

# Towards SSH3: how HTTP/3 improves secure shells

François Michel  
UCLouvain, Belgium  
francois.michel@uclouvain.be

Olivier Bonaventure  
UCLouvain & WELRI, Belgium  
olivier.bonaventure@uclouvain.be

## ABSTRACT

The SSH protocol was designed in the late nineties to cope with the security problems of the telnet family of protocols. It brought authentication and confidentiality to remote access protocols and is now widely used. Almost 30 years after the initial design, we revisit SSH in the light of recent protocols including QUIC, TLS 1.3 and HTTP/3. We propose, implement and evaluate SSH3, a protocol that provides an enhanced feature set without compromise compared to SSHv2. SSH3 leverages HTTP-based authorization mechanisms to enable new authentication methods in addition to the classical password-based and private/public key pair authentications. SSH3 users can now configure their remote server to be accessed through the identity provider of their organization or using their Google or Github account. Relying on HTTP/3 and the QUIC protocol, SSH3 offers UDP port forwarding in addition to regular TCP forwarding as well as a faster and secure session establishment. We implement SSH3 over quic-go and evaluate its performance.

## 1 INTRODUCTION

Remote access to distant computers was one of the motivations for the creation of the ARPANET. It was usually realized through a TCP connection between the client and distant servers. This connection was used to authenticate the user and then exchange commands. Several protocols use this approach including telnet [1] and rsh/rlogin [26]. These protocols suffered from various security issues [8] and are now deprecated. In the mid 1990s, the Secure Shell (SSH) protocol [43] was proposed as a secure alternative. It provides several remote services over an authenticated and encrypted channel using a single TCP connection: (i) execution of a program, (ii) access to a shell session, (iii) TCP port forwarding and (iv) forwarding of X11 graphical sessions. Standardized in 2006, the specification describes the version 2.0 of the protocol, also known as SSHv2. The design of SSHv2 is complex and spread over four documents: the high-level architecture design [31], the user authentication protocol [29], the transport layer protocol [32] and the connection protocol [30]. SSHv2 builds its own secure channel and provides mechanisms for user authentication and authorization.

With the growing use of the web for critical operations such as bank transactions or e-commerce, the need for secure communications democratized the use of HTTPS [15] with strong security guarantees coming from the Transport Layer Security (TLS) protocol. TLS 1.3 [36] now provides comparable security features to SSHv2 with a shorter session establishment. A lot of effort has also been put in HTTP to enhance user authorization to control the access to protected web resources. This led to the design of a generic authorization protocol for HTTP [16] allowing the use of newer techniques such as the *Bearer* [25] mechanism in addition to the older password-based authentication [17]. This enabled the development of advanced mechanisms such as OAuth 2.0 [22]. OAuth opened the door to the OpenID Connect standard [37] extensively

used by major Internet actors such as Google and Microsoft to provide user authentication through their identity platforms. All this makes HTTP authorization more flexible than SSH while keeping applications simple. With its use for multiple scenarios, HTTP authorization is still evolving with new schemes being developed [38].

Finally, the diversity of contents available on a remote computer has evolved. SSHv2 provides the X11 forwarding service to use graphical programs and access visual content. However, this service is not suited to highly dynamic content such as videos and livestreams (e.g. a high framerate camera) available on the remote server. An alternative is to use the TCP port forwarding feature of SSH to tunnel a streaming application. This is however not sufficient if low latency is needed since TCP provides a fully reliable bytestream at the cost of added latency. Users wanting such features have to use alternatives such as the Remote Desktop Protocol (RDP) [4] that includes UDP-based transport for displaying a low latency remote graphical desktop. Furthermore, with the significant growth of QUIC traffic on the Internet [20], the TCP forwarding of SSHv2 is not sufficient anymore and the lack of UDP port forwarding may become progressively more visible in the future.

We start from the above observations and consider the following question: “*What would the design of the SSH protocol look like if it were invented today?*”. We revisit SSH from the ground-up and propose SSH3, a new version of the protocol based on the strong features of modern web protocols. Running on top of HTTP/3, SSH3 is as easy to deploy as any HTTP application. We describe the technical details of SSHv2 and HTTP/3 in Section 2 and show how SSH3 can provide either similar or stronger features than its predecessor in Section 3. We present our implementation in Section 4 and show performance improvements of SSH3 in Section 5. We then discuss other appealing aspects of SSH3 in Section 6.

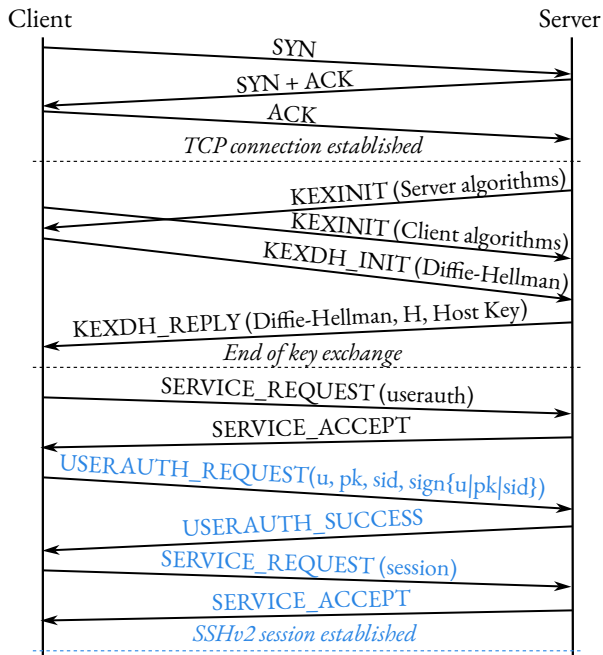
## 2 BACKGROUND

This section covers the concepts required to explain the design choices of SSH3. Section 2.1 describes how SSHv2 operates above TCP. Section 2.2 presents the HTTP/3 protocol and how modern applications can leverage its protocol stack to build advanced and easily deployable services with only little implementation effort.

