CADMUS LABS
DESIGN PAPER

Technology Follows Purpose

The Architectural Requirements Behind BOCA Prime

Cadmus Labs
August 2026

Executive Summary

BOCA Prime is a proposed dental coverage system designed to make coverage clearer before care begins while reducing the administrative machinery required to coordinate it. Its architecture combines conventional and emerging technologies—including artificial intelligence, digital identity, privacy-preserving cryptography, programmable systems, blockchain infrastructure, and modern payment capabilities—with proprietary systems and technologies being developed by Cadmus Labs to address the particular requirements of the coverage model.

Each technology has a specific job to do. The architecture begins by identifying those jobs: establishing trustworthy facts, protecting sensitive information, preserving useful identity and history, applying defined rules consistently, coordinating financial obligations, and creating meaningful accountability around how the system operates.

Many of these responsibilities are well served by conventional technology. Private clinical records require secure information systems. Dental practices need familiar interfaces and integrations. Compliance, customer support, analytics, security operations, administrative functions, and many payment activities can continue to use established software and infrastructure.

BOCA Prime is attempting something more fundamental than placing a new interface over the existing coverage process. It is being designed to move more of the coordination itself into the architecture—establishing necessary facts earlier, automating routine determinations, connecting verified coverage decisions more directly to payment, and allowing different parts of the system to work together with less repetitive administration.

That shift creates additional architectural requirements. As more responsibility moves into the system itself, Prime needs reliable ways to establish which facts can be trusted while protecting the sensitive information used to establish them. Identity and relevant history need continuity

when that continuity can eliminate unnecessary repetition. Established rules need to produce consistent outcomes, and consequential system activity needs to remain accountable.

The more coordination the architecture assumes, the more important trust, privacy, continuity, verification, and accountability become.

BOCA Prime's current architectural hypothesis brings together different technologies and Cadmus-developed components to address these responsibilities. Artificial intelligence may support clinical verification, anomaly detection, analysis, and increasingly sophisticated coordination, including through DentAI, a proprietary AI capability Cadmus Labs is developing for verification and system integrity within the BOCA architecture.

Digital identity and privacy-preserving cryptography may support persistent identity, selective disclosure, credentials, permissions, and verifiable facts. Cadmus Labs is developing **Rootprint™**, its patent-pending identity and trust architecture, to explore how verified identity, participation, relationships, credentials, and relevant history can persist where continuity creates value.

Programmable systems may make selected rules and obligations more consistent and inspectable. Blockchain infrastructure is designed to provide a shared foundation for selected forms of verification, accountability, programmable coordination, and persistent trust across organizational boundaries.

These capabilities can operate alongside conventional systems rather than replacing them. The resulting architecture is deliberately hybrid: institutional responsibility remains where legal, financial, clinical, security, and operational accountability require it, while selected forms of verification, execution, and participation can become more distributed where doing so creates meaningful value.

This paper explains the reasoning behind that architecture. It examines the requirements BOCA Prime creates, the technologies being considered to address them, the tradeoffs those choices introduce, and the standards each technology will ultimately have to meet.

Technology earns responsibility by creating a property the system actually needs and performing that responsibility well enough to justify its place in the architecture.

Scope and Status

This Design Paper documents Cadmus Labs' current architectural reasoning for BOCA Prime.

BOCA Prime remains a proposed future coverage system. Its technical architecture will continue to evolve through legal and regulatory analysis, economic and actuarial modeling, technical development, security testing, and evidence from real-world participation. The technologies

discussed here represent current design directions and working hypotheses rather than finalized implementation commitments.

The paper focuses specifically on the relationship between **system requirements and technology choices**. Detailed coverage economics, actuarial assumptions, benefit structures, regulatory classifications, and commercial implementation are addressed elsewhere in the BOCA research and development program.

External technical capabilities described in this paper are supported by references where appropriate. Proposed applications of those capabilities within BOCA Prime represent Cadmus Labs' architectural reasoning and remain subject to validation.

Technology Follows Purpose

BOCA Prime creates a demanding coordination problem.

The system needs to reliably establish who the participants are—including the identity of the patient and the verified status of the provider. A patient and dental practice then need to know whether coverage is active, whether a proposed service qualifies, what evidence is required, how the applicable rules affect the benefit, what the coverage system will contribute, and what financial responsibility may remain with the patient. The system needs to reach those determinations while protecting sensitive information, preserving appropriate clinical judgment, managing exceptions, maintaining financial integrity, and operating within legal and regulatory requirements.

Once those responsibilities are considered together, the technology architecture becomes easier to reason about.

The system needs to establish trustworthy facts.

It needs to protect information that should remain private.

It needs to preserve relevant identity and history where continuity creates value.

It needs to apply established rules consistently.

It needs to maintain meaningful accountability around consequential rules, financial obligations, and system activity.

It needs to move routine coordination toward automation where that automation can operate reliably.

And when a coverage obligation has been established, the architecture should help connect that determination efficiently to what happens next.

Those needs can be organized around several architectural properties.

Privacy. Clinical, personal, and financial information should remain appropriately protected even when selected facts derived from that information must be verified.

Verification. Important facts, credentials, rules, and actions should be capable of establishing trust among participants who may belong to different organizations and maintain different systems.

Continuity. Relevant identity, participation, relationships, credentials, and trust should be capable of persisting where continuity creates value.

Participant agency & control. Where appropriate, identity, credentials, permissions, and participant-linked information should remain under meaningful participant control. The architecture should support informed permission around how selected information is used or shared and preserve the ability to recognize participant contribution when that participation creates value.

Accountability. Selected rules, financial obligations, reserve activity, and important system changes should become easier to inspect and more difficult to obscure.

Execution. Once necessary facts have been established and an obligation determined, the architecture should support a direct and reliable path toward fulfilling that obligation.

Appropriate boundaries. Routine coordination should move toward automation wherever the system can perform it reliably, while human judgment and institutional accountability remain available where the nature of the responsibility requires them.

These properties provide a practical way to evaluate technology choices. A technology earns architectural responsibility when its capabilities materially improve one or more of these properties and when the resulting benefit justifies the complexity, risk, and operating requirements it introduces.

That standard applies equally to conventional and emerging technology.

Conventional Technology Remains Essential

A large portion of BOCA Prime can be built using familiar systems.

Patient and practice interfaces, internal administration, customer support, private databases, practice-management integrations, analytics, compliance workflows, protected clinical records, security operations, artificial-intelligence infrastructure, manual review, and many payment functions all have established technological approaches.

Those systems matter because reliability and usability are themselves architectural requirements.

A dental practice interacting with BOCA Prime should be able to work through interfaces designed around familiar clinical and administrative workflows. A patient should be able to understand coverage without learning the infrastructure beneath it. Protected records should remain in environments designed for protected records. Compliance and security functions should retain clear institutional ownership. Existing payment infrastructure should remain available wherever it provides the appropriate combination of speed, cost, legal certainty, consumer protection, and ease of use.

This produces an important design discipline: **technological complexity should be visible only when the participant actually benefits from seeing it.**

A blockchain-based mechanism, for example, can operate behind an ordinary patient action such as reviewing and accepting treatment. The underlying interaction may use blockchain infrastructure or programmable execution, but the patient should not need to understand cryptographic keys, smart contracts, or the technical process occurring beneath the interface. A cryptographic proof can support a coverage process without asking a dental office to understand the mathematics behind it. A programmable payment mechanism can operate beneath an ordinary interface. An AI-assisted verification system can present its conclusion and supporting information through familiar software while more complicated analysis occurs underneath.

