

TECH & AI CASE STUDY - HEALTHCARE SAAS - 2026

Cloud Migration That Delivered 99.99% Uptime for Healthcare SaaS

Based on a delivery engagement with CareBridge Health

99.99%

PLATFORM UPTIME

-34%

INFRASTRUCTURE COST

**monthly to
weekly**
RELEASE CADENCE**-96%**

UNPLANNED DOWNTIME

A fillable practice edition: the full story, 8 actions,
6 reflections and a 90-day plan - work it with your team.

Free - no email required
emersontechnology.com

MAKE THIS PRACTICE BOOK YOURS

Work it, do not just read it

This practice book turns the CareBridge Health case study into a working session for your team. Read the full story first - what happened, what was built and what it changed. Then tick the 8 actions you can commit to, write short answers to the 6 reflections, and finish with a 90-day plan. Every field in this PDF is fillable: open it in Adobe Acrobat Reader, Chrome, Edge or macOS Preview, type directly into the boxes and save.

YOUR NAME	ROLE
COMPANY	TEAM / DEPARTMENT
DATE	EMAIL (OPTIONAL)

MY GOAL FOR THIS PRACTICE SESSION

- 01 The story** - What happened at CareBridge Health - before, during and after the project.
- 02 The results** - The numbers that changed, plus industry benchmarks from public research.
- 03 The 8-action checklist** - Tick boxes as your organization completes each step. Owner and date fields included.
- 04 The 6 reflections** - Fillable questions that pressure-test your own roadmap.
- 05 My 90-day plan** - Three fillable phase cards that turn this case study into your roadmap.

RUN IT AS A 90-MINUTE TEAM SESSION

25 min READ Read the story together. Ask everyone to note the one number that surprised them most.	40 min PRACTICE Work the 8-action checklist and the 6 reflections. Assign an owner and a date to every action you commit to.	25 min COMMIT Fill in the 90-day plan. One goal per phase, one number that tells you it worked. Book a review date before you leave the room.
--	--	---

THE STORY

What happened, and why it worked

01 The Company

CareBridge Health is a 90-employee healthcare SaaS company whose care-coordination platform serves 120 clinics. The platform sits at the center of its customers' daily operations: patient scheduling, secure messaging between care teams, telehealth hand-offs and billing all flow through it, which means every clinic receptionist, nurse coordinator and practice manager who opens a browser is touching CareBridge's infrastructure. The company sells to small and mid-sized practices on annual contracts with availability commitments written into the agreements, and its product roadmap depends on shipping clinical workflow improvements faster than the practice-management incumbents it competes against.

Because the platform handles protected health information end to end, CareBridge lives under HIPAA obligations that shape every infrastructure decision: encryption at rest, access logging, audit trails and the ability to produce evidence on demand when a risk assessment or a customer security review arrives. In healthcare, the stakes of getting this wrong are unusually high. IBM's Cost of a Data Breach report has repeatedly found healthcare carries the industry's highest average breach costs — reportedly exceeding \$10 million per incident in recent years. Reliability failures are cheaper than breaches, but as the team would soon learn, they are not cheap.

02 Where Things Stood

When Emerson Technology first talked with CareBridge's leadership, the entire platform ran on two racks of aging servers in a leased data center. The hardware was past end-of-life on several critical components, power feeds had no redundancy, and the network uplink was a single path whose failure meant a total outage. The operations burden fell on a small infrastructure group who spent their weeks applying patches by hand, babysitting storage arrays and negotiating maintenance windows with the business. Quarterly patching cycles forced weekend downtime that clinicians had to plan around, and engineers privately dreaded each one because manual runbooks meant the outcome was never fully certain until Monday morning.

The defining incident came in March, when a single disk failure took patient scheduling offline for nine hours during a business day. The consequences cascaded exactly as you would expect: SLA credits were issued, the three largest clinic customers raised churn threats, and the incident triggered a formal HIPAA risk assessment finding that leadership now had to remediate on a timeline. Beyond that acute failure, the chronic symptoms were wearing the team down. Flu-season traffic spikes — predictable every winter — regularly degraded response times because capacity was fixed at what the racks could hold. A vendor quote to refresh the data center arrived at \$1.4M in capex, with no meaningful improvement in reliability, no improvement in release speed and the same single-site risk profile as before.

03 The Challenge

The challenge was not "move to the cloud" as a slogan; it was a three-sided trap. On one side, reliability: the platform's availability posture was one hardware fault away from another nine-hour outage, and Uptime Institute's Annual Outage Analysis reports that a majority of operators say their most recent significant outage cost more than \$100,000, with a meaningful share exceeding \$1 million — figures that match CareBridge's own March experience once credits and crisis labor were tallied. On the second side, compliance: the HIPAA risk assessment finding meant every future architecture decision needed documented controls, not best-effort intentions. On the third side, cost: a \$1.4M refresh would lock five more years of capex into infrastructure that still could not scale for flu season or ship a release faster than monthly.

