

# Spoiled Onions: Exposing Malicious Tor Exit Relays

Philipp Winter  
*Karlstad University*

Stefan Lindskog  
*Karlstad University*

## Abstract

Several hundred Tor exit relays together push more than 1 GiB/s of network traffic. However, it is easy for exit relays to snoop and tamper with anonymised network traffic and as all relays are run by independent volunteers, not all of them are innocuous. In this paper, we seek to expose malicious exit relays and document their actions. First, we monitored the Tor network after developing a fast and modular exit relay scanner. We implemented several scanning modules for detecting common attacks and used them to probe all exit relays over a period of four months. We discovered numerous malicious exit relays engaging in different attacks. To reduce the attack surface users are exposed to, we further discuss the design and implementation of a browser extension patch which fetches and compares suspicious X.509 certificates over independent Tor circuits. Our work makes it possible to continuously monitor Tor exit relays. We are able to detect and thwart many man-in-the-middle attacks which makes the network safer for its users. All our code is available under a free license.

## 1 Introduction

As of January 2014, nearly 1,000 exit relays [24] distributed all around the globe serve as part of the Tor anonymity network [7]. As illustrated in Figure 1, the purpose of these relays is to establish a bridge between the Tor network and the “open” Internet. A user’s Tor circuits, which are encrypted tunnels, terminate at exit relays and from there, the user’s traffic proceeds to travel over the open Internet to its final destination. Since exit relays can see traffic as it is sent by a Tor user, their role is particularly sensitive compared to entry guards and middle relays; especially because traffic frequently lacks end-to-end encryption.

By design, exit relays act as a “man-in-the-middle” (MitM) in between a user and her destination. This

renders it possible for exit relay operators to run various MitM attacks such as traffic sniffing, DNS poisoning, and SSL-based attacks such as HTTPS MitM and sslstrip [19]. An additional benefit for attackers is that exit relays can be set up quickly and anonymously, making it very difficult to trace attacks back to their origin. While it is possible for relay operators to specify contact information such as an email address<sup>1</sup>, this is optional. As of January 2014, only 56% out of all 4,962 relays publish contact information. Even fewer relays have *valid* contact information.

To thwart a number of popular attacks, TorBrowser [23]—the Tor Project’s modified version of Firefox—ships with extensions such as HTTPS-Everywhere [8] and NoScript [14]. While HTTPS-Everywhere provides rules to rewrite HTTP traffic to HTTPS traffic, NoScript attempts to prevent many script-based attacks. However, there is little users can do if web sites implement poor security such as the lack of site-wide TLS, session cookies being sent in the clear, or using weak cipher suites in their web server configuration. Often, such bad practices enable attackers to spy on users’ traffic or, even worse, hijack accounts. Besides, TorBrowser cannot protect against attacks targeting protocols such as SSH.

All these attacks are not just of theoretical nature. In 2007, a security researcher published 100 POP3 government credentials he captured by sniffing traffic on a set of exit relays under his control [22]; supposedly to show the need for end-to-end encryption when using Tor. In Section 2, we will discuss additional attacks which were found in the wild.

### 1.1 What Happens to Bad Exits?

The Tor Project has a way to prevent clients from selecting bad exit relays as the last hop in their three-hop

<sup>1</sup>Contact information can be useful to get in touch with relay operators, e.g., if they misconfigured their relay.

circuits. After a suspected relay is communicated to the project, the reported attack is first reproduced. If the attack can be verified, a subset of two (out of all nine) directory authority operators manually blacklist the relay using Tor’s `AuthDirBadExit` configuration option. Every hour, the directory authorities vote on the *network consensus* which is a signed list of all relays, the network is comprised of. Among other information, the consensus includes the *BadExit flag*. As long as the majority of the authorities responsible for the BadExit flag, i.e., two out of two, agree on the flag being set for a particular relay, the next network consensus will label the respective relay as BadExit. After the consensus was then signed by a sufficient number of directory authorities, it propagates through the network and is eventually used by all Tor clients after a maximum of three hours. From then on, clients will no longer select relays labelled as BadExit as the last hop in their circuits. Note that this does not mean that BadExit relays become effectively useless. They keep getting selected by clients as their entry guards and middle relays. All the malicious relays we discovered were assigned the BadExit flag.

Note that the BadExit flag is not only given to relays which are proven to be malicious. It is also assigned to relays which are misconfigured or are otherwise unable to fulfil their duty of providing unfiltered Internet access. A frequent cause of misconfiguration is the use of third-party DNS resolvers which block certain web site categories.

Apart from the BadExit flag, directory authorities can blacklist relays by disabling its *Valid* flag which prevents clients from selecting the relay for *any* hop in its circuit. This option can be useful to disable relays running a broken version of Tor or are suspected to engage in end-to-end correlation attacks.

## 1.2 Contributions

The three main contributions of this paper are as follows.

- We discuss the design and implementation of `exitmap`; a flexible and fast exit relay scanner which is able to detect several popular MitM attacks.
- Using `exitmap`, we monitored the Tor network over a period of four months. We analyse the attacks we discovered in the wild during that time period.
- We propose the design and prototype of a browser extension patch which fetches and compares X.509 certificates over diverging Tor circuits. That allows our patch to detect MitM attacks against HTTPS.

The remainder of this paper is structured as follows. Section 2 begins by giving an overview of related work. It is followed by Section 3 which discusses the design




Figure 1: The structure of a three-hop Tor circuit. Exit relays constitute the bridge between encrypted circuits and the open Internet. As a result, exit relay operators can see—and tamper with—the anonymised traffic of users.

and implementation of `exitmap`. Section 4 then presents the attacks we discovered in the wild. Next, Section 5 proposes the design and implementation of a browser extension patch which can protect against HTTPS MitM attacks. Finally, Section 6 concludes this paper.

## 2 Related Work

While MitM attacks have generally received considerable attention in the literature [12, 30], their occurrence in the Tor network remains largely unexplored. This is unfortunate as the Tor network enables the study of real-world MitM attacks which are rare and poorly documented outside the Tor network.

In 2006, Perry began developing the framework “Snakes on a Tor” (SoaT) [25]. SoaT is a Tor network scanner whose purpose—similar to our work—is to detect misbehaving exit relays. Decoy content is first fetched over Tor, then over a direct Internet connection, and finally compared. Over time, SoaT was extended with support for HTTP, HTTPS, SSH and several other protocols. However, SoaT is no longer maintained and makes use of deprecated libraries. Compared to SoaT, our design is more flexible and significantly faster.

