Zeroconf

A new way of life?

Ing. R.M. Visser
Ing. A.A. Ahrouch

February 2004

UNIVERSITEIT VAN AMSTERDAM

Zeroconf
A new way of life?

A Comprehensive Guideline for Zeroconf

The report is submitted in accordance to the requirements for the degree of:
Master of Science, System & Network Engineering

The graduate internship is approved by:
Drs. J. Scheerder

Educational Institute:
Universiteit van Amsterdam,
Department of Information Technology
Weesperzijde 190
1079 DZ Amsterdam
The Netherlands

Phone: +31 (0) 20 595 1111
Web: http://www.os3.nl

Authors:
Ing. R.M. Visser
Ing. A.A. Ahrouch

Date:
February 2004

Available versions of this report:
The report is available in PDF format on the following websites:

http://www.os3.nl/~ahrouch
http://www.os3.nl/~fear

Abstract

The Internet Engineering Task Force (IETF) developed Zeroconf around 1999 when it became obvious that basic network configuration can be done without manual configuring at all. Automatic assignment of DNS and IP address with the power of service discovery would be the key to the success of Zeroconf.

Zeroconf should reduce the costs of configuration by system administrators when new services are implemented in the existing network environment. When adding a simple network resource like a printer requires every workstation to be configured, it is obvious much can be gained with service discovery.

Home networks without administration requirements should make the life of everyone allot easier. A digital camera should be able to find available printers as soon as the camera is plugged into the network allowing everyone to print pictures from the camera without asking someone to configure the printer.

Not all these beautiful things come without a price and certain risks should be considered. Discovery protocols can claim a high toll on network resources when allot of services and clients are on the same network, each broadcasting messages all day long to discover new services. Security is also something that has to be looked at very carefully; when everything works automatically, it may work for people you do not want it to work for.

Foreword

Intended audience

This document is the result of four weeks of research required by the University of Amsterdam for the study System and Network Administration. Knowledge of network technologies like DNS, DHCP, Multicast and limited knowledge of TCP/IP is required to fully understand this document.

Plan of the text

In chapter one we describe the research scope as we have defined it with our mentor drs. J. Scheerder.

Chapter two gives an introduction in Zeroconf and describes what the functionalities are that Zeroconf offers.

Chapter three is devoted to various Service Discovery Protocols, which are functionalities of Zeroconf. This chapter gives a technical overview of service location protocols with a description of their protocol overview, strenghts and weaknesses.

Chapter four describes an analysis of the benefits, requirements and security aspects of Zeroconf.

Chapter five takes the reader through the various applicabilities that Zeroconf provides to system administrators. This chapter handles also the ethical discussion behind Zeroconf.

Conclusive we will describe our findings and results of our research.

Contents

1 Research scope

1.1 Research fields

The research was performed on the following aspects of Zeroconf.

1. Understanding of Zeroconf.
2. Applicabilities of Zeroconf.
3. Security issues of Zeroconf.

To gather the relevant information for this paper focus was put on the following research questions.

1. In which form offers Zeroconf advantages for system management?
2. What are the security aspects that have to be considered when using Zeroconf in a Local Area Network environment?
3. What are the problems that can appear when using Zeroconf on a cross-platform network?

1.2 Terminology

The terminology used in this paper is similar to the terminology used in the internet draft *draft-ietf-zeroconf-ipv4-linklocal-10.txt* [1] is available at the IETF website.

This document describes link-local addressing, limited to IPv4 address communication between two hosts on a single link. A set of hosts is considered to be "on the same link", if:

- a host A from that set sends a packet to any other host B in that set, using unicast, multicast, or broadcast, the entire link-layer packet payload arrives unmodified, and
- a broadcast sent over that link by any host from that set of hosts can be received by every other host in that set.

[1] See reference: http://files.zeroconf.org/draft-ietf-zeroconf-ipv4-linklocal.txt

1.3 IETF Zeroconf Working Group

Back in 1997, the idea began to form a networking mailing list on *net-thinkers* directed to Macintosh users. That idea evolved for a couple of years before ending in semi-formal meetings and finally the establishment of a full Zeroconf Internet Engineering Task Force Working Group (IETF). The Zeroconf Working Group of the IETF is chartered September 1999 and held its first official meeting at the 46th IETF in Washington, D.C., in November.

