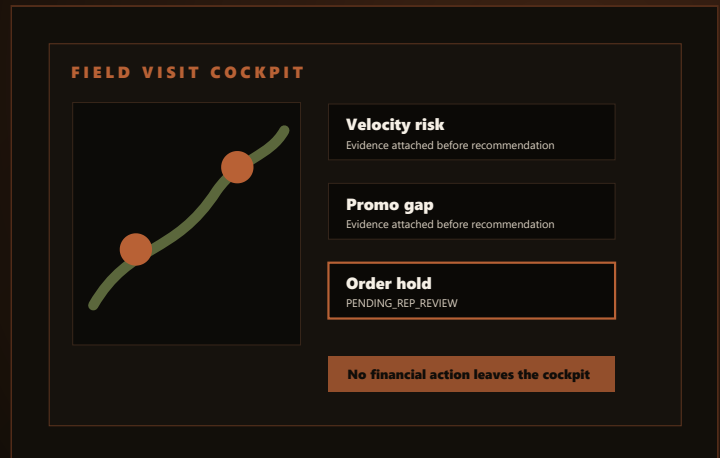


## PHANTOM VSA / RESEARCH BRIEF

# PHANTOM Demo Script

A commercial demo script for showing field-sales value without pretending the agent autonomously places orders.



01 1 visit packet

02 3 next actions

03 0 submitted orders

# Field value depends on the hold, not only the answer.

Retail and CPG teams want faster visit preparation, but order proposals touch commercial risk. The demo must prove workflow value and control at the same time. PHANTOM should be positioned as a governed commercial cockpit: it prepares the rep, explains evidence, proposes next best actions, and stops before any financial commitment.

**01**

The fastest path to trust is showing the hold, not hiding it.


**02**

Synthetic data is enough for demo value if the boundary is explicit.

**03**

The first commercial promise is better preparation and review, not autonomous ordering.

### Field-sales proof loop



The proof is a better rep decision, not autonomous order submission.

**GARTNER / CUSTOMER SERVICE AGENTIC AI PREDICTION**

Customer-facing agentic automation is expected to handle more routine service work, raising buyer expectations for fast, governed experiences.

**BOSTON CONSULTING GROUP / MANAGING THE MACHINES THAT MANAGE THEMSELVES**

Agentic AI behaves partly like software and partly like a colleague, which forces new supervision and decision-rights models.


**PARLOA / AI AGENT LIFECYCLE MANAGEMENT FOR CUSTOMER SERVICE**

Lifecycle, analytics, compliance, and audit logging are part of the production operating model for customer-facing agents.

# Rep-facing proof sequence.

The PHANTOM cockpit turns account, velocity, promo, and planogram signals into a reviewed visit package. Read the artifact as a visit-package storyboard: PHANTOM prepares the rep, explains the evidence, proposes the action, and stops before order commitment.

## Visit package storyboard



**ACCOUNT**

Load account context, visit history, SKU set, and relevant commercial constraints before analysis.

**SIGNAL**

Highlight velocity, stock, promo compliance, and next best actions with traceable justifications.

**PROPOSAL**

Generate a proposed order as a review artifact, not as a committed financial action.

**HOLD**

Stop at rep approval, record audit events, and expose the financial boundary in the UI.

# Show commercial lift and commercial restraint in one flow.

Use the demo to make synthetic-data limits, rep approval, FINANCIAL boundaries, and first KPI baselines visible instead of buried.

## INSPECTION QUESTIONS

- Can the rep see why each SKU action is recommended?
- Does the UI clearly show PENDING\_REP\_REVIEW before any financial action?
- Which source fields would need real client integration after the demo?
- What KPI baseline would prove the workflow improved visit preparation?

## DO NOT SHIP UNTIL

- Do not let the demo imply autonomous order submission.
- Do not hide synthetic-data limits.
- Do not skip the rep approval narrative; it is the trust signal.

<b>DEMO</b>	Run the synthetic visit package and narrate evidence before productivity.	Audit rail, identity panel, visit package, and proposed order hold.
<b>DIAGNOSTIC</b>	Replace synthetic data with a client sample and validate data access.	Source map for accounts, SKU, stock, promo, and planogram inputs.
<b>PILOT</b>	Measure preparation time, rep confidence, suggested order quality, and approval friction.	CLEAR baseline and rep review notes.

# Source ledger and demo boundary

The brief combines external research signals with HARNEXA's implementation boundary: identity, permission, approval, audit, CLEAR evaluation, and no autonomous public execution.

<p><b>GARTNER</b></p> <p><b>Customer service agentic AI prediction</b></p> <p>Gartner predicts agentic AI will autonomously resolve most common customer service issues by 2029 and reduce operational cost.</p> <p><a href="https://www.gartner.com/en/newsroom/press-releases/2025-03-05-gartner-predicts-agentic-ai-will-autonomously-resolve-80-percent-of-common-customer-service-issues-without-human-intervention-by-20290">https://www.gartner.com/en/newsroom/press-releases/2025-03-05-gartner-predicts-agentic-ai-will-autonomously-resolve-80-percent-of-common-customer-service-issues-without-human-intervention-by-20290</a></p>	<p><b>BOSTON CONSULTING GROUP</b></p> <p><b>Managing the machines that manage themselves</b></p> <p>BCG frames agentic AI as both software and colleague, requiring redesigned workflows, governance, roles, and investment logic.</p> <p><a href="https://www.bcg.com/publications/2025/machines-that-manage-themselves">https://www.bcg.com/publications/2025/machines-that-manage-themselves</a></p>	<p><b>PARLOA</b></p> <p><b>AI agent lifecycle management for customer service</b></p> <p>Parloa emphasizes lifecycle management, versioning, analytics, live orchestration, compliance, and audit logging for customer-facing AI agents.</p> <p><a href="https://www.parloa.com/blog/ai-agent-lifecycle-customer-service/">https://www.parloa.com/blog/ai-agent-lifecycle-customer-service/</a></p>
<p><b>RESEARCH STANCE</b></p> <p>External research is cited as market signal; HARNEXA interpretation is kept separate from source claims.</p>	<p><b>BUYER ARTIFACT</b></p> <p>Every report includes a concrete visual artifact that can be inspected in a sales, risk, or architecture review.</p>	<p><b>OPERATING BOUNDARY</b></p> <p>Synthetic proof, rep review, no autonomous order submission. The asset is for governed review, not autonomous production execution.</p>

**HARNEXA BOUNDARY**

Technical governance aid only. Not legal advice or compliance certification. HARNEXA AI builds audit-ready foundations for qualified review. No PDF in this library claims EU AI Act certification, legal advice, payment readiness, or autonomous production execution.