### 2.1 SSHv2

Running solely atop TCP, SSHv2 defines its own mechanisms to build a secure channel and authenticate users. Upon successful key exchange and user authentication, users can access their remote server and start using the services provided by SSH.

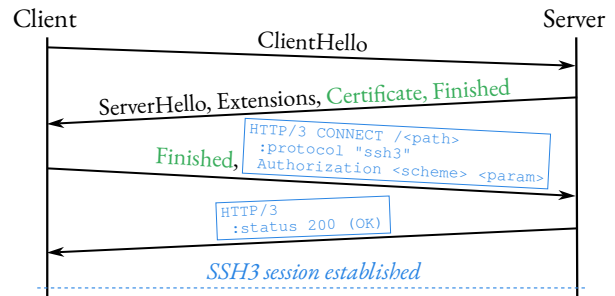
*2.1.1 Secure channel establishment.* As they carry critical and sensitive information, the confidentiality and authenticity of SSH communications is crucial. The design of SSH predates TLS. For this reason, SSH defines its own mechanisms for establishing a secure channel to encrypt and authenticate the session data. Figure 1



**Figure 1: SSHv2 session establishment using Diffie-Hellman.** SSH version exchange has been omitted to save space. Messages in blue are protected using the keys derived from the key exchange. OpenSSH uses an additional round-trip for user authentication by first sending the USERAUTH\_REQUEST(none) message to discover the available authentication methods.

illustrates the establishment of an SSHv2 session. Such a session starts with a TCP handshake depicted at the top of the Figure. The client and server exchange KEXINIT messages to negotiate the algorithms used for key exchange, encryption and hashing [32]. Then, the actual key exchange is performed (here, using Diffie-Hellman) with the client sending the KEXDH\_INIT message and the server replying with the KEXDH\_REPLY message. This message also contains its host key used for server authentication and H, a hash applied over several fields including the negotiated Diffie-Hellman values and the resulting shared secret. Upon reception of these messages, both ends can derive the cryptographic keys to secure the connection. Additionally, a session identifier is derived from the key exchange to uniquely identify the SSH session. This session identifier is later used to prevent replay attacks when using public key user authentication. The bottom of Figure 1 illustrates user authentication described in the next section.

**2.1.2 User authentication.** The SSHv2 protocol ensures that services are provided only to authenticated and authorized users. The SSHv2 standard defines the user-auth protocol to this end [29]. In order to start the user-auth protocol, the client first sends a SERVICE\_REQUEST message and waits for the answer of the server. SSHv2 supports several authentication methods, the main ones being password-based and public key authentication. The password-based mechanism simply consists in the client sending the username



**Figure 2: SSH3 session establishment.** Key exchange is performed during the QUIC handshake. Messages in green are protected using QUIC handshake keys. Messages in blue are protected using the keys from the key exchange (1-RTT keys). User authentication is performed by setting the Authorization header in the HTTP/3 CONNECT request. The <scheme> and <param> variables depend on the used HTTP authentication scheme.

and password through the secure channel and the server verifying that they match with the credentials on the system. The public key authentication illustrated at the bottom of Figure 1 is achieved with the user performing a signature over several fields, including the username (u), the public key (pk) and the session identifier (sid). The server then checks whether the public key proposed by the client is suitable for authentication (e.g. it matches the user's public key locally installed on the server's filesystem) and then verifies the signature. Signing the session identifier prevents replay attacks.

Some SSH implementations such as OpenSSH or Tectia support other ways to authenticate users. Among them is the certificate-based user authentication: only users in possession of a certificate signed by a trusted certificate authority (CA) can gain access to the remote server [12]. Available for more than 10 years, this authentication method requires setting up a CA and distributing the certificates to new users and is still not commonly used nowadays.

**2.1.3 Session flow.** Once the secure channel is established and the user is successfully authenticated, the latter can open new sessions and perform operations on the remote server. This is done by opening SSH channels with specific uses such as running a shell or executing programs. Channels are bidirectional message streams multiplexed over the single TCP connection used by SSHv2. The messages exchanged over a channel serve several purposes. For instance, messages can be used by a client to request the allocation of a new pseudo-terminal (PTY), the execution of new commands or simply sending the user-typed input to the remotely running processes. Messages can also be sent by the server to announce the termination of the remote program along with its status or signal code or to forward its standard and error outputs to the user. Finally, SSH users can open channels to perform TCP port forwarding. By doing so, the client listens on a specified local TCP port. Every connection initiated towards this port on the client will be tunneled to the server through the SSH channel. The server then initiates a new TCP connection towards a remote host whose IP and port number are chosen by the client. The bytes exchanged on the TCP connections are simply placed in new messages and

sent over the channel. Port forwarding has however its limits in SSHv2. First, as SSHv2 channels are all multiplexed on the only TCP connection carrying the whole session, head-of-line blocking can occur upon network losses if several TCP connections are forwarded. Second, SSHv2 only provides TCP forwarding and has no support for carrying UDP packets. SSHv2 is therefore unable to tunnel UDP-based DNS, RTP or even the QUIC protocol, which now constitutes a large part of Internet traffic nowadays [20].

