THE COMPOSABLE AGENCY

Architecting the Future of Autonomous Public Service

A Strategic White Paper for US State Governments

Mike O'Brien| 2026

CONFIDENTIAL — For Executive Distribution

Executive Summary

The 2026 Point of Peril

By late 2026, three converging forces will overwhelm agencies still running monolithic architectures: the **HIPAA Security Rule overhaul** mandates encryption, multi-factor authentication, and continuous audit capabilities that legacy mainframes—designed for closed, terminal-to-mainframe networks—cannot support without fundamental re-architecture; the **California Delete Act** creates a de-facto national standard for automated data-deletion on citizen request that siloed, undocumented databases cannot honor at scale; and the **Talent Cliff**—the retirement wave of the last COBOL-fluent workforce—removes the only professionals who understand the code keeping 50-year-old systems alive. Agencies that have not begun decomposing monoliths into composable, governable components by Q1 2026 face simultaneous compliance failure, workforce collapse, and cascading service disruption.

The compact between the governed and the government is fraying. Citizens navigating the apparatus of the state—securing food assistance, renewing a license, accessing healthcare—encounter a fragmented landscape of disconnected agencies, redundant data entry, and opaque decision-making. This "friction tax" disproportionately burdens the most vulnerable, erodes public trust, and undermines the legitimacy of democratic institutions. Cosmetic upgrades—a new website, a centralized portal—are insufficient if the underlying machinery remains monolithic and siloed.

This white paper presents the **Composable Agentic Agency**—a unified architectural and operational blueprint that answers these converging threats. The model rests on three interlocking pillars:

Packaged Government Capabilities (PGCs) decompose monolithic systems into autonomous, interchangeable modules—Identity, Eligibility, Claims, Notifications—each owned by a dedicated product team and swappable without system-wide risk. A **Data Mesh** replaces fragile centralized data lakes with a marketplace of domain-owned, API-accessible data products, enabling a whole-person view on demand while encoding HIPAA and deletion compliance directly into the platform as **compliance-as-code**. And a **Multi-Agent System (MAS)** orchestrates these capabilities through specialized AI agents—operating inside a sandboxed, auditable architecture with strict Human-in-the-Loop (HITL) guardrails—to deliver proactive, life-event-driven services at digital speed.

The result is a government that shifts from *renting intelligence* from monolithic vendors to **owning sovereign digital assets**; from reactive bureaucracy to a sensing, adaptive public-service nervous system. The pages that follow provide the technical blueprint, governance framework, operational models drawn from Code for America and the Harvard Government Performance Lab, and a tactical **3-3-3 Roadmap** that delivers measurable value in 90-day pulses. Above all, the vision insists that AI agents **augment** public servants rather than replace them in critical decisions, and that proactive digital services are built on a foundation of consent, privacy, and human oversight.

1. The Modernization Imperative

The crisis facing state HHS agencies is structural, not cosmetic. Three compounding forces demand immediate architectural action.

Strategic Debt

Roughly **80% of the $100B+ federal IT budget** is consumed by operating and maintaining legacy systems—many built on COBOL mainframes over 50 years old. This is not technical debt (a coding shortfall) but **Strategic Debt**: an organizational condition in which the cost of maintaining the status quo absorbs all resources, preventing any forward motion. The financial drain creates a vicious cycle: high maintenance costs starve the budget for modernization, ensuring legacy systems become older, more expensive, and more brittle with each passing year.

The security implications are equally severe. Originally designed for closed, mainframe-to-terminal networks, these systems were never architected for the modern threat landscape. The GAO has confirmed that critical legacy systems harbor **known cybersecurity vulnerabilities that cannot be remediated without modernization**. HHS's own information-security program was rated **"Not Effective" for FY 2024**, placing PHI and PII at direct risk. With the HIPAA Security Rule overhaul mandating encryption, multi-factor authentication, and continuous audit capabilities by 2026, agencies running legacy infrastructure face an unavoidable compliance gap that no amount of patching can close.

