CYBERSECURITY™
MADE IN EUROPE

Arbit TRUST Gateway

Whitepaper

arbit

Table of contents

1. Concept 3
2. The Arbit TRUST Gateway 5
2.1. Arbit Trust Gateway Architecture 6
2.2. VOID Network 7
3. Arbit Trust Gateway Deployment 8
3.1. A unidirectional CDS deployed with one ATG 8
3.2. A unidirectional CDS deployed with one ATG and one ADD 9
3.3. A bidirectional CDS deployed with two ATG's 9
4. Arbit Trusted Information Release 10
5. Arbit Trust Gateway Protocol Support 10
6. Arbit Trust Gateway Hardware 11
6.1. Dell COTS servers 11
6.2. Ruggedized 12
7. Proxy Servers 13
8. Software 13
9. Filtering 14
10. Data Integrity 14

1. Concept

The Challenge of Trusted Information Exchange

Across defence, government, intelligence, and other highly regulated industries, organizations increasingly rely on isolated networks to protect their most valuable information. These environments are intentionally separated from less secure networks to reduce cyber-attacks, protect classified information, and maintain operational integrity. While this isolation in the form of air-gapping significantly improves security, it also creates a fundamental operational challenge: valuable information still needs to move between environments.

Information is rarely created for a single network. Intelligence reports, command and control data, operational data, software updates, sensor outputs, and business documentation to give a few examples, often originate in one security domain but are required in another. As organizations become increasingly data-driven, the demand for timely and secure information exchange continues to grow.

This creates a dilemma between two equally important objectives:

- Maintaining the highest possible level of security through network isolation.
- Enabling users to exchange trustworthy information to make timely operational decisions.

Neither objective can be compromised.

That is where Cross Domain Solutions enter the picture. They enable users to safely move approved information from isolated air-gapped networks that are highly classified to less secure networks.

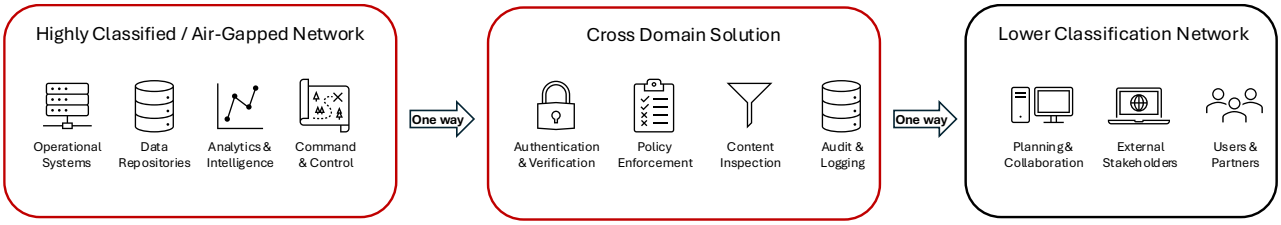

The Human Workflow

For the users operating within secure environments, transferring information is rarely their primary responsibility. Analysts, engineers, operators, and administrators simply need to move approved information from one domain to another so they can continue their work.

From their perspective, the desired workflow is straightforward:

1. Create and/or select the information that needs to be transferred.
2. Submit it for approval.
3. Have it verified against organizational security policy.
4. Deliver it safely into the destination environment.
5. Continue working without unnecessary delay.

The complexity behind this seemingly simple process is substantial. Every transfer represents a potential security risk and must be assessed to ensure that unauthorized information is not disclosed, malicious content is not introduced, and organizational policy is consistently enforced.

Consequently, organizations establish rigorous governance around cross domain information exchange. Every transfer may involve automated inspection, policy validation, human review, audit logging, and formal authorization before data is permitted to cross security boundaries.

Beyond Technology

Cross domain information exchange is therefore not simply a networking problem. It is an operational process involving multiple stakeholders with different responsibilities, that happen to rely on technology.

- Security teams define the policies that determine what information may be transferred.
- Approvers assess whether individual transfers satisfy organizational and regulatory requirements.
- System administrators maintain the infrastructure that performs inspection, filtering, auditing, and delivery.
- End users generate the business value by exchanging information that enables collaboration, decision-making, and mission execution.

Each participant has different objectives, but all of them depend on a transfer process that must be trusted and is secure, transparent, auditable, and efficient.