The Zeroconf Working Group is engaged with methodologies for interface configuration in assigning IP addresses to hosts, translation between domain names and addresses and the discovery of offered services on the network on basis of name or characteristics.

The Zeroconf Working Group has the goal to enable networking without any manually configuration and administration of network protocols and services. With the current internet protocols the idea of self-configuring and self-administrating hosts cannot be implemented[2], what actually would be ideal; the applicability's would be indefinably.

[2] As noted in RFC 1122, STD 3

2 Introduction to Zeroconf

Hosts in a network can be assigned static or dynamic network configurations. Static hosts are usually been used when hosts are permanently connected to administered networks, while dynamic hosts are often used in large networks, such as corporate local-area networks or dial-in accounts, when computers are frequently reconfigured, or when a limited number of IP addresses is available to share between many computers.

When configuring hosts in a network, there is a system administrator necessary to manage a simple form of centralized network system such as a DHCP (Dynamic Host Configuration Protocol) server to assign the necessary IP configurations to the host. Adding a host to a network requires additional configuration steps to be taken that could be time consuming to implement before the host can make use of the offered services. In local-link networks the use of centralized network systems can be managed, but are not always easy to maintain.

Ad-hoc networking turns out to be interesting when a host can connect to a network in a fast way to make use of the offered network services. A step further is to take up hosts in the network infrastructure without letting the user perform any necessary network configuration.

A solution for this can be found in the so-called Zeroconf method. Zeroconf is a method, which can be used to enable configuration- and administration-less networks.

2.1 Definition of Zeroconf

For a consistent use of the term “Zeroconf”, we handle the following definition:

> *“Zeroconf is short for zero configuration IP networking, a method of network devices via an Ethernet cable without requiring configuration and administration. Zeroconf is able to allocate addresses without a DHCP server, translate between domain names and IP addresses without a DNS server, and find services, such as a printer, without a directory service.”*[3]

The technology is intended for use in small networking situations where a low priority on security and where it is inappropriate or impossible to establish a working IP network using traditional technologies, such as DHCP and DNS.

[3] Source definition: http://www.webopedia.com/TERM/Z/Zeroconf.html

2.2 Zeroconf functionality

In a local-link network, Zeroconf functionality can be implemented in four main areas:

1. Allocate addresses without a DHCP server.
2. Translate between names and IP addresses without a DNS server.
3. Find services, like printers and applications, without a directory server.
4. Allocate IP Multicast addresses without a MADCAP server.

When implementing Zeroconf in these areas, keep in mind that network security should be no worse than when these goals were not implemented.

2.2.1 Allocate addresses without a DHCP server

In a centralized network hosts receive an IP configuration from a DHCP (Dynamic Host Configuration Protocol) server. DHCP is a communication protocol that manages and automates centrally the assignment of IP configuration to hosts. Without a DHCP server, each host has to be assigned an IP address manually, but the approach of network configuration is still based on a centralized network structure.

Address assignment with Zeroconf happens randomly without any involvement of a system administrator. Advantage of this approach is that there is no centralized server, because every host knows about every host in the local-link network. Disadvantage is that the assigned IP address can only be deployed in the local-link network and cannot be routed out of the network.

Zeroconf is actually intended for use with small ad-hoc networks and local-link networks. Theoretically the amount of 65024 (2^{16}-512, the first and last 256 addresses are reserved for future use.) link-local Ipv4 addresses are available for randomly assignment. Attempting to use all those addresses in a single local-link network would result in a high probability of an address conflict because a host would have to take an excessive amount of time to find an available address.

The algorithm to assign randomly an IP address to a client is as follows:

1. Generate a random number by using the MAC address.
2. Select with this random number an IP address from the IP range of 169.254/16 (excluding the first and last 256 addresses)
3. Do an ARP request for the selected IP address. If the ARP request response, return to step two.
4. Do gratuitous ARPs to claim the address.
5. Observe for other hosts that want to claim the same IP address, and preserve the own IP address if needed.

Linux users with IPv4 networks can add the support of automatic address assignment to their systems by installing the *zcip* package from the SourceForge Zeroconf project (this package is described in chapter 5.1.1). Users of any operating system with IPv6 support do not need any additional packages because IPv6 already has IP auto-configuration built in.

2.2.2 Translate between names and IP addresses without a DNS server

Normally a network has a DNS server to translate names to IP addresses and back again. However, when you are in a network without a DNS server you will need to find anther way to translate between IP addresses and names. This works the same way as IP assignment without DHCP server. You claim a name and broadcast your claim. When nobody answers the name is free and you have your name. When someone replies saying he has the name, you will need to try again with a new name.

