Australian Government
Australian Signals Directorate

ASD AUSTRALIAN SIGNALS DIRECTORATE
ACSC Australian Cyber Security Centre

Defending against the malicious use of the Tor network

First published: October 2020
Last updated: October 2021

Introduction

The Tor network is a system that facilitates anonymous communication by concealing a user's Internet Protocol (IP) address through encryption and a series of self-described anonymous and private connections. The Tor network receives its name from the original software project it is based upon, 'The onion router', and is maintained by the Tor Project. The Tor network can have legitimate uses, however, in practice traffic from the Tor network is overwhelmingly malicious.

Blocking traffic from the Tor network will prevent malicious actors from using the Tor network to easily conduct anonymous reconnaissance and exploitation of systems and typically has minimal, if any, impact on legitimate users. This publication provides guidance on the prevention and detection of traffic from the Tor network.

Background of the Tor network

The core principles of onion routing were created by the United States Navy in the mid-1990s with the intent of protecting United States intelligence communications. However, the Tor project is now a not-for-profit initiative and there is no cost associated with accessing Tor software and services.

The Tor network operates by using open-source software that utilises the onion routing technique for anonymous communications over computer networks. Onion routing includes the use of multiple layers of encryption around a message. Each layer is decrypted by an onion router which will expose the message's next destination. Each onion router in the path can only know the previous onion router and the next onion router. Using this process, it is impossible for a single onion router to see both the original source and destination addresses.

The Tor network consists of thousands of onion routers (known as Tor nodes) which are used to conceal a user's location from the destination, usually a website or web server. The anonymity that this concealment provides may assist users to stay hidden from trackers and censorship filters, although it should be noted that some tracking technologies, such as website cookies, continue to be effective.

Tor exit nodes, relays and bridges

Tor nodes fall into four categories:

- **Tor exit nodes:** A Tor exit node is the last Tor node that traffic passes through in the Tor network before exiting onto the internet.
- **Tor guard nodes:** A Tor guard node is the point of entry into the Tor network.

- **Tor middle nodes:** A Tor middle node is a Tor node that sits in the middle of the Tor network between a Tor guard node and a Tor exit node. A message can interact with multiple Tor middle nodes before reaching a Tor exit node.
- **Tor bridge nodes:** A Tor bridge node is a specific type of Tor guard node that is not listed on the public directory of Tor nodes.

Involvement of the Tor network in cyberattacks

When conducting malicious activities, malicious actors often take care in planning the infrastructure that will be used to conduct their campaigns. By carefully choosing infrastructure from which to conduct reconnaissance and attacks, malicious actors reduce the risk of detection and attribution. The Tor network provides malicious actors with these capabilities for free and with minimal setup. Further, the activities of other users helps to obfuscate malicious actors' activities.

The Tor network provides malicious actors with a multitude of source locations from which to conduct malicious activities against their targets. By ensuring that different Tor exit nodes are used, malicious actors are able to make it more difficult for defenders to correlate activity (as malicious actors may use thousands of different Tor exit nodes to conduct their operations), block activity (blocking individual IP addresses is ineffective when malicious actors can counter the move easily), avoid takedown and other lawful mechanisms aimed at disrupting malicious activity, and create barriers to attribution.

The Australian Signals Directorate (ASD) recommends organisations block traffic from Tor exit nodes to their internet-exposed services provided this will not meaningfully impact accessibility for significant numbers of legitimate users.

The use of the Tor network by malicious actors can be aligned with MITRE's *ATT&CK for Enterprise* framework. For example:

- **TA0043 - Reconnaissance:** The Tor network can be used by malicious actors to select targets for further activities. This includes using the Tor network to scan large ranges of potential targets to prevent attribution to a specific malicious actor.
- **TA0001 - Initial Access:** The Tor network can be used to gain an initial foothold on a network. This is usually accomplished by using the Tor network to send malicious traffic to exploit internet-exposed services.
- **TA0011 - Command and Control:** The Tor network can be used for outbound communications to command and control servers.
- **TA0010 - Exfiltration:** The Tor network can be used to exfiltrate information out of an environment.
- **TA0040 - Impact:** The Tor network can be used in denial-of-service attacks.

