Bringing Healthcare AI to PHI — Without the Compromise

Solution Brief · AI-Vendor Egress with Pseudonymization

Synthology Healthcare Solutions Group · 2026 Edition · v6 (2026-08-19)

At a glance

Healthcare organizations want the best AI medical vendors for their clinical and operational workflows: imaging AI for radiology and pathology, clinical NLP for documentation and coding, decision-support engines, foundational LLMs for narrative tasks, and specialty-workflow automation. Today they face a hard choice: lock into the small subset of vendors that sign a HIPAA Business Associate Agreement, or stay out of the market entirely because the AI vendor that solves their problem won't sign one.

Synthology AI-Vendor Egress is being built to remove the choice. At general availability, your on-prem Synthology Router will pseudonymize every PHI element before any byte leaves your hospital network. The pseudonymized request will flow through a Synthology cloud routing tier — operating under our BAA — to the AI vendor of your choice. The response will return the same way, and the on-prem Router will re-identify it for the radiologist, clinician, or care team. **The cloud will never see plaintext PHI. The vendor will never see plaintext PHI. Only your network will.**

This brief explains how the architecture is designed, what is shipping today, what is generally available, and what remains forward-looking, what the architecture does and does not protect against, and how it compares to the cloud-native AI options.

> **Status:** The customer-keypair architecture, DICOM-side reversible pseudonymization, hash-chained audit trail, and on-prem Router infrastructure are shipping today. The AI-vendor egress flow — cloud routing tier, HL7/FHIR pseudonymization, free-text PHI detection, and vendor adapters — reached **general availability** and is delivered by SynthCloudConnect together with the on-prem Router. A small number of forward-looking items remain, tagged **[Planned]**.

The problem

AI medical vendors are increasingly the table-stakes capability for radiology and pathology imaging interpretation, prior authorization, clinical-documentation improvement, coding assistance, structured-report generation, decision support, and dozens of other healthcare workflows. Across the imaging-AI, clinical-NLP, decision-support, and foundational-LLM categories, most of these vendors are not willing to operate as Business Associates under HIPAA. Some offer BAA-eligible enterprise tiers; many do not. Capabilities and contractual posture change quarter by quarter.

The two compromises healthcare organizations make today:

Compromise 1 — Lock into a cloud's BAA-eligible AI roster. Azure OpenAI, AWS Bedrock, Google Vertex AI all offer BAA-covered access to a curated subset of models. This is the most common path. The cost: vendor lock-in (you cannot use a competing model that isn't on that cloud's menu), single-vendor risk (a model deprecation, a price change, or a policy shift is non-negotiable), and continued plaintext PHI exposure inside the cloud provider's perimeter — under their BAA, but still inside their reach.

Compromise 2 — Don't use the AI. Forgo workflows your competitors are deploying because the model you actually want isn't BAA-eligible. This is increasingly an unsustainable position as patient and physician expectations rise.

Neither compromise is satisfying. Both are increasingly common.

The Synthology approach

Synthology's pseudonymization architecture is built on a simple principle: **the AI vendor doesn't need to see real PHI to do its job — it needs to see structurally faithful text.** A radiology report with patient names replaced by reversible token pseudonyms, MRNs replaced by reversible IDs, and dates shifted by a per-patient offset still produces an excellent impression summary. A prior-auth narrative with provider names tokenized still drafts well. The model gets what it needs; the patient identity stays in your network.

The architecture has four boundaries:

1. **On-prem ingress.** PHI enters from your EMR, PACS, or modality. Synthology's on-prem Router pseudonymizes every PHI element using your customer-managed keypair, encrypts the result, and writes a hash-chained audit row. DICOM path shipping; HL7 / FHIR / free-text paths generally available.
2. **Cloud transit.** The pseudonymized payload flows to Synthology's cloud routing tier, which orchestrates routing, archival, and vendor calls. **The cloud routing tier holds no decryption key.** Even Synthology employees with full cloud access cannot read PHI from this layer. Generally available.
3. **AI-vendor egress.** The cloud tier forwards the pseudonymized request to your chosen AI vendor over the vendor's standard API. The vendor processes pseudonymized text and returns a pseudonymized response. The cloud tier returns the response to your on-prem Router. Generally available.
4. **Re-identification.** The on-prem Router will re-identify the response using the same keypair and deliver plaintext to the workstation, EMR, or PACS that originated the request. For DICOM-side workflows, this is shipping today via the reversible pseudonymization pipeline.

