CYGNA LABS

DHCP Threats and Mitigations

By Timothy Rooney

About Cygna Labs

Cygna Labs is a software developer and one of the top three global DDI vendors. Many Fortune 100 customers rely on Cygna Labs' DDI products and services, in addition to its industry-leading security and compliance solutions to detect and proactively mitigate data security threats, affordably pass compliance audits, and increase the productivity of their IT departments. For more information, visit https://cygnalabs.com.

Introduction

Dynamic Host Configuration Protocol (DHCP) is a core network service for initializing devices for operation on an Internet Protocol (IP) network. DHCP is typically deployed within an organization and is not typically exposed to Internet access. However, malicious administrators or malware could infiltrate DHCP to deny service to users or to manipulate device initializations and configurations.

Like any risk analysis process, the initial step in formulating a security strategy, each organization must assess the relative risk of occurrence of such attacks with the associated outage costs against mitigation costs. Outage costs for DHCP could be severe, as any network devices requiring a new or renewed DHCP lease could be left unserved and therefore unconnected. For this reason, most organizations mitigate server or network outage risks by deploying redundant DHCP servers. For this and other DHCP risks, various mitigation strategies are discussed in this paper.

DHCP Service Availability Risks

Your users and devices rely on DHCP for basic IP network initialization, so if DHCP services are unavailable or are performing inadequately, end users or subscribers may be unable to access the network.

DHCP Service Availability Vulnerabilities

Risks associated with the availability of DHCP services include:

- Inadequate DHCP service capacity due to too few servers deployed and/or servers deployed with inadequate processing, storage or input/output specifications. If a DHCP server is functioning as a failover DHCP server, its load would presumably increase above nominal levels during an outage of a primary server, so its capacity must be engineered for the heaviest expected loads.
- Unavailability of DHCP services due to network unreachability, possibly due to relay agent misconfigurations, poor or asymmetric routing, or suboptimal network placement of DHCP servers.
- The failure of a DHCP server due to hardware failure, power failure, natural disaster, nefarious activity, or innocuous human error can cause unreachability.

DHCP deployment approaches to maximizing availability

Most DHCP server deployment strategies face the omnipresent trade-off of budget dollars vs. server quantities. The general goal is to deploy DHCP servers where end users will always be able to obtain these services in a timely manner, while minimizing the total dollars spent on servers deployed and associated server lifecycle expenses. Budget amounts must account not only for server purchases, but for ongoing support and maintenance, which includes server hardware upgrades, operating system (OS) patches and upgrades, as well as DHCP upgrades for new features, bug fixes, or security measures.

DHCP Server Platforms

DHCP servers can be deployed in a variety of platforms from physical hardware servers or appliances, or as virtual servers on a virtualization platform. When we discuss deployment options, we'll generically use the term platform, which can generally be interpreted as either one of these options in each case.

DHCP Servers

The traditional "old school" model for deploying DHCP servers entails deploying a physical server supporting the recommended processing components and operating systems supported by the corresponding DHCP vendor. Often other applications are installed on such servers to maximize hardware utilization.

DHCP Appliances

DHCP appliances are pre-installed DHCP services on secure hardware platforms, typically with a hardened Linux operating system. Appliances are "hardened" in that the base Linux kernel installed on the platform has been stripped of any unnecessary services. This results in a customized kernel and OS that supports only DHCP services (and other services supported by the vendor such as DNS). Underlying file system, users, permissions, and network ports should also be pared down accordingly by the appliance vendor to limit the attack surface.

Appliances offer simplified deployment with one-stop shopping, instead of having to coordinate and acquire server hardware, install the proper OS version and patch levels, then install DHCP services software. Appliances can simplify the ongoing upgrade process by pre-packaging upgrades with compliant OS and services versions with corresponding hardware platforms. Depending on your vendor, these upgrades may be applied from a single centralized console, eliminating the need to physically deploy staff to perform upgrades. In addition, most vendors support centralized monitoring of deployed

appliances, enabling proactive detection of outages or degradations. Of course, appliances generally cost more than general purpose server hardware, and most incorporate open source DHCP services, which are freely available for most leading OSs.

Virtualized DHCP Deployment

Deployment of virtualized DHCP servers enables organizations to instantiate and destroy DHCP services on demand for supported virtualization platforms. Deployment on virtual machines or containers saves on hardware costs, rack space and cooling/power draw, while enabling better segregation and feature improvements over the installation of a DHCP software daemon on a generic hardware server or router. Major appliance vendors offer their appliance products as containers or virtual machines, which combines the hardening, upgrading and other benefits of hardware appliances without the hardware, i.e., as virtualized network functions (VNFs).

Cloud DHCP Deployment

Deployment of virtualized DHCP appliances or containers within a public cloud service is another deployment option supported by some IPAM vendors though not by all cloud services. Most cloud services perform address assignment functions either using DHCP or other means that are native to the cloud services operator. For example, Microsoft Azure prohibits DHCP traffic by blocking IP broadcasts and IP traffic on well-known DHCP ports (UDP ports 67, 68); Azure operates its own DHCP service within Azure vnets to assign IP addresses to VNFs.

DHCP Deployment Approaches

Centralized DHCP Server Deployment

The deployment of DHCP servers generally boils down to a trade-off between wide distribution of a large number of servers "closer" to clients vs. narrow distribution of a fewer number of DHCP servers serving clients from a variety of locations. The extremes of this trade-off consist of a DHCP server on every subnet vs. one or more DHCP servers centrally located serving all of the organization's clients. The key is to balance availability and reasonable performance of the DHCP service between clients and servers while remaining within budget constraints for servers and ongoing management thereof. Your deployment will likely fall between these two extremes.

Figure 1 illustrates a hypothetical centralized DHCP deployment approach scenario where DHCP servers are deployed within regional offices, or the upper tier of a two (or more) tier inter-regional routing network serving subtending branch and remote tiers. This scenario features the deployment of a pair of DHCP servers per region, one functioning as the primary and the other as failover or backup. All DHCP traffic must be funneled to the regional headquarters sites, imposing higher reliance on robust network connectivity to these sites from the respective regions. This architecture also implies the DHCP service platform is sufficiently sized to meet performance and capacity requirements. Note that DHCP primary and failover servers should generally be deployed in separate physical locations for disaster resilience. An outage at one site would not interrupt all DHCP services for a region.

