Bit2: Enabling Agentic Commerce in a Bitcoin-Native Economy

Fairgate

MARCH 2026

Table of Contents

Executive Summary

Autonomous systems are rapidly becoming economic actors. They consume computation, data, energy, and services continuously, coordinate programmatically, and operate at scales and timescales that differ fundamentally from human economic activity.

This shift implies a qualitative change in economic infrastructure requirements. Agentic commerce demands neutral settlement, deterministic enforcement, high-frequency coordination, and verifiable fulfillment. These are properties that existing financial systems and general-purpose blockchains are structurally ill-equipped to provide.

Bitcoin is uniquely positioned to serve as the settlement layer for agentic economies. Its rules are deterministic, its operation is neutral, and its enforcement does not depend on discretionary governance or institutional trust. However, Bitcoin Layer 1 is intentionally constrained and cannot serve as a high-frequency execution environment.

Bit2 introduces a new category of Bitcoin Layer 2 designed explicitly for sovereign agent-scale economic activity. It separates execution, validation, and enforcement: autonomous systems execute and coordinate locally; correctness and fulfillment are verified cryptographically; and Bitcoin enforces outcomes and provides final settlement when required.

By enabling verifiable computation, conditional clearing, configurable disclosure, and irreversible settlement to Bitcoin Layer 1, Bit2 allows agentic economic activity to occur at scale and without expanding Bitcoin's trust model or relying on centralized operators.

Bit2 is not designed as a retail payment network or a generalized execution platform. Its primary purpose is to serve as foundational infrastructure for autonomous markets—initially exemplified by AI compute and agent-to-agent services, but generalizable to a broad class of sovereign agent-native economic interactions. As agentic commerce increases Bitcoin's transactional velocity and adoption, human-facing payment systems built on Bit2 may emerge naturally as a downstream evolution rather than as a primary design objective.

The architecture reflects a progression of Bitcoin Layer 2 research and experience focused on preserving Bitcoin's neutrality while extending its usability. Rather than optimizing for rapid feature expansion, Bit2 prioritizes correctness, durability, and alignment with Bitcoin's long-term role as a global settlement layer.

As agentic economies grow, infrastructure that enables verifiable, enforceable, and scalable economic coordination becomes essential. Bit2 represents a fundamental step extending Bitcoin's enforcement guarantees to enable a future of sovereign agent-native economic activity.

The Emergence of Agentic Economies

Autonomous Systems as Economic Actors

Artificial intelligence systems are transitioning from passive tools into autonomous agents capable of initiating actions, procuring resources, executing tasks, and interacting with other systems without continuous human direction.

These systems already participate in economically meaningful activity. They consume computation, data, bandwidth, energy, storage, and access to specialized services, and they do so continuously rather than episodically.

Unlike traditional software, autonomous systems operate globally, at machine timescales, and under deterministic control loops. As autonomy increases, these systems cease to be merely instruments of human economic activity and instead become economic actors in their own right.

Order-of-Magnitude Scale of Agentic Economies

Agentic commerce differs from human commerce not only in degree, but in kind.

A single large-scale model training run can involve millions of dollars in compute expenditure. Inference and downstream services involve billions of fine-grained interactions, each individually small but economically relevant in aggregate. When multiplied across fleets of agents, infrastructure providers, and specialized services, the number of economic interactions increases by several orders of magnitude.

Agentic economies are characterized by:

- continuous rather than episodic activity,
- fine-grained rather than coarse value transfer,
- and parallel interactions across large populations of agents.

Even conservative assumptions imply transaction volumes and coordination frequencies that substantially exceed those of human-facing payment systems. The challenge is therefore not only throughput, but the ability to coordinate value transfer at machine cadence and granularity.

Why Existing Financial Infrastructure Breaks

Existing financial infrastructure is designed around assumptions that do not hold for autonomous systems: identifiable account holders, jurisdictional identity, business-hour settlement, discretionary reversibility, and institutional arbitration.

Autonomous systems cannot open bank accounts, wait for batch settlement, or rely on legal enforcement processes. They require continuous availability, deterministic settlement behavior, and rules that can be reasoned about programmatically.

This mismatch is structural. Existing systems cannot be adapted to support agentic economies without reintroducing trust, discretion, and coordination overhead that undermine automation.

Why This Pressure Is Emerging Now

The emergence of agentic economies is the result of converging trends:

- rapid advances in autonomous AI systems,
- commoditization and global distribution of compute and data,
- and increasing programmability of economic interactions.

