Raising the bar for Cross Domain data transfer in KSA

Fawaz Al Malki
Chief Information Security Officer (CISO)

Sarah Al Rabiah
Graduate Trainee, CISO Office

With appreciation to colleagues across BAE Systems worldwide.

Publish Date: 13/04/2026, v1.0
ISE Website: https://ise.sa

Document Disclaimer: *This paper is intended solely for informational and discussion purposes and does not constitute a final or binding commitment. All insights and proposed approaches remain subject to further discussions, validation, review, and refinement.*

Executive summary

Secure data transfer is no longer a specialist or exceptional requirement. Across defence, government, and critical national infrastructure, it has become a routine operational necessity. Organisations increasingly depend on the ability to move information between systems operating at different security levels, while maintaining assurance, compliance, and trust.

In this paper, a Cross Domain Solution (CDS) refers to a set of controls appropriate for high threats environments that enables the controlled transfer of information between systems operating under different security policies or classification domains, while robustly preventing unauthorised disclosure, modification, or bypass of policy enforcement.

In the Kingdom of Saudi Arabia, this challenge is becoming more pronounced. Digital transformation is accelerating, sensitive data volumes are growing, and collaboration across organisations and partners is expanding. Meanwhile the cyber threat to organisations is continuously increasing, and critical national Information Technology systems require greater levels of protection from attack than ever before. As a result, information assurance has moved beyond a technical concern and is now a strategic responsibility shared by leadership, regulators, and operators alike.

Yet the mechanisms used to transfer data between security domains have not always evolved at the same pace as the environments they support. A wide range of approaches are commonly grouped under the broad term "Cross Domain", despite offering very different levels of assurance. In lower-risk settings, this may be sufficient. In high-assurance environments, where compromise is unacceptable, these distinctions matter.

Global best practice increasingly recognises that not all Cross Domain Solutions are equal, and that assurance is determined not by functionality alone, but by how trust is enforced. In the most sensitive environments, this has led to a preference for hardware-based approaches that reduce reliance on software and enforce security at the most fundamental level.

At present, such capabilities are not yet widely available locally in the Kingdom. That reality makes it more important, not less, to establish clarity now. By aligning early on what good looks like, stakeholders can help ensure that future standards are shaped deliberately and in line with national priorities.

This paper sets out the context, distinctions, and principles that underpin secure Cross Domain data transfer in high-assurance environments. Its purpose is to support informed dialogue and shared understanding, ahead of wider adoption.

From exception to strategic reality

Assured Cross Domain data transfer was once treated as an edge case for the most technologically advanced nationally critical platforms. Critical systems have typically either designed to operate largely in isolation, where any movement of information between security domains was limited or has very tightly controlled interfaces into enterprise IT.

That operating model no longer reflects reality. Organisations across defence, government, and critical national infrastructure now generate and rely on significantly greater volumes of sensitive data whilst also operating in a more sophisticated and hostile cyber threat landscape. Mission effectiveness increasingly depends on the timely sharing of information between systems operating at different classification levels, security postures, or trust boundaries.

At the same time, collaboration has expanded. Joint operations, shared platforms, external partners, and integrated supply chains all require controlled data exchange as well as AI platforms now consume data faster than humans and can manually transfer files in order to offer the autonomy that organisations aspire to engineer. Users in classified domains have an absolute need to work together in real time on enterprise applications. In these environments, the inability to move information securely can constrain decision-making, delay operations, and reduce overall effectiveness.

As a result, the risk profile has shifted. Risk no longer sits solely within individual systems. It increasingly exists between systems, at the points where information is transferred, transformed, or re-used.

This shift has elevated secure data transfer from a technical concern to a strategic one. Assurance is no longer just about protecting individual platforms. It is about ensuring that the movement of data between them can be trusted, governed, and defended at the appropriate level without disrupting data-centric decision making and interrupting AI powered operational efficiencies.

How Cross Domain transfer is handled today

In many environments, Cross Domain data transfer has evolved pragmatically rather than deliberately. Existing systems are adapted to support information exchange, often through additional controls layered onto general-purpose platforms.

