Package: future.batchtools (via r-universe)

November 23, 2026

Version 0.32.0-9000

Depends R (>= 3.2.0), future (>= 1.80.0)

Imports parallelly, batchtools (>= 0.9.17), utils, checkmate, stringi

Suggests future.apply, listenv, snow

Title Parallel Future Backend for HPC Schedulers using 'batchtools'

Description Implementation of the Future API <doi:10.32614/RJ-2021-048> on top of the 'batchtools' package. The backend is fully compliant with the Future API, as validated by the 'future.tests' conformance test suite, meaning that code using futures gives identical results and behaves the same regardless of which future backend is used. Futures are resolved on the local machine, on an ad-hoc cluster of machines, or via job schedulers on high-performance compute ('HPC') clusters, such as 'Slurm', 'SGE', 'LSF', 'OpenLava', and 'TORQUE' / 'PBS'. A common use is to parallelize existing map-reduce calls via the 'futurize' package <doi:10.48550/arXiv.2601.17578>, e.g. 'plan(future.batchtools::batchtools_slurm)' followed by 'lapply(xs, fcn) |> futurize()' or 'purrr::map(xs, fcn) |> futurize()'.

License LGPL (>= 3)

Copyright file inst/COPYRIGHTS

LazyLoad TRUE

URL https://future.batchtools.futureverse.org, https://github.com/futureverse/future.batchtools

BugReports https://github.com/futureverse/future.batchtools/issues

Language en-US

Encoding UTF-8

Roxygen list(markdown = TRUE)

Config/roxygen2/version 8.1.1

Config/pak/sysreqs cmake make libicu-dev libuv1-dev

Repository https://futureverse-devel.r-universe.dev

Date/Publication 2026-11-22 23:58:36 UTC

RemoteUrl https://github.com/futureverse/future.batchtools

RemoteRef develop

RemoteSha 331257a683885beaaa0c2b99d0fdaa1eb7cbc806

Contents

batchtools_bash 2
batchtools_interactive 7
batchtools_local 8
batchtools_lsf 10
batchtools_multicore 16
batchtools_openlava 17
batchtools_sge 23
batchtools_slurm 29
batchtools_torque 36
find_template_file 42
future.batchtools 43
makeClusterFunctionsSlurm2 44
render_template 45
zzz-future.batchtools.options 46

Index **48**

`batchtools_bash` *A batchtools Bash backend that resolves futures sequentially via a Bash template script*

Description

The `batchtools_bash` backend was added to illustrate how to write a custom **future.batchtools** backend that uses a templated job script. Please see the source code, for details.

Usage

```
batchtools_bash(
  ...,
  template = "bash",
  fs.latency = 0,
  resources = list(),
  delete = getOption("future.batchtools.delete", "on-success"),
  registry = list(),
  finalize = getOption("future.finalize", TRUE)
)
```

```
makeClusterFunctionsBash(template = "bash", fs.latency = 0, ...)
```

Arguments

`template`	(optional) Name of job-script template to be searched for by `batchtools::findTemplateFile()`. If not found, it defaults to the `templates/bash.tmpl` part of this package (see below).
`fs.latency`	[`numeric(1)`] Expected maximum latency of the file system, in seconds. Set to a positive number for network file systems like NFS which enables more robust (but also more expensive) mechanisms to access files and directories. Usually safe to set to `0` to disable the heuristic, e.g. if you are working on a local file system.
`resources`	(optional) A named list passed to the **batchtools** job-script template as variable `resources`. This is based on how `batchtools::submitJobs()` works, with the exception for specially reserved names defined by the **future.batchtools** package; • `resources[["details"]]`, if TRUE, results in the job script outputting job details and job summaries at the beginning and at the end. It also reports on the number of packages installed in each R package library, which can be slow for large libraries on slow file systems. • `resources[["startup"]]` and `resources[["shutdown"]]` are character vectors of shell code to be injected to the job script as-is. • `resources[["modules"]]` is character vector of Linux environment modules to be loaded. • `resources[["envs"]]`, is an optional names character vector specifying environment variables to be set. • `resources[["rscript"]]` is an optional character vector specifying how the 'Rscript' is launched. The `resources[["rscript_args"]]` field is an optional character vector specifying the 'Rscript' command-line arguments. • `resources[["asis"]]` is a character vector that are passed as-is to the job script and are injected as job resource declarations. • All remaining `resources` named elements are injected as named resource specification for the scheduler.
`delete`	Controls if and when the batchtools job registry folder is deleted. If `"on-success"` (default), it is deleted if the future was resolved successfully *and* the expression did not produce an error. If `"never"`, then it is never deleted. If `"always"`, then it is always deleted.
`registry`	(optional) A named list of settings applied to each future's **batchtools** Registry. This is a more convenient alternative to using argument `conf.file`.
`finalize`	If TRUE, a future's **batchtools** Registry is automatically deleted when the future is garbage collected, otherwise not.
`...`	Not used.

Details

Batchtools Bash futures use **batchtools** cluster functions created by `makeClusterFunctionsBash()` and requires that `bash` is installed on the current machine and the `timeout` command is available.

The default template script `templates/bash.tmpl` can be found in:

```
system.file("templates", "bash.tmpl", package = "future.batchtools")
```

and comprises:

```
#!/bin/bash
######################################################################
# A batchtools launch script template
#
# Author: Henrik Bengtsson
######################################################################

## Bash settings
set -e          # exit on error
set -u          # error on unset variables
set -o pipefail # fail a pipeline if any command fails
trap 'echo "ERROR: future.batchtools job script failed on line $LINENO" >&2; exit 1' ERR

## Redirect stdout and stderr to the batchtools log file
exec > <%= log.file %> 2>&1

<%
  ## Work on a copy of 'resources', because this template is evaluated
  ## in the environment of the batchtools JobCollection, and modifying
  ## 'resources' would modify the resources of the job itself
  job_resources <- resources

  ## Should job "details" be seen?
  details <- isTRUE(job_resources[["details"]])
  job_resources[["details"]] <- NULL

  ## Shell "startup" code to evaluate
  startup <- job_resources[["startup"]]
  job_resources[["startup"]] <- NULL

  ## Shell "shutdown" code to evaluate
  shutdown <- job_resources[["shutdown"]]
  job_resources[["shutdown"]] <- NULL

  ## Environment modules specifications
  modules <- job_resources[["modules"]]
  job_resources[["modules"]] <- NULL
```

```
  ## Environment variables to be set
  envs <- job_resources[["envs"]]
  if (length(envs) > 0) {
    stopifnot(is.character(envs), !is.null(names(envs)))
  }
  job_resources[["envs"]] <- NULL

  ## Custom "Rscript" command and Rscript arguments
  rscript <- job_resources[["rscript"]]
  if (is.null(rscript)) {
    rscript <- "Rscript"
  } else if (length(rscript) == 0 || !nzchar(rscript)[1]) {
    stop("Argument 'resources' specifies an empty 'rscript' field")
  }
  job_resources[["rscript"]] <- NULL

  ## Maximum runtime?
  timeout <- job_resources[["timeout"]]
  if (is.null(timeout)) timeout <- job_resources[["runtime"]]
  job_resources[["timeout"]] <- NULL
  job_resources[["runtime"]] <- NULL
  if (length(timeout) > 0) {
      rscript <- c("timeout", timeout, rscript)
  }
  rscript_args <- job_resources[["rscript_args"]]
  job_resources[["rscript_args"]] <- NULL
  rscript_call <- paste(c(rscript, rscript_args), collapse = " ")
%>

<% if (length(startup) > 0) {
  writeLines(startup)
} %>

<% if (length(modules) > 0) {
  writeLines(c(
    'echo "Load environment modules:"',
    sprintf('echo "- modules: %s"', paste(modules, collapse = ", ")),
    sprintf("module load %s", paste(modules, collapse = " ")),
    "module list"
  ))
} %>

<% if (length(envs) > 0) {
  writeLines(c(
    sprintf("echo 'Setting environment variables: [n=%d]'", length(envs)),
    sprintf("echo ' - %s=%s'", names(envs), shQuote(envs)),
    sprintf("export %s=%s", names(envs), shQuote(envs))
  ))
```

```
} %>

echo "Session information:"
echo "- timestamp: $(date +"%Y-%m-%d %H:%M:%S%z")"
echo "- hostname: $(hostname)"
echo "- Rscript: <%= paste(rscript, collapse = " ") %>"
echo "- Rscript args: <%= paste(rscript_args, collapse = " ") %>"
echo "- Rscript call: <%= rscript_call %>"
if ! command -v <%= rscript[1] %> &> /dev/null; then
   >&2 echo "ERROR: Argument 'resources' specifies a non-existing 'Rscript' launch command: <%= rscript[
    exit 1
fi
echo "- Rscript version: $(<%= paste(rscript, collapse = " ") %> --version)"
echo "- R_LIBS_USER=${R_LIBS_USER:-<not set>}"
echo "- R_LIBS_SITE=${R_LIBS_SITE:-<not set>}"
echo "- R_LIBS=${R_LIBS:-<not set>}"
echo "- Rscript library paths:"
<% if (details) { %>
<%= rscript_call %> -e "for (p in .libPaths()) cat('  - ', shQuote(p), ' [n=', sum(file.exists(file.path
<% } else { %>
<%= rscript_call %> -e "cat(sprintf('  - %s\n', shQuote(.libPaths())), sep = '')"
<% } %>
echo

# Launch R and evaluate the batchtools R job
echo "Calling 'batchtools::doJobCollection()' ..."
echo "- job name: '<%= job.name %>'"
echo "- job log file: '<%= log.file %>'"
echo "- job uri: '<%= uri %>'"
<%= rscript_call %> -e 'batchtools::doJobCollection("<%= uri %>")'
res=$?
echo " - exit code: ${res}"
echo "Calling 'batchtools::doJobCollection()' ... done"
echo

<% if (length(shutdown) > 0) {
  writeLines(shutdown)
} %>

echo "End time: $(date +"%Y-%m-%d %H:%M:%S%z")"

## Relay the exit code from Rscript
exit "${res}"
```

Value

Nothing.

`makeClusterFunctionsBash()` returns a ClusterFunctions object.

