Network Working Group
Internet-Draft
Intended status: Informational
Expires: May 2, 2017

E. Lear

October 29, 2016

Time To End The War on Network Protection
draft-lear-network-helps-01

Abstract

Since the Edward Snowden's release of secret information, some in the IETF have taken an approach that the network is such a useful tool that it is also an enemy. With several high visibility attacks that have been based on low end systems (Things), it is now clear that not only is the network not the enemy, but that it is required to protect the system as a whole. When the network has at least some information about a device, we get a second chance to limit attacks against the device and, in some cases, a third chance to limit attacks from the device. This memo discusses ways in which network protection assists in protection of devices, and some caveats around that protection, and suggests considerations implementers and protocol developers should consider as connectivity continues to expand to new applications.

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). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

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."

This Internet-Draft will expire on May 2, 2017.

Copyright Notice

Table of Contents

1. Introduction

In June of 2013 Edward Snowden released a vast trove of secret NSA documents that demonstrated numerous vulnerabilities of the Internet architecture, that included collection of aggregate information, tapping of communication lines, hacking of devices in transit, and other means. Many of these vulnerabilities were known to be possible in theory, although the scale of such an attack was unprecedented.

The Internet Architecture Board held a plenary meeting in November 2013 in which we openly discussed these attacks, and what we would do

about them. The result was [RFC7258], which states that pervasive surveillance should be treated formally as a form of attack.

Since that time HTTP2 has been released, and work has begun on QUIC [QUIC], a transport protocol that reside atop UDP.

The premise of much of this work has been that if the network has visibility to ANY meta-information, then it is possible for a government or other similarly well-funded entity to effect a pervasive surveillance attack. The conclusion in some minds has been that the network has aided and abetted attackers to the point that its indistinguishable from an attacker. A natural, yet flawed, conclusion was endpoints alone can and must be responsible for their own protection.

Since 2013, the Internet of Things has come into its own, as connection capabilities have developed on everything from dolls to door bells. While the ability to connect to the Internet has developed, ability to maintain a secure device has not kept up. If the network cannot be part of the solution, and the device is unable to secure itself, then the device by definition will be open to attack.

This document is structured as follows. First Section 2 provides a general overview of the value and risks of sharing at various layer. Next Section 3 provides an overview of different classes of devices and the forms of attacks that occur where some amount of sharing might have provided some useful defense. We then review the role of encryption in Section 4. A basic principle is that information sharing should take place as a matter, and not as an accident, of design.

Finally we make recommendations for how devices and networks should collaborate under several different use cases.

This document considers how to address privacy considerations [RFC6973] in the context of Things. While we do not pull terminology from that document, the concepts should be readily identifiable.

2. What might be shared (and why)

Within the Internet architecture it is possible for any piece of

information to be shared. This includes application-layer information, TCP/UDP port information, source/destination IP addresses, intended communication direction, and L2 information. In addition to information shared in flight, profile information can also be shared. The following discussion motivates why these pieces of information might be shared.

2.1. Application-layer information sharing in flight

When application information is shared with a firewall or similar system, that system is in a position to validate application layer exchanges for correctness. For example, an end-to-end banking

transaction authorized by a trader may be open to audit in order to avoid fraud, or the setting of a valve in an oil well should be validated to be within a set of parameters in order to avoid a spill or worse. Some content discrimination may be necessary. For instance, some parameters may be transparent and others encrypted. The trader's order might be clear, but authentication information would be encrypted.

The challenge with this approach are threefold:

- The network access point and the end points must have an identical and up-to-date semantic understanding of the information being shared.

- In addition, the level of trust conferred to the network access point is absolute. If it is compromised, all information is revealed.

- This level of sharing also presents scaling problems whereby the network must expend processing power to determine appropriate actions.

When such an approach is used, it must be explicitly configured, and there must be an automated means to update both end points and network access points such that transactions are always properly interpreted by all parties. In addition, appropriate resources must be available on the network access points.

2.2. Transport Layer information

At this layer service information is revealed. This might indicate what applications are running in some instances, but not the content being exchanged. Layer 4 is generally considered to be so-called "meta-information", although it is information that is exposed, nonetheless. Sharing of Layer 4 information generally provides network access points with a basic understanding of services the device is using. Combined with directional information, sharing at this layer usually is sufficient to indicate which end has initiated a conversation. The simplest example is TCP packets that have or lack the ACK flag. More advance forms retain flow state.