Similar to SoaT, Marlinspike implemented `tortunnel` [20]. The tool exposes a local SOCKS interface which accepts connections from arbitrary applications. Incoming data is then sent over exit relays using one-hop circuits. By default, `exitmap` does not use one-hop circuits as that could be detected by attackers which could then act innocuously.

A first attempt to detect malicious exit relays was made in 2008 by McCoy *et al.* [21]. The authors established decoy connections to servers under their control. They further controlled the authoritative DNS server responsible for the decoy hosts’ domain names. As long as an attacker on an exit relay sniffed network traffic with

reverse DNS lookups being enabled, the authors were able to map reverse lookups to exit relays by monitoring the authoritative DNS server’s traffic. Using that side channel, McCoy *et al.* were able to find one exit relay sniffing POP3 traffic at port 110. However, attackers could avoid that side channel by disabling reverse lookups. The popular tool tcpdump implements the command line switch `-n` for that exact purpose.

In 2011, Chakravarty *et al.* [3] attempted to detect exit relays sniffing Tor users’ traffic by systematically transmitting decoy credentials over all active exit relays. Over a period of ten months, the authors uncovered ten relays engaging in traffic snooping. Chakravarty *et al.* could verify that the operators were sniffing exit traffic because they were later found to have logged in using the snooped credentials. While the work of Chakravarty *et al.* represents an important first step towards monitoring the Tor network, their technique only focused on SMTP and IMAP. At the time of writing, only 20 out of all ~1,000 exit relays allow exiting to port 25. HTTP appears to be significantly more popular [13, 21]. Also, similar to McCoy *et al.*, the authors only focused on traffic snooping attacks which are passive. Active attacks remain entirely unexplored until today.

The Tor Project used to maintain a web page documenting misbehaving relays which were assigned the BadExit flag [15]. As of January 2014, this page lists 35 exit relays which were discovered in between April 2010 and July 2013. Note that not all of these relays engaged in attacks; almost half of them ran misconfigured anti virus scanners or used broken exit policies<sup>2</sup>.

Since Chakravarty *et al.*, no systematic study to spot malicious exits was conducted. Only some isolated anecdotal evidence emerged [28]. Our work is the first to give a comprehensive overview of *active attacks*. We further publish our code under a free license<sup>3</sup>. By doing so, we enable and encourage continuous and *crowd-sourced* measurements rather than one-time scans.

### 3 Probing Exit Relays

We now discuss the design and implementation of `exitmap` which is a lightweight Python-based *exit relay scanner*. Its purpose is to create custom circuits to exit relays which are then probed by modules which establish decoy connections to various destinations. We seek to provoke exit relays to tamper with our connections, thus revealing their malicious intent. By doing so, we

<sup>2</sup>An exit relay’s *exit policy* determines to which addresses and ports the relay forwards traffic to. Often, relay operators choose to not forward traffic to well-known file sharing ports in order to avoid copyright infringement.

<sup>3</sup>See: [http://www.cs.kau.se/philwint/spoiled\\_onions](http://www.cs.kau.se/philwint/spoiled_onions).




Figure 2: The design of `exitmap`. Our scanner invokes a Tor process and uses the library Stem to control it. Using Stem, circuits are created “manually” and attached to decoy connections which are initiated by our probing modules.

hope to discover and remove all “spoiled onions” which might be part of the Tor network.

We will also show that our scanner’s *modular design* enables quick prototyping of new scanning modules. Also, its *event-driven architecture* makes it possible to scan the entire Tor network within a matter of only seconds while at the same time sparing its resources.

#### 3.1 The Design of `exitmap`

The schematic design of our scanner is illustrated in Figure 2. Our tool is run on a single machine and requires the Python library Stem [26]. Stem implements the Tor control protocol [27] and we use it to initiate and close circuits, attach streams to circuits as well as to parse the network consensus. Upon starting `exitmap`, it first invokes a local Tor process which proceeds by fetching the newest network consensus in order to know which exit relays are currently online.

Next, our tool is fed with a set of exit relays. This set can consist of a single relay, all exit relays in a given country, or the set of all Tor exit relays. Random permutation is then performed on the set so that repeated scans do not probe exit relays in the same order. This is useful while developing and debugging new scanning modules as it equally distributes the load over all selected exit relays.

Once `exitmap` knows which exit relays it has to probe, it initiates circuits which use the respective exit relays as last hop. All circuits are created asynchronously in the background. Once a circuit to an exit relay is established, Tor informs `exitmap` about the circuit by sending an asynchronous circuit event over the control connection. Upon receiving the notification about a successfully created circuit, `exitmap` invokes the desired probing module which then proceeds by establishing a connection to a decoy destination (see § 3.3). Tor creates

stream events for new connections to the SOCKS port which are also sent to exitmap. At this point, we attach the stream of a probing module to the respective circuit. Note that stream-to-circuit attaching is typically done by Tor. In order to have control over this action, our scanner invokes Tor with the configuration option `__LeaveStreamsUnattached` which instructs Tor to leave streams unattached.

For performance reasons, Tor builds circuits preemptively, i.e., a number of circuits are kept ready even if there is no data to be sent yet. Since we want full control over all circuits, we prevent Tor from creating circuits preemptively by using the configuration option `__DisablePredictedCircuits`.

Probing modules can either be standalone processes or Python modules. Processes are invoked over the `torsocks` wrapper [29] which hijacks system calls such as `socket()`, `connect()`, and `gethostbyname()` in order to redirect them to Tor’s SOCKS port. We used standalone processes for our HTTPS and SSH modules. In addition, probing modules can be implemented in Python. To redirect Python’s networking API over Tor’s SOCKS port, we extended the `Socksipy` module [10]. We used Python for our `sslstrip` and `DNS` modules.

### 3.2 Performance Hacks

A naive approach to probing exit relays could cause non-trivial costs for the Tor network; mostly computationally but also in terms of network throughput. We implemented a number of tweaks in order for our scanning to be as fast and cheap as possible.

