

Framework for Security and Privacy in Automotive Telematics

Sastry Duri, Marco Gruteser, Xuan Liu,
Paul Moskowitz, Ronald Perez, Moninder Singh, Jung-Mu Tang

IBM Thomas J. Watson Research Center

19 Skyline Drive

Hawthorne, New York 10532

ABSTRACT

Automotive telematics may be defined as the information-intensive applications that are being enabled for vehicles by a combination of telecommunications and computing technology. Telematics by its nature requires the capture of sensor data, storage and exchange of data to obtain remote services. In order for automotive telematics to grow to its full potential, telematics data must be protected. Data protection must include privacy and security for end-users, service providers and application providers. In this paper, we propose a new framework for data protection that is built on the foundation of privacy and security technologies. The privacy technology enables users and service providers to define flexible data model and policy models. The security technology provides traditional capabilities such as encryption, authentication, non-repudiation. In addition, it provides secure environments for protected execution, which is essential to limiting data access to specific purposes.

Categories and Subject Descriptors

D.4.6 [Operating Systems]: Security and Protection – Access controls, Information flow controls

General Terms

Security.

Keywords

Automotive Telematics, Privacy, Privacy Policies, Security

1. INTRODUCTION

Automotive telematics may be defined as the information-intensive applications that are being enabled for vehicles by a combination of telecommunications and computing technology. The automobile is, in effect, a computing platform to which

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WMC '02, September 28, 2002, Atlanta, Georgia.

Copyright 2002 ACM 1-58113-000-0/00/0000...\$5.00.

mobile commerce services may be delivered. The services being delivered today on a regular basis and projected for the near future include navigation information, emergency roadside assistance, location-based services, delivery of digital information such as e-mail, entertainment, diagnostics and prognostics, and pay-for-use rental and insurance. These applications are enabled by the collection and use of data which may include information on the location of a vehicle as a function of time, emergency situations including accidents and personal health emergencies, diagnostic data on the many systems within the vehicle, services and entertainment that are selected by the vehicle occupants, the demographics of the driver and passengers, and the behavior of the vehicle driver.

We can compare the growing automotive e-commerce telematics industry with that of the Web. The growth of e-commerce on the World Wide Web has been limited by the reluctance of consumers to release personal information. In “Building Consumer Trust in Online Environments” [1] the authors find that “Fully 94 percent of Web users have declined to provide personal information to Web sites at one time or another when asked and 40 percent who have provided demographic data have gone to the trouble of fabricating it”. If potential automotive telematics users share the concerns of Web users, then a large segment of the potential telematics market, perhaps as much as fifty percent may be lost.

There is a significant potential for the misuse of collected data. End users or consumers may substitute false data or hack into in-vehicle applications. Telematics service providers and application providers may sell consumers’ data to third parties without the permission of the consumers. Although, there are no current US regulations in place to “safeguard” the information collected, certain existing European regulations, and pending US and European statutes may soon impose strict controls on the collection, use, and storage of information about individuals. In general, telematics applications will be successful if providers know that the data that they receive is accurate and if end users know that their privacy is assured. Thus, data must be protected. Users must be assured that their privacy is respected and the security is in place to protect data from being divulged to unauthorized entities. Data protection consists of providing both privacy and security protection. Our goal is to achieve that protection while enabling the sharing of data.

Privacy protection today at a fundamental level requires a user to trust service providers to handle personal data according to stated terms. There is a certain degree of goodwill that is at stake to

