Notes on Network Security: DHCP attacks[1]

DHCP attacks

DHCP Starvation

A DHCP starvation attack works by broadcasting DHCP requests with spoofed MAC addresses. This is easily achieved with attack tools such as “the gobbler”. If enough requests are sent, the network attacker can exhaust the address space available to the DHCP servers for a period of time. This is a simple resource starvation attack just like a synchronization (SYN) flood attack. Network attackers can then set up a rogue DHCP server on their system and respond to new DHCP requests from clients on the network.

Rogue DHCP Server

A rogue DHCP server is a DHCP server set up on a network by an attacker, or by an unaware user, and is not under the control of network administrators. An accidental rogue device is commonly a modem or home wireless router with DHCP capabilities which a user has attached to the network unaware of the consequences of doing so. Rogue DHCP servers are also commonly used by attackers for the purpose of network attacks such as Man in the Middle, Sniffing, and Reconnaissance attacks.

By placing a rogue DHCP server on the network, a network attacker can provide clients with addresses and other network information. Because DHCP responses typically include default gateway and Domain Name System (DNS) server information, network attackers can supply their own system as the default gateway and DNS server resulting in a man-in-the-middle attack.

Man in the Middle DHCP attacks can be used to forge network resources. The Rogue DHCP reply will offer an IP address and information that may designate the attacker’s machine as the default gateway or Domain Name System (DNS) server. If the attacker is designated default gateway, any clients with addresses assigned from the Rogue DHCP Server will forward packets to the attacking device, which may in turn send them to the desired destination, or possibly elsewhere. If the attacker also designates its own Rogue DNS Server(s), they may design phishing websites to obtain other confidential information, such as credit card details and passwords.

Mitigation

DHCP Snooping

DHCP snooping is a DHCP security feature that provides network security by filtering untrusted DHCP messages and by building and maintaining a DHCP snooping binding database, which is also referred to as a DHCP snooping binding table.

DHCP snooping is a Cisco Catalyst switch feature that determines which switch ports can respond to DHCP requests. Ports are identified as trusted and untrusted. Trusted ports can source all DHCP messages, while untrusted ports can source requests only. Trusted ports host a DHCP server or can be an uplink toward the DHCP server. If a rogue device on an untrusted port attempts to send a DHCP response packet in to the network, the port is shut down. This feature can be coupled with DHCP option 82, in which switch information, such as the port ID of the DHCP request, can be inserted into the DHCP request packet.

1 http://itsecurity.telelink.com/dhcp-attacks/

Untrusted ports are those not explicitly configured as trusted. A DHCP binding table is built for untrusted ports. Each entry contains the client MAC address, IP address, lease time, binding type, VLAN number, and port ID recorded as clients make DHCP requests. The table is then used to filter subsequent DHCP traffic. From a DHCP snooping perspective, untrusted access ports should not send any DHCP server responses, such as DHCPOFFER, DHCPACK, or DHCPNAK.

When a switch receives a packet on an untrusted interface and the interface belongs to a VLAN in which DHCP snooping is enabled, the switch compares the source MAC address and the DHCP client hardware address. If addresses match (the default), the switch forwards the packet. If the addresses do not match, the switch drops the packet.

Port Security

Port Security can be used to mitigate DHCP starvation attack by limiting the number of MAC addresses allowed on a port. You can use the port security feature to restrict input to an interface by limiting and identifying the MAC addresses of the stations allowed to access the port. When you assign secure MAC addresses to a secure port, the port does not forward packets with source addresses outside the group of defined addresses.

Port security allows you to specify MAC addresses for each port or to permit a limited number of MAC addresses. When a secure port receives a packet, the source MAC address of the packet is compared to the list of secure source addresses that were manually configured or autoconfigured (learned) on the port. If a MAC address of a device attached to the port differs from the list of secure addresses, the port eithershuts down permanently (default mode) or drops incoming packets from the insecure host. The behavior of the port depends on how you configure it to respond to a security violator.