

Kore.ai, delivered to *production.*

Pronix is a Kore.ai reseller and strategic implementation partner for the full lifecycle: adopt pre-built vertical agents with AI for X, migrate legacy XO Platform V1/V2 estates to the Agent Platform, or stand up a greenfield Agent Platform (Artemis) program end-to-end.

From strategy and pilots to production outcomes.

AI for X

XO → Agent Platform migration

Agent Platform (Artemis)

Contact Center AI

Search AI

Managed run

3

Delivery pillars under one certified partner

60–80%

Intent coverage on day one with AI for X

99%+

Regression parity on XO → Agent Platform

0

Downtime with canary cutover and rollback

Kore.ai adoption is no longer one program — it is three.

Kore.ai has evolved from the XO Platform to the Agent Platform (codename Artemis) — an LLM-native, agentic foundation with orchestration, tools, memory and unified governance. Most enterprises need a mix of all three tracks, run under one partner.

1

AI for X moves fastest

Pre-built BankAssist, HealthAssist, RetailAssist, ITAssist, HRAssist and SearchAssist agents cover 60–80% of common intents on day one — integration, tuning and governance are the real work.

2

XO V1/V2 estates cannot lift-and-shift

Legacy dialogs, entities, knowledge graphs and custom scripts must be assessed, mapped and rebuilt on the Agent Platform without breaking live conversations.

3

Artemis is architecturally different

The Agent Platform is LLM-native — agents, tools, memory, evals and guardrails. Greenfield builds need a new blueprint, not an XO clone.

WHAT GOOD LOOKS LIKE

The outcomes these programs are measured on.

3

Delivery pillars under one certified partner

60–80%

Intent coverage on day one with AI for X

99%+

Regression parity on XO → Agent Platform

0

Downtime with canary cutover and rollback

Ranges reflect outcomes observed across Pronix Kore.ai engagements. Your baseline is measured from your own conversation data in the first two weeks, and every figure we report afterwards is a movement against it.

Our position. AI for X, migration and Artemis implementation share governance, integrations and operating model. We run them as three coordinated tracks under one accountable partner so investments compound instead of colliding.

PARTNER STATUS

Kore.ai reseller and strategic implementation partner.

Certified engineering bench

Kore.ai-certified developers, architects and admins across XO Platform, Agent Platform (Artemis), AI for X packs, SmartAssist and AgentAssist.

Licensing and delivery in one place

As a reseller we can transact Kore.ai licensing alongside implementation and managed run — one commercial relationship, one accountable team.

CERTIFIED PARTNER PLATFORMS

Kore.ai delivered alongside the platforms your estate already runs.

kore.ai

GENESYS

NICE
CXone

Amazon Connect

Five9



Microsoft



ServiceNow

OpenAI

Three Kore.ai delivery pillars, one certified team.

Certified Kore.ai engineers deliver each pillar end-to-end — with shared architecture, governance and managed services underneath.

01

Pillar 1 — AI for X · Vertical pre-built agents

Adopt BankAssist, HealthAssist, RetailAssist, ITAssist, HRAssist and SearchAssist with your branding and policies · Pre-packaged flows for account servicing, claims, appointments, order status, IT tickets, HR self-service and enterprise search · Connect Salesforce, ServiceNow, Workday, SAP, Epic, Guidewire and core CRM/ERP/ITSM/HRIS; tune NLU, prompts and retrieval · Guardrails, PII redaction and audit trails applied before go-live

02

Pillar 2 — XO migration · V1/V2 to Agent Platform

Bot inventory, dialog and entity mapping, integration rebuild and knowledge migration with a costed wave plan · FAQs, knowledge graphs and documents moved into the Agent Platform retrieval stack with chunking, embeddings and citation controls · Automated regression suites compare XO responses to rebuilt agents across intents, entities and transactional flows · Canary cutover with rollback plans — targeting 99%+ parity and zero downtime

03

Pillar 3 — Agent Platform · Greenfield Artemis implementation

Agent design, tool framework, task memory, orchestration and model routing across Kore · Multi-agent workflows with human-in-the-loop handoff for voice and digital contact centers · Prompt and response evals, topic and safety guardrails, audit trails and enterprise policy controls · Production-grade observability and eval regression in CI from the first release

SHARED ACROSS ALL THREE PILLARS

Integration and run, engineered once.

