

Every flight you save is written as a CSV file, a plain text table you can open in any spreadsheet program or read with your own code. This page describes that format in detail. It covers a Revision 4 or later board, which records the full set of motion data; a Revision 3 or earlier board uses the same layout with the motion columns left out, and there is a note on the differences at the end.

The whole file, the header and every reading, is the cleaned True Path result, and it is signed so it can be checked for tampering. There is more on those on the True Path and apogee page and the security page.

TIP

The easiest way to look at a flight is to upload the file straight to Altimeter Cloud. You do not need to make sense of the raw CSV yourself: the website lays your flight out in clear charts and figures that are simple to read. This page is really for anyone who wants to work with the raw data directly. There is a full walkthrough on the [Altimeter Cloud](#) page.

How the file is laid out

The file has three parts, in this order:

First comes a header of fourteen rows holding a summary of the flight and the exact conditions and settings it was recorded under. Then a single blank row. Then one row of column names, followed by the flight data itself, one row per recorded sample. In other words the header takes up sixteen rows, and the readings start on row seventeen.

The header is arranged as label and value pairs sitting side by side in columns, so it reads tidily when opened in a spreadsheet: a pair on the left, a pair in the middle, and on Rev4 and later a further pair on the right carrying the motion related figures. If you are writing a parser, the simplest approach is to read the header rows for the fields you want, skip the blank row, and treat everything from the column names onward as the data table.

The header

The header carries everything you need to understand the flight without reading a single data row. The fields fall into a few groups.

The **flight summary** gives the apogee in metres and the time it happened, the landing height and time, the maximum upward velocity, the landing velocity, and the average samples per second. The **identification** fields are your device tag, your competitor tag, and the altimeter's serial number, which are the details you use to add a flight to your account at altimetercloud.com.

The **conditions and settings** cover the sea level pressure and static temperature used to work out height, whether sensor temperature was used, the board temperature at the start and end of the flight, the battery voltage at the start and end, the firmware version, the board revision, and a number of individual settings such as orientation, launch protection, hybrid mode, sample speed and recording stop. There is also a single *Settings string* field that packs the complete configuration used for the flight into one line, and an *Adjusted log* field that reads "no" for an original flight and tells you if the log has since been re-derived with different pressure or temperature.

Because the Settings string is written in shorthand, here is what each key holds. The configurable ones match the options described on the settings page.

Key	Holds
fp	Sea level pressure used to work out height (hPa)
ft	Static temperature used (°C)

ut	Whether sensor temperature was used (0 or 1)
ss	Sample speed (Hz)
sr	Sample ratio
ms	Maximum samples
ld	Launch detect threshold (m)
lp	Launch protection (mG)
sl	Startup lock
rs	Recording stop
mp	Maintain power
hm	Hybrid mode
or	Orientation
se	Sync enable (Rev4 and later)
os, ii, kf	Internal pressure handling settings (oversampling and filter parameters)
rcp	The ground reference captured at launch
rmi	Launch time, in milliseconds
cgx, cgy, cgz	Stored gyro calibration offsets (Rev4 and later)
cax, cay, caz	Stored accelerometer calibration offsets (Rev4 and later)

Finally the **verification hash** is the HMAC-SHA256 signature that lets Altimeter Cloud confirm the file has not been altered. On Rev4 and later the right hand side of the header adds a set of **motion statistics**: burnout and ejection times, burnout and descent velocities, the maximum and average accelerations during ascent, descent and the burn, the launch pitch, roll, yaw and tilt, the dominant axis, and timing details.

! IMPORTANT

Do not edit a flight log if you want to keep it valid. The verification signature is worked out from the exact contents of the file, so changing anything at all, even a single number or an added space, means it no longer matches. Altimeter Cloud only accepts original, unedited logs, so an edited file will fail to verify and cannot be uploaded. If you want to work with the data, copy the file first and edit the copy, and keep the original safe.

The data columns (Revision 4 and later)

After the blank row and the row of column names, each row is one recorded sample. The columns appear in this order:

Column	Units	Format	What it is
Time(ms)	ms	whole numbers	Time from launch, with launch as zero
Altitude(m)	m	3 decimals	Height above the launch point
Velocity(m/s)	m/s	2 decimals	Vertical speed, worked out from the height trace
acceleration_x(mG)	mG	whole numbers	Acceleration along the board's X axis
acceleration_y(mG)	mG	whole numbers	Acceleration along the Y axis

acceleration_z(mG)	mG	whole numbers	Acceleration along the Z axis
acceleration_total(mG)	mG	1 decimal	Overall acceleration magnitude with the 1 G of gravity taken out, so it sits near zero at rest
gyro_x	°/s	whole numbers	Rotation rate about the X axis
gyro_y	°/s	whole numbers	Rotation rate about the Y axis
gyro_z	°/s	whole numbers	Rotation rate about the Z axis
Board temp(C)	°C	2 decimals	Sensor temperature, written about once a second and left blank on the rows in between
pitch	degrees	2 decimals	Orientation angle from the fusion filter
roll	degrees	2 decimals	Orientation angle from the fusion filter
yaw	degrees	2 decimals	Orientation angle from the fusion filter
tilt	degrees	2 decimals	Angle away from vertical

A couple of things are worth knowing. Time is measured from launch, so launch is zero; because the Nano keeps a short buffer of samples from just before lift off, the first few rows can show small negative times. The acceleration and gyro axes are the board's own axes, which line up with the rocket according to your orientation setting. There is one row per sample up to the maximum you set, so a full log can be tens of thousands of rows.

A short example using real flight data, showing the column names and a few rows from around the moment of launch:

```
Time(ms),Altitude(m),Velocity(m/s),acceleration_x(mG),acceleration_y(mG),acceleration_z(mG),acceleration_total(mG),gyro_x,gyro_y,gyro_z,Board temp(C),pitch,roll,yaw,tilt
-10,0.085,4.54,-993,35,76,-3.5,376,-12,-68,,3.40,-26.21,38.18,26.41,
0,0.141,5.76,-994,37,74,-2.6,330,-22,-74,,1.63,-26.97,35.28,27.02,
10,0.224,7.07,-992,36,75,-4.5,299,-28,-71,,0.06,-27.61,32.71,27.61,
20,0.334,8.44,-1376,99,120,384.8,256,-29,-68,-1.59,-28.15,30.55,28.19,
```

You can see the blank temperature cells, the small negative time just before launch, and that each data row finishes with a trailing comma, so a strict parser will see one empty field at the end of every data row.

Revision 3 and earlier

Boards without the motion sensor use the same file in the same shape, just simpler. The motion columns, that is the three acceleration columns, the total acceleration, the three gyro columns and the pitch, roll, yaw and tilt, are all absent, and the header does not carry the motion statistics on its right hand side.

There is one addition rather than a removal: these boards include a **Pressure(hPa)** column, written to six decimals, straight after Altitude. So a Revision 3 or earlier log has just five columns, in this order: Time(ms), Altitude(m), Pressure(hPa), Velocity(m/s) and Board temp(C). Everything else on this page, the header groups, the time being measured from launch and the temperature being written about once a second, applies in the same way.