Cybersecurity Best Practice Is a Baseline, Not a Blueprint

By: Jaiz Anuar

14 July 2026 · Cybersecurity · Security Architecture · 10 min read

Best practices provide a useful starting point. But good security architecture begins by understanding the risk—not by prescribing the same control for every environment.

“Industry best practice” is one of the most frequently used phrases in cybersecurity.

Consultants use it to support recommendations. Auditors use it to describe control gaps. Vendors use it to position their solutions. Management may rely on it when deciding whether an organisation is sufficiently protected.

Best practices are valuable. They represent established knowledge, lessons learned, and controls that have worked across many organisations.

But a best practice is not automatically the right architecture. That distinction matters.

Start With the Security Objective

Consider an organisation that allows third parties to access a core transaction system.

A consultant may recommend multi-factor authentication, or MFA. The recommendation sounds reasonable because passwords alone are no longer sufficient for protecting important systems.

However, the discussion should not begin with: “How do we implement MFA?”

It should begin with: What security risk are we trying to prevent?

The real objective may be to prevent third parties from obtaining direct and uncontrolled access to the core system.

Once that objective is understood, MFA becomes one possible control—not the entire answer.

The organisation can then examine several architectural options:

A web server or reverse proxy

A middleware or integration server

An SSO termination point

A tightly controlled VPN

A zero-trust access service

An API gateway

Privileged access management

Native MFA within the application

A combination of several controls

Each option may contribute to the same security outcome in a different way. The purpose of security architecture is to understand those differences.

MFA Protects the Login, Not the Entire Access Path

MFA is an important security control. It makes it more difficult for an attacker to access a system using only a stolen password.

But MFA mainly strengthens authentication. It does not automatically control everything that happens before and after the user logs in.

A user may successfully complete MFA but still have excessive privileges. The user's device may already be compromised. A valid session may be stolen. The application may contain exploitable vulnerabilities. Another network path may allow the control to be bypassed entirely.

MFA alone does not necessarily:

Prevent direct connectivity to the core environment

Restrict access to approved application functions

Inspect malicious requests

Isolate third-party sessions

Protect insecure application interfaces

Detect misuse by an authorised user

Provide transaction-level monitoring

Prevent alternative access paths

This does not make MFA ineffective. It simply means MFA should not be mistaken for a complete security architecture.

One Objective Can Have Multiple Control Options

There is rarely only one way to achieve a security objective.

If external users must not connect directly to a core system, a controlled web server could become the access point. Requests can be inspected, authenticated, and logged before they are allowed to reach the internal application.

A middleware server could expose only the functions required by the third party while keeping the core system isolated. An SSO termination point could centralise authentication, access policies, and user

revocation.

A VPN could provide encrypted access, although it must be carefully segmented to avoid giving users unnecessary network reach. An API gateway could be used when one organisation’s application needs to communicate with another organisation’s system.

Native MFA within the application may provide immediate protection when larger architectural changes cannot be implemented quickly.

None of these options is automatically correct in every situation. Each creates different dependencies, operational requirements, and residual risks.

A web server can still be bypassed if backend access remains open. A VPN can create excessive network exposure. An API gateway can become a simple forwarding proxy if it does not enforce authorisation and validate transactions. Native MFA may still leave the application directly exposed.

The control must therefore be assessed as part of the complete access path.

Human Access Is Not the Same as System Integration

Another problem with generic recommendations is that they sometimes treat every type of access in the same way. A person accessing an application is different from one system communicating with another.

For human access, appropriate controls may include:

SSO

MFA

Role-based authorisation

Device-posture validation

Session monitoring

Time-bound access

For system-to-system integration, interactive MFA may not be relevant because no person is waiting to approve an authentication request.

The appropriate controls may instead include:

Application or workload identity

Mutual TLS

Managed certificates and secrets

OAuth tokens

API-level authorisation

Payload validation

Rate limiting

Transaction logging

Administrative support access creates another scenario. It may require privileged access management, approval workflows, session recording, and just-in-time privileges.

The security objective may be similar, but the appropriate control pattern can be very different. Good architecture recognises those differences before recommending the solution.

Attackers Also Understand Best Practices

Best practices are not secret. Sophisticated attackers study the same technologies, frameworks, and reference architectures used by defenders. They understand how organisations commonly implement MFA, VPNs, firewalls, endpoint protection, and identity systems.

They also know where organisations frequently stop.

The firewall has been deployed, but the rules are rarely reviewed. MFA has been enabled, but recovery procedures remain weak. The API gateway exists, but an alternative backend connection bypasses it. The VPN is protected, but it provides broader network access than necessary.

The organisation may believe it has followed best practice. The attacker looks for the assumptions surrounding that practice.

This does not mean established practices should be rejected. It means they must be adapted, layered, tested, and continuously challenged.

A control becomes valuable not because it appears in a framework, but because it remains effective against the threats facing the organisation.

The Entire Access Path Must Be Examined

Security architecture should examine the complete journey from the requester to the protected system. It should establish:

Who or what is requesting access?

Where does the connection enter?

How is identity verified?

What device or system is initiating the request?

Where does the connection terminate?

What functions are permitted?

What information is inspected and recorded?

Can access be revoked immediately?

Is there another path that bypasses the control?

What happens if the control fails?

This is more valuable than asking only whether MFA, SSO, or an API gateway has been implemented.

A control can exist and still fail to achieve its intended outcome. The presence of a control demonstrates implementation. Its ability to prevent, detect, and contain an attack demonstrates effectiveness.

Consultants Should Present Options, Not Only Prescriptions

Consultants provide valuable external experience. They can identify weaknesses, introduce new perspectives, and help organisations avoid mistakes observed elsewhere.

But “this is industry best practice” should not be the end of the recommendation.

A strong recommendation should explain:

What specific risk is being addressed?

What security outcome must be achieved?

What architectural options were considered?

Why is the selected option suitable for this environment?

What assumptions must remain true?

How could the control be bypassed or misused?

What residual risks will remain?

How will effectiveness be tested and monitored?

This allows management to evaluate risk reduction alongside cost, complexity, operational capability, and business impact.

The consultant may recommend the preferred option, but the organisation must retain ownership of the decision.

Interim Controls Can Still Be Reasonable

Not every organisation can implement the ideal architecture immediately. A legacy application may not support modern SSO. An API gateway may require significant redesign. A new access platform may take months to implement.

In such situations, native MFA or another compensating control may still be reasonable as an interim measure.

But the arrangement should have:

A clearly defined security objective

Documented limitations

Supporting compensating controls

An accountable owner

A specific target date

A migration path to the strategic solution

An interim control should reduce immediate exposure while the stronger architecture is being developed. It should not quietly become permanent simply because the original audit finding has been closed.

Final Thought

Best practices matter. They provide organisations with a baseline and help prevent known security mistakes. But they cannot replace threat modelling, architectural analysis, or professional judgement.

The question is not simply: “Are we following best practice?”

The better question is: Does this architecture achieve the required security outcome, and can we prove that it works?

MFA may be the right control. An API gateway may be more appropriate. A termination point, middleware layer, privileged-access platform, or combination of controls may provide stronger protection.

The answer depends on the user, the system, the access path, the threat, and the business requirement.

Cybersecurity best practice tells us where to begin. It should never tell us where to stop thinking.

Question assumptions.

Share knowledge.

Build trust.