Chang Hu Foo
August 15, 2004
GIAC Security Essentials Certification (GSEC)
Practical Assignment Version 1.4b, Option 1

Should Your Organization Implement a Data Diode?

Summary

A data diode placed between two networks is one mechanism that could be used to secure valuable information in a higher classification network, against the threats posed by the presence of the cross-network link. This essay attempts to help an organization in determining if it should be implementing a data diode.

The essay starts off by establishing the importance of securing an organization’s valuable information, against the risks posed by a connection to the Internet. For organizations that have two separate networks, a High Security Network (High Network) and a Low Security Network (Low Network), the key would lie on securing the cross-network link. An overview of the various mechanisms to secure the link is explored, to have a preliminary determination of a mechanism that would best suit the needs of the organization. This is achieved by performing a risk analysis utilizing the C-I-A triad for each of the methods and conducting a brief assessment of the impact that each would have on the operational environment. Then, the article delves deeper into the implementation issues related to the data diode and suggests some ways to address the raised concerns.

Risks of Being Connected to the Internet

Over the years, the Internet has grown to become an essential tool for organizations as they rely on it to perform a myriad of important activities, from sending and receiving electronic messages, researching for business-related information on the Internet to performing online commercial transactions.

On the flip side, by being connected to the Internet, the organization also exposes itself to threats posed to the critical information that it needs to judiciously protect. Malicious persons may capitalize on the network links to the Internet to compromise the security of such information. These information, which may include manufacturing processes, in-house research results and even employees’ compensation, must be vigilantly protected against. Otherwise, trade secrets may be compromised, manufacturing processes go awry or normal business operations made impossible. In severe cases, these consequences would deal a fatal blow to the very existence of the organization.

Securing the Link from Low Network to High Network is the Key

In many organizations, two networks are created to facilitate the mitigation of such threats. One, which shall be known as Low Network, is directly connected to the Internet and contains information that is of a lower value. The other, which shall be known as the High Network, would be indirectly connected to the Internet via the Low Network and contains information that is highly important to the organization.

In such a set-up, the protection of valuable information against threats as a result of being online, would thus lie on securing the link between the High Network and the Low Network.

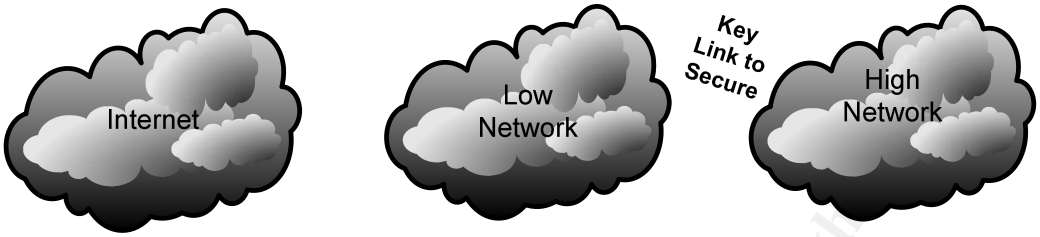

Figure 1. Key Link to Secure

The data diode has been touted as a useful mechanism to secure such a link. However, is your organization really suitable for the implementation of a data diode? This article aims to help the reader in answering this question.

Methods to Secure the Cross-Network Link

In the following section, an overview of the options available to secure such a link is conducted. This is to have a preliminary determination if there exists mechanisms other than the data diode that are more suitable for the organization. In particular, an analysis of the protection that each mechanism offers to the High Network against threats of being online is performed using the C-I-A (Confidentiality, Integrity and Availability) triad. Additionally, a brief review of the limitations that each of these mechanisms would have on the operations environment is done.

Network Isolation

Traditionally, the best method would be to separate the two networks completely by eliminating all network links between the High Network and the Low Network. Therefore, no nodes in the High Network would have any direct or indirect (via Low Network) connections to the Internet anywhere and at anytime.

Obviously, the High Network would be fully assured of confidentiality, integrity and availability of its contents against the threat posed by the cross-network links as these links are totally eliminated.

However in most organizations, this implementation would be impractical as there exists a need to transfer information from the lower classification network to the higher classification network conveniently.

Removable storage media like USB disk drives and floppy diskettes may be used in place. However, such a manual means of transfer is less convenient compared to the network means. This is even more so if we would want to transfer huge data files, especially when the transfer is to take place at frequent intervals.

Moreover, the employment of removable storage media as an information transfer mechanism would mean that the High Network terminals are accessible by such media and thus exposed to the associated risks. However, there exists measures to mitigate the risks if it is decided that the removable storage media are needed. These measures are covered in the second paragraph of a later section titled "Inability to Transfer Information from High Network via Network Links".