The architecture can therefore incorporate emerging capabilities without requiring the user experience to become technologically exotic.

Conventional technology provides much of the operational foundation. The additional architectural questions arise where multiple participants need to rely on information, rules, or actions that cross organizational boundaries.

That is where coordination begins to become a question of trust.

When Coordination Becomes a Trust Problem

Coverage requires authoritative information.

Participant identity and provider status must be reliably established. Eligibility must be established. Benefit rules must be known. Participation must be recorded. Clinical requirements may need verification. Contributions and payments must be tracked. Coverage determinations need to be preserved. Financial obligations and reserve activity require records.

A conventional database can perform all of those storage functions.

The architectural challenge appears when several independent participants need to rely upon what those records say.

Patients need to trust that their coverage has been represented accurately. Practices need to trust that the rules being applied to treatment and payment are the rules under which they agreed to participate. Administrators need reliable information from patients and practices. Auditors or regulators may need evidence of financial or operational activity. Different technical systems may need to exchange facts without sharing every underlying record.

The system therefore needs more than stored information. It needs a way to establish **which information can be trusted, who established it, what authority they had to establish it, and whether consequential records or rules have changed.**

Blockchain infrastructure offers useful properties for selected parts of that problem. The National Institute of Standards and Technology (NIST) describes blockchain as a distributed digital ledger of cryptographically signed transactions grouped into blocks, with each block cryptographically linked to the previous one so that records become tamper-evident and increasingly tamper-resistant as the chain grows.[1]

For BOCA Prime, the relevant capability is shared verification.

A selected record can potentially become independently checkable by authorized participants without requiring one organization's private database to provide the only authoritative account. Cryptographic signatures can establish provenance. Recorded state changes can create an auditable history. Shared infrastructure can allow separately administered systems to coordinate around selected facts.

The value of those properties depends on where they are applied.

Clinical records require privacy. Internal administrative data may have no reason to leave Cadmus-controlled systems. Many routine operations benefit from clear centralized ownership. Other information—such as selected credentials, attestations, rules, system parameters, proofs, or records of consequential actions—may gain value from being independently verifiable.

This distinction allows institutional responsibility and independent verification to coexist.

Cadmus Labs can remain accountable for operating BOCA Prime while the architecture reduces the number of important claims that participants must accept solely because Cadmus says they are true.

That matters over time as well as at launch. Organizations change. Leadership changes. Economic pressures and incentives evolve. Rules accumulate. A system designed around verifiable behavior can preserve selected forms of accountability even as the institution operating it develops.

Trust can therefore become an architectural property rather than relying exclusively on institutional reputation.

Verification Without Unnecessary Exposure

Stronger verification can improve trust and accountability, but it must coexist with the privacy expected of healthcare and financial information.

BOCA Prime may need to establish that a patient is eligible, that a provider is authorized, that a required clinical condition has been verified, that a particular rule applies, or that an obligation has been established. None of those requirements automatically creates a need to expose the complete underlying record.

This distinction is central to the privacy architecture.

A clinical record can remain within an appropriately protected environment while another part of the system receives only the information necessary for a specific coverage decision. An identity can be established without repeatedly disclosing every identifying attribute. A credential can communicate that an authorized party has verified something without forcing every subsequent participant to repeat the original verification.

Where the nature of the information and the interaction permit it, the participant should also have meaningful control over what information is shared, with whom, and for what authorized purpose. Privacy is therefore not only a question of protecting information from exposure; it is also a question of preserving appropriate agency over how participant-linked information moves through the system.

Modern digital-credential systems already provide structures for this kind of separation. The **World Wide Web Consortium (W3C)** Verifiable Credentials model allows claims issued by one party to be cryptographically verifiable by another, while current credential standards include mechanisms intended to support selective disclosure and data minimization.[2]

Zero-knowledge techniques extend the same architectural principle. Under appropriate implementations, they can allow a party to demonstrate that a statement is true or that a defined condition has been satisfied without revealing all of the underlying information used to establish it.[3]

For dental coverage, the practical value is straightforward.

The system may need to know that a requirement has been satisfied.

It may not need every component of the system to possess the private information that established it.

A protected radiograph could remain in an appropriate clinical environment while an authorized verification process establishes the fact required by the coverage rules. A participant could establish eligibility without repeatedly exposing unrelated identity information. A provider credential could remain verifiable across interactions without distributing the underlying documentation every time it is needed.

Blockchain infrastructure can potentially provide a shared verification layer for selected proofs or attestations while the sensitive information supporting them remains elsewhere.

This creates a different relationship between privacy and trust.

The architecture can seek stronger verification while reducing the amount of sensitive information that has to move through the system—and, where appropriate, giving participants greater agency over that movement.

The need for this kind of verification does not necessarily end with a single coverage decision. A verified identity, credential, participation history, or relationship may remain relevant across future interactions. Requiring those facts to be repeatedly reconstructed can reintroduce the administrative burden the architecture is intended to reduce.

That creates the next architectural question: how can appropriate forms of verified information remain trustworthy and useful over time?

And that is where the architecture begins to move from isolated facts toward **identity, continuity, and Rootprint**.

Identity, Continuity, and Rootprint

Verification establishes whether a fact can be trusted. A continuing coverage relationship creates a second requirement: **that trust may need to remain useful over time.**

Dental care rarely occurs as a single isolated interaction. Patients return to practices. Providers establish relationships with patients and coverage systems. Participation continues across years. Clinical history accumulates. Credentials remain relevant. Circumstances change.

Yet many administrative systems repeatedly reconstruct information that has already been established elsewhere.

A patient may change dental practices and repeat portions of an identity or coverage process. A provider may have credentials verified independently by multiple organizations. Information established during one interaction may have limited ability to remain useful in another because each organization maintains its own records for its own purposes.

That fragmentation creates administrative work, but it also creates a deeper architectural question:

Can selected facts remain trustworthy beyond the interaction in which they were originally established?

BOCA Prime explores that possibility through Rootprint.

Rootprint™
Patent Pending

Rootprint is being designed to give participants a trusted digital identity that can remain useful over time and across different interactions.

Selected information—such as verified credentials, participation history, permissions, and trusted relationships—can remain connected to that identity so trusted information does not have to be repeatedly reconstructed or verified as participants move through different interactions.

Continuity is only useful if the identity beneath it can be trusted. Rootprint therefore begins with the problem of establishing with high confidence that the participant interacting with the system is who they claim to be—and, for providers and organizations, that the relevant identity, status, or authority has been appropriately verified.

As coverage coordination becomes increasingly automated, that foundation becomes increasingly important. The requirement applies to patients, providers, organizations, and eventually software acting with authority delegated by those participants.

Rootprint would not be the first system to rely on trusted identity or credentials across organizational boundaries. Existing digital-identity standards already demonstrate that one organization can establish information about a participant and another organization can later rely on that information without independently recreating the entire verification process.

NIST describes identity federation as an approach in which a credential service can authenticate a participant and provide verified information about that participant to a separately administered system that relies on it.[4] In practical terms, one trusted system can establish identity or an attribute once and allow another authorized system to rely on that result.

