
ActivityWatch Documentation

Release v0.9.0

Erik Bjäreholt and contributors

Mar 15, 2020

1	Introduction	3
1.1	What ActivityWatch is	3
1.2	Reason for existence	3
1.3	Data philosophy	4
2	Getting started	5
2.1	Installation	5
2.1.1	Windows	5
2.1.2	macOS	5
2.1.3	Linux	5
2.2	Usage	6
2.3	Autostart	6
2.4	Config	6
2.5	Installing on GNOME	7
3	Features	9
3.1	User Interface	9
3.1.1	Web Interface	9
3.1.2	Tray icon	9
3.2	Exporting data	10
3.3	Pausing logging	10
3.4	Filtering data	10
4	FAQ	11
4.1	How does ActivityWatch know when I am AFK?	11
4.2	Why is the active window logged as “unknown” when using Wayland?	11
4.3	How do I programmatically use ActivityWatch?	12
4.4	How do I understand the data that is stored?	12
4.5	What happens if it is down or crashes?	12
4.6	What happens when my computer is off or asleep?	12
4.7	Some events have 0 duration. What does this mean?	12
5	History	15
5.1	Present	15
5.2	Future	16
5.2.1	Building new types of privacy-aware services which require data collection	16
5.2.2	Ubiquitous recording for meaningful information about the past	16

6	Watchers	17
6.1	Browser watchers	17
6.2	Editor watchers	17
6.3	Media watchers	18
6.4	Custom watchers	18
7	Importers	19
8	Architecture	21
8.1	Dependency graph	21
8.2	Server	22
8.3	Clients (watchers, importers, and observers)	22
8.4	User interfaces	22
8.5	Libraries	22
8.5.1	aw-core	22
8.5.2	aw-client	22
8.5.3	aw-analysis	23
9	Data model	25
9.1	Buckets	25
9.2	Events	25
9.2.1	Event types	26
10	API Reference	29
10.1	aw_core	30
10.1.1	aw_core.models	30
10.1.2	aw_core.log	30
10.1.3	aw_core.dirs	30
10.2	aw_client	30
10.3	aw_transform	31
10.4	aw_query	33
10.5	aw_server	36
10.5.1	aw_server.api	36
11	Development	39
11.1	Working with submodules	39
11.2	Making a release	39
12	Extending ActivityWatch	41
12.1	Collecting more data	41
12.2	Fetching Data	41
13	Syncing	43
14	Security	45
14.1	ActivityWatch is only as secure as your system	45
14.2	Deleting sensitive data	45
14.3	Encrypting data	45
14.4	Reproducible builds	46
14.5	CORS configuration	46
14.6	More?	46
15	Writing your first watcher	47
15.1	Minimal client	47
15.2	Reference client	48

16 Querying Data	51
16.1 Writing a Query	51
16.2 Fetching Raw Events	53
17 Installing from source	55
17.1 Cloning the submodules	55
17.2 Checking dependencies	55
17.3 Using a virtualenv	56
17.4 Building and installing	56
17.5 Running	56
17.6 Updating from source	57
17.7 Packaging your changes	57
18 Installing from source (on Windows)	59
19 REST API	61
19.1 REST Security	61
19.2 REST Reference	61
19.2.1 Buckets API	62
19.2.2 Events API	62
19.2.3 Heartbeat API	62
19.2.4 Query API	63
20 Changelog	65
20.1 v0.9.0	65
20.2 v0.8.4	66
20.3 v0.8.3	66
20.4 v0.8.2	67
20.5 v0.8.1	67
20.6 v0.8.0b9	67
20.7 v0.8.0b8	68
20.8 v0.8.0b7	68
20.9 v0.8.0b2 - v0.8.0b6	68
20.10 v0.8.0b1	68
20.11 v0.7.1	69
20.12 v0.7.0b4	69
20.13 v0.7.0b3	69
20.14 v0.7.0b2	70
20.15 v0.7.0b1	70
20.16 v0.6.0 and older	70
21 Indices and tables	71
Python Module Index	73
Index	75



If you are new to ActivityWatch, check out the *Getting started guide*.

ActivityWatch is a bundle of software that tracks your computer activity. You are, by default, the sole owner of your data.

It also offers an ecosystem of software to work around it, including ways to collect more data and do different kinds of analysis,

1.1 What ActivityWatch is

- A set of watchers that record relevant information about what you do and what happens on your computer (such as if you are AFK or not, or which window is currently active).
- A way of storing data collected by the watchers.
- A dataformat accomodating most logging needs due to its flexibility.
- An ecosystem of tools to help users extend the software to fit their needs.

1.2 Reason for existence

There are plenty of companies offering services which do collection of Quantified Self data with goals ranging from increasing personal productivity to understanding the people that managers manage (organizational productivity). However, all known services suffer from a significant disadvantage, the users data is in the hands of the service providers which leads to the problem of trust. Every customer of these companies have their data in hands they are forced to trust if they want to use their service.

This is a significant problem, but the true reason that we decided to do something about it was that existing solutions were inadequate. They focused on short-term insight, a goal worthy in itself, but we also want long-term understanding. We made it completely free and open source so anyone can use, improve and extend it.

1.3 Data philosophy

Data in its raw form is always the most valuable.

Quantified self data doesn't take much space by today's standards, but for services such as RescueTime which have over a thousand of customers, every megabyte per user counts.

For the users however, every megabyte of data is worth it. It is therefore of importance that we collect and store data in the highest reasonable resolution such that we later don't have to "fill the gaps" in lower resolution data with lossy heuristics.

Many services doing collection and analysis of QS data today don't actually store the raw data but instead store only summaries (such as only storing how long you used an application during a given hour, instead of storing the individual uses). This is a problem with existing services: they store summarized data instead of the raw data.

This is indicative of that they actually lack a long-term plan. They want to provide a certain type of analysis *today* but we expect to want to do some unknown analysis in the future, and for that we might need the raw data.

Simply put: It is of importance that we start collecting raw data now, because if we don't it will be forever lost.

Getting started with ActivityWatch is as easy as installing, starting it, and setting up autostart (if your installation method doesn't do it for you).

2.1 Installation

2.1.1 Windows

Download and run the Windows installer for the [latest release from GitHub](#).

2.1.2 macOS

Note: macOS 10.15 (Catalina) introduced some complications for running ActivityWatch on macOS, see [issue #334](#).

Download the `.dmg` for the [latest release from GitHub](#) and drag the `.app` to your Applications folder as usual, then add it to your autostart applications.

2.1.3 Linux

Note: If you are using Arch Linux you can install ActivityWatch directly from the [AUR](#).

Download the [latest release from GitHub](#), unzip the archive into an appropriate directory, and add the `aw-qt` executable to your autostart applications.

2.2 Usage

The `aw-qt` application is the easiest way to use ActivityWatch. It creates a trayicon and automatically starts the server and the default watchers.

If you've installed by extracting a zip archive, simply run the `./aw-qt` binary in the installation directory (either from your terminal or on Windows by double-clicking). You now should see an icon appear in your system tray.

You should now also have the web interface running at `localhost:5600` and will in a few minutes be able to view your data in the Activity view!

If you want more advanced ways to run ActivityWatch (including running it without `aw-qt`), check out the “Running” section of *Installing from source*.

Note: If you are running GNOME 3 or another desktop that does not support system trays, or if for some reason Qt can't be used on your machine, read the *Installing on GNOME* section.

Note: If you are using a proxy, activitywatch will not work by default. To circumvent this you can set the environment variable `HTTP_PROXY` before starting `aw-qt`. How to set an environment variable depends on your operating system, use Google if you are unsure how to do this.

2.3 Autostart

Note: Autostart is set up automatically by the Windows installer and for Arch Linux by the AUR package (if your desktop environment supports [XDG Autostart](#)).

You might want to make `aw-qt` start automatically on login using the We hope to automate this for you in the future but for now you'll have to do it yourself. Searching the web for “autostart application <your operating system>” should get you some good results that don't take long.

2.4 Config

Configuration files for ActivityWatch can be found at the following default locations:

- Unix: `~/.config/activitywatch` or the path defined by the `$XDG_CONFIG_HOME` environment variable.
- Mac OS X: `~/Library/Application\ Support/activitywatch/`
- Windows: `%LocalAppData%\activitywatch\activitywatch`

Config options for the server, client, and default watchers are listed below:

- `aw-server`
- `host` Hostname to start the server on. Currently only `localhost` or `127.0.0.1` are supported.
- `port` Port number to start the server on.
- `storage` Type of storage for holding buckets and events. Supported types are `memory`, `mongodb`, or `peewee`.

- `aw-client`
- `hostname` Hostname of the server to connect to.
- `port` Port number of the server to connect to.
- `aw-watcher-afk`
- `timeout` Time in seconds with no activity required to become afk.
- `poll_time` Time in seconds between checks for activity.
- `update_time` Not yet implemented.
- `aw-watcher-window`:
- `poll_time` Time in seconds between window checks.
- `exclude_title` Don't track window titles
- `update_time` Not yet implemented.

2.5 Installing on GNOME

As an alternative for users of GNOME 3 and other DEs that don't support app trays, or simply to avoid depending on Qt, you can place two simple workaround scripts in your ActivityWatch install folder:

`start.sh`:

```
#!/bin/bash

cd ~/.local/opt/activitywatch          # Put your ActivityWatch install folder here

./aw-server/aw-server &
./aw-watcher-afk/aw-watcher-afk &
./aw-watcher-window/aw-watcher-window &          # you can add --exclude-title_
↪here to exclude window title tracking for this session only

notify-send "ActivityWatch started"      # Optional, sends a notification when_
↪ActivityWatch is started
```

`kill.sh`:

```
#!/bin/bash

pkill aw-
notify-send "ActivityWatch killed"      # Optional, sends a notification when_
↪ActivityWatch is killed
```

Don't forget to `chmod +x start.sh` and `chmod +x kill.sh`.

