

Common Home Network Attacks, Security Principles and Risks

Mark Baugher, Cisco
mbaugher@cisco.com

ABSTRACT: This paper describes some attacks on home networks in general, and UPnP™ home networks in particular. The paper analyzes attacks in terms of the home network assets, risks and threats. In some cases, the home network is vulnerable to attack because access cannot be controlled and the service protocol violates basic principles of secure design. Although the paper focuses on UPnP, most of what is said about UPnP applies to other protocols such as Apple Bonjour™.

1 Introduction

The UPnP Internet Gateway Device (IGD) standard for residential gateways is increasingly important to “connected home” products. IGD version 1 is a good example of an attack-prone protocol for the home network. In particular, IGD has a port-mapping service called “NAT Traversal”, which is needed by many home-network applications. For example, AppleTV™ currently uses UPnP IGD NAT traversal protocol, which allows Apple’s servers to send unsolicited packets to an AppleTV on a through non-Apple gateways that feature UPnP NAT Traversal but not Bonjour™ NAT-PMP.

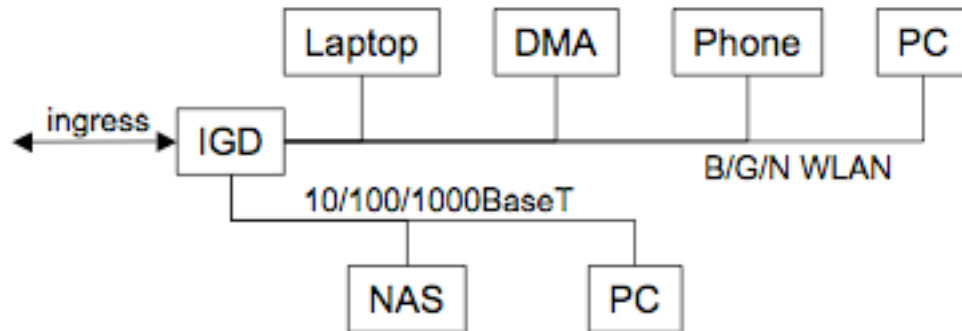


Figure 1-1: Common Home Network Configuration

The “IGD” in Figure 1-1 is the UPnP Internet Gateway Device service in a residential gateway/router that connects the residential network to the service provider network over the interface labeled “ingress”; the gateway/router typically hosts a firewall and a NAT service in addition to an IGD service. When running dual-stack, IPv4 and IPv6, the gateway/router has no NAT for IPv6 traffic but likely has an IPv6 firewall instead.

The “B/G/N WLAN” in Figure 1-1 is a wireless local area network (WLAN) service. Most residential networks in North America and other parts of the world include a WLAN running the IEEE 802.11 B, G or N services. Thus, a WLAN access point (AP) is shown as one of two IGD interfaces in Fig 1-1. Fixed and mobile devices such as laptops, Wi-Fi phones, and other devices use the WLAN. Digital media adapters (DMA), media extenders, and game consoles also typically connect over the WLAN to the home TV, the home server and to the Internet. Home servers are often configured on the “10/100/1000 BaseT” Ethernet, which is show in the figure. These interior devices might use the IGD NAT Traversal service to install a port mapping in the gateway’s NAT; the port mapping would allow unsolicited packets from computers that are exterior and remote to the home network.

Unfortunately, IGD NAT Traversal, network configuration and other services are deemed insecure by many security experts, who recommend that all IGD services be disabled by default. Most vendors of home-networking appliances currently ship gateway products with IGD *enabled* by default.

Shipping products that leave IGD off by default would require an additional configuration step for the user; this additional complexity needs to be weighed against fact that many users will skip a confusing or unwanted security step. Certainly the potential for attacks exists, and the claim that IGD is insecure is credible for particular attacks, including some potentially very serious attacks.

