` because you're interested in the performance of Elasticsearch 2.x, then you can specify also the distribution version:

```
esrally list teams --team-repository=private --distribution-version=2.0.0
```

Rally will follow the same branch fallback logic as described above.

Adding an already existing team repository

If you want to add a team repository that already exists, just open `~/.rally/rally.ini` in your editor of choice and add the following line in the section `teams`:

```
your_repo_name.url = <<URL_TO_YOUR_ORIGIN>>
```

After you have added this line, have Rally list the tracks in this repository:

```
esrally list cars --team-repository=your_repo_name
```

2.13 Using Elasticsearch Plugins

You can have Rally setup an Elasticsearch cluster with plugins for you. However, there are a couple of restrictions:

- This feature is only supported from Elasticsearch 5.0.0 onwards

- Whereas Rally caches downloaded Elasticsearch distributions, plugins will always be installed via the Internet and thus each machine where an Elasticsearch node will be installed, requires an active Internet connection.

2.13.1 Listing plugins

To see which plugins are available, run `esrally list elasticsearch-plugins`:

```
Available Elasticsearch plugins:

Name                               Configuration
-----                               -
analysis-icu
analysis-kuromoji
analysis-phonetic
analysis-smartcn
analysis-stempel
analysis-ukrainian
discovery-azure-classic
discovery-ec2
discovery-file
discovery-gce
ingest-attachment
ingest-geoip
ingest-user-agent
lang-javascript
lang-python
mapper-attachments
mapper-murmur3
mapper-size
repository-azure
repository-gcs
repository-hdfs
repository-s3
store-smb
```

Rally supports plugins only for Elasticsearch 5.0 or better. As the availability of plugins may change from release to release we recommend that you include the `--distribution-version` parameter when listing plugins. By default Rally assumes that you want to benchmark the latest master version of Elasticsearch.

Let's see what happens if we run `esrally list elasticsearch-plugins --distribution-version=2.4.0`:

```
No Elasticsearch plugins are available.
```

As mentioned before, this is expected as only Elasticsearch 5.0 or better is supported.

2.13.2 Running a benchmark with plugins

In order to tell Rally to install a plugin, use the `--elasticsearch-plugins` parameter when starting a race. You can provide multiple plugins (comma-separated) and they will be installed in the order to that you define on the command line.

Example:

```
esrally --distribution-version=5.5.0 --elasticsearch-plugins="analysis-icu,analysis-
↳phonetic"
```

This will install the plugins `analysis-icu` and `analysis-phonetic` (in that order). In order to use the features that these plugins provide, you need to write a *custom track*.

Rally will use several techniques to install and configure plugins:

- First, Rally checks whether directory `plugins/PLUGIN_NAME` in the currently configured team repository exists. If this is the case, then plugin installation and configuration details will be read from this directory.
- Next, Rally will use the provided plugin name when running the Elasticsearch plugin installer. With this approach we can avoid to create a plugin configuration directory in the team repository for very simple plugins that do not need any configuration.

As mentioned above, Rally also allows you to specify a plugin configuration and you can even combine them. Here are some examples (requires Elasticsearch < 6.3.0 because with 6.3.0 x-pack has turned into a module of Elasticsearch which is treated as a “car” in Rally):

- Run a benchmark with the x-pack plugin in the security configuration:
`--elasticsearch-plugins=x-pack:security`
- Run a benchmark with the x-pack plugin in the security and the graph configuration:
`--elasticsearch-plugins=x-pack:security+graph`

Note: To benchmark the security configuration of x-pack you need to add the following command line options: `--client-options="use_ssl:true,verify_certs:false,basic_auth_user:'rally',basic_auth_password:'rally-password'"`

You can also override plugin variables with `--plugin-params` which is needed for example if you want to use the `monitoring-http` configuration in order to export monitoring data. You can export monitoring data e.g. with the following configuration:

```
--elasticsearch-plugins="x-pack:monitoring-http" --plugin-params="monitoring_type:  
↪'https',monitoring_host:'some_remote_host',monitoring_port:10200,monitoring_user:  
↪'rally',monitoring_password:'m0n1t0r1ng'"
```

The `monitoring_user` and `monitoring_password` parameters are optional, the other parameters are mandatory. For more details on the configuration options check the [Monitoring plugin documentation](#).

If you are behind a proxy, set the environment variable `ES_JAVA_OPTS` accordingly on each target machine as described in the [Elasticsearch plugin documentation](#).

2.13.3 Building plugins from sources

Plugin authors may want to benchmark source builds of their plugins. Your plugin is either:

- built alongside Elasticsearch
- built against a released version of Elasticsearch

Plugins built alongside Elasticsearch

To make this work, you need to manually edit Rally’s configuration file in `~/.rally/rally.ini`. Suppose, we want to benchmark the plugin “my-plugin”. Then you need to add the following entries in the `source` section:

```
plugin.my-plugin.remote.repo.url = git@github.com:example-org/my-plugin.git  
plugin.my-plugin.src.subdir = elasticsearch-extra/my-plugin
```

(continues on next page)

(continued from previous page)

```
plugin.my-plugin.build.command = ./gradlew :my-plugin:plugin:assemble
plugin.my-plugin.build.artifact.subdir = plugin/build/distributions
```

Let's discuss these properties one by one:

- `plugin.my-plugin.remote.repo.url` (optional): This is needed to let Rally checkout the source code of the plugin. If this is a private repo, credentials need to be setup properly. If the source code is already locally available you may not need to define this property. The remote's name is assumed to be "origin" and this is not configurable. Also, only git is supported as revision control system.
- `plugin.my-plugin.src.subdir` (mandatory): This is the directory to which the plugin will be checked out relative to `src.root.dir`. In order to allow to build the plugin alongside Elasticsearch, the plugin needs to reside in a subdirectory of `elasticsearch-extra` (see also the [Elasticsearch testing documentation](#)).
- `plugin.my-plugin.build.command` (mandatory): The full build command to run in order to build the plugin artifact. Note that this command is run from the Elasticsearch source directory as Rally assumes that you want to build your plugin alongside Elasticsearch (otherwise, see the next section).
- `plugin.my-plugin.build.artifact.subdir` (mandatory): This is the subdirectory relative to `plugin.my-plugin.src.subdir` in which the final plugin artifact is located.

Warning: `plugin.my-plugin.build.command` has replaced `plugin.my-plugin.build.task` in earlier Rally versions. It now requires the **full** build command.

In order to run a benchmark with `my-plugin`, you'd invoke Rally as follows:

```
esrally --revision="elasticsearch:some-elasticsearch-revision,
my-plugin:some-plugin-revision" --elasticsearch-plugins="my-plugin" where you
need to replace some-elasticsearch-revision and some-plugin-revision with the appropriate git
revisions. Adjust other command line parameters (like track or car) accordingly. In order for this to work, you need to
ensure that:
```

- All prerequisites for source builds are installed.
- The Elasticsearch source revision is compatible with the chosen plugin revision. Note that you do not need to know the revision hash to build against an already released version and can use git tags instead. E.g. if you want to benchmark against Elasticsearch 5.6.1, you can specify `--revision="elasticsearch:v5.6.1,my-plugin:some-plugin-revision"` (see e.g. the [Elasticsearch tags on Github](#) or use `git tag` in the Elasticsearch source directory on the console).
- If your plugin needs to be configured, create a proper plugin specification (see below).

Note: Rally can build all [Elasticsearch core plugins](#) out of the box without any further configuration.

Plugins based on a released Elasticsearch version

To make this work, you need to manually edit Rally's configuration file in `~/.rally/rally.ini`. Suppose, we want to benchmark the plugin "my-plugin". Then you need to add the following entries in the `source` section:

```
plugin.my-plugin.remote.repo.url = git@github.com:example-org/my-plugin.git
plugin.my-plugin.src.dir = /path/to/your/plugin/sources
plugin.my-plugin.build.command = /usr/local/bin/gradle :my-plugin:plugin:assemble
plugin.my-plugin.build.artifact.subdir = build/distributions
```

Let's discuss these properties one by one:

- `plugin.my-plugin.remote.repo.url` (optional): This is needed to let Rally checkout the source code of the plugin. If this is a private repo, credentials need to be setup properly. If the source code is already locally available you may not need to define this property. The remote's name is assumed to be "origin" and this is not configurable. Also, only git is supported as revision control system.
- `plugin.my-plugin.src.dir` (mandatory): This is the absolute directory to which the source code will be checked out.
- `plugin.my-plugin.build.command` (mandatory): The full build command to run in order to build the plugin artifact. This command is run from the plugin project's root directory.
- `plugin.my-plugin.build.artifact.subdir` (mandatory): This is the subdirectory relative to `plugin.my-plugin.src.dir` in which the final plugin artifact is located.

