From: Kelly DeYoe
To: Roger Dingledine
Subject: Re: Contract status update
Date: Thursday, March 23, 2006 10:56:30 AM

Once the contract is awarded, it will be for 8 months starting from the award date, if it overlaps into next fiscal year, it doesn't matter, the money has already been obligated for 8 months and carries over until the end of the contract.

In the end, it is probably better to have it overlap a bit after the next fiscal year begins anyway, since just because the government fiscal year starts on October 1, doesn't necessarily mean Congress has actually approved our agency's budget before then (it took them until December this year I think?), and our spending is limited during these periods when operating under a continuing resolution rather than a budget passed by law.

-k

Roger Dingledine wrote:
> On Wed, Mar 22, 2006 at 06:18:14PM -0500, Kelly DeYoe wrote:
>
>>Ok, so the latest update on the contract is that it is just awaiting the
>>CFO's approval before it can go to the Contracts office for award...
>>which sounds not much further along than it was 3 weeks ago I'm afraid.
>> Unfortunately, the administrative person who is in charge of this from
>>our office is out this week, so I was receiving the info secondhand from
>>another admin. officer, so I will check with the primary contact again
>>on Monday.
>
>
> Exciting. Thanks for sticking with this.
>
> Just so we're on the same page, the money for the current contract is
> intended to be spent in FY06 on work in FY06? So that means that once
> the contract finally does come through, we're still planning for me to
> bill all the hours in the current round by the end of FY06?
>
> Thanks,
> --Roger
>

From: Roger Dingledine
To: Chris Walker; Kelly DeYoe; Ken Berman; (b) (6)
Subject: Re: December GeoIP stats
Date: Sunday, April 13, 2008 3:26:45 AM

Four months later produces some very interesting stats. The percentage of German Tor users has gone up relative to the other users. But the total number of users in each country has gone up. Compared to the 97125 Tor clients that we glimpsed in mid December, we're now glimpsing a bit over 150000.

Does that mean the number of running Tor clients has gone up by 50% in the past four months? Remember that my measurements are still pretty rough, but it seems that there's a serious increase across the board, and a huge increase in Germany.

Note that we changed our geoip code since December, so this new count is leaving out about 1200 IP addresses that our geoip db doesn't recognize. Manually checking a few of them makes me think these are still mostly African countries. (These were lumped into country-code "UN" in previous batches.)

I'm going to try to think up some smarter tests, to see if we can figure out if this is a fluke or if we really did grow that much. :)

--Roger

```
Apr 13 02:00:00.643 [notice] Clients seen:
de=42875(28.030%)
us=24099(15.755%)
cn=18617(12.171%)
it=6022(3.937%)
fr=4904(3.206%)
gb=4778(3.124%)
ca=3449(2.255%)
pl=3378(2.208%)
jp=3033(1.983%)
ru=2709(1.771%)
es=2594(1.696%)
at=2286(1.494%)
tr=2270(1.484%)
br=1918(1.254%)
se=1748(1.143%)
tw=1726(1.128%)
au=1720(1.124%)
nl=1503(0.983%)
ch=1460(0.954%)
mx=1078(0.705%)
ir=885(0.579%)
in=874(0.571%)
ro=859(0.562%)
ar=834(0.545%)
cz=811(0.530%)
fi=789(0.516%)
be=782(0.511%)
no=763(0.499%)
ua=742(0.485%)
hk=699(0.457%)
```

```
vn=618(0.404%)
gr=587(0.384%)
dk=580(0.379%)
cl=576(0.377%)
sg=573(0.375%)
hu=572(0.374%)
il=570(0.373%)
pt=562(0.367%)
my=467(0.305%)
th=457(0.299%)
id=390(0.255%)
sk=360(0.235%)
ie=333(0.218%)
co=324(0.212%)
lt=320(0.209%)
ph=320(0.209%)
bg=284(0.186%)
nz=282(0.184%)
ve=265(0.173%)
sa=260(0.170%)
hr=231(0.151%)
qa=222(0.145%)
si=207(0.135%)
pe=180(0.118%)
cs=177(0.116%)
kr=177(0.116%)
lv=166(0.109%)
kw=163(0.107%)
ee=136(0.089%)
pk=132(0.086%)
ge=118(0.077%)
by=112(0.073%)
md=93(0.061%)
pr=88(0.058%)
lu=87(0.057%)
gt=85(0.056%)
kz=85(0.056%)
sv=79(0.052%)
cr=73(0.048%)
do=72(0.047%)
rs=71(0.046%)
jo=70(0.046%)
mk=67(0.044%)
ae=63(0.041%)
uy=58(0.038%)
sy=57(0.037%)
om=56(0.037%)
pa=49(0.032%)
ps=49(0.032%)
ec=48(0.031%)
cy=44(0.029%)
ba=43(0.028%)
jm=41(0.027%)
uz=40(0.026%)
bo=36(0.024%)
is=36(0.024%)
mt=32(0.021%)
az=31(0.020%)
bh=30(0.020%)
lb=30(0.020%)
```

```
py=28(0.018%)
bd=25(0.016%)
hn=21(0.014%)
lk=20(0.013%)
ni=19(0.012%)
mo=18(0.012%)
mv=15(0.010%)
al=14(0.009%)
me=14(0.009%)
tt=14(0.009%)
am=13(0.008%)
iq=11(0.007%)
bs=10(0.007%)
mn=10(0.007%)
ye=10(0.007%)
cu=9(0.006%)
li=9(0.006%)
bb=8(0.005%)
bm=7(0.005%)
ky=7(0.005%)
aw=6(0.004%)
bn=6(0.004%)
fj=6(0.004%)
gu=6(0.004%)
af=5(0.003%)
an=5(0.003%)
bz=5(0.003%)
gl=5(0.003%)
pf=5(0.003%)
ag=4(0.003%)
ax=4(0.003%)
dz=4(0.003%)
fo=4(0.003%)
kh=4(0.003%)
mc=4(0.003%)
mq=4(0.003%)
cd=3(0.002%)
dm=3(0.002%)
ng=3(0.002%)
np=3(0.002%)
re=3(0.002%)
sm=3(0.002%)
tj=3(0.002%)
ad=2(0.001%)
eg=2(0.001%)
ht=2(0.001%)
la=2(0.001%)
ne=2(0.001%)
sr=2(0.001%)
tm=2(0.001%)
zm=2(0.001%)
zw=2(0.001%)
bj=1(0.001%)
cg=1(0.001%)
ck=1(0.001%)
cv=1(0.001%)
fk=1(0.001%)
gh=1(0.001%)
gi=1(0.001%)
gq=1(0.001%)
```

```
gy=1(0.001%)
ke=1(0.001%)
kg=1(0.001%)
lc=1(0.001%)
mh=1(0.001%)
mm=1(0.001%)
mp=1(0.001%)
mz=1(0.001%)
nc=1(0.001%)
sb=1(0.001%)
tc=1(0.001%)
tv=1(0.001%)
tz=1(0.001%)
vc=1(0.001%)
vg=1(0.001%)
vi=1(0.001%)
za=1(0.001%)
```