As these trends intersect, the lack of suitable economic infrastructure becomes a binding constraint. The issue is no longer whether autonomous systems will engage in economic activity, but whether infrastructure exists that can support coordination at machine scale.

This constraint defines the problem space addressed in the following section.

Infrastructure Selection for Agentic Economies

Settlement Requirements at Machine Scale

Autonomous systems operate continuously, transact at machine frequency, and cannot rely on discretionary intervention to resolve ambiguity. As a result, settlement infrastructure for agentic economies must be evaluated under stricter constraints than those applied to human-facing systems.

At machine scale, uncertainty introduced by governance processes, emergency intervention, or rule changes cannot be mitigated operationally. Autonomous systems must be able to reason about settlement guarantees purely through protocol behavior.

Infrastructure is therefore evaluated according to three core properties:

- **Determinism**, ensuring rules remain predictable over long time horizons
- **Continuous availability**, without reliance on coordinated intervention
- **Enforceability**, through protocol-level mechanisms rather than institutional authority

These requirements prioritize minimizing long-term dependency risk under adversarial and evolving conditions.

Neutrality as a Functional Requirement

For agentic economies, neutrality is not ideological. It is a functional constraint.

Any settlement system whose operation depends on identifiable authorities, governance bodies, or coordinated intervention introduces uncertainty that autonomous systems cannot model or price. When economic interactions occur continuously and at global scale, this uncertainty becomes unacceptable.

As a result, agentic economies will converge on settlement layers that minimize governance, discretion, and reliance on identifiable operators.

Neutrality in this context means that the rules governing settlement cannot be selectively modified, suspended, or enforced by institutions.

Bitcoin as the Only Credibly Neutral Settlement Layer

Bitcoin uniquely satisfies these requirements.

Its monetary rules are deterministic and enforced by protocol rather than by institutions. Its operation does not depend on foundations, committees, or roadmap-driven governance. Settlement guarantees arise from accumulated proof-of-work and decentralized validation rather than discretionary enforcement.

For autonomous systems, Bitcoin's resistance to change is a critical feature. Stable rules minimize long-term dependency risk and allow systems to operate without embedding assumptions about future human coordination.

This positioning is not philosophical. It follows directly from Bitcoin's design as a settlement system rather than an execution platform.

Why Competing Blockchains Fail the Selection Criteria

Most general-purpose blockchains prioritize global execution, rapid feature evolution, and developer flexibility.

These properties introduce governance-mediated rule changes, shared global state, extractive dynamics, and operational complexity. While these trade-offs may be acceptable for human-driven applications, they introduce unbounded risk for autonomous systems.

From the perspective of sovereign agents, these platforms function primarily as **execution environments**, not neutral settlement layers.

As a result, they cannot serve as the monetary base for large-scale autonomous economies.

Why Bitcoin Layer 1 Is the Right Base and the Wrong Execution Layer

Bitcoin as Settlement, Not Computation

Bitcoin Layer 1 is intentionally constrained in what it can execute and how frequently it can process economic activity.

These constraints define its role in agentic economies: not as a global execution environment, but as a final arbiter of economic outcomes. Bitcoin provides enforcement and irreversibility rather than continuous coordination.

Importantly, these constraints are not limitations to be overcome. They are the source of Bitcoin's credibility as a settlement system.

Settlement Finality as a Coordination Primitive

For autonomous systems, settlement finality is not merely an accounting property. It is a coordination mechanism.

When settlement is deterministic and irreversible, economic actors can coordinate without relying on trust, reputation, or discretionary enforcement. Bitcoin encodes this finality mechanically through protocol rules and accumulated proof-of-work.

Because these guarantees are protocol-enforced rather than institutionally mediated, autonomous systems can reason about them programmatically.

Why Bitcoin Layer 1 Cannot Support Machine-Scale Execution

Machine-scale economic activity requires properties that Bitcoin Layer 1 is not designed to provide:

- High-frequency interaction
- Fine-grained value transfer
- Predictable transaction costs
- Immediate local confirmation

Increasing throughput, reducing confirmation latency, or embedding complex execution logic directly into the base layer would weaken decentralization and long-term security.

This limitation is therefore **structural**, not a temporary scaling problem.

The Necessity of Layering

The correct response to this constraint is not to modify Bitcoin Layer 1, but to build additional layers on top of it.

