

Policy as Code for Healthcare Security and Compliance

By IntuitionLabs.ai • 8/27/2025 • 50 min read

policy as code

healthcare it

healthcare compliance

security automation

regulatory compliance

devsecops

patient data privacy





Policy as Code in Healthcare

Introduction: Definition and Importance

Policy as Code (PaC) is the practice of expressing organizational policies (security, privacy, compliance, etc.) as machine-readable code crowdstrike.com. In healthcare, this approach is especially critical due to the industry's stringent regulations and high stakes for security and privacy. Unlike manual policies (which can be inconsistent or slow to enforce), Policy as Code allows healthcare organizations to **codify rules and regulations** into automated checks and controls. This means policies (e.g. who can access patient data, how systems must be configured) are defined in code and automatically enforced across systems. According to CrowdStrike, PaC "represents policies and regulations as code to improve and automate policy enforcement and management," reducing the time-consuming, error-prone nature of manual policy administration crowdstrike.com.

Healthcare faces unique challenges that make PaC particularly important. The sector is one of the most regulated, dealing with **hundreds of privacy, security, and compliance standards**. It's also a prime target for cyberattacks due to sensitive patient data (which can fetch up to \$1,000 per record on the dark web) cleardata.com. Data breaches in healthcare are extremely costly, averaging \$9.42 million per incident cleardata.com. Traditional policy management (via written guidelines and periodic audits) struggles to keep up with the **complex, evolving healthcare regulatory environment** cleardata.com. Policy as Code addresses this by **automating compliance and security controls**, ensuring that requirements are consistently enforced in software systems. For example, ClearDATA's healthcare-specific PaC engine translates thousands of lines of privacy regulations ([HIPAA](https://www.hhs.gov/hipaa), GDPR, etc.) into technical controls and cloud configurations cleardata.com. This enables *continuous compliance*: policies are checked and remediated in real-time rather than only during annual audits dataart.com. In short, Policy as Code is increasingly viewed as a foundation for **** healthcare IT governance****, because it helps manage the "sheer volume" of rules and reduces the risk of human error in applying them higson.io higson.io.

Enforcing Compliance, Security, and Access Control with PaC

One of the biggest advantages of Policy as Code is how it enforces **compliance and security requirements automatically** across complex environments. In healthcare, compliance standards like **HIPAA, GDPR, and HITRUST** mandate strict controls over patient data confidentiality, integrity, and availability. By encoding these controls as code, organizations can



ensure they are uniformly applied. For instance, a policy that “all Protected Health Information (PHI) must be encrypted at rest and in transit” can be codified and continuously evaluated in [cloud deployments](#). AWS provides services (Config rules, etc.) that, when combined with PaC scripts, automatically check resources for HIPAA compliance and flag misconfigurations [dataart.com](#). Using PaC, organizations achieve **continuous compliance** – the system is always monitoring and enforcing policies, rather than relying on periodic manual checks [dataart.com](#). This greatly reduces the chance of drifting out of compliance in between audits.

Security and privacy are also strengthened by PaC. Policies for data handling (like who can access or share patient records) can be expressed in code and evaluated on each access request. This ensures **fine-grained access control** and prevents unauthorized disclosures. For example, Immuta (a data security platform) allows defining data access policies as code (in YAML) that automatically **mask or restrict personal data** across multiple databases [immuta.com](#) [immuta.com](#). This means if a regulation or internal policy says “Only clinical staff can see full patient identifiers, others see masked data,” the rule is centrally codified and enforced by the system in every query. The result is a consistent implementation of privacy rules, supporting principles like HIPAA’s minimum necessary access and GDPR’s data minimization by design.

PaC also helps enforce **access control policies** such as **least privilege and segregation of duties**. By treating access rules as code, organizations can implement complex conditions that go beyond traditional role-based access. For instance, policies can require that *only on-duty clinicians can access a patient’s record, and only if they are assigned to that patient’s care team*. With Policy as Code, such conditions are encoded and evaluated dynamically, blocking access that doesn’t meet the criteria. This approach enables **real-time checks against compliance and security rules** during system operations [sentinelone.com](#) [crowdstrike.com](#). As an example, Datagrid’s healthcare policy engine monitors HIPAA and other rules “through real-time enforcement capabilities,” meaning if someone attempts an action that violates policy, the system can immediately deny it or trigger an alert [datagrid.com](#). This real-time enforcement significantly improves security posture by preventing violations (or catching them instantly) rather than after the fact.

In summary, Policy as Code brings the benefits of **automation, consistency, and speed** to healthcare compliance and security. It allows policies to be **version-controlled, tested, and automatically deployed** similarly to software code [drata.com](#). This yields a more proactive stance on compliance: policies for configuration, access, [auditing](#), etc., are baked into the IT infrastructure and applications. Healthcare organizations using PaC have reported **greater automation, faster remediation, and reduced costs** in keeping up with evolving security standards [cleardata.com](#) [cleardata.com](#). In essence, PaC serves as a critical tool to enforce the myriad of healthcare rules (from internal policies to laws like HIPAA) in a uniform, error-resistant way, thereby strengthening both compliance and overall cybersecurity.



Rule Engines and Policy Frameworks in Healthcare

To implement Policy as Code, healthcare IT teams often leverage **rule engines and policy frameworks**. These are specialized software components that evaluate rules/policies written in a high-level language and determine whether certain conditions are met (often to allow or deny an action, or to trigger a decision). Several rule engine technologies are prominent in healthcare applications:

- **JBoss Drools:** Drools is a popular open-source *business rule engine* that uses an “if-then” rule language (DRL) and the Rete algorithm for efficient rule matching. It has been widely used in healthcare for implementing complex decision logic. For example, Drools can encode clinical guidelines or billing rules, and then automatically apply them to patient data or claims. One study used Drools to translate clinical criteria (Quality Data Model rules) for execution on [electronic health record data](#) [pmc.ncbi.nlm.nih.gov](#). Drools is known to be **powerful and flexible**, capable of handling a wide range of business logic [nected.ai](#). Healthcare organizations have used it for tasks like **clinical decision support systems (CDSS)** – e.g. to alert providers if a prescribed drug conflicts with a patient’s allergies [medium.com](#) – and for **revenue cycle management** – e.g. to validate insurance claims and coding. A vendor report notes that Drools (and similar engines) can manage tens of thousands of rules and process millions of transactions (claims) to ensure they meet payer policies and coding guidelines [careers.athenahealth.com](#). **Trade-off:** Drools excels at complex, interdependent rules and can perform inferencing (chaining rules based on facts). However, it typically runs as part of an application’s JVM or a rules server, which means **policies are often tightly integrated with application logic**. Maintaining Drools rules may require developer skills (Java/DRL), though business analysts can often author rules using guided rule editors.



