e-Business and the Need for “Air Gap” Technology

A Whale Communications White Paper[1]

Frederick M. Avolio
Avolio Consulting
<fred@avolio.com>

The Challenge

As use of computers has grown over the years, moving from being instruments of highly specialized users to a ubiquitous tool in the office and home, the business requirement for sharing of information has increased exponentially. Information sharing and information security are often at odds with each other for the simple reason that distribution and free flow of critical business information must be secured and controlled to protect against malicious data. Today, this need is most critical in the area of e-business.[2]

The fundamental requirement for successful e-business deployment is the ability to integrate front-office/Internet access to the businesses bedrock back office applications. The obvious question that arises during an e-business implementation though is, “How can we do this without putting the whole back office at risk?”

A successful e-business strategy depends on “externalizing” back office applications to users outside the trusted corporate network. Unfortunately, most experts rightly consider this kind of access to be high risk, and thus many companies limit what on-line services they offer to customers, partners and other external users. Custom-made solutions, such as data mirroring or protocol switching, have been used to minimize the risk. Yet the essential problem remains.

Firewalls are the most prevalent security technology for classical Internet applications, however firewalls are not designed to provide access to crucial back office data. Firewall vendors themselves recommend against this sort of connectivity. Vulnerabilities in the underlying TCP/IP protocols, as well as the firewall software and hardware itself, make web-to-back office integration via a firewall a dangerous proposition for organizations.

This white paper discusses the options and products security professionals have been using to attempt to address the 2-sided coin of information security and information sharing. It then talks about a new class of gateways — the “air gap.”

[1] Much of the text of this paper was taken from the author’s product review of Whale Communications e-Gap™ product.

[2] There are many definitions of e-business. e-business includes “e-commerce,” but it is more than selling and support to customers. It also includes utilizing extranets, intranets, and the Internet to manage business transactions with business partners and suppliers, and to carry on business intelligence.

Connecting the Trusted to the Untrusted

Consistent throughout history has been this challenge: how to protect the trusted and sensitive — that which is highly valued — from the unknown and untrusted, while allowing for interaction, communication, or commerce. The Great Wall of China might have been more formidable a barrier to attackers if it did not have gatehouses. Medieval castles would have been more secure if they did not have drawbridges over the moats. Moreover, an organization would have less cause for concern about computer virus infections or intrusions launched from the Internet if it had no connection to the web. But then the Chinese would have limited their ability to exit out past the wall, the castle-dwellers would have been stuck inside — just as Greeks bearing gifts would have been stuck outside, and the enterprise would lose the ability to use the Internet for business transactions and communication.

In the internetworked world, as organizations open up to the web, an outside attacker can use these lines of communication to gain access to sensitive information. A hacker can potentially use the same avenues that a business partner uses to access critical systems and data.

Sneakernet

Initially file transfer between computers was done via an extremely secure avenue -- “sneakernet.” Before the advent of widely networked IT infrastructures, if one wanted to transfer files or information between two computers, one would copy the files from one computer to a floppy disk or magnetic tape, and physically carry the media to the computer the data was to be uploaded into. Someone walked the files across a gap between the computers. Because the computers were not physically connected, there was no attack path to the system.

Obviously, sneakernet has severe limitations. This follows an axiom about the relationship between security and usability: Security and usability are inversely proportional. Another way of saying this is, complete security yields zero usability and complete usability results in zero security. Since connectivity is the norm in today’s digital economy, sneakernet is not an option.

Electronic Business

The advent of e-business driven by Internet accessibility. In order to do this effectively, customers, partners, suppliers and other interested third parties need to be able to access information traditionally stored on back office systems. Potential customers are accessing product information, purchasing those products or communicating with corporate resources, from an organization's web site. Business suppliers, partners, and remote sales forces need access to critical back office data such as inventory, pricing, customer information and other sensitive information.

