ARCHITECTURE / CRITICAL INFRASTRUCTURE

OCTOBER 2026

THE BANKING INDUSTRY IS DEBATING THE WRONG PROBLEM

The issue is not who owns the cloud.
The issue is why critical banking functions depend on it.

MATEUSZ CHRZANOWSKI

SYSTEM OF SYSTEMS ARCHITECT

BANKING ARCHITECTURE / TRUST / RESILIENCE / CONTROL

The recent debate in the United Kingdom about banks becoming increasingly dependent on external cloud providers has reignited an important discussion. Unfortunately, I believe it starts with the wrong question.

The issue is not who owns the cloud.

The issue is why critical banking functions depend on it in the first place.

Cloud computing is one of the most important technologies developed over the last two decades. It transformed the economics of computing, accelerated innovation and allowed organisations of every size to build digital services at a scale that was once reserved for governments and the world's largest corporations.

None of that is the problem.

There are no good or bad technologies. There are only technologies that are well chosen—or poorly chosen—for a specific purpose.

The problem begins when technology stops being a tool and becomes an ideology.

Today, far too many projects start with the same assumptions: cloud-first, AI-first, everything online. It sounds modern, looks impressive in investor presentations and creates the impression of technological maturity. Yet architecture should never be driven by fashion. It should be driven by purpose.

This distinction becomes particularly important when we move from convenience to responsibility.

A social media platform, a travel application or an e-commerce website can tolerate a fundamentally different risk profile than a bank. They do not carry the same consequences when something goes wrong. Critical infrastructure should never be designed according to the same architectural philosophy as consumer software simply because both happen to use digital technologies.

This is where I fundamentally disagree with much of today's discussion.

The real question is not whether cloud computing is good or bad. That debate is far too simplistic. The real question is whether the technology chosen for a system reflects the function that system is expected to perform.

Architecture starts with the mission.

Technology comes afterwards.

That principle has quietly disappeared from many conversations about digital transformation. Instead of asking what level of control, resilience and independence a critical system requires, we increasingly ask which technologies are currently considered modern.

Those are two completely different design philosophies.

As an infrastructure architect, I have never started a project by choosing technologies. I start by defining the critical function of the system. Only then do I decide which technologies belong inside it and, equally importantly, which technologies should never become part of its core.

Because the question is not who owns the cloud.

The question is why the critical core of a bank depends on it at all.

Trust Cannot Be Outsourced

The role of a bank is fundamentally different from that of almost any other company.

A technology company primarily protects its own intellectual property. A retailer protects its inventory. A manufacturer protects its production lines.

A bank protects something that, in most cases, does not belong to the bank at all.

The money stored in current accounts, savings accounts and investment portfolios belongs to millions of individuals, businesses and institutions that have entrusted the bank with one of the most fundamental responsibilities in modern society.

Trust.

That single fact changes everything.

Banks should never optimise critical architecture for convenience. They should optimise it for responsibility.

This is precisely why I believe many institutions are asking the wrong question. The discussion should not revolve around whether cloud providers are secure enough. The discussion should begin much earlier.

How much of a bank's critical infrastructure should depend on third parties at all?

For many organisations, outsourcing infrastructure is a rational business decision. It reduces operational complexity, accelerates deployment and allows engineering teams to focus on products rather than facilities.

Banks are different.

A bank is not simply another software company.

It is part of a country's financial infrastructure.

In my opinion, every major bank should own and operate its own cloud infrastructure, distributed across at least two highly redundant data centres designed specifically for its own operations.

Many people instinctively perceive such an investment as enormous.

It is not.

Building two world-class data centres capable of operating a bank's private cloud may require an investment measured in hundreds of millions of euros. For a major financial institution, this is not an extraordinary expense. It is the cost of owning the infrastructure on which public trust depends.

More importantly, it fundamentally changes the architecture.

CONTROL

The bank controls its own facilities.

Its own physical security.

Its own operational procedures.

Its own engineers.

