From: Andrew Lewman, Executive Director
To: the tor community
Date: April 10, 2011

This report documents progress in February 2011.

New releases, new hires, new funding

We contracted Runa Sandvik to work on moving the torouter https://trac.torproject.org/projects/tor/wiki/TheOnionRouter/Torouter project forward, translations, integration of tor web server log analysis.

New Releases

1. On February 23rd, we released an updated Tor -stable. Tor 0.2.1.30 fixes a variety of less critical bugs. The main other change is a slight tweak to Tor's TLS handshake that makes relays and bridges that run this new version reachable from Iran again. We don't expect this tweak will win the arms race long-term, but it buys us time until we roll out a better solution.

```
Changes in version 0.2.1.30
  o Major bugfixes:
    - Stop sending a CLOCK_SKEW controller status event whenever
      we fetch directory information from a relay that has a wrong clock.
      Instead, only inform the controller when it's a trusted authority
      that claims our clock is wrong. Bugfix on 0.1.2.6-alpha; fixes
      the rest of bug 1074.
    - Fix a bounds-checking error that could allow an attacker to
      remotely crash a directory authority. Bugfix on 0.2.1.5-alpha.
      Found by "piebeer".
    - If relays set RelayBandwidthBurst but not RelayBandwidthRate,
      Tor would ignore their RelayBandwidthBurst setting,
      potentially using more bandwidth than expected. Bugfix on
      0.2.0.1-alpha. Reported by Paul Wouters. Fixes bug 2470.
    - Ignore and warn if the user mistakenly sets "PublishServerDescriptor
      hidserv" in her torrc. The 'hidserv' argument never controlled
      publication of hidden service descriptors. Bugfix on 0.2.0.1-alpha.

  o Minor features:
    - Adjust our TLS Diffie-Hellman parameters to match those used by
      Apache's mod_ssl.
    - Update to the February 1 2011 Maxmind GeoLite Country database.

  o Minor bugfixes:
    - Check for and reject overly long directory certificates and
      directory tokens before they have a chance to hit any assertions.
      Bugfix on 0.2.1.28. Found by "doorss".
    - Bring the logic that gathers routerinfos and assesses the
      acceptability of circuits into line. This prevents a Tor OP from
      getting locked in a cycle of choosing its local OR as an exit for a
      path (due to a .exit request) and then rejecting the circuit because
      its OR is not listed yet. It also prevents Tor clients from using an
      OR running in the same instance as an exit (due to a .exit request)
      if the OR does not meet the same requirements expected of an OR
      running elsewhere. Fixes bug 1859; bugfix on 0.1.0.1-rc.

  o Packaging changes:
    - Stop shipping the Tor specs files and development proposal documents
      in the tarball. They are now in a separate git repository at
      git://git.torproject.org/torspec.git
    - Do not include Git version tags as though they are SVN tags when
      generating a tarball from inside a repository that has switched
      between branches. Bugfix on 0.2.1.15-rc; fixes bug 2402.
```

Design, develop, and implement enhancements that make Tor a better tool for users in censored countries.

- Arm development has stayed relatively on track, with the revised connection panel very nearly achieving parity with its predecessor (and in most respects surpassing it). Most of what remains are refinements and tasty new features. Arm has also been added to Debian (Sid) and Ubuntu (Natty) with backports pending. Many thanks to Peter for his help.

- Tom spent some time assisting Jacob with a satellite test. The test wound up breaking due to flaky hardware, however they were able to collect some usable data.

- Created the trac ticket around hidden service improvements, https://trac.torproject.org/projects/tor/ticket/2552 We need to focus on improving hidden services and fixing some of the performance and reliability issues within.

- Mike fixed a bunch of torbutton bugs. His summary iteration results are at https://trac.torproject.org/projects/tor/ticket/2591.

- Mike helped fix the bandwidth authority on salsa that exploded due to a reinstall.

Architecture and technical design docs for Tor enhancements related to blocking-resistance.

