

Subverting Stateful Firewalls with Protocol States

Amit Klein
Bar Ilan University
aksecurity@gmail.com

Abstract—We analyzed the generation of protocol header fields in the implementations of multiple TCP/IP network stacks and found new ways to leak information about global protocol states. We then demonstrated new covert channels by *remotely* observing and modifying the system’s global state via these protocol fields. Unlike earlier works, our research focuses on hosts that reside in *firewalled* networks (including source address validation – SAV), which is a very common scenario nowadays. Our attacks are designed to be *non-disruptive* – in the exfiltration scenario, this makes the attacks stealthier and thus extends their longevity, and in case of host alias resolution and similar techniques – this ensures the techniques are ethical. We focused on ICMP, which is commonly served by firewalls, and on UDP, which is forecasted to take a more prominent share of the Internet traffic with the advent of HTTP/3 and QUIC, though we report results for TCP as well.

The information leakage scenarios we discovered enable the construction of practical covert channels which directly pierce firewalls, or indirectly establish communication via hosts in firewalled networks that also employ SAV. We describe and test three novel attacks in this context: exfiltration via the firewall itself, exfiltration via a DMZ host, and exfiltration via co-resident containers. These are three generic, new use cases for covert channels that work around firewalling and enable devices that are not allowed direct communication with the Internet, to still exfiltrate data out of the network. In other words, we exfiltrate data from isolated networks to the Internet. We also explain how to mount known attacks such as host alias resolution, de-NATting and container co-residence detection, using the new information leakage techniques.

I. INTRODUCTION

Our research focuses on how to exploit global state in the UDP and ICMP protocol stacks to mount attacks against firewalled networks and hosts. The attacks we are interested in are data exfiltration, host alias resolution and de-NATting, and co-residency detection – see Table I. In some cases, the same attacks can be carried out over TCP, in a more efficient manner than existing techniques.

A. Motivation

Our motivating example is a highly secured network, wherein internal machines are not allowed to communicate directly with the Internet. This is enforced by a network firewall, which prevents the network hosts from accessing the Internet, but may allow them to access services on a

TABLE I: Attack Use Cases

Use Case	Sender Location	Target Host	Receiver Location
Firewall Piercing (exfiltration)	Isolated network	Firewall	Internet
Exfiltration (via DMZ)	Isolated network	Firewalled host in DMZ	Internet
Co-Resident Container Exfiltration	Isolated network	Two co-resident containers (internal+external)	Internet
Alias Detection and de-NATting	Internet	Firewalled/NATted host	Internet (=Sender)

neighboring internal network and/or management access to the firewall itself. Such *isolated* networks can be found in financial institutions, defense industries, hospitals, sensitive infrastructure sites, etc. The attacker’s goal is to exfiltrate data out of an isolated network, to the Internet. The compromised data may consist of small “strategic” pieces of data, such as cryptographic keys, or compressed text files.

We assume that an attacker controls a compromised internal machine. For example, the attacker may have installed malware on that machine, by means out of scope for this research. The attacker’s goal is to exfiltrate data from this machine. It should be noted that the machine may be part of an attacker network of compromised machines, such that the attacker funnels the data to be exfiltrated from other machines to the machine from which the exfiltration is to take place.

Our techniques are able to circumvent the firewall’s enforcement, and allow the attacker to exfiltrate data from a machine in an isolated network to an Internet host the attacker controls. The exfiltration throughput is hundreds to thousands bits per hour (except for Windows), which suffices to transmit e.g. a 128-bit key in less than 22 minutes.

It should be stressed that this is an example to one use case (firewall circumvention); in the paper we also outline other attacks, such as host alias resolution and de-NATting, and co-residency detection. The motivation for these use cases is to provide operational intelligence as part of a larger attack effort. Detecting that two services reside on the same host may enable an attack against a weaker (vulnerable) service to take over a host which also runs a more strategic (but not vulnerable in itself) service.

B. Organization

Next in this section we provide background on terms and concepts we use throughout the paper (stateful firewalls, host alias resolution, and co-residency), followed by a summary of our contributions. In Section II we discuss related work. Section III describes the current algorithms that generate the

protocol fields of interest. Section IV outlines our attack concepts, Section V then explains how the concepts are implemented per-OS, and Section VI describes our experiments. A discussion of remediation options is provided in Section VII. We draw our conclusions in Section VIII, and finally we report on the vendor status in Section IX.

C. Stateful Firewalls and SAV

A *stateful firewall* is a “network-based firewall that individually tracks sessions of network connections traversing it” [44]. According to [28], “[a] stateful inspection firewall is the de facto standard for network protection”. SAV, and some features of stateful firewalls can eliminate known data leakage attacks. We explain this below, and we note that our techniques are not affected by SAV and by stateful firewalls.

In response to an incoming invalid packet – a TCP packet or a UDP packet destined to a closed port, or a TCP packet (other than SYN-only to an open port) that does not belong to an already established TCP session – a firewall may either silently *drop* the packet, or actively *reject* it. Rejection is carried out by sending back a TCP RST packet in the TCP case, or by sending back an ICMP Type 3 Code 3 packet (“port unreachable”) in the UDP case.

Dropping invalid packets is considered a somewhat more secure choice (at the expense of a slight overall performance degradation due to a delay in responsiveness at the remote party when no rejection packet is sent), and is oftentimes set as the default firewall behavior. In Appendix A, we report the results of an experiment in which we measured and estimated the portion of widely-used Internet hosts that do not send any response for UDP and TCP packets arriving to closed ports. We found that the vast majority of the Internet hosts silently drop such UDP/TCP packets. Therefore, attacks that assume a response for an invalid, connection-less UDP/TCP packet will fail for most servers. Our attacks, on the other hand, only use valid and acceptable packets, and are thus unaffected by stateful firewalls.

A firewall can also employ Source Address Validation (SAV), such that it does not allow traffic from a first network into a second network with source addresses outside the first network. This security measure can prevent some trivial information leaking attacks. For example, a host *A* on an isolated, internal network which is prevented from sending packets to the Internet can send a spoofed packet whose source address *C* is an Internet address, to host *B* in the DMZ network which is allowed to send data to the Internet. Host *B* will then respond by sending a packet to *C*, an Internet address. This can be used as a covert channel signal. SAV prevents this by dropping *A*’s spoofed packets since their spoofed source address *C* does not match their origin (the internal network).

Henceforth, we use the term *firewall* as a shorthand for “a stateful firewall that silently drops invalid packets, with SAV between all its networks”.

D. Host Alias Resolution, De-NATting and Co-Residency Detection

Note: we use the term *host* throughout this paper to denote a single instance of an operating system kernel. This can be

TABLE II: Information Leaking Fields and Their Channel Properties

OS	Protocol	Field	Sender Capabilities	Channel Bandwidth
Linux	UDP/IPv4, ICMP/IPv4	IPv4 ID	IP Spoofing	400 b/h (UDP), 3000 b/h (ICMP)
Windows	UDP/IPv4, TCP/IPv4	IPv4 ID	IP Spoofing	47.4 b/h (UDP), 41.4 b/h (TCP)
macOS	UDP/IPv4	IPv4 ID	Web Traffic Emission	1800 b/h
macOS	ICMP/IPv4	(rate limit)	Web Traffic Emission	1800 b/h (estimated)
OpenBSD	UDP/IPv4, TCP/IPv4, ICMP/IPv4	IPv4 ID	Web Traffic Emission	720 b/h (UDP, throttled), 360 b/h (TCP, throttled), 4200 b/h (ICMP)
NetBSD	TCP/IPv4, TCP/IPv6	TCP ISN	Web Traffic Emission	7200 b/h (throttled)
NetBSD	TCP/IPv6	IPv6 Flow Label	Web Traffic Emission	Untested

a physical host, or – if virtual machines are used – a single virtual machine. We use the term *physical host* to explicitly exclude the latter case.

A *host alias resolution* [40] (also *host alias detection* [10]) technique enables a remote party to infer whether two given IP addresses map to the same host or to two different hosts. We can extend this into inferring whether two services (defined by IP address, the listening port and the transport protocol – TCP or UDP) map to the same host or not. For example, one can ask whether 1.2.3.4:80 (over TCP) and 5.6.7.8:443 (over UDP) map to the same host.

De-NATting [4] [35] is a use case of host alias detection, wherein the host(s) in questions reside behind a NAT, and thus the two services are mapped from their external endpoints into internal endpoints by a NAT device.

In this research, we focus on *server de-NATting*, meaning, discerning whether two external endpoints refer to the same host or to two different hosts. This is in contrast to *client de-NATting*, which refers to discerning whether two outbound connections (designated by their external address endpoints) are made by a single client or by two different clients. The material difference is that a NATted client is likely to use a single internal IP address, whereas a NATted server may use two IP addresses for two different services. Some operating systems use the (IP_{SRC}, IP_{DST}) tuple to generate network fields, making it trivial to de-NAT clients, but not servers.

Co-Residency Detection for Containers [39] is a concept similar to host alias detection, which aims to detect whether two endpoints on two different containers are hosted by the same physical host. If virtualization is used, then the containers may run on two different virtual machines. We are interested in a variant of this concept, which is to detect whether two endpoints on two different containers are hosted by the same *operating system kernel*.

Container technology (e.g. Docker) is available for Linux and Windows. Our co-residency detection attacks apply to Linux, but not to Windows.

E. Our Contribution

- New **information leakage techniques** based on the statefulness of IPv4 ID, TCP ISN (IPv4 and IPv6) and IPv6 flow label in popular server operating systems: Linux, Windows Server, macOS Server, OpenBSD and NetBSD. In the IPv4 ID and TCP ISN cases, this is done without accurately predicting the full values of these protocol fields. Our results are summarized in Table II. These techniques are non-disruptive, and can be used even when the target host is firewalled, including SAV. The attacks are algorithmic, and therefore are not hardware-specific, nor do they require fine-grained timing or rely on race conditions. And since they rely on kernel implementations, they are agnostic to the nature of applications/services the operating system runs.
- Exfiltration from a firewalled network (including SAV) and exfiltration between co-resident containers – a new **covert channel use case** for exploitation of stateful IPv4 ID, TCP ISN and IPv6 flow label. In this use case, an internal device is prevented from directly connecting to the Internet, by a firewall, but is allowed to connect to the firewall itself, or to a DMZ host, or to an internal container co-residing with a 3rd party container. Our attacks exploit the host’s global protocol states to exfiltrate data from the internal device to the Internet. The covert channels’ bandwidth (in bits/hour) depends on the information leakage techniques’ bandwidth listed in Table II.
- Constructing a covert channel based on the new information leakage techniques and exploiting it for well known attack use cases: server alias resolution, de-NATting, container co-residency detection, and idle scanning. Previously, these attacks were inapplicable either due to the underlying information technique being mitigated, or due to the disruptive nature of the underlying technique or other limitations.
- A description of the present day algorithm used by Windows to manage (IP_{SRC}, IP_{DST}) tuples, including IPv4 ID management.

II. RELATED WORK

A. IPv4 Global State and Side Channels

1) *Prediction of IPv4 ID values:* The IPv4 ID field is a 16 bit identifier used to reassemble IP packets from smaller fragments. This field is populated by the OS regardless of whether fragmentation is actually needed. If the IPv4 ID field is generated in a predictable manner, e.g. sequential increment per packet, then it is possible to sample the field twice and determine whether the target host sent another packet in between. For this approach, the IPv4 ID field must be highly predictable, i.e with a near-zero entropy. Note of course that if the ID is static (e.g. 0), which is legitimate for atomic (DF=1)¹ packets [43, Section 4], then no information is leaked. However, some operating systems (e.g. Linux) do set DF=1 for

¹The IPv4 DF (“Don’t Fragment”) header flag controls fragmentation by routers. Setting DF=1 prohibits routers from fragmenting the packet.

short packets, yet still populate the ID field with a non-static value.

The original idle-scan technique [38] exploited sequential IP IDs which were very common in the operating systems of that era (1998). Later in 2008, [23] cryptanalyzed the IPv4 ID mechanism of OpenBSD and macOS Server and showed it to be predictable. In response, OpenBSD and macOS Server changed their algorithm. In 2019, [25] attacked the IP ID generation in Windows TCP/IPv4 and UDP/IPv4, and in Linux UDP/IPv4. In response, Windows and Linux changed their algorithms. Recently, [24] showed that the Linux TCP/IPv4 ID field is also predictable, and in response Linux changed that algorithm as well.

The only operating systems nowadays that generate predictable (sequential) IPv4 IDs are FreeBSD (UDP/IPv4 only) and NetBSD (only for packets longer than 68 bytes, thus excluding e.g. TCP RST packets). With the exception of FreeBSD and NetBSD, all operating systems we surveyed upgraded their IPv4 ID generation scheme to a non-predictable algorithm, hence attacks that rely on IPv4 ID prediction are no longer applicable to them.

Moreover, in general, idle-scanning, which relies on the pivot server to send an RST packet for “unexpected” SYN+ACK packet, does not apply to a firewalled pivot server. A firewall is likely to silently drop the SYN+ACK packet, or alternatively the firewall may respond with an RST packet, without forwarding it to the pivot server. We do, however, note that our NetBSD flow label attack can be used for *efficient* idle scanning TCP/IPv6 hosts with NetBSD as a pivot server, when the latter is not firewalled.