Warning: `plugin.my-plugin.build.command` has replaced `plugin.my-plugin.build.task` in earlier Rally versions. It now requires the **full** build command.

In order to run a benchmark with `my-plugin`, you'd invoke Rally as follows:

```
esrally --distribution-version="elasticsearch-version"
--revision="my-plugin:some-plugin-revision" --elasticsearch-plugins="my-plugin"
```

where you need to replace `elasticsearch-version` with the correct release (e.g. `6.0.0`) and `some-plugin-revision` with the appropriate *git revisions*. Adjust other command line parameters (like `track` or `car`) accordingly. In order for this to work, you need to ensure that:

- All prerequisites for source builds are installed.
- The Elasticsearch release is compatible with the chosen plugin revision.
- If your plugin needs to be configured, create a proper plugin specification (see below).

2.13.4 Anatomy of a plugin specification

Simple plugins

You can use Rally to benchmark community-contributed or even your own plugins. In the simplest case, the plugin does not need any custom configuration. Then you just need to add the download URL to your Rally configuration file. Consider we want to benchmark the plugin "my-plugin":

```
[distributions]
plugin.my-plugin.release.url=https://example.org/my-plugin/releases/{{VERSION}}/my-
plugin-{{VERSION}}.zip
```

Then you can use `--elasticsearch-plugins=my-plugin` to run a benchmark with your plugin. Rally will also replace `{{VERSION}}` with the distribution version that you have specified on the command line.

Plugins which require configuration

If the plugin needs a custom configuration we recommend to fork the [official Rally teams repository](#) and add your plugin configuration there. Suppose, you want to benchmark "my-plugin" which has the following settings that can be configured in `elasticsearch.yml`:

- `myplugin.active`: a boolean which activates the plugin
- `myplugin.mode`: Either `simple` or `advanced`

We want to support two configurations for this plugin: `simple` which will set `myplugin.mode` to `simple` and `advanced` which will set `myplugin.mode` to `advanced`.

First, we need a template configuration. We will call this a “config base” in Rally. We will just need one config base for this example and will call it “default”.

In `$TEAM_REPO_ROOT` create the directory structure for the plugin and its config base with `mkdir -p myplugin/default/templates/config` and add the following `elasticsearch.yml` in the new directory:

```
myplugin.active: true
myplugin.mode={{my_plugin_mode}}
```

That’s it. Later, Rally will just copy all files in `myplugin/default/templates` to the home directory of the Elasticsearch node that it configures. First, Rally will always apply the car’s configuration and then plugins can add their configuration on top. This also explains why we have created a `config/elasticsearch.yml`. Rally will just copy this file and replace template variables on the way.

Note: If you create a new customization for a plugin, ensure that the plugin name in the team repository matches the core plugin name. Note that hyphens need to be replaced by underscores (e.g. “x-pack” becomes “x_pack”). The reason is that Rally allows to write custom install hooks and the plugin name will become the root package name of the install hook. However, hyphens are not supported in Python which is why we use underscores instead.

The next step is now to create our two plugin configurations where we will set the variables for our config base “default”. Create a file `simple.ini` in the `myplugin` directory:

```
[config]
# reference our one and only config base here
base=default

[variables]
my_plugin_mode=simple
```

Similarly, create `advanced.ini` in the `myplugin` directory:

```
[config]
# reference our one and only config base here
base=default

[variables]
my_plugin_mode=advanced
```

Rally will now know about `myplugin` and its two configurations. Let’s check that with `esrally list elasticsearch-plugins`:

Available Elasticsearch plugins:

Name	Configuration
-----	-----
analysis-icu	
analysis-kuromoji	
analysis-phonetic	
analysis-smartcn	
analysis-stempel	
analysis-ukrainian	
discovery-azure-classic	
discovery-ec2	

(continues on next page)

(continued from previous page)

```

-----
--
jit          JIT Compiler Profiler  Enables JIT compiler logs.
gc           GC log                  Enables GC logs.
jfr          Flight Recorder        Enables Java Flight Recorder (requires an Oracle_
--JDK)
perf         perf stat              Reads CPU PMU counters (requires Linux and perf)
node-stats  Node Stats              Regularly samples node stats

Keep in mind that each telemetry device may incur a runtime overhead which can skew_
--results.

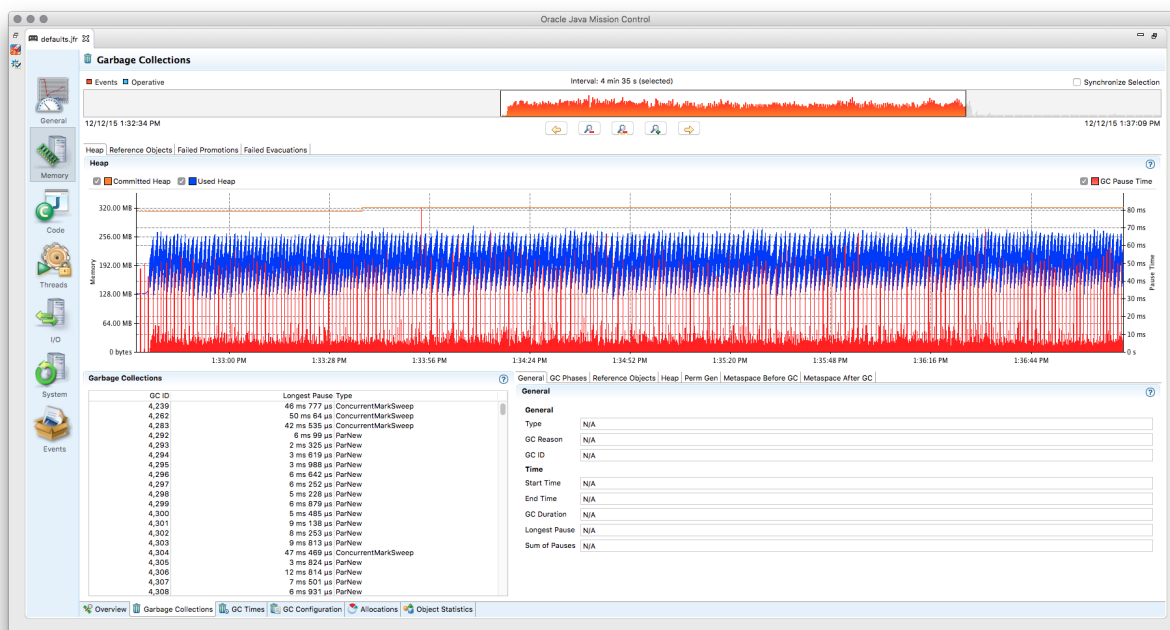
```

You can attach one or more of these telemetry devices to the benchmarked cluster. Except for `node-stats`, this only works if Rally provisions the cluster (i.e. it does not work with `--pipeline=benchmark-only`).

2.14.1 jfr

The `jfr` telemetry device enables the [Java Flight Recorder](#) on the benchmark candidate. Java Flight Recorder ships only with Oracle JDK, so Rally assumes that Oracle JDK is used for benchmarking.

To enable `jfr`, invoke Rally with `esrally --telemetry jfr`. `jfr` will then write a flight recording file which can be opened in [Java Mission Control](#). Rally prints the location of the flight recording file on the command line.



Supported telemetry parameters:

- `recording-template`: The name of a custom flight recording template. It is up to you to correctly install these recording templates on each target machine. If none is specified, the default recording template of Java flight recorder is used.

Note: The licensing terms of Java flight recorder do not allow you to run it in production environments without a valid license (for details, refer to the [Oracle Java SE Advanced & Suite Products page](#)). However, running in a QA

environment is fine.

2.14.2 jit

The `jit` telemetry device enables JIT compiler logs for the benchmark candidate. If the HotSpot disassembler library is available, the logs will also contain the disassembled JIT compiler output which can be used for low-level analysis. We recommend to use [JITWatch](#) for analysis.

`hsdis` can be built for JDK 8 on Linux with (based on a [description by Alex Blewitt](#)):

```
curl -O -O -O -O https://raw.githubusercontent.com/dmlloyd/openjdk/jdk8u/jdk8u/
↳hotspot/src/share/tools/hsdis/{hsdis.c,hsdis.h,Makefile,README}
mkdir -p build/binutils
curl http://ftp.gnu.org/gnu/binutils/binutils-2.27.tar.gz | tar --strip-components=1 -
↳C build/binutils -z -x -f -
make BINUTILS=build/binutils ARCH=amd64
```

After it has been built, the binary needs to be copied to the JDK directory (see README of `hsdis` for details).

