



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Dynamic Host Configuration Protocol:
Security Implications and Possible Safeguards

Matthew P. Harvey

05 February 2004

Submitted in partial fulfillment of the requirements for
the SANS GIAC Security Essentials Certification

Abstract

The Dynamic Host Configuration Protocol is a very widely-used protocol that provides a fundamental service in many IP networks, both large and small. The service provided by DHCP is a critical one since the settings provided, IP address, default gateway, DNS server address, and the like, define how hosts communicate on the network. If this service can be compromised, the possibilities for further exploitation by a patient and knowledgeable attacker are nearly limitless.

Because DHCP runs over UDP and because one side of all UDP communication does not have an IP address during the conversation, DHCP is an inherently insecure protocol. A hostile host can cause significant interference or compromise on a network by acting in the role of an illicit host or server (or both).

While current DHCP implementations have many vulnerabilities, there are many current proposals for improving the security of the protocol. Furthermore, there are several steps that can be taken using the current state of DHCP to operate in a more secure fashion. This paper discusses the security limitations of DHCP and several present and future steps for security the protocol.

© SANS Institute 2004, Author retains full rights.

The DHCP Protocol – Its Purpose and Functioning

The Dynamic Host Configuration Protocol is a very widely-used protocol that provides a fundamental service in many IP networks, both large and small. DHCP is a reliable protocol for on-demand remote configuration of clients on an IP-based network. DHCP was defined in RFC 1531 in October of 1993. The current authoritative RFC is RFC 2131, dating to March of 1997. Today the majority of IP devices include configuration by DHCP as an option, and obtaining IP addressing information by DHCP is the default setting for most desktop operating systems.

The primary purpose of DHCP is to allow IP configuration information to be passed to hosts on an on-demand basis. This allows an unconfigured host to be attached to a network and to obtain a valid IP address and other basic configuration information. The most common pieces of information provided by DHCP are IP address, subnet mask, default gateway, and DNS server, but there are hundreds of other options which may be set. Furthermore, there is a