IGD attacks are described in section 2 of this paper. Section 3 considers home network security independently of UPnP and identifies the assets that need protection, the risks to those assets, and the threats that lead to successful attacks. Section 4 returns to UPnP Device Control Protocols (DCPs) and considers protocol design¹ in light of the set of security principles. Section 5 concludes the paper with a brief summary.

2 Attacks on UPnP IGD

There are several documented attacks on UPnP IGD services. Armijn Hemel authored a paper on UPnP security flaws that describes several problems in IGD [Hemel]. Hemel gives an account of problems with the port mapping features found in the IGD WANIPConnection and WANPPConnection profiles, and he mentioned a problem with the LANHostConfigManagement profile that was later exploited and published on the GNUCitizen website [GNUCitizen].

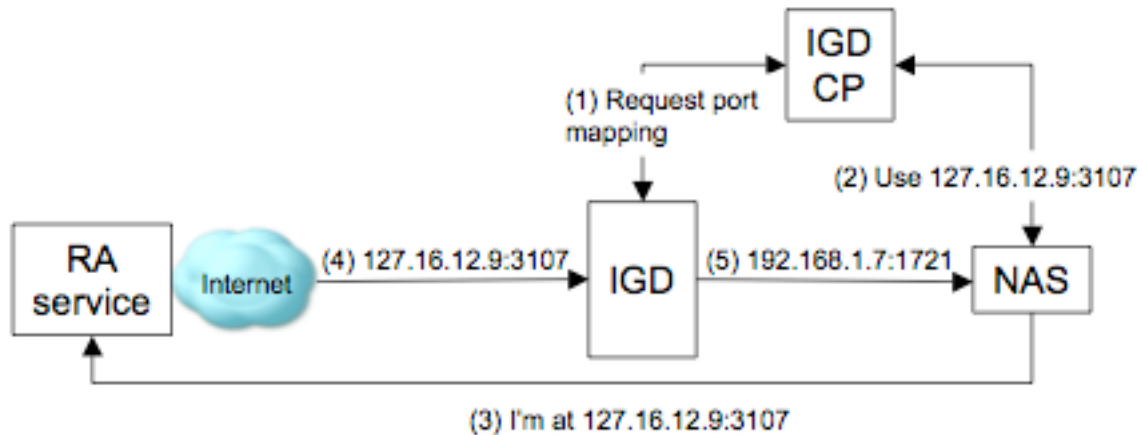


Figure 2-1: Example IGD Port Mapping Flow

Figure 2-1 illustrates the process of IGD port mapping. An application on the residential network requests a UPnP IGD control point (CP) to identify the public internet address used by the IGD and to establish a port mapping for a 3rd device (Step 1). In this case, a network attached storage (NAS) device receives a port mapping to its home-network IPv4 address. The IGD CP is a logical entity that could reside in the NAS or some other device. For example, an IGD CP is incorporated into some NAS and network media products (like AppleTV). Not shown in Fig. 2-1 is the application process that directs the IGD CP to request the port mapping, which might be initiated by a configuration script in a NAS device, for example, or through a user interface on a portable device. The request to the CP must specify the local address and port (of the NAS) to use on the residential network side of the mapping. When the IGD CP request is successful, The CP reports the public address and public port that the IGD has mapped to the local address and port (Step 2). With the mapping between the public address and port to the local address and port in place, the NAS or other device can publish that address (Step 3) for use by a remote access service on the Internet or by an external device (Step 4).

¹ Often, the “protocol” is an XML representation of an Action against a state Variable and carried “on the wire” in an HTTP Get.

