


Standards-grounded agents

TM Forum, MCP and governed execution — how a loop knows what it is looking at.

Connecting an agent to a BSS estate is an integration problem, and telecom has spent decades learning how hard those are. The harder problem sits on top of it: an agent that reaches exactly what it is permitted to reach, understands what it finds, and can prove afterwards why it acted.



Exosphere sits between the agent platform and the systems of record, inside the customer-controlled environment. The systems of record stay authoritative for their own data.

In short

Industry standards provide the common language. Customer context makes it operational. The tool gateway decides what any of it is allowed to touch.

This paper covers the three of those in order, then the two things that make the result reviewable: where evidence comes from, and where execution stops.

Three layers of understanding

Every loop inherits all three. They are separated because they change at different rates and are owned by different people — the first by the industry, the second by you, the third by the loop's designer.

Layer	Contributes	Owned by
Industry foundation	Shared telecom meaning — information concepts, process context, modular architecture, standards-aligned access	TM Forum
Customer operating model	Products, contracts, policies, processes, terminology, system relationships	The operator
Loop context	Relevant entities and relationships, applicable rules, required evidence, authority, permitted actions	The loop

The standards foundation

TM Forum contributes four distinct things, and conflating them is a common source of confusion in agentic proposals.

SID	The shared information model — what a product, a service, a customer and a party actually are, and how they relate.
eTOM	Business process context — where a given activity sits in the operator's process landscape.
ODA	The modular operating architecture — how capabilities are bounded and composed.
Open APIs	Standards-aligned access to the systems that already run, rather than a bespoke integration per estate.

Alongside these sits plain industry knowledge: the terminology and practices that are never written down in a specification but are assumed by everyone who works in the domain.

“This is connected meaning, not a document-search repository.”

That distinction is worth being precise about, because it sets expectations correctly. Retrieval over documents is how a loop *cites* the industry. The ontology is how it *relates* your products, contracts and policies to one another. A system that only does the first can quote a tariff; it cannot tell you which tariff governs this account, which order changed it, or whether the resulting charge is permitted.

Where Exosphere sits

The stack on the cover shows the arrangement. The agent platform is Microsoft's: Copilot Studio for low-code authoring and for publishing into the surfaces people already use, and Microsoft Foundry for the models, the agent service and the tool catalogue behind them. Exosphere provides telecom meaning and governed execution between that platform and the systems of record. The boundary is deliberate: it is what allows the agent layer and the BSS estate to evolve independently of each other.

One clarification, because these names are used interchangeably more often than they should be. Microsoft Fabric is an analytics platform — the unified data and reporting estate built on OneLake — and Microsoft Foundry is the model and agent platform. Fabric is adjacent to a loop rather than inside it: the evidence a loop reasons over comes from the systems of record through the tool gateway, not from a reporting lake. Where Fabric becomes relevant is measurement — how the outcome a loop moved is reported afterwards, against the baseline the business already uses.

State the dependency plainly, because it is a procurement question as much as an architectural one: this arrangement assumes a Microsoft agent platform. Exosphere is the telecom layer, not the agent runtime. Where an operator has already chosen that platform the boundary is an advantage; where it has not, that is a conversation to have at the start rather than at integration.

MCP and the tool gateway

Agents are good at reasoning and poor at restraint. The gateway is the component that decides what a given loop may actually reach, and Model Context Protocol is how that catalogue is *offered* to the agent ecosystem, alongside the Open APIs, events and adapters already in place.

It is worth being precise about the division of labour, because the two get conflated. MCP is a way to expose tools, resources and prompts to an agent host. It carries no opinion about which loop may use which tool, no read/write permission model and no attribution. Those are decisions the gateway makes; the protocol simply carries the result. The four properties below are gateway properties, not protocol properties — which matters, because swapping the transport would not change any of them.

- **Capability is granted per loop, not per platform.** A collections loop cannot reach what a quoting loop uses, because it was never given it.
- **Reads and writes are separate permissions.** Retrieving a bill and issuing a credit are not the same grant, and no loop holds the second by accident.
- **Systems of record stay authoritative.** Exosphere proposes and records; it does not take custody of data that belongs elsewhere.
- **Every call is attributable** — which loop, under which authority, on whose behalf, against which system.

The practical consequence is that the blast radius of a new loop is knowable before it runs, and reviewable afterwards, without reading the prompt.

Where evidence comes from

A *cited* answer is not taken from a model's training data. Any model brings general telecom understanding with it and that cannot be switched off; what is controlled is the evidence a loop puts its name to, which is retrieved from the published record and traceable back to it. The reference estate spans several families:

- **TM Forum** — Open API REST and component specifications, the ODA Canvas reference implementation, the framework assets — Business Process Framework (eTOM), Information Framework (SID) and Application Framework (TAM) — and the published industry guides.
- **Beyond TM Forum** — CAMARA network APIs, MEF LSO, ETSI NFV and MEC, Broadband Forum USP and CWMP, 3GPP 5G core and SA5 charging and management, and IETF and OpenConfig YANG. ONAP sits alongside these as an open-source orchestration platform rather than a specification, carried because estates are implemented against it.
- **Charging and the commercial layer** — converged charging and CDR handling, mediation and rating concepts, and public product and API documentation from commercial BSS and billing vendors.
- **Operational systems** — ERP, CRM, helpdesk and service-management documentation, because a billing dispute rarely stays inside billing.

Two constraints govern how this is used. Licensed and member-gated material is used for grounding and citation and is never republished. And on the prospect-facing surface, retrieval returns cited excerpts with source links and no generated answer at all — a response carrying one is rejected rather than displayed, and every returned source is re-checked against the approved corpus before it reaches the page.

The output contract

Every loop response is structured identically. This is what makes agent reasoning reviewable by someone who was not watching when it ran.

Section	Contains	Governance check
Facts	Account, bill, payment, order, usage or service facts actually retrieved	No unsupported facts; gaps named
Assessment	Case classification, root cause or anomaly	Confidence and missing data visible
Recommendation	Next best action and its alternatives	Policy and approval thresholds applied
Action	Only permitted steps drafted or executed	Human gate for high-impact actions
Audit	Reason code, data sources, proposed or actual outcome	Traceable decision

Where execution stops

Disconnection, large credits, write-offs, payment-plan exceptions and invoice holds stop at an approval gate by design rather than by configuration someone has to remember to switch on. A declined approval is recorded with its reason and an alternate route, which makes refusals as auditable as approvals — usually the more

interesting half.

Autonomy is then widened deliberately, loop by loop, where the evidence supports it. Nothing about the architecture requires a loop to act without a person; what it requires is that the choice be explicit.

What this is not

- It is not a replacement for your BSS. The systems of record remain authoritative and unchanged.
- It is not a data lake or a document search product. Retrieval supports citation; the ontology carries meaning.
- It is not a general-purpose assistant with access to everything. A loop holds the capabilities it was granted and no others.

Exos Systems is a TM Forum member and aligns to ODA, and operates to ISO/IEC 27001:2022. These describe the context we work in and do not imply endorsement by those bodies.

Demonstrations run on synthetic accounts. The loops referenced are examples of possible outcomes rather than a standard catalogue.

Companion paper: *Pick a loop. Improve an outcome.* — an executive briefing on the operating model and choosing a first loop.