CCaaS & channel integration

IVR, web, mobile, Teams, Slack, WhatsApp and contact-center channels — Genesys, NICE, Amazon Connect, Five9, Google CCAI — with agent handoff and context passing.

Managed services & optimization

Post-launch NLU tuning, prompt and eval maintenance, new use-case rollout and 24×7 operations across AI for X, migrated bots and Artemis agents under one SLA.

Governance built into delivery, not bolted on afterwards. Prompt and response evals, PII redaction, topic and safety guardrails, human-in-the-loop on regulated content, retained evaluation evidence per model version and full audit trails — in place before the first intent goes live.

Where Kore.ai pays in the enterprise.

Named programs we deliver across AI for X, migration and greenfield Artemis — each with an evaluation harness, guardrails and audit trail from day one.

AI for Banking

Balance, transactions, disputes, card controls and servicing intents live in weeks with pre-built BankAssist flows integrated to core banking and CRM.

AI for Healthcare

Appointment scheduling, Rx refills, benefits and triage flows connected to EHR, payer and scheduling systems with healthcare-grade guardrails.

AI for Retail / IT / HR

Order status, returns, IT tickets, password resets, HR self-service and policy Q&A launched from pre-built packs and extended to your systems.

XO 11 → Agent Platform migration

Legacy XO 11 bots with rule-heavy dialogs and custom scripts rebuilt as agentic workflows with modern tooling and governance.

Agent Platform V1 → V2 migration

V1 agents mapped, improved and migrated with minimal business disruption and 99%+ regression parity.

Multi-bot consolidation

Overlapping XO bots consolidate into fewer, more capable agents that share tools, knowledge and governance.

Agentic contact center

Greenfield multi-agent workflows for voice and digital with tool-use, memory and human-in-the-loop handoff.

Enterprise search & knowledge agents

Retrieval-grounded agents over SharePoint, Confluence, ServiceNow and document stores with citations and continuous eval.

Internal copilots

Employee-facing copilots for IT, HR, sales and operations with tool-use across CRM, ERP, ITSM and HRIS.

TYPICAL TIMELINES

4–8 wks

Per AI for X agent, launched and integrated

3–6 mo

Typical phased XO → Agent Platform migration

8–16 wks

Greenfield Artemis implementation

24×7

Managed run under a single SLA


Timelines assume access to systems of record, a named business owner per intent and a decision cadence for guardrail and escalation policy.

How the first agent reaches production.

A sequenced program across the three pillars, with parity gates, fallback and the evidence pack a steering committee needs to fund the next wave.


DELIVERY SEQUENCE

Discover, blueprint, three waves, then run.



Pilot intents run behind fallback, then A/B by team, then production — with guardrails agreed before build and no SLA movement without evidence.

FIRST 90 DAYS



Outcomes our clients audited.

Selected Kore.ai engagements — the baseline, what we shipped and the measured result.

CASE STUDY · HEALTHCARE PROVIDERS

38% call containment in 90 days for a regional healthcare provider — healthcare-safe conversational AI

Regional healthcare provider, 42 clinics · Kore.ai Agent Platform · Epic · Amazon Connect

Call volume was up 34% year over year, driven by appointment scheduling, medication refills and pre-visit intake. Hiring couldn't keep pace and abandon rates were climbing. Any AI rollout had to stay inside healthcare guardrails, integrate with Epic and never mis-triage a clinical symptom.

- Bounded intents, clinical guardrails
- Voice + chat from day one
- Epic-native tool layer
- PHI guardrails and HITL

38%

Tier-1 call containment

-24%

Average handle time on transferred calls

+0.3pt

Patient CSAT

0

healthcare audit findings

"Our access center finally has room to breathe — and patients are getting appointments booked at 11pm on a Sunday."

— VP Patient Access

CASE STUDY · FINANCIAL SERVICES

Retail banking self-service at a regional bank — conversational AI across servicing journeys

Regional bank, 2.3M retail customers · Kore.ai Agent Platform · Amazon Connect · Fiserv DNA

Retail servicing was 62% of contact volume with four repetitive intents. Prior IVR self-service capped at 14% containment. FFIEC and enterprise-security posture meant every new self-service intent had to be reviewed for authentication rigor and dispute-window handling.

