Experiences and Recommendations from Operating a Tor Exit Node at a University

Michael Sonntag[a] and René Mayrhofer[b]
Institute of Networks and Security, Johannes Kepler University Linz, Altenbergerstr. 69, 4040 Linz, Austria

Keywords: Anonymisation, Tor Exit Node, Legality, Public Perception, Recommendations.

Abstract: We report on a multi-year operation of a Tor exit node at a public university and provide recommendations for running other instances. These include legal issues, such as permissions perhaps required in advance, and where potential pitfalls are, like blocking content/DNS resolution or monitoring/logging requirements. We also discuss organizational aspects including preparations for inquiries and problem reports, how to avoid issues with potential legal enforcement, or who should have access to which systems. Technical issues are discussed in detail, including lessons learnt from DoS attacks both on the university as well as the exit node in particular. Finally, we provide technical and organizational recommendations on longitudinal data collection and other research on exit node traffic without compromising anonymity.

1 INTRODUCTION

Our university institute operated a high-bandwidth Tor (Dingledine et al., 2004) exit node for several years and performed research – particularly statistical traffic analysis – on it. Tor is an anonymisation systems based on routing encrypted traffic across multiple nodes, each of which strips away an encryption layer. Observers in the Internet (or server operators) therefore see (only) the IP address of the last (=exit) node as the system requesting some data, but not the actual originator. Misuse can be tracked back to the exit node, but not any further

For reasons detailed below, this has now stopped. At the beginning, such a project was considered highly controversial and a big risk by many relevant stakeholders, but our experiences showed that these were mostly unfounded, although we were not able to convince everyone of the relatively low risk of operating a Tor exit node. Issues that led to the highest amount of discussions were spread throughout different concerns: legal (*Is this allowed at all? Is it "contributing to illegal acts" and therefore exposing the university to potential liability?*), organizational (*How to ensure that in case legal problems arise there is no impact on the rest of the university? How not to compromising anonymity and still perform research?*), technical (*Might we become an attack target? How does the normal operation – or any attacks – impact the university network?*), and public perception (*What media opinion pieces might this generate? Does the general public perceive anonymity as positive or negative?*).

In this paper we report on the project and provide recommendations and lessons learnt from it, so that others can more easily start and operate a Tor exit node.

2 LEGAL ISSUES

Operating on solid legal ground is an obvious precondition; therefore, we proactively verified the current legal situation on multiple fronts: First, we checked whether it is legal at all to provide such a service in Austria, and if yes, whether some kind of registration/permission/etc is needed. We could confirm that a) it is legal to operate an anonymizing network packet forwarding service, and b) that no explicit permission or registration is a precondition for such operation (Sonntag, 2015). This may be different in other jurisdictions, but is probably consistent within

[a] https://orcid.org/0000-0002-2506-2350
[b] https://orcid.org/0000-0003-1566-4646

Sonntag, M. and Mayrhofer, R.
Experiences and Recommendations from Operating a Tor Exit Node at a University.
DOI: 10.5220/0010156802830290
In *Proceedings of the 7th International Conference on Information Systems Security and Privacy (ICISSP 2021)*, pages 283-290
ISBN: 978-989-758-491-6

the EU. Other typical legal "dangers" are, and should therefore be answered (in our case negatively):

* Operating a Telecommunication Service: Such a service transmits third party data from an external source to an external destination, i.e. it is no endpoint but a relay – similar to an upstream ISP. This "carrier" service could be regulated and require registration or licensing or be subject to oversight etc. The important part here is, that there is no contractual binding to any of the communicating parties (additionally, the end-user is explicitly not known to the service).
Recommendation: No explicit action needed in Austria, but registration or an operator license may be required in other jurisdictions.

* Monitoring Obligations: The service might be required to investigate all or certain traffic for some undesirable content. This might be independent of whether this is possible at all – although the outgoing traffic of an exit node is "cleartext", it can still be HTTPS and therefore encrypted. As we are not the operator e.g. of a cloud service, respective keys to decrypt and investigate are not in our sphere of influence. Note that aside of a MitM attack no "retention" and therefore disclosure of cryptographic keys is possible – this is end-to-end encryption. The type (e.g. chat, web browsing, or file download) of content might still be recoverable, e.g. from time and traffic amount analysis (Lashkari et al., 2017).

