

The KISS principle in Software-Defined Networking: a framework for secure communications

Diego Kreutz, Jiangshan Yu, Paulo Esteves-Verissimo

Cátia Magalhães, Fernando M. V. Ramos

Abstract—Security is an increasingly fundamental requirement in Software-Defined Networking (SDN). However, the pace of adoption of secure mechanisms has been slow, which we estimate to be a consequence of the performance overhead of traditional solutions and of the complexity of their support infrastructure.

To address these challenges we propose KISS, a secure SDN control plane communications architecture that includes innovative solutions in the context of key distribution and secure channel support. Core to our contribution is the integrated device verification value (iDVV), a deterministic but indistinguishable-from-random secret code generation protocol that allows local but synchronized generation/verification of keys at both ends of the control channel, even on a per-message basis.

We show that our solution, while offering the same security properties, outperforms reference alternatives, with performance improvements up to 30% over OpenSSL, and improvement in robustness based on a code footprint one order of magnitude smaller.

Keywords—*software-defined networking, SDN, security, cryptographic primitives, integrated device verification value (iDVV), perfect forward secrecy, performance, system architecture.*

I. INTRODUCTION

In Software-Defined Networking (SDN), network control is separated from the forwarding devices and logically centralised in a controller. This separation is achieved by means of a protocol (typically, OpenFlow) that enables the SDN controller to remotely populate the forwarding tables of network switches. The OpenFlow standard includes Transport Layer Security (TLS) (see IETF RFC 5246) as an *optional* security feature for authenticating forwarding devices and controllers and for encrypting the communication channel. However, to date most reported deployments still use TCP for control traffic, and SDN controllers and switching hardware with TLS support are still rare [1]. This makes the control plane communication vulnerable to different attacks [1], [2].

Four fundamental issues can slow down the rate of adoption of secure mechanisms in SDN. First, securing communications has a non-negligible cost in terms of increased communication latency and reduced performance. Several recent studies have analysed this overhead in various contexts [3]. Second, the computing capabilities of commodity switches are typically weak. The typical SDN switch is equipped with a single or dual-core CPU running at approximately 1GHz, which compares unfavourably with the multi-core CPUs found in typical commodity servers. Imposing the additional cost of TLS to these computing-constrained networking devices is a problem. Third, poor choice of cryptographic primitive implementations can also have a significant impact on the performance of

the control plane communications handled by the controller. Finally, the Public Key Infrastructure (PKI) on which TLS relies is complex and thus vulnerability prone [4], opening a large surface for successful attacks [5].

In order to meet these challenges, we propose a modular secure SDN control plane communications architecture KISS (Section II), which aims to increase the robustness of control communications whilst enhancing their performance, by decreasing the complexity of the support infrastructure, as an alternative to current approaches based on classic configurations of TLS and PKI.

A core novel component of our architecture is the integrated device verification value (iDVV), a deterministic but indistinguishable-from-random secret code generation protocol (Section III). The concept was inspired by the iCVVs (integrated card verification values) used in credit cards to authenticate and authorize transactions in a secure and inexpensive way. We develop and extend the idea for SDN, proposing a flexible method of generating iDVVs by adapting proven one-time password-like techniques. iDVV codes allow the safe decentralized generation/verification of keys at both ends of the channel, at will, even on a per-message basis.

To understand and minimize the cost of security, we quantify (Section IV) the impact of secure primitives on the performance and scalability of control plane communications, through a performance study of different implementations of TCP vs. TLS, complemented by a deeper study of underlying hashing and message authentication code (MAC) primitives. Those experiments confirm our intuition that the choice of protocols and primitives used in secure communication may well be one strong reason behind the slow adoption of these mechanisms in SDN. This in-depth study leads to the selection of the NaCl cryptographic library [6], and the best performing MAC and strong hash primitives — Poly1305 and SHA512 OpenSSL — as the baseline secure channel technologies for KISS.

iDVVs team-up with NaCl, in order to safely replace the cryptographic primitives and key-exchange protocols and key derivation functions commonly used in TLS. As a result, the NaCl-iDVV compound, while achieving the same functional level of security, is simpler, potentially leading to a higher level of implementation robustness by vulnerability reduction. In fact, we estimate the proposed security architecture footprint to be smaller than TLS-PKI alternatives with traditional protocols, by an order of magnitude, in terms of the number of lines of code (LOC). Such a differential also points to reducing the cyclomatic complexity. These metrics are typically used to assess the robustness and estimate verifiability of software systems.

Finally, in Section V we discuss aspects related to the security and robustness of our solution. We close the paper with related work and some directions to further work.

II. KISS ARCHITECTURE

In this section we present our proposal of KISS, a modular secure control plane communications architecture for SDN offering alternatives to classic configurations of secure channel and authentication protocols and subsystems followed in TLS and PKI. We assume a typical SDN architecture, as illustrated in Figure 1, composed of controllers and forwarding devices. We further assume that device registration and association services are in place. For lack of space, we do not discuss them in detail, but for self-containment, we discuss some properties and their interface below.

The two components encapsulated by the KISS boxes are the crucial components of the architecture, and the main subject of our study: a secure channel protocol suite, composed of a judicious choice of state-of-the-art mechanisms and protocols, which we dub SC for convenience of description, and a novel deterministic but indistinguishable-from-random secret code generation protocol, which we call iDVV.

