WHITE PAPER

CISCO CATALYST INTEGRATED SECURITY—ENABLING THE SELF-DEFENDING NETWORK

Network security has become the primary concern of most enterprise and commercial network administrators. Whereas the greatest concern in the past was "physical security," the almost-weekly release of new Internet-based viruses, worms, and attack tools has vastly increased the risk to the very fabric of business productivity. Network security was once viewed merely as a firewall between the network and the Internet. Now, network managers are looking to integrate security end to end throughout the network, where security policies and capabilities permeate every place and device in the network. With the network becoming a critical integration point for IP-enabled applications and communications systems, such as Voice over IP, security is more important than ever. The place to start is the switching infrastructure.

Cisco Systems® has innovated and integrated industry-leading security capabilities into the Cisco® Catalyst® Intelligent Switching portfolio. Today, enterprises are buying LAN switches with the expectation that it will offer built-in capabilities to identity users. The campus switching infrastructure represents the first line of defense: it is here that networked devices attach to the network; therefore, it is the first point of entry of an attack into the network.

Cisco Catalyst Integrated Security offers three systems that integrate with the Cisco Self-Defending Network architecture.

- Trust and identity—Gives the network manager control over who and what attaches to the network and what kinds of policies are applied; these capabilities include 802.1x and associated identity-based networking services as well as Network Admission Control (NAC)
- Threat defense—Allows the network to prevent and mitigate attacks if they are launched at other users on the network or at the network devices themselves; these capabilities include control-plane rate limiting, access control lists (ACLs), man-in-the-middle attack mitigation, and firewalling
- Secure connectivity—Provides privacy for communication, either within the campus network or between sites connected over the Internet

This paper explores Cisco Catalyst Integrated Security and discusses the capabilities that Cisco has embedded in the campus, branch, and data-center infrastructure.

TODAY'S SECURITY CHALLENGES

In the past, the network was designed to be an open utility. Network security, although always a consideration, was rarely of primary concern and was often limited merely to physical security (for example, preventing someone from getting access to the switch). The first driver for change came as enterprises began accessing the Internet through their networks. These connections opened a door to the outside world, allowing limitless possibilities for network users and limitless opportunities for intruders, hackers, data thieves, or anyone with malicious intent. Firewalls, which secure the perimeter of the network, and policies, which govern access to resources, protect the interior while still maintaining reasonable access and mobility. With the advent of increasingly sophisticated worms and viruses, those policies and devices have proven insufficient. Figure 1 shows how attacks have changed over the years.

Figure 1
Understanding Security Attacks—Past, Present, and Future

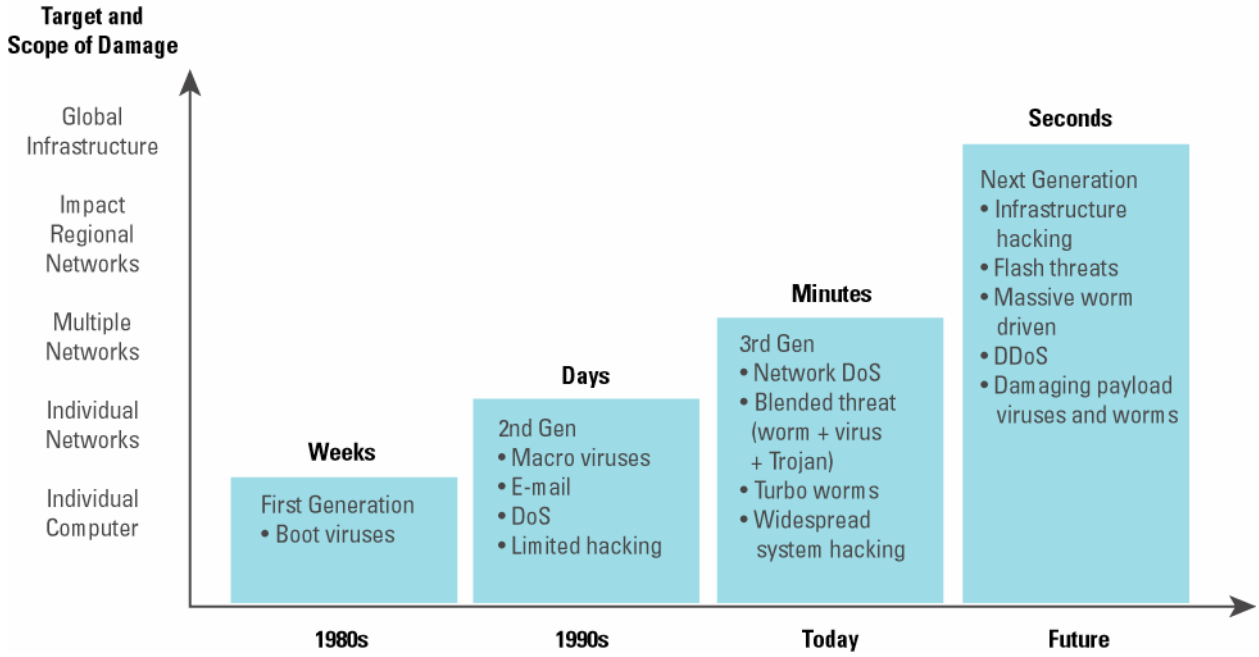

The speed and unpredictability of these attacks require a combination of products and technologies, consisting of both proactive and reactive elements, in an overall system. This system must first function at multiple layers in the network, from the teleworker to the campus to the data center. These layers must be able to address sophisticated, blended threats and defend against multiple avenues of attack. Secondly, the system must be automated such that the network can authenticate users and address “zero hour,” when the attack is first launched. Finally, the system must be fully integrated into all aspects of the network, so that the network manager can facilitate a coordinated response to the attack.

The switching infrastructure must be a key part of the overall security strategy of an enterprise. Typically, a “security operations” manager is most concerned with the overall security of the network. However, the network operations manager has a greater stake now in security and is often concerned about threats to the devices themselves and attacks that can be launched from the end stations connected to the network. Because these devices affect the switches, routers, and end stations under network operations control, now security is a consideration in every aspect of the network. Enterprise LAN managers expect the device to have inherent security capabilities to protect the switches themselves and the users attached to them. As shown in Figure 2, a self-defending network consists of three main systems.