A typical example would be blocking pornography (e.g. the UK Digital Economy Act, 2017). This is potentially problematic, as e.g. the UK law applies to ISPs and "ancillary service providers" (see section 21 5 b) giving access to such content to persons in the United Kingdom. Depending on jurisdiction, this could therefore also apply to Tor nodes in other countries – although these do not know in any way in which country the end-user who uses the exit node is located. Also, an exit node probably doesn't profit from the mere-conduit-exception of liability in Art 12 of the E-Commerce directive (still optimistic Minarik and Osula, 2015). The CJEU judgment from 15.9.2016, C-484/14 – Mc Fadden, made clear that it only applies to services provided for remuneration, even though these could be "unrelated" (e.g. a restaurant providing free Internet access to its paying customers). As exit nodes do not have *any* business regarding the unknown end-users – especially if operated by a publicly funded university – this does not apply (note that in Austria this exemption was explicitly extended to free services: ECG § 19 Abs 2).
Recommendation: For exit node operation no monitoring should be required, as resulting logs would be (mostly) meaningless to usual monitoring cases.

* Content Blocking Requirements: Blocking of certain content might be required, both from fixed lists (updated e.g. by courts) as well as ad-hoc (court judgments, information from rights holders etc) in various ways (Lodder and Polter, 2017). This can also include the requirement to start ongoing monitoring for identical or similar content. Note that while a large amount of exit traffic is encrypted, according to our experiences this does not cover the entirety (25-35% remains unencrypted web traffic at the time of this writing). Additionally, the appropriate infrastructure might be required anyway (for the possibility that such a problem may arise in the remaining clear-text traffic). This could perhaps be reduced/avoided by limiting traffic to encrypted services only – so that any kind of inspection and therefore blocking is impossible anyway. However, that would significantly reduce the utility of an exit node at this time, and we therefore decided to explicitly allow cleartext traffic to be forwarded. Note that blocking and monitoring might be combined: blocking "inappropriate" traffic and simultaneously logging all such requests.
Recommendation: We believe that within the EU no a priori content blocking needs to be implemented. In specific cases as ordered by a court of law, blocking selected targets may become necessary. Avoiding this may be easiest implemented using a narrow exit policy disallowing any cleartext traffic (with the obvious disadvantages for the network).

* DNS Blocking: Blocking might not take place on the content level, but also on the DNS resolution level, which is the typical case today for e.g. copyright infringements (Geiger and Izyumenko, 2019). As exit nodes also perform DNS requests for end-users, "lying" to block specific (web)sites might be necessary. This, like the previous category, could perhaps be passed to the upstream provider: if it is in the same country, then all exit node traffic has to go through their network too. But such obligations might only apply to them for servicing "end-users" – which an exit node could be argued not to be (i.e. an ISP might have to filter traffic of private homes/businesses, but perhaps not traffic from other ISPs, especially international ones). Otherwise, implementing DNS blocking is relatively easy on the technical level if the exit runs its own DNS resolver as recommended (Tor-DNS 2019), because DNS resolution can otherwise be a significant privacy problem (Greschbach et al., 2017).
Recommendation: We believe that no a-priori DNS blocking infrastructure needs to be implemented by a Tor exit node operator, and that relevant legal requests may typically be better addressed by the respective upstream provider if necessary.

Another aspect is preparation for "legal" inquiries, i.e. if a national law enforcement agency contacts the service operator because its IP address was found on a child pornography server, someone distributed/downloaded/shared a movie, or some offensive message was sent (mail, forum post, comment…). These might originate from public officials (police, courts, administration), but also from private persons.
Recommendation: Prepare an easy-to-understand (for non-experts) explanation of the Tor system and why you are technically unable to answer any requests for end-user data, and what their other options are (if any). This should be readily available, e.g. on a website. In our experience this was not a problem, as the few official persons contacting us knew perfectly well what Tor is, and lost all interest as soon as they were informed about this fact. This was similarly true for the private persons that contacted us, which either knew what Tor is or were sufficiently informed to not contact us again by our (automated) response.