We have considered using TLS implementations (e.g. OpenSSL) as the baseline protocol for SC. However, the experiments in Section IV have alerted us to: the sheer performance cost of cryptographic communication; and the further impact of sub-optimal choices of cryptographic primitives. This motivated us to adopt NaCl [6], a high performance yet secure cryptographic library, as the substrate of SC, complemented by the MAC and strong hash primitives with best performance according to our experiments – Poly1305 and SHA512 OpenSSL. SHA-512 is used by the iDVV generator while Poly1305 is a fast MAC algorithm.

The iDVV, a novel component we propose, helps to further enhance the security of SC, through strong crypto material generated at a low cost (e.g. one-time keys, per-message authentication and authorization codes) to be used by NaCl ciphers. The indistinguishability-from-random allied to the determinism allow the safe decentralized generation/verification of per-message keys at both ends of the channel.

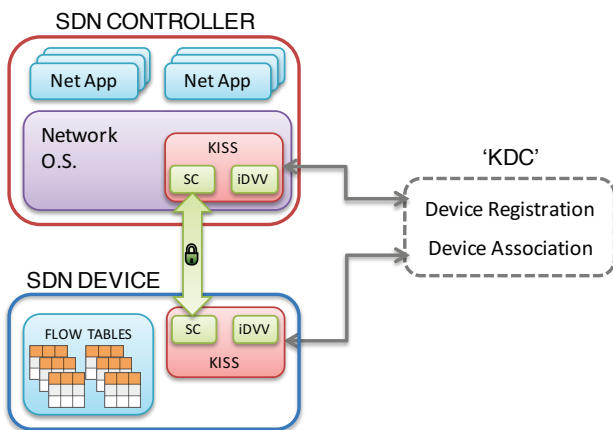


Fig. 1. General architecture

A. System and threat model

For simplicity and without loss of generality, we assume that the controllers and forwarding devices are registered and associated through a secure and robust key distribution service provided by a key distribution center (KDC), which for space reasons is out of the scope of this paper, but can be readily secured by state-of-the-art KDCs like Kerberos Key Distribution Center [7].

The device registration process is by default invoked by network administrators to the KDC, to register new devices. As the result of device registration, the device and the KDC securely share a symmetric key. We denote K_{kc} the shared key between the KDC authority and a registered controller, and K_{kf} the shared key between the KDC authority and a registered forwarding device.

Registered controllers and forwarding devices must be securely associated, also through the KDC authority, as a precondition to communicate securely. The most common case is a forwarding device f_i requesting an association to a controller c_j , through the KDC. After associating, a controller and a forwarding device share two symmetric secrets (of size 256 bits), namely a $seed_{ij}$ and a key_{ij} . The key is generated by the KDC and the seed is generated by the KDC in cooperation with the controller. These secrets will be used to bootstrap the iDVV module (see Section III-A).

As threat model, we consider a Dolev-Yao style attacker, who has complete control of the network, namely the attacker logs all messages, and can arbitrarily delay, drop, re-order, insert, or modify messages. In addition, this strong attacker is able to compromise any network device (e.g. a controller or a forwarding device) at any time. We assume the security of the used cryptographic primitives, including MAC (i.e. Poly1305), hash function (i.e. SHA-512), and symmetric encryption algorithm (e.g. AES).

B. Security goals

The main goal of KISS is to provide security properties including authenticity, integrity, and confidentiality for control plane communications, while minimizing cost and complexity.

The secure communication between participants can be easily guaranteed when a secure encryption algorithm is used, as long as the shared secret key is kept secure. To provide a robust SDN system, we focus on advanced security guarantees for the situation when the shared key is exposed to an attacker, as this might happen in practice. In particular, if an attacker has compromised a device and learnt its shared keys, then we are aiming at providing “perfect forward secrecy” (PFS) of communications. That is, the secrecy of a device’s past communications should be protected when the device is compromised and its shared keys are exposed to an attacker. It is important to emphasize that PFS is an essential requirement for SDN. The lack of it can lead to information disclosure, i.e., reveal different aspects of the network’s state and the controller’s strategy (e.g., proactive or reactive flow setup).

Established KDC technologies like Kerberos have robust implementations and are intensely used by industry, which makes us consider the logical single-point-of-failure they present as moderate, and an acceptable option for the current

state of the art. Even though, as we said, the KDC is out of the scope of the paper, we present mitigation measures to achieve PFS in case of compromise of the KDC. We also plan, as future work, to investigate the development of SDN KDCs resilient to accidental and malicious faults, drawing from fault and intrusion tolerance techniques [8].

On the devices side, we make no claim about their sheer resilience, since this is largely dependent on vendors. More precisely, when a controller and/or a forwarding device is compromised, we consider that the attacker is able to obtain all knowledge of the victim device(s), including all stored secrets and the session status. However, it is our goal to guarantee the confidentiality of all past communications through measures that allow us to achieve perfect forward secrecy.