Examples

```
library(future)

# Limit runtime to 30 seconds per future
plan(future.batchtools::batchtools_bash, resources = list(runtime = 30))

message("Main process ID: ", Sys.getpid())

f <- future({
  data.frame(
    hostname = Sys.info()[["nodename"]],
          os = Sys.info()[["sysname"]],
       cores = unname(parallelly::availableCores()),
         pid = Sys.getpid(),
     modules = Sys.getenv("LOADEDMODULES")
  )
})
info <- value(f)
print(info)
```

---

batchtools_interactive

A batchtools backend that resolves futures sequentially in the current R session

---

Description

The batchtools interactive backend is useful for verifying parts of your **batchtools** setup locally, while still being able to do interactive debugging.

Usage

```
batchtools_interactive(
  ...,
  fs.latency = 0,
  delete = getOption("future.batchtools.delete", "on-success"),
  registry = list(),
  finalize = getOption("future.finalize", TRUE)
)
```

Arguments

`fs.latency`	`[numeric(1)]` Expected maximum latency of the file system, in seconds. Set to a positive number for network file systems like NFS which enables more robust (but also more expensive) mechanisms to access files and directories. Usually safe to set to `0` to disable the heuristic, e.g. if you are working on a local file system.

`delete`	Controls if and when the batchtools job registry folder is deleted. If `"on-success"` (default), it is deleted if the future was resolved successfully *and* the expression did not produce an error. If `"never"`, then it is never deleted. If `"always"`, then it is always deleted.
`registry`	(optional) A named list of settings applied to each future's **batchtools** Registry. This is a more convenient alternative to using argument `conf.file`.
`finalize`	If TRUE, a future's **batchtools** Registry is automatically deleted when the future is garbage collected, otherwise not.
`...`	Not used.

Details

Batchtools interactive futures use **batchtools** cluster functions created by `batchtools::makeClusterFunctionsInteracti` with `external = FALSE`.

An alternative to the batchtools interactive backend is to use `plan(future::sequential)`, which is a faster way process futures sequentially and that also can be debugged interactively.

Value

Nothing.

Examples

```
library(future)
plan(future.batchtools::batchtools_interactive)

message("Main process ID: ", Sys.getpid())

f <- future({
  data.frame(
    hostname = Sys.info()[["nodename"]],
          os = Sys.info()[["sysname"]],
       cores = unname(parallelly::availableCores()),
         pid = Sys.getpid(),
     modules = Sys.getenv("LOADEDMODULES")
  )
})
info <- value(f)
print(info)
```

`batchtools_local` *A batchtools backend that resolves futures sequentially in transient background R sessions*

Description

The batchtools local backend is useful for verifying parts of your **batchtools** setup locally, before using a more advanced backend such as the job-scheduler backends.

Usage

```
batchtools_local(
  ...,
  fs.latency = 0,
  delete = getOption("future.batchtools.delete", "on-success"),
  registry = list(),
  finalize = getOption("future.finalize", TRUE)
)
```

Arguments

`fs.latency`	[`numeric(1)`] Expected maximum latency of the file system, in seconds. Set to a positive number for network file systems like NFS which enables more robust (but also more expensive) mechanisms to access files and directories. Usually safe to set to `0` to disable the heuristic, e.g. if you are working on a local file system.
`delete`	Controls if and when the batchtools job registry folder is deleted. If `"on-success"` (default), it is deleted if the future was resolved successfully *and* the expression did not produce an error. If `"never"`, then it is never deleted. If `"always"`, then it is always deleted.
`registry`	(optional) A named list of settings applied to each future's **batchtools** Registry. This is a more convenient alternative to using argument `conf.file`.
`finalize`	If TRUE, a future's **batchtools** Registry is automatically deleted when the future is garbage collected, otherwise not.
`...`	Not used.

Details

Batchtools local futures use **batchtools** cluster functions created by `batchtools::makeClusterFunctionsInteractive()` with `external = TRUE`.

The batchtools local backend is synchronous, that is, the main R process blocks while a job is running. This means that nothing can be relayed from a job until it is done. For example, immediate conditions, such as progress updates signaled by the **progressr** package, are relayed first when the job has finished. To get them as soon as possible, use an asynchronous backend, e.g. `batchtools_multicore()` with two or more workers, or a job scheduler.

An alternative to the batchtools local backend is to use `plan(future::cluster, workers = I(1))`.

Value

Nothing.

Examples

```
library(future)
plan(future.batchtools::batchtools_local)

message("Main process ID: ", Sys.getpid())
```

```
f <- future({
  data.frame(
    hostname = Sys.info()[["nodename"]],
          os = Sys.info()[["sysname"]],
       cores = unname(parallelly::availableCores()),
         pid = Sys.getpid(),
     modules = Sys.getenv("LOADEDMODULES")
  )
})
info <- value(f)
print(info)
```

`batchtools_lsf` *A batchtools LSF backend resolves futures in parallel via a Load Sharing Facility (LSF) job scheduler*

Description

A batchtools LSF backend resolves futures in parallel via a Load Sharing Facility (LSF) job scheduler

Usage

```
batchtools_lsf(
  ...,
  template = "lsf",
  scheduler.latency = 1,
  fs.latency = 65,
  resources = list(),
  delete = getOption("future.batchtools.delete", "on-success"),
  workers = getOption("future.batchtools.workers", default = 100L),
  registry = list(),
  finalize = getOption("future.finalize", TRUE)
)
```

Arguments

`template` (optional) Name of job-script template to be searched for by `batchtools::findTemplateFile()`. If not found, it defaults to the `templates/lsf.tmpl` part of this package (see below).

`scheduler.latency` [`numeric(1)`]
Time to sleep after important interactions with the scheduler to ensure a sane state. Currently only triggered after calling `submitJobs`.

`fs.latency` [numeric(1)]
Expected maximum latency of the file system, in seconds. Set to a positive number for network file systems like NFS which enables more robust (but also more expensive) mechanisms to access files and directories. Usually safe to set to `0` to disable the heuristic, e.g. if you are working on a local file system.

`resources` (optional) A named list passed to the **batchtools** job-script template as variable `resources`. This is based on how `batchtools::submitJobs()` works, with the exception for specially reserved names defined by the **future.batchtools** package;

- `resources[["details"]]`, if TRUE, results in the job script outputting job details and job summaries at the beginning and at the end. It also reports on the number of packages installed in each R package library, which can be slow for large libraries on slow file systems.
- `resources[["startup"]]` and `resources[["shutdown"]]` are character vectors of shell code to be injected to the job script as-is.
- `resources[["modules"]]` is character vector of Linux environment modules to be loaded.
- `resources[["envs"]]`, is an optional names character vector specifying environment variables to be set.
- `resources[["rscript"]]` is an optional character vector specifying how the 'Rscript' is launched. The `resources[["rscript_args"]]` field is an optional character vector specifying the 'Rscript' command-line arguments.
- `resources[["asis"]]` is a character vector that are passed as-is to the job script and are injected as job resource declarations.
- All remaining `resources` named elements are injected as named resource specification for the scheduler.

`delete` Controls if and when the batchtools job registry folder is deleted. If `"on-success"` (default), it is deleted if the future was resolved successfully *and* the expression did not produce an error. If `"never"`, then it is never deleted. If `"always"`, then it is always deleted.

`workers` The maximum number of workers the batchtools backend may use at any time, which for HPC schedulers corresponds to the maximum number of queued jobs. The default is `getOption("future.batchtools.workers", 100)`.

`registry` (optional) A named list of settings applied to each future's **batchtools** Registry. This is a more convenient alternative to using argument `conf.file`.

`finalize` If TRUE, a future's **batchtools** Registry is automatically deleted when the future is garbage collected, otherwise not.

`...` Not used.

Details

Batchtools Load Sharing Facility (LSF) futures use **batchtools** cluster functions created by `batchtools::makeClusterFunc` which are used to interact with the LSF job scheduler. This requires that LSF commands `bsub`, `bjobs`, and `bkill` are available on the current machine.

The default template script `templates/lsf.tmpl` can be found in:

```
system.file("templates", "lsf.tmpl", package = "future.batchtools")
```