The keypair lives only in your environment — today on the on-prem Router host (DPAPI machine scope on Windows, or AES-GCM under a deployment-managed KEK on Linux); customer-managed HSM and cloud key-vault providers are on the roadmap. If Synthology's cloud were entirely compromised tomorrow, the attacker would obtain pseudonymized data and the encryption envelopes around it; without your keypair, the data is mathematically unrecoverable.

How it works

> **Section convention.** Components in production today are tagged **[GA]** — this covers both the long-shipping DICOM-side pipeline and the AI-vendor egress components that reached general availability. The few remaining forward-looking items are tagged **[Planned]**. Where a subsection mixes both, the individual paragraphs are tagged.

Inbound — pseudonymization at the edge

DICOM path is shipping today. Your modality, EMR, or PACS sends DICOM over your existing integration. The on-prem Synthology Router intercepts the relevant fields and applies pseudonymization rules per the canonical PHI tag policy (mirrors PS3.15 §E "Basic Application-Level Confidentiality Profile" with reversibility added). Tag-by-tag actions:

- **Pseudonymize** — replace value with HMAC-SHA256(customer_key, value) truncated to a `PSN-{12hex}` token. Used for PatientName, PatientID, AccessionNumber, ReferringPhysiciansName, InstitutionName, and other free-form PHI strings.
- **Date shift** — shift dates (DA / DT / TM tags) by a per-patient deterministic offset derived from HMAC of the patient identifier. Stable across all dates for the same patient so temporal relationships are preserved within a patient timeline.
- **Keep** — pass through demographically interesting fields (PatientSex) and structurally required indexing fields (Modality).
- **Zap** — delete tags that should never leave the workstation under any circumstance (operator notes, free-text comments).

The reversal map (the bidirectional table that lets the on-prem Router undo the transformation) is encrypted with your customer keypair using AES-256-GCM and embedded in a DICOM private group on the egressing study. The map travels WITH the pseudonymized study; the cloud and the AI vendor receive the ciphertext only.

HL7 segment pseudonymization (PID, MSH, OBR, OBX identifier fields) is generally available.

FHIR resource pseudonymization (Patient.name, Patient.identifier, Practitioner.name, and corresponding identifier slots in encounters, observations, and reports) is generally available.

Free-text PHI detection (patient names, MRNs, dates of birth, addresses, and phone numbers detected inside narrative text such as radiology impressions, clinical notes, and prior-auth letters) is generally available. The free-text pipeline combines regex matching with healthcare-tuned named-entity recognition. Precision/recall benchmarks against published gold-standard healthcare corpora are still being prepared for publication, and will be reproducible against customer evaluation datasets.

Pseudonyms are **deterministic in context**: the same patient identifier in two requests produces the same pseudonym, so the AI vendor sees consistent references across conversations. Pseudonyms are **per-customer salted**: a pseudonym derived from your data cannot be linked to a pseudonym derived from another Synthology customer's data, because the customer keypair is the HMAC key and is held by you alone.

Transit — Synthology cloud under BAA

The cloud routing tier is generally available. Pseudonymized payloads flow to the Synthology cloud routing tier over a Cloudflare-fronted TLS connection. The cloud tier is designed as a thin orchestration layer with three

responsibilities:

1. Receive the pseudonymized request from your on-prem Router and identify which vendor it is destined for (based on the vendor allowlist your administrator configures).
2. Apply the appropriate per-vendor request transformer so the request matches the vendor's API contract — whether that's an imaging-AI vendor expecting a DICOM study, a clinical-NLP vendor expecting an HL7 message, a decision-support engine expecting a coded encounter, or a foundational-LLM platform expecting a chat-style payload.
3. Forward the request, capture the response, and return it to your on-prem Router.

By design, **the cloud tier will hold no copy of your customer keypair**, no copy of the pseudonymization map, and no plaintext PHI. AI-vendor API keys themselves will live in customer-managed escrow and be passed through encrypted-at-rest with your keypair; the cloud tier will hold them as ciphertext blobs only.

Synthology will operate the cloud routing tier under a Business Associate Agreement with you. We make a precise contractual claim for the GA architecture: **our BAA will cover transit and storage of pseudonymized data only; once the cloud routing tier deploys, we will be contractually and technically incapable of producing plaintext PHI from cloud-side state.** That is the architectural difference from AWS HealthLake or Azure Health Data Services, both of which hold PHI plaintext-readable inside their cloud perimeters under their BAAs.

