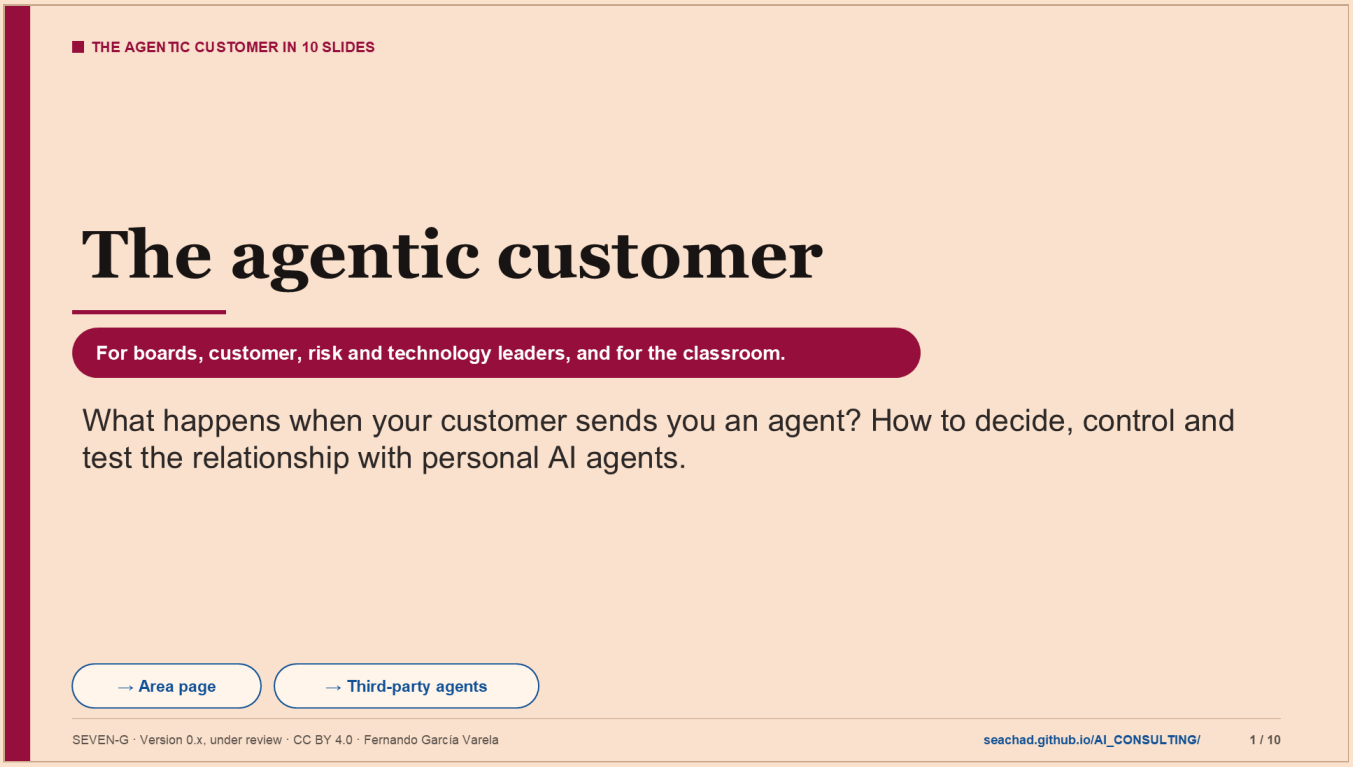


Notes

This presentation explains the SEVEN-G area "The agentic customer" in ten slides. It is meant for boards, customer service, risk and technology leaders, and for the classroom of a business school.

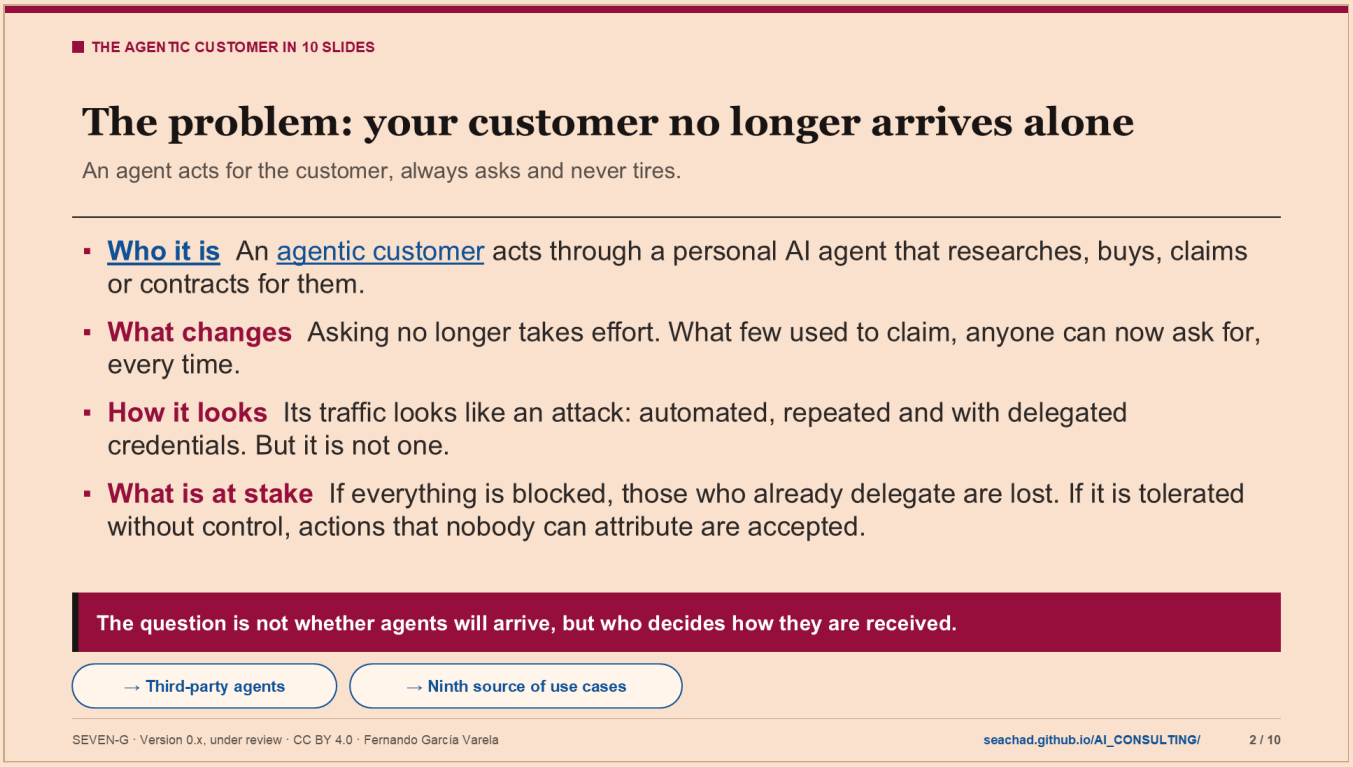
An agentic customer is a customer who acts through a personal artificial intelligence agent: an assistant that researches, buys, claims or contracts on their behalf. The presentation walks through the decision this forces. First, the problem and the four faces of the agentic customer. Then, the posture by journey: block, tolerate or serve through an identified channel. Then, what that channel requires, the split of responsibility proposed by the European financial sector, what regulation does and does not say, the economics of concessions, tests with synthetic customers and, at the end, how to start.