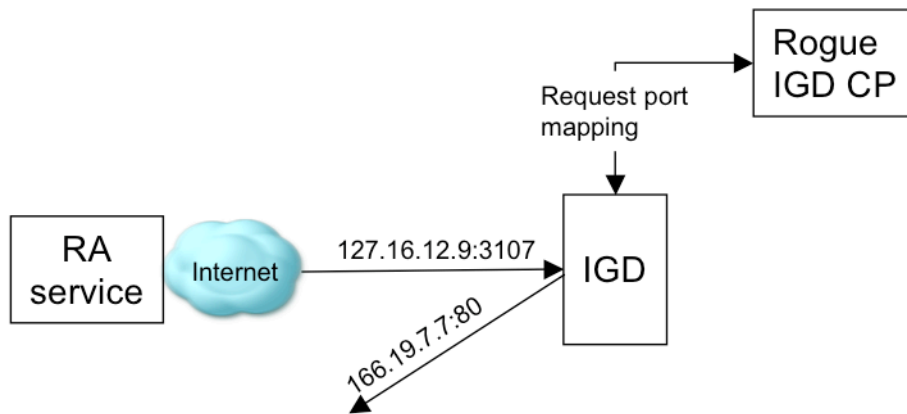


Figure 2-2: Unwanted mapping to an external address

UPnP IGD does not restrict the mapping from a public address and port to an address and port on the residential network: It is allowed for the mapping to be from the public address and port to another public address and port because of an error in the UPnP IGD:1 specification². Consider what happens if the CP is under the control of malware, such as from a virus or Trojan horse that gets installed on a home PC. This “Rogue Control Point” could easily redirect the mapping to go to a criminal’s web server or some other nefarious site on the Internet. In UPnP, the intended destination device is not necessarily notified when a mapping installed or changed.

Port mapping is an essential service for the next generation of home networking applications such as remote access (RA) service, and the attack against this service is relatively easy. We can expect to witness such attacks as use of IGD port mapping becomes more common and criminals discover uses for this particular type of attack³.

A second type of attack is against a more sophisticated and potent service: A Domain Name Service that is hosted on the IGD. The IGD LANHostConfigurationManagement profile hosts this service and allows an IGD control point to alter entries in the DNS.

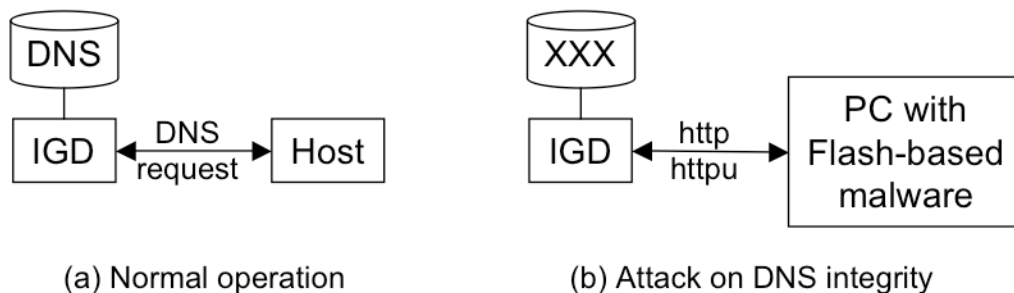


Figure 2-3: Flash-based attack on IGD LANHostConfigurationManagement

The attack against LANHostConfig is shown in Fig. 2-3 and could cause serious losses to users. As depicted in the Fig 2-3(a), an IGD that supports the LANHostConfigurationMangement profile with a DNS server was at great risk from Adobe Flash malware that can be easily deployed onto a host PC

² These problems are being corrected in IGD:2.

³ The Conficker virus uses IGD NAT Traversal, for example.

from an Internet website, email or other source⁴. The Flash application has access to the Flash API, which allows HTTPU⁵ messages to be sent to the IGD. In this attack, the malware formats IGD CP commands in the HTTPU messages, sends them to the IGD; the message reconfigures the DNS service to direct the unsuspecting home user to a website suitable for a serious phishing attack. This attack has been documented and demonstrated against a commercial product [GNUCitizen].

The Flash-based attack on LANHostConfiguration services poses a very great risk where the unsuspecting user could turn over usernames and passwords for online banking and other accounts. This attack and the other threats described in this paper often do not have simple or even certain remedies. In some cases, the attack can come from a “war driver” who gets access to an unprotected WLAN from the street or adjacent residence. Although Wifi Protected Access can keep a war driver out of the residential WLAN, there is the potential for mis-configuration or misuse of the protection mechanisms. Moreover, It is much more difficult to offer protection against viruses, Trojans, and other types of malware that execute on PCs and other devices on the residential network.

