WHITE PAPER

INFOSYS LABORFORCE AI: BUILDING THE AGENTIC STATE OF WORK

A Visionary White Paper for State Labor Leaders

Abstract

State labor agencies are under mounting pressure from rising caseloads, sophisticated fraud, and citizens demanding seamless digital experiences. Entering the "Agentic Era," governments are adopting AI systems that act autonomously under human oversight to deliver outcome-driven services. Infosys LaborForce, built on Salesforce, exemplifies this shift by leveraging intelligent agents to orchestrate work at machine speed while preserving human judgment for fairness and trust.

Unlike traditional automation, Agentic AI augments rather than replaces human decision-making, integrating Salesforce, Agentforce, Data Cloud, and a Trust Layer for transparency, security, and auditability. By aligning operational efficiency with ethical AI governance, LaborForce provides a resilient blueprint for proactive, equitable digital government.

Infosys®
Public Services

Executive Summary

State labor agencies face converging pressures: surging caseloads, increasingly sophisticated fraud, and citizens who expect digital services as intuitive as the private sector. Governments worldwide are entering what thought leaders call the 'Agentic Era' where AI systems perceive, reason, and act autonomously under human guidance. Infosys LaborForce, built natively on Salesforce, shows what an Agentic State can look like in practice. It shifts Labor systems from process-driven routines to outcome-driven work, where intelligent agents handle tasks at machine speed and human staff focus on judgment, equity, and trust.

Unlike previous waves of automation, Agentic AI augments rather than replaces human judgment. LaborForce's architecture integrates Salesforce Einstein, Data Cloud, and the Trust Layer to ensure transparency, auditability, and security at every step. By aligning operational excellence with ethical AI governance, LaborForce offers states a blueprint for digital government that is proactive, fair, and resilient.

The following framework defines how LaborForce operationalizes the Agentic State model across interconnected layers.

- **The Implementation Layers** represent how AI transforms the daily work of government, from intake and adjudication to fraud prevention and constituent engagement.
- **The Enablement Layers** provide the foundation for responsible adoption, ensuring data integrity, privacy, cybersecurity, and continuous learning across people and systems.

Vision: The Agentic LaborForce

The future of state labor administration will be defined by governments that embrace outcome-driven AI as a partner in governance. The Agentic LaborForce embodies this future, a connected ecosystem where AI agents support eligibility, adjudication, integrity, and citizen service with empathy and speed. By uniting technology, transparency, and human oversight, LaborForce becomes not only a platform but a model for accountable innovation in public service.

Every decision made in LaborForce, from each claim reviewed to every fraud prevented and every citizen interaction, contributes to a more adaptive, trusted, and human government.

Agentic LaborForce: Definition

The Agentic LaborForce represents a transformative workforce model where human talent and intelligent agents collaborate seamlessly to achieve organizational objectives. This concept emphasizes autonomy, adaptability, and decision-making capability enabling both people and AI-driven systems to act proactively within defined governance frameworks.

Key characteristics include:

- **Human-AI Synergy:** Combining human judgment with machine efficiency to optimize outcomes.
- **Autonomy with Accountability:** Empowering human or digital agents to make informed decisions while maintaining compliance and oversight.
- **Dynamic Scalability:** Leveraging automation and augmentation to respond rapidly to evolving business needs.
- **Outcome-Oriented Execution:** Aligning every action with measurable goals, ensuring agility without sacrificing control.

The Agentic LaborForce is not just a technological shift, it is a strategic evolution that redefines productivity, resilience, and innovation for the modern enterprise.

The Agentic LaborForce: AI-Layer Framework

Implementation Pillars

Claims & Adjudication	Program Integrity	Constituent Support	Security & Compliance
• Conversational intake • Document processing • Rules-guided support	• Real-time fraud detection • Cross-program analysis	• 24/7 virtual agents • Seamless human handoff	• Trust-layer guardrails • Audit logging • Access controls

• Future proofing • De-coupled AI continuous platform

Enablement Pillars

AI Governance	Data & Privacy	Tech Stack	Cyber-security
• Accountability redress • Transparency	• Auditability • Role-based access	• Salesforce integration • APIs	• Zero data retention • Resilience

Together, these layers illustrate how LaborForce scales trustworthy, outcome-driven automation while maintaining human oversight and accountability which is the essence of an Agentic government.