Common approaches include software-based inspection and filtering, manual or semi-automated processes, and bespoke workflows designed to balance operational need with perceived risk. Speed, flexibility, and convenience are often key drivers in these decisions.

Over time, a wide range of solutions have come to be described collectively as "Cross Domain", despite significant differences in architecture, assurance, and suitability. This broad labelling can obscure important distinctions, particularly in regulated or high-assurance settings.

Security controlling functions that exist in CDS systems focus on two different attack vectors – that only permitted traffic flows can traverse the network boundaries, and that only permitted data can traverse the network boundaries.

Traditionally, two classical implementation patterns for network segmentation are frequently encountered. They are often components within a broader network architecture, rather than CDS designs in their own right:

- Layer 7 firewalls are fully software based commercial-off-the-shelf (COTS) platforms control both traffic flows and data flows bidirectionally, applying policy and inspection logic to the traffic flows moving between domains. They are flexible and familiar but rely heavily on software enforcement running on general-purpose systems. From an assurance perspective, firewall-based approaches introduce inherent complexity. Policy enforcement, inspection logic, operating systems, and underlying software dependencies all form part of the trusted computing base. As security requirements increase, maintaining confidence across this stack becomes progressively more challenging – both because accidental configuration errors can expose systems to threats, and also because complex software stacks are prone to zero-day vulnerabilities through which adversaries can bypass the security controlling functions. For this reason, Layer 7 firewalls are not considered CDS as there is no mechanism for guaranteed enforcement.

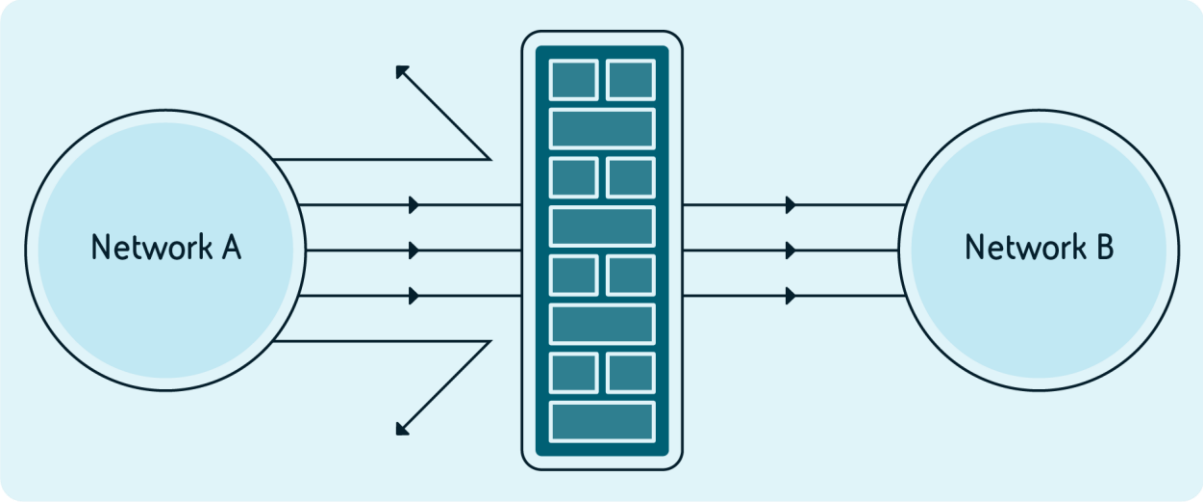

- Diode-based approaches enforce one-way traffic flow at the physical level, preventing return paths by design. They reduce certain classes of risk by guaranteeing that insecure networks can be connected to secure networks with absolute guarantees that data cannot flow back out of the secure network. While diode-based approaches are designed to mitigate specific categories of risk by enforcing physical one-way flow, they do not, on their own, address content-level data assurance. Malicious content can traverse a diode into the secure environment, as the only security enforcing function is traffic flow. Furthermore, given the single direction of traffic, it is impossible to have transactional interactions where the success of the data transfer is reported back to the originator, which is not ideal for environments where such failures are going to cause economic or safety related impacts.

