Before the

Department of Transportation

National Highway Traffic and Safety Administration (NHTSA)

NPRM: Federal Motor Vehicle Safety Standards; V2V Communications

Docket No. NHTSA-2016-0126

Comments of Electronic Frontier Foundation

April 12, 2017

Submitted by:

Lee Tien
Jamie Williams
Seth Schoen
Roland Shoemaker
Electronic Frontier Foundation
815 Eddy Street
San Francisco, CA 94109
Telephone: (415) 436-9333
jamie@eff.org

The Electronic Frontier Foundation (EFF) submits the following comments in response to NHTSA's Notice of Proposed Rulemaking (NPRM) regarding "Federal Motor Vehicle Safety Standards; V2V Communications" (Docket No. NHTSA-2016-0126). EFF is a member-supported, nonprofit, public interest organization dedicated to protecting privacy, civil liberties, and innovation in the digital age. Founded in 1990, EFF represents tens of thousands of dues-paying members, including consumers, hobbyists, computer programmers, entrepreneurs, students, teachers, and researchers. EFF and its members are united in their commitment to ensuring that new technologies are not used to undermine privacy and security.

EFF has four significant concerns regarding the V2V system outlined in the NPRM (each discussed in more detail below):

(1) The proposal tries unsuccessfully to mitigate the privacy risk presented by V2V and will not prevent vehicle tracking;

(2) The proposed application of a Public Key Infrastructure (PKI) is unduly complicated and will create potentially dangerous misconceptions regarding whether the contents of basic safety messages (BSMs) are "safe" and to be trusted;

(3) The proposal fails to address the serious security concerns presented by V2V—leaving drivers and passengers at potentially grave risk; and

(4) The proposal is inefficient from a common sense, cost-benefit perspective; the technology is expensive and, if implemented, will be outpaced by other communications technology by the time it is fully deployed.

Unless and until the privacy and security concerns with the NPRM's V2V proposal are resolved, it would be irresponsible for NHTSA to recommend—let alone mandate—the technology. Furthermore, given the high likelihood that V2V will be soon be outpaced by other technology, it would be imprudent for NHTSA to mandate V2V even if the privacy and security concerns were resolved. And should NHTSA proceed with any V2V proposal—either now or in the future after the privacy and security concerns raised herein and by other commenters are addressed—consumers must be provided with the option and right to disable V2V in order to protect their privacy, security, and safety.

I. The Proposal Fails to Mitigate the Privacy Risks of V2V.

First, while NHTSA attempted to address the significant privacy risk presented by V2V—that V2V transmissions could be linked to specific vehicles (or persons) and thus used to track their physical location—the V2V system outlined in the NPRM is not sufficient to prevent vehicle tracking. It will not, as the agency suggests, "make it difficult to track through space and time specific vehicles, owners or drivers on a persistent basis." 82 FR 3869. The proposal would in fact give parties motivated to track vehicles a straightforward way to do so. This is a very real threat; it is already clear from other contexts—specifically, Automated License Plate Reader (ALPR) technology—that people will go to great lengths to track and amass vehicle location data. The privacy issues must be adequately resolved before any V2V proposal moves forward.

A. Physical Location is Highly Sensitive.

Data regarding an individual's physical location—including details regarding the particular route taken or the physical start and end points of a trip—is extraordinarily sensitive. It can paint an intimate portrait of a person's daily life and reveal private information, such as confidential personal and professional relationships, medical information, religious affiliation, participation in stigmatized activities, and more. *See United States v. Jones*, 132 S.Ct. 945, 955 (2012) (Sotomayor, J., concurring) (noting that location information, such as GPS data, can "generate[] a precise, comprehensive record of a person's public movements that reflects a wealth of detail about her familial, political, professional, religious, and sexual associations"). Disclosing a person's physical location also facilitates stalking and a wide range of crimes against people or their property (such as burglary when a person is known to be out of town).

B. The Proposed System Would Make it Possible for a Sensor Network to Track a Vehicle Over the Course of One Day.

The proposed system relies on message authentication via certificates—*i.e.*, the initiating V2V device would sign its basic safety messages (BSMs) with a cryptographic key, attested as belonging to a vehicle by a certificate, thereby authenticating itself as an accurate source of information. Meanwhile, the receiving V2V device would verify the authenticity of the certificate and confirm the validity of the signature, thereby ensuring that the sender is a vehicle certified as participating in the system. NHTSA recognizes that using a single certificate per vehicle would enable an observer collecting V2V transmissions to track that vehicle simply by associating the BSMs coming from a single V2V sender with a single vehicle. NHTSA thus proposes a system of rotating certificate credentials, whereby each V2V-enabled vehicle would have 20 valid certificates per week, which would change at least once every 5 minutes. Each car

would thus rotate through 1,040 certificates over the course of a year. NHTSA suggests that this proposal strikes the appropriate balance "between the competing interests of maximizing privacy protections and technological practicability." 82 FR 3911.