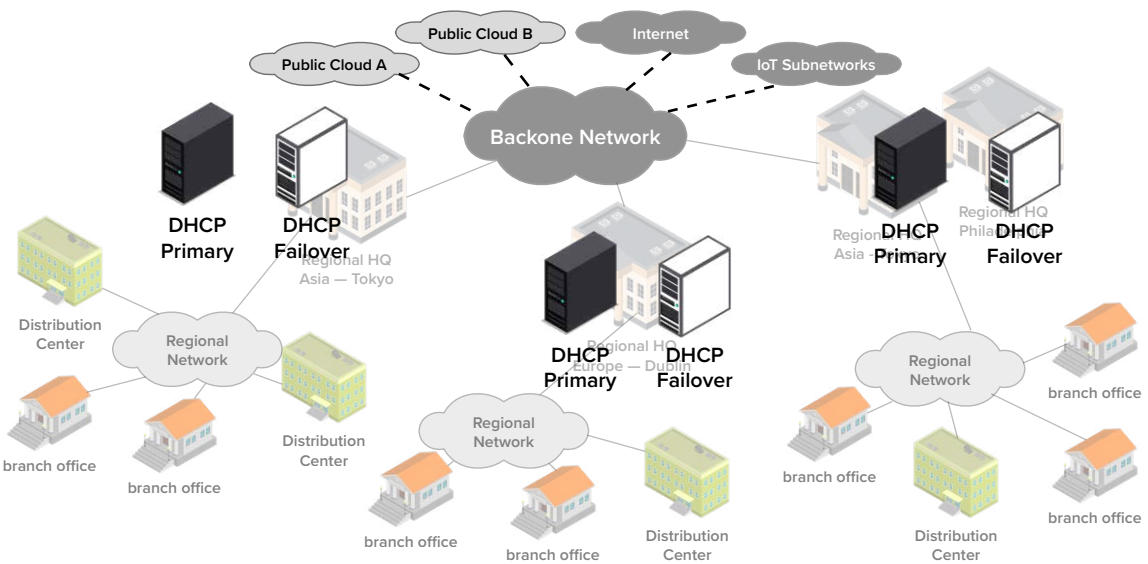

Figure.1: Centralized DHCP Server Deployment Example

Distributed DHCP Server Deployment

At the other end of the deployment continuum, the decentralized deployment approach is illustrated in Figure 2. In this figure, a primary DHCP server is located at [nearly] every branch office and remote site. This localizes DHCP traffic, affording deployment of less stringently sized DHCP platforms. Network connectivity to the regional headquarters is still required however due to the deployment of DHCP failover services there. These services act as failover servers for the servers deployed within each erspective region, though more than one per region may be required for load sharing. Consider the load and redundancy capabilities of your chosen DHCP vendor to identify viable alternative architectures for your network.

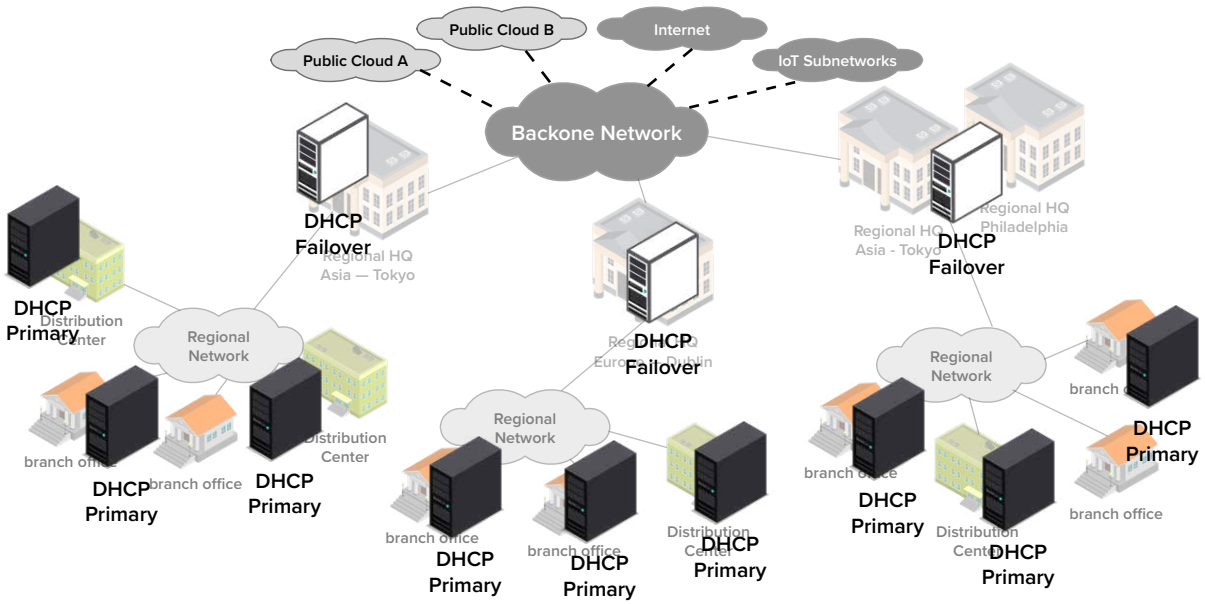

Figure.2: Distributed DHCP Server Deployment Example

Contrasting the two extremes of Figures 1 and 2, the former requires fewer, albeit more powerful DHCP servers and rock solid network connectivity to the regional headquarters sites. The latter requires many more DHCP servers, though of more modest specifications, providing localized services with a network reachable shared backup. You may be wondering, if the network link to a site goes down, what good is having an IP address from a DHCP server? Without a redundant link, other than providing IP access to local network resources, it may indeed be of limited value though even connecting to a local network printer could be challenging. As always, the trade-off must be considered and generally a mixed approach of centralized with at least partial distribution often minimizes overall outage risk.

While many reference DHCP server implementations are single-threaded applications, their performance is usually sufficient for most environments. If you have several thousand DHCP clients attempting to obtain leases at about the same time however, some delays will be likely. If this occurs frequently, you may want to consider deploying additional servers and partitioning finer networks-per-server granularity to reduce the load per server. Again, this is usually not a major concern unless you provide a service that utilizes DHCP to initialize devices like customer premises modems for paying subscribers. After recovery from a neighborhood power outage, devices will come back up and inundate the DHCP server for addresses. In such environments, we recommend you consider a commercial performance-oriented DHCP server.

Relay Agent Configuration

Prepare your routers to support DHCP by configuring the IP addresses of your DHCP servers within your routers' relay agent lists. These lists within each router enable the router to terminate received Discover packet broadcasts, then retransmit them as unicast packets to each configured DHCP sever IP address on its relay agent list. If you partition your network such that address pools for certain subnets are served by a given DHCP server, while those for other subnets are served by another DHCP server, make sure you configure routers serving those subnets accordingly. You could add all DHCP servers to all routers, but this will result in needless relay agent traffic, especially if you have several DHCP servers.

DHCP for IPv6 networks utilizes well-known multicast addresses, obviating the need to configure relay agent lists on routers, though such configuration may alternatively be performed on the relay agent to control which DHCPv6 servers are to process relayed DHCP transactions and not just any DHCPv6 server listening on this multicast address.

DHCP Services Deployment Design Summary

Key considerations when formulating the DHCP server deployment design including the following.

