Global Cloud Token

Tokenomics White Paper — Version 1.3

Global Cloud Token · gct.to

February 2026

Table of Contents

1. Executive Summary . . . 3
2. The Problem . . . 3
   2.1 The Agent Economy Is Infrastructure-Starved . . . 3
   2.2 The Credit Card Bottleneck . . . 4
   2.3 Agents Cannot Hold Credit Cards . . . 4
   2.4 Existing Crypto Compute Projects Miss the Mark . . . 4
3. The Solution . . . 5
   3.1 GCT + Mosaic: The Agent Infrastructure Layer . . . 5
   3.2 Full Service Catalog . . . 5
   3.3 Why Solana . . . 8
4. Token Mechanics . . . 9
   4.1 Supply . . . 9
   4.2 Allocation . . . 9
   4.3 Vesting Schedules . . . 9
   4.4 Staking Reward Emission Schedule . . . 11
   4.5 Human Authorization Smart Contract . . . 12
   4.6 Conversational Commerce — Approve from Your Messaging App . . . 12
   4.7 GCT-Native Gas (Gas Kitty) . . . 14
5. Compute Pricing Model . . . 15
   5.1 Design Principles . . . 15
   5.2 Base Pricing Formula . . . 15
   5.3 Staking Discount Tiers . . . 15
   5.4 Spot Pricing . . . 16
   5.5 Indicative Pricing (at TGE) . . . 16
   5.6 Prepaid Compute Credits . . . 16
6. Agent Wallet Architecture . . . 17
   6.1 Overview . . . 17
   6.2 Three Wallet Types . . . 17
   6.3 Wallet Hierarchy . . . 18
   6.3 Agent Wallet Program . . . 19
   6.4 Programmatic Spending Flow . . . 19
   6.5 Multi-Agent Coordination . . . 19
7. Deflationary Mechanics . . . 20

7.1 Burn on Compute Purchase . . . 20
7.2 Projected Burn . . . 20
7.3 Revenue Buyback and Burn . . . 20
7.4 Burn Floor . . . 21
7.5 Burn Safeguards . . . 21
8. Governance . . . 22
8.1 Path to Decentralization . . . 22
8.2 Voting Mechanics . . . 22
8.3 Proposal Parameters . . . 22
9. Security . . . 23
9.1 Smart Contract Security . . . 23
9.2 Multi-Sig Configuration . . . 23
9.3 Oracle Security . . . 23
9.4 Infrastructure Security . . . 24
10. Roadmap . . . 24
Phase 0 — Foundation (Q1 2026) . . . 24
Phase 1 — Devnet Launch (Q2 2026) . . . 24
Phase 2 — Mainnet Launch (Q3 2026) . . . 24
Phase 3 — Ecosystem Growth (Q4 2026 – Q2 2027) . . . 24
Phase 4 — DAO & Marketplace (Q3 2027 – Q4 2027) . . . 25
Phase 5 — Full Decentralization (2028+) . . . 25
11. Team & Backing . . . 25
11.1 Founder . . . 25
11.2 Backing Infrastructure . . . 25
11.3 Why This Matters . . . 26
12. Legal Disclaimers . . . 26
12.1 Token Classification . . . 26
12.2 Risk Factors . . . 26
12.3 Forward-Looking Statements . . . 26
12.4 Jurisdictional Restrictions . . . 27

---

1. Executive Summary

Global Cloud Token (GCT) is a Solana-based SPL utility token designed to serve as the native unit of account for purchasing cloud infrastructure, compute resources, and telecommunications services across the Mosaic platform (mosaic.site). GCT enables autonomous AI agents — and the humans and organizations that deploy them — to programmatically acquire, pay for, and manage a full spectrum of infrastructure resources without intermediary approval, credit checks, or fiat payment rails.

The token is backed by real, operating infrastructure: data centers with significant GPU and general-purpose compute capacity, the Mosaic cloud platform, and Fibernetics — a licensed Canadian telecommunications provider. This is not a speculative protocol seeking product-market fit. GCT tokenizes access to infrastructure that already exists, already serves customers, and already generates revenue.

Key parameters:

Parameter	Value
Token Name	Global Cloud Token
Ticker	GCT
Public Equity Ticker	GCT.TO
Blockchain	Solana (SPL Token)
Total Supply	1,000,000,000 (1 billion)
Supply Model	Fixed supply, deflationary via burn
Decimal Precision	9 (Solana standard)

GCT solves a concrete problem: AI agents are the fastest-growing category of compute consumers, yet they have no native, permissionless mechanism to acquire infrastructure. Every existing cloud provider requires a human with a credit card, an identity verification process, and a billing agreement. GCT removes these barriers entirely. An agent with GCT in a Solana wallet can provision a GPU cluster, spin up a database, register a phone number, or deploy a container — all through on-chain transactions settled in under 400 milliseconds.

---

2. The Problem

2.1 The Agent Economy Is Infrastructure-Starved

The number of autonomous AI agents in production is growing exponentially. By early 2026, millions of agents operate across customer service, code generation, research, trading, content creation, and enterprise automation. Each of these agents requires infrastructure: compute to run on, storage to persist data, networking to communicate, and increasingly, telecommunications services to interact with the physical world.

Yet acquiring this infrastructure remains a fundamentally human process.

2.2 The Credit Card Bottleneck

Every major cloud provider — AWS, GCP, Azure, and their downstream resellers — requires:

1. **Human identity verification** — government-issued ID, corporate registration, or KYC documentation
2. **Credit instruments** — credit cards, wire transfer agreements, or enterprise procurement contracts
3. **Billing cycles** — monthly invoices, net-30/60/90 payment terms, manual reconciliation
4. **Approval workflows** — human-in-the-loop provisioning for high-value resources (GPU clusters, dedicated servers, telecom services)

For a single developer managing a handful of agents, this friction is tolerable. For an ecosystem of thousands of agents — each needing to independently acquire, scale, and release resources in real time — it is a structural bottleneck.

2.3 Agents Cannot Hold Credit Cards

An AI agent is not a legal entity. It cannot sign contracts, pass KYC, or hold a credit card in its own name. Today, every agent's infrastructure is provisioned through a human or corporate proxy. This creates:

- **Scaling friction** — provisioning new resources requires human approval at every step
- **Counterparty risk** — the agent depends entirely on its operator's continued willingness and ability to pay
- **Latency** — what should be a sub-second infrastructure decision becomes a minutes-to-hours human workflow
- **Lack of composability** — agents cannot delegate infrastructure budgets to sub-agents without complex, fragile API key sharing

2.4 Existing Crypto Compute Projects Miss the Mark

Several blockchain-based compute projects exist (Akash, Render, Golem, io.net). They share common limitations:

- **No integrated telecom** — none offer phone numbers, SIP trunking, or SMS as part of the compute layer
- **Limited service breadth** — most focus narrowly on GPU rental or container hosting, ignoring databases, storage tiers, CI/CD, monitoring, and the dozens of services a real application needs
- **No real infrastructure backing** — most operate as marketplace protocols matching anonymous providers with consumers, offering limited reliability guarantees
- **Poor agent ergonomics** — designed for human users clicking through web interfaces, not for agents making programmatic Solana transactions

GCT is different because it is backed by owned and operated infrastructure, offers a full-stack service catalog, and is designed from the ground up for agent-native purchasing.