A layered architecture separates concerns:

- **Bitcoin Layer 1** provides global settlement and enforcement
- **Higher layers** provide execution, coordination, and transaction velocity

This separation preserves Bitcoin's settlement guarantees while enabling the economic activity required for machine-scale systems.

For agentic economies, such layering is not optional. It is the only architecture capable of combining credible neutrality with high-frequency autonomous interaction.

Why Existing Bitcoin Layer 2s Do Not Solve the Agentic Economy Problem

The Constraint Is Not Scaling Alone

Most Bitcoin Layer 2 designs frame the problem narrowly as one of throughput or transaction cost reduction. For machine-scale economic activity, however, the dominant constraint is not average cost but predictability under load. Autonomous systems require settlement paths whose fees and timing do not degrade unpredictably under congestion—a property that shared blockchain execution environments structurally fail to provide.

While scalability is a necessary condition for agentic commerce, it is not sufficient. Autonomous systems impose additional constraints around availability, predictability, trust minimization, and independence from human-operated infrastructure.

Layer 2 solutions that address throughput while reintroducing coordination, governance, or operational dependencies fail to meet these constraints.

Payment Channels and Routed Networks

Payment-channel–based systems move transactions off-chain while relying on pre-allocated liquidity and continuous availability of channel participants.

These designs are effective for certain payment flows but impose structural limitations for autonomous systems: capital must be locked in advance, routing depends on intermediary availability, failures propagate through network topology, and recovery often requires human intervention.

Payment channels optimize for bilateral or routed payments between known participants. They do not generalize well to open-ended, multi-party machine coordination.

Federated Sidechains and External Validator Sets

Federated systems improve performance by delegating validation and execution to a defined group of operators.

Recent designs attempt to reduce explicit trust assumptions through cryptographic enforcement and stronger security models. However, these systems still introduce external execution and availability dependencies that operate outside Bitcoin's consensus.

For autonomous systems, reliance on identifiable operator sets, coordinated availability, or off-chain recovery mechanisms can introduce exposure to censorship, downtime, or regulatory intervention that cannot always be mitigated programmatically.

As a result, architectures intended to support agentic economies will aim to minimize reliance on specific organizations or operational coordination wherever possible. Reducing these dependencies increases robustness under adversarial conditions and improves long-term compatibility with autonomous economic systems.

Rollup-Style Architectures on Bitcoin

Rollup-style designs attempt to import scaling techniques from other blockchain ecosystems.

These architectures typically rely on shared global execution environments, centralized or semi-centralized sequencers, explicit data availability assumptions, and off-chain coordination for correctness and liveness.

While settlement may ultimately occur on Bitcoin, these architectures remain tightly coupled to Bitcoin Layer 1 for data availability. As a result, execution cost and liveness inherit Bitcoin's fee variance under congestion. For machine-scale coordination, this coupling degrades Layer 2 predictability precisely when demand increases, undermining cost stability and timing guarantees required for autonomous operation.

Client-Side Validation: Necessary but Not Sufficient

Client-side validation represents an important conceptual advance by eliminating the need for global execution.

However, early implementations lack robust Bitcoin-enforced outcome guarantees: mechanisms that allow objectively verifiable execution results to be unilaterally enforced or settled on Bitcoin Layer 1 under adversarial conditions. They also lack native support for configurable disclosure and the level of expressiveness required for machine-scale coordination.

The Missing Properties

Across existing Bitcoin Layer 2 approaches, a consistent pattern emerges. Systems either preserve Bitcoin's trust model while remaining too constrained for machine-scale economic activity, or expand functionality by reintroducing trust, coordination, or governance assumptions.

What is missing is a Layer 2 architecture that simultaneously provides local execution aligned with autonomous systems, strong Bitcoin-enforced outcome guarantees, privacy-preserving execution with configurable disclosure, and the ability for participants to exit unilaterally to Layer 1.

At machine scale, unconditional global transparency is not equivalent to verifiability. Autonomous systems require the ability to prove correctness and fulfillment while selectively disclosing information based on policy, counterparties, or regulatory context.

Implication: A New Category of Bitcoin Layer 2 Is Required

Agentic economies do not need Bitcoin to become more expressive at the base layer. They require an execution environment that remains strictly subordinate to Bitcoin's settlement authority.

This implies a Layer 2 design where execution is local, validation is independent, enforcement is Bitcoin-native, and scalability emerges from parallelism rather than shared throughput.

Existing Bitcoin Layer 2s do not satisfy this combination of requirements.

