PahiLabs

MAIA-PQ benchmark white paper

Why clients can trust MAIA-PQ secure communication.

A public-safe benchmark interpretation for buyers evaluating post-quantum readiness, secure workflow behavior, deployment evidence, and replay rejection.

Download PDF

Back to resources

Benchmark evidence snapshot

What buyers should notice first.

Benchmark evidence is organized across library-level, use-case workflow, Cloud Run deployment, and validation materials.

The public trust story focuses on secure message delivery, authenticated response, deployment realism, and replay rejection.

MAIA-PQ is positioned as a selective secure communication layer, not a claim that clients must replace an entire application.

The benchmark narrative avoids unsupported numeric claims where source PDFs are image-heavy or not text-extractable.

The right client path is a scoped post-quantum migration pilot around selected messages, APIs, or field workflows.

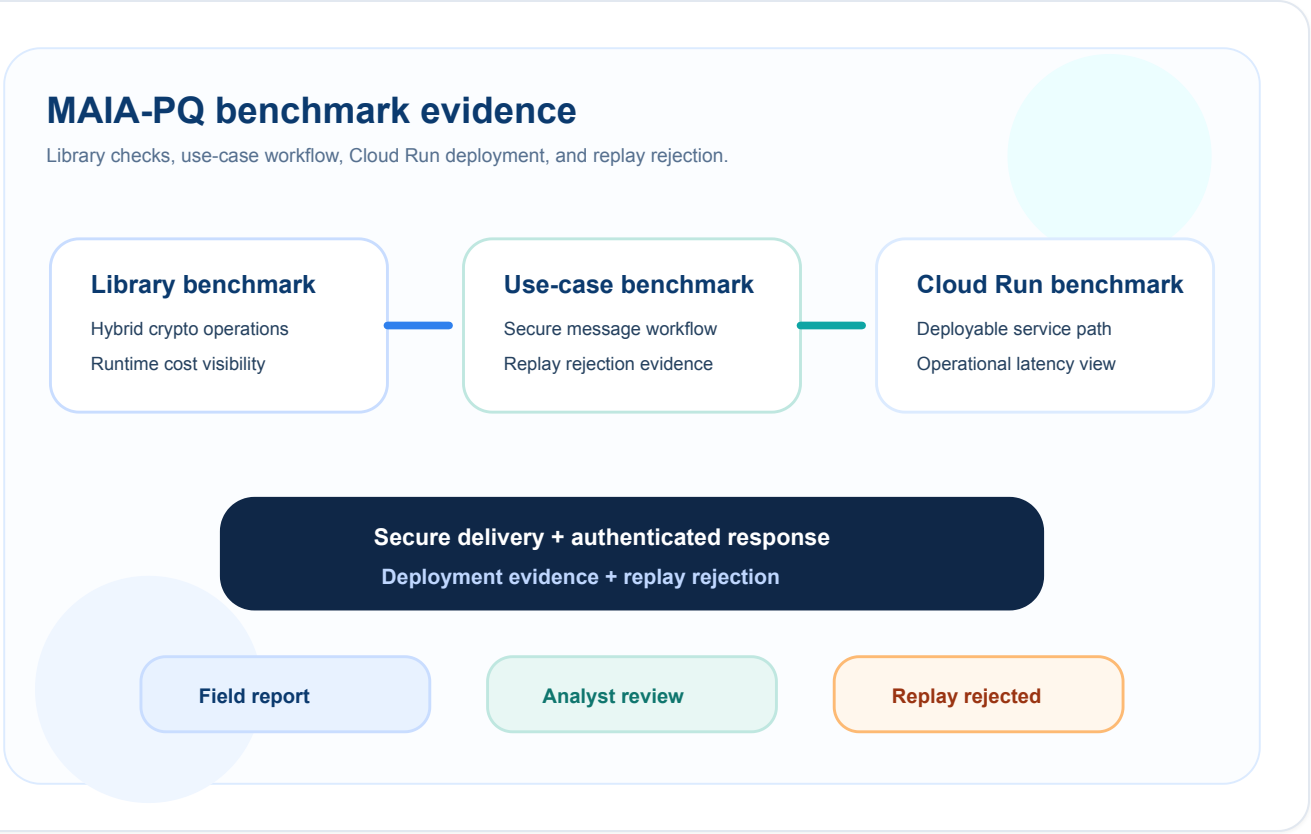

Executive summary

MAIA-PQ is PahiLabs' post-quantum secure communication direction for clients that need to begin migration planning before quantum-era risk becomes a last-minute procurement crisis. This white paper explains how the benchmark and validation materials should be communicated on the website. The goal is not to publish proprietary protocol detail or source-level implementation mechanics. The goal is to help a potential client understand why the technology is credible enough for a demo, architecture briefing, or controlled pilot.

The available benchmark package is organized around several evidence categories: library benchmark material, use-case benchmark material, Cloud Run use-case benchmark material, validation benchmark material, workflow screenshots, and product trust documents. Together, those materials support a public-safe story: MAIA-PQ has been exercised as more than an abstract cryptography discussion. It has a demonstrable communication workflow, benchmark artifacts, deployment-oriented evidence, and a replay rejection story that buyers can understand without reading internal code.

The website should be careful with numeric claims. Some MAIA-PQ source PDFs are image-heavy and do not expose reliable text extraction in the local environment. That means the public page should avoid inventing exact latency, throughput, or percentile values unless those values are explicitly available in a source that can be verified. The trust message can still be strong. A buyer does not only need numbers; they need to understand what categories were tested, what workflow was demonstrated, what failure behavior was observed, and how the technology can be evaluated inside their own environment.

Why post-quantum evidence matters now

Post-quantum migration is different from ordinary feature adoption. A client cannot wait until every system is already exposed to future cryptanalytic pressure before deciding what to do. Sensitive communications, long-lived