In 2008, Danezis described how to construct covert channels from incremental IPv4 IDs [11]. But Danezis’s use case is different than ours. Danezis assumes the monitored party can still connect to the Internet, and as such, he focuses on channel robustness, rather than on the technicalities of implementing the technique for firewalled networks. Additionally, that work does not describe UDP-based techniques, IPv6-based techniques and container-based techniques.

In contrast, the global IPv4 ID states we describe exist in modern IPv4 ID generator implementations. These global states are typically subtle and not easily observed.

2) *Attacks on the Linux IPv4 ID Generation Algorithm for UDP and TCP Connection-Less Packets:* The Linux IPv4 ID generation algorithm for *connection-less packets* (ICMP, UDP packets unless `connect()` is used, and TCP RST packets not in a connection context) is described in Section III-A. This construction was attacked in several recent papers.

Some attacks against Linux TCP connections exploited this algorithm, either when used for TCP RST packets [46] [3], or through a “downgrade attack” from the more secure TCP IP ID generation scheme [14]. However, these attacks do not affect Linux UDP and ICMP traffic, which is the object of our research. Additionally, these attacks practically require the attacker to own thousands of IPv4 addresses.

In 2019, [25] exploited the Linux IP ID generation structure to find partial hash collisions and extract the hashing key. In response, the hash function was modified to use a larger key.

A recent DNS cache poisoning attack [47] involving Linux-based recursive resolvers hinges on the fact that when two machines are behind the same NAT, their external IPv4 address is identical, and hence the packets sent by a Linux host to both machines share the IPv4 ID counter. But to exploit this for exfiltration, the sender and receiver must be on the same network and behind the same NAT, which is inapplicable to an exfiltration scenario.

Similarly to the above attacks, our attacks exploit the counter array of the Linux IPv4 ID algorithm, but in contrast to the previous attacks, our attacks target ICMP and UDP services, and they do not require ownership of thousands of IPv4 addresses, or having the sender and the receiver reside on the same network, and can target any host.

3) *IPv4 De-fragmentation Cache*: Another global IPv4 state is the memory limit for the de-fragmentation cache, a kernel data structure that maintains the incoming IPv4 fragments. The limit can be applied per source IP address, or globally. Overflowing the de-fragmentation cache can be detected remotely, as explained in [16], and can thus be used as a side channel. This technique requires IPv4 fragments to traverse the Internet unmodified. However, the prevalence of Carrier Grade NATs (CGNATs) reduces the attack surface of this technique. In CGNATs, the attacker cannot space the arrival time to the target host of fragments that belong to the same IP packet. Rather, the first fragment will await virtual de-fragmentation on the CGNAT device, and will only be transmitted when the second fragment arrives at the CGNAT. In fact, anecdotal evidence from our lab’s ISP indicates that CGNATs may employ an even more destructive strategy, wherein they reassemble the IPv4 packet and forward the complete packet.

Additionally, since this attack overflows the de-fragmentation cache, it is disruptive, as organic inbound fragments may be evicted from the cache before re-assembly, resulting in dropping of organic inbound packets.

B. UDP Global State and Side Channels

The UDP protocol does not contain global states per-se. However, [2] describes a global state based on a rate limit imposed by some kernels on ICMP protocol “port unreachable” messages, sent in response to incoming UDP packets destined to closed ports. It should be noted though, that firewalls handle UDP packets sent to closed ports themselves (without forwarding them to the destination host), and are likely not to send back ICMP messages, due to security reasons. In Appendix A we report that an estimated 85% of the widely-used Internet hosts/firewalls do not send ICMP “port unreachable” messages.

C. TCP Global States and Side Channels

Our attack on NetBSD is TCP-based, and our attacks on Windows and OpenBSD IPv4 ID can also be used in TCP protocols. Therefore, we also survey TCP-based side channels, but only if they are relevant to any of these operating systems.

In 2005, Kohno et al. [26] demonstrated device identification based on the device clock skew, observed in the TCP timestamp field. This work was later improved in [31].

Windows, OpenBSD and NetBSD do not implement TCP timestamps useful for the [26] and [31] attacks (Windows Server does not send a TCP timestamp at all, and OpenBSD and NetBSD have TCP timestamp at 0.5s resolution). In the past, TCP Timestamps could be used as a “passive” side channel for host alias resolution, as described in [7] [20], but nowadays this only applies to macOS.

In 2010, Ensafi et al. [13] used model checking to discover information leakage due to the rate limit on a global TCP RST counter. Exploiting this requires the sender client to send the server either SYN packets for a closed port, or otherwise packets that purportedly belong to a non-existing TCP circuit. In both cases, a stateful firewall may silently drop the packets and thwart the attack. In Appendix A, we report that an estimated 92% of the widely-used Internet hosts/firewalls drop such packets.

Additional side channels are described by [13], [8] and [9], which are based e.g. on TCP SYN cookies and TCP challenge ACKs. But these mechanisms are not employed (at least not by default) by Windows, OpenBSD and NetBSD.

We see, therefore, that for firewalled Windows Server, OpenBSD and NetBSD hosts, there are no currently known remotely exploitable TCP global states.

D. Host Alias Detection and De-NATting

Early host alias detection and de-NATting relied on global states that are trivial to observe and follow, and compared such global states between two endpoints. Bellovin’s seminal 2002 de-NATting paper [4] is based on a globally incrementing IPv4 ID as a global state. This concept was later applied to host alias resolution, e.g. with IPv6 ID [5]. Client de-NATting based on the IPv4 ID of outbound DNS queries over UDP was described in [35]. This attack relies on the incremental nature of the IP ID for a fixed (IP_{SRC}, IP_{DST}) tuple, and thus cannot be applied to *server* de-NATting where a single host may use multiple source IP addresses.

Nowadays, IP IDs are not generated via a global counter (except for FreeBSD and NetBSD), and TCP timestamps are usually generated per TCP connection (except for macOS), thus these techniques are no longer in effect. Conceptually, it is still possible to mount host alias detection and de-NATting attacks using the more sophisticated global states described in the previous sub-sections, but here too, the severe limitations described above for these techniques apply.

A host alias resolution technique based on packet delay sequences is described in [41], however this technique aims at *router* IP aliasing, which does not address the question of whether two IP addresses are served from a single host or from two distinct hosts on the same network.

E. Container Co-residency Detection and Cross Container Information Leakage

A container co-residency detection technique is described in [39], but this attack is local, i.e. the attacker’s container must co-reside with the target container. Our attack can *remotely* detect that two containers reside on the same host.

Cross-container leaks are reviewed in [15], but the attacks described there cannot be remotely mounted, whereas our

TABLE III: Exfiltration from Networks

Target Network	Typical Throughput	Typical Requirement	Examples
Air-Gapped	8b/h-4Kb/s	Proximity (0.4m-30m)	[18], [19]
Isolated	1b/h-10Kb/h	IP spoofing (in some cases)	This paper , [31]
Regular (Firewalled)	100Kb/s and above	None	[30], [6], [45], [22]

attack is remote, and does not require the attacker to control a co-resident container.

F. Other Side Channels and Exfiltration Techniques

Our exfiltration attacks are focused on isolated networks. Isolated networks are less secure than air-gapped (or nearly air-gapped) networks, but are more secure than ordinary networks. As can be expected, exfiltration techniques that work on less secure networks generally tend to have better throughput, but cannot be applied to more secure networks. Specifically, techniques that apply to less secure networks typically cannot be applied to isolated networks, and therefore cannot be compared, throughput-wise, to our techniques, that do work against isolated networks. Likewise, our exfiltration techniques do not work against air-gapped networks, or nearly air-gapped networks (networks behind a network pump).

Exfiltration from air-gapped networks requires close proximity (0.4m-30m) [17], which is a serious limitation. The throughput varies drastically among the techniques, from 8b/h to 4Kb/s [19]. A typical example is [18] which describes an acoustical covert channel based on hard drive noise, yielding 3b/s throughput.

Our work is a rare example of exfiltration from an isolated network (without the proximity constraint). Another example is [31], which is effective against Linux TCP services only, with 2b/h-8b/h throughput.

Finally there are examples of exfiltration techniques that do not apply to isolated networks. For example, [30] uses a web counter as a covert channel, but this requires the DMZ target host to have (1) a shared web application between the internal and the external networks; and (2) to have a web counter in that application. [6] uses public cloud service as a covert channel, and [45] uses multiple public services as a covert channel. All these covert channels are very unlikely to apply for isolated networks as their prerequisites are obviously inconsistent with network isolation. In 2004, [22] demonstrated a covert channel via DNS queries (“DNS tunneling”). A throughput analysis can be found in [1]. Though it focuses on the downlink, it can be inferred that the uplink throughput is in the hundreds of Kb/s, many orders of magnitude better than our techniques. However, since DNS tunneling is well-known and well-understood for almost two decades, it is very unlikely to find isolated networks with unrestricted access to an Internet-connected DNS resolver.

Table III compares exfiltration techniques for the three network types.

III. CURRENT PROTOCOL HEADER FIELD GENERATION ALGORITHMS

A. Linux IPv4 ID for Connection-Less Protocols

Starting with kernel v3.16, the Linux IPv4 ID generation algorithm for connection-less packets (e.g. UDP, ICMP) uses an array β of 2048 counters, and a corresponding array τ of their last access times. Algorithm 1 describes how an IPv4 ID is generated. Note that the time is measured in “jiffies” (a clock whose frequency is set at kernel compile time: $f=250\text{Hz}$ by default, $f=1000\text{Hz}$ for some MIPS platforms), h is a hash function, $proto$ is the IANA protocol number (1 for ICMP and 17 for UDP) and net_key adds a container-dependent quantity, if applicable. Also define $\text{RANDOM}(\emptyset) = 0$.

Algorithm 1 Linux IPv4 ID Generation for Connection-less Protocols

```

1: procedure GENERATE-IPID
2:    $i \leftarrow h(IP_{DST}, IP_{SRC}, proto, net\_key) \bmod 2048$ 
3:    $hop \leftarrow 1 + \text{RANDOM}(\{0, \dots, t_{now} - \tau[i] - 1\})$ 
4:    $\beta[i] \leftarrow (\beta[i] + hop) \bmod 2^{16}$ 
5:    $\tau[i] \leftarrow t_{now}$ 
6:   return  $\beta[i]$ 

```

B. Windows Server IPv4 ID

Note: since Windows is closed-source, extracting via reverse engineering, and documenting the Windows Server IPv4 ID generation logic is part of our contribution.

As of version 1903, Windows Server generates the IPv4 ID of outbound IPv4 packets as follows. Windows implements the IPv4 ID as a counter per each (IP_{SRC}, IP_{DST}) tuple that corresponds to outbound traffic, where IP_{SRC} is an address on the local machine, and IP_{DST} is a peer address. Such tuple-specific data is kept in an object called `Path`. The set of `Path` objects is maintained in a hash table called `PathSet`. Windows keeps a `PathSet` per compartment, i.e. one IPv4 `PathSet` instance per container. A new `Path` object is created upon the first attempt to send a packet from IP_{SRC} to IP_{DST} , where there is no existing `Path` object with these indices. When a new `Path` object is created, the IPv4 ID field is initialized with random data. Then, for each packet sent from IP_{SRC} to IP_{DST} , the IPv4 ID in the corresponding `Path` object is incremented and used as ID. Naïvely speaking, this setup provides no leakage, since each tuple runs its own independent IPv4 ID counter. However, an important aspect of the `PathSet` life-cycle was left out so far – how Windows handles the potentially almost infinite growth of `PathSet`. Obviously `PathSet` cannot be allowed to grow indefinitely, so `PathSet` objects must occasionally be purged. Indeed, Windows removes objects from the `PathSet` in batches, which we term *purge sequences*. Generally speaking, a purge sequence can be triggered either by `PathSet`’s size exceeding some (pretty high) thresholds, or by a `PathSet` growth rate of 10,000 new `Path` objects per second (“flood detection”). In both cases, a purge sequence is initiated, which goes over all the `Path` objects in the `PathSet` at a rate of 2000/sec and removes “stale” `Path` objects – objects that were last accessed over 60 seconds² ago.

²10 seconds in Windows 10, determined by `TcpipIsServerSKU`.

C. macOS Server (Connection-Less Protocols) and OpenBSD IPv4 ID

Both macOS (for connection-less protocols only – e.g. UDP and ICMPv4) and OpenBSD implement their IPv4 ID generation along the concept of keeping a list of the M recently used IPv4 ID values, picking a random ID value out of the remaining $65536 - M$ values, and updating the list of used values accordingly (popping the oldest value and pushing the freshly generated value). For macOS, $M = 4096$, and for OpenBSD, $M = 32768$. Thus, it is guaranteed that the IPv4 ID value is unique in intervals of length M , at the cost of reducing the entropy from 16 bits to $\log_2(65536 - M)$ bits.

Note: the IPv4 ID in macOS is generated at the macOS kernel – XNU. Therefore, the attack techniques we describe are actually inherent to XNU. However, macOS Server is the only widely-used operating system based on the XNU kernel, therefore we name macOS as the attack target, rather than XNU.