2.14.3 gc

The `gc` telemetry device enables GC logs for the benchmark candidate. You can use tools like [GCViewer](#) to analyze the GC logs.

2.14.4 perf

The `perf` telemetry device runs `perf stat` on each benchmarked node and writes the output to a log file. It can be used to capture low-level CPU statistics. Note that the `perf` tool, which is only available on Linux, must be installed before using this telemetry device.

2.14.5 node-stats

Warning: This telemetry device will record a lot of metrics and likely skew your measurement results.

The `node-stats` telemetry devices regularly calls the `node-stats` API and records metrics from the following sections:

- Indices stats (key `indices` in the `node-stats` API)
- Thread pool stats (key `jvm.thread_pool` in the `node-stats` API)
- JVM buffer pool stats (key `jvm.buffer_pools` in the `node-stats` API)
- Circuit breaker stats (key `breakers` in the `node-stats` API)
- Network-related stats (key `transport` in the `node-stats` API)
- Process cpu stats (key `process.cpu` in the `node-stats` API)

Supported telemetry parameters:

- `node-stats-sample-interval` (default: 1): A positive number greater than zero denoting the sampling interval in seconds.

- `node-stats-include-indices` (default: `false`): A boolean indicating whether indices stats should be included.
- `node-stats-include-thread-pools` (default: `true`): A boolean indicating whether thread pool stats should be included.
- `node-stats-include-buffer-pools` (default: `true`): A boolean indicating whether buffer pool stats should be included.
- `node-stats-include-breakers` (default: `true`): A boolean indicating whether circuit breaker stats should be included.
- `node-stats-include-mem` (default: `true`): A boolean indicating whether JVM heap stats should be included.
- `node-stats-include-network` (default: `true`): A boolean indicating whether network-related stats should be included.
- `node-stats-include-process` (default: `true`): A boolean indicating whether process cpu stats should be included.

2.15 Rally Daemon

At its heart, Rally is a distributed system, just like Elasticsearch. However, in its simplest form you will not notice, because all components of Rally can run on a single node too. If you want Rally to *configure and start Elasticsearch nodes remotely* or *distribute the load test driver* to apply load from multiple machines, you need to use Rally daemon.

Rally daemon needs to run on every machine that should be under Rally's control. We can consider three different roles:

- **Benchmark coordinator:** This is the machine where you invoke `esrally`. It is responsible for user interaction, coordinates the whole benchmark and shows the results. Only one node can be the benchmark coordinator.
- **Load driver:** Nodes of this type will interpret and run *tracks*.
- **Provisioner:** Nodes of this type will configure an Elasticsearch cluster according to the provided *car* and *Elasticsearch plugin* configurations.

The two latter roles are not statically preassigned but rather determined by Rally based on the command line parameter `--load-driver-hosts` (for the load driver) and `--target-hosts` (for the provisioner).

2.15.1 Preparation

First, *install* and *configure* Rally on all machines that are involved in the benchmark. If you want to automate this, there is no need to use the interactive configuration routine of Rally. You can copy `~/.rally/rally.ini` to the target machines adapting the paths in the file as necessary. We also recommend that you copy `~/.rally/benchmarks/data` to all load driver machines before-hand. Otherwise, each load driver machine will need to download a complete copy of the benchmark data.

Note: Rally Daemon will listen on port 1900 and the actor system that Rally uses internally require access to arbitrary (unprivileged) ports. Be sure to open up these ports between the Rally nodes.

2.15.2 Starting

For all this to work, Rally needs to form a cluster. This is achieved with the binary `esrallyd` (note the “d” - for daemon - at the end). You need to start the Rally daemon on all nodes: First on the coordinator node, then on all others. The order does matter, because nodes attempt to connect to the coordinator on startup.

On the benchmark coordinator, issue:

```
esrallyd start --node-ip=IP_OF_COORDINATOR_NODE --coordinator-ip=IP_OF_COORDINATOR_
↳NODE
```

On all other nodes, issue:

```
esrallyd start --node-ip=IP_OF_THIS_NODE --coordinator-ip=IP_OF_COORDINATOR_NODE
```

After that, all Rally nodes, know about each other and you can use Rally as usual. See the *tips and tricks* for more examples.

2.15.3 Stopping

You can leave the Rally daemon processes running in case you want to run multiple benchmarks. When you are done, you can stop the Rally daemon on each node with:

```
esrallyd stop
```

Contrary to startup, order does not matter here.

2.15.4 Status

You can query the status of the local Rally daemon with:

```
esrallyd status
```

2.16 Pipelines

A pipeline is a series of steps that are performed to get benchmark results. This is *not* intended to customize the actual benchmark but rather what happens before and after a benchmark.

An example will clarify the concept: If you want to benchmark a binary distribution of Elasticsearch, Rally has to download a distribution archive, decompress it, start Elasticsearch and then run the benchmark. However, if you want to benchmark a source build of Elasticsearch, it first has to build a distribution using the Gradle Wrapper. So, in both cases, different steps are involved and that’s what pipelines are for.

You can get a list of all pipelines with `esrally list pipelines`:

```
Available pipelines:
Name                Description
-----
↳-----
from-distribution   Downloads an Elasticsearch distribution, provisions it, runs
↳a benchmark and reports results.
from-sources-complete Builds and provisions Elasticsearch, runs a benchmark and
↳reports results.
```

(continues on next page)

(continued from previous page)

```
benchmark-only          Assumes an already running Elasticsearch instance, runs a
↳benchmark and reports results
from-sources-skip-build Provisions Elasticsearch (skips the build), runs a benchmark
↳and reports results.
```

2.16.1 benchmark-only

This is intended if you want to provision a cluster by yourself. Do not use this pipeline unless you are absolutely sure you need to. As Rally has not provisioned the cluster, results are not easily reproducible and it also cannot gather a lot of metrics (like CPU usage).

To benchmark a cluster, you also have to specify the hosts to connect to. An example invocation:

```
esrally --pipeline=benchmark-only --target-hosts=search-node-a.intranet.acme.com:9200,
↳search-node-b.intranet.acme.com:9200
```

2.16.2 from-distribution

This pipeline allows to benchmark an official Elasticsearch distribution which will be automatically downloaded by Rally. The earliest supported version is Elasticsearch 1.7.0. An example invocation:

```
esrally --pipeline=from-distribution --distribution-version=1.7.5
```

The version numbers have to match the name in the download URL path.

You can also benchmark Elasticsearch snapshot versions by specifying the snapshot repository:

```
esrally --pipeline=from-distribution --distribution-version=5.0.0-SNAPSHOT --
↳distribution-repository=snapshot
```

However, this feature is mainly intended for continuous integration environments and by default you should just benchmark official distributions.

Note: This pipeline is just mentioned for completeness but Rally will autoselect it for you. All you need to do is to define the `--distribution-version` flag.

2.16.3 from-sources-complete

You should use this pipeline when you want to build and benchmark Elasticsearch from sources. This pipeline will only work from Elasticsearch 5.0 onwards because Elasticsearch switched from Maven to Gradle and Rally only supports one build tool in the interest of maintainability.

Remember that you also need git installed. If that's not the case you'll get an error and have to run `esrally configure` first. An example invocation:

```
esrally --pipeline=from-sources-complete --revision=latest
```

You have to specify a *revision*.

Note: This pipeline is just mentioned for completeness but Rally will automatically select it for you. All you need to do is to define the `--revision` flag.

2.16.4 from-sources-skip-build

This pipeline is similar to `from-sources-complete` except that it assumes you have built the binary once. It saves time if you want to run a benchmark twice for the exact same version of Elasticsearch. Obviously it doesn't make sense to provide a revision: It is always the previously built revision. An example invocation:

```
esrally --pipeline=from-sources-skip-build
```

2.17 Metrics

2.17.1 Metrics Records

At the end of a race, Rally stores all metrics records in its metrics store, which is a dedicated Elasticsearch cluster. Rally stores the metrics in the indices `rally-metrics-*`. It will create a new index for each month.