First, we expose a configuration option for avoiding the default of three-hop circuits. Instead, we only use *two hops* as illustrated in Figure 3. Tor’s motivation for three hops is anonymity but since our scanner has no need for strong anonymity, we only select a static entry relay—ideally operated by `exitmap`’s user—which then directly forwards all traffic to the respective exit relays. We offer no option to use one-hop circuits as that would make it possible for exit relays to isolate scanning connections: A malicious exit relay could decide not to tamper with a circuit if it originates from a non-Tor machine. Since we use a static first hop which is operated by us, we concentrate most of the scanning load on a single machine which is well-suited to deal with the load. Other entry and middle relays do not have to “suffer” from scans. However, note that over time malicious exit relays are able to correlate scans with relays, thus determining which relays are used for scans. To avoid this problem, `exitmap`’s first hop could be changed periodically and we hope that by crowd-sourcing our scanner, isolating middle relays is no longer a viable option for attackers.

Another computational performance tweak can be




Figure 3: Instead of establishing a full three-hop circuit, our scanner is able to use a static middle relay; preferably operated by whoever is running our scanner. By doing so, we concentrate the load on one machine while making our scanning activity slightly more obvious.

achieved on Tor’s authentication layer. At the moment, there are two ways how a circuit handshake can be conducted; either by using the *traditional TAP* or the *newer NTor* handshake. TAP—short for Tor Authentication Protocol [9]—is based on Diffie-Hellman key agreement in a multiplicative group. NTor, on the other hand, uses the more efficient elliptic curve group `Curve25519` [2]. A non-trivial fraction of a relay’s computational load can be traced back to computationally expensive circuit handshakes. By preferring NTor over TAP, we slightly reduce the computational load on exit relays. Since NTor supersedes TAP and is becoming more and more popular as Tor clients upgrade, we believe that it is not viable for attackers to “whitelist” NTor connections.

### 3.3 Scanning Modules

After discussing the architecture of `exitmap`, we now present several probing modules we developed in order to detect specific attacks. When designing a module, it is important to consider its *indistinguishability* from genuine Tor clients. As mentioned above, malicious relay operators could closely inspect exit traffic (e.g., by examining the user agent string of browsers) and only attack connections which appear to be genuine Tor users.

#### 3.3.1 HTTPS

McCoy *et al.* [21] showed that HTTP is the most popular protocol in the Tor network, clearly dominating other protocols such as instant messaging or e-mail<sup>4</sup>. While HTTPS lags behind, it is still widely used and unsurprisingly, several exit relays were documented to have tampered with HTTPS connections [15] in the past.

<sup>4</sup>This is particularly true based on *connections* but not so much based on *bytes transferred*.



We implemented an HTTPS module which fetches a decoy destination's X.509 certificate and extracts its fingerprint. This fingerprint is then compared to the expected fingerprint which is hard-coded in the module. If there is a mismatch, an alert is triggered. Originally, we began by fetching the certificate using the command line utility `gnutls-cli`. We later extended the module to send a TLS client hello packet as it is sent by TorBrowser to make the scan less distinguishable from what a real Tor user would send.

Note that an attacker might become suspicious after observing that a Tor user only fetched an X.509 certificate without actually browsing the web site. However, at the point in time an attacker would become suspicious, we already have what we need; namely the X.509 certificate. Also, our module could be extended to simulate simple browsing activity.

### 3.3.2 sslstrip

Instead of *interfering* with TLS connections, an attacker can seek to *prevent* TLS connections. This is the purpose of the tool `sslstrip` [19]. The tool achieves this goal by transparently rewriting HTML documents sent from the server to the client. In particular, it rewrites HTTPS links to HTTP links. A secure login form such as `https://login.example.com` is subsequently rewritten to HTTP which can cause a user's browser to submit her credentials in the clear. While the HTTP Strict Transport Security policy [11] prevents `sslstrip`, it is still an effective attack against many large-scale web sites with Yahoo! being one of them as of January 2014. From an attacker's point of view, the benefit of `sslstrip` is that it is a comparatively silent attack. Browsers will not show certificate warnings but vigilant users might notice the absence of browser-specific TLS indicators such as a green address bar.

We implemented a probing module which can detect `sslstrip` attacks. Our module fetches web sites containing HTTPS links over unencrypted HTTP. Afterwards, the module simply verifies whether the fetched HTML document contains the expected HTTPS links or if they were "downgraded" to HTTP. After experiments in a lab setting showed our module to work, we began `sslstrip` scans on October 24, 2013.

### 3.3.3 SSH

The Tor network is also used to transport SSH traffic. This can easily be done with the help of tools such as `torsocks` [29]. Analogous to HTTPS-based attacks, malicious exit relays could run MitM attacks against SSH. In practice, this is not as easy as targeting HTTPS given SSH's "trust on first use" model. As long as the very first

```
1 function probe( fingerprint, command ) {
2
3     ssh_public_key = "11:22:33:44:55:66:77:88" +
4                     "99:00:aa:bb:cc:dd:ee:ff";
5
6     output = command.execute("ssh -v 1.2.3.4");
7
8     if (ssh_public_key not in output) {
9         print("Possible MitM attack by " + fingerprint);
10    }
11 }
```

Figure 4: Pseudo code illustrating a scanning module which tests SSH. It establishes an SSH connection to a given host and verifies if the fingerprint is as expected. If the observed fingerprint differs, an alert is raised.

connection to an SSH server with a given key was secure, the public key is then stored by the client and kept as reference for subsequent connections. That way, SSH is able to print a warning whenever the server's public key is unexpected. As a result, a MitM attack has to target a client's very first SSH connection where the server's public key is not yet known.

Nevertheless, this practical problem might not stop attackers from attempting to interfere with SSH connections. Our SSH module, conceptually similar to the pseudo code shown in Figure 4, makes use of OpenSSH's `ssh` and `torsocks` to connect to a decoy server. Again, the server's key fingerprint is extracted and compared to the hard-coded fingerprint. However, compared to the HTTPS module, it is difficult to achieve indistinguishability over time. After all, a malicious relay operator could monitor an entire SSH session. If it looks suspicious, e.g., it only fetches the public key, or it lasts only one second, the attacker could decide to whitelist the destination in the future. Alternatively, we could establish SSH connections to random hosts on the Internet. This, however, is often considered undesired scanning activity and does not constitute good Internet citizenship. Instead, we again seek to solve this problem by publishing our source code and encouraging people to crowdsource exitmap scanning. Every exitmap user is encouraged to use her own SSH server as decoy destination. That way, we can achieve destination diversity without bothering arbitrary SSH servers on the Internet.

