

The Handbook of Information Security

John Wiley & Sons

Chapter 25. Routers and Switches

Mar 31, 2005

- FINAL VERSION -

OUTLINE

INTRODUCTION

Principles of Routing and Switching

ROUTERS AND SWITCHES IN A NUTSHELL

How Routers work

Router Architecture and Performance

How Switches work

Switch Architecture and Performance

ROUTER AND SWITCH SECURITY

Best Practices in Securing Routers and Switches

Router Vulnerabilities and Attacks

Router Configuration and Deployment Practices

Switch Vulnerabilities and Attacks

Switch Configuration and Deployment Practices

Case Studies

CONCLUSION

GLOSSARY

REFERENCES

KEYWORDS

Access Control Lists, Denial-of-Service attack, Firewall, Internet Infrastructure Security, Router Security, Packet Filtering, Packet switching, Routing Protocol, Store-and-Forward, Switch Security

ABSTRACT

Routers and switches perform the essential task of transporting packets through the Internet by creating dynamic communication paths between hosts based on application demands. As key components of the Internet infrastructure, routers calculate and maintain forwarding tables to determine the best possible paths where packets are to be sent based on their header information. Switches connect successive segments of a network, providing the actual link for a packet from input port to output port. Routers and switches can be found in local access networks as well as in wide-area backbones, catering to local and global communication needs. Through decades of Internet evolution, digital switches and routers have become increasingly more complex, reflecting the need for networking services at greater speed, quality, and increasingly security.

Greater network complexity, ubiquity and pervasiveness of Internet services go in parallel with rising concerns about cyber terrorism and the sophistication of security breaches and attacks. Internet security concerns trusted information exchange as well as securing the infrastructure itself. Attacks on Internet infrastructure can have disastrous consequences, as different components have mutual, implicit trust relationships. This article presents the working principles of switches and routers as pillars of a dependable Internet infrastructure, with special regard for security concerns. The main focus will be on generic IP router and switch architecture, functionality, and protocols in wired, non-optical networks.

INTRODUCTION

Communication networks are systems in hardware and software to facilitate information exchange in a broader arrangement than a single point-to-point link. The telephone network is the most familiar and ubiquitous communication network, designed primarily for voice transmission. A computer network is a communication infrastructure between computing devices in different locations to enable digital information exchange and sharing of resources such as messages, software, compute time, storage space, or peripherals. Finding efficient ways to share communication links is one of the main problems in designing networks. Network design traditionally seeks to optimize several criteria at the same time: minimizing the cost of deployment; maximizing the aggregate bandwidth between two end-points; avoiding "hot spots" where a small number of network nodes and links handle a large percentage of the total traffic; minimizing the latency between any sender and receiver, and maximizing scalability, that is, performance should scale with the number of nodes in the network.

Routers and switches are the key devices in meeting these demands and are traditionally attributed to Layer 2 and 3 in the 7-layer OSI model and the 5-layer TCP/IP protocol stack, with the latter shown in Figure 1 together with respective processing units. Routers process packets, whereas Layer-2 switches operate with frames. Layers indicate which device uses specific header information pieces to decide how to process and forward a message with the goal to cooperatively move frames and packets from one network segment to another to ultimately reach the designated receiver.

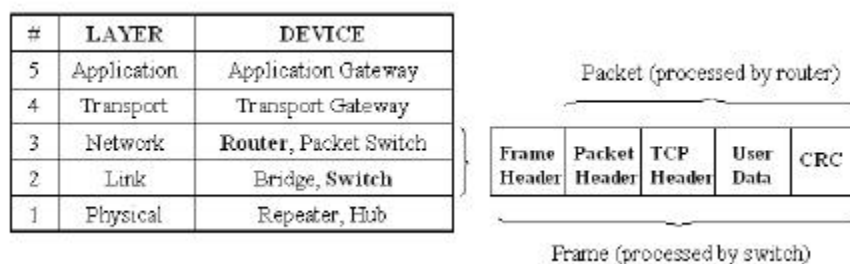


Figure 1. Routers and switches in the TCP/IP protocol stack and processing units.

Today's routing and switching technology is a legacy from the early packet switching networks created for the ARPANET in the 1970s. With increasing Internet population and complexity of service provision, infrastructure security for more dependable communication is a major concern. A plethora of new services creates new loopholes for compromising the security of a computer network and require steady countermeasure investigation and innovation. This article describes the working principles of routers and switches, their prominent security vulnerabilities, and suggests some best practices for router and switch security.

Principles of Routing and Switching

Although the various types of digital communication networks differ in how they are used and transmit information, they have routers and switches as architectural components in common. Routing is provided in a layer above switching and concerned with establishing an optimal communication path. Routing may also take Quality-of-Service concerns into account (Shenker, 1995). In the Internet, many protocols work together to accomplish this goal to establish end-to-end connectivity. For interior routing within administratively autonomous domains, protocols such as the Routing Information Protocol (RIP) and Open Shortest Path First (OSPF) are predominant. RIP is used in small to medium-sized networks and is simple to configure but slow in responding to network failures, which in itself can be a threat to network security. OSPF scales well in mid-size to very large networks, but incurs greater messaging and computation overhead. For inter-domain, policy-based routing the exterior path-finding protocol BGP (Border Gateway Protocol) is used as a standard in large routers (Huitema, 2000).

On the other hand, switching is mainly concerned with data relay and the associated policing and output scheduling. Switching is the mechanics of creating a connecting path between input and output in a switching device, with the goal to effectively relay a message from an incoming link to a chosen destination designated by an output port. It has played a vital role in telecommunication networks since first-generation telephone exchanges were installed. Switching technology seeks to maximize capacity for a given cost and reliability, and to minimize blocking and packet loss. With switching, data can be forwarded across dynamically created communication paths, whose sections are multiplexed among various senders and receivers. Switching is performed by telephone circuit switches for voice calls in the Public Switched Telephone Network (PSTN), "virtual circuit" switches in Asynchronous Transfer Mode (ATM) networks with fixed-sized cells, or packet switches with variable-sized messages.

Circuit-switching, based on telecommunication services, pre-establishes a path from one end-system to another via a sequence of switches, where device resources are reserved to guarantee a dedicated circuit for the duration of the transmission. A circuit switch must reject a call if it cannot reserve a path from input to output, which is referred to as call blocking. Circuit switching is implemented as connection-oriented service in the Public Switched Telephone Network (PSTN), where voice samples are switched from a source to a receiver, or with virtual circuit-switching in ATM networks. Recent progress in peer-to-peer telephony has shown that it is also possible to establish connection-oriented communication in a self-organizing peer-based logical network.

In contrast, device resources are not reserved in packet-switched networks, but shared and utilized on demand by using statistical multiplexing among data packets, which are composed of a header with control information and the payload. Packet switching can be implemented connectionless using the Internet Protocol (IP), or in connection-oriented mode using for example ATM. Routers performing connectionless packet switching are also called datagram or packet switches. In connectionless switching, packets are self-contained and independently routed towards a destination, whereas connection-oriented switching associates data to a specific, fixed forwarding path. Table 1 shows networking vs. switching modes with their prototypical implementation (Keshav, 1997). Besides ATM there are other connection-oriented packet switching technologies such as frame relay or X.25.

	Networking modes	Connection-less	Connection-oriented
Switching modes	Packet Switching	IP	ATM
	Circuit Switching	—	PSTN

Table 1. Switching vs. Networking Modes.