Its own redundancy strategy.

Its own disaster recovery.

Its own priorities.

Control itself becomes a security mechanism.

This philosophy is not new.

Swiss banking has long been associated with maintaining a high degree of direct control over critical infrastructure and operational security. Even today, for selected high-value operations, additional human verification may complement standard authentication procedures. The principle is remarkably simple: when responsibility increases, so should the level of control.

Poland offers another interesting example.

BLIK, a mobile payment system developed by the Polish banking sector, has become one of Europe's most convenient payment methods. It allows users to make instant payments, withdraw cash, shop online and transfer money using short one-time codes directly linked to their banking applications.

Its success demonstrates something that is often forgotten in today's technology discussions.

Innovation is not defined by public cloud dependency.

Innovation is defined by solving the right problem.

Owning infrastructure, however, is only the beginning.

The Core Should Never Be Online

Owning the infrastructure is only the first step. A bank may own every server, every network device and every square metre of every data centre. It may employ its own engineers, build redundant facilities and implement the highest operational standards available today. Yet one fundamental question still remains unanswered.

Should the critical core of a banking system remain permanently connected to the rest of the infrastructure?

I do not believe it should.

This is where architecture becomes more important than technology. Modern systems are built around continuous connectivity. Information flows constantly between services, databases, authentication systems and external platforms. This architecture delivers extraordinary performance and flexibility, but it also creates an unavoidable reality: anything that remains permanently connected inevitably becomes part of the same attack surface.

The traditional response is straightforward. Add another security layer. Deploy another monitoring platform. Implement another cybersecurity solution. Those measures certainly improve resilience, but they do not change the architecture itself. They reduce risk within the existing design rather than questioning whether the design is fundamentally correct.

The more important question is much simpler.

Why should the most sensitive part of the system remain continuously reachable at all?

For decades the technology industry has become obsessed with immediacy. Everything must happen instantly. Payments. Communication. Artificial intelligence. Decisions. We have gradually accepted the idea that speed is always the highest priority.

I fundamentally disagree when critical infrastructure is involved.

Critical infrastructure should optimise for certainty before it optimises for speed. If validating a financial transaction takes five additional seconds instead of completing instantly, nothing meaningful changes in people's lives. If exceptionally sensitive operations require thirty seconds rather than three, society will continue functioning perfectly well. The purpose of critical infrastructure is not to be the fastest. Its purpose is to remain trustworthy.

Only after accepting that principle does a different architecture become possible.

The operational environment should continue processing transactions, serving customers and managing millions of requests every day. The validation core, however, should operate under an entirely different philosophy. Instead of maintaining a permanent connection, transaction requests should first pass through a dedicated buffer layer. Before any communication with the validation environment begins, the networking module responsible for external communication should be physically disconnected from power—not logically, but physically.

I would take this separation even further. Rather than allowing the buffer server to communicate directly with the validation servers, both systems should exchange information through a dedicated shared storage layer functioning as a secure digital mailbox. The buffer writes the request and disconnects. The validation environment connects, reads the request, validates it, writes the response and disconnects. Only then does the buffer reconnect to the operational network and return the validated result.

At no point does the operational environment maintain a live communication path with the validation core.

That is not simply another cybersecurity mechanism.

It is a completely different architectural philosophy.

