



Accelerating healthcare transformation through intelligent delivery





The challenge facing healthcare payers

Healthcare payers must modernize core administration platforms faster while navigating greater business complexity, regulatory change, cloud transformation and cost pressure. Product proliferation across Medicare Advantage, Medicaid and commercial lines, evolving mandates, legacy platform constraints and integration dependencies can limit growth and scalability. Traditional implementation models are constrained by:

Dependence on tribal knowledge and scarce subject matter experts (SMEs)

Lengthy requirements and discovery cycles

Manual benefit and provider contract configuration

Complex system integration development


Data migration risks and prolonged conversion cycles

Extensive testing and validation effort

Cloud installation complexity

If left unresolved, these constraints delay revenue from new products and markets, strain member and provider experiences, and increase regulatory and compliance risk.

At the same time, business leaders expect:



Moving AI from productivity gains to technology transformation

AI operates as an accountable execution layer within the delivery toolchain rather than a collection of disconnected point tools.

For healthcare payers, AI adoption is moving beyond incremental productivity gains toward enterprise transformation. The greatest value comes from embedding intelligence across the implementation lifecycle to create reusable assets, modernize delivery workflows, accelerate decisions and improve the workstreams that determine program speed, cost, quality, predictability and business readiness.

In practice, a consistent set of capabilities is applied across every implementation track: document and data intelligence that converts unstructured plans, contracts and legacy specifications into structured, machine-usable configuration; automated validation that reconciles configured output against source intent before issues reach production; and agentic orchestration—multiple specialized AI models, each with a defined role, coordinated by a supervising orchestrator—that compresses work once passed sequentially between teams. The same intelligence


extends into data migration and quality engineering. Critically, these capabilities are governed centrally through versioned prompt patterns, curated knowledge assets, controlled model selection and managed consumption, so clients move from experimentation to scaled execution quickly.

Beyond time-to-market and effort reduction, an equally critical dimension of this challenge is protecting intellectual property: As TriZetto® and Facets® product knowledge, configuration logic and delivery assets increasingly pass through gen AI tools operated by other vendors or clients, unmanaged exposure creates real IP risk. Equally important is the ability to capture gen AI's productivity benefits even in engagements where clients restrict the use of external LLM or gen AI tools within their environment. The vendor that resolves both challenges—protecting proprietary IP while still delivering AI-driven productivity regardless of client-side tool restrictions—stands to capture disproportionate market share and valuation advantage.



Cognizant is addressing these challenges through an **AI-infused TriZetto implementation factory** that embeds AI across the implementation lifecycle to improve speed, quality, consistency and scalability. The roadmap targets **30%–35% effort optimization across the greenfield implementation lifecycle**, with gains across platform foundation, configuration, system integration, data migration and quality engineering.

Effort optimization by workstream



Ranges are indicative delivery-side effort optimization targets and should be confirmed against the reference programs cited in this paper.

Figure 1. Effort optimization by workstream. Ranges are indicative delivery-side targets and should be confirmed against the reference programs cited in this paper. PMO and program management carries some AI enablement; governance, client-side decisions, sign-off, reviews and UAT are not AI-addressable.

Outcomes for a typical greenfield implementation

- 30%–35% effort optimization across the greenfield implementation lifecycle
- Up to 30% faster time to market for new products, markets and lines of business
- Savings reflected in the commercial construct at proposal stage
- Reduced dependency on scarce SMEs, with Facets domain knowledge codified into reusable assets


AI infusion across the implementation lifecycle

The five sections that follow describe each delivery tower in turn—the client challenge it addresses, how AI is infused and the outcomes clients can expect. Figure 4 at the close of this paper summarizes all five on a single page.

AI-powered platform foundation

Client challenge

Platform provisioning, environment setup, deployment orchestration and infrastructure standardization often create delays at the beginning of transformation programs. This typically impacts the program's ability to demonstrate early wins and sustain executive confidence in the transformation investment.



AI infusion

An AI-powered TriZetto application deployment tool supports infrastructure-as-code (IaC) generation to standardize and simplify infrastructure creation. The DevOps-agnostic automation generates deployment scripts that integrate with the client's existing enterprise CI/CD tools, avoiding the need to build new automation from scratch. This accelerates cloud readiness while improving repeatability and operational control.