- **Response time requirements** – Most devices tolerate response times on the order of seconds, but certain applications may be more demanding. The more stringent your requirements, the more important will be server performance and perhaps client proximity.
- **Load requirements** – For broadband service providers utilizing DHCP as a customer premises equipment initialization technology, load spikes may occur upon recovery from a residential power outage or equipment installation or reboot. For enterprise environments, such a spike could occur at the start of the work day if several associates arrive at or near the same time, though many devices will simply attempt to renew an IP address previously used by default.
- **Traffic expectations** – The provision of short lease times to minimize overbooking, which causes more frequent renewal attempts, also impacts deployment design. Generally, the shorter the lease time, the shorter the interval between obtaining the lease and subsequent lease renewal attempts. This drives increasing traffic on the network to and from the DHCP server(s) and must be considered when designing to the aforementioned response time and load requirements for server quantities and associated bandwidth.
- **Availability requirements** – do your clients positively have to be able to obtain an IP address or configuration via DHCP 24X7 or is the service "best effort" based? Most will answer that high availability is critical, but with devices growing increasingly multi-networked, as long as one network's address assignment mechanism is available this may be acceptable*. Mean time to repair (MTTR) is another consideration in meeting DHCP services availability objectives. Having a spare server locally or the ability to instantly instantiate a virtual DHCP instance can shorten MTTR (though lease information may not be available) while having to order a replacement would delay this process.

The first three considerations above relate to deploying sufficient quantities of servers of given lease distribution rate to meet respective performance objectives. A good starting point is to identify the number of expected DHCP clients at each site on your network. This number should account for all devices requiring DHCP, including data devices, voice devices, and all IP devices requiring DHCP at each site. Don't forget to account for "peak" quantities of users and devices so that everyone, even associates visiting on temporary basis, may obtain a valid lease.

After accounting for peak quantities of DHCP clients, consider the frequency of DHCP transactions. This will be dependent on your lease times, as well as client lease release configuration. Most clients will "remember" a prior lease and attempt to request it upon power-up, e.g., when an employee returns to work the next day, though this is not always the case.

The fourth consideration listed above relates to providing high availability DHCP services for DHCP clients. Once you've designed your deployment based on performance requirements, total or selective high availability may be planned. Based on server technologies you plan to deploy, implementation of high availability will impact not only the number of servers required, but potentially your address space plan.

Most DHCP servers support a DHCP redundancy or failover protocol such that for a given address pool, one DHCP server will act as the primary, while a second DHCP server will act as the backup or failover server. This basic configuration is illustrated in Figure 3.

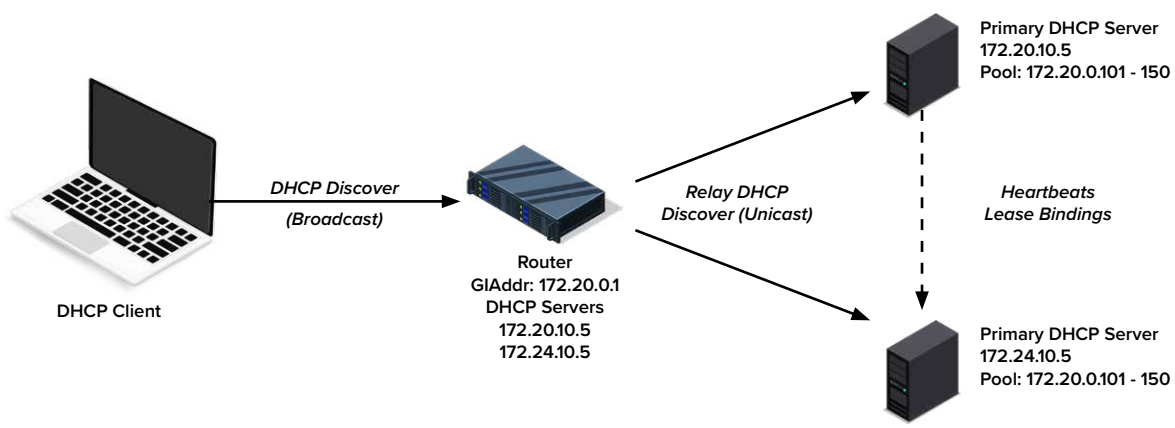

Figure.3: DHCP Failover Configuration

* Of course this statement assumes different DHCP services serve these different interfaces which may not be the case.

Each relay agent must be configured to unicast received DHCP [for IPv4] broadcast packets to both the primary and failover DHCP servers, 172.20.10.5 and 172.24.10.5 in Figure 3. DHCPv6 relay agents can likewise be configured with DHCPv6 server addresses or may utilize the well-known multicast address, ff05::1:3. The DHCP servers utilize a failover protocol such that the primary sends heartbeat messages as well as lease binding information to the failover server. The failover server utilizes user-settable parameters to determine that the primary is down and begins processing the unicast DHCP packets from the relay agent(s). Thus, clients are able to continue receiving IP address and parameter assignments despite the primary server being down. Upon recovery, the primary server obtains the current lease database from the failover server, then assumes its role as primary once again.

Another technique is referred to as a split scopes approach, which entails deploying two DHCP servers with complementary address pools, not the same address pools. In this way, either server can process DHCP transactions without worry of duplicate assignment. In Figure 4, we illustrate the splitting of 50-address pool 172.20.0.101-150 into two non-overlapping pools and deploy each on two different DHCP servers, respectively. Like the failover configuration in Figure 3, each relay agent needs to be configured with both DHCP server addresses on which the split scopes are provisioned. Both DHCP servers should receive all DHCP transactions within the given subnet and provide a lease if capacity exists. Thus, clients on the subnet should have access to the same address pool, albeit split across two DHCP servers in this case.

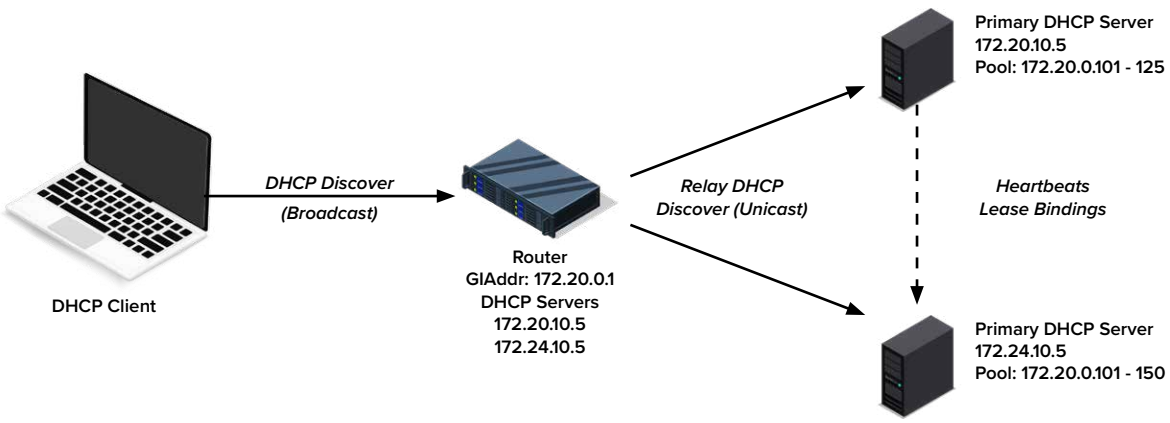

Figure.4: Split Scopes Configuration

Since both servers are required to meet the capacity needs, you may end up with an inability to meet IP address demands should one fail. Another alternative is to allocate double the number of addresses such that each DHCP server is configured with 100% of the required capacity. In this manner, each server can handle the capacity needs should one fail. Referring to Figure 4 for this scenario, one DHCP server would be configured with address pool 172.20.0.101-150 and other server with 172.20.0.151-200 to provide the full pool capacity, fifty IP addresses, on each server. This solution provides full redundancy at the expense of doubling the required address space.