A key design choice is that AI ***augments rather than replaces*** human judgment. Conversational interviews, document understanding, and decision support accelerate routine steps and surface risks early, while human reviewers retain control over complex or adverse decisions. Because LaborForce is native to Salesforce, AI capabilities arrive with platform guardrails (permissions, audit logs, trust layer policies) so programs can scale intelligence without compromising compliance or due process.

Value Proposition to States

1

Reduced Risk

Proven Salesforce technology backed by experienced state labor implementers.

2

Accelerated Delivery

Typical workforce implementations take 18–24 months; our integrated approach targets 12–15 months.

3

Future-Ready

Cloud-native, AI-enabled platform expandable to UI, Paid Leave, Disability, and Licensing & Regulation.

4

Customer Success

Dedicated post-implementation support framework to ensure adoption and drive continuous value.

Outcomes at a glance

The table below summarizes how LaborForce applies AI to each executive priority. Use it as a map: the left column anchors the business objective; the middle columns explain the capability and the reason it matters operationally; and the right columns translate improvement into outcomes and the Salesforce components that make the change sustainable.

Priority	What AI Does	Why It Matters	Example Outcome	LaborForce Capability
Faster claims & adjudication	• Conversational intake • Document AI • Case summaries • Rules-guided decision support • Selective auto-adjudication	Reduces backlogs while preserving human judgment for complex issues	• Shorter time to initial determination • Steadier quality and fewer reversals	Agentforce, Einstein, Flows, OmniStudio, Business Rules. AI and Business Rule Driven Auto Adjudication and Fraud identification.
Program integrity & employer tax	• Real-time anomaly scoring • ID/bank orchestration • Cross-program link analysis	• Prevents improper payments before disbursement • Improves investigator yield	• More dollars prevented/recovered • Lower false-positive load	Data Cloud, Einstein ML; connectors (IDH, DNP, IDV). AI and Business Rule driven integrity cross-reference.
Constituent experience	• 24/7 virtual agents that guide forms, answer FAQs, send reminders • Seamless human handoff with context	• Deflects routine contacts • Increases completion rate • Improves satisfaction	• Higher self-service containment • Better FCR/CSAT	Agentforce Digital Engagement, Knowledge, Messaging. Award winning User Experience with AI driven help and system navigation.
Security & compliance	• Trust Layer guardrails - dynamic grounding, masking, zero data retention • Role-based access • Audit trails	Adopt AI within strict privacy and audit boundaries	• Audit-ready AI artifacts and explanations • Faster approvals	Einstein Trust Layer; Gov Cloud; Shield. Designed with Infosys' Responsible AI guidelines.
Future-proof investment	• Decoupled AI layer • Platform-delivered upgrades without rewriting business logic	• Adds new AI safely over time • Lowers cost-of-change	Quarterly feature adoption with minimal rework	Agent Builder; Platform; APIs; Flows. LaborForce AI Layer protects business logic from change in tools and technology.

In practice, these outcomes are achieved by reshaping the order of work. Interviews become adaptive, documents are interpreted as they arrive, and risk scoring runs in the background so staff see complete, organized case context sooner. The result is less time lost to triage and rework, more time spent on the complex, and a clearer audit story of how decisions were made.

1) Faster Claims & Adjudication

Claims operations are intrinsically sequential: collect facts, validate evidence, apply policy, draft a determination. During surges, that sequence creates bottlenecks and quality variation. AI changes the order of work. A conversational agent begins with open-ended narratives from the claimant and employer, then asks targeted follow-ups tied to rule elements, not a one-size-fits-all form.

