Secure Infrastructure Design

Scott C. Zimmerman, CERT® Coordination Center

Information technology (IT) and enterprise networks have become the core of many organizations. Critical business functions often depend on a fully functioning IT infrastructure: no network means no ability to generate revenue. To this end, an organization's growth and evolution should be reflected in the growth and evolution of its network. Organizational changes can include new or expanded missions, new factors such as mobile workers, and growth or downsizing in response to purely external factors. Infrastructure changes that stem from these factors can include additional network components (of a type already present), new types of components, and additional subnets or Internet connections.

This paper describes the fundamental components of infrastructure design, provides an overview of risk management concepts, and illustrates samples of network topologies.

Background

A brief discussion of some common network entities will lay the groundwork for later discussion.

1. **Servers**
   Servers are probably the most visible network systems. Most people—even non-techies—are familiar with the concept of a web server if nothing else. Each server consists of an underlying platform, which includes the system hardware and an operating system such as Unix or Windows. Running on top of the operating system is a server application to support services such as email, web service, or a database.

2. **Hubs**
   Hubs are passive devices used to connect multiple computers. The interior of the hub is wired in such a way that traffic destined to one machine on the hub is available to all machines on the hub. Hubs are inexpensive and easy to deploy but there is a drawback: if multiple machines are located on a hub and the local network segment is very busy, the broadcast nature of the hub may cause a prohibitively high amount of packet collision. A broadcast network also allows users to capture traffic destined for others on the local segment.

3. **Switches**
   Switches look like hubs, but there is an important difference: switches are designed so that traffic for Machine A goes only to Machine A; Machines B and C will not receive A's traffic. Using individual circuits—vice broadcasting packets to all machines—allows more efficient allocation of the network's locally available bandwidth: machines on the subnet will not be broadcasting all communication to all other machines on the subnet. Switches may be passive or managed. Some managed switches are shipped with rudimentary operating systems, allowing an administrator to control traffic flow more effectively.

4. **Routers**
   Simply put, routers control the direction of traffic on the Internet. A fair parallel would be the telephone network equipment that routes communications from the caller to the callee. A data packet that is sent across the Internet is routed based on its destination address; the sender need not know how the packet gets from point A to point B. The default gateway router on the sending network examines the packet, and either sends the packet directly to its destination or sends it to another router that will direct the packet to its next hop. Eventually the packet will arrive at the gateway of the network to which it is addressed, and the gateway will direct the packet to the destination machine.

   Many routers can do more than route packets. Newer models contain a fair number of security functions and can act as a simple firewall. Routers can filter packets on individual interfaces based on source and destination IP address and port. For example, a router with a web server on interface 2 can block all traffic except that addressed to port 80 on that interface. It is important to note that routers generally will not filter based on packet payload; an actual firewall is necessary in this case.

5. **Firewalls**
   A network firewall should not be confused with the "personal firewalls" that have become popular on desktop machines. A firewall in the network sense has two (or more) interfaces, and traffic comes in one side and goes out another. While passing through the firewall, packet source and destination IP address, source and destination port, packet payload, and other characteristics will be examined to determine whether the packet should be allowed through the firewall. To extend the example from the router analogy above, only packets addressed to www.example.com:80 will be permitted to reach the web server. However, what if the packet contains the signature of a particular buffer overflow attack? If the firewall has been configured with an appropriate rule set, this attack will be blocked and the malicious packets will not reach the web server.

Risk Management

Risk management is a field unto itself, and a detailed treatise is beyond the scope of this paper. However, there are two basic tenets that can guide personnel and organizations when making network infrastructure changes.

Risk Management Concept #1 – The Principle of Least Privilege
The Principle of Least Privilege has been around for many years, and it can be summed up rather succinctly: **anything that is not expressly permitted is denied**. This applies particularly well to the world of network security, and is the impetus behind the minimalist approach to publicly accessible systems, that is, do not run any services that are not absolutely necessary. In the realm of physical security, for example, this concept is demonstrated in the distribution of keys and other access devices. Does Bob have a legitimate business need to have a key to the lab in which he works? Of course he does. Does he need a key to the CEO's office? He almost certainly does not. To put this in IT terms, do all workers need access to the web? In many cases such access does help the employees do their jobs. Do these workers need access to peer-to-peer file-sharing services such as Napster? Almost certainly not, and the firewall rule set should reflect this position. Personnel should have access to the resources they need to do their jobs—no more and certainly no less.