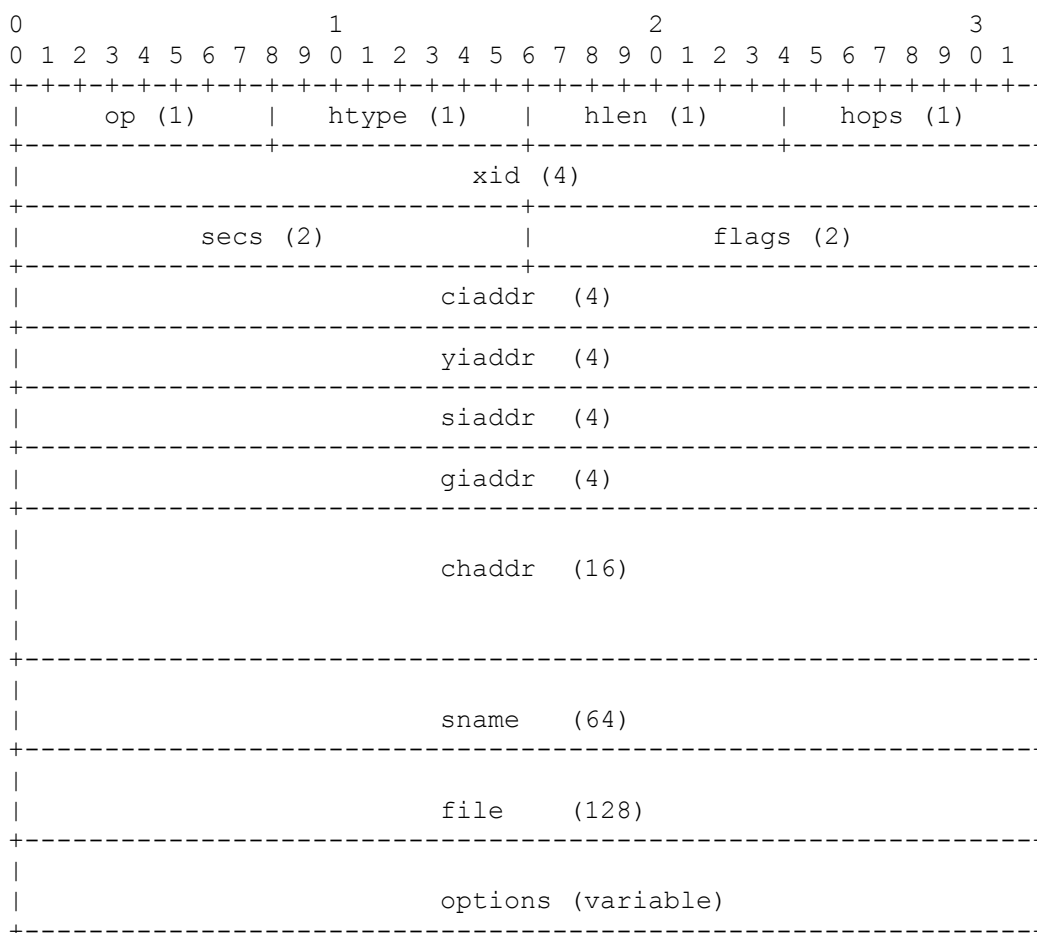


Figure 1: Format of a DHCP message (Droms, RFC 2131, 1997)

FIELD -----	OCTETS -----	DESCRIPTION -----
op	1	Message op code / message type. 1 = BOOTREQUEST, 2 = BOOTREPLY
htype	1	Hardware address type, see ARP section in "Assigned Numbers" RFC; e.g., '1' = 10mb ethernet.
hlen	1	Hardware address length (e.g. '6' for 10mb ethernet).
hops	1	Client sets to zero, optionally used by relay agents when booting via a relay agent.
xid	4	Transaction ID, a random number chosen by the client, used by the client and server to associate messages and responses between a client and a server.
secs	2	Filled in by client, seconds elapsed since client began address acquisition or renewal process.
flags	2	Flags (see figure 2).
ciaddr	4	Client IP address; only filled in if client is in BOUND, RENEW or REBINDING state and can respond to ARP requests.
yiaddr	4	'your' (client) IP address.
siaddr	4	IP address of next server to use in bootstrap; returned in DHCPOFFER, DHCPACK by server.
giaddr	4	Relay agent IP address, used in booting via a relay agent.
chaddr	16	Client hardware address.
sname	64	Optional server host name, null terminated string.
file	128	Boot file name, null terminated string; "generic" name or null in DHCPDISCOVER, fully qualified directory-path name in DHCPOFFER.
options	var	Optional parameters field. See the options documents for a list of defined options.

Table 1: Description of fields in a DHCP message (Droms, RFC 2131, 1997)

provision in the protocol that allows companies to develop their own vendor-specific extensions (Droms, RFC 2131, 1997).

The basic process of client receiving configuration information from a DHCP server works is a four-part interaction. The client initiates the process with a DHCPDISCOVER packet, which it broadcasts to its local subnet. This packet MAY include a suggestion for a specific IP address the client would like to be assigned (usually the one it was assigned most recently) and a desired lease duration. If the network is so configured, DHCP relay agents present on the local subnet may forward this request to DHCP servers on other subnets. The client's MAC address is included in the 'chaddr' field, and this is used to uniquely identify the client throughout the conversation (Droms, RFC 2131, 1997).

Each DHCP server will respond with a DHCPOFFER message, assuming it is configured to do so and has some currently available addresses to offer. This message includes in the 'yiaddr' field the IP address that the server is offering to the client, as well as any other configuration options recommended, which would

be carried in the 'options' field. At this point the server may or may not reserve the address in its internal repository of addresses; doing so cuts down on potential conflicts but makes the system more vulnerable to denial of service attacks, as we will see below. Before offering an address, the server should ping the offered address to confirm that it is not in use, although this option may be disabled by the administrator. The DHCPOFFER may be relayed through a DHCP agent as necessary (Droms, RFC 2131, 1997).