This is some of what is holding back the growth of e-business today. E-business requires connections to the back office, wherein is customer information, inventory information, and price sheets. However, a connection from the outside to the back office puts too many assets at risk. Traditionally a human go-between (sneakernet) has been used. A prospective buyer telephones an operator, the operator has access to the back office and, thus, this maintains separation between the outside user and inside information — at a cost. This does a good job of keeping the prospective buyer from attacking the back officer, but is insufficient for the fast-paced world of the Internet and e-business.

e-Business Architectures and Air Gaps

Typical e-business architectures utilize a web server that is strategically placed on a DMZ that affords the server some corporate firewall protection. This mode works fairly well if the web server serves up static content and does not require back office access.

Corporations who are aligning their e-business initiatives with their business objectives normally require access to back office systems to properly implement them. Ideally, an application web server would be placed behind the firewall, on the trusted side of the network, and users would access the various back office systems. However, in an e-business scenario, many of the users are beyond the corporate firewall and are accessing these mission critical applications via the Internet. Therefore businesses often end up placing their applications servers in the DMZ, which affords limited security, as depicted in Figure 1. In this scenario, clients connect from the Internet to the application web server. Requests are then funneled through the firewall to the appropriate back office systems.

Since the application server has access to the back office, it represents a weakness in the security architecture. If anyone gets control of this system on the DMZ, an attacker could potentially have complete access to the back office. Furthermore, if there are other systems on the DMZ — an e-mail relay or a dial-up server, for example — this heightens the risk of an attack through these systems.

Figure1

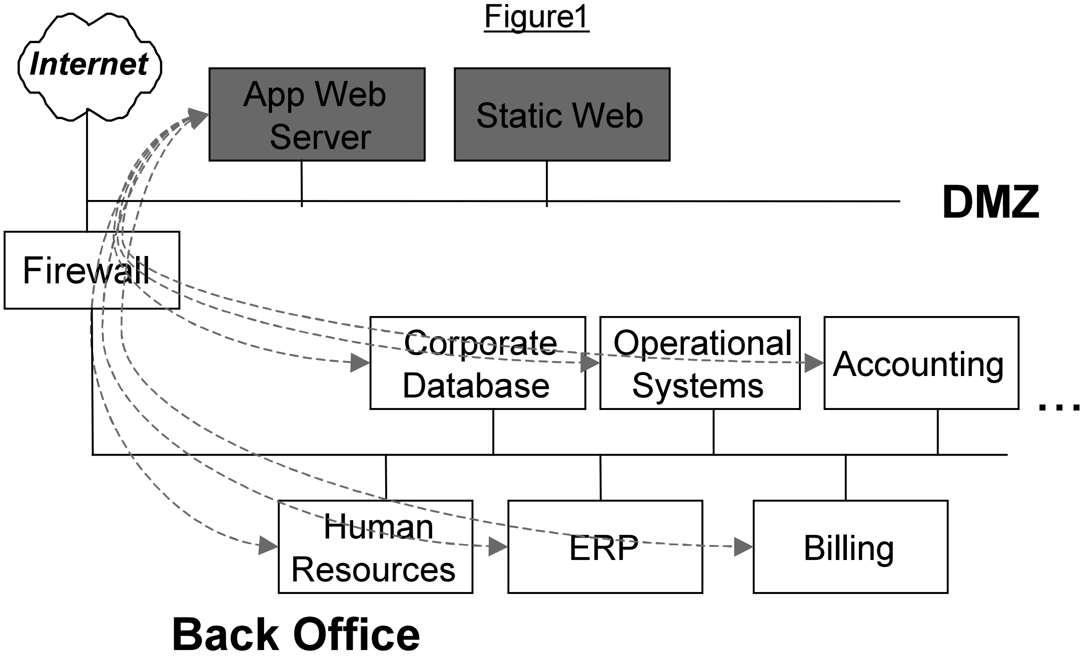

e-Business Connectivity Based on Physical Disconnection.

Overview of the Technology

An answer to the challenges of deploying an effective e-business architecture is a high speed air gap. An air gap secures the transaction path between e-business servers and internal corporate applications and databases, protecting an organization's back office from attack or compromise. In addition, it secures the path of data transfer between any two networks, trusted and untrusted. It accomplishes this by physically disconnecting the two networks, yet shuttling data between them. This enables selective transmission of certain highly focused applications between networks, while maintaining absolute physical disconnection between them.

