

Security Requirements for the DHCP protocol
<[draft-ietf-dhc-security-requirements-00.txt](#)>

Status of this Memo

This document is an Internet-Draft. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

To view the entire list of current Internet-Drafts, please check the "1id-abstracts.txt" listing contained in the Internet-Drafts Shadow Directories on ftp.is.co.za (Africa), ftp.nordu.net (Northern Europe), ftp.nis.garr.it (Southern Europe), munnari.oz.au (Pacific Rim), ftp.ietf.org (US East Coast), or ftp.isi.edu (US West Coast).

Distribution of this memo is unlimited.

Abstract

This document addresses the general security requirements of both DHCPv4 and DHCPv6. This document lists security requirements and the reasons for each requirement. This document does not address how to implement the security requirements.

1. Background

This section presents some concepts and definitions used throughout this document.

1.1. Authentication, confidentiality, data integrity

[RFC-1825](#)[\[RFC1825\]](#) contains a great description of these terms and their uses. Authentication is the process of establishing the identity of some entity. Once identity has been authenticated, that identity can be used for access control, accounting etc. There are number of authentication technologies available.

Public key cryptography is a powerful tool that relies on complex mathematical operations to provide information that only the holder of the private key could have generated.

Shared secret authentication is the process of digesting the data transmitted and obfuscating the digest by applying a transformation by a key that is only used by the two entities. This technology can be used to provide both authentication and data integrity. Each pair can share multiple shared secrets, it is important that each secret have an identifier attached to it.

Confidentiality can be accomplished by encrypting the data contents of the outgoing packet. Shared secrets can be used as keys for symmetric encryption.

1.2. Shared secrets

Shared secrets are between two entities; there is NEVER a need to share these secrets with other entities. The hosts storing the secrets MUST protect the secrets as well as possible.

1.3 Terminology and DHCP v6 considerations

This document uses DHCPv4 terminology as it is more familiar than the new DHCPv6~~[DHCPv6]~~ terminology. When this document talks about DHCP v4 DISCOVER messages the same will apply to DHCP v6 Solicit Message. Subsequent v6 messages are similar to v4 messages. Both v4 and v6 take advantage of RELAY agents, in some cases these agents can add to the messages from servers, it is important that the added information is treated the same way as data from servers.

Gudmundsson, Droms

Expires October 1998

[Page 2]

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#).

2. Proposed DHCP security requirements

The proposed requirements can be summarized in the following rules.

- Initial Client/Server Authentication

1. Server MUST be able to authenticate client identity.
2. Client MUST be able to authenticate the server identity as an authorized server.

- Initial Relay Agent/Server Authentication

3. Server MUST be able to authenticate relay agent identity as an authorized relay agent.
4. Relay Agent MUST be able to authenticate server identity as an authorized server.

- Successive Client to Server and Relay Agent to Server Communication

5. Client and Server MUST agree on security model for protecting future communication.

- Server/Relay agent advertisements

6. Advertisements MUST be verifiable by all recipients.

- Server/Server communication

7. All communication MUST be protected for data integrity. Servers MAY request that communication be encrypted.

DHCP security cannot be accomplished in a vacuum; as DHCP is not a general purpose communication protocol. Fortunately there are available (or soon will be) protocols that DHCP can take advantage of. First and foremost DNSSEC[RFC2065] or some other key distribution mechanism must be available. IPSEC[RFC1825,IPSEC] will be able to handle requirements 5 and 7.

2.1. DHCP Identity

In order to secure DHCP all clients MUST have an identity, this identity can possibly be one of the following: host name, user identity, account code. The "prime" identity MUST have a public key stored in the key distribution mechanism. The client MUST know its identity before contacting the server. Each client MUST have access to the correct private key before contacting the DHCP server.

If the identity selected for a host is its host name and the key distribution mechanism is DNS, then the public key used to authenticate the host is stored under the host name in its home zone. The private key needs to be stored in the computer at all times. If the identity selected is the user then the key is stored under the user name in DNS (e.g.: ogud.tis.com for me), and the user needs to load the computer with the private key before the host can contact the server. If the identity of the host is just there to uniquely identify the host, the host still needs a private key.

Traditional identifiers such as MAC addresses, are not suitable in all cases in identify clients, first there is no database of what MAC goes with what host, secondly some MACs are portable and can easily migrate between hosts such as Ethernet PC CARDS.

2.2. DHCP communication protection

DHCP is a protocol that carries publicly known information, thus there is limited need for confidentiality. DHCP requires data integrity protection for communication. The option to allow DHCP servers/clients to request confidentiality SHOULD be part of any security architecture.

2.3. Policy issues.

This document does not address access control issues as that is a policy issue for each site. Effective access control depends on correct authentication, thus this work will make access control simpler. This document does not address the issue of protecting the private key on either server, agent or client.

2.4 Security threats to DHCP

2.4.1 Attacks against servers

There are many possible attacks possible against servers, including denial of service by exhausting the servers allocated address space. Another denial of service attack is to overload the server causing it not to respond to clients.

