

Building Automotive Security on Internet Standards: An Integration of DNSSEC, DANE, and DANCE to Authenticate and Authorize In-Car Services

Timo Salomon^{1,2}, Mehmet Mueller², Philipp Meyer², and Thomas C. Schmidt²

¹ Faculty of Computer Science, Dresden University of Technology, 01187 Dresden, Germany

² Department of Computer Science, Hamburg University of Applied Sciences, 20099 Hamburg, Germany

Corresponding Author: Timo Salomon (e-mail: timo.salomon@haw-hamburg.de).

Abstract—The automotive industry is undergoing a software-as-a-service transformation that enables software-defined functions and post-sale updates via cloud and vehicle-to-everything communication. Connectivity in cars introduces significant security challenges, as remote attacks on vehicles have become increasingly prevalent. Current automotive designs call for security solutions that address the entire lifetime of a vehicle. In this paper, we propose to authenticate and authorize in-vehicle services by integrating DNSSEC, DANE, and DANCE with automotive middleware. Our approach decouples the cryptographic authentication of the service from that of the service deployment with the help of DNSSEC and thereby largely simplifies key management. We propose to authenticate in-vehicle services by certificates that are solely generated by the service suppliers but published on deployment via DNSSEC TLSA records solely signed by the OEM. Building on well-established Internet standards ensures interoperability with various current and future protocols, scalable management of credentials for millions of connected vehicles at well-established security levels. We back our design proposal by a security analysis using the STRIDE threat model and by evaluations in a realistic in-vehicle setup that demonstrate its effectiveness.

I. INTRODUCTION

Vehicles comprise a distributed system of software-defined and hardware-controlled functions. Software innovations significantly enhance vehicle performance, safety, and comfort and contribute substantially to the value of modern cars [1], [2]. Online, connected vehicles open innovation options further and enable Advanced Driver Assistance Systems (ADASs), which continuously access infrastructure services and environmental data via Vehicle-to-X (V2X) communication. In general, connectivity introduced new software life cycles characterized by dynamic updates and functional upgrades that enhance agility and expand business models within the automotive industry.

Many suppliers contribute in-car software, which is organized in a loosely coupled Service-Oriented Architecture (SOA) [2]. Vehicular services can be dynamically orchestrated based on customer-selected configurations and available network and compute resources [3], [4]. However, this desired flexibility combined with global Internet exposure largely increases the attack surface of automotive systems. The risk of remote manipulation by third parties [5]–[7] became apparent from cyber-attacks in the field [8].

Vehicles are long-lasting products that require durable security solutions, which can be managed and maintained throughout their lifetime—correspondingly reflected in current

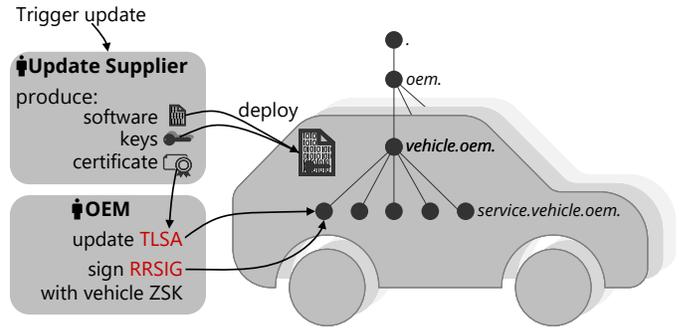


Fig. 1. Distributed and secure credential management for automotive services based on the DNSSEC ecosystem that decouples the software signing by the update supplier from the record signing by the OEM.

automotive security standards, such as ISO/SAE 21434 [9]. However, many automotive communication protocols, originally designed for closed systems, completely lack security features [6], [10], [11]. This includes the *Scalable service-Oriented Middleware over IP* (SOME/IP) [12], [13], the current choice of the AUTomotive Open System ARchitecture (AUTOSAR) for publish-subscribe communication, which has been shown vulnerable to various attacks [11], [14]. Previous attempts to secure the protocol [14]–[17] have focused on tailored, tightly integrated solutions for message encryption, authentication, or authorization. These solutions reinvent crypto protocols and build on certificates pre-shared across all communication partners.

The complexity and dynamics of in-vehicle services demand for (i) a scalable and robust key management to (ii) accommodate secure authentication and authorization of services and clients. In the automotive ecosystem, the Original Equipment Manufacturer (OEM) is responsible for the overall security lifecycle of a car—including its services. For this, it must coordinate with various stakeholders including tier-1 suppliers that provide critical components, software vendors that deliver specialized services, and aftermarket developers that offer add-on features. The heterogeneity of this ecosystem poses extraordinary challenges on key provisioning and robust credential management.

In this paper, we address these challenges by leveraging the established Domain Name System Security Extensions (DNSSEC) [18] ecosystem, as illustrated in Figure 1. Our approach decouples software and keys of the update supplier from the signing credentials of the OEM. While the soft-

ware supplier signs its product and supplies a corresponding certificate, the OEM signs the deployment of the certificate in a DNSSEC TLSA record using its own, independent credentials. A vehicle-delegated DNSSEC zone cached within the car enables offline verification of security credentials and accommodates mid-term scenarios without Internet access.

Based on this DNSSEC credential management, we derive a security layer for automotive service architectures. We integrate name-based service discovery and invocation in automotive software using the new SVCB service endpoint records in DNS [19]. DNSSEC provides verification of data integrity and authenticity for all of its records. The DNS then offers a scalable and distributed database, which stores the relevant data for service endpoint security within a resolver infrastructure that subsumes protocols and procedures for a remote update of records and for key rollover. A comprehensive ecosystem of tools and guidance supports the deployment and maintenance of DNSSEC in the Internet, which can be leveraged for the automotive domain. In detail, our main contributions read:

- We derive automotive security requirements and show how the DNSSEC ecosystem can comply with them (§III).
- We introduce DNSSEC-based authentication and authorization for publishers *and* subscribers in the automotive system (§IV).
- We propose deployment scenarios that support offline verification of security credentials and a secure key lifecycle management, addressing real-world requirements for automotive systems (§V).
- We integrate our approach with SOME/IP (§V-C) and evaluate its security (§VI) and performance (§VII) to demonstrate its effectiveness in practice.

II. BACKGROUND AND RELATED WORK

The automotive industry is embracing the Software as a Service (SaaS) transformation to enhance flexibility, upgradability, and customization [1], [2]. A SOA enables dynamic service orchestration based on customer configurations and available resources [3], [4]. Examples of such dynamic services include driving state-dependent features (*e.g.*, parking assistance), aftermarket software (*e.g.*, infotainment apps, ADAS), and hardware add-ons (*e.g.*, trailers with sensors, lights, and brakes). Consequently, this trend drives In-Vehicle Network (IVN) complexity with more optional features and variants [1]–[4].

Modern IVNs are transitioning from legacy Electrical/Electronic (E/E) architectures with dedicated Electronic Control Units (ECUs) organized in functional domains for powertrain, chassis, and infotainment, to a more integrated approach with centralized High-Performance Controllers (HPCs) and a unified communication backbone [2], [3]. Figure 2 illustrates a future zone topology [2], [4], [20] characterized by multi-service compute units in physical zones. Traditional bus systems like Controller Area Network (CAN) are gradually being replaced by a high-speed Ethernet backbone, working towards

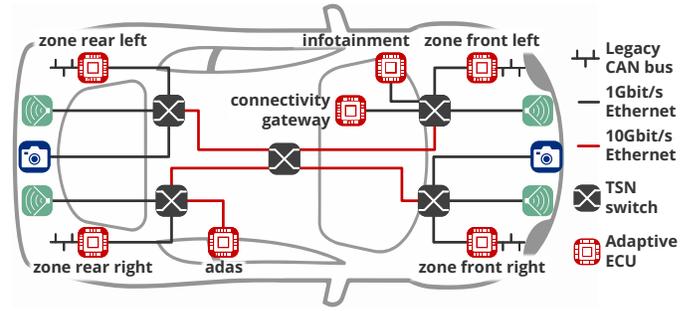


Fig. 2. Modern IVN in a zone topology with adaptive multi-service compute units and a unified Ethernet backbone.

a flat topology with a standardized IP stack [2], [4]. IEEE Time-Sensitive Networking (TSN) [21] enables concurrent best-effort and real-time communication over Ethernet [4], [20], [22].

Middleware layers facilitate service discovery, connection setup, and Quality-of-Service (QoS) provisioning. Prominent candidates include Data Distribution Service (DDS) [23], originating from robotics, and SOME/IP [12], [13], specifically tailored for automotive needs. Both support event-based publish-subscribe models, operating on top of the established transport protocols UDP or TCP [12], [23].

This paper focuses on SOME/IP due to its widespread adoption in the automotive domain and as the current choice of the AUTOSAR consortium [12], [24]. However, we emphasize that our concept is protocol-independent and can be readily adapted to DDS. To demonstrate the feasibility of our approach, we evaluate it in a realistic IVN scenario.

A. Security Goals, Threats, and Protection

Runtime-discovered services in software-defined vehicles include critical driving functions [2]–[4], making robust cybersecurity protection essential as vehicle connectivity increases [7], [25]. The attack surface of a modern car spans physical, limited range wireless, and globally accessible remote interfaces [5], [25], with threats capable of propagating across the IVN [6], [25]. Rising incidents [7] and demonstrated remote exploits [8] threaten passenger safety, economic damage, and violate privacy [6].

Attacks on IVNs include data alteration, monitoring, disabling, and forging [26]. Securing the IVN involves ensuring resource and service *availability*, safeguarding data *integrity*, and verifying *authenticity* of sources and sinks [4], [25], [26]. Dynamic service discovery requires authentication and authorization mechanisms that prevent spoofing service announcements, subscriptions, or cancellations [11], [14].

Layered security measures for automotive services include attack prevention, detection, and mitigation [27]. Prevention techniques involve firewalls [28], network flow separation [4], and state-aware defense [29]. Intrusion detection systems can identify ongoing attacks [30], *e.g.*, using machine learning to detect anomalies [31]. Message integrity can be ensured with authentication codes and encryption methods for IVNs [15], [32], [33]. Finally, secure boot [34] and Over-the-Air (OTA)

updates [35] protect against unauthorized software modifications and patch vulnerabilities.

We evaluate our authentication and authorization scheme against the STRIDE [36] threat model (Spoofing, Tampering, Repudiation, Information disclosure, Denial of service, Elevation of privilege) and specific automotive attack scenarios [11], [14]. We focus on remote attackers with control over ECUs and communication channels, capable of interception, modification, replay, blocking, or injection. Note that interactions with other security measures are beyond the scope of this work.

B. Service Authentication and Authorization

Service authentication establishes trust by verifying endpoints through certificates containing a public key and a signature derived from a trust anchor. Entities confirm authenticity via a challenge-response protocol, ensuring the service holds the corresponding private key.