In addition to this physical disconnection, a properly designed air gap eliminates all the protocol layers responsible for transmitting the critical data across the untrusted/trusted network boundary, including the physical layer. In this way, it can eliminate clandestine attacks on a network. This is in contrast to firewall technology, for example, which *uses* the current infrastructure and tries to enforce legitimate usage of it. The use of an air gap resembles the security system at an after-hours gas station, where the cashier sits behind bulletproof plastic. At no time is there any direct contact between the cashier and the client, however they can exchange money and payment receipt via a metal drawer. This setup fully protects the cashier and the cash register against external threats, since there is a complete physical disconnection between the two, yet the cash transaction can easily take place.

Description of the Technology

An air gap maintains a physical disconnection between two networks while enabling selective data transfer between them. An implementation may use a switching device that switches memory storage (an electronic disk, or e-disk) access between two computer hosts, where each server host is on a different network.

Figure 2 pictures one possible solution. The data transfer channel between the networks utilizes the following path: an incoming web-based transaction is passed to a server residing on the untrusted side of a corporate network. The server proceeds to write the data on the e-disk. Once the data has been written to the e-disk, the switch rapidly disconnects the e-disk from the untrusted server and connects it to the server on the trusted side of the network, which then reads and passes the data on to the back office systems. (See Figure 2.)

Figure 2

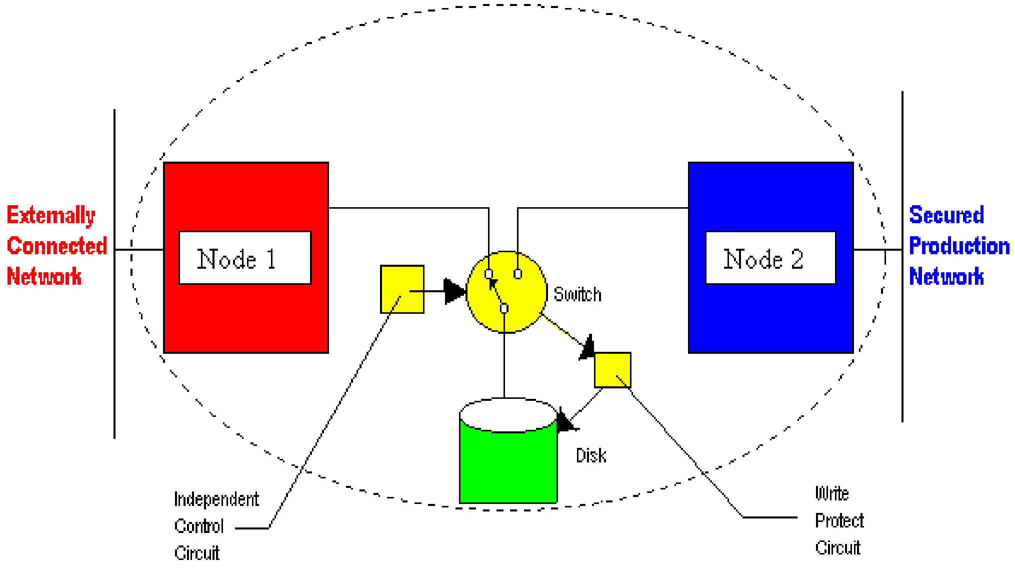

E-business: Linking the Application Web Server with the Back Office

Figure 3 shows a suggested configuration: a virtual internal web host — the only system that communicates with the back office — and a virtual external web host — the only system with which clients can communicate. Note, there is no real web server on the external machine. The air gap sits and operates between to two virtual servers.

This configuration also protects authentication servers and certificate servers, keeping them are out of hackers' reach. Further, an encrypted session — in support of an e-commerce transaction (typically an SSL connection) — remains encrypted all the way to the trusted internal back office. Currently, most encrypted transactions terminate at a server residing in the DMZ, thus leaving a potential security vulnerability in place.