Introducing Bit2: The Sovereign Agent-Compatible Bitcoin Usage Layer

Bit2 in One Sentence

Bit2 is a Bitcoin-native coordination, clearing, and settlement layer that enables agentic economic activity by combining local execution, verifiable computation, and Bitcoin-enforced outcomes.

From Payments to Economic Coordination

Bit2 is not designed as a retail payment network or a generalized execution platform. Its primary purpose is to enable autonomous systems to coordinate economically without relying on trust, human-operated infrastructure, or discretionary enforcement.

In agentic economies, value transfer is not an isolated action. Payments are a consequence of verified fulfillment: the delivery of computation, services, or predefined outcomes. Bit2 is designed around this principle, treating payments as the result of objective verification rather than discretionary authorization.

Verifiable Computation as the Basis for Clearing

A central requirement of agentic commerce is the ability to validate that work has been performed correctly.

Bit2 enables value transfer to be conditioned on verifiable outcomes. Autonomous systems can produce cryptographic proofs attesting to correct execution or delivery of computation and services. These proofs allow counterparties to verify fulfillment without requiring shared execution environments, trusted intermediaries, or full disclosure of underlying data or methods.

Zero-knowledge techniques play a key role in this process, enabling correctness to be proven while selectively controlling what information is revealed. In this way, Bit2 supports verification without global observability.

Local Execution, Global Enforcement

Execution and validation in Bit2 occur locally. Each participant verifies only what is relevant to them, using cryptographic proofs rather than relying on centralized infrastructure or shared global state.

Bitcoin's role is precise and limited: it enforces outcomes. When disputes arise or settlement is required, Bitcoin serves as the final authority, ensuring that predefined economic consequences are applied correctly. This separation allows Bit2 to scale with the number of autonomous agents rather than with global transaction throughput.

Configurable Disclosure and Regulatory Compatibility

Bit2 is designed to support privacy-preserving execution with configurable disclosure. Autonomous systems can selectively reveal information based on policy, counterparty requirements, or regulatory context, without sacrificing verifiability or enforcement guarantees.

This approach recognizes that different agent interactions require different levels of transparency. Bit2 enables auditability and compliance where required, while avoiding unconditional global disclosure that would expose strategies or create unnecessary economic attack surfaces.

What Bit2 Is and Is Not

Bit2 is:

- A coordination and clearing layer for autonomous economic agents
- A Bitcoin-native Layer 2 that preserves Bitcoin's trust model
- An infrastructure layer where payments follow verifiable fulfillment

Bit2 is not:

- A global execution environment
- A custodial system or operator-driven network
- A replacement for Bitcoin's base layer

Transition to Architecture

The effectiveness of Bit2 depends on how its core principles are realized technically.

The next section describes the architectural lineage and system design that allow local execution, verifiable computation, configurable disclosure, and Bitcoin-enforced outcomes to function together coherently—building on more than a decade of Bitcoin Layer 2 experimentation and research.

Architecture Overview: From Bitcoin Settlement to Machine-Scale Coordination

Architectural Lineage and Design Philosophy

Bit2 is not a clean-slate design. Its architecture emerges from successive generations of Bitcoin Layer 2 research that progressively clarified what *does not* work at scale.

Early federated bridges and sidechains demonstrated that richer execution environments could be anchored to Bitcoin, but at the cost of expanded trust assumptions and operational dependencies. Payment-channel networks showed that off-chain execution could scale, but revealed coordination and liquidity constraints. Client-side validation introduced locality and scalability, yet initially lacked strong enforcement and expressiveness.

Bit2 synthesizes these lessons into a single design philosophy: **execution must be local, validation must be independent, and enforcement must remain Bitcoin-native**.

Separation of Execution, Validation, and Enforcement

A core architectural principle of Bit2 is the explicit separation of responsibilities across layers.

- **Execution** occurs off-chain and locally, within the context of the autonomous systems participating in an interaction.
- **Validation** is performed independently by participants using cryptographic proofs, without reliance on shared global state or centralized infrastructure.

- **Enforcement** is delegated to Bitcoin, which serves as the final arbiter when economic outcomes must be settled or disputes resolved.

This separation allows Bit2 to scale horizontally while remaining strictly subordinate to Bitcoin's settlement authority.

Local Execution and Client-Side Validation

Client-side validation (CSV), originally proposed by Peter Todd, enables Bitcoin-based systems to scale without requiring global consensus on application state. Instead of validating all activity on-chain or within a shared execution environment, participants independently verify only the state transitions relevant to them.

