9/29/2025

Decoupling Data and Connection based Risks in CDS

This page is intentionally blank

Change History

Version	Name	Description	Date
1.0	Ben Rosenfeld	Final Draft	9/29/2025

TABLE OF CONTENTS

1.0 INTRODUCTION

Cross Domain Solutions (CDSs) are designed with a singular, mission-critical objective: to enable the secure and efficient transfer of data across security boundaries in support of operational goals. Without a mission, the rationale for sharing information, and by extension, for implementing a CDS, ceases to exist. In this context, CDSs must be engineered not only for technical robustness (to align with acceptable risk tolerances) but also to ensure that security policies do not unnecessarily hinder mission execution. A CDS should enhance, not impede, the effectiveness of mission operations. As mission architectures evolve and technological advancements reshape operational environments, a comprehensive and adaptive risk strategy becomes even more critical.

Importantly, the risks associated with using a CDS extend beyond cybersecurity or technical concerns. Inadequate data transfer capabilities can directly impact mission success, potentially resulting in loss of life or providing adversaries with a strategic advantage. Therefore, cross-domain risk must be understood as the possibility of any negative outcome, from technical failure to operational disruption. Identifying, quantifying, and prioritizing these risks based on both likelihood and impact allows decision-makers to make informed choices about which risks to mitigate and which may be acceptable considering mission objectives.

This paper explores the multifaceted risks inherent in cross domain solutions and emphasizes that disparate risks do not always need to be tightly coupled within system architectures to achieve effective mitigation. Unnecessarily linking distinct risk areas can lead to avoidable lifecycle challenges and architectural complexity.

2.0 UNDERSTANDING THE RISKS OF CDS

Connecting multiple networks that were initially designed to operate independently introduces inherent risks. While cybersecurity and technical vulnerabilities are often the primary concerns, the full spectrum of risks includes significant operational threats.

From an operational perspective, a CDS architecture can pose a serious risk if it creates a bottleneck for mission-critical data. Delays or failures in data transmission can lead to mission degradation, loss of equipment, financial costs, and, in extreme cases, loss of life, all of which can provide an adversary with a strategic advantage.

On the cybersecurity side, risks can generally be categorized into two main areas, data-based risk and connection-based risk. Understanding the distinction between data-based risk and connection-based risk is crucial for effective cybersecurity strategies.

