

Independent Submission
Internet Draft
Intended status: Informational
Expires: April 22, 2023

M. McFadden
internet policy advisors
Jim Reid
RTFM, llp
October 23, 2022

Internet Threat Model - A Reconsideration
draft-mcfadden-threat-model-00.txt

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>

This Internet-Draft will expire on April 22, 2023.

Copyright Notice

Copyright (c) 2022 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Simplified BSD License.

[RFC3552](#)/BCP72 describes an Internet Threat model that has been used in Internet protocol design. More than twenty years have passed since [RFC3552](#) was written and the structure and topology of the Internet have changed dramatically. With those changes comes a question: has the Internet Threat Model changed? Or, is the model described in [RFC3552](#) still mostly accurate? This draft attempts to describe a non-exhaustive list of the most likely updates and changes in the current threat environment. This paper has the goal of suggesting a way forward for describing the contemporary threat model and how it might inform security aspects of protocol design.

Table of Contents

1.Introduction.....	2
1.1.Scope.....	2
2.The Established Model in BCP72	3
2.1.BCP72 Passive Attacks.....	3
2.2.BCP72 Active Attacks.....	3
3.Changes to the Attack Landscape.....	4
3.1.Quantifiable Changes.....	4
3.2.Qualitative Changes.....	5
3.3.Data at Rest and Intermediaries.....	6
3.4.The Evolution of Endpoints and Applications.....	7
4.Path Forward.....	7
5.Security Considerations.....	8
6.Privacy Considerations.....	8
7.IANA Considerations.....	8
8.References.....	9
8.1.Informative References.....	9
9.Acknowledgments.....	9

1.Introduction

[RFC3552] describes an Internet threat model. According to that RFC, the threat model "describes the capabilities that an attacker is assumed to be able to deploy against a resource. It should contain such information as the resources available to an attacker in terms of information, computing capability, and control of a system."

Has the Internet Threat Model really changed; or, is the model described in [RFC3552](#) still mostly accurate?

1.1.Scope

The purpose of this draft is to examine the threat landscape of the contemporary Internet and answer those questions. If the answer is

clearly that the threat model has changed, then the draft will suggest a way forward toward documenting those changes.

Reconsideration of the guidelines for writing Security Considerations sections of RFCs is not in scope for this memo.

2.The Established Model in [BCP72](#)

[BCP72](#)'s threat model divides attacks based on the capabilities required to mount the attack. In particular, it divides attacks into two groups: passive attacks where an attacker has only limited, or read-only, access to the network; and active attacks where the attacker has the resources available to write to the network. [BCP72](#) is careful not to locate the attack. The attacks can come from arbitrary endpoints. Dividing the threat model in this way also allows for the model to incorporate attacks that come from resources not at endpoints. In fact, an entire subsection of the BCP discusses on-path versus off-path attacks.

2.1.BCP72 Passive Attacks

[BCP72](#) describes passive attacks as those in which an attacker "reads packets off the network but does not write them." It then gives some specific examples including password sniffing, attacks on routing infrastructure, and unprotected wireless channels.

The description in [BCP72](#) tacitly assumes that the attacker is in control of a single resource. For example, the first type of passive attack considered is one in which an attacker uses read-only access to packets to extract otherwise private information. [BCP72](#) discusses the problems encountered when packets are transported without some form of transport or application layer security.

[BCP72](#) also describes offline cryptographic attacks in which an attacker has made offline copies of packets that have been read off the network. The attacker then mounts a cryptographic attack on those packets in order to extract confidential information from them offline.

2.2.BCP72 Active Attacks

[BCP72](#) says, "when an attack involves writing data to the network, we refer to this as an active attack." In this case, the BCP discusses spoofing packet replay attacks, message insertion, deletion and insertion, man-in-the-middle, as well as a Denial-of-Service attack.

In each of these cases, the BCP suggests either mitigations or descriptions of what technologies could have been used to avoid the weakness.

3.Changes to the Attack Landscape

3.1.Quantifiable Changes

In the period since 2003, one dramatic change is the number of attacks seen. Published studies show orders of magnitude increases in the number of devices compromised, scale of privacy breach, and the number of attacks taking place. Recent studies show that the vast majority of attacks come from attackers using automated, distributed tools. This makes a threat model that is built around the notion of a single attacker inapplicable in the current Internet. [BCP72](#) does reference the concept of distributed denial of service (DDoS), however its focus is on single attackers either on or off-path.

Studies also show that certain well-known ports [[IANA-WKP](#)] are the primary targets for this large jump in automated attacks. Ports 445, 22, 23, 53 and 1433 make up 99% of the targets.

The growth in the attacks has in part resulted from endpoints that are not capable of supporting endpoint protection software, lack effective encryption, or proper authentication. These have proliferated on the public Internet. That many of these devices do not have facilities for either self-protection or protecting against becoming a threat on their own has been documented in an IAB Workshop [[IAB-IOT](#)]. The greater number of improperly protected devices has the potential to amplify attacks that use them as sources for attacks on the rest of the Internet ecosystem.

