Description of the typical setup of a Tor relay:

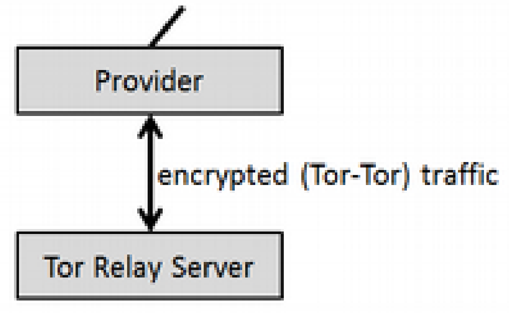

Description of the typical setup of a Tor exit node:

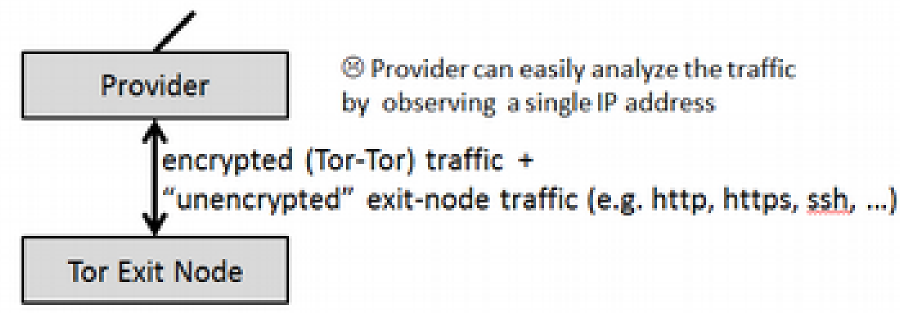
Provider
encrypted (Tor-Tor) traffic +
"unencrypted" exit-node traffic (e.g. http, https, ssh, ...)
Tor Exit Node

☹ Provider can easily analyze the traffic by observing a single IP address

Ideal setup:

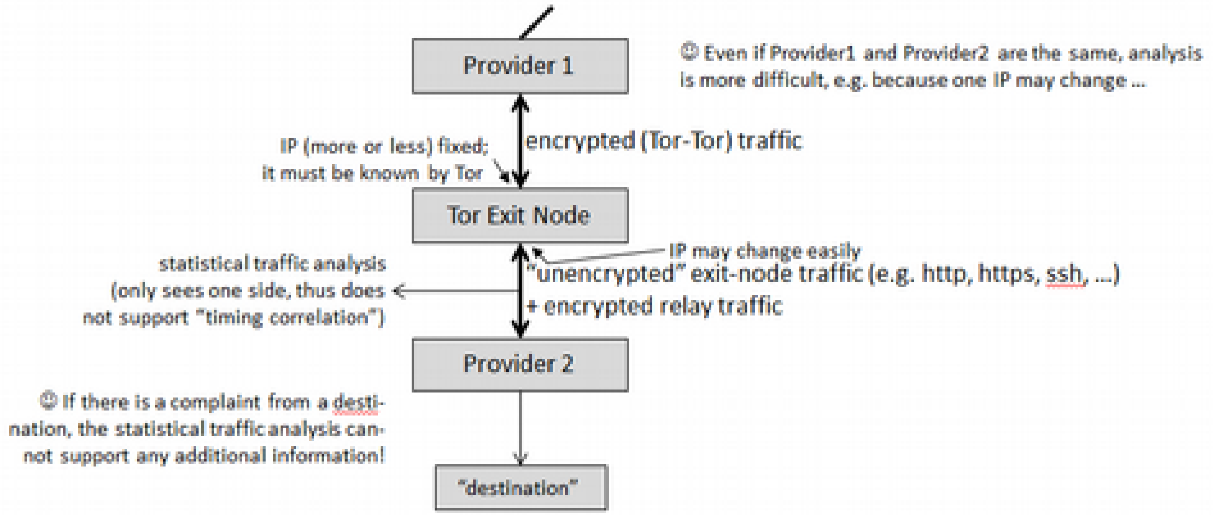
Provider 1
encrypted (Tor-Tor) traffic
Tor Exit Node
IP may change easily
"unencrypted" exit-node traffic (e.g. http, https, ssh, ...)
+ encrypted relay traffic
Provider 2
"destination"

☺ Even if Provider1 and Provider2 are the same, analysis is more difficult, e.g. because one IP may change ...

IP (more or less) fixed; it must be known by Tor

statistical traffic analysis (only sees one side, thus does not support "timing correlation")

☺ If there is a complaint from a destination, the statistical traffic analysis cannot support any additional information!

Reasons for this setup:

- Clear separation of input and output traffic
- Potential for multiple providers, further complicating attacks as more parties would need to get involved.
- Output interface could easily/frequently change its IP address or use multiple ones, hindering blocking.
- Output traffic can e.g. easily be sent through a tunnel to reappear on a completely different part of the Internet, thus more easily supporting a second provider.
- When statistics are generated they can be limited to the only ones potentially useful, i.e. the “unencrypted” traffic.
- If inquiries are made, no data on input traffic is present at all, preventing trace-back across the exit node based on information that is already available (if only for a brief time).
- Any orders to correlate input and output require changes to the network structure, modification of the exit node source code, or additional systems. Merely reconfiguring a single system or removing a “delete” or “filter” statement is not sufficient.
- Should the monitoring system get hacked, no correlation is possible as no input traffic is ever seen by that system.
- The patch further enables us to distinguish between exit and relay traffic, even more reducing monitoring for statistics/data used briefly/”readily” available to the only traffic of interest – which is available to anyone on the further route to the server in exactly this form anyway.
- A separate output simplifies IDS filtering of exit traffic to reduce attacks/abuse reports