W3C's Verifiable Credentials architecture extends a related idea to digital credentials: one party can issue a credential, the participant can hold or present it, and another party can cryptographically verify that credential.[2]

These established models help demonstrate the technical foundation for an important Rootprint objective: trusted information does not necessarily have to be recreated from the beginning every time a participant interacts with a different authorized system.

Rootprint explores how those capabilities could support a continuing dental-coverage relationship by strengthening the connection between a digital identity and the person that identity represents over time. The objective is for confidence in that relationship to strengthen as multiple independent forms of trusted evidence accumulate over time.

That evidence can draw upon multiple forms of verification rather than depending on a name, account, or single identifier. For patients, professionally attested clinical evidence—including radiographs—could provide one particularly useful source of continuity. Dental radiographs can capture a distinctive combination of anatomical features, restorations, treatment history, pathology, and other characteristics that can be compared across time for human identification.[5]

A dentist could attest that a set of radiographs belongs to a particular patient. Future radiographs and other professionally verified records could then add additional evidence to that relationship. Changes visible across time—including treatment that has occurred and anatomy viewed from different images or angles—could strengthen confidence that those records belong to the same person and provide further continuity rather than requiring the identity relationship to begin again with each interaction.

Once Rootprint has established sufficient confidence that a digital identity belongs to a particular person, that trusted identity can become a foundation for future interactions.

Rootprint brings together three complementary forms of trust: **biological continuity, professional attestation, and cryptographic continuity.**

The radiographic evidence described above can help establish biological continuity. Professional attestation connects that evidence to legitimate clinical interactions. Cryptographic mechanisms can then preserve selected proofs, permissions, or attestations in a form that authorized participants and systems can verify without requiring access to the underlying clinical record. This combination—rather than any one form of evidence alone—is central to the Rootprint hypothesis.

Participant agency and control are also part of the design. Where appropriate, participants should be able to authorize how credentials, permissions, and participant-linked information associated with their Rootprint identity are shared or used, while clinical records remain protected and maintained by the institutions responsible for them.

The result is **continuity with selective portability**: relevant trust can follow the participant without requiring the underlying clinical record to be shared simply to establish that trust. This separation between the evidence used to establish something and the proof another system actually needs is part of the privacy-preserving architecture Rootprint is intended to provide.

Rootprint's exact implementation remains subject to technical, privacy, security, and regulatory validation.

For BOCA Prime, Rootprint is intended to provide a persistent, privacy-preserving identity foundation that becomes more trustworthy through real-world interactions while giving participants meaningful control over the information the architecture allows them to control.

With identity and continuity established, the architecture can turn to another question: **can the rules and financial obligations applied to that identity become similarly trustworthy?**

Rules, Financial Activity, and Accountability

Coverage depends on rules.

Eligibility criteria determine who participates. Benefit parameters determine which obligations the system accepts. Contribution requirements affect financial participation. Reserve policies influence system resilience. Coverage conditions determine when a benefit applies. Exceptions determine when routine execution should stop and additional review should begin.

BOCA Prime is being designed to establish more of those rules before individual treatment decisions occur.

That creates a corresponding requirement: **participants should have meaningful confidence that the rules governing an interaction are the rules actually being applied.**

Programmable systems may help with that responsibility.

A rule that can be expressed clearly enough for reliable software execution can potentially produce greater consistency than a process that repeatedly requires individual interpretation—or changes with the whims of a new CEO, board, or administrator. Eligibility conditions, benefit parameters, contribution rules, payout conditions, selected system limits, and financial obligations may therefore become candidates for increasingly programmable execution.

The appropriate degree of programmability will vary.

Some rules may translate readily into deterministic logic. Others will depend upon clinical interpretation, incomplete information, unusual circumstances, legal requirements, or judgment that cannot responsibly be reduced to a binary condition.

BOCA Prime's architecture therefore needs both **reliable execution and explicit exception paths**.

Routine cases can move through defined logic when the necessary facts are established with sufficient confidence. Cases involving uncertainty, conflicting information, unusual circumstances, or responsibilities requiring professional judgment can move into appropriate review.

This structure can make automation more useful because it defines its boundaries.

The same approach can strengthen accountability.

If a coverage rule is established before treatment, the system can preserve which version of the rule applied, which verified facts were used, what determination resulted, and when consequential parameters changed.

Programmable execution can then provide more than speed. It can create a clearer relationship among **rule, evidence, decision, and outcome**.

The same need for accountability extends to financial activity.

A coverage system may coordinate contributions, reserves, benefit obligations, payments, and other movement of value. Participants, auditors, regulators, or other authorized parties may need confidence that selected financial responsibilities are being handled as represented.

Blockchain infrastructure can help make selected financial and system activity independently verifiable and more difficult to alter after the fact. NIST identifies tamper-evidence and tamper-resistance as core properties of blockchain records.[1] For BOCA Prime, those properties could strengthen accountability around consequential activity such as changes to system parameters, defined financial obligations, and reserve movements.

Privacy remains important here as well.

Financial accountability does not require publishing individual patient transactions, exposing protected information, or removing appropriate participant control over how participant-linked information is shared. Cryptographic proofs, attestations, access controls, and appropriately designed audit mechanisms can allow different participants to verify different facts according to the authority and information each legitimately requires.

For example, an authorized reviewer may need evidence that a reserve requirement has been satisfied without needing access to every participant transaction. A practice may need confirmation of a specific payment obligation without receiving unrelated financial information about the system. A participant may need evidence that the applicable coverage rule was the rule actually used without gaining access to protected records belonging to someone else.

This is another reason the architecture is unlikely to rely on one technical mechanism for every responsibility.

Accountability can be shared while information remains appropriately bounded.

There is also an important economic boundary.

A technically verifiable financial record can establish what occurred.

It cannot establish that the underlying financial policy was wise.

A proof that assets exist does not establish that reserves are sufficient. An auditable transaction history does not determine appropriate contribution levels. A programmable reserve rule cannot decide whether the actuarial assumptions beneath that rule are sound.

BOCA Prime will still require actuarial analysis, reserve modeling, stress testing, financial management, and appropriate regulatory oversight.

Verifiability can strengthen financial accountability. Sound economics determines whether the financial system itself is viable.

The architecture needs both.

From Coverage Determination to Payment

Establishing a coverage obligation earlier can reduce uncertainty before treatment. The administrative benefit depends partly on what happens after that obligation has been established.

Consider the sequence BOCA Prime is attempting to create for a routine case.

The participant is identified. Participation and eligibility are established. The necessary clinical information is verified. The applicable coverage rules are known. The system determines its contribution, and the patient's remaining responsibility becomes clearer.

At that point, the system has already performed much of the work necessary to determine the financial obligation.

The next architectural objective is therefore to make fulfillment of that obligation **a continuation of the coverage process rather than the beginning of another administrative process.**

This matters because payment friction has consequences beyond settlement speed.

Every additional handoff creates another opportunity for delay, reconciliation, repeated information requests, inconsistent records, manual follow-up, and administrative cost. Dental practices devote staff time to tracking payments for care already delivered. Patients can remain uncertain about balances while transactions are reconciled. Coverage administrators spend resources processing information that may already have been established elsewhere in the workflow.

A more coordinated architecture can shorten that distance.

Conventional payment infrastructure will remain important.