Since 2003, there have been a variety of studies examining the growth in the number of devices connected to the Internet. At the time of writing, one estimate is that the difference between the number of devices connected in 2003 and 2021 is in the region of 22 billion. The sheer quantity of devices means that the Internet's attack surface is significantly expanded. Quantitative surveys also indicate that the greatest growth is in so-called enterprise IoT and household automation. The security properties of these endpoints are substantially different from hosts that made up the majority of the Internet in 2003.

Another important quantitative change to the structure of the Internet is the consolidation of its infrastructure. While [BCP72](#) is certainly correct in its focus on the technologies and protocols that can be exploited by attackers, it is hard to ignore the fact that the threat landscape has been affected by the emergence of consolidation. One example of this would be commercial or governmental surveillance capabilities. In an environment where there are a small number of very large entities that control the fabric of connectivity and content, the threat landscape is affected

by the fact that it may be easier to exert control and implement attacks on a small number of organizations.

3.2. Qualitative Changes

The Internet in 2003 had a relatively small number of types of hosts. The client/server model of computing was dominant at that time and endpoints were relatively homogeneous.

The diversity of deployment is an important part of the contemporary Internet landscape. Not only is there a measurable and huge increase in the number of endpoints (greatly increasing the attack surface), but there is a rich diversity in the capacity and purpose of those endpoints. As a result, while the number of protocols may not have increased exponentially, the kinds and quantities of devices that can be sources or targets of exploits has increased significantly.

The threat landscape is also affected by the balance between convenience versus protection from threats. Applications and services fight for market and mind share by being the easiest to adopt, install and use. Many users treat security and protection in the same way that they treat personal health - they ignore it until there is a serious problem and then expect the problem to be mitigated quickly. There are also market incentives which discourage the deployment and adoption of security features, for instance support costs, ease of use, hardware constraints, key management, algorithm agility and so on.

The class of attackers has changed as well. In 2003, advanced persistent attacks hadn't yet been given that name and the estimated monetary loss to attackers was estimated to be less than \$1 billion USD. The emergence of scripted and other automated tools has changed the landscape dramatically. In 2019, one estimate of losses due to network-based attacks was in excess of \$315 billion. This is the direct result of the speed, financing and flexibility of those doing the attacking.

It is true that, since [BCP 72](#) was published there have been significant improvements to communications security. This includes securing the transport layer through protocols such as TLS 1.3, HTTP/2 and secure SMTP. However, secure transport does not prevent rogue applications from executing attacks, even when secure transport is in place. An example of this happens when VPNs themselves examine or exploit traffic rather than do what they are advertised to do.

Recent experience tells us that the Internet has evolved from primarily supporting unidirectional, two-party data flows to

supporting both two-party and multi-endpoint communications. This trend is especially seen in the move toward large-scale, work from

home models where multiparty communication is taken as a fundamental use case. The implications of this evolution on the threat model should be a part of any reconsideration of [BCP72](#).

One of the other crucial changes to the Internet is the rise of the application. Apps do everything for themselves that they can so they do, for example, DoH [[RFC8484](#)], encrypt on their own and make changes to the way the application interfaces with the Internet. It used to be that applications simply relied on lower layers of the stack for their services. This is no longer always the case, and the implications of this on the threat model may be that the nature and platforms for attacks has significantly changed.

3.3.Data at Rest and Intermediaries

The Internet Threat model in [BCP72](#) primarily speaks to data being transmitted, transited or received over the network. More recent approaches to providing services over the Internet involve intermediate nodes that may redirect, manipulate or store traffic. While technologies such as exchange points may be seen to simply part of the fabric between senders and receivers, the insertion of content networks, caches and traffic analyzers has become ubiquitous.

These middleboxes play an important role in content provision, analysis and security in today's Internet. They were in limited use when [BCP72](#) was published. The importance of middleboxes is such that, when protocols are developed that effectively route around them, operators and content providers sometimes object.

One view of these intermediaries is that they are on the path between source and destination and receive and forward information for the benefit of one (or, both) of the endpoints. This is different from network resources that facilitate on connection, such as shared recursive DNS servers.

A helpful example of an intermediary has been provided by Martin Thomson. He says, "in WebRTC there are signaling servers, who intermediate the signaling stuff (control plane if you will), but do not intermediate the media (data plane). Media is intermediated in different ways by selective forwarding units (SFUs) or bridges or mixers or focuses (there are lots of names for these and lots of ways to build them). There are also relays, which are intermediaries that help with NAT and sometimes firewall traversal."

It is important to see that intermediaries, and their security properties are also a matter of perspective. Support for end-to-end, human-to-human communications is one aspect of the threat model. Today's internet also supports large-scale deployment of objects and

"things" which have different intermediaries - and different threat models.