On Wed, Dec 12, 2007 at 08:41:40PM -0500, Roger Dingledine wrote:

```
> Dec 12 20:36:06.977 [notice] 97125 Total
> Dec 12 20:36:06.977 [notice] 21672 DE (22.314%)
> Dec 12 20:36:06.977 [notice] 17028 US (17.532%)
> Dec 12 20:36:06.977 [notice] 16679 CN (17.173%)
> Dec 12 20:36:06.977 [notice] 3666 FR (3.775%)
> Dec 12 20:36:06.977 [notice] 3528 IT (3.632%)
> Dec 12 20:36:06.978 [notice] 2746 GB (2.827%)
> Dec 12 20:36:06.978 [notice] 2561 CA (2.637%)
> Dec 12 20:36:06.978 [notice] 2139 JP (2.202%)
> Dec 12 20:36:06.978 [notice] 2082 PL (2.144%)
> Dec 12 20:36:06.978 [notice] 1911 TW (1.968%)
> Dec 12 20:36:06.978 [notice] 1482 ES (1.526%)
> Dec 12 20:36:06.978 [notice] 1328 BR (1.367%)
> Dec 12 20:36:06.978 [notice] 1303 RU (1.342%)
> Dec 12 20:36:06.978 [notice] 1120 AU (1.153%)
> Dec 12 20:36:06.978 [notice] 1090 AT (1.122%)
> Dec 12 20:36:06.978 [notice] 959 UN (0.987%)
> Dec 12 20:36:06.978 [notice] 952 SE (0.980%)
> Dec 12 20:36:06.978 [notice] 899 NL (0.926%)
> Dec 12 20:36:06.978 [notice] 778 CH (0.801%)
> Dec 12 20:36:06.978 [notice] 685 TR (0.705%)
> Dec 12 20:36:06.978 [notice] 598 CZ (0.616%)
> Dec 12 20:36:06.978 [notice] 543 MX (0.559%)
> Dec 12 20:36:06.978 [notice] 528 NO (0.544%)
> Dec 12 20:36:06.978 [notice] 524 IN (0.540%)
> Dec 12 20:36:06.978 [notice] 522 IR (0.537%)
> Dec 12 20:36:06.979 [notice] 509 BE (0.524%)
> Dec 12 20:36:06.979 [notice] 463 RO (0.477%)
> Dec 12 20:36:06.979 [notice] 451 AR (0.464%)
> Dec 12 20:36:06.979 [notice] 440 FI (0.453%)
> Dec 12 20:36:06.979 [notice] 409 DK (0.421%)
> Dec 12 20:36:06.979 [notice] 404 TH (0.416%)
> Dec 12 20:36:06.979 [notice] 383 PT (0.394%)
> Dec 12 20:36:06.979 [notice] 358 IL (0.369%)
> Dec 12 20:36:06.979 [notice] 351 SG (0.361%)
> Dec 12 20:36:06.979 [notice] 336 GR (0.346%)
> Dec 12 20:36:06.980 [notice] 326 UA (0.336%)
> Dec 12 20:36:06.980 [notice] 316 MY (0.325%)
> Dec 12 20:36:06.980 [notice] 315 HK (0.324%)
> Dec 12 20:36:06.980 [notice] 311 VN (0.320%)
> Dec 12 20:36:06.980 [notice] 282 LT (0.290%)
> Dec 12 20:36:06.980 [notice] 280 SK (0.288%)
> Dec 12 20:36:06.980 [notice] 276 HU (0.284%)
> Dec 12 20:36:06.980 [notice] 225 CL (0.232%)
> Dec 12 20:36:06.980 [notice] 201 PH (0.207%)
> Dec 12 20:36:06.980 [notice] 199 NZ (0.205%)
> Dec 12 20:36:06.980 [notice] 191 IE (0.197%)
> Dec 12 20:36:06.980 [notice] 165 CO (0.170%)
> Dec 12 20:36:06.980 [notice] 163 ID (0.168%)
> Dec 12 20:36:06.980 [notice] 143 BG (0.147%)
> Dec 12 20:36:06.980 [notice] 142 KR (0.146%)
> Dec 12 20:36:06.980 [notice] 138 SA (0.142%)
> Dec 12 20:36:06.980 [notice] 129 HR (0.133%)
> Dec 12 20:36:06.980 [notice] 129 VE (0.133%)
> Dec 12 20:36:06.980 [notice] 104 PE (0.107%)
> Dec 12 20:36:06.980 [notice] 99 SI (0.102%)
> Dec 12 20:36:06.980 [notice] 92 CS (0.095%)
> Dec 12 20:36:06.980 [notice] 83 QA (0.085%)
> Dec 12 20:36:06.980 [notice] 80 LV (0.082%)
> Dec 12 20:36:06.981 [notice] 72 BY (0.074%)
> Dec 12 20:36:06.981 [notice] 72 KW (0.074%)
> Dec 12 20:36:06.981 [notice] 71 EE (0.073%)
> Dec 12 20:36:06.981 [notice] 62 PK (0.064%)
> Dec 12 20:36:06.981 [notice] 50 SY (0.051%)
> Dec 12 20:36:06.981 [notice] 48 PR (0.049%)
> Dec 12 20:36:06.981 [notice] 46 LU (0.047%)
> Dec 12 20:36:06.981 [notice] 40 JO (0.041%)
> Dec 12 20:36:06.981 [notice] 40 GT (0.041%)
> Dec 12 20:36:06.981 [notice] 39 OM (0.040%)
> Dec 12 20:36:06.981 [notice] 39 DO (0.040%)
> Dec 12 20:36:06.981 [notice] 37 CR (0.038%)
> Dec 12 20:36:06.981 [notice] 35 MK (0.036%)
> Dec 12 20:36:06.981 [notice] 31 AE (0.032%)
> Dec 12 20:36:06.981 [notice] 29 KZ (0.030%)
> Dec 12 20:36:06.981 [notice] 29 SV (0.030%)
> Dec 12 20:36:06.981 [notice] 29 UY (0.030%)
> Dec 12 20:36:06.981 [notice] 28 EC (0.029%)
> Dec 12 20:36:06.981 [notice] 28 RS (0.029%)
> Dec 12 20:36:06.981 [notice] 28 BA (0.029%)
> Dec 12 20:36:06.981 [notice] 28 PA (0.029%)
> Dec 12 20:36:06.981 [notice] 24 BO (0.025%)
> Dec 12 20:36:06.982 [notice] 21 BH (0.022%)
> Dec 12 20:36:06.982 [notice] 20 CY (0.021%)
> Dec 12 20:36:06.982 [notice] 19 LK (0.020%)
> Dec 12 20:36:06.982 [notice] 17 PY (0.018%)
> Dec 12 20:36:06.982 [notice] 17 MD (0.018%)
> Dec 12 20:36:06.982 [notice] 16 UZ (0.016%)
> Dec 12 20:36:06.982 [notice] 16 MT (0.016%)
> Dec 12 20:36:06.982 [notice] 15 PS (0.015%)
> Dec 12 20:36:06.982 [notice] 14 NI (0.014%)
> Dec 12 20:36:06.982 [notice] 13 LB (0.013%)
> Dec 12 20:36:06.982 [notice] 13 YE (0.013%)
> Dec 12 20:36:06.982 [notice] 12 HN (0.012%)
> Dec 12 20:36:06.982 [notice] 12 MO (0.012%)
> Dec 12 20:36:06.982 [notice] 12 BD (0.012%)
> Dec 12 20:36:06.982 [notice] 11 IQ (0.011%)
> Dec 12 20:36:06.982 [notice] 9 BB (0.009%)
> Dec 12 20:36:06.982 [notice] 9 BN (0.009%)
> Dec 12 20:36:06.982 [notice] 9 GE (0.009%)
> Dec 12 20:36:06.982 [notice] 9 AL (0.009%)
> Dec 12 20:36:06.982 [notice] 8 IS (0.008%)
> Dec 12 20:36:06.983 [notice] 6 TT (0.006%)
> Dec 12 20:36:06.983 [notice] 6 CU (0.006%)
> Dec 12 20:36:06.983 [notice] 6 PF (0.006%)
> Dec 12 20:36:06.983 [notice] 6 MV (0.006%)
> Dec 12 20:36:06.983 [notice] 5 BM (0.005%)
> Dec 12 20:36:06.983 [notice] 5 KH (0.005%)
> Dec 12 20:36:06.983 [notice] 5 BS (0.005%)
> Dec 12 20:36:06.983 [notice] 5 GU (0.005%)
> Dec 12 20:36:06.983 [notice] 5 AN (0.005%)
> Dec 12 20:36:06.983 [notice] 5 JM (0.005%)
> Dec 12 20:36:06.983 [notice] 5 AZ (0.005%)
> Dec 12 20:36:06.983 [notice] 4 LI (0.004%)
> Dec 12 20:36:06.983 [notice] 4 FJ (0.004%)
> Dec 12 20:36:06.983 [notice] 4 AX (0.004%)
> Dec 12 20:36:06.983 [notice] 3 NG (0.003%)
> Dec 12 20:36:06.983 [notice] 3 AF (0.003%)
> Dec 12 20:36:06.983 [notice] 3 KY (0.003%)
> Dec 12 20:36:06.983 [notice] 3 MC (0.003%)
> Dec 12 20:36:06.983 [notice] 3 MN (0.003%)
> Dec 12 20:36:06.983 [notice] 3 SM (0.003%)
> Dec 12 20:36:06.983 [notice] 3 NC (0.003%)
> Dec 12 20:36:06.984 [notice] 2 KG (0.002%)
> Dec 12 20:36:06.984 [notice] 2 ME (0.002%)
> Dec 12 20:36:06.984 [notice] 2 NP (0.002%)
> Dec 12 20:36:06.984 [notice] 2 AW (0.002%)
> Dec 12 20:36:06.984 [notice] 2 VC (0.002%)
> Dec 12 20:36:06.984 [notice] 2 GY (0.002%)
> Dec 12 20:36:06.984 [notice] 2 AD (0.002%)
> Dec 12 20:36:06.984 [notice] 2 AM (0.002%)
> Dec 12 20:36:06.984 [notice] 2 AG (0.002%)
> Dec 12 20:36:06.984 [notice] 2 EG (0.002%)
> Dec 12 20:36:06.984 [notice] 2 VI (0.002%)
> Dec 12 20:36:06.984 [notice] 1 LC (0.001%)
> Dec 12 20:36:06.984 [notice] 1 GL (0.001%)
> Dec 12 20:36:06.984 [notice] 1 CK (0.001%)
> Dec 12 20:36:06.984 [notice] 1 DM (0.001%)
> Dec 12 20:36:06.984 [notice] 1 MZ (0.001%)
> Dec 12 20:36:06.984 [notice] 1 LY (0.001%)
> Dec 12 20:36:06.984 [notice] 1 PG (0.001%)
> Dec 12 20:36:06.984 [notice] 1 GI (0.001%)
> Dec 12 20:36:06.984 [notice] 1 ST (0.001%)
> Dec 12 20:36:06.984 [notice] 1 DZ (0.001%)
> Dec 12 20:36:06.984 [notice] 1 KE (0.001%)
> Dec 12 20:36:06.985 [notice] 1 GN (0.001%)
> Dec 12 20:36:06.985 [notice] 1 BZ (0.001%)
> Dec 12 20:36:06.985 [notice] 1 ZM (0.001%)
> Dec 12 20:36:06.985 [notice] 1 FO (0.001%)
> Dec 12 20:36:06.985 [notice] 1 MA (0.001%)
> Dec 12 20:36:06.985 [notice] 1 FK (0.001%)
> Dec 12 20:36:06.985 [notice] 1 TL (0.001%)
```

From: Andrew Lewman
To: Sho Ho
Cc: Roger Dingledine; Ken Berman; Kelly DeYoe
Subject: Re: Defcon17, FOE, and Sho
Date: Thursday, June 25, 2009 10:07:26 AM

On 06/18/2009 09:54 AM, Sho Ho wrote:
> FOE stands for "Feed Over Email". In very simple terms, FOE uses email
> to transport data to end-users. The data that FOE transport can be
> anything from RSS feeds to normal files to proxy addresses. FOE
> messages are compressed and encoded so normal keyword-filtering
> technologies won't be able to censor FOE messages (data can also be
> encrypted if necessary.)