There was also a sequencing problem that made the situation genuinely difficult. CareBridge could not pause the product while a migration happened; clinics do not tolerate a frozen scheduling platform, and the sales team was actively in renewal conversations with the very customers who had issued churn threats. The migration therefore had to happen underneath a live, revenue-critical, regulated system — one service at a time, with rollback options at every step. Leadership framed the decision bluntly: either commit to a cloud exit that removed the single-site risk and the capex cycle, or spend \$1.4M to stand still.

04 Discovery and Diagnosis

Emerson Technology opened with a two-week discovery sprint across CareBridge's systems, dependencies and operations rituals. The team inventoried every service — scheduling, messaging, telehealth hand-offs, billing and the supporting batch jobs — and mapped how data moved between them, tracing which jobs read from which databases and which nightly processes silently depended on each other. Three findings shaped everything that followed. First, several components documented as independent were in fact coupled through shared file shares and a cron scheduler nobody had touched in years. Second, capacity had been sized for average load, so winter spikes had nowhere to go but response-time graphs. Third, compliance evidence existed mostly as tribal knowledge inside two senior engineers rather than as retrievable artifacts.

The diagnosis reframed the vendor's \$1.4M quote as the wrong baseline. McKinsey research reportedly estimates that migrating on-prem workloads to the cloud can reduce total cost of ownership by roughly 30-40% for well-optimized environments — and CareBridge's environment, with its low utilization and manual operations, was the kind of environment that optimization helps most. The discovery sprint closed with a wave plan: service domains grouped by data dependency and blast radius, ordered so that each wave could be rehearsed, executed and rolled back independently, with the patient-facing scheduling core deliberately scheduled last when the team's cloud operating muscle was proven.

05 The Solution

Emerson Technology executed a strangler-pattern migration to AWS over seven months, moving each service domain in planned waves while the platform stayed live throughout. The destination architecture was designed for three properties: no single point of failure, no manually provisioned components, and compliance controls that existed as code from day one.

The landing zone

Everything in the new environment is defined in Terraform across three Availability Zones. Networks, subnets, security groups and IAM roles are reusable modules that get reviewed as code rather than configured by hand. EC2 Auto Scaling groups sit behind Application Load Balancers, so capacity rises and falls with demand instead of being bolted to a rack. Amazon RDS for PostgreSQL runs with a Multi-AZ standby, giving database failover that the old single-array setup could never offer, and encrypted EBS volumes plus S3 lifecycle policies handle clinical document storage with predictable, tiered cost. Because the whole environment is code, a failed region-level experiment can be torn down and rebuilt in an afternoon — a property the old data center never had.

The migration waves

Each service domain moved in its own wave: dependency mapping first, a staging rehearsal second, a monitored cutover third, and a documented rollback path throughout. GitHub Actions pipelines run automated test suites and rehearsed database migrations before every promotion to production, which is what turned releases from a monthly event into a weekly routine. Load tests modeled 5x flu-season peaks, deliberately harsher than any winter the platform had seen, so the autoscaling behavior was proven before the patients depended on it. The final workloads moved during a monitored cutover weekend run from a shared checklist, with dashboards for latency,

error rates and sync lag visible to everyone in the room — and it closed with zero unplanned downtime.

The compliance backbone

PHI is protected end to end by design rather than by policy documents. KMS-managed encryption covers every volume and bucket at rest, TLS 1.3 protects data in transit, application tiers live in private subnets with no direct internet path, and IAM roles grant least-privilege access per service. CloudTrail audit logging captures every API call, and each control is mapped directly to a named HIPAA requirement in a living evidence index. That mapping is the quiet hero of the project: when the formal risk assessment was re-run, and later when CareBridge pursued SOC 2 Type II, the auditors drew from the same artifacts the engineers already produced, so audits became a review exercise instead of an archaeology project.

