



Conduit & ConduitX

A Unified Integration Fabric for Modern Healthcare Platforms

Reducing System Integration to Days Not Weeks

An AllyGPO White Paper prepared by:

Ravi Seshadri
Chief Technology Officer

Veena Shetty
VP, Data & Analytics Architecture

Govind Bundawala
Principal Architect

Suresh Nakeeran
Data Architect

Executive Summary

Systems integration in healthcare has been treated as a development problem for far too long. Historically, every new vendor means a new codebase, a new authentication implementation, a new data model, and a new maintenance burden. The result is integration effort that scales semi-linearly with the number of connected systems, which may drive cost and consume engineering capacity that could be spent on value-added innovation – not connecting pipes.

Conduit and ConduitX address this structural inefficiency across configuration, testing, and production-ready environments. Together, they form a unified integration fabric that often reduces vendor onboarding from weeks or months to days – reducing integration complexity, accelerating deployment timelines, and enabling a flexible ecosystem without having to build costly point-to-point integrations over and over again.

The Result

Integration is no longer a bottleneck. It becomes an accelerator. **Conduit and ConduitX enable AllyGPO and its member practices to scale faster, integrate smarter, and operate with confidence across a complex and fragmented healthcare ecosystem.**

Conduit is a configuration-driven data pipeline engine that standardizes how systems connect, communicate, and exchange data.

ConduitX is a hardened webhook runtime that ensures reliable, durable, and observable ingestion of external events. Together, they transform integration from a bottleneck into an accelerator.

Real-World Application & Value Demonstrated

“In the nearly 6 years that our organization has been using our current EMR and Practice Management system, we have undergone many integrations with various software solutions. With each integration, there have been countless problems; the process has taken many months, or the software in question has not been fully integrated. Until AllyIQ. This integration was seamless and completed in less time than projected, making the transition from our previous Inventory Management System as painless as it could have been.”

- Clinic Manager

“In a seemingly overwhelming practice start-up period, AllyGPO was one of the few easy decisions and integrations. The flow of data between my EMR/PM and AllyGPO's AllyIQ platform just works - AllyIQ knows which of my patients are scheduled, and my EMR receives data on my injectables with a click of a button. Beyond just inventory management, AllyIQ helps me stay current and ahead of constantly changing injectable-related practice economics. The best systems are those that you don't need to think twice about; in my practice, AllyIQ is one of those.”

- Founder & Practicing Physician

The Integration Challenge for Community Practice

AllyGPO operates across specialty healthcare segments in oncology and retina, and must connect with dozens of practice management systems, specialty distributors, scheduling platforms, payer systems, and clearinghouses. Each system speaks a different language, and many leverage use unique protocols or API keys. Layered on top are healthcare data standards, such as HL7, FHIR, EDI, XML, JSON, with individual parsing requirements and edge cases.

The complex structure of these environments and platforms often render existing integration tools inadequate.

Configuration

Configuration is never enough. Typically, existing integration technologies require scripting beyond simple configuration. Integration efforts do not grow with scale – it grows disproportionately because each new vendor introduces not just new configuration, but new custom logic, new test cases, and new failure modes.

Tool Fragmentation

Tool fragmentation forces complexity. Handling REST APIs, webhooks, SFTP file drops, and healthcare messaging protocols within a single platform requires either complex extensions to general-purpose tools or multiple specialized tools stitched together. Both approaches create operational overhead and architectural fragility.

Agentic AI

Agentic AI changes the equation. Within AllyIQ, AI-assisted workflow agents support configurable orchestration of operational and financial workflows, subject to user permissions, system controls, and human review where required. Integration is no longer limited to static, human-configured pipelines; authorized agents can help identify, configure, and invoke approved data flows within defined governance parameters.

Design Principles

Conduit and ConduitX were designed around four non-negotiable principles that directly address the failures of existing approaches.

Configuration-first, code-minimized

Adding a new vendor requires a YAML file describing its base URL, authentication method, endpoint definitions, and response mapping rules. This approach reduces the need for custom engineering, deployment cycles, and regression risk. The configuration is the integration.