- **Open Policy Agent (OPA):** OPA is a newer, general-purpose policy engine, widely used in cloud and microservices environments for access control and compliance. OPA decouples policy decisions from application code: you write policies in a declarative language called Rego, and OPA evaluates these policies given input data (like a user's role, resource, action). In healthcare, OPA is emerging as a tool to enforce fine-grained **authorization policies** in APIs and systems. For instance, a healthcare API (such as a FHIR server) can deploy OPA as a sidecar; when a client makes a data request, the API queries OPA to decide if the request is allowed [medium.com](#) [medium.com](#). OPA enables implementing **Attribute-Based Access Control** easily: policies can consider attributes like the user's job role, the patient's consent flags, time of day, location, etc., all in one Rego policy. It's "lightweight, general-purpose," and supports deployment as a sidecar, host agent, or library [medium.com](#). One Medium case study describes using OPA to safeguard a FHIR API, where it allowed "fine-grained, context-aware access control" critical for protecting patient data [medium.com](#). Similarly, a recent example by Ferguson (2025) showed OPA integrated with a SMART on FHIR authorization setup: the FHIR requests go through an API Gateway plugin which consults OPA policies before allowing access [rob-ferguson.me](#) [rob-ferguson.me](#). **Trade-off:** OPA policies (Rego) are extremely expressive for access decisions and OPA is designed for high performance in evaluating lots of small queries (e.g. many API calls). It cleanly separates policy from code and makes policies **versionable and auditable (since they are text files)** [rob-ferguson.me](#). On the other hand, OPA is stateless (each decision is made fresh based on input and policy, without built-in memory of past decisions) and is focused on **allow/deny style decisions** (though it can return arbitrary data). It might not replace a full business rules engine for sequential logic or calculations – instead, it shines as a centralized authorization service. Many healthcare microservice architectures adopt OPA for **unified policy management**, improving consistency across services [rob-ferguson.me](#). (Notably, Styra, the company behind OPA, cites healthcare as one industry benefiting from PaC to "smooth out the shift to cloud-native" healthcare apps [styra.com](#).)
- **XACML and Attribute-Based Policy Engines:** Before OPA, the eXtensible Access Control Markup Language (XACML) was a widely used standard for ABAC policies. XACML defines a common language (XML or JSON) for policies and an architecture (Policy Decision Point, Policy Enforcement Point, etc.). In healthcare, XACML-based engines (e.g., Axiomatics, Oracle Entitlements, WSO2 Balana) have been used to enforce fine-grained access rules, especially where regulatory requirements demand complex logic. For example, the U.S. Veterans Health Administration piloted ABAC to enforce patient consent directives: a policy could stipulate that mental health records require a special consent attribute, and the XACML engine would evaluate that before allowing access. **Trade-off:** XACML provides standardization and powerful policy constructs (conditions, obligations, combining algorithms). However, it has a reputation for complexity (policies can be verbose), and it requires mapping system attributes into the XACML framework. Modern approaches like OPA or domain-specific languages are sometimes favored for simplicity. That said, XACML or similar **policy-based access control (PBAC)** systems are still used in healthcare for enterprise scenarios requiring rich policies and audit trails. They excel in environments needing *policy federation* (sharing policies across institutions) or *externalized authorization* with formal compliance.

- Domain-Specific Rule Engines:** Healthcare also employs rule engines specialized for clinical or billing domains. For example, **FICO Blaze Advisor** (a commercial engine) is used by some large healthcare enterprises for claims rules and fraud detection. **Arden Syntax** engines and **Clinical Decision Support Services (CDSS)** are domain-specific rule systems for clinical guidelines (Arden Syntax is a standard for medical logic modules). Additionally, new tools like **HL7 CQL (Clinical Quality Language)** engines are used to express clinical criteria for quality measures or decision support. These engines often integrate with EHRs to provide real-time guidance (for example, alerting if a patient meets criteria for a sepsis protocol). While such systems may not be termed “Policy as Code” in the DevOps sense, they **function as policy engines for clinical policies** – ensuring care protocols are followed consistently via coded logic.

Integration and Comparison: It’s common to use multiple types of engines in concert. A hospital might use Drools or a proprietary engine for **clinical workflow rules**, and OPA or XACML for **security and privacy rules**. The choice depends on use case: Drools/BPM engines are ideal for *process-oriented* rules (e.g., determine insurance reimbursement based on complex criteria), whereas OPA/XACML are ideal for *access control decisions* (yes/no gates based on attributes). All these frameworks share the benefit of **centralizing policy logic** outside of hard-coded application software, which improves maintainability and consistency. As summarized in one industry comparison, open-source Drools offers a powerful, flexible rules engine for general business logic; OPA provides a unified, context-aware policy enforcement across stacks; and commercial engines like Blaze or cloud-based services offer scalability and domain-specific features [nected.ai](#) [nected.ai](#). Most modern tools (open or commercial) support healthcare requirements like **HIPAA compliance** (e.g., audit logs, role segregation) as a baseline [nected.ai](#). The table below outlines a brief comparison of popular policy-as-code tools:

Tool/Engine	Type & Language	Healthcare Use Cases	Strengths	Trade-offs
JBoss Drools	Production rule engine; DRL (Java-based)	Clinical decision support, claims processing, billing rules medium.com decisionrules.io	Powerful inferencing; handles complex logic and large rule sets; open-source nected.ai	Requires Java expertise; not specialized for distributed authZ (embedded in apps)
Open Policy Agent	Policy decision engine; Rego (declarative)	API/Microservice authorization (e.g. FHIR APIs), cloud config compliance medium.com	Decoupled, unified policy service; fine-grained ABAC; easy integration as sidecar; strong community/CNCF support rob-ferguson.me rob-ferguson.me	Learning curve for Rego; stateless decisions (no built-in session memory); primarily outputs allow/deny decisions
XACML (various vendors)	ABAC framework; XACML (XML/JSON) policies	Enterprise-wide access control, cross-organization data sharing policies (consent, security labels)	Very expressive attribute rules; standardized; mature in compliance-focused environments	Verbose policy format; complex setup; performance can be an issue if not optimized; less developer-friendly
FICO Blaze Advisor	Business rule management system (commercial)	Insurance plans, reimbursement and eligibility rules; healthcare revenue cycle	High performance; user-friendly rule authoring UI; good for large-scale decision tables	Expensive licensing; proprietary; needs domain experts to configure rules
Custom (Embedded)	Inline code checks (e.g. in	Simple site-specific policies (e.g. hard-coded	Quick to implement within app; no new tools required	Scattered logic, hard to audit or update consistently; prone to

Tool/Engine	Type & Language	Healthcare Use Cases	Strengths	Trade-offs
code or scripts)	EHR or SQL procedures)	role checks or filters in an EHR)		human error and drift from policy changes

As the table suggests, there is no one-size-fits-all solution. Many healthcare organizations adopt a **hybrid approach**: for example, using RBAC within applications for broad role permissions, and an overlay of ABAC via a policy engine for finer conditions. What's crucial is that whichever tool is used, it supports the ability to **externalize and automate policy enforcement**, and integrates with healthcare systems (EHRs, databases, APIs). Integration with standards and data sources is a key consideration – which we explore next.

Attribute-Based Access Control (ABAC) vs. Role-Based Access Control (RBAC)

Access control in healthcare has traditionally been dominated by Role-Based Access Control (RBAC), where permissions are assigned to roles (e.g. doctor, nurse, billing clerk) and users acquire permissions by being assigned those roles. While RBAC is straightforward and aligns with job functions, it has limitations in complex, dynamic environments like healthcare.

Attribute-Based Access Control (ABAC) has emerged as a more flexible model, using policies that consider many attributes (user attributes, resource attributes, environmental/contextual attributes) to make access decisions [hhs.gov](https://www.hhs.gov).

