


OFFICIAL ACTUAL · 2024

U.S. healthcare spending reached \$5.3T in 2024.

\$5.3T

US healthcare spending · 2024 actual



● Green: administrative expenses · 19% of total spending

U.S. healthcare spending reached \$5.3T in 2024.

U.S. healthcare spending reached \$5.3 trillion in 2024. This report examines administrative costs and the use of AI in revenue cycle management.

Green dots represent the report’s administrative-spending scenario, with an illustrative distribution across service categories. The report covers 2024 spending data, 2026 projections, RCM workflows, workforce distribution, and AI adoption. It also examines automation costs, performance, governance, and customer results.

INTERACTIVE REPORT ↓ [Scroll from actuals to projections](#)

[CMS · 2024 actual expenditures ↗](#)

CHART CONTEXT


The opening is a CMS actual; the next view is a CMS projection. Administrative and workflow amounts are GenHealth scenarios, not independently measured 2026 national allocations. Dots approximate shares; cloud area scales with total spending. See the spending methodology below for normalization and payer/provider assumptions.

OFFICIAL PROJECTION · 2026

CMS projects U.S. healthcare spending will reach \$6.0T in 2026.

\$6.02T

2026 CMS projection · +14.0% versus 2024



● Green: administrative expenses · 19% of total spending

CMS projects U.S. healthcare spending will reach \$6.0T in 2026.

CMS projects \$6.02T in US healthcare spending in 2026, up 14.0% from 2024.

The June 2026 release puts spending at \$6,017.4 billion, compared with \$5,278.6 billion in 2024. These are nominal dollars, including both price and utilization changes.

From this point forward, the report uses the 2026 spending outlook.

Service spending comes directly from CMS. Administrative and workflow amounts are explicitly labeled GenHealth scenarios built on that outlook.

[CMS · 2025–2034 projection tables, June 2026 ↗](#)

CHART CONTEXT


The opening is a CMS actual; the next view is a CMS projection. Administrative and workflow amounts are GenHealth scenarios, not independently measured 2026 national allocations. Dots approximate shares; cloud area scales with total spending. See the spending methodology below for normalization and payer/provider assumptions.

OFFICIAL PROJECTION · 2026

Healthcare spending spans hospitals, physicians, medicines, and other services.

\$6.02T


2026 projected spending · by service




Hospital care · 31.1%



Physician / clinical · 20.7%



Retail prescriptions · 9.3%



All other spending · 38.9%

● Green: administrative expenses · 19% of total spending

Healthcare spending spans hospitals, physicians, medicines, and other services.

In 2026, CMS projects \$1870.5B for hospital care, \$1244.5B for physician and clinical services, and \$561.5B for retail prescriptions.

These represent 31.1%, 20.7%, and 9.3% of projected spending. The remaining 38.9% covers other services and products, public health, investment, and insurance-related costs.

Procedures and medications administered in a hospital sit inside the relevant service category. They are not separate, additive buckets.

[CMS · Table 2, 2026 projected expenditures ↗](#)

CHART CONTEXT

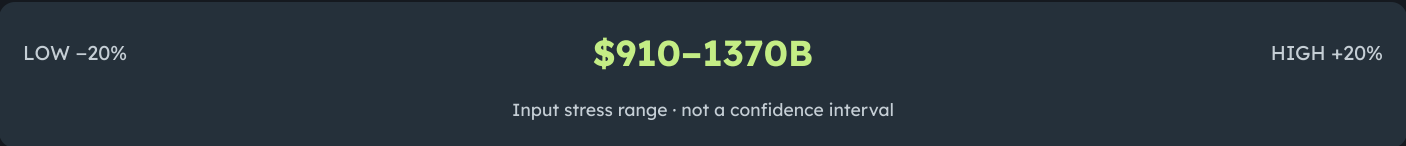
The opening is a CMS actual; the next view is a CMS projection. Administrative and workflow amounts are GenHealth scenarios, not independently measured 2026 national allocations. Dots approximate shares; cloud area scales with total spending. See the spending methodology below for normalization and payer/provider assumptions.

GENHEALTH MODEL · 2026

Our model puts 2026 healthcare administration at \$1.1T.

\$1.1T

2026 administrative spending · GenHealth scenario



● Green: administrative expenses · 19% of total spending

Our model puts 2026 healthcare administration at \$1.1T.

Every visit also creates scheduling, eligibility, documentation, authorization, claims, and payment work.

Our 2026 scenario is approximately**\$1140B in administrative spending**. It holds the administrative share constant while national spending grows14.0% from 2024 to 2026.

For this scenario, we normalize the roughly \$1T estimate discussed by Penn LDI in 2025 to a 2024 reference year. This is a modeling assumption, not a CMS administrative forecast or a claim that Penn dated that estimate to 2024.

Administration is embedded in healthcare spending. Adding it on top would count the same dollars twice.

[2026 scenario assumptions and sources ↗](#)

CHART CONTEXT


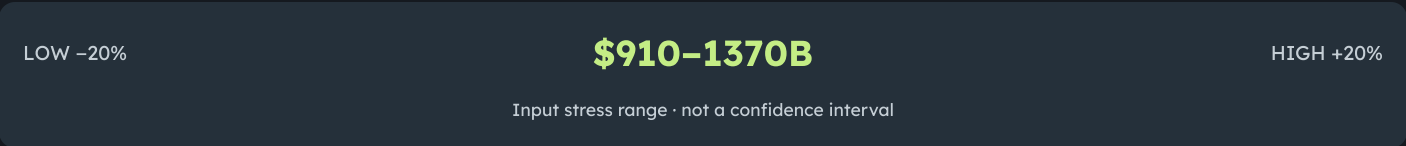
The opening is a CMS actual; the next view is a CMS projection. Administrative and workflow amounts are GenHealth scenarios, not independently measured 2026 national allocations. Dots approximate shares; cloud area scales with total spending. See the spending methodology below for normalization and payer/provider assumptions.

GENHEALTH MODEL · HISTORICAL RATIO


Providers account for most administrative spending in our model.

\$1.1T

2026 admin scenario · historical payer/provider mix



PAYER · \$312B



PROVIDER · \$828B

Providers account for most administrative spending in our model.

Payers verify coverage, review requests, and adjudicate claims. Providers prepare evidence, deliver care, submit claims, and pursue payment.

Applying McKinsey’s historical payer/provider mix to our 2026 administrative scenario gives approximately**\$312B for payers** and**\$828B for providers**—27% and 73%.

The ratio comes from McKinsey’s 2019 baseline of \$260B payer and \$690B provider-side administration. Only the ratio carries forward; these displayed dollar amounts are 2026 scenario estimates, not newly measured totals.

[McKinsey · historical allocation underlying the scenario ↗](#)

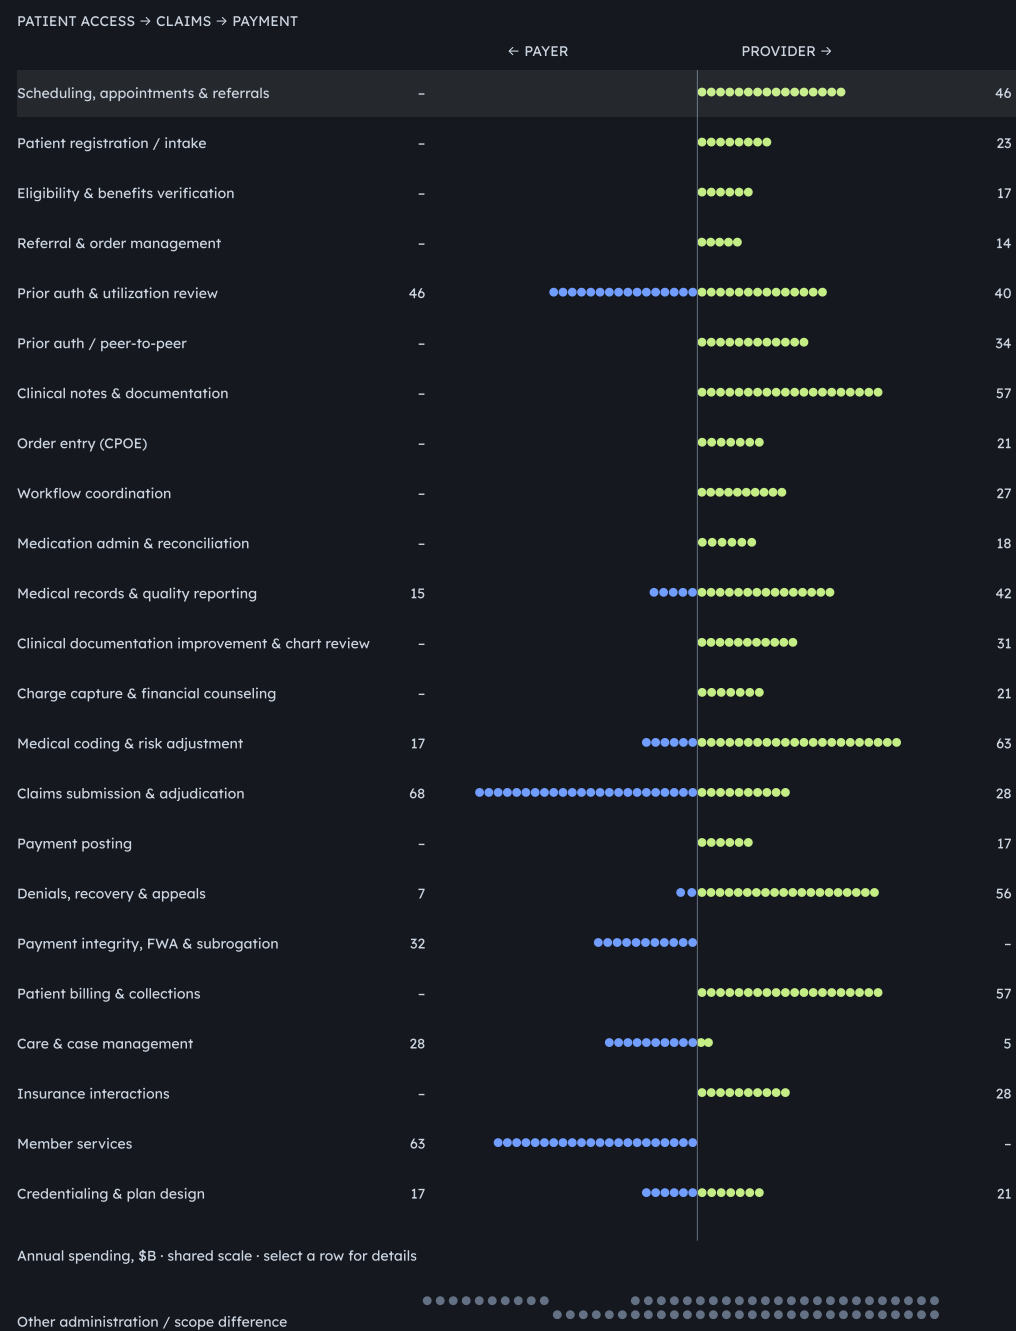
CHART CONTEXT

The opening is a CMS actual; the next view is a CMS projection. Administrative and workflow amounts are GenHealth scenarios, not independently measured 2026 national allocations. Dots approximate shares; cloud area scales with total spending. See the spending methodology below for normalization and payer/provider assumptions.

Administrative work extends from patient access through payment.

\$1.1T

Administrative work from patient access to payment



Ongoing support functions appear last; they span the cycle.

Scheduling, appointments & referrals

Turn a referral or appointment request into a scheduled visit with the right clinician, service, and location.

Administrative work extends from patient access through payment.

One view of the administrative work surrounding an encounter: scheduling and coverage checks, authorization, documentation and coding, claims, payment posting, denials, and collections.

Read from top to bottom. Payer work stays on the left and provider work on the right. Select any row to read what the work involves. Rows use the same spending scale.

The \$1.1T total includes administration beyond these workflows. The detailed rows come from the narrower GenHealth workflow model; gray dots retain the remaining scope difference. These are illustrative allocations, not measured national workflow costs.

Member services, care management, and credentialing support many encounters rather than taking place only after payment. They appear together at the bottom.