Unified integration surface

One platform designed to handle APIs, webhooks, SFTP, messaging systems, and healthcare protocols, reducing the need for protocol-specific middleware or stitched-together integration tools.

Canonical data model

Every external payload is normalized into AllyGPO's internal schema before it reaches a consuming application. Downstream systems are fully decoupled from vendor-specific data structures.


Production-grade resilience

Durable ingestion, retry management, replay capability, dead letter handling, and full observability are built into the core – not bolted on as operational afterthoughts.

Conduit: The Data Pipeline Engine

Conduit is the core integration fabric. It connects any supported source to any supported destination through a unified, configuration-driven pipeline that processes data through five discrete stages. A Job Manager creates jobs with the required parameters, triggers execution, and monitors completion. A Config Store provides vendor-specific and customer-specific parameters at each stage.

Figure 1: Conduit Pipeline Architecture — Five-stage sequential pipeline with Job Manager orchestration



Request Flows Through Conduit

- When a request enters Conduit from an application, an AI agent, or a scheduled job, the caller specifies the vendor, the endpoint, and the parameters using Conduit's unified request format — it speaks one language regardless of the destination.
- Conduit loads the vendor's YAML configuration: base URL, auth method, endpoint definitions, and any customer-specific overrides for the requesting practice. No code executes—the configuration fully describes the integration.
- Authentication is handled invisibly. Whether the vendor requires OAuth2 password grant, OAuth2 client credentials, API key, bearer token, or basic auth, Conduit manages it automatically, including token refresh and secure credential storage.
- The response is validated, parsed via JSONPath, mapped to AllyGPO's canonical schema, and delivered to the destination. The consuming system never encounters vendor-specific data structures.

Format Translation

Conduit supports bidirectional translation across common healthcare data formats through pluggable parsers: HL7 to JSON, FHIR to JSON, EDI to JSON, and XML to JSON. Format translation is configuration-driven—adding support for a new format variant requires a parser plugin, not a code change to the pipeline itself.

Multi-Tenant Configuration


Each integration supports two layers of configuration: a vendor-level baseline that defines authentication, endpoints, and response mapping, and a customer-specific overlay that customizes behavior per practice. This enables a single Conduit deployment to serve hundreds of practices with vendor-specific and practice-specific variations without code divergence.

ConduitX: The Webhook Runtime

While Conduit handles orchestrated and outbound data flows, ConduitX is purpose-built for a fundamentally different problem: inbound event ingestion from external partners sending webhooks. Shipping carriers, pharmacy systems, specialty distributors and others all push events into AllyGPO's ecosystem and each requires durable, validated, observable receipt.

ConduitX is not a general-purpose webhook receiver. It is a hardened runtime engineered around a single design principle: never acknowledge before durability. The HTTP 2xx response is returned only after the raw payload and headers have been persisted to blob storage.

Figure 2: ConduitX Webhook Runtime Architecture — Layered ingestion with partner isolation and replay



The Canonical Envelope

Every event processed by ConduitX is wrapped in a canonical internal envelope before it reaches any downstream consumer. The envelope includes the partner identifier, event type, version, a reference to the raw payload in blob storage, and a correlation ID for end-to-end tracing. Downstream systems within the AllyGPO ecosystem consume normalized events, never partner-specific formats.

Retry Domain Separation

ConduitX explicitly separates two retry domains that most webhook systems conflate. External retries – when a partner re-sends an event because they didn't receive a 2xx – are the partner's concern. Internal retries – when a worker fails to process or publish an event – are handled by the Retry Policy Executor within the WorkerPool. This separation prevents cascading failures and ensures that a downstream outage never causes a partner to flood the ingestion layer with duplicate events.

Preparing for Agentic Integration

Conduit was designed with a capability that distinguishes it from conventional integration engines: it is agent-configurable. AllyIQ's AI agents can discover available pipelines, provision new ones, and invoke data flows dynamically and much less human intervention.

This becomes increasingly important because the next generation of healthcare platform operations will not be driven by static, pre-configured integrations. Agents performing demand forecasting will need to pull inventory data from distributors in real time. Agents optimizing rebate tiers will need to invoke contract data from multiple sources simultaneously. Agents managing shortage response will need to trigger alternative sourcing workflows that span multiple vendor systems.

