

Tor: An Anonymous Routing Network for Covert On-line Operations

By Nicholas A. Fraser, Captain, USAF, Richard A. Raines, Major, USAF (Retired) and Rusty O. Baldwin, Major, USAF (Retired)

Center for Information Security Education and Research
Air Force Institute of Technology
Wright Patterson AFB

Editorial Abstract: The authors discuss the functionality of Tor, an anonymous Internet communication system with potential applications for both friend and foe in the area of computer network operations.

In 2002, a computer hacker belonging to a sophisticated hacker group gained access to an unclassified USAF computer system located at an unnamed Air Force base. A review of the logs from the victim computer system disclosed the hacker sent emails from the victim to an address registered to an U.S. Internet Service Provider (ISP). After obtaining necessary court orders, the ISP turned over to investigators the connection logs indicating the IP addresses used to check the account and also the content of emails still residing on the ISP's servers. An analysis of the connection logs disclosed the intruder likely resided in Romania and was using the account to record and save the IP addresses for over a hundred other compromised computers. Clearly, the success of the investigation, in terms of attribution and intelligence, was due to the hacker's ignorance of the ISP's policy of recording account connections. Had he known better, the hacker could have shielded from the ISP his true IP addresses and avoided arrest. All he needed to do was check his web email using the anonymous Internet communication system Tor.

The Tor system [1,5] provides anonymity to individuals using interactive Internet services like the World Wide Web (WWW), Internet Relay Chat (IRC), or secure shell (SSH). Tor is an overlay network, first introduced in 2002 and originally sponsored by the Naval Research Laboratory. It is now sponsored by the Electronic Frontier Foundation and developed under the Free Haven Project [4]. Tor provides anonymous message delivery with minimal latency by routing messages

through special servers called onion routers (ORs). These ORs are administered by volunteers with over 200 currently online in more than 20 countries. Users connect to Tor via an onion proxy (OP) that is installed on individual computer systems. When a user, Alice, wishes to send messages to a server, Bob, on the Internet, her OP constructs a circuit or path through three ORs (see Figure 1), the last of which is responsible for connecting to Bob. Once the circuit is established, Alice sends her messages.

The rest of this paper is organized as follows: First, a detailed description of the Tor circuit establishment and hidden services is discussed in Section 2. This is followed by a discussion of the application of Tor to enemy cyber operations and American Information Operations (IO) in Section 3.

How does Tor Work?

Tor is based on the concept of onion routing where messages are wrapped in layers of encryption before being sent. When a message arrives at the first OR on a circuit, the outer layer of encryption is removed and the message is forwarded to the next OR. This process is repeated until it reaches the final OR on the circuit. At this point, the message is decrypted (revealing clear text) and forwarded to the destination address. Layered encryption is a common technique used in anonymous communication systems, however, to achieve low-latency various other techniques that strengthen anonymity are omitted. Mixing is one such feature. A server is said to be a mix if it takes in a collection of messages called a pool, transforms (usually through encryption or decryption), reorders, and delays the pool, and finally flushes the pool by forwarding each message according to the new order. Mixing increases latency to the extent that interactive services cannot be supported. Instead, mixing is most commonly used by anonymous communication systems specializing in email delivery.

Tor operates using fixed 512 byte cells (or packets) for stronger anonymity and the Transport Layer Security (TLS) protocol for authentication and privacy. Figure 2 is a timeline showing the steps Alice takes to communicate with Bob. First, Alice's OP randomly selects three ORs. Next, Alice establishes a TLS connection to OR1. Alice then sends a create cell containing the first half of a secret key exchange. This message is encrypted using OR1's public key. OR1 responds to Alice with the second half of the secret key and a securely hashed value of the symmetric key. Alice can be

