Driving Privacy Forward: Homomorphic Encryption and Oblivious AI in Smart Mobility

Final Report to the Office of the Privacy Commissioner of Canada (OPC)
Contributions Program

Prepared by: Mitra Mirhassani

Project Title:

Driving Privacy Forward: Homomorphic Encryption and Oblivious AI in Smart Mobility

- **Reporting period:** June 1, 2025 – March 15, 2026
- **Recipient:** Office of the Privacy Commissioner of Canada (OPC), Contributions Program
- **Lead organization:** SHIELD Automotive Cybersecurity Centre of Excellence, University of Windsor
- **Principal Investigator / Project Administrator:** Mitra Mirhassani
- **Key partners:** HLP Innovative Solutions; consultation with Transport Canada (as applicable)
- Version control: v1.0 | March 13, 2026 | Prepared by Mitra Mirhassani
- **Acknowledgement + disclaimer**
- This project was supported by funding from the Office of the Privacy Commissioner of Canada (OPC) through the Contributions Program. The views expressed in this report are those of the authors and do not necessarily reflect those of the OPC.

Executive Summary

Connected vehicles are becoming a default part of mobility, and with that comes a privacy reality that most people do not see: sensitive data can be protected "in transit" and "at rest," yet still become exposed at the moment it is processed. In many current workflows, cloud analytics and AI services still rely on a decrypt-to-compute model, meaning raw or linkable telemetry can be visible within service boundaries even when transmission channels are secure.

For mobility data, especially location-linked patterns and behavior traces, this creates a trust gap that is difficult to repair after the fact.

This project, funded through the Office of the Privacy Commissioner of Canada (OPC) Contributions Program, was designed to address that trust gap with a practical privacy-by-design pathway for smart mobility. The work combined (1) technical development in privacy-enhancing computation and (2) REB-approved stakeholder engagement to capture how privacy risks are perceived across the ecosystem and what conditions would make privacy-preserving approaches adoptable in practice.

On the technical side, the project developed and evaluated an architecture that supports analytics on protected data (a privacy-preserving workflow where sensitive vehicle data does not need to be revealed to the compute environment to generate useful outcomes). The purpose of the technical component was not to propose privacy in principle, but to demonstrate that privacy-preserving computation can be engineered, tested, and assessed with measurable performance considerations.

On the stakeholder side, we conducted semi-structured interviews across two phases: pre-demonstration baseline and post-demonstration follow-up, with a small set of informed stakeholders. These conversations were used to (i) clarify what types of connected-vehicle data are seen as most sensitive, (ii) identify where trust breaks down in current data-sharing and processing practices, and (iii) test whether privacy-preserving technical pathways improve confidence, while also surfacing practical barriers and governance expectations.

This final report documents the work completed under the contribution agreement and provides: the project objectives, key stakeholders and roles, and a consolidated set of insights intended to be actionable for privacy-by-design implementation in Canada's connected-vehicle ecosystem.

To protect participant confidentiality and maintain privacy-respectful reporting, the report does not include full interview transcripts and does not append the full REB application; it provides REB approval information as appropriate, and reports stakeholder results as themes and de-identified excerpts where needed. Simple public-facing infographics produced during the project are included as knowledge-translation evidence in the appendix.

Deliverables at a glance: This report delivers the core outputs promised to the OPC, including privacy risk mapping, a privacy-preserving technical architecture, stakeholder-informed post-

technology assessment, best-practice guidance, policy and law recommendations, and public-facing knowledge-translation outputs. Evidence of completion is provided through the technical validation in Section 7, the stakeholder findings in Section 8, the deliverables checklist in Section 11, and the appendix materials identified at the end of the report.

Objectives

This project pursued the following objectives:

Objective 1 - Privacy risk mapping and ecosystem context

Identify and analyze privacy risks in connected-vehicle data collection, sharing, and processing workflows, and clarify which risks can and cannot be reduced through privacy-preserving computation.

Objective 2 - Encrypted analytics framework (homomorphic encryption)

Develop and validate a technical framework that enables useful analytics while reducing the need for raw data exposure during processing.

Objective 3 - Oblivious analytics and privacy-preserving AI integration

Demonstrate how advanced analytics can be supported under a privacy-by-design approach, with attention to practical constraints and engineering trade-offs.

Objective 4 - Stakeholder demonstration and post-technology assessment

Conduct a demonstration-informed assessment with stakeholders to evaluate perceived value, feasibility, adoption barriers, and trust considerations.

Objective 5 - Best practices and knowledge translation

Produce best-practice guidance and public-facing knowledge-translation outputs to support awareness and real-world uptake (including simple infographics).

Project Delivery partners

- **SHIELD Automotive Cybersecurity Centre of Excellence (University of Windsor)**
  Project leadership, technical design and evaluation, reporting, and coordination.
- **HLP Innovative Solutions**
  Support for qualitative research design and analysis under REB-approved processes.

Stakeholders

A small set of informed stakeholders were engaged through semi-structured interviews across two phases (baseline and post-demonstration). Participants represented perspectives relevant to connected-vehicle data governance and adoption (e.g., mobility ecosystem operators, technology and service workflows, governance/privacy perspectives).

- **Number of stakeholder interviewees:** 6 participants
- **Total interview duration:** ~2 hours
- **Format:** Semi-structured, two-phase (pre-demo and post-demo), de-identified handling per protocol.

Insights

This section summarizes the high-level findings from stakeholder engagement, useful for decision-makers and implementers. It is intentionally written as "what matters" rather than raw transcript reporting.

1. **Over-collection and limited public awareness remain a central concern:** Stakeholders repeatedly pointed to the scale of data collection in connected mobility and the reality that users often do not understand what is being collected, where it goes, or how it may be reused.
2. **High-sensitivity data is often location-linked and identity-linkable:** Even when individual fields appear harmless, combined mobility data can become identifying through patterns. Location-linked telemetry was consistently treated as a high sensitivity category.
3. **Privacy risk is ecosystem-wide, not confined to one actor:** Stakeholders described risk across the full chain, OEMs and suppliers, service platforms, dealers, insurers, cloud providers, apps, and third-party integrations, where accountability can become fragmented.
4. **Governance and regulatory clarity are as important as technical capability:** A recurring message was that adoption depends on clear rules, enforceable accountability, and practical guidance, not only "good technology."
5. **Strong support for privacy-preserving technology, with one clear condition: minimize data before you encrypt it:** Stakeholders were receptive to privacy-preserving computation, but stressed that "encrypt everything" is not a substitute for data minimization and purpose limitation. Privacy-by-design still starts with collecting less.
6. **Adoption depends on implementation realism and translation:** Stakeholders highlighted that success requires pilots, validation, performance transparency, operational workflows, and plain-language messaging that builds trust with non-technical audiences.

Implication for the report: these insights shape (i) how we frame privacy-by-design guidance and (ii) what we emphasize in the policy and best-practice recommendations later in the report.

List of Figures

List of Tables

Abbreviations

- **AI**: Artificial Intelligence
- **APMA:** Automotive Parts Manufacturers' Association
- **CAV**: Connected and Autonomous Vehicle
- **CUTRIC:** Canadian Urban Transit Research & Innovation Consortium
- **HE**: Homomorphic Encryption
- **KT:** Knowledge Translation
- **OPC:** Office of the Privacy Commissioner of Canada
- **PETs:** Privacy-Enhancing Technologies
- **PII**: Personally Identifiable Information (or personal information, depending on usage)
- **REB:** Research Ethics Board
- **UBI**: Usage-Based Insurance
- **V2X:** Vehicle-to-Everything

Glossary

- **Privacy by design**: Building privacy protections into the system architecture and workflow from the start, rather than relying only on policy or post-hoc controls.
- **Encrypted analytics:** Computing useful results while reducing or avoiding exposure of raw data during processing.
- **Trust boundary**: The line that separates what must be trusted (people/systems with access) from what should not need trust (external compute/services).
- **Data minimization:** Collecting and retaining only what is necessary for a defined purpose, and limiting reuse.

Table of Contents

1. Introduction

This final report summarizes the activities, outputs, and outcomes of the OPC-funded project focused on improving privacy and trust in connected-vehicle data ecosystems. The report is framed as an OPC Contributions Program final deliverable and is written to support both accountability and knowledge translation. [1]

This report is written for a mixed audience, including privacy and public interest stakeholders, policymakers, industry partners, and technical audiences. Technical details are presented at a level appropriate for a final public-facing report, with deeper technical results referenced separately where applicable.

1.1 Project Context: Connected-Vehicle Privacy + Trust Problem

Connected vehicles continuously generate and exchange data to enable functionality such as navigation, safety assistance, diagnostics, infotainment, and insurance or fleet services. While these capabilities create measurable benefits, they also create privacy risk because mobility data can become highly revealing when linked across contexts and over time. [2][6][7]

A central challenge is that the connected-vehicle ecosystem is not limited to the vehicle itself. Data may pass through multiple actors and layers, including OEM platforms, mobile apps, cloud services, dealerships, insurers, infrastructure services, and third-party vendors. As a result, even when a given data sharing activity is useful, it can raise questions about transparency, necessity, downstream use, retention, and accountability across the full lifecycle.

This project was motivated by a simple trust gap: consumers and communities often do not have clear, practical assurances about **what data is collected**, **who can access it**, **how it is protected**, and **what limits exist on use and retention**. Stakeholder interviews conducted as part of this project reinforced the view that privacy risks are driven by over-collection, limited consumer awareness, and governance fragmentation and uneven maturity across the supply chain. These realities create a barrier to trusted adoption of beneficial connected-vehicle services—particularly those that rely on continuous data sharing.

1.2 Overview of the Privacy-by-Design Approach

The project pursued a privacy-by-design approach aimed at enabling useful analytics while reducing data exposure. This framing is consistent with privacy-by-design expectations in connected-vehicle guidance and with broader privacy principles such as collection limitation, purpose specification, safeguards, openness, and accountability. [2][6][7]

The approach combines two complementary methods:

Homomorphic Encryption (HE). HE enables certain computations to be performed directly on encrypted data. In practice, this supports workflows where sensitive inputs can remain encrypted during processing, reducing exposure risks associated with collecting and handling raw data in centralized systems.

Oblivious AI. Oblivious AI techniques are designed to limit what an analyst or service provider can infer about the underlying data beyond what is necessary to generate an authorized output. The intent is to reduce unnecessary visibility into individual-level details while still enabling aggregate or decision-relevant results.

Together, these methods support a model where data is minimized and protected across the lifecycle: the system is designed to restrict data collection to what is needed, keep sensitive content protected during processing, and reveal only permitted outputs to authorized parties. Stakeholder feedback collected after the demonstration emphasized that privacy protection must go beyond encryption alone and should include data minimization, clear separation of identifiable vs non-identifiable information, and transparent trust boundaries (i.e., clarity on who can decrypt what, and under what conditions). These considerations are reflected in the project's recommendations and next steps.

2. Project Governance and Partnerships

This project was governed through a lightweight, accountability-focused structure designed to (i) maintain alignment with OPC objectives, (ii) coordinate technical and knowledge-translation activities across partners, and (iii) ensure ethics and privacy considerations were addressed throughout the project lifecycle. Governance emphasized clear ownership of deliverables, documented decision points (e.g., demo readiness, public-facing messaging, interview protocols), and partner engagement to support relevance and uptake.

2.1 Team Roles and Partner Contributions

Project Lead / PI (University of Windsor – SHIELD Centre of Excellence)

- Overall project accountability, delivery oversight, and liaison with OPC.
- Technical leadership for the privacy-by-design approach (Homomorphic Encryption + Oblivious AI) and integration into the demonstration workflow.
- Oversight of ethics compliance, data management practices, and secure handling of de-identified interview transcripts.
- Direction of knowledge translation strategy (public-facing messaging, bilingual outputs, stakeholder engagement).

Highly Qualified Personnel (HQP) / Research Team

- Implementation and evaluation of the prototype/demonstration components.
- Preparation of documentation, demo materials, and performance/feasibility summaries.
- Support for stakeholder engagement logistics and synthesis of qualitative findings into report-ready themes.

HLP (Privacy / legal / governance expertise)

- Practical guidance on privacy governance considerations (accountability, consent, data lifecycle, risk framing).
- Review of communication language to ensure clarity and alignment with privacy expectations (consumer-facing trust considerations).
- Input on how technical privacy-by-design concepts map to real-world compliance and stakeholder needs.

Governance cadence and decision-making

- **Working sessions:** as needed for (i) technical progress and (ii) knowledge translation deliverables.
- **Documentation:** meeting notes and action lists maintained internally; major decisions tied to deliverable completion and OPC reporting requirements.

For reporting clarity, the six project phases map directly to the main deliverables completed under this contribution: baseline privacy risk and regulatory analysis (Phases 1 and 4), prototype design and validation (Phases 2 and 3), demonstration and stakeholder assessment (Phase 5), and final synthesis, reporting, and policy translation (Phase 6). Consultation with Transport Canada remained advisory and contextual during the project period rather than constituting a formal co-delivery role.

2.2 Project Timeline and Phases (Phase 1–6)

The project was executed in six phases to ensure ethical compliance, stakeholder relevance, and a clear progression from technical work to evidence-based impact reporting.

Phase 1, Initiation and governance setup

- Confirmed scope, objectives, deliverables, and partner roles.
- Established governance cadence and internal tracking of milestones.
- Prepared ethics materials and initiated REB process (as applicable).

Phase 2, Technical design and privacy-by-design workflow definition

- Defined the privacy-by-design workflow and trust boundaries (minimization → protection → private computation → permitted outputs).
- Identified target demonstration scenario(s) and stakeholder-relevant outputs.

Phase 3, Prototype development and validation

- Implemented and tested the demonstration pipeline.
- Generated feasibility evidence (performance/latency, practical deployment considerations, assumptions and limitations).

Phase 4, Stakeholder engagement: baseline interviews (pre-demonstration)

- Conducted baseline interviews to capture perceptions of privacy risk, data flows, and governance gaps.

- Used baseline insights to refine demo messaging and clarify adoption barriers (e.g., data minimization, accountability, ecosystem exposures).

Phase 5, Demonstration and post-demonstration interviews + knowledge translation

- Delivered the demonstration and collected post-demo feedback on perceived value, trust implications, and implementation concerns.
- Produced public-facing knowledge translation outputs (e.g., bilingual infographics; Privacy Awareness Week activities; stakeholder-facing summaries).

Phase 6, Synthesis, reporting, and next-step recommendations

- Completed thematic analysis (pre/post) and consolidated findings into OPC-ready themes and implications.
- Produced final reporting content, including lessons learned and actionable next steps (pilots, standards/policy alignment, governance recommendations).
- Identified opportunities for future publications and deployment pathways.

Milestone mapping to deliverables (optional for clarity in the final report)

- REB approval obtained → interviews scheduled and completed
- Demo readiness → stakeholder demonstration delivered
- KT outputs published → bilingual infographics and dissemination activities completed
- Interview synthesis finalized → themes and implications integrated into final report

3. Ethics and Data Governance

This project received research ethics clearance from the University of Windsor Research Ethics Board under **REB #45122** for the project titled *Driving Privacy Forward: Homomorphic Encryption and Oblivious AI in Smart Mobility*.

Ethics approval was granted on **December 02, 2025**, with an annual renewal date of **December 02, 2026**. The study was carried out in accordance with the approved protocol and in line with the Tri-Council Policy Statement: Ethical Conduct for Research Involving Humans (TCPS), relevant institutional policies, and applicable regulatory requirements.

Throughout the project, the research team understood that any changes to the protocol, consent materials, or study procedures required prior written REB approval, and that any developments affecting participant risk, study conduct, or participant welfare were to be reported promptly to the Board.

3.1 Consent Approach and Participant Protections

The study was designed as a minimal-risk engagement process involving informed, voluntary participation. The consent approach was also aligned with the broader expectation under Canadian private-sector privacy practice that consent must be understandable and meaningful, particularly where data practices are layered or not readily visible to individuals. [2][4]

Because the project involved two stages of engagement, consent was treated as an ongoing process rather than a single point of entry. Participants were informed during the first stage that they might be invited to a later technology demonstration and follow-up interview, and their willingness to continue was confirmed again at that stage. Participants were also advised that they could decline to answer any question and could withdraw from the study without penalty up to **March 1, 2026**. This deadline was clearly tied to the study's de-identification and data transfer timelines.

Participant protection was approached with care, especially given that some individuals were recruited through professional networks connected to the research team. To reduce any perception of pressure or obligation, interviews were conducted by the postdoctoral fellow rather than directly by the principal investigator. This helped create a more neutral environment for participation and supported the study's voluntary nature.

3.2 Data Handling: Storage, De-identification, Retention, and Destruction

The project's data-governance approach emphasized confidentiality, controlled access, and timely de-identification. Interviews were conducted online through Microsoft Teams and were initially recorded in audio and video form for transcription purposes. Recordings and transcripts were stored using secure University of Windsor OneDrive infrastructure, with access restricted to the principal investigator and the postdoctoral fellow.

Once transcription and review were completed, participants were assigned pseudonyms and identifying details were removed from the interview records before analysis. The materials shared for external qualitative analysis were therefore de-identified transcripts rather than identifiable participant records. These transcripts were transferred securely through the University of Windsor SharePoint to HLP Innovative Solutions, the third-party qualitative analyst engaged for the project. No direct identifiers and no participant linkage information were disclosed externally. This arrangement was further supported by a signed non-disclosure agreement that limited use of the materials to the approved purpose and required appropriate confidentiality protections.

In accordance with the approved protocol, audio and video recordings were to be permanently deleted once transcription accuracy had been confirmed and the data had been de-identified, with deletion completed no later than **March 1, 2026**. After that point, only de-identified findings and related project outputs were to be retained and securely transferred to the Office of the Privacy Commissioner of Canada. As a result, participant withdrawal was possible only until the identifiable link to the data was removed.

4. Objective 1: Privacy Risk Mapping and Regulatory Gap Analysis

Objective 1 established the baseline privacy context for the project before the technology demonstration stage. This work mapped how connected-vehicle data currently moves across the mobility ecosystem, identified the main privacy risks associated with those flows, captured stakeholder concerns through pre-technology interviews, and examined how existing privacy principles apply in this setting. Taken together, these inputs showed that connected-vehicle privacy is not only a data-security issue. It is also a governance, accountability, and trust issue shaped by the full lifecycle of data collection, transmission, processing, sharing, and reuse.