Legacy IVNs relied on protocol-specific authentication and authorization with limited interoperability. CAN bus security approaches, for instance, face bandwidth and frame size constraints and are tailored to shared media [32]. Moreover, constrained resources of ECUs hinder robust security solutions [37], as established and verified ECUs are used across many vehicle generations. More powerful centralized platforms enable advanced security mechanisms [4], [37].

Dedicated protection for SOA protocols is essential for authenticating and authorizing access between services. SOME/IP, lacking embedded security features [6], [10], [11], could theoretically rely on TLS or IPSec, but these face restrictions in the automotive sector [15], [37]. IPSec is application-independent and cannot authenticate services. TLS and DTLS add complex authentication handshakes without authorization and lack support for multicast channels, which account for 75% of in-vehicle control flows [4].

Related work has exposed severe vulnerabilities in SOME/IP including timing-based Denial of Service (DoS) attacks, spoofing of discovery messages, and unauthorized access [11], [14]. Previous works [14]–[17] addressed these with message encryption, authentication, and access control. Iorio *et al.* [15] add a challenge-response handshake for authenticating, authorizing, and encrypting SOME/IP payloads. Zelle *et al.* [14] secure service discovery broadcasts employing ECU authentication and message validation based on an authorization server. Ma *et al.* [16] employ Message Authentication Codes (MACs) to prevent tampering and use a domain controller for mutual authentication of ECUs. Lee *et al.* [17] introduce a ticket system where an authentication server grants permissions for subscriptions. However, existing approaches depend on custom handshakes often using external authentication servers with little prior deployment experience, limiting interoperability with other protocols and requiring dedicated security hardening. None specifically authenticate and authorize the service discovery and subscription process, which is crucial for service availability and integrity [11], [14].

In this work, we design protocol-independent authentication and authorization for automotive services, enabling interoperability, longevity, and scalability based on the DNSSEC ecosystem. We demonstrate effectiveness through an integration with SOME/IP, given its widespread automotive use and lack of robust security mechanisms. Our approach incorporates a challenge-response scheme directly into SOME/IP discovery messages, securing service announcements and subscriptions.

C. Security Credential Management

While OTA software updates are becoming standard in the automotive domain, security credential updates have received little attention. Static and pre-deployed certificates pose significant risks over the long lifetime of a car, as NIST recommends using authentication keys for no more than one or two years [38]. Efficient and secure key rotation and distribution processes are missing to maintain authentication and authorization over the vehicle lifetime.

Previous SOME/IP security extensions [14]–[17] rely on pre-deployed certificates and static access control lists with no means of updating or managing security credentials. This approach creates problems in the automotive supply chain: (i) OEMs must pre-configure all potential peer credentials during manufacturing, and face inflexibility when adding or updating services due to cascading credential updates for related services. (ii) The system does not provide a process for suppliers to update software or credentials, which is crucial for maintaining security over the vehicle lifetime.

DDS leaves certificate management open, relying on static certificates or a Public Key Infrastructure (PKI) [39]. The X.509 PKI [40] attests certificate authenticity of Internet applications through trusted public Certification Authorities (CAs). However, the public CA model allows any trusted CA to issue certificates for any domain name, potentially leading to attestation collisions [41]. Further, the model does not account for the supplier-OEM relationships in automotive manufacturing, and certificate management becomes fragmented across different suppliers and components.

Recently, Püllen *et al.* [42] presented a credential deployment process utilizing a central authentication server to distribute access tokens across ECUs. The solution relies on a custom top-down procedure to deploy credentials to this central entity that involves system designers in credential management, which may impact longevity, interoperability, and scalability. Additionally, the procedure couples the in-vehicle trust chain provided by the authentication server with the trust in the software supply chain, which weakens the security model.

Mueller *et al.* [43] introduced a DNSSEC-based authentication approach for publisher validation, replacing the distributed discovery mechanism of SOME/IP with DNSSEC lookups without operational overhead. Their approach leaves access control management, client authentication, and in-vehicle deployment details unresolved. We fill these gaps to provide a comprehensive authentication and authorization framework for automotive services. As the discovery serves multiple

purposes, including service availability, service mobility, and service selection, we do not replace but rather enhance it with secure authentication and authorization.

In this work, we tackle authentication and authorization with integrated certificate management, designed to complement existing security measures. Our goal is to unify credential management across current and future automotive protocols using established DNSSEC mechanisms, providing a scalable, secure, and standardized solution with decades of deployment experience. We decouple the supplier security of software and credential production from the in-vehicle chain of trust provided by the OEM, offering a solution for the multi-supplier automotive ecosystem.

III. THE CASE FOR DNSSEC, DANE, AND DANCE

Vehicular communication involves domain-specific functions requiring appropriate protective measures for the security and safety of a connected car. Most notably, the heterogeneous multi-supplier SOA relies heavily on group communication, which is vulnerable to message fraud and unauthorized access during dynamic service discovery, instantiation, and migration. Implementing security in vehicles must carefully address domain-specific requirements.

A. Requirements for Authentication and Authorization of Vehicular Services

From the software-defined vehicle setting, security goals, threats, and known issues, we derive the following requirements for secure authentication and authorization in automotive systems.

Service continuity: Cars must always operate securely—unaffected by (lack of) remote connectivity, as well as software or security lifecycle management.

Long-term, agile credential management: Proven procedures for signing, certificate renewal, and key rollover must be in place in the multi-supplier ecosystem for managing large, heterogeneous fleets over the lifecycle of all cars.

Replicable, cacheable service credentials: Security-related service credentials must be self-consistent and locally verifiable—dependent of any (remote) attestation service. Security guarantees should be inherent to content objects and must not rely on third parties.

Scalability: Security solutions must scale with the growing number of services, clients, vehicles, and backend systems.

Interoperability and distribution: Security solutions must cater for distributed suppliers. Well established standards should ensure the robust interoperability of security protocols across suppliers.

Performance: While the service discovery itself is not time-critical, the authentication and authorization process should not introduce significant delays (< 200 ms [4]).

B. The DNSSEC Ecosystem

The DNS is the largest directory service on the Internet. Over the years, it has evolved into a rich ecosystem of information, application, and tooling. Various DNS records

can be transparently accessed, cached, and delegated, while DNSSEC [18] protects their integrity and authenticity via signature records (RRSig). As such, DNSSEC is interoperable and performant on Internet scale and caters for state-of-the-art cryptography. More than 15 years of operational experience have proven robust key management [44].

DNS-Based Authentication of Named Entities (DANE) and DANE Authentication for Network Clients Everywhere (DANCE) [45] build on DNSSEC to store and sign keys and certificates in TLSA records. They bind a certificate to a named service endpoint within the DNS hierarchy, establishing a chain of trust from the root zone to delegated zones. This not only avoids the well-known attestation collisions problem present in the WebPKI PKI [40], but in particular allows for independent signing of the service resource by its certificate owner and the service deployment by the owner of the DNS TLSA record. We will use this aspect to decouple the signing keys of the third-party service provider from the secure deployment of a service in a car by the OEM.

The security of DNSSEC, DANE, and DANCE has been thoroughly discussed in the IETF standardization [18], [41], [45], [46], [46], [47] and analyzed in extensive research [44], [48]–[51]. It has proven resilient over many years on the Internet [44].

In the automotive domain, OEMs can control fleet-wide deployment and backend systems, enabling full DNSSEC support while ensuring compatibility across clients and servers. Delegate zones and private namespaces can separate credentials for different vehicles. Individual suppliers can produce software updates and generate security credentials, while the OEM maintains the DNS hierarchy and ensures namespace integrity. Local in-car resolvers can cache records for secure offline operations, crucial for remote and disconnected scenarios, making this framework a robust, scalable solution for automotive security.

RFC 6394 [52] covers DANE use cases and requirements that we believe to meet with our protocol design for IVN service discovery, authentication, and authorization. Leveraging the DNSSEC ecosystem will enable secure key management while allowing services to authenticate clients and servers following a name hierarchy for authorization.

In the following, we present an adaptation of key automotive protocol primitives that integrate a challenge-response scheme into the SOME/IP service discovery, demonstrating feasibility in a realistic automotive scenario. We will show and evaluate how the DNS ecosystem can ensure a full in-car service authentication and authorization while meeting the specific requirements enumerated above.

IV. DNSSEC-BASED AUTHENTICATION AND AUTHORIZATION FOR PUBLISH-SUBSCRIBE SYSTEMS

Our automotive publish-subscribe security concept leverages the DNSSEC ecosystem for authentication and authorization. Figure 3 provides an overview of the involved modules and communication. The DNSSEC infrastructure is a distributed system that supports zone delegation and replication

across regions. Records for backend and vehicle services can be controlled and provisioned by the OEM.

A local DNSSEC resolver in the car caches the necessary records which can be pre-loaded and queried from the backend when they are missing on a client request. Thus, during normal operation, the car can operate offline. The local resolver verifies authenticity and integrity for all records through the DNSSEC chain of trust, eliding the need for validation on ECUs with limited capabilities. Still, applications can independently verify the chain. During service discovery, publishers and subscribers query the local resolver for required records.

A. Acquiring Valid Endpoint Information

Optional find messages allow subscribers to trigger service discovery via multicast. Publishers respond to find requests or periodically offer their services to signal availability. This prompts subscribers to renew their subscriptions.

During the discovery or after the offer is received, subscribers query the local DNSSEC resolver for publisher endpoint information, which we store in Service Binding (SVCB) [19] records due to their flexibility and extensibility. SVCB supports various Application Layer Protocol Numbers (ALPNs) (e.g., indicating HTTPS over TCP), and includes the corresponding address and port.

A service query may return multiple SVCB records, such as when a service supports multiple protocols or has multiple instances. The subscriber compares the DNSSEC-secured SVCB information with the offered publisher to ensure the endpoint is valid before issuing a subscription.

B. Authenticating Publishers and Subscribers

DANE [41] and DANCE [45] bind certificates to endpoints and clients using TLSA records. We integrate a challenge-response mechanism into the existing publish-subscribe handshake, inspired by current transport solutions, such as QUIC, following the principles of TLS 1.3.

Publishers can embed an authentication challenge in their offer. Subscribers respond with a subscribe message including a signed authentication response using their private key. They can request authentication from publishers by including a challenge in the subscribe message, which the publisher addresses in the subscription acknowledgment.

To verify the presented signatures, publishers and subscribers query the local DNS for the TLSA records of the other party, deriving the name from information provided in discovery messages. We describe a SOME/IP implementation in section V-C. TLSA records can contain either the full certificate or a hash to be verified against a presented signature. The **DANE-EE** (End Entities) option does not require additional CAs, relying solely on the DNSSEC trust chain. Again, a query may return multiple records, allowing for certificate migration.