[Workflow model and spending methodology ↗](#)

CHART CONTEXT

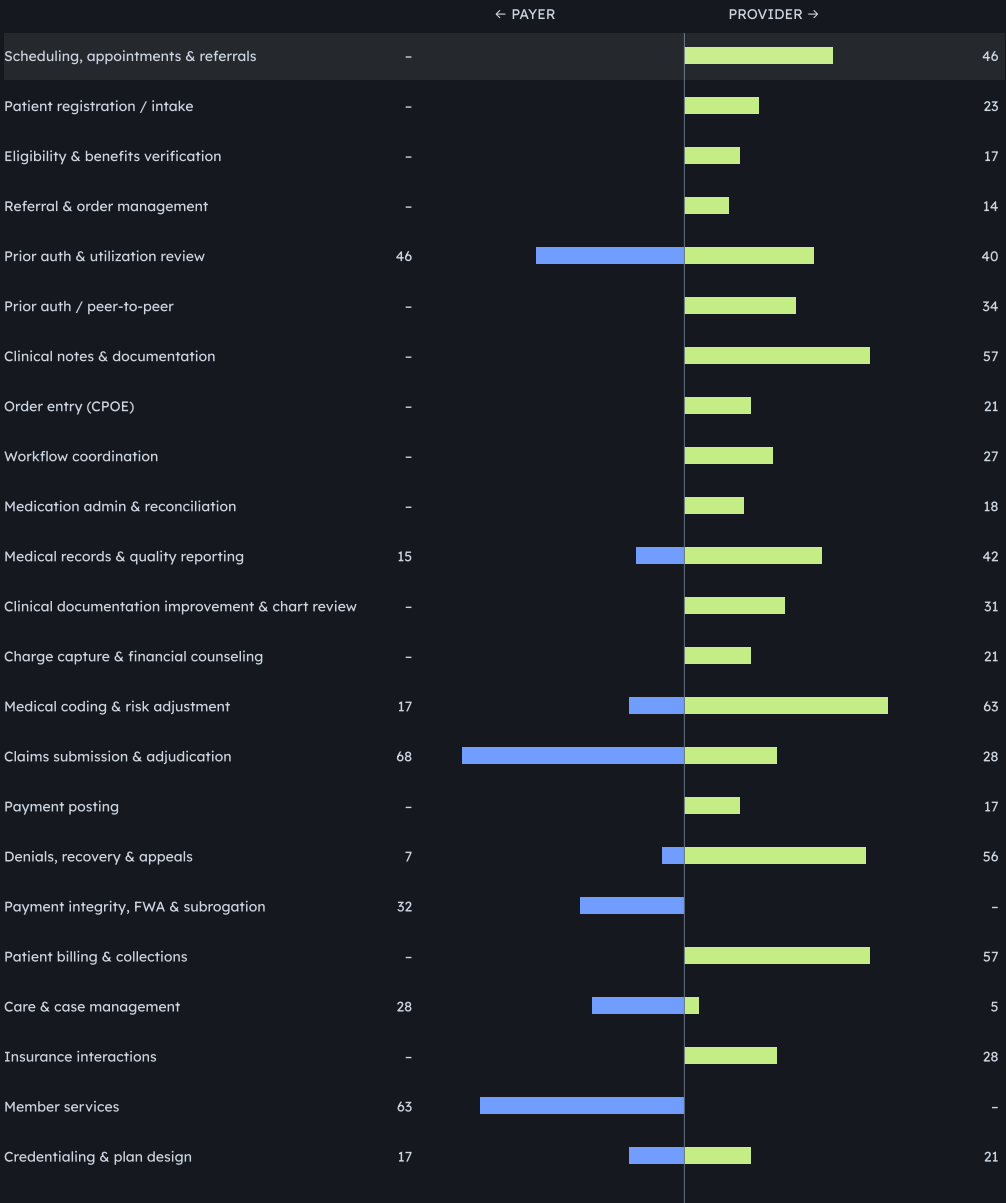
The opening is a CMS actual; the next view is a CMS projection. Administrative and workflow amounts are GenHealth scenarios, not independently measured 2026 national allocations. Dots approximate shares; cloud area scales with total spending. See the spending methodology below for normalization and payer/provider assumptions.

Payers and providers share the cost of administrative workflows.

\$1.1T

Administrative work from patient access to payment

PATIENT ACCESS → CLAIMS → PAYMENT



Annual spending, \$B · shared scale · select a row for details

Other administration / scope difference

Ongoing support functions appear last; they span the cycle.

Scheduling, appointments & referrals

Turn a referral or appointment request into a scheduled visit with the right clinician, service, and location.

Payers and providers share the cost of administrative workflows.

The bars compare annual spending on each workflow. Their lengths use the same scale on both sides, so payer and provider effort can be compared directly.

Provider-heavy work includes scheduling, documentation, coding, and patient billing. Payer-heavy work includes claims adjudication, payment integrity, and member services. Authorization and appeals require work on both sides.

Select a row to read its definition. The gray dots below the chart continue to represent administration outside the detailed workflow model.

[Allocation assumptions and sources ↗](#)

CHART CONTEXT

The opening is a CMS actual; the next view is a CMS projection. Administrative and workflow amounts are GenHealth scenarios, not independently measured 2026 national allocations. Dots approximate shares; cloud area scales with total spending. See the spending methodology below for normalization and payer/provider assumptions.

Each workflow combines software, outsourced labor, and employee costs.

\$1.1T

Administrative work from patient access to payment



Annual spending, \$B · shared scale · select a row for details

Other administration / scope difference

Ongoing support functions appear last; they span the cycle.

Scheduling, appointments & referrals

Turn a referral or appointment request into a scheduled visit with the right clinician, service, and location.

In-house labor:75% · \$34B

Outsourced labor:10% · \$5B

Software / IT:15% · \$7B

Each workflow combines software, outsourced labor, and employee costs.

Each bar now separates in-house labor, outsourced labor, and software / IT. The total length stays the same. Lighter shades represent employees, middle shades represent outsourced labor, and darker shades represent software and IT.

These allocations use GenHealth’s Series A workflow model. It assigns a different mix to each workflow and applies that mix to both payer and provider spending. They are planning assumptions, not independently measured national spending shares.

The model assigns more outsourcing to coding and patient collections, more software to payment posting and eligibility checks, and more in-house labor to clinical documentation and coordination. Select a row to inspect its mix and dollar amounts.

Rounded source shares are normalized to 100% so segments preserve every bar’s total. The gray remainder is outside the detailed workflow model and has no assigned mix.

[Model methodology ↗](#)

CHART CONTEXT

The opening is a CMS actual; the next view is a CMS projection. Administrative and workflow amounts are GenHealth scenarios, not independently measured 2026 national allocations. Dots approximate shares; cloud area scales with total spending. See the spending methodology below for normalization and payer/provider assumptions.

Claims and collections are the most familiar parts of RCM.

\$1.1T

Claims, payment, denials, and collections

Payer RCM
\$107B

Provider RCM
\$158B

Combined RCM
\$266B



Annual spending, \$B - shared scale - select a row for details

Other administration / scope difference

Ongoing support functions appear last; they span the cycle.

Scheduling, appointments & referrals

Turn a referral or appointment request into a scheduled visit with the right clinician, service, and location.

Claims and collections are the most familiar parts of RCM.

When people talk about RCM, they often mean the work from claim submission onward: adjudication, payment posting, denial follow-up, appeals, patient billing, collections, and payment integrity. Those rows are highlighted first.

This part of the cycle pursues and reconciles payment for care already delivered. In the report’s model, the highlighted rows total \$107B for payers and \$158B for providers.

But the claim depends on work completed much earlier. Coverage errors, missing authorization, and incomplete documentation can become payment problems downstream.

CHART CONTEXT

The opening is a CMS actual; the next view is a CMS projection. Administrative and workflow amounts are GenHealth scenarios, not independently measured 2026 national allocations. Dots approximate shares; cloud area scales with total spending. See the spending methodology below for normalization and payer/provider assumptions.

RCM runs from patient access through final payment.

\$1.1T

The full revenue cycle starts before the claim



Annual spending, \$B - shared scale - select a row for details

Other administration / scope difference

Ongoing support functions appear last; they span the cycle.

Scheduling, appointments & referrals

Turn a referral or appointment request into a scheduled visit with the right clinician, service, and location.

RCM runs from patient access through final payment.

RCM starts before a claim exists. Scheduling and registration establish the account; eligibility and authorization establish coverage requirements; documentation and coding support the claim. The highlight now extends to this earlier work as well as payment and recovery. RCM requires the full process from start to finish.

The highlighted rows cover patient access, eligibility, authorization, documentation improvement, coding, insurer communications, claims, payment posting, denials, patient collections, and payment integrity. The totals include software, outsourcing, and in-house labor on each side.

Within this report’s 2026 allocation model, these rows total \$170B for payers and \$474B for providers, or \$644B combined. They are a subset of the \$1.1T administrative scenario, not an additional expense.

Dimmed rows cover broader clinical documentation, order entry, medication administration, workflow coordination, records and quality reporting, care management, member services, and credentialing / plan design. Those functions can support RCM, but their full costs are outside this report’s selected RCM scope.

This is an editorial scope applied to the GenHealth model, not a measured national RCM total. Combined rows such as coding / risk adjustment and referrals / order management are counted whole because the source does not separate their component costs. Payer totals describe corresponding financial-administration work.

[Definitions and model assumptions ↗](#)

CHART CONTEXT

The opening is a CMS actual; the next view is a CMS projection. Administrative and workflow amounts are GenHealth scenarios, not independently measured 2026 national allocations. Dots approximate shares; cloud area scales with total spending. See the spending methodology below for normalization and payer/provider assumptions.

CAQH estimates a further \$21B opportunity in transaction automation.

2025 CAQH INDEX / RELEASED FEBRUARY 2026

\$258B

Estimated 2024 cost avoidance from electronic transactions and data exchange

\$21B

Estimated remaining automation opportunity in tracked transactions

Not AI ROI

These figures do not measure savings caused by generative AI

CAQH estimates a further \$21B opportunity in transaction automation.

CAQH’s 2025 Index, released February 2026, estimates \$258B of administrative cost avoidance in 2024 from electronic transactions and data exchange, with a further \$21B opportunity from automating manual and partially manual transactions. Its release cites more than 600 provider organizations and plans representing 63% of insured lives.

These are research estimates grounded in transaction activity, not audited realized cash savings or an experiment isolating AI. They cover a narrower transaction scope than our administrative model. Neither amount should be added to that model or labeled AI revenue.

The same release reports administrative AI use at more than 50% of plans and 25% of provider organizations. These figures are distinct from HFMA’s deployment stages and NAIC’s enterprise AI/ML definition; they do not establish autonomous throughput.

[CAQH · 2025 Index release, February 19, 2026 ↗](#)

CMS still identifies 2024 as the latest historical spending year. We retain the 2026 projection rather than relabeling an estimate as an actual. Older specialty benchmarks remain dated because this release does not replace those specialty-level observations.

[CMS · Latest historical national expenditure accounts ↗](#)

CHART CONTEXT

2025 Index refers to the report edition, released in February 2026 and discussing 2024 activity. Cost avoidance is estimated rather than realized cash savings; administrative AI use does not mean autonomous completion.

HISTORICAL STUDY · SCOPE-SPECIFIC

Eligibility checks lead provider transaction spending in CAQH’s analysis.

PROVIDER ECONOMICS / HISTORICAL EVIDENCE



[Source and methodology ↗](#)

Volume × effort × cost

A costly individual case is not necessarily the largest spending pool. Aggregate specialty RCM spend requires volumes and fully loaded costs measured on the same basis.

Eligibility checks lead provider transaction spending in CAQH’s analysis.

CAQH’s 2023 Index specialty analysis identifies eligibility, claim submission, and claim-status inquiries as the three largest medical transaction spending categories. Applying its reported provider shares gives approximately **\$41.2B**, **\$18.4B**, and **\$12.1B**, respectively.

[CAQH · specialty analysis, pages 2, 4, 6 ↗](#)

This is a narrower, historical transaction-cost lens—not an allocation of the report’s [2026 core workflow scenario](#). The numbers must not be added together across those two models.

Total cost depends on transaction volume as well as the cost of each transaction.

CHART CONTEXT


CAQH figures describe 2022 activity in its 2023 Index. Provider totals are derived from rounded industry totals and reported provider shares. Manual transaction costs exclude preparation, follow-up, and systems. JAMA encounter estimates come from one academic health system and include professional and hospital billing; neither dataset is a 2026 national specialty forecast.

HISTORICAL STUDY · SCOPE-SPECIFIC

Specialists face higher costs for several manual billing tasks.

PROVIDER ECONOMICS / HISTORICAL EVIDENCE

- Eligibility
- Prior authorization
- Claim status



[Source and methodology ↗](#)

Volume × effort × cost

A costly individual case is not necessarily the largest spending pool. Aggregate specialty RCM spend requires volumes and fully loaded costs measured on the same basis.

Specialists face higher costs for several manual billing tasks.

For manual eligibility, authorization, and claim-status work, CAQH reports higher unit costs for specialists and behavioral-health providers than generalists. Use the selector to compare tasks.

[CAQH · specialty cost table, page 10 ↗](#)

Operationally, a provider-focused product needs more than one generic queue. Useful segments include primary care, procedural specialties, behavioral health, and hospital-based services. Specific workflows should be validated with each practice rather than inferred from its specialty name.

- **Primary care:** repeated coverage checks, routine billing, and patient balances.
- **Procedural specialties:** authorization evidence, procedure documentation, and claim edits.
- **Behavioral health:** coverage details, authorization requirements, and recurring-visit billing.
- **Hospital services:** documentation, coding, and coordination of facility and professional billing.

These are workflow-design examples, not quantified specialty market shares. Named fields such as oncology, cardiology, orthopedics, radiology, and DME need their own volume and cost data before assigning national RCM dollars.


CHART CONTEXT

CAQH figures describe 2022 activity in its 2023 Index. Provider totals are derived from rounded industry totals and reported provider shares. Manual transaction costs exclude preparation, follow-up, and systems. JAMA encounter estimates come from one academic health system and include professional and hospital billing; neither dataset is a 2026 national specialty forecast.

HISTORICAL STUDY · SCOPE-SPECIFIC

Complex encounters cost more to bill in one health-system study.

