TEACUP v0.8 - Command Reference

Sebastian Zander
Centre for Advanced Internet Architectures, Technical Report 150210B
Swinburne University of Technology
Melbourne, Australia
szander@swin.edu.au

Abstract

This technical report lists all the TEACUP tasks implemented and their parameters as well as the environment variables that can be used to control the look of graphs.

Index Terms

TCP, Testbed experiments

I. Introduction

TEACUP[1] [1] is a software package designed to run TCP experiments. In this report we list all the tasks implemented by TEACUP version 0.8 and their parameters as well as the environment variables that can be used to control the look of graphs. This report is an update of the earlier reports [2], [3].

The tasks are listed in alphabetical order. Each task is explained in its own sub-section. For each task we list and explain all parameters. Note, that all task parameters are strings (as of Fabric version 1.8 and lower).

II. Environment variables for plotting

TEACUP uses environment variables to pass information from the tasks (e.g. analyse_throughput) to the underlying plot functions (R scripts). Here we provide an overview of some environment variables that can be used by a user to customise the plotting of graphs [1]. The user simply needs to define one of these variables in a shell before executing a TEACUP analyse task.

There are many other environment variables used by TEACUP internally. A description of all existing environment variables can be found in the R plot scripts.[2]

Variable	Default Value	Explanation
AGGR_INT_FACTOR	4	Interpolation factor for throughput calculation windows (must be an integer equal or larger than 1). If set to an integer greater than 1, time windows actually overlap with the gap between windows being AGGR_WIN_SIZE / AGGR_INT_FACTOR seconds. This meas we get interpolated points.
AGGR_WIN_SIZE	1	Time window size in seconds (can be fractional value) over which a single value of throughput is calculated
NICER_XLABS	'0'	The boolean variable changes how x-axis labels are plotted for analyse_cmpexp. By default variable names and values are plotted at each x-axis tick. If set to '1' variable names are only plotted once on the left side and only variable values are plotted at each x-axis tick.

[1] "TCP Experiment Automation Controlled Using Python"

[2] Since TEACUP version 0.8 all analyse_* tasks have a parameter called `plot_params` that can be used by a user to set and overrule *any* of the environment variables to customise the plots.

NO_NOMINAL	'0'	If set to '1' analyse_dash_goodput will not plot the nominal goodput line. By default the nominal goodput line will be plotted.
OUTLIER_QUANT	0	Remove outliers before plotting with analyse_cmpexp. Any points in the lowest OUTLIER_QUANT and highest OUTLIER_QUANT quantiles are removed from the plot. For example, specifying OUTLIER_QUANT=0.01 will remove all data points that fall in the <0.01 quantile and all data points that fall in the >0.99 quantile.
POINT_SIZE	0.5	The point size in graphs can be controlled with a variable POINT_SIZE. Note that POINT_SIZE does not specify an absolute point size, but it is a scaling factor that is multiplied with the actual default point size. Hence, if POINT_SIZE is set to 1.0 the size of points will be the default size, if POINT_SIZE is set to 0.5 the size of points will be half the default size and so on.
PTHIN_DIST	0	With PTHIN_DIST one can set the minimum (Euclidean) distance between plotted data points. Any data points within the minimum distance are not plotted. For example, PTHIN_DIST=0.25 means the minimum distance between two plotted points is 0.25 and any data points in-between are not plotted. By default point thinning is disabled. Use only PTHIN_DIST or PTHIN_DIST_FAC.
PTHIN_DIST_FAC	0	PTHIN_DIST_FAC controls the minimum distance of points in x-direction and y-direction separately, relative to the x-range and y-range plotted. A point is plotted if either the distance in x-direction is larger or equal PTHIN_DIST_FAC*<xrange> or the distance in y-direction is larger or equal PTHIN_DIST_FAC*<yrange>, where <xrange> and <yrange> are given by the data to plot and the settings applied by the user through specifying ymin, ymax, stime, etime. Use only PTHIN_DIST or PTHIN_DIST_FAC.
YMAX_INC	0.09	The variable YMAX_INC controls the space for the legend. It assumes the legend is plotted at the top, which is the default. The actual y-axis maximum for the plot will be $y_{\max}\,(1 + \mathrm{YMAX\_INC})$, where $y_{\max}$ is the maximum based on the data (or the maximum specified by the user using the ymax parameter).

III. ANALYSE_ALL

This task computes Round Trip Time (RTT), TCP congestion window (CWND) and throughput statistics.