Strategies for risk management

ASD recommends blocking traffic from the Tor network, however, this may not be practical in all circumstances. In some cases, organisations may decide to allow traffic from the Tor network to access specific internet-exposed services. In decreasing order of preference, ASD recommends:

- blocking trafficking from the Tor network
- monitoring traffic from the Tor network

- logging traffic from the Tor network.

Blocking traffic from the Tor network

Blocking is ASD's recommended strategy to avoid reconnaissance and exploitation from malicious actors using the Tor network. In ASD's estimation, this has the lowest ongoing cost (apart from doing nothing) and the highest impact on malicious actors' activities.

Although ASD's advice is to block traffic from the Tor network, a policy decision can be made to allow traffic from the Tor network to access specific internet-exposed services where it is believed that the use of the Tor network is a specific business requirement. It is suggested that in these cases systems allowing traffic from the Tor network be logged and monitored.

Monitoring traffic from the Tor network

Traffic from Tor exit nodes to Australian infrastructure is disproportionately malicious when compared to other internet traffic. If an organisation has the capability, and cannot block traffic from the Tor network for policy or accessibility reasons, then subjecting the traffic to additional scrutiny can be an effective use of security resources.

Logging traffic from the Tor network

As noted above, traffic from Tor exit nodes to Australian infrastructure is disproportionately malicious. While logging in isolation provides no security barrier, it can assist organisations in responding to compromises. In order to enrich logging, noting that traffic from the Tor network has a temporal aspect, it should be done when, or soon after, traffic from the Tor network is detected accessing internet-exposed services.

Identifying any legitimate use

Organisations may want to establish the extent of legitimate use of the Tor network before blocking it. If the legitimate access of internet-exposed services from the Tor network is low, organisations can proceed to a blocking strategy with minimal impact. However, if the legitimate access of internet-exposed services from the Tor network is significant, organisations may need to seek a policy decision. In either case organisations may want to prepare a change strategy and communications plan for managing any affected users.

For legitimate users who may be impacted by a blocking strategy, it may be worth informing them that if their intent is to avoid identification and attention then:

- their use of the Tor network is likely to subject their communications to greater scrutiny, not less
- if they are using authenticated access they have already been identified.

Alternatively, if there is an organisational requirement to allow users to access internet-exposed services via the Tor network, then alternate mitigation strategies can be implemented to limit the potential attack vectors. Some alternate mitigation strategies include:

- **Implement challenge-response tests for users:** One method of blocking automated scans from reaching internet-exposed services is to apply a challenge-response test, such as a Completely Automated Public Turing Test to Tell Computers and Humans Apart (CAPTCHA).
- **Use secure Transport Layer Security (TLS) connections:** Without robust and secure TLS configurations, it is possible that a user's information and data could be accessed by a compromised or malicious Tor exit node.

Ensuring that internet-exposed services have secure TLS will limit the information that malicious actors are able to gain from the organisation's client traffic. Although this does not reduce the risk from malicious traffic from Tor exit nodes, it does protect user information.

Planning to block traffic from the Tor network

ASD suggests blocking anonymity networks, such as the Tor network, under the web content filtering strategy within the *Strategies to mitigate cybersecurity incidents* publication.

The Tor nodes participating in the Tor network at any time are constantly changing. Consequently, the task of blocking traffic from the Tor network comes down to three key requirements:

- a reliable means to identify Tor exit nodes
- Policy Enforcement Points (PEPs) where traffic can be filtered based on IP addresses
- the ability to automate the configuration of PEPs with the information from the identified Tor exit nodes.

Exactly how to meet these requirements will vary highly based on an organisation's existing network and system architectures. The following sections provide possible approaches to these challenges.

Identifying Tor exit nodes

The Tor Project provides a list of current Tor exit nodes as a file, which can be downloaded from their website, and as a Domain Name System (DNS) lookup service. Note, the downloadable file contains a list of all known Tor exit nodes at a point in time, and may provide a suitable source if an organisation's plan is to implement an IP address filtering list in a layer 3 or 4 network device such as a router or firewall.

The DNS lookup service provides an advantage in that it is self-updating, such that if a PEP is configured to query the list it will always get the latest information on whether an IP address is, or is not, a Tor exit node at that time. This approach may be useful for layer 7 implementation in systems such as web servers and mail servers, but is less likely to be useful in high-flow network devices such as firewalls and routers. Organisations should be aware that network-added delay from DNS query time may also impact response times for interactive services such as websites.

