gart.

Compliance-by-Design in HealthTech

Why Shifting Left on HIPAA, GDPR & NIS2 is the Ultimate Growth Lever for HealthTech Startups

Gart Solutions

AGENDA

1
The Illusion of the Fast Launch
Why skipping compliance guarantees legal failure at scale

2
Blueprinting Compliant Infrastructure
VPC isolation, PHI encryption, Zero-Trust by construction

3
Continuous Compliance in CI/CD
Automated security gates — compliance as code

4
AI-Ready Data Foundations
Provenance, reproducibility, and bias mitigation

5
Zero-Trust Healthcare AI Architecture
EHDS alignment, GDPR, NIS2 & EU AI Act readiness

Why Postponing Compliance Guarantees Legal Failure

The Trap

- Raw prototype built on low-code (Lovable, etc.) feels like progress
- PHI scattered across unencrypted DB columns & S3 buckets
- No network isolation — dev and prod share a flat VPC
- Zero audit trail for data access or processing

The Reality

- Enterprise legal review triggers a full architecture rewrite
- US hospitals & EU medical networks block deployment
- Expensive re-platforming can cost 3–10× the original build
- Compliance debt compounds — the later you fix it, the costlier

"Compliance isn't a blocker for teams that build it into their architecture. It only becomes a wall when bolted on afterward."

The True Cost of Compliance Debt

3–10×

Cost multiplier

Retrofitting compliance costs 3 to 10 times more than building it in from day one

6–18

months delayed

Average enterprise sales cycle extension when legal review exposes architecture gaps

$1.5M

average HIPAA fine

Average penalty per HIPAA violation category (OCR enforcement data, 2023–2025)

100%

of hospital CIOs

Will conduct a security review before signing any vendor contract. Zero exceptions.

"Compliance debt compounds — every sprint you run without it, the rewrite gets larger and the contract gets further away."

From VPC Isolation to Cryptographic Safeguards

Infrastructure as Code (Terraform) makes every environment reproducible, auditable, and version-controlled

Network Segregation

- ✓ Private subnets for all PHI workloads
- ✓ VPC peering with least-privilege rules
- ✓ Security Groups as soft firewalls per tier
- ✓ No public IPs on DB or processing nodes

Data Encryption

- ✓ AES-256 encryption at rest (RDS, S3)
- ✓ TLS 1.3 in transit across all services
- ✓ KMS-managed keys with per-service policies
- ✓ CloudFront signed URLs for PHI assets

Zero-Trust Access

- ✓ IAM roles, not long-lived credentials
- ✓ mTLS for service-to-service auth
- ✓ MFA enforced on all human access paths
- ✓ Continuous access verification per request

Layer 1: Network Segregation — Lock PHI Into Private Subnets

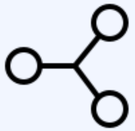

Private Subnets for All PHI Workloads

Every service that touches Protected Health Information runs in a private subnet with no internet-facing route. Only a tightly controlled API gateway is public.

VPC Peering with Least-Privilege Rules

When multiple VPCs must communicate (e.g., analytics ↔️ clinical), VPC peering is configured with explicit allow-lists. Default deny everything.

Security Groups as Software Firewalls

Each service tier (web, app, DB) has its own security group. DB nodes allow inbound only from the app tier's security group — not from the internet, not from dev subnets.

No Public IPs on Databases or Processing Nodes

Zero exceptions. RDS, ElastiCache, Kafka — all deployed in private subnets. Access only via VPN or bastion host with MFA and session recording.

Layer 2: Data Encryption — PHI at Rest & In Transit

Encryption AT REST

AES-256 encryption for all RDS instances (PostgreSQL, MySQL) — enable at instance creation, not afterward.

S3 bucket encryption enforced via bucket policy — reject PutObject requests without server-side encryption header.

KMS Customer Managed Keys (CMK) per service — separate key for DB, separate key for S3, separate key for backups.

CloudFront signed URLs for any PHI assets served via CDN — time-limited, user-specific access tokens.

EBS volume encryption for all EC2 instances handling PHI — enforce via IAM policy denying unencrypted volume launches.

Encryption IN TRANSIT

TLS 1.3 enforced across all service-to-service communication — no TLS 1.0 or 1.1 permitted in security groups.

Certificate management via ACM (AWS Certificate Manager) or Let's Encrypt — automated renewal, no manual certificate handling.