III. iDVV: KEEP IT SIMPLE AND SECURE

Integrated device verification values (iDVs) are sequentially generated to protect and authenticate requests between two networking devices. The generator is conceived so that its output sequence has the indistinguishability-from-random and determinism properties. In consequence, the same sequence of random-looking secret values is generated on both ends of the channel, allowing the safe decentralized generation/verification of per-message keys at both ends. However, if the seed and key initial values and the state of the generator are kept secret, there is no way an adversary can know, predict or generate an iDVs.

In other words, an iDVs is a unique secret value generated by a device A (e.g. a forwarding device), which can be locally verified by another device B (e.g. a controller). The iDVs generation is made flexible to serve the needs of SDN. iDVs can therefore be generated: (a) on a per message basis; (b) for a sequence of messages; (c) for a specific interval of time; and (d) for one communication session. The main advantages of iDVs are their low cost and the fact that they can be generated locally, i.e., without having to establish any previous agreement.

Different from standard KDF algorithms such as HKDF, which assumes that keying material is not uniformly random or pseudorandom, our keying material (i.e. seed and key) are random symmetric secrets (each of size 256 bits), generated by the KDC, with high entropy. In such cases, a strong hash function can be safely used to derive a key (RFC 4880). As shown by the results in Section IV, the iDVs generation is simpler and faster than standard KDF algorithm such as HKDF (RFC 5869) and similar solutions.

A. iDVs bootstrap

As discussed before, the association between two SDN devices, e.g., forwarding device f_i and controller c_j , happens through the help of KDC, under the protection of the long-term secret keys obtained from registration (K_{kf} , resp. K_{kc}). The outcome of the association protocol is the distribution of two random secrets to both devices: a seed $seed_{ij}$, and an association key key_{ij} . The iDVs mechanism is bootstrapped by installing these two secret values in both the controller and the switch, to animate the iDVs generation algorithms, which we describe next.

Note that the set-up and generation of the iDVs values are performed in a deterministic way, so that they can be done locally at both ends. However, as iDVs will be used as keys by cryptographic primitives such as MAC or encryption functions, they have to be indistinguishable from random. Hashing primitives are natural choices for our algorithms, since they provide indistinguishable-from-random values if one or more of the input values are known only by the sender and the receiver. This explains why it is crucial that seed and association key are sent encrypted and therefore known only to the communicating devices. Moreover, in order to prevent information leakage, all variables $seed$, key , and $idvv$ in the algorithms below should have the same length, which we chose to be 256 bits in our design. This length is commonly considered robust. From our experiments discussed in Section IV, the hashing primitive to be used is SHA512, which yields 512 bits, of which we will use the most-significant q bits if we need to reduce the output length to q (as recommended by IETF RFC 4880). For example, we use the most-significant 256 bits of the SHA512 output as the key for symmetric ciphers.

The initial iDVs value is deterministically created at both ends of the association between two devices¹, by calling function `idvv_init`, which performs hashing on the concatenation of the initial $seed$ and key , as illustrated by algorithm 1. After set-up, the generator is ready for first use, as described in the following section.

Algorithm 1: iDVs set-up

```
1: idvv_init()
2: idvv ← H(seed || key)
```

B. iDVs generation

After the bootstrap with the initial $idvv$ value, the `idvv_next` function is invoked on-demand (again, synchronously at both ends of the channel) to autonomously generate authentication or encryption keys that will be used for securing the communications, as illustrated by algorithm 2.

The key remains the only constant shared secret between the devices. The $seed$ evolves to a new indistinguishable-from-random value each time `idvv_next` is invoked to generate a new iDVs. The new seed is the outcome of a hashing primitive H over the current $seed$ and current $idvv$ (line 2). The $new\ idvv$, output of function `idvv_next`, is the outcome of a hashing primitive H over the concatenation of the $new\ seed$ and association key key .

Algorithm 2: iDVs generation

```
1: idvv_next()
2: seed ← H(seed || idvv)
3: idvv ← H(seed || key)
```

C. iDVs synchronization

The iDVs mechanism is agnostic w.r.t. secure communication protocols, and can be used in a number of ways, in a number of protocols, as a key-per-message or key-per-session, etc. The only key issue about iDVs generation, is to keep it

¹For readability, we omit the device-identifying subscripts in the variables.

synchronized in both ends of the channel. So, we present some recommendations in this regard.

The most general style of iDVV use is *Indexed iDVV*: iDVs are indexed by the generation number, and they are operated in "one key per direction" mode, i.e., at each end, one iDVV is generated for each communication direction. This way, they support competitive, non-synchronized correspondents. This mode also supports unreliable, connectionless protocols like UDP. Each iDVV generated is indexed by a sequence number (the initial iDVV being $idvv^0$) and the sequence number is included in the message where the respective $idvv$ is used. This way, each receiving end (this works in either direction, as we have two pairs of iDVs) can know the exact $idvv$ number that should be used and, for example, detect and recover from omissions, by generating $idvv$'s the necessary number of times to resynchronize.

iDVs can get out of sync for a number of reasons, such as speed differences, omission errors, or even DoS attacks. When de-synchronization happens, a baseline technique consists of advancing the iDVV of the "slower" end, to catch up. The process is made robust by two techniques. First, communication should be authenticated (encrypt-then-MAC recommended), such that any messages failing crypto (decryption or MAC verification), can be simply discarded. Second, when say, $idvv^k$ is advanced to $idvv^l$ ($k < l$) to re-synchronize, and the operation is not successful (crypto fails), the old $idvv^k$ is restored (and the message motivating the recovery, is discarded, as per above). This restoration does not affect the PFS of communications because the $idvv^k$ (or newer) has not yet been used to secure the traffic between the two communicating devices. Finally, in the case of attacks, these robustness techniques also help to foil them, since the attacker cannot mimic valid crypto, so the message is discarded, and the node returns to the original iDVV state.