ROUTERS AND SWITCHES IN A NUTSHELL

How Routers Work

Routers are Layer 3+ devices (the + means that these devices may perform higher-layer routing functions.) A router forwards packets based on IP addresses and routing tables. Routers are available for small and large scale networks, for wired and wireless networks, and for non-optical and optical communication networks. Home networks typically are equipped with DSL or Cable modems connected to access network routers. At the network edge we find Enterprise Wide-Area-Network (WAN) access switches and edge routers. Backbones use Carrier Class routers besides ATM Switches or Frame Relay Switches. Unlike switches, routers always require some configuration.

Routers pass traffic between two different IP networks which may be either LANs or WANs, forwarding packets on the best-possible path towards a destination, in terms of the number of hops, the cumulative delay, or some other optimization criterion. When a packet arrives at a router, the frame header and trailer are stripped off and the packet contained in the payload field of the frame becomes input for the routing software, which uses the packet header to choose an output line. Packets may be queued in a router until they can be switched through to the appropriate port based on a routing table lookup to find the best-possible path. An IP packet processed by a router will contain a 32-bit (IPv4) or 128-bit (IPv6) address. The routing software is oblivious of the frame address and

whether the packet arrived from a LAN or point-to-point link. A router examines the destination IP address of each incoming packet and sends data through an egress port based upon a routing table entry. Routing tables can be manually configured or incrementally learned using discovery procedures in routing protocols. Such routing algorithms are either static or dynamic. Static routing is applicable when network topologies do not change significantly and route calculations can be performed offline. Typically, however, the dynamics of networks (Paxson, 1997) necessitates more sophisticated algorithms to allow routers to adapt to rapidly changing network conditions, where nodes and links can be congested, go down or come up at any instant. Routers hence store and forward packets with the goal to adaptively optimize path-finding between sources and destinations in dynamic networks. While Internet routers until recently serviced packets in a first-come-first-served “best-effort” manner, differentiated service qualities are now required to cater to application specifics, which necessitates that routers support new mechanisms such as admission control, per-flow queuing, resource reservation, and fair scheduling. These mechanisms require routers to use packet classification algorithms to distinguish and isolate traffic in different flows for suitable processing (Gupta and McKeown, 2001).

Router Architecture and Performance Routers incorporate a hierarchical packet forwarding logic in hardware plus computational capabilities to maintain routing tables. A router consists of a set of input ports, a backplane, and a set of output ports, including a CPU and memory to maintain a routing table for shortest path calculations. In a black box view, a router is a device that switches data packets arriving at an input port to an output port according to a routing criterion and after some queuing delay. From an inside view, data are shifted in traditional design from input ports ideally collision-free through an internal fabric to an output port, as shown in Fig. 2a.

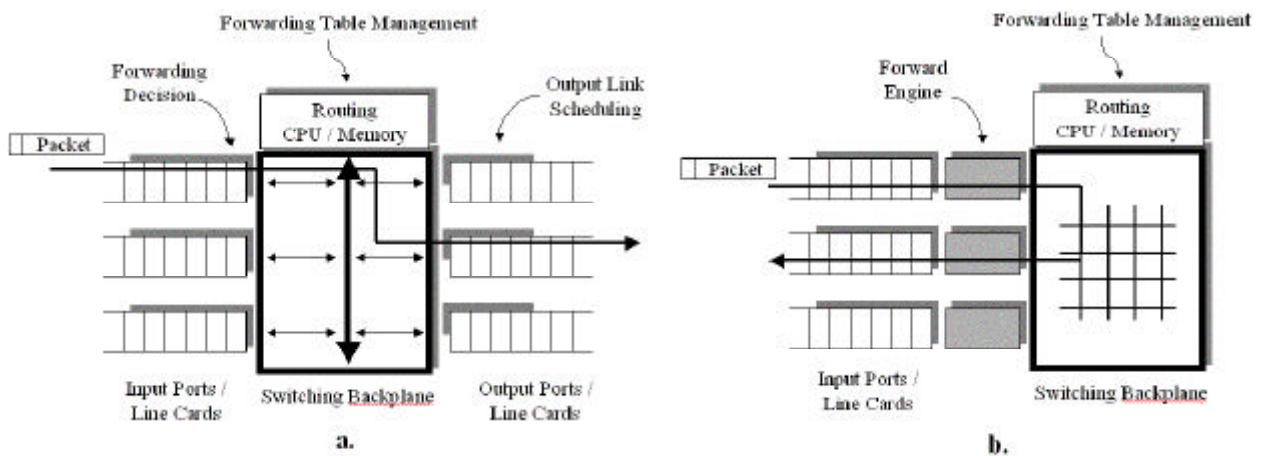


Figure 2. Traditional and Modern Router Design.

In high-performance systems, datapath functions such as forwarding decision making, backplane switching, and output link scheduling are typically implemented in special purpose hardware, while control functions such as routing table exchanges with neighbors or management are software-based. Instead of shared, congested backplanes, modern router design now incorporates switched backplanes, which results in a compact, parallel design with extremely high throughput for both unicast and multicast traffic. Fig 2b. depicts the use of more complex hardware on the main datapath and a crossbar design in the backplane to achieve more parallelism for high-speed routers (McKeown, 1997).

The main performance characteristics of a router are the number of ports or links, the throughput possible at each link, the maximum rate of packets that can be switched through, and the delay across the switch. Crucial for the performance of a router are fast table lookups and calculations of routes based on available network topology information (Keshav and Sharma, 1998). A current bare-bones router operates in software only without routing protocols and is simple to configure with built-in NAT. Low-end routers support around 200 Mbps full-duplex with Fast Ethernet and offer simple hardware-based operation, built-in NAT and firewall functionality, but lack fault tolerance. With an increasing number of channels transmitted on a single fiber, routers must scale port densities to handle these channels. Midrange routers add support for Virtual LANs (VLANs) and Virtual Private Networks (VPNs) with a broader range of WAN connectivity. High-end routers add component redundancy and resilience, scaling from 2.5 Gbps to 40 Gbps per slot with multi-terabit switching capabilities and several OC-192/STM-64 interfaces. A top-end ISP carrier router offers a theoretical capacity up to 92 Tbps with support for up to over 1000 OC-768c interfaces on the data plane and multimillion packet-per-second performance.

The advent of broadband access for end-users has raised expectations in IP reliability and performance. While earlier switches and routers offered single-purpose functionality, more recent products are convergence devices incorporating several functions in one box. Examples are combined firewall and VPN appliances (Cisco, 2004), wireless broadband routers with print server functionality, routers supporting VoIP and multimedia, or Layer 4+ content routers relaying packets based on URL semantics rather than IP-based semantics. As a result, there is great demand for rich-featured convergence devices used in SOHO and small business networks, and for Gigabit and Terabit (Singhal and Jain, 2002) electronic and optical routers used in backbones. These routers do not only forward packets at a rate of billions of packets per second, but must also provide quality guarantees for differentiated services with data, voice, or video, work with a wide variety of interface types, offer scalability in terms of port density and capacity, and offer backward compatibility with various packet formats and routing protocols.

How Switches work

A switch can be compared to a train station or airport dynamically interconnecting different travel pathways. Switches essentially interlink physical segments of a network and allow data to be exchanged between these segments. On a microscopic scale, the intrinsic design of a switch determines how effective this interlinking of communication paths is performed within each switch stage. On a macroscopic scale, the interlinking of switches in the overall network topology determines the aggregate performance of the network.

Switching functions vary depending on the operational needs of applications and may include higher layer switching functions. A typical switch operates at Layer 2 in the OSI model and TCP/IP stack. It checks address information in link and interface-level headers in hardware, and directs frames to the next hop on the path to some destination host based on data link layer addresses, for example Ethernet MAC addresses. Layer 2 switches are adequate for high-volume traffic generated between local devices, such as workstations and servers. A flat address space in broadcast domains, the possibility of broadcast storms, and the limited number of links supported by a Layer-2 switch led to the development of Layer-3 devices able to optimally route traffic in a hierarchy of nodes. Layer 3 switching is functionally identical to routing and typically added in large switches. Layer 4-7 switches deliver intelligent traffic and bandwidth management based on the application content of a session, not only based on network connections (Srinivasan et. al., 1998) and can provide load balancing, Denial-of-service attack protection, intelligent application scanning and Virtual Local Area Network (VLAN) configuration.