Parameter	Default Value	Explanation
etime	duration of experiment	End time for plot window (x-axis ends at this time).
io_filter	'o'	Specify whether TCP statistics are plotted based on incoming (set to 'i'), outgoing (set to 'o') or incoming and outgoing packets (set to 'io'). Only works with SIFTR logs (FreeBSD).
exp_list	experiments_completed.txt	Specifies the file that contains the test ID list. Statistics will be computed for all experiments listed. Only used if test_id is an empty string.

link_len	‘0’	If set to ‘0’ throughput is calculated based on the length of the IP packets. If set to ‘1’ throughput is calculated based on the link-layer frame length.
lnames	‘’	List of names for legend (separated by semicolons). The list must have the same length as the data series plotted.
min_values	‘3’	Only data series with more than min_values data values are plotted.
omit_const	‘0’	If set to ‘1’ any data series that are constant for the the duration of the experiment are not plotted.
out_dir	‘’	Extracted data files and plots are generated in this directory. By default the files are generated in the experiment directory.
out_name	‘’	A user-defined string that is used as prefix for the generated plot file, it can be used to describe the plot file.
pdf_dir	out_dir	Store generated plots (PDF files) and plot logs (Rout files) in this directory. By default the files are created in the out_dir directory.
replot_only	‘0’	If set to ‘1’ the data extraction is skipped and the plots are regenerated based on the data previously extracted.
resume_id	‘’	If a test ID is specified, the analysis will resume this test ID. The parameter implies that a list of test IDs is used, i.e. test_id is empty and exp_list points to a file of test IDs.
source_filter	‘’	Specify the sources or destinations to filter on. Metrics will only be plotted for the sources and destinations on the list. See [1] for how to specify the list
smoothed	‘1’	If set to ‘1’ (default) smoothed TCP RTTs are plotted. If set to ‘0’ unsmoothed TCP RTT estimates are plotted and for SIFTR data the ERTT [4] estimates are plotted.
stime	‘0.0’	Start time for plot window (x-axis starts at this time).
test_id	‘’	Specifies the test ID of the experiment to be analysed. Specifying multiple test IDs with semicolons as separators allows to plot the results of multiple experiments in the same graphs. If an empty string the test IDs will be read from exp_list.
ts_correct	‘0’	If set to ‘1’ plot data with timestamps corrected based on estimated clock offsets. This requires to have a clock offset data file that can be created with the get_clock_offsets task. If set to ‘0’ (default) plot data based on timestamps as they are in the log files.
web10g_version	‘2.0.9’	Explicitely specific web10g version, since the log file format depends on the version. By default TEACUP tries to guess the version. Note, this only for the TCP RTT plotting. Note, that the Windows estats logger produces output equivalent to version 2.0.7.

IV. ANALYSE_CMPEXP

This tasks allows to compare one of the metrics, such as RTT, CWND or throughput, for experiments with different settings.

Parameter	Default Value	Explanation
etime	duration of experiment	End time for filtering window. For each experiment (test ID) only data inside the specified time window is used for plotting.
exp_list	experiments_completed.txt	Specifies the file that contains the test ID list. All listed experiments will be potentially included in the comparison. variables allows to further filter out experiments.
group-by-experiment	'0'	If set to '0' each group is a flow identified by the tuple source IP, source port, destination IP, destination port. If set to '1' each group is a series of experiments identified by a test ID prefix. In this case the flows of different test ID prefixes can have different flow tuples but they should be comparable, e.g. same type of traffic.
lnames	''	Semicolon-separated list of legend names to use for the flows filtered with source_filter. Must be of the same length as the source filter list.
metric	'throughput'	The metric to use. Currently supported metrics are 'throughput', 'spprtt' or 'tcprtt'.
min_values	'3'	Only data series with more than min_values data values are plotted.
omit_const	'0'	If set to '1' any data series that are constant for the the duration of the experiment are not plotted.
omit_const_xlab_vars	'0'	If set to '1' any variables that have been constant are omitted from the x-axis labels.
out_dir	''	Newly extracted data files and plots are generated in this directory (defined relative to the experiment directory).
out_name	''	A user-defined string that is used as prefix for the generated plot file, it can be used to describe the plot file.
pdf_dir	out_dir	Store generated plots (PDF files) and plot logs (Rout files) in this directory. Defaults to out_dir if not otherwise set.
ptype	'box'	Specifies the type of plot. Must be either 'box', 'median' or 'mean'.
res_dir	''	Directory that contains previously extracted data for the experiments. If this is an empty string, first analyse_all is executed and newly extracted data is placed in out_dir before proceedings with generating the comparison plot.
smoothed	'1'	If set to '1' (default) smoothed TCP RTTs are plotted. If set to '0' unsmoothed TCP RTT estimates are plotted and for SIFTR data the ERTT [4] estimates are plotted.
source_filter	''	Specify the sources or destinations to filter on. Metrics will only be plotted for the sources and destinations on the list. See [1] for how to specify the list
stime	'0.0'	Start time for filtering window. For each experiment (test ID) only data inside the specified time window is used for plotting.