- Karsten and Sebastian tried to improve the database schema in metrics-db to speed up relay search performance. Unfortunately, the required updates from the old schema took forever, so we don't just need a better schema, but also a better migration strategy to go from one schema to the next.

- Karsten started moving code from metrics-db to metrics-web to make the metrics website a self-contained unit that's independent of aggregating descriptors. The idea is that people can take the metrics-web code and improve it or replace it with a better metrics website written in the web language of their choice.
- Karsten started working on better visualizations of Tor data using the Thematic Mapping API together with Rachel Binx.

Hide Tor's network signature.

- Collaborated with George K on obfsproxy, a generic protocol obfuscator. It seems to work ok but needs more testing.
- Nick worked on improving the pluggable-transport design.
- Jacob did another revision on what is now prop 179, `https://gitweb.torproject.org/torspec.git/blob/HEAD:/proposals/179-TLS-cert-and-parameter-normalization.txt`
- Jacob looked at the EFF SSL data and have some improvements for how we can get better data for future research questions.

Grow the Tor network and user base. Outreach.

Measures of the Tor Network

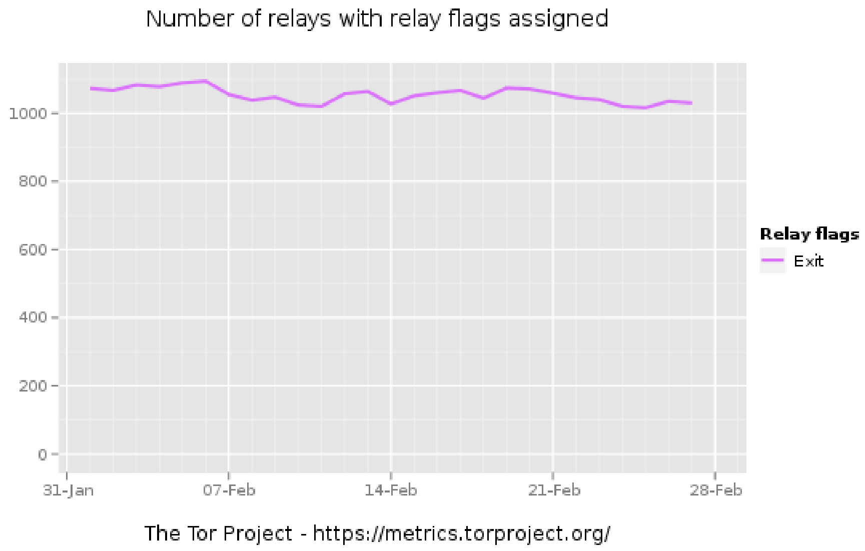

This graph shows the total quantity of exit relays in February 2011. We seem to have kept most of our relays since the bump due to Tunisia and Egypt.

