

# Nginx Access Logs

You can tail Nginx access logs from Treasure Agent (td-agent), to continuously import the access logs into the cloud.

Continue to the following topics:

- [Installing td-agent](#)
- [Modifying td-agent.conf](#)
- [Confirm Data Import](#)
- [Next Steps](#)

## Installing td-agent

'td-agent' must be installed on your application servers. td-agent is a daemon program dedicated to the streaming upload of any kind of the time-series data. td-agent is developed and maintained by Treasure Data.

To set up td-agent, refer to the following articles; we provide deb/rpm packages for Linux systems.

If you have...	Refer to...
MacOS X	<a href="#">Installing td-agent on MacOS X</a>
Ubuntu System	<a href="#">Installing td-agent for Debian and Ubuntu</a>
RHEL / CentOS System	<a href="#">Installing td-agent for Redhat and CentOS</a>
AWS Elastic Beanstalk	<a href="#">Installing td-agent on AWS Elastic Beanstalk</a>

td-agent is under the [fluentd project](#). td-agent extends fluentd with custom plugins for Treasure Data.

## Modifying td-agent.conf

Specify your authentication key by setting the `apikey` option. You can view your API key from the [TD Console](#).

Access `/etc/td-agent/td-agent.conf` to set the `apikey` option in your `td-agent.conf` file.

Note: `YOUR_API_KEY` should be your API key string.

```
# Tailing the Apache Log
<source>
  type tail
  path /var/log/nginx/access.log
  pos_file /var/log/td-agent/nginx-access.pos
  tag td.production.access
  format nginx
</source>

# Treasure Data Input and Output
<match td.*.*>
  type tdlog
  endpoint api.treasuredata.com
  apikey YOUR_API_KEY
  auto_create_table
  buffer_type file
  buffer_path /var/log/td-agent/buffer/td
  use_ssl true
</match>
```

Restart your agent when the following lines are in place.

```
$ sudo /etc/init.d/td-agent restart
```

td-agent keeps tailing the log, buffers the log (`/var/log/td-agent/buffer/td`), and automatically uploads the log into the cloud.

# Confirm Data Import

Sending a SIGUSR1 signal flushes td-agent's buffer; upload starts immediately.

```
# generate access logs
$ curl http://host:port/

# flush the buffer
$ kill -USR1 `cat /var/run/td-agent/td-agent.pid`
```

To confirm that your data has been uploaded successfully, issue the `td tables` command as follows.

```
$ td tables
+-----+-----+-----+-----+
| Database | Table   | Type | Count |
+-----+-----+-----+-----+
| production | access | log  | 1     |
+-----+-----+-----+-----+
```

Check `/var/log/td-agent.log` if it's not working correctly. `td-agent:td-agent` must have permission to access the logs.

**td-agent handles log-rotation.** td-agent keeps a record of the last position of the log, ensuring that each line is read exactly once even if the td-agent process goes down. However, because the information is kept in a file, the "exactly once" guarantee breaks down if the file becomes corrupted.

## Next Steps

We offer a schema mechanism that is more flexible than that of traditional RDBMSs. For queries, we leverage the Hive Query Language.

- [Schema Management](#)
- [Hive Query Language](#)
- [Programmatic Access with REST API and its Bindings](#)