ts_correct	'0'	If set to '1' plot data with timestamps corrected based on estimated clock offsets. This requires to have a clock offset data file that can be created with the get_clock_offsets task. If set to '0' (default) plot data based on timestamps as they are in the log files.
variables	''	Semicolon-separated list of the form <var>=<value>[;<var>=<value]*, where <var> is an experiment variable name (the name as it appear in the file names) and value is a value. Only experiments where the variables listed had the values listed will be included in the comparison.
ymax	'0'	Maximum y-axis value. By default (if ymax set to '0') the maximum will be determined automatically. The parameter can be used to enforce a certain maximum, i.e. to generate different plots with the same scale.
ymin	'0'	Minimum y-axis value. The parameter can be used to enforce the specified minimum.

V. ANALYSE_CWND

This tasks extracts TCP CWND data and plots it over time.

Parameter	Default Value	Explanation
etime	duration of experiment	End time for plot window (x-axis ends at this time).
io_filter	'o'	Specify whether TCP statistics are plotted based on incoming (set to 'i'), outgoing (set to 'o') or incoming and outgoing packets (set to 'io'). Only works with SIFTR logs (FreeBSD).
lnames	''	List of names for legend (separated by semicolons). The list must have the same length as the data series plotted.
min_values	'3'	Only data series with more than min_values data values are plotted.
omit_const	'0'	If set to '1' any data series that are constant for the the duration of the experiment are not plotted.
out_dir	''	Extracted data files and plots are generated in this directory. By default the files are generated in the experiment directory.
out_name	''	A user-defined string that is used as prefix for the generated plot file, it can be used to describe the plot file.
pdf_dir	out_dir	Store generated plots (PDF files) and plot logs (Rout files) in this directory. By default the files are created in the out_dir directory.
replot_only	'0'	If set to '1' the data extraction is skipped and the plots are regenerated based on the data previously extracted.
source_filter	''	Specify the sources or destinations to filter on. Metrics will only be plotted for the sources and destinations on the list. See [1] for how to specify the list
stime	'0.0'	Start time for plot window (x-axis starts at this time).
test_id	''	Specifies the test ID of the experiment to be analysed. Must be specified. Specifying multiple test IDs with semicolons as separators allows to plot the results of multiple experiments in the same graph.

ts_correct	'0'	If set to '1' plot data with timestamps corrected based on estimated clock offsets. This requires to have a clock offset data file that can be created with the get_clock_offsets task. If set to '0' (default) plot data based on timestamps as they are in the log files.
ymax	'0'	Maximum y-axis value. By default (if ymax set to '0') the maximum will be determined automatically. The parameter can be used to enforce a certain maximum, i.e. to generate different plots with the same scale.
ymin	'0'	Minimum y-axis value. The parameter can be used to enforce the specified minimum.

VI. ANALYSE_DASH_GOODPUT

This tasks allows to compare the goodput of DASH-like flows over time.

Parameter	Default Value	Explanation
dash_log_list	''	Name of a file with a list of DASH logs (*_httperf_dash.log.gz), one name per line (file name only, path information is not required). For each log goodput is plotted over time. If this parameter is not specified, the list of DASH log files is set to all DASH log files for the specified experiment(s) (test_id).
etime	duration of experiment	End time for plot window (x-axis ends at this time).
lnames	''	Semicolon-separated list of legend names to use for the flows filtered with source_filter. Must be of the same length as the number of DASH-like sources.
out_dir	''	Extracted data files and plots are generated in this directory. By default the files are generated in the experiment directory.
out_name	''	A user-defined string that is used as prefix for the generated plot file, it can be used to describe the plot file.
pdf_dir	out_dir	Store generated plots (PDF files) and plot logs (Rout files) in this directory. By default the files are created in the out_dir directory.
replot_only	'0'	If set to '1' the data extraction is skipped and the plots are regenerated based on the data previously extracted.
stime	'0.0'	Start time for plot window (x-axis starts at this time).
test_id	''	Specifies the test ID(s) of the experiment to be analysed.
ts_correct	'0'	If set to '1' plot data with timestamps corrected based on estimated clock offsets. This requires to have a clock offset data file that can be created with the get_clock_offsets task. If set to '0' (default) plot data based on timestamps as they are in the log files.
ymax	'0'	Maximum y-axis value. By default (if ymax set to '0') the maximum will be determined automatically. The parameter can be used to enforce a certain maximum, i.e. to generate different plots with the same scale.
ymin	'0'	Minimum y-axis value. The parameter can be used to enforce the specified minimum.

VII. ANALYSE_INCAST

This tasks plots the response times for queries over time for incast experiment (response times are taken from httperf files).

