


Submitted by
Thomas Raab, BSc

Submitted at
**Institute of Networks and
Security**

Thesis Supervisor
**Univ. -Prof. Priv.-Doz.
Dr. René Mayrhofer**

April 2021

Unlinkable Onion Services: Improved Resilience against Metadata Analysis



Master Thesis
to obtain the academic degree of
Diplom-Ingenieur
in the Master's Program
Computer Science

Sworn Declaration

I hereby declare under oath that the submitted Master's degree thesis has been written solely by me without any third-party assistance, information other than provided sources or aids have not been used and those used have been fully documented. Sources for literal, paraphrased and cited quotes have been accurately credited. The submitted document here present is identical to the electronically submitted text document.

City, Date

Signature

Kurzfassung

In unserer digitalisierten Gesellschaft, in der verschiedenste Interessengruppen versuchen die Internetnutzung zu kontrollieren und zu überwachen, ist Anonymität eine der wesentlichen Eigenschaften um Privatsphäre im Internet zu gewährleisten. Eine der Techniken, die zur Wahrung dieser Eigenschaft eingesetzt wird, ist das Anonymisierungsnetzwerk *Tor*, welches die Verbindungsdaten einer Kommunikation so verschleiert, dass deren Initiator nicht mehr identifiziert werden kann. Da dies aber nur den Initiator, nicht aber weitere Kommunikationsteilnehmer schützt, wurden *Tor Onion Services* entwickelt, welche die Anonymität des Senders und Empfängers sicherstellen. Bedingt durch die Metadaten die bei der Nutzung dieser Onion Services entstehen, könnten Angreifer durch die Nutzung zusätzlicher Informationsquellen dennoch in der Lage sein, die Teilnehmer einer Kommunikation zu identifizieren.

Im Zuge dieser Arbeit wurde ein Protokoll entwickelt, welches Metadaten, die zur Identifikation von Kommunikationsteilnehmern führen können, weitestgehend reduziert. Dazu wird ein zweistufiges Adressierungsschema eingesetzt, das Benutzern ermöglicht, eine individuelle Adresse für einen Service über eine öffentliche Service Adresse zu erhalten, deren Verwendung nicht zurückverfolgt werden kann. Um die technische Machbarkeit dieses Protokolls unter Beweis zu stellen, wurde ein Prototyp auf Basis von Python entwickelt. Da die Latenzzeit eines Services eine der ausschlaggebenden Kriterien in der Nutzungsentscheidung darstellt, wurde eine Performance-Analyse durchgeführt, um die Provisionzeit von Onion Services zu messen, da diese einen maßgeblichen Einfluss auf die Dauer der Adressausstellung hat. Die Architektur und Vorgehensweise dazu musste eigens konzipiert und implementiert werden, da zum Zeitpunkt der Erstellung dieser Arbeit keinerlei Forschung zur Provisionzeit von Onion Services in ihrer aktuellen Version existierte.

Eine statistische Analyse der Ergebnisse zeigte, dass die Dauer der Ausstellung individueller Adressen mithilfe des entwickelten Protokolls die Akzeptanzgrenze von Benutzern mit 6.35 Sekunden überschreitet. Das trifft jedoch nicht auf den Service-Zugriff unter Verwendung der individuellen Adresse zu, was impliziert, dass die Verwendung des Protokolls nach einer Verbesserung des Adressausstellungsmechanismus möglich ist. Das würde die Metadaten beim Zugriff auf einen Onion Service reduzieren und somit zur Verbesserung der Anonymität der Kommunikationsteilnehmer beitragen.

Abstract

In our digitized society, in which different organizations attempt to control and monitor Internet use, anonymity is one of the most desired properties that ensures privacy on the Internet. One of the technologies that can be used to provide anonymity is the anonymization network *Tor*, which obfuscates the connection data of communications in a way that its initiator cannot be identified. However, since this only protects the initiator without protecting further communication participants, *Tor Onion Services* were developed, which ensure the anonymity of both the sender and the recipient. Due to the metadata created when using these Onion Services, adversaries could still be able to identify participants in a communication by using additional sources of information.

In the course of this thesis, a protocol was developed that reduces metadata leading to the identification of communication participants as far as possible. For this purpose, a two-staged addressing scheme was employed that allows users to obtain an individual address for a service via its public service address, which cannot be traced back. To prove its technical feasibility, a prototype of the protocol was implemented based on Python. Since latency is one of the decisive criteria in the usage decision of services, a performance analysis was carried out to measure the provisioning time of onion services, since this has a significant influence on the duration of address issuing. The architecture and procedure for this had to be specially designed and implemented, as at the time of writing no research existed on the provisioning time of onion services in their current version.

A statistical analysis of the results revealed that the duration of issuing individual addresses using the proposed protocol exceeds the acceptance threshold of users with 6.35 seconds. However, this does not apply to service access using the individual address, implying that the use of the protocol is possible after improving the address issuance procedure. This would reduce the metadata when accessing an Onion service and thus help improve the anonymity of communication participants.

Sworn Declaration	i
1 Introduction	1
1.1 Problem Statement	3
1.2 Objectives	3
2 Related Work	4
2.1 Performance Measurement	4
2.1.1 Tor Performance	4
2.1.2 Onion Service Performance	5
2.2 Unlinkability	6
3 Background	8
3.1 Tor	8
3.2 Onion Services	11
3.2.1 Service Initialization	12
3.2.2 Service Access	18
3.2.3 Key Types	22
3.2.4 Time Periods and Shared Random Values	24
3.2.5 Security Enhancements	25
4 Unlinkable Onion Services	27
4.1 Methodological Approach	28
4.2 Custom Addressing	29
4.3 Unlinkable Onion Service Protocol	31
4.3.1 Concept	31
4.3.2 Architecture	32
4.3.3 Identity Assignment	33
4.3.4 Protocol Definition	35
4.4 Proof of Concept	38
4.4.1 Technical Building Blocks	38
4.4.2 Proxy Module Implementation	41
4.4.3 Server Module Implementation	45
5 Performance Analysis	50
5.1 Methodological Approach	50
5.2 Architecture	53
5.3 Implementation	55
5.3.1 Analysis Module	55
5.3.2 Container Environment	61
5.3.3 Analysis Server	62
5.4 Results	63
5.5 Discussion	67
6 Conclusion	71
6.1 Summary	71
6.2 Future Work	72
References	73
Appendices	77

A	Unlinkable Onion Service Protocol	77
A.1	Database Utils	77
A.2	Connection Utils	80
A.3	Server State Machine	81
A.4	Server Module	82
A.5	Server Tor Configuration	86
A.6	Proxy State Machine	86
A.7	Proxy Module	87
A.8	Proxy Tor Configuration	93
B	Performance Analysis	94
B.2	Circuit Analyzer	99
B.3	Log Analyzer	101
B.4	Descriptor Analyzer	102
B.5	Start Analysis	103
B.6	Utils	104
B.7	Docker File	105
B.8	Analysis Server	106
B.9	Tor Configuration	107
B.10	Initialization Script	107

1 Introduction

Communication is the cornerstone of our modern society. Although the demand for communication has not changed, the technological innovations in the last three decades caused major changes in the nature of communication. Messages and data are now no longer transmitted in analogue form, but exchanged via the Internet. The convenience and usefulness of the Internet made it increasingly popular among the population, allowing it to become an integral part of both the private and the professional daily life [1].

The foundation of communication over the Internet is the *TCP/IP* stack, which is a collection of network protocols that enable end-to-end communication between network participants. The stack is composed of four layers that describe the physical data transmission (Link Layer), the forwarding and routing of packets (Internet Layer), the end-to-end communication (Transport Layer) and the exchange of application-specific data (Application Layer) [2, pp. 45-48]. The Internet Protocol (IP) is one of the eponymous protocols of this stack, whose function is the routing of packets across network boundaries. Routing is often performed by independent devices, which are referred to as routers. In order for the routers to determine where to forward a packet, each IP packet contains an IP header in addition to the payload data. In order to make a routing decision, the routers must know the destination of the IP packet. Apart from the sender address and other routing-related information, this recipient address is part of the IP header. Due to the stateless concept of IP, in order to enable reliable data transmission, the Transmission Control Protocol (TCP) is used to transfer data between applications. Besides partitioning the byte stream to be sent into segments, TCP offers error detection as well as flow control and ensures that the segments sent are eventually reassembled in the correct order [3].

In case an IP packet is transmitted without any protection, every router on the path from the sender to the recipient would theoretically be able to inspect and manipulate its content. To overcome this issue, an encryption protocol such as Transport Layer Security (TLS) can be applied to TCP segments, which provides authenticity, integrity and confidentiality for data transmission. Since TLS operates on a higher layer in the TCP/IP stack compared to IP and the IP address of the recipient is needed for the routing procedure, the IP header is not protected by TLS. At first glance, this may not seem like an issue, as the actual content of the message is encrypted, nevertheless, the following information can be obtained from the non-encrypted metadata of the communication: [2, pp. 853-857]

- Both the sender and the recipient of the communication can be identified from the IP header. Although it is not possible to determine what is being communicated, it is possible to discover who is communicating with whom.
- The TCP header, which is also not encrypted when using TLS, reveals the source and destination port of the communication. This allows conclusions to be drawn about the application that is being used in the communication (e.g. port 443 indicates that a website is accessed).

- The frequency and time of communication can be determined from the IP header and the intercepting system.

For example, a company could use this information to monitor whether employees are browsing online job portals or social media during working hours, provided that only one website is running on the monitored servers at a time or that Server Name Indication (SNI) is available. Another example are state actors that could use the meta-information provided in the IP and TCP header to enforce censorship and oppress their opponents' freedom of speech by disturbing or preventing access to their means of communication, e.g. Internet forums.

In order to protect not only the payload but also the metadata of the communication and thus the communication participants, anonymization services like the *Tor* network were developed [4]. Tor uses a concept referred to as *Onion Routing* to disguise the identity of the client based on TCP. A schematic illustration of general functioning of Onion Routing is shown in Figure 1.1.



Figure 1.1: Tor Onion Routing Network

Using Onion Routing, a message is not sent directly to the recipient, but routed through several randomly selected Onion Routers (OR). In order to prevent an OR from accessing the data of the TCP segment, the initiator establishes a shared key with each OR in advance. It encrypts the segment with the shared keys to generate a layered encrypted object that conceptually resembles an onion, due to which this procedure is named. On the path from the initiator to the responder, each OR decrypts one layer with its shared key. The decrypted data contains information on the next hop in the path, which implies that each OR only knows its predecessor and successor. The reply of the responder is encrypted by each OR with its shared key with the initiator, so that ultimately only the initiator can decrypt the response. [4]

Onion Routing protects only the anonymity of the initiator, but not that of the responder, whose IP address must be known to enable communication. To overcome this issue, Tor onion service were developed to ensure both initiator and responder anonymity within the Tor network. In order to achieve this level of anonymity, services are no longer addressed through their IP address, but through an *onion address*. This address is based on the key material that is generated for an onion service during its creation. Each service chooses several *Introduction Points* (IPs), which are the first point of contact for communication with the onion service. Information about these IPs and further service-related data are consolidated in service descriptors. These descriptors are uploaded to specific *Hidden Service Directories* (HSDirs) depending on the initial created key material and the time of upload. A user who wants to access the service needs to know its onion address. With this address, the user can download the service descriptors from the HSDirs, which contain the addresses of the IPs. Once the user obtains this information, they choose a random node in the Tor network that will acts as its *Rendezvous Point* (RP) and establishes a connection to it. In principle,

this RP eventually becomes the middleman between the service and the user. Subsequently, the user contacts one of the IPs and provides the chosen RP. The IP relays the request to the service, which in turn decides whether it wants to communicate with the client. If yes, it connects to the RP provided by the client, so that both client and service are now connected to the RP, which enables communication between client and service.

1.1 Problem Statement

Since the Tor network is run by volunteers, anyone can contribute computational resources in the form of computers or servers, that can eventually become RP, IP or HSDir of an onion service. This does not seem like an issue at first glance, because even if an adversary were to control one of the onion service components, it would only be able to collect minimal amounts of metadata that on their own can be linked to neither a server nor a client.

However, the situation is different if an adversary is able to obtain additional metadata from another source and correlate this data with the metadata of the onion service. This scenario is particularly effective if an adversary can control one of the HSDirs of a service. Since the descriptors are stored encrypted on the HSDir, no information can be gained from them, but it is possible for the adversary to determine when and how often a particular onion service is accessed, because it can monitor the HTTP requests associated with the descriptor downloads. If the adversary is able to establish a temporal link between the download time of the descriptor from the HSDir and other information such as Tor traffic on the corporate network the user is connected to, it may be possible to identify the user, or at least the device that submitted the HTTP request.

1.2 Objectives

The goal of this thesis is the proposal of a novel protocol that allows the access to an onion service, while preventing its linkability through HSDir descriptor downloads. In order to demonstrate its technical feasibility, a Proof of Concept (PoC) shall be implemented, that enables clients to communicate with a Hypertext Transfer Protocol (HTTP) server. In order to assess whether the proposed protocol is usable in practice, a performance analysis shall be carried out to measure the time it takes to provisioning an onion service, as this will be an integral part of the proposed protocol. Thereby, the different variants of onion service shall be compared. The results of the measurements shall be examined by means of a thorough statistical analysis to uncover any correlations. Using the results from the prototype implementation and the performance analysis, it shall finally be determined whether the developed protocol can reduce the metadata of an onion service access with reasonable additional latency.

2 Related Work

Since the research in the area of onion services is very focused, the related work section is structured according to the two main components of this thesis.

2.1 Performance Measurement

Köpsell et al. [5] discovered that an increase of delay in terms of latency in anonymization services is linearly related to the drop-out rate of users, as less security-affine participants question the utility of this service and stop using it. This implies that the number of users decreases as the delay in the anonymity service increases, which has a negative impact on the quantity of anonymity provided by the service. Besides usability, this is the major motivation for conducting performance measurements in anonymization services. The types of performance measurements can roughly be divided into measurements of the performance of the Tor network and measurements of the performance of onion services.

2.1.1 Tor Performance

Wendolsky et al. [6] compared the performance of the two low-latency anonymization services *Tor* and *AN.ON* regarding the performance indicators latency and bandwidth. Based on the work of Köpsell et al. [5] they conclude that the tolerance limit for latency is four seconds for inexperienced users. This outcome can only partially be used for comparison with the results of this thesis, because the architecture of Onion services is much more complex and involves more nodes. Furthermore, Wendolsky et al. [6] noticed performance variations in the use of Tor depending on the time of day.

To measure long-term performance trends in the Tor network, Custura et al. [7] developed *Onion-Perf*. This software tool measures Tor network performance by tracking the bulk download of random data produced by traffic generators. This random data is either hosted on a local machine to emulate access via the Internet, or the data is hosted via onion service to simulate access to an onion service [8]. In contrast to the previously described work, Panchenko et al. [9] focused their study on the examination of node limitations and their relations to latency and throughput. They carried out several experiments in a private tor network as this represents an optimal environment without external disruptive influences. One of their findings indicates that the more clients use a node to build a circuit, the lower is its throughput. An interesting fact is that the throughput remains constant from about 14 clients on, no matter how many circuits the node is involved in.

Furthermore, Panchenko et al. [9] analysed the effect of different parameters of path selection on the *Round Trip Time* (RTT) of a circuit. The shortest mean RTT is achieved when the selection of nodes on the path is not carried out uniformly, but weighted according to the advertised bandwidth of the nodes. Moreover, the study indicates that the RTT of a circuit also depends on the geographical diversity of the nodes in the circuit.

2.1.2 Onion Service Performance

Lösing et al. [10] were the first to analyse the performance of onion services in terms of latency in a scientific publication. In this paper, a special focus is laid on overall response time, as the work of Köpsell et al. [5] suggests that latencies and connection setup times have a much greater effect on service usability than bandwidth. For the measurements, two Tor instances are used. The first instance serves as an onion service, while the other instance is configured as a relay, acting as an *Introduction Point* (IP). Changes in the source code ensure that both the onion service and the client use the IP controlled by the authors. The *Rendezvous Point* (RP) is not controlled by the authors because the selection of a specific RP does not work in case of circuit cannibalization. Clients for accessing the onion service were periodically created to avoid any caching of the *Onion Proxy* (OP) on the client. Every client performs only a single access-attempt before it is discarded. The mean overall response time for the measurements was about 24 seconds. Considering the overall duration in relation to the individual sub-steps for accessing an onion service, it is noticeable that building the circuits from the client to the introduction point and fetching the service descriptor takes almost half of the overall duration.

Similarly to Lösing et al. [10], Wilms [11] measures the overall response time of the access to an onion service by controlling the individual instances that are needed to operate the service. In contradiction to Lösing et al. [10], Wilms [11] found a way of also controlling the RP. Despite the *RendNodes* option in the Tor client configuration allowing the selection of a specific RP, a different RP was previously selected because three internal three-hop circuits are generated during bootstrapping, which are cannibalized when the connection to the RP is established. Since these internal circuits were built without considering the desired RP, a different RP was chosen. By changing the Tor source code, an internal circuit with the RP as last node is built at bootstrapping, which is then cannibalized as RP circuit. The RP, the onion service and the OP are operated on the same machine, as no direct communication between the components can introduce bias to the connection times. Wilms [12] use the tool *PuppeTor* [13] to deploy the Tor clients and measure the overall response time, which was on average about 39 seconds. This is considerably longer than the overall response time measured by Lösing et al. [10]. However, the reason for this increase could not be determined.

Lehnhard et al. [14] focus their work on measuring the performance of onion services in low-bandwidth access networks, because the communication overhead of anonymization services becomes even more apparent in this context. As in the work of Wilms [11], the IP, RP and Onion services are controlled by the authors and hosted on the same server. In the first measurement phase the client uses a low-bandwidth access network (analog modulation via the telephone network and *Enhanced Data Rates for GSM Evolution* (EDGE)), while the server uses a broadband network access. In the second measurement phase, this concept is reversed, such that the measurement server hosting the RP, IP and the onion service uses the low-bandwidth access network and the clients use the broadband network. Similar to the other studies, new clients for access to onion services are regularly created, which are discarded after one access attempt in each case, in order

to avoid incorrect measurements through caching. In addition to measuring the overall response time, a measurement of the bootstrap time of each Tor client instance is carried out. The measurement results show that the bootstrap time is a major problem when accessing onion services from a low-bandwidth access network, since it takes about five times longer (232.9 seconds for EDGE and 249.0 seconds for modem compared to 22.9 seconds for broadband) than using a broadband network. Considering the overall response time, the access of a client to an onion service using a low-bandwidth access network takes about twice as long as using a broadband network. In the reversed scenario where the server is using the low-bandwidth access network, the difference shrinks to eight seconds between EDGE and broadband and only one second between broadband and modem. The analysis of the circuit creation times shows that they represent a bottle-neck when accessing onion services, especially when using a low-bandwidth access network.

There is currently little to no research on the provisioning time of onion services. Furthermore, all the above mentioned work is based on the outdated onion service V2 protocol, which will no longer be supported in the foreseeable future [15]. This work fills this gap by providing a comprehensive performance analysis of the provisioning of a V3 onion service including all necessary sub-steps.

2.2 Unlinkability

Oeverlier et al. [16] were the first to publish an attack on the anonymity of onion services. For this attack it is assumed that the attacker controls both a relay in the network and a client of an onion service. Any relay in the Tor network that offers stability can be used by the onion service to build a circuit to the RP. The attacker creates several rendezvous circuits to the onion service and sends a specific traffic pattern over these circuits. With the help of logged time values and direction information from the client and the relay, the attacker can determine whether his relay is part of a circuit to the onion service. If this is the case, the attacker can further ascertain at which position his node is located in this circuit. Since the attacker has selected the RP, he knows its address. This enables him to determine if his node is closest to the RP. If the node has unknown addresses on each side, it can be either the middle node in the circuit or the first node of the path to the onion service. This distinction can be made using a timing analysis or the *predecessor attack* [17]. If the attacker's node is the first node on the path to the onion service, the attacker has the IP address of the onion service and thus can conclude its position. To prevent this type of attack, *Guard nodes* have been introduced, which are described in section 3.2.5.

Biryukov et al. [18] focus on the discovery of onion services through the collection of onion service descriptors. They claim to have collected all service descriptors available at the time within two days at the cost of only 100 USD. To collect the service descriptors of all onion services, it would require a very large number of relays (approximately half the number of hosted onion services at a given time point) acting as HSDir. Since the Tor specification limits the number of relays per IP address to two, it would also require a high number of unique IP addresses apart from the systems. To be able to collect the descriptors more efficiently, they use so-called *shadow relays*. These are relays that do not appear in the consensus but are still monitored by the directory authorities and

thus obtain flags. If one of the two running relays under one IP address fails, a shadow relay takes over, which then also acts as HSDir, since it has already received the required flags. This way it is possible to fade in and fade out HSDirs. Using this method, it was possible to operate 600 Tor relays on just 50 IP addresses. Under each IP address 24 relays were operated, which were gradually disabled over a period of 24 hours to cover the complete range of the Tor *distributed hash table* (DHT). The described attack is now no longer possible because the flag assignment vulnerability has been fixed and the distribution of onion services V3 on the DHT works in a different way such that positions can not be pre-calculated.

Owen et al. [19] similarly to Biryukov et al. [18], collect the addresses of onion services by placing nodes in the DHT. As the use of shadow relays was no longer possible due to a fix in the Tor source code, the resources in terms of time and relays are considerably higher. They were able to collect approximately 80.000 unique onion addresses, which they subsequently classified. From the results it can be concluded that about 80% of the requests at the time were related to onion services that offered abusive sites (sites where the title suggests some form of sexual abuse).

Elices et al. [20] propose a method for the identification of a hidden server under the assumption that the adversary operates on an ISP level and is able to monitor the traffic between the onion server and its guard node. With access to this information, they are able to correlate the increment of TCP traffic on the link between the hidden server and its guard node with a client request and thus get the IP address of the hidden server.

Biryukov [18] et al. and Ling et al. [21] both propose a method to determine whether a relay controlled by the attacker is the guard-node of an onion service based on characteristics of the Tor protocol. These attacks offer a high accuracy in the identification of guard nodes, but are limited to one onion service per attack. Chen et al [22] resolve this problem by embedding a identifier for each attacked onion service as watermark to distinguish them. According to the authors, this speeds up the attack process by a factor of 11.6.

Current research is focused on the identification of the IP address of onion services and on harvesting onion service addresses. Almost every study contains measures to increase the anonymity of onion services, but metadata is often disregarded. In the course of this work, a protocol was developed that minimizes the linkability of an onion service through its metadata.

3 Background

3.1 Tor

Tor is an open-source *Onion Routing* network that enables bi-directional anonymous communication for TCP-based applications [4]. The goal of this network is to enable users to hide their identity from services they use and limit the ability of traffic analysis [23]. This technique can be used, amongst others, to circumvent censorship measures and access policies, to disguise the web-surfing habit of users or to anonymously provide services to other users. In an ordinary network architecture, the routing information provides details about the identity and location of the requesting user. Even if the payload is encrypted, the routing information has to be transmitted in plain text, as the routers need to be able to read it to find the right path [23].

This issue is addressed by means of Onion Routing. Each user runs a local client referred to as *Onion Proxy* (OP), which handles TCP connections from applications and fetches directory information which consists, amongst others, of a list of known Onion Routers (ORs) and their current state. As Tor utilizes the standardized *SOCKS* proxy protocol as specified in the *RFC1928* [24], most applications can use it without major modifications. The OP selects a path through the network referred to as *circuit*, which consists of OR nodes that only know their direct successor and predecessor. An example of such a circuit is depicted in Figure 3.1.




Figure 3.1: Onion Routing Network [23]

In this example the *Initiator* wants to transmit data to the *Responder*. For this purpose its OP builds a circuit through OR_1 , OR_3 and OR_4 . Each OR holds a long-term identity key, which is mainly used to sign TLS-certificates, router information and directories and a short-term onion key [4]. This onion key consists of a public part, which is made available to all network participants, and a private part, which is only known to the OR. The public part of the onion key is used to encrypt requests during circuit creation, which can only be decrypted by the OR, who is in possession of the corresponding private onion key. Communication in the Tor network is based on fixed-sized *cells* of 512 bytes, each consisting of a header and a payload. The header contains a circuit identifier *circID* and a command, that specifies how the data in the payload is processed. A fundamental distinction is made between *control* cells and *relay* cells. While control cells are interpreted directly by the receiving node, relay cells are used to relay end-to-end data streams. A graphical representation of

the structure of those two cell types is given in Figure 3.2. In comparison to control cells, relay cells have an additional header containing a stream identifier (*StreamID*), which allows the multiplexation of multiple streams over a circuit, an integrity checksum, the payload length and a relay command. This relay command is utilized for the management of data streams over the circuit. [4]




Figure 3.2: Tor Cell Types [4]

The subsequent information on the Tor protocol procedure is taken from the Tor specification [25]. In order to create a new circuit, the initiator’s OP needs to send a control cell containing the *create* command to the first hop in the chosen path. For ease of reference, the path from the example in Figure 3.1 is used. The OP chooses a circID not currently in use and sends the first half of its handshake to OR_1 . This handshake is necessary for the establishment of a shared symmetric key between the OP and each OR on the path. Currently, two different types of handshakes are supported, which are both based on the *Diffie-Hellman* (DH) key exchange:

- **The *Tor Authentication Protocol* (TAP) Handshake**

Within the course of this handshake option, a regular DH handshake in \mathbb{Z}_p is performed. For this purpose, the initiator generates a private key x randomly. The initiator constructs an onion-skin, which refers to the DH public key g^x . This onion-skin is encrypted using the *legacy hybrid encryption* algorithm, in the course of which an additional symmetric key K is generated, that is then encrypted together with the first part of the public DH key (g^x) and additional padding with the public short-term onion key of the second participant using *RSA*. The second part of the public DH key (g^x) is encrypted using the symmetric key K . Both encrypted parts are transmitted to the second participant as the payload of a create cell.

The second participant first decrypts the ciphertext containing the symmetric key and the first half of the DH handshake using its private key. Using the decrypted symmetric key K , it can now decrypt the second part of the DH handshake and can obtain g^x . After receiving the DH public key, the participant randomly generates a private key y and is now able to compute the shared key g^{xy} . It sends its DH public key g^y along with data that demonstrates knowledge of the computed shared key to the initiator. Once the handshake is completed, the initiator can calculate the shared key g^{xy} . This shared key is used as an input to a *Key Derivation Function* (KDF) that outputs keys to secure the connection between the two participants.

- **The *ntor* Handshake**

With this handshake option, multiple DH handshakes are used to generate a shared key. The following calculations are carried out on the group *curve25519* as specified by Bernstein [26].

In order to perform this handshake, the initiator needs to know the public ntor onion key (hereinafter referred to as B) and the digest of the identity key of the second participant (Hereinafter referred to as ID). In the first step the initiator calculates a temporary keypair x, X using a curve25519 key generation algorithm and generates a client-side handshake containing the identity digest of the other participant, its ntor onion key and the public part of the newly generated key pair (X). This handshake is referred to as client handshake and is sent to the other participant, which in turn generates a keypair y, Y and computes the following values using its private ntor onion key b :

$$\begin{aligned}
 secret_input &= EXP(X, y) | + EXP(X, b) | ID | B | X | Y | PROTOID \\
 KEY_SEED &= H(secret_input, t_key) \\
 verify &= H(secret_input, t_verify) \\
 auth_input &= verify | ID | B | Y | X | PROTOID | "Server"
 \end{aligned} \tag{3.1}$$

The values $PROTOID$, t_key and t_verify are statically defined strings. The participant generates its handshake consisting of the public key of the newly generated key pair (Y) and a hash of $auth_input$ and sends it to the initiator. This is referred to as the server handshake. The initiator checks whether the received public key is a group member and calculates:

$$\begin{aligned}
 secret_input &= EXP(Y, x) | EXP(B, x) | ID | B | X | Y | PROTOID \\
 KEY_SEED &= H(secret_input, t_key) \\
 verify &= H(secret_input, t_verify) \\
 auth_input &= verify | ID | B | Y | X | PROTOID | "Server"
 \end{aligned} \tag{3.2}$$

The initiator can now generate a hash of $auth_input$ and compare this value with the value it received from the other participant. If these match, both parties are now in possession of a key seed, which serves as input for a KDF that generates the keys needed to secure the communication.

As the TAP handshake is rather slow due to the usage of RSA, the ntor handshake is preferred. After completing the handshake, relay cells between OP and OR_1 are encrypted using the established key. To extend the circuit, the OP sends a relay cell with the command extend to OR_1 , which contains the next hop on the path and the ntor client handshake with OR_3 . OR_1 creates a circuit to OR_3 and decrypts the relay cell it received from OP. It copies the received half client handshake into a create cell and passed this cell to OR_3 . OR_3 completes the handshake and responds with a control cell with the *created* command, which includes the server handshake. OR_1 receives the response and forwards it to OP using a relay cell. To extend the circuit, the OP acts as described above, always asking the current last node in the path to extend one hop further. Once the circuit is established, the OP shares a session key with each node on the path. To transmit data anonymously, the OP encrypts the data layer by layer with the established session keys of the respective OR starting with the last hop on the route [23]. This layered encryption scheme inspired the name of this mechanism,

as it is reminiscent of the layers of an onion, as illustrated in 3.3. The ellipses show the encryption layers of the onion using the session keys of the ORs (PK_{OR_x}). On the way from the OP to the responder, each node removes a layer of encryption so that the last node in the path can read the plaintext and make a request to the desired resource. The responder's reply to the OP is secured in the same way. Each OR on the path adds a layer of encryption using its shared key with the OP, such that only the OP can decrypt it.




Figure 3.3: Onion Structure

3.2 Onion Services

Onion services, formerly known as hidden services, offer a method for providing bidirectional TCP-based services on the internet while maintaining the anonymity of the provider. This is achieved through the use of an intermediary node through which the communication is handled. Both the client and the service use the Tor network to establish an anonymous connection to this intermediary, thus allowing communication to take place without the client knowing the location of the service and the service knowing the location of the user. As the IP address specifies the location of a network device, the service is now no longer contacted via its IP address, but instead via an address that is specific to the Tor network. This address is referred to as *onion address* and is unambiguously associated with the service. Despite the need to install specialized software for the set-up of such onion services, no changes to the service to be provided are necessary.

The behaviour and functioning of onion services are specified in the *Tor Rendezvous Specification*, which is currently available in version three [25] and is the basis for the information provided in this chapter. The following subsections describe the content of this specification and the process for providing and accessing onion service in detail. Figure 3.4 provides an illustrated overview of the steps involved in this process. The numbered labels match the respective textual descriptions in the subsections. [4, 25]




Figure 3.4: Onion Service Architecture

3.2.1 Service Initialization

1. Generate Long-term Identity Key

Initially, the provider generates a long-term key pair, which identifies the service. This key pair is kept offline and is solely used for the generation of blinded signing keys. A detailed description of the menagerie of keys used and their interrelationships is provided in section 3.2.3.

2. Choose Introduction Points

Once the identity keypair is created, the hidden service selects several Tor nodes as its introduction points. These introduction points represent the initial contact points for clients to contact an onion service. The number of introduction points a service uses ranges between zero and twenty. If an onion service has zero introduction points, it cannot be reached by any client. The default number of introduction points defined in the specification is three [25]. As first part of the selection process, the onion service builds an anonymous three-hop circuit to its randomly chosen introduction points. In case of a single onion service being used, which trades service-side location privacy for improved performance [27], non-anonymous one-hop circuits are utilized for that purpose.