This sounds very much like an email2web gateway I wrote back in the 1990s to get around my employer's stupid policy of "no ncsa mosiac or gopher".

> The main difference between regular email and FOE is that, instead of
> the user reading the email directly, the FOE client program will
> decompress and decode FOE messages and present the data in meaningful
> ways (e.g. displaying RSS feed, downloading files/applications,
> providing latest proxy server addresses, etc.)

Sounds good.

> Since FOE is based on email (SMTP and POP3), it is relatively
> straightforward to create a version for mobile devices or other
> computing platforms.

Does it support/require ssl-enabled smtp/pop3? What about imap?

> I will release the full documentation and source code of FOE at around
> DEFCON.

Great. I look forward to seeing all of this released. We'll review and provide feedback.

Thanks for the quick overview!

--
Andrew Lewman
The Tor Project
pgp 0x31B0974B
(b)(6)

Website: https://torproject.org/
Blog: https://blog.torproject.org/
Identica/Twitter: torproject

From: Sho Ho
To: [redacted]: Roger Dingledine
Cc: [redacted]ly DeYoe
Subject: Re: Defcon17, FOE, and Sho
Date: Thursday, June 18, 2009 10:54:33 AM

Hi Andrew, Roger,

Nice to meet you. It is my pleasure to give a brief introduction on FOE.

FOE stands for "Feed Over Email". In very simple terms, FOE uses email to transport data to end-users. The data that FOE transport can be anything from RSS feeds to normal files to proxy addresses. FOE messages are compressed and encoded so normal keyword-filtering technologies won't be able to censor FOE messages (data can also be encrypted if necessary.)

The main difference between regular email and FOE is that, instead of the user reading the email directly, the FOE client program will decompress and decode FOE messages and present the data in meaningful ways (e.g. displaying RSS feed, downloading files/applications, providing latest proxy server addresses, etc.)

Each FOE user will need an email account outside his country in order for FOE to achieve the maximum success rate. Users can use any email providers that support SMTP and POP3 such as Gmail.

The main goals of FOE are (1) allow users to receive the latest news and (2) complement existing anti-censorship tools. For example, FOE can push new proxy addresses to users so they can browse the uncensored Internet via one of our web-based proxy servers. Users can also use FOE to download software applications such as Tor, Freegate, or Ultrasurf. Most importantly, users will be able to receive the latest news using FOE.

Since FOE is based on email (SMTP and POP3), it is relatively straightforward to create a version for mobile devices or other computing platforms.

I will release the full documentation and source code of FOE at around DEFCON. However, FOE is still in its infancy so there are plenty of things that need to be worked on, and there are many ways you can help:

- Send us your honest opinions on what you think will or will not work.
- Help us to test and identify potential vulnerabilities.
- Help us to improve the codes, algorithms, architecture, etc.
- Create improved versions of FOE for different platforms.
- Provide suggestions on just about anything that you can think of.

I hope this gives you a better idea of FOE. Please feel free to contact me if you have any questions.

Thanks,
Sho

Ken Berman wrote:
> Why not talk to the author herself?!

> Sho - meet the Tor folks; you've been reading their progress reports for several months now....
>
> Ken
>
> -----Original Message-----
> From: Andrew Lewman [mailto:[redacted] (b) (6)
> Sent: Wednesday, June 17, 2009 4:39 PM
> To: Kelly DeYoe; Ken Berman
> Cc: Roger Dingledine
> Subject: Defcon17, FOE, and Sho
>
> Hello Ken and Kelly,
>
> I noticed there is a talk at Defcon17 by Sho about FOE, http://www.defcon.org/html/defcon-17/dc-17-speakers.html#Ho. Can you tell us anything about FOE? Are there ways we can help?
>
> Thanks!
>
> --
> Andrew Lewman
> The Tor Project
> pgp 0x31B0974B
> [redacted] (b) (6)
>
> Website: https://torproject.org/
> Blog: https://blog.torproject.org/
> Identica/Twitter: torproject
>
>

From: Sho Ho
To: Andrew Lewman
Cc: Roger Dingledine; Ken Berman; Kelly DeYoe
Subject: Re: Defcon17, FOE, and Sho
Date: Thursday, June 25, 2009 10:52:46 AM

Hi Andrew,

Yes, it supports SMTP and POP3 over SSL; and no, it doesn't support IMAP, yet, but is on the wish list. Thanks.

Regards,
Sho

Andrew Lewman wrote:
> On 06/18/2009 09:54 AM, Sho Ho wrote:
>
>> FOE stands for "Feed Over Email". In very simple terms, FOE uses email
>> to transport data to end-users. The data that FOE transport can be
>> anything from RSS feeds to normal files to proxy addresses. FOE
>> messages are compressed and encoded so normal keyword-filtering
>> technologies won't be able to censor FOE messages (data can also be
>> encrypted if necessary.)
>>
>
> This sounds very much like an email2web gateway I wrote back in the
> 1990s to get around my employer's stupid policy of "no ncsa mosiac or
> gopher".
>
>
>> The main difference between regular email and FOE is that, instead of
>> the user reading the email directly, the FOE client program will
>> decompress and decode FOE messages and present the data in meaningful
>> ways (e.g. displaying RSS feed, downloading files/applications,
>> providing latest proxy server addresses, etc.)
>>
>
> Sounds good.
>
> > Since FOE is based on email (SMTP and POP3), it is relatively
>
>> straightforward to create a version for mobile devices or other
>> computing platforms.
>>
>
> Does it support/require ssl-enabled smtp/pop3? What about imap?
>
>
>> I will release the full documentation and source code of FOE at around
>> DEFCON.
>>
>
> Great. I look forward to seeing all of this released. We'll review and
> provide feedback.
>
> Thanks for the quick overview!
>
>

From: Ken Berman
To: Andrew Lewman; Kelly DeYoe
Cc: Roger Dingledine; Sho Ho
Subject: RE: Defcon17, FOE, and Sho
Date: Wednesday, June 17, 2009 5:44:16 PM

Why not talk to the author herself?!
Sho - meet the Tor folks; you've been reading their progress reports for several months now....

Ken

-----Original Message-----
From: Andrew Lewman [mailto:(b) (6)
Sent: Wednesday, June 17, 2009 4:39 PM
To: Kelly DeYoe; Ken Berman
Cc: Roger Dingledine
Subject: Defcon17, FOE, and Sho

Hello Ken and Kelly,

I noticed there is a talk at Defcon17 by Sho about FOE, http://www.defcon.org/html/defcon-17/dc-17-speakers.html#Ho. Can you tell us anything about FOE? Are there ways we can help?

Thanks!

--
Andrew Lewman
The Tor Project
pgp 0x31B0974B
(b) (6)

Website: https://torproject.org/
Blog: https://blog.torproject.org/
Identica/Twitter: torproject

From: Ken Berman
To: Roger Dingledine
Cc: Kelly DeYoe; Shava Nerad; Hiu Ho; Betty Pruitt
Subject: Re: First draft of blocking-resistance design
Date: Tuesday, February 13, 2007 2:34:44 PM

Let's do a conf call later in the week or early next week, ladies and gents. Ken

Roger Dingledine wrote:

