

## BUIP126: Planet-on-a-LAN stress test model network

Submitted by: Jonathan Toomim

Date: 2019/5/17

### Summary

I'd like to set up a LAN-based test apparatus for simulating a planet-wide network of nodes, using the Linux `netem` module to add network delays, packet loss, and bandwidth caps as needed. This will provide a controlled environment in which developers can test out new code or new configurations quickly and efficiently, and rapidly collect and collate performance data.

The suggested nature of this network will be a server rack full of used servers, all interconnected via Gigabit Ethernet, operated in my datacenter in Moses Lake, WA on a separate LAN with a dedicated 100 Mbps connection to the internet. These servers will likely cost **around \$800** each after adding SSDs and HDDs, and will come with 16 to 40 CPU cores and roundabout 128 GB RAM. Each physical machine can run multiple nodes (either in separate VMs or just on different ports of the same machine). For around \$8,000, we can set up 10 servers and get around 40 to 100 network nodes.

Operating costs for a 10 machine setup will be about \$60/month for the 100 Mbps internet connection plus around \$80/month for electricity costs, plus some undetermined amount for labor from my employees. In contrast, renting 10 dedicated servers with this kind of HW specs would cost **around \$2,000/month** if we ordered it from standard cloud hosting providers. Owning will be cheaper than renting if this network is in operation for more than 4 months.

I expect to be able to keep this rig in operation at least until April 1st, 2020, at which time it will likely need to be relocated.

These machines can either be set up with one static global IP address per machine or behind NAT with forwarded ports for SSH and other services. I personally prefer the NAT/port forward concept, as I expect that will be cheaper, more scalable, and more easily relocated.

Experiments will probably usually use `regtest` mode. Use in `testnet` mode as part of an actual global network can also be done with machines in this network, but the 100 Mbps outbound pipe may become a bottleneck.

For `ssh`, non-simulation admin tasks, and performance data collection, we can set up a separate LAN (using the secondary Ethernet ports) without any artificial latency or packet loss.

This test setup is intended to be made available by all developers in the BCH community, not just BU developers. Experiments using a heterogenous mixture of node implementations will be allowed, as will experiments using a homogenous set of nodes that are not BU (e.g. `bchd` alone, or `ABC` alone). I expect BU's developers will make heavier use of it, though, as BU's developers tend to be more scaling-focused than the other implementations.

I will probably end up just buying this gear outright and setting it up with or without BU's financial support. However, if BU's membership wishes to reimburse me for the hardware costs by passing this BUIP, I will accept that support.

I am also interested in hearing how big our members think we should make this network. 10

machines? 5? 50? Small networks will outperform big ones, and many performance problems might not be apparent unless we have high per-node peer counts or high hop counts for total tx and block propagation paths. But also, no nodes no money. So. How much?

## **Budget**

I would recommend setting a budget of \$10k for this project to allow for some future expansion should we find the need for it (e.g. upgrading to 10 Gbps ethernet, or adding more servers).

Discussion of this project can happen in [https://t.me/BCH\\_stress\\_testnet](https://t.me/BCH_stress_testnet).

*edit 2019/8/15 solex:  
heading & budget*