

UDOC

UNIFIED DIGITAL OVERSIGHT & COORDINATION

Technical Whitepaper

Architecture, Security Model, Data Governance, and Implementation Specification

v1.0 Document Version	PQC Quantum-Ready Encryption Standard	99.9% Uptime SLA Target	Zero Foreign Cloud Dependency
---------------------------------	--	-----------------------------------	---

G.O.D.S Holdings (Pty) Ltd · Johannesburg, South Africa · 2026
Prepared by: Sashin J. Singh, Founder & Architect

EXECUTIVE SUMMARY

What UDOC Is and Why It Matters

UDOC is not a software product. It is sovereign digital governance infrastructure — the technical backbone that allows governments and regulated enterprises to govern AI systems, protect national data sovereignty, and generate independently auditable compliance records, without dependency on foreign-owned cloud platforms.

The problem UDOC addresses is structural. As of 2026, the majority of governments globally rely on infrastructure owned and operated by three American cloud providers — Amazon AWS, Microsoft Azure, and Google Cloud — for functions ranging from tax administration to national identity systems. This arrangement transfers operational sovereignty from the state to commercial entities operating under foreign legal frameworks.

Simultaneously, the deployment of AI systems in consequential decision-making — across welfare administration, border management, criminal justice, and financial regulation — has outpaced the governance infrastructure designed to oversee it. Audit trails do not exist. Explainability is inconsistent. Bias detection is ad hoc. Compliance frameworks are aspirational rather than technical.

UDOC addresses both problems with a single, integrated architecture: a sovereign-first, quantum-ready, independently auditable governance platform that can be deployed as enterprise SaaS for regulated industries and as national digital backbone for governments. It is built on four technical pillars: data sovereignty, AI transparency, quantum security, and operational auditability.

Version	Technical Whitepaper v1.0 — Architecture Specification
Classification	Confidential — Government & Institutional Partners Only
Stack	Python / FastAPI backend · React/TypeScript frontend · PostgreSQL + Apache Kafka · Kubernetes on sovereign cloud · NIST PQC encryption
Deployment	Sovereign cloud (government-owned infrastructure) · Private enterprise cloud · Hybrid on-premise/cloud
Encryption	CRYSTALS-Kyber (KEM) · CRYSTALS-Dilithium (signature) · AES-256-GCM (symmetric) · SHA-3 (hashing) — full NIST PQC compliance
Data Jurisdiction	All client data remains under client jurisdiction at all times. G.O.D.S retains zero data after contract termination.
Audit Architecture	Cryptographically verified, tamper-evident, chronologically immutable audit trail on every AI decision event
Uptime SLA	99.9% availability guaranteed · 4-hour RTO / 1-hour RPO for government backbone deployments
Interoperability	EU AI Act · OECD AI Principles · ISO/IEC 42001 · NIST AI RMF · South African POPIA · GDPR-compatible

CONTENTS

Technical Whitepaper Structure

1.	Executive Summary	2
2.	The Digital Sovereignty Problem	4
3.	UDOC System Architecture	5
4.	Core Platform Components	7
5.	Data Architecture & Sovereignty Model	10
6.	AI Governance Engine	12
7.	Quantum-Ready Security Architecture	15
8.	Audit Trail & Cryptographic Verification	17
9.	Deployment Architecture	19
10.	API Design & Integration Specifications	21
11.	Compliance Framework Alignment	23
12.	Performance Specifications & SLAs	25
13.	Implementation Roadmap	27
14.	Commercial Architecture	29
15.	Appendices	30

SECTION 2

The Digital Sovereignty Problem

Why existing infrastructure is structurally inadequate

2.1 The Foreign Infrastructure Dependency

As of 2026, the majority of national governments in the African Union, Middle East, and parts of the Asia-Pacific region operate critical public services on infrastructure owned by Amazon Web Services (AWS), Microsoft Azure, or Google Cloud Platform (GCP). These three providers collectively hold an estimated 65–70% of global public cloud infrastructure.

The implications of this concentration are structural rather than theoretical. When a government stores citizen data on AWS infrastructure, the legal framework governing that data — including law enforcement access provisions under the US CLOUD Act — is American law, regardless of where the servers are physically located. The government's data sovereignty guarantee to its citizens is, in practice, limited by its dependency on a foreign-jurisdiction cloud provider.