Figure 2
Self-Defending Network Framework

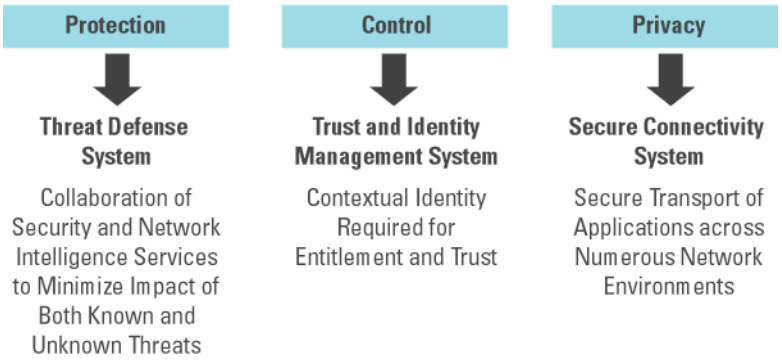

So what do these systems protect against? Most people are familiar with the viruses that they receive through e-mail, and the evening news is certainly full of stories of a company significantly disrupted by an attack, but many security threats are much more insidious.

- The tailgater—This individual penetrates the physical security of the building by simply walking in behind an honest employee. When inside, the person can attach to the network, download confidential information, or launch attacks.
- The infected laptop—The prevalence of laptop computers, and the variety of networks to which they can connect, allow for rapid infection. A laptop taken from a public network to the corporate network might inadvertently spread viruses or worms. When the worm spreads, network productivity can be severely impacted.
- The unauthorized device—Many users may connect devices that open up security holes in the network. The most obvious is the unsecured wireless access point, which can allow anyone with a wireless device in range access to the corporate network.
- The mass infection—Denial-of-service (DoS) attacks flood the network with superfluous traffic such that the networking devices cannot respond to legitimate network requests and traffic.
- The disgruntled employee (also known as the man in the middle)—In this attack a malicious user activates publicly available software to "spy" on other employees' data, including IP phone calls, passwords, etc. This includes Dynamic Host Configuration Protocol (DHCP) spoofing, IP spoofing, and Address Resolution Protocol (ARP) poisoning attacks.

The security mechanisms embedded within the Cisco Catalyst hardware and software provide tools and protection against these types of attacks so that an attack can be prevented or, if an attack is under way, it can be mitigated quickly and successfully.

TRUST AND IDENTITY SYSTEM

The first line of defense in the campus or branch infrastructure is to determine who or what is accessing the network, what resources these users should have access to, what the state of the device is. In addition to the tailgater, the rapid adoption of wireless networking in campus and branch offices has added to the security headaches of the network manager because a malicious or random user can sit in a parking lot and gain access to the network. At best, these unauthorized users are drawing on network bandwidth that does not belong to them; at worst, they could be launching attacks or spying on other users. Additionally, companies have large interconnections with suppliers, partners, and customers, all of whom may need access to the enterprise network to maintain productivity—but none of whom should have access to all the network. Cisco offers a comprehensive suite of capabilities to help ensure that trust and identity are maintained within the switched network.

Trust and identity comprises two fundamental components:

- Identity-based networking services, which identifies and validates the network user or device credentials prior to granting physical access to the network and can be used to ensure access to the correct network resources
- Network Admission Control (NAC), which identifies the posture (or compliance) of the device to ensure that the device can connect to the network without undue hazard

The trust and identity system provides the foundation for network access and policy control. Figure 3 shows how the different steps in a trust and identity system work together.

Identity-Based Networking Services (IBNS)

Most companies require employees to log into the network to access servers, e-mail, and other resources; Identity-Based Networking Services (IBNS) uses the same principles. To use an analogy, IBNS is similar to a person having a combination passport and airline frequent flyer card. The passport provides the authenticator with the identity of the person (often a security operations concern), whereas the frequent flyer card designates how the person is to be treated (often a network operations concern). This capability has some obvious benefits. First, because the switched network is the first line of defense, IBNS provides the first level of that first line of defense by first allowing or disallowing the user onto the network. Another benefit is the ability to treat that user according to the predefined policies regardless of where they are in the network. Finally, in some network architectures, the identity of the user can map directly to a particular workgroup or VPN that can be segmented from other workgroups.

IEEE 802.1x

The first steppingstone of IBNS is the IEEE 802.1x protocol, a MAC-layer protocol that communicates with a RADIUS server—such as the Cisco Secure ACS—to map the end station to the username and password. The 802.1x standard operates between the end station and the RADIUS server. Because this is an end-station protocol, an identity-enabled network functions only if the operating system supports 802.1x. Fortunately, most operating systems, including Windows XP, 2000, and NT as well as MacOS, support this capability (this list is not exhaustive). The switch acts as an intelligent middleman to the transaction and enables the port based on successful completion of the authentication process. The actual authentication mechanism used is called Extensible Authentication Protocol (EAP). EAP is essentially carried in the 802.1x frame, passed through the network by the switched infrastructure, and sent off to the authentication server.

The switch has the responsibility of querying the end station as soon as the end station connects to ask for login credentials. If the client supports 802.1x, then the end station replies with its credentials and the switch forwards that information off to the Cisco Secure ACS. If the client does not support 802.1x or does not authenticate, the client can be placed in a “guest VLAN” (discussed later). Upon successful authentication, the switch fully enables that port and allows access to the networked resources. Based on the information provided by the Cisco Secure ACS, however, the network can enable other policies.

VLAN Assignment for Management

One of the first policies that can be implemented on an IBNS is the VLAN name in which the user belongs. Suppose John Smith is an authenticated user who works in the marketing department. After authentication, the Cisco Secure ACS server returns to the switch with the VLAN name in which that user belongs, namely VLAN Marketing. The switch maps that name to a predefined VLAN number. This aids the network manager in tracking and accounting for users in the network. It is important to note that this does not mean that a VLAN number is following a particular user throughout the enterprise (a setup that would create a VLAN management nightmare as well as a significantly more complex spanning-tree domain). Only the configured VLAN name is provided.