As documents arrive, document AI extracts key dates, wages, and conditions and flags gaps automatically. The system drafts a case synopsis, timeline, issues, and citations so adjudicators start with context rather than raw inputs. Rules-guided decision support highlights edge conditions and requests missing items before review even begins.

Customer Problem	Surges create backlogs; quality varies across adjudicators; too much manual triage
AI Approach	Conversational intake; document AI; evidence summaries; rules-guided decision support; selective auto-adjudication on clear, high-confidence issues
Business Benefit	Shorter cycle time and steadier outcomes; staff attention shifts to complex cases and coaching
KPIs	Initial determination timeliness; first-payment timeliness; appeals reversal %; examiner throughput
Relevant Capabilities	Agentforce, Einstein, Flows, OmniStudio, Business Rules, Document AI

Implementation pattern: define explicit confidence thresholds that separate auto-adjudicable issues from those requiring human review; store the AI-generated rationale with references in the case record; and create a closed-loop learning path from appeals and supervisor feedback. These practices preserve due process while continuously improving guidance.

Change management: publish a short "adjudicator co-pilot" playbook, run side-by-side evaluations for two pay cycles, and calibrate thresholds before scaling statewide. Focus first on high-volume, low-discretion scenarios to earn trust and capacity.

2) Program Integrity & Employer Tax

Rule-only integrity approaches are brittle against adaptive actors and coordinated rings. AI introduces continuous risk scoring that weighs behavior, device fingerprints, geospatial patterns, and relationship signals alongside identity and bank verification. Rather than "pay and chase," the system pauses or routes risks before funds move.

Data Cloud unifies data under governance so link analysis can reveal shared devices, payment instruments, email patterns, or employer relationships across claims and programs. Investigators receive cases pre-assembled with the most relevant evidence and a suggested next-step checklist, boosting yield and shortening cycle time.

Customer Problem	Static rules overwhelm investigators with alerts and false positives; organized fraud slips through
AI Approach	Real-time anomaly scoring; identity and bank verification orchestration; cross-program link analysis and clustering; suppression of low-value alerts
Business Benefit	Prevention-first posture; better hit rate; faster investigations; lower burden on legitimate claimants
KPIs	Improper payment rate; dollars prevented/recovered; false-positive rate; investigation cycle time; precision/recall of alerts
Relevant Capabilities	Data Cloud, Einstein ML; connectors for IDH, Do Not Pay, IDV; case management and audit logs

Governance essentials: log every action (hold, release, escalate) with model version and salient features; require human confirmation for adverse actions; and monitor models for drift and fairness with periodic back-testing. This balances assertive prevention with due process and equity.

Operating model: designate an integrity product owner, stand up a weekly analytics huddle with investigators, and tune thresholds based on yield and appeal outcomes. Small, steady threshold moves often outperform sweeping model changes.

3) Constituent Experience Powered by Agents

Constituents value clarity and immediacy. AI agents provide 24/7 assistance across web, chat, and SMS: explaining eligibility in plain language, guiding users through only the questions that apply to them, checking status, and reminding them about documents or hearings. The goal is fewer abandoned applications and fewer avoidable calls.

When a conversation exceeds the agent's scope, it hands off to a human with the full history and case context, so staff don't need to ask people to repeat themselves. Knowledge is grounded in approved content and constrained by permissions, so answers match the user's role and the state's policy.

Customer Problem	Long hold times, incomplete applications, and repeat contacts during peaks
AI Approach	Multilingual virtual agents grounded in approved knowledge; workflow completion and proactive reminders; seamless human handoff with full context
Business Benefit	Higher containment of routine contacts; better first-contact resolution; improved completion quality and satisfaction
KPIs	Self-service containment; FCR; average handle time; completion rate; satisfaction score
Relevant Capabilities	Agentforce Digital Engagement, Knowledge, Messaging, Omnichannel Routing

Design for adoption: start with the top five question clusters (basic eligibility, status, documentation, payments, appointments). Instrument containment and satisfaction, and establish a lightweight editorial workflow so policy updates propagate quickly to agent responses.