D. NetBSD TCP ISN

A TCP Initial Sequence Number (ISN) is generated by each TCP connection party at the beginning of the connection, as a 32-bit random starting point for the sequence numbers it generates for the connection. A TCP client generates and sends ISN with the TCP SYN packet, and a TCP server generates and sends ISN with the TCP SYN+ACK packet.

NetBSD generates the most significant 8 bits of the ISN as a sum of a 2Hz timer and a TCP connection counter, modulo 256.

IV. ATTACKS

A. Threat Model

All our attacks are built around the same concept. In our attacks, there is a sender, a receiver and a target host (or hosts), identified by one or more endpoint designations, where an endpoint is defined as a combination of IP address, protocol – TCP, UDP or ICMP, and port (for TCP and UDP only). The attacker controls the sender and the receiver entities, but has no control whatsoever over the target host(s). Therefore, one of the main concerns of this research is to determine what attacks can be mounted in various implementations (operating systems) and roles of the target host. Thus, when we refer to an attack against an operating system, it is a shorthand to an attack against a target host running that operating system.

In one use case, the target host is a firewall preventing the sender and receiver from communicating directly. In other use cases, the target host is firewalled. The receiver and sender cooperate to fulfill their task, and are furthermore synchronized to a sub-second accuracy.

We use a covert channel through the target host wherein its global state can be read and altered by sending and receiving packets, as described below. In Section V we will see that such a covert channel exists for many target host operating systems, though the covert channel details differ between the operating systems.

The sender entity signals a single bit by sending packets to one of the target host’s designated endpoints. In most of

our attacks, the sender can be a non-privileged process at the system it runs on, sending conventional packets. In fact, for web protocols, such traffic can be emitted by an HTML page rendered by a browser. However, in the Windows and Linux attacks, the sender must be able to spoof the source IP address of its packets, thus the sender must run as a privileged process on its machine.

The receiver entity reads this bit by sending packets to one of the target host’s designated endpoints (not necessarily the one used by the sender), receiving packets from the target and analyzing them. The receiver entity typically needs access to the raw packets at the IP level, and as such may need to run as a privileged process on its system. This is a reasonable assumption since the receiver is an attacker-controlled machine on the Internet.

In general, the attack proceeds by having the receiver send and receive some packets, thus recording the global state of the target host. Then the sender either sends packets to the target host (thereby changing its global state) to signal the bit “1”, or does nothing (thereby retaining the global state in the target host) to signal the bit “0”. Finally, the receiver reads the global state again, and determines whether it changed or not, thereby reading “1” or “0” respectively.

Naturally, there are timing constraints involved. For example, the receiver must conclude its first global state extraction before the sender is allowed to send its packets, and the sender must conclude its global state modification (if it takes place) before the receiver is allowed to send packets for its second global state extraction. Additional timing constraints are required for Windows covert channels.

Our attacks work best against network-wise idle (or almost idle) target hosts. Organic traffic may interfere with our techniques, although a minimal amount of traffic can be easily mitigated by adding some margins to the threshold parameters. For example, our Linux technique (Section V-A) expects an increase of no more than $\Delta t \cdot f$ in some counter when the machine is idle. When some traffic happens to use that counter, it is incremented further, say by up to n packets per Δt time slice. In such case, we choose $M > n + \Delta t \cdot f$ (instead of $M > \Delta t \cdot f$). A mild amount of traffic can be mitigated by introducing redundancy (e.g. majority over three same-bit transmissions) and/or error correction codes (e.g. Reed-Solomon [36]).

We demonstrate our UDP attacks using the HTTP/3 protocol. Specifically, the sender sends QUIC Initial packets [21] to the target host, and the receiver sends Initial packets and receives QUIC response packets. We demonstrate our TCP attacks using the HTTP/1.1 and SSH protocols. In both cases, the attack uses open UDP/TCP port on the target device. In our ICMP attacks, we use the standard, built-in ICMP Echo Request-Reply mechanism, which is commonly open³ to the Internet. In Appendix A we report that 89% of the widely-used Internet hosts/firewalls surveyed responded to ICMP Echo requests.

In the firewall bypassing scenarios, we focused on data exfiltration attacks, but in general, the opposite direction –

³We do not consider opening ICMP to the Internet a security vulnerability. The security issue is with the implementation of the underlying IP layer.

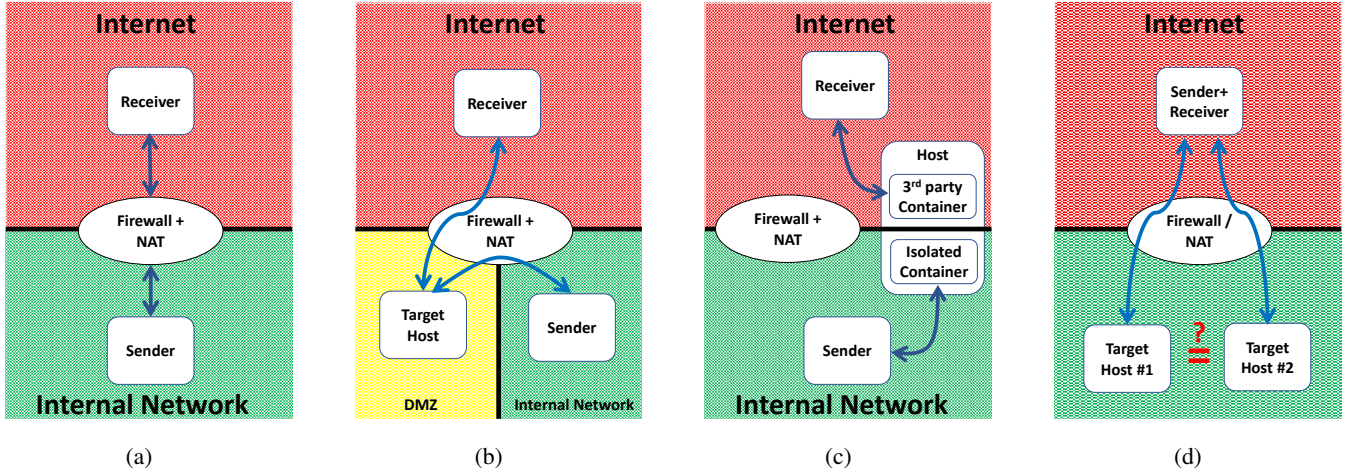


Fig. 1: Network Diagram for (a) Firewall Piercing Exfiltration; (b) Exfiltration (via DMZ); (c) Exfiltration between Co-Resident Containers; (d) Host Alias Detection and de-NATting

sending data into an internal network – is also possible. This can facilitate e.g. a Command & Control channel for the malware. We demonstrated this direction in one test (LF-1).

B. Source Address Spoofing and Sender Privileges

For the Linux and Windows attacks, the sender needs to be able to spoof the source IP address of its packets. This requires the sender to have root privileges on a machine which is on an internal network. With root privileges, spoofing is typically trivial, using e.g. raw sockets or libpcap. The firewall’s SAV protection still allows a device on the internal network (the sender) to send packets with a source address from the same (internal) network to the firewall or to the DMZ. This suffices for our firewall subversion and exfiltration attacks. For the host alias resolution attacks where the sender is on the Internet, SAV employed by ISPs may prevent address spoofing. According to [29, Fig. 4], 11.9%-30.5% of the ASes do not employ such filtering on their customers (*outbound* filtering), hence can be used by the attacker to send spoofed packets from. Naturally, the attacker is free to choose such a network, thus having enough non-filtering networks guarantees that the attacker can mount the attack. We conclude therefore that packet spoofing over the Internet is feasible, hence the Linux and Windows host alias resolution attacks over the Internet are practical.

C. Firewall Piercing Exfiltration

This use case places minimal requirements on the network topology. The attacker’s goal is to exfiltrate information from a compromised machine (the sender) in a well guarded enterprise network. It is assumed that the sender cannot send packets to the Internet. Fig. 1a illustrates the network topology (each network has its own background color, blue arrows represent direct communication, bold black lines define network boundaries enforced by the firewall).

The attacker uses an ICMP-based covert channel to exfiltrate one bit at a time from the sender to the receiver. The only requirement is that the firewall responds to ICMP Echo requests both from the Internet and from the internal network.

D. Firewall Subversion and Exfiltration (via DMZ)

In this use case, the attacker’s goal is to exfiltrate information from a compromised machine (the sender) in a well guarded enterprise network. It is assumed that the sender cannot send packets to the Internet directly. Fig. 1b illustrates the network topology. The sender can send and receive packets to/from a server host in a DMZ segment of the enterprise network (the target host). This server serves both internal clients such as the sender, and external Internet clients such as the receiver. The target host is protected by a stateful firewall both from the internal network and from the Internet. The firewall also enforces SAV between the three networks (internal network, DMZ and Internet).

The attack uses the covert channel to exfiltrate one bit from the compromised machine (the sender) through the target host (the DMZ server) to the attacker machine on the Internet (the receiver).

E. Exfiltration via Co-Resident Containers

A special case of the exfiltration technique described in Section IV-D is exfiltration between containers. Fig. 1c illustrates this attack scenario. In this case, there are two co-resident containers on the same host (kernel), which are supposedly completely isolated. The “isolated container” resides in a completely isolated local network, shown as the green network in Fig. 1c. The other container, depicted as “3rd party container” can belong to any entity, and is only required to run an Internet-facing service in the protocol designated for exfiltration. None of these containers is assumed to be controlled by the attacker.

A malicious sender that resides on the local network can exfiltrate data to an Internet receiver using the same technique described in Section IV-D, assuming the global state is shared between the containers.

In a variation of this use case (not depicted in Fig. 1c), the sender may be malware running on the isolated container, without root privileges. Such a sender can use 127.x.y.z as

local address destinations, instead of spoofing packet source addresses.

F. Host Alias Resolution, De-NATting and Co-Residency Detection

In this use case, the attacker’s goal is to determine whether two endpoints reside on the same host (even across containers, if possible) or not. The attacker is an Internet machine, and the endpoints are located behind a stateful firewall. The attacker runs both the sender logic and the receiver logic, thus there is no importance in distinguishing between the receiver and the sender, and we will use the term “attacker” for both. The attack applies the receiver logic to one endpoint, and the sender logic, specifically writing the bit “1”, to the other endpoint. If both endpoints reside on the same host, the receiver will read the bit “1”, but if the two endpoints reside on different hosts, then writing to one host does not affect the other host, and the receiver will read the bit “0”. This use case is illustrated in Fig. 1d. If the specific covert channel works across containers, then so will this attack.

Since the attack implementations for host alias resolution and de-NATting are identical, we will henceforth use the term “host alias resolution” to denote both, and we will note whether the attack also applies to container co-residency detection, per case.

V. INFORMATION LEAKAGE AND COVERT CHANNELS: GLOBAL STATES IN PROTOCOL FIELDS

We now show the existence of information leakage in the ICMP/IPv4 and UDP/IPv4 implementations of some popular server operating systems, and in the TCP implementation of NetBSD. The information leakage enables the receiver – a client of the server – to obtain information on the state of the server. We then show how to exploit this weakness to construct a covert channel, in which a sender client, typically running on a different machine than the receiver, signals a bit by forcing the server state to meet a certain condition or fail to meet the said condition. Thus, the receiver and sender engage in a covert channel protocol in which the sender sends an arbitrary bit to the receiver. Our attacks share some common principles, but differ in their technical details between operating systems. In some cases, our techniques can also be applied to TCP/IPv4.

A. Linux Connection-Less IPv4 ID

The global state that enables a covert channel is the IP ID counter used by the target host to send packets to the receiver, i.e. $\beta[R]$ (defined in Section III-A) where

$$R = h(IP_{RCVR}, IP_{TARGET}, proto, net_key) \bmod 2048$$

The sender needs to send packets to the target rapidly. When the sender can force the target to send more than f packets per second from $\beta[R]$, the receiver will be able to detect this by sampling packets generated by the target host using the $\beta[R]$ counter, and observe that $\beta[R]$ incremented at a rate $> f$. Whereas when the target host does not receive sender packets, $\beta[R]$ will increment at a rate $\leq f$.

There is a problem with the above approach, though. The probability of the sender’s IP address hash to collide with R is

$\frac{1}{2048}$. The solution is simple: note that the sender needs only to force the target host to send packets whose IP ID is generated using $\beta[R]$. The sender does not need to actually receive and read those packets. And so, the sender can simply craft packets that elicit the target host’s response, spoof the source address of these packets (to an unused internal network IP address), go over enough source addresses to practically guarantee that one of the responses uses $\beta[R]$ (as explained below), and thereby fulfill the condition for the attack.