PROVIDER ECONOMICS / HISTORICAL EVIDENCE



[Source and methodology ↗](#)

Volume × effort × cost

A costly individual case is not necessarily the largest spending pool. Aggregate specialty RCM spend requires volumes and fully loaded costs measured on the same basis.

Complex encounters cost more to bill in one health-system study.

A 2018 JAMA study at one academic health system estimated billing and insurance costs from **\$20.49 per primary-care visit** to **\$215.10 per inpatient surgical encounter**. The analysis included professional and hospital billing. The graphic compares all five encounter types.

[Tseng et al. · JAMA 2018, Table 1 ↗](#)

This supports a complexity gradient, not a current national ranking. High case cost, high claim volume, and a high percentage of revenue consumed by billing are three different measures.

CMS publishes Medicare Part B expenditure and service counts by specialty. Those are payments for care, not RCM expense. They can supply a scoped volume lens, but multiplying them by an unrelated administrative percentage would produce an unsupported estimate.

[CMS · expenditures and services by specialty ↗](#)

Our takeaway: target repeatable volume in routine workflows and expensive exceptions in complex care. We have evidence for those two opportunities; we do not yet have a defensible all-payer national RCM spending leaderboard by individual specialty.

CHART CONTEXT

CAQH figures describe 2022 activity in its 2023 Index. Provider totals are derived from rounded industry totals and reported provider shares. Manual transaction costs exclude preparation, follow-up, and systems. JAMA encounter estimates come from one academic health system and include professional and hospital billing; neither dataset is a 2026 national specialty forecast.

U.S. RCM operations rely on domestic and overseas teams.

GLOBAL DELIVERY / DOCUMENTED LOCATIONS

All RCM spend

Outsourced labor only

\$644B

2026 RCM model · all spending types


City spend estimates

Documented locations

Spending scenario

Company footprint

All organizations · 33 hubs / 6 countries



Spikes count documented organizations, not city spending. Totals above come from the national RCM model; no geographic allocation is available.

- Dallas
- Hackensack
- Atlanta
- Cranbury
- El Paso
- Chennai
- Coimbatore
- Salem
- Pune
- Noida
- Navi Mumbai
- Bengaluru
- Thiruvananthapuram
- Hyderabad
- Indore
- Metro Manila
- Iloilo
- Tagbilaran
- Kingston
- St. Catherine
- Barranquilla
- Murray / Salt Lake City
- Chicago
- Blue Ash / Cincinnati
- Boca Raton
- Nashville
- Rochester
- Cleveland
- Pittsburgh
- Los Angeles
- New York City
- Buenos Aires
- Bogotá

Dallas · United States

[Access Healthcare · organization location / hiring evidence](#)

[Conifer · organization location / hiring evidence](#)

City headcount and payroll are not disclosed. Provider job locations are separate from vendor offices; neither measures local RCM capacity.

U.S. RCM operations rely on domestic and overseas teams.

The globe starts over the United States, where the healthcare spending modeled in this report originates. The CMS 2026 projection is \$6.02T; the report’s selected RCM workflow scenario is \$644B. Those totals are national, not city-level estimates.

The initial view includes provider-employed teams and service providers. U.S. examples include Mayo Clinic in Rochester, Cleveland Clinic, UPMC in Pittsburgh, UCLA Health in Los Angeles, NYC Health + Hospitals, and HCA / Parallon in Nashville. Job postings establish team locations; some roles are hybrid or remote.

[Mayo Clinic · revenue-cycle hiring](#)

[Cleveland Clinic · patient access](#)

[UPMC · corporate revenue cycle](#)

[UCLA Health · revenue capture](#)

[NYC Health + Hospitals · patient accounts](#)

[HCA / Parallon · RCM shared services](#)

RCM can be performed by a provider’s own staff, a domestic service partner, or an offshore team. Outsourced does not necessarily mean abroad; a U.S. vendor can employ people in several countries.

U.S. RCM operations rely on domestic and overseas teams.

GLOBAL DELIVERY / DOCUMENTED LOCATIONS


All RCM spend Outsourced labor only

\$644B

2026 RCM model · all spending types

City spend estimates Documented locations Spending scenario

Company footprint All organizations · 33 hubs / 6 countries



Spikes count documented organizations, not city spending. Totals above come from the national RCM model; no geographic allocation is available.

- Dallas Hackensack Atlanta Cranbury El Paso Chennai Coimbatore Salem Pune Noida
- Navi Mumbai Bengaluru Thiruvananthapuram Hyderabad Indore Metro Manila Iloilo Tagbilaran
- Kingston St. Catherine Barranquilla Murray / Salt Lake City Chicago Blue Ash / Cincinnati Boca Raton
- Nashville Rochester Cleveland Pittsburgh Los Angeles New York City Buenos Aires Bogotá

Dallas · United States

[Access Healthcare · organization location / hiring evidence ↗](#)

[Conifer · organization location / hiring evidence ↗](#)

City headcount and payroll are not disclosed. Provider job locations are separate from vendor offices; neither measures local RCM capacity.

Vendor disclosures show that RCM teams span multiple countries.

A disclosed example: roughly 41% U.S. / 59% international. R1’s 2023 filing reports 12,300 U.S. employees (including 500 part-time) and 17,600 international employees. This is company headcount, not FTEs or a census of U.S. RCM.

[R1 · 2023 Form 10-K, Human Capital Management ↗](#)

A newer, different business mix: Sagility’s FY2024–25 workforce. Its sustainability report lists 38,754 employees across five countries. The country figures imply 6.3% in the U.S. and 93.7% abroad; India and the Philippines together account for 85.0%. Sagility serves both health plans and providers, so this is broader healthcare operations employment, not provider RCM alone.

Country	Employees	Share
India	17,885	46.2%
Philippines	15,062	38.9%
Jamaica	3,123	8.1%
United States	2,430	6.3%
Colombia	254	0.7%

[Sagility · FY2024–25 sustainability report, p. 17 ↗](#)

These figures use the sustainability report’s reporting perimeter. Sagility’s FY25 earnings release separately reports 39,409 year-end employees; we do not combine that denominator with this country breakdown. The contrast with R1 demonstrates why one vendor’s staffing mix cannot stand in for the whole market.

[Sagility · FY25 earnings release, May 14, 2025 ↗](#)

We did not identify a public, comprehensive count covering in-house provider teams and every RCM service firm. The industry-wide domestic/offshore share remains unquantified here. Historical workforce evidence is dated explicitly, alongside the report’s 2026 spending scenario.

CHART CONTEXT

Location evidence combines company directories, public company profiles and provider job postings across 33 hubs and six countries. Spikes count documented organizations, not staff or local spending. The city-spending view uses a separate supplied allocation of \$68.6B across 100 cities. Its spike heights are proportional to estimated dollars, unlike the documented-locations view. City values are unverified and have no specified reference year; the cited BPO market reports do not establish those allocations. The separate 1,000-person staffing scenario retains its explicitly illustrative inputs.

The supplied city estimates allocate \$68.6B to outsourced administration.


All RCM spend Outsourced administration

\$68.60B

Annual city allocation · 100 supplied estimates

Region All regions

City spend estimates Documented locations Spending scenario



Spike height is linear in annual outsourced-administration spend. City figures are supplied estimates, not independently verified municipal data.

- BengaluruChennaiMumbaiHyderabadNoidaGurugramPuneCoimbatoreKochi
- TrivandrumAhmedabadKolkataJaipurChandigarhIndoreBhubaneswarVisakhapatnam
- NagpurMaduraiMysuruTrichySuratVadodaraLucknowPatnaTaguig (Metro Manila)
- Quezon City (Metro Manila)Makati (Metro Manila)Cebu CityDavao CityIloilo CityClark / Angeles City
- BacolodCagayan de OroDumagueteBaguioSanta Rosa (Laguna)General Santos CityTacloban
- ZamboangaChicago, ILNashville, TNDallas, TXMinneapolis, MNBoston, MAAtlanta, GA
- New York, NYLos Angeles, CAHouston, TXPhoenix, AZDenver, COCharlotte, NCOrlando, FL
- Tampa, FLSalt Lake City, UTPhiladelphia, PASan Francisco, CASSeattle, WAAustin, TXSt. Louis, MO
- Miami, FLCleveland, OHPittsburgh, PAIndianapolis, INLouisville, KYSan José, Costa Rica
- Santo Domingo, Dominican RepublicGuadalajara, MexicoMonterrey, MexicoBogota, Colombia
- Medellin, ColombiaKingston, JamaicaMontego Bay, JamaicaGuatemala City, Guatemala
- San Salvador, El SalvadorTegucigalpa, HondurasSan Pedro Sula, HondurasManagua, Nicaragua
- Panama City, PanamaLima, PeruSantiago, ChileBuenos Aires, ArgentinaSao Paulo, Brazil
- Quito, EcuadorPort of Spain, Trinidad & TobagoKrakow, PolandWarsaw, PolandBucharest, Romania
- Sofia, BulgariaBudapest, HungaryPrague, Czech RepublicCairo, EgyptCape Town, South Africa
- Johannesburg, South AfricaNairobi, KenyaCasablanca, MoroccoTunis, TunisiaDubai, UAE
- Amman, JordanBelgrade, Serbia

Chicago, IL · United States · \$3.50B / year
5.10% of the supplied city allocation.

[Download all 100 city estimates ↗](#)

The supplied city estimates allocate \$68.6B to outsourced administration.

The spike map uses the supplied city allocation: \$68.6B in annual outsourced healthcare administration across 100 cities. Each spike is proportional to its city’s dollar value; select a city to see its amount and share. The map starts over the United States.

Chicago is assigned \$3.5B, Taguig \$3.2B, Nashville \$3.1B, Quezon City \$2.9B, and Bengaluru \$2.6B. Buenos Aires is assigned \$70M. These are the supplied estimates, not figures derived from our location directories.

Region	City total
India	\$19.69B
Philippines	\$16.40B
United States	\$22.10B
Latin America & Caribbean	\$7.60B
Europe, Africa & Middle East	\$2.81B

Regional totals are calculated directly from the city rows. They differ from the regional summary in the supplied text, although the grand total is the same \$68.6B. Each original city value is preserved.

The supplied scope covers outsourced RCM, claims, coding, and clinical voice work. It does not cover all in-house administration and is separate from the report’s \$133B outsourced-labor workflow scenario. Amounts represent estimated activity routed to or managed through cities, not verified local payroll or municipal accounts.

The linked market reports describe broader healthcare BPO markets; their public pages do not substantiate this city allocation. The city figures remain a supplied illustrative dataset with no documented city-level estimation methodology or reference year.

[MarketsandMarkets · broader healthcare BPO market context ↗](#)

[Fortune Business Insights · broader healthcare BPO market context ↗](#)

[Download the 100 city values](#)

CHART CONTEXT

Location evidence combines company directories, public company profiles and provider job postings across 33 hubs and six countries. Spikes count documented organizations, not staff or local spending. The city-spending view uses a separate supplied allocation of \$68.6B across 100 cities. Its spike heights are proportional to estimated dollars, unlike the documented-locations view. City values are unverified and have no specified reference year; the cited BPO market reports do not establish those allocations. The separate 1,000-person staffing scenario retains its explicitly illustrative inputs.

GENHEALTH SCENARIO · ILLUSTRATIVE SPENDING

Staffing location changes the cost of an RCM team.

GLOBAL DELIVERY / EDITABLE SCENARIO

- City spend estimates
- Documented locations
- Spending scenario



Spike height shows illustrative annual spending. These are editable assumptions.

- Dallas
- Hackensack
- Boca Raton
- Chennai
- Noida
- Pune
- Hyderabad
- Coimbatore
- Metro Manila
- Bogotá

Chennai · India · \$3.0M scenario cost (40% of its country bucket).[Location source ↗](#)

▼ Model a 1,000-person RCM team

U.S. staffing share:50%



Country	People	Cost/person	Annual spend
United States	500	80000	\$40.0M 79.0%
India	375	20000	\$7.5M 14.8%
Philippines	100	24000	\$2.4M 4.7%
Colombia	25	30000	\$0.8M 1.5%

Total annual staffing cost:\$50.6M

Staffing location changes the cost of an RCM team.

Switch to the globe’s Spending scenario layer to explore a 1,000-person team. Its illustrative starting case puts 500 people in the U.S., 375 in India, 100 in the Philippines, and 25 in Colombia.

At assumed annual fully loaded costs of \$80,000 / \$20,000 / \$24,000 / \$30,000, respectively, spending is**\$40M / \$7.5M / \$2.4M / \$0.75M**. Half the people are abroad, but just 21% of the \$50.65M staffing budget is abroad. Change the inputs to see the difference.

These are planning assumptions, not researched wage estimates or a forecast of national RCM spending. City weights are illustrative too; the tallest spikes show the largest allocations in this scenario, not a measured ranking.

Back in the broader[\\$958B core workflow model](#), the assumed 65% employee / 18% outsourced / 17% technology mix corresponds to\$622B /\$172B /\$163B. That model includes work beyond narrow RCM. Neither labor bucket has a verified geographic split, so we do not multiply these national totals by the example team’s shares.

Use the current team’s costs and performance as the baseline for an AI evaluation.

CHART CONTEXT

Location evidence combines company directories, public company profiles and provider job postings across 33 hubs and six countries. Spikes count documented organizations, not staff or local spending. The city-spending view uses a separate supplied allocation of \$68.6B across 100 cities. Its spike heights are proportional to estimated dollars, unlike the documented-locations view. City values are unverified and have no specified reference year; the cited BPO market reports do not establish those allocations. The separate 1,000-person staffing scenario retains its explicitly illustrative inputs.