## 2.2 HTTP/3

During the last years, the IETF has finalized the standardization of several protocols that provide a more secure and flexible transport than TCP. The latest version of TLS, TLS 1.3 [36] uses Diffie-Hellman and provides perfect forward secrecy by design. QUIC [23, 27] integrates the TLS 1.3 handshake to provide encryption and authentication of both control and application data as a transport feature and reduces the connection establishment duration.

The HTTP/3 protocol provides the features of HTTP/2 on top of QUIC instead of the classical combination of TCP and TLS. Aside from the security features discussed above, QUIC provides modern transport features such as stream multiplexing that allows applications to send data independently on several streams. In addition to these byte streams, QUIC also supports the exchange of datagrams [35]. QUIC provides seamless connection migration [23] and soon multipath communication [28], enabling smooth network handovers that would have disrupted the connection with TCP.

The Extended CONNECT HTTP extension allows the application to directly use the underlying QUIC stream for sending arbitrary protocol data. The application can also open new streams and send QUIC datagrams. This method is already extensively used by WebTransport [18] to bootstrap new transport connections in web browsers.

The most appealing aspect of HTTP/3 is its native support for user authentication [15]. A critical part of the SSH protocol resides in the session establishment and especially the user authentication process. HTTP already provides a solid set of mechanisms to perform user authentication that have been implemented and used for years for sensitive use-cases such as banking and e-commerce.

## 3 SSH3

In this article, we entirely reconsider the SSH protocol stack and propose SSH3, a modern iteration of its architecture. Several initiatives have already considered running SSH over the QUIC protocol, but these propositions limit themselves to carrying classical SSH mechanisms inside QUIC streams [11, 41]. In contrast with these propositions, SSH3 is built above HTTP/3, not directly over QUIC, and reconsiders the whole protocol architecture. One might wonder why building SSH3 over HTTP/3 instead of QUIC. An initial benefit is that HTTP/3 provides multiplexing at the URL level: SSH3 instances can be accessible through specific URLs. First, it allows the use of HTTP/3 proxies to act as SSH3 gateways leading to different physical servers depending on the URL path specified in the CONNECT request. Second, it can make SSH3 robust to scanning attacks as discussed in Section 6.1. This architecture also enables a wider feature set than SSHv2 over TCP and reuses existing HTTP/3 mechanisms. The secure channel establishment is entirely performed by the QUIC layer and the SSH3 session is initiated by

sending an HTTP/3 Extended CONNECT request. Furthermore, user authentication material is attached to the request using the HTTP authorization mechanisms. This significantly simplifies the design and implementation aspects of SSH3 and allows implementers to focus on the protocol features and actual system security.

### 3.1 Secure connection establishment

Figure 2 illustrates how an SSH3 session is established. The first three arrows correspond to the QUIC handshake, but the Finished message already carries the HTTP/3 Extended CONNECT request. With TLS 1.3, algorithms negotiation and key exchange are performed concurrently with the client sending a ClientHello message and the server responding with a ServerHello message, a few extensions required for key exchange and its certificate. These operations occur during the QUIC handshake and can be completed during the first round-trip [39], reducing significantly the connection establishment time. The first SSH3 data can already be exchanged after one round-trip while the first round-trip was dedicated to the TCP handshake in SSHv2, as depicted in Figure 1. To derive a unique session identifier, SSH3 relies on *TLS exporters* to export cryptographic material from the QUIC session to the upper layer protocol. Thanks to TLS exporters, both endpoints generate a same pseudorandom value uniquely tied to the current TLS session and the running protocol. This allows obtaining a unique session identifier similarly to SSHv2 as discussed in Section 2.1. Finally, TLS natively allows authenticating both clients and servers through certificates, similarly to the certificate-based authentication of OpenSSH. It also works through HTTP proxies as recent HTTP mechanisms allow forwarding client certificates through dedicated HTTP headers [10]. As client and server certificate-based authentication is part of TLS, SSH3 can benefit from this authentication method without implementation or maintenance effort. The use of TLS certificates on web servers has been democratized thanks to the efforts of Let's Encrypt [5]. Server certificates can now be installed in a matter of seconds and provide a stronger security than SSHv2 host keys currently used for authenticating servers. For a private or personal use of SSH3, self-signed certificates can also be deployed, providing a security level comparable to host keys.