RBAC: In RBAC, access rights are predefined for each role. For example, a hospital might define that *"Nurses can view and edit nursing notes; Physicians can view all patient records and edit if they are the attending; Billing clerks can view billing info but not clinical notes."* Users are granted one or more roles, and thereby inherit the role's permissions [hhs.gov](https://www.hhs.gov). The simplicity of RBAC (users → roles → permissions) makes it easy to implement and understand. However, RBAC is **static** – it doesn't easily take into account context (time, location, patient consent, etc.). In a large hospital, RBAC can also lead to a role explosion problem: to capture nuanced differences (e.g., different access for an on-call physician vs. primary physician, or for emergency vs. routine access), admins might need to create many specific roles. Indeed, in complex environments like healthcare, "RBAC may lead to an explosion of newly created roles to accommodate every possible scenario" blog.lastpass.com. This becomes hard to manage and risks either too many privileges or an administrative nightmare of role management.

ABAC: Attribute-Based Access Control defines policies that evaluate attributes of the *subject* (who is requesting), the *object/resource* (what data or system is being accessed), the *action* (read, write, delete, etc.), and sometimes *environment* (time, location, etc.) [hhs.gov](https://www.hhs.gov). For example, an ABAC policy could be: *"Allow access to a patient's record if the requester's role is 'Physician' and the requester is assigned as that patient's care provider and the access is during an active treatment episode."* Here, *role = Physician* is one attribute, *relationship = provider for*



patient is another attribute (often determined by data like the patient's care team list), and possibly *context = during treatment* is an environmental attribute. All conditions must be true for access to be granted. ABAC policies can get very granular: **boolean logic combining arbitrary attributes** is allowed blog.lastpass.com. In fact, NIST defines ABAC as using policies that can express complex Boolean rules over attributes to determine access csrc.nist.gov.

Use Cases and Examples: Consider a scenario from the **LastPass blog** which illustrates ABAC in a hospital context. Dr. Adams is a cardiologist (user attribute: specialty) who wants to access cardiac patient records (resource attribute: record type = cardiac) during his shift (environment attribute: time) using a hospital-approved device (environment attribute: device compliance). An ABAC policy can ensure **all** those conditions are met: *IF user.department = Cardiology AND resource.type = "cardiac record" AND time ∈ [8am-8pm] AND device.trusted = true, THEN allow access* blog.lastpass.com. This fine-grained rule ensures Dr. Adams sees only the pertinent records and only when appropriate, enforcing least privilege. Moreover, ABAC can handle exceptions dynamically: if there's an emergency at 3am, an **"acute care ABAC"** extension might allow the cardiologist temporary access outside normal hours, provided an emergency flag is set blog.lastpass.com blog.lastpass.com. This kind of context-sensitive access is very hard to implement with pure RBAC, which would have needed a separate "OnCallAfterHours" role or similar for every permutation.

Another example: ABAC can include patient consent in the decision. A policy could say: *"Deny access to behavioral health notes unless the patient has given consent for this provider or it's an emergency."* Here the patient's consent flag is an attribute on the data, and emergency status might be an environmental attribute. In fact, Rob Ferguson's OPA example for SMART on FHIR suggests a policy: *allow medication data only if the app is launched by a clinician and the patient consented to data sharing* rob-ferguson.me – a clear ABAC scenario combining a user role attribute and a patient consent attribute. In traditional RBAC, you cannot easily incorporate such patient-specific conditions; ABAC handles it by design.

Comparing Benefits: ABAC's main benefit is **flexibility and granularity**. It can enforce **context-aware** and **content-based** restrictions, which is crucial for compliance with laws like HIPAA (which might require differing access based on treatment relationship) and for implementing principles like need-to-know. According to NIST, ABAC provides "fine-grained flexibility and security" and is well-suited for enterprises with complex structures such as hospitals with evolving requirements blog.lastpass.com. NIST SP 800-162 also notes that ABAC tends to be more **scalable and resilient** than RBAC in large, dynamic environments blog.lastpass.com. This is because adding a new scenario in ABAC often means adding a new attribute or policy rule, rather than proliferating roles. Indeed, ABAC can mimic RBAC (roles themselves can be just another attribute), but it can go further by adding more conditions.

RBAC's advantage is **simplicity and clarity**: it's easy for auditors and admins to understand that "User X has Role Y, which grants these permissions." ABAC policies, being arbitrary logic, require careful governance to avoid unintended consequences (policies can conflict or be hard to understand if not managed well). However, tools and practices are improving; for instance,



policies as code can be put under version control and peer-reviewed to maintain trust and clarity, and testing can be done to ensure ABAC rules do what is intended.

In practice, many healthcare organizations implement a combination: RBAC to set broad **eligibility (baseline permissions)** and ABAC to refine **contextual conditions**. A 2023 article on healthcare access control suggests a unified model where “RBAC layer enforces baseline access, and an ABAC layer applies granular, adaptive rules on top” [mdpi.com](https://www.mdpi.com). For example, RBAC might ensure only licensed clinicians can access any clinical data at all, while ABAC ensures each clinician only accesses records of their own patients. This hybrid approach can offer the best of both: RBAC for ease of administration, ABAC for fine-tuned enforcement. Indeed, one can view RBAC as a subset of ABAC (an attribute “role” can be just one attribute in a policy). The key is that **healthcare’s dynamic environment (shifts, emergency overrides, patient consent, varying state laws)** often necessitates ABAC capabilities. ABAC avoids the role-explosion problem by using attributes instead of ever-more specific roles blog.lastpass.com, thereby simplifying management in the long run for complex policies.

Integration with EHR Systems and Healthcare Standards

For Policy as Code to be effective in healthcare, it must integrate with core health IT systems and standards. Two crucial areas of integration are **Electronic Health Records (EHRs)** and **health data interoperability standards (like HL7 FHIR)**.

Integration with EHRs: Modern EHR systems (Epic, Cerner, etc.) typically have built-in role-based permissions and audit controls. Augmenting these with Policy as Code often means extending or interfacing with the EHR’s authorization mechanisms. This can be done in a few ways:

- *EHR APIs and Hooks:* Many EHRs expose APIs or “hooks” (for example, SMART on FHIR apps or CDS Hooks) that allow external systems to intervene. A policy engine can intercept requests to the EHR via these APIs. For instance, if an EHR uses the SMART on FHIR standard for apps, an external OPA service can be called during the OAuth authorization process to add extra checks beyond the standard scope-based (role-based) checks rob-ferguson.me. Ferguson’s 2025 example demonstrates using an API Gateway plugin for a FHIR server: the gateway forwards requests to OPA which enforces policy decisions before the request hits the FHIR data store rob-ferguson.me. This effectively inserts a policy enforcement point in front of the EHR’s FHIR API.
- *Database-level or Service-level Policies:* Some healthcare organizations implement policy checks at the database or service layer underlying the EHR. For example, an ABAC engine could sit between the EHR application and the database, evaluating queries. If a query violates policy (e.g., a user tries to access a record they shouldn’t), the engine can block it or sanitize results. A real-world approach is using SQL policy tagging or Immuta-like solutions on the data warehouse that the EHR uses, to ensure even back-end data access obeys policies (like automatically filtering out records from a certain department if not allowed).



- *Custom EHR Modules:* Large EHR platforms often allow custom rules or scripts. PaC can be introduced by writing these custom rules in a controlled, code-driven way. For example, in Epic one might use their native programming (like rules or their workflow engine) to enforce an organizational policy, but treat those rule definitions as code (with version control and testing outside of Epic, if possible).