> On Tue, Jan 30, 2007 at 02:20:05PM -0500, Kelly DeYoe wrote:
>
> > Ken didn't realize it, but I'm actually going to be out
> > on Friday and
> > all week next week, so we'll probably need to postpone
> > this until the
> > week of 2/12.
>
> Hi folks,
>
> I am now back from the dead. Sorry for the scare. :(
>
> I'm out of town from 2/12 to 2/22, and also on 2/26 and 2/27.
> Should
> we schedule for 2/23 or 3/1 or 3/2?
>
> In the mean time here are the development release notes for the
> last few
> releases, to give you a sense of what we've been up to. We're
> trying
> to stabilize the Tor 0.1.2.x release, which has lots of new
> features
> (and still a few bugs), before we go on.
>
> Thanks,
> --Roger

```
Changes in version 0.1.2.7-alpha - 2007-02-06
  o Major bugfixes (rate limiting):
    - Servers decline directory requests much more aggressively
when
      they're low on bandwidth. Otherwise they end up queueing
more and
      more directory responses, which can't be good for latency.
    - But never refuse directory requests from local addresses.
    - Fix a memory leak when sending a 503 response for a
networkstatus
      request.
    - Be willing to read or write on local connections (e.g.
controller
      connections) even when the global rate limiting buckets are
empty.
    - If our system clock jumps back in time, don't publish a
negative
      uptime in the descriptor. Also, don't let the global rate
limiting
      buckets go absurdly negative.
    - Flush local controller connection buffers periodically as
we're
      writing to them, so we avoid queueing 4+ megabytes of data
before
      trying to flush.

  o Major bugfixes (NT services):
    - Install as NT_AUTHORITY\LocalService rather than as SYSTEM;
add a
      command-line flag so that admins can override the default by
saying
      "tor --service install --user "SomeUser"".  This will not
```

```
affect
      existing installed services.  Also, warn the user that the
service
      will look for its configuration file in the service user's
      %appdata% directory.  (We can't do the 'hardwire the user's
appdata
      directory' trick any more, since we may not have read access
to that
      directory.)

  o Major bugfixes (other):
    - Previously, we would cache up to 16 old networkstatus
documents
      indefinitely, if they came from nontrusted authorities. Now
we
      discard them if they are more than 10 days old.
    - Fix a crash bug in the presence of DNS hijacking (reported
by Andrew
      Del Vecchio).
    - Detect and reject malformed DNS responses containing
circular
      pointer loops.
    - If exits are rare enough that we're not marking exits as
guards,
      ignore exit bandwidth when we're deciding the required
bandwidth
      to become a guard.
    - When we're handling a directory connection tunneled over
Tor,
      don't fill up internal memory buffers with all the data we
want
      to tunnel; instead, only add it if the OR connection that
will
      eventually receive it has some room for it. (This can lead
to
      slowdowns in tunneled dir connections; a better solution
will have
      to wait for 0.2.0.)

  o Minor bugfixes (dns):
    - Add some defensive programming to eventdns.c in an attempt
to catch
      possible memory-stomping bugs.
    - Detect and reject DNS replies containing IPv4 or IPv6
records with
      an incorrect number of bytes. (Previously, we would ignore
the
      extra bytes.)
    - Fix as-yet-unused reverse IPv6 lookup code so it sends
nybbles
      in the correct order, and doesn't crash.
    - Free memory held in recently-completed DNS lookup attempts
on exit.
      This was not a memory leak, but may have been hiding memory
leaks.
    - Handle TTL values correctly on reverse DNS lookups.
    - Treat failure to parse resolv.conf as an error.

  o Minor bugfixes (other):
    - Fix crash with "tor --list-fingerprint" (reported by
seeess).
    - When computing clock skew from directory HTTP headers,
consider what
      time it was when we finished asking for the directory, not
what
      time it is now.
    - Expire socks connections if they spend too long waiting for
the
      handshake to finish. Previously we would let them sit around
for
      days, if the connecting application didn't close them
either.
    - And if the socks handshake hasn't started, don't send a
      "DNS resolve socks failed" handshake reply; just close it.
    - Stop using C functions that OpenBSD's linker doesn't like.
    - Don't launch requests for descriptors unless we have
```

```
networkstatuses
      from at least half of the authorities.  This delays the
first
      download slightly under pathological circumstances, but can
prevent
      us from downloading a bunch of descriptors we don't need.
    - Do not log IPs with TLS failures for incoming TLS
      connections. (Fixes bug 382.)
    - If the user asks to use invalid exit nodes, be willing to
use
      unstable ones.
    - Stop using the reserved ac_cv namespace in our configure
script.
    - Call stat() slightly less often; use fstat() when possible.
    - Refactor the way we handle pending circuits when an OR
connection
      completes or fails, in an attempt to fix a rare crash bug.
    - Only rewrite a conn's address based on X-Forwarded-For:
headers
      if it's a parseable public IP address; and stop adding extra
quotes
      to the resulting address.

  o Major features:
    - Weight directory requests by advertised bandwidth. Now we
can
      let servers enable write limiting but still allow most
clients to
      succeed at their directory requests. (We still ignore
weights when
      choosing a directory authority; I hope this is a feature.)

  o Minor features:
    - Create a new file ReleaseNotes which was the old ChangeLog.
The
      new ChangeLog file now includes the summaries for all
development
      versions too.
    - Check for addresses with invalid characters at the exit as
well
      as at the client, and warn less verbosely when they fail.
You can
      override this by setting ServerDNSAllowNonRFC953Addresses to
1.
    - Adapt a patch from goodell to let the contrib/exitlist
script
      take arguments rather than require direct editing.
    - Inform the server operator when we decide not to advertise a
      DirPort due to AccountingMax enabled or a low BandwidthRate.
It
      was confusing Zax, so now we're hopefully more helpful.
    - Bring us one step closer to being able to establish an
encrypted
      directory tunnel without knowing a descriptor first. Still
not
      ready yet. As part of the change, now assume we can use a
      create_fast cell if we don't know anything about a router.
    - Allow exit nodes to use nameservers running on ports other
than 53.
    - Servers now cache reverse DNS replies.
    - Add an --ignore-missing-torrc command-line option so that we
can
      get the "use sensible defaults if the configuration file
doesn't
      exist" behavior even when specifying a torrc location on the
command
      line.

  o Minor features (controller):
    - Track reasons for OR connection failure; make these reasons
      available via the controller interface. (Patch from Mike
Perry.)
    - Add a SOCKS_BAD_HOSTNAME client status event so controllers
      can learn when clients are sending malformed hostnames to
Tor.
    - Clean up documentation for controller status events.
    - Add a REMAP status to stream events to note that a stream's
      address has changed because of a cached address or a
MapAddress
      directive.


Changes in version 0.1.2.6-alpha - 2007-01-09
  o Major bugfixes:
    - Fix an assert error introduced in 0.1.2.5-alpha: if a single
TLS
      connection handles more than 4 gigs in either direction, we
crash.
    - Fix an assert error introduced in 0.1.2.5-alpha: if we're an
      advertised exit node, somebody might try to exit from us
when
      we're bootstrapping and before we've built our descriptor
yet.
      Refuse the connection rather than crashing.

  o Minor bugfixes:
    - Warn if we (as a server) find that we've resolved an address
that we
      weren't planning to resolve.
    - Warn that using select() on any libevent version before 1.1
will be
      unnecessarily slow (even for select()).
    - Flush ERR-level controller status events just like we
currently
      flush ERR-level log events, so that a Tor shutdown doesn't
prevent
      the controller from learning about current events.

  o Minor features (more controller status events):
    - Implement EXTERNAL_ADDRESS server status event so controllers
can
      learn when our address changes.
    - Implement BAD_SERVER_DESCRIPTOR server status event so
controllers
      can learn when directories reject our descriptor.
    - Implement SOCKS_UNKNOWN_PROTOCOL client status event so
controllers
      can learn when a client application is speaking a non-socks
protocol
      to our SocksPort.
    - Implement DANGEROUS_SOCKS client status event so controllers
      can learn when a client application is leaking DNS
addresses.
    - Implement BUG general status event so controllers can learn
when
      Tor is unhappy about its internal invariants.
    - Implement CLOCK_SKEW general status event so controllers can
learn
      when Tor thinks the system clock is set incorrectly.
    - Implement GOOD_SERVER_DESCRIPTOR and
ACCEPTED_SERVER_DESCRIPTOR
      server status events so controllers can learn when their
descriptors
      are accepted by a directory.
    - Implement CHECKING_REACHABILITY and
REACHABILITY_{SUCCEEDED|FAILED}
      server status events so controllers can learn about Tor's
progress in
      deciding whether it's reachable from the outside.
    - Implement BAD_LIBEVENT general status event so controllers
can learn
      when we have a version/method combination in libevent that
needs to
      be changed.
    - Implement NAMESERVER_STATUS, NAMESERVER_ALL_DOWN,
DNS_HIJACKED,
      and DNS_USELESS server status events so controllers can
learn
      about changes to DNS server status.

  o Minor features (directory):
    - Authorities no longer recommend exits as guards if this
```

```
would shift
      too much load to the exit nodes.


Changes in version 0.1.2.5-alpha - 2007-01-06
  o Major features:
    - Enable write limiting as well as read limiting. Now we
sacrifice
      capacity if we're pushing out lots of directory traffic,
rather
      than overrunning the user's intended bandwidth limits.
    - Include TLS overhead when counting bandwidth usage;
previously, we
      would count only the bytes sent over TLS, but not the bytes
used
      to send them.
    - Support running the Tor service with a torrc not in the same
      directory as tor.exe and default to using the torrc located
in
      the %appdata%\Tor\ of the user who installed the service.
Patch
      from Matt Edman.
    - Servers now check for the case when common DNS requests are
going to
      wildcarded addresses (i.e. all getting the same answer), and
change
      their exit policy to reject *:* if it's happening.
    - Implement BEGIN_DIR cells, so we can connect to the
directory
      server via TLS to do encrypted directory requests rather
than
      plaintext. Enable via the TunnelDirConns and
PreferTunneledDirConns
      config options if you like.

  o Minor features (config and docs):
    - Start using the state file to store bandwidth accounting
data:
      the bw_accounting file is now obsolete. We'll keep
generating it
      for a while for people who are still using 0.1.2.4-alpha.
    - Try to batch changes to the state file so that we do as few
      disk writes as possible while still storing important things
in
      a timely fashion.
    - The state file and the bw_accounting file get saved less
often when
      the AvoidDiskWrites config option is set.
    - Make PIDFile work on Windows (untested).
    - Add internal descriptions for a bunch of configuration
options:
      accessible via controller interface and in comments in saved
      options files.
    - Reject *:563 (NNTPS) in the default exit policy. We already
reject
      NNTP by default, so this seems like a sensible addition.
    - Clients now reject hostnames with invalid characters. This
should
      avoid some inadvertent info leaks. Add an option
      AllowNonRFC953Hostnames to disable this behavior, in case
somebody
      is running a private network with hosts called @, !, and #.
    - Add a maintainer script to tell us which options are missing
      documentation: "make check-docs".
    - Add a new address-spec.txt document to describe our special-
case
      addresses: .exit, .onion, and .noconnnect.

  o Minor features (DNS):
    - Ongoing work on eventdns infrastructure: now it has dns
server
      and ipv6 support. One day Tor will make use of it.
    - Add client-side caching for reverse DNS lookups.
    - Add support to tor-resolve tool for reverse lookups and
SOCKS5.
    - When we change nameservers or IP addresses, reset and re-
```

```
launch
      our tests for DNS hijacking.

  o Minor features (directory):
    - Authorities now specify server versions in networkstatus.
This adds
      about 2% to the side of compressed networkstatus docs, and
allows
      clients to tell which servers support BEGIN_DIR and which
don't.
      The implementation is forward-compatible with a proposed
future
      protocol version scheme not tied to Tor versions.
    - DirServer configuration lines now have an orport= option so
      clients can open encrypted tunnels to the authorities
without
      having downloaded their descriptors yet. Enabled for morial,
      moria2, tor26, and lefkada now in the default configuration.
    - Directory servers are more willing to send a 503 "busy" if
they
      are near their write limit, especially for v1 directory
requests.
      Now they can use their limited bandwidth for actual Tor
traffic.
    - Clients track responses with status 503 from dirservers.
After a
      dirserver has given us a 503, we try not to use it until an
hour has
      gone by, or until we have no dirservers that haven't given
us a 503.
    - When we get a 503 from a directory, and we're not a server,
we don't
      count the failure against the total number of failures
allowed
      for the thing we're trying to download.
    - Report X-Your-Address-Is correctly from tunneled directory
      connections; don't report X-Your-Address-Is when it's an
internal
      address; and never believe reported remote addresses when
they're
      internal.
    - Protect against an unlikely DoS attack on directory servers.
    - Add a BadDirectory flag to network status docs so that
authorities
      can (eventually) tell clients about caches they believe to
be
      broken.

  o Minor features (controller):
    - Have GETINFO dir/status/* work on hosts with DirPort
disabled.
    - Reimplement GETINFO so that info/names stays in sync with
the
      actual keys.
    - Implement "GETINFO fingerprint".
    - Implement "SETEVENTS GUARD" so controllers can get updates
on
      entry guard status as it changes.

  o Minor features (clean up obsolete pieces):
    - Remove some options that have been deprecated since at least
      0.1.0.x: AccountingMaxKB, LogFile, DebugLogFile, LogLevel,
and
      SysLog. Use AccountingMax instead of AccountingMaxKB, and
use Log
      to set log options.
    - We no longer look for identity and onion keys in
"identity.key" and
      "onion.key" -- these were replaced by secret_id_key and
      secret_onion_key in 0.0.8prel.
    - We no longer require unrecognized directory entries to be
      preceded by "opt".

  o Major bugfixes (security):
    - Stop sending the HttpProxyAuthenticator string to directory
      servers when directory connections are tunnelled through
```

```
Tor.
    - Clients no longer store bandwidth history in the state file.
    - Do not log introduction points for hidden services if
SafeLogging
      is set.
    - When generating bandwidth history, round down to the nearest
      1k. When storing accounting data, round up to the nearest
1k.
    - When we're running as a server, remember when we last
rotated onion
      keys, so that we will rotate keys once they're a week old
even if
      we never stay up for a week ourselves.

  o Major bugfixes (other):
    - Fix a longstanding bug in eventdns that prevented the count
of
      timed-out resolves from ever being reset. This bug caused us
to
      give up on a nameserver the third time it timed out, and try
it
      10 seconds later... and to give up on it every time it timed
out
      after that.
    - Take out the '5 second' timeout from the connection retry
      schedule. Now the first connect attempt will wait a full 10
      seconds before switching to a new circuit. Perhaps this will
help
      a lot. Based on observations from Mike Perry.
    - Fix a bug on the Windows implementation of tor_mmap_file()
that
      would prevent the cached-routers file from ever loading.
Reported
      by John Kimble.

  o Minor bugfixes:
    - Fix an assert failure when a directory authority sets
      AuthDirRejectUnlisted and then receives a descriptor from an
      unlisted router. Reported by seeess.
    - Avoid a double-free when parsing malformed DirServer lines.
    - Fix a bug when a BSD-style PF socket is first used. Patch
from
      Fabian Keil.
    - Fix a bug in 0.1.2.2-alpha that prevented clients from
asking
      to resolve an address at a given exit node even when they
ask for
      it by name.
    - Servers no longer ever list themselves in their "family"
line,
      even if configured to do so. This makes it easier to
configure
      family lists conveniently.
    - When running as a server, don't fall back to 127.0.0.1 when
no
      nameservers are configured in /etc/resolv.conf; instead,
make the
      user fix resolv.conf or specify nameservers explicitly.
(Resolves
      bug 363.)
    - Stop accepting certain malformed ports in configured exit
policies.
    - Don't re-write the fingerprint file every restart, unless it
has
      changed.
    - Stop warning when a single nameserver fails: only warn when
_all_ of
      our nameservers have failed. Also, when we only have one
nameserver,
      raise the threshold for deciding that the nameserver is
dead.
    - Directory authorities now only decide that routers are
reachable
      if their identity keys are as expected.
    - When the user uses bad syntax in the Log config line, stop
      suggesting other bad syntax as a replacement.
    - Correctly detect ipv6 DNS capability on OpenBSD.

  o Minor bugfixes (controller):
    - Report the circuit number correctly in STREAM CLOSED events.
Bug
      reported by Mike Perry.
    - Do not report bizarre values for results of accounting
GETINFOs
      when the last second's write or read exceeds the allotted
bandwidth.
    - Report "unrecognized key" rather than an empty string when
the
      controller tries to fetch a networkstatus that doesn't
exist.
```

