V2X Service Provisioning with 5G V2N2V Communications with Cross-Stakeholder Information Sharing

B. Coll-Perales[1], M.C. Lucas-Estañ[1], O. Altintas[2], J. Gozalvez[1], M. I. Khan[2], M. Sepulcre[1]

[1]Uwicore laboratory, Universidad Miguel Hernandez de Elche, Elche (Alicante), Spain.

[2]InfoTech Labs, Toyota Motor North America R&D, Mountain View, CA, U.S.A.

{bcoll, m.lucas, j.gozalvez, msepulcre}@umh.es; {onur.altintas, mohammad.irfan.khan, sergei.avedisov}@toyota.com

***Abstract*— The support of Cooperative Intelligent Transport Systems (C-ITS) services requires seamless interoperability between involved stakeholders. To this aim, the 5G Automotive Association has recently endorsed a Vehicle-to-Network-to-Everything (V2N2X) architecture trialed at national initiatives to support road traffic management V2X services. The architecture enables interoperability at the application level through a cloud-federated Information Sharing Domain (ISD) that supports data sharing and interoperability among stakeholders. This study analyses the possibility to support critical and latency-sensitive V2X services using 5G-based Vehicle-to-Network-to-Vehicle (V2N2V) communications over the federated cloud-based V2N2X architecture. The analysis considers the intersection collision avoidance (ICA) service as a case study and scenarios involving multiple Mobile Network Operators (MNOs) and Original Equipment Manufacturer (OEM) clouds. We show that the ICA requirements can be supported, provided connections with controlled latencies (under Service Level Agreements or SLAs) are established between the OEM clouds and the ISD. However, the small tolerance to latency variations can compromise the support of the critical and latency-sensitive V2X services over the federated cloud-based V2N2X architecture, and solutions are necessary to ensure the scalability of the system.**

Keywords—5G, cloud, C-ITS, E2E latency, Information Sharing, ISD, V2N2V, V2N2X architecture, V2X

I. INTRODUCTION

Cooperative Intelligent Transport Systems (C-ITS) services rely on the exchange of information among various stakeholders. These stakeholders can include Service Providers (SP) and their clients (normally on smartphone apps), infrastructure owner-operators (IOO, e.g. road operator or authority) and their connected infrastructure, and original equipment manufacturer (OEM) and their connected vehicles. 5GAA has recently published a whitepaper [1] supporting a V2N2X (Vehicle-to-Network Communication to Everything Communication) architecture for C-ITS service provisioning that enables interoperability among stakeholders at the application level through a federated Information Sharing Domain (ISD) on the cloud. This architecture interconnects all stakeholders, and facilitates data sharing and interoperability, resilience and scalability compared to centralized solutions, e.g. the so-called national access points (NAP) trialed in Europe. The technical and operational feasibility of this cloudified V2N2X architecture has been piloted in different projects and national initiatives (e.g. NordicWay, Mobildata and Talking Traffic) for the support of road traffic management V2X services (e.g. traffic light information, speed advice, lane information and parking space availability) and it is unclear if this architecture could also support critical and latency-sensitive V2X services. These services have been traditionally designed considering sidelink or direct V2V (Vehicle to Vehicle) connections. However, the authors demonstrated in [2] that critical V2X services could also be supported with 5G network-based Vehicle-to-Network-to-Vehicle (V2N2V) communications under certain deployment conditions. This study considered the case of a single OEM and multiple MNOs connected directly or through the Internet. However, 5G-based V2N2V communications can also be supported over the federated and cloud-based V2N2X architecture presented in [1] that provides a more holistic framework to support a large range of C-ITS or V2X-based services.

This work was supported in part by MICIU/AEI/10.13039 /501100011033 a (grant PID2020-115576RB-I00), by the "European Union NextGenerationEU/PRTR" (TED2021-130436B-I00), by Generalitat Valenciana, and UMH's Vicerrectorado de Investigación.