2.2.3 Find services, like printers, without a directory server

To use a service a user first needs to know where the service is located and which protocol it requires. There are three ways to discover a service. The first and most used way is to ask a host that is known where for example a printer service is located. This request can be more specific depending on the protocols available. Maybe you only want to know printers that are reachable via LPR? This requires configuration, which is exactly what we want to avoid. There are two ways to locate a service without configuration. You can broadcast a request for a service through the network and the required service can respond to it. Another way is that the service tries to advertise itself and the service broadcasts his location. It is like going to a party with a blind date and you only know the name of your date. You could ask the host of the party, but you first need to find him/her and ask where your date is. The Zeroconf way is to scream across the room for your date (or your date does the screaming part) and you locate each other as well.

2.2.4 Allocate IP Multicast addresses without a MADCAP server.

The IETF has existing or developmental work in the first three functionality of Zeroconf. For the multicast address assignment functionality, the general acceptance is though very low. Work on multicast address assignment in a Zeroconf environment has therefore effectively been abandoned due to a lack of interest.

Multicast Address Dynamic Client Allocation Protocol (MADCAP) is an extension to the DHCP standard used to support dynamic assignment and configuration of IP multicast addresses on TCP/IP-based networks.

Typical applications for multicast are conferencing and audio, which usually require users to configure multicast addresses specifically. Unlike IP broadcasts, which are received by all computers or other hosts on the network, a multicast address is a group of computers, using the concept of a group membership to identify the computers to which the message is to be sent.

3 Service discovery protocols

3.1 Introduction to service discovery protocols

When a service is introduced into a network of computers, every computer that wants to use the service needs to know the location of the service and some kind of configuration to use the service. This requires allot of work for system administrators in large computer networks and can be very expensive. In the year 1999, several parties started on the development of protocols that would reduce the configuration costs.

The main task of a service discovery protocol is to locate services for users and to advertise services for service providers. Currently three protocols could be used for Zeroconf.

- The Internet Engineering Task Force (IETF) developed the Service Location Protocol version 1 and 2(SLPv1 and SLPv2) which is copyright by The Internet Society.
- Apple Computer Inc. developed Domain Name System based Service Discovery (DNS-SD) to be used in Apple's Implementation of Zeroconf named Rendezvous.
- The Internet Engineering Task Force (IETF) developed Simple Service Discovery Protocol (SSDP). Microsoft used it in Universal Plug and Play, Microsoft's answer to Zeroconf.

3.2 Service Location Protocol

SLP provides a dynamic configuration mechanism for applications in local area networks. Applications act as clients that need to find services attached to the available networks. When there are many different clients and/or services available, SLP is adapted to make use of Directory Agents that act as a centralised repository for advertised services.

SLP is intended to function in networks that work under a cooperative administrative control. Such networks permit policies to be implemented regarding routing, security, etc. These policies are not possible on the scale of the internet.

3.2.1 Protocol overview

In the Service Location Protocol, a client is viewed as a User Agent (UA) and services are advertised by Service Agents (SA). A Directory Agent (DA) can act as a broker between a UA and a SA.
A UA issues a request for a service, specifying the characteristics of the service, which the client requires. SLP allows a UA to directly send its requests to a SA. In this case, the request is send via multicast. The SA then advertises its service via unicast containing the location of the service.

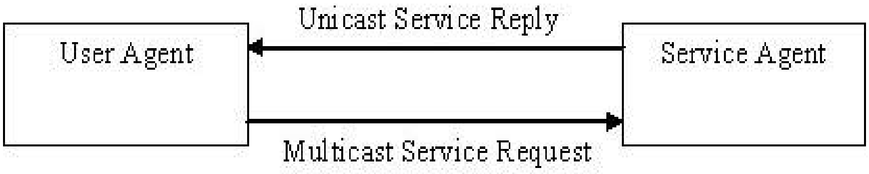

In larger networks, DAs can be used to function as cache. Server Agents send Service Registry messages to a DA containing all the services they advertise and receive acknowledgements in reply. These advertisements must be refreshed with the DA before they expire. User Agents send via unicast requests to the DA instead of the SA if any DA is known.