Segmentation

One of the most interesting aspects of IBNS is the ability to segment users into different workgroups or VLANs based on who they are. For example, suppose an enterprise's customers come into the building and want to connect their laptop to the network to download their e-mail remotely. The enterprise would understandably not want those customers to access confidential information, so a guest VLAN can be created. By default, these customers would not be authenticated by the Cisco Secure ACS server and could be placed, for example, on a segmented VLAN that has access only to the Internet. In that way, the enterprise network is secured from unauthorized access while the customers are still able to access their resources over the Internet.

Segmentation has even greater applications, depending on the level of granularity and security required. If many disparate workgroups reside on a common IP infrastructure, segmentation based on identity could become very useful in providing secure VPNs. To use another example, imagine a large enterprise that houses many contractors and suppliers in addition to its employees. The network manager may want to treat each group (contractors, suppliers, and employees) differently, allowing them different network access, different quality-of-service (QoS) policies, or different performance levels on the network. Another example would be an airport, which has a common infrastructure (gates, ticket counters, etc.) and a common IP network. Based on identity, the network could distinguish a United Airlines employee from an American Airlines employee and provide access to each according to its own secure workgroups.

Segmentation also has implications for transmission over the common backbone. It is unlikely that the entire network would be running on a single switch, so secure transmission of segmented data is important. This is discussed in greater detail in the "Secure Connectivity System" section of this paper.

Network Admission Control

Ensuring that the user's identity is verified prior to network access is an important component of the trust and identity system. However, identifying the user solves only part of the problem. Although users may be allowed on the network based on the overall security policy, the computers they are using may not be desired on the network—a situation that becomes a network operations issue. Why is that? The pervasiveness of laptop computers in today's environment has the benefit of increasing users' productivity because they can have their computer with them in any location they choose. This, however, has a disadvantage: these computers are far more likely to become infected with a virus or worm, which may be unintentionally carried into the corporate environment.

NAC is an important part of the Cisco Self-Defending Network initiative. Whereas IBNS verifies the identity of the user, NAC identifies the "posture" of the device. NAC on the switching platforms works as a system in conjunction with the Cisco Trust Agent. The Cisco Trust Agent collects security state information from multiple security software clients, such as antivirus clients, and communicates this information to the connected Cisco network where access control decisions are enforced. Application and operating system status, such as antivirus and operating system patch levels or credentials, can be used to determine the appropriate network admission decision. Cisco and NAC cosponsors will integrate the Cisco Trust Agent with their security software clients. Cisco is partnering with such vendors as McAfee Security, Symantec, Trend Micro, and IBM for integration of their virus-checking software into the Cisco Trust Agent.

The switches demand host credentials from the Cisco Trust Agent and relay this information to policy servers where NAC decisions are made. Based on customer-defined policy, the network enforces the appropriate admission control decision: permit, deny, quarantine, or restrict. These ACLs are configured automatically in the edge switches based on the policy returned to the switch. If clients do not authenticate correctly, they can be placed in the "quarantine VLAN" so that they can update their virus-checking software or client-based security agents. It is possible that, based on 802.1x authentication, the port is enabled, only to be restricted or denied because a device is not considered "safe."

The policy server evaluates the endpoint security information relayed from network devices and determines the appropriate access policy to apply. Cisco Secure ACS server, an authentication, authorization, and accounting (AAA) RADIUS server, is the foundation of the policy server system. It may work in concert with NAC cosponsor application servers that provide deeper credential validation capabilities, such as antivirus policy servers.

Using the Cisco Trust and Identity System

Working together, the components of the trust and identity system provide the foundation for network access and policy control in the network. Figure 3 shows how the different steps in a trust and identity system work together.

Figure 3
Trust and Identity System

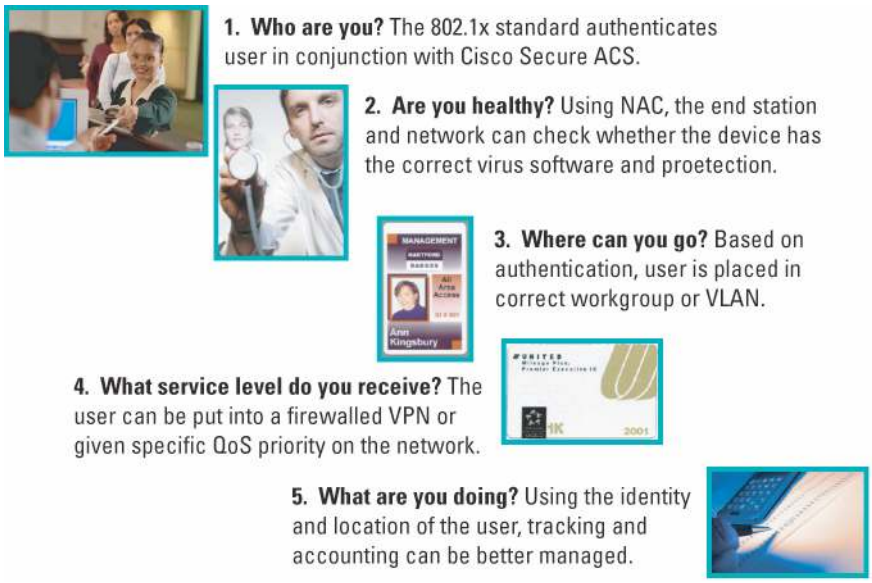

THREAT DEFENSE SYSTEM

The threat defense system is designed to protect the network from attack as well as ensure that the network can recover if an intrusion occurs or an attack is launched. The switching infrastructure is responsible for a large part of the threat defense system because the switch could be the target of an attack and also would aid, unwillingly of course, in the perpetuation of the attack as it connects users and servers to each other. The Cisco Catalyst switch offers important functions to protect itself and the users connected to it.

Protecting the Switch Configuration and Control Packets

First, the switch itself can be the victim of an attack. This attack can come in many forms, such as hackers attempting to access the switch to change the configuration, worms that flood the MAC forwarding table or the Layer 3 flow table, or an attempt to manipulate control information, such as routing updates or spanning-tree bridge protocol data units (BPDUs). Cisco Catalyst switches have mechanisms built into hardware and software to make the network self-defending.