Parameter	Default Value	Explanation
boxplot	'0'	If set to '0' plot one line per flow. If set to '1' plot boxplots over all flows.
etime	duration of experiment	End time for plot window (x-axis ends at this time).
lnames	''	List of names for legend (separated by semicolons). The list must have the same length as the data series plotted.
min_values	'3'	Only data series with more than min_values data values are plotted.
omit_const	'0'	If set to '1' any data series that are constant for the the duration of the experiment are not plotted.
out_dir	''	Extracted data files and plots are generated in this directory. By default the files are generated in the experiment directory.
out_name	''	A user-defined string that is used as prefix for the generated plot file, it can be used to describe the plot file.
pdf_dir	out_dir	Store generated plots (PDF files) and plot logs (Rout files) in this directory. By default the files are created in the out_dir directory.
replot_only	'0'	If set to '1' the data extraction is skipped and the plots are regenerated based on the data previously extracted.
slowest_only	'0'	If set to '0' plot one per flow. If set to '1' at each point in time only plot the slowest response time over all flows.
source_filter	''	Specify the sources or destinations to filter on. Metrics will only be plotted for the sources and destinations on the list. See [1] for how to specify the list
stime	'0.0'	Start time for plot window (x-axis starts at this time).
test_id	''	Specifies the test ID of the experiment to be analysed. Must be specified. Specifying multiple test IDs with semicolons as separators allows to plot the results of multiple experiments in the same graph.
ts_correct	'0'	If set to '1' plot data with timestamps corrected based on estimated clock offsets. This requires to have a clock offset data file that can be created with the get_clock_offsets task. If set to '0' (default) plot data based on timestamps as they are in the log files.
ymax	'0'	Maximum y-axis value. By default (if ymax set to '0') the maximum will be determined automatically. The parameter can be used to enforce a certain maximum, i.e. to generate different plots with the same scale.
ymin	'0'	Minimum y-axis value. The parameter can be used to enforce the specified minimum.

VIII. ANALYSE_RTT

This tasks computes RTT using SPP [5], [6] and plots the RTT over time.

Parameter	Default Value	Explanation
etime	duration of experiment	End time for plot window (x-axis ends at this time).

lnames	''	List of names for legend (separated by semicolons). The list must have the same length as the data series plotted.
min_values	'3'	Only data series with more than min_values data values are plotted.
omit_const	'0'	If set to '1' any data series that are constant for the the duration of the experiment are not plotted.
out_dir	''	Extracted data files and plots are generated in this directory. By default the files are generated in the experiment directory.
out_name	''	A user-defined string that is used as prefix for the generated plot file, it can be used to describe the plot file.
pdf_dir	out_dir	Store generated plots (PDF files) and plot logs (Rout files) in this directory. By default the files are created in the out_dir directory.
replot_only	'0'	If set to '1' the data extraction is skipped and the plots are regenerated based on the data previously extracted.
source_filter	''	Specify the sources or destinations to filter on. Metrics will only be plotted for the sources and destinations on the list. See [1] for how to specify the list
stime	'0.0'	Start time for plot window (x-axis starts at this time).
test_id	''	Specifies the test ID of the experiment to be analysed. Must be specified. Specifying multiple test IDs with semicolons as separators allows to plot the results of multiple experiments in the same graph.
ts_correct	'0'	If set to '1' plot data with timestamps corrected based on estimated clock offsets. This requires to have a clock offset data file that can be created with the get_clock_offsets task. If set to '0' (default) plot data based on timestamps as they are in the log files.
udp_map	''	This parameter allows to specify a map that defines how to combine unidirectional UDP flows, as SPP needs bidirectional flows. The format is: <ip1>:<port1>:<ip2>:<port2>[;<ip3>:<port3>:<ip4>:<port4>] Each entry specifies the two sources (in terms of IP address and port) that are then linked to each other and treated as a bidirectional flow. This parameter is useful if UDP flows are not symmetric, i.e. the sending and receiving ports differ.
ymax	'0'	Maximum y-axis value. By default (if ymax set to '0') the maximum will be determined automatically. The parameter can be used to enforce a certain maximum, i.e. to generate different plots with the same scale.
ymin	'0'	Minimum y-axis value. The parameter can be used to enforce the specified minimum.

IX. ANALYSE_TCP_RTT

This task extracts the TCP RTT estimate and plots the estimate over time.