The client receives any DHCPOFFER messages, and may wait a defined period of time so as to receive multiple offers. The client will choose one of the offers to accept; this choice may be based on the parameters offered, which usually means that if the client requested a specific address in its original DHCPDISCOVER message it will prefer to accept an offer of that address. Otherwise, and most often, the client will accept the first offer which it receives. It accepts by broadcasting a DHCPREQUEST message, setting the 'server identification' option to that of the server whose offer wishes to accept. It is required to set the 'requested IP address' option equal to the 'yiaddr' value from the server's DHCPOFFER. If the client receives no DHCPOFFERS within a defined period of time it will time out and broadcast a new DHCPDISCOVER (Droms, RFC 2131, 1997).

DHCP servers receive the DHCPREQUEST from the client. Any server other than the selected server treats this message as notification that the client will not be using the address it offered and de-allocates the address as necessary. The selected server will commit the allocated address in its repository and respond with a DHCPACK message containing the same 'yiaddr' and configuration information as in the DHCPOFFER. If for some reason the offered address is no longer available, or if there is some other error in the client's DHCPREQUEST, the server replies with a DHCPNAK. If the server does not receive any DHCPREQUEST in a defined period of time, any address allocated at the time of the DHCP offer should be marked as available (Droms, RFC 2131, 1997).

When the client receives a DHCPACK it should perform a final check by ARPing for the allocated network address. If the address is not in use, the process is completed and the client begins using the configuration information provided. If the address responds to the ARP, the client sends a DHCPDECLINE message and restarts the process; a ten second wait is required before restarting the process in order to avoid excessive traffic. DHCP servers receiving a DHCPDECLINE must mark the relevant address as "not available" (Droms, RFC 2131, 1997).

When the IP address is no longer required, the client may send a DHCPRELEASE message to the server. The client unbinds the address and the server marks the address as available for assignment, but the server should keep a record of the client's configuration info for use in future interactions with that client (Droms, RFC 2131, 1997).

