Monthly Progress report for April 2012

The Tor Project, Inc.

Contents

1 Research

1.1 Anonymous Communications

- On April 23rd, we released a new alpha version of Tor.

```
Tor 0.2.3.14-alpha fixes yet more bugs to get us closer to a release
candidate. It also dramatically speeds up AES: fast relays should
consider switching to the newer OpenSSL library.

Changes in version 0.2.3.14-alpha - 2012-04-23
  o Directory authority changes:
    - Change IP address for ides (v3 directory authority), and rename
      it to turtles.

  o Major bugfixes:
    - Avoid logging uninitialized data when unable to decode a hidden
      service descriptor cookie. Fixes bug 5647; bugfix on 0.2.1.5-alpha.
    - Avoid a client-side assertion failure when receiving an INTRODUCE2
      cell on a general purpose circuit. Fixes bug 5644; bugfix on
      0.2.1.6-alpha.
    - If authorities are unable to get a v2 consensus document from other
      directory authorities, they no longer fall back to fetching
      them from regular directory caches. Fixes bug 5635; bugfix on
      0.2.2.26-beta, where routers stopped downloading v2 consensus
      documents entirely.
    - When we start a Tor client with a normal consensus already cached,
      be willing to download a microdescriptor consensus. Fixes bug 4011;
      fix on 0.2.3.1-alpha.

  o Major features (performance):
    - When built to use OpenSSL 1.0.1, and built for an x86 or x86_64
      instruction set, take advantage of OpenSSL's AESNI, bitsliced, or
      vectorized AES implementations as appropriate. These can be much,
      much faster than other AES implementations.

  o Minor bugfixes (0.2.2.x and earlier):
    - Don't launch more than 10 service-side introduction-point circuits
      for a hidden service in five minutes. Previously, we would consider
      launching more introduction-point circuits if at least one second
      had passed without any introduction-point circuits failing. Fixes
      bug 4607; bugfix on 0.0.7pre1.
    - Change the BridgePassword feature (part of the "bridge community"
      design, which is not yet implemented) to use a time-independent
      comparison. The old behavior might have allowed an adversary
      to use timing to guess the BridgePassword value. Fixes bug 5543;
```

bugfix on 0.2.0.14-alpha.
- Enforce correct return behavior of tor_vsscanf() when the '%%' pattern is used. Fixes bug 5558. Bugfix on 0.2.1.13.
- When sending an HTTP/1.1 proxy request, include a Host header. Fixes bug 5593; bugfix on 0.2.2.1-alpha.
- Don't log that we have "decided to publish new relay descriptor" unless we are actually publishing a descriptor. Fixes bug 3942; bugfix on 0.2.2.28-beta.

o Minor bugfixes (0.2.3.x):
- Fix a bug where a bridge authority crashes (on a failed assert) if it has seen no directory requests when it's time to write statistics to disk. Fixes bug 5508. Bugfix on 0.2.3.6-alpha.
- Fix bug stomping on ORPort option NoListen and ignoring option NoAdvertise. Fixes bug 5151; bugfix on 0.2.3.9-alpha.
- In the testsuite, provide a large enough buffer in the tor_sscanf unit test. Otherwise we'd overrun that buffer and crash during the unit tests. Found by weasel. Fixes bug 5449; bugfix on 0.2.3.12-alpha.
- Make sure we create the keys directory if it doesn't exist and we're about to store the dynamic Diffie-Hellman parameters. Fixes bug 5572; bugfix on 0.2.3.13-alpha.
- Fix a small memory leak when trying to decode incorrect base16 authenticator during SAFECOOKIE authentication. Found by Coverity Scan. Fixes CID 507. Bugfix on 0.2.3.13-alpha.

o Minor features:
- Add more information to a log statement that might help track down bug 4091. If you're seeing "Bug: tor_addr_is_internal() called with a non-IP address" messages (or any Bug messages, for that matter!), please let us know about it.
- Relays now understand an IPv6 address when they get one from a directory server. Resolves ticket 4875.
- Resolve IPv6 addresses in bridge and entry statistics to country code "??" which means we at least count them. Resolves ticket 5053; improves on 0.2.3.9-alpha.
- Update to the April 3 2012 Maxmind GeoLite Country database.
- Begin a doc/state-contents.txt file to explain the contents of the Tor state file. Fixes bug 2987.

