DEMO: VPKIaaS: A Highly-Available and Dynamically-Scalable Vehicular Public-Key Infrastructure

Hamid Noroozi
Networked Systems Security Group
KTH, Stockholm, Sweden
hnoroozi@kth.se

Mohammad Khodaei
Networked Systems Security Group
KTH, Stockholm, Sweden
khodaei@kth.se

Panos Papadimitratos
Networked Systems Security Group
KTH, Stockholm, Sweden
papadim@kth.se

ABSTRACT

The central building block of secure and privacy-preserving Vehicular Communication (VC) systems is a Vehicular Public-Key Infrastructure (VPKI), which provides vehicles with multiple anonymized credentials, termed *pseudonyms*. These pseudonyms are used to ensure message authenticity and integrity while preserving vehicle (and thus passenger) privacy. In the light of emerging large-scale multi-domain VC environments, the efficiency of the VPKI and, more broadly, its scalability are paramount. In this extended abstract, we leverage the state-of-the-art VPKI system and *enhance* its functionality towards a highly-available and dynamically-scalable design; this ensures that the system remains operational in the presence of benign failures or any resource depletion attack, and that it dynamically *scales out*, or possibly *scales in*, according to the requests' arrival rate. Our full-blown implementation on the Google Cloud Platform shows that deploying a VPKI for a large-scale scenario can be cost-effective, while efficiently issuing pseudonyms for the requesters.

KEYWORDS

VPKI, Identity and Credential Management, Security, Privacy, Availability, Scalability, Micro-service, Container Orchestration, Cloud

1 INTRODUCTION

In Vehicular Communication (VC) systems, vehicles beacon Cooperative Awareness Messages (CAMs) periodically, at a high rate, to enable transportation safety and efficiency. Vehicle-to-Vehicle (V2V)/Vehicle-to-Infrastructure (V2I) (V2X) communication is protected with the help of Public Key Cryptography: a set of short-lived anonymized certificates, termed *pseudonyms*, are issued by a Vehicular Public-Key Infrastructure (VPKI), e.g., [7], for registered vehicles. Vehicles switch from one pseudonym to a non-previously used one towards message unlinkability as pseudonyms are per se inherently unlinkable.

With emerging large-scale multi-domain VC environments, the efficiency of the VPKI and, more broadly, its scalability are paramount. Deploying a VPKI differs from a traditional PKI in different aspects. One of the most important factors is the dimension of the PKI, i.e., the number of registered ''users'' (vehicles) and the multiplicity of certificates per user. According to the US Department of Transportation (DoT), a VPKI should be able to issue pseudonyms for more that 350 million vehicles across the Nation [1]. Considering the average daily commute time to be 1 hour [1] and a pseudonym lifetime of 5 min, the VPKI should be able to issue at least 1.5×10^{12} pseudonyms per year[1], i.e., 5 orders of magnitude more than what the largest current PKI issues (10 million certificates per year [10]).

Each vehicle is expected to interact with the VPKI regularly, e.g., once or a few times per day, not only to refill its pseudonym pool, but also to fetch the latest revocation information. As shown in [5, 7], the performance of a VPKI system can be drastically degraded under a clogging Denial of Service (DoS) attack: adversaries could compromise the availability of the VPKI entities with spurious requests. The cost of unavailability[2] is twofold: security (degradation of road safety) and privacy. An active malicious entity could prevent other vehicles from accessing the VPKI to fetch the latest revocation information. Moreover, signing CAMs with the private keys corresponding to expired pseudonyms, or the Long Term Certificate (LTC), is insecure and harms user privacy. Even though one can refill its pseudonym pool by relying on other anonymous authentication primitives, e.g., [8], the performance of the safety-related applications could be degraded if the majority of vehicles leverage such schemes, i.e., causing 30% increase in cryptographic processing overhead in order to validate CAMs [8].

In this work, we leverage and *enhance* the state-of-the-art VPKI towards a highly-available, dynamically-scalable, and fault-tolerant (highly-reliable) design to ensure that the system remains operational in the presence of benign failures or any resource depletion attack (clogging DoS). Moreover, we show how to dynamically scale out, or possibly scale in[3], based on the workload on the VPKI system, so that it can comfortably handle any demanding load while being cost-effective by systematically allocating/deallocating resources.

2 VPKI as a SERVICE (VPKIaaS)

We leverage the state-of-the-art VPKI system [7] that provides pseudonyms in an *on-demand* fashion: each vehicle *''decides''* when to trigger the pseudonym acquisition process based on various factors [6]. Pseudonyms have a lifetime (a validity period), typically ranging from minutes to hours; in principle, the shorter the pseudonym lifetime (τ_P) is, the higher the unlinkability and thus the higher the privacy protection that can be achieved.

