

NAC@ACK

Hacking the Cisco NAC Framework

Blackhat USA 2007 Whitepaper



Michael Thumann, mthumann@ernw.de

&

Dror-John Roecher, droecher@ernw.de

Version 1.0 – March 9th 2007

Marketing Buzz – What is Cisco NAC?	2
A look under the hood – NAC technical primer	3
Terminology	4
More technical Details.....	6
A simple sample Lab-setup for NAC-L3-IP	10
Flawed by Design – Analysis of Cisco NAC	16
Approaching NAC@ACK	18
Thoughts on Mitigation	27
List of Figures	28
References & Further Reading	29
Acknowledgements	29

Marketing Buzz – What is Cisco NAC?

Looking back at the evolution of security one can easily make out technologies that become en vogue for some time and then transform from “hot-stuff” into standard-security technologies. Just think of firewall, S2S-VPNs, RAS-VPNs, IDS, IPS or SSL-VPN solutions, to name a few. Moving away from perimeter-defense and taking security into the depth of networks, security has been concentrating on endpoints for quite some time now. Virusscanner, Anti-Spyware/Anti-Phishing and Desktop-Firewalls are some popular examples of Endpoint-Security-Technologies nowadays commonly applied to Enterprise-Desktops.

Besides Endpoint Security there has been a big fuzz on “compliance solutions”, though everyone seems to understand compliance in a slightly different way. Compliance to BCP, or to external regulations (just think about SOX or Basel-II) or even compliance with internal policies (is the antivirus up-to-date?) – as you can see there are many different flavours of compliance.

Now try to imagine what a vendor could achieve by combining “Endpoint Security” and “Compliance” into a single concept or solution – as a vendor you would have two cows to milk: The “IT” department and the “Compliance/Audit” department. Additionally you could try to make this a really big source of additional revenue, because a solution like that could be as pervasive as you wish – creating the need to upgrade/replace all network-equipment, installing new management-servers, selling client-software et cetera. Of course you would need to be a big vendor to be able to push your proprietary solution upon the market, like e.g. Cisco Systems. As you may have noticed, Cisco has been quite heavy on the marketing of the “Self Defending Network” [we will refrain from commenting on that term] and a core component of that is “Network Admission Control”, or simply NAC.

Cisco’s NAC solution has the goal to improve security by allowing only compliant clients admission to the network – the compliance is defined by a set of self-defined policies (e.g. antivirus-signatures are current and antivirus is active). Clients failing to comply with the policies can be granted restricted access or denied access altogether. Sounds pretty cool – it actually is and it works quite well, but the design is flawed which enables an attacker to access a NAC-protected networks without being compliant under certain circumstances. Under what circumstances and how to get that access is what will be detailed in this paper. But first a thorough understanding of the technical operations and components involved is needed – so let’s stop the marketing buzz now and dig into the technology.

A look under the hood – NAC technical primer

The Cisco NAC solution encompasses four distinct components playing together to check a client for compliance and to enforce access-restrictions on the client based on the results of the checks. Naturally a new solution like NAC will come along with a new acronym-soup and terminology. Figure 1 depicts the relationship between the components and the protocols used for communication.

1. **Client-side software:** For NAC to check the state of a client, a piece of software is required to run on the client and this software, in its most basic form, is the “Cisco Trust Agent” (CTA). The CTA communicates via EAP and supplies so called “credentials” – basically attribute-values-pairs conveying information about the state of the client to a backend policy-server. The CTA itself can provide some basic credentials (OS-Type & Version, etc – see Table 3 for a complete list) but if further credentials are supposed to be checked, additional NAC-aware software is needed on the client to plug-in to the CTA and provide the CTA with the needed information.
2. **The “Enforcing Device”/Network Access Device (NAD):** The Network Access Device is the enforcing point on the network and it has two tasks. It relays the EAP-messages from the CTA as RADIUS-messages to the Cisco Secure ACS and back to the client and it enforces the policy determined by the Cisco Secure ACS. Routers, Switches, Firewalls, VPN-Concentrators and WLAN-Access-Points can act as NAC-enabled devices as long as they run up-to-date software and are “made by Cisco”.
3. **Cisco Secure ACS:** The core component of Cisco’s NAC is the “Cisco Secure Access Control Server” (ACS). The ACS is a RADIUS server which serves multiple functions.
 - It accepts client-credentials via the NAD.
 - It checks the credentials against a local policy
 - It may check 3rd-party policy servers, if additional NAC-enabled applications are running on the client.
 - From the credentials provided and policies configured, access-restrictions are derived and communicated to the NAD and to the client.
4. **Optional Backend-Servers:** If additional security-software, e.g. antivirus, is to be included in NAC, then not only need it be installed on the client, but a policy-server for that application may be needed, too. The ACS needs to be made aware of this external policy server and it queries it using the “Host Credential Authorization Protocol” (HCAP) – basically “XML over HTTPS”. This third-party policy server is not looked at within this paper or the accompanying talk.