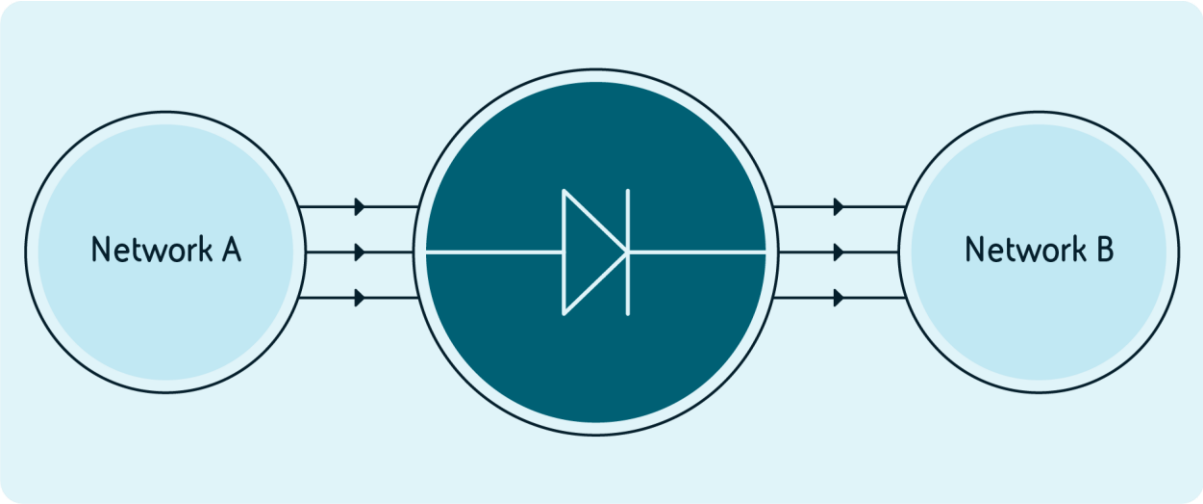

Both approaches have valid applications, and both are used globally. However, neither should be assumed to be inherently suitable for nationally critical environments. Their effectiveness depends on how they are implemented, assured, and governed within an appropriate security context.

In some environments, friction introduced by security controls can also influence human behaviour. Where approved mechanisms are perceived as slow or restrictive, users may seek informal workarounds to meet operational needs. These can include manual processes or physical media transfers that bypass intended controls entirely, increasing risk rather than reducing it. This human factor reinforces the importance of solutions that balance assurance with usability, as informal workarounds can undermine even well-designed technical controls if they sit outside the assured boundary.

Assurance, regulation, and why implementation matters

In high-assurance environments, the acceptability of a network segmentation solution is not determined solely by what it does, but by how trust is established and maintained. Cross Domain Solutions represent the approach to trusted interconnectivity in high threat environments.

Assurance is defined externally, through regulatory frameworks, accreditation processes, and risk acceptance by designated authorities. These mechanisms exist to ensure that systems behave predictably, enforce policy reliably, and do not introduce unacceptable risk, particularly when operating across security boundaries.

In many jurisdictions, this distinction is already well established. In the United Kingdom, the United States, and Australia, High-Assurance Cross Domain Solutions are explicitly recognised within regulatory guidance as a necessary control for connecting systems operating at different security levels. In certain circumstances their use may also extend to interconnections between systems of the same security levels where strict control, inspection, and assurance of data flows are required. While the specific frameworks vary, the underlying principle is consistent: secure information exchange in high-assurance environments requires deliberate design, independent assurance, and clear differentiation between approaches.

Within this context, implementation choices matter. Solutions that rely heavily on software running on general-purpose platforms inherit the complexity, dependencies, and attack surface of those environments. As assurance requirements increase, these characteristics can become limiting factors.

Technical assurance considerations in high-assurance Cross Domain environments

From a technical assurance perspective, regulators place emphasis not just on functional capability, but on the properties of the system enforcing security controls.