In this context, this study progresses the current state-of-the-art by assessing for the first time the potential of the V2N2X architecture with a federated Information Sharing Domain (ISD) on the cloud to support latency-sensitive and safety-critical V2X services using 5G-based V2N2V communications. We use the Intersection Collision Avoidance (ICA) service as a case study and consider various deployment scenarios involving multiple MNOs and OEM clouds with different type of connections between the clouds and the ISD. The study shows that it is possible to support the ICA service with 5G network-based V2N2V communications utilizing the federated and cloudified V2N2X architecture with ISD, provided that controlled connections with Service Level Agreements (SLA) are established between the OEM clouds and the ISD. However, the capacity to support this type of critical V2X services can be compromised by variations in conditions that impact any of the latency components in the E2E V2N2V connection. This is illustrated in this study by analyzing the impact of the scalability of the message queuing protocol utilized in the ISD for exchanging messages between cloud nodes and IS instances.

II. V2N2X ARCHITECTURE

The V2N2X architecture includes an ISD in the cloud that interconnects all stakeholders at the application level and serves as a dedicated trust domain for V2X data sharing and interoperability. The ISD follows a federated system architecture with loosely coupled information sharing (IS) instances and aims to support:

- Independence: IS instances can operate or be modified without affecting others. They interact through well-defined interfaces and their internal operations are isolated.
- Flexibility: changes to one IS instance are less likely to require changes in another instance. This allows for easier maintenance, updates, and enhancements over time.
- Scalability: it is possible to add new IS instances without disrupting the entire system.
- Resilience: failures can be isolated (independence) and IS instances can back each other up.

The ISD facilitates data sharing and interoperability between stakeholders through well-defined interfaces (see Fig. 1) for data and control planes specified by C-ROADs [3]. The Basic Interface (BI) is used for real-time data exchange of C-ITS or standardized V2X application messages (e.g. -Decentralized Environmental Notification Message -DENMs)

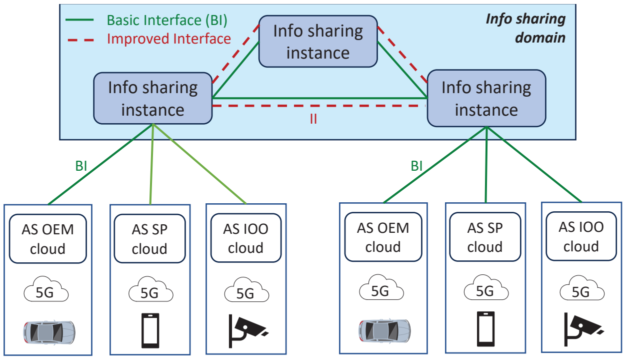

Fig. 1. V2N2X architecture with federated ISD in the cloud

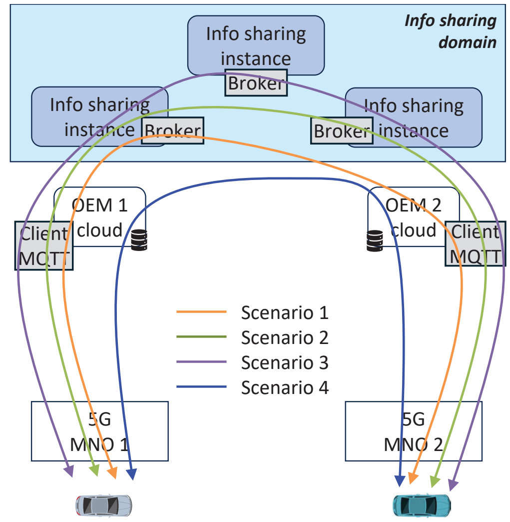

Fig. 2. Deployment scenarios

in the backend, i.e., between stakeholders' clouds and IS instances and between IS instances (see Fig. 1). The BI encapsulates the V2X application messages as payload, and adds meta-data as headers. The meta-data can include, for example, the payload's message type, location of the event related to the message (e.g. latitude/longitude, or area identified by quadtree tiles), and the originator of the message. The meta-data can be used to process and filter messages without accessing the payload (that could be signed), and select those that should be shared with each stakeholder connected to the ISD. This transforms the ISD into a dedicated trust-domain for data sharing across stakeholders. The Improved Interface (II) provides control plane support for data sharing over the BI. It automates ISD data access and sharing and federates the exchange of information. The II operation allows ISD clients (e.g. OEM clouds) to interact with a single IS instance, while providing access to all data available in the ISD (whether managed by the same or different IS instances). IIs are established between IS instances, and are utilized to request data and redirect requests on behalf of the ISD client they support through the entire ISD domain.