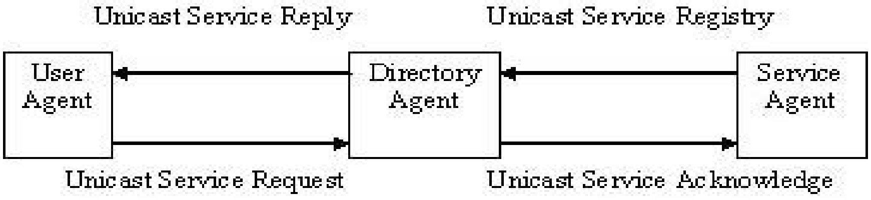

User and Server Agents discover Directory Agents in two ways. First, they send a multicast Service Request for the DA service when they start up. Second, the DA sends an unsolicited advertisement via multicast on random times, which the SA and UA listen for. In both cases, the Agents receive a DA advertisement.

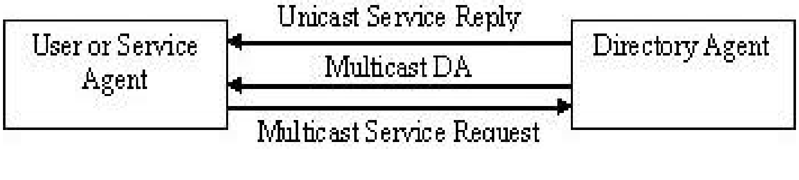

Services are grouped together in scopes. These scopes are strings that identify services that are administratively identified. A scope could indicate a location (first floor, Amsterdam, etc), network name (Toscana, Sicilia) or some other category. When a UA is assigned a scope, it can only discover services in that scope. When a UA does not get a scope assigned, it can discover every service available on the network. This allows an administrator to control the access to services.

Service Agents and User Agents may verify digital signatures provided with DA adverts. User Agents and Directory Agents may verify service information registered by Service Agents. SLP Security Parameter Index (SLP SPI) is used with keying material to verify digital signatures. Every host configured to generate digital signatures includes the SLP SPI used to verify it in the Authentication Block of the protocol when transmitting.

3.2.2 Strengths

SLP is a very complete protocol with allot of options. Extensive search options are available and it can be managed to limit the scope of a search. Security is built inside the protocol with SLP SPI. It supports allot of standard protocols to allow communication between clients and servers.

3.2.3 Weaknesses

SLP is a very large protocol that basically requires Directory Agents to function in large networks. Without a DA, SLP will overload the network with all the requests and replies. Even with DAs Scalability is a big problem that SLP does not handle very well. When a network goes down, due to power failure, SLP stresses the network allot re-establishing all the service locations for every User Agent. A so-called Net storm Basically SLP shows four general types of problems.

- Extensibility – SLP prevents the creation of extensions without approval from IANA/IETF, thus limits the possibilities in which this protocol can be used.
- Re-inventing Wheels – SLP designs functionality that already exists without providing and additional value to it. SLP has its own protocol for every message thus replacing HTTP, Directory Agents replace LDAP and the scopes replace URIs.
- Security – SLP Mandates use of X.509, which requires a complicated administrative structure and forces a centralized implementation for which Zeroconf is not designed.
- Low-end Devices – SLP is very large and low-end products such as Cameras do not require such a big protocol and require more on the network to provide certain services for them. For example, take a digital camera connected to a home network with maybe two printers. The small LCD on the camera can easily display both printers. However when the same camera is connected to a large corporate network with maybe 200 printers it gets allot harder to find the printer you want.

3.3 DNS based Service Discovery

Domain Name System based Service Discovery (DNS-SD) finds its origin loosely in AppleTalk Name Binding Protocol (NBP). The requirements of NBP where used to develop DNS-SD. It is currently used in Apple's implementation of Zeroconf named Rendezvous.

3.3.1 Protocol overview

As the name suggests DNS-SD uses DNS technology to allow for service discovery. SRV Resource Records (RFC 2782) are used to locate the desired service. SVR Records with for example the (hypothetical) name "_http._tcp.os3.nl." would allow a client to discover a list of all servers implementing the "_http._tcp" service (i.e. Web servers) for the "os3.nl." domain. The assumption is that all servers offer an identical set of Web pages and it does not matter which server is used as long as the total load is balanced. When however a printer service is needed, getting a random printer is not what a user usually wants. To get a specific service, PTR Records are used which list zero or more Service Instance Names (SIN) of the form:

SIN = <Instance> . <Service> . <Domain>

<Instance> is the name that is given to the service and mostly the only part displayed to allow easy reference.
<Service> consists of two parts, the first part is the protocol name used to communicate with the service and the second part is "_udp" or "_tcp" to indicate which network protocol is used.
<Domain> is the domain name in which the service is located.