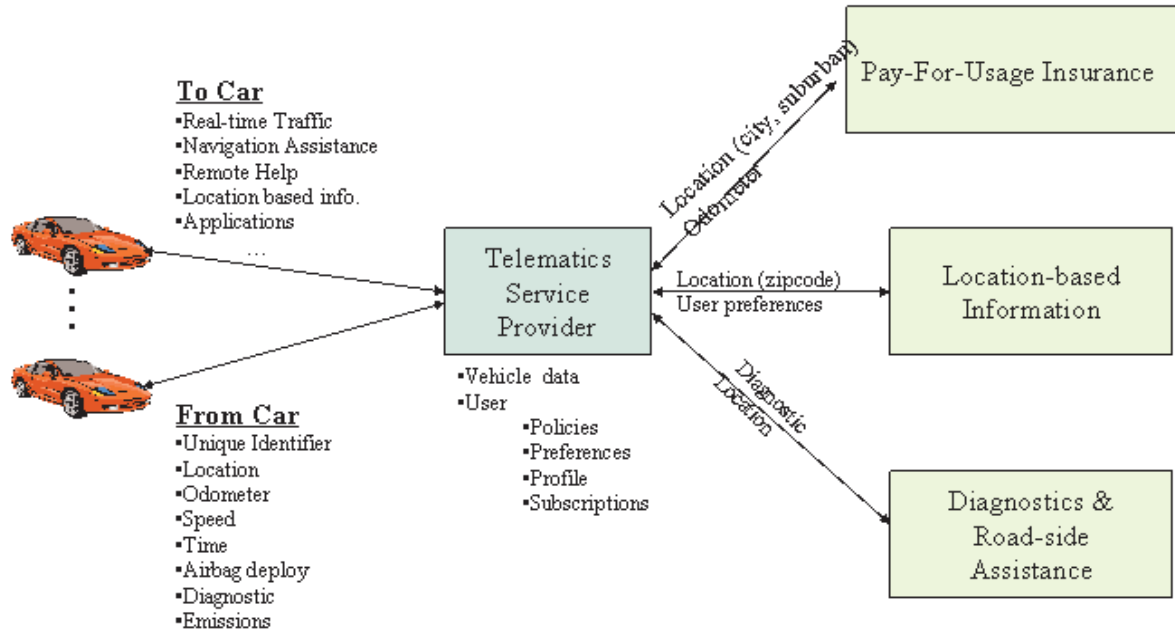


Figure 1 Automotive Telematics System Overview

prevent a service provider from using the data in an inappropriate manner. However, there are no safeguards in place to prevent inappropriate use of data; and no protection from insider abuse.

Likewise, there are no protections to assure a vehicle user that applications that are running in the car are secure. Kingpin and Mudge [2][3] analyze the susceptibility of portable devices, primarily PDAs, to attack by malicious code. They make the point that you cannot have a secure application without a secure foundation. As with PDAs, it is key to the future of automotive telematics that end users, telematics and application service providers be assured of the security of their systems from end-to-end. *Security* is a broad term encompassing many concepts and elements including confidentiality/secretcy (including privacy), integrity, and availability [15]. Security and privacy threats to systems similar to those used and being proposed for automotive telematics infrastructures have been studied for quite some time (e.g., see [16,23]). Here, our security focus will be on assuring the privacy and integrity of telematics information – user data, vehicle data, time and location information, and even executable software – that is generated or stored in, or transmitted to/from, the in-vehicle client platform during its life cycle.

In the following sections of this paper we provide a description of an automotive telematics application and a scenario, the challenges posed by automotive telematics data, and an overview of privacy technology used in the proposed framework. We then

detail proposed data protection framework, and conclude the paper with a summary of our work.

2. AUTOMOTIVE TELEMATICS APPLICATION

Figure 1 shows an overview of a typical automotive telematics application. Cars shown in the picture are equipped with a wireless communication device, variety of sensors, and a car computer that has a display, sufficient memory, storage, and processing to run complex embedded applications and middleware. The car computer interfaces to car bus and other car sensors, for example, Global Positioning System (GPS) sensor, and collects car engine performance data, safety information, and car location.

Car users subscribe to a telematics service provider (TSP) to get variety of services from application service providers (ASP) which include Pay-for-Use Insurance, Information, and Car Care and Emergency Assistance as shown in Figure 1. In order to get services from a ASP, a car user needs to send some or all the information collected by the car computer to the ASP. In the setup shown above each car transmits data as necessary to telematics service provider which then provides data to different ASPs as needed. In this case, the telematics service provider acts as a service aggregator and a data broker. In addition to the data transmitted by cars the TSP stores user preferences and user subscriptions to services.

As shown in Figure 1 different ASPs need different user data and use it for different purposes. The Pay-for-Use Insurance ASP needs user identification data, GPS data, miles driven to compute premiums and perform risk analysis. The Information ASP needs user location, and user preferences to send back information on local attractions. The data identifying user need not be sent to this service provider. The Car Care and Emergency Assistance ASP needs car engine performance and safety information on regular basis, and car location in case of emergency.

2.1 Pay for Use Insurance Scenario

The following scenario, taken from the point of view of a user, illustrates how a customer may choose among a set of privacy policies and how data may be aggregated by a telematics service provider and used to calculate the customer's bill.

