

TULIP Landmarks Selection of Tier=0 hosts


Tier=0 landmarks are the ones that are used for the initial estimate of what region a target is located in. There are always one or two tier1 landmarks per major region that can be used for region identification. This idea is to first find roughly the regions of the target is in by launching a few network ping probes from just the tier=0 landmarks. Not only does this reduce network traffic compared to pinging the target from all landmarks, but it is hoped to be faster since once we know the region we can choose a small set of nearby landmarks to make more accurate measurements.

Selection of the tier=0 landmarks is based on several criteria:

- How many other landmarks it can see
- How well the RTT correlates with the distance to other known targets
- How well the RTT from the tier1 landmark separates targets in the same region from those in other regions.

Europe:


Thus for Europe if we choose a landmark roughly in the middle of Europe we may be able to define an RTT circle that tells us whether the host is in the European region or not. The chart below shows the RTT and distance from the `pinger.cern.ch` host in Geneva Switzerland to other hosts with known locations both in (green) and outside Europe (pink).



Inspection of the chart ([see here for the spreadsheet](#)) shows that for these hosts at least there is a clear dividing line at around an RTT of 50ms. For RTTs < 50ms the target is in Europe, for RTTs > 50 ms the target is elsewhere. The dividing line for distance is about 2000km. Using this metric Algeria falls inside Europe.

N. America

A similar chart ([the spread sheet can be found here](#)) is shown below for a landmark in Chicago making ping RTT measurements to hosts with known locations around the world.



This time the gap is not so clear, however there is a definite gap such that targets within 120ms are in N.America and those >120ms are elsewhere. The North America target at 119ms is in Vancouver Canada. The [route](#) goes via Chicago, Kansas City, Boise, Denver, Hanford (Washington). Then it goes through 207.231.240.21 to Vancouver, Canada. According to Geotool 207.231.240.21 is in Los Angeles. However data based location methods are typically poor for routers so the confidence in this is poor. At the same time 207.231.240.21 does not respond to pings so we are unsure of its location. At the cost of increasing the tier1 regional estimation traffic, it is possible we could have a clearer separation using more tier1 landmarks in N. America. This may deserve more exploration

Reliability selection

Look at the analyzed logs using <http://www-wanmon.slac.stanford.edu/cgi-wrap/reflector.cgi?function=analyze&days=5>. Copy and paste them to a file. Select out the success table and import into Excel. Then sort by descending success rate ([example](#)).

One wants to choose 2 reliable tier 0 landmarks for each region. THis can be done by looking at the above spreadsheet.

You can find out the current tier = 0 landmarks using <http://www-wanmon.slac.stanford.edu/cgi-wrap/reflex.cgi?target=www.cern.ch>.