The Talent Cliff

The last generation of COBOL programmers and mainframe specialists is reaching retirement. When they leave, no one remains who understands the code. Simultaneously, the outdated technology environment repels modern engineers who expect cloud-native tools and agile workflows. The result is a **workforce dependency on a vanishing talent pool** and an inability to recruit replacements—a staffing crisis that compounds every year of delay.

The "Part-Person Data" Problem

HHS programs operate through siloed systems, each with its own database and schema. The Medicaid system knows the mother has diabetes; SNAP knows the family is food insecure; Housing knows they face eviction. But **the government as a whole knows none of this collectively**. Caseworkers cannot assemble a 360-degree view of a family's needs, leading to duplicative intake, uncoordinated interventions, and punitive responses to problems rooted in systemic gaps. A caseworker addressing a child's truancy might not know the family recently lost their housing—leading to punitive measures rather than supportive ones.

The Trust Deficit and Administrative Burden

Families in crisis are forced to act as the manual integration layer for the state—carrying physical documents from office to office, re-entering the same data into multiple portals, and navigating a labyrinth of eligibility rules. Research shows that **40% of SNAP-eligible non-participants cite burdensome paperwork** as a primary barrier. This is not inconvenience; it is a systemic failure that prevents critical aid from reaching those it was designed to serve, undermining the policy intent of the programs themselves.

The gap between seamless private-sector experiences (real-time package tracking, one-click ordering) and opaque government processes creates a **Constituent Experience Gap**. While 51% of Americans prefer digital public services, **74% rate local government websites as user-unfriendly**. This erosion of experience drives a trust deficit that undermines not just technology adoption but institutional legitimacy. The Seamless Government Experience is therefore not merely a technical project but a **democratic imperative** to restore the social contract in the digital age.

2. The Composable Blueprint

To escape the Strategic Debt cycle, agencies need a new architectural paradigm—not another monolithic replacement project, but a fundamentally different approach to how government technology is built, funded, and governed.

2.1 Four Principles of Composability

Modularity: Decompose massive systems (e.g., an entire MMIS) into discrete, capability-scoped components. **Autonomy:** Each component can be developed, tested, and deployed independently, decoupling release cycles. **Orchestration:** Components are assembled dynamically into workflows—a "Disaster Relief Application" composed from existing Identity, Payment, and Notification modules. **Discovery:** A catalog makes existing capabilities findable and reusable, eliminating duplicative builds.

2.2 Packaged Government Capabilities (PGCs)

The fundamental unit is the **Packaged Government Capability (PGC)**—a business-level grouping that delivers a complete piece of agency functionality. PGCs are the "verbs" of government: *Verify Identity, Process Claim, Determine Eligibility*. Unlike microservices (a technical pattern), PGCs are **business products** owned by persistent teams.

PGC	Function	Consuming Programs
Identity & Access Mgmt	Verifies identity; controls access via MFA, RBAC	Medicaid, SNAP, Child Welfare, DMV
Eligibility & Enrollment	Rules-engine determination; manages enrollment/renewal	Medicaid, SNAP, TANF, WIC
Provider Management	Credentialing, directory, network adequacy	Medicaid, CHIP, Child Care
Case Management	Client records, service plans, referral tracking	Child Welfare, Behavioral Health
Claims Processing	Intake, adjudication, remittance	Medicaid, CHIP
Communications	Multi-channel outreach (SMS, email, mail)	All HHS Programs
Consent Management	Captures/manages data-sharing consent; revocation	All HHS Programs, HIEs

Table 1: Core PGCs for HHS. Each is independently deployable, replaceable, and governed.

This modularity enables the **Strangler Fig pattern**: a legacy system is replaced piece by piece—swap the Notification module for a modern alternative without touching Claims or Eligibility—until the monolith is safely decommissioned.