The VPKI consists of a set of Certification Authorities (CAs) with distinct roles: the Root CA (RCA), the highest-level authority, certifies other lower-level authorities; the Long Term CA (LTCA)

[1] Note that this number could be even greater by considering the envisioned vehicular ecosystem, i.e., pedestrian and cyclist being part of the Intelligent Transport Systems (ITSs) with a gamut of services, e.g., Location Based Services (LBSs).
[2] Note that the VPKI could be unreachable for other reasons, e.g., intermittent coverage of sparsely-deployed Roadside Units (RSUs), that are orthogonal to this investigation.
[3] Scaling in/out, termed *horizontal* scaling, refers to replicating a new instance of a service, while scaling up/down, termed *vertical* scaling, refers to allocating/deallocating resources for an instance of a given service.

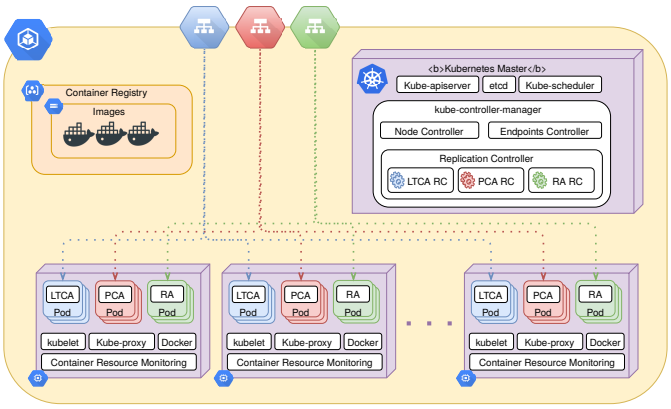

Figure 1: A high-level VPKIaaS architecture.

is responsible for the vehicle registration, the Long Term Certificate (LTC) issuance, as well as (authorization) *ticket* issuance, used for obtaining pseudonyms. The Pseudonym CA (PCA) issues pseudonyms for the registered vehicles and the Resolution Authority (RA) can initiate a process to resolve and revoke all pseudonyms of a misbehaving vehicle. Vehicles can cross into other *domains* [9]; trust between two domains can be established with the help of an RCA, or through cross certification between them [9]. The efficiency and robustness of the VPKI system is systematically investigated in [6, 7] and the VPKI can handle large workloads. A detailed protocol description can be found in [5, 7].

Towards ensuring viability as VC systems grow, we deploy the VPKI on the Google Cloud Platform (GCP) (cloud.google.com), and evaluate the availability, reliability, and dynamic scalability of our scheme under various circumstances. Fig. 1 illustrates a high-level abstraction of the VPKIaaS architecture on a managed Kubernetes cluster (kubernetes.io) on GCP.[4] A set of Pods will be created for each micro-service, e.g., LTCA or PCA, from their corresponding container images, facilitating their horizontal scalability. When the rate of pseudonym requests increases, the Kubernetes master, shown on the top, schedules new Pods or kills a running Pod in case of benign failures, e.g., system faults or crashes, or resource depletion attacks, e.g., a DoS attack. The Pods could be scaled out to the number set in the deployment configuration, or the amount of available resources enabled by Kubernetes nodes.

Each Pod publishes two types of metrics: *load* and *health*. The load metric values are generated by a resource monitoring service, which facilitates horizontal scaling of a micro-service, i.e., upon reaching a threshold of a defined load, replication controller replicates a new instance of the micro-service to ensure a desired Service Level Agreement (SLA). Health metric ensures correct operation of a micro-service by persistently monitoring its status: a faulty Pod is killed and a new one will be created. In our system, we defined CPU and memory usage as the load metric. In order to monitor the health condition of a micro-service, dummy requests (tickets for LTCA micro-services and pseudonyms for PCA micro-services) are queried (locally by each micro-service).

[4]The RCA entity is assumed to be off-line, thus not included in the architecture.

Note: Multiple replicas of a micro-service interact with the same database to accomplish their operations, e.g., all replicas of PCAs interact with a single database to validate an authorization ticket and store information corresponding to the issued pseudonyms. This could be a bottleneck in our architecture, possibly a single point of failure; how to mitigate this would be part of our future investigation. Moreover, the information corresponding to the issued pseudonyms is stored asynchronously, i.e., a PCA micro-service delivers the pseudonym response without confirmation of its successful storage in the database. If there are multiple replicas of a micro-service, e.g., a PCA, a *''malicious''* vehicle, repeatedly requesting to obtain pseudonyms, might be provided with more than a set of pseudonyms per ticket.[5] As future work, we plan to utilize alternative storage solutions, e.g., *NoSQL* databases, and apply zonal resource synchronization to prevent issuing multiple pseudonyms per ticket, thus fully mitigating the vulnerability.

