

Terrapin Attack: Breaking SSH Channel Integrity By Sequence Number Manipulation

Fabian Bäumer
Ruhr University Bochum

Marcus Brinkmann
Ruhr University Bochum

Jörg Schwenk
Ruhr University Bochum

Abstract

The SSH protocol provides secure access to network services, particularly remote terminal login and file transfer within organizational networks and to over 15 million servers on the open internet. SSH uses an authenticated key exchange to establish a *secure channel* between a client and a server, which protects the confidentiality and integrity of messages sent in either direction. The secure channel prevents message manipulation, replay, insertion, deletion, and reordering. At the network level, SSH uses the SSH *Binary Packet Protocol* over TCP.

In this paper, we show that as new encryption algorithms and mitigations were added to SSH, the SSH Binary Packet Protocol is no longer a secure channel: SSH channel integrity (INT-PST) is broken for three widely used encryption modes. This allows *prefix truncation attacks* where some encrypted packets at the beginning of the SSH channel can be deleted without the client or server noticing it. We demonstrate several real-world applications of this attack. We show that we can fully break SSH extension negotiation (RFC 8308), such that an attacker can downgrade the public key algorithms for user authentication or turn off a new countermeasure against keystroke timing attacks introduced in OpenSSH 9.5. We also identified an implementation flaw in AsyncSSH that, together with prefix truncation, allows an attacker to redirect the victim's login into a shell controlled by the attacker.

In an internet-wide scan for vulnerable encryption modes and support for extension negotiation, we find that 77% of SSH servers support an exploitable encryption mode, while 57% even list it as their preferred choice.

We identify two root causes that enable these attacks: First, the SSH handshake supports optional messages that are not authenticated. Second, SSH does not reset message sequence numbers when encryption is enabled. Based on this analysis, we propose effective and backward-compatible changes to SSH that mitigate our attacks.

1 Introduction

Secure Shell (SSH). While TLS is commonly used to secure user-facing protocols such as web, email, or FTP, SSH is used by administrators to deploy and maintain these servers, often with high privilege (root) access and a large attack surface for lateral movement within an organization's infrastructure. As such, SSH was developed by Tatu Ylonen in 1995 as a secure alternative to telnet and rlogin/rcp and has since become a critical component of internet security.

In 1996, SSHv2 was developed to fix severe vulnerabilities in the original version. In February 1997, the IETF formed the SECSH working group to standardize SSHv2. After a decade, it published five core RFCs [25–29]. SSHv2 provides cryptographic agility and protocol agility without breaking backward compatibility. Since its original release, dozens of standardized and informal updates to the protocol have been published. Because of this, SSHv2 remains relevant after 25 years without major redesign, but it has also become difficult to analyze. There is a significant risk that these extensions of SSH interact to undermine its security goals.

SSH Connections. An SSH connection between a client and a server begins with the *Transport Layer Protocol* [29], which defines the handshake messages for key exchange and server authentication and how messages are exchanged over TCP using the SSH *Binary Packet Protocol* (BPP). After the handshake, SSH provides a *secure channel* for application data. At the application level, the client chooses a sequence of *services* to run. In practice, the client will run precisely two services: the *Authentication Protocol* [26] for user authentication with a password or public key, followed by the *Connection Protocol* [27] for the bulk of SSH's features like terminal sessions, port forwarding, and file transfer.

1.1 SSH Channel Security

In this work, we focus on the integrity of the SSH handshake and the resulting secure channel, as shown in [Figure 1](#). Af-

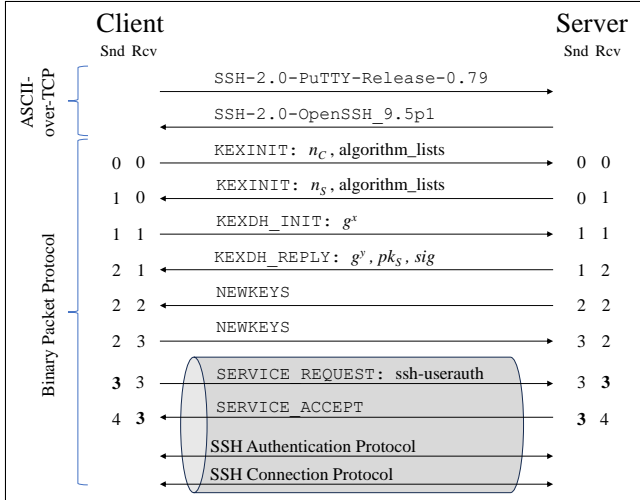


Figure 1: SSH handshake using a finite-field Diffie-Hellman key exchange. Included sequence numbers are implicit and maintained by the BPP. Snd denotes the counter for sent packets and Rcv for received packets. Sequence numbers verified using authenticated encryption are in **bold**.

ter an initial exchange of version information directly over TCP, the BPP exchanges *packets*, each containing precisely one *message*. Initially, the BPP is used without encryption or authentication for the duration of the handshake until the NEWKEYS message. Afterward, the encryption and authentication keys are used to form a secure channel, with the intent to protect the confidentiality and integrity of the *ordered stream of all following messages*. Note that technically, the secure channel consists of two separate cipher streams, one for each direction, and that the order of message arrival is only guaranteed for each direction separately.

Message Authentication Codes. As SSH is an interactive protocol, the integrity of each packet must be verified when it is received so that it can be promptly decrypted and processed. For this, the BPP appends a Message Authentication Code (MAC) to each packet. A cipher mode and a MAC form an authenticated encryption scheme [4]. SSH historically uses *Encrypt-and-MAC* (EaM), where the MAC is computed over the plaintext, but this is vulnerable to oracle attacks [2]. Later, *Encrypt-then-MAC* (EtM) was added, where the MAC is computed over the ciphertext instead. SSH has recently adopted the AEAD ciphers AES-GCM and ChaCha20-Poly1305, where ciphertext integrity is built into the encryption scheme [38].

A Trivial Example: Suffix Truncation Attacks. Note that a per-packet MAC cannot fully protect the channel’s integrity, as packets are verified and decrypted before the end of communication has been seen. This allows for a trivial *suffix*

truncation attack, where the attacker interrupts the message flow at some point during the communication. This is an inherent limitation of interactive protocols and an accepted trade-off in the design of SSH, but also, e.g., the TLS Record Layer. Although this attack cannot be prevented, it can be detected by requiring “end-of-communication” messages as the last messages in both directions. Unfortunately, the SSH standard does not define such messages, while TLS defines “close_notify” alerts for this purpose [37]. For SSH, it is left to the user to notice suspicious interruptions of SSH sessions.

Implicit Sequence Numbers. If the MAC was only computed over the payload of each packet, an attacker could still delete, replay, or reorder packets. Therefore, a *sequence number* is included in the MAC computation, corresponding to the position of the message in the stream. Each peer maintains two distinct counters (starting at 0), one for each direction. The Snd counter is incremented after a packet has been sent, and the Rcv counter is incremented after a received packet has been processed. Once the secure channel has been established, the current value of the Snd counter is used to compute the MAC of an outgoing packet and the current Rcv to verify the MAC of an incoming packet. If packets in the secure channel are deleted, replayed, or reordered, the sequence numbers get out of sync, and MAC verification will fail.

Because TCP is a reliable transport, accidental reordering of SSH packets cannot occur on the network. Thus, SSH (like other TCP-based protocols) uses *implicit sequence numbers* that are not transmitted as part of the packet.

Security Guarantees of Secure Channels. For TLS, the security guarantees of the Record Layer were formalized as stateful length-hiding encryption [35], where the state mainly consists of the implicit sequence number. The security of the BPP and implicit sequence numbers was analyzed by Bellare et al. in [3] and later refined and extended by Paterson and Watson [36] and Albrecht et al. [1]. These works define, in slightly idealized scenarios, the following informal security goal for a secure channel:

When a secure channel between A and B is used, the data stream received by B should be identical to the one sent by A and vice versa (INT-PST, [15]).

Within their idealization, all three works confirm that the BPP is indeed a secure channel. The difference between the models is that [36] also includes the encrypted length fields of the CBC Encrypt-and-MAC modes, and [1] considers the more recent cipher modes ChaCha20-Poly1305, AES-GCM, and generic Encrypt-then-MAC. Our attacks show that the models underlying the proofs in [1] are only partially accurate. We will explain the discrepancies between the proofs and our findings in [Section 2](#).

1.2 Overview of Our Attacks on SSH

In this paper, we show that SSH fails to protect the integrity of the encrypted message stream against meddler-in-the-middle (MitM) attacks. More precisely, we present novel *prefix truncation attacks* against SSH:

We show that the SSH Binary Packet Protocol is *not* a secure channel because a MitM attacker can delete a chosen number of integrity-protected packets from the beginning of the channel in either or both directions without being detected (Figure 2).

Attacker Model. We consider a MitM attacker who can observe, change, delete or insert bytes at the TCP layer. We assume that the attacker can *not* break the confidentiality of the session keys, i.e. the attacker has no information about the derived encryption keys, MAC keys, or IV.