Then you can create two desktop files for these scripts to show up among your apps:

`~/.local/share/applications/aw-start.desktop`:

```
[Desktop Entry]
Name=Start ActivityWatch
Comment=Start AW
Exec=~/.local/opt/activitywatch/start.sh
Hidden=false
Terminal=false
```

(continues on next page)

(continued from previous page)

```
Type=Application
Version=1.0
Icon=activitywatch
Categories=Utility;
```

~/.local/share/applications/aw-kill.desktop:

```
[Desktop Entry]
Name=Kill ActivityWatch
Comment=Kill AW
Exec=~/.local/opt/activitywatch/kill.sh
Hidden=false
Terminal=false
Type=Application
Version=1.0
Icon=activitywatch
Categories=Utility;
```

Here we will document a few features.

3.1 User Interface

3.1.1 Web Interface

ActivityWatch comes with a web interface which currently has the following features:

- **Activity overview**
 - Most used applications by day
 - Timeline
 - Most time spent on a website (*requires the ActivityWatch browser extension*)
- **Bucket overview**
 - When a bucket was last updated
 - Listing of the latest events

More advanced and configurable visualization (such as the ones found in Zenobase and RescueTime) is not a priority and is unlikely to get implemented as a part of the core ActivityWatch project anytime soon.

3.1.2 Tray icon

The tray icon (aw-qt) manages the core ActivityWatch services (server + watchers) and offers:

- Manage which ActivityWatch services to run
- Popup when a service crashes

3.2 Exporting data

If you go to the “Raw Data” page in the ActivityWatch webui you can download any of the buckets which contain every collected datapoint in ActivityWatch as a single file.

You can also export data programatically using the REST API, but we do not have a guide for that yet **todo: explain how**

3.3 Pausing logging

The possibility to pause logging is a low-tech solution to filter sensitive data. You can do it easily by simply unchecking the watcher module you want to pause in the aw-qt trayicon menu, this will stop the watcher until you check it again.

So, if you for example want to pause the logging of window titles:

- Click the ActivityWatch trayicon
- Uncheck ‘Modules -> aw-watcher-window’

3.4 Filtering data

Note: This is a planned feature.

ActivityWatch was born out of a frustration with the privacy issues of existing life logging solutions. We feel that it’s important that some things that are exceptionally sensitive shouldn’t be logged at all. This way the cost of data breach is bounded, and the barrier to sharing your own data will hopefully become smaller.

This is expected to be almost impossible to perfect since what someone considers exceptionally sensitive might not be for someone else (due to e.g. culture and law). But the basics are easy to get right (such as not logging private browser tabs).

For the ones who believe they can adequately protect their data, they should be offered the option to disable the filter.

Currently, the only way to do this is by manually [pausing logging](#).

Note: Some of these questions are technically not frequently asked.

4.1 How does ActivityWatch know when I am AFK?

On Windows and macOS, we use functionality offered by those platforms that gives us the time since last input.

On Linux, we monitor all mouse and keyboard activity so that we can calculate the time since last input. We do not store what that activity was, just that it happened.

With this data (seconds since last input) we then check if there is less than 3 minutes between input activity. If there is, we consider you not-AFK. If more than 3 minutes passes without any input, we consider that as if you were AFK from the last input until the next input occurs.

4.2 Why is the active window logged as “unknown” when using Wayland?

The Wayland protocol does not have a notion of an active window, and it is unlikely to ever have. Wayland is also developed in security in mind, so access should be handed out on an app-by-app basis. This is a good idea, any application shouldn't just give that privacy-sensitive information away freely.

Unfortunately, in Wayland compositors like Gnome's Mutter there is no way at all to get the current window, this leaves the window watcher completely disabled in Wayland.

Solution: Switch to using X11 (the best option), and if you can't: bother the developer of your Wayland compositor.

You can see the general status of the ability of [getting the active window in Wayland on StackOverflow](#) or follow the [issue for ActivityWatch tracking the problem](#).

4.3 How do I programmatically use ActivityWatch?

See the documentation for *Extending ActivityWatch* or checkout the aw-client repository.

4.4 How do I understand the data that is stored?

All ActivityWatch data is represented using *Data model*.

All events from have the fields `timestamp` (ISO 8601 formatted), `duration` (in seconds), and `data` (a JSON object).

You can programmatically get some events yourself to inspect with the following code:

```
ac = aw_client.ActivityWatchClient("")

# Returns a dict with information about every bucket
buckets = ac.get_buckets()

# Get the first bucket
bucket_id = next(buckets.keys())
events = ac.get_events(bucket_id)
```

As an example for AFK events: The data object contains has one attribute `status` which can be `afk` or `not-afk`.

No two events in a bucket should cover the same moment, if that happens there is an issue with the watcher that should be resolved.

4.5 What happens if it is down or crashes?

Since ActivityWatch consists of several modules running independently, one thing crashing will have limited impact on the rest of the system.

If the server crashes, all watchers which use the heartbeat queue should simply queue heartbeats until the server becomes available again. Since heartbeats are currently sent immediately to the server for storage, all data before the crash should be untouched.

If a watcher crashes, its bucket will simply remain untouched until it is restarted.

4.6 What happens when my computer is off or asleep?

If your computer is off or asleep, watchers will usually record nothing. i.e. one events ending (`timestamp + duration`) will not match up with the following event's beginning (`timestamp`).

4.7 Some events have 0 duration. What does this mean?

Watchers most commonly use a polling method called heartbeats in order to store information on the server. Heartbeats are received regularly with some data, and when two consecutive heartbeats have identical data they get merged and the duration of the new one becomes the time difference between the previous two. Sometimes, a single heartbeat doesn't get a following event with identical data. It is then impossible to know the duration of that event.

The assumption could be made to consider all zero-duration events actually have a duration equal to the time of the next event, but all such assumptions are left to the analysis stage.

It all started in 2013 when I, the founder of this project, was just about to start university. I had been soaked in hacker and transhumanism culture for most of my adolescent life, and was eager to put my programming abilities to good use. I had many interests, among others in neuroscience, biohacking, and Quantified Self. A pivotal moment in my interest was when I one day, long after first reading about [brain computer interfaces](#), suddenly realized the implications of the future generations of the technology: We would be able to record our own thoughts and frictionlessly communicate them to others. I wrote a private note to myself, stating an intent that once the field has advanced sufficiently for research to start getting interesting I should make it a priority to contribute as much as I can.

Around the same time, I was obsessively collecting data on my behavior (then known as [lifelogging](#), now more commonly known as [Quantified Self](#)). This included automated time-trackers (like ActivityWatch), a massive spreadsheet, a diary, location tracking, a step/sleep-tracker (Fitbit at the time), extensive use of version control, etc.

While unexpected, the similarities between brain computer interfaces and the behavioral aspects of Quantified Self became apparent over time. After all, the best approximation of our thoughts is our behavior. While we aren't yet able to automatically record our own thoughts, we are able to record our behaviors and what occupies our attention, such as which projects we work on, the ideas we read about, and the culture we consume.

So on Dec 30th, 2014, I started building a prototype. In April 2016 I started working on a rewrite (that included the client-server model) which became the foundation for what ActivityWatch is today. Some time in 2016, my brother [@johan-bjareholt](#) became a regular contributor, and has since become the second largest contributor to the project by a wide margin.

5.1 Present

Development is slowly but steadily moving forward as lead developer Erik Bjäreholt finishes his degree.

Focus currently lies on building tools for data exploration, building an [Android app](#), as well as making it easier to import and export data to and from ActivityWatch.

5.2 Future

There's much to be said here, and while the future is inherently unpredictable we've slowly started outlining our vision for ActivityWatch.

Among other things, we're trying to [secure funding](#) to ensure financial sustainability and accelerate development. In the meanwhile, we get some support from our wonderful users through [donations](#).

5.2.1 Building new types of privacy-aware services which require data collection

Many services rely on the collection of data in order to function, but the more data they need to collect the greater the privacy implications. One way to get around this is to never have a third party get access to the data at all, and keep the user in exclusive control of their data.

Examples:

- [Thankful](#), an application that tracks the users culture consumption, and allows them to automatically donate cryptocurrency to the creators of it.
- Proposal for a [self-hosted newsfeed aggregator](#), with a highly customizable recommendation engine.

5.2.2 Ubiquitous recording for meaningful information about the past

“We live in an interesting time when more and more of our actions can be in some way recorded and played back without our intervention. [...] There's voice recording technology. Web browsing history. Live desktop video recording and playback. Heck, some folks [...] have shown us a taste of the future as power-users of autonomous or assisted self-recording technologies. Go-pro and other consumer tech products are thriving as they discover / cherry-pick / surface compelling use cases. I haven't experienced general-purpose AI which is quite up to the job of organizing my notes for me. But we're close to having ubiquitous recording (and storing bits is the important part – facebook didn't start with entity tags on day 1 but has been able to retroactively infer and index these). There's no way to record everything with perfect fidelity, because that would require us to preserve as many bits as there are in reality (which violates physical constraints) but there's a lot we can do to improve. There are still unexplored frontiers, like recording, transmitting, and playing back one's thoughts (which I don't think we should consider science fiction, just somewhat expensive and contentious to make viable). Suffice to say, interfaces (there aren't great memex-like ways to create graph based notes with semantic, taggable entities), politics and logistics of services competing to silo our information, and insufficient AI to infer our meaning and, in fact, to de-duplicate our thoughts and those of others (read: <https://distill.pub/2017/research-debt>) are major barriers which conspire against making mind-mapping and organizing one's life's work frictionless.”

[@mekarpeles](#) in a [comment on Facebook](#).

Watchers are the parts of ActivityWatch that do all the data collecting.

ActivityWatch comes bundled with two watchers by default:

- [aw-watcher-afk](#) - Watches for mouse & keyboard activity to detect if the user is active.
- [aw-watcher-window](#) - Watches the active window and its title.

The default watchers are collecting some of the most important data. But there is more to collect, so here are some other watchers that let you do so.