D. iDVV implementation and application

iDVs require minimal resources, which means that they can be implemented on any device, from a simple and very limited smart card to most existing devices. In other words, they are a simple and viable solution that can be embedded in any networking device. Just three values per association have to be securely stored — the seed, the association key and the iDVV itself — in order to use iDVV continuously. Furthermore, only hash functions, simple to implement and with a very small code base, are required to generate iDVs. Such kind of resource is already available on all networking devices that support traditional network protocols and basic security mechanisms.

We advocate (and demonstrate in Section IV-B) that iDVs are inexpensive and, as a result, can be used on a per-message basis to secure communication. It is worth emphasizing that, from a security perspective, one fresh iDVV per message makes it much harder for attacks such as key recovery, advanced side channel attacks, among other general HMAC attacks, to succeed. In fact, the one-time key approach was initially used for generating MACs. Yet, it was set aside (i.e. replaced by keys with a longer lifetime) due to performance reasons. However, as the iDVV generation has a low cost, we incur a lower penalty.

Finally, iDVs can have further practical applications. For instance, the TLS handshake can be used to bootstrap the iDVV. After that, iDVs can be used as session keys, i.e., in security mechanisms such as encrypt-then-MAC.

IV. ON THE COST OF SECURITY

In this section we provide a quantitative analysis of the impact of cryptographic primitives on control plane communication. Although the number of use cases is expanding, SDN has been mainly targeting data centers. As such, SDN controllers have to be capable of dealing with the challenging workloads of these large-scale infrastructures. In these environments new flows² can arrive at a given forwarding device every 10 μ s, with a great majority of mice traffic lasting less than 100ms [9]. This means that current data centers need to handle peak loads of tens of millions of new flows/s. The control plane has to meet both the network latencies and throughputs required to sustain these high rates. Current controllers are capable of achieving a throughput of up to 20M flows/s using TCP [1].

So any effort to systematically secure control plane communications has to meet these challenges. In the following we try to put the problem in perspective, by analysing the effect of including even the most basic security primitives to ensure authenticity, confidentiality and integrity when considering peak loads of this magnitude. We start by analyzing the latency impact of TLS, relative to TCP, and then we focus on hashes and MACs as they are the essential primitives for authenticity and integrity of communication.

A. The cost of secure channels

Our first experiments assess the compared average latency of TCP and TLS on control plane communication. We analyse the latency of connection setup and of OpenFlow PACKET_IN/FLOW_MOD messages. The OpenFlow PACKET_IN message is used by switches to send packets to the controller (e.g. when there is no rule matching the packet received in the switch). FLOW_MOD messages allow the controller to modify the state of an OpenFlow switch.

The connection setup time for TLS is two orders of magnitude higher than for TCP, since TLS has a more elaborate handshake protocol between the devices [10]. Also, PolarSSL (a library used in systems from companies such as Gemalto, ARM, and Linksys) induces nearly twice the overhead of OpenSSL. However important, a high connection cost can be amortized by maintaining persistent connections. As such, we focus on the communications cost. Figure 2 shows the latency of FLOW_MOD messages, averaged over 10k messages. The results with PACKET_IN messages were similar so we omit them for clarity. The costs of TCP, OpenSSL and PolarSSL grow nearly linearly with the number of forwarding devices. OpenSSL latency is approximately 3x higher than TCP. This is explained by the high overhead of cryptographic primitives, as we further analyse in the next section. PolarSSL is significantly worse, increasing the latency by up to 7x when compared with TCP.

Conclusions: The main findings of this analysis can be summarised in two points. First, different implementations of

²In spite of the fact that there are several definitions of flow in SDN [1], we equate SDN flow with TCP flow for the sake of simplicity.

TLS present very different performance penalties. Second, the additional computation required by the cryptographic primitives used in TLS leads to a non-negligible performance penalty in the control plane. In consequence, we turn to lightweight cryptographic libraries, such as NaCl [6] and TweetNaCl [11], which are starting to be used in different applications. NaCl has been designed to be secure [12], [6] and to be embedded in any system [11], taking a clean slate approach and avoiding most of the pitfalls of other libraries (e.g. OpenSSL – misuse issues). First, it exposes a simple and high-level API, with a reduced set of functions for each operation. Second, it uses high-speed and highly-secure primitives, carefully implemented to avoid side-channel attacks. Third, NaCl is less error-prone because low-security options are eliminated and it also provides a limited number of cryptographic primitives. In other words, users do not need deep knowledge regarding security to use it correctly. This is one of the major differences between it and other libraries such as OpenSSL. For instance, it has been recurrently shown that developers have been using OpenSSL in incorrect ways, leading to several security issues. Fourth, it has already been shown that secure and high-performance network protocols, outperforming OpenSSL, can be designed and implemented using NaCl [13].