The Need for Better User Experience

Historically, cross domain transfer solutions have prioritized security above usability. While this focus is entirely justified, many existing solutions require users to navigate complex workflows, manually prepare transfer packages, or rely on specialist administrators to complete routine operations.

The result is often increased operational overhead, longer processing times, reduced productivity, and user frustration. In some cases, cumbersome processes encourage users to seek alternative methods of moving information, such as USB sticks and removable hard drives, which introduce additional security risks.

Modern Cross Domain Solutions must therefore achieve more than regulatory compliance and technical assurance. They must provide an experience that aligns with how users naturally work while maintaining uncompromising security controls.

The objective is not simply to move data securely, but to enable trusted information exchange that is intuitive for users, enforceable by security teams, and fully auditable for the organization.

This whitepaper introduces the Arbit TRUST Gateway designed around that principle: combining robust cross domain security with a user-centric workflow that simplifies information exchange without compromising security.

2. The Arbit TRUST Gateway

Part of Arbits Cross Domain Solution portfolio, which also includes the Arbit Data Diode, the Arbit TRUST Gateway (ATG) solves the aspects described in the above chapter. The ATG eliminates the need to rely on manual processes and insecure data transport methods which may circumvent security procedures due to lack of time or even malicious intent.

Compared with other methods of data release, the ATG provides the most secure way of releasing data and offers the most user-centric approach to cater for work-flow efficiency.

Data Release Method Comparison	USB	ATG
Avoids two-way connection when releasing data	No	Yes
Verifies source of released data	No	Yes
Verifies data content according to security policy	No	Yes
Ensures release work-flow policy is upheld 24x7	No	Yes
Low latency in release of data	No	Yes
Automatic threat & content scanning	No	Yes

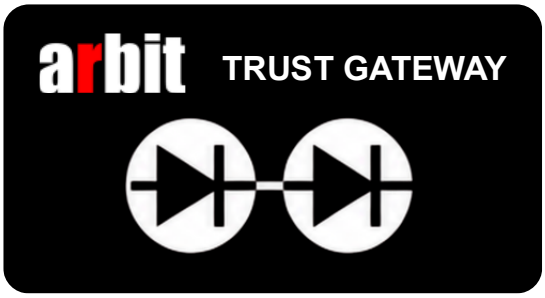

The fundamental building block of the Arbit Trust Gateway is the Arbit Data Diode, and it comprises of two Arbit Data Diodes connected in series. While the Arbit Data Diode is powerful in its own way, it is rather promiscuous in terms of what it allows to pass through.

In its basic mode of operation, the Arbit Data Diode simply monitors a certain folder on the sending side for data and ensures that it is transferred to the receiving side.

It can also perform necessary protocol conversions for that operation to happen and it can also pass the data through filters for added content checks, but other than that, the Arbit Data Diode does not authenticate the origin of the source and assumes that who ever provided the data is allowed to pass it along.

However, the Arbit TRUST Gateway is a more advanced version of the Arbit Data Diode, in the sense that it not only inherits the strengths and capabilities of the Arbit Data Diode which is CC EAL7+ certified, providing galvanic separation and enforcing unidirectional traffic on the hardware level to ensure data is only transmitted in one direction. It also provides data release capability through source identification and data verification on top.

By validating the identity of the source, we can make sure that only those that are authorized to send data are allowed to do so. That is done through secure key management and digital signing technology and only identities with a correct key are allowed to send data through.

Additionally, verifying the data content according to security policies before it is sent ensures that the data being released is allowed to leave the highly certified environment as well.

Furthermore, the ATG configuration is safeguarded against unauthorized changes when compared to the standard Arbit Data Diode configuration, as the ATG is configured through a separate tool running on a separate computer. The configuration is then stored on Smartcards, that are read by the ATG for added security hardening, which cannot be modified on the fly.

2.1. Arbit Trust Gateway Architecture

The Arbit Trust Gateway architecture, with its two Arbit Data Diodes in series, comprises of four servers, each connecting to an input and output network interface. Each of the four servers plays a role in transporting data through the ATG and follow the pitcher – catcher relationship of the Arbit Data Diode, creating galvanic separation and unidirectional data flows.