Conduit's YAML-driven, API-first architecture makes this possible. An agent does not need to understand vendor-specific protocols because it interacts with Conduit's unified interface and lets the engine handle the complexity beneath.

Operational Impact

The combined effect of Conduit and ConduitX on AllyGPO's integration operations is measurable across four dimensions.

- 01 Integration Velocity**
Vendor onboarding moves from weeks or months to a few days. Configure, test, deploy. No custom adapters, no bespoke code, no extended engineering cycles.
- 02 Engineering Leverage**
Integration work shifts from development to configuration. Engineers can now spend time on value-added innovation instead of writing custom vendor adapters.
- 03 Operational Reliability**
Durable ingestion, dead letter queues, replay capability, and per-event observability ensure that data movement is auditable, recoverable, and transparent.
- 04 Scaling Independence**
Adding more vendors, more practices, or higher data volume does not require major architectural changes. The same platform, same configuration model, and same operational playbook apply at any scale.

Integration via Conduit and ConduitX now becomes an accelerator and enables AllyGPO and its member practices to scale faster, integrate smarter, and operate with confidence across a complex and fragmented healthcare ecosystem.


About AllyGPO

AllyGPO® is a leading innovator and next-generation specialty group purchasing organization (GPO) equipping independent, community-based practices with advanced, technology-driven solutions and proprietary tools to navigate the complex specialty drug market. By combining comprehensive drug contracting services with intelligent inventory management, predictive analytics, and hands-on operational support, AllyGPO delivers actionable insights that enhance practice sustainability and performance. Its AI-enabled AllyIQ® platform, purpose-built for specialty care, integrates data from multiple practice systems to enable clinical, operational, and financial teams to work from a shared view and act with greater speed and confidence. Headquartered in Frisco, Texas, AllyGPO partners with community-based specialty practices and pharmaceutical manufacturers across the U.S. [Learn more at www.AllyGPO.com](http://www.AllyGPO.com).

Content is provided for general informational purposes only and does not constitute medical, legal, regulatory, reimbursement, coding, billing, financial, technical or other professional advice. AllyGPO does not practice medicine, and participating providers remain solely responsible for all clinical, prescribing, purchasing, and patient-care decisions. Savings, rebates, pricing, contract opportunities, operational improvements, and financial results vary by practice, eligibility, utilization, payer mix, product availability, and other factors, and are not guaranteed. Any data, analytics, benchmarking, education, promotional communications, or other insights made available through AllyGPO are intended to support independent decision-making and are subject to applicable law, contract terms, and program availability.

Appendix


Figure 1: Conduit Pipeline Architecture — Five-stage sequential pipeline with Job Manager orchestration



Stage	Responsibility
DataSource	Connects to vendor APIs, SFTP endpoints, and messaging systems. Handles authentication transparently across OAuth2, API key, and bearer token methods.
Parsing	Translates inbound formats—HL7, FHIR, EDI, XML—into a normalized JSON representation using pluggable parsers.
Mapping	Applies JSONPath-based field mapping and customer-specific overlays to produce canonical data structures.
Formatting	Reshapes output into the format required by the destination system, including schema validation.
Output	Delivers data to the target: APIs, databases, message bus, SFTP, or local filesystem. Supports persistence and streaming modes.

Appendix

Figure 2: ConduitX Webhook Runtime Architecture — Layered ingestion with partner isolation and replay



Layer	Function
HTTP Ingress	Webhook endpoints with partner resolution via configuration. Signature validation (HMAC, timestamps) ensures only authentic events are accepted.
Durable Receipt	Raw payload and headers persisted to blob storage immediately. HTTP 2xx acknowledgment returned only after durability is confirmed.
Queue Decoupling	Events enqueued to a durable job queue, fully separating ingestion velocity from processing throughput.
WorkerPool	Concurrent, partner-isolated processing: event type resolution, idempotent deduplication, canonical envelope construction, and downstream publishing.
Resilience	Dead letter queue for failed events, isolated retry policies (internal vs. partner retries), and full replay from the envelope store.
Observability	Per-event traceability: event ID, partner, timestamps, latency, errors, and correlation IDs across the entire chain.