JOB POSTINGS · SELECTED EXAMPLES, NOT EMPLOYMENT TOTALS

Job postings reveal where RCM teams hire and what they do.

PUBLIC POSTING SNAPSHOT / SEPTEMBER 8, 2026

Chennai, India · office-based

Accounts Receivable Callers

Pay not stated

Claim status → denials → follow-up

[Access Healthcare · source posting ↗](#)

Iloilo, Philippines

Healthcare Associates

Pay not stated

Healthcare support · broader than RCM

[Sagility · source posting ↗](#)

Philippines · remote, city unspecified

Internal Billing & Revenue Cycle Specialist

\$1,200–\$1,300 / month

Full cycle · behavioral health

[Assist World · source posting ↗](#)

Philippines · remote, city unspecified

Hospice Billing Manager

\$4–\$7 / hour

Specialty claims · hospice

[Moments Hospice · source posting ↗](#)

United States · remote / hybrid near named offices

Revenue Cycle Advisor

\$70,000–\$80,000 / year

RCM performance · customer operations

[HHAeXchange · source posting ↗](#)

Selected examples. Different duties, seniority, and pay periods; not a wage index.

Job postings reveal where RCM teams hire and what they do.

Employer career pages and employer-hosted Lever boards add role-level evidence to the location map. This five-example snapshot was reviewed September 8, 2026. It establishes advertised roles and requirements, not completed hires, city payroll, or net job growth.

Chennai’s A/R role combines payer calls, IVR, portals, and follow-up documentation. Assist World’s behavioral-health role spans claims, posting, appeals, and patient collections; Moments Hospice asks for Medicare and Medicaid billing expertise. Sagility’s broader healthcare-associate listing establishes recruiting in Iloilo, without proving that every advertised position performs RCM.

Access Healthcare’s posting dates to October 18, 2023 and remained in its careers directory when reviewed. The other four pages do not display posting dates. Continued visibility does not establish a newly opened vacancy. The hospice page uses both “manager” and “specialist”; we retain its title without inferring seniority.

Automation requirements: Offshore work includes specialty knowledge, appeals, patient conversations, and ownership of a full billing cycle. An AI replacement must cover that actual task bundle, including exceptions and quality checks. The U.S. advisor role emphasizes coordination and analysis; comparing its pay directly with a billing specialist would overstate geographic savings.

Pay-data limitations: Pay ranges retain the posting’s dollar symbol and original time unit. They are neither verified paid wages nor fully loaded costs. We do not annualize hourly contractor offers or use these examples to reset the globe’s cost assumptions. Remote national postings are not allocated to a city.

Sagility’s board repeats some requisitions and displays “500” jobs on many different listings. We do not add those fields into an employment estimate. A repeatable hiring dataset would deduplicate employer and requisition IDs, distinguish evergreen listings, record first and last observed dates, and separate RCM from broader healthcare support.


CHART CONTEXT

Location evidence combines company directories, public company profiles and provider job postings across 33 hubs and six countries. Spikes count documented organizations, not staff or local spending. The city-spending view uses a separate supplied allocation of \$68.6B across 100 cities. Its spike heights are proportional to estimated dollars, unlike the documented-locations view. City values are unverified and have no specified reference year; the cited BPO market reports do not establish those allocations. The separate 1,000-person staffing scenario retains its explicitly illustrative inputs.

VENDOR CAPABILITIES · ILLUSTRATIVE ACTIONS

AI can automate routine steps across the revenue cycle.

AUTOMATION TODAY / ACTION BY ACTION



Eligibility verification

Transaction automation + validation rules

1. Check coverage before billing
2. Flag missing or inconsistent insurance details
3. Route unresolved coverage to staff

Human handoff: Conflicting coverage and unusual benefit questions still need resolution.

[Product evidence ↗](#)

AI can automate routine steps across the revenue cycle.

Automation today includes checking eligibility, validating claims, retrieving claim status, and posting payments. Much of this uses structured transactions and rules. It does not require a conversational agent for every step.

Waystar advertises automated claim-status checks, payment posting, and reconciliation. Its claim-management offering also connects coverage checks with claim validation. These are product capabilities, not evidence that the entire revenue cycle runs unattended.

[Waystar · available automation packages ↗](#)

Waystar also offers denial prioritization, generated appeal letters, payer forms, submission, and delivery tracking. These automate parts of recovery without guaranteeing payment.

[Waystar · denial and appeal automation ↗](#)

Select an example in the graphic to see the actions, the technology involved, and the point where a person takes over. Each maps to the [front-to-back RCM timeline](#).


CHART CONTEXT

Sources establish offered capabilities, not industry penetration. Action sequences and exception handoffs illustrate workflow design; they do not claim that every vendor executes every action autonomously.

VENDOR CAPABILITIES · ILLUSTRATIVE ACTIONS

AI can turn clinical documentation into inputs for billing.

AUTOMATION TODAY / ACTION BY ACTION



Eligibility verification

Transaction automation + validation rules

1. Check coverage before billing
2. Flag missing or inconsistent insurance details
3. Route unresolved coverage to staff

Human handoff: Conflicting coverage and unusual benefit questions still need resolution.

[Product evidence ↗](#)

AI can turn clinical documentation into inputs for billing.

Language models can extract relevant evidence, summarize documentation, and identify possible coding gaps. That extends automation beyond copying fields and checking transaction formats.

AKASA describes re-analysis of clinical records to surface supported coding opportunities, followed by human expert review before billing. Its authorization-assistance offering is another example of helping staff prepare the work rather than independently deciding whether care should be approved.

[AKASA · clinical record analysis and human review ↗](#)

[AKASA · Authorization Advisor announcement, March 2024 ↗](#)

The practical distinction is between generating a suggestion, executing a permitted action, and closing an account correctly. A generated summary is useful, but it is not the same outcome as an accepted claim or reconciled payment.

CHART CONTEXT

Sources establish offered capabilities, not industry penetration. Action sequences and exception handoffs illustrate workflow design; they do not claim that every vendor executes every action autonomously.

Outsourced work is easier to automate; in-house labor offers larger savings per hour.

AUTOMATION TODAY / ACTION BY ACTION

Front end Eligibility verification	Front end Authorization preparation
Mid-cycle Coding and documentation review	Back end Claim preparation
Back end Claim-status follow-up	Back end Payment posting
Back end Denials and appeals	

Eligibility verification

Transaction automation + validation rules

1. Check coverage before billing
2. Flag missing or inconsistent insurance details
3. Route unresolved coverage to staff

Human handoff: Conflicting coverage and unusual benefit questions still need resolution.

[Product evidence ↗](#)

Outsourced work is easier to automate; in-house labor offers larger savings per hour.

Standardized outsourced queues are likely to be automated first. Work that has already been packaged for an external team often has documented steps, repeatable inputs, measurable outputs, and clear handoffs. Eligibility checks, status follow-up, and routine posting are examples of bounded actions that can be easier to automate.

This is our sequencing thesis, not a measured industry-wide order. Outsourcing does not itself make a task easy: vendors also handle complex coding and appeals, while many in-house teams run highly standardized processes.

Higher-cost in-house labor creates a larger incentive to reduce labor spend per hour saved. When a domestic employee’s fully loaded cost exceeds that of an offshore delivery role, the same successful automation has more potential dollar value in-house. Actual savings depend on the role, service contract, review burden, and whether capacity can be redeployed or expense reduced.

In the [illustrative workforce scenario](#), an \$80,000 annual role costs four times a \$20,000 role. Saving the same fraction of time therefore has four times the gross labor value—before implementation, oversight, and exceptions. Those costs are editable assumptions, not national wage estimates.

Prioritize repeatable tasks. Evaluate more complex workflows against their labor costs.

The existing model also allocates 65% to employee labor versus 18% to outsourced labor. Those are broad administrative scenario shares, not measured RCM staffing shares or immediately realizable savings.

CHART CONTEXT

Sources establish offered capabilities, not industry penetration. Action sequences and exception handoffs illustrate workflow design; they do not claim that every vendor executes every action autonomously.

AI can keep routine RCM queues moving outside office hours.

The advantage

Keep routine queues moving after hours, process work in parallel, and absorb volume spikes.

The tradeoff

Availability depends on system uptime, valid credentials, and payer office hours.

In an RCM workflow

Recheck pending claims overnight; escalate unresolved accounts when the team returns.

AI can keep routine RCM queues moving outside office hours.

AI does not need sick time, vacation coverage, or a shift change. For bounded work, it can reduce scheduling constraints and keep a backlog moving outside office hours. Parallel execution can add capacity without recruiting and training a new person for every increment of volume. But compute, rate limits, maintenance, integration support, and human oversight still cost money. Around-the-clock availability is a system design objective, not an inherent uptime guarantee.

Recheck pending claims overnight; escalate unresolved accounts when the team returns.

[Compare the full cost per completed task →](#)

CHART CONTEXT

These are illustrative operating examples. AI handles defined work; people retain exception ownership and accountability. Availability, document coverage, and repeatability depend on implementation.

AI can review long patient records without fatigue.

The advantage

Apply the same evidence checklist across lengthy records without boredom or fatigue.

The tradeoff

OCR errors and retrieval gaps can omit important evidence from long records.

In an RCM workflow

Track every page, extract relevant evidence with citations, and flag unreadable scans or conflicting dates.

AI can review long patient records without fatigue.

An AI workflow can be designed to process an entire 200-page authorization or appeal packet, rather than relying on a quick skim. That makes comprehensive review more economically feasible. It should account for all pages, check extraction quality, and connect each conclusion to its source. A large context window alone does not guarantee exhaustive comprehension: missing pages, poor scans, retrieval omissions, and degraded recall can all change the answer.

Track every page, extract relevant evidence with citations, and flag unreadable scans or conflicting dates.

[Anthropic · long-context limitations and context engineering ↗](#)

[Compare the full cost per completed task →](#)

CHART CONTEXT

These are illustrative operating examples. AI handles defined work; people retain exception ownership and accountability. Availability, document coverage, and repeatability depend on implementation.

Automated workflows need validation and ongoing monitoring.

The advantage

Reuse approved checks, retain structured logs, and apply a validated workflow update across many accounts.

The tradeoff

Faulty rules can affect many accounts before detection. Language-model outputs also require repeatability testing.

In an RCM workflow

Verify the patient, policy version, and submission outcome; prevent duplicate claims before retrying.

Automated workflows need validation and ongoing monitoring.

Organizations want consistent checklist execution and visibility into what happened on each account. Software can make versioned rules and audit trails easier to apply across a large queue. These controls require explicit implementation. Scaling requires validation, outcome monitoring, and a way to stop or roll back a faulty workflow. Track completion quality alongside throughput.

Verify the patient, policy version, and submission outcome; prevent duplicate claims before retrying.

[Compare the full cost per completed task →](#)

CHART CONTEXT

These are illustrative operating examples. AI handles defined work; people retain exception ownership and accountability. Availability, document coverage, and repeatability depend on implementation.

Missing patient context can lead AI to take the wrong RCM action.

The advantage

Handle known cases and assemble context so a person can spend less time searching.

The tradeoff

The system may not know about a side conversation, a special arrangement, or a patient’s circumstances—and may not recognize that it should ask.

In an RCM workflow

A patient says a financial counselor paused collections. Stop outreach and verify the arrangement instead of following the next reminder step.

Missing patient context can lead AI to take the wrong RCM action.

Experienced staff notice when the literal instruction is no longer the right action. A payer representative may mention an exception on the phone; a clinician may change the treatment plan; a patient may reveal a hardship absent from the record. An AI can reason about unfamiliar situations, but it cannot reliably reconstruct facts it never received. Out-of-band events need a path into shared context, and material ambiguity needs a human handoff rather than a confident guess.

A patient says a financial counselor paused collections. Stop outreach and verify the arrangement instead of following the next reminder step.

[Explore repeatability and exception handling →](#)

CHART CONTEXT

These are illustrative operating examples. AI handles defined work; people retain exception ownership and accountability. Availability, document coverage, and repeatability depend on implementation.

RCM AI needs current payer rules and provider-specific context.

The advantage

Use current policy documents, local procedures, and approved examples to adapt general capabilities to a defined task.

The tradeoff

Rare exceptions, changing payer rules, and specialty-specific cases may be poorly represented. Historical outcomes can encode errors or inconsistent practice.

In an RCM workflow

A new denial code or conflicting policy date should trigger review—not an invented explanation or an unsupported appeal.

RCM AI needs current payer rules and provider-specific context.

General model training may not include a provider’s contracts, local workflows, recent payer changes, or uncommon encounter types. Retrieving current sources helps supply facts; it does not substitute for representative evaluation of the actual task. Useful labeled data must show correct actions and verified outcomes, including difficult cases—not only whether a claim happened to be paid. Teams need authorized data, expert review, and ongoing checks as policies and systems change. Evaluate unfamiliar cases separately to identify gaps in task performance.

A new denial code or conflicting policy date should trigger review—not an invented explanation or an unsupported appeal.

[NIST · Generative AI Profile: reliability, data, and oversight ↗](#)

[Explore repeatability and exception handling →](#)

CHART CONTEXT

These are illustrative operating examples. AI handles defined work; people retain exception ownership and accountability. Availability, document coverage, and repeatability depend on implementation.