o Default torrc changes:
- Stop listing "socksport 9050" in torrc.sample. We open a socks port on 9050 by default anyway, so this should not change anything in practice.
- Stop mentioning the deprecated *ListenAddress options in

```
    torrc.sample. Fixes bug 5438.
  - Document unit of bandwidth related options in sample torrc.
    Fixes bug 5621.

o Removed features:
  - The "torify" script no longer supports the "tsocks" socksifier
    tool, since tsocks doesn't support DNS and UDP right for Tor.
    Everyone should be using torsocks instead. Fixes bugs 3530 and
    5180. Based on a patch by "ugh".

o Code refactoring:
  - Change the symmetric cipher interface so that creating and
    initializing a stream cipher are no longer separate functions.
  - Remove all internal support for unpadded RSA. We never used it, and
    it would be a bad idea to start.
```

- Roger believes "yes, we can get rid of stream-level congestion windows." The main result would be a) slightly improved performance for clients who do one thing at once, and b) unknown performance hits for clients who do two or more streams at once.

  The simulators need more work before we can answer this question well. Roger says he should help Rob with his CSET submission and move forward the various tickets on simulator fail.

  Roger hopes that removing stream-level congestion windows might be unnecessary because we will decide N23 is better.

1.2 Metrics

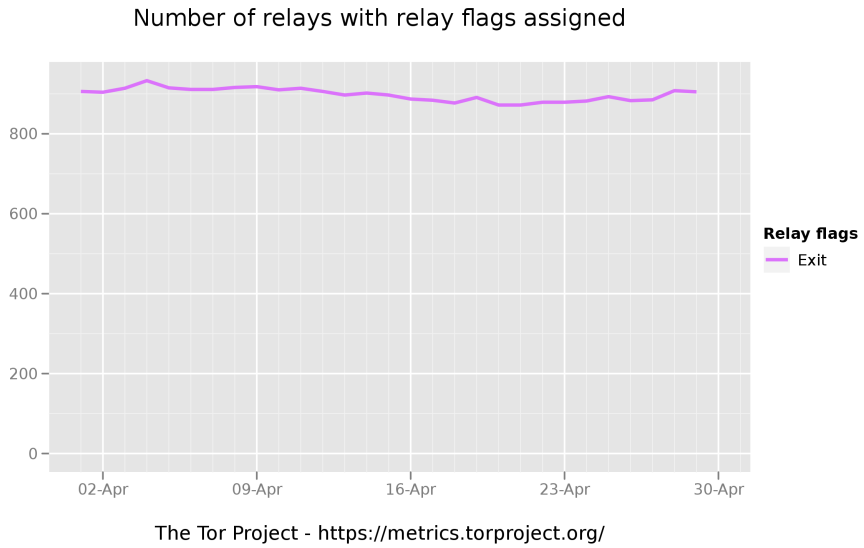

This graph shows the total quantity of exit relays in April 2012.

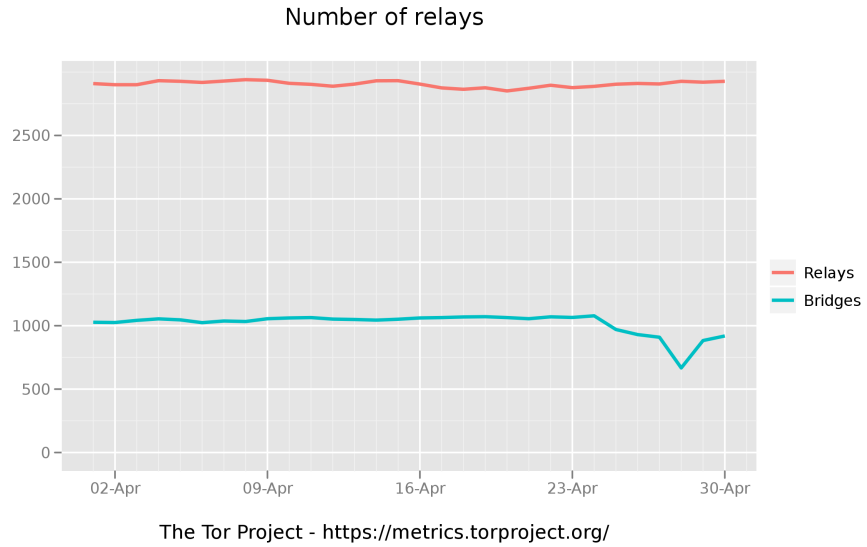

This graph shows the total quantity of relays and the total quantity of bridges in April 2012.