Bit2 adopts this model for autonomous economic systems. Participants execute interactions locally and exchange only the data required for verification.

Because validation occurs independently, the system does not require global consensus on the intermediate state. This eliminates congestion, extractive ordering dynamics, and shared failure modes associated with shared execution environments.

By limiting validation scope to economically relevant state, Bit2 enables parallelism that scales with the number of interacting agents rather than with network-wide throughput.

Verifiable Computation and Selective Cryptographic Proofs

For agentic commerce, economic claims must be tied to objective fulfillment rather than trust or attestation.

Bit2 enables this by supporting mechanisms through which participants can demonstrate that services were performed according to predefined rules or that measurable outcomes were achieved. Verification may rely on different approaches depending on the nature of the service, including deterministic checks, cryptographic commitments, external attestations, or cryptographic proofs.

Zero-knowledge techniques can play an important role when verification must occur without revealing sensitive information. They allow correctness to be proven without disclosing proprietary data, internal models, or full execution traces. However, because such proofs introduce computational overhead, their use is reserved for cases where privacy or trust minimization justifies the cost.

By supporting selective disclosure and multiple verification methods, Bit2 enables participants to choose verification mechanisms appropriate to the economic context.

These verifiable outcomes form the basis for clearing: value transfer can be conditioned on objective evidence of fulfillment rather than discretionary authorization.

Bitcoin-Enforced Outcomes and Unilateral Exit

While execution and validation occur off-chain, Bit2's security model ultimately relies on Bitcoin for enforcement.

Using Bitcoin-enforced programmability, predefined outcome conditions can be asserted and, if necessary, enforced on Bitcoin Layer 1. This ensures that economic consequences can be applied correctly even under adversarial conditions.

A key design objective of Bit2 is enabling unilateral exit: the ability for participants to settle or recover their positions directly on Bitcoin without relying on trusted operators or cooperative counterparties. Achieving this property for complex off-chain interactions remains an active area of research.

Accordingly, Bit2's enforcement mechanisms are designed with the goal of making settlement conditions verifiable and, where possible, enforceable through Bitcoin itself rather than through off-chain authorities. Fully trustless unilateral exit represents a long-standing goal of decentralized systems, and Bit2 aims to advance toward this property as the architecture evolves.

Configurable Disclosure and Auditability

Bit2 treats disclosure as a policy decision rather than a protocol mandate.

Participants can selectively reveal information to counterparties, auditors, or regulators while preserving cryptographic verifiability of outcomes. This enables regulatory compatibility and auditability without imposing unconditional global observability.

By decoupling verification from disclosure, Bit2 avoids the false dichotomy between transparency and confidentiality that characterizes many existing systems.

Architectural Coherence

Taken together, these components form a coherent system in which:

- sovereign agents execute and coordinate locally,
- correctness is proven cryptographically,
- economic outcomes are enforced by Bitcoin,
- and scalability emerges from parallelism rather than shared execution.

This architecture reflects a progression from earlier Bitcoin Layer 2 approaches toward a design explicitly suited for autonomous economic activity, while remaining fully aligned with Bitcoin's role as a neutral settlement layer.

Agent-to-Agent Use Cases

The emergence of autonomous systems capable of economic coordination creates a new class of machine-native markets. These markets arise wherever agents must acquire resources, services, or capabilities in order to complete tasks on behalf of users or other systems.

Agentic commerce can broadly be divided into two categories depending on the nature of the counterparties involved.

Agent-to-Web2 interactions occur when autonomous agents consume services provided by existing digital infrastructure. These include APIs, cloud services, data providers, AI compute platforms, and other conventional online services. In these interactions, agents participate in traditional service markets but require mechanisms for automated payment, verification of delivery, and programmatic coordination.

Agent-to-Agent interactions arise when autonomous systems transact directly with one another. In this environment, agents provide specialized services—such as compute, data access, optimization, or domain-specific capabilities—to other agents. These markets operate continuously, often without persistent relationships or shared trust assumptions, and require mechanisms for objective verification, conditional payment, and robust settlement.

Bit2 provides infrastructure that supports both interaction models. In Agent-to-Web2 scenarios, it enables autonomous systems to interact economically with existing digital services. In Agent-to-Agent markets, it provides the primitives required for sovereign agents to coordinate directly and form decentralized service economies.

The following sections illustrate several representative use cases enabled by this architecture.