One key consideration is determinism. In high-assurance environments, regulators and accreditors expect policy enforcement and content-handling logic (including parsing, validation, and reconstruction) to behave in predictable and repeatable ways under all operating conditions.

Another critical factor is the trusted computing base. As reliance on software increases, so too does the volume of code, libraries, and operating system components that must be trusted to behave correctly at all times. Reducing the size and complexity of this trusted base is a common regulatory objective, as it directly affects the feasibility of independent assessment and assurance.

Regulators also consider attack surface and enforcement boundaries. Controls implemented at higher layers of the system stack are inherently exposed to a broader range of potential failure modes and exploitation techniques. Enforcing policy closer to the physical boundary between security domains limits the opportunities for bypass, misconfiguration, or unintended interaction.

For these reasons, higher levels of assurance are typically associated with architectures that minimise reliance on software, restrict functionality to what is strictly necessary, and enforce security controls as close to the physical layer as possible. These characteristics do not eliminate risk, but they make risk more visible, more measurable, and more defensible within an accreditation process.

To support clearer discussion, High-Assurance Cross Domain Solutions are often grouped into three broad categories:

- **Software-based solutions** rely primarily on highly assured software controls for inspection, filtering, and enforcement. They offer confidence and flexibility but depend on the correct operation of complex software stacks and are therefore theoretically vulnerable to sophisticated attacks as they are reliant on software with no guaranteed enforcement through hardware.
- **Hybrid solutions** combine software controls with some hardware-enforced elements (often in a complex bidirectional architecture). They can reduce certain risks through the use of diode-like directional control and through the use of robust use of software-based content verification (usually performing identical verification on both sides of the diodes) but still retain significant reliance on theoretically vulnerable software for decision-making and enforcement. While these offer a good balance of security and efficiency, their hypothetical vulnerability remains unacceptable for organisations facing the greatest threat.
- **Hardware-enforced solutions** place non-bypassable security enforcement of both flows and data within FPGA hardware. They still typically include software for parsing, transformation, and policy verification, but that software is deliberately constrained, minimised, and engineered to be independently assessable; furthermore, the code is in housed in FPGAs that are inaccessible to attackers. However, the verification in FPGA of all content is exceptionally complex in terms of hardware design and software development costs.

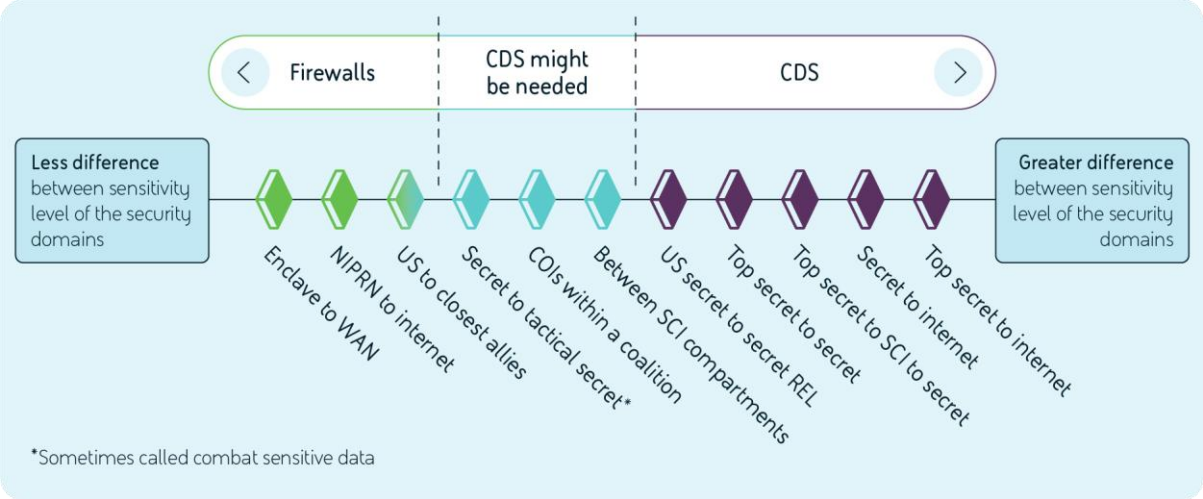

Ref: U.S. Department of War, Cross Domain Solutions Overview, Defense Acquisition University Webinar, 9 February 2023

This is not a judgement of quality or intent. Each category has appropriate use cases. However, treating them as interchangeable obscures the relationship between assurance requirements and implementation realities.