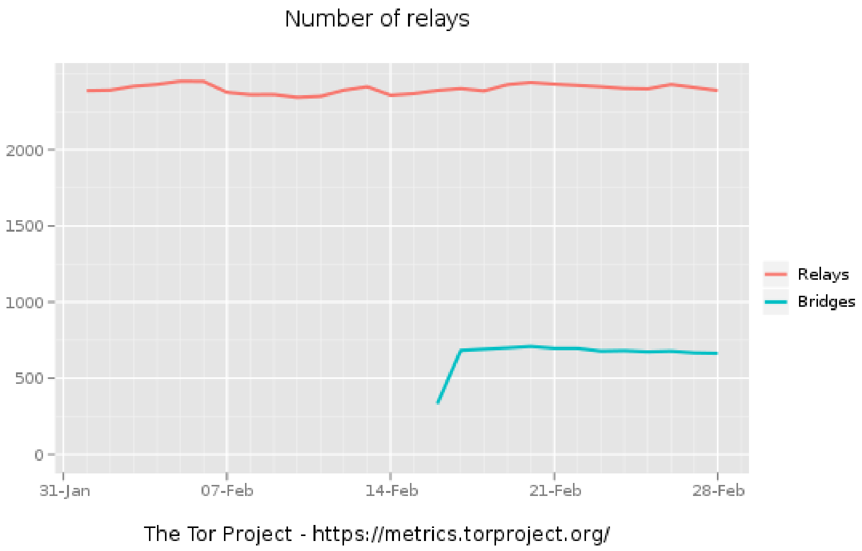

This graph shows the total quantity of relays and the total quantity of bridges in February 2011. We seem to have kept most of our relays since the bump due to Tunisia and Egypt. Due to a data collection error in February, we're missing two weeks of data for bridges.

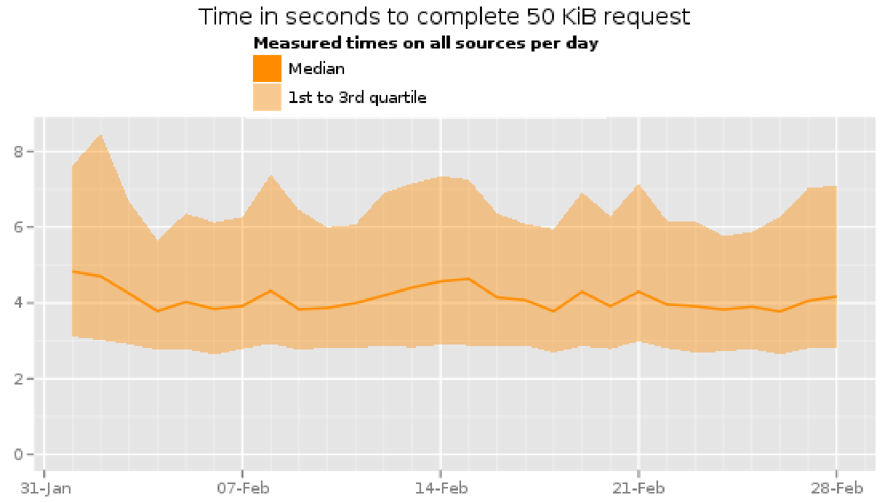

The Tor Project - https://metrics.torproject.org/

This graphs shows how many seconds it took to complete a 50KB download from a standard Tor client. This is an average of all measurements from servers located in Illinois, Massachusetts, and Sweden. Performance remains relatively steady at 4 seconds.

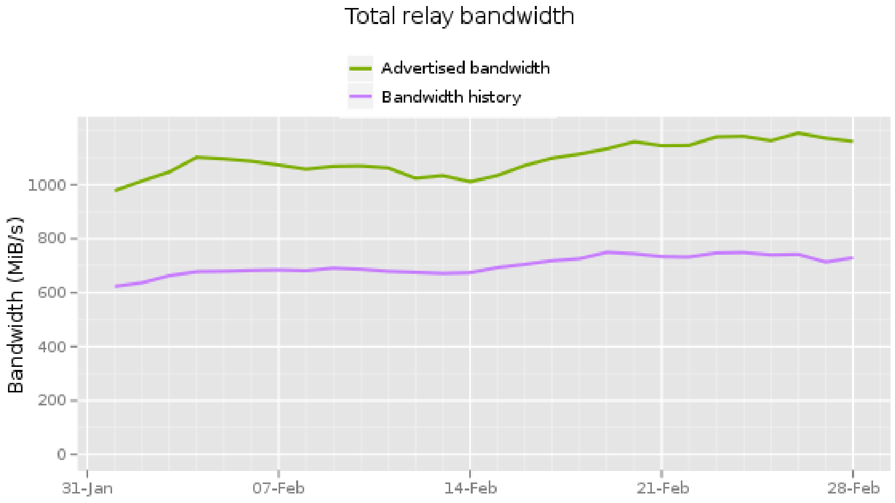

The Tor Project - https://metrics.torproject.org/

This graph shows the total available bandwidth available to clients and how much was actually used throughout the month. The influx of relays from last month creates almost 1.4GBps (11.2 Gbps) of bandwidth available.