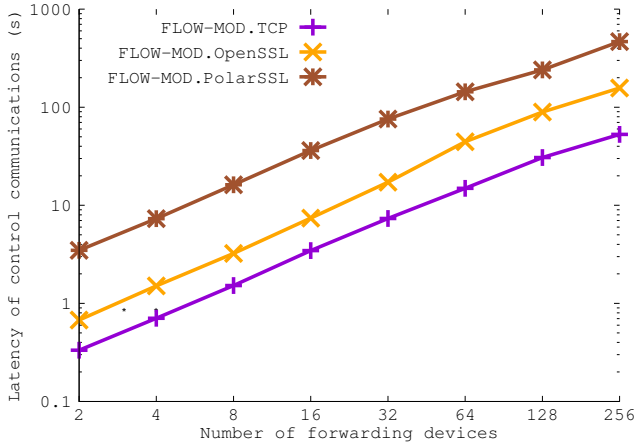


Fig. 2. FLOW_MOD latency (in log scale)

B. A closer look at the cost of cryptography

To understand in more detail the cause of the previous findings we now perform a fine-grained analysis of two main classes of security primitives used in secure channel protocols: hashing and MAC.

We analyse the performance of nine hashing primitives. The results are presented in Figure 3. The red bars represent primitives that are provided by OpenSSL, while white bars (BLAKE and KECCAK) indicate the original implementation of primitives that are not part of OpenSSL. From Figure 3, we observe that the primitives with smaller digest sizes (SHA-1 and MD5) achieve better performance, as expected. The stronger versions of the SHA and BLAKE families achieve comparable performance (slightly slower), with higher security guarantees. Interestingly, SHA-512 outperforms SHA-256. This behavior is explained by the fact that on a 64-bit processor each round can process twice as much data (64-bit words

instead of 32-bit words). However, SHA-256 is faster on a 32-bit processor.

To understand the variance between different implementations, we present in Figure 4 the costs of the five hashing primitives for which different implementations were available. The OpenSSL implementation shows the best performance for hashing primitives. With the exception of RIPEMD160, the PolarSSL implementation always presented higher message latencies.

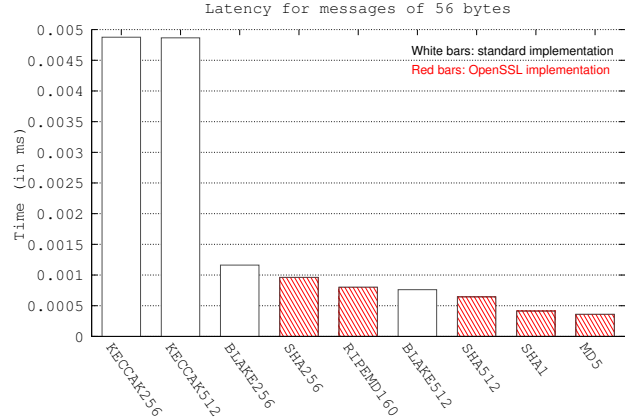


Fig. 3. Hashing primitives

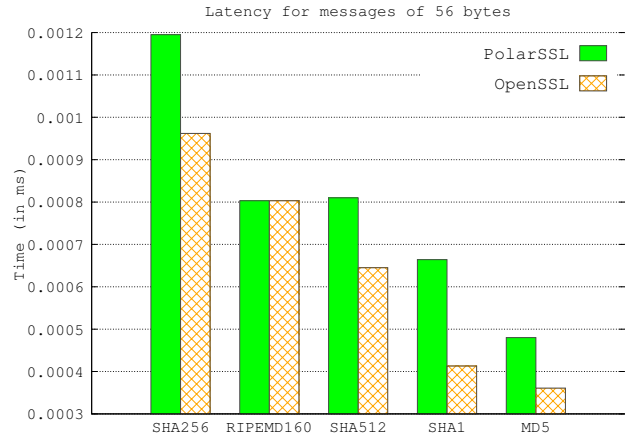


Fig. 4. Implementations of hashing primitives

Finally, Figure 5 shows the results of the latency analysis of six MAC primitives. It is clear that Poly1305 outperformed all other primitives, being approximately two times faster than OpenSSL’s HMAC-SHA1, and close to four times faster than HMAC-SHA512, for instance.

Conclusions: From the results of Figure 5, considering the MAC primitive with best performance in the analysis (Poly1305 with 0.001ms per message), around 20 dedicated cores are needed to compute a MAC in order to maintain a rate of 20M flows/s. To understand the importance of judiciously selecting the security primitives implementation, the HMAC-SHA512 OpenSSL (worst case performance in the analysis) would require over three times more cores (up to 65) to compute MACs at these rates. From the hashing primitive




Fig. 5. MAC primitives

analysis in Figures 3 and 4, of the strong primitives (i.e. all except SHA1 and MD5), SHA-512 performs the best. However, concerning MAC primitives, the performance of HMAC-SHA512 disappoints, and it is clear that Poly1305 outperformed all other primitives, providing security with high speed and low per-message overhead.