This proposal, while well intentioned, will not actually protect privacy. Vehicles transmitting V2V communications will still be trackable. Specifically, NHTSA has failed to adequately account for the need to protect privacy against *systematic attempts* that will undoubtedly be made to monitor and record BSMs for the purpose of tracking vehicles.

Even assuming rotating certificates were an appropriate approach (as outlined in the next section, the proposed PKI also presents major implementation challenges), rotating through a mere 20 different identities, every 100 minutes, over the course of one week will not protect a vehicle's privacy. While a human being might find it confusing to remember 20 different identities for the same vehicle, it would be straightforward for a computer to analyze data collected via a sensor network and identify a vehicle over the course of one day—including associating the full set of certificates assigned to the vehicle.

After a sensor network has determined the identity of a vehicle over the course of a day—via its 20 rotating certificates—it would be able to immediately identify the vehicle for the remainder of the week. The vehicle would be *completely deanonymized* for the course of that week, and for the corresponding week in any subsequent year. The sensor network would merely have to complete the same process every week, but this would be feasible given the straightforward nature of the process. And because "human mobility traces are highly unique,"[1] it would be easy, in the case of a vehicle used in its typical way, to recognize and track a vehicle from week to week, even as the vehicle's list of 20 assigned certificates changed. Indeed, a 2009 study by the Palo Alto Research Center (PARC) showed that 5% of Americans—*i.e.*, more than 15 million people—could be uniquely identified by simply pairing data regarding their home and work areas.[2] And this does not even take into account other locations that people routinely or habitually visit—such as schools, daycare facilities, yoga studios, or grocery stores—or patterns regarding the times at which they visit these locations. And by combining vehicle location history aggregated over time with other information or data sources (such as databases tracking employment and home addresses), it will likely even be possible to identify who exactly is behind the wheel.

Recent computer science research has achieved remarkable success at classifying and recognizing entities from noisy data, readily finding the nearest matches even when data sets are not precisely identical. This has often enabled practical deanonymization of pseudo-identifiers—

---

[1] *See* Yves-Alexandre de Montjoye et al., Unique in the Crowd: The privacy bounds of human mobility, Scientific Reports 3, Article Number 1376 (Mar. 23, 2013), available at http://www.nature.com/articles/srep01376 (finding that "four spatio-temporal points are enough to uniquely identify 95% of the individuals" in a study of 15 months of human mobility data for 1.5 million individuals, with a dataset where the location of an individual is specified hourly).

[2] Philippe Golle & Kurt Partridge, Palo Alto Research Center, *On the Anonymity of Home/Work Location Pairs* (2009), https://crypto.stanford.edu/~pgolle/papers/commute.pdf.

an entire field of research in its own right. In an age of machine learning and artificial intelligence, NHTSA should be aware that computer algorithms will be able to *quickly and easily* analyze V2V data to track vehicles over time—notwithstanding attempts to prevent this with key rotation. While this form of deanonymization is not 100% foolproof, early attempts at deanonymizing location-related data sets have shown striking success.

The proposal also fails to account for the privacy implications of both (a) the vast amount of metadata contained in the cleartext[3] of BSMs, and (b) the metadata that can be derived from the chain of trust built between the leaf certificate and the root certificate pursuant to the proposed Public Key Infrastructure (PKI).[4] For instance, in a design (like the one proposed) involving multiple root certificate authorities (CAs), intermediate CAs, and pseudonym CAs—all controlled by different manufacturers and departments within these manufacturers—it would be trivial to significantly narrow down the number of vehicles from which any given BSM could have possibly transmitted by simply analyzing the path between the leaf certificate and the root certificate.

In addition to metadata, BSM content—with information such as a vehicle's speed or lane changes—could itself be used to help identify vehicles. By "eavesdropping" on unencrypted BSMs, a sensor network or anyone else with access to the BSMs could glean information about how and where the vehicle is moving and use that information to try to identify it.

Combining BSM content and metadata with the relative simplicity of tracking a single vehicle that is transmitting numerous messages per second over a short period, it is evident that simply observing a vehicle for the relatively short rotation period (*i.e.*, 100 minutes) would likely be sufficient to fully deanonymize it. There are many parties who would be interested in vehicle tracking, and this proposal provides such parties with a straightforward way of doing so.