Everything said comes from the SEVEN-G documents: 35 (AI and agent security), 13 (thesis and risk appetite), 40 (value measurement), 33 (risks), 34 (regulation), 52 (operations), 54 (organisation) and the adversarial testing template. The example figures on the site are fictitious. Agent protocols are drafts and are cited only as guidance.

Question for the classroom: does your company know today how many of its customer contacts are already made by an agent?

SOURCES AND LINKS

- **The agentic customer area page** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/areas/cliente-agentico.html
- **Document 35 · Third-party agents acting for the customer** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/35_SEVEN-G_Seguridad_de_IA_y_agentes.html#9-4-third-party-agents-acting-for-the-customer
- **Document 13 · Posture towards customer agents** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/13_SEVEN-G_Tesis_de_IA_ambicion_y_apetito_de_riesgo.html#5-6-posture-towards-customer-agents-by-journey
- **Area page** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/areas/cliente-agentico.html
- **Third-party agents** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/35_SEVEN-G_Seguridad_de_IA_y_agentes.html#9-4-third-party-agents-acting-for-the-customer



Notes

An agentic customer is a customer who delegates to a personal AI agent. The agent researches products, compares, buys, asks for service and claims. It does so on the customer's behalf and with the permissions the customer gives it.

Why this is a governance problem and not only a technology one. Until now, friction protected the company: asking for a refund, a discount or a concession took time, and for that reason many people entitled to one did not ask. An agent removes that friction. The gap between who is entitled to a benefit and who asks for it closes, and the company bears the cost.

An agent's traffic looks like an attack: it is automated, repeated and carries delegated credentials. That is why section 9.4 of SEVEN-G document 35 separates this case from the one in section 9.1, where agents are a threat. Here they act legitimately for the customer.

In document 54, the customer's agent is the ninth source of use cases: it arrives as traffic, not as a proposal, and nobody asked for it or sponsors it. That is why a business owner is needed to set the posture and answer for it.

Why it matters: if the company does not decide, each team decides on its own, and the consequences show up in the income statement.

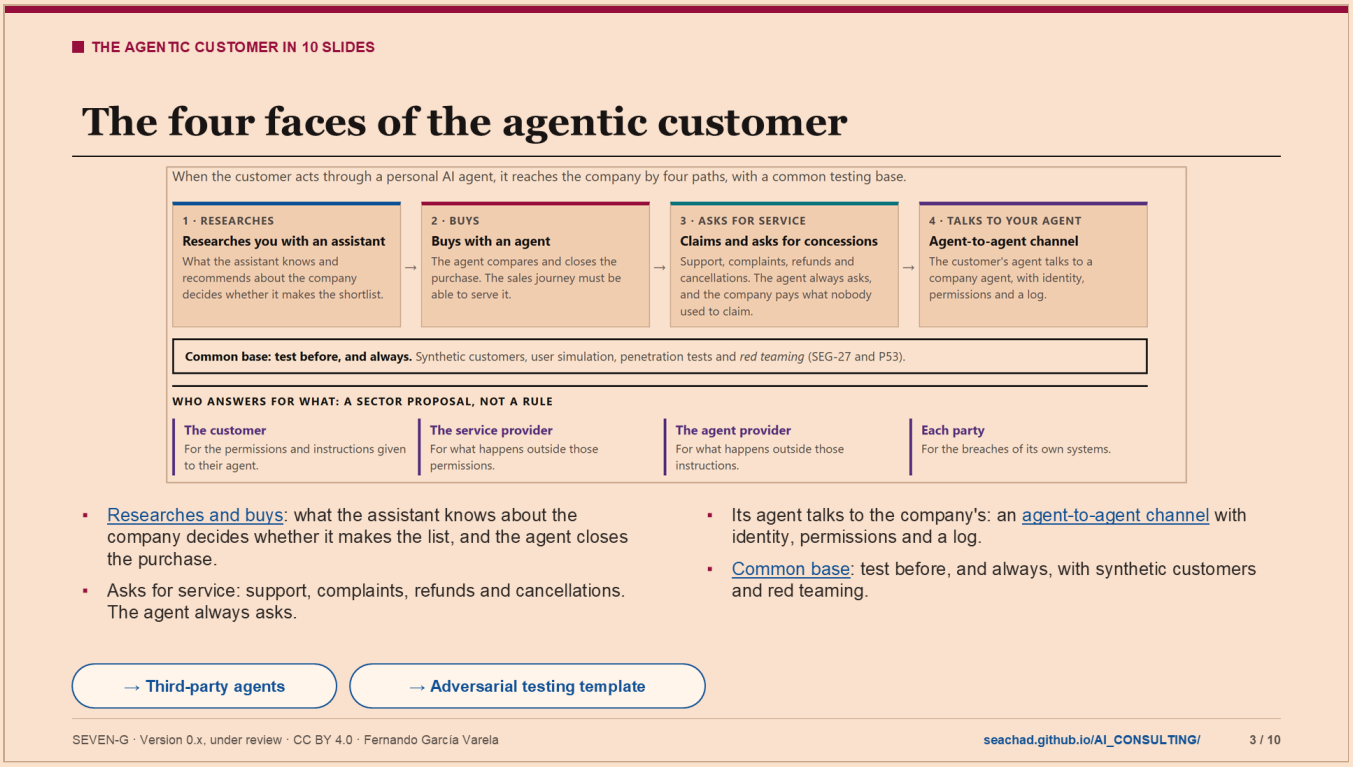
Question for the classroom: which benefits or concessions in your company exist today only because few customers claim them?