HL7 FHIR and SMART on FHIR: HL7 FHIR (Fast Healthcare Interoperability Resources) has become the dominant standard for data exchange. It includes basic security mechanisms: primarily **SMART on FHIR**, an OAuth2-based protocol for apps to get access tokens with certain **scopes**. Out-of-the-box, SMART scopes implement a form of RBAC: e.g., a token might have scope "patient/Observation.read" meaning the app can read Observation resources for a patient. The FHIR server will enforce that scope (deny any request outside it). However, these scopes are fairly broad or role-oriented (user-level vs system-level, patient-specific vs global). To achieve true ABAC, additional context may need to be considered that isn't captured in a static scope string.

Policy as Code can enhance FHIR/SMART security by adding context-aware decisions:

- An OPA sidecar for a FHIR server can look at the full request (the user's roles, the patient ID, tags on the data) and apply organization-specific rules before the FHIR server processes it. The Medium article "Leveraging OPA to safeguard your FHIR API" describes exactly this: define policies on who can access which patient data under what conditions, then deploy OPA with the FHIR service such that every data request triggers a policy query medium.com medium.com. The outcome is that **fine-grained access rules (beyond SMART scopes) are enforced**. For example, you might allow Dr. Jones to access Patient X's oncology notes only if Dr. Jones is part of oncology team or if a consent resource exists linking Dr. Jones to Patient X.
- HL7 FHIR has a concept of **security labels** on resources (meta.security). These can encode confidentiality levels or categories (e.g., "MentalHealth" or "Restricted"). A Policy as Code system can leverage these labels as attributes in ABAC decisions. For instance, a policy could say: *deny access if resource.securityLabel = "Psychiatry" and user.department != "Behavioral Health"*. The FHIR server or an intermediary service would consult that policy whenever a request for a FHIR resource with that label is made groups.google.com. In practice, some FHIR implementations or extensions have used XACML or OPA to honor security labels and consent directives. In the HL7 Data Segmentation for Privacy (DS4P) standard, certain sensitive data is tagged and the receiving system is expected to enforce policies accordingly – an automated policy engine is well-suited for that enforcement.
- **SMART on FHIR ABAC extensions:** The SMART framework itself has begun to consider ABAC. For example, there's a concept of "launch context" where certain attributes (like which patient is being accessed, or purpose of use) can be factored into authorization. Some discussions (HL7 security workgroup) mention using the "launch" scope as a means to convey contextual ABAC decisions kodjin.com groups.google.com. In general, adding a Policy as Code layer means that even if an app gets an OAuth token with fairly broad scopes, the actual data access can be filtered or constrained by additional logic.

Beyond FHIR, other interoperability standards like **HL7 v2**, **DICOM**, and **HIE frameworks** can also integrate with policy engines. For instance:

- In Health Information Exchanges (HIEs), where data is shared across organizations, a policy engine can enforce consent and jurisdictional rules before allowing data to flow. E.g., if a patient opts out of sharing, the policy engine can prevent their data from being included in HIE queries.
- The emerging **FHIR Bulk Data APIs** and analytics platforms (like AWS HealthLake, Google Healthcare API) often include policy-as-code capabilities to ensure that large data exports still obey privacy rules (like excluding certain data or de-identifying it based on policy).

Standards Alignment: It's important that Policy as Code tools align with healthcare data standards. For example, a rule engine should be able to understand FHIR resource structures (which are JSON) – OPA can naturally handle JSON inputs, and Drools can parse JSON or use Java objects mapped to FHIR resources (HAPI FHIR library integration). Similarly, policy frameworks need to support healthcare vocabularies (roles, confidentiality codes, etc.). The good news is many policy-as-code implementations are **extremely flexible about input data** – they just need attributes, no matter the domain. It's up to healthcare IT architects to feed the relevant data (user roles, patient attributes, resource labels) into the policy queries. This often involves integration work: for instance, configuring the API gateway or middleware to supply the user's department, or the patient's consent flags, to the policy engine at runtime rob-ferguson.me.

SMART on FHIR App Authorization Example: To illustrate integration, imagine a SMART on FHIR app that wants to pull a patient's medication list. Normally, with SMART, the app gets a token with `patient/MedicationRequest.read` scope after the user (and possibly patient) authorize it. With Policy as Code in place, when the app calls the FHIR API for `/MedicationRequest?patient=123`, an OPA policy could additionally check:

- Is the requesting app a trusted application (e.g., registered in the hospital's system)? (attribute of client)
- Has patient 123 given consent for sharing medication data with this app or this type of user? (attribute of patient or a Consent resource)
- Is the user who authorized the app a clinician involved in patient 123's care, or the patient themselves? (user attributes and context)

Only if all conditions pass does OPA return `allow = true`, and the FHIR server then returns the data rob-ferguson.me. If not, OPA returns `deny`, and the FHIR server can return an error or filtered result. The **audit log** will also capture that decision from OPA, which is important for compliance review.

This level of integration shows how Policy as Code can augment standard healthcare auth protocols to achieve **contextual, patient-centered security controls**. It aligns with the trend of **"consent and policy enforcement as a service"** in healthcare IT, which is envisioned by



initiatives like ONC's rules for information blocking (which require honoring a patient's sharing preferences, something easiest done by a rules engine checking those preferences at runtime).

Policy Enforcement Use Cases: Real-Time Decisions in Healthcare

To make the discussion concrete, here are some **practical examples of policy enforcement** in real-time within healthcare scenarios:

- **Real-Time Clinical Decision Support (CDS):** In patient care, clinicians rely on up-to-date protocols and safety checks. A rule engine (policy as code for clinical rules) can enforce these in real-time. **Example:** A hospital deploys a **clinical rule engine** that encodes medication prescribing guidelines. When a doctor enters a medication order in the EHR, the engine automatically checks it against policies: "Is the dose within safe range for this patient's age and kidney function? Is the patient allergic or on a contraindicated medication?" If a rule triggers, the system can alert the doctor or even prevent the order until overridden (a "soft" vs "hard" stop policy). Such rules are often derived from internal policy (e.g., "No NSAIDs for patients with Stage 4 kidney disease" or "Follow CDC opioid prescribing guidelines"). By implementing them as code, the hospital ensures these policies are applied **consistently, for every patient and every order, in real-time**, rather than relying on memory or manual checks. This improves patient safety and standardizes care. Studies have shown that rule engines can reduce medication errors by generating **instant alerts for critical conditions**, prompting timely interventions [nected.ai](#) [nected.ai](#). These CDS rules need constant updating as evidence evolves – PaC makes updates easier to deploy uniformly.
- **Claims Processing and Billing Compliance:** On the administrative side, insurers and health systems use policy rules to automate claims and billing decisions. **Example:** An insurance company uses a rule engine (like Drools or DecisionRules) to automatically adjudicate claims. As a claim is received (with procedure codes, diagnoses, patient info), the engine applies a series of policy checks: *Does the patient's plan cover this procedure? Is the provider in-network? Is the claimed amount within allowed limits? Are there any coding edits or potential fraud indicators?* These checks are all codified as rules. The engine might approve straightforward claims or flag ones that violate a policy for manual review. According to one vendor, such a rules engine can "verify the accuracy of submitted claims, check for compliance with insurance policies and regulations, and determine reimbursement amounts" automatically [decisionrules.io](#). This real-time (or near real-time) processing speeds up reimbursements and reduces human error or bias. It also ensures **consistency** – every claim is evaluated against the same policy criteria, which is important for fairness and compliance (e.g., with Medicare rules). Because healthcare billing rules (like CMS's National Correct Coding Initiative or state Medicaid policies) change frequently, using Policy as Code means the insurer can update the rule set centrally and instantly apply it to all incoming claims. This contrasts with older manual processes where policy changes might not propagate to all adjusters quickly.