To access a DNS service which can identify if a connection is coming from a Tor exit node, use the instructions located on the Tor Project website.

Other threat intelligence sources and services that organisations use may also have further information on identifying Tor exit nodes.

Blocking traffic from the Tor network

As noted above, organisations will need to identify the most appropriate PEP in their environments to block traffic from the Tor network. The following is a non-exhaustive list of potential approaches with a summary of advantages and disadvantages. Each approach is discussed in further detail below:

- Blocking traffic from the Tor network using firewalls and routers.
    - **Advantage:** Depending on network topology, a smaller number of devices may need to be configured and maintained.

- **Disadvantage:** These high-flow devices typically need to work with IP addresses, and will not be able to implement a DNS lookup-based method. Automation will be required to regularly update block lists.

- Blocking traffic from the Tor network using web application firewalls and proxies.
  - **Advantage:** Can potentially use DNS lookup-based methods, removing the burden of automating updates.
  - **Disadvantage:** Will likely need to be implemented across a broader set of internet-exposed services creating different management challenges, such as ensuring filtering is enabled when commissioning new services.

Blocking traffic from the Tor network using firewalls and routers

Blocking traffic from the Tor network is appropriate in most scenarios. Blocking traffic from the Tor network can usually be achieved by implementing the correct control at network boundaries. An example of this is a block list within firewalls and routers.

Firewalls, and some security appliances and internet-exposed services, can be configured to block connections from lists of IP addresses. These lists are typically statically entered and must be configured or scripted to be updated on a regular basis.

Blocking traffic from the Tor network using web application firewalls and proxies

Where an environment does not have access to a configurable gateway or firewall, most web application firewalls and proxies can be configured to block traffic from the Tor network.

This approach is also appropriate for organisations that have decided only specific internet-exposed services should be accessible from the Tor network. This could occur where it is determined there is a legitimate business requirement to allow traffic from the Tor network to access specific internet-exposed services.

Planning to monitor or log traffic from the Tor network

Some organisations may find it necessary to detect if the Tor network is being used to access their environment. This could be to determine if a blocking strategy is appropriate, or to allow for greater awareness of traffic associated with the Tor network by security personnel. Traffic from Tor exit nodes can be detected using the same block lists mentioned previously.

Detecting traffic from the Tor network using firewalls and gateways

Traffic from the Tor network can be detected by configuring a firewall or gateway to audit and log connections from Tor exit nodes. This can be achieved by using an up-to-date list of Tor exit nodes in a block list that has been configured in audit mode instead of enforcement mode. In doing so, it is important to ensure the logs are stored in a centralised logging facility, protected from unauthorised access and maintained for review by a qualified analyst.

Detecting traffic from the Tor network in web application logs

Tor exit nodes can be detected in a web application's log of connections that have been made to the server, if they include the public source IP address of the transaction initiator.

As Tor exit nodes can be temporal, it is possible that traffic from the Tor network identified in a log may have come from a Tor exit node that has been removed prior to the creation of the list being used for comparison. This can be prevented by ensuring logs are interrogated or enriched as close as possible to creation.

For use cases where historical logs need to be analysed, there is a database created by the Tor Project to identify whether a Tor node was running on a given IP address on a given date.

Further information

The *Information security manual* is a cybersecurity framework that organisations can apply to protect their systems and data from cyberthreats. The advice in the *Strategies to mitigate cybersecurity incidents*, along with its Essential Eight, complements this framework.

Contact details

If you have any questions regarding this guidance you can write to us or call us on 1300 CYBER1 (1300 292 371).

Disclaimer

The material in this guide is of a general nature and should not be regarded as legal advice or relied on for assistance in any particular circumstance or emergency situation. In any important matter, you should seek appropriate independent professional advice in relation to your own circumstances.

The Commonwealth accepts no responsibility or liability for any damage, loss or expense incurred as a result of the reliance on information contained in this guide.

For more information, or to report a cybersecurity incident, contact us:

cyber.gov.au | 1300 CYBER1 (1300 292 371)

Australian Government

Australian Signals Directorate