records, field reports, operational messages, regulated data, and critical infrastructure workflows often need a planning horizon that is longer than the procurement cycle. The practical question is not whether every organization must replace every cryptographic primitive immediately. The practical question is where the organization should start learning, testing, and reducing migration risk.

MAIA-PQ should be positioned as a way to make post-quantum readiness tangible. Many buyers have heard warnings about quantum risk, but fewer have seen a working communication flow that can be mapped to their own environment. A white paper, by itself, rarely changes a roadmap. A working demo with benchmark evidence and replay rejection behavior is more useful because it lets a technical team ask specific questions: where does the secure layer sit, what traffic should be protected first, what operational cost is introduced, and how does the system behave when an invalid message is replayed?

The benchmark package therefore matters because it shifts the conversation from abstract risk to concrete evaluation. It does not need to claim final production certification. It needs to show that the technology can be discussed as a deployable pattern. For buyers, that is often the missing step between reading about post-quantum migration and approving an actual pilot.

Evidence categories in the benchmark package

The MAIA-PQ materials contain several public-safe evidence categories. The library benchmark category helps frame the cost and behavior of the cryptographic building blocks in isolation. The use-case benchmark category connects the technology to a realistic communication path rather than treating it as a standalone library exercise. The Cloud Run use-case benchmark category shows that the technology has been considered in a deployable service

environment. The validation benchmark category connects the workflow to screenshots and proof points such as a secure message being reviewed and a replay attempt being rejected.

These categories are important because they answer different buyer questions. A cryptography engineer may want to know whether the building blocks are practical. An application architect may want to know where the client and server responsibilities sit. A cloud architect may want to know whether the pattern can run in a managed service context. A security leader may want to know whether the system rejects invalid or replayed traffic. A product buyer may want to know whether the story can be explained to non-specialists.

The public website should not collapse these categories into one vague claim. It should explicitly state that MAIA-PQ evidence spans library behavior, workflow behavior, deployment behavior, and rejection behavior. That structure makes the resource page more credible and gives clients a reason to request the underlying trust package during evaluation.