3 DEMONSTRATION

In this work, we demonstrate three scenarios: pseudonym acquisition using Nexcom vehicular On-Board Units (OBUs) (S1), Pseudonym acquisition for a large-scale urban vehicular mobility dataset (S2), and the performance of the VPKIaaS system, notably its *high-availability*, *robustness*, *reliability*, and *dynamic-scalability* (S3). For the first two scenarios, our metric is the end-to-end pseudonym acquisition latency, measured at the OBU side. For the last scenario, we aim at demonstrating the performance of our VPKIaaS system by emulating a large volume of workload. For each scenario, the authors (presenters) will explain the underlying concepts behind different components of our scheme along with the achieved results.

Experimental setup: We created and pushed Docker images for LTCA, PCA, RA, and MySQL to the Google Container Registry (cloud.google.com/container-registry/). Isolated namespaces and deployment configuration files are defined before Google Kubernetes Engine v1.9.6 (cloud.google.com/kubernetes-engine) cluster runs the workload. We configured a cluster of three Virtual Machines (VMs), each with eight vCPUs and 10GB of memory. To emulate a large volume of workload, we created another Kubernetes cluster of four VMs (in a different data center), each with 10 vCPUs and 10GB of memory. Our full-blown implementation is in C++ and we use FastCGI [4] to interface Apache webserver. We use XML-RPC (xmlrpc-c.sourceforge.net) to execute a remote procedure call on the cloud. Our VPKIaaS interface is language-neutral and platform-neutral, as we use Protocol Buffers (developers.google.com/protocol-buffers) for serializing and deserializing structured data. For the cryptographic protocols and primitives (Elliptic Curve Digital Signature Algorithm (ECDSA) and TLS), we use OpenSSL with ECDSA-256 public/private key pairs according to the ETSI (TR-102-638) and IEEE 1609.2 standards; other algorithms and key sizes are compatible in our implementation.

3.1 S1: Pseudonym Acquisition by an OBU

Fig. 3.a shows two Nexcom vehicular OBUs (Dual-core 1.66 GHz, 1GB memory) from PRESERVE project (www.preserve-project.eu), which support IEEE 802.11p. In this scenario, we consider one of them to be an RSU, connected to the VPKI via Ethernet. The second

[5]This vulnerability is also relevant to the ticket acquisition process.

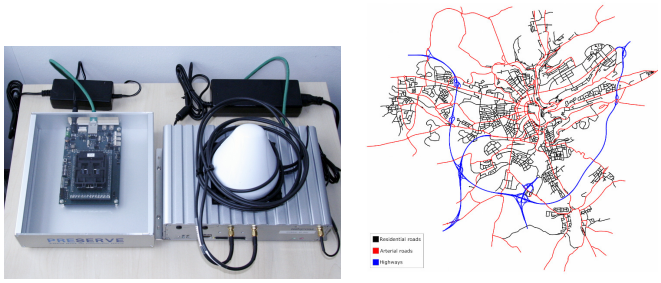

(a) (b)

Figure 2: (a) Nexcom boxes from the PRESERVE project, used in S1. (b) LuST Topology [3], used in S2.

OBU requests pseudonyms from the VPKI via the ''RSU'' and the communication is over IEEE 802.11p.

3.2 S2: Large-scale Pseudonym Acquisition

Fig. 3.b shows the LuST [3] scenario topology, a full-day realistic mobility pattern in the city of Luxembourg. We use OMNET++ (omnetpp.org) and the Veins framework to simulate this large-scale scenario using SUMO. In our simulation, we placed 100 RSUs in a region (50KM × 50KM). Each vehicle requests pseudonyms for its actual trip duration and V2I communication is IEEE 802.11p.

Fig. 3.a illustrates the CDF of the actual end-to-end latencies for obtaining pseudonyms with different pseudonyms lifetimes (τ_P) during the rush hours (7-9 am and 5-7 pm). For example, with τ_P = 1 minute, 95% of the vehicles are served within less than 286 ms. Fig. 3.b shows the average end-to-end latency with different pseudonyms lifetimes. Obviously, the shorter the τ_P, the higher the workload on the VPKI, thus the higher the end-to-end latency. The results confirm that our scheme is efficient and scalable: the pseudonym acquisition process incurs low latency and it efficiently issues pseudonyms for the requesters.