### 3.3.4 DNS

While the Tor protocol only transports TCP streams, clients can ask exit relays to do DNS resolution by wrapping domain names in a `RELAY_BEGIN` cell [6]. This cell is then sent to the exit relay, once a circuit was established. In the past, some exit relays were found to inadvertently censor DNS queries, e.g., by using an OpenDNS config-

uration which blocks certain domain categories such as “Pornography” or “Proxy/Anonymiser” [15]. Recall that while such behaviour is not intentionally malicious, it is certainly enough to get the BadExit flag assigned.

Our probing module maintains a whitelist of domains together with their corresponding IP addresses and raises an alert if the DNS A record of a domain name is unexpected. This approach works well for sites with a known set of IP addresses but large sites frequently employ a diverse—and sometimes geographically load-balanced—set of IP addresses which is difficult to enumerate. Our module probes several domains in the categories finance, social networking, political activism, and pornography.

### 3.4 Ethical Considerations

Due to exitmap’s modular architecture, it can be used for various unintended—and even unethical—purposes. For example, modules for web site scraping or online voting manipulation come to mind. All sites which naively bind identities to IP addresses might be an attractive target. While we do not endorse such actions, we point out that these activities are hard to stop and will continue to happen and already happen regardless; with or without scanner. If somebody decides to abuse our scanner for such actions, it will at least spare the Tor network’s resources more than a naive design. As a result, we believe that by publishing our code, the benefit to the public outweighs the damage caused by unethical use.

## 4 Experimental Results

On September 19th, we ran our first full scan over all ~950 exit relays which were part of the Tor network at the time. From then on, we scanned all exit relays several times a week. Originally, we began our scans while only armed with our HTTPS module but as time passed, we added additional modules which allowed us to scan for additional attacks. In this section, we will discuss the results we obtained by monitoring the Tor network over a period of several months.

### 4.1 Scanning Performance

The performance of our probing modules is illustrated in Figure 5. The ECDF’s  $x$ -axis shows the time it takes for a module to finish successfully. The  $y$ -axis shows the cumulative fraction of all exit relays. The diagram shows that all modules are able to scan at least 98% of all Tor exit relays under 50 seconds.

Our data further shows that for all modules, 84%–88% of circuit creations succeeded. The remaining circuits




Figure 5: The performance of our probing modules. The DNS module is slower because it resolves several domain names at once. All other modules can scan at least 98% of all Tor exit relays under 40 seconds.

either timed out or were torn down by the respective exit relay using a DESTROY cell.

### 4.2 Malicious Relays

Table 1 contains the 25 malicious and misconfigured exit relays we found. We discovered the first two relays “manually” before we had developed exitmap. All the data illustrated in the table was gathered on the day we found the respective attack. The columns are, from left to right:

**Fingerprint** The first 4 bytes of the relay’s unique 20-byte SHA-1 fingerprint.

**IP addresses** All IPv4 addresses or netblocks, the relay was found to have used over its life time.

**Country** The country in which the relay resided. The country was determined with the help of Max-Mind’s GeoIP lite database.

**Bandwidth** The advertised bandwidth, the relay was willing to contribute to the network.

**Attack** The attack, the relay was running or its configuration problem.

**Sampling rate** The sampling rate of the attack, i.e., how many connections were affected.

**First active** The day, the relay was set up.

**Discovery** The day, we discovered the relay.

Apart from all the conspicuous HTTPS MitM attacks which we will discuss later, we exposed two relays running sslstrip for a short time. The relay 5A2A51D4 injected custom HTML code into HTTP traffic (see Appendix B). While the injected code seemed harmless

Table 1: All 25 malicious and misconfigured exit relays we discovered over a period of 4 months. The data was collected right after a relay was discovered. We have reason to believe that all relays whose fingerprint ends with a † were run by the same attacker.

Fingerprint	IP addresses	Country	Bandwidth	Attack	Sampling rate	First active	Discovery
F8FD29D0†	176.99.12.246	Russia	7.16 MB/s	HTTPS MitM	<i>unknown</i>	2013-06-24	2013-07-13
8F9121BF†	64.22.111.168/29	U.S.	7.16 MB/s	HTTPS MitM	<i>unknown</i>	2013-06-11	2013-07-13
93213A1F†	176.99.9.114	Russia	290 KB/s	HTTPS MitM	50%	2013-07-23	2013-09-19
05AD06E2†	92.63.102.68	Russia	5.55 MB/s	HTTPS MitM	33%	2013-08-01	2013-09-19
45C55E46†	46.254.19.140	Russia	1.54 MB/s	SSH & HTTPS MitM	12%	2013-08-09	2013-09-23
CA1BA219†	176.99.9.111	Russia	334 KB/s	HTTPS MitM	37.5%	2013-09-26	2013-10-01
1D70CDED†	46.38.50.54	Russia	929 KB/s	HTTPS MitM	50%	2013-09-27	2013-10-14
EE215500†	31.41.45.235	Russia	2.96 MB/s	HTTPS MitM	50%	2013-09-26	2013-10-15
12459837†	195.2.252.117	Russia	3.45 MB/s	HTTPS MitM	26.9%	2013-09-26	2013-10-16
B5906553†	83.172.8.4	Russia	850.9 KB/s	HTTPS MitM	68%	2013-08-12	2013-10-16
EFF1D805†	188.120.228.103	Russia	287.6 KB/s	HTTPS MitM	61.2%	2013-10-23	2013-10-23
229C3722	121.54.175.51	Hong Kong	106.4 KB/s	sslstrip	<i>unsampled</i>	2013-06-05	2013-10-31
4E8401D7†	176.99.11.182	Russia	1.54 MB/s	HTTPS MitM	79.6%	2013-11-08	2013-11-09
27FB6BB0†	195.2.253.159	Russia	721 KB/s	HTTPS MitM	43.8%	2013-11-08	2013-11-09
0ABB31BD†	195.88.208.137	Russia	2.3 MB/s	SSH & HTTPS MitM	85.7%	2013-10-31	2013-11-21
CADA00B9†	5.63.154.230	Russia	187.62 KB/s	HTTPS MitM	<i>unsampled</i>	2013-11-26	2013-11-26
C1C0EDAD†	93.170.130.194	Russia	838.54 KB/s	HTTPS MitM	<i>unsampled</i>	2013-11-26	2013-11-27
5A2A51D4	111.240.0.0/12	Taiwan	192.54 KB/s	HTML Injection	<i>unsampled</i>	2013-11-23	2013-11-27
EBF7172E†	37.143.11.220	Russia	4.34 MB/s	SSH MitM	<i>unsampled</i>	2013-11-15	2013-11-27
68E682DF†	46.17.46.108	Russia	60.21 KB/s	SSH & HTTPS MitM	<i>unsampled</i>	2013-12-02	2013-12-02
533FDE2F†	62.109.22.20	Russia	896.42 KB/s	SSH & HTTPS MitM	42.1%	2013-12-06	2013-12-08
E455A115	89.128.56.73	Spain	54.27 KB/s	sslstrip	<i>unsampled</i>	2013-12-17	2013-12-18
02013F48	117.18.118.136	Hong Kong	538.45 KB/s	DNS censorship	<i>unsampled</i>	2013-12-22	2014-01-01
2F5B07B2	178.211.39	Turkey	204.8 KB/s	DNS censorship	<i>unsampled</i>	2013-12-28	2014-01-06
4E2692FE	24.84.118.132	Canada	52.22 KB/s	OpenDNS	<i>unsampled</i>	2013-12-21	2014-01-06

```
1 C=US
2 ST=Nevada
3 L=Newbury
4 O=Main Authority
5 OU=Certificate Management
6 CN=main.authority.com
7 EMAIL=cert@authority.com
```