Given the risks of LANHostConfigurationManagement and the limited benefits of hosting DNS on the residential network, this paper recommends against deployment of this IGD profile. The port mapping of WANCommonInterfaceConfiguration, WANIPConnection and WANPPPConnection, however, is arguably far more useful, less risky, and easier to protect.

3 The Risks of Home Networking

The previous section described some real-world attacks on a specific home-network protocol. Home network protocols actually predate home networks: IBM/Microsoft NetBIOS and AppleTalk™ appeared on small, private local area networks (LAN) for enterprises and schools in the 1980s and 90s. UPnP and Bonjour™ protocols are the descendents of these earlier ones and share their properties: A low need for configuration but a great need for one-to-many packet communications across a local network. This class of protocols is “chatty” compared to what’s typically found on the Internet or in the enterprise; a chatty protocol sends a lot of information directed to any and all devices that get access to the network. UPnP and Bonjour devices automatically discover one another with multicast and describe services over subsequent unicast message exchanges – all with little or no user intervention. UPnP was originally intended to mean “Universal Plug and Play”, although it is not a trademark of the UPnP Forum™. Nonetheless, it is a “little p” plug-and-play network that’s intended to make it easy for any person or computer program to request services. This ease of access creates new vulnerabilities on a typical home network.

The risks of plug-and-play vary by device. An attack against a gateway router has arguably more risks than an attack against a personal device, to the extent that many devices can be attacked through a compromised gateway router. The section considers assets, risks and threats [PHB] to those assets, and it considers basic security techniques to mitigate these risks.

3.1 Assets, Risk and Threats

Residential networks have critical assets such as gateway devices, personal computers, firewalls and network storage. Among the biggest risks to these assets are the re-configuration of network devices and theft of personal passwords. By re-configuring the DNS server name, for example, an attacker can do a pharming attack. A phishing attack steals passwords to get access to online banking accounts and password-protected devices. Malware is a well-known attack vector for pharming and phishing attacks on home networks [GnuCitizen]. Computer viruses, trojan horses

⁴ Adobe has since fixed the problem and UPnP may randomize URLs in the future.

⁵ Unacknowledged HTTP on UDP

and other types of malware get routinely downloaded and installed on programmable devices in the residential network.

Another attack vector is "war driving", by which an interloper typically uses an open wireless LAN to gain access to a residential network device, typically for non-intrusive Internet access but a potential threat nonetheless. The easiest way to stop a "war driver" is to ensure that Wi-Fi Protected Access is enabled on the WLAN, and similar steps are needed for other shared-media networks such as powerline networks and even MoCA networks where signals can overlap between neighbors.

Unlike war drivers, there is no straightforward way to limit the effects of malware once it is established on a privileged device on the home network. One possible remedy is to not give privileges to computers that run scripts. On some operating systems, simply putting the user in the loop of a privileged request (like install a program) prevents a great deal of malware, but the problems with known techniques are discussed below.

3.2 Authentication and Authorization

Strong identification, authentication and authorization can prevent threats to residential networks from war drivers, visitors, and other interlopers who gain access through an open wireless LAN or other means. Nonetheless, malware can gain execution privileges on an authorized end system, such as a personal computer user account having the privilege to set the DNS name in a residential gateway. Thus, automated methods of authentication using public-key or secret key cryptography are sometimes insufficient. In the case of malware, multi-factor authentication such as device public-key authentication coupled with a user passphrase puts the user in the loop. Multi-factor authentication can potentially prevent malware from executing its actions on the host device. But there are human-factors problems when the user is in the authorization loop: The user might be conditioned to approve every action and type in a password whenever prompted to do so, for example. As discussed below, password-based authentication comes with additional risks.