Subsequently, the onion service generates a short-term introduction point authentication key pair, which is used to ensure the authenticity of messages passed between the onion service and the introduction point during the establishment of the introduction point. Furthermore, the public component of this key is included in requests from clients to introduction points to ensure that the intended introduction point was reached. In order to register an introduction point, the onion service first builds a circuit to that node and sends an *ESTABLISH_INTRO* relay cell, which contains the public part of the generated introduction point authentication key, the length and type of the introduction point authentication key, possible extension (e.g. DoS defence), a *Message Authentication Code* (MAC) and a signature, which is signed with the private part of the introduction point authentication key. When a Tor node receives an *ESTABLISH_INTRO* cell, it checks whether:

- The key type is recognized
- The key is well formatted
- The signature is correct (using the public key included in the request)
- The MAC is valid
- The circuit isn't already a rendezvous circuit or an introduction circuit
- The key isn't already in use

If any of these checks fail, the Tor node rejects the cell and tears down the circuit. If all above verifications pass, the introduction point reports its status back to the onion service using an *INTRO_ESTABLISHED* relay cell. From this point on, the Tor node acts as an introduction point for the onion service.

3. Build Service Descriptors

In the next step the onion service generates a set of descriptors, which contain, among other things, information on the introduction points of the onion service, which is necessary in order to contact them. As these descriptors are going to be encrypted twice, the first structure generated by an onion service is referred to as the second layer plaintext. The format of this second layer plaintext is depicted in Figure 3.5.

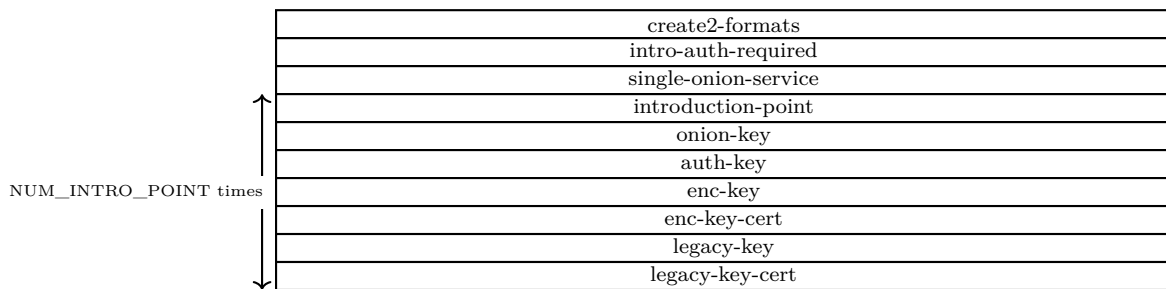


Figure 3.5: Onion Service Descriptor Second Layer Plaintext Format

The field *create2-formats* contains a list of accepted *CREATE2* cell types. Thereby, at least the current default handshake type *ntor* must be supported. As an onion service may restrict access to authorized users, the client must be aware of the introduction-layer authentication types the onion service supports. These types are stored in the field *intro-auth-required*. Currently, authentication is possible by the use of an *Ed25519* key pair or a password. The field *single-onion-service* indicates whether the onion service is provisioned as a single-onion service using non-anonymous circuit connections to introduction points and rendezvous points. After this block specific to the onion service the second-layer-plaintext contains the following items for each of its introduction points.

- *introduction-point*: Base64 encoded link specifiers and identities of the Tor node acting as introduction point
- *onion-key*: Base64 encoded medium-term ntor onion key of the introduction point used when a client extends to it.
- *auth-key*: A certificate cross-certifying the public introduction point authentication key using the descriptor signing key.
- *enc-key*: Base64 encoded public key used to encrypt requests to the onion service.
- *enc-key-cert*: A certificate cross-certifying the public encryption key using the descriptor signing key.
- *legacy-key*: RSA public key used for a legacy introduction point.
- *legacy-cert*: A certificate cross-certifying the legacy key using the descriptor signing key.

The second-layer plaintext is encrypted as specified in equation 3.3, whose design is intended to protect the confidentiality of the descriptor against unauthorized clients.

$$ENCRYPTED = STREAM(SECRET_IV, SECRET_KEY) XOR PLAINTEXT \quad (3.3)$$

The initialization vector $SECRET_IV$ and the key $SECRET_KEY$ used for this encryption are computed as denoted in equation 3.16 in section 3.2.3.6. If client authentication is enabled, the onion service generates 32 random bytes as a *descriptor_cookie*. In case client authentication is not used, the *descriptor_cookie* is left blank. The following parameters are used for the key derivation:

- $SECRET_DATA = \text{blinded-public-key} \mid \text{descriptor_cookie}$
- $STRING_CONSTANT = \text{"hsdir-encrypted-data"}$

The ciphertext is stored in the *encrypted* field of the first-layer plaintext, whose structure is depicted in Figure 3.6.

desc-auth-type
desc-auth-ephemeral-key
auth-client
encrypted

Figure 3.6: Onion Service Descriptor First Layer Plaintext Format

The field *desc-auth-type* describes the type of authorization which is used to protect the descriptor. Currently, "x25519" is the only valid value for this field, which represents the multi-staged cookie encryption scheme described below. For this the onion service generates an ephemeral x25519 key pair, whose public key is stored in the *desc-auth-ephemeral-key*. It is assumed that each authorized client holds a shared x25519 public key with the onion service. If client authentication is enabled, the *descriptor_cookie* is encrypted for each authorized clients. As multiple clients could be authorized for one onion service, the *auth-client* section contains one encrypted *descriptor_cookie* per client. The keys for this encryption are computed as follows:

$$\begin{aligned}
 SECRET_SEED &= x25519(\text{hs_y}, \text{client_X}) \\
 KEYS &= KDF(\text{subcredential} \mid SECRET_SEED, 40) \\
 CLIENT - ID &= \text{first 8 bytes of KEYS} \\
 COOKIE - KEY &= \text{last 32 bytes of KEYS}
 \end{aligned} \quad (3.4)$$

where *hs_y* denotes the associated ephemeral *x25519* private key for the public key in *desc-auth-ephemeral-key*, *client_X* is referred to as the pre-shared *x25519* public key of the authorized client and *subcredential* is generated during the key blinding phase for each period and can be generated by clients with the knowledge of the onion service public identity key.

Each line in *auth-clients* contains the base64 encoded *CLIENT_ID*, a base64 encoded initialization vector *iv* and the base64 encoded encrypted cookie in the *encrypted-cookie* field. The cookie is encrypted as described in equation 3.5.

$$\text{encrypted-cookie} = \text{STREAM}(iv, \text{COOKIE-KEY}) \text{ XOR } \text{descriptor_cookie} \quad (3.5)$$

The first-layer plaintext is eventually encrypted using the same cipher as in equation 3.3. The following parameters are used for the key derivation, based on equation 3.16 in section 3.2.3.6:

- SECRET_DATA = blinded-public-key
- STRING_CONSTANT = "hsdir-superencrypted-data"

This can be seen as a first step towards access control, as a client needs to know the unblinded onion address in order to calculate the blinded onion address for the respective time period. The encrypted data blob is base64 encoded and stored in the field *superencrypted* of the onion service descriptor structure, which is illustrated in Figure 3.7. The format follows the same meta structure as other Tor directory objects [25].




Figure 3.7: Hidden Service Descriptor Format

The field *hs-descriptor* contains the version number of the descriptor, which currently is three. The lifetime of the descriptor is specified in the field *descriptor-lifetime* and ranges between 30 and 720 minutes. A HSDir invalidates a descriptor when its lifetime is exceeded. The field *descriptor-signing-key-cert* contains a certificate that cross-certifies the short-term descriptor signing key with the blinded public key. The format of the certificate differs from standardized certificate formats and is specified in the *tor proposal 220* concerning the migration from *RSA* to *Ed25519* keys [28]. Its structure is depicted in Figure 3.8.

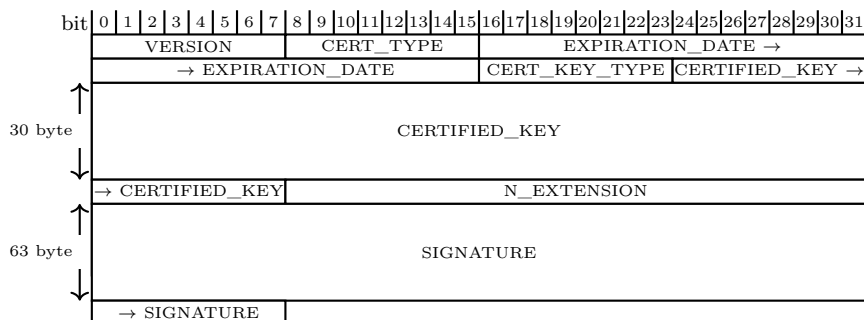


Figure 3.8: Tor Certificate Format

If an HSDir receives multiple descriptors for a descriptor-singing-key, the descriptor with the higher revision number specified in the *revision-counter* field is sent to the requester. The field *superencrypted* contains the twice encrypted and *base64* encoded blob, whose creation is described in detail above.

4. Publish Service Descriptors

In order to publish and fetch onion service descriptors, both services and clients must be able to examine the directories responsible for the service in the given time period. This is accomplished by the usage of *HSDir hash rings*, whose structure is depicted in Figure 3.9.

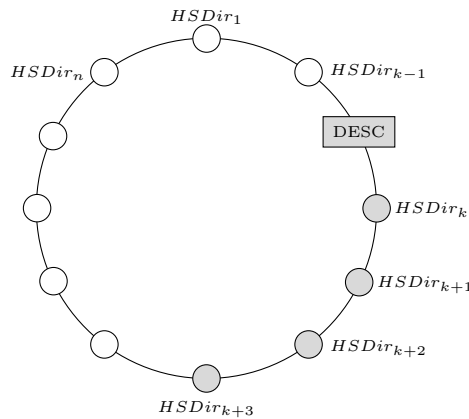


Figure 3.9: Onion Service Hash Ring Structure

For each HSDir in the consensus a hash value is calculated that includes information about the identity of the directory, as well as information about the current time period. This hash value specifies the position of the directory in the hash ring, as indicated by $HSDir_x$ in the figure above, where x ranges from one to n depending on the position of the descriptor in the hash ring. This position is computed on the basis of the blinded public key of the service and information of the current time period and is denoted as $DESC$ in Figure 3.9. By default, the descriptor is uploaded to the first four directories ($HSDir_k$ to $HSDir_{k+3}$), which numerically directly follow the position of the descriptor in the hash ring. A detailed description of the calculation of indices is provided later in this chapter.

To prevent a set of descriptor directories from censoring a service, the onion service uploads the descriptor to several different directories that vary over time. The frequency of directory changes (i.e. the length of the time period) depends on the *hmdir-interval* parameter, which is a consensus parameter established between directories. The current default value for *hmdir-interval* is 1140 minutes, which equals to one day. The time period number, which is necessary for the calculation of the index of the onion service descriptors and its responsible directories, is determined as shown in equation 3.6.

$$period_num = \frac{minutes_since_epoch - rotation_time_offset}{time_period_length} \quad (3.6)$$

The parameter *minutes_since_epoch* stands for minutes since the Unix epoch¹. As the developers of *hidden services* wanted to start the time period at *12:00 UTC*, an offset is deducted (*rotation_time_offset*). Finally, the difference is divided by the period length, which is specified in *hmdir_interval*. The result of the calculation is the time period number (*period_num*). To calculate the position of a service descriptor within the hash ring, the following parameters from the *consensus* are required:

- *hmdir_n_replica*: Number of replicas used for uploading the onion service descriptor in the range of [1,16] with a default value of 2.
- *hmdir_spread_fetch*: Number of nodes a client chooses from, when fetching a onion service descriptor in the range of [1,128] with a default value of 3.
- *hmdir_spread_store*: Number of nodes to which an onion service descriptor is uploaded in the range of [1,128] with a default value of 4.
- *valid-after*: Uniform time value for calculating the current period number to minimize the desynchronization of clients and services by asynchronous clocks.
- *router_status*: List of onion router nodes including their identity, address, port and flags.

Now for every time period, the onion service calculates:

$$\begin{aligned}
 &\text{for replicanum in } 1 \dots \text{hmdir_n_replicas:} \\
 &hs_index(replicanum) = H(\text{"store-at-idx"} \mid \\
 &blinded_public_key \mid \\
 &INT_8(replicanum) \mid \\
 &INT_8(period_length) \mid \\
 &INT_8(period_num))
 \end{aligned} \tag{3.7}$$

The parameter *replicanum* is in the range of *hmdir_n_replica*. The *blinded_public_key* is derived from the public master identity key as described in section 3.2.3.2. The resulting *hs_index* is the index of the descriptor in the hash ring. In the next step, for every entry in the *router_status* consensus parameter that has the *HSDirV3* flag set, the directory index is calculated as follows:

$$\begin{aligned}
 &hmdir_index(node) = H(\text{"node-idx"} \mid node_identity \mid \\
 &shared_random_value \mid \\
 &INT_8(period_num) \mid \\
 &INT_8(period_length))
 \end{aligned} \tag{3.8}$$

As for the calculation of the *period_num* in equation 3.7, *period_length* is taken from the consensus. To limit the predictability of the position of a onion service descriptor in the

¹1970-01-01 00:00 UTC

hash ring, the authorities generate an additional random value (*shared_random_value*(*srv*)), which is added to the consensus and used in equation 3.8. As the transition of time period and random values does not happen simultaneously, at every point in time there are two valid random values. A detailed description of this behaviour can be found in chapter 3.2.4. When an onion service wants to publish its descriptors, it calculates *hs_index* for every replica and *hsdir_index* for every directory in the consensus. The service uploads the descriptors to the first *hsdir_spread_store* (denoted as *HSDir_k* to *HSDir_{k+3}* in Figure 3.9) nodes that immediately follow the numerical value of *hs_index*. The actual upload is performed using an *HTTP POST* request to the selected directory.

3.2.2 Service Access

5. Fetch Service Descriptor

In order to retrieve the descriptor from the responsible HSDirs, the client requires the onion address of the services, which must be transmitted using any kind of out-of-the-band communication. Once the client is in possession of the onion address, it is used to calculate the blinded onion address for the current time period, as described in 3.2.3. For each node in the consensus that obtained the HSDir (V3) flag, the *hsdir_index* is calculated as described in 3.2.1. Subsequently, the *hs_index* for both replicas is computed, based on the blinded public key and the current time period. Now the client randomly chooses one directory randomly from the set of responsible directories, which is composed of the first *hsdir_spread_fetch* directories which numerically follow the *hs_index*. To fetch the descriptor, a *HTTP GET* request for the URL */tor/hs/<version>/<z>* is initiated, where *<z>* represents the base64 encoded blinded public key of the onion service and *version* stands for the onion service protocol version. If the connection attempt fails, or the descriptor is not found by the directory, other directories from the previously chosen subset are queried.

6. Establish Rendezvous Point

The client decrypts the descriptor and now knows the introduction points of the service, which can be considered as initial contact points. Prior to connecting to the introduction point of an onion service, the client must establish a connection to a rendezvous point, which acts as an intermediary between client and service. For this purpose, the client selects a random Tor node, establishes a circuit to that node and sends an *ESTABLISH_RENDEZVOUS* cell containing a *RENDEZVOUS_COOKIE*. This cookie is an arbitrary 20-byte value and is later used when the onion service connects to the rendezvous point. The rendezvous points acknowledges the establishment of the rendezvous point by responding to the client with a *RENDEZVOUS_ESTABLISHED* cell.

7. Connect to Introduction Point

To contact the service, the client builds an anonymous circuit to one of the introduction points contained in the descriptor. Then it sends a *INTRODUCE1* cell with the following content to the introduction point:

- *LEGACY_KEY_ID*: Used to differentiate between legacy and new style introduction cells. If new style cells are used, this field is zeroed.
- *AUTH_KEY_TYPE*: Type of the authentication key used (currently only the authentication using an *Ed25519* public key is possible).
- *AUTH_KEY_LEN*: Length of the authentication key.
- *AUTH_KEY*: Introduction point authentication key (taken from the second-layer plaintext of the descriptor)
- *EXTENSIONS*: Possible extensions.
- *ENCRYPTED*: Encrypted plaintext, which provides information about the RP selected by the client. Dependent on the type of encryption key used, there are multiple ways of decrypting this ciphertext.

In addition to exchanging information about the RP of the client, handshake data is exchanged using the *INTRODUCE1* cell. Thereby, the following requirements are specified for a handshake as part of the introduction protocol:

- The onion service must be able to decrypt additional information, that contains the rendezvous token and the information required to extend to this RP.
- During the handshake, a set of shared keys is established, which can later be used to secure traffic between onion service and client.
- The onion service must be able to verify the authenticity of the cell, such that any modification or manipulation since the creation of the cell would be noticed.

The Tor Authentication Protocol (TAP) handshake, which is now considered obsolete, would meet the first two requirements, but since it does not make use of a MAC, the last requirement would not be fulfilled. An example for a handshake that satisfies all requirements is the *ntor* handshake, whose description is provided in section 3.1. In the context of onion service, several modifications of the handshake protocol are necessary. The modified protocol is shown in equation 3.9.

$$\begin{aligned}
 \text{secret_input} &= \text{EXP}(B, x) \mid \text{AUTH_KEY} \mid X \mid B \mid \text{PROTOID} \\
 \text{info} &= m_hsexpand \mid \text{subcredential} \\
 \text{hs_keys} &= \text{KDF}(\text{intro_secret_hs_input} \mid t_hsenc \mid \\
 &\quad \text{info}, S_KEY_LEN + MAC_LEN) \\
 \text{ENC_KEY} &= \text{hs_keys}[0 : S_KEY_LEN] \\
 \text{MAC_KEY} &= \text{hs_keys}[S_KEY_LEN : S_KEY_LEN + MAC_KEY_LEN]
 \end{aligned} \tag{3.9}$$

Initially, the client generates a curve25519 keypair consisting of the private key x and the public key X . The *ntor* onion key referred to as B in the regular *ntor* handshake is replaced by the public introduction point encryption key. Furthermore, the node id is substituted by

the *AUTH_KEY* generated previously by the onion service for the specific introduction point. The fields *t_hsend*, *t_hverify*, *t_hsmac* and *m_hsexpand* are text constants which are not discussed further. As the encrypted part of the *INTRODUCE1* cell, the client sends:

- CLIENT_PK
- ENCRYPTED_DATA
- MAC

where the *CLIENT_PK* is the public key X , *ENCRYPTED_DATA* corresponds to the original structure of the encrypted field, which contains information about the RP and is encrypted with the symmetric *ENC_KEY*. The *MAC* field contains an MAC of the entire cell using the computed *MAC_KEY*.

8. Relay Introduction Cell

The introduction point checks whether the authentication key matches an auth key of an active introduction circuit. If that is the case, it sends an *INTRODUCE2* cell with exactly the same content as the *INTRODUCE1* cell to the os and notifies the client of the successful referral with an *INTRODUCE_ACK* cell.

9. Join Rendezvous Point

The onion service receives the *INTRODUCE2* cell and checks if the *AUTH_KEY* or the *LEGACY_KEY_ID* matches the keys for this introduction circuit. If the key match, the next step is to verify whether the onion service has already received a cell with the same content. If this is the case, the cell was replayed and is silently dropped by the onion service. Otherwise, the service generates a curve25519 keypair y, Y , where y denotes the private key and Y represent the public key. The onion service computes:

$$\begin{aligned}
 \text{secret_input} &= \text{EXP}(X, b) \mid \text{AUTH_KEY} \mid X \mid B \mid \text{PROTOID} \\
 \text{info} &= m_hsexpand \mid \text{subcredential} \\
 \text{hs_keys} &= \text{KDF}(\text{intro_secret_hs_input} \mid t_hsenc \mid \\
 &\quad \text{info}, S_KEY_LEN + \text{MAC_LEN}) \\
 \text{ENC_KEY} &= \text{hs_keys}[0 : S_KEY_LEN] \\
 \text{MAC_KEY} &= \text{hs_keys}[S_KEY_LEN : S_KEY_LEN + \text{MAC_KEY_LEN}]
 \end{aligned}
 \tag{3.10}$$

The onion service is now able to decrypt the ciphertext using the *ENC_KEY* and check the authenticity and integrity of the cell using the *MAC_KEY*. The plaintext of the *ENCRYPTED* field is structured as listed below:

- RENDEZVOUS_COOKIE: Rendezvous cookie provided by the client.
- EXTENSIONS: Possible extension.

- `ONION_KEY_TYPE`: Type of the onion key (either TAP or ntor).
- `ONION_KEY_LEN` : Length of the onion key.
- `ONION_KEY`: Onion key that must be used when extending to the rendezvous point.
- `NSPEC`: Refers to the number of link specifiers included in the cell.
- `NSPEC times: LSTYPE, LSLEN, LSPEC`: Entry for each link specifier provided including its type and length.
- `PAD`: Optional padding .

In the next step the onion service finishes the ntor handshake by calculating:

$$\begin{aligned}
 \mathit{rend_secret_hs_input} &= \mathit{EXP}(X,y) \mid \mathit{EXP}(X,b) \mid \mathit{AUTH_KEY} \mid B \mid \\
 &\quad X \mid Y \mid \mathit{PROTOID} \\
 \mathit{NTOR_KEY_SEED} &= \mathit{MAC}(\mathit{rend_secret_hs_input}, t_hsenc) \\
 \mathit{verify} &= \mathit{MAC}(\mathit{rend_secret_hs_input}, t_hsverify) \\
 \mathit{auth_input} &= \mathit{verify} \mid \mathit{AUTH_KEY} \mid B \mid Y \mid X \mid \mathit{PROTOID} \mid \text{"Server"} \\
 \mathit{AUTH_INPUT_MAC} &= \mathit{MAC}(\mathit{auth_input}, t_hsmac)
 \end{aligned} \tag{3.11}$$

As mentioned for the client ntor handshake, t_hsenc , $t_hsverify$, t_hsmac and $m_hsexpand$ are text constants, which are not described further. The onion service generates a handshake reply that contains its generated public key (Y) and the $\mathit{AUTH_INPUT_MAC}$. Subsequently the onion service builds a circuit to the RP and sends a $\mathit{REDENZVOUS1}$ cell, which consists of the handshake reply in the form of $\mathit{HANDSHAKE_INFO}$ and the rendezvous cookie. If this cookie corresponds to a rendezvous cookie assigned to a circuit that is not yet connected, the RP connects the two circuits and sends a $\mathit{RENDEZVOUS2}$ cell to the client containing the $\mathit{HANDSHAKE_INFO}$. The client is now able to finish the ntor handshake as follows:

$$\begin{aligned}
 \mathit{rend_secret_hs_input} &= \mathit{EXP}(Y,x) \mid \mathit{EXP}(B,x) \mid \mathit{AUTH_KEY} \mid B \mid \\
 &\quad X \mid Y \mid \mathit{PROTOID} \\
 \mathit{NTOR_KEY_SEED} &= \mathit{MAC}(\mathit{ntor_secret_input}, t_hsenc) \\
 \mathit{verify} &= \mathit{MAC}(\mathit{ntor_secret_input}, t_hsverify) \\
 \mathit{auth_input} &= \mathit{verify} \mid \mathit{AUTH_KEY} \mid B \mid Y \mid X \mid \mathit{PROTOID} \mid \text{"Server"} \\
 \mathit{AUTH_INPUT_MAC} &= \mathit{MAC}(\mathit{auth_input}, t_hsmac)
 \end{aligned} \tag{3.12}$$

After the validation of the authenticity of the received data by comparing $\mathit{AUTH_INPUT_MAC}$ to the AUTH field received in the $\mathit{HANDSHAKE_INFO}$, both the client and the onion service are in possession of handshake output to derive shared keys that are used encrypt and authenticate data end-to-end between them.

3.2.3 Key Types

To ensure the integrity and authenticity of onion services, a variety of different keys are used, whose hierarchy is shown in Figure 3.10. The only key needed to identify and access an onion service is the public master identity key and its derived blinded public key. The remaining keys are used internally for authentication and encryption. The following section describes the use and origin of each key species.

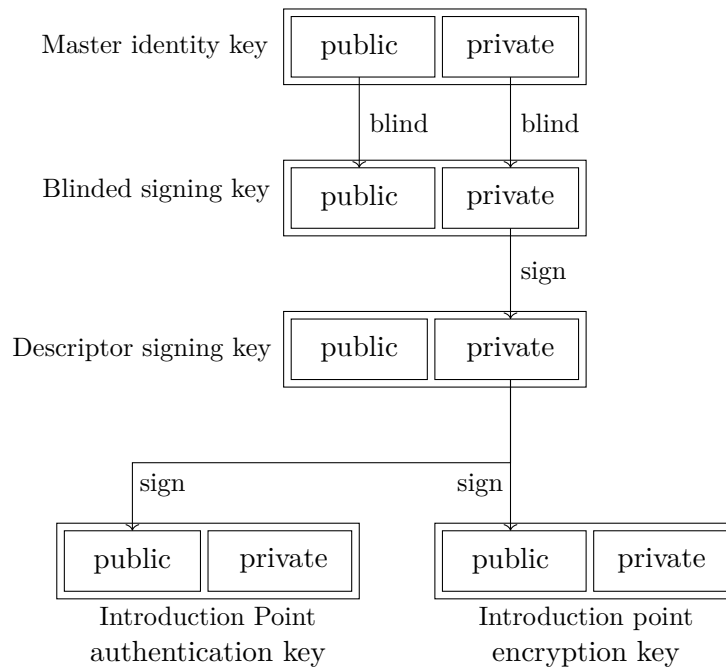


Figure 3.10: Onion Service Key Hierarchy

3.2.3.1 Master Identity Key

The long-term identity key-pair is used for the generation of the blinded signing keys and can be kept offline. As proposed in the Tor specification [29], the *Ed25519* key signature system based on the elliptic curve *Curve25519* is used.

3.2.3.2 Blinded Signing Key

For every time period the onion service uses a different blinded signing key to sign directory information. Given a nonce n , from the keypair (signing_key , public_key) a new keypair (signing_key_n , public_key_n) can be derived. Without the knowledge of public_key , the key public_key_n cannot be derived. Furthermore it is possible to check a signature signed with signing_key_n with public_key_n . The complete procedure for blinding a key works as follows:

Given is a master identity key pair (a, A) , where a denotes the private key, A refers to the public key and B is the base point of the elliptic curve as specified in [30], such that $A = aB$ and l is

the prime order of B . To derive a blinded key pair (a', A') for a nonce N and an optional secret s , several blinding factors are calculated:

$$\begin{aligned}
 BLIND_STRING &= \text{"Derive temporary signing key"} \mid INT_1(0) \\
 N &= \text{"key-blind"} \mid INT_8(\text{period-number}) \mid INT_8(\text{period-length}) \\
 B &= Ed25519 \text{ base point as specified in [30]} \\
 h &= (BLIND_STRING \mid A \mid s \mid B \mid N)
 \end{aligned}
 \tag{3.13}$$

Additionally, according to the specification of *Ed25519*, some factors of h are clamped to specific values. Using the calculated blinding factors, the blinded private key a' for a period is calculated as:

$$\begin{aligned}
 a' &= h a \text{ mod } l \\
 RH' &= SHA-512(RH_BLIND_STRING \mid RH)[:32] \\
 RH_BLIND_STRING &= \text{"Derive temporary signing key hash input"}
 \end{aligned}
 \tag{3.14}$$

The blinded public key A' is calculated based on the blinded private key as follows:

$$A' = h A = (ha)B \tag{3.15}$$

3.2.3.3 Descriptor Signing Key

The descriptor signing key is used to sign the onion service descriptors. In contrast to blinded signing keys, these keys must be stored online at the onion service host. The key itself is signed using the blinded signing key. The descriptor signing key pair is a newly generated *X25519* key pair and is not dependent on any other key used in this cryptographic scheme.

3.2.3.4 Introduction Point Authentication Key

The onion service generates a unique *X25519* key pair for each introduction point it establishes. When a client connects to an introduction point in order to contact the onion service, the introduction point verifies that it maintains a circuit to the service corresponding with that introduction point authentication key. This allows the client to ensure that it is connected to the correct introduction point and the onion service to check whether the contact of the client took place through a known introduction point. The public part of the key is included in the signing-key extension of the onion service descriptors and is cross-certified by the descriptor signing key.

3.2.3.5 Introduction Point Encryption Key

The introduction point encryption key is used to encrypt information about the rendezvous point chosen by the client in the introduction protocol. Furthermore, the key is used as the public key of the onion service in case the advanced ntor handshake is used to exchange shared keys material. The *X25519* key pair is generated by the onion service for every introduction point and is cross-certified by the descriptor signing key.

3.2.3.6 Descriptor Encryption Keys

The key for encrypt hidden service descriptors is derived from secret data provided by the *SECRET_DATA* parameter and a string constant given by the parameter *STRING_CONSTANT* as follows:

$$\begin{aligned}
 SALT &= 16 \text{ bytes from } H(\text{random}) \\
 secret_input &= SECRET_DATA \mid subcredential \mid INT_8(revision_counter) \\
 keys &= KDF(secret_input \mid salt \mid STRING_CONSTANT, S_KEY_LEN \\
 &\quad + S_IV_LEN + MAC_KEY_LEN) \\
 SECRET_KEY &= \text{first } S_KEY_LEN \text{ bytes of keys} \\
 SECRET_IV &= \text{next } S_IV_LEN \text{ bytes of keys} \\
 MAC_KEY &= \text{last } MAC_KEY_LEN \text{ bytes of keys}
 \end{aligned} \tag{3.16}$$

3.2.4 Time Periods and Shared Random Values

As mentioned in 3.2.1, onion services maintain two descriptors at any time to be reachable by clients that use the current consensus or the previous one. The descriptors are periodically published to the responsible HSDirs. After each publication of a descriptor, a random timer between 60 and 120 minutes is set. When the timer expires, the service uploads its descriptors again to the HSDirs responsible for that *Time Period* (TP) and *Shared Random Value* (SRV). Each consensus contains the SRV for the current and the previous TP. The procedure for choosing the right SRV along with the TP differs between client and service. The differences can best be explained by using an example, which is illustrated in Figure 3.11. *TP#X* indicates, that from this point on a new TP *X* begins. The symbol \$ indicates a descriptor rotation. *SRV#X* means that a new SRV *X* will be available from that point.




Figure 3.11: Tor Hidden Service Key Hierarchy [25]

The two descriptors uploaded by the service are referred to as *first* and *second* service descriptors and follow different upload logics and purposes:

- **First Descriptor:**

The first descriptor is intended for clients that have a consensus that is still in the previous TP. To create this descriptor, the previous TP and SRV is used when the service is in the time segment between a new TP and a new SRV (marked with '-' in Figure 3.11). For example, if the service is at 14:00, shortly after *TP#1*, it would still upload the descriptor with *TP#0* and *SRV#0*. If the service is in the time segment between a new SRV and a new time period (marked with '=' in Figure 3.11) it uses the current TP and the previous SRV. For instance,

if the service is at 02:00, shortly after $SRV\#2$, it would upload the descriptor with $TP\#1$ and $SRV\#1$.