2.1 Improving Legal "State of the Art"

Whether something is "state of the art" is an important concept from the legal point of view: "State of science" (also called "best available techniques") aspects can easily be disregarded and do not have to be implemented in most areas of business (see BVerfG, 1978 for the source of this differentiation). Companies might not have to actually implement all of the "state of the art" either, but it is mandatory to "consider" it. This means, if you choose not to follow recommendations from this level, you need an explanation as to why not (which may typically include "it is too expensive" or "it would disrupt core business"). "State of the art" still requires "practically tested" and "sufficiently proven", i.e. it must have been used in a wider area or to a larger extent (not only tested in a laboratory or a single special instance) and has worked sufficiently well for a longer span of time (Weidenhammer and Gundlach, 2018). At least in Austria, Tor was (because of the world's first court proceedings on operating an exit node, where the operator was found guilty) seen as neither: (practically) nobody operated an exit node (at least not openly), and its operation (effort required, potential problems etc; but not that it is technically possible and works as intended, which was accepted) was not proven either.

In this sense we provided a significant contribution to the "state of the art" in Austria, and perhaps even to a wider (e.g. EU) jurisdiction: operating a Tor exit node was proven to be practically possible for a long time without significant drawbacks or problems, neither technical, nor legal, nor organizational. Through our successful operation and the publicity it is now proven that it works not only on the computer science level, but on a practical level too, transitioning the operation of a Tor exit node from the concept of "state of science" to "state of the art". Consequently, if discussing techniques to improve privacy, considering the Tor system will now be a mandatory element that can no longer be ignored. It might not be used, but then arguments are needed why. These will usually have to be technical (latency etc), as our explicit lack of bad experience in terms of legal/organization aspects renders arguing on those grounds difficult. Financial reasons are limited too, as the necessary hardware is very cheap (at least for any kind of business venture), and personnel costs are low because little work is needed for installation as well as operation.
Recommendation for Operators: No specific action is needed, because our successful multi-year project with very little actual risk or damage has set a precedent that operating a Tor exit node in Austria (and therefore the EU) is now "state of the art".
Recommendation for Other Services with Anonymity Needs: Because the Tor system is now "state of the art", it needs to be considered as one potential solution for anonymous communication and service access, e.g. by providing an Onion service (Goulet et al., 2013) or relaying client/server connections through the Tor network.

3 ORGANIZATIONAL ISSUES

From the organizational point of view, several precautions should be taken for an exit node that are not necessary for an entrance or middle node. Some of them might be easier at a university, where personnel of the same organizational level are more common than in a company with a strict hierarchy, but these should still be considered. These recommendations serve to increase security both in a technical sense as well as from a legal point of view.

First, the group of persons with access to the hardware should be kept as small as possible: hardware manipulation can remain undetected electronically, and most servers have ample space inside to introduce additional small devices. This requires physical separation of the exit node from other systems. This applies to the whole "system", i.e. a firewall dedicated to it, a switch/router, DNS server, the exit node itself, and any statistical monitoring or other components recording research data. It does not apply to the "general" parts, i.e. everything where merely data is transported but not acted upon – which are identical to the

untrusted "public Internet" anyway. Therefore, the exit node hardware up to and including its boundary gateway to public systems needs to be secured. In our case this was a physically separate room, but variants are of course possible too, like wire cages or locked racks. This has the additional advantage that in case of a legal order to confiscate hardware equipment for analysis the damage can be isolated to only the Tor exit node components by clearly marking that the relevant IP addresses and data are all handled inside this single room with no external dependencies. We did not have to face such a situation during our multi-year operation, but, based on prior legal action against a Tor exit node in Austria, were prepared for it. We also believe that this isolation (=mitigation against risk to other systems) was a positive factor in gaining initial permission to run the system within the university.
Recommendation: If possible, we strongly recommend to physically isolate the whole system, either in a separate room or at least a separate, locked rack.