Beyond automation, AI contributes a decision and assurance layer: recommending

environment topology and sizing based on comparable prior implementations, generating and reviewing IaC templates and deployment runbooks, detecting configuration drift and environment inconsistencies before they surface as downstream defects and flagging provisioning bottlenecks from historical program telemetry. The result is a foundation that is not only automated, but also self-validating and consistently reproducible across development, test, and production environments.

Client outcomes

8%–10% TriZetto application platform setup effort optimization potential

Faster project mobilization and earlier demonstrable progress

Improved deployment consistency, governance and repeatability across environments

Accelerated cloud readiness

AI-powered business configuration design and build

Client challenge

Benefits, contracts, claims workflows, billing, provider setup, and Medicare and Medicaid configurations demand extensive manual interpretation and repetitive effort, elevating the risk of member and provider abrasion and exposing payers to compliance and audit risk under CMS and state regulatory timelines. Pressure peaks around fixed program milestones—the AEP chief among them—where configuration, testing and certification must converge within a nonnegotiable window, and any slippage cascades into delayed market readiness.



AI infusion

Benefits configuration

AI extends benefits configuration through the TriZetto® Configuration Orchestrator, which converts benefit product data into core, Facets-ready configuration. AI-driven extraction and classification—interpret benefit terms, and utilization management logic, liability and limit values into a structured configuration matrix that is mapped and validated before loading into the target platform. A built-in validation layer checks accuracy on load, and an outbound reconciliation extract confirms configured data stays aligned with source intent—shortening build cycles, strengthening data accuracy and reducing manual effort.

Provider contract configuration

AI-enabled provider contract validation strengthens pricing accuracy by extracting reimbursement terms, carve-outs and amendment history from contract documents that vary widely in format and language, then translating them into structured pricing configuration. Cognizant has operationalized this as a generative AI assurance routine: it reads rate and service-level detail from base provider contracts held as unstructured PDFs, parses compensation grids maintained in spreadsheet form, reconciles the two sources and reports every discrepancy for resolution before configuration is loaded. For payers managing large, frequently amended provider networks, this compresses contract-to-configuration turnaround while strengthening pricing integrity and provider trust.

Industry configuration accelerators

Pre-built configuration assets help clients accelerate implementation by reusing standardized product, system, code and claims-rule patterns instead of starting every configuration cycle from scratch.

Client outcomes

25%–30%
configuration effort
optimization potential

Faster benefit design
and build cycles

Fewer configuration defects
and better configuration
quality, improving member
and provider experience

Stronger compliance readiness
and greater confidence in
meeting open enrollment and
regulatory deadlines

Accelerated annual
renewals



AI-powered system integration and custom development

Client challenge

Integration programs often require significant effort across requirements, design, code development, testing, deployment and project coordination. This effort-intensive cycle extends time to market, increases delivery cost and limits the organization's agility to respond to emerging business, regulatory and market opportunities.



AI infusion

Industrialized AI delivery, anchored in core platform domain expertise

Cognizant brings two complementary AI-enabled delivery models to system integration and custom development: the **system integration (SI) AI factory** and the **multi-agent delivery model**. The factory model suits large-scale, enterprise-wide programs where governance, standardized execution and scale across many interfaces and product lines are the dominant requirements. The multi-agent model targets well-bounded work products where autonomous agents can compress a repetitive lifecycle—mappings, technical specifications, code generation, validation artifacts, and inbound and outbound extract development.

The choice is driven by engagement size, technical and functional complexity, delivery scope, client context and the level of automation the client is ready to adopt. On most enterprise programs the two operate together: The factory model provides the governed delivery backbone, with multi-agent accelerators deployed inside it wherever a work product is a strong automation candidate.

Cognizant system integration AI factory model

Cognizant's SI AI factory turns system integration from a manual, document-heavy delivery model into an industrialized, AI-powered execution engine,

bringing together client-facing TriZetto implementation teams, Facets SMEs, AI-enabled accelerators, engineering copilots and reusable delivery assets.


The model establishes clear ownership at every layer. Client-facing teams manage priorities, business decisions and stakeholder alignment. Platform SMEs translate payer processes, legacy behavior, configuration nuances and platform constraints into implementation-ready guidance. The factory team then converts that guidance into standardized, AI-assisted outputs, including user stories, interface specifications, mapping logic, build components, test assets and deployment runbooks.