C. The Proposal is Premised on a Flawed Threat Model.

The NPRM's proposal seems to be based on a very weak—and flawed—threat model. Specifically, NHTSA seems to be making two patently false assumptions:

(1) that broadcasts will typically only be received by other vehicles—or receivers—in the immediate physical vicinity of a vehicle; and

(2) that efforts to track or identify vehicles would be casual, sporadic, or incidental.

First, based on experience, it is likely that V2V transmissions will be received at a distance far beyond their designed or intended read range.[5] Thus, it is likely that these

---

[3] "Cleartext" is readable data transmitted or stored "in the clear," *i.e.*, unencrypted. *See* Dan Cornell, Cleartext vs. Plaintext vs. Ciphertext vs. Plaintext vs. Clear Text, Denim Group (Oct. 19, 2007), https://www.denimgroup.com/blog/2007/10/cleartext-vs-pl/.

[4] The flaws of the proposed use of PKI are outline in Section II, *infra*.

[5] There are many examples of radio signals being successfully received and interpreted at distances far greater than expected. *See, e.g.,* Peter Shipley, Open WLANS: The early results of

transmissions will be received by distant receivers—from fixed receiving stations to receivers attached to drones or aircraft—which could potentially track a large number of vehicles over a wide area at once, all via a single V2V interception device.

Second, it has been demonstrated *in many other contexts* involving the exposure of personal information via some technical means that people will go to *considerable* lengths to collect, store, and share such information.[6] NHTSA has failed to account for the threat of long-term, systematic, large-scale efforts to re-identify vehicles. The gap between what can be accomplished with casual observation and what can be accomplished with a dedicated, organized, large-scale effort is a wide one. For a computer (or for an organized project), the projected 1,040 identities (over the course of one year) is a small number. Observable reuse of identities will be commonplace.[7]

This is particularly troubling given that efforts to track vehicles over time are already underway via the use of Automated License Plate Reader (ALPR) technology. The extensive use of this technology demonstrates that governments and commercial entities will spend a great deal of time and money to track vehicles' whereabouts. Indeed, the collection and sale of ALPR data is growing into an entire industry. There are currently two main private companies—DRN and MVTrac—that hire contractors to collect license plate data from cars across the United States.[8] DRN's database contains over 2 billion records, and MVTrac said in 2012 that it has data on a "large majority" of the vehicles in the United States, while ALPR tracking seems to have grown significantly since then.[9] Both companies have lobbied and litigated against legislative attempts

---

WarDriving (2001), https://blyx.com/public/wireless/open_wireless_lans.pdf (researchers were able to make a connection to a network with an intended range of around 150 feet from around 25 miles away).

[6] This includes government surveillance programs seeking access to data in bulk and commercial tracking efforts, such as online behavior profiling (for which data brokers spend a great deal of money and resources gathering personal information for use ways that the consumer never expected and that are unrelated to the purpose for which the data was originally collected) and retail analytics (which has involved placing sensors or tracking equipment at many locations and then trying to correlate the information collected with information from different databases and data sources, in order to track individuals' whereabouts over time—again in unexpected and unanticipated ways).

[7] The likelihood of observing the same identifier twice is very high and typically requires far fewer observations than one would tend to expect. Take the well-known "birthday paradox": In a room of just 23 people there's a 50% chance of two people having the same birthday; in a room of 70, that chance increases to 99.9%. *See* Wikipedia, Birthday Problem (last updated Apr. 8, 2017), https://en.wikipedia.org/wiki/Birthday_problem.

[8] Julia Angwin & Jennifer Valentino-Devries**,** New Tracking Frontier: Your License Plat*es*, *Wall Street Journal (*Sept. 29, 2012), https://www.wsj.com/articles/SB10000872396390443995604578004723603576296.

[9] EFF, Street Level Surveillance, Automated License Plate Readers: Frequently Asked Questions, https://www.eff.org/sls/tech/automated-license-plate-readers/faq#faq-How-many-law-

to limit collection and retention of data.[10] DRN and MVTrac share their data with banks, insurance companies, credit reporting agencies, and "auto recovery" (*i.e.*, repo) companies and assert that their data can help these companies find fraud and identity theft.[11] This data may also be of interest for marketing and advertising purposes, *e.g.*, a grocery store sending coupons or directing ads to individuals or households who's vehicle was recorded as being parked in the lot of a competitor. It may be of interest to divorce attorneys. DRN's parent company is Vigilant Solutions, which not only provides ALPR tools to law enforcement but also shares its vast database of privately-collected data with the government.[12]

NHTSA's V2V proposal would amplify the privacy threat presented by ALPR by making it significantly cheaper to get more reliable information about a vehicle's whereabouts, more of the time, in more situations, in a clandestine manner, and without requiring a line-of-sight to a vehicle's license plate. The prospect of more detailed, more reliable, and cheaper-and-easier-to-collect data about vehicle movements would surely be of interest to anyone interested in ALPR data. Indeed, it would likely be lucrative to combine V2V data with other sources of data, like ALPR data, mobile device data, or web-browsing data, to build a more complete profile on vehicles and the individuals associated with them. It would even be possible to combine V2V readers and ALPR readers into a single device—such as at a toll plaza or equipped on a vehicle, such as a police car or a car owned by a vehicle tracking company.[13]

---

enforcement-agencies-use-ALPRs.