Prefix Truncation Attacks. While our attacks on SSH are novel, the idea of prefix truncation attacks against network protocols by sequence number manipulation is not. To the best of our knowledge, the first and only description of such an attack is by Fournet (on behalf of miTLS) in an email to the TLS working group in 2015, targeting a draft version of TLS 1.3 [16]. Fournet’s attack increases sequence numbers in TLS by message fragmentation rather than message injection and remains theoretical, as “*prefix truncations will probably cause the handshake to fail.*” Subsequently, the draft was modified, and no prefix truncation attacks against the final version of TLS 1.3 are known. In contrast, we present the first real-world, practical prefix truncation attack against a mature, widely used protocol.

Root Cause Analysis. Our results depend on two technical observations about how SSH protects the integrity of the handshake and channel:

1. *SSH does not protect the full handshake transcript.* Although server authentication uses a signature to verify the integrity of the handshake, the signature is formed over a fixed list of handshake messages rather than the complete transcript. This gap in authentication allows an attacker to insert messages into the handshake and thereby manipulate sequence numbers.
2. *SSH does not reset sequence numbers at the beginning of the secure channel.* Instead, SSH increases sequence numbers monotonically, independent of the encryption state. Any manipulation of sequence numbers before the secure channel carries over into the channel.

Based on these two key observations, we present a series of novel attacks on SSH with increasing complexity and impact.

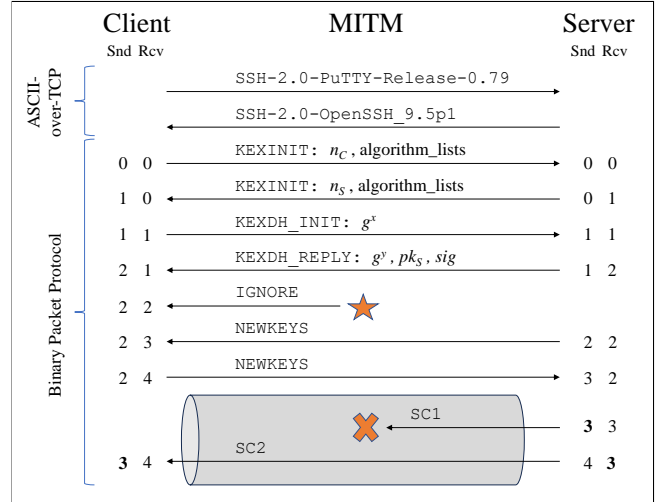


Figure 2: A novel prefix truncation attack on the BPP. The server sends SC1 and SC2, but the client only receives SC2.

Sequence Number Manipulation. We show that an attacker can *increase the receive counters* of the server and the client by inserting messages into the handshake (Section 4.1). Although not required for any of our attacks, we also show that, for some implementations, an attacker can *fully control the receive and send counters*, setting them to arbitrary, attacker-chosen values (Appendix A).

A Prefix Truncation Attack on the BPP. An attacker can use sequence number manipulation to *delete a chosen number of packets at the beginning of the secure channel*. Neither the client nor the server detects this change. This breaks the channel integrity of SSH (Section 4.2).

Extension Negotiation Downgrade Attack. As a practical example, we show an attack that uses prefix truncation to *break extension negotiation* [9], thereby downgrading the security of the connection. The attacked client might mistakenly believe that the server does not support recent signature algorithms for user authentication or does not implement certain countermeasures to attacks (Section 5.2). Specifically, we can turn off protection against keystroke timing attacks in the recently released OpenSSH 9.5.

Rogue Extension Attack and Rogue Session Attack. As another example, we show two attacks on the AsyncSSH client and server. In the first attack, *the victim’s extension info message is replaced* with one chosen by the attacker (Section 6.1). For the second attack, the attacker must have a user account on the same server as the victim. The attacker injects a malicious user authentication message so that *the victim logs into a shell controlled by the attacker* rather than the victim’s shell, thereby giving the attacker complete control

over the victim’s terminal input (Section 6.2). These attacks rely on implementation flaws of the AsyncSSH server, in addition to the prefix truncation attack.

Limitations. Our attacks critically depend on the SSH encryption mode negotiated between the client and the server. The attack works best with the AEAD cipher ChaCha20-Poly1305 (added in 2013). The attack also works with significant probability with any CBC-EtM mode (added in 2012), while CTR-EtM is vulnerable but not exploitable. On the other hand, CBC-EaM, CTR-EaM, and GCM modes are not vulnerable. See Section 4.3 for a complete analysis.

In an internet-wide scan, we show that despite these limitations, 77% of all SSH servers on the internet support a vulnerable encryption mode, and 57% even list it as their preferred choice (Section 7).

1.3 Our Contributions

We contribute the following novel results:

- An analysis of the integrity of SSH channels, where we identify two previously unknown flaws in the SSH specification, namely gaps in the handshake authentication and the use of sequence numbers across key activation.
- A novel prefix truncation attack on SSH channel integrity, where we show that an attacker can manipulate the sequence numbers and delete several messages from the beginning of the secure channel.
- A first security analysis of SSH extension negotiation, including a novel downgrade attack that disables extension negotiation completely. Consequently, support for public key signature algorithms or, in the case of OpenSSH 9.5, protection against keystroke timing attacks is disabled.
- As a practical demonstration, two novel attacks on AsyncSSH. First, a rogue extension attack, where the attacker can insert a chosen extension negotiation message. Second, a rogue session attack that allows the attacker to log the victim into an attacker-controlled shell. Both escalate implementation flaws in AsyncSSH using the prefix truncation attack.
- An internet-wide scan with up-to-date information on the distribution of SSH encryption modes and extensions.

Artifacts. Proof-of-concept implementations for our attacks and the results of our internet-wide scan in an aggregated form are available under the Apache-2.0 open-source license. See: <https://github.com/RUB-NDS/Terrapin-Artifacts>

Ethics Consideration and Responsible Disclosure. We started a coordinated disclosure and are in discussion with the maintainers of OpenSSH and AsyncSSH, who acknowledge our findings. Other vendors of major SSH implementations will be contacted later in the process. For our internet-wide scans, we provide an opt-out option and an email address for inquiries. Additionally, we deployed a blocklist to exclude networks that did opt out of previous scans. Scan results will only be published in aggregated form, without any information that could identify individual servers or networks.

2 Related Work

Secure Channels. In 2001, Canetti and Krawczyk [11] established the first model for secure channels, which only requires protection against adversarial insertion of messages. Bellare et al. [3] use stateful authenticated decryption to establish a more robust model, which, as the authors state, should also prevent replay and out-of-order delivery attacks. They use implicit sequence numbers as a shared state between sender and receiver. Paterson et al. [35] define stateful length-hiding authenticated encryption (sLHAE) to model the TLS record layer as a secure channel. This definition was used in [21, 22] to define authenticated and confidential channel establishment (ACCE) to analyze the combination of the TLS handshake and record layer. These steps to formalize secure channels were summarized by Fischlin et al. in [15]. Our attacks show that SSH BPP, when instantiated with ChaCha20-Poly1305, CBC-EtM, or CTR-EtM does not provide integrity of plaintext streams (INT-PST) as defined there.

Formally, SSH BPP security was modelled as stateful decryption [1, 3, 36]. Implicitly, this state was associated with SSH sequence numbers, and it was assumed that this state cannot be manipulated by an adversary. These models can be extended in two directions: (1) Include a broader definition of state. By including chained IVs, key stream state and GCM invocation counters, these models can be used to show why certain cipher modes resist our attacks, and that they indeed achieve INT-PST security. (2) Introduce a novel adversarial query `ModifyState` to model the attacks described here.

Truncation Attacks. Suffix truncation attacks against web services using TLS have been demonstrated by Smyth and Pironti in [39]. A prefix truncation attack against a draft version of TLS 1.3 was described by Fournet (on behalf of miTLS) in an email to the TLS working group in 2015 [16]. Fournet’s attack increases TLS sequence numbers by message fragmentation rather than injection to avoid breaking handshake authentication. The attack remained theoretical as “*prefix truncations will probably cause the handshake to fail.*” As a countermeasure, the draft was changed back to reset sequence numbers to 0 when activating keys.

Attacks on SSH. The most severe attack on SSH was presented by Albrecht, Paterson, and Watson [2] in 2009. It exploited the encrypted length field, using the length of the ciphertext accepted by the server from the network as a decryption oracle for a ciphertext block. In [36], this peculiarity of the BPP was formalized, and in [1] a variant of this attack was presented. Other attacks on SSH include a timing attack on SSH keystrokes by Song, Wagner, and Tian [40], a theoretical attack on SSH CBC cipher modes by Wei Dai [12], and a SHA-1 chosen prefix collision attack on the handshake transcript by Bhargavan and Leurent [7]. The weakness of some SSH host keys presented by Heninger et al. [18] was caused by a lack of entropy and faulty implementations and is not an inherent weakness of the protocol.

Formal Proofs for SSH. The SSH handshake was analyzed in [6, 42]. [3] presents a generic security model for SSH BPP, and [36] a specific, more detailed one for SSH-CTR. Albrecht et al. [1] analyze encryption modes in SSH, provide deployment statistics for them, and present a variant of the attack on the length field from [2]. They also include security statements for using ChaCha20-Poly1305, generic Encrypt-then-MAC, and AES-CTR in SSH, claiming the indistinguishability and integrity of the ciphertext. Careful analysis of their proofs reveals some assumptions that do not hold. In particular, they assume that the sequence counter is initialized to 0, which is false for ChaCha20-Poly1305 in SSH. They also assume that the symmetric encryption is a pseudorandom permutation, which is not valid for CBC with IV-chaining. This latter assumption is not apparent from the paper, which omits the proof, but it is discussed carefully in [17] by Hansen, one of the authors.