Figure 1 Components of a Cisco NAC Network

Terminology

Before delving into more technical details the used terminology shall be introduced. The information provided by the client to the ACS is called “posture credentials”. For each NAC-enabled application the ACS derives an “application posture token” (APT). This posture token is one out of “healthy”, “checkup”, “transition”, “quarantine”, “infected” or “unknown”. These tokens have a well-defined meaning:

- “Healthy”: fully compliant with the admission policy for the specified application.
- “Checkup”: partial but sufficient compliance with the admission policy, no need to restrict access, a warning to the user may be issued.
- “Transition”: either during boot-time, when not all necessary services have been started or during an audit-process for clientless hosts¹, temporary access-restrictions may be applied.
- “Quarantine”: insufficient compliance with the admission policy, network access is usually restricted to a quarantine/remediation segment.
- “Infected”: active infection detected, usually most restrictive network access even up to complete isolation.
- “Unknown”: a token can not be determined or no CTA installed on client. This may lead to partial access (guest-vlan & internet-access for example).

From all APTs a system posture token (SPT) is deduced, this SPT corresponds to the APT which will grant the least access to the client. Based on the SPT the access-restrictions are chosen and sent back to the NAD to be enforced. All APTs and the SPTs are also communicated back to the client, where software may use this information to inform or warn the user or to change an applications’ behavior.

There are 3 different “flavours” of NAC relating to different access-vectors of the client bringing along different methods for restricting access to the network.

1. Layer3-IP: With NAC Layer3-IP the access-restrictions are implemented as IP-ACLs on a Layer-3 device (e.g. a Router or a VPN-Concentrator/Firewall). The communication takes place using PEAP over EAP over UDP (EoU).
2. Layer2-IP: Layer2-IP NAC enforces access-restrictions as IP-ACLs on a VLAN-interface of a switch. The communication takes place using PEAP over EAP over UDP (EoU).
3. Layer2-802.1x: Layer2-802.1x uses 802.1x port control to restrict network access – obviously the device enforcing these restrictions is a switch. In this case the EAP-FAST is used in conjunction with 802.1x. This is the only NAC flavour where the client is:
 - a. authenticated before being allowed on the network
 - b. restricted from communicating with its local subnet

This is the most secure (and the most elaborate) NAC setup possible, because a working 802.1x infrastructure is assumed to be in place, but because of its “Layer2-nature” it is not feasible on all setups.

Table 1 shows some of the features of the three different variants of NAC and serves as a quick comparison.

¹ Clientless hosts are hosts which have no CTA installed or where the CTA is not operating properly. These fall into three categories: (1) Legitimate systems where the CTA can not be installed, e.g. network-printers, (2) external hosts, e.g. guests or consultants and (3) legitimate systems where the CTA is not operating properly. These systems can be either statically whitelisted (this is done with systems from the first category) or can be “audited” by a vulnerability-scanner, or they may be treated with a “default”-restriction like guest-access.