---

3. The Solution

3.1 GCT + Mosaic: The Agent Infrastructure Layer

GCT is the native payment token for the Mosaic cloud platform (mosaic.site). Mosaic is a full-stack infrastructure platform offering compute, storage, networking, databases, GPU clusters, developer tools, and telecommunications services — all available for purchase with GCT via Solana transactions.

The architecture is straightforward:

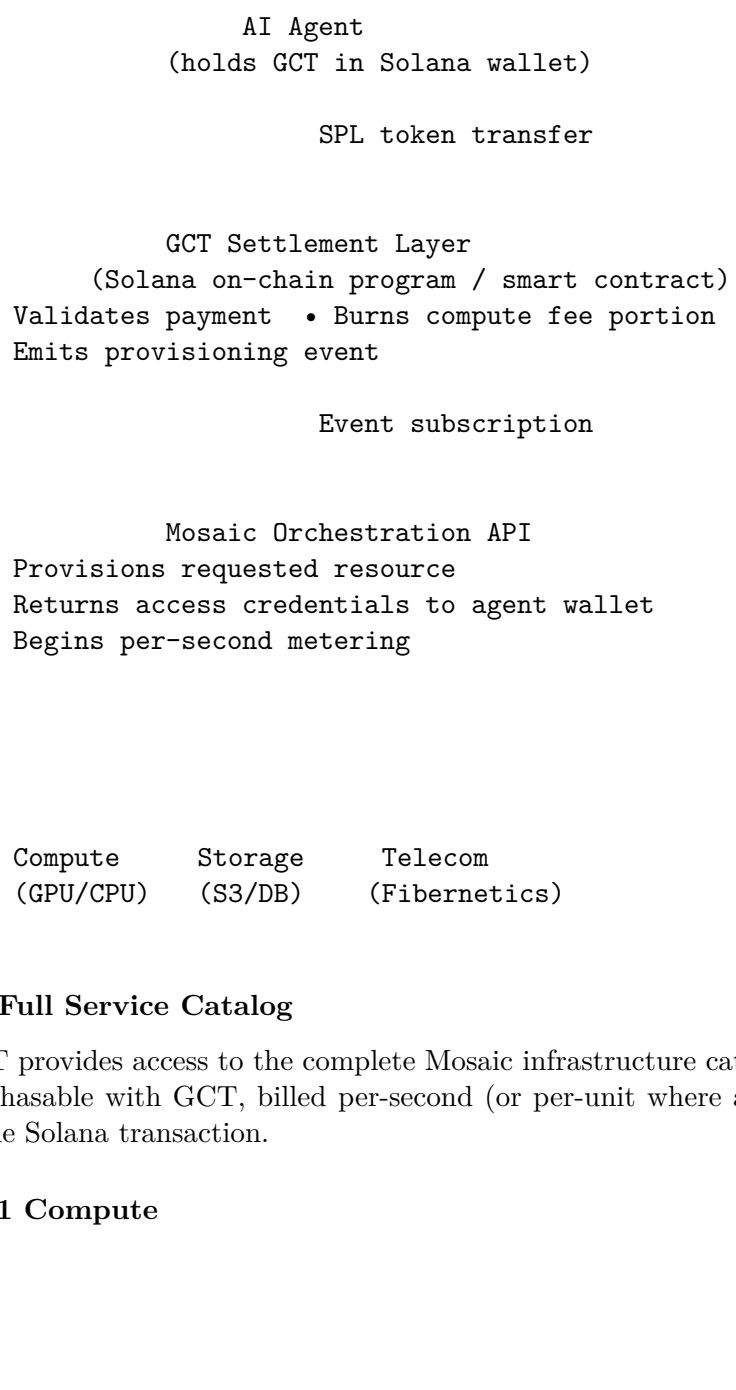

3.2 Full Service Catalog

GCT provides access to the complete Mosaic infrastructure catalog. Every resource listed below is purchasable with GCT, billed per-second (or per-unit where applicable), and provisionable via a single Solana transaction.

3.2.1 Compute

Resource	Description	Billing Unit
Containers	Docker containers, 0.25–128 vCPU, 256 MiB–512 GiB RAM	per-second
Virtual Machines	Full VMs, KVM-based, Linux/Windows	per-second
Bare Metal Servers	Dedicated physical servers, custom configs	per-hour
Kubernetes Clusters	Managed K8s, auto-scaling node pools	per-second (per node)
Serverless Functions	Event-driven execution, 0 to scale	per-invocation + per-ms

3.2.2 GPU & AI

Resource	Description	Billing Unit
GPU Instances (A100)	NVIDIA A100 80GB, single or multi-GPU	per-second
GPU Instances (H100)	NVIDIA H100 SXM5, NVLink interconnect	per-second
GPU Instances (L40S)	NVIDIA L40S, inference-optimized	per-second
Inference Endpoints	Managed model serving, auto-scaling	per-request + per-second
Training Jobs	Distributed training, checkpoint management	per-GPU-second

3.2.3 Storage

Resource	Description	Billing Unit
Object Storage (S3-compatible)	Unlimited capacity, multi-region replication	per-GiB-month + per-request
Block Storage	High-IOPS SSD volumes, attachable to VMs/containers	per-GiB-month
Archive Storage	Cold storage, retrieval within 1–12 hours	per-GiB-month

3.2.4 Databases

Resource	Description	Billing Unit
PostgreSQL	Managed, HA, point-in-time recovery	per-second
MySQL	Managed, read replicas, automated backups	per-second
Redis	In-memory, clustered, persistence options	per-second

Resource	Description	Billing Unit
MongoDB	Document store, sharded, managed	per-second
ClickHouse	Columnar analytics, high-throughput ingest	per-second

3.2.5 Telecommunications (via Fibernetics)

Resource	Description	Billing Unit
Phone Numbers (DID)	Local, toll-free, international numbers	per-number-month
PBX Service	Hosted PBX, IVR, call routing, voicemail	per-extension-month
SIP Trunking	Enterprise-grade SIP, concurrent call channels	per-channel-month + per-minute
SMS / MMS	Programmable messaging, long code and short code	per-message
WebRTC	Browser/app-based real-time communication	per-session-minute

3.2.6 Networking

Resource	Description	Billing Unit
Bandwidth	Egress data transfer	per-GiB
CDN	Global edge caching, custom origins	per-GiB + per-request
Load Balancers	Layer 4/7, health checks, auto-scaling	per-second
DNS	Managed zones, anycast, low-latency resolution	per-zone-month + per-query
Custom Domains	Domain registration, WHOIS privacy	per-domain-year
SSL/TLS Certificates	Auto-provisioned, managed renewal	included with domain
Static IPs	Dedicated IPv4 addresses	per-IP-month
DDoS Protection	Always-on volumetric and application-layer mitigation	per-resource-month