To reduce response time, the subscriber can query SVCB and TLSA records in parallel, ignoring the latter if the SVCB record is insecure [47]. If a TLSA response is insecure (or

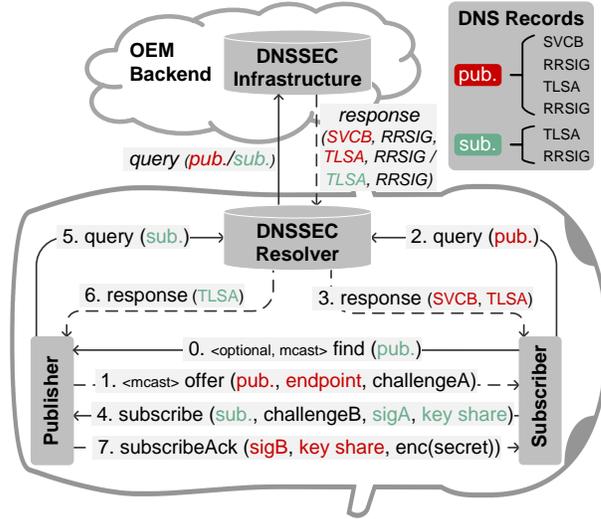


Fig. 3. DNSSEC-based authentication and authorization scheme for automotive publish-subscribe systems. The publisher and subscriber information are color coded.

absent), both parties may establish an insecure connection subject to implementation policies [41].

C. Name-based Authorization for Services

Traditional authorization mechanisms often rely on access control lists or complex infrastructure, which provoke cascading updates and error-prone maintenance. Instead, we reuse existing TLSA record names for service authorization. Authorization by name supports fine-grained policies while remaining simple to manage and deploy, without need for additional infrastructure or cryptographic overhead.

Publisher record names include the service name/id, domain, and vehicle information, resulting in a name like *service.domain.vehicle.oem.*, with prepended protocol information. Subscribers verify the presence of a DNSSEC-secured TLSA record matching this exact name to implicitly authorize the publisher, relying on the DNSSEC trust chain.

Subscribers can be authorized along the same mechanism. Further, we introduce a hierarchical naming scheme to allow for scoped access policies, e.g., per domain or vehicle-wide. Subscriber names can reflect these authorization scopes for certificates with optional fields: *_proto-client.sub[.pub][.domain][.vehicle].oem.* Publishers define accepted scopes depending on application requirements, e.g., requiring service specific names for tighter authorization of critical services.

D. Session Establishment and Key Exchange

We establish a secure session between the publisher and subscriber using Diffie-Hellman (DH) key exchange [53]. The subscriber includes its public share in the subscription, and the publisher in the subscription acknowledgement. Both parties can then compute the symmetric session key.

For multicast endpoints, the session key must be agreed upon between multiple parties, though a dedicated group key agreement protocol is beyond the scope of this work. In

our scheme, the publisher is the key sponsor that creates a symmetric group key and provides it to subscribers through the established secure channel. As the publisher is the only party that knows all subscribers, it can securely distribute the group key to each subscriber individually without overhead.

For both unicast and multicast, the symmetric key can be used to encrypt and decrypt published messages, ensuring confidentiality and integrity. The publisher updates the session key when required using the secure channel. Furthermore, secure and fast session resumption mechanisms, *e.g.*, following DTLS, may reduce overhead and improve performance, but are not in the scope of this work.

V. IN-VEHICLE DEPLOYMENT AND INTEGRATION

Our solution addresses the specific requirements of in-vehicle operations and the automotive industry, particularly for OEMs and suppliers. Safety and security are paramount, with OEMs bearing responsibility for the entire vehicle lifecycle, including third-party components. Repair and maintenance in independent workshops are integral to this framework. We decouple the cryptographic authentication of service suppliers from that of OEMs using DNSSEC. Our approach involves authenticating in-vehicle services with certificates generated by service suppliers but published via DNSSEC TLSA records signed solely by the OEM.

A. Certificate Management

Effective certificate management is a cornerstone of the proposed solution, with the OEM overseeing the entire fleet and many suppliers providing software updates and services. Leveraging DNSSEC infrastructure for storing and distributing service security credentials offers significant benefits, streamlining processes and enhancing security. The DNSSEC ecosystem includes robust tools, best practices, and in-depth security analysis [44], [49].

Service credential updates: Regular rotation of credentials is critical for maintaining security. NIST recommends that authentication keys be used for no more than one or two years [38]. Figure 1 illustrates our credential management approach for automotive services in a multi-supplier environment. To ensure a secure and cohesive system, the DNSSEC hierarchy is maintained by the OEM, which retains control over in-vehicle service deployment. A delegate zone for each vehicle stores service and client credentials (*e.g.*, `adas.vehicle42.vw.`) in TLSA records, as per DANE and DANCE. Each delegated zone uses a unique DNSSEC Zone Signing Key (ZSK) per vehicle signed by the OEM's Key Signing Key (KSK), building the chain of trust.

When a software or credential update is required, the update supplier produces and signs a software binary, a public-private key pair, and a corresponding certificate. The OEM updates the TLSA record with the new public certificate and adds an RRSIG record signed with the vehicle ZSK. TLSA and RRSIG records are deployed to the DNS tree, where they can be validated and cached by the local in-car resolver. Multiple TLSA records can coexist for the same service, allowing

a smooth transition to new credentials before expiration. A secure OTA update process deploys the software into the car, storing the private key locally, *e.g.*, in a secure element of a Hardware Security Module (HSM).

Our approach separates the software signing by the update supplier from the record signing by the OEM, ensuring the OEM maintains full control over the DNS. Furthermore, software and credentials are handled independently without involvement of a third party, reducing risks for private key exposure. OEMs can provide tooling to enable suppliers and independent garages to securely update software and credentials, deploying DNSSEC records on demand, honoring the right to repair.

Certificate uniqueness: Certificates can be issued per car, per service, or per car and service, depending on security requirements and supply-chain relationships. However, reusing certificates across vehicles poses significant risks. Attackers could extract private keys from one vehicle and exploit them across multiple cars.

Using unique key pairs per service in each vehicle presents a conceptually straightforward and secure solution but poses scalability challenges. For example, assuming 10 million cars manufactured per year over one decade, each with 500 services, the system would need to handle at least 50 billion entries. This problem is not unique to our DNSSEC approach but also affects other solutions where an OEM maintains all vehicle credentials.

Here, the distributed and delegated DNS provides a distinct advantage. We deploy a dedicated DNSSEC zone for each car with only the required few thousand records making it easily manageable, compared to Internet DNSSEC deployment. A well-structured DNSSEC delegation hierarchy ensures scalability while maintaining security with unique certificates for each vehicle.

Differentiation from pre-deployed certificates: Our certificate management process simplifies the deployment scenario of public certificates and allows for a more dynamic approach to service deployment. Pre-deployed certificates are not well-suited for agile service deployments due to their static nature requiring all parties to know all public certificates in advance – necessitating cascading credential updates whenever a service is modified. When a service is updated or added, it must be deployed not only with its own keys but also with public certificates for all other services that may interact with it. The same goes for interacting services, which must be updated with the new public certificates.

In contrast, our approach allows for dynamically requesting certificates via DNSSEC when a new communication partner is discovered or when a new service is deployed, without requiring prior knowledge of the full in-vehicle service deployment. Afterwards, credentials can be used in the same way as with pre-deployed certificates over their Time to Live (TTL). Both pre-deployed and dynamic certificates can exist in parallel.

B. Online vs. Offline Operation

Our approach is primarily designed for **regular online operation**, facilitating certificate updates and changes in service deployments. In the automotive domain, connectivity has become a de facto standard, with features like eCall requiring online capabilities. Many other features, including interactions with backend services, charging stations, and other infrastructure, also necessitate online capabilities. Our approach enables comprehensive integrations by providing authentication and authorization between in-car, OEM backend, and third-party services and clients.

Fallback offline operation: Vehicles may operate in offline mode for extended periods, such as during long-distance travel or in remote areas with limited connectivity. It is important to note that the attack surface is reduced in offline scenarios, particularly concerning remote exploits. Our system still supports offline verification using essential DNSSEC records cached and previously validated by the local resolver. However, software updates or the installation of new components are not possible unless the necessary credentials have been pre-loaded.

To ensure system reliability and security during typical offline periods (*i.e.*, a few weeks or months), the lifetime of records and certificates must be sufficient. Degradation modes may be implemented to allow prolonged offline operability with the extended use of expired certificates. However, users should be notified of the reduced security level. Furthermore, session resumptions using the last known session key could be implemented, restricting rekeying to online operations. It is essential to carefully design such fallback mechanisms to avoid creating new attack vectors targeting system degradation.

End of life considerations: OEMs or third-party suppliers may discontinue support for older vehicle and service generations, leading to issues with certificate expiration. Countermeasures include the use of self-signed certificates, creating an in-car trust chain or deployment of fallback certificates with unlimited lifetime credentials. Vehicles with insecure certificate should not be permitted to operate in a connected environment.

While the degradation of functionality and potential decrease in operational lifetime of cars going out of manufacturer support is a concern, it is not new. Future (possibly autonomous) connected cars will require software and security maintenance to continue safe and secure operation. DNSSEC provides a long-living and well-established solution for long-term security maintenance.

C. Example Integration with SOME/IP

Our security framework is designed to be compatible with a wide range of existing and emerging communication protocols by leveraging open standards. To demonstrate this, we focus on AUTOSAR SOME/IP [12], [13], integrating DNS-based authentication and authorization while maintaining compatibility with existing SOME/IP implementations.

Messages and configuration options: SOME/IP messages follow a publish-subscribe model, beginning with a find-offer mechanism to discover services and then establishing subscriptions. Appendix A provides an overview of the SOME/IP message structure (*cf.*, Figure A.1). Service discovery operates as a SOME/IP service supporting unicast and multicast, using known UDP-IP endpoint, *Service ID*, *Method ID*, and *Client ID*.

Service discovery messages contain a list of entries and options (*cf.*, Figure A.2). Entries, such as **Find**, **Offer**, **Subscribe**, or **SubscribeAck** reference options that provide, for example, endpoint information. **Offer** and **Subscribe** entries with a TTL of zero serve as **StopOffer** and **StopSubscribe** messages.

Configuration Options (*cf.*, Figure A.3) are standard-conform, structured, application-specific strings. We implement four custom options: (i) a *challenge* includes a random 32-bit nonce, (ii) the *response* contains the signed nonce, (iii) *key exchange* contains the public share of the DH exchange and cryptographic parameters, and (iv) a *session key* distributes the encrypted symmetric key for secure multicast communication.