ACH, bank transfers, cards, and regulated payment providers offer established integrations, consumer familiarity, legal frameworks, and operational infrastructure. Newer conventional systems are also reducing settlement delays. The Federal Reserve's FedNow Service, for example, enables participating financial institutions to provide instant-payment services around the clock, with funds generally available within seconds.[6]

That capability helps clarify the standard BOCA Prime should apply to newer payment technologies: speed alone is not enough. Digital payment rails need to provide additional advantages—such as programmable execution, interoperability, auditability, or more direct integration with conditions established elsewhere in the coverage architecture—to justify their role.

Digital payment rails introduce additional capabilities.

Tokenized money and stable-value digital assets can support forms of programmable settlement in which payment logic is more directly connected to conditions established elsewhere in the system. The Bank for International Settlements has examined tokenisation as a means of integrating messaging, reconciliation, and asset transfer on programmable platforms, while also emphasizing the institutional, regulatory, and monetary foundations required for trustworthy financial systems.[7]

For BOCA Prime, the important question is what those capabilities can improve in the coverage workflow.

A payment could move more quickly once an obligation has been established. Programmable execution could connect settlement to conditions the coverage architecture has already verified. Transaction records could remain more closely connected to the obligation they satisfy, simplifying reconciliation. Interoperable digital infrastructure could also make it easier to coordinate payment across separately administered systems.

Different payment mechanisms will offer different combinations of these capabilities.

The architecture can therefore select payment rails according to the obligation being fulfilled.

A conventional payment rail may provide the best combination of usability, integration, settlement characteristics, protections, and legal or regulatory fit in one circumstance. A programmable digital rail may provide the better combination of those same requirements while adding advantages in areas such as conditional execution, interoperability, programmability, or auditability in another.

The participant experience can remain simple in either case.

A patient does not need a lesson in payment infrastructure before understanding or using a benefit. A dental practice should not need to redesign its financial operations around a particular technology simply because Prime uses that technology somewhere beneath the interface.

At the same time, the architecture can preserve choice. Where legally, technically, and operationally appropriate, patients could choose among supported methods of payment, and practices could choose how they prefer to receive funds—including through conventional payment methods or supported digital rails. Participants who want access to newer financial capabilities should be able to use them without making those capabilities a requirement for everyone else.

The architecture should absorb unnecessary complexity while preserving useful optionality.

Once a valid obligation has been established, payment should follow through the mechanism that fulfills it most reliably, efficiently, transparently, and appropriately.

The larger sequence then becomes straightforward: verified identity establishes the parties; verified clinical information establishes the relevant facts; defined rules establish the obligation; programmable systems can coordinate execution; payment infrastructure fulfills the obligation; and recorded activity provides evidence that fulfillment occurred.

The objective is a shorter distance between **knowing what the system owes** and **the system doing what it has already determined it should do**.

Governance as Architecture

Technology can help establish facts, apply rules, preserve records, and execute obligations.

It cannot decide by itself who should have the authority to define or change the rules governing the system.

That responsibility belongs to governance.

BOCA Prime will require clear institutional authority throughout development and operation. Cadmus Labs and any regulated or otherwise responsible entities involved in the system must be able to satisfy legal obligations, maintain financial stability, protect participants, respond to security events, administer operations, and remain accountable for responsibilities that require clearly defined institutional authority.

Those requirements establish one side of the governance architecture.

The other comes from the nature of the system itself.

Patients contribute to and depend upon coverage. Dental practices provide care and interact with the administrative architecture. Financial and operational participants may support functions upon which the system relies. Rules governing contributions, benefits, verification, reserves, participation, and other consequential matters can affect these groups differently.

A mature system therefore needs a disciplined way to decide who has authority over which decisions, how that authority is constrained, and how consequential changes become visible to the people affected by them.

Cadmus Labs' current design direction preserves strong institutional responsibility while exploring where transparency, programmable constraints, independent verification, and carefully bounded participant governance could improve the system.

One objective is to make selected commitments harder to abandon as the institution operating BOCA changes over time. Leadership, ownership, financial pressures, and institutional priorities can change. Where a rule or constraint is important enough to the integrity of the system, the architecture may be able to limit how easily any future administrator—including Cadmus Labs itself—can alter it unilaterally.

Blockchain-based governance and programmable constraints create possibilities for moving selected protections from institutional promises toward enforceable system rules. The appropriate mechanisms remain to be determined, but the principle matters: some protections may be stronger when participants do not have to rely entirely on the continuing judgment or incentives of whoever happens to control the organization.

Governance may also provide a way to incorporate information and preferences distributed across the people participating in the system. Patients, practices, and other participants experience different parts of the coverage relationship and bear different consequences when rules change. Carefully designed participation could allow that information to influence selected decisions while preserving the institutional authority required for legal, financial, security, and operational responsibilities.

That creates the possibility of governance serving both as a constraint on concentrated power and as a mechanism for better collective decision-making. Participants may sometimes support decisions that impose costs on themselves when the tradeoffs are understandable and the decision strengthens the system they depend upon. A contribution increase, benefit adjustment,

reserve requirement, or other consequential change could potentially be evaluated with greater visibility into why it is necessary and what effects it is expected to have.

The same principle could extend to the economic relationships surrounding care. Greater transparency around costs, contributions, benefits, and system constraints may allow patients and practices to make better-informed decisions without requiring the coverage system to dictate every price or clinical choice. That is consistent with BOCA Prime's broader objective of preserving participant and provider autonomy while making the information surrounding those choices more understandable.

Governance therefore has a larger purpose than accountability alone. It can help preserve the system's foundational commitments, constrain inappropriate concentrations of authority, incorporate participant knowledge, and create a more durable balance among constituencies whose interests overlap without always being identical.

That could take several forms. Important proposals or rule changes could become visible before implementation, while authority over different types of decisions could be explicitly defined. Consequential changes could leave a reliable history. Programmable guardrails could constrain selected actions even when initiated by an authorized party. And where participant involvement is appropriate, defined groups could eventually gain a formal voice in decisions that directly affect their participation.

This is another expression of participant agency and control. Participation in a system does not necessarily justify authority over every decision, but people and practices affected by consequential rules may have legitimate reasons to understand how those rules change and, in appropriate circumstances, to have a defined role in that process.

Distributed-governance mechanisms provide additional tools for designing that participation, along with important risks. Token-based voting can concentrate influence among large holders. Participation may be low or unrepresentative. Strategic behavior and conflicts of interest can affect outcomes. Decision-making can become slow or vulnerable to capture. Research on decentralized autonomous organization governance has documented several of these tensions.[8]

Rootprint may provide another option for how participation is structured. If the system can establish persistent, difficult-to-duplicate participant identities with sufficient confidence, selected governance processes could potentially assign influence according to verified participation rather than simply according to the number of tokens or wallets a participant controls. In appropriate circumstances, that could support models closer to one verified participant, one vote—or other structures designed to prevent economic ownership alone from determining governance power.

That does not determine how BOCA Prime should ultimately govern itself. It expands the range of governance models the architecture can credibly investigate.

Those risks matter particularly in a coverage system involving healthcare and financial responsibilities.

Governance must support accountability without making the system unable to act. It must allow appropriate evolution while protecting financial stability, legal compliance, security, and participants from poorly constrained decisions.

This leads toward **bounded participation rather than indiscriminate decentralization**.