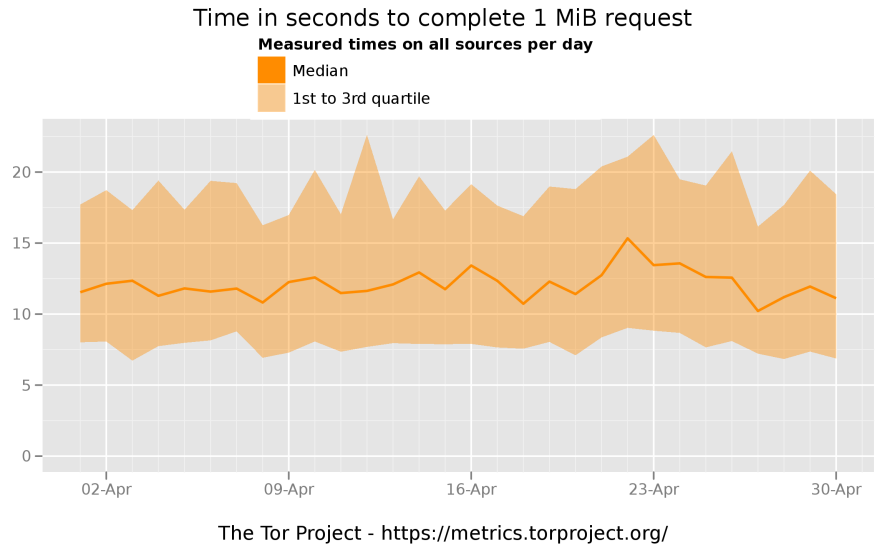

This graphs shows how many seconds it took to complete a 1 megabyte download from a standard Tor client. This is an average of all measurements from servers located in Illinois, Massachusetts, and Sweden.

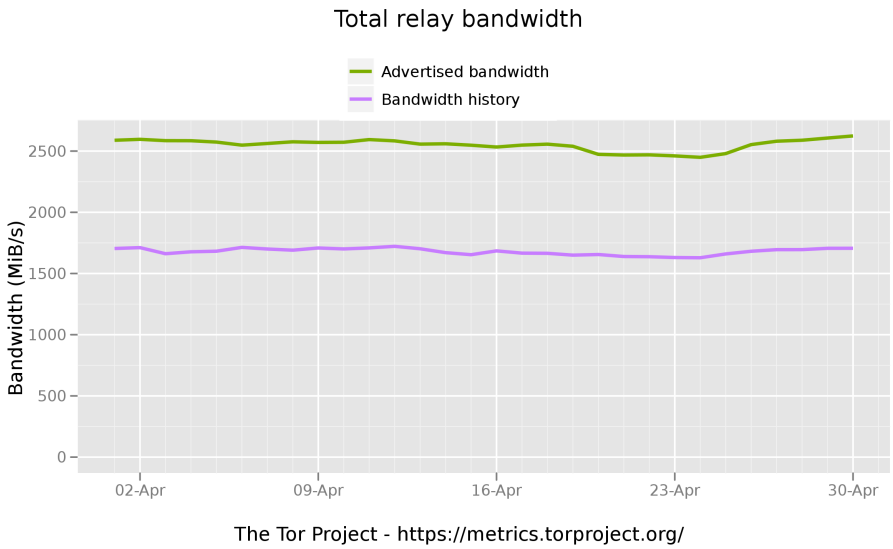

This graph shows the total available bandwidth available to clients and how much was actually used throughout the month.

1.3 Censorship & Circumvention

- The OONI project reports on Internet censorship in Bethlehem, Palestine at https://blog.torproject.org/blog/politically-motivated-censorship-bethlehem-palestine

- We prepared a server for migration of bridges.torproject.org from the current server, run by a volunteer, to one run by Tor. In testing, some of the patches for reCAPTCHA support and other enhancements made bridgedb consume excessive amounts of ram. Reverted the patches in preparation for a migration in May.