Securing Remote Access

It is often important for a network manager to configure a device remotely. This has the obvious benefit of centralizing configuration of multiple switches in one location. Traffic sent from the network manager's configuration station to the switch device can potentially be snooped

for passwords or malicious intent. One of the “table stakes” aspects of security is the ability to encrypt this traffic to guarantee privacy. This is accomplished using the Secure Shell (SSH) Protocol. SSH provides a means to secure remote connections across an unsecured network. Data is sent through an encrypted tunnel to secure transmission and integrity of data. Both ends authenticate users and ensure secure file transfer and copying. This has the benefit of mitigating man-in-the-middle attacks and ensuring that critical management information is not compromised.

Changing the Switch Configuration

Changing the switch configuration is one of the basic types of attacks that can be launched against the infrastructure. Using password snooping, a malicious user can access the switch with the intent of modifying the configuration. The change can be simple (changing from 100 to 1000 Mbps, for example) to destructive, such as modifying control information, such as the routing table. The Cisco Catalyst Family provides encrypted passwords as well as *multi-level account privileges* to help ensure that the device itself cannot be accessed. This helps ensure that, even within the IT staff, access can be tiered such that not all individuals have access to all switch configurations. This is coupled with *RADIUS and TACACS+ authentication* to ensure trust and identity and is often the first level of switch protection.

Spoofing Control Information

Another common attack is for an attacker to spoof control information and send it into the network. Control information, such as routing updates or spanning-tree BPDUs, are critical to network operation. These protocols ensure loop-free, deterministic operation that allows the network to function properly. A switch, seeing those packets, would innocently respond to them. The possible result is that the switch thinks that a network resides on an interface that it does not reside on, or that another spanning-tree root bridge exists, causing catastrophic consequences and seriously disrupting network operation.

To prevent routing information from being spoofed, Cisco Catalyst Layer 3 switches use *Message Digest Algorithm 5 (MD5) route authentication*. MD5 is a protocol by which routing updates between neighboring Layer 3 switches or routers exchange information such that the message can be exchanged and acknowledged between only those two routing peers. Other routes injected into the network are ignored, ensuring that the routed stability of the network is not compromised.

Spanning tree lacks the intelligence inherent in Layer 3 routing protocols. Therefore, disruption to a spanning-tree network is relatively straightforward and can be done either intentionally or unintentionally. For example, a user with multiple PCs might use a switch to connect them together. That switch could send BPDUs into the network saying “I am the root bridge.” At best, troubleshooting the network could be problematic if a problem arises. At worst, the switched network could become unstable, affecting all users in the Layer 2 domain.

To prevent this scenario, Cisco has developed numerous enhancements to the Spanning Tree Protocol: the port can either accept BPDUs or not. If the device connected to that port sends an unauthorized BPDU into the network, that port is disabled. This feature is best used on the user-facing interfaces. Root Guard is enabled on a per-port basis to ensure the integrity of the root bridge. Ports enabled with Root Guard become designated ports, meaning that surrounding switches cannot become the root bridge. This prevents the Layer 2 network from an unexpected network reconvergence.

Attacks Against the Switch Processor and Forwarding Table

DoS attacks, launched by worms or viruses, are often launched against the network infrastructure, causing significant network instabilities by disabling the ability of the switch to learn addresses and process control information. These attacks are among the most common seen in enterprise and commercial networks today, and they can have a huge impact. According to a 2003 CSI/FBI Computer Crime and Security survey of 400 companies, more than 90 percent of the participants reported security breaches. Their estimated losses totaled more than $200 million. In addition, worms can spread alarmingly fast. For example, the recent Structured Query Language (SQL) Slammer Worm doubled every 8.5 seconds. In the span of 3 minutes, it could perform 55 million scans per second, bringing a 1-Gbps link to a standstill in just 1 minute.

Because each attack can affect the switch in a different way, different tools are available to prevent and mitigate these attacks. These attacks basically flood the network with bogus traffic, thus shutting out legitimate traffic that the switch needs to process, such as routing updates, unknown MAC addresses, or spanning-tree BPDUs. Each of these attacks is thwarted by specific Cisco Catalyst features, a summary of which is shown in Table 1. Each of these attacks and mitigating features will now be explored in more detail.

Table 1. Mitigating Attacks on the Switch

Worm or Virus Behavior	Cisco Catalyst Feature to Solve
Flooding the network links	QoS Scavenger Class
Flooding the content-addressable-memory (CAM) table	Port Security
Flooding the Layer 3 flow cache	Cisco Express Forwarding
Flooding the switch CPU	Control Plane Rate Limiting

Flooding the Network with Bogus Traffic

A DoS or worm attack can flood network links, filling up device buffers and filling up the Ethernet link between switches. This in turn can destroy voice quality and slow application "goodput" down considerably (the term “goodput” refers to the true application throughput achieved, which is independent from the raw throughput of the switch). Many of these worms scan the network, using many different TCP or User Datagram Protocol (UDP) flows. This creates a huge volume of data over the network. The next section discusses problems and solutions caused by an excess of flows, but the problem of link congestion is an aggregate problem: no one flow is causing the problem, so all flows on a link must be restricted. This problem can be mitigated and the links protected by using *Cisco Catalyst QoS Scavenger Class*.

QoS is an intelligent delivery system that enforces a trust boundary, classifying traffic and prioritizing that traffic based on policies. This allows the network manager to manage rates and reclassify specific traffic types, if required, as well as schedule and transmit based on configurable delivery requirements. The scavenger class is traditionally intended to provide deferential services, or *less-than best-effort* services, to certain applications. For *best-effort* traffic, an implied good-faith commitment states that at least some network resources are available. In this model, however, good faith cannot be assumed and scavenger class services are used to *deprioritize* traffic from systems with abnormally high traffic rates.

This solution uses the multiple-queues-per-port architecture in the Cisco Catalyst switches. During a worm attack, a large volume of traffic is sent by an end station into the network. This traffic is compared against the policies defined, including how to determine scavenger class traffic. Scavenger class traffic is defined as traffic that, over a sustained period of time, bursts or transmits above a defined threshold, something that a properly functioning end station would not do. The first detection is not enough to drop packets because it may be a legitimate burst. However, it is unlikely that a port would see user traffic sustained over time. That traffic is then marked as “scavenger” and placed in the lowest-priority queue. That queue is configured to be shallow, so that queue overflow, or tail drops, are frequent. TCP-based flows throttle back, and UDP flows are dropped. In this way, voice traffic and other traffic, such as legitimate best-effort traffic, is protected.