(1) AI Compute as a Verifiable Economic Service

AI compute represents one of the earliest domains where agentic commerce may emerge at scale. It combines high capital intensity, global distribution, continuous demand, and measurable outputs—properties that make it economically significant and, in some cases, technically amenable to verification.

Training and inference workloads increasingly involve heterogeneous providers of compute, data, models, and specialized acceleration. These providers interact programmatically, often across jurisdictions and without preexisting trust relationships. In such environments, economic coordination benefits from mechanisms that allow participants to verify aspects of service delivery and link compensation to observable outcomes.

While fully cryptographic verification of arbitrary computation remains computationally expensive, practical approaches can verify specific properties of execution—such as resource usage, completion of defined workloads, or adherence to agreed performance bounds. As verification technologies improve, a broader range of compute services may become economically verifiable.

From Attestation to Proof

In conventional systems, delivery of compute services is typically validated through attestations, usage reports, or trusted intermediaries. These mechanisms rely on trust in providers or platforms and often introduce reconciliation delays, counterparty risk, and discretionary dispute resolution.

In an agentic environment, participants benefit from stronger forms of verification. Rather than relying solely on attestations, service delivery can be supported by verifiable evidence tied to predefined service conditions.

Depending on the nature of the workload, verification may rely on different mechanisms. These include deterministic checks of observable outputs, cryptographic commitments to execution traces, hardware-backed attestations, or—when economically justified—cryptographic proofs of correct execution.

Zero-knowledge techniques are one possible tool in this spectrum. They can allow providers to prove compliance with agreed execution conditions without disclosing proprietary models, training data, or internal execution details. However, because such proofs introduce computational overhead, their use is most appropriate in scenarios where privacy or trust minimization justifies the cost.

By enabling multiple forms of verification, Bit2 allows participants to select mechanisms that balance assurance, privacy, and computational efficiency.

Clearing and Conditional Payment

Once a proof of correct execution is produced, payment is not authorized by an intermediary but cleared automatically based on predefined conditions.

Bit2 enables value transfer to be conditioned on the acceptance of valid proofs. Autonomous agents verify proofs locally and release payment only when contractual conditions are satisfied. This transforms payment from a discretionary act into a deterministic consequence of verified delivery.

At machine scale, this model supports fine-grained compensation, including micropayments, streaming payments, and incremental settlement aligned with ongoing computation rather than coarse, batch-based billing cycles. In many cases, small payments correspond to immediately observable or economically useful results—such as an inference response or data query—allowing agents to validate delivery directly without expensive verification mechanisms. More advanced cryptographic proofs can be used selectively for higher-value interactions or when stronger guarantees or privacy are required.

Settlement and Enforcement

While clearing occurs off-chain, settlement guarantees ultimately rely on Bitcoin.

If disputes arise —such as challenges to proof validity or payment conditions—Bitcoin-enforced outcomes provide a final resolution path. Predefined enforcement mechanisms allow parties to settle claims or exit positions on Bitcoin Layer 1 without reliance on trusted arbitrators or centralized operators.

This ensures that even in adversarial scenarios, economic outcomes remain enforceable under Bitcoin's security model.

Why AI Compute Is a Leading Indicator

AI compute is not a special case but an early manifestation of a broader class of agentic services. Its scale, measurability, and economic importance make it a leading indicator of the requirements future agentic economies will impose.

By enabling verifiable compute delivery, conditional clearing, and Bitcoin-enforced settlement, Bit2 provides infrastructure that generalizes beyond AI compute to other forms of autonomous service exchange, including data provision, model access, bandwidth, storage, and energy-aware coordination.

Implications for Scale

Even modest deployment of this model implies transaction volumes and interaction frequencies that exceed those of human-facing payment systems by orders of magnitude.

This reinforces the necessity of an architecture that scales through local execution, independent validation, and parallel coordination rather than through shared global throughput—precisely the design principles underlying Bit2.

(2) Agent-to-Agent Services and Autonomous Markets

Beyond compute provisioning, autonomous systems increasingly require access to specialized services provided by other agents. These include data acquisition and curation, model inference and fine-tuning, optimization services, monitoring, bandwidth, storage, and domain-specific APIs.

Unlike traditional SaaS markets, these interactions occur programmatically, continuously, and often transiently. Counterparties may interact for seconds or minutes, with no persistent relationship and no shared trust assumptions. Economic coordination in this environment requires objective verification of service delivery and deterministic compensation.