- Karsten is now more convinced that the way how bridges report their statistics to the bridge authority is not our problem. It's the way how we derive user numbers from unique IP addresses which is totally broken. We should try an approach that's similar to how we count directory requests on directory mirrors.

  The tech report is now available.

- We will research Internet drafts for one or two TLS features to improve its traffic-analysis resistance. In particular, TLS needs better support for link padding, and for a mode where it does not send its record headers as plaintext.

  Here's what Nick thinks the stages are:

  1. Sketch out some initial designs. 2. Show these to the people I know who are involved with the TLS working group, and see if they think we're fundamentally insane. Iterate until the basic designs seem right. 3. Write up the designs in internet-draft form. Get help from others as we do so. 4. Submit them. 5. receive comments, revise as needed, iterate.

We fully suspect that one or both of these ideas might get shot down at some point in the process, especially at step 4. But that's why we committed to internet drafts, not to RFCs.

Now, as for the timeline on these steps, obviously we should get to stage 5 as soon as we can for maximal utility, but I believe that only stage 3 or 4 is what we promised for November.

Marsh Ray appeared in #tor-dev to tell us that the IETF TLS working group is worrying that Google's TLS NextProto proposal will increase the blocking surface of TLS, since it advertises the transport layer protocol in plaintext (in the TLS extensions in the Hello records).

He told us that it's a good chance to influence the future of TLS to be more traffic-resistant. Here is the relevant mailing list thread:https://www.ietf.org/mail-archive/web/tls/current/msg08685.html

For our records, Marsh is lobbying for replacing the plaintext NPN that Google currently uses with a DH handshake as part of ClientHello and ServerHello. The server would use the DH key to encrypt the cert chain for the client in the ServerHello reply itself. IUC, the client would then put the NPN bits as part of its Finished message, also encrypted with the key.

So they seem to be totally open to redoing the TLS handshake to provide less data on the wire for blocking. Mike thinks the server cert chain will be a major issue for us, unless we want to do gymnastics like providing fake unused certs, so the plan seems like a step in the right direction.

- The current BridgeDB code contains a function to accept a list of blocked bridges by country to give out non-blocked bridges to users (1837). This code needs testing. Also, there is a bug with how BridgeDB learns which country a user is interested in getting un-blocked bridges for, which currently conflicts with the language selection.

  There are at least three approaches for giving out non-blocked bridges to users: a) exclude blocked bridges from results, b) replace blocked bridges with non-blocked once, or c) include blocked bridges in results and write "maybe blocked" next to them. We're currently doing a), but we should do c) to improve usability. There are variants of b), e.g., b1) Christian suggests to give out exactly one non-blocked bridge if otherwise we'd give out zero bridges and b2) Roger suggests to give out the first three non-blocked bridges in a fixed set of five bridges. One part of this deliverable should be to discuss these alternatives and agree on one of them that shall be implemented.

  Roger is worried that people have wildly different ideas of what algorithms BridgeDB should use. We agreed that doing a simple version to start is fine. If we have time left we should rather spend it on working on reachability testing.

  In the context of algorithms that BridgeDB uses, Roger asked about the status of the BridgeDB spec (1606). It was merged into bridgedb.git in April 2011 and is waiting for a review by Roger, Aaron, and Christian. It also wasn't updated since then, which may be due to lack of significant code changes. It should probably be updated once the changes for this deliverable are implemented.

- Nick is going to write a proposal for using a front-end proxy like Apache for bridge scanning resistance. Roger is most interested in the question "what changes on the Tor side will make it easier for people who want to stick a $foo in front of Tor to make Tor look like $foo unless

you authenticate right;" where Apache is the first $foo. George added that the other question is "what changes on the $foo side will be needed."

George asked how the Apache to Tor transition on correct password is going to happen. Roger mentioned a master's student at UT who wrote an Apache module to proxypass Tor traffic if you GET the right URL first.