An Example could be: "Printer._ipp._tcp.os3.nl.". This shows a Printer with the Internet Printer Protocol that is reachable via TCP on the os3.nl. domain.

When a client requires a service, he asks for the right SRV record and gets all the information needed to connect to the service. This information consists of an IP address, the protocol and a port number.

When a service requires more information then a SRV record can handle, the user can ask for a TXT record. A TXT record can hold the additional data like a queue name, which is often used with LPR printer services. RFC 1035 section 3.3.14 specifies the layout of TXT records and is not specific to DNS-SD.

3.3.2 Strengths

All services are nicely sorted with this system and all the information is spread out throughout the network. You only get services for which you have the protocol to use them. It also has no application specific information in it thus it in theory it could be used with every program and on every Operating System (OS) as it is completely open source. Another big strength is that it does not use multi- and unicast messages, which require additional routing to be able to cross subnets.

3.3.3 Weaknesses

DNS-SD does not have any search method that allows a user to search without knowing instance names. A program like iTunes that uses DNS-SD to search for play lists shared by other iTunes users knows the instance name of other iTunes clients. Thus, iTunes can search for other clients while a normal user cannot. The biggest weakness of DNS-SD is that a working DNS is required. This can be a regular DNS service with a central server in big corporate networks or multicast DNS to allow DNS services without a central server. To be completely free of human intervention multicast DNS is needed and provided in Rendezvous. The DNS also needs to be automatically updated when new services become available.

3.4 Simple Service Discovery Protocol

SSDP is a Protocol developed to allow clients to locate which http resources provide the desired service. It should work without configuration, management or administration. It provides the bare minimum in search capabilities and should be switched on by default to allow total carefree networking.

3.4.1 Protocol overview

Discovery occurs when a SSDP client multicasts a HTTP discovery request to the SSDP multicast channel/Port. SSDP services listen to the SSDP multicast channel/Port in order to hear such discovery requests. If a SSDP service hears a HTTP discovery request that matches the service it offers then it will respond using a unicast HTTP response.

SSDP services may send HTTP notification announcements to the SSDP multicast channel/port to announce their presence this is used when a service comes online or when the caching time expires. The information is sent in the ST header of HTTP and the MAN header is used to mark the message as a discover message with ssdp:discover.

Services are identified by a unique pairing of a service type URI and a Unique Service Name (USN). Service types identify a type of service, such as a printer service. USNs are used to differentiate between two services with the same service type for example two printers; the first is a HP4100 and the second a HP4200.

In addition to providing both a service type and a USN, discovery results and presence announcements also provide expiration and location information. Location information identifies how one should contact a particular service. One or more location URIs may be included in a discovery response or a presence announcement.

Expiration information identifies how long a SSDP client should keep information about the service in its cache. Once the entry has expired, it is to be removed from the SSDP client's cache.

Thus, a SSDP client service cache might look like:

USN URI	Service Type URI	Expiration	Location
upnp:uuid:k91...	upnp:clockradio	3 days	http://foo.com/cr
upnp:uuid:x7z...	ms:wince	1 week	http://msce/win

In the previous example both USN URIs are actually UUIDs such as upnp:uuid:k91d4fae-7dec-11d0-a765-00a0c91c6bf6.

3.4.2 Strengths

Using standards SSPD does not have to reinvent the wheel again. URIs and URNs are used to identify services, Caching via UUIDs and all messages are sent via HTTP. It provides the bare minimum in search functions, which is good for low-end products like digital cameras. It does not take much space and with caching, it does not put a heavy load on the network as there are only two messages needed to discover a service provided they are both received the first time. An announcement is sent when a service becomes online, and a discovery request is sent when a client becomes online which results in a reply from the service. SSPD inside Universal Plug and Play (UPnP) is available for Windows and Linux.

3.4.3 Weaknesses

SSDP uses Device Control Protocols (DCP) to allow communication between a service and a client once the service is discovered. These DCPs are supervised by the Universal Plug and Play (UPnP) Steering Committee. Every DCP needs to be registered by the UPnP Steering Committee, which does slow down the development of new devices and DCPs. SSDP does have some form of protection against network overloads with caching. However when every client comes down during a power shortage, it can overload the network if the number of services and clients is too large.