- **Patient Data Access Control (Privacy Enforcement):** Ensuring that only authorized individuals access patient data is a core security requirement (HIPAA's Privacy Rule calls this "minimum necessary" access). **Example:** A large hospital implements an ABAC policy engine to mediate all attempts to access patient records. When a staff member tries to open a patient's chart, the system checks a policy: *"Is this staff member involved in the patient's care or do they have a job function that permits this access?"* If not, access is denied or requires an override with justification (often called a "break-glass" scenario). For instance, a psychiatrist should not freely access oncology patients' records if they are not treating them. A policy engine can enforce this by checking the relationship attribute (treatment team vs not) and perhaps emergency context. If it's truly an emergency, a break-glass override could be allowed but logged (the policy might allow an emergency department physician to break-glass into any record **but** immediately log the access for audit). This dynamic decision-making is essentially real-time privacy enforcement. It prevents the common problem of *curiosity browsing* (staff peeking at records of high-profile patients without cause) by implementing the policy "access only if you have a need." In practice, some healthcare organizations have built such ABAC systems – for example, the Veterans Affairs ABAC system was designed so that **every access is evaluated against policies considering user role, patient relationship, and purpose**. If any condition fails, the user gets a "access denied" message or is prompted for additional authorization. The **auditability** of this approach is crucial: all decisions (grant or deny) can be logged with the attributes that were evaluated rob-ferguson.me, which helps in compliance audits and incident investigations (proving that, say, only authorized personnel accessed a VIP patient's data).
- **Real-Time Auditing and Remediation:** Policy as Code isn't only about allow/deny decisions; it can also automate enforcement actions. **Example:** A hospital has a policy that all servers containing PHI must have disk encryption enabled and latest patches applied. Using an **Infrastructure as Code** and PaC approach, they write policies for their cloud environment or data center (using tools like HashiCorp Sentinel, AWS Config rules, or OPA). If a server is launched without encryption or a baseline configuration, the policy engine immediately flags it and can even **auto-remediate** (shut it down or apply encryption). ClearDATA's CyberHealth platform, for instance, encodes such cloud security policies and will automatically **remediate misconfigurations** to maintain HIPAA compliance cleardata.com. This is a form of real-time operational decision: the system "decides" to quarantine or fix a resource because it violates code-based policy. In effect, the policy engine acts as a continuous auditor and guard for security compliance.



- **Consent and Data Sharing Decisions:** As healthcare data sharing increases (patient portals, health information exchanges, research data sharing), honoring patient consent and data use policies is vital. **Example:** A research database has data from multiple hospitals. A Policy as Code system governs queries to this database: *Researchers can only see de-identified data by default*. If someone has higher access (with IRB approval to see identified data), the policy might require that they only access data for which they have patient consent or a waiver. When a researcher queries, the system checks attributes like project ID, consent forms, and data sensitivity. If the query violates any rule (say trying to get identified data without proper approval), the query is denied or the identifiers are stripped automatically. Blockchain-based policy enforcement has even been proposed for this: one paper describes using smart contracts to encode consent policies that are **dynamically enforced in data sharing** – every data request triggers a blockchain check that the requester's attributes match policy, granting or denying accordingly, and logging the outcome immutably [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/31111111/). While blockchain is not widely used for this yet, it shows the forward-looking approach to **distributed trust and enforcement** of policies in real-time across organizational boundaries.

These examples demonstrate how Policy as Code enables **automated decision-making in live workflows**: whether it's a clinical action, an admin transaction, or an access event, the policy engine applies the rules and delivers an instant decision. This not only improves efficiency (decisions are made in milliseconds that might take a human committee days to review) but also ensures every decision is **consistent with defined policies**. Consistency is key for fairness and compliance: for instance, all claims are paid by the same rules, all accesses are granted by the same criteria – no favoritism or oversight gaps. And when exceptions are needed (like an emergency override), those too are governed by explicit policy (with safeguards like requiring proper documentation). This level of rigor and agility is increasingly necessary in healthcare as processes speed up and go digital.

Regulatory Considerations: HIPAA, GDPR, and Beyond

Healthcare organizations must navigate a web of data protection regulations. Policy as Code can be a powerful ally in achieving and demonstrating compliance with these laws and standards, as policies can be directly aligned to regulatory requirements:



- **HIPAA (Health Insurance Portability and Accountability Act):** In the US, HIPAA sets forth rules for protecting electronic Protected Health Information (ePHI). It mandates safeguards like access controls, audit controls, and transmission security. PaC allows organizations to enforce HIPAA technical safeguards systematically. For example, HIPAA requires that users have access only to the minimum necessary information for their job. An ABAC policy engine can enforce “minimum necessary” by dynamically restricting data fields or records based on user role and purpose. HIPAA also requires auditing of access – a policy engine automatically logs every decision (with attributes) providing an audit trail for compliance reviews rob-ferguson.me. Another aspect is configuration management: HIPAA's Security Rule (45 CFR 164.306) expects secure configurations; Infrastructure-as-Code combined with PaC can continuously enforce that all systems have required encryption, timeouts, etc., essentially **codifying HIPAA security policies**. As one practitioner notes, **Policy-as-Code can codify key HIPAA compliance policies within CI/CD pipelines** – for instance, automatically checking that no cloud storage bucket is open to public or that all data is in approved regions gartsolutions.com cleardata.com. If a policy violation occurs, the pipeline can fail the deployment, preventing non-compliant infrastructure changes. This approach moves compliance from a periodic audit to an ongoing DevOps process. In essence, PaC helps implement “security by design” as called for by HIPAA, since compliance checks are baked in at every step.
- **GDPR (General Data Protection Regulation):** GDPR in the EU imposes strict rules on personal data processing, including health data. Key principles include data minimization, purpose limitation, consent, and the right of individuals to control their data. Policy as Code can support GDPR compliance by enforcing *purpose-based access* and *consent management*. For instance, under GDPR a patient's data should only be used for the purpose they consented to. A policy engine can include “purpose of use” as an attribute in access decisions. If a user tries to access data for a purpose not allowed (say a researcher trying to see clinical data without patient consent for research), the policy can deny it. GDPR also grants the right to restrict processing – which could be implemented as a policy that if a patient flags their record as restricted, **no one can access it unless overriding legal justifications are present**. These rules can get complex, but ABAC is suited for complexity. Another area is **data retention/deletion**: policies can enforce that data is not retained longer than permitted. For example, a policy script could periodically purge or archive records that exceed the retention period unless flagged for an active patient. Moreover, GDPR's emphasis on accountability means organizations must *demonstrate* compliance. Policy as Code provides a clear mapping from regulation to system behavior – you can show auditors the code that enforces “only doctors in EU can access EU patient data” for data transfer restrictions, etc. This transparency and the audit logs of decisions help fulfill GDPR's accountability requirement.
- **Other Regulations and Frameworks:** There are many others: **HITRUST CSF** (a common security framework mapping HIPAA, NIST, etc.), **NIST 800-53** (security controls for federal systems), **ISO 27001**, **PCI-DSS**, and regional laws (like CCPA in California, or various countries' health data laws). Policy as Code allows organizations to unify compliance with multiple frameworks by writing rules that cover overlapping requirements. For example, both HIPAA and PCI require certain encryption and access controls – a single set of coded policies can ensure those across the environment. ClearDATA's PaC engine incorporates a wide array of frameworks (HIPAA, GDPR, GxP, NIST, PCI, etc.) and continuously updates with new regulatory enforcement actions cleardata.com cleardata.com. This highlights how PaC can be kept current with the regulatory landscape: when a new guideline or an incident occurs (say an OCR penalty for some misconfiguration), rules can be adjusted to prevent that scenario.