2 Relay

2.1 Tor Network

- Roger says we have UPnP and NAT-PMP, but it would be great to let relays and bridges function behind NATs without doing port forwarding, e.g., by UDP tricks or TCP tricks or third-party tricks or something. There are NAT-piercing libraries out there, but most of them are poorly written. There are also a bunch of NAT-piercing techniques, some of them are probably even good ideas.

  Jake and Karsten agreed that a good first step would be to write a tech report that compares the different NAT-piercing options we have. It could be titled "Overview of NAT-piercing approaches for Tor relays and bridges" and include UPnP, NAT-PMP, and whatever we come up with. Jake says he's going to write such a tech report.

  Jake has sent a paper draft to Steven who's going to review it soon. Jake plans to send the revised and then final paper to tor-dev mailing list to kick off a discussion. Jake expects this to happen in April; it should happen before June to have some time for the discussion. People who should be made aware of the tor-dev mailing list discussion to allocate time for it are: Roger, Nick, and package-building people.

- Roger said "simulating tor networks by sampling random relays" has turned more into Rob's CSET paper. The bigger question is: "how do you model a big Tor network when all you can run is a little Tor network?" It has to do with selecting relays and relay capacities and rate limiting (a.k.a. this deliverable), but it also has to do with selecting client load. 5398 is very related.

- Runa created 14 new Tor Cloud images, updated the website, emailed tor-talk, and wrote a blog post about it (5568).

2.2 Console Client

- On April 29th, Arm 1.4.5 was released.

  A new release of arm [1] is now available which includes numerous fixes of mounting importance [2]. In particular this corrects several issues around arm's connection panel, terminal glitches due to disruption of the curses module by readline [3], and incompatibility with tor's new development releases [4][5].

  [1] http://www.atagar.com/arm/ [2] http://www.atagar.com/arm/releaseNotes.php [3] https://gitweb.torpro
  [4] https://lists.torproject.org/pipermail/tor-talk/2012-April/023961.html

2.3 Tor router

No progress this month.

3 Client

Andrew and Dr. Angela Sasse of University College of London started the qualitative phase of their privacy and circumvention usability study. Eight of Eleven individuals agreed to be interviewed.

3.1 Tor Browser

1. On April 28th we released a new version of Tor Browser. The Tor Browser Bundles have all been updated to the latest Firefox 12.0 as well as a number of other software updates, bugfixes, and new features. We've re-branded Firefox so it should now be more easy to distinguish between it and your normal Firefox. We've also added Korean and Vietnamese to the available languages.

   Tor Browser Bundle (2.2.35-9)

   ```
   Update Firefox to 12.0
   Update OpenSSL to 1.0.1b
   Update Libevent to 2.0.18-stable
   Update Qt to 4.8.1
   Update Libpng to 1.5.10
   Update HTTPS Everywhere to 2.0.2
   Update NoScript to 2.3.9
   Rebrand Firefox to TorBrowser (closes: #2176)
   New Firefox patches
       Make Download Manager memory-only (closes: #4017)
       Add DuckDuckGo and Startpage to Omnibox (closes: #4902)
       Add Steven Michaud's OS X crash fix patch. It doesn't fix #5021 but will hopefully
       https://bugzilla.mozilla.org/show_bug.cgi?id=715885#c35
   Make the 32-bit Tor Browser Bundle compatible with OS X 10.5
   ```

3.2 Obfsproxy

- Sebastian made new obfsproxy packages for users (with hard-coded obfsproxy bridge addresses) and uploaded them on April 12.

- George outlined a few possible next steps:

  George knows a couple of people who are coding Python pluggable transports; one of them is wiretapped with banana phone, another one is the Portland university guys. He wants to prepare a pluggable transport library for them. Roger adds that flash proxy also has a Python part. In general, Roger thinks that looking at other Python pluggable transports and trying to help them all use our framework should be part of this deliverable.

  George thinks that most of the future transports we are thinking of can be handled (performance-wise) by a Python program. He wants to prepare an obfs2-like thing and benchmark it by

feeding it with thousands of connections. George wants to see how many of them it will handle (he expects many).

George also wants to prepare a very lite managed proxy library, which will read environment variables etc. He also wants to write a very simple pluggable transport in Python which will use that managed proxy library.

George wants to evaluate py2exe (and the other similar things).

Nick suggests that we could do something minimalistic with twisted. George hopes to be able to use all the other web protocols baked in twisted for transports.