Protocol state machine: SOME/IP must incorporate DNS interactions, authentication, and authorization. Appendix B shows the adapted protocol state machines with simplifications such as omitting the initial wait phases, retransmissions, and timeout handling. We concentrate on scenarios where applications actively request services.

Servers periodically send **Offer** messages upon startup, a behavior unchanged by the proposed security mechanisms. The client-side discovery process (*cf.*, Figure B.1) starts with an initial **Find** and a DNSSEC query. Clients wait for both the DNS lookup and find to complete. Upon receiving an **Offer**, the client validates it against the SVCB record. Invalid **Offers** are ignored, and the client awaits the next.

Client-side subscription handling (*cf.*, Figure B.2) starts with a **Subscribe** message and then waits for a **SubscribeAck**. The acknowledgment is authenticated against the previously retrieved TLSA record. A valid acknowledgment establishes the subscription; otherwise, the client awaits another. Clients respond to cyclic **Offers** by **Subscribeing** again.

Servers (*cf.*, Figure B.3) cache active subscriptions. Incoming **Subscribes** are first authorized by client name. The server awaits the client TLSA record to authenticate the subscription and acknowledges it with a **SubscribeAck**.

VI. SECURITY ANALYSIS

We evaluate security properties of our DNSSEC-based authentication and authorization scheme for automotive publish-subscribe systems against common threats and known attacks. Our approach relies on established cryptographic algorithms and protocols (*e.g.*, DH key exchange); no new crypto primitives are introduced. Hence, we do not consider attacks on the primitives and algorithms themselves. Instead, we focus on the security properties of the protocol design and its

integration with the automotive middleware SOME/IP. We first identify inherent properties of the DNSSEC ecosystem that are critical for automotive security. Subsequently, we analyze protocol state transitions with STRIDE threats and identify open limitations.

A. Adversary and Threat Model

In this analysis for automotive cybersecurity threats [6], [26], we assume a group of attackers with critical knowledge and all necessary tools, which represents the worst-case scenario. Given the large attack surface of vehicles, as explored in various studies [5]–[8], [10], [26], any device in the vehicle could be compromised. We also assume that attackers could bypass any protection mechanisms of the connectivity gateway of the car, gaining remote access to all in-vehicle communication channels.

We apply the Dolev-Yao adversary model [54] for public key protocols, previously used for automotive networks [32]. Here, the adversary has full control over all communication channels and can intercept, modify, replay, block or inject messages. However, we assume that an attacker cannot access (read/write) internally protected information of devices (e.g., securely stored keys), or change the physical network connections, including those of the attacker [11]. While additional security measures, such as secure flow control [4], intrusion detection systems [30], and firewalls [28], might be in place, they are not considered in our evaluation.

We do not consider asymmetries in computational power as would be relevant for crypto attacks, since we assume crypto-primitives, certificates, and signatures are secure, e.g., with the help of protected key-storage.

B. Inherent Properties of the DNSSEC Ecosystem

The following properties are inherent to DNSSEC, DANE, and DANCE and remain altered by our integration.

TLSA certificate security: The security of certificates relies on the DANE TLSA record, which binds certificates to DNSSEC-secured names, ensuring integrity and authenticity [41], [46]. By applying the certificate usage **DANE-EE**, we bypass third-party CA checks, as the TLSA record directly references the certificate or its public key. This lean verification avoids additional checks, reducing potential inconsistencies and ensuring conflict-free validation [47], provided the DNSSEC infrastructure remains secure.

DNSSEC validators and caching: Validation of DNSSEC records is critical to prevent spoofing and service disruptions, particularly during events like certificate rollovers [41], [49]. For highest endpoint security, DNSSEC validation should occur on the host system, which is possible in our setting. In cases of threat alerts or as routinely security monitoring, in-car ECUs can check and validate TLSA records to prevent failures or misbehavior [48]. In normal DNSSEC operation, the (on-board) recursive resolver performs the record validation, which is the preferred setting in the automotive environment with constrained in-car ECUs.

The in-car DNSSEC resolver not only delivers DNSSEC-validated data but also pre-loads and caches the respective zone records of the car to support offline operation. A correct resolver implementation will honor the DNS TTLs and certificate lifetimes, thus avoiding issues such as stale data exploitation or denial-of-service attacks.

Services rely on the in-vehicle DNSSEC resolver to ensure the integrity of data that is resolved non-locally, e.g., after an update or certificate rollover. DNS communication is transmitted over a secure channel to prevent man-in-the-middle attacks. However, if the resolver itself is compromised, falsified data could still be validated, making DNSSEC resolver protection critical for maintaining data integrity. Protection could include system security such as firewalls, secure boot, etc., and intrusion detection systems actively monitor resolver activity, probe DNSSEC validation, or create event logs to detect potential attacks.

DNSSEC key management: DNSSEC uses cryptographic keys to sign records, with the security of DNSKEYs depending on technical protection, process controls, and personnel expertise [49]. The separation of KSKs and ZSKs enhances security, as KSKs are suitable and advised for offline storage. A robust DNSSEC infrastructure requires automated and secure management processes [50], [55], which are widely available. We assume that automotive OEMs have the knowledge and capabilities to professionally control DNSSEC infrastructure.

Attacks on the DNSSEC infrastructure: The impact of infrastructure attacks is significantly reduced due to the constraints on key usage. DNSKEYs are restricted to the zones they manage, so a compromise would not affect other zones. In our setting, each vehicle operates in its individual zone, thus preventing the escalation of a compromise across cars. In contrast, other certification infrastructures, such as the public CAs, can issue certificates for any service and thus enable escalation of compromises. While neither DNSSEC nor public CA models can fully prevent invisible attacks from compromised keys, the automated DNSKEY rollover and revocation mechanisms provide additional security layers to reduce risks and limit attack opportunities.

C. STRIDE Threat Analysis

We use the STRIDE threat model [36] for our threat analysis, as it covers relevant attack vectors such as spoofing, tampering, repudiation, information disclosure, denial of service, and elevation of privilege. Based on the example of securing the SOME/IP middleware, we describe how attacks apply to the automotive SOA and how our authentication and authorization mechanisms protect the service discovery against such attacks, summarized in Table I.

1) *Spoofing publisher or subscriber identities:* Spoofing publisher or subscriber identities is a major issue for automotive service discovery [11], [14]. We protect the previously unverified SOME/IP **Offer**, **Subscribe**, and **SubscribeAck** messages with a challenge-response handshake. Signatures are validated against the certificate stored in TLSA records, ensuring the authenticity of the publisher and

TABLE I
STRIDE SECURITY ANALYSIS OF THE PROPOSED DNSSEC-BASED AUTHENTICATION AND AUTHORIZATION SCHEME.

Sec.	STRIDE threat	Automotive attack vector	Protective measure
VI-C1	Spoofing identities	Impersonation of publisher or subscriber identities, <i>e.g.</i> , conflicting offers or cancellation of subscriptions	Challenge response handshake; signatures validated against TLSA records
VI-C2	Tampering with messages	Manipulation of discovery or subscription messages, <i>e.g.</i> , man-in-the-middle, replay attacks	Signing critical discovery messages; MACs as proposed in [15], [16]; encrypted data from derived session key
VI-C3	Repudiation of actions	Denying involvement in a transaction or communication, <i>e.g.</i> , a subscription request or data exchange	Name-based authorization proves the origin of messages; logging actions prevents denial
VI-C4	Information disclosure	Unauthorized access to sensitive information, <i>e.g.</i> , service relationships and data exchange	Authorization along name hierarchy; DNSSEC protects against false certificate associations; secure channel
VI-C5	Denial of service	Disrupt service availability, <i>e.g.</i> , exhausting resources, or exploiting state transitions and timings in protocols	Spoofing mitigation; SOME/IP stack configurations; rate limiting; DNSSEC resolver protections
VI-C6	Elevation of privilege	Gaining higher-level access or permissions than authorized, <i>e.g.</i> , access to safety critical domains	Scoped name hierarchy; TLSA association to specific names, protocols, ports; OEM controlled trust chain

subscriber. DNSSEC protects against false certificate associations (*cf.*, section VI-B), ensuring that only valid certificates are authenticated. We do not secure **Find** messages, and every endpoint can discover all services as per the SOME/IP and DNS designs.

To prevent false service offers, subscribers verify the service endpoint using the DNS SVCB records and send subscribe messages only to matching endpoints. The subscriber signs the nonce included in the offer challenge and sends it with the subscribe message. After retrieving the TLSA record, the publisher can securely authenticate and authorize the subscriber. For subscription acknowledgments, publishers sign the nonce in the subscribe message and return it, enabling subscribers to authenticate responses against the publisher TLSA record. **Stop Offer** and **Stop Subscribe** messages are authenticated along the same mechanism. Thereby, invalid publishers or subscribers can be rejected, which prevents spoofing attacks.

2) *Tampering with messages*: Tampering with discovery or subscription messages violates authenticity, which is mitigated by our protection. While not the focus of this work, protection against tampering, *e.g.*, man-in-the-middle attacks, for SOME/IP has been studied before [15], [16], introducing MACs to ensure the integrity of messages. Additionally, security credentials are validated against secured DNSSEC records. No confidential information is shared before the other party is authenticated and authorized. With our integrated DH key exchange mechanism [53], subsequent communication can be encrypted using the derived session key, ensuring data confidentiality and integrity.

3) *Repudiation of actions*: Before services can be accessed, the publisher and subscriber identities are authenticated and authorized, directly linking interactions to applications. Logging all actions helps to associate actions or changes, preventing repudiation attacks.

4) *Information disclosure*: The service itself determines the name hierarchy authorized for access, and only names with valid TLSA records are authorized. Although information about service relationships within the car can be inferred, this data is non-sensitive and can also be obtained via the SOME/IP service discovery.

No data is transmitted through the communication channel until both parties are authenticated and authorized. A secure channel with perfect forward secrecy is established between the publisher and subscriber during the subscription. For multicast scenarios, publishers share the session key over this secure channel with valid subscribers. Event notifications can be encrypted using the symmetric session key. Publishers should update multicast session keys regularly to prevent access of expired subscribers from previous sessions. Though, considering the automotive use case, subscribers are likely to re-subscribe with each cruise.

5) *Denial of service*: Service availability is also a safety requirement for vehicle operation, especially while driving. The service discovery phase, however, is not time-critical and service availability is not guaranteed within the design of current cars. Current automotive systems have little redundancy but shall only operate after critical services are started and communication is established.

For the SOME/IP service discovery, DoS attacks have been identified [11], [14], mainly abusing state transitions or timings in the protocol by spoofing messages. Spoofing, however, is mitigated by our authentication scheme (*cf.*, section VI-C1). Prevention by suitable configurations of the SOME/IP stack is possible [11].