Data localization laws (like storing EU data in EU) can also be enforced by policy code in cloud deployments (checks on resource region). **Consent and sensitive category laws** (42 CFR Part 2 for substance abuse records in the US, for instance) can be encoded as attributes on data that policies use to restrict access unless certain consent is present. ABAC is actually very good at modeling complex regulatory rules because those rules often have conditional clauses ("if X and Y, then must do Z"), which is essentially policy logic.

One crucial regulatory concept is **auditing and reporting**. Regulations want evidence of compliance. A huge benefit of PaC is that it makes policies **auditable**: since policies are code, you can inspect them, version-control them, and test them. It's clear what rules are in place at any given time (unlike human procedures that might be applied inconsistently). Also, because enforcement is automated, you reduce the risk of an employee bypassing a policy out of convenience. And if they do (say via an emergency override), that event is logged and can be reviewed. This kind of before-the-fact control and after-the-fact logging greatly aids compliance. As noted, ABAC engines can even support "*before-the-fact audit*" – meaning policies can be simulated or analyzed to prove they meet certain compliance criteria before being deployed blog.lastpass.com. For example, you can query the policy base to answer "Could any unauthorized role ever access patient SSN under these policies?" and if the answer is no, that's evidence of compliance.

In summary, **Policy as Code operationalizes regulatory compliance**. It shifts organizations from a reactive posture (finding policy violations at audit time) to a proactive one (policies encoded to prevent violations). By aligning code with laws, healthcare IT can ensure that systems inherently enforce legal requirements (HIPAA, GDPR, etc.) and can adapt quickly to new rules. This is increasingly vital as regulations evolve (for example, new privacy laws or changes in HIPAA or cross-border data sharing rules). PaC provides the agility to update one code module and instantly have new restrictions in effect system-wide, which is much faster than retraining staff or manually reconfiguring dozens of systems.

Industry Adoption, Case Studies, and Best Practices

Policy as Code in healthcare is gaining traction as organizations seek to improve compliance and agility. A variety of healthcare stakeholders – from cloud service providers to hospitals and payers – have started sharing success stories and best practices:



- **Cloud Healthcare Platforms:** Companies like ClearDATA specialize in cloud services for healthcare and have made Policy-as-Code a core feature. ClearDATA's CyberHealth platform uses a **Policy-as-Code engine** to enforce healthcare-specific compliance in cloud environments cleardata.com. They report that this approach significantly reduces the time and cost for clients to maintain compliance, as the automation handles continuous monitoring and updates cleardata.com. By encoding frameworks like HIPAA and HITRUST into cloud configuration rules, ClearDATA helped over 200 healthcare organizations maintain secure cloud operations cleardata.com. This indicates broad adoption among those moving to cloud infrastructure – rather than manually checking each S3 bucket or VM for compliance, the policies do it automatically.
- **Health Tech and Startups:** Many digital health startups (handling sensitive data or AI in healthcare) adopt Policy as Code early on for security. For instance, anecdotally, healthcare API providers use OPA to manage authorization across microservices, since it's easier to update a Rego policy than to recode each microservice when a new access rule is needed. The Open Policy Agent project's own survey noted strong uptake in sectors like healthcare and finance blog.openpolicyagent.org. This suggests that healthcare organizations, aware of compliance demands, are often early adopters of tools like OPA or Sentinel. Companies providing FHIR-as-a-service or healthcare data platforms often highlight their use of PaC for things like data masking, multi-tenant isolation policies, etc., as a selling point.
- **Enterprises and Case Studies:** Large healthcare enterprises (hospitals, insurers) have begun publicly discussing internal projects:
 - For example, **Anthem (Elevance Health)** was known to use advanced rule engines for claims processing and fraud detection, processing billions of transactions with coded policies (though specifics might not be public, the scale indicates heavy rule automation).
 - **Athenahealth**, a major EHR and billing services provider, noted they maintain tens of thousands of billing rules in an engine to ensure claims are correct and compliant careers.athenahealth.com. They process tens of millions of claims through this engine, suggesting robust, high-scale use of Policy as Code in a critical financial workflow.
 - Government healthcare organizations, like the **U.S. Department of Veterans Affairs**, piloted an enterprise ABAC system to replace an old RBAC model. The goal was to better enforce patient consent and need-to-know. That project (part of VA's OSEHRA initiative) essentially is a case study in moving from coarse role-based control to fine-grained policy control across one of the largest health networks.
- **Best Practices Emerging:** As more projects implement PaC, some best practices have crystallized:
 - **Involve cross-functional teams:** Policy writing in healthcare should involve IT, compliance officers, and clinical/business stakeholders. A best practice is to have a **policy review board** that approves any new code-based policies, ensuring they align with actual regulatory interpretation and clinical workflow needs. This parallels how one would review written policies, but now reviewing code logic too.



- **Use source control and CI/CD:** Store policy definitions (Rego files, Drools rule files, etc.) in a version-controlled repository (e.g., Git). This provides traceability (who changed what policy when) and helps in audits. Automated tests can be written – for example, if a policy should allow certain access and deny others, you can have test cases to verify the code does that. This testing catches errors before deployment (crucial because a buggy policy could lock out needed access or leave a hole). Upon passing tests, integrate policy deployment into continuous delivery pipelines so that updates go live in a controlled, consistent manner drata.com drata.com.
- **Start with clear requirements and frameworks:** A common best practice is to map your policies to the regulatory or business requirements first in plain language drata.com. For example, list out: “HIPAA: must audit access to PHI; GDPR: must log consent for use; Internal policy: doctors in Dept A can’t see Dept B’s patients.” Then implement these in code. This ensures you don’t miss something and that the code is traceable to a requirement. Drata’s guidance suggests clarifying requirements and scope as a first step to successful PaC adoption drata.com.
- **Monitor and update continuously:** Policy as Code is not a set-and-forget. Healthcare regulations and organizational policies change (e.g., new COVID-19 emergency rules, or annual changes to coding guidelines). A best practice is to designate owners for policies who periodically review them. Monitoring is also key: use the logs from policy decisions to detect anomalies or policy failures. For instance, if you see repeated denials for a certain access that clinicians actually need, that might signal the policy is too strict and needs adjustment (or conversely, attempts to violate policy might signal an insider threat or training issue).
- **Performance and Scaling:** In large deployments, optimize the policy engine for performance. Best practices here include caching frequent decisions (if possible), using efficient data structures for attributes (e.g., indexing users by role for quick lookup), and load testing the policy engine. In critical healthcare operations, latency from a policy check needs to be minimal to not impede care. OPA, for example, can evaluate typical policies in microseconds and can be scaled horizontally if needed. Drools can handle large rule sets, but one must be mindful of rule design to avoid slowdowns (e.g., avoid overly broad conditions that trigger too many evaluations).
- **Fallback and Manual Override:** Design policies with safe fallbacks. In healthcare, availability is also critical – you don’t want the policy layer failing and blocking all access to EHR. Best practice is to have a fail-open or fail-closed strategy defined. Many go with fail-safe defaults: if the policy engine is unreachable, default to *deny sensitive actions but allow critical ones with logging*. Also, provide a mechanism for humans to override in dire situations (with proper approval or logging) – e.g., a break-glass code that, when used, logs that policy X was bypassed by Dr. Y at time Z with justification. Policy as Code can incorporate these as explicit rules (like “if override token provided and user is emergency role, then allow and flag for audit”).