[10] *See, e.g.*, Reuters/Mike Segar, License plate reader makers sue Arkansas for banning their tech, RT.com (June 18, 2014), https://www.rt.com/usa/166916-vigilant-drn-arkansas-alpr-lawsuit/.

[11] *See, e.g.*, DRN: Vehicle Location Data for Auto Lenders, Insurance Carriers and Recovery Professionals, http://drndata.com/ ("Our data helps lenders make right party contact to reduce charge-offs, insurers improve pricing at underwriting and claims investigations, and gives recovery agents the technology they need to recover more vehicles.").

[12] Cyrus Farivar, NYPD to conduct "virtual stakeouts," get alerts on wanted cars nationwide, *Ars Technic* (Mar. 2, 2015), https://arstechnica.com/tech-policy/2015/03/nypd-to-conduct-virtual-stakeouts-get-alerts-on-wanted-cars-nationwide/.

[13] Note that, in addition to the situations outlined above, monitoring infrastructure and databases can be used to re-associate the NPRM's proposed rotating identities with a license plate or other vehicle identity in several ways: (1) by combining V2V readers with readers for other vehicular identifiers, such as electronic toll collection tokens or tire pressure monitoring systems (TPMS); (2) by noting the moment at which a vehicle switches from using one identity to another, extrapolating from the vehicle's position and movement that the two identities in fact *must* correspond to the same vehicle; and (3) by noting patterns in a vehicle's presence over time at specific locations or along specific courses of travel that suggest that certain repeatedly observed identities do, in fact, represent the same vehicle. In addition, an attacker could simply try to compromise a central authority associated with the V2V's PKI to gain direct access to desired information that has not yet even been collected.

Other technologies might be able to achieve better privacy protections in the context of a system of this kind. Cryptography researchers have proposed measures involving revocable group signatures, among other technologies, which can make it harder for the general public to know whether or not two observations related to the same vehicle, while still allowing an authority to perform revocations. Multiparty computation has also been applied to collision detection without revealing objects' exact locations or trajectories. (Effectively, the former approach reveals a vehicle's location but not its identity, while the latter approach reveals a vehicle's identity but not its location.) Thus, the use of a "large" number of fixed pseudonyms represents a comparatively antiquated and unsophisticated approach against the backdrop of ongoing technical research in this area.

NHTSA's current proposal effectively relies on the false hope that people will not try very hard or very often to correlate observed pseudonymous identities in order to track or identify vehicles. This is not a realistic or reasonable solution to the privacy threat presented by V2V. The United States doesn't have a general-purpose data protection law. As we've seen in other contexts, we are therefore reliant on technology as a first line of defense against unauthorized uses of personal data—by preventing collection from being feasible in the first place. NHTSA's proposal of rotating identities is inadequate to mitigate the privacy threat presented by V2V location broadcasts. NHTSA must work with cryptographers to find better technical solutions that will actually protect privacy and prevent vehicle identification and tracking before moving forward with any V2V proposal.

II. The Proposed Use of a PKI Will Not Achieve its Purported Purpose of Ensuring that Messages are "Safe"; It Will Only Create Unnecessary Complexity and Confusion.

NHTSA proposes that V2V devices sign and verify BSMs using a Public Key Infrastructure (PKI). The proposal envisions a Root Certificate Authority (CA) as the trusted entity that will distribute digital certificates, which will delegate authority to Intermediate CAs to protect the Root CA from unnecessary exposure. The Intermediate CAs will in turn authenticate individual vehicles. The proposed application of PKI creates a chain of trust from the Root CA to the leaf certificates (*i.e.*, end-entity certificates) used by vehicles and enables revocation of trust in these leaf certificates. The proposal's implicit suggestion is that use of such a PKI will enable vehicles to assess whether a message is "safe"—*i.e.*, whether a receiver can view the contents of a message as truthful and rely on the message as a basis for decisions.

This proposal's application of PKI for determining whether a message is "safe" is misguided. It seems to misunderstand the limits and weaknesses of PKI—both in terms of the current understanding of technical functionality and the policy stability of large distributed systems. The proposal not only envisions a PKI larger and more complicated than anything in existence, but it also misses the entire purpose of distributed PKI systems—associating a key with an identity for the *sole purpose* of determining who or what produced a validly signed message. As a result, the proposal is not only unduly complex, but it will create widespread—and potentially quite dangerous—confusion about the level of confidence that should be placed in the contents of validly signed messaged.