3.5 Service discovery and network administration

Each of the three protocols is being used in the field now. SLP is used with a KDE version of VNC to locate VNC clients, DNS-SD is being used within MacOS X and iTunes (and many more apple applications) and SSDP is being used in Microsoft's UPnP and Intel's open source implementation of UPnP for Linux. Each however has their own strengths and weaknesses however each share the same weakness that makes Zeroconf in general unfit for very large corporate networks. Service discovery protocols can be a huge drain on network resources. SLP uses Directory Agents to compensate for this and SSDP uses caching and a hybrid of announcement and discovery request messages. DNS-SD only uses discover messages and thus run the risk of overloading the network. However, because only discover messages are used, it should be a low risk. Not every user is requesting for printers at the same time after a power loss.

Low-end devices have the most benefit of DNS-SD and SSDP. SLP has large messages and allot of options most devices will not use or require from network services themselves. This can be seen by the choice of Apple Computing and Microsoft, which both avoid SLP.

4 Zeroconf Analysis

4.1 Benefits of Zeroconf

The benefits of Zeroconf protocols over existing configured protocols are an increase in the ease-of-use for end-users and a simplification of the infrastructure necessary to operate protocols. Zeroconf can also offer advantages on maintenance points in an IT infrastructure, though this is not the primary goal of Zeroconf. Since it was developed with the goal to enable the creation of entirely new kinds of networked products, products that with the current technologies would simply not be commercially viable because of the inconvenience and support costs involved in setting up, configuring, and maintaining a network to allow them to operate.

In assigning IP configuration to hosts in a network Zeroconf is not meant to compete with DHCP, but to work together. If a host detects a DHCP server, it receives a dynamic IP configuration. If there is no DHCP server present, the host will assign itself an IP configuration by use of Zeroconf.

Zeroconf can also be useful for a DHCP server in solving the DHCP bootstrap problem. A DHCP server can assign an own IP configuration with Zeroconf, start-up and then assign the other hosts in the network a dynamic IP configuration.

4.2 Security issues

We are talking about the networks that do not require an administrator to configure and maintain the network. But what happens to the security of the network?

When connecting a host on an IPv4 network and assigning an IP configuration using Zeroconf, the host will be vulnerable for potential IP-based attacks, such as ARP poisoning attacks. IP over networks use Address Resolution Protocol (ARP) to resolve IP addresses into hardware, or MAC (Medium Access Controllers) addresses. All the hosts in the LAN keep a cache of resolved addresses. ARP resolution is invoked when a new IP address has to be resolved or an entry in the cache expires. The ARP poisoning attack consists of maliciously modifying the association between an IP address and its corresponding MAC address.

Furthermore, when performed on two different hosts at the same time, ARP poisoning enables an adversary to launch a Man In The Middle (MITM) attack. With MITM attacks, traffic between two hosts is redirected through a third one, which acts as the man in the middle, without the other two knowing it. The MITM may simply relay the traffic after inspecting it or modify it before resending it. MITM attacks are possible at various layers of the OSI stack. ARP poisoning allows hackers to perform such an attack at the data link layer.

At the network layer, the attack exploits DNS poisoning. The attacker first modifies the DNS tables to associate its own IP address with the symbolic names of both victim hosts. Thus, when the victims will query the DNS asking for the each other's IP address, they will receive the attacker's IP address. At this point, all the traffic between the two hosts will first be received by the attacker that will forward it to the respective destination, after possibly modifying it.

There are some other concerns in using Zeroconf address assignment. When a host in a network gives up an address and reconfigures, another host allows the possibility to easily successfully hijack existing TCP connections. Therefore, a host should reset any existing connections that are using the address, before abandon the address due to a conflict.

Zeroconf enables legitimate users to discover what services are available on the network. For hackers it does not make it more difficult to discover the running services, but the fact is that a hacker can already find out what is on the network with widely available tools on the internet.

Zeroconf protocols must not be any less secure the related current IETF-standard protocols. Zeroconf protocols are intended to operate in local link networks in parallel with standard configured network protocols.

Not using Zeroconf does not make life harder for hackers; they already know how to find machines, it only helps your employees. Not using Zeroconf "is like hitting everybody who comes to work with a baseball bat because it will help discourage thieves".