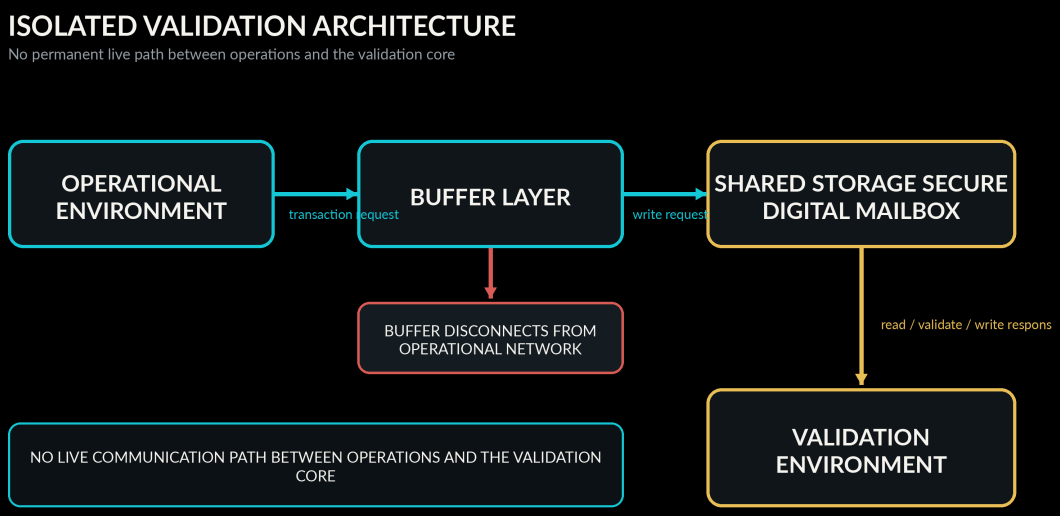

Physical disconnection changes the attack surface before another cybersecurity layer is added.

Architecture Changes the Economics of an Attack

At this point someone will inevitably ask the obvious question.

Can such a system still be compromised?

Of course it can.

Every system can.

Every safe can eventually be opened.

Every bank can be robbed.

Even a national mint can be breached. La Casa de Papel built an entire global phenomenon around exactly that idea.

The objective of architecture has never been to create a system that is impossible to break. Such systems do not exist. The objective is to design a system in which breaking it becomes progressively more difficult, more expensive and significantly less rational than attacking another target.

One of the first warning signs for any architect should be hearing someone claim that a system is impossible to break.

It is not a sign of exceptional engineering.

It is often a sign of excessive confidence.

Every complex system contains assumptions. Every architecture contains trade-offs. Every security model has boundaries that may eventually be discovered, challenged or bypassed.

The most dangerous engineer is not the one who admits uncertainty.

It is the one who believes uncertainty no longer exists.

That is precisely why systems presented as “unbreakable” deserve even greater scrutiny. Confidence should never replace verification, because architecture is built on disciplined scepticism, not on certainty.

This is where architecture fundamentally changes the economics of an attack.

> Compromising a cloud environment is no longer enough.

> Compromising operational servers is no longer enough.

> Intercepting network traffic is no longer enough.

An attacker now faces multiple independent layers that were intentionally designed to operate under different architectural assumptions. The operational environment continues serving customers. The validation layer remains isolated. The critical trust mechanism never exposes itself through a permanent communication path. Even a successful compromise of one environment does not automatically grant access to the next.

The attack has not become impossible.

It has become exponentially more complex.

At some point the attacker no longer needs better malware. They need physical access. They need coordinated operations. They need resources that move the scenario away from everyday cybercrime and closer to something resembling a state-level operation.

That distinction matters.

Cybersecurity is often presented as an endless race between attackers and defenders. Architecture changes that relationship. Instead of endlessly adding new defensive products, it changes the battlefield itself. The attacker is forced to solve an entirely different problem.

And that is precisely what good architecture should achieve.

There is another lesson here that extends far beyond banking.

Imagine that a military command system depended entirely on a public cloud operated by an external provider. Most people immediately recognise how absurd that sounds. Critical national defence infrastructure must remain under the complete control of the organisation responsible for operating it.

So why should we apply a different standard to institutions entrusted with the savings of millions of people?

The technologies may be different.

The architectural principle is exactly the same.

Critical functions should never depend on infrastructure that the owner does not fully control.

Cloud computing remains one of the greatest technological achievements of our time. Artificial intelligence will transform entire industries. Both deserve their place in modern architecture.

But neither should become an ideology.

Technology should serve architecture.

Architecture should serve trust.

And trust should never be outsourced.