

Insights 2026

Four briefings on sovereignty, Secure by Design, zero trust in operational technology, and the governance model that joins them.

TIS-BRF-001 • Sovereign Cloud

Five questions that separate genuine sovereignty from a region setting

TIS-BRF-002 • Secure by Design

Secure by Design, three years in: what assurance teams actually look for

TIS-BRF-003 • Zero Trust

Zero trust for critical infrastructure: assume breach where downtime isn't an option

TIS-BRF-004 • Governance

One governance model, three disciplines

Five questions that separate genuine sovereignty from a region setting

Jurisdiction, control plane, personnel, keys and exit

Cloud location is important, but it is not sovereignty. A region marker tells you where a provider intends to process or store particular data. It does not, by itself, tell you which law can reach the provider, who can operate the service, where privileged instructions are executed, who can recover encryption keys, or whether the customer can leave on workable terms.

Sovereignty is better understood as the ability of an organisation or nation to make and enforce decisions about its data, technology and operations without an unacceptable external dependency. It is not an absolute state. It is a risk position that varies by workload, threat, classification, sector and mission consequence.

A simple distinction Data residency answers “where?” Genuine sovereignty also answers “under whose authority, operated by whom, using which dependencies, with whose keys, and with what route out?”

1. Which jurisdictions can reach the service?

Start with the complete legal and corporate chain, not the address of the data centre. The contracting entity, parent company, subcontractors, support organisations and technology licensors may sit under different legal regimes. A locally incorporated supplier may still depend on a foreign parent for software, updates, finance or operational authority. Conversely, an international supplier can sometimes build strong local controls. The point is to test the arrangement, not the flag on the website.

The analysis should identify applicable law, routes for lawful access, conflict-of-law scenarios, insolvency exposure, export controls and the contractual response to government demands. For higher-consequence workloads, counsel should test realistic scenarios rather than merely list statutes.

- Which entity contracts with us, and who ultimately owns or controls it?
- Which jurisdictions apply to the provider, its parent, material subcontractors and support chain?
- Can a foreign authority compel access, secrecy, service restriction or transfer? What would the provider do?
- Which functions rely on licences, updates or approvals that could be withdrawn?

2. Who controls the control plane?

The control plane is the machinery used to create, configure, monitor and administer the service. If data sits in the UK but administrators, orchestration services, telemetry, software repositories or global support systems remain outside the customer's effective control, the service may be resident without being operationally sovereign.

Map the management path end to end: identity, privileged access, orchestration, software supply chain, monitoring, backup, recovery, billing and support. Ask what still works if international connectivity is lost or the parent service withdraws cooperation. Sovereignty becomes real when local operators have the authority, tooling, knowledge and supply chain required to sustain the agreed mission.

- Where are management services and administrative logs processed?
- Can global engineering or support staff change the platform, directly or through automation?
- Can the service be patched, recovered and safely operated during isolation from global dependencies?
- Who can suspend tenants, revoke accounts, alter policy or push software?

3. Who are the people behind privileged access?

A sovereign operating model needs more than local help-desk coverage. It needs explicit control over who can exercise privilege, from where, under whose employment and vetting, using which devices and approval paths. Follow-the-sun support can be valuable, but it should not be presented as sovereign if sensitive actions can move outside the defined personnel boundary.

The strongest patterns combine named roles, least privilege, just-in-time elevation, separation of duties, approved access locations, session recording and locally accountable decision-makers. They also address the harder question: whether local personnel possess enough engineering depth to operate and recover the service without an undeclared dependency on overseas experts.

- Are all privileged roles identified, including vendor engineers and automated service identities?
- What nationality, residency, clearance, employment and location constraints apply?
- Can any support route bypass local approval, recording or monitoring?
- Is critical knowledge genuinely held in-country, or merely fronted by an in-country team?

4. Who controls the keys – and the ability to use them?

Customer-managed encryption keys are valuable, but the phrase can conceal important details. A provider may host the key service, retain recovery material, control the code that requests decryption or hold privileged access to the plaintext after decryption. Key ownership is meaningful only when key generation, storage, use, backup, rotation, revocation and destruction form a coherent control boundary.

