

A Case Study in Practical Security of Cable Networks

Amir Alsbih¹, Felix C. Freiling², and Christian Schindelbauer¹

¹ Albert-Ludwigs-Universität Freiburg, Germany

² Universität Mannheim, Germany

Abstract. Cable networks are complex systems that have evolved over years and in which new features like Internet access and Voice over IP (VoIP) have been integrated. We argue that threat models must evolve together with such systems and show that inadequate threat models can be used to explain known and unknown vulnerabilities in today's cable networks. We do this by demonstrating an attack on the DOCSIS provisioning standard in cable networks. By exploiting this weakness, an attacker can hijack VoIP accounts. We also show how to mitigate the attack.

1 Introduction

1.1 Motivation

Cable networks were initially deployed as a low-cost means to broadcast television programs to customers, first analog and then also digital. Today, almost all cable networks have been re-engineered so that they are able to additionally transport arbitrary digital data both to and from the end user. The number of Internet users in Germany and elsewhere that access the network in this way is rising sharply [5]. Within the system of cable networks there are three main stakeholders: First of all, there is the user (customer) who wishes to enjoy digital media and network access in high quality at low cost. Second, there is the cable network provider (CNP), a company running the physical cable network infrastructure. Third, there is the Internet service provider (ISP) who provides access to the global Internet. Traditionally, the context and mindset of the CNP is a closed, physical network with static functionality. Since the services of an ISP were added to the portfolio of the CNP only rather recently, the original mindset of the CNP may be reflected in the way the ISP is managed. Since the Internet is an open, virtualizable and dynamic network, there is a potential for misconceptions regarding security assumptions that can lead to many surprises. In this paper we show that this is in fact the case.

1.2 Context

Technically, the user's *cable modem* acts as bridge between the customers home network and the backbone of the ISP. The cable modem is a rather simple device

that downloads its configuration after every reboot from a server at the ISP. This is called *provisioning* and the relevant standard is the *Data Over Cable Service Interface Specification* (DOCSIS) [7]. So there is a way to access the ISP servers from the customer's home networks and therefore the ISP has to secure the provisioning process to restrict the possibilities of a service abuse. In this paper, we document a weakness in the way the provisioning process is handled today.

1.3 Related Work

There is relatively little work that investigated cable networks security from an academic viewpoint. Existing work mainly comes from industrial or rather applied forums and deals with possibilities of service theft, e.g., possibilities of achieving higher service bandwidth without paying for it (so-called “uncapping”) by manipulating the cable modem [1–3, 11, 15]. Other threats to ISPs have been formulated as well [1] and refer to attacks on confidentiality and weak endpoints: The classical attack to confidentiality of network traffic through eavesdropping is omnipresent in shared medium networks like the cable network is. Since the downstream traffic is broadcast across the shared medium, network providers have to enforce rigid access control techniques at the endpoints of the network to ensure confidentiality. A related threat is network access through stolen authentication credentials, e.g., cloning of MAC addresses, a problem that is hard to tackle in networks where endpoints are under complete control of the customer [2]. In this paper, we give another example for this fact. Since one of the main applications of cable networks is digital telephony (voice over IP, VoIP), attacks on the corresponding protocols like the Session Initiation Protocol (SIP) have also been investigated in the cable network environment [3, 8]. These attacks include tampering of SIP message bodies like malformed SIP messages, hijacking dedicated SIP accounts or interrupting sessions by injecting fake messages into the network traffic. While being relevant to VoIP technology, they are only specific to cable networks as far as these attacks use access techniques that only exist in cable networks. In this paper, we present such an attack on SIP that is specific to cable networks.

1.4 Contributions

In this paper, we describe the context and specifics of the system of cable networks as an evolving complex system that is in regular use all over the world. We show that different mindsets and threat models can be used to explain past and present vulnerabilities in such networks. As a confirmation, we present a new attack on VoIP in cable networks that allows an attacker to extract SIP credentials and therefore misuse the VoIP system in such networks. More specifically, our attack exploits the DOCSIS provisioning requirement that every cable modem needs a provisioning file that contains the configuration of SIP credentials. By gaining access to the management network and partial exhaustive search of the namespace of configuration file names, we are able to extract SIP credentials and take over telephone accounts of other customers. We also show how this attack

can be mitigated by adjusting network management policies in cable networks. At the time of writing, our attack was possible in one major (cooperating) cable network in Germany. Since many other CNPs all over the world use the same hardware and software configurations, we believe that the attack is relevant not only in Germany. However, the wish to point out that the attack is just a vehicle to transport the another insight, namely that threat models must be checked regularly. And if they turn out to be unrealistic, they (and all affected security procedures) must be adapted.