The additional resolution of DNSSEC records may delay the service discovery process. However, current DNSSEC resolvers are optimized and can handle many requests in parallel. Still, mass requests may lead to a DoS attack. Source address spoofing in DNS queries can cause a reflection attack, which rate limiting or firewalling in the network, as suggested in [4], can mitigate, making each client ECU detectable. Time-Sensitive Networking protocols [21], as foreseen in future cars, will additionally protect time-critical streams from delays by non-time-critical DNS requests.

Attacks targeting the DNSSEC resolver are possible, *e.g.*, requesting *existing* records with long DNSSEC chains leading to a larger number of crypto operations. This can be prevented in cars by restricting resolver access to the pre-fetched zone relevant for car operations. Furthermore, the in-car resolver is only accessible from inside the car; thus,

external services cannot exploit it for attacks. Finally, most resolvers have defense strategies that abort high-load or long-verification lookups.

6) *Elevation of privilege*: The publisher uses the DNS name of the subscriber to decide whether to authorize the subscription during authentication. TLSA certificate associations are constrained to specific names, protocols, and ports, which further limits the attack surface. DNSSEC-protected records prevent elevation of privileges, as the correct private key is required to impersonate a service and use its identity. By strict separation of concerns, the OEM can manage the DNSSEC infrastructure independently from update suppliers, maintaining the DNSSEC chain of trust.

D. Discussions and Limitations

Our approach successfully addresses authentication and authorization using the established Internet standards DNSSEC, DANE, and DANCE for robust credential lifecycle management. With this, we protect SOME/IP against common STRIDE threats. However, some limitations remain, which we discuss in the following.

Mixed security levels: Our analysis focuses on providing the highest security level for automotive service discovery, which may be unnecessary for some services. Allowing for mixed security status improves backward compatibility with existing systems. In a vehicle, different levels of security are likely [15]; *e.g.*, not all services require data encryption. Subscribers and publishers might request authentication, authorization, and encryption depending on service requirements [15], the adaptation of which is straightforward with our approach. Nevertheless, for services that accept mixed security levels, downgrade attacks should be considered and mitigated [15], [41], which is not the focus of this work.

Supply chain attacks: While credential management is secured by the DNSSEC chain of trust, our analysis does not cover specific attacks on software production, credential generation, or secure OTA distribution. We decouple the software production from credential management. Update suppliers sign new credentials and binaries, enabling the OEM to verify integrity and authenticity before deployment. The OEM itself signs the DNSSEC records; thus, vehicles fully rely on trust in the OEM via the DNSSEC chain, removing dependencies on third-party trust anchors. Secure build systems, credential production, OTA updates [35], and bootstrapping [34] can protect the software supply chain in the automotive domain. As such, they require an independent detailed security analysis, which is beyond the scope of this work.

Private key compromise: Private key compromise, though not considered a significant threat when HSMs are in use, could occur if an attacker gains physical access to an ECU along with the necessary tools, knowledge, and time. If a publisher's private key is compromised, the attacker could offer services authenticated with the stolen certificate. Our approach restricts fraudulent offers to services with valid endpoint information in the SVCB records and a matching TLSA record. Similarly, a compromised private key of a

subscriber enables the attacker to access services associated with TLSA records tied to the compromised client key. We emphasize that each vehicle uses unique key pairs for its service zone, meaning that the compromise of a private key affects only the specific vehicle in which it is deployed.

Discovery delays: While spoofing and the majority of DoS attacks are mitigated, delays in discovery remain possible. Conflicting messages may arrive before DNSSEC records are retrieved, causing subscribers to receive multiple **Offer** or **Stop Offer** messages for the same publisher service. Once DNSSEC records are resolved, only the latest offer message is validated, leaving a window for timing attacks as described in [11]. In the worst case, the subscriber must wait for the next valid cyclic offer message, delaying discovery by up to one cycle. Similarly, publishers may receive multiple subscribe messages for the same client before the client certificate is obtained. Although caching these messages could mitigate delays, it may introduce new attack vectors. Finally, all conflicting requests are resolved once DNSSEC records are retrieved, resulting in a minor vulnerability window for discovery delays. Initializing secure connections before the vehicle starts mitigates the impact of these delays.

Complexity attacks: Complexity attacks exploit the system state or inflate cryptographic operations of a protocol to slow down or overload the system. In its regular mode, all asymmetric crypto operations of DNSSEC are asynchronously executed by the internal recursive resolver, leaving nonce generation, signature generation, and verification only for the initial phase of service session establishment. Corresponding complexity attacks in this phase will delay or suspend the startup of the car. Vehicles could mitigate remote attacks by closing the gateway or tightening the firewall. In the case of an attacking internal component, the vehicle would likely require a software reset—either via OTA or offline service maintenance.

State attacks target memory held for service offers, unauthenticated subscribe requests, and active subscriptions. While the state for offered services and active subscriptions remains similar to SOME/IP, our solution increases the memory footprint due to additional nonces, certificates, and signatures. For unauthenticated subscription requests, DNSSEC-based validation requires storing incoming requests until their validation, which increases the state compared to static pre-deployed certificates. Such attacks require a high volume of malicious requests and can be mitigated by monitoring request patterns.

VII. EVALUATION IN A REALISTIC IN-CAR SCENARIO

We compare the performance of our DNSSEC-based authentication and authorization scheme to the vanilla SOME/IP service discovery and a pre-deployed certificate approach. Therefore, we recreate a realistic in-vehicle scenario from Meyer *et al.* [31] in a Mininet environment, focusing on service discovery performance, cryptographic operations, and DNSSEC resolution time.

Our implementation builds upon open-source projects, including the SOME/IP reference implementation *vsomeip* [56].

TABLE II
 PUBLISHER (SERVICE 42 , INSTANCE 1 , MAJOR 2 , MINOR 3), AND SUBSCRIBER (CLIENT 17) RECORDS IN *vehicle1.oem*. DELEGATE ZONE.

Name	Type	Value
._someip.3.2.1.42.service.vehicle1.oem.	IN SVCB	1 . ipv4hint=10.0.0.2 port=5000 instance=1 major=3 minor=2 ip_proto=17
._5000._someip.3.2.1.42.service.vehicle1.oem.	IN TLSA	3 0 0 (308205...)
._someip-client.2.1.42.17.client.vehicle1.oem.	IN TLSA	3 0 0 (308204c...)

Extensions with our authentication scheme rely on libraries for DNS lookups (*c-ares* [57]) and cryptography (*Crypto++* [58]). All services run on a VM (*Ubuntu 22.04.5 LTS*, 32 GB RAM, 16 Intel®13900K p-cores), with each service utilizing its own SOME/IP stack to prevent local connections without service discovery. We employ *Mininet* [59] as a virtual network and *NSD* [60] as our DNSSEC resolver. Following the anonymous review, we will make our implementation available as open-source, including configurations and the evaluation setup.

The SOME/IP service discovery is configured with the default settings from the *vsomeip* subscription example: an initial delay of 10 ms to 100 ms to scatter requests, a maximum of 3 find/offer repetitions with increasing delays starting from 200 ms and a cyclic offer delay of 2 s.

We measure service discovery latencies, DNSSEC lookup times, and the computational cost of relevant cryptographic operations. Our evaluation compares three variations to assess the impact of DNSSEC-based authentication and authorization: (i) *Vanilla vsomeip*: The reference implementation with added statistics to measure performance. (ii) *Pre-deployed certificates*: Full publisher and subscriber authentication using static certificates. (iii) *DNSSEC, DANE, and DANCE*: DNSSEC-based authentication and authorization with local DNS resolution.

A. Configuring a Realistic IVN

We base our setup on a realistic open-source IVN model [31], adapted from a real vehicle transformed to a zone topology. We simplify the Mininet configuration by replacing the redundant ring backbone with a star topology (see Figure 2). The network comprises five switches, four lidar systems, two cameras, four zone control ECUs, and three HPCs for ADAS, infotainment, and telematics services. We convert all traffic sources and sinks to SOME/IP publishers and subscribers while preserving original communication relations and application hosts. This configuration deploys 212 publishers and 448 subscribers, with publishers having an average of 2.1 subscribers (minimum 1, maximum 4), utilizing only multicast endpoints. Nodes have an average of 16 local publishers (min 0, max 79) with 35 remote subscribers (min 0, max 174), and 35 local subscribers (min 0, max 131) for 35 remote publishers (min 0, max 79).

Certificate and key generation: For secure communication, all publishers and subscribers use unique certificates and keys generated with OpenSSL [61] (SHA-256, 2048-bit RSA keys, approx. equivalence with elliptic curve prime256v1). The *vanilla vsomeip* scenario does not use any certificates. *Pre-deployed certificates* are linked directly in the application configuration, so each publisher knows the certificates of its

subscribers. For the *DNSSEC, DANE, and DANCE* scenario, publishers and subscribers only store their private keys, with public certificates resolved via DNSSEC.

DNSSEC records: We create a DNSSEC zone for our car *vehicle1.oem*. divided into two zones: *.client.* and *.service.* The complete IVN requires 872 records, each signed with a corresponding RRSIG record. DNSSEC clients rely on the resolver for validation, however, in our evaluation we do not validate the chain of trust to a root zone.

The *.service.* zone contains SVCB and TLSA records for each of the 212 publishers. The SVCB data describes the publisher information matching SOME/IP offer messages, using custom key-value pairs supported by RFC 9460 [19]. Different SVCB records may coexist under more general names, e.g., *_someip.42.service.*, linking to the same publisher data but allowing for requesting all records for any version and instance, as proposed by Mueller *et al.* [43]. The TLSA record is always specific to the full service information, and we use the **DANE-EE** certificate association type, eliminating the need for additional CAs.

The *.client.* zone has TLSA records for all 448 subscribers, which contain the full certificate of the subscriber. These client-service specific TLSA records must be present to authenticate and authorize access to a service. We do not utilize wildcards, as we cannot predict them realistically from the anonymized services in the original network [31].

B. Performance Evaluation

We evaluate performance in a demanding scenario with 660 applications initialized on a single host system. This setup introduces significant variability due to random influences such as start times, offer cycles, initialization delays, and concurrent processing on a 16-core system. Launching all applications on the system already takes approximately 300 ms. We collect metrics across 25 runs for each scenario (see Table III), analyzing minimum, mean, and maximum values. Additionally, we conduct a simple scalability analysis with 1 publisher and 1 to 50 subscribers to assess the impact on subscription latency.

Setup time: Setup time for each service measures the duration from a publisher’s first offer until all subscribers are acknowledged and validated. This is the most critical factor for the performance of the system, as it directly impacts the time when applications are functional.

Table III shows the min, mean, and max setup time per service for all three scenarios, revealing significant jitter likely caused by message loss, retransmission, and parallel processing. This time includes message passing, all cryptographic operations, DNSSEC lookups and possible retransmissions if messages are lost. Subscribers that miss an offer message

TABLE III
LATENCIES OF THE IVN STARTUP IN THE THREE SCENARIOS.

