Distribution Sovereignty in Tourism and Mobility in the Age of AI Agents

Strategic note

October 2026

With the participation of

ACCOR · AIRFRANCE KLM GROUP · beti MOBILITÉ AUTOMATISÉE · Capgemini · THINK DO TANK DIGITAL NEW DEAL

G7 · Inria · instant system · SNCF GROUPE

Contents

Summary 3

Introduction 6

1. Issues and context 9
   Distribution put to the test by agentic AI

2. Vision & guiding principles 12
   Towards sovereign, open and AI-agent-ready distribution

3. Use cases 17
   3.0 Traveller journey · 3.1 Sovereignty & resilience · 3.2 Attested Offers · 3.3 Multimodal booking journey

4. Operational risks and protection mechanisms 28

5. Reference architecture 30

6. Governance and operating model 33

7. Benefits and strategic value 34

8. Conditions for collective success & recommendations 35

Annexes: glossary & regulatory overview 38

Disclaimer

This strategic note is a sector contribution resulting from the work of the working group "Sovereignty of service distribution in the age of AI agents", facilitated by EONA-X. It is neither an official standard, nor a normative specification, nor legal advice: it aims to frame a vision, principles, operational flows, protection mechanisms and concrete recommendations.

It may contain forward-looking statements relating to technological, regulatory or market developments. These statements are based on the state of knowledge at the time of writing and remain subject to change. Each organisation is invited to carry out its own checks before making any decision based on this document.

Composition of the working group behind this reflection

Role	Organisations
Coordination & governance	EONA-X (working group facilitation)
Business players	Accor · Air France-KLM · Beti · G7 · Instant System · SNCF
Research & consulting	Capgemini · Digital New Deal · Inria
Technical partner	Linagora

Strategic summary

The note's central message

The challenge is not to reject AI agents, but to define the conditions under which they can access offers and act, without turning European mobility and tourism providers into invisible players.



"The new era of AI agents challenges our collective ability to preserve our technological and economic sovereignty. In this context, the challenge is to enable users to access mobility and tourism services through AI agents, while ensuring that industry players control the conditions under which these agents use and distribute them."

Zineb Essahli Director of Strategy and Transformation, Taxis G7

The diagnosis

AI is transforming distribution: agents can now search, compare, recommend, book and modify services on the traveller's behalf. Market access is gradually shifting towards a model in which it is the customer's agent that queries providers. This reversal profoundly alters established balances and makes it necessary to know, at every interaction, who is acting, for what purpose and with what rights.

This transformation also creates an offensive opportunity: trusted cooperation between players can make it possible to compose services that no provider can produce on its own, by drawing, with consent, on private data, loyalty entitlements, inventories and operational capabilities that are fragmented today.

More broadly, the rise of AI agents represents an unprecedented opportunity for tourism and mobility players. While distribution today remains heavily dependent on search engines and large booking platforms, these technologies could open up new channels of access to offers, foster more direct interactions and enable new forms of cooperation between providers. Within a shared trusted framework, they thus make it possible to build distribution that is more open, interoperable and value-creating for the sector as a whole.

The risk

Without a common framework, European providers risk becoming invisible behind AI assistants: loss of brand visibility, capture of the customer relationship, opaque recommendations, outdated or re-aggregated data, circumvention of commercial rules, dependence on uncontrolled platforms and capture of value by the orchestration layer. Combined, these effects would reduce players to mere executors.

The risk is not limited to dependence on a model: it concerns the entire chain (model, cloud, protocols, orchestration layer, identity, payment and distribution interfaces). The sovereignty sought must therefore be understood as the ability to manage one's interdependencies and to preserve switching options.

The proposed response

The working group proposes a governed agent-to-agent architecture, designed to enable more automated distribution without depriving providers of control over their offers, their data and their commercial rules. It is based on a federated and decentralised data space model, in which each player retains control of its data and of the conditions of access to it, while allowing that data to be used within a trusted framework.

This architecture is organised around three core mechanisms:

1. **The verified agent**
   Every agent is identified, qualified and governable before accessing an offer or triggering an action.

2. **The attested offer**
   The offer is served from the source, fresh, attributable, enforceable and cited to the traveller.

3. **Coordinated booking**
   The multi-party journey is orchestrated without capturing the customer relationship, the data or the value.

These mechanisms apply to two types of interaction. On the one hand, external AI agents can access partners' data, offers or capabilities within a trusted framework. On the other hand, the internal agents of a given player, or those of different organisations, can interact with one another and draw on their respective capabilities. The aim is therefore not simply to open APIs to external assistants, but to organise governed interactions between agents.

The roadmap

The working group structures its work around different use cases covering the main dimensions of sovereign agentic distribution. Three operational use cases explore the mechanisms described above in depth: sovereignty and resilience, attested offers and booking journey. A shared data foundation, cutting across these three use cases, forms the fourth dimension of the work.

The shared data foundation addresses in particular governance at source, connection to local data, standardisation of exchanges and compliance with the European framework. The cross-analysis of the three use cases and of this cross-cutting foundation thus makes it possible to identify the conditions required for controlled agent-to-agent distribution: preserving providers' visibility, their ability to enforce their commercial rules and their relationship with the traveller, while enabling fluid, interoperable and secure orchestration of services.

This work directly feeds into the working group's target architecture and operational recommendations.

Introduction

From the web of search to the web of action

Why this note

The travel industry, whether mobility or tourism, is entering a profound transformation of its distribution models. After the web of search, followed by the web of apps and comparison sites, a web of action is emerging: travellers no longer simply browse websites or interfaces; they entrust an intention to an artificial intelligence system capable of searching, comparing, recommending, booking, modifying and optimising an entire journey.

This shift is not merely an evolution of the user experience. It potentially alters the entire distribution chain: who discovers the offer, who organises the comparison, who recommends, who carries out the transaction, who owns the customer relationship and how value is shared among the various players.

For European mobility and tourism players, this movement carries both a promise and a risk. The promise is that of smoother, more personalised and more accessible distribution, capable of coordinating offers that are fragmented today. The risk is the emergence of a new layer of intermediation, operated in particular by large non-European general-purpose assistants, which would position itself between providers and their customers and gradually capture attention, data, knowledge of the traveller and a share of the value.

The challenge is therefore not to oppose AI agents, but to determine the conditions under which they can access the offers, data and action capabilities of travel players. This note starts from this observation and proposes an operational response.

Origin and mission of the EONA-X working group

The working group “Sovereignty of service distribution in the age of AI agents” was set up by EONA-X and its partners to build this response collectively.

Its starting point is simple: to prevent AI agents from becoming a new layer of intermediation that captures the customer relationship, data and value, without closing the door on a technology set to shape market access for the long term.

The sovereignty sought is therefore neither a defensive posture nor a desire to close off. It refers here to the ability of players to define the conditions under which their offers can be discovered, compared, recommended and activated by agents, while retaining control over their data, their commercial rules, their identity and their relationship with the customer.

The group brings together business players from mobility, tourism and insurance, technology partners and contributors from research. Its core conviction is that no single player can, alone, define a credible trust framework in the face of the emergence of AI agents. The response must be interoperable across partners, sector-wide in scope and European in its roots.

The working group's work

The group's work was structured around four dimensions: three operational use cases (sovereignty and resilience, attested offers, booking journey) and a shared data foundation that cuts across the other three.

Their analysis rests on a common starting condition: relying on data space governance, so that each player retains control over its data, the conditions of access to it and its uses. This governance is the foundation from which the group examines the conditions for agentic distribution that is open, interoperable and controlled.

Gen4Travel: building a new generation of agentic travel services

In the broader context of the transformation of travel by artificial intelligence, EONA-X also coordinates Gen4Travel, a project supported under France 2030 and dedicated to developing a new generation of travel services based on generative and agentic artificial intelligence.

The working group and Gen4Travel are complementary but distinct initiatives, each with its own objectives and governance. Gen4Travel thus provides a concrete field of application for the agentic distribution, interoperability and orchestration issues addressed by the working group.

Gen4Travel responds to a major shift in the uses of AI: the move from assistants mainly capable of informing and recommending to systems capable of understanding an intention, interpreting it, breaking it down into different tasks, searching for services, drawing on several business capabilities and coordinating their execution.

In the travel sector, this shift paves the way for assistants able to support users across a much larger part of their journey, from the initial expression of a need through to the completion or modification of certain operations.

Gen4Travel's ambition is to develop a shared technology foundation enabling tourism and mobility players to design and operate their own intelligent assistants. The aim is not to create a new platform centralising all offers, data and the customer relationship, but to make different capabilities accessible to, and orchestrable by, agents through common, interoperable mechanisms.

The project's architecture thus combines several complementary building blocks: artificial intelligence models, specialised agents, planning and orchestration mechanisms, search and verification tools, providers' information systems, and

mechanisms enabling the various components to communicate and act with one another.

The intended operation can be summarised as a chain of action: collecting the necessary information, adapting it to the context, understanding the traveller's intention, interpreting it, breaking it down into tasks, checking constraints, selecting the appropriate capabilities, executing the actions and then returning the result to the user.

The project also relies on the EONA-X data space, which provides an environment for organising exchanges of data and services between different players. This approach aims in particular to allow providers to retain control over their resources and the conditions under which they can be used, rather than transferring all information to a centralised infrastructure.