Outbound — AI-vendor calls

The vendor-egress flow is generally available. The cloud tier calls the AI vendor over the vendor's standard API surface. Pseudonymized prompt arrives; pseudonymized response returns. Capabilities the architecture preserves:

- **No-training-data flags.** Where the vendor offers an enterprise setting that excludes your data from their training pipeline (organization-level toggles, enterprise tiers, contractual exclusions, etc.), Synthology's per-vendor configuration will document the setting your administrator must apply on the vendor side. Synthology does not — and cannot — verify the vendor's compliance with their own setting; that's a vendor-side trust decision your CISO makes with the vendor.
- **Streaming responses.** For vendors that support streaming, the cloud tier is designed to stream chunks back to your on-prem Router, which will re-identify and stream to your workstation. End-to-end latency will be preserved.
- **Tool / function calling.** Vendor function-calling features are being designed to work with symmetrically-applied pseudonymization on tool inputs and outputs.

Vendors that won't sign a downstream BAA-style instrument with Synthology will **not** be routable. The vendor allowlist will be the regulatory boundary, gated by a SynthQMS administrative permission (final permission name to be determined before GA), and every change will be audit-logged.

Inbound — re-identification at the edge

DICOM-side re-identification is shipping today. Pseudonymized studies that arrive back at the on-prem Router are decrypted using your keypair, the embedded PHI map is consulted, and real patient names, MRNs, dates, and identifiers are substituted back. Plaintext is delivered to the originating workstation, EMR, or PACS.

Every re-identification event is recorded in the audit log: who requested the AI workflow, which vendor was called, what feature flag triggered it, what the request and response hashes were, and what the latency and token-cost were. The audit log is hash-chained for tamper evidence.

Who this is for

Chief Information Security Officers evaluating whether to allow third-party AI medical vendors inside their hospital network. The pseudonymization architecture, at GA, will let you say "yes" to a clinical workflow your radiologists or clinicians are asking for, while honestly answering an audit committee that no PHI ever reaches the AI vendor's servers.

Chief Medical Information Officers who need their clinical-documentation-improvement, coding-assist, decision-support, and imaging-AI tools sourced from the best vendor for the workflow — not just the BAA-eligible subset offered by their primary cloud vendor.

RIS and PACS architects who want to send DICOM Structured Reports to specialty-grading services or research-grade ML models without giving up control of patient identity. The DICOM-side pipeline is shipping today.

Compliance and privacy officers who need a per-request audit trail showing what was sent to which vendor, when, by whom, and what came back. The audit trail is hash-chained at the on-prem Router; tamper detection is built in and shipping today.

Health-system CIOs evaluating multi-vendor AI strategy without committing to a single cloud's roster, and who need procurement leverage when vendor pricing or capability changes shift.

Security and compliance posture

Threat model

The architecture is designed to defeat a specific set of threats. Those tagged **[GA]** apply to the pipeline in production today, including the generally-available cloud routing tier.

- **Cloud-side insider threat.** A Synthology employee with full operational access to the cloud routing tier obtains, at most, pseudonymized payloads and ciphertext envelopes. Without your keypair — which never leaves your environment — they cannot decrypt or re-identify anything.
- **AI-vendor breach.** If an AI vendor suffers a breach exposing the requests or imaging studies they have processed for you, the attacker obtains pseudonymized data. Without your re-identification map (held in your environment, not theirs), the attacker cannot link pseudonyms to real patients.
- **Stolen backup tape from Synthology cloud.** Same outcome as the cloud-side insider threat: pseudonymized payloads + ciphertext envelopes are not recoverable PHI.
- **AI-vendor logs requests for training.** Where vendors offer no-training-data settings (organizational toggles, enterprise tiers, contractual exclusions, etc.), customer documentation specifies the required configuration. Where vendors are unwilling to honor opt-out, Synthology does not allow them on the vendor allowlist.

- **Cross-customer pseudonym linkage** (one customer's pseudonyms vs another's). The customer keypair is the HMAC key for pseudonym generation, so pseudonyms from different customers are cryptographically independent. This is true today for all DICOM workflows.
- **Unauthorized vendor enabling.** The vendor allowlist is gated by an administrative permission. Every change is audit-logged. An ordinary user — including a radiologist, technologist, or even a non-administrator IT operator — cannot route data to a new vendor on their own.