A. The Proposed PKI is Overly Complicated and Fails to Account for Considerable Technical Challenges Presented in Much Smaller Systems.

The closest PKI system the proposal can be compared to is the WebPKI, which is arguably the largest existing PKI system. This system currently has *considerable technical challenges* surrounding issues such as (a) deprecating[14] insecure methods or designs that were previously considered secure, and (b) management of root trust stores[15] across various platforms. The NPRM's proposed PKI would be *magnitudes larger* than the WebPKI, which would necessitate introducing numerous significant components that do not exist in any other existing PKI system.

One widely cited challenge with the WebPKI involved deprecating the partially broken hash algorithm SHA-1.[16] The process took *four years*, and it was ultimately only possible because web browsers in the WebPKI (*i.e.*, the PKI's validating clients which act on behalf of end users) had the power to explicitly force any certificates using the SHA-1 algorithm to be considered invalid. We are extremely concerned that the NPRM's proposed PKI would have even greater problems rejecting certificates based on algorithms shown to be insecure after the PKI's establishment. In a self-regulatory environment—where vehicle manufacturers represent both the certificate authorities (CAs) and the clients—there would be no driving factors to force agility in deprecation of systems considered insecure that are already widely deployed.

This is particularly troubling given concerns regarding quantum computing. The proposed PKI is intended for long-term use, but it doesn't make any mention of the upcoming transition away from the signing algorithms current in use due to concerns regarding their strength in a post-quantum world.[17] The NSA's Information Assurance Directorate (IAD) has itself announced that system owners and developers should move away from RSA and elliptic

---

[14] "Deprecation" refers to the practice of discouraging the use of some feature, design, or practice, typically because it has been superseded or is no longer considered efficient or safe. Wikipedia, Deprecation (last modified Mar. 31, 2017), https://en.wikipedia.org/wiki/Deprecation.

[15] A trust store is used to store certificates from trusted CAs and are used to verify certificates. In WebPKI, if a server's certificate is signed by a recognized CA, the default trust store will already trust it (because it already trusts the recognized CA). *See, e.g.*, Javin Paul, Difference between trustStore vs keyStore in Java SSL, Java67 (Dec. 2012), http://www.java67.com/2012/12/difference-between-truststore-vs.html.

[16] *See* Chris Palmer & Ryan Sleevi, Gradually sunsetting SHA-1, Google Security Blog (Sept. 5, 2014), https://security.googleblog.com/2014/09/gradually-sunsetting-sha-1.html (explaining how Chrome began to sunset the SHA-1 cryptographic hash algorithm due to its insecurity).

[17] "Post-quantum cryptography refers to cryptographic algorithms (usually public-key algorithms) that are thought to be secure against an attack by a quantum computer." Wikipedia, Post-Quantum Cryptography (last updated Apr. 10, 2017), https://en.wikipedia.org/wiki/Post-quantum_cryptography.

curve cryptography (ECC) and instead adopt PQ algorithms in the near future,[18] in order to prevent the need for rapid transition away from these algorithms if and when they are broken by quantum computers.

Such a rollover of algorithms would require both senders and receivers to update their validation and signing logic in order for the system to properly transition. For the NPRM's proposed PKI, this would require *every vehicle* to update its cryptographic software. Is unlikely that this level of agility would exist, and it would be significantly complicated by manufacturers using hardware-accelerated cryptography, which will require the replacement of physical cryptographic modules in order to enable continued participation. In this scenario, if a vehicle did not receive an update and an individual manufacturer pre-rollover RSA/ECC key is cracked, an attacker would be able to create undetectable valid leaf certificates for use in attacks against these un-upgraded vehicles—and revocation would no longer be a theoretical defense. Given current quantum computing concerns, it would be imprudent to roll out a PKI with such a long expected lifetime—especially without any plan or possibility for algorithmic agility.

B. The Proposed PKI Fundamentally Misunderstands the Entire Purpose of PKI Systems.

Existing PKI systems are used for a very specific function—*i.e.*, to associate some form of identity (such as, in WebPKI, a domain name) with a cryptographic key so that a third party can verify that a message claiming to be created by that identity was in fact created by that identity. In other words, its purpose is to authenticate the message as being from that sender. This is its only purpose. PKI cannot and does not guarantee anything about the contents of the message; it only verifies that the message was generated by the identity associated with its key. It does not, for instance, provide any guarantee that the contents of the message are in any way "safe"—*i.e.*, truthful and thus the sound basis for decisions.