Switches are usually "plug-and-play", which simplifies installation. Switches operate either in store-and-forward or in cut-through mode. A store-and-forward switch accepts a frame on an input line, buffers it briefly in the input port, the fabric or output port, before forwarding it to the next processing stage ready to handle the entire frame. When frames arrive faster than they can be forwarded, buffers overflow and packet loss occurs. In contrast, cut-through switching is based on the idea to start forwarding a frame as soon as its destination header field has arrived, which is usually handled entirely in hardware. It boosts throughput at the possibility of forwarding bad frames, since the Cyclic Redundancy Check (CRC) cannot be checked before transmission. The switching process is executed in hardware at wire-speed with effectively zero latency.

Switches originally linked segments with multiple devices but with dropping switch prices it became normal to attach a single device to each port, which is known as "switched" Ethernet. No packet collisions are possible with only one device per port, which improves network performance by allowing devices to run in full-duplex to achieve maximum throughput.

Depending on the layout of the physical space, multiple smaller switches or a few very large switches may be chosen in the design of a network, which is a matter of manageability, cost, traffic characteristics, and reliability, among other reasons. Fig. 3 shows an unsecured sandbox network with switches connecting Ethernet cluster, which are again connected via the root switch R together to access a server farm and linked to a router connecting the Intranet with the Internet. By using switches A, B, and C, traffic in leaf subnetworks is localized and degradation in one subnet does not affect the other parts of the network.

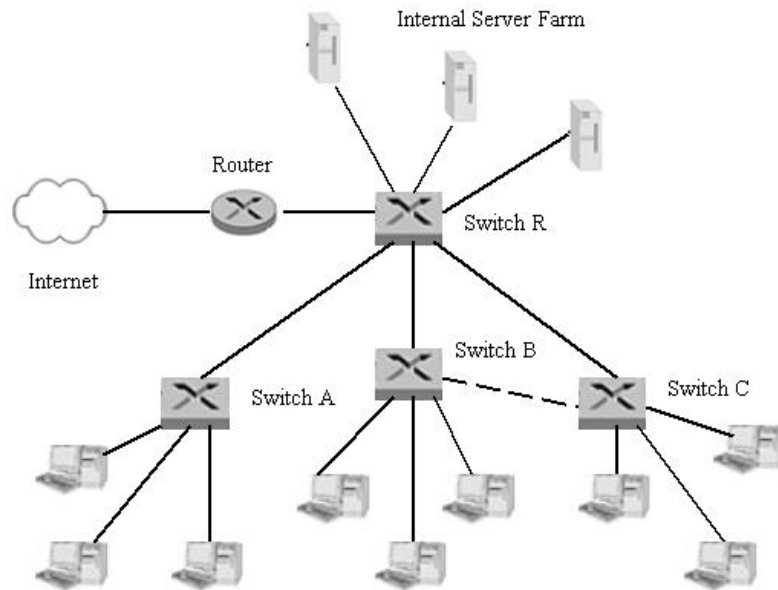


Figure 3. Switching in a LAN environment.

In the context of Ethernet LANs, switches are often called bridges, if used to connect multiple Ethernets, and run a Spanning Tree Protocol to form a loop-free forwarding topology to relay traffic between LAN segments. Looping occurs when devices are linked in way that creates a loop in the topology, for example by linking switches B and C directly (dashed link.) This means that packets could circulate between switch B and C, which can be avoided by forming a spanning tree if bridges are being used instead for B and C. The terms switch and bridge are often used synonymously, because both devices route on frame addresses at Layer 2. While bridges are used to interconnect two or more LANs and perform frame forwarding in software, switches are typically used to connect individual computers, which requires many more line cards with buffer space for frames arriving on switch ports.

Instead of a bridge or switch in a switched Ethernet, devices could also be interlinked via a hub in a "shared Ethernet" architecture. A hub repeats any incoming transmission on all of its other ports, regardless of packet content and addressing information. If two or more packets arrive at the hub at the same time, packets from different

network sections ("domains") collide. Unlike switches, which are capable of transmitting multiple packets full-duplex via different ports by examining source and destination addresses, hubs do not separate collision domains. The more transmissions are placed between disjoint senders and receivers, the larger the throughput advantage using a switch, which makes a hub an inefficient and outdated choice.

State-of-the-art switches come with a small number of ports in desktop form for Small Office / Home Office (SOHO) environments, or in modular, stackable form with hundreds of ports for wiring closets and enterprise deployment. They include support for the IEEE 802.3af standard for end devices, and Ethernet support for in-line power devices beyond IP phones and wireless LAN access points to connect more power-hungry devices such as IP-powered video surveillance cameras, security systems, or fire protection and motion detection devices.

Switch Architecture and Performance Switches are essentially the distributor nodes, across which various communication lines fan in from different sources and fan out to different locations. First generation switching networks were purely electromechanical and blocking, i.e., an existing connection prevents links between other input and output ports. They were invented to automate the interlinking of segments in communication exchange due to the fact that building networks by laying out point-to-point links between each pair of users would be cost prohibitive and an enormous waste of resources. Over decades switch technology has evolved from electromechanical through electronic to optical or photonic switching (Veeraraghavan et. al., 2001), and from using space-division switching in first-generation systems, through time-division circuit switching and multi-channel and multi-rate switching, to wavelength-division switching in photonic networks. The predominant relay mechanics in switches are consequently called Time-Division Multiplexing (TDM), Frequency-Division Multiplexing (FDM), and Wavelength-Division Multiplexing (WDM).

The three prevalent implementation paradigms for switches are based on shared memory, time-division or space-division fabrics. Shared-memory switches can be built simply from off-the-shelf components, and access shared memory space to carry out forwarding of packets from input port to output port. Time-division switches rearrange the order of incoming data on a multiplexed line according to a time schedule, whereas space-division switches, as common building blocks in high-performance networks, forward packets across a spatial fabric layout in a single stage or in multiple stages.

Crossbar switches are the most prominent example of a single-stage switch, where every input link is directly connected to every output port in a matrix with cross points that can be selectively enabled. With this simple

principle of connectivity, such switches have no internal blocking and low latency due to the minimal number of connecting points between arbitrary input and output links, and can be easily implemented in VLSI. However, the number of cross points required grows quadratically with the number of input and output ports, which increases the capacitive loading and signal distribution time, and hence slows transmission. In addition, crossbars have no fault tolerance and are difficult to expand and too costly. Small crossbar switches are often used to as components in more complex switch designs.

Multistage switches are composed of networks of smaller switches, which works well if each stage in the pipeline takes the same time. The interlinking of low-complexity switches in intermediate stages increases reliability, because input and output can still be connected in case of failure of a component switch. An example is the knockout switch, which is a crossbar with concentrators modeled on a knockout tournament to reduce space complexity. Multistage switching can be more effectively implemented by recirculation or with spatially separate switches. This class is represented by Clos networks, Banyan networks, and Batcher sorting networks (Jajszczyk, 2003). In Clos networks, a strictly non-blocking switching network can be constructed with fewer cross points than a crossbar of the same capacity. Each switch of one stage has an output feeding into each switch of the next stage. With a minimum number of cross points, a Clos switch becomes strictly non-blocking, otherwise it is blocking. While Clos networks have been primarily of theoretical interest, they are considered again today as a promising architecture for optical cross-connect systems and for high-capacity and high-performance backbone IP routers. Banyan switches are self-routing using a binary representation of output ports. The direction at each stage is specified by a bit corresponding to that stage. Half of the traffic is blocked in the worst case, and multiple planes of latter stages are used to buffer blocked traffic. Batcher networks are based on the idea that sorting and switching are essentially the same, except that outputs need to be spread after sorting to account for idle ports (Ahmadi et. al., 1989). Buffer placement and scheduling turn out to be a greater challenge in switch design than creating the fabric.