Flooding the Switch Layer 2 Forwarding Table

A Layer 2 forwarding table, also known as a CAM table, builds the switch forwarding table based on MAC addresses. This is how a switch or bridge performs the forwarding, filtering, and learning mechanisms at Layer 2. The table, however, has only a finite space. This attack forces a switch to learn bogus MAC addresses, thereby overloading the CAM table. When a CAM table is full, it cannot learn additional addresses that it needs to learn , so the data from those addresses is flooded throughout the Layer 2 VLAN domain. Although this is standard behavior for Layer 2 switches, it can result in poor network and end-station performance.

To prevent this kind of attack, the Cisco Catalyst switch offers the *Port Security* feature, which limits the number of MAC addresses that can be learned at a given port. If more addresses enter the switch, the Cisco Catalyst switch puts those ports into error-disable mode to protect the

network in case of an attack. This helps ensure that only a small number of MAC addresses are learned at a given port, locking down the port if other addresses are learned. In this way, an attack is stopped immediately. Another feature is *Broadcast Suppression*, or Storm Control. This feature sets a threshold for the number of broadcast or multicast packets an interface can transmit into the network. If that threshold is crossed, those packets are dropped.

Flooding the Switch Layer 3 Flow Table

A flow is defined as a Layer 3 connection between a source and destination IP address or between corresponding TCP or UDP port numbers. Many worms vary the source and destination IP, TCP, or UDP, forcing the switch to learn hundreds of thousands of new flows. Many switch architectures use a flow-based architecture, which requires the packet in a flow to be sent to the CPU (or route processor) for a lookup against the routing table. After that lookup is performed, the flow entry is cached in the hardware forwarding table so that high-speed switching can occur. It is this "first-packet processing" that this type of attack exploits. By varying the IP, TCP, or UDP information, the switch CPU can be overloaded and thereby crippled. Legitimate flows are no longer switched or learned, and the routed network can be seriously impacted. Recent worms that had this effect were Code Red, Slammer, and Witty.

The Cisco Catalyst switches use *Cisco Express Forwarding (CEF)* a switching mechanism based on the topology—not the traffic—of the network. As the routing table is populated by the routing protocol (Open Shortest Path First [OSPF], Enhanced Interior Gateway Routing Protocol [EIGRP], etc.), a hardware-based table is populated based on network prefix. For example, instead of IP destination address 192.24.20.25 residing in interface *GigabitEthernet 0/1*, the network 192.24.20.0/24 would reside off that interface. Originally designed to make the network more resilient against route flaps, it also has direct applicability to worm attacks. Because Cisco Express Forwarding does not care about the flows in the network, the Layer 3 forwarding table is not affected. When coupled with QoS scavenger class, Cisco Express Forwarding can mitigate the effects of a worm in the network.

Attacking the Switch CPU

Besides attacking the forwarding tables of the switch, the CPU of the device itself can be attacked by sending control information, such as ARP packets, to the CPU for processing. With a finite amount of processing power, the CPU can become overloaded, with a possible result that real control packets, such as routing updates or BPDUs, are dropped. Many network managers are familiar with the consequences of a CPU running at 100-percent usage—the device and the network become unstable.

The Cisco Catalyst 6500 supports *Control Plane Rate Limiting*, which limits the rate at which packets can be sent to the CPU. It is very unlikely that, in normal operation, a large amount of traffic would be sent to the CPU unless software-based features are enabled. It is even more unlikely that traffic of the type used in these attacks is sent at a high data rate to the CPU. These types of packets include *Internet Control Message Protocol (ICMP) unreachable, ICMP redirect, Cisco Express Forwarding no route, time-to-live (TTL) failures, and ingress or egress ACLs.* If rate limiting on these types of packets is enabled to the CPU, excess packets above a certain rate are dropped, ensuring that the CPU can handle (most likely drop) the malicious packets without failing. This allows the CPU to continue handling normal system tasks—even while under attack.

The Cisco Catalyst 4500 Series supports similar capabilities with *CPU QoS*, which provides multiple queues to the CPU. Each queue receives a certain amount of bandwidth up to the CPU, which is governed by a round-robin scheduler. This limits the amount of bandwidth available to any particular traffic type and helps ensure that one queue (or type of traffic) cannot overrun the processor. In normal operation, the queues are sufficient to handle any and all control packets with the processing they need. Additionally, rate limiters can be applied to *ARP* and *DHCP* packets to help ensure that, if an attack is launched at the Cisco Catalyst 4500 CPU, priority traffic will still be processed.

Protecting the End Devices

Besides protecting itself, the Cisco Catalyst switch must also protect the users and servers attached to it. Unlike attacks that are disruptive to the network, many attacks against users and servers can go undetected. These attacks, often called man-in-the-middle attacks, use common tools that can be downloaded from the Internet. These tools use a variety of mechanisms that allow a malicious user to spy on other employees, managers, or executive staff. This can result in the theft of proprietary information as well as privacy violations.

According to the 2003 CSI/FBI Computer Crime and Security survey mentioned previously, this accounted for more than $70 million in loss, with 75 percent of the respondents citing insider attack by disgruntled employees as the likely source of the attacks. In addition, new privacy legislation can result in severe penalties if confidential data from your network falls into the wrong hands. For example, the Health Insurance Portability and Accountability Act (HIPAA) law in the United States allows for a fine of up to $250,000 and 5 years in jail, per incident, if the security of confidential electronic health information is compromised. The Cisco Catalyst switches offer numerous capabilities to mitigate these attacks and protect data and user integrity.

Table 2. Man-in-the-Middle Attacks and Solutions

Attack	Cisco Catalyst Feature to Solve
Spoofing the DHCP server	DHCP Snooping and Port Security
Spoofing the default gateway	Dynamic ARP Inspection (DAI)
Spoofing an IP address	IP Source Guard

Spoofing the DHCP Server

One of the best ways for a hacker to gain control of a switched network is to spoof the DHCP server, thereby sending out false addresses and default gateway information. (It is worth noting that this can be done accidentally, such as if a misconfigured user starts up a DHCP server and begins serving addresses to other users.) By spoofing the DHCP server, users who are sending out DHCP requests for addresses do not actually reach the DHCP server and are given an address by the malicious user. This is the first step in malicious users' plans to gain access to information they should not have.