Messaging queuing protocols (e.g., Advanced Message Queuing Protocol -AMQP- or Message Queuing Telemetry Transport -MQTT- as advocated by 5GAA and C-ROADS) can be used for efficient and scalable data-sharing with(in) the ISD. These protocols rely on a broker server that distributes the information to the interested clients connected to the server. The information is organized in topics. Clients publish the information through the broker using topics, and receive the topics they are subscribed to through the broker. The messaging queuing protocols can use the meta-data added on the BI interface to manage the subscription and dissemination of messages.

III. Deployment Scenarios

The V2N2X architecture has been used so far for road traffic management services where the vehicles or (smartphone) clients interact with the connected infrastructure. However, the architecture can also support V2N2V communications between vehicles of the same or different OEMs. In the latter case, each vehicle communicates with its OEM cloud, and OEMs' clouds are interconnected through IS instances in the ISD. The support of 5G-based V2N2V communications over the federated and cloud-based V2N2X architecture is represented in Fig. 2. The Fig. 2 represents the scenario where each vehicle communicates with its OEM cloud through a different MNO, but the architecture could also support the communications between vehicles operating over the same MNO. The integration depicted in Fig. 2 is not influenced by the specific 5G MNO network deployments [2] since the 5G network is connected to the OEMs' clouds through an Internet connection from the 5G core network's User Plane Function (UPF).

Fig. 2 represents different V2N2V communication scenarios over the V2N2X architecture:

1) Scenario 1: the involved OEM clouds are served by the same IS instance. The IS instance hosts the message broker to which clients running on the OEMs' clouds would publish to and subscribe to for the message exchange.
2) Scenario 2: OEM clouds are each supported by a different IS instance, and each instance has its own broker. They must then publish/subscribe to the broker running on the IS instance they are connected to, and each IS instance subscribes/publishes to the other IS instance (over the II interface) on behalf of the OEM cloud they provide support to.
3) Scenario 3: OEM clouds are supported by different IS instances and these instances are connected through another IS instance. Each instance has its own broker. This scenario adds to the previous case the publication/subscription processes from the IS instances that are directly connected to the OEM clouds to the IS instance that interconnects them.
4) Scenario 4: OEMs clouds are directly connected without using the ISD. This scenario requires the establishment of interfaces and agreements between OEM clouds.

IV. 5G V2N2V E2E Latency Model

We utilize and extend the 5G E2E (end-to-end) latency model l_{E2E} introduced by the authors in [4] to assess the support of latency sensitive V2X services over the different deployment scenarios presented in Section III.

The latency model introduced in [4] for centralized 5G network deployments (l_{5Gcent}) accounts for the latency experienced at the radio (l_{radio}), transport (l_{TN}) and core (l_{CN}) networks, as well as the latency generated by the Internet connection between the CN's UPF node and the OEM cloud that hosts the V2X Application Server ($l_{UPF\text{-}AS}$). The model derived in [4] considers multi-MNO and single OEM scenarios. In multi-OEM scenarios, each vehicle communicates via its OEM V2X AS and V2X packets are processed by each V2X AS involved in the V2N2V connection. This adds an additional latency component (l_{AS}) that we estimate following [5] and assuming that the V2X ASs only forward the received packets. The processing power of the ASs is dimensioned to avoid backlogging of packets at the AS queues [4].

The 5G E2E latency model is further extended to account for the interactions between the OEMs V2X AS through the IS instances of the ISD. We consider these interactions are based on publish-subscribe messaging protocols [1]. We have computed the latency introduced by the publish-subscribe messaging protocols (l_{PubSub}) through experimental tests using the MQTT protocol. For modelling scenarios with controlled connections (with SLAs) between the OEM clouds and the IS instances, the MQTT clients and broker used in the tests are deployed in the same node. We then experimentally quantify the latency involved in the local publish-subscribe processes (l_{PubSub_local}), and add the latency that would be introduced in the interconnection links between MQTT clients and the broker. These links are called peering points, and their latencies are modeled based on the empirical study reported in [6]. We consider local (or private) peering points (with latency l_{pp_local}) for modelling controlled connections with SLA on the links between OEM clouds and IS instances. When such connections are not controlled through SLAs, the links are modeled as remote (or public) peering points (l_{pp_remote}). For this scenario, the MQTT clients connect through the Internet to a broker deployed on the cloud, and we experimentally quantify the latency experienced in the remote publish-subscribe processes (l_{PubSub_remote}). This latency component also includes the latency introduced by the non-controlled remote peering point connection through the Internet.