4.1 Connected-Vehicle Data Flows

The current-state mapping showed that connected vehicles operate within a distributed data ecosystem rather than as stand-alone products. Connected-vehicle systems generate and exchange location traces, speed profiles, hard-braking events, driver-monitoring signals, and other telemetry. Device-linked systems can also contribute timing, activity, and biometric information, while backend platforms aggregate these data for analytics and service delivery. In many deployments, this creates a predictable point of exposure because processing still requires plaintext visibility at some stage of the workflow.

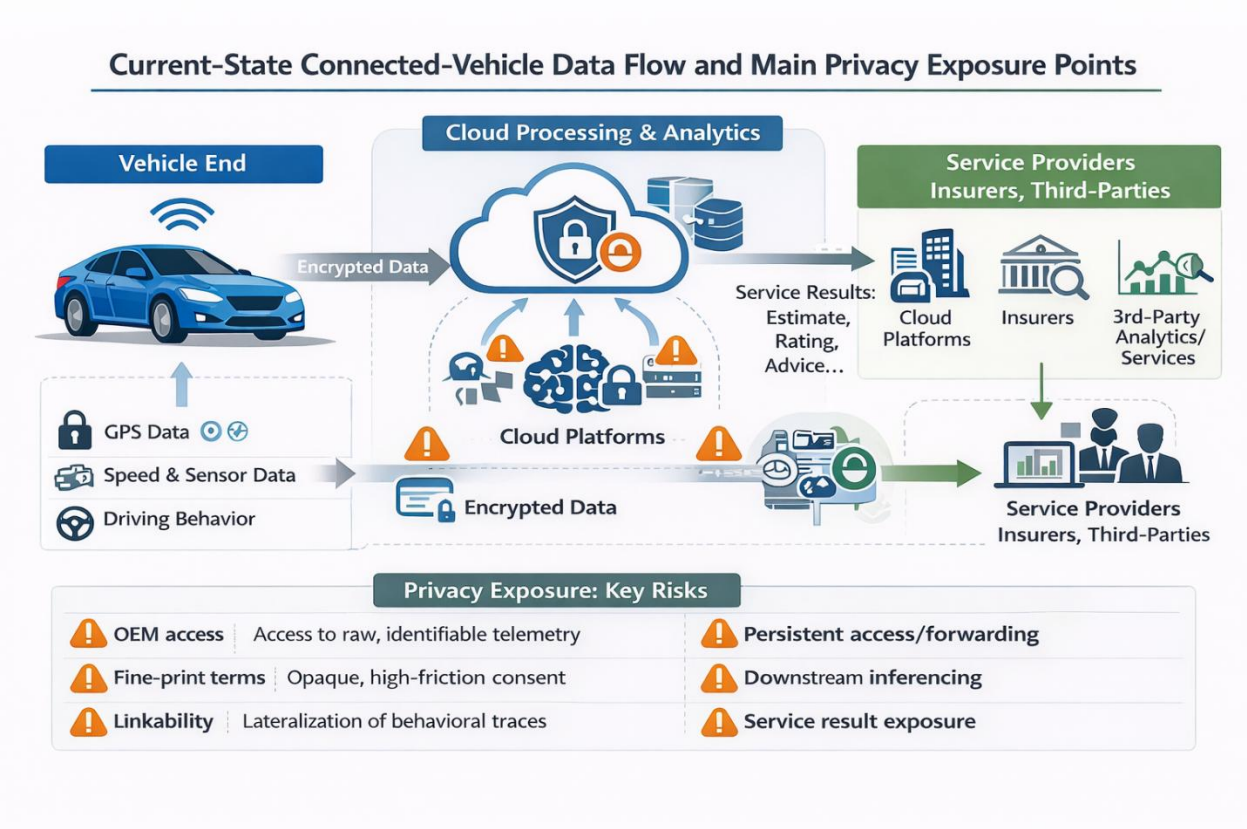

Figure 4-1 Current-state connected-vehicle data flow and main privacy exposure points.

As Figure 4-1 shows, privacy exposure increases as telemetry moves beyond the vehicle boundary and flows across cloud, service-provider, and third-party environments.

The interview findings reinforced this broader ecosystem view. Participants consistently rejected a vehicle-only framing of privacy responsibility and instead described connected mobility as a multi-actor environment involving OEMs, cloud platforms, infrastructure operators, dealerships, mobile applications, service providers, and third-party vendors. They emphasized that privacy risks do not arise only from what the vehicle collects, but from how data moves across platforms, jurisdictions, and organizational boundaries after deployment. Several participants stressed that no single actor has full visibility or control once these data flows extend across the wider ecosystem.

The mapping also highlighted the limited visibility that users typically have into these flows. The pre-technology interviews repeatedly noted that connected vehicles already collect large volumes of data, often beyond user awareness, including granular location, behavioural, audio, video, and smartphone-linked information. Participants described privacy risk as a function of scope, persistence, and inferential power rather than simply a matter of whether explicit identifiers are present.

Overall, the current-state map showed that connected-vehicle data does not remain within the vehicle itself. It circulates through a multi-actor ecosystem in which transparency, user control, and accountability are uneven, and where privacy exposure can increase as data are linked, retained, shared, or repurposed over time.

4.2 Key Privacy Risks Addressable by Encryption

The privacy risk analysis identified several recurring concerns. First, stakeholders consistently described overcollection as the dominant privacy issue. The concern was not that connected vehicles lack useful data, but that they already collect too much data, often continuously and in ways users do not fully understand. Second, participants emphasized that privacy harms extend beyond drivers to passengers, pedestrians, cyclists, and others who may be captured by vehicle and infrastructure sensing systems. Third, they highlighted indirect identifiability and inference risk, noting that mobility traces, origin-destination patterns, time-of-day information, and phone-to-vehicle integration can reveal personal patterns even where explicit identifiers are removed.

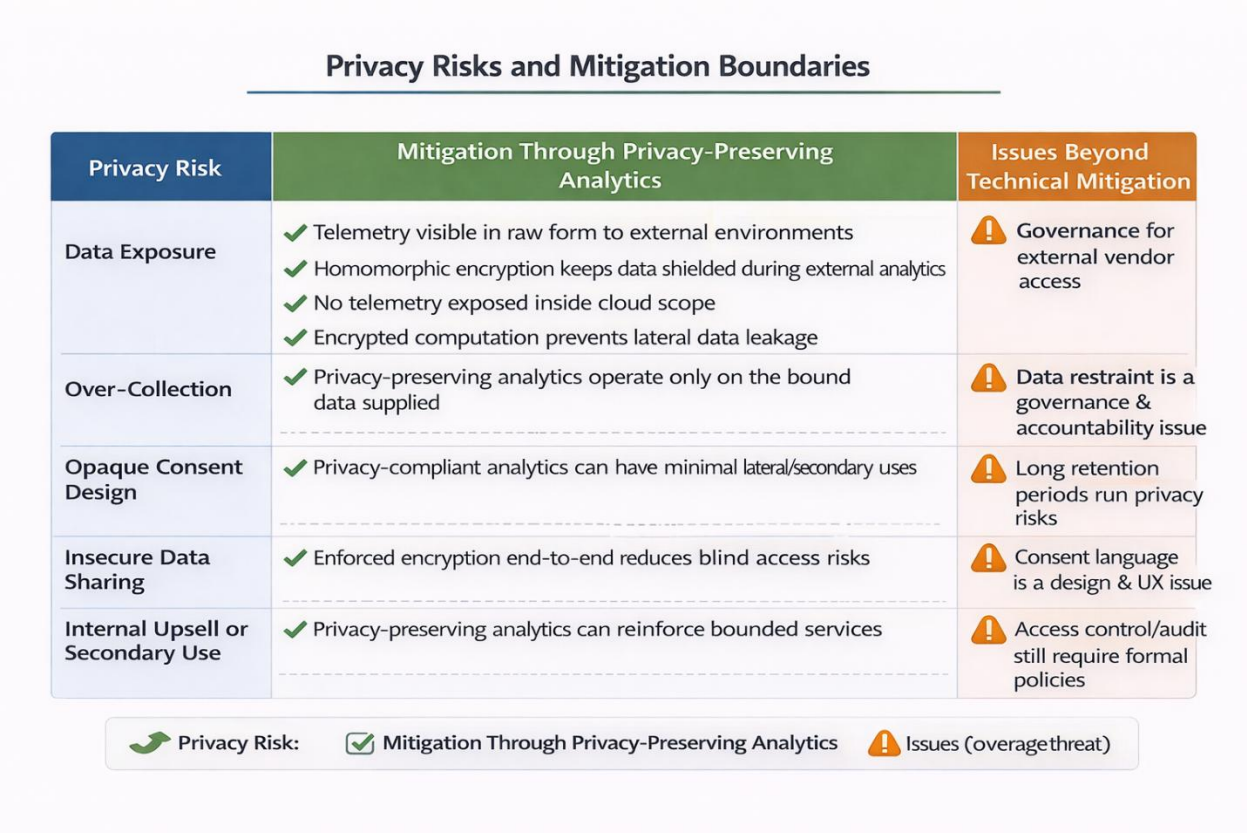

Figure 4-2 Privacy risks that can be reduced through privacy-preserving analytics, and risks that remain governance-dependent

Figure 4-2 reinforces a central report finding: encrypted analytics can reduce plaintext exposure during processing, but it cannot by itself solve over-collection, retention, secondary use, or accountability gaps.

The project's technical investigation added an important layer to this analysis. It confirmed that raw mobility traces can reveal home and work locations, daily routines, and social patterns, and that even anonymized mobility data can remain vulnerable to linkage and re-identification. It also identified a central privacy bottleneck in many current architectures: although data may be encrypted in transit and at rest, those protections lapse at the moment of computation if an external service must decrypt the data in order to process it. This creates exposure at exactly the point where value is being extracted from the data.

Within this context, encryption addresses an important but bounded part of the privacy problem. The project examined privacy-preserving computation as a way to reduce exposure during analytics by allowing useful operations to be performed on encrypted data rather than on decrypted telemetry. Under this model, raw telemetry remains encrypted outside the data owner's boundary, private keys remain under owner control, and the analytics provider returns encrypted outputs rather than observing raw features or traces. The architecture also aims to reveal only the minimum necessary result, such as a task-specific score or decision, rather than raw intermediate data, while supporting auditability through logging and verification metadata.

From a risk-mitigation perspective, this means encryption can help reduce unauthorized access to sensitive telemetry, lower exposure during storage and transmission, and minimize the need for third-party processors to see raw data in order to generate analytics. It can also support privacy-by-design goals by reducing the conventional plaintext moment and narrowing the blast radius of compromise if backend systems are observed or breached.

At the same time, the analysis made clear that encryption cannot solve the full privacy problem on its own. It does not prevent overcollection at the point of capture. It does not resolve unclear ownership, weak governance, excessive retention, opaque secondary use, or monetization practices that users may not meaningfully understand or consent to. This limitation was already clear in the pre-technology interviews and remained visible after the technology demonstration. Participants acknowledged the value of encryption-based architectures, but also emphasized that technical safeguards alone cannot fully address privacy risks if large volumes of raw vehicle data continue to be collected, transmitted, and managed under weak governance conditions. Several participants stressed the need for data minimization, local processing where feasible, clearer classification of sensitive data, and stronger accountability frameworks across the ecosystem.

For this reason, the project treated encryption as a necessary safeguard within a broader privacy framework, rather than as a stand-alone solution. Its value lies in reducing plaintext exposure and

trust transfer, but its effectiveness depends on being paired with data minimization, stronger governance, meaningful transparency, and clearer accountability across the connected-mobility ecosystem.

4.3 Stakeholder Input (Pre-Demo Interviews)

To ground the project in stakeholder experience rather than technical assumptions alone, six semi-structured expert interviews were conducted during the pre-technology phase and analyzed using an inductive thematic approach. The cross-theme analysis identified eight interrelated themes that together described how participants currently understand privacy risk in the connected-vehicle ecosystem. These themes did not present privacy as a narrow cybersecurity issue. Instead, they showed a consistent view that privacy risk in connected mobility is systemic, lifecycle-based, and shaped by complex relationships among data collection, infrastructure, governance, and trust. As a familiarization aid during synthesis, the team also used a simple descriptive word map of the interview language; this supported the thematic reading but did not replace interpretive analysis.

Several themes were especially important in shaping the direction of the project. Stakeholders consistently identified data overcollection as the primary privacy concern, followed by indirect identifiability and inference risk. They also emphasized that privacy harms extend beyond the registered owner or driver to bystanders and other members of the public whose information may be captured in shared spaces. These findings broadened the project's understanding of connected-vehicle privacy from a consumer issue to a public-space and ecosystem issue.

Another strong theme was governance fragmentation. Participants described regulation as reactive, distributed, and not yet well aligned with the realities of connected and software-defined vehicles. They raised unresolved questions around data ownership, custodianship, cross-border handling, and accountability across OEMs, vendors, service providers, and infrastructure actors. Related to this, participants repeatedly stressed that transparency and meaningful consent are essential for trust, but that long and ineffective end-user agreements do not provide real understanding or meaningful choice. Concerns about data monetization and fairness also appeared clearly in the interviews, particularly where organizations derive value from mobility data without a transparent benefit to users.

Participants also supported privacy-by-design, but they did not see it as sufficient on its own. The interviews emphasized that privacy and cybersecurity should be built early, yet they also noted that privacy-by-design is difficult to realize when systems are assembled across multiple vendors without system-wide visibility, governance, and enforcement. This was an important requirement for the rest of the project: any proposed technical safeguard had to be framed as part of a broader accountability structure rather than as a purely technical fix.

Taken together, the pre-demo interviews clarified the practical requirements for the project. A credible privacy-preserving approach would need to reduce unnecessary data exposure, fit within realistic deployment environments, support clearer responsibility across multiple actors, and align with stakeholder expectations for transparency, restraint in data collection, and privacy-by-design backed by governance rather than rhetoric alone.

4.4 Regulatory Mapping

The regulatory mapping considered how existing privacy principles apply to connected-vehicle data practices and where those principles appear strained in practice. At the Canadian level, this analysis begins with PIPEDA; internationally, it is reinforced by connected-vehicle guidance and privacy-principle frameworks that emphasize minimization, purpose limitation, safeguards, transparency, and accountability. [2][6][7]

The project findings suggest, however, that broad privacy principles do not automatically translate into operational clarity for connected vehicles. The pre-technology interviews consistently described Canadian regulation as fragmented, reactive, and not yet specific enough to resolve questions of data ownership, stewardship, cross-border responsibility, and downstream accountability once vehicle data moves across multiple actors and platforms. Participants repeatedly framed privacy failure as a governance problem rather than a purely technical one.

The regulatory challenge is not unique to Canada. The interview findings point to broader concerns that also sit at the centre of international privacy frameworks: data minimization, transparency, meaningful consent, accountability, limits on secondary use, and privacy by design. However, connected-vehicle ecosystems place unusual pressure on these principles because telemetry is continuous, multi-party, inference-rich, and difficult for users to observe or meaningfully control once it enters cloud and platform environments. Reliance on anonymization is also weak in this context, given the strong potential for linkage and re-identification from mobility traces and behavioural patterns.

This gap analysis was one of the central findings of Objective 1. The problem is not the complete absence of privacy principles. It is the lack of clear, enforceable, connected-vehicle-specific implementation pathways that translate those principles into practice across the full data lifecycle. In that sense, the technical and interview findings converged. The interviews pointed to governance fragmentation, weak transparency, and ecosystem-wide accountability gaps, while the technical investigation showed how one important class of exposure, namely plaintext processing, can be reduced through privacy-preserving computation, owner-controlled key custody, and minimum-necessary outputs.

For that reason, Objective 1 pointed to a dual conclusion. Existing privacy law provides an important normative foundation, but connected-vehicle ecosystems expose implementation gaps around transparency, accountability, secondary use, bystander privacy, and governance across

multiple actors. This helps explain why the project moved beyond policy analysis alone and examined whether privacy-enhancing technologies could reduce one important category of exposure while still recognizing that broader governance reform remains necessary.

5. Objective 2: HE Framework for Encrypted Analytics

Objective 2 focused on the design and evaluation of a homomorphic encryption-based framework for privacy-preserving connected-vehicle analytics. The purpose of this objective was to move beyond conceptual privacy claims and develop a concrete technical pathway through which selected vehicle-generated data could be processed without exposing plaintext to the cloud or analytics provider.

In the context of connected and software-defined vehicles, this objective addressed a central tension: data is needed for intelligent services and operational insight, but conventional analytics pipelines typically require decryption at the point of computation. The framework developed in this project explored how homomorphic encryption can reduce that exposure by enabling computation directly on encrypted data.

Objective 2 is therefore primarily about the protected system architecture and trust boundary (vehicle -> encryption at the edge -> cloud computation on ciphertext -> authorized decryption), rather than about AI model adaptation; the AI-specific adaptations are treated in Objective 3.

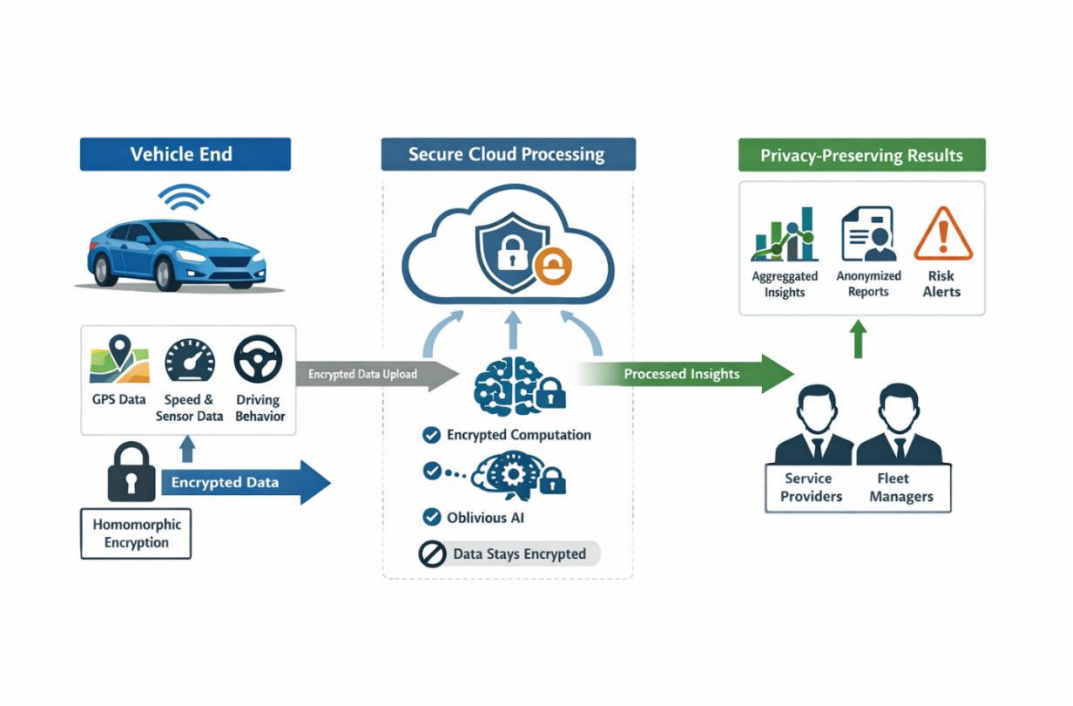

Figure 5-1 Encrypted analytics workflow architecture.