1.5 Paper Outline

This paper is structured as follows: We give a brief introduction into cable network technology in Sect. 2. In Sect. 3 we show how threat models for cable networks have evolved in the past and the effect of this process on the provisioning process. We present the resulting attack in Sect. 4 and possible countermeasures in Sect. 5. We conclude in Sect. 6.

2 Background

In this section we give a brief introduction into the basics of modern cable networks.

2.1 System Overview

The cable network is a complex system consisting of multiple interconnected networks (see Fig. 1). The customer network connecting the end user devices

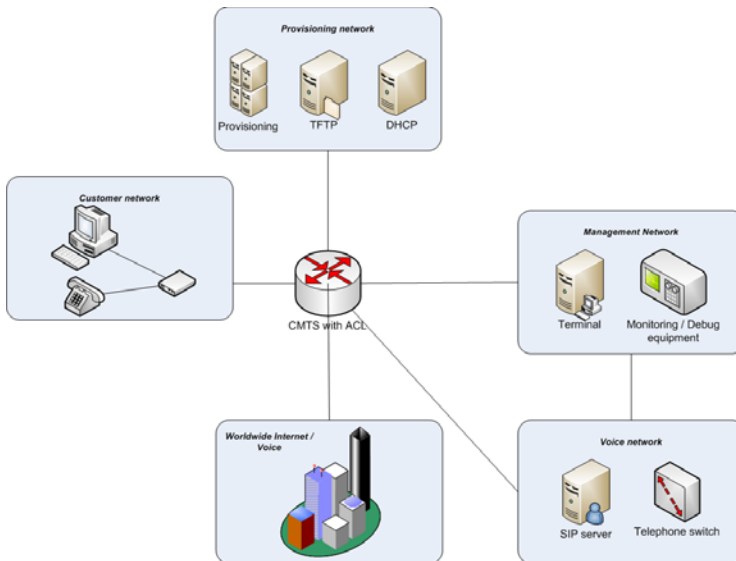


Fig. 1. Cable network reference figure

with each other is a broadcast network that can be used to transport analog and digital signals in an integrated fashion over distinguishable frequency bands. It can be thought of as a long wire to which many receiving stations can be connected, similar to the early Ethernet technology (10BaseT). Technically, the customer network is a hybrid network consisting of optical fiber and coaxial cables that is connected in a tree-like topology. Consequently, this network is called *hybrid fiber coax* (HFC).

2.2 Physical Aspects and Frequency Bands

Inside the HFC network, the frequency spectrum is divided into channels. The channels can be divided into *downstream* (towards end user) and *upstream* (from end user). Originally, all channels were downstream and carried either radio or television signals. Subsequently, frequencies were defined for data (i.e., Internet) communication. Since this communication is bidirectional, the CNP will need to provide at least one channel in each direction. On top of this digital channel, digital telephony services (VoIP) can be offered. For Internet communication, the digital signal has to be modulated on to and demodulated from the physical medium at each connected station. This is done at the side of the end user by a *cable modem*. At the side of the cable network provider, this is done by the *cable modem termination system* (CMTS). This is similar to how DSL technologies work over the telephone network.

2.3 The Interfaces of the Cable Modem

Internally, the cable modem consists at least of two interfaces, each with its own MAC address:

- The first interface is used for remote management of the cable modems from the side of the CNP. It has a unique MAC address called *C-MAC*.
- The second interface is used for realizing the Internet service for the customer. It's MAC address is called *E-MAC*.

After the cable modem has been connected to the cable and turned on, the cable modem will configure itself. This process is called *provisioning* and will be explained later, since it is in the center of our attack. The result of a correct provisioning process is that both interfaces will receive its own IP address, each IP address being in a separate IP address range.