What this architecture does **not** protect against

We are explicit about the boundaries:

- **A live root attacker on your on-prem Router.** If an adversary has root on the host running the on-prem Router, they have the keypair and the pseudonymization map. This is true of any encryption architecture; the architectural job ends at the boundary of your administered network.
- **A human acting maliciously inside your network.** The architecture cannot defeat a credentialed insider with appropriate workflow access. Audit logs make malicious activity discoverable; they do not prevent it.
- **Vendor-side BAA non-compliance.** If a vendor agrees to no-training-data settings and violates the agreement, that's a vendor-side breach of contract — Synthology can document it for you and exclude the vendor from your allowlist going forward, but cannot guarantee vendor-side behavior.
- **Inversion of weakly-anonymized fields.** Free-text PHI detection has a precision/recall envelope. The pipeline provides a customer-side allowlist for known-non-PHI tokens that look like PHI; reproducible benchmarks against gold-standard healthcare corpora are still being prepared for publication. We do not claim mathematical perfection on free-text detection.

HIPAA posture

Pseudonymized data with a re-identification key in customer custody is **still PHI under HIPAA**. We will operate the cloud routing tier under a Business Associate Agreement once it deploys. The architectural distinction from AWS HealthLake or Azure HDS is not that we exit HIPAA's scope — we won't — but that **at GA, our BAA will cover transit and storage of pseudonymized data only**, and the cloud tier will be technically incapable of producing plaintext PHI from cloud-side state.

AI vendors will be subcontractors in the BAA chain. Synthology will document the chain to your compliance officer and refuse to route to vendors who won't sign downstream BAA-style instruments.

Audit trail

Shipping today. Every re-identification event in the DICOM-side pipeline generates an audit-log row containing the timestamp (UTC, hash-chain integrity), the user and session that initiated the request, the workflow context, the request and response hashes (SHA-256 of pseudonymized payloads), and operational metadata.

The audit log is hash-chained at the on-prem Router using SHA-256 over a canonical serialization of each row. Any subsequent modification of a row is detectable on a read because the recomputed hash will not match the stored hash. Synthology provides a **POST /api/audit-log/verify** endpoint that walks the entire chain and returns one of three statuses: chain intact, chain broken at a specific row, or no chain (pre-deployment row). This is how you demonstrate audit-trail integrity to a SOC 2 or 21 CFR Part 11 reviewer.

Every AI-vendor request is additionally audit-logged with the vendor identity, the feature flag that triggered the AI workflow, the token count, the estimated cost, and the latency.

Vendor support

Vendor categories planned for support

The architecture is being designed to treat every AI vendor as an HTTPS-accessible inference endpoint, agnostic to what the vendor does with the request. That means the same egress tier, at GA, will support the full range of medical AI vendors a healthcare organization might want to use:

- **Imaging AI** — radiology AI for chest X-ray, CT, MR, and ultrasound; pathology AI for whole-slide image analysis; cardiology AI for echo, cardiac MR, and ECG; dermatology, ophthalmology, and other imaging specialties. Pseudonymization runs at DICOM-tag granularity. DICOM-side pseudonymization is ready today.
- **Clinical NLP** — clinical-documentation improvement, computer-assisted coding, structured-data extraction from free text, prior-authorization narrative generation, scribe assistance. Pseudonymization handles HL7 / FHIR identifier fields plus free-text PHI inside narrative content. Generally available.
- **Decision support** — risk stratification, protocol selection, predictive analytics for sepsis or deterioration, drug-interaction checking, evidence-based guideline matching. Generally available.
- **Foundational LLMs** — general-purpose large language models for narrative summarization, query answering, and unstructured workflow tasks. Streaming responses, function/tool calling, and chat history are preserved symmetrically across the pseudonymization boundary. Generally available.
- **Specialty workflow AI** — peer review, quality measures, prior authorization, claim adjudication assistance. Generally available.
- **On-prem inference** — for organizations that prefer air-gapped end-to-end, the same architecture supports local inference servers. Pseudonymization isn't strictly required when the model never leaves your network, but the consistent operational pattern (audit, policy allowlist, key custody) applies anyway.

Active vendor program

Synthology is in active conversations with AI vendors across all of the categories above. Our published vendor list will grow as we confirm, for each vendor, three things:

1. The vendor's API surface is stable enough to support a long-lived integration.
2. The vendor is willing to operate with pseudonymized inputs — some inference pipelines have downstream dependencies that don't tolerate token-substitution well, and we want to know that before a customer learns it the hard way.
3. The vendor will sign a downstream BAA-style instrument with Synthology, OR the architectural posture is acceptable to your compliance officer without one (vendor-by-vendor).