Metric	Min [ms]	Mean [ms]	Max [ms]
<i>Vanilla vsomeip</i>			
Service setup	3.261	747.074	1851.737
Network setup	1726.756	1881.639	2015.111
<i>Pre-deployed certificates</i>			
Service setup	5.351	1098.015	2965.049
Network setup	2066.790	2332.539	3368.344
Create signature	0.529	2.876	195.803
Verify signature	0.076	0.380	31.121
<i>DNSSEC, DANE, and DANCE</i>			
Service setup	6.129	731.618	1644.674
Network setup	1301.157	1482.990	2053.484
Create Signature	0.529	3.465	156.164
Verify Signature	0.089	0.413	41.707
Resolve Pub SVCB	0.171	1320.750	3222.964
Resolve Sub TLSA	0.092	185.644	1345.219
Resolve Pub TLSA	0.082	16.688	149.494

have to wait for the next repetition. The total setup time of all services on the network ranges from 1.3 s to 3.3 s across scenarios. Over the different runs, however, random factors may have significant impact, for example, *pre-deployed certificates* perform slower than the scenario with DNSSEC lookups.

To further analyze the subscription setup time, we create a small study with 1 publisher and 1 to 50 subscribers. Figure 4 shows the mean setup time for all simultaneous subscriptions. All three approaches can establish the 50 subscriptions for the service within milliseconds. Security operations add approximately 2 ms compared to vanilla SOME/IP, with DNSSEC resolution adding another 1 ms.

Cryptographic operations: Both secure scenarios require cryptographic operations including nonce generation, message signing/verification, and hash generation. These operations contribute significantly to setup delays (see Table III), with similar performance between scenarios. The DNSSEC variant also validates the endpoint, comparing the SVCB entry with the offer message, which has a negligible impact with an average of 0.000,3 ms. Again, we observe a large jitter in the results, which is likely caused by the parallel execution of many crypto operations on the same system. In average, signatures take longer to generate (2.6 ms to 4.2 ms) than to verify (0.3 ms to 0.5 ms.). While cryptographic overhead is low compared to the 10 ms to 100 ms default for the initial delay, ECUs with multiple services may require HSM support for efficient operation.

DNSSEC resolution: Our authentication and authorization scheme requires two DNSSEC lookups per publisher and one per subscriber. These lookups take a minimum of 0.1 ms which is expected for local resolution on the same machine. Higher mean and maximum times likely result from message loss and the c-ares library’s timeouts. Limitations on the number of channels for parallel DNSSEC resolution may cause delays at peak times, with failed resolutions retried in the next offer

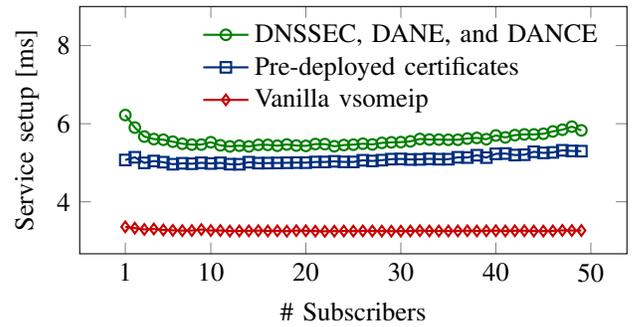


Fig. 4. Mean subscription setup time in a synthetic study for 1 publisher with 1 to 50 subscribers across the implementations.

cycle. Publisher resolution can begin before the first offer is sent, explaining why maximum publisher DNS resolution times can exceed total setup times.

C. Discussion and limitations

Our evaluation successfully demonstrates DNSSEC-based authentication and authorization for a realistic IVN. The results provide valuable insights into real-world deployment for an OEM fleet, with each vehicle zone requiring approximately 1000 records. Unlike alternative solutions, DNSSEC zones enable secure delegation of record management and certificate lifecycle within a scalable infrastructure without causing certificate association conflicts.

Limitations of the evaluation setup: The automotive network exhibited high performance jitter, likely due to parallel execution of all services on a single system. However, our scalability study revealed that setup time for individual services is not significantly affected by subscriber count and remains well below the 200 ms requirement 200 ms [4]. As service discovery is not real-time critical, this delay is acceptable for initial service setup.

Network transmission times were not accounted for as switches were emulated using Open vSwitch in Mininet. Measuring performance within the application may also affect results; alternatives such as tracing the communication and measuring delays between messages could provide additional insights. Additionally, decoupling applications from service discovery could reduce discovery instances and potentially improve performance.

Message loss may have occurred during initialization due to message floods, with numerous DNSSEC resolutions, **Find**, and **Offer** messages processed simultaneously. Optimizations such as improved scattering of discovery messages, DNS resolutions, or parallel execution configurations could enhance performance. These challenges are not unique to our approach and also affect the *vanilla vsomeip* scenario.

DNSSEC resolution: Resolution times for single requests have shown to be very small in a local environment. DNSSEC resolvers are well-optimized for resolving multiple queries, which makes the delay for a single query negligible. However, our implementation is not optimized for parallel resolution of mass requests, resulting in significant jitter. In real-world

scenarios, optimizing message loss handling and discovery process retransmissions would be crucial for reducing overall setup time.

Our evaluation did not include remote DNSSEC resolution, for which performance has been studied before – in our tests it took 10 ms to 40 ms to resolve `cloudflare.com` with DNSSEC validation. We anticipate DNSSEC resolution to occur during update procedures rather than vehicle startup, minimizing its impact on performance.

Credential management: The evaluation setup indicates that nodes resolve an average of 70 TLSA certificates via DNSSEC for remote publishers and subscribers. With the SaaS transformation of the automotive industry, we expect this number to grow as more services become optional, are regularly updated, or installed post-sale.

The traditional approach using pre-deployed certificates presents several management challenges. Nodes must securely store credentials even for applications that are not currently installed, to ensure compatibility with potential future updates or newly added services. This requirement creates significant management overhead, as already discussed in section V-A. While reducing the number of certificates (*e.g.*, to one per ECU) could ease deployment, it would significantly limit service mobility and reduce the granularity of access control options.

Our DNSSEC-based approach streamlines certificate management. By securely storing certificates in the DNSSEC zone, the system enables on-demand fetching of required credentials, substantially reducing the need for nodes to maintain comprehensive sets of unused remote service credentials. Publishers and subscribers can fetch the required certificates on demand, minimizing unnecessary maintenance. Caching public certificates on nodes can improve performance to levels comparable with pre-deployed certificate systems. However, this requires careful management of certificate lifetimes to prevent use of outdated or revoked credentials. The architecture supports parallel existence of multiple SVCB or TLSA records during service updates or certificate transitions, ensuring continuous secure operation throughout update processes and facilitating efficient certificate revocation when necessary.

Optimizations with session resumptions: Further optimization opportunities exist through session resumption, as publishers and subscribers often interact repeatedly across multiple vehicle operations. This would reduce cryptographic handshake overhead by reusing session keys or IDs. Nodes can also cache DNS responses, honoring the TTL value of records, to further minimize setup delays.

Given that IVNs remain stable over consecutive trips, a full cryptographic handshake might only be necessary once per trip. For instance, all sessions could be reinitialized during vehicle charging. During operation, periodic offer and subscribe messages could maintain session keys, improving efficiency without compromising security.

VIII. CONCLUSIONS AND OUTLOOK

Automotive service architectures lack robust security. The increasingly connected vehicles, though, require authentication and authorization—including a robust and secure credential management. Currently proposed static and pre-deployed certificates pose significant risks over the long lifespan of vehicles, as NIST recommends authentication keys to be replaced every one or other year [38].

In this paper, we introduced a key management solution for the automotive ecosystem that decouples service provider keys from service deployment keys using DNSSEC TLSA records. This approach largely simplifies security management between third party service providers and OEMs. Our solution leverages DNSSEC, DANE, and DANCE to authenticate automotive publish-subscribe systems and introduces implicit authorization through the presence of TLSA records. We identified deployment strategies for automotive applications, including offline verification of security credentials and secure key lifecycle management.

We designed DNSSEC-based authentication for SOME/IP middleware. Our security analysis revealed effective protection against threats such as spoofing identities, unauthorized access, and denial-of-service attacks previously apparent in SOME/IP. To mitigate the impact of private key compromises, each vehicle uses unique key pairs authenticated against its own DNSSEC zone. This ensures that a key compromise affects solely one vehicle, preserving the security of others.

We conducted a performance evaluation in a realistic in-vehicle setting, confirming the feasibility of our solution for authenticating and authorizing automotive services. Unlike pre-deployed certificates, our approach reduces dependence on locally stored credentials for all interacting parties while maintaining offline operability through a local DNSSEC resolver. This is a significant advantage over traditional PKI solutions, which require connectivity to external CAs.

This promising approach opens three future research directions. First, performance evaluations in real-world vehicle setups using dedicated hardware and switches will reveal additional insights into the scalability and robustness of our solution. These could also investigate optimizations through different key and signature algorithms and post-quantum cryptography. Additional work on session resumption could further optimize the solution by reducing overhead during vehicle startup, in particular as services are likely to communicate with the same peers across multiple trips. Finally, even though this work focuses on automotive applications, the proposed approach could extend to other local networks, such as industrial control systems or smart home environments, and demonstrate its broader applicability.

REFERENCES

- [1] R. N. Charette, “How Software Is Eating the Car,” *IEEE Spectrum*, Jun. 2021. [Online]. Available: <https://spectrum.ieee.org/software-eating-car>
- [2] D. F. Blanco, F. L. Mouël, T. Lin, and M.-P. Escudié, “A comprehensive survey on Software as a Service (SaaS) transformation for the automotive systems,” *IEEE Access*, pp. 73 688–73 753, 2023.