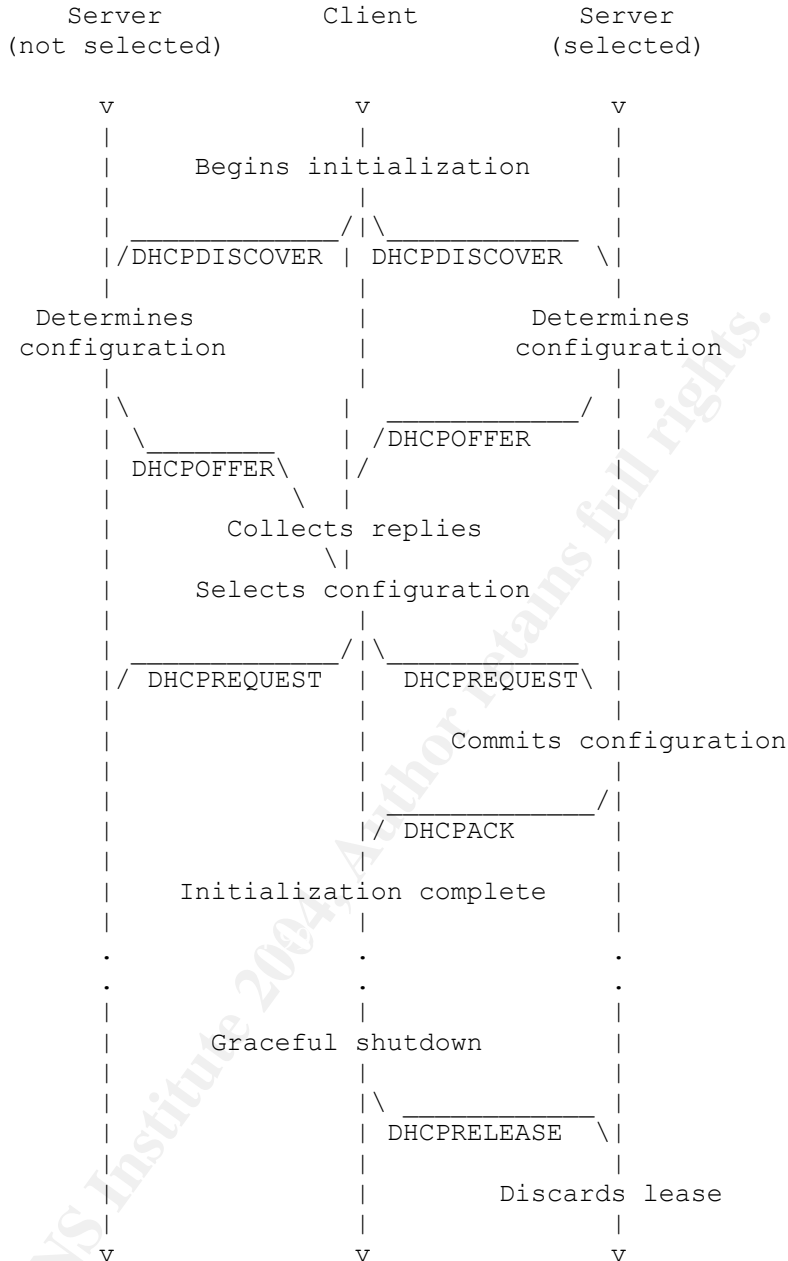


Figure 3: Timeline diagram of messages exchanged between DHCP client and servers when allocating a new network address (Droms, RFC 2131, 1997)

An abbreviated version of this process may be used by clients which have previously been issued configuration information and are now being reinitialized. In this case the client begins the process by broadcasting a DHCPREQUEST for the same configuration which it previously held. The server can reply with a DHCPACK if the configuration is still acceptable, or with a DHCPNAK; in the case of a DHCPNAK the client will start over with the normal DHCP process to obtain a new IP address (Droms, RFC 2131, 1997).

One of the configuration details set in the DHCP process is the length of time that the client is being given an IP address for. This is referred to as a “lease” duration. The client can request a specific lease duration, and generally the server will grant that request unless the request is for a longer lease than the maximum length the server is set to allow. The limitation on length of IP address leases is designed to simplify the process of recovering addresses that are no longer being used and adding them back to the pool of available addresses. A configured client will contact the server with a DHCPREQUEST after a defined period of time, usually $\frac{1}{2}$ the length of the lease, in order to “renew” the lease for a further period (Droms, RFC 2131, 1997).

Security Issues Involving Rogue DHCP Servers

The most well-known security issues with DHCP involve the potential harm to be done by the presence of unauthorized DHCP servers on the network. Because DHCP clients depend on the information sent by the DHCP server or servers to set their most basic IP configuration information, false DHCP configuration messages can interfere with or compromise DHCP-configured hosts at the most fundamental levels. Furthermore, the host usually has no way of knowing that it is being attacked or suborned, since it depends on the DHCP server to define what the network is supposed to look like and what its place in the network is supposed to be.

The simplest type of attack involving a rogue DHCP server is a denial of service (DoS) attack. The goal of this attack would be to prevent clients that depend on DHCP to obtain their configuration from joining the network. In order to do this, you would simply need to be the quickest server on the network to respond to every DHCPDISCOVER packet broadcast. You would simply respond with a DHCPOFFER that was valid but never respond with a DHCPACK to the resulting DHCPREQUEST from the client. Since the rogue server is not handing out real, valid configurations, it does not need to go through any tables to ensure that an address is not already taken; as a result, it should be easy to be the first to answer, every time. Some DHCP client implementations, however, will remember servers that initiate the lease process but don't complete it, and will ignore further DHCPOFFERS from these servers. (Hibbs, et al, 2003).

A more complex DoS attack would give out configuration information to clients, but make it invalid in such a way that clients that accepted it would not be able to communicate correctly. Offering addresses in an entirely different subnet often works for a more subtle effect one could provide a valid IP address and subnet mask but an incorrect gateway, thus restricting hosts' access to their own subnet and no farther; frequently this will prevent the user from logging on to their login server or domain controller (Jones).

As a matter of fact, the above scenario often happens on networks entirely by accident when a user or administrator innocently but erroneously configures their machine to act as a DHCP server. This problem has become so common that Microsoft built a safeguard into Windows 2000 Active Directory to prevent this type of thing from occurring by accident. When a Windows 2000 server has

the DHCP service turned on it first users DCHPDISCOVER messages to find out if there are any other DHCP servers on the network; if not, it continues to operate. If there are other DHCP servers operating, the server will attempt to find an Active Directory domain controller and determine whether the server is an authorized DHCP server in the Active Directory. This feature only works with the built-in DHCP service in Windows 2000 (and 2003), so it is really just a safeguard against an accidental rogue DHCP server, not a malicious one (Shinder, 2000).