Figure 6 shows the performance of different primitives for generating cryptographic material. We compare the iDVV generator using SHA512 (iDVV-S5), with an implementation of a common key derivation function (KDFx) with different values for the exponent c (128, 64, 32, and 16, respectively), the Diffie-Hellman implementation used by OpenSSL (DH-OSSL), and the `randombytes()` function (NaCl-R) provided by NaCl. The latencies of the several primitives are significantly higher than iDVV. Even the `randombytes()` primitive of NaCl, the second fastest after iDVV, still presents a latency at least 2.6x higher.

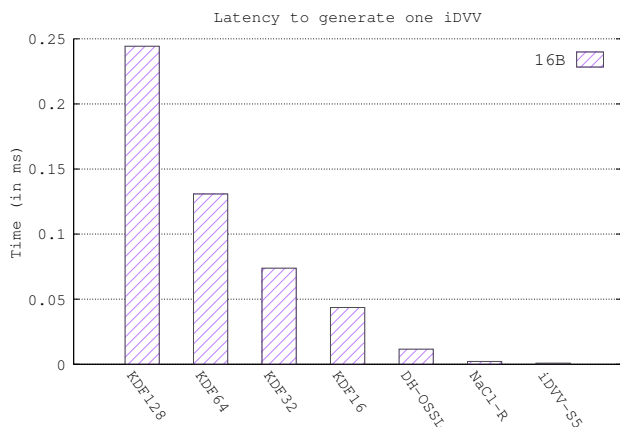


Fig. 6. Latency to generate keys

In summary, our findings in this section indicate that (i) the inclusion of cryptographic primitives results in a non-negligible performance impact on the latency and throughput of the control plane; and that (ii) a careful choice of the primitives used and their respective implementations can significantly contribute to reduce this performance penalty and enable feasible solutions in certain scenarios. Taking

the outcome of our analysis into consideration, and given the benefits of NaCl described in Section IV-A, we have selected the NaCl lightweight cryptographic library, and the MAC and strong hash primitives with best performance – Poly1305 and SHA512 OpenSSL – as the baseline SC secure channel component technologies. NaCl is complemented in our architecture with the iDVV mechanism to generate crypto material (e.g. keys) used by NaCl ciphers. Taken together they provide, as per our evaluation, the best trade-off between security and performance for control plane communications in SDN.

V. DISCUSSION

A. On the security of iDVV

With respect to the secrecy of iDVs, it is ensured from initialization and so long as neither the KDC, controller, nor forwarding device is compromised. Our scheme also achieves perfect forward secrecy in the face of compromise of either KDC, controller, or forwarding device. In short, when the KDC is compromised, then the attacker would be able to obtain all the shared secrets (between the authority and registered devices), decrypt the past communication that delivered the initial *seed* and *key* to the associated devices, re-generate iDVs and, in consequence, decrypt past conversations.

We provide a simple mechanism for providing PFS even when the authority is compromised: we update the shared key each time a forwarding device is associated with a controller. The key is updated as follows: $K_{kc} \leftarrow H(K_{kc})$ and $K_{kf} \leftarrow H(K_{kf})$. This way, a shared key captured cannot decrypt any past messages, since they have been encrypted with previous generations of that key, which have been “forgotten” in the system, given the irreversible nature of hashes.

As far as devices are concerned, when they are compromised, the current values of *seed*, *key* and *idvv* are captured. Note that *key* stays as the original secret, but *seed* is rolled forward everytime a new iDVV is generated. So, the attacker will be unable to synthesize any past iDVs since day one and so, cannot decrypt past conversations, achieving PFS, as we desired.

B. On the solution robustness

Our proposal compares well with traditional solutions such as EJBCA (<http://www.ejbca.org/>) and OpenSSL, two popular implementations of PKI and TLS, respectively.

The first interesting take away is that our solution has nearly one order of magnitude less LOC (85k) and uses four times less external libraries and only four programming languages. This makes a huge difference from a security and dependability perspective. For instance, to formally prove more than 717k LOC (OpenSSL + EJBCA) is by itself a tremendous challenge. And it gets considerably worse if we take into account eighty external libraries and eleven development languages. Moreover, it is worth emphasizing that libraries such as OpenSSL suffer from different fundamental issues such as too many legacy features accumulated over time, too many alternative modes as result of tradeoffs made in the standardization, and too much focus on web and DNS names.

Second, OpenSSL is complex and highly configurable. This has been also the source of many security incidents, i.e., developers and users frequently use the library in an inappropriate way. It has also been shown that the majority of the security incidents are still caused by errors and misconfiguration of systems. Lastly, recent research has uncovered new vulnerabilities on TLS implementations [14].

In contrast, our proposed architecture exhibits gains in both performance and robustness, contributing to solving the dilemma we outlined in the introduction. By having less LOC, we significantly reduce the threat surface – by one order of magnitude – and by combining NaCl and the iDVV mechanism, we provide a potentially equivalent level of security, but quite increased performance/robustness product, as keys can be rolled even on a per message basis.

