

1 Dont get caught, keep your Onions in a Vault

2 **Humza Ikram***  

3 Carnegie Mellon University, Pennsylvania, US

4 **Rumaisa Habib***  

5 Stanford University, California, US

6 **Muaz Ali***  

7 University of Arizona, Arizona, US

8 **Zartash Afzal Uzmi**  

9 LUMS University, Pakistan

10 — Abstract

11 When web applications wish to operate anonymously, they routinely host themselves as ‘Hidden
12 Services’ in the Tor network. However, these services are frequently threatened by deanonymization
13 attacks, whereby their IP address and location may be inferred by the authorities. We present
14 **VaulTor**, a novel architecture for the Tor network that ensures an extra layer of security for the
15 Hidden Services against deanonymization attacks. In this new architecture, a volunteer (vault) is
16 incentivized to host the web application content on behalf of the Hidden Service. The vault runs the
17 hosted application in a Trusted Execution Environment (TEE) and becomes the point of contact
18 for interested clients. This setup can substantially reduce the uptime requirement of the original
19 Hidden Service provider, thereby significantly decreasing the chance of deanonymization attacks
20 against them. Using a vault node in place of the hidden service node does not cause any noticeable
21 performance degradation when accessing the hosted content.

22 **2012 ACM Subject Classification** Security and privacy → Pseudonymity, anonymity and untrace-
23 ability

24 **Keywords and phrases** Tor, anonymity, Hidden Services, Trusted Execution Environments

25 **Digital Object Identifier** 10.4230/OASIcs.NINeS.2026.17

26 **Funding** *Rumaisa Habib**: Stanford Graduate Fellowship

27 1 Introduction

28 The enormous expansion of the world wide web is coupled with growing demands for anonym-
29 ity and privacy. Besides a huge end-user client base, an increasing number of web services—
30 legal and illegal—also choose to remain anonymous [2], fearing closure, or even prosecution,
31 by the government and law enforcement agencies [8]. The Onion Router (Tor) [5] network
32 has emerged as one of the most popular solutions for providing anonymity: nearly 3 million
33 clients connect to Tor daily [9], and hundreds of thousands of anonymized web addresses are
34 published each day, with over 150,000 currently serving traffic to end users [17].

35 Tor Hidden Services (aka. Onion services) aim to uphold freedom of speech in repressive
36 regimes and offer circumvention in regions of undue and excessive internet censorship, thus
37 bringing benefits to the public [58]. At the same time, Hidden Services pave the way for
38 criminal activities such as selling illegal drugs and weapons [28]. All in all, there are incent-
39 ivies for governments and law enforcement to deanonymize Hidden Service (HS) providers¹
40 and curb their operations. In 2014, the authorities of 6 European countries and the United

* These authors contributed equally.

¹ Hidden Service Provider is an individual that owns the Hidden Service. They also create and serve the content of the Hidden Service.



© Humza Ikram, Rumaisa Habib, Muaz Ali, and Zartash Afzal Uzmi;

licensed under Creative Commons License CC-BY 4.0

1st New Ideas in Networked Systems (NINeS 2026).

Editors: Katerina J. Argyraki and Aurojit Panda; Article No. 17; pp. 17:1–17:24

 OpenAccess Series in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

17:2 Dont get caught, keep your Onions in a Vault

41 States collaborated in *Operation Onymous* through which they cracked down on 416 Hidden
42 Services, accused them of foul play, and arrested 17 individuals [8].

43 We focus on the *technical* challenges of reinforcing the anonymity of hidden service
44 providers. We present **VaulTor**, a novel overlay architecture that builds upon Tors existing
45 HS design by introducing a *vault* node between the HS provider and clients. The HS provider
46 offloads service operation to the vault, which serves clients on the providers behalf. This
47 setup strengthens the anonymity of HS providers in three key ways:

- 48 1. After offloading their service, the HS provider may sporadically connect to the vault to
49 update content as and when needed, without the need to be online all the time. This
50 significantly reduces the attack surface against long-running deanonymization attacks on
51 the HS provider.
- 52 2. The HS provider may connect to the vault from widely varying locations, further reducing
53 deanonymization risk.
- 54 3. An adversary cannot make an on-demand connection with an HS provider. Communica-
55 tion opportunities with the HS provider are now solely available to the vault, and even
56 those channels are initiated by the HS provider.

57 The service offloaded by the HS provider to the vault is hosted inside a Trusted Execution
58 Environment (TEE), such as the one provided by Intel-SGX [40]. Using a TEE brings
59 two additional benefits: (i) the vault owner (even with root access) cannot snoop on the
60 application code/data while it is being uploaded, served, or stored in the vault, and (ii) since
61 client and the HS provider contact the vault through similar encrypted channels, the vault
62 is never sure if it is an HS provider or a client it is communicating with, ruling out a myriad
63 of attacks against the HS provider (details in Sec 9.1).

64 The **VaulTor** architecture, which has no bearing on the client anonymity (discussed in
65 Sec 8.2), promises enhanced anonymity for the HS provider by shifting the deanonymization
66 risk to the vault owner (see Sec 8.3 for details). This design choice is suitable as the HS
67 providers may only be interested in content creation without taking the deanonymization
68 risk that leaps up when serving the content, even as a Tor Hidden Service. In contrast,
69 the vault owners are willing to serve the content in exchange for monetary benefits (such as
70 crypto payments²) or social incentives (serving content they wish to support and propagate).
71 A motivating example for our scenario might be a journalist in an oppressive regime who can
72 anonymously upload content banned in their regime to a vault located in a neutral country.
73 Another example may be providing ‘on-the-ground’ information from within a region where
74 internet outages frequently occur (and the HS provider cannot remain online for extended
75 periods).

76 **VaulTor** offers various attractive features not ubiquitous in alternate design choices: (a)
77 allows for hosting dynamic services, in contrast to data hosting services such as IPFS or
78 pastebin.com, (b) offers content isolation from the server administrator, in contrast to a
79 simple (non-TEE based) virtual machine-based approach, and (c) ensures that the server
80 administrator cannot snoop on private keys in order to serve arbitrary content on behalf of
81 the HS provider.

82 **VaulTor** also brings robustness to the HS operation. If a vault node is shut down by the
83 action of authorities, the HS provider can use a different one to make the content available.
84 Alternatively, if the HS provider goes offline permanently, the TEE can continue hosting the
85 content as long as the vault owner remains incentivized.

86 In **VaulTor**, the vault simply replaces and acts on behalf of the HS provider, responding

² We discuss a secure way to do crypto transactions in Section 10.3

87 to clients through network paths containing the same number of Tor nodes as if the content
 88 is served by the HS provider in the existing architecture. Thus, the `VaulTor` design bears
 89 no negative impact on the network latency. We do observe a slight average increase in the
 90 time to first byte and time to last byte (a maximum of 5.7% and 2.9% respectively when
 91 using Intel-SGX as the TEE), attributed to hosting the content inside a TEE which incurs
 92 its own performance overhead (9.2.2).

93 To understand the operation of `VaulTor`, we note that a typical Tor circuit is an overlay
 94 path through three volunteer relay nodes (the entry guard node, a middle node and the exit
 95 node) on the Tor network [32]. Such a circuit provides *one-way anonymity* to a client trying
 96 to connect via the circuit to any non-anonymous server on the web. If the server wants to
 97 remain anonymous as well, it must also choose another Tor circuit with three nodes [24].
 98 The existing HS architecture anonymizes both the client and the server (two-way anonymity)
 99 using a mechanism that stitches the two Tor circuits together (see Sec 2.2). With the
 100 `VaulTor` architecture, a server that desires to be anonymous is facilitated to replicate, and
 101 infrequently update, its web content, and service on the vault. This is similar to a CDN
 102 node offering content replication to a server. The two-way anonymity continues to exist
 103 between the vault node and the client, delivering similar delay performance as observed in
 104 the existing hidden service architecture. Indeed, our results in Sec 9.2.2 confirm this. We
 105 also consider a sample of popular deanonymization attacks and provide an outline of how
 106 `VaulTor` offers enhanced HS provider anonymity under those attacks (see Sec 9.1). We also
 107 discuss legal and deployability considerations for both HS providers and vault operators in
 108 Sec 10.5 Altogether, this paper makes the following contributions:

- 109 ■ A novel architecture of Tor nodes, used by `VaulTor` to provide robust anonymity to host
 110 Hidden Services (HS).
- 111 ■ Step-wise description of the protocol `VaulTor` uses to ensure anonymity of the HS provider,
 112 vault, and the clients.
- 113 ■ A thorough security analysis and description of the attacks that are mitigated by `VaulTor`.
- 114 ■ A working prototype and its performance measurement over the actual Tor network. The
 115 prototype is available at: <https://github.com/RumaisaHabib/vaultor>

116 2 Background

117 2.1 Trusted Execution Environments (TEEs)

118 TEEs provide a platform for *secure remote computation* which allows *securely* executing
 119 an application in a remote *untrusted* system without compromising the *Integrity* and *Con-*
 120 *fidentiality* of the application data. Several hardware architectures provide TEE imple-
 121 *mentations* [40, 31, 16, 48]. Our prototype implementation leverages the TEE provided by
 122 Intel-SGX [40] to create a secure execution channel between the vault and the HS provider.
 123 We host an HS inside the TEE which itself is set-up inside a vault (Sec 7). An instance of
 124 a program running inside a TEE is called an enclave.