and comprises:

```
#!/bin/bash
##############################################################################
# A batchtools launch script template for LSF and OpenLava
#
# Author: Henrik Bengtsson
##############################################################################

## Job name
#BSUB -J <%= job.name %>

## Direct streams to logfile
#BSUB -o <%= log.file %>

## Resources needed
<%
  ## Work on a copy of 'resources', because this template is evaluated
  ## in the environment of the batchtools JobCollection, and modifying
  ## 'resources' would modify the resources of the job itself
  job_resources <- resources

  ## Should scheduler "details" be seen?
  details <- isTRUE(job_resources[["details"]])
  job_resources[["details"]] <- NULL

  ## Shell "startup" code to evaluate
  startup <- job_resources[["startup"]]
  job_resources[["startup"]] <- NULL

  ## Shell "shutdown" code to evaluate
  shutdown <- job_resources[["shutdown"]]
  job_resources[["shutdown"]] <- NULL

  ## Environment modules specifications
  modules <- job_resources[["modules"]]
  job_resources[["modules"]] <- NULL

  ## Environment variables to be set
  envs <- job_resources[["envs"]]
  if (length(envs) > 0) {
    stopifnot(is.character(envs), !is.null(names(envs)))
  }
  job_resources[["envs"]] <- NULL

  ## Custom "Rscript" command and Rscript arguments
  rscript <- job_resources[["rscript"]]
```

```
  if (is.null(rscript)) {
    rscript <- "Rscript"
  } else if (length(rscript) == 0 || !nzchar(rscript)[1]) {
    stop("Argument 'resources' specifies an empty 'rscript' field")
  }
  job_resources[["rscript"]] <- NULL
  rscript_args <- job_resources[["rscript_args"]]
  job_resources[["rscript_args"]] <- NULL
  rscript_call <- paste(c(rscript, rscript_args), collapse = " ")

  ## As-is resource specifications
  job_declarations <- job_resources[["asis"]]
  job_resources[["asis"]] <- NULL

  ## Remaining resources are assumed to be of type '<key>=<value>'
  opts <- unlist(job_resources, use.names = TRUE)
  opts <- sprintf("%s=%s", names(opts), opts)
 job_declarations <- sprintf("#BSUB %s", c(job_declarations, sprintf("-%s", opts)))
  writeLines(job_declarations)
%>

## Bash settings
set -e          # exit on error
set -u          # error on unset variables
set -o pipefail # fail a pipeline if any command fails
trap 'echo "ERROR: future.batchtools job script failed on line $LINENO" >&2; exit 1' ERR

<% if (length(job_declarations) > 0) {
  writeLines(c(
    "echo 'Job submission declarations:'",
    sprintf("echo '%s'", job_declarations),
    "echo"
  ))
} %>

<% if (details) { %>
if command -v bjobs > /dev/null; then
  echo "Job information:"
  bjobs -l "${LSB_JOBID}"
  echo
fi
<% } %>

<% if (length(startup) > 0) {
  writeLines(startup)
} %>

<% if (length(modules) > 0) {
```

```
  writeLines(c(
    "echo 'Load environment modules:'",
    sprintf("echo '- modules: %s'", paste(modules, collapse = ", ")),
    sprintf("module load %s", paste(modules, collapse = " ")),
    "module list"
  ))
} %>

<% if (length(envs) > 0) {
  writeLines(c(
    sprintf("echo 'Setting environment variables: [n=%d]'", length(envs)),
    sprintf("echo ' - %s=%s'", names(envs), shQuote(envs)),
    sprintf("export %s=%s", names(envs), shQuote(envs))
  ))
} %>

echo "Session information:"
echo "- timestamp: $(date +"%Y-%m-%d %H:%M:%S%z")"
echo "- hostname: $(hostname)"
echo "- Rscript call: <%= rscript_call %>"
if ! command -v <%= rscript[1] %> &> /dev/null; then
   >&2 echo "ERROR: Argument 'resources' specifies a non-existing 'Rscript' launch command: <%= rscript[
    exit 1
fi
echo "- Rscript version: $(<%= paste(rscript, collapse = " ") %> --version)"
echo "- R_LIBS_USER=${R_LIBS_USER:-<not set>}"
echo "- R_LIBS_SITE=${R_LIBS_SITE:-<not set>}"
echo "- R_LIBS=${R_LIBS:-<not set>}"
echo "- Rscript library paths:"
<% if (details) { %>
<%= rscript_call %> -e "for (p in .libPaths()) cat('  - ', shQuote(p), ' [n=', sum(file.exists(file.path
<% } else { %>
<%= rscript_call %> -e "cat(sprintf('  - %s\n', shQuote(.libPaths())), sep = '')"
<% } %>
echo


## Launch R and evaluate the batchtools R job
echo "Calling 'batchtools::doJobCollection()' ..."
echo "- job name: '<%= job.name %>'"
echo "- job log file: '<%= log.file %>'"
echo "- job uri: '<%= uri %>'"
<%= rscript_call %> -e 'batchtools::doJobCollection("<%= uri %>")'
res=$?
echo " - exit code: ${res}"
echo "Calling 'batchtools::doJobCollection()' ... done"
echo
```

```
<% if (details) { %>
if command -v bjobs > /dev/null; then
  echo "Job summary:"
  bjobs -l "${LSB_JOBID}"
fi
<% } %>

<% if (length(shutdown) > 0) {
  writeLines(shutdown)
} %>

echo "End time: $(date +"%Y-%m-%d %H:%M:%S%z")"

## Relay the exit code from Rscript
exit "${res}"
```

Value

Nothing.

References

- https://en.wikipedia.org/wiki/IBM_Spectrum_LSF

Examples

```
library(future)

# Limit runtime to 10 minutes and total memory to 400 MiB per future,
# request a parallel environment with four slots on a single host.
# Submit to the 'freecycle' queue. Load environment modules 'r' and
# 'jags'. Report on job details at startup and at the end of the job.
plan(future.batchtools::batchtools_lsf, resources = list(
  W = "00:10:00", M = "400",
  asis = c("-n 4", "-R 'span[hosts=1]'", "-q freecycle"),
  modules = c("r", "jags"),
  details = TRUE
))

f <- future({
  data.frame(
     hostname = Sys.info()[["nodename"]],
           os = Sys.info()[["sysname"]],
    osVersion = utils::osVersion,
        cores = unname(parallelly::availableCores()),
      modules = Sys.getenv("LOADEDMODULES")
  )
})
info <- value(f)
print(info)
```

batchtools_multicore *A batchtools backend that resolves futures in parallel via forked background R processes*

Description

A batchtools backend that resolves futures in parallel via forked background R processes

Usage

```
batchtools_multicore(
  ...,
  workers = availableCores(constraints = "multicore"),
  fs.latency = 0,
  delete = getOption("future.batchtools.delete", "on-success"),
  registry = list(),
  finalize = getOption("future.finalize", TRUE)
)
```

Arguments

workers	The number of multicore processes to be available for concurrent batchtools multicore futures. If a single worker is requested, which is the default when only one CPU core is available, the batchtools local backend is used instead. It is synchronous, that is, the main R process blocks while a job is running. To use an asynchronous batchtools multicore backend also with a single worker, protect the number of workers using base::I(), e.g. workers = I(1).
fs.latency	[numeric(1)] Expected maximum latency of the file system, in seconds. Set to a positive number for network file systems like NFS which enables more robust (but also more expensive) mechanisms to access files and directories. Usually safe to set to 0 to disable the heuristic, e.g. if you are working on a local file system.
delete	Controls if and when the batchtools job registry folder is deleted. If "on-success" (default), it is deleted if the future was resolved successfully *and* the expression did not produce an error. If "never", then it is never deleted. If "always", then it is always deleted.
registry	(optional) A named list of settings applied to each future's **batchtools** Registry. This is a more convenient alternative to using argument conf.file.
finalize	If TRUE, a future's **batchtools** Registry is automatically deleted when the future is garbage collected, otherwise not.
...	Not used.

Details

Batchtools multicore futures use **batchtools** cluster functions created by `batchtools::makeClusterFunctionsMulticore(` with `ncpus = workers`.

Since batchtools multicore futures are asynchronous, conditions, such as progress updates signaled by the **progressr** package, are relayed while the jobs are running. Note, however, that a single worker, unless protected using `base::I()`, falls back to the synchronous batchtools local backend, in which case they are relayed first when a job has finished.

An alternative to the batchtools multicore backend is to use `plan(future::multicore)`.

Value

Nothing.

Examples

```
library(future)
plan(future.batchtools::batchtools_multicore, workers = 2)

message("Main process ID: ", Sys.getpid())

f <- future({
  data.frame(
    hostname = Sys.info()[["nodename"]],
          os = Sys.info()[["sysname"]],
       cores = unname(parallelly::availableCores()),
         pid = Sys.getpid(),
     modules = Sys.getenv("LOADEDMODULES")
  )
})
info <- value(f)
print(info)
```

`batchtools_openlava` *A batchtools OpenLava backend resolves futures in parallel via an OpenLava job scheduler*

Description

A batchtools OpenLava backend resolves futures in parallel via an OpenLava job scheduler

Usage

```
batchtools_openlava(
  ...,
  template = "openlava",
  scheduler.latency = 1,
  fs.latency = 65,
```

```
  resources = list(),
  delete = getOption("future.batchtools.delete", "on-success"),
  workers = getOption("future.batchtools.workers", default = 100L),
  registry = list(),
  finalize = getOption("future.finalize", TRUE)
)
```

Arguments

`template` (optional) Name of job-script template to be searched for by `batchtools::findTemplateFile()`. If not found, it defaults to the `templates/openlava.tmpl` part of this package (see below).

`scheduler.latency`
[`numeric(1)`]
Time to sleep after important interactions with the scheduler to ensure a sane state. Currently only triggered after calling `submitJobs`.

`fs.latency` [`numeric(1)`]
Expected maximum latency of the file system, in seconds. Set to a positive number for network file systems like NFS which enables more robust (but also more expensive) mechanisms to access files and directories. Usually safe to set to `0` to disable the heuristic, e.g. if you are working on a local file system.

`resources` (optional) A named list passed to the **batchtools** job-script template as variable `resources`. This is based on how `batchtools::submitJobs()` works, with the exception for specially reserved names defined by the **future.batchtools** package;

- `resources[["details"]]`, if TRUE, results in the job script outputting job details and job summaries at the beginning and at the end. It also reports on the number of packages installed in each R package library, which can be slow for large libraries on slow file systems.
- `resources[["startup"]]` and `resources[["shutdown"]]` are character vectors of shell code to be injected to the job script as-is.
- `resources[["modules"]]` is character vector of Linux environment modules to be loaded.
- `resources[["envs"]]`, is an optional names character vector specifying environment variables to be set.
- `resources[["rscript"]]` is an optional character vector specifying how the 'Rscript' is launched. The `resources[["rscript_args"]]` field is an optional character vector specifying the 'Rscript' command-line arguments.
- `resources[["asis"]]` is a character vector that are passed as-is to the job script and are injected as job resource declarations.
- All remaining `resources` named elements are injected as named resource specification for the scheduler.

`delete` Controls if and when the batchtools job registry folder is deleted. If `"on-success"` (default), it is deleted if the future was resolved successfully *and* the expression did not produce an error. If `"never"`, then it is never deleted. If `"always"`, then it is always deleted.

workers	The maximum number of workers the batchtools backend may use at any time, which for HPC schedulers corresponds to the maximum number of queued jobs. The default is `getOption("future.batchtools.workers", 100)`.
registry	(optional) A named list of settings applied to each future's **batchtools** Registry. This is a more convenient alternative to using argument `conf.file`.
finalize	If TRUE, a future's **batchtools** Registry is automatically deleted when the future is garbage collected, otherwise not.
...	Not used.

Details

Batchtools OpenLava futures use **batchtools** cluster functions created by `batchtools::makeClusterFunctionsOpenLava()` which are used to interact with the OpenLava job scheduler. This requires that OpenLava commands `bsub`, `bjobs`, and `bkill` are available on the current machine.