To this end, Gen4Travel brings together expertise from tourism, mobility, distribution, artificial intelligence, digital infrastructure and data governance. The project unites travel and technology players around a common objective: making interaction possible between assistants, information systems and business services in an open and interoperable environment.

The developments are being tested in particular in different areas, including Inspiration, Disruption and Accessibility. The Inspiration area focuses on an assistant's ability to understand a traveller's expectations and build a personalised proposal by drawing on different information and services. The Disruption area addresses the management of disruptions and the ability to coordinate several solutions when an event affects the journey. Finally, the Accessibility area aims to better take into account travellers' specific needs and to coordinate the necessary services throughout the trip.

Beyond these experiments, Gen4Travel thus pursues a broader ambition: developing the technical building blocks needed for an environment in which several agents, information systems and service providers can interact to respond to the same user intention.

1. Issues and context

Distribution put to the test by agentic AI

From digital distribution to agentic distribution

For two decades, digital distribution has consisted of guiding customers between websites, apps and comparison sites. Agentic distribution reverses this logic: the traveller expresses an intention to an assistant, which then mobilises agents capable of acting on their behalf. The system no longer merely displays information; it composes a journey, selects options, triggers bookings and manages changes. This shift turns distribution into a conversation between machines, in which the quality of access to offers and the trust placed in each counterpart become decisive.

The AI assistant as the new entry point

Consumer AI assistants and enterprise agents are becoming the traveller's first point of contact. They capture intent very early on, structure the search and steer the recommendation. For providers, the risk is that their offer is no longer consulted directly, but through an interface that controls its presentation, its ranking and its narrative. As these assistants gain autonomy, they no longer merely inform: they decide, book and pay.

The challenge thus becomes twofold: remaining visible in the interfaces to which travellers' attention is shifting, while preventing this visibility from leading to a lasting loss of control over the offer, data and customer relationship.

Distinguishing between AI assistant and AI agent

In this note, the AI assistant refers to the interface with which the traveller interacts: it captures the intention, presents the results and, where necessary, requests validation. The AI agent refers to a software component capable of executing or coordinating a task on behalf of a user or an organisation: querying a source, checking availability, comparing, booking, modifying or triggering an action.

A single assistant may mobilise several specialised agents, and an agent may act without being directly visible to the traveller. This distinction is important for correctly attributing identity, mandate, rights and responsibility: the assistant carries the user experience; the agent carries a defined capacity for action, subject to the permissions granted to it.

Some market benchmarks

Several signals are converging. The major conversational assistants now count hundreds of millions of active users and handle a growing volume of transactional queries. Traffic directed to merchant sites by agents and assistants is rising sharply year on year, particularly during peak commercial periods. According to some industry projections, a significant share of travel bookings could be initiated by an AI agent by the end of the decade. Above all, interoperability protocols between agents are already spreading, often at the initiative of non-European players: this is precisely the ground on which a sovereign response must position itself without delay.

The shift towards a customer-to-business model

Agentic distribution establishes a so-called "customer-to-business" model: it is the customer's agent that queries providers, and not the other way round. This shift changes market access and makes it necessary to know, for every request, who is asking, for what purpose and with what rights. It also shifts where trust lies: it is no longer only the traveller who must be reassured, but the agent that must be authenticated and supervised.

The chain of trust must be able to distinguish at least four levels: the technology player at the origin of the action; the organisation that operates or mandates the agent; any distribution intermediary; and the end traveller, within the strict scope required for the purpose pursued. The travel context (leisure or corporate, group or family, loyalty entitlements, specific needs) can then be drawn on in accordance with consent and the applicable rules.

The risks of disintermediation

Agentic disintermediation covers several converging risks. The brand may lose visibility as soon as the offer is reformulated by the interface. The direct customer relationship may be captured by the assistant. Recommendations may become opaque, based on undisclosed criteria. The data used may be scraped, re-aggregated or outdated. Commercial and pricing rules may be circumvented. Dependence on uncontrolled platforms or models may set in. Finally, value may be captured by the orchestration layer. Taken individually, each of these risks is manageable; combined, they sketch a scenario in which providers gradually fade away.

Uncontrolled exposure also creates an asymmetry that is difficult to reverse: once price, inventory or behavioural histories have been collected and stored by a third party, the provider cannot simply "take back" that knowledge. Hence the importance of targeted, purpose-bound and revocable exposure rather than indiscriminate opening of inventories.

Why the mobility and tourism sectors are particularly exposed

Mobility and tourism journeys are multi-party, dynamic and highly sensitive to real time: variable prices, fluctuating availability, disruptions, ancillary services and compliance requirements constantly interact. This complexity makes the freshness and reliability of data critical, and disintermediation therefore particularly damaging. An agent relying on approximate information would expose both the traveller and the provider.

Personalisation is also richer than in many other sectors: loyalty status, business travel, accessibility needs, pets, bicycles, group composition, schedules, preferences and unforeseen events can change how relevant an offer actually is. The ability to draw on this context with consent becomes a source of differentiation and service quality.

Why a collective response is needed

No single player can, alone, impose a credible trust framework on AI agents: the value of an agent registry, an attested offer format or audit rules lies precisely in their shared adoption. The response must therefore be interoperable across partners, sector-wide in scope and European in its roots. This is the purpose of the framework developed in the following chapters, which combines a vision, an architecture and demonstrable use cases.

From protection to value creation

The collective response must not be purely defensive. By linking identity, consent, private data and business capabilities, the industry can create an experience more useful than the mere aggregation of offers: recognition of loyalty entitlements without re-entry, seamless chaining of several services, proactive management of a disruption, relevant additional services and end-to-end continuity. Sovereignty then becomes a means of delivering a better experience, rather than a constraint imposed on the traveller.

2. Vision & guiding principles

Towards sovereign, open and AI-agent-ready distribution

The working group's target vision

The vision is to identify the levers that will allow AI agents to access partners' offers within a trusted framework, without capturing the customer relationship or causing providers to lose control. Sovereignty is not closure: it is the concrete ability to choose, control, trace, negotiate and audit the conditions under which one's offers are exposed. A sovereign player does not prohibit access by agents; it sets the rules and retains the ability to change or revoke them.

This vision adds a requirement: provider control must not come at the expense of traveller choice. The framework must allow providers to retain control of their offer and their conditions, while guaranteeing fair comparison, explainable recommendation criteria and the possibility for the traveller to request a specific brand as well as a neutral search.

Defining distribution sovereignty

Distribution sovereignty is defined here as the ability to decide who accesses the offer, with what rights, for what purpose, under what rules, with what traceability and with what possibilities of revocation. It applies simultaneously to data, to recommendation criteria, to commercial conditions and to the choice of technical intermediaries. It cannot be decreed: it must be equipped, through identity, attestation, orchestration and audit building blocks.

It is also measured by the ability to keep operating when a link in the chain changes or disappears. Resilience therefore becomes a verifiable property of the architecture: known dependencies, tested switchover mechanisms, portable formats and continuity of rights and mandates.

Principle 1: trust by design

Identity, rights, consent and audit must be built in from the design of the framework, not added after the fact. Trust is an architectural property, not an option: a system that merely added safeguards a posteriori would remain structurally vulnerable. In practice, this means that every agentic flow must, by default, be identified, authorised and logged.

This principle explicitly extends to the traveller's identity and to the mandate: an agent must be able to prove on whose behalf it is acting, on what basis, for what operation, for how long and within what limits. The identity and consent building blocks developed in

related initiatives therefore provide a relevant testing ground for linking identity, consent and data access.

Principle 2: a model-independent and resilient architecture

The framework must work regardless of the model or assistant used, whether Mistral, Claude, ChatGPT, Gemini, an enterprise agent or a future solution. Model independence is a condition of sovereignty and reversibility: it prevents any model provider from becoming an unavoidable gateway and guarantees the portability of the system over time.

This independence must be demonstrated through testing: unavailability of a model or cloud, a change in a provider's terms, failure of an orchestrator or evolution of a protocol. Switchover must be possible without losing journey continuity, partner connections, proof of consent or mandates that are still valid. Dependency mapping inspired by approaches such as the Digital Resilience Index (IRN), led by the Digital New Deal think tank, can help to assess this capability objectively.

Principle 3: the offer served from the source

The offer must remain fresh, attributable, enforceable and linked to its provider. Any sensitive recommendation must be able to cite its source and rely on verifiable provenance, rather than on re-aggregated data whose origin and validity are no longer guaranteed. Serving the offer from the source both protects the traveller from erroneous information and preserves the provider's attribution.

Exposure must remain proportionate to the intention: an agent does not need to receive an entire inventory to respond to a specific request. The target is a targeted query, with identified rights and purpose, freshness metadata and machine-readable terms of use.

Principle 4: orchestration without capture

The orchestration layer coordinates without owning: it must become neither an exclusive channel, nor the holder of the customer relationship, nor the enforcer of its own commercial rules. Coordinating without dispossessing is the framework's red line. This requires each partner to execute its own part of the transaction and to retain the relationship with its customer.