2.3 MACH Architecture & Digital Sovereignty

PGCs are realized through **MACH architecture** (Microservices, API-first, Cloud-native, Headless). The API-first principle is the critical enabler: standardized contracts that let PGCs communicate, mirroring Estonia's **X-Road** exchange, which connects 900+ institutions and saves the country over **1,345 years of working time annually**. The Headless principle decouples front-end presentation from back-end logic, enabling true **no-wrong-door delivery** across web, mobile, SMS, and voice.

Critically, this architecture shifts agencies from **renting intelligence** (locked inside a vendor's monolith) to **owning sovereign digital assets**. When the state controls the API contracts, it controls the ecosystem. Vendors compete on

component quality within an open marketplace; the state retains architectural authority and can replace any underperforming component without triggering a system-wide migration. The outdated dependency on a single "Prime Systems Integrator"—an approach that reliably recreates vendor lock-in under a different label—is replaced by disciplined, modular procurement under open standards. The state itself serves as orchestrator, maintaining coherence across the composable ecosystem through strong program management and technical oversight capabilities built in-house. Modern modular procurement increases control and flexibility rather than diminishing it: clear interface contracts and API standards allow multiple specialized vendors or internal teams to contribute without a heavy-handed gatekeeper. States are already piloting these frameworks, proving that a well-governed modular approach reduces risk by avoiding single points of failure and enabling incremental improvement.

2.4 From Projects to Products

Composability demands a parallel shift in operating model. The traditional approach—fund a project, build a system, disband the team, enter "maintenance" decay—is replaced by **persistent, cross-functional product teams** that own a capability indefinitely. Their metric is not on-time delivery but continuous value: reduced processing time, improved accuracy, higher user satisfaction. As Code for America emphasizes, if digital strategy is fully outsourced, the government loses the ability to iterate on its own citizens' needs.

3. The Data Engine

PGCs provide the skeleton; **data is the circulatory system**. Centralized "one big database" approaches have largely failed—they create governance bottlenecks, strip domain context, and become compliance nightmares. The Composable Agency requires a **Data Mesh**.

3.1 Four Principles of the Government Data Mesh

Domain Ownership. The Medicaid division owns Claims Data; the Department of Labor owns Wage Data. Domain experts—not a central IT team—manage quality, regulation, and context. **Data as a Product.** Each domain publishes discoverable, documented, API-accessible data products with Service Level Objectives for quality and availability—e.g., a "Medicaid Eligibility Data Product." **Self-Serve Infrastructure.** Central IT shifts from gatekeeper to platform provider, offering tooling (storage, pipelines, cataloging, security) that lets domain teams publish products without deep infrastructure expertise. **Federated Computational Governance.** Global rules—HIPAA compliance, data masking, role-based access—are codified into the platform as **compliance-as-code**, applied automatically and consistently across all data products. PII leaving the Health domain is encrypted or tokenized by default.

3.2 Solving Interoperability

The Georgetown Beeck Center identifies data interoperability as the primary barrier to integrated benefits delivery. The Data Mesh resolves this by replacing dangerous raw-data dumps with a **marketplace of verified data products**. When SNAP needs income verification, it queries the Labor Department's "Verified Income" product via a secure API—never touching the internal database.

Dimension	Traditional Data Lake	Government Data Mesh
Architecture	Monolithic, centralized	Decentralized, distributed
Ownership	Central IT (bottleneck)	Domain teams (Medicaid, SNAP, Labor)
Data Quality	Poor; context lost centrally	High; owned by subject-matter experts
Access	Slow; manual requests	Fast; self-serve via catalog & APIs
Governance	Bureaucratic, manual audits	Automated compliance-as-code
Interoperability	Rigid global schemas	Flexible product interfaces

Table 2: Traditional data sharing vs. the Data Mesh paradigm.

This architecture enables a caseworker—or an AI agent—to assemble a **whole-person view dynamically** from distributed, trusted data products without creating a massive centralized citizen database. The "Part-Person Data" problem is solved not by aggregating everything into one risky repository, but by federating access across sovereign domains. Interoperability is achieved *on demand*, preserving privacy while enabling unprecedented coordination.