C. On the cost of iDVV

Similarly to iCVVs, iDVVs are a low overhead solution that requires minimal resources. This solution is thus feasible to be integrated into compute-constrained devices as commodity switches. Our preliminary evaluation has revealed that the iDVV mechanism is faster than traditional solutions, namely, the key-exchange algorithms embedded in the OpenSSL implementation. Considering a setup with 128 switching devices, sending `PACKET_IN` messages to and receiving `FLOW_MOD` messages from the controller, our results shows our proposed solution (iDVV + NaCl's ciphers) to be more than 30% faster than an OpenSSL-based implementation using AES256-SHA (the most common high performance cipher suite, used by IT companies such as Google, Facebook, Microsoft, and Amazon). Importantly, we were able to outperform OpenSSL-based deployments while still providing the same security properties: authenticity, integrity, and confidentiality. In addition, we achieved this result not only while offering the same properties, but also with stronger security guarantees: the tests were made by generating one iDVV *per packet*, while the OpenSSL-based implementation uses a single key (for symmetric ciphering) for the entire communication session.

VI. RELATED WORK

There are several feasible attacks against the SDN control plane [2]. Most of them explore vulnerabilities such as the lack of authentication, authorization and other essential security properties. However, almost no attention has been paid to the security requirements of control plane associations and communication between devices. For instance, only recently, the use of secrecy through obscurity has been proposed to protect SDN controllers from DoS attacks [15]. In this case, the switch authentication ID is hidden in a specific field in the IP protocol. It is assumed that the devices share a look-up table and unique IDs. However, in spite of being capable of mitigating DoS attacks, this technique does not address the security issues of control plane communications.

VII. CONCLUDING REMARKS

In this paper, we set out to explore and confirm our intuition for the possible reasons behind a slower than expected adoption of security mechanisms in SDN, and based on those findings,

we proposed KISS, a modular secure SDN control plane communications architecture.

We started by investigating the impact of essential cryptographic primitives and TLS implementations on the control plane performance. We showed that whilst even the most basic security primitives add a non-negligible degradation of performance, a judicious choice of these primitives and their specific implementations can mitigate the penalty significantly. This is particularly important for the typical SDN scenario that resorts to commodity hardware, sometimes with modest computing capabilities.

The second problem we explored in this paper was the complexity of the centralized support infrastructure for authentication and key distribution. We proposed iDVV, a simple and robust decentralized mechanism for generating and verifying the secrets necessary for secure communications between network devices. As future work, we are also investigating the reduction of single-point-of-failure syndromes: architectures for SDN KDCs resilient to accidental and malicious faults, drawing from fault and intrusion tolerance techniques.

Our results are encouraging in terms of an increase of performance — 30% improvement over OpenSSL — and robustness — an order of magnitude reduction in the number of LOC, and implied cyclomatic complexity. This also means that formal verification is more tractable, which is one of our future goals for iDVV, for instance.

An extended report of this work can be found in [10], extending discussion of the iDVV performance, forward secrecy, randomness, and proofs of its security properties. As ongoing work we are extensively evaluating the NaCl-iDVV compound. Our results are clear in showing the solution to outperform OpenSSL on control plane communications.

We believe that this is one first step towards lightweight but effective security for control plane communication, and potentially for SDN in general. We make a “call to arms” to foster developments on securing SDN communications without impairing performance, a fundamental pre-condition for widespread adoption by future SDN deployments.

ACKNOWLEDGMENT

We would like to thank the anonymous reviewers and Christian Esteve Rothenberg and Marcus Völp for the insightful comments. This work is partially supported by the Fonds National de la Recherche Luxembourg (FNR) through PEARL grant FNR/P14/8149128, by European Commission funds through the H2020 programme, namely by funding of the SUPERCLOUD project, ref. H2020-643964, and by Portuguese national funds through Fundação para a Ciência e a Tecnologia (FCT), namely by funding of LaSIGE Research Unit, ref. UID/CEC/00408/2013.

REFERENCES

- [1] D. Kreutz, F. Ramos, P. Esteves Verissimo, C. Esteve Rothenberg, S. Azodolmolky, and S. Uhlig, “Software-defined networking: A comprehensive survey,” *Proceedings of the IEEE*, vol. 103, no. 1, Jan 2015.
- [2] S. Scott-Hayward, S. Natarajan, and S. Sezer, “A survey of security in software defined networks,” *IEEE Communications Surveys Tutorials*, vol. 18, no. 1, pp. 623–654, Firstquarter 2016.