Should physical and electronic access be separated? Anyone with el. access can get at the data, but this would also produce digital traces. However, separating access gives an additional person an opportunity for malicious behaviour (or, of course) simply mistakes): voluntarily, or through bribes or extortion.
Recommendation: Merge all kinds of access (physical and digital) in as few persons as possible.

For both electronic and physical access, the four-eyes-principle should be followed. For physical access this results in double locks or in access to the room separated from access to the rack/cage. Digitally this is much harder, as typical access is via SSH, which does not readily support this principle. A simple "circumvention" is to employ multi-factor authentication and give each person only a certain factor, e.g. one person knows the password, the other controls the token/authenticator app etc. This works for a simple "two-person" requirement, but not for more complex "at least two out of N persons". Unless special custom software is used, we therefore recommend this "simple" two-factor authentication to ensure that no single person can modify the configuration, introduce additional software, or extract data. While this might be possible for the exit node itself, this can be more complicated (or even impossible) on other devices. Firewalls might support such functionality, but for switches/routers it is uncommon. If the setup is possible without such devices, this should be preferred. This might not necessarily be easy, as a webserver (information about Tor/the exit node), a DNS server (DNS lookups; support for blocking) and a firewall (protecting these systems) need to be connected. However, modifying the switch/router should only rarely be necessary, e.g. in case of a firmware update. Therefore, an option is to forbid any electronic access and only allow physical access, e.g. through a serial interface, porting the problem to the physical four-eyes principle. Note that we did not strictly enforce the four-eyes principle on the digital level for all components, but that different administrative roles were assigned to different people.

Another issue is top-level permission to operate an exit node. We did perform research on the node and published the results in several venues, but simultaneously it was also a service for the public (unavoidably and intentionally, as we could not produce realistic traffic ourselves). We therefore obtained prior, explicit permission from the university, which turned out to be the biggest problem of all: technical requirements could be solved easily, legal issues required some work and research, but posed no real hindrance either. This university permission was always only granted for a single year (which is problematic, as we had to design an experiment, prepare it and implement it/collect traffic within this timeframe). The last time, explicit permission was not granted anymore until the time of this writing; we applied for it, but, repeatedly, no decision was made. No specific other data, plans etc were requested (already included in application), so we couldn't "remove" the barrier for a permission. As we found out later, the university's biggest fear was bad publicity (see below).

A contributing factor could have been that we did not manage to obtain third-party funding for running the exit node – we did not require any expensive hardware (a 5-year old server that was phased out by the central IT department was deemed more than sufficient, reducing add. hardware cost to close to zero) or software. The university also did not have to pay for the traffic (which was donated by the network provider, so in a sense we *did* obtain external funding). Direct costs involved were trivial and borne by the institute (no university funding), but no external grant money was flowing into the university either. The publications and publicity originating from the operation were sufficient and useful for researchers and department – and hopefully the scientific community.
Recommendation: Engage with all stakeholders (IT services, upstream provider (in our case ACOnet, the Austrian academic network), legal and public media outreach departments) and apply for explicit permission to run an exit node before commencing its setup and operation. If permission is only granted for a limited duration, apply early for extensions.

Another recommended precaution is preparing and testing to block all traffic to/from the exit node from the outside. While there was a DoS attack

against the university (different from the one described below; this was targeted at the whole university, not specifically the Tor node), we were cut off first: "unnecessary" traffic is shed first if bandwidth gets limited. Technically, this can be done in various ways, e.g. shutting down the exit node or unplugging it. As the persons capable/allowed to do this might not be available (or as quickly as necessary), interrupting traffic on routers/firewalls should be prepared. This is technically easy, but should be tested and guidelines for when this can/should be done instituted.
Recommendation: Prepare a "crisis situation" plan to quickly deactivate network access to the whole system while the situation is analysed and other steps can be prepared. Documenting how to turn off central power or unplug the boundary gateway from the main network for admin staff is a perfectly valid approach.

