

# 1 Periscope Frontend

The Periscope Frontend starts the analysis agent hierarchy.

It can also start the application (Parameter `apprun` specifies the command line usually used in `mpirun` for the application) and optimize the mapping of application processes and agents to processors. The number of processes and threads are specified via `ompnumthreads` and `mpinumprocs`.

If the frontend does not start the application, it requires that the application was started manually. It queries the registry for entries with the application name specified via the command line parameter `appname`.

The frontend has two modes: automatic and manual. In automatic mode it also starts the search process. In manual mode, a prompt is shown and commands can be specified by the user.

## 1.1 Environment Variables

|                                   |   |
|-----------------------------------|---|
| PERISCOPE_ROOT                    | Root directory of the Periscope installation. It includes Periscope's configuration file.<br>Default: <code>/usr/local/periscope</code>                                   |
| PSC_REGISTRY<br><hostname>:<port> | Specifies the host and port of the registry service.  |
| PSC_APPNAME                       | Specifies the name of the application. It is either set by the frontend if it starts the application or has to be set by the programmer before starting the application.  |
| PERISCOPE_DEBUG                   | 0..5<br>0=quiet<br>1=startup, found properties in each search<br>2=candidate properties and found properties in each strategy step<br>3=details on refinement<br>4=<br>5= |
|                                   |   |
|                                   |   |
|                                   |   |
|                                   |   |

## 1.2 Periscope Configuration File

The configuration of Periscope can be loaded from a configuration file. Its name is `.periscope`. It is located in Periscope's root directory.

Example:

```
MACHINE      = hlr2
SITE         = LRZ
REGSERVICE_HOST = hlr2
REGSERVICE_PORT = 31337
PERISCOPE_APP = heat
```

### 1.3 Agent Host File

The host to be used by the analysis agents can be specified via a file. It contains a number of lines each specifying a hostname optionally followed by “:<n>” specifying a multiplicity. Lines with # are comment lines. Preceding tabs are ignored.

### 1.4 Arguments

|                              |   |
|------------------------------|---|
| -help                        | Help information  |
| - registry=<Hostname>:<port> | If registry is not specified on the command line, the information is taken from the Periscope configuration file. An error message is generated if it does not exist.<br>Default: Periscope configuration file                                    |
| -port=<port>                 | The port to be used by the frontend.<br>Default: 30000  |
| -maxfan=<n>                  | Determines the fan-out of the tree of high-level agents.<br>Default: 4  |
| -maxcluster=<n>              | Maximum number of MPI processes or OMP threads in a cluster determined in the mapping algorithm.<br>Default: maximum of specified MPI processes and OpenMP threads.<br>Default: 3   |
| -appname=<name>              | It specifies the application to be searched for in the registry. If the value is defined, it will be copied to PSC_APPNAME and used by mrmonitor for the registry entries.<br>Default: appl   |
| -apprun=<appl cmdline>       | This is the command line used by pbsdsh to start an application process. It should be the same as in mpirun -np procs <applcmdline><br>If aprun is not specified, the frontend searches the registry for entries with the given application name. |
| -ompnumthreads=<n>           | Number of OMP threads to be started per MPI process.  |
| -mpinumprocs=<n>             | Number of MPI processes to be started.  |
| -agenthostfile=<filename>    | Full path of the agent host file.   |
| -hlagenthosts=<host_list>    | Currently only a single host name can be specified.<br>Default: localhost   |
| -nodeagenthosts=<host_list>  | Only a single host name can be specified. If none is specified and start_app is requested an error message is generated.  |
| -masterhost=<host_name>      | The host of the master agent.<br>Default: localhost   |
| -timeout=<secs>              | Timeout for startup of agent hierarchy.<br>Default: 20  |
| -debug=level                 | Level of debugging.<br>Default: PERISCOPE_DEBUG or 0  |
| -manual                      | Start frontend in manual mode<br>Default: automatic   |
| -force_localhost             | Passed to master agent.   |
| -dontcluster                 | Passed to master agent.   |
| -interval=<n>                | Currently ignored. Interval between check messages to analysis agents.  |
| -property=<propname>         | Property to be search by analysisagent-<br>Default: L2DMissRate   |
| -strategy=<strategyname>     | Strategy used by analysisagent.   |

|                 |   |
|-----------------|---|
|                 | Default: RegionNestingStrategy  |
| -sir=<filename> | SIR file of the application to be analyzed. Path is relative to home directory.<br>Default: appl_name.sir |
| -threads=<n>    | Number of threads or each MPI process.  |