Firewall

Due to the above mentioned limitations, most implementations would have these cross-network links preserved. The higher classification network could then be protected by placing a firewall between the two networks.

Applying the Principle of Least Privilege, the firewall is to be configured such that it only allows the necessary application traffic between nodes in the higher classification network and nodes in its external networks to pass through. A great resource for choosing the appropriate firewall and configuring it securely would be NIST's "Guidelines on Firewalls and Firewall Policies"[1] which can be downloaded from http://csrc.nist.gov/publications/nistpubs/800-41/sp800-41.pdf.

Such an implementation offers great operational convenience for users as the firewall could be configured flexibly to allow the terminals in the High Network to communicate, via various necessary applications, with the terminals in the Low Network and the Internet. This also weakens the case for the opening up of terminals in the High Network to access by removable storage media as the network links provide a convenient mechanism for information transfer to the High Network from its external networks.

However, even with the elimination of unnecessary cross-network application traffic, the application traffic that is not blocked by the firewall remains as an avenue for malicious persons to compromise the security of information in the High Network. An insider could make use of the cross-network links to leak information out from the High Network and import tools from the external networks to perform integrity and availability attacks on the High Network. On the other hand, an outsider could make use of the application traffic that is not blocked by the firewall to gain access to terminals in the High Network to steal information, compromise the integrity of systems or launch Denial-of-Service attacks.

Firewall with Protocol Isolation

Applying the defense-in-depth strategy, a more secure way would be to have an additional protection mechanism in the form of protocol isolation. Protocol isolation techniques involve using network devices that utilize protocols other than TCP/IP as the means of network communication. These protocols are typically non-routable, like NETBEUI.[2]

The below figure shows how a variation of the method, protocol separation with server replication[3], can be used in conjunction with a firewall. Each of the border servers would be communicating with its counterpart in the opposite network using a non-routable protocol**,** like NETBEUI while communicating with other nodes in the same network using TCP/IP.

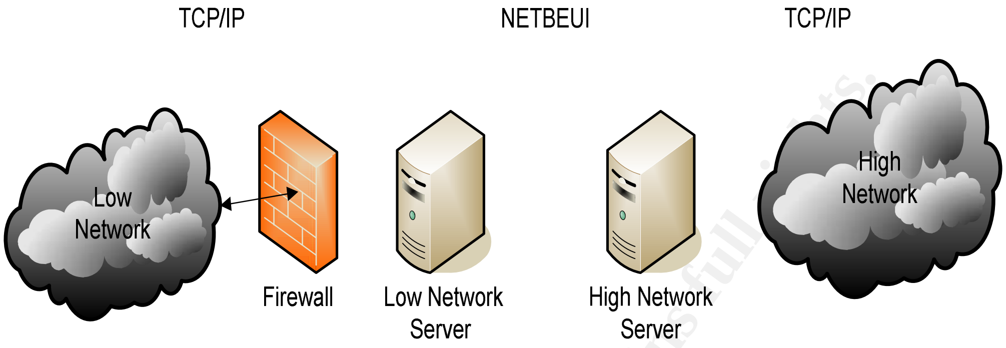

Figure 2. Protocol Separation with Server Replication, Used with Firewall

The addition of the protocol isolation with server replication mechanism further hampers any potential external hackers. The TCP/IP protocol that hackers ride on to access the organization's network could only bring them to the Low Network Server as the Low Network Server could not perform protocol conversion. Thus, compared to the case where only a firewall is used, the external hacker would need to be familiar with NETBEUI hacking also and compromise the two NETBEUI servers. Unfortunately, this addition does little to strengthen the security against attacks from insiders, who have legitimate access to the High Network terminals.

With the higher overall security assurance, comes also the trade-off. Users in the High Network would no longer be able to connect to the Internet to surf or send e-mails. But this set-up still allows file sharing between the two border servers such that information downloaded from the internet could be brought into the internal network via network means and vice-versa.

Data Diode

If the above method is still not secure enough, a data diode may be what is needed. A data diode can be defined as a physical layer device that allows information to flow in one direction but prevents any information flow in the opposite direction[4]. Examples of data diode implementations available in the market include Tenix's Data Diode[5], Owl Computing Technologies' Secure Information Transfer Systems[6] and QinetiQ's SyBard/Diode[7].

An example of how a data diode is implemented between the two networks is shown in the following figure.