3.3 S3: VPKIaaS Performance

In this scenario, we aim at demonstrating the performance of our VPKI, notably its reliability and dynamic scalability. To emulate a large volume of workload, we generated synthetic workload with up to 14 containers with 1-4 vCPUs and 1-4 GB of memory. Each container generates 80,000 requests in the span of one hour leveraging 16 threads. One pseudonym request encapsulates 100 Certificate Signing Requests (CSRs) according to the standard (ETSI TR-102-638 and IEEE 1609.2). Fig. 4 shows how our VPKI system dynamically scales out/in according to the rate of pseudonym requests. The numbers next to the arrows show the number of PCA Pods at a specific system time.

We achieve a 5-fold improvement over prior work [2]: the processing delay to issue a pseudonym for [2] is 20 ms, while it is approx. 4 ms in our system. Moreover, unlike the VPKI system in [2], our implementation supports dynamic scalability, i.e., the VPKI scales in/out based on the arrival rate of pseudonym requests.

4 CONCLUSION

Paving the way for the deployment of a secure and privacy-preserving VC system relies on deploying a special-purpose VPKI. However,

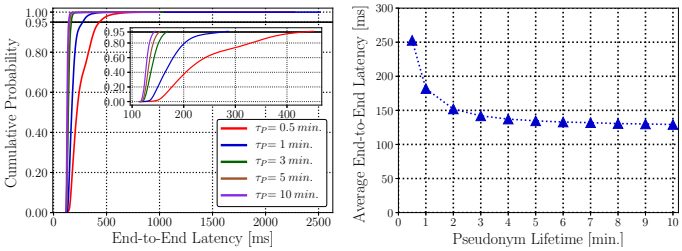

(a) (b)

Figure 3: (a) End-to-end latency for pseudonym acquisition. (b) Average end-to-end latency.

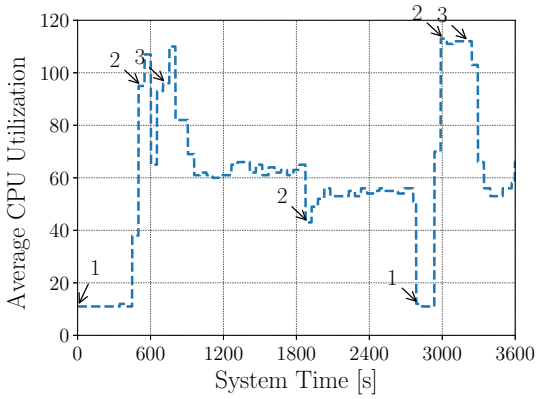

Figure 4: Dynamic scalability of the VPKIaaS.

its success requires extensive experimental evaluation, to ensure viability (in terms of performance and cost). We leverage the state-of-the-art VPKI and show its availability, resiliency, and scalability towards a cost-effective VPKI deployment.

ACKNOWLEDGEMENT

This work has been partially supported by the Swedish Foundation for Strategic Research (SSF).

REFERENCES

[1] 2014. V2V Communications: Readiness of V2V Technology for Application. (Aug. 2014). National Highway Traffic Safety Administration, DOT HS 812 014.
[2] P. Cincilla and et al. 2016. Vehicular PKI Scalability-Consistency Trade-Offs in Large Scale Distributed Scenarios. In *IEEE VNC*. Columbus, Ohio, USA.
[3] L. Codeca and et al. 2015. Luxembourg SUMO Traffic (LuST) Scenario: 24 Hours of Mobility for Vehicular Networking Research. In *IEEE VNC*. Kyoto, Japan.
[4] Paul Heinlein. 1998. FastCGI. *Linux journal* 1998, 55es (1998), 1.
[5] M. Khodaei and et al. 2014. Towards deploying a scalable & robust vehicular identity and credential management infrastructure. In *VNC*. Paderborn, Germany.
[6] M. Khodaei and et al. 2016. Evaluating On-demand Pseudonym Acquisition Policies in Vehicular Communication Systems. In *IoV/VoI*. Paderborn, Germany.
[7] M. Khodaei and et al. 2018. SECMACE: Scalable and Robust Identity and Credential Management Infrastructure in Vehicular Communication Systems. *IEEE Transactions on Intelligent Transportation Systems* 19, 5 (May 2018), 1430--1444.
[8] M. Khodaei, A. Messing, and P. Papadimitratos. 2017. RHyTHM: A Randomized Hybrid Scheme To Hide in the Mobile Crowd. In *IEEE VNC*. Torino, Italy.
[9] M. Khodaei and P. Papadimitratos. 2015. The Key to Intelligent Transportation: Identity and Credential Management in Vehicular Communication Systems. *IEEE VT Magazine* 10, 4 (Dec. 2015), 63--69.
[10] W. Whyte, A Weimerskirch, V. Kumar, and T. Hehn. 2013. A Security Credential Management System for V2V Communications. In *IEEE VNC*. Boston, MA.