DHCP Deployment on Edge Devices

Most router and access point products provide a DHCP service as a component of their router platforms. This may lead one to question whether a separate server is needed to support DHCP services. As with most design questions, the answer is, «it depends.» Small environments with a few sites with local routers serving up to 100 or so monolithic clients each may be well served by configuring the router to provide DHCP services. However, larger organizations or those requiring more advanced DHCP services, e.g., for discriminating voice vs. data clients for address and option parameter assignment, would be better served deploying discrete (non-router-integrated) DHCP servers.

The advantages of running DHCP on a router device include:

- **Lower hardware cost** – no need to procure a server or set of servers
- **Single user interface** – the same command line interface can be used to configure the router and the DHCP server, and no relay agent configuration is needed
- **« Fewer moving parts»** – one less communication link and server required to perform DHCP functions, which in general can increase the overall solution reliability

The main disadvantages of running DHCP on a router are:

- **Options support** – most router based DHCP servers are primitive, supporting address assignment but little in the way of options support.
- **Client class support** – major vendors do not support client classes, which is required for discriminatory address/option assignment to different devices, e.g., VoIP vs. data devices.
- **No failover** – if a router fails, you've probably lost connectivity in any case but if there are two routers serving a subnet for redundancy a split scopes approach would have to be employed, increasing management complexity.
- **No centralized management** – router based DHCP services are configured via command line and unless a centralized tool is employed, each router DHCP server must be configured manually with respect to the IP addressing plan; less likely support is possible if multiple router vendor products are in use.

Of course, if your edge device or router supports a virtualization platform, you can deploy a centrally managed virtual appliance or container to attain all of the advantages of edge DHCP support with none of the disadvantages!

DHCP Server/OS Threats

As with all network servers, vulnerabilities within the server operating system (OS) and applications may be exploited by attackers in order to severely hamper or crash the server. These attacks can be of the following forms:

- **Hardware** – Physical access to DHCP servers enables the attacker to unplug, disconnect or physically remove the server, literally removing the server from service, thereby reducing the availability of the DHCP service and possible capture of configuration information. Physical removal of a server affords the attacker an opportunity to hack the server for IP address and associated configuration parameters that may relate to broader network details such as NTP, DNS or FTP server IP addresses.
- **Operating System attacks** – An attacker may attempt to gain local or remote console access to the server by hacking passwords or overflowing the code execution stack or buffer. In general, an attack may exploit a known vulnerability of the operating system or version of DHCP software running on the hardware server or virtualized platform.
- **DHCP service attacks** – An attacker may attempt to exploit a known vulnerability for a given vendor and version of DHCP server software running on the victim server to shut it down or otherwise corrupt and/or disrupt service.
- **API channel attack** –The DHCP server may support an API interface, which if exploited, could provide a convenient mechanism to remotely control or configure the DHCP server. Such power may entice an attacker to attempt to access the API channel to perform nefarious functions such as stopping the DHCP service thereby denying IP address assignment services to network clients.

DHCP Server/OS Attack Mitigation

The following approaches may be employed to defend against DHCP server attacks:

- DHCP server host access controls including ACLs, identity/password access, encrypted transport, and least privilege permissions
- DHCP server operating system hardening to prevent access via unauthorized protocols
- Monitor security advisories (e.g., CERT) for operating system or DHCP service vulnerabilities and keep systems updated to prevent exploitation
- Protect the DHCP service API channel using available controls such as ACLs, authentication, and encryption.
- Monitor DHCP server configuration for changes to detect unauthorized changes; this could consist of periodically checking configuration file checksums and to verify changed values with authorized changes.
- Physical access controls to constrain access to data centers or rooms housing DHCP servers.

DHCP Service Threats

Within enterprise environments, most threats to DHCP services are posed by internal (i.e., intra-organizational) clients. DHCP servers should not be reachable by external clients by simply not deploying DHCP servers on external subnets nor relaying DHCP packets from external sources. For service providers that initialize subscriber devices using DHCP, whether cell phones, cable or fiber routers, etc., threats to DHCP services can originate externally to the network due to the necessity of exposure. In short, all organizations using DHCP are vulnerable. The degree of vulnerability and the impacts of compromise should drive the response in the form of securing DHCP to minimize such impacts.

Like all network services, DHCP is vulnerable to traffic sniffing and denial of service (DOS) attacks. Those with access to the data path over which DHCP packets traverse have the ability to sniff DHCP traffic for use in identifying network devices, their IP addresses and other parameters. A DOS attack involves an attacker flooding a given server with requests too numerous for the server to handle, so the server spends all of its cycles attempting to deal with the flood and not on legitimate client requests; thus leaving these legitimate clients unserved, and thereby denying service.

Another type of DHCP attack involves a rogue client attempting to obtain a valid IP address and configuration to access the network. This could be malicious, e.g., theft of broadband service, or merely accidental, e.g., a visitor plugging into the wall jack in the conference room. But this relates to broader network security so we'll cover this later in this paper.

A third form of attack features a rogue DHCP server which responds to lease requests from clients with an invalid or inappropriate IP address and/or option parameter information. This "man in the middle" type of attack may attempt to set improper configuration parameters on the client, such as the default gateway or DNS server address(es) to use. Note that with IPv4, a rogue DHCP server attack is generally only applicable when the server is on the same subnet as the client; relay agents should be configured to relay DHCP packets to authorized DHCP servers. A remote rogue DHCPv6 server may be reachable via the DHCP servers well-known multicast address.

The client may receive Advertisements or Offers from both the legitimate DHCP server(s) and the rogue server. Many clients will select the first such offer that includes its requested parameters. If the rogue server is on the same subnet as the client, and legitimate servers are not, then it's likely the rogue server may be able to specify the IP configuration of the client.

DHCP Threat Mitigation

For an enterprise network, protection against DOS attacks could be implemented via packet rate limiting, e.g., via iptables albeit in a reactionary fashion, but DOS protection should be considered in a broader context comprising all key network services, not just DHCP. Other potential targets within an organization including DNS servers or web servers imply that a gateway-based packet filtering approach be considered to protect all servers with a common solution. Such a solution typically involves packet analysis and pre-defined thresholds to limit the number of outstanding packets in process, though care must be taken with DHCP since most clients' transactions are funneled through DHCP relay agents, concentrating packets from a given set of source addresses.