Document AI can extract evidence from medical records.

01	02	03	04
PDF (image)	Voice	Web	Computer



RELATIVE DEPLOYMENT COMPLEXITY

Lower

Document AI can extract evidence from medical records.

Healthcare work arrives as PDFs, scanned orders, referral packets, and faxes. Document AI turns those inputs into structured information.

For a bounded task, this is usually the easiest starting point: the input is fixed and extracted fields can be checked against the original page.

Validation still matters. Missing pages, handwriting, ambiguous identifiers, and contradictory evidence need explicit handling. Extraction accuracy is not the same as a correctly completed authorization.

[Explore GenHealth document technology ↗](#)

CHART CONTEXT

The deployment-complexity ordering is GenHealth’s heuristic for bounded administrative tasks. Access, integration, and risk can change the order.

Voice agents can handle routine RCM calls.

01	02	03	04
PDF (image)	Voice	Web	Computer



LISTEN → UNDERSTAND → CONFIRM

“Let’s verify the status of this authorization.”

RELATIVE DEPLOYMENT COMPLEXITY

Moderate

Voice agents can handle routine RCM calls.

Voice is often the next easiest modality to deploy for bounded workflows: status checks, scheduling, reminders, and information collection.

The conversation can follow a defined objective, confirm what was heard, and escalate when the other party cannot provide the answer. Identity checks, consent requirements, background noise, IVR menus, and write-back still shape the implementation.

Our practical ordering is **PDF → voice → web → computer**. This is a deployment heuristic, not a universal benchmark.

[Example: Hippocratic AI · voice agents ↗](#)

CHART CONTEXT

The deployment-complexity ordering is GenHealth’s heuristic for bounded administrative tasks. Access, integration, and risk can change the order.

Web agents can complete RCM tasks across payer and provider portals.

01
PDF (image)

02
Voice

03
Web

04
Computer



RELATIVE
DEPLOYMENT
COMPLEXITY

Higher

Web agents can complete RCM tasks across payer and provider portals.

Portals and EHRs require more than understanding language. The agent must authenticate, find the correct record, navigate permissions, enter data, and verify that the system accepted the change.

Sessions expire. Forms change. Different payers expose different workflows. The difficulty comes from reliably executing the whole sequence.

Structured interfaces and authorized network requests can reduce reliance on repeated visual interpretation. Where those interfaces are unavailable, browser and human fallbacks still matter.

[GenHealth · automated workflows ↗](#)

CHART CONTEXT

The deployment-complexity ordering is GenHealth’s heuristic for bounded administrative tasks. Access, integration, and risk can change the order.

Computer-use automation still faces cost and reliability limits.

01 PDF (image) 02 Voice 03 Web 04 Computer



RELATIVE DEPLOYMENT COMPLEXITY ■ ■ ■ ■ Highest

Computer-use automation still faces cost and reliability limits.

A visual computer-use agent reads screenshots and moves a cursor through applications. It can reach legacy systems that lack a convenient integration.

But screen interpretation, coordinate selection, changing layouts, and long action sequences introduce failure points. Repeated screenshots, reasoning, and retries also consume tokens and time.

Computer-use models are capable, but a successful demonstration does not establish repeatability or favorable economics on a production workflow. Both need to be measured.


[Anthropic · computer-use evaluation methodology ↗](#)

CHART CONTEXT

The deployment-complexity ordering is GenHealth’s heuristic for bounded administrative tasks. Access, integration, and risk can change the order.

AI can extract evidence for deterministic policy checks.

01 / Understand the evidence



AI can extract evidence for deterministic policy checks.

Medical policies contain criteria, thresholds, exclusions, and evidence requirements. Their application needs explicit logic and a traceable policy version.

AI can help extract and organize the evidence. A validated decision tree then applies the approved rules and records which criteria were met, missing, or ambiguous.

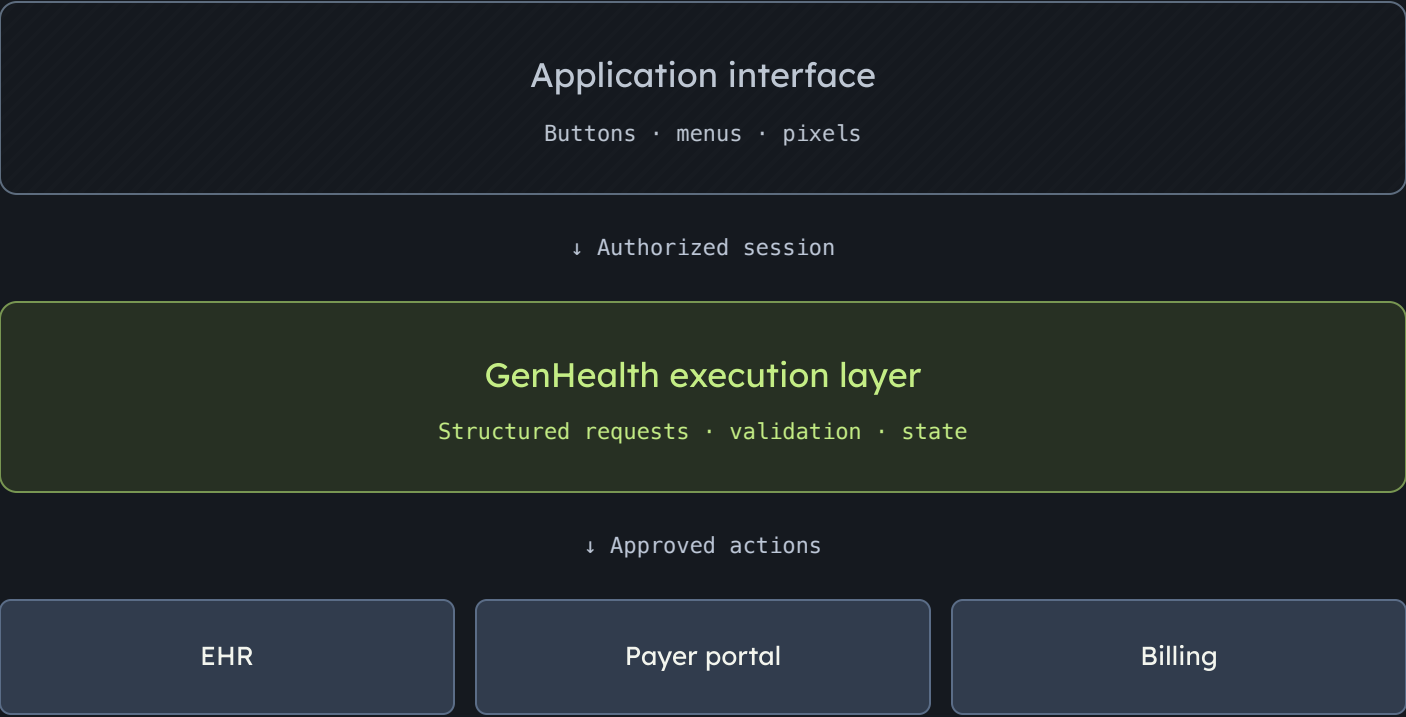
Consistency alone does not establish correctness. Rules need clinical and operational validation; missing evidence must route to review rather than become an invented answer.

[GenHealth · policy decision trees ↗](#)

CHART CONTEXT

This is GenHealth’s architecture and product approach. Deterministic execution does not make extracted facts infallible or prevent external systems from changing.

GenHealth uses network interactions to automate browser workflows.



GenHealth uses network interactions to automate browser workflows.

GenHealth’s approach is to translate a workflow into repeatable execution through a purpose-built browser and the authorized network interactions behind the application.

Instead of asking a vision model to rediscover each button on every run, the workflow uses structured actions, known state, and explicit checks.


This architecture is designed to reduce repeated model calls and improve consistency. Its advantage must still be demonstrated on the customer’s systems, including authentication, changing endpoints, and error recovery.

[GenHealth · behind-the-scenes execution ↗](#)

CHART CONTEXT

This is GenHealth’s architecture and product approach. Deterministic execution does not make extracted facts infallible or prevent external systems from changing.

RCM automation needs repeatable results and explicit error handling.



RCM automation needs repeatable results and explicit error handling.

Test repeated execution using the same approved inputs again under the same policy, workflow version, and system state.

Look for validated preconditions, safe retries, duplicate prevention, outcome checks, version control, and an audit record. When the environment changes, the workflow should stop or escalate deliberately.

Define completion criteria and the conditions that require a workflow to stop or escalate.

Deterministic execution is a design property. Accuracy, coverage, and availability remain separate measurements.

CHART CONTEXT

This is GenHealth’s architecture and product approach. Deterministic execution does not make extracted facts infallible or prevent external systems from changing.

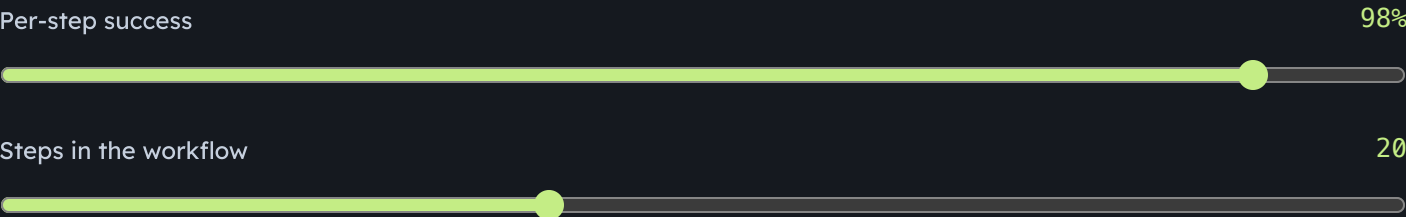

MATHEMATICAL SCENARIO · NOT AN AI BENCHMARK

Errors compound across the steps of an automated workflow.

THE COMPOUNDING PROBLEM

66.8%

Illustrative end-to-end success



Errors compound across the steps of an automated workflow.

At 98% success per step, a 20-step workflow has only about 67% end-to-end success under a simple independent-error model. Adjust the sliders to explore the difference.

Published benchmark scores also need context. Anthropic’s February 2026 system card reports OSWorld-Verified scores of **61.4% for Sonnet 4.5** and **72.5% for Sonnet 4.6**. These are general computer-task results under specified evaluation conditions, not RCM accuracy rates or a current leaderboard.

Document extraction, coding accuracy, task completion, and claim acceptance have different denominators. This report does not assign unsupported, comparable accuracy percentages to vendors or modalities.

[Anthropic · Sonnet 4.6 system card, February 2026 ↗](#)

CHART CONTEXT

The calculator is a scenario, not a model benchmark: p^n assumes independent steps, equal success probabilities, no retries, and failure if any step fails. Actual errors can be correlated.

EDITABLE COST SCENARIO • NOT MEASURED PRICES

AI savings depend on the full cost of completing each task.

COST PER COMPLETED TASK

AI + review

\$6.00

Manual labor

\$5.00

AI costs more in this scenario.



AI savings depend on the full cost of completing each task.

For tasks people complete quickly, repeated model calls, visual processing, retries, and human rework can make AI more expensive than manual execution.

That is a scenario to test, not an inevitable property of computer use. Model choice, caching, workflow length, supervision, and labor rates change the result.

Use the calculator to compare **cost per completed task**. The starting values are illustrative assumptions, not current model prices or measured GenHealth results.

Include review, rework, and exception handling in cost-per-task calculations.

CHART CONTEXT

Calculator inputs are editable assumptions, not vendor prices. AI cost includes tokens × blended rate plus review labor. Token counts include all attempts per completed task. Platform, integration, telephony, maintenance, and downstream error costs are excluded and must be added for a full business case.

Chatbots answer questions when people ask.

ESTABLISHED PRODUCT PATTERNS



Answers a question

Chatbots answer questions when people ask.

The first familiar interface was a conversation: summarize a policy, explain a denial, draft a letter. A person still carried the answer into the operational system.


This remains useful work. But generating an answer and owning an outcome are different capabilities.

CHART CONTEXT

This is a capability framework, not a claim that the whole industry has reached a particular stage. People retain accountability, access control, and exception ownership.

AI agents can complete assigned tasks.

ESTABLISHED PRODUCT PATTERNS



AI agents can complete assigned tasks.

An agent adds tools and execution: retrieve a record, check eligibility, fill a form, or submit an approved request. It can carry out an ad hoc task across several steps.

The important boundary is the handoff. What can the agent actually complete, and what does it return to a person?

Adoption is uneven: HFMA’s February 2026 survey of 95 finance professionals found 27% deploying AI at scale across multiple functions and 53% piloting selected areas.


[HFMA · Revenue Cycle of the Future, April 2026 ↗](#)

CHART CONTEXT

This is a capability framework, not a claim that the whole industry has reached a particular stage. People retain accountability, access control, and exception ownership.

AI coworkers can manage defined work queues.

EMERGING OPERATING MODEL



AI coworkers can manage defined work queues.

The next step is an AI coworker: persistent context, assigned responsibilities, a work queue, and clear escalation rules.


It does not wait for a new prompt for each account. It monitors pending work, takes the next authorized action, and hands exceptions to the right person.

“Coworker” describes an operating model, not independent legal or clinical accountability. People set the permissions, policies, and quality standards.