3.1 Contact with Public Authorities

Contact with public authorities was very rare. As discussed above, they seemed to know exactly what Tor is, and what data we therefore could potentially provide to them (i.e. nothing useful). This also led to no additional contacts like physical visits, searches, confiscated hardware, etc. However, this might partially be caused by being a public university (which does have a legal department – and a legal faculty, can reach out to the public, has experience with publications, and has connections to politics). Experiences of private person as operators might vary.
Recommendation: A large, publicly funded, and generally considered trustworthy institution is a perfect place to run a Tor exit node. Smaller organizations may potentially face more pressure.

In terms of an added organizational precaution we initially implemented, that we don't know how effective it was: a separate domain name/IP address range, with matching WhoIs entries. These included that this was a "Tor research" project. If our IP address was found, querying the WhoIs database might be a very early and quick investigative step. Whether this was enough to discourage actual inquiries (or they were still made in spite of a "suspicion" of this being a Tor exit node), is unknown. Still, we would recommend doing this, as it could reduce the number of contacts and doesn't impose direct cost if separate IP ranges are available. The only drawback is, that the node could more easily/quickly end up on a list of "exit nodes", which are blocked/delayed/require additional confirmation etc. But as exit nodes are public anyway and their IP addresses can trivially be downloaded (while the WhoIs is more restricted and has no fixed format for noting such data!), this seems irrelevant.
Recommendation: If organizationally available, assign a distinct IP subnet and reverse lookup domain to the Tor exit node. Set WhoIs information appropriately (noting in text fields that this is a Tor exit node) and define an "abuse" email contact separate from the main IP address range used for the rest of the network.
Recommendation: Use a ticketing system to systematically handle incoming abuse/inquiry email requests with an auto-reply explaining that this is a Tor exit node (e.g. with the template explanation of what it does) and that there is no personal data available. The auto-reply could mention that another request (by replying to this auto-reply) would trigger a manual response by the node operator team. Although the total number of "manual" inquiries was extremely low - so this might not be significant - we did not every receive a repeated inquiry from the same institution.

3.2 Continuous Effort Required

Apart from research, continuous effort required to run the exit node was moderate: the node itself needed little attention (e.g. SW updates), but answering every single abuse report required significant time. This we promised to do to the university admin., to ensure that we didn't run into any problems (and the number/type of inquiries we received was in scope of our research topic). But practically all reports were created fully automatically and we did not ever receive any meaningful replies. This could be because nobody read our standard replies, or nobody cared because they were about a Tor exit node, and so they could not pursue their inquiries anyways. Additionally, nearly all of them were reports about password tries to SSH servers. Such attempts are extremely common even on the normal network - and not really a cause for serious concern. After we removed port 22 from our exit policy, such reports, and the associated work, disappeared for all practical considerations.
Recommendation: Either not allow port 22, or at least filter out (and ignore) all reports about this port. Note that SSH traffic is only a tiny part of the traffic regarding bandwidth, but produces practically all complaints. Additionally, contacting a service where you have to explicitly identify yourself securely is only rarely useful anonymously (and so the utility of a Tor exit node is not limited significantly by disallowing SSH traffic): hiding the source IP address or the fact of using that server/service. Most of these reports originated from *fail2ban*, which was often misconfigured too (e.g. no IP address/hostname of the server included, or no E-Mail address to reply to).

While no effort was expended by the team at the university, there was an incident where a DoS attack

was performed against the exit node. This was detected and averted at the boundary of the Austrian academic network (ACOnet). As it was a naïve DoS attack, the effort was probably also small at that institution – however this is not something that has to be the case every time (see Jansen et al., 2019 for more complex attacks on both individual nodes as well as the whole Tor network). A more sophisticated (or larger/longer) attack would cause problems or potentially require more effort to reduce it, or even commercial services to help against it. And while this network of all Austrian universities (and other educational institutions) could come under attack for many reasons, a Tor exit node could be an additional cause.