Infrastructure Layer	Current Provider	Sovereignty Risk	UDOC Response
Identity Systems	AWS GovCloud / Azure	Foreign jurisdiction access	Sovereign key management
Tax Administration	Microsoft Azure	Data residency uncertainty	In-country compute guarantee
Social Grant Systems	AWS	CLOUD Act exposure	Zero-transfer architecture
Border Management	Mixed foreign cloud	Intelligence access risk	Air-gapped option available
National Health Records	Google Health Cloud	Commercial data mining risk	Sovereign encryption at rest

2.2 The AI Governance Gap

The deployment of AI systems in consequential government decision-making has accelerated dramatically since 2022. Predictive policing systems, welfare eligibility algorithms, credit scoring models, and tax audit selection tools now operate at scale across most developed and many developing economies. The governance infrastructure designed to oversee these systems has not kept pace.

Current state: most AI systems deployed in government contexts operate with minimal audit infrastructure. When a welfare application is denied by an algorithmic system, the applicant has no access to an explanation. When a predictive policing model produces a false positive, the data trail required to identify the error may not exist. When a regulator wishes to audit an AI-assisted financial decision, the technical infrastructure to support that audit has often not been designed into the system.

UDOC treats AI governance not as a compliance checkbox but as a technical requirement embedded in the architecture from inception. Every AI decision event produces a cryptographically

signed, tamper-evident record. Every model produces an explainability output. Every system is registered in the AI Registry before deployment.

SECTION 3

UDOC System Architecture

Technical overview of the platform layers

3.1 Architectural Philosophy

UDOC's architecture is governed by five design principles that are non-negotiable across all deployments and cannot be overridden by commercial or operational considerations:

- Sovereignty First: No client data traverses infrastructure not under client legal control
- Audit by Default: Every system event produces a cryptographically verifiable record
- Zero Trust: No component trusts any other without continuous verification
- Quantum-Ready: All cryptographic primitives are NIST PQC-compliant from inception
- Interoperable: All interfaces publish to open standards; no proprietary lock-in

3.2 Platform Layer Overview

Layer 1 — Data Ingestion	Secure API gateway · Event streaming (Apache Kafka) · Cryptographic signing at ingestion · Schema validation and data lineage tagging
Layer 2 — Storage	PostgreSQL primary (ACID compliant) · Apache Cassandra for audit logs (immutable write) · Redis caching layer · All data encrypted at rest (AES-256-GCM) with sovereign key management
Layer 3 — Processing	Python/FastAPI microservices · AI model execution sandbox · Explainability engine · Risk classification pipeline · Async event processing
Layer 4 — Governance Engine	AI Registry · Bias detection · Compliance rule engine · Workflow management · Alert and escalation framework
Layer 5 — Audit & Crypto	Merkle tree audit log · CRYSTALS-Dilithium signing · Tamper-evident chain · Export to regulatory authority format · Key rotation management
Layer 6 — Presentation	React/TypeScript SPA · REST + GraphQL APIs · Role-based access control (RBAC) · Mobile-responsive · Screen-reader accessible
Layer 7 — Infrastructure	Kubernetes orchestration · Sovereign cloud deployment · Horizontal autoscaling · Multi-zone redundancy · 99.9% uptime architecture

3.3 Technology Stack Specification

Component	Technology	Version	Justification
API Framework	Python / FastAPI	0.111+	Async performance · Type safety · OpenAPI auto-docs
Frontend	React / TypeScript	18.x	Component reuse · Type safety · Government audit trail UI

Primary Database	PostgreSQL	16.x	ACID compliance · Row-level security · Audit extension
Audit Log Store	Apache Cassandra	5.x	Immutable write · Horizontal scale · 10-year retention
Event Streaming	Apache Kafka	3.x	Tamper-evident log · Event replay · Compliance trail
Caching	Redis	7.x	Session management · Rate limiting · Temp data isolation
Orchestration	Kubernetes	1.30+	Sovereign cloud portability · Autoscaling · Zero-downtime deploy
Search & Analytics	OpenSearch	2.x	Full-text audit search · Compliance dashboards · Alerting
Monitoring	Prometheus + Grafana	Latest	Real-time SLA monitoring · Anomaly detection · Alerting
CI/CD	GitLab CI / ArgoCD	Latest	Air-gapped deployment · GitOps · Compliance pipeline

SECTION 4

Core Platform Components

The five functional modules

4.1 AI System Registry

The AI Registry is the authoritative catalogue of all AI systems operating within the client's jurisdiction. Every AI system — regardless of vendor, deployment model, or operational context — must be registered in UDOC before it may be used in any consequential decision.