3 Background

SSH Handshake (Figure 1). To initiate an SSH connection, both peers exchange a version banner. The Binary Packet Protocol (see below) is used from the third message on but without encryption and authentication. In the KEXINIT messages, nonces and ordered lists of algorithms are exchanged: One list for key exchange, one for server signatures, and two (one per direction) each for encryption, MAC, and compression. For each list, the negotiated algorithm is the first client algorithm in that list which is also offered by the server.

In the KEXDHINIT and KEXDHREPLY messages, a finite-field Diffie-Hellman key exchange is performed. SSH also supports elliptic curves (ECDH) and hybrid schemes with post-quantum cryptography (PQC) as alternatives. The server authenticates itself with a digital signature as part of the handshake. The signature is computed over the contents of the previously exchanged messages, in a specified order.




Figure 3: CBC-EaM in SSH (diagram based on [2, Fig. 1]).

The Exchange Hash: A Partial Handshake Transcript.

In contrast to TLS, SSH uses only a selection from the handshake transcript for authentication. The hash value computed from this selection is called *exchange hash* H , defined as:

$$H = \text{HASH}(V_C \parallel V_S \parallel I_C \parallel I_S \parallel K_S \parallel X \parallel K)$$

where HASH is the hash function of the negotiated key exchange, V_C and V_S are the version banners of the client and server, I_C and I_S are the KEXINIT messages, K_S is the server’s public host key, and K the shared secret derived from the key exchange. The value of X depends on the key exchange and contains a composition of negotiated parameters (if any) and the ephemeral public keys of the key exchange [29, Sec. 8]. Each field includes a length field defined by the encoding.

Although the exchange hash contains everything that may influence the negotiation of algorithms or computation of the shared secret, it excludes seemingly ‘unimportant’ messages or message parts, such as IGNORE messages and unrecognized messages. This authentication gap allows a MitM attacker to inject messages into the handshake.

Sequence Numbers. Each sequence number is stored as a 4-byte unsigned integer initialized to zero upon connection. After a binary packet has been sent or received, the corresponding sequence number Snd or Rcv is incremented by one. Sequence numbers are never reset for a connection but roll over to 0 after $2^{32} - 1$. To avoid replay attacks, rekeying must occur at least once every 2^{32} packets [33, Sec. 6.1].

We illustrate the use of sequence numbers in Figure 1: After the banner exchange, the counters Snd and Rcv are initialized with (0, 0) on both sides. During algorithm negotiation and key exchange, sequence numbers are increased but not used in any MAC computation or verification. Only after keys are activated the secure channel is established, and sequence numbers are used for MAC computation and verification. For each BPP packet, the sequence numbers in bold must match at both peers; otherwise, the BPP packet is rejected.

SSH Binary Packet Protocol. The BPP is used to encrypt and authenticate messages. First, a message is prefixed by a 4-byte message length and a 1-byte padding length. Then, at least 4 bytes of padding are added to the message so that the

total plaintext length is a multiple of the block size or 8, whatever is larger. On the secure channel, the packet is encrypted by the cipher mode, and a MAC is added. The details depend on the chosen cipher mode, which makes use of an implicit initialization vector IV_{KDF} derived from the session key:

CBC-EaM [29] is part of the original SSH specification. Here, the MAC is computed over the implicit sequence number and the packet plaintext (Figure 3). The IV of the first packet is IV_{KDF} , and IV chaining is used (i.e., the IV of packet i is the last ciphertext block of packet $i - 1$).

CBC-EtM [32] was added to OpenSSH in 2012. Here, the packet length is *not* encrypted to allow checking the MAC before decryption. The MAC is computed over the sequence number, the unencrypted packet length, and the ciphertext of the packet payload. The IVs are handled as with CBC-EaM.

CTR [33] mode was proposed by Bellare, Kohno, and Namprempe [3] as a countermeasure to attacks on CBC with IV chaining. Here, IV_{KDF} is used as the initial counter value and incremented after encrypting a plaintext block. CTR can be used with EaM or EtM, with identical implications for the length field and MAC computation as above.

GCM [20] mode was specified by the NSA for Suite B-compliant SSH implementations [19]. Here, ciphertext integrity is part of the encryption scheme. The length field is *not* encrypted (solely authenticated) to allow verification of the authentication tag before returning any plaintext. Internally, GCM uses an invocation counter that is initialized to IV_{KDF} and incremented by one for each message. The sequence number is not used but is always offset by a constant from the invocation counter.

ChaCha20-Poly1305 [30] was added to OpenSSH in 2013, inspired by a similar proposal for TLS by Langley and Chang [23, 24]. Here, two different encryption keys are derived, one for the length field and one for the packet payload. This prevents the length field from being used as a decryption oracle for the payload. The MAC is computed over the concatenation of the two ciphertexts. Internally, the AEAD construction uses the sequence number as IV for each packet.

We note that the SSH specification says that the length field is encrypted [29, Sec. 6], and that the sequence number is used for integrity checks [29, Sec. 6.4]. This is only true for CBC-EaM, CTR-EaM, and ChaCha20-Poly1305. The modes CBC-EtM, CTR-EtM, and GCM do not encrypt the length field, and GCM also does not use the sequence number.

4 Breaking SSH Channel Integrity

In this section, we present a novel *prefix truncation attack* on SSH. The basic idea is that the attacker injects messages into the handshake to increase the implicit sequence number in one of the peers and then deletes a corresponding number of messages to that peer at the beginning of the secure channel. Two key insights about the SSH protocol enable this attack:

SSH Does Not Protect the Full Handshake Transcript.

As detailed in Section 3, the exchange hash signed by the server during the handshake only authenticates some parts of the handshake transcript, while other parts are left unauthenticated. This allows an attacker to inject some messages into the handshake, which cannot affect the key exchange but can affect the implicit sequence numbers of the peers.

SSH Does Not Reset Sequence Numbers at the Beginning if the Secure Channel.

In SSH, sequence numbers are only incremented and never reset to 0, even when the encryption state changes. This allows an attacker to manipulate the sequence number counters in the secure channel before encryption and authentication keys are activated.

Comparison to Other Protocols.

In IPsec/IKE, only a portion of the handshake transcript is signed, but unlike SSH, sequence numbers are reset to 0 when encryption and MAC keys are activated. In TLS, FINISHED messages are exchanged at the beginning of the secure channel to verify the integrity of the complete handshake transcript, and sequence numbers are reset to 0 after every CHANGECIPHERSPEC message.

4.1 Sequence Number Manipulation

In this section, we show how a MitM attacker can arbitrarily increase the receive sequence numbers C.Rcv and S.Rcv in the client and the server. This will be the basis for our prefix truncation attack and its applications, allowing the attacker to compensate for messages deleted from the secure channel.

Technique RcvIncrease (Figure 6a). During the handshake, a MitM attacker can increase C.Rcv (resp. S.Rcv) by N , while not changing any other sequence number, by sending N IGNORE messages to the client (resp. server).

The correctness is evident from the fact that the SSH standard requires for IGNORE that “*All implementations MUST understand (and ignore) this message at any time.*” [29, Sec. 11.2]. The intended purpose of this message is to protect against traffic analysis, so it is considered a security feature, although there is no benefit from it during the handshake phase. We note that the attacker may also use any other message type that does not generate a response.

In addition, we found that an attacker can set the sequence numbers to arbitrary values by abusing the rollover after 2^{32} messages. These advanced techniques require that the implementation allows handshakes with many messages, a large amount of data, and a long operating time. As we do not require these advanced techniques for our attacks, a description can be found in Appendix A.

Evaluation. We verified that, as expected, this technique works with OpenSSH 9.4p1 and 9.5p1, Dropbear 2022.83,

PuTTY 0.79, AsyncSSH 2.13.2, and libssh 0.10.5.

4.2 Prefix Truncation Attack on the BPP

Single Message Prefix Truncation Attack. We assume the attacker wants to delete the first message SC1 sent from the server (Figure 2). The attack takes two steps:

1. The attacker uses the RcvIncrease technique to increase C.Rcv by one, e.g., by injecting an IGNORE message to the client before NEWKEYS.
2. The attacker deletes the first message SC1 sent by the server.

We first analyze this attack with regard to handshake authentication and sequence numbers. As the key exchange does not protect the handshake transcript from inserting IGNORE messages (Section 3), handshake authentication is not broken. Before the first step, we have $C.Rcv = C.Snd$. After the first step, we have $C.Rcv = S.Snd + 1$, but during the handshake this manipulation is not detected. After the second step, we have $C.Rcv = S.Snd$, and sequence numbers are back in sync.

It remains to be shown that the attacker can delete the message from the channel, which requires knowing its length, and that its deletion does not affect the MAC verification and decryption output for the following messages. This analysis depends on the encryption mode and will be given in Section 4.3. Here, we conclude by describing the general attack.

(N_S, N_C) -Prefix Truncation Attack. In a single attack, the attacker can generally delete an arbitrary number of N_S initial messages sent from the server and N_C initial messages sent from the client. This is straightforward: Instead of inserting one IGNORE message to the client before NEWKEYS, the attacker inserts N_S such messages to the client and N_C to the server. Consequently, instead of deleting the first message from the server, the attacker deletes N_S initial messages from the server and the N_C initial messages from the client.