2.1.1 Enrollment

Jane is a working professional who uses her automobile only to commute a short twenty miles to work and for local shopping. She uses a rental car for company business trips. Thus, she is interested in the new pay-for-use (PFU) program that is offered by her insurance company, Giant Inc. The description of the program that she received in the mail indicates that she can enroll by calling an 800 number or by using the company's web site. Jane chooses the web site.

Jane enters the URL of the site on her laptop at home and quickly sees the page for the Giant PFU program. The page explains that PFU subscribers will be charged only when they use their car. Rates will be based upon miles driven and whether the driving is done in an urban area or a suburban area such as the one in which Jane lives. The page also explains that there are several privacy policies available.

Policy 1 – This policy provides the greatest degree of personal protection. Only Jane's cumulative data, not detailed location data, will be available to the insurance company without Jane's explicit consent.

Policy 2 – This policy allows Giant full access to Jane's driving data after all personal identification information has been stripped from the data. Only summary reports of total cumulative mileage are sent to Giant with Jane's ID attached. It also allows Giant to sell anonymous data to third parties. This policy is offered at a five percent discount with respect to policy 1.

Policy 3 – This policy offers the protection of the Location Privacy Protection Act with respect to the disclosure of Jane's data to third parties. However, it allows Giant full access to Jane's driving and personal information to enable Giant to provide Jane with special offers. This policy is offered at a ten percent discount with respect to policy 1.

Policy 4 – This policy allows Giant and third parties full access to Jane's driving data and personal information. This policy is offered at a fifteen percent discount with respect to policy 1.

Jane chooses Policy 2. She does not mind having her anonymous driving data used by Giant and third parties. The enrollment web page asks Jane to enter her insurance ID number to confirm her choice. Jane installs necessary software in her car and is ready to go.

2.1.2 Driving - Data Aggregation

That evening when Jane starts her car, she is pleased to see a message appear on the navigation screen- "PFU" system now running - press # 1 for charges incurred this month". Jane presses # 1, only to see the message "Cumulative Charges for January 2003 - \$0.00". Of course, she has yet to drive any distance. She tries # 1 again after returning home. This time the screen reads "Cumulative Charges for January 2003 - \$1.00". Jane does a quick calculation; at 5 cents per mile, her yearly insurance bill for the 15,000 miles that she normally drives will be only \$750. This represents a savings of more than \$250 per year over her previous insurance rates.

As Jane drives, her data is accumulated at the CarAid center in a trusted computing system that is not directly controlled by Giant. CarAid is a telematics service provider that delivers a variety of services to Jane's vehicles: emergency assistance, navigation, concierge services. Monthly reports on total mileage for urban and suburban areas where Jane has driven are sent by CarAid to the Giant billing computer. Specific location information is divulged to Giant and third parties with personal information deleted in conformance with policy 2. The Giant billing computer calculates charges based upon cumulative mileage and sends bills to Jane. Jane is pleased to see that the charges in the bills correspond to the charges that she has been informed of by her in-car device.

3. PRIVACY

In a general sense, privacy may be defined as the ability of individuals to decide when, what, and how information about them is disclosed to others. Privacy principles [27][28]demand that systems minimize personal data collection, for example through anonymization [28]. Before personal data can be collected, consent from the data subject needs to be obtained by notifying about the nature and purpose of their data-collection and offering policy choices. Furthermore, it also requires the application of privacy preferences, either through technology, business practices, laws, or some combination thereof, in the use and further dissemination of the disclosed information

Several approaches to handling privacy preferences during personal information exchanges have been proposed in the past (the interested user is referred to Bohrer *et al.* [4][5]for a more detailed discussion of these methods). Some of these techniques, such as e-Wallet and Data Vault products and services, provide individuals with the ability to store, and sometimes share, personal information, along with tools to enable them to drag and drop their stored data onto Web forms as needed. Examples of such products include Microsoft's Internet Explorer and Passport service, Novell's digitalMe, Lumeria's SuperProfile and ZeroKnowledge's Freedom[6][7][8]. Microsoft's .NET MyServices offering [9]is an extension to its Passport service that provides individuals a repository to store their personal data, and allows them to grant permission to third party services and applications to access that data. Other approaches, such as the AT&T Privacy Minder [10], provide Web-privacy enforcing