CHART CONTEXT

This is a capability framework, not a claim that the whole industry has reached a particular stage. People retain accountability, access control, and exception ownership.

AI operators would manage entire processes under human supervision.



AI operators would manage entire processes under human supervision.

Our longer-term thesis is an operating layer that coordinates whole administrative processes. One person—or a small team—could supervise work previously distributed across many queues and applications.

The shift is from executing tasks on demand to automating their initiation, completion, verification, and follow-up.

The operator model would fail to scale if exception volume, supervision, or maintenance costs rise as fast as the work it takes on. Persistent dependence on human judgment could favor bounded assistants instead of a small team supervising an entire operation.


This is a forecast, not today’s industry baseline. It depends on dependable integrations, repeatable execution, monitoring, and humans who can resolve the cases the system cannot.

CHART CONTEXT

This is a capability framework, not a claim that the whole industry has reached a particular stage. People retain accountability, access control, and exception ownership.

Cross-system agents can connect work across multiple applications.

TWO PATHS TO AN AGENTIC SYSTEM



Cross-system agents can connect work across multiple applications.

Application creators can add native agents with deep access to their own product. That can make a single system much more effective.

Third parties can build an orchestration layer across the EHR, payer portal, billing platform, documents, and phone. This reaches the handoffs where healthcare work often stalls.

GenHealth’s thesis: the cross-system approach is more likely to deliver cross-system automation because RCM workflows involve several vendors’ systems. Native agents can still be valuable components; the outcome is not settled.

Native agents could win in a single-vendor environment with deep EHR access, supported write operations, and fewer external handoffs. Cross-system automation can lose its advantage when integration maintenance, fragmented permissions, or exception work outweigh the benefits of broader coverage.

We would revise our thesis if matched real-world evaluations showed native agents delivering better verified completion, lower all-in cost, and fewer safety incidents across the same multi-system case mix. A hybrid approach may outperform either alone.

CMS’s prior authorization rule also pushes interoperability forward, with API requirements generally beginning in 2027 for affected payers.


[CMS · interoperability and prior authorization rule ↗](#)

CHART CONTEXT

The 17% software/IT, 18% outsourced labor, and 65% employee labor split is an illustrative rounded reference-model mix, not audited market shares or a job-displacement forecast. Vendors may combine native agents and cross-system execution.

AI can automate administrative labor as well as software tasks.

WHAT ARE YOU AUTOMATING?



AI can automate administrative labor as well as software tasks.

Adding a copilot to SaaS improves the application. Automating labor means completing the work people perform across applications and organizations.

The reference model illustrates the distinction: approximately **17% SaaS/IT**, **18% outsourced labor**, and **65% employee labor**. These are directional estimates, not audited national accounts.

Toggle the graphic to compare the two scopes. GenHealth’s ambition is to address the cross-system administrative workload, including the work done by internal teams and service providers.

Spend is not equivalent to automatable savings, and automating a task does not mean eliminating an entire job.

CHART CONTEXT


The 17% software/IT, 18% outsourced labor, and 65% employee labor split is an illustrative rounded reference-model mix, not audited market shares or a job-displacement forecast. Vendors may combine native agents and cross-system execution.

VENDOR DISCLOSURES · NOT MARKET SHARES

Established outsourcing vendors already operate large RCM businesses.

RCM VENDOR REVENUE DISCLOSURES

Reported service revenue · different years



Broader totals · RCM portion undisclosed



Additional vendors · select for available evidence

- Guidehouse
- Conifer
- Omega Healthcare
- GeBBS
- Parallon / HCA

R1 RCM

\$2.254B

Net services revenue · FY2023

Enterprise revenue-cycle operations plus modular services. Reported revenue includes operating, incentive, and modular fees; it is not a pure end-to-end outsourcing revenue figure.

[Read the source ↗](#)

Established outsourcing vendors already operate large RCM businesses.

The established RCM firms take responsibility for people, work queues, and financial operations. Their footprint can span the [entire revenue cycle](#), or a specific function such as coding, payment posting, or collections.

R1, Ensemble, Optum, and Guidehouse are the four firms covered in KLAS’s 2025 end-to-end revenue cycle outsourcing report. Conifer is another established enterprise operator. Access Healthcare, Omega, and GeBBS illustrate the global delivery side of the market.

[KLAS · End-to-End Revenue Cycle Outsourcing 2025 ↗](#)

Parallon belongs in the operating landscape too, but its HCA shared-services role is different from an independent third-party outsourcing business. A market-share calculation needs to decide whether captive operations are in or out.

[Parallon · services for HCA Healthcare ↗](#)

These firms connect directly to the report’s [global workforce](#) and [outsourced labor opportunity](#). They are potential adopters, partners, and competitors for AI workflow companies.

CHART CONTEXT


These firms are representative examples, not a market-share ranking. No consistent public numerator and denominator were established. “Not isolated” means unavailable in this review, not zero. Vendor fees, patient revenue managed, and employee counts measure different things.

VENDOR DISCLOSURES · NOT MARKET SHARES

Vendor disclosures do not support a comparable RCM market-share ranking.

RCM VENDOR REVENUE DISCLOSURES

Reported service revenue · different years



Broader totals · RCM portion undisclosed



Additional vendors · select for available evidence

- Guidehouse
- Conifer
- Omega Healthcare
- GeBBS
- Parallon / HCA

R1 RCM

\$2.254B

Net services revenue · FY2023

Enterprise revenue-cycle operations plus modular services. Reported revenue includes operating, incentive, and modular fees; it is not a pure end-to-end outsourcing revenue figure.

[Read the source ↗](#)

Vendor disclosures do not support a comparable RCM market-share ranking.

A defensible percentage ranking is not available from the public evidence reviewed here. The missing piece is comparable U.S. outsourced-RCM service revenue for every firm and a matching total market denominator.

The service-revenue chart shows R1’s \$2.254B FY2023 net services revenue and Ensemble’s \$600.0M FY2020 net revenue. Ensemble also disclosed \$401.1M for the first half of 2021. The chart retains the annual historical figure instead of annualizing a partial period or projecting it to 2026. These are RCM-related company revenues with different dates and service mixes, so bar lengths do not establish current market shares.

[R1 · FY2023 Form 10-K ↗](#)

[Ensemble · 2021 SEC prospectus ↗](#)

The separate, hatched chart shows Optum Insight’s \$19.4B FY2025 total and the greater than \$800M combined-revenue outlook announced for Smarter Technologies in May 2025. Optum includes payment integrity, analytics, government, and life-sciences services. Smarter combines Access Healthcare, SmarterDx, and Thoughtful.ai. Neither disclosure isolates the revenue of a standalone provider-RCM outsourcing business. The two charts use different scales, shown above their bars.

[UnitedHealth · original FY2025 results ↗](#)

[UnitedHealth · FY2025 business descriptions ↗](#)

[New Mountain · combination announcement ↗](#)

CHART CONTEXT


These firms are representative examples, not a market-share ranking. No consistent public numerator and denominator were established. “Not isolated” means unavailable in this review, not zero. Vendor fees, patient revenue managed, and employee counts measure different things.

VENDOR DISCLOSURES · NOT MARKET SHARES

Several major vendors do not disclose standalone RCM revenue.

RCM VENDOR REVENUE DISCLOSURES

Reported service revenue · different years



Broader totals · RCM portion undisclosed



Additional vendors · select for available evidence

- Guidehouse
- Conifer
- Omega Healthcare
- GeBBS
- Parallon / HCA

R1 RCM

\$2.254B

Net services revenue · FY2023

Enterprise revenue-cycle operations plus modular services. Reported revenue includes operating, incentive, and modular fees; it is not a pure end-to-end outsourcing revenue figure.

[Read the source ↗](#)

Several major vendors do not disclose standalone RCM revenue.

GeBBS’ India arm reported ₹1,938 crore of consolidated FY2025 revenue, according to VCCircle’s report citing VCCEdge. This is secondary reporting on a legal-entity consolidation, with no verified global or RCM-only reconciliation. It stays outside the dollar chart. Guidehouse’s RCM revenue and Omega’s service revenue were not established from the reviewed sources; employee counts cannot supply those amounts.

[VCCircle · GeBBS India-arm financials ↗](#)

Tenet’s FY2025 filing includes Conifer within Hospital Operations and describes RCM and value-based care together. The reviewed tables do not isolate Conifer’s RCM fees. Parallon is HCA’s shared-services operation. Treating HCA hospital revenue as Parallon sales would count payments for care as outsourcing fees.

[Tenet · FY2025 Form 10-K ↗](#)

Ensemble’s more than \$55B of client revenue managed in June 2026 and Conifer’s \$25B website figure remain useful coverage measures. Neither is vendor revenue, and the sources do not provide a supported fee rate for converting them into vendor sales.

[Ensemble · June 2026 managed revenue ↗](#)

[Tenet · Conifer managed revenue ↗](#)

KLAS satisfaction scores and purchasing consideration are also not revenue market share. Nor is the report’s broad administrative or outsourced-labor scenario a valid denominator for narrow RCM vendor sales.

CHART CONTEXT


These firms are representative examples, not a market-share ranking. No consistent public numerator and denominator were established. “Not isolated” means unavailable in this review, not zero. Vendor fees, patient revenue managed, and employee counts measure different things.

VENDOR DISCLOSURES · NOT MARKET SHARES

Established RCM vendors are investing in AI automation.

RCM VENDOR REVENUE DISCLOSURES

Reported service revenue · different years



Broader totals · RCM portion undisclosed



Additional vendors · select for available evidence

- Guidehouse
- Conifer
- Omega Healthcare
- GeBBS
- Parallon / HCA

R1 RCM

\$2.254B

Net services revenue · FY2023

Enterprise revenue-cycle operations plus modular services. Reported revenue includes operating, incentive, and modular fees; it is not a pure end-to-end outsourcing revenue figure.

[Read the source ↗](#)

Established RCM vendors are investing in AI automation.

The 2025 announcement combining Access Healthcare with SmarterDx and Thoughtful.ai illustrates the convergence of delivery operations and AI. “Old guard” describes an established customer and operating footprint, not an absence of technology investment.

For buyers, compare the partner’s actual responsibility: staffing a queue, managing the full cycle, licensing software, or completing work through automation. Bundled contracts can contain several of these at once.

For GenHealth, the strategic question is how repeatable cross-system automation fits into those operations—and whether it improves the full cost per completed outcome.

[Explore the economics of execution →](#)

CHART CONTEXT

These firms are representative examples, not a market-share ranking. No consistent public numerator and denominator were established. “Not isolated” means unavailable in this review, not zero. Vendor fees, patient revenue managed, and employee counts measure different things.

Early RCM automation can improve collections before competitors adapt.

A COMPETITIVE HYPOTHESIS TO TEST

Potential benefit

Reach recoverable work at lower cost

Competitive response

Plans and other providers can adopt AI too

Required evidence

Compare net outcomes, costs, and case mix over time

Not established

Early adoption by itself causes superior returns

Early RCM automation can improve collections before competitors adapt.

Our hypothesis: early RCM automation can create an operating advantage, but timing alone does not establish superior returns. Providers can close documentation gaps, pursue recoverable balances, and process work that their existing teams could not economically reach.

When a provider’s automation is ahead of a plan’s review capabilities, it may have an operational advantage: stronger evidence, more complete submissions, faster follow-up, and greater capacity to challenge incorrect denials. Collecting legitimately owed revenue at lower cost can improve profitability.

That relative advantage can narrow. Plans can deploy their own automation, strengthen review and audit processes, and adapt contracting and reimbursement approaches. As productivity gains become expected, some value may be priced into vendor fees and future contract negotiations. Late adopters may still gain efficiency but retain less of the initial upside.

Use measured net financial benefit to decide whether to expand a deployment.

This is a competitive thesis, not a universal law or a claim that payer countermeasures will erase every gain. Plan sophistication varies, and better tools create value only when they produce correct, compliant, economically useful outcomes.

CHART CONTEXT

These are GenHealth-reported customer results with different scopes and measures, not guaranteed returns or a controlled early-versus-late adoption comparison.

CUSTOMER-REPORTED · EARLY SIGNALS

GenHealth customers report gains in collections, workload, and billing speed.

EARLY CUSTOMER SIGNALS / DIFFERENT MEASURES

Piedmont Medical Solutions

+34.2% estimated collections association

Only 3 treated paid months; not a controlled causal estimate.

[Read the case study ↗](#)

GuideHealth

Up to 60% reported UM workload reduction

\$1.2M annual savings are projected, not realized savings.

[Read the case study ↗](#)

MedExpress

62% lower billed days reported

Realized administrative-cost savings were not quantified in dollars.

[Read the case study ↗](#)

GenHealth customers report gains in collections, workload, and billing speed.

Higher revenue collections — Piedmont. GenHealth’s case study reports a +34.2% paid-to-date collections regression effect relative to the January 2025 baseline, based on a short observation window. The analysis adjusts for trend and seasonality; it is an association, not proof that AI alone caused the increase. Only three treated paid months were available.

[Piedmont · collections case study ↗](#)

Prior-authorization adjudication support — GuideHealth. The deployment reports up to 60% lower utilization-management workload and approximately \$1.2M in projected annual cost savings. The savings are projected, and the workflow supports clinical review and decision recommendations.

