



Specialized AI & CX Systems Integrator.

From strategy and pilots to production outcomes.

Leading with AI. Innovating with Digital.


Executive capabilities overview · pronix.ai · 2026

WHO WE ARE

An enterprise AI and CX implementation partner — since 2010.

Pronix Inc (operating as pronix.ai) is a systems integrator specialized in AI and CX — improving customer experience, automating operations and accelerating AI adoption through production.

Fixed-scope implementation · managed run · certified talent pods — one accountable delivery model.



2010

Founded · Plainsboro, NJ

3

Delivery regions: US · India · EMEA

24×7

Managed run and hypercare

15+

Platforms certified across CCaaS, cloud and LLM

AI pilots are easy. Production is where programs stall.



Fragmented CX estate

Legacy IVR, multiple CCaaS tenants and CRM debt make every AI use case a bespoke integration project.



No integration path

Models demo well in isolation and then fail against core systems, entitlements and real customer data.



Governance added last

Evals, redaction and audit trails are retrofitted after the fact — so risk and legal slow the rollout.



Outcomes unmeasured

Containment, AHT and cost per resolution are never instrumented, so the business case never closes.

Pronix.ai exists to close the distance between the pilot and the P&L.

Three primary practices.

Enterprise AI & Agentic AI

- Agentic workflows with bounded scope and HITL
- RAG, evaluation harnesses and model routing
- LLM cost, FinOps and production observability

CX & Contact Center AI

- CCaaS modernization and IVR-to-voice-AI cutovers
- Agent assist, auto-QA and post-call analytics
- Journey orchestration across voice and digital

AI Business Automation

- Document intelligence and straight-through processing
- Back-office workflow agents on existing systems
- Process mining and outcome instrumentation

Three ways to start — each with a defined exit.

Diagnostic

2–6 weeks

- Readiness and estate assessment
- Use-case portfolio scored on value, feasibility, risk
- Target-state architecture
- Business case with payback timeline

Fixed fee

Pilot

8–12 weeks

- One bounded use case into production
- Integrated to your core systems
- Evaluation harness and guardrails
- Measured against agreed success metrics

Fixed-scope SOW

Scale

Ongoing

- Multi-use-case delivery roadmap
- Managed run with 24x7 operations
- Continuous tuning and cost optimization
- Quarterly outcome and roadmap reviews

Managed service or T&M pods

Commercials are scoped to the use case and estate — no packaged list pricing. You are never locked into the next phase.

Every model has a defined scope, a defined exit and an agreed success metric.

THE FIRST 90 DAYS

What a pilot-to-production engagement looks like.

Weeks 0-2



Discovery

Estate mapping, intent and volume analysis, use-case scoring, success metrics agreed.

Scored use-case portfolio

Weeks 2-6



Design & build

Target-state architecture, integrations to core systems, guardrails and retrieval design.

Working build in your tenant

Weeks 6-10



Evaluate & UAT

Golden datasets, eval harness, red-teaming, business and compliance sign-off.

Evaluation report + sign-off

Weeks 10-12



Cutover & hypercare

Phased traffic ramp, 24x7 hypercare, outcome instrumentation live in dashboards.

Production + measured baseline

Illustrative plan for a single bounded use case; adjusted to estate complexity at scoping.

OUTCOMES WE DELIVER

The benchmark range enterprises should hold us to.

45–60%

Voice containment on transactional intents

15–25%

Handle-time reduction from agent assist

60–85%

Cheaper per contained interaction vs human-handled

20–40%

Blended cost-per-contact reduction, year over year

100%

Of interactions scored by automated QA

60–75%

Of LLM traffic routed to the cheap tier, no quality loss

PROOF IN PRODUCTION

Anonymized outcomes from delivered programs.