For the most sensitive use cases, consider whether keys are held in customer-controlled hardware security modules; whether quorum or split knowledge is required; whether the provider can alter the workload or policy to obtain plaintext; and how keys are recovered without creating a universal override. The desired result is cryptographic control tied to operational governance, not simply a checkbox labelled “bring your own key”.

- Who generates, stores, backs up, rotates, revokes and destroys each class of key?
- Can the provider or a subcontractor decrypt data, change the code around it, or compel the key service to act?
- What happens to availability if the external key service, network or authority is unavailable?
- Can the customer independently prove destruction or denial of access?

5. Can we leave – in time, at known cost, with a working service?

Exit is where many sovereignty claims fail. Contractual ownership of data is not the same as practical portability. A credible exit position covers data, metadata, identities, configurations, automation, logs, security history, keys, software licences and the knowledge needed to rebuild operations elsewhere.

The exit plan should be exercised before it is needed. Measure extraction time and egress cost; test format compatibility; rebuild a representative service; validate deletion; and identify functions that cannot be moved without redesign. For critical services, include a continuity route for supplier failure, geopolitical restriction, prolonged disconnection and hostile withdrawal – not just orderly contract expiry.

- What must be extracted besides primary data?
- How long would exit take at realistic bandwidth and scale, and what would it cost?
- Which proprietary services create the greatest rebuild burden?
- When was the exit plan last tested, and what evidence was retained?

A practical sovereignty scorecard

Dimension	Evidence to request	Warning sign
Jurisdiction	Corporate chain; applicable-law analysis; subcontractor register; compelled-access process	The answer stops at “UK data centre”
Control plane	Management architecture; dependency map; isolation mode; local recovery procedure	Global services or staff can act without local control
Personnel	Role matrix; vetting/location rules; PAM records; skills and succession plan	Local service desk, overseas engineering authority
Keys	Key hierarchy; HSM ownership; ceremonies; recovery and destruction evidence	“Customer-managed” but provider can still reach plaintext
Exit	Tested exit plan; formats; timings; costs; deletion evidence; continuity scenario	Data export offered, operational rebuild ignored

Board-level question If a foreign parent, overseas support function or external control service became unavailable or unwilling tomorrow, which missions would stop – and for how long?

Secure by Design, three years in: what assurance teams actually look for

The artefacts that pass first time, the gaps that stall programmes, and how to build the evidence chain as you go

The MOD's 2023 Secure by Design shift was deliberately more than a change of paperwork. It moved accountability for cyber risk toward senior responsible owners, capability owners and delivery teams, and replaced a late accreditation event with assurance through the lifecycle. Public MOD guidance describes security as a shared responsibility, supported by independent assessment at key stages rather than a final gate bolted onto delivery.[1][2]

That changes what "good" looks like. Assurance teams do not primarily need a large document set. They need confidence that the programme understands its mission, architecture, threat and risk; has made deliberate design choices; has implemented and tested the resulting controls; and has put residual risk in front of the person authorised to own it.

The evidence chain

Link	Question it must answer	Typical evidence
Mission and context	What must the capability achieve, protect and survive?	Service outcomes; critical functions; impact tolerances; data and safety context
Threat	Who might act, with what capability, against which paths?	Threat assessment; abuse cases; attack paths; intelligence assumptions
Risk	What could happen, why does it matter and who owns it?	Risk scenarios; likelihood/impact rationale; named owners; treatment decisions
Architecture	Where are trust boundaries, dependencies and failure modes?	Current and target diagrams; data flows; identity and admin paths; dependency map
Controls	Which measures reduce each material risk?	Control design; requirements traceability; configuration standards; supplier obligations
Verification	Did the controls work in the real build?	Test results; code/configuration evidence; vulnerability findings; exercise outcomes
Decision and operation	What remains, who accepted it and how is it maintained?	Residual-risk decision; operating model; monitoring; change and incident records

Artefacts that create confidence

A concise security case

The security case should make a structured argument: the capability is acceptably secure for a defined context because stated claims are supported by evidence. It should acknowledge uncertainty and exceptions. A clear twenty-page argument with live evidence is more useful than two hundred pages of generic control text.