6.1 Browser watchers

Watches properties of the active tab like title, URL, and incognito state.

- [aw-watcher-web](#) - The official browser extension, supports Chrome and Firefox.

6.2 Editor watchers

Watches the actively edited file and associated metadata like path, language, and project name (folder name of git root)

- [aw-watcher-vim](#) - vim extension, by [@johan-bjareholt](#) and [@ahnlab](#).
- [aw-watcher-vscode](#) - Visual Studio Code extension, by [@Otto-AA](#).
- [pauldub/activity-watch-mode](#) - emacs mode forked from [wakatime-mode](#), by [@pauldub](#).
- [OlivierMary/aw-watcher-jetbrains](#) - JetBrains IntelliJ plugin, by [@OlivierMary](#).
- [LaggAt/ActivityWatchVS](#) - Visual Studio extension, by [@LaggAt](#)
- [pascalwhoop/aw-idea](#) - (WIP) JetBrains IntelliJ IDEA/PyCharm/WebStorm/etc extension forked from [waka-time](#), by [@pascalwhoop](#)
- [kostasdzizas/aw-watcher-sublime](#) - Sublime Text 3, by [@kostasdzizas](#)

6.3 Media watchers

If you want to more accurately track media consumption.

- `aw-watcher-spotify` - (Beta) Uses the Spotify Web API to get the active track.
- `aw-watcher-chromecast` - (not working yet) Watches what is playing on you Chromecast device.
- `aw-watcher-openvr` - (not working yet) Watches active VR applications.

6.4 Custom watchers

For help on how to write your own watcher, see *Writing your first watcher*.

Have you written one yourself? Send us a PR to have it included!

CHAPTER 7

Importers

ActivityWatch can't track everything, so sometimes you might want to import data into ActivityWatch from another source.

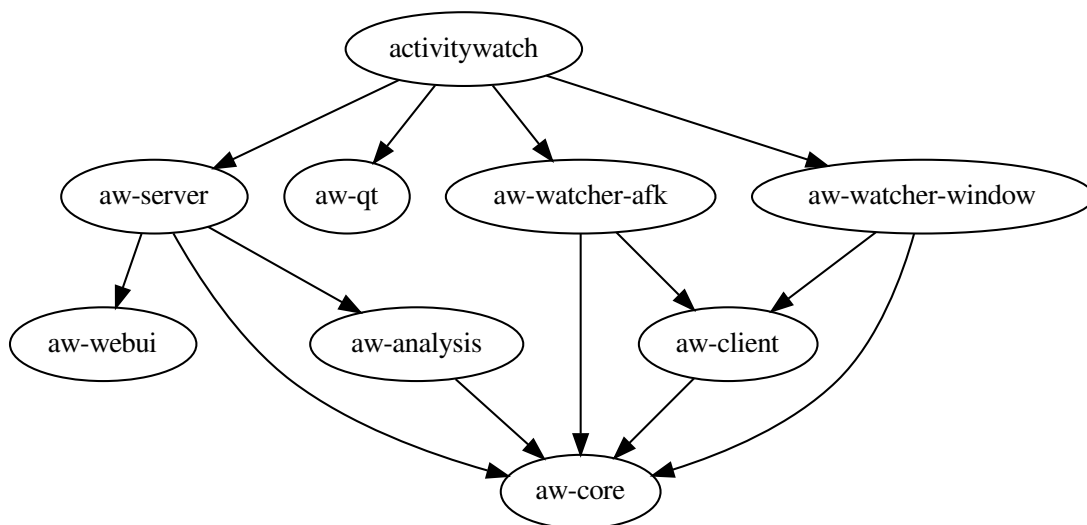
There aren't many yet, but here are some attempts:

- [aw-importer-smartertime](#), imports from [smartertime](#) (Android time tracker).
- LastFM importer, [@ErikBjare](#) has code for it somewhere, ask him if you're interested.

Here we hope to clarify the architecture of ActivityWatch for you. Please file an issue or pull request if you think something is missing.

8.1 Dependency graph

The below is a graph of the fundamental dependencies between projects, these do not reflect the folder structure.



8.2 Server

Known as aw-server, it handles storage and retrieval of all activities/entries in buckets. Usually there exists one bucket per watcher.

The server also hosts the Web UI (aw-webui) which does all communication with the server using the REST API.

8.3 Clients (watchers, importers, and observers)

Since aw-server doesn't do any data collection on it's own, we need watchers that observe the world and sent the data off to aw-server for storage.

These utilize the aw-client library for making requests to the aw-server.

For a list of watchers, see *Watchers*. For a list of importers see *Importers*.

8.4 User interfaces

ActivityWatch currently has two user interfaces, aw-qt and aw-webui.

- **aw-qt** - Manages the server and watchers to make ActivityWatch easy to use for end-users.
- **aw-webui** - Offers visualization and an overview of the database. Hosted by aw-server in the bundle.

8.5 Libraries

Some of the logic of ActivityWatch is shared across the server and clients, for these cases we moved some logic into separate libraries.

8.5.1 aw-core

The aw-core library contains many of the essential parts of ActivityWatch, notably:

- The *Data model*
- The datastore layer
- Event transformation and queries
- Utilities (configuration, logging, decorators)

8.5.2 aw-client

Writing these clients is something we've tried to make as easy as possible by creating client libraries with a clear API. A client could both be a watcher which sends data as well as a visualizer which fetches and presents data from the aw-server.

Currently the primary client library is written in Python (known simply as aw-client) but a client library written in JavaScript is on the way and is expected to have the same level of support in the future.

- **aw-client** (Python)
- **aw-client-js** (TypeScript/JavaScript, beta)

- `aw-client-rust` (Rust, work in progress)

8.5.3 `aw-analysis`

There are also plans to create a library called `aw-analysis` to aid in different types of analysis and transformation one might want to make using ActivityWatch data.

9.1 Buckets

The fundamental datacontainer in ActivityWatch, a bucket contains events and common metadata for those events (such as which type of events they are, where they were collected, and by what).

It is recommended to have one bucket per watcher and host. A bucket should always receive data from the same source.

For example, if we want to write a watcher that should track the currently active window we would first have to create a bucket named ‘example-watcher-window_myhostname’ and then start reporting events to that bucket (using heartbeats).

```
bucket = {
  "id": "aw-watcher-test_myhostname",
  "created": "2017-05-16T13:37:00.000000",
  "name": "A short but descriptive human readable bucketname",
  "type": "com.example.test",          // Type of events in bucket
  "client": "example-watcher-test",    // Identifier of client software used to report_
  ↪data
  "hostname": "myhostname",           // Hostname of device where data was collected
}
```

For information about the “type” field, see examples at *Event types*.

9.2 Events

The event model used by ActivityWatch is pretty simple, here is the JSON representation:

```
event = {
  "timestamp": "2016-04-27T15:23:55Z", // ISO8601 formatted timestamp
  "duration": 3.14,                    // Duration in seconds
```

(continues on next page)

(continued from previous page)

```
"data": {"key": "value"}, // A JSON object, the schema of this depends on the_
↪event type
}
```

It should be noted that all timestamps are stored as UTC. Timezone information (UTC offset) is currently discarded.

The content in the “data” field could be any JSON object, but it is recommended that every event in a bucket should follow some format depending on the buckettype so the data is easy to analyze.

9.2.1 Event types

To separate different types of data in ActivityWatch there is the event type. A buckets event type specified the schema of the events in the bucket.

By creating standards for watchers to use we enable easier transformation and visualization.

web.tab.current

An event type for the currently active webbrowser tab.

```
{
  url: string,
  title: string,
  audible: bool,
  incognito: bool,
}
```

app.editor.activity

An event type for tracking the currently edited file.

```
{
  file: string, // full path to file
  project: string, // full path of cwd
  language: string, // name of language of the file
}
```

currentwindow

Note: There are suggestions to improve/change this format (see [issue #201](#))

```
{
  app: string,
  title: string,
}
```


afkstatus

Note: There are suggestions to improve/change this format (see [issue #201](#))

```
{
  status: string // "afk" or "not-afk"
}
```


CHAPTER 10

API Reference

Here's an API reference for some of the most central components in `aw_core`, `aw_client` and `aw_server`. These are the most important packages in ActivityWatch. A lot of it currently lacks proper docstrings, but it's a start.

Contents

- *API Reference*
 - *aw_core*
 - * *aw_core.models*
 - * *aw_core.log*
 - * *aw_core.dirs*
 - *aw_client*
 - *aw_transform*
 - *aw_query*
 - *aw_server*
 - * *aw_server.api*

10.1 aw_core

10.1.1 aw_core.models

```
class aw_core.models.Event (id: Union[int, str, None] = None, timestamp: Union[datetime.datetime, str] = None, duration: Union[datetime.timedelta, int, float] = 0, data: Dict[str, Any] = {})
```

Used to represents an event.

data

duration

id

timestamp

to_json_dict () → dict

Useful when sending data over the wire. Any mongodb interop should not use do this as it accepts date-times.

to_json_str () → str

10.1.2 aw_core.log

aw_core.log.get_latest_log_file (name, testing=False) → Optional[str]

Returns the filename of the last logfile with name. Useful when you want to read the logfile of another ActivityWatch service.

aw_core.log.get_log_file_path () → Optional[str]

DEPRECATED: Use get_latest_log_file instead.

```
aw_core.log.setup_logging (name: str, testing=False, verbose=False, log_stderr=True, log_file=False, log_file_json=False)
```