- **Second Descriptor:**

The second descriptor is intended for clients who have an up-to-date consensus in the same TP as the service. In order to create this descriptor, the current TP and the current SRV are always used to upload the descriptor if the service is in the time segment between a new TP and a new SRV (marked with '=' in Figure 3.11). For example, if the service is at 14:00, shortly after $TP\#1$, it would upload the descriptor with $TP\#1$ and $SRV\#1$. If the service is in the time segment between a new SRV and a new TP (marked with '=' in Figure 3.11) it uses the next time period and the current SRV. For instance, if the service is at 02:00, shortly after $SRV\#2$, it would upload the descriptor with $TP\#2$ and $SRV\#2$.

If services receive a consensus with a *valid_after* time past the next SRV calculation time, they rotate their descriptors by discarding the first descriptors, pushing the second descriptors to the first descriptors and rebuilding the second descriptors. The client aims to synchronize TPs and SRV when fetching descriptors, so it always tries to use $TP\#X$ with $SRV\#X$.

3.2.5 Security Enhancements

Tor in its original implementation was vulnerable to traffic correlation attacks. Within these attacks, adversaries attempt to de-anonymise users by monitoring both the entrance of user traffic to the anonymity network and its exit [31]. In order to do this, the adversary must control both the first node and the last node in a circuit. If it is able to do so, it knows both the client (because it is directly connected to it) and the resource that the client uses. This type of attack on the Tor network was first demonstrated by Oeverlier et al. [16] in the course of the de-anonymisation of onion servers [31]. Based on ideas of Wright et al. [32] they introduced the idea of fixing the first nodes in a circuit to prevent adversaries from gaining control over them. These first nodes are referred to as *entry guards* or *guard nodes*. An illustration of a three-hop circuit utilizing this guard nodes is provided in Figure 3.12.



Figure 3.12: Security Enhanced Onion Routing Network [23]

An enhanced circuit consists of the client, in this example referred to as initiator, an entry guard as first hop, a *middle relay* as second hop and an *exit relay* as last hop. The only node that is directly connected to the initiator is the entry guard. The middle relay knows neither initiator nor responder. At the exit relays, traffic exits the Tor network. These are the only nodes that are directly connected to the responder and thus know their location. The entry guards are chosen from a small subset of all relays in the consensus, which are in possession of the *Guard* flag. This flag indicates,

that the node is stable, fast and has a specific minimal uptime [33]. This subset is further filtered based on several performance and utility parameters to obtain a set of *primary guards*. Currently, three primary guards are selected, nevertheless it is attempted to use the same guard each time [34]. To keep the probability of success of a traffic correlation attack as low as possible, the entry guard must be persistent. Currently, the lifetime of a guard is 120 days [34]. The middle relay of a circuit is chosen randomly from the relays available in the latest consensus weighted by bandwidth. Similarly, the exit relays are selected, except that their relay must support an appropriate exit policy [34].

Since this makes it very complex to become the guard node of a service, *guard discovery* attacks aim to identify the entry guard of an onion service. Once the entry guard is known to an adversary, an attempt is made to compromise it using other methods. To complicate this, the *Vanguards* extension was developed, which is written in Python. It utilizes the Tor control protocol and consists of three components: [35]

1. **Vanguards:**

This is the core function of Vanguards. In addition to the entry guards, also middle relays and exit relays are pinned for a longer period of time. With this, the time to discover an entry guard increases from five minutes to weeks or months, even if the attacker is able to control a significant portion of the Tor network. To protect the newly pinned nodes, the rendezvous circuits, directory circuits, and introduction circuits are lengthened. [36]

2. **Bandguards:**

The Bandguard component checks whether there are traces of bandwidth side channel attacks within circuits, that are used by adversaries to amplify traffic analysis attacks. If such traces are found, an alarm is generated and the affected circuit is optionally closed.

3. **Rendguards:**

This component tracks the frequency of rendezvous point used on the onion service side. Since some attacks rely on the aid of custom malicious rendezvous point for traffic analysis, an overuse of certain nodes indicates a possible malevolent utilization. If certain threshold values are exceeded, an alarm is generated and the rendezvous circuit is optionally torn down.

4 Unlinkable Onion Services

“The unlinkability of two or more items of interest from an attacker’s perspective means that within the system (comprising these and possibly other items), the attacker cannot sufficiently distinguish whether these items of interest are related or not.” [37]

Following this definition, three specific requirements for unlinkable onion services emerge:

1. Recurrent connections to the same onion services must not be linked to a client.
2. Multiple connections of a client to different onion services must not be relatable.
3. Connections of multiple clients to specific onion services must not be associable.

Considering the functional behaviour of onion services in version three, these requirements are only partially met. An adversary controlling one of the Hidden Service Directories (HSDirs) of the onion service cannot identify which users want to access the service, since the Tor clients are connected to the Rendezvous Point (RP) of the service using anonymous circuits, but it can draw conclusions about the number of users and the frequency of access to the service. This can become detrimental for users in case an attacker additionally has access to further sources of information, e.g. the network traffic of the client. Through the correlation of multiple information sources based on the chronological sequence of traffic, the user could be identified and associated with the service. However, for a targeted attack, an adversary would need to know the address of the onion service. With onion services in the previous version two, it was possible to collect addresses on the HSDir, because the required information to determine the address of an onion service was stored unencrypted in the descriptor. This weakness was fixed in onion service version three by securing the descriptor with a multi-layered encryption scheme. Nevertheless, if the onion address is known to the adversary, an analysis of the usage pattern of this service is possible.

Another threat to the unlinkability of onion services is constituted by the monitoring of the host server of the onion service, which typically is a web server or an application server. Theoretically, the usage of an anonymous circuit between the RP and the onion service protects the location of this server, however, it has been shown that revealing the location of the onion service is possible for instance by means of protocol-level attacks [22] [21] [18]. These attacks enable the determination of whether a relay controlled by an adversary is the guard relay of an onion service or the guard relay is adjacent to the controlled relay. If the adversary in fact controls the guard node, it can monitor the entire traffic between the onion service and its clients and can thus find out how many connections take place and when. Additionally, the adversary would know about the IP address of the onion service. If the guard relay is adjacent to the relay controlled by the adversary, the identity of the guard node is known and the relay can be compromised, surveilled or coerced [35]. Therefore, the Tor project announced these so-called *guard discovery attacks* to be among the most dangerous threats to onion service as of 2018 [38]. The Vanguard add-on was developed to address

this problem [35], but it is not currently used in the default configuration. Similarly to the monitoring of the HSDir, the adversary needs to know the address of a service to carry out such attacks.

As in both of the above described threats to unlinkability, the identity (i.e. the onion address) of a service plays a key-role, the obvious prevention consists in the decoupling of the onion service and its identity. The following sections provide detailed information on the architecture and efficient implementation of such a decoupled approach.

4.1 Methodological Approach

Uncoupling an onion service from its identity means that the service needs to hand out different custom onion addresses to every client that wants to access it. In order to accomplish that, it is necessary to change the public master key of the service as, beside the version, this key is the only identifier of the service that is encoded in its address [39, 25]. The amendment of the public master key of a onion service can be performed using two different approaches, partly depending on the version of the onion service:

1. Multiple Onion Services

The straight forward way to change the identity of a service is to provision a fresh onion service pointing to the same application. Technically this doesn't correspond to a change of the key, but produces a similar result. The issue with this solution is scalability. In order to provision a new onion service, new keys have to be generated, three additional connections to Introduction Points (IPs) have to be established and maintained and 16 descriptors have to be created and uploaded to the HSDirs. This can best be illustrated by the example of the social media network Facebook, which in 2016 announced, that up to 1.000.000 clients utilized the onion service version of Facebook in a period of 30 days [40]. If each user would be provided with a custom service address, 1.000.000 unique onion services would have to be provisioned. In comparison, the number of unique V2 onion service has never exceeded 240.000 [41].

2. Multiple Descriptors

Rather than provisioning a separate onion service per user, a new descriptor is created for each participant pointing to the same service. The idea of creating custom service descriptors is not new, as it was introduced in the onion service V2 protocol as *stealth* mode for client authorization to hide the activity of the service from unauthorized users. For this purpose, the service generates an asymmetric *client key* and a symmetric *descriptor cookie* per user. After their creation, the service communicates these two values to the authorized user in form of a contact information string. The client key replaces the public master key (referred to as permanent-key in V2), which causes the onion address to change since this is the *base32*-encoded hash value of this key. The descriptor cookie prevents the prediction of the storage location of the descriptor on the HSDir, as this would require the knowledge of the descriptor. Additionally, the descriptor cookie is used to encrypt the IPs of the service, so that only the authorized user can contact the service. This approach is a clear advantage in terms of

unlinkability. The onion service now has full control over the dissemination of its descriptors and thus over its accessibility. If the service stops publishing descriptors for a certain user and said user is not able to link descriptors issued for other users to the service, it will no longer get any information on the activities of the onion service. Moreover, the service could theoretically only be found by authorized users, as only they know the onion address of the service. If a non-authorized user nevertheless succeeds in obtaining one of the onion addresses of the service (e.g. by monitoring HSDirs), it cannot use it to contact the service, as the IPs to the service are encrypted. However, this concept does not only have advantages. As the number of descriptors increases, the load on both the onion services, which have to manage and regularly renew these descriptors and the HSDirs, grows. In order to deal with this problem, the maximum number of descriptors an onion service can manage in V2 has been set to 16. [39]

With onion service V3, among other innovation, a novel structure for service descriptors was implemented, which is protected with a double-layered encryption. This prevents adversaries from obtaining onion addresses by monitoring HSDirs. With this new structure and the concept of key blinding, which is described in section 3.2.3.2, the stealth client authorization is no longer needed, since an adversary is not able to find out the onion address of a service, if it is not published. For this reason, the creation of new onion addresses using the client authorization in V3 is no longer possible.

As onion services V2 will no longer be supported in the foreseeable future [15] and onion services V3 don't support stealth client authorization, creating multiple descriptors to generate custom onion addresses and thus provide unlinkability is not feasible at the moment. Therefore, the first prototype implementation of unlinkable onion services will create a custom onion service per user, which implies that the scaling problem still remains. A potential solution could be to not provision a custom service for each user, but instead manage a pool of custom onion addresses and randomly hand out onion addresses from this pool. However, due to time restrictions the solution for this problem is not part of this thesis and is therefore considered future work.

In the next step it must be determined how the generated addresses can be exchanged in the most effective way. In principle, a distinction can be drawn between two approaches, which differ in terms of the size of the expected user pool of a service. These two approaches are described in the following two sections.

4.2 Custom Addressing

Services that have a rather small user pool can take advantage of a simple implementation of unlinkability referred to as custom addressing, whose purpose is the provisioning of a per-client ephemeral onion service and the exchange of the resulting onion address. For the sake of simplicity, the application server, which provides the actual service, and the onion service in combination are referred to as server. The idea of this approach is that the server is aware of all its communication partners due to the limited number of users. The service provisions an onion service for every user it wants to communicate with and sends the custom onion addresses of these services to the respective

clients using a secure channel. As no communication from the client to the server is required, only uni-directional communication between server and client is intended. The sequence diagram of the custom addressing procedure is depicted in Figure 4.1.

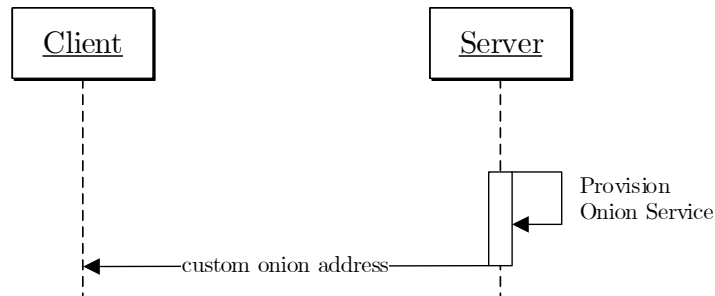


Figure 4.1: Custom Addressing

The major advantage of this approach is its simplicity. No additional complexity is added through the use of a protocol. If the client only uses the provided custom onion address to contact the server, unlinkability is ensured. Furthermore, a server would have full control over its own accessibility since if a server want to be accessed by a client, it simply tears down its custom onion service. In case this client is not able to obtain the custom onion address of another client, it will no longer receive any information about the onion service and its activity. The major disadvantage of this approach is constituted in its scalability. Although there is no theoretical limit of onion services a tor instance can create [25], both the HSDirs and the Tor client are not designed for operating hundreds of thousands of onion services, as for instance the Facebook onion service would require [40].

Furthermore, this approach is based on the assumption that the first contact originates from the server, which excludes any new clients not yet known to the server. This means that a client must be aware of the communication channel for receiving its custom address previous to the first contact. The communication channel cannot utilize the Tor network, as the server has no point of contact with the client, which implies that the server would need to know the identity (e.g. the IP address) of the client. This eliminates the concept of anonymity. An adversary that operates on the Internet Access Provider (ISP) tier would be able to monitor the initial communication and link the client to the server. A remedy could be the use of an anonymous communication service such as *ProtonMail* [42] for the initial communication between the server and the client. ProtonMail uses public-key cryptography to ensure the confidentiality of messages sent between users. For this to work, the server would need to be in possession of a public key of all its users. Due to the restricted user pool, this is feasible. However, the use of communication based on public-key cryptography raises another problem. Both the server and the client must be able to verify the authenticity of the key material used. This could be solved, for example, by using certificates, but this in turn would undermine the principle of anonymity, since the client would need to authenticate itself again a central Certificate Authority (CA) in order to obtain the certificate. To simplify the exchange of custom onion addresses, the *Unlinkable onion Service Protocol* (UOSP) was designed.

4.3 Unlinkable Onion Service Protocol

4.3.1 Concept

As exemplified in section 4.2, the exchange of custom onion addresses between server and client is not a trivial problem. For the sake of usability, accessing an unlinkable service must not differ drastically from accessing a normal onion service from a client's perspective. This implies, that the client must be able to retrieve a custom onion address for a service without additional knowledge and without the need of using a separate secure communication channel. For this purpose, a standardized approach in the form of a protocol must be defined, which specifies the exact procedure from the first request of the client to the receipt of the custom onion address. With the *Unlinkable Onion Service Protocol* (UOSP), such a standardized approach was designed and is described in this section in detail.

The basic conceptual difference of the UOSP compared to *Custom Addressing* consists in the types of addresses a clients uses to contact a service. In a regular onion service access, the service's onion address is used for the initial request and then data is exchanged between client and service. In the UOSP, the address used for exchanging data between client and server is referred to as *data address*. In addition, the UOSP uses a second type of address, whose only purpose is the initial contact of a client to a service and the exchange of each client's custom data onion address. Therefore, this address type is called *introduction address*. Beside their different purposes, the two address types also have different levels of confidentiality. As the introduction address is needed in order to establish an initial contact to a service, this address needs to be publicly available. It could for instance be posted in internet forums or handed out to the clients who should be able to reach a certain service. The data address on the other hand, is personal to a specific client and needs to be kept secret in order to remain unlinkable. If the address were to be published, adversaries could try to manipulate the service's Hidden Service Directories (HSDirs) and thereby gain information about the utilization of the service. This would lead to linkability and render the reason for using the UOSP obsolete.

For this reason, in order to ensure the unlinkability of onion services, the following environmental assumptions are considered to be met:

1. The introduction address of an onion service is published in a way that all clients that need to communicate with the service can obtain it. The address does not have to be kept secret.
2. The data address of an onion service is private to a single client or a client group and is therefore not published in a way that other clients or other client groups that are not authorized can obtain it.

In this context *authorized* means that the onion service has designated the client or the client group to use its data address and is not associated with a cryptographic meaning. Authorization in the cryptographic sense would require that each authorized client obtains information in advance (e.g. a key pair) with which it can unambiguously prove that it is authorized to use a service. However, this would lead to the same implications as with Custom Addressing, that are described in section 4.2.

4.3.2 Architecture

Since each unlinkable onion service needs to provide at least one public introduction address and probable multiple private data addresses, it must provision services for responding to requests to introduction addresses as well as for responding to requests to data addresses. Both of these services must be implemented as onion services, otherwise their location (i.e. the IP address of the service) cannot be hidden. The minimal set of components needed for the provisioning of an unlinkable onion service is illustrated in Figure 4.2.




Figure 4.2: Minimal Set of Components Required for Unlinkable Onion Services

The Tor client of the users acts as its onion proxy, which establishes a connection to the Tor network. The introduction service is hosted as a single onion service, whose address represents the UOSP introduction address. Although the introduction service and the data service are illustrated as separate servers, both services must be hosted on a single server because the introduction service must be in control of the provisioning of data services. The application to which the data services point does not necessarily have to be operated on the same server. The basic procedure of a UOSP request is structured as follows:

1. The user obtains the introduction address of an unlinkable onion service through out-of-band-communication (e.g. through an email) or other communication platforms. It is assumed that the creator of the introduction address is trusted.
2. The user then sends a request to the introduction address. The protocol utilized for this request is not limited to web protocols, in principle any TCP-based protocol can be used. Moreover, the environment with which the request is submitted is not specified; it could be a browser, for example. It has to be noted that this request is linkable since a publicly known introduction address of a service is used for this purpose.
3. The request reaches the introduction service, where it is interpreted. At this point, authorization mechanisms such as the client authorization embedded in onion service could be used, which decide whether the service communicates with the user or not. Subsequently, the service provisions a custom onion address and responds to the user.

4. The user checks the well-formedness of the received custom onion address. This can easily be achieved, since the length of a valid onion address is defined in the specification [25]. The obtained custom onion address must be memorized for further requests, as each request would otherwise result in a request to the introduction address, which in turn would lead to linkability.
5. At this point, the user can submit a request to the custom onion address received. The data service responsible for that custom onion address will respond with the requested data.

As indicated in the basic sequence, the user has to perform a number of tasks by himself. This includes the manual query of the custom onion address, the check of its well-formedness and the memorization of previously queried addresses. This might be feasible for technically proficient persons, but is rather unusable for normal users. For this reason, apart from client and server, a third main component in the form of a proxy is introduced. This proxy handles the tasks that previously had to be performed by the user and must be placed between the user and the Tor client, since the data traffic between the Tor client and the Tor network is already encrypted. For reasons of usability, requirements are placed on the transparency and performance of the proxy. If the latency introduced by the proxy is too high, users will not use it [5]. The situation is similar in regard to transparency. If the user experiences a significant increase in the complexity of using a service through the use of the proxy, it will not be used.

4.3.3 Identity Assignment

Essentially, there are two approaches for assigning individual onion addresses to users. These approaches and their advantages and disadvantages are described below.

- **Key Derivation:**

With this method, a key exchange between proxy and server is initially performed. After the key exchange, both parties are in possession of the same symmetric key, from which further keys can be derived. These derived keys can be used to provision onion services, for which only the client and the server know the address, as its partly based on the symmetric key. To implement this approach, a two-phased protocol was considered in reference to the phase-based mode of operation of Transport Layer Security (TLS) [43]:

1. Key Exchange Phase:

This phase serves the establishment of a shared symmetric master key, which could be implemented using an asymmetric key exchange protocol, such as the Diffie-Hellman key exchange. The key exchange protocol requires a number of domain parameters to be agreed upon in advance.

2. Usage Phase:

In the second phase, both parties can derive keys from the shared symmetric master key established in phase one. For the key derivation, a double ratchet algorithm is used, which offers future secrecy. The derived key is used to establish an onion service on

the server side, which can only be accessed by the client due to the knowledge of the symmetric key used.

This approach offers the possibility of generating ephemeral onion services, which can be changed at a defined interval by deriving a new key to avoid a correlation of service usage. Furthermore, it provides an additional mechanism for access control. If the server tears down the previous custom onion service of a client each time it derives a new temporary key and provisions new custom onion service, an adversary that gets into possession of that temporary key can only monitor the service that is bound to it. The adversary can not obtain the temporary key of a new time interval, neither can he obtain the temporary key of a previous time interval due to future secrecy.

Nevertheless, this approach does have a number of implications. As mentioned in the description of phase one, both the client and the server must agree on domain parameters in advance. This goal can be achieved by specifying the parameters in the protocol definition since they must not be kept secret. A more serious issue lies in the synchronization. Both participants must synchronize the derivation of new temporary keys, otherwise the client might use an address for a service that is no longer available or the server might offer a service that is no longer used. This synchronization is partly accomplished by specifying a time interval for the derivation of new temporary keys in the protocol definition. The interval by itself does not solve the issue. Fixed time points that define the start of an interval are required, which have to be identical for both systems. Furthermore, additional synchronization messages are needed in case one participant is offline during the beginning of a new interval.

- **Random Assignment:**

An alternative approach is the randomized assignment of custom onion addresses. This is based on the premise that when a new onion service is provisioned, a new master keypair is created for each service, which cannot be associated with the keypair of other services. In the context of the UOSP this means that each time a new data address is requested, a dedicated onion service is provisioned, whose onion address is only dependent on the master identity public key and the version of the onion service and cannot be linked with previously created data addresses or to the introduction address of a service. This onion address is used as data address and can be considered random, as there are, beside the version, no parameters that influence its generation.

This requires neither key exchange nor key derivation or synchronization, which reduces the complexity of the protocol since the assignment, which can be reduced to an onion service provisioning, is entirely handled by Tor.

Both options offer similar properties in terms of unlinkability. In both approaches, data addresses cannot be linked to other data addresses or introduction addresses. The *Key derivation* approach has the advantage that new data addresses can be generated without re-contacting the introduction address, but it has a significantly higher inherent complexity. As the requirements for unlinkability are met by both approaches, the *Random Assignment* is used due to its simplicity and lightweights.

4.3.4 Protocol Definition

The individual components of the protocol, as well as the basic procedure for the creation of unlinkable onion services have been described in the previous sections. In this section, a formal definition of the protocol will be presented. According to Holzmann et al. [44] a protocol specification consists of the five following components:

- “The service to be provided by the protocol”
- “The assumptions about the environment in which the protocol is executed”
- “The vocabulary of messages used to implement the protocol”
- “The encoding (format) of each message in the vocabulary”
- “The procedure rules guarding the consistency of message exchanges”

In the course of this section these components are defined for the UOSP.

- **Service Specification**

The purpose of this protocol is the provisioning of one or multiple per-client ephemeral onion services, whose addresses are referred to as data addresses. The protocol is intended for the use in a client-server architecture. A proxy arranged between client and server abstracts the complexity of the protocol from the user. The initial connection from client to proxy takes place unidirectionally using a fixed onion address referred to as introduction address, whose sole purpose is the initial communication between proxy and server. No payload data must be exchanged over this connection. The proxy maintains a cache to store the results of previous requests to the introduction addresses of services. This is mandatory in order to ensure unlinkability for future requests. Each entry in the cache consists of the issued data address, its validity and the introduction address that was used to issue this data address. Once the proxy receives a request from the client, it first checks whether there is a valid data address in the cache for the introduction address in the request. If this is the case, the valid data address is structured in a *UOSP Client Response* and is sent back to the client.

If the proxy does not hold a cached address, it extracts the required data from the request and generates a *UOSP Request*, which it then forwards to the server. The server provisions a new ephemeral onion service and responds to the proxy with the public key of the newly provisioned onion service encoded as its onion address. Additionally, the validity of the generated address is transferred as metadata. This validity set by the server controls how long a service can be accessed via its address. After the validity period has expired, the server tears down the onion service, which is thus no longer accessible. The proxy confirms the reception of the public key and its metadata by sending an acknowledgement message in form of a *UOSP Ack* to the server. Attached to this acknowledgement is a hash of the previously received public key. The server in turn hashes the public key of the previously provisioned onion service and is now able to compare both values.

If these values do not match, the integrity of the public key sent cannot be verified and it can be concluded that an error has occurred either in the transmission path from server to proxy or from proxy to server. To alert the proxy about this error, the server sends a *UOSP Error* to the proxy and tears down the previously provisioned service. The proxy relays this error to the client. If the two hash values match, the integrity of the public key is assured. To end the protocol cycle, the server sends a finalization note in the form of a *UOSP Fin* message to the proxy. The proxy caches the data address for future requests and sends a *UOSP Client Response* to the client. This response can either contain textual information about the newly provisioned data address or a protocol-specific redirect.

- **Environmental Assumptions**

The environment in which the protocol is executed consists of a client, a proxy, a server and a transmission channel. It is assumed that the proxy is local to the client, which implies that the client's onion proxy and the proxy are located on the same system. Nevertheless, the operation of a remote proxy would be possible, but would not serve its purpose in terms of unlinkability, since information on the client's service usage could no longer be managed locally. As this protocol operates on the application layer in the Open Systems Interconnection model (OSI model), segmentation, flow control, congestion control and transport of protocol messages is performed by the underlying layers. This underlying layers assure that packets arrive at their destination in the correct order without data corruption. It is assumed that all components are online during a protocol cycle. Any archiving of protocol messages for later transmission is not considered. Both the proxy and the server must have a mutual consent on the time zone, which is used for temporal functions that are dependent on time.

- **Protocol Vocabulary**

The protocol vocabulary specifies eight distinct types of messages:

- **UOSP Request:**

Requests the provisioning of a new temporary data address. The required introduction address is extracted from the UOSP client request.

- **UOSP Response:**

Response to a UOSP request. The payload of this response contains the encoded public key of the newly created data service, as well as metadata.

- **UOSP Client Request:**

Request from the client to the proxy. This requests must contain a valid destination address. The destination address can be an arbitrary address including introduction addresses and data addresses.

- **UOSP Client Response:**

Response to a UOSP client request that includes the data address of the newly provisioned data service. This response can either take a textual form or be implemented as a protocol-specific redirection.

- **UOSP Ack:**

Acknowledgement of the reception of a UOSP response, which includes the hash of the encoded public key received with the UOSP response.

- **UOSP Fin:**

Indicates the successful completion of a protocol cycle.

- **UOSP Error:**

Signals a corruption of the integrity of the transmitted encoded public key on the transmission path between proxy and server.

- **Message Format**

Every UOSP message contains at least the UOSP specifier, which indicates whether the message is UOSP related or not and the UOSP version, which currently is 0.1. This basic message structure is referred to as standard message format. Dependent on the protocol message sent, the extended validity UOSP message format or the extended data UOSP message format is used. The extended validity format includes the validity of a data address transferred through the message, whereas the extended data format contains an indicator of whether data is transferred over a UOSP connection or not. The menagerie of message formats is depicted in Figure 4.3. The optional payload in all three message formats is stored in the data field.

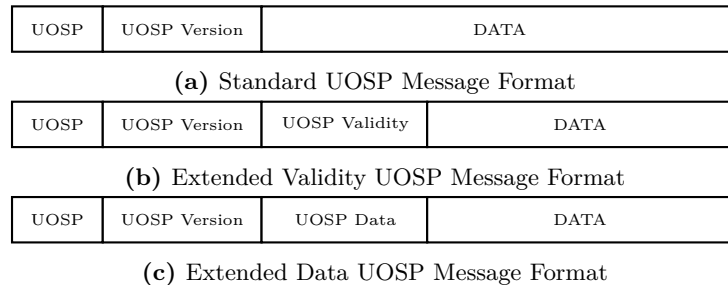


Figure 4.3: UOSP Message Formats

The standard message format is used by the UOSP Request, UOSP Ack, UOSP Fin and UOSP Error. The extended validity format is used in UOSP responses. The UOSP client request and UOSP client response are not specified, as the client request is dependent on the protocol used (e.g. HTTP) and the client response should be adjustable to the users' environment. The extended data UOSP format is solely used for data transfer using data addresses.

- **Procedure Rules**

1. The resource a client wants to access must be provided by an onion service and thus have an onion address.
2. The proxy must store previous requests to an introduction address and the responses containing the data addresses for these requests.

3. If the client requests an introduction address for which a data address has been provisioned previously, the cached address must be used. If the proxy does not hold a cached address, it must send a request to the server.
4. The cache of the proxy must delete entries that are invalid. This validity check must be carried out either periodically or with each request.
5. If the server receives a valid request for an introduction address associated with the server, it must provision a new onion service with its associated data address.
6. The response to a valid UOSP request must contain a valid data address for a temporary data service.
7. Every UOSP response must be acknowledged. This acknowledgement is used to ensure that the response has been received. To each acknowledgement the hash value of the data address must be appended in order to assure the integrity of the received address.
8. If the integrity check fails, the server must alarm the other participants by sending an error message.
9. The server must maintain all provisioned data services within their validity period.
10. If services exceed their validity periods, they must be teared down by the server. It is therefore mandatory that also the server maintains a cache storage for provisioned services and their validity.
11. The protocol ends by either a UOSP Fin or a UOSP Error message.

A complete overview of the UOSP can be obtained in the Unified Modeling Language (UML) sequence diagram in Figure 4.4.

4.4 Proof of Concept

To demonstrate the feasibility of the unlinkable onion service protocol proposed in section 4.3, a Proof of Concept (PoC) is implemented, which consists of a proxy module and a server module. In principle, the protocol is applicable for all TCP-based communication, however, for this implementation, the Hypertext Transfer Protocol (HTTP) was chosen, as this is apart from Skynet (botnet) the most common used protocol for onion services [45].

In the subsequent section, the technical components used for both the proxy and the server module are delineated. The implementations of the server and proxy module are described in detail in separate sections.