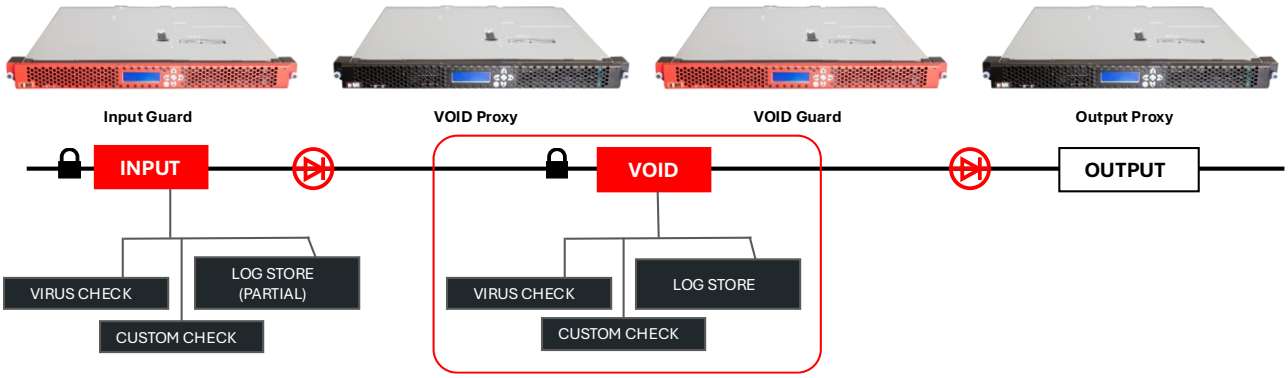

The four servers as depicted in the above diagram are connected as follows:

Arbit Trust Gateway Server	Connected Input	Connected Output
Input Guard	Sending Network	VOID Proxy
VOID Proxy	Input Guard	VOID Guard
VOID Guard	VOID Proxy	Output Proxy
Output Proxy	VOID Guard	Receiving Network

Furthermore, the Input Guard and VOID Guard can send data to external verification servers or filters as they are also called, for added content checking and virus scanning.

The ATG receives data from a source system using a custom protocol which enables the data and corresponding metadata to be signed by up to five different digital signatures. When the Input Guard receives the data, it can send it to an external filter for additional verification and content checking before the signatures are checked and the data gets sent to the VOID Proxy.

The connection between the VOID Proxy and the VOID Guard creates an isolated network with no interactive access from either the sending or the receiving side. This isolated network called the VOID network, which stands for Verification Of Identity and Data, performs the same verification done by the Input Guard a second time. If all signatures and filters accept the data, it passes through to the Output Proxy to be released to the receiving side, and a log entry is created. Otherwise, the data is only logged, and possible counter measures may be applied.

In other words, the ATG can verify data requested to pass through in two fundamental ways:

1. Verify that the data originates from an approved source system using digital signing technology.
2. Apply additional filters providing verdicts on whether the data can pass through or not.

The ATG supports several data channels called services. Each service can have different keys associated and use different filters as well, ensuring that various source systems and users can access different services. One service might be a system service, exporting very specific data which can be checked automatically. Another service may allow export of PDF-files, which are much more complex, and therefore require two-key authentication before the release with specially appointed qualified personnel.

Each approved source system that is allowed to transmit data through the ATG will sign the data to be released. This feature is used by certain applications, like the Arbit Trusted Information Release (ATIR), to enable three-key authorization. In the case of ATIR, a Drafting Officer, a Releasing Officer and the ATIR system itself sign the data, ensuring that it originates from the ATIR platform and that two different individuals have approved the release.

Each service may also have several filters attached. When data is passing through a certain service, all attached filters for that service are presented with the data. They then each inspect the data and cast a verdict: YES or NO. If any filter answers NO, then the data is rejected. If all filters answer YES, then the data is released.

2.2. VOID Network

As explained earlier, the Arbit Trust Gateway creates a VOID Network between the VOID Proxy and the VOID Guard, that is isolated from the Sending network and the Receiving network, and that all checks (signatures and filters) are performed twice. Once on the Input Guard and then repeated in the VOID Network. The first verification on the Input Guard also acknowledges to the sender that the ATG has received the transaction. The design rationale behind performing the verification checks of signatures and filters twice is the fact that the second verification done in the VOID Network is uncompromisable and that the configuration of it cannot be modified without proper authorization. This means that the configuration in the VOID Network is immutable and always holds the truth so to say.