2.4 CMTS and Access Control

If a cable modem is provisioned and tries to communicate, every communication request of the cable modem is handled by the CMTS. The CMTS plays the role of a central router, guiding network packets from the customer network to the Internet and vice versa. At the same time, the CMTS acts as a firewall, enforcing filter rules to, for example, separating the network traffic from the different IP address ranges belonging to the different interfaces of the cable modem. Filtering

is even performed on packets that are “routed back” into the customer network. This happens, for example, if two cable customers communicate with each other. So even if the cable network environment is a shared medium every upstream communication is only possible over the CMTS. The filtering rules of the CMTS are one of the most important parts of cable network security. For example, one of the most important rules is one that forbids normal users (via their E-MAC) access via SNMP to the management interface of the cable modem (C-MAC) of other customers. Without this rule it would be possible for every customer to access (and manage) the cable modem of other customers via SNMP. Since the filtering rules on the CMTS are the only reliable way to restrict the customer’s communication abilities, it is important to invest a lot of time into the right setup, making sure that the customer is only able to act in the way intended by the CNP [9, 14].

2.5 IP Layer

As mentioned above, the cable modem and the CMTS function as endpoints to transport data over the physical HFC network. The CMTS routes the IP data from the fast Ethernet backbones of the ISP to the cable network and vice versa. The cable modem works as a bridge between the network of the ISP and the local area network of the customer (see Fig. 2).

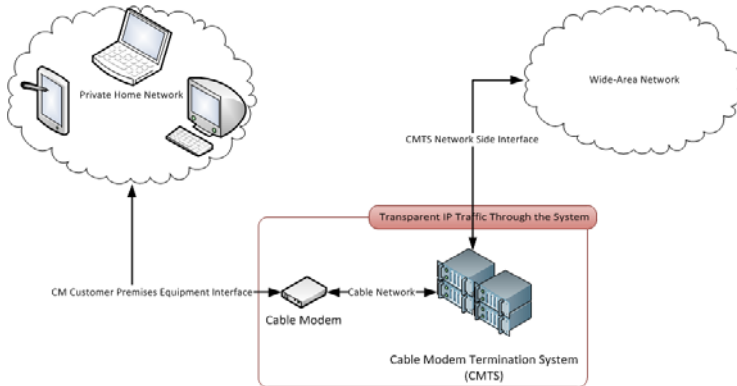


Fig. 2. IP traffic via cable modem and CMTS based on [6]

2.6 VoIP via SIP

The *Session Initiation Protocol* (SIP) is the most commonly used VoIP protocol today. SIP phone calls use two different protocols, SIP for the connection handling of the calls, and *Realtime Transport Protocol* (RTP) for the voice stream. SIP was designed for the IP world, and since IP addresses are not as static as telephone numbers, the clients have to register themselves at a SIP server, to let the SIP server know where it should route the incoming call to. To register an account on a SIP server, the customer needs (1) his own phone number, as well

as (2) the corresponding username and (3) the password for that phone number. He also needs to know the SIP server managing the connection. Note that these three items have to be known to the cable modem in order to allow seamless telephone service over VoIP for the cable network customer. Therefore, these credentials are contained within the configuration file during the provisioning process (see below).

3 Different Threat Models and Their Effects on the Provisioning Process

A *threat model* is a precise description of the possible threats to the system [16]. It usually consists of a set of security issues a system designer cares about together with a set of expected attacks. Often, threat models only exist implicitly in the mindset of the people working at the network operator and are therefore not documented within organizations. In such cases the threat model used in an organization can only be inferred through interviews and from analyzing existing security mechanisms.

3.1 Traditional Threat Model of the CNP

Traditionally, the context and mindset of the CNP is a closed, physical network with rather static functionality. Before upstream data communication was possible, the cable network was a pure “broadcast” network. The endpoints (antenna sockets in houses) were usually protected by physical means like tamper-evident seals. This was also the way how access control to the cable network worked. Since the transmitted data was the same for everyone and the selection of which channel to watch was performed at the endpoint (the television set), there were also no real privacy or confidentiality problems. Possible attacks involved only physically breaking the seal of the endpoint and accessing the service without paying.

3.2 Adapted Threat Model of the CNP/ISP

The threat situation changes dramatically if individual communication is handled via the cable network both upstream and downstream. The typical threat model used by CNP/ISP in these scenarios, however, is very similar to the original threat model. As mentioned above and from our experiences, the threat model is usually not explicitly documented. So we inferred the following assumptions from interviews and an analysis of the literature on known attacks [1–3, 11, 15]:

- The endpoint is physically protected. This means that only original cable modems are attached to the cable endpoints and these cable modems always correctly follow the provisioning process.
- The end users are untrustworthy, i.e., they may send and receive arbitrary packets via their E-MAC to/from the Internet. This implies that the management network (accessed using the C-MAC) needs good protection from the user network (accessed using the E-MAC).