Feature	NAC-L2-802.1x	NAC-L2-IP	NAC-L3-IP
Trigger	Data Link / Switchport	DHCP / ARP	Routed Packet
Machine ID	Yes	No	No
User ID	Yes	No	No
Posture	Yes	Yes	Yes
VLAN Assignment	Yes	No	No
URL Redirection	No	Yes	Yes
Downloadable ACLs	Cat65k only	Yes	Yes
Posture Status Queries	No	Yes	Yes
802.1x Posture Change	Yes	No	No
Transport EoU/PEAP	No	Yes	Yes
Transport 802.1x / EAP-FAST	Yes	No	No

Table 1 Feature Comparison

Each of these different flavours is suitable for different situations and they definitely differ in the security level they provide and how the secure the mechanisms themselves are (as e.g. is obvious from the lack of authentication in two flavours).




Figure 2 Sample communication-flows in Cisco NAC

More technical Details

Now that the technology has been roughly sketched it is time for some technical details which will become interesting or important later on.

Host Credential Information

The credentials collected on the host and communicated to the ACS can be roughly divided into two different classes. (1) Information gathered through NAC-enabled applications and (2) arbitrary information gathered by the CTA through a scripting interface.

The credentials are transmitted as attribute-value-pairs and are of different value-types:

- Octet-Array, Integer32, Unsigned32, String (UTF-8), IPv4 Address, IPv6 Address, Time (4 Octets), Version (4-x-2 octet-sets)

The attribute-part of the AV-pair is defined by a set of parameters:

- Numerical Vendor ID / Vendor ID Name. The IANA SMI private enterprise ID assignments are used.² For example Cisco has ID 9, McAfee 1230, Trendmicro 6101.
- Numerical identifier of the application type (16 Bit value). Some values are reserved:

Application-Type	Application-Type Name	Usage
1	PA	Posture Agent
2	Host / OS	Host information
3	AV	Anti Virus
4	FW	Firewall
5	HIPS	Host IPS
6	Audit	Audit
32768 – 65536		Reserved for “local use” (custom plug-ins or scripts)

Table 2 Application Identifiers

- Numerical Attribute-ID, Attribute Name & Attribute-Type

The following table lists the credentials available through the CTA (version 2.0) without installation of any other NAC-enabled application. This is meant to give an impression on the type of information available for policy-enforcement.

Application-Type	Attribute Number	Attribute Name	Value-Type
Posture Agent	3	Agent-Name (PA-Name)	String
	4	Agent-Version	Version
	5	OS-Type	String
	6	OS-Version	Version
	7	User-Notification	String
	8	OS-Kernel	String
	9	OS-Kernel-Version	Version

² For SMI private enterprise assignments check: <http://www.iana.org/assignments/enterprise-numbers>

Host	11	Machine-Posture-State	1 – Booting, 2 – Running, 3 – Logged in.
	6	Service Packs	String
	7	Hot Fixes	String
	8	Host-FQDN	String

Table 3 CTA included credentials

So a posture-credential can be identified in shorthand as “Vendor : Application-Type : Attribute”, e.g. “9:1:8” would translate to Vendor: Cisco (9), Application: Posture Agent (1), Attribute: OS-Kernel and with reference to Table 3 it would expect it to be a string (e.g. “Linux-2.6.4-8-i386-custom”).

As mentioned the CTA has two distinct ways of enabling credentials to be collected and transmitted. The first one is “regular posture plug-ins” and the second one is via a “scripting interface”. Let’s have a short look on how they differ and what they do.

Plug-ins (on Windows) are realized through dll-files (an so-files on Linux) which are located in “%CommonProgramFiles%\PostureAgent\Plugins” (again on Windows) and are configured through corresponding ini-files where the available credentials for the plug-in are described. When installing CTA 2.0 three plug-ins are installed by default:

- Host Posture Plug-in
- CTA Plug-in
- Scripting Plug-in

Third-party applications will install additional plug-ins.

The scripting interface (realized through ctasi.exe and ctascriptPP.dll on the client) can be used to add third-party-applications and self-written assessments to NAC which are not realized as posture plug-ins. Again an ini-file is used to describe the available credentials and an external program is referenced which collects the credentials - the format of the inf-file is the same as for plug-ins. The program is assumed to write the collected credentials into a plaintext “posture-data” file which is read by the CTA (actually the program is expected to call ctasi.exe) and parsed using the information from the inf-file. The inf-files for plug-ins and scripts take the same format.