### 3.2 HTTP-based authentication

SSH3 uses the generic HTTP authorization mechanism [15] and puts user authentication material in the Authorization header of the CONNECT request. This comes down to setting the <scheme> and <param> variables of Figure 2 to appropriate values. If the provided header suffices for authenticating and granting access to the user, the server responds with the *200 OK* HTTP response. Otherwise, the server returns a *401 Unauthorized* response. The latter includes the WWW-Authenticate header indicating at least one authentication method that can be used for subsequent requests. Our SSH3 prototype implements three authentication techniques.

**3.2.1 Password authentication.** Password authentication can be achieved using the HTTP *Basic* authentication scheme [17]. In Figure 2, the client sets the <scheme> variable to *Basic* and <param> to a base64-encoded value containing the username and password. Similarly to SSHv2, the server then compares the username and password with the local credentials and grant access if they match.

**3.2.2 OpenID Connect authentication.** The flexibility of HTTP authentication mechanisms opens the door to new authentication methods. For instance, the OpenID Connect protocol provides modern ways for authenticating users using an external identity provider. This protocol is implemented by popular Single Sign-On (SSO) services such as Google Identity or Microsoft Entra. For instance, using this authentication method, an SSH3 server can be configured to grant access to users that successfully logged in to their company's identity provider using their professional email address. The identity provider can itself implement advanced authentication methods such as multi-factor authentication without impacting the SSH3 server configuration. This method is entirely based on standard HTTP mechanisms such as OAuth2 and the *Bearer* authentication scheme [25]. To get access to the SSH3 server, the user first requests a base64-encoded ID token from the identity provider by logging in on its login platform [37]. This ID token issued by the identity provider certifies that the user successfully logged in and actually owns its email address. The token is passed to the SSH3 server by placing it in the `<param>` variable of Figure 2 and setting the `<scheme>` variable to *Bearer*. When decoded, the ID token is in the JSON Web Token (JWT) format [24]. In simple terms, JWTs can be seen as a JSON object (payload) associated with a header and a digital signature performed by the identity provider over the payload and the header. The header contains descriptive information such as the signature algorithm or the identifier of the key used to sign the JWT. The payload contains several fields such as a date of issue, an expiry time, the url of the identity provider and the email address of the user. The SSH3 server then parses the token and verifies its authenticity using the identity provider's public key. It also checks the validity of the JWT payload itself (e.g. ensuring the token has not expired). Once the token has been verified and if the user's email address is authorized by the SSH3 server, the server sends a *200 OK* response to the client's *CONNECT* request and the session can start. ID tokens can be stored on the client (e.g. using a keyring) and automatically renewed to avoid the user from manually logging in at each session establishment.

**3.2.3 Public key authentication.** SSH3 also provides mechanisms for public key authentication like SSHv2. Similarly to Section 3.2.2, this method relies on JWT tokens passed to the server using the *Bearer* authentication scheme. The difference is that the JWT token is signed using the user's private key instead of an external identity provider. The JWT token payload contains the public key as well as the unique session identifier obtained through the use of TLS exporters described in Section 3.1. To authenticate the user, the SSH3 server ensures that the public key in the JWT matches the user's locally installed public key and then verifies the signature of the JWT. Including the session identifier in the JWT payload provides a protection against replay attacks equivalent to the public key authentication of SSHv2 described in Section 2.1.2. This however makes the token deviate from the definition of a *Bearer* token as it does not, by design, grant the same access right to any other party in possession of it [25]. We see this as a minor and temporary issue. There is an ongoing standardisation effort towards a new HTTP *Signature* authentication scheme serving the exact purpose of asymmetric key authentication. Our plan is for SSH3 to use this new authentication scheme once it becomes standard [38].

### 3.3 Session flow

The SSH3 session can start once the user is authenticated. SSH3 integrates the same concepts as SSHv2: bidirectional channels can be used to exchange SSH messages serving different purposes: spawning new processes, carrying process standard input and output, signalling process termination, etc. In SSHv2, different channels are multiplexed on the same TCP connection, leading to head-of-line blocking in case of network packet losses. With SSH3, each channel is carried over a dedicated bidirectional QUIC stream. Messages transmitted over a same SSH3 channel still need to be delivered in-order to preserve the ordering of programs' output and the keystrokes typed by the user. Finally, SSH3 can provide both TCP and UDP port forwarding. Each TCP connection runs over a dedicated channel. UDP forwarding is done by tunnelling UDP packets coming in both direction using QUIC datagrams.

## 4 IMPLEMENTATION

We implemented the different concepts presented in the article and provide a first open source implementation of SSH3 [33]. Our implementation is written in Go. The Go language has a large network library making it easy to implement new client and server applications on top of the HTTP semantics. Our SSH3 implementation relies on the `quic-go` library [13] that provides support for both the QUIC and HTTP/3 protocols. Our implementation consists in 4000 lines of Go code in total, including channels and messages handling, session establishment, user authentication, unit tests as well as the client and server command-line programs. The implementation provides classical SSH services such as opening interactive shells attached to a pseudoterminal, executing single commands when interactivity is not needed and TCP port forwarding. It also provides the UDP port forwarding feature discussed in Section 3.3. It integrates the three authentication methods described in this article. Depending on the user-specific configuration on the server, users can authenticate themselves using passwords, OpenID Connect and public key authentication. The OpenID Connect authentication method currently opens a browser window for the user to log in to the identity provider. This authentication method thus currently needs a graphical environment on the client to be used.