3.2.7 Developer Tools

Resource	Description	Billing Unit
Git Hosting	Private repositories, branch protection, webhooks	per-repo-month
CI/CD Pipelines	Build minutes, parallel runners, caching	per-build-minute

Resource	Description	Billing Unit
Container Registry	Private Docker/OCI image storage	per-GiB-month
Dev Environments	Cloud-based development workspaces	per-second

3.2.8 Observability & Security

Resource	Description	Billing Unit
Logging	Centralized log ingestion, search, retention	per-GiB ingested
Metrics	Time-series collection, dashboards, alerting	per-series-month
Alerting	Threshold and anomaly-based alerts, multi-channel notification	included with metrics
Uptime Monitoring	HTTP/TCP/ICMP checks, global probe network	per-monitor-month
Secrets Management	Encrypted key-value store, rotation, audit logs	per-secret-month
Agent Identity & Auth	Solana wallet-based identity, RBAC, API key issuance	per-agent-month

3.2.9 Email & Messaging

Resource	Description	Billing Unit
Email Hosting	Custom domain mailboxes, IMAP/SMTP	per-mailbox-month
Transactional Email	API-driven sending, deliverability monitoring	per-email
Webhooks	Inbound/outbound HTTP event delivery	per-delivery

3.3 Why Solana

GCT is built on Solana for specific technical reasons:

1. **Finality speed** — Solana achieves transaction finality in ~400ms. An agent can pay for and receive infrastructure credentials in under one second.
2. **Transaction cost** — Solana transactions cost fractions of a cent. Micro-payments for per-second compute billing are economically viable.
3. **Throughput** — Solana processes thousands of transactions per second, supporting a marketplace of millions of agents transacting simultaneously.
4. **SPL token standard** — Mature, well-audited token standard with broad wallet and exchange support.

5. **Program composability** — Solana programs (smart contracts) can be composed, allowing third-party agent frameworks to integrate GCT spending natively.

---

4. Token Mechanics

4.1 Supply

GCT has a fixed total supply of **1,000,000,000** (one billion) tokens. No minting function exists in the token contract. The mint authority is permanently revoked after initial token generation. Supply can only decrease over time through the burn mechanism described in Section 7.

4.2 Allocation

Allocation	Tokens	Percentage	Purpose
Compute Reserve (Mining Pool)	400,000,000	40%	Proof of Compute mining rewards for infrastructure providers
Ecosystem & Partnerships	150,000,000	15%	Integrations, developer grants, agent platform partnerships
Team & Founders	150,000,000	15%	Core team compensation and retention
Liquidity	100,000,000	10%	DEX liquidity pools and market-making operations
Community & Rewards	100,000,000	10%	Staking rewards, early adopters, airdrops
Treasury	100,000,000	10%	Operational reserve for unforeseen needs

4.3 Vesting Schedules

4.3.1 Compute Reserve — Proof of Compute Mining (400M) The 400M Compute Reserve is a **mining pool** — tokens are not sold, they are **earned** by providers who contribute infrastructure to the GCT network. This is **Proof of Compute** mining.

How It Works

Providers contribute real infrastructure — GPUs, VMs, storage, phone numbers, bandwidth, and other resources — to the GCT network. They earn GCT tokens proportional to the useful compute work they serve to agents and users. This is not proof-of-work hash mining. There is no wasted energy solving arbitrary puzzles. Instead, real agent workloads validate every contribution. If an agent provisions a GPU instance, runs a database query, or sends an SMS through a provider's infrastructure, that provider earns GCT.

Verification Mechanism: Agents that consume resources sign on-chain attestations confirming that work was delivered. Providers cannot fake mining because real workloads validate every contribution. In disputed cases, staked arbiters review attestation evidence and resolve conflicts. This creates a trustless, verifiable proof of useful work.

Two-Sided Marketplace: Agents **spend** GCT to purchase compute and infrastructure. Providers **earn** GCT by serving that demand. This creates a real circular token economy where GCT flows from consumers to producers, backed by actual infrastructure delivery. Anyone with a server can participate — the network starts with GCT's own datacenter infrastructure, then opens to third-party providers.

Sub-Allocation of the 400M Mining Pool:

Sub-Allocation	Tokens	% of Reserve	Purpose
Compute Mining Rewards	280,000,000	70%	Earned by providers for serving infrastructure
Early Provider Bonus	60,000,000	15%	Boosted rewards for early providers (decays over time)
Uptime & Quality Bonuses	60,000,000	15%	Rewards for 99.99% uptime, low latency, high reliability

- **Compute Mining Rewards (280M):** The core mining pool. Providers earn tokens proportional to the compute work they serve, measured in GCT-equivalent resource units consumed by agents.
- **Early Provider Bonus (60M):** Bootstrapping incentive. The first providers to join the network receive boosted reward multipliers that decay over time as the provider base grows. This incentivizes early supply-side participation when the network needs it most.
- **Uptime & Quality Bonuses (60M):** Providers that maintain 99.99%+ uptime, sub-millisecond latency, and high reliability scores earn additional GCT from this pool. Quality of service directly translates to additional earnings.

Emission Curve

Mining rewards follow a halving-style emission over approximately 10 years. Rewards are high early to bootstrap provider supply and taper over time as the network matures — similar to Bitcoin's halving schedule, but tied to actual compute served rather than arbitrary block intervals.

Period	Annual Mining Emission (est.)	Cumulative Emitted
Year 1	80,000,000	80,000,000
Year 2	60,000,000	140,000,000
Years 3–4	40,000,000/yr	220,000,000
Years 5–6	30,000,000/yr	280,000,000
Years 7–10	30,000,000/yr (quality + uptime pools)	400,000,000

Emissions are demand-responsive: if compute demand is low, fewer tokens are emitted (providers

earn less because less work is served). If demand surges, emission accelerates toward the cap but never exceeds it. The total 400M pool is a hard cap.

4.3.2 Ecosystem & Partnerships (150M) Released over 36 months: - **Months 1–6:** 25,000,000 (initial integration incentives) - **Months 7–18:** 50,000,000 (developer grants, agent framework partnerships) - **Months 19–36:** 75,000,000 (strategic partnerships, ecosystem growth)

All ecosystem disbursements require multi-sig approval (see Section 9).

4.3.3 Team & Founders (150M)

- **1-year cliff:** No tokens vest during the first 12 months
- **Linear vesting:** After the cliff, tokens vest linearly over the remaining 36 months (4,166,667 tokens per month)
- **Full vesting:** 48 months from token generation event (TGE)

Team tokens are held in a time-locked Solana program. The vesting contract is immutable once deployed. Vested but unclaimed tokens remain in the contract until explicitly claimed.