Internal microservice traffic encrypted with mTLS — both client and server present certificates; prevents MITM.

Database connections require SSL — enforce via DB parameter group (rds.force_ssl=1 for PostgreSQL).

VPN with WireGuard or AWS Client VPN for human access to private subnets — no raw SSH to production.

Layer 3: Zero-Trust Access — Never Trust, Always Verify

IAM Roles, Not Long-Lived Keys

Eliminate static API keys entirely. Every service assumes an IAM role with the minimum permissions required. Keys rotate automatically. No key in Git = no breach vector.

MFA on All Human Access Paths

Every engineer, every admin, every contractor — MFA enforced at the IAM policy level. A stolen password alone cannot access production PHI.

mTLS for Service-to-Service Authentication

Every microservice presents a certificate when calling another service. No certificate = no response. Implemented with SPIFFE/SPIRE or AWS ACM PCA for automated cert rotation.

Continuous Access Verification + Audit Logs

Access is re-verified per request, not per session. All API calls, DB queries, and S3 operations logged to CloudTrail / CloudWatch with immutable retention. Audit-ready from day one.

Infrastructure as Code: Terraform Makes It Reproducible

Version-Controlled Infrastructure

Every security group rule, every subnet CIDR, every KMS policy — stored in Git. Full history of who changed what and when. Rollback in minutes.

Reproducible Environments

Spin up an identical compliance-certified environment for a new hospital partner in hours, not weeks. The Terraform module is already compliant.

Audit-Ready by Default

Terraform state files and plan outputs provide a clear record of infrastructure configuration at any point in time. Regulators can verify.

Policy as Code (Sentinel/OPA)

Enforce compliance rules at the infrastructure level: 'No S3 bucket without encryption.' 'No security group allowing 0.0.0.0/0 on port 5432.'

 # Compliant VPC Module (Terraform)

```
module "vpc" {
  source = "terraform-aws-modules/vpc/aws"
  enable_dns_hostnames = true
  enable_flow_log      = true

  private_subnets = [
    "10.0.1.0/24",  # PHI workloads
    "10.0.2.0/24",  # Analytics
  ]
  public_subnets = ["10.0.100.0/24"]

  # No public IPs on private subnets
  map_public_ip_on_launch = false
}

resource "aws_s3_bucket_server_side_encryption
_configuration" "phi" {
  rule { apply_server_side_encryption
    _by_default { sse_algorithm = "aws:kms" } }
}
```

Automating Security Gates in Your CI/CD Pipeline

Shift Left

Compliance and security checks happen at commit time — not after deployment. Fail fast, fix cheap.

Automated Enforcement

Jenkins / GitHub Actions / GitLab CI integrations with Checkov, Trivy, OWASP ZAP, and OPA policies.

Failsafe Deployments

Non-compliant code is blocked from reaching production. Perpetual audit-ready state — zero manual audit prep.

Health data is not the new oil — it's shale.

Abundant. Fragmented.
Expensive to refine.
AI models fail not because the model is bad —
but because the data pipelines are broken.

Break Vendor Lock

Architect infrastructure that cleanly separates clinical data from any single legacy EMR vendor.

Data Provenance

Implement structured, machine-checkable provenance and reproducibility in every pipeline stage.

Bias Mitigation

Address real-world failures (Epic sepsis model, Obermeyer et al.) by shifting operational environments into core infrastructure design.

IaC & Reproducibility

Terraform-driven environments — every data environment is version-controlled, auditable, and reproducible on demand.

Real-World AI Failures — And How Infrastructure Prevents Them

Epic Sepsis Prediction Model

Finding

The model had an AUROC of 0.63 in external validation — barely better than chance — despite widespread deployment across US hospitals.

Root Cause:

Training data was not representative of the hospitals deploying it. No data provenance tracking meant nobody could identify when the distribution drift occurred. No production monitoring detected degraded performance.

Infrastructure Fix:

Data provenance pipeline to track data lineage from source → training → model. Automated distribution shift detection in production. Per-cohort performance dashboards reviewed quarterly.

Obermeyer et al. — Racial Bias in Risk Scoring

Finding

A widely-used commercial algorithm was shown to be significantly less likely to refer Black patients for care management, despite similar or worse health status.

Root Cause:

Healthcare cost was used as a proxy for health need. Cost data reflects historical access disparities, not actual health status. No bias audit was conducted before deployment.