The framework does not assume that a central orchestrator is required for every interaction. Depending on the use case, two agents may communicate directly under a common trust framework; in other scenarios, multi-party orchestration brings real value. The role of the data space is to provide identity, rules, usage contracts, traceability and trust services, without requiring all data to pass through a central platform.

Principle 5: human control over sensitive decisions

The traveller and the provider remain in control of binding actions. The mandate given to the agent is explicit, consent is revocable, and humans remain in the loop for high-stakes decisions, particularly those involving payment, cancellation or personal data. The agent's autonomy ends where the risk to the user begins.

Protecting the provider relationship must not limit the traveller's freedom of choice. Recommendations must be traceable, the main recommendation criteria must be explainable and selection rules must be non-discriminatory. A preference expressed by the traveller may be respected, whereas a bias introduced by the orchestrator must be detectable and auditable.

Principle 6: interoperability through open standards

The framework relies on open standards, existing APIs, connectors and data space principles, to guarantee interoperability between partners and portability over time.

Building on what already exists rather than reinventing it accelerates adoption and lowers the cost of entry for each player.

The standardisation sought concerns first and foremost the exchange framework: identity and permissions, offer semantics, ontologies, purpose, usage rules, traceability and error handling. The protocols and technical mechanisms used to implement these exchanges are likely to evolve. The framework must therefore remain open, interoperable and independent of any particular technology.

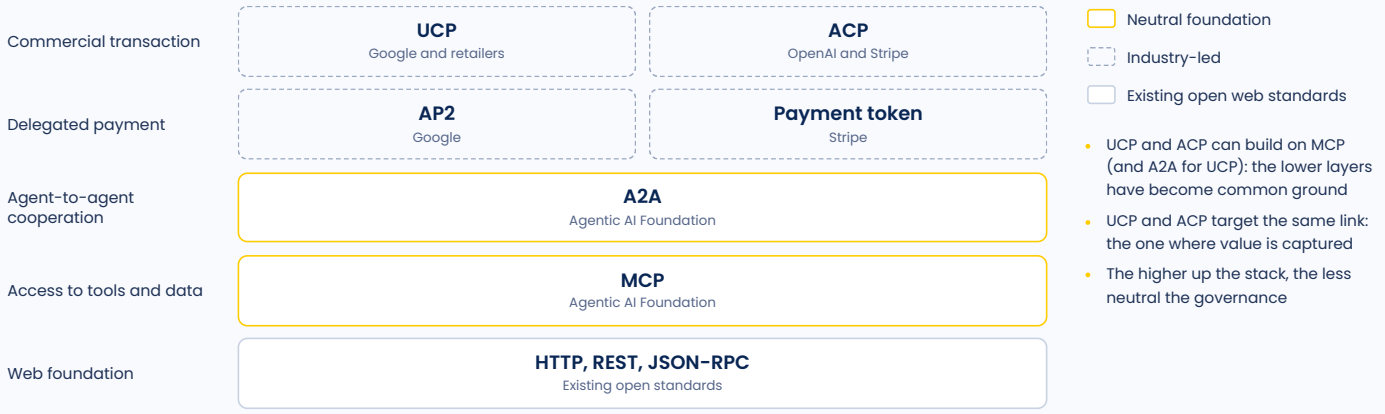
Principle 7: cooperating to create services that are impossible in isolation

The framework must allow the industry to demonstrate positive value: combining several business capabilities, recognising a customer and their entitlements, preserving their preferences from one service to the next, managing a change or disruption in a coordinated way and offering relevant associated services. The data space then becomes a trust infrastructure serving the creation of collective value.

3. Agentic protocols: from open standards to new control points

As part of our work, LINAGORA presented an analysis of the main agentic protocols. This perspective clarifies the roles of their different layers, their governance models and the new control points they give rise to.

Figure 1: *A stack of complementary protocols*



Distinct but composable layers

MCP connects the agent to the operator's tools and data, while A2A enables task delegation between agents from different organisations. UCP and ACP, for their part, cover the stages of the commercial transaction, with complementary payment mechanisms such as AP2 or payment tokens.

The point is therefore not to choose a single protocol, but to compose interfaces suited to each function in the chain.

Increasingly contested governance

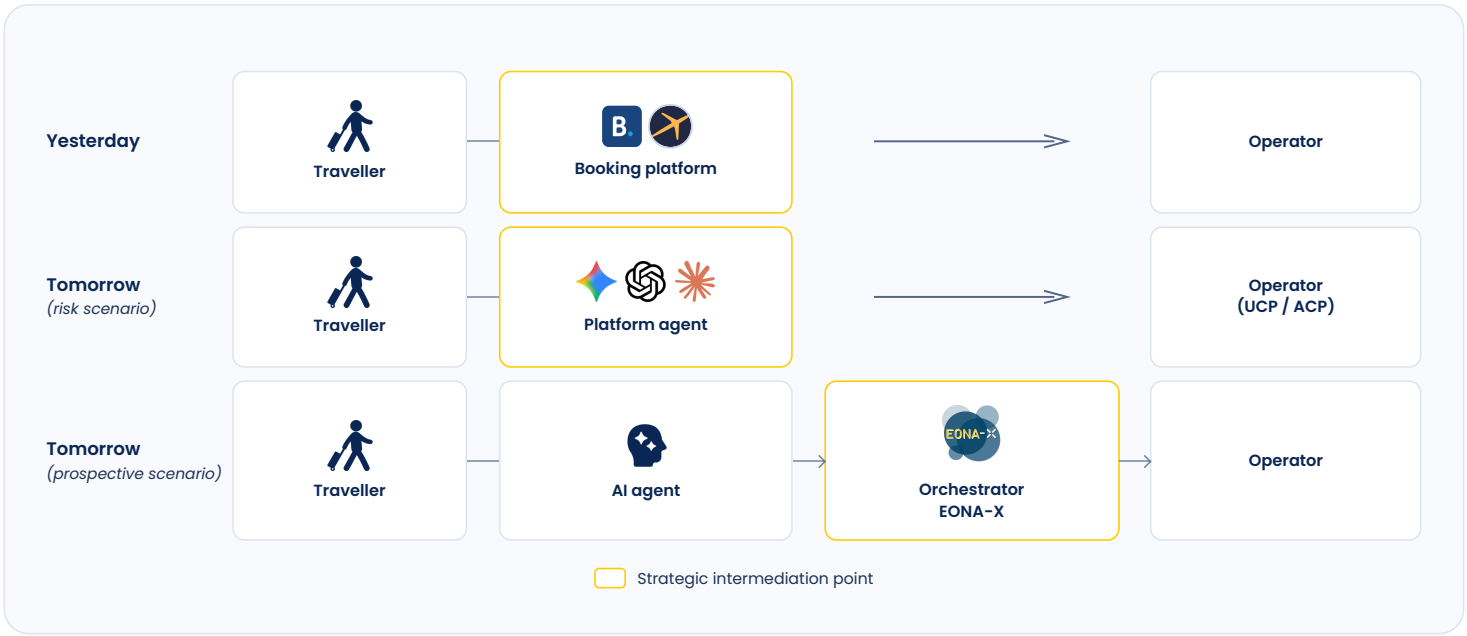
MCP and A2A fall under the Agentic AI Foundation, whereas commerce standards remain more closely steered by their initiators: UCP (Google and partners) and ACP (OpenAI and Stripe). Publishing an open specification does not, in itself, guarantee the neutrality of the body that develops it.

LINAGORA highlights the emergence of a UCP variant dedicated to the hotel industry ("Lodging"), in which Amadeus and several industry players are taking part.

Agentic commerce: where is intermediation shifting?

LINAGORA offers a reading of how the distribution chain is evolving: from the booking site to the platform agent, and then to the hypothesis of collective orchestration. The diagram below illustrates this progression.

Figure 2: *Evolution of distribution interfaces*



What this shift changes in agentic commerce

With UCP and ACP, the agent can become the interface for discovery and purchase initiation, while the provider retains commercial execution. The point of attention then becomes the concrete allocation of responsibilities between the party presenting the offer, the party transmitting the order and the party assuming the obligations linked to the sale. A transaction protocol does not, in itself, settle this allocation.

The collective orchestrator scenario presented by LINAGORA illustrates one possible way of connecting several agents and operators without imposing a proprietary integration per platform. For it to be a credible alternative, its listing arrangements, request routing rules and sharing of responsibilities should remain explicit and governed by the stakeholders.

European requirements to be built in from the design stage

The presentation highlights two concrete points for journeys combining mobility and tourism: strong customer authentication requirements at payment, which may call for direct validation by the traveller, and the rules applicable to package travel when an agent assembles several services. These constraints must be taken into account when defining transactional journeys and their validation points.

3. Use cases

Three use cases to structure the framework

This section is the operational core of the note. It applies the sequence “verified agent, attested offer, coordinated booking” through three complementary use cases. Each diagram follows the working group’s visual identity (actors, layers and process steps) and is followed by a detailed analysis, concluding with directly actionable recommendations.

3.0 Reference traveller journey

The journey below illustrates an experience in which a traveller states an intention to an assistant, authorises a scope of action, receives reliable offers and then coordinates a booking and its follow-up, without any shared player capturing the customer relationship. It can be demonstrated with two complementary partners before being extended to more players.