10.1.3 aw_core.dirs

aw_core.dirs.ensure_path_exists (path: str) → None

aw_core.dirs.get_cache_dir (module_name: Optional[str]) → str

aw_core.dirs.get_config_dir (module_name: Optional[str]) → str

aw_core.dirs.get_data_dir (module_name: Optional[str]) → str

aw_core.dirs.get_log_dir (module_name: Optional[str]) → str

10.2 aw_client

The `aw_client` package contains a programmer-friendly wrapper around the servers REST API.

```
class aw_client.ActivityWatchClient (client_name: str = 'unknown', testing=False, host=None, port=None, protocol='http')
```

connect ()

create_bucket (*bucket_id: str, event_type: str, queued=False*)
delete_bucket (*bucket_id: str*)
disconnect ()
export_all () → dict
export_bucket (*bucket_id*) → dict
get_buckets ()
get_eventcount (*bucket_id: str, limit: int = -1, start: datetime.datetime = None, end: datetime.datetime = None*) → int
get_events (*bucket_id: str, limit: int = -1, start: datetime.datetime = None, end: datetime.datetime = None*) → List[aw_core.models.Event]
get_info ()
 Returns a dict currently containing the keys ‘hostname’ and ‘testing’.
heartbeat (*bucket_id: str, event: aw_core.models.Event, pulsetime: float, queued: bool = False, commit_interval: Optional[float] = None*) → Optional[aw_core.models.Event]

Args: *bucket_id:* The bucket_id of the bucket to send the heartbeat to *event:* The actual heartbeat event
pulsetime: The maximum amount of time in seconds since the last heartbeat to be merged with the previous heartbeat in aw-server
queued: Use the aw-client queue feature to queue events if client loses connection with the server
commit_interval: Override default pre-merge commit interval

NOTE: This endpoint can use the failed requests retry queue. This makes the request itself non-blocking and therefore the function will in that case always returns None.

import_bucket (*bucket: dict*) → None
insert_event (*bucket_id: str, event: aw_core.models.Event*) → aw_core.models.Event
insert_events (*bucket_id: str, events: List[aw_core.models.Event]*) → None
query (*query: str, start: datetime.datetime, end: datetime.datetime, name: str = None, cache: bool = False*) → Union[int, dict]
send_event (*bucket_id: str, event: aw_core.models.Event*)
send_events (*bucket_id: str, events: List[aw_core.models.Event]*)
setup_bucket (*bucket_id: str, event_type: str*)

10.3 aw_transform

The aw_transform package contains transforms used in the query language.

Note: Their function signatures and return types may deviate from how the transforms are actually implemented in the query language. For more details, see [aw_query.functions](#)

aw_transform.**flood** (*events: List[aw_core.models.Event], pulsetime: float = 5*) → List[aw_core.models.Event]

Takes a list of events and “floods” any empty space between events by extending one of the surrounding events to cover the empty space.

For more details on flooding, see this issue:

- <https://github.com/ActivityWatch/activitywatch/issues/124>

`aw_transform.concat (events1, events2) → List[aw_core.models.Event]`
Concatenates two lists of events

`aw_transform.categorize (events: List[aw_core.models.Event], classes: List[Tuple[List[str], aw_transform.classify.Rule]])`

`aw_transform.tag (events: List[aw_core.models.Event], classes: List[Tuple[str, aw_transform.classify.Rule]])`

`class aw_transform.Rule (rules: Dict[str, Any])`

`match (e: aw_core.models.Event) → bool`

`aw_transform.period_union (events1: List[aw_core.models.Event], events2: List[aw_core.models.Event]) → List[aw_core.models.Event]`

Takes a list of two events and returns a new list of events covering the union of the timeperiods contained in the eventlists with no overlapping events.

Warning: This function strips all data from events as it cannot keep it consistent.

Example:

```
events1 | ----- |
events2 | ----- |
result  | ----- |
```

`aw_transform.filter_period_intersect (events: List[aw_core.models.Event], filterevents: List[aw_core.models.Event]) → List[aw_core.models.Event]`

Filters away all events or time periods of events in which a filterevent does not have an intersecting time period.

Useful for example when you want to filter away events or part of events during which a user was AFK.

Usage: `windowevents_notafk = filter_period_intersect(windowevents, notafkevents)`

Example:

```
events1 | ===== |
events2 | ----- |
result  | ===== |
```

A JavaScript version used to exist in `aw-webui` but was removed in [this PR](#).

`aw_transform.union (events1: List[aw_core.models.Event], events2: List[aw_core.models.Event]) → List[aw_core.models.Event]`

Concatenates and sorts union of 2 event lists and removes duplicates.

Example: Merges events from a backup-bucket with events from a “living” bucket.

```
events = union(events_backup, events_living)
```

`aw_transform.concat (events1, events2) → List[aw_core.models.Event]`
Concatenates two lists of events

`aw_transform.sum_durations (events) → datetime.timedelta`
Sums the durations for the given events

`aw_transform.sort_by_timestamp (events) → List[aw_core.models.Event]`
Sorts a list of events by timestamp

`aw_transform.sort_by_duration(events)` → List[aw_core.models.Event]
Sorts a list of events by duration

`aw_transform.heartbeat_reduce(events: List[aw_core.models.Event], pulsetime: float)` → List[aw_core.models.Event]
Merges consecutive events together according to the rules of [heartbeat_merge](#).

`aw_transform.heartbeat_merge(last_event: aw_core.models.Event, heartbeat: aw_core.models.Event, pulsetime: float)` → Optional[aw_core.models.Event]
Merges two events if they have identical data and the heartbeat timestamp is within the pulsetime window.

`aw_transform.merge_events_by_keys(events, keys)` → List[aw_core.models.Event]

`aw_transform.chunk_events_by_key(events: List[aw_core.models.Event], key: str, pulsetime: float = 5.0)` → List[aw_core.models.Event]

`aw_transform.limit_events(events, count)` → List[aw_core.models.Event]
Returns the count first events in the list of events

`aw_transform.filter_keyvals(events: List[aw_core.models.Event], key: str, vals: List[str], exclude=False)` → List[aw_core.models.Event]

`aw_transform.filter_keyvals_regex(events: List[aw_core.models.Event], key: str, regex: str)` → List[aw_core.models.Event]

`aw_transform.split_url_events(events: List[aw_core.models.Event])` → List[aw_core.models.Event]

`aw_transform.simplify_string(events: List[aw_core.models.Event], key: str = 'title')` → List[aw_core.models.Event]

10.4 aw_query

The `aw_query` package contains the interpreter for the query language and registers the standard functions, usually based on Python implementations of them available in `aw_transform`.

`aw_query.functions.q2_categorize(events: list, classes: list)`

`aw_query.functions.q2_chunk_events_by_key(events: list, key: str)` → List[aw_core.models.Event]

`aw_query.functions.q2_concat(events1: list, events2: list)` → List[aw_core.models.Event]

Note: Documentation automatically copied from underlying function `aw_transform.concat`

Concatenates two lists of events

`aw_query.functions.q2_exclude_keyvals(events: list, key: str, vals: list)` → List[aw_core.models.Event]

`aw_query.functions.q2_filter_keyvals(events: list, key: str, vals: list)` → List[aw_core.models.Event]

`aw_query.functions.q2_filter_keyvals_regex(events: list, key: str, regex: str)` → List[aw_core.models.Event]

`aw_query.functions.q2_filter_period_intersect(events: list, filterevents: list)` → List[aw_core.models.Event]

Note: Documentation automatically copied from underlying function `aw_transform.filter_period_intersect`

Filters away all events or time periods of events in which a filterevent does not have an intersecting time period.

Useful for example when you want to filter away events or part of events during which a user was AFK.

Usage: `windowevents_notafk = filter_period_intersect(windowevents, notafkevents)`

Example:

events1		=====		=====		
events2		-----	---	---	----	
result		====	=	====		

A JavaScript version used to exist in aw-webui but was removed in [this PR](#).

`aw_query.functions.q2_find_bucket` (*datastore: aw_datastore.datastore.Datastore, filter_str: str, hostname: str = None*)

Find bucket by using a filter_str (to avoid hardcoding bucket names)

`aw_query.functions.q2_flood` (*events: list*) → List[aw_core.models.Event]

Note: Documentation automatically copied from underlying function `aw_transform.flood`

Takes a list of events and “floods” any empty space between events by extending one of the surrounding events to cover the empty space.

For more details on flooding, see this issue:

- <https://github.com/ActivityWatch/activitywatch/issues/124>

`aw_query.functions.q2_limit_events` (*events: list, count: int*) → List[aw_core.models.Event]

Note: Documentation automatically copied from underlying function `aw_transform.limit_events`

Returns the count first events in the list of events

`aw_query.functions.q2_merge_events_by_keys` (*events: list, keys: list*) → List[aw_core.models.Event]

`aw_query.functions.q2_nop` ()

No operation function for unittesting

`aw_query.functions.q2_period_union` (*events1: list, events2: list*) → List[aw_core.models.Event]

Note: Documentation automatically copied from underlying function `aw_transform.period_union`

Takes a list of two events and returns a new list of events covering the union of the timeperiods contained in the eventlists with no overlapping events.

Warning: This function strips all data from events as it cannot keep it consistent.

Example:

```
events1 | | ----- |
events2 | | ----- |
result  | | ----- |
```

```
aw_query.functions.q2_query_bucket (datastore:          aw_datastore.datastore.Datastore,
                                   namespace: Dict[str, Any], bucketname: str) →
                                   List[aw_core.models.Event]
```

```
aw_query.functions.q2_query_bucket_eventcount (datastore:
                                                aw_datastore.datastore.Datastore,
                                                namespace: Dict[str, Any], bucketname:
                                                str) → int
```

```
aw_query.functions.q2_simplify_window_titles (events: list, key: str) →
                                              List[aw_core.models.Event]
```

```
aw_query.functions.q2_sort_by_duration (events: list) → List[aw_core.models.Event]
```

Note: Documentation automatically copied from underlying function `aw_transform.sort_by_duration`

Sorts a list of events by duration

```
aw_query.functions.q2_sort_by_timestamp (events: list) → List[aw_core.models.Event]
```

Note: Documentation automatically copied from underlying function `aw_transform.sort_by_timestamp`

Sorts a list of events by timestamp

```
aw_query.functions.q2_split_url_events (events: list) → List[aw_core.models.Event]
```

```
aw_query.functions.q2_sum_durations (events: list) → datetime.timedelta
```

Note: Documentation automatically copied from underlying function `aw_transform.sum_durations`

Sums the durations for the given events

```
aw_query.functions.q2_tag (events: list, classes: list)
```

10.5 aw_server

10.5.1 aw_server.api

The `ServerAPI` class contains the basic API methods, these methods are primarily called from RPC layers such as the one found in `aw_server.rest`.

class `aw_server.api.ServerAPI` (*db, testing*)

create_bucket (*bucket_id: str, event_type: str, client: str, hostname: str, created: Optional[datetime.datetime] = None*) → bool

Create bucket. Returns True if successful, otherwise false if a bucket with the given ID already existed.

create_events (*bucket_id: str, events: List[aw_core.models.Event]*) → Optional[aw_core.models.Event]

Create events for a bucket. Can handle both single events and multiple ones.

