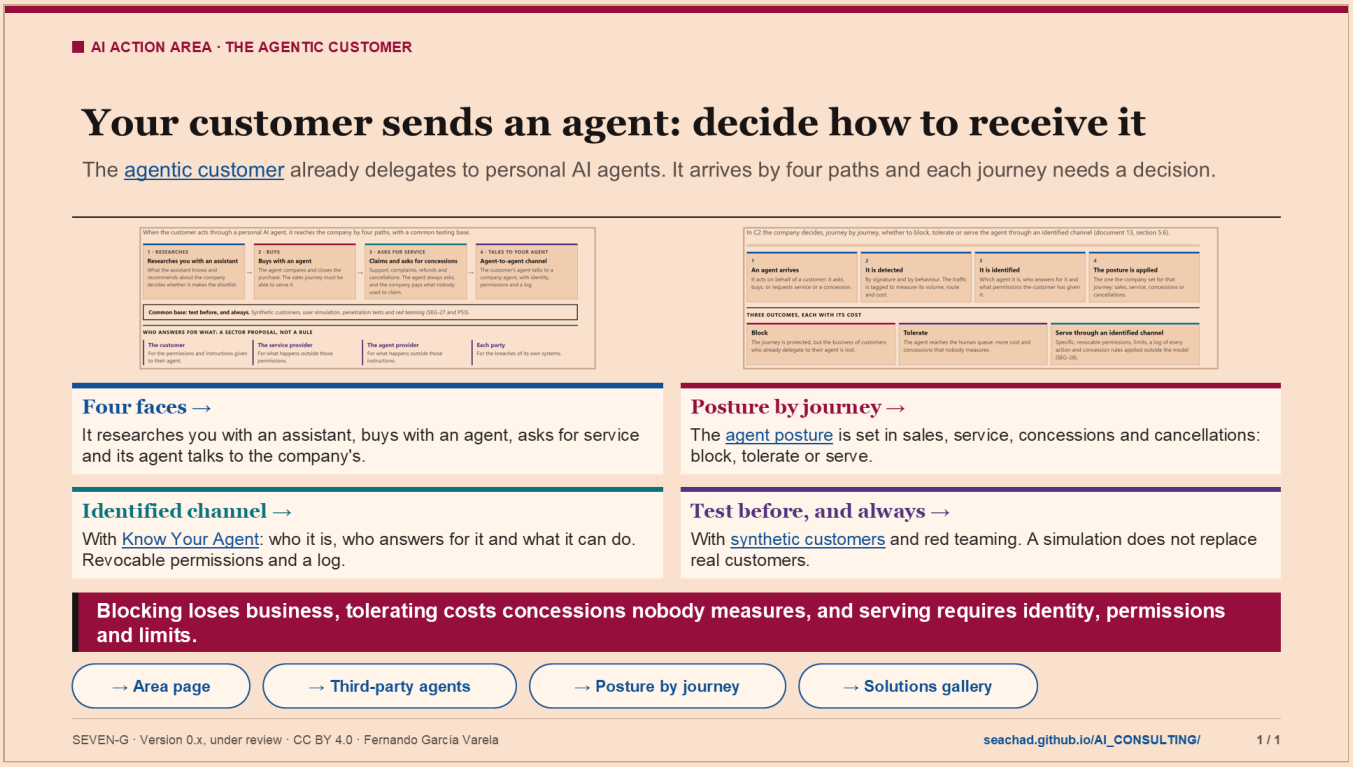


Notes

This slide summarises the SEVEN-G area "The agentic customer". It answers a question that customer service, sales and risk leaders are already asking: what happens when your customer sends you an agent.

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. It reaches the company by four paths. First, it researches you with an assistant: what the assistant knows and recommends about the company decides whether it makes the shortlist. Second, it buys with an agent: the agent compares and closes the purchase, and the sales journey must be able to serve it. Third, it asks for service: support, complaints, refunds and cancellations. The agent always asks, and the company ends up paying benefits that hardly anyone used to claim. Fourth, its agent talks to the company's agent: an agent-to-agent channel that needs identity, permissions and a log.

The central decision is the posture by journey. For each journey (sales, service, concessions and cancellations) the board sets whether the company blocks the agent, tolerates it or serves it through an identified channel. Each option has its cost. Blocking protects the journey, but the business of customers who already delegate to their agent is lost. Tolerating lets the agent reach the human queue, with a higher cost and with concessions that nobody measures. Serving requires designing, testing and maintaining a channel with verifiable identity of the agent and of whoever answers for it, specific and revocable customer permissions, limits, a log of every action and concession rules that are applied outside the AI model. SEVEN-G asks for this posture when the company has direct exposure to customers. It is not part of the core rules or of the Express intensity.

Everything is tested before, and always. Synthetic customers are simulated customer agents that ask for concessions, probe limits, chain requests and impersonate another customer. They are combined with user simulation, penetration tests and red teaming. A simulation does not replace real customers: it is checked against a sample of real contacts.

Why it matters. Customer agents are already arriving, with or without the company's permission. Setting the posture by journey, instead of letting each team improvise, avoids two opposite failures: losing sales by blocking and paying unmeasured concessions by tolerating. In addition, economic exposure is calculated as a separate scenario, with the company's own data, and it never adds to the portfolio's net value.

A note of caution. The split of responsibility that is proposed follows control: the customer answers for the permissions they give, the service provider for what happens outside those permissions, the agent provider for what happens outside the instructions, and each party for the breaches of its own systems. It is a proposal from the European financial sector, not a rule. Agent protocols are drafts and do not underpin any control.

Question for the classroom: in which journey of your company would it be most expensive to tolerate a customer's agent, and in which would it be most expensive to block it?

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
- **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. 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
- **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
- **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

- **The agentic customer area page** https://seachad.github.io/AI_CONSULTING/SEVEN-G/html/en/areas/cliente-agentico.html
- **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
- **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/