-62%

Top-5 US health insurer

Agentic claim triage · Amazon Bedrock, Kore.ai

Claim-triage cycle time across a multi-million-claim annual book — zero audit findings

90 days

National retailer, 2,400 agents

Legacy IVR cutover · Amazon Connect

Legacy IVR to live across 2,400 agent seats — AHT down 28%, CSAT up through peak

-22%

Global BPO, 1,200 seats

Agent assist · NICE CXone, Kore.ai, Azure AI

AHT across 1,200 seats and three client programs, with a 14% QA score lift

20–40%

Fortune 100 insurer

LLM FinOps · multi-model routing

LLM spend reduction on production traffic, with no measured quality regression

Where we fit against the alternatives on your shortlist.

Global SI

THE GAP

Deep bench, but multi-quarter mobilization, generalist CX teams and pyramid pricing on a use case that needs eight weeks.

WITH PRONIX

Senior-only pods, fixed-scope SOW, first production use case in one quarter.

AI boutique

THE GAP

Strong models and prompts, no integration depth into CCaaS, core systems or the run model after go-live.

WITH PRONIX

Integration, governance and 24x7 managed run under one accountable contract.

CCaaS partner / reseller

THE GAP

Optimizes the platform they resell; AI stops at the vendor's roadmap and licence line.

WITH PRONIX

Platform-neutral across 15+ CCaaS, cloud and LLM platforms — architecture chosen on fit.

Specialist depth of a boutique, delivery discipline of an SI — with the run model included.



Start with a 30-minute strategy session.

We will walk your estate, name the two or three use cases worth funding first, and show you the delivery path from pilot to production.

WEB

pronix.ai/book

EMAIL

info@pronix.ai

HQ

666 Plainsboro Rd, Suite 1361
Plainsboro, NJ 08536

This deck: pronix.ai/capabilities.pdf



APPENDIX

Delivery detail, assurance and industry priorities.

Delivery model · Enabling capabilities · Platform ecosystem · Governance · Security & compliance · Six industry views

DELIVERY MODEL

Advise → Implement → Run. Staff where capacity is the gap.

01	Advise Strategy & Consulting	Readiness assessment, scored use-case portfolio, target-state architecture, CFO-ready business case.	2–6 weeks Fixed-fee assessment or advisory retainer
02	Implement Implementation Services	Integrated build on your CCaaS, cloud and AI platforms, with evals, cutover and hypercare.	8–16 weeks Fixed-scope statement of work
03	Run Managed Services	24×7 operations, tuning, model and cost optimization, quarterly outcome reviews.	Ongoing Monthly managed service against SLAs

Staff — certified talent pods embedded in your delivery organization, billing within two weeks.

What makes the practices land in production.

-  **Enterprise Integration**
APIs, event buses and core-system connectivity into CRM, claims, policy and banking platforms.
-  **Digital Engineering**
Cloud, data and application engineering that carries the AI workload rather than fighting it.
-  **Data & AI Foundations**
Retrieval, knowledge curation, entitlements and the evaluation infrastructure AI depends on.
-  **AI Governance**
Model registry, evals, red-teaming, PII/PHI guardrails and an auditable decision trail.
-  **Managed Services**
24x7 run, tuning and continuous optimization after go-live, against agreed SLAs.
-  **Certified Talent Pods**
CCaaS, cloud and AI engineers on the bench to extend your team where capacity is the constraint.

PLATFORM ECOSYSTEM

Certified on the platforms your estate already runs on.

Contact center & CCaaS	Conversational & agentic AI	Cloud, LLM & data
 Amazon Connect	 Kore.ai XO	 AWS Bedrock
 Genesys Cloud CX	 Copilot Studio	 Microsoft Azure AI
 NICE CXone	 Google CCAI / Dialogflow	 Google Vertex AI
 Five9	 Salesforce Agentforce	 OpenAI Enterprise
 Talkdesk	 Retell AI	 Anthropic Claude

AWS Advanced Tier Services Partner · Genesys, NICE CXone and Five9 implementation partner · Kore.ai strategic partner · Microsoft Gold · Salesforce consulting partner

Regulators and risk teams are stakeholders from day one.