Returns the inserted event when a single event was inserted, otherwise None.

delete_bucket (*bucket_id: str*) → None

Delete a bucket

delete_event (*bucket_id: str, event_id*) → bool

Delete a single event from a bucket

export_all () → Dict[str, Any]

Exports all buckets and their events to a format consistent across versions

export_bucket (*bucket_id: str*) → Dict[str, Any]

Export a bucket to a dataformat consistent across versions, including all events in it.

get_bucket_metadata (*bucket_id: str*) → Dict[str, Any]

Get metadata about bucket.

get_buckets () → Dict[str, Dict[KT, VT]]

Get dict {bucket_name: Bucket} of all buckets

get_eventcount (*bucket_id: str, start: datetime.datetime = None, end: datetime.datetime = None*) → int

Get eventcount from a bucket

get_events (*bucket_id: str, limit: int = -1, start: datetime.datetime = None, end: datetime.datetime = None*) → List[aw_core.models.Event]

Get events from a bucket

get_info () → Dict[str, Dict[KT, VT]]

Get server info

get_log ()

Get the server log in json format

heartbeat (*bucket_id: str, heartbeat: aw_core.models.Event, pulsetime: float*) → aw_core.models.Event

Heartbeats are useful when implementing watchers that simply keep track of a state, how long it's in that state and when it changes. A single heartbeat always has a duration of zero.

If the heartbeat was identical to the last (apart from timestamp), then the last event has its duration updated. If the heartbeat differed, then a new event is created.

Such as:

- Active application and window title - Example: aw-watcher-window

- Currently open document/browser tab/playing song - Example: wakatime - Example: aw-watcher-web - Example: aw-watcher-spotify
- Is the user active/inactive? Send an event on some interval indicating if the user is active or not. - Example: aw-watcher-afk

Inspired by: <https://wakatime.com/developers#heartbeats>

import_all (*buckets: Dict[str, Any]*)

import_bucket (*bucket_data: Any*)

query2 (*name, query, timeperiods, cache*)

Note: This part is a work in progress, reach out to the maintainers if you have any questions!

We recommend you follow Kenneth Reitz folder structure guide when writing Python programs which will be under the control of the ActivityWatch organisation: <http://docs.python-guide.org/en/latest/writing/structure/>

11.1 Working with submodules

Working with submodules comes with some complexity, here are a few neat tricks to make things easier:

- We recommend configuring git to include submodule changes in `git status`, you can do so with the following: `git config --global status.submoduleSummary true`
- If you want the latest committed version of all submodules, use: `git submodule update --recursive`
- If you want the latest master branch on all submodules, use: `git submodule update --recursive --remote`
- If you want to ensure you've pushed all commits in the submodules, use: `git submodule foreach 'git push'`

A longer guide to git submodules can be found [here](#).

11.2 Making a release

1. Close [milestone on GitHub](#) if one exists.
2. Ensure that all the tests pass: `make test && make test-integration`
3. **Test the latest build and check that it works correctly**

- Travis artifacts are available in S3 at the base URL: <https://activitywatch-builds.s3.amazonaws.com/>
 - Appveyor artifacts are available on Appveyor.
4. Write a changelog entry in `docs/changelog.rst`
 5. Sign the commit: `git commit -a -S -m "bumped version"`
 6. Create a signed tag: `git tag -s v0.7.1`
 7. Push the commit and tag: `git push origin refs/tags/v0.7.1`
 8. **Create a release on GitHub**
 - Generate commit changelog with `scripts/release_notes.sh`
 - Read it from `commit_summary.txt` and clean it a bit (remove non-user-affecting changes, merge commits etc)
 9. Wait for the builds to finish
 10. Post about it online: Twitter, the forum, mailinglist (if major)

Extending ActivityWatch

So, you want to do something more with ActivityWatch? Great!

We've tried to make things easy for you (and ourselves) so here's some advice on how to get started.

12.1 Collecting more data

ActivityWatch is written to be flexible to be able to gather most types of data. Except for the included `aw-watcher-window` and `aw-watcher-afk` which tracks your application usage, there are additional so-called *Watchers* for activity-watch. Watchers are small programs that collect data and send it off to the server. The only requirement for what kind of data is sent to `aw-server` as an event is that it has to contain a `starttime` (and preferably a `duration` aswell) so it can fit on a timeline.

If you want to write a watcher of your own, see *Writing your first watcher*.

12.2 Fetching Data

If you want to fetch data from `aw-server` for visualization, exporting, backup or something we have not yet thought of, there are a few ways you can do this:

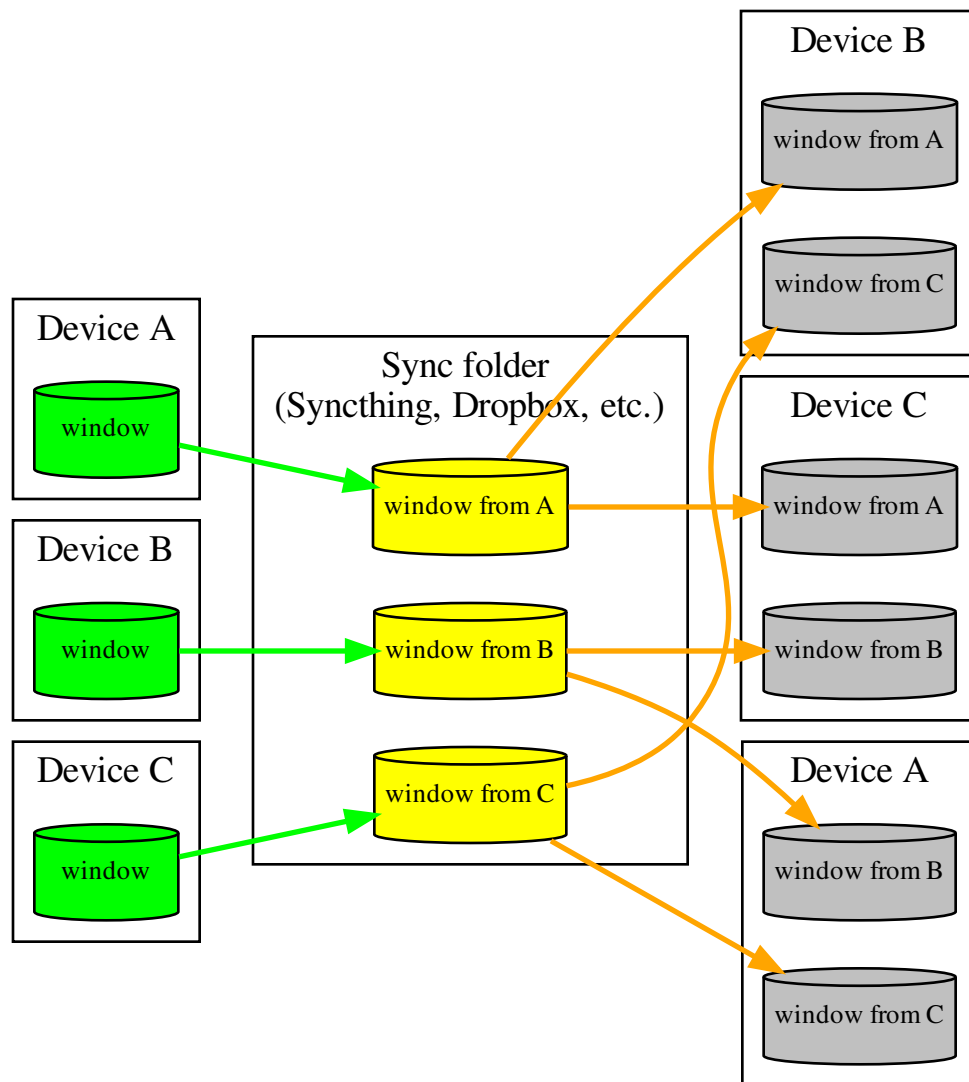
- *Exporting a Bucket* If you want a complete dump of all events of bucket
- *Bucket REST API* If you want to export raw events in a specific time interval from a bucket
- *Writing a Query* If you want to summarize/aggregate one or more buckets into more easily readable data

CHAPTER 13

Syncing

There isn't much written about syncing yet since it's not yet implemented in a stable release. However, there does exist a working proof-of-concept prototype which should be easy to implement once details have been finalized. You can read what has been discussed in this issue: <https://github.com/ActivityWatch/activitywatch/issues/35>

Here's a graph showing how data flows in the current syncing prototype:



Green boxes are source buckets (only written to and read from by the owner). Yellow boxes are the synced version of buckets (written to by the owner, read by consumers). Gray boxes are local copies of remote buckets.

It can be briefly described as follows:

Device A takes its buckets to sync and puts the data in the synced copy, the synced copy gets distributed to device B, device B takes the synced copy and imports it to its local datastore.

ActivityWatch deals with highly sensitive data, and the security of it is therefore of paramount importance.

Unfortunately, we don't have a lot of resources, so things like security audits are currently out of reach for us.

We do try our best to keep security in mind. In this section of the documentation we'll outline some important security considerations, including risks and possible improvements.

14.1 ActivityWatch is only as secure as your system

Some things we can't protect against. Examples are malware running on the same host and anything that can access the database file.