4.3.4 Liquidity (100M)

- **At TGE:** 50,000,000 tokens deployed to Raydium and Orca liquidity pools, paired with SOL and USDC
- **Months 1–12:** 30,000,000 tokens available for market-making operations
- **Months 13–24:** Remaining 20,000,000 tokens for additional exchange listings and pool depth

4.3.5 Community & Rewards (100M)

- **Staking Rewards:** 60,000,000 tokens distributed over 48 months (declining emission schedule — see Section 4.4)
- **Early Adopter Airdrops:** 20,000,000 tokens for first 10,000 agent wallets that transact on the platform
- **Community Programs:** 20,000,000 tokens for bug bounties, content creation, ambassador programs

4.3.6 Treasury (100M) Held in a 4-of-7 multi-sig wallet. Disbursements require a governance proposal and vote once the DAO is established (see Section 8). Until DAO activation, the multi-sig signers control treasury disbursements. Treasury tokens have no vesting schedule but are subject to a self-imposed quarterly spending cap of 5,000,000 tokens.

4.4 Staking Reward Emission Schedule

Staking rewards follow a halving-inspired declining emission:

Year	Annual Emission	Monthly Emission	APY (est. at 30% staked)
Year 1	24,000,000	2,000,000	~27%
Year 2	18,000,000	1,500,000	~18%
Year 3	12,000,000	1,000,000	~10%
Year 4	6,000,000	500,000	~4%

Total staking emissions: 60,000,000 GCT over 4 years. After Year 4, staking rewards transition to a fee-sharing model funded by a portion of compute transaction fees.

4.5 Human Authorization Smart Contract

GCT supports two distinct transaction modes, giving agents and their operators the flexibility to choose the right level of oversight:

Autonomous Mode (Default) In autonomous mode, agents spend GCT directly from their wallets with no intermediary. Transactions settle on Solana in ~400ms at near-zero cost (standard Solana transaction fees only). This is the default mode for agents operating within pre-set budgets and allowlists. It is permissionless, fast, and frictionless — ideal for high-frequency, low-value infrastructure purchases.

Human-Authorized Mode (Guardian Smart Contract) For enterprises, regulated environments, or any scenario where human oversight is required, agents can be configured to route transactions through the **Guardian Authorization Smart Contract**. In this mode:

1. The agent submits a purchase intent (resource type, quantity, cost) to the Guardian contract
2. The contract notifies the designated human approver (via the Guardian Wallet — see Section 6)
3. The human reviews and approves or denies the transaction
4. If approved, the contract executes the GCT transfer and provisions the resource
5. If denied, the transaction is cancelled and no GCT is spent

Human-authorized transactions incur a **1% authorization fee**, distributed to ledger and validator nodes that process and mediate the approval workflow. This fee compensates the network for the additional computation and state management required by the authorization flow.

Mode	Cost	Latency	Use Case
Autonomous	Near-zero (Solana base fees)	~400ms	High-frequency agent spending within budgets
Human-Authorized	1% authorization fee	Seconds to hours (depends on human response)	Enterprise oversight, large purchases, compliance

The two modes are not mutually exclusive. An agent can be configured to operate autonomously for purchases under a threshold (e.g., <100 GCT) and require human authorization for larger amounts. This hybrid approach provides the best of both worlds: speed for routine operations and control for significant expenditures.

4.6 Conversational Commerce — Approve from Your Messaging App

4.6.1 The Problem with Traditional Approval UIs Most crypto and DeFi approval flows require a human to connect a wallet to a dApp, open a dedicated website, or launch a separate application. This creates friction and latency at the worst possible moment — when an agent needs a resource *now*. The human has to context-switch away from whatever they are doing, navigate to an unfamiliar interface, authenticate, review, and confirm. For a Guardian overseeing multiple agents across different tasks, this approval tax compounds rapidly.

4.6.2 The GCT Solution — Messaging-Native Approvals GCT integrates with AI agent runtimes like **OpenClaw**, which operate natively inside messaging platforms — Telegram, Signal, Discord, Slack, and others. When an agent requests a purchase that exceeds the Guardian's auto-approve threshold, the Guardian receives an **inline notification in their existing chat** — the same conversation where the agent already reports its work.

Example flow:

1. An agent determines it needs 85 GCT of GPU time to complete a training job
2. The agent submits a purchase intent to the Guardian Authorization Smart Contract (Section 4.5)
3. The Guardian receives a Telegram message:

```
Agent-7 requests approval:
 Resource: 4x A100 GPU (2 hours)
 Cost: 85 GCT (~$127.50)
 Reason: Fine-tuning embedding model

 [ Approve]  [ Deny]
```

4. The Guardian taps **Approve** — one tap, inline, no context switch
5. The Guardian contract releases funds from escrow
6. Mosaic provisions the GPU cluster
7. The agent continues working — total human interaction time: **under 3 seconds**

There is no app to download, no wallet to connect, no website to visit. The Guardian's wallet **is** the messaging conversation. Approval happens where the human already is — in the same chat thread where the agent communicates its progress, asks questions, and delivers results.

4.6.3 Spending Rules + Conversational Approval Guardians can configure **auto-approve thresholds** to eliminate notification noise for routine purchases:

- **Auto-approve under threshold** — e.g., "auto-approve any purchase under 20 GCT." Small, routine infrastructure purchases (a database query, a brief container spin-up, an SMS) flow automatically without any human involvement.
- **Notify above threshold** — only purchases exceeding the threshold trigger a message to the Guardian. This ensures humans are only interrupted for meaningful spending decisions.
- **Category-based rules** — Guardians can set different thresholds for different resource types (e.g., "auto-approve storage under 50 GCT, but always ask for GPU purchases over 10 GCT").
- **Time-based rules** — optional restrictions like "no auto-approvals after business hours" or "require approval for any purchase on weekends."

Rules are configured via simple chat commands (e.g., `/set-limit 20 GCT`) or through a lightweight dashboard. The system defaults to conservative settings — all purchases require approval until the Guardian explicitly relaxes the rules.

4.6.4 Why This Matters AI agents already live inside messaging platforms. Telegram bots, Discord bots, Slack integrations — the agent's primary interface with its human operator is a chat window. Yet today, the moment money is involved, the interaction jumps to an entirely different context: a wallet app, a browser extension, a dApp website. This is a UX failure.

GCT eliminates this discontinuity. The agent communicates in chat. The agent requests resources in chat. The human approves spending in chat. The agent reports results in chat. **The entire loop stays in one place.**

This is **conversational commerce for the agent economy** — zero friction between "agent requests" and "human approves." It transforms the Guardian role from a burdensome oversight task into a natural part of the ongoing conversation between humans and their agents.

4.7 GCT-Native Gas (Gas Kitty)

A critical UX problem with Solana-based tokens is that every transaction requires a small amount of SOL to pay network fees. For AI agents, managing SOL dust alongside GCT creates unnecessary complexity: agents would need to acquire SOL, monitor SOL balances, and handle SOL top-ups — none of which relates to their actual task of purchasing infrastructure.

The Gas Kitty is a protocol-owned SOL liquidity pool that abstracts away Solana gas fees entirely. Agents never need to hold or manage SOL.