Here is a typical metrics record:

```
{
  "environment": "nightly",
  "trial-timestamp": "20160421T042749Z",
  "trial-id": "6ebc6e53-ee20-4b0c-99b4-09697987e9f4",
  "@timestamp": 1461213093093,
  "relative-time": 10507328,
  "track": "geonames",
  "track-params": {
    "shard-count": 3
  },
  "challenge": "append-no-conflicts",
  "car": "defaults",
  "sample-type": "normal",
  "name": "throughput",
  "value": 27385,
  "unit": "docs/s",
  "task": "index-append-no-conflicts",
  "operation": "index-append-no-conflicts",
  "operation-type": "Index",
  "lap": 1,
  "meta": {
    "cpu_physical_cores": 36,
    "cpu_logical_cores": 72,
    "cpu_model": "Intel(R) Xeon(R) CPU E5-2699 v3 @ 2.30GHz",
    "os_name": "Linux",
    "os_version": "3.19.0-21-generic",
    "host_name": "beast2",
    "node_name": "rally-node0",
    "source_revision": "a6c0a81",
    "distribution_version": "5.0.0-SNAPSHOT",
    "tag_reference": "Github ticket 1234"
  }
}
```

(continues on next page)

(continued from previous page)

```
}
}
```

As you can see, we do not only store the metrics name and its value but lots of meta-information. This allows you to create different visualizations and reports in Kibana.

Below we describe each field in more detail.

environment

The environment describes the origin of a metric record. You define this value in the initial configuration of Rally. The intention is to clearly separate different benchmarking environments but still allow to store them in the same index.

track, track-params, challenge, car

This is the track, challenge and car for which the metrics record has been produced. If the user has provided track parameters with the command line parameter, `--track-params`, each of them is listed here too.

If you specify a car with mixins, it will be stored as one string separated with “+”, e.g. `--car="4gheap,ea"` will be stored as `4gheap+ea` in the metrics store in order to simplify querying in Kibana. Check the *cars* documentation for more details.

sample-type

Rally runs warmup trials but records all samples. Normally, we are just interested in “normal” samples but for a full picture we might want to look also at “warmup” samples.

trial-timestamp

A constant timestamp (always in UTC) that is determined when Rally is invoked.

trial-id

A UUID that changes on every invocation of Rally. It is intended to group all samples of a benchmark trial.

@timestamp

The timestamp in milliseconds since epoch determined when the sample was taken.

relative-time

The relative time in microseconds since the start of the benchmark. This is useful for comparing time-series graphs over multiple trials, e.g. you might want to compare the indexing throughput over time across multiple benchmark trials. Obviously, they should always start at the same (relative) point in time and absolute timestamps are useless for that.

name, value, unit

This is the actual metric name and value with an optional unit (counter metrics don't have a unit). Depending on the nature of a metric, it is either sampled periodically by Rally, e.g. the CPU utilization or query latency or just measured once like the final size of the index.

task, operation, operation-type

`task` is the name of the task (as specified in the track file) that ran when this metric has been gathered. Most of the time, this value will be identical to the operation's name but if the same operation is ran multiple times, the task name will be unique whereas the operation may occur multiple times. It will only be set for metrics with name `latency` and `throughput`.

`operation` is the name of the operation (as specified in the track file) that ran when this metric has been gathered. It will only be set for metrics with name `latency` and `throughput`.

`operation-type` is the more abstract type of an operation. During a race, multiple queries may be issued which are different `operation`s` but they all have the same `operation-type` (Search). For some metrics, only the operation type matters, e.g. it does not make any sense to attribute the CPU usage to an individual query but instead attribute it just to the operation type.

lap

The lap number in which this metric was gathered. Laps start at 1. See the [command line reference](#) for more info on laps.

meta

Rally captures also some meta information for each metric record:

- CPU info: number of physical and logical cores and also the model name
- OS info: OS name and version
- Host name
- Node name: If Rally provisions the cluster, it will choose a unique name for each node.
- Source revision: We always record the git hash of the version of Elasticsearch that is benchmarked. This is even done if you benchmark an official binary release.
- Distribution version: We always record the distribution version of Elasticsearch that is benchmarked. This is even done if you benchmark a source release.
- Custom tag: You can define one custom tag with the command line flag `--user-tag`. The tag is prefixed by `tag_` in order to avoid accidental clashes with Rally internal tags.
- Operation-specific: The optional substructure `operation` contains additional information depending on the type of operation. For bulk requests, this may be the number of documents or for searches the number of hits.

Note that depending on the “level” of a metric record, certain meta information might be missing. It makes no sense to record host level meta info for a cluster wide metric record, like a query latency (as it cannot be attributed to a single node).

2.17.2 Metric Keys

Rally stores the following metrics:

- `latency`: Time period between submission of a request and receiving the complete response. It also includes wait time, i.e. the time the request spends waiting until it is ready to be serviced by Elasticsearch.
- `service_time`: Time period between start of request processing and receiving the complete response. This metric can easily be mixed up with `latency` but does not include waiting time. This is what most load testing tools refer to as “latency” (although it is incorrect).
- `throughput`: Number of operations that Elasticsearch can perform within a certain time period, usually per second. See the [track reference](#) for a definition of what is meant by one “operation” for each operation type.
- `merge_parts_total_time_*`: Different merge times as reported by Lucene. Only available if Lucene index writer trace logging is enabled.
- `merge_parts_total_docs_*`: See `merge_parts_total_time_*`
- `disk_io_write_bytes`: number of bytes that have been written to disk during the benchmark. On Linux this metric reports only the bytes that have been written by Elasticsearch, on Mac OS X it reports the number of bytes written by all processes.
- `disk_io_read_bytes`: number of bytes that have been read from disk during the benchmark. The same caveats apply on Mac OS X as for `disk_io_write_bytes`.
- `cpu_utilization_1s`: CPU usage in percent of the Elasticsearch process based on a one second sample period. The maximum value is $N * 100\%$ where N is the number of CPU cores available.
- `node_startup_time`: The time in seconds it took from process start until the node is up.
- `node_total_old_gen_gc_time`: The total runtime of the old generation garbage collector across the whole cluster as reported by the node stats API.
- `node_total_young_gen_gc_time`: The total runtime of the young generation garbage collector across the whole cluster as reported by the node stats API.
- `segments_count`: Total number of segments as reported by the indices stats API.
- `segments_memory_in_bytes`: Number of bytes used for segments as reported by the indices stats API.
- `segments_doc_values_memory_in_bytes`: Number of bytes used for doc values as reported by the indices stats API.
- `segments_stored_fields_memory_in_bytes`: Number of bytes used for stored fields as reported by the indices stats API.
- `segments_terms_memory_in_bytes`: Number of bytes used for terms as reported by the indices stats API.
- `segments_norms_memory_in_bytes`: Number of bytes used for norms as reported by the indices stats API.
- `segments_points_memory_in_bytes`: Number of bytes used for points as reported by the indices stats API.
- `merges_total_time`: Total runtime of merges as reported by the indices stats API. Note that this is not Wall clock time (i.e. if M merge threads ran for N minutes, we will report $M * N$ minutes, not N minutes). These metrics records also have a `per-shard` property that contains the times per primary shard in an array.
- `merges_total_throttled_time`: Total time within merges have been throttled as reported by the indices stats API. Note that this is not Wall clock time. These metrics records also have a `per-shard` property that contains the times per primary shard in an array.

- `indexing_total_time`: Total time used for indexing as reported by the indices stats API. Note that this is not Wall clock time. These metrics records also have a `per-shard` property that contains the times per primary shard in an array.
- `indexing_throttle_time`: Total time that indexing has been throttled as reported by the indices stats API. Note that this is not Wall clock time. These metrics records also have a `per-shard` property that contains the times per primary shard in an array.
- `refresh_total_time`: Total time used for index refresh as reported by the indices stats API. Note that this is not Wall clock time. These metrics records also have a `per-shard` property that contains the times per primary shard in an array.
- `flush_total_time`: Total time used for index flush as reported by the indices stats API. Note that this is not Wall clock time. These metrics records also have a `per-shard` property that contains the times per primary shard in an array.
- `final_index_size_bytes`: Final resulting index size on the file system after all nodes have been shut-down at the end of the benchmark. It includes all files in the nodes' data directories (actual index files and translog).
- `store_size_in_bytes`: The size in bytes of the index (excluding the translog) as reported by the indices stats API.
- `translog_size_in_bytes`: The size in bytes of the translog as reported by the indices stats API.

2.18 Summary Report

At the end of each *race*, Rally shows a summary report. Below we'll explain the meaning of each line including a reference to its corresponding *metrics key* which can be helpful if you want to build your own reports in Kibana. Note that not every summary report will show all lines.

