Independent Submission
Internet Draft
Intended status: Informational
Expires: September 9, 2020

M. McFadden
internet policy advisors

March 9, 2020

BCP72 - A Problem Statement
draft-mcfadden-smart-threat-changes-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 September 9, 2020.

Copyright Notice

Copyright (c) 2020 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.

Abstract

RFC3552/BCP72 describes an Internet Threat model that has been used in Internet protocol design. More than sixteen years have passed since RFC3552 was written and the structure and topology of the Internet has changed. With those changes comes a question: has the Internet Threat Model really changed? Or, is the model described in RFC3552 still largely accurate? This draft attempts to describe an non-exhaustive list of changes in the current threat environment. It suggests that there are both qualitative and quantitative differences from the environment described in RFC3552 and is intended as input to the IAB program on the Internet threat model started in 2020.

Table of Contents

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."

In 2020, the IAB approved an IAB program on the Internet threat model. One of its goals was to explore how the world has changed in terms of threats experienced and how protocol endpoints are implemented and deployed. During early discussions for that IAB program - called model-t - a natural question was raised: has the

Internet Threat Model really changed? Or, is the model described in RFC3552 still largely accurate?

The purpose of this draft is to examine the threat landscape of the contemporary Internet and answer those questions. The draft might then be used as input into the IAB's model-t process for documenting why an update to BCP72 might be needed.

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

2. BCP72 Threat Model

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. It's worth noting that 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 details describes passive attacks as those in which an attacker "reads 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 makes note of offline cryptographic attacks in which an attacker has made offline copies of packets that have 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 an an active attack." In this case, the BCP discusses spoofing packets 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 {1} show orders of magnitude increases in the size of attacks. 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.

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, and 1433 make up 99% of the targets.

The growth in the attacks on Telnet [RFC854] is a reflection of another development in the public Internet: the growth in numbers of constrained devices. Endpoints that are not capable of supporting endpoint protection software, effective encryption, or proper authentication 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].

Since 2003, there have been a variety of studies examining the growth in the number of devices connected to the Internet[2]. At the time of writing, one estimate is that the difference between the number of devices connected in 2003 and 2020 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 potentially 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 host. The client/server model of computing remained important 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 rich diversity in the capacity, connectivity, purpose of those endpoints. As a result, while the number of protocols may have not increased exponentially, the kinds 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. Today's landscape is affected by the conflict between protection from attackers and threats, and convenience. 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.

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 the less than $1 billion ISD. 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. Another example of this happens when VPNs themselves examine or exploit traffic rather than do what they are advertised to do.

3.3. Data at Rest

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 middle boxes 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 these middleboxes is such that, when protocols are developed that effectively route around them, operators and content providers sometimes object.

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.

4. Observations

If the IAB's Model T program finds that there have been both quantitative and qualitative changes to the Internet threat model, then perhaps it would be time to consider revising BCP72 to reflect those changes. In this case the IAB should provide some initial assistance to the IETF on how to proceed with the revision. Others have argued that the end-to-end architecture model of the Internet cannot be understood by just considering all of the protocol layers up to the application layer.[Arkko]

In addition, 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 connect.

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).

Finally, protection from phishing attacks in the presence of certain implementations of IDNA means that applications are implementing 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.

5. Problem Statement

BCP72 is an accurate reflection of the security threat landscape at the time which it was written. While the work of the IAB program on the Internet threat model is essential, a revision to RFC3552 is in the remit of the IETF.

BCP72 represents a too narrow view of the Internet's threat landscape. An update is needed to:

- 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;
- Reflecting on the how the growth of the number of devices connected affects the attack surface for the Internet at large;
- Research by the IAB and others on 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 revise BCP72.

6. Security Considerations

This document is entirely about security on the Internet and is intended as input into the IAB's Model T work.

7. IANA Considerations

The memo has no actions for IANA

8. References

8.1. Normative References

8.2. Informative References

Informational references are to be added to a later version of this draft.

9. Acknowledgments

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

Authors' Addresses

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

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