Registration Fields	System name and version · Owner organisation · Training data description and provenance · Model type and architecture · Intended use cases · Prohibited use cases · Risk classification (Low / Medium / High / Critical) · Human oversight requirements · Performance metrics and validation results · Vendor contact and SLA
Automated Scanning	API-based registration trigger detects new model deployments. Unregistered AI system detection alerts compliance team within 4 hours. Automated de-registration on model retirement.
Audit Integration	Every registration event is cryptographically signed and written to immutable audit log. Registry state at any historical point can be reconstructed for regulatory inquiry.

4.2 Risk Classification Engine

All AI systems are automatically classified by risk level on registration and continuously re-evaluated as operational context changes. The classification framework is aligned with EU AI Act taxonomy and extended for African regulatory contexts.

Risk Levels	CRITICAL (prohibited uses: social scoring, real-time biometric mass surveillance) · HIGH (criminal justice, welfare eligibility, credit scoring, border management) · MEDIUM (recruitment screening, insurance pricing, diagnostic assistance) · LOW (chatbots, spam filtering, personalisation)
Classification Logic	Rule-based classification from registration data + ML-assisted risk scoring from operational data. Human review required for all HIGH and CRITICAL classifications. Quarterly re-assessment for all registered systems.
Override Protocol	Any classification may be appealed through a defined process. Appeals are reviewed by UDOC compliance committee. Override decisions are permanently logged with justification.

4.3 Compliance Reporting Engine

UDOC generates regulatory audit reports automatically from operational data. Reports are cryptographically signed, formatted to regulatory standards, and available on demand.

Supported Frameworks	EU AI Act (Articles 9, 13, 14, 17, 18, 29) · OECD AI Principles · ISO/IEC 42001:2023 · NIST AI RMF · South African POPIA · GDPR Article 22 (automated decision-making)
Report Types	Conformity Assessment (annual) · Incident Report (on-trigger) · Audit Trail Extract (on-demand) · Risk Register (quarterly) · Human Oversight Verification (continuous)

Generation Time	Standard compliance report: <60 seconds. Full audit trail extract (12-month): <10 minutes. Emergency incident report: <5 minutes.
-----------------	--

4.4 Workforce Impact Monitoring (SETHS Integration)

For UDOC deployments integrated with S.E.T.H.S, the platform includes an AI displacement assessment module that tracks AI-driven workforce changes and models reskilling pathways.

Displacement Tracking	Role-level analysis of AI automation exposure · Department-level impact reporting · Organisation-wide transition planning dashboard
Pathway Modelling	Integration with S.E.T.H.S certified training pathways · Personalised reskilling recommendations · Progress tracking against transition plan
Government Reporting	National-level AI displacement data aggregated across all enterprise clients (anonymised) · Policy planning dashboards for Department of Employment and Labour

4.5 Audit Trail Service

The audit trail is UDOC's most technically distinctive component. Every decision event produces a cryptographically signed record stored in an append-only, tamper-evident log.

Event Capture	Model ID + version · Input data hash (not raw data) · Decision output · Confidence score · Timestamp (nanosecond precision) · Human reviewer ID (if applicable) · Digital signature (CRYSTALS-Dilithium)
Storage Architecture	Apache Cassandra append-only cluster · Merkle tree integrity verification · 10-year minimum retention · Cross-region replication for government deployments · WORM (Write Once Read Many) compliance
Verification API	Any historical decision can be verified in <500ms · Merkle proof generated on demand · Batch verification for regulatory audits · Export formats: JSON-LD, XML, PDF with digital signature

SECTION 5

Data Architecture & Sovereignty Model

Technical data governance specification

5.1 Data Sovereignty Architecture

Data sovereignty is the foundational constraint of every UDOC architectural decision. The sovereignty model is not a policy document — it is a technical guarantee embedded in the architecture. It is impossible by design, not merely by policy, for client data to leave the client's jurisdictional control.

Key Management	Client holds master encryption keys. G.O.D.S engineers cannot access client data at rest. Key rotation is client-controlled. Emergency access requires signed client authorisation.
Data Residency	All client data processed and stored within client-defined geographic boundary. No cross-border replication without explicit client instruction. Country-specific data residency certificates available.
Compute Sovereignty	Processing occurs on client-owned or client-contracted sovereign cloud infrastructure. G.O.D.S never operates on AWS/Azure/GCP for government backbone deployments.
Network Isolation	Government deployments support full network air-gap with USB-based secure update mechanism. All external communications require client firewall authorisation.
Termination Protocol	On contract termination: all client data deleted within 30 days · Deletion certificate issued with cryptographic proof · G.O.D.S retains no copies · Client receives full data export in open format before deletion
Third-Party Access	No third-party access to client data without documented client consent. Law enforcement requests handled per client jurisdiction law. G.O.D.S will not comply with foreign law enforcement requests for client data.