A clever attacker may use a rogue DHCP server as a means of gaining further access to the network or to individual hosts, rather than to conduct something so crude and obvious as a denial of service attack. The most obvious way to do this would be to configure the attacker's IP address as the default gateway for client machines. By doing this, the attacker can force all of the target machine's traffic that is outbound from the local subnet to travel through the attacker's machine. If the attacker reads the packets and then passes them on to the correct gateway, the target is unlikely to notice the situation. However, the attacker is now able to read all of the target's outbound traffic, even on a switched network that might foil most attempts at traffic sniffing. The downside of this exploit is that the attacker is only able to read the outbound side of any conversations, since the inbound reply traffic will still be passed directly from the real gateway to the target machine (Jones).

Another exploit would be to set an incorrect DNS server on network clients. The attacker sets his own machine's IP address as the DNS server, and then answers DNS requests for client machines just like a real DNS server. However, for certain sites she can provide an incorrect IP address, directing the client to her own web server. With a clever false homepage for a service like Hotmail, this would make for a virtually undetectable type of password fishing (Jones).

Security Issues Involving Rogue DHCP Clients

Even without setting up a rogue DHCP server, there are a number of ways an attacker can exploit the DHCP service on a network. On the client side, the most obvious attack again is a denial of service. There are many rogue DHCP client programs available for download that will generate a massive number of DCHPDISCOVER requests, spoofing a different MAC address for each request. The rogue client answers the resulting DHCP OFFERS and quickly the network's DHCP servers run out of available addresses to assign to new clients. Some DHCP servers use ARP requests or PINGs periodically to query the addresses they've given out and see which ones can be de-allocated and added back to the pool of available addresses. It would be possible for a rogue client to listen on all of the IPs it has been assigned and answer such requests, but in practice DHCP servers generally cannot recover non-answering addresses nearly as fast as a rogue client can request new ones (Jones). Some networks may restrict DHCP clients to a list of specific MAC addresses. Because all DHCP clients broadcast their MAC addresses when they request service, many rogue client programs

“harvest” these MAC addresses for later use, providing them in the ‘chaddr’ field of their falsified DHCP requests (Hibbs, et al, 2003).

A more subtle move by a rogue client would be to attempt to usurp the place of an existing machine on the network, especially a machine such as a server that gets its IP via DHCP. The attack begins when the target machine is shut down, or the attacker may use a conventional DoS attack to render the target temporarily incapable of communicating on the network. The attacker then requests to “renew” the IP lease of the target machine, effectively usurping its IP address. By combining this move with a DHCP DoS attack, the attacker can prevent the target from re-taking its accustomed IP address (Jones).

A rogue DHCP client may also be used simply to gain service on a network where it does not belong. If the network has no security in place to restrict access to DHCP, this would seem to be more like taking advantage of their generosity than a true network exploit. But if a client listens for valid MAC addresses and then spoofs one in the ‘chaddr’ field in order to gain access, that level of sophistication is clearly an attack (Hibbs, et al, 2003).

Another method of exploiting the DHCP service would be to modify the DHCP packets traveling on the wire. This is made easier because DHCP is a UDP-based service and thus has no error checking provided by the transport layer. This type of attack could be used to prevent clients from getting or renewing a lease on an address, or could be used in combination to facilitate other types of attacks (Hibbs, et al, 2003).

Current Proposals for Improving DHCP Security

The Internet Engineering Task Force’s (IETF) DHCP working group released RFC 3118 in March of 2003. This RFC describes two extensions of DHCP to allow for authentication of DHCP messages. The first technique is token-based, with servers and clients exchanging passwords or “tokens” in plain text over the wire (Droms and Arbaugh, 2003). This technique is an improvement over the present state of affairs, but is so weak as to raise a concern that it would only lead to a false sense of security (Glazer, Hussey, and Shea, 2003).