■ Model governance

Registry, versioning and promotion gates — every model change passes an evaluation harness before production.

■ Evaluation & red-teaming

Golden datasets, LLM-as-judge plus human calibration, adversarial testing on every release.

■ Data protection

PII/PHI redaction, entitlement-aware retrieval and tenant isolation designed into the architecture.

■ Auditability

Every AI action logged, replayable and mapped to your control library for internal audit and regulators.



Security questionnaires, controls documentation and named client references are available under NDA.

What procurement and risk will ask for.

Information security	[SOC 2 / ISO 27001 STATUS — confirm]	Controls documentation and completed security questionnaires shared on request.
Healthcare data	HIPAA-aligned handling	PHI redaction by default, BAA available, entitlement-aware retrieval on clinical content.
Privacy & data residency	GDPR / DPA ready	DPA and sub-processor list provided; US, EU and India delivery regions with residency options.
AI-specific risk	NIST AI RMF-aligned practice	Model registry, eval gates, red-teaming and a replayable audit trail per model version.
Access & tenancy	Least-privilege by design	Tenant isolation, client-owned cloud accounts where required, no training on client data.

Financial services

Priority use cases

- Servicing deflection on balance, payment and status intents
- Fraud and AML copilots inside the agent workflow
- Advisor assist grounded in product and compliance content

Integration note

Core banking, card and servicing platforms integrated via APIs, with full redaction and audit trails.

45–60%

Containment on transactional intents

15–25%

AHT reduction with agent assist

Ranges are published pronix.ai benchmarks; figures marked (delivered) are anonymized client outcomes.

Healthcare providers

Priority use cases

- AI scheduling and access with real EHR write-back
- Ambient clinical documentation support
- Nurse and agent copilots on approved clinical content

Integration note

Epic and Cerner-adjacent integration patterns, HIPAA-aligned handling and PHI redaction by default.

30–45%

Call volume deflected from access centers

25–40%

Faster patient intake and triage cycles

Ranges are published pronix.ai benchmarks; figures marked (delivered) are anonymized client outcomes.

Health payers

Priority use cases

- Member service IVA for claims, benefits and EOB
- Prior-authorization automation with clinical human-in-the-loop
- Appeals, grievances and attachment intelligence

Integration note

CMS-0057-F aligned workflows, decision timelines and defensible audit trails per model version.

-62%

Claim-triage cycle time (delivered)

100%

Interactions scored by automated QA

Ranges are published pronix.ai benchmarks; figures marked (delivered) are anonymized client outcomes.

Insurance

Priority use cases

- FNOL conversational triage and claims intake
- Underwriting submission extraction and appetite checks
- Policy and endorsement copilots with state-aware retrieval

Integration note

Guidewire, Duck Creek and legacy policy-admin integration via APIs and message-bus patterns.

20–35%

Straight-through processing on eligible claims

25–40%

Faster quote turnaround for SMB submissions

Ranges are published pronix.ai benchmarks; figures marked (delivered) are anonymized client outcomes.

Retail & e-commerce

Priority use cases

- Order status, returns and WISMO deflection
- Peak-season readiness and CCaaS modernization
- Conversational commerce and post-purchase care

Integration note

Commerce, OMS and Service Cloud integration, with warm site-by-site cutovers ahead of peak.

90 days

Legacy IVR to live on Amazon Connect (delivered)

-28%

AHT through holiday peak (delivered)

Ranges are published pronix.ai benchmarks; figures marked (delivered) are anonymized client outcomes.

BPO & outsourcing

Priority use cases

- Agent assist across multi-client programs
- Auto-QA replacing 5% sampling with full coverage
- Multilingual voice AI and outcome-based delivery

Integration note

Shared platform with per-client isolation — SLA-safe shadow-mode rollout before production suggestions.

-22%

AHT across 1,200 seats (delivered)

+14%

QA score lift across three programs

Ranges are published pronix.ai benchmarks; figures marked (delivered) are anonymized client outcomes.