Services as Verifiable Outcomes

In a sovereign agent-native market, services are not validated through reputation systems or manual reconciliation. They are validated through proofs.

Agents offering services can define measurable delivery conditions—such as data freshness bounds, accuracy thresholds, completion of specified tasks, or other observable service properties—and provide verifiable evidence attesting to fulfillment. Depending on the nature of the service, this evidence may take different forms, including deterministic checks of outputs, cryptographic commitments, hardware-backed attestations, or other verification mechanisms.

Certain performance metrics, such as response latency, may require additional measurement infrastructure—such as synchronized clocks or external monitoring—to be evaluated reliably. Consuming agents verify the relevant evidence locally before releasing payment according to predefined service conditions.

Where services involve proprietary methods or sensitive data, zero-knowledge techniques allow agents to prove compliance with service conditions without revealing internal logic or raw inputs. This enables markets where verification is possible without commoditizing intellectual property.

Continuous Clearing and Micropayments

Agent-to-agent services often involve continuous provision rather than discrete jobs. Examples include streaming data feeds, adaptive inference services, or ongoing optimization processes.

Bit2 supports continuous clearing by allowing value transfer to occur incrementally as verified service delivery progresses. Payments can be streamed, metered, or settled in fine-grained increments aligned with objective performance metrics rather than coarse billing cycles.

This model reduces counterparty risk, improves capital efficiency, and enables autonomous systems to coordinate economically in real time.

Settlement, Dispute Resolution, and Exit

As with compute delivery, final settlement and dispute resolution rely on Bitcoin.

If disagreements arise regarding service fulfillment or proof validity, predefined enforcement paths allow agents to escalate to Bitcoin for resolution. Participants retain the ability to exit positions unilaterally and settle claims under Bitcoin’s security model, without reliance on trusted intermediaries or off-chain governance.

This ensures that autonomous markets remain robust even in adversarial conditions.

Generalization to Agentic Economies

Agent-to-agent services illustrate how Bit2 generalizes beyond any single domain.

The same primitives—local execution, cryptographic verification, configurable disclosure, and Bitcoin-enforced outcomes—support a wide range of sovereign agent-native economic interactions. As autonomous systems proliferate, these interactions form dense networks of economic activity that operate continuously and at scale.

In this context, Bit2 functions not as an application layer, but as foundational infrastructure for autonomous markets.

(3) Agent Execution

Beyond acquiring individual services, autonomous systems must execute multi-step economic workflows that involve both **Agent2Web2** and **Agent2Agent** interactions.

Examples include automated trading strategies, AI model orchestration, infrastructure management, and supply-chain coordination. In these scenarios, agents continuously acquire services, evaluate results, and initiate subsequent actions based on predefined policies or learned strategies.

Local Workflow Execution

In Bit2, these workflows are executed locally by each agent rather than within a shared execution environment. Agents maintain their own internal state, evaluate inputs from prior interactions, and determine subsequent actions independently.

This model allows autonomous systems to coordinate complex activity without requiring global consensus on an intermediate state.

Conditional Economic Coordination

Agent workflows frequently involve conditional interactions across multiple counterparties. An agent may acquire data from a provider, use the result to perform computation, and subsequently purchase additional services from other agents or platforms.

Bit2 enables these interactions to be coordinated through verifiable service outcomes and conditional payments. Agents verify evidence of fulfillment locally and release payment or initiate further actions only when predefined conditions are satisfied.

Bitcoin-Anchored Settlement

While execution and coordination occur off-chain, economic outcomes ultimately anchor to Bitcoin.

If disagreements arise regarding fulfillment conditions or payment obligations, predefined enforcement paths allow participants to escalate claims to Bitcoin Layer 1. This provides a secure settlement layer while allowing the vast majority of interactions to occur off-chain.

Through this architecture, Bit2 enables autonomous agents to execute complex economic workflows that combine **Agent2Web2 services, Agent2Agent markets, and Bitcoin-secured settlement**.

Why Fairgate

Continuity Rather Than Reinvention

Bit2 is not the result of a sudden change in direction or an isolated insight. It represents a continuation of a long trajectory of work focused on extending Bitcoin's utility without compromising its trust model.

Over more than a decade, successive generations of Bitcoin infrastructure have explored how to enable richer economic activity while remaining anchored to Bitcoin's settlement guarantees. Each iteration clarified both what was possible and where fundamental limits lay.

Bit2 emerges from this accumulated understanding rather than from first principles alone.