5.2 Data Classification Framework

Classification	Description	Handling	Encryption Level
SOVEREIGN-CRITICAL	National security, identity, defence	Air-gapped only; no cloud	AES-256 + CRYSTALS-Kyber
SOVEREIGN-HIGH	Citizen PII, health, financial	Sovereign cloud; no foreign provider	AES-256 + forward secrecy
GOVERNANCE	AI decision records, audit logs	Sovereign or private cloud	AES-256, immutable store
OPERATIONAL	Anonymised system performance	Any approved cloud	AES-128 minimum
PUBLIC	Open governance reports	Any infrastructure	TLS in transit only

5.3 Data Lineage Tracking

Every data item that enters UDOC is assigned a cryptographic lineage tag at ingestion. This tag tracks the data item's entire journey through the system: source, transformation operations, processing events, storage locations, and access events. The lineage record is stored separately from the data and cannot be modified.

```
# UDOC Data Lineage Tag — Example Schema JSON
{
  "lineage_id": "udoc-ln-20260310-f4a8e291",
  "source": { "system_id": "SETHS-placement-v2.1", "jurisdiction": "ZA" },
  "ingestion_timestamp": "2026-03-10T08:42:16.000003Z",
  "data_hash": "sha3-256:a7f4c829...",
  "classification": "GOVERNANCE",
  "jurisdiction_lock": "ZA",
  "transformations": [
    { "op": "anonymise-pii", "ts": "2026-03-10T08:42:16.000521Z", "actor": "pipeline-v3" }
  ],
  "signature": "dilithium3:3a8f...",
  "expiry": "2036-03-10T00:00:00Z"
}
```

Figure 5.1 — UDOC data lineage tag structure. Every data item receives this record at ingestion.

SECTION 6

AI Governance Engine

The technical core of UDOC

6.1 Governance Pipeline Architecture

The AI Governance Engine is the technical core of UDOC. It intercepts every AI decision event, validates it against registered governance rules, produces an explainability output, checks for bias signals, and writes a cryptographically signed record to the immutable audit log — all within the latency budget allocated to the decision event.

Design constraint: The governance pipeline adds a maximum of 50ms latency to any decision event. For most use cases this is imperceptible. For latency-critical applications (real-time credit scoring, fraud detection), a configurable "async governance mode" writes the record asynchronously while the decision proceeds, with a short hold on execution for CRITICAL-classified systems.

6.2 Explainability Engine

UDOC's explainability engine produces a human-readable explanation for every AI decision. The explanation format is adapted to the audience: technical detail for internal audit, plain language for affected citizens, regulatory format for compliance authorities.

```
# UDOC Explainability Output — Welfare Eligibility Example
{
  "decision_id": "udoc-dec-20260310-welfare-8842",
  "model_id": "welfare-eligibility-v3.2",
  "decision": "APPROVED",
  "confidence": 0.94,
  "explanation": {
    "primary_factors": [
      { "feature": "household_income", "direction": "approve", "weight": 0.42 },
      { "feature": "employment_status", "direction": "approve", "weight": 0.31 },
      { "feature": "prior_grant_compliance", "direction": "approve", "weight": 0.21 }
    ],
    "plain_language": "Application approved. Income below threshold. Compliant history.",
    "regulatory_ref": "Social Assistance Act s.5(1)(a)"
  },
  "bias_check": { "status": "PASS", "protected_attrs_checked": ["race", "gender", "age"] },
  "human_review_required": false,
  "signature": "dilithium3:9c2a...",
  "timestamp": "2026-03-10T09:15:42.000018Z"
}
```

Figure 6.1 — UDOC explainability output structure. Every decision event generates this record.

6.3 Bias Detection Framework

UDOC implements three complementary bias detection approaches that operate continuously on live decision data:

Statistical Parity	Continuous monitoring of decision rates across protected attribute groups. Alert threshold: >5% differential in approval/denial rates for any protected group with >1000 decision sample.
Counterfactual Testing	Automated counterfactual pairs generated for sample of decisions. Tests whether changing protected attribute (race, gender, age) while holding other factors constant changes the decision.
Individual Fairness	Embedding distance analysis ensures similar individuals receive similar decisions. Outlier detection flags cases where dissimilar outcomes occur for similar individuals.
Alerting Protocol	ALERT: Statistical parity breach · ESCALATION: Persistent breach >72 hours → compliance committee → mandatory model review · SUSPENSION: Persistent bias in HIGH/CRITICAL system → human-only decisions until model is retrained and validated

6.4 Human Oversight Enforcement

For AI systems classified as HIGH or CRITICAL risk, UDOC enforces mandatory human oversight — not as a policy aspiration, but as a technical constraint. The system cannot route decisions below the required human review threshold without triggering a compliance alert and suspending the AI system.