Parameter	Default Value	Explanation
etime	duration of experiment	End time for plot window (x-axis ends at this time).
io_filter	'o'	Specify whether TCP statistics are plotted based on incoming (set to 'i'), outgoing (set to 'o') or incoming and outgoing packets (set to 'io'). Only works with SIFTR logs (FreeBSD).

lnames	‘’	List of names for legend (separated by semicolons). The list must have the same length as the data series plotted.
min_values	‘3’	Only data series with more than min_values data values are plotted.
omit_const	‘0’	If set to ‘1’ any data series that are constant for the the duration of the experiment are not plotted.
out_dir	‘’	Extracted data files and plots are generated in this directory. By default the files are generated in the experiment directory.
out_name	‘’	A user-defined string that is used as prefix for the generated plot file, it can be used to describe the plot file.
pdf_dir	out_dir	Store generated plots (PDF files) and plot logs (Rout files) in this directory. By default the files are created in the out_dir directory.
replot_only	‘0’	If set to ‘1’ the data extraction is skipped and the plots are regenerated based on the data previously extracted.
smoothed	‘1’	If set to ‘1’ (default) smoothed TCP RTTs are plotted. If set to ‘0’ unsmoothed TCP RTT estimates are plotted and for SIFTR data the ERTT [4] estimates are plotted.
source_filter	‘’	Specify the sources or destinations to filter on. Metrics will only be plotted for the sources and destinations on the list. See [1] for how to specify the list
stime	‘0.0’	Start time for plot window (x-axis starts at this time).
test_id	‘’	Specifies the test ID of the experiment to be analysed. Must be specified. Specifying multiple test IDs with semicolons as separators allows to plot the results of multiple experiments in the same graph.
ts_correct	‘0’	If set to ‘1’ plot data with timestamps corrected based on estimated clock offsets. This requires to have a clock offset data file that can be created with the get_clock_offsets task. If set to ‘0’ (default) plot data based on timestamps as they are in the log files.
web10g_version	‘2.0.9’	Explicitely specific web10g version, since the log file format depends on the version. By default TEACUP tries to guess the version. Note, that the Windows estats logger produces output equivalent to version 2.0.7.
ymax	‘0’	Maximum y-axis value. By default (if ymax set to ‘0’) the maximum will be determined automatically. The parameter can be used to enforce a certain maximum, i.e. to generate different plots with the same scale.
ymin	‘0’	Minimum y-axis value. The parameter can be used to enforce the specified minimum.

X. ANALYSE_TCP_STAT

This task allows to extract an arbitrary TCP statistic and plot that statistic over time.

Parameter	Default Value	Explanation
etime	duration of experiment	End time for plot window (x-axis ends at this time).

io_filter	‘o’	Specify whether TCP statistics are plotted based on incoming (set to ‘i’), outgoing (set to ‘o’) or incoming and outgoing packets (set to ‘io’). Only takes affect for SIFTR tcp logs (FreeBSD).
lnames	‘’	List of names for legend (separated by semicolons). The list must have the same length as the data series plotted.
min_values	‘3’	Only data series with more than min_values data values are plotted.
omit_const	‘0’	If set to ‘1’ any data series that are constant for the the duration of the experiment are not plotted.
out_dir	‘’	Extracted data files and plots are generated in this directory. By default the files are generated in the experiment directory.
out_name	‘’	A user-defined string that is used as prefix for the generated plot file, it can be used to describe the plot file.
pdf_dir	out_dir	Store generated plots (PDF files) and plot logs (Rout files) in this directory. By default the files are created in the out_dir directory.
replot_only	‘0’	If set to ‘1’ the data extraction is skipped and the plots are regenerated based on the data previously extracted.
siftr_index	‘9’	Index (column number starting with 1) of the statistic in SIFTR log files.
source_filter	‘’	Specify the sources or destinations to filter on. Metrics will only be plotted for the sources and destinations on the list. See [1] for how to specify the list
stime	‘0.0’	Start time for plot window (x-axis starts at this time).
test_id	‘’	Specifies the test ID of the experiment to be analysed. Must be specified. Specifying multiple test IDs with semicolons as separators allows to plot the results of multiple experiments in the same graph.
ts_correct	‘0’	If set to ‘1’ plot data with timestamps corrected based on estimated clock offsets. This requires to have a clock offset data file that can be created with the get_clock_offsets task. If set to ‘0’ (default) plot data based on timestamps as they are in the log files.
web10g_index		Index (column number starting with 1) of the statistic in web10g log files.
ylabel	‘’	Y-axis label for the graph.
ymax	‘0’	Maximum y-axis value. By default (if ymax set to ‘0’) the maximum will be determined automatically. The parameter can be used to enforce a certain maximum, i.e. to generate different plots with the same scale.
ymin	‘0’	Minimum y-axis value. The parameter can be used to enforce the specified minimum.
yscaler	‘1.0’	Scaling factor for the extracted values.

XI. ANALYSE_THROUGHPUT

This tasks extracts the packet sizes from the tcpdump files and plots throughput over time.