Note that the single message attack above is the specific case of a $(1, 0)$ -prefix truncation attack.

4.3 Analysis of Encryption Modes

In this section we first discuss how an attacker can determine the byte-length of SSH messages. Then we analyze which encryption modes are vulnerable to our attacks, and which can be exploited. An encryption mode is *vulnerable* if, after prefix truncation, all following packets on the secure channel are decrypted, i.e. an AEAD mode does not generate the distinguished symbol INVALID or an authenticated encryption mode successfully verifies the MAC. Note that we allow decryption to a different plaintext for probabilistic attacks. To capture this, we define that an encryption mode is *exploitable* for an attack, if the message stream after decryption is well-formed and supports that attack.

Determining the Byte-Length of Messages. To successfully delete packets from the secure channel, the attacker has to know their length. For some encryption modes (CBC-EtM, CTR-EtM, GCM), the length is sent unencrypted so the attacker can observe it. For others (CBC-EaM, ChaCha20-Poly1305), the length is encrypted, and the attacker has to guess the length, either from known plaintext based on the used SSH implementations or from network side channels such as TCP segment sizes. For the following analysis, we assume that the attacker always knows the message lengths.

4.3.1 Not Vulnerable

GCM. GCM [20] mode does not use the implicit sequence number. Instead, it uses an invocation counter to derive encryption and MAC keys, initialized to IV_{KDF} and incremented after each message. The authors justify this by stating that the resulting nonce is always a fixed offset from the sequence number. By deviating from the SSH standard, GCM stops our attack, as the attacker cannot manipulate the invocation counter during the handshake.

CBC-EaM and CTR-EaM. CBC uses IV chaining, and CTR uses a key stream. When the attacker deletes any prefix of the ciphertext in either mode, the first ciphertext block received will be decrypted as pseudo-random. Because EaM computes the MAC over the plaintext, MAC verification will fail with probability close to 1, and our attack is stopped.

4.3.2 Vulnerable And Perfectly Exploitable

ChaCha20-Poly1305. ChaCha20-Poly1305 [30] directly uses the sequence number in its key derivation, which makes it vulnerable to our prefix truncation attack. All messages following the truncated prefix are decrypted to their original plaintext, because the integrity check of the AEAD cipher is only over the ciphertext and the sequence number that has been manipulated by the attacker to match. Under the assumption that the attacker can correctly guess the packet length, the prefix truncation attack always succeeds.

Note that the fault is not with ChaCha20-Poly1305 as an AEAD encryption scheme, but with its integration into the SSH secure channel construction.

4.3.3 Vulnerable, But Not Exploitable

CTR-EtM. With CTR-EtM, the MAC is computed over the unencrypted length, the sequence number, and the ciphertext. So, removing some packets from the beginning of the channel does not cause a MAC failure, and cryptographically, the attack succeeds. However, CTR uses a block counter initialized to IV_{KDF} , which is incremented after each block. After prefix truncation, the key stream is desynchronized, so *all* following ciphertexts are decrypted as pseudo-random packets. Each

corrupted packet has a significant probability of causing a critical failure, which eventually stops our attack.

Remark: Decryption Oracle for CTR-EtM Using Prefix Truncation. For CTR-EtM, prefix truncation of k blocks (which exactly contain one or more messages) provides a very limited *decryption oracle* on the ciphertext c_1, \dots, c_k where $c_i := \text{Enc}(\text{IV}_{\text{KDF}} + i) \oplus p_i, 1 \leq i \leq k$. After deleting the first k blocks, MAC verification for the following message of length l blocks will succeed because the length, sequence number, and ciphertext are correct. The blocks c_{k+1}, \dots, c_{k+l} will be decrypted as $p'_j := \text{Enc}(\text{IV}_{\text{KDF}} + j) \oplus c_{k+j}, 1 \leq j \leq l$, and processed as a pseudo-random SSH message SC1'. Due to format oracle side channels in SSH at the BPP layer, e.g. the padding length, but also at the protocol layer, e.g. if a message is ignored or triggers a response, the attacker can get some information about the bits in p'_j . This reveals information about the first l key stream blocks, and thus also about p_1, \dots, p_l , potentially leaking confidential information like passwords in user authentication. If processing SC1' does not cause a critical failure, the attack can even continue further with the decryption of the next message, revealing more about the following key stream and, thus, plaintext. Exploiting this requires a careful study of format oracles in SSH, which is outside the scope of this work.

4.3.4 Vulnerable And Probabilistically Exploitable

CBC-EtM. With CBC-EtM, the MAC is computed from the unencrypted length, the sequence number, and the ciphertext. The IV is not required because IV_{KDF} is implicit, and all other IVs are authenticated before use. Consequently, prefix truncation does not cause a MAC failure, and cryptographically, the attack succeeds. Nevertheless, we need to consider the impact that IV chaining has on the immediately following packet to see if this attack is practically exploitable.

Recall that the decryption of the first block is $p_1 := \text{Dec}(c_1) \oplus \text{IV}_{\text{KDF}}$ and for block i it is $p_i := \text{Dec}(c_i) \oplus c_{i-1}$. We assume the attacker uses prefix truncation to remove blocks c_1, \dots, c_k . The following block c_{k+1} will now be decrypted as $p'_1 := \text{Dec}(c_{k+1}) \oplus \text{IV}_{\text{KDF}}$. We are interested in how SSH implementations process the resulting pseudo-random block p'_1 as the first block in the decrypted packet. Intuitively, it should result in a corrupted packet that causes a critical failure.¹

Surprisingly, there is a significant probability that the attack can continue, although it is highly implementation-dependent. For a corrupted packet, there are three possible outcomes:

1. *Critically Corrupt:* If the corruption is detected at the BPP or application level, e.g., some length field exceeds the packet length, the connection should be closed.

¹ Similarly to CTR-EtM, any format oracle side channel for p'_1 reveals a relationship between IV_{KDF} and p_{k+1} via $\text{IV}_{\text{KDF}} \oplus p_{k+1} = c_k \oplus p'_1$, which is a marginal information leak for the (secret) IV given information on p_{k+1} , and vice versa. Again, we do not explore this further here.

2. *Marginally Corrupt:* If the packet happens to be similar enough to the original, e.g., if the corruption is limited to optional fields, it should be processed without error and have the same effect as the original would have had.
3. *Evasively Corrupt:* If the packet is well-formed (i.e., has valid padding length) but has an unrecognized message ID, an UNIMPLEMENTED response must be sent, and the connection continues normally [29, Sec. 11.4].

Clearly, the first outcome stops any attack from going forward. However, the second and third outcome may be beneficial for the attacker. We will now present two instructive scenarios for these two outcomes and estimate the success probability of an attack relying on that outcome. Later, we will see some experimental verification for these estimates.

Scenario 1: CBC-EtM Prefix Truncation Of a Single Message, Second Message Has Format Flexibility. In this scenario, the attacker wants to remove the first message, and the second, corrupted message needs to be functionally preserved but has some format flexibility. For example, the second message might be SERVICEACCEPT (see Section 5.2), which is critical to start user authentication. The encrypted part of the packet looks like this, where p is the padding length, m the message id, and n the service name length:

p	m	n	Service Name
0e	06	00 00 00 0c	s s h - u s e r a u
t	h	Random Padding	

The probability that the first block decrypts exactly as shown is only 2^{-128} for a 128-bit block cipher. However, for some clients, the service name string is optional. These clients accept a 1-byte message with $p = 30$ (0x1E) and $m = 6$ as *marginally corrupt*, which has a success probability of 2^{-16} , independent of the block size.

Although SERVICEACCEPT may be a lucky case, there are structural reasons for this result: First, SSH messages are often short and can be smaller than a single block. Second, the padding is random and cannot be verified. Third, some messages have redundant fields that implementations ignore (e.g., the service name above).

We experimentally verified that OpenSSH, Dropbear, PuTTY, and libssh allow empty SERVICEACCEPT messages from the server, enabling this attack. At the same time, AsyncSSH is very strict and requires the correct service name.

Scenario 2: CBC-EtM Prefix Truncation Attack On More Than One Message. In this scenario, we assume the attacker wants to remove the first $N > 1$ messages and perfectly preserve all of the following messages. Then, the attacker can use prefix truncation to delete the first $N - 1$ messages, and take a bet on the N -th message to be *evasively corrupt*.

Let ℓ be the length of the ciphertext of the N -th message, with padding length p , message ID m , and random padding. The attack succeeds regardless of the content of the corrupted packet as long as it is well-formed and unrecognized: A packet is *well-formed* if $4 \leq p \leq \ell - 2$ (accounting for the padding length and message ID). A packet is *unrecognized* if m is a message ID not known by the implementation.

Because the message is well-formed, it is not rejected at the BPP layer. Furthermore, because the message is unrecognized, the peer must respond with UNIMPLEMENTED and otherwise ignore it [29, Sec. 11.4], so our attack succeeds.

The probability that a packet is well-formed depends on ℓ . The padding length is between 4 and 255, and ℓ is a multiple of $\max(8, \text{block size})$, so the number of valid padding length values is $\min(252, \ell - 5)$ out of 2^8 .