[GuideHealth · prior authorization case study ↗](#)

Less manual administration — MedExpress. Automated intake, eligibility checks, and order creation shift staff toward exceptions. The case study reports 62% lower billed days and more than 35,000 orders built. These describe operational activity; the study does not quantify a realized administrative-cost reduction in dollars.

[MedExpress · administrative workflow case study ↗](#)

These are GenHealth-reported customer outcomes, not typical-return guarantees or direct evidence comparing early and late adopters. Each case study provides the scope and methodology behind its figures.

CHART CONTEXT


These are GenHealth-reported customer results with different scopes and measures, not guaranteed returns or a controlled early-versus-late adoption comparison.

SURVEY RESPONSES • SCOPE-SPECIFIC

27% of HFMA respondents are scaling AI; 53% are piloting it.

PROVIDER REVENUE CYCLE / HFMA / FEBRUARY 2026

% of 95 healthcare finance respondents



Most HFMA respondents are piloting AI rather than deploying it at scale.

HFMA’s February 2026 survey of 95 healthcare finance professionals found 27% reporting AI deployment at scale across multiple functions and 53% conducting pilots in select areas. That makes the implementation gap a central RCM opportunity: buying or testing AI is much more common than operating it broadly.

These are respondent-reported organizational stages, not the percentage of claims handled autonomously or a nationally representative estimate of all providers. The survey does not establish achieved savings or collections uplift.

[HFMA: The Revenue Cycle of the Future, April 2026, p. 4 ↗](#)

CHART CONTEXT

HFMA reports provider-side revenue-cycle maturity in February 2026; NAIC reports broader insurer AI/ML use from November 2024–January 2025. These are historical survey observations, not 2026 projections or directly comparable adoption rates. Payer workflow bars show company counts on a common 0–50 scale; organizations may appear in multiple rows. Neither survey measures the share of national RCM spending automated.

SURVEY RESPONSES · SCOPE-SPECIFIC

84% of surveyed insurers already use AI/ML.

HEALTH INSURERS / NAIC / NOVEMBER 2024–JANUARY 2025

78 of 93 responding insurers use AI/ML



84% of surveyed health insurers use AI or machine learning.

NAIC’s May 2025 report found that 78 of 93 responding health insurers used AI/ML in some capacity. Sixteen states surveyed companies from November 2024 through January 2025, selecting larger insurers or those with significant state market share.

This is broad enterprise adoption, including statistical models, rather than a measure of generative AI agents or autonomous RCM. The surveyed business covers comprehensive major medical and student health products; it is not a census of every plan.

Do not subtract HFMA’s 27% from NAIC’s 84% to measure a payer lead. The dates, populations, technologies, and deployment thresholds differ. Our inference: an early-mover advantage must come from better execution in a particular workflow; providers should not assume the payer is starting without AI.


[NAIC: Health AI/ML Survey Report, methodology and Table 1 ↗](#)

CHART CONTEXT

HFMA reports provider-side revenue-cycle maturity in February 2026; NAIC reports broader insurer AI/ML use from November 2024–January 2025. These are historical survey observations, not 2026 projections or directly comparable adoption rates. Payer workflow bars show company counts on a common 0–50 scale; organizations may appear in multiple rows. Neither survey measures the share of national RCM spending automated.

Surveyed insurers already use AI in claims and prior authorization.

Companies reporting AI/ML already in production



Surveyed insurers already use AI in claims and prior authorization.

NAIC’s Table 5 reports 31 companies with AI/ML already in production for claims adjudication, 18 for prior authorization, and 42 for utilization, severity, or quality management. These are current-production responses, separate from planned deployments.

For the provider, these touch the front-end authorization and back-end payment stages in our [RCM timeline](#). The practical response is to improve evidence completeness, policy alignment, submission quality, and exception resolution, then measure acceptance and collections on matched case mixes.

Production use does not establish autonomous approval or denial, the share of claims touched, or accuracy. The report describes Table 5’s respondent base inconsistently; we retain its published company counts without calculating penetration percentages. These surveys also do not quantify how much of our modeled administrative spending is already automated.

[NAIC: Table 5, printed pp. 14–15 ↗](#)

CHART CONTEXT


HFMA reports provider-side revenue-cycle maturity in February 2026; NAIC reports broader insurer AI/ML use from November 2024–January 2025. These are historical survey observations, not 2026 projections or directly comparable adoption rates. Payer workflow bars show company counts on a common 0–50 scale; organizations may appear in multiple rows. Neither survey measures the share of national RCM spending automated.

RCM teams should validate AI before delegating work.

FROM PILOT TO OPERATIONS /ADOPTION

Increasing delegated responsibility

Conceptual stages; bar heights do not represent adoption rates.



RCM teams should validate AI before delegating work.

Adoption has several meanings: a purchased license, a pilot, a tool staff use, or a workflow that reliably completes work. Report the fraction of eligible accounts actually processed and the human touches still required, not only the number of AI projects.

The 2026 Guidehouse/HFMA survey page reports 69% outsourcing some or all of RCM and 88% identifying payer issues among their top three concerns. These are respondent findings, not AI penetration or market-share estimates.

[Guidehouse/HFMA · 2026 RCM trends ↗](#)


The path in the graphic is our proposed rollout sequence. Advance based on evidence, with explicit scope and a named operations owner. A mature eligibility workflow can coexist with a coding pilot inside the same organization.

CHART CONTEXT

The maturity path, roles, and failure modes are proposed operating guidance, not surveyed stages or measured frequencies. A/R cash release equals annual revenue ÷ 365 × days reduced. It is a one-time working-capital effect, not annual profit, and assumes stable revenue and collections.

RCM automation creates value through savings, collections, and faster cash.

FROM PILOT TO OPERATIONS /ECONOMICS

\$13.7M in earlier cash receipts

Annual net patient revenue: \$1000M



Reduction in A/R days:5

**Cash released: \$13.7M**

RCM automation creates value through savings, collections, and faster cash.

Four benefits need separate accounting: staff capacity released, expense actually avoided, additional revenue collected, and cash received sooner. They have different implications for a CFO.

Subtract integration, licensing, model usage, supervision, rework, and ongoing maintenance. Labor time saved may reduce overtime or hiring needs; it does not automatically reduce current payroll. Evaluate recurring benefits separately from one-time implementation costs.

The cash calculator illustrates how five fewer A/R days on \$1B annual net patient revenue could release approximately \$13.7M. That is accelerated cash, not new revenue or recurring profit. The estimate assumes stable daily revenue and collectible receivables.


[Compare cost per completed task →](#)

CHART CONTEXT

The maturity path, roles, and failure modes are proposed operating guidance, not surveyed stages or measured frequencies. A/R cash release equals annual revenue ÷ 365 × days reduced. It is a one-time working-capital effect, not annual profit, and assumes stable revenue and collections.

RCM automation depends on usable data, clear ownership, and exception handling.

FROM PILOT TO OPERATIONS /SUCCESS & FAILURE



RCM automation depends on usable data, clear ownership, and exception handling.

Successful rollout starts with a narrow queue, a measured baseline, reliable system access, and staff who can handle exceptions. Validate on ordinary cases and difficult ones, including payer variation, missing documents, conflicting records, and failed writes.

A common failure mode to test for is a strong demo followed by expensive review: the model looks accurate, but staff still check every field, repair integrations, and chase unresolved accounts. Another is optimizing throughput while increasing denials or patient complaints.

Track first-attempt completion, verified outcome, review minutes, backlog age, final denials, and full cost per account. Compare matched case mixes over sufficient time for claims to settle. Agree on stop conditions before the pilot, and retain a staffed fallback.


Expand deployment after verifying improvements in operational and financial results.

CHART CONTEXT

The maturity path, roles, and failure modes are proposed operating guidance, not surveyed stages or measured frequencies. A/R cash release equals annual revenue ÷ 365 × days reduced. It is a one-time working-capital effect, not annual profit, and assumes stable revenue and collections.

Automated RCM decisions need source evidence and verified outcomes.

FROM PILOT TO OPERATIONS /TRUST



Automated RCM decisions need source evidence and verified outcomes.

Trust comes from inspectable evidence and correction mechanisms. Staff should see which record and policy supported an action, what changed in the destination system, and why an exception was escalated.

For an appeal, keep links to the supporting chart passages and approved argument; verify delivery separately from eventual payment. For coding, preserve the expert review where required by the workflow. Measure error severity, not only average accuracy.

Test performance across payer and specialty groups and, where appropriate and lawful, relevant patient populations. Track complaints and overrides. Give reviewers time and authority to disagree, so human oversight does not become a rubber stamp.


[NIST · AI Risk Management Framework ↗](#)

CHART CONTEXT

The maturity path, roles, and failure modes are proposed operating guidance, not surveyed stages or measured frequencies. A/R cash release equals annual revenue ÷ 365 × days reduced. It is a one-time working-capital effect, not annual profit, and assumes stable revenue and collections.

Every RCM agent needs defined permissions and an accountable owner.

FROM PILOT TO OPERATIONS / GOVERNANCE



Every RCM agent needs defined permissions and an accountable owner.

Governance answers concrete questions: Who can authorize a claim submission? Which actions require review? Who owns a wrong payment adjustment? Who can pause the system when a payer changes its rules?

Use task-specific permissions and separate reading records from submitting claims, changing codes, adjusting balances, or initiating collections. Version models, policies, and workflows; revalidate material changes; log actions with appropriate access controls and retention. Treat instructions embedded in outside documents as untrusted content.

Map where patient data travels, which vendors and subcontractors process it, permitted uses, retention, and incident responsibilities. HHS states that covered entities and business associates using a cloud provider to handle ePHI on their behalf must establish an appropriate BAA and satisfy applicable HIPAA requirements, including risk analysis. A signed agreement alone is not proof of secure operation or clinical accuracy.

[HHS · HIPAA and cloud computing ↗](#)


NIST's framework is voluntary risk-management guidance, not certification. The ownership model shown here is our operational recommendation, not a claim that any vendor has implemented every control.

CHART CONTEXT

The maturity path, roles, and failure modes are proposed operating guidance, not surveyed stages or measured frequencies. A/R cash release equals annual revenue ÷ 365 × days reduced. It is a one-time working-capital effect, not annual profit, and assumes stable revenue and collections.

Model training requires controls against patient-data exposure.

FROM PILOT TO OPERATIONS / TRAINING DATA & PRIVACY



Also audit retention, logs, access, and subprocessors.

Model training requires controls against patient-data exposure.

Training-data handling deserves particular scrutiny in RCM procurement. Models can memorize sensitive training examples and disclose them in later outputs or extraction attacks. Training or fine-tuning on identifiable patient records creates a serious privacy risk; disclosure is possible, not a guaranteed or unavoidable outcome of every training run.

[NIST · Generative AI Profile: data privacy ↗](#)

Our recommended purchasing standard: prohibit vendors from using identifiable patient data to train shared models. If patient-derived data will be used for training, require documented de-identification before it enters the training pipeline, including free-text notes, attachments, call transcripts, and feedback. Simply masking names or replacing patient IDs does not establish de-identification.

HHS recognizes Safe Harbor and Expert Determination as methods for HIPAA de-identification. Properly de-identified information is no longer PHI under HIPAA; calling it “anonymized” does not establish that standard. HHS also notes that residual identification risk is not zero. This is our recommended vendor requirement, not a claim that HIPAA universally prohibits every authorized use of PHI for training.

[HHS · De-identification guidance ↗](#)

Processing a patient record to perform an authorized RCM task is different from using it to update model weights. Verify both paths: training and fine-tuning permissions, prompt and output retention, logs, retrieval stores, subprocessors, and cross-customer access. A no-training promise does not by itself prevent leakage through those other systems. Governance must also cover action accuracy, access controls, and human accountability.

CHART CONTEXT

The maturity path, roles, and failure modes are proposed operating guidance, not surveyed stages or measured frequencies. A/R cash release equals annual revenue ÷ 365 × days reduced. It is a one-time working-capital effect, not annual profit, and assumes stable revenue and collections.

Federal policy aims to reduce healthcare administrative work.



4.4M hours / year

CMS estimated paperwork reduction from the 2019 rule

Federal policy aims to reduce healthcare administrative work.

Administrative burden is a recognized federal policy problem. In 2019, CMS estimated that its Omnibus Burden Reduction rule would remove 4.4 million hours of provider paperwork annually by eliminating unnecessary, obsolete, or excessively burdensome requirements. This was a forecast for a specific rule, not measured savings from AI.

[CMS · 2019 burden-reduction rule ↗](#)


CMS also publishes concrete changes to simplify Medicare claims documentation. These efforts acknowledge that government requirements themselves can create work that can be reduced while preserving patient protection and payment integrity.

[CMS · Simplifying documentation ↗](#)

CHART CONTEXT

Federal examples reviewed September 8, 2026. The 2019 hours figure is CMS's projected paperwork reduction for one rule, not observed AI savings. Dates and obligations vary by program. The protection/work/simplification diagram is an editorial framework, not a quantified cost allocation.

Healthcare regulations also create documentation and reporting work.



Reuse across authorized workflows with standardized exchange

Healthcare regulations also create documentation and reporting work.

Documentation supports accountable payment and oversight, but assembling evidence, maintaining records, and communicating decisions also consume staff time. Our reading: regulation is both a source of administrative requirements and a tool for simplifying them. It is not the only source; payer policies and fragmented systems also create work.