4.4.1 Technical Building Blocks

4.4.1.1 Tor Stem

Both the proxy and the server module are implemented with *Python 3.7*, due to the availability of the Tor *Stem* library [46]. Stem is a controller library written in Python, that enables communication with a Tor binary using the Tor control protocol. The library is composed of three different API modules at its core [46]:

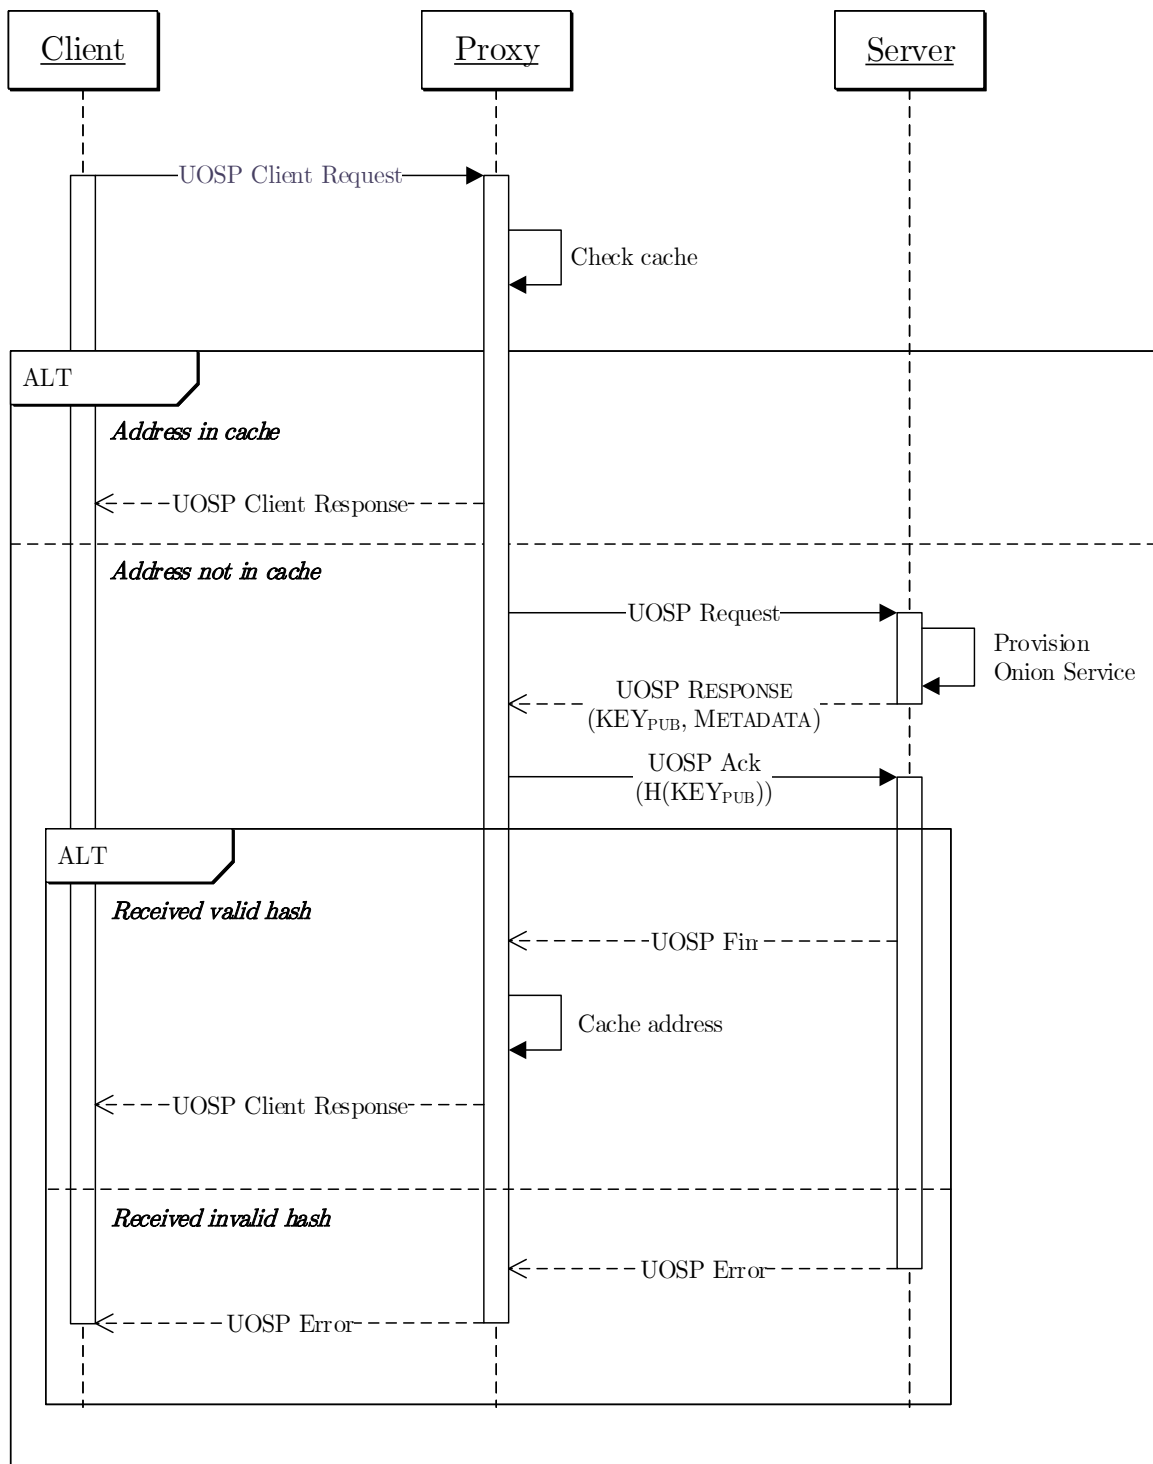


Figure 4.4: Unlinkable Onion Service Protocol Sequence

- **Controller:** The controller module enables authenticated communication with the Tor control socket and abstracts low-level connection details.
- **Descriptor:** The descriptor module can be used to retrieve, manage and generate a variety of data structures used in Tor. These include server descriptors, micro descriptors, network status documents, directories information and many more.
- **Utilities:** Helper functions of the two modules controller and descriptor are bundled in the utilities module. These include functions for retrieving system information, handlers for the textual configuration files, helper functions related to IP communication, string tools and functions for working with the underlying system.

In the course of this thesis, only the Controller module and its helper functions from the utilities module are used.

4.4.1.2 Socket Server

To implement the communication between proxy and server, the Python module *socketserver* is used. This module is a framework for creating network servers, which provides classes for the synchronous processing of network requests for different protocols and connection types. The processing of a request is thereby performed by two different classes. The *server* classes take care of the low-level communication tasks like listening on a socket or accepting connections. The protocol-specific processing, is performed by the *handler* classes. Their tasks include, for example, processing and interpreting incoming data and sending responses back to the requester. The advantage of dividing the responsibility for processing a request between communication-specific and protocol-specific classes is their re-usability. This re-usability enables users to utilize existing server classes for basic communication (e.g. TCP) without major changes and to adapt the protocol-specific handler to the respective application. [47]

A socket server is constructed by instantiating an object of a server class. Thereby, the address to be listened to by the server and a *request handler* class, which is called when a request is processed, must be passed as parameters. All request handler classes are based on the *BaseRequestHandler* super-class, which defines three interfaces, which must be implemented by handler classes inheriting from it: [47, 48]

- **Setup:** The setup method is called prior to the handle method. It prepares the request handler for the request by performing all initial actions required.
- **Handle:** The handle method carries out all the real work required to handle a request. This contains parsing the request, processing it and sending a response.
- **Finish:** The finish method is called after the handle method completed and performs a clean-up of anything created during the setup.

For both stream-based and datagram-based traffic, there are pre-defined base classes that inherit from the *BaseRequestHandler* and have already implemented the setup and finish methods so that only the actual request handling must be supplemented. As mentioned before, requests are processed synchronously. This means that a request must be completed before the next request can be processed, which implies that parallel processing of requests is not possible. To overcome this limitation, mix-in classes can be used to enable asynchronous processing in the form of a separate thread per request (*ThreadingMixIn*), or a separate process per request (*ForkingMixIn*). [47, 48]

4.4.1.3 Caching Database

As stated in the protocol definition in section 4.3.4, the data address of a newly provisioned temporary data service must be stored on both the server and the proxy. To avoid complications through inter thread communication or inter process communication, a central database is used to store the addresses. This database should generate as little overhead as possible. Therefore, renowned database management systems such as *MySQL* or *PostgreSQL* are excluded, as only a fraction of the functions offered are required.

The alternative lightweight database approach used for this prototype is the Zope Object Database (ZODB) [49]. As the name suggests, this is an object database management system, which was originally developed for the *Zope* web application server. ZODB is entirely written in Python and can therefore be seamlessly integrated in the code, as no separate language for database operations is used. Furthermore, no object-rational mapping is necessary, since the database can directly persist Python objects. This persistence does not have to be manually triggered by the user, nor any code is needed to read or write objects. The persistent objects are attached to a container, which works similarly to a Python dictionary. This dictionary is referred to as the *root* of the database. Every object that can be pickled into a standard serial format can be stored in the database. [49]

4.4.1.4 State Machine

Both the proxy and server module of the PoC are implemented using a finite state machine. This enforces the compliance of the implementation to the protocol specification, since only defined transitions are possible. In this prototype the light-weight, object-oriented Python implementation *transitions* is used [50].

4.4.2 Proxy Module Implementation

The basic operation of the proxy module entails that the client sends an HTTP request to the proxy, which evaluates the request. If the requested address is a UOSP introduction address, the proxy checks whether it has already stored an associated UOSP data address from previous requests. In this case, it sends this address as a response to the client. Otherwise, the proxy forwards the request to the server. The server provisions a new UOSP data service and sends its address to the proxy. The proxy forwards this address to the client and caches it for future requests. A detailed description of the protocol sequence can be found in section 4.3.4. The implementation of the proxy

module is described on the basis of its finite state machine in Figure 4.5. For the sake of readability, generic names were used for states and transitions, which are described in Table 4.1.

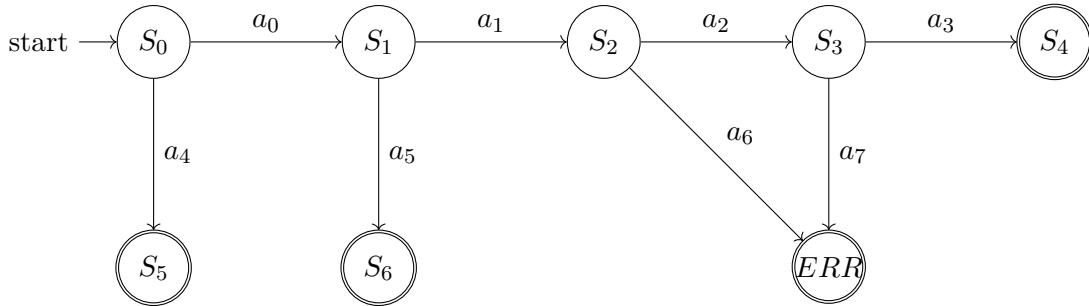


Figure 4.5: Finite State Machine of the UOSP Proxy Implementation

S_0	S_1	S_2	S_3	S_4	S_5	S_6	ERR
Initial state	Client request sent	Server public key received	Client ack sent	Server Fin received	Cache address used	Data response received	Error
a_0	a_1	a_2	a_3	a_4	a_5	a_6	a_7
Send request	Receive public key	Send UOSP Ack	Receive server Fin	Use cache address	Receive data response	Public key malformed	Server hash invalid

Table 4.1: Proxy Finite State Machine Legend

In the subsequent chapters, the implementation of the individual states and their transitions are described in detail. The complete source code of the proxy module can be found on GitLab [51] and in Annex A.

4.4.2.1 Initial State

To be able to receive requests from the client, the proxy uses the TCP server class of the socktserver module described in section 4.4. As this class processes request synchronously, the optional mixin class `ForkingMixIn` is used. As a result, a new process is forked for each request, which can be handled asynchronously. The `ThreadMixIn` class, which would create a new thread for each request, was not used because each individual request needs to be isolated to avoid protocol failures. During the forking of a process, a state machine is automatically created for each request handler. This state machine is used throughout the protocol execution to keep track of the state of the process. Starting from the initial state, there are two possible transitions.

In order to decide which of the two possible transitions shall be carried out, the HTTP request received from the client must be analysed. For this purpose, the Python module `http-parser` [52] is utilized, which automatically parses the request and maps its data structure to an object. From the request header, the original destination address and its port are required. In this proof of concept implementation, it is assumed that the proxy only handles client requests whose top-level domain of the destination address is `.onion` and whose destination is port 80. HTTP requests whose pattern differ from this are forwarded to their intended destination, but are not considered part of the protocol.

In the next step, the proxy examines whether a UOSP data address has previously been issued for the UOSP introduction address used in the request and whether this data address is still valid. For this purpose, the proxy maintains an object database as described in section 4.4. This database is referred to as *cache* in the future. The class diagram of the objects stored in this cache is depicted in Figure 4.6:

UospIntroAddress
<ul style="list-style-type: none"> - uosp_intro_address : str - uosp_data_addresses : PersistentList
<ul style="list-style-type: none"> + add_data_address(uosp_data_address: str, uosp_validity : datetime) : void + remove_data_address(uosp_data_address : str) : void + revise_validity() : void - get_latest_valid_data_address() : str - get_data_address_count() : int - get_address_index(uosp_data_address : str) : int - check_validity(uosp_data_address : str) : bool

Figure 4.6: Proxy Database Objects

For each UOSP introduction address not known to the proxy, an instance of the class *UospIntroAddress* is created. The properties of these instances are the introduction address and a list of associated data addresses. The data type used for *uosp_data_addresses* is specific to the implementation of the ZODB object database used. If the Python list datatype were used for this property, changes to the list would not be recognized by the ZODB and thus would not be persisted. The *UospIntroAddress* class provides methods for adding, removing and revising the validity of UOSP data addresses. To check whether valid UOSP data addresses are stored on the proxy for a UOSP introduction address, first a revision of the existing addresses is carried out using *revise_validity*. Subsequently, the number of stored UOSP data addresses is checked using *get_data_address_count*. If this number is greater than zero, the cache of the proxy contains a valid UOSP data address. In this case, transition a_4 is triggered, otherwise the proxy has not stored a valid UOSP data address and must make a request to the server. For this purpose, a socket connection is established with the server. Although for this prototype implementation HTTP is used on the proxy side, the traffic between proxy and server module is encrypted using the python *ssl* module to avoid the leakage of data addresses. As the server functionality is provided as an onion service, the request must be done within the Tor network. In this proxy implementation Tor version 0.4.3.5 is used, whose configuration can be found in Appendix A.8.

By default, Tor automatically starts a SOCKS proxy on port 9050 if it is not manually deactivated [53]. This SOCKS proxy is used by the proxy module to be able to communicate over the Tor network. The proxy module sends a request to the server using the standard UOSP message format illustrated in Figure 4.3a, which triggers transition a_0 .

4.4.2.2 Cache Address Used

The program execution transitions into this state if the proxy has a valid UOSP data address for a UOSP introduction address in its cache. This UOSP data address is sent to the client using the extended validity UOSP message format illustrated in Figure 4.3b. The payload of this message is an HTML page advising the user to use the UOSP data address provided. Alternatively, a HTTP redirect to the UOSP data address can be sent, which is carried out automatically by a browser, for example. As this is a final state, the protocol ends after submitting the response to the client and the process is stopped.

4.4.2.3 Client Request Sent

After sending a UOSP request the proxy enters an exchange loop where it waits for the response of the server. If a response is received, it is parsed in the first step using the Python module *http_parser*. This separates the HTTP headers from the payload and examines which type of response was retrieved, which results in three options:

1. The response has the form of an extended data UOSP message illustrated in 4.3c and thus is the response to a request that used a UOSP data address. The exchange of user data, although included in this state machine, is not part of the UOSP.
2. The response has the form of an extended validity UOSP message illustrated in 4.3b and thus is the response to a request that utilized a UOSP introduction address.
3. The response does not have a UOSP header and thus is a response to an onion service not capable of performing UOSP. The data received is relayed to the client, although the client data exchange is not part of the UOSP and is therefore not shown in the state machine.

In the first case, transition a_4 is triggered, whereby the state *Data response received* is reached. The reception of a public key as in the second case triggers transition a_1 and leads to the state *Server public key received*.

4.4.2.4 Data Response Received

In case of a response to a request that used a UOSP data address, the payload of the response must be forwarded. The payload is extracted from the response and sent to the client in the form of a standard UOSP message, whose structure is shown in Figure 4.3a. As this is a final state, the protocol ends after submitting the response to the client and the process is stopped.

4.4.2.5 Server Public Key Received

Once the public key has been received, its validity is checked in the first step. This can be implemented very simply, as the public key is available in its encoded form as onion service V3 address, which must contain exactly 56 characters. According to the rendezvous specification [25], it should additionally be checked whether the key contains a torsion component, as this would allow attackers to create multiple onion addresses with the same key for a single service. As this is only a prototype

implementation, this check is not implemented and can be considered as future work. If the public key is malformed, transition a_6 is triggered which results in the error state. Otherwise, it is verified in the next step that the integrity of the transferred key was not violated during the transmission. For this purpose, the obtained key is hashed with the *SHA-256* hash function, which is provided by the Python module *hashlib* [54]. The resulting hash value is appended as payload to a standard UOSP message, whose structure is illustrated in Figure 4.3a. This message is referred to as UOSP Ack in the protocol and triggers transition a_2 .

4.4.2.6 Client Ack Sent

After the UOSP Ack has been sent to the server, the proxy waits for its response in the form of an extended Data UOSP message. If the integrity validation fails, meaning that the hash value the server created doesn't match the hash value provided by the proxy, the response contains the payload *ERR*, which triggers transition a_7 and leads to the error state. Otherwise, the payload is *FIN*, which triggers transition a_3 and results in the state *Server Fin received*.

4.4.2.7 Server Fin Received

If the proxy receives the protocol message *Fin*, a new UOSP data address has been successfully generated. This address is sent to the client in the form of an extended validity UOSP message. As mentioned in the description of the state *Cache address used*, an HTTP redirect can be sent alternatively. Finally, the UOSP data address is cached. The datastructure used for the storage of the data addresses is depicted in diagram 4.6. If no *UOSPIntroAddress* object with the UOSP introduction address exists yet, a new object is instantiated. Subsequently, the UOSP data address is added using the *add_data_address* method. As this is a final state, the protocol ends after submitting the response to the client and the process is stopped.

4.4.3 Server Module Implementation

The primary function of the server module in the UOSP consists in the provisioning of temporary data services. These provisions are triggered by the proxy module through a specially formatted HTTP request to a UOSP introduction address. After validating the responsibility for the received request, the server module connects to the Tor binary, which is used to provision an ephemeral onion service. The onion address of this newly created onion service is stored in the server cache to be able to manage all the services provisioned. Subsequently, the address is transmitted to the proxy module using a UOSP message. The proxy module receives the address and responds with its hash value. The server module compares the received hash value with the hash value of the onion address generated by itself and determines whether the integrity has been compromised on the transmission path. If this is the case, it sends an error message to the proxy module. Otherwise, the server responds with a message that indicates the end of the protocol cycle. The implementation of the server module is described on the basis of its finite state machine in Figure 4.7. For the sake of readability, generic names were used for states and transitions, which are described in Table 4.2.




Figure 4.7: Finite State Machine of the UOSP Server Implementation

S_0	S_1	S_2	S_3	S_4	S_5	S_6	ERR
Initial state	Client request received	UOSP response sent	UOSP Ack received	Server Fin sent	Data response sent	Legacy HTTP response sent	Error
a_0	a_1	a_2	a_3	a_4	a_5	a_6	
Receive request	Send UOSP response	Receive UOSP Ack	Send server Fin	Send data response	Send legacy data response	Public key malformed	

Table 4.2: Server Finite State Machine Legend

In the following, the implementation of the individual states and their transitions are described in detail. The complete source code of the server module can be found on GitLab [51] and in Annex A.

4.4.3.1 Initial State

In order to be able to receive and respond to HTTP requests from the proxy, the server module must maintain an HTTP server, which needs to be provided as onion service. For the provisioning of the HTTP server, the Python module *http.server* [54] is used, which is a sub-class of the *socketserver* base classes described in section 4.4. This module includes a special sub-class of the *BaseRequestHandler* that allows to handle GET and POST requests in a straightforward way. As with the proxy module, the mix-in class *ForkingMixIn* is used, which creates a new process for each request. In addition, the python module *ssl* is used to encrypt the traffic between proxy and server module. It must be noted, that the ssl configuration used is intended for experimental purposes only and should not be used in a production environment. Since the server must be reachable via the tor network, the server module needs to launch a Tor binary. In this server implementation, Tor version 0.4.3.5 is used whose configuration excerpt can be found in Listing 1. The complete configuration is provided in Annex A.5.

```

1 ControlPort 9051
2 HashedControlPassword 16:0717EF71CDF228DE60E93242A4AFE681FB525080DC1A330A103FAD9782

```

Listing 1: Excerpt from the Tor Configuration of the Server Module

In contrast to the Tor configuration of the proxy, the configuration of the server activates the control port and protects it with a hashed password. To make the HTTP server accessible via an onion

service, the Tor Stem library is utilized. For the usage of Stem, auxiliary functions were developed, whose source code can be found on GitLab [51] and in Annex A.2. The headers of these auxiliary methods are provided in Listing 2.

```

1 def establish_connection(address=CONTROL_ADDRESS, port=CONTROL_PORT, password=None)
2 def create_ephemeral_hidden_service(controller, source_port, target_port, version)
3 def create_hidden_service(controller, path, port)

```

Listing 2: Header of the Tor Stem Auxiliary Methods

The method *establish_connection* is used to connect to the Tor control socket, while the methods *create_ephemeral_hidden_service* and *create_hidden_service* are used to create onion services. Ephemeral onion services are bound to the control connection through which they were created, which means that the ephemeral onion services in this implementation would be torn down after the protocol cycle is finished. To overcome this issue, the Tor Stem flag *detached=True* is used, which detaches the service from the controller so that it can continue to exist after the end of the control connection.

After the provisioning of its introduction service, the server module resides in its initial state and waits for a client GET request to arrive. A list of UOSP introduction addresses provided by the server module is maintained in order to be able to distinguish them from UOSP data addresses later on. The reception of a request triggers the transition a_0 , which leads to the state *Client request received*.

4.4.3.2 Client Request Received

Once the UOSP request has been received, it must be ensured that the connection is not terminated after the request has been processed and the response has been sent, since the server module must wait for the UOSP Ack. In order to avoid the connection tear-down, the argument *close_connection* of the handler class must be set to false. In the next step, the server examines whether the request received is a UOSP message, which is characterized by the presence of specific UOSP headers, as illustrated in Figure 4.3. If this is not the case, the message is a legacy HTTP request and the server sends an according response in the form of an extended data UOSP message to the proxy, which triggers transition a_5 . In this prototype, the data addresses can only be accessed via the proxy, since their direct utilization could lead to a dis-balanced usage of certain data addresses in terms of frequency and could potentially render the protocol useless.

If, on the other hand, the request contains the required UOSP headers, a further distinction must be made:

1. The requested address is one of the UOSP data addresses of the server module.
2. The requested address is one of the UOSP introduction addresses of the server module.

In the first case, the server must respond with the web page the client intended to reach. Therefore, an extended data UOSP message with the service's user data is created and sent to the proxy. This triggers transition *a4*, which results in state *Data response sent*. It must be pointed out that the exchange of data (both legacy and data requests) is not part of the UOSP, but has been implemented in this prototype.

In the second case, a new UOSP data service must be provisioned. For this purpose, a connection with the Tor control socket is established and an ephemeral onion service is provisioned using the auxiliary method excerpted in Listing 2. Each provisioned address is automatically assigned a validity by the server, which is currently set to 24 hours. A detailed analysis of the impact of the length of the validity period on the unlinkability of onion service is considered future work. After the validity period has expired, the service is torn down, which requires the server module to store information on all provisioned UOSP data services. For this purpose, the server maintains an object database as described in section 4.4.1. This database will be referred to as *cache* in the future. The class diagram of the object stored in this cache is depicted in Figure 4.8.

UospDataAddress
- uosp_data_addresses : PersistentMapping
+ add_data_address(uosp_data_address: str, uosp_validity : datetime) : void
+ remove_data_address(uosp_data_address : str) : void
+ revise_validity(controller: Tor Stem Controller) : void
- check_validity(uosp_data_address : str) : bool

Figure 4.8: Server Database Objects

In contrast to the proxy, the server module does not create a new object for each UOSP introduction address, as the distinction is not relevant when invalidating UOSP data services. For this reason, there is only one instance of the *UospDataAddress* class that is used by all processes. The only property of this object is a dictionary like data type that contains the provisioned UOSP data addresses as keys and their validity as values.

In a separate thread, the server periodically checks whether the validity period of one of the data services in the *uosp_data_addresses* list has exceeded by utilizing the method *revise_validity*. A handle for communication with the Tor control socket is passed to the method so that exceeded services can be teared down immediately. Finally, the service is deleted from the list using *remove_data_address*.

In case of a newly provisioned data service, its onion address is added to the database using the method *add_data_address*. Subsequently, the server prepares an extended validity UOSP message, whose structure is illustrated in Figure 4.3b and appends this address as payload. Finally, this message is sent to the proxy, which triggers transition *a1*.

4.4.3.3 Data Response Sent & Legacy HTTP Request Sent

After sending the data response to the proxy, the work for the server module is done. As this is a final state, the protocol ends and the process for handling the request is stopped.

4.4.3.4 UOSP Response Sent

Once the UOSP response has been sent to the proxy, the server module waits for its acknowledgement. Since data is transmitted with this response, the server expects the usage of the HTTP POST method rather than the HTTP GET method. The reception of that acknowledgement triggers transition a_2 , which results in the state *UOSP Ack received*.

4.4.3.5 UOSP Ack Received

In the first step, the received acknowledgement is divided into header and body. Since the *http.server* Python module was used for providing the HTTP server, no additional tools or modules are necessary in order to achieve this separation. It is assumed that the body of the acknowledgement solely contains the hashed public key that has been sent to it in the previous state. To assure the integrity of the transmission channel, the server also calculates the hash value of the address of the provisioned UOSP data service. Similarly to the proxy, a *SHA-256* hash is generated using the Python module *hashlib*. The hash value contained in the acknowledgement is then compared with the generated hash value. If these values differ, the encoded public key of the UOSP data service has been changed during transmission. This triggers the transition a_6 , which leads to the error state. Otherwise, the server sends a response in the form of an extended data UOSP message to the proxy that indicates the successful completion of the protocol cycle. This triggers transition a_3 , which results in the final state *Server Fin sent*.

4.4.3.6 Server Fin Sent

After sending the finalize message to the proxy, the work for the server module is done. As this is a final state, the protocol ends and the process for handling the request is stopped.

5 Performance Analysis

As stated in the related work section, Köpsell et al. [5] discovered that the increase in latency of an anonymized service is linearly related to the drop-out rate of users. Based on these findings, Wendolsky et al. [6] were able to define the tolerance level of inexperienced users with regard to the latency of a service at about four seconds. This implies that an initial response to a user request must not exceed this duration, otherwise the usability of the service is limited. That also applies to the *Unlinkable Onion Services Protocol* (UOSP) specified in section 4. Additional latency is added to the communication through the provisioning of UOSP data services and the exchange of their addresses.

This section examines whether the increase in the latency of an onion service through the use of the UOSP can be reconciled with the latency limit defined by Wendolsky et al. [6]. Since the provisioning of UOSP data services will most likely account for the majority of the additional latency, it is necessary to find a method to measure the duration of the provisioning of an onion service. Several temporal measurements of onion services have already been carried out [10, 12, 14], but these measurements focused primarily on the access-time of services and were designed for the out-of-date V2 onion services. This thesis fills this gap by designing and implementing a system to measure the performance of onion services with respect to their provisioning time. In the following sub-sections, the methodological approach, the architecture and the implementation of this system are described in detail. To conclude, the results of the implementation and its impact on the usability of the UOSP will be discussed.

5.1 Methodological Approach

In order to be able to measure the provisioning time of an onion service it must be defined from which point on a service is considered to be provisioned. For this purpose, it is necessary to look at the individual steps that are involved in the provisioning. Initially, the service randomly selects Tor nodes that shall act as its introduction points. It builds anonymous three-hop circuits to those nodes and utilizes the *Extensible ESTABLISH_INTRO* protocol [25] to register them as introduction points for the service. Once it receives the acknowledgements for this registrations in the form of *INTRO_ESTABLISHED* cells, the onion service builds its descriptors, which contain, among other things, information on the established introduction points. After the completion of descriptor creation, anonymous circuits to the *Hidden Service Directories* (HSDirs) responsible for the service for the respective time period and replica are established. Taking into account both replicas and the two time periods under consideration (current and previous), there are a total of 16 responsible directories, to which a descriptor upload is initiated. Since the uploads are conducted in parallel, the order in which the uploads are started need not correspond to the order in which they are finished. Therefore, the successful completion of a specific upload event is not sufficient to determine whether the provisioning of the service has been completed. In principle, there are three different approaches to identifying the finalization of an onion service provisioning:

- The provisioning of an onion service is completed if all of its descriptor uploads are finished. This would imply that provisioning fails if one of the 16 uploads is not successful.

- The provisioning of an onion service is completed if the service can be accessed by clients, which indicates that at least one upload has successfully finished. This approach is used by the Tor stem library to determine whether the service was published or not.
- The provisioning of an onion service is completed if at least 50 percent of all uploads finished successfully. The sequence of the uploads is thereby not relevant.

In this work, a blend of two of the approaches described above is used. All service provisions in which one of the uploads has failed are counted as failed provisions. This is necessary to preserve the comparability of individual measurements. On the other hand, for the calculation of the overall provisioning time of an onion service, only the time required to complete 50 percent of the uploads is taken into account.

Another methodological consideration concerns the granularity of the measurements. For the measurement of the provisioning time, it would be sufficient to measure the total time from the creation of the service until the completion of the required uploads. Nevertheless, this would prevent the detection of any performance bottlenecks, which could be mitigated by measures and would significantly complicate troubleshooting in case of performance issues. For this reason, all sub-steps of the provisioning of an onion service shall be observed. A list of all measurable steps is provided in Table 5.1.

#	Source	Target	Event	Measurable	Measured
1	Service	Service	Started onion service provisioning	Yes	Yes
2	Service	Introduction Points	Started the creation introduction circuits	Yes	Yes
3	Service	Introduction Points	Finished the creation of introduction circuits	Yes	Yes
4	Service	Introduction Points	Sended <i>ESTABLISH_INTRO</i> cells	Yes	Yes
5	Introduction Point	Introduction Point	Registered introduction points	Yes	No
6	Service	Introduction Points	Received <i>INTRO_ESTABLISHED</i> cells	Yes	Yes
7	Service	Service	Started the creation of service descriptors	Yes	Yes
8	Service	Service	Finished the creation of service descriptors	Yes	Yes
9	Service	HSDirs	Started the descriptor uploads	Yes	Yes
10	Service	Service	Started the creation of upload circuits	Yes	No
11	Service	HSDirs	Finished the creation of upload circuits	Yes	Yes
12	Service	HSDirs	Finished the descriptor uploads	Yes	Yes

Table 5.1: Measurable Events During the Provisioning of an Onion Service

All actions in this table are referred to as events. The first event that takes place in the course of an onion service provisioning is the service creation. Thereby, two different types of services can be created. Non-ephemeral onion services are generated by either changing the in-memory Tor configuration or the Tor configuration file (*torrc*). In the latter case, the Tor binary must be reloaded in order to recognize the modification. For each non-ephemeral onion service, a new folder is created in the file system of the service host, which contains, among other things, the key material of the service, keys for client authentication and the host name of the service. On a restart of the Tor binary, all onion services contained in the Tor configuration are automatically re-provisioned. In contrary, ephemeral onion services are not created via the configuration file, but via the control

socket of the Tor binary. For this purpose, for example, the Tor *Stem* library can be used. Ephemeral services exist only as long as their control connection exists. If the process through which the service was created using the control protocol is terminated, the service is also torn down. To decouple the services from their process, the attribute *detached* must be set to *True* when the service is provisioned. However, this has no effect on the persistence of ephemeral services. Restarting the Tor binary will not re-provisioning ephemeral onion services. Since both types of services can be created using the Tor control protocol, the creation timestamp of an onion service corresponds to the time the required control command was issued.

The second measurable event is the start of the creation of the circuits from the service to the chosen introduction points. If available, Tor tries to use existing pre-built circuits for that purpose, which only need to be extended by one hop [14]. This technique is referred to as *cannibalization* and decreases the time for the service provisioning, since the usage of a cannibalized circuit does not introduce further delay [14]. The completion of introduction circuit creation is measured by means of the third event. Through the measurement of the start and end times of introduction circuit creations, their duration can be computed. The fourth measurable event after the creation of the introduction circuits is the sending of the *ESTABLISH_INTRO* cells, which are needed to register the chosen Tor nodes as introduction points of the service. Once the Tor nodes receives the requests from the service, they carry out cryptographical validations and sanity checks and register themselves as introduction points. The duration of this registration can only be measured on the introduction points themselves, which means that, for the measurement, an introduction point would have to be operated and the Tor source code would have to be modified such that this introduction point is selected by the service. As this would not provide any significant new insights, this specific duration is not measured in the course of this thesis. The introduction points acknowledge the registration by sending an *INTRO_ESTABLISHED* cell to the service. This is understood as the completion of the establishment of the introduction points and is measured as the sixth event. Since the start and end times of the introduction point establishments are measured, its duration can be computed.