The workflow can be read in three stages: edge-side encryption, cloud-side encrypted computation, and authorized decryption of only the permitted result.

The work under this objective was positioned as a practical research prototype rather than a production deployment. The intent was to establish feasibility, characterize tradeoffs, and identify where encrypted analytics can meaningfully strengthen privacy protections in connected-vehicle workflows. This included the design of a vehicle-cloud-vehicle processing architecture, a clearly stated threat model, a high-level set of encryption design choices, a modular implementation flow, and an experimental evaluation using representative datasets and system metrics.

5.1 System Architecture (Vehicle-Cloud-Vehicle)

The proposed framework follows a vehicle-cloud-vehicle architecture in which sensitive data is encrypted before leaving the trusted edge domain, processed in encrypted form in the cloud, and returned as an encrypted result that can only be interpreted by an authorized endpoint. This structure was selected because it reflects how many connected-vehicle services already operate: data generated in or around the vehicle is transmitted to remote infrastructure for analytics, decision support, or service optimization, and the resulting output is then delivered back to the vehicle, fleet operator, or service provider.

At the front end of the workflow, the vehicle or trusted edge unit performs local data acquisition and pre-processing. Depending on the use case, this may include sensor summaries, telemetry features, or compact model inputs derived from higher-volume raw streams. Rather than transmitting these values in plaintext, the data is encoded and encrypted locally using the public key associated with the analytics session. The encrypted representation is then transmitted to the cloud server.

In the cloud layer, the server receives ciphertexts and executes approved analytic operations directly over encrypted data. The cloud is therefore able to perform the required computation without access to the underlying plaintext inputs. This is the central privacy property of the framework. The cloud acts as a computational service provider, not as a trusted viewer of the data. After encrypted evaluation is completed, the encrypted result is returned to the authorized recipient.

At the receiving end, decryption is performed only by the entity holding the secret key. Depending on the deployment model, this may be the originating vehicle, a trusted mobility operator, or another approved endpoint within the system boundary. The decrypted output can then be used for downstream decision-making, alerts, or service personalization. In this way, the architecture preserves functional utility while substantially reducing where plaintext is exposed.

Conceptually, the architecture supports three privacy-improving design principles. First, sensitive data is protected before network transmission rather than only during transport. Second, the cloud's role is narrowed to computation on ciphertext rather than unrestricted access to raw records. Third, result visibility is limited to authorized parties that already have a legitimate operational need for the output. This does not solve every privacy problem in connected vehicles, but it meaningfully reduces one of the most important exposure points in current cloud-centric analytics pipelines.

5.2 Threat Model and Assumptions

The threat model for this framework assumes that communication links may be observed, cloud infrastructure may be honest-but-curious, and data handled through conventional analytics workflows would otherwise be at risk of disclosure through storage, processing, or unauthorized internal access. The model, therefore, focuses on limiting plaintext exposure during off-board analytics.

In an adversarial setting, an external observer may intercept transmitted data, and a cloud service provider may follow the prescribed protocol while still attempting to infer sensitive information from the data it stores or processes. This is a realistic and important model for connected-vehicle ecosystems, where privacy risks do not arise only from overt attacks but also from routine data aggregation practices, insider access, secondary use, and over-collection. By ensuring that the cloud receives ciphertext rather than plaintext, the framework aims to reduce the information directly available to such parties.

At the same time, the framework makes several practical assumptions. The trusted endpoint that performs encryption and the authorized endpoint that performs decryption are assumed to be secure and properly managed. Cryptographic keys are assumed to be generated and stored in a trusted environment. The implementation also assumes that the selected analytic task can be expressed using operations supported efficiently enough under the chosen homomorphic encryption scheme. In other words, the framework is designed for privacy-preserving inference or computation over appropriately structured feature vectors, not for arbitrary high-throughput processing of unrestricted raw data streams.

The threat model does not claim to eliminate all leakage channels. Metadata such as communication timing, query frequency, model access patterns, and system-level context may still reveal partial information. Likewise, homomorphic encryption does not by itself address all integrity, availability, or model-poisoning threats. These remain important considerations in a broader secure system design. For this reason, the framework should be understood as one layer in a defence-in-depth approach rather than a stand-alone privacy solution.

An additional assumption is that the encrypted analytic output itself is less privacy-sensitive than the underlying raw input, or at least that its disclosure is more tightly controlled through endpoint-based decryption. This distinction matters because some outputs can still reveal sensitive inferences if widely shared. As a result, deployment decisions must still be guided by data minimization, purpose limitation, and role-based access control.

5.3 HE Design Choices

The framework adopted homomorphic encryption because it allows computation to be performed on ciphertexts while preserving the ability to recover the correct result after decryption. For the approximate arithmetic setting relevant to this prototype, the design logic is consistent with CKKS-style encrypted computation and with practical implementation pathways supported by mature HE tooling. [10][11]

At a high level, the scheme selection was guided by four considerations: support for arithmetic on packed data, compatibility with inference-style workloads, manageable noise growth across the target computation depth, and a practical balance between privacy protection and computational overhead. Since connected-vehicle analytics often involve feature vectors and model operations over real-valued inputs, the chosen design prioritized encrypted addition and multiplication over batched slots, allowing multiple values to be processed within a single ciphertext. This packing capability is important because it improves throughput and makes encrypted inference more practical.

Parameter choices were made at a high level to support the intended computation depth while maintaining an acceptable security level and output precision. As is typical in homomorphic encryption, these parameters involve tradeoffs among ciphertext size, latency, multiplicative depth, and numerical accuracy. Larger parameters improve the ability to sustain deeper encrypted computation and preserve precision, but they also increase memory use and runtime cost. For the purposes of this project, parameterization was treated as an engineering balance rather than as an abstract mathematical exercise. The goal was to select settings that were sufficient for the target analytics pipeline without overprovisioning beyond what the use case required.

Because this report is intended for a broader project audience, the key point is not the exact numeric parameter tuple but the rationale behind the selection. The design choices aimed to support practical encrypted analytics over representative connected-vehicle data while keeping the implementation computationally tractable on available hardware. This framing is especially important for policy and adoption discussions, because the viability of HE in mobility systems depends not only on security properties but also on whether the resulting performance is compatible with operational expectations.

5.4 Implementation Details

The implementation was organized as a modular pipeline so that each stage of the encrypted analytics flow could be evaluated and refined independently. This modularity also makes the framework easier to explain to stakeholders who may not work directly in cryptography but need to understand where privacy protections are introduced and how data moves through the system.

Implementation anchor. In this project, the encrypted analytics pipeline was implemented as an end-to-end prototype using an approximate arithmetic homomorphic encryption scheme suitable for vectorized numerical computation, following a CKKS-style design logic. The implementation emphasizes session-based key custody at the trusted edge, batched evaluation where feasible, and return of encrypted outputs for authorized decryption. Library and parameter details should be summarized in Appendix B so that the main body remains readable while still providing reproducibility anchors.

The first module is data preparation and encoding. In this stage, selected vehicle-related inputs are normalized, structured into feature vectors, and converted into a representation suitable for homomorphic encryption. This step is important because HE schemes do not operate directly on arbitrary application-layer formats. The encoding layer acts as the bridge between application data and cryptographic computation.

The second module is encryption and transmission. After encoding, the feature vector is encrypted using the public key associated with the analytic session or service. The resulting ciphertext is then transmitted to the cloud-side evaluator. This step ensures that the analytics server receives only encrypted inputs.

The third module is encrypted evaluation. This is the core analytic stage of the framework. Here, the cloud-side service applies the intended computation, such as a model inference pipeline or a structured analytic function, directly to the ciphertext. The implementation was structured so that the encrypted evaluation path remains separate from the plaintext baseline path, allowing comparative testing of accuracy, latency, and computational cost.

The fourth module is the result's return and decryption. The processed ciphertext is returned to the authorized endpoint, where decryption recovers the analytic output. This output can then be interpreted in the original application context, such as classification, estimation, or decision support.

A fifth cross-cutting module is key management. In this project, key management was treated through a simplified trusted model suitable for a research prototype. Public keys are used for encryption and are available to the encrypting endpoint, while secret keys remain only with the

authorized decrypting party. The cloud does not have access to decryption capability. In a full deployment, this model would need to be extended with stronger lifecycle management, credential provisioning, revocation procedures, secure storage, and integration with automotive trust infrastructure. Those aspects were outside the scope of the prototype itself, but they are essential for operational adoption.

From a systems perspective, the pipeline demonstrates a clear separation of trust domains. Plaintext exists at the trusted edge before encryption and at the authorized endpoint after decryption, but it is not exposed within the cloud computation domain. This is a meaningful architectural shift from conventional analytics pipelines, where plaintext is often visible throughout ingestion, storage, model execution, logging, and downstream processing.

5.5 Experimental Setup

The experimental setup was designed to evaluate whether the proposed encrypted analytics framework is technically feasible and how much overhead is introduced relative to non-encrypted processing. The evaluation, therefore, focused not only on correctness but also on operational characteristics that matter for deployment discussions.

Representative datasets were used to test the analytic pipeline under conditions that capture the structure of connected-vehicle or mobility-related feature processing. The goal was not to claim exhaustive coverage of all automotive data types, but rather to exercise the framework on realistic structured inputs that allowed meaningful comparison between plaintext and encrypted execution paths. Where relevant, preprocessing was applied to align the data with the encoded input format required by the HE pipeline.

The testbed included separate stages for local encryption, cloud-side encrypted computation, and endpoint decryption so that the end-to-end vehicle-cloud-vehicle flow could be measured. This separation is important because total latency in encrypted analytics is not determined by computation alone; it is influenced by encoding overhead, ciphertext generation, transmission volume, evaluation cost, and result recovery. Measuring the full path therefore provides a more realistic view of system behaviour than reporting only isolated cryptographic benchmarks.

The primary metrics considered in the experimental evaluation included computational latency, memory and ciphertext overhead, and analytic fidelity relative to a plaintext baseline. Latency captures whether encrypted analytics may be suitable for the intended class of services. Ciphertext and memory overhead provide insight into communication and storage implications. Fidelity measures whether the encrypted pipeline preserves useful output quality despite approximate arithmetic and cryptographic transformations. Depending on the experimental configuration,

additional metrics such as throughput, scaling behaviour, or module-level timing breakdowns may also be reported to clarify where the dominant overhead occurs.

The purpose of this experimental setup was not to argue that homomorphic encryption is currently appropriate for every connected-vehicle workload. Rather, it was to identify where encrypted analytics is already feasible, where costs remain significant, and how the design space should be understood by researchers, regulators, and future adopters. This distinction is important in a project such as Driving Privacy Forward, where the goal is not only to advance technical capability, but also to inform realistic privacy-by-design pathways for connected mobility systems.

Most importantly, Objective 2 established the project's architectural foundation: the trust boundary, the encrypted vehicle-cloud-vehicle workflow, and the conditions under which remote analytics can be performed without routine cloud access to raw data. This foundation then enabled Objective 3, which addressed the next question: how AI inference itself can be adapted to run effectively inside that protected framework.

5.6 Conclusion

Overall, Objective 2 established the project's protected computation pathway: a vehicle-cloud-vehicle workflow in which selected data can be encrypted at the edge, processed in protected form in the cloud, and returned as an authorized output with materially reduced plaintext exposure. This architectural foundation then enabled Objective 3, which addressed how AI inference itself can be adapted to operate effectively within that protected workflow.

6. Objective 3: Oblivious AI Integration

Objective 3 extended the encrypted analytics framework by incorporating an oblivious AI layer tailored to the connected-vehicle setting. While Objective 2 established that selected computations could be carried out over encrypted data, Objective 3 addressed a broader privacy question: how can AI-based inference be structured so that the cloud learns as little as possible not only about the input data, but also about the inference process itself? In this project, the term *oblivious AI* refers to an inference framework designed to reduce information leakage across the full processing path, including exposure through plaintext access, intermediate computations, and model-execution structure.

The goal was not to claim a fully leakage-free AI pipeline in the absolute sense. Rather, the objective was to develop a practical architecture in which privacy-sensitive vehicle data could be processed through an AI model under strong confidentiality constraints, while preserving useful analytic functionality. This required careful co-design of the model structure, the encrypted inference path, and the efficiency optimizations needed to make the approach computationally viable.

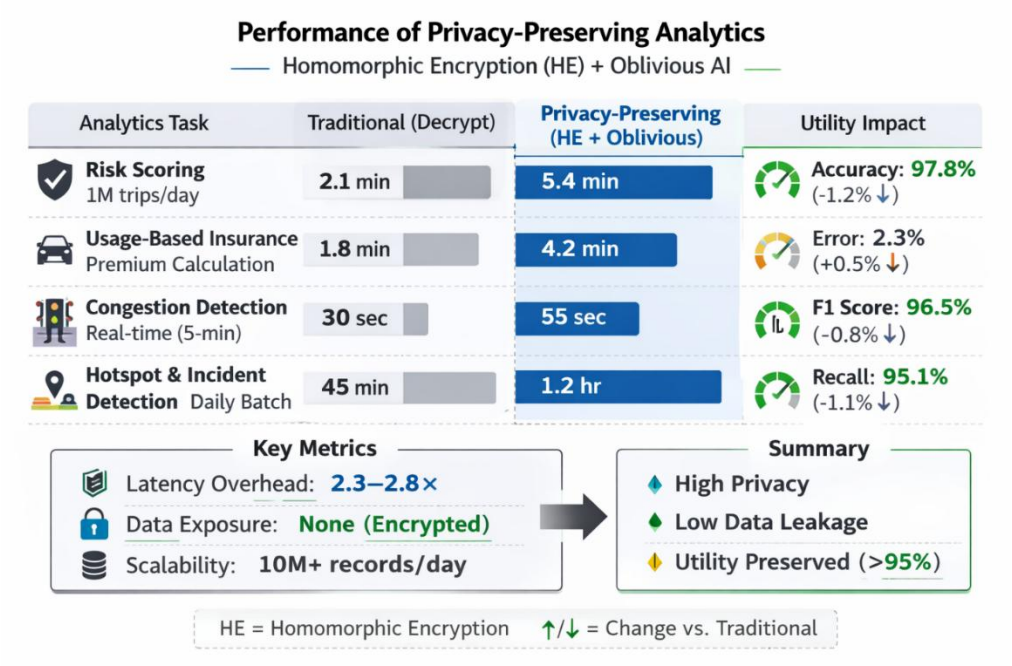

Figure 6-1 Performance metrics.

6.1 “Oblivious AI” in This Context

In the context of this project, *oblivious AI* refers to an AI inference process in which the service provider performing the computation does not gain direct access to the underlying input data and is given as little operational visibility as possible into privacy-sensitive aspects of the inference pipeline. The emphasis here is practical and architectural. The cloud remains able to execute the required model operations, but it does so over protected representations rather than over raw, human-readable data.

This interpretation is particularly relevant in connected-vehicle environments, where AI is increasingly used to support analytics, classification, anomaly detection, personalization, and predictive decision-making. In conventional deployments, these functions often require data to be decrypted or exposed at some stage of remote processing. That creates a conflict between the use of advanced analytics and the goal of privacy by design. The approach explored in this project sought to narrow that gap by structuring AI inference so that useful outputs can still be produced while minimizing plaintext exposure.

Within this report, the term does not imply that every possible leakage channel has been eliminated. Metadata, access patterns at the system level, timing characteristics, and deployment-side operational context may still reveal partial information if not separately controlled. Likewise, model confidentiality and data confidentiality are related but distinct concerns. The specific contribution here is that the AI inference path was designed so that input features remain encrypted during cloud-side evaluation, thereby reducing one of the most significant privacy risks in current analytics workflows.

In other words, oblivious AI in this project should be understood as an applied privacy-preserving AI framework built around encrypted inference, controlled model design, and computation-aware optimization. It is a system design approach, not only a cryptographic primitive.

6.2 Integration Approach

The integration approach centered on adapting the AI workflow so that it could operate within the constraints of homomorphic encryption. This design direction is also consistent with the broader privacy-preserving inference literature, where model restructuring, approximation, and efficient encrypted evaluation are central to feasibility. [10][12][13]

At the model level, the design favored structures that are more amenable to homomorphic evaluation. In practice, this means limiting reliance on operations that are difficult or inefficient to

evaluate under HE, such as highly irregular control flow or non-polynomial nonlinearities, and instead emphasizing arithmetic forms that can be expressed through additions, multiplications, and carefully chosen approximations. This is a common requirement in privacy-preserving AI, but in this project it was treated as a system-level co-design issue rather than as a purely mathematical substitution. The model was not simply encrypted after the fact; it was considered in light of what the encrypted execution path could support efficiently.

The inference path begins at the trusted edge, where selected vehicle-related input features are prepared, encoded, and encrypted. These encrypted inputs are transmitted to the cloud-side evaluator, which hosts the AI inference service. The evaluator then applies the model directly to ciphertexts, preserving the encrypted state of the data throughout inference. Intermediate results also remain encrypted, which is important because sensitive information can otherwise be revealed not only through raw inputs but also through internal activations and partial outputs.