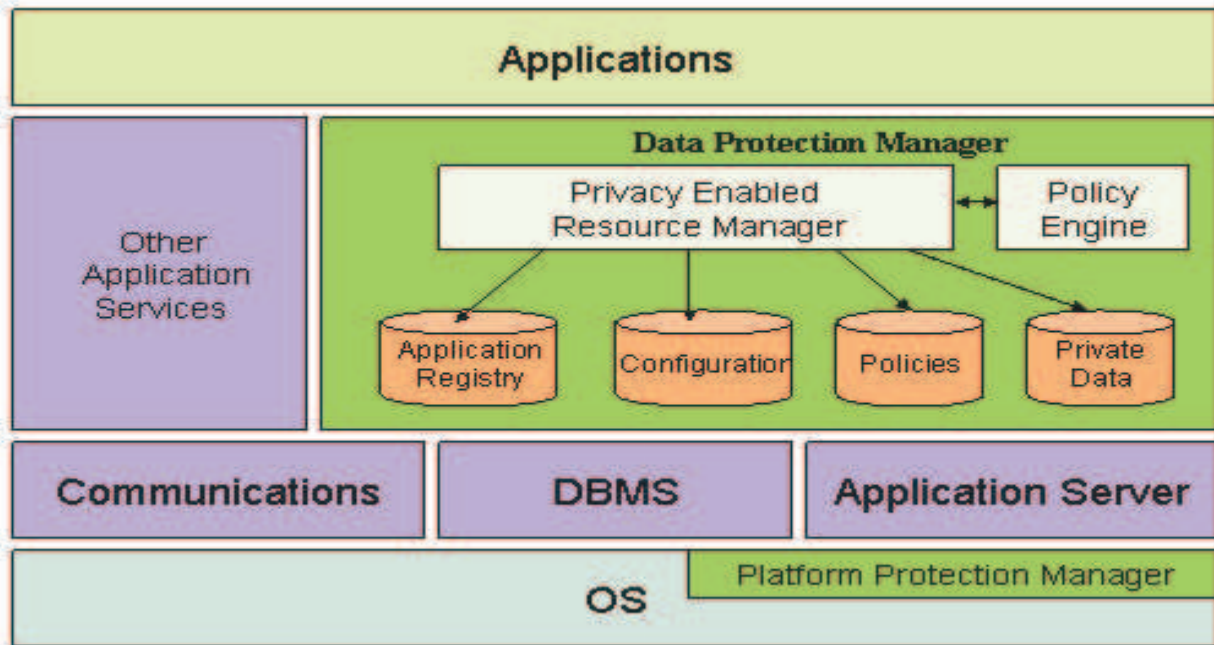


Figure 2 Generic Data Protection Platform Architecture

agents that enable individuals to formally express their privacy preferences in P3P (Platform for Privacy Preferences) [11], and automatically match them to the privacy policy of any Web sites visited by the individual. Standards have also been developed that promote the exchange of data through non-Web messaging systems. The Customer Profile Exchange Specification or CPEXchange [12] is a standard that defines how a P3P policy can be associated with personal data in an XML message. This provides a general way for an enterprise to include the privacy policy when exchanging personal data.

The IBM Privacy Services (IPS) system [4][5] provides a set of core components, based on IBM's Enterprise Privacy Architecture (EPA), [13][14] to provide individuals with greater flexibility in specifying their own privacy preferences as well as greater control over the distribution of their data. Of all the privacy-related products and services in the market, IPS is best suited for handling privacy concerns for automotive telematics applications. First, an individual can specify relatively complex privacy policies over data that is captured and stored by a smart-client within an automobile, as well as over data that is released to one or more external parties, such as service providers. Second, IPS provides the means for automatic and manual authorization for release of this data by matching the individual's privacy policies with those of data-requesters, automatic response to such requests for information, logging requests, as well as interaction with the individual to obtain manual authorization, if required.

4. CHALLENGES

There are security and privacy issues which are unique to automotive telematics. Automobiles are sensor-rich environments, thus in addition to static data such as vehicle

identification information, a significant amount of data generated in the vehicle is dynamic. There are a large and growing number of electronic control units (ECUs) which constantly monitor and adjust vehicle parameters, and the data generated by an ECU is available to external monitoring by way of the car bus. Examples of dynamic data may include parameters for emission controls, engine operation, brake application, and the speed of the vehicle. This information may be linked to position data obtained from GPS sensors and to personal information to provide a detailed picture of the operation of the vehicle and the actions of the vehicle driver. Such use of information can be desirable. For instance, General Motors' OnStar® uses the deployment of an automobile's airbags to alert a call center that emergency assistance may be needed and uses the GPS data from that vehicle to inform the call center where to send emergency assistance. On the other hand the GPS data obtained from a vehicle may be used inappropriately to track individuals as they go about their daily business.