Using the established introduction points, the service generates its service descriptors in the next step. The start and the end time of this generation are captured in the seventh and eighth event, which allows the duration of this process to be measured. In the last step, the created descriptors are uploaded to the responsible HSDirs. For this purpose, circuits to the responsible directories must be established. Similarly to the introduction circuits, cannibalization can be used to speed up this process. It may seem odd that the event that indicates the start of the descriptor upload chronologically precedes the event for the creation of the upload circuits, but this event triggers the whole upload process, without which the circuits would not be created. Since the time gap between the ninth and tenth event corresponds to a median of about *0.00144* seconds, the start of the upload circuit creation is not measured, instead, the start of the descriptor uploads also marks the start of the circuit creation. The completion of the upload circuit creation is measured in event eleven and simultaneously indicates the start of the descriptor uploads. The uploads are asynchronous, which means that the sequence in which the uploads start does not necessarily influence the sequence in

which they end. Completion of uploads is measured in the twelfth event. As the start and end times of the descriptor uploads are measured, its duration can be calculated. Additionally, the HSDirs for the individual descriptor uploads are recorded.

Finally, it must be determined how the events specified in Table 5.1 can actually be measured. Essentially, there are two possible approaches:

1. Tor offers an extensive log system, which can be enabled by appropriate entries in the *torrc* configuration file. Individual severity levels or functionality domains can be selected, whose log messages are then redirected to a file, the standard output or syslog, for example [53]. Therefore, one possible approach is to configure custom logging, which saves the required events to the chosen medium. On completion of provisioning, the log events from the selected medium are read out and evaluated.
2. All required events can be obtained through the use of the Tor control protocol. Thereby, a user or a program can register for specific event classes and then receive all log events corresponding to this class.

In the course of this work, the events are obtained via the Tor control protocol using the Tor Stem library, as this provides a considerably better handling in Python than reading from files or reading the standard output. In addition, each event class can be handled in its own thread, which enables parallel processing.

In principle, the focus of the measurement is set on V3 onion services, as V2 onion services will no longer be supported in the foreseeable future [15]. Nevertheless, currently more than 170,000 V2 onion services are still running, which renders them a non-negligible group in terms of their high prevalence [41]. Therefore, V2 onion services are included in the measurements in order to be able to compare their temporal behaviour and performance with the current version. In addition, the influence of the method used to create the onion service (ephemeral and non-ephemeral) and the use of additional security features, as described in 3.2.5, is measured.

5.2 Architecture

In order to obtain measurement results that are ideally unaffected by possible sources of interference, these interferences must be identified in the first step. Since the components used to perform the measurements are most likely operated within an university or corporate network, they are subject to fluctuations in network performance dependent on the time of day and number of users. Furthermore, if the underlying hardware of the server on which the components are operated is shared with other applications, their utilization can influence the hardware's resource consumption and thus influence the measurement performance. Similarly, the Tor network, which is used for the provisioning of onion services, is subject to fluctuations in terms of bandwidth and circuit build times that could potentially have an impact on the measurement results. To prevent these short-term fluctuations in the networks and systems from having an effect on the measured values, numerous measurements must be performed at different points in time.

In order to achieve this, an architecture is required that automatically provisions onion services, measures the events that occur in this process and outputs them in a format that can be used for analysis. To avoid measurements influencing each other (e.g. through caching), each measurement must be carried out completely isolated from the other measurements. This requires an environment that can be reset or recreated for each measurement with no significant overhead. In order to achieve this, containers are used, which are inherently isolated from each other through operating system features. In addition to avoiding interfering influences, a significant advantage of using containers consists in the fact that faulty measurements do not influence other measurements in terms of a temporal delay.

The minimal set of components required for measuring the provisioning time of onion services using containers are depicted in Figure 5.1. In the following, the individual components and their mode of operation are described.

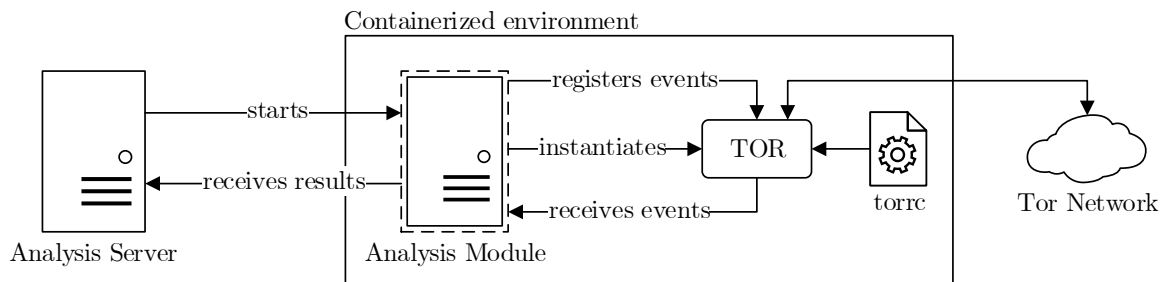


Figure 5.1: Performance Measurement Architecture

- **Analysis Server**

The *Analysis server* constitutes the backbone of the measurement system. It periodically initiates a new measurement instance by starting a new measurement container. This container performs all necessary steps for the measurement and returns the result to the analysis server in a pre-defined format. Irrespective of the wording, the analysis server does not necessarily have to be deployed on a separate physical server, but can be located on the same server as the host of the measurement containers.

- **Analysis Module**

All steps necessary for measuring the provisioning time of an onion service are performed by the *Analysis Module*. In the first step, a new Tor process is created and Tor is bootstrapped using the configuration in the provided *torrc* file. Subsequently, the module connects to the Tor control socket using the Tor stem library and registers for several protocol events. When these events occur, the module receives notifications containing all information about the events. Once all event listeners are registered, a new onion service is created. Now the time values relevant for the measurement are received by means of the previously registered event

listeners. As soon as the generation of the service and its measurement are completed, these time values are analysed and correlated. The result of this analysis is returned to the analysis server in a pre-defined format.

- **Tor Binary**

The Tor binary represents the connection between the analysis module and the components of the onion service to be provisioned in the Tor network. To avoid distorting the measurements, it is important that no caching of any kind is performed on the side of the Tor binary. This is ensured by creating a new Tor process for each measurement.

- **Tor Network**

Since onion services can only be provisioned and accessed in the Tor network, the analysis module and respectively the Tor binary must be connected to it.

The advantages of this architecture are its modularity and its ease of use. Since the analysis module is implemented as a container, it includes all components required for the measurement. A measurement process is initiated simply by starting the container, which returns the result in a format that can be read by a wide range of analysis tools. The most important configuration settings can be set with parameters that are specified when the container is started, without having to change the container. In addition, it is possible to replace the analysis server with another solution and thus integrate the analysis module into another software project.

5.3 Implementation

As described in section 5.2, the architecture for measuring the provisioning time of onion services consists of the three main components *Analysis Server*, *Analysis Module* and the *Tor* binary. In the following subsections, the implementation and usage of these components and all decisions made in their development process are described in detail. Since space is limited, only excerpts from the source code can be shown in the solution description. The complete source code of the implementation is provided on GitLab [55] and in Annex B.

5.3.1 Analysis Module

The analysis module forms the core of the measurement system. It is implemented in *Python 3.7*, due to the dependency on the Tor Stem library [46], which is used to measure the events listed in Table 5.1. A description of the structure of this library can be found in section 4.4.1. The analysis module consists of several Python files, whose structure is shown in Figure 5.2.

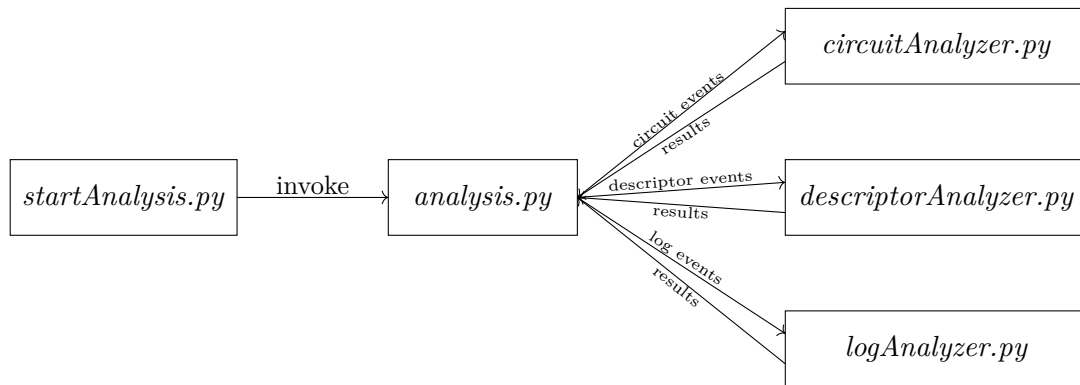


Figure 5.2: Analysis Module Program Structure

In order to be able to connect to the Tor control socket and ensure that there is no timing corruption, it is important to verify that the Tor process is fully bootstrapped. This verification is realized within the file *startAnalysis.py* by matching a predefined string sequence that indicates the completion of the bootstrap using regular expressions. The source code is provided in Annex B.5.

Once a match is found, the actual measurement activity is started by invoking the function *start* in the file *analysis.py*. The header of this function is shown in Listing 3.

```
1 def start(version, ephemeral, timeout, num_intro_points)
```

Listing 3: Measurement Invocation

Thereby the main measurement settings are passed as function parameters. The parameter *version* specifies the version of the onion service to be provisioned and measured. Currently, the only recognized values are version two and version three. If *ephemeral* is set to true, the module provisions an ephemeral onion service. Otherwise, a non-ephemeral onion service is generated. The parameter *timeout* defines the maximum overall duration a measurement process may last before it is aborted. This is intended to prevent a failed or incomplete measurement from delaying the measurement procedure. The number of introduction points an onion service uses can be specified in the parameter *num_intro_points*. However, this value is only used to inform the module of how many introduction points are expected and has no influence on the number of the actual introduction points used, which can be set by manipulating the Tor source code.

In the first step of the measurement process, a connection with the Tor control socket is established using the *from_port* function of the Stem library, which returns a connection handle referred to as *controller*. Subsequently, this controller is used to register event listeners for certain groups of Tor events using the Stem function *add_event_listener*. As parameters, this function is provided with the event or group of events to be listened to and the name of the function to which the event or group of events should be relayed. The following groups of events are required to measure the onion service performance: [46]

- **CIRC**

A *CIRC* event takes place when a circuit has changed. The content layout of the information contained in the event is shown in Figure 5.3. One of the most valuable fields in this structure with regard to the comprehensibility of the measured information is the *id*. This represents the unique identifier of a circuit, which can be used to track it. The *status* field provides information about the status of a circuit. The fingerprints and nicknames of all Tor relays currently used by the circuit are contained in the field *path*. The field *purpose* contains a description of the purpose of the circuit. Information on the other fields can be obtained from the Stem manual [46].

<i>id</i>	<i>status</i>	<i>path</i>	<i>build_flags</i>
<i>purpose</i>	<i>hs_state</i>	<i>rend_query</i>	<i>created</i>
<i>reason</i>	<i>remote_reason</i>	<i>socks_username</i>	<i>socks_password</i>

Figure 5.3: Structure of a CircuitEvent

All circuit events are forwarded to the function *circCallback* provided in the file *circuitAnalyzer.py*, whose source code is provided in Annex B.2. As described in Table 5.1, the start and end time of the introduction circuit creation and the end time of the upload circuit creation are measured. For this purpose, all circuits elements in relation to the onion service must be recorded. This is achieved by evaluating the fields of the circuit event. Due to the gradual nature of circuit building, it is necessary to extract the individual circuit relays from a circuit event and assign them to their corresponding circuit according to the circuit id in the *id* field. This is accomplished by extracting the circuit fingerprints from the *path* field of the event and assign the relay to the circuit it extended.

In order to be able to effectively access all the nodes of a circuit, a class for storing the individual circuit elements was developed, whose class diagram is provided in Figure 5.4.




Figure 5.4: Circuit Element Class Diagram

Each node in the circuit is identified by its unique *fingerprint*. The time at which the node was added to the circuit is stored in the attribute *time_added*. This attribute is decisive for the subsequent time measurement. Both attributes can be obtained using a getter method.

To be able to assign the individual circuit elements to their circuits, a dictionary entry is created for each circuit, whose key is the unique circuit id. The values of these dictionary entries are lists to which the individual circuit elements are appended. In the end, all elements of a circuit can be retrieved by accessing the dictionary using the specific circuit id.

- **CIRC_MINOR**

As the name suggests, the *CIRC_MINOR* events is triggered when minor changes to circuits are performed. The structure of the information contained in the event is illustrated in Figure 5.5. Since this is also a circuit event, the fields *id*, *path*, *build_flags*, *purpose*, *hs_state*, *rend_query* and *created* are identical to the *CIRC* event. The *event* field provides information about the minor change that has been carried out on the circuit. Thereby it is distinguished whether an existing circuit is cannibalized or the purpose of the circuit has changed. Information on the other fields can be obtained from the Stem manual [46].

<i>id</i>	<i>event</i>	<i>path</i>	<i>build_flags</i>
<i>purpose</i>	<i>hs_state</i>	<i>rend_query</i>	<i>created</i>
<i>old_purpose</i>		<i>old_hs_state</i>	

Figure 5.5: Structure of a CircuitMinorEvent

This event is used to distinguish between a newly built or a cannibalized circuit by using only the fields *id* and *event*. For this, all circuit minor events are forwarded to the callback handler *circMinorCallback* provided in the file *circuitAnalyzer.py*. The procedure of the callback handler is identical to the procedure of *circCallback* described above.

- **HS_DESC**

Any operations that affect the descriptors of an onion service trigger the *HSDescEvent* event, whose structure is depicted in Figure 5.6. The field *action* contains the description of the action carried out. An identifier of the Hidden Service Directory (HSDir) to which the descriptor is uploaded is contained in the field *directory_fingerprint*. Information on the other fields can be obtained from the Stem manual [46].

<i>action</i>	<i>address</i>	<i>authentication</i>	<i>directory</i>
<i>directory_fingerprint</i>	<i>directory_nickname</i>	<i>descriptor_id</i>	<i>reason</i>
<i>replica</i>		<i>index</i>	

Figure 5.6: Structure of a HSDescEvent

All descriptor events are relayed to the method *descCallback* provided in the file *descriptorAnalyzer.py*, whose source code is provided in Annex B.4. The events are necessary for the measurement of the duration of the descriptor creation and the descriptor upload. To distinguish the individual events, the field *action* is utilized. The action *CREATED* indicates that the onion service descriptor was just created, whereas the action *UPLOAD* and *UPLOADED* mark the beginning and the end of the descriptor upload. Since numerous descriptor events take place during the provisioning of an onion service, the events are stored separately and are correlated and analysed later on. For storing the events, an object structure similar to the structure of the circuit events is used, which is illustrated in Figure 5.4. Instead of the circuit fingerprint, the fingerprint of the HSDir the descriptor is uploaded to is stored in the attribute *fingerprint*.

- **LogEvent**

With the help of *LogEvents* one gains access to all log events of the Tor process. Each time the process generates a log entry, this event is triggered. The structure of a log event is illustrated in Figure 5.7. The *runlevel* refers to the importance rating of the event, while the actual log message is contained in the field *message*.




Figure 5.7: Structure of a LogEvent

Using the log events, the completion of the introduction point establishment can be recognized. For this purpose, all log events are forwarded to the methods *logCallback_v2* or *logCallback_v3* provided in *logAnalyzer.py*, depending on the version of the onion service to be measured. The source code of this file is provided in Annex B.3. The duration of the introduction point establishment is calculated in the analysis using the end time of the introduction circuit creation and the introduction point established notification from the log events. To be able to link these two types of events, the unique circuit id must be extracted from the log event. This is achieved by using the regular expressions shown in Listing 4.

```

1 INTRO_ESTABLISHED_REGEX_V3 = 'service_handle_intro_established\(\):
  ↳ Successfully received an INTRO_ESTABLISHED cell on circuit ([0-9]{1,})
  ↳ \(\id: ([0-9]{1,})\)\'
2 INTRO_ESTABLISHED_REGEX_V2 = 'rend_service_intro_established\(\): Received
  ↳ INTRO_ESTABLISHED cell on circuit ([0-9]{1,}) \(\id: ([0-9]{1,})\)\'

```

Listing 4: Introduction Point Establishment Log Event Regular Expressions

Depending on the service version, a different regular expression is used. In addition to the circuit ids, the time at which the event was received is stored.

In the next step, all results have to be collected from the individual event handlers. Since each event handler is operated in its own thread, a solution for safe inter-thread communication has to

be used to exchange the results. For this purpose, synchronized queues are used. As the threads are implicitly created by Stem and are not aware of when the provisioning of an onion service is completed, it is difficult to obtain the results from the queues, because it is not possible to check the content of a queue element without removing it from the queue. For this reason, a workaround was developed in which the queue elements are removed from the queue for checking at periodic intervals and are added to the queue again by the handlers at the same frequency. This procedure is terminated and the Tor event listeners are detached once all required values are available. In order to determine when this is the case, the onion service protocol values from its specification are utilized (e.g. number of introduction points). A timer prevents the measurement process from taking too long due to missing or faulty values. Currently, this time-out is set to 180 seconds.

Finally, the collected results have to be correlated and processed in order to be analysed and stored in an uniform format. This necessary post-processing is individual and depends on the data to be processed. The following processing steps were carried out in the course of this thesis:

- **Descriptor Events:**

In case of the descriptor analyzer, the individual events have already been divided into three lists based on their appropriate action (*created*, *upload* and *uploaded*), which simplifies their processing. One further advantage lies in the storage of the individual events, which now enables an effective correlation. To obtain the descriptor creation time, the object-method *getTime* is used to receive the creation time stamp of the descriptor event objects in the *created* list. These time values are assigned a label, which allows them to be associated with the respective descriptor. The assignment of upload events to their corresponding uploaded events has proven to be more complex, since both events are stored in separate lists. In the first step, all upload event objects are read from the list of upload events and placed in the dictionary, such that the individual upload procedure can be accessed using a label. Subsequently, the relevant uploaded event is extracted from the list of all uploaded events on the basis of its unique circuit fingerprints. Finally, the timestamps of these uploaded events are obtained using the object method *getTime* and assigned to the upload procedure using its label. In addition, the path of the upload circuit and the directory used for each upload are stored. Thereby, the path of an upload circuit is obtained from circuit events and the responsible directory is obtained by querying the Tor network status using the unique relay fingerprint from the upload events.

- **Circuit Events:**

To obtain the start and end time of the introduction circuit creation, initially the circuit id of the circuit to be analysed has to be obtained. This is done by firstly parsing the log events that indicate the establishment of an introduction point. These events contain, amongst other information, the circuit ids of all introduction circuits. This circuit ids can then be used as key in the dictionary of introduction circuits to get all circuit elements that correspond to a specific circuit. The temporally first and last element can now be selected from this set of elements, which results in the start and end events of the introduction circuit. Finally, the time stamps of these elements are extracted using the *getTime* method of the circuit element objects.

- **Log Events:**

The last time values to be obtained are the introduction point establishment times. As mentioned above, the log events are utilized for this purpose. Since these events are not stored within objects, the time stamps are passed in the form of lists, which contain, apart from the timestamp, the ids of the introduction circuits. The time stamps are extracted and saved in a labeled data structure, to be then assigned to their corresponding introduction circuits.

All measurement fields for which no value could be determined due to a timeout or any other error are assigned the value *FAIL* in order to be able to distinguish them from regular fields later in the evaluation. Finally, the measurement results are saved in a csv file using the Python module *csv*, as this format is supported by the most common analysis tools.

5.3.2 Container Environment

As described in section 5.2, the measurement is carried out in a containerized environment to avoid falsification of the results by any influences. The virtualization software Docker is used for this purpose, which uses operating system on-board techniques like *control groups* or *namespaces* to isolate the individual environments [56]. Its architecture consists primarily of three parts. The *Docker Client* communicates with the *Docker Daemon*, which is responsible for the building, running and distribution of Docker containers. All containers are based on an image, which is a template that contains all necessary building blocks for the application that shall be hosted. *Docker Registries* are services that store and provide such images. Besides the predefined images, one can also create customized images, which are either written from scratch or based on other images. All the steps that are necessary to create a custom image are collected in a so-called *Dockerfile*. [56]

Since the requirements for the onion service measurement cannot be fulfilled by a standard environment, a custom image based on *Debian* was created. The most important excerpts from the Dockerfile created for this purpose are described below. The complete file is provided on GitLab [57] and in Annex B.7.

One of the main components of the resulting image is the Tor binary, which is required for the communication with the Tor network. As described in section 5.3.1, in addition to circuit and descriptor events, log events are also required to measure the provisioning time of an onion service. Thereby, the logs contain the notifications of the establishments of introduction points, which are then linked to the corresponding introduction circuit events. This linking takes place on the basis of the unique circuits identifiers. Since circuit events and log events use different storage structures that have different identifiers in the context of introduction circuits, the Tor source code had to be adapted to establish a shared identifier. Specifically, the log outputs of the introduction point establishment notification for V2 services (see Listing 5) and V3 (see Listing 6) services were changed in a way that, in addition to the structure-specific identifier, the global circuit identifier is also provided. This enables the linking of circuit and log events.

```

1 //rendservice.c:3443
2 log_debug(LD_REND, "Received INTRO_ESTABLISHED cell on circuit %u (id: %" PRIu32 ")
  ↳ for service %s", (unsigned)circuit->base->n_circ_id, circuit->global_identifier,
  ↳ serviceid);

```

Listing 5: Tor Source Code Log Extensiveness Extension V2

```

1 //hs_services.h:3324
2 log_debug(LD_REND, "Successfully received an INTRO_ESTABLISHED cell on circuit %u
  ↳ (id: %" PRIu32 ") for service
  ↳ %s", TO_CIRCUIT(circ)->n_circ_id, circ->global_identifier,
  ↳ safe_str_client(service->onion_address));

```

Listing 6: Tor Source Code Log Extensiveness Extension V3

The modified Tor source code is added to the container and built using build automation tools. The second important component of the container is the Python measurement environment required to measure the provisioning time of onion services. In order to be able to use it in the container, the Python environment and a package-management system are installed in the first step. The package-management system is used to install the Python modules *Stem* and *Vanguards*. *Stem* provides the basis for the measurement procedure and is described in section 4.4.1. The *Vanguards* module adds additional security features to onion service, which are delineated in section 3.2.5. These security features can be additionally adjusted by a configuration file. To obtain comparable results, the default configuration of *Vanguards* was used in the course of this work. In order to make the analysis module usable, finally its source code is added into the container.

Additionally, in order to automatize the measuring process as far as possible, a script was developed which performs all necessary steps to start the measurement procedure based on the measurement parameters provided. These parameters include the version and type of the onion service to be measured, a timeout which is awaited before the measurement is aborted and the number of introduction points used. The script starts the Tor process using the containers *torrc* file and subsequently starts the Python measurement procedure that generates a csv file with the measurement results. To ensure that this script is invoked when the container is executed, the script is set as the entry point of the Docker image. The source code of the script is provided in Annex B.10.

5.3.3 Analysis Server

If a container based on the containerized measurement environment is executed, one measurement cycle is performed. To obtain representative results, it is necessary to automatically carry out numerous measurements cycles. For this purpose, the analysis server in the form of a bash script was developed, whose source code can be found in Annex B.8.

The following parameters can be set in this script:

- **Target directory:** Path to the directory where the results of the measurements are stored.
- **Docker image:** Name and tag of the custom docker image to be executed.
- **Timeout:** Individual measurement timeout.
- **Introduction points:** Number of introduction points to be expected.

In order to be able to conduct the measurements of the individual versions and types of onion services under comparable conditions, all measurement variants are carried out serially one after the other within one measurement cycle. Finally, the individual files of the measurement cycles are merged with the help of a Python script in order to obtain one single csv file for each measurement type.

5.4 Results

The provisioning times of onion services were measured in the four different variants V2, V3, V3 non-ephemeral and V3 with Vanguard using the implementation described in section 5.3. To obtain the results presented in this section, two measurement sessions were conducted:

1. In the course of the first measurement session, 1520 individual measurements per variant were carried out. This session covered a time period of about three days, which reduces the sensitivity of the measurement results to errors due to temporal fluctuations in the performance of the server or the Tor network.
2. As the first measurement session revealed interesting irregularities, which could not be analysed without additional data material, a second measurement was carried out with an extended measurement scope. In the course of this second session approximately 500 measurements were carried out per variant.

Table 5.2 shows the most important statistical characteristics of the first measurement session. Note that all values in the table are specified in seconds.

<i>Type</i>	<i>Samples</i>	<i>Mean</i>	<i>Median</i>	<i>SD</i>	<i>25%</i>	<i>75%</i>	<i>Min</i>	<i>Max</i>	<i>Error</i>
V2 default	1520	32.82	32.53	3.88	31.83	32.95	26.93	133.53	4.21%
V3 default	1520	5.13	4.23	3.76	3.10	4.23	0.57	93.35	12.69%
V3 non-ephemeral	1520	5.04	4.19	3.85	3.07	6.02	0.79	104.53	14.27%
V3 vanguards	1520	5.49	4.36	5.09	3.25	6.30	1.15	94.08	15.26%

Table 5.2: Total Provisioning Time of Different Onion Service Types

A graphical representation of these numbers in the form of a box plot is provided in Figure 5.8.




Figure 5.8: Total provisioning Time of different Onion Service Types

The provisioning time of a V2 service averaged at 32.82 seconds, that of V3 services at 5.13 seconds, that of non-ephemeral V3 services at 5.04 seconds and that of V3 services with Vanguards plugin enabled at 5.49 seconds. These numbers indicate that the average provisioning time of an onion service V3 is equivalent to about one sixth of the provisioning time of an onion service V2. Also, the minimum provisioning time of V2 services is significantly higher compared to V3 services. The maximum values of the provisioning time also differ, but not as drastically as the minimum values. The error rate, which indicates whether one of the sub-steps in the provisioning of an onion service has failed, is about two-thirds higher for onion service based on V3 than for V2 services. Since the median and the mean of the individual measurements variants are close to each other, the results appear to be evenly distributed. In order to be able to use statistical tests to compare the provisioning times of the individual variants, it is necessary to ascertain whether the data are normally distributed. In the first step, this assessment is accomplished with the help of a *quantile-quantile* (q-q) diagram, which is shown in Figure 5.9.

The q-q diagram is an efficient way to compare the distributions of two data sets. Thereby, the quantiles of a distribution are plotted against the quantiles of another distribution. If the quantiles of the two distributions are approximately equal, the quantile-quantile points in the diagram should form a straight line, represented by the red line in the graph. In order to check if the measured data is normally distributed, it is plot against generated normally distributed data. [58]

As can be seen in the individual graphs in Figure 5.9, the results of none of the measurement variants are normally distributed, since q-q points significantly deviate from the ideal line. To confirm the

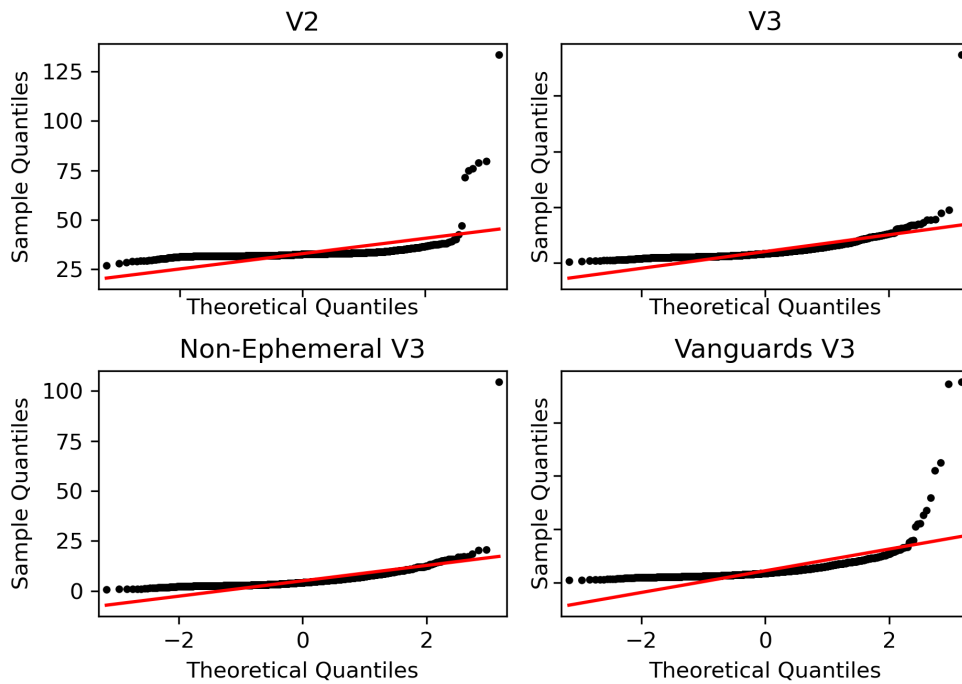


Figure 5.9: Quantile-Quantile (Q-Q) Diagram

graphical results quantitatively, a *Shapiro-Wilk* normality test with a significance level of 1% was carried out for all measurement variants, which also came to the conclusion that the results are not normally distributed. Therefore, the non-parametric *Mann-Whitney U* test is used to compare the provisioning times. In this analysis, the overall provisioning time is the dependent variable and the measurement variable is independent. The following hypotheses are used in the tests:

- H_0 : The samples of both measurement variants are from the same distribution. This means that there are no significant deviations in the provisioning times of the two onion service variants considered.
- H_1 : The samples of both measurement variants are not from the same distribution. This means that there are significant deviations in the provisioning times of the two onion service variants considered.

The results of the Mann-Whitney U test are shown in Table 5.3.

<i>Type one</i>	<i>Type two</i>	<i>p</i>	<i>Result</i>
V2 default	V3 default	< 0.0001	reject H0
V3 default	V3 non-ephemeral	0.0125	can't reject H0
V3 default	V3 vanguards	0.5255	can't reject H0
V3 non-ephemeral	V3 vanguards	0.0016	reject H0

Table 5.3: Results of the Mann-Whitney U Test

Based on a significance level of 1%, H_0 is rejected if the p value of the Mann-Whitney U test is less than 0.01. Comparing the measured values of the provisioning times of V2 and V3 onion services, the null hypothesis is rejected, which means that the two distributions differ significantly. The same applies to the comparison of non-ephemeral V3 services and V3 services that use the Vanguard plugin. There is no significant difference between V3 onion services and the two onion services V3 based variants. From this results in can be concluded, that there are significant differences in the provisioning times of V2 onion services and V3 onion services while the different variants of the V3 onion services differ only slightly.

The bar chart in Figure 5.10 provides a visual comparison of the duration of the individual sub-steps of the provisioning of the different onion service variants. The mean values of the provisioning sub-steps used for this plot are also provided in Table 5.4. The term *SD* stands for standard deviation and all values in this table are provided in seconds.