Any contemporary Internet threat model must go beyond the threats to traffic as it moves from Alice to Bob. Beyond intermediaries, the more personal digital devices there are, the more difficult it is to control and protect them. The threat model should also include attacks that take place when the data is at rest or being manipulated for operational reasons. Observations

It seems that there are significant changes in the architecture and service model of the Internet. Those significant changes suggest the threat model documented in [RFC3552](#) may need to be reassessed.

3.4.The Evolution of Endpoints and Applications

[BCP72](#)'s concentration on the communication channel fails to account for two of the central developments of the Internet in the last ten years: the rise of the application as the endpoint and the diversity of endpoints that are publicly connected.

It might also be observed that there have already been limited attempts to reconsider [BCP72](#)'s threat model. As an example, the Same-Origin Policy detailed in [[RFC6454](#)] shows how an application-layer protocol can protect itself against certain kinds of attacks based on the concept of origin (the determination and use of an origin URI).

Another change is the emergence of state-sponsored attacks on both endpoints and infrastructure. These attacks are quite different in both capability and intensity compared to the threats seen in 2003.

Finally, protection from phishing attacks in the presence of certain implementations of IDNA means that applications themselves are implementing their own protections against certain types of attacks. This is another example of how the application layer imposes controls on an otherwise secure communication channel.

These are intended as only examples of how the landscape has changed. It seems clear that many more changes exist and need to be researched and documented.

4.Path Forward

[BCP72](#) is an accurate reflection of the security threat landscape at the time which it was written. It is also clear that protocol work since the development of [BCP72](#) has not been impacted by the fact that there has been no update to that RFC.

Still, it seems clear that the existing model does not reflect current operational reality. The IETF could choose to continue to not update [BCP72](#). This memo accepts that is a possible option.

However, this memo suggests that a fresh approach to documenting the Internet's threat model would be valuable to protocol designers, network operators and application implementers. This memo does not suggest replacing [BCP72](#). Instead, it suggests supplementing that RFC with new, more current information. It's worth emphasizing that this work is clearly in the remit of the IETF.

[BCP72](#) represents a too narrow view of the Internet's threat landscape as it exists today. It should be complimented by a new document that would:

- Reflect the diversity of endpoint deployment on the Internet;
- Document the impact of application-based security on the more narrow communication channel model (possibly: consideration of data in use in addition to data in motion);
- Account for data at rest as part of the model as well as data in motion;
- Reflect on the how the growth of the number of devices connected affects the attack surface for the Internet at large;
- Research how a new, contemporary threat model might be described and communicated to protocol designers and others; and,
- Make constructive suggestions for an approach (or, methodology) for the IETF to supplement the threat model [BCP72](#).

5.Security Considerations

This document is entirely about security on the Internet. An earlier version of this draft was intended as input into the IAB's Model-T activity which has since closed.

6.Privacy Considerations

This document does not discuss how [RFC3552](#) might be revised or replaced with an additional emphasis on privacy or trust issues. Taking on privacy and trust seems out of scope for a discussion that is focused on the Internet's treat model.

This memo is not intended to address privacy or related issues in relation to protocol design.

7.IANA Considerations

This memo contains no instructions or requests for IANA.

8. References

8.1. Informative References

- [RFC3552] Rescorla E., Korver, B., IAB, Guidelines for Writing RFC Text on Security Considerations, [BCP 72](https://tools.ietf.org/html/rfc3552), [RFC 3552](https://tools.ietf.org/html/rfc3552), <https://tools.ietf.org/html/rfc3552>
- [RFC6454] Barth, A., "The Web Origin Concept," ISSN: 2070-1721, [RFC 6454](https://tools.ietf.org/html/rfc6454), <https://tools.ietf.org/html/rfc6454>
- [RFC8484] Hoffman, P., McManus, P., "DNS Queries over HTTPS (DoH)," ISSN: 2070-1721, [RFC 8484](https://tools.ietf.org/html/rfc8484), <https://tools.ietf.org/html/rfc8484>
- [IAB-IOT] Jimenez, J., Tschofenig, H., Thaler, D., "Report from the Internet of Things (IoT) Semantic Interoperability (IOTSI) Workshop 2016," <https://tools.ietf.org/html/draft-iab-iotsi-workshop-02> (work in progress), July 2018.
- [IANA-WKP] "Service Name and Transport Protocol Port Number Registry," <https://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xhtml>

9. Acknowledgments

This document was prepared using 2-Word-v2.0.template.dot.

The authors are happy to acknowledge the comments of participants in the IAB's Model-T program. In particular, the comments of Martin Thomson, Dominique Lazanski and Jari Arkko have been helpful in building a first version of this draft.

Authors' Addresses

Mark McFadden
Internet policy advisors llc
513 Elmside Blvd
Madison WI 53704 US

Phone: +1 608 504 7776
Email: mark@internetpolicyadvisors.com

Jim Reid
RTFM llp
395 Hillington Road
Glasgow G51 4BL
Scotland

Email: jim@rfc1035.com