TERMS ON THIS SLIDE

- **Agentic customer**: A customer who acts through a personal AI agent: an assistant that researches, buys, claims or contracts on their behalf. It reaches the company by four paths: it researches with an assistant, buys, asks for service and its agent talks to the company's. The link opens the section of document 35 that covers it. https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/35_SEVEN-G_Seguridad_de_IA_y_agentes.html#9-4-third-party-agents-acting-for-the-customer

SOURCES AND LINKS

- **Document 54 · Nine sources of use cases** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/54_SEVEN-G_Organizacion_de_la_compania_para_la_IA.html#5-1-nine-sources
- **Bain & Company: When Your Customer Sends an Agent (10 Nov 2026)** https://www.bain.com/insights/when-your-customer-sends-an-agent/
- **Ninth source of use cases** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/54_SEVEN-G_Organizacion_de_la_compania_para_la_IA.html#5-1-nine-sources



Notes

The chart shows the four paths by which an agentic customer reaches the company, and the common base that supports them. It is a SEVEN-G model of its own.

First face, researching. The customer asks an AI assistant and it answers with what it knows. If the company does not appear, or is poorly described, it does not make the shortlist. This is a matter of visibility and of the quality of the information the company publishes; it has no security control of its own, but it does need a business owner.

Second face, buying. An agent compares and closes the purchase for the customer. The sales journey must be able to serve an agent: prices and terms readable by a machine, identification and limits on what the agent may buy.

Third face, asking for service. Support, complaints, refunds, concessions and cancellations. This is where the economics change most, because the agent always asks.

Fourth face, the agent-to-agent channel. The customer's agent talks to a company agent. Verifiable identity, specific permissions and a log of every action are needed.

The common base is to test before, and always. Synthetic customers, user simulation, penetration tests and red teaming are used, with the corporate test bench and the adversarial testing template.

Why it matters: the four faces are not managed by the same person. Seeing them together prevents each department from solving only its own piece.

Question for the classroom: which department in your company is responsible today for each of the four faces?

TERMS ON THIS SLIDE

- **Agent-to-agent channel** (SEG-28): Channel through which a customer's personal agent talks to a company agent, with verifiable identity, specific and revocable customer permissions, limits, a log of every action and traffic

tagging. It is the "serve" option of the agent posture. The link opens the section of document 35. https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/35_SEVEN-G_Seguridad_de_IA_y_agentes.html#9-4-third-party-agents-acting-for-the-customer

SOURCES AND LINKS

- **Adversarial testing template · Synthetic customers** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/plantillas/P53_SEVEN-G_Plan_e_informe_de_pruebas_adversarias.html#5-3-synthetic-customers-and-user-simulation
- **Bain & Company: When Your Customer Sends an Agent (10 Nov 2026)** https://www.bain.com/insights/when-your-customer-sends-an-agent/
- **Adversarial testing template** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/plantillas/P53_SEVEN-G_Plan_e_informe_de_pruebas_adversarias.html#5-3-synthetic-customers-and-user-simulation



Notes

The agent posture is a decision taken by a company with direct exposure to customers, journey by journey: sales, service, concessions and cancellations. The board takes it with management, together with the AI thesis and the risk appetite, and it is reviewed every year.

There are three options and none is free. Blocking requires detecting the agent and not serving it: the journey is protected, but the business of customers who already delegate to their agent is lost. Tolerating is letting it through as if it were a person: it reaches the human queue at a higher cost, and the concessions it asks for are not measured by anyone. Serving through an identified channel is offering it a channel of its own with identity, permissions, limits and a log: the channel must be designed, tested and maintained, and eligibility and concession rules set outside the AI model.

For the decision to be applicable, agent traffic is detected by signature, where one exists, and by behaviour, and it is tagged. That way its volume, route and cost are measured, and it is seen apart from human traffic and from attacks.

The decision is taken with data: the economic exposure in document 40, the tagged traffic and the appetite for economic, legal and security risk. In SEVEN-G it applies to whoever has direct exposure to customers. It is not part of the core or of the Express intensity.

Why it matters: without a written decision, sales blocks, support tolerates and nobody adds up the cost.

Question for the classroom: which posture would you choose for the cancellations journey, and what data would you need to defend it before the board?