2.18.1 Total indexing time

- **Definition:** Total time used for indexing as reported by the indices stats API. Note that this is not Wall clock time (i.e. if M indexing threads ran for N minutes, we will report M * N minutes, not N minutes).
- **Corresponding metrics key:** `indexing_total_time`

2.18.2 Indexing time per shard

- **Definition:** Minimum, median and maximum time used for indexing per primary shard as reported by the indices stats API.
- **Corresponding metrics key:** `indexing_total_time` (property: `per-shard`)

2.18.3 Total indexing throttle time

- **Definition:** Total time that indexing has been throttled as reported by the indices stats API. Note that this is not Wall clock time (i.e. if M indexing threads ran for N minutes, we will report M * N minutes, not N minutes).
- **Corresponding metrics key:** `indexing_throttle_time`

2.18.4 Indexing throttle time per shard

- **Definition:** Minimum, median and maximum time used that indexing has been throttled per primary shard as reported by the indices stats API.
- **Corresponding metrics key:** `indexing_throttle_time` (property: `per-shard`)

2.18.5 Total merge time

- **Definition:** Total runtime of merges as reported by the indices stats API. Note that this is not Wall clock time.
- **Corresponding metrics key:** `merges_total_time`

2.18.6 Merge time per shard

- **Definition:** Minimum, median and maximum time of merges per primary shard as reported by the indices stats API.
- **Corresponding metrics key:** `merges_total_time` (property: `per-shard`)

2.18.7 Total refresh time

- **Definition:** Total time used for index refresh as reported by the indices stats API. Note that this is not Wall clock time.
- **Corresponding metrics key:** `refresh_total_time`

2.18.8 Refresh time per shard

- **Definition:** Minimum, median and maximum time for index refresh per primary shard as reported by the indices stats API.
- **Corresponding metrics key:** `refresh_total_time` (property: `per-shard`)

2.18.9 Total flush time

- **Definition:** Total time used for index flush as reported by the indices stats API. Note that this is not Wall clock time.
- **Corresponding metrics key:** `flush_total_time`

2.18.10 Flush time per shard

- **Definition:** Minimum, median and maximum time for index flush per primary shard as reported by the indices stats API.
- **Corresponding metrics key:** `flush_total_time` (property: `per-shard`)

2.18.11 Total merge throttle time

- **Definition:** Total time within merges have been throttled as reported by the indices stats API. Note that this is not Wall clock time.
- **Corresponding metrics key:** `merges_total_throttled_time`

2.18.12 Merge throttle time per shard

- **Definition:** Minimum, median and maximum time that merges have been throttled per primary shard as reported by the indices stats API.
- **Corresponding metrics key:** `merges_total_throttled_time` (property: `per-shard`)

2.18.13 Merge time (x)

Where X is one of:

- postings
- stored fields
- doc values
- norms
- vectors
- points
- **Definition:** Different merge times as reported by Lucene. Only available if Lucene index writer trace logging is enabled (use `--car-params="verbose_iw_logging_enabled:true"` for that).
- **Corresponding metrics keys:** `merge_parts_total_time_*`

2.18.14 Median CPU usage

- **Definition:** Median CPU usage in percent of the Elasticsearch process during the whole race based on a one second sample period. The maximum value is $N * 100\%$ where N is the number of CPU cores available
- **Corresponding metrics key:** `cpu_utilization_1s`

2.18.15 Total Young Gen GC

- **Definition:** The total runtime of the young generation garbage collector across the whole cluster as reported by the node stats API.
- **Corresponding metrics key:** `node_total_young_gen_gc_time`

2.18.16 Total Old Gen GC

- **Definition:** The total runtime of the old generation garbage collector across the whole cluster as reported by the node stats API.
- **Corresponding metrics key:** `node_total_old_gen_gc_time`

2.18.17 Index size

- **Definition:** Final resulting index size on the file system after all nodes have been shutdown at the end of the benchmark. It includes all files in the nodes' data directories (actual index files and translog).
- **Corresponding metrics key:** `final_index_size_bytes`

2.18.18 Store size

- **Definition:** The size in bytes of the index (excluding the translog) as reported by the indices stats API.
- **Corresponding metrics key:** `store_size_in_bytes`

2.18.19 Translog size

- **Definition:** The size in bytes of the translog as reported by the indices stats API.
- **Corresponding metrics key:** `translog_size_in_bytes`

2.18.20 Totally written

- **Definition:** number of bytes that have been written to disk during the benchmark. On Linux this metric reports only the bytes that have been written by Elasticsearch, on Mac OS X it reports the number of bytes written by all processes.
- **Corresponding metrics key:** `disk_io_write_bytes`

2.18.21 Heap used for x

Where X is one of:

- doc values
- terms
- norms
- points
- stored fields
- **Definition:** Number of bytes used for the corresponding item as reported by the indices stats API.
- **Corresponding metrics keys:** `segments_*_in_bytes`

2.18.22 Segment count

- **Definition:** Total number of segments as reported by the indices stats API.
- **Corresponding metrics key:** `segments_count`

2.18.23 Throughput

Rally reports the minimum, median and maximum throughput for each task.

- **Definition:** Number of operations that Elasticsearch can perform within a certain time period, usually per second.
- **Corresponding metrics key:** `throughput`

2.18.24 Latency

Rally reports several percentile numbers for each task. Which percentiles are shown depends on how many requests Rally could capture (i.e. Rally will not show a 99.99th percentile if it could only capture five samples because that would be a vanity metric).

- **Definition:** Time period between submission of a request and receiving the complete response. It also includes wait time, i.e. the time the request spends waiting until it is ready to be serviced by Elasticsearch.
- **Corresponding metrics key:** `latency`

2.18.25 Service time

Rally reports several percentile numbers for each task. Which percentiles are shown depends on how many requests Rally could capture (i.e. Rally will not show a 99.99th percentile if it could only capture five samples because that would be a vanity metric).

- **Definition:** Time period between start of request processing and receiving the complete response. This metric can easily be mixed up with `latency` but does not include waiting time. This is what most load testing tools refer to as “latency” (although it is incorrect).
- **Corresponding metrics key:** `service_time`

2.18.26 Error rate

- **Definition:** The ratio of erroneous responses relative to the total number of responses. Any exception thrown by the Python Elasticsearch client is considered erroneous (e.g. HTTP response codes 4xx, 5xx or network errors (network unreachable)). For specific details, check the [reference documentation of the Elasticsearch client](#). Usually any error rate greater than zero is alerting. You should investigate the root cause by inspecting Rally and Elasticsearch logs and rerun the benchmark.
- **Corresponding metrics key:** `service_time`. Each `service_time` record has a `meta.success` flag. Rally simply counts how often this flag is `true` and `false` respectively.

2.19 Migration Guide

2.19.1 Migrating to Rally 1.0.0

Handling of JDK versions

Previously the path to the JDK needed to be configured in Rally’s configuration file (`~/.rally/rally.ini`) but this is too inflexible given the increased JDK release cadence. In order to keep up, we define now the allowed runtime JDKs in `rally-teams` per Elasticsearch version.

To resolve the path to the appropriate JDK you need to define the environment variable `JAVA_HOME` on each targeted machine.

You can also set version-specific environment variables, e.g. `JAVA7_HOME`, `JAVA8_HOME` or `JAVA10_HOME` which will take precedence over `JAVA_HOME`.

Note: Rally will choose the highest appropriate JDK per Elasticsearch version. You can use `--runtime-jdk` to force a specific JDK version but the path will still be resolved according to the logic above.

Custom Parameter Sources

In Rally 0.10.0 we have deprecated some parameter names in custom parameter sources. In Rally 1.0.0, these deprecated names have been removed. Therefore you need to replace the following parameter names if you use them in custom parameter sources:

Operation type	Old name	New name
search	use_request_cache	cache
search	request_params	request-params
search	items_per_page	results-per-page
bulk	action_metadata_present	action-metadata-present
force-merge	max_num_segments	max-num-segments

In Rally 0.9.0 the signature of custom parameter sources has also changed. In Rally 1.0.0 we have removed the backwards compatibility layer so you need to change the signatures.