Risk Management Concept #2 – Compartmentalization of Information
Readers who have worked with the Department of Defense or other government entities in a security capacity will likely recognize this idea; it is often described as the **need to know**. This is a subset of the Principle of Least Privilege that deals primarily with information. For example, does the engineering staff need access to their department's AutoCAD archive? One can make a strong case that they do. Do they need access to the Human Resources or Payroll databases? To put it another way, will the engineering staff's work suffer if they do not have access to the HR or payroll systems? Of course not, and this should be reflected in the network design and configuration.

Sample Network Topologies

We will discuss three network topologies, ranging from simple to complex, in the context of risk management.

Simple Topology

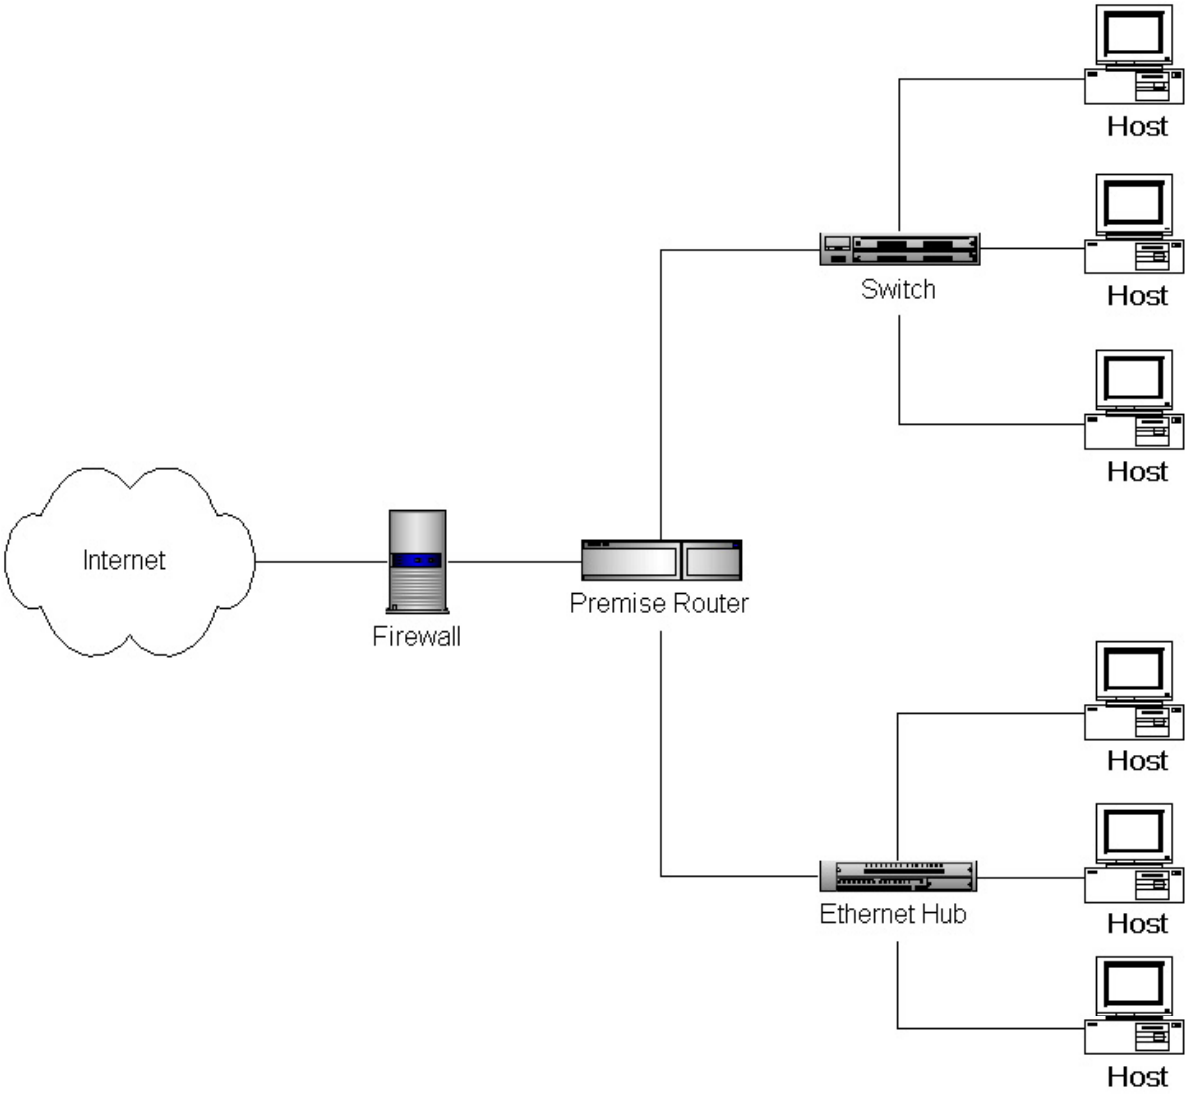

In this network, the entire enterprise is protected by a firewall. The premise router separates the two workgroups but can allow filtered connectivity as needed. The upper workgroup is using a switched segment to discourage packet sniffing on the local subnet and to provide better network bandwidth utilization. The lower workgroup is shown using a hub for demonstration purposes.