- [3] P. Laclau, S. Bonnet, B. Ducourthial, X. Li, and T. Lin, "Enhancing Automotive User Experience with Dynamic Service Orchestration for Software Defined Vehicles," 2024.
- [4] T. Häckel, P. Meyer, F. Korf, and T. C. Schmidt, "Secure Time-Sensitive Software-Defined Networking in Vehicles," *IEEE Transactions on Vehicular Technology*, vol. 72, no. 1, pp. 35 – 51, January 2023. [Online]. Available: <https://doi.org/10.1109/TVT.2022.3202368>
- [5] S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, S. Savage, K. Koscher, A. Czeskis, F. Roesner, T. Kohno *et al.*, "Comprehensive experimental analyses of automotive attack surfaces." in *USENIX Security Symposium*. San Francisco, 2011, pp. 77–92.
- [6] D. Zelle, C. Plappert, R. Rieke, D. Scheuermann, and C. Krauß, "Threat-Surf: A method for automated Threat Surface assessment in automotive cybersecurity engineering," *Microprocessors and Microsystems*, vol. 90, p. 104461, Apr. 2022.
- [7] Upstream, "2024 Global Automotive Cybersecurity Report," Upstream, Tech. Rep., 2024. [Online]. Available: https://info.upstream.auto/hubfs/Security_Report/Security_Report_2024/Upstream_2024_Global_Automotive_Cybersecurity_Report.pdf
- [8] C. Miller and C. Valasek, "Remote Exploitation of an Unaltered Passenger Vehicle," *Black Hat USA*, vol. 2015, p. 91, 2015.
- [9] International Organization for Standardization (ISO), "Road vehicles – Cybersecurity engineering," ISO, Geneva, CH, Standard ISO/SAE 21434:2021, 2021.
- [10] J. Yu and F. Luo, "A Systematic Approach for Cybersecurity Design of In-Vehicle Network Systems with Trade-Off Considerations," *Security and Communication Networks*, vol. 2020, pp. 1–14, Aug. 2020.
- [11] K. Iehira and H. Inoue, "Feasibility Assessment of Denial-of-Service Attacks by Analyzing SOME/IP-SD State Transition Models," in *2024 IEEE 48th Annual Computers, Software, and Applications Conference (COMPSAC)*, vol. 7. IEEE, Jul. 2024, pp. 1732–1738.
- [12] AUTOSAR, "SOME/IP Protocol Specification," AUTOSAR, Standard 696, 2024, part of Adaptive Platform, Standard Release R24-11.
- [13] —, "SOME/IP Service Discovery Protocol Specification," AUTOSAR, Standard 802, 2024, part of Foundation, Standard Release R24-11.
- [14] D. Zelle, T. Lauser, D. Kern, and C. Krauß, "Analyzing and Securing SOME/IP Automotive Services with Formal and Practical Methods," in *The 16th International Conference on Availability, Reliability and Security*. ACM, Aug. 2021.
- [15] M. Iorio, M. Reineri, F. Risso, R. Sisto, and F. Valenza, "Securing SOME/IP for In-Vehicle Service Protection," *IEEE Transactions on Vehicular Technology*, vol. 69, no. 11, pp. 13450–13466, Nov. 2020.
- [16] B. Ma, S. Yang, Z. Zuo, B. Zou, Y. Cao, X. Yan, S. Zhou, and J. Li, "An Authentication and Secure Communication Scheme for In-Vehicle Networks Based on SOME/IP," *Sensors*, vol. 22, no. 2, p. 647, Jan. 2022.
- [17] S. Lee, W. Choi, and D. H. Lee, "Protecting SOME/IP Communication via Authentication Ticket," *Sensors*, vol. 23, no. 14, p. 6293, Jul. 2023.
- [18] 3rd D. E. Eastlake, "Domain Name System Security Extensions," IETF, RFC 2535, March 1999. [Online]. Available: <https://doi.org/10.17487/RFC2535>
- [19] B. Schwartz, M. Bishop, and E. Nygren, "Service Binding and Parameter Specification via the DNS (SVCB and HTTPS Resource Records)," IETF, RFC 9460, November 2023. [Online]. Available: <https://doi.org/10.17487/RFC9460>
- [20] J. Walrand, M. Turner, and R. Myers, "An Architecture for In-Vehicle Networks," *IEEE Transactions on Vehicular Technology*, vol. 70, no. 7, pp. 6335–6342, Jul. 2021.
- [21] Institute of Electrical and Electronics Engineers, "IEEE Standard for Local and Metropolitan Area Networks—Bridges and Bridged Networks," IEEE, Tech. Rep. IEEE Std 802.1Q-2022 (Revision of IEEE Std 802.1Q-2018), Dec. 2022.
- [22] L. Lo Bello and W. Steiner, "A Perspective on IEEE Time-Sensitive Networking for Industrial Communication and Automation Systems," *Proceedings of the IEEE*, vol. 107, no. 6, pp. 1094–1120, Jun. 2019.
- [23] Object Management Group (OMG), "Data Distribution Service," OMG, Standard DDS 1.4, 2015. [Online]. Available: <https://www.omg.org/spec/DDS/1.4>
- [24] A. Ioana, A. Korodi, and I. Silea, "Automotive IoT Ethernet-Based Communication Technologies Applied in a V2X Context via a Multi-Protocol Gateway," *Sensors*, vol. 22, no. 17, p. 6382, Aug. 2022.
- [25] M. Dibaei, X. Zheng, K. Jiang, R. Abbas, S. Liu, Y. Zhang, Y. Xiang, and S. Yu, "Attacks and defences on intelligent connected vehicles: a survey," *Digital Communications and Networks*, vol. 6, no. 4, pp. 399–421, Nov. 2020.
- [26] J.-P. Monteuiis, A. Boudguiga, J. Zhang, H. Labiod, A. Serval, and P. Urien, "SARA: Security Automotive Risk Analysis Method," in *Proceedings of the 4th ACM Workshop on Cyber-Physical System Security*. ACM, May 2018. [Online]. Available: <https://dl.acm.org/doi/pdf/10.1145/3198458.3198465>
- [27] V. L. L. Thing and J. Wu, "Automonomous Vehicle Security: A Taxonomy of Attacks and Defences," in *2016 IEEE International Conference on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData)*. Piscataway, NJ, USA: IEEE Press, Dec. 2016.
- [28] M. D. Pesé, K. Schmidt, and H. Zweck, "Hardware/Software Co-Design of an Automotive Embedded Firewall," in *SAE Technical Paper*. SAE International, Mar. 2017. [Online]. Available: <https://doi.org/10.4271/2017-01-1659>
- [29] L. Xue, Y. Liu, T. Li, K. Zhao, J. Li, L. Yu, X. Luo, Y. Zhou, and G. Gu, "SAID: State-aware Defense Against Injection Attacks on In-vehicle Network," in *31st USENIX Security Symposium (USENIX Security 22)*. Boston, MA: USENIX Association, Aug. 2022, pp. 1921–1938. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity22/presentation/xue-lei>
- [30] G. K. Rajbahadur, A. J. Malton, A. Walenstein, and A. E. Hassan, "A Survey of Anomaly Detection for Connected Vehicle Cybersecurity and Safety," in *2018 IEEE Intelligent Vehicles Symposium (IV)*. IEEE Press, Jun. 2018.
- [31] P. Meyer, T. Häckel, T. Lübeck, F. Korf, and T. C. Schmidt, "A Framework for the Systematic Assessment of Anomaly Detectors in Time-Sensitive Automotive Networks," in *Proc. of the 15th IEEE Vehicular Networking Conference (VNC)*. IEEE, May 2024, pp. 57–64. [Online]. Available: <https://doi.org/10.1109/VNC61989.2024.10576017>
- [32] P. Mundhenk, A. Paverd, A. Mrowca, S. Steinhorst, M. Lukasiewicz, S. A. Fahmy, and S. Chakraborty, "Security in Automotive Networks: Lightweight Authentication and Authorization," p. 1–27, Mar. 2017. [Online]. Available: <https://doi.org/10.1145/2960407>
- [33] Q. Hu and F. Luo, "Review of Secure Communication Approaches for In-Vehicle Network," *International Journal of Automotive Technology*, vol. 19, no. 5, pp. 879–894, Sep. 2018.
- [34] S. Sanwald, L. Kaneti, M. Stöttinger, and M. Böhner, "Secure Boot Revisited: Challenges for Secure Implementations in the Automotive Domain," *SAE International Journal of Transportation Cybersecurity and Privacy*, vol. 2, no. 2, Aug. 2020.
- [35] S. Mahmood, H. N. Nguyen, and S. A. Shaikh, "Systematic threat assessment and security testing of automotive over-the-air (OTA) updates," *Vehicular Communications*, vol. 35, p. 100468, Jun. 2022.
- [36] M. Howard and S. Lipner, *The Security Development Lifecycle*. Microsoft Press, 2006.
- [37] D. Zelle, C. Krauß, H. Strauß, and K. Schmidt, "On Using TLS to Secure In-Vehicle Networks," in *Proceedings of the 12th International Conference on Availability, Reliability and Security*. ACM, Aug 2017.
- [38] NIST, "Recommendation for Key Management: Part 1 – General," NIST, Washington, D.C., Tech. Rep. SP 800-57 Part 1 Rev. 5, 2020.
- [39] Object Management Group, "DDS Security," Online, OMG, Standard DDS-SECURITY 1.1, Jul. 2018. [Online]. Available: <https://www.omg.org/spec/DDS-SECURITY/1.1>
- [40] D. Cooper, S. Santesson, S. Farrell, S. Boeyen, R. Housley, and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile," IETF, RFC 5280, May 2008. [Online]. Available: <https://doi.org/10.17487/RFC5280>
- [41] P. Hoffman and J. Schlyter, "The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA," IETF, RFC 6698, August 2012. [Online]. Available: <https://doi.org/10.17487/RFC6698>
- [42] D. Püllen, F. Frank, M. Christl, W. Liu, and S. Katzenbeisser, "A Security Process for the Automotive Service-Oriented Software Architecture," *IEEE Transactions on Vehicular Technology*, vol. 73, no. 4, pp. 5036–5053, Apr. 2024.
- [43] M. Mueller, T. Häckel, P. Meyer, F. Korf, and T. C. Schmidt, "Authenticated and Secure Automotive Service Discovery with DNSSEC and DANE," in *Proc. of the 14th IEEE Vehicular Networking Conference (VNC)*. IEEE, April 2023, pp. 231–238. [Online]. Available: <https://doi.org/10.1109/VNC57357.2023.10136332>