The first step in preventing this type of an attack is *DHCP Snooping*, a feature on Cisco Catalyst switches that helps ensure that only certain ports on a switch can process DHCP information other than a DHCP request. This feature defines trusted ports (which can send DHCP requests and acknowledgements) and untrusted ports, which can forward only DHCP requests. It is assumed that trusted ports are those that connect to either the DHCP server itself or switched ports, such as uplinks, that connect the switch to the rest of the network. If a malicious user attempts to send a DHCP acknowledgement (ACK) packet into the network, the port is shut down. This feature enables the switch to build a *DHCP Binding Table* that maps a client's MAC address, IP address, VLAN, and port ID. This table is used as a foundation for other Cisco Catalyst features that further prevent attacks.

This feature is coupled with *DHCP Option 82*, which allows the switch to insert information about itself into the DHCP packet. The most common type of information to insert is the physical port ID of the DHCP request. This provides the network with considerable intelligence to prevent rogue device and server attachment.

Identity Spoofing with Gratuitous ARP Packets

End stations send an ARP packet to discover the MAC address of its default gateway and map it to the gateway IP address. This is normal operating procedure in an IP network. However, there is a security hole in ARP—a router or end station is allowed to send out a gratuitous ARP, which is an unsolicited ARP reply. This is often referred to as *ARP poisoning*. By sending out a gratuitous ARP, malicious users can spoof the default gateway or any other device in the network, placing themselves between the user and the true default gateway. This allows malicious users to sit between the innocent user and the default gateway and spy on all the data being sent on the network by the user. The

problem is that neither the default gateway nor the end user is aware that this attack is taking place. Therefore, malicious users can continue spying on private information for as long as they want, allowing them time to access passwords, e-mail messages, transactions, etc.

Cisco Catalyst switches offer a feature called DAI to stop this attack. Like DHCP Snooping, DAI uses the concept of trusted and untrusted ports to decide which ARP packets need to be inspected. To do this, DAI intercepts all ARP packets and examines them for proper MAC-to-IP bindings. This is done by using the DHCP binding table that was built by enabling DHCP Snooping. If an ARP packet arrives on a trusted port, then no examination is made. If it arrives on an untrusted port, the ARP is examined and compared against the table.

For instance, user John Smith has an IP address of 172.20.24.45, a MAC address of 00aa.0062.c609, and resides on port *FastEthernet* 3/47 (an untrusted port). This information is recorded in the DHCP binding table. If John Smith sends a gratuitous ARP packet into the network, DAI examines the ARP packet and compares its information with the information in the table. If there is not a match, the ARP packet is dropped.

Spoofing an IP Address

Another potential attack is for a malicious user to spoof an innocent user's IP address. This is very simple to do in most operating systems and can make the malicious user appear to be on another subnet or bypass ACLs. Using the DHCP binding table, *IP Source Guard* checks the binding of an IP address to a MAC address, its port, and its associated VLAN. It does this by automatically configuring an ACL in the Cisco Catalyst ternary content addressable memory (TCAM) that limits a port to the configured IP address handed to the end station by the DHCP server. If an IP address appears that is inconsistent with the information in the binding table, the port is disabled.

Mitigating Man-in-the-Middle Attacks Summary

Cisco Catalyst switches provide inherent capabilities to stop man-in-the-middle attacks. The attacks themselves typically involve several steps: accessing the port, sending out bogus DHCP information and gratuitous ARPs, or hijacking another DHCP-served IP address. To thwart these attacks, a combination of all these features, working together, is recommended. At the core of these features is the DHCP binding table, which is first populated by DHCP Snooping. DAI and IP Source Guard use this table to further prevent other means of attacks. In this way, privacy within the network can be provided and theft of valuable or confidential information prevented.

Integrated Firewall and Intrusion Detection

The increase in threats to the enterprise and commercial business has forced network managers to be much more proactive and deploy threat-defense appliances throughout the network. Because campus and branch networks run at gigabit speeds, integration of these capabilities into the LAN switching infrastructure is important. The Cisco Catalyst 6500 offers the industry's leading services platform by integrating dedicated processing for high-touch services directly onto the platform.

Firewalls are specialized devices that act as an access control system, inspecting every packet entering or exiting the network segment. Based on security policies, the firewall can permit or deny traffic coming into or out of the segment, providing a very granular level of network control. Firewalls are generally associated with the Internet edge, an obvious place in the network in which security policies are critical. The need for securing connectivity within the network infrastructure has driven the requirement for firewalls within the enterprise network itself.

Intrusion detection systems (IDSs) also are used for threat defense. IDS appliances and modules sit in the data path and look for specific patterns that indicate an attack might be under way. This can be done by matching patterns, or signatures, of a particular attack or by looking for network behavior that is anomalous. The sooner the attack or network intrusion can be detected, the faster it can be stopped.

The *Firewall Services Module (FWSM)* for the Cisco Catalyst 6500 is the first (and only) integrated firewall in a multilayer switching platform. Firewalls are often associated with "keeping intruders out" of the network and are typically deployed at the Internet edge. This is usually a security-operations consideration. However, the inclusion of firewall functions at high-speed data rates provides unique abilities to segment users into secure workgroups (a network operations consideration). This capability has helped enable enterprises to provide greater levels of

security in their network, better map security policies to workgroups, and ensure that those policies are available to the user regardless of the user's location in the network.

The *Intrusion Detection Service Module (IDSM)*, also for the Cisco Catalyst 6500, integrates intrusion detection directly into the switch data stream, connecting into the Cisco Catalyst 6500 high-speed switching fabric. The IDSM spans across all the VLANs configured in the switch and scans for any intrusion signatures.

SECURE CONNECTIVITY SYSTEM

In any enterprise, privacy is an important consideration. Individuals expect that their desks and personal space will be respected and not disturbed, and that their computers will not be entered and their data viewed by unauthorized persons. That expectation of privacy extends to the network itself. Many of the capabilities already discussed in this paper help provide privacy in the network by controlling who has access to the network as well as protecting it from a variety of attacks. However, some enterprises, as well as unique implementations such as airports or universities, require more advanced levels of privacy that are guaranteed within the campus switched infrastructure.