We now expand the above technique into a covert channel. For simplicity, we assume that the receiver owns a single Internet IP address, however practically, we expect the attack to be mounted with a receiver that owns multiple IP addresses. Assume the receiver can sample the target host at Δt intervals, i.e. at a frequency of $\frac{1}{\Delta t}$. Denote by M an attack parameter (threshold) $\Delta t \cdot f \leq M \ll 2^{16}$. Ideally, M should be minimal to reduce the required bandwidth of both parties (and also to avoid wrapping around), yet should take into consideration possible noise and jitter, i.e. M should be slightly higher than $\Delta t \cdot f$. The sender chooses a list of IP addresses that will be used for spoofing. The list length L is a function of the acceptable probability p of bit read error (for bit “1”), which happens when all L addresses do not elicit an answer that uses $\beta[R]$. We choose a parameter Λ , such that $p = (1 - \frac{1}{2048})^\Lambda \approx e^{-\frac{\Lambda}{2048}} < 0.001$, i.e. $\Lambda \geq 14148$. For a receiver with a single IP address, we set $L = \Lambda$ (in our experiments, we rounded this down to $L = 14000$). Assume the sender can send B packets per second. Then transmitting one bit takes $\frac{(2M-1) \cdot L}{B}$ seconds. The attack works as follows:

- Receiver: During the transmission time of a single bit, i.e. for $\frac{(2M-1) \cdot L}{B}$ seconds, the receiver samples the target host at intervals Δt and records the IPv4 ID field of the server packets – ID_i . It then checks, for each $i > 0$, if $(ID_i - ID_{i-1}) \bmod 2^{16} \geq M + 1$. If this is met for any $i > 0$, then the bit transmitted is determined to be “1”. Otherwise the bit is determined to be “0”. We show below that it is guaranteed that the sender increments the IP ID counter by at least M for one interval, and the receiver’s first sample increments it at least by 1, hence a signalled interval should have an increment of at least $M + 1$.
- Sender: To send a bit “1”, the sender goes through the L addresses, and for each, spoofs $2M - 1$ packets originating from the address and whose destination is the target host, during a Δt time interval, over ICMP or the UDP protocol of choice. To send a bit “0”, the sender does nothing.

The sender needs to send $\frac{2M-1}{\Delta t}$ packets per second from each spoofed address, since there is no expectation to synchronize the time the sender’s packets arrive at the target host, with the time the receiver’s packets arrive at the host. All $2M - 1$ packets from the same spoofed address should be sent in rapid succession. By sending $2M - 1$ packets at an arbitrary Δt interval, the sender ensures that at least one Δt interval sampled by the receiver will have at least M sender packets in it. In other words, the sender can send a burst of $2M - 1$ packets from one spoofed address (assuming this takes less than Δt) and move on immediately to the next spoofed address.

Note that for this technique to yield correct results, no other network traffic emitted by the target host, other than to the sender and to the receiver, must be using the $\beta[R]$ counter. This is a reasonable assumption, since the probability of the host’s outbound traffic to a random address to use $\beta[R]$ is $\frac{1}{2048}$.

This covert channel works across containers. Even though the *net_key* introduces dependency in the container for which the hash is calculated, the β table itself is shared among all containers. Since the attack does not rely on having the same hash function for the receiver and the sender, it follows that the attack covers the case wherein the receiver is served by one container, and the sender is served by another container.

The covert channel also works across protocols (*protocol=1* for ICMP, 17 for UDP), by the above argument.

Ideally, the attack is carried out with a receiver that owns $n \gg 1$ IPv4 addresses. In such a case the sender needs to spoof only $L = \frac{\Delta}{n}$ IPv4 addresses (similar to the birthday paradox), and thus the attack speeds up by a factor of n . This requires the receiver to send an aggregate of $\frac{n}{\Delta t}$ packets per second to the host. The receiver needs to first ensure that each IP address it owns uses an exclusive counter. For this, the receiver can employ the techniques described in [25, Section 6] to detect collisions in pairs of IP addresses, and remove one address from each such pair, which results in a collision-free list. Also, once the receiver knows which IP addresses have counters that collide with those of the sender’s addresses, the receiver can send packets only to these IP addresses, which saves bandwidth. Moreover, for the firewall piercing attack, the sender (assuming it is on the internal network) cannot spoof too many IP addresses, since the addresses it spoofs are on the firewall’s local network⁴, and thus the firewall has to have them in its ARP cache, which is limited in size. As long as there are less than `gc_thrash1` (defaults to 128) ARP entries, garbage collection will not run [27], therefore the sender must keep the number of spoofed addresses below 128 (e.g. $L = 127$). This is feasible as explained above when the receiver can listen on $n \geq \lceil \frac{\Delta}{127} \rceil$ IP addresses.

The limiting factors in this attack are the sender’s bandwidth (vs. $M \cdot L$), and the network jitter between the receiver and the host (vs. Δt) as a jitter value close to Δt may cause the receiver’s probe packets to arrive at the host out of order.

Note that it is possible to improve this attack if an opposite direction channel exists. In such a case, the receiver and the sender can iteratively find out which (single) spoofed IP address used by the sender is effective (collides with one of the receiver’s IP addresses), and use this to significantly reduce the traffic, bandwidth, and number of addresses required for the channel.

Linux generates the IPv6 Identification header extension field in a very similar manner, and using the same table. IPv6 fragmentation attacks are out of scope for this research.

B. Windows Server IPv4 ID

1) *Information Leakage*: Suppose at time t , a receiver at IPv4 address IP_1 received a packet from a Windows Server

target host at IP_2 , with IPv4 identifier ID , and at time $t'' > t$, the receiver received a second packet from the target host, with IPv4 identifier ID'' . ID originated from a `Path` object indexed by (IP_1, IP_2) , at time t . ID'' originated from a `Path` object indexed by (IP_1, IP_2) , at time t'' . If the `Path` object indexed by (IP_1, IP_2) remained intact in the time window $[t, t'']$ then $ID'' = ID + 1 \pmod{2^{16}}$. However, if, at some time t' , $t < t' < t''$, this `Path` object was purged, then the second packet originated from a freshly created `Path` object, whose IPv4 ID is initialized with a random value, i.e. $ID'' = ID + 1 \pmod{2^{16}}$ with probability 2^{-16} . The leakage then is whether a purge occurred at time t' or not.

The purge timing is essential for the leakage to occur. One restriction is $t' - t \geq 60s$, as only `Path` objects whose last access time is over 60 seconds ago are purged. Another restriction relates to the exact purge timing of the `Path` object. As we will see below, there may be up to 20000 `Path` objects in the `PathSet` (or even a bit higher, if the target host is not completely idle). Since the purge rate is 2000/sec, it may take up to 10 seconds for the specific `Path` object to get purged. Therefore, the additional condition is $t'' - t' \geq 10s$.

From the above $t'' - t \geq 70s$, i.e. the minimal time between receiver samplings is 70 seconds. We also need take into account the fact that the flood detection triggering consumes another second, i.e. the probes have to be 71s apart. Practically, we need to pad this a bit to accommodate for `Path` objects that organically get inserted into `PathSet` during the normal system run. An additional 1s accommodates 2000 such objects.

Interestingly, due to a bug in Windows, oftentimes there is an additional purge sequence right after the first purge sequence, taking additional 5 seconds in our use case. This additional purge sequence does not interfere with our covert channel, because if it extends a few seconds beyond the next probe, that probe will not be affected, as it is not stale.

It should be noted that a leakage does not occur across containers (compartments), since each compartment maintains its own IPv4 `PathSet`.

2) *Covert Channel*: The receiver and the sender have synchronized transmission cycles. The receiver samples the IPv4 ID of the host at the beginning and at the end of every transmission cycle, by sending K UDP requests that elicit UDP answers from the host, and observing the IPv4 ID of the host packets. Multiple requests are needed to cater for packet loss. The receiver determines that the transmitted bit is “0” if the difference (modulo 65536) between the IDs is less than the number of packets sent to the host during the cycle – $2K$. Otherwise, the bit is determined to be “1”. The choice of K can be arbitrary, as long as the value is very small, with very little impact on the channel’s performance.

The sender signals a bit “0” by not doing anything during the transmission cycle. This guarantees that the `Path` object in the host, that belongs to the receiver, is not purged, and therefore the IPv4 IDs will remain sequential for the end-of-cycle probe. On the other hand, for the bit “1”, the sender needs to force that `Path` object to get purged. In the end-of-cycle probe, the receiver will be assigned a new `Path` object, with a random IPv4 ID base, which is very likely to fall outside the range for bit “0”.

⁴This restriction does not apply to other attack scenarios, since there, the sender and the target host reside on different networks.

Thus, to transmit the bit “1”, the sender needs to inflict a purge cycle on the target host. For this, the sender can trigger the `PathSet` flood detection, by forcing the target host to send packets to 10000 different IP addresses in a single second. Since the flood detection is measured in intervals of 0.5 seconds (5000 new `Path` objects per 0.5 second), this ensures that flood detection will be triggered. In the QUIC use case, the sender can achieve this by sending QUIC Initial packets (1228 bytes at the IP layer) from spoofed IP addresses to a UDP port on the target host listening to a protocol that is on top of QUIC, e.g. HTTP/3. The target host is then obliged to send a QUIC response to the IP address from which the packet allegedly originated. The required bandwidth is therefore 98.24Mb/sec, well below the limit of modern WiFi and Ethernet networks, and an achievable upload speed by fiber Internet connections.

This attack’s performance is conditional upon the sender having bandwidth of at least 10000 packets/s in the chosen protocol, and takes at least 71s per bit.

A “1” bit can be interpreted as a “0” when the new `Path` object contains an IP ID that accidentally falls into the range allowed for bit “0”. This happens with probability $2^K/65536$, and only for “1” bits. Therefore, the bit error rate is $K/65536$.

Note that this attack is not disruptive. Organic `Path` objects are indeed removed from the `PathSet`, but only if their last access time was over 60 seconds ago. This means that two IPv4 packets/fragments generated over 60 seconds apart may have the same IPv4 ID, with probability $1/65536$. This is not an issue since it is highly unlikely that the first fragment will not be reassembled at the destination for this long.

An Internet attack cannot be mounted over ICMP, because Windows only responds to ICMP Echo requests from the local network.

C. macOS Server (Connection-Less Protocols) and OpenBSD IPv4 ID

For two IPv4 ID values ID_0, ID_m sampled $m > 0$ packets apart, we have (recall that \mathcal{M} is the uniqueness interval):

$$\Pr\{ID_m = ID_0\} = \begin{cases} 0 & m \leq \mathcal{M} \\ \frac{1}{(65536+m-2\mathcal{M}-1)} & \mathcal{M} < m \leq 2\mathcal{M} \\ \approx 1/65536 & 2\mathcal{M} < m \end{cases}$$

Thus when we sample two IPv4 ID values, we can statistically distinguish between a case where $m' \geq \mathcal{M}$ packets were generated in between (in such case, $\Pr\{ID_{m'+1} = ID_0\} \geq 1/65536$) and a case where $m' < \mathcal{M}$ packets were generated ($\Pr\{ID_{m'+1} = ID_0\} = 0$). Practically, the receiver first gets K packets from the target host, as rapidly as possible. Then the sender forces the target host to send \mathcal{M} packets to signal the bit 1, and no packets to signal the bit 0, and finally the receiver gets additional K packets. K should be picked such that $K^2 \gg 65536$ and $2K < \mathcal{M}$ (ideally $K \ll \mathcal{M}$). The receiver now checks whether there are any collisions between the two sets. When a bit “0” is transmitted, we clearly have $\Pr\{C = 0\} = 1$. When a bit “1” is transmitted, the probability of *no* collisions is:

$$\Pr\{C = 0\} = \prod_{0 \leq i, j < K} (1 - \Pr\{ID_{j+\mathcal{M}+K} = ID_i\})$$

$$\leq \prod_{0 \leq i, j < K} (1 - 1/65536) \approx e^{-\frac{K^2}{65536}}$$

Since $\frac{K^2}{65536} \gg 1$ this probability is negligible.

To summarize, we described a covert channel, in which the receiver forces the target host to send back K packets, then the sender forces the target host to send back \mathcal{M} packets (or none at all), and finally the receiver forces the target host to send back additional K packets. The receiver looks for collisions in the IPv4 IDs of the two K packet sets. If collisions are found, the receiver infers that the sender interacted with the target host (bit “1”). If no collisions are found, the receiver infers the sender did not interact with the target host (bit “0”).

The limiting factors for this attack is the receiver’s bandwidth (vs. K) and the sender’s bandwidth (vs. \mathcal{M}). The jitter (assuming low values) has almost no significance here, as the order of packet arrival is unimportant neither for the sender, nor for the receiver. And since both the sender and the receiver are required to send thousands of packets, which takes a lot more time than the typical jitter, the bandwidth overshadows the jitter as a dominant factor.

See Appendix B for a stealthier variant of this attack.

In macOS 11.3 and above, an ICMP Echo rate limit is imposed, with a threshold value in the range 251-500 chosen at random once per boot, and applied to 1-2 second intervals. When the rate limit is exceeded in an interval, the ICMP Echo response is generated with probability inversely proportional to the excess ICMP Echo request packets. That is, if the random rate limit is R and $N > R$ ICMP Echo requests were already received in the interval then the next ICMP Echo request will be answered with probability $\frac{1}{N-R}$. The above attack can be mounted almost as-is by throttling it to under 125 packets/s. But a more efficient attack can simply take advantage of the rate limit to signal bits. To signal “1”, the sender sends a burst of e.g. 600 ICMP Echo requests (which guarantees an excess of at least 100 requests, thus ensuring that subsequent ICMP Echo requests in that time period has a probability $\leq \frac{1}{100}$ to be answered by macOS). The sender does nothing for “0”. The receiver then sends an ICMP Echo request to the host, which is always answered if the sender sent nothing (“0”), but is not answered (with probability > 0.99) if the sender sent 600 requests. This procedure needs to be fortified with some error correction logic (to take care of the occasional response from the server even in the case the rate limit was exceeded). Thus a single bit can theoretically be signalled in two seconds. Due to logistic reasons, we could not test this attack.