Architecture that matches the deployed service

Assurers need more than a high-level solution diagram. Useful views show data movement, trust boundaries, administrative paths, external dependencies, identity flows, logging, key management, backup and recovery. The names in the diagrams should reconcile with inventories, tests and operating procedures. If the diagram and the build disagree, the build wins.

Threat and risk that drive engineering

Threat assessments should be specific enough to change a design choice. Risk statements should describe cause, event and consequence, identify the affected mission or asset, and have accountable owners. A register of generic cyber risks, all scored "medium", rarely demonstrates understanding.

Traceable security requirements

Each important security requirement should be traceable backwards to a risk, policy or operational need and forwards to design, implementation and test evidence. This does not require a perfect toolchain. It requires stable identifiers and disciplined links. Traceability allows an assurer to sample a claim and follow it without guesswork.

Evidence from the delivery pipeline

The strongest programmes make evidence a by-product of engineering: reviewed code, signed builds, configuration scans, infrastructure-as-code checks, component inventories, test outputs, privileged-access logs and release approvals. Evidence is dated, attributable and tied to the version that will operate. Screenshots without provenance are weak; repeatable results are strong.

An operating model, not just a launch position

Security at entry to service is temporary. Assurance needs to see how vulnerabilities, changes, suppliers, identities, keys, logs, incidents, backups and recovery will be managed. It should be obvious which team acts, within what time, using which authority and information.

The gaps that stall programmes

Gap	Why it stalls assurance	Practical correction
Security starts after architecture	Material trust and dependency choices are already expensive to change	Embed security architecture and threat work in discovery and option selection

Gap	Why it stalls assurance	Practical correction
Documents describe an ideal build	Assurers cannot rely on evidence that does not match deployment	Generate diagrams and evidence from authoritative inventories and pipelines
Risk has no empowered owner	Treatment and residual-risk decisions cannot be made	Name owners early; define escalation and acceptance thresholds
Supplier assurance stops at a certificate	Service-specific dependencies and privileged paths remain untested	Obtain architecture, control and incident evidence for the actual service
Testing proves presence, not effectiveness	A control can exist yet fail against the threat scenario	Use misuse cases, attack paths, recovery tests and adversarial validation
Temporary exceptions become permanent	The live risk position diverges from the approved position	Time-box waivers; fund remediation; report ageing and expiry
Operational teams arrive late	Controls cannot be sustained after go-live	Bring operators into design, exercises, handover and acceptance

Build the evidence as you go

1. Define the mission, critical functions, data and impact tolerances before selecting controls.
2. Create an architecture baseline and record material decisions, alternatives and assumptions.
3. Model credible threats and express risks in operational language that owners can decide.
4. Map risks and obligations to security requirements with stable identifiers.
5. Attach implementation and verification evidence to each requirement during delivery.
6. Run assurance reviews at decision points while change remains affordable.
7. Transfer the evidence chain into operations and keep it alive through change, incidents and testing.

What “pass first time” really means It does not mean no findings. It means the programme can explain its system and risk position, produce trustworthy evidence quickly, show ownership of gaps and make proportionate decisions without reconstructing its history.

Ten-minute assurance test

Choose one material risk at random. Ask the programme to show the threat or failure scenario, affected mission outcome, owner, design response, implemented control, latest test result, outstanding finding and residual-risk decision. If this chain cannot be followed in ten minutes, the programme has an evidence-management problem even if every expected document exists.

Zero trust for critical infrastructure: assume breach where downtime isn't an option

Applying identity-first architecture to OT estates that cannot be re-platformed - sequencing, segmentation and what to leave alone

Zero trust removes implicit trust from network location and requires access decisions to be based on verified identity, device or service condition, policy and context. NCSC guidance emphasises knowing the architecture and identities, assessing health and behaviour, authorising through policy, authenticating and authorising across the environment, monitoring users/devices/services and not trusting any network by default.[4]

Those principles remain useful in operational technology, but their implementation must respect a different order of priorities. OT controls physical processes. A security change that interrupts protection, control or visibility can create consequences more serious than the cyber risk it was intended to reduce. The objective is therefore not to force every controller, sensor or safety system to behave like a modern enterprise endpoint. It is to remove unnecessary trust at the safest enforceable points around the process.