George will try to learn the truth about bytearrays and cryptographic material overwriting. The idea is that by using bytearrays in Python one can actually overwrite cryptographic material "securely".

3.3 Anonymous Computing

On April 25, Tails released a new version of their anonymous operating system, https://tails.boum.org/news/version_0.11/.

3.4 Mobile

We work closely with The Guardian Project to keep tor working on Android and related devices. Learn more about Guardian at https://guardianproject.info.

4 Community

4.1 Support

```
A total of 461 tickets were created in the Tor help desk system this
month: 387 tickets in the queue called help, 25 tickets in the queue
called help-fa, and 49 tickets in spam.

The help queue

  461 tickets were created in the help queue this month. 376 of them
  have been marked as resolved, 11 are currently open and waiting on a
  reply from a support assistant.

The help-fa queue

  25 tickets were created in the help-fa queue this month. All 25 have
  been marked as resolved.

The support assistants

  The previous status reports have the wrong numbers for tickets
```

```
  resolved by each support assistant. The numbers in the previous
  reports are tickets created *and* resolved within the same month, not
  tickets marked as resolved that month (regardless of when the tickets
  were created). Here are the correct numbers for March 2012:

  Ardeshir: 58 tickets resolved, 0 new, 0 open
  Nasim: 7 tickets resolved, 0 new, 0 open
  Runa: 530 tickets resolved, 0 new, 12 open

  Total number of tickets resolved by each support assistant since we
  started help@rt.tpo six months ago:

  Andrew: 4 tickets
  Ardeshir: 129 tickets
  Kaveh: 27 tickets
  Kim: 89 tickets
  Nasim: 102 tickets
  Nathan: 2 tickets
  Runa: 2247 tickets

The support requests and other things

  Users with OS X 10.5 continue to email saying they can't get the
  latest Tor Browser Bundle to work.

  A ton of users asked a question that has been answered in the short
  user manual. What is the status on including this in future TBBs?

  A lot of users in China are asking for TBB. I send them the Obfsproxy
  TBB, and they the majority seem happy with that. A few reply asking
  for a few bridges, but no one has emailed saying it doesn't work at
  all. I think that only some of the bridges we hardcoded are blocked.

  A few users needed help with extracting TBB on the Desktop to allow
  Firefox to launch properly. I should add a sentence about this to the
  short user manual.
```

4.2 Education & Training

- Andrew trained a few journalists and human rights workers on Tor, Internet safety and security, and leaving data trails online while in Stockholm.
- Andrew wrote up his experiences with first-time Tor/Tails users at https://lists.torproject.org/pipermail/tor-dev/2012-April/003472.html

4.3 Outreach

1. Jacob attended http://www.soros.org/initiatives/fellowship/events/privacy-overrated-20120416
2. Andrew attended the Stockholm Internet Forum 2012. His trip report is available.
3. Jacob coached a team at UW on cybersecurity stuff - we took nationals (!) this month for the second year in a row: http://seattletimes.nwsource.com/html/localnews/2018035088_cyber21m.html
4. We released a report on Ultrasurf, https://blog.torproject.org/blog/ultrasurf-definitive-review.
5. Jacob participated in the WSJ Data Transparency event and won an award with Nadim for Cryptocat: http://datatransparency.wsj.com/
6. Jacob and others produced the first spec for what an mpOTR implementation might look like if we expand on Ian's paper (review requested): https://github.com/kaepora/cryptocat/blob/master/spec/mpOTR.txt
7. Andrew was interviewed by The Guardian about the UK proposal to record all Internet activities, http://www.guardian.co.uk/media/2012/apr/02/internet-companies-warn-government-email-s
8. Swedish researchers helped elucidate the new blocking techniques seen in China, http://www.v3.co.uk/v3-uk/news/2165733/swedish-researchers-uncover-key-chinas-tor-blocking
9. Andrew was interviewed about online anonymity by The Guardian, http://www.guardian.co.uk/technology/2012/apr/19/online-identity-authenticity-anonymity.
10. Forbes talked about our latest project, OONI, http://www.forbes.com/sites/andygreenberg/2012/04/30/the-tor-projects-new-tool-aims-to-map-out-internet-censorship/.