The second point is realistic and the main reason for the complex filtering rules within the CMTS. The first point, however, does not hold in today's networks and can be exploited in most cable networks today, as we now explain.

3.3 The Provisioning Process and Its Weaknesses

The DOCSIS standard describes the steps each modem has to fulfill to register itself on the cable network. If a step in the process fails, the modem has to repeat the step until it succeeds. Since it is up to the CNP where he will place the digital channels that are in use for realizing the Internet service, the first step of the cable modem is a large frequency scan to search for the downstream channel. After that, the cable modem will get the parameters for the upstream by searching for a special packet in the downstream channel called *upstream channel descriptor*. Since the cable modem now has knowledge about both the down- and the upstream channels, the modem now has to synchronize itself to the channels in a step called ranging. In this step, the cable modem adjusts the timing, power, and frequency to balance the network delay. Subsequently, the cable modem establishes IP connectivity. Therefore it sends a Dynamic Host Configuration Protocol (DHCP) discover packet with option code 60 (vendor class identifier) for the C-MAC and a normal DHCP (without option 60) for the E-MAC interface. The cable modem listens for a DHCP offer packet that contains the needed data. Option 60 of DHCP allows the interfaces to tell the DHCP server which kind of network devices they are [4] by attaching a message to the DHCP request. This is used to ensure that every interface gets an IP address in a separate IP address pool. The DHCP offer for the cable modem contains an IP address that is assigned to the cable modem management part (C-MAC). Usually, this IP address comes from the cable modem address pool 10.61.0.0/16 [7]. The DHCP offer also includes the IP address of a TFTP server in the management network and the name of a configuration file residing on the TFTP server. The content of a typical DHCP offer is shown in Fig. 3. The client IP address (assigned to the C-MAC) is the entry in the field "your client IP address", in this case it is 10.61.151.101. The name of the TFTP server that hosts the configuration file for this cable modem is contained in the field "Next server IP address" and is 172.30.*.* in this case. The name of the configuration file itself is encoded in the "Boot file name" field.

As next step, the cable modem will download the configuration file from the TFTP server. After this step the modem has to send the file to the CMTS in a step that is called *transferring the operational parameters*. This has to be done to authenticate the modem as modem of the CNP. If the modem is in the database of the CNP, the CMTS sends a message to the modem that it has passed registration. Now the modem is fully provisioned and able to act as bridge between the cable network and the LAN of the customer. A sample configuration file is shown in Fig. 4 where critical data (like passwords) has been sanitized. The SIP username ("0305338890") and the SIP password ("ABCDE123456") are directly stored in cleartext. This shows that access to configuration files opens complete access to a SIP account.

```

Message type: Boot Reply (2)
Hardware type: Ethernet
Hardware address length: 6
Hops: 1
Transaction ID: 0xe6d50d0c
Seconds elapsed: 8
☒ Bootp flags: 0x0000 (Unicast)
Client IP address: 0.0.0.0 (0.0.0.0)
Your (client) IP address: 10.61.153.101 (10.61.153.101)
Next server IP address: 172.30.█ (172.30.█)
Relay agent IP address: 10.61.128.1 (10.61.128.1)
Client MAC address: █
Server host name not given
Boot file name: █_d_h.cfg
Magic cookie: (OK)
☒ Option: (t=53,l=1) DHCP Message Type = DHCP Offer

```

Fig. 3. Excerpt from the DHCP offer sent by the DHCP server within the configuration process (captured and visualized using Wireshark, identifiable data is obfuscated)

```

PCMA
comm1
public
comm1
@mtaprov
comm1
comm1
comm2
5g21wm7sdl
comm2
@mtaprov
comm2
comm2
*x[0-9]*.[t#]||1|[025]||[*#][2-9]x|[*#]1xx
&My_Small_Company
My_Small_Company.com
SIP.Registrar.IP
SIP.Registrar.IP
ABCDE123456
0305338890
0305338890
0305338890

```

Fig. 4. Example DOCSIS configuration file extracted with strings

4 Attacking the Provisioning Process

There are many reasons for an attacker to steal configuration files, but only the SIP credentials are profitable for an attacker. Therefore, we aim to steal the configuration files that contain SIP credentials from the provisioning servers. With the SIP credentials, the attacker can make free telephone calls, hijack and spoof phone calls, and do anything that the real owner of the SIP account can do.