3.3 Problems with Password-based Authentication

In general, passwords are a poor authentication method; this has been true for some time [Neumann, RT79]. And it is truer today given advances in hardware speeds and password cracking [Elcomsoft]. It is possible that advances in password security engineering can improve how people use passwords in an unmanaged environment such as the home [Anderson]. Practically speaking, however, there is no proven, simple method to ensure that passwords are strong and unique across unmanaged residential-network devices. Use of identical and similar passwords for a variety of purposes such as for firewall control and online banking, increases the risks of password compromise. A combination of techniques such as public-key cryptography, passwords with password checkers, strong pre-shared symmetric keys, hardware token devices and other means are referenced in current standards [WPS] [UDS1.0]. These methods have gained little use on home networks, which are by and large unmanaged. Home network services such as NAT traversal and firewall control continue to be unauthenticated and risky.

The problems of controlling access to personal and shared devices have no easy solution on the home network. Thus far, conventional security techniques using public-key cryptography, authenticated key establishment, and secret key encryption have not proven effective for individual users in a personally managed or unmanaged home network⁶.

⁶ UPnP Device Security 1.0 is a state of the art security system that has not been deployed in the same way that PGP mail and personal voice security systems are rarely used for personal mail and VoIP.

3.4 Unmitigated Risks

Authentication and authorization are hard problems in unmanaged networks of personal devices. A one-time procedure is usually needed for a human user to prove locality or control as a precondition for an authorization [WE]. An initial authorization for a firewall control interface might be an authorization for packet forwarding for an internal IPv6 GUA or an IPv4 NAT port mapping, for example. A more privileged authorization might be to request packet forwarding for another device, such as a visitor to the residential network.

To establish cryptographic security associations, some explicit action or actions by the user is necessary. This is sometimes called “secure introduction”. Thus far, secure introduction over the wireless network has large-scale user adoption, but anecdotal evidence suggests that this relies on the PC installation script and not on the many methods offered in Wi-Fi Protected Setup [WPS]. The PC script approach usually relies upon the device being shipped with a well-known password, e.g. “admin”. Most commercial WLAN devices are at risk from the moment they are installed until they are configured using secure introduction, something which many PC installation scripts lack. Furthermore, secure introduction as done in WPS and other WLAN systems is not enough for other home-network devices. The WLAN secure-introduction method does not scale to the entire home network because it works only on an individual device (e.g. an 802.11 AP), not on a service made up of multiple devices.

Thus, there are at least three unmitigated risks on the home network for the user who wishes to have exclusive access to home-network devices.

1. People today commonly place strong access controls on their Wi-Fi networks but not using the secure-introduction methods of WPS [WPS]. More commonly, PC-based installation scripts are used for Wi-Fi access points, servers, etc. that rely upon a well-known password and have a window of vulnerability when an attacker can get control of the AP or enroll the WLAN client. This results from human-factors problems that even the best vendors have not solved⁷.
2. The well-known password problem is ubiquitous on the home network. In some cases, a WLAN that is left open will allow an attacker to use a well-known password to reconfigure the gateway. The most popular home NAS products also have this problem. It is less a human factors problem than it is a technical failing of the industry.
3. Password-based authentication, WPS, and similar techniques do not scale beyond a single device: If one has N devices and a controller that works on all of them, for example, that user will in all likelihood refuse to run N secure introductions. Use of an epidemic or gossip protocols [Ford] is one approach to scale a secure introduction to a service, i.e. to multiple devices. But these techniques have not appeared in the marketplace yet.

4 Security Principles for Home Networking

The set of risks and vulnerabilities that are characteristic of most home network protocols are found in UPnP device control protocols (DCP). Section 3 describes problems with implementing access controls in home network protocols, and this applies to UPnP, which suffers the attacks described in section 2. Thus, certain attacks are endemic to all or most home network products. But many of the attacks described in section 2 are the result of specific flaws in a particular UPnP DCP. This section analyzes particular DCP flaws as violations of recommended security practices and principles.

