



Technical Portfolio

Francisco Porcel Vila · Selected production systems and published work

ceo@optimaquantum.com | WhatsApp +34 680 402 180 | linkedin.com/in/cescoporcel

Everything here is publicly reachable or demonstrable on request. Each entry carries a live URL you can open now, or metrics taken from the production database. This is a deliberate selection of nine systems from a portfolio of more than a hundred.

AI SYSTEMS IN PRODUCTION

WhatsIA

whatsia.io

Multi-tenant conversational AI platform. Founded, architected and built end to end.

- **11 autonomous agents** running 24/7 as independent services: prospecting, qualification, profiling, message generation, reply detection and two-level response.
- **143,718 executions** over 100 days of continuous operation, **267,962 records processed**, on multi-provider routing rather than a single-vendor contract.
- Isolated multi-tenant architecture: dedicated database per client, per-tenant rate limiting, independent AI personas. Live in Spain, the UAE and Thailand.
- Every run logged with duration, outcome and spend, so the unit economics of the fleet are measured rather than estimated.

InfoBitcoin

infobitcoin.es

On-chain intelligence platform combining data engineering with generative AI, running on Bitcoin Core full nodes I operate.

- Real-time analytics: mempool radar detecting large transfers before block inclusion, wallet attribution across **680+ identified entities**, dormant-coin detection.
- Autonomous LLM pipeline that ingests market and news feeds, scores them through model-based evaluation and publishes to web, Telegram and X without human intervention.
- Time-series platform on InfluxDB at 15-second resolution, Grafana dashboards, and a subscription product with authenticated premium panel.
- **Cut operating API cost by 92%** by instrumenting usage telemetry and removing spend that produced no published output.
- Rebuilt observability to alert on **silence rather than errors**, after a 30-day outage passed undetected by conventional error monitoring.

RETRIEVAL AND MEMORY SYSTEMS

Business Brain

Internal · [architecture on request](#)

Retrieval-augmented generation over a live document corpus, in daily operational use.

- RAG service over **1,744 indexed documents**: chunking, embeddings, vector storage and semantic retrieval, wired to a document management backend as source of truth.
- Built on sentence-transformer embeddings with an embedded vector store, chosen so the system runs on modest hardware with no external vector database and no data leaving the estate.

Elephant Memory

Internal · [MCP server](#)

- Indexes **~1,750 AI conversations** into searchable semantic memory, exposed to agents through the Model Context Protocol.
- Solves the problem every enterprise AI programme hits at month six: institutional knowledge trapped in conversation history that nobody can retrieve.

Multi-provider LLM gateway

Internal infrastructure

- Routing layer between my agents and multiple model vendors, with automatic failover when a provider degrades, errors or exhausts credit.
- Migrated production agent traffic off paid single-vendor inference onto a routed chain with validated fallback, without changing the agents themselves.
- This is the layer most enterprises skip, and the reason their AI programmes end up hostage to one vendor's pricing and uptime.

OPEN SOURCE

github.com/optimaquantum

7 public repositories

claude-critical-rules-mcp

MCP server enforcing operational guardrails on an AI coding agent, built from analysis of real recurring failures so the agent cannot repeat them.

win11-admin-skill

Agent skill for Windows 11 administration and hardening: registry, services and security configuration.

bitcoin-quantum-checker

Client-side tool that checks whether a Bitcoin address has exposed its public key to a future quantum adversary. No key material leaves the browser.

claude-ai-best-practices

Failure-driven operating instructions for AI agents, derived from analysis of real production incidents.

PUBLISHED RESEARCH — POST-QUANTUM BITCOIN

Eight-article research series in Spanish

infobitcoin.es/bitcoin-quantum/

- Covers BIP-360 and BIP-361, ML-DSA signatures, NIST FIPS 203/204/205, comparative exposure across chains, and the Satoshi-era P2PK coin problem.
- Sourced from primary material: NIST standards, arXiv papers from Google Quantum AI, Ethereum Foundation research and IBM hardware roadmaps. Every citation verified live.
- Paired with a working checker tool, so the research is testable rather than only readable.
- **Position taken:** the real path is BIP-360 as a soft fork inside Bitcoin, not a separate corporate fork. Most coverage conflates the two.

SECURITY PRACTICE

Optima Quantum

optimaquantum.com

- **Proprietary five-phase audit methodology:** OSINT and reconnaissance, port scanning, service enumeration, parallelised vulnerability assessment, and dual executive and technical reporting.
- Executed only under **written authorisation**, with a documented seven-step ethical protocol: no production impact without notice, no exfiltration, immediate disclosure of critical findings, full cleanup.
- Penetration testing, incident response, wireless assessment and vulnerability remediation for private and public sector clients.
- Applied internally too: a dependency audit taking an agent platform from 18 known vulnerabilities to zero.

PORTFOLIO AT A GLANCE

Systems built and documented	100+ across AI, blockchain, security, data and automation. This document is a selection of nine.
Production servers operated	7 across Europe and the UAE, plus Bitcoin and Lightning nodes on three sites.
Live public products	optimaquantum.com · whatsapp.io · infobitcoin.es
Autonomous agents in production	11 in the commercial pipeline, plus scheduled agents for market intelligence, content, monitoring and research.

ML and quantitative work

Time-series forecasting, volatility modelling, foundation models for market prediction, and ML pipelines with documented negative results.

A note on source code

The internal systems above are not published. They index client documents, business records and personal data belonging to third parties, and their configuration carries production credentials, so putting them on a public or shared repository would be the wrong call regardless of how good the code is. That is a deliberate decision, not a limitation. What I prefer instead: walking through the architecture live — schema, retrieval strategy, agent handoffs, evaluation criteria, failure modes and the reasoning behind each trade-off — and demonstrating the running systems. What I will not do is move other people's data to prove a point in a hiring process, and I would apply exactly that judgement to a client's data.

On knowing when to stop

Several trading models were built, measured honestly and shut down when the evidence said they had no edge. I include that deliberately. Knowing when to kill a model matters more in enterprise AI than knowing how to build one, and it is the failure mode I most often see in adoption programmes.