The default template script `templates/openlava.tmpl` can be found in:

```
system.file("templates", "openlava.tmpl", package = "future.batchtools")
```

and comprises:

```
#!/bin/bash
##############################################################################
# A batchtools launch script template for LSF and OpenLava
#
# Author: Henrik Bengtsson
##############################################################################

## Job name
#BSUB -J <%= job.name %>

## Direct streams to logfile
#BSUB -o <%= log.file %>

## Resources needed
<%
  ## Work on a copy of 'resources', because this template is evaluated
  ## in the environment of the batchtools JobCollection, and modifying
  ## 'resources' would modify the resources of the job itself
  job_resources <- resources

  ## Should scheduler "details" be seen?
  details <- isTRUE(job_resources[["details"]])
  job_resources[["details"]] <- NULL

  ## Shell "startup" code to evaluate
  startup <- job_resources[["startup"]]
  job_resources[["startup"]] <- NULL
```

```
  ## Shell "shutdown" code to evaluate
  shutdown <- job_resources[["shutdown"]]
  job_resources[["shutdown"]] <- NULL

  ## Environment modules specifications
  modules <- job_resources[["modules"]]
  job_resources[["modules"]] <- NULL

  ## Environment variables to be set
  envs <- job_resources[["envs"]]
  if (length(envs) > 0) {
    stopifnot(is.character(envs), !is.null(names(envs)))
  }
  job_resources[["envs"]] <- NULL

  ## Custom "Rscript" command and Rscript arguments
  rscript <- job_resources[["rscript"]]
  if (is.null(rscript)) {
    rscript <- "Rscript"
  } else if (length(rscript) == 0 || !nzchar(rscript)[1]) {
    stop("Argument 'resources' specifies an empty 'rscript' field")
  }
  job_resources[["rscript"]] <- NULL
  rscript_args <- job_resources[["rscript_args"]]
  job_resources[["rscript_args"]] <- NULL
  rscript_call <- paste(c(rscript, rscript_args), collapse = " ")

  ## As-is resource specifications
  job_declarations <- job_resources[["asis"]]
  job_resources[["asis"]] <- NULL

  ## Remaining resources are assumed to be of type '<key>=<value>'
  opts <- unlist(job_resources, use.names = TRUE)
  opts <- sprintf("%s=%s", names(opts), opts)
 job_declarations <- sprintf("#BSUB %s", c(job_declarations, sprintf("-%s", opts)))
  writeLines(job_declarations)
%>

## Bash settings
set -e          # exit on error
set -u          # error on unset variables
set -o pipefail # fail a pipeline if any command fails
trap 'echo "ERROR: future.batchtools job script failed on line $LINENO" >&2; exit 1' ERR

<% if (length(job_declarations) > 0) {
  writeLines(c(
    "echo 'Job submission declarations:'",
```

```
    sprintf("echo '%s'", job_declarations),
    "echo"
  ))
} %>

<% if (details) { %>
if command -v bjobs > /dev/null; then
  echo "Job information:"
  bjobs -l "${LSB_JOBID}"
  echo
fi
<% } %>

<% if (length(startup) > 0) {
  writeLines(startup)
} %>

<% if (length(modules) > 0) {
  writeLines(c(
    "echo 'Load environment modules:'",
    sprintf("echo '- modules: %s'", paste(modules, collapse = ", ")),
    sprintf("module load %s", paste(modules, collapse = " ")),
    "module list"
  ))
} %>

<% if (length(envs) > 0) {
  writeLines(c(
    sprintf("echo 'Setting environment variables: [n=%d]'", length(envs)),
    sprintf("echo ' - %s=%s'", names(envs), shQuote(envs)),
    sprintf("export %s=%s", names(envs), shQuote(envs))
  ))
} %>

echo "Session information:"
echo "- timestamp: $(date +"%Y-%m-%d %H:%M:%S%z")"
echo "- hostname: $(hostname)"
echo "- Rscript call: <%= rscript_call %>"
if ! command -v <%= rscript[1] %> &> /dev/null; then
   >&2 echo "ERROR: Argument 'resources' specifies a non-existing 'Rscript' launch command: <%= rscript[
    exit 1
fi
echo "- Rscript version: $(<%= paste(rscript, collapse = " ") %> --version)"
echo "- R_LIBS_USER=${R_LIBS_USER:-<not set>}"
echo "- R_LIBS_SITE=${R_LIBS_SITE:-<not set>}"
echo "- R_LIBS=${R_LIBS:-<not set>}"
echo "- Rscript library paths:"
<% if (details) { %>
```

```
<%= rscript_call %> -e "for (p in .libPaths()) cat('  - ', shQuote(p), ' [n=', sum(file.exists(file.path
<% } else { %>
<%= rscript_call %> -e "cat(sprintf('  - %s\n', shQuote(.libPaths())), sep = '')"
<% } %>
echo


## Launch R and evaluate the batchtools R job
echo "Calling 'batchtools::doJobCollection()' ..."
echo "- job name: '<%= job.name %>'"
echo "- job log file: '<%= log.file %>'"
echo "- job uri: '<%= uri %>'"
<%= rscript_call %> -e 'batchtools::doJobCollection("<%= uri %>")'
res=$?
echo " - exit code: ${res}"
echo "Calling 'batchtools::doJobCollection()' ... done"
echo

<% if (details) { %>
if command -v bjobs > /dev/null; then
  echo "Job summary:"
  bjobs -l "${LSB_JOBID}"
fi
<% } %>

<% if (length(shutdown) > 0) {
  writeLines(shutdown)
} %>

echo "End time: $(date +"%Y-%m-%d %H:%M:%S%z")"

## Relay the exit code from Rscript
exit "${res}"
```

Value

Nothing.

References

- https://en.wikipedia.org/wiki/OpenLava

Examples

```
library(future)

# Limit runtime to 10 minutes and total memory to 400 MiB per future,
# request a parallel environment with four slots on a single host.
# Submit to the 'freecycle' queue. Load environment modules 'r' and
# 'jags'. Report on job details at startup and at the end of the job.
```

```
plan(future.batchtools::batchtools_openlava, resources = list(
  W = "00:10:00", M = "400",
  asis = c("-n 4", "-R 'span[hosts=1]'", "-q freecycle"),
  modules = c("r", "jags"),
  details = TRUE
))

f <- future({
  data.frame(
     hostname = Sys.info()[["nodename"]],
           os = Sys.info()[["sysname"]],
    osVersion = utils::osVersion,
        cores = unname(parallelly::availableCores()),
      modules = Sys.getenv("LOADEDMODULES")
  )
})
info <- value(f)
print(info)
```

`batchtools_sge` *A batchtools SGE backend resolves futures in parallel via a Sun/Son of/Oracle/Univa/Altair Grid Engine job scheduler*

Description

A batchtools SGE backend resolves futures in parallel via a Sun/Son of/Oracle/Univa/Altair Grid Engine job scheduler

Usage

```
batchtools_sge(
  ...,
  template = "sge",
  scheduler.latency = 1,
  fs.latency = 65,
  resources = list(),
  delete = getOption("future.batchtools.delete", "on-success"),
  workers = getOption("future.batchtools.workers", default = 100L),
  nodename = "localhost",
  registry = list(),
  finalize = getOption("future.finalize", TRUE)
)
```

Arguments

`template`	(optional) Name of job-script template to be searched for by `batchtools::findTemplateFile()`. If not found, it defaults to the `templates/sge.tmpl` part of this package (see below).

scheduler.latency
: [numeric(1)]
Time to sleep after important interactions with the scheduler to ensure a sane state. Currently only triggered after calling submitJobs.

fs.latency
: [numeric(1)]
Expected maximum latency of the file system, in seconds. Set to a positive number for network file systems like NFS which enables more robust (but also more expensive) mechanisms to access files and directories. Usually safe to set to 0 to disable the heuristic, e.g. if you are working on a local file system.

resources
: (optional) A named list passed to the **batchtools** job-script template as variable resources. This is based on how batchtools::submitJobs() works, with the exception for specially reserved names defined by the **future.batchtools** package;

- resources[["details"]], if TRUE, results in the job script outputting job details and job summaries at the beginning and at the end. It also reports on the number of packages installed in each R package library, which can be slow for large libraries on slow file systems.
- resources[["startup"]] and resources[["shutdown"]] are character vectors of shell code to be injected to the job script as-is.
- resources[["modules"]] is character vector of Linux environment modules to be loaded.
- resources[["envs"]], is an optional names character vector specifying environment variables to be set.
- resources[["rscript"]] is an optional character vector specifying how the 'Rscript' is launched. The resources[["rscript_args"]] field is an optional character vector specifying the 'Rscript' command-line arguments.
- resources[["asis"]] is a character vector that are passed as-is to the job script and are injected as job resource declarations.
- All remaining resources named elements are injected as named resource specification for the scheduler.

delete
: Controls if and when the batchtools job registry folder is deleted. If "on-success" (default), it is deleted if the future was resolved successfully *and* the expression did not produce an error. If "never", then it is never deleted. If "always", then it is always deleted.

workers
: The maximum number of workers the batchtools backend may use at any time, which for HPC schedulers corresponds to the maximum number of queued jobs. The default is getOption("future.batchtools.workers", 100).

nodename
: [character(1)]
Nodename of the master host. All commands are send via SSH to this host. Only works iff

1. Passwordless authentication (e.g., via SSH public key authentication) is set up.
2. The file directory is shared across machines, e.g. mounted via SSHFS.
3. Either the absolute path to the file.dir is identical on the machines, or paths are provided relative to the home directory. Symbolic links should work.

registry	(optional) A named list of settings applied to each future's **batchtools** Registry. This is a more convenient alternative to using argument conf.file.
finalize	If TRUE, a future's **batchtools** Registry is automatically deleted when the future is garbage collected, otherwise not.
...	Not used.

Details

Batchtools SGE futures use **batchtools** cluster functions created by batchtools::makeClusterFunctionsSGE(), which are used to interact with the SGE job scheduler. This requires that SGE commands qsub, qstat, and qdel are available on the current machine.