## 1.5 Frontend Commands

|            |                               |
|------------|-------------------------------|
| graph      | Show the agent network graph. |
| start      | Start search.                 |
| check      | Check for properties.         |
| help       | Show help information.        |
| quit       | Quit Periscope                |
| properties |                               |
|            |                               |
|            |                               |

## 1.6 Example

```
~/psc/frontend/frontend.ia64 --appname=loop --apprun=loop
--sir=mritest/loop.sir --mpinumprocs=1 --
strategy=StallCycleAnalysisBreadthFirst
```

## 1.7 Startup steps

1. If frontend was started from within a PBS job, the list of partitions (PBS\_NODEFILE) is analyzed whether the specified masterhost is part of the list. Otherwise an error message is printed.

If it was started outside of a PBS job, the specified masterhost will be used for the master agent.

2. The master agent is either started via pbsdsh or via ssh.

Its parameters are

- Registry
- Parent host and port
- Executable of the analysis agent
- Either list of application hosts if application is already running or --start\_app plus list of application hosts if it should be started.
- Maxfan
- Maxthreads
- Timeout
- Parameter
- High level agent hosts
- Tag of the masteragent in the registry and in debug messages "fe[%d]:0"
- Force localhost
- Dontcluster
- SIR file
- Property
- Strategy
- Threads
- Quiet

3. Performs the mapping computation. (File frontend.cc)
4. Starts the application processes either with mpirun or as OpenMP application. The applcmdline specified as argument to the frontend is used in mpirun for the application with parameters.
5. Checks whether application components registered.
6. Either enters automatic mode or waits for commands.

## 2 Registry

The registry collects information about the application processes and analysis agents. It is started via `regsrv.ia64&`

The default port is 31337.

## 3 Commands

|              |                                       |
|--------------|---------------------------------------|
| List         | Show the entries                      |
| Clean        | Removes all entries                   |
| Help         | Shows list of commands                |
| Liststr <id> | Shows strings attached with entry id. |

## 4 Master Agent (High-level Agent)

### 4.1 Arguments

|   |   |
|---|---|
| -help   | Help information  |
| - registry=<Hostname>:<port>                  | If registry is not specified on the command line, the information is taken from the Periscope configuration file. An error message is generated if it does not exist.<br>Default: Periscope configuration file                                |
| -port=<port>                                  | The port to be used by the frontend.<br>Default: 30000  |
| -tag =<tag>                                   | All debug messages and the registry entry are marked by tag.  |
| -parent=<Hostname>:<port>                     | Port of the parent agent in the agent hierarchy.  |
| -apphosts=<host1>:<id>:<id>,<host2>:<id>:<id> | Registry IDs of application processes controlled by analysis agents in this high-level agents subtree   |
| -maxfan=<n>                                   | Determines the fan-out of the tree of high-level agents.<br>Default: 4  |
| -maxcluster=<n>                               | Maximum number of MPI processes or OMP threads in a cluster determined in the mapping algorithm. Will be corrected if to next smaller multiple of OMP threads, since they have to be controlled by a single analysis agent.<br><br>Default: 4 |
| -maxthread=<n>                                |   |
| -appname=<name>                               | It specifies the application to be searched for in the registry. If the value is defined, it will be copied to PSC_APPNAME and used by mrimonitor for the registry entries.<br>Default: appl  |
| -apprun=<appl cmdline>                        | This is the command line used by pbsdsh to start an   |

|   |  |
|---|--|
|   | application process. It should be the same as in<br>mpirun -np procs <applcmdline><br>If apprun is not specified, the frontend searches the registry for entries with the give application name.   |
| -dontcluster  | Properties reported to agent are not clustered.  |
| -gatherproperties   | Switches method for combining reported properties.   |
| -ompnumthreads=<n>  | Number of OMP threads to be started per MPI process.   |
| -mpinumprocs=<n>  | Number of MPI processes to be started.   |
| -agenthostfile=<filename>   | Full path of the agent host file.  |
| -hlagenthosts=<host_list>   | Currently only a single host name can be specified.<br>Default: localhost  |
| -nodeagenthosts=<host_list>   | Only a single host name can be specified. If none is specified and start_app is requested an error message is generated.   |
| -masterhost=<host_name>   | The host of the master agent.<br>Default: localhost  |
| -timeout=<secs>   | Timeout for startup of agent hierarchy.<br>Default: 20   |
| -debug=level  | Level of debugging.<br>Default: PERISCOPE_DEBUG or 0   |
| -force_localhost  | Next level of high-level agents is started on localhost independent of the specification in placement.   |
| -dontcluster  | Passed to master agent.  |
| -interval=<n>   | Currently ignored. Interval between check messages to analysis agents.   |
| -property=<propname>  | Property to be search by analysisagent-<br>Default: L2DMissRate  |
| -strategy=<strategyname>  | Strategy used by analysisagent.<br>Default: RegionNestingStrategy  |
| -nodeagentname=<agent>  | Executable of the analysis agent.  |
| -placement:<br>#<analysisagenthost>#<rel. cpu<br>number of anal. agent>,<br><entryID:first thread<br>number>,...'.' <rel.cpu number of<br>anal. Agent>....<br>#a01#1,4:0,6:0.#a09#2,5:0,7:0,8:0 | Specification of mapping of application IDs to analysis agents.<br>List of hosts including list of cpus for analysis agents with its entryIDs<br>First thread number is always zero. The monitor does not allow processes to be controlled by multiple agents. |
| -sir=<filename>   | SIR file of the application to be analyzed.<br>Default: appl_name.sir  |
| -threads=<n>  | Number of threads or each MPI process.   |