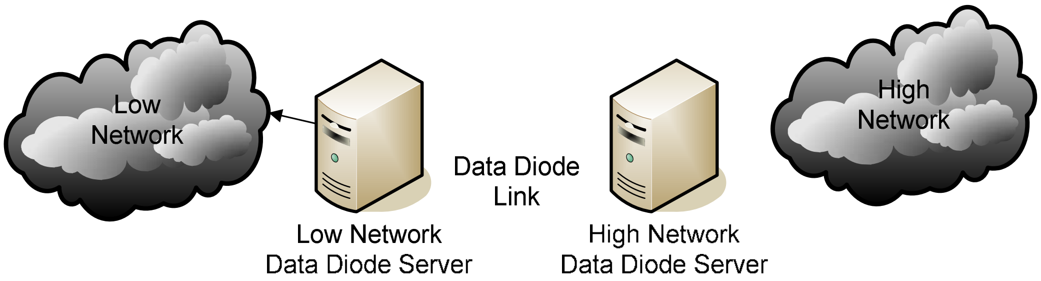

Figure 3. Data Diode between Two Networks

The data diode guarantees that neither external hackers nor insiders could compromise the confidentiality of information in the High Network via the network link. This is achieved by means of an optic fiber link, in which only the receiving terminal of the High Network Data Diode Server's fiber optic card is connected and not the transmitting terminal.

An illustration of how the fiber optic cards could be connected in a data diode is shown in the figure below. Two NIC cards are needed for the Low Network Data Diode Server as a carrier signal is required to the receive line of NIC1 for it to transmit data.[8]

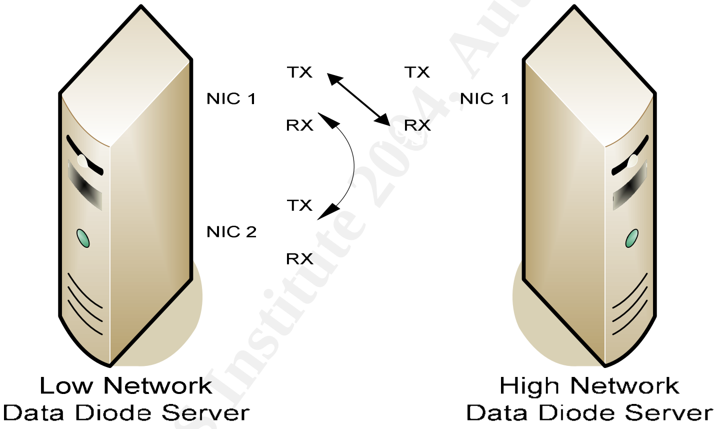

Figure 4. Fiber Optic Connection of Data Diode Servers

Besides ensuring a one-way connection, a fiber optic link has another advantage in that the data in transmission is not as susceptible to TEMPEST attacks as a copper link. This is because the electromagnetic radiation emitted is considerably lower.

Apart from addressing the confidentiality issue, the integration of a data diode makes it impossible for hackers from the Internet to perform reconnaissance on the High Network as there would be no information fed back from the network. Thus, it raises the difficulty to perform purposeful integrity and Denial-of-Service attacks on the High Network terminals as the hackers are not aware of the systems and the services that are running. [8] Insiders, though, would not be hampered, as they would still be able to bring malicious tools into the High Network via the data diode link.

However, the mere presence of the data diode does not completely address the issues of integrity and availability attacks by external hackers. Worms and viruses, which may be attached to the files that are transported through the data diode, could still infiltrate the High Network.

As with the previous method of securing the link with protocol separation and a firewall, this implementation would not allow the terminals in the High Network to communicate with the Internet to perform surfing and e-mailing. In addition, a legitimate transfer of information from the High Network via the cross-network link would no longer be possible.

Implementation Considerations

After the above comparisons, the data diode may have emerged to be a potential candidate to secure the link. If so, what are the other generic implementation issues to consider before deciding if a data diode really is the most appropriate security mechanism? These considerations are followed by some suggestions on how to address the raised concerns.

While the discussion of the issues is attempted to apply to all implementations of the data diode as far as possible, many issues are discussed with reference to a particular commercial implementation, to facilitate a more in-depth discussion. Tenix's data diode has been chosen it has the most amount of technical literature published for the public's reference.

Inability to Transfer Information from High Network via Network Links

As mentioned earlier, the data diode does not allow data to be transported via the network link from the High Network to an external network. Unfortunately, practicality demands that sometimes documents (like presentation slides or e-mails) are still needed to be sent out from the higher classification network. Further effort must be invested to establish an acceptable means (a balance between security and operational convenience) for the organization to legitimately transfer information out from the High Network before the data diode is to be implemented.