Minimum Review Rate	CRITICAL: 100% human review. HIGH: 20% minimum random sampling + all edge cases. MEDIUM: 5% minimum. LOW: Optional.
Review Interface	Dedicated human reviewer dashboard. Decision presented with full explainability output and model confidence score. Reviewer cannot see their own prior decisions to prevent anchoring bias.
Override Logging	Every human override of AI decision is logged with reviewer ID, timestamp, reason code, and cryptographic signature. Override rate is tracked and reviewed quarterly.
Timeout Enforcement	If required human review is not completed within configured window (default: 48 hours for welfare decisions), decision is escalated and system performance alert is triggered.

SECTION 7

Quantum-Ready Security Architecture

NIST PQC-compliant cryptographic framework

The transition to quantum computing poses a definitive threat to all current RSA, ECC, and Diffie-Hellman based cryptographic systems. Systems deployed today that use classical cryptography will be vulnerable to "harvest now, decrypt later" attacks — where adversaries collect encrypted data today intending to decrypt it when quantum computers become sufficiently powerful. UDOC is built on NIST PQC standards from inception, eliminating this risk for all clients.

7.1 NIST Post-Quantum Cryptography Standards

Algorithm	Type	Use Case in UDOC	Security Level	Status
CRYSTALS-Kyber (ML-KEM)	Key Encapsulation	TLS session keys · Data encryption keys	NIST Level 3 (128-bit PQ)	FIPS 203 Finalized
CRYSTALS-Dilithium (ML-DSA)	Digital Signature	Audit log signing · API auth · Document signing	NIST Level 3 (128-bit PQ)	FIPS 204 Finalized
FALCON	Digital Signature	Compact signatures for IoT/sensor endpoints	NIST Level 1 (64-bit PQ)	FIPS 206 Finalized
SPHINCS+ (SLH-DSA)	Stateless Sig	Long-term archive signing (10-year records)	NIST Level 3	FIPS 205 Finalized
AES-256-GCM	Symmetric Encryption	Data at rest · Transport encryption	Classical 128-bit equiv.	Quantum-secure
SHA-3 (SHAKE256)	Hash Function	Data integrity · Merkle tree nodes	Classical 128-bit equiv.	Quantum-secure

7.2 Hybrid Transition Architecture

During the transition period (2026–2035), UDOC operates a hybrid cryptographic architecture that maintains backward compatibility with classical systems while providing quantum-resistant security for all new data. This is implemented through a dual-algorithm approach: every cryptographic operation uses both the classical algorithm (for compatibility) and the PQC algorithm (for security). The classical algorithm can be deprecated when the client's ecosystem is ready.

Hybrid Key Exchange — UDOC TLS Configuration

TLS 1.3 with hybrid key exchange:

ECDH (X25519) — classical

+

CRYSTALS-Kyber-768 — post-quantum

=

Combined_key = HKDF(classical_secret || kyber_secret)

Both must be compromised to break session security.

Classical break: session protected by Kyber.

Quantum break (future): session protected by Kyber.

Both broken simultaneously: theoretically impossible with current cryptanalysis.

TLS CONFIG

7.3 Key Management Architecture

Master Key Storage	Client-owned Hardware Security Module (HSM). Government deployments: FIPS 140-3 Level 3 HSM required. Enterprise deployments: FIPS 140-2 Level 2 minimum.
Key Hierarchy	Master Key (HSM) → Data Encryption Keys (DEK) → Key Encryption Keys (KEK) → Session Keys. Three-level hierarchy limits blast radius of any single key compromise.
Key Rotation	Session keys: per-session. DEKs: 90-day automatic rotation. KEKs: annual rotation with 30-day overlap. Master key: manual rotation, requires 2-of-3 custodian authorisation.
G.O.D.S Access	G.O.D.S engineers have zero access to client master keys. Support access to decrypted data requires signed client authorisation, is time-limited (4 hours), and is fully logged.
Breach Response	Automated detection of key material exposure triggers immediate key rotation, session invalidation, audit notification, and incident response workflow — target: <15 minutes from detection.

SECTION 8

Audit Trail & Cryptographic Verification

Tamper-evident record architecture

8.1 Merkle Tree Audit Architecture

The UDOC audit log uses a Merkle tree structure to provide cryptographic proof of audit trail integrity without requiring full log disclosure. Each audit event produces a leaf node in the Merkle tree. The tree root hash is published at regular intervals (default: every 15 minutes) to provide a publicly verifiable checkpoint. Any modification to any historical audit event will produce a detectable change in the Merkle root.