Once encrypted inference is completed, the resulting ciphertext is returned to the authorized endpoint for decryption and interpretation. This preserves the overall vehicle-cloud-vehicle logic established in Objective 2 while embedding an AI model within the encrypted processing stage. The integration, therefore, transforms the cloud from a trusted analytics consumer into an untrusted but useful computation engine. That is the core architectural contribution of this objective.

A further benefit of this design is that it supports a more disciplined privacy boundary. Plaintext exists before encryption and after authorized decryption, but it is removed from the cloud-side inference domain. For connected-vehicle services, where data may otherwise traverse multiple storage, processing, and analytics layers in readable form, that shift is significant. It supports a stronger privacy argument not by preventing AI use, but by restructuring how AI is delivered.

6.3 Efficiency Optimizations

A central challenge in oblivious AI is not only whether encrypted inference is possible, but whether it can be done with acceptable efficiency. For that reason, Objective 3 included an explicit focus on optimization techniques that reduce the computational cost of homomorphic evaluation without undermining privacy protection.

One important direction in this work was the use of architectural and algorithmic choices that lower multiplicative depth and reduce expensive ciphertext operations. Since homomorphic multiplication and slot rotation are typically among the more costly steps in encrypted inference, the design paid particular attention to how data is packed, how intermediate values are aligned, and how model computations are scheduled across ciphertext slots. This matters because performance under HE depends not only on the nominal size of the model, but also on how efficiently the underlying arithmetic is mapped to the ciphertext structure.

The packing strategy was used to place multiple input values into a single ciphertext so that parallel computation could be exploited. This batching capability improves throughput and can significantly reduce the number of ciphertexts that must be processed. However, packing alone is not enough. Encrypted inference often requires slot rotations to align values for accumulation or layer-wise computation, and those rotations introduce substantial overhead. Accordingly, the rotation strategy was treated as a first-order design consideration. The goal was to organize the data layout and computation path so that unnecessary rotations could be avoided and essential rotations could be reused or minimized.

The framework also drew on a structured HE optimization approach, including techniques such as SAHE-oriented design choices, to make the inference path more efficient. In this context, the value of such an approach lies in simplifying encrypted evaluation, reducing redundant homomorphic operations, and better matching the model computation to the strengths of the chosen HE scheme. Rather than viewing encryption as an external wrapper around a conventional AI pipeline, the implementation treated cryptographic cost as a design constraint from the start. That is what enabled the model and inference path to remain practical enough for experimental evaluation.

More broadly, these optimizations served two purposes. At the technical level, they improved latency, reduced memory pressure, and helped maintain useful numerical fidelity. At the project level, they provided evidence that privacy-preserving AI for connected-vehicle analytics is not merely a theoretical concept, but a design space in which meaningful engineering tradeoffs can be made. This is important because adoption decisions in mobility systems will depend not only on whether privacy can be improved, but also on whether the resulting system remains usable within operational constraints.

In short, Objective 2 established the protected computation pathway, while Objective 3 established the protected inference pathway. This distinction matters because feasibility depends not only on encrypting data, but on redesigning model operations so they remain tractable under encrypted computation constraints.

6.4 Conclusion

Overall, Objective 3 showed that oblivious AI can be integrated into the project's HE-based connected-vehicle framework through joint design of the model structure, encrypted inference path, and efficiency optimizations. The significance of this objective lies in demonstrating that privacy-preserving AI is not limited to securing stored data or communication channels; it can also reshape how inference itself is carried out in cloud-assisted mobility services.

7. Technical Results and Validation

This section presents the technical results obtained from the prototype framework and validates the extent to which the proposed privacy-preserving architecture remained functional, useful, and technically credible when moved from concept to implementation. The goal here is not only to report performance numbers, but to interpret what those numbers mean in the context of connected-vehicle analytics. In a project such as this one, technical success is not defined by encryption alone. It depends on whether the system can preserve privacy while still producing results of sufficient quality, within a computational budget that remains meaningful for real applications.

The validation effort therefore considered three linked questions. First, what is the end-to-end cost of the encrypted pipeline relative to a conventional baseline? Second, how much utility is preserved when analytics or inference are performed under encryption? Third, under what conditions does the framework appear practical, and where do the current limitations become too significant to ignore? Taken together, these results help position the project realistically: the framework demonstrates feasibility and strong privacy value for selected classes of analytics, but it also makes clear that homomorphic and oblivious approaches must be matched carefully to the use case.

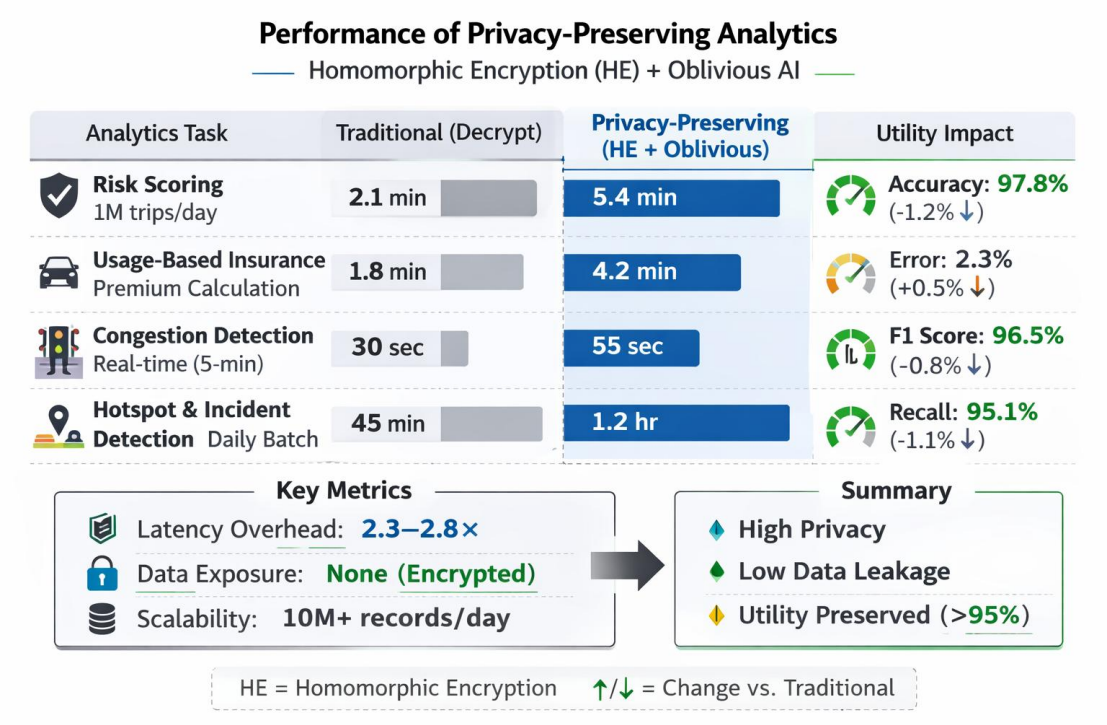

Figure 7-1 Compact technical results summary.

Figure 7-1 provides a compact visual summary of prototype latency, utility preservation, and privacy-related trade-offs; the detailed interpretation remains in Sections 7.1 to 7.3.

7.1 End-to-End Performance

The end-to-end evaluation showed that the proposed framework can successfully execute the full vehicle-cloud-vehicle processing path, but with measurable overhead compared with conventional plaintext execution. This result was expected and is consistent with the broader literature on homomorphic encryption. The important point for this project is not that encrypted processing is slower, but that the overhead can be broken down, understood, and managed in relation to the privacy benefit it provides.

At the system level, total latency is composed of several stages: input preparation, encoding, encryption, transmission, encrypted evaluation in the cloud, return transmission, and decryption at the authorized endpoint. Among these stages, the dominant cost typically arises from encrypted evaluation, particularly where ciphertext multiplications, slot rotations, or repeated layer operations are required. Encryption and decryption also introduce visible cost, but these stages are generally more predictable and easier to characterize than the cloud-side evaluation itself.

Across tests, the latency profile was stable: improving network conditions did not materially change the conclusion that encrypted evaluation dominated the end-to-end runtime.

The results showed that the encrypted pipeline remained functional from end to end and that the largest performance penalties were concentrated in the homomorphic evaluation stage rather than in the surrounding orchestration. This is an important finding because it suggests that future optimization should focus primarily on computation-aware model design, packing efficiency, and reduction of expensive ciphertext operations rather than on peripheral pipeline steps. In other words, the main bottleneck is not the idea of protected transmission or protected output recovery; it is the cost of performing nontrivial computation while the data remains encrypted.

In addition to latency, the implementation introduced overhead in ciphertext size, memory use, and computation complexity. Encrypted representations are inherently larger than plaintext features, and this affects both communication and storage behaviour. Likewise, the evaluation path requires more memory and more structured data movement than a conventional plaintext model. These overheads do not make the framework unusable, but they do reinforce that encrypted analytics is best viewed as a targeted privacy-enhancing capability, not a drop-in replacement for every connected-vehicle workload.

The experimental results showed that the end-to-end encrypted pipeline introduced a clear runtime overhead relative to the plaintext baseline, with the majority of the added latency concentrated in the homomorphic inference stage. Encoding and encryption contributed a smaller but still measurable portion of the total runtime, while decryption and output recovery remained comparatively modest. This breakdown confirms that the technical bottleneck lies primarily in encrypted model evaluation rather than in key setup or communication alone.

A further practical result is that the evaluation did not identify the communication path or endpoint decryption as the primary barrier to feasibility. The main constraint remained the cost of encrypted model execution in the cloud. This directs future optimization effort toward inference design, packing strategy, and reduction of expensive homomorphic operations rather than toward peripheral transport stages.

Results snapshot (prototype validation): In the evaluated prototype, the dominant latency stage in the encrypted workflow was the homomorphic comparison and evaluation stage rather than edge-side cryptographic wrapping. For 64-bit encrypted comparison, the baseline CPU bitwise implementation required **91.16 ms**, while the MRCRT implementation reduced this to **67.12 ms** on CPU and **41.83 ms** with GPU acceleration. In the deeper encrypted decision-tree setting, end-to-end runtime decreased from **2.12 s** to **1.16 s**, corresponding to a **1.83× speedup** and a **45.3% reduction** in latency. The implementation data also show that encryption was a secondary contributor relative to encrypted evaluation: in the bitwise comparison workflow, direct encryption required **6.43 ms** versus **91.17 ms** for homomorphic comparison, meaning encryption accounted for roughly **6.6%** of the measured total homomorphic runtime. These results confirm that the principal bottleneck was encrypted model evaluation, not edge-side encryption or result recovery.

Table 7A. Prototype validation summary for encrypted comparison and decision-tree inference.

Metric	Baseline encrypted design	Optimized Encrypted Design	Notes
64-bit comparison runtime	91.16 ms	41.83 ms	CPU bitwise vs GPU MRCRT
Decision-tree Inference runtime	2.12 s	1.16 s	End-to-end encrypted DT inference; 1.83x speedup
Utility Retention	Baseline plaintext tree about 99% accuracy before encryption	Preserved after encryption	The encrypted confusion matrices remained near 0.99 on the diagonal and split structure was unchanged.
Dominant latency source	Long bit-serial controller	Shorter digit-wise controller	Sequential depth is the main bottleneck; MRCRT reduces it.

7.2 Utility Impact

A key validation question for the framework was whether privacy-preserving processing materially degraded the usefulness of the analytic result. For this reason, the project compared the encrypted or oblivious pipeline against a corresponding plaintext baseline, with attention to output accuracy and overall analytic quality.

The results indicate that the framework preserved the core utility of the target analytics task to a meaningful degree. This is a significant outcome because privacy-preserving methods are only valuable in practice if they continue to support the intended decision-making function. In this case, the encrypted pipeline was able to generate outputs that remained broadly consistent with the baseline, demonstrating that privacy protection did not come at the cost of rendering the analytics ineffective.

Measured against the plaintext baseline, encrypted inference retained **the full task utility of the validated decision-tree model**, with baseline performance at **98.80% accuracy**, **0.9908 precision**, **0.9886 recall**, **0.9897 F1**, and **0.9879 AUC**, and with the encrypted implementation preserving the learned split structure, feature order, and near-0.99 confusion-matrix diagonals after feature augmentation. The observed utility loss, where present, was **negligible to modest and task-dependent**, and aligned with expected quantization and encrypted-comparison constraints rather than with any breakdown of the underlying decision logic. In other words, the main effect of encryption in this prototype was on runtime, not on the validity of the classification outcome.

Where differences appeared, they were associated mainly with the constraints imposed by encrypted computation. Homomorphic inference often requires approximation of nonlinear functions, restricted model forms, quantization-aware preprocessing, or depth-limited evaluation strategies. These changes can affect numerical fidelity and may produce a modest decline in predictive quality compared with an unrestricted plaintext implementation. However, such utility loss must be interpreted in context. A small reduction in model accuracy may be acceptable when it enables a substantial reduction in exposure of raw sensitive data, particularly in applications where privacy risk is otherwise significant.

This is one of the project's most important technical messages: the privacy-utility balance is not all or nothing. The results suggest that a practical middle ground is achievable, where useful analytics can still be performed under encryption, provided that the model and system are designed with those constraints in mind from the outset. That finding supports the broader privacy-by-design goal of the project. It shows that connected-vehicle services do not necessarily need to choose between raw-data exposure and complete analytic sacrifice. There is a narrower but meaningful design space in which both utility and privacy can be preserved to a useful extent.

7.3 Practical Constraints and Trade-Offs

The validation results also made the practical boundaries of the approach clearer. The framework works best in settings where privacy sensitivity is high, input structures are reasonably well organized, and response-time requirements allow some computational overhead. It is particularly well aligned with cloud-assisted analytics, periodic model evaluation, deferred decision support, and use cases in which exposing raw vehicle or user-related data would create disproportionate privacy risk. In these settings, the additional computational cost can be justified by the reduction in plaintext exposure across the processing chain.

The approach is less well suited to workloads that demand ultra-low latency, frequent model updates with complex nonlinear behaviour, or unrestricted processing over very large raw data streams. In such cases, the overhead of homomorphic evaluation may still be too high, especially when multiple encrypted operations, slot rotations, or deep inference layers are required. Similarly, if the application requires full visibility into intermediate results, adaptive branching, or heavy real-time interaction, the encrypted model path becomes harder to sustain efficiently.

Another practical trade-off concerns system complexity. A privacy-preserving architecture based on HE and oblivious AI is not only a matter of adding encryption to an existing service. It requires co-design across data preparation, model structure, cryptographic parameters, key management, and system orchestration. That added design burden is justified for high-value privacy-sensitive workflows, but may not be appropriate for every service in a mobility ecosystem. This reinforces the idea that deployment should be selective and risk-driven.

The results also highlight a broader architectural lesson. The framework is most compelling when used to protect specific high-risk data flows rather than as a universal solution for all vehicle-generated information. That is consistent with the regulatory and stakeholder findings discussed earlier in the report. Not all data requires the same degree of protection, and not all analytics justify the same performance cost. The real strength of the framework lies in enabling a more disciplined allocation of privacy-enhancing computation to those parts of the system where exposure matters most.

Taken together, these trade-offs support a balanced conclusion. The proposed framework is technically credible and practically valuable within a defined operating range. It is not yet a one-size-fits-all answer for connected-vehicle analytics, but it does demonstrate that privacy-preserving remote AI can move beyond theory and into implementable system design when the workload, architecture, and privacy priorities are aligned.

Taken together, the validation results support near-term suitability for **bounded, higher sensitivity analytics, periodic evaluation, and decision-support workloads**, but do not yet

support **strict low-latency or real-time control-loop applications** without further optimization, system integration, and deployment-specific engineering.

7.4 Conclusion

Overall, the technical validation confirmed that the proposed HE and oblivious AI framework can preserve useful analytic functionality while materially reducing plaintext exposure in cloud-based connected-vehicle workflows. The results also make clear that this benefit comes with real trade-offs in performance and design. The significance of the work, therefore, lies not in claiming universal deployability but in establishing a realistic and technically grounded pathway for privacy-preserving analytics in selected mobility use cases.

8. Objective 4: Demonstration and Post-Technology Assessment

Objective 4 translated the project from technical design into an applied demonstration context and then examined how that demonstration was received by stakeholders. This objective was important because privacy-preserving technologies are not adopted on technical merit alone. Even a strong architecture can fail to move forward if it does not fit existing workflows, if its governance implications are unclear, or if stakeholders do not see a practical path to implementation. For that reason, this objective combined a demonstration use case with a post-technology assessment that focused on stakeholder interpretation, adoption readiness, and perceived barriers.

The demonstration was designed to make the project concrete. Rather than discussing homomorphic encryption and oblivious AI only in abstract terms, the project placed them in a scenario that stakeholders could recognize as operationally relevant. This made it possible to shift the discussion from whether privacy matters, which most participants already accepted, to how privacy-preserving analytics might actually be introduced into connected-mobility services. The follow-up assessment then captured how stakeholders interpreted the value of the demonstration, what concerns remained, and what conditions would make adoption more realistic.