Old:

```
# for parameter sources implemented as functions
def custom_param_source(indices, params):

# for parameter sources implemented as classes
class CustomParamSource:
    def __init__(self, indices, params):
```

New:

```
# for parameter sources implemented as functions
def custom_param_source(track, params, **kwargs):

# for parameter sources implemented as classes
class CustomParamSource:
    def __init__(self, track, params, **kwargs):
```

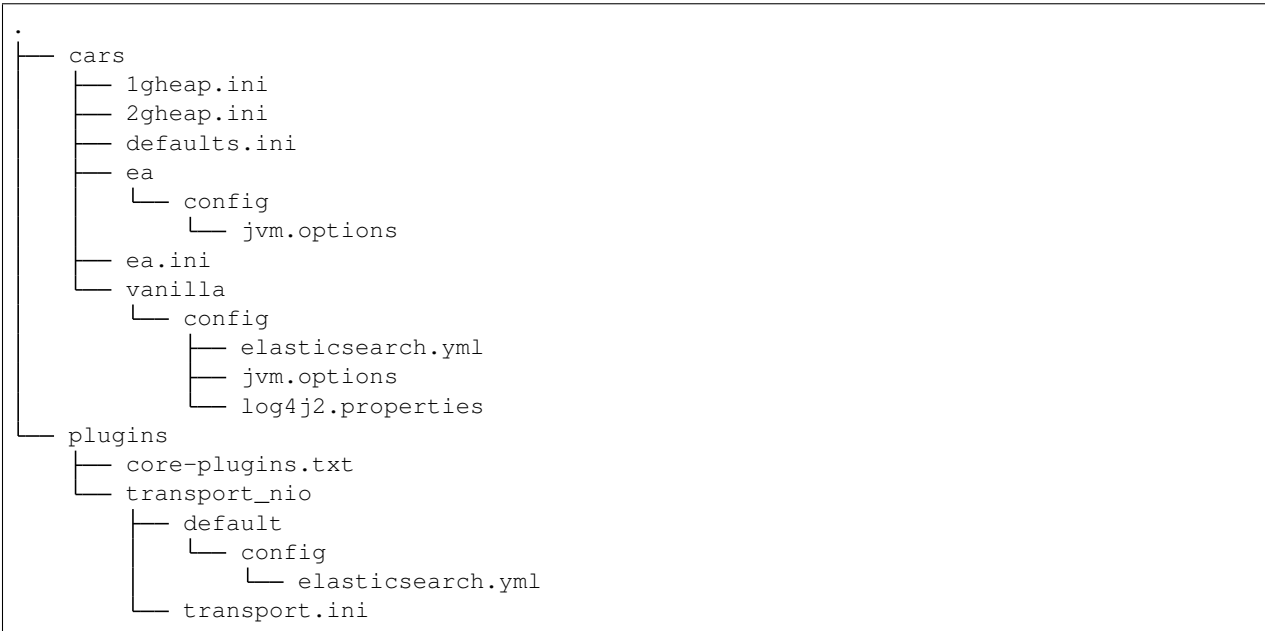
You can use the property `track.indices` to access indices.

2.19.2 Migrating to Rally 0.11.0

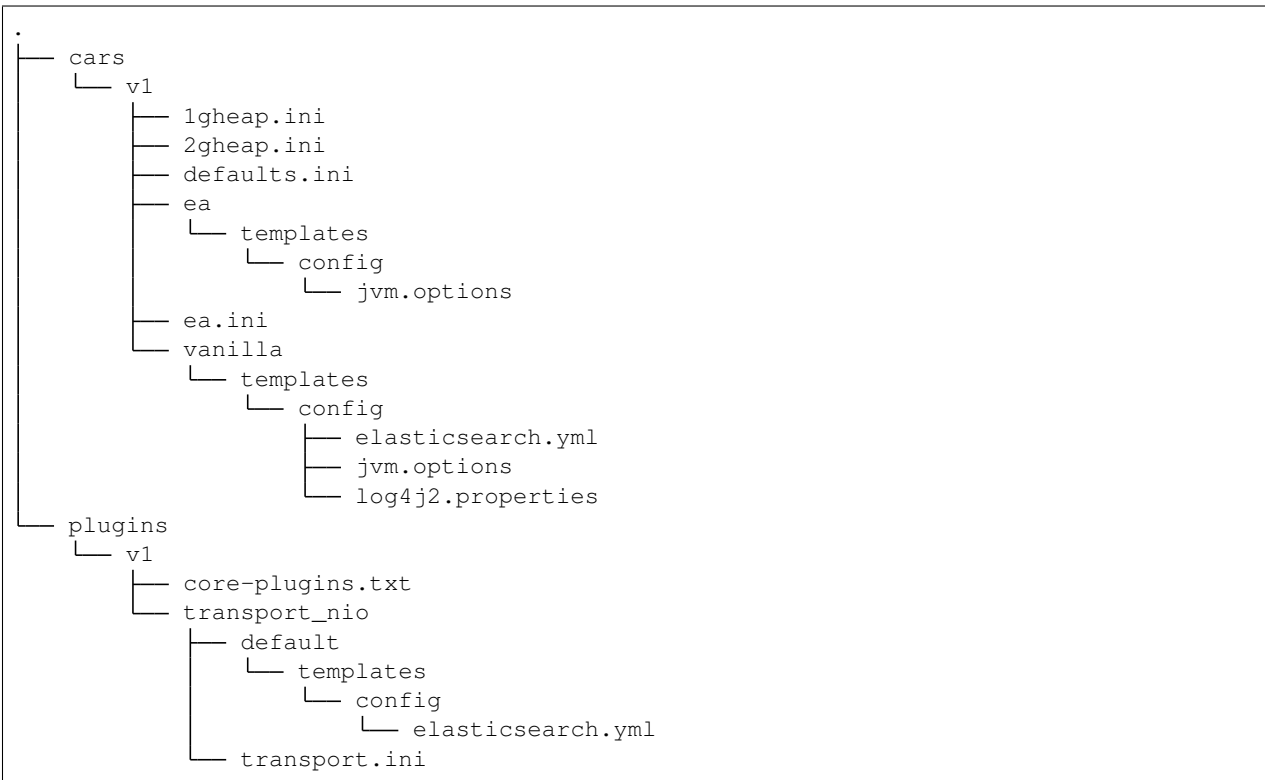
Versioned teams

Note: You can skip this section if you do not create custom Rally teams.

We have introduced versioned team specifications and consequently the directory structure changes. All cars and plugins need to reside in a version-specific subdirectory now. Up to now the structure of a team repository was as follows:

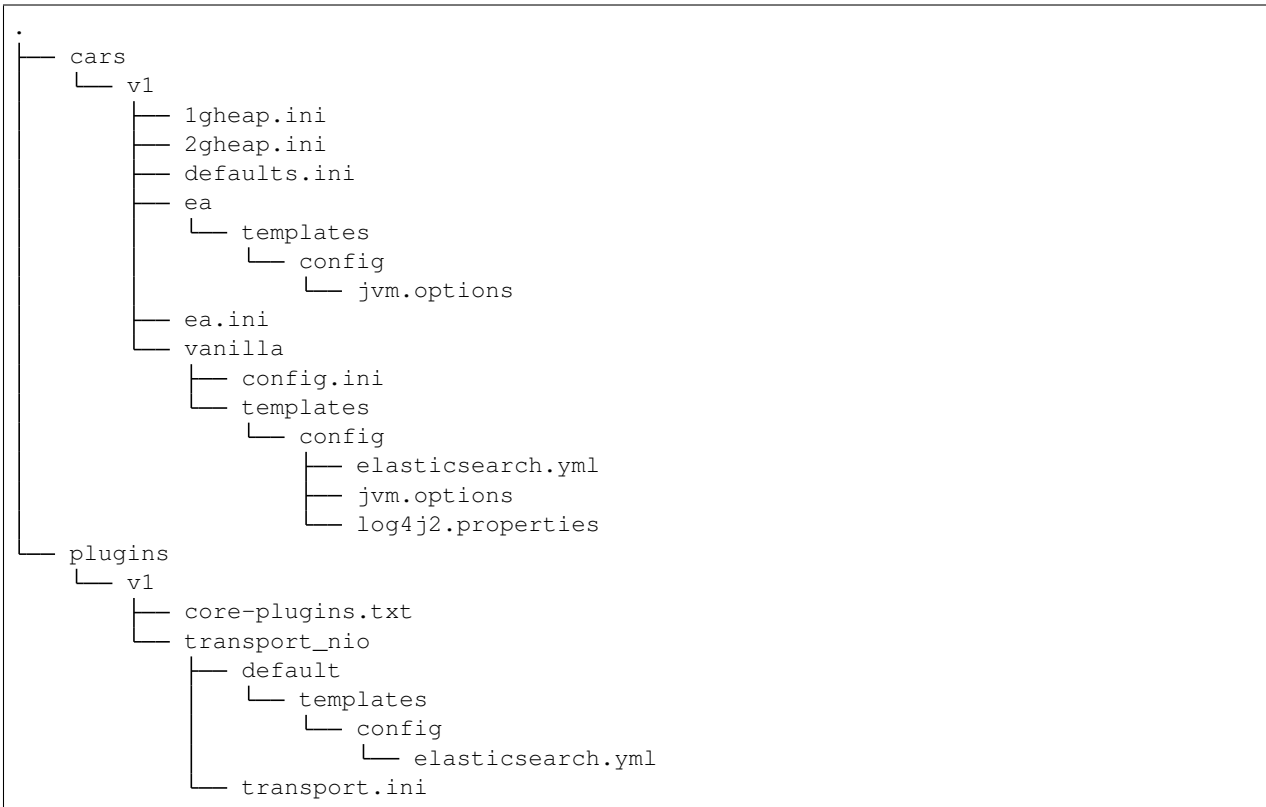


Starting with Rally 0.11.0, Rally will look for a directory “v1” within cars and plugins. The files that should be copied to the Elasticsearch directory, need to be contained in a templates subdirectory. Therefore, the new structure is as follows:



It is also required that you create a file `variables.ini` for all your car config bases (optional for mixins). There-

fore, the full directory structure is:



For distribution-based builds, `config.ini` file needs to contain a section `variables` and a `release_url` property:

```
[variables]
release_url=https://artifacts.elastic.co/downloads/elasticsearch/elasticsearch-oss-
↔{VERSION}.tar.gz
```

2.19.3 Migrating to Rally 0.10.0

Removal of auto-detection and dependency on Gradle

We have removed the auto-detection and dependency on Gradle, required until now to build from source, in favor of the [Gradle Wrapper](#) which is present in the [Elasticsearch repository](#) for all branches $\geq 5.0.0$.

Use full build command in plugin configuration

With Rally 0.10.0 we have removed the property `build.task` for plugin definitions, in the source section of the Rally configuration file. Instead, a new property `build.command` has been introduced where the **full build command** needs to be supplied.

The earlier syntax, to build a hypothetical plugin called `my-plugin` alongside [Elasticsearch](#), required:

```
plugin.my-plugin.build.task = :my-plugin:plugin:assemble
```

This needs to be changed to the full command:

```
plugin.my-plugin.build.command = ./gradlew :my-plugin:plugin:assemble
```

Note that if you are configuring [Plugins based on a released Elasticsearch version](#) the command specified in `build.command` will be executed from the plugins root directory. It's likely this directory won't have the Gradle Wrapper so you'll need to specify the full path to a Gradle command e.g.:

```
plugin.my-plugin.build.command = /usr/local/bin/gradle :my-plugin:plugin:assemble
```

Check [Building plugins from sources](#) for more information.

Removal of operation type `index`

We have removed the operation type `index` which has been deprecated with Rally 0.8.0. Use `bulk` instead as operation type.

Removal of the command line parameter `--cluster-health`

We have removed the command line parameter `--cluster-health` which has been deprecated with Rally 0.8.0. When using Rally's standard tracks, specify the expected cluster health as a track parameter instead, e.g.: `--track-params="cluster_health:'yellow'".`

Removal of `index-automanagement`

We have removed the possibility that Rally automatically deletes and creates indices. Therefore, you need to add the following definitions explicitly at the beginning of a schedule if you want Rally to create declared indices:

```
"schedule": [
  {
    "operation": "delete-index"
  },
  {
    "operation": {
      "operation-type": "create-index",
      "settings": {
        "index.number_of_replicas": 0
      }
    }
  },
  {
    "operation": {
      "operation-type": "cluster-health",
      "request-params": {
        "wait_for_status": "green"
      }
    }
  }
]
```

The example above also shows how to provide per-challenge index settings. If per-challenge index settings are not required, you can just specify them in the index definition file.

This behavior applies similarly to index templates as well.

Custom Parameter Sources

We have aligned the internal names between parameter sources and runners with the ones that are specified by the user in the track file. If you have implemented custom parameter sources or runners, adjust the parameter names as follows:

Operation type	Old name	New name
search	use_request_cache	cache
search	request_params	request-params
search	items_per_page	results-per-page
bulk	action_metadata_present	action-metadata-present
force-merge	max_num_segments	max-num-segments

2.19.4 Migrating to Rally 0.9.0

Track Syntax

With Rally 0.9.0, we have changed the track file format. While the previous format is still supported with deprecation warnings, we recommend that you adapt your tracks as we will remove the deprecated syntax with the next minor release.

Below is an example of a track with the previous syntax:

```
{
  "description": "Tutorial benchmark for Rally",
  "data-url": "http://benchmarks.elasticsearch.org.s3.amazonaws.com/corpora/geonames",
  "indices": [
    {
      "name": "geonames",
      "types": [
        {
          "name": "type",
          "mapping": "mappings.json",
          "documents": "documents.json",
          "document-count": 8647880,
          "uncompressed-bytes": 2790927196
        }
      ]
    }
  ],
  "challenge": {
    "name": "index-only",
    "index-settings": {
      "index.number_of_replicas": 0
    },
    "schedule": [
      {
        "operation": {
          "operation-type": "bulk",
          "bulk-size": 5000
        },
        "warmup-time-period": 120,
        "clients": 8
      }
    ]
  }
}
```

(continues on next page)

(continued from previous page)

```
}
}
```

Before Rally 0.9.0, indices have been created implicitly. We will remove this ability and thus you need to tell Rally explicitly that you want to create indices. With Rally 0.9.0 your track should look as follows:

```
{
  "description": "Tutorial benchmark for Rally",
  "indices": [
    {
      "name": "geonames",
      "body": "index.json",
      "auto-managed": false,
      "types": [ "type" ]
    }
  ],
  "corpora": [
    {
      "name": "geonames",
      "documents": [
        {
          "base-url": "http://benchmarks.elasticsearch.org.s3.amazonaws.com/corpora/
↪geonames",
          "source-file": "documents.json",
          "document-count": 8647880,
          "uncompressed-bytes": 2790927196
        }
      ]
    }
  ],
  "challenge": {
    "name": "index-only",
    "schedule": [
      {
        "operation": "delete-index"
      },
      {
        "operation": {
          "operation-type": "create-index",
          "settings": {
            "index.number_of_replicas": 0
          }
        }
      },
      {
        "operation": {
          "operation-type": "cluster-health",
          "request-params": {
            "wait_for_status": "green"
          }
        }
      },
      {
        "operation": {
          "operation-type": "bulk",
          "bulk-size": 5000
        }
      }
    ]
  }
}
```

(continues on next page)

(continued from previous page)

```

    },
    "warmup-time-period": 120,
    "clients": 8
  }
]
}
}

```

Let's go through the necessary changes one by one.

Define the document corpus separately

Previously you had to define the document corpus together with the document type. In order to allow you to reuse existing document corpora across tracks, you now need to specify any document corpora separately:

```

"corpora": [
  {
    "name": "geonames",
    "documents": [
      {
        "base-url": "http://benchmarks.elasticsearch.org.s3.amazonaws.com/corpora/
→geonames",
        "source-file": "documents.json",
        "document-count": 8647880,
        "uncompressed-bytes": 2790927196
      }
    ]
  }
]

```

Note that this is just a simple example that should cover the most basic case. Be sure to check the *track reference* for all details.

Change the index definition

The new index definition now looks as follows:

```

{
  "name": "geonames",
  "body": "index.json",
  "auto-managed": false,
  "types": [ "type" ]
}

```

We have added a `body` property to the index and removed the `mapping` property from the type. In fact, the only information that we need about the document type is its name, hence it is now a simple list of strings. Just put all type mappings now into the `mappings` property of the index definition; see also the [create index API documentation](#).

Secondly, we have disabled index auto-management by setting `auto-managed` to `false`. This allows us to define explicit tasks below to manage our index. Note that index auto-management is still working in Rally 0.9.0 but it will be removed with the next minor release Rally 0.10.0.

Explicitly delete and recreate the index

We have also added three tasks at the beginning of the schedule:

```
{
  "operation": "delete-index"
},
{
  "operation": {
    "operation-type": "create-index",
    "settings": {
      "index.number_of_replicas": 0
    }
  }
},
{
  "operation": {
    "operation-type": "cluster-health",
    "request-params": {
      "wait_for_status": "green"
    }
  }
}
```

These tasks represent what Rally previously did implicitly.

The first task will delete all indices that have been declared in the `indices` section if they existed previously. This ensures that we don't have any leftovers from previous benchmarks.