The default template script templates/sge.tmpl can be found in:

```
system.file("templates", "sge.tmpl", package = "future.batchtools")
```

and comprises:

```
#!/bin/bash
##############################################################################
# A batchtools launch script template for SGE
#
# Author: Henrik Bengtsson
##############################################################################
## Shell
#$ -S /bin/bash

## Job name
#$ -N <%= job.name %>

## Direct streams to logfile
#$ -o <%= log.file %>

## Merge standard error and output
#$ -j y

## Tell the queue system to use the current directory
## as the working directory
#$ -cwd

## Resources needed
<%
  ## Work on a copy of 'resources', because this template is evaluated
  ## in the environment of the batchtools JobCollection, and modifying
  ## 'resources' would modify the resources of the job itself
  job_resources <- resources

  ## Should scheduler "details" be seen?
  details <- isTRUE(job_resources[["details"]])
```

```
  job_resources[["details"]] <- NULL

  ## Shell "startup" code to evaluate
  startup <- job_resources[["startup"]]
  job_resources[["startup"]] <- NULL

  ## Shell "shutdown" code to evaluate
  shutdown <- job_resources[["shutdown"]]
  job_resources[["shutdown"]] <- NULL

  ## Environment modules specifications
  modules <- job_resources[["modules"]]
  job_resources[["modules"]] <- NULL

  ## Environment variables to be set
  envs <- job_resources[["envs"]]
  if (length(envs) > 0) {
    stopifnot(is.character(envs), !is.null(names(envs)))
  }
  job_resources[["envs"]] <- NULL

  ## Custom "Rscript" command and Rscript arguments
  rscript <- job_resources[["rscript"]]
  if (is.null(rscript)) {
    rscript <- "Rscript"
  } else if (length(rscript) == 0 || !nzchar(rscript)[1]) {
    stop("Argument 'resources' specifies an empty 'rscript' field")
  }
  job_resources[["rscript"]] <- NULL
  rscript_args <- job_resources[["rscript_args"]]
  job_resources[["rscript_args"]] <- NULL
  rscript_call <- paste(c(rscript, rscript_args), collapse = " ")

  ## As-is resource specifications
  job_declarations <- job_resources[["asis"]]
  job_resources[["asis"]] <- NULL

  ## Remaining resources are assumed to be of type '<key>=<value>'
  opts <- unlist(job_resources, use.names = TRUE)
  opts <- sprintf("%s=%s", names(opts), opts)
 job_declarations <- sprintf("#$ %s", c(job_declarations, sprintf("-l %s", opts)))
  writeLines(job_declarations)
%>

## Bash settings
set -e          # exit on error
set -u          # error on unset variables
set -o pipefail # fail a pipeline if any command fails
```

```
trap 'echo "ERROR: future.batchtools job script failed on line $LINENO" >&2; exit 1' ERR

<% if (length(job_declarations) > 0) {
  writeLines(c(
    "echo 'Job submission declarations:'",
    sprintf("echo '%s'", job_declarations),
    "echo"
  ))
} %>

<% if (details) { %>
if command -v qstat > /dev/null; then
  echo "Job information:"
  qstat -j "${JOB_ID}"
  echo
fi
<% } %>

<% if (length(startup) > 0) {
  writeLines(startup)
} %>

<% if (length(modules) > 0) {
  writeLines(c(
    "echo 'Load environment modules:'",
    sprintf("echo '- modules: %s'", paste(modules, collapse = ", ")),
    sprintf("module load %s", paste(modules, collapse = " ")),
    "module list"
  ))
} %>

<% if (length(envs) > 0) {
  writeLines(c(
    sprintf("echo 'Setting environment variables: [n=%d]'", length(envs)),
    sprintf("echo ' - %s=%s'", names(envs), shQuote(envs)),
    sprintf("export %s=%s", names(envs), shQuote(envs))
  ))
} %>

echo "Session information:"
echo "- timestamp: $(date +"%Y-%m-%d %H:%M:%S%z")"
echo "- hostname: $(hostname)"
echo "- Rscript call: <%= rscript_call %>"
if ! command -v <%= rscript[1] %> &> /dev/null; then
   >&2 echo "ERROR: Argument 'resources' specifies a non-existing 'Rscript' launch command: <%= rscript[
    exit 1
fi
echo "- Rscript version: $(<%= paste(rscript, collapse = " ") %> --version)"
```

```
echo "- R_LIBS_USER=${R_LIBS_USER:-<not set>}"
echo "- R_LIBS_SITE=${R_LIBS_SITE:-<not set>}"
echo "- R_LIBS=${R_LIBS:-<not set>}"
echo "- Rscript library paths:"
<% if (details) { %>
<%= rscript_call %> -e "for (p in .libPaths()) cat('  - ', shQuote(p), ' [n=', sum(file.exists(file.path
<% } else { %>
<%= rscript_call %> -e "cat(sprintf('  - %s\n', shQuote(.libPaths())), sep = '')"
<% } %>
echo


## Launch R and evaluate the batchtools R job
echo "Calling 'batchtools::doJobCollection()' ..."
echo "- job name: '<%= job.name %>'"
echo "- job log file: '<%= log.file %>'"
echo "- job uri: '<%= uri %>'"
<%= rscript_call %> -e 'batchtools::doJobCollection("<%= uri %>")'
res=$?
echo " - exit code: ${res}"
echo "Calling 'batchtools::doJobCollection()' ... done"
echo

<% if (details) { %>
if command -v qstat > /dev/null; then
  echo "Job summary:"
  qstat -j "${JOB_ID}"
fi
<% } %>

<% if (length(shutdown) > 0) {
  writeLines(shutdown)
} %>

echo "End time: $(date +"%Y-%m-%d %H:%M:%S%z")"

## Relay the exit code from Rscript
exit "${res}"
```

This template and the built-in `batchtools::makeClusterFunctionsSGE()` have been verified to work on a few different Grid Engine HPC clusters;

1. SGE 8.1.9 (Son of Grid Engine), Rocky 8 Linux, BeeGFS global filesystem, UCSF Wynton HPC environment (June 2026)
2. AGE 2024.1.0 (8.9.0), Rocky 9 Linux, NSF global filesystem (August 2025)

Value

Nothing.

References

- https://en.wikipedia.org/wiki/Oracle_Grid_Engine

Examples

```
library(future)

# Limit runtime to 10 minutes and memory to 400 MiB per future,
# request a parallel environment with four slots on a single host.
# Submit to the 'freecycle' queue. Load environment modules 'r' and
# 'jags'. Report on job details at startup and at the end of the job.
plan(future.batchtools::batchtools_sge, resources = list(
  h_rt = "00:10:00", mem_free = "100M",  ## memory is per process
  asis = c("-pe smp 4", "-q freecycle.q"),
  modules = c("r", "jags"),
  details = TRUE
))

f <- future({
  data.frame(
     hostname = Sys.info()[["nodename"]],
           os = Sys.info()[["sysname"]],
    osVersion = utils::osVersion,
        cores = unname(parallelly::availableCores()),
      modules = Sys.getenv("LOADEDMODULES")
  )
})
info <- value(f)
print(info)
```

`batchtools_slurm` *A batchtools Slurm backend resolves futures in parallel via a Slurm job scheduler*

Description

A batchtools Slurm backend resolves futures in parallel via a Slurm job scheduler

Usage

```
batchtools_slurm(
  ...,
  template = "slurm",
  scheduler.latency = 1,
  fs.latency = 65,
  resources = list(),
  delete = getOption("future.batchtools.delete", "on-success"),
  workers = getOption("future.batchtools.workers", default = 100L),
```

```
  array.jobs = TRUE,
  nodename = "localhost",
  registry = list(),
  finalize = getOption("future.finalize", TRUE)
)
```

Arguments

`template`	(optional) Name of job-script template to be searched for by `batchtools::findTemplateFile()`. If not found, it defaults to the `templates/slurm.tmpl` part of this package (see below).
`scheduler.latency`	[`numeric(1)`] Time to sleep after important interactions with the scheduler to ensure a sane state. Currently only triggered after calling `submitJobs`.
`fs.latency`	[`numeric(1)`] Expected maximum latency of the file system, in seconds. Set to a positive number for network file systems like NFS which enables more robust (but also more expensive) mechanisms to access files and directories. Usually safe to set to `0` to disable the heuristic, e.g. if you are working on a local file system.
`resources`	(optional) A named list passed to the **batchtools** job-script template as variable `resources`. This is based on how `batchtools::submitJobs()` works, with the exception for specially reserved names defined by the **future.batchtools** package; • `resources[["details"]]`, if TRUE, results in the job script outputting job details and job summaries at the beginning and at the end. It also reports on the number of packages installed in each R package library, which can be slow for large libraries on slow file systems. • `resources[["startup"]]` and `resources[["shutdown"]]` are character vectors of shell code to be injected to the job script as-is. • `resources[["modules"]]` is character vector of Linux environment modules to be loaded. • `resources[["envs"]]`, is an optional names character vector specifying environment variables to be set. • `resources[["rscript"]]` is an optional character vector specifying how the 'Rscript' is launched. The `resources[["rscript_args"]]` field is an optional character vector specifying the 'Rscript' command-line arguments. • `resources[["asis"]]` is a character vector that are passed as-is to the job script and are injected as job resource declarations. • All remaining `resources` named elements are injected as named resource specification for the scheduler.
`delete`	Controls if and when the batchtools job registry folder is deleted. If `"on-success"` (default), it is deleted if the future was resolved successfully *and* the expression did not produce an error. If `"never"`, then it is never deleted. If `"always"`, then it is always deleted.

`workers` The maximum number of workers the batchtools backend may use at any time, which for HPC schedulers corresponds to the maximum number of queued jobs. The default is `getOption("future.batchtools.workers", 100)`.

`array.jobs` [`logical(1)`]
If array jobs are disabled on the computing site, set to `FALSE`.

`nodename` [`character(1)`]
Nodename of the master host. All commands are send via SSH to this host. Only works iff

1. Passwordless authentication (e.g., via SSH public key authentication) is set up.
2. The file directory is shared across machines, e.g. mounted via SSHFS.
3. Either the absolute path to the `file.dir` is identical on the machines, or paths are provided relative to the home directory. Symbolic links should work.

`registry` (optional) A named list of settings applied to each future's **batchtools** Registry. This is a more convenient alternative to using argument `conf.file`.

`finalize` If TRUE, a future's **batchtools** Registry is automatically deleted when the future is garbage collected, otherwise not.

`...` Not used.