Dynamically generating data within an automobile creates unique challenges. The sheer amount of the data generated makes it difficult, if not impossible, to store it within the automobile itself. Thus, decisions about what to store, and where, become very important. This issue is amplified by the privacy concern of data storage. More importantly, in cases where certain pieces of data are not stored within the automobile (or by a trusted third party on behalf of the individual), the retention aspect of the individual's privacy policies becomes important. Once the data is destroyed, there is no way to recover it later.

Moreover, unlike static data, which has to be collected only once by any interested party, dynamic data has to be collected repeatedly by a service provider to keep it up-to-date. Thus, there has to be a continuous transfer of dynamic data from many

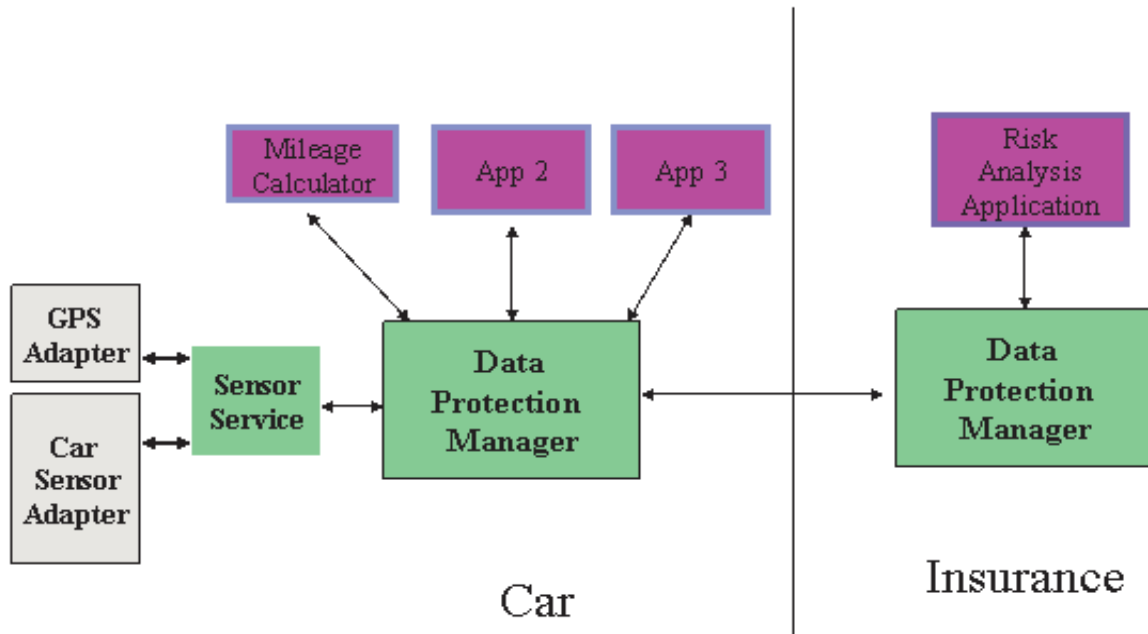


Figure 3 Example Blackboard Interactions

vehicles through the telematics service provider to application service providers. This requires an efficient and scalable evaluation of constraints in the privacy policies.

Furthermore, telematics location data is very precise. For example, location information for vehicles can be collected from a GPS receiver with 3m accuracy [29], compared to the 125m accuracy required for E-911 mobile phone services [25]. Such accuracy challenges privacy-enhancing techniques like anonymization and pseudonym switching [24]. If identifiers are removed from vehicle GPS data, an attacker ostensibly could still identify vehicles based on their overnight parking location (at least in suburban areas). If a car switches its pseudonym, an attacker can correlate new and old pseudonym based on the cars location. However, not all applications require such accuracy, which motivates flexible data aggregation mechanisms.

In-vehicle applications, providing services to the vehicle occupant(s) on behalf of the telematics service providers, may need access to data from particular vehicle sensors (e.g., GPS coordinates), and/or actuators (e.g., navigation display unit). However, direct access by applications to sensors and actuators is undesirable, for safety and liability reasons as well as for security and privacy reasons. Therefore, the challenge for the data protection framework will be to provide authenticated access to sensors and actuators in a manner that can be agreed to in advance such that each access can easily be verified and logged.