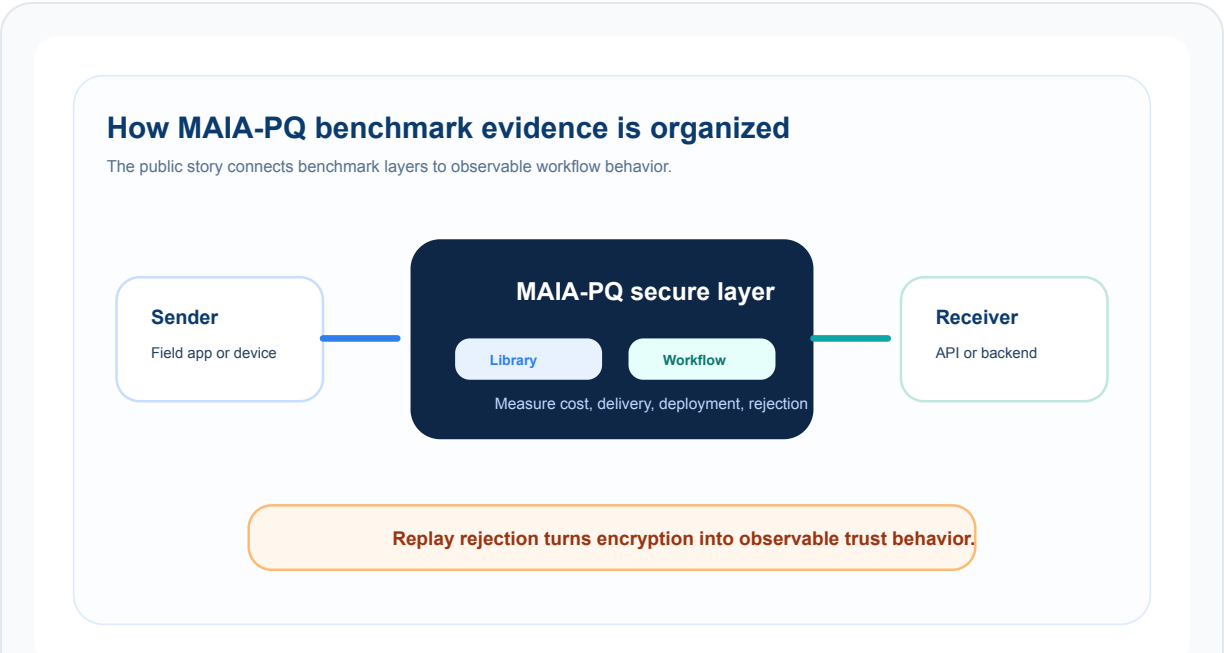

Benchmark layers tied to visible workflow evidence

This visual explains how library, workflow, deployment, and replay-rejection evidence fit together for a client evaluation.

What was demonstrated in the workflow

The workflow assets show a secure communication narrative that is easy for buyers to understand. A field user submits a protected report. An analyst reviews a secure message. A field user sends an update. A command or receiving side issues a response. A replay attempt is rejected. This sequence is valuable because it demonstrates both successful communication and failure behavior. Many demos show only the happy path. A trust product becomes more credible when it can show what happens when an invalid path is attempted.

The workflow is also useful because it maps to several client environments. A field report can represent a utility technician update, a defense or public-sector status message, a healthcare operational report, an industrial sensor exception, a supply-chain event, or a critical infrastructure instruction. The exact vertical can change, but the communication pattern remains similar: a sender needs to deliver sensitive information, the receiving side needs to verify it, and invalid repeated traffic should not be accepted as new truth.

The website should communicate this workflow in business language. The public claim should not be a detailed internal protocol trace. It should say that MAIA-PQ can wrap selected sensitive workflows with a secure communication layer and that the demonstration includes accepted messages and rejected replay behavior. That is enough for a buyer to understand the value while keeping the implementation details protected.