Industry Collaboration: There are also collaborative efforts in the industry. For instance, HL7 and OAuth communities have been discussing **standardized approaches to convey attributes** (like user roles, patient context) in authorization protocols – which will make implementing ABAC easier. The Healthcare Identity Management community has interest in standards like UMA (User-Managed Access) for patient-mediated access control, which essentially is policy as code for patients (patients set rules about who can access their data, and a UMA service enforces it).



Case Study – UCSF (hypothetical example based on common practice): A large academic medical center implemented PaC for its multi-cloud infrastructure. They used Terraform (Infrastructure as Code) to provision systems and integrated OPA policies to ensure every infrastructure change meets security rules (e.g., no open security groups, proper network segmentation between research and clinical systems). During this process, they developed a library of policies corresponding to controls in their HITRUST certification. This not only kept them compliant continuously, but when it came time for audit, they could show auditors the live dashboard of compliance (with all controls coded and their current status) rather than assembling documents. The audit preparation time dropped significantly. Furthermore, when Log4j vulnerability hit, they quickly wrote a policy to scan all systems for the vulnerable library and flag any that hadn't been patched, demonstrating how PaC can assist in security incident response as well.

Overall, the trend is clear: **leading healthcare organizations are moving toward automated, code-driven policy enforcement** to cope with the scale and complexity of modern healthcare IT. Those who have adopted it report improvements in consistency, a reduction in compliance drift, and better confidence that they can pass audits or withstand security incidents. It aligns well with broader industry movements like DevSecOps and digital transformation in healthcare.

Challenges and Future Directions

While Policy as Code offers many benefits, there are challenges that healthcare organizations must address:

1. Complexity of Policy Management: As policies grow in number and complexity, managing them as code can become challenging. Writing correct and comprehensive rules requires skilled personnel who understand both the technical side and the legal/clinical side. There is a risk of **policy misconfiguration** – a subtle bug in a rule could inadvertently deny access to someone who needs it (impacting patient care) or allow a loophole. To mitigate this, organizations need rigorous testing and possibly policy simulation before deployment. In healthcare, testing policies can be complex because of the many corner cases (e.g., testing every combination of attributes like every role, every data category, every scenario like emergency vs non-emergency). **Tooling improvements** are needed to make policy authoring more intuitive, such as higher-level abstractions or GUIs that non-developers can use safely. We're likely to see more **user-friendly policy authoring interfaces** so that compliance officers or clinicians can contribute to policy definitions without needing to code in Rego or XML. Some vendors already tout "natural language" policy writing or low-code approaches for rules [nected.ai](#) [nected.ai](#).

2. Integration with Legacy Systems: Healthcare has many legacy systems (older EHRs, lab systems, etc.) that were not designed with external policy engines in mind. Integrating a modern policy as code approach might require significant plumbing – for example, wrapping an old system with an API gateway or modifying it to call out to a PDP (Policy Decision Point). In some

cases, you might not be able to enforce fine-grained policies without upgrading the system. This can slow adoption. One future direction is the development of **adapters or middleware** that can sit on top of legacy systems to enforce policies (perhaps intercepting database calls or using RPA techniques). As standards like HL7 FHIR gain traction, older systems will gradually be replaced or interfaced via FHIR, making it easier to apply external policies uniformly.

3. Performance and Scalability: As noted, policy engines need to handle potentially **high volumes of requests with low latency** – e.g., every EHR screen load might trigger dozens of policy checks for different data elements. Ensuring performance at scale is an ongoing challenge. Caching and indexing help, but one must be careful that caching doesn't serve stale decisions when attributes change (e.g., if a patient revokes consent, you need that to immediately reflect in decisions). Future directions might include more use of **event-driven updates** (push changes to the policy engine when attribute data changes, to invalidate caches) and **distributed policy evaluation** closer to the data (for instance, compiling policies down to run at the database level). Research is also looking at using **AI to optimize policy evaluation** – analyzing which rules are frequently triggered and optimizing those, or using predictive techniques to pre-evaluate some policies. For now, thorough performance testing and scaling horizontally are the best practices.

4. Policy Conflict Resolution: With many policies, conflicts can occur. For example, one policy might say "Allow research data access to Dr. Smith" but another says "Deny any access to oncology records without oncology role" – what if Dr. Smith (not oncologist) tries to access an oncology record for research? Which policy wins? In ABAC engines, typically a combining algorithm is defined (deny-overrides, permit-overrides, etc.). Deciding on the right strategy and communicating the outcome to end users can be difficult. We need better **policy analysis tools** that can detect overlaps or conflicts in policies before they cause an incident. The future might see **policy simulators** that can take a set of policies and a set of hypothetical requests and highlight any inconsistencies or unintended results, guiding policy authors to refine them. This is an active area of research in access control (some academic works focus on policy verification and anomaly detection).

5. Organizational Buy-in and Culture: Moving to Policy as Code can change workflows and responsibilities. Compliance teams and IT teams must collaborate more closely. Clinicians might need to be educated about automated enforcement ("the system will stop you if you violate policy X"). There can be resistance or workarounds if not handled well (e.g., if a policy blocks something a doctor needs, they might find unofficial ways to get the info, which is worse). Thus, an ongoing challenge is ensuring policies encoded in code still allow appropriate flexibility for patient care. *Governance committees* should continually review whether policies are too rigid or need exceptions. Also, from a cultural perspective, treating policies like software means adopting DevOps practices in a traditionally siloed environment. Training and change management are key – some staff might fear that automation is taking away judgment calls, so organizations need to emphasize that PaC is there to assist, not replace, professional judgment (except in clearly defined forbidden cases). In the long run, as successes accumulate (e.g., "We

prevented 50 improper accesses last quarter thanks to the policy engine, without blocking any legitimate care”), confidence and acceptance grow.

6. Evolving Threats and Requirements: The future will bring new challenges like increased ransomware attacks, more stringent privacy laws, and the expansion of healthcare beyond hospital walls (telehealth, IoT medical devices streaming data, etc.). Policy as Code will need to extend to these domains. For example, **medical IoT devices** might need policy enforcement at the edge – “a insulin pump should accept commands only from authorized hospital servers” – which could be a policy enforced in the device firmware via a small rules engine. We might see **micro-PDPs** running in devices or at edge gateways. Also, **AI in healthcare** may require policy checks (e.g., an AI algorithm can only use certain data if bias mitigation policies are satisfied). Interestingly, some are exploring using AI to help manage policies: perhaps learning from logs to suggest new policies or adjustments (like noticing a pattern of repeated override requests in emergencies and suggesting a refined rule for it). However, AI itself would need to be carefully validated in this context to avoid introducing risk.