Accessibility: ensure agents handle plain language, Spanish/English at a minimum, and fall back gracefully to human support. Proactive reminders should include SMS and email with accessible formatting.

4) Security, Privacy & Compliance

Public-sector AI must meet strict privacy, security, and records requirements. Salesforce's Trust Layer provides guardrails such as dynamic grounding (only the minimum necessary context goes to the model), data masking, toxicity filters, and zero data retention for third-party LLM calls. Combined with Gov Cloud controls, role-based access, and full audit trails, agencies can scale AI responsibly.

Equally important is explainability. For summaries, recommendations, or risk scores that influence staff decisions, LaborForce captures the input sources, key factors, and the human's accept/reject action, creating an audit-ready chain of reasoning that supports hearings and oversight.

Customer Problem	Adopt AI without compromising privacy, security, or records obligations
AI Approach	Trust Layer guardrails; permission-aware grounding; comprehensive logging and explanations; government-grade cloud controls
Business Benefit	Confidence to scale; simplified audits; faster Authority to Operate (ATO) updates for new AI features
KPIs	Security findings; audit exceptions; time-to-approve features; incident Mean Time to Resolution (MTTR)
Relevant Capabilities	Einstein Trust Layer; Gov Cloud; Shield; audit logs and retention controls

Operating guardrails: define an AI review board, catalog datasets approved for grounding, and set red-lines (e.g., no autonomous adverse decisions). Train users on when AI outputs are advisory versus binding, and require explicit acceptance before any recommendation becomes part of the record.

Records & retention: store generated artifacts with linkage to the user action and underlying sources; apply existing retention schedules so AI artifacts follow standard legal hold and discovery processes.

5) Protecting Your AI Investment

The AI landscape evolves daily; your business logic should not. Technology cycles move fast, what feels cutting-edge today can become outdated and unsupported in a matter of months. New AI models, integration patterns, and security requirements emerge constantly, often rendering bespoke solutions brittle or costly to maintain. Agencies that hardwire AI into core business logic risk expensive rewrites every time the market shifts. A decoupled, configuration-first approach ensures you can adopt the latest capabilities without destabilizing operations, preserving both agility and compliance.

LaborForce separates AI experiences from key areas such as core case, eligibility, and payment logic using configuration-first patterns and the One Code / One Action principle for shared functions. This minimizes rework as platform features advance and lets teams adopt new capabilities behind existing governance.

Practically, this means interview components, decision patterns, and alerting thresholds are modular assets with owners, not ad-hoc scripts. As platform releases ship, you can upgrade the intelligence while keeping policies, workflows, and security boundaries stable.