Switching throughput is measured either in bits-per-second (bps) or packets-per-second (pps), which is correlated with link speeds, support for half or full-duplex connections, support for jumbo frames, and non-blocking behavior. The goal is to achieve an overall throughput equal to the sum of the full theoretical bandwidth of each switch port. Current low-end switches offer connectivity to a small user base, typically with 5 to 16 ports at 10 to 100Mbps, and lack Spanning Tree protocol support, VLAN support or remote management. A current typical midrange switch offers support for 16 to 250 users with a bandwidth of 10 Gbps (non-blocking), a packet forwarding throughput of 2 Mpps and a packet buffer memory of 1 GB. Top-end midrange switches may include

redundant power supply and forwarding engines to greater availability, a greater number of Ethernet ports, remote access and configuration, and support for VLANs, Spanning Trees and Layer 3 switching. These switches are either unmanaged Layer-2 devices with auto-sensing, or managed Layer 2 or 3 devices, which can be used to inexpensively off-load routers segmenting a network to lay the groundwork for innovative IP-based services such as IP Telephony or video conferencing.

High performance switches (Newman et. al., 1997) offer the same features as mid-tier switches, plus Layer 4-7 switching, security features, modular gigabit connectivity for several hundred full-duplex ports, integrated IP address lookup (Waldvogel et. al. 1997), multicast support, an address database size of several thousand media access control (MAC) addresses per system, and a bandwidth of at least 10 Gbps. Multicast support inherently impacts blocking properties and the routing strategy to set up multicast connections via a switching network. At a packet size of 240 bytes, a router with a T1 line rate of 1.5 Mbps requires a lookup performance of 0.78 Kpps, or 5.21 Mpps at an OC-192 line rate of 10 Gbps. Current carrier-class switches can be scaled to 100 Terabit per second (Tbps) capacity, processing hundreds of millions of packets per second (pps) depending on the switch fabric, the line cards and their performance, and port speeds ranging from OC-48 to OC-192.

ROUTER AND SWITCH SECURITY

Various taxonomies on Internet infrastructure security have been proposed to categorize vulnerabilities and propose remedies. Generally, threats can be aimed at the devices themselves, or attack the services provided. At more detail, four major areas of protection can be considered to secure routers and switches from external intrusions, as well as violations from within a network: physical access, administrative control over a switch or router, intrusions against a switch or router, or intrusions through a switch or router (Microsoft, 2004). Alternatively, Pfleeger (2003) classifies threats to computing systems into four kinds: interceptions, interruptions, modifications or fabrications. Interception allows an attacker to read traffic and threatens confidentiality, interruption leads to lack of availability through service degradation or denial, modification threatens data integrity, and fabrication may cause traffic or service alteration and enable future attacks.

However, the majority of successful remote attacks on switches and routers can be traced to exploitation of a small number of security flaws. Chakrabarti and Manimaran (2000) identify four main categories of concrete concern for Internet infrastructure security: Domain Name System (DNS) hacking, routing table poisoning, packet mistreatment, and denial-of-service attacks.

The implications of these attacks are as diverse as the nature of their approach. IP packets are misrouted, confidential information is disclosed, deceptive or incorrect information is injected into the network through message modification, wrong address-name translations corrupt the routing process, network activities are disrupted through denial of service attacks, and network partitions or congestion can occur. Countermeasures can be preventive or reactive, and may be implemented at the network edge or core, with the latter generally being more costly due to deployment issues.

Best Practices in Securing Routers and Switches

Router and switch security concerns device treatment aspects such as physical access, administrative access and access control, availability and reliability, and traffic treatment aspects such as authentication among routers, confidential exchange of control and regular messages, as well as the integrity of messages transferred. Device mistreatment can be divided into physical and network intrusions. Entry-level switches, as plug-and-play devices, are relatively immune to network born violations, while mid-range switches and routers are typically most vulnerable. High-end routers and switches need various measures such as access control to hold off intrusion. Physical intrusions can be trivially prevented by limiting access to device, for example with a locked wiring closet. Network-based intrusions may be curtailed by adding firewalls to the network perimeter or to routers themselves and employing intrusion detection intelligence. Availability and reliability can be boosted through full duplication of equipment, but this solution adds considerable cost and necessitates automatic switchover. Traffic mistreatment is more complex due to the large spectrum of possible attacks. The next sections discuss notorious vulnerabilities (ISS, 2004) of routers and switches and suggest best practices for keeping attacks in reign.

Router Vulnerabilities and Attacks

Routing protocols to date have largely remained unprotected and open to attacks. Routing attacks can be targeted at the intra-domain and inter-domain level to disrupt correct routing in terms of the ability to reach the destination in compliance with a given forwarding policy. Generally, threats to routing protocols can be external or internal. External threats come from outside attackers who are non-participants in the protocol. Internal threats come from compromised protocol participants, which are referred to as Byzantine due their unpredictable nature. These attackers may corrupt, forge, or delay messages, or send conflicting messages (Puig et. al., 2004).

These threats can be exerted in various ways. Deliberate exposure occurs when routing information is revealed to unauthorized parties, for example by taking control over a router. By sniffing on control exchanges, an attacker can observe or record routing information. Traffic analysis is a more systematic method to detect patterns and vulnerabilities in router exchanges. Spoofing is an identify change problem. Falsification includes misadvertising network resources, tampering with protocol header fields, and misstating route attributes. Finally, attackers can interfere with exchanges between legitimate routers, or place excess burden on such routers.

Threat actions may result in usurpation, deception, disruption, or disclosure, listed with decreasing impact. Usurpation lets an attacker gain control over legitimate router functions. Deception happens when a forged routing message is accepted as authentic by a legitimate router and result in similar damages as usurpation. Disruption, as in a Denial-of-Service attack, causes temporary service outage and becomes more pressing when its frequency, duration or range increases. Disclosure allows attackers to monitor a link due to lack of confidentiality in routing exchanges. Table 2 summarizes this taxonomy (cf. Babir et. al., 2005).

THREAT SOURCE	CONCERNS	THREAT ACTIONS	OUTCOMES	CONSEQUENCES	
<ul style="list-style-type: none"> • Outside • Byzantine 	<ul style="list-style-type: none"> • Access Control • Authentication • Availability • Confidentiality • Data Integrity • Physical Access • Reliability 	<ul style="list-style-type: none"> • Deliberate Exposure • Sniffing • Traffic Analysis • Spoofing • Falsification • Interference • Overload 	<ul style="list-style-type: none"> • Usurpation • Deception • Disruption • Disclosure 	<u>Infrastructure:</u> <ul style="list-style-type: none"> ▪ Blackholing ▪ Churning ▪ Clog ▪ Congestion ▪ Instability ▪ Looping ▪ Overcontrol ▪ Partition 	<u>Hosts:</u> <ul style="list-style-type: none"> ▪ Cut ▪ Delay ▪ Eavesdrop ▪ Looping ▪ Starvation

Table 2. Summary of threats to routers and their implications.