Service providers may experience what appears to be a DHCP DOS attack when a power outage is restored and thousands of residential customers, i.e., DHCP clients, power back up and request IP addresses. Many large service providers deploy carrier-class DHCP servers in a manner that load balances traffic across multiple servers in such a scenario.

Mitigation steps for the threat of unknown DHCP clients accessing the IP network by illicitly obtaining an IP address from DHCP requires identification of clients based on various network access control techniques we'll discuss later.

Rogue DHCP servers may be difficult to detect, especially for clients on the same subnet as the rogue server. Periodic IP address sweeps or discoveries can help identify rogue devices including illicit DHCP servers. And both ISC and Microsoft implementations provide means to mitigate rogue servers. For ISC, use the authoritative directive, which configures the server to issue a negative acknowledgement if a client requests a lease for an address for which the server is authoritative yet for which the server has no record. Microsoft requires DHCP servers to be authorized within Active Directory; thus when a Windows DHCP server boots it verifies its authorization in Active Directory before processing DHCP packets.

DHCP Authentication and Encryption

To mitigate snooping and rogue device admittance, authenication and encryption techniques may be used. Unfortunately, no major DHCP clients natively support such functions so these mitigations may be challenging to implement. Nevertheless, we'll discuss them here for completeness.

The DHCPv6 protocol supports authentication and encryption of DHCPv6 messages between relay agents and servers through the use of standard IP security archtiecture (IPSec) in accordance with RFC 8213. Encryption effectively combats traffic sniffing while authentication prevents imposter relay agents from interfacing with a given DHCP server. RFC 8213 also applies to DHCPv4 as a means to secure relay agent-to-DHCP server communications.

In a nutshell IPSec enables authentication and data integrity validation through the inclusion of an additional Authentication Header (AH), which enables a recipient to validate the identity of the sender and the message integrity. The Encapsulating Security Protocol (ESP) also supports authentication and data integrity validation but adds the ability to encrypt data as well. Both AH and ESP process packet data using a pairwise shared secret key or key shared via a key exchange mechansim such as Internet Key Exchange (IKE).

RFC 8213 describes security for the relay-server link, though the IETF has also defined an end-to-end (client-server) DHCP authentication mechanism for IPv4 in RFC 3118. This scheme provides simple validation of the sender of DHCP packets via the use of shared tokens or keys. A token is simply a fixed value that is inserted into the DHCP Authentication option field. The receiver of the packet examines the token and if the token matches its configured token, the packet is accepted; otherwise it is dropped. This method provides weak endpoint authentication and no message verification. The use of shared keys can provide stronger endpoint authentication with message verification. However, shared keys must be configured on each client, with each client's key configured on each DHCP server through which the client obtains leases. The DHCP Authentication specification does not define the mechanism for key distribution. Mobile clients for example would need to be configured with tokens for each DHCP server with which they may interact and vice versa.

The client creates an HMAC-MD5 hash of its Discover packet and signs it using the shared key. The resulting digest is placed in the DHCP Authentication option and transmitted within the Discover packet to the server. For the purposes of the hash computation, the hash portion of the DHCP Authentication option must be set to zero. The DHCP server would then compute a hash of the received message utilizing the shared key associated with the client (identified by the secret ID field of the DHCP Authentication option). The server zeroes out the hash value, hops and GIAddr fields for the purposes of the hash computation. If the calculated hash matches that transmitted in the original DHCP Authentication option, the client and the contents of the packet are considered authenticated. The DHCP server utilizes the same shared key to compute the hash value of its DHCP Authentication option when it prepares its Offer and future packets to the client.

As previously mentioned, there have been very few implementations of DHCP Authentication. The challenges of key distribution and management as well as processing delays due to hash computation have been deemed too heavy a price to pay for the perceived benefits. Security of the DHCP service then typically falls on DHCP server administrators to monitor networks and servers and react to incidents as they occur.

Securing Network Access with DHCP

Network access control or "NAC" is a broad term which comprises a set of technologies which seek to identify who is attempting to access your network prior to providing such access. Various techniques are available offering various levels of access control. We'll start by analyzing DHCP-based access control, which admittedly is among the weaker approaches to network access control. We'll then touch on more wide-reaching techniques.

Discriminatory Address Assignment with DHCP

Let's focus first on DHCP services and some approaches to implement discriminatory address assignment. There are several levels of policies or controls most DHCP solutions provide for discrimination of "who's asking" for an IP address via DHCP. The first is to simply filter requests by an available form of client identifier such as the MAC address of the client requesting an address. The MAC address is found in the client hardware address (chaddr) field of a DHCPv4 packet. DHCPv6 device identifiers consist not of an interface's MAC address but the Device Unique ID (DUID) and Identity Associations (IAs) which identify each client and interface respectively.

If the DHCP server has a list of acceptable (and/or unacceptable) device identifiers, it can be configured to provide a certain IP address and associated parameters to those clients with an acceptable identifier, and either no IP address or a limited function IP address to those without an acceptable device identifier. By limited function IP address, we mean that the network routing infrastructure is pre-configured to route IP packets with such source IP addresses (source address dependent routing) to only certain networks, such as to the Internet only. An IP packet with source address A may be routable across the enterprise while one with source address B may be routable only to the Internet for example.

The vendor class identifier, user class identifier or even the option request option, which some vendors refer to a "DHCP fingerprint," can be used to discriminate clients. If your DHCP server can be configured to recognize particular option values provided by the client, then you can selectively configure the client's IP address and configuration parameters. Addresses can be assigned from a certain pool and/or additional configuration parameters can be provided to the client via standard or vendor-specific DHCP options.

Another level of discriminating IP address assignment is possible by authenticating the user of the machine requesting an IP address. This function can be used in conjunction with discriminatory address assignment. For example, if a client with an unknown or unacceptable device identifier attempts to obtain an IP address, one option is to completely deny an address; another option is to require the user of the client to login via a secure access web portal page.

This enables easier capture of new device identifiers for legitimate users of your network. Solutions ranging from perl scripts such as NetReg, to sophisticated integrated software solutions are available to direct such users to a login/password requesting webpage. A simple lookup against a database of legitimate users then allows access or denial of the client to a production IP address. These systems typically work in accordance with the following packet flow in Figure 5.

Walking through this flow, the process begins with a device connecting to the network, attempting to obtain an IP address via DHCP. The DHCP server, employing device identifier or client class type filtering, determines if the device is a known user device*. If the device is known or has otherwise already authenticated, the DHCP process may continue with an Offer/Advertise for a production IP address, followed by a Request and Ack/Reply. However, if the device is not known or is required to authenticate, the DHCP server can still provide an IP address though the IP address assigned in this case would be a captive portal, walled garden or quarantined IP address.