<i>Type</i>	<i>Introduction Point Establishment</i>	<i>SD</i>	<i>Descriptor Creation</i>	<i>SD</i>	<i>Upload</i>	<i>SD</i>
V2 default	0.3955	0.3594	30.9890	0.6475	1.4383	3.8020
V3 default	0.4173	0.6143	0.7665	0.2463	3.9725	3.7332
V3 non-ephemeral	0.3585	0.3319	0.8178	0.2775	3.9200	3.8398
V3 vanguards	0.6350	0.3921	0.6220	0.4095	4.2403	5.0283

Table 5.4: Sub-steps of the Provisioning Time of Different Onion Service Types

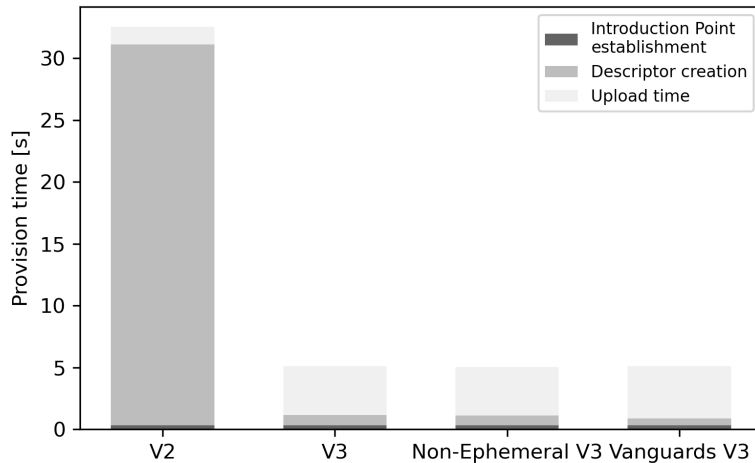


Figure 5.10: Onion Service Provisioning Sub-step Duration

In comparison to the mean duration of the other sub-steps involved in the provisioning of an onion service, the time required for the establishment of introduction points had the least impact on

the overall provisioning time. This applies to all four variants, for which the introduction point establishment took approximately the same amount of time. This contrasts with the descriptor creation time that took about 30.98 seconds for onion service V2, which is about 38 times longer than for the other variants. Also the upload time differs significantly between V2 and V3 services. Since the standard deviation of the upload times is very high compared to the standard deviation of the other sub-steps, it can be assumed that the upload times are spread out and not clustered around the median. The distribution of the upload times is shown in the bar chart in Figure 5.11.




Figure 5.11: Onion Service Upload Time Distribution

From the graph it can be concluded that the majority of uploads took about 5 seconds, while almost all uploads were completed after 20 seconds. An interesting result of the analysis is that some uploads required more than 100 seconds. In order to further investigate this issue, additional measurement values were collected during the second measurement session, which allow the subdivision of the upload time into the duration of the creation of the upload circuits and the actual HTTP upload time. The resulting distributions are depicted in Figure 5.12 and Figure 5.13. From these bar plots it can be concluded that the conspicuously high total upload time is due to issues with the creation of the upload circuits.

5.5 Discussion

The results of the performance measurements provided in section 5.4 indicate that the usage of the Unlinkable Onion Services Protocol (UOSP) in its current version in principle is not feasible with regard to its additional latency. According to the work of Wendolsky et al. [6], the duration inexperienced users of anonymization services are willing to wait is approximately four seconds. Assuming that the data service provisioned in the course of the UOSP is an onion service version




Figure 5.12: Onion Service HTTP Upload Duration

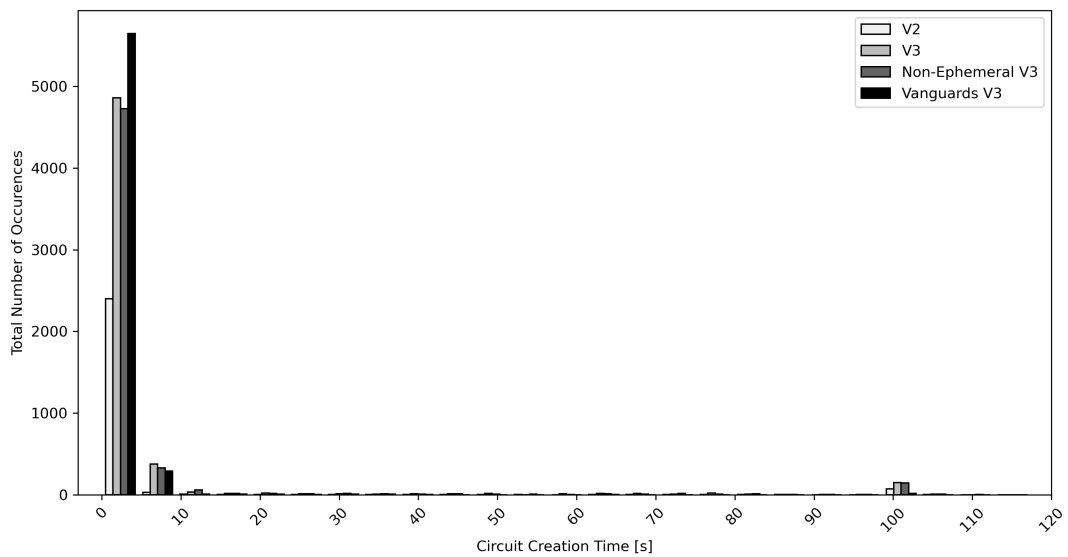


Figure 5.13: Onion Service Upload Circuit Creation Duration

three, this provisioning would require 5.13 seconds on average. In addition to provisioning, the overall response time is composed of the following three HTTP round trips:

1. Client UOSP request to the proxy and its response
2. Proxy UOSP request to the server and its response
3. Proxy UOSP acknowledgement to the server and its response

To be able to approximate the overall response time, HTTP Round Trip Times (RTTs) from the measurement server to an external HTTP server were conducted, whose results are shown in Table 5.5. In total 1000 measurements were performed. The payload of the HTTP responses was kept minimal to simulate UOSP messages.

<i>Mean</i>	<i>Median</i>	<i>SD</i>	<i>25%</i>	<i>75%</i>	<i>Min</i>	<i>Max</i>
0.4097	0.3964	0.0408	0.3845	0.4216	0.3665	0.8124

Table 5.5: HTTP Round Trip Times (RTT)

Based on the average RTT, a client would have to wait approximately 6.35 seconds for the first response from the proxy, which clearly exceeds the four-second limit. In contrast, if a UOSP data service has already been provided and its address is in the cache of the proxy, the overall response time consists solely of an HTTP round trip and the query of the cache of the proxy. This would significantly fall below the four-second limit and thus enable the use of the UOSP.

As can be observed from the data in Table 5.2, the average provisioning of onion service V2 is considerably slower compared to V3 services. After combining this finding with the durations of the provisioning sub-steps in Table 5.4, it becomes apparent that the creation of the service descriptors takes up almost the entire provisioning time. This is particularly remarkable because the generation of V2 descriptors requires much less cryptographic effort than V3 descriptors. The first assumption to explain this discrepancy was that the cryptography used in V2 performed poorly on the measurement system. Since at the time of descriptor creation the RSA key pair of the V2 service is already generated, this cannot have any impact on the duration of the descriptor creation. The cryptographic operations performed during the creation of the descriptors are the hashing of values and the signing of the descriptor. For hashing, the outdated hash algorithm *SHA1* is used, whose execution on the measurement system took place in less than 0.0001 seconds. The signature is an RSA signature that uses the *PKCS.1* padding. Such a signature could be created on the measurement system in about 0.59 seconds on average. These results suggest that the cryptographic operations are not decisive for the long creation time of the descriptors of V2 onion service. The real reason for the significantly increased descriptor creation times was finally found in the source code. The first time a descriptor for a V2 onion service is published, a random delay is added, to ensure that the descriptor is stable before it is uploaded. The minimal delay is 30 seconds, which explains the clustering of descriptor creation times at around 31 seconds. Nevertheless, the V2

descriptors required less upload time, which is most likely due to their reduced size compared to the V3 descriptors. [39]

It was anticipated that the provisioning times would not differ drastically between the different V3 services. In principle, this can be confirmed on the basis of the measurement results, but there were significant differences between non-ephemeral onion services and onion services that use the *Vanguards* plugin. This becomes particularly evident in the comparison of introduction point establishment times, where V3 services that used *vanguards* required about twice the time of non-ephemeral onion services. This is probably due to the fact that introduction circuits are extended by one node when using the *Vanguards* plugin. Surprisingly, services using *Vanguards* built their descriptors faster than the rest of the variants. This seemed initially unreasonable, since there is no difference in how the descriptors are built. In order to provide an explanation for this phenomenon, a number of measurements were performed with different numbers of introduction points. These measurements yielded interesting results, as measurements that used more introduction points took longer to establish their introduction points, but completed the creation of their descriptor at the same time as services that used fewer introduction points. This implies that parts of the descriptors are created in parallel to the establishment of the introduction points, which in turn indicates that the number of introduction points has no effect on the overall provisioning time. As expected, the descriptor upload took the longest with the services that use *Vanguards*, as the upload circuits also contain an additional hop.

A remarkable finding in the results was the high standard deviation of the upload duration compared to the durations of the other sub-steps. The bar plot in Figure 5.11 revealed that a large proportion of the descriptor uploads are completed within a reasonable time frame (1-10 seconds), whereas several of the descriptor uploads required more than 100 seconds. To determine the cause of this behaviour, a further measurement was conducted that allows the distinction between the upload circuit creation time and the mere HTTP upload time. An analysis of the results of these measurements showed that, contrary to expectations, the upload circuit creation was the main issue. A further source code analysis revealed that there is a 100 second timeout for the creation of circuits to the *HSDirs*, after which a new connection attempt is made, which is the reason for the delay.

6 Conclusion

This chapter provides a summary of the main findings of this thesis and their implications for Tor onion services. As this is a very broad field of research and the time limitations of a Master's thesis do not allow to investigate all aspects in detail, several open research topics are noted in the future work section.

6.1 Summary

The functional principles of Hidden Service Directories (HSDirs) provide adversaries with the ability to determine when and how often a service is accessed if the attacker is able to control one or more HSDirs of that service. Combining this information with temporal metadata from an additional source (e.g. the log data of a company's network traffic) could allow the identification of the service client or at least its device. As the behaviour of the HSDirs does not in itself constitute a vulnerability, it cannot be assumed that their functional principle will be changed in the foreseeable future. This is a particular threat to those who use the Tor network and Tor onion service to protect their identity against government repression and surveillance.

This thesis aimed to address this problem by developing a protocol that reduces and, if possible, prevents the collection of metadata at the HSDirs. The fundamental problem consists in the fact that a service can be uniquely identified by its onion address and thus tracked by a potential adversary. In order to dissolve the relation between the service and its address, a two-tier addressing concept was developed. This concept allows clients to issue temporal limited individual onion addresses for a service via a publicly known service address referred to as introduction address. An adversary is not able to establish a connection between this individual address and an onion service, provided that the client does not share the address. However, this only applies to service requests issued using the individual onion address. The initial request to the public introduction address constitutes the same threat to anonymity as it would without the protocol. To prove the feasibility of the proposed protocol, a Proof of Concept (PoC) was implemented in Python for Hypertext Transfer Protocol (HTTP) services.

Since the latency of a service is a crucial factor in terms of its usability, a comprehensive performance analysis was carried out, measuring the provisioning time of onion services in different variants, as this is the most time-consuming component of the protocol. The results reveal that the latency involved in issuing an individual onion addresses is approximately 6.35 seconds. According to the previous investigation on latency acceptance of Wendolsky et al. [6], this latency exceeds the acceptable limit, which would lead to the protocol not being used. However, this limitation does not apply to the usage of the issued individual onion addresses, whose latency is significantly below the limit. This implies that the usage of the protocol is in principle feasible, but the commissioning process of the individual onion addresses must be expedited.

The use of the protocol minimizes the ability of collecting metadata on the HSDirs. However, the proposed protocol has two known limitations:

1. Each client must cache the individual addresses issued to it in order to avoid the necessity of making a request to the introduction address every time it accesses an onion service as this would eliminate the advantage of non-linkable addresses. Nevertheless, this has the effect that the cache of the client allows a complete associability of individual addresses to the introduction addresses, which an adversary could use to monitor which services a client accesses. This trade-off is necessary and the risk of compromise is considered low as the cache is under the control of the client in the current protocol proposal.
2. The approach does not scale. As a new onion service must be provisioned for each individual address, the number of addresses is limited. One possible solution to this problem is to no longer provide each user with their individual address, but to randomly select an address from a pool of addresses.
3. The additional latency introduced through the use of UOSP leads to an average Round Trip Time of 6.35 seconds, which exceeds the acceptable latency limit of 4 seconds proposed by Wendolsky et al. [6]. This circumstance is mitigated in practical use, as the caching of the UOSP data addresses on the proxy side considerably reduces this Round Trip Time.

6.2 Future Work

The following research topics regarding the reduction of metadata when accessing onion services remained unprocessed due to temporal limitations of the master thesis:

- **Quantitative Onion Service Anonymity Measurement**

In order to be able to better compare and classify the improvements of the proposed protocol or any changes to the onion service specification, it is necessary to develop a method to quantify the anonymity of onion services.

- **Scalability Issues**

The current protocol proposal suffers from scaling issues, as a new onion service must be provisioned for each individual onion address, which limits the number of possible clients and addresses. A possible solution is to no longer issue an individual address to each user, but to randomly choose an already provided address from a pool. For example, this could be implemented with *OnionBalance* [59], which combines introduction points of several services in a super descriptor, which are then randomly selected by the client.

- **Pool Size**

In case a solution with a pool of onion addresses is considered, the size of this pool and its impact on the anonymity of the clients has to be investigated.

- **Validity Period**

Currently, the validity period of individual onion addresses is set to 24 hours. In order to determine an ideal period, the impact of the validity period on the anonymity of onion services needs further investigation.

References

- [1] Christina Gardner and Donald L. Amoroso. ‘Development of an instrument to measure the acceptance of internet technology by consumers’. In: *Proceedings of the Hawaii International Conference on System Sciences* 37.C (2004), pp. 4143–4152. ISSN: 10603425. DOI: 10.1109/hicss.2004.1265623.
- [2] Andrew S Tanenbaum and David J Wetherall. *Computer Networks*. 5th. Upper Saddle River, NJ: Prentice Hall Press, 2010. ISBN: 0132126958.
- [3] J. Postel. *Transmission Control Protocol*. USA, Sept. 1981. DOI: 10.17487/rfc0793. URL: <https://www.rfc-editor.org/info/rfc0793>.
- [4] Roger Dingledine, Nick Mathewson and Paul Syverson. ‘Tor: The second-generation onion router’. In: *Proceedings of the 13th Conference on USENIX Security Symposium - Volume 13*. San Diego, CA: USENIX Association, 2004, p. 21. DOI: 10.5555/1251375.1251396.
- [5] Stefan Köpsell. ‘Low Latency Anonymous Communication – How Long Are Users Willing to Wait?’ In: *Emerging Trends in Information and Communication Security*. Ed. by Günter Müller. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 221–237. ISBN: 978-3-540-34642-5.
- [6] Rolf Wendolsky, Dominik Herrmann and Hannes Federrath. ‘Performance Comparison of Low-Latency Anonymisation Services from a User Perspective’. In: *Privacy Enhancing Technologies*. Ed. by Nikita Borisov and Philippe Golle. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 233–253. ISBN: 978-3-540-75551-7.
- [7] Ana Custura et al. *Onionperf*. 2020. URL: <https://gitlab.torproject.org/tpo/metrics/onionperf> (visited on 03/12/2020).
- [8] Ana Custura, Iain Learmonth and Gorry Fairhurst. ‘Measuring mobile performance in the tor network with onionperf’. In: *TMA 2019 - Proceedings of the 3rd Network Traffic Measurement and Analysis Conference* (2019), pp. 233–238. DOI: 10.23919/TMA.2019.8784601.
- [9] Andriy Panchenko, Lexi Pimenidis and Johannes Renner. ‘Performance analysis of anonymous communication channels provided by tor’. In: *ARES 2008 - 3rd International Conference on Availability, Security, and Reliability, Proceedings* (2008), pp. 221–228. DOI: 10.1109/ARES.2008.63.
- [10] Karsten Loesing et al. ‘Performance measurements and statistics of Tor hidden services’. In: *Proceedings - 2008 International Symposium on Applications and the Internet, SAINT 2008* (2008), pp. 1–7. DOI: 10.1109/SAINT.2008.69.
- [11] Christian Wilms. ‘Improving the Tor Hidden Service Protocol Aiming at Better Performances’. Master Thesis. Otto-Friedrich Universität Bamberg, 2008.
- [12] Christian Wilms. ‘Performance Evaluation of Tor Hidden Services’. In: (2007).
- [13] Helge Rausch. *Puppetor*. 2013. URL: <https://github.com/tsujigiri/puppetor> (visited on 05/12/2020).

- [14] Jörg Lenhard, Karsten Loesing and Guido Wirtz. ‘Performance measurements of tor hidden services in low-bandwidth access networks’. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 5536 LNCS.March (2009), pp. 324–341. ISSN: 03029743. DOI: 10.1007/978-3-642-01957-9_20.
- [15] David Goulet. *Onion Service version 2 deprecation timeline*. 2020. URL: <https://blog.torproject.org/v2-deprecation-timeline?page=1> (visited on 05/12/2020).
- [16] Lasse Øverlier and Paul Syverson. ‘Locating hidden servers’. In: *Proceedings - IEEE Symposium on Security and Privacy* 2006 (2006), pp. 100–114. ISSN: 10816011. DOI: 10.1109/SP.2006.24.
- [17] Matthew K. Wright et al. ‘The predecessor attack’. In: *ACM Transactions on Information and System Security* 7.4 (2004), pp. 489–522. ISSN: 1094-9224. DOI: 10.1145/1042031.1042032.
- [18] Alex Biryukov, Ivan Pustogarov and Ralf Philipp Weinmann. ‘Trawling for tor hidden services: Detection, measurement, deanonymization’. In: *Proceedings - IEEE Symposium on Security and Privacy* (2013), pp. 80–94. ISSN: 10816011. DOI: 10.1109/SP.2013.15.
- [19] Gareth Owen and Nick Savage. ‘Empirical analysis of Tor hidden services’. In: *IET Information Security* 10.3 (2016), pp. 113–118. ISSN: 17518717. DOI: 10.1049/iet-ifs.2015.0121.
- [20] Juan A. Elices and Fernando Pérez-González. ‘Locating Tor hidden services through an interval-based traffic-correlation attack’. In: *2013 IEEE Conference on Communications and Network Security, CNS 2013* 2 (2013), pp. 385–386. DOI: 10.1109/CNS.2013.6682740.
- [21] Zhen Ling et al. ‘Protocol-level hidden server discovery’. In: *Proceedings - IEEE INFOCOM* (2013), pp. 1043–1051. ISSN: 0743166X. DOI: 10.1109/INFCOM.2013.6566894.
- [22] Muqian Chen et al. ‘SignalCookie: Discovering Guard Relays of Hidden Services in Parallel’. In: *2019 IEEE Symposium on Computers and Communications (ISCC)*. Vol. 2019-June. IEEE, June 2019, pp. 1–7. ISBN: 978-1-7281-2999-0. DOI: 10.1109/ISCC47284.2019.8969639. URL: <https://ieeexplore.ieee.org/document/8969639/>.
- [23] David M. Goldschlag, Michael G. Reed and Paul F. Syverson. ‘Hiding routing information’. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 1174 (1996), pp. 137–150. ISSN: 16113349. DOI: 10.1007/3-540-61996-8_37.
- [24] M Leech et al. *SOCKS Protocol Version 5*. USA, Mar. 1996. DOI: 10.17487/rfc1928. URL: <https://www.rfc-editor.org/info/rfc1928>.
- [25] Cristopher Baines et al. *Tor Rendezvous Specification - Version 3*. 2020. URL: <https://gitweb.torproject.org/torspec.git/tree/rend-spec-v3.txt>.
- [26] Daniel J Bernstein. ‘Curve25519: New Diffie-Hellman Speed Records’. In: *Public Key Cryptography - PKC 2006*. Ed. by Moti Yung et al. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 207–228. ISBN: 978-3-540-33852-9.
- [27] Tim Wilson-Brown et al. *Proposal 260 - Rendezvous Single Onion Services*. 2017. URL: <https://github.com/torproject/torspec/blob/master/proposals/260-rend-single-onion.txt>.

- [28] Nick Mathewson. *Proposal 220 - Migrate server identity keys to Ed25519*. 2013. URL: <https://gitweb.torproject.org/torspec.git/plain/proposals/220-ecc-id-keys.txt>.
- [29] Roger Dingledine and Nick Mathewson. *Tor Protocol Specification*. 2020. URL: <https://gitweb.torproject.org/torspec.git/plain/tor-spec.txt>.
- [30] Simon Josefsson and Ilari Liusvaara. *Edwards-Curve Digital Signature Algorithm (EdDSA)*. Tech. rep. Jan. 2017. DOI: 10.17487/RFC8032. URL: <https://tools.ietf.org/html/rfc8032><https://www.rfc-editor.org/info/rfc8032>.
- [31] Aaron Johnson et al. ‘Users get routed: Traffic correlation on Tor by realistic adversaries’. In: *Proceedings of the ACM Conference on Computer and Communications Security* (2013), pp. 337–348. ISSN: 15437221. DOI: 10.1145/2508859.2516651.
- [32] M Wright et al. ‘Defending anonymous communications against passive logging attacks’. In: *2003 Symposium on Security and Privacy, 2003*. 2003, pp. 28–41. DOI: 10.1109/SECPRI.2003.1199325.
- [33] George Kadianakis et al. *Tor directory protocol*. 2021. URL: <https://gitweb.torproject.org/torspec.git/tree/dir-spec.txt>.
- [34] Isis Lovecruft et al. *Tor Guard Specification*. 2020. URL: <https://gitweb.torproject.org/torspec.git/plain/guard-spec.txt>.
- [35] George Kadianakis. *Announcing the Vanguard Add-On for Onion Services*. 2018. URL: <https://blog.torproject.org/announcing-vanguards-add-onion-services> (visited on 13/01/2021).
- [36] George Kadianakis and Mike Perry. *Mesh-based vanguards*. 2021. URL: <https://gitweb.torproject.org/torspec.git/commit/proposals/292-mesh-vanguards.txt>.
- [37] Andreas Pfitzmann and Marit Hansen. *A terminology for talking about privacy by data minimization: Anonymity, Unlinkability, Undetectability, Unobservability, Pseudonymity, and Identity Management*. Aug. 2010. URL: http://dud.inf.tu-dresden.de/literatur/Anon_Terminology_v0.34.pdf.
- [38] Mike Perry. *Tor’s Open Research Topics: 2018 Edition*. 2018. URL: <https://blog.torproject.org/tors-open-research-topics-2018-edition> (visited on 11/12/2020).
- [39] Roger Dingledine, Nick Mathewson and Paul Syverson. *Tor Rendezvous Specification - Version 2*. 2019. URL: <https://gitweb.torproject.org/torspec.git/tree/rend-spec-v2.txt>.
- [40] Alec Muffett. *1 Million People use Facebook over Tor*. 2016. URL: <https://www.facebook.com/notes/facebook-over-tor/1-million-people-use-facebook-over-tor/865624066877648/> (visited on 17/12/2020).
- [41] Tor Metrics. *Unique .onion addresses (version 2 only)*. 2020. URL: <https://metrics.torproject.org/hidserv-dir-onions-seen.csv> (visited on 17/12/2020).
- [42] Andy Yen. *ProtonMail*. 2020. URL: <https://protonmail.com> (visited on 19/12/2020).
- [43] Eric Rescorla. *The Transport Layer Security (TLS) Protocol Version 1.3*. USA, Aug. 2018. DOI: 10.17487/RFC8446. URL: <https://www.rfc-editor.org/info/rfc8446>.

- [44] Gerard J Holzmann and William Slattery Lieberman. *Design and Validation of Computer Protocols*. Vol. 94. New Jersey: Prentice hall Englewood Cliffs, 1991, p. 500. ISBN: 0135399254.
- [45] Alex Biryukov et al. ‘Content and popularity analysis of tor hidden services’. In: *Proceedings - International Conference on Distributed Computing Systems* 30-June-20 (2014), pp. 188–193. DOI: 10.1109/ICDCSW.2014.20. arXiv: 1308.6768.
- [46] Damian Johnson. *Stem*. 2020. URL: <https://stem.torproject.org/> (visited on 01/01/2021).
- [47] Doug Hellmann. *SocketServer - Creating network servers*. 2010. URL: <https://bip.weizmann.ac.il/course/python/PyMOTW/PyMOTW/docs/SocketServer/index.html> (visited on 01/01/2021).
- [48] Python Software Foundation. *socketserver - A framework for network servers*. 2021. URL: <https://docs.python.org/3/library/socketserver.html> (visited on 02/01/2021).
- [49] Zope Foundation. *ZODB - a native object database for Python*. 2016. URL: <http://www.zodb.org/en/latest/> (visited on 02/01/2021).
- [50] Tal Yarkoni. *Transitions*. 2020. URL: <https://github.com/pytransitions/transitions> (visited on 02/01/2021).
- [51] Thomas Raab. *Prototype Unlinkability*. 2021. URL: <https://git.ins.jku.at/proj/digidow/prototype-unlinkability> (visited on 13/02/2021).
- [52] Benoit Chesneau. *http-parser*. 2020. URL: <https://github.com/benoitc/http-parser/> (visited on 03/01/2021).
- [53] Tor Project. *Configuration file for a typical Tor user*. 2019. URL: <https://raw.githubusercontent.com/torproject/tor/master/src/config/torrc.sample.in> (visited on 04/01/2021).
- [54] Python Software Foundation. *http.server - HTTP servers*. 2021. URL: <https://docs.python.org/3/library/http.server.html> (visited on 05/01/2021).
- [55] Thomas Raab. *Prototype Timing-Analysis*. 2021. URL: <https://git.ins.jku.at/proj/digidow/prototype-timing-analysis> (visited on 13/02/2021).
- [56] Docker Inc. *Docker overview*. 2020. URL: <https://docs.docker.com/get-started/overview/> (visited on 17/01/2021).
- [57] Thomas Raab. *tor-docker*. 2021. URL: <https://git.ins.jku.at/students/tor-docker> (visited on 14/02/2021).
- [58] John M Chambers et al. *Graphical Methods for Data Analysis*. London: Chapman and Hall/CRC, Jan. 2018. ISBN: 9781351072304. DOI: 10.1201/9781351072304.
- [59] George Kadianakis. *Cooking with Onions: Finding the Onionbalance*. 2016. URL: <https://blog.torproject.org/cooking-onions-finding-onionbalance> (visited on 07/02/2021).

Appendix A: Unlinkable Onion Service Protocol

A.1 Database Utils

```

1 import persistent
2 import persistent.list
3 from datetime import datetime, timezone
4 import pytz
5 import logging
6 from ZODB import DB, FileStorage
7
8 utc = pytz.UTC
9
10 TIME_FORMAT = '%Y-%m-%d %H:%M:%S'
11
12 """
13 Class for connecting to the ZODB database file
14 """
15
16
17 class MyZODB(object):
18     def __init__(self, path):
19         self.storage = FileStorage.FileStorage(path)
20         self.db = DB(self.storage)
21         self.connection = self.db.open()
22         self.dbroot = self.connection.root()
23
24     def close(self):
25         self.connection.close()
26         self.db.close()
27         self.storage.close()
28
29
30 """
31 Data representation for storing and managing custom UOSP data addresses on the server
32 ↪ side
33 """
34
35 class UospDataAddresses(persistent.Persistent):
36     def __init__(self):
37         self.__uosp_data_addresses = persistent.mapping.PersistentMapping()
38
39     """
40     Add a custom UOSP data address to the servers database
41
42     :param uosp_data_address: Custom UOSP data address to be added to the servers
43     ↪ database
44     :param uosp_validity: Validity of the newly provisioned address
45     """
46
47     def add_data_address(self, uosp_data_address, uosp_validity):
48         self.__uosp_data_addresses[uosp_data_address] = uosp_validity
49         logging.info(f"Added data address {uosp_data_address}")
50
51     """
52     Remove a custom UOSP data address from the servers database

```

```

53     :param uosp_data_address: Custom UOSP data address to be removed from the servers
↪ database
54     """
55
56     def remove_data_address(self, uosp_data_address):
57         if uosp_data_address in self.__uosp_data_addresses:
58             del self.__uosp_data_addresses[uosp_data_address]
59             logging.info(f"Removed data address {uosp_data_address}")
60
61     """
62     Check the validity of a custom UOSP data address
63
64     :param uosp_data_address: Custom UOSP data address whose validity shall be
↪ checked
65     :return: True, iff the custom address provided is still valid, false otherwise
66     """
67
68     def __check_validity(self, uosp_data_address):
69         if uosp_data_address in self.__uosp_data_addresses:
70             # get time stamp
71             uosp_validity = self.__uosp_data_addresses[uosp_data_address]
72             # compare time stamps
73             if datetime.now(timezone.utc) > uosp_validity:
74                 return False
75             else:
76                 return True
77
78     """
79     Revise the validity of all currently provisioned custom UOSP addresses
80
81     :param controller: Handle for the stem tor controller
82     """
83
84     def revise_validity(self, controller):
85         for uosp_address in self.__uosp_data_addresses.copy():
86             if not self.check_validity(uosp_address):
87                 # remove invalid entries from the database
88                 self.remove_data_address(uosp_address)
89                 # discontinue invalid services (tear them down)
90                 controller.remove_ephemeral_hidden_service(uosp_address)
91
92     """
93     Data representation for storing and managing custom UOSP data addresses on the client
↪ side
94     """
95
96
97
98     class UospIntroAddress(persistent.Persistent):
99
100         def __init__(self, uosp_intro_address):
101             self.__uosp_intro_address = uosp_intro_address
102             self.__uosp_data_addresses = persistent.list.PersistentList()
103
104     """
105     Add a custom UOSP data address to the proxy's database
106
107     :param uosp_data_address: Custom UOSP data address to be added to the proxy's
↪ database
108     :param uosp_validity: Validity of the newly provisioned address

```

```

109     """
110
111     def add_data_address(self, uosp_data_address, uosp_validity):
112         self.__uosp_data_addresses.append([uosp_data_address, uosp_validity])
113         logging.info(f"Added data address {uosp_data_address} for intro address
114             ↪ {self.__uosp_intro_address}")
115
116     """
117     Remove a custom UOSP data address from the proxy's database
118
119     :param uosp_data_address: Custom UOSP data address to be removed from the
120     ↪ proxy's database
121     """
122
123     def remove_data_address(self, uosp_data_address):
124         for address_entry in self.__uosp_data_addresses:
125             if uosp_data_address == address_entry[0]:
126                 self.__uosp_data_addresses.remove(address_entry)
127                 logging.info(f"Removed data address {uosp_data_address} for intro
128                 ↪ address {self.__uosp_intro_address}")
129
130     """
131     Revise the validity of all currently provisioned custom UOSP addresses of the
132     ↪ objects intro address
133     """
134
135     def revise_validity(self):
136         for uosp_address in self.__uosp_data_addresses:
137             if not self.__check_validity(uosp_address[0]):
138                 self.remove_data_address(uosp_address[0])
139
140     """
141     Get the index of a custom UOSP address in the list of addresses for the objects
142     ↪ intro address
143
144     :param uosp_data_address: Custom UOSP address
145     :return: Index of the entry if available, None otherwise
146     """
147
148     def __get_address_index(self, uosp_data_address):
149         for uosp_address in self.__uosp_data_addresses:
150             if uosp_address[0] == uosp_data_address:
151                 return self.__uosp_data_addresses.index(uosp_address)
152         return None
153
154     """
155     Check the validity of a custom UOSP data address
156
157     :param uosp_data_address: Custom UOSP data address whose validity shall be
158     ↪ checked
159     :return: True, iff the custom address provided is still valid, false otherwise
160     """
161
162     def __check_validity(self, uosp_data_address):
163         index = self.__get_address_index(uosp_data_address)
164         if index is not None:
165             uosp_validity = datetime.strptime(self.__uosp_data_addresses[index][1],
166                 ↪ TIME_FORMAT)
167             uosp_validity = utc.localize(uosp_validity)
168             # compare time stamps+