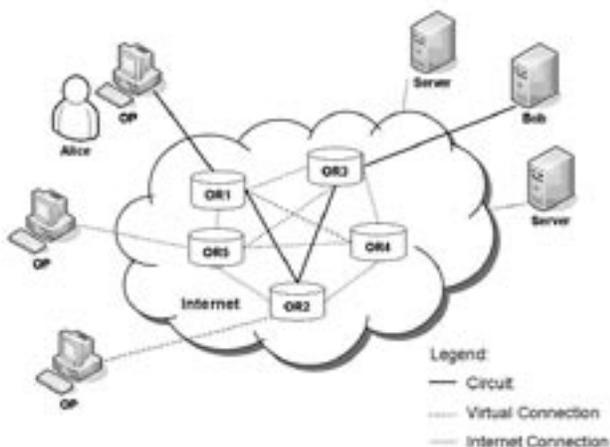
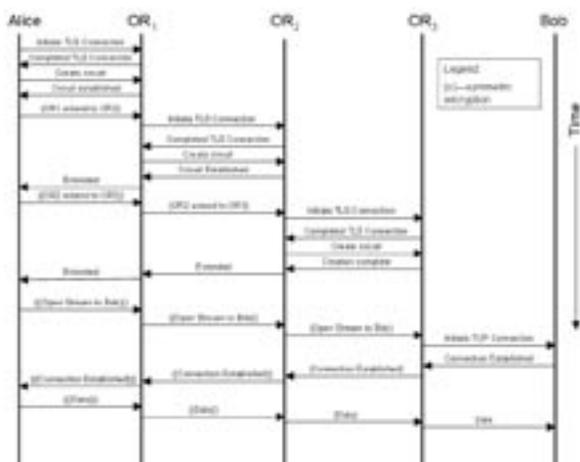


Figure 1: The Tor Overlay Network



Alice builds a three-hop circuit and sends a message to Bob

assured that OR1 is who he says he is and has done so without providing any information as to her identity or location. The above protocol creates a distinct circuit between Alice and OR1. This allows for multiple circuits to be established over the same TLS connection. Additionally, messages between Alice and OR1 can only be decrypted during the lifetime of the circuit (usually measured in minutes) and so if the circuit is destroyed (or recycled), messages can never be decrypted again. This capability is called forward secrecy.

If Alice wishes to extend the circuit to OR2, she sends OR1 a message containing the address of OR2 and the first half of the secret key exchange between Alice and OR2. This content of the message to OR1 is encrypted using OR2's public key. OR1 then establishes a TLS connection with OR2 and sends the encrypted key to OR2. OR2 responds to OR1 with the second half of the key exchange and a hash of the key. OR1 cannot determine the key because he was unable to observe the first half of the key exchange. Next, OR1 forwards the reply from OR2 to Alice. If the message is authentic, Alice and OR2 have a set of secret symmetric keys but without OR2 knowing he is communicating with Alice. This same process is similar for OR3 and once complete forms a circuit through Tor. Next, Alice tells OR3 to open a connection to Bob. Finally, when Alice has formed a message for Bob she wraps it in layered encryption first using OR3's secret key and concluding with OR1's. The encrypted message is then sent through the circuit, begin unwrapped as it traverses each OR until it reaches OR3 and is forwarded to Bob.

Tor also allows for hidden services to be established on the Internet. This means a web server for example can be accessed without the publisher identifying readers and also without readers identifying the publisher's IP address. If a user, Alice, wishes to establish a hidden service, she first advertises a collection of ORs as contact points for users wishing to access the service. If another user, Bob, wishes to connect to the hidden service, he must first protect his own identity by connecting to an OR that will serve as a rendezvous point (RP). Next, Bob connects to one of Alice's contact points and tells her of the rendezvous point he wishes to use to connect to the hidden service. If Alice agrees to provide service to Bob, she connects to the RP. An additional advantage to this capability is

that anyone wishing to bring down the service via a distributed denial of service (DDoS) attack is forced to target Tor.

Information Operations Applications

Tor is a tool that can aid unsophisticated hackers, terrorist organizations, and foreign information operators. URL-based attacks that take advantage of simple vulnerabilities in web servers, like the Unicode vulnerability in Microsoft's Internet Information Server 4.0 [2], can be effectively launched using Tor. Additionally, Tor could be used by an adversary to control botnets. Such attack networks are typically established by exploiting computer systems via malicious email attachments or web scripts. They are specifically targeted because they have constant Internet connectivity and their users are unlikely to notice additional connections and CPU degradation. The use of botnets is a growing problem as most present uses target the propagation of spam and adware. Future uses may be much more malicious. The combined use of botnets and systems such as Tor can have severe negative impacts on national defense and law enforcement. With Tor the adversary can send instructions to his "zombie" computers without concern that his command and control location will be discovered.