Moderately Complex Topology

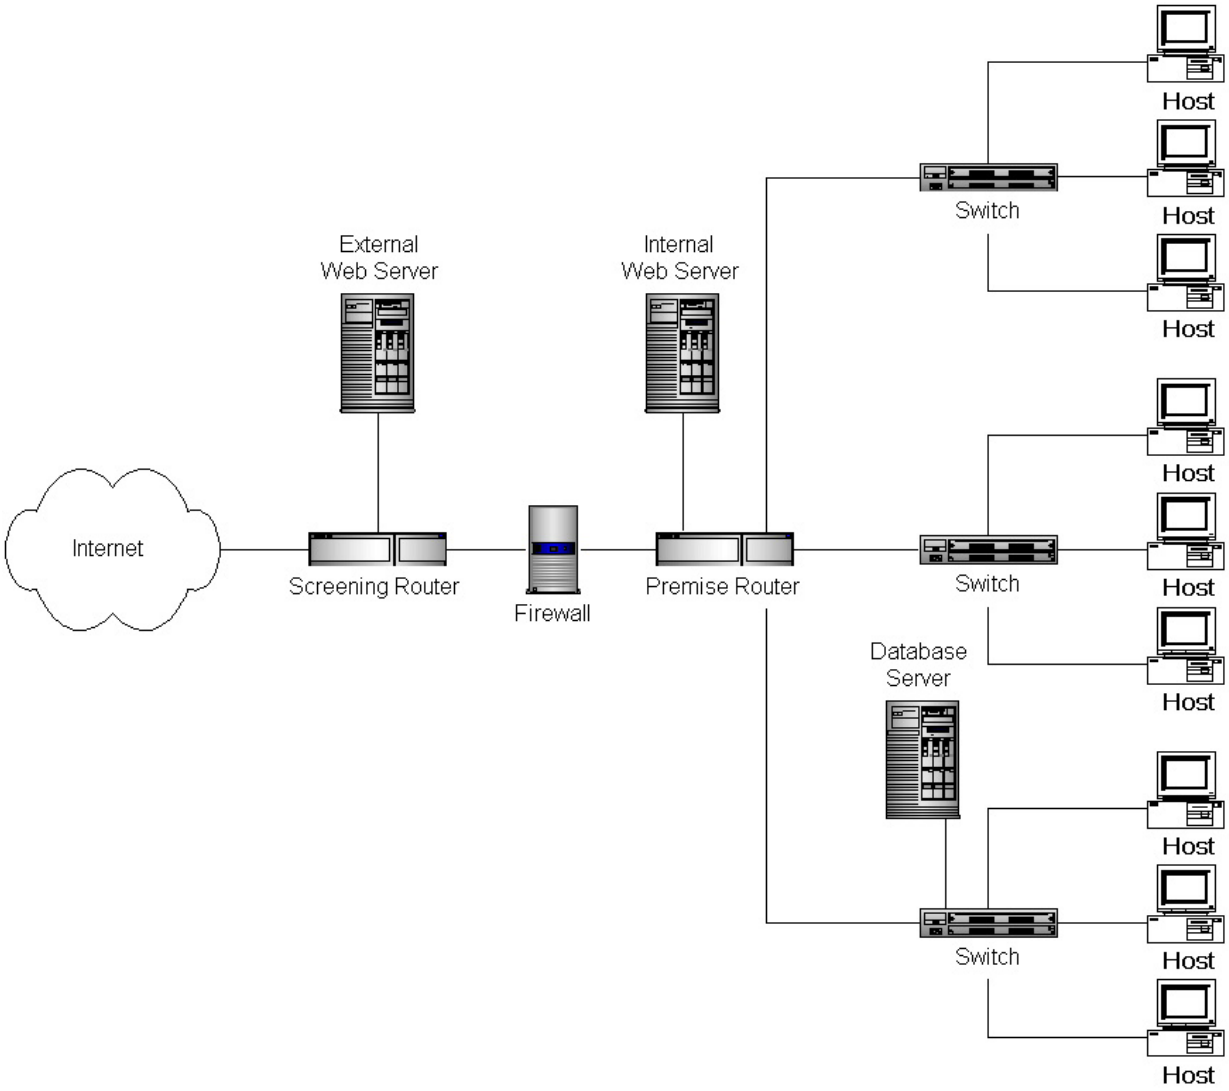

In this network model, various network functions have been isolated to minimize exposure. The external and internal web servers are on their own router interfaces with unique rule sets to govern traffic flow. The firewall filters traffic to **and from** the internal network to ensure no malicious payload can enter or leave the organization. Deploying additional firewalls between the routers and the web servers can prevent malicious traffic from reaching the servers. Behind the firewall the premise router separates the various workgroups according to mission. All three workgroups are on switched segments, but the bottom workgroup has an additional component: a database server. This server can be isolated from the other two workgroups using filtering rules on the premise router. Since the segment is on a switch the members of the workgroup cannot easily capture the packets of other users. Traffic to and from the database server can be limited to only individuals with authorization.