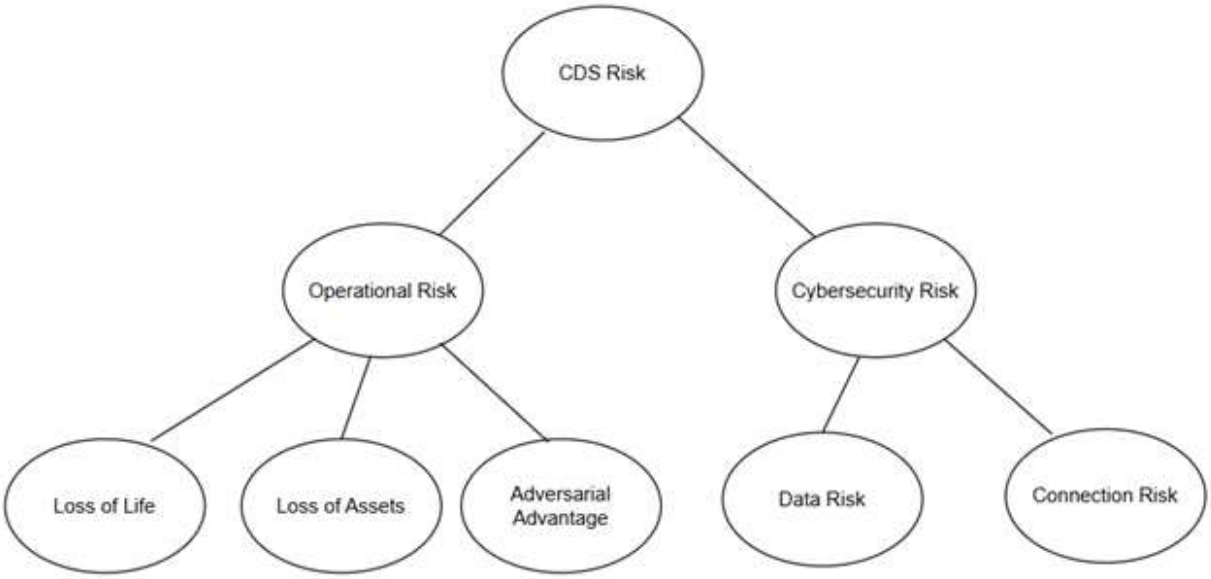

Figure 1: CDS Risk Breakdown

2.1 Data Risks

Data risks refer to risks resulting from threats against the integrity or confidentiality of data and the downstream systems that consume that data. This threat can be present at rest, in use, or in transit and can be generally categorized as being associated with data hiding, data disclosure, and data attack.

- **Data Hiding** refers to the deliberate concealment of information by malicious actors who attempt to pass along sensitive or unauthorized data without detection. Techniques like data encoding and steganography are often used to mask the true content or purpose of the data being transmitted.
- **Data Disclosure** occurs when confidential information is exposed to unauthorized individuals, either accidentally or intentionally. For example, a malicious insider might transfer classified data from a secure internal network to the public internet using a poorly understood or overly complex data type.
- **Data Attack** involves malicious actions taken by attackers to use data maliciously in order to gain unauthorized access to, manipulate, or destroy systems or other data. These attacks can take many forms, such as parser attacks, malware deployment, or exploiting system vulnerabilities.

There is a broad and critical dimension to data risk. Not only is it important to protect the data, it's equally important to protect higher classification networks or higher security domains from the data itself. This includes preventing malicious or inappropriate data from introducing vulnerabilities or serving as a vector for attack. As historical risk frameworks for cross-domain systems have emphasized, it is the high-side system, not the CDS, that is the exploitation goal. In this context, Data risk is the risk that data traversing a particular CDS might either introduce or exploit

vulnerabilities in the high-side environment or inadvertently cause data to be disclosed to the low-side. This perspective reinforces the need to treat data not only as a valuable asset, but also as a potential threat vector that must be rigorously managed.

2.2 Connection Risks

Connection risks involve the compromise of the channel through which data is transmitted or stored. It is the network-based risk created by physically connecting the associated networks. For example, an adversary might monitor network ports, intercept traffic, or exploit vulnerabilities in one network to gain unauthorized access to another, ultimately compromising the connection itself. This can manifest in several ways, such as Man-in-the-Middle (MitM) attacks, where an attacker intercepts or alters communications between systems without detection, and Denial-of-Service (DoS) attacks, which flood network channels to disrupt legitimate traffic or render the system inoperable.

One key aspect to emphasize here is that concerns related to connection threat are not limited to the CDS itself, rather, they are about protecting the connected networks from network-based exploits that originate from other networks they are linked to through the CDS. A compromise in one connected domain can serve as a launchpad for attacks on another, making the integrity of the connection a critical security concern.

3.0 DECOUPLING DATA AND CONNECTION RISK IN CDS ARCHITECTURES

With this distinction between connection risk and data risk, it would seem reasonable to consider how tightly coupled (or not) their mitigation functions should be within a cross-domain architecture. As with the presence of these risks, the nature of these risks is largely independent. For example, simply unplugging a multi-domain connection would eliminate the network threat, but it would not address the data threat for any data shared through some other mechanism.

Despite this, traditional CDS designs often bundle mitigation functions of both risk domains into a single, tightly integrated, and monolithic solution. With this approach, components are interdependent in lifecycle and architecture; any modification to one element inevitably triggers a cascade of consequences across the entire system. While such coupling may offer advantages in centralized control and uniform security guarantees, it also introduces significant limitations, particularly regarding system flexibility, scalability, and speed of evaluation; all of which arguably increase the level of operational risk to missions through decreased agility, throughput, loss of schedule, and increased cost.

Alternatively, connection risk and data risk could be considered as independent concerns, with each addressed independently without demanding the mitigation of both using the same mechanisms. To illustrate this, consider two simplistic yet contrasting thought exercises:

First, consider a scenario where a cross-domain system is used to transfer sensitive, low-bandwidth data between networks that are spread out over large geographical footprints. Because this data is so significant, it makes sense to have multiple sites capable of sending this data between the networks to ensure geographical redundancy. In this scenario, the physical network connections associated with the cross-domain capability must scale to multiple locations, representing increased

connection risk. However, as the required data types and bandwidth are not affected by this requirement, the data risk is relatively unchanged. From an operational perspective, there is no need to scale or re-evaluate the data-risk mitigation already in place.