The probability that a packet is unrecognized depends on the implementation. The attack requires at least one unknown message ID. Through source code review, we identified 43 IDs that are in active use, so we estimate up to 213 unknown message IDs out of 2^8 .

In total, assuming a block size of at least 128-bit (i.e., $\ell \geq 16$), we estimate that the success probability of this attack is between $11 \cdot 2^{-16} \approx 0.0002$ and $252 \cdot 213 \cdot 2^{-16} \approx 0.8190$ for vulnerable implementations. Our experiments show success probabilities from 0.0003–0.8383, in good agreement with our analysis (Section 5.2). Increasing the block size increases the lower bound, while the upper bound stays the same.

5 Breaking SSH Extension Negotiation

While the fact that BPP does not implement a secure channel is troublesome enough, exploiting this vulnerability requires an analysis of the SSH protocol at the application layer.

As our attack achieves prefix truncation, it is only natural to ask which SSH messages can occur at the beginning of a secure channel. Historically, the first messages exchanged are `SERVICEREQUEST` and `SERVICEACCEPT`. Removing either causes the connection to go stale, as the client will not begin the user authentication protocol. Then our attack, while cryptographically successful, fails at the application layer.

However, the SSH Extension Negotiation mechanism [9] introduces a new message, `EXTINFO`, which can occur immediately after `NEWKEYS` as the first message on the secure channel. Some of the extensions that can be negotiated are security-relevant, providing a relevant attack surface for our prefix truncation attack and raising its impact.

In this section, we will first describe SSH Extension Negotiation and then demonstrate how an attacker can downgrade the security of a connection by removing the `EXTINFO` message from the secure channel in a prefix truncation attack.

5.1 SSH Extension Negotiation

Even though the original SSH RFCs were designed with extensibility in mind, they do not provide any mechanism to negotiate protocol extensions securely. RFC 8308 [9] closes this gap. The RFC describes a signaling mechanism enabling extension negotiation, the extension negotiation mechanism itself, and a set of initially defined extensions.

Support for extension negotiation is signaled as part of the `KEXINIT` message. To not pose compatibility issues, the structure of the message is not being altered, and the reserved field is not used. Instead, each peer may include an indicator name within the list of key exchange algorithms. The indicator name differs depending on the role of the peer (`ext-info-c` vs. `ext-info-s`) to avoid accidental negotiation.

Whenever a peer signals support for extension negotiation, the other side may send an `EXTINFO` message as the first message after `NEWKEYS`. Additionally, the server can send a second `EXTINFO` later to authenticated clients to avoid disclosing extension support to unauthenticated clients. Each `EXTINFO` message contains several extension entries. Negotiation requirements are defined on a per-extension level.

RFC 8308 defines an initial set of four protocol extensions, and vendors have proposed and implemented additional extensions. We detail those relevant to our attacks here and describe the others in Appendix B.

`server-sig-algs` [9] is a server-side extension that informs the client about all supported signature algorithms when using a public key during client authentication.

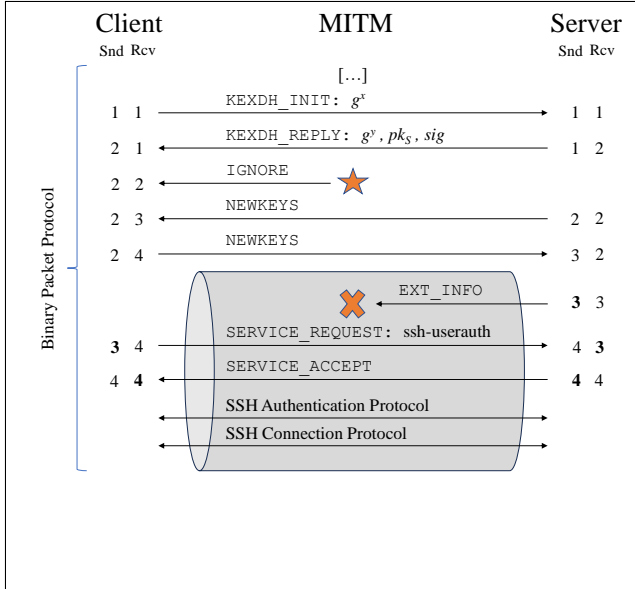
`publickey-hostbound@openssh.com` [31,32] is a server-side extension to advertise support for host-bound public key authentication, which deviates from public key authentication by also covering the server’s host key. This allows the enforcement of per-key restrictions when allowing remote servers to access local secret keys (i.e., when using SSH Agent).

`ping@openssh.com` [32] is a server-side extension to advertise support for a transport level ping message similar to the Heartbeat extension in TLS [41].

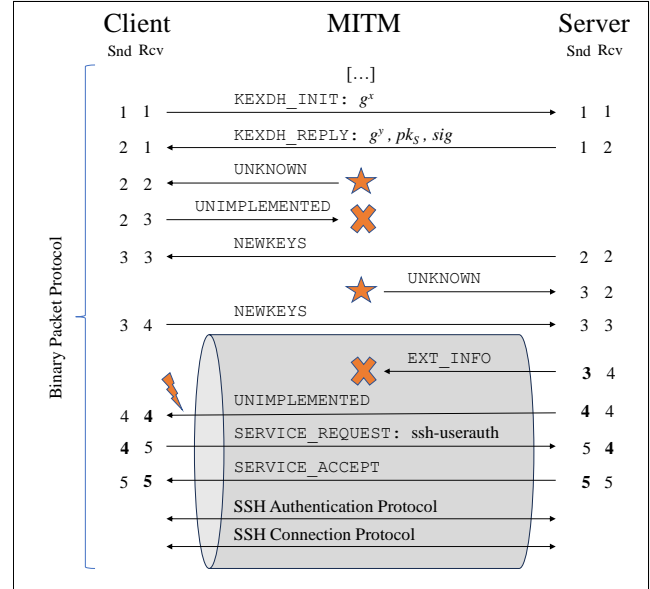
5.2 Extension Downgrade Attack

We now show how the prefix truncation attack can be applied to delete the `EXTINFO` message sent by the client, server, or both parties without either noticing. Our attack differs depending on the encryption mode. For ChaCha20-Poly1305, we can use the basic attack strategy. For CBC-EtM, we show two strategies to generate additional messages in the secure channel, so that the attacker can use the “evasively corrupt” outcome of Scenario 2 in Section 4.3.

Extension Downgrade for ChaCha20-Poly1305. The downgrade attack for ChaCha20-Poly1305 against the client is depicted in Figure 4a. It is identical to the single message prefix truncation attack from Section 4.2, with `EXTINFO` now



(a) Extension Downgrade Attack for ChaCha20-Poly1305: The MitM injects an IGNORE message before the handshake concludes. The change in sequence numbers allows the MitM to strip the EXTINFO from within the secure channel.



(b) Extension Downgrade Attack for CBC-EtM: The MitM injects UNKNOWN before the NEWKEYS is sent by the client. As the server already sent NEWKEYS, the provoked UNIMPLEMENTED message will be sent within the secure channel after EXTINFO. The corrupted UNIMPLEMENTED message has a significant probability of being ignored (see Scenario 2 in Section 4.3).

Figure 4: Variants of the extension downgrade attack for ChaCha20-Poly1305 and CBC-EtM.

taking the place of SC1 in Figure 2. If the attack should be directed against the server instead, a $(0, 1)$ -prefix truncation attack should be performed instead. This allows an attacker to delete any EXTINFO sent immediately after NEWKEYS.

While the server may send a second EXTINFO just before signaling successful client authentication, stripping the EXTINFO message sent after NEWKEYS renders most publicly specified extensions unusable. This is because they are either scoped to the authentication protocol, sent by the client only, or must be sent by both parties to take effect. Solely the ping@openssh.com extension may be sent in the second EXTINFO to enable keystroke timing countermeasures inside the connection protocol. However, OpenSSH 9.5 does not implement any facility to send a second extension negotiation message. As we show in Section 7, extensions scoped to the authentication protocol are the most common among SSH servers on the internet by a significant margin.

Successfully performing the extension downgrade can directly impact the security level of the connection. Most notably, the recently introduced keystroke timing countermeasures by OpenSSH 9.5 will remain disabled when the server has not sent ping@openssh.com. Furthermore, stripping an EXTINFO containing the server-sig-algs extension can lead to a signature downgrade during client authentication, as the client has to resort to trial-and-error instead.

Extension Downgrade for CBC-EtM. In Figure 4b, we show how the attack can also work with CBC-EtM. Suppose an attacker injects an UNKNOWN message to the server after the server sends NEWKEYS and EXTINFO, but before the client’s NEWKEYS message (and also injects UNKNOWN to the client to realign sequence numbers). In that case, the server sends the response UNIMPLEMENTED as the second message in the secure channel immediately after the EXTINFO message. The attacker now wants to remove two messages from the channel and can benefit from the “evasively corrupt” in Scenario 2 in Section 4.3. The attacker removes EXTINFO from the secure channel, which causes the decryption of the first block of UNIMPLEMENTED to be pseudo-random. Because UNIMPLEMENTED messages are relatively small ($\ell = 16$ for AES), the upper estimate for the success probability is only $11 \cdot 213 \cdot 2^{-16} \approx 0.0358$.