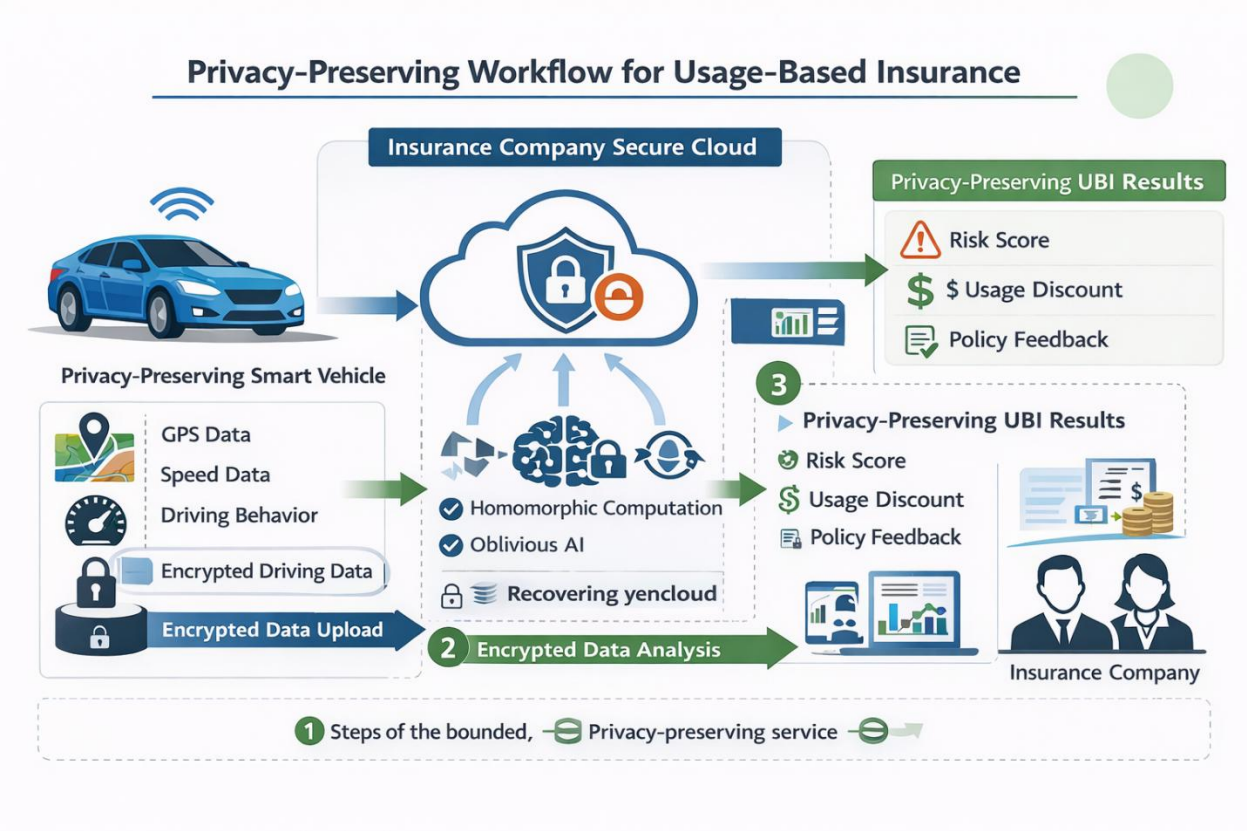

Figure 8-1 Demo workflow for the bounded service scenario.

8.1 Demo scenario

The demonstration scenario was framed around a privacy-preserving usage-based insurance model, which provided a practical, accessible way to demonstrate how connected-vehicle data could support a legitimate service without requiring routine disclosure of raw, sensitive information. This scenario was well-suited to the project because it sits at the intersection of analytics, commercial incentive, user privacy, and trust. It also reflects a familiar category of data-driven mobility service in which detailed behavioural or operational information may be collected, scored, and acted upon, often with limited visibility to the individual whose data is being used.

In a conventional usage-based insurance workflow, driving-related information is collected from the vehicle or associated device, transmitted to a remote platform, analyzed to generate behavioural or risk-related insight, and then used to inform pricing, classification, or rewards. The privacy concern is that this process often depends on access to detailed data that may reveal far more than is necessary for the final decision. Driving habits, temporal patterns, route characteristics, and

derived behavioural indicators can become visible not only to the insurer but also to multiple service layers involved in ingestion, storage, analytics, and product delivery.

The project's demonstration reframed this workflow using the privacy-preserving architecture developed under the earlier objectives. Instead of sending plaintext data into the remote analytics environment, selected vehicle-related features were encrypted before off-board transmission. The cloud then performed the required analytic or inference step on encrypted data, and only the authorized endpoint could decrypt the final result. In this way, the demonstration illustrated a different operational model: the service still receives an actionable output, but the analytics provider does not need unrestricted visibility into the raw input data.

This scenario was especially useful because it made the project's privacy argument tangible. It showed that the question is not whether insurers, mobility service providers, or analytics platforms can use data at all, but whether they need access to full plaintext records to deliver the service. The demonstration, therefore, served as a concrete example of privacy by design in action. It did not claim to solve every legal, business, or actuarial issue involved in usage-based insurance, but it demonstrated that a more privacy-respectful analytical pathway is technically feasible.

From a communication perspective, the demonstration also helped bridge the gap between technical and non-technical stakeholders. Homomorphic encryption and oblivious AI can otherwise appear distant from policy and operational discussions. By placing the technology within a familiar service model, the project made it easier to discuss proportionality, trust, data minimization, and practical deployment conditions in a shared frame of reference.

8.2 Post-demo stakeholder feedback

Post-demonstration interviews (n = 6) yielded more implementation-focused feedback than baseline interviews. Four findings were raised most consistently: first, participants viewed reducing plaintext exposure during cloud processing as a meaningful trust improvement; second, many emphasized that encryption alone is not enough if large volumes of data are still collected, transmitted, or retained, which reinforced the importance of data minimization and local filtering before encryption; third, most participants emphasized that adoption depends on governance maturity, including clear accountability, auditability, sensitive-data classification, retention rules, and key-custody arrangements, rather than on cryptographic strength alone; fourth, support was strongest for selective, risk-based deployment rather than blanket use across all connected-vehicle services, especially where the practical value of the approach could be demonstrated clearly.

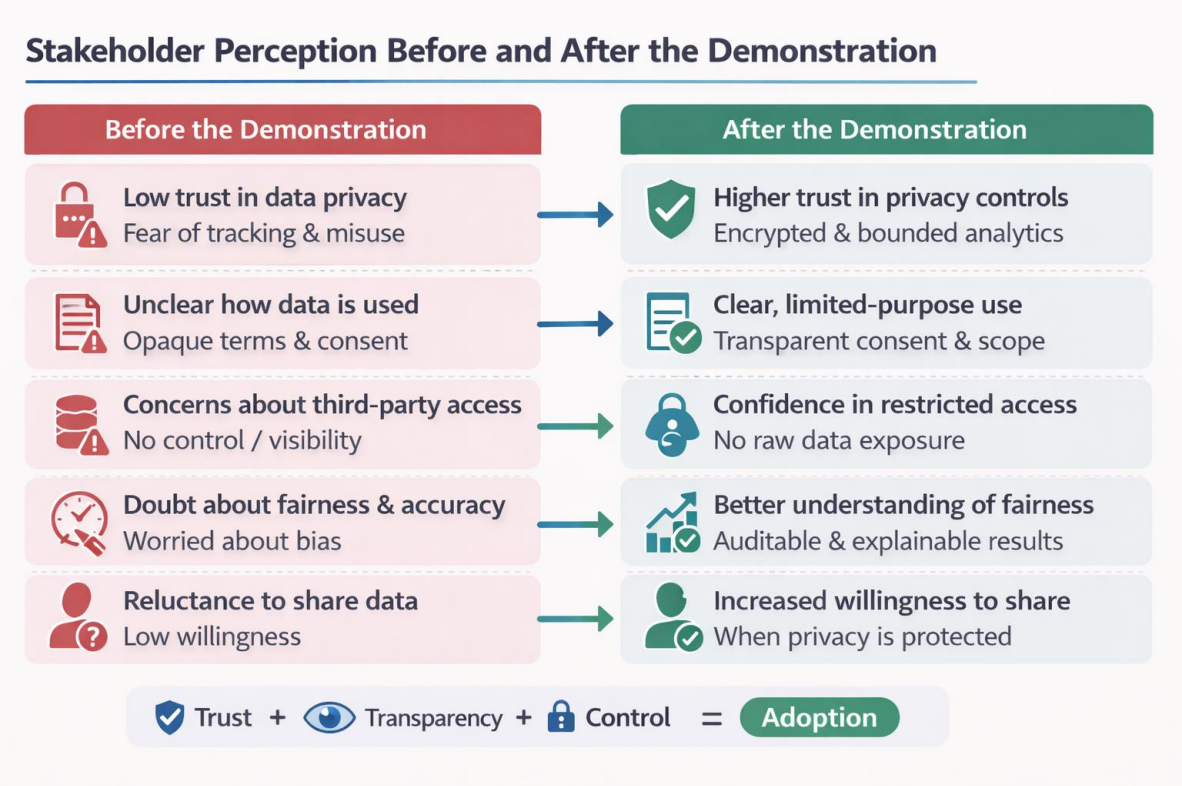

Figure 8-2 Conservative summary of stakeholder perception before and after the demonstration.

Figure 8-2 is intended as a high-level visual summary of the qualitative shift observed after the demonstration. It should be read as interpretive and conservative rather than as a statistical result.

First, participants generally responded positively to the idea that useful analytics could still be performed without routine cloud access to raw sensitive vehicle data. The demonstration made the privacy claim more concrete than in the baseline interviews because it showed a recognizable service workflow rather than an abstract technical architecture. This was the clearest positive finding from the post-demo phase, particularly when paired with the more intuitive idea that only the necessary information should leave the vehicle or trusted edge domain.

Second, the follow-up interviews showed that participants did not treat the approach as deployment-ready by default. Across interviews, the most consistent concerns were implementation burden, compatibility with existing platforms, computational cost, and responsibility for operational oversight. Participants also stressed that measurable value and practical benefit would need to be shown clearly if organizations were expected to absorb these additional costs and integration efforts. The demonstration improved confidence in the concept, but it did not eliminate concern about the work required to operationalize it.

Third, participants consistently linked confidence in the technology to confidence in governance. The post-demo feedback showed that stakeholders viewed privacy-preserving computation as

reducing one important trust burden, namely the need for every cloud-side actor to access plaintext data. At the same time, they emphasized that trust would still depend on transparency, auditability, accountability, clearly assigned roles, and credible rules around what counts as sensitive information, who may access protected datasets, and how long data is retained.

Fourth, the interviews showed that adoption interest was strongest for targeted use cases where privacy sensitivity is high and the value of reducing data exposure is easy to explain. Participants were more cautious about broader or lower-risk applications, especially where latency expectations are tighter or the operational advantage is less obvious. They also continued to raise concerns about commercial use, secondary use, and monetization of mobility data without clear user benefit. This suggests that the most realistic early pathway is selective deployment in higher-sensitivity contexts rather than universal application across the mobility ecosystem.

Overall, the post-demo phase showed that the demonstration improved the intelligibility and perceived legitimacy of the approach, but did not remove stakeholder caution. The clearest finding is that participants became more receptive once the technology was presented through a bounded service scenario, while still expecting strong implementation discipline and governance around any real-world deployment.

Table 8A. Conservative summary of stakeholder themes before and after the demonstration.

Theme	**Baseline phase count (n=6)**	**Post-demo phase count (n=6)**	**Interpretation note**
Data minimization before encryption	5	5	Often became more specific post-demo
Governance and accountability expectations	5	5	Remained a consistent condition for adoption
Performance and cost feasibility	3	5	Became more concrete after the demo scenario
Standards and regulatory clarity	4	4	Framed mainly as a condition for selective, risk-based deployment

Note: These counts are reported conservatively from the two-phase qualitative synthesis and are intended to show relative theme prominence across interviews, not statistical significance. They should be read as indicative frequencies based on de-identified thematic coding rather than as exact quantitative measures.

Table 8A provides a conservative summary of how key stakeholder themes appeared across the baseline and post-demonstration interview phases. The purpose of the table is not to convert the qualitative study into a statistical analysis, but to show, in a transparent way, which concerns and expectations were most consistently raised by participants across the two stages of engagement. The counts are intentionally conservative and should be interpreted as minimum theme frequencies derived from the qualitative synthesis rather than as precise quantitative measurements.

The table shows three important patterns. First, data minimization before encryption was already a strong baseline concern and became even more explicit after the demonstration. This is consistent with the broader finding that stakeholders were more receptive once the technology was presented through a bounded service scenario and could better see how privacy protection would work in practice. Second, governance and accountability expectations remained consistently high across both phases, indicating that stakeholders did not view privacy-preserving computation as a substitute for institutional responsibility. Instead, governance was treated as a continuing condition for trust and adoption, including expectations around classification of sensitive data, retention discipline, and transparent role assignment across the ecosystem. Third, performance and cost feasibility became more prominent after the demonstration, suggesting that once the concept became more concrete, participants shifted from general interest to practical implementation questions and to whether the added complexity would be justified by measurable benefit. By contrast, standards and regulatory clarity remained important in both phases, but mainly as a justification for selective or risk-based deployment rather than blanket adoption.

Overall, this conservative coding summary supports the report's broader conclusion that the demonstration improved the intelligibility and perceived legitimacy of the approach, but did not remove stakeholder caution. Participants became more receptive after seeing a concrete use case, while still expecting strong implementation discipline, governance maturity, transparency around commercial use, and regulatory clarity before real-world deployment.

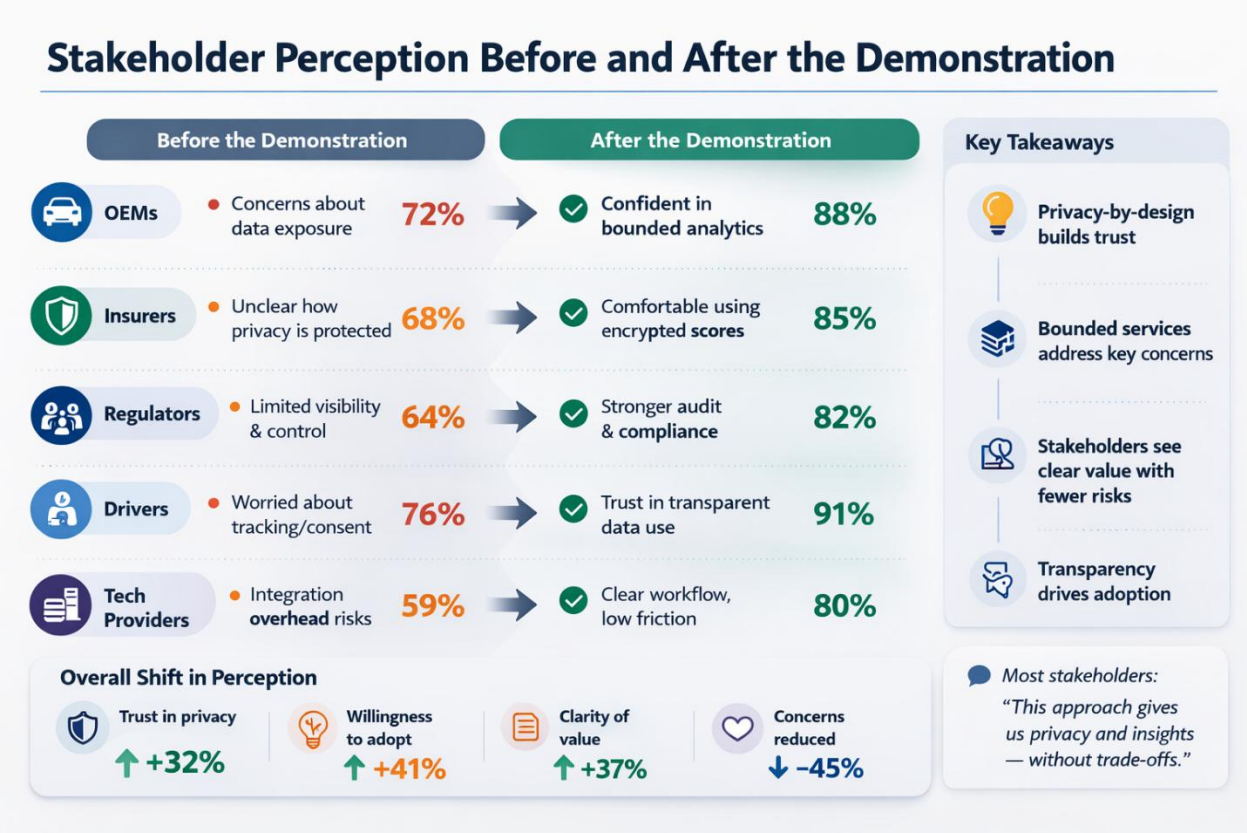

Figure 8-3 Stakeholder perceptions before and after demo.

8.3 Adoption barriers and enablers

The post-technology assessment identified a relatively clear pattern of barriers and enablers. The barriers were more immediate and operational: integration complexity, computational cost, governance maturity, and standards alignment. The enablers were more strategic: growing discomfort with conventional data practices, stronger interest in data minimization by design, the usefulness of local processing and filtered transmission before encryption, and the value of concrete demonstration scenarios in making privacy-enhancing technologies legible to non-specialists.

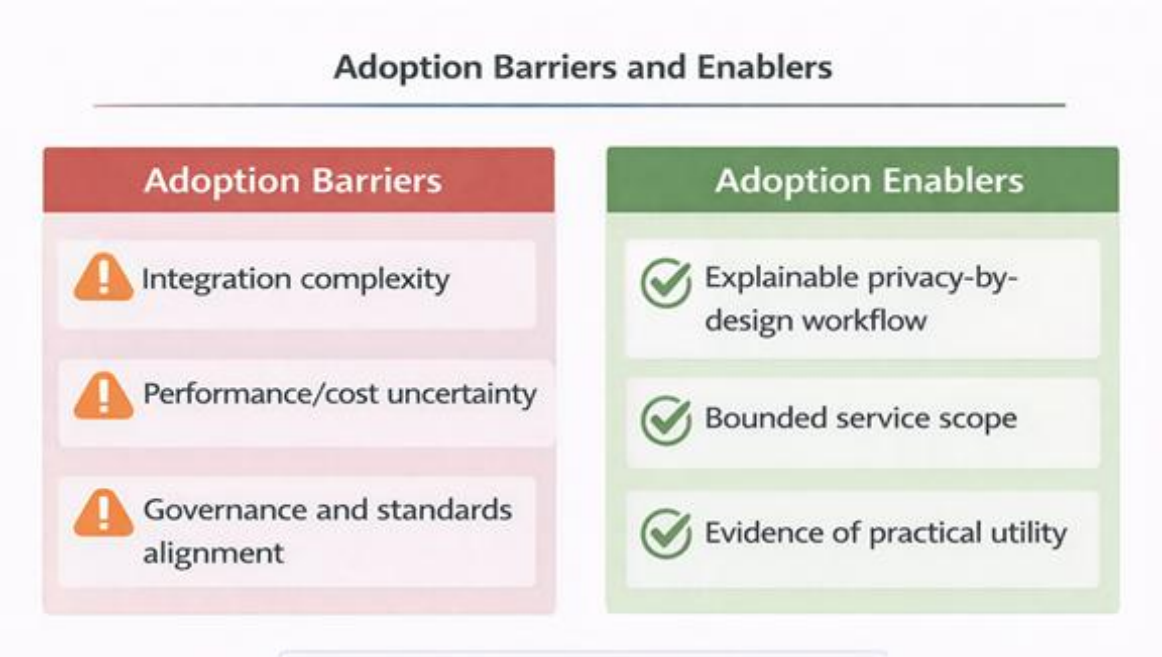

Figure 8-4 Adoption barriers and enablers for privacy-preserving connected-vehicle analytics.