Details

Batchtools Slurm futures use **batchtools** cluster functions created by `makeClusterFunctionsSlurm2()`, which are used to interact with the Slurm job scheduler. This requires that Slurm commands `sbatch`, `squeue`, `sacct`, and `scancel` are available on the current machine.

The default template script `templates/slurm.tmpl` can be found in:

```
system.file("templates", "slurm.tmpl", package = "future.batchtools")
```

and comprises:

```
#!/bin/bash
##############################################################################
# A batchtools launch script template for Slurm
#
# Author: Henrik Bengtsson
##############################################################################

## Job name
#SBATCH --job-name=<%= job.name %>
## Direct streams to logfile
#SBATCH --output=<%= log.file %>

## Resources needed
<%
  ## Work on a copy of 'resources', because this template is evaluated
```

```
## in the environment of the batchtools JobCollection, and modifying
## 'resources' would modify the resources of the job itself
job_resources <- resources

## Should scheduler "details" be seen?
details <- isTRUE(job_resources[["details"]])
job_resources[["details"]] <- NULL

## Shell "startup" code to evaluate
startup <- job_resources[["startup"]]
job_resources[["startup"]] <- NULL

## Shell "shutdown" code to evaluate
shutdown <- job_resources[["shutdown"]]
job_resources[["shutdown"]] <- NULL

## Environment modules specifications
modules <- job_resources[["modules"]]
job_resources[["modules"]] <- NULL

## Environment variables to be set
envs <- job_resources[["envs"]]
if (length(envs) > 0) {
  stopifnot(is.character(envs), !is.null(names(envs)))
}
job_resources[["envs"]] <- NULL

## Custom "Rscript" command and Rscript arguments
rscript <- job_resources[["rscript"]]
if (is.null(rscript)) {
  rscript <- "Rscript"
} else if (length(rscript) == 0 || !nzchar(rscript)[1]) {
  stop("Argument 'resources' specifies an empty 'rscript' field")
}
job_resources[["rscript"]] <- NULL
rscript_args <- job_resources[["rscript_args"]]
job_resources[["rscript_args"]] <- NULL
rscript_call <- paste(c(rscript, rscript_args), collapse = " ")

## As-is resource specifications
job_declarations <- job_resources[["asis"]]
job_resources[["asis"]] <- NULL

## Remaining resources are assumed to be of type '<key>=<value>'
opts <- unlist(job_resources, use.names = TRUE)
opts <- sprintf("%s=%s", names(opts), opts)
job_declarations <- sprintf("#SBATCH %s", c(job_declarations, sprintf("--%s", opts)))
writeLines(job_declarations)
```

```
%>

## Bash settings
set -e          # exit on error
set -u          # error on unset variables
set -o pipefail # fail a pipeline if any command fails
trap 'echo "ERROR: future.batchtools job script failed on line $LINENO" >&2; exit 1' ERR

<% if (length(job_declarations) > 0) {
  writeLines(c(
    "echo 'Job submission declarations:'",
    sprintf("echo '%s'", job_declarations),
    "echo"
  ))
} %>

<% if (details) { %>
if command -v scontrol > /dev/null; then
  echo "Job information:"
  scontrol show job "${SLURM_JOB_ID}"
  echo
fi
<% } %>

<% if (length(startup) > 0) {
  writeLines(startup)
} %>

<% if (length(modules) > 0) {
  writeLines(c(
    "echo 'Load environment modules:'",
    sprintf("echo '- modules: %s'", paste(modules, collapse = ", ")),
    sprintf("module load %s", paste(modules, collapse = " ")),
    "module list"
  ))
} %>

<% if (length(envs) > 0) {
  writeLines(c(
    sprintf("echo 'Setting environment variables: [n=%d]'", length(envs)),
    sprintf("echo ' - %s=%s'", names(envs), shQuote(envs)),
    sprintf("export %s=%s", names(envs), shQuote(envs))
  ))
} %>

echo "Session information:"
echo "- timestamp: $(date +"%Y-%m-%d %H:%M:%S%z")"
echo "- hostname: $(hostname)"
```

```
echo "- Rscript call: <%= rscript_call %>"
if ! command -v <%= rscript[1] %> &> /dev/null; then
   >&2 echo "ERROR: Argument 'resources' specifies a non-existing 'Rscript' launch command: <%= rscript[
    exit 1
fi
echo "- Rscript version: $(<%= paste(rscript, collapse = " ") %> --version)"
echo "- R_LIBS_USER=${R_LIBS_USER:-<not set>}"
echo "- R_LIBS_SITE=${R_LIBS_SITE:-<not set>}"
echo "- R_LIBS=${R_LIBS:-<not set>}"
echo "- Rscript library paths:"
<% if (details) { %>
<%= rscript_call %> -e "for (p in .libPaths()) cat('  - ', shQuote(p), ' [n=', sum(file.exists(file.path
<% } else { %>
<%= rscript_call %> -e "cat(sprintf('  - %s\n', shQuote(.libPaths())), sep = '')"
<% } %>
echo

## Launch R and evaluate the batchtools R job
echo "Calling 'batchtools::doJobCollection()' ..."
echo "- job name: '<%= job.name %>'"
echo "- job log file: '<%= log.file %>'"
echo "- job uri: '<%= uri %>'"
<%= rscript_call %> -e 'batchtools::doJobCollection("<%= uri %>")'
res=$?
echo " - exit code: ${res}"
echo "Calling 'batchtools::doJobCollection()' ... done"
echo

<% if (details) { %>
if command -v sstat > /dev/null; then
  echo "Job summary:"
 sstat --format="JobID,AveCPU,MaxRSS,MaxPages,MaxDiskRead,MaxDiskWrite" --allsteps --jobs="${SLURM_J
fi
<% } %>

<% if (length(shutdown) > 0) {
  writeLines(shutdown)
} %>

echo "End time: $(date +"%Y-%m-%d %H:%M:%S%z")"

## Relay the exit code from Rscript
exit "${res}"
```

This template and the built-in makeClusterFunctionsSlurm2() have been verified to work on a few different Slurm HPC clusters;

1. Slurm 21.08.4, Rocky Linux 8, NFS global filesystem, UCSF C4 HPC environment (June 2026)

2. Slurm 22.05.10, Rocky Linux 9, Lustre global filesystem (September 2025)
3. Slurm 22.05.11, Rocky Linux 8, NFS global filesystem (September 2025)
4. Slurm 23.11.4, Ubuntu 24.04 LTS, NFS global filesystem, UC Berkeley SCF HPC environment (June 2026)
5. Slurm 24.11.3, AlmaLinux 9, Lustre global filesystem (September 2025)*
6. Slurm 24.11.5, Rocky Linux 9, VAST global filesystem, UCSF CoreHPC environment (June 2026)

(*) Verified with **future.batchtools** 0.20.0, which used `batchtools::makeClusterFunctionsSlurm()` since superseeded by `makeClusterFunctionsSlurm2()`.

Value

Nothing.

References

- https://en.wikipedia.org/wiki/Slurm_Workload_Manager

Examples

```
library(future)

# Limit runtime to 10 minutes and memory to 400 MiB per future,
# request a parallel environment with four slots on a single host.
# Submit to the 'freecycle' partition. Load environment modules 'r' and
# 'jags'. Report on job details at startup and at the end of the job.
plan(future.batchtools::batchtools_slurm, resources = list(
  time = "00:10:00", mem = "400M",
  nodes = 1, ntasks = 4, "cpus-per-task" = 1,
  partition = "freecycle",
  modules = c("r", "jags"),
  details = TRUE
))

f <- future({
  data.frame(
     hostname = Sys.info()[["nodename"]],
           os = Sys.info()[["sysname"]],
    osVersion = utils::osVersion,
        cores = unname(parallelly::availableCores())
  )
})
info <- value(f)
print(info)

# As above, but use R from the Rocker 'r-base' Linux container;
#
#   mkdir -p ~/lxc
#   apptainer build ~/lxc/rocker_r-base.sif docker://rocker/r-base
#
```

```
# Example assumes that 'future.batchtools' has already been installed in
# the container to the 'R_LIBS_USER' package folder living on the host;
#
#   R_LIBS_USER=~/R/rocker-%p-library/%v ~/lxc/rocker_r-base.sif
#   ...
#   > chooseCRANmirror(ind = 1)
#   > install.packages("future.batchtools")
#
plan(future.batchtools::batchtools_slurm, resources = list(
  time = "00:10:00", mem = "400M",
  nodes = 1, ntasks = 4, "cpus-per-task" = 1,
  partition = "freecycle",
  details = TRUE,
  envs = c(R_LIBS_USER = "~/R/rocker-%p-library/%v"),
  rscript = c("apptainer", "exec", "~/lxc/rocker_r-base.sif", "Rscript")
))

f <- future({
  data.frame(
     hostname = Sys.info()[["nodename"]],
           os = Sys.info()[["sysname"]],
    osVersion = utils::osVersion,
        cores = unname(parallelly::availableCores()),
      modules = Sys.getenv("LOADEDMODULES")
  )
})
info <- value(f)
print(info)
```

`batchtools_torque` *A batchtools TORQUE backend resolves futures in parallel via a TORQUE/PBS job scheduler*

Description

A batchtools TORQUE backend resolves futures in parallel via a TORQUE/PBS job scheduler

Usage

```
batchtools_torque(
  ...,
  template = "torque",
  scheduler.latency = 1,
  fs.latency = 65,
  resources = list(),
  delete = getOption("future.batchtools.delete", "on-success"),
  workers = getOption("future.batchtools.workers", default = 100L),
  registry = list(),
  finalize = getOption("future.finalize", TRUE)
)
```