Complex Topology

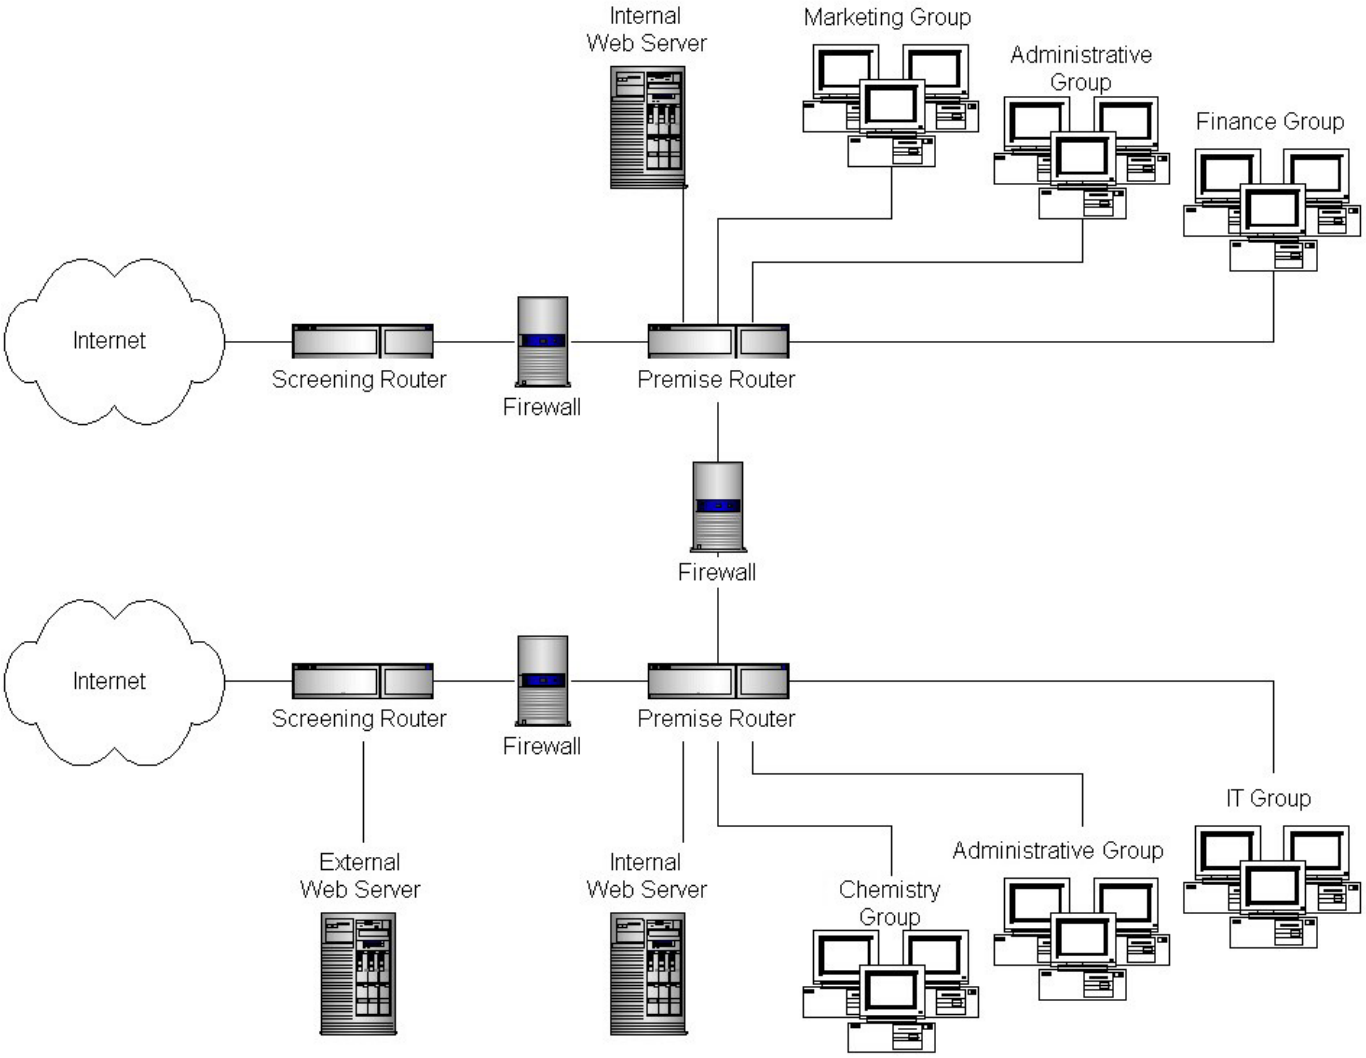

Here is a hypothetical large corporation. The mission is complex, and so is the layout of the network. Dual Internet connections provide redundancy for business continuity. Each connection has its own screening router and firewall, and the external web server is on its own appropriately filtered router interface. Behind the firewalls are the premise routers, where each workgroup is connected using its own—again, appropriately filtered—interface. This means, for example, that the chemistry group's access to the IT group's segment can be controlled in a broad fashion. Note the additional firewall between the two premise routers. This firewall controls the type of access that the upper workgroups may have to the lower workgroups' resources, and vice versa. Using a compartmented approach to the application of technology, this organization can control access to its many assets.

Conclusion

Secure infrastructure design can be a tricky proposition, but a methodical approach to planning will pay dividends in implementation. The first step is a thorough evaluation of the organization's current and potential business needs and assets, grouped by mission, capability, and requirements. The network design should support compartmentalization of information, and should allow room for expansion as the organization grows. Mission and security requirements can change rapidly and should be re-evaluated on a regular basis; once a month would not be too often. Personnel should keep abreast of industry developments: new technologies may be helpful in solving old problems. Finally, the two principles of risk management are a good yardstick for evaluating proposed changes to a mission-critical infrastructure.