What makes this a true factory model is the repeatable rhythm of intake, SME enrichment, AI-assisted production, expert validation and client review. Every engagement strengthens the next by converting delivery learnings into reusable assets, reducing reliance on scarce expert capacity.


Equally important, the AI layer sits inside the Cognizant-managed factory rather than as fragmented, project-by-project tooling. That gives centralized control over AI platform usage, prompt patterns, reusable knowledge assets, model selection and token consumption—and gives clients a more predictable cost model with less exposure to uncontrolled AI spend.

The Facets implementation factory model in practice

Client-facing teams own the client. The factory, with AI inside it, owns industrialized production.



One governed pipeline from business requirement to production-ready code




Observed on a large US health plan implementation

- Two months of schedule recovered against a three-month requirement delay
- 40% development effort saving across 40 interface builds
- About 25% reduction in overall system integration development effort

Figure 2. The Facets implementation factory model in practice. Client-facing teams retain ownership of the client relationship and production deployment; the factory, with the governed AI layer inside it, owns industrialized production. Results shown were observed on a large US health plan implementation and are indicative rather than guaranteed.

Critical AI factory value points



Multi-agent solution delivery

Where the constraint is limited platform SME availability, manual coding bottlenecks and a steep learning curve on unfamiliar technologies, Cognizant deploys a crew of specialized AI agents that mirrors how expert teams collaborate—with the Model Context Protocol (MCP), an open standard for connecting AI models to enterprise systems and data, grounding every agent in the real product data model.

The pipeline runs in three orchestrated stages. A Map2Model agent takes the business requirement document (BRD) and produces the layout mapping document. A Model2TSD agent consumes the BRD and that mapping to generate the technical specification document (TSD), unit test document and requirement traceability matrix (RTM). A TSD2Code agent then produces standards-compliant C# and SQL components with control totals—record and value reconciliation counts that prove


an extract moved the data it was supposed to move—checked into the repository and deployed through the pipeline.

Within each stage, business analyst, system analyst, SME, developer and reviewer agents iterate under a group orchestrator that collaborates, validates and refines before the stage releases its output, with human review retained as the final governance checkpoint.

Outbound extract development on SQL Server and .NET is the strongest early candidate, where agents automate much of the requirement-to-code lifecycle and engineers validate the result. Inbound interfaces and extensions—typically spanning Oracle and other heterogeneous stacks—are the next frontier, where agents assist with inventory analysis, complexity classification, mapping and specification generation.

Multi-agent solution delivery

One business requirement document in, a deployable solution out—three orchestrated agent stages



Every agent is anchored through the MCP in the actual Facets data model, scaffold metadata and Cognizant coding standards.

Figure 3. Multi-agent solution delivery. One business requirement document in, deployable and traceable components out: three orchestrated agent stages, each running a crew of specialized agents under a group orchestrator, with human review retained at every stage boundary.

Critical multi-agent solution value points

~40% design and build dividend:

Roughly 40% less design and build effort on extract development, with reduced SME dependency and consistent, standards-aligned code

Traceability built in, not bolted on:

Requirement traceability matrices, unit test documents and control totals are generated as part of the flow, supporting audit and validation from the outset



One BRD in, deployable solution out:

A single business requirement document drives layout mapping, technical specifications, unit tests, traceability and production-ready C# and SQL without manual handoffs between stages

An accelerator, not a substitute:

Deploys inside broader modernization and implementation programs as a targeted accelerator, complementing rather than replacing the factory model

Client outcomes

30%–35% system
integration effort
optimization potential

Faster solution
delivery and
improved agility
to respond to
evolving market and
regulatory demands


Higher development
quality and greater
engineering
productivity

Reduced
documentation
effort

AI-powered data migration factory

Client challenge

Data migration remains one of the highest-risk implementation workstreams because of extensive profiling, mapping, code generation and validation activities. Migration setbacks can jeopardize go-live timelines, disrupt continuity of claims and member data, and create downstream business risk if data quality issues surface after cutover, impacting member trust, regulatory reporting and operational continuity.