Instead of implementing PKI for its true conceptual purpose—associating a key with an identity—the NPRM envisions using PKI as a way to establish that the contents of a message can be "trusted." In this way, the NPRM seems to imply that the message is factually true if the sender is verified as the true sender. This relies on the incorrect conclusion that if a third-party can prove that a chain of trust exists between a leaf certificate and a root authority, then any messages signed by the key can be "trusted." This is a dangerous misconception when the contents of the messages are information about vehicle movements destined for use in safety-critical applications.

The NPRM's proposal is dangerous in this way: it will create false assumptions about the level of trust a system can assign messages it receives based on the mere fact that it can verify

---

[18] National Security Agency/Central Security Service, Information Assurance Directive, Commercial National Security Algorithm Suite and Quantum Computing FAQ (Jan. 2016), https://www.iad.gov/iad/customcf/openAttachment.cfm?FilePath=/iad/library/ia-guidance/ia-solutions-for-classified/algorithm-guidance/assets/public/upload/CNSA-Suite-and-Quantum-Computing-FAQ.pdf&WpKes=aF6woL7fQp3dJiH77RNPPc4f28rNgYFpYQcHrp.

that the messages were properly signed. The contents of a message are not necessarily "safe" just because a message is signed. For instance, even if private keys were protected in such a way that an attacker with physical access to a vehicle's systems could not extract the keys required to sign a malicious message, they could simply inject their own messages into the existing system—either by communicating directly with the cryptographic module or by injecting false sensor data into the system that constructs BSMs to be signed. This is only one example.

The solution that seems to be suggested for this problem—the maintenance of a centralized authority that tracks malfunctioning or malicious systems using valid keys and the publication of Certificate Revocation Lists (CRLs) covering these bad actors—is infeasible when considered inside of the framework of the proposed PKI. It is infeasible for two reasons:

First, individual vehicles that may be targets of attacks are themselves responsible for detecting and reporting these malicious actions. Where a bad actor is attempting to influence the behavior of a vehicle's internal systems based on malicious BSMs, it is very likely that these messages would be formed in such a way that the car's own anomaly detection systems would not be triggered—and would thus consider the malicious messages completely valid. It may be possible to identify and collect these malicious messages after the fact—after a vehicle has detected that it was being told something strange—but the practical means of doing so, and the institutional responsibility for performing this kind of investigation, remain underspecified. Such an investigation is also of limited use if a successful attack has already taken place. Furthermore, the time between detection and the centralized misbehavior authority (MA) being able to generate and distribute a CRL containing the bad actor's key pair (or the full set of key pairs that a vehicle was provisioned with) is likely to be long enough that a attack could be repeated multiple times before being negated. And where an attack is detected, the attacker could simply move on to another vehicle as soon as they observed a published CRL containing their credentials.

Second, sending out CRLs to all V2V vehicles would be impracticably inefficient, and CRLs are thus not a realistic way managing revocation of trust. Given the tremendous size of the NPRM's proposed PKI, the size of CRLs alone would overwhelm the system. Assuming, based on publicly available data on vehicle recall rates,[19] that even a mere 1% of all recalled vehicles had issues with their sensors (or some other system used to generate BSMs), given the number of certificates per week that would be provisioned to each individual vehicle, the size of a CRL required to cover just these vehicles would contain—conservatively—tens of gigabytes (GB) of data, which would need to be distributed regularly to every vehicle participating in the system.

This is *orders of magnitude* larger than the largest CRLs used in the WebPKI, where distribution of CRLs is already an extremely traffic-intensive process. Even in the WebPKI, CRLs are "difficult to maintain" and widely viewed as "an inefficient method of distributing

---

[19] *See* Thomas Lee, Which Automakers Have the Best Recall Rate and Timeliness?, iSeeCars.com (2016), http://blog.iseecars.com/which-automaker-has-the-best-recall-rate-or-timeliness/ (listing recall rate by manufacturer from Jan. 1985 through Sept. 2016, as well as industry average, based on NHTSA's new vehicle sales data and recall data).

critical information in real time."[20] Requiring each vehicle to stay up to date with the proposed CRLs would require terabytes of data transfered over what are otherwise low traffic systems. The size of these CRLs would cause serious implementation and performance issues for the systems required to verify BSMs, even with the relatively efficient data structures web browser vendors have developed with their experience using CRLs in the WebPKI.

III. The V2V Proposal Opens Up Cars to an Entirely New Surface of Attack But Fails to Address This Serious Security Concerns, Putting People at Risk of Potentially Grave Harm.