Finally, for a smooth and progressive transition from SSHv2 to SSH3, our prototype provides secure OpenSSH agent forwarding. For instance, this allows users to open a shell session from their laptop to a remote server using SSH3 and then connect from that server to legacy SSHv2 servers using the SSH keys locally installed on the laptop. This also enables the use of SSH3 servers as gateways towards SSHv2 servers and still benefit from SSH3 features on the first hop such as client connection migration.

## 5 EVALUATION

In this section, we evaluate our implementation of SSH3 through diverse scenarios. It can be complicated to evaluate SSH3 and compare it with SSHv2 as some SSH3 features are simply not present on SSHv2. This section focuses on three elements. In Section 5.1, we show that SSH3 significantly reduces the session establishment time compared to SSHv2. In Section 5.2, we show that our implementation improves the user experience of an interactive shell session. In Section 5.3 we show that our implementation provides




Figure 3: Completion time of non-interactive sessions.

port forwarding capabilities suitable for the WAN scenario and that UDP port forwarding can be used to tunnel UDP-based latency-sensitive applications. Unless otherwise specified, the SSHv2 and SSH3 clients used in this section run in machines wired to the UCLouvain university campus network while the SSHv2 and SSH3 servers run on a server located in France with a 15 milliseconds round-trip time towards our clients.

### 5.1 Session establishment

We first evaluate the completion time of short non-interactive sessions. We define the *session completion time* as the time needed to establish a new session, run a single command, display the output to the user and close the session. Figure 3 illustrates the average session completion time for three different commands shown on the  $x$  axis, comparing SSH3 with the OpenSSH implementation of SSHv2. We also modified the source code of OpenSSH to provide *SSHv2-nodelay*, a variant enabling the TCP\_NODELAY socket option prior the SSH handshake to lower the handshake time of OpenSSH that is increased by the Nagle algorithm [9]. Each command is run 400 times with each solution. We show average values since the standard deviation is low for SSH3 (14.2) and SSHv2-nodelay (8.4) and hard to present graphically. It is significantly larger for the vanilla SSHv2 (more than 120 for each command) due to the Nagle algorithm inflating the completion time. The three commands have different output sizes on average: 582 bytes for `df`, 35kB for `sysctl` and 131kB for `ls`. As we can see, SSH3 drastically reduces the session completion time compared to SSHv2 and SSHv2-nodelay for every run command. This is mostly due to the significantly shorter session establishment time discussed in Section 3.1. SSHv2 consumes additional round-trips when opening new SSH channels as it has to negotiate channel flow control parameters before sending data. This can be avoided by SSH3 as it can reuse QUIC’s per-stream flow control mechanism. Finally, the executed command in itself has only a negligible impact on the session completion time as they all complete rapidly and have a relatively short output.

### 5.2 Terminal responsiveness

We now analyze the latency experienced using our SSH3 prototype and the OpenSSH implementation of SSHv2 in an interactive session. In this section, we want to show that SSH3 can at least provide a latency comparable to SSHv2. To rule out the impact of session establishment already discussed in Section 5.1, we perform our experiments over established sessions. The metric we evaluate is the *keystroke latency*, computed using `typometer` [14]. `typometer`




Figure 4: Keystroke latency in the Vim editor, either running locally (Local), over SSHv2 or over SSH3.

records the screen while writing keystrokes in an open text editor. The latency is the duration between the instant when a character was typed and the instant it is actually displayed on the screen. In these experiments, we compare the keystroke latency of the Vim text editor [34] opened on the remote host over SSHv2 and SSH3 sessions. Both OpenSSH and our SSH3 implementation only display a typed character when the server actually echoes it to the client. This means that the keystroke latency is impacted by the network latency. To study the impact of network latency on our scenarios, half of the experiments are performed towards our server in France and the other half are performed with our SSHv2 and SSH3 servers running on a Cloudlab server located in the United States, having a 170 milliseconds round-trip time with our clients. Figure 4 shows the keystroke latency distribution of SSH3 and SSHv2 when an interactive session is established. We do not experiment with our SSHv2-nodelay implementation here as OpenSSH enables the TCP\_NODELAY socket option for interactive sessions after the SSHv2 handshake. We report under the *Local* label the keystroke latency that `typometer` recorded with Vim running locally on the client, without SSH, as a baseline. It therefore measures the unavoidable latency induced by the graphical desktop environment and the text editor, also present in the SSHv2 and SSH3 curves. As we can see on the Figure, SSH3 shows a slightly lower keystroke latency than SSHv2. This is not surprising as the OpenSSH implementation being more than 20 years old, it includes a series of features and complex mechanisms inducing a non-negligible CPU overhead. The experiments performed with the server located in the USA help us to ensure that this difference between SSHv2 and SSH3 is not due to the network latency as the difference in keystroke latency between the two solutions stays similar between the US and France while the round-trip time is tenfold. Our prototype being significantly simpler than OpenSSH, we limit ourselves to concluding that SSH3 can provide a similar keystroke latency compared to SSHv2.

### 5.3 TCP and UDP port forwarding

We also perform experiments evaluating the TCP and UDP port forwarding features offered by SSH3. We first study the pure throughput that can be obtained through TCP and UDP port forwarding. We then show how real-time media content can be accessed using the UDP port forwarding feature of SSH3.