125 TEEs provide many guarantees. These include:

- 126 1. **Sealing:** A program running inside a TEE can encrypt and write data to the disk for
 127 persistent storage. Only the *same program* running on *the same device* in a TEE can
 128 decrypt this data.
- 129 2. **Isolation:** A program inside a TEE can not access the memory of its host and vice versa.
- 130 3. **Remote Attestation:** A piece of code running inside a TEE can prove, to an outside
 131 observer, what piece of code it is and that it is running inside a TEE.

17:4 Dont get caught, keep your Onions in a Vault

132 A local trusted device may use the following *simplified flow* for remote attestation to
133 verify that a remote untrusted device is running the desired piece of code:

- 134 1. The local device builds the code and gets a measurement of the memory space³.
- 135 2. The local device sends the code to the remote device.
- 136 3. The remote device runs this code inside a TEE.
- 137 4. When running inside a TEE, the code can then request the CPU to generate a crypto-
138 graphic hash of the program's memory space. This hash is then signed by the CPU using
139 a secret hardware attestation key embedded in the CPU.
- 140 5. The local device can request this signed measurement from the remote device. It then
141 verifies this signature by matching the measurement hash against one computed locally
142 and by verifying the signature against its public key.

143 We will stick to Intel-SGX terminology, in which the signed cryptographic hash is called
144 a *quote*. Furthermore, it should be noted that a small amount of arbitrary data (produced
145 by the code inside the TEE) can be embedded in the quote as well. This data is called the
146 REPORTDATA field. Different TEE implementations employ different means for signature
147 verification. Intel-SGX, for example, requires sending the quote to Intel's online Attestation
148 Service which verifies that the quote is valid. The local device also has the option to perform
149 this attestation via a proxy (such as Tor).

150 The REPORTDATA field is crucial for establishing a secure connection with an enclave
151 (program running within the TEE). This field routinely contains a public key whose corre-
152 sponding private key is known only to the enclave. This public key can then be used to create
153 a secure connection with the enclave using any form of key exchange such as Diffie-Hellman.

154 2.2 Conventional Tor architecture

155 In the current Tor protocol, there are 6 Tor nodes between a client and a hidden service (HS).
156 This ensures two-way anonymity, i.e., both the client accessing the HS and the HS provider
157 itself remain anonymous. Hidden services are identified using an onion URL. Figure 1 and
158 the following points detail the protocol for the establishment of communication between an
159 HS and a client in the *current* architecture.

- 160 1. An HS provider contacts a relay and asks them to act as an Introduction Point (IN)⁴.
161 The HS provider receives an acknowledgment from the IN.
- 162 2. The HS provider creates an "Onion service descriptor" which includes the public key for
163 the HS and the IP addresses of its INs. This descriptor is signed by the public key of the
164 HS. The HS provider sends this to a HSDir. The HS provider gets an acknowledgement
165 from one of many Hidden Service Directories (HSDirs)⁵.
- 166 3. A client asks the HSDir for the service descriptor of the HS. They receive and verify the
167 signature of the HS.
- 168 4. The client picks a Tor relay to act as a Rendezvous Point (RP) and establishes a Tor
169 circuit to it. The client gives the RP a Rendezvous cookie.
- 170 5. The client sends the same Rendezvous cookie and the IP address of the RP to an IN.
- 171 6. The IN forwards the cookie and Rendezvous address to the HS provider.
- 172 7. The HS provider makes a Tor circuit with the Rendezvous Point and sends the cookie.
- 173 8. The Rendezvous Point compares the two cookies and, if they match, relays communica-
174 tion from both sides to each other.

³ This includes the code itself, the data, the stack, and the heap.

⁴ IN is chosen to be distinct from Internet Protocol (IP)

⁵ HSDirs are special relays that store and provide hidden service descriptors to the clients.




Figure 1 The current implementation of hidden services. Red paths represent information flow through Tor circuits with 6 nodes. Green paths represent information flow through Tor circuits with three nodes. Blue paths represent information flow through Tor circuits with only two nodes.

175 3 Related work

176 There are two notable prior works that leverage TEE technology to improve Tor. In SGX-
 177 Tor [45], different nodes (such as Directory authority, HSDir, and middle nodes) run the Tor
 178 software in TEEs so that they cannot be modified to collude and launch deanonymization
 179 attacks. In contrast, **VaultTor** reduces the deanonymization threats by not requiring the
 180 HS provider to be online. Furthermore, **VaultTor** is an architectural solution that does not
 181 require any Tor node to run their service within TEE (unless they volunteer to be a vault),
 182 thus maintaining backward compatibility for Tor nodes.

183 Another approach, SmarTor [14], utilizes TEEs and smart contracts to decentralize directory
 184 authorities (which are currently a few trusted and centralized servers) in the Tor
 185 network. The authors argue that this would increase the security of the Tor network as
 186 it would reduce the need to trust particular servers and make it more difficult for govern-
 187 mental authorities to crack them down. **VaultTor**, on the other hand, leverages the privacy
 188 guarantees provided by the TEEs to create trust by allowing a separate, new entity (the
 189 vault) to host web application content, thus making it harder to deanonymize the Hidden
 190 Service provider.

191 Additional prior works [7, 20, 55, 62] have also suggested modifications to the Tor network
 192 to enhance HS anonymity. Sec 10.1 discusses these solutions and their potential use cases.
 193 There exists a whole body of work aiming to improve the Tor network in general. Security
 194 improvements include a layered mesh topology [33] for circuit formation and prevention of
 195 route capture attacks [51]. There has also been work to improve client-side stability in
 196 censored regions using WebRTC [25, 35, 19].

197 Performance improvements include optimizing the Tor path selection algorithm [12, 15,
 198 59, 65] and load balancing techniques [13, 30, 44]. In particular, CenTor [17] considers
 199 serving content from multiple CDN-like content replication nodes to improve the content
 200 delivery time. These improvements can be applied in conjunction with **VaultTor** to further
 201 optimize the anonymity network (see Sec 10.1).

202 **4 Design Goals**

203 Our design goals are meant to ensure enhanced anonymity for HS-provider, while ensuring
 204 that anonymity guarantees for other parties (i.e., client, vault node) remain intact.

205 We will outline our design goals succinctly here:

- 206 1. Ensure that the anonymity guarantees for the HS provider in `VaultTor` are *at least* as
 207 strong as an HS provider in the conventional Tor architecture.
- 208 2. Minimize the uptime for the HS provider, shrinking the attack surface against them.
- 209 3. Ensure that the anonymity guarantees for the newly introduced Vault node are *at least*
 210 as strong as those for an HS provider in the conventional Tor architecture.
- 211 4. Ensure that the anonymity guarantees for the client remain *at least* as strong as those
 212 for a client in the conventional Tor architecture.
- 213 5. Ensure that the HS provider retains full control of any content that is being served on
 214 its behalf.
- 215 6. Ensure that the system is able to serve rich content (not just limited to static content).
- 216 7. Ensure that the system is backward compatible and is built using off-the-shelf components
 217 and technology.

218 **5 Alternate hosting options**

219 The `VaultTor` architecture meets all our design goals specified in Sec 4. Other simple alterna-
 220 tives for anonymous content hosting may not satisfy this requirement. This section considers
 221 two such possible alternatives and describes why they are unable to achieve the full set of
 222 our design goals.

223 **5.1 Using a static content hosting service**

224 An HS provider may upload static content to a data hosting service like an IPFS. While
 225 it is possible that the HS provider is able to upload content securely, this content can not
 226 be modified dynamically and, thus, fails to meet our fifth design goal. For example, an
 227 HS provider trying to run a forum online will not be able to get posts or comments by
 228 clients. `VaultTor` permits dynamic content to be served which allows the HS provider to host
 229 complicated websites on an external server.

230 **5.2 Hosting on the cloud**

231 An HS provider may host data dynamically on a virtual machine present on an external server
 232 (or a public cloud). They may use a remote login tool such as SSH (and the Tor proxy) to
 233 anonymously access the server and upload their content. However, the administrator of the
 234 server (with root access, or the cloud operator) may be able to access and modify this data
 235 even if the HS provider wishes to restrict access. This means the provider does not maintain
 236 full control over the content, violating the fourth design goal.

237 In `VaultTor`, the administrator cannot access data that is present inside the TEE. This
 238 data exists either in an encrypted manner on the disk or is only accessible to the TEE if the
 239 data is in the main memory. Lastly, as both the client and HS provider make an encrypted
 240 channel with the TEE, content is safe while in transit.