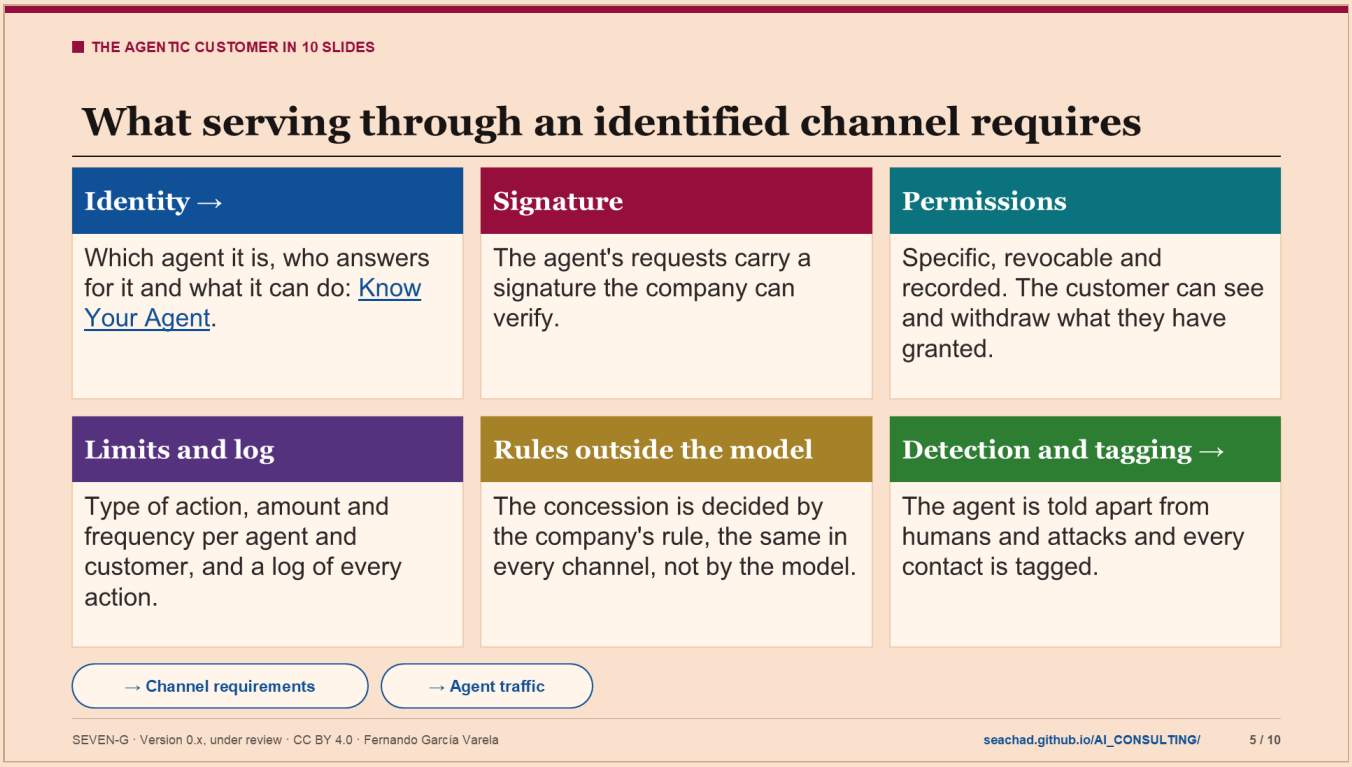
TERMS ON THIS SLIDE

- **Agent posture**: A decision a company with direct exposure to customers takes, journey by journey (sales, service, concessions, cancellations): block the customer's agent, tolerate it or serve it through an identified channel. Each option has its cost. It is required with direct exposure to customers. The link opens the section of document 13. http

s://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/13_SEVEN-G_Tesis_de_IA_ambicion_y_apetito_de_riesgo.html#5-6-posture-towards-customer-agents-by-journey

SOURCES AND LINKS

- **Document 52 · Customer agent traffic** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/52_SEVEN-G_Manual_de_operacion_de_IA.html#customer-agent-traffic
- **Document 40 · Economic exposure** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/40_SEVEN-G_Reglas_de_medicion_del_valor.html#5-4-economic-exposure-to-customer-agents
- **Bain & Company: When Your Customer Sends an Agent (10 Nov 2026)** https://www.bain.com/insights/when-your-customer-sends-an-agent/
- **Agent traffic** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/52_SEVEN-G_Manual_de_operacion_de_IA.html#customer-agent-traffic



Notes

Serving a customer's agent through an identified channel requires eight things, which SEVEN-G gathers into a new security control in document 35. It applies only to use cases with direct exposure to customers.

Identity answers three questions: which agent it is, who answers for it and what it can do. The company does not accept an anonymous agent in a channel with effect. It is the same requirement the framework already places on the company's own agents, applied to external ones.

The signature lets the company verify that requests come from that agent. Permissions are granted by the customer to their agent before the company: they must be specific, revocable and recorded. Limits set which actions, of what amount and how often each agent may take for each customer, and the log records each action with the agent, the customer and the permission applied.

Eligibility and concession rules are applied outside the AI model: the model does not decide which concession is paid, the company's rule does, the same in every channel. Detection and tagging make it possible to measure the traffic and the concessions each agent resolves.

The agent protocols being discussed in the sector today, including bot authentication at the IETF, are drafts. SEVEN-G cites them as guidance and does not use them to underpin any control.

Why it matters: without these elements, a dispute over who authorised an action cannot be resolved. That is why the risk catalogue includes a dispute over an action taken by the customer's agent.

Question for the classroom: with what information could you prove today what a customer authorised their agent to do?

TERMS ON THIS SLIDE

- **Know Your Agent**: Principle by which the company knows which third-party agent acts in a channel, who answers for it and what it can do, before accepting its actions. The link opens the section of document 35. https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/35_SEVEN-G_Seguridad_de_IA_y_agentes.html#9-4-third-party-agents-acting-for-the-customer

SOURCES AND LINKS

- **Document 52 · Customer agent traffic** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/52_SEVEN-G_Manual_de_operacion_de_IA.html#customer-agent-traffic
- **IETF · Web Bot Auth (draft)** https://datatracker.ietf.org/wg/webbotauth/about/
- **Agent traffic** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/52_SEVEN-G_Manual_de_operacion_de_IA.html#customer-agent-traffic



Notes

A clear split of responsibility avoids each party pointing at the others when something goes wrong. The proposal SEVEN-G takes up ties responsibility to control: whoever controls answers.