However, the success probability can be increased significantly by exploiting the new ping extension in OpenSSH 9.5. To make use of this, the attacker replaces the UNKNOWN message sent to the server with a PING message containing at least 255 bytes of payload. As per specification, the server will reflect this data in the PONG response. This yields $\ell \geq 264$, maxing out the probability of the packet being well-formed. Consequently, the upper estimate for the success probability is now $252 \cdot 213 \cdot 2^{-16} \approx 0.8190$.

Evaluation. We successfully evaluated the attack on ChaCha20-Poly1305 and CBC-EtM against OpenSSH 9.5p1 and PuTTY 0.79 clients, connecting to OpenSSH 9.4p1 (UNKNOWN only) and 9.5p1. For CBC-EtM, our success rate in practice was 0.0003 (OpenSSH) resp. 0.03 (PuTTY), improved to 0.0074 (OpenSSH) resp. 0.8383 (PuTTY) when sending PING instead of UNKNOWN.

6 Message Injection Attacks on AsyncSSH

Going beyond the SSH specifications, we now demonstrate how prefix truncation attacks can also be used to exploit implementation flaws. Specifically, we target AsyncSSH,² an SSH implementation for Python with an estimated 60k daily downloads.³ We present two attacks that exploit weaknesses in the handling of unauthenticated messages during the handshake. These attacks are enabled by prefix truncation and sequence number manipulation.

Note that we describe these attacks only for ChaCha20-Poly1305. Adjusting them for CBC-EtM is straightforward, injecting appropriate IGNORE and UNKNOWN messages, but requires some of the advanced techniques described in [Appendix A](#). These advanced techniques only work against some SSH implementations.

6.1 Rogue Extension Negotiation Attack

The rogue extension negotiation attack targets an AsyncSSH client connecting to any SSH server sending an EXTINFO message. The attack exploits an implementation flaw in the AsyncSSH client to inject an EXTINFO message chosen by the attacker and a prefix truncation against the server to delete its EXTINFO message, effectively replacing it.

The attack is a variant of the extension downgrade attack in [Section 5.2](#), but instead of IGNORE, the attacker sends a chosen EXTINFO packet to the client. Similar to IGNORE, EXTINFO does not trigger a response from the client. A correct SSH implementation should not process an unauthenticated EXTINFO message. However, the injected message is accepted due to flaws in AsyncSSH.

AsyncSSH clients support the `server-sig-algs` and `global-requests-ok` extensions. Hence, the attacker can try to downgrade the algorithm used for client authentication by meddling with the value of `server-sig-algs`.

Evaluation. We successfully evaluated the attack against AsyncSSH 2.13.2 as a client, connecting to AsyncSSH 2.13.2.

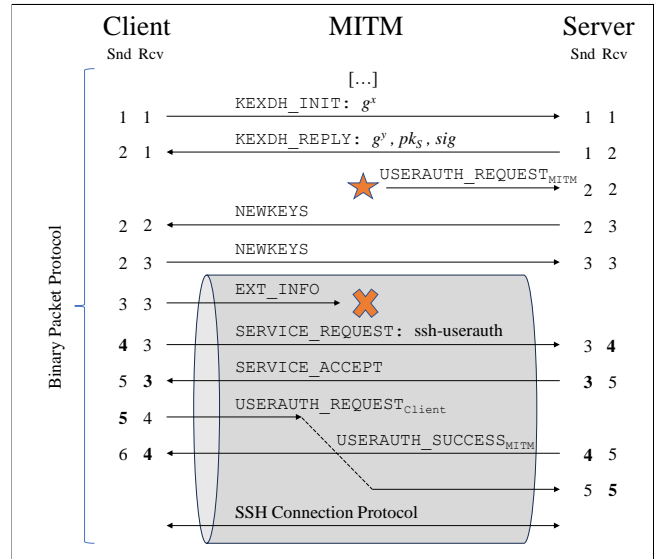


Figure 5: Rogue Session Attack on AsyncSSH: The MitM injects a malicious authentication request before the handshake completes. The client-side extension information message is deleted to account for the change in sequence numbers. By delaying the authentication request sent by the client, the MitM ensures that the malicious one is being processed. Any additional authentication requests are silently ignored.

6.2 Rogue Session Attack

The rogue session attack targets any SSH client connecting to an AsyncSSH server, on which the attacker must have a shell account. The goal of the attack is to log the client into the attacker’s account without the client being able to detect this. At that point, due to how SSH sessions interact with shell environments, the attacker has complete control over the remote end of the SSH session. The attacker receives all keyboard input by the user, completely controls the terminal output of the user’s session, can send and receive data to/from forwarded network ports, and is able to create signatures with a forwarded SSH Agent, if any. The result is a complete break of the confidentiality and integrity of the secure channel, providing a strong vector for a targeted phishing campaign against the user. For example, the attacker can display a password prompt and wait for the user to enter the password, elevating the attacker’s position to a MitM at the application layer and enabling impersonation attacks.

The messages exchanged during the attack are depicted in [Figure 5](#). The attacks work by the attacker injecting a chosen USERAUTHREQUEST before the client’s NEWKEYS. The USERAUTHREQUEST sent by the attacker must be a valid authentication request containing his credentials. The attacker can use any authentication mechanism that does not require

²<https://github.com/ronf/asyncssh>

³<https://pypi.org/packages/asyncssh>

exchanging additional messages between client and server, such as `password` or `publickey`. Due to a state machine flaw, the AsyncSSH server accepts the unauthenticated `USERAUTHREQUEST` message and defers it until the client has requested the authentication protocol.

To avoid a race condition between the `USERAUTHREQUEST` sent by the client and `USERAUTHREQUEST` injected by the attacker, the attacker delays the client’s `USERAUTHREQUEST` until after the server signaled a successful authentication in response to the injected `USERAUTHREQUEST`. The AsyncSSH server silently ignores any additional authentication request after a successful authentication.

To complete the attack, the attacker has to fix the sequence numbers using one of two strategies (note that [Figure 5](#) only shows the first strategy):

1. Suppose the client sends an extra message before `SERVICEREQUEST`. In that case, the attacker can delete that message from the channel, effectively performing the (0,1)-prefix truncation attack with `USERAUTHREQUEST` instead of the usual `IGNORE` message.
2. Suppose the server sends an extra message before `SERVICEACCEPT`. In that case, the attacker can delete that message from the channel after injecting an additional `UNKNOWN` message to the client before `NEWKEYS`, triggering a `UNIMPLEMENTED` response. This increases both `C.Snd` and `C.Rcv`, moving the send count deficit from the client to the server.

Evaluation. We successfully evaluated the attack against AsyncSSH 2.13.2 as a server, connecting to with AsyncSSH 2.13.2 and OpenSSH 9.4p1.

7 SSH Deployment Statistics

To estimate the impact of the prefix truncation attacks, we scan for the SSH servers preferring or supporting the vulnerable encryption modes. Similarly, to estimate the impact of the extension downgrade attack, we scan for servers sending `EXTINFO` messages.

Methodology. For scanning, we used ZMap [14] and ZGrab2 [13] on port 22 of the entire IPv4 address space. The first scan was performed over two days in early October 2023, totaling 15.164M SSH servers.

As ZGrab2 cannot capture SSH extensions, we performed a second scan at the end of June 2023 using a custom tool on a subset of 2^{20} open ports discovered by ZMap. The scan covered a total of 830k servers. All data relating to the use of extension negotiation in SSH is sourced from this scan.

In SSH, the algorithm order of the client determines which algorithm is preferred. However, we cannot scan for actual client use. Assuming that servers and clients are bundled in

Cipher Family	Preferred		Supported	
ChaCha20-Poly1305	8,739k	57.64%	10,247k	67.58%
AES-CTR	4,785k	31.56%	14,866k	98.04%
AES-GCM	1,219k	8.04%	10,450k	68.92%
AES-CBC	236k	1.56%	4,069k	26.84%
Other	147k	0.97%	-	-
Unknown / No KEXINIT	34k	0.23%	-	-
Total	15,164k	100%		

Table 1: Preferred SSH cipher families as of October 2023.

a single product and share algorithm preference and support, we use the server’s lists as a surrogate, as was also done in [1].

Symmetric Encryption Algorithms. In [Table 1](#), we show the number of servers preferring and supporting various encryption modes. A cipher is preferred if placed first in the list of supported algorithms.

We find that, by far, the most preferred encryption cipher is ChaCha20-Poly1305, with 57.64% listing this algorithm first. This is followed by AES-CTR (31.53%) and, with some distance, by AES-GCM (8.04%) and AES-CBC (1.56%).

Authenticated Encryption Modes. As non-AEAD ciphers must be combined with a MAC, we also evaluate which *authenticated* encryption modes are preferred and supported by the servers. The numbers for the AEAD modes ChaCha20-Poly1305 (57.64%) and GCM (8.04%) are identical to those for encryption modes, as the MAC is already integrated. Preference for CTR modes is split into a majority for CTR-EaM (26.14%) and a minority for CTR-EtM (5.46%). Preference for CBC modes is mostly CBC-EaM (2.37%), while only a marginal share of servers prefers CBC-EtM (0.09%).

In summary, 57.73% of all servers prefer an authenticated encryption mode vulnerable to our attacks.

Looking at the support for authenticated encryption modes vulnerable to our attacks, we find that 67.58% of all servers support ChaCha20-Poly1305, while 17.24% support CBC-EtM. In total, 77% support at least one vulnerable mode, and 7% support both (these numbers are not shown in the table).