Figure 3

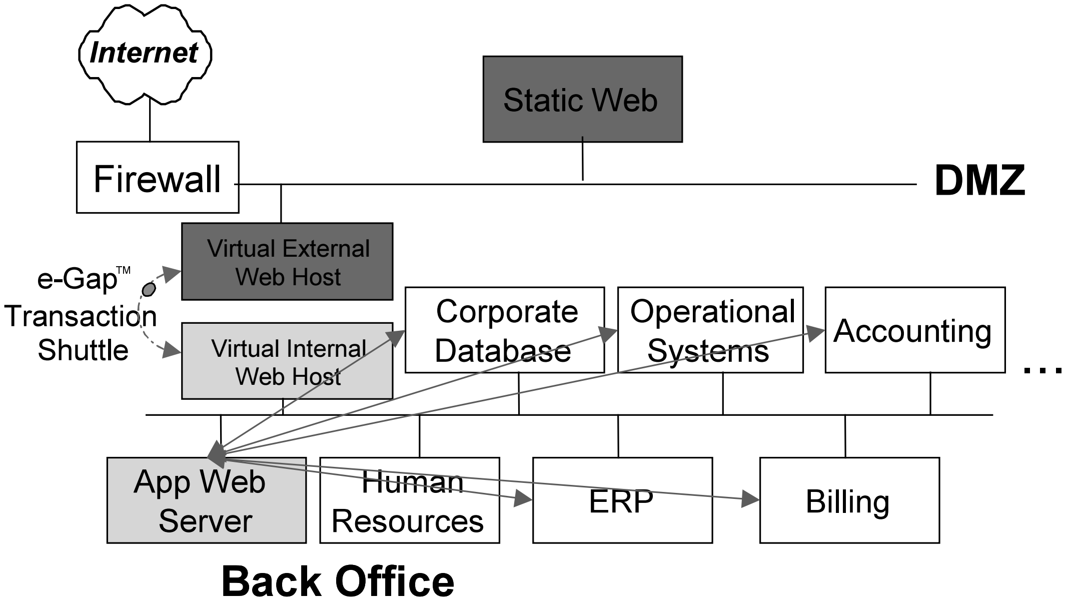

Objections and Concerns

Immediately, a few concerns should come to mind. First, traditional air gaps have demonstrated poor performance. The picture of someone running across a room with a magnetic tape recently removed from one computer to load into another computer is not one that speaks of speed. Can an air gap be fast enough to support e-commerce solutions? The answer "yes," assuming you use the right hardware and software. We have seen this technology using an extremely fast Ultra-wide SCSI standard run at greater than 80Mbit/second with no noticeable communications latency. E-commerce sites report that performance issues lie primarily in the applications themselves, rather than in the network infrastructure that transmits the data from the corporate boundary to these applications.

But is this system too restrictive? Upon careful review, we would have to say "no". In these business critical environments, an air gap would a high level of security while maintaining an open channel from untrusted to trusted network. It is not meant to be applied to every application. It's very appropriate for sensitive applications that require communication to back office, or where disconnection must be provable.

Others have raised the question as to whether an air gap is redundant if a firewall is in place. It is not. It is used in conjunction with an Internet firewall (as in Figure 3). It is for very particular services where tight control that enforces disconnection is required. This is especially true of the needs of e-business to connect securely to the back office. It is ideal for extranet or intranet connectivity where one-time configuration and minimum administration is important.

Conclusion

Providing service and security has been and continues to be the challenge for any organization wanting to use the Internet for business today. Simple mechanisms have been replaced by ones that are more complex as Internet services have become more elaborate, and less secure, and as people have started to demand these services.

The most secure kind of internetwork connection is *no* connection — an air gap. Air gaps have been used for centuries in physical security, and decades in communications. As the requirements have become more sophisticated, as organizational needs have grown, and as the threats have become greater, the requirements for both communication and a provable air gap have grown. Connecting the organization's external servers — servers for business partners and customers — to the organization's operational systems, requires secure enforcement of the separation between the networks.

An air gap can provide high-speed virtual online access, while enforcing physical separation between trusted and untrusted networks, and it can do it so with simplicity.