The customer answers for the permissions and instructions given to their agent, because they decide what to delegate and to whom. The service provider, that is, the company, answers for what happens outside those permissions, because it applies them. The agent provider answers for what happens outside the customer's instructions, because it controls how the agent acts. And each party answers for the breaches of its own systems.

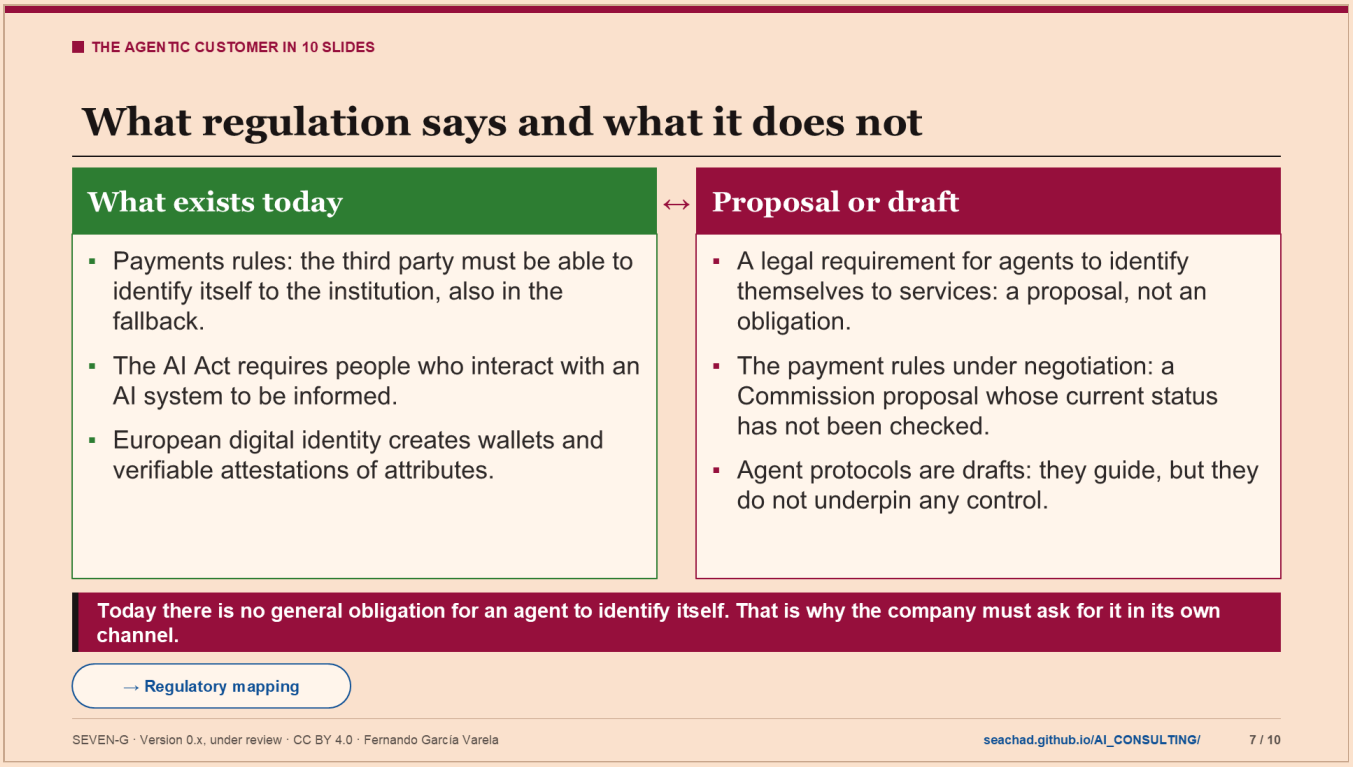
This comes from an opinion piece by a bank chairman published in 2026, which also proposes that the law require agents to identify themselves to services and prohibit impersonation after a transition period, and that an agent's credential be based on European digital identity. These are proposals by its author, not rules. SEVEN-G takes them only as a design criterion for the identified channel and proposes writing the split into the terms of the channel, with legal review.

Why it matters: if the split is not written down, a dispute over an action taken by a customer's agent is settled without evidence. SEVEN-G's risk catalogue includes that case.

Question for the classroom: what would change in your company's terms of use if a customer could delegate to an agent?

SOURCES AND LINKS

- **Document 35 · Who answers for what** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/35_SEVEN-G_Seguridad_de_IA_y_agentes.html#9-4-third-party-agents-acting-for-the-customer
- **Torres Vila: AI agents should not have to pretend to be us (opinion, 2026)** https://www.expansion.com/economia/financial-times/2026/11/13/agentes-ia-deberian-necesitar-hacerse-pasar-nosotros.html
- **Who answers for what** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/35_SEVEN-G_Seguridad_de_IA_y_agentes.html#9-4-third-party-agents-acting-for-the-customer



Notes

Regulation says less about customer agents than many believe. This slide separates what exists from what is a proposal or a draft.

What exists. The European payments rules and their delegated regulation require third-party providers to be able to identify themselves to the institution that holds the account, and require the institution to ensure that they can be identified also in the fallback mechanism. SEVEN-G limits itself to those articles: it does not claim that the rules prohibit or allow screen access in themselves. Article 50 of the European AI Act requires people who interact with an AI system to be informed. It does not require an agent to identify itself to a service. The European digital identity regulation creates the framework for identity wallets and electronic attestations of attributes, and does not regulate the identity of AI agents.