D. NetBSD TCP ISN

The TCP ISN -based global state is the number of TCP connections established since boot (modulo 256).

The covert channel works as follows: the receiver sends a TCP SYN packet to the NetBSD target host, and records the most significant 8 bits of the ISN in the target host’s SYN+ACK response – m . The sender then sends two or more TCP SYN packets to the host, for bit “1”, or no packets for bit “0”. Finally, after Δt time, the receiver sends a TCP SYN packet to the host, and records m' from the host’s SYN+ACK response. Then if $(m' - m) \bmod 256 > 1 + \lceil 2\Delta t \rceil$, there was

another connection served by the target host, i.e. the sender signalled the bit “1”, otherwise the sender signalled “0”.

The attack requires that the jitter be much smaller than Δt .

Note that unlike the well known sequential IPv4 ID attack for NetBSD, our attack also works with TCP/IPv6, and can combine TCP/IPv6 and TCP/IPv4 senders and receivers.

E. Additional Attacks

We also found additional new attacks against NetBSD, that are disruptive and/or inferior in performance (packets required) to the attacks described above. For sake of completeness, we describe these attacks in Appendix C. Note that if/when the attacks described above are addressed by NetBSD, the attacks in Appendix C may take their place as the most effective attacks (unless they are fixed as well).

F. Practical Considerations

1) *RTT*: Our attacks require some synchronization between the time of arrival of the sender packets and the receiver packets, at the target host. Interestingly, such synchronization is not affected by the RTT itself, assuming each party knows its RTT to the target host beforehand. For simplicity, let us assume a *fixed, symmetric* RTT. In order for a packet to arrive at time t to the target host, the sending party needs to send it at time $t - \frac{RTT}{2}$, assuming the party’s machine time itself is synchronized. In reality, RTT is not fixed, and the RTT variance, or jitter, determines the synchronization granularity. Likewise, in reality, $\frac{RTT}{2}$ is only an approximation to the end-to-end time between two parties – the actual value may differ somewhat due to asymmetry in routing paths. This unknown quantity can also be accumulated into the synchronization granularity.

Moreover, in our attacks, each packet sent by the sender and by the receiver is independent of the previous packets, i.e. the attacks are entirely non-adaptive. This means that neither party needs to wait for any response (to a previous packet) to arrive from the target host, before sending the next packet. Therefore, there is no significance to the RTT in timing the attacks. The only exception is some TCP/IP attack scenarios wherein one or two parties send an RST in response to the host’s SYN+ACK answer. The RST packet contains the host-generated TCP sequence number from the host’s SYN+ACK, and thus it relies on a previous packet from the host, i.e. subject to RTT. However, sending the RST packets can be done in parallel with the attack’s main packet stream, and thus the RTT does not affect the overall attack time in this case as well.

2) *Jitter and Bandwidth*: The limiting factor for the attack speed is the sender’s bandwidth (which we throttle for some attacks) and the jitter, which is typically at milliseconds granularity. In fact, due to its typically low value, jitter is a limiting factor only in Linux and NetBSD, and the sender bandwidth is a limiting factor in Linux, OpenBSD and macOS. In Windows, the attack time is dominated by the 60s time-to-live limit imposed on stale `Path` objects, and is thus unaffected by jitter and bandwidth (except that it requires the sender to be able to send 10000 packets/s bursts).

3) Synchronization between the Sender and the Receiver:

We assume that the sender and receiver machines are synchronized at 0.1s resolution. This is a fair assumption, given the accuracy of present day hardware clocks, combined with NTP which keeps the clocks from drifting [34]. Thus the sender and receiver can decide on a UTC time (in seconds) when they begin the exfiltration session. Since the exfiltration techniques are based on the receiver polling the state before and after the sender, the receiver can poll 0.1s before the bit transmission start time, and poll again 0.1s after the bit transmission end time (all times are in the target’s time frame, i.e. both sender and receiver should compensate based on their respective RTTs, by transmitting at their $t - \frac{RTT}{2}$). This compensates up to 0.1s synchronization offset between the sender and the receiver. A coarser-grained synchronization (e.g. 0.2s) can be accommodated via increasing the bit transmission time by 0.2s.

4) Interference (Organic Outbound Traffic from the Target Host):

In the firewall piercing scenario, using ICMP, the firewall can be considered idle in terms of network traffic, since the vast majority of the traffic it handles is as a router, where the firewall is not required to generate e.g. IPv4 IDs.

In general, our techniques are not sensitive to a mild level of interference. The macOS and OpenBSD techniques only require that the organic outbound traffic bandwidth (in terms of packets/s) be lower than the bandwidth generated by the sender (or more precisely – lower than the threshold that the reader interprets as a signalled bit “1”). Raising this threshold can compensate for interference. The Windows technique can easily compensate several thousands IP tuples (our experiment setup actually already has such compensation in place) by allowing a bit more time for purging them (purging 2000 tuples takes 1s). The Linux technique is affected by interference only if the specific bucket used for the communication happens to also serve organic traffic, which is very unlikely (probability $\frac{1}{2048}$ per destination IP address).

5) *Packet Loss*: Our techniques are not sensitive to a mild level of packet loss. In general, the attacks need a certain number of packets, or a certain rate of packets in order to succeed. Packet loss can be compensated by increasing the number/rate of packet sent accordingly.

6) *Stealth vs. Throughput*: We throttle the sender’s bandwidth to 100Mb/s-150Mb/s in order not to disrupt the normal network usage on a 1Gb/s LAN. Note that in all cases (except Windows) we could get an almost order of magnitude improvement in the exfiltration bit rate by consuming almost all the 1Gb/s LAN bandwidth, though this can disrupt the network functionality for other LAN users and end up in exposing the channel. On the other hand, we could throttle the sender and receiver much further, reducing the throughput, but also the probability to get flagged by network IPS/IDS. Specifically, the NetBSD TCP ISN technique only requires a few TCP connections per second, and the Linux technique’s sender bandwidth can be reduced by a factor of n if the receiver controls n IP addresses. Additionally, see Appendix B for stealthier attack variants against macOS and OpenBSD.

VI. EXPERIMENTS

A. A Setup for the Exfiltration (via DMZ) and Firewall Piercing Experiments

The exfiltration experiments have the target host and the sender on 1Gb/s Ethernet network (the *lab network*). **This network is connected to the Internet via a firewall+NAT device.** Where possible, we ran two tests per attack: one with the receiver in a transcontinental location *A*, 3300Km from the target host, and one with the receiver in a transatlantic location *B*, 9600Km from the target host.

We also measured the RTT average and standard deviation (the latter is indicative of the network jitter) between the two locations and the lab network. The average RTT for *A* is 79.9ms, with standard deviation 0.3ms, and the average RTT for *B* is 146.9ms with standard deviation 1.4ms. This is slightly better than the results reported in 2014 [37] which is expected in light of the Internet technology improvement over the last seven years.

A single exfiltration test transmits 128 bits from the sender on the internal network, to the receiver on the Internet. This message size was chosen as a benchmark since it demonstrates that AES-128 keys can be exfiltrated in a reasonable amount of time.

Where needed, we used libpcap to craft arbitrary network packets, particularly for spoofing the source IP address of packets, and for accurate timing of packets and packet bursts.

For HTTP/3 server (over QUIC and UDP), we used Caddy Server with the `experimental_http3` directive. For SSH server (over TCP), we used OpenSSH, and for HTTP/1.1 server (over TCP) we used Apache Web Server. We used Docker version 19.03.8 for the Linux containers.

B. A Setup for the Host Alias Resolution Experiments

Given two services (designated by IP address and port), host alias resolution (and de-NATing) provides a single bit of information – whether the two services reside on the same host (kernel). To implement this, the attacker runs a sender logic for bit “1” against one service, and a receiver logic against the other service. If the receiver logic reads “1”, then the two services run on the same host. If the receiver logic reads “0”, then the services run on different hosts.

To demonstrate that our results are not randomly correct, we conducted a series of 10 bit “1” transfers, for each of the following configurations:

- Target addresses on the same host
- (Where applicable) target addresses in different containers on the same host (kernel)
- Target addresses in different hosts

The expected result for the first two configurations, where the addresses are served by the same kernel, is an all-“1” readout. The last configuration is expected to result in an all-“0” readout, since the sender changes the state of one kernel, but the receiver probes another kernel. In other words, assuming correct results in the above tests, we are able to tell apart a situation in which the two target addresses are served

from the same host (kernel), and a situation in which they are served from two different hosts. This demonstrates host alias resolution, including across containers, as well as de-NATing and container co-residency detection.

We could only test host alias resolution with Linux, Windows and OpenBSD, because we only had one instance of macOS and NetBSD in our lab. We tested Linux and Windows with an attacker in the lab, because the test involves the attacker spoofing packet source IP addresses, and we categorically do not run such tests over the Internet. With OpenBSD, there is no need to spoof packet source IP addresses, and thus it was tested with an attacker in location *A*.

Other aspects of the experiment are identical to the ones described in Section VI-A.

C. Description of the Hosts and the Network

The target hosts used in the experiments are listed in Table IV. All hosts are located in our lab network, which is a 1Gb/s Ethernet network connected to the Internet over IPv4, through a firewall + NAT device.

D. Experiment Descriptions

Our results are summarized in Table V. Only the sender and receiver locations are listed – the target host is always on the lab net (except for OF-1). Column “Section” refers to the below sections for further information. An empty cell indicates the results are self explanatory.

1) *Linux Exfiltration (via DMZ) over UDP/IPv4*: In tests LE-1 and LE-2 (locations *A* and *B*, resp.), the target is an Intel x64 machine, therefore $f=250\text{Hz}$. The covert channel was defined with $\Delta t = 10\text{ms}$ (considerably lower values are not usable due to the jitter). Theoretically, we could have used $M = 3$ with this Δt , but we set $M = 6$ to eliminate the occasional timing deviation in sampling caused by jitter. Due to logistic and budgetary constraints, the receiver in our experiments used only a single IP address.

During the preliminary tests, we noticed that the target machine’s CPU was mildly loaded during the transmissions of the bit “1”. This is not expected to happen in a production system, where the server hardware is typically much more powerful than our Dell machine. However, given the Dell machine’s CPU load, one could argue that the effects we observe have to do with CPU load (either due to application processing, or due to network processing). Ironically, this points at another potential covert channel – measuring the effects of CPU load on the network traffic shape. But as we mention above, this is unlikely to happen in a production network, where the CPU power is more in balance with the network capacity.

To make sure the effect we measure is only related to the IPv4 ID information leakage, we altered the channel definition for the purpose of the experiment. Instead of staying idle (not transmitting anything) for bit “0”, in the experiment, we transmit the same amount of packets, but using a small subset of 90 spoofed source IP addresses whose counters do not collide with the receiver’s counter. A-priori, choosing 90 counters out of 2048 yields a good likelihood of this condition to be met. A-posteriori, we can verify this easily.