Outreach and Advocacy

1. Jacob continued working on Egypt related issues.
2. Jacob did a training for people in Bahrain.
3. Tor, the ACLU and the OPC launched our privacy challenge: http://www.develop4privacy.org/.
4. Jacob did a bit of looking at the Libyan Internet.
5. Jacob gave the keynote speech at the Ctrl-X-Ethics Workshop in Toronto on ethics of security research.
6. We ran a successful hackfest with the help of MITs Center for Future Civic Media, https://blog.torproject.org/blog/tor-open-hackfest-february-19-2011 and the followup at https://blog.torproject.org/blog/hackfest-thanks.
7. Roger was the keynote speaker for Workshop on Free and Open Communication on the Internet (FOCI), http://www.gtisc.gatech.edu/foci.html.
8. Andrew talked to the Wesleyan HFOSS team about Tor and classwork for their Summer 2011 session. http://hfoss.org/.
9. Roger and Steven presented at Financial Crypto and the Workshop on the Ethics of Computer Security Research.
10. Andrew spoke at a few panels under Chatham House Rules. He published his speech notes as a blog post, https://blog.torproject.org/blog/five-minutes-speak.

Roger wrote a blog post about using our data archive as input to new safety metrics: https://blog.torproject.org/blog/research-problem-measuring-safety-tor-network.

Roger talked to the Philly FBI for the Philadelphia Infragard chapter about Tor and anonymity online.

Roger taught a Tor lecture for Drexel's security class.

Andrew was interviewed by Discovery News about Tor's role in the unrest in Tunisia and Egypt, http://news.discovery.com/tech/egypt-internet-online-protesters-110201.html

Andrew was interviewed by the Walpole Times about Tor and what we do, http://www.wickedlocal.com/walpole/news/x95296113/Tor-Project-a-Walpole-based-company-helps-Egyptians-avo

Damian started thinking about our various projects in a more streamlined and easy-to-understand way. The results are at https://www.torproject.org/getinvolved/volunteer.html.en#Projects.

Preconfigured privacy (circumvention) bundles for USB or LiveCD.

- Jacob did some testing of Gibberbot's Tor and OTR integration. Gibberbot is an XMPP chat client for Android designed to work over Tor.

- Jacob did a bunch of work on ttdnsd - some important (but not critical) bug fixes and he's planning on pushing out a release in the future. Jacob and Robert did some work on torsocks integration and in the process hammered out a reasonable torsocks API for people who want to have auto-magically Torified sockets without understanding Tor internals.

- Jacob worked on OpenWRT packaging issues - as well as other work on the Torouter project.

- Jacob worked on Tahoe (`http://tahoe-lafs.org/trac/tahoe-lafs`) and Tor related Hidden Service documentation; after moderate amount of Tor testing with Tahoe now and it seems to be partially functional.

Bridge relay and bridge authority work.

- Karsten prepared a patch for BridgeDB to export bridge pool assignments to a local file. This patch needs some cleanup before being deployed on BridgeDB.

- Karsten wrote a first draft of a BridgeDB specification that Nick commented on. The next step is to include Nick's comments and change the writing style, so that the specification describes what the current BridgeDB code does, not what a generic BridgeDB implemention should do.

- Karsten extended the bridge descriptor sanitizing algorithm to include IP address hashes in the sanitized descriptors. Sanitized all existing bridge descriptors using this new algorithm. Instead of 127.0.0.1, bridges now have 10.x.y.z addresses with x.y.z being stable for a given bridge fingerprint in a given month. This allows analyses of how often bridges change their IP addresses in a given month.