Reference traveller journey

TRUSTED AGENTIC DISTRIBUTION

An end-to-end view linking identity, mandate, offer, choice, execution and journey continuity, without any shared player capturing the customer relationship.

Traveller · Assistant · Trust framework · Providers

1 Intent & context

The traveller expresses their need, preferences and context: leisure or corporate, group, accessibility.

WHAT IS EXPRESSED

- Stated need
- Preferences
- Leisure or corporate
- Group · accessibility

2 Identity, mandate & consent

The assistant establishes on whose behalf it acts, and the purpose, duration and limits of the mandate. Loyalty entitlements can be recognised.

WHAT IS ESTABLISHED

- On whose behalf the agent acts
- Purpose of the mandate
- Duration and limits
- Loyalty entitlements

3 Identity qualification

The framework verifies the identity of the traveller and that of the agent acting on their behalf, as well as the link with the mandate and associated permissions.

WHAT IS QUALIFIED

- Traveller identity
- Agent identity
- Agent ↔ mandate link
- Associated permissions

4 Targeted search

The agent only queries the inventories relevant to the stated intention, in accordance with the applicable access rules.

WHAT IS QUERIED

- Relevant inventories
- Access rules
- Applicable contracts
- Minimal scope

journey continues

5 Attested offers

Providers serve fresh, attributed and contextualised offers, together with their conditions.

WHAT IS SERVED

- Fresh offers
- Attributed offers
- Contextualisation
- Conditions & services

6 Comparison & choice

The assistant makes sources and criteria explicit. Commercial rules are balanced with the traveller's freedom of choice.

WHAT IS MADE EXPLICIT

- Offer sources
- Ranking criteria
- Commercial rules
- Freedom of choice

7 Coordinated booking

7.1 Booking with each provider

Each provider retains control of its offer and itself executes the booking of its service, through to confirmation.

7.2 Separate functions

The functions can be decoupled: **inventory access → booking → payment**. They can be performed by different players while remaining coordinated within the journey.

8 Follow-up, changes & disruption

Confirmations, voluntary changes and disruptions are coordinated; the mandate remains revocable at any time.

WHAT REMAINS OPEN

- Confirmations
- Voluntary changes
- Revocable mandate
- Auditable actions

What this journey must preserve

THE FIVE JOURNEY INVARIANTS

Whatever the implementation (direct agent-to-agent connection or shared orchestration), these five properties must hold.



01

The traveller stays in control

They retain control over their mandate and their choice at every stage of the journey.



02

Each provider remains in charge

It retains control of its offer and itself executes its share of the booking.



03

Only necessary data flows

Minimisation applies to every exchange, from mandate to follow-up.



04

Sources and decisions remain auditable

The origin of offers and the decisions taken can be traced and verified.



05

The journey survives a failure

No component is an unavoidable gateway: no central platform is recreated.

WHY THESE INVARIANTS

They are the condition for agentic distribution to remain sovereign: the journey can be demonstrated with two complementary partners, then extended to more players, without a shared intermediary standing between the traveller and providers.

3.1 Use case 1: Sovereignty & resilience

Identifying the AI agent seeking to distribute the offer

Executive summary

An AI agent must be verified, qualified and governable before it can access offers or trigger an action. Identification comes before access: it is the first condition for sovereign distribution, because one can only govern what one can identify.

Identification does not concern only the software executing the action. The chain of trust must be able to link the agent to the technology player that operates it, to the organisation that mandates it, to any distribution intermediary and, where the purpose so requires, to the traveller or principal.

Problem

Today, an agent or a simple script can query an API and trigger a transaction without any shared framework. The provider does not always know precisely who is acting, for what purpose, with what rights and under whose responsibility. This opacity opens the door to uncontrolled uses, non-compliant actions and a dilution of responsibility in the event of an incident.

Vision

The target is for every agent to be identifiable, qualifiable and governable, by means of a model-independent agent-to-agent trust mechanism. The agent introduces itself, proves its identity and affiliation, and is granted a level of access proportionate to what it is authorised to do.

Target process

The traveller expresses an intention to their assistant, which mobilises an agent. The agent presents itself to the trust layer; its identity, publisher and organisation are verified; its rights are qualified according to the authorised level of action. When the action requires personal or transactional context, the traveller's mandate is verified in parallel and only the attributes necessary for the purpose are exposed. Any distribution intermediary is identified as a player distinct from the AI agent. Every sensitive action remains traced, auditable and revocable.

Key capabilities

The use case draws on a registry of verified agents, a verifiable identity based on attestations, fine-grained rights management distinguishing reading, recommending and acting, access revocation, auditing of agent behaviour, as well as mechanisms for controlling hallucinations and self-preferencing.

Anti-disintermediation mechanisms

Protection relies on several cumulative levers: a registry of verified and qualified agents; a ban on any sensitive action without explicit rights; distinct authorisation levels according to the type of action; revocation possible at any time; logging of all requests; and behavioural auditing designed to detect self-preferencing and drift.

Questions to be decided

Several points remain to be decided collectively: who attests the identity of an AI agent; which minimum rights should be distinguished; which uses should be prohibited from the first version; how to guarantee independence from the model used; and who bears responsibility in the event of a non-compliant action.

Recommendations

The working group recommends creating a common agent registry, formalising explicit authorisation levels (read, recommend, act) and prohibiting any sensitive action without verified identity. It also recommends systematically logging every request and making revocation immediate and enforceable, so that access can be withdrawn without delay in the event of misuse.

Use case 1: Sovereignty & resilience

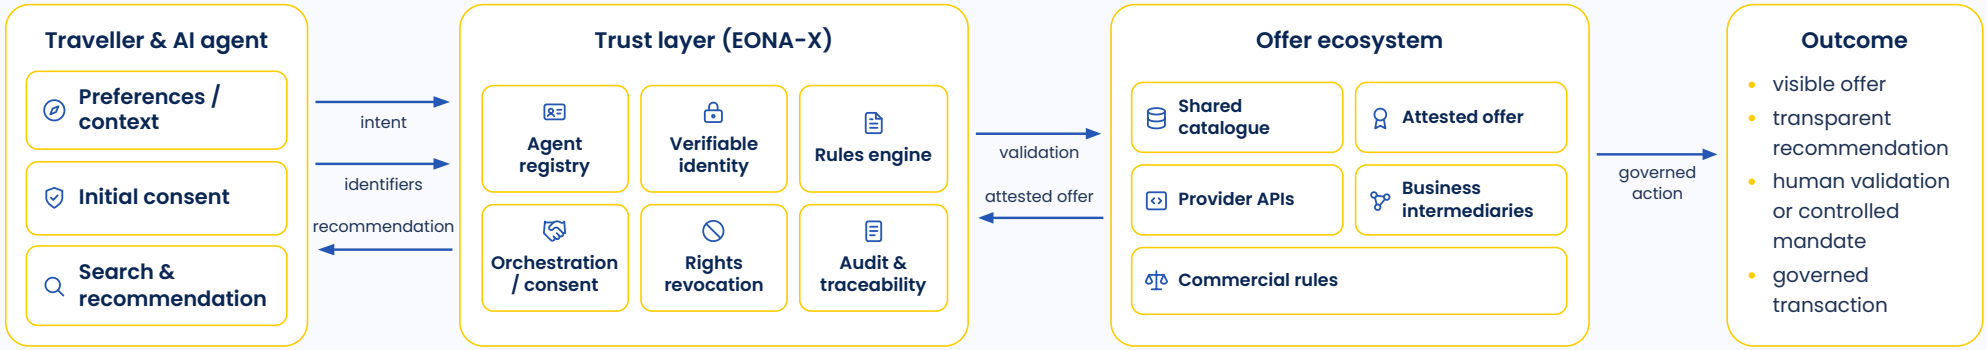
Identifying the AI agent seeking to distribute the offer

Challenge

Problem: any script claiming to be an "AI agent" can query an API and trigger a transaction without a shared framework.

Vision: an agent must be identifiable, qualifiable and governable before acting, via a model-agnostic agent-to-agent trust mechanism.

Actor / component	Role in the scenario
Traveller	Expresses their intention, shares preferences, consents, then validates the offer.
Distributing AI agent	Searches, rephrases, recommends and acts according to its permissions.
EONA-X	Verifies identity, role, organisation, rights and applicable rules.
Offer provider	Exposes an attested, visible and enforceable offer.
Business intermediary	Shares APIs and contributes to orchestration.
Shared catalogue	Provides provenance, freshness and multi-company data.



Process steps

1. Travel intent: the traveller expresses their preferences via an AI agent.
2. Verification: the agent presents a verifiable identity and credentials.
3. Trust decision: role, organisation, rights and usage contracts are verified.
4. Exposure: providers and intermediaries expose attested offers.
5. Recommendation: the agent makes its sources explicit and the traveller stays in control.
6. Governed action: the transaction is executed with consent and traceability.

3.2 Use case 2: Attested Offers

Exposing and attesting a reliable, fresh and attributable offer

Executive summary

An offer must remain fresh, attributable, enforceable and served from the source. Any sensitive recommendation must cite its provider, so that the traveller always knows where the offer comes from and under what conditions it applies. This is the condition for the agent to enrich distribution without distorting it.