4.3 Requirements in using Zeroconf

Zeroconf is ambitious in its network interface configuration-less solution, but there are some fixed difficulties in the implementation of Zeroconf. Requirements are set up for Zeroconf to make networking as easy as possible. In some cases, other concerns may control ease of use. For example, network security requires some configuration that may not be as easy as the unacceptable alternative of 'no security.'

The following requirements are defined by the IETF.

- Sufficient security features to prevent networks from being any less secure than networks that do not use Zeroconf protocols.
- When Zeroconf networks or hosts, which are configured using Zeroconf protocols, are connected to the big Internet, they should not automatically become vulnerable to new security threats.
- Zeroconf protocols MUST minimise their impact on existing networks.
- Zeroconf protocols SHOULD minimise their impact on existing applications.

- Zeroconf protocols MUST NOT be any less secure than related current IETF-Standard protocols.

There has been lots of research in the Zeroconf technology, but there are still some problems unsolved, mostly due to the following reasons:

- Non-existence of real products that use the Zeroconf technology.
- Could not keep things simple.
- Failed to understand basic principles like browsing for services instead for devices

5 Applicabilities of Zeroconf

Zeroconf is a solution for connecting equipment in a network environment where it is impractical difficult or almost impossible to configure the equipment. An example is to connect a couple of different laptops on a conference without any form of centralized network structure such as a MADCAP-, DHCP- or DNS-server. Another usage of Zeroconf is to connect a PDA in a study classroom in an educational institute to download digital study information without configuring the PDA on the network.

Zeroconf can also prove it is usability in situations where it is necessary to configure in short time computer equipment in a complete network infrastructure. In humanitarian operations, for instance it is necessarily to couple fast and efficient medical equipment and GPS-systems that work together in a local link network.

In the area of grid computing, Apple uses its Rendezvous technology (Apple's implementation of Zeroconf) with the software Xgrid[4] to cluster a set of Macintosh computers. Xgrid manages the set of computers and generate for each computer an own IP number with Rendezvous.

5.1 System administration purposes

In the field of system administration Zeroconf can add benefits to simplify the work of a system administrator. The following applications are already implemented with Zeroconf technology, and more will definitely follow.

5.1.1 ZCIP

Zcip is an implementation of the ad-hoc link-local IP auto configuration algorithm described in the IETF Draft "*Dynamic Configuration of IPv4 link-local address*[5]". It is the IPv4 address assignment implementation of Zeroconf on Linux systems.

5.1.2 KDE VNC

KDE VNC is a program for Linux systems that is included inside KDE. It gives the KDE desktop the ability to remotely use a desktop. The Service Location Protocol is used to discover VNC clients that allow remote login via VNC. This saves the user the trouble of remembering host addresses.

[4] Xgrid can be found on: http://www.apple.com/acg/xgrid/
[5] See Reference: http://files.zeroconf.org/draft-ietf-zeroconf-ipv4-linklocal.txt

5.1.3 OpenSLP

Service Location Protocol is an IETF standards track protocol that provides a framework to allow networking applications to discover the existence, location, and configuration of networked services in enterprise networks. Traditionally, in order to locate services on the network, users of network applications have been required to supply the host name or network address of the machine that provides a desired service. Ensuring that users and applications are supplied with the correct information has, in many cases, become an administrative nightmare.

SLP can eliminate the need for users to know the names of network hosts. With SLP, the user only needs to know the description of the service he is interested in. Based on this description, SLP is then able to return the URL of the desired service.

5.1.4 CUPS

CUPS is an implementation of the Internet Printing Protocol, and is rapidly becoming the industry standard for network printing. Internet Printing Protocol does not specify a way of locating network printers, but CUPS has its own protocol. Unfortunately, the CUPS printer discovery protocol is broadcast based and uses a lot of network bandwidth, to the extent where some organisations turn it off.

However, CUPS also includes a Service Location Protocol option, which is built if the configuration process finds suitable system libraries (nearly always OpenSLP). This can offer a standardised way of locating any printer on the network, without excessive network traffic, potentially including printer location across routers.

5.1.5 Rendezvous applications

Rendezvous is Apple's implementation of Zeroconf. In the operating system Mac OS X Panther of Apple there are for example a number of interesting applications supplies that make use of Rendezvous.