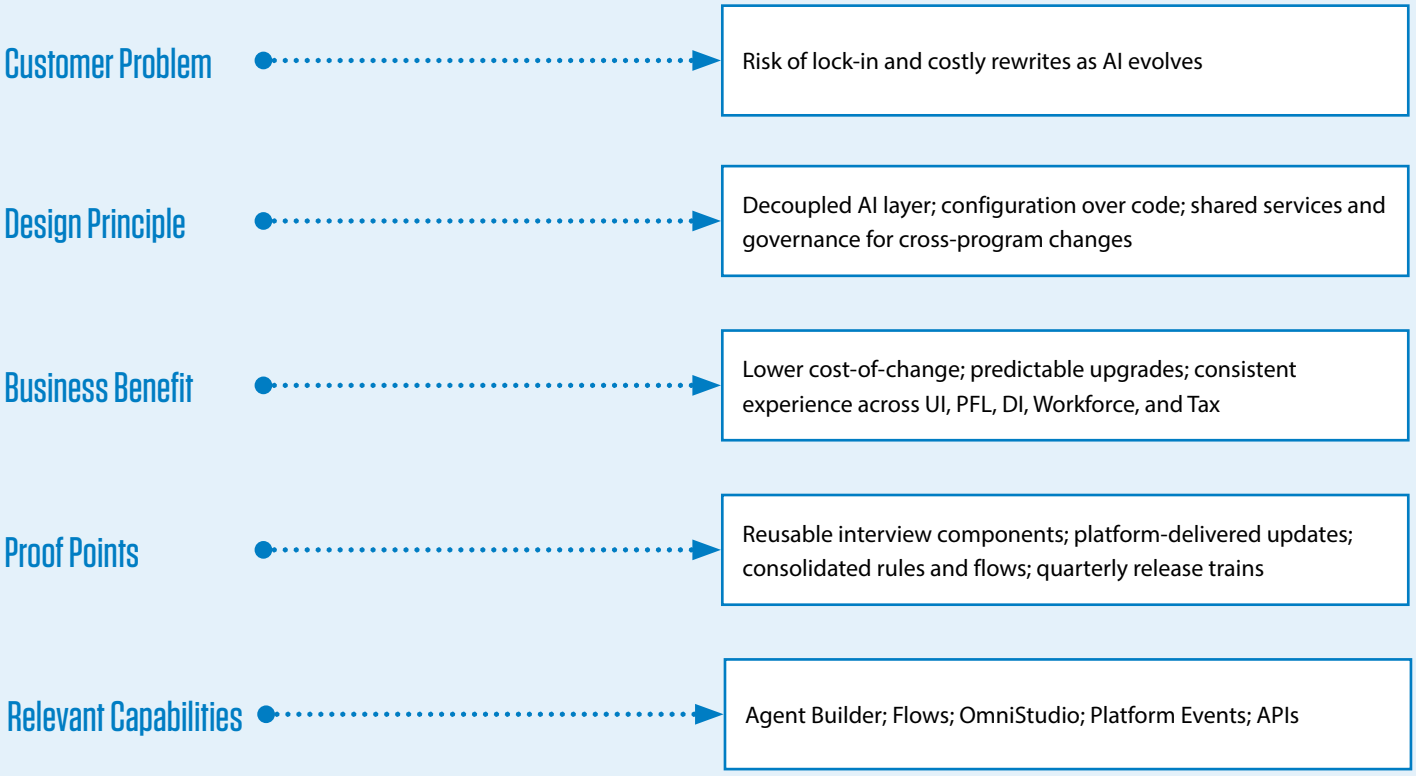

Governance pattern: maintain a component catalog (prompts, flows, summaries), assign product owners, and run a quarterly AI release train with pre-production trials. Validate downstream effects across programs before each release using scripted regression tests and sample audits.

Vendor strategy: prefer standards-based connectors and keep fine-tuning data portable. Avoid bespoke models when a governed platform capability meets the need to reserve custom modeling for genuinely unique problems with clear ROI.

6) 90–180–365 Day Adoption Roadmap

A phased roadmap reduces risk and builds confidence. Start with low-risk, high-value scenarios to establish guardrails and change patterns; then expand into adjudication support and integrity orchestration as skills mature. Each phase includes technical, operational, and governance tasks.

Phase	Focus	Example Deliverables	Success Measures
0–90 days	AI readiness & quick wins	Chat/FAQ agent; document AI for intake; integrity pilot on subset	Containment rate; intake cycle time; pilot precision/recall; staff adoption
90–180 days	Claims acceleration	Conversational intake + AI summaries; adjudicator co-pilot; link analysis dashboards	Determination timeliness; backlog trend; investigator productivity; reversal %
180–365 days	Scale & harden	Selective auto-adjudication for simple issues; proactive outreach; expanded fraud orchestration	Improper payment rate; satisfaction; cost per claim; audit readiness

Governance & change: in Phase 1, establish the AI review board, catalog permissible data and prompts, and publish an acceptable-use policy. In Phase 2, build human-in-the-loop patterns with clear acceptance steps and rationale storage. In Phase 3, validate automated actions with audit artifacts and rollback plans before broadening scope.