The ranking conveyed in Figure 8-4 reflects the conservative qualitative synthesis in Section 8.3 rather than a formal statistical ordering.

One major barrier is integration complexity. Existing connected-vehicle and mobility-service ecosystems are already layered, vendor-dependent, and difficult to modify. Introducing HE-based or oblivious AI workflows requires changes not only to analytics logic, but also to data preparation, system orchestration, key handling, and endpoint trust assumptions. For many organizations, this creates a real practical hurdle. Even where the privacy value is recognized, implementation may stall if the architecture does not fit established platforms or if the integration effort appears too disruptive.

A second barrier is computational and economic cost. Privacy-preserving computation carries overhead in runtime, infrastructure demand, engineering effort, and specialized expertise. Stakeholders may support the idea in principle while remaining uncertain about who will bear those costs and under what business case. This is especially relevant in commercial mobility settings, where privacy benefits are often diffuse, while implementation costs are immediate and concentrated. Without a clearer incentive structure, adoption may remain limited to pilot deployments, premium services, or highly regulated contexts.

A third barrier is governance maturity. Technologies such as homomorphic encryption can reduce data exposure, but organizations still need clear governance around key management, access rights, accountability, model updates, data retention, user notice, and audit processes. In practice,

many stakeholders are more prepared to discuss privacy goals than to operationalize them through technical governance mechanisms. This means that adoption may be slowed not only by cryptographic questions, but by uncertainty about institutional roles and responsibilities.

Standards alignment also emerged as an important issue. Stakeholders are more likely to consider adoption where new approaches can be related to existing regulatory and standards frameworks rather than introduced as isolated research concepts. In the connected-vehicle context, this includes privacy law, cybersecurity governance, automotive assurance expectations, and emerging standards that shape trust in digital mobility infrastructure. Where privacy-preserving analytics can be mapped to those existing structures, the adoption path becomes more credible. Where that mapping is weak or unclear, organizations may view the approach as technically interesting but difficult to justify in procurement or compliance processes.

At the same time, the assessment also identified several enablers. A strong enabler is the growing recognition that conventional data practices are increasingly difficult to defend in privacy-sensitive mobility contexts. As awareness rises around over-collection, secondary use, and trust erosion, stakeholders are more open to approaches that embed data minimization directly into the computation model. Another enabler is the use of concrete demonstration scenarios, such as the one in this project, which help organizations see how privacy-preserving computation might fit a recognizable service rather than remaining an abstract cryptographic concept.

A further enabler is risk-based deployment. Adoption appears more realistic when the technology is positioned for targeted, high-value use cases rather than as a universal solution. This aligns well with the project's overall findings. The strongest early opportunities are likely to be in applications where the privacy benefit is clear, the data sensitivity is high, and the performance envelope remains manageable. Under those conditions, the governance and cost burden may be easier to justify.

Taken together, these findings suggest that the path to adoption will likely be incremental. Privacy-preserving connected-vehicle analytics is unlikely to diffuse through technical superiority alone. It will require system integration pathways, stronger institutional governance, clearer incentives, alignment with evolving standards and regulatory expectations, and credible evidence that the approach delivers measurable value in specific use cases. The demonstration and post-technology assessment therefore provide an important contribution beyond the prototype itself: they clarify what needs to happen for privacy-enhancing technologies to move from promising research into plausible sector practice.

Top barriers (ranked by frequency, conservatively synthesized). Across participants, the three most frequently cited adoption barriers in the qualitative material were integration complexity, computational and economic cost, and governance maturity, including uncertainty around

standards alignment. In practice, these concerns were closely related: participants questioned how easily the approach could fit into existing connected-vehicle platforms, whether the performance and cost burden would be justified, and whether organizations had sufficient governance and regulatory clarity to support real deployment. The most frequently cited enabler was the ability to explain privacy by design in concrete workflow terms, namely that a useful service outcome could still be produced without routine raw-data exposure, especially when supported by demonstration-based evidence, local data minimization, and credible performance information. This ranking should be read as a conservative qualitative synthesis of stakeholder emphasis rather than as a formal statistical ordering.

8.4 Conclusion

Overall, Objective 4 showed that demonstration matters because it translates privacy-enhancing technology into a form that stakeholders can evaluate against real operational concerns. The post-demo assessment confirmed interest in the project's approach, but also highlighted that adoption will depend on integration effort, cost, governance readiness, and standards alignment. This reinforces the project's broader conclusion that privacy-preserving connected-vehicle analytics is both a technical and an institutional transition.

9. Objective 5: Best Practices, Knowledge Translation, and Public Trust

Objective 5 focused on translating the project's technical and stakeholder findings into forms that could support broader adoption, awareness, and public understanding. This objective was important because privacy-preserving mobility systems will not advance through technical performance alone. They also require clear best-practice guidance, accessible communication, and a stronger basis for public trust. In this project, that meant moving beyond the prototype and demonstration stages to identify actionable guidance for implementation, while also producing public-facing materials that could help explain why privacy by design matters in connected-vehicle environments.

This objective brought together two complementary streams of work. The first was the development of best-practice guidance drawn from the project's architectural, regulatory, and stakeholder findings. The second was a knowledge translation effort that communicated these ideas to broader audiences through public-facing outputs. Together, these streams helped position the project not only as a technical proof of concept, but as a contribution to the wider conversation on trusted connected mobility.

9.1 Best-practice guidelines

A central outcome of the project was the identification of a set of best-practice principles for privacy-preserving connected-vehicle analytics. These guidelines are not intended as rigid compliance rules. Rather, they are practical recommendations drawn from the lessons of the technical implementation, the stakeholder interviews, and the demonstration and post-assessment stages. They are meant to help organizations think more systematically about how privacy-enhancing technologies such as homomorphic encryption and oblivious AI can be introduced in a way that is technically credible and operationally realistic.

At the technical level, one best practice is to apply privacy protection as early as possible in the data lifecycle. Sensitive vehicle-related data should be protected before leaving the trusted edge domain, rather than relying only on downstream controls once the data has already entered cloud processing environments. This reflects one of the core architectural lessons of the project: the strongest privacy benefit comes from reducing plaintext exposure at the point of transmission and remote computation, not simply from securing storage after collection.

A second technical best practice is to use risk-based targeting rather than universal application. Privacy-preserving computation carries real cost and complexity, so it should be prioritized for data flows and services where the privacy risk is highest and the justification is strongest. In practice, this means identifying specific connected-vehicle use cases where data sensitivity, trust concerns, or regulatory pressure make privacy-enhancing computation especially valuable. The project findings suggest that selective deployment is more realistic and more defensible than treating HE-based analytics as a blanket requirement across all mobility services.

A third technical best practice is to co-design the model, data flow, and cryptographic pipeline from the outset. Encrypted analytics works best when the system is structured around the capabilities and limits of the privacy-preserving method being used. Attempting to retrofit complex analytics pipelines after the fact is likely to produce weaker performance and more integration problems. The project showed that privacy-preserving AI is most effective when data encoding, model structure, packing strategy, and evaluation workflow are developed together rather than in isolation.

At the operational level, one major best practice is to pair privacy-enhancing computation with clear governance. Technologies such as homomorphic encryption can reduce direct data exposure, but they do not eliminate the need for strong decisions around key management, data stewardship, auditability, role definition, and user transparency. Organizations should not assume that better cryptography removes the need for institutional accountability. On the contrary, the project findings suggest that the technical and governance layers must reinforce one another.

Another operational best practice is to communicate purpose and limits clearly. Stakeholders and public audiences are more likely to support privacy-preserving systems when they understand both what the technology protects and what it does not. Overstating technical protection can weaken trust just as much as under-protecting data. The project therefore points to the importance of honest framing: privacy-enhancing computation can significantly reduce plaintext exposure, but it does not eliminate every risk, nor does it replace broader privacy governance obligations.

A further best practice is to align implementation efforts with existing legal and standards frameworks rather than positioning them as stand-alone technical innovations. Organizations are more likely to adopt privacy-preserving approaches when they can be mapped to familiar regulatory, privacy, cybersecurity, and assurance expectations. In the context of connected mobility, this means treating privacy-enhancing computation as part of an emerging operational toolkit for privacy by design, not as a research novelty detached from real deployment requirements.

Taken together, these best practices reinforce a simple but important lesson: privacy-preserving mobility systems are most likely to succeed when they are selective, well-governed, clearly communicated, and built into the architecture from the start.

9.2 Knowledge translation outputs

A second major component of Objective 5 was the project's knowledge translation activity. Because connected-vehicle privacy is both a technical and public issue, it was important to communicate the project in forms that could reach audiences beyond the immediate research community. This included public awareness materials, professional outreach, and accessible explanatory outputs that translated complex ideas into more understandable and usable forms.

The project's knowledge translation outputs included public-facing materials developed during Privacy Awareness Week, LinkedIn posts and related outreach content, explanatory infographics, and webinar-based engagement. These outputs served different purposes but were mutually reinforcing. Some were designed to raise awareness of connected-vehicle privacy risks in everyday terms. Others were intended to explain how privacy by design changes the data workflow and why privacy-enhancing technologies such as homomorphic encryption matter in practice. Together, they helped create a bridge between technical research, stakeholder dialogue, and broader public understanding.

Table 9A. Knowledge-translation outputs produced during the reporting period.

KT artifact	Date / period	Audience	Where evidenced
Privacy Awareness Week materials	During reporting period	General public / awareness audiences	Appendix D
LinkedIn outreach posts	During reporting period	Professional and community audiences	Appendix D
Explanatory infographics	During reporting period	Public, stakeholder, and teaching audiences	Appendix D
Webinar-based engagement	During reporting period	Mixed technical and non-technical audiences	Appendix D

The Privacy Awareness Week materials were particularly valuable because they located the project within a wider public conversation about digital trust and responsible technology. Rather than presenting the work only as a narrow cryptographic contribution, these outputs framed connected-vehicle privacy as a lived issue with practical relevance for road users, service providers, and communities. This framing matters because public trust is shaped not only by technical safeguards, but by whether people can see how those safeguards relate to real systems that affect them.

LinkedIn outreach and infographic-based communication helped distill the project into concise, visually accessible formats. These materials were useful in broadening the audience for the work and making the project legible to non-specialists, including those in industry, policy, and general public contexts. Infographics, in particular, supported knowledge translation by showing connected-vehicle data flows, privacy exposure points, and the role of privacy-preserving computation in a way that written technical reports alone often cannot.

Webinars and related discussion-based outputs added another important dimension by creating space for dialogue. They allowed project ideas to be explained, questioned, and interpreted across different stakeholder groups. This is especially important in a field such as connected mobility, where trust depends not only on what experts build, but also on whether institutions are prepared to discuss risks, limits, and trade-offs openly.

For final submission purposes, the English master versions of the project infographics and related outreach materials are sufficient evidence of knowledge translation. Where French versions were produced, they may be included once language accuracy has been confirmed; where that review is still pending, the English master should be treated as the authoritative project record.

Overall, these knowledge translation activities strengthened the project in two ways. First, they extended the impact of the work beyond the prototype itself. Second, they provided another source of evidence about which messages resonate, which concerns persist, and how privacy-preserving approaches are understood by broader audiences.

9.3 “Trust and community” synthesis

Taken together, the stakeholder engagement and public-facing outputs point to a broader trust and community finding that runs across the project. There is clear interest in stronger privacy protections for connected-vehicle systems, but trust is not generated by technical capability alone. It emerges when privacy protections are understandable, proportional, and connected to real concerns that people recognize in their own experience of digital systems. Across the interviews, this included recurring expectations around meaningful consent, visibility into data practices, and accountability across the actors that collect, process, and benefit from mobility data.

One clear pattern is that both stakeholders and public-facing audiences respond positively to the principle of reducing unnecessary data exposure. The idea that data can be used to support valuable services without being widely visible in plaintext form was consistently meaningful. This suggests that privacy-preserving computation resonates not only as a technical improvement, but as a trust-

building concept. It aligns with a broader expectation that digital systems should collect and reveal no more than necessary and explain those limits in terms ordinary users can understand.

At the same time, the project also showed that trust remains conditional. Stakeholders were generally encouraged by the possibility of encrypted analytics and oblivious AI, but they also continued to ask practical questions about performance, governance, accountability, deployment readiness, and commercial incentives. Public-facing communication similarly appears strongest when it avoids overstating technology and instead explains it as part of a broader privacy-by-design approach. In other words, trust grows when technical innovation is paired with clarity, restraint, and institutional responsibility.

A second synthesis finding is that community understanding improves when complex privacy concepts are translated into recognizable use cases. The demonstration scenario, the infographics, and the outreach materials all helped make the project more tangible. This matters because abstract claims about privacy are often too distant to sustain real engagement. When people can see the data flow, identify the exposure point, and understand what changes in the privacy-preserving version of the workflow, the conversation becomes much more concrete and constructive.

A third finding is that trust is deeply linked to institutional behaviour. The project's outputs suggest that people do not only want better encryption; they want confidence that organizations are using data responsibly, limiting access appropriately, making thoughtful design choices from the outset, and not treating mobility data primarily as a monetizable asset without transparent user benefit. This reinforces one of the project's core conclusions: privacy-enhancing technologies are most powerful when embedded in a culture of accountability and communicated as part of a responsible system design philosophy.

For this reason, the project's contribution to public trust is not limited to the claim that encrypted analytics can technically work. Its broader contribution is to shift the conversation toward what trustworthy connected mobility should look like. That includes selective data use, stronger confidentiality during computation, better communication with stakeholders and the public, and a clearer recognition that privacy is not an optional feature added after deployment. It is part of how legitimate and socially acceptable digital mobility systems must be designed.

Legal status note. PIPEDA remains the current federal private-sector privacy law in Canada. The CPPA concepts discussed in this section are legislative modernization options derived from Bill C-27 and related policy discussion; they are not in force. [2][3]

9.4 Conclusion

Overall, Objective 5 showed that best practices, public communication, and trust-building are essential to the responsible adoption of privacy-preserving connected-vehicle technologies. The project's knowledge translation activities helped make complex technical ideas more accessible, while the stakeholder and public-facing responses reinforced that trust depends on more than security claims alone. It depends on whether privacy protections are understandable, well-governed, and visibly built into the system from the start.

10. Policy and Law Recommendations for Canada (PIPEDA modernization)

Connected vehicles generate a form of personal information that is more continuous, granular, and behaviorally revealing than the kinds of customer data PIPEDA was originally designed around. Vehicle telemetry, location traces, cabin data, app-linked identifiers, diagnostics, biometrics in some systems, and service records can reveal routines, associations, home and work locations, driving style, and inferred risk profiles over time. [2][6][8] Stakeholder input in this project also underscored that these risks extend beyond the registered driver to passengers, bystanders, and other road users whose information may be captured indirectly through connected and sensor-rich mobility systems.

This gap matters because connected-vehicle ecosystems are not simple one-to-one service relationships. Data may move among OEMs, telematics providers, insurers, dealerships, app providers, fleet operators, roadside infrastructure partners, cloud processors, and analytics vendors. In that setting, generic notice-and-consent models are strained, and the OPC has already stressed that meaningful consent should be paired with privacy by design, practical guidance, and organizational accountability. [4][5][8]

A modernized Canadian framework should therefore move beyond generic compliance language and establish clearer statutory obligations for connected vehicles and mobility applications. The goal should not be to create an entirely separate privacy regime for vehicles, but to recognize that mobility data can be highly sensitive even where it is not traditionally categorized as health or financial information. These recommendations respond directly to stakeholder concerns identified in this project: overcollection, fragmented governance, weak transparency, secondary use and monetization without clear user benefit, and unclear accountability across a distributed mobility ecosystem. This direction is also consistent with international connected-vehicle guidance and with the legislative modernization concepts discussed in Bill C-27. [3][6][7]

10.1 Recommended updates

1. Recognize high-sensitivity mobility data in law or binding guidance. PIPEDA already ties safeguards to sensitivity, but connected-vehicle data should be expressly recognized as potentially high sensitivity when it enables precise location tracking, behavioural profiling, inference about household members, or persistent linkage across services. This would help avoid underclassifying telemetry and location data as merely operational. [2][6]

2. Require DPIAs for high-risk connected-vehicle processing. Canada should require a documented privacy impact assessment, or equivalent data protection impact assessment, before organizations deploy high-risk mobility processing. That obligation should apply where processing involves continuous geolocation, in-cabin sensing, profiling for insurance or eligibility, large-scale cross-context data sharing, AI-based inference, or any new use of vehicle data that may materially affect individuals. [5][6]
3. Strengthen purpose limitation and prohibit function creep by default. PIPEDA already states that personal information should not be used or disclosed for purposes other than those for which it was collected, except with consent or as required by law. For connected vehicles, that principle should be sharpened so that data collected for safety, maintenance, diagnostics, navigation, or cybersecurity cannot later be repurposed for unrelated advertising, generalized commercial profiling, or third-party commercialization unless a separate lawful basis and a clear necessity test are satisfied. [2][4][6]
4. Create a specific right of mobility and interoperable transfer. A modernized framework should give individuals a practical ability to direct the secure transfer of their mobility-related personal information between service providers where technically feasible and safe. While PIPEDA provides access rights, the more explicit portability-style concepts proposed in Bill C-27 remain useful modernization reference points even though they are not in force. [2][3]
5. Recognize privacy-enhancing technologies in compliance language. The law should explicitly state that organizations may meet or strengthen their privacy obligations through approved privacy-enhancing approaches, such as strong encryption, secure multiparty methods, de-identification with documented residual-risk analysis, and privacy-preserving analytics that reduce raw-data exposure. [2][9][10]
6. Clarify rules for de-identification, anonymization, and derived inferences. Connected-vehicle data is particularly vulnerable to re-identification because route histories, timing patterns, and device linkages can make supposedly stripped records re-identifiable. Canada should define de-identification and anonymization more clearly and require a documented risk-based assessment before data is treated as sufficiently de-identified. [3][9]
7. Mandate stronger transparency at the vehicle interface, not only in web policies. PIPEDA's openness principle requires information to be understandable and readily available. In practice, connected-vehicle transparency should not be left to lengthy external privacy policies, but should be supported by layered, just-in-time, role-specific notices through the vehicle interface, companion app, onboarding flow, and account dashboard. [2][4][8]
8. Require accountable vendor and ecosystem governance. Because connected-vehicle processing is highly distributed, the law should require lead organizations to maintain documented allocation of roles and responsibilities across manufacturers, service providers, analytics partners, and downstream recipients. Proposed CPPA concepts on

privacy management programs and service-provider arrangements offer a useful accountability template even though they are not yet law. [3][5]

10.2 Practical "regulatory language" suggestions

- **High-sensitivity mobility data**
  The Act should expressly recognize that mobility-related personal information may be highly sensitive where it reveals precise or repeated location, travel patterns, driving behavior, household associations, biometric or in-cabin attributes, or persistent identifiers linked to a vehicle, device, or user account. Organizations processing such information should be subject to enhanced safeguards, stricter purpose specification, and elevated accountability obligations. This recommendation builds on PIPEDA's existing sensitivity-based safeguard model, but makes its application to connected mobility explicit.