The consequences are varied, and may damage the infrastructure of the whole network, or damage communication for a particular host or network. Blackholing occurs when one router is overburdened with redirected traffic. Churning happens when network forwarding patterns change rapidly and cause large variation in data delivery. Clogging describes the situation, when a router runs out of resources to handle excessive load, as caused by blackholing or congestion. Congestion is caused by overburdening a network portion with traffic. When routing becomes unstable, a global forwarding state is not reached. Looping may result in data never delivered. Overcontrol happens when protocol control overhead dominates actual message exchanges, and partitioning designates artificial breaks in the network topology such that routers do not communicate with each other when they could. Consequences of attacks on particular routers include cuts, delays, eavesdropping, looping and starvation.

When routers are cut off, they do not communicate when in fact they could. Through eavesdropping, rogue routers receive and see traffic when they should not.

Frequent exchanges of routing updates among neighboring routers in distance vector routing, and flooding of updates in link state protocols make intra-domain routing protocols vulnerable to attacks. Various strategies have been proposed to protect routing updates, including adding sequence and predecessor information to updates, or introducing authentication and cryptographic measures for router exchanges such as digitally signing updates. Implied drawbacks are increased traffic volume and more processing overhead at routers. Although BGP is the de facto standard for inter-domain routing, it does not ensure the integrity, freshness and authentication of messages or the authenticity of path attributes and permits forging of path vectors, with the consequences that packets often get misrouted and BGP operation can be compromised. Butler et. al. (2004) present a comprehensive overview on the various approaches to secure BGP, which include encryption of session information and message attributes, introduction of a public key mechanism into the routing infrastructure, route validation, using certificates for authenticating and authorizing network entities, and using a routing registry and building various forms of cryptography into BGP communication.

From a more concrete and practical perspective, at the administrative level a remote attacker could bypass *access control lists (ACLs)* in a router due to configuration errors. In an otherwise properly configured ACL, this problem could allow a remote attacker to connect through the switch onto the 'protected' side. Edge routers with an integrated DHCP server can be susceptible to a *BOOTP denial of service attack*. When a specially-crafted BOOTP packet is sent to the router, a remote attacker could obtain sensitive information using a DHCP reply or cause the device to crash. An ADSL Router Integrated Switch can also be vulnerable to *cross-site scripting*, when a remote attacker creates a malicious URL link containing embedded code, which would be executed in the victim's Web browser within the security context of the hosting site, once the link is clicked. Since any device in a network could send Address Resolution Protocol (ARP) messages, an attacker on a LAN can easily spoof the gateway in a *gratuitous ARP attack*. Using *TCP Sequence Prediction*, an attacker could hijack another user's session or perform an IP spoofing attack to manipulate connections to the TCP services or modify the configuration of the router.

Router Configuration and Deployment Practices Router security includes firmware and software patches and updates, administrative access, auditing and logging, intrusion detection, routing protocol configuration, and additional services. The National Security Agency (NSA, 2003) makes various recommendations to secure routers against intrusions.

First and foremost, most devices have back-door access and should therefore be physically locked up. A written router security policy should be maintained to define management and logging practices, in particular, who can access, configure and update a router. Updates should be always tested before deploying them in a production environment. Administrative access must be restricted to specific locations and interfaces and use encryption to prevent hijacking, which entails strong password policies, using an administration access control system, controlling physical access, shutting down web-based configuration, and disabling unused interfaces. Passwords should be encrypted and configured securely for console and virtual terminal lines, and SSH should be adopted for remote administration. Consistent deployment of the same services across large scale networks can be provided by including the ability to rapidly configure routers for specific services with macro configuration templates comprising series of command line configuration commands, however, these templates and master versions of configuration files should be well commented, stored offline and kept in sync with running configurations to diagnose attacks and enable fast recovery. Router security should be audited regularly, especially after any reconfiguration.

Protocol-level vulnerabilities are often the target of denial-of-service attacks, for example by flooding the network, and can be countered, among other measures (SANS, 2004) by using ingress and egress filtering, screening ICMP traffic, and limiting broadcast traffic or other unnecessary traffic. Unneeded services in the router such as BOOTP, Finger, SNMP, source routing, trace route, or protocols such as the Cisco Discovery Protocol (CDP) should be disabled. If SNMP is required for a network, it should be configured with ACLs and hard-to-guess community strings.

Loop-back addresses, source addresses from any internal network, and packets having the same source and destination address should be blocked in order to prevent TCP sequence number guessing of packets with obviously fake or reserved addresses from untrusted networks. Blocking illegal addresses not only prevents attackers from hijacking a router but also helps to diagnose poorly configured internal networks and hosts. ICMP echo and redirect may be used by remote attackers to scan a network and manipulate routing behavior, and should also be blocked. Broadcast packets, as generated by DHCP and BOOTP, should not be used on external interfaces or cross border routers and should be blocked. On border routers, only internal addresses should be allowed to enter the router from

internal interfaces, and only traffic headed to internal addresses should be permitted from external interfaces.

Multicast packets should only be allowed in networks supporting IP multicast.

Access control lists prevent certain traffic to enter or exit a network and should be implemented to allow only those protocols, ports and IP addresses required by users and services and explicitly disallow any others. Previous access lists should be cleared out when implementing a new list and traffic address restrictions should be enforced at all times. Logging should be enabled for all denied traffic, and unusual traffic patterns should be audited, with logs being centrally stored and secured. Errors and blocked packets should be logged to an internal, trusted SYSLOG host, and SYSLOG traffic from untrusted networks should be blocked. Logs should include event timing and at least two Network Time Protocol (NTP) servers should be configured to ensure availability and accuracy of timing information to enable precise tracking of network attacks.

Standard routers perform typically little filtering of data and therefore inherently may allow forwarding of compromised traffic. Stateful packet inspection (SPI) is a network layer mechanism to examine packet contents up through the application layer rather than merely looking at packet header information, and compiling connection state information in a table. Filtering decisions are hence established based on a context of prior packets having passed through the router and allow to establish dynamic filtering criteria beyond static administrator-defined rules. SPI also thwarts port scanning by closing off ports until a connection to a specific port is requested. Packet filtering is fairly effective and transparent to users, but difficult to configure and susceptible to IP spoofing.

Routers become application gateways when applying security policies to specific applications. Although a firewall is only a first line of defense against network intrusion, modern routers incorporate firewall features and contain numerous mechanisms in hardware and software to strengthen their lines of defense against attacks, with the idea to provide plug-and-play security. A variety of router security features handle Denial-of-Service attacks via detection and logging, dropped packet logging, time-based usage control, URL filtering, deep packet filtering, trusted user identity management.

Router manufactures have therefore integrated such services increasingly in their hardware as a more efficient layer of defense for the network infrastructure. If a router acts as circuit-level gateway by applying security mechanisms to a transport level connection, packets could flow between hosts without further checking. For greater security, information exchanges should be encrypted. Current routers typically offer 56-bit to 168-bit encryption with DES, Triple DES, RC4, MD5, or SHA-1 encryption algorithms (Stallings, 2000).

As edge devices in local networks, routers can also shield private networks from outside access by mapping external, global addresses to internal addresses. Routers hence often include support for NAT, PPTP and L2TP, DHCP, DMZ, IPsec, and VRRP.

NAT (Network Address Translation) is a popular but also controversial (Phifer 2000) mechanism to counter IP address depletion, and makes a network appear to the outside world with a single IP address by hiding its internals. NAT serves the three main purposes of providing a type of firewall by hiding internal IP addresses, enabling a company to use more internal IP addresses, and combining multiple broadband connections into a single Internet connection. A translation table is maintained to map application-specific outside requests to internal IP addresses and associated port numbers. In static NAT, a private IP address is mapped to a public IP address, where the public address is always the same. This allows an internal host, such as a Web server, to have an unregistered, private IP address and still be reachable over the Internet. With dynamic NAT, a private IP address is mapped to a public IP address drawing from a pool of registered, public IP addresses. Dynamic NAT helps to secure a network as it masks the internal configuration of a private network and makes it difficult for someone outside the network to monitor individual usage patterns. A router using the Point-to-Point Tunneling Protocol (PPTP) or the Layer 2 Tunneling Protocol (L2TP) is able to establish secure, tunneled connections through an open network as a baseline service to establish VPN connections.