We're intentionally not publishing a roadmap of named vendors yet. Vendor-side BAA posture, no-training-data settings, and de-identification readiness are moving targets that change quarter by quarter. The list we publish will be the list our customers can actually run, validated end-to-end.

Adding a new vendor is designed to be a small engineering effort. Once a vendor's API surface and BAA posture are confirmed, the per-vendor adapter will typically be a few days' work. If your organization has a preferred vendor in mind, tell us — customer-driven prioritization is how the list grows.

How this compares

| Capability | AWS HealthLake | Azure Health Data Services | 3verest | Synthology |

|---|---|---|---|---|

| Customer holds the decryption key | No (AWS-managed KMS by default) | Optional (customer-managed key) | Customer-controlled (healthcare host) | **Always** (per-customer keypair, shipping) |

| Cloud provider can read plaintext PHI | Yes (under their BAA) | Yes (under their BAA) | No (hosting platform; not a tenant of your data) | **No, by design** (architecturally incapable without your key; cloud routing tier GA) |

| Use AI vendors that don't sign your BAA | Bedrock only (in-cloud) | Azure OpenAI only (in-cloud) | Open (you deploy any compatible AI integration) | **Any** (generally available) |

| Vendor lock-in | High (AWS-only) | High (Azure-only) | None (open hosting) | **None** (pluggable destinations) |

| ISO 13485-aligned firm | No | No | No (infrastructure host) | **Yes** (ISO 13485-aligned; non-device software under FD&C Act §520(o)(1)(D)) |

| DICOM-native + HL7 / FHIR / X12 | FHIR-only | FHIR + DICOMweb | Customer-deployed stack | DICOM; HL7 / FHIR |

AWS HealthLake and Azure Health Data Services are managed AI/data services. 3verest is a healthcare-specialized hosting platform — the column reflects deployment flexibility rather than a turnkey managed service. Synthology supports deployment on customer-chosen infrastructure including 3verest as a hosting partner.

The architectural difference is structural, not contractual. AWS and Azure can — and do — improve their privacy postures over time, but the model where the cloud holds the keys is fundamentally different from the model where the customer does.

Frequently asked questions

Is pseudonymized data still PHI?

Under HIPAA, yes. A pseudonymized dataset with a re-identification key held by the original custodian is still PHI; only Safe Harbor or Expert Determination de-identification (irreversible) removes the data from HIPAA's scope. The architectural difference Synthology offers is not that we exit HIPAA — at GA we will operate the cloud tier under a BAA — but that the cloud tier will be technically incapable of producing plaintext PHI inside our cloud, regardless of what the BAA permits.

If your use case requires fully de-identified data (research datasets, public-health reporting, etc.), Synthology's Research VNA tier offers irreversible anonymization — that's a separate product, with a different architectural posture, and addresses a different problem.

What happens if Synthology's cloud is compromised (post-GA)?

The attacker obtains pseudonymized payloads and ciphertext envelopes. Without your keypair, no re-identification is possible. We don't claim the cloud is unbreachable; we claim that a breach of our cloud cannot expose plaintext PHI, because plaintext PHI is never in our cloud.

What happens if I lose my keypair?

You lose the ability to re-identify any future vendor responses. Existing data in your environment remains intact (the re-identification map is held by you, not Synthology). We strongly recommend regular key-rotation procedures and offline cold-storage of recovery material; customer-side HSM / key-vault custody is on the roadmap. Detailed key-management guidance is part of customer onboarding.

Can I use multiple AI vendors at the same time (post-GA)?

Yes. Per-workflow vendor selection is the design model: radiology AI on one vendor, prior-auth drafting on another, coding-assist on a third — all simultaneously. The routing will be configured per workflow in the SynthQMS administrative UI, and a workflow will be re-routable to a different vendor without code changes.

What's the latency overhead from pseudonymization?

DICOM-tag pseudonymization itself is sub-millisecond per study (HMAC-SHA256 + AES-256-GCM are fast on modern CPUs). The cloud-tier round-trip and free-text PHI detection are the dominant added components. End-to-end **egress** latency — the added round-trip through the cloud routing tier to an AI vendor and back — has not yet been measured and published; the design target is <200 ms p99 for typical radiology-report-sized payloads (~5 KB). We quote measured results rather than targets, so that figure appears here only once the run is done. Router and SynthIQ *throughput* are separately measured and published: see WP-2026-002 (SynthIQ capacity on 2 vCPU / 8 GB) and WP-2026-003 (Router + SynthIQ at 4 vCPU / 16 GB, 4 vCPU / 32 GB and 8 vCPU / 64 GB).