The 2024 prior authorization rule illustrates the tradeoff. It requires specific denial reasons and public reporting of authorization metrics, alongside standardized electronic exchange. Existing notice requirements remain. Greater transparency therefore requires operational work even as standardization aims to reduce repeated calls and submissions.

[CMS · CMS-0057-F requirements ↗](#)

This work sits inside the documentation, authorization, and payment stages of our RCM model. We do not assign a national dollar share of administrative spending to regulation: the sources here do not establish that allocation or the net cost of all U.S. regulation.

CHART CONTEXT

Federal examples reviewed September 8, 2026. The 2019 hours figure is CMS’s projected paperwork reduction for one rule, not observed AI savings. Dates and obligations vary by program. The protection/work/simplification diagram is an editorial framework, not a quantified cost allocation.

New prior-authorization standards require operational changes.



Evidence → standardized APIs

Requirements generally begin in these years; scope varies

New prior-authorization standards require operational changes.

Under CMS-0057-F, operational requirements generally start in 2026 and API requirements in 2027; exact dates depend on payer type. Covered payers include Medicare Advantage, Medicaid and CHIP programs and managed care, and federally facilitated exchange plans. This is not a mandate covering every commercial payer.

For non-drug prior authorization, impacted payers other than those exchange plans face 72-hour expedited and seven-calendar-day standard decision limits. The rule’s prior authorization provisions exclude drugs. Teams should confirm their applicable program, deadline, and current rule before changing workflows.

[CMS · Scope, deadlines, and exceptions ↗](#)

Our recommendation: capture source evidence once, reuse it for authorized submissions, retain policy versions and acknowledgments, and route missing evidence to an accountable person. Measure total staff time, rework, and reporting effort—not just faster data entry. AI can help perform required documentation; it does not waive the underlying obligation or justify inventing evidence to complete a record.

CHART CONTEXT

Federal examples reviewed September 8, 2026. The 2019 hours figure is CMS’s projected paperwork reduction for one rule, not observed AI savings. Dates and obligations vary by program. The protection/work/simplification diagram is an editorial framework, not a quantified cost allocation.

GenHealth still needs a comparable, verified performance benchmark.

This report does not publish a verified, comparable GenHealth performance dataset. The customer case studies are not a substitute for one. We therefore cannot establish GenHealth's relative accuracy, repeatability, or all-in cost advantage from this report.

A useful next evaluation would take a bounded workflow such as eligibility verification, define consecutive eligible cases and exclusions, record payer and system mix, and compare against a matched manual baseline over disclosed dates. Count failed and escalated cases, not just successfully completed work.

Publish first-attempt completion with its denominator; repeat identical cases against fixed policies and system state; report exception rate, human review minutes, and cost per verified result including integration, retries, oversight, and rework. State sample sizes, uncertainty, and who independently checked the outcomes.

Until those data are available, the scorecard says “not reported.” Missing evidence is not zero performance, and a proposed evaluation is not a completed benchmark.

Available data leaves gaps in national RCM spending estimates.

We have no measured 2026 national allocation by RCM workflow, consistent outsourcing market shares, or national city-level staffing costs. The spending and globe models illustrate assumptions; they cannot establish the addressable market or automatable savings.

Customer results have different metrics and observation windows. They do not establish typical returns, realized GuideHealth dollar savings, MedExpress cost savings, or a causal advantage from adopting early. No new raw customer dataset was available for this revision; the short Piedmont observation window remains a material limitation.

Survey adoption is not autonomous task volume. Generic model benchmarks do not measure RCM outcomes. Cross-system orchestration and the operator model are GenHealth theses, with alternatives and failure conditions described in the architecture section.

Use the source dates, scope labels, and calculators to formulate a local business case. Validate it against your own case mix before treating it as a spending or staffing plan.

Follow the evidence.

This report combines published research, an illustrative GenHealth workflow model, and GenHealth’s product perspective. These are different kinds of evidence. It is not an independent vendor evaluation.

▼ Spending and model methodology

The opening uses CMS 2024 actuals (\$5,278.6B). The first scroll advances to the June 2026 CMS projection release: 2026 total \$6,017.4B; hospital care \$1,870.5B; physician/clinical services \$1,244.5B; retail prescriptions \$561.5B. Table 1 cells M5/O5 and Table 2 cells J4/J7/J9/J16 supply these inputs. All following spending displays use 2026. The other-services bucket is the remainder; shares and dots are rounded.

CMS does not project the broad administrative and 23-workflow categories used here. For a consistent scenario, GenHealth normalizes the approximate \$1T administrative estimate and the undated \$840B Series A model to a 2024 reference year, then multiplies them by $6,017.4 / 5,278.6 = 1.139961$. This holds spending shares fixed and assumes every workflow grows at the national rate. The normalization is our assumption, not a source date. The broad payer/provider split retains McKinsey’s historical 260:690 ratio; the core workflow model retains its distinct 257:583 ratio. These are GenHealth scenarios, not official CMS administrative forecasts.

The original 23-row model sums to \$583B provider + \$257B payer = \$840B. It excludes corporate overhead and includes clinical administrative work beyond narrow RCM. Projected row values use unrounded calculations; model headlines round to about \$5B and component bars to \$1B. Rounded amounts may not sum exactly. The 17% / 18% / 65% labor mix remains a fixed illustrative assumption. The download includes both reference and projected values so the calculation can be reproduced.

The lifecycle groups all 23 reference components into six sequential stages plus an ongoing support lane. Every row appears exactly once, so stage totals reconcile to the core workflow total. Combined rows can span stages; their placement is an editorial grouping, not an additional spending estimate. Denials can loop back to earlier stages. Each mirrored bar uses a common \$0–70B scale. A zero means no allocation in this model, not proof that the activity does not exist. The original reference shows providers on the left; this report reverses that orientation as labeled.

[Download the 2026 workflow scenario \(.csv\)](#) ↓

▼ Early-mover thesis and customer outcomes

The early-adoption, payer-response, and pricing discussion represents GenHealth’s outlook. Customer case studies demonstrate reported outcomes within particular deployments, not an early-versus-late adoption experiment. Piedmont’s +34.2% measure is an estimated regression coefficient relative to its January 2025 baseline, with 17 paid-month observations and three treated paid months. GuideHealth annual savings are projected. MedExpress’s billed-days reduction is operational evidence, not a measured payroll savings percentage. Collections, workload reduction, cash timing, and profit are distinct outcomes.

▼ Adoption, economics and governance

The Guidehouse/HFMA survey percentages refer to respondents and are not population estimates. Its public page contains conflicting AI/automation adoption figures, so this section does not reproduce a single AI adoption percentage from it. The maturity path, success criteria, failure modes, and ownership model are editorial recommendations, not measured prevalence. The A/R calculator is an illustrative steady-state working-capital calculation, not a savings forecast.

[Guidehouse/HFMA survey](#) · [NIST AI RMF](#) · [HHS cloud guidance](#). Reviewed September 8, 2026. NIST guidance is voluntary; the HHS discussion is scoped to the relationships described in its guidance. No certification or vendor compliance claim is made.

▼ Provider specialty and care-setting costs

CAQH’s 2023 Index specialty brief covers 2022 activity. Provider dollar estimates use rounded inputs: \$42B × 98%, \$19B × 97%, and \$12.5B × 97%. The unit-cost selector reproduces manual transaction costs from page 10. This limited transaction scope excludes information gathering, follow-up, and systems. JAMA’s 2018 study supplies encounter-level billing costs from one academic system, including overhead; the two sources therefore cannot be combined directly. Neither is inflated into a 2026 forecast.

Clinical specialties and care settings are separate dimensions. The provider workflow examples are editorial illustrations; CMS specialty spending measures care payments, not administrative expense. None of these datasets establishes a national all-payer RCM spending allocation by named specialty.

[CAQH specialty brief](#) · [JAMA study](#) · [CMS specialty data](#)

▼ AI workforce benefits and limitations

The workforce discussion is operational reasoning, not a measured AI-versus-human productivity study. The 200-page packet is an illustrative task size. Benefits depend on implemented coverage checks, reliable integrations, approved rules, logging, and review. Anthropic’s context-engineering discussion supports the distinction between context capacity and reliable recall; NIST’s July 2024 Generative AI Profile provides a framework for evaluating model reliability and data risks. RCM scenarios are editorial examples, not incidents attributed to a vendor.

▼ Automation examples and labor sequencing

Automation examples reference Waystar’s packages and claim-management pages and AKASA’s product explanations and announcements, reviewed September 8, 2026. These establish offered capabilities, not independent outcome validation or industry-wide penetration. Sequences and handoffs are illustrative workflow design examples. Structured transactions, rules-based automation, and generative AI are distinguished.

The expectation that standardized outsourced queues will be automated early is GenHealth’s thesis. Ease depends on process structure and system access, not employer or country alone. The incentive to automate higher-cost in-house labor is a conditional economic comparison; time released is not automatically a reduction in payroll. The 4× illustration inherits the editable geography scenario’s \$80,000 and \$20,000 annual cost assumptions.

▼ Outsourcing vendors and market-share limitations

Vendor evidence was reviewed September 8, 2026. KLAS’s 2025 end-to-end outsourcing report defines the enterprise comparison group, not a complete vendor census. The roster adds selected global delivery firms and an explicitly separate captive operation. It is not ordered by size or quality.

No verified percentage shares are assigned. A valid share requires same-period, same-geography, same-service vendor revenue divided by the corresponding outsourced RCM market total. R1’s FY2023 revenue includes modular services; Smarter Technologies’ >\$800M is the May 2025 announcement’s expected combined revenue; managed patient revenue and employee counts are different measures. None are extrapolated to 2026 using CMS spending growth. Do not add overlapping vendor claims, count Access and its announced combination twice, or use broad parent-company revenue as outsourced RCM sales.

Each vendor card links its primary source. The KLAS page supplies independent category coverage; company and investor disclosures supply scale and service descriptions. Unavailable measures are not zero. Reported scores, client counts, employees, contract wins, valuation, and patient revenue managed do not establish revenue market share.

▼ Workforce geography and globe methodology

Research reviewed September 8, 2026. R1’s 2023 headcount yields 12,300 / 29,900 = 41.1% U.S. and 58.9% international; approximate company counts are not an industry estimate. No national or city-level RCM payroll dataset was established in this research. Office directories support location presence only.

The expanded location layer covers 33 selected city / metro hubs, including Access Healthcare and Sagility directories. It groups multiple local offices and counts company presence per hub, not headcount or spending. Coordinates are approximate city centers. Sagility’s country table uses its FY2024-25 sustainability reporting perimeter of 38,754 employees, with no allocation of national totals to cities.

The separate spending layer models an illustrative 2026 annual staffing budget for 1,000 people. Offshore headcount is allocated 75% India, 20% Philippines, 5% Colombia. Annual cost equals people × editable cost per person. Country shares use total modeled staffing cost. Costs include an assumed labor burden, exclude vendor margins and technology spending, and are not sourced compensation benchmarks.

City weights within countries: U.S.—Dallas 50%, Hackensack 30%, Boca Raton 20%; India—Chennai 40%, Noida 20%, Pune 15%, Hyderabad 15%, Coimbatore 10%; Philippines—Metro Manila 100%; Colombia—Bogotá 100%. These arbitrary scenario weights reconcile country totals but do not rank actual cities. Spikes scale linearly with modeled dollars, subject to globe projection. Land dots derive from public-domain Natural Earth 1:110m land polygons.

[R1 2023 Form 10-K](#) · [Access Healthcare locations](#) · [Omega country offices](#) · [Natural Earth map attribution](#)

▼ Benchmarks, calculators and forward-looking claims

Computer-use scores are dated, vendor-reported general-purpose benchmarks, not healthcare accuracy. The reliability calculator assumes independent steps and no retries. The cost calculator uses editable illustrative assumptions and excludes platform, maintenance and integration costs. Neither calculator is a performance claim about GenHealth or another vendor.

The modality ordering, network-layer execution discussion, coworker/operator evolution, and cross-system orchestration thesis represent GenHealth’s approach and outlook. No verified comparative GenHealth accuracy or cost dataset was supplied for this report.

1. [CMS — 2025–2034 National Health Expenditure Projections](#)
June 2026 release; Tables 1 and 2, 2024 actuals and 2026 projections.
2. [CMS — National Health Expenditures 2024 Highlights](#)
US spending and service categories; actual expenditures.

3. [Penn LDI — Unpacking the Paradox of Health Care’s GDP Percentage](#)

February 6, 2025; approximate administrative spending context.

4. [McKinsey — Administrative simplification](#)

October 20, 2021; historical payer/provider split and functional definitions. Uses a 2019 baseline, not a current-year estimate.

5. [HFMA — Revenue Cycle of the Future](#)

April 2026; February survey, n = 95 healthcare finance professionals.

6. [Anthropic — Claude Sonnet 4.6 System Card](#)

February 2026; OSWorld-Verified evaluation and model comparison.

7. [CMS — Interoperability and Prior Authorization Final Rule](#)

Operational requirements generally begin in 2026; API requirements generally begin in 2027 for affected payers.

8. Vendor product descriptions: [AKASA](#), [Waystar](#),[Hippocratic AI](#), [GenHealth](#).

Referenced September 2026; descriptions are not independent validation.