The Dynamic Host Configuration Protocol (DHCP, RFC 2131) enables routers to automatically obtain a different IP address when reconnecting mobile hosts with intermittent connectivity. It typically operates in conjunction with NAT and supports hybrid use of static and dynamic IP addresses. Due to the dynamic nature of the address assignment, attackers will regularly have to rescan a network to obtain current addresses.

A DMZ (DeMilitarized Zone) inserts a neutral zone between a private network and public networks. Routers supporting setup of a DMZ allow users from the public network to access only the DMZ host. This "exposed host" provides limited access to company data, and can operate as proxy server for inside hosts to request data from the public network, but is not able to initiate sessions back into the private network.

IPsec operates with two different encryption modes, transport and tunnel, to introduce tighter security at the network layer. Transport mode encrypts only the data portion (payload) of each packet, but leaves the header untouched. In the more secure tunnel mode a router encrypts both the header and the payload. On the receiving side, an IPsec-compliant device is needed to decrypt each packet. For IPsec to work, the sending and receiving devices must share a public key through the Internet Security Association and Key Management Protocol/Oakley

(ISAKMP/Oakley) protocol, which allows the receiver to obtain a public key and authenticate the sender using digital certificates.

Finally, the Virtual Router Redundancy Protocol (VRRP) runs on link in a fault-tolerant arrangement between one device operating as master and another device operating in standby mode, so each router knows whether the "alter ego" is alive and is able to react in failover mode if necessary.

Securing Ad Hoc routing presents difficulties not existent in traditional networks. Neither centrally administrated secure routers nor strict policies exist in ad hoc networks, where nodes communicate without regulation through centralized access points. If routers go mobile, a new class of security mechanisms needs to be introduced. So far, security provisions for wireless networks entail support for Wi-Fi Protected Access (WPA), Pre-Shared Key (PSK), MAC authentication, SSID hiding, and Wired Equivalent Privacy (WEP) (Miller, 2001). However, software-only defense solutions are not sufficient in the face of distributed denial-of-service attacks, spoofing and other attacks bogging down the network. For example, while the WEP protocol was designed to provide confidentiality, access control and data integrity, it fails to fulfill these goals due to keystream reuse and message authentication flaws (Borisov et. al., 2001).

Greater service complexities in routers, such as differentiated quality or voice telephony support, create new loopholes and necessitate constant updating of security measures. Mechanisms such as IP traceback, while still a topic of research (Savage et. al., 2000), could be useful in the future in pinpointing attackers. Newer development in Security Information Management (SIM) tools and Intrusion Detection Systems (IDS) may help to correlate feeds from various audits and logs to recognize and diagnose attacks, using database and visualization tools with "dashboard" functionality for more effective monitoring and securing of networks.

Switch Vulnerabilities and Attacks

Switch operation can be compromised at various service layers. An attacker, who gains administrative access to a vulnerable switch in a *password attack*, could take complete administrative control of the device. Proper password management, in particular the avoidance of default passwords, and restricted network and console access to these devices is therefore critical. In a *MAC address sniffing attack*, an attacker connected to a vulnerable switch could use a packet sniffing tool to obtain the MAC addresses of connected systems and can cause traffic to be broadcast to all systems connected to the switch. In a *frame injecting attack*, an attacker could inject 802.1q frames into a VLAN and transfer data to unauthorized network segments. In a *Spanning Tree Protocol (STP) attack*, an attacker

connected to two different switches could send Bridge Protocol Data Unit (BPDU) messages to gain root switch privileges over the tree and intercept transmissions among subtrees. In a *frame forwarding attack*, a switch receiving an 802.1x packet frame on a Spanning Tree Protocol (STP) blocked port does not drop the packet but instead forwards it in the VLAN. This can cause an 802.1x frames network storm, which slows the performance of the network. By sending a malformed Internet Control Message Protocol (ICMP) packet to a switch or to a computer behind it, a remote attacker can cause the switch to crash in a *Denial-of-Service attack*.

A remote attacker can also cause a denial of service by sending a flood of TCP SYN packets to the vulnerable device. In both cases, the power to the switch must be completely shut off and restored for the switch to regain normal functionality. A content switch can be vulnerable to a *denial of service attack of HyperText Transfer Protocol (HTTP) requests* to the Web management interface. A remote attacker could send a malformed HTTP POST request or eXtended Markup Language (XML) data to the Web management interface to cause the device to reboot, and hence deny services to legitimate users. A content service switch can also be vulnerable to a *UDP Denial-of-Service attack* when malformed UDP packets are sent to the management port (e.g., UDP port 5002) and a remote attacker can cause the device to reload. Finally, in a *web-based switch management vulnerability* a remote attacker could access a switch via its administrative web interface, and gain access to sensitive information without authentication by bookmarking the Web management URL.

Switch Configuration and Deployment Practices The National Security Agency (NSA, 2004) recommends for threat defense that switches should be covered in an institution's network security policy, including operating systems, port management, passwords, network services, spanning tree protocol, logging and debugging, access control lists, authentication, authorization and accounting. First and foremost, physical access to a switch must be restricted, and switches should be securely configured so that sessions automatically time out and only necessary network services are enabled. The switch configuration file should be well commented and kept offline with limited access permissions. Software on switches should be regularly patched to run the latest stable release. Password access could be configured at different privilege levels, and SSH should be used instead of TELNET to administer a switch.

Securing a switch configuration entail various aspects, including the proper configuration of services and VLANs; the disabling of unused ports, the use of an access control system, and employing encryption, in particular on wireless links. Traditionally encryption is not implemented in a switch, but it ensures that intercepted packets are useless when an intruder sniffs on the same switched segment or where the switch is compromised. Switch management should be either "out-of-band" or on a separate VLAN for in-band management. Access control lists

(ACLs) allow controlling of inter-VLAN traffic between IP subnets by restricting the flow of traffic between different segments of the network, and must be configured correctly. Typically, a simple static packet filter is used, in contrast to functions such as stateful packet inspection or application-layer proxying performed by dedicated firewall devices.

Port security should disable unused ports and limit access based on MAC addresses. The spanning tree protocol and auto-trunking should be disabled in loop-free topologies. A static VLAN configuration should be used when possible, the number of VLANs that can be transported over a trunk should be limited and trunk ports should be assigned to an otherwise unused native VLAN number. In addition, 802.1x and the Extensible Authentication Protocol (EAP) can be used to leverage edge-security by authenticating a machine before it is allowed to access the network. Finally, logging should be enabled and logs should be sent to a dedicated, secure log server. Logs should be reviewed for incidents and archived in accordance with the general security policy. While these NSA recommendations are based on Cisco IOS products, they apply to switches in general.

Current software-based enhancement services of high-end switches include the ability to prevent MAC address flooding attacks by locking down ports, prevent attacks from false DHCP servers, and limit network access through port-level ACLs. Current switches also include authentication capabilities in standard IEEE 802.1x to be able to attribute authenticated traffic to a specific virtual LAN or add Quality of Service features as well as prevent denial-of-service attacks by dynamically inspecting Address Resolution Protocol traffic and binding MAC and IP addresses to specific ports. Layer 2 and 3 switches also typically support standards such as IEEE 802.1Q (Static VLAN groups), port-based VLAN (any one port can belong to different VLAN groups), IEEE 802.1p Class of Service (CoS), IEEE 802.1D Spanning Tree Protocol, and manual Port Trunking as per IEEE802.3ad. In addition, IGMP Snooping, Port Mirroring, RFC 1157 (SNMP), RFC 2819 (RMON), RFC 1213 MIB II, RFC 1643 (Ethernet Managed Objects), RFC 1493 (Bridge Managed Objects), RFC 951 (BOOTP), RFC 2998 (Differentiated Services), and RFC 2865 (RADIUS) are often supported.