241 An HS provider that uploads their content to an external server (via SSH or through
 242 other means) is not guaranteed that this content will be served without modification. The
 243 administrator of this server may modify this content (while keeping the same URL and x509

244 certificate) and serve content that the administrator desires. This is possible because the
 245 administrator can snoop the private key from main memory or storage, presenting a major
 246 threat to client anonymity. On the other hand, **VaulTor** guarantees that the content served
 247 to a client is the content intended by the HS provider. The enclave creates a certificate and
 248 is the only entity (alongside the HS provider) that can serve this content.

249 An HS provider may choose to provision cloud services that enable TEEs and host
 250 their content inside one. In this scenario, the cloud would act as an intermediary, reducing
 251 the uptime requirement for the HS provider. However, provisioning cloud services isn't
 252 anonymous. Know-your-customer (KYC) and/or payments via fiat currency can compromise
 253 the anonymity of the HS provider. It would also necessitate online transactions, which may
 254 be untenable in regions of instability, where volunteer vaults may be the only option.

255 6 Threat Model

256 We make realistic assumptions about the objectives and capabilities of adversaries. An
 257 adversary may have access to a small fraction of Tor relays and sufficient resources to
 258 qualify as a guard relay, HSDir, or a vault. Moreover, a malicious vault may be able to
 259 target specific HSes to host in their machine for the explicit purpose of deanonymizing the
 260 HS provider and clients of that service or for the purpose of modifying that service. Even a
 261 strong adversary will have a limited ability to acquire network traffic from ISPs and monitor
 262 traffic patterns between an end user and a guard node. We wish to protect the anonymity
 263 of the HS provider from this adversary and retain the control of the HS provider over any
 264 content that the HS is providing. In addition, a malicious HS may attempt to deanonymize
 265 a specific vault by hosting their web content on a TEE that the vault runs. As discussed in
 266 sec 2.1, we do not consider side-channel attacks.

267 We also assume that any two entities (amongst the vault, HS provider and client) can
 268 collude to launch an attack against the third. For example, an HS provider and a client may
 269 work together in an attempt to deanonymize a vault.

270 Moreover, we assume a careful HS provider. That is, we assume that the HS provider
 271 does not leak identifying information in the content it provides. Furthermore, at time t , it
 272 will refuse to offload any content to a vault unless it verifies that, at time t , it is connecting
 273 to a valid enclave (this is a realistic assumption as verification is an easy task). Similarly, we
 274 assume a careful vault. That is, a vault inspects the program provided by an HS provider
 275 for malicious code, and will not run this program if it deems it malicious.

276 7 **VaulTor**

277 7.1 Architecture

278 Building upon the conventional Tor hidden services architecture (Sec 2.2), **VaulTor** intro-
 279 duces three new entities: a device which we refer to as a *vault*, a TEE which will host an
 280 *enclave* and an optional external attestation service. The enclave is present inside the vault
 281 (see Figure 2) and utilizes its computational resources.

282 In **VaulTor**, a willing device may offer to host content by advertising itself as a vault. This
 283 can be achieved without modifications to the current Tor architecture. The vault creates
 284 an onion address for itself (we shall refer to this onion website as the *Vault Contact Hidden*
 285 *Service* (VCHS)) to facilitate correspondence with potential HS providers. Vaults may ask
 286 for compensation for hosting an HS (further discussion on incentives is given in Sec 10.3).




Figure 2 Our proposed implementation. Red paths represent information flow through Tor circuits with 6 nodes. The data inside the enclave is secure and information flow through the arrows is encrypted i.e., the vault owner can not interpret it.

287 In this scenario, the vault has the same privacy guarantees that the Hidden Services have
 288 in the current Tor architecture.

289 The HS provider can reach out to a vault through the vault's VCHS and provide a basic
 290 program which we shall refer to a *host program* (HP) that 1) hosts a web server and 2)
 291 provides an interface through which (only) the HS provider can upload or remove content
 292 when they wish to update their website⁶. This program must be running inside an enclave
 293 to provide the privacy guarantees detailed in this architecture. Thus, the HS provider must
 294 verify that the enclave has been correctly set up (i.e., the program is running without any
 295 modifications to the code and within a TEE) *before* it provides all of the content it wishes
 296 to host and any incentives to the vault owner. Future content is also provided through the
 297 same update interface described earlier. This process is outlined in detail in Sec 7.2.2.

298 The TEE guarantees security and privacy for any content the HS provider transfers into
 299 the vault. In addition, it removes the requirement for continuous up-time of the original
 300 HS provider as the enclave can continue servicing the clients. While the vault must remain
 301 online to serve web content, the HS provider does not have this obligation and is hence
 302 protected from various deanonymization attacks (Sec 8).

303 Our architecture ensures increased anonymity and flexibility for the HS provider, minimal
 304 decrease in performance for a client (while maintaining the same security guarantees that
 305 Tor provides), and a level of anonymity for the vaults that is comparable to that of HSes in
 306 the current Tor architecture.

307 7.2 Protocol

308 7.2.1 Host Program (HP) Creation

309 Before an HS provider contacts a vault, it must write a host program that is intended to
 310 run within the TEE. This program should provide the following key functionalities:

⁶ The interface asks for a secret and after it is verified, the content can be uploaded or removed as desired. Since the secret verification occurs within the script that is running within an enclave, the vault cannot tamper with it without being noticed by the HS provider.

311 └─ It should host a server⁷ bound to a port. This server should be able to handle POST
 312 and GET requests. Furthermore, the code should be able to handle requests dynamically.
 313 For example, it should be able to store files uploaded by a client in separate directories.
 314 └─ The HP should, by default, provide an interface on the server to input an authentication
 315 secret. If the secret (which is known to the HS provider) passes the hardcoded verification
 316 in the HP, the HS provider should be allowed to modify the contents of the enclave
 317 directory. The authentication key can either be a password that is hashed and compared,
 318 or it could be the HS provider's private key whose corresponding public key is written in
 319 the HP.
 320 └─ Upon starting the server, the program should generate a *quote* file, that can be verified
 321 by the HS provider or a client (details given in Sec 7.2.2). This quote file is available in
 322 the web directory and can be accessed by the HS provider or the client for verification
 323 through remote attestation as described in Sec 2.1.
 324 └─ The HP should provide some functionality to create backups of the hosted content in
 325 case the vault crashes. Backups are stored on the disk but encrypted using the enclave's
 326 *sealing key*. This sealing key is deterministically generated by the CPU depending on
 327 the program running inside the TEE and the key burnt into the CPU. This key is only
 328 available to the enclave.
 329 Running an arbitrary application inside a TEE is not straightforward. However, a library
 330 OS (libOS), such as Gramine-SGX, facilitates this process with minimal modifications to the
 331 application. Furthermore, since the entire libOS is contained inside the TEE, no inspection
 332 of the application code is necessary. To this end, we assume that a libOS like Gramine-SGX
 333 is available to the vault.

334 7.2.2 Bootstrapping

335 The vault owner hosts and advertises the onion URL of its VCHS. The HS provider, vault
 336 owner and the HP (running inside a TEE) take the following steps to host their service in
 337 the vault (also shown in Figure 3):
 338 (i) The HS provider creates the HP (with the functionality described in Sec 7.2.1).
 339 (ii) The HS provider uploads the HP at the VCHS. It is important to note that the host
 340 program is uploaded in plaintext (either as a script or a binary).
 341 (iii) The vault owner runs the HP inside a TEE, hosts a hidden service (we shall refer to
 342 this as *yourHS.onion*), and binds it to the network port on which the host program will
 343 handle incoming requests.
 344 (iv) The host program will generate a public-private key pair (P_{srv} , S_{srv}) for the hidden
 345 service it provides. P_{srv} is made available to anyone who connects to *yourHS.onion* while
 346 S_{srv} only exists as a variable inside the enclave's memory (and is sealed to the disk for
 347 persistent storage).
 348 (v) The host program generates a quote which represents P_{srv} ⁸ and the program. This quote
 349 is available to anyone accessing *yourHS.onion*.
 350 (vi) The HS provider accesses *yourHS.onion*, retrieves the quote, and verifies that the quote
 351 is legitimate⁹.

⁷ Common servers such as Apache or Nginx can be used.

⁸ If the size of P_{srv} is too large to fit inside the REPORTDATA of the quote, a hash of P_{srv} may be used. If P_{srv} is embedded in a certificate, then typically the certificate hash is used.

⁹ In Intel-SGX, this may involve sending this quote to Intel's online attestation service. In RISC-V Keystone, the end user can verify the quote themselves.

17:10 Dont get caught, keep your Onions in a Vault




Figure 3 Our proposed implementation. Red paths represent 6-node circuits and green paths represent 3-node circuits.