Some responsibilities require clear authority. A security team needs to respond to an active threat. Financial managers need defined responsibility for maintaining required reserves and controls. Legal and regulatory obligations need accountable institutions capable of satisfying them. Clinical judgment requires appropriate professional boundaries.

Other decisions may benefit from greater transparency, recorded change processes, defined consultation, programmable constraints, or eventually formalized participant involvement.

The governance architecture can distinguish among these responsibilities.

The goal is a system in which authority is clear, consequential power is appropriately constrained, and the people who share the system's consequences can gain greater visibility into—and where appropriate influence over—how it evolves.

Why the Architecture Is Hybrid

The requirements described so far do not converge on a single technological model.

They converge on a **division of responsibilities**.

Some functions benefit from clear institutional ownership. Others benefit from independent verification, participant control, programmable execution, or shared infrastructure. Some require both.

That is why BOCA Prime is being designed around a **hybrid architecture**.

Conventional software can provide the interfaces, protected records, integrations, administration, compliance functions, and security operations that benefit from clear institutional responsibility.

Artificial intelligence has a more specialized role within the BOCA architecture. DentAI is being developed as a proprietary AI capability to support clinical verification, anomaly detection, and system integrity. As Prime becomes more automated, DentAI could help evaluate whether clinical information, utilization patterns, submissions, and other activity are consistent with what

the system would reasonably expect—helping identify potential fraud, waste, abuse, errors, or unusual cases that warrant additional review.

That makes DentAI an important counterpart to Rootprint. Rootprint helps establish who is participating and preserve trustworthy identity over time; DentAI can help evaluate what is occurring across those identities and interactions. Together, they form part of the verification architecture Prime will need if routine coordination is going to become more automated without sacrificing appropriate safeguards.

Other parts of the architecture address requirements that become more important across organizational boundaries. Digital identity and cryptographic systems can support permissions, attestations, selective disclosure, and independently verifiable facts. Rootprint can provide continuity around identity and trust. Programmable systems can connect established facts and defined rules to execution. Blockchain infrastructure can provide shared verification, tamper-evident history, and coordination that does not depend entirely on one organization's private database.

Payments can use conventional or digital rails according to the requirements of the obligation being fulfilled. Governance can preserve institutional accountability while using transparency, programmable constraints, independent verification, and appropriately bounded participant involvement to limit concentrated authority and protect important system commitments.

The design challenge is to place each responsibility where it can be performed most effectively and responsibly.

A clinical record may remain inside a protected conventional system while an authorized proof establishes the fact another part of the architecture needs. Cadmus Labs may remain accountable for operating the coverage system while selected rules and consequential actions become independently verifiable. Artificial intelligence may evaluate information while permissions and defined rules constrain its authority. Routine obligations may move toward automated execution while explicit exception paths preserve appropriate review.

The same principle applies to participant control and financial infrastructure. A participant can exercise appropriate authority over credentials or permissions without assuming responsibility for institutional records. A conventional payment rail can fulfill one obligation while a programmable digital rail fulfills another.

The architecture can therefore distribute trust, verification, and execution selectively while keeping responsibility identifiable.

Centralized and decentralized systems provide different properties. Institutional control can provide coherent administration, legal accountability, security responsibility, and operational authority. Decentralized infrastructure can strengthen independent verification, persistence, programmable constraints, and forms of coordination that should not depend entirely upon one institution.

BOCA Prime can use both.

Centralized responsibility and decentralized verification can be complementary parts of the same system.

Hybrid does not necessarily describe the architecture's final destination.

BOCA Prime is being designed to begin with clearly accountable institutions because the system must operate safely, legally, and reliably while its assumptions are tested. Over time, capabilities that demonstrate they can operate reliably through programmable rules, independent verification, automated execution, or appropriately distributed governance may require less continuing intervention from a central administrator.

The longer-term architectural objective is to explore how much of the system can become reliably autonomous without sacrificing the protections, accountability, financial integrity, and human judgment its responsibilities require.

Evidence will determine how far that transition should go. Some responsibilities may remain institutional indefinitely. Others may become increasingly automated or decentralized as the architecture proves that they can be.

The destination therefore remains deliberately open, but the direction is intentional: **reduce unnecessary dependence on centralized administration wherever the system can responsibly replace that dependence with verifiable rules, reliable automation, and appropriately distributed authority.**

That combination also creates obligations of its own. Every additional technology introduces complexity. Every integration creates dependencies. Distributed functions create questions about authority, security, recovery, interoperability, and responsibility when something fails.

Hybrid architecture therefore requires clear boundaries.

Each component needs a defined responsibility. Each connection needs a reason to exist. Automation needs conditions under which it proceeds and conditions under which it stops. Distributed mechanisms need an accountable relationship to the institutions and participants relying upon them.

And each technology needs evidence that **the improvement it produces is significant enough to justify the complexity and risk it introduces.**

That discipline becomes even more important as software moves beyond assisting people and begins increasingly to **act on their behalf**.

That is the next architectural horizon.

Designing for an Agentic Environment

The architecture described so far is intended to support BOCA Prime as a coverage system.

It may also become increasingly relevant as software assumes a more active role in coordinating healthcare, financial, and administrative activity.

Artificial intelligence is already moving beyond analysis and content generation toward systems capable of interpreting information, using tools, completing multistep tasks, communicating with other systems, and taking actions on behalf of people and organizations. Future healthcare and financial environments may include increasingly capable software agents operating within authority delegated by patients, providers, administrators, and institutions.[9]

That development creates an important extension of BOCA Prime's existing architectural requirements.

Consider a patient trying to understand and use dental coverage.

A future patient agent might be authorized to review available benefits, identify applicable coverage rules, communicate permitted information to a dental practice, compare treatment-related financial options, coordinate scheduling, or arrange payment according to preferences established by the patient.

A dental practice might use another agent to establish participation, obtain relevant coverage information, submit required evidence, coordinate a coverage determination, identify exceptions requiring staff attention, or reconcile financial activity.

A coverage system may use specialized agents to assemble verified information, apply established rules, coordinate downstream processes, communicate with authorized systems, and identify cases requiring additional review.

Financial institutions and payment providers may operate their own automated systems.

Those agents could eventually need to coordinate with one another across organizational boundaries.

That environment depends on more than intelligence.

It depends on **authority**.

A useful agentic system needs reliable ways to establish **who an agent represents, what authority it has been given, which information it may access or disclose, and which rules**

constrain its actions. That authority may need to expire or be revoked. The system also needs ways to resolve conflicting information, identify decisions that require human or institutional review, and preserve enough evidence to establish what was authorized and what actually occurred.

These requirements connect directly to the architecture BOCA Prime is already exploring.

Rootprint could become particularly important in that environment. An agent may have a cryptographic identity of its own, but the more consequential question is the relationship behind it: which real person or organization does this agent represent, and what authority has that participant actually delegated to it?

Rootprint could provide a persistent, biometric, longitudinal, and verifiable foundation for that relationship. An agent acting for a patient could be connected to the patient's established Rootprint identity along with explicit permissions defining what the agent may access, disclose, decide, or authorize. Those permissions could potentially be limited by purpose, duration, transaction type, financial amount, or other defined conditions—and revoked when the participant no longer wants the agent to act.

That would extend Rootprint from establishing continuity of identity to helping establish continuity of authorized agency. As software becomes capable of acting across healthcare, financial, and administrative systems, reliably connecting an agent to the real participant behind it—and proving the boundaries of that delegated authority—could become an important piece of trustworthy machine-to-machine coordination.