What is a proposal or a draft. A legal requirement for agents to identify themselves to services is a sector proposal. The payment rules under negotiation are a Commission proposal; their current status has not been checked. Agent protocols are drafts. They are cited as guidance and underpin no control, no risk and no gate.

Why it matters: a company that waits for the law to require agents to identify themselves will wait a long time. Meanwhile, it can require identity, permissions and a log in its own channel, which is what SEVEN-G proposes.

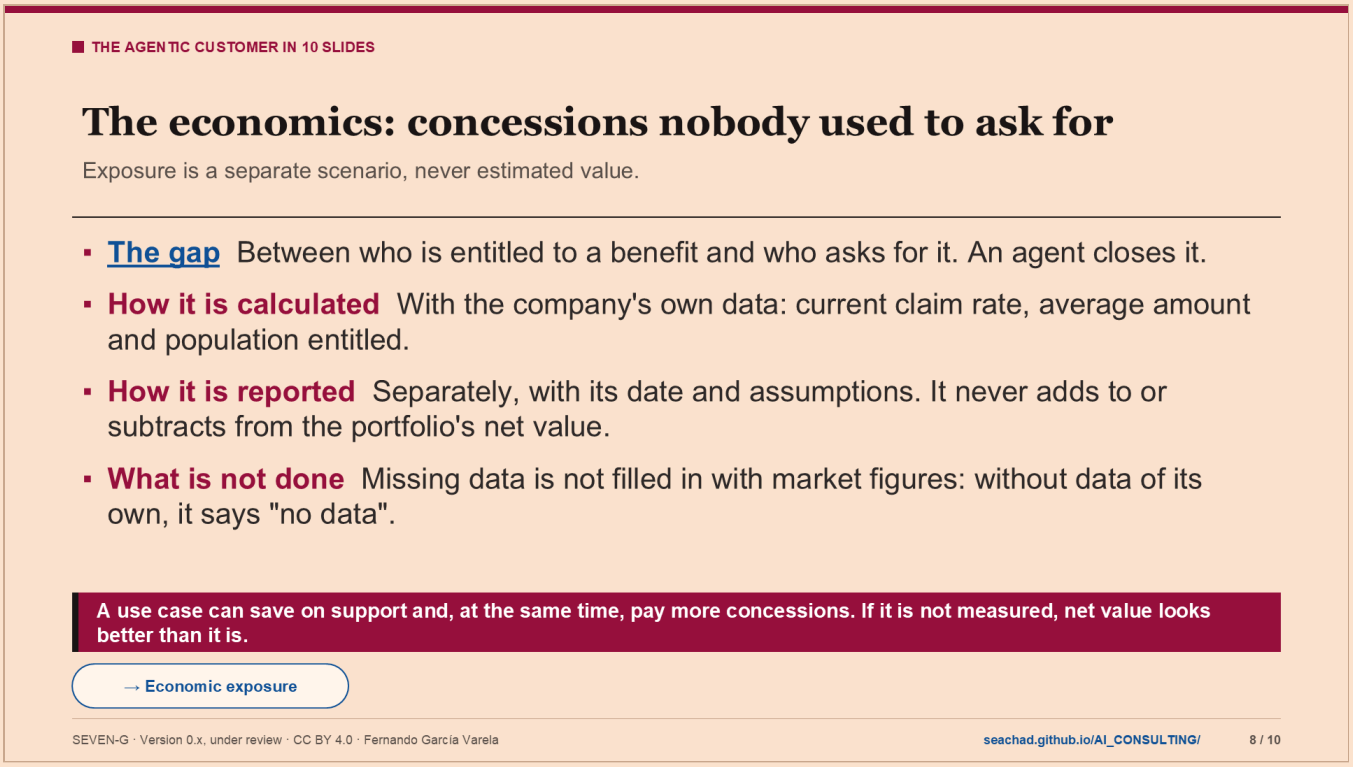
Notice: this slide is not legal advice. Each organisation must verify the rules that apply to it.

Question for the classroom: which part of your service channel would have to change if the law required agents to identify themselves tomorrow?

SOURCES AND LINKS

- **Document 34 · Regulatory mapping (transparency and annex B)** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/34_SEVEN-G_Mapeo_regulatorio.html#3-11-transparency-art-50
- **Directive (EU) 2015/2366 (PSD2) on EUR-Lex** https://eur-lex.europa.eu/eli/dir/2015/2366/oj/spa
- **Delegated Regulation (EU) 2018/389 on EUR-Lex** https://eur-lex.europa.eu/eli/reg_del/2018/389/oj/spa
- **Regulation (EU) 2024/1183 (eIDAS 2) on EUR-Lex** https://eur-lex.europa.eu/eli/reg/2024/1183/oj/spa

- **Regulatory mapping** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/34_SEVEN-G_Mapeo_regulatorio.html#3-11-transparency-art-50



Notes

This slide covers the economic effect of customer agents. SEVEN-G's value measurement rules say that every amount has a formula and a status, that what cannot be put into euros is reported separately and that "no data" is not zero. Exposure to customer agents applies those rules.

The gap is the distance between who is entitled to a benefit (a concession, a refund, a discount) and who asks for it. Until now friction kept that gap open. An agent closes it because it always asks.

Exposure is calculated as a scenario, with the company's own data: the current claim rate, the average benefit amount and the population entitled to it. The result is the cost of granting what would be asked for if asking took no effort. It is reported separately, with its date and assumptions. It neither adds to nor subtracts from the portfolio's net value, and it is never estimated value.

There is a common error that SEVEN-G expressly records: net value inflated by not deducting the concessions that an agent resolves. A use case that serves customer agents shows the support saving and omits the increase in concessions paid. The correction is to deduct the additional concessions and compare them with the human baseline.