352(vii) The HS provider creates a secure connection with the enclave using P_{srv} and any form
 353 of key-exchange such as Diffie-Hellman. This public key (P_{srv}) may be embedded in a
 354 self-signed x509 certificate in order to facilitate *https* connections.
 355(viii) The HS provider supplies the authentication secret over this secure connection, after
 356 which it can securely upload content (such as HTML, CSS, PHP, and JavaScript files).
 357 (ix) The enclave hosts the uploaded web application content.
 358 (x) Any client that connects to `yourHS.onion` can interact with the hosted web application.
 359 (xi) The enclave regularly encrypts and backs up these files into non-volatile storage using
 360 the sealing key.
 361 This procedure ensures increased anonymity for the HS provider. Their data is hosted
 362 in the vault and the vault owner can not access data inside the enclave or read the traffic
 363 in or out of the enclave. Figure 4 provides a flow diagram of this process detailing some
 364 decisions the HS provider must make while uploading to the vault.

365 7.3 Enclave Isolation

366 To ensure the protection of the vault, the enclave must have limited privileges. The following
 367 conditions, at minimum, are necessary:

368 1. The enclave has access to a fixed, and limited, amount of RAM. One possible way to
 369 achieve this is by running the enclave within a virtual machine¹⁰ configured with limited
 370 memory. This prevents an enclave from occupying all available RAM, thus safeguarding
 371 the performance of other programs on the vault.

372 2. The enclave can only use a fixed amount of persistent memory. This can be achieved
 373 by isolating the enclave in a separate disk partition. This prevents an enclave from
 374 completely filling up persistent disk space that should be available to other programs on
 375 the vault.

376 3. The enclave can only use a fixed amount of network resources. This is possible by
 377 controlling the network traffic rate through Trickle [6].

378 4. Any connections going out from the enclave *must* be restricted to only go via the Tor
 379 Proxy. This can be done by creating rules in iptables [4]. This is necessary to ensure

¹⁰The maximum ram must be set in Intel-SGX at enclave creation time.




Figure 4 The steps taken in order for an HS provider to trust and upload to a vault.

380 that the IP of the vault is not made available to the HS provider. Any traffic that is
 381 *not* going out through port 9050 (the default Tor Proxy port) is blocked by a firewall.
 382 Furthermore, any traffic going to the Host Program must originate from the Tor client
 383 ¹¹.
 384

7.4 Client connection

385 A client must ensure that it is connected to the correct HS identified by its P_{srv} (embedded
 386 in an x509 certificate). To this end, the HS provider distributes not only the onion URL
 387 but also the hash of the x509 certificate when it wishes to advertise its service (similar to
 388 how conventional onion URLs are advertised). When connecting to an HS hosted on a vault,
 389 the client only needs to verify that the x509 certificate supplied by the service matches
 390 its advertised hash to ensure that it is connected to the appropriate entity. The client may
 391 maintain a list of valid certificate hashes¹². Note that the client *never* has to perform remote
 392 attestation themselves.

393 The client in **VaulTor** is exactly the same as a client in the conventional Tor architecture
 394 – they use the same connection protocol. Layered on top of this is the ability to match
 395 an x509 certificate (or its hash) with the one that is advertised by the HS provider. This
 396 matching can be done trivially with a browser extension to the Tor browser.

8 Attack surfaces

397 In this section, we specify the anonymity guarantees the **VaulTor** provides to each of the
 398 three entities: the client, the vault, and the HS provider. We consider the scenarios where

¹¹This is to done to preserve vault anonymity from a malicious HS provider as discussed in 8.3.

¹²Checking of certificate hash is trivial and may be added as a subroutine in the client's Tor browser.

17:12 Dont get caught, keep your Onions in a Vault

400 each of these can be malicious as well as the scenario in which two of them collaborate
401 to deanonymize the third. **VaulTor** enhances HS provider anonymity and leaves the client
402 anonymity as it is. We further show that our new actor—the vault—is as protected as an
403 HS provider in the current Tor Hidden Services architecture.

404 8.1 HS provider Anonymity

405 406 **Scenario 1: Malicious client**

407 The HS provider no longer interacts with the clients directly. To access the web application
408 content, the clients now establish a connection with the vault instead. Thus, unless the
409 HS provider leaks identifying information in their content, they are safe from deanonymization
410 at the hands of a client.

411 **Scenario 2: Malicious vault**

412 In the traditional Tor design, where the client and HS provider communicate over a two-
413 way anonymous channel initiated by the client, the attack scenarios in Table 1 render the
414 client a harder anonymity target by a malicious HS provider than an honest HS provider by
415 a malicious client. A mirror situation exists in the **VaulTor** design where the HS provider
416 uploads and updates the content on a vault over a two-way anonymous channel. Thus,
417 even in the worst case, an HS provider in **VaulTor** is as anonymous as an HS provider in
418 the traditional architecture. Furthermore, the minimal uptime requirement enhances the
419 anonymity of the HS provider in **VaulTor** architecture. The use of TEE at the vault offers
420 additional guarantees of data integrity and data confidentiality to the HS provider.

421 **Scenario 3: Vault and client collude**

422 The attack opportunities open to a client are a subset of the attack opportunities possible
423 for a vault (since a vault has the same privileges as a client and more). Thus, the protection
424 guaranteed for an HS provider from the vault applies in a scenario where the vault and client
425 may collude.

426 8.2 Client privacy

427 We now show that a client is as protected in **VaulTor** as they are in the current Tor architecture.

429 **Scenario 1: Malicious HS provider**

430 The HS provider is completely disconnected from the client, and hence is unable to
431 launch attacks on the client directly.

432 **Scenario 2: Malicious vault**

433 A malicious vault may attempt to a) serve modified content or b) launch a deanonymization
434 attack on the clients. We now show why these attacks are not feasible in our architecture:
435 a) As the content is being hosted inside an enclave, clients can ensure that any content
436 being served by the vault has not been maliciously modified. Since the x509 certificate
437 is generated by the HP running inside a TEE and the corresponding private key (S_{srv})
438 is *only* available to the TEE and the HS provider, a secure connection established using
439 the certificate is guaranteed to be serving content vetted by the HS provider. Another
440 consideration is that the content being hosted inside a vault is regularly encrypted and
441 backed up to the disk. While this backed-up data can not be modified, a vault owner can
442 selectively delete this backed-up data and restart the program in the enclave. This may
443 result in the enclave accidentally serving outdated data to clients. However, if pieces of

444 content are properly timestamped, the TEE can refuse to serve content that is outdated
 445 or add warnings while serving this content.

446 b) In **VaulTor**, a client's perspective of the Hidden Service architecture remains the same. A
 447 client still accesses content through a 6-node connection—except that instead of connecting
 448 to the HS provider, it connects to a vault. Thus the client enjoys the same privacy
 449 guarantees as a client in the conventional Tor architecture. Furthermore, as discussed in
 450 sec 10.8, the client may enjoy enhanced data privacy.

451 **Scenario 3: Vault and HS provider collude**

452 A colluding vault and HS provider in the **VaulTor** architecture have the same attack
 453 opportunities against a client as a malicious HS provider in the conventional Tor architecture.
 454 Thus, the deanonymization risk for the client is the same as when only the vault is malicious.

455 **8.3 Vault privacy**

456 **VaulTor** introduces a new entity in the Tor architecture: a vault that assumes a role similar
 457 to that of an HS in the current Tor design, thus maintaining a similar level of protection
 458 against deanonymization attacks.

459 **Scenario 1: Malicious client**

460 A vault is as vulnerable to a client as a Hidden Service is in the traditional Tor network.
 461 One may even argue that the vault has stronger anonymity guarantees due to the fact that
 462 the web application content is hosted within an enclave (a secret, protected environment
 463 the vault cannot modify). This may provide plausible deniability to the vault as it would
 464 be blind to the traffic that enters and leaves its enclave.

465 **Scenario 2: Malicious HS provider**

466 A vault is protected from attacks launched by an HS provider through the Host Program
 467 by ensuring that the safety criteria specified in Sec 7.3 are satisfied.

468 Consider a malicious HS provider who uploads code that tries to obtain identifying
 469 information about the vault by leveraging the fact that the enclave is utilizing the vault's
 470 hardware. For example, this code may try to read files belonging to the vault or the vault's
 471 OS in order to *directly* find identifying information or it may try to obtain its IP *indirectly* by
 472 pinging an external server. The security guarantees provided by TEEs make direct attempts
 473 impossible; the host is *also* isolated from the TEE just as the TEE is isolated from the host.
 474 Furthermore, most applications for TEEs run in a VM (as is the default in gramine-SGX [1]),
 475 adding to the isolation. Indirect attacks are also mitigated using a firewall. By ensuring
 476 that all outgoing traffic is ported through port 9050 (the default port for Tor), only the IP
 477 of the exit node is available to the Host Program.

478 Furthermore, by ensuring that all requests originate from the Tor client software on the
 479 vault, the vault is protected from pinging based attacks. If this is not done, a malicious HS
 480 provider may upload a simple Host Program that replies with a unique phrase to the HS
 481 provider's IP when this Host Program is pinged. The HS provider may then ping various
 482 candidate IPs in the hopes of stumbling upon the vault's IP which would reply with the
 483 phrase. This sort of attack is only possible if the Host Program can be pinged directly, from
 484 outside the Tor client software.