TABLE IV: Hosts

OS	Hardware	Protocol	Service	Server SW
Ubuntu 20.04 Linux (Linux kernel 5.8.12)	Dell E7450: Intel i7-5600U CPU, 16GB RAM, 1Gb/s Ethernet (USB)	UDP/IPv4	HTTP/3	Caddy 2.2.0
Windows 20H2 build 9042.685	Intel i7-3770 CPU, 16GB RAM, 1Gb/s Ethernet	UDP/IPv4	HTTP/3	Caddy 2.2.0
		TCP/IPv4	HTTP/1.x	Apache 2.2.25
macOS Big Sur 11.1 (xnu-7195.60.75-1)	MacBook Air Model A1466: Intel i5-5250U CPU, 8GB RAM, 1Gb/s Ethernet (Thunderbolt)	UDP/IPv4	HTTP/3	Caddy 2.2.0
OpenBSD 6.8	Dell E7450: Intel i7-5600U CPU, 16GB RAM, 1Gb/s Ethernet (USB)	UDP/IPv4	HTTP/3	Caddy 2.2.0
		TCP/IPv4	SSH	OpenSSH 8.4
NetBSD 9.1	Intel NUC7CJYH: Intel Celeron J4005 CPU, 8GB RAM, 1Gb/s Ethernet	TCP/IPv4	SSH	OpenSSH 8.0
		TCP/IPv6	SSH	OpenSSH 8.0

TABLE V: Experiments

Test	OS	Test Scenario	Protocol	Sender	Receiver	Bits/Repetitions	Success	Bit rate	Section	
LE-1	Linux	Exfiltration (via DMZ)	UDP/IPv4	Lab Net	Loc. A	128	128/128	400 b/h	VI-D1	
LE-2		Exfiltration (via DMZ)	UDP/IPv4	Lab Net	Loc. B	128	128/128	400 b/h	VI-D1	
LE-3		Exfiltration (container)	UDP/IPv4	Lab Net	Loc. A	128	128/128	400 b/h	VI-D2	
LE-4		Exfiltration (container)	UDP/IPv4	Lab Net	Loc. B	128	128/128	400 b/h	VI-D2	
LF-1		Firewall Piercing (IPfire, slow)	ICMP/IPv4	Lab Net	Lab Net	128	128/128	400 b/h	VI-D3	
LF-2		Firewall Piercing	ICMP/IPv4	Lab Net	Lab Net	128	128/128	3000 b/h	VI-D3	
LA-1a		Alias Resolution (1 host)	UDP/IPv4	Lab Net	Lab Net	10	10/10	N/A		
LA-1b		Alias Resolution (2 containers)	UDP/IPv4	Lab Net	Lab Net	10	10/10	N/A		
LA-1c		Alias Resolution (2 hosts)	UDP/IPv4	Lab Net	Lab Net	10	10/10	N/A		
WE-1		Windows	Exfiltration (via DMZ)	UDP/IPv4	Lab Net	Loc. A	128	128/128	47.4 b/h	VI-D4
WE-2	Exfiltration (via DMZ)		UDP/IPv4	Lab Net	Loc. B	128	128/128	47.4 b/h	VI-D4	
WE-3	Exfiltration (via DMZ)		TCP/IPv4	Lab Net	Loc. A	128	128/128	41.4 b/h	VI-D5	
WE-4	Exfiltration (via DMZ)		UDP/IPv4 (Sender), TCP/IPv4 (Receiver)	Lab Net	Loc. A	128	128/128	47.4 b/h		
WA-1a	Alias Resolution (1 host)		UDP/IPv4	Lab Net	Lab Net	10	10/10	N/A		
WA-1b	Alias Resolution (2 hosts)		UDP/IPv4	Lab Net	Lab Net	10	10/10	N/A		
ME-1	macOS		Exfiltration (via DMZ)	UDP/IPv4	Lab Net	Loc. A	128	128/128	1800 b/h	VI-D6
ME-2			Exfiltration (via DMZ)	UDP/IPv4	Lab Net	Loc. B	128	128/128	1800 b/h	VI-D6
OE-1	OpenBSD		Exfiltration (via DMZ)	UDP/IPv4	Lab Net	Loc. A	128	128/128	720 b/h	VI-D7
OE-2			Exfiltration (via DMZ)	UDP/IPv4	Lab Net	Loc. B	128	128/128	720 b/h	VI-D7
OE-3		Exfiltration (via DMZ)	TCP/IPv4	Lab Net	Loc. A	128	128/128	360 b/h	VI-D8	
OE-4		Exfiltration (via DMZ)	UDP/IPv4 (Sender), TCP/IPv4 (Receiver)	Lab Net	Loc. A	128	128/128	720 b/h		
OF-1		Firewall Piercing (slow)	ICMP/IPv4	Loc. B	Lab Net	128	128/128	720 b/h	VI-D9	
OF-2		Firewall Piercing (Esdenera ruleset, slow)	ICMP/IPv4	Lab Net	Lab Net	128	128/128	1200 b/h	VI-D9	
OF-3		Firewall Piercing	ICMP/IPv4	Lab Net	Lab Net	128	128/128	4200 b/h	VI-D9	
OA-1a		Alias Resolution 1 host	UDP/IPv4	Loc. A	Loc. A	10	10/10	N/A		
OA-1b	Alias Resolution 2 hosts	UDP/IPv4	Loc. A	Loc. A	10	10/10	N/A			
NE-1	NetBSD	Exfiltration (via DMZ)	TCP/IPv4	Lab Net	Loc. A	128	128/128	7200 b/h	VI-D10	
NE-2		Exfiltration (via DMZ)	TCP/IPv4	Lab Net	Loc. B	128	128/128	7200 b/h	VI-D10	
NE-3		Exfiltration (via DMZ)	TCP/IPv6	Lab Net	Lab Net	128	128/128	7200 b/h	VI-D11	
NE-4		Exfiltration (via DMZ)	TCP/IPv6 (Sender), TCP/IPv4 (Receiver)	Lab Net	Loc. A	128	128/128	7200 b/h		

If the receiver’s IP address IP ID counter does collide with one of the 90 counters, the vector of bits the receiver reads will be all 1’s. We verified this is not the case. This revised experiment ensures that the host’s CPU load is comparable for both “0” and “1” bits.

Given the above results, we estimate that the bit error rate in this channel is below 0.004. In this case, simple error correction code can be used to correct any sparse errors.

2) *Linux Exfiltration via Containers over UDP/IPv4*: The LE-3 and LE-4 tests (locations *A* and *B* resp.) are very similar to the previous section, with two co-resident containers running on a single host (kernel). The sender resides in the internal (isolated) network, communicating with the internal (isolated) container over HTTP/3, and the receiver resides on the Internet, communicating with the second container over HTTP/3. In all other aspects, these tests are identical to the one in the previous section.

3) *Linux Firewall Piercing over ICMP/IPv4*: Due to technical reasons, we could only conduct the Linux firewall piercing attacks in the lab. For test LF-1, we installed a typical Linux-based firewall image (IPfire 2.25 (x86_64) running Linux kernel 4.14.112) on a dedicated device, and configured it to block communications between the internal “green” network and the external “red” network (an Internet-facing lab network). We then ran our attack, but for simplicity of testing, we sent data from the red network to the green network. The data transfer rate was limited by the sender machine’s effective bandwidth – we could only send up to $\approx 14,000$ ICMP packets/s. A stronger sender machine would have achieved a higher throughput.

In order to compare the speed of ICMP exfiltration to that of QUIC/UDP exfiltration, we conducted another experiment (LF-2), this time using the same target machine as in VI-D1. We used a stronger sender that was able to send $\approx 130,000$ ICMP packets/s, and the throughput was 1.2 bits/s (3000 b/h), which is 7.5 times faster than the QUIC/UDP variant. Here too,

we could use a stronger sender to get even higher exfiltration throughput.

4) *Windows Server Exfiltration (via DMZ) over UDP/IPv4*: We did not have at our disposal a Windows Server on a network that allows spoofing, so instead, for tests WE-1 and WE-2 (locations *A* and *B*, resp.), we simulated Windows Server on a Windows 10 platform. This was done by setting `TcpipIsServerSKU` to 1 and the default compartment's `IPv4 PathSet` purge threshold (32 bit quantity at offset 332 of `PathSet`) to 32768. This made the system behave identically to Windows Server for the purpose of our experiments.

For the channel definition, we used $K = 6$. A single bit transfer took 76s (theoretically this could be done with 71s per bit, but due to technical issues, and to make sure the channel can handle some organic network traffic, we padded it by 5s).

5) *Windows Exfiltration (via DMZ) over TCP/IPv4*: In experiment WE-3, we demonstrated that the Windows covert channel can be used over TCP. In essence, this is a very similar experiment to the one described in Section VI-D4, but this time we used TCP. The sender sent SYN packets with spoofed source IP address to the target host. These packets are short (40 bytes at the IPv4 level) and can be transmitted rapidly. Assuming the spoofed IP addresses are not associated with any machines, there will not be any transmission from these IP addresses following the spoofed SYN packets. Windows responds immediately with a SYN+ACK packet, with additional two retransmissions (the default “Max SYN Retransmissions” is 2) at exponential backoff, so the first retransmission is sent after an “Initial Retransmission Time-Out” seconds, henceforth *IRTO* (in Windows Server, by default $IRTO=3s$), and the second retransmission packet is sent $2 \cdot IRTO$ after the first retransmission. These packets are not answered of course, so finally Windows sends an RST packet, $4 \cdot IRTO$ seconds after the third SYN+ACK packet is sent ($7 \cdot IRTO$ seconds after the SYN packet is received). Therefore, Windows sends the last packet to a spoofed IP address 21s after a SYN is received from it. This overlaps with the 10 seconds the receiver needs to wait after the sender's burst in order for the purge to be concluded, so only 11 seconds are actually added to the time-per-bit compared to the UDP case. In a lab experiment, we found that Windows' SYN cache is not exhausted even by 13000 half-open connections, thus our technique for TCP/IPv4 in Windows is non-disruptive.

As for the receiver, we implement it to send an RST packet to the target host as soon as the target host's SYN+ACK arrives at the receiver, thus the `Path` object's last access time is only extended by the RTT between the target host and the receiver, having a negligible impact on the attack timing.

We carried out the experiment similar to Section VI-D4, with the modified attack timing. Since we used Windows 10 as a target host, we had to adjust its `Initial RTO` value to 3000ms to simulate Windows Server (the default Windows 10 value is 300ms).

6) *macOS Exfiltration (via DMZ) over UDP/IPv4*: In tests ME-1 and ME-2 (locations *A* and *B*, resp.), the channel was defined with $K = 700$, thus ensuring that the probability for a sent bit “1” to be received as “0” is very low ($1/1767$). The sender transmitted 5000 packets per bit. Theoretically, the sender needed to send 4096 packets to signal the bit “1”, but

in practice, due to server responsiveness issues, we had the sender send some extra packets. With throttling, this took 0.39 seconds. Due to some overheads incurred by our specific test implementation, each bit transmission took two seconds.

7) *OpenBSD Exfiltration (via DMZ) over UDP/IPv4*: In tests OE-1 and OE-2 (locations *A* and *B*, resp.), the channel was defined with $K = 700$, thus ensuring that the probability for a sent bit “1” to be received as “0” is very low ($1/1767$). The sender transmitted 40000 packets per bit. Theoretically, the sender needed to send 32768 packets to signal the bit “1”, but in practice, due to server responsiveness issues, we had the sender send 40000 packets. This took 3.32 seconds. Due to some overheads incurred by our specific test, each bit transmission took 5 seconds.

8) *OpenBSD Exfiltration (via DMZ) over TCP/IPv4*: In test OE-3, the sender sent SYN packets to the target host. Both the sender and the receiver were implemented to send an RST packet to the target host as soon as the target host's SYN+ACK arrives at the sender/receiver, to ensure no retransmissions clutter the network and that the host's SYN cache is not overflowed, since the pending connection is removed from it once the RST packet for it arrives.

We throttled the sender to about 6000 packets/sec in order to make sure the host's TCP SYN cache is not exhausted (the OpenBSD SYN cache size limit is 10255 entries). Due to the throttling, the time per bit was extended to 10s.

9) *OpenBSD Firewall Piercing*: Test OF-1 was conducted over the Internet, with a very weak machine running OpenBSD 6.7 in Location *B*. That machine could only handle $\approx 9,000$ ICMP Echo Requests per second, which set the constraint on the throughput. We achieved a single bit transfer in 5s.

We also ran in-the-lab tests with firewalls (we could not deploy Internet-facing firewalls). In test OF-2, we simulated an OpenBSD-based firewall. For this, we used a PF ruleset we extracted from an Esdenara Firewall-3 installation [33]. This ruleset explicitly allows ICMP traffic to and from the firewall, on all interfaces. We tested OpenBSD 6.9 in this case. In test OF-3, we measured the speed of the attack with a sender capable of sustaining $\approx 68,000$ ICMP Echo Requests per second, against the usual Dell target running OpenBSD 6.8, which can serve $\approx 260,000$ Echo Requests per second. We obtained a 0.85s bit transfer time, but using a stronger sender machine can significantly increase this figure.

10) *NetBSD Exfiltration (via DMZ) over TCP/IPv4*: Our covert channel for NetBSD over IPv4 is not in itself interesting, because NetBSD already has a TCP/IPv4 covert channel in the form of its global IPv4 ID counter. However, we wanted to emphasize the applicability of our TCP/IPv4 technique over the Internet since we cannot demonstrate our TCP/IPv6 technique over the Internet due to technical constraints. By demonstrating our TCP/IPv4 covert channel over the Internet, we infer that the same holds for our TCP/IPv6 channel.

In our NE-1 and NE-2 experiments (with senders at *A* and *B*, resp.), we set $\Delta t=0.2s$, with total time per bit 0.5s (this is due to some limitations of our testing code, and we note that theoretically NetBSD's covert channel can be run at much higher speeds), thus theoretically threshold for $m - m'$ should

have been 2, but we set it to 3 to avoid false positives from occasional TCP connections established with the target host.

11) *NetBSD Exfiltration (via DMZ) over TCP/IPv6*: In test NE-3 we had to use a local receiver since our lab does not have IPv6 connectivity to the Internet.

VII. REMEDIATION AND SUGGESTIONS

We begin with the easy parts, in which the attack can be completely eliminated:

- TCP/IPv4 packets, and IPv4 packets with DF=1 – these packets can have ID=0 or a completely random ID. In Linux, all TCP packets, and UDP packets whose length is below the PMTU are sent with DF=1. In Windows and macOS, all TCP packets are sent with DF=1. Moreover, in macOS, the TCP packets have ID=0, which demonstrates that simply setting ID=0 suffices for these cases. Note that e.g. QUIC *requires* setting DF=1 in UDP [21, Section 14], but Caddy Server leaves the choice of DF to the operating system.
- TCP ISN – this field should not contain state information, i.e. should contain random data and client information only.
- IPv6 Flow label (see Appendix C1 and Appendix C2) – can be chosen at random (i.e. using a “regular” PRNG). There is no RFC requirement for the flow label to have a unique value in any time duration.