Directionality is a key ingredient to being able to stop unwanted traffic, including malware. Simply put, "if I didn't ask for it, I don't want it." Now we apply this axiom in the context of the firehose we call the Internet. Directionality can be detected in the transport layer and as a function of the first packet seen from a particular interface. Each of these mechanisms has limitations, but each provides some level of protection. Directionality is

particularly important in environments where highly constrained devices can have their resources overwhelmed or drained, a simple example being an energy-harvesting light switch. Only the network can enforce an approach if an end node listens to any traffic at all.

A common pattern of communication for devices is that they would need DNS, NTP and perhaps either outbound or inbound web services. Use of protocols like Port Control Protocol (PCP) [RFC6887] is predicated on the assumption that meaningful protection is provided by restricting access to other ports.

This information is not quite as sensitive as application layer information, in that usernames, passwords, and other private pieces of information usually are not available to be exchanged. Processing requirements at this layer vary based on the transport protocol used and the level of protection required.

2.3. IP Layer Information

The IP layer consists of source and destination addresses. This information can be considerably more revealing than transport layer information. Based on this information, an observer may often

discern who the parties of a communication are, based on reverse address lookups or by examination of the IPv6 Interface Identifier(IID). IP addresses are, conversely, the primary discriminators that many firewalls use to determine whether a communication should be allowed. A common design pattern is that an system may offer a certain set of inbound services, perhaps even from anywhere, communicate outbound to a certain set of devices, and then not require any other communications. Many firewall rule sets are built upon this premise.

Cloud-based applications that make use of short TTL values of DNS records for load distribution have changed the nature of this game somewhat in that it is no longer sufficient to simply attend to IP addresses to authorize a service- one must also pay attention to an ever-changing mapping between address and name.

IP addresses also provide some hint at geographic location. This function is used today for many purposes, such as determining timezones or rights to certain content. That location information may be abused to track individuals.

2.4. Sharing of Device Profile Information

A device profile consists of information that describes what a device does. That information may be of a general nature shared by a manufacturer such as Manufacturer Usage Descriptions (MUD)

[I-D.ietf-opsawg-mud] or it may be of a more specific nature unique to an individual deployment or owner. The more specific the information revealed, the more sensitive. For instance, telling the network that a device is a printer may reveal very little. Telling the network that it is Eliot Lear's printer reveals ownership, and that he may have some relationship to the location in which the printer resides (like perhaps owning the or renting it).

General information along the lines of MUD provides no information about who owns the device, but does reveal what the device is. However, with the information, a network access point is in a positon to apply an appropriate set of access lists to limit the scope of attack against the end node.

Who has access to this information will depend on the means in which

the profile is communicated. For instance, if a device inventory system makes use of TLS, the information is shared only with that system. The same can be used if information is shared over EAP-TLS [RFC5216].

3. How Information Sharing Could Stop Some Attacks

3.1. DVRs

One recent attack based on the Mirai code [Krebs-MIRAI] that was made available on Github address itself to digital video recorders, cameras, and home Internet routers. Some of these devices are said to be old and not upgradable. Attempts to take over the device occurred through known telnet, SSH, and HTTP where known passwords were used. It is also said that these devices, in their normal function, make use of one or two ports.

Had the device manufacturer made use of Manufacturer Usage Descriptions (MUD), an access point could have blocked them from accessing the DVR, even though it had old firmware. An Example MUD file is given in Appendix A. Note that this file may not have stopped an already-infected device from attacking, and it requires that local deployment information be filled in for the class named "http://dvr264.example.com/controller".

As [I-D.ietf-opsawg-mud] specifies, there are numerous ways for a device to indicate the URL by which to retrieve that file. Some of those methods might reveal to an observer the type of device. To generalize guidance in this space we might say the following of network devices:

- Information about a device should not be volunteered in insecure environments.

- Where possible, such information should be encrypted to an authorized recipient.

- Information that is intended for a router, such as DHCP requests, should only be forwarded to authorized DHCP servers, and not to all ports on a network.