Problem

AI agents may rely on scraped, outdated or re-aggregated data. Travellers no longer always know who is issuing the offer, whether it is still valid, or what conditions apply to it. This uncertainty undermines trust and exposes providers to seeing their offers distorted, presented out of context or compared in a misleading way.

Vision

The target is to create an attested offer that can be consumed by an AI agent within a trusted framework, with provenance, verifiable freshness and machine-readable exposure rules. The offer ceases to be raw data and becomes an attributable, enforceable object, durably linked to its provider.

Target process

The provider publishes an offer together with its metadata; provenance, freshness and exposure rules are attached to it; the agent queries the trust layer; the offer is returned with its source, its conditions and its level of reliability; the traveller understands the origin of the offer before validating; the booking is then initiated without loss of attribution.

Key capabilities

The use case draws on offer provenance, verified freshness, mandatory citation of the provider, machine-readable exposure rules, brand control, standardised metadata and recommendation auditing.

Anti-disintermediation mechanisms

Protection involves mandatory citation of the provider in any sensitive recommendation, attaching freshness and provenance to each offer, prohibiting any misleading rephrasing, preserving brand control, machine-readable exposure rules and systematic logging of recommendations.

Frictions identified by the partners

The partners identified several concrete frictions: pricing complexity and dynamic conditions, the need for personalisation, recognition of loyal customers, the requirement for transparency on the final price, and the role of additional services and reassurance.

Added to these are the lack of stabilised standards, the speed at which protocols evolve, interaction performance, volumes at the top of the funnel and the need to reach a critical mass of offers and content for an agentic channel to be truly relevant.

Finally, traditional distribution already faces difficulties in contracting access to offers and fare inventories. Access to availability, fares and booking capabilities often remains conditional on commercial agreements, specific authorisations and each provider's own rules. The involvement of AI agents therefore adds further complexity, particularly in determining access rights, responsibilities and the conditions under which these agents may query an offer or initiate a transaction. Technical interoperability does not, in itself, create a commercial right of access or booking.

Questions to be decided

Still to be clarified collectively: what exactly constitutes an attested offer; the minimum metadata that must accompany it; the means of discouraging scraping; how to guarantee brand visibility; and the handling of dynamic prices, availability and conditions.

Recommendations

The working group recommends defining a minimal attested offer format, shared between partners, and making source citation mandatory in any sensitive recommendation. It also recommends measuring and exposing the freshness of the offer, so that an agent can discard any outdated information and the traveller always has a valid, clearly attributed offer.

Use case 2: Attested Offers

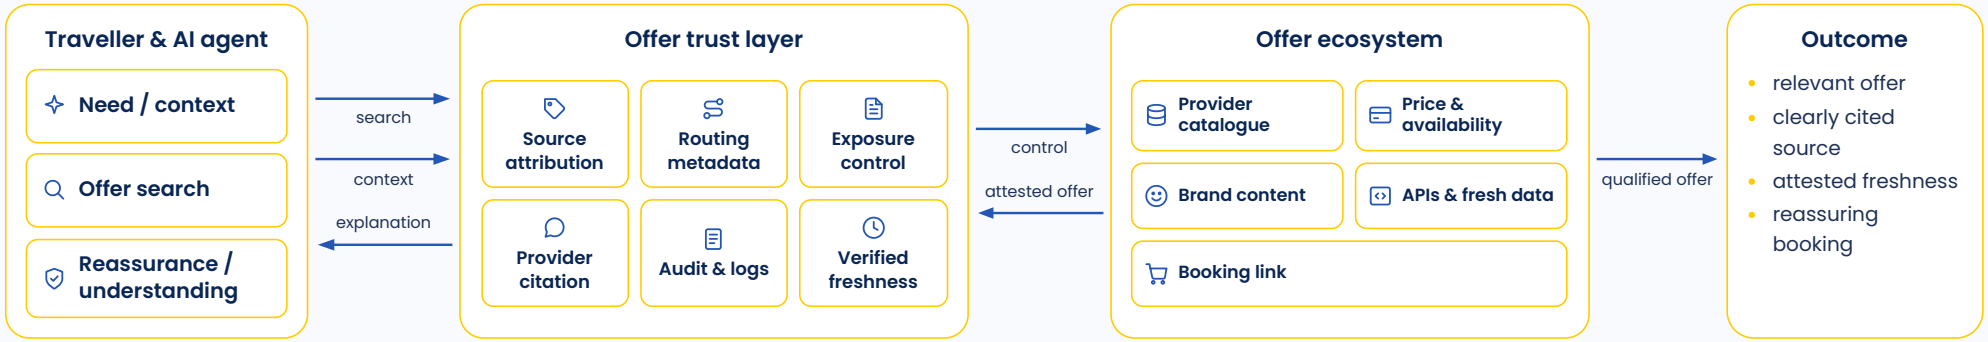
Exposing and attesting a reliable offer from a European travel / mobility provider

Challenge

Problem: the AI agent receives data scraped or re-aggregated by third parties, sometimes several hours old. The traveller no longer knows who is issuing the offer or whether it is still valid.

Vision: an offer must remain attributable, traceable and served from the source, with provider citation, verifiable freshness and exposure rules.

Actor / component	Role in the scenario
Traveller	Looks for a relevant, reliable, understandable and bookable offer.
Distributing AI agent	Queries sources, contextualises the request and recommends.
Offer Trust Layer	Attests provenance, applies exposure rules and ensures auditability.
Provider	Publishes a fresh, attributable and governed offer.
Business intermediary	Relays APIs, enriches context and orchestrates.
Shared catalogue	Manages availability, prices, metadata and source citation.



Process steps

1. Publication: the provider exposes the offer, its metadata and its access rules.
2. Attestation: provenance, freshness and citation are attached to the offer.
3. Contextualised search: the agent queries the catalogue according to the need.
4. Recommendation: the offer is presented with source, criteria and reliability.
5. Validation: the traveller understands the origin and stays in control before booking.
6. Booking: the selected offer is booked without loss of attribution.

3.3 Use case 3: Booking journey

Coordinating a multi-party journey without recreating a predatory platform

Executive summary

The booking journey must reduce friction for the traveller without recreating a platform that would capture the customer relationship, the data or the margin. The objective is clear: coordinate, not centralise.

Problem

A multimodal journey often involves several websites, accounts, forms and repeated data entry. Intermediation platforms reduce this friction, but at the cost of heavy intermediation and value capture. The difficulty lies in offering the same fluidity without reproducing this pattern of dispossessing providers.

How a booking API works today

In practice, a booking API is not open access to a provider's inventory. It exposes functions for searching, creating and managing a booking, but transactional use is generally conditional on a contractual and commercial relationship with the provider or with an authorised intermediary. Depending on the sector and the channel, this may involve technical credentials, access rights, pricing and distribution rules, security or certification prerequisites, as well as specific arrangements for payment, ticketing, settlement and servicing.

The existence of an API, a standard or a protocol therefore does not, in itself, create a right to book. An AI agent must act within the scope of an authorised player and under a verifiable mandate. The proposed framework adds a layer of trust and orchestration to existing mechanisms. It replaces neither commercial contracts nor providers' transactional authorisations.

For a multi-party journey, two approaches can coexist: each provider retains and executes its own booking; or orchestration separately sequences inventory access, booking and payment. In both cases, the access rules and the responsibility of each provider remain enforceable.

Vision

The target is to coordinate without centralising or capturing: a journey distributed among several partners, with single consent, minimisation of shared data, full traceability and value-sharing rules clearly established from the outset.

Target process

The traveller expresses an intention to the assistant; the assistant mobilises one or more agents to compose the journey; the traveller gives single consent; calls are orchestrated towards providers; each partner directly executes its share of the transaction; confirmations, changes and disruptions are tracked in a coherent and readable flow.

Key capabilities

The use case draws on orchestration without centralisation, single and revocable consent, minimisation of shared data, management of action permissions, distributed booking, traceability of flows and value-sharing rules.

Anti-disintermediation mechanisms

Protection relies on orchestration without an exclusive customer relationship, each provider retaining its customer relationship, single and revocable consent, minimisation of shared data, direct execution by partners and traceability of both flows and value sharing.

Questions to be decided

Still to be decided: which minimal journey to demonstrate first; which partners to involve; the scope of payment in an initial phase; how to trace each player's contribution; and the safeguards to prevent orchestration from itself becoming a dominant platform.

Recommendations

The working group recommends building on Gen4Travel's orchestration component, explicitly preserving each provider's customer relationship, tracing every action by the agent and defining, from the outset, transparent value-sharing rules among the partners involved.