Benchmark metrics that matter to buyers

For MAIA-PQ, the most useful public metrics are categories rather than unsupported exact numbers. Buyers should know that PahiLabs considers library behavior, use-case workflow behavior, deployment environment behavior, and invalid-message rejection. Within a controlled pilot, these categories can become concrete measurements: setup time, secure message completion, request latency, message processing cost, replay rejection, error handling, operator experience, and deployment complexity.

This framing is more responsible than publishing numbers that cannot be verified from the available text extraction. It also reflects how serious clients evaluate security technology. A bank, public agency, utility, or enterprise security team will not decide based on one generic benchmark number. They will ask whether the technology works in their workflow, whether the performance cost is acceptable, whether the operational model is understandable, and whether the security behavior can be observed during testing.

The website should therefore introduce MAIA-PQ benchmarks as a readiness framework. The framework says: first measure the cryptographic operations, then measure the secure workflow, then measure a deployable environment, then verify rejection behavior. This is a strong message because it shows engineering discipline without disclosing protected implementation mechanics.

Replay rejection as trust evidence

Replay rejection is one of the most important public-safe claims in the MAIA-PQ evidence package. A secure communication system must not simply accept any message that looks encrypted. It must distinguish new valid communication from repeated or invalid traffic. The benchmark and validation materials include a replay rejection asset, and that should be used carefully as a trust proof point.

The public claim should be simple: the demonstration includes rejected replay behavior. It should not disclose internal replay-tracking mechanics, key schedule details, or state transition internals. A buyer does not need those details on the public website. What the buyer needs is confidence that PahiLabs understands the difference between encryption and secure communication. Encryption protects content. Secure communication also needs authentication, freshness, integrity, and clear rejection behavior.

Replay rejection also makes the demo more memorable. A visitor can understand success and failure in one sequence: a valid report is delivered, the receiving side reviews it, a response is sent, and an attempted replay is rejected. That story is much more compelling than a static statement that the product uses modern cryptography. It shows security as behavior.

How clients can integrate MAIA-PQ

MAIA-PQ should be presented as a selective secure communication layer. Clients do not need to replace every application to begin a post-quantum migration pilot. They can choose one sensitive workflow, one message path, one API exchange, or one field system and wrap that path with MAIA-PQ for evaluation. This lowers adoption friction and gives the client a practical way to learn.

The integration story has four public steps. First, identify the sensitive workflow that needs quantum-ready secure communication. Second, add the MAIA-PQ client layer to the sending application, device, or field interface. Third, add the MAIA-PQ receiving layer to the API service, backend, or command system. Fourth, route protected messages through the secure layer while existing business logic continues behind it. This is simple enough for the website and safe enough for public communication.

The pilot should measure the client's own environment. Useful pilot evidence includes message completion, operational latency, deployment complexity, error handling, operator clarity, replay rejection, logging needs, and integration impact. The existing benchmark package gives PahiLabs a starting point, while the client's pilot turns the evaluation into deployment-specific evidence.

Questions clients should ask during evaluation

A client evaluating MAIA-PQ should begin with data classification. Which messages are sensitive enough to deserve early protection? Which records have a long confidentiality horizon? Which field workflows, service-to-service exchanges, or operational APIs would create the greatest risk if captured today and decrypted in the future? These questions help a buyer avoid an unfocused migration project. MAIA-PQ is strongest when it is introduced around a selected workflow that can be measured and explained.

The second question is where the secure layer should sit. In some cases, the sender is a desktop application or web interface. In other cases, it is an IoT device, gateway, field tool, service, or backend process. The receiving side may be an API, command system, operational dashboard, or storage service. A good evaluation maps these boundaries before implementation. The goal is to protect the communication path while letting the client's existing business logic continue to operate behind it.

The third question is what evidence will convince the buyer's internal stakeholders. A cryptography team may want details about algorithms and key agreement. A cloud team may care about deployment and operational latency. A CISO may care about replay rejection, logs, and migration governance. An executive sponsor may need a simple explanation of why this project reduces future risk. The MAIA-PQ benchmark structure is useful because it supports all

of those conversations without putting protected implementation detail on the public website.