How High-Assurance Cross Domain Solutions enforce trust

Despite their implementational differences, High-Assurance Cross Domain Solutions typically follow a common operational pattern designed to support data assurance across security boundaries, in a form that Layer 7 firewalls do not.

Data is first **transformed** into a form that can be safely inspected. This step is not simply about format conversion but about reducing ambiguity and complexity before any policy decisions are made. By normalising data into a simple form and removing complex structures that cannot be reliably analysed, the transform stage constrains the scope of what must be trusted and ensures that subsequent verification is performed against a well-defined and fully understood representation of the information. For a Microsoft Word document, for example, this may extract the text, images and styles only, therefore excluding malicious content such as code and technical vulnerabilities. At this point the normalized content from the Word document is also assured to not contain classified information, prohibited character sets or sensitive data.

The transformed data is then **verified** against defined security policies. In high-assurance environments, this verification typically involves strict validation of structure, syntax, and allowable content, rather than reliance on pattern matching or heuristic analysis. The objective is to ensure that only data which can be fully understood and explicitly permitted by policy is allowed to cross the domain boundary. This verification determines whether the data complies with permitted formats, content rules, and contextual constraints. This stage verifies that the data being transferred is indeed a valid Word document format of data, and not other potentially malicious data masquerading as a document to evade detection.

Only once verification is complete is the data **reconstructed** in the destination domain. Reconstruction is deliberately constrained to produce output that is appropriate to the new security context, without reintroducing hidden behaviour, implicit functionality, or unverified structures. This ensures that the act of rebuilding the data does not undermine the assurance gained through earlier stages. This recreates the Word document in the native docx format, built from only the verified content.

Whilst many of the common data formats can be exceptionally challenging to verify. - proprietary enterprise data formats, web browsing, software updates etc, the approach to verification and normalization of data can be embedded in a nationally relevant standard. The standard can then be applied conceptually in an identical way regardless of whether the underlying hardware is software or hardware based. A local market of software developers can easily create a marketplace of these components, encouraging innovation in how some of the more complex data formats can be verified.

The Kingdom context and the importance of timing

The Kingdom is undergoing rapid digital transformation across defence, government, and critical national infrastructure. As systems become more connected and data-driven, the need for secure Cross Domain transfer will continue to grow.

At present, High-Assurance Cross Domain capabilities that are independently evaluated and routinely accredited for the most sensitive use cases do not appear to be widely established or locally available within the Kingdom.

This creates a window of opportunity. Decisions made now will influence how standards are set, how risk is managed, and how future capabilities are adopted. Without early alignment, there is a risk that approaches become entrenched through retrofit rather than intent.

Establishing shared understanding at this stage helps ensure that future implementations are guided by principle, proportion, and national priorities, rather than convenience alone.

Looking ahead

Assurance requirements for Cross Domain data transfer are expected to continue rising, driven by increasing data sensitivity, regulatory maturity, and operational dependence on secure information exchange.

As assurance expectations increase, greater emphasis is likely to be placed on demonstrable properties such as deterministic behaviour, reduced trusted computing bases, and enforceable separation between security domains. Understanding the distinctions between different approaches early enables more informed planning and reduces the likelihood of costly redesign or compromise later.

Looking ahead also involves investing in knowledge, not only implementation. ISE is interested in exploring how they can collaborate with universities and academic institutions to encourage research into the practical engineering challenges of High-Assurance Cross Domain Solutions. Supporting research and education helps build local capability, deepen understanding, and lay the foundation for potential future adoption in the Kingdom.

By promoting informed dialogue today, stakeholders can help shape a future where secure data transfer supports national objectives with clarity, confidence, and trust.