A sample ini-file (tmabpp.inf) from Trendmicros' OfficeScan will serve to illustrate this process:

```
[main]
dll=tmabpp.dll
PluginName=tmabpp.dll
VendorID=6101
VendorIDName=TrendMicro, Inc
AppList=av

[av]

AppType=3
AppName=Antivirus
AttributeList=attr1,attr2,attr3,attr4,attr5,attr6,attr7,attr8,attr9,attr10,attr11,attr12,attr13,attr14
attr1= 1, Unsigned32, Application-Posture-Token
attr2=2, Unsigned32, System-Posture-Token
attr3=3, String, Software-Name
attr4=4, Unsigned32, Software-ID
attr5=5, Version, Software-Version
attr6=6, Version, Scan-Engine-Version
attr7=7, Version, Dat-Version
attr8=8, Time, Dat-Date
attr9=9, Unsigned32, Protection-Enabled
attr10=10, String, Action

attr11=32768, String, OSCE-Srv-Hostname
attr12=32769, OctetArray, Client-GUID
attr13=32770, Ipv4Address, Client-IP
attr14=32771, OctetArray, Client-MACddd
```

The name of the plug-in. In case of a script this would be ctascriptPP.dll and the vendor-id would be "Cisco" for scripts.

Official Credentials

Private Credentials from the Vendor

And a sample posture-data file which is created as output of a script and passed to the CTA:

```
[attr#0]
vendor-id=9
vendor-name=Cisco
application-id=61440
application-name=ERNW-NAC-Script-1
attribute-id=32768
attribute-name=Script-Name
attribute-profile=in
attribute-type=string
attribute-value=Script "posture_file_01"
[attr#1]
vendor-id=9
vendor-name=Cisco
application-id=61440
application-name=ERNW-NAC-Script-1
attribute-id=32769
attribute-name=Script-Run-Counter
attribute-profile=in
attribute-type=unsigned integer
attribute-value=17
[attr#2]
vendor-id=9
vendor-name=Cisco
application-id=61440
application-name=ERNW-NAC-Script-1
attribute-id=32770
attribute-name=Host-IP-Address
attribute-profile=in
attribute-type=ipaddr
attribute-value=196.168.100.200
[attr#3]
vendor-id=9
vendor-name=Cisco
application-id=61440
application-name=ERNW-NAC-Script-1
attribute-id=32772
attribute-name=Script-Date
attribute-profile=in
attribute-type=date
attribute-value=1077771601
[attr#4]
vendor-id=9
vendor-name=Cisco
application-id=61440
application-name=ERNW-NAC-Script-1
attribute-id=32773
attribute-name=Script-Version
attribute-profile=in
attribute-type=version
attribute-value=0.1.0.0
```

A list containing detailed descriptions of all attributes and their actual values. These must match the corresponding inf-file of the script.

Vendor is always Cisco for scripts and Application-ID is from "local use" range.

This data-file is created by the script which then invokes ctasi.exe.

Ctasi.exe parses the script and relays the credentials via the ctascriptPP.dll plug-in to the regular CTA communication process.

The architecture of the CTA (as shown in Figure 3) clarifies the relationship between the different components. The communication layer accepts posture-credentials (EAP-TLV) from the broker and passes this information to lower variant-specific transport-oriented modules (EoU/802.1x). The posture-plugin API and scripting interface supply information from plug-ins or from scripts to the broker/communication-layer.