The second, and preferred, method proposed in RFC 3118 is known as “Delayed Authentication.” This method utilizes a shared symmetric key and sends only a hash based on a varying part of the key; the key itself is never sent out over the wire. It also incorporates a one time nonce or the current time in the hash to limit the exchange’s vulnerability to “replay” attacks (Droms and Arbaugh, 2003).

In the DHCPDISCOVER and DHCPOFFER packets the client and server each send a SID or secret ID that references a specific part of the shared-secret key; this is the part that should be used in hashing the reply to that packet. Also included in these packets is a unique nonce or “replay” string (often a timestamp) that is to be included as part of the hashed value in the reply. Since each participant includes its own SID and its own replay string that the counterpart must use to create the hash that it returns, the exchange proves both that the

replying party possesses the key and that the reply is not simply being replayed from an earlier exchange that has been monitored (Glazer, Hussey, and Shea, 2003).

Neither of the authentication methods proposed in RFC 3118 has seen significant acceptance or implementation by any major software vendor. The primary obstacle appears to be that adding a key to the client configuration requires some pre-configuration of the client; avoiding this type of pre-configuration is exactly the reason why most networks that use DHCP use it. The Internet Software Consortium has created a working proof of concept build of a Delayed Authentication DHCP server and client, which is available as open source software on their website at <http://www.isc.org/products/DHCP>. The DHCP working group's 2003 Internet Draft suggests that the solution would be some sort of certificate-based authentication system similar to that commonly used to authenticate websites to users and, less commonly, to authenticate users to websites (Hibbs, et al).

Glazer, Hussey, and Shea, graduate students at UCLA, proposed such a system in their 2003 paper. Their proposed system of Certificate-Based DHCP Authentication (CBDA) would utilize certificates with a multi-level signing structure. This system is more scalable and requires less administrative effort than the Delayed Authentication method, but in order to truly provide trusted authentication of servers to clients it still requires some pre-configuration by administrators before deploying client machines. Alternatively, machines could be factory-configured to trust certificates signed by certain commercial vendors, much as web browsers commonly trust certificates signed by, for instance, VeriSign or Thawte (Glazer, Hussey, and Shea, 2003). This type of system would provide less security than a certificate or key provided by the using enterprise, but would present a significant hurdle to most attackers attempting to introduce a rogue DHCP server into the network.

Suggestions for Securing Networks that use DHCP

The following suggestions are tailored for a medium to large enterprise running Microsoft workstations and servers in a Windows 2000 or 2003 Active Directory network. However, many of the elements are relevant to a broader range of network configurations.

- Implement switches in place of hubs and ensure that your switches prevent MAC address spoofing. Many of the possible DHCP attacks depend on or are facilitated by MAC spoofing techniques (Jones).
- Ensure that the proper DHCP servers (and only the proper DHCP servers) are authorized in your Active Directory (Shinder, 2000).
- Utilize active detection techniques to identify potential rogue DNS servers. The Snort IDS, for instance, has a plug-in available for this purpose (Jones).

- Block DHCP (UDP ports 67 and 68) at the firewall separating your network from the Internet (NSA, 2002).
- Consider using a list of MAC addresses on the DHCP server of machines that are allowed to draw a lease (Hibbs, et al, 2003).
- It is common practice to use DHCP to set the DNS address of clients. The National Security Agency's Systems and Network Attack Center, however, recommends against this practice as having too great a potential for abuse (2002).
- Uninstall any DHCP/BOOTP relay agents on the firewall and on any routers where it is not required (NSA, 2002).
- Assign fixed IP addresses to DMZ servers, critical internal servers, and critical Internet clients. What is meant by assigning a fixed IP in this case is not creating a DHCP reservation for the host but manually hard-coding the IP address on the client itself. The DHCP client service should also be disabled on these machines (NSA, 2002).
- If any DHCP server is multi-homed, disable the service binding from any interfaces that will not be used for servicing DHCP client requests (NSA, 2002).
- Use a member server that is NOT a domain controller for your DHCP server (NSA, 2002).
- By default, Windows 2000 DHCP servers will contact the Windows 2000 DNS server and update the DNS records of any clients that cannot perform this function on their own (e.g. Windows 95 clients). The DNS server should be set to allow "only secure updates" so that only authenticated clients can have their records updated automatically, and only on their own behalf (NSA, 2002).
- Regularly check the DHCP audit logs, located in the %SystemRoot%\system32\DHCP directory (NSA, 2002).
- Consider using PPPoE or some other type of access protocol that requires authentication to a RADIUS server before an IP address will be assigned. This type of measure is especially important in an environment with little physical security and/or with a high degree of mobility and change in legitimate clients (Graham, 2002).