Crucially, this architecture makes compliance with emerging data-rights legislation tractable. The 2026 **California Delete Act**—which creates a de-facto national standard requiring automated data deletion upon citizen request—becomes manageable because each domain manages its own products and can honor deletion requests at the source, without needing to locate and purge records scattered across a monolithic data lake. Similarly, the HIPAA overhaul's enhanced audit and encryption requirements are met through the **federated computational governance** layer, which applies compliance rules automatically across all data products regardless of which domain publishes them.

4. The Agentic Leap

Composability provides the body. The Data Mesh provides the circulatory system. **Agentic AI provides the brain**—not a single monolithic "AI" but a managed ecosystem of specialized agents that reason, plan, and execute workflows to achieve high-level goals.

4.1 From Automation to Autonomy

Automation is script-based: "If form X is submitted, send email Y." It is rigid and breaks when conditions change. **Autonomy** is goal-based: "The constituent lost their job. Identify all eligible benefits, retrieve necessary proofs from data products, pre-fill applications, and schedule a career counseling session." The agent determines the sequence of steps, handles exceptions, requests missing information, and adapts its path based on what it discovers.

This distinction maps to a critical architectural choice. **Deterministic systems** (rules engines, workflow engines) produce identical outputs from identical inputs—ideal for routine, well-defined processes like calculating a Medicaid income threshold or processing a standard renewal. **Probabilistic systems** (ML models, LLMs) reason in open-ended domains with flexibility at the cost of some unpredictability—better suited for parsing a citizen's free-text request, identifying cross-program eligibility from ambiguous circumstances, or detecting anomalous patterns that warrant investigation. The Composable Agentic Agency deploys each where appropriate: deterministic precision for stable, high-volume tasks and probabilistic intelligence for nuanced, cross-cutting challenges that exceed the reach of static rules. This dual-mode approach—using the right tool at the right layer—is what distinguishes a production-grade agentic architecture from a proof of concept.

4.2 The Multi-Agent System (MAS)

The core architecture is not a single, monolithic "AI" but a **Multi-Agent System (MAS)** in which specialized agents collaborate to deliver services—mirroring the specialization of human teams but operating at digital speed:

Orchestrator Agent —the citizen's digital concierge. It serves as the primary interface, parsing natural-language requests ("I need help feeding my family") and decomposing them into a set of sub-tasks distributed across the agent ecosystem. The Orchestrator maintains state across the entire interaction, ensuring coherence as multiple agents execute their responsibilities in parallel. **Data Agent** —interacts exclusively with the Data Mesh, knowing how to discover, authenticate against, and retrieve information from various data products (Wage Data, Enrollment Data, Foster Care Data) via secure APIs. It never accesses raw databases. **Policy Agent** —specialized in interpreting regulatory logic. It leverages the Eligibility PGC's rules engine to apply complex policy calculations (SNAP income limits, Medicaid thresholds, TANF time limits) to the specific case, translating legislative intent into deterministic outcomes. **Communication Agent** —drafts and dispatches notifications via the citizen's preferred channel, ensuring messages are clear, empathetic, and compliant with accessibility standards. **Validation Agent** —a dedicated "critic" whose sole purpose is to review other agents' outputs against policy rules, historical precedents, and anomaly thresholds before any action reaches the citizen or triggers a financial transaction.

These agents coordinate under cognitive architectures like **ReAct (Reason + Act)**. In this framework, an agent engages in an iterative loop: it generates a Thought ("I should check the user's income data"), takes an Action (queries the Data Agent for income information), and observes the result. Based on the observation, it generates the next Thought and continues until the goal is achieved or an exception triggers escalation. This reasoning trace is logged in its entirety, creating a **transparent, auditable decision trail**—a critical feature for public-sector accountability and due-process requirements. If a citizen appeals a determination, the agency can produce the exact sequence of logic the agent followed, in plain language.