Privacy-preserving credentials and proofs could then allow an agent to establish necessary facts without gaining access to every piece of sensitive information beneath them. Programmable rules could constrain execution, while tamper-evident records could preserve evidence of consequential actions.

DentAI and other specialized AI capabilities could operate within those boundaries to perform verification, analysis, anomaly detection, and increasingly sophisticated coordination. Payment infrastructure could then allow properly authorized obligations to move toward settlement once the required conditions have been satisfied.

Blockchain infrastructure may become particularly useful when autonomous agents need to coordinate or transact across organizational boundaries. Agents can operate continuously, interact directly with other software, and execute transactions without waiting for the business hours or manual processes around which many traditional systems were designed. Programmable blockchain infrastructure and stable-value digital assets can provide machine-readable rules, always-available settlement, and payment mechanisms that software can execute directly when authorized conditions have been satisfied.[10]

This is already moving beyond theory. Research from Visa and Artemis examining live onchain activity describes AI agents beginning to transact using blockchain-based payment infrastructure and identifies stablecoins and low-cost blockchain settlement as emerging

components of agentic commerce.[10] The significance for BOCA Prime is broader than payment alone: a shared programmable environment could eventually allow agents representing patients, practices, coverage systems, and financial institutions to coordinate identity, permissions, verified facts, obligations, and settlement without requiring every organization to operate inside the same private system.

These capabilities become especially important when agents belonging to different organizations need to work together.

A patient agent and a practice agent do not necessarily share a database. A practice agent and a coverage agent do not necessarily share the same incentives. A coverage agent and a financial institution's systems may operate under different rules, permissions, security models, and legal obligations.

Reliable coordination among them requires common ways to establish **identity, authority, permissions, trustworthy facts, and outcomes**.

Existing standards provide useful foundations for parts of that problem. W3C's Verifiable Credentials architecture defines roles for issuers, holders, and verifiers and provides mechanisms through which cryptographically secured credentials can be presented and verified across organizational boundaries.[2] NIST's digital-identity guidance similarly addresses identity federation and assertions among separately administered systems.[4]

Artificial intelligence creates another requirement: the amount of authority assigned to an AI system should reflect how reliably and safely it can perform the responsibility it has been given. NIST's AI Risk Management Framework identifies characteristics including validity and reliability, safety, security and resilience, accountability and transparency, explainability, privacy, and fairness as relevant dimensions of trustworthy AI systems.[9]

Those concerns become more consequential as software gains authority to act rather than merely provide information.

BOCA Prime's approach to an agentic environment therefore begins with **constrained authority**.

An agent should be able to do what an authorized participant has permitted it to do, using information it is permitted to access, within rules appropriate to the responsibility it has been given.

More capable automation can build on that foundation.

For DentAI, that could eventually mean assuming more routine coordination as its capabilities and safeguards are validated, while cases involving uncertainty, conflicting evidence, clinical judgment, security concerns, or other high-stakes conditions move through appropriate review.

The result could be a system in which human attention becomes increasingly concentrated where judgment and accountability are actually required, while software handles more of the repetitive coordination surrounding routine cases.

The timing and extent of that transition remain open.

BOCA Prime can operate without autonomous agents coordinating every interaction. Its near-term architecture does not depend on a particular prediction about how quickly agentic software will mature or how healthcare regulation will respond.

Designing for that possibility today means building around properties Prime already needs: **identity, permissions, privacy, verifiable facts, explicit rules, bounded authority, and accountable execution.**

Those properties support BOCA Prime's current architecture. An increasingly agentic environment would make them even more important.

BOCA Prime can prepare for that future without depending upon it. The architecture is being designed to remain trustworthy if more of the participants interacting through it eventually become software acting on someone's behalf.

Blockchain and the Token Question

Blockchain infrastructure occupies a meaningful place in BOCA Prime's current architectural hypothesis.

It can support shared verification, tamper-evident records, programmable coordination, persistent attestations, and forms of accountability that cross organizational boundaries.

A native BOCA-related digital asset presents a separate design decision.

Blockchain infrastructure can provide those capabilities without requiring BOCA Prime to create a native token.

The token question should therefore be evaluated on its own merits:

What additional capability could a native digital asset create for BOCA that the rest of the architecture cannot provide as effectively?

Several possibilities deserve investigation.

A native asset could potentially support incentives for forms of participation that strengthen the system. It could compensate participants or infrastructure providers for useful contributions, support selected governance mechanisms, or provide a programmable means of transferring value across a broader BOCA network.

It could also create new ways to recognize participant contribution. If participants voluntarily contribute information, verification, governance participation, infrastructure, or other activity that helps create measurable benefit for the broader system, the architecture could explore whether some portion of that benefit should return to the people whose participation helped produce it.

That possibility fits BOCA Prime's broader principle of participant agency: participation should not automatically mean surrendering control while all resulting economic benefit accrues to the institution operating the system.

A native asset might also help align economic incentives among participants whose actions affect the health of the network. Depending on legal, financial, and technical design, digital assets could eventually interact with payment, collateral, reserves, or other financial mechanisms.

Each proposed role would need to answer a different question.

An incentive mechanism needs to identify which behavior deserves encouragement and whether rewarding it actually improves the system. A compensation mechanism needs a defensible way to determine who contributed and what contribution occurred. A value-transfer mechanism needs to perform some function better than conventional money or contractual arrangements.

Governance requires particular care.

A native asset could potentially support governance, but **token ownership does not have to determine governance power.** As discussed earlier, Rootprint may make it possible to structure selected governance processes around verified participants or other criteria rather than simply giving the participant with the most tokens the most influence.

A token's economic function does not need to determine how governance power is distributed.

Any role involving collateral, reserves, or other coverage-related financial resources would face an even higher standard and require careful actuarial, legal, liquidity, custody, security, and regulatory analysis.

Tokenisation more broadly is receiving serious institutional attention. The Bank for International Settlements has described tokenisation as a potentially transformative way to combine assets, information, rules, and transaction execution on programmable platforms, including the possibility of reducing reconciliation and enabling new financial arrangements.[7]

That supports investigation of the underlying capabilities.

It does not establish that BOCA needs a native asset.

A BOCA-related asset would introduce its own economic system. Its supply and distribution would matter. So would liquidity, incentives, custody, security, tax treatment, usability, and its relationship to the financial obligations of the coverage system. Regulatory treatment could also depend heavily on the specific functions the asset performs.

Speculation introduces another design challenge.

A useful BOCA-related asset may develop a market value, and market participants may form expectations about that value. The architecture therefore needs to distinguish **utility created by the system from demand created primarily by expectations of price appreciation**.

That produces a useful test:

A native asset should strengthen BOCA Prime even when speculative appreciation is removed from the analysis.

If an asset improves coordination, rewards legitimate contribution, enables a useful form of participation, supports a defensible governance mechanism, improves value transfer, or creates another capability that materially improves the system, that function can be evaluated on its own merits.

If its usefulness depends primarily on people wanting the asset because they expect someone else to pay more for it later, the architecture has not established a sufficient reason for it to exist.

The same discipline applies to stable-value digital assets used as payment infrastructure.