14.2 Deleting sensitive data

Some data may wish to be deleted/filtered/redacted, or simply never logged at all. Making this easy should be one of the most basic privacy features we can add.

This is actually issue #1 in the ActivityWatch repository: <https://github.com/ActivityWatch/activitywatch/issues/1>

14.3 Encrypting data

Encrypting old data with a password would minimize the amount of sensitive data that would be leaked in case of a breach.

The easiest way to build this would be to write a client that takes all events older than some duration and moves it into an encrypted container. This way it wouldn't add complexity to the server code.

14.4 Reproducible builds

It's important that our builds are reproducible, such that we can ensure the integrity of a built package.

We currently lack tests for it, so we don't actually know if they are (they should be, at least some of them).

14.5 CORS configuration

CORS is configured such that origins can only be `localhost:5600` or match the ActivityWatch WebExtension URL for Chrome, or **any extension** on Firefox.

This is due to that on Chrome, the origin of a WebExtension is always a fixed URL. In Firefox however the URL changes for each install, in order to prevent fingerprinting which extensions are installed. It's mentioned here: <https://github.com/ActivityWatch/aw-server-rust/issues/24#issuecomment-520802579>

This means that on Firefox, a malware WebExtension could easily fetch the entire datastore and do what it wants with it.

Ways to solve this:

- Short term: Restrict what we let those origins do (i.e. only send heartbeats, maybe even only to a certain bucket)
- Long term: Use an OAuth2 authentication flow when first installing the extension (this also adds many opportunities for integrations)

14.6 More?

This is an early version of this document. There might be more things mentioned in issues (search for "security").

Writing your first watcher

Writing watchers for ActivityWatch is pretty easy, all you need is the `aw-client` library.

Note: These examples runs the client in *testing* mode, which means that it will try to connect to a `aw-server` in testing mode on the port 5666 instead of the normal 5600.

15.1 Minimal client

Below is a minimal template client to quickly get started. This example will:

- create a bucket
- insert an event
- fetch an event from an `aw-server` bucket
- delete the bucket again

```
#!/usr/bin/env python3

from datetime import datetime, timezone

from aw_core.models import Event
from aw_client import ActivityWatchClient

# We'll run with testing=True so we don't mess up any production instance.
# Make sure you've started aw-server with the `--testing` flag as well.
client = ActivityWatchClient("test-client", testing=True)

bucket_id = "{}_{}".format("test-client-bucket", client.hostname)
client.create_bucket(bucket_id, event_type="dummydata")

shutdown_data = {"label": "some interesting data"}
```

(continues on next page)

(continued from previous page)

```

now = datetime.now(timezone.utc)
shutdown_event = Event(timestamp=now, data=shutdown_data)
inserted_event = client.insert_event(bucket_id, shutdown_event)

events = client.get_events(bucket_id=bucket_id, limit=1)
print(events) # Should print a single event in a list

client.delete_bucket(bucket_id)

```

15.2 Reference client

Below is an example of a watcher with more in-depth comments. This example will describe how to:

- how to create buckets
- how to send events by heartbeats
- how to insert events without heartbeats
- how to do synchronous as well as asynchronous requests
- fetch events from an aw-server bucket
- delete buckets

```

#!/usr/bin/env python3

from time import sleep
from datetime import datetime, timedelta, timezone

from aw_core.models import Event
from aw_client import ActivityWatchClient

# We'll run with testing=True so we don't mess up any production instance.
# Make sure you've started aw-server with the `--testing` flag as well.
client = ActivityWatchClient("test-client", testing=True)

# Make the bucket_id unique for both the client and host
# The convention is to use client-name_hostname as bucket name,
# but if you have multiple buckets in one client you can add a
# suffix such as client-name-event-type or similar
bucket_id = "{}_{}".format("test-client-bucket", client.hostname)
# A short and descriptive event type name
# Will be used by visualizers (such as aw-webui) to detect what type and format the
# events are in
# Can for example be "currentwindow", "afkstatus", "ping" or "currentsong"
event_type = "dummydata"

# First we need a bucket to send events/heartbeats to.
# If the bucket already exists aw-server will simply return 304 NOT MODIFIED,
# so run this every time the client starts up to verify that the bucket exists.
# If the client was unable to connect to aw-server or something failed
# during the creation of the bucket, an exception will be raised.
client.create_bucket(bucket_id, event_type="test")

# Asynchronous loop example

```

(continues on next page)

(continued from previous page)

```

with client:
    # This context manager starts the queue dispatcher thread and stops it when done,
    ↪always use it when setting queued=True.
    # Alternatively you can use client.connect() and client.disconnect() instead if
    ↪you prefer that

    # Create a sample event to send as heartbeat
    heartbeat_data = {"label": "heartbeat"}
    now = datetime.now(timezone.utc)
    heartbeat_event = Event(timestamp=now, data=heartbeat_data)

    # Now we can send some events via heartbeats
    # This will send one heartbeat every second 5 times
    sleeptime = 1
    for i in range(5):
        # The duration between the heartbeats will be less than pulsetime, so they
        ↪will get merged.
        # TODO: Make a section with an illustration on how heartbeats work and insert
        ↪a link here
        print("Sending heartbeat {}".format(i))
        client.heartbeat(bucket_id, heartbeat_event, pulsetime=sleeptime+1,
        ↪queued=True)

        # Sleep a second until next heartbeat
        sleep(sleeptime)

        # Update timestamp for next heartbeat
        heartbeat_event.timestamp = datetime.now(timezone.utc)

    # Give the dispatcher thread some time to complete sending the last events.
    # If we don't do this the events might possibly queue up and be sent the
    # next time the client starts instead.
    sleep(1)

# Synchronous example, insert an event
event_data = {"label": "non-heartbeat event"}
now = datetime.now(timezone.utc)
event = Event(timestamp=now, data=event_data)
inserted_event = client.insert_event(bucket_id, event)

# The event returned from insert_event has been assigned an id by aw-server
assert inserted_event.id is not None

# Fetch last 10 events from bucket
# Should be two events in order of newest to oldest
# - "shutdown" event with a duration of 0
# - "heartbeat" event with a duration of 5*sleeptime
events = client.get_events(bucket_id=bucket_id, limit=10)
print(events)

# Now lets clean up after us.
# You probably don't want this in your watchers though!
client.delete_bucket(bucket_id)

# If something doesn't work, run aw-server with --verbose to see why some request
↪doesn't go through
# Good luck with writing your own watchers :-)
```


There are a couple of ways to query data in activitywatch.

aw-server supplies an “/query” endpoint (also accesible via aw-client’s query method) which supplies a basic scripting language which you can utilize to do transformations on the server-side. This option is good for basic analysis and for lightweight clients (such as aw-webui).

Another option is to fetch events from the “/buckets/bucketname/events” endpoint (also accesible via aw-client’s get_events method) and either program your own transformations or use transformation methods available in the aw-analysis python library (which includes all transformations available in the query endpoint). This require a lot of more work since you will likely have to reprogram transformations already available in the query API, but on the other hand it is much more flexible.

16.1 Writing a Query

Note: This section is still WIP. There is still no documentation of all the transform functions, but for most simple queries these examples should be enough.

Queries are the easiest yet advanced way to get events from aw-server buckets in a format which fits most needs. Queries can be done by doing a POST request to aw-server either manually or with the aw-client library.

For an incomplete API reference of the transform functions, see the API reference for *aw_transform* and *aw_query*.

In a query you start by getting events from a bucket and assign that collection of events to a variable, then there are multiple transform functions which you can use to for example filter, limit, sort, and merge events from a bucket. After that you assign what you want to receive from the request to the RETURN variable.

Minimal example: Minimal query which only gets events from a bucket and returns it:

```
events = query_bucket("my_bucket");  
RETURN = events;
```

Example which arranges a hierarchy: A query which merges events from a bucket in a key1->key2 hierarchy:

```

events = query_bucket("my_bucket");
events = merge_events_by_keys(events, "merged_key1", "merged_key2");
RETURN = events;

```

Example combining window and AFK events: A simplified query example of how to summarize what programs used while not afk. The query intersects the not-afk events from the afk bucket with the events from the window bucket, merges keys from the result and sorts by duration.

```

window_events = query_bucket("window_bucket");
not_afk_events = query_bucket("afk_bucket");
not_afk_events = filter_keyvals(not_afk_events, "status", ["not-afk"]);
window_events = filter_period_intersect(window_events, not_afk_events);
events = merge_events_by_keys(window_events, "appname");
events = sort_by_duration(events);
RETURN = events;

```

Example including aw-client: This is an example of how you can do analysis and aggregation with the query method in python with aw-client

```

#!/usr/bin/env python3

from time import sleep
from datetime import datetime, timedelta, timezone

from aw_core.models import Event
from aw_client import ActivityWatchClient

client = ActivityWatchClient("test-client", testing=True)

now = datetime.now(timezone.utc)
start = now

query = "RETURN=0;"
res = client.query(query, "1970-01-01", "2100-01-01")
print(res) # Should print 0

bucket_id = "{}_{}".format("test-client-bucket", client.hostname)
event_type = "dummydata"
client.create_bucket(bucket_id, event_type="test")

def insert_events(label: str, count: int):
    global now
    events = []
    for i in range(count):
        e = Event(timestamp=now,
                  duration=timedelta(seconds=1),
                  data={"label": label})
        events.append(e)
        now = now + timedelta(seconds=1)
    client.insert_events(bucket_id, events)

insert_events("a", 5)

query = "RETURN = query_bucket('{}');".format(bucket_id)

res = client.query(query, "1970", "2100")
print(res) # Should print the last 5 events

```

(continues on next page)

(continued from previous page)

```
res = client.query(query, start + timedelta(seconds=1), now -  
↳timedelta(seconds=2))  
print(res) # Should print three events  
  
insert_events("b", 10)  
  
query = """  
events = query_bucket('{}');  
merged_events = merge_events_by_keys(events, 'label');  
RETURN=merged_events;  
""".format(bucket_id)  
res = client.query(query, "1970", "2100")  
# Should print two merged events  
# Event "a" with a duration of 5s and event "b" with a duration of 10s  
print(res)  
  
client.delete_bucket(bucket_id)
```

16.2 Fetching Raw Events

TODO: Write this section

Bucket REST API

Installing from source

Here's the guide to installing ActivityWatch from source. If you are just looking to try it out, see the getting started guide instead.