**5.3.1 Throughput tests.** We first analyze pure throughput tests. We run the `iperf3` tool between two c6525-25g Cloudlab servers (AMD 7302P 3GHz CPUs and 25Gbps network interfaces) with a round-trip time below 30 $\mu$ s. The UDP send and receive buffers are configured with a default size of 25MB to avoid packet loss due to buffer size. Without any tunnelling, `iperf3` reports 22Gbps of TCP throughput and 3.2Gbps of UDP throughput (no UDP packet loss was experienced). The TCP throughput falls down to 4.62Gbps when `iperf3` is tunnelled using OpenSSH port forwarding. We obtained 1.91Gbps of TCP throughput when using SSH3 TCP port forwarding. This result is not surprising since the QUIC and HTTP/3 stacks as well as the Go language are more computationally intensive than the mature OpenSSH implementation. Our prototype has not been optimized for throughput and there is room for improvement in that matter. The obtained UDP throughput was 583Mbps when tunnelled through UDP port forwarding, which is also significantly lower than the performance without tunnelling. This is however sufficient for a good part of UDP utilizations, including real-time UDP-based applications that generally favour latency to throughput. UDP port forwarding is simply not provided by SSHv2.

**5.3.2 UDP port forwarding and real-time streaming.** We finally study the use of latency-sensitive video applications through UDP port forwarding. We use FFmpeg [2] to send a real-time video stream from our client in Belgium to a GStreamer [3] receiver running on our server in France. The sent videos are remote drone piloting videos recorded by Baltaci *et al.* [7] encoded with real-time parameters. The GStreamer receiver is configured with a 20 milliseconds playback buffer to handle a potential small jitter induced by QUIC pacing. Frames lost or delayed by more than 20ms will therefore cause undesirable video artifacts on the playback. Before starting the stream, a 5MB file is sent to probe the link bandwidth. 200 video streams have been performed, half of it being tunnelled using the UDP port forwarding feature of SSH3. Every received video frame was then scored using the SSIM metric to assess its fidelity to the originally sent video frame. In every performed experiment, every frame of every sent video has been received in its entirety and the video has been displayed totally unaltered by the GStreamer receiver. The sent and received videos as well as packet captures are provided with the paper artifacts. This shows that SSH3 can be used for accessing low-latency media resources such as camera live recordings available on the remote server.

## 6 DISCUSSION

In this section, we discuss additional use-cases for SSH3.

### 6.1 Robustness to scanning and RST attacks

SSHv2 is subject to port scanning attacks like many TCP-based applications. Attackers can easily discover public SSH servers by scanning every TCP port and finding the ones answering to SSH session establishment. Once the attackers have discovered a public SSH endpoint, they can try dictionary attacks on passwords. Based over HTTP/3, SSH3 servers can avoid being publicly discovered by only answering to SSH3 clients putting a specific `<path>` value in their HTTP CONNECT request. Placing protected resources behind secret links like this is a common behaviour in web applications. It is however only complementary to the user authentication process.

Since SSHv2 runs above TCP, it is susceptible to attacks with spoofed TCP RST packets [6, 42]. Long-lived SSHv2 connections are particularly vulnerable to such packet-injection attacks. With SSHv2, the only possibilities to counter these attacks are to tunnel SSHv2 above IPsec or DTLS or use TCP-AO [40] or TCP MD5. Using a secure tunnel adds operational complexity and TCP-AO relies on pre-shared keys that are impossible to use at a large scale. As SSH3 runs above QUIC, it is not susceptible to these attacks since all QUIC packets and encrypted and authenticated.

### 6.2 Using SSH3 to monitor devices

The Simple Network Management Protocol (SNMP) is often used to monitor various types of devices. In a nutshell, the state of each device is exposed as a set of metrics collected in a Management Information Base (MIB) which can be queried using SNMPv3 [19]. The main usage of SNMP is that when a device receives an SNMP GET request for a specific variable, it returns its current value.

Could SSH3 replace SNMPv3 to query MIB variables? On many devices, the value of a MIB variable is usually the output of a single command. For example, the `sysUpTime` MIB variable is the uptime of the device measured in seconds. On Linux, the same information can be obtained by using the `uptime` command. With SSH3, a management station could simply execute this command on a remote device to query its uptime. Since QUIC supports 0-RTT, we could obtain this information in a single round-trip time by sending an HTTP/3 GET request containing the command to execute. This would not initiate a long-lasting SSH3 session as it is done using a CONNECT request but would still go through the SSH3 authentication process. Since HTTP/3 does not provide replay protection, we must ensure that the 0-RTT SSH3 connections are only used to carry idempotent requests. SSH3 could restrict the utilization of 0-RTT for specific users that only execute explicitly white-listed idempotent commands. 0-RTT would be disabled for other users and other commands. Another benefit of SSH3 from a device monitoring viewpoint is that `syslog` messages [21] can be transported by the SSH3 session as QUIC datagrams.