- Strong-auth first
- Dispute intake as guided form
- Core banking tool layer
- Model risk management pack

36%

Containment on top-4 servicing intents

-52%

Dispute intake AHT

0

enterprise-security / FFIEC findings

+8

Digital NPS

Containment, AHT and quality — measured.

Agent Platform and Contact Center AI programs on regulated retail and healthcare floors.

CASE STUDY · RETAIL & ECOMMERCE

41% self-service containment for a national retailer — order status and returns on conversational AI

National retailer, 3,100 agents · Kore.ai Agent Platform · Salesforce Commerce Cloud · Genesys Cloud

70% of contacts were three intents — order status, returns and loyalty balance — but the previous chatbot only handled 12% end-to-end. Leadership wanted containment without a wave of angry loyalty-member escalations at peak.

- Intent-first, tier-based rollout
- Search AI grounding on the product catalog
- Voice + chat single-source authoring
- Peak-week guardrails

41%

Containment on top-3 intents

+0.5pt

CSAT

-19%

AHT on escalated calls

12 wk

First intent live to peak-ready

CASE STUDY · HEALTHCARE PROVIDERS

Agent assist across 900 seats at a healthcare provider — 22% AHT reduction with healthcare guardrails

Integrated healthcare provider, 900 patient-access seats · Kore.ai Contact Center AI · Genesys Cloud · Epic

Agents were flipping between five apps per call and after-call notes into Epic were inconsistent — CMIO flagged clinical-documentation risk. Prior agent-assist pilots had stalled because summaries were too generic to paste into an Epic encounter note.

- Overlay, not replacement
- Epic-shaped auto-summary
- Real-time next-best-action on approved KB
- Supervisor coaching loop

-21%

AHT

+18%

First-contact resolution

-63%

After-call work

900

Seats live in 10 weeks

ACROSS THESE FOUR PROGRAMS

36–41%

Containment on targeted intents

-19 to -63%

AHT and after-call work reduction

10–12 wks

From first intent to production scale

0

Regulatory or audit findings raised

Full case studies, reference calls and the evaluation evidence behind these figures are available under NDA at pronix.ai/resources/case-studies.

Engagement models and how to reach us.

Start with one intent. Prove it against baseline. Then scale across the estate.

HOW WE ENGAGE

Four commercial models, one accountable partner.

Fixed-scope implementation

A defined outcome, a fixed price and a dated plan — used for AI for X launches, migration waves and greenfield Artemis builds.

Managed run

A named run pod owning tuning, evals, releases and 24×7 operations against agreed availability, containment and quality SLAs.

Staff augmentation

Kore.ai-certified developers, architects and admins embedded into your program under your delivery leadership.

T&M Agile Teams

A blended pod for evolving roadmaps — sprint cadence, shared backlog and monthly value reviews.

One accountable delivery model: US-based architecture and program leadership with global engineering pods running 24×7 build, cutover and hypercare.

CONTACT

Talk to the team that ships these programs.

HEADQUARTERS — US

666 Plainsboro Rd, Suite# 1361, Plainsboro, NJ 08536, USA

GLOBAL DELIVERY CENTER

Hyderabad, India

SATELLITE OFFICES

Dubai, UAE · London, UK

pronix.ai/book · info@pronix.ai · pronix.ai/contact

PRONIX AT A GLANCE

200+

Certified engineers & architects

15+

Average years of enterprise delivery experience

9

Named CCaaS/CX, Conversational AI & Cloud partnerships

100+

Enterprise Clients Worldwide

3

Global delivery regions

24×7

Managed operations across three regions

NEXT STEP

Model the business case before you commit.

Run the contact center AI ROI model, or book a working session with the Kore.ai-certified practitioners who deliver AI for X, migration and Artemis programs.

pronix.ai/book
pronix.ai/roi/contact-center-ai
pronix.ai/platforms/kore-ai

Pronix Inc (pronix.ai) is a specialized AI & CX Systems Integrator. Leading with AI. Innovating with Digital. Kore.ai and third-party names and marks belong to their owners. Outcomes are drawn from delivered engagements and vary by scope, platform and operating model. Security questionnaires, controls documentation and named client references are available under NDA.