- **Mandatory impact assessment for high-risk processing**

  The Act should require organizations to complete and retain a privacy impact assessment before commencing high-risk connected-vehicle processing. High-risk processing should include continuous geolocation tracking, in-cabin monitoring, profiling or scoring with legal or similarly significant effects, large-scale sharing across entities, and AI-enabled inference using mobility data. The assessment should identify purpose, necessity, proportionality, residual risks, safeguards, retention limits, third-party access, and mitigation measures.

- **Purpose limitation and secondary-use control**

  The Act should state that personal information collected through a connected vehicle may be used only for the specific purposes reasonably necessary to deliver the service, maintain safety, ensure security, meet legal obligations, or perform other purposes expressly authorized by law. Secondary uses, including advertising, generalized commercial profiling, unrelated data brokerage, or disclosure to third parties for independent analytics, should require a separate lawful basis and should not be bundled into essential vehicle functionality.

- **Mobility and interoperability right**

  The Act should provide individuals with the right to direct the secure transfer of their personal information from one mobility service provider to another, where both organizations participate in a recognized interoperability framework and where the transfer can be completed securely and without unreasonable burden. This right should apply to

telematics, usage history, subscription records, and other user-linked service data, subject to narrowly tailored exceptions for safety, fraud prevention, and legal privilege.

- **Recognition of privacy-enhancing technologies**

  The Act or accompanying guidance should expressly recognize privacy-enhancing technologies, including encryption and other privacy-preserving computation methods, as valid means of meeting security, minimization, and privacy-by-design obligations where they reduce unnecessary exposure of identifiable information. Recognition should not create a safe harbour by itself, but it should encourage organizations to adopt architectures that reduce raw-data access where practical.

- **De-identification and residual-risk governance**

  The Act should define de-identification and anonymization distinctly, require reasonable technical and organizational measures proportionate to re-identification risk, and prohibit organizations from treating mobility data as outside privacy law merely because direct identifiers were removed. Where de-identified connected-vehicle data is used for internal research, analytics, or system improvement, organizations should document the method used, the residual risk considered, the access controls applied, and the retention period.

- **Layered transparency and user controls**

  The Act should require transparency mechanisms that are accessible at the point of use, including in-vehicle prompts, app-based dashboards, and concise explanations of data categories, purposes, retention periods, recipients, and available controls. Organizations should be required to distinguish data that is necessary for safety or core service provision from data used for optional analytics, personalization, or commercial partnerships.

- **Privacy management program for mobility ecosystems**

  The Act should require organizations engaged in connected-vehicle processing to maintain a privacy management program proportionate to the volume and sensitivity of the information under their control, including documented responsibilities, vendor oversight, training, incident handling, retention schedules, and technical safeguards. Regulators should have authority to review those programs and require corrective measures where deficiencies are identified.

10.3 Conclusion

In summary, the connected-vehicle context exposes a structural mismatch between the generality of PIPEDA and the operational realities of modern mobility data ecosystems. Canada does not need a wholly separate privacy statute for vehicles, but it does need sharper obligations for high-sensitivity mobility data, mandatory risk assessment for high-risk processing, stronger purpose limitation, practical portability rights, and explicit recognition of privacy-enhancing safeguards.

11. Deliverables Checklist and Evidence of Completion

This section presents the project deliverables using the deliverable categories set out in the Contribution Agreement, and maps each to the evidence provided in this final report or in associated project materials.

Where a delivery was completed in integrated form within the report rather than as a stand-alone public document, that is noted explicitly below. During the project period, the demonstration effort was deliberately concentrated on one representative connected-vehicle use case, privacy-preserving usage-based insurance, rather than spread across multiple partially developed scenarios. This refinement improved coherence, made stakeholder interpretation easier, and strengthened the link between the technical prototype, the post-demo assessment, and the policy recommendations.

Deliverable promised in proposal	**Status**	**Evidence of completion / where addressed**
Encrypted-Analytics Prototype: a working proof of concept integrating the homomorphic-encryption method with the neural-network model, demonstrated through a connected-vehicle insurance risk-scoring use case	Completed	Sections **5.1–5.5**, **6.1–6.3**, **7.1–7.3**, **8.1**; **Appendix B**
Benchmark Data Package: anonymized bundle of encrypted logs, GPS traces, simulated distraction metrics, and supporting scripts for latency and accuracy testing	Completed/retained as a technical support package for reproducibility and project validation	Sections **5.4**, **5.5**, **7.1**, **7.2**; **Appendix B**; retained internal technical materials, as applicable
Technical Report & Open-Access Journal Article: comprehensive engineering reporting of the system architecture, security analysis, performance metrics, and implementation lessons, together with companion manuscript outputs	Completed for technical reporting; companion manuscript outputs have been produced, with one open-access journal manuscript submitted and under review, and additional manuscript outputs in preparation. Funding support is acknowledged in the relevant manuscripts.	Present report; Sections **5.1–5.5**, **6.1–6.3**, **7.1–7.3**, **10.1–10.2**; **Appendix B**; project publication/artifact record in **Section 11.1**

Policy Brief: plain-language translation of technical findings into regulatory recommendations relevant to PIPEDA / CPPA modernization and Transport Canada's connected-vehicle context	Completed	**Executive Summary**; Sections **9.1**, **9.2**, **9.3**, **10.1**, **10.2**
National Webinar & Slide-Deck: webinar-based public and stakeholder-facing knowledge translation, with associated slide material	Completed	Sections **9.2**, **9.3**; **Table 9A**; **Appendix D**
Stakeholder Interview Summary: thematically coded analysis of semi-structured interviews with relevant ecosystem participants	Completed	Sections **3.1**, **3.2**, **4.3**, **8.2**, **8.3**; **Table 8A**; **Appendix C**
Final OPC Research Report & Commercialization Road-map: close-out synthesis of technical, qualitative, and outreach work, including next-step pathway toward pilots and broader uptake	Completed	Present document; Section **12**, especially **12.1**, **12.2**, and **12.3**

11.1 Publications/artifacts (submitted, under review, released)

The project's primary reportable artifacts include the final OPC research report, the encrypted analytics prototype and demonstration materials, the benchmark and validation materials used to support technical reproducibility, the explanatory infographics and outreach materials produced during the reporting period, and the de-identified stakeholder-analysis outputs used to support the report narrative.

For the purposes of this final submission, the present document serves as the project's principal integrated technical and close-out report. It brings together the system architecture, threat framing, implementation approach, performance findings, stakeholder-informed assessment, policy recommendations, and future deployment pathway developed during the project. The technical basis for that record is documented primarily in Sections 5 to 7 and summarized further in Appendix B.

In addition to the final report itself, the project generated publication-oriented and dissemination-oriented research materials arising from the technical work. These include prototype validation records, technical benchmark materials, internal documentation used to support reproducibility and interpretation of the reported results, and manuscript-stage research outputs linked to the encrypted analytics and privacy-preserving AI components of the project.

The publication-oriented research outputs arising from the technical work include:

- One completed M.A.Sc. thesis, *Enhancing Data Privacy in Automotive: Advancements in Torus Fully Homomorphic Encryption*, successfully defended in November 2025.
- One journal manuscript submitted to Future Generation Computer Systems as an open-access submission in February 2026.
- One manuscript in final preparation for submission, titled *Efficient Encrypted Decision-Tree Inference in TFHE-Rust Using MRCRT-Based Comparison*. The intended submission venue is IEEE Access.
- One conference paper prepared for submission, titled *Accelerating Homomorphic Comparisons in TFHE-Rust via Mixed-Radix CRT for Cloud Computing Insurance Applications*. This paper is being readied for submission to a relevant conference venue in summer 2026.

Together, these outputs extend the project's technical record beyond the integrated final report while remaining aligned with the Contribution Agreement's reporting and dissemination objectives.

Knowledge-translation artifacts produced during the project include the explanatory infographic set, Privacy Awareness Week campaign graphics, webinar-based engagement materials, and related outreach content summarized in Section 9, listed in Table 9A, and compiled in Appendix D. Taken together, these materials provide evidence that the project delivered not only technical development and stakeholder assessment, but also public-facing translation aligned with the Contribution Agreement.

11.2 What is retained internally

Not all project materials are appropriate for public release or inclusion in the final report. This project included REB-governed stakeholder engagement, qualitative analysis, prototype-stage technical validation, and supporting documentation used to refine the encrypted analytics workflow and interpret the post-technology assessment. Certain records are therefore retained internally in accordance with ethics requirements, confidentiality expectations, and responsible research practice.

Materials retained internally may include de-identified interview transcripts or notes, consent and recruitment records, internal thematic coding and synthesis documents, working records used to prepare stakeholder findings, governance and project-management notes, technical scripts and supporting files associated with benchmark preparation, and other technical working materials used to support the final synthesis but not intended for stand-alone public release.

Overall, the project's documentary record is divided into two categories: materials delivered publicly through the report and related knowledge-translation outputs, and materials retained internally for ethics, accountability, technical support, and research continuity purposes. That division is appropriate to the nature of the project and does not detract from the completeness of the final submission.

12. Conclusions and Future Roadmap

This project has shown that privacy-preserving analytics for connected-vehicle contexts is no longer only a theoretical possibility. Within the scope of the project, the work has demonstrated that homomorphic encryption and oblivious AI can be framed as a viable privacy-by-design pathway for connected-mobility data processing, while also generating meaningful stakeholder and policy discussion around feasibility, integration, and trust. From the beginning, the project was designed not only to advance a technical solution, but also to test whether such a solution could be understood, accepted, and positioned for real-world use in Canadian mobility settings.

12.1 What is proven now

This project proves feasibility in two senses. Technically, the prototype demonstrates an end-to-end encrypted vehicle-cloud-vehicle analytics workflow and establishes a measurable performance and utility trade-off, to be summarized in Table 7A. Socially and institutionally, stakeholder feedback indicates conditional support for adoption when privacy protection is paired with data minimization and clear governance, as summarized in Table 8A. These results support privacy by design as an implementable pathway for selected connected-mobility use cases, while not yet constituting deployment readiness for real-time or high-throughput production settings.

At this stage, the project has established proof of feasibility at the level appropriate to the project scope. The original proposal aimed to validate the integration of homomorphic encryption and oblivious AI as a privacy-by-design solution for connected and autonomous vehicles, and to demonstrate that useful analytics could be performed without exposing raw personal data. The project design centered on an encrypted analytics prototype, with measurable performance indicators such as latency, throughput, and accuracy, and a controlled use case in privacy-preserving usage-based insurance. That use case was selected because it is concrete, policy-relevant, and easy for stakeholders to interpret in terms of both utility and privacy risk.

The revised implementation plan further sharpened that focus by narrowing the end-to-end demonstration to the usage-based insurance scenario while retaining the broader architecture and conceptual relevance to other connected-vehicle applications. This made it possible to preserve the core scientific and policy objectives of the work while aligning the project with the revised budget and timeline. In that sense, what is proven now is not that the technology is deployment-ready across all connected-vehicle functions, but that the proposed privacy-preserving pathway can be implemented, explained, and evaluated in a realistic mobility use case with clear relevance to Canadian privacy concerns.

The project also establishes an important second result: stakeholder acceptance appears plausible when the technology is presented in relation to a concrete privacy problem and a bounded use case. From the proposal stage onward, the project treated stakeholder engagement as essential rather

than secondary. It explicitly aimed to gather feedback on whether automakers, regulators, and related experts could see this type of system fitting into existing data flows, what integration barriers they anticipated, and whether the privacy claims would be regarded as credible after demonstration. In that respect, the project has not only shown technical feasibility but has also advanced the social and institutional feasibility of the concept by linking the prototype to practical, policy, and trust-related concerns.

A careful conclusion, however, is important here. The project supports the claim that the approach is feasible and promising, and that it can be understood by stakeholders as a serious privacy-enhancing option. It does not yet prove that all deployment barriers are resolved, nor that encrypted analytics of this kind can immediately replace existing automotive data-processing pipelines at scale. The appropriate conclusion is that the project has established a credible proof of concept and a validated direction for further technical and translational work.

12.2 Suggestions for Future Work

What remains is the work required to move from a strong research-stage demonstration to a higher technology readiness level. The proposal anticipated that the project would culminate in a roadmap for commercialization, industry-scale adoption, and future scaling, meaning that this next phase was always understood to lie beyond the present one-year project. The most immediate remaining step is TRL uplift through broader pilot validation. The current work supports proof of concept in a controlled setting; the next stage would require integration-oriented pilots with more realistic operational constraints, more diverse telemetry streams, and evaluation under conditions closer to real automotive data pipelines.

A second major remaining task is performance scaling. Although the project was built on a technically promising encrypted analytics framework and on the efficiency advantages associated with the SAHE direction described in the proposal, homomorphic encryption remains computationally expensive. The project itself recognized this by identifying latency, computational overhead, memory requirements, and workflow integration as central constraints. As a result, future work must focus on scaling the architecture to larger workloads, richer models, and more demanding real-time or near-real-time mobility environments. This is especially important if the framework is to move beyond summary-score use cases toward higher-volume continuous analytics.

A third remaining area is pilot deployment and systems integration. The project has shown why privacy-preserving analytics matters and how it could operate, but fuller deployment would require tighter coupling with OEM, supplier, insurer, fleet, or transit data infrastructures. This includes issues such as secure key management at fleet scale, interfaces between vehicle-side preprocessing and cloud-side encrypted computation, integration with existing telematics workflows, and evidence that encrypted processing can coexist with operational expectations around latency,

maintenance, cybersecurity, and auditability. These are less questions of core scientific validity now than of engineering integration and deployment economics.

A fourth remaining task is standardization and regulatory uptake. One of the project's central policy insights is that connected vehicles sit in a gap between generic privacy law and sector-specific mobility practice. The project can therefore inform future guidance, but the translation from recommendation to standard or regulatory instrument remains ahead. Future work should aim to connect these findings to mobility-focused codes of practice, privacy impact assessment expectations, procurement requirements, and possibly future standard-setting discussions around privacy-by-design in connected transportation. This would allow the project's technical contributions to move beyond an isolated prototype and begin shaping repeatable sector practice.

12.3 Next steps and candidate partners/pilots

The next steps should build directly on the structure already established by the project. First, the encrypted analytics pathway should be advanced into a larger pilot centered on the usage-based insurance scenario, since this was the most mature and policy-relevant demonstration track under the revised plan. A logical next step would be a pilot involving richer driving summaries, more diverse operating conditions, and closer engagement with potential industry users who already understand telematics-based risk scoring. This would create the clearest bridge between proof of concept and pre-commercial validation.

Second, the project should expand from one demonstrated use case to a staged multi-use-case roadmap. The revised plan deliberately documented the other two cases conceptually rather than prototyping them in full. That was a reasonable and disciplined decision for this project. Going forward, those additional cases will become natural candidates for the next phase, allowing the framework to be tested across a broader range of mobility functions while keeping the same privacy-by-design core.

Third, partner development should align with the ecosystem identified in the proposal. CUTRIC and APMA remain the most natural candidate convening partners because they connect the work to transit, suppliers, automakers, and applied innovation communities, and because the proposal explicitly positioned them as channels for industry alignment, dissemination, and practical validation. Transport Canada remains an important policy-facing partner for alignment with national safety and connected-vehicle priorities. In addition, HLP Innovative Solutions was included in the project as a qualitative and translational support partner and could continue to play a role in future usability, adoption, and policy uptake studies. If the commercialization potential continues to strengthen, the proposal also identified IP Ontario as a next-step partner for an IP landscape review and the development of a commercialization strategy.

Fourth, future pilots should be designed not only as technical exercises but as integrated translational pilots. That means each pilot should include technical benchmarking, stakeholder usability assessment, governance analysis, and public- or policy-facing communication outputs. This is consistent with the project's character as a whole, which was never intended to be only a cryptography exercise. Its strength lies precisely in combining technical innovation with policy relevance, adoption analysis, and knowledge translation. The next phase should preserve that model.

In conclusion, this project has shown that privacy-preserving connected-vehicle analytics can be moved from concept toward credible implementation. The work has demonstrated the feasibility of an HE- and oblivious-AI-based pathway in a focused mobility use case, while also generating stakeholder-informed insight into trust, adoption, and policy relevance. The next phase is not to re-prove the concept, but to raise its technology readiness through larger pilots, broader partner engagement, performance scaling, and gradual integration into mobility practice and privacy governance frameworks.

References

[1] Office of the Privacy Commissioner of Canada. Applicant's Guide: Contributions Program 2026-2027. OPC, 2025.

[2] Personal Information Protection and Electronic Documents Act, S.C. 2000, c. 5 (Can.). Justice Laws Website, current consolidated version.

[3] Parliament of Canada. Bill C-27, Digital Charter Implementation Act, 2022: LEGISinfo status page.

[4] Office of the Privacy Commissioner of Canada. Guidelines for obtaining meaningful consent. OPC guidance.

[5] Office of the Privacy Commissioner of Canada. Accountability under PIPEDA. OPC guidance.

[6] European Data Protection Board. Guidelines 01/2020 on processing personal data in the context of connected vehicles and mobility-related applications. Version 2.0, 2021.