AI infusion

Migration discovery and strategy— Deliberately human-led

The migration strategy phase remains a governance-heavy, human-led activity where expert judgment is required to define migration scope, sequencing and cutover approach in close alignment with business timelines, member communication windows and regulatory commitments. AI supports this phase with analysis and options; it does not set the strategy.

Mapping modernization

AI-powered profiling and mapping reduce requirements and design effort by generating data insights, accelerating SME mapping sessions and recommending intelligent source-to-target mappings. Profiling exposes actual data distributions, null and default patterns, orphaned relationships and value domains in the legacy estate—the conditions that most often derail a conversion cycle when discovered late.

Migration code generation

AI transforms migration builds by converting mapping rules into executable SQL logic, reducing manual code generation effort and accelerating validation readiness. Generated logic is version-controlled and reviewed by migration engineers before it is executed against any client data.

Validation intelligence—Control retained with the delivery team

Data load and validation remain control-focused activities, with AI value concentrated upstream in profiling, mapping and code generation to improve readiness and reduce downstream risk. Reconciliation, sign-off and cutover authorization stay with named delivery and client leads.

Client outcomes

**Targeted 40%–45%
effort optimization
opportunity**

**Faster conversion
cycles**

**Improved migration
quality and reduced risk of
member and provider data
disruption post-cutover**

AI-powered quality engineering

Client challenge

Testing often consumes one of the largest portions of implementation effort due to test creation, data generation, execution, defect management and migration validation. Extended testing cycles delay go-live, increase implementation cost and elevate the risk of undetected defects reaching production, potentially disrupting claims processing, member and provider servicing, and regulatory compliance after launch.



AI infusion

Test planning and design

AI-enabled test design converts business requirements and prior test assets into structured high-level and detailed test cases with greater consistency and traceability, improving coverage of critical business flows and freeing quality engineering teams to focus on risk-based prioritization rather than case-by-case authoring.

Test data—Query building and file generation

AI-assisted data and query generation automates the creation of complex, scenario-specific test data—transaction files, simulated claims scenarios, database validation and migration queries—shortening the data-readiness bottleneck that often delays execution. Test data is masked or synthetically generated so protected health information is not exposed to AI processing.

Test execution, defect logging and reporting

AI-enabled execution and defect intelligence automate system and interface validation, regression execution, real-time transaction verification, file comparison, extract validation and end-to-end workflow orchestration—expanding coverage, reducing manual effort and accelerating confidence in production readiness.

Client outcomes

30%–35% quality engineering effort optimization potential

Broader test coverage, including migration validation

Faster defect identification and accelerated regression cycles

Higher confidence in production readiness and business continuity at go-live

Conclusion

Across platform foundation, configuration, system integration, data migration and quality engineering, AI transforms TriZetto core claims implementations from a linear, effort-intensive process into a faster, more predictable and more intelligent delivery model. The AI-infused factory approach compounds gains across the lifecycle by reducing reliance on scarce SMEs, industrializing repetitive build and validation work, and preserving human judgment at critical decision points.

The result is a **30%–35% effort optimization opportunity** across the greenfield implementation lifecycle and a **three- to four-month reduction in elapsed time**. These gains can accelerate time to market, improve delivery predictability and lower total cost of ownership, with savings reflected in the commercial construct at proposal stage. The exhibit below summarizes the five towers, the role of AI in each and the outcomes they deliver.



AI infusion across the implementation lifecycle

Five AI-powered towers—each delivering measurable business and IT outcomes

AI-powered

Platform foundation	Configuration design and build	System integration and development	Data migration factory	Quality engineering
---------------------	--------------------------------	------------------------------------	------------------------	---------------------

Optimization potential

8–10%	25–30%	30–35%	40–45%	30–35%
-------	--------	--------	--------	--------

AI infusion

laC and DevOps-agnostic automation accelerate cloud readiness, with an AI decision and assurance layer that recommends topology and sizing, generates deployment runbooks, and detects environment drift early.	TriZetto Configuration Orchestrator (TCO) turns benefit product data into validated, Facets-ready configuration. Gen AI validation reconciles contract PDFs to compensation grids, and pre-built accelerators reuse proven configuration patterns.	SI AI factory industrializes specs, code and tests with Facets SME depth and governed AI. Multi-agent delivery turns one BRD into mappings TSDs, and C#/SQL, MCP-grounded, ~40% dividend, with human review at every stage boundary.	AI profiling and source-to-target mapping cut discovery effort, and mapping rules are converted into executable SQL—with strategy, cutover and load validation kept expert-led.	AI test design converts requirements and prior assets into test cases, generates scenario-specific data and validation queries and speeds regression, extract comparison and defect intelligence.
---	--	--	--	--