4.3 The Sandboxed Architecture

Architecture Principle: Containment by Design

Every AI agent operates inside an isolated **sandbox** with tightly scoped permissions. Agents interact with databases and transaction systems *only* via secure API gateways that enforce policy rules. The sandbox orchestrator logs every transaction and can throttle or halt agents that deviate from expected behavior. Even if an agent leverages an LLM for natural-language understanding, it cannot exceed its authorized bounds. Agents are treated as **plug-in modules under constant supervision**—autonomy never equals uncontrolled access.

The sandboxed architecture is the critical differentiator between responsible agentic deployment and the "wild AI" scenarios that rightfully concern regulators and the public. The sandbox layer serves as a safety buffer between the probabilistic intelligence of AI models and the deterministic requirements of government transactions. All agent actions—queries, updates, communications, financial disbursements—are mediated by the sandbox orchestrator, which enforces role-based permissions, rate limits, and policy constraints in real time. If an agent attempts an action outside its contract, the orchestrator blocks the action, logs the deviation, and alerts the governance team.

Each agent instance operates with an **ephemeral digital identity** that expires once its task completes. When an agent initiates a workflow, it is issued temporary credentials scoped only to the required systems and data for that specific workflow. Upon completion, credentials are immediately revoked. This design ensures agents cannot accumulate unchecked access over time. If an agent malfunctions, its blast radius is automatically contained to the scope of a single session. Ephemeral identities also create an unambiguous audit trail, since every action maps to a short-lived ID tied to a specific agent session and purpose.

4.4 The Sentient Agency: Proactive Governance

The ultimate evolution of this model moves government from reactive service delivery to **proactive service delivery**. By continuously and ethically monitoring sanctioned data products—aggregate layoff notices, public-health indicators, economic signals—agents detect events and initiate rapid-response workflows *before* citizens must navigate bureaucratic channels themselves.

Consider the contrast. In the **reactive state**, government waits for laid-off workers to discover the unemployment website, navigate confusing rules, and file a claim—often after weeks of delay and financial hardship. In the **sentient state**, an agent monitors a "Regional Economic Health" data product. It detects a WARN Act notice at a local factory. It identifies affected workers via secure data matching with employment records, verifies recent wages, and proactively contacts them: "We are aware of the layoff at Factory X. You are likely eligible for **$400/week in unemployment benefits**. Click here to confirm your details and activate your claim." Benefits reach people within **minutes, not weeks**. Government shifts from a passive vending machine to an **active nervous system** that senses events and responds—dramatically reducing the time-to-value for vulnerable populations.

It is essential to clarify what proactive governance is *not*. These agents are not omniscient surveillance systems. They operate within strict legal and ethical bounds, monitoring only sanctioned data streams (official layoff notices, eligibility data agencies already collect), and acting only on triggers defined by policy and public need. A robust **Just-in-Time Consent Flow** ensures the citizen controls engagement: the outbound message itself serves as a consent prompt. By clicking to proceed, the citizen grants permission for the agent to access the necessary personal records and complete enrollment. If they ignore or decline, the agent takes no further action with their data. Citizens experience outreach as timely assistance *which they control*, not as intrusion.

4.5 Taming the Demon: HITL Governance

Autonomous agents in government require governance as rigorous as any human workforce. Seven layered safeguards ensure accountability:

1. Agent Contracts. Every agent operates under a strict digital contract defining its bounds, resource budget, and ethical guardrails. **2. Human-in-the-Loop (HITL).** For high-stakes decisions—benefit denial, fraud investigation, child

removal—the agent prepares the case and makes a recommendation, but **a human caseworker must execute the final action**. AI augments; it does not replace. **3. Validation Agents.** Dedicated "critic" agents review outputs of other agents, flagging errors, inconsistencies, or bias for human review. **4. Explainability.** The ReAct reasoning trace serves as the legal record; if a citizen appeals, the agency produces the exact logic in plain language. **5. Just-in-Time Data Access.** Agents retrieve personal data only at the precise moment needed, tied to an active case or user consent, with ephemeral queries that minimize privacy exposure. **6. Ephemeral Identity.** Temporary credentials that expire on task completion. **7. User Consent and Control.** No proactive enrollment or cross-agency data sharing without the citizen's knowledge and permission.