Merkle Proof Verification — ExamplePYTHON

```
def verify_decision_in_audit(decision_id, audit_root_hash):  
    """  
    Returns (True, path) if decision is in audit log.  
    Raises TamperDetected if Merkle path is invalid.  
    """  
    leaf = audit_db.get_leaf(decision_id)  
    path = audit_db.get_merkle_path(decision_id)  
    computed_root = compute_merkle_root(leaf, path)  
    if computed_root != audit_root_hash:  
        raise TamperDetected(f"Audit record {decision_id} has been modified")  
    return True, path
```

Figure 8.1 — Merkle tree verification function. Any modification to any log entry is detectable.

8.2 Audit Event Schema

Event ID	Globally unique, cryptographically derived from event content + timestamp
Actor	System ID for automated events · Authenticated user ID for human actions · Both for human-reviewed decisions
Action Type	DECISION · OVERRIDE · REGISTRATION · CONFIGURATION · ACCESS · EXPORT · ROTATION · BREACH · ESCALATION
Object	The entity acted upon — AI system ID, decision ID, data record ID, or configuration key
Timestamp	Nanosecond precision UTC · Signed with CRYSTALS-Dilithium · NTP-synchronised with government time server
Outcome	APPROVED · DENIED · ESCALATED · SUSPENDED · REGISTERED · MODIFIED · EXPORTED
Metadata	Risk classification at time of event · Jurisdiction · Regulatory framework applicable · Human oversight flag
Integrity Hash	SHA-3-256 of all other fields + previous event hash (chain linking)
Digital Signature	CRYSTALS-Dilithium signature from signing service · Verifiable against published UDOC public key registry

SECTION 9

Deployment Architecture

Sovereign, enterprise, and hybrid deployment models

9.1 Deployment Models

Model	Description	Target Client	Sovereignty Level
National Backbone	Deployed on government-owned infrastructure. Full air-gap option. G.O.D.S has no ongoing access without client authorisation.	National governments	Maximum — full sovereign control
Sovereign Cloud	Deployed on client-contracted sovereign cloud (in-country data centre, government-approved operator). G.O.D.S manages software layer only.	Regulated industries, parastatal	High — data residency guaranteed
Private Enterprise	Deployed on client's private cloud or on-premise. Client manages infrastructure. G.O.D.S provides software and support.	Banks, insurers, large enterprise	High — client-controlled
SaaS (Restricted)	G.O.D.S-managed infrastructure. South Africa-based sovereign data centres only. Not available for SOVEREIGN-CRITICAL data classification.	SME · emerging market entry	Standard — SA jurisdiction only

9.2 Infrastructure Requirements — National Backbone

Compute (minimum)	32 vCPU · 128GB RAM per node · 3-node minimum cluster · 10Gbps networking
Storage (minimum)	10TB NVMe SSD for hot data · 100TB archival for audit log (10-year retention) · Geographic redundancy for disaster recovery
Network	Air-gap option: dedicated MPLS or physical isolation · Internet-connected option: TLS 1.3 only · No public internet exposure for SOVEREIGN-CRITICAL data
HSM	FIPS 140-3 Level 3 Hardware Security Module · Dual HSM for redundancy · Annual FIPS recertification
Power	Dual power feed · UPS 4-hour minimum · Generator backup · N+1 redundancy on all critical components
Physical Security	ISO 27001 certified data centre · CCTV · Access logs · Two-person integrity for critical operations · Annual physical penetration test

9.3 High Availability Architecture

UDOC's 99.9% uptime SLA (maximum 8.76 hours downtime per year) is achieved through multi-zone redundancy, automated failover, and zero-downtime deployment pipelines. For government backbone deployments, the target is 99.95% (maximum 4.38 hours per year).

Load Balancing	HAProxy with health checks · Automatic unhealthy node removal · Session affinity where required
Database HA	PostgreSQL streaming replication · Patroni automatic failover · <30 second RTO for database failover
Kubernetes	Multi-zone cluster · Pod disruption budgets prevent simultaneous node maintenance · Rolling updates with automatic rollback
Backup Strategy	Continuous WAL archiving · 4-hour snapshot frequency · Cross-zone replication · Monthly restore test · 1-hour RPO target
Monitoring	Prometheus metric collection · Grafana dashboards · PagerDuty alerting · 5-minute SLA breach notification · Incident commander activation

SECTION 10

API Design & Integration

REST and event-based integration specifications

10.1 API Design Principles

UDOC exposes all platform functionality through a versioned REST API and a real-time event stream. The API is designed for integration-first: every capability available in the UDOC interface is available via API, with no functionality restricted to the UI.