Eq (1) reports the total end-to-end latency l_{E2E} for the V2N2V connection over the first 3 deployment scenarios reported in Section III (i.e. those for which the OEM clouds are interconnected through the ISD) when the connections between OEM clouds and IS instances and between IS instances are based on controlled connections with SLA.

$$l_{E2E} = l_{5Gcent} + 2l_{AS} + (x+1)l_{pp\_local} + xl_{PubSub\_local} \quad (1)$$

where l_{5Gcent} is the latency over a 5G centralized network, $2l_{AS}$ is the latency added by the processing of the packets at the two OEM clouds (in particular, at their V2X AS servers), $(x+1)l_{pp\_local}$ is the latency introduced by the local peering point connections used to exchange messages, and xl_{PubSub_local} is the latency introduced by the publish-subscribe processes. The three scenarios differ on the number of publish-subscribe messages exchanged between OEM clouds and IS instances, and therefore on the value of x in eq (1). In scenario 1, all OEMs clouds are supported by the same IS instance, and when a OEM cloud client publishes a message on a topic, the other OEM cloud client receives it via the broker running on the same IS instance. In this case, x in eq (1) is equal to 1. In scenario 2 and 3, x in eq (1) is equal to 2 and 3 respectively.

Eq (2) computes the total end-to-end latency l_{E2E} for the first 3 deployment scenarios, but for the case when the connections between OEM clouds and IS instances, and between IS instances, are not based on controlled connections with SLA. In this case, x={1, 2, 3} for scenarios 1), 2) and 3), respectively. Please note that l_{PubSub_remote} accounts not only for the latency created by the publish-subscribe processes but also the latency introduced by the remote connection through the Internet.

$$l_{E2E} = l_{5Gcent} + 2l_{AS} + xl_{PubSub\_remote} \quad (2)$$

Scenario 4 considers a direct connection between OEM clouds. Eqs. (3) and (4) report the E2E latency for this scenario with local or remote peering point connections between the clouds depending on whether controlled connections with SLA are established or not, respectively.

$$l_{E2E} = l_{5Gcent} + 2l_{AS} + l_{pp\_local} \quad (3)$$

$$l_{E2E} = l_{5Gcent} + 2l_{AS} + l_{pp\_remote} \quad (4)$$

V. E2E Latency Analysis

A. Scenario

We evaluate the possibility to support critical V2X services with 5G-based V2N2V communications over the V2N2X architecture with ISD considering the ICA service defined in 3GPP TR37.885 and TR36.885. Following those, vehicles (from different OEMs) exchange 350-byte Cooperative Awareness Messages (CAM) while approaching a 4-way intersection served by a gNB with 300m cell radius. Following 5GAA guidelines in [7], the ICA service requirements for the exchange of CAM messages include a 100 ms latency budget with 99.99% reliability. We consider two vehicle densities to analyze the impact of variable network loads. In the first setting, vehicles drive at 50 km/h and exchange CAM messages at 3.47Hz according to ETSI TS 101-539-2. In the second setting, vehicles drive at 15 km/h and exchange CAM messages at 2.2Hz according to real traces from C2C-CC [8]. The gNB receives 506 pkt/s and 779 pkt/s, respectively, for the two vehicle density settings.