KPIs & Business Case

KPIs convert intent into measurable progress. We recommend setting directional targets, validating them with pilot data, and then hardening to operational goals. Claims metrics focus on timeliness and quality; integrity metrics balance prevention with fairness; CX metrics track containment and satisfaction; and Ops/IT metrics ensure the platform can absorb change predictably.

Area	Leading KPIs (pilot →scale)	Target Benchmarks (set collaboratively)
Claims	Timeliness; backlog; appeals reversal %; examiner throughput	15–30% faster initial determinations; 10–20% fewer reversals (directional)
Integrity	Prevented/recovered $; false positives; precision/recall	+$ prevented; reduced false positives vs. rules-only; improved precision/recall
CX	Containment; FCR; AHT; completion rate; satisfaction	Higher containment and satisfaction; lower AHT; better completion quality
Ops/IT	Release velocity; rework %; ATO time; incident MTTR	Quarterly AI drops with minimal rework; quicker approvals; stable operations

Business case levers typically include avoided overtime during surges, reduced improper payments, lower average cost-per-contact via containment, and faster delivery of policy changes via low-code updates. We recommend capturing a baseline now so gains are attributable and defensible.

About Infosys LaborForce

LaborForce unifies all major labor programs under a secure portal, reducing redundancy and improving the user experience. Infosys LaborForce is a next-generation enterprise platform purpose-built to support the full spectrum of labor programs including workforce services, unemployment insurance, paid leave, and more within a unified and scalable architecture.

Developed on the Salesforce platform, LaborForce delivers the flexibility of modular design with the rigor of enterprise-grade governance, enabling agencies to modernize at their own pace while maintaining operational continuity.

What sets LaborForce apart is its ability to provide best-in-class functionality across every domain, without sacrificing the integrated experience today's customers expect. Whether a constituent is applying for services, reporting employment, managing appeals, or navigating job search activities, they can do so through a secure, user-friendly portal that aggregates all case data and interactions in one place eliminating silos and confusion.

Designed to evolve with agency needs, LaborForce offers:

- Cross-program visibility and data sharing through a single pane of glass
- Advanced security and privacy protections, meeting state and federal compliance standards
- Role-based workflows that streamline staff operations while preserving policy flexibility
- API-first design that simplifies integrations with external systems and partner organizations

LaborForce is more than just a configurable solution, it is a strategic foundation for labor agencies looking to improve service delivery, increase transparency, and empower both customers and staff through modern, cloud-based technology.

- Unified Data Cloud integrates program data, verification sources, and fraud prevention systems.
- Adaptable AI Layer evolves independently, allowing rapid deployment of new capabilities.
- Government-grade security with FedRAMP High certification, encryption, role-based access, and audit logging.

These strengths allow agencies to scale during surges, implement new programs quickly, and deliver consistent service.

References

- Salesforce Newsroom: Agentforce & Data Cloud FedRAMP High authorization (June 4, 2025); https://www.salesforce.com/news/stories/fedramp-high-agentforce-salesforce-platform/

- Trusted AI: The Einstein Trust Layer; https://www.salesforce.com/artificial-intelligence/trusted-ai/

- NASWA Integrity Data Hub; https://www.naswa.org/integrity-center/integrity-data-hub

- U.S. Treasury Press Release: Enhanced fraud detection prevented/recovered $4B in FY2024; https://home.treasury.gov/news/press-releases/jy2650

- Route Fifty: New Hampshire's benefits program embraces AI (Mar 4, 2025); https://www.route-fifty.com/artificial-intelligence/2025/03/new-hampshires-benefits-program-embraces-ai-amid-modernization-push/403448/

- Digital.gov: Using Chatbots to Improve Customer Experience; https://digital.gov/2021/04/07/using-chatbots-to-improve-customer-experience/

- Kaspar Korjus: The Agentic State; https://agenticstate.org/

To learn how Infosys LaborForce can help your agency navigate it's next email Darryl Scott at **darryl.scott@infosys.com**

For more information, contact askus@infosyspublicservices.com

Infosys® Public Services

Infosyspublicservices.com

Stay Connected