7. Standards and Interoperability of Policy: Right now, each tool (OPA, Drools, etc.) has its own policy language. In the future, the industry might benefit from more standardized policy representations for common healthcare scenarios. Efforts like HL7’s basic consent and security labels are a start, but higher-level standard policies (e.g., a standard way to express “break glass” or “care team access”) could emerge, which organizations can adopt and then tweak. Projects like **HCAC (Healthcare Context-Aware Access Control)** models or the mentioned “AC-ABAC” in research blog.lastpass.com aim to formalize some of this. If successful, one could imagine a library of proven policy code modules that any hospital could plug in (with adjustments for their local policy) – reducing the burden of writing from scratch and ensuring best practices are shared.

In spite of these challenges, the trajectory for Policy as Code in healthcare is very positive. Many challenges are being actively worked on by the community. For example, open-source contributions to OPA include integrations with healthcare audit systems; the XACML community has tools for policy editing and testing; and healthcare standards bodies are aware of the need for dynamic access control standards. **Future directions** likely include:

- **Greater adoption of hybrid RBAC/ABAC (PBAC):** where systems provide intuitive RBAC management for broad strokes and ABAC for specifics, possibly with interfaces that hide the complexity (some call this Policy-Based Access Control or PBAC, effectively bridging RBAC and ABAC).
- **Policy-as-Code for patients:** giving patients accessible ways to specify their data sharing preferences (perhaps via smartphone apps) which then directly feed into policy engines. This patient-centric access control will be crucial as interoperability expands and as laws like GDPR emphasize individual rights.



- **Integration with blockchain and distributed ledgers:** to provide tamper-proof logs of policy decisions and possibly to decentralize trust (useful in multi-organization health networks where no single party fully trusts the others to enforce rules).
- **Continuous learning and adaptation:** eventually, policy enforcement might become more adaptive using risk analytics. For instance, if an unusual access request comes in, a risk engine (maybe AI-driven) could temporarily tighten policy or require step-up authentication. Some research (like risk-aware RBAC/ABAC models mdpi.com) is heading this way, blending static policies with dynamic risk scores.

In conclusion, Policy as Code is transforming how healthcare IT ensures that the right policies are in place and followed. It brings a level of precision, consistency, and agility that manual policies alone cannot achieve in the modern era. Though challenges exist in complexity, integration, and human factors, the ongoing advancements and collaborative best practices are steadily overcoming these. As healthcare continues to digitize and interconnect, Policy as Code will likely become a **standard practice**, much like quality checks and safety protocols are standard in clinical care – an integral part of the system that continuously guards compliance, security, and ethical use of sensitive health data.

Sources:

- ClearDATA, "Policy-as-Code™: Springing Compliance Frameworks into Action," Dec. 2021 cleardata.com cleardata.com cleardata.com cleardata.com cleardata.com cleardata.com.
- CrowdStrike, "What is Policy As Code (PaC)?," Feb. 2024 crowdstrike.com.
- Ferguson, R., "Use Open Policy Agent to enforce access control policies in SMART on FHIR applications," July 2025 rob-ferguson.me rob-ferguson.me rob-ferguson.me rob-ferguson.me rob-ferguson.me rob-ferguson.me.
- FHIRFLY (Medium), "Leveraging Open Policy Agent to Safeguard Your FHIR API," June 2023 medium.com medium.com.
- LastPass Blog, "Attribute-Based Access Control (ABAC)," Oct. 2023 blog.lastpass.com blog.lastpass.com blog.lastpass.com blog.lastpass.com blog.lastpass.com.
- HHS HC3 Cybersecurity Brief, "Access Control on Health Information Systems," Apr. 2020 hhs.gov.
- Nected.ai Blog, "Rule Engine for Healthcare: Transforming Medical Decision-Making," Jan. 2025 nected.ai nected.ai nected.ai.
- DecisionRules.io, "Business Rule Engine for Healthcare," 2023 decisionrules.io decisionrules.io decisionrules.io.
- Drata, "Policy as Code: Best Practices + Examples," 2023 drata.com drata.com drata.com.



IntuitionLabs - Industry Leadership & Services

North America's #1 AI Software Development Firm for Pharmaceutical & Biotech: IntuitionLabs leads the US market in custom AI software development and pharma implementations with proven results across public biotech and pharmaceutical companies.

Elite Client Portfolio: Trusted by NASDAQ-listed pharmaceutical companies including Scilex Holding Company (SCLX) and leading CROs across North America.

Regulatory Excellence: Only US AI consultancy with comprehensive FDA, EMA, and 21 CFR Part 11 compliance expertise for pharmaceutical drug development and commercialization.

Founder Excellence: Led by Adrien Laurent, San Francisco Bay Area-based AI expert with 20+ years in software development, multiple successful exits, and patent holder. Recognized as one of the top AI experts in the USA.

Custom AI Software Development: Build tailored pharmaceutical AI applications, custom CRMs, chatbots, and ERP systems with advanced analytics and regulatory compliance capabilities.

Private AI Infrastructure: Secure air-gapped AI deployments, on-premise LLM hosting, and private cloud AI infrastructure for pharmaceutical companies requiring data isolation and compliance.

Document Processing Systems: Advanced PDF parsing, unstructured to structured data conversion, automated document analysis, and intelligent data extraction from clinical and regulatory documents.

Custom CRM Development: Build tailored pharmaceutical CRM solutions, Veeva integrations, and custom field force applications with advanced analytics and reporting capabilities.

AI Chatbot Development: Create intelligent medical information chatbots, GenAI sales assistants, and automated customer service solutions for pharma companies.

Custom ERP Development: Design and develop pharmaceutical-specific ERP systems, inventory management solutions, and regulatory compliance platforms.

Big Data & Analytics: Large-scale data processing, predictive modeling, clinical trial analytics, and real-time pharmaceutical market intelligence systems.

Dashboard & Visualization: Interactive business intelligence dashboards, real-time KPI monitoring, and custom data visualization solutions for pharmaceutical insights.

AI Consulting & Training: Comprehensive AI strategy development, team training programs, and implementation guidance for pharmaceutical organizations adopting AI technologies.

Contact founder Adrien Laurent and team at <https://intuitionlabs.ai/contact> for a consultation.



DISCLAIMER

The information contained in this document is provided for educational and informational purposes only. We make no representations or warranties of any kind, express or implied, about the completeness, accuracy, reliability, suitability, or availability of the information contained herein.

Any reliance you place on such information is strictly at your own risk. In no event will IntuitionLabs.ai or its representatives be liable for any loss or damage including without limitation, indirect or consequential loss or damage, or any loss or damage whatsoever arising from the use of information presented in this document.

This document may contain content generated with the assistance of artificial intelligence technologies. AI-generated content may contain errors, omissions, or inaccuracies. Readers are advised to independently verify any critical information before acting upon it.

All product names, logos, brands, trademarks, and registered trademarks mentioned in this document are the property of their respective owners. All company, product, and service names used in this document are for identification purposes only. Use of these names, logos, trademarks, and brands does not imply endorsement by the respective trademark holders.

IntuitionLabs.ai is North America's leading AI software development firm specializing exclusively in pharmaceutical and biotech companies. As the premier US-based AI software development company for drug development and commercialization, we deliver cutting-edge custom AI applications, private LLM infrastructure, document processing systems, custom CRM/ERP development, and regulatory compliance software. Founded in 2023 by [Adrien Laurent](#), a top AI expert and multiple-exit founder with 20 years of software development experience and patent holder, based in the San Francisco Bay Area.

This document does not constitute professional or legal advice. For specific guidance related to your business needs, please consult with appropriate qualified professionals.

© 2025 IntuitionLabs.ai. All rights reserved.