We adopt the hierarchical network topology recommended by ITU (analyzed in [4]) and consider the same network load in all gNBs. The ITU's hierarchical network topology consists of 1728 gNBs served by 3 levels of multiplexing nodes (M1, M2, M3) at the TN before the V2X traffic is passed to the 5G's CN. The network is configured with distances of 3, 12, 60 and 200 km for the links gNB-M1, M1-M2, M2-M3 and M3-UPF (that connects to the Internet), respectively. The network is configured with link capacities of 10, 300, 6000 and 6000 Gb/s for the links gNB-M1, M1-M2, M2-M3 and M3-UPF, respectively. We dimension the network and determine the fraction (α) of the link capacities that should be allocated to support the V2X CAM traffic in order to avoid backlogs at the TN and CN nodes following the methodology in [4] and network planning practices described in [9] for the highest considered network load (i.e., 779 pkts/s received at each gNB). The dimensioning results in an α equal to 1.125%.

B. Latency evaluation

Table I.a reports the average and 99.99th percentile of the round-trip latency for each of the 5G components in the latency model[1] as well as the peering point connections. A range of latency values are shown when applicable for the lowest and highest network traffic loads (i.e., 506 pkts/s/gNB and 779 pkts/s/gNB). Tables I.b-d also report the experimental latency values measured for the publish-subscribe messaging protocol. Tests were conducted for low, medium and high messaging queueing loads to analyze the impact of the number of messages that a broker processes on the l_{PubSub}[2].

[1] The radio network latency l_{radio} is calculated following [4] and considering a common FDD reference configuration with Sub-Carrier Spacing (SCS) of 30 kHz and a cell bandwidth of 20 MHz.

[2] Our {low, medium. high} load settings consider {1, 10, 100} MQTT clients publish 350-byte messages at 20 Hz, and 1 MQTT subscriber receives them through the MQTT broker. Each client publishes more than 1000 messages.

TABLE I. LATENCY IN MS FOR THE DIFFERENT E2E LINK COMPONENTS

a) 5G network latency components

Link	Average	99.99th pctile
l_{radio}	1.502 – 1.502	2.745 – 2.782
l_{TN}	2.851 – 2.937	7.162 – 7.542
l_{CN}	2.001 – 2.001	2.005 – 2.006
$l_{UPF\text{-}AS}$	10.3	43
l_{AS}	0.5	0.75
l_{pp_local}	0.306	1.493
l_{pp_remote}	13.001	99.212

b) Messaging queuing protocol latency – low load

l_{PubSub_local}	1.18	12.42
l_{PubSub_remote}	26.24	235.1

c) Messaging queuing protocol latency – medium load

l_{PubSub_local}	1.636	42.0
l_{PubSub_remote}	35.572	608.03

d) Messaging queuing protocol latency – high load

l_{PubSub_local}	4.823	169.002
l_{PubSub_remote}	49.175	892.0

Table II.a reports the E2E latency experienced in the 4 deployment scenarios discussed in Sections III and IV. Table II.a reports the average and 99.99th percentile of l_{E2E}, considering whether controlled connections (CC) with SLA are established between OEM clouds and IS instances, or not (Non-Controlled Connection – NCC). Table II.a considers the lowest load setting for the messaging queuing protocol (Table I.b). Supporting the ICA service requires vehicles to exchange 99.99% of CAM messages in less than 100ms (total E2E latency). The capacity to support or not the ICA service is then analyzed considering the percentile values in Table II. Table II.a shows that it is not possible to support the ICA service (i.e. the 99.99th percentile of l_{E2E} is above 100ms) when there are no controlled connections with SLA between the cloud elements (i.e., NCC); this is highlighted with red colored cells. This is the case even for scenario 4) in which OEM clouds connect directly to each other without the support of the ISD using a remote peering point link (l_{pp_remote}). Table II.a shows that the ICA requirements can be fulfilled only when CC are established between OEM clouds and IS instances; highlighted with green colored cells. The orange-colored cell highlights the case in which the ICA requirements are fulfilled only for the lowest vehicle density analyzed (see Section V.A).

Table II.a shows that the ICA requirements can only be fulfilled with CC between the OEM clouds and the IS instances. However, the E2E latency values are close to the latency limit established by the ICA requirements, and therefore have a low tolerance to potential latency variations in any of the components of the E2E V2N2V connection. In addition, we should note that the results reported in Table II.a have been obtained with the lowest load in the messaging queuing protocol. Table II.b reports the impact of varying such load on the E2E latency results for the scenarios 1-3 when CC are established between OEM clouds and IS instances[3]. Table II.b shows that an increase in the messaging queuing load can compromise the capacity to fulfill the ICA requirements even with CC connections between OEM clouds and IS instances due to the increase of the publish-subscribe latency reported in Tables I.b and I.c.