Use case 3: Booking flow

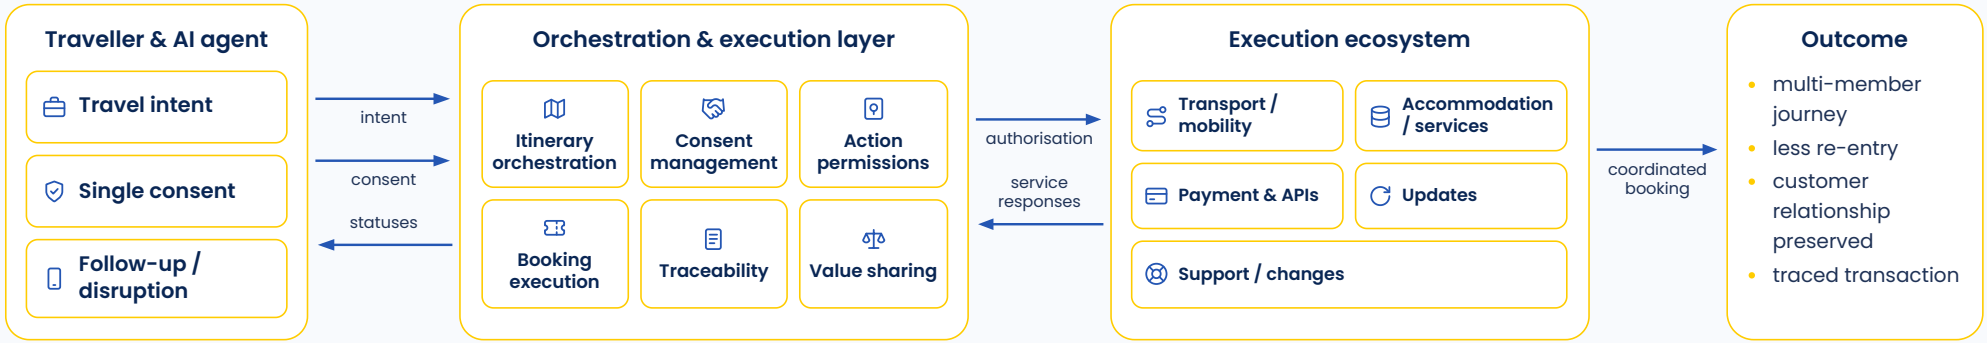
Coordinating a multi-party journey without an OTA

Challenge

Problem: a multimodal trip spans 5 websites, 5 accounts and 5 rounds of data entry. OTAs solve the friction by inserting themselves, capturing the customer relationship and the margin.

Vision: coordinate without centralising or capturing the customer relationship, through a distributed multi-member flow with consent, minimisation, traceability and value sharing.

Actor / component	Role in the scenario
Traveller	States their intention, gives single consent and stays in control.
AI agent / orchestrator	Assembles the journey and triggers calls according to permissions.
Orchestration Layer	Coordinates calls, traces flows and applies rules.
Business members	Execute their part of the transaction directly.
Booking / payment	Manage booking, payment, changes and cancellation.
Audit & governance	Verify consent, traceability and value sharing.



Process steps

1. Travel intent: the traveller expresses their need and preferences.
2. Assembly: the agent composes a multi-party proposal from the offers.
3. Consent: the traveller authorises minimal sharing of the necessary data.
4. Orchestration: calls are coordinated without centralising the customer relationship.
5. Booking & payment: each member executes its part with traceability.
6. Follow-up: confirmations, changes and disruptions are reported in a coherent flow.

4. Operational risks and protection mechanisms

The safeguards to build into the three use cases

Moving from a world in which customers browse to one in which their agent acts introduces specific risks. The proposed framework addresses them through cross-cutting mechanisms, applicable to all three use cases and designed to be reusable. The following paragraphs detail the main risks and the associated mitigations.

Risk of capture

The first risk is that the orchestrator becomes a new dominant platform. Mitigation relies on shared governance, transparency about EONA-X's role and each provider retaining its customer relationship: orchestration coordinates, but appropriates neither the customer nor the value.

Neutrality must be verifiable: documented routing rules, no exclusivity, management of conflicts of interest, common access criteria and auditing of any preferences on the part of the orchestrator.

Risk of poor recommendation

The agent may recommend an offer that is unsuitable, biased or non-compliant. The response is to make sources explicit, display recommendation criteria, audit recommendations and detect self-preferencing, so that no player can unduly favour its own offers.

Fairness must be assessed from the traveller's point of view as much as from the provider's: a legitimate commercial rule may limit an offer's eligibility, but it must not be confused with an opaque ranking or an imposed preference.

Risk of hallucination

The agent may invent an offer, a condition or availability. Attested offers, fresh data, access served from the source and limiting unsourced answers greatly reduce this risk by anchoring every answer in verifiable information.

Risk of loss of data control

Exposed data may be reused or used to train models without authorisation. Machine-readable usage rules, a no-training-by-default principle and access revocation allow providers to retain control over their information assets.

Risk of technological dependence

A player may become dependent on a model, a cloud or a protocol. A model-independent architecture, open standards, interoperability and reversibility preserve freedom of choice and prevent any lock-in.

Cybersecurity, volume and instability risk

Automation can multiply requests, create retry loops or amplify an attack. The framework must incorporate strong authentication, quotas, request rate limiting, prevention of duplicate execution for transactional actions, anomaly detection, emergency suspension mechanisms, segmentation of rights and recovery arrangements. Performance and stability objectives must be measured from the POC stage, particularly at the upstream end of the journey.

Legal and reputational risk

A pricing error, incorrect information, a non-compliant transaction or poorly managed consent exposes the players. Keeping humans in the loop (human-in-the-loop), an explicit mandate, complete logs, auditability and a clear allocation of responsibilities provide the best protection.

The legal classification of multi-service distribution, the role of any authorised intermediary and the scope of payment must be validated for each demonstrator before a proof of concept is turned into a commercial service.

Table 1: *Summary of operational safeguards*

Risk	Protection mechanism	Owner
Unidentified access	Verifiable identity, agent registry, authorisation levels	EONA-X & technical partners
Unclear consent or mandate	Explicit, revocable consent, with purpose and action specified	Partners & EONA-X
Loss of attribution	Mandatory source citation, brand control	Business partners
Outdated or scraped data	Attested offer: freshness, provenance, served from the source	Business partners
Opaque actions	Traceability and auditability of every request and recommendation	Technical & research partners
Capture by the orchestrator	Orchestration without an exclusive customer relationship or rules of its own	EONA-X
Data reuse	No model training by default, unless authorised	Partners & EONA-X

5. Reference architecture

Governed agent-to-agent distribution

Overview

The target architecture connects the traveller to their AI assistant, then to a specialised agent, which communicates with a trust layer operated at ecosystem level (EONA-X), itself connected to providers via connectors. The end-to-end flow runs as follows: the traveller addresses their assistant, which mobilises a specialised agent, which goes through the trust layer before reaching the partners' connectors and, finally, their offers, APIs and information systems. On output, every action is traced and auditable.

This architecture can be read on two complementary levels: an open discovery layer, which promotes visibility, and a governed trust and execution layer, which frames sensitive actions. The components below specify the role of each building block.

Discovery layer

The discovery layer provides access to open content: inspiration, public information, initial browsable catalogues and brand content. It maximises players' visibility without engaging in any sensitive action, and is the natural gateway for AI assistants into the ecosystem.

Trust and execution layer

The trust and execution layer frames high-stakes actions: qualified recommendation, booking, modification, cancellation, consent and transaction. This is where identity, rights, traceability and audit apply. It is the heart of the sovereignty mechanism.

Agent registry

The agent registry makes it possible to identify the agent, its publisher, its organisation, its role, its rights and its trust level. It is the reference directory that makes the agent attributable and governable: without registration and without explicit rights, no agent should be able to trigger a sensitive action.

Verifiable identity and access rights

Verifiable identity makes it possible to link each agent to its publisher, its operator or the organisation that mandates it, and then to prove the attributes and permissions associated with it. The system may align with eIDAS 2.0 and the European digital identity framework, in particular through electronic attestations of attributes and trust services, in order to have a root of trust recognised at European level. eIDAS 2.0 does not, however, by itself confer a legal identity on the agent: proofs of identity, responsibility

and mandate will need to be translated by the EONA-X registry into distinct operational rights (read, recommend, act) that can be revoked at any time. This granularity prevents an agent authorised to consult from being able, in practice, to book or pay.

Offer attestation layer

The offer attestation layer attests the provenance, freshness, attribution and exposure rules of an offer. It makes the offer enforceable and durably links it to its provider, so that a recommendation can always cite its source and the traveller has reliable information.