Once servers start updating DNS and other directory services, DHCP servers can be spoofed to register incorrect information in those services.

Another possible attack is to gain unauthorized access to some resources, such as network access.

2.4.2 Attacks against clients

Fake servers[DHCPVERSERY] can provide clients with partially correct information that allows the attacker to route traffic through certain host where critical information can be collected. This becomes important to detect and prevent when encrypted traffic is allowed to pass through firewalls.

Clients can be configured with bogus data, so that they will assume that the network is down. In some cases it is hard to get a client to reconfigure itself. Clients can also be configured with addresses of other clients, causing address conflicts.

The bright side of this problem is that fake servers are easy to detect by monitoring the network for DHCP traffic.

2.5. Complications in implementing the security models

A Client that issues DISCOVER message does not have any IP address that works outside the local network, and may not even work on the local network. This prevents the clients from checking with outside information sources. Servers on the other hand are fully configured and can use any information sources accessible.

Clients will not wait long for OFFER message, some security checks may take longer than the DHCP retransmission timeout. If DHCP servers had an option to inform clients that DISCOVER messages are

being worked on and client should expect an answer in short order,

Gudmundsson, Droms

Expires October 1998

[Page 5]

then this problem would be solved.

Some DHCP servers do not have CPU cycles to spare to do security checks. Computational load on server in verifying the identity of client can be significant. Different authentication mechanisms have different computational overhead, similarly network delays have to be taken into account if DHC server needs to query remote data source for more data.

3. DHCP Security components

3.1 Authentication services

In order for DHCP servers to be able to determine if a client request should be serviced it is essential for the server to be able to establish the client's identity. There are two kinds of identities that are possible, local mutually agreed upon identities and global identities.

Local identity is sufficient if the client will only be configured from a small set of servers, and if there are no expectations that the client will migrate to another location. This is an acceptable solution for a site where all computers are stationary but are configured from DHCP for administrative reasons. Solutions of this kind have certain scaling problems.

Global identity on the other hand is needed when a client can connect to multiple servers and it provides some of its identity.

An example of local identity is a name or number that is configured in the client and server. This could be the name of the client on the local network. An example of global identity is a DNS name.

3.2 Replay prevention

In order to protect replay attacks, all communication to servers should contain some variable data that never repeats and both server and client can agree on. A simple approach is to use time of day information that clients and servers check and act upon. Clock synchronization service can be provided by an outside entity[RFC1305] once a client is configured, but bounds must be placed on acceptable skew while a client is off line or migrates between locations. Clients SHOULD not trust time information from

servers until after servers have been validated as such. Clients

Gudmundsson, Droms

Expires October 1998

[Page 6]

should always assume that the network is insecure.

Another approach is to use counters but this requires clients to keep state for each server they talk and has synchronization issues.

3.3 Data confidentiality

This is not desired service at this point, but it can be added at a later point for all communication except DISCOVER and possibly OFFER and REQUEST. All subsequent communication can be encrypted.

3.4 Possible security models for DHCP

This section will present a few possible security models and the reasons why each one may be useful. This section IS NOT an advocacy for any of the described models there may be other models possible. It is expected that any security proposal put forward state which model is used as its bases.

3.4.1 The ``No security'' model

This is the current situation. Below are few arguments that can be made for the status quo.

Security is hard. Considering that DHCP for IPv4 is a hack built on an older hack (BOOTP), there is not enough flexibility in the protocol to add security.

A smart client attached to a broadcast network can learn everything it needs to know to configure itself by listening to network traffic. The client can either monitor DHCP traffic and/or all network traffic to find gateways, servers and unused addresses. There is no protection against this.

DHCPv6 can on the other hand be extended and modified to fit any security model selected. Sites will migrate to IPv6 soon, and the ones that do not deserve what they get.

In this model DHCP clients will be able to do harm and be harmed by bogus servers. This model is not acceptable when DHCP servers perform DNS update operations on a client's behalf. Sites MAY select this model but this is strongly discouraged.

3.4.2 The ``Simple'' model

A DHCP client is configured with a token that allows it to authenticate itself to the servers in the DHCP DISCOVER message. If servers can authenticate the token and the client associated with the token is allowed to communicate with the server the server will reply with OFFER message.

In this model servers will know with which client they are dealing, and that should be sufficient protection against most of the attacks against the servers. If a client is able to authenticate the server response, the client might be protected against bad servers.

With minor extensions to DHCP, all subsequent communication can be protected.

3.4.3 The ``Comprehensive'' Model

In this model DHCP servers and clients have the ability to authenticate each other. The requirement here is that clients must be able to authenticate the server without any communication as they can not trust the information from the server. This model also must prevent replay attacks.

This model protects all traffic between clients and servers, making it impossible to stage any attacks other than denial of service attacks due to CPU overload of servers.

4. Client Authentication

Initial authentication is the most important step. Once server and client have established each other's identity the remaining problems can be solved.