TABLE II. V2N2V E2E LATENCY OVER THE V2N2X ARCHITECTURE WITH ISD ON THE CLOUD (IN MS)

a) Messaging protocol– low load

Scenario	$\overline{l_{E2E}}$		99.99th pctl l_{E2E}	
	CC	NCC	CC	NCC
Scenario 1: all OEM clouds supported by the same IS instance	19.39 - 19.50	43.84 - 43.94	71.82 - 72.19	291.51 - 291.89
Scenario 2: OEM clouds supported by different IS instances	20.88 - 20.98	70.08 - 70.18	85.73 - 86.11	526.61 - 526.99
Scenario 3: OEM clouds supported by different IS instances which are connected through another IS instance	22.37 - 22.47	96.32 - 96.42	99.64 - 100.02	761.71 - 762.09
Scenario 4: Direct interaction between OEM clouds	17.91 - 18.01	30.60 - 30.70	57.90 - 58.28	155.62 - 156.0

b) CC & Messaging protocol - medium load and high load

Scenario	$\overline{l_{E2E}}$		99.99th pctl l_{E2E}	
	Medium load	High load	Medium load	High load
1) All OEM clouds supported by the same IS instance	19.9 - 20.01	23.09 - 23.17	101.4 - 101.8	228.40 - 228.82
2) OEM clouds supported by different IS instances	21.84 - 21.93	28.22 - 28.34	144.89 - 145.31	398.89 - 399.31
3) OEM clouds supported by different IS instances which are connected through another IS instance	23.78 - 23.87	33.35 - 33.43	188.38 - 188.80	569.39 - 539.81

VI. CONCLUSIONS

We have evaluated the feasibility of the 5GAA's recommended V2N2X architecture with ISD on the cloud for the support of the latency-sensitive and safety-critical ICA V2X service in multi-OEM scenarios. Our analysis has shown that the 5G-based V2N2V communications relying on this architecture can meet the ICA requirements if controlled connections with SLA are established between the OEM clouds and the IS instances. However, access to the cloud of the V2X messages allows for only a limited time margin with respect to the ICA requirements. We have identified that variations in the latency experienced in the messaging queuing protocols utilized to interconnect the OEMs clouds and ISD due to an increase on the V2X messages processed by the broker may challenge the support of the ICA service.

REFERENCES

[1] 5GAA, "Road traffic operation in a digital age: A holistic cross-stakeholder approach", Whitepaper, January 2024.

[2] M.C. Lucas-Estañ, et al., "Direct-V2X Support with 5G Network-Based Communications: Performance, Challenges and Solutions", *IEEE Network, vol. 37, no. 4*, pp. 200-207, July/August 2023.

[3] C-ROADS Working Group 2,"C-ITS IP Based Interface Profile (v2.0.8)", June 2023.

[4] B. Coll-Perales, et al., "End-to-End V2X Latency Modeling and Analysis in 5G Networks", *IEEE Transactions on Vehicular Technology, vol 72, no. 4*, pp. 5094 - 5109, April 2023.

[5] M. Emara, et al., "MEC-Assisted End-to-End Latency Evaluations for C-V2X Communications", *Proc. EuCNC*, Slovenia, Jun. 2018.

[6] George Nomikos, et al., "O Peer, Where Art Thou? Uncovering Remote Peering Interconnections at IXPs", *Proc. ACM IMC*, Oct. 2018.

[7] 5GAA, "C-V2X Use Cases and Service Level Requirements Volume I", Technical Report, 2020.

[8] Car-to-Car Communications Consortium, "Survey on ITS-G5 CAM statistics", Tech. Report, Dec. 2018.

[9] Cisco, "Best Practices in Core Network Capacity Planning", WhitePaper, Sept. 2020.

[3] Results are not depicted in Table II.b for scenario 4 as the OEM clouds are directly connected and do not rely on a messaging queuing protocol to exchange the messages.