Orchestration layer

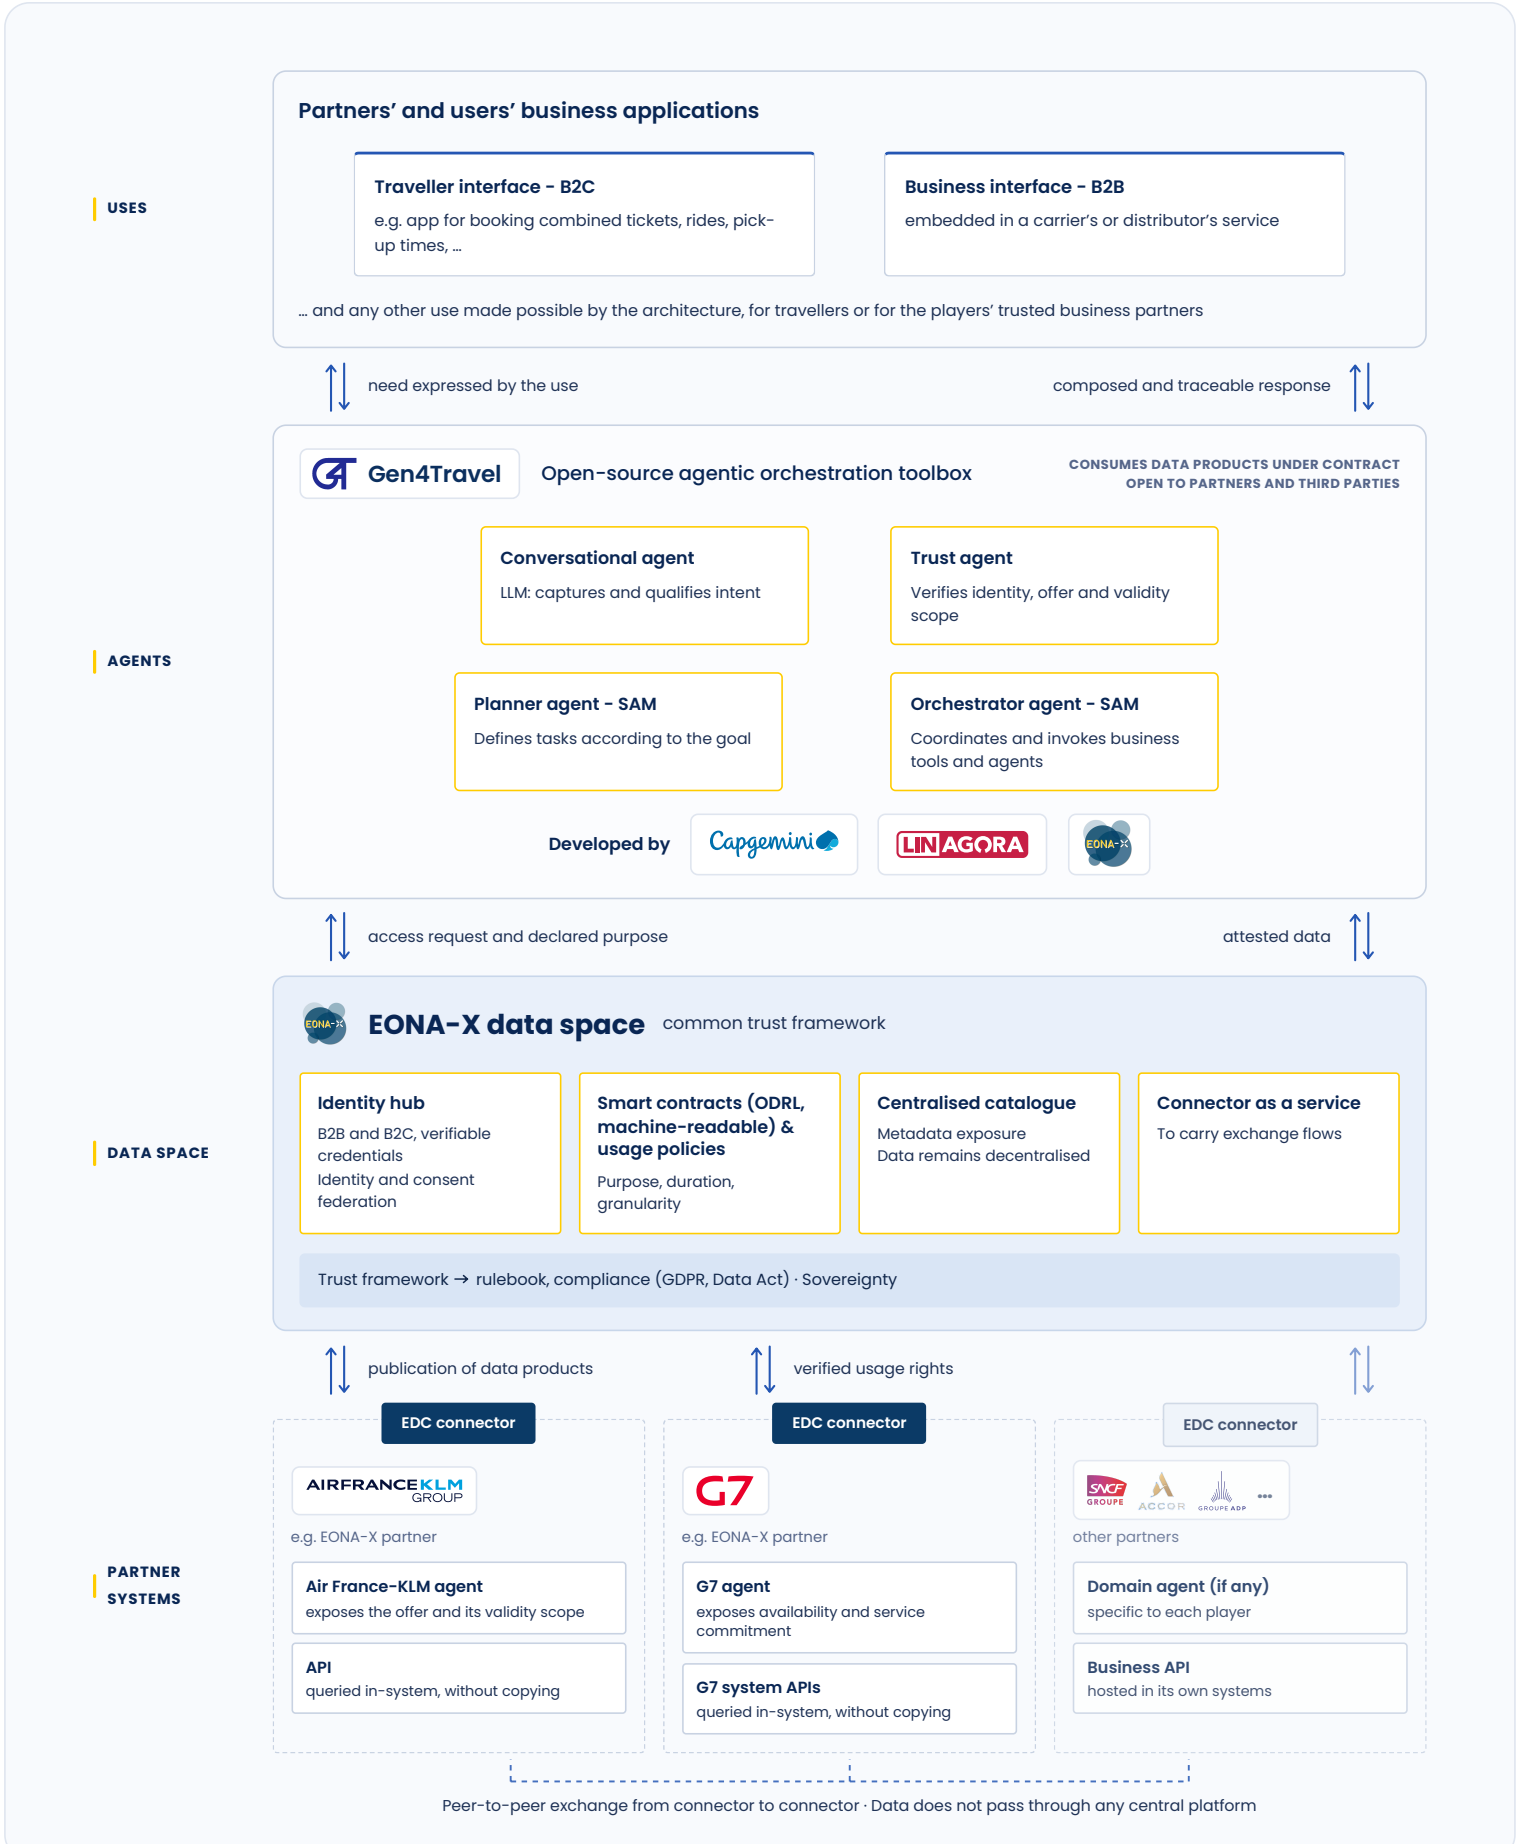
The orchestration layer coordinates several partners without centralising data, the customer or value. It assembles a journey and triggers calls according to permissions, with each partner directly executing its part of the transaction. It is designed to coordinate, never to position itself permanently between the provider and its customer.

Shared data foundation and audit layer

The shared data foundation manages fresh data, provenance, access rules, consents and logs; it provides the common substrate on which the other building blocks rely. Finally, the audit and compliance layer traces interactions, audits recommendations, detects self-preferencing, limits hallucinations and ensures regulatory compliance. Together, they make the framework verifiable from end to end.

The three use cases that follow instantiate this architecture. Each is preceded by its reference diagram, presented full-page, and then detailed according to a consistent structure: summary, problem, vision, target process, key capabilities, anti-disintermediation mechanisms, questions to be decided and recommendations.

Figure 3: *Example of a functional architecture for the orchestration and trust management of AI agents within the EONA-X Data Space*



6. Governance and operating model

How to organise trust

Roles and responsibilities

Governance distinguishes what must be common from what must remain federated. It clarifies the roles of EONA-X, business partners, technology partners, research and audit players, and institutions. This allocation is the condition for a system that is both interoperable and respectful of each party's autonomy.

What must be common

The agent registry and authorisation rules, the attested offer repository and its minimum formats, traceability rules and audit principles, as well as the minimum exchange formats between partners, must be pooled. These are the elements whose value depends on shared adoption and which guarantee interoperability.