485 **Scenario 3: Client and HS Provider collude**

486 The HS provider is in a unique position as it directly provides code that the vault runs
 487 within an enclave. If an entity controls both the HS provider and a client node, we consider
 488 an attempt to launch a watermarking attack (described in Sec 9.1). This attack, in particular,
 489 only requires the control of two entities connected to the third. Moreover, the fact that the

490 HS provider and client can make repeated on-demand requests to the vault further benefits
 491 the viability of this attack.

492 To attempt to launch a watermarking attack (similar to what [38, 39] describe), the HS
 493 provider would add some watermark to the Tor traffic that can be identified at the client
 494 end. Despite controlling both the HS provider and a client, an attacker would not have
 495 the ability to successfully launch a watermarking attack on the vault. This is because of a
 496 missing component that this attack requires: the control of a guard relay. If we also assume
 497 the control of a guard relay, the HS provider is no longer necessary, as the guard relay can
 498 be the entity that watermarks the traffic. A guard relay and client could potentially launch
 499 this attack on their own without the requirement of an HS provider. Hence, the vault is as
 500 protected from a watermarking attack as an HS is in the current architecture. Scenario 3 is
 501 thus akin to having two malicious clients in the conventional Tor architecture.

502 9 Evaluation

503 In this section, we will qualitatively evaluate the effect of `VaulTor` in deflecting various
 504 families of existing attacks on HSes from the HS provider to our new Vault node. Afterwards,
 505 we will quantitatively measure the performance impact of `VaulTor` on client side network
 506 performance.

507 9.1 Known Attacks Deflected

508 We will list various attacks and briefly explain how the `VaulTor` architecture deflects them
 509 from the HS provider to the vault. This list is non-exhaustive yet exemplifies the prominent
 510 attacks in recent years.

511 **Clock Skew:** These attacks rely on repeatedly sending requests to an HS in order to
 512 heat up its CPU which has some tangible effect on the timestamp of incoming packets [60, 52].
 513 In `VaulTor`, this is impossible. *No one* can repeatedly send packets to the HS provider.

514 **Congestion:** This attack relies on an adversary congesting existing guard nodes in the
 515 network [42], forcing the HS to connect to their compromised guard node long enough for
 516 the adversary to correlate traffic. This attack is completely deflected in `VaulTor`; the HS
 517 provider is sporadically online for limited periods and an adversary would have to congest
 518 the network indefinitely.

519 **Fingerprinting:** These attacks rely on learning the traffic patterns of an HS and ref-
 520 erencing this against the traffic of a candidate set of guard nodes [46, 57, 53, 36]. In our
 521 architecture, the vault is serving the traffic while the HS provider is taciturn. Thus, this
 522 type of attack will not work on the HS provider. Similar attacks that rely on compromised
 523 middle nodes [41] are similarly deflected.

524 **Guard Node Discovery:** The Tor developers currently consider this the most potent
 525 threat against hidden services [7]. This attack relies on making multiple connections with
 526 the HS provider such that their malicious middle node is next to the HS provider's guard
 527 node. Repeated, on-demand connections with the HS provider are impossible in `VaulTor`.
 528 As such, this attack is eliminated.

529 **Location Leaks:** Such attacks rely on the negligence of the HS provider and are out of
 530 the scope of this paper [49].

531 **Watermarking:** In this type of attack, an adversary watermarks traffic on the client
 532 side in order to detect it at the malicious guard node of the HS provider [38, 39]. If a
 533 malicious vault node tries to launch this attack on the HS provider, this attack would be

534 rendered less effective because HS-provider would have minimal uptime connection with the
 535 vault code instead of a constant connection.

536 Table 1 shows the various scenarios in which an adversary can launch a deanonymizing
 537 attack on Tor Hidden Services along with the impact of **VaultTor** on these attacks.

Scenario	Attack Categories	VaultTor impact
Adversary can send arbitrary requests to the HS provider	Clock Skew, Watermarking, Guard Node Discovery, Fingerprinting	Scenario eliminated
HS Provider has high uptime	Clock Skew, Watermarking, Congestion, Fingerprinting	Scenario diminished
High volume of traffic coming from the HS provider	Watermarking, Congestion, Fingerprinting	Scenario diminished

538 **Table 1** Various scenarios that lead to categories of contemporary attacks on Tor Hidden Services along with the impact of **VaultTor** on these scenarios.

9.2 Performance

539 It is important that the security improvements **VaultTor** brings do not significantly degrade
 540 client side network performance. Important client side performance metrics include the
 541 network latency and the throughput. In this section, we detail our experimental setup
 542 (which includes our implementation of a vault) and our evaluation of these metrics. Our
 543 experiments measure the time experienced by the client to retrieve the data from an HS; the
 544 registration and bootstrapping processes in **VaultTor** occur only once and have a negligible¹³
 545 performance impact in the overall lifetime of the HS.

9.2.1 Experimental setup

546 To measure and compare the network performance of Hidden Services when hosted within
 547 and outside a TEE, we ran two instances of a Host Program on the same machine. One of
 548 these HPs ran inside a TEE (facilitated by the gramine-SGX library OS [1]) while the other
 549 HP (which we shall refer to as a *vanilla* HP) ran outside a TEE. A Tor client¹⁴ was also
 550 launched on the same machine which generated two onion URLs: one for the enclave and
 551 one for the vanilla HP. The Tor client directs traffic for each of these onion URLs to their
 552 respective HPs, allowing them to serve content via Tor.

553 Both the enclave and the vanilla HP ran web servers and, in order to ensure consistency,
 554 served the same landing webpage simultaneously. In addition, the HP running inside a TEE
 555 had the ability to generate a quote in order to facilitate remote attestation. Both web servers
 556 were written in Python3 and regularly backed up data to persistent memory. Moreover, these
 557 webpages were hosted on the same device with an SGX-enabled Intel processor (Core-i5
 558 10210U).

559 We conducted these experiments with three webpages, each with varying page sizes
 560 (0.5kB, 50kB, and 5000kB). The content on these webpages included HTML and JavaScript.

¹³This performance impact will only be negligible if uploading content does not significantly increase the uptime of HS provider.

¹⁴This client is not the same as a client in Tor architecture which we have discussed above. This is a program necessary to interact with the Tor network.

562 We measured the performance using three methods:

563 1. *Random Relays*: We restarted the Tor application between each measurement to establish
564 fresh circuits. This gave us three random relays for every measurement.


565 2. *Fixed Relays*: We used fixed/constant relays¹⁵ across webpages for both **VaultTor** and
566 Tor. We report the average performance of three different fixed circuits.

567 3. *Local*: We locally accessed the webpages.

568 Each webpage was loaded 250 times in each of the methods, save for the Fixed Relays method,
569 for which we loaded each webpage 250 times on each circuit (a total of 750 measurements)
570 and took the average of the results. Methods (1) and (2) were conducted on the actual Tor
571 network.

572 We thus quantified any overall changes in performance caused specifically by hosting an
573 HS within a TEE.

574 9.2.2 Results



575 **Figure 5** Time to first byte (TTFB) and time to last byte (TTLB) for webpages with varying page
576 sizes without and within a TEE. Error bars represent 99% confidence intervals.

577 Figure 5 shows the time to first byte (TTFB) and time to last byte (TTLB) for the 3
578 webpages hosted in the 2 architectures (the current architecture and **VaultTor**).

579 We note a minimal difference in performance across webpages and testbeds. Note that,
580 for most of the results collected over Tor, the average TTFB and TTLB in the **VaultTor**
581 architecture fall within the 99% confidence interval of those of the current architecture. The
582 only result (collected over the Tor network) that lied outside the confidence interval was for
583 the TTLB of a 5000kB webpage routed through fixed relays. This had an average increase
584 of 2.9%.

585 If we consider all the results, including those that lie within the confidence intervals, we
586 note a maximum increase in TTFB and TTLB of 5.7% (5000kB, Random Relays) and 2.9%
587 (5000kB, Fixed Relays), respectively.

588 It should be noted that running an arbitrary program inside an Intel-SGX TEE may have
589 a non-negligible computation overhead. When accessing the webpages locally (and hence,
590 not over the Tor network), we note a maximum percentage increase in time in the case of
591 the TTFB of a webpage of size 50kB (15.9%). However, this delay is negligible compared
592 to delays caused by Tor's network latency. As such, it is not surprising that the percentage

593 ¹⁵These were chosen randomly from advertised Tor relays here: <https://www.dan.me.uk/tornodes>

591 performance overhead of **VaulTor** over the conventional Tor architecture is minimal when
592 measured over the real Tor network.

593 We believe this nominal decrease in performance is justified considering the major an-
594 onymity benefits **VaulTor** brings to the HS provider.