Figure 6: X.509 information which is part of the malicious certificates used for the MitM attacks. The full certificate is shown in Appendix A.

during our tests, we cannot rule out malicious intent. Two more relays—02013F48 and 2F5B07B2—were subject to their country’s DNS censorship. The Turkish relay blocked many pornography web sites and redirected the user to a government-run web server which explained the reason for the redirection. The second relay seemed to have fallen prey to the Great Firewall of China’s DNS poisoning; perhaps, the relay made use of a DNS resolver in China. Several domains such as torproject.org, facebook.com and youtube.com returned invalid IP addresses which were also found in previous work [18]. Finally, 4E2692FE was misconfigured because it used an OpenDNS policy which would censor web sites in the category “pornography”.

All the remaining relays engaged in HTTPS and/or SSH MitM attacks. Upon establishing a connection to the decoy destination, these relays exchanged the destination’s certificate with their own, self-signed version. Since these certificates were not issued by a trusted authority contained in TorBrowser’s certificate store, a user falling prey to such a MitM attack would be redirected to the `about:certerror` warning page.

Interestingly, we have reason to believe that all relays whose fingerprint ends with a  $\dagger$  were run by the same person or group of people. This becomes evident when analysing the self-signed certificates which were injected for the MitM attacks. In every case, the certificate chain consisted of only two nodes which both belonged to a “Main Authority” and the root certificate—partially shown in Figure 6—of all chains was *identical*. This means that these attacks can be traced back to a common origin even though it is not clear where or what this origin is as we will discuss later.

Apart from the identical root certificate, these relays had other properties in common. First, with the exception of 8F9121BF which was located in the U.S., they were *all located in Russia*. Upon investigating their IP addresses, we discovered that most of the Russian relays were run in the network of a virtual private system (VPS) provider. Several IP addresses were also located in the same netblock, namely 176.99.12.246, 176.99.9.114, 176.99.9.111, and 176.99.11.182. All

these IP addresses are part of the netblock GlobaTel-net which spans 176.99.0.0/20. Furthermore, the malicious exit relays all used Tor version 0.2.2.37<sup>5</sup>. Given its age, this is a rather uncommon version number amongst relays. In fact, we found only two benign exit relays—in Switzerland and the U.S.—which are running the same version. We suspect that the attackers might have a pre-compiled version of Tor which they simply copy to newly purchased systems to spawn new exit relays. Unfortunately, we have no data which would allow us to verify when this series of attacks began. However, the full root certificate shown in Appendix A indicates that it was created on February 12, 2013.

### 4.3 Connection Sampling

Whenever our hunt for malicious relays yielded another result, we strived to confirm the attack by rerunning the scan on the newly discovered relay. However, in the case of the Russian relays, this did not always result in the expected HTTPS MitM attack. Instead, we found that only every  $n$ th connection seemed to have been attacked. We estimated the exact *sampling rate* by establishing 50 HTTPS connections over every relay. We used randomly determined sleep periods in between the scans in order to disguise our activity. The estimated sampling rate for every relay is shown in Table 1 in the column “Sampling rate”. For all Russian relays, it varies between 12% and 68%. We do not have an explanation for the attacker’s motivation to sample connections. One theory is that sampling makes it less likely for a malicious exit relay to be discovered; but at the cost of collecting fewer MitM victims.

Interestingly, the sampling technique was implemented *ineffectively*. This is due to the way how Firefox (and as a result TorBrowser) reacts to self-signed certificates. When facing a self-signed X.509 certificate, Firefox displays its `about:certerror` page which warns the user about the security risk. If a user then decides to proceed, the certificate is *fetched again*. We observed that the malicious exit relays treat the certificate re-fetching as a separate connection whose success depends on the relay’s sampling rate. As a result, a sampling rate of  $n$  means that a MitM attack will only be successfully with a probability of  $n^2$ .

### 4.4 Who is the Attacker?

An important question is where on the path from the exit relay to the destination the attacker is located. At first glance, one might blame the exit relay operator. However, it is also possible that the actual attack happens *after*

<sup>5</sup>For comparison, as of January 2014, the current stable version is 0.2.4.20. Version 0.2.2.37 was declared stable on June 6th, 2012.



the exit relay, e.g., by the relay’s ISP, the network backbone, or the destination ISP. In fact, such an incident was documented in 2006 for a relay located in China [5].

With respect to our data, we cannot entirely rule out that the HTTPS MitM attacks were actually run by an upstream provider of the Russian exit relays. However, we consider it unlikely for the following reasons: 1) the relays were located in diverse IP address blocks and there were numerous other relays in Russia which did not exhibit this behaviour, 2) one of the relays was even located in the U.S., 3) there are no other reported cases on the Internet involving a certification authority called “Main Authority”, and 4) the relays frequently disappeared after they were assigned the BadExit flag.

The identity of the attacker is difficult to ascertain. The relays did not publish any contact information, nicknames, or revealed other hints which could enable educated guesses regarding the attacker’s origin.