Parameter	Default Value	Explanation
etime	duration of experiment	End time for plot window (x-axis ends at this time).

link_len	‘0’	If set to ‘0’ throughput is calculated based on the length of the IP packets. If set to ‘1’ throughput is calculated based on the link-layer frame length.
lnames	‘’	List of names for legend (separated by semicolons). The list must have the same length as the data series plotted.
min_values	‘3’	Only data series with more than min_values data values are plotted.
omit_const	‘0’	If set to ‘1’ any data series that are constant for the the duration of the experiment are not plotted.
out_dir	‘’	Extracted data files and plots are generated in this directory. By default the files are generated in the experiment directory.
out_name	‘’	A user-defined string that is used as prefix for the generated plot file, it can be used to describe the plot file.
pdf_dir	out_dir	Store generated plots (PDF files) and plot logs (Rout files) in this directory. By default the files are created in the out_dir directory.
replot_only	‘0’	If set to ‘1’ the data extraction is skipped and the plots are regenerated based on the data previously extracted.
source_filter	‘’	Specify the sources or destinations to filter on. Metrics will only be plotted for the sources and destinations on the list. See [1] for how to specify the list
stime	‘0.0’	Start time for plot window (x-axis starts at this time).
test_id	‘’	Specifies the test ID of the experiment to be analysed. Must be specified. Specifying multiple test IDs with semicolons as separators allows to plot the results of multiple experiments in the same graph.
ts_correct	‘0’	If set to ‘1’ plot data with timestamps corrected based on estimated clock offsets. This requires to have a clock offset data file that can be created with the get_clock_offsets task. If set to ‘0’ (default) plot data based on timestamps as they are in the log files.
ymax	‘0’	Maximum y-axis value. By default (if ymax set to ‘0’) the maximum will be determined automatically. The parameter can be used to enforce a certain maximum, i.e. to generate different plots with the same scale.
ymin	‘0’	Minimum y-axis value. The parameter can be used to enforce the specified minimum.

XII. AUTHORIZE_KEY

This task can be used to append the current user’s public RSA key to the ~./ssh/authorized_keys file of the remote user. The user can then login via SSH without having to enter a password. The task has no parameters.

XIII. CHECK_CONFIG

This tasks performs a number of sanity checks for the given config.py file. It will terminate with an error message if there is an error in the config file. Otherwise, it will terminate with an OK message. The task has no parameters.

XIV. CHECK_CONNECTIVITY

This task checks the connectivity between each pair of hosts using ping. The task only checks connectivity on the test network, it does not check connectivity on the control network. The task has no parameters.

XV. CHECK_HOST

The task checks if all necessary tools are installed on a host. If a required tool is missing the task will terminate with an error. The task has no parameters.

XVI. COPY_FILE

The task will copy a file from the local file system to the remote host(s). If hosts are not explicitly specified, the task will copy the file to all hosts listed in the config.py file including the router (TPCONF_router plus TPCONF_hosts). The file will be copied as the user env.user, which must be specified in config.py (or on the command line).

Parameter	Default Value	Explanation
file_name	''	Name of the file on the local file system.
method	'put'	Method used for copying. Must be either 'put' to use Fabric's put or 'scp' to use the scp tool (assuming scp is installed).
remote_path	''	Path on the remote where the file shall be copied to.

XVII. EXEC_CMD

The task will execute a command on the remote host(s). If hosts are not explicitly specified, the task will copy the file to all hosts listed in the config.py file including the router (TPCONF_router plus TPCONF_hosts).

Parameter	Default Value	Explanation
cmd	''	Command to be executed. Will be passed to sh -c.

XVIII. GET_CLOCK_OFFSETS

This task will estimate the offsets between host clocks during an experiment. The task can only be used if broadcast/multicast ping traffic was enabled during the experiment (see [1]).

Parameter	Default Value	Explanation
baseline_host	Router specified in config.py	Specify the host's clock we use as reference/baseline clock. By default the host is the router specified in the config.py file (TPCONF_router).
exp_list	'experiments_completed.txt'	List of experiments for which to compute the clock offsets. This variable is ignored if test_id is specified.
out_dir	''	Clock offset estimates file is generated in this directory. By default the files are generated in the experiment directory.
pkt_filter	'icmp and dst host <addr>'	tcpdump filter that specifies the packets used for the clock offset calculations. By default we select the ICMP packets send to the broadcast or multicast address <addr> specified in the config.py file (TPCONF_bc_ping_address).
test_id	''	List of test IDs for which to compute the clock offsets.

XIX. GET_NETINT

This task will return the network interface name(s).

Parameter	Default Value	Explanation

int_no	‘0’	The interface number starting with 0.
windump	‘0’	On Windows there are two names: 1) the name Windows uses and 2) the name Windump uses. If this parameter is set to ‘0’ the task will return the Windows name, if this parameter is set to ‘1’ it will return the Windump name.