595 **9.2.3 Ethics**

596 We had ethical considerations while conducting our performance measurements. We solely
597 collected timing information and the size of the files we downloaded. We did not store the
598 IP addresses of the entry and exit nodes, so as to preserve their anonymity. In addition,
599 our load on the Tor network was negligible. We ensured this by conducting the experiments
600 serially, and not in parallel, to minimize the load at any given time. To the best of our
601 knowledge, we did not hinder any other users' experiences on the Tor network.

602 **10 Discussion**

603 **10.1 Additional Anonymity Measures**

604 There exist a number of proposals (such as [7, 20, 55, 62]) that enhance the anonymity of
605 the HS provider. These solutions, however, result in degraded network performance (longer
606 delays and lower throughput) for the client, when used in the conventional HS architecture.
607 This reduction in network performance renders these solutions less attractive today. With
608 the **VaulTor** architecture, the vault serves the content to the clients, and any retrofitting at
609 the HS provider side has no bearing on the network performance experienced by the client.

610 For example, the Vanguard add-on [7] (which inserts additional hops in the connection)
611 may be used by the HS provider without affecting client-side performance. Similarly, privacy-
612 preserving path selection methods [20, 55, 62] may incur latency costs but are a non-issue
613 for an HS provider that connects with the vault infrequently.

614 Furthermore, techniques like temporary proxies are completely compatible with our sys-
615 tem and may be used by clients and the HS provider to connect to vaults to obfuscate their
616 traffic [25, 35, 19].

617 **10.2 HS Provider Flexibility**

618 In the current Tor architecture, the HS provider must remain static in order to serve content.
619 The flexibility offered by **VaulTor** can be leveraged by the HS provider to communicate from
620 secure and variable locations. This would especially be beneficial in the context of activists or
621 journalists who want to report their content from secure *intermediate* locations in oppressive
622 regimes without the risk of getting caught. Moreover, **VaulTor** would allow the HS provider's
623 content to remain accessible during Internet outages, which is commonplace in regions with
624 political instability and censorship [21, 22].

625 **10.3 Incentives for Vault Node**

626 **10.3.1 Monetary Incentives**

627 The vault owner proxies for the HS provider and, on its behalf, serves content to the clients.
628 This act must be incentivized for the vault owner. These incentives may be social incentives
629 – similar to how users of Tor run relays and nodes.

17:18 Dont get caught, keep your Onions in a Vault

630 However, if incentives are monetary, they must be exchanged in a secure and private
631 manner. Towards this end, a blockchain may be used to ensure that the vault owner receives
632 cryptocurrency rewards for hosting content for the HS provider. One approach to this is
633 that the vault owner supplies the HS provider with their address on a public blockchain such
634 as Ethereum [67]. Only if regular cryptocurrency payments are made to this public address
635 does the vault owner continue hosting. This allows both the HS provider and the clients to
636 “crowdsource” an HS on a vault.

637 This previous approach does necessitate timely payments from the HS provider. This
638 requirement can be removed via the use of smart contracts. A smart contract can lock
639 the cryptocurrency that it receives from the HS provider and clients. The smart contract
640 can then use an oracle to verify that the HS is being hosted properly and perform remote
641 attestation. If the HS is being hosted properly and the remote attestation is successful,
642 the smart contract releases the cryptocurrency to the vault’s blockchain address. In order
643 to preserve privacy, zero knowledge enabled cryptocurrency such as Zcash [23] can be used.
644 Furthermore, a Decentralized Exchange (DEX) may be used to trade cryptocurrency. These
645 measures reduce the possibility of profiling based attacks.

646 For example, a vault owner may make their zero-knowledge blockchain address (such
647 as for monero [56]) available on VCHS and the HS provider may anonymously transfer this
648 cryptocurrency by publishing their transaction by connecting to an external blockchain node
649 through the Tor proxy. In a zero-knowledge blockchain, the transaction itself will have no
650 identifying information present in it that may be used for social engineering (such as, by
651 using chain analysis tools).

652 10.3.2 Altruistic Reasons

653 Vault operation may also be motivated by altruistic considerations rather than direct fin-
654 ancial compensation. This model closely mirrors the operation of Tor relays, where parti-
655 cipants voluntarily contribute resources to enhance privacy, censorship resistance, and overall
656 network resilience without receiving monetary rewards. While operating a Tor relay may
657 incidentally support unlawful activities, it also enables many socially beneficial uses, such
658 as protecting free expression and providing access to information under restrictive condi-
659 tions. Similarly, vault nodes may be used for both benign and potentially abusive purposes;
660 however, their primary value lies in supporting privacy preserving and censorship resistant
661 services.

662 Similar altruistic participation models exist in other systems, including running VPN
663 or proxy servers for community use, and operating public Network Time Protocol (NTP)
664 servers. In these cases, operators are motivated by a desire to support open infrastructure,
665 enhance collective security, or contribute to public-good Internet services.

666 10.4 Plausible Deniability for Vault

667 In the **VaulTor** architecture, a vault owner is not privy to the content present inside the TEE.
668 We believe that this adds an extra layer of plausible deniability, greater than the plausible
669 deniability of conventional data hosting services that are aware of the content being served.
670 When hosting content on behalf of an HS provider, the only thing the vault owner knows
671 is the Onion URL of the Hidden Service. This has a parallel with Guard Nodes in the
672 conventional Tor architecture that know what Onion URL’s traffic is routed through them.
673 Both the vault and the guard node can not read this traffic or compromise its integrity, only
674 help move this content. The only difference is that the physical storage resources of the

675 Vault owner are being used. However, even this physical storage is encrypted and opaque to
 676 the Vault owner which is not privy to the information being served, just like a guard node.

677 In our future work, we can enable vault owners to run a Tor client inside an enclave and
 678 run the HP inside another enclave on the same machine. The Tor client generates a new
 679 onion URL and shares it (only) with the Host Program. The HP then serves content on
 680 this onion URL using the Tor client as a proxy. The HP also forwards the onion URL to
 681 the HS provider who can then advertise it as before. In this scenario, the client does not
 682 need to modify their browser. As such, the vault owner is not aware of the content they are
 683 serving. In this scenario, they can not be held liable for the content they are serving as the
 684 information about which machine is serving what content is available to “*no one*”. And “*no*
 685 *one*” includes the vault owner and the HS provider.

686 10.5 Legal and deployability considerations

687 **VaulTor** is a technical design that aims to improve the anonymity and availability of hidden
 688 services. While it does not change Tors underlying trust or threat model, legal and regulatory
 689 constraints may nonetheless become a practical barrier to deployment. For example, a vault
 690 operator could face legal risk for hosting third-party content that is alleged to be unlawful
 691 (e.g., infringing or prohibited content), even if the operator cannot readily inspect the hosted
 692 state. In addition, hidden service operators remain responsible for the services and content
 693 they publish, regardless of whether hosting is delegated to a TEE-hosted vault node.

694 For vault operators, legal exposure is jurisdiction-dependent and may hinge on how local
 695 law treats third-party hosting, infrastructure provision, and duties arising from notice or
 696 investigation. Thus, even though **VaulTor** is designed so that the vault cannot inspect the
 697 encrypted hidden service content state, prospective vault operators should treat legal risk
 698 as non-negligible and evaluate deployability under applicable local law (e.g., jurisdictional
 699 choice, operational policies, and whether participation is restricted to vetted deployments).

700 10.6 Incremental deployment:

701 The **VaulTor** architecture supports incremental deployment (albeit its strength is fully utilized
 702 when there are many vaults present on Tor). Vaults can register their VCHS themselves
 703 to an HSDir similar to how Hidden Services are currently already registered. This allows
 704 for a slow, optional adoption of **VaulTor**.

705 The client *does* need to install a small extension (as discussed in Sec 7.4) that compares
 706 P_{srv} (or its hash) with the one advertised by the HS provider but this is a trivial add-on
 707 and does not affect traditional HSes.

708 10.7 Multiple Vaults:

709 An HS provider may commission multiple vaults to hold their data. To this end, they may
 710 download the S_{srv} and the x509 certificate from the HP of one vault and upload it to an HP
 711 they have hosted on another vault. As clients use the certificate hash provided by an HS to
 712 validate its identity (as described in Sec 7.4), they can be certain they are being served by
 713 the same HS provider even if the onion URL of the HS is different. This will add redundancy
 714 and fault tolerance to the HS provider’s content: if one vault becomes inactive, the other
 715 vaults can continue to serve content.

716 **10.8 Strengthening client data privacy**

717 In the **VaulTor** architecture, the host program is present inside a TEE and its measurement
 718 (see Sec 2.1) is available to the client (in addition to the HS provider). The HS provider
 719 may elect to make the code itself available to the client, allowing the client to inspect this
 720 code. If the code is simple (for example, the code only stores and serves content to password
 721 authenticated requests), then the client can upload private data to the server *without having*
 722 *to trust the HS provider* as is necessary in the current Tor architecture.