## 4.5 Destination Targeting

While Tor’s nature as an anonymity tool renders targeting individuals difficult<sup>6</sup>, an attacker can target classes of users based on their communication *destination*. For example, an attacker could decide to only tamper with connections going to the fictional `www.insecure-bank.com`. Interestingly, we found evidence for exactly that behaviour; at some point the Russian relays began to target at least `facebook.com`. We tested the https version of the Alexa top 10 web sites [1] but were unable to trigger MitM attacks despite numerous connection attempts. Popular Russian web sites such as the mail provider `mail.ru` and the social network `vk.com` also remained unaffected. Note that it is certainly possible that the relays targeted additional web sites we did not test for. Answering this question comprehensively would mean probing for thousands of different web sites.

We have no explanation for the targeting of destinations. It might be another attempt to delay the discovery by vigilant users. However, according to previous research [13], social networking appears to be as popular over Tor as it is on the clear Internet. As a result, limiting the attack to `facebook.com` might not significantly delay discovery.

## 5 Thwarting HTTPS MitM Attacks

The discovery of destination targeting made us reconsider defence mechanisms. Unfortunately, we cannot rule out that there are additional, yet undiscovered exit relays which target low-profile web sites. If we wanted

<sup>6</sup>We assume, of course, that users do not somehow reveal their real identity when using Tor, e.g., by posting on Internet forums under their real name.

to achieve high coverage, we would have to connect to millions of web sites; and given the connection sampling discussed in Section 4.3, this even has to be done repeatedly! After all, an attacker is able to *arbitrarily reduce the scope* of the attack but we are *unable to arbitrarily scale* our scanner. This observation motivated another defence mechanism which is discussed in this section.

### 5.1 Threat Model

We consider an adversary who is controlling the upstream Internet connection of a small fraction of exit relays<sup>7</sup>. The adversary’s goal is to run HTTPS-based MitM attacks against Tor users. We further expect the adversary to make an effort to stay under the radar in order to delay discovery. The actual MitM attack is conducted by injecting self-signed certificates in the hope that users are not scared off by the certificate warning page.

Our threat model does not cover adversaries who control certificate authorities which would enable them to issue valid certificates to avoid TorBrowser’s warning page. This includes several countries as well as organisations which are part of TorBrowser’s root certificate store. Furthermore, we cannot defend against adversaries who control a significant fraction of Tor exit bandwidth.

### 5.2 Multi Circuit Certificate Verification

As long as an attacker is unable to tamper with all connections to a given destination<sup>8</sup>, MitM attacks can be detected by fetching a public key over *differing paths in the network*. This approach was picked up by several projects including Perspectives [30], Convergence [16] and Crossbear [4]. In this section, we discuss a patch for TorBrowser which achieves the same goal but is adapted to the Tor network.

Apart from NoScript and HTTPS-Everywhere, TorBrowser contains another important extension: *Torbutton*. This extension provides the actual interface between TorBrowser and the local Tor process. It directs TorBrowser’s traffic to Tor’s SOCKS port and exposes a number of features such as the possibility to create a new identity.

Torbutton already contains rudimentary code to talk to Tor over the local control port. The control port—typically bound to `127.0.0.1:9151`—provides local applications with an interface to control Tor. For example, Torbutton’s “New Identity” feature works by sending the NEWNYM signal which instructs Tor to switch to clean circuits so that new application requests do not share circuits with old requests. Torbutton already implements

<sup>7</sup>By “fraction”, we mean a relay’s bandwidth as it determines how likely a client is to select the relay as part of its circuit.

<sup>8</sup>This would be the case if an attacker controls the destination.

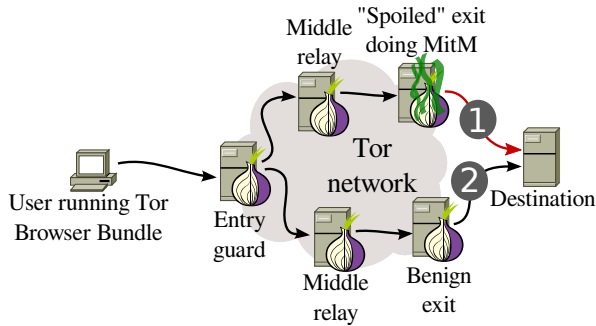


Figure 7: A user stumbles across a self-signed certificate ❶ which could be an indication for a HTTPS MitM attack ran by a malicious exit relay. To verify if the certificate is genuine, the client re-fetches it over an independent exit relay ❷ and checks if the certificate matches or not.

a useful code base for us which made us decide to implement our extension as a set of patches for Torbutton rather than build an independent extension.

### 5.3 Extension Design

Our patch set hooks into the browser event `DOMContentLoaded` which is triggered whenever a document (but not necessarily stylesheets and images) is loaded and parsed by the browser. We then check if the URI of the page contains `“about:certerror”` because whenever TorBrowser encounters a self-signed certificate, it displays this page. However, it is not clear whether the certificate is genuinely self-signed or part of an attack.

In order to be able to distinguish between these two cases, our patch now attempts to re-fetch the certificate over at least one additional and distinct Tor circuit as illustrated in Figure 7. We create a fresh circuit by sending `SIGNAL NEWNYM` to Tor’s control port. Afterwards, we re-fetch the certificate by issuing an `XMLHttpRequest`. If the SHA-1 fingerprints of both certificates match, the certificate is probably<sup>9</sup> genuine. Otherwise, the user might have fallen prey to a MitM attack. False positives are possible, though: large sites could have different certificates for different geographical regions. Note that we are not very likely to witness many such false positives as our code only gets active upon observing self-signed certificates or certificates which somehow trigger the `about:certerror` warning page.

Our extension also informs the user about a potential MitM attack: In case of differing certificates, we open a browser dialogue which informs the user about the situ-

<sup>9</sup>Note that powerful adversaries might be able to control multiple exit relays, network backbones, or even the destination.

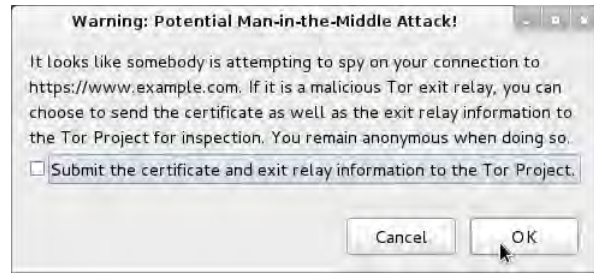


Figure 8: The popup window in TorBrowser which informs the user about the potential HTTPS MitM attack. The user can agree to submitting the gathered information to the Tor Project for further inspection.

ation. A screenshot of our design prototype is shown in Figure 8. We point out that this is likely an attack and we ask the user for permission to send the data to the Tor Project for further inspection. The submitted data contains the *exit relays* used for certificate fetching as well as the *observed certificates*. We transmit no other data which could be used to identify users; as a result, certificate submission is anonymous. While it would be technically possible to transmit the data silently, we believe that users would not appreciate it and consider it as “phoning home”. As a result, we seek to obtain informed consent.