From: Roger Dingledine
To: Kelly DeYoe
Cc: Ken Berman; Shava Nerad; Hiu Ho
Subject: Re: First draft of blocking-resistance design
Date: Thursday, February 08, 2007 4:45:08 PM

On Tue, Jan 30, 2007 at 02:20:05PM -0500, Kelly DeYoe wrote:
> Ken didn't realize it, but I'm actually going to be out on Friday and
> all week next week, so we'll probably need to postpone this until the
> week of 2/12.

Hi folks,

I am now back from the dead. Sorry for the scare. :(

I'm out of town from 2/12 to 2/22, and also on 2/26 and 2/27. Should we schedule for 2/23 or 3/1 or 3/2?

In the mean time here are the development release notes for the last few releases, to give you a sense of what we've been up to. We're trying to stabilize the Tor 0.1.2.x release, which has lots of new features (and still a few bugs), before we go on.

Thanks,
--Roger

Changes in version 0.1.2.7-alpha - 2007-02-06

o Major bugfixes (rate limiting):

- Servers decline directory requests much more aggressively when they're low on bandwidth. Otherwise they end up queueing more and more directory responses, which can't be good for latency.
- But never refuse directory requests from local addresses.
- Fix a memory leak when sending a 503 response for a networkstatus request.
- Be willing to read or write on local connections (e.g. controller connections) even when the global rate limiting buckets are empty.
- If our system clock jumps back in time, don't publish a negative uptime in the descriptor. Also, don't let the global rate limiting buckets go absurdly negative.
- Flush local controller connection buffers periodically as we're writing to them, so we avoid queueing 4+ megabytes of data before trying to flush.

o Major bugfixes (NT services):

- Install as NT_AUTHORITY\LocalService rather than as SYSTEM; add a command-line flag so that admins can override the default by saying "tor --service install --user "SomeUser"". This will not affect existing installed services. Also, warn the user that the service will look for its configuration file in the service user's %appdata% directory. (We can't do the 'hardwire the user's appdata directory' trick any more, since we may not have read access to that directory.)

o Major bugfixes (other):

- Previously, we would cache up to 16 old networkstatus documents indefinitely, if they came from nontrusted authorities. Now we discard them if they are more than 10 days old.
- Fix a crash bug in the presence of DNS hijacking (reported by Andrew

Del Vecchio).