One possible way is to transfer such information via removable storage media such as floppy diskettes and USB drives. However, if left uncontrolled, the risks of information outflow through such means would be unacceptable. The control needed can be performed by an I/O access control software such as SecureWave's Sanctuary Device Control. Such a software is able to control the use of I/O devices according to different users. At the same time, it provides a log that includes the file name and the file itself to audit if authorized users have abused their rights. [9]

Reliability

Another implementation consideration would be the reliability of the data diode in delivering data successfully. Due to the very nature that it is a one-way link and hence the lack of a feedback path, data cannot be guaranteed to be transported to the destination side successfully. Data loss would occur when the High Network Data Diode Server could not cope with the high data rate that is coming from the Low Network Data Diode Server.

Such concerns are not ignored by the vendors. In the case of Tenix, it has addressed this issue by utilizing a few mechanisms. One of them is a choke mechanism that is used to balance the bandwidth across the data diode at any instance in time to minimize the occurrences of data packet losses.[10] Another is the allowing to configure the number of times the data diode would send a data packet.[11]

Despite the above measures, it is still not possible to ensure that there is complete reliability for data packets to be sent across successfully. Thus, there needs to be contingency measures to address the occasions where there is data loss. These measures are especially important when the data may be part of an important e-mail or critical document that demands to be sent across the data diode reliably. A mutual comfort level on the effectiveness of these contingency measures should be achieved with the relevant system owners prior to the implementation. Some possible measures are suggested in the following.

Users could be advised on that due to the one-way nature of the data diode, there would be occasional loss of e-mails or documents. Thus, if a user has been expecting to receive a document or e-mail from the external network but has yet to receive it, one could help himself by getting the originator to re-send the missing document or email.

For the cases where the recipient could not help himself, there should be logs in the High Network Server when file loss is detected. In Tenix's data diode, the Data Diode Servers could notify administrators of the occurrence of file loss via the SNMP trap mechanism.[11] Ideally, there should also be a mechanism available, so that upon the notification of a file loss, the administrator could help to re-send the file that is lost. Unfortunately, the presence of such a mechanism could neither be confirmed nor denied due to the absence of relevant literature.

Another reliability related concern would be the single points of failure presented at either of the two Data Diode Servers as neither of the servers could be part of a Cluster group to allow for failover.

An alternative, albeit one that incurs more down-time, monetary costs and is more manpower intensive, would be to have redundant hardware installed with the software on standby. They could then be manually switched over in the event of failure.[10] To minimize the downtime, it is advised that some server monitoring tool be utilized to help alert administrators quickly when a server is down so that swift action could be taken.

Throughput

Though the data diode may be rated to have a maximum throughput of say 100Mbps[9] as in the Tenix data diode, this is not possible in the light of overheads such as the choke mechanism and allowing the data diode to send data a few times. Thus in deployment, the actual throughput of the data diode would be considerably lower.

The limited throughput of the data diode, with the time criticality of some of the data (real-time monitoring would require data to be sent across immediately), may demand that more than one data diode to be deployed. An estimation of the number needed should be performed prior to implementation, as it would affect the setup and maintenance costs which may affect the decision on whether any data diodes are to be implemented.

Operating System

The operating system used by the data diode servers could also pose an implementation issue. For Tenix's data diode servers, the OS used could be either Redhat 7.1 or Solaris 8.[10] These operating systems are not as commonly used as the Windows operating systems. Consequently, one of the implications is that the organization may not have personnel who have the expertise in maintaining the chosen operating system. If so, it must be prepared to invest resources in developing its personnel to gain the relevant expertise.

Another similar implication is that the organization may not have an established set of guidelines for hardening the chosen operating system. The vendor may provide a set of guidelines but this may be only a baseline security configuration and not provide the level of security that fulfills the organization's needs. Thus, expertise needs to be developed also in the area of hardening the operating system. A good source to start off for securing Linux systems would be the book "Securing Linux: A Survivial Guide for Linux Security" by David Koconis, et al[12] while a good guide for securing Sun Solaris systems would be Hal Pomeranz's "Solaris Security: Step by Step"[13].

When faced with the choice of two operating systems for the data diode terminals, one consideration would be if the operating system has been discontinued to be supported by the vendor or if there are near term plans to do so. This is an important consideration as vendors would then be not obliged to release security patches for new exploits to their operating systems. In the case of Tenix's Data Diode, Redhat has ceased its support for Redhat Linux 7.1. [14] while Solaris 8 is still supported by SUN.[15]

Access Control of Data Diode Servers

As with all security devices, it is important to control the physical access to the data diode servers. With physical access alone, one could physically damage the terminals to render them non-functional. The functionality of the data diode servers could also be compromised by altering the physical connection on them such that it becomes technically possible to communicate from the High Network to an external network via the network links.