How it works:

1. When an agent submits a GCT transaction, a small **GCT gas surcharge** is added to the transaction amount (fractions of a cent, calibrated to cover Solana network fees plus a small buffer)
2. The Gas Kitty program receives the GCT surcharge and converts it to SOL via on-chain swap (Raydium/Jupiter)
3. The converted SOL pays the Solana network fee for the transaction
4. The agent's transaction is submitted and settled — the agent only ever interacts with GCT

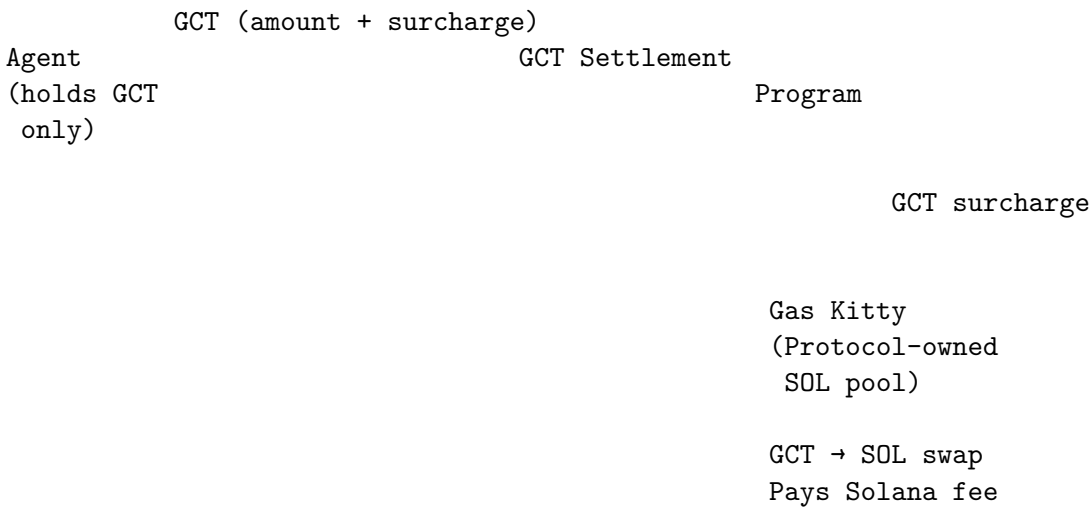

Funding the Gas Kitty:

- **Initial seeding:** Funded from the Treasury allocation (Section 4.3.6)
- **Ongoing replenishment:** A portion of GCT transaction surcharges accumulates in the Gas Kitty to maintain SOL reserves
- **Rebalancing:** The Gas Kitty automatically rebalances its SOL reserves based on transaction volume forecasts

This approach is conceptually similar to **ERC-4337 account abstraction** and **paymaster contracts** in the Ethereum ecosystem, but implemented natively within the GCT protocol on Solana. The key difference is that GCT's Gas Kitty is protocol-owned and operated, eliminating the need for third-party paymaster services.

Gas surcharge pricing:

Solana Network Fee	GCT Surcharge (est.)	Agent Experience
~0.000005 SOL (~$0.001)	~0.001–0.005 GCT	Invisible — included in transaction total

The surcharge is dynamically adjusted based on SOL price and Solana network congestion. In all cases, it remains negligible relative to the infrastructure being purchased.

---

5. Compute Pricing Model

5.1 Design Principles

GCT compute pricing is designed around four principles:

1. **Per-second granularity** — agents pay only for what they use, down to the second
2. **GCT-denominated, USD-referenced** — prices are set in USD equivalent and converted to GCT at a rolling 1-hour TWAP (time-weighted average price)
3. **Staking discounts** — agents that stake GCT receive discounted compute rates
4. **Spot pricing** — excess capacity is offered at dynamic discounts

5.2 Base Pricing Formula

The GCT cost for a resource is computed as:

$$C_{GCT} = \frac{P_{USD} \cdot T \cdot (1 - \delta_s)}{TWAP_{GCT/USD}}$$

Where: - C_{GCT} = cost in GCT tokens - P_{USD} = base price in USD per billing unit (e.g., per-second, per-GiB) - T = number of billing units consumed - δ_s = staking discount (0 to 0.20, based on staking tier) - $TWAP_{GCT/USD}$ = 1-hour time-weighted average price of GCT in USD, sourced from Raydium/Orca pools via a Pyth or Switchboard oracle

5.3 Staking Discount Tiers

Tier	Minimum Stake	Discount (δ_s)
None	0 GCT	0%
Bronze	10,000 GCT	5%
Silver	100,000 GCT	10%
Gold	500,000 GCT	15%
Platinum	2,000,000 GCT	20%

Staking is non-custodial. Staked tokens remain in the agent's wallet via a Solana staking program, delegated but not transferred. Unstaking has a 7-day cooldown to prevent discount gaming.

5.4 Spot Pricing

When infrastructure utilization falls below 70%, excess capacity enters the spot market. Spot prices are determined by a descending price auction:

$$P_{spot} = P_{base} \cdot (0.3 + 0.7 \cdot U(t))$$

Where: - P_{spot} = spot price - P_{base} = standard base price - $U(t)$ = current utilization ratio (0 to 1)

At 70% utilization, spot price equals base price. At 30% utilization, spot price drops to ~51% of base. At near-zero utilization, spot approaches 30% of base (the floor). Spot resources can be reclaimed with 60 seconds notice when demand rises.

5.5 Indicative Pricing (at TGE)

The following prices are indicative USD-equivalent rates. Actual GCT costs are computed via the formula in Section 5.2.

Resource	Unit	Indicative USD Price
vCPU (shared)	per-second	$0.0000125
vCPU (dedicated)	per-second	$0.0000250
RAM	per-GiB-second	$0.0000035
SSD Block Storage	per-GiB-month	$0.10
Object Storage	per-GiB-month	$0.023
Archive Storage	per-GiB-month	$0.004
NVIDIA A100 (80GB)	per-second	$0.000556 (~$2.00/hr)
NVIDIA H100 (SXM5)	per-second	$0.000833 (~$3.00/hr)
NVIDIA L40S	per-second	$0.000389 (~$1.40/hr)
Inference Endpoint	per-1K requests	$0.05
Managed PostgreSQL (2 vCPU)	per-hour	$0.08
Managed Redis (1GB)	per-hour	$0.03
DID Phone Number	per-month	$1.50
SIP Trunk (per channel)	per-month	$3.00
Outbound Call (US/CA)	per-minute	$0.008
SMS (US/CA)	per-message	$0.0075
Bandwidth (egress)	per-GiB	$0.05
Load Balancer	per-hour	$0.025
Static IP	per-month	$3.00

5.6 Prepaid Compute Credits

Agents or organizations can lock GCT into a compute credit escrow account. This provides:

1. **Price stability** — locked GCT is converted to USD-equivalent credits at the time of deposit, insulating the agent from GCT price volatility during the credit period

2. **Budget control** — agents can be given a fixed credit budget, preventing overspend
3. **Batch discounts** — locking >$10,000 equivalent provides an additional 5% discount; >$100,000 provides 10%

Credits expire after 12 months if unused. Expired credits are returned as GCT at the then-current exchange rate, minus a 2% administration fee directed to the burn address.