Having the VOID network in isolation means:

1. It cannot be breached from the Receiving Network as the second diode in the ATG prevents any data from coming in. This means that the configuration in the VOID Network (signatures and filters) cannot be altered/read from the outside.
2. Any data received from the Sending Network will not get any acknowledgements back, due to the first diode, and therefore any attempt to read or alter the VOID Network configuration is also not possible.

This isolation provides a high degree of protection for the VOID Network and thus for the integrity of the software running in it. However, this also means that the software is unable to detect threats occurring on the Sending Network and the Receiving Network.

Threats on the Receiving Network are not a problem for the ATG itself as this network has no influence on the process of releasing data, which is the critical part of the ATG. However, threats on the Sending Network are of more serious concern. Therefore, the ATG performs the checks twice. Once on the Input Guard (on the Sending Network) and once on the VOID Guard (on the VOID network). This provides the VOID Guard with the assurance that all data passed to it may be released as it would otherwise have been blocked by the Input Guard.

If the VOID Guard encounters any non-releasable data, this means that either the Input Guard is misconfigured, or it has been compromised. Both cases are equally critical for a gateway system, and the VOID Guard can act on this information even if it has very limited information about the Sending Network.

The VOID Guard can be configured to act in the following modes in case it encounters non-releasable data:

Mode	Action
LOOSE	Incident is logged
PATIENT	Incident is logged and if 3 incidents occur the service is shut down
STRICT	Incident is logged and the service is shut down
POWEROFF	Incident is logged and the VOID Guard is powered off

This makes the Arbit TRUST Gateway ideal for exporting data from highly classified networks to networks with similar or lower classification levels, as we can control all aspects of the data to be released.

3. Arbit Trust Gateway Deployment

The Arbit Trust Gateway is a flexible and modular solution allowing for several deployment options and configurations depending on the use case. There are 3 main deployment types in which the Arbit Trust Gateway can be used to design and build Cross Domain Solutions with.

1. A unidirectional CDS deployed with one ATG using one unidirectional channel.
2. A bidirectional CDS deployed with one ATG and one ADD using two separate unidirectional channels.
3. A bidirectional CDS deployed with two ATG's using two separate unidirectional channels.

Together with Arbit applications and other custom-made applications, the 3 deployment types offer secure and trusted data transfer to enable information exchange across multiple domains and classifications, where the ATG in combination with the Arbit Data Diode provide the fundamental cyber security provisions of Cross Domain Information exchange.

Depending on the use cases to be deployed, it is important to have planned and designed what type of data is allowed to flow through the CDS, which people are authorized to send data through, and what external filters and content checks are to be performed. The Arbit Trust Gateway configuration is stored on Smartcards, which is programmed externally using the Arbit Configurator tool. The Configurator tool writes the configuration on Smartcards, one for each server of the ATG. This added security measure ensures that the configuration of the ATG is locked on the Smartcards and that it cannot be changed or modified without proper access to the Configurator tool.

3.1. A unidirectional CDS deployed with one ATG

This basic solution of the ATG caters for exporting data from High to Low networks using one outgoing channel and is deployed using the 4 servers of the ATG. It comprises of the Input Guard, VOID Proxy, VOID Guard and Output Proxy, with the Arbit CC EAL7+ 10GbE SFP+ modules installed on the servers.

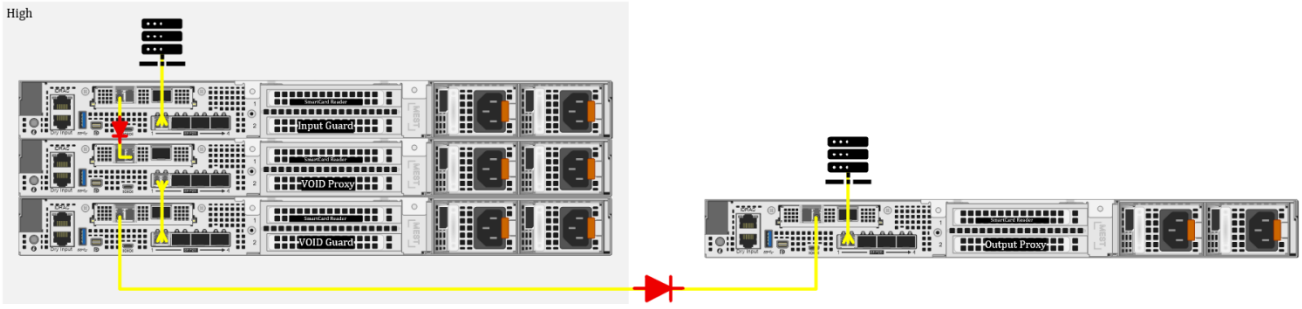

This deployment is used for example in the Arbit Trusted Information Release (ATIR) application.
Filter and logging servers are not depicted in the above diagram, and if included would add another 2 servers.