Design rule In critical infrastructure, “assume breach” must be paired with “preserve safe operation”. Security controls should reduce blast radius, improve visibility and constrain access without turning identity, inspection or cloud connectivity into a new single point of operational failure.

Identity first does not mean identity everywhere

Many legacy devices cannot support modern authentication, endpoint agents, strong cryptography or rapid certificate rotation. Some protocols were designed for deterministic behaviour on isolated networks. Retrofitting controls directly onto them can be unsupported, untestable or unsafe.

The practical pattern is to strengthen identity at the points where people, services and zones meet: privileged access workstations, jump hosts, remote-access gateways, engineering workstations, management servers, API brokers and controlled conduits. These become policy enforcement points. Behind them, older assets can remain stable while access to them becomes attributable, time-bound, least-privileged and observable.

Sequence matters

Phase	Purpose	Actions
1. Discover	Know what must not be interrupted	Map assets, flows, owners, safety dependencies, vendors, remote access and recovery constraints; observe passively where possible
2. Stabilise access	Remove unmanaged entry paths	Inventory accounts; close direct vendor links; introduce controlled remote access, MFA where supportable, approvals and session recording
3. Segment	Limit lateral movement and fault propagation	Define zones and conduits by process consequence; enforce only required flows; preserve safety and deterministic communications
4. Add identity context	Make access decisions attributable	Use named identities, PAM, device certificates or workload identity at gateways and management tiers
5. Monitor safely	Detect misuse without destabilising control	Centralise gateway, identity and network telemetry; baseline behaviour; tune alerts with operators
6. Exercise and improve	Prove isolation and recovery	Test isolation plans, break-glass access, manual operation, backups and staged restoration during agreed windows

Segment by consequence, not by organisational chart

Good segmentation groups assets and communications according to process function, safety consequence, exposure and required interaction. The resulting zones may not align with sites, departments or IP subnets. Conduits between zones should carry only known flows, through controlled and monitored paths.

Start with the highest-consequence boundaries: enterprise-to-OT access, internet or supplier connectivity, inter-site links, safety-related systems, engineering access and management tiers. NCSC secure-connectivity guidance for OT emphasises balancing risk and opportunity, limiting exposure, centralising and standardising connections, hardening boundaries, limiting the impact of compromise, logging connectivity and establishing an isolation plan.[5]

Microsegmentation can add value, but it is not the opening move in an estate with incomplete asset knowledge. Enforcing a policy against an unknown dependency can stop a process. Build visibility, validate flows with engineers and introduce controls in monitor-only or staged modes before blocking.

Design for degraded identity and disconnected operation

A central identity provider can strengthen control and simultaneously become an operational dependency. The architecture must state what happens when identity, DNS, time, certificate validation, logging or the WAN is unavailable. Critical functions may need locally cached decisions, bounded offline credentials, resilient local authorities or carefully controlled break-glass procedures.

These mechanisms should not become permanent bypasses. They need constrained scope, dual authorisation where proportionate, strong recording, periodic exercise and automatic review after use. The question is not whether exceptions exist; it is whether they are deliberate, bounded and recoverable.

What to leave alone

Asset or condition	Why direct change may be wrong	Safer compensating pattern
Unsupported PLC, RTU or protection device	Firmware or agent change may invalidate support or safety evidence	Isolate the zone; restrict engineering paths; monitor passively; retain tested spares
Deterministic or latency-sensitive control flow	Inline inspection may introduce jitter, delay or failure modes	Enforce at an upstream conduit; allowlist known flows; use out-of-band monitoring
Safety instrumented system	Cyber change can affect certified safety function and independence	Preserve separation; tightly govern maintenance access; validate with safety authority
Vendor-dependent legacy system	Unapproved hardening can prevent support or recovery	Contractual controls; mediated remote access; session recording; tested recovery image
Stable isolated enclave with no business requirement to connect	Connectivity may add more risk than value	Keep it isolated; control media and maintenance; document and periodically verify the boundary

