Using Multiple Sauce Connect Tunnels

In most situations, you will only need to use one Sauce Connect tunnel for your tests, and all traffic from jobs that are run under your account will use that tunnel automatically and transparently. However, there may be situations in which you want to run multiple Sauce Connect tunnels at the same time. For example, you may want to use a tunnel on your local machine and another on your CI server, or perhaps run multiple jobs on the same CI server. This topic describes how to use tunnel identifiers to run multiple tunnels on the same machine or different machines, and how identified and unidentified tunnels respond to starting up a new tunnel. It also describes an option for running multiple Sauce Connect tunnels as a pool that is treated as single tunnel.

Using Tunnel Identifiers with Multiple Tunnels

If you want to use multiple Sauce Connect tunnels, you need to create tunnel identifiers to manage them. This is because Sauce Connect only supports running one tunnel with the same identifier at at a time. This table shows how tunnels running on different machines with different tunnel identifiers will interact with each other when they start up.

	An Already Running	**Unidentified Tunnel**	**Tunnel `Alpha`**	**Tunnel `Beta`**
When you start a new				
Unidentified Tunnel		Will shut down	Will remain up	Will remain up
Tunnel `Alpha`		Will remain up	Will shut down	Will remain up
Tunnel `Beta`		Will remain up	Will remain up	Will shut down

To create a tunnel identifier, start Sauce Connect using the `--tunnel-identifier` flag (or `-i`) and provide your own unique identifier string. Once the tunnel is up and running, any tests that you want going through this tunnel will need to provide the correct identifier using the `tunnelIdentifier` desired capability. This is how you make sure that a job uses the correct tunnel, as jobs will also automatically make use of unidentified tunnels. This table shows what to expect when you are running jobs with both identified and unidentified tunnels.

	When you have	**No tunnels running**	**An unidentified tunnel running**	**Tunnel `Alpha` running**	**An unidentified tunnel AND tunnel `Alpha` running**
When tunnel-identifier is					
Not provided		No tunnel is used	The unidentified tunnel is used	No tunnel is used	The unidentified tunnel is used
`Alpha`		Your tests don't work	Your tests don't work	Tunnel `Alpha` is used	Tunnel `Alpha` is used
`Beta`		Your tests don't work	Your tests don't work	Your tests don't work	Your tests don't work

Using Tunnels with Some Tests But Not Others

If you want to have some tests use tunnels, and some not use any tunnels at all, use tunnel identifiers for all your tunnels.

Additional Flags for Multiple Tunnels on the Same Machine

If you want to run different tunnels on the same machine, you need to use additional flags to start them, because multiple tunnels on the same machine require different ports, logfiles and PID files. The required flags are `--pidfile`, `--logfile`, `--scproxy-port`, `--se-port`, and `-i`, as shown in this example:

```
sc --pidfile /tmp/sc2.pid --logfile /tmp/sc2.log --scproxy-port 29999 --se-port 4446 -i my-tun2
```

Flag	**Description**

`--pidfile`	A file used to record the process identifier. It can be erased or reused once the tunnel is shut down.
`--logfile`	The location where Sauce Connect writes its logs. This can be reused or erased once the tunnel is shut down, but since Sauce Labs Support will sometimes request this file if you're having test issues, deleting it automatically isn't a great idea.
`--scproxy-port`	An internal Sauce Connect port. This simply needs to be a free port.
`--se-port`	The port for Sauce Connect's inbuilt Selenium Relay. This can be any free port if your test code sends Selenium Commands directly to Sauce Labs. Otherwise, you'll need to match this port to the port used by the test code that will be using this tunnel.
`-i`	The tunnel identifier

Using Sauce Connect Tunnel Pools

If you're using the Sauce Connect High Availability Configuration options, you can run multiple Sauce Connect tunnels as a pool in which traffic to the Sauce Labs browser is shared among tunnels. In this case, the pool of tunnels is considered to be a single tunnel, and you do not need to use tunnel identifiers for your tests in order to prevent collisions between tunnels. For example, if you had three unnamed Sauce Connect tunnels, or you had three tunnels with the same identifier, and you have set the High Availability options for each of your Sauce Connect instance, this set of tunnels will function as a single tunnel.