

20180709 PingER Meeting on Blockchains

Time & date

12:00 noon 7/9/2018 at SLAC

Attendees:

Bebo White, Umar Kalim, Les Cottrell

Discussion

See also [Feasibility of a Blockchain model for the collection and distribution of PingER data](#), and <https://blockgeeks.com/guides/what-is-blockchain-technology/> for an introduction.

The idea is to reduce/eliminate the dependence on SLAC.

We identified two possibilities for PingER:

1. The PingER Oracle meta database of host coordinates (NODEDETAILS)
2. The actual raw measurement data. Typically found on the PingER Measurement Agent (MA) cached under /usr/local/share/pinger/data. This data is gathered on a daily basis (by ping_data.pl) from each active MA by SLAC and archived as /nfs/slac/g/net/pinger/pinger2/data/ping-<YYYY>-<MM>.txt. The data flow is described in [PingER data flow at SLAC](#). It is already publicly available via anonymous FTP.

1). NODEDETAILS

Enabling MAs to update NODEDETAILS independently would enable a richer sharing of both Beacons (kept in <Beacons> and target hosts (kept in <HostList>). Currently, only Beacons are shared. This sharing could be a big advantage to MAs such as SLAC, GZHU and UBRU which have large local <HostList>s. The amount of data in the database relatively small. There are about 3500 hosts in the database including active: ~ 127 Beacons, 40 MAs; 2200 Disabled (no longer active hosts). Each host has about 20 columns of information, each of which is up to 100Bytes long. As envisioned each snapshot would be a complete representation of the database. The database is only updated occasionally, e.g. say once a week on average so the number of snapshots is not large.

Maybe while one is learning about about Blockchains this might be a place to start. Once this is better understood then move on to item 2.

2). Raw measurement data

For the current data storage, just from 2016-01 thru 2018-06 there are about 32 GBytes or ~0.4GBytes/month or ~0.012GBytes/day. The data is updated on a daily basis.

- If each new snapshot is to be complete then this could get huge, e.g. if we keep the snapshots going back only for the most recent 24 months each snapshot is ~24 (months) * 0.4GBytes = ~10 GBytes and there are 2 * 365 (days in a year) of them, i.e. ~ 7TBytes/participating MA.
- If on the other hand, each snapshot is just the daily measurement from all the MAs, then each snapshot is ~0.012GBytes.
 - In this case, the analysis will need to add all these snapshots together. Some thought will be needed to figure out how to save and access the data.
- Another alternative would be for each MA (or maybe just a subset of MAs) to save just its data each 30 minutes into the blockchain. I.e. a transaction is a set of measurements made each 30 minutes by each MA. The amount of data to be saved varies depending on the number of targets. For each target, for each day, there are 96 measurements (each 30 minutes for 100 and 1000 Byte pings) of each target and each target measurement is ~200Bytes, or ~ 100*200Bytes / target/day or ~400Bytes/target/30 min measurement
 - SLAC MA (~900 targets): 18MBytes/day or ~ 400KBytes/target/30 minute measurement
 - Typical MA (just monitoring ~ 100 Beacons): 2MBytes or ~40KBytes/target/30 minute measurement
 - There are currently about 40 active MAs or about 80MBytes/day or 1.8MBytes/all targets/all typical MAs/30 minutes.
- Another alternative would be to save each set of ~10 pings to the ledger. I.e. a transaction is a set of roughly 10 nByte pings where n is 100 or 1000.
 - Each addition would be ~ 200Bytes
 - Since pinger2.pl make the measurements by multiple threads, it is possible for multiple measurement results being added to the blockchain simultaneously.
 - This would provide give real-time results but would require changes to the guts of pinger2.pl

If we take the model that each blockchain ledger entry is the result of a 30-minute measurement by an MA, and each MA make its own entry, then we have:

- Near real-time raw data for all MAs is current in the blockchain as of < 30 minutes is accessible publicly.
- So we will need:
 - A modification or a wrapper for the script that makes the measurements (pinger2.pl) to save the measured raw data to the blockchain ledger after each 30 min set of measurements are completed
 - provide an interface to a service that makes the data easily accessible

Questions

- How complicated is setting up a Blockchain
 - If the effort is too high then we may not have resources to implement
 - Do we have the resources?

- Take a look at <https://hyperledger.org/projects/fabric> - good students should be able to handle this - since it's an open source project, there should be a good resource for questions/support;
 - Also see: <https://www.ccn.com/google-cloud-launches-blockchain-toolkit-for-app-developers/>
- The transition cost could be large:
 - we would need to demonstrate how access to blockchain content is accomplished - clearly adding data to a blockchain is only half the model;
- Do all MAs have to participate,
 - this is probably not practical since many MAs have little or no resources for this type of effort/transition.
 - If only say Gzhu, Ubru and SLAC MAs participate, is this sufficient redundancy?
 - On the otherhand maybe we could provide a blockchain script as part of the pinger2.pl measurement script, so each MA would automatically save the data in the blockchain each 30 minutes.
- It might simplify the data deployment
- Tying PingER to Blockchain could increase the interest and resources in the PingER project
 - Blockchains are a hot topic today.
 - Stanford has created a Blockchain Institute, see <https://www.bitrates.com/news/p/stanford-has-announced-their-new-world-class-center-for-blockchain-research>
 - Many universities are pursuing Blockchain see <https://www.accounting-degree.org/college-cryptocurrency-blockchain-courses/>
 - Get students interested and writing papers
 - get access to MS and PhD students
- What is Saqib's situation (Saqib can you weigh in here):
 - Duration at Gzhu: Post-doc finishes Feb 2019
 - Access to students to work on blockchain for PingER: He does not have any students. However, he will try and find someone interested in blockchain.
 - Interest in working on Items 1 or 2 (or both) above? He is interested in both items.
- How to transition from today's centralized on SLAC to a more distributed Blockchain implementation
 - Will need to continue current PingER while new Blockchain implementation is being developed, made robust and complete
 - Will need web interfaces to the data and new mechanisms
- What about the analysis, presentation?
- How long does it take to validate a transaction? This may be important if each transaction contains the results from a set of 10xnByte pings, and each transaction needs validating to save the block. For SLAC, in 30 minutes, the number of transactions would be ~1600 and today it takes about 20 minutes.
 - From <https://blockgeeks.com/guides/what-is-blockchain-technology> it appears that the bitcoin network reconciles every transaction that happens in ten-minute intervals. Presumably, the results are not available until the transactions are validated. Thus if we were to validate say every 30 minutes, the individual sets of 10 pings are anyway not available until the 30 minutes are over and the validation completed. So it appears there is no advantage in the complexity of individually adding each ~10 nByte ping result to the blockchain. Instead simply make the complete set of ping measurements made each 30 minutes by each MA into a transaction.
 - Bebo suggests using permissioned networks see https://monax.io/learn/permissioned_blockchains/ since unpermissioned blockchain networks are not very performant and are public spaces, that are slow to innovate.
 - How is the permissioning achieved for first class citizens in a permissioned network? It says it is out of band and hints at using VPN, or possibly with public/private keys? Presumably this information (public key?) has to be shared somehow. Typically how?
 - Also for a permissioned blockchain network is the validation of a transaction before it is put into a blockchain just the out-of-band information and thus very fast? Or is some consensus required? What is the consensus, how is it achieved?
 - Do we start off with say 3 nodes (SLAC, GZHU and possibly UBRU)? Is this enough nodes for consensus.