06 The First 90 Days

The first six weeks belonged to foundations. The landing zone was built and peer-reviewed as Terraform, the CI/CD pipelines were stood up with automated tests and migration rehearsals, and the team ran its first full staging rehearsal of a migration wave — which promptly surfaced a hidden dependency between messaging attachments and a legacy file share that no design review had caught. Catching it in staging, with a rollback available and no customers watching, was exactly what the rehearsal discipline exists for. By week six, the messaging domain was live on AWS, serving real clinic traffic, with the on-prem path still warm as a fallback. Weeks seven through ten moved billing and telehealth hand-offs, and with them came the operational habits that make a platform auditable: structured logging forwarded centrally, CloudTrail mapped into the HIPAA evidence index, and SLA instrumentation — availability targets, alert thresholds and error budgets — defined and reviewed before the spend grew large enough to blur them. Reserved-capacity purchases were deliberately deferred until roughly the 60-day mark, after utilization data showed the real steady-state footprint, so commitments were bought against observed usage rather than guessed sizing. Weeks eleven and twelve carried the last pre-scheduling wave and a full 5x load-test rehearsal of the scheduling service, including a timed cutover drill that the team improved twice before the real weekend. By day 90, seven months of work had a proven rhythm, and only the final cutover lay ahead. It went according to the checklist — the highest compliment a migration can be paid.

07 The Results

In the twelve months since the final cutover, CareBridge has sustained 99.99% availability while absorbing two full flu-season peaks that previously degraded response times every winter. Monthly infrastructure spend fell 34% against the old data center, and the \$1.4M hardware refresh simply evaporated as a line item. Weekly releases replaced the monthly cycle, unplanned downtime fell roughly 96% year over year, and the HIPAA risk assessment was closed with no critical findings. CareBridge then passed a SOC 2 Type II audit drawing on the same evidence pipeline the engineers already maintained — no audit-specific scramble, no documentary heroics.

METRIC	BEFORE	AFTER
Platform uptime	-	99.99%
Infrastructure cost	On-prem baseline	-34%
Release cadence	Monthly	Weekly
Unplanned downtime	Baseline	-96%

The operational texture of the company changed along with the numbers. Release weekends disappeared from engineers' calendars, replaced by routine weekly promotions that a single on-call engineer can supervise.

Capacity conversations shifted from "what can the racks hold" to "what does this week's forecast need," and the answer arrives by changing an autoscaling parameter instead of issuing a purchase order. Sales conversations changed too: availability commitments that once required caveats are now made from evidence, because the monitoring that proves 99.99% is the same monitoring the team watches daily.

08 What Your Business Should Do

The sequence CareBridge followed is repeatable, and its ordering matters more than its tooling. Start by mapping workloads into migration waves grouped by data dependency and blast radius rather than by department, because wave planning is what lets you rehearse each move, catch coupling problems early and keep rollback options open. Write Terraform before you migrate anything: codify networks, IAM, databases and storage as reusable modules, since reviewing infrastructure-as-code is faster and safer than reviewing screenshots, and it makes compliance evidence collection dramatically easier. Encrypt regulated data with managed keys from the start — KMS at rest, TLS in transit, least-privilege IAM — and map every control to a named requirement so audits stop being archaeology projects. Then rehearse the cutover in staging with production-like load tests and timed drills; teams that rehearse discover their real dependencies weeks earlier than runbook reviews ever will. Handle the economics with the same discipline. Automate releases before you move, because migrating a manually deployed system recreates every bad habit you had on-prem, just on more expensive infrastructure. Wait 60 to 90 days after cutover to observe real utilization, and only then commit to reserved capacity — buying commitments against guessed sizing locks in the wrong footprint for years. Instrument SLAs before your first cloud invoice: define availability targets, alert thresholds and error budgets up front and review them quarterly, because a 99.99% commitment you cannot measure is a marketing claim rather than an operations policy. Finally, close the audit loop once: build a single evidence pipeline of logs, access reviews and change records that serves HIPAA, SOC 2 and customer security reviews together, so every future audit reuses the same artifacts. As you plan, sit with the questions this engagement raised — what a nine-hour outage would actually cost you, how many of your controls exist as evidence versus tribal knowledge, which workload you would move first if you could only rehearse one wave this quarter, and what standing still truly costs. Work through these eight steps with your team using the action checklist in this practice book, and revisit the reflection questions each quarter to keep the migration honest long after the last server is switched off.

REAL-WORLD PROOF & BENCHMARKS

IBM's Cost of a Data Breach report has repeatedly found healthcare carries the industry's highest average breach costs — reportedly exceeding \$10 million per incident in recent years. - IBM, 2023

Uptime Institute's Annual Outage Analysis reports that a majority of operators say their most recent significant outage cost more than \$100,000, with a meaningful share exceeding \$1 million. - Uptime Institute, 2022

McKinsey research reportedly estimates that migrating on-prem workloads to the cloud can reduce total cost of ownership by roughly 30-40% for well-optimized environments. - McKinsey & Company

THE PRACTICE PAGES

The 8-action checklist

Work through the list with your team. Tick each box as your organization completes the step, and fill in an owner and a target date so every action has a name and a deadline.