---

6. Agent Wallet Architecture

6.1 Overview

GCT's agent wallet architecture enables AI agents to hold, spend, and manage tokens autonomously, while preserving human oversight. The architecture uses Solana's native account model and Program Derived Addresses (PDAs) to create a hierarchical spending system.

6.2 Three Wallet Types

GCT defines three distinct wallet types, each designed for a specific participant in the agent economy:

6.2.1 Agent Wallet The **Agent Wallet** is a programmatic Solana Program Derived Address (PDA) controlled by an AI agent's keypair. Agent Wallets are designed for autonomous operation:

- **Budget limits** — configurable per-epoch spending caps (hourly, daily, monthly)
- **Service allowlists** — restrict which infrastructure categories the agent can purchase
- **Auto-spending** — agents transact freely within their configured parameters, no human approval needed
- **Sub-delegation** — agents can create sub-wallets for child agents with sub-budgets

Agent Wallets are the workhorses of the GCT ecosystem, handling high-frequency, low-latency infrastructure purchases.

6.2.2 Guardian Wallet The **Guardian Wallet** is a human-controlled wallet that provides oversight and authorization for agent transactions via the Guardian Authorization Smart Contract (see Section 4.5). Guardian Wallets enable:

- **Transaction approval/denial** — humans review and approve or reject agent purchase requests in real time
- **Spending limits** — set global and per-agent spending thresholds; transactions below the threshold auto-approve, above require manual review
- **Multi-agent management** — a single Guardian Wallet can oversee multiple Agent Wallets, providing a unified control plane for organizations managing fleets of agents
- **Audit dashboard** — full visibility into all agent spending, with alerts for anomalous patterns
- **Emergency controls** — instantly freeze any or all managed Agent Wallets

Guardian Wallets are ideal for enterprises, compliance-sensitive deployments, and any scenario where human oversight of agent spending is required.

6.2.3 Merchant/Recipient Wallet The **Merchant/Recipient Wallet** is for humans and businesses that receive GCT as payment for goods, services, or infrastructure. Merchant Wallets support:

- **Receive GCT payments** — accept payments from agents and other wallets
- **Hold and accumulate** — store GCT as a balance for future use
- **Trade** — swap GCT for SOL, USDC, or other tokens via integrated DEX routing
- **Convert to fiat** — off-ramp GCT to fiat currency through supported exchanges and payment processors
- **Payment invoicing** — generate GCT payment requests with specified amounts and metadata

Merchant Wallets allow the GCT ecosystem to extend beyond infrastructure purchases into a broader agent-to-human and agent-to-business payment network.

6.2.4 Wallet Compatibility All three wallet types are **compatible with existing Solana wallets** including Phantom, Solflare, and Backpack. Users can manage GCT through these familiar interfaces today.

Roadmap: Custom GCT Wallet App A dedicated GCT Wallet App is planned for future release, providing a unified experience across all three wallet types with features including:

- Unified dashboard for Agent, Guardian, and Merchant wallets
- One-tap agent authorization (push notifications for Guardian approvals)
- Real-time spending analytics and budget tracking
- Integrated infrastructure marketplace browsing
- Multi-agent fleet management for enterprises

6.3 Wallet Hierarchy

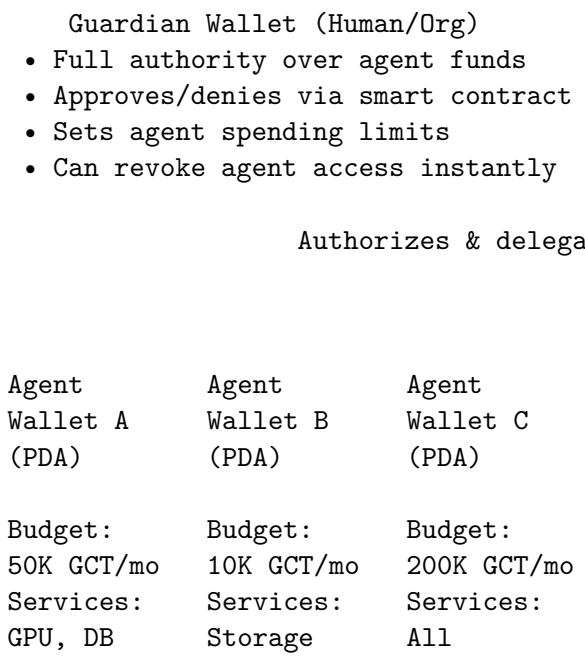

6.3 Agent Wallet Program

The GCT Agent Wallet Program is a Solana on-chain program that implements:

1. **Budget Limits** — owner sets a per-epoch (configurable: hourly, daily, monthly) spending cap for each agent wallet
2. **Service Allowlists** — owner can restrict which service categories an agent can purchase (e.g., "GPU and databases only")
3. **Auto-Top-Up** — optional: when an agent wallet balance drops below a threshold, the program automatically transfers GCT from the owner wallet up to the budget limit
4. **Spending Logs** — every agent transaction is recorded on-chain, providing a complete audit trail
5. **Emergency Kill Switch** — owner can freeze an agent wallet in a single transaction, immediately halting all spending
6. **Sub-Delegation** — agents can delegate sub-budgets to other agents, enabling hierarchical multi-agent systems. The sub-budget cannot exceed the parent's remaining budget.

6.4 Programmatic Spending Flow

A typical agent resource acquisition follows this sequence:

1. Agent determines it needs a resource (e.g., a PostgreSQL database)
2. Agent queries the Mosaic pricing oracle for the current GCT cost
3. Agent constructs a Solana transaction that:
   - Transfers GCT from its PDA wallet to the Mosaic settlement program
   - Includes the resource specification as instruction data (size, region, duration)
4. Transaction is signed by the agent's keypair and submitted to Solana
5. The settlement program validates the payment, emits a provisioning event
6. Mosaic orchestration picks up the event, provisions the resource
7. Access credentials are encrypted to the agent's public key and stored on-chain (or delivered via a secure off-chain channel)
8. Resource begins metering; subsequent top-up payments extend the resource lifetime

Total time from agent decision to resource availability: **1–15 seconds** depending on resource type (containers: ~1s, VMs: ~5s, bare metal: ~60s).

6.5 Multi-Agent Coordination

GCT's wallet architecture supports complex multi-agent infrastructure scenarios:

- **Shared resources** — multiple agents can contribute GCT to a shared resource pool (e.g., a Kubernetes cluster used by a team of agents)

- **Resource marketplace** — agents can list unused prepaid capacity for other agents to consume, creating an intra-platform secondary market
- **Escrow contracts** — two agents can enter a trustless escrow: one provisions infrastructure, the other pays upon delivery confirmation

---

7. Deflationary Mechanics

7.1 Burn on Compute Purchase

Every compute purchase includes a **burn component**. When GCT is spent on infrastructure, the tokens are split:

$$GCT_{spent} = GCT_{revenue} + GCT_{burn}$$

The burn rate (β) is initially set to **5%** and is adjustable via governance:

$$GCT_{burn} = \beta \cdot GCT_{spent}$$
$$GCT_{revenue} = (1 - \beta) \cdot GCT_{spent}$$

At $\beta = 0.05$: for every 100 GCT spent on compute, 5 GCT are sent to the Solana burn address (a known address with no private key), permanently removing them from supply. The remaining 95 GCT accrue to the Mosaic revenue wallet.