Conclusions

DHCP is a commonly used protocol that provides a high level of service and convenience for both users and administrators. It has few security safeguards built into it and is quite susceptible to abuse. Most of the security difficulties with DHCP stem from the desire to have the ability to plug an out-of-the-box system into the network and be able to use it immediately. If one desires this type of ability, there will always be serious limitations as to how much security can be implemented. For networks with conventional security

requirements, there are steps that can be taken to provide a reasonable level of confidence that DHCP will not be abused or suborned with the network.

For those with stricter security requirements, it seems that they would do well to consider avoiding DHCP when practical and configuring hosts by hand until such a time as some of the schemes for security and authentication or implemented by system vendors. Alternatives such as the use of an authentication protocol such as RADIUS prior to issuing IP addresses should also be considered.

© SANS Institute 2004, Author retains full rights.

Bibliography

- Alexander, S. and Ralph Droms. "DHCP Options and BOOTP Vendor Extensions." Internet Engineering Task Force RFC 2132, March 1997. URL: <http://www.ietf.org/rfc/rfc2132.txt> (28 January 2004).
- Desmond, Paul. "CERT Warns Of DHCP Threat." ESecurityPlanet.com, May 15, 2002. URL: http://www.esecurityplanet.com/trends/article.php/10751_1122631 (28 January 2004).
- Droms, Ralph et al. "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)." Internet Engineering Task Force RFC 3115, July 2003. URL: <http://www.ietf.org/rfc/rfc3115.txt> (28 January 2004).
- Droms, Ralph. and William Arbaugh. "Authentication for DHCP Messages." Internet Engineering Task Force RFC 3118, March 2003. URL: <http://www.ietf.org/rfc/rfc3118.txt> (29 January 2004).
- Droms, Ralph. "Dynamic Host Configuration Protocol." Internet Engineering Task Force RFC 2131, March 1997. URL: <http://www.ietf.org/rfc/rfc2131.txt> (28 January 2004).
- Glazer, Glenn, Cora Hussey, and Roy Shea. "Certificate-Based Authentication for DHCP." March 20, 2003. URL: http://www.cs.ucla.edu/~chussey/proj/dhcp_cert/cdba.pdf (03 February 2004).
- Graham, Joseph. "Authenticating Public Access Networking." *Proceedings of the 30th Annual SIGUCCS Conference on User Services* (2002): 247-248.
- Hibbs, Richard, et al. "Dynamic Host Configuration Protocol for Ipv4 (DHCPv4) Threat Analysis." (work in progress) Internet Engineering Task Force, June 2003. URL: <http://www.ietf.org/proceedings/03nov/I-D/draft-ietf-dhc-v4-threat-analysis-00.txt> (28 January 2004).
- Jones, Steven. "Flaws within the Dynamic Host Configuration Protocol." NetworkPenetration.com, URL: http://networkpenetration.com/dhcp_flaws.html (29 January 2004).
- Perkins, Charles and Kevin Luo. "Using DHCP With Computers That Move." *Wireless Networks* 1(1995): 341-353.
- Shinder, Thomas. "Back To Basics: Windows 2000 Rogue DHCP Server Detection?" October 8, 2000, ServerWatch. URL: <http://www.serverwatch.com/tutorials/article.php/2193001> (03 February 2004).
- Takamichi Saito, Komori Tadashi, Mizoguchi Fumio. "A Secure DHCP System with User Authentication." *IPSI JOURNAL* 43(8): 2002.

United States National Security Agency (NSA). "Guide to Securing Microsoft Windows 2000 DHCP." July 2002. URL:
<http://nsa2.www.conxion.com/win2k/guides/w2k-18.pdf> (22 January 2004).

© SANS Institute 2004, Author retains full rights.