As we will discuss below, it may not always be possible to encrypt

information. Thus a risk tradeoff must be made: will the information cause substantial harm through leakage. In the case of DHCP, the risk is that a local device is eavesdropping. However, in this circumstance, even if the device emitted a DHCP option that was broadcast to all local devices, there would have been no additional damage, because a probe was used to determine that a device was vulnerable. As we raise the bar, however, we may wish to consider how to better protect such information through the use of other mechanisms.

3.2. Electrical Grid Attacks

A large country has already seen its electrical grid attacked. The attack was multifaceted, but specifically targeted the industrial control systems (ICSes). In one attack, breakers were opened to cause a failure. If the network between the control system and the circuit breaker actuator were gatewayed by a firewall observing commands, that attack may well have been thwarted. As discussed above, such protection comes at a high cost. In particular, the firewall itself becomes a point of attack. It also requires that the firewall understands not only the commands, but how and when it is safe for them to be executed. A poorly configured firewall might prevent a necessary emergency shutdown.

Thus we might derive some general rules of thumb regarding use of application information:

- These mechanisms should, when at all possible, be explicitly authorized, where encryption is used between all components.

- Where encryption is not possible, substantial additional levels of security should be placed around the control system so as to otherwise limit unauthorized access. This might include, for instance, a limitation on remote connectivity or use of VPNs, where access of the physical communication path cannot be controlled.

- There must be clear parameters as to what reasonable values are, and what to do in exceptional circumstances.

3.3. Mobile Medical Devices

A number of medical devices that have transceivers may now be implanted in humans. These devices are as mobile as their patients are. Such devices may be subject to nearly all classes of attack, such as unauthorized access to denial of service. All such attacks could prove deadly to the patient. The problem is complicated by the fact that these devices are generally battery operated where intended communication is expected to be rare, but responsive when needed. One approach used to address this limitation is to only enable near field communications, so that remote attacks are not possible. Another is to require a magnetic field to enable remote access, as is done with pacemakers. In these cases, ad hoc connectivity is then established. Examining the threat model, if someone is going to attack a person with a magnet, they may well have other ways to effect an attack. The key is that the magnet is essentially used as an electrical switch to enable communication. Having some local activation mechanism can prevent resource drain, where no information is gratuitously shared.

Such an approach may not be practicable in all circumstances. When that is the case, the network should be used to detect and prevent denial of service attacks, without the need to reveal identifying about the patient.

3.4. Mobile Phones

Mobile phones have been well studied. Risks associated with these devices often involve users taking actions not in their best interests, such as installing malware, or permitting excessive rights to an app.[EGEL12]. Mobile Service providers typically already have information as to devices attached to the network, in part because they often sell those devices at reduced prices with contracts.

4. Encryption and Sharing

When data is obscured via encryption, then it must be shared explicitly with intended recipients. When practicable, this is a preferred approach, but a number of problems often arise:

- Trust between parties. While some ongoing work is exploring trusted introduction [I-D.ietf-anima-bootstrapping-keyinfra], due to memory and connectivity constraints it is often difficult to establish trust between two or more parties.

- Even once a trusted introduction has occurred, ongoing key management and algorithm selection remains a challenge. The entire device lifecycle must be taken into account.

Because of poor interactions between network components and devices, many services now reside on TCP port 443, meaning that when encryption is possible, it is not possible for a firewall to filter based on service as it has been in the past.

In order to avoid tracking, a number of mobile devices are now regularly changing their L2 MAC addresses, where possible. This makes filtering based on MAC address impracticable. Similarly, devices deploying IPv6 have the ability to make use of different IPv6 Interface Identifier (IID). [RFC7721] discusses the privacy implications and threats of using stable IIDs. As that document mentions, if the IID is part of an authentication paradigm, its change means that the device itself must be reauthenticated, and may add to system fragility.

5. Conclusions

When networks take on certain functions there are some risks that must be considered. The same is true when only devices attempt to provide for their own security. An systemic and architectural approach is needed that makes use of both device and network capabilities in good measure. Such an approach must take into account both privacy and security requirements, where appropriate balances can be made.

6. Security Considerations

This document discusses the security of the Internet.

7. IANA Considerations

This section may be removed upon publication. There are no IANA considerations.