SSH Extensions. We also looked at SSH extensions offered by servers before user authentication; see [Table 3](#). We can see that 76.81% of all servers send the `server-sig-algs` extensions to indicate support for better signature schemes for client public key authentication. Furthermore, 8.8% send the `publickey-hostbound` extension, improving security for multi-hop authentication authentication using SSH agent. Both extensions provide opportunities for downgrade attacks, as their absence can weaken the strength of the authentication.

Cipher Mode and MAC	Preferred		Supported	
ChaCha20-Poly1305	8,739k	57.64%	10,247k	67.58%
CTR-EaM	3,964k	26.14%	4,200k	27.70%
GCM	1,219k	8.04%	10,450k	68.92%
CTR-EtM	828k	5.46%	10,685k	70.46%
CBC-EaM	359k	2.37%	1,585k	10.46%
CBC-EtM	14k	0.09%	2,614k	17.24%
Other	4k	0.28%	-	-
Unknown / No KEXINIT	34k	0.23%	-	-
Total	15,164k	100%		

Table 2: Distribution of supported authenticated encryption modes as of October 2023.

Extension name	Times Offered	
server-sig-algs	637,466	76.81%
publickey-hostbound@	73,040	8.80%
delay-compression	283	0.03%
no-flow-control	283	0.03%
global-requests-ok	283	0.03%

Table 3: SSH extensions offered by servers after the initial handshake, @openssh.com abbreviated to @. Extensions sent by clients, and by servers upon successful client authentication are not included.

8 Suggested Countermeasures

As a stop-gap measure, the vulnerable cipher modes can be disabled. Suitable alternatives are AES-GCM or AES-CTR, which are widely supported. However, the root cause analysis shows that the real problems are in the SSH specification. We therefore suggest two changes to the specification, and a way to negotiate support for these changes without breaking backward compatibility.

Sequence Number Reset. Resetting sequence numbers to zero when encryption keys are activated ensures that these sequence number manipulations during the handshake can no longer affect the secure channel. Resetting the sequence numbers is a major break in compatibility. To avoid connection failures due to only one peer resetting their sequence numbers, we suggest that an implementation signals the support for this countermeasure by including an identification string in the list of supported key exchange algorithms. The SSH extension negotiation mechanism is already employing this method. If and only if both peers signal support for this countermeasure, the sequence numbers will be reset.

Full Transcript Hash. The second countermeasure is to authenticate the full handshake transcript, as seen by the client and the server. This can detect any attempts of handshake

manipulation by a MitM attacker, including sequence number manipulation through our techniques. It is not possible to simply extend the scope of the exchange hash, as the server signature is transmitted before the new keyset is taken into use. Therefore, any messages sent after the key exchange, but before NEWKEYS, can not be included. We suggest that the both peers send a MAC of the full transcript at the start of the secure channel, similarly to TLS FINISHED messages. Signaling support should be done as above.

Other Issues. We suggest that SSH specifies “end-of-communication” messages to detect suffix truncation attacks. Also, AsyncSSH should be hardened to disallow unauthenticated, application-layer messages during the SSH handshake.

9 Conclusion

We have shown that the complexity of SSHv2 has increased over its 25 years of development to a point where the addition of new algorithms and features has introduced new vulnerabilities. The root cause analysis has shown that the potential for our attacks was already present in the original specification. Handshake transcripts were never fully authenticated, and sequence numbers were never reset to 0. However, as new authenticated encryption modes and extension messages were added, these weaknesses grew into exploitable vulnerabilities.

We introduced the novel sequence number manipulation and prefix truncation attacks for secure channels, which invalidate the INT-PST [15] security of SSH BPP for certain ciphers. We extended this vulnerability to real-world exploits like disabling SSH extension negotiation. This yields novel insights into the complex interplay between a practical security mechanism (sequence numbers) and abstract security notions (INT-PTXT vs. INT-CTXT vs. AEAD, [5]): Since implicit sequence numbers are not transmitted, they cannot be part of the ciphertext – therefore generic EtM modes fail, and INT-CTXT does not protect against prefix truncation. INT-PTXT and AEAD modes, on the other hand, do include them, either as part of the plaintext, or as associated data.

Our close look at the Extension Negotiation mechanism reveals its design weaknesses: First, sending EXTINFO is optional even if both parties signal support for extension negotiation during the handshake. Second, SSH extension negotiation cannot be used to negotiate extensions affecting the SSH handshake itself, e.g., the countermeasures proposed in this paper. As a consequence, all extension negotiations should be done within the KEXINIT.

Although we suggest backward-compatible countermeasures to stop our attacks, we note that the security of the SSH protocol would benefit from a redesign from scratch. This redesign should be guided by all findings and insights from both practical and theoretical security analysis, in a similar manner as was done for TLS 1.3.

References

- [1] Martin R. Albrecht, Jean Paul Degabriele, Torben Brandt Hansen, and Kenneth G. Paterson. A surfeit of SSH cipher suites. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016*, pages 1480–1491. ACM Press, October 2016.
- [2] Martin R. Albrecht, Kenneth G. Paterson, and Gaven J. Watson. Plaintext recovery attacks against SSH. In *2009 IEEE Symposium on Security and Privacy*, pages 16–26. IEEE Computer Society Press, May 2009.
- [3] Mihir Bellare, Tadayoshi Kohno, and Chanathip Namprempre. Authenticated encryption in SSH: Provably fixing the SSH binary packet protocol. In Vijayalakshmi Atluri, editor, *ACM CCS 2002*, pages 1–11. ACM Press, November 2002.
- [4] Mihir Bellare and Chanathip Namprempre. Authenticated encryption: Relations among notions and analysis of the generic composition paradigm. In Tatsuaki Okamoto, editor, *ASIACRYPT 2000*, volume 1976 of *LNCS*, pages 531–545. Springer, Heidelberg, December 2000.
- [5] Mihir Bellare and Chanathip Namprempre. Authenticated encryption: Relations among notions and analysis of the generic composition paradigm. *Journal of Cryptology*, 21(4):469–491, October 2008.
- [6] Florian Bergsma, Benjamin Dowling, Florian Kohlar, Jörg Schwenk, and Douglas Stebila. Multi-ciphersuite security of the Secure Shell (SSH) protocol. In Gail-Joon Ahn, Moti Yung, and Ninghui Li, editors, *ACM CCS 2014*, pages 369–381. ACM Press, November 2014.
- [7] Karthikeyan Bhargavan and Gaëtan Leurent. Transcript collision attacks: Breaking authentication in TLS, IKE and SSH. In *NDSS 2016*. The Internet Society, February 2016.
- [8] Denis Bider. Extended authentication information in Secure Shell (SSH). Internet-Draft draft-ssh-ext-auth-info-01, Internet Engineering Task Force, March 2018. Work in Progress.
- [9] Denis Bider. Extension Negotiation in the Secure Shell (SSH) Protocol. RFC 8308, March 2018.
- [10] Denis Bider. Sending and Handling of Global Requests in Secure Shell (SSH). Internet-Draft draft-ssh-global-requests-ok-00, Internet Engineering Task Force, December 2018. Work in Progress.
- [11] Ran Canetti and Hugo Krawczyk. Analysis of key-exchange protocols and their use for building secure channels. In Birgit Pfitzmann, editor, *EUROCRYPT 2001*, volume 2045 of *LNCS*, pages 453–474. Springer, Heidelberg, May 2001.
- [12] Wei Dai. email to IETF mailing list. <https://www.ietf.org/ietf-ftp/ietf-mail-archive/secsh/2002-02.mail>, 2002. Accessed: 2023-10-11.
- [13] Zakir Durumeric, David Adrian, Ariana Mirian, Michael Bailey, and J. Alex Halderman. A search engine backed by internet-wide scanning. In Indrajit Ray, Ninghui Li, and Christopher Kruegel, editors, *ACM CCS 2015*, pages 542–553. ACM Press, October 2015.
- [14] Zakir Durumeric, Eric Wustrow, and J. Alex Halderman. ZMap: Fast internet-wide scanning and its security applications. In Samuel T. King, editor, *USENIX Security 2013*, pages 605–620. USENIX Association, August 2013.
- [15] Marc Fischlin, Felix Günther, Giorgia Azzurra Marson, and Kenneth G. Paterson. Data is a stream: Security of stream-based channels. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part II*, volume 9216 of *LNCS*, pages 545–564. Springer, Heidelberg, August 2015.
- [16] Cédric Fournet. email to IETF mailing list. <https://mailarchive.ietf.org/arch/msg/tls/exto09ETJLnEm3MRDT023x70DFM>, 2015. Accessed: 2023-10-16.
- [17] Torben Brandt Hansen. *Cryptographic Security of SSH Encryption Schemes*. Phd thesis, University of London, 2020.
- [18] Nadia Heninger, Zakir Durumeric, Eric Wustrow, and J. Alex Halderman. Mining your ps and qs: Detection of widespread weak keys in network devices. In Tadayoshi Kohno, editor, *USENIX Security 2012*, pages 205–220. USENIX Association, August 2012.
- [19] Kevin Igoe. Suite B Cryptographic Suites for Secure Shell (SSH). RFC 6239, May 2011.
- [20] Kevin Igoe and Jerome Solinas. AES Galois Counter Mode for the Secure Shell Transport Layer Protocol. RFC 5647, August 2009.
- [21] Tibor Jager, Florian Kohlar, Sven Schäge, and Jörg Schwenk. On the security of TLS-DHE in the standard model. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 273–293. Springer, Heidelberg, August 2012.