XX. GET_NETMAC

This task returns the MAC address of a host’s control network or experimental network network interface.

Parameter	Default Value	Explanation
internal_int	‘0’	If set to ‘0’ the MAC for the testbed interface is returned (works only for hosts but not the router). If set to ‘1’ the MAC address of the control interface is returned.

XXI. GET_TYPE

This task returns the type of host(s), e.g. ‘Linux’, ‘FreeBSD’ or ‘CYGWIN’. The task has no parameters.

XXII. INIT_CC_ALGO

This task configures the congestion control for a host.

Parameter	Default Value	Explanation
algo	‘default’	The name of the algorithm. Currently, this can be ‘newreno’, ‘cubic’, ‘cdg’, ‘htcp’, ‘vegas’ on FreeBSD or Linux and ‘compound’ on Windows.

XXIII. INIT_ECN

This task enables or disables explicit congestion notification (ECN) for a host.

Parameter	Default Value	Explanation
ecn	‘0’	If set to ‘0’ ECN is disabled, if set to ‘1’ ECN is enabled.

XXIV. INIT_HOST

This tasks performs basic initialisation for a host (other than the router), including disabling the TCP host cache and disabling various NIC offloading mechanisms, such as TCP segmentation offloading (TSO). The task has no parameters.

XXV. INIT_HOST_CUSTOM

This task executes custom initialisation commands on hosts based on the config.py settings. The task has no parameters.

XXVI. INIT_OS

This tasks initialises the OS on host(s), i.e. it reboots hosts into the desired OS.

Parameter	Default Value	Explanation
boot_timeout	‘100’	Number of seconds to wait for host to reboot. After this timeout give up or power cycle host if do_power_cycle is set to ‘1’
do_power_cycle	‘0’	If set to ‘0’ do not power cycle. If set to ‘1’ power cycle host if it does not come up after boot_timeout seconds.
file_prefix	‘’	Prefix for generated PXE configuration file.
force_reboot	‘0’	If set to ‘0’ host is not rebooted if the current OS equals the desired OS. If set to ‘1’ the host is always rebooted.
os_list	‘’	Comma-separated list of OS names (‘Linux’, ‘FreeBSD’ or ‘CYGWIN’), one name for each host to reboot. The order must be the same as the order of the hosts specified with the fab -H command line parameter. If the number of entries in the list is smaller than the number of hosts it will be padded to the same length by duplicating the last entry. This allows to specify a single name for booting many hosts into the same OS.

XXVII. INIT_PIPE

This task configures a pipe on the router.

Parameter	Default Value	Explanation
attach_to_queue	‘’	This parameter works on Linux only! It allows to direct matching packets into an existing queue referenced by the specified queue ID (counter), but to emulate flow-specific delay/loss (different from the delay and loss of other traffic going through the same queue). If attach_to_queue is specified, the matching traffic will go through the already existing queue but the emulated delay/loss is set by the current init_pipe.
bidir	‘0’	If set to ‘0’ the pipe is unidirectional (packets going from source to dest only). If set to ‘1’ the pipe is bidirectional (packets going from source to dest and packets going from dest to source). Note that in the bidirectional case there are completely different buffers in both directions.
counter	‘1’	Unique ID of pipe/queue (must be an integer).
delay	‘’	Emulated delay in milliseconds. By default if empty string, the emulated delay is zero.
dest		Destination IP or destination network (<ip>[/<prefix>]). Must be specified.
loss	‘’	Emulated packet loss rate. By default if empty string, the emulated loss rate is zero.
queue_disc	‘’	The queuing discipline / AQM mechanism used. This can be the same of any of the queuing disciplines supported by Linux, such as ‘fq_codel’, ‘codel’, ‘red’, ‘choke’, ‘pfifo’, ‘pie’ etc. On FreeBSD the only queuing disciplines available are ‘fifo’ and ‘red’. For compatibility, with FreeBSD one can specify ‘fifo’ on Linux, which is mapped to ‘pfifo’ (‘pfifo’ is the default for HTB classes, which we use for rate limiting). Must be specified explicitly.

queue_disc_params	‘’	String of AQM parameters passed unchanged to Linux tc or FreeBSD Dummynet.
queue_size	‘’	Queue size in packets or bytes (depending on AQM used). Can be set to ‘bdp’ which will set the size according to the nominal BDP. If ‘bdp’ is specified and queue_size must be in packets, then the BDP size is configured based on the assumption that the average packet length is 600 bytes. The default depends on the operating system and possible also on the queuing discipline; hence this should be explicitly specified.
queue_size_mult	‘1.0’	A multiplier for queue size. This should only be used if queue_size if set to ‘bdp’. This allows to vary the queue size in multiples of the nominal BDP.
rate	‘’	Rate limit. Must be specified in bytes or with unit specifiers allowd by Linux tc or FreeBSD Dummynet. For example, Linux tc allows to specify ‘kbit’ or ‘mbit’.
rtt	‘’	Emulated RTT in milliseconds. This parameter only needs to be specified if queue_size is set to ‘bdp’ and the RTT is not 2·delay of the current pipe (e.g. if we set up asymmetric delay with attach_to_queue).
source	‘’	Source IP or source network (<ip>[/<prefix>]). Must be specified.