Years of practice and research into operating system, file system, and networks present the engineer with a vast amount of relevant security theory and practice. Some excellent books are available on

⁷ This problem is found in practically all commercial home gateways.

security engineering [Anderson]. This section has the more modest aim of trying to draw lessons from the published attacks on UPnP and in particular what was wrong with the particular UPnP Device Control Protocol that led to the attack.

4.1 Least Privilege

Some UPnP DCPs allow too much access with too little control. A DCP offers “too much access” when it gives a requester of its service more access than the requester needs to do what it needs to do. UPnP Internet Gateway Device protocol, for example, implicitly authorizes a control point that needs a port mapping to its own address to similarly request a port mapping for another device, such as a UPnP renderer [Hemel]. A third-party port mapping has effects on systems other than the requester and thus violates the “principle of least privilege” by giving more access privileges than needed to a device control point that needs only a port mapping only to itself.

A DCP would have more control if it restricted an unauthenticated control point to perform an action that affects only itself, such as to direct a flow or open an external port only to its own address⁸. There is greater risk when a control point is permitted actions that affect other devices. In this case, there is an obvious need to separate privileges between those that only affect the requester and those that can affect other devices.

This security issue is intrinsic to the UPnP design, which separates the control point function, such as remote control, from the device function, such as a UPnP TV. If external access is needed to the TV device⁹, then the control point needs to have the privilege of establishing a NAT port mapping for the TV. But the UPnP design gives the control point the privilege of establishing NAT port mappings for any device on the home network even if it is dedicated to controlling only one particular device. A more secure design embeds the control point in the device, so it shares an address with the device. An even more secure design would use strong authentication of control point action requests, although cryptographic security is hard to do on an unmanaged network.

4.2 Privilege Escalation

UPnP control points and devices typically share a residential network with other systems, including non-UPnP systems. UPnP implementations often are sub-systems in shared devices such as gateway/routers and media servers that have their own access-control mechanisms. In the case of UPnP IGD, for example, an unauthenticated control point can perform actions on the gateway/router that configure a DNS server name, shut off Internet access, and perform other actions that are normally accessed through a passphrase-protected administrative interface. In this way, a poor DCP implementation can escalate the privileges of a control point without the vendor-defined level of authentication. That has happened in at least one commercial product [GnuCitizen]. The DCP in this case becomes a “back door” to unauthorized administrative access through what is essentially an implementation error, i.e. the vendor should not require authenticated access to the same data through one interface but unauthenticated access through another.

4.3 Privilege Separation

As described above, a privilege granted to a UPnP control point that affects one particular UPnP device usually needs to be separate from a privilege that affects other UPnP devices. Similarly, a

⁸ It is noted that checking the IP or medium access control address can be easily spoofed by a source, and a check of the source address is a very weak form of authentication. Since UPnP has already shipped its NAT traversal function, backward compatibility of future releases makes it impossible to do much more than a source-address check.

⁹ Apple TV uses UPnP NAT traversal for allowing unsolicited packets from external sources to reach it.

privilege for a UPnP sub-system in a device should not affect non-UPnP devices. In particular, the privilege of administering a shared resource such as a gateway/router or a media server needs to be separate from administering the UPnP subsystem that is embedded in that gateway/router or media server. If the system vendor provides an authentication method for administrator authorization, the UPnP DCP authentication should use this method and be at least as strong if it is to modify administrative variables.

Privilege separation is needed within a UPnP DCP. Within a DCP, the privilege to *grant* access to a service or data is higher privilege than *getting* access to it. For example, a UPnP Audio/Video control point will obtain data from a UPnP server device only to send it to a UPnP renderer device. A device that grants a device access to its data needs to separate two cases. The first case is when the receiver keeps the data to itself. The second case is when the receiver shares the data with another device.