3.2. A unidirectional CDS deployed with one ATG and one ADD

The second solution caters for exporting data from High to Low and importing data from Low to High using an ATG for the outgoing channel and an ADD for the incoming channel.

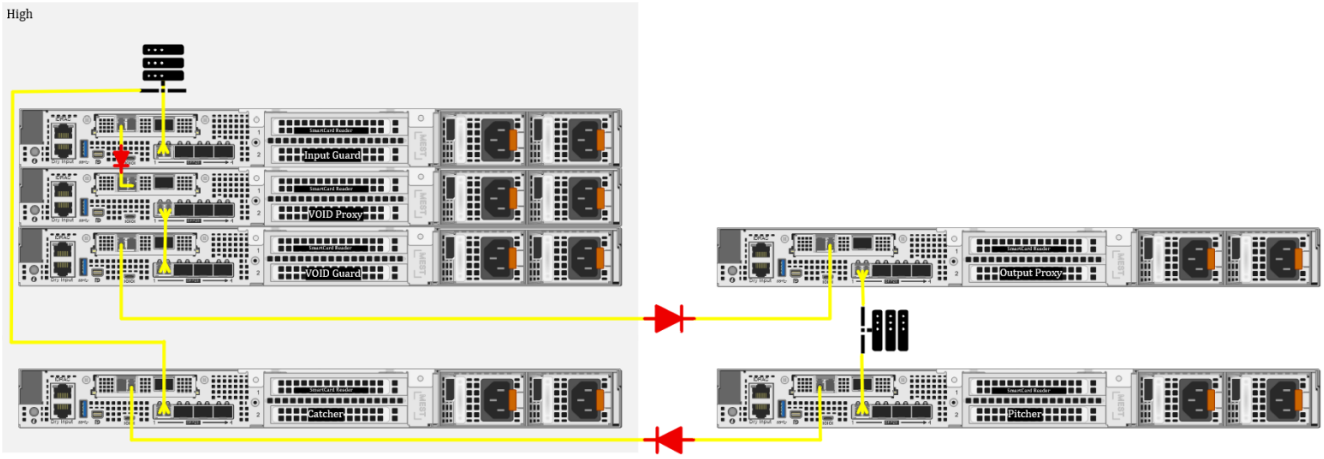

This deployment is typically used for example in the Arbit Web Gateway (AWG) or the Arbit Desktop Gateway (ADG) applications.

3.3. A bidirectional CDS deployed with two ATG's

The third solution caters for exporting data from High to Low and importing data from Low to High using an ATG for the outgoing channel and an ATG for the incoming channel.