XXVIII. INIT_ROUTER

This task performs basic initialisation of the router, for example it sets up the root for queuing disciplines and disables NIC offloading mechanisms on Linux. The task has no parameters.

XXIX. KILL_OLD_PROCESSES

This tasks kills any possible old processes on the host(s). The task has no parameters.

XXX. LOG_QUEUE_STATS

This task log the queue statistics from the router.

Parameter	Default Value	Explanation
file_prefix	‘’	Prefix for generated log files.
local_dir	‘.’	Local directory where log files are stored.
remote_dir	‘’	Directory on the remote host where the log files are initially created and stored before they are copied and removed. By default if remote_dir is empty, the log files are created in the home directory of the user (env.user).

XXXI. LOG_SYSDATA

This task logs various information from the host(s), such as the output of uname, the list of currently running processes, the list of all sysctl variables. The logged information is described in more detail in [1].

Parameter	Default Value	Explanation

file_prefix	''	Prefix for generated log files.
local_dir	'.'	Local directory where log files are stored.
remote_dir	''	Directory on the remote host where the log files are initially created and stored before they are copied and removed. By default if remote_dir is empty, the log files are created in the home directory of the user (env.user).

XXXII. POWER_CYCLE

This task power cycles the host(s). It requires a config.py that specifies the power controller(s) for the host(s). The task has no parameters.

XXXIII. RUN_EXPERIMENT_MULTIPLE

This task runs a series of experiments based on the parameters to vary (specified in config.py).

Parameter	Default Value	Explanation
test_id	''	The test ID prefix used.
resume	'0'	If set to '0' do all experiments. If set to '1' do not repeat experiments that have been completed already according to experiments_completed.txt.

XXXIV. RUN_EXPERIMENT_SINGLE

This task runs a single experiment with the default parameters inf config.py.

Parameter	Default Value	Explanation
test_id	''	The test ID prefix used.

XXXV. SANITY_CHECKS

This tasks executes the check_host, check_connectivity and kill_old_processes tasks for the host(s). The task has no parameters.

XXXVI. SHOW_PIPES

This task shows the current pipe setup and statistics on the router. The task has no parameters.

ACKNOWLEDGEMENTS

TEACUP v0.8 was developed as part of a project funded by Cisco Systems and titled "Study in TCP Congestion Control Performance In A Data Centre". This is a collaborative effort between CAIA and Fred Baker of Cisco Systems.

REFERENCES

[1] S. Zander, G. Armitage, "TEACUP v0.8 – A System for Automated TCP Testbed Experiments," Centre for Advanced Internet Architectures, Swinburne University of Technology, Tech. Rep. 150210A, 2015. [Online]. Available: http://caia.swin.edu.au/reports/150210A/CAIA-TR-150210A.pdf

[2] S. Zander, "TEACUP v0.4 - Command Reference," Centre for Advanced Internet Architectures, Swinburne University of Technology, Tech. Rep. 140314C, 2014. [Online]. Available: http://caia.swin.edu.au/reports/140314C/CAIA-TR-140314C.pdf

[3] ——, "TEACUP v0.6 - Command Reference," Centre for Advanced Internet Architectures, Swinburne University of Technology, Tech. Rep. 140918B, 2014. [Online]. Available: http://caia.swin.edu.au/reports/140918B/CAIA-TR-140918B.pdf

[4] D. Hayes, "Timing enhancements to the FreeBSD kernel to support delay and rate based TCP mechanisms," Centre for Advanced Internet Architectures, Swinburne University of Technology, Melbourne, Australia, Tech. Rep. 100219A, 19 February 2010. [Online]. Available: http://caia.swin.edu.au/reports/100219A/CAIA-TR-100219A.pdf

[5] S. Zander and G. Armitage, "Minimally-Intrusive Frequent Round Trip Time Measurements Using Synthetic Packet-Pairs," in *The 38th IEEE Conference on Local Computer Networks (LCN 2013)*, 21-24 October 2013.

[6] A. Heyde, "SPP Implementation," August 2013. [Online]. Available: http://caia.swin.edu.au/tools/spp/downloads.html