- [22] Hugo Krawczyk, Kenneth G. Paterson, and Hoeteck Wee. On the security of the TLS protocol: A systematic analysis. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 429–448. Springer, Heidelberg, August 2013.
- [23] Adam Langley and Wan-Teh Chang. ChaCha20 and Poly1305 based Cipher Suites for TLS. Internet-Draft draft-agl-tls-chacha20poly1305-04, Internet Engineering Task Force, November 2013. Work in Progress.
- [24] Adam Langley, Wan-Teh Chang, Nikos Mavrogiannopoulos, Joachim Strombergson, and Simon Josefsson. ChaCha20-Poly1305 Cipher Suites for Transport Layer Security (TLS). RFC 7905, June 2016.
- [25] Chris M. Lonvick and Sami Lehtinen. The Secure Shell (SSH) Protocol Assigned Numbers. RFC 4250, January 2006.
- [26] Chris M. Lonvick and Tatu Ylonen. The Secure Shell (SSH) Authentication Protocol. RFC 4252, January 2006.
- [27] Chris M. Lonvick and Tatu Ylonen. The Secure Shell (SSH) Connection Protocol. RFC 4254, January 2006.
- [28] Chris M. Lonvick and Tatu Ylonen. The Secure Shell (SSH) Protocol Architecture. RFC 4251, January 2006.
- [29] Chris M. Lonvick and Tatu Ylonen. The Secure Shell (SSH) Transport Layer Protocol. RFC 4253, January 2006.
- [30] Damien Miller. This document describes the chacha20-poly1305@openssh.com authenticated encryption cipher supported by openssh. <https://cvsweb.openbsd.org/cgi-bin/cvsweb/src/usr.bin/ssh/PROTOCOL.chacha20poly1305?rev=1.5>. Accessed: 2023-10-18.
- [31] Damien Miller. SSH agent restriction. <https://www.openssh.com/agent-restrict.html>, 2022. Accessed: 2023-10-17.
- [32] Damien Miller, Markus Friedl, Mike Frysinger, Todd C. Miller, and Darren Tucker. This documents openssh’s deviations and extensions to the published ssh protocol. <https://cvsweb.openbsd.org/cgi-bin/cvsweb/src/usr.bin/ssh/PROTOCOL?rev=1.49>. Accessed: 2023-10-18.
- [33] Chanathip Namprempre, Tadayoshi Kohno, and Mihir Bellare. The Secure Shell (SSH) Transport Layer Encryption Modes. RFC 4344, January 2006.
- [34] Richard Ogier. OSPF Database Exchange Summary List Optimization. RFC 5243, May 2008.
- [35] Kenneth G. Paterson, Thomas Ristenpart, and Thomas Shrimpton. Tag size does matter: Attacks and proofs for the TLS record protocol. In Dong Hoon Lee and Xiaoyun Wang, editors, *ASIACRYPT 2011*, volume 7073 of *LNCS*, pages 372–389. Springer, Heidelberg, December 2011.
- [36] Kenneth G. Paterson and Gaven J. Watson. Plaintext-dependent decryption: A formal security treatment of SSH-CTR. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 345–361. Springer, Heidelberg, May / June 2010.
- [37] Eric Rescorla. The Transport Layer Security (TLS) Protocol Version 1.3. RFC 8446, August 2018.
- [38] Phillip Rogaway. Authenticated-encryption with associated-data. In Vijayalakshmi Atluri, editor, *ACM CCS 2002*, pages 98–107. ACM Press, November 2002.
- [39] Ben Smyth and Alfredo Pironti. Truncating TLS connections to violate beliefs in web applications. In *7th USENIX Workshop on Offensive Technologies (WOOT 13)*, Washington, D.C., August 2013. USENIX Association.
- [40] Dawn Xiaodong Song, David A. Wagner, and Xuqing Tian. Timing analysis of keystrokes and timing attacks on SSH. In Dan S. Wallach, editor, *USENIX Security 2001*. USENIX Association, August 2001.
- [41] Michael Williams, Michael Tüxen, and Robin Seggelmann. Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS) Heartbeat Extension. RFC 6520, February 2012.
- [42] Stephen C. Williams. Analysis of the SSH key exchange protocol. In Liqun Chen, editor, *13th IMA International Conference on Cryptography and Coding*, volume 7089 of *LNCS*, pages 356–374. Springer, Heidelberg, December 2011.

A Advanced Sequence Number Manipulation

Using that sequence numbers rollover to zero after $2^{32} - 1$, the attacker can also attempt to *decrement* C.Rcv and S.Rcv:

Technique RcvDecrease (Figure 6b). During the handshake, a MitM attacker can decrease C.Rcv (resp. S.Rcv) by N , while not changing any other sequence number, by sending $2^{32} - N$ IGNORE messages to the client (resp. server).

A single IGNORE message is only 5 bytes, so it fits into a single block even for a 128-bit block cipher. Sending $2^{32} - N$ such messages transfers $\approx 2^{32} \cdot 16B \approx 69GB$ of data. Consequently, this technique can fail on implementations with timeouts or restrictions to the amount of data or the number of messages transferred during the handshake.

We can also combine these techniques to manipulate the C.Snd and S.Snd sequence numbers. For this, we require a message that generates a response message but is otherwise ignored. Conveniently, the SSH standard requires this for all messages with unrecognized message IDs [34, Sec. 11.4]. Let UNKNOWN be a message with an unrecognized message ID.

Technique SndIncrease (Figure 6c). During the handshake, a MitM can increase C.Snd (resp. S.Snd) by N while not changing any other sequence number by sending N UNKNOWN and $2^{32} - N$ IGNORE messages to the client (resp. server) and deleting all generated UNIMPLEMENTED messages.

Technique SndDecrease (Figure 6d). During the handshake, a MitM can decrease C.Snd (resp. S.Snd) by N while not changing any other sequence number by sending $2^{32} - N$ UNKNOWN and N IGNORE messages to the client (resp. server), and deleting all generated UNIMPLEMENTED messages.

Here, the total data transfer required is $\approx 69GB$ for Snd-Increase and twice as much, or $\approx 138GB$ for SndDecrease. Again, these techniques may fail on implementations that have timeouts or restrict the amount of data or number of messages exchanged during the handshake.

Evaluation. We verified all techniques successfully against PuTTY 0.79. Additionally, our experiments show that OpenSSH 9.5p1 recognizes a rollover of sequence numbers and terminates the connection, thus not being affected by any advanced technique. AsyncSSH 2.13.2 and libssh 0.10.5 terminate the connection due to handshake timeouts before these techniques conclude. Dropbear 2022.83 disconnects on UNKNOWN messages instead of responding with UNIMPLEMENTED, but allows Rcv to roll over, therefore being affected by RcvDecrease only.

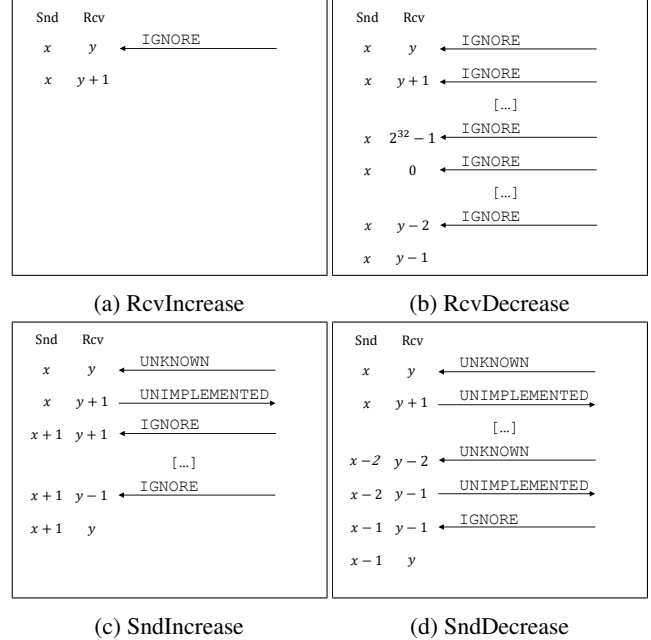


Figure 6: Techniques for sequence number manipulation as a MitM in the SSH protocol. All techniques can target either client or server before the initial handshake concludes. The MitM deletes all generated UNIMPLEMENTED messages.

B More Extensions

These extensions are defined in addition to those listed in Section 5.

- `no-flow-control` - Negotiable extension to disable flow control in the SSH connection layer [9].
- `delay-compression` [9] - Negotiable extension to renegotiate compression algorithms as delayed compression without a key re-exchange. Delayed compression in the context of SSH refers to any compression that will take effect after successful client authentication.
- `elevation` - Informational client-side extension to specify the desired security context in case of administrator logins in Windows operating systems [9].
- `global-requests-ok` [10] - Informational extension to announce global requests within the SSH connection layer will be handled in accordance with RFC 4254.
- `ext-auth-info` [8] - Client-side extension to signal support for extended authentication failure messages during client authentication containing an arbitrary set of key-value pairs.