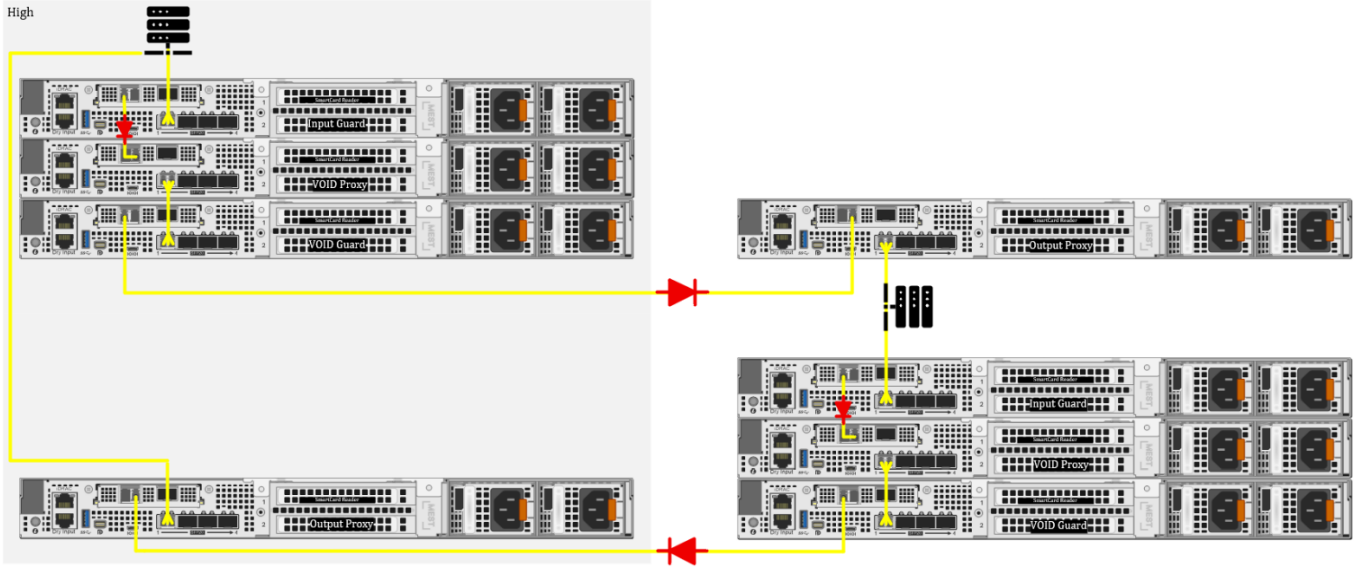

This deployment is used for example in the Arbit Sitaware Security Gateway (SSG) and can also be used for the other Arbit applications as well.

4. Arbit Trusted Information Release

The Arbit Trusted Information Release (ATIR) application works perfectly well with the ATG to cater for the human workflow we mentioned in chapter 1. In situations where you have highly classified networks withholding secrets that you must guard, and at the same time have a need to exchange releasable information with outside networks, you must ensure that no information is sent out unintentionally or maliciously. With ATIR a two-factor release process is used to verify and control that the information you release is indeed, what is allowed to pass through, abiding by the Trust but Verify principle. Additionally, the releasable data is passed through filters for inspection and added security control.

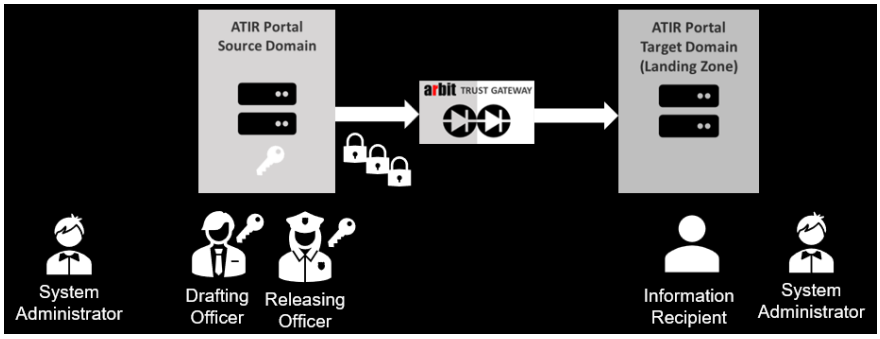

This is just one example of how the Arbit Trust Gateway can be used as a unidirectional CDS.

5. Arbit Trust Gateway Protocol Support

The Arbit Trust Gateway supports a wide range of protocols, which the Arbit Data Diode also supports. However, the ATG with its primary goal of transferring data securely and in a controlled and trusted manner, relies on an internal Arbit CDS Protocol, that is transported over UDP for its internal communication. This leaves the job of file handling, storage, protocol conversion of the originating files and protocols to be delivered to the ATG for proxy servers running either Arbit applications or other 3rd party applications. These applications perform file handling and protocol conversion before relaying the data through the ATG. Each of the 4 ATG servers handles the packets passing through, performing their different tasks, such as validating and authenticating signatures and content checking, and when data leaves the ATG, the receiving proxy server converts the data back to its original form and delivers it to the receiving system or target domain.