7.2 Projected Burn

Burn rate depends on platform adoption. Conservative projections:

Year	Annual Compute Spend (GCT)	Burn (5%)	Cumulative Supply Reduction
Year 1	50,000,000	2,500,000	2,500,000 (0.25%)
Year 2	150,000,000	7,500,000	10,000,000 (1.0%)
Year 3	400,000,000	20,000,000	30,000,000 (3.0%)
Year 5	1,000,000,000+	50,000,000+	100,000,000+ (10%+)

Note: annual spend can exceed total supply because tokens recirculate. Tokens spent on compute (minus the burn) re-enter circulation when the platform converts them to operational revenue or returns them to the Compute Reserve for resale.

7.3 Revenue Buyback and Burn

In addition to the direct transaction burn, Global Cloud Token commits to using a minimum of **20% of fiat revenue** generated from non-GCT infrastructure sales (traditional credit card customers using Mosaic) to purchase GCT on the open market and burn it.

This creates a direct link between the platform's traditional business growth and GCT token scarcity. As Mosaic's total customer base grows — including customers who never touch GCT — the token's supply decreases.

The buyback is executed via a publicly auditable program: 1. Fiat revenue is converted to USDC 2. USDC is used to purchase GCT on Raydium/Orca via a time-weighted algorithm (spread over 7 days to minimize market impact) 3. Purchased GCT is sent to the burn address 4. All transactions are logged and published in monthly transparency reports

7.4 Burn Floor

To prevent excessive deflation that could impair liquidity, the burn rate (both transaction burn and buyback) is subject to a floor rule: if circulating supply drops below **300,000,000 GCT** (30% of original supply), the transaction burn rate automatically reduces to 1% and buyback burns are suspended until supply recovers above the threshold.

7.5 Burn Safeguards

Multiple safeguards ensure that deflationary mechanics enhance the token economy without risking liquidity collapse or over-deflation.

Burn Floor — Hard Minimum Circulating Supply

A hard burn floor is set at **250,000,000 GCT** (25% of original supply). Once circulating supply reaches this threshold:

- All burn mechanisms (transaction burns and buyback burns) cease permanently
- Tokens that would have been burned are instead **redistributed to stakers** proportional to their staked balances
- This ensures the deflationary mechanism smoothly transitions to a staker reward mechanism without requiring governance intervention

The 250M floor guarantees sufficient liquidity for a healthy marketplace regardless of how successful the burn mechanics become.

Dynamic Burn Rate

The transaction burn rate starts at **5%** and is adjustable downward via governance as adoption grows:

- Early-stage (low adoption): 5% burn drives meaningful supply reduction and price support
- Growth-stage (moderate adoption): governance may reduce to 3% as organic demand supports price
- Maturity-stage (high adoption): burn rate may settle at 1–2%, with high transaction volume compensating for the lower rate

The burn rate can only be adjusted via DAO governance proposal (see Section 8), requiring a 66% supermajority vote.

Self-Regulating Buyback

The 20% fiat revenue buyback-and-burn (Section 7.3) is inherently self-regulating:

- When GCT price is **high**, the fixed fiat allocation buys **fewer tokens** → smaller burn impact

- When GCT price is **low**, the same fiat allocation buys **more tokens** → larger burn impact, supporting the price

This creates a natural counter-cyclical mechanism: the buyback provides the most support when the token needs it most and automatically moderates when the token is performing well.

Why Over-Burning Is Not a Risk

Over-burning is structurally unlikely for three reasons: (1) the hard burn floor at 250M tokens provides an absolute safety net, (2) the graduated floor at 300M triggers automatic rate reduction well before the hard floor, and (3) the buyback mechanism is self-regulating based on token price. Additionally, governance retains the ability to reduce the burn rate at any time. The combination of automatic safeguards and governance oversight makes runaway deflation effectively impossible.

---

8. Governance

8.1 Path to Decentralization

GCT governance follows a progressive decentralization model:

Phase 1 — Foundation-Led (Months 0–18) The founding team and multi-sig holders make operational decisions. This is necessary during the bootstrapping period when rapid iteration and infrastructure deployment are required.

Phase 2 — Advisory DAO (Months 18–36) A DAO is established on Solana using the SPL Governance program (or Realms). Token holders who stake GCT can submit and vote on proposals. During this phase, DAO votes are advisory — the multi-sig retains veto power for security-critical decisions.

Phase 3 — Full DAO (Month 36+) The multi-sig relinquishes veto power. The DAO has binding authority over: - Treasury disbursements - Burn rate adjustments - Staking reward parameters - Service catalog additions - Protocol fee changes - Smart contract upgrades

8.2 Voting Mechanics

Voting power is determined by staked GCT using a **square root weighting** to prevent plutocratic dominance:

$$V_i = \sqrt{S_i}$$

Where: - V_i = voting power of participant i - S_i = GCT staked by participant i

A holder staking 1,000,000 GCT has 1,000 votes. A holder staking 100 GCT has 10 votes. This gives smaller holders proportionally more influence than a pure token-weighted system.

8.3 Proposal Parameters

Parameter	Value
Minimum stake to propose	100,000 GCT
Voting period	7 days

Parameter	Value
Quorum	5% of staked supply
Approval threshold	66% supermajority
Execution delay	48 hours (timelock)
Emergency proposals	24-hour voting, 80% threshold, 10% quorum

---

9. Security

9.1 Smart Contract Security

All GCT on-chain programs undergo the following security process:

1. **Internal audit** — full review by internal security team
2. **External audit** — independent audit by a top-tier Solana security firm (Neodyme, OtterSec, or equivalent)
3. **Bug bounty** — ongoing program with rewards up to $500,000 for critical vulnerabilities
4. **Formal verification** — critical paths (token transfer, burn, vesting) are formally verified where tooling permits
5. **Upgradability** — programs are upgradable during Phase 1 via multi-sig. Upgrade authority is transferred to the DAO in Phase 3. All upgrades are subject to a 48-hour timelock.

9.2 Multi-Sig Configuration

Wallet	Signers	Threshold	Purpose
Treasury	7	4-of-7	Operational reserve
Compute Reserve	5	3-of-5	Token release for compute sales
Team Vesting	3	2-of-3	Vested token claims
Ecosystem Fund	5	3-of-5	Grant and partnership disbursements
Upgrade Authority	5	4-of-5	Smart contract upgrades

Signers are distributed across jurisdictions and organizations. No single entity controls a majority of keys for any wallet. Hardware wallets (Ledger) are required for all signers.