- Christian deployed a new version of BridgeDB, the one that's i18n enabled (1613) and also can dump bridge pool assignments to files. We can now assign an amount of unassigned bridges to someone/something and dump them to file buckets. See 1612 for more infos. In theory, we can now have an amount of Twitter assigned bridges that we pump out over Twitter.

- Christian also started writing a python script that is able to dump stuff to Twitter.

- After deployment of the new BridgeDB, some issues came up that were fixed (2556 and others). It seems to run smoothly now. We'll be even more happy about it when we have important (read: Chinese and Farsi) translations ready and deployed.

- Christian and Karsten discussed about whether his planned "dump bridge pool assignments to files" feature can use the bucket mechanism of 1612. Turns out it can't since both have a different set of goals and would be to painful to sync with every change.

- Mike helped Karsten with improving the output of Torperf for future experiments involving circuit build timeouts.

Scalability, load balancing, directory overhead, efficiency.

- Improved Torperf and finally deployed it to collect data about used paths and to measure performance with custom guard node selections. This is still work in progress together with Mike and Tom as part of our first Scrum iteration that ends on March 5.
- Worked on Florian's and Björn's token bucket patch some more together with Sebastian. The current state of the patch is that it needs some more love before it can be merged into 0.2.3.x.
- Nick collaborated a little with two volunteers on what we think at this point must be the 5th generation of a "launch a private network" tool. This one is called "chutney".
- Nick reviewed a bunch of patches, reviewed a bunch of bugs, fixed a bunch of bugs, merged many people's code, got 0.2.2.x closer to done.
- Sebastian wrote a proposal for a safer voting process for consensus parameters, and wrote an implementation for it. `https://gitweb.torproject.org/torspec.git/blob/HEAD:/proposals/178-param-voting.txt`

Incentives work.

1. Christian cleaned up the rather hackish installation of Weather on bahri. The stable installation now lives under '/home/weather/opt/current' and actually is update-able through 'git pull'. There's also a testing installation to test stuff and play around at `https://weather2.torproject.org/`. He's tried to update the documentation with all the stuff that is necessary to install and run Weather.
2. Christian tried looking into 2467. Some people complained that Weather didn't know their relay fingerprint. On Sebastian's and Mike's idea, Christian changed the torrc to include 'FetchDirInfoEarly 1' and 'FetchUselessDescriptors 1'. Since that no one complained again about Weather not knowing a certain relay (except for one time, when the Weather process had silently crashed and therefore the database wasn't updated for a day).
3. After Tor 0.2.1.30 was tagged and made it to the recommended versions', people running 0.2.1.29 started complaining about getting "Node out of date!" emails from Weather. It turned out that Weather was actually doing the right thing, namely mailing them that they were not running the latest recommended stable version anymore. No one seemed to have read the text near the checkbox in the signup process. After discussing this intensely with Sebastian, we decided to go for a more simple solution: People now get email when they don't run one of the recommended versions or a more recent dev version of Tor.

More reliable (e.g. split) download mechanism.

- Christian did a rather large GetTor overhaul. The way GetTor manages its packages is now much easier to understand and enhance. GetTor moved from a ini-style configuration file and parser to a more BridgeDB-like configuration management. Also, packages are now configured rather than hard coded. In addition, he cleaned up the i18n management of GetTor to something similar to what we use in BridgeDB. Not only are the translation strings cleaner now, but the translation and installation is smoother. Also, the logging was simplified because it had too many features that no one used and generally was polluting the log file with too much useless information. Furthermore, the MakeStat.py script that creates GetTor's package statistics was simplified a lot.

- Christian fixed 1586, users requesting non-existent split packages now are informed about that fact.
- Nick worked on a thandy packaging spec with erinn

Footprints from Tor Browser Bundle.

Nothing to report.

Translation work, ultimately a browser-based approach.

- Sebastian Started figuring out a way how translations can be pulled from transifex and used in their respective products in a more automated fashion.
- New or updated website translations in French, Russian, Italian, Japanese, Spanish, Mandarin Chinese, and Greek.