In general, a control point that receives data from a device is receiving one level a privilege that does not necessarily authorize it to become a sending device for those data. It may be implicitly understood in a DCP that data sent to a control point will be shared to another device, but this separation of function in the UPnP design means that one device is providing its data to another device without authentication. Secure operation of a DCP might require changes to the protocol to accommodate needed access controls, which probably were not considered during DCP design and development.

4.4 Security Review

Some UPnP DCPs appear to have had no effective security review. It is not always possible to identify all risks in protocols before they are released, let alone fix all problems. But past practice has shown that most flaws become apparent when security analysis accompanies the protocol design. And good engineering practice is to review the security of implementations of the protocol in products prior to shipping those products. This step needs to be formalized inside standards development organizations like UPnP and within companies that ship the protocols in products.

4.5 Security Audit

UPnP like many standards development organizations do not have a formal process for security audit of milestone documents and specifications. Unlike a security review, which can be done with resources internal to the organization, third-part audits of security experts is needed for home networking protocols and products. An audit is best done early in the process of specification, such as requirements, prior to implementation such that the architecture and design are evaluated.

5 Postscript and Summary

Since this paper was originally written, the Conficker virus was found to use UPnP NAT Traversal for file transfers and “bot” control [Conficker]. In this case, the NAT traversal service makes it easier for the Conficker botnet to operate, just as it makes it easier for users to support remote access to their home networks. This is true for Bonjour as it is for UPnP: If a service cannot discriminate between a legitimate user program and a bot, then both will have the benefits of the service.

Access controls can separate the legitimate user from the interloper in the case of a war driver: Most Wi-Fi networks are protected against this avenue of potential attack. But regardless of the home network protocols that it supports, most commercial home gateways are highly vulnerable to war drivers when those gateways operate with a well-known administrative password. And all home networks are vulnerable to malware that executes from a privileged account. Home network access controls are not up the task of stopping malware.

Although superior product features or secure protocol design cannot mitigate all potential or even serious security risks, the security situation can certainly be improved over the current state of the art. This paper recommends a set of principles and practices that should inform the design of UPnP

and other home networking protocols to improve the current level of security in UPnP. Most of these improvements apply equally well to Bonjour and other home networking protocols.

6 References

[Anderson] Anderson, R., "Security Engineering", 2001.

[Elcomsoft] "ElcomSoft Breaks Wi-Fi Encryption Faster with GPU Acceleration", October 2008.

[Ford] Ford, B., UIA: A Global Connectivity Architecture for Mobile Personal Devices, PhD Thesis, MIT, September 2008,
<http://www.brynosaurus.com/pub/net/phd.pdf>

[GNUCitizen08] <http://www.gnucitizen.org/blog/flash-upnp-attack-faq/>

[Hemel06] Hemel, Armijn, Universal Plug and Play: Dead simple or simply deadly, 5th System Administrator and Network Engineering Conference, SANE 2006, May 15-19, 2006,
<http://www.sane.nl/sane2006/program/final-papers/R6.pdf>

[Neumann] Neumann, P., "Risks of Passwords (<http://portal.acm.org/citation.cfm?id=175289>)", April 1994.

[PHB] P. Hallam-Baker, Asset, Risk, Threat (ART) Analysis,
<http://www.mbaugher.com/StructuredSecurityPolicy.pdf>

[RT79] Morris, R. and K. Thompson, "Password Security: A Case History", November 1979.

[UDS1.0] Ellison, C., "DeviceSecurity:1
(http://www.upnp.org/standardizeddcps/documents/DeviceSecurity_1.0cc_001.pdf)", November 2003.

[WPS] Wikipedia, "Wi-Fi Protected Setup (http://en.wikipedia.org/wiki/Wi-Fi_Protected_Setup)", February 2009.

[WE] Walker, J. and C. Ellison, "UPnP[TM] Security Ceremonies Design Document
(www.upnp.org/download/standardizeddcps/UPnPSecurityCeremonies_1_0secure.pdf)", October 2003.