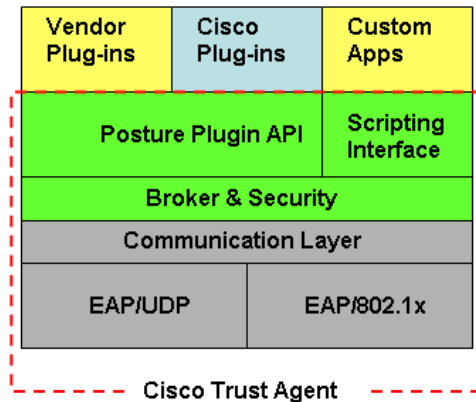


Figure 3 Architectural Overview of the CTA

A simple sample Lab-setup for NAC-L3-IP

To get a feeling for the operations of NAC a brief description of a NAC-L3 setup will be described. This is the description of the LAB actually used for the research presented in this paper and for developing the attacks shown in our presentation. Figure 4 depicts a generalized view of the environment, the ACS is running in the “NAC protected Core”. Users can access the Office-Environment through RAS-VPN but need to pass “NAC protection” in order to access the “holy core”. The choice of “NAC-flavour” for this scenario is easy – obviously Layer2-802.1x will not work as the clients are on the internet and are not connected to a local switch in the office. The same holds true for L2-IP and additionally the NAD never sees ARP-requests or DHCP-requests from the client. So the only possible flavour for this setup is NAC-L3-IP.

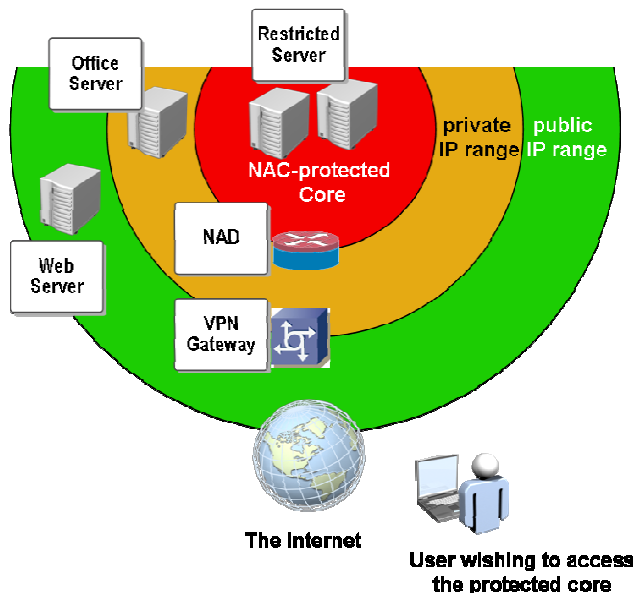


Figure 4 Sample NAC-L3 Setup

The setup was tested with a multitude of different policies in order to understand the packet-format and the behaviour of the application. These policies were kept as simple as possible. Please see Figure 5 for an example policy used in the test-scenarios..




Figure 5 A sample policy which will result in an APT for PA of „healthy“ if the client is running Windows XP and has a hostname of „vm-xp-nocta“. The APT will be „quarantine“ if at least one of the tests fail.

The client can be notified of the results of the assessments. The notification is dependent on the SPT and is configured on the ACS (see Figure 6).

Group Mappings for - Network Admission Control - IIR DEMO

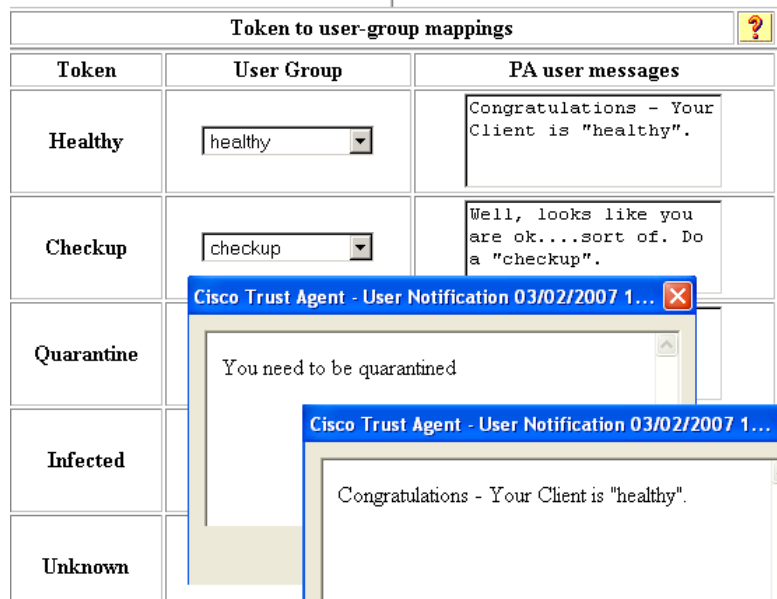