## 7 CONCLUSION

In this article, we presented SSH3, a new version of the SSH protocol rethinking its design with the modern features of the HTTP/3 protocol stack in mind. SSH3 relieves the design and implementation complexity of the protocol by reusing the secure mechanisms of TLS 1.3 and standard HTTP authentication mechanisms. We showed that it significantly reduces the connection establishment time compared to SSHv2, natively provides flexible and new ways to authenticate users and provides new features such as UDP port forwarding and connection migration with no compromise on the features proposed by its predecessor. This article is a first step towards SSH3. Rethinking SSH requires the feedback and support from the community. We therefore encourage thoughtful comments and collaborations to move forward and come up with a design document describing the protocol details of SSH3. We also plan on studying how SSH3 can be integrated to existing SSH implementations while maintaining and improving our prototype in the meantime. Our prototype and artefacts are publicly available [33].

## ACKNOWLEDGMENTS

We warmly thank Thomas Wirtgen for executing the typometer experiments on his computer. We also sincerely thank Jonathan Hoyland, Lucas Pardue and Olivier Pereira for reading and providing thoughtful comments on the article. We thank Marten Seemann and the contributors of `quic-go` for its API and HTTP abstractions making it easy and fast to implement our SSH3 prototype.

## REFERENCES

- [1] 1983. Telnet Protocol Specification. RFC 854. (May 1983). <https://doi.org/10.17487/RFC0854>
- [2] 2023. FFmpeg. (2023). <http://www.ffmpeg.org>.
- [3] 2023. GStreamer. (2023). <https://gstreamer.freedesktop.org/>.
- [4] 2023. Remote Desktop Protocol: Basic Connectivity and Graphics Remoting. (2023). [https://learn.microsoft.com/en-us/openspecs/windows\\_protocols/ms-rdpbcgr/5073f4ed-1e93-45e1-b039-6e30c385867c](https://learn.microsoft.com/en-us/openspecs/windows_protocols/ms-rdpbcgr/5073f4ed-1e93-45e1-b039-6e30c385867c)
- [5] Josh Aas, Richard Barnes, Benton Case, Zakir Durumeric, Peter Eckersley, Alan Flores-López, J Alex Halderman, Jacob Hoffman-Andrews, James Kasten, Eric Rescorla, et al. 2019. Let's Encrypt: an automated certificate authority to encrypt the entire web. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. 2473–2487.
- [6] Martin Arlitt and Carey Williamson. 2005. An analysis of TCP reset behaviour on the internet. *ACM SIGCOMM Computer Communication Review* 35, 1 (2005), 37–44.
- [7] Aygün Baltaci, Hendrik Cech, Nitinder Mohan, Fabien Geyer, Vaibhav Bajpai, Jörg Ott, and Dominic Schupke. 2022. Analyzing real-time video delivery over cellular networks for remote piloting aerial vehicles. In *Proceedings of the 22nd ACM Internet Measurement Conference*. 98–112.
- [8] Steven M Bellovin. 1989. Security problems in the TCP/IP protocol suite. *ACM SIGCOMM Computer Communication Review* 19, 2 (1989), 32–48.
- [9] Robert T. Braden. 1989. Requirements for Internet Hosts - Communication Layers. RFC 1122. (Oct. 1989). <https://doi.org/10.17487/RFC1122>
- [10] Brian Campbell and Mike Bishop. 2023. Client-Cert HTTP Header Field. RFC 9440. (July 2023). <https://doi.org/10.17487/RFC9440>
- [11] denis bider. 2020. *QUIC-based UDP Transport for Secure Shell (SSH)*. Internet-Draft draft-bider-ssh-quic-09. Internet Engineering Task Force. <https://datatracker.ietf.org/doc/draft-bider-ssh-quic/09/> Work in Progress.
- [12] Marlon Dutra. 2016. Scalable and secure access with SSH. (September 2016). <https://engineering.fb.com/2016/09/12/security/scalable-and-secure-access-with-ssh/>.
- [13] M. Seemann et al. 2023. A QUIC implementation in pure go. (2023). <https://github.com/quic-go/quic-go>
- [14] Pavel Fatin. 2023. typometer. (2023). <https://github.com/pavelfatin/typometer>
- [15] Roy T. Fielding, Mark Nottingham, and Julian Reschke. 2022. HTTP Semantics. RFC 9110. (June 2022). <https://doi.org/10.17487/RFC9110>
- [16] Roy T. Fielding and Julian Reschke. 2014. Hypertext Transfer Protocol (HTTP/1.1): Authentication. RFC 7235. (June 2014). <https://doi.org/10.17487/RFC7235>
- [17] Professor John Franks, Phillip Hallam-Baker, Lawrence C. Stewart, Jeffery L. Hostetler, Scott Lawrence, Paul J. Leach, and Ari Luotonen. 1999. HTTP Authentication: Basic and Digest Access Authentication. RFC 2617. (June 1999). <https://doi.org/10.17487/RFC2617>