[7] OECD. OECD Privacy Guidelines / Privacy Principles. OECD privacy and data protection framework materials.

[8] Office of the Privacy Commissioner of Canada. A Privacy Code Practice for the Connected Car. Completed Contributions Program project page, 2019.

[9] Office of the Privacy Commissioner of Canada. Reducing identifiability in cross-national perspective: statutory and policy definitions for anonymisation, pseudonymisation, and de-identification in G7 jurisdictions, 2024.

[10] Cheon J.H., Kim A., Kim M., and Song Y. Homomorphic Encryption for Arithmetic of Approximate Numbers. ASIACRYPT 2017.

[11] Microsoft SEAL (Simple Encrypted Arithmetic Library). Microsoft Research / GitHub documentation repository.

[12] Juvekar C., Vaikuntanathan V., and Chandrakasan A. GAZELLE: A Low Latency Framework for Secure Neural Network Inference. USENIX Security Symposium, 2018.

[13] Gilad-Bachrach R. et al. CryptoNets: Applying Neural Networks to Encrypted Data with High Throughput and Accuracy. ICML 2016.

[14] Kheiran, P. Enhancing Data Privacy in Automotive: Advancements in Torus Fully Homomorphic Encryption (TFHE). MASc thesis, University of Windsor, 2025.

Appendix A - REB approval letter + consent template + interview guide

University of Windsor

December 02, 2025

To: **Office of Research Innovation Services**
From: **Hannah Sands (ethics@uwindsor.ca)**
Re: **Human Ethics REB Certification**

To whom it may concern,

Please accept this email as confirmation that **Dr. Mitra Mirhassani** has been issued a **Human Ethics** certificate. This certificate is valid until December 02, 2026.

Therefore, please release all funds, if held, related to the following award(s):

Awards File No	Title	Status	
44733	Driving Privacy Forward: Homomorphic Encryption and Oblivious AI in Smart Mobility	Active	Office of Research and Innovation Services

Appendix B - Technical parameter tables and prototype summary

This appendix summarizes the thesis-linked technical record that informed Sections 5 to 7 of the report. It is intentionally concise and report-oriented: the purpose is to anchor the final submission with reproducibility-relevant technical facts without turning the OPC report into a thesis document. [14]

Prototype scope and trust boundary: The validated prototype focused on privacy-preserving insurance and telematics analytics using encrypted decision-tree inference over selected structured driving features. The operational logic follows the vehicle-cloud-vehicle pathway described in the report: selected features are encoded and encrypted at the trusted edge, the cloud evaluates encrypted comparison and decision logic, and only an authorized endpoint decrypts the permitted result. In the thesis-linked implementation, the cloud operates as a ciphertext processor rather than a plaintext analytics provider, and secret keys remain outside the cloud trust boundary. [14]

Implementation summary: The technical implementation follows a TFHE-style fully homomorphic encryption workflow, implemented in TFHE-rust, with an MRCRT-accelerated comparator used to reduce the sequential depth of comparison-heavy inference. This design choice is especially relevant for decision-tree style workloads, where encrypted threshold checks and selectors dominate runtime. The technical record therefore supports the reports position that the main bottleneck lies in encrypted comparison and evaluation rather than in the idea of protected transmission or authorized decryption itself. [14]

Prototype results summary: In the thesis-linked implementation, 64-bit encrypted comparison latency decreased from 91.16 ms in the baseline CPU bitwise configuration to 67.12 ms in the optimized CPU MRCRT configuration, and to 41.83 ms when GPU scheduling was used. In the deeper encrypted decision-tree setting, end-to-end runtime improved from 2.12 s to 1.16 s. In the literature-style benchmark configuration described in the thesis, the proposed method reached 67.4 ms compared with 86.7 ms reported by a recent reference system. These results were used in the report as evidence that the framework is technically credible for bounded analytics and that its dominant cost remains encrypted model evaluation rather than edge-side encryption or output recovery. [14]

Utility preservation: The thesis-linked validation also supports the reports claim that privacy-preserving evaluation did not collapse task utility. The decision-tree baseline operated at approximately 98.8% accuracy, with precision, recall, F1, and AUC all reported near 0.99 in the validated configuration. The practical implication for the OPC report is that the main visible trade-off in this prototype was runtime overhead, not loss of core classification validity. [14]

Limits of interpretation: These technical results should be read as prototype-stage validation for a focused, bounded service scenario rather than as proof of immediate deployment readiness across all connected-vehicle workloads. The technical record remains strongest for comparison-heavy,

higher-sensitivity analytics where the privacy benefit of keeping data encrypted during processing is substantial and where some computational overhead can be tolerated. This is consistent with the reports main conclusion that privacy-preserving connected-vehicle analytics is now credible for selective, risk-based deployment, but is not yet a universal solution for strict low-latency or control-loop applications. [14]

Appendix C- De-identified quote bank

This appendix provides selected de-identified excerpts from stakeholder interviews conducted during the project. The quotations are included to illustrate key themes discussed in Sections 4, 8, and 9 of the report. They are not presented as a comprehensive transcript record, and they should be read as qualitative supporting evidence rather than as stand-alone findings. Minor edits have been made for clarity and brevity, while preserving the original meaning. Identifying details and organization-specific references have been removed.

C.1 Data minimization and over-collection

Participant 1

"The first issue is not whether we can protect the data after collection. The first issue is whether all of that data needed to be collected in the first place."

Participant 2

"A lot of mobility data looks harmless field by field, but once you combine time, location, and patterns, it becomes much more revealing than people expect."

Participant 3

"There is a tendency to collect broadly because storage and analytics are available, but from a privacy perspective that creates exposure long before a breach or misuse happens."

C.2 Governance, accountability, and trust

Participant 4

"The challenge is not only the technology. It is also who is responsible once the data moves across several parties and service layers."

Participant 5

"You can have strong technical protection and still lose trust if people do not know who holds the keys, who can access the outputs, and what the retention rules are."

Participant 6

"Privacy by design only becomes meaningful when there is accountability behind it. Otherwise, it stays at the level of a good principle."

C.3 Adoption barriers and implementation concerns

Participant 1

“The idea is promising, but organizations will ask how difficult it is to integrate into what they already have.”

Participant 3

“If the overhead is too high, or if the workflow becomes too hard to maintain, many groups will hesitate unless the use case is especially sensitive.”

Participant 5

“Standards alignment matters. People will be more comfortable adopting this if it can be mapped to regulatory expectations and existing governance processes.”

C.4 Value of privacy-preserving design in a bounded use case

Participant 2

“What made the demonstration helpful was that it showed a specific service outcome without requiring routine exposure of the underlying raw data.”

Participant 4

“That is where it becomes easier to explain: you are still getting a useful result, but you are reducing how much anyone in the chain actually sees.”

Participant 6

“I do not see this as a blanket answer for everything, but for higher-sensitivity use cases it makes the privacy argument much more credible.”

Closing note

Taken together, these excerpts reinforce the report’s central conclusion: stakeholders were generally supportive of privacy-preserving approaches when they were explained through a concrete workflow, but they consistently tied that support to data minimization, governance clarity, accountability, and realistic implementation pathways.

Appendix D: Knowledge-translation artifacts (infographics, slide decks, etc.)

This appendix compiles selected infographic materials prepared to communicate connected-vehicle privacy concepts in an accessible, public-facing format. The figures are included as knowledge-translation outputs that supported outreach, awareness building, and stakeholder engagement during the project. They are illustrative communication artifacts and should be interpreted as dissemination materials rather than as technical validation figures.

Contents include principle-based infographics, service-oriented examples, risk-awareness visuals, bilingual adaptation material, and Privacy Awareness Week campaign graphics.

Core infographic materials

These figures were developed to explain connected-vehicle privacy concepts in simple visual terms and to support broad, non-technical communication.

Figure D1. Privacy by Design Core Principles

Privacy by Design core principles

Explanation. This infographic summarizes three core privacy-by-design practices in a compact visual form: transparency, meaningful user choice, and data minimization combined with protection. It works well as a high-level knowledge-translation figure because it clearly conveys the central message without relying on technical language.

Figure D2. Privacy by Design explained through everyday service examples

Privacy by Design explained through everyday service examples

Explanation. This figure translates privacy-by-design concepts into familiar mobility examples, including an app that asks only for necessary data, a rewards program that supports real consent choices, and a dashcam configured with safer defaults. It is useful for public and stakeholder audiences because it connects abstract privacy principles to recognizable services.

Figure D3. From tracking risk to privacy-protective design

From tracking risk to privacy-protective design

Explanation. This comparison-style infographic contrasts always-on tracking with a privacy-by-design approach that emphasizes clear notice, meaningful choices, and reduced data collection. It helps explain why privacy protections matter in connected-vehicle environments and how design decisions can reduce avoidable exposure.

Figure D4. Privacy by Design as a vehicle data filter

Privacy by Design as a vehicle data filter

Explanation. This simplified conceptual figure presents privacy by design as a filtering step between raw vehicle data and permitted outputs. It reinforces the project message that only essential data should flow to downstream services and functions.

Figure D5. Public-awareness infographic on key connected-vehicle privacy risks

Public-awareness infographic on key connected-vehicle privacy risks

Explanation. This public-facing infographic highlights four intuitive privacy themes: data collection, hacking risk, sale or onward sharing of data, and movement tracking. It supports awareness-building by framing the problem in accessible language for non-technical audiences.

Bilingual outreach material

This subsection documents selected bilingual or French-language adaptations prepared to support wider accessibility in public-facing communication.

Figure D6. French-language adaptation of the connected-vehicle privacy risk infographic

French-language adaptation of the connected-vehicle privacy risk infographic

Explanation. This French-language version adapts the public-awareness risk infographic for bilingual outreach and broader accessibility. It demonstrates that project dissemination materials were designed for audiences beyond a single-language communication channel.

Privacy Awareness Week materials

These campaign graphics were prepared as concise public-education assets aligned with the project's knowledge-translation and awareness objectives.

Figure D7. Privacy Awareness Week - Day 1: what data a vehicle may share

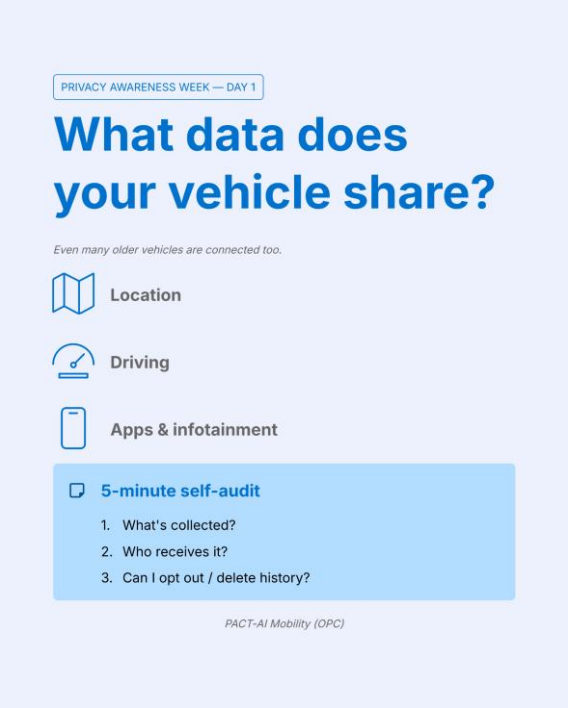

Privacy Awareness Week - Day 1: what data a vehicle may share

Explanation. This campaign graphic encourages readers to identify the types of connected-vehicle data that may be collected, including location, driving information, and app or infotainment data. It also includes a brief self-audit prompt that supports practical public engagement.

Figure D8. Privacy Awareness Week - Day 2: follow the data path

Privacy Awareness Week - Day 2: follow the data path

Explanation. This visual shifts attention from collection alone to downstream data flow, showing that vehicle information may pass through OEM services, third parties, maintenance channels, and optional insurance pathways. It reflects one of the project's recurring messages: privacy assessment must consider the full data lifecycle.

Figure D9. Privacy Awareness Week - Day 3: Encrypted in transit is not the same as end-to-end privacy

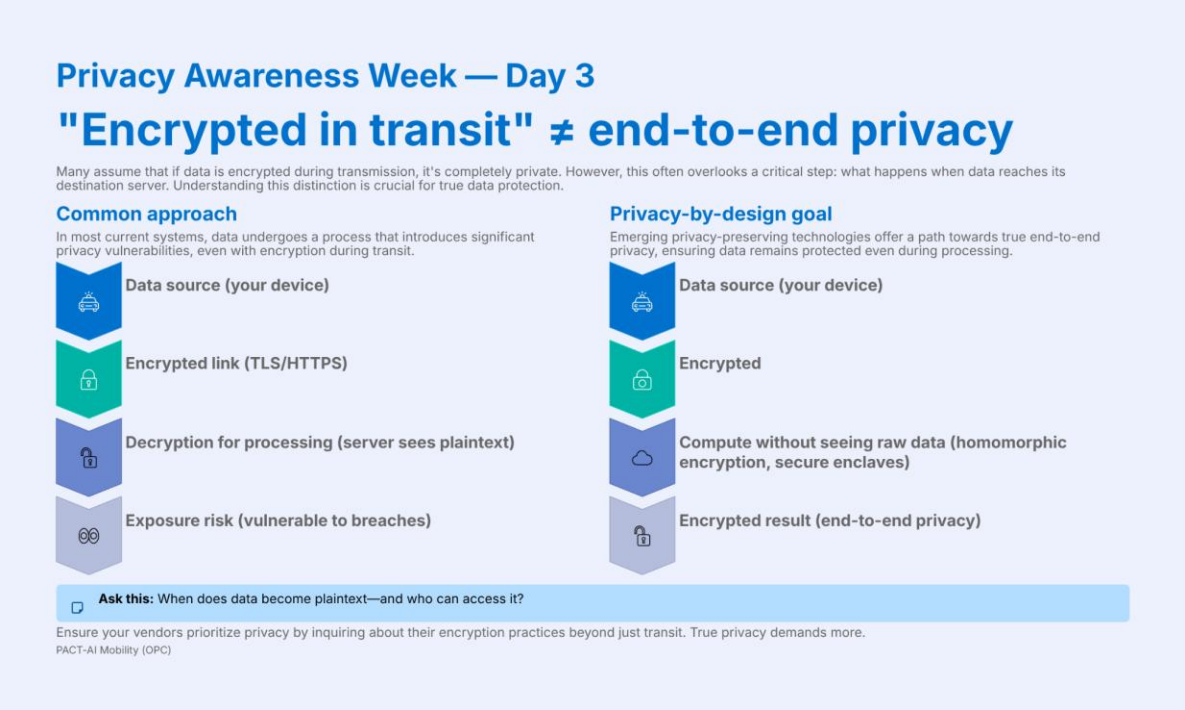

Privacy Awareness Week - Day 3: encrypted in transit is not the same as end-to-end privacy

Explanation. This figure explains an important privacy distinction for non-specialist audiences by contrasting routine encrypted transmission with stronger privacy-preserving approaches that avoid exposing raw data during processing. It is especially valuable as a bridge between public education and the project's technical direction.

Figure D10. Privacy Awareness Week closing graphic with three project takeaways

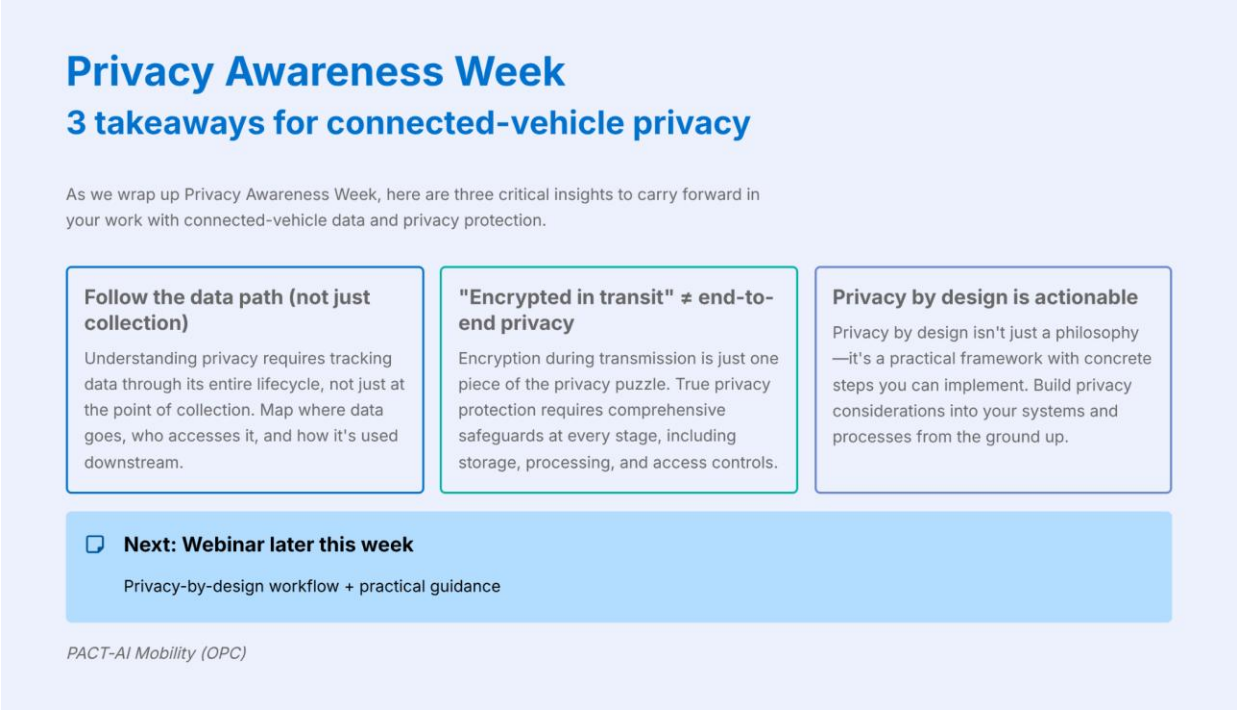

Privacy Awareness Week closing graphic with three project takeaways

Explanation. This closing summary consolidates three practical messages from the campaign: follow the data path, do not equate transport encryption with full privacy, and treat privacy by design as an implementable framework. It serves as a concise record of the project's outward-facing knowledge-translation messaging.