723 **10.9 Attacks against TEEs:**

724 A wide variety of side-channel attacks exist that can target TEEs [68, 64, 61, 27, 54]. These
 725 attacks aim to discern secrets contained inside the TEE, such as private keys, through
 726 various means such as leveraging page faults. Vendors are prompt in mitigating side-channel
 727 attacks [10] as the community uncovers those. Considering the research and development
 728 that focuses on mitigating side-channel attacks [66, 43, 47, 37], architectural designs discount
 729 such attacks from their threat models [63, 11, 50]. We also follow take the same course of
 730 action (Sec 6).

731 **10.10 Advancements/variations in TEE technology:**

732 Although our current implementation utilizes Intel-SGX, **VaulTor** is a generic solution that
 733 could theoretically support any TEE service. As newer TEE services (such as Intel-TDX [3])
 734 emerge and improve, the strength and flexibility of **VaulTor** improve as well. In the future,
 735 a diverse set of vaults utilizing differing TEE technologies could be built and tested. Other
 736 promising implementations of TEEs are also being proposed [18, 26, 29, 34, 48]. We believe
 737 that TEEs will become increasingly resistant to side-channel attacks.

738 **11 Conclusion**

739 We present **VaulTor** as an architectural solution that leverages TEE technology to reduce
 740 the threat of deanonymization attacks against HS providers on the Tor network. To this end,
 741 **VaulTor** introduces a new actor: the vault, which serves content on the HS provider's behalf.
 742 We show that **VaulTor** prevents several HS deanonymization attacks by utilizing the vault,
 743 whilst preserving the same level of client anonymity as in the current architecture. This
 744 is achieved without any noticeable performance degradation experienced by the client. We
 745 also argue that vaults have the same security guarantees as HS providers in the conventional
 746 Tor architecture.

747

 References

748 1 Gramine, <https://gramineproject.io/>.

749 2 How Do Onion Services Work?, <https://community.torproject.org/onion-services/overview/>.

750 3 Intel's Trust Domain Extensions (Intel's TDX), <https://www.intel.com/content/www/us/en/developer/articles/technical/intel-trust-domain-extensions.html>.

751 4 iptables(8) - Linux man page, <https://linux.die.net/man/8/iptables>.

752 5 Tor, <https://www.torproject.org/>.

753 6 trickle(1) - Linux man page, <https://linux.die.net/man/1/trickle>. 2002.

754 7 Announcing the Vanguards Add-On for Onion Services, <https://blog.torproject.org/announcing-vanguards-add-onion-services/>. 2014.

755 8 Global action against dark markets on Tor network, <https://www.europol.europa.eu/newsroom/news/global-action-against-dark-markets-tor-network>. 2014.

756 9 Users - Tor Metrics, <https://metrics.torproject.org/userstats-relay-country.html>. 2022.

757 10 2023.

758 11 Adil Ahmad, Juhee Kim, Jaebaek Seo, Insik Shin, Pedro Fonseca, and Byoungyoung Lee. Chancel: Efficient multi-client isolation under adversarial programs. 01 2021.

759 12 Masoud Akhoondi, Curtis Yu, and Harsha V. Madhyastha. Lastor: A low-latency as-aware tor client. *IEEE/ACM Trans. Netw.*, 22(6):17421755, dec 2014.

760 13 Mashaal AlSabah, Kevin Bauer, Tariq Elahi, and Ian Goldberg. The path less travelled: Overcoming tor's bottlenecks with traffic splitting. In Emiliano De Cristofaro and Matthew Wright, editors, *Privacy Enhancing Technologies*, pages 143–163, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.

761 14 Greubel Andre, Dmitrienko Alexandra, and Kounev Samuel. Smartor: Smarter tor with smart contracts: Improving resilience of topology distribution in the tor network. In *Proceedings of the 34th Annual Computer Security Applications Conference*, ACSAC '18, page 677691, New York, NY, USA, 2018. Association for Computing Machinery.

762 15 Robert Annessi and Martin Schmiedecker. Navigator: Finding faster paths to anonymity. In *2016 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 214–226, 2016.

763 16 ARM Limited. Security technology: building a secure system using trustzone technology http://infocenter.arm.com/help/topic/com.arm.doc.prd29-genc-009492c/PRD29-GENC-009492C_trustzone_security_whitepaper.pdf. 2018.

764 17 Arushi Arora and Christina Garman. Improving the performance and security of tor's onion services. *Proceedings on Privacy Enhancing Technologies*, 2025:531–552, 01 2025.

765 18 Raad Bahmani, Ferdinand Brasser, Ghada Dessouky, Patrick Jauernig, Matthias Klimmek, Ahmad-Reza Sadeghi, and Emmanuel Stapf. CURE: A security architecture with CUstomizable and resilient enclaves. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 1073–1090. USENIX Association, August 2021.

766 19 Diogo Barradas, Nuno Santos, Luís Rodrigues, and Vítor Nunes. Poking a hole in the wall: Efficient censorship-resistant internet communications by parasitizing on webrtc. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, CCS '20, page 3548, New York, NY, USA, 2020. Association for Computing Machinery.

767 20 Armon Barton and Matthew Wright. Denasa: Destination-naive as-awareness in anonymous communications. *Proceedings on Privacy Enhancing Technologies*, 2016(4):356–372, 2016.

768 21 David Belson. A recent spate of Internet disruptions, <https://blog.cloudflare.com/a-recent-spate-of-internet-disruptions-july-2024/>. 2024.

769 22 David Belson. Q2 2024 Internet disruption summary, <https://blog.cloudflare.com/q2-2024-internet-disruption-summary/>. 2024.

770 23 Eli Ben-Sasson, Alessandro Chiesa, Christina Garman, Matthew Green, Ian Miers, Eran Tromer, and Madars Virza. Zerocash: Decentralized anonymous payments from bitcoin. Cryptology ePrint Archive, Paper 2014/349, 2014. <https://eprint.iacr.org/2014/349>.

17:22 Dont get caught, keep your Onions in a Vault

798 24 Alex Biryukov, Ivan Pustogarov, and Ralf-Philipp Weinmann. Trawling for tor hidden ser-
799 vices: Detection, measurement, deanonymization. In *2013 IEEE Symposium on Security and*
800 *Privacy*, pages 80–94, 2013.

801 25 Cecylia Bocovich, Arlo Breault, David Fifield, Serene, and Xiaokang Wang. Snowflake, a
802 censorship circumvention system using temporary WebRTC proxies. In *33rd USENIX Se-*
803 *curity Symposium (USENIX Security 24)*, pages 2635–2652, Philadelphia, PA, August 2024.
804 USENIX Association.

805 26 Ferdinand Brasser, David Gens, Patrick Jauernig, Ahmad-Reza Sadeghi, and Emmanuel Stafp.
806 Sanctuary: Arming trustzone with user-space enclaves. 01 2019.

807 27 Jo Van Bulck, Nico Weichbrodt, Rüdiger Kapitza, Frank Piessens, and Raoul Strackx. Telling
808 your secrets without page faults: Stealthy page Table-Based attacks on enclaved execution.
809 In *26th USENIX Security Symposium (USENIX Security 17)*, pages 1041–1056, Vancouver,
810 BC, August 2017. USENIX Association.

811 28 Nicolas Christin. Traveling the silk road: a measurement analysis of a large anonymous online
812 marketplace. *Proceedings of the 22nd international conference on World Wide Web*, 2013.

813 29 Victor Costan, Ilia Lebedev, and Srinivas Devadas. Sanctum: Minimal hardware extensions
814 for strong software isolation. In *Proceedings of the 25th USENIX Conference on Security*
815 *Symposium*, SEC'16, page 857874, USA, 2016. USENIX Association.

816 30 Hussein Darir, Hussein Sibai, Chin-Yu Cheng, Nikita Borisov, Geir Dullerud, and Sayan
817 Mitra. Mleflow: Learning from history to improve load balancing in tor. *Proceedings on*
818 *Privacy Enhancing Technologies*, 2022:75–104, 01 2022.

819 31 Tom Woller David Kaplan, Jeremy Powell. AMD memory encryption, https://developer.amd.com/wordpress/media/2013/12/AMD_Memory_Encryption_Whitepaper_v7-Public.pdf. 2016.

821 32 Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The Second-Generation onion
822 router. In *13th USENIX Security Symposium (USENIX Security 04)*, San Diego, CA, August
823 2004. USENIX Association.

824 33 Joan Feigenbaum, Aaron Johnson, and Paul Syverson. Preventing active timing attacks in
825 low-latency anonymous communication. pages 166–183, 07 2010.

826 34 Andrew Ferraiuolo, Andrew Baumann, Chris Hawblitzel, and Bryan Parno. Komodo: Using
827 verification to disentangle secure-enclave hardware from software. In *Proceedings of the 26th*
828 *Symposium on Operating Systems Principles*, SOSP '17, page 287305, New York, NY, USA,
829 2017. Association for Computing Machinery.