```

```

162         if datetime.now(timezone.utc) > uosp_validity:
163             return False
164         else:
165             return True
166
167         """
168         Revise the validity of all custom UOSP addresses of the objects intro address and
↪ return the latest valid entry
169
170         :return: Latest valid custom UOSP address for the objects intro address
171         """
172
173     def get_latest_valid_data_address(self):
174         self.revise_validity()
175         if len(self.__uosp_data_addresses) > 0:
176             # return last = latest entry
177             return self.__uosp_data_addresses[-1]
178
179         """
180         Return the count of custom UOSP addresses for the objects into address
181         """
182
183     def get_data_address_count(self):
184         return len(self.__uosp_data_addresses)

```

A.2 Connection Utils

```

1 import sys
2 import stem.control
3 import stem.connection
4 import logging
5
6 """
7 Establish a connection to the control port of tor and authenticate
8 """
9 CONTROL_PORT = 9051
10 CONTROL_ADDRESS = '127.0.0.1'
11
12 logging.basicConfig(level=logging.INFO)
13 logging.getLogger("stem").setLevel(logging.ERROR)
14
15
16 def establish_connection(address=CONTROL_ADDRESS, port=CONTROL_PORT, password=None):
17     # get connection to tor via the control port
18     try:
19         # the explicit declaration of the port isn't necessary since it's the default
↪ port
20         controller = stem.control.Controller.from_port(address, port)
21     except stem.SocketError as exc:
22         logging.error(f"Unable to connect to port control port 9051 ({exc})")
23         sys.exit(1)
24
25     if controller:
26         # authenticate
27         try:
28             if password != None:
29                 controller.authenticate(password=password)
30             else:

```

```

31         controller.authenticate()
32         logging.info("Connection established and authenticated")
33         return controller
34     except stem.connection.PasswordAuthFailed:
35         logging.error("Unable to authenticate, password is incorrect")
36         sys.exit(1)
37     except stem.connection.AuthenticationFailure as exc:
38         logging.error(f" * Unable to authenticate: {exc}")
39         sys.exit(1)
40
41
42 """
43 Create new ephemeral Hidden Service and return its onion id
44 """
45
46
47 def create_ephemeral_hidden_service(controller, source_port, target_port, version):
48     try:
49         if version == 2:
50             hs = controller.create_ephemeral_hidden_service({source_port:
51                 ↪ target_port}, await_publication=False,
52                                                         key_type="NEW",
53                                                         ↪ key_content="RSA1024",
54                                                         ↪ detached=True)
55
56         else:
57             hs = controller.create_ephemeral_hidden_service({source_port:
58                 ↪ target_port}, await_publication=False,
59                                                         detached=True,
60                                                         ↪ key_type='NEW',
61                                                         ↪ key_content='BEST')
62
63         if hs.service_id:
64             return hs.service_id
65         else:
66             return None
67
68     except stem.ControllerError as exc:
69         logging.error(f"Unable to create the hidden service ({exc})")
70         sys.exit(1)
71
72 def create_hidden_service(controller, path, port):
73     try:
74         controller.create_hidden_service(path, port, target_address='127.0.0.1')
75     except stem.ControllerError as exc:
76         logging.error(f"Unable to create the hidden service ({exc})")
77         sys.exit(1)

```

A.3 Server State Machine

```

1 from transitions import Machine
2 import logging
3
4
5 class UOSPServerStateMachine(object):
6     states = ['initial_state', 'client_req_received', 'legacy_http_resp_sent',
7             ↪ 'uossp_resp_sent', 'data_resp_sent',
8             ↪ 'uossp_ack_received', 'server_fin_sent', 'ERROR']

```

```

9     # Functions
10
11     def __init__(self):
12         logging.getLogger('transitions').setLevel(logging.ERROR)
13
14         # Initialize the state machine
15         self.machine = Machine(model=self, states=UOSPServerStateMachine.states,
16                               ↪ initial='initial_state')
17
18         # Transitions
19
20         # s0 -> s1
21         self.machine.add_transition('receive_request', source='initial_state',
22                                   ↪ dest='client_req_received')
23
24         # s1 -> s2
25         self.machine.add_transition('send_legacy_response',
26                                   ↪ source='client_req_received', dest='legacy_http_resp_sent')
27
28         # s1 -> s3
29         self.machine.add_transition('send_uosp_response',
30                                   ↪ source='client_req_received', dest='uosp_resp_sent')
31
32         # s1 -> s4
33         self.machine.add_transition('send_data_response',
34                                   ↪ source='client_req_received', dest='data_resp_sent')
35
36         # s3 -> s5
37         self.machine.add_transition('receive_uosp_ack', source='uosp_resp_sent',
38                                   ↪ dest='uosp_ack_received')
39
40         # s5 -> ERR
41         self.machine.add_transition('pub_key_malformed', source='uosp_ack_received',
42                                   ↪ dest='ERROR')
43
44         # s5 -> s6
45         self.machine.add_transition('send_server_fin', source='uosp_ack_received',
46                                   ↪ dest='server_fin_sent')

```

A.4 Server Module

```

1 from stem.control import EventType, Controller
2 from utils.connect import establish_connection, create_ephemeral_hidden_service
3 from http.server import BaseHTTPRequestHandler, HTTPServer
4 from socketserver import ForkingMixIn
5 import time
6 import hashlib
7 import uosp_state_machine_server
8 import logging
9 import signal
10 import sys
11 from datetime import datetime, timezone, timedelta
12 from utils.db_util import UospDataAddresses, MyZODB
13 import transaction
14 import threading
15 import ssl
16
17 # password for control port authentication

```



```
18 PASSWORD = 'tor'
19 # uosp version
20 UOSP_VERSION = '0.1'
21 VALIDITY_MINUTES = 1440
22 VALIDITY_REVISION_MINUTES = 10
23
24 DB_PATH = 'data_addresses.fs'
25
26 # connect and authenticate
27 intro_addresses = {}
28 controller = None
29
30
31 class ForkedHTTPServer(ForkingMixIn, HTTPServer):
32     """Handle requests in a separate process"""
33
34
35 class RequestHandler(BaseHTTPRequestHandler):
36     uosp_address = None
37     uosp_validity = None
38     close_connection = False
39     controller = None
40
41     # server state machine
42     ssm = uosp_state_machine_server.UOSPServerStateMachine()
43
44     def do_GET(self):
45         # trigger for uosp data protocol header
46         uosp_data = False
47
48         # state changed from initial to request received
49         state = self.ssm.state
50         self.ssm.receive_request()
51         logging.info('State changed from %s to %s' % (state, self.ssm.state))
52
53         # don't close connection after first response
54         self.close_connection = False
55
56         # check if UOSP header is set and version is supported
57         url = self.headers.get('Host').split('.')[0]
58         state = self.ssm.state
59         if self.headers.get('UOSP') and self.headers.get('UOSP_VERSION') == '0.1':
60             if url in intro_addresses:
61                 # UOSP
62                 if self.controller is None:
63                     controller = establish_connection(password=PASSWORD)
64
65                 message = create_ephemeral_hidden_service(controller, 80, 80, 3)
66                 self.uosp_validity = datetime.now(timezone.utc) +
67                     → timedelta(minutes=VALIDITY_MINUTES)
68                 self.uosp_address = message
69
70                 db = MyZODB(DB_PATH)
71                 # no db entry yet
72                 if len(db.dbroot) == 0:
73                     db.dbroot['addresses'] = UospDataAddresses()
74                 db.dbroot['addresses'].add_data_address(self.uosp_address,
75                     → self.uosp_validity)
76                 transaction.commit()
77                 db.close()
```

```
76
77     self.ssm.send_uosp_response()
78     else:
79         # UOSP, but no into URL -> deliver web page
80         self.ssm.send_data_response()
81         message = "<HTML><HEAD><TITLE>Pretty Web
82         ↪ Page</TITLE></HEAD><BODY><h3>Web Page</h3>This is the web " \
83         ↪ "page the client intended to reach</BODY></HTML>"
84         uosp_data = True
85
86     else:
87         # No UOSP, deliver legacy page
88         self.ssm.send_legacy_response()
89         message = "<HTML><HEAD><TITLE>UOSP
90         ↪ Response</TITLE></HEAD><BODY><h3>UOSP</h3>This page supports the " \
91         ↪ "Unlinkable Onion Service Protocol (UOSP). When you receive this
92         ↪ message, your client/browser " \
93         ↪ "doesn't support this protocol. </BODY></HTML>"
94         uosp_data = True
95
96     logging.info('State changed from %s to %s' % (state, self.ssm.state))
97
98     self.protocol_version = "HTTP/1.1"
99     self.send_response(200)
100    self.send_header("Content-Length", str(len(message)))
101    self.send_header("UOSP_DATA", str(uosp_data))
102    self.send_header("UOSP", 'True')
103    self.send_header("UOSP_VERSION", UOSP_VERSION)
104
105    # add validity if intro url used
106    if self.uosp_validity is not None:
107        self.send_header("UOSP_VALIDITY", self.uosp_validity.strftime("%Y-%m-%d
108        ↪ %H:%M:%S"))
109    self.end_headers()
110    self.wfile.write(bytes(message, "utf8"))
111    return
112
113    def do_POST(self):
114        self.close_connection = True
115        content_len = int(self.headers.get('Content-Length'))
116        post_body = self.rfile.read(content_len).decode()
117
118        state = self.ssm.state
119        self.ssm.receive_uosp_ack()
120        logging.info('State changed from %s to %s' % (state, self.ssm.state))
121
122        # calculate hash value
123        hash_pub = hashlib.sha256()
124        hash_pub.update(self.uosp_address.encode())
125        hash_pub = hash_pub.hexdigest()
126
127        state = self.ssm.state
128        if str(hash_pub) == post_body:
129            message = "FIN"
130            self.ssm.send_server_fin()
131        else:
132            self.ssm.pub_key_malformed()
133            message = "ERROR"
134
135        logging.info('State changed from %s to %s' % (state, self.ssm.state))
```

```
132
133     self.protocol_version = "HTTP/1.1"
134     self.send_response(200)
135     self.send_header("Content-Length", str(len(message)))
136     self.send_header("UOSP", 'True')
137     self.send_header("UOSP_DATA", 'FALSE')
138     self.send_header("UOSP_VERSION", UOSP_VERSION)
139     self.end_headers()
140     self.wfile.write(bytes(message, "utf8"))
141     return
142
143
144 """
145 Called when the service shuts down to delete all created ephemeral onion services
146 """
147
148
149 def exit_gracefully(signum, frame):
150     service_list = controller.list_ephemeral_hidden_services(default=None,
151     ↪ detached=True)
152     for service in service_list:
153         controller.remove_ephemeral_hidden_service(service)
154         logging.info(f"Service with uri {str(service)} deleted")
155     logging.info("Shutting down server")
156     sys.exit(0)
157
158 def revise_uoasp_data_addresses(ctrl):
159     while True:
160         time.sleep(VALIDITY_REVISION_MINUTES * 60)
161         db = MyZODB(DB_PATH)
162         if 'addresses' in db.dbroot:
163             db.dbroot['addresses'].revise_validity(ctrl)
164         logging.info("Revised UOSP data addresses")
165         transaction.commit()
166         db.close()
167
168
169 if __name__ == '__main__':
170     controller = establish_connection(password=PASSWORD)
171     service = create_ephemeral_hidden_service(controller, 80, 80, 3)
172     intro_addresses[str(service)] = datetime.now(timezone.utc)
173     logging.info(f"Service with uri {str(service)} created")
174     # start thread that periodically revises the validity of provisioned custom onion
175     ↪ services
176     threading.Thread(target=revise_uoasp_data_addresses, args=(controller,),
177     ↪ daemon=True).start()
178     ForkedHTTPServer.allow_reuse_address = True
179     signal.signal(signal.SIGINT, exit_gracefully)
180     server = ForkedHTTPServer(('localhost', 80), RequestHandler)
181     server.socket = ssl.wrap_socket(
182     ↪ server.socket,
183     ↪ keyfile="key.pem",
184     ↪ certfile='cert.pem',
185     ↪ server_side=True)
186     server.serve_forever()
```

A.5 Server Tor Configuration

```

1  ## Logging
2  Log notice file /var/log/tor/notices.log
3  Log info file /var/log/tor/info.log
4  Log debug file /var/log/tor/debug.log
5
6  ## Control port
7  ControlPort 9051
8  HashedControlPassword 16:0717EF71CDF228DE60E93242A4AFE681FB525080DC1A330A103FAD9782
9
10 ## Run as daemon in background
11 RunAsDaemon 1

```

A.6 Proxy State Machine

```

1  from transitions import Machine
2  import logging
3
4  class UOSPClientStateMachine(object):
5      states = ['initial_state', 'cache_addr_used', 'client_req_sent',
6              ↪ 'server_pub_key_received', 'client_pub_ack_sent',
7              ↪ 'server_fin_received', 'data_resp_received', 'ERROR']
8
9      # Functions
10     def __init__(self):
11         logging.getLogger('transitions').setLevel(logging.ERROR)
12
13         # Initialize the state machine
14         self.machine = Machine(model=self, states=UOSPClientStateMachine.states,
15                               ↪ initial='initial_state')
16
17         # Transitions
18
19         # s0 -> s1
20         self.machine.add_transition('send_request', source='initial_state',
21                                   ↪ dest='client_req_sent')
22
23         # s0 -> s6
24         self.machine.add_transition('use_cache_addr', source='initial_state',
25                                   ↪ dest='cache_addr_used')
26
27         # s1 -> s2
28         self.machine.add_transition('receive_pub_key', source='client_req_sent',
29                                   ↪ dest='server_pub_key_received')
30
31         # s2 -> s3
32         self.machine.add_transition('send_uosp_ack',
33                                   ↪ source='server_pub_key_received', dest='client_pub_ack_sent')
34
35         # s2 -> ERR
36         self.machine.add_transition('pub_key_malformed',
37                                   ↪ source='server_pub_key_received', dest='ERROR')
38
39         # s3 -> ERR
40         self.machine.add_transition('server_hash_invalid',
41                                   ↪ source='client_pub_ack_sent', dest='ERROR')

```

```

35     # s3 -> s4
36     self.machine.add_transition('receive_server_fin',
    ↪     source='client_pub_ack_sent', dest='server_fin_received')
37
38     # s1 -> s5
39     self.machine.add_transition('receive_data', source='client_req_sent',
    ↪     dest='data_resp_received')

```

A.7 Proxy Module

```

1  import logging
2  import select
3  import socket
4  from socketserver import ForkingMixIn, TCPServer, StreamRequestHandler
5  import tldextract
6  import socks
7  import validators
8  from utils.db_util import UospIntroAddress, MyZODB
9  import transaction
10
11 import uosp_state_machine_client
12 import hashlib
13
14 import ssl
15
16 # try to import C parser then fallback in pure python parser.
17 try:
18     from http_parser.parser import HttpParser
19 except ImportError:
20     from http_parser.pyparser import HttpParser
21
22 logging.basicConfig(level=logging.INFO)
23
24 # running tor instance
25 UPSTREAM_HOST = '127.0.0.1'
26 UPSTREAM_PORT = 9050
27
28 # proxy configuration
29 PROXY_HOST = '0.0.0.0'
30 PROXY_PORT = 9011
31
32 UOSP_VERSION = '0.1'
33
34 DB_PATH = 'intro_addresses.fs'
35
36
37 # base class for threaded TCP server
38 class ForkingTCPServer(ForkingMixIn, TCPServer):
39     pass
40
41
42 """
43 For every new client connection, a new thread is started , which instantiates the
    ↪ SocksProxy class as handler for its requests
44 """
45
46
47 class SocksProxy(StreamRequestHandler):

```

```

48     # client state machine
49     csm = uosp_state_machine_client.UOSPClientStateMachine()
50
51     # placeholder variables which are used to store the received intro host and
52     ↪ temporary uosp address received from the server
53     intro_host = None
54     uosp_address = None
55     uosp_validity = None
56     legacy_request = False
57     cache_request = False
58
59     """
60     Handle state changes
61
62     :param data Data received on a socket connection (can be request/response)
63     ↪ :param recv_socket Handle for the socket used to send a response to the remote
64     ↪ server (avoid separate socket connections for each message pair)
65     """
66
67     def handle_state_change(self, data, recv_socket=None):
68         # print(self.uosp_data_addresses)
69         # get current state
70         state = self.csm.state
71
72         # transition from initial state when the client request is received
73         if state == 'initial_state':
74             if not self.cache_request:
75                 self.csm.send_request()
76             else:
77                 self.csm.use_cache_addr()
78
79             logging.info('State changed from %s to %s' % (state, self.csm.state))
80
81         # transition from client_req sent to server_pub_key_received
82         if state == 'client_req_sent':
83             if data[0]:
84                 # this is a uosp data response
85                 self.csm.receive_data()
86             else:
87                 self.csm.receive_pub_key()
88                 logging.info('State changed from %s to %s' % (state, self.csm.state))
89
90                 # store received onion address for later use
91                 self.uosp_address = data[2]
92
93                 # the pub key was received, send its hash value back to the server to
94                 ↪ ensure error-free transmission
95                 if self.intro_host is not None:
96                     # calculate hash value
97                     hash_pub = hashlib.sha256()
98                     hash_pub.update(self.uosp_address.encode())
99                     hash_pub = hash_pub.hexdigest()
100
101                     request_string = 'POST /session HTTP/1.1\r\nHost: ' +
102                                     ↪ self.intro_host + '\r\nContent-Type:
103                                     ↪ text/plain\r\nContent-Length: ' + str(
104                                         len(hash_pub)) + '\r\nUOSP: True\r\nUOSP_VERSION: ' +
105                                     ↪ UOSP_VERSION + '\r\n\r\n' + hash_pub

```

```

101         # use the socket handle passed within the parameters to send the
102         ↪ response (avoid new socket connection for every message pair)
103         recv_socket.sendall(request_string.encode())
104         self.csm.send_uosp_ack()
105         logging.info('State changed from %s to %s' % (state,
106         ↪ self.csm.state))
107
108     # transition from client_pub_ack_sent to server_fin_received or ERROR
109     if state == 'client_pub_ack_sent':
110         if data[2] == 'FIN':
111             self.csm.receive_server_fin()
112             logging.info('State changed from %s to %s' % (state, self.csm.state))
113         else:
114             self.csm.server_hash_invalid()
115             logging.info('State changed from %s to %s' % (state, self.csm.state))
116
117     """
118     Receive HTTP request and parse its header
119
120     :return Dictionary containing the decoded relevant header fields
121     """
122
123     def parse_http_header(self):
124         p = HttpParser()
125         header = {}
126
127         # read data from socket
128         data = self.connection.recv(1024)
129
130         recved = len(data)
131         nparsed = p.execute(data, recved)
132         assert nparsed == recved
133
134         # header completely received
135         if p.is_headers_complete():
136             http_header = p.get_headers()
137
138             # extract host address
139             header['host'] = http_header['Host'].split(':')[0]
140
141             # extract suffix from host (must be a public ICANN TLD)
142             no_fetch_extract = tldextract.TLDExtract(suffix_list_urls=None)
143             header['suffix'] = no_fetch_extract(http_header['host']).suffix
144
145             # extract host port
146             if len(http_header['Host'].split(':')) > 1:
147                 header['port'] = http_header['Host'].split(':')[1]
148             else:
149                 header['port'] = 80
150
151         if len(header) > 0:
152             return header
153         else:
154             return None
155
156     """
157     Parse the body of a HTTP response
158
159     :param data The data is read from the socket, alternatively the already read
160     ↪ HTTP

```

```

158     message can be passed
159     :return List containing information whether its an uosp data response and the
↪ parsed decoded http body
160     """
161
162     def parse_http_body(self, sock_conn):
163         p = HttpParser()
164         body = b""
165         uosp_data = False
166         uosp = False
167
168         while True:
169             data = sock_conn.recv(1024)
170             if not data:
171                 break
172
173             recved = len(data)
174             nparsed = p.execute(data, recved)
175             assert nparsed == recved
176
177             if p.is_headers_complete():
178                 headers = p.get_headers()
179                 # check if uosp data response
180                 if 'UOSP' in headers:
181                     uosp = True
182
183                 if 'UOSP_DATA' in headers:
184                     if headers['UOSP_DATA'] == 'True':
185                         uosp_data = True
186                 if 'UOSP_VALIDITY' in headers:
187                     self.uosp_validity = headers['UOSP_VALIDITY']
188
189             if p.is_partial_body():
190                 body += (p.recv_body())
191
192             if p.is_message_complete():
193                 break
194
195             return [uosp_data, uosp, body.decode().rstrip()]
196
197     """
198     Function called when a new request is received
199     """
200
201     def handle(self):
202         logging.info('Accepting connection from %s:%s' % self.client_address)
203
204         # parse http header of the received request to extract host and port
205         http_header = self.parse_http_header()
206
207         if http_header is not None:
208             # only onion traffic triggers the uosp, normal http traffic is proxied
↪ normally
209             if http_header['port'] == 80 and http_header['suffix'] == 'onion':
210                 # store the initial host and port for later use (UOSP ACK)
211                 self.intro_host = str(http_header['host']) + ':' +
↪ str(http_header['port'])
212
213                 # check if uosp data adress is already in the cache db
214                 db = MyZODB(DB_PATH)

```



```

215
216     if self.intro_host.split('.')[0] in db.dbroot:
217         print(db.dbroot[self.intro_host.split('.')[0]])
218         db.dbroot[self.intro_host.split('.')[0]].revise_validity()
219
220     if self.intro_host.split('.')[0] not in db.dbroot or
221     ↪ db.dbroot[self.intro_host.split('.')[0]].get_data_address_count()
222     ↪ < 1:
223         logging.info("No cached uosp data address for intro address " +
224         ↪ self.intro_host.split('.')[0])
225         # create socket connection to requested resource using the
226         ↪ upstream socks proxy
227         remote = socks.socksocket()
228         remote.set_proxy(socks.SOCKS5, UPSTREAM_HOST, UPSTREAM_PORT) #
229         ↪ SOCKS4 and SOCKS5 use port 1080 by default
230     else:
231         cache_address = db.dbroot[self.intro_host.split('.')[0]].
232         ↪ get_latest_valid_data_address()
233         self.cache_request = True
234         self.handle_state_change(None)
235         message = "<HTML><HEAD><TITLE>UOSP
236         ↪ Response</TITLE></HEAD><BODY><h3>UOSP</h3>This page supports
237         ↪ the " \
238             "Unlinkable Onion Service Protocol (UOSP). Please use
239             ↪ the following link <a href='http://' +
240             ↪ cache_address[0] + ".onion"> USOP
241             ↪ Link</a></BODY></HTML>"
242
243         data = 'HTTP/1.1 200 OK\r\nServer: BaseHTTP/0.6
244         ↪ Python/3.6.9\r\nContent-Length: ' + str(
245             len(
246                 message)) + '\r\nUOSP: True\r\nUOSP_VERSION:
247             ↪ 0.1\r\nUOSP_VALIDITY:' + cache_address[1] +
248             ↪ '\r\n\r\n' + message
249
250         self.connection.send(data.encode())
251         transaction.commit()
252         db.close()
253     if not self.cache_request:
254         # Can be treated identical to a regular socket object
255         if validators.url('http://' + http_header['host']):
256
257             remote.connect((http_header['host'], http_header['port']))
258             ssl_remote = ssl.wrap_socket(remote, ca_certs=None,
259             ↪ cert_reqs=ssl.CERT_NONE)
260
261             # send request to the requested resource with UOSP header
262             ↪ added
263             request_string = 'GET / HTTP/1.1\r\nHost: ' +
264             ↪ http_header['host'] + ':' + str(
265                 http_header['port']) + '\r\nUOSP: True\r\nUOSP_VERSION:'
266             ↪ + UOSP_VERSION + '\r\n\r\n'
267             ssl_remote.sendall(request_string.encode())
268
269             # trigger and handle state change
270             self.handle_state_change(request_string)
271             self.exchange_loop(self.connection, ssl_remote)
272         else:
273             logging.error("The URL specified wasn't valid")
274     else:

```

```

257         # relay request to remote http server
258         try:
259             logging.info('Non-onion traffic, relay HTTP traffic as a normal
                ↪ proxy would do')
260             remote = socket.socket(socket.AF_INET, socket.SOCK_STREAM)
261             remote.connect((http_header['host'], http_header['port']))
262             logging.info('Connected to %s %s' % (http_header['host'],
                ↪ http_header['port']))
263             request_string = 'GET / HTTP/1.1\r\nHost: ' + http_header['host']
                ↪ + ':' + str(
264                 http_header['port']) + '\r\n\r\n'
265             remote.sendall(request_string.encode())
266             self.legacy_request = True
267             self.exchange_loop(self.connection, remote)
268         except Exception as e:
269             print(e)
270     else:
271         logging.error("Couldn't parse request header")
272
273     """
274     Loop for exchanging data between client and remote server
275
276     :param client Handle for client socket
277     :param remote Handle for server socket
278     """
279
280     def exchange_loop(self, client, remote):
281         while True:
282             # wait until client or remote is available for read
283             r, w, e = select.select([client, remote], [], [])
284
285             # there is data to be read from the client socket
286             if client in r:
287                 # receive data and parse its body
288                 parsed_data = self.parse_http_body(client)
289                 if parsed_data is not None:
290                     # trigger protocol state change2
291                     self.handle_state_change(parsed_data)
292
293             # there is data to be read from the server socket
294             if remote in r:
295                 # check if legacy request
296                 if not self.legacy_request:
297                     # receive data and parse its body
298                     parsed_data = self.parse_http_body(remote)
299                     if parsed_data is not None:
300                         # check if UOSP response
301                         if not parsed_data[1]:
302                             # no UOSP response, just send back the data received
303                             data = 'HTTP/1.1 200 OK\r\nServer: BaseHTTP/0.6
                                    ↪ Python/3.6.9\r\nContent-Length: ' + str(
304                                 len(parsed_data[2])) + '\r\n\r\n' + parsed_data[2]
305                             client.send(data.encode())
306                             break
307                         else:
308                             # UOSP response, trigger protocol state change
309                             self.handle_state_change(parsed_data, remote)
310             # don't send anything back to the client, till the UOSP has
                ↪ finished, then reply with the temporary onion address
311             if self.csm.state == 'server_fin_received':

```

```

312     # data = 'HTTP/1.1 301 Moved Permanently\r\nLocation: ' +
    ↪ 'http://' + str(self.uosp_address) + '.onion\r\n\r\n'
313     message = "<HTML><HEAD><TITLE>UOSP
    ↪ Response</TITLE></HEAD><BODY><h3>UOSP</h3>This page
    ↪ supports the " \
314         "Unlinkable Onion Service Protocol (UOSP). Please
    ↪ use the following link <a href='http://' + str(
315     self.uosp_address) + ".onion"> USOP
    ↪ Link</a></BODY></HTML>"
316
317     data = 'HTTP/1.1 200 OK\r\nServer: BaseHTTP/0.6
    ↪ Python/3.6.9\r\nContent-Length: ' + str(
318     len(
319     message)) + '\r\nUOSP: True\r\nUOSP_VERSION:
    ↪ 0.1\r\nUOSP_VALIDITY:' + self.uosp_validity +
    ↪ '\r\n\r\n' + message
320
321     client.send(data.encode())
322     # add uosp data address to db for the given intro address
323     db = MyZODB(DB_PATH)
324     if self.intro_host.split('.')[0] not in db.dbroot:
325         db.dbroot[self.intro_host.split('.')[0]] =
    ↪ UospIntroAddress(
326             self.intro_host.split('.')[0])
327     db.dbroot[self.intro_host.split('.')[0]].
    ↪ add_data_address(self.uosp_address, self.uosp_validity)
328     transaction.commit()
329     db.close()
330     break
331     # handle uosp data response
332     if self.csm.state == 'data_resp_received':
333         data = 'HTTP/1.1 200 OK\r\nServer: BaseHTTP/0.6
    ↪ Python/3.6.9\r\nContent-Length: ' + str(
334         len(parsed_data[2])) +
    ↪ '\r\nUOSP: True\r\nUOSP_VERSION: 0.1\r\n\r\n' +
    ↪ parsed_data[2]
335         client.send(data.encode())
336     else:
337         # legacy request send back the response to the client
338         data = remote.recv(4096)
339         if client.send(data) <= 0:
340             break
341
342
343 if __name__ == '__main__':
344     ForkingTCPServer.allow_reuse_address = True
345     with ForkingTCPServer((PROXY_HOST, PROXY_PORT), SocksProxy) as server:
346         server.serve_forever()

```

A.8 Proxy Tor Configuration

```

1 ## Logging
2 Log notice file /var/log/tor/notices.log
3 Log info file /var/log/tor/info.log
4 Log debug file /var/log/tor/debug.log
5 ## Run as daemon in background
6 RunAsDaemon 1

