IVPN Privacy Guides

Expert Guides from the leaders in VPN Privacy

Back to Privacy Guides

An Introduction to Tor vs I2P

Basicby Ed Holden

Darknet. The word in and of itself brings to mind visions of the seedy underbelly of the internet; a virtual red-light district, back alley, and digital ghetto all rolled into one. Despite this threatening image that the media and many governments would like to imprint on the public consciousness, privacy-aware individuals know that in todays world of ISP data retention being measured in petabytes and massive supercomputing resources being thrown at traffic analysis by both governments and private industry alike, individuals must take it upon themselves to ensure the freedoms that come with anonymous information access and communication. Two of the most popular tools for doing so on the internet are Tor and I2P. Both will be compared and contrasted below.

TOR

We will begin by examining the underlying technology of the Tor network with an eye towards how it works to protect your anonymity online. The Tor network is comprised of three different types of nodes: directory servers, exit points (also referred to as exit relays), and internal relays. When you connect to Tor, the first thing your client does is acquire a current list of relays from one of the trusted directory servers. The addresses of these servers are included with the basic configuration files shipped with the client (of course, as with any reputable privacy tool, you have the option to alter what directory servers you trust to provide you with valid relays).

After retrieving a list of currently operational relays from the directory servers, your client then determines the optimal route for your traffic across the Tor network and finally terminating (from the Tor network perspective) at an exit node. This circuit created consists of your computer, the relay to which you are connecting and multiple internal relays before reaching an exit node. Note that this is substantially different that the traditional IP forwarding that occurs between routers on the internet. Traditional IP routers follow a best possible route on a per-packet basis, there are no “stateful” circuits from an IP perspective (as a qualifier to this statement, it is necessary to grant that it is within the technical realm of possibility that every router between you and the computer you are connecting to could have single, static routes to one another, though in practice this is a near impossibility). In short, for the life of a circuit, all of your traffic will follow the same route within the Tor network and exit at the same point. Later, we will see how this is fundamentally different that the way the I2P network operates.

During the circuit creation process, your client exchanges cryptographic keys with the first relay it connects to and begins encrypting traffic back and forth. Further each hop in transit between the various relays is encrypted using those relays' cryptographic keys. You can visualize this as layers of encryption being wrapped around your data: this is where the phrase "onion routing" comes from when describing the type of network Tor establishes. Finally, your encrypted traffic is decrypted at the exit relay where it is then forwarded out onto the "regular" internet. This is one of the ways that Tor helps maintain your privacy online – each exit node is aggregating traffic from many other Tor users and putting it out onto the internet all at once. Your traffic becomes a small stream in the giant swath of data coming from and entering back into any given exit node. It is also important to note that your exit node only knows which intermediate node to send receiving data back to (this is also true for each internal to internal leg of the circuit). What this means is that your identity and the content of your traffic are cryptographically bifurcated – your entry node knows who you are but not what you are doing and your exit node knows what you are doing but not who you are. All the relays in between only know to forward the encrypted payload to the next relay on the circuit. Assuming that the content of your traffic does not reveal your identity, this permits you to browse the internet completely anonymously.

As a side note, Tor also allows you to run and access what are called "hidden services." These are servers that are accessible only from within the Tor network itself. While this is not the primary purpose for Tor, it does provide an opportunity for one to use dedicated in-network services in a cryptographically secure manner. Among the various hidden services are various blogs, email servers, and forums. We will see later how I2P provides a better framework for providing these hidden services, but if one's primary goal is to access "regular" internet services in a anonymous fashion, Tor is a vital tool in one's arsenal.

I2P

On the surface, I2P appears to provide many of the same benefits that Tor does. Both allow anonymous access to online content, both make use of a peer-to-peer-like routing structure, and both operate using layered encryption. However, I2P was designed from the ground up to provide a different set of benefits. As we saw above, the primary use case for Tor is enabling anonymous access of the public internet with hidden services as an ancillary benefit. I2P on the other hand, was designed from day one to be a true "darknet." Its primary function is to be a "network within the internet," with traffic staying contained in its borders. Very few outbound relays exist in the I2P network, and the few that do exist are rarely usable.

As mentioned above, I2P routes traffic differently than Tor. At its heart, I2P performs packet based routing as opposed to Tor's circuit based routing. This has the benefit of permitting I2P to dynamically route around congestion and service interruptions in a manner similar to the internet's IP routing. This provides a higher level of reliability and redundancy to the network itself. Additionally, I2P does not rely on a trusted directory service to get route information. Instead, network routes are formed and constantly updated dynamically, with each router constantly evaluating other routers and sharing what it finds. Finally, I2P establishes two independent simplex tunnels for traffic to traverse the network to and from each host as opposed to Tor’s formation of a single duplex circuit. This provides the additional benefit of only disclosing half the traffic in the case of an in-network eavesdropper.

From an application-level perspective there is a fundamental difference between the I2P and Tor networks as well. Tor functions by providing a proxy on your local machine that you must configure your applications to use (of download specially configured application bundles). In contrast, I2P is generally used by applications that are written specifically to run on the I2P network. These include, but are not limited to, instant message, file sharing, email, and distributed storage applications (yes, you can store encrypted data in the I2P “cloud,” similar to Freenet).

Conclusion

We see that both Tor and I2P provide cryptographically sound methods to anonymously access information and comunicate online. Tor provides one with better anonymous access to the open internet and I2P provides one with a more robust and reliable “network within the network,” a true darknet, if you will. Of course, when implementing either of these two tools, one must always be aware that one’s ISP can see that he or she is using Tor or I2P (though they cannot determine the content of the traffic itself). In order to hide this knowledge from one’s ISP, one should make use of a high-quality VPN service to act as an entry point to either one’s anonymous network of choice or to the internet at large.

Next Articles

- Applying Risk Management to Privacy
- Creating a VM within a hidden truecrypt partition
- How to perform a VPN leak test