Authentication	OAuth 2.0 + JWT · mTLS for government backbone deployments · API key + IP allowlist for service accounts
Authorisation	Role-based access control (RBAC) · Attribute-based for fine-grained data access · All permission checks logged
Versioning	/v1/ · /v2/ · Breaking changes require 12-month deprecation notice · No breaking changes within a major version
Rate Limiting	Per-client quotas · Burst allowance for compliance exports · Government backbone: no rate limit on internal network
Documentation	OpenAPI 3.1 spec auto-generated · Interactive docs at /docs · Postman collection published · Changelog maintained

10.2 Key API Endpoints

Endpoint	Method	Description	Auth Level
/v1/registry/systems	GET/POST	List all AI systems or register new system	ADMIN, COMPLIANCE
/v1/registry/systems/{id}	GET/PATCH	Retrieve or update registered system details	ADMIN, COMPLIANCE
/v1/decisions	POST	Submit decision event for governance processing	SERVICE
/v1/decisions/{id}	GET	Retrieve decision record with explainability output	COMPLIANCE, CITIZEN
/v1/decisions/{id}/verify	GET	Cryptographic verification of audit record	PUBLIC
/v1/audit/trail	GET	Query audit log with filters	COMPLIANCE, ADMIN
/v1/compliance/report	POST	Generate compliance report (specify framework)	COMPLIANCE, ADMIN
/v1/bias/alerts	GET	Retrieve active bias detection alerts	COMPLIANCE, ADMIN

/v1/oversight/queue	GET	Human review queue for required oversight	REVIEWER
/v1/data/lineage/{hash}	GET	Retrieve full data lineage record	COMPLIANCE, ADMIN

SECTION 11

Compliance Framework Alignment

Regulatory mapping across key jurisdictions

Framework	Jurisdiction	Key Requirements	UDOC Coverage
EU AI Act	European Union	Risk classification, transparency, human oversight, technical documentation, post-market monitoring	Full — all High/Critical requirements addressed by governance engine
GDPR Art.22	European Union	Right to explanation for automated decisions, human review right	Full — explainability engine + human oversight enforcement
ISO/IEC 42001	International	AI management system, risk assessment, continuous improvement	Full — registry, risk engine, bias monitoring, continuous reporting
NIST AI RMF	United States	Govern, Map, Measure, Manage framework	Full — all four function categories implemented
OECD AI Principles	International	Transparency, accountability, security, inclusive growth	Full — aligned in governance engine design
POPIA	South Africa	Data privacy, consent, breach notification, data minimisation	Full — sovereign key management, lineage, breach alerting
AU AI Continental Strategy	Africa	Sovereignty, inclusive AI, capacity building	Full — sovereignty-first architecture, SETHS integration

SECTION 12

Performance Specifications & SLAs

Technical performance guarantees

<50ms Governance pipeline latency overhead	99.9% Platform uptime SLA (enterprise)	99.95% Platform uptime SLA (govt backbone)	<500ms Audit record verification time	<60s Compliance report generation
---	--	--	--	--

Metric	Target	Measurement Method	SLA Penalty
API response time (p95)	<200ms	Prometheus histogram	Service credit if >500ms sustained
Governance pipeline latency	<50ms	Decision event tracing	Review if >100ms sustained
Audit write latency	<100ms	Kafka consumer lag monitoring	Alert if >500ms
Compliance report generation	<60s standard	Automated test suite	Service credit if >300s
Platform uptime	99.9% monthly	External uptime monitoring	5% monthly credit per 0.1% below target
Recovery Time Objective	<4 hours	DR test (quarterly)	Priority support activation
Recovery Point Objective	<1 hour	Backup verification	Priority support activation
Security patch deployment	<72 hours (critical)	Patch log	Mandatory escalation if exceeded