830 35 Gabriel Figueira, Diogo Barradas, and Nuno Santos. Stegozoa: Enhancing webrtc covert
831 channels with video steganography for internet censorship circumvention. In *Proceedings of*
832 *the 2022 ACM on Asia Conference on Computer and Communications Security*, ASIA CCS
833 '22, page 11541167, New York, NY, USA, 2022. Association for Computing Machinery.

834 36 Jamie Hayes and George Danezis. k-fingerprinting: A robust scalable website fingerprinting
835 technique. In *25th USENIX Security Symposium (USENIX Security 16)*, pages 1187–1203,
836 Austin, TX, August 2016. USENIX Association.

837 37 Zhou Hongwei, Ke Zhipeng, Zhang Yuchen, Wu Dangyang, and Yuan Jinhui. Tsgx: De-
838 feating sgx side channel attack with support of tpm. In *2021 Asia-Pacific Conference on*
839 *Communications Technology and Computer Science (ACCTCS)*, pages 192–196, 2021.

840 38 Alfonso Iacovazzi, Daniel Frassinelli, and Yuval Elovici. The DUSTER attack: Tor onion
841 service attribution based on flow watermarking with track hiding. In *22nd International*
842 *Symposium on Research in Attacks, Intrusions and Defenses (RAID 2019)*, pages 213–225,
843 Chaoyang District, Beijing, September 2019. USENIX Association.

844 39 Alfonso Iacovazzi, Sanat Sarda, and Yuval Elovici. Inflow: Inverse network flow watermarking
845 for detecting hidden servers. In *IEEE INFOCOM 2018 - IEEE Conference on Computer*
846 *Communications*, pages 747–755, 2018.

847 40 Intel. Intel software guard extensions programming reference, <https://software.intel.com/sites/default/files/managed/48/88/329298-002.pdf/>. 2014.

849 41 Rob Jansen, Marc Juarez, Rafael Galvez, Tariq Elahi, and Claudia Diaz. Inside job: Applying
850 traffic analysis to measure tor from within. 01 2018.

851 42 Rob Jansen, Florian Tschorsch, Aaron Johnson, and Björn Scheuermann. The sniper attack:
852 Anonymously deanonymizing and disabling the tor network. In *NDSS*, 2014.

853 43 Yuan Jinhui, Zhou Hongwei, and Zhang Laishun. F-sgx: Next generation sgx for trusted
854 computing. In *2021 IEEE Asia-Pacific Conference on Image Processing, Electronics and*
855 *Computers (IPEC)*, pages 673–677, 2021.

856 44 Aaron Johnson, Rob Jansen, Nicholas Hopper, Aaron Segal, and Paul Syverson. Peerflow:
857 Secure load balancing in tor. *Proceedings on Privacy Enhancing Technologies*, 2017, 04 2017.

858 45 Seongmin Kim, Juhyeng Han, Jaehyeong Ha, Taesoo Kim, and Dongsu Han. Sgx-tor: A
859 secure and practical tor anonymity network with sgx enclaves. *IEEE/ACM Transactions on*
860 *Networking*, 26(5):2174–2187, 2018.

861 46 Albert Kwon, Mashael AlSabah, David Lazar, Marc Dacier, and Srinivas Devadas. Circuit
862 fingerprinting attacks: Passive deanonymization of tor hidden services. In *24th USENIX*
863 *Security Symposium (USENIX Security 15)*, pages 287–302, Washington, D.C., August 2015.
864 USENIX Association.

865 47 Fan Lang, Wei Wang, Lingjia Meng, Jingqiang Lin, Qiongxiao Wang, and Linli Lu. Mole:
866 Mitigation of side-channel attacks against sgx via dynamic data location escape. In *Pro-
867 ceedings of the 38th Annual Computer Security Applications Conference*, ACSAC '22, page
868 978988, New York, NY, USA, 2022. Association for Computing Machinery.

869 48 Dayeol Lee, David Kohlbrenner, Shweta Shinde, Krste Asanović, and Dawn Song. Keystone:
870 An open framework for architecting trusted execution environments. In *Proceedings of the*
871 *Fifteenth European Conference on Computer Systems*, EuroSys '20, New York, NY, USA,
872 2020. Association for Computing Machinery.

873 49 Srdjan Matic, Platon Kotzias, and Juan Caballero. Caronte: Detecting location leaks for
874 deanonymizing tor hidden services. In *Proceedings of the 22nd ACM SIGSAC Conference*
875 *on Computer and Communications Security*, CCS '15, page 14551466, New York, NY, USA,
876 2015. Association for Computing Machinery.

877 50 J. Menetrey, M. Pasin, P. Felber, and V. Schiavoni. Twine: An embedded trusted runtime
878 for webassembly. In *2021 IEEE 37th International Conference on Data Engineering (ICDE)*,
879 pages 205–216, Los Alamitos, CA, USA, apr 2021. IEEE Computer Society.

880 51 Prateek Mittal and Nikita Borisov. Shadowwalker: Peer-to-peer anonymous communica-
881 tion using redundant structured topologies. In *Proceedings of the 16th ACM Conference on*
882 *Computer and Communications Security*, CCS '09, page 161172, New York, NY, USA, 2009.
883 Association for Computing Machinery.

884 52 Steven Murdoch. Hot or not: Revealing hidden services by their clock skew. pages 27–36, 01
885 2006.

886 53 Milad Nasr, Alireza Bahramali, and Amir Houmansadr. Deepcorr: Strong flow correlation
887 attacks on tor using deep learning. In *Proceedings of the 2018 ACM SIGSAC Conference*
888 *on Computer and Communications Security*, CCS '18, page 19621976, New York, NY, USA,
889 2018. Association for Computing Machinery.

890 54 Alexander Nilsson, Pegah Nikbakht Bideh, and Joakim Brorsson. A survey of published
891 attacks on intel sgx, 2020.

892 55 Rishab Nithyanand, Oleksii Starov, Phillipa Gill, Adva Zair, and Michael Schapira. Measuring
893 and mitigating as-level adversaries against tor. *ArXiv*, abs/1505.05173, 2016.

894 56 Shen Noether. Ring signature confidential transactions for monero. *IACR Cryptology ePrint*
895 *Archive*, 2015:1098, 2015.

896 57 Andriy Panchenko, Asya Mitseva, Martin Henze, Fabian Lanze, Klaus Wehrle, and Thomas
897 Engel. Analysis of fingerprinting techniques for tor hidden services. In *Proceedings of the*
898 *2017 on Workshop on Privacy in the Electronic Society*, WPES '17, page 165175, New York,
899 NY, USA, 2017. Association for Computing Machinery.

17:24 Dont get caught, keep your Onions in a Vault

900 58 phobos. Using Tor hidden services for good, <https://blog.torproject.org/using-tor-hidden-services-good/>. 2012.

902 59 Florentin Rochet, Ryan Wails, Aaron Johnson, Prateek Mittal, and Olivier Pereira. *CLAPS: Client-Location-Aware Path Selection in Tor*, page 1734. Association for Computing Machinery, New York, NY, USA, 2020.

905 60 Komang Oka Saputra, Wei-Chung Teng, and Yi-Hao Chu. A clock skew replication attack detection approach utilizing the resolution of system time. In *2015 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT)*, volume 3, pages 211–214, 2015.

909 61 Shweta Shinde, Zheng Leong Chua, Viswesh Narayanan, and Prateek Saxena. Preventing page faults from telling your secrets. In *Proceedings of the 11th ACM on Asia Conference on Computer and Communications Security*, ASIA CCS '16, page 317328, New York, NY, USA, 2016. Association for Computing Machinery.

913 62 Yixin Sun, Anne Edmundson, Nick Feamster, Mung Chiang, and Prateek Mittal. Counter-raptor: Safeguarding tor against active routing attacks. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 977–992, 2017.

916 63 Yuanyuan Sun, Sheng Wang, Huorong Li, and Feifei Li. Building enclave-native storage engines for practical encrypted databases. *Proc. VLDB Endow.*, 14(6):10191032, feb 2021.

918 64 Jo Van Bulck, Frank Piessens, and Raoul Strackx. Sgx-step: A practical attack framework for precise enclave execution control. In *Proceedings of the 2nd Workshop on System Software for Trusted Execution*, SysTEX'17, New York, NY, USA, 2017. Association for Computing Machinery.

922 65 Tao Wang, Kevin Bauer, Clara Forero, and Ian Goldberg. Congestion-aware path selection for tor. In Angelos D. Keromytis, editor, *Financial Cryptography and Data Security*, pages 98–113, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.

925 66 Jan Wichelmann, Anna Pätschke, Luca Wilke, and Thomas Eisenbarth. Cipherfix: Mitigating ciphertext side-channel attacks in software, 2023.

927 67 Gavin Wood et al. Ethereum: A secure decentralised generalised transaction ledger. *Ethereum project yellow paper*, 151(2014):1–32, 2014.

929 68 Yuanzhong Xu, Weidong Cui, and Marcus Peinado. Controlled-channel attacks: Deterministic side channels for untrusted operating systems. In *2015 IEEE Symposium on Security and Privacy*, pages 640–656, 2015.