Through these seven layered safeguards, AI agents in government are given clear rules of engagement and are monitored and directed with the same rigor as a human workforce. They become **force-multipliers for public servants**—handling the drudgery and complexity at digital speed while **ultimate accountability and judgment remain human**. The architecture does not merely tolerate oversight; it is *designed for it*. Every decision, every data access, every recommendation exists as a readable, auditable record.

This governance framework directly addresses the 2026 regulatory cliff. The HIPAA overhaul's new audit trail requirements are met by the ReAct reasoning trace logged at every agent decision point. The California Delete Act's data-rights mandates are honored through Just-in-Time data access and ephemeral agent identity—agents never retain data beyond the scope of a single session, and deletion requests can be propagated cleanly through the Data Mesh's domain-owned products. The framework ensures that as agencies deploy increasingly autonomous systems, they do so within a governance structure that regulators, legislators, and the public can verify and trust.

5. Operationalizing Innovation

Technology alone cannot fix broken government. The Composable Agency requires new operational "software"—people, culture, and process models drawn from proven social-innovation programs.

5.1 Delivery-Driven Policy (Code for America)

The lethal gap between policy intent and technical execution is closed by embedding delivery teams in the policy-drafting process. Technical feasibility, user experience, and on-the-ground realities are considered *during* legislation—not thrown over the wall afterward. The traditional waterfall policy model—legislators write a law, then throw it to IT for implementation—is replaced by iterative feedback loops where real user data shapes both policy and delivery simultaneously.

Case Study — Clear My Record (California). California passed a law to expunge certain marijuana convictions, but the manual, paper-based process was so burdensome that few people benefited. Code for America did not merely build a website; they partnered directly with District Attorneys to reimagine the delivery process from end to end. They developed an algorithm to automatically identify and clear eligible records in bulk—bypassing the case-by-case adjudication that would have taken years. Result: **144,000+ convictions cleared in months**, restoring rights and reducing barriers to employment and housing for tens of thousands. In the Agentic Agency, the Orchestrator Agent extends this model: every digital interaction feeds real-time instrumentation data back to policymakers ("40% of applicants drop off at Question 5"), enabling continuous, evidence-based policy refinement that was previously impossible.

5.2 Active Contract Management (Harvard Kennedy School GPL)

As government relies increasingly on modular vendors (providing PGC components) and autonomous agents, the **Active Contract Management (ACM)** framework developed by the Harvard Kennedy School Government Performance Lab becomes essential. ACM shifts contract management from a compliance mentality (checking boxes, enforcing terms after the fact) to an outcomes orientation (managing performance in real time through high-frequency, data-informed collaboration).

Case Study — Seattle Homelessness. Seattle's Human Services Department used results-driven contracting to transform its homelessness response. Instead of siloed contracts measuring activities ("shelter beds filled"), they shifted to outcome metrics: **"exits to permanent housing"** and **"returns to homelessness."** By actively managing contracts—meeting regularly with providers to review data, identify barriers, and adapt strategies collaboratively—Seattle consolidated programs, eliminated underperforming interventions, and focused resources on approaches with demonstrated results. The outcome: **dramatically improved system performance** and more people reaching stable housing. Applied to the Agentic Agency: leadership holds **"Agent Performance Stat" meetings**—structured reviews of the digital workforce's dashboards. Are eligibility agents reducing backlogs? Is the virtual assistant resolving citizen inquiries? Are there outcome disparities across demographics? This active management ensures the Agentic Agency stays aligned with public values and enables rapid course-correction when data reveals unintended consequences.

5.3 The Impact Lab Model

Innovation should not happen on live production populations without rigorous testing. Modeled after Stanford Impact Labs and Georgetown's Beeck Center, **Public Impact Labs** create structured government-academic-technology partnerships where new PGCs and AI-driven workflows are piloted in controlled environments before scaling. Embracing academic rigor, these partnerships test interventions via randomized controlled trials (RCTs) or robust A/B testing—measuring accuracy, speed, and equity against traditional processes.