Does Synthology see my prompts and responses (post-GA)?

The cloud routing tier will see pseudonymized prompts and pseudonymized responses — as ciphertext envelopes that we route, log, and observe for operational metrics. We will not see plaintext PHI. We will have neither the keypair nor the pseudonymization map; we will not be able to produce plaintext even if compelled to. This is by architectural design, and we'd rather not have to.

How do I evaluate this in my environment?

Synthology offers design-partner engagements: a small set of health-system customers will run the architecture in production with limited workflows, real-vendor calls, and full instrumentation. Design partners receive engineering and clinical-workflow support throughout the engagement and pricing aligned with the program's evaluation goals. Contact Synthology to discuss whether your organization is a fit.

What does pricing look like at general availability?

Pricing aligned with vendor usage patterns — the dominant cost component for most customers is the underlying AI-vendor cost, not the Synthology layer. We will publish standard pricing at general availability; the design-partner program operates on a different commercial model. Reach out for specifics.

What to do next

If you're evaluating an AI workflow that's currently blocked by a vendor's lack of BAA, or you've taken a single-cloud lock-in compromise that you'd like to undo:

- **Talk to Synthology.** Request a 30-minute call with our solutions engineering team. We'll walk through your specific workflow and the pseudonymization rules that would apply at GA. We can usually tell within that call whether the architecture is a fit for your problem or whether something else (a single-cloud BAA AI menu, or a fully de-identified research model) is better matched.

- **Request the design-partner package.** If the architecture is a fit and your clinical workflow is concrete, the design-partner program is the fastest path to running it in production.
- **Read the public materials.** Our marketing site at `synthologysolutions.net/pseudonymization` has the architecture diagram, comparison content, and links to product pages for the on-prem Router and the VNA family.

Roadmap reference (what ships when)

Component	Status	Target
Per-customer keypair generation, escrow, distribution	Shipping	—
HMAC-SHA256 deterministic pseudonyms with per-customer salting	Shipping	—
AES-256-GCM PHI envelope embedded in DICOM private group	Shipping	—
Reversible DICOM-tag pseudonymization (PS3.15 §E + reversibility)	Shipping	—
Per-patient deterministic date shift	Shipping	—
Hash-chained audit log + verify endpoint	Shipping	—
On-prem Router (Windows + Linux)	Shipping	—
Regulatory posture (non-device software under FD&C Act §520(o)(1)(D))	Confirmed	—
HL7 segment pseudonymization (PID, MSH, OBR, OBX)	GA	—
FHIR resource pseudonymization	GA	—
Free-text PHI detection (regex + healthcare-tuned NER) with published benchmarks	GA	—
Cloud routing tier under BAA	GA	—
AI-vendor adapter framework + allowlist	GA	—
Streaming responses + function/tool calling preservation	GA	—
Customer-managed HSM / cloud key-vault providers	Roadmap	Planned
End-to-end egress latency (p99 added)	Roadmap	Measurement pending

---

Synthology Healthcare Solutions Group, LLC builds the integration layer between healthcare modalities, EMRs, PACS, AI vendors, and the cloud — with privacy and regulatory posture as architectural constraints, not afterthoughts. Synthology products are ISO 13485-aligned, general-purpose non-device software under FD&C Act §520(o)(1)(D).

Copyright © 2026 Synthology Healthcare Solutions Group, LLC. All rights reserved. This solution brief contains a small number of forward-looking statements. Forward-looking statements are tagged [Planned] and reflect current development plans; they carry no committed delivery date, and feature sets may change. AWS, AWS HealthLake, Azure, Azure Health Data Services, Azure OpenAI, AWS Bedrock, Google Vertex, OpenAI, 3verest, and other named products are trademarks of their respective owners and are referenced for comparison purposes only; Synthology is not affiliated with or endorsed by any of them. This solution brief reflects publicly documented capabilities and roadmap as of 2026 Q2 and is subject to change as products, vendor partnerships, and regulatory posture evolve.

Regulatory classification. All Synthology and XyDromatics software described here is non-device software under the FD&C Act §520(o)(1)(D) software exclusion — not a medical device, and no clinical diagnostic function. Published by Synthology Healthcare Solutions Group, LLC as a controlled document; the authoritative version is the approved record in the Synthology quality-management system.