In some operating systems, a built-in filtering functionality facilitates overriding a protocol field value with a fixed/random value.

The challenging case is UDP/IPv4 (and ICMPv4/IPv4) ID with DF=0. In this case, there is a real need for the ID field to be non-repeating to some extent. We can make the following suggestions to reduce the attack surface and eliminate some attack variants, but the underlying issue is still in effect even with our suggestions:

- Maintain a separate IPv4 ID state according to the destination address: one state for RFC1918 destination addresses (internal networks) and another state for external destinations. This eliminates the exfiltration attacks where the sender is on an internal network and can only spoof internal network addresses.
- Maintain an IPv4 ID generator state per service (i.e. per a listening IP address and port). This eliminates the host alias resolution attacks. A weaker version is to maintain a state per listening IP address.
- Maintain an IPv4 ID generator state per container – this eliminates the cross-container attacks.

Another approach is to ensure that connection-less protocols emit IPv4 packets whose length is below a reasonably chosen minimal PMTU, and simply set DF=1 and ID=0 for *all* these packets. For example, DNS over UDP/IPv4 can function with message length limit of 512 bytes (540 bytes including UDP and IPv4 headers), QUIC over UDP/IPv4 can function with message length limit of 1200 bytes (1228 bytes including

UDP and IPv4 headers), and ICMP Echo Requests whose size exceeds PMTU can be discarded.

Finally, it is possible to simply randomize the IP ID, which guarantees maximum security, at the expense of somewhat increased malformed/dropped packet rate (due to the possible ID collisions and their effect on fragment reassembly).

Note that ICMP rate limits do not eliminate the covert channels: even an aggressive per-destination rate limit is ineffective against the attack, as the attacker may use multiple, possible spoofed, source IP addresses. A global rate limit, on the other hand, facilitates a covert channel via overflowing (or not) the packet counter.

VIII. CONCLUSIONS

We conclude that global states in network protocols of server hosts can be exploited in various ways to weaken the security of networks where such hosts reside, and to glean information on the implementation of services which is otherwise abstracted from their consumers.

While a global state can be as obvious as an incremental counter, we also demonstrated that it can take a more subtle form. For example in the case of OpenBSD and macOS IPv4 ID, the observation of collisions in IPv4 ID values (or lack thereof) is an indication of the internal global state, which is not directly exposed to the attacker. Other such non trivial examples are the Linux IPv4 ID generation scheme, and the Windows IPv4 ID generation scheme.

Our work focused on *firewalled* hosts, which is a very common scenario for Internet-facing servers. A stateful firewall eliminates many known information leakage attacks, and therefore poses a serious obstacle for techniques that rely on information leakage via a global protocol state. We described three new attack scenarios based on information leakage via a global protocol state: firewall piercing, exfiltration from a secure network, and exfiltration across containers. We demonstrated our three new scenarios, as well as a host alias resolution scenario and a host alias resolution across containers scenario. All these scenarios were demonstrated with the sender (and the host, where applicable) behind a firewall, in realistic conditions, and over the Internet where possible. We also demonstrated some cross protocol attacks, e.g. where the sender uses UDP/IPv4 and the receiver uses TCP/IPv4, and where the sender uses TCP/IPv6 and the receiver uses TCP/IPv4. We measured the exfiltration bandwidth achieved, and found it to be sufficient for e.g. key material exfiltration in hours or minutes.

The problem of *secure* IPv4 ID generation seems to be theoretically intractable, due to the conflicting requirements – generating long sequences of unique values, at least per (IP_{SRC}, IP_{DST}) and leaking no information, while maintaining a “reasonably sized” state, which is much smaller than the number of expected concurrent tuples. Intuitively, this is due to the pigeonhole principle, as it is impossible to maintain an independent state per tuple with less bits than the number of concurrent tuples. However, we do list some recommendations that considerably reduce the attack surface and specifically address the attacks scenarios we described.

It is important to understand that information leakage from network protocol global states can affect the security of

networks and systems in surprising and unexpected ways. In our research, we highlight three novel attack scenarios. It is quite possible that additional attack scenarios exist, and thus there is room for future research on this topic.

IX. VENDOR STATUS

All vendors were notified on January 27th, 2021. **Linux** issued a patch [12] for kernel version 5.13-rc1 that significantly increases the size of the β table. This patch was back-ported to versions 5.12.4, 5.11.21, LTS branches 5.10.37, 5.4.119, 4.19.196, 4.14.238, 4.9.274 and 4.4.274. Back-ports to all 4.x branches were provided by us based on the back-port to the 5.x versions. Linux also issued a patch replacing the existing IPv6 ID algorithm (similar to the IPv4 ID algorithm) with a completely randomized IPv6 ID [42] which is included in version 5.14-rc1, and was back-ported to kernel versions 5.13.3, 5.12.18, 5.10.51, 5.4.133, 4.19.198, 4.14.240, 4.9.276 and 4.4.276. **Microsoft** informed us on March 15th, 2021, that “[Microsoft] determined your finding does not meet our bar for servicing”. **Apple** informed us on November 8th, 2021 that “We [Apple] are planning to address this issue in a security update by the end of 2021”. The ICMP rate limit was reported to Apple separately on June 11th, 2021, and regarding it, Apple informed us on July 21st, 2021, that “while we [Apple] do not see any security implications, we have forwarded it on to the appropriate team to investigate for potential future enhancements”. **OpenBSD** has not responded to multiple requests for status, has not informed us about any decision/resolution, and has not fixed the issue in their public source code repository. **NetBSD** included a patch [32] that randomizes all the affected protocol fields in NetBSD 9.2.

ACKNOWLEDGMENT

We would like to thank Benny Pinkas, the anonymous NDSS 2022 reviewers and Haixin Duan (our NDSS shepherd) for their valuable feedback, and Ehud Porat for his help in reverse engineering some parts of Windows code.

REFERENCES

- [1] M. Aiello, A. Merlo, and G. Papaleo, “Performance assessment and analysis of dns tunneling tools,” *Logic Journal of IGPL*, vol. 21, 03 2012.
- [2] M. Albrecht, “Udp idle scanning,” Oct. 2019. [Online]. Available: <https://martinralbrecht.wordpress.com/2019/10/25/udp-idle-scanning/>
- [3] G. Alexander, A. M. Espinoza, and J. R. Crandall, “Detecting tcp/ip connections via ipid hash collisions,” *Proceedings on Privacy Enhancing Technologies*, vol. 2019, no. 4, pp. 311 – 328, 01 Oct. 2019. [Online]. Available: <https://content.sciendo.com/view/journals/popets/2019/4/article-p311.xml>
- [4] S. M. Bellovin, “A Technique for Counting Natted Hosts,” in *Proceedings of the 2Nd ACM SIGCOMM Workshop on Internet Measurment*, ser. IMW ’02. New York, NY, USA: ACM, 2002, pp. 267–272. [Online]. Available: <http://doi.acm.org/10.1145/637201.637243>
- [5] R. Beverly, W. Brinkmeyer, M. Luckie, and J. P. Rohrer, “Ipv6 alias resolution via induced fragmentation,” in *Passive and Active Measurement*, M. Roughan and R. Chang, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 155–165.
- [6] C. Brubaker, A. Houmansadr, and V. Shmatikov, “Cloudtransport: Using cloud storage for censorship-resistant networking,” in *Privacy Enhancing Technologies*, ser. Lecture Notes in Computer Science, E. D. Cristofaro and S. J. Murdoch, Eds., vol. 8555. Springer, 2014, pp. 1–20. [Online]. Available: <http://dblp.uni-trier.de/db/conf/pet/pets2014.html#BrubakerHS14>
- [7] E. Bursztein, “Time has something to tell us about network address translation,” in *Proc. 12th Nordic Workshop on Secure IT Systems (NORSEC 2007)*, Úlfar Erlingsson and A. Sabelfeld, Eds. Iceland: University of Reykjavik, Oct. 2007.
- [8] Y. Cao, Z. Qian, Z. Wang, T. Dao, S. V. Krishnamurthy, and L. M. Marvel, “Off-path tcp exploits of the challenge ack global rate limit,” *IEEE/ACM Transactions on Networking*, vol. 26, no. 2, pp. 765–778, 2018.
- [9] Y. Cao, Z. Wang, Z. Qian, C. Song, S. V. Krishnamurthy, and P. Yu, “Principled unearthing of tcp side channel vulnerabilities,” in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS ’19. New York, NY, USA: ACM, 2019, pp. 211–224. [Online]. Available: <http://doi.acm.org/10.1145/3319535.3354250>
- [10] W. Chen, Y. Huang, B. F. Ribeiro, K. Suh, H. Zhang, E. de Souza e Silva, J. Kurose, and D. Towsley, “Exploiting the ipid field to infer network path and end-system characteristics,” in *Passive and Active Network Measurement*, C. Dovrolis, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2005, pp. 108–120.
- [11] G. Danezis, “Covert communications despite traffic data retention,” in *Security Protocols XVI - 16th International Workshop, Cambridge, UK, April 16-18, 2008. Revised Selected Papers*, ser. Lecture Notes in Computer Science, B. Christianson, J. A. Malcolm, V. Matyas, and M. Roe, Eds., vol. 6615. Springer, 2008, pp. 198–214. [Online]. Available: https://doi.org/10.1007/978-3-642-22137-8_27
- [12] E. Dumazet, “inet: use bigger hash table for IP ID generation,” <https://git.kernel.org/pub/scm/linux/kernel/git/stable/linux.git/commit/net/ipv4/route.c?id=aa6dd21e4b1dde9d5dc25d699d35f789ae7eeba>, Mar. 2021.
- [13] R. Ensafi, J. C. Park, D. Kapur, and J. R. Crandall, “Idle port scanning and non-interference analysis of network protocol stacks using model checking,” in *Proceedings of the 19th USENIX Conference on Security*, ser. USENIX Security’10. USA: USENIX Association, 2010, pp. 257–272.
- [14] X. Feng, C. Fu, Q. Li, K. Sun, and K. Xu, “Off-path tcp exploits of the mixed ipid assignment,” in *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS ’20. New York, NY, USA: Association for Computing Machinery, 2020, p. 1323–1335. [Online]. Available: <https://doi.org/10.1145/3372297.3417884>
- [15] X. Gao, Z. Gu, M. Kayaalp, D. Pendarakis, and H. Wang, “Containerleaks: Emerging security threats of information leakages in container clouds,” in *2017 47th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. USA: IEEE Computer Society, 2017, pp. 237–248.
- [16] Y. Gilad and A. Herzberg, “Fragmentation Considered Vulnerable,” *ACM Trans. Inf. Syst. Secur.*, vol. 15, no. 4, pp. 16:1–16:31, Apr. 2013. [Online]. Available: <http://doi.acm.org/10.1145/2445566.2445568>
- [17] M. Guri, A. Kachlon, O. Hasson, G. Kedma, Y. Mirsky, and Y. Elovici, “Gsmem: Data exfiltration from air-gapped computers over GSM frequencies,” in *24th USENIX Security Symposium (USENIX Security 15)*. Washington, D.C.: USENIX Association, Aug. 2015, pp. 849–864. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity15/technical-sessions/presentation/guri>
- [18] M. Guri, Y. Solewicz, A. Daidakulov, and Y. Elovici, “Acoustic data exfiltration from speakerless air-gapped computers via covert hard-drive noise (‘diskfiltration’),” in *Computer Security – ESORICS 2017*, S. N. Foley, D. Gollmann, and E. Snekkenes, Eds. Cham: Springer International Publishing, 2017, pp. 98–115.
- [19] M. Guri, B. Zadov, A. Daidakulov, and Y. Elovici, “xled: Covert data exfiltration from air-gapped networks via switch and router leds,” in *2018 16th Annual Conference on Privacy, Security and Trust (PST)*, 2018, pp. 1–12.
- [20] V. Hailperin, “Misusing tcp timestamps,” Mar. 2015. [Online]. Available: <https://www.scip.ch/en/?labs.20150305>
- [21] J. Iyengar and M. Thomson, “QUIC: A UDP-Based Multiplexed and Secure Transport,” RFC 9000, May 2021. [Online]. Available: <https://rfc-editor.org/rfc/rfc9000.txt>
- [22] D. Kaminsky, “Black ops of dns – the black hat briefings 2004,” 2004. [Online]. Available: <https://www.blackhat.com/presentations/bh-usa-04/bh-us-04-kaminsky/bh-us-04-kaminsky.ppt>