Arguments

`template`	(optional) Name of job-script template to be searched for by `batchtools::findTemplateFile()`. If not found, it defaults to the `templates/torque.tmpl` part of this package (see below).
`scheduler.latency`	[`numeric(1)`] Time to sleep after important interactions with the scheduler to ensure a sane state. Currently only triggered after calling `submitJobs`.
`fs.latency`	[`numeric(1)`] Expected maximum latency of the file system, in seconds. Set to a positive number for network file systems like NFS which enables more robust (but also more expensive) mechanisms to access files and directories. Usually safe to set to `0` to disable the heuristic, e.g. if you are working on a local file system.
`resources`	(optional) A named list passed to the **batchtools** job-script template as variable `resources`. This is based on how `batchtools::submitJobs()` works, with the exception for specially reserved names defined by the **future.batchtools** package; • `resources[["details"]]`, if TRUE, results in the job script outputting job details and job summaries at the beginning and at the end. It also reports on the number of packages installed in each R package library, which can be slow for large libraries on slow file systems. • `resources[["startup"]]` and `resources[["shutdown"]]` are character vectors of shell code to be injected to the job script as-is. • `resources[["modules"]]` is character vector of Linux environment modules to be loaded. • `resources[["envs"]]`, is an optional names character vector specifying environment variables to be set. • `resources[["rscript"]]` is an optional character vector specifying how the 'Rscript' is launched. The `resources[["rscript_args"]]` field is an optional character vector specifying the 'Rscript' command-line arguments. • `resources[["asis"]]` is a character vector that are passed as-is to the job script and are injected as job resource declarations. • All remaining `resources` named elements are injected as named resource specification for the scheduler.
`delete`	Controls if and when the batchtools job registry folder is deleted. If `"on-success"` (default), it is deleted if the future was resolved successfully *and* the expression did not produce an error. If `"never"`, then it is never deleted. If `"always"`, then it is always deleted.
`workers`	The maximum number of workers the batchtools backend may use at any time, which for HPC schedulers corresponds to the maximum number of queued jobs. The default is `getOption("future.batchtools.workers", 100)`.
`registry`	(optional) A named list of settings applied to each future's **batchtools** Registry. This is a more convenient alternative to using argument `conf.file`.
`finalize`	If TRUE, a future's **batchtools** Registry is automatically deleted when the future is garbage collected, otherwise not.
`...`	Not used.

Details

Batchtools TORQUE/PBS futures use **batchtools** cluster functions created by `batchtools::makeClusterFunctionsTORQUE` which are used to interact with the TORQUE/PBS job scheduler. This requires that TORQUE/PBS commands `qsub`, `qselect`, and `qdel` are available on the current machine.

The default template script `templates/torque.tmpl` can be found in:

```
system.file("templates", "torque.tmpl", package = "future.batchtools")
```

and comprises:

```
#!/bin/bash
##############################################################################
# A batchtools launch script template for TORQUE/PBS
#
# Author: Henrik Bengtsson
##############################################################################

## Job name
#PBS -N <%= job.name %>

## Direct streams to logfile
#PBS -o <%= log.file %>

## Merge standard error and output
#PBS -j oe

## Resources needed
<%
  ## Work on a copy of 'resources', because this template is evaluated
  ## in the environment of the batchtools JobCollection, and modifying
  ## 'resources' would modify the resources of the job itself
  job_resources <- resources

  ## Should scheduler "details" be seen?
  details <- isTRUE(job_resources[["details"]])
  job_resources[["details"]] <- NULL

  ## Shell "startup" code to evaluate
  startup <- job_resources[["startup"]]
  job_resources[["startup"]] <- NULL

  ## Shell "shutdown" code to evaluate
  shutdown <- job_resources[["shutdown"]]
  job_resources[["shutdown"]] <- NULL

  ## Environment modules specifications
  modules <- job_resources[["modules"]]
```

```
  job_resources[["modules"]] <- NULL

  envs <- job_resources[["envs"]]
  if (length(envs) > 0) {
    stopifnot(is.character(envs), !is.null(names(envs)))
  }
  job_resources[["envs"]] <- NULL

  ## Custom "Rscript" command and Rscript arguments
  rscript <- job_resources[["rscript"]]
  if (is.null(rscript)) {
    rscript <- "Rscript"
  } else if (length(rscript) == 0 || !nzchar(rscript)[1]) {
    stop("Argument 'resources' specifies an empty 'rscript' field")
  }
  job_resources[["rscript"]] <- NULL
  rscript_args <- job_resources[["rscript_args"]]
  job_resources[["rscript_args"]] <- NULL
  rscript_call <- paste(c(rscript, rscript_args), collapse = " ")

  ## Environment variables to be set
  ## As-is resource specifications
  job_declarations <- job_resources[["asis"]]
  job_resources[["asis"]] <- NULL

  ## Remaining resources are assumed to be of type '<key>=<value>'
  opts <- unlist(job_resources, use.names = TRUE)
  opts <- sprintf("%s=%s", names(opts), opts)
 job_declarations <- sprintf("#PBS %s", c(job_declarations, sprintf("-l %s", opts)))
  writeLines(job_declarations)
%>

## Bash settings
set -e          # exit on error
set -u          # error on unset variables
set -o pipefail # fail a pipeline if any command fails
trap 'echo "ERROR: future.batchtools job script failed on line $LINENO" >&2; exit 1' ERR

<% if (length(job_declarations) > 0) {
  writeLines(c(
    "echo 'Job submission declarations:'",
    sprintf("echo '%s'", job_declarations),
    "echo"
  ))
} %>

<% if (details) { %>
if command -v qstat > /dev/null; then
```

```
  echo "Job information:"
  qstat -f "${PBS_JOBID}"
  echo
fi
<% } %>

<% if (length(startup) > 0) {
  writeLines(startup)
} %>

<% if (length(modules) > 0) {
  writeLines(c(
    "echo 'Load environment modules:'",
    sprintf("echo '- modules: %s'", paste(modules, collapse = ", ")),
    sprintf("module load %s", paste(modules, collapse = " ")),
    "module list"
  ))
} %>

<% if (length(envs) > 0) {
  writeLines(c(
    sprintf("echo 'Setting environment variables: [n=%d]'", length(envs)),
    sprintf("echo ' - %s=%s'", names(envs), shQuote(envs)),
    sprintf("export %s=%s", names(envs), shQuote(envs))
  ))
} %>

echo "Session information:"
echo "- timestamp: $(date +"%Y-%m-%d %H:%M:%S%z")"
echo "- hostname: $(hostname)"
echo "- Rscript call: <%= rscript_call %>"
if ! command -v <%= rscript[1] %> &> /dev/null; then
   >&2 echo "ERROR: Argument 'resources' specifies a non-existing 'Rscript' launch command: <%= rscript[
    exit 1
fi
echo "- Rscript version: $(<%= paste(rscript, collapse = " ") %> --version)"
echo "- R_LIBS_USER=${R_LIBS_USER:-<n/a>}"
echo "- R_LIBS_SITE=${R_LIBS_SITE:-<n/a>}"
echo "- R_LIBS=${R_LIBS:-<n/a>}"
echo "- Rscript library paths:"
<% if (details) { %>
<%= rscript_call %> -e "for (p in .libPaths()) cat('  - ', shQuote(p), ' [n=', sum(file.exists(file.path
<% } else { %>
<%= rscript_call %> -e "cat(sprintf('  - %s\n', shQuote(.libPaths())), sep = '')"
<% } %>
echo

## Launch R and evaluate the batchtools R job
```

```
echo "Calling 'batchtools::doJobCollection()' ..."
echo "- job name: '<%= job.name %>'"
echo "- job log file: '<%= log.file %>'"
echo "- job uri: '<%= uri %>'"
<%= rscript_call %> -e 'batchtools::doJobCollection("<%= uri %>")'
res=$?
echo " - exit code: ${res}"
echo "Calling 'batchtools::doJobCollection()' ... done"
echo

<% if (details) { %>
if command -v qstat > /dev/null; then
  echo "Job summary:"
  qstat -f "${PBS_JOBID}"
fi
<% } %>

<% if (length(shutdown) > 0) {
  writeLines(shutdown)
} %>

echo "End time: $(date +"%Y-%m-%d %H:%M:%S%z")"

## Relay the exit code from Rscript
exit "${res}"
```

This template and the built-in batchtools::makeClusterFunctionsTORQUE() have been verified to work on the following PBS/TORQUE HPC cluster;

1. PBSPro 2024.1.2, Rocky 8 Linux, Lustre global filesystem (September 2025)

Value

Nothing.

References

- https://en.wikipedia.org/wiki/TORQUE

Examples

```
library(future)

# Limit runtime to 10 minutes and total memory to 400 MiB per future,
# request a parallel environment with four slots on a single host.
# Submit to the 'freecycle' queue. Load environment modules 'r' and
# 'jags'. Report on job details at startup and at the end of the job.
plan(future.batchtools::batchtools_torque, resources = list(
  walltime = "00:10:00", mem = "100mb",  ## memory is per process
  asis = c("-l nodes=1:ppn=4", "-q freecycle"),
```

```
  modules = c("r", "jags"),
  details = TRUE
))

f <- future({
  data.frame(
     hostname = Sys.info()[["nodename"]],
           os = Sys.info()[["sysname"]],
    osVersion = utils::osVersion,
        cores = unname(parallelly::availableCores()),
      modules = Sys.getenv("LOADEDMODULES")
  )
})
info <- value(f)
print(info)
```

`find_template_file` *Locate a batchtools job-script template file*

Description

Locate a batchtools job-script template file

Usage

```
find_template_file(template)
```

Arguments

`template`	(character string) The name of a batchtools template, e.g. `"sge"` or `"slurm"`, or the full pathname to a template file.

Value

(character string) The absolute pathname to the located template file.

Examples

```
pathname <- find_template_file("sge")
print(pathname)
```

future.batchtools *future.batchtools: A Future for batchtools*

Description

The **future.batchtools** package implements the Future API on top of **batchtools** such that futures can be resolved on for instance high-performance compute (HPC) clusters via job schedulers. The Future API is defined by the **future** package.

Details

To use batchtools futures, attach **future**, and select the type of future you wish to use via `future::plan()`.