Stanford's Regulation, Evaluation, and Governance Lab (RegLab) partnered with the IRS to modernize tax audit selection. They used machine learning not just to catch evasion but to **identify and mitigate racial and income bias**

in audit algorithms—demonstrating that AI can increase both fairness and effectiveness in government when deployed with empirical oversight. These labs also serve as vital **talent pipelines**: by drawing data scientists, designers, and engineers into high-impact public problems, they inspire a new generation of civic technologists who often move into permanent government roles—helping agencies address the talent cliff by infusing fresh skills and perspectives.

5.4 Michigan's "Project One Day"

The power of combining these operational models is exemplified by the Michigan Department of Health and Human Services (MDHHS) partnership with Civilla on an initiative nicknamed **Project One Day**.

The Problem: Michigan's public benefits application was **40 pages long—the longest in America**. Processing took weeks. Both applicants and staff were overwhelmed by complexity and bureaucracy, and the sheer volume of paperwork created backlogs that delayed aid to families in crisis. **The Intervention:** Using human-centered design—a core social innovation practice—the team radically redesigned the application, cutting it by **80%**. They then reimagined the entire benefits enrollment workflow under a "One Day" vision: the audacious goal that an application could be processed in a single day instead of weeks. **The Result: 90% application registration within two hours** of submission, as opposed to applications languishing for days. Time-to-determination for benefits was drastically reduced. Critically, this success was not purely a technology upgrade—it was a combination of **composable, headless interface principles** with **delivery-driven process reform**. The result delivered massive gains in both operational efficiency and human dignity: people got help faster, and the process felt respectful rather than degrading. Project One Day demonstrates that when architectural modernization and social innovation operate in concert, the impact is transformative.

The Seamless Experience: Maria

Proof of Concept: The Composable Agentic Agency in Action

This narrative integrates every architectural layer—PGCs, Data Mesh, Multi-Agent System, HITL governance—into a single constituent journey.

The Trigger. Maria, a single mother, loses her job. She texts a government help line: "I lost my job and I need help."

Agentic Orchestration. An Intake-Orchestrator Agent parses her natural-language request, recognizes intent, and initiates a Holistic Support workflow spanning multiple programs.

Identity Verification. The agent calls the **Identity PGC**, which sends a secure biometric challenge to her phone. Maria verifies with fingerprint/face ID.

Data Mesh Discovery. Authenticated, the Data Agent queries the Department of Labor's "Wage Data" product through the secure mesh. It finds a layoff notice filed yesterday and retrieves Maria's wage history—her recent salary, the abrupt termination—*without asking Maria to upload a termination letter, dig through her email, or re-enter information the government already has*. The agent also queries a "Household Composition" data product to confirm she has dependent children, and checks the "Health Coverage Status" product to see whether her employer-provided insurance is about to lapse.

Composable Assembly. The Orchestrator identifies eligibility for Unemployment Insurance, SNAP, and Medicaid. It invokes the **Eligibility & Enrollment PGC** to run Maria's data against all three programs' policy rules simultaneously.

Proactive Delivery. Instead of three separate applications, the agent pre-fills and presents a single unified confirmation: "You are eligible for **$400/week in UI, $300/month in SNAP, and Medicaid coverage** for you and your children. Tap to accept." Maria reviews, consents, and accepts.

Active Governance. A Validation Agent reviews the benefit package against policy rules, anomaly thresholds, and historical precedents. No issues are flagged; verification is logged. Had there been a low-confidence score or red flag, the case would have been routed to a human caseworker with a full reasoning trace.

Outcome. Within minutes, Maria has UI processing, an EBT card en route, and confirmed health coverage. She did not visit three offices, navigate three portals, or wait in bureaucratic limbo. The entire interaction—from "I lost my job" to full enrollment across three programs—took less time than ordering a meal delivery. The state, on its end, delivered critical aid through an efficient, integrated digital concierge while maintaining a complete audit trail of every agent decision, every data access, and every consent granted. The interaction was **dignified, fast, and effective**—and it was **accountable** at every step.