Stablecoins and other forms of tokenized money can provide useful capabilities around programmability, interoperability, around-the-clock settlement, and machine-readable financial execution. At the same time, their design can raise questions involving reserve quality, redemption, custody, financial integrity, interoperability, regulatory treatment, and their relationship to the broader monetary system. BIS analysis continues to examine both the capabilities of tokenised finance and weaknesses or policy risks associated with current stablecoin arrangements.[11]

BOCA Prime can evaluate those tradeoffs according to the financial responsibility being performed.

The architecture should also remain adaptable as digital-payment infrastructure evolves. Stablecoins, regulated tokenized deposits, tokenized commercial-bank money, conventional instant-payment systems, and other programmable financial infrastructure may develop substantially before Prime reaches later stages of implementation. BIS is already examining models that combine tokenized central-bank and commercial-bank money with programmable platforms, reinforcing how quickly the available design space is expanding.

BOCA Prime should preserve the ability to use whichever financial rail best serves the obligation rather than tying the architecture unnecessarily to one form of money.

A native BOCA asset therefore remains an **architectural hypothesis**.

Cadmus Labs intends to investigate whether such an asset can perform a real function within the system that other mechanisms cannot perform as effectively.

Blockchain has defined responsibilities within the current architecture. A native asset must earn responsibilities of its own.

Technology Must Earn the Responsibility We Give It

BOCA Prime's architecture asks technology to perform consequential work.

Identity, clinical verification, coverage decisions, financial coordination, payment, governance, privacy, and increasingly automated execution can all affect real people, real practices, and real financial obligations.

Technical capability alone is therefore not enough.

The relevant question is whether a technology can perform the responsibility assigned to it reliably, securely, economically, legally, and understandably enough to improve the system that depends upon it.

Artificial intelligence illustrates the point.

AI can evaluate complex information, identify patterns, detect anomalies, assist verification, and coordinate increasingly sophisticated workflows. Within BOCA Prime, DentAI could use those capabilities to support clinical verification and system integrity, including identifying patterns that may indicate fraud, waste, abuse, error, or cases requiring additional review.

Those capabilities also introduce risks. Model error, bias, manipulation, unreliable inputs, security vulnerabilities, explainability limitations, and changing performance can all become more consequential as AI assumes greater responsibility.

NIST's AI Risk Management Framework treats trustworthy AI as a combination of characteristics—including validity and reliability, safety, security and resilience, accountability

and transparency, explainability, privacy enhancement, and fairness—rather than a single measure of model capability.[9] DentAI should therefore receive responsibility according to what it can reliably perform, the evidence supporting its conclusions, the consequences of error, and the safeguards available when uncertainty remains.

Cryptographic systems create a different balance.

They can strengthen privacy, verification, credential integrity, and auditability. Their effectiveness depends on correct implementation, secure key management, usable recovery mechanisms, reliable standards, interoperability, and appropriate integration with the systems holding the underlying information.

Mathematical strength does not automatically produce operational reliability.

Programmable execution creates another balance.

Encoding a rule can make its application more consistent and inspectable. It can also allow a flawed rule to execute consistently at scale.

Its reliability therefore depends on the quality of the rule, the information supplied to it, the ability to recognize exceptions, and the mechanisms available to stop or correct execution when necessary.

Blockchain infrastructure can strengthen shared verification, persistence, programmable coordination, and historical integrity. Those capabilities must justify the additional requirements the infrastructure introduces around privacy, scalability, interoperability, governance, security, recovery, and integration with accountable institutions.

Governance mechanisms can constrain concentrated authority and create meaningful participant involvement. They can also reproduce concentration in different forms. Research on decentralized autonomous organization governance has documented problems involving concentrated voting power and other governance tensions.[8] Rootprint-enabled participant identity may allow BOCA to investigate alternatives to governance based primarily on token ownership, but those alternatives will need evidence of their own.

Digital assets face the same standard. Incentives, participant compensation, governance, programmable value transfer, or other functions may justify a role, but only if the capability provided is strong enough to warrant the economic, regulatory, security, custody, usability, liquidity, and speculative risks introduced with it.

Interoperability and agentic coordination add another layer. Connecting participants and systems can reduce repetitive administrative work and allow information, authority, and obligations to move more efficiently across organizational boundaries. Every connection also creates dependencies and requires clear permissions, security boundaries, and responsibility when information conflicts or something fails.

These tradeoffs are not arguments against the architecture. They are part of designing it responsibly.

None of these technologies earns responsibility merely because it is new or emerging.

The evaluation begins with the system rather than the technology: **What does BOCA Prime need to accomplish? What properties are required to accomplish it? Which technology can perform the resulting responsibility best?**

From there, the standard becomes practical. The architecture must consider the advantages a technology provides, the risks and dependencies it introduces, whether those risks can be appropriately bounded, whether ordinary participants can use the resulting system without unnecessary complexity, and whether implementation can satisfy the legal, regulatory, security, privacy, and economic requirements surrounding the responsibility.

The improvement has to justify the complexity and responsibility assigned to the technology.

That process may produce different answers across Prime.

Conventional software may remain the strongest solution for many functions. Artificial intelligence may earn increasing responsibility as performance and safeguards improve. Cryptographic identity and verification may provide important advantages where trust needs to cross organizational boundaries. Programmable systems may make selected rules and obligations more consistent and inspectable. Blockchain infrastructure may provide shared verification, accountability, persistence, and coordination. Digital payment rails may improve particular forms of settlement. Participant governance may improve selected decisions while remaining inappropriate for others. A native digital asset may ultimately establish sufficient utility to earn a role.

The architecture can accommodate those possibilities without requiring every possibility to survive.

That flexibility is deliberate.

BOCA Prime will develop while technology, regulation, financial infrastructure, artificial intelligence, cryptographic standards, and participant expectations continue to change. The strongest solution available today may not remain the strongest solution later. A capability that appears unnecessary today may become important as the system evolves. A design that works technically may fail economically or operationally. A compelling theoretical mechanism may prove too difficult for ordinary participants to use.

Evidence should be able to change those decisions.

The architectural commitment is to the **properties the system needs and the outcomes those properties are intended to produce.**

Technology earns responsibility by demonstrating that it can help deliver those outcomes—and that the benefits of assigning it that responsibility justify the risks, complexity, and constraints that come with it.

Conclusion: Whether It Creates a Better System

BOCA Prime is an attempt to redesign the coordination surrounding dental coverage.

Its architecture begins with the experience the system is intended to create.

Patients should be able to understand more of their coverage before making treatment decisions. Dental practices should spend less time administering coverage and more time running their practices and caring for patients. Necessary information should become easier to verify without requiring unnecessary exposure of sensitive information. Identity and relevant trust should be able to persist when doing so prevents the same relationships from being repeatedly reconstructed. Participants should have meaningful agency over the information, permissions, and credentials the architecture allows them to control.

Coverage rules should become clearer and more consistent. Routine coordination should become increasingly automated where automation can perform it reliably. Financial obligations should become easier to establish, verify, reconcile, and fulfill. And the institutions responsible for operating the system should remain accountable while the architecture makes selected forms of verification, execution, governance, and trust less dependent on institutional discretion alone.

Those objectives lead to a hybrid architecture.