How this white paper should be used by buyers

This document should be used as a decision aid before a technical briefing. It tells the buyer that MAIA-PQ has a benchmark and validation story, but it does not ask the buyer to accept unsupported universal claims. It explains the evidence categories, the workflow narrative, the integration pattern, and the responsible limits of the public message. That is the right amount of information for a website resource because it builds confidence while preserving the deeper details for direct evaluation.

The best next step is to request a MAIA-PQ architecture and benchmark briefing. In that session, PahiLabs can map the client's sensitive communication path, explain the pilot structure, identify which benchmark categories should be repeated in the client's environment, and define what evidence should be collected. That turns the public white paper into a concrete evaluation plan.

Why clients can trust the technology at the evaluation stage

Clients can trust MAIA-PQ at the evaluation stage because it is not being presented as a vague post-quantum slogan. It has a product page, a working communication narrative, benchmark categories, validation assets, and a clear integration path. That is the right level of maturity for a controlled pilot and technical review.

Trust does not require public disclosure of intellectual property. In fact, public disclosure of protocol internals can make a website less appropriate for buyers. Enterprise visitors need enough information to decide whether the product is

worth evaluating, not enough detail to reconstruct protected implementation ideas. The right balance is to publish outcomes, evidence categories, integration patterns, and responsible limitations.

The most credible public message is this: MAIA-PQ helps organizations start quantum-readiness work around selected secure communication paths. It has evidence across library, workflow, deployment, and replay rejection categories. The next step is a scoped pilot that measures the client's own workflow. That statement is strong, accurate, and appropriate for enterprise buyers.

Limitations and responsible claims

The MAIA-PQ website should avoid claiming that the product is a certified replacement for every enterprise cryptographic system. It should avoid claiming universal performance numbers unless those numbers are published from a verified benchmark source. It should avoid implying that post-quantum migration is solved by a single library or one demo. Serious buyers know that migration requires inventory, prioritization, integration planning, operational testing, governance, and long-term maintenance.

Responsible claims make the product more credible. MAIA-PQ can say that it provides a practical demonstrator and integration path for quantum-ready secure communication. It can say that evidence exists across multiple benchmark categories. It can say that replay rejection is part of the demonstrated trust behavior. It can say that PahiLabs can support a scoped pilot. Those claims are useful and defensible.

A good public white paper should leave room for the next conversation. It should make buyers want to request the detailed trust package, not replace that package. The website version should explain why the technology deserves trust, what kind of evidence exists, and what a client can validate next.

Conclusion

MAIA-PQ should be trusted because it turns post-quantum readiness from an abstract warning into a practical evaluation path. The benchmark package shows evidence categories that matter: library behavior, use-case workflow, deployable service context, and replay rejection. The workflow assets show a story that buyers can understand. The integration model shows how a client can start without replacing an entire application.

The strongest public communication is disciplined. It should not overstate certification, universal performance, or final production maturity. It should say that MAIA-PQ is ready for a scoped technical evaluation and that PahiLabs can help clients measure the right workflow. That is the message a serious enterprise buyer needs.

The practical takeaway is simple: MAIA-PQ gives clients a way to begin post-quantum secure communication pilots now. It protects selected workflows, gives teams measurable evaluation categories, and demonstrates that invalid repeated traffic is not part of the accepted path. That is enough to earn the next meeting and begin a real technical review.

About Us

Products

Consultancy & Services

Resources & Blog

Contact & Support

Products

MAIA SSO

MAIA-IOT

MAIA-PQ

LENS

Tok2DBs

Contact Us

support@pahilabs.com
info@pahilabs.com

IPN - Building C,
Rua Pedro Nunes, 3030-199, parish of Santo António dos Olivais,
municipality of Coimbra, Portugal

Supported by

1 2 9 0 UNIVERSIDADE D COIMBRA

BUSINESS®