What must remain federated

By contrast, information systems, commercial data and the customer relationship, pricing policies and operational execution, as well as internal technology choices, must remain under each partner's control. Each player's sovereignty over its assets and its customer relationship is never delegated to the common system.

EONA-X's role

EONA-X acts as guarantor of the common framework: trust, governance, standardisation, facilitation and coordination of implementation. It is the custodian of the framework, not an exclusive commercial operator. Its neutrality is the guarantee that orchestration does not turn into capturing intermediation.

Role of participants

Partners expose offers, define their business rules, identify their APIs and contribute to the use cases. They retain control over their customer relationship, their commercial terms and the arrangements for exposing their services.

Technical access to an API does not amount to commercial authorisation. Each provider retains its distribution contracts, its eligibility conditions, its security policies and the rights granted to its partners. The trust layer can verify the identity, permissions and mandate of the player requesting a service, but it replaces neither existing contracts nor the authorisations required to book, issue or modify a service.

Role of technical, research and audit partners

Technical partners prototype the building blocks: agent registry, agent-to-agent orchestration, offer attestation, connectors, logs and auditability. Research and audit players define evaluation methods covering fairness, hallucinations, bias, self-preferencing, robustness and traceability, thereby equipping trust through measurement.

7. Benefits and strategic value

What the framework brings to the ecosystem

For travellers

Travellers benefit from a smoother journey, more reliable offers and better personalisation. Above all, they gain in transparency: they know where recommended offers come from and retain control over their consent, which strengthens their trust in the use of agents.

For providers

Providers preserve their visibility, see their brand cited and maintain a direct customer relationship. Their commercial rules remain enforceable throughout the agentic chain, which prevents them from being reduced to mere executors behind an interface they do not control.

For EONA-X and the European ecosystem

For EONA-X, the framework consolidates its positioning as a trust infrastructure and European benchmark for sovereign agentic distribution. For the ecosystem as a whole, it offers a credible alternative to closed platforms, strengthens sovereignty, guarantees interoperability and opens up the capacity to scale.

For technology partners

Technology partners gain a concrete testing ground for AI agents, orchestration, data spaces and verifiable identity. They can thus showcase their expertise within a sovereign framework and help shape tomorrow's standards.

8. Conditions for collective success and recommendations

Final recommendations

Following the work carried out, the group makes eight recommendations to enable the development of open, interoperable and controlled agentic distribution.

1. **Place identity, mandate and action rights at the heart of the system**

Identifying the agent is not enough: the framework must also make it possible to know on whose behalf it is acting, on what basis, for what purpose and within what limits. Rights to read, recommend and act must be explicit, verifiable and revocable.

2. **Keep offers, prices and inventories at the source**

Providers must retain control over their data and their commercial terms. Agents access only the information needed for the stated intention, with guarantees of freshness, provenance and attribution of the offer.

3. **Preserve both provider control and the traveller's freedom of choice**

Commercial rules must remain enforceable without leading to restrictions on the user's choice. Recommendations must be explainable and auditable, and any bias or self-preferencing mechanisms must be detectable.

4. **Standardise the exchange framework rather than freezing a technology**

Interoperability must rest on common elements: identity, permissions, mandate, offer semantics, ontologies, usage rules, provenance, traceability and error handling. Protocols and technical mechanisms may evolve without calling this common foundation into question.

5. **Align orchestration with existing distribution and booking rules**

Technical interoperability does not, in itself, create a right of transactional access. Providers' distribution contracts, authorisations, security rules and commercial terms remain applicable. Depending on the scenario, inventory access, booking, ticketing and payment may fall to separate building blocks or players.

6. **Make security, resilience and reversibility demonstrable**

The framework must be able to keep operating if a model, cloud, orchestrator or connector becomes unavailable or changes. These capabilities must be tested, as must the management of volumes, duplicate executions, anomalies and recovery mechanisms.

7. **Establish neutral, verifiable and economically sustainable governance**

Access, routing and audit rules must be transparent and verifiable. The common system must become neither an exclusive channel nor a new source of capture of the customer relationship. Its business model must fund the common functions without recreating a platform rent.

8. **Demonstrate the framework end to end before extending it**

Gen4Travel can serve as a testing ground by bringing several partners together around a real journey, from the traveller's intention through to booking and then, where relevant, to modification or disruption management. Scaling up must then be guided by measurable results: service quality, reliability, adoption, integration costs, resilience and value created.

Outlook

The framework can only be sustainable if it funds connector maintenance, orchestration, audit, governance and the evolution of standards. The business model must nevertheless avoid reproducing that of the platforms the system precisely seeks to move beyond. The first avenues discussed (affiliation, lead-based remuneration, a compensation mechanism or co-investment) must be assessed on the basis of the value actually created: additional sales, lower distribution costs, access to new partners, improved data quality, reduced friction and preservation of the customer relationship. In this context, the Myerson value provides a useful analytical framework. It makes it possible to remunerate not only direct commercial contribution, but also the functions that make cooperation possible: interoperability, trust, governance, consent management, data quality and network access. It can thus help to assess objectively the value of EONA-X as a relational infrastructure and trusted third party. At this stage,

however, it does not constitute a definitive allocation key. The hypotheses will need to be tested on the use cases and then compared with the actual costs and effective contributions of the partners.

Conclusion

Agentic AI is set to transform the conditions of market access in mobility and tourism. The challenge for European players is not to oppose it, but to define collectively the rules that enable AI agents to act within a reliable, traceable, interoperable and sovereign framework. The roadmap is concrete: identify agents, attest offers, coordinate journeys and build on Gen4Travel to demonstrate value, before gradually extending the framework to the European level. The priority is no longer to produce an exhaustive institutional document, but to give partners a decision-support tool and a shared direction.

Annexes

Glossary and regulatory overview

Annex 1: Glossary

Term	Definition
AI agent	Autonomous system acting on behalf of a user: searching, comparing, recommending, booking, modifying.
AI assistant	Conversational or application interface through which the user expresses their intention, receives the results and validates actions. It may mobilise one or more specialised AI agents.
Agentic distribution	Distribution model in which the customer's agent queries and solicits providers.
Customer-to-business	Model in which it is the customer's agent that queries providers, and not the other way round.
Attested offer	Offer with verifiable provenance, freshness and exposure rules.
Agent registry	Directory identifying agents, their publishers, their rights and their trust levels.
Orchestration	Coordination of a multi-party journey without centralising data or the customer.
Data space	Trust framework for sharing and using data between players, without a centralised database.
Self-preferencing	Bias whereby an intermediary favours its own offers or partners.
Model-independent	Property of a system that works regardless of the underlying AI model.
Travel agency / distribution intermediary	Business player distinct from the AI agent, which may be involved in selling, booking or assembling services depending on its role and the applicable legal framework.
Traveller's mandate	Proof that an agent is acting on behalf of a traveller or an organisation for a given purpose and duration, with authorised operations and explicit, revocable limits.
Reversibility	Ability to change a technical component or a provider without loss of rights, connections, traceability or operational continuity.

Annex 2: Regulatory overview (key references)

Instrument	Relevance to the framework	Status (indicative)
AI Act	Risk-based approach; transparency, governance and oversight of AI systems	Entered into force on 1 August 2024; largely applicable since 2 August 2026, with certain later deadlines
Data Act	Data access and sharing; conditions of use	Applicable since September 2025
Data Governance Act	Framework for European data spaces	Applicable since 2023
eIDAS 2.0	European digital identity, identity wallet (EUDI Wallet), authentication and electronic trust attestations	In force since May 2024; gradual roll-out of EUDI Wallets from 2026
GDPR	Personal data protection, minimisation, consent	In force
DMA	Regulation of large platforms (gatekeepers)	In force
LOM (Art. 28) & mobility data	Opening up and sharing of mobility data	National and sector-specific framework

Acknowledgements to Contributors

This strategic note is the result of collective work carried out within the working group “Sovereignty of the distribution of mobility and tourism services in the age of AI agents”, led by EONA-X.

EONA-X warmly thanks all the partners who contributed to the working group’s various workshops. By sharing their expertise, their experience and their business and technological perspectives, their contributions enriched the discussions and progressively built the analyses and directions presented in this note.

Accor
- Gilles Estord

Air France-KLM
- Virgile Boëssé
- Maxime Boussard
- Sixtine Thirion

Beti
- Pierre Chehwan

Capgemini
- Emmanuelle Brodard
- Nicolas Brunel

Digital New Deal
- Olivier Dion
- Arno Pons

G7
- Claude Boinnot
- Zineb Essahli

Inria
- Sylvain Gitton

Instant System
- Mohammed-Amine Ait Ouahmed
- Clémence Castell
- Yann Hervouet
- Xavier Lecomte

LINAGORA
- Benjamin Bellamy
- Florian Mondon
- Andrzej Neugebauer

SNCF
- Angélique Bidault-Verliac

EONA-X
- Titouan Fortin
- Simon Girardeau
- Nicolas Karageuzian
- Mélanie Veissier

We also thank all the representatives of member organisations who took part in the discussions, the preparatory work and the various stages of consolidating this strategic note.

EONA-X
21, avenue de Messine
75008 Paris
contact@eona-x.eu
eona-x.eu