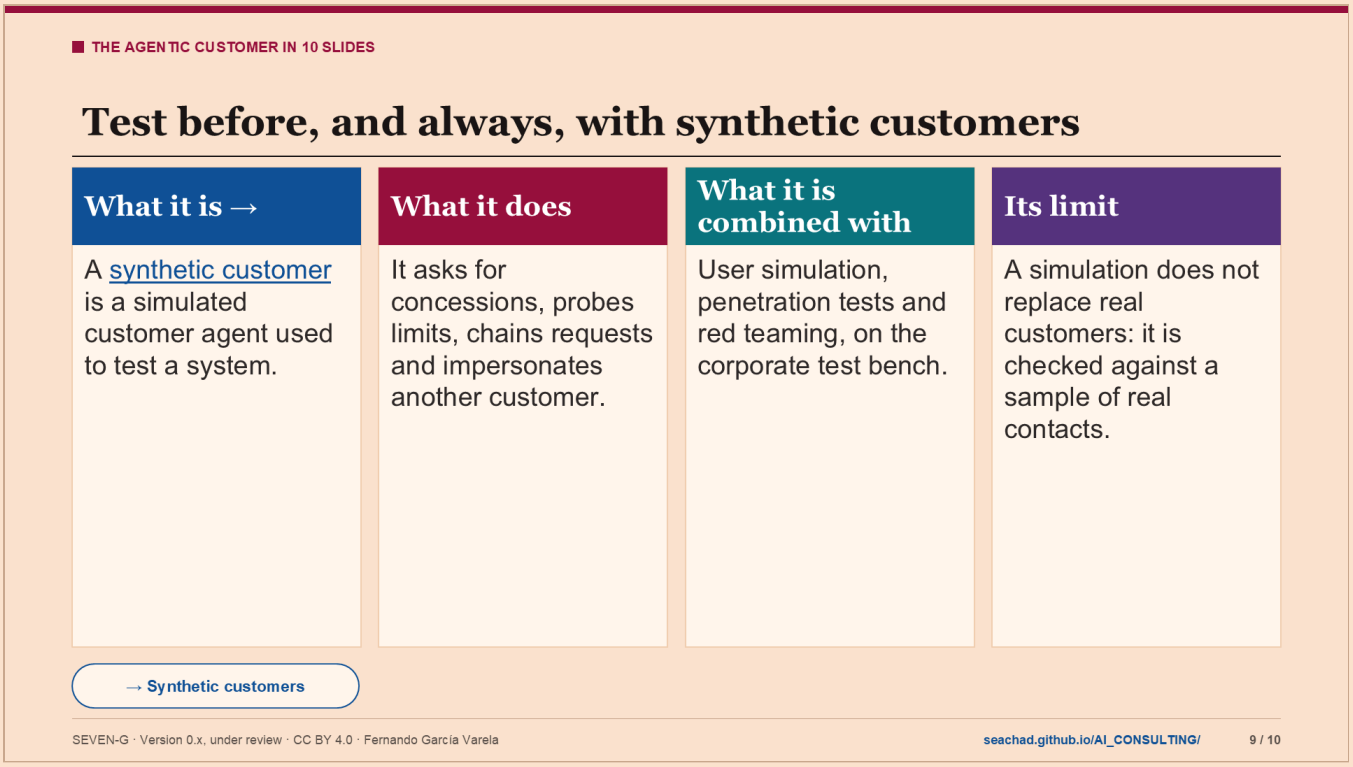
Why it matters: the board sets the posture with this figure in front of it. Without it, the decision is made blind. What then actually happens is measured with an indicator of concessions resolved by agents.

Question for the classroom: what data would you need to estimate how many customers entitled to a benefit do not ask for it today?

SOURCES AND LINKS

- **Document 40 · Economic exposure to customer agents** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/40_SEVEN-G_Reglas_de_medicion_del_valor.html#5-4-economic-exposure-to-customer-agents
- **Document 35 · Indicators for the board** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/35_SEVEN-G_Seguridad_de_IA_y_agentes.html#indicators-for-the-board
- **Bain & Company: When Your Customer Sends an Agent (10 Nov 2026)** https://www.bain.com/insights/when-your-customer-sends-an-agent/

- **Economic exposure** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/40_SEVEN-G_Reglas_de_medicion_del_valor.html#5-4-economic-exposure-to-customer-agents



Notes

Synthetic customers are the way to test the agent channel before real agents arrive. They are simulated customer agents. They ask for concessions, probe the limits of the channel, chain requests to see whether a rule breaks and impersonate another customer.

They are used together with user simulation, penetration tests and red teaming. Scenarios are versioned in the corporate agent test bench and recorded in the adversarial testing template, by journey: how many scenarios succeeded, how many concessions were granted outside the rule, which limits were exceeded and which actions were left without a log. Each finding has an owner and a date, and is retested after correction.

The limit is clear: a simulation does not replace real customers. An academic study with generative agents built from the answers of more than a thousand real people reproduced their survey answers with high consistency. But that is not extrapolated to purchasing behaviour, which is why a simulation is always checked against a sample of real contacts.

The common base of the four faces is to test before, and always: before go-live and after every relevant change.

Why it matters: a channel that has not been tested against agents trying to abuse it is, in fact, tested on customers.