These terms refer to the fact that the IP address assigned to the client will only be routed to the subnet or VLAN that has the authentication web server and associated servers running. This quarantined VLAN enables IP communications but only to this restricted set of devices. This approach seeks to cordon off the device from infiltrating the rest of the network until the corresponding user can be authenticated. The routing infrastructure must be configured to route packets with a source address from the quarantined address pool to the quarantined VLAN and/or the client must be configured with the classless static route option. Thus, walled garden address X as shown in the figure above is a member of the quarantined VLAN, on which only limited network resources are available. Figure 6 illustrates an example network topology of this captive portal configuration.

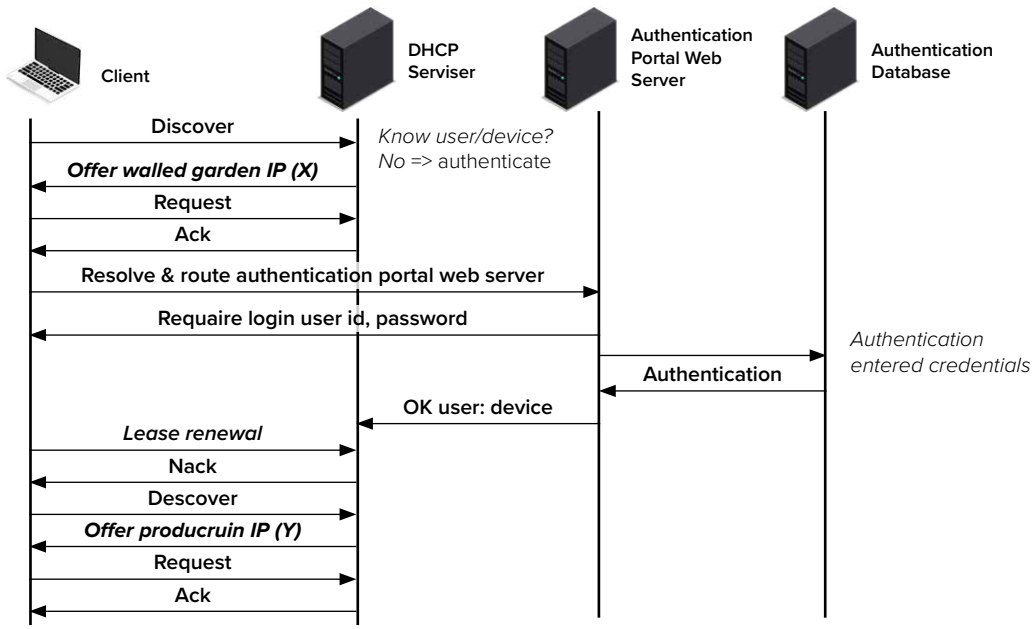

Figure 5: Basic DHCP Captive Portal Flow

Now when the user opens up a web browser, he/she can type in any web address. A limited configuration DNS server is required on the quarantined VLAN; limited in the sense that it will resolve any and every query to the IP address of the authentication web server. Thus, no matter what website address is entered in the web browser, the web address is resolved to the captive portal web server. The authentication web server presents the login page. You may have seen something similar to this if you travel and use a hotel's broadband or wireless service. Once the requested credentials are entered, which for an enterprise environment, would typically comprise a user ID and password, the web page can pass the entered credentials to a back-end database for authentication.

Based upon the results of the authentication, the requesting device would then be deemed authorized or not, and if authorized, optionally at a particular service class. The service class of authorization provides more granularity than a simple boolean "authorized or not," where different authorized users can be assigned a different production IP address, which in turn can provide access to different network resources. For example, basic level users may be granted access to a basic set of resources, while advanced level users may be granted access to additional resources, e.g., IT resources. Once again this requires the routing topology be configured with multiple source-routed or VLAN segments, with these networks and corresponding routing plan mapped to DHCP server configurations in terms of associating address pools with service levels.

The manner in which the production IP address is assigned follows from expiration or denial of renewal of the quarantined IP address. The quarantined IP address lease time is generally configured as a short lease time (~1-5 minutes). This enables the device to attempt to renew quickly. Should the device still be in the process of authentication, its renewal attempt would be ACK'd, extending the lease. Once authentication is completed successfully, the authentication system updates the DHCP server to add the device identifier to the "known" or "allow" pool. The renewal attempt for the quarantined address would then be NAK'd, enabling a fresh DHCP process to provide a "production" IP address (address Y in Figure 5). Should the device fail authentication, the renewal can be NAK'd and subsequent address attempts denied; alternatively, the quarantined address renewal attempt can be granted in order to provide access only to resources on the quarantined network if desired.

Beyond these device and user identification measures based on device identifiers, this general flow can also provide additional validation on the machine requesting the IP address. The DHCP process can be used to invoke an external security scanning system like openVAS or another third-party application to scan the requesting client for viruses or to validate use of acceptable virus protection software. This device scanning step can be used alone or in conjunction with the device identification measures to provide a robust access security solution via DHCP.

One example network configuration for DHCP based secure access is depicted in Figure 6. The DHCP server shown in the diagram would be configured with a number of client class sets. We refer to the client class as the matching criteria in the DHCP packet, which links to the associated network accessibility. For example, we would need a client class set for at least each of the following in our example:

* In some cases, even known user devices may require periodic re-authentication as a security precaution.

Securing Network Access with DHCP

Network access control or “NAC” is a broad term which comprises a set of technologies which seek to identify who is attempting to access your network prior to providing such access. Various techniques are available offering various levels of access control. We’ll start by analyzing DHCP-based access control, which admittedly is among the weaker approaches to network access control. We’ll then touch on more wide-reaching techniques.

Discriminatory Address Assignment with DHCP

- Captive portal network (Remediation VLAN)
- Production network 1 network
- Production network 2 network