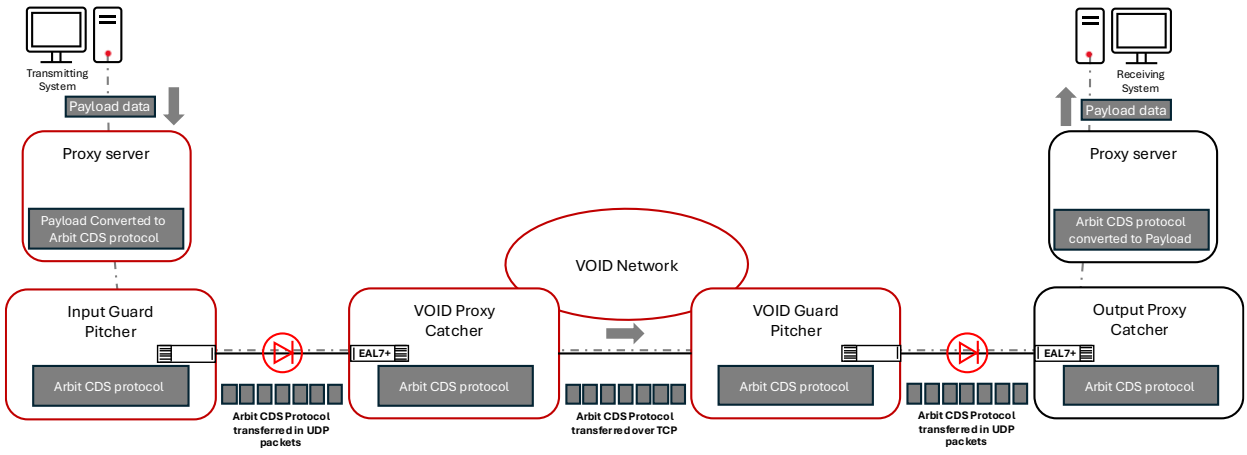

This means that applications can be built to cater for almost any type of protocol or file type. The Arbit SDK provides an open API for enabling customers and application developers to create their own services and use cases, making the workflow of Cross Domain Information Exchange more user friendly and simple.

6. Arbit Trust Gateway Hardware

The Arbit Trust Gateway is available in two form factors:

- Dell COTS servers (10GbE): This deployment option is designed for server room usage.
- Ruggedized (1GbE): This deployment option is designed for battlefield usage and is limited to 1GbE.

6.1. Dell COTS servers

This hardware configuration designed for data center deployments relies on four Dell COTS servers, to build the ATG consisting of its 4 servers (Input Guard, VOID Proxy, VOID Guard, Output Proxy) equipped with 2 Arbit SimLink SFP+ modules and with 2 EAL7+ certified Arbit Data Diode 10GbE unidirectional SFP+ modules.

The supported server configuration is the Dell PowerEdge XR5610 and is delivered with a 16GB read-only SSD for storing the Operating Software and all data processing happens in the 32 GB of RAM.

ATG servers

Configuration of the Units is performed using the Arbit Configurator and the configuration is stored on Smartcards for enhanced security hardening. This ensures added security against tampering and unwanted and unlawful configuration changes.

This configuration also supports Filters, which may perform a wide variety of custom-made actions and content checks, for additional verification. Refer to section 7. Filtering for more information.

Dell COTS Server Specifications

HW Specifications	Value
Configuration Type	Arbit Trust Gateway Server
Supported server	Dell PowerEdge XR5610
Memory and Storage	32 GB RAM / 16 GB Read Only SSD
Diode principle	Hardware based (optical)
Backflow	None (secured by the laws of physics)
Certification	Pending Common Criteria EAL4 for the VOID Guard
Fiber technology	Single mode 1310 nm
Diode NIC speed	10 Gigabit
Server Size	1U, 19” (4 units required for a complete Gateway: Input Guard, VOID Proxy, VOID Guard, Output Proxy)
Server Weight	Max. 11 kg
Power supply	2 x 100-240 VAC 50/60 Hz (Hot swap redundant)
Dimensions (W x H x D)	19” x 1U x 463 mm

6.2. Ruggedized

For battlefield and other harsh environments, the Ruggedized Arbit Trust Gateway is engineered and tested to international standards, it delivers unparalleled operational reliability, flexibility, and environmental resistance in mobile and stationary installations. Designed for size optimization and to support a range of military standards, to withstand the extreme environments it can be deployed in.

It comes pre-installed with the Arbit Data Diode for unidirectional communication. The main difference between the Dell COTS server version and the ruggedized version, besides the physical casing, is the throughput value which is 1GbE for the ruggedized version.

The ruggedized Arbit Trust Gateway enhances operational efficiency on the Battlefield allowing for seamless C2 Integration and has a smaller footprint, compared to the COTS server version.

A CDS with an ATG for exporting data and an ATG for importing data in the COTS form factor needs 8 x 19” 1U servers. In the ruggedized form factor, it needs only 3 x 19”/2 1U servers.