- Detect and reject malformed DNS responses containing circular pointer loops.
- If exits are rare enough that we're not marking exits as guards, ignore exit bandwidth when we're deciding the required bandwidth to become a guard.
- When we're handling a directory connection tunneled over Tor, don't fill up internal memory buffers with all the data we want to tunnel; instead, only add it if the OR connection that will eventually receive it has some room for it. (This can lead to slowdowns in tunneled dir connections; a better solution will have to wait for 0.2.0.)

o Minor bugfixes (dns):

- Add some defensive programming to eventdns.c in an attempt to catch possible memory-stomping bugs.
- Detect and reject DNS replies containing IPv4 or IPv6 records with an incorrect number of bytes. (Previously, we would ignore the extra bytes.)
- Fix as-yet-unused reverse IPv6 lookup code so it sends nybbles in the correct order, and doesn't crash.
- Free memory held in recently-completed DNS lookup attempts on exit. This was not a memory leak, but may have been hiding memory leaks.
- Handle TTL values correctly on reverse DNS lookups.
- Treat failure to parse resolv.conf as an error.

o Minor bugfixes (other):

- Fix crash with "tor --list-fingerprint" (reported by seeess).
- When computing clock skew from directory HTTP headers, consider what time it was when we finished asking for the directory, not what time it is now.
- Expire socks connections if they spend too long waiting for the handshake to finish. Previously we would let them sit around for days, if the connecting application didn't close them either.
- And if the socks handshake hasn't started, don't send a "DNS resolve socks failed" handshake reply; just close it.
- Stop using C functions that OpenBSD's linker doesn't like.
- Don't launch requests for descriptors unless we have networkstatuses from at least half of the authorities. This delays the first download slightly under pathological circumstances, but can prevent us from downloading a bunch of descriptors we don't need.
- Do not log IPs with TLS failures for incoming TLS connections. (Fixes bug 382.)
- If the user asks to use invalid exit nodes, be willing to use unstable ones.
- Stop using the reserved ac_cv namespace in our configure script.
- Call stat() slightly less often; use fstat() when possible.
- Refactor the way we handle pending circuits when an OR connection completes or fails, in an attempt to fix a rare crash bug.
- Only rewrite a conn's address based on X-Forwarded-For: headers if it's a parseable public IP address; and stop adding extra quotes to the resulting address.

o Major features:

- Weight directory requests by advertised bandwidth. Now we can let servers enable write limiting but still allow most clients to succeed at their directory requests. (We still ignore weights when choosing a directory authority; I hope this is a feature.)

o Minor features:

- Create a new file ReleaseNotes which was the old ChangeLog. The new ChangeLog file now includes the summaries for all development versions too.
- Check for addresses with invalid characters at the exit as well as at the client, and warn less verbosely when they fail. You can override this by setting ServerDNSAllowNonRFC953Addresses to 1.
- Adapt a patch from goodell to let the contrib/exitlist script take arguments rather than require direct editing.
- Inform the server operator when we decide not to advertise a DirPort due to AccountingMax enabled or a low BandwidthRate. It was confusing Zax, so now we're hopefully more helpful.
- Bring us one step closer to being able to establish an encrypted directory tunnel without knowing a descriptor first. Still not ready yet. As part of the change, now assume we can use a create_fast cell if we don't know anything about a router.
- Allow exit nodes to use nameservers running on ports other than 53.
- Servers now cache reverse DNS replies.
- Add an --ignore-missing-torrc command-line option so that we can get the "use sensible defaults if the configuration file doesn't exist" behavior even when specifying a torrc location on the command line.

o Minor features (controller):

- Track reasons for OR connection failure; make these reasons available via the controller interface. (Patch from Mike Perry.)
- Add a SOCKS_BAD_HOSTNAME client status event so controllers can learn when clients are sending malformed hostnames to Tor.
- Clean up documentation for controller status events.
- Add a REMAP status to stream events to note that a stream's address has changed because of a cached address or a MapAddress directive.

Changes in version 0.1.2.6-alpha - 2007-01-09

o Major bugfixes:

- Fix an assert error introduced in 0.1.2.5-alpha: if a single TLS connection handles more than 4 gigs in either direction, we crash.
- Fix an assert error introduced in 0.1.2.5-alpha: if we're an advertised exit node, somebody might try to exit from us when we're bootstrapping and before we've built our descriptor yet. Refuse the connection rather than crashing.

o Minor bugfixes:

- Warn if we (as a server) find that we've resolved an address that we weren't planning to resolve.
- Warn that using select() on any libevent version before 1.1 will be unnecessarily slow (even for select()).
- Flush ERR-level controller status events just like we currently flush ERR-level log events, so that a Tor shutdown doesn't prevent the controller from learning about current events.

o Minor features (more controller status events):

- Implement EXTERNAL_ADDRESS server status event so controllers can learn when our address changes.
- Implement BAD_SERVER_DESCRIPTOR server status event so controllers can learn when directories reject our descriptor.
- Implement SOCKS_UNKNOWN_PROTOCOL client status event so controllers can learn when a client application is speaking a non-socks protocol to our SocksPort.
- Implement DANGEROUS_SOCKS client status event so controllers

```
      can learn when a client application is leaking DNS addresses.
    - Implement BUG general status event so controllers can learn when
      Tor is unhappy about its internal invariants.
    - Implement CLOCK_SKEW general status event so controllers can learn
      when Tor thinks the system clock is set incorrectly.
    - Implement GOOD_SERVER_DESCRIPTOR and ACCEPTED_SERVER_DESCRIPTOR
      server status events so controllers can learn when their descriptors
      are accepted by a directory.
    - Implement CHECKING_REACHABILITY and REACHABILITY_{SUCCEEDED|FAILED}
      server status events so controllers can learn about Tor's progress in
      deciding whether it's reachable from the outside.
    - Implement BAD_LIBEVENT general status event so controllers can learn
      when we have a version/method combination in libevent that needs to
      be changed.
    - Implement NAMESERVER_STATUS, NAMESERVER_ALL_DOWN, DNS_HIJACKED,
      and DNS_USELESS server status events so controllers can learn
      about changes to DNS server status.

  o Minor features (directory):
    - Authorities no longer recommend exits as guards if this would shift
      too much load to the exit nodes.


Changes in version 0.1.2.5-alpha - 2007-01-06
  o Major features:
    - Enable write limiting as well as read limiting. Now we sacrifice
      capacity if we're pushing out lots of directory traffic, rather
      than overrunning the user's intended bandwidth limits.
    - Include TLS overhead when counting bandwidth usage; previously, we
      would count only the bytes sent over TLS, but not the bytes used
      to send them.
    - Support running the Tor service with a torrc not in the same
      directory as tor.exe and default to using the torrc located in
      the %appdata%\Tor\ of the user who installed the service. Patch
      from Matt Edman.
    - Servers now check for the case when common DNS requests are going to
      wildcarded addresses (i.e. all getting the same answer), and change
      their exit policy to reject *:* if it's happening.
    - Implement BEGIN_DIR cells, so we can connect to the directory
      server via TLS to do encrypted directory requests rather than
      plaintext. Enable via the TunnelDirConns and PreferTunneledDirConns
      config options if you like.

  o Minor features (config and docs):
    - Start using the state file to store bandwidth accounting data:
      the bw_accounting file is now obsolete. We'll keep generating it
      for a while for people who are still using 0.1.2.4-alpha.
    - Try to batch changes to the state file so that we do as few
      disk writes as possible while still storing important things in
      a timely fashion.
    - The state file and the bw_accounting file get saved less often when
      the AvoidDiskWrites config option is set.
    - Make PIDFile work on Windows (untested).
    - Add internal descriptions for a bunch of configuration options:
      accessible via controller interface and in comments in saved
      options files.
    - Reject *:563 (NNTPS) in the default exit policy. We already reject
      NNTP by default, so this seems like a sensible addition.
    - Clients now reject hostnames with invalid characters. This should
      avoid some inadvertent info leaks. Add an option
```

AllowNonRFC953Hostnames to disable this behavior, in case somebody is running a private network with hosts called @, !, and #.

- Add a maintainer script to tell us which options are missing documentation: "make check-docs".
- Add a new address-spec.txt document to describe our special-case addresses: .exit, .onion, and .noconnnect.

o Minor features (DNS):

- Ongoing work on eventdns infrastructure: now it has dns server and ipv6 support. One day Tor will make use of it.
- Add client-side caching for reverse DNS lookups.
- Add support to tor-resolve tool for reverse lookups and SOCKS5.
- When we change nameservers or IP addresses, reset and re-launch our tests for DNS hijacking.

o Minor features (directory):

- Authorities now specify server versions in networkstatus. This adds about 2% to the side of compressed networkstatus docs, and allows clients to tell which servers support BEGIN_DIR and which don't. The implementation is forward-compatible with a proposed future protocol version scheme not tied to Tor versions.
- DirServer configuration lines now have an orport= option so clients can open encrypted tunnels to the authorities without having downloaded their descriptors yet. Enabled for moria1, moria2, tor26, and lefkada now in the default configuration.
- Directory servers are more willing to send a 503 "busy" if they are near their write limit, especially for v1 directory requests. Now they can use their limited bandwidth for actual Tor traffic.
- Clients track responses with status 503 from dirservers. After a dirserver has given us a 503, we try not to use it until an hour has gone by, or until we have no dirservers that haven't given us a 503.
- When we get a 503 from a directory, and we're not a server, we don't count the failure against the total number of failures allowed for the thing we're trying to download.
- Report X-Your-Address-Is correctly from tunneled directory connections; don't report X-Your-Address-Is when it's an internal address; and never believe reported remote addresses when they're internal.
- Protect against an unlikely DoS attack on directory servers.
- Add a BadDirectory flag to network status docs so that authorities can (eventually) tell clients about caches they believe to be broken.

o Minor features (controller):

- Have GETINFO dir/status/* work on hosts with DirPort disabled.
- Reimplement GETINFO so that info/names stays in sync with the actual keys.
- Implement "GETINFO fingerprint".
- Implement "SETEVENTS GUARD" so controllers can get updates on entry guard status as it changes.

o Minor features (clean up obsolete pieces):

- Remove some options that have been deprecated since at least 0.1.0.x: AccountingMaxKB, LogFile, DebugLogFile, LogLevel, and SysLog. Use AccountingMax instead of AccountingMaxKB, and use Log to set log options.
- We no longer look for identity and onion keys in "identity.key" and "onion.key" -- these were replaced by secret_id_key and secret_onion_key in 0.0.8pre1.
- We no longer require unrecognized directory entries to be

preceded by "opt".