Controls that usually deliver early value

- One controlled route for remote and third-party access, with named identities, approval, time limits and recording.
- Dedicated privileged access workstations or hardened jump services for engineering and administration.
- Removal of shared or default accounts where technically safe; vaulting and rotation where removal is not possible.
- Network zones and allowlisted conduits around high-consequence processes and management tiers.
- Passive asset discovery and network monitoring, validated with site engineers.
- Protected offline configuration backups, golden images, spares and rehearsed restoration.
- An isolation plan that states who can disconnect what, under which conditions, and how safe operation continues.

Measures that show whether it is working

Measure	What it reveals
Percentage of OT access paths under managed identity and recording	Whether unmanaged entry routes are actually disappearing
Number and age of shared, default and dormant privileged accounts	Identity debt and accountability gaps
Approved versus observed flows across critical conduits	Policy drift, hidden dependencies and anomalous connectivity
Time to isolate a zone while maintaining safe operation	Whether containment is operationally credible
Successful restoration of representative controller/configuration backups	Whether recovery evidence exists beyond backup completion
Third-party sessions attributable to an individual and work order	Whether supplier access is governed in practice

The implementation test If the new security control fails, does the process fail safely, continue in a bounded mode, or stop unexpectedly? If the answer is unknown, the control is not ready for enforcement.

One governance model, three disciplines

Why sovereignty, Secure by Design and zero trust are better governed as one decision than three separate workstreams

Sovereignty, Secure by Design and zero trust are often managed as separate workstreams. In high-consequence programmes they should meet in one control model.

Decision	Sovereignty lens	Secure by Design lens	Zero-trust / OT lens
Identity service	Who controls and operates it? Which jurisdiction applies?	Which risks and requirements does it address? How is it tested?	What happens to safe operations when it is unreachable?
Privileged support	Where are staff, authority and tooling located?	Is the path represented in architecture, risk and evidence?	Is access attributable, time-bound, mediated and recoverable?
Encryption and keys	Who can use, recover, revoke or compel the keys?	Are key controls traced to data risks and verified?	Will loss of key services interrupt essential control?
Supplier dependency	Can service continue if the supplier or nation withdraws?	Is the dependency owned, tested and accepted?	Can the affected zone isolate and operate safely?
Exit or recovery	Can capability move or continue under local control?	Has the plan been exercised with retained evidence?	Are restoration order and operational safety understood?

Questions for the next programme review

1. Which external authority or dependency has the greatest power to interrupt the mission?
2. Can we follow a material cyber risk from threat to tested control and accountable decision?
3. Which operational process would be endangered by a failed security or identity control?
4. Which privileged access path is least visible or least attributable?
5. What would we need to operate, isolate, recover or exit if normal external support disappeared?
6. When did we last test those claims, and what did the evidence show?

Final thought Mature security is not the absence of dependency, exception or residual risk. It is the ability to see them clearly, decide them at the right level, constrain their effect and prove that the service can continue when assumptions fail.

About TIS Global

TIS Global helps organisations design, assure and deliver secure digital transformation in complex and regulated environments. We bring together strategy, enterprise architecture, cyber security, cloud and operational delivery to turn policy into services that work in the real world.

References and further reading

[1] Ministry of Defence, "Secure by Design - a new way to manage cyber risk in capabilities" (30 May 2023).

[2] Ministry of Defence, "Secure by Design: User-centred guidance on Digital MOD.UK" (10 February 2025).

[3] Ministry of Defence, "Secure by Design: Two Years of Building a More Resilient MOD" (13 August 2025).

[4] National Cyber Security Centre, "Zero Trust" collection, reviewed 16 January 2026.

[5] National Cyber Security Centre, "Secure connectivity principles for operational technology (OT)".

[6] National Cyber Security Centre, "Cyber Assessment Framework".

© TIS Global 2026 | These papers provide general guidance and does not replace legal, safety or system-specific assurance advice.