- [23] A. Klein, "OpenBSD DNS Cache Poisoning and Multiple O/S Predictable IP ID Vulnerability," https://dl.packetstormsecurity.net/papers/attack/OpenBSD_DNS_Cache_Poisoning_and_Multiple_OS_Predictable_IP_ID_Vulnerability.pdf, Feb. 2008.
- [24] —, "Cross layer attacks and how to use them (for dns cache poisoning, device tracking and more)," in *2021 IEEE Symposium on Security and Privacy (SP)*. Los Alamitos, CA, USA: IEEE Computer Society, May 2021, pp. 927–944. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/SP40001.2021.00054>
- [25] A. Klein and B. Pinkas, "From IP ID to Device ID and KASLR Bypass," in *28th USENIX Security Symposium (USENIX Security 19)*. Santa Clara, CA: USENIX Association, Aug. 2019, pp. 1063–1080. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity19/presentation/klein>
- [26] T. Kohno, A. Brodido, and k. Claffy, "Remote physical device fingerprinting," *IEEE Transactions on Dependable and Secure Computing*, vol. 2, no. 2, pp. 93–108, May 2005.
- [27] Linux, "arp(7) — linux manual page," <https://man7.org/linux/man-pages/man7/arp.7.html>.
- [28] M. Lucas, A. Singh, C. Cantrell, and A. Henmi, *Firewall Policies and VPN Configurations*. Rockland, MA: Syngress Publishing, 2006.
- [29] M. Luckie, R. Beverly, R. Koga, K. Keys, J. A. Kroll, and k. claffy, "Network hygiene, incentives, and regulation: Deployment of source address validation in the internet," in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 465–480. [Online]. Available: <https://doi.org/10.1145/3319535.3354232>
- [30] X. Luo, E. W. Chan, and R. K. C. Chang, "Crafting web counters into covert channels," in *New Approaches for Security, Privacy and Trust in Complex Environments*, H. Venter, M. Eloff, L. Labuschagne, J. Eloff, and R. von Solms, Eds. Boston, MA: Springer US, 2007, pp. 337–348.
- [31] S. J. Murdoch, "Hot or not: Revealing hidden services by their clock skew," in *Proceedings of the 13th ACM Conference on Computer and Communications Security*, ser. CCS '06. New York, NY, USA: Association for Computing Machinery, 2006, p. 27–36. [Online]. Available: <https://doi.org/10.1145/1180405.1180410>
- [32] NetBSD, "NetBSD Security Advisory 2021-001: Predictable ID disclosures in IPv4 and IPv6," <http://ftp.netbsd.org/pub/NetBSD/security/advisories/NetBSD-SA2021-001.txt.asc>, Mar. 2021.
- [33] E. Networks, "Esdenera firewall 3," <https://www.esdenera.com/firewall-3/>.
- [34] ntp.org, "Ntp faq – how accurate will my clock be?" 2003. [Online]. Available: <http://www.ntp.org/ntpfaq/NTP-s-algo.htm#Q-ACCURATE-CLOCK>
- [35] L. Orevi, A. Herzberg, H. Zlatokrilov, and D. Sigron, "DNS-DNS: DNS-based De-NAT Scheme," Jan. 2017. [Online]. Available: https://www.researchgate.net/publication/320322146_DNS-DNS_DNS-based_De-NAT_Scheme
- [36] I. Reed and G. Solomon, "Polynomial codes over certain finite fields," *Journal of The Society for Industrial and Applied Mathematics*, vol. 8, pp. 300–304, 1960.
- [37] SamKnows, "Quality of broadband services in the eu," European Commission, Tech. Rep., Oct. 2014. [Online]. Available: http://ec.europa.eu/newsroom/dae/document.cfm?action=display&doc_id=10816
- [38] S. Sanfilippo, "new tcp scan method," <http://seclists.org/bugtraq/1998/Dec/79>, Dec. 1998.
- [39] S. Shringarputale, P. McDaniel, K. Butler, and T. La Porta, "Co-residency attacks on containers are real," in *Proceedings of the 2020 ACM SIGSAC Conference on Cloud Computing Security Workshop*, ser. CCSW'20. New York, NY, USA: Association for Computing Machinery, 2020, p. 53–66. [Online]. Available: <https://doi.org/10.1145/3411495.3421357>
- [40] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson, "Measuring isp topologies with rocketfuel," *IEEE/ACM Trans. Netw.*, vol. 12, no. 1, p. 2–16, Feb. 2004. [Online]. Available: <https://doi.org/10.1109/TNET.2003.822655>
- [41] Y. Tao, G. Hu, B. Hou, Z. Cai, J. Xia, and C. C. Fong, "An alias resolution method based on delay sequence analysis," *Computers, Materials & Continua*, vol. 63, no. 3, pp. 1433–1443, 2020. [Online]. Available: <http://www.techscience.com/cmc/v63n3/38885>
- [42] W. Tarreau, "ipv6: use prandom_u32() for id generation," May 2021. [Online]. Available: <https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/commit/?id=62f20e068ccc50d6ab66fdb72ba90da2b9418c99>
- [43] D. J. D. Touch, "Updated Specification of the IPv4 ID Field," RFC 6864, Feb. 2013. [Online]. Available: <https://rfc-editor.org/rfc/rfc6864.txt>
- [44] Wikipedia contributors, "Stateful firewall — Wikipedia, the free encyclopedia," 2020, [Online; accessed 25-November-2020]. [Online]. Available: https://en.wikipedia.org/w/index.php?title=Stateful_firewall&oldid=987652216
- [45] A. Zarras, "Leveraging internet services to evade censorship," in *Proceedings of the 19th Information Security Conference (ISC)*, 2016.
- [46] X. Zhang, J. Knockel, and J. R. Crandall, "ONIS: Inferring TCP/IP-based Trust Relationships Completely Off-Path." in *INFOCOM*. Los Alamitos, CA, USA: IEEE, 2018, pp. 2069–2077. [Online]. Available: <http://dblp.uni-trier.de/db/conf/infocom/infocom2018.html#ZhangKC18>
- [47] X. Zheng, C. Lu, J. Peng, Q. Yang, D. Zhou, B. Liu, K. Man, S. Hao, H. Duan, and Z. Qian, "Poison over troubled forwarders: A cache poisoning attack targeting DNS forwarding devices," in *29th USENIX Security Symposium (USENIX Security 20)*. USENIX Association, Aug. 2020, pp. 577–593. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity20/presentation/zheng>

APPENDIX

A. Experiment: A Survey of Responses from Widely-Used Internet Servers to UDP, TCP and ICMP Probing

We estimate the portion of widely-used Internet hosts which do not respond to a UDP packet to a closed port, i.e. hosts that do not send back an ICMP message of any type (or any non-ICMP message) for such invalid packets. We also estimate the portion of widely-used Internet hosts which silently drop TCP SYN packets to closed TCP ports. Finally, we estimate the portion of widely-used Internet hosts that respond to ICMP Echo Request (via the Ping utility).

We use the top 1000 domains in the Alexa top 1M domain list obtained on December 13th, 2020 as a starting point. We then query for the DNS NS (name server) record-set for each domain, take the first NS record from the answer, and resolve its IPv4 address (by sending a DNS A query and take the first record from the answer). This yields an IP address of an authoritative DNS server for the domain at hand. This host answers UDP queries to port 53 (DNS). As such, it serves as a model for a reliable host that runs a UDP-based service, similar to a QUIC server. Moreover, since the host is an authoritative DNS server for a popular domain, we expect it to be highly available, i.e. we can assume there are no downtime/service issues for this host. We eliminated the duplicates in the IP address list, and ended up in a list of 748 unique IP addresses.

UDP Closed Ports: For each IP address, we sent a UDP packet for three randomly chosen high ports. These UDP ports are un-assigned by IANA, and have no popular service associated with them, hence are very likely to be closed. We observe whether we receive any response from the host, either over UDP or over ICMP. As expected, the only responses we received were ICMP Type 3 messages ("Destination Unreachable"), mostly with Code 3 ("Port Unreachable"), but in three cases, we received Code 10 ("Communication with Destination Host is Administratively Prohibited") and in two cases we received Code 13 ("Communication Administratively Prohibited").

TCP Closed Ports: We sent three TCP SYN packets to each port used above, and recorded whether we received any response back. As expected, most servers that responded, did so with an RST packet, although one server responded with SYN+ACK on two ports, which is likely a deceptive practice. Regardless, we count this server as responsive to TCP probes for closed ports.

ICMP Echo Requests: We sent each server three ICMP Echo Request messages, and we recorded whether we received any ICMP Echo Reply messages back.

The results are as follows: 116 (15.5%) servers responded to the UDP probing, and 59 (7.9%) servers responded to the TCP probing. The overlap was significant – 42 servers responded to both UDP and TCP probing. Overall, 133 (17.8%) servers responded to either UDP or TCP. 667 (89.2%) servers replied to ICMP Echo Requests. This set almost completely covers the set of servers that responded to either UDP or TCP probing – the overlap was 125 servers. In other words, the ICMP probing is almost a complete superset of the TCP or UDP probing – it misses about 1% of the space covered by UDP or TCP, and contributes additional 72.5% of coverage.

B. A Stealthier IPv4 ID Attack for macOS and OpenBSD

For macOS and OpenBSD, the IPv4 ID attack can be made stealthier by having the sender and receiver run “native” QUIC sessions with the target host. For example, the sender can request a 4.9152MB (macOS) or 39.3216MB (OpenBSD) file (or larger) from the target host. This ensures that at least 4096 (macOS) or 32768 (OpenBSD) QUIC packets are sent from the target host to the sender (as each QUIC packet carries at most 1200 payload bytes). Likewise the receiver can request a 840KB file (or larger) file from the target host to force the host to send the receiver at least 700 QUIC packets. Of course, instead of requesting one file (or resource), the sender and the receiver can request multiple resources so that the total number of packets sent by the host is above the respective thresholds.

This can also be applied to OpenBSD with TCP traffic. In this case, there is an additional benefit, which is that there is no pressure on the TCP SYN cache in the target host.

C. Additional Information Leakage Attacks

1) *A Covert Channel via NetBSD TCP/IPv6 Flow Label – Cryptanalysis:* NetBSD’s TCP/IPv6 flow label field (20 bits) is generated by the A2/20 PRNG (following the naming convention introduced in [23]). The flow label field comprises a most significant bit (MSB) part, and 19 least significant bits part. The PRNG is reseeded every 180 seconds or 200000 internal PRNG steps (whichever comes first). As part of the reseeding procedure, the parameters for generating the least significant 19 bits are sampled from a strong kernel PRNG, and the MSB is flipped. The cryptanalysis itself is too long to be included in this paper, and will be provided in an extended version. It results in the ability, given ≈ 100 consecutive flow label values, to predict the next internal PRNG steps.

This cryptanalytic result can be employed in order to implement a covert channel between a sender and a receiver. The receiver first establishes around 100 consecutive TCP connections with the target host, extracts the IPv6 flow labels,

and obtains the next PRNG internal outputs using the cryptanalytic techniques. The sender then either establishes four TCP connections with the target host, to signal the bit “1”, or does nothing, to signal the bit “0”. Finally, the receiver establishes another TCP connection with the target host, and extracts the final IPv6 flow label. The receiver then matches this value to the series of eight possible future values (corresponding to the PRNG internal steps). If a match is found in positions 2-8 of the list, then the receiver infers that the sender did not establish four TCP connections with the target host (because each TCP connection incurs at least two internal steps, and the receiver’s final connection also consumes at least two internal steps – altogether 10 internal steps at the minimum). This is interpreted by the receiver as bit “0”. Otherwise, the receiver assumes the bit sent was “1”.

For host alias resolution use case, the attacker can improve the basic technique a bit, since the attacker can observe the flow label received from both endpoints, and match them to the same series (for endpoints that are mapped to the same host) or not (for endpoints that are mapped to different hosts).

2) *A Covert Channel via NetBSD TCP/IPv6 Flow Label Most Significant Bit:* We now describe an attack against the NetBSD TCP/IPv6 flow label generation, which does not make any use of the specific algorithm that generates the least significant 19 bits of the flow label field. This is a more generic attack, compared to Appendix C1, but at the price of being much less efficient, in terms of number of packets (TCP connections) required.

The attack is based on the observation that a “passive” side channel exists in the flow label generator, in the form of the timing of the MSB flip. When two TCP/IPv6 services are served from a single NetBSD host, the MSB in their flow label fields will flip simultaneously. But when the services are served from different NetBSD hosts, this is unlikely to happen.

We now calculate how many TCP connections are needed to force a reseed. In general, let S be the internal steps the flow label PRNG executed while generating n flow labels, then S is a sum of $2n$ IIDs whose distribution is uniform on $\{1, 2, 3, 4\}$. Therefore, $S \sim N(5n, \frac{\sqrt{10n}}{2})$, and $Prob(S \geq 200000) = \Phi(\frac{5n-200000}{\sqrt{10n}/2})$. We can thus force an immediate algorithm reseed and MSB flip with probability ≈ 1 by establishing $n = 41000$ TCP connections with the target host.

This can be turned into an “active” covert channel. The sender and receiver agree on time t_0 . Shortly before t_0 , the sender establishes 41000 TCP connections with the target host and forces an algorithm reseed. Then every Δt seconds the sender either establishes 41000 TCP connections with the target host to send the bit “1”, or does nothing to send the bit “0”. Every Δt seconds the receiver establishes a single TCP connection with the target host, and observes whether the MSB bit flipped (signalling the bit “1”) or not (“0”). Δt should be chosen such that the sender can establish 41000 TCP connections with the target host during this time period. Special care must be taken if 180 seconds elapse since the last transmission of the bit “1”, as this will cause the bit to flip due to a 180s timeout, and may thus result in a false positive (reading bit “1” by the receiver). To counteract this, the sender and receiver can decide to skip the time slice in which 180s elapsed since the previous bit “1” was transmitted.