- The Safari browser uses Rendezvous to find any web addresses on a local network for printer, router or web cam setup and administration.
- iChat AV lets you see which people are available for chatting or video conferencing on the local network, and automatically removes them when they leave.
- Rendezvous is used in iTunes and iPhoto to facilitate sharing music and photos on local networks
- SubEthaEdit is a document-editing environment where a group of people can work together in real-time on a document. With Rendezvous, you do not have to edit long list of preferences, configuring servers and clients. Rendezvous detects the users on the network, and people can be invited to participate in the editing of the document with a click on the share button.

Printer manufacturers are integrating Zeroconf into new printers to enable them to be added and removed from networks without configuration. When a computer is added to a network, the available Zeroconf-enabled printers will automatically be discovered and connected too.

5.2 Ethical results of using Zeroconf

Zeroconf can potentially contribute to ease the division between technically able, and those who cannot access technology. The ability to use and maintain the local network is now the domain of the technical elite. It would be unethical if users with no knowledge of the technology behind a local network cannot have the opportunity to maintain there own network, especially in using ad-hoc networks.

Conversely, with a technology like Zeroconf there is no human intervention in network administration, which can result in a lack of innovation. With a smaller group of people involved in network administration, we may want to consider where the next generation of network designers and programmers will come from.

6 Conclusion

With Zeroconf, a system administrator can save a lot of time when installing for example 100 computers with four printers. A disadvantage of the usage of Zeroconf for configuring so many hosts is the huge amount of network traffic it generates for every shared resource.

Zeroconf is now in its early stages of development. However, if Zeroconf networking is suitable for medium or large networks with enough security, it can change the world of networking with no network administrators to configure and maintain the networks.

According to our results, we can say:

- Zeroconf could work for a hundred or even a thousand hosts, but if you have a hundred computers on your network, you can probably afford to install a DHCP server.
- The use of local link networks without physical security, like LANs with wireless base stations, is an irresponsibly danger for users that connect to network resources. The security of Zeroconf is actually based on the security that underlying protocols provide.
- Service discovery protocols have uses when it comes to minimizing the costs of maintaining large computer networks. It can save allot of time, configuring every workstation in a dynamic network environment. However, service discovery protocols have risks that should be fully understood. Networks can be completely overloaded when large amounts of services and clients are searching the network after a power loss. Division of the network is needed when the amount of services becomes too much. It is probably not needed to be able to print documents on 200 different printers located throughout the building or throughout the world.

- In all cases, whether or not link-local IPv4 addresses are used, it is necessary for implementers of devices supporting the Internet Protocol to analyze the known and credible threats to which a specific host or device might be subjected. In addition, provide security mechanisms, which ameliorate or reduce the risks associated with such threats.

7 References

The charter website
http://www.ietf.org/html.charters/zeroconf-charter.html

The official website of the Zeroconf Workgroup
http://www.zeroconf.org

Dynamic configuration of link-local IPv4 addresses
http://files.zeroconf.org/draft-ietf-zeroconf-ipv4-linklocal.txt

Updated Draft: Dynamic Configuration of link-local IPv4 Addresses
http://www.ietf.org/internet-drafts/draft-ietf-zeroconf-ipv4-linklocal-11.txt

Understanding Zeroconf and Multicast DNS
http://www.oreillynet.com/pub/a/wireless/2002/12/20/zeroconf.html

Understanding Zeroconf and Multicast DNS by Heath Johns
http://www.oreillynet.com/pub/a/wireless/2002/12/20/zeroconf.html

DNS Service Discovery Draft
http://files.dns-sd.org/draft-cheshire-dnsext-dns-sd.txt

RFC2608 SLPv2
http://www.ietf.org/rfc/rfc2608.txt

SLPv2 analysis by Yaron Y. Goland
http://www.goland.org/Tech/slpv2.htm

Zeroconf overview
http://www.linux-mag.com/2003-08/zeroconf_01.html

Zeroconf Official Issue List
http://www.ietf.org/html.charters/zeroconf-charter.html

Apple's Xgrid software for grid computing
http://www.apple.com/acg/xgrid/

Zcip homepage
http://zeroconf.sourceforge.net/?selected=zcip

SSDP draft v1.03
http://upnp.org/download/draft_cai_ssdp_v1_03.txt

Rendezvous applications
http://www.apple.com/macosx/features/rendezvous/

Rendezvous Technical FAQ
http://developer.apple.com/macosx/rendezvous/faq.html

Security Considerations
http://files.zeroconf.org/draft-ietf-zeroconf-ipv4-linklocal.txt