The DoS attack was solely directed at the exit node (no other systems or nearby IP addresses were attacked) and had a maximum of 15,1 Gbit/s (1,8 Mpps), so was quite large even for a "normal" commercial network, but not internationally significant. It lasted for 31 minutes and only consisted of UDP packets targeted at various unprivileged ports. It was uniquely simple in the sense that approx. 75% of all traffic originated from only four IP addresses (and there mostly from two), while the rest was highly distributed. It also seems that every source IP address used only a single source port. Moreover, 91% of all traffic originated in a single AS in Poland. Almost all traffic was simply UDP packets of 1050-1200 bytes length; only tiny parts were IP fragmentation or CLDAP amplification attacks. These properties made it comparatively easy to protect against at the network ingress. There was no communication regarding the attack, i.e. no demand for anything, no "warning" or anything else: it simply started and soon after ended.

4 TECHNICAL ISSUES

Because of the potential for problems, both legal and technical, as described above, we recommend a separation from other aspects of a company's organization (such as the university) as far as possible. This includes a separate domain name (easy and cheap) and a separate IP range (difficult and potentially expensive with IPv4, but should not be a problem with IPv6). Physical separation, or at least very clear marking, is also recommended to avoid potential collateral damage. This might be problematic with "shared" equipment, like a switch or a firewall. We therefore recommend to employ either completely general systems (e.g. a firewall for everything, with no special handling for the exit node), or dedicated systems. This might be less useful for a switch/router, as there no meaningful content can be expected, but this depends on the technical expertise of the personnel potentially tasked with investigating/impounding hardware. But operating an exit node as a virtual machine on a hardware together with other business-critical VMs is carries high risk – the police might not know/care and inspect/impound the whole physical server. To avoid having to set aside a physically separate firewall, OS built-in firewalls can be used and configured, and e.g. a direct cable for the communication with a separate DNS resolver. The easiest version to implement is however, if the exit node is physically the same system as all other elements, either because everything is a separate service on one OS, or virtualization is used (i.e. the exit node and DNS/web/… servers are VMs on a dedicated hardware server). We do not recommend the latter, as it provides a chance to obtain data from inside the exit node, without the node itself being able to protect against. This might also take place inadvertently, e.g. by virus scanners integrated into the hypervisor and scanning the content of the virtual machines (and potentially passing it on to cloud services for analysis).

A basic precaution is to ensure that no personal or content data is stored at all and not longer than necessary, i.e. log files should not be created or stored. This applies also to security devices like firewalls and can be problematic as not all products support turning off all kinds of logging completely.
Recommendation: Explicitly disable all logging and monitoring not required for the operation of the Tor exit node or obtaining the research data in scope of the current experiment. This will lead to a non-standard configuration on most systems, but mitigates other legal and organizational risks and effort.

Additionally, while devices dedicated solely to the exit node can be adapted comparatively easily – they are after all close by and usually under the same control – the same might not apply to organization-wide elements. For instance, all traffic of the university has to pass through an IDS system, and turning it off for the exit node was not possible. This was seen as acceptable, as it only blocks attacks according to a list of known attacks or known bad systems (no heuristics). Moreover, it does not log any "normal" traffic, only a few statistics on detected (and blocked) attacks. If such systems exist, it must be closely verified whether they pose a danger to privacy regarding the exit node, or whether they can be turned off for this part of the network – which is typically a policy issue.
Recommendation: If central systems can be influenced as part of operating an exit node, either apply policies of no monitoring or no specific policies for the node itself, treating it like all other client systems.

The most important technical question for running a Tor exit node is the bandwidth required, which is a problem for the whole Tor network (Panchenko et al., 2010). Unlike normal end-user systems (typically much more download than upload, or in case of uploading backups to remote storage exactly the reverse) a Tor exit node always produces symmetric traffic: the amount of data going in is the same as the outgoing traffic, as it merely acts as a relay (HTTP requests might be small, but they enter encrypted from "inside" the Tor system and exit unencrypted to the Internet, and while web content is large, it enters in clear from "outside" and exits encrypted inside). Additionally, as a Tor node should serve a large number of users to improve anonymity, its bandwidth should be large. All elements combined, symmetric traffic of many users with high bandwidth results in a permanent traffic without interruptions. The amount of traffic can therefore quickly become significant to the whole organization, even for institutions with many employees. According to information we received informally, the network traffic of our university to/from outside of the Austrian university network (inside, i.e. to other universities, very large data transfers of various research data seem to take place) doubled because of our 200 Mbit/s exit node, i.e. producing the same "external" traffic as approx. 15,000 students and 6,000 employees. If ACOnet had not classified the traffic caused by our Tor exit node as research project traffic, the university would not have granted permission to run it because of the significant additional traffic cost. We therefore (again) thank ACOnet for it support of our multi-year project.