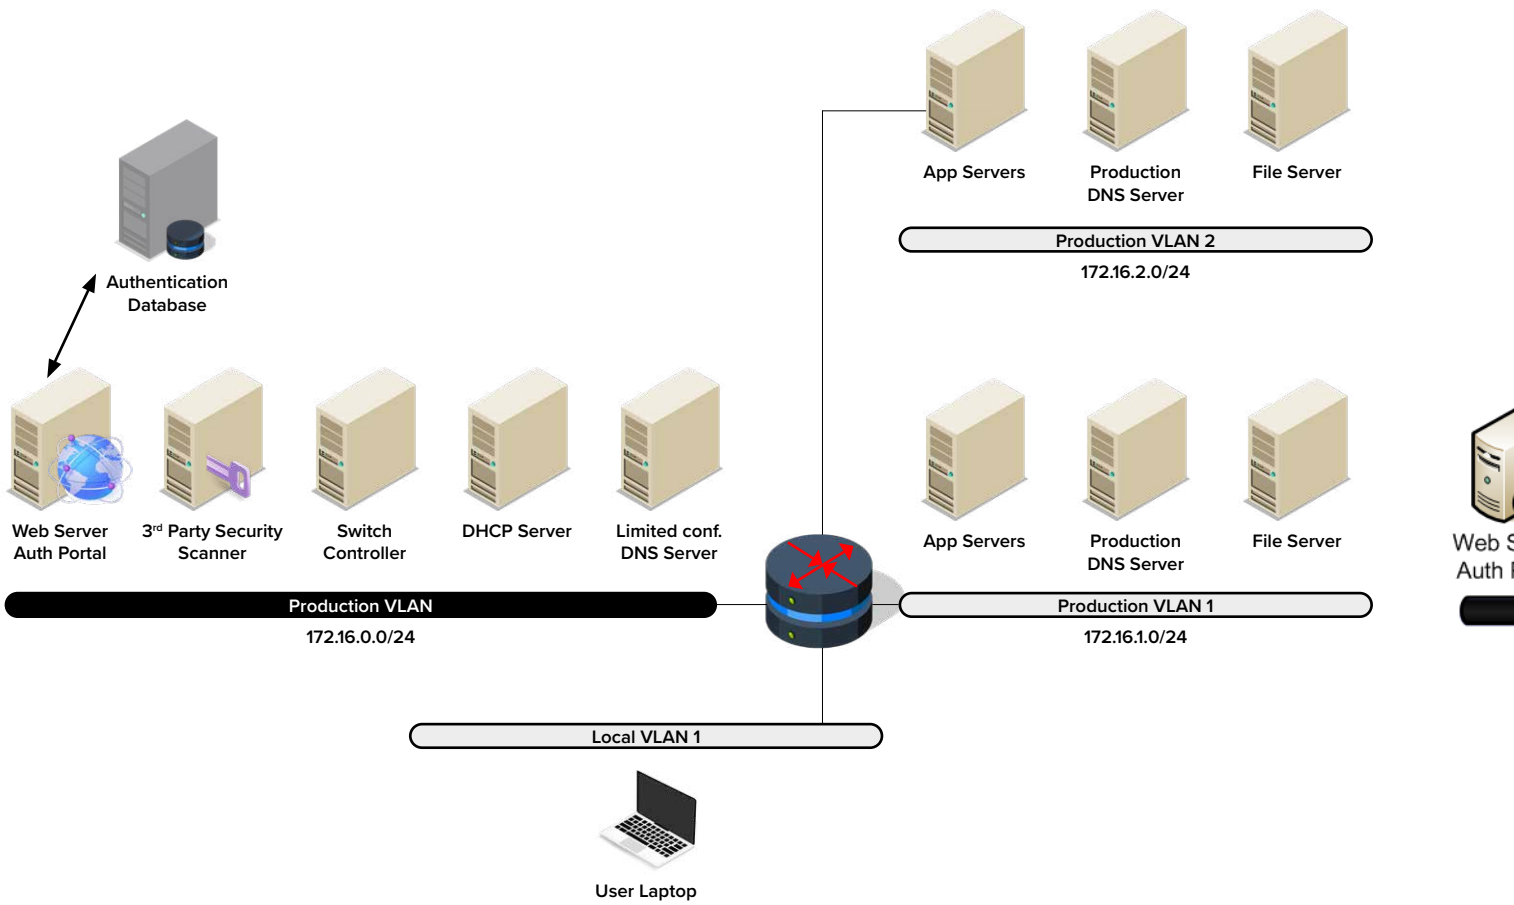

Figure 6: Captive Portal Network Diagram

Think of these client class sets as bins into which individual clients are placed based on the linking of their authentication state to the device’s client class. Thus, client class members would be categorized by the DHCP server in accordance with defined client classes as they appear on the network and users authenticate. These client classes would generally map to pool definitions on the DHCP server as shown in the following simple example ISC server configuration . Note that additional options can be defined for each of the pools to provide additional configuration granularity to clients falling into each set or pool.

```
subnet 172.16.0.0 netmask 255.255.252.0 {
# subnet level options here...
        pool{                                   #captive portal pool
                range 172.16.0.10 172.16.0.254;
                option domain-name-servers 172.16.0.5;   #limited config DNS server
                default-lease-time 150;          #short lease time
                allow unknown clients;           #clients not predefined.
        }
        pool {                                  #Prod Net 1
                range 172.16.1.10 172.16.1.254;
                option domain-name-servers 172.16.1.5;   #production DNS server
                default-lease-time 14400;        #normal lease time
                deny unknown clients;            #clients must be predefined.
                allow members of "net1";         #client class net1 allowed
```

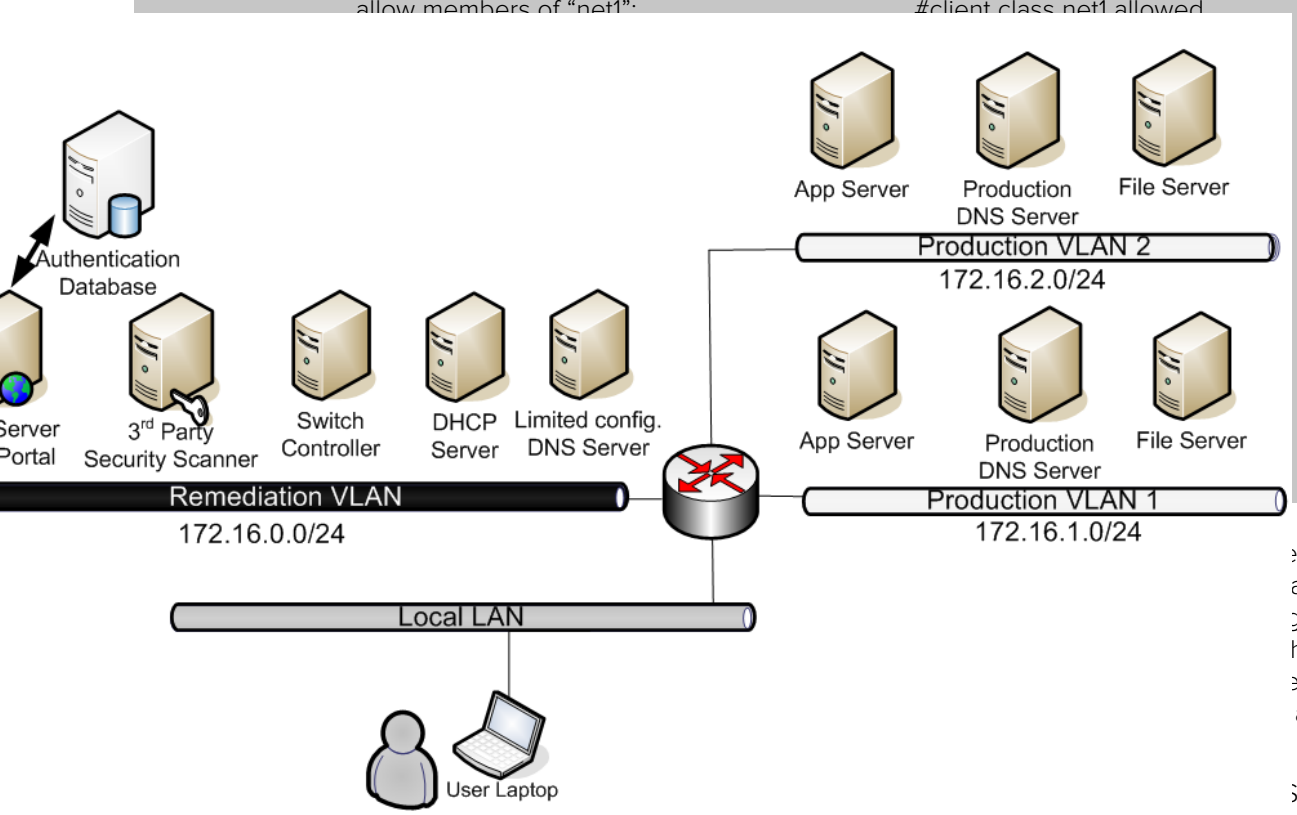

e DHCP configuration
access to production
group for production
he ISC DHCP server
eclaration can define
along with any other