- [18] Alan Frindell, Eric Kinnear, and Victor Vasiliev. 2023. *WebTransport over HTTP/3*. Internet-Draft draft-ietf-webtrans-http3-08. Internet Engineering Task Force. <https://datatracker.ietf.org/doc/draft-ietf-webtrans-http3/08/> Work in Progress.
- [19] Rob Frye, David B. Levi, Bert Wijnen, and Shawn A. Routhier. 2003. Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework. RFC 3584. (Aug. 2003). <https://doi.org/10.17487/RFC3584>
- [20] APNIC Blog Geoff Houston. [n. d.]. A look at QUIC use. <https://stats.labs.apnic.net/quic>. ([n. d.]). Accessed: 2022-02-09.
- [21] Rainer Gerhards. 2009. The Syslog Protocol. RFC 5424. (March 2009). <https://doi.org/10.17487/RFC5424>
- [22] Dick Hardt. 2012. The OAuth 2.0 Authorization Framework. RFC 6749. (Oct. 2012). <https://doi.org/10.17487/RFC6749>
- [23] Jana Iyengar and Martin Thomson. 2021. QUIC: A UDP-Based Multiplexed and Secure Transport. RFC 9000. (May 2021). <https://doi.org/10.17487/RFC9000>
- [24] Michael B. Jones, John Bradley, and Nat Sakimura. 2015. JSON Web Token (JWT). RFC 7519. (May 2015). <https://doi.org/10.17487/RFC7519>
- [25] Michael B. Jones and Dick Hardt. 2012. The OAuth 2.0 Authorization Framework: Bearer Token Usage. RFC 6750. (Oct. 2012). <https://doi.org/10.17487/RFC6750>
- [26] Brian Kantor. 1991. BSD Rlogin. RFC 1282. (Dec. 1991). <https://doi.org/10.17487/RFC1282>
- [27] Adam Langley, Alistair Riddoch, Alyssa Wilk, Antonio Vicente, Charles Krasnic, Dan Zhang, Fan Yang, Fedor Kouranov, Ian Swett, Janardhan Iyengar, et al. 2017. The quic transport protocol: Design and internet-scale deployment. In *Proceedings of the conference of the ACM special interest group on data communication*. 183–196.
- [28] Yanmei Liu, Yunfei Ma, Quentin De Coninck, Olivier Bonaventure, Christian Huitema, and Mirja Kühlewind. 2023. *Multipath Extension for QUIC*. Internet-Draft draft-ietf-quic-multipath-06. Internet Engineering Task Force. <https://datatracker.ietf.org/doc/draft-ietf-quic-multipath/06/> Work in Progress.
- [29] Chris M. Lonvick and Tatu Ylonen. 2006. The Secure Shell (SSH) Authentication Protocol. RFC 4252. (Jan. 2006). <https://doi.org/10.17487/RFC4252>
- [30] Chris M. Lonvick and Tatu Ylonen. 2006. The Secure Shell (SSH) Connection Protocol. RFC 4254. (Jan. 2006). <https://doi.org/10.17487/RFC4254>
- [31] Chris M. Lonvick and Tatu Ylonen. 2006. The Secure Shell (SSH) Protocol Architecture. RFC 4251. (Jan. 2006). <https://doi.org/10.17487/RFC4251>
- [32] Chris M. Lonvick and Tatu Ylonen. 2006. The Secure Shell (SSH) Transport Layer Protocol. RFC 4253. (Jan. 2006). <https://doi.org/10.17487/RFC4253>
- [33] François Michel. 2023. SSH3: Redesigning SSH over the HTTP/3 protocol stack. (2023). <https://github.com/francoismichel/ssh3>
- [34] Bram Moolenaar. [n. d.]. Vim, the editor. <https://www.vim.org/>. ([n. d.]). Accessed: 2023-11-30.
- [35] Tommy Pauly, Eric Kinnear, and David Schinazi. 2022. An Unreliable Datagram Extension to QUIC. RFC 9221. (March 2022). <https://doi.org/10.17487/RFC9221>
- [36] Eric Rescorla. 2018. The Transport Layer Security (TLS) Protocol Version 1.3. RFC 8446. (Aug. 2018). <https://doi.org/10.17487/RFC8446>
- [37] Natsuhiko Sakimura, John Bradley, Mike Jones, Breno De Medeiros, and Chuck Mortimore. 2014. Openid connect core 1.0. *The OpenID Foundation* (2014), S3. [https://openid.net/specs/openid-connect-core-1\\_0.html](https://openid.net/specs/openid-connect-core-1_0.html)
- [38] David Schinazi, David Oliver, and Jonathan Hoyland. 2023. *The Signature HTTP Authentication Scheme*. Internet-Draft draft-ietf-httpbis-unprompted-auth-05. Internet Engineering Task Force. <https://datatracker.ietf.org/doc/draft-ietf-httpbis-unprompted-auth/05/> Work in Progress.
- [39] Martin Thomson and Sean Turner. 2021. Using TLS to Secure QUIC. RFC 9001. (May 2021). <https://doi.org/10.17487/RFC9001>
- [40] Dr. Joseph D. Touch, Ron Bonica, and Allison J. Mankin. 2010. The TCP Authentication Option. RFC 5925. (June 2010). <https://doi.org/10.17487/RFC5925>
- [41] Manfred Touron. 2023. quicssh. (2023). <https://github.com/moul/quicssh>
- [42] Nicholas Weaver, Robin Sommer, and Vern Paxson. 2009. Detecting Forged TCP Reset Packets.. In *NDSS*.
- [43] Tatu Ylonen. 1996. SSH—secure login connections over the Internet. In *Proceedings of the 6th USENIX Security Symposium*, Vol. 37. 40–52.