Finally, one of the most important and difficult challenges facing security and privacy in automotive telematics is *trust*. Trust must be established by both the users and service providers that the end-to-end system is doing the "right thing" at all times. This means establishing trust in the hardware and software that make up the in-vehicle client and service provider platforms

themselves. It also means providing user access to all logs and repositories concerning user data. Trust can be achieved in part by avoiding "security through obscurity", developing an architecture based on open standards and accepted practices where they exist, by insisting on openness where new innovations are necessary, and by subjecting the architecture and its components to appropriate review and security evaluations.

5. DATA PROTECTION FRAMEWORK

5.1 Approach

The primary goal of the Data Protection Framework (DPF) is to enable building telematics computing platforms that can be trusted by both users and service providers. For example, users need to trust them to protect privacy of their personal information and service providers need to trust them to protect integrity of the data. The framework employs three key concepts to build this trust. First, it uses defense-in-depth approach to build secure platform from the ground up. Second, the framework enables data aggregation close to source on the computing system trusted by the user. Third, the framework uses user defined privacy policies for obtaining user consent before data collection and usage.

5.2 Data Protection Platform Framework

Figure 2 shows the generic data protection platform architecture. This architecture can be instantiated in vehicle, in telematics service provider and in application service provider settings by choosing appropriate implementations of the two bottom layers.

In a car environment, we expect a real-time operating system such as QNX, whereas the TSP and ASP will use server operating systems such as Linux. For application server, in-car environments typically use the OSGi-based [30]platforms while server provider platforms use a typical Web Application Server. The Platform Protection Manager, which is a part of OS, monitors the integrity of all system software including the Data Protection Manager and provides security functions such as verifying signatures on applications. The Communications layer handles encrypted, authenticated, and monitored network connections. For example, it supports protocols like SSL or IPSec. The DBMS layer provides basic storage capabilities for the Data Protection Manager.

Applications follow the blackboard architectural style shown in Figure 3 for communicating with data sources, with other applications, and with external world. The Data Protection Manager provides an interface for information producers such as sensors or aggregation applications to publish data on the blackboard. Information consumers access this data through periodic queries or through a subscription/notification mechanism. We also extend the blackboard paradigm across the network. That is, applications at the TSP or ASP can submit queries to or receive notifications from the in-car blackboard mechanism.

The example illustrates how applications are composed in this framework. The GPS sensor periodically publishes location data items in the Data Protection Manager. The Classified Mileage Calculator can subscribe to the GPS data and compute with the help of a road map the total mileage driven on different types of roads. The results are again published in the Data Protection Manager. A Risk Analysis application running on the insurance server remotely subscribes to the aggregated and classified mileage data.

Blackboard-based architectures provide a simple paradigm for composing sensor-based applications. It is a common choice for building ubiquitous computing smart spaces, which depend on aggregated and interpreted sensor data. However, blackboards exhibit another key advantage for our privacy protection framework. Every data access passes through the central Data Protection Manager. This simplifies verifying that data accesses comply with the privacy policies.

5.2.1 Defense in Depth

To build a computing system that is trustworthy for both the data subject (driver) and application service providers, we take a bottom-up approach. Each layer of hardware and software provides its own security functions. This approach is often called *defense-in-depth*.

Ideally, physically and logically secure systems would be used for the in-vehicle clients as well as services and solutions providers' servers -- i.e., systems that would resist most physical and logical attacks (e.g., physical penetration, voltage or temperature attacks, power analysis, monitoring of electromagnetic emissions), and sensing and responding to all others before a system compromise (e.g., by rendering sensitive data inaccessible). However, such systems do not currently exist commercially. What does exist are *secure coprocessors*:

physically and logically secure subsystems that operate in conjunction with a local host system, employing cryptographic acceleration hardware, and providing a secure execution environment for the programs that are supposed to be run. Such secure coprocessors exist today, as exemplified by the IBM 4758 PCI Cryptographic Coprocessor [17], a product used extensively in servers for applications requiring the highest levels of assurance (e.g., banking and financial applications, electronic commerce systems). Furthermore, the near term future promises similar devices for mobile and client platforms, at prices commensurate with such client devices, and offering performance capabilities surpassing the current generation of server-oriented secure coprocessors [18]. Pervasive low-end secure coprocessors (e.g., smart cards, secure tokens), used for key storage and user authentication, are also currently available and may provide limited security assurances in lieu of more comprehensive and capable devices.