The problem of initial client authentication cannot be solved by IPSEC, as the client does not have an IP identity when it requests service for the first time from the server. Once a client has been configured it can enter IPSEC security associations with other DHCP servers during the lifetime of the IP address lease.

4.1 Identification of DHCP clients and scaling issues

From a DHCP servers perspective it needs a 'handle' that can be used to uniquely identify each client, to the server it should not matter what kind of handle is used. From a security point of view, it is important that the 'handle' be always the same and no possibility of confusion. In DHCP there are at least two types of clients: clients that request some of their net identity from DHCP, and clients that request all of their net identity from DHCP, but from the security requirements standpoint these are identical.

MAC addresses are frequently proposed as 'handles', but in many cases they are not suitable. For example most laptop computers have network connectivity via a PCCARD, these cards are easy to swap and thus are not static. Similarly laptops at different times connect via Ethernet, modem, infrared or wireless all with different MAC addresses but the laptop may ask for the same Identity regardless of connection.

Previous DHCP security proposals[DHCAUTH] have suggested the use of shared secrets and passwords to identify clients. It is also possible to use some form of challenge/response system to identify clients. These approaches have limited scaling ability and require a server to server protocol. But in many environments these weaker authentication mechanisms are adequate.

The most general case is the identification of a computer that connects to a world wide ISP network and expects the same identity regardless of location. In this case it is unlikely that the same DHCP server serves both India and Iceland. A network of this kind can have a collaborative agreement between a number of different ISPs, with multiple administrative domains. It is not reasonable/scalable that all DHCP servers in this network know shared secrets, or passwords for all computers that are allowed to connect. From a security standpoint it is a bad practice to distribute shared secrets or passwords to many places.

4.2. Motivation for single strong authentication schema.

DHCP SHOULD require all servers and clients to support at least one mandatory authentication protocol, and allow other ones. This will ensure interoperability of all servers and clients.

4.3 Motivation for global DNS identities for DHCP clients

Once the global identity is registered with an information service, this identity is available within the limits of the information service. DNS is the most common information service used by computers.

DNSSEC[RFC2065] strengthens DNS[RFC1035] against information corruption and provides distribution of public keys. If every host that is configured by DHCP has a public key stored in DNS then servers can verify digital signatures generated by that key. Once clients are configured it is possible for client to verify that the server it was configured by is a good DHCP server. In order to do this, DHCP servers for each domain must be listed.

IPSEC can be preconfigured with SPI's but there is no definition for the format of the 'destination address'. If it is DNS format, DHCP entities MAY enter IPSEC relationship without a key exchange once client has received DHCP ACK message.

5. Server verification by clients

When a client receives an DHCP OFFER message it should try to authenticate the server. For stationary clients this can be as simple as verifying that this is one of the servers it knows about, and trusts. For mobile clients and in adverse networks this is more difficult, there must be a mechanism for identifying the servers that are authorized to allocate addresses in a range. This could be accomplished by adding an RP record at delegation points in the inverse DNS tree or at every node that points the authorities for that address(es), the <mbox-dname> is the mail address of the responsible party and the <txt-dname> is the authorized server. There can be as many RP records as there are servers. If the inverse address map is protected by DNSSEC then this is a convenient mechanism to authenticate this is a good server. For clients that have host name configured they should perform similar lookup to make sure the server is authorized to allocate names in that space.

6. Security considerations

This document addresses how to add security features to the unsecured DHCP protocol.

References

- [DHCP] R. Droms, "Dynamic Host Configuration Protocol", [RFC 2131](#), Bucknell University, April 1997.
- [DHCAUTH] R. Droms, "Authentication for DHCP Messages", Internet Draft <[draft-ietf-dhc-authentication-04.txt](#)> August 1997
- [DHCPv6] J. Bound, C. Perkins, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", Internet Draft <[draft-ietf-dhc-dhcpv6-10.txt](#)> May 1997
- [DHCPVERSERV] R. Watson, O. Gudmundsson, "DHCP Server verification by client via DNSSEC", <[draft-watson-dhc-serv-ver-00.txt](#)> July 1997.
- [IPSEC] R. Atkinson, "Security Architecture for the Internet Protocol", Internet Draft <[draft-ietf-ipsec-arch-sec-03.txt](#)>, February 1998.
- [RFC1035] P. Mockapetris, "Domain Names - Implementation and Specification," [RFC 1034](#), ISI, November 1987.
- [RFC1305] Mills, D., "Network Time Protocol (v3)", [RFC 1305](#), March 1992.
- [RFC1825] R. Atkinson, "Security Architecture for the Internet Protocol", [RFC 1825](#), September 1995.
- [RFC2065] D. Eastlake, C. Kaufman, "Domain Name System Security Extensions", [RFC 2065](#), January 1997.

9. Author address

Olafur Gudmundsson
Trusted Information System
3060 Washington Road

Glenwood, MD 21738
+1 301 854 6889
ogud@tis.com

Ralph Droms
Computer Science Department
323 Dana Engineering
Bucknell University
Lewisburg, PA 17837
+1 717 524 1145
droms@bucknell.edu