```

Appendix B: Performance Analysis

B.1 Analysis Module

```
1 import time
2 import csv
3 import re
4 import datetime
5
6 # default timeout after 3 minutes (180s)
7 TIMEOUT = 180
8
9 # default number of intro points is currently 3
10 NUM_OF_INTRO_POINTS = 3
11
12 from connection.utils import establish_connection
13 from stem.control import EventType
14 from circuitAnalyzer import circCallback
15 from logAnalyzer import logCallback_v3
16 from logAnalyzer import logCallback_v2
17 from descriptorAnalyzer import descCallback
18 from logAnalyzer import thread_intro_circuit_q
19 from circuitAnalyzer import thread_circuit_list_q
20 from circuitAnalyzer import thread_upload_circuit_q
21 from descriptorAnalyzer import thread_descriptors_q
22 from circuitAnalyzer import circMinorCallback
23 from connection.utils import create_ephemeral_hidden_service
24 from connection.utils import create_hidden_service
25
26 # password for control port authentication
27 PASSWORD = 'tor'
28
29 # regex for ipv4 address in desc event
30 UPLOAD_REGEX = '([0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}) ([0-9]{3,4})'
31
32 # V3
33 NUM_OF_CREATE_EVENT_V3 = 2
34 NUM_OF_UPLOAD_EVENTS_V3 = 16
35 NUM_OF_UPLOADED_EVENTS_V3 = 16
36
37 # V2
38 NUM_OF_CREATE_EVENT_V2 = 2
39 NUM_OF_UPLOAD_EVENTS_V2 = 6
40 NUM_OF_UPLOADED_EVENTS_V2 = 6
41
42 controller = None
43
44 # filepath of output csv file
45 CSV_PATH = "/tmp/provisioning_time_" + time.strftime("%Y_%m_%d_%H_%M",
46 ↪ time.localtime()) + ".csv"
47
48 def checkDictEntries(dictionary, csv_field_name):
49     for entry in csv_field_name:
50         if entry not in dictionary:
51             dictionary[entry] = "FAIL"
52
53
54 def start(version, ephemeral, timeout, num_intro_points):
```

```

55 global controller, TIMEOUT, NUM_OF_INTRO_POINTS
56 csv_dict = {}
57 startTime = time.time()
58
59 # placeholder for the lists
60 intro_point_list = []
61 descriptor_create_list = []
62 descriptor_upload_list = []
63 descriptor_uploaded_list = []
64 hmdir_upload_circuit_list = []
65
66 # if timeout is provided by parameters use this one, otherwise the default
67 ↪ timeout (360s) is used
68 if timeout is not None and isinstance(timeout, int):
69     TIMEOUT = timeout
70
71 # if a custom number of intro points is provided, use it, otherwise use the
72 ↪ default (3)
73 if num_intro_points is not None and isinstance(num_intro_points, int):
74     NUM_OF_INTRO_POINTS = num_intro_points
75
76 # csv header names
77 CSV_FIELD_NAME_V3 = ['service_initiated']
78
79 for i in range(1, NUM_OF_INTRO_POINTS + 1):
80     CSV_FIELD_NAME_V3.append('intro_' + str(i) + '_build_start')
81     CSV_FIELD_NAME_V3.append('intro_' + str(i) + '_build_end')
82
83 for i in range(1, NUM_OF_INTRO_POINTS + 1):
84     CSV_FIELD_NAME_V3.append('intro_' + str(i) + '_established')
85
86 CSV_FIELD_NAME_V3.extend(['desc_1_created', 'desc_2_created', 'upload_hs_1',
87     'uploaded_hs_1',
88     'upload_hs_2', 'uploaded_hs_2', 'upload_hs_3',
89     ↪ 'uploaded_hs_3', 'upload_hs_4',
90     'uploaded_hs_4',
91     'upload_hs_5', 'uploaded_hs_5', 'upload_hs_6',
92     ↪ 'uploaded_hs_6', 'upload_hs_7',
93     'uploaded_hs_7',
94     'upload_hs_8', 'uploaded_hs_8', 'upload_hs_9',
95     ↪ 'uploaded_hs_9', 'upload_hs_10',
96     'uploaded_hs_10', 'upload_hs_11', 'uploaded_hs_11',
97     ↪ 'upload_hs_12', 'uploaded_hs_12',
98     'upload_hs_13',
99     'uploaded_hs_13', 'upload_hs_14', 'uploaded_hs_14',
100    ↪ 'upload_hs_15', 'uploaded_hs_15',
101    'upload_hs_16',
102    'uploaded_hs_16', 'upload_hs_1_hmdir_addr',
103    ↪ 'upload_hs_2_hmdir_addr',
104    'upload_hs_3_hmdir_addr',
105    'upload_hs_4_hmdir_addr', 'upload_hs_5_hmdir_addr',
106    ↪ 'upload_hs_6_hmdir_addr',
107    'upload_hs_7_hmdir_addr', 'upload_hs_8_hmdir_addr',
108    ↪ 'upload_hs_9_hmdir_addr',
109    'upload_hs_10_hmdir_addr', 'upload_hs_11_hmdir_addr',
110    ↪ 'upload_hs_12_hmdir_addr',
111    'upload_hs_13_hmdir_addr', 'upload_hs_14_hmdir_addr',
112    ↪ 'upload_hs_15_hmdir_addr',
113    'upload_hs_16_hmdir_addr'])

```

```

103     for i in range(1, 17):
104         CSV_FIELD_NAME_V3.append(f"upload_hs_{i}_path")
105         CSV_FIELD_NAME_V3.append(f"upload_hs_{i}_circuit_built")
106
107     CSV_FIELD_NAME_V2 = ['service_initiated']
108
109     for i in range(1, NUM_OF_INTRO_POINTS + 1):
110         CSV_FIELD_NAME_V2.append('intro_' + str(i) + '_build_start')
111         CSV_FIELD_NAME_V2.append('intro_' + str(i) + '_build_end')
112
113     for i in range(1, NUM_OF_INTRO_POINTS + 1):
114         CSV_FIELD_NAME_V2.append('intro_' + str(i) + '_established')
115
116     CSV_FIELD_NAME_V2.extend(['desc_1_created', 'desc_2_created', 'upload_hs_1',
117     ↪ 'uploaded_hs_1',
118     ↪ 'upload_hs_2', 'uploaded_hs_2', 'upload_hs_3',
119     ↪ 'uploaded_hs_3', 'upload_hs_4',
120     ↪ 'uploaded_hs_4',
121     ↪ 'upload_hs_5', 'uploaded_hs_5', 'upload_hs_6',
122     ↪ 'uploaded_hs_6', 'upload_hs_1_hsdire_addr',
123     ↪ 'upload_hs_2_hsdire_addr', 'upload_hs_3_hsdire_addr',
124     ↪ 'upload_hs_4_hsdire_addr', 'upload_hs_5_hsdire_addr',
125     ↪ 'upload_hs_6_hsdire_addr'])
126
127     for i in range(1, 7):
128         CSV_FIELD_NAME_V2.append(f"upload_hs_{i}_path")
129         CSV_FIELD_NAME_V3.append(f"upload_hs_{i}_circuit_built")
130
131     # connect and authenticate
132     controller = establish_connection(password=PASSWORD)
133     csv_dict['service_initiated'] = datetime.datetime.now()
134
135     # add event listeners
136     controller.add_event_listener(circCallback, EventType.CIRC)
137     controller.add_event_listener(descCallback, EventType.HS_DESC)
138     controller.add_event_listener(circMinorCallback, EventType.CIRC_MINOR)
139     if version == 3:
140         controller.add_event_listener(logCallback_v3, EventType.DEBUG)
141     else:
142         controller.add_event_listener(logCallback_v2, EventType.DEBUG)
143
144     # create hidden service
145     if ephemeral == False and version == 3:
146         create_hidden_service(controller, '/var/lib/tor', 80)
147     else:
148         create_ephemeral_hidden_service(controller, 80, 80, version)
149
150     print(" * Onion Service created")
151
152     # number of required events for V2/V3
153     if version == 3:
154         createEvent = NUM_OF_CREATE_EVENT_V3
155         uploadEvents = NUM_OF_UPLOAD_EVENTS_V3
156         uploadedEvents = NUM_OF_UPLOADED_EVENTS_V3
157     else:
158         createEvent = NUM_OF_CREATE_EVENT_V2
159         uploadEvents = NUM_OF_UPLOAD_EVENTS_V2
160         uploadedEvents = NUM_OF_UPLOADED_EVENTS_V2
161
162     # loop while the required data is received

```

```

159     while True:
160         if time.time() - startTime > TIMEOUT:
161             print(" * One of many of the uploads timed out")
162             break
163
164         if thread_circuit_list_q.qsize() > 1:
165             circuit_list = thread_circuit_list_q.get()
166             thread_circuit_list_q.queue.clear()
167
168         if thread_intro_circuit_q.qsize() > 1:
169             intro_point_list = thread_intro_circuit_q.get()
170             thread_intro_circuit_q.queue.clear()
171
172         if thread_descriptors_q.qsize() > 3:
173             descriptor_create_list = thread_descriptors_q.get()
174             descriptor_upload_list = thread_descriptors_q.get()
175             descriptor_uploaded_list = thread_descriptors_q.get()
176             thread_descriptors_q.queue.clear()
177
178         if not thread_upload_circuit_q.empty():
179             hmdir_upload_circuit_list.append(thread_upload_circuit_q.get())
180
181         if len(intro_point_list) >= 3 and len(descriptor_create_list) == 2 and len(
182             descriptor_upload_list) == uploadEvents and
183             ↪ len(descriptor_uploaded_list) == uploadedEvents:
184             print(" * All required data available, starting correlation")
185             break
186
187         # clean up, remove event listeners (commented, because it was blocking some
188         ↪ times, don't know why)
189         controller.remove_event_listener(circCallback)
190         controller.remove_event_listener(descCallback)
191         controller.remove_event_listener(circMinorCallback)
192         if version == 3:
193             controller.remove_event_listener(logCallback_v3)
194         else:
195             controller.remove_event_listener(logCallback_v2)
196
197         """
198         correlates an element form the uploaded list with a provided fingerprint and
199         ↪ returns the matched element
200
201         :param uploaded_list: List of uploaded descriptors
202         :param fingerprint: Fingerprint from the upload list to correlate
203         :return: Element from the uploaded list that has the same fingerprint as the
204         ↪ element from the upload list
205         """
206
207         def correlateFingerprint(uploaded_list, fingerprint):
208             for elem in uploaded_list:
209                 if elem.fingerprint == fingerprint:
210                     return elem
211             return None
212
213         # read data from the various lists and put the into a dict uniformly that their
214         ↪ values can be written to a CSV file
215         for i in range(0, createEvent):
216             csv_dict["desc_" + str(i + 1) + "_created"] =
217                 ↪ (descriptor_create_list[i]).getTime()
218             i += 1

```

```

213
214 for i in range(0, uploadEvents):
215     if i <= len(descriptor_upload_list) - 1:
216         csv_dict["upload_hs_" + str(i + 1)] = descriptor_upload_list[i]
217         try:
218             hmdir_addr = re.search(UPLOAD_REGEX, str(
219                 controller.get_network_status
220                 ↪ (relay=descriptor_upload_list[i].getFingerprint())))
221             if hmdir_addr:
222                 csv_dict["upload_hs_" + str(i + 1) + "_hmdir_addr"] =
223                     ↪ str(hmdir_addr.group(1)) + ":" + str(
224                         hmdir_addr.group(2))
225             else:
226                 csv_dict["upload_hs_" + str(i + 1) + "_hmdir_addr"] = "FAIL"
227         except:
228             csv_dict["upload_hs_" + str(i + 1) + "_hmdir_addr"] = "FAIL"
229     else:
230         csv_dict["upload_hs_" + str(i + 1)] = 'FAIL'
231         csv_dict["upload_hs_" + str(i + 1) + "_hmdir_addr"] = 'FAIL'
232
233 for i in range(0, len(hmdir_upload_circuit_paths)):
234     # to upload circuit path is contained in the fingerprint attribute, while the
235     ↪ upload circuit built timestamp is contained in the objects time stamp
236     csv_dict[f"upload_hs_{i + 1}_circuit_built"] =
237     ↪ hmdir_upload_circuit_list[i].getTime()
238     csv_dict[f"upload_hs_{i + 1}_path"] =
239     ↪ str(hmdir_upload_circuit_list[i].getFingerprint())
240
241 for i in range(0, uploadedEvents):
242     if i <= len(descriptor_uploaded_list) - 1:
243         if csv_dict["upload_hs_" + str(i + 1)] != 'FAIL' and
244             ↪ correlateFingerprint(descriptor_uploaded_list, str(
245                 csv_dict["upload_hs_" + str(i + 1)].getFingerprint())):
246             csv_dict["uploaded_hs_" + str(i + 1)] =
247             ↪ correlateFingerprint(descriptor_uploaded_list, str(
248                 csv_dict["upload_hs_" + str(i + 1)].getFingerprint()).getTime())
249             csv_dict["upload_hs_" + str(i + 1)] = csv_dict["upload_hs_" + str(i +
250             ↪ 1)].getTime()
251         else:
252             csv_dict["upload_hs_" + str(i + 1)] = "FAIL"
253     else:
254         if csv_dict["upload_hs_" + str(i + 1)] != 'FAIL':
255             csv_dict["upload_hs_" + str(i + 1)] = csv_dict["upload_hs_" + str(i +
256             ↪ 1)].getTime()
257             csv_dict["uploaded_hs_" + str(i + 1)] = 'FAIL'
258
259 for i in range(0, NUM_OF_INTRO_POINTS):
260     if i <= len(intro_point_list) - 1:
261         csv_dict["intro_" + str(i + 1) + "_established"] = intro_point_list[i][0]
262         csv_dict["intro_" + str(i + 1) + "_build_start"] =
263         ↪ circuit_list[str(intro_point_list[i][1])][0].getTime()
264         csv_dict["intro_" + str(i + 1) + "_build_end"] =
265         ↪ circuit_list[str(intro_point_list[i][1])][-1].getTime()
266     else:
267         csv_dict["intro_" + str(i + 1) + "_established"] = "FAIL"
268         csv_dict["intro_" + str(i + 1) + "_build_start"] = "FAIL"
269         csv_dict["intro_" + str(i + 1) + "_build_end"] = "FAIL"
270
271     # write output to csv file
272 print(" * Correlation performed, writing results to CSV file")

```



```

262
263 # check if all mandatory fields are set, if not set them to FAIL
264 if version == 3:
265     checkDictEntries(csv_dict, CSV_FIELD_NAME_V3)
266 else:
267     checkDictEntries(csv_dict, CSV_FIELD_NAME_V2)
268
269 if version == 3:
270     with open(CSV_PATH, 'w', newline='') as csvfile:
271         logwriter = csv.DictWriter(csvfile, delimiter=',', quotechar='"',
272     ↪ quoting=csv.QUOTE_MINIMAL,
273                                     fieldnames=CSV_FIELD_NAME_V3)
274         logwriter.writeheader()
275         logwriter.writerow(csv_dict)
276 else:
277     with open(CSV_PATH, 'w', newline='') as csvfile:
278         logwriter = csv.DictWriter(csvfile, delimiter=',', quotechar='"',
279     ↪ quoting=csv.QUOTE_MINIMAL,
280                                     fieldnames=CSV_FIELD_NAME_V2)
281         logwriter.writeheader()
282         logwriter.writerow(csv_dict)

```

B.2 Circuit Analyzer

```

1 import datetime
2 import re
3 import queue
4
5 # global dictionary for circuits
6 circuits = {}
7
8 # regex for fingerprint
9 FINGERPRINT_REGEX = '[A-z0-9]{40}'
10
11 # synchronized queue for inter thread communication
12 thread_circuit_list_q = queue.Queue()
13
14 # synchronized queue for tor upload circuits
15 thread_upload_circuit_q = queue.Queue()
16
17 """
18 class for circuit elements
19 """
20
21 class circuitElement:
22     def __init__(self, time_added, fingerprint):
23         self.__time_added = time_added
24         self.__fingerprint = fingerprint
25
26     def getTime(self):
27         return self.__time_added
28
29     def getFingerprint(self):
30         return self.__fingerprint
31
32
33 """
34 callback handler when a CIRC event takes place

```

```

35 """
36
37
38 def circCallback(event):
39     global circuits
40     # save timestamp of the event
41     timestamp = datetime.datetime.now()
42
43     # write to queue such that the main thread has access to the circuit list
44     if len(circuits) > 0:
45         # put new version into queue
46         thread_circuit_list_q.put(circuits)
47
48     # circuits for finding introduction points
49     if event.purpose == "HS_SERVICE_INTRO":
50         # Circuit is new, add new entry to global circuit dictionary
51         if event.status == "LAUNCHED":
52             circuits[event.id] = []
53         # circuit is extended
54         elif event.status == "EXTENDED":
55             circuits[event.id].append(circuitElement(timestamp,
56             ↪ extractFingerprint(event.path)))
57         # circuit failed or closed, delete entry from dictionary
58         elif event.status == "FAILED" or event.status == "CLOSED":
59             circuits.pop(event.id)
60
61     if event.purpose == "HS_SERVICE_HSDIR":
62         if event.status == "LAUNCHED":
63             print(f"[{datetime.datetime.now()}] CircuitEvent: {event.id} | status:
64             ↪ {event.status} | purpose: {event.purpose} | path: {event.path}")
65         if event.status == "BUILT":
66             # in this case the fingerprint attribute is misued for the upload circuit
67             ↪ path
68             thread_upload_circuit_q.put(circuitElement(timestamp, event.path))
69             print(f"[{datetime.datetime.now()}] CircuitEvent: {event.id} | status:
70             ↪ {event.status} | purpose: {event.purpose} | path: {event.path}")
71         elif event.status == "FAILED":
72             # One Upload failed, so a new attempt will be made.
73             # thread_upload_circuit_q.put("FAIL")
74             print(f"[{datetime.datetime.now()}] CircuitEvent: {event.id} | status:
75             ↪ {event.status} | purpose: {event.purpose} | path: {event.path}")
76
77
78 """
79 callback handler when a CIRC MINOR event takes place
80 """
81
82
83 def circMinorCallback(event):
84     global circuits
85     if not event.id in circuits:
86         if event.event == 'CANNIBALIZED':
87             path_list = extractFingerprint(event.path, cannibalized=True)
88             circuits[event.id] = []
89             for elem in path_list:
90                 circuits[event.id].append(circuitElement(str(datetime.datetime.now())
91                 ↪ + " (CAN)", elem))
92
93 """

```

```

89 extracts the fingerprint from a path string that is not yet in the specific circuits
    ↪ lists
90
91 :param path: Path string of the CIRC event
92 :return: Fingerprint
93 """
94
95
96 def extractFingerprint(path, cannibalized=False):
97     fingerprint_re = re.compile(FINGERPRINT_REGEX)
98     path_list = re.findall(fingerprint_re, str(path))
99
100     if len(path_list) > 0 and not cannibalized:
101         return path_list[-1]
102     elif len(path_list) > 0 and cannibalized:
103         return path_list
104
105
106 """
107 print all the circuits and their paths
108 """
109
110
111 def printCircuitList():
112     for circuit in circuits:
113         print(f"# {circuit}")
114         for relay in circuits[circuit]:
115             print(f" {relay.getTime()}: [{relay.getFingerprint()}]")
116         print("\\n")

```

B.3 Log Analyzer

```

1 """
2 callback handler when a LOG event takes place
3 """
4
5 import datetime
6 import queue
7 import re
8
9 # regex for intro points
10 INTRO_ESTABLISHED_REGEX_V3 = 'service_handle_intro_established\\(\\): Successfully
    ↪ received an INTRO_ESTABLISHED cell on circuit ([0-9]{1,}) \\(id: ([0-9]{1,})\\)'
11 INTRO_ESTABLISHED_REGEX_V2 = 'rend_service_intro_established\\(\\): Received
    ↪ INTRO_ESTABLISHED cell on circuit ([0-9]{1,}) \\(id: ([0-9]{1,})\\)'
12
13 intro_circuit_list = []
14
15 # synchronized queue for inter thread communication
16 thread_intro_circuit_q = queue.Queue()
17
18 """
19 callback handler when a LOG event takes place
20 """
21
22 def logCallback_v2(event):
23     global intro_points_extracted
24

```

```

25     timestamp = datetime.datetime.now()
26     # extractIntroEstablishedCircuit(event.message)
27     int_tmp = extractCircuitID(event.message, timestamp, 2)
28     if int_tmp is not None:
29         intro_circuit_list.append(int_tmp)
30         # put new version into queue
31         thread_intro_circuit_q.put(intro_circuit_list)
32
33
34 def logCallback_v3(event):
35     global intro_points_extracted
36
37     timestamp = datetime.datetime.now()
38     # extractIntroEstablishedCircuit(event.message)
39     int_tmp = extractCircuitID(event.message, timestamp, 3)
40     if int_tmp is not None:
41         intro_circuit_list.append(int_tmp)
42         # put new version into queue
43         thread_intro_circuit_q.put(intro_circuit_list)
44
45
46 """
47 extract the circuit ID from a intro established log message and return a list with
48 ↪ its timestamp and id
49 :param line: Log string
50 :param timestamp: timestamp when the controller received the log message
51 :return: list of timestamp and ID of intro circuit (global ID)
52 """
53 def extractCircuitID(line, timestamp, version):
54     if version == 3:
55         circuit_id = re.search(INTRO_ESTABLISHED_REGEX_V3, str(line))
56     else:
57         circuit_id = re.search(INTRO_ESTABLISHED_REGEX_V2, str(line))
58
59     # if there is a match, extract the circuit ID
60     if circuit_id is not None:
61         return [timestamp, circuit_id.group(2)]
62     else:
63         return None

```

B.4 Descriptor Analyzer

```

1 import queue
2 import datetime
3
4 # synchronized queue for inter thread communication
5 thread_descriptors_q = queue.Queue()
6
7 descUploadEvents = []
8 descUploadedEvents = []
9 descCreatedEvents = []
10
11 """
12 class for descriptor elements
13 """
14
15

```

```

16 class descElement:
17     def __init__(self, time_added, fingerprint):
18         self.time_added = time_added
19         self.fingerprint = fingerprint
20
21     def getTime(self):
22         return self.time_added
23
24     def getFingerprint(self):
25         return self.fingerprint
26
27
28 """
29 callback handler when a DESC event takes place
30 """
31
32
33 def descCallback(event):
34     global descEvents
35     # collected all required data, put it in a queue such that the main thread can
36     ↪ read them
37     # put new version into queue
38     thread_descriptors_q.put(descCreatedEvents)
39     thread_descriptors_q.put(descUploadEvents)
40     thread_descriptors_q.put(descUploadedEvents)
41
42     if event.action == "CREATED":
43         descCreatedEvents.append(descElement(datetime.datetime.now(), None))
44     elif event.action == "UPLOAD":
45         descUploadEvents.append(descElement(datetime.datetime.now(),
46     ↪ event.directory_fingerprint))
47         print(f"[{datetime.datetime.now()}] action: {event.action} | descriptor:
48     ↪ {event.descriptor_id} | replica: {event.replica}")
49     elif event.action == "UPLOADED":
50         descUploadedEvents.append(descElement(datetime.datetime.now(),
51     ↪ event.directory_fingerprint))
52         print(f"[{datetime.datetime.now()}] action: {event.action} | descriptor:
53     ↪ {event.descriptor_id} | replica: {event.replica}")

```

B.5 Start Analysis

```

1 import re
2 from analysis import start
3 from connection.utils import establish_connection
4 import sys
5
6 NOTICES_LOG_PATH = '/var/log/tor/notices.log'
7 BOOTSTRAP_REGEX = 'Bootstrapped 100\%'
8 PASSWORD = 'tor'
9
10 version = int(sys.argv[1])
11 ephemeral = sys.argv[2] == "true" or sys.argv[2] == "True" or sys.argv[2] == "TRUE"
12 timeout = int(sys.argv[3])
13 num_intro_points = int(sys.argv[4])
14
15 bootstrapped = False
16
17

```

```

18 def parseLog(file_path):
19     global bootstrapped
20     with open(file_path) as infile:
21         for line in infile:
22             if checkIfBoostrapped(line):
23                 bootstrapped = True
24
25 def checkIfBoostrapped(line):
26     if re.search(BOOTSTRAP_REGEX, line) is not None:
27         return True
28     return False
29
30 while not bootstrapped:
31     parseLog(NOTICES_LOG_PATH)
32
33 controller = establish_connection(password=PASSWORD)
34 print(f" * Number of circuits after bootstrapping: {len(controller.get_circuits())}")
35 print(" * Tor completely bootstrapped, starting analysis")
36
37 start(version, ephemeral, timeout, num_intro_points)

```

B.6 Utils

```

1 import sys
2 import stem.control
3 import stem.connection
4
5 """
6 Establish a connection to the control port of tor and authenticate
7 """
8 CONTROL_PORT = 9051
9 CONTROL_ADDRESS = '127.0.0.1'
10
11
12 def establish_connection(address=CONTROL_ADDRESS, port=CONTROL_PORT, password=None):
13     # Get connection to tor via the control port
14     try:
15         # the explicit declaration of the port isn't necessary since it's the default
16         ↪ port
17         controller = stem.control.Controller.from_port(address, port)
18     except stem.SocketError as exc:
19         print(f" * Unable to connect to port control port 9051 ({exc})")
20         sys.exit(1)
21
22     if controller:
23         # authenticate
24         try:
25             if (password != None):
26                 controller.authenticate(password=password)
27             else:
28                 controller.authenticate()
29             print(" * Connection established and authenticated")
30             return controller
31         except stem.connection.PasswordAuthFailed:
32             print(" * Unable to authenticate, password is incorrect")
33             sys.exit(1)
34         except stem.connection.AuthenticationFailure as exc:
35             print(f" * Unable to authenticate: {exc}")

```

```

35         sys.exit(1)
36
37
38 """
39 Create new ephemeral Hidden Service and return its onion id
40 """
41
42
43 def create_ephemeral_hidden_service(controller, source_port, target_port, version):
44     try:
45         if version == 2:
46             hs = controller.create_ephemeral_hidden_service({source_port:
47                 ↪ target_port}, await_publication=False,
48                                                             key_type="NEW",
49                                                             ↪ key_content="RSA1024",
50                                                             ↪ detached=False)
49         else:
50             hs = controller.create_ephemeral_hidden_service({source_port:
51                 ↪ target_port}, await_publication=False,
52                                                             detached=False)
51         if (hs.service_id):
52             return hs.service_id
53         else:
54             return -1
55
56     except stem.ControllerError as exc:
57         print(f"Unable to create the hidden service ({exc})")
58         sys.exit(1)
59
60
61 def create_hidden_service(controller, path, port):
62     try:
63         controller.create_hidden_service(path, port, target_address='127.0.0.1')
64     except stem.ControllerError as exc:
65         print(f"Unable to create the hidden service ({exc})")
66         sys.exit(1)

```

B.7 Docker File

```

1 # image based on debian
2 FROM debian:latest
3
4 # update package lists
5 RUN apt-get update -y
6
7 # workaround for avoiding debconf failure messages
8 RUN echo 'debconf debconf/frontend select Noninteractive' | debconf-set-selections
9
10 # install all packages required to compile and run tor within the container
11 RUN apt-get install libssl-dev -y -qq > /dev/null
12 RUN apt-get install zlib1g-dev -y -qq > /dev/null
13 RUN apt-get install python3 -y -qq > /dev/null
14 RUN apt-get install python3-pip -y -qq > /dev/null
15 RUN apt-get install libevent-dev -y -qq > /dev/null
16
17 # create directories
18 RUN mkdir -p /home/script
19 RUN mkdir -p /home/tor

```

```
20 RUN mkdir -p /etc/tor
21 RUN mkdir -p /var/log/tor
22 RUN mkdir -p /var/lib/tor
23 RUN mkdir -p /run/tor
24
25 # init file, tor binaries and torrc file
26
27 # the init handles execution parameters (which analysis should be started, V2, V3,
  ↪ V3VN, V3NE?)
28 ADD init_tor.sh /home/script/init_tor.sh
29
30 # tor source code, which is compiled within the container (should be kept up-to-date)
31 ADD tor /home/tor
32
33 # tor config file
34 ADD torrc /etc/tor/torrc
35
36 # vanguards config file
37 ADD vanguards.conf /etc/tor/vanguards.conf
38
39 # python requirements (currently only stem)
40 ADD requirements.txt /home/requirements.txt
41
42 # install python package manager
43 RUN pip3 install --upgrade pip
44
45 # install required packages (currently only stem)
46 RUN pip3 install -r /home/requirements.txt
47
48 # install vanguards (could also be done within the requirements file)
49 RUN pip3 install vanguards
50
51 # create log files
52 RUN touch /var/log/tor/debug.log
53 RUN touch /var/log/tor/info.log
54 RUN touch /var/log/tor/notices.log
55
56 # build custom tor binary
57 RUN sh ./home/tor/configure
58 RUN chmod +x /home/tor/install-sh
59 RUN make && make install
60
61 RUN chmod +x /home/script/init_tor.sh
62 RUN chmod 700 /var/lib/tor
63
64 # add timing analysis prototype
65 ADD prototype-timing-analysis/ /home/timing-analysis/
66
67 # run init script
68 ENTRYPOINT ["/bin/bash", "/home/script/init_tor.sh"]
```

B.8 Analysis Server

```
1 #!/bin/bash
2
3 target="/path/to/results"
4 image="tor-docker:latest"
5 timeout=180
```



```

6 introduction_points=3
7 for i in {1..N}
8 do
9     echo "Iteration $i: V2"
10    docker run -it --rm -v${target}V2:/results -v/etc/localtime:/etc/localtime:ro
11    ↪ "${image}" -v 2 -t ${timeout} -i ${introduction_points}
12    echo "Iteration $i: V3"
13    docker run -it --rm -v${target}V3:/results -v/etc/localtime:/etc/localtime:ro
14    ↪ "${image}" -v 3 -t ${timeout} -i ${introduction_points}
15    echo "Iteration $i: V3NE"
16    docker run -it --rm -v${target}V3NE:/results -v/etc/localtime:/etc/localtime:ro
17    ↪ "${image}" -v 3NE -t ${timeout} -i ${introduction_points}
18    echo "Iteration $i: V3VN"
19    docker run -it --rm -v${target}V3VN:/results -v/etc/localtime:/etc/localtime:ro
20    ↪ "${image}" -v 3VN -t ${timeout} -i ${introduction_points}
21 done

```

B.9 Tor Configuration

```

1 ## Logging
2 ## Send all messages of level 'notice' or higher to /var/log/tor/notices.log
3 Log notice file /var/log/tor/notices.log
4 ## Send all messages of level 'info' possible message to /var/log/tor/info.log
5 Log info file /var/log/tor/info.log
6 ## Send every possible message to /var/log/tor/debug.log
7 Log debug file /var/log/tor/debug.log
8
9 ## Control port
10 ControlPort 9051
11 HashedControlPassword 16:CCA169830431CAB5608D58317D89A87D314B84B42FCCF3EDC82FD8B4E4
12
13 ## Run as daemon in background
14 RunAsDaemon 1
15
16 ## The directory for keeping all the keys/etc. By default, we store
17 ## things in $HOME/.tor on Unix, and in Application Data\tor on Windows.
18 DataDirectory /var/lib/tor
19
20
21 SafeLogging 0

```

B.10 Initialization Script

```

1 #!/bin/bash
2
3 # get parameter
4 while getopts ":v:t:i:" opt; do
5     case $opt in
6         v) version="$OPTARG"
7             ;;
8         t) timeout="$OPTARG"
9             ;;
10        i) num_intro="$OPTARG"
11            ;;
12        \?) echo "Invalid option -$OPTARG" >&2
13            ;;
14    esac

```

```
15 done
16
17 if [[ -z ${version} ] || [-z ${timeout}] || -z ${num_intro}];
18 then
19     echo -e "\t\nUsage:\nbash init_tor.sh -v VERSION -t TIMEOUT -i
    ↪ NUM_INTRO_POINTS\n"
20     echo -e "\t3: Ephemeral Onion Service Version 3\n"
21     echo -e "\t3NE: Non-ephemeral Onion Service Version 3\n"
22     echo -e "\t3VN: Ephemeral Onion Service with Vanguard Version 3\n"
23     echo -e "\t2: Ephemeral Onion Service Version 2\n"
24     exit 1
25 else
26     # start tor
27     tor -f /etc/tor/torrc --quiet &
28
29     export PYTHONPATH="/home/timing-analysis"
30     case "$version" in
31         3)         python3 /home/timing-analysis/analysis/startAnalysis.py 3 "True"
    ↪     ${timeout} ${num_intro}
32         ;;
33         3NE)       python3 /home/timing-analysis/analysis/startAnalysis.py 3 "False"
    ↪     ${timeout} ${num_intro}
34         ;;
35         3VN)       vanguards --config /etc/tor/vanguards.conf &
36         python3 /home/timing-analysis/analysis/startAnalysis.py 3 "True"
    ↪     ${timeout} ${num_intro}
37         ;;
38         2)         python3 /home/timing-analysis/analysis/startAnalysis.py 2 "True"
    ↪     ${timeout} ${num_intro}
39         ;;
40         *)         echo -e "Version can only be [2]|[3]|[3NE]|[3VN]"
41         exit 1
42         ;;
43     esac
44     # copy result file
45     cp /tmp/provisioning_time* /results
46 fi
```