8. Acknowledgments

The author wishes to thank Brian Weis and Lee Howard for their comments. These individuals may or may not support the views contained herein.

9. Informative References

[EGEL12] Egelman, S., Felt, A., and D. Wagner, "Choice Architecture and Smartphone Privacy: There's A Price for That", 2012, <http://www.econinfosec.org/archive/weis2012/papers/Egelman_WEIS2012.pdf>.

[I-D.ietf-anima-bootstrapping-keyinfra]
Pritikin, M., Richardson, M., Behringer, M., and S. Bjarnason, "Bootstrapping Remote Secure Key Infrastructures (BRSKI)", draft-ietf-anima-bootstrapping-keyinfra-03 (work in progress), June 2016.

[I-D.ietf-opsawg-mud]
Lear, E., Droms, R., and D. Romascanu, "Manufacturer Usage Description Specification", draft-ietf-opsawg-mud-01 (work in progress), September 2016.

[Krebs-MIRAI]
"Source Code for IoT Botnet 'Mirai' Released", October 2016, <https://krebsonsecurity.com/2016/10/source-code-for-iot-botnet-mirai-released/>.

[QUIC] "QUIC Working Group Charter", 2016, <https://datatracker.ietf.org/wg/quic/charter/>.

[RFC5216] Simon, D., Aboba, B., and R. Hurst, "The EAP-TLS Authentication Protocol", RFC 5216, DOI 10.17487/RFC5216, March 2008, <http://www.rfc-editor.org/info/rfc5216>.

[RFC6887] Wing, D., Ed., Cheshire, S., Boucadair, M., Penno, R., and P. Selkirk, "Port Control Protocol (PCP)", RFC 6887, DOI 10.17487/RFC6887, April 2013, <http://www.rfc-editor.org/info/rfc6887>.

[RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", RFC 6973, DOI 10.17487/RFC6973, July 2013, <http://www.rfc-editor.org/info/rfc6973>.

[RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", BCP 188, RFC 7258, DOI 10.17487/RFC7258, May 2014, <http://www.rfc-editor.org/info/rfc7258>.

[RFC7721] Cooper, A., Gont, F., and D. Thaler, "Security and Privacy

Considerations for IPv6 Address Generation Mechanisms",
RFC 7721, DOI 10.17487/RFC7721, March 2016,
<http://www.rfc-editor.org/info/rfc7721>.

Appendix A. Example MUD File for a DVR

```
{
  "ietf-mud:meta-info": {
    "lastUpdate": "2016-10-23T14:11:52+02:00",
```



```
    "systeminfo": "DVR H.264",
    "cacheValidity": 1440
  },
  "ietf-acl:access-lists": {
    "ietf-acl:access-list": [
      {
        "acl-name": "mud-10387-v4in",
        "acl-type": "ipv4-acl",
        "ietf-mud:packet-direction": "to-device",
        "access-list-entries": {
          "ace": [
            {
              "rule-name": "clout0-in",
              "matches" : {
                 "ietf-mud:direction-initiated" : "from-device"
                 },
              "actions": {
                "permit": [
                  null
                ]
              }
            },
            {
              "rule-name": "entin0-in",
              "matches": {
                "ietf-mud:controller":
                 "http://dvr264.example.com/controller",
                 "ietf-mud:direction-initiated" : "to-device"
              },
              "actions": {
                "permit": [
                  null
```

```
                ]
              }
            }
          ]
        }
      },
      {
        "acl-name": "mud-10387-v4out",
        "acl-type": "ipv4-acl",
        "ietf-mud:packet-direction": "from-device",
        "access-list-entries": {
          "ace": [
            {
              "rule-name": "clout0-in",
              "matches": {
                 "ietf-mud:direction-initiated" : "from-device"
```



```
              },
              "actions": {
                "permit": [
                  null
                ]
              }
            },
            {
              "rule-name": "entin0-in",
              "matches": {
      "ietf-mud:controller":  "http://dvr264.example.com/controller",
                 "ietf-mud:direction-initiated" : "to-device"
              },
              "actions": {
                "permit": [
                  null
                ]
              }
            }
          ]
        }
      }
    ]
  }
}
```

Author's Address

Eliot Lear

Email: lear@ofcourseimright.com