Terrorist organizations can also make good use of the Tor network. Tor could serve as a conduit to Internet communication channels known to be used by terrorist organizations like web pages and web-based email. Furthermore, Tor can be used to research targets and weapons construction techniques without fear of being located or identified. Finally, Tor's hidden services, intended to aid in countering government censorship, can be used as a digital drop box where terrorist leaders can secretly execute command and control.

Tor is vulnerable to a number of attacks aimed at both denying service and degrading anonymity. DDoS attacks targeting an OR's CPU are possible due to Tor's dependence on TLS. Such attacks force an OR to execute so many public key decryptions that it can no longer route messages. An additional attack, available to organizations with enough resources and reach, populates the Tor network with ORs that can be used to monitor communication habits. Although targeting specific users would be difficult given a circuit's ORs are randomly selected, an attack exists wherein a malicious server can "mark" a user's path through Tor [3]. This means if information operators can mislead Tor using cyberterrorists into becoming dependent on "front company" services, they can deny the service at a time critical moment like immediately before an attack and without revealing the service as a front.

Tor is clearly advantageous to organizations opposing U.S. information superiority, but it can also be used by U.S. intelligence organizations. Cyber operations, either proactive or initiated in response to a hostile action or an intelligence need, require specialized tools/services like Tor to establish believable cover stories for cyberspace operations. For example, many organizations are collecting intelligence using open source web pages. Some organizations go even further and engage individuals in chat rooms. These operations often identify reliable sources capable of providing IO operators early warning of network attacks and hacker tool development. Such operations can use Tor to ensure sound operations security,

thus hiding from adversaries U.S. targets, IO techniques, and tools. To protect against DDoS attacks like the one already mentioned we are studying the use of client puzzles as a mitigation technique. Servers employing client puzzles force a requesting client, whether they are honest or malicious, to complete a puzzle before it allocates a resource. As a result, the attacker is forced to find additional resources, i.e. more zombies, in order to degrade service. In the case of Tor, puzzles are used to keep ORs from having to complete the large number of decryption operations forced upon it during an attack. Of additional concern is the impact this defense has on latency and anonymity.

Summary

Tor was developed so individuals could hide their Internet activity from snooping governments and corporations. Given that this tool is free and the source code is available for download, it is likely that adversarial governments have added this technology to their asymmetric arsenal. On the other hand, cyberspace intelligence collection is still very much dependent on human interaction. It may not be possible to locate sources or verify their true identity, but that does not exclude them from being valuable resources to U.S. information operations.

Access to the individuals developing “zero-day” exploits or participating in state-sponsored operations does not occur without first establishing trust and credibility. Oddly enough, establishing trust often requires an individual demonstrate an ability to use concealment tools like Tor. U.S. information operators, particularly those specializing in online collection, should consider Tor for future operations and if an offensive purpose cannot be identified, they should, at a minimum, acknowledge its possible use by American adversaries.

Endnotes:

- [1] Dingledine, Roger, Nick Mathewson, and Paul Syverson. “Tor: The Second-Generation Onion Router”. Proceedings of the 13th USENIX Security Symposium. August 2004.
- [2] Microsoft IIS 4.0 / 5.0 vulnerable to directory traversal via extended unicode in url (MS00-078). Vulnerability Note VU#111677, US-CERT, October 2000.
- [3] Murdoch, Steven J. and George Danezis. “Low-Cost Traffic Analysis of Tor”. Proceedings of the 2005 IEEE Symposium on Security and Privacy. IEEE CS, May 2005.
- [4] “The Freehaven Project”, August 2005. <http://freehaven.net>
- [5] “Tor: An Anonymous Internet Communication System”, November 2004. <http://tor.eff.org>