Key outcomes

• Faster mobilization • Accelerated cloud readiness • Deployment consistency • Better governance	• Faster benefit build cycles • Fewer configuration defects • Accelerated renewals • Compliance readiness	• Faster solution delivery • Higher development quality • Engineering productivity • Improved market agility	• Faster conversion cycles • Improved migration quality • Go-live confidence • Lower disruption risk	• Broader test coverage • Faster defect identification • Accelerated regression • Production readiness
---	--	---	---	---

Compounding gains across the full lifecycle—faster time-to-market, predictable delivery, lower cost of ownership

Governed AI inside a Cognizant-managed factory—protecting TriZetto and Facets IP and delivering AI productivity even where client policy restricts external gen AI tools.

Figure 4. AI infusion across the implementation lifecycle—five AI-powered towers, each delivering measurable business and IT outcomes.

Trademarks and legal notices

TriZetto, Facets and the TriZetto Configuration Orchestrator are trademarks or registered trademarks of Cognizant Technology Solutions Corporation or its subsidiaries in the United States and other countries. Cognizant and the Cognizant logo are trademarks or registered trademarks of Cognizant Technology Solutions Corporation. All other product names, service marks and trademarks referenced in this document—including Microsoft, SQL Server and .NET (Microsoft Corporation) and Oracle (Oracle Corporation)—are the property of their respective owners and are used for identification purposes only. Model Context Protocol is an open specification and is referenced descriptively.

Forward-looking and performance statements

The optimization ranges, timeline reductions and outcomes described in this document reflect results observed on specific Cognizant engagements and internal measurement of comparable delivery baselines. They are provided for illustrative purposes and are not a guarantee, warranty or contractual commitment of future performance. Actual results depend on program scope, platform version and configuration complexity, data quality and legacy estate condition, regulatory context, client-side readiness and resource availability, and the extent of AI-enabled accelerator adoption agreed for the engagement. Statements regarding planned or future capability are forward-looking and subject to change without notice. Nothing in this document constitutes legal, regulatory or compliance advice.

© 2026 Cognizant. All rights reserved.



About the authors

Author



Avinash Kumar Loganathan
Go-to-Market India Executive
and Strategic Consultant,
TriZetto Products and Services

Co-author



Sai Sriram Janardhanan
Business Development India
Lead, TriZetto Products
and Services

Reviewers



Gopalakrishnan Iyer
Head of Healthcare
Product Consulting



Vadivelu Nagasundaram
Head of TriZetto
Product Engineering



Cognizant (Nasdaq: CTSI) is an AI Builder and technology services provider, bridging the gap between AI investment and enterprise value by building full-stack AI solutions for our clients. Our deep industry, process and engineering expertise enables us to build an organization's unique context into technology systems that amplify human potential, drive tangible outcomes and keep global enterprises ahead in a fast-changing world. See how at cognizant.ai or @cognizant.

World Headquarters

300 Frank W. Burr Blvd.
Suite 36, 6th Floor
Teaneck, NJ 07666, USA
Tel: (201) 801-0233

European Headquarters

280 Bishopsgate
London
EC2M 4AG
England
Tel: +44 (0) 20 7297 7600

India Corporate Office

Siruseri-Software Technology Park of India (STPI)
SDB Block – Ground Floor North Wing
Plot No H4, SIPCOT IT Park
Chengalpattu District
Chennai 603103, Tamil Nadu
Tel: 1800 208 6999

APAC Headquarters

1 Fusionopolis Link,
Level 5 NEXUS@One-North,
North Tower, Singapore 138542
Phone: + 65 6812 4000

© 2025–2027, Cognizant. All rights reserved. No part of this document may be reproduced, stored in a retrieval system, transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise without the express written permission of Cognizant. The information contained herein is subject to change without notice. All other trademarks mentioned herein are the property of their respective owners.

WF 5325300