## 5 Analysis Agent

The Periscope analysis agent is searching for performance bottlenecks in a subset of the application's MPI processes. It can be started from the hierarchy of agents but also be run as a standalone tool.

If used as a standalone tool, the application has to be running already and the processes have to be registered in the Periscope registry. The tool searches for entries tagged with "MRIMONITOR" and the hostname. It then attaches to those application processes and starts the bottleneck search. The agent itself does not have the ability to restart the application. Therefore a *user region* has to mark an iterative phase of the program.

If the analysis agent is started within the hierarchy, the ids of the processes are passed via a program argument to the agent. It connects to the processes and starts the analysis. If a restart of the application is required to continue the search, a request is propagated to the frontend, the frontend restarts the application and informs the agent of the ids of the same MPI processes. Thus the agent will be responsible for the processes with the same ranks.

## 5.1 Arguments

|                              |   |
|------------------------------|---|
| -help                        | Help information  |
| - registry=<Hostname>:<port> | If registry is not specified on the command line, the information is taken from the Periscope configuration file. If registration is required, i.e., dontregister is not specified, an error message is generated if it does not exist. |
| -dontregister                | Suppresses registration of the agent in the Periscope registry.   |
| -port=<port>                 | The port to be used by the agent.<br>Default: 30000   |
| -appname=<name>              | It specifies the application to be searched for in the registry.<br>Default: appl-loop  |
| - parent=<Hostname>:<port>   | High level agent which is the parent of this analysis agent.  |
| -tag=<string>                | Tag to be used in debug or error messages.<br>Default: local  |
| -debug=level                 | Level of debugging.<br>Default: PERISCOPE_DEBUG or 0  |
| -property=<propname>         | Property to be search by the analysis agent.<br>Default: L2DMissRatio   |
| -strategy=<strategyname>     | Strategy used by the analysis agent.<br>Default: RegionNestingStrategy  |
| -sir=<filename>              | SIR file of the application to be analyzed.   |
| -threads=<n>                 | Number of threads for application startup. In standalone mode it is used to instruct the agent to search in this number of threads.   |
| -id=<id1>, <id2>...          | List of MPI process ids from the registry. If missing, the agent searches for processes in the registry with tag "MRIMONITOR".  |
| - searches                   | Analysis agent performs this number of successive searches. The results of the searches are compared and additional and missing properties are highlighted.   |

## 5.2 Example

```
../psc/analysisagent/analysisagent.ia64 --sir=loop.sir
--strategy=StallCycleAnalysisBreadthFirst
```

Oder

```
../psc/analysisagent/analysisagent.ia64 --sir=loop.sir
--strategy=StallCycleAnalysisBreadthFirst -id=1
```

## 6 Sample PBS Batch file

```
#!/bin/bash
#PBS -j      oe
#PBS -S      /bin/bash
#PBS -l      select=512:ncpus=1
#PBS -l      walltime=00:10:00
#PBS -N      Periscope
#PBS -M      user@email.com
#PBS -m      e
. /etc/profile

# Change to the project's folder
cd psctest/add

# Start the Periscope Registry on port 31390 (default port: 31337)
regsrv.ia64 31390 &

# Wait a bit for the registry to initialize
sleep 10

#Start the front-end and the application
~/psc/frontend/frontend.ia64 --registry=$HOSTNAME:31390 --appname=appl --
apprun=~ /psctest/add/add --sir= psctest/add/add.sir --mpinumprocs=450 --maxcluster=40 --
maxfan=2 --strategy=StallCycleAnalysis --debug=5 --dontcluster
```