Both client and server platforms should allow for secure configuration, update, and execution (booting) of system and application software. Typically, this functionality must exist primarily in the firmware/software that is initially executed upon power-on (e.g., BIOS or system boot firmware). This power-on software layer is often in read-only memory, it should have minimal complexity and size, and should be able to [cryptographically] authenticate/verify a minimal set of commands and data that enable the configuration and update of the subsequent software layer (e.g., the system software or operating system layer). Once the platform system software has been securely configured/updated, the power-on software layer is responsible for authenticating the system software before each execution/instantiation. This is often called *secure boot* [19].

The operating system, like the power-on software and physical platform before, must also provide certain security features, such as access control, in order to support overall system and data protection. There is a great deal of ongoing work in the area of secure operating systems (e.g., see [20]and [21]). Elements of the application and application support layer may be highly integrated with the operating system. Together, these layers may provide support for cryptographic programming libraries, secure communication protocols, encrypted file systems or databases, firewall and intrusion detection capabilities, and even virtual machine application authentication, and execution. Further, because application isolation is important when applications potentially come from competing or otherwise mutually hostile parties, the application support layer itself may provide virtual environments/machines for the purpose of protecting these applications from interfering with each other or the operating system.

As alluded to above, the same hardware and software layers, and their respective security features, described for the in-vehicle client platform are also required for the various telematics service provider servers. We are able to assure end-to-end and life cycle protection of relevant data only when the same level of security is employed across the entire system.

5.2.2 Data Aggregation Close to Source

User trust can be further enhanced by minimizing the amount of private data that leaves the computing system trusted by the user.

To this end service providers who need access to private data deploy data aggregation applications inside the computing system. Only the aggregated results can be sent back to the service provider. However, it difficult to ensure that the aggregation applications do not misbehave i.e., leak private data, or carryout denial of service attacks.

We use the following mechanisms to monitor and control application behavior. First, the computing system verifies that each deployed application has proper credentials. Second, it places each application in a sandbox to protect itself and other applications. This sandbox allows us to define individual application access privileges for system resources such as files, sockets. It also prevents direct communication between different applications. In order to prevent any malicious transmission of private data all application modules are denied network privileges. All local and network communication is through framework data protection manager which checks privacy policies and generates an audit trail for later verification.

5.2.3 User Privacy Policies

Privacy principles require notifying users and obtaining consent before data collection. User-defined policies specifying personal data handling preferences, and solution provider policies attesting to user data handling practices will together form virtual contracts between users and solutions providers. The framework will enable enforcement of these policies by classifying data and defining data handling rules according to classifications and policies, and by assuring application/solution compliance to the rules. Enforcement of policies and compliance assurance will extend from the in-vehicle client to the solution or service provider back-end systems, and can be extended to third-party interactions within the domain of the framework.

Internally, the Privacy Enabled Resource Manager (PERM) component, shown in Figure 2 handles requests for private data. A typical request for data includes application credentials, privacy policy concerning data, and description of data items. The PERM first verifies application credentials. Upon successful verification of credentials the PERM compares application privacy policy with the user's privacy policy to determine whether to grant access or not. For more details please refer to [4][5].

6. SUMMARY

In this paper, we have dealt with the protection of private data in the automotive telematics domain. As we have stated, our goal is to enable the controlled sharing of private data according to policies agreed to by the owner of the data. Further, we would like to assure services providers that the data is not tampered with at its point of origin or anywhere in the processing chain. We have outlined the various challenges to protecting automotive telematics data, and have presented a framework to address these challenges. Next, we intend to implement the proposed data protection framework within an end-to-end solution in order to enable real world applications.

7. ACKNOWLEDGMENTS

The authors thank their colleagues: Charles Tressor for challenging us to explore issues of privacy and security for

automotive telematics; George Salmi and Barbara Churchill for providing us with their automotive telematics domain expertise; Paul B. Chou for his support and encouragement.