Configuration of the Ruggedized hardware is performed using the Arbit Configurator and the configuration is stored on Smartcards for enhanced security hardening. This ensures added security against tampering and unwanted and unlawful configuration changes.

Ruggedized Server Specifications

HW Specifications	Value
Processor	3× Intel Atom® x7835RE (each with 8 cores)
Power	Input 12–36 VDC
Operating Temperature	-40°C to +60°C
Storage Temperature	-40°C to +71°C
Humidity	95 ± 4% RH (MIL-STD-810H)
IP Protection	IP65 (Solid: 6X, Water: X5)
Shock Resistance	40 g / 11 ms (MIL-STD-810H)
Vibration	Wheeled (20), tracked (20), helicopter (14) (MIL-STD-810H)
EMI/EMC Compliance	MIL-STD-461F (RE102, CE102, CS114)
Dimensions	(WxH) 220 × 43.4 mm
Rack Depth Compatibility	400 mm including cable clearance
Weight	3,5 kg

7. Proxy Servers

Proxy Servers are used together with the Arbit Trust Gateway and depending on the use case can either be deployed as Filter servers, Logging Servers or Application servers. Proxy servers must fulfill the following minimum requirements in case of physical HW:

Feature	Minimum requirement (industrial)
CPU Cores	4 (not counting hyper-threading)
CPU architecture	x86 64 bit (AMD or Intel)
RAM	32 GB
Disk type	SSD Mixed Use
Disk configuration	JBOD
RAID controller	Linux OS compatible
Disk space	*Depending on the use case of the Filter, Log or Applicatio*.
Network interfaces	3 (one of which must be either: 1000BASE-T, 10GBASE-T or 10GBASE-LR)

Proxy Servers can also be deployed virtually, giving more freedom of choice. Arbit provides appropriate server configurations according to the specific use-case.

8. Software

The license includes the following software features:

Feature	Value
Multiple services (data channels)	Yes
Signature check of all data	Yes
Number of individual signatures on data	1 – 5
Key verification	Yes
Complete transaction log	Yes (add-on)
Content filters	Yes (add-on)
Content filters configurable by service	Yes
Content filter API	Yes (java based)
Latency	< 0.1 sec (not including filters)
Max transaction size	Depends on Max size of RAM.
Source system integration	Open (java based)
Supported Arbit systems	ATIR, ADG, AWG
Configuration and key generation	Stand-alone laptop
Running configuration changeable	No (runnable read-only boot media must be created on site)
Domain data flow	HIGH to LOW, LOW to HIGH, HIGH to HIGH, LOW to LOW
Regular maintenance	None (the system runs unattended)

9. Filtering

The Arbit Open API for added Security and Content Filtering is an inherent part of the secure by design philosophy and architecture of the Arbit Cross Domain Solutions.

The Arbit Open API provides users total control over their data that passes through the Arbit Trust Gateway, simplifying integration towards for example Command and Control systems, Battle Management Systems or any other software application. It allows customers to flexibly control and shape their security profile as they see fit, as it provides an added layer of verification when importing data through the ATG.

Arbit CDS solutions deployed in Multi Domain Operation scenarios, require integration between several classifications, physical domains and even across borders, where multiple protocols and standards are in play. The Arbit Open API enables secure integration of connectivity and communication when dealing with for example Federated Mission Networks using the Tactical Assault Kit platform with the Cursor on Target protocol.

The Arbit Open API can be used for creating filters or labelling services and the like that verify and check data against certain standards, such as STANG 4774 / 4778.2 which then gives a go/no-go decision to determine if that data should be allowed to pass through the Arbit Trust Gateway or not.

The Arbit Open API also enables easy integration towards external applications, such as virus scanners like ClamAV or MetaDefender from our partner OPSWAT, to ensure malicious code doesn't get imported and spread between servers in a secure network.

10. Data Integrity

The Arbit Trust Gateway utilizes the same mechanisms as the Arbit Data Diode and uses several layers of data integrity:

1. The unidirectional transport is based on UDP which includes a checksum ensuring that the individual packages are not corrupted.
2. On top of the UDP transport the ATG maintains strict transaction control, which is able to detect any lost packets, which ensures that all packages are processed in the correct order.
3. When transporting files and supporting TCP-like streams, the ATG uses a forward error correction algorithm which prevents almost all network transmission errors.
4. All files are verified by checksum after arrival on the Catchers/Guards.