- [3] D. Naylor, A. Finamore, I. Leontiadis, Y. Grunenberger, M. Mellia, M. Munafo, K. Papagiannaki, and P. Steenkiste, "The cost of the "S" in HTTPS," in *Proceedings of the Tenth ACM Conference on Emerging Networking Experiments and Technologies*, ser. CoNEXT '14. New York, NY, USA: ACM, 2014, p. 7.
- [4] A. Wazan, R. Laborde, F. Barrere, A. Benzekri, and D. Chadwick, "PKI interoperability: Still an issue? a solution in the x.509 realm," in *Info. Assurance and Sec. Education and Training*. Springer, 2013, vol. 406.
- [5] N. van der Meulen, "DigiNotar: Dissecting the first dutch digital disaster," *Journal of Strategic Security*, vol. 6, no. 2, 2013.
- [6] D. Bernstein, T. Lange, and P. Schwabe, "The security impact of a new cryptographic library," in *Progress in Cryptology - LATINCRYPT*, ser. L. N. in CS. Springer, 2012, vol. 7533.
- [7] B. C. Neuman and T. Ts'o, "Kerberos: An authentication service for computer networks," *IEEE Communications magazine*, vol. 32, no. 9, pp. 33–38, 1994.
- [8] P. Verissimo, M. Correia, N. F. Neves, and P. Sousa, "Intrusion-resilient middleware design and validation," in *Information Assurance, Security and Privacy Services*, ser. Handbooks in Information Systems. Emerald Group Publishing Limited, May 2009, vol. 4, pp. 615–678.
- [9] T. Benson, A. Akella, and D. A. Maltz, "Network traffic characteristics of data centers in the wild," in *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement*, ser. IMC '10. New York, NY, USA: ACM, 2010, pp. 267–280.
- [10] D. Kreutz, P. Esteves-Verissimo, C. Magalhaes, and F. M. V. Ramos, "The KISS principle in Software-Defined Networking: An architecture for Keeping It Simple and Secure," *ArXiv e-prints*, Jun. 2017. [Online]. Available: <https://arxiv.org/abs/1702.04294>
- [11] D. Bernstein, B. van Gastel, W. Janssen, T. Lange, P. Schwabe, and S. Smetsers, "TweetNaCl: A crypto library in 100 tweets," in *Progress in Cryptology - LATINCRYPT 2014*, ser. Lecture Notes in Computer Science, D. F. Aranha and A. Menezes, Eds. Springer International Publishing, 2015, vol. 8895, pp. 64–83.
- [12] J. B. Almeida, M. Barbosa, J. S. Pinto, and B. Vieira, "Formal verification of side-channel countermeasures using self-composition," *Science of Computer Programming*, vol. 78, no. 7, pp. 796 – 812, 2013, special section on Formal Methods for Industrial Critical Systems (FMICS 2009 + FMICS 2010) & Special section on Object-Oriented Programming and Systems (OOPS 2009), a special track at the 24th ACM Symposium on Applied Computing.
- [13] W. M. Petullo, X. Zhang, J. A. Solworth, D. J. Bernstein, and T. Lange, "MinimalLT: Minimal-latency networking through better security," in *Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security*, ser. CCS '13. New York, NY, USA: ACM, 2013, pp. 425–438. [Online]. Available: <http://doi.acm.org/10.1145/2508859.2516737>
- [14] B. Beurdouche, K. Bhargavan, A. Delignat-Lavaud, C. Fournet, M. Kohlweiss, A. Pironti, P.-Y. Strub, and J. K. Zinzindohoue, "A messy state of the union: Taming the composite state machines of TLS," in *2015 IEEE Symposium on Security and Privacy*. IEEE, 2015, pp. 535–552.
- [15] O. I. Abdullaziz, Y. J. Chen, and L. C. Wang, "Lightweight authentication mechanism for software defined network using information hiding," in *2016 IEEE Global Communications Conference (GLOBECOM)*, Dec 2016, pp. 1–6.

Diego Kreutz is a member of the CritiX group at the Interdisciplinary Centre for Security, Reliability and Trust (SnT), working toward the Ph.D. degree in computer science at the University of Luxembourg, and an Adjunct Professor at Federal University of Pampa. His main research interests are in software-defined networking, intrusion tolerance, system security and dependability, and cloud computing. He is also member of learned societies such as IEEE, IEEE Computer Society, ACM, and SIGOPS. Contact him at diego.kreutz@uni.lu.

Jiangshan Yu is currently a research fellow at the Interdisciplinary Centre for Security, Reliability and Trust (SnT), University of Luxembourg, and an honorary research fellow at the University of Birmingham, where he received his Ph.D

in computer science. The focus of his research has been on applied cryptography, post-compromised security, cryptographic key management, and blockchains. Contact him at jiangshan.yu@uni.lu.

Paulo Esteves-Verissimo is a Professor of the Univ. of Luxembourg, and head of the CritiX research lab at UL's SnT Centre. Previously, he has been with the Univ. of Lisbon. He is Fellow of IEEE and of ACM, Chair of the IFIP WG 10.4 on Dependable Computing and Fault-Tolerance and vice-Chair of the Steering Committee of the DSN conference, as well as associate editor of the IEEE Transactions on Computers. He is author of over 180 peer-refereed publications and co-author of 5 books. Contact him at paulo.verissimo@uni.lu.

Cátia Magalhães received her Computer Science degree and MSc degree in Computer Science from University of Lisbon, Faculty of Sciences in 2013 and 2015, respectively. Currently, she is working in the Digital Marketing area as Software Developer. For the past two years, she has worked with Java technologies and has participated in development projects, mainly in banking and telcos areas. Contact her at catiamagalhaes27@gmail.com.

Fernando Ramos is an assistant professor in the Department of Computer Science and Engineering at the Faculty of Sciences of the University of Lisbon. His research interests lie at the intersection of networking, systems, and security. Fernando has a PhD in computer science from the University of Cambridge. He is a member of the ACM, the IEEE, and the MEF research council. Contact him at fvramos@ciencias.ulisboa.pt.