☐

1. Map workloads into migration waves

Group applications by data dependencies and blast radius rather than by department. Wave planning lets you rehearse each move, catch coupling problems early and keep rollback options open.

OWNER

TARGET DATE

STATUS

--	--	--

☐

2. Terraform everything before you migrate

Codify networks, IAM, databases and storage as reusable modules. Reviewing infrastructure-as-code is faster and safer than reviewing screenshots, and it makes HIPAA evidence collection dramatically easier.

OWNER

TARGET DATE

STATUS

--	--	--

☐

3. Encrypt PHI with managed keys

Use KMS or equivalent for encryption at rest, TLS in transit, and least-privilege IAM. Map every control to a named compliance requirement so audits stop being archaeology projects.

OWNER

TARGET DATE

STATUS

--	--	--

☐

4. Rehearse the cutover in staging

Run full production-like load tests and timed cutover drills before migration weekend. Teams that rehearse cutovers discover their real dependencies weeks earlier than runbook reviews ever will.

OWNER

TARGET DATE

STATUS

--	--	--

☐

5. Buy reserved capacity after baselining

Wait 60-90 days to observe real utilization, then commit to reserved instances or savings plans. Committing before you understand usage locks in the wrong footprint for years.

OWNER

TARGET DATE

STATUS

--	--	--

☐

6. Automate releases before the move

Stand up CI/CD with automated tests and rehearsed database migrations first. Migrating manually deployed systems recreates every bad habit you had on-prem, just on more expensive infrastructure.

OWNER

TARGET DATE

STATUS

☐

7. Instrument SLAs before your first invoice

Define availability targets, alert thresholds and error budgets up front, and review them quarterly. A 99.99% commitment you cannot measure is a marketing claim, not an operations policy.

OWNER

TARGET DATE

STATUS

☐

8. Close the audit loop once

Build one evidence pipeline — logging, access reviews, change records — that serves HIPAA, SOC 2 and customer security reviews together, so every future audit reuses the same artifacts.

OWNER

TARGET DATE

STATUS

AFTER THE CHECKLIST: OUR TOP 3 PRIORITIES

From the eight actions above, write down the three you will start first - and the single number you expect each one to move. Three priorities with owners beat eight without any.

PRIORITY 1	PRIORITY 2	PRIORITY 3
NUMBER IT SHOULD MOVE	NUMBER IT SHOULD MOVE	NUMBER IT SHOULD MOVE

THE PRACTICE PAGES

Reflections for your team

Use these questions in your next planning session and type short answers into the boxes. The gaps you find are usually the fastest wins.

Q1. What would a nine-hour outage tonight actually cost you in SLA credits, churned accounts and sales-team credibility?

Q2. How many of your compliance controls exist as evidence today versus as tribal knowledge inside two senior engineers?

Q3. Which workloads would you move first if you could only rehearse one migration wave this quarter?

Q4. Is your infrastructure documented in code that survives employee turnover, or in diagrams that already have?

Q5. What percentage of your release process would still work if the person who built it left tomorrow?

Q6. If hardware refresh quotes keep rising, what is the real annual cost of standing still?

THE PRACTICE PAGES

My 90-day plan

Finish the session by committing to a plan. Three phases, one goal each, one number that tells you it worked. Bring this page to your next leadership review - or to a free consultation with Emerson Technology, and we will pressure-test it with you.

Days 1-30 Focus: audit and alignment

PRIMARY GOAL

KEY ACTIONS

HOW WE MEASURE SUCCESS

OWNER

Days 31-60 Focus: build and prove

PRIMARY GOAL

KEY ACTIONS

HOW WE MEASURE SUCCESS

OWNER

Days 61-90 Focus: scale and measure

PRIMARY GOAL

KEY ACTIONS

HOW WE MEASURE SUCCESS

OWNER

Tip: save a copy per team so each group keeps its own answers. Revisit the reflections every quarter - the questions stay the same even as the numbers move.

BOOK THE REVIEW BEFORE YOU CLOSE THE BOOK

Plans without a review date drift. Before you close this practice book, agree on when you will come back to these pages, who will run that session, and when you will check the numbers again. If you want an outside pair of eyes, bring this page to a free consultation with Emerson Technology and we will pressure-test the plan with you.

REVIEW DATE	REVIEWED BY	NEXT CHECK-IN
------------------------------------	------------------------------------	--------------------------------------

Ready for a transformation like CareBridge Health's?

Emerson Technology designs and builds the data, AI, cloud and automation systems behind results like these. On a free consultation call we will map this exact playbook to your business - your data, your constraints, your numbers.

Book a free consultation

support@emersontechnology.com

or start online: emersontechnology.com/contact