Case Studies

Deploying a switched network entails choosing which class of routers or switches is needed, what functionality is required, and where those devices should be placed in relation to each other. Decisions on current connectivity and future growth, communication between devices, the required bandwidth and acceptable latency, and what VLAN architecture may be required are all interdependent parameters affecting overall security. Among the many possible

organizational structures, SOHO and multilevel switching architectures are most common. A simple, secured SOHO setup, shown in Fig. 4, may include a modem, usually a cable or DSL modem, a router and a switch, which are often combined in a low-cost broadband router acting as an Ethernet switch with a firewall function.




Figure 4. Simple secure SOHO Network.

In an example for a basic secure multilevel design in Fig. 5, routers separate the public Internet from an intranet through a perimeter network. The first segment uses a border router facing the Internet with initial firewall capability, switched to a perimeter firewall that in turn connects to a web server cluster in the perimeter network via a switch. Another switch connects the server farm to an internal firewall, which is switched to backend PCs and internal servers. Aside from the border switch, which is located in a more vulnerable zone, all other switches could be separate VLANs on the same switch, which could be one larger unit or consist of multiple smaller devices. Scenarios far more complex than presented are conceivable, but composition and the idea of isolation of critical components are principally the same.

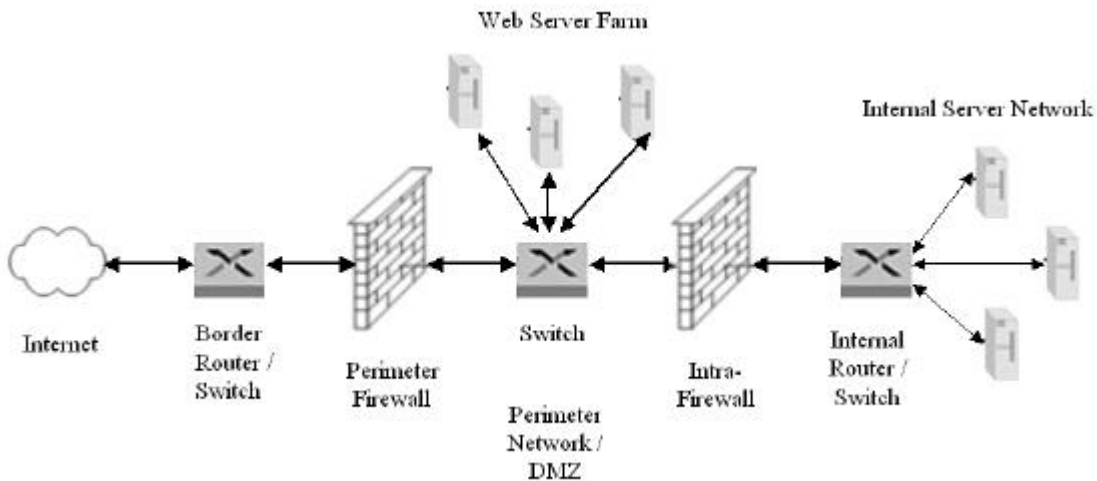


Figure 5. Secure multi-level switched network.

CONCLUSION

The phenomenal growth in digital communication traffic has transformed the Internet from a research testbed into both a daily convenience and a mission-critical service. As critical network infrastructure components, routers and switches are key devices in controlling network traffic and linking together computer networks at greater scale. Routers, from small to very large scale, perform primarily path finding optimization based on interior routing protocols such as RIP and OSPF, or exterior policy-centric protocols such as BGP. In security terms, routers can be used to segregate intranets from internetworks through packet filtering and address translation. While attacks on access network routers may only compromise the local network, vulnerabilities in backbone routers could create large-scale damage to a communication infrastructure. Switches can radically improve network throughput by reducing packet collisions in the forwarding process. Switches segment networks according to the operational and security policies of an organization. Physical and administrative loopholes can make switches vulnerable to attacks that may impact the greater network around the switch. Tight physical, administrative, and protocol-centric security measures need hence to be enacted upon routers and switches alike to minimize network downtime and to fend off attacks on the integrity of distributed communication.

At the dawn of the next generation of massively ubiquitous communication tools and interactive applications exchanging high-volume traffic, such as video conferencing, video-on-demand, multiplayer simulations and distance learning, switching and routing technologies may once again have to change to adapt to the next wave of innovation in communication, and the implied greater needs in security. In this section we discussed the essentials of router and switch mechanics, their vulnerabilities, and some current best practices in securing these devices, with main focus on wired networks. Issues such as content switching and routing, Quality-of-Service provision in its interplay with security, or wireless security need to be addressed separately. The ultimate vision is to equip routers and switches with the inherent adaptive intelligence to make networks self-organizing, self-defending, and self-healing against any conceivable attack at the hardware or software level, while offering uninterrupted high-performance service.

GLOSSARY

ACL – Access Control List, a predefined set of rules that configure router packet filtering capabilities for routed network protocols based on an organization's security policy.

ARP – Address Resolution Protocol, a distributed protocol to convert an IP address into a physical (Data Link Control) address such as an Ethernet address. A host broadcasts an ARP request with an IP address to a TCP/IP network and the host with this IP address replies with its physical address. With the Reverse ARP protocol (RARP) a host broadcasts its physical address to receive an RARP reply with the host's IP address.

Autotrunking – Multiple ports are configured to be one the same trunk for data transferring from one network switch to another network switch with the same trunking features and settings, in order to increase bandwidth.

Crossbar – Simple non-blocking switch with inputs and outputs arranged as rows and columns in a matrix, where data are switched by activating cross points in the matrix.

DHCP – Dynamic Host Configuration Protocol, a protocol for assigning IP addresses dynamically from a pre-configured pool of IP addresses to devices on a network. With dynamic addressing, a device can have a different IP address every time it connects to the network. Dynamic addressing simplifies network administration because the software keeps track of IP addresses rather than requiring an administrator to manage the task.

DMZ – DeMilitarized Zone, a computer or small subnetwork placed between a trusted internal network, such as a corporate private LAN, and an untrusted external network, such as the public Internet.

DoS – Denial-Of-Service, a type of attack on a network to bring the network to its knees by transmission of useless traffic. Many DoS attacks, such as the Ping of Death and Teardrop attacks, exploit limitations in the TCP/IP protocols.

Firewall – A system, implemented in both hardware and/or software, and designed to prevent unauthorized access to or from a private network as a first line of defense. A firewall examines each incoming or leaving message and blocks those that do not meet specific pre-configured rules.

Gateway – Node on a network serving as entrance to another network. This term is often associated with both a router and a switch but modern uses of the term indicate higher-layer functionality. For example, an application-specific type of firewall is also called an application-level gateway.

IPsec – IP Security, a set of protocols developed by the IETF to support secure exchange of packets at the IP layer and implement VPNs.

Ingress Filtering – Examination of packets going through a router from the outside network to the inside. The opposite is **egress filtering**.

L2TP – Layer 2 Tunneling Protocol, an extension to the PPP protocol that enables ISPs to operate VPNs.

Line Card – Contains physical layer components to interface external data link to the switch fabric.

NAT – Network Address Translation, an Internet protocol standard that enables a local-area network (LAN) to use one set of IP addresses for internal traffic and a second set of addresses for external traffic. A NAT box located where the LAN meets the Internet makes all necessary IP address translations.

