

How to configure remote execution



FEATURE AVAILABILITY ENDING WITH RELEASE 1.9 The information in this article applies to JobScheduler releases up to 1.9

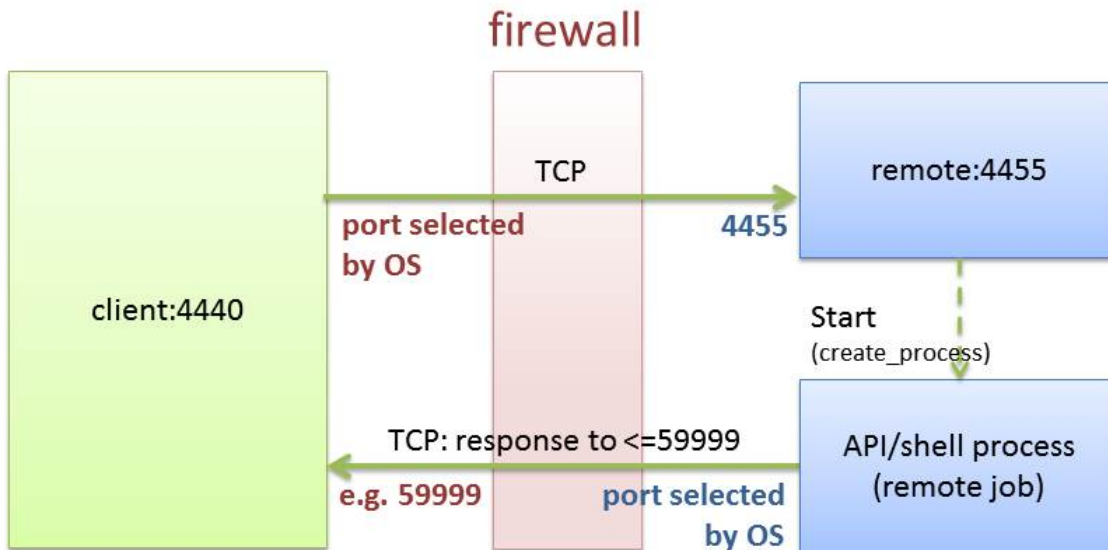
FEATURE AVAILABILITY STARTING FROM RELEASE 1.10 JobScheduler releases starting from 1.10 and later do not use remote execution but use the [JobScheduler Universal Agent](#) and establish a single http connection from the Master to the Agent.

- [Remote Execution](#)
- [Notes](#)

Remote Execution

To trigger a job or an order on a remote JobScheduler you have to open the configured JobScheduler port at the remote side (port 4455 in the example below) and ports 59999 and lower at the client side:

remote execution



incoming port 59999 and before has to be opened at client side

Only incoming ports should be restricted when configuring a firewall for JobScheduler (see the above diagram). The outgoing ports are allocated by operating system functions at random and therefore it is not worthwhile restricting them in the firewall.

The maximum number of remote processes that can run in parallel should be considered when setting the range of the ports to be opened below 59999. It is also important to consider that operating systems generally take around two minutes before they finally release a port and allow it to be used to create a new TCP connection.



Unknown macro: 'mscgen'

Notes

- Use of ports as explained in this article applies to the communication between JobScheduler Master instances for remote execution.
- The communication between JobScheduler Master and JobScheduler Universal Agents makes use of a single HTTP or HTTPS connection, see [JobScheduler Universal Agent - Architecture](#).