8. REFERENCES

- [1] Hoffman, D.L., Novak, T.P., and Peralta, M.A. "Building Consumer Trust Online," Communications of the ACM, Volume 42 (4), 80-85, April, 1999.
- [2] Kingpin and Mudge, "Analysis of Potable Devices and Their Weaknesses Against Malicious Code Threats", RSA Conference, San Francisco, CA, April 11, 2001.
- [3] Kingpin and Mudge, "Security Analysis of the Palm Operating System and its Weaknesses Against Malicious Code Threats", Proceedings of the 10th USENIX Security Symposium, Washington, DC, August 13-17, 2001.
- [4] Bohrer, K., Liu, X., Kesdogan, D., Schonberg, E., Singh, M. and Spraragen, S.L. "Personal Information Management and Distribution", Proceedings of the 4th International Conference on Electronic Commerce Research, Dallas, TX, November 8-11, 2001.
- [5] Bohrer, K., Kesdogan, D., Liu, X., Podlaseck, M., Schonberg, E., Singh, M. and Spraragen, S.L. "How to Go Shopping On the World Wide Web Without Having Your Privacy Violated", Proceedings of the 4th International Conference on Electronic Commerce Research, Dallas, TX, November 8-11, 2001.
- [6] Novell Inc, "digitalMe: Making life easier on the net," <http://www.digitalme.com>
- [7] Lumeria, "An Informediary Approach to Privacy Problem," <http://www.lumeria.com/whitepaper.shtml>
- [8] Zero-Knowledge-Systems, Inc. "The Freedom Network Architecture," <http://www.freedom.net/>
- [9] Microsoft Inc., "A platform for user-centric application," <http://www.microsoft.com/myservices/>
- [10] AT&T, "Privacy Minder," <http://www.research.att.com/projects/p3p/pm>
- [11] "The Platform for Privacy Preferences 1.0 (P3P1.0) Specification". April, 2002 <http://www.w3.org/TR/P3P>
- [12] CPExchange, "Global standards for privacy-enabled customer data exchange," <http://www.cpexchange.org/standard/>
- [13] IBM, "Enterprise Privacy Architecture (EPA)," <http://www.ibm.com/services/security/epa.html>
- [14] Karjoth, G., Schunter, M., and Waidner, M., "Platform for Enterprise Privacy Practices: Privacy-enabled Management of Customer Data", Proceedings of the 2nd Workshop on Privacy Enhancing Technologies, 2002.
- [15] Russell, D., and Gangemi Sr., G.T. "Computer Security Basics", O'Reilly & Associates, Inc. 1991.
- [16] Karger, P., Yair, F., "Security and Privacy Threats to ITS", The Second World Congress on Intelligent Transport Systems, pp. 2452-2458, November 1995

- [17] Smith, S.W., and Weingart, S.H. (April 1999) "Building a High-Performance, Programmable Secure Coprocessor", *Computer Networks (Special Issue on Computer Network Security)*, 31: 831-860
- [18] Dyer, J., Perez, R., Sailer, R., Van Doorn, L., "Personal Firewalls and Intrusion Detection Systems", 2nd Australian Information Warfare and Security Conference 2001, November 2001.
- [19] Arbaugh, W.A., Farber, D.J., and Smith, J.M. "A Secure and Reliable Bootstrap Architecture." 1997.
- [20] Security-Enhanced Linux, <http://www.nsa.gov/selinux>
- [21] Bastille Linux, <http://www.bastille-linux.org>
- [22] Langheinrich, M., "Privacy by Design -- Principles of Privacy-Aware Ubiquitous Systems," *ACM UbiComp*, 2001.
- [23] Agre, P., "Looking Down the Road: Transport Informatics and the New Landscape of Privacy Issues", *CPSR Newsletter* 13(3). 1995.
- [24] Samfat, D., Molva, R., Asokan, N., "Untraceability in Mobile Networks", *Mobicom*. 1995.
- [25] Reed, J., Krizman, K., Woerner, B., Rappaport, T., "An Overview of the Challenges and Progress in Meeting the E-911 Requirement for Location Service," *IEEE Communications Magazine*, 30-37, April 1998.
- [26] Covington, M.J., Long, W., Srinivasan, S., Dey, A.K., Ahamad, M., and Abowd, G.D.. Securing context-aware applications using environment roles. In 6th ACM Symposium on Access Control Models and Technologies (SACMAT 2001). 2001.
- [27] Simone Fischer-Hubner, editor. *IT-Security and Privacy - Design and Use of Privacy-Enhancing Security Mechanisms*. LNCS. Springer. 2001.
- [28] Pfitzmann, A., and Koehntopp, M.. Anonymity, unobservability, and pseudonymity - a proposal for terminology. In *Workshop on Design Issues in Anonymity and Unobservability*. 2000.
- [29] Blacksher, S., Foley, T. Boulder HOPs Aboard GPS Tracking. In *GPS World*, January 01, 2002
- [30] Open Services Gateway Initiative, <http://www.osgi.org/>