Recommendation: Prior to activating a Tor exit node, estimate the total traffic bandwidth and volume and ensure that traffic cost can be covered long-term. This is the single most important cost factor for the operation of a Tor exit node.

5 PUBLIC PERCEPTION ISSUES

Public perception of anonymization services is highly ambivalent. On the one hand, privacy concerns have become far more significant in the last 2-3 years due to well-publicized abuses of surveillance and data collection (Szoldra, 2016; Gibbs, 2015; Collins, 2020; Puig, 2020). On the other hand, some public media articles have linked anonymization services in general and Tor in particular to scary-sounding topics like the "Darknet", "terrorism", "child pornography", "drug markets", "illegal weapons", etc (Power, 2020; Fariva and Blankstein, 2019). Common (mis-)conceptions of "nothing to hide, nothing to fear", law enforcement "going blind", or that the Internet should not become a "lawless place" make decent headline material but do not generally promote an informed, nuanced, and helpful public debate. Therefore, public perception of an organization running a public anonymization service is, for all practical matters, unpredictable and can change quickly – e.g. any local newspaper could write that the university is actively helping criminals and transporting child pornography.

Permission to continue to operate the Tor exit node by the university would probably be granted without a time limit if we could guarantee that this kind of publicity would never happen. Unfortunately, this is impossible to guarantee or merely estimate. We provided a list of reasons, why such an exit node is important not only for research, but also society, so that a "response" for a bad-press incident was already prepared – but this was not considered to be sufficient to mitigate the publicity risks to the university. After more than 7 months and multiple tries with additional arguments in favour of continuing the project, a decision on the permission to operate the exit node was again postponed. Even if the request was not declined or the operation prohibited, no decision is a decision too, and we therefore shut the project down (the servers were already deactivated earlier when the last permission ran out). It should be noted that no such bad publicity did occur during operating the exit node – the mere fear of such an incident potentially occurring at some point in time was enough.

We explicitly note that we do not blame the university administration for the hesitancy in accepting the unpredictable risk of public perception, but document this reason for discontinuing our project to help other operators prepare for this issue. In a period of post-factual (social) media reporting, every organization must decide on its own prioritization of how (pot.) public perception influences its goals, including research and open debate of controversial topics.

6 SUMMARY

Although we cannot operate the Tor exit node anymore, there are now several other exit nodes in Austria, which were started based on our positive experiences. Based on the communication with them and their informal feedback we deem it very unlikely that these would have been created without our precedence, and testing the legal and technical operation. These can serve as an example of moving a technology from *state of science* to *state of the art*, and towards "normal business", i.e. something that can be

operated without difficulties. This broadens the base of relays and additionally moves operating an exit node closer to end-users, which would improve overall performance and anonymity (Ngan et al., 2010). If all end users simultaneously act as exit-nodes with their excess bandwidth, that would provide plausible deniability for their own traffic too.

It is also important to note that we only offered an exit node, i.e. a relay to the public Internet. There was never any question of hosting hidden services for third parties, as this would have exposed the university to much larger legal dangers as a hosting provider. Similarly, although we of course participate in the operation of hidden services as any other Tor node, we did not provide any directory or list of onion URLs or similar: while not hosting the content, providing directions to them would still have required investigating/checking each of them for legality.

We can therefore conclude that operating a Tor exit node is easily possible from the technical point of view, but the biggest problem seems to remain on the organizational level. Nonetheless, our multi-year project successfully established precedent to now make operation of Tor exit nodes "state of the art" within Austria and therefore within the EU. Legal ramifications in terms of having to consider Tor exit nodes as a viable network anonymization techniques remain a topic for future research.