Note: This is written for Linux and macOS. For Windows the build process is more complicated and we therefore suggest using the pre-built packages instead on that operating system (but if you really have to, see *this guide*).

17.1 Cloning the submodules

Since the ActivityWatch bundlerepo uses submodules, you first need to clone the submodules.

This can either be done at the cloning stage with:

```
git clone --recursive https://github.com/ActivityWatch/activitywatch.git
```

Or afterwards (if you've already cloned normally) using:

```
git submodule update --init --recursive
```

17.2 Checking dependencies

You need:

- Python 3.6 or later, check with `python3 -V` (required to build the core components)
- Node 8 or higher, check with `node -v` and `npm -v` (required to build the web UI)

17.3 Using a virtualenv

Note: If you don't want to use a virtualenv you could instead set the environment variable `PIP_USER=true` when building in the next step. But make sure that the folder `~/.local/bin` (on Linux) or `~/Library/Python/<version>/bin` (on macOS) is in your `PATH`.

It is recommended to use a virtualenv in order to avoid polluting your system with ActivityWatch-specific Python packages. It also makes it easier to uninstall since all you have to do is remove the virtualenv folder.

```
python3 -m venv venv
```

Now activate the virtualenv in your current shell session:

```
# For bash/zsh users:
source ./venv/bin/activate
# For Windows git bash users:
source ./venv/Scripts/activate
# For fish users:
source ./venv/bin/activate.fish
```

17.4 Building and installing

Build and install everything into the virtualenv:

```
make build
```

Note: If you're building from source to develop we suggest building/installing using `make build DEV=true` which installs all Python packages with pip's handy `--editable` flag. By doing this you won't have to reinstall everything whenever you want to try out a code change.

17.5 Running

Now you should be able to start ActivityWatch **from the terminal where you've activated the virtualenv**. Or, if you were using the `PIP_USER` trick, from any terminal with a correctly configured `PATH`. You have two options:

1. Use the tray icon manager (Recommended for normal use)
 - Run from your terminal with: `aw-qt`
2. Start each module separately (Recommended for developing)
 - Run from your terminal with: `aw-server`, `aw-watcher-afk`, and `aw-watcher-window`

Both methods take the `--testing` flag as a command line parameter to run in testing mode. This runs the server on a different port (5666) and uses a separate database file to avoid mixing your important data with your testing data.

Now everything should be running! Check out the web UI at <http://localhost:5600/>

If anything doesn't work, let us know!

Note: On Linux, if you want to run from source using a `.desktop` file launcher, see [issue #176](#).

17.6 Updating from source

First pull the latest version of the repo with `git pull` then get the updated submodules with `git submodule update --init --recursive`. All that's needed then is a `make build`.

If it doesn't work, you can first try to run `make uninstall` and then do a fresh `make build`. If that fails as well, remove the `virtualenv` and start over.

Please report all issues you might have so we can make things easier for future users.

17.7 Packaging your changes

If you made some changes and want to create a proper build with portable executables (like normal ActivityWatch releases) you need to install `pyinstaller` (and on Debian-like distros `python3-dev`).

```
apt install python3-dev # Or equivalent for your Linux distribution
pip3 install --user pyinstaller
```

Then simply run the following to package it:

```
make package
```

When the packaging is done you will have `./dist` folder where you can find a zipped version and an unzipped `activitywatch` folder, you can move or copy that folder anywhere you need and set `aw-qt` to run from startup.

Installing from source (on Windows)

This was a guide hastily written by [@ErikBjare](#) when he had to build on Windows once, it is not complete.

- Install Git for Windows (including Git Bash)
- Install MinGW
- Rename C:/MinGW/mingw-make.exe to C:/MinGW/make.exe
 - `cp C:\\MinGW\\mingw32-make.exe C:\\MinGW\\make.exe`
- Set PATH to use MinGW
 - `SET PATH=C:\\MinGW\\bin;%PATH%`
- Install Python 3.5.4
- Install PyInstaller
 - `pip install pyinstaller`
 - Add PyInstaller script to PATH: `SET PATH=C:\\Users\\User\\AppData\\Roaming\\Python\\Python35\\Scripts`

ActivityWatch uses a REST API for all communication between aw-server and clients. Most applications should never use HTTP directly but should instead use the client libraries available. If no such library yet exists for a given language, this document is meant to provide enough specification to create one.

Warning: The API is currently under development, and is subject to change. It will be documented in better detail when first version has been frozen.

Note: Part of the documentation might be outdated, you can get up-to-date API documentation in the API browser available from the web UI of your aw-server instance.

19.1 REST Security

Note: Our current security consists only of not allowing non-localhost connections, this is likely to be the case for quite a while.

Clients might in the future be able to have read-only or append-only access to buckets, providing additional security and preventing compromised clients from being able to cause a severe security breach. All clients will probably also encrypt data in transit.

19.2 REST Reference

Note: This reference is highly incomplete. For an interactive view of the API, try out the API playground running on your local server at: <http://localhost:5600/api/>

19.2.1 Buckets API

The most common API used by ActivityWatch clients is the API providing read and append access buckets. Buckets are data containers used to group data together which shares some metadata (such as client type, hostname or location).

Get Bucket Metadata

Will return 404 if bucket does not exist

```
GET /api/0/buckets/<bucket_id>
```

List

```
GET /api/0/buckets/
```

Create

Will return 304 if bucket already exists

```
POST /api/0/buckets/<bucket_id>
```

19.2.2 Events API

The most common API used by ActivityWatch clients is the API providing read and append *Events* to buckets. Buckets are data containers used to group data together which shares some metadata (such as client type, hostname or location).

Get events

```
GET /api/0/buckets/<bucket_id>/events
```

Create event

```
POST /api/0/buckets/<bucket_id>/events
```

19.2.3 Heartbeat API

The heartbeat API is one of the most useful endpoints for writing watchers.

```
POST /api/0/buckets/<bucket_id>/heartbeat
```

19.2.4 Query API

TODO: Add link to writing queries once that page is done

20.1 v0.9.0

Released 2020-03-15

- add Rust server as an option. It has better performance and will be the main supported server soon.
- Almost all python components have been moved to poetry / pyproject.toml from old pip / setup.py building.

main activitywatch

- Add aw-server-rust binary to pyinstaller bundles (a567003)
- Automatic stale issue detection / closing (issues marked as stale after 180 days of inactivity, closed after 14 stale days)
- Added aw-server-rust
- Add .app and .dmg building

aw-client:

- switched to using pyproject/poetry (d66f319)
- Change default limit from 100 to return all by default (d35449a)

aw-core:

- added strict-rfc3339 (c5e9d34)
- upgraded jsonschema (a803443)

aw-qt

- Use manager to stop all modules (ac40452)
- Remove confirm dialog on exit (246ff05)
- Remove startup notification (#45) (b1f66c3)
- added aw-server-rust to possible modules (b4e9411)

- Add dark mode detection to MacOS logo logic (502b77d)

aw-server

- migrated from flask-restplus to flask-restx, updated submodule (36e970e)
- Makefile improvements: use pip3 (47a208b), fixed make build DEV=true (817723b)

aw-server-rust:

- Added in this release.

aw-watcher-window

- readed script that disappeared in previous PR (1e424e4)
- revise exclude_title logic (b1bd00d)

aw-webui:

- added ability to edit start and end in EventEditor, added ability to edit events from Bucket view (10dc80f)
- Add vuex support for editor activity (47fed8b)
- added setting for 'show last' duration default value (20ca6fa)
- buckets: Fix broken export bucket button (0a60cf7)
- fixed long lines (c502787)
- fixed bugs in event editor and stopwatch (d121fcb)
- fixed editing events in bucket view (b6c1405)
- stopwatch: Fixes broken "stop" and "edit" (06be96b)
- Fix firefox developer edition browser data not showing up in activity view. (4a05131)
- views/activity: Fix so subview is kept between date changes (c0ae191)
- Improvements to Stopwatch and EventEditor (#166) (2040a92)
- package.json: Update to fix vulnerabilities (a2f48a8)
- Swich from timeline-plus to vis-timeline (98cd372)
- queries: Removed trailing lines (d25cf09)

20.2 v0.8.4

Released 2019-12-10

- Adds the ability to filter the activity view by category.
- Fixes a bug where setting rule for a category to 'None' breaks the UI [issue #313](#).
- Fixes the broken favicon for aw-webui.

20.3 v0.8.3

Released 2019-11-13 (yes, the same day as v0.8.2)

- Fixes the Windows builds

20.4 v0.8.2

Released 2019-11-13

Just minor fixes to the v0.8.1 that was released the day before.

- Added a Windows installer
- Fixes the Windows builds (it did not)
- Fixes a bug in the web UI where weeks didn't always start on Mondays and months didn't always start on the first

20.5 v0.8.1

Released 2019-11-12

The v0.8 versions are finally leaving beta, and we celebrate this by giving you a truly awesome release!

It's the culmination of a lot of behind-the-scenes work that has been going on for quite a while. Several of our most requested features have made it into this release (categorization, better weekly/monthly visualizations, preparations for syncing), making it our best release yet!

Web UI:

- Categorization is finally here! Including visualizations and settings. The UX still leaves some things to be desired, but it's a great start.
- Daily and Summary views are now merged into one so that you get all the goodies of the Daily view but for arbitrary timeperiods like days/weeks/months!
- Updated the start page, including links to resources like the [user survey](#).