The secure connectivity system is often associated with VPNs, an appropriate association. VPNs over the public Internet require the setup of a private, encrypted, and protected connection between the client and the enterprise. In many ways, though, the concept of VPN has been around in campus switching for many years. More stringent security needs have added more technology requirements to the architecture, but the concepts are still the same.

Secure connectivity is based on the concept of segmentation. Segmentation essentially divides the network into subworkgroups based on any number of criteria: network addressing, location of users and the wiring closet, physical workgroup-to-user mapping, or identity of the user or device connecting to the network. This section discusses how the network is segmented and how additional security is provided. Following are the primary components of the secure connectivity system:

- Virtual LANs—Traffic segmentation at Layer 2
- Virtualized firewalls—Policy enforcement at the VLAN boundary
- Secure transport—Segmentation of workgroup traffic over the common backbone using Multiprotocol Label Switching (MPLS)

Virtual LANs and Private VLANs

VLANs have been a part of Layer 2 switching since its infancy in 1994. The concept is simple: a common switching infrastructure is segmented into different broadcast domains, each domain being its own LAN. With Cisco Inter-Switch Link (ISL) and the later standard, IEEE 802.1Q, those individual VLANs could transit the entire switched infrastructure. The introduction of Layer 3 switching coupled with the centralization of data servers and resources eliminated that need for those workgroups to span the entire enterprise. Many enterprises, however, are reevaluating the need for workgroups that span the enterprise, and VLANs are the first step in that process. It is important to note that secure connectivity does not require VLANs to span end to end in the enterprise; they are merely the first means of segmentation on the switched infrastructure.

Another VLAN feature used for privacy is *private VLANs*, essentially smaller, sub-VLANs within a larger VLAN that each requires privacy. This requirement was first seen in hosted data centers, but also is often seen in networks that have many disparate groups connecting into a common infrastructure. Private VLANs designate certain port types: a promiscuous port, an isolated port, and a community port. A promiscuous port is one that can communicate with all ports in all private VLANs, such as an uplink port. An isolated port is one that can communicate only with the promiscuous port and no other. Finally, a community port is one that can communicate with the promiscuous port as well as other ports within its own private VLAN. In no case can ports in one private VLAN communicate with ports in other private VLANs at Layer 2; to accomplish this, the packet must be sent to the default gateway, routed at Layer 3, and forwarded back. In this way, privacy and policy can be assured.

Virtualized Firewall

Many enterprise networks are essentially behaving as service providers by offering secure service to each of its "customers." In some large enterprises, network usage is billed back to each department based on the level of service required. In other networks, such as airports, different companies are actively using a common backbone and must be secured from each other. Coupled with the need for workgroup privacy, the virtualized nature of the network has lead to *firewall virtualization*.

Firewall virtualization uses the FWSM in the Cisco Catalyst 6500 to segment off different workgroups in the network. In this case, the FWSM looks to the network like multiple firewalls, each with its own policies (NAT, ACLs, etc). When placed at the network distribution layer, which is aggregating multiple wiring closets, firewall virtualization can segment the network into multiple security domains, allowing very granular policy control per "customer." VLANs can now be used as a first-level means of segmentation, with each VLAN mapping to a firewall instance on the FWSM. This secures each VLAN from other VLANs in the switching domain while also securing the traffic for transport over the common backbone.

Transport over the Backbone Using MPLS

The final piece of secure connectivity in the switched infrastructure is how to actually transport the secured workgroup over the common backbone. In a Layer 3 routed network, VLANs no longer exist and the traffic is switched from source to destination in an open manner. Securing that communication over the backbone is important. Fortunately, the IP backbone provides for secure transmission mediums that can be overlaid on the common backbone. A mechanism that enables transmission of secure workgroups is MPLS.

Long used in the service provider environment, emerging segmentation requirements in the enterprise are now making MPLS an attractive campus technology. MPLS VPN, specified in RFC 2547, allows for specific workgroups, each with its own virtual routing and forwarding (VRF) table, to exist within the MPLS network. Each VRF table maps to a specific "customer" whose traffic is being sent over the common backbone infrastructure. MPLS designated two types of routers: the provider edge (PE) and the provider (P) router. Each has specific functions that enable the VPN to function properly.

In a campus MPLS implementation, provider edge routers receive routes from the "customer edge" router, which is located at the edge of the network and extends the VPN space to the edge of the routed domain. In a campus VPN implementation, the customer edge is often a LAN switch or a multilayer switch that is centralizing security services. The provider edge router then transports these routes to other provider edge routers across the MPLS backbone. In the middle of the network are the provider [NOT provider edge?] routers, also called label switch routers (LSRs), which implement the transport service. It is important to understand that VPN information is required only at the provider edge. It is at the provider edge that the VPNs are partitioned; only provider edges that belong to a particular VPN need to have any knowledge of that VPN. Additionally, because the P routers are switching based on the label appended to the packet, they also do not need to be aware of the VPN. This allows for greater scalability in the routing table of the P-router.

Entry to the MPLS backbone is governed by the provider edge, allowing the network engineer to control the positioning, the number of access points into the VPN, and the security policy. This point of access provides centralization of services and policy for a campus, whether large or small. Centralizing these services at the provider edge provides several important benefits. First, the network engineer can reduce the number of ACLs in the network. Secondly, at this point, transit areas between VPNs and access to shared resources, such as the data center, WAN, or Internet, can be managed centrally for a potentially large portion of the network. Finally, centralization allows for the sharing of service modules or appliances that provide firewalling and intrusion detection. A primary component of this system, for example, is the Cisco Catalyst 6500 Firewall Virtualization solution, which can firewall VPNs into the MPLS cloud for greater security and privacy.

CISCO CATALYST INTELLIGENT SWITCHING PORTFOLIO

Cisco offers the most complete range of switching products in the networking industry today. These products can be deployed across enterprise, service provider, and commercial applications and provide security capabilities designed to prevent, mitigate, and recover from security breaches, attacks, or intrusions.