REFERENCES

BVerfG 8.8.1978, 2 BvL 8/77 - Kalkar I, https://dejure.org/dienste/vernetzung/rechtsprechung?Gericht=BVerfG&Datum=08.08.1978

Collins, K., 2020. US-EU Privacy Shield data-sharing pact invalidated https://www.cnet.com/news/us-eu-privacy-shield-data-sharing-pact-invalidated-over-surveillance-fears/

Dingledine, R., Mathewson, N., Syverson, P., 2004. Tor: The second-generation onion router. In: *Proc. of the 13th USENIX Security Symposium* (August 2004)

Puig, A., 2020. Equifax Data Breach Settlement, https://www.consumer.ftc.gov/blog/2019/07/equifax-data-breach-settlement-what-you-should-know

Fariva, C., Blankstein, A., 2019. Feds take down world's 'largest dark web child porn marketplace' https://www.nbcnews.com/news/crime-courts/feds-take-down-world-s-largest-dark-web-child-porn-n1066511

Geiger, C., Izyumenko, E., 2020. Blocking Orders: Assessing Tensions with Human Rights, In: Frosio, G. (Ed.), *The Oxford Handbook of Intermediary Liability Online* (OUP, 2020), 566

Gibbs, S., 2015. What is 'safe harbour' and why did the EUCJ just declare it invalid? https://www.theguardian.com/technology/2015/oct/06/safe-harbour-european-court-declare-invalid-data-protection

Goulet, D., Kadianakis, G., Mathewson, N, 2015. Next-Generation Hidden Services in Tor, https://gitweb.torproject.org/torspec.git/tree/proposals/224-rend-spec-ng.txt

Greschbach, B., Pulls, T., Roberts, L. M., Winter, P., Feamster, N., 2017. The Effect of DNS on Tor's Anonymity. *NDSS '17*, Internet Society

Jansen, R., Vaidya, T., Sherr, M., 2019, Point Break: A Study of Bandwidth DoS Attacks against Tor, *Proc. of the 28th USENIX Security Symp.* 2019

Lashkari, A., Gil, G., Mamun, M., Ghorbani, A., 2017. Characterization of Tor Traffic using Time based Features. *3rd Int. Conference on Information Systems Security and Privacy*, 253-262

Lodder, A. R., Polter, P., 2017. ISP blocking and filtering: on the shallow justification in case law regarding effectiveness of measures, *European Journal of Law and Technology*, Vol 8, No 2

Minárik, T., Osula, A.-M., 2016. Tor does not stink: Use and abuse of the Tor anonymity network from the perspective of law, *Computer Law & Security Review*, Volume 32, Issue 1, 111-127

Ngan, T.-W., Dingledine, R., Wallach, D.S., 2010. Building Incentives into Tor, In: Sion R. (Ed.) *Financial Cryptography and Data Security. FC 2010*. LNCS 6052, Springer, Berlin

Panchenko, A., Lanze, F., Engel, T., 2012. Improving performance and anonymity in the Tor network, *2012 IEEE 31st International Performance Computing and Communications Conference (IPCCC)*, 1-10

Power, M., 2020. Online Drug Markets Are Entering a 'Golden Age', https://www.vice.com/en_us/article/dyz3v7/online-drug-markets-are-entering-a-golden-age

Sonntag, M., 2015. Rechtsfragen im Zusammenhang mit dem Betrieb eines Anonymisierungsdienstes, *JusIT* 6, 2015, 215-222 (ISSN 1996-8228)

Szoldra, P., 2016. This is everything Edward Snowden revealed in one year of unprecedented top-secret leaks, https://www.businessinsider.com/snowden-leaks-timeline-2016-9

TorRelayGuide, https://trac.torproject.org/projects/tor/wiki/TorRelayGuide#DNSonExitRelays

Weidenhammer, D., Gundlach, R., 2018. Wer kennt den „Stand der Technik"? *DuD* 2/2018, 106-110.