Infrastructure Fix:

Mandatory fairness audits with disaggregated performance metrics by demographic subgroup. Feature selection review to identify and exclude biased proxies. Bias monitoring dashboards in production.

Healthcare Data Shouldn't Move to Be Useful

HIPAA | GDPR | NIS2 | EU AI Act | EHDS

1

Localised Analysis

Run AI inference inside the hospital's own network boundary. Raw records never leave the facility — only aggregated, anonymised outputs travel.

3

EHDS-Aligned Services

Service topologies designed to satisfy the European Health Data Space requirements. Secure APIs for cross-border research without raw record transfer.

3

Compliant by Construction

Every service enforces zero-trust: mTLS auth, attribute-based access control, immutable audit logs. Legal reviews for studies are reduced from months to days.

Case Study: Brain.key — Secure MRI & X-Ray AI at Hospital Scale

THE CHALLENGE

AI-powered neuroimaging platform needed to integrate with multiple hospital networks for MRI and X-ray analysis, handling sensitive imaging data under peak clinical loads — without ever moving raw patient images outside hospital network boundaries.

HIPAA-Aligned VPC Design

Isolated VPC per hospital network with dedicated private subnets for imaging workloads. VPC peering with explicit security group rules — no cross-hospital data flow without explicit allow-list.

Zero-Trust Integration Layer

mTLS-authenticated API gateway for hospital system integration. Each hospital's PACS system receives a unique client certificate. Access revocable per-certificate without system downtime.

In-Network AI Inference

AI inference engine deployed inside the hospital's own network boundary. Only analysis results (structured reports) leave the network — never raw MRI data. GDPR and HIPAA compliant by design.

Peak Load Architecture

Auto-scaling GPU inference clusters with SRE-monitored SLOs. Elastic scaling absorbs sudden radiology demand spikes. Achieved 100% uptime during highest-volume hospital deployments.

The Playbook:

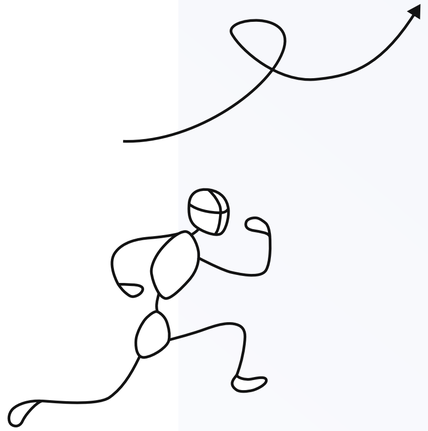

01

360° IT Audit

Identify production-readiness gaps, security holes, and PHI exposure in early-stage code before a single hospital sees it.

02

Architecture Replatform

Replace ad-hoc prototype infrastructure with HIPAA/GDPR-aligned VPCs, private subnets, encryption layers, and access controls.

03

Automated CI/CD Gates

Build security-validated pipelines. Code never touches production without passing compliance, vulnerability, and container scans.

04

Elastic Scale for Medical Institutions

Auto-scaling groups, read replicas, and SRE-monitored SLOs to absorb sudden demand spikes without downtime.

Key Takeaways

- **The 'fast launch' trap is real.** Retrofitting compliance costs 3–10× more and delays enterprise contracts by 6–18 months. Build it in from day one.
- **Compliant infrastructure = 3 layers:** network segregation (private subnets, VPC isolation), data encryption (AES-256, TLS 1.3, KMS), and zero-trust access (IAM roles, mTLS, MFA).
- **Compliance as code is your force multiplier.** Checkov, Trivy, OWASP ZAP, and OPA in CI/CD means every commit is validated against hundreds of controls automatically.
- **AI models fail on broken data pipelines, not bad algorithms.** Build data provenance, bias audits, and production monitoring before you train your first model.
- **For EU markets:** GDPR, NIS2, EU AI Act, and EHDS are not optional. Zero-trust architecture — where data never leaves the hospital — is the compliance pattern that satisfies all four.

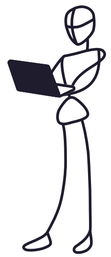

GART is your digital transformation partner

We work to solve your tech challenges on time and on a budget

200x	24x	3x
more frequent deployments	faster recovery from failures	lower change failure rate

Contact Us

Web

gartsolutions.com

Phone

+38 (093) 210 34 71 (UA)

Linkedin

Gart Solutions

Address

Nyzhnokliuchova St, 14
Kyiv, Ukraine