Cisco Catalyst 6500 Series

The Cisco Catalyst 6500 Series is the industry's most flexible and innovative switching platform and is the premier Cisco switching platform, setting the standard for IP Communications and application delivery in campus and service provider environments. Since the inception of the platform in 1999, it has provided investment protection by supporting all line cards and supervisor engines. The Cisco Catalyst 6500 Series has evolved from a 32-Gbps system with Supervisor Engine 1 to a 720-Gbps system featuring more than 400 million packets per second of performance while maintaining backward compatibility. Today, the Cisco Catalyst 6500 Series provides the highest port densities for 10/100, 10/100/1000 Gigabit Ethernet, and 10-Gigabit in the industry today.

The Cisco Catalyst 6500 supports the greatest range of hardware- and software-based security capabilities, including identity and man-in-the-middle attack mitigation features. Network managers, who previously had to purchase specialized appliances to implement firewall, intrusion detection, and Secure Sockets Layer (SSL) termination, now can integrate this function directly into the Cisco Catalyst 6500 through feature modules. Unique in the industry, these specialized modules provide high-speed services that can be managed with the switch while providing multigigabit performance. In addition to integrated trust and identity and threat defense, the Cisco Catalyst 6500 leads in securing connectivity, in terms of IP VPN aggregation, through the VPN service module; it also integrates secure transport technologies such as MPLS and generic routing encapsulation (GRE) into hardware. The modules can be deployed in campus environments to create virtual, secure workgroups or in data centers to provide an integrated services platform.

Finally, the Cisco Catalyst 6500 implements a wide range of features to mitigate against DoS attacks. This includes CPU-based rate limiting, which gives the network manager a full range of filters with which to protect the switch processor. Per-port scavenger class queuing can limit the scope of attacks in progress by discarding excess traffic that is out of profile for a particular port.

Cisco Catalyst 4500 Series

Cisco Catalyst 4500 Series switches are designed for enterprise wiring closets, branch offices, and smaller Layer 3 distribution points. The Cisco Catalyst 4500 Series comprises a series of chassis that support three choices of supervisor engines and an extensive set of line cards, including high-density, 10/100, 10/100/1000 (both with 802.3af Power over Ethernet [PoE] options), 100BASE-FX, and 1000BASE-X. The Cisco Catalyst 4500 Series scales to support up to 336 10/100/1000 ports over copper or fiber interfaces, and is designed for deployment in enterprise wiring closets, small network backbones, and branch office environments. In the Cisco Catalyst 4500 Series, all the system intelligence resides in the supervisor engine—incremental, hardware-based capabilities that could be offered in the future, such as IPv6 and 10 Gigabit Ethernet, can be added with a supervisor engine upgrade. Today's supervisor engine provides up to 96-Gbps switching capacity with 72 million packets per second of switching performance.

The focus on the wiring closet has allowed the Cisco Catalyst 4500 to specialize in wiring-closet features and functionality. Like the Cisco Catalyst 6500, the Cisco Catalyst 4500 supports a full suite of trust and identity features, including 802.1x and many of its extensions. Also, the Cisco Catalyst 4500 mitigates man-in-the-middle attacks at the ingress of the network, stopping these types of attacks before they can make their way deeper into the network. Finally, the Cisco Catalyst 4500 implements multiple CPU-based queues, each of which is limited in bandwidth. This limits the effects of DoS attacks based on DHCP and ARP packets.

Cisco Catalyst 3750 Series

The innovative Cisco Catalyst 3750 Series switches improve LAN operating efficiency by combining industry-leading ease of use and the highest resiliency available for stackable switches. This product series represents the next generation in desktop switches, and features Cisco StackWise™ technology, a 32-Gbps stack interconnect that allows customers to build a unified, highly resilient switching system—one switch at a time. For midsize organizations and enterprise branch offices, the Cisco Catalyst 3750 Series provides integrated security functions designed to ensure identity access and control as well as mitigate many common man-in-the-middle attack scenarios**.**

Cisco Catalyst 3560

Cisco Catalyst 3560 Series switches complement the Cisco Catalyst 3750 product line by providing standalone, fixed-configuration switches supporting IEEE 802.3af and Cisco prestandard PoE switches. In Fast Ethernet configurations, the Cisco Catalyst 3560 provides availability, security, and QoS to enhance network operations. These switches are best suited for standalone deployments in small wiring closets and branch offices that require low density, but still need intelligent switching services such as availability and security. Like the other fixed-configuration switches, the Cisco Catalyst 3560 Series embeds the Cisco Cluster Management Suite (CMS) Software, which allows users to simultaneously configure and troubleshoot multiple Cisco Catalyst desktop switches using a standard Web browser.

Cisco Catalyst 3550 Series

Cisco Catalyst 3550 Series intelligent Ethernet switches bring enterprise-class capabilities to midsize businesses and enterprise branch offices. Available in wire-speed Fast Ethernet and Gigabit Ethernet configurations, these switches can support intelligent services such as advanced QoS and rate limiting, security ACLs, multicast, and IP routing. These switches serve as access and distribution layer switches for midsize enterprise and branch office wiring closets. The Cisco Catalyst 3550 provides a baseline of security functions for 802.1x authentication and guest VLANs, port security to prevent MAC flooding attacks, and SSH and SSL for secure remote switch management.

SELF-DEFENDING YOUR NETWORK

Cisco Catalyst LAN switching is a primary foundational component of the Cisco Self-Defending Network initiative. By securing the infrastructure and protecting users, servers, and the switch itself, Cisco Catalyst Integrated Security provides the first line of defense in an enterprise or commercial company’s overall security strategy. Cisco is at the forefront of security leadership in LAN switching, driving the development and implementation of features designed for protection, control, and privacy.

Cisco innovative integrated security systems help organizations deploy trusted and protected business applications and services through modular, rich security services that can be flexibly deployed on Cisco routers, switches, endpoint software, and specialized security appliances. Cisco integrated security systems incorporate a comprehensive selection of feature-rich security services (VPN, threat defense, identity) that can be flexibly deployed on standalone appliances (including the Cisco PIX® security appliance, Cisco IDS sensors, and Cisco VPN concentrators), security hardware modules and software for routers and switches, security software agents for desktops and servers, and trust and identity management systems, along with a centralized management and analysis of all security technologies for the most complete system in the marketplace.

For more information, refer to http://www.cisco.com/go/security and http://www.cisco.com/en/US/products/hw/switches/index.html.

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on **the Cisco Web site at www.cisco.com/go/offices.**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

204064_ETMG_CC_08.04

Printed in the USA