Network Processor – Runs the routing protocol, computes the routing tables that are copied into each forwarding engine, handles network management, and processes special handling for unusual packets.

OC – Optical Carrier, a frame format and speed metric in the SONET multiplexing hierarchy used to assess high-performance network throughput. OC-1 offers a data rate 51.84 Mbps.

OSI Model – The Open Systems Interconnection Reference Model defines a 7-Layer architecture of communication functions for cooperating network devices.

Pps – Packets per second, a metric to measure switch and router performance.

PPTP – Point-to-Point Tunneling Protocol, used to ensure that messages transmitted from one VPN node to another are secure. With PPTP, users can dial in to their corporate network via the Internet.

Router – An internetworking layer 3 device connecting two or more networks running the same routing protocol. It uses the Internet Protocol to address devices and maintains routing tables to determine the next hop on the optimal path to forward packets based on header information and routing policies.

Spanning Tree Protocol – Link management protocol in the IEEE 802.1 standard for bridges, which provides path redundancy in a network with several possible paths between hosts while preventing undesirable loops.

SPI – Stateful Packet Inspection is a firewall architecture that works at the network layer. Unlike static packet filtering, which examines a packet based on the information in its header, stateful inspection tracks each connection traversing all interfaces of the firewall and makes sure they are valid. Also referred to as dynamic packet filtering.

Spoofing – A variety of ways in which hardware and software can be fooled. In IP spoofing a message is altered in its header to appear as if it came from the authorized IP address of a trusted host.

Switch Fabric – Interconnects the various components of router and offers higher aggregate capacity than the more conventional backplane bus.

VLAN – Virtual LAN technology is used to create logically separate LANs on the same physical switch and apply

access control based on security rules to devices in logical network segments. A VLAN is a software-configured network of computers which behave as if they are connected to the same wire even though they may actually be physically located on different segments of a LAN.

VPN – Virtual Private Network, a logical topology of computers across the public Internet.

Tunneling – Encapsulation of packets to enable one network to send its data via another network's connections.

Packets from one network protocol are embedded within TCP/IP packets carried by through public Internet.

REFERENCES

- H. Ahmadi, W.E. Denzel (1989).** A survey of modern high-performance switching techniques, *IEEE Journal of Selected Areas in Communications*, Vol. 7, No. 7, pp. 1091-103.
- A. Barbir, S. Murphy, Y. Yang (2004).** Generic Threats to Routing Protocols, Internet-Draft draft-ietf-rpsec-routing-threats-07, Available at www.ietf.org/internet-drafts/draft-ietf-rpsec-routing-threats-07.txt, Date of access (DOA): Mar 30, 2005.
- N. Borisov, I. Goldberg, and D. Wagner (2001).** Intercepting mobile communications: The insecurity of 802.11. Proc. MOBICOM, Rome, Italy.
- K. Butler, T. Farley and P. McDaniel (2004).** A Survey of BGP Security Issues and Solutions, Tech. Report TD-5UGJ33, *AT&T Labs - Research*, Florham Park, NJ, June 2004.
- A. Chakrabarti, G. Manimaran (2000).** Internet Infrastructure Security - A Taxonomy. *IEEE Network*, Vol. 16, No. 6, pp. 13-21.
- Cisco Systems (2004).** Security and VPN. Available at www.cisco.com/cgi-bin/Support/browse/index.pl?i=Products&f=753&viewall=true , DOA: Nov 27, 2004.
- P. Gupta and N. McKeown (2001).** Algorithms for Packet Classification, *IEEE Network*, Vol. 15, No. 2, pp.24-32.
- C. Huitema (2000).** Routing in the Internet. Prentice Hall.
- Internet Security Systems (2004).** Security Alerts and Advisories. Available at xforce.iss.net , DOA: Nov 27, 2004.
- A. Jajszczyk (2003).** Nonblocking, repackable, and rearrangeable Clos networks: fifty years of the theory evolution. *IEEE Communications Magazine*, Vol. 41, No. 10, pp. 28-33.
- S. Keshav (1997).** An Engineering Approach to Computer Networking, Addison-Wesley.
- S. Keshav and R. Sharma (1998).** Issues and Trends in Router Design. *IEEE Communications Magazine*, Vol.36, No.5, pp. 144-51.
- S. K. Miller (2001).** Facing the Challenge of Wireless Security. *IEEE Computer*, Vol.34, No. 7, pp. 16-8.
- N. McKeown (1997).** A Fast Switched Backplane for a Gigabit Switched Router, *Business Communications Review*, Dec. 1997.
- Microsoft Technet (2004).** Router and Switch design. Available at www.microsoft.com/technet/security/guidance/secmod40.aspx, DOA: Dec 10, 2004.
- National Security Agency (2003).** Router Security Configuration Guide. Available at www.nsa.gov/snac/downloads_cisco.cfm?MenuID=scg10.3.1 , DOA: Nov 27, 2004.

- National Security Agency (2004).** Switch Security Configuration Guide. Available at www.nsa.gov/snac/downloads_switches.cfm?MenuID=scg10.3.1 , DOA: Nov 27, 2004.
- P. Newman, G. Minshall, T. Lyon, L. Huston (1997).** IP switching and gigabit routers, IEEE Communications Magazine, Vol. 35, No. 1, pp. 64-9.
- V. Paxson (1997).** End-to-End Routing Behavior in the Internet. IEEE/ACM Transactions on Networking, Vol. 5, No. 5, pp. 601-15.
- R. Perlman (1992).** Interconnections: Bridges and Routers. Addison-Wesley.
- C. P. Pfleeger and S. L. Pfleeger (2003).** Security in Computing, 3rd ed. Prentice-Hall.
- L. Phifer (2000).** The Trouble with NAT. The Internet Protocol Journal, Available at www.cisco.com/warp/public/759/ipj_3-4/ipj_3-4_nat.html , DOA: Nov 30, 2004.
- JJ. Puig, M. Achemlal, E. Jones, and D. McPherson (2005).** Generic Security Requirements for Routing Protocols, Internet-Draft draft-ietf-rpsec-generic-requirements-01, January 2005, Available at www.ietf.org/internet-drafts/draft-ietf-rpsec-generic-requirements-01.txt, DOA: Mar 30, 2005.
- S. Savage, D. Wetherall, A. Karlin, and T. Anderson (2000).** Practical Network Support For IP Traceback, Proc. ACM SIGCOMM, Stockholm, Sweden, pp. 295-306.
- SANS Institute (2004).** The Twenty Most Critical Internet Security Vulnerabilities. Available at www.sans.org/top20/, DOA: Dec 10, 2004.
- S. Shenker (1995).** Fundamental Design Issues for the Future Internet. IEEE Journal of Selected Areas in Communication, Vol. 13, No. 7, pp. 1176-88.
- A. Singhal and R. Jain (2002).** Terabit Switching: A Survey of Techniques and Current Products. Computer Communications, Vol. 25, No. 6, pp. 547-556.
- V. Srinivasan, G. Varghese, S. Suri, and M. Waldvogel (1998).** Fast and scalable layer four switching. Proc. ACM SIGCOMM, pp. 191-202.
- W. Stallings (2000).** Network Security Essentials - Application and Standards. Prentice Hall.
- M. Veeraraghavan, M. Karol, R. Grobler, R. Karri, and T. Moors (2001).** Architectures and Protocols that enable new Applications on Optical Networks. IEEE Communications Mag., Vol. 39, No. 3, pp. 118-27.
- M. Waldvogel, G. Varghese, J. Turner, B. Plattner (1997).** Scalable high speed IP routing lookups, Proc. ACM SIGCOMM, Cannes, France, pp. 25-36.