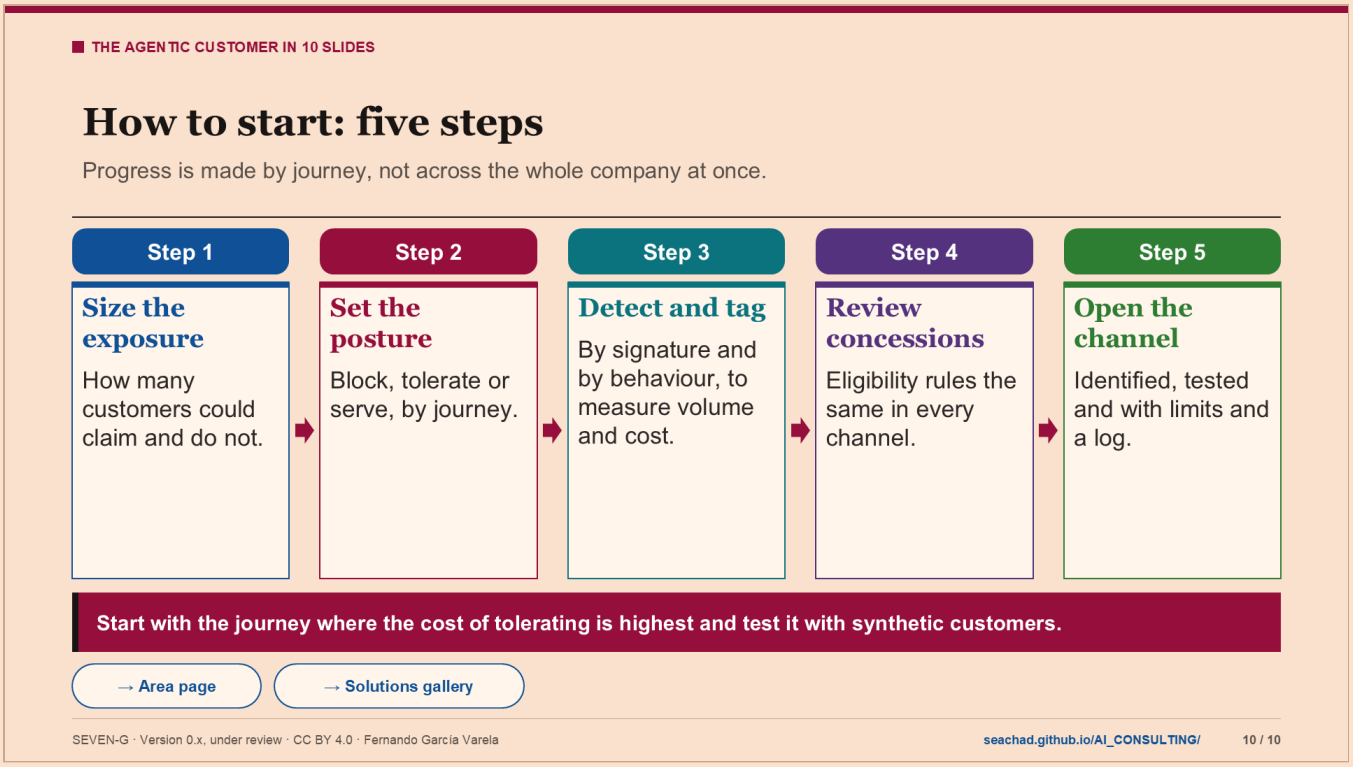
Question for the classroom: which three scenarios would you ask a synthetic customer to try in your company's concessions journey?

TERMS ON THIS SLIDE

- **Synthetic customer**: A simulated customer agent used to test a system before and during its operation: it asks for concessions, probes limits, chains requests and impersonates another customer. A simulation does not replace real customers. The link opens the section of the adversarial testing template. https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/plantillas/P53_SEVEN-G_Plan_e_informe_de_pruebas_adversarias.html#5-3-synthetic-customers-and-user-simulation

SOURCES AND LINKS

- **Document 35 · Common testing base** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/35_SEVEN-G_Seguridad_de_IA_y_agentes.html#9-4-third-party-agents-acting-for-the-customer
- **Park et al. (2024): simulating individuals with generative agents** https://arxiv.org/abs/2411.10109



Notes

These five steps follow the five actions proposed by an international consultancy in a November 2026 study, adapted to SEVEN-G's language and controls. They are a starting guide, not a recipe.

First, size the exposure: estimate how many customers could claim a benefit today and do not, and calculate the cost of handling those claims if asking cost little. Second, set the posture by journey, with the board and management. Third, build detection: identify agent contacts by behavioural signals and by identity verification where it exists, and tag them to follow their volume, route and cost. Fourth, review concession policies: where entitlement exceeds what is actually asked for, what it would cost if more were asked and which discretionary benefits to adjust, applying them the same in every channel. Fifth, open an agent-to-agent channel with an authenticated interface, defined permissions and an action log, routed apart from the human queues and with automatic decisions within limits.

The site's gallery lists market companies and solutions for each face of the agentic customer, with what their official page says. It is a market reading, not a recommendation: before using any of them with company data, even to test it, the supplier assessment applies.

Why it matters: starting with one journey lets the company learn at low risk and with real data before deciding for the whole company.

Question for the classroom: which of the five steps would your company start with, and who would be the business owner of the posture?

SOURCES AND LINKS

- **The agentic customer area page** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/areas/cliente-agentico.html
- **Gallery of companies and solutions** https://seachad.github.io/AI_CONSULTING/SEVEN-G/herramientas/T17_panel_consejo/galeria/index.html#en-cliente-agentico
- **Bain & Company: When Your Customer Sends an Agent (10 Nov 2026)** https://www.bain.com/insights/when-your-customer-sends-an-agent/
- **Area page** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/areas/cliente-agentico.html

- **Solutions gallery** https://seachad.github.io/AI_CONSULTING/SEVEN-G/herramientas/T17_panel_consejo/galeria/index.html#en-cliente-agentico

LEGAL NOTICE

Legal notice: a generic, free methodological aid provided "as is" for information purposes; it is not legal advice and has not been legally reviewed for any case. Each organisation adapts it and is responsible for its legal review and its regulatory compliance; it may not reflect recent regulatory changes. To the fullest extent permitted by law, the author accepts no liability whatsoever for the effects of its application. It does not grant certification of any kind. Details in document 93.

CC BY 4.0 (content) · MIT (code) · Fernando García Varela · seachad.github.io/AI_CONSULTING/