SC DHCP server. The captive portal network (the remediation VLAN in the figure) is deployed including the limited configuration DNS server, web server as the authentication portal, with access to an authentication database, and optionally a security scanning server and any other required pre-access services.

More than one DHCP server may be deployed for high availability and/or for scaling for larger networks. This approach does complicate things, as the DHCP server configurations need to be consistent on both servers to route unknown clients or clients requiring authentication to the captive portal net.

DHCP Lease Query

To verify alignment of authorized DHCP assignments with actual addresses attempting to initiate IP connections, DHCP lease query can be used. Given most or all addresses on a subnet are configured using DHCP by policy, each IP address should have a corresponding DHCP lease. The DHCP Lease Query is a DHCP protocol message that enables an edge router to query the DHCP server regarding the lease status of a particular device or set of devices. This provides some assurance that a device attempting to communicate via the router has not spoofed an address that should have been assigned by the DHCP server.

When the router receives IP traffic within a layer 2 frame from a particular device for example, it can issue a DHCPLeaseQuery message to its configured DHCP servers (i.e., in its role as relay agent) to determine the state of a DHCP lease, querying by IP address, DUID or MAC address. If a DHCP server had previously provided a lease for the client, it will respond to the router, and the router will give the green light and route the device's packets. If not, the device does not have a lease and the router can drop the device's packets. The router can cache this information as well to constrain the Lease Query rate. Of course, this form of access control applies only when all clients on a subnet use DHCP such as in broadband access networks, not when other statically addressed devices communicate on the subnet.

Alternative Access Control Approaches

You may be thinking that the DHCP-based approach is fine for clients utilizing DHCP; but what about those "clever users" who figure out the subnet address, then manually encode a static IP address on their machines to access the network? These clever users may after all be those of most concern from a secure access perspective. In addition, for devices using SLAAC, no DHCP interaction is required for address assignment.

There are two basic alternative approaches for enabling detection and associated remediation action of devices without relying solely on the DHCP-based approach.

- **Layer 2 switch alerting** – as a device boots up, it will synch up its Ethernet layer (layer 2), which its connected switch can detect its literal initial network access. If your switch can alert on such "link up" state transitions, a system could configure the switch to connect the corresponding switch port to an authenitcation VLAN, similar to the operation described previously for DHCP client classes. As the device authenticates and the user authorized, a command to the switch would be needed to place the port on a production VLAN.
- **802.1X – IEEE 802.1X** is the foundation of most network access control systems, including those used for zero trust networking. This protocol specificatoin enables switch (or generically, edge device) capture of new access attempts with the use of Radius authentication and dynamic switch port configuration as described above. This scheme requires each client to be conifugred with an agent called a supplicant, which interacts with the network autehtnication server.

Diamond IP DHCP Security Capabilities

Diamond IP offers a rich set of DHCP security features and capabilities to enable you to implement these mitigation approaches within a centralized, holistic IP address management solution. These are summarized as follows:

- Our virtual and hardware Sapphire appliances are built from scratch to mitigate operating system, server, or poisoning attacks
    - Non-commercial Linux distribution built from scratch in a secure environment with a non-modular kernel, uninterruptible boot, and protections against networking attacks such as spoofing, route and ICMP redirections, and more. The file system includes only necessary binaries which run in a sterile jailed environment and have non-privileged attributes. Please contact us for a full summary of Sapphire appliance security measures.
    - Sapphire DNS appliances support standard DNS query logging to log collectors or full DNS query and response capture with transmission to our Sapphire A30 IPAM Auditor appliance.

  - Sapphire DHCP appliances support multiple resiliency features including hardware clustering (TwinMirror), split scope and DHCP failover server deployments.
  - Sapphire DHCP appliances support SNMP MIBs and traps to view and report address pool capacity exhaustion, alerting administrators to a possible threat or to supplement address capacity.

- IPControl software or Sapphire Executive (EX) appliances provide centralized configuration, monitoring and management of deployed Sapphire DHCP appliances.
  - IPControl provides a web graphical user interface (GUI) to configure all Sapphire, ISC or Microsoft DHCP and DHCPv6 attributes on all deployed DHCP servers, including pools, shared subnets, manual DHCP ("reservations"), options, polices, client classes and more.
  - IPControl provides threshold and alert definitions to enable administrators to be notified in advance of address pool exhaustion. Our prediction models facilitate planning and can dictate urgency for proactive actions. Beyond notification, alert conditions can trigger automated actions, such as provisioning of additional address space for added capacity.
  - IPControl supports centralized staging and distribution of updates and patches to deployed Sapphire appliances.
  - IPControl's Appliance Dashboard provides a centralized summary of each deployed appliance's service status and enables drill-down for appliance level configuration and diagnostics.
- The Sapphire A30 IPAM Auditor appliance provides visibility to and reporting for DHCP transactions on your servers.
  - The IPAM Auditor appliance aggregates inputs from deployed Sapphire DHCP appliances to enable graphical reporting of all levels of DHCP traffic from time-series summary trends to drill down to specific DHCP packets for forensics analysis.
  - Graphical dashboards for appliance and DHCP statistics provide easily consumable DHCP information to enable rapid detection and investigation of issues.
  - The IPAM Auditor also aggregates DHCP and IPAM data for full DDI reporting.

Conclusion

As a critical network service that your clients relay on to access your network, DHCP services must be highly available, performant and resilient. While typically not exposed beyond the confines of the enterprise network, erroneous or nefarious DHCP activity can create a nuisance if not an outage for clients attempting to access your network. By implementing protections against most common risks along with reporting to provide visibility to DHCP activity and capacity, you can reduce the risk of IP address unavailability. Diamond IP offers a rich set of security features for monitoring and protecting your DHCP services.

Toll Free: **(844) 442-9462**
International: **+1 (305) 501-2430**
Fax: **+1 (305) 501-2370**

Sales: **sales@cygnalabs.com**
Support: **support@cygnalabs.com**
Billing: **finance@cygnalabs.com**

cygnalabs.com

CYGNA LABS