### 5.4 Limitations

In our threat model, we mentioned that our design does not protect against adversaries with the ability to issue valid certificates. While our extension could easily be extended to conduct certificate comparison for all observed certificates, it would flood the Tor network with certificate re-fetches. To make matters worse, the overwhelming majority of these re-fetches would not even expose attacks. There exist other techniques to foil CA-capable adversaries such as certificate pinning [17].

By default, our patch re-fetches a self-signed X.509 certificate only once. An attacker who is controlling a significant fraction of exit relays might be able to conduct a MitM attack for the first as well as for the second fetch. Nevertheless, we would eventually catch the adversary; it would simply be a matter of time until a user selects two independent exit relays.

## 6 Conclusions

In this paper, we revisited the trustworthiness of Tor exit relays. After developing a scanner, we closely monitored all  $\sim 1,000$  exit relays over a period of four months. We discovered 25 relays which were either outright malicious or simply misconfigured. Interestingly, the ma-

majority of the attacks were coordinated instead of being isolated actions of independent individuals. Our results further suggest that the attackers make an active effort to remain under the radar and delay detection.

To make the Tor network safer, we first developed exitmap; an easily extensible scanner which is able to probe exit relays for a variety of MitM attacks. Furthermore, we developed a set of patches for the Tor Browser Bundle which is capable of fetching self-signed X.509 certificates over different network paths to evaluate their trustworthiness. We believe that by being armed with these two tools, the security of the Tor network can be greatly increased. Finally, all our source code is freely available: [http://www.cs.kau.se/philwint/spoiled\\_onions](http://www.cs.kau.se/philwint/spoiled_onions).

## Acknowledgements

We want to thank Internetfonden whose research grant made this work possible. Furthermore, we want to thank Aaron Gibson, Georg Koppen, Harald Lampesberger, and Linus Nordberg for helpful feedback and suggestions.

## References

- [1] Alexa. *The top 500 sites on the web*. 2013. URL: <http://www.alexa.com/topsites>.
- [2] Daniel J. Bernstein. “Curve25519: new Diffie-Hellman speed records”. In: *Public Key Cryptography*. Springer, 2006. URL: <http://cr.yp.to/ecdh/curve25519-20060209.pdf>.
- [3] Sambuddho Chakravarty et al. “Detecting Traffic Snooping in Tor Using Decoys”. In: *RAID*. Springer, 2011. URL: <http://www.cs.columbia.edu/~mikepo/papers/tordecoys.raid11.pdf>.
- [4] *Crossbear*. URL: <http://www.crossbear.org>.
- [5] Roger Dingledine. *Re: Holy shit I caught 1*. 2006. URL: <http://archives.seul.org/or/talk/Aug-2006/msg00262.html>.
- [6] Roger Dingledine and Nick Mathewson. *Tor Protocol Specification*. URL: [https://gitweb.torproject.org/torspec.git?a=blob\\_plain;hb=HEAD;f=tor-spec.txt](https://gitweb.torproject.org/torspec.git?a=blob_plain;hb=HEAD;f=tor-spec.txt).
- [7] Roger Dingledine, Nick Mathewson, and Paul Syverson. “Tor: The Second-Generation Onion Router”. In: *USENIX Security*. USENIX Association, 2004. URL: [http://static.usenix.org/event/sec04/tech/full\\_papers/dingledine/dingledine.pdf](http://static.usenix.org/event/sec04/tech/full_papers/dingledine/dingledine.pdf).
- [8] Electronic Frontier Foundation. *HTTPS Everywhere*. 2013. URL: <https://www.eff.org/https-everywhere>.
- [9] Ian Goldberg. “On the Security of the Tor Authentication Protocol”. In: *PETS*. Springer, 2006. URL: <http://freehaven.net/anonbib/cache/tap:pet2006.pdf>.
- [10] Dan Haim. *Socksipy - A Python SOCKS client module*. 2006. URL: <http://socksipy.sourceforge.net>.
- [11] Jeff Hodges, Collin Jackson, and Adam Barth. *RFC 6797: HTTP Strict Transport Security (HSTS)*. 2012. URL: <https://tools.ietf.org/html/rfc6797>.
- [12] Ralph Holz et al. “X.509 Forensics: Detecting and Localising the SSL/TLS Men-in-the-Middle”. In: *ESORICS*. Springer, 2012. URL: [http://www.net.in.tum.de/fileadmin/bibtex/publications/papers/holz\\_x509forensics\\_esorics2012.pdf](http://www.net.in.tum.de/fileadmin/bibtex/publications/papers/holz_x509forensics_esorics2012.pdf).
- [13] Markus Huber, Martin Mulazzani, and Edgar Weippl. “Tor HTTP Usage and Information Leakage”. In: *Communications and Multimedia Security*. Springer, 2010. URL: <http://freehaven.net/anonbib/cache/huber2010tor.pdf>.
- [14] InformAction. *NoScript*. 2013. URL: <http://noscript.net>.
- [15] *Known Bad Relays*. URL: <https://trac.torproject.org/projects/tor/wiki/doc/badRelays>.
- [16] Thoughtcrime Labs. *Convergence*. 2011. URL: <http://convergence.io>.
- [17] Adam Langley. *Public key pinning*. 2011. URL: <https://www.imperialviolet.org/2011/05/04/pinning.html>.
- [18] Graham Lowe, Patrick Winters, and Michael L. Marcus. *The Great DNS Wall of China*. Tech. rep. New York University, 2007. URL: <http://cs.nyu.edu/~pcw216/work/nds/final.pdf>.
- [19] Moxie Marlinspike. *sslstrip*. URL: <http://www.thoughtcrime.org/software/sslstrip/>.
- [20] Moxie Marlinspike. *tortunnel*. URL: <http://www.thoughtcrime.org/software/tortunnel/>.
- [21] Damon McCoy et al. “Shining Light in Dark Places: Understanding the Tor Network”. In: *PETS*. Springer, 2008. URL: [http://homes.cs.washington.edu/~yoshi/papers/Tor/PETS2008\\_37.pdf](http://homes.cs.washington.edu/~yoshi/papers/Tor/PETS2008_37.pdf).
- [22] Ryan Paul. *Security expert used Tor to collect government e-mail passwords*. 2007. URL: <http://arstechnica.com/security/2007/09/security-expert-used-tor-to-collect-government-e-mail-passwords/>.