After that we will create all indices that have been declared in the `indices` section. Note that we have also removed the special property `index-settings` and moved it to the `settings` parameter of `create-index`. Rally will merge any settings from the index body definition with these settings. This means you should define settings that are always the same in the index body and settings that change from challenge to challenge in the `settings` property.

Finally, Rally will check that the cluster health is green. If you want to be able to override the cluster health check parameters from the command line, you can leverage Rally's track parameter feature:

```
{
  "operation": {
    "operation-type": "cluster-health",
    "request-params": {
      "wait_for_status": "{{ cluster_health|default('green') }}"
    }
  }
}
```

If you don't specify anything on the command line, Rally will use the default value but you can e.g. specify `--track-params="cluster_health:'yellow'"` so Rally will check for (at least) a yellow cluster health status.

Note that you can *customize these operations*.

Custom Parameter Sources

With Rally 0.9.0, the API for custom parameter sources has changed. Previously, the following syntax was valid:

```
# for parameter sources implemented as functions
def custom_param_source(indices, params):

# for parameter sources implemented as classes
class CustomParamSource:
    def __init__(self, indices, params):
```

With Rally 0.9.0, the signatures need to be changed to:

```
# for parameter sources implemented as functions
def custom_param_source(track, params, **kwargs):

# for parameter sources implemented as classes
class CustomParamSource:
    def __init__(self, track, params, **kwargs):
```

Rally will issue a warning along the lines of Parameter source 'custom_param_source' is using deprecated method signature if your track is affected. If you need access to the indices list, you can call `track.indices` to retrieve it from the track.

2.20 Frequently Asked Questions (FAQ)

2.20.1 A benchmark aborts with Couldn't find a tar.gz distribution. What's the problem?

This error occurs when Rally cannot build an Elasticsearch distribution from source code. The most likely cause is that there is some problem building the Elasticsearch distribution.

To see what's the problem, try building Elasticsearch yourself. First, find out where the source code is located (run `grep src ~/.rally/rally.ini`). Then change to the directory (`src.root.dir + elasticsearch.src.subdir` which is usually `~/.rally/benchmarks/src/elasticsearch`) and run the following commands:

```
./gradlew clean
./gradlew :distribution:tar:assemble
```

By that you are mimicking what Rally does. Fix any errors that show up here and then retry.

2.20.2 Where does Rally get the benchmark data from?

Rally comes with a set of tracks out of the box which we maintain in the [rally-tracks repository on Github](#). This repository contains the track descriptions. The actual data are stored as compressed files in an S3 bucket.

2.20.3 Will Rally destroy my existing indices?

First of all: Please (please, please) do NOT run Rally against your production cluster if you are just getting started with it. You have been warned.

Depending on the track, Rally will delete and create one or more indices. For example, the [geonames track](#) specifies that Rally should create an index named “geonames” and Rally will assume it can do to this index whatever it wants. Specifically, Rally will check at the beginning of a race if the index “geonames” exists and delete it. After that it creates a new empty “geonames” index and runs the benchmark. So if you benchmark against your own cluster (by

specifying the `benchmark-only pipeline`) and this cluster contains an index that is called “geonames” you will lose (all) data if you run Rally against it. Rally will neither read nor write (or delete) any other index. So if you apply the usual care nothing bad can happen.

2.20.4 What does *latency* and *service_time* mean and how do they related to the *took* field that Elasticsearch returns?

Let’s start with the *took* field of Elasticsearch. *took* is the time needed by Elasticsearch to process a request. As it is determined on the server, it can neither include the time it took the client to send the data to Elasticsearch nor the time it took Elasticsearch to send it to the client. This time is captured by *service_time*, i.e. it is the time period from the start of a request (on the client) until it has received the response.

The explanation of *latency* is a bit more involved. First of all, Rally defines two benchmarking modes:

- **Throughput benchmarking mode:** In this mode, Rally will issue requests as fast as it can, i.e. as soon as it receives a response, it will issue the next request. This is ideal for benchmarking indexing. In this mode `latency == service_time`.
- **Throughput-throttled mode:** If you define a specific target throughput rate in a track, for example 100 requests per second (you should choose this number based on the traffic pattern that you experience in your production environment), then Rally will define a schedule internally and will issue requests according to this schedule regardless how fast Elasticsearch can respond. To put it differently: Imagine you want to grab a coffee on your way to work. You make this decision independently of all the other people going to the coffee shop so it is possible that you need to wait before you can tell the barista which coffee you want. The time it takes the barista to make your coffee is the service time. The service time is independent of the number of customers in the coffee shop. However, you as a customer also care about the length of the waiting line which depends on the number of customers in the coffee shop. The time it takes between you entering the coffee shop and taking your first sip of coffee is latency.

If you are interested in latency measurement, we recommend you watch the following talks:

“How NOT to Measure Latency” by Gil Tene:

Benchmarking Elasticsearch with Rally by Daniel Mitterdorfer:

2.20.5 Where and how long does Rally keep its data?

Rally stores a lot of data (this is just the nature of a benchmark) so you should keep an eye on disk usage. All data are kept in `~/.rally` and Rally does not implicitly delete them. These are the most important directories:

- `~/.rally/logs`: Contains all log files. Logs are rotated daily. If you don’t need the logs anymore, you can safely wipe this directory.
- `~/.rally/benchmarks/races`: telemetry data, Elasticsearch logs and even complete Elasticsearch installations including the data directory if a benchmark failed. If you don’t need the data anymore, you can safely wipe this directory.
- `~/.rally/benchmarks/src`: the Elasticsearch Github repository (only if you had Rally build Elasticsearch from sources at least once).
- `~/.rally/benchmarks/data`: the benchmark data sets. This directory can get very huge (way more than 100 GB if you want to try all default tracks). You can delete the files in this directory but keep in mind that Rally may needs to download them again.
- `~/.rally/benchmarks/distributions`: Contains all downloaded Elasticsearch distributions.

There are a few more directories but the ones above are the most disk-hogging ones.

2.20.6 Does Rally spy on me?

No. Rally does not collect or send any usage data and also the complete source code is open. We do value your feedback a lot though and if you got any ideas for improvements, found a bug or have any other feedback, head over to Rally's [Discuss forum](#) or [raise an issue on Github](#).

2.20.7 Do I need an Internet connection?

You do NOT need Internet access on any node of your Elasticsearch cluster but the machine where you start Rally needs an Internet connection to download track data sets and Elasticsearch distributions. After it has downloaded all data, an Internet connection is not required anymore and you can specify `--offline`. If Rally detects no active Internet connection, it will automatically enable offline mode and warn you.

We have a dedicated documentation page for [running Rally offline](#) which should cover all necessary details.

2.21 Glossary

track A *track* is the description of one or more benchmarking scenarios with a specific document corpus. It defines for example the involved indices, data files and which operations are invoked. List the available tracks with `esrally list tracks`. Although Rally ships with some tracks out of the box, you should usually *create your own track* based on your own data.

challenge A challenge describes one benchmarking scenario, for example indexing documents at maximum throughput with 4 clients while issuing term and phrase queries from another two clients rate-limited at 10 queries per second each. It is always specified in the context of a track. See the available challenges by listing the corresponding tracks with `esrally list tracks`.

car A *car* is a specific configuration of an Elasticsearch cluster that is benchmarked, for example the out-of-the-box configuration, a configuration with a specific heap size or a custom logging configuration. List the available cars with `esrally list cars`.

telemetry *Telemetry* is used in Rally to gather metrics about the car, for example CPU usage or index size.

race A *race* is one invocation of the Rally binary. Another name for that is one "benchmarking trial". During a race, Rally runs one challenge on a track with the given car.

tournament A tournament is a comparison of two races. You can use Rally's *tournament mode* for that.

2.22 Community Resources

Below are a few community resources about Rally. If you find an interesting article, talk or custom tracks, [raise an issue](#) or [open a pull request](#).

2.22.1 Talks

2.22.2 Articles

[Using Rally to benchmark Elasticsearch queries](#) by Darren Smith

CHAPTER 3

License

This software is licensed under the Apache License, version 2 (“ALv2”), quoted below.

Copyright 2015-2018 Elasticsearch <<https://www.elastic.co>>

Licensed under the Apache License, Version 2.0 (the “License”); you may not use this file except in compliance with the License. You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an “AS IS” BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

C

car, [117](#)
challenge, [117](#)

R

race, [117](#)

T

telemetry, [117](#)
tournament, [117](#)
track, [117](#)