4.1 The Attack

The steps of an attacker are as follows:

1. In a first step the attacker fakes the MAC address of his own computer using standard tools [13] for anonymity and maybe also to avoid access control restrictions in the CMTS (for example, if only certain vendors are allowed).

Then he attaches his own computer to the HFC network to the LAN port of the cable modem.

2. Now the attacker accesses the provisioning network. The attacker has to spoof the device information of his computer in a way that the DHCP server that is responsible for the cable modems “thinks” that the attacker is in fact a part of the cable modem. This can be realized by configuring the DHCP client to send the right form of option 60 with his DHCP request. One correct form of this option is for example to set the vendor class identifier to the value "pktc1.0" [7]. This tells the provisioning server that the device sending the request is a cable modem that is only capable of the DOCIS 1.0 provisioning process. After issuing this request using appropriate tools [4], the attacker will receive the corresponding DHCP response including an IP address inside the HFC access network. The DHCP response also includes the IP address of the TFTP server hosting the configuration files, and a route through the CMTS to the provisioning network.
3. Depending on the MAC address of the attacker, the name of the configuration file within the DHCP response will be some default configuration file name pointing to a location on the TFTP server that will not contain anything important. If by chance, the MAC address is known to the TFTP server, it will point to a configuration file containing the credentials of the corresponding user. Since the name scheme of the configuration files is up to the CNP, it could require some reverse engineering to find out the mapping of MAC addresses to configuration file names. One approach, for example, is to sniff the provisioning process by using the real MAC address of the attacker. In practice we observed different naming schemes. One scheme concatenated the MAC address with the suffix `d.u.cfg`. Using this knowledge about the configuration file naming scheme and with access to the TFTP server, the attacker can easily enumerate configuration file names by using tools like *TFTP brute* to brute force MAC.
4. After the successful download of one or more configuration files, the attacker can extract the SIP credentials from the configuration files easily, e.g., by using the Linux command `strings`.

Now it is possible to abuse that SIP account.

4.2 Why Is the Attack Possible?

There are two specific points in the attack that are critical for its success. First, the names of configuration files are deterministic and can be enumerated. This problem can be easily fixed by giving configuration files a new random name when they are distributed. In a sense, this also ties a specific configuration file to a specific MAC address. But even if this is done, the second problem remains, namely that end users can impersonate any other device by spoofing their MAC address. In summary, the attack exploits exactly those weaknesses that result from an inadequate threat model. The main point is the inadequate assumption that endpoints are physically protected and therefore impersonation is impossible. That this attack exists is surprising since it is well known how easy it is

to attach computers instead of cable modems to the network endpoints. Similarly, it is usually possible to reprogram cable modems by installing manipulated firmware. As shown in the attack, this opens the path for MAC address spoofing and impersonation of specifically weak end devices.

5 Countermeasures

There are many obvious technical countermeasures to the attack described above. The basis for good countermeasures, however, is a more realistic threat model.

5.1 Adapted Threat Model

The adequate threat model for cable networks takes into consideration that the CMTS is the last (physically) controllable point in the cable network and that the cable modem is an untrustworthy device that could do anything and act different from the way that is expected. In particular, it is not possible to bind service usage to particular physical endpoints as it can be done in classical telephone networks.

5.2 Technical Countermeasures

Binding service usage to particular users is known as the problem of authorization. A prerequisite for authorization is authentication. The DOCSIS standard specifies a set of features that can enforce authentication within the provisioning process (BPI+ and DOCSIS shared secret [1, 10]). Most of these features are, however, not used or only optional and not enforced in practice. These features at least prevent other known attacks such as uncapping of cable modems, clear text network traffic on the downstream side and normal protocol attacks that are not specific to cable networks by not only enforcing authentication but also encryption. The fact that they are almost never turned on points again to an insufficient threat model but possibly also can be explained by the cumbersome effort to set up a public key infrastructure and equip modems with certificates. While enforcing authentication, both BPI+ and DOCSIS Shared Secret do not prevent the mass downloading of configuration files. One approach to prevent this is to hide the location of the TFTP server from the customer. By using *cable dynamic-secret mode* (DMIC), the CMTS will change the DHCP offer in the way that it will point to the CMTS and not to the TFTP server. The CMTS then will download the configuration file from the real TFTP server and insert a HMAC based on a onetime password and a cryptographic hash of the file. The modem then only is allowed to register itself if it contains the correct HMAC. The important aspect is that the TFTP server is hidden from the customer. This makes it impossible for the attacker to brute-force configuration files and extract the SIP credentials [12].