Vehicle information and communications security is a challenging problem. Even sophisticated parties often commonly fail to mitigate vulnerabilities and attacks, and vehicle makers are often unfortunately relatively unsophisticated in comparison to the IT industry in this context. Security researchers in the field have demonstrated successful attacks—even remotely—on existing vehicle systems.[21] These attacks establish that the computerization of vehicles has already rendered them significantly vulnerable in a range of threat scenarios—and that manufacturers need to pay more attention to security.

As one commenter, Alex Kreilein, a former Department of Homeland Security lead cybersecurity strategist and the cofounder and managing partner of SecureSet, has already noted, "the addition of DSRC exposes a *new, additional* attack surface to vehicles which may already be vulnerable through different means."[22] EFF agrees with Kreilein's concerns regarding V2V security and the shortfalls of the NPRM's proposal. If V2V were to be deployed, it is imperative that it be deployed—and used—with extreme caution. The current proposal fails to lay out either a security framework or a compliance regime, putting the safety and lives of individuals at risk. Indeed, while some attackers may aim to merely block traffic, frighten motorists, or damage a competing manufacturers' business or reputation, other may seek to cause catastrophic car crashes.

---

[20] Margaret Rouse, Certificate Revocation Lists (CRLs), TechTarget (last updated May 2016), http://searchsecurity.techtarget.com/definition/Certificate-Revocation-List; *see also* Pawel Szalachowski, Laurent Chuat, & Adrian Perrig, PKI safety net (PKISN): Addressing the too-big-to-be-revoked problem of the TLS ecosystem (Feb. 2016) (Noting that one of the drawbacks of the CRL approach is that "[i]t is inefficient, since the entire CRL must be downloaded to verify a single certificate"), available at http://www.netsec.ethz.ch/publications/papers/PKISN2016.pdf; Adam Langley, Revocation Still Doesn't Work, ImperialViolet (Apr. 29, 2014), https://www.imperialviolet.org/2014/04/29/revocationagain.html.

[21] Charlie Miller & Chris Valasek, Remote Exploitation of an Unaltered Passenger Vehicle, IOActive (Aug. 10 2015), http://illmatics.com/Remote%20Car%20Hacking.pdf (discussing successful remote attack of an unaltered 2014 Jeep Cherokee).

[22] Alex Kreilein, Dedicated Short Range Communications (DSRC) Expose Critical Gaps in Security and Privacy, SecureSet, 2 (Mar. 29, 2017) (emphasis in original), http://glenechogroup.isebox.net/securesetaccelerator/dedicated-short-range-communications-dsrc-expose-critical-gaps-in-security-and-privacy.

EFF is concerned about the issues and scenarios laid out in Kreilein's comment—including but not limited to the lack of segmentation and system interdependencies inherent in vehicles and the potential for non-safety commercial uses of DSRC to compound the security concerns presented by V2V by enabling communication with a wider range of unknown third parties. In addition:

- As noted, certificate revocation is not an adequate solution for *preventing* spoofed location announcements. The first sign of a spoofed announcement may be a real-world attack, which may or may not be investigated successfully after-the-fact. Attackers do not have to announce themselves or leave evidence ahead of time. And serious harm could happen on the first instance. Furthermore, sophisticated actors could simply obtain new credentials following revocation.

- The work of the information security community shows that it is extraordinarily difficult to prevent attackers from modifying devices, or from extracting secrets from devices, that are under their physical control. Effort must be expended to put in place protections for situations in which attackers have gained physical control of a vehicle.

- Applications that parse and/or act on incoming V2V transmissions may be a vector for attack. Given the ease of spoofing and the inadequacy of revocation to address it, vehicles should *never* act autonomously solely in response to other vehicles' self-reported position or velocity. Autonomous evasive action based on incoming reports that are not corroborated by sensor data is an unacceptable risk. NHTSA must adopt a security framework that acknowledges and addresses the fact that spoofing will never be impossible or implausible—and that does not ever automatically rely on other vehicles' reports as trustworthy absent corroboration.

These issues, along with each of the vulnerabilities and scenarios addressed in Kreilein's comment, must be addressed before NHTSA moves further with any V2V proposal. NHTSA must also expressly caution implementers about the limitations of the trustworthiness of V2V transmissions—and about the need to improve security in computer systems built into vehicles.

IV. The NPRM's Proposal is Inefficient and Does Not Make Sense From a Cost-Benefit Perspective.

Finally, the V2V proposal is inefficient from a common sense, cost-benefit perspective. As more explained in detail in the comments submitted by Brad Templeton, a developer of and commentator on self-driving cars, software architect, and Internet entrepreneur, it will take a great deal of time and resources before there is any "payoff" in terms of increased safety. By that time, give the exponential rate of technological development in mobile data networks alone, it's likely the technology will become obsolete. EFF agrees with the concerns and objections raised in Templeton's submission.