In fact, the physical access control requirements for the data diode servers should be as strict as the terminals in the High Network since they are used to protect them. These requirements could include locking them up in racks and rooms that require two-factor authentication, like a smart card and a complex password. Similarly, likewise access control requirements for the operating system and the data diode applications may also apply. These additional access control measures could add on to the costs of implementation.

Measures to Mitigate Residual Risks Posed by Cross-Network Link

As mentioned in the earlier C-I-A analysis of the security provided by the data diode, it does not fully address the integrity and availability risks to the High Network that results from the presence of the cross-network link. This section looks at how these risks could be mitigated.

One approach would be to content filter and conduct virus scans to the data before it enters into the High Network via the data diode. Tenix's data diode provides an interface for content filtering software such as Mailsweeper for SMTP[10], which in turns caters for the integration of an anti-virus software. By so doing, malicious code, worms and Trojans are less likely to pass into the High Network via the data diode and thus risks of integrity and availability attacks from these elements are significantly reduced.

Another approach would be to apply the Defense-in-Depth Strategy again, and position another protection mechanism before the data diode, in the form of a firewall. This helps to make sure that the data diode servers and thus the High Network is not accessed by non-authorized terminals, or authorized terminals via non-authorized applications.

Policies could also play a role here. It could be stated that staff bringing materials into the High Network Link via the data diode link are to ensure that these materials are strictly work-related. And if they are downloaded, they are to be from trustworthy websites, such as the websites of reputable software vendors. To ensure compliancy, penalties for non-conformance could be spelt out.

Conclusions

In this article, security mechanisms that help secure the High Network against the threats resulting from the presence of the cross-network link to the Low Network are compared. These mechanisms include having no cross-network links at all, using a firewall, using a firewall with protocol isolation and utilizing a data diode. A data diode addresses such threats significantly though its ability to prevent any data leakage from the High Network. However, this comes at the expense of limiting what users in the High Network could access to.

In addition, there are some other implementation issues to consider. They include the inability of the data diode in guaranteeing that data would be sent across successfully and the inability to allow legitimate transfer of data from the High Network to external networks via network means. Some suggestions on how these concerns could be addressed are discussed.

Should your organization implement a data diode? It is hoped that the article has placed the reader in a better stead to answer this question.

References

1. Wack, John, Ken Cutler, and Jamie Pole. "Guidelines on Firewalls and Firewall Policies". January 2002.
   URL: http://csrc.nist.gov/publications/nistpubs/800-41/sp800-41.pdf

2. Edwards, Mark J. "Understanding Network Security"
   URL**:** http://www.windowsitlibrary.com/Content/121/02/2.html

3. "The organization and implementation of WMO Global Data Processing Systems."
   URL: http://www.wmo.ch/web/www/reports/nyoni_part2-4.html

4. Mansfield, Aubrey. "Security's Last Stand: Isolation", January 2004
   URL: http://www.naspa.com/PDF/2004/0104/T0401004.pdf

5. "Tenix: Veto DD"
   URL: http://www.tenix.com/Main.asp?ID=723

6. Owl Computing Technologies, Inc., Secure Information Transfer System
   URL: http://www.owlcti.com/

7. "SyBard::Sentry"
   URL:
   http://www.qinetiq.co.uk/home/markets/security/securing_your_business/information_and_network_security/secure_products/sybard_sentry.html

8. Westmacott, Jason. “Uni-Directional Networking”. 06 March 03.
URL:
http://www.giac.org/practical/GSEC/Jason_Westmacott_GSEC.pdf

9. “SecureWave | Sanctuary Device Control”
URL:
http://www.securewave.com/turcana/securewave/sanctuary_DC.jsp

10. Veto Technical FAQs – Frequently Asked Questions
URL: http://www.tenix.com/Print.asp?ID=752

11. “Veto Uni-directional Network Bridge and Data Pump Applications White Paper”. 2002
URL: http://www.tenix.com/PDFLibrary/130.pdf

12. Koconis, David, et al. Securing Linux. A Survival Guide for Linux (Version 1.0). SANS Press, February 2003.

13. Pomeranz, Hal. Solaris Security: Step-by-Step (Version 2.0). SANS Institute, February 2001.

14. “Redhat.com | End of Life Products”
URL: https://www.redhat.com/security/archives.html

15. “Solaris Operating Systems Releases
URL: http://wwws.sun.com/software/solaris/fcc/releases.html