Example

```
library(futurize)                          # also attaches 'future'
plan(future.batchtools::batchtools_slurm)  # parallelize via Slurm jobs

# A slow function
slow_fcn <- function(x) {
  Sys.sleep(1.0)
  x^2
}

# Example data
xs <- 1:10

# Sequential and parallel base R apply
ys <- lapply(xs, slow_fcn)
ys <- lapply(xs, slow_fcn) |> futurize()

# Sequential and parallel purrr map
library(purrr)
ys <- xs |> map(slow_fcn)
ys <- xs |> map(slow_fcn) |> futurize()

# Sequential and parallel foreach calls
library(foreach)
ys <- foreach(x = xs) %do% slow_fcn(x)
ys <- foreach(x = xs) %do% slow_fcn(x) |> futurize()
```

Author(s)

Maintainer: Henrik Bengtsson <henrikb@braju.com> (ORCID) [copyright holder]

Authors:

- Henrik Bengtsson <henrikb@braju.com> (ORCID) [copyright holder]

Other contributors:

- Michel Lang (ORCID) (Code adapted from the 'batchtools' package) [contributor, copyright holder]
- Bernd Bischl (ORCID) (Code adapted from the 'batchtools' package) [contributor, copyright holder]

See Also

Useful links:

- https://future.batchtools.futureverse.org
- https://github.com/futureverse/future.batchtools
- Report bugs at https://github.com/futureverse/future.batchtools/issues

---

`makeClusterFunctionsSlurm2` *ClusterFunctions for Slurm Systems (patched)*

Description

This function enhances `batchtools::makeClusterFunctionsSlurm()` with a few patches. Firstly, it patches the `listJobsQueued()` cluster function such that it falls back to querying Slurm's account database (`sacct`), if the future was *not* found in the Slurm job queue (`squeue`), which might be the case when Slurm provisions a job that was just submitted to the scheduler. Secondly, it patches the `submitJob()` cluster function such that the system call to `sbatch` does not capture stderr together with stdout, but rather separately such that any extra INFO messages from `sbatch` do not corrupt the output intended to come from stdout only.

Usage

```
makeClusterFunctionsSlurm2(
  template = "slurm",
  array.jobs = TRUE,
  nodename = "localhost",
  scheduler.latency = 1,
  fs.latency = 65
)
```

Arguments

`template`	`[character(1)]` Either a path to a **brew** template file (with extension "tmpl"), or a short descriptive name enabling the following heuristic for the file lookup: 1. "batchtools.[template].tmpl" in the path specified by the environment variable "R_BATCHTOOLS_SEARCH_PATH". 2. "batchtools.[template].tmpl" in the current working directory.

3. "[template].tmpl" in the user config directory (see `user_config_dir`); on linux this is usually "~/.config/batchtools/[template].tmpl".
4. ".batchtools.[template].tmpl" in the home directory.
5. "[template].tmpl" in the package installation directory in the subfolder "templates".

`array.jobs` [`logical(1)`]
If array jobs are disabled on the computing site, set to `FALSE`.

`nodename` [`character(1)`]
Nodename of the master host. All commands are send via SSH to this host. Only works iff

1. Passwordless authentication (e.g., via SSH public key authentication) is set up.
2. The file directory is shared across machines, e.g. mounted via SSHFS.
3. Either the absolute path to the `file.dir` is identical on the machines, or paths are provided relative to the home directory. Symbolic links should work.

`scheduler.latency` [`numeric(1)`]
Time to sleep after important interactions with the scheduler to ensure a sane state. Currently only triggered after calling `submitJobs`.

`fs.latency` [`numeric(1)`]
Expected maximum latency of the file system, in seconds. Set to a positive number for network file systems like NFS which enables more robust (but also more expensive) mechanisms to access files and directories. Usually safe to set to `0` to disable the heuristic, e.g. if you are working on a local file system.

Value

A batchtools::ClusterFunctions object.

`render_template` *Render a batchtools job script without submitting it*

Description

Renders the job script that **batchtools** would submit to the scheduler for a given batchtools future strategy, without actually submitting (or even creating) a real job. This is useful for inspecting and troubleshooting a template, e.g. to see exactly how `resources` are expanded into scheduler directives.

Usage

```
render_template(strategy = NULL, ...)
```

Arguments

`strategy`	A batchtools future strategy as accepted by `future::plan()`, e.g. `future.batchtools::batchtools_sg` or `future::tweak(future.batchtools::batchtools_sge, resources = list(slots = 8))`. If `NULL` (default), the currently set future plan is used.
`...`	Named resources overriding those already set by `strategy`, e.g. `render_template(future.batchtools time = "00:05:00")`.

Value

A character string holding the rendered job script.

Examples

```
script <- render_template(
  future.batchtools::batchtools_slurm,
  time = "00:05:00", mem = "200M"
)
cat(script)
```

---

zzz-future.batchtools.options
Options used for batchtools futures

Description

Below are the R options and environment variables that are used by the **future.batchtools** package. See future::future.options for additional ones that apply to futures in general.

WARNING: Note that the names and the default values of these options may change in future versions of the package. Please use with care until further notice.

Settings for batchtools futures

'`future.batchtools.workers`': (a positive numeric or `+Inf`) The default number of workers available on HPC schedulers with job queues. (Default: `100`)

'`future.batchtools.output`': (logical) If TRUE, **batchtools** will produce extra output. If FALSE, such output will be disabled by setting **batchtools** options '`batchtools.verbose`' and '`batchtools.progress`' to FALSE. (Default: `getOption("future.debug", FALSE)`)

'`future.batchtools.expiration.tail`': (a positive numeric) When a **batchtools** job expires, the last few lines will be relayed by batchtools futures to help troubleshooting. This option controls how many lines are displayed. (Default: `48L`)

'`future.cache.path`': (character string) An absolute or relative path specifying the root folder in which **batchtools** registry folders are stored. This folder needs to be accessible from all hosts ("workers"). Specifically, it must *not* be a folder that is only local to the machine such as `file.path(tempdir(), ".future"` if a job scheduler on an HPC environment is used. (Default: `.future` in the current working directory)

'`future.batchtools.relay.immediate`': (logical) If TRUE, conditions of class `immediateCondition` are relayed as soon as possible. To do this, the worker saves each of them as a small file in the **batchtools** registry folder of the future, and the main R process reads these files when it polls the future, e.g. using `future::resolved()`. If FALSE, no such files are written, and immediate conditions are treated as any other condition, that is, they are relayed at the end, when the value of the future is collected. Disabling this can be useful if a future signals so many immediate conditions that the file system gets overloaded. (Default: `TRUE`)

'`future.batchtools.script.validate`': (logical) If TRUE, each generated job script that has a Bash or POSIX-shell shebang is checked for syntax errors using `bash -n` before it is submitted, which gives an immediate error instead of a failed job on the compute node. Set to FALSE to disable, e.g. if a false positive is reported. (Default: `TRUE`)

'`future.batchtools.delete`': (character string) Controls whether or not the future's **batchtools** registry folder is deleted after the future result has been collected. If `"always"`, it is always deleted. If `"never"`, it is never deleted. If `"on-success"`, it is deleted if the future resolved successfully, whereas if it failed, it is left as-is to help with troubleshooting. (Default: `"on-success"`)

Environment variables that set R options

All of the above R '`future.batchtools.*`' options can be set by corresponding environment variable `R_FUTURE_BATCHTOOLS_*` *when the* **future.batchtools** *package is loaded*. This means that those environment variables must be set before the **future.batchtools** package is loaded in order to have an effect. For example, if `R_FUTURE_BATCHTOOLS_WORKERS="200"` is set, then option '`future.batchtools.workers`' is set to `200` (numeric).

Examples

```
# Set an R option:
options(future.cache.path = "/cluster-wide/folder/.future")
```

Index

base::I(), *16, 17*
batchtools::ClusterFunctions, *45*
batchtools::findTemplateFile(), *3*, *10*, *18*, *23*, *30*, *37*
batchtools::makeClusterFunctionsInteractive(), *8, 9*
batchtools::makeClusterFunctionsLSF(), *11*
batchtools::makeClusterFunctionsMulticore(), *17*
batchtools::makeClusterFunctionsOpenLava(), *19*
batchtools::makeClusterFunctionsSGE(), *25*, *28*
batchtools::makeClusterFunctionsSlurm(), *35*, *44*
batchtools::makeClusterFunctionsTORQUE(), *38*, *41*
batchtools::submitJobs(), *3*, *11*, *18*, *24*, *30*, *37*
batchtools_bash, 2
batchtools_interactive, 7
batchtools_local, 8
batchtools_lsf, 10
batchtools_multicore, 16
batchtools_multicore(), *9*
batchtools_openlava, 17
batchtools_sge, 23
batchtools_slurm, 29
batchtools_torque, 36

ClusterFunctions, *6*

find_template_file, 42
future.batchtools, 43
future.batchtools-package *(*future.batchtools*)*, 43
future.batchtools.delete *(*zzz-future.batchtools.options*)*, 46
future.batchtools.expiration.tail *(*zzz-future.batchtools.options*)*, 46
future.batchtools.options *(*zzz-future.batchtools.options*)*, 46
future.batchtools.output *(*zzz-future.batchtools.options*)*, 46
future.batchtools.relay.immediate *(*zzz-future.batchtools.options*)*, 46
future.batchtools.script.validate *(*zzz-future.batchtools.options*)*, 46
future.batchtools.workers, *11*, *19*, *24*, *31*, *37*
future.batchtools.workers *(*zzz-future.batchtools.options*)*, 46
future::future.options, *46*
future::plan(), *43*, *46*
future::resolved(), *47*

makeClusterFunctionsBash *(*batchtools_bash*)*, 2
makeClusterFunctionsBash(), *4*
makeClusterFunctionsSlurm2, 44
makeClusterFunctionsSlurm2(), *31*, *34, 35*

R_FUTURE_BATCHTOOLS_DELETE *(*zzz-future.batchtools.options*)*, 46
R_FUTURE_BATCHTOOLS_EXPIRATION_TAIL *(*zzz-future.batchtools.options*)*, 46
R_FUTURE_BATCHTOOLS_OUTPUT *(*zzz-future.batchtools.options*)*, 46

R_FUTURE_BATCHTOOLS_RELAY_IMMEDIATE
(zzz-future.batchtools.options), 46
R_FUTURE_BATCHTOOLS_SCRIPT_VALIDATE
(zzz-future.batchtools.options), 46
R_FUTURE_BATCHTOOLS_WORKERS
(zzz-future.batchtools.options), 46
R_FUTURE_CACHE_PATH
(zzz-future.batchtools.options), 46
Registry, *3*, *8*, *9*, *11*, *16*, *19*, *25*, *31*, *37*
render_template, 45

submitJobs, *10*, *18*, *24*, *30*, *37*, *45*

user_config_dir, *45*

zzz-future.batchtools.options, 46