SECTION 13

Implementation Roadmap

18-month technical delivery plan

Phase	Timeline	Deliverables	Milestone
Phase 0 — Foundation	Month 1–3	Legal entity · IP registration · Core team hire (6 engineers + architect) · Cloud infrastructure provisioned · CI/CD pipeline · Security framework	Infrastructure live
Phase 1 — MVP	Month 3–6	AI Registry MVP · Basic risk classification · Compliance report generator (SA POPIA) · Audit trail v1 · REST API v1 · Frontend dashboard v1	First enterprise demo
Phase 2 — Governance Engine	Month 6–9	Explainability engine · Bias detection (statistical parity + counterfactual) · Human oversight enforcement · EU AI Act compliance module · Advanced audit (Merkle tree)	First paying client
Phase 3 — Sovereign Ready	Month 9–12	PQC encryption full deployment · Sovereign deployment package · Government procurement documents · National backbone pilot · SETHS integration API	Govt pilot live
Phase 4 — Scale	Month 12–18	Multi-client SaaS platform · Self-service onboarding · Automated compliance testing · Performance optimisation · SOC2 Type II audit · ISO 27001 certification	Series A ready

SECTION 14

Commercial Architecture

Revenue model and pricing specification

R3M–R7M ACV enterprise SA market	\$250K–\$500K ACV enterprise International	75–85% Gross margin at scale	8–12× Revenue valuation multiple
--	--	--	--

Product Tier	Target	Pricing Model	Core Entitlements
UDOC Enterprise	Banks, insurers, large regulated corps	Annual contract R3M–R7M base + usage	10 AI systems · 1M decisions/yr · Full compliance suite · Dedicated support
UDOC Government	National and provincial government	Annual contract R15M–R50M	Unlimited systems · Unlimited decisions · Sovereign deployment · National backbone ready
UDOC National Backbone	National government — full deployment	Infrastructure + licensing R50M–R500M	Full sovereign deployment · Air-gap option · All modules · 5-year minimum
UDOC International	Regulated industries outside SA	SaaS subscription \$250K–\$500K/yr	Cloud SaaS · 5 AI systems · GDPR/EU AI Act modules · Standard support
UDOC API (Starter)	SME, startups, government pilots	Usage-based R15K–R80K/yr	API access only · 100K decisions/yr · Basic compliance module

Gross Margin Target	Year 1–2: 50–60% (development overhead). Year 3+: 75–85% (SaaS leverage). Year 5: 80–85%.
ARR Projection Yr 3	R45M–R210M (15–30 enterprise clients at average R7M ACV)
Govt Backbone Yr 5	R150M–R500M from 3–10 national backbone contracts
Licensing Revenue	UDOC methodology licensing to international system integrators — Year 5+. Target: 10–15% of total revenue.
Valuation Multiple	8–12× revenue — consistent with enterprise governance SaaS at institutional scale. Comparable: Palantir (FY2021), Archer (compliance SaaS), Diligent Corp.

SECTION 15

Appendices

Reference material

Appendix A — Glossary

Term	Definition
CRYSTALS-Kyber (ML-KEM)	NIST-standardised post-quantum key encapsulation mechanism. Replaces RSA and ECC for key exchange.
CRYSTALS-Dilithium (ML-DSA)	NIST-standardised post-quantum digital signature algorithm. Replaces RSA/ECDSA for signing.
Merkle Tree	Cryptographic tree structure where each parent node is a hash of its children. Allows efficient tamper detection.
PQC	Post-Quantum Cryptography. Cryptographic algorithms resistant to attacks by quantum computers.
SOVEREIGN-CRITICAL	UDOC data classification for national security, identity, and defence data. Requires air-gap deployment.
Zero Trust	Security model requiring continuous verification of every entity attempting to access resources.
HSM	Hardware Security Module. Dedicated hardware for secure cryptographic key storage and operations.
WORM	Write Once Read Many. Storage architecture preventing modification or deletion of written data.
RTO / RPO	Recovery Time Objective (time to restore) and Recovery Point Objective (acceptable data loss window).

Appendix B — Standards References

Standard	Reference	UDOC Alignment
NIST FIPS 203	ML-KEM (Kyber)	Full — primary KEM algorithm
NIST FIPS 204	ML-DSA (Dilithium)	Full — primary signature algorithm
NIST FIPS 205	SLH-DSA (SPHINCS+)	Full — long-term archive signing
NIST AI RMF 1.0	AI risk management	Full — all 4 functions implemented
ISO/IEC 42001	AI management systems	Full — registry and governance engine
EU AI Act 2024	AI regulation	Full — risk classification and audit
ISO/IEC 27001	Information security	Target certification Month 18

SOC 2 Type II	Security & availability	Target certification Month 18
---------------	-------------------------	-------------------------------

This document is a technical specification. G.O.D.S reserves the right to update technical specifications as the platform develops. The governance principles — sovereignty, auditability, quantum-readiness, zero foreign dependency — are non-negotiable constraints that will not change across platform versions.

For full technical due diligence, including architecture review sessions, security assessment, and proof-of-concept deployment terms, contact G.O.D.S Holdings (Pty) Ltd.