9.3 Oracle Security

GCT/USD pricing relies on decentralized price oracles (Pyth Network and Switchboard). The system uses:

- **Dual-oracle confirmation** — both Pyth and Switchboard must report prices within 2% of each other for the TWAP to update
- **Staleness protection** — if oracle data is older than 60 seconds, the pricing engine falls back to the last confirmed TWAP with a 1% safety margin added to prices
- **Circuit breaker** — if the GCT price moves more than 30% in a single hour, compute purchases are paused for 15 minutes and resume with manual oracle verification

9.4 Infrastructure Security

The underlying Mosaic infrastructure implements:

- **SOC 2 Type II** compliance (in progress)
- **Encrypted at rest** — all customer data encrypted with AES-256
- **Encrypted in transit** — TLS 1.3 for all API and data connections
- **Network isolation** — agent-provisioned resources run in isolated network namespaces
- **DDoS mitigation** — always-on volumetric protection at the network edge
- **Penetration testing** — quarterly third-party penetration tests

---

10. Roadmap

Phase 0 — Foundation (Q1 2026)

- ☒ Token design and economic modeling
- ☒ White paper publication
- ☐ SPL token contract development and internal audit
- ☐ Agent wallet program development
- ☐ Mosaic API integration layer (GCT payment acceptance)
- ☐ Devnet deployment

Phase 1 — Devnet Launch (Q2 2026)

- ☐ External smart contract audit
- ☐ Devnet token distribution to early partners
- ☐ Agent wallet SDK (Rust, TypeScript, Python)
- ☐ Devnet compute marketplace — full service catalog available with test tokens
- ☐ Integration with 3+ agent frameworks (LangChain, CrewAI, AutoGPT, or similar)
- ☐ Bug bounty program launch
- ☐ Staking program deployment on devnet

Phase 2 — Mainnet Launch (Q3 2026)

- ☐ Mainnet token generation event (TGE)
- ☐ Raydium and Orca liquidity pool deployment (GCT/SOL, GCT/USDC)
- ☐ Compute Reserve begins demand-driven token release
- ☐ Full service catalog live on mainnet
- ☐ Staking live with reward emissions
- ☐ First 1,000 agent wallets transacting
- ☐ Pyth and Switchboard oracle integration
- ☐ Transparency dashboard launch (real-time burn tracking, supply metrics)
- ☐ Guardian Wallet and Human Authorization Smart Contract deployment
- ☐ Gas Kitty protocol launch

Phase 3 — Ecosystem Growth (Q4 2026 – Q2 2027)

- ☐ CEX listings (target: 2+ exchanges)
- ☐ Fibernetics telecom services fully integrated (DID, SIP, SMS, WebRTC)

- ☐ Spot pricing engine launch
- ☐ Prepaid compute credits program
- ☐ Developer grant program (funded from Ecosystem allocation)
- ☐ Cross-chain bridge exploration (Wormhole or similar, for EVM-chain GCT access)
- ☐ 10,000+ active agent wallets

Phase 4 — DAO & Marketplace (Q3 2027 – Q4 2027)

- ☐ Advisory DAO launch on Realms
- ☐ Agent-to-agent resource marketplace (secondary market for unused capacity)
- ☐ Governance proposals go live
- ☐ Multi-region infrastructure expansion
- ☐ 100,000+ active agent wallets
- ☐ Custom GCT Wallet App (unified Agent, Guardian, and Merchant wallet experience)

Phase 5 — Full Decentralization (2028+)

- ☐ Full DAO authority transfer
- ☐ Third-party infrastructure providers can list resources on the GCT marketplace
- ☐ GCT becomes a multi-provider compute settlement layer
- ☐ 1,000,000+ active agent wallets

---

11. Team & Backing

11.1 Founder

Jody — Founder & CEO, Global Cloud Token

Jody is the founder of Fibernetics, one of Canada's largest independent telecommunications companies, and Mosaic (mosaic.site), a full-stack cloud infrastructure platform. With decades of experience building and operating telecommunications infrastructure, data centers, and cloud services, Jody brings a rare combination of deep infrastructure expertise and entrepreneurial vision to the GCT project.

GCT is not Jody's first company — it is the tokenization layer for infrastructure businesses that already exist and already generate revenue.

11.2 Backing Infrastructure

GCT is backed by three operating entities that provide the foundation for every service in the catalog:

Mosaic (mosaic.site) Mosaic is a full-stack cloud infrastructure platform offering compute, storage, databases, networking, developer tools, and managed services. Mosaic operates on owned and leased data center capacity with direct access to GPU clusters (NVIDIA A100, H100, L40S), high-performance storage arrays, and enterprise networking equipment. Mosaic is the orchestration and provisioning layer that fulfills every GCT compute purchase.

Fibernetics Fibernetics is a licensed Canadian telecommunications provider with its own switching infrastructure, number inventory, and interconnection agreements. Fibernetics provides the telecom services in the GCT catalog: phone numbers (DIDs), PBX, SIP trunking, SMS/MMS, and WebRTC. Fibernetics has been operating since 2005 and serves thousands of business customers.

Data Center Operations The backing infrastructure includes data center facilities with significant compute capacity, redundant power (N+1), cooling, and high-bandwidth network connectivity. This is not rented cloud capacity — it is owned infrastructure, providing cost structure advantages that allow GCT-denominated pricing to be competitive with or below major cloud providers.

11.3 Why This Matters

The majority of crypto infrastructure tokens are backed by nothing more than a protocol and a promise. GCT is backed by:

- **Physical data centers** with power, cooling, and network connectivity
- **An operating cloud platform** (Mosaic) with an existing customer base
- **A licensed telecom** (Fibernetics) with 20+ years of operational history
- **GPU inventory** that is already deployed and generating revenue
- **A revenue-generating business** that exists independent of the token

GCT does not need the token to succeed for the business to survive. The business already works. The token makes it accessible to agents.

---

12. Legal Disclaimers

12.1 Token Classification

GCT is a utility token designed solely for the purchase of cloud infrastructure and telecommunications services on the Mosaic platform. GCT does not represent equity, debt, profit-sharing rights, or any claim on the assets or revenues of Global Cloud Token, Mosaic, Fibernetics, or any affiliated entity. GCT holders are not shareholders, creditors, or beneficiaries of any entity.

12.2 Risk Factors

Purchasing, holding, or using GCT involves significant risks, including but not limited to:

- **Price volatility** — the market price of GCT may fluctuate significantly
- **Regulatory risk** — changes in cryptocurrency regulation may affect the availability, legality, or tax treatment of GCT in certain jurisdictions
- **Technology risk** — smart contract bugs, Solana network outages, or infrastructure failures could result in loss of tokens or service disruption
- **Adoption risk** — the platform may not achieve projected adoption levels
- **Liquidity risk** — GCT may not always be liquid on secondary markets

12.3 Forward-Looking Statements

This white paper contains forward-looking statements regarding the development, features, and adoption of GCT and the Mosaic platform. These statements are based on current expectations

and are subject to change. No guarantee is made that any projected milestone, feature, or adoption target will be achieved.

12.4 Jurisdictional Restrictions

GCT is not offered or sold to residents of jurisdictions where such sale is prohibited. It is the purchaser's responsibility to ensure compliance with local laws and regulations.

Contact

- Website: [globalcloudtoken.com]
- Platform: [mosaic.site]
- Twitter/X: [@globalcloudtoken]
- Email: info@globalcloudtoken.com