o Major bugfixes (security):

- Stop sending the HttpProxyAuthenticator string to directory servers when directory connections are tunnelled through Tor.
- Clients no longer store bandwidth history in the state file.
- Do not log introduction points for hidden services if SafeLogging is set.
- When generating bandwidth history, round down to the nearest 1k. When storing accounting data, round up to the nearest 1k.
- When we're running as a server, remember when we last rotated onion keys, so that we will rotate keys once they're a week old even if we never stay up for a week ourselves.

o Major bugfixes (other):

- Fix a longstanding bug in eventdns that prevented the count of timed-out resolves from ever being reset. This bug caused us to give up on a nameserver the third time it timed out, and try it 10 seconds later... and to give up on it every time it timed out after that.
- Take out the '5 second' timeout from the connection retry schedule. Now the first connect attempt will wait a full 10 seconds before switching to a new circuit. Perhaps this will help a lot. Based on observations from Mike Perry.
- Fix a bug on the Windows implementation of tor_mmap_file() that would prevent the cached-routers file from ever loading. Reported by John Kimble.

o Minor bugfixes:

- Fix an assert failure when a directory authority sets AuthDirRejectUnlisted and then receives a descriptor from an unlisted router. Reported by seeess.
- Avoid a double-free when parsing malformed DirServer lines.
- Fix a bug when a BSD-style PF socket is first used. Patch from Fabian Keil.
- Fix a bug in 0.1.2.2-alpha that prevented clients from asking to resolve an address at a given exit node even when they ask for it by name.
- Servers no longer ever list themselves in their "family" line, even if configured to do so. This makes it easier to configure family lists conveniently.
- When running as a server, don't fall back to 127.0.0.1 when no nameservers are configured in /etc/resolv.conf; instead, make the user fix resolv.conf or specify nameservers explicitly. (Resolves bug 363.)
- Stop accepting certain malformed ports in configured exit policies.
- Don't re-write the fingerprint file every restart, unless it has changed.
- Stop warning when a single nameserver fails: only warn when _all_ of our nameservers have failed. Also, when we only have one nameserver, raise the threshold for deciding that the nameserver is dead.
- Directory authorities now only decide that routers are reachable if their identity keys are as expected.
- When the user uses bad syntax in the Log config line, stop suggesting other bad syntax as a replacement.
- Correctly detect ipv6 DNS capability on OpenBSD.

o Minor bugfixes (controller):

- Report the circuit number correctly in STREAM CLOSED events. Bug reported by Mike Perry.

- Do not report bizarre values for results of accounting GETINFOs when the last second's write or read exceeds the allotted bandwidth.
- Report "unrecognized key" rather than an empty string when the controller tries to fetch a networkstatus that doesn't exist.

From: Shava Nerad
To: Ken Berman; Roger Dingledine
Cc: Kelly DeYoe; Shava Nerad; Hiu Ho
Subject: Re: First draft of blocking-resistance design
Date: Tuesday, January 30, 2007 1:18:40 PM

At 08:43 AM 1/30/2007, Ken Berman wrote:
>Is it time for a status phone call? Th Feb 1st at 3:00??

Hi, Ken et al! Could we schedule this for early next week? Roger is on the mend from being ill, and probably won't be up and about until late this week, at best. His doctor has ordered him off the computer entirely at the moment -- obviously drastic measures...:) Still he does anticipate being back among the living in a couple/few days.

As always, anything administrative, y'all can always tap me directly.

Thanks!

--

Shava Nerad
Executive Director
The Tor Project
http://tor.eff.org/
http://blogs.law.harvard.edu/anonymous/
(b) (6)
(b) (6)
(b) (6) (cell)
skype: shava23

From: Kelly DeYoe
To: Roger Dingledine
Cc: Ken Berman; Shava Nerad; Hiu Ho
Subject: Re: First draft of blocking-resistance design
Date: Monday, November 20, 2006 5:29:51 PM

Thanks for the updated version, Roger, we will be sure to review it as well.

I'm afraid Bennett is not currently under contract with us, so if we had
a need for his services at this time, one of us would have to start a
new contract with him.

-k

Roger Dingledine wrote:
> On Mon, Nov 20, 2006 at 09:16:35AM -0500, Ken Berman wrote:
>
>>Roger - our team went thru this last Friday and were quite impressed.
>
>
> Great. This new version has an ending, and other fixes and ideas sprinkled
> throughout, compared to the one you read last Friday.
>
>
>>We have some thoughts as to how this fits into your overall road map and
>>our funding desires. Next time we talk we can explore this.
>
>
> Great. There's lots to be done. I figure my job is to lay out what
> needs to be done and figure out how quickly we can do it, and then
> everything else will follow. :)
>
>
>>btw - this guy "fred", aka "farid pouya" reports on the translation to
>>Farsi:
>>In a week. We have been preparing launch of new site then everything
>>was and still quite crazy. Sorry for delay. In a wee k it will be done.
>>
>>We'll see...........
>
>
> When he starts the translation, he should check out the latest Vidalia
> from the repository -- somebody gave us a partial Farsi translation last
> week, so it is at least a starting point.
>
>
>>>Should we bring Bennett back in at this point?
>
>
> Any thoughts on this one? I don't want to cut him totally out of the
> loop; I just realized that I needed to write the "how Tor offers now"
> section before he could be more useful. But I figure I should check with
> you first, since I haven't heard from him in a few months and you might
> be using him for other things at this point.
>
> Thanks!
> --Roger
>

From: Roger Dingledine
To: Kelly DeYoe
Cc: Shava Nerad; Ken Berman; Hiu Ho
Subject: Re: First draft of blocking-resistance design
Date: Tuesday, January 23, 2007 3:56:17 AM

On Mon, Nov 20, 2006 at 08:17:03AM -0500, Roger Dingledine wrote:
> Take a look at
> http://tor.eff.org/svn/trunk/doc/design-paper/blocking.pdf
> for our first draft on how to adapt Tor to have a blocking-resistance
> component.

Hi folks,

I presented this design a few weeks ago at the CCC congress in Berlin, and there's a snazzy video here:

http://freehaven.net/~arma/23C3-1444-en-tor_and_china.m4v
http://freehaven.net/~arma/slides-23c3.pdf

Ken will be happy to learn that I have started using VoA as one of the example websites in my talks. :)

Thanks,
--Roger

From: Roger Dingledine
To: Ken Berman
Cc: Kelly DeYoe; Shava Nerad; Hiu Ho
Subject: Re: First draft of blocking-resistance design
Date: Monday, November 20, 2006 10:59:15 AM

On Mon, Nov 20, 2006 at 09:16:35AM -0500, Ken Berman wrote:
> Roger - our team went thru this last Friday and were quite impressed.

Great. This new version has an ending, and other fixes and ideas sprinkled throughout, compared to the one you read last Friday.

> We have some thoughts as to how this fits into your overall road map and
> our funding desires. Next time we talk we can explore this.

Great. There's lots to be done. I figure my job is to lay out what needs to be done and figure out how quickly we can do it, and then everything else will follow. :)

> btw - this guy "fred", aka "farid pouya" reports on the translation to
> Farsi:
> In a week. We have been preparing launch of new site then everything
> was and still quite crazy. Sorry for delay. In a wee k it will be done.
>
> We'll see..........

When he starts the translation, he should check out the latest Vidalia from the repository -- somebody gave us a partial Farsi translation last week, so it is at least a starting point.

> >Should we bring Bennett back in at this point?

Any thoughts on this one? I don't want to cut him totally out of the loop; I just realized that I needed to write the "how Tor offers now" section before he could be more useful. But I figure I should check with you first, since I haven't heard from him in a few months and you might be using him for other things at this point.

Thanks!
--Roger

From: Ken Berman
To: Roger Dingledine
Cc: Kelly DeYoe; Shava Nerad; Hiu Ho
Subject: Re: First draft of blocking-resistance design
Date: Monday, November 20, 2006 9:16:35 AM

Roger - our team went thru this last Friday and were quite impressed. We have some thoughts as to how this fits into your overall road map and our funding desires. Next time we talk we can explore this.

btw - this guy "fred", aka "farid pouya" reports on the translation to Farsi:
In a week. We have been preparing launch of new site then everything was and still quite crazy. Sorry for delay. In a wee k it will be done.

We'll see...........

Ken

Roger Dingledine wrote:

>Hi Kelly, Ken, Hiu,
>
>Take a look at
>http://tor.eff.org/svn/trunk/doc/design-paper/blocking.pdf
>for our first draft on how to adapt Tor to have a blocking-resistance
>component.
>
>Please don't publicize it widely yet (we want to go through a few more
>iterations first). Comments and thoughts appreciated -- on which parts
>don't make sense and need more explanation, on which parts claim wrong
>things or make bad assumptions, on which parts need better solutions,
>etc.
>
>Should we bring Bennett back in at this point?
>
>A roadmap for how to start building-and-deploying this in 2007 will
>follow soon, hopefully sometime this week.
>
>Thanks!
>--Roger
>
>
>

From:	Shava Nerad
To:	Ken Berman; Roger Dingledine
Cc:	Kelly DeYoe; Shava Nerad; Hiu Ho; Betty Pruitt
Subject:	Re: First draft of blocking-resistance design
Date:	Tuesday, February 13, 2007 2:47:25 PM

At 02:34 PM 2/13/2007, Ken Berman wrote:
>Let's do a conf call later in the week or early next week, ladies
>and gents. Ken

Alas, now Roger is in Ecuador with his fiance, and pretty much unreachable. Let's set this up for as soon as he gets back, which I think is the end of next week. If there are any administrative details I can take care of them?