Conversely, consider the use of human-based transmission of data (aka "sneakernet") where data is physically transported (e.g., via paper or optical media) between isolated networks. In this scenario, the traditional network-based connection risk is effectively eliminated, as there is no persistent or electronic link between the networks. However, this does not eliminate the data risk.

The data remains a potential threat even if the human courier is trustworthy and not a malicious insider (an assumption made for this example). Such data could have been inadvertently or intentionally modified before transfer to introduce integrity risk (i.e. the data possesses malicious malware, etc) or confidentiality concerns (i.e., containing information that should not be transferred.) Without mitigating these risks, the receiving system could be compromised through a data attack or subject to inappropriate release, even though no network connection exists.

4.0 BENEFITS OF DECOUPLING RISK

These examples raise questions related to the possibility of decoupling risk domains within CDS architectures. Doing so would enable efficient risk mitigation strategies that could be more readily tailored to the specific threat landscape of a mission or operational context.

Moreover, decoupling these risks presents substantial practical benefits for the CDS lifecycle, including the assessment process. CDS evaluation labs are overburdened and under-resourced, leading to significant delays in certifying CDSs for operational deployment. Because of the interdependent nature of monolithic architectures, even minor tweaks to a system intended to optimize performance or meet mission-specific requirements often involve re-examining broader portions of the solution. This is both time-consuming and resource-intensive.

By contrast, a decoupled architecture would allow for more focused assessments. If a change is made that only affects data risk, for example, then only that portion of the architecture and its associated mitigations need to be reconsidered. This modularity could significantly shorten assessment cycles, allowing for faster deployment of only those functions, increasing agility in response to evolving mission needs.

Separating data and connection risk concerns fosters architectural and technological innovation. Engineers who have effectively mitigated connection risk could focus on developing additional or novel solutions for data risk without being constrained by unrelated architectural dependencies and vice versa. Decoupling also supports operational adaptability. For instance, in time-sensitive missions where data becomes obsolete quickly yet requires a large amount of processing for filtration, mitigation functions could be scaled independently without unnecessarily scaling the entirety of the solution. There should be no need to scale physical network boundary devices when only increased compute for content threat is required.

5.0 RISK DECISIONS VS MITIGATION ARCHITECTURES

It should be noted in this discussion that there is a distinct difference between the risk decision that authorizes the use of a CDS (or data flow) and the implementation and evaluation of architectures that implement a CDS. In an authorization decision, connection and data risks will both be considered, contextualized by the networks, mission imperatives, and current threats. However, while these are evaluated as components of a multivariate decision, that consideration should not require the technical mitigations of those separate concerns to be unnaturally coupled. Stated differently, the risk model used to facilitate cross-domain risk decisions should be careful not to assume or demand unnecessary coupling of mitigations for disparate risks. To the contrary, higher levels of independence between these mitigation functions would provide finer controls for tailoring the authorization decisions for any given mission.

6.0 CONCLUSION

Based on the discussion, it seems reasonable that any unnecessary coupling of risk mitigations (data and connection) within current CDS architecture introduces significant lifecycle and architectural complications. By intertwining these distinct risk domains, CDSs become increasingly rigid, harder to maintain, and more burdensome to assess, leading to longer evaluation times and a slower path to operational deployment. The current approach stifles innovation by forcing developers to navigate a complex web of interdependencies when designing or updating security mechanisms. Even incremental improvements require broad reassessment considerations, discouraging experimentation and delaying integration of novel security solutions tailored to evolving mission requirements.

In contrast, working to decouple how we consider connection and data risk offers clear benefits. It enables increased modularity within CDS architectures, allowing for targeted risk mitigation, faster assessment cycles, and better alignment with mission-specific needs.

By rethinking how we approach these disparate risks, we would be freer to create more adaptable systems and better positioned to leverage emerging technologies. This would open the door for innovation by providing a flexible foundation upon which vendors can implement emerging technologies without being constrained by unrelated architectural components.