- [44] E. Osterweil, P. F. Tehrani, T. C. Schmidt, and M. Wählisch, "From the Beginning: Key Transitions in the First 15 Years of DNSSEC," *Transactions on Network and Service Management (TNSM)*, vol. 19, no. 4, pp. 5265–5283, December 2022. [Online]. Available: <https://doi.org/10.1109/TNSM.2022.3195406>
- [45] A. Wilson, S. Huque, O. E. Johansson, and M. Richardson, "An Architecture for DNS-Bound Client and Sender Identities," Internet Engineering Task Force, Internet-Draft draft-ietf-dance-architecture-07, Oct. 2024, work in Progress. [Online]. Available: <https://datatracker.ietf.org/doc/draft-ietf-dance-architecture/07/>
- [46] V. Dukhovni and W. Hardaker, "The DNS-Based Authentication of Named Entities (DANE) Protocol: Updates and Operational Guidance," IETF, RFC 7671, October 2015. [Online]. Available: <https://doi.org/10.17487/RFC7671>
- [47] T. Finch, M. Miller, and P. Saint-Andre, "Using DNS-Based Authentication of Named Entities (DANE) TLSA Records with SRV Records," IETF, RFC 7673, October 2015. [Online]. Available: <https://doi.org/10.17487/RFC7673>
- [48] H. Lee, I. Ashiq, M. Müller, R. van Rijswijk-Deji, T. Kwon, and T. Chung, "Under the Hood of DANE Mismanagement in SMTP," in *Proc. of the USENIX Security Symposium*. Boston, MA: USENIX Association, Aug. 2022. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity22/presentation/lee>
- [49] M. Müller, M. Thomas, D. Wessels, W. Hardaker, T. Chung, W. Toorop, and R. v. Rijswijk-Deij, "Roll, Roll, Roll Your Root: A Comprehensive Analysis of the First Ever DNSSEC Root KSK Rollover," in *Proceedings of the Internet Measurement Conference*, ser. IMC '19. New York, NY, USA: ACM, 2019, p. 1–14.
- [50] H. Lee, A. Gireesh, R. van Rijswijk-Deij, T. T. Kwon, and T. Chung, "A Longitudinal and Comprehensive Study of the DANE Ecosystem in Email," in *Proc. of the USENIX Security Symposium*. USENIX Association, Aug 2020, pp. 613–630. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity20/presentation/lee-hyeonmin>
- [51] P. F. Tehrani, R. Hiesgen, T. Lübeck, T. C. Schmidt, and M. Wählisch, "Do CAA, CT, and DANE Interlink in Certificate Deployments? A Web PKI Measurement Study," in *Proc. of Network Traffic Measurement and Analysis Conference (TMA)*. Piscataway, NJ, USA: IEEE Press, May 2024, pp. 1–11. [Online]. Available: <https://doi.org/10.23919/TMA62044.2024.10559089>
- [52] R. Barnes, "Use Cases and Requirements for DNS-Based Authentication of Named Entities (DANE)," IETF, RFC 6394, October 2011. [Online]. Available: <https://doi.org/10.17487/RFC6394>
- [53] W. Diffie and M. Hellman, "New Directions in Cryptography," *IEEE Transactions on Information Theory*, vol. 22, no. 6, pp. 644–654, Nov. 1976.
- [54] D. Dolev and A. C. Yao, "On the Security of Public Key Protocols," in *22nd Annual Symposium on Foundations of Computer Science (sfcs 1981)*. IEEE, Oct. 1981.
- [55] T. Chung, R. van Rijswijk-Deij, B. Chandrasekaran, D. Choffnes, D. Levin, B. M. Maggs, A. Mislove, and C. Wilson, "A Longitudinal, End-to-End View of the DNSSEC Ecosystem," in *26th USENIX Security Symposium (USENIX Security 17)*. Vancouver, BC: USENIX Association, Aug. 2017, pp. 1307–1322. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/chung>
- [56] BMW AG, "vsomeip," Online, version 3.4.10. [Online]. Available: <https://github.com/COVESA/vsomeip/tree/3.4.10>
- [57] "c-ares," Online, version 3.4.10. [Online]. Available: https://github.com/c-ares/c-ares/tree/cares-1_18_1
- [58] "Crypto++," Online, version 8.9.0. [Online]. Available: https://github.com/weidai11/cryptopp/tree/CRYPTOPP_8_9_0
- [59] "Mininet," Online, version 2.3.0. [Online]. Available: <https://github.com/mininet/mininet/tree/2.3.0>
- [60] NLnetLabs, "NSD," Online, version 4.8.0. [Online]. Available: https://github.com/NLnetLabs/nsd/tree/NSD_4_8_0_REL
- [61] OpenSSL Software Foundation Inc, "OpenSSL," Online, version 3.0.2. [Online]. Available: <https://github.com/openssl/openssl/tree/openssl-3.4.0>

APPENDIX A
SOME/IP HEADER FORMAT, SERVICE ENTRIES AND CONFIGURATION OPTIONS

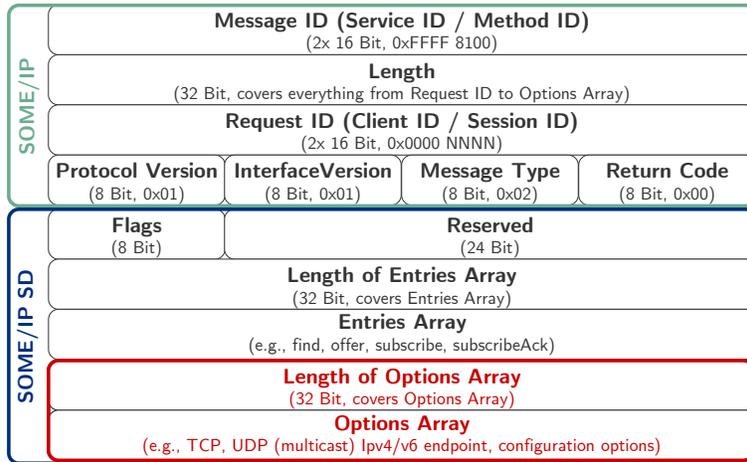


Fig. A.1. SOME/IP message layout for service discovery messages [13].

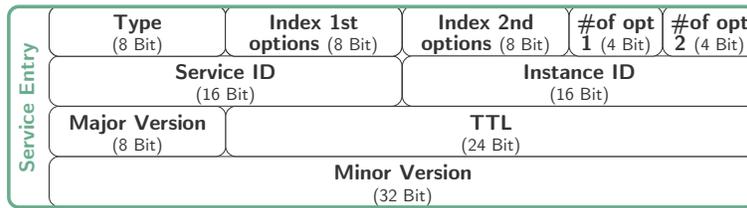


Fig. A.2. Structure of SOME/IP service discovery entry included in service discovery messages. [13]

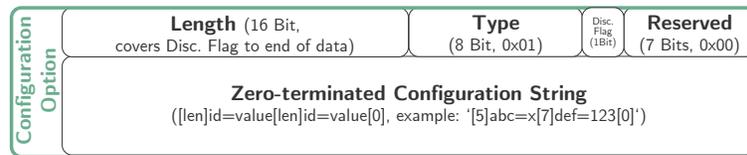


Fig. A.3. SOME/IP configuration option layout included in service discovery messages [13].

APPENDIX B
SOME/IP SERVICE DISCOVERY STATE MACHINES WITH INTEGRATIONS FOR DNSSEC OPERATIONS

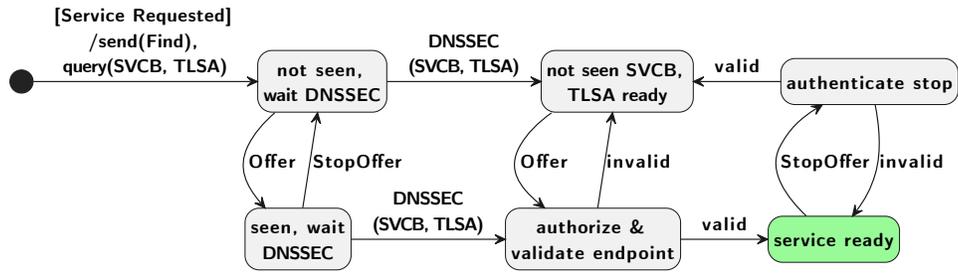


Fig. B.1. State machine for SOME/IP discovery of requested services at the client side including DNS interactions not including the initial wait phases, retransmissions, and timeout handling.

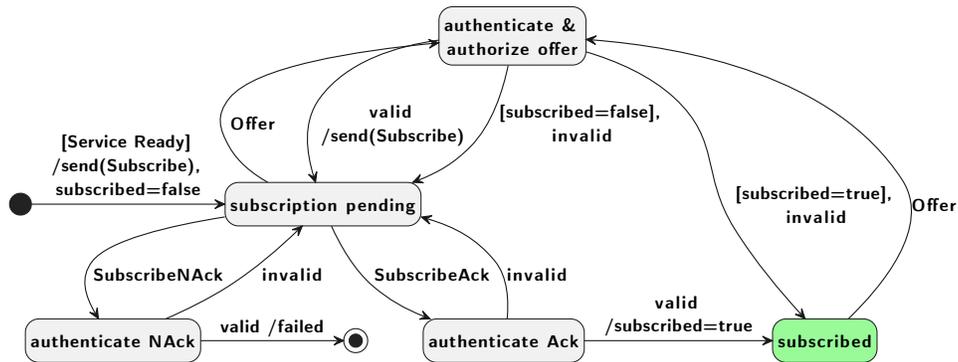


Fig. B.2. State machine for SOME/IP subscription handling with authentication and authorization at the client side not including, retransmissions, and timeout handling.

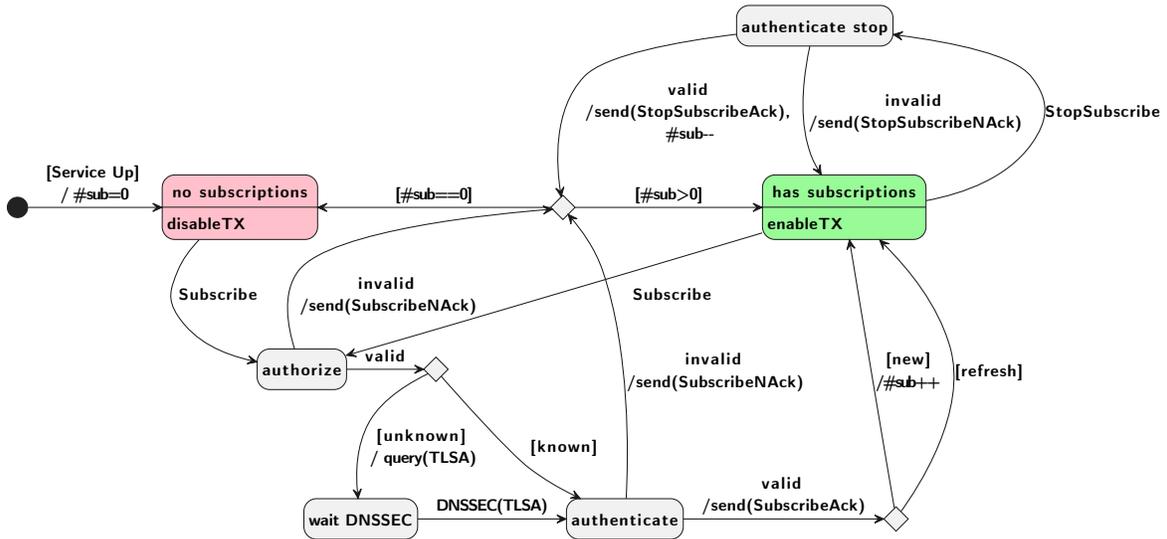


Fig. B.3. State machine for SOME/IP subscription handling at the server side including DNS interactions, authentication, and authorization, but not including, retransmissions, and timeout handling.