Operating at Bitcoin's Trust Boundary

A recurring theme in Fairgate's work has been operating close to Bitcoin's trust and enforcement boundaries.

This includes early experimentation with federated models to expand functionality, long-term experience with Bitcoin-anchored execution environments, and direct involvement in advancing Bitcoin-enforced programmability. These efforts consistently confronted the same underlying question: how to enable off-chain activity while preserving Bitcoin as the ultimate arbiter.

Bit2 reflects a synthesis of these experiences. Its architecture assumes neither perfect coordination nor benevolent operators. Instead, it treats adversarial conditions, partial failure, and uncooperative participants as baseline design constraints.

From Execution to Enforcement

One of the clearest lessons from prior Bitcoin Layer 2 efforts is that scalability does not come from moving execution elsewhere alone, but from minimizing what must be executed and who must be trusted.

Fairgate's work on Bitcoin-enforced computation frameworks helped clarify a critical distinction: Bitcoin does not need to execute complex logic, but it can enforce the outcomes of such logic when correctness can be proven.

This insight directly informs Bit2's design. Local execution, cryptographic verification, and Bitcoin-enforced outcomes are not independent ideas; they are the product of iterative refinement across multiple systems and generations of research.

Research-Driven Engineering

Bit2 reflects an approach in which research and engineering are tightly coupled.

Design choices are informed by formal models, adversarial analysis, and empirical experience operating real systems, rather than by abstract optimization or short-term market considerations. This bias toward correctness and durability over rapid feature expansion is intentional.

As a result, Bit2 is designed as infrastructure meant to persist, evolve conservatively, and remain compatible with Bitcoin's long-term trajectory rather than as a platform optimized for near-term experimentation.

Implicit Alignment

The alignment between Fairgate and Bit2 is not expressed through ownership or control, but through continuity of problem focus.

Bit2 addresses the same core challenge that has guided prior work: enabling economic activity at scale while preserving Bitcoin's neutrality, enforceability, and resistance to discretionary change.

This continuity reduces execution risk. Bit2 is not an attempt to retrofit Bitcoin to a new use case, but a natural extension of a line of inquiry that has been refined over many years.

Transition to Open Questions

Despite this continuity, Bit2 does not assume that all challenges are solved.

The next section outlines open questions, risks, and areas of active research that remain as agentic economies and Bitcoin-native infrastructure continue to evolve.

Risks, Open Questions, and Outlook

Technical Maturity and Research Risk

Bit2 builds on emerging techniques in client-side validation, cryptographic proof systems, and Bitcoin-enforced programmability. While these components are grounded in active research and early implementations, their composition at machine scale introduces open questions.

Key areas of ongoing research include proof generation efficiency for large-scale computation, enforcement mechanisms under worst-case adversarial conditions, and the operational boundaries of unilateral exit guarantees. Progress in these areas is expected to be incremental and conservative, favoring correctness over rapid expansion.

Operational and Adoption Risk

Agentic economies will not emerge uniformly. Adoption is likely to be uneven across domains, driven initially by use cases where verification and automation provide clear advantages over existing systems.

Integrating Bit2 into real-world sovereign agent workflows requires tooling, standards, and interfaces that do not yet exist at scale. Adoption therefore depends not only on protocol design, but on ecosystem development and gradual operational validation.

Regulatory and Disclosure Considerations

While Bit2 is designed to support configurable disclosure and auditability, regulatory frameworks governing autonomous systems and agentic commerce remain underdeveloped.

Future requirements around reporting, compliance, or transparency may influence how Bit2 is used in practice. The architecture is intended to accommodate such constraints without undermining verifiability or enforcement, but regulatory evolution remains an external uncertainty.

Economic and Systemic Dynamics

Large-scale agentic economies may introduce new economic dynamics, including feedback loops, concentration effects, or resource contention across compute, energy, and data markets.

Bit2 does not attempt to preempt these dynamics through protocol-level policy. Instead, it provides neutral infrastructure within which such interactions can be observed, analyzed, and, where appropriate, constrained by higher-layer rules or external governance.

Outlook

Despite these open questions, the direction is clear.

As autonomous systems proliferate, the absence of neutral, enforceable, and sovereign agent-compatible economic infrastructure will become a limiting factor. Systems that rely on trust, discretion, or global execution will encounter scaling and coordination failures.

Bit2 represents an attempt to address this gap conservatively, by extending Bitcoin's settlement guarantees into a domain where sovereign agents, rather than humans, are the primary economic actors.