6. The 3-3-3 Roadmap

Transformation at this scale requires disciplined execution in **90-day pulses**—each delivering measurable value, building institutional confidence, and de-risking the next phase.

Pulse 1: Foundation (Days 1–90)

Operating Model: Stand up 2–3 empowered Product Teams for high-reuse PGCs (Identity, Notifications, Consent). Shift funding from siloed projects to product-centric teams. **Architecture:** Deploy the first PGCs in production using the Strangler Fig pattern against the highest-friction legacy touchpoints. **Data:** Launch Data Mesh platform infrastructure (cloud environment, catalog, security layers). Select one high-impact domain—Medicaid recommended—to develop and publish the first certified data product. **Governance:** Establish an Impact Lab partnership (university or civic-tech group) to advise on ethical AI deployment and independently evaluate early outcomes.

Pulse 2: Pilot (Days 91–180)

Agentic Deployment: Deploy an Orchestrator Agent for a specific, high-volume life event ("Having a Baby," "Starting a Small Business," or "Losing a Job") that requires multi-program coordination. Select a life event where the friction tax is measurable and the cross-program touchpoints are well understood—this ensures the pilot demonstrates the full stack (PGCs, Data Mesh, MAS) rather than just a chatbot layer. **Instrumentation:** Instrument the pilot heavily—collect metrics on time-to-determination, completion rates, drop-off points, agent accuracy, and demographic equity. Feed insights in real time to both product teams and policy teams for iterative refinement. Use A/B testing where possible, comparing agent-assisted flows against traditional processes to build an evidence base for expansion. **Active Management:** Launch Agent Performance Stat meetings involving product teams, agency leadership, and Impact Lab partners. Review dashboards monthly; adjust policies, agent behavior, and PGC configurations based on evidence. Establish clear escalation protocols for when agents encounter edge cases or low-confidence scenarios, ensuring human caseworkers are seamlessly integrated into the workflow.

Pulse 3: Scale (Days 181–270)

Expansion: Scale PGC library to cover 80% of agency functions, adding Fraud Detection, Advanced Analytics, and Inter-Agency Data Sharing modules. Encourage cross-department and cross-state reuse to drive down costs and improve quality through shared learning. **Proactive Workflows:** Activate sentient workflows on pre-approved triggers—each subjected to ethical review and consent design before enablement. Start with high-value, low-risk scenarios (e.g., proactive renewal reminders) before expanding to rapid-response workflows. **Ecosystem:** Open the PGC and agent marketplace beyond the agency. Foster a cross-jurisdictional ecosystem where states contribute and adopt shared modules, creating network effects that accelerate innovation nationwide. A rich marketplace of government-tested APIs and agents that every state can draw from—reducing duplication, increasing quality, and lowering the barrier to modernization for smaller agencies.

Conclusion

The rigid, monolithic systems of the past are collapsing under their own weight. The 2026 regulatory cliff—the HIPAA Security Rule overhaul demanding encryption and audit capabilities, the California Delete Act requiring automated data-deletion infrastructure, and the talent exodus removing the last generation capable of maintaining legacy mainframes—will break agencies that have not begun this transition. The question is not whether to modernize, but whether to do so on your terms or under duress.

By embracing the **Composable Agency**, government sheds the crushing weight of Strategic Debt, replaces vendor lock-in with digital sovereignty, and unlocks the architectural agility to respond to policy changes in days rather

than years. By deploying the **Agentic Agency**, it scales its capacity to care—moving from reactive bureaucracy to proactive, life-event-driven support that meets citizens where they are, when they need help, with the dignity they deserve. By anchoring transformation in the human-centered rigor of social innovation—delivery-driven policy, active contract management, impact labs—and enforcing strict consent and HITL guardrails at every layer, we ensure this new machine remains a **servant of the people, not a master**.

The future of government is **Composable, Agentic, and Human**. The blueprint is here. The regulatory clock is ticking. The time to build is now.