Thanks!

--

Shava Nerad
Executive Director
The Tor Project
http://tor.eff.org/
http://blogs.law.harvard.edu/anonymous/
(b) (6)
(b) (6)
[redacted] (cell)
skype: shava23

From:	Roger Dingledine
To:	Kelly DeYoe
Cc:	Ken Berman; Shava Nerad; Hiu Ho
Subject:	Re: First draft of blocking-resistance design
Date:	Monday, November 27, 2006 7:33:38 AM

On Tue, Nov 21, 2006 at 04:24:20PM -0500, Roger Dingledine wrote:
> I will aim to get a copy of the roadmap to you (most of the tasks,
> broken down by when we'd like to aim to do them) by then too.

http://freehaven.net/~arma/roadmap-blocking.pdf

It's still a work in progress as you can see, but it includes all of the topics that we need to tackle, and has some notion of timeframe and priorities. As you can see, there's a lot to do. :)

Let me know if you have questions or want more details on any part of it.

Thanks!
--Roger

From:	Demetria. Anderson
To:	Roger Dingledine
Cc:	Kelly DeYoe
Subject:	Re: First invoice for Moria Research Labs, BBGCON1806S6149
Date:	Friday, August 04, 2006 12:26:27 PM

Hey Roger,

Please add the following to make a complete invoice for the finance office as follows:

Invoice number

Your telephone number

Specify the services rendered

Add Fiscal Data: 9568-06-0206-E009601067-454000-4335-2544

In your chart change hours to months

Total: $20,000 or $40,000 not clear to me

(All else remains the same)

Thanks!

From: Roger Dingledine
To: Demetria Anderson
Cc: Kelly DeYoe
Subject: Re: First invoice for Moria Research Labs, BBGCON1806S6149
Date: Friday, August 04, 2006 12:44:01 PM
Attachments: mrl-ibb1.pdf

Hi Demetria,

Attached is the second iteration of the invoice. Please let me know if there's anything more I should add. :)

Thanks!
--Roger

From: Kelly DeYoe
To: Andrew Lewman
Cc: Ken Berman; Melissa Gilroy
Subject: Re: Fw: IBB
Date: Friday, April 08, 2011 3:01:51 PM

I'm checking over my own records right now, and have also passed this along to Malita Dyson, who handles processing of our invoices for payment to check against her records and in our accounting system to figure out what happened here.

As you may know, funding for the federal government expires tonight, so unless another continuing resolution or the full budget is passed, we may be shutdown starting on Monday, and further investigation of this would be delayed. (Payments have already been obligated for your contract though of course, so you will be paid, the shutdown may just cause delays since we have to investigate the missing payments.)

-k

Andrew Lewman wrote:
> Hello Ken and Kelly,
>
> Melissa, Tor's CFO, noticed that we're missing 2 payments from your
> organization.
>
> Anything I can do to help?
>
> Thanks!
>
> Begin forwarded message:
>
> Date: Fri, 8 Apr 2011 11:03:17 -0400
> From: "Melissa Gilroy" <(b) (6)
> To: "'Andrew Lewman'" <(b) (6)
> Subject: IBB
>
>
> IBB paid us:
>
> 1/11/2011 - $15k for invoice # 31 (November's service)
> 3/3/2011 - $15k for invoice #33 (January service)
>
> We are missing payment for invoice #32 for December 2010 and #34 for
> February services.
>
> Melissa
>

From: Kelly DeYoe
To: Andrew Lewman
Cc: Ken Berman; Melissa Gilroy
Subject: Re: Fw: IBB
Date: Tuesday, April 12, 2011 4:38:21 PM

Andrew, based on my own review and the review of our admin. officer who processes the payments, Malita Dyson, we have no record of receiving invoice #32 (12/2010) from Tor. We have processed invoices #31, #33, #34, but #32 was skipped. If you can please submit it to us, it will be processed for payment.

When submitting invoices, you may wish include both Malita and myself on the email in the future, as I don't have any email trail of invoices since they have been sent to Malita exclusively, and she just presents me with the hardcopy to approve and sign.

-k

Andrew Lewman wrote:
> On Fri, Apr 08, 2011 at 02:01:51PM -0400, [redacted] (b) (6) wrote 1.2K bytes in 43 lines about:
> : I'm checking over my own records right now, and have also passed
> : this along to Malita Dyson, who handles processing of our invoices
> : for payment to check against her records and in our accounting
> : system to figure out what happened here.
>
> It appears something arrived today. I'm not sure which month was paid,
> so now we have 1 payment outstanding, and 1 new payment submitted for
> March.
>
> Thanks!
>

From: Ken Berman
To: Roger Dingledine
Cc: Kelly DeYoe; (b) (6)
Subject: RE: FW: Small Sister project anonymizes mail using Tor
Date: Friday, January 02, 2009 8:14:23 AM

Yes, completely.... thx,

-----Original Message-----
From: Roger Dingledine [mailto:(b) (6)
Sent: Friday, January 02, 2009 6:16 AM
To: Ken Berman
Cc: Kelly DeYoe; (b) (6)
Subject: Re: FW: Small Sister project anonymizes mail using Tor

On Tue, Dec 30, 2008 at 01:14:01PM -0500, Ken Berman wrote:
> Roger - any thoughts on this? Ken

We like the idea behind Small Sister. We're not sure that it's ready for actual users though.

I talked to Peter Roozemaal, the developer, at 25C3 in Berlin this week.
I tried to stress (as usual I guess :) the importance of transparency, writing down the protocol, and getting feedback from the broader security community. Peter really does want to get feedback, but he is taking the "first I'll build it, then I'll try to get people to use it, then I'll ask people if what I've built is any good" approach.

I asked Steven Murdoch to take a look at it, and this was his response:

"I looked at the SmallSister code. It's not good. No input validation, but I think I can only crash it. Man-in-the-middle'able (though since they use Tor hidden services, they get partial resistance). The challenge response protocol uses a two byte challenge (which they then add stuff like the time to, but that doesn't make a big deal). You can also get the client to sign arbitrary text."

He summarizes with: "The whole design is weird. They have invented their own mail protocol. If I were doing this, I'd use IMAP/SMTP-TLS."

When I was talking to Peter, I was also confused by their approach. One of his points was that smtp is hard to sanitize (that is, even if you use smtp over Tor, the smtp headers themselves are going to give away all sorts of sensitive info). So they wanted a new protocol. I told him what they really wanted was a converter -- dangerous smtp to acceptable smtp.

Steven adds "It uses GPGME, which is scary. This is a library that parses the output of GnuPG to work out what's going on. I worry that if you send it malformed messages the library will get confused. Say you give it a message with a wrong signature, but add some control characters that make the output parser think the message validated."

So, on the one hand I don't want to be too negative -- it's certainly not right to give up on it at this point. But on the other hand, they need to think harder about the security of their design as well as implementation -- and it's not clear that the team they have so far includes all the right people for doing that.

Hope that helps,
--Roger

From: Roger Dingledine
To: Ken Berman
Cc: Kelly DeYoe; [redacted] (b)(6)
Subject: Re: FW: Small Sister project anonymizes mail using Tor
Date: Friday, January 02, 2009 6:15:36 AM

On Tue, Dec 30, 2008 at 01:14:01PM -0500, Ken Berman wrote:
> Roger - any thoughts on this? Ken

We like the idea behind Small Sister. We're not sure that it's ready for actual users though.

I talked to Peter Roozemaal, the developer, at 25C3 in Berlin this week. I tried to stress (as usual I guess :) the importance of transparency, writing down the protocol, and getting feedback from the broader security community. Peter really does want to get feedback, but he is taking the "first I'll build it, then I'll try to get people to use it, then I'll ask people if what I've built is any good" approach.

I asked Steven Murdoch to take a look at it, and this was his response:

"I looked at the SmallSister code. It's not good. No input validation, but I think I can only crash it. Man-in-the-middle'able (though since they use Tor hidden services, they get partial resistance). The challenge response protocol uses a two byte challenge (which they then add stuff like the time to, but that doesn't make a big deal). You can also get the client to sign arbitrary text."

He summarizes with: "The whole design is weird. They have invented their own mail protocol. If I were doing this, I'd use IMAP/SMTP-TLS."

When I was talking to Peter, I was also confused by their approach. One of his points was that smtp is hard to sanitize (that is, even if you use smtp over Tor, the smtp headers themselves are going to give away all sorts of sensitive info). So they wanted a new protocol. I told him what they really wanted was a converter -- dangerous smtp to acceptable smtp.

Steven adds "It uses GPGME, which is scary. This is a library that parses the output of GnuPG to work out what's going on. I worry that if you send it malformed messages the library will get confused. Say you give it a message with a wrong signature, but add some control characters that make the output parser think the message validated."

So, on the one hand I don't want to be too negative -- it's certainly not right to give up on it at this point. But on the other hand, they need to think harder about the security of their design as well as implementation -- and it's not clear that the team they have so far includes all the right people for doing that.

Hope that helps,
--Roger