- [23] Mike Perry, Erinn Clark, and Steven Murdoch. *The Design and Implementation of the Tor Browser [DRAFT]*. 2013. URL: <https://www.torproject.org/projects/torbrowser/design/>.
- [24] The Tor Project. *Relays with Exit, Fast, Guard, Stable, and HSDir flags*. 2013. URL: <https://metrics.torproject.org/network.html#relayflags>.
- [25] The Tor Project. *Snakes on a Tor*. URL: <https://gitweb.torproject.org/torflow.git/tree/HEAD:/NetworkScanners/ExitAuthority>.
- [26] The Tor Project. *Stem Docs*. 2013. URL: <https://stem.torproject.org>.
- [27] The Tor Project. *TC: A Tor control protocol (Version 1)*. URL: <https://gitweb.torproject.org/torspec.git/blob/HEAD:/control-spec.txt>.
- [28] *TOR exit-node doing MITM attacks*. URL: <http://www.teamfurry.com/wordpress/2007/11/20/tor-exit-node-doing-mitm-attacks>.
- [29] *Torsocks: use socks-friendly applications with Tor*. URL: <https://code.google.com/p/torsocks/>.
- [30] Dan Wendlandt, David G. Andersen, and Adrian Perrig. "Perspectives: Improving SSH-style Host Authentication with Multi-Path Probing". In: *USENIX Annual Technical Conference*. USENIX Association, 2008. URL: [http://perspectivesecurity.files.wordpress.com/2011/07/perspectives\\_usenix08.pdf](http://perspectivesecurity.files.wordpress.com/2011/07/perspectives_usenix08.pdf).

## A Malicious X.509 Root Certificate

Below, the root certificate which was shared by all Russian and the single U.S. exit relay is shown. While the domain [authority.com](http://authority.com) does exist, it seems unrelated to the CA "Main Authority", the issuer.

```

1 Certificate:
2   Data:
3     Version: 3 (0x2)
4     Serial Number: 16517615612733694071 (0xe53a5be2bd702077)
5     Signature Algorithm: sha1WithRSAEncryption
6     Issuer: C=US, ST=Nevada, L=Newbury, O=Main Authority,
7           OU=Certificate Management,
8           CN=main.authority.com/emailAddress=cert@authority.com
9     Validity
10    Not Before: Feb 12 08:13:07 2013 GMT
11    Not After : Feb 10 08:13:07 2023 GMT
12    Subject: C=US, ST=Nevada, L=Newbury, O=Main Authority,
13           OU=Certificate Management,
14           CN=main.authority.com/emailAddress=cert@authority.com
15    Subject Public Key Info:
16    Public Key Algorithm: rsaEncryption
17    Public-Key: (1024 bit)
18    Modulus:
19      00:da:5d:5f:06:06:dc:8e:f1:8c:70:b1:58:12:0a:
20      41:0e:b9:23:cc:0e:6f:bc:22:5a:05:12:09:cf:ac:
21      85:9d:95:2c:3a:93:5d:c9:04:c9:4e:72:15:6a:10:
22      f1:b6:cd:e4:8e:ad:5a:7f:1e:d2:b5:a7:13:e9:87:
23      d8:aa:a0:24:15:24:84:37:d1:69:8e:31:8f:5c:2e:
24      92:e3:f4:9c:c3:bc:18:7d:cf:b7:ba:b2:5b:32:61:

```

```

25      64:05:cd:1f:c3:b5:28:e1:f5:a5:1c:35:db:0f:e8:
26      c3:1d:e3:e3:33:9c:95:61:6d:b7:a6:ad:de:2b:0d:
27      d2:88:07:5f:63:0d:9c:1e:cf
28      Exponent: 65537 (0x10001)
29    X509v3 extensions:
30      X509v3 Subject Key Identifier:
31      07:42:E0:52:A7:DC:A5:C5:0F:C5:
32      AF:03:56:CD:EB:42:8D:96:00:D6
33      X509v3 Authority Key Identifier:
34      keyid:07:42:E0:52:A7:DC:A5:C5:0F:C5:
35      AF:03:56:CD:EB:42:8D:96:00:D6
36      DirName:/C=US/ST=Nevada/L=Newbury/O=Main Authority
37      /OU=Certificate Management
38      /CN=main.authority.com/emailAddress=cert@authority.com
39      serial:E5:3A:5B:E2:BD:70:20:77
40
41      X509v3 Basic Constraints:
42      CA:TRUE
43    Signature Algorithm: sha1WithRSAEncryption
44      23:55:73:1b:5c:77:e4:4b:14:d7:71:b4:09:11:4c:ed:2d:08:
45      ae:7e:37:21:2e:a7:a0:49:6f:d1:9f:c8:21:77:76:55:71:f9:
46      8c:7b:2c:e8:a9:ea:7f:2f:98:f7:45:44:52:b5:46:a4:09:4b:
47      ce:88:90:bd:28:ed:05:8c:b6:14:79:a0:f3:d3:1f:30:d6:59:
48      5c:dd:e6:e6:cd:3a:a4:69:8f:2d:0c:49:e7:df:01:52:b3:34:
49      38:97:c5:9a:c3:fa:f3:61:b8:89:0f:d2:d9:a5:48:e6:7b:67:
50      48:4a:72:3f:da:28:3e:65:bf:7a:c2:96:27:dd:c0:1a:ea:51:
51      f5:09

```

## B Injected HTML Code

The following HTML code was injected by the relay 5A2A51D4 (see Table 1). It was appended right in front of the closing HTML tag.

```

1 <br>
2 

```

When requesting the image link inside the HTML code, the server responds with another HTML document. The full HTTP response is shown below.

```

1 HTTP/1.1 200 OK
2 Date: Tue, 14 Jan 2014 17:12:08 GMT
3 Server: Apache/2.2.22 (Ubuntu)
4 Vary: Accept-Encoding
5 Transfer-Encoding: chunked
6 Content-Type: text/html
7
8
9 <HTML>
10 <HEAD>
11 <TITLE>No Title</TITLE>
12 </HEAD>
13 <BODY>
14
15 </BODY>
16 </HTML>

```