V2V comes with a significant network effect—*i.e.*, the value of the technology is highly dependent on the number of others using it. The technology is of no use if none of the surrounding vehicles have it. AT&T's Picturephone failed in the 1970s due to a similar network effect: the phone was only useful if both the caller and the recipient possessed the device; with

only a few hundred Picturephones in the wild, the ability to use the device was extremely limited.[23]

In the case of V2V, even if millions of cars in the United States contained interoperable V2V devices, the chance that any two cars on the road would both have the technology and thus be able to communicate to prevent an accident would be incredibly low. As the NPRM itself acknowledges, it would take *decades* before a significant percentage of vehicles on the road were equipped with V2V, not even taking into account that some of these devices will go unrepaired or un-updated. *See* RT 82 FR 3989-3990.

Meanwhile, the cost of implementing V2V would be great. NHTSA estimates that the total annual cost to comply with its proposed V2V mandate would range from $2.2 billion on the low end to $5.0 billion on the high end, corresponding to a cost per new vehicle of roughly $135 to $301. 82 FR 398X. Not taking into account opportunity costs, NHTSA estimates that under its proposed rule—assuming that a final rule was issued in 2019, the phase-in period began in 2021, and compliance was required by 2023—the breakeven year would be between 2029, at its most liberal estimate, and 2036, at its most conservative estimate—*i.e.*, over 15 years after the final rule was issued. 82 FR 3859. The investment NHTSA is proposing is no small matter: at an annual cost of between $2.2 billion and $5.0 billion, the proposal would cost $33 to $75 billion over the 15-year-period during which the monetary investment in the technology surpassed any gains. In other words, we would see no benefit from this technology until after $33 to $75 billion was already spent. And these estimates do not even take into account the costs associated with the privacy and security risks introduced by the technology, which will include accounting for security breaches and identify theft.

The NPRM suggests the investment is worth it, estimating that by Model Year (MY) 2050, V2V would prevent 261,241 to 453,138 crashes, save 587 to 1,006 lives, reduce 181,408 to 307,409 injuries, and eliminate up to 549,803 PDOVs (property damage only vehicle crashes). 82 FR 3994.

But the NPRM has not adequately considered whether, by that time, V2V will be outpaced by other forms of technology. NHTSA has dismissed concerns regarding the potential obsolescence of V2V, stating that it views other technologies "as complementary and not competing." 82 FR 3866. With all due respect, technologies can be technically complementary yet compete for real-world dollars. The billions of dollars spent on V2V during the 15-year ramp-up to societal benefit is money that cannot be spent on other technologies—which translates into a substantial opportunity cost.

That money could be spent on one of the V2V alternatives—such as use of 5G cellular networks, either through vehicle-to-cloud or phone-to-phone communications. Given the rapid pace of innovation in the wireless device space, where people upgrade smartphones far more frequently than they replace cars, and given the extra costs associated with upgrading the

---

[23] Wikipedia, Videotelephony last modified Apr. 7, 2017), https://en.wikipedia.org/wiki/Videotelephony.

information systems in cars, NHTSA's proposed system is likely to fall behind long before it can be useful.

Currently, the one true "benefit" of V2V—as compared to *current* alternatives—is in situations that require low latency. But given the fact that V2V messages should never be trusted absent corroboration, for the reasons discussed previously, this benefit is as a practical matter limited. It may be useful in very specific situations—such as via platoon where trusted vehicles can follow behind one another with very short headways (as road cyclist are known to do) to decrease drag and increase fuel efficiency—but NHTSA should not mandate a technology with known limited utility. This is especially true where both the cost and opportunity cost of implementation is so substantial. And given the pace of current technological development, other technologies may outpace V2V as far as latency is concerned as well, negating the benefit all together before it was even of any use.

V. Conclusion

It is concerning how far the NPRM's proposal comes from actually protecting the privacy and security of vehicles that would be equipped with V2V—and thus the privacy and security of the public. Until these serious privacy and security concerns are adequately resolved, NHTSA should not recommend—let alone mandate—this technology. Furthermore, given the likelihood that any potential benefit will be achievable via other means in the near future, as V2V is outpaced by other communications technology, NHTSA should decline to mandate V2V even if the privacy and security concerns were somehow adequately resolved. Should NHTSA proceed with any V2V proposal—either now or in the future after the privacy and security concerns are addressed—consumers must be provided with the option and right to disable V2V in order to protect their privacy, security, and safety.

Respectfully submitted,

Lee Tien
Jamie Williams
Seth Schoen
Roland Shoemaker
Electronic Frontier Foundation
815 Eddy Street
San Francisco, CA 94109
Telephone: (415) 436-9333
jamie@eff.org