

# Timelord

Timelords support the network by creating sequential proofs of time (using a Verifiable Delay Function) and broadcasting them approximately every nine seconds. This provides "deterministic randomness", which is used to decide the winning proofs of space.

Since this computation is sequential, very little energy is consumed, as opposed to proof-of-work systems, where computation is parallelizable. For example, if 100 timelords are doing the same computation on a proof of time, they will all create the exact same output.

A timelord is required to connect to exactly one beacon client, typically on the same machine. This connection is verified with a certificate. This 1:1 architecture has a large security benefit: it keeps the timelord sandboxed in its own private network. That way, the beacon client protocol is the only protocol that requires total security. If more than one beacon client could connect to the same timelord, it would add a potential attack vector to the network.

Timelords do not directly earn rewards. Furthermore, only the fastest timelord on the network will broadcast proofs at any given time. Therefore, only one timelord is required to keep the network running, and most farmers will not feel compelled to run one. However, farmers with multi-PiB farms may want to run a timelord, both for redundancy and for protection against temporary local latency issues.

If someone controls the fastest timelord in the world, it doesn't give them much of an advantage at winning rewards. However, they could potentially orphan or censor other farmers, depending on how much faster their timelord is.

Furthermore, an attacker with a significantly faster timelord than anyone else could potentially run a long-range attack against the network with less than 42.7% of the total netspace. For security purposes, it is very important to maintain open designs of VDF hardware.

## Types of Timelords

There are two primary types of Timelords: Regular and Blueboxes.

The first is the core Timelord that takes in Proofs of Space and uses a single fastest core to perform repeated squaring in a class group of unknown order as fast as possible. Beside each running VDF (referred to as a `vdf_client` in the application and source) is a set of proof generation threads that accumulate the proof that the time calculation's number of iterations was done correctly.

The second are Bluebox Timelords. Blueboxes are most any machine - especially things like old servers or gaming machines - that scour the historical chain looking for uncompressed proofs of time. So that the chain moves quickly, the regular Timelords use a faster method of generating proofs of time but the proofs are larger, which takes your Raspberry Pi a lot more time and effort to validate and sync the blockchain. A Bluebox picks up an uncompressed Proof of Time and recreates it, but this time with the slower and more compact proofs generated at the end. Those are then gossiped around to everyone so they can replace the large and slow to verify Proofs of Time with the compact and much quicker to validate version of exactly the same thing.

## Running a Timelord

The network only requires one running Timelord to keep moving (liveness.) The way Timelords race is like they are on a series of 100 meter dashes. Each one takes off with the last good Proof of Space and tries to get to the total number of iterations required to complete a given Proof of Space. Better Proofs of Space require less iterations to prove. When the fastest Timelord announces the Proof of Time for this Proof of Space all of the other Timelords stop racing and are magically teleported to the starting line of the next 100 meter dash to start it all over again.

It's good to have a few Timelords out there. There can be things like routing flaps or the overzealous backhoe that takes large swaths of the internet offline. If the fastest Timelord was just about to win the current dash when its internet blinked off in a fury of construction misadventure, then the second fastest will win that dash and the next dashes - until the fastest returns. One of the key qualities about Proofs of Time is that given the same Proof of Space, their output and proof are always the same (though the proofs can be larger or smaller and harder or easier to validate - they all end up with the same outcome.)

BPX developers plans to run a few Timelords around the world - and some backups too - to ensure that all Farmers and nodes can hear the beat that the Timelords are calling.

## Installing a Timelord

If you want to run a Timelord on Linux/macOS, first follow the [Install from Source](#) instructions here. Then run:

```
. ./activate
sh install-timelord.sh
bpx start timelord
```

Timelords execute sequential verifiable delay functions (proofs of time or VDFs), that get added to blocks to make them valid. This requires fast CPUs and a few cores per VDF.

Due to restrictions on how MSVC handles 128 bit numbers and how Python relies upon MSVC, it is not possible to build and run Timelords of all types on Windows.

## Regular Timelords

On MacOS x86\_64 and all Linux distributions, building a Timelord is as easy as running `bpx start timelord` in the virtual environment. You can also run `./vdf_bench square_asm 400000` once you've built Timelord to give you a sense of your optimal and unloaded ips. Each run of `vdf_bench` can be surprisingly variable and, in production, the actual ips you will obtain will usually be about 20% lower due to load of creating proofs. The default configuration for Timelords is good enough to just let you start it up. Set your log level to INFO and then grep for "Estimated IPS:" to get a sense of what actual ips your Timelord is achieving.

## Bluebox Timelords

Once you build the Timelord with `sh install-timelord.sh` in the virtual environment, you will need to make two changes to `~/bpxchain/beacon/config/config.yaml`. In the `timelord` section, set `bluebox_mode` to `True`. Then you need to proceed to the `beacon` section and set `send_uncompact_interval` to something greater than 0. We recommend 300 seconds there so that your Bluebox has some time to prove through a lot of the un-compacted Proofs of Time before the node drops more into its lap. The default settings may otherwise work but if the total effort is a little too much for whatever machine you are on you can also lower the `process_count`: from 3 to 2, or even 1, in the `timelord_launcher` section. You know it is working if you see `VDF Client: Sent proof` in your logs at INFO level.

# Timelords and Attacks

One of the things that is great about BPX consensus is that it makes it almost impossible for a Farmer with a maliciously faster Timelord to selfishly Farm. Due to the way the consensus works, a Farmer with a faster Timelord is basically compelled to prove time for all the farmers winning blocks around him also. Maliciously running a faster Timelord can give a benefit when attempting to 51% attack the network, so it is still important that over time we push the Timelord speeds as close to the maximum speeds of the silicon processes available. We expect to have the time and the resources to do that right and make open-source hardware versions widely available.

---

Revision #4

Created 5 June 2023 11:31:04 by Admin

Updated 7 November 2024 15:42:20 by Admin