A similar result can be achieved by using randomization in the following way: Whenever a cable modem sends a DHCP request, the name of the configuration

file is chosen in a random fashion and uploaded to the TFTP server. The space from which the filename is chosen must be large enough so that it is hard to guess real filenames (e.g., a 64 bit number). But this method has not been standardized in DOCSIS yet. The DOCSIS standard also does not specify any procedures for the secure provisioning of cable modems that use VoIP since there is no well-tested and accepted procedure for this yet. It is not clear whether there will be a clean and working solution in the near future. Therefore, at least a fast detection approach for such attacks has to be implemented.

6 Conclusion

As shown in this paper and in the literature, the security problems of cable networks are not only the result of a weak DOCSIS standard, it is the way how DOCSIS is “implemented” within an organization that causes the attack vectors to persist. Many security features such as BPI+ are often not enabled. Since such actions are consistent with the current threat model that assumes the user endpoint as trustworthy, we have argued that the risks in modern cable networks are mainly due to inadequate threat models at the side of the CNP. While this point is true in general for all systems that allow attacks, cable networks offer a particularly interesting case study because they show how adding certain features (Internet access) to a secure system (TV cable networks) results in an insecure system if the threat model does not evolve too.

Acknowledgments. We thank Andreas Dewald and Martin Mink for helpful comments on a previous version of this paper.

References

1. Security on Data-over-Cable Systems: DOCSIS, BPI+ and Beyondm (November 2000), http://www.3com.com/other/pdfs/infra/corpinfo/en_US/50301102.pdf
2. Hacking the Cable Modem: What Cable Companies Don't Want You to Know. No Starch Press, San Francisco (2006)
3. PacketCable 2.0: Security Technical Report. Technical Report PKT-TR-SEC-V05-080425, Cable Television Laboratories, Inc. (April 2008)
4. Alexander, S., Droms, R.: DHCP Options and BOOTP Vendor Extensions. RFC 2132 (1997)
5. Bundesnetzagentur. Tätigkeitsbericht 2008/2009 Telekommunikation (December 2009)
6. Cable Television Laboratories, Inc., Cable Modem to Customer Premise Equipment Interface. Technical Report CM-SP-CMCI-C01-081104 (November 2008)
7. Cable Television Laboratories Research Consortium. DOCSIS Website (2010), <http://www.cablelabs.com/cablemodem/>
8. Endler, D., Collier, M.: Hacking Exposed VoIP: Voice Over IP Security Secrets & Solutions, 1st edn. McGraw-Hill, Inc., New York (2007)
9. Johns, M.S.: DOCSIS Cable Device MIB Cable Device Management Information Base for DOCSIS compliant Cable Modems and Cable Modem Termination Systems. RFC 2669 (1999)

10. Latini, P.S.: Avoiding Piracy in DOCSIS Networks. Canitec Conference and Exhibition (April 2010)
11. McKelvey, J.: Combating security risks on the cable IP network. Cisco Systems, Inc., Whitepaper (June 2002)
12. Millet, M.: Theft of Service — Inevitable? CableFAX: The Magazine (December 2005)
13. Pahwa, P., Tiwari, G., Chhabra, R.: Spoofing Media Access Control (MAC) and its Counter Measures. International Journal of Advanced Engineering & Application (January 2010)
14. Raftus, D., Cardona, E.: Radio Frequency (RF) Interface Management Information Base for Data over Cable Service Interface Specifications (DOCSIS) 2.0 Compliant RF Interfaces. RFC 4546 (2006)
15. Shah, N., Kouvatso, D., Martin, J., Moser, S.: A Tutorial on DOCSIS: Protocol and Performance Models. In: International Working Conference on Performance Modeling and Evaluation of Heterogeneous Networks (July 2005)
16. Swiderski, F., Snyder, W.: Threat modeling. Microsoft Press, Redmond (2004)