Server:

- New unique device ID is now exposed through the info API endpoint, a pre-requisite for building the much requested sync feature.
- Now contains the transforms needed for categorization.

Other:

- There is now a Windows installer available, and it automatically sets up autostart!

20.6 v0.8.0b9

Released 2019-07-03

Note: Changelog incomplete

Web UI:

- Now includes the Summary view for summarizing activity across weekly/monthly/yearly timeperiods!

20.7 v0.8.0b8

Released 2019-03-09

Note: Changelog incomplete

Server:

- Import and export APIs are now usable

Web UI:

- Added Stopwatch functionality
- Added ability to import buckets from export
- Bucket export button now does a full export that includes metadata

Other:

- The lowest version of Python supported for building ActivityWatch is now 3.6.
- Fixed PyInstaller-built releases on Windows

20.8 v0.8.0b7

Released 2018-11-03

Web UI:

- Fix broken editor bucket visualization

Misc:

- CI Improvements

20.9 v0.8.0b2 - v0.8.0b6

No changelog written.

20.10 v0.8.0b1

Released 2018-05-07

Server:

- New query2 API for querying and transforming data
- Added `version` field to `/info` endpoint
- Set stricter allowed CORS origins in testing mode
- Added `--cors-origins` CLI argument

Web UI:

- Added datepicker to the activity view

- Moved the today/clock visualization into the activity view
- New visualization for most-visited domains
- New visualization for previous days active time
- New query explorer
- Now displays version and hostname in bottom-right corner
- Now uses aw-client-js for all API calls

Watchers:

- Improved stability of client event queues ([see this PR](#))

Other:

- Windows: Console window and taskbar icon now hidden by default ([issue #139](#))
- All issues assigned to the v0.8 milestone can be found [on GitHub](#)

20.11 v0.7.1

Released 2017-11-06

- Actually fixed the timezone issue in the web UI ([issue #117](#)).
- All issues assigned to the v0.7 milestone can be found [on GitHub](#).

20.12 v0.7.0b4

Released 2017-10-22

- The ActivityWatch WebExtension is officially supported from this version forward, see the announcement [on the forum](#).
- (Not really, see v0.7.1) Fixed pesky timezone issue in web UI ([issue #117](#)).
- Fixed bug on macOS where keyboard activity would not be used to detect AFK state.
- Fixed packaging bugs (macOS, PyInstaller).
- The web extension now has a better look and notifies if connection to server failed.

20.13 v0.7.0b3

Released 2017-08-25

- Even more improvements to the web UI.
- Major improvements to the documentation, notably instructions on how to install from builds and sources.

20.14 v0.7.0b2

Released 2017-08-09

- Improvements to the web UI: a new visualization method (the “today” view) and information for users about the state of the project on the first page.

20.15 v0.7.0b1

Released 2017-06-14

There have been several major changes since v0.6. Much of it wont end up here but hopefully the major things will.

Note: If you are upgrading from a previous version, you might want to stop all loggers for the duration of your UTC offset to prevent issues which we’ve had difficulty debugging (or you can just start right away and expect your first hours to end up a bit weird).

- Now works on Windows.
- Working standalone packages. (edit: not reliable on all systems, but a lot easier to get running in many cases)
- All timestamps are now in UTC.
- Updated outdated parts of the documentation.
- Makefiles are now used throughout the projects to manage building, testing, and CI.
- A lot of bug fixes (and hopefully not too many new bugs).
- Vastly improved code quality.

20.16 v0.6.0 and older

We haven’t been keeping track of changes very well for older versions. Please refer to the git history.

CHAPTER 21

Indices and tables

- `genindex`
- `modindex`
- `search`

a

- aw_client, 30
- aw_core, 30
- aw_core.dirs, 30
- aw_core.log, 30
- aw_core.models, 30
- aw_query, 33
- aw_query.functions, 33
- aw_server, 36
- aw_server.api, 36
- aw_transform, 31

A

ActivityWatchClient (class in *aw_client*), 30
aw_client (module), 30
aw_core (module), 30
aw_core.dirs (module), 30
aw_core.log (module), 30
aw_core.models (module), 30
aw_query (module), 33
aw_query.functions (module), 33
aw_server (module), 36
aw_server.api (module), 36
aw_transform (module), 31

C

categorize() (in module *aw_transform*), 32
 chunk_events_by_key() (in module *aw_transform*), 33
 concat() (in module *aw_transform*), 31, 32
 connect() (*aw_client.ActivityWatchClient* method), 30
 create_bucket() (*aw_client.ActivityWatchClient* method), 30
 create_bucket() (*aw_server.api.ServerAPI* method), 36
 create_events() (*aw_server.api.ServerAPI* method), 36

D

data (*aw_core.models.Event* attribute), 30
 delete_bucket() (*aw_client.ActivityWatchClient* method), 31
 delete_bucket() (*aw_server.api.ServerAPI* method), 36
 delete_event() (*aw_server.api.ServerAPI* method), 36
 disconnect() (*aw_client.ActivityWatchClient* method), 31
 duration (*aw_core.models.Event* attribute), 30

E

ensure_path_exists() (in module *aw_core.dirs*),

30

Event (class in *aw_core.models*), 30
 export_all() (*aw_client.ActivityWatchClient* method), 31
 export_all() (*aw_server.api.ServerAPI* method), 36
 export_bucket() (*aw_client.ActivityWatchClient* method), 31
 export_bucket() (*aw_server.api.ServerAPI* method), 36

F

filter_keyvals() (in module *aw_transform*), 33
 filter_keyvals_regex() (in module *aw_transform*), 33
 filter_period_intersect() (in module *aw_transform*), 32
 flood() (in module *aw_transform*), 31

G

get_bucket_metadata() (*aw_server.api.ServerAPI* method), 36
 get_buckets() (*aw_client.ActivityWatchClient* method), 31
 get_buckets() (*aw_server.api.ServerAPI* method), 36
 get_cache_dir() (in module *aw_core.dirs*), 30
 get_config_dir() (in module *aw_core.dirs*), 30
 get_data_dir() (in module *aw_core.dirs*), 30
 get_eventcount() (*aw_client.ActivityWatchClient* method), 31
 get_eventcount() (*aw_server.api.ServerAPI* method), 36
 get_events() (*aw_client.ActivityWatchClient* method), 31
 get_events() (*aw_server.api.ServerAPI* method), 36
 get_info() (*aw_client.ActivityWatchClient* method), 31
 get_info() (*aw_server.api.ServerAPI* method), 36
 get_latest_log_file() (in module *aw_core.log*), 30

get_log() (*aw_server.api.ServerAPI method*), 36
 get_log_dir() (*in module aw_core.dirs*), 30
 get_log_file_path() (*in module aw_core.log*), 30

H

heartbeat() (*aw_client.ActivityWatchClient method*), 31
 heartbeat() (*aw_server.api.ServerAPI method*), 36
 heartbeat_merge() (*in module aw_transform*), 33
 heartbeat_reduce() (*in module aw_transform*), 33

I

id (*aw_core.models.Event attribute*), 30
 import_all() (*aw_server.api.ServerAPI method*), 37
 import_bucket() (*aw_client.ActivityWatchClient method*), 31
 import_bucket() (*aw_server.api.ServerAPI method*), 37
 insert_event() (*aw_client.ActivityWatchClient method*), 31
 insert_events() (*aw_client.ActivityWatchClient method*), 31

L

limit_events() (*in module aw_transform*), 33

M

match() (*aw_transform.Rule method*), 32
 merge_events_by_keys() (*in module aw_transform*), 33

P

period_union() (*in module aw_transform*), 32

Q

q2_categorize() (*in module aw_query.functions*), 33
 q2_chunk_events_by_key() (*in module aw_query.functions*), 33
 q2_concat() (*in module aw_query.functions*), 33
 q2_exclude_keyvals() (*in module aw_query.functions*), 33
 q2_filter_keyvals() (*in module aw_query.functions*), 33
 q2_filter_keyvals_regex() (*in module aw_query.functions*), 33
 q2_filter_period_intersect() (*in module aw_query.functions*), 33
 q2_find_bucket() (*in module aw_query.functions*), 34
 q2_flood() (*in module aw_query.functions*), 34
 q2_limit_events() (*in module aw_query.functions*), 34

q2_merge_events_by_keys() (*in module aw_query.functions*), 34
 q2_nop() (*in module aw_query.functions*), 34
 q2_period_union() (*in module aw_query.functions*), 34
 q2_query_bucket() (*in module aw_query.functions*), 35
 q2_query_bucket_eventcount() (*in module aw_query.functions*), 35
 q2_simplify_window_titles() (*in module aw_query.functions*), 35
 q2_sort_by_duration() (*in module aw_query.functions*), 35
 q2_sort_by_timestamp() (*in module aw_query.functions*), 35
 q2_split_url_events() (*in module aw_query.functions*), 35
 q2_sum_durations() (*in module aw_query.functions*), 35
 q2_tag() (*in module aw_query.functions*), 35
 query() (*aw_client.ActivityWatchClient method*), 31
 query2() (*aw_server.api.ServerAPI method*), 37

R

Rule (*class in aw_transform*), 32

S

send_event() (*aw_client.ActivityWatchClient method*), 31
 send_events() (*aw_client.ActivityWatchClient method*), 31
 ServerAPI (*class in aw_server.api*), 36
 setup_bucket() (*aw_client.ActivityWatchClient method*), 31
 setup_logging() (*in module aw_core.log*), 30
 simplify_string() (*in module aw_transform*), 33
 sort_by_duration() (*in module aw_transform*), 32
 sort_by_timestamp() (*in module aw_transform*), 32
 split_url_events() (*in module aw_transform*), 33
 sum_durations() (*in module aw_transform*), 32

T

tag() (*in module aw_transform*), 32
 timestamp (*aw_core.models.Event attribute*), 30
 to_json_dict() (*aw_core.models.Event method*), 30
 to_json_str() (*aw_core.models.Event method*), 30

U

union() (*in module aw_transform*), 32