Conventional software provides much of the operational foundation. Rootprint is intended to provide persistent, privacy-preserving identity and continuity. DentAI can support clinical verification, anomaly detection, fraud, waste, and abuse protection, and increasingly sophisticated coordination. Cryptographic systems can protect sensitive information while allowing necessary facts to remain verifiable. Programmable systems can connect established facts and defined rules to execution. Blockchain infrastructure can provide shared verification, tamper-evident history, programmable coordination, and infrastructure that does not depend entirely upon one organization's private database.

Conventional and digital payment rails can fulfill different obligations according to their strengths and participant preferences. Governance can preserve clear institutional responsibility while constraining concentrated authority and creating appropriate opportunities for transparency and participant involvement. Agentic systems may eventually use these same foundations to act on behalf of patients, practices, and institutions within explicitly delegated authority. A native digital

asset may eventually earn a role if it provides a function the rest of the architecture cannot perform as effectively.

The longer-term direction is toward a system capable of performing more of its routine coordination reliably through the architecture itself. As technologies, rules, safeguards, and governance mechanisms prove themselves, some responsibilities may require less continuing intervention from a central administrator.

How far that transition should go must be determined by evidence.

The technological ambition therefore remains substantial, but technological sophistication alone does not improve dental coverage.

Improvement has to appear in the experience and operation of the system.

A patient understands the financial circumstances of treatment sooner. A practice avoids repetitive administrative work. A verified fact can be trusted without exposing unnecessary information. Identity established through trusted interactions becomes stronger rather than beginning again each time. Established rules produce consistent and inspectable outcomes. Financial obligations move more directly toward fulfillment. Automation removes repetitive coordination while preserving appropriate safeguards. Participants gain greater visibility into consequential parts of the system and meaningful control where that control properly belongs.

And if the architecture requires fewer resources to administer the same necessary functions, more of those resources can remain available for the purposes the coverage system exists to serve.

Those are the outcomes that give the technology a reason to be there.

They also provide the standard for deciding when a particular technology should expand, remain limited, change, or disappear from the architecture altogether.

BOCA Prime's current design represents Cadmus Labs' working hypothesis for how conventional, emerging, and proprietary technologies can be combined to produce those outcomes.

The hypothesis will continue to be tested.

Some technologies may assume larger roles as evidence develops. Others may remain narrow. Implementation choices will change. New capabilities may emerge. Components that fail to justify their complexity should not survive simply because they were part of an earlier design.

The ambition can remain constant while the architecture improves.

Technology follows purpose because the technology is ultimately accountable to what the system is supposed to accomplish.

For BOCA Prime, that standard is straightforward:

Whether it creates a better system.

References and Notes

1. Yaga, Dylan; Mell, Peter; Roby, Nik; Scarfone, Karen. *Blockchain Technology Overview.* National Institute of Standards and Technology, NISTIR 8202, October 2018. DOI: 10.6028/NIST.IR.8202. Provides a technical overview of blockchain as a distributed digital ledger in which cryptographically linked records become tamper-evident and tamper-resistant, supporting the paper's discussion of shared verification and independently checkable system history.

2. World Wide Web Consortium. *Verifiable Credentials Data Model v2.0.* W3C Recommendation, May 15, 2025. Defines a standards-based model for cryptographically secure, privacy-respecting, machine-verifiable credentials involving issuers, holders, and verifiers, including privacy and data-minimization principles relevant to credential presentation across organizational boundaries.

3. National Institute of Standards and Technology. *Privacy-Enhancing Cryptography: Zero-Knowledge Proof.* NIST Cryptographic Standards and Guidelines project. Describes zero-knowledge proofs as cryptographic techniques that can establish the truth of a statement without revealing additional underlying information, supporting the paper's discussion of verification without unnecessary exposure.

4. Richer, Justin; Fenton, James L.; Lefkovitz, Naomi; Temoshok, David; Galluzzo, Ryan; Regenscheid, Andrew; Choong, Yee-Yin. *Digital Identity Guidelines: Federation and Assertions.* National Institute of Standards and Technology, NIST SP 800-63C-4, 2025. DOI: 10.6028/NIST.SP.800-63C-4. Describes identity federation in which credential service providers provide authentication and subscriber attributes to separately administered relying parties.

5. Girijan, Preeji; Boedi, Rizky; Mânica, Scheila; Franco, Ademir. *The Radiographic Diversity of Dental Patterns for Human Identification — Systematic Review and Meta-Analysis. Journal of Forensic and Legal Medicine*, Vol. 95, April 2023, Article 102507. DOI: 10.1016/j.jflm.2023.102507. Reviews nine eligible studies involving 5,700 panoramic radiographs and evaluates the distinctiveness of radiographic dental patterns for evidence-based human identification.

6. Board of Governors of the Federal Reserve System. *What Is the FedNow® Service?* Describes the Federal Reserve's instant-payment infrastructure, through which participating financial institutions can enable individuals and businesses to send and receive payments within seconds, at any time of day and on any day of the year, with funds available for immediate use.

7. Bank for International Settlements. *The Next-Generation Monetary and Financial System.* BIS Annual Economic Report, Chapter III, June 24, 2025. Examines tokenisation and programmable financial infrastructure, including new arrangements in payments and financial markets, while emphasizing the institutional and monetary foundations required for a trustworthy financial system.

8. Appel, Ian; Grennan, Jillian. *Control of Decentralized Autonomous Organizations. AEA Papers and Proceedings*, Vol. 113, May 2023, pp. 182–185. DOI: 10.1257/pandp.20231119. Examines 10,639 proposals across 151 DAOs and finds that control over decisions is frequently concentrated among a small number of entities, supporting the paper's discussion of concentration risk in decentralized governance.

9. Tabassi, Elham. *Artificial Intelligence Risk Management Framework (AI RMF 1.0).* National Institute of Standards and Technology, NIST AI 100-1, January 2023. DOI: 10.6028/NIST.AI.100-1. Establishes a voluntary framework for managing AI risk and identifies multiple characteristics relevant to trustworthy AI, including validity and reliability, safety, security and resilience, accountability and transparency, explainability and interpretability, privacy enhancement, and fairness.

10. Conrad, Tim; Visa and Artemis. *Agentic Payments: What Onchain Data Reveals About Commerce.* Visa, July 14, 2026. Reports findings from Visa and Artemis examining protocols, standards, and live onchain data associated with emerging AI-agent payments, including the use of stablecoins and blockchain infrastructure in agentic commerce.

11. Garratt, Rodney; Shin, Hyun Song. *Stablecoins Versus Tokenised Deposits: Implications for the Singleness of Money.* BIS Bulletin No. 73, April 11, 2023. Examines stablecoins and tokenised deposits, including implications for settlement, redemption and the singleness of money, and discusses the expanded functionality available through programmable ledgers and tokenised deposits.

Publication Status

This Design Paper reflects Cadmus Labs' architectural reasoning for BOCA Prime as of August 2026, drawing upon research and design work developed over the course of the BOCA program.

BOCA Prime remains a proposed future coverage system. The technologies, components, financial mechanisms, governance approaches, identity systems, payment infrastructure, digital-asset functions, and other architectural elements discussed here remain subject to continuing legal, regulatory, actuarial, economic, technical, security, privacy, usability, and real-world validation.

References to technologies or technical capabilities describe areas Cadmus Labs is evaluating for possible application within BOCA Prime. Their inclusion does not represent a commitment that any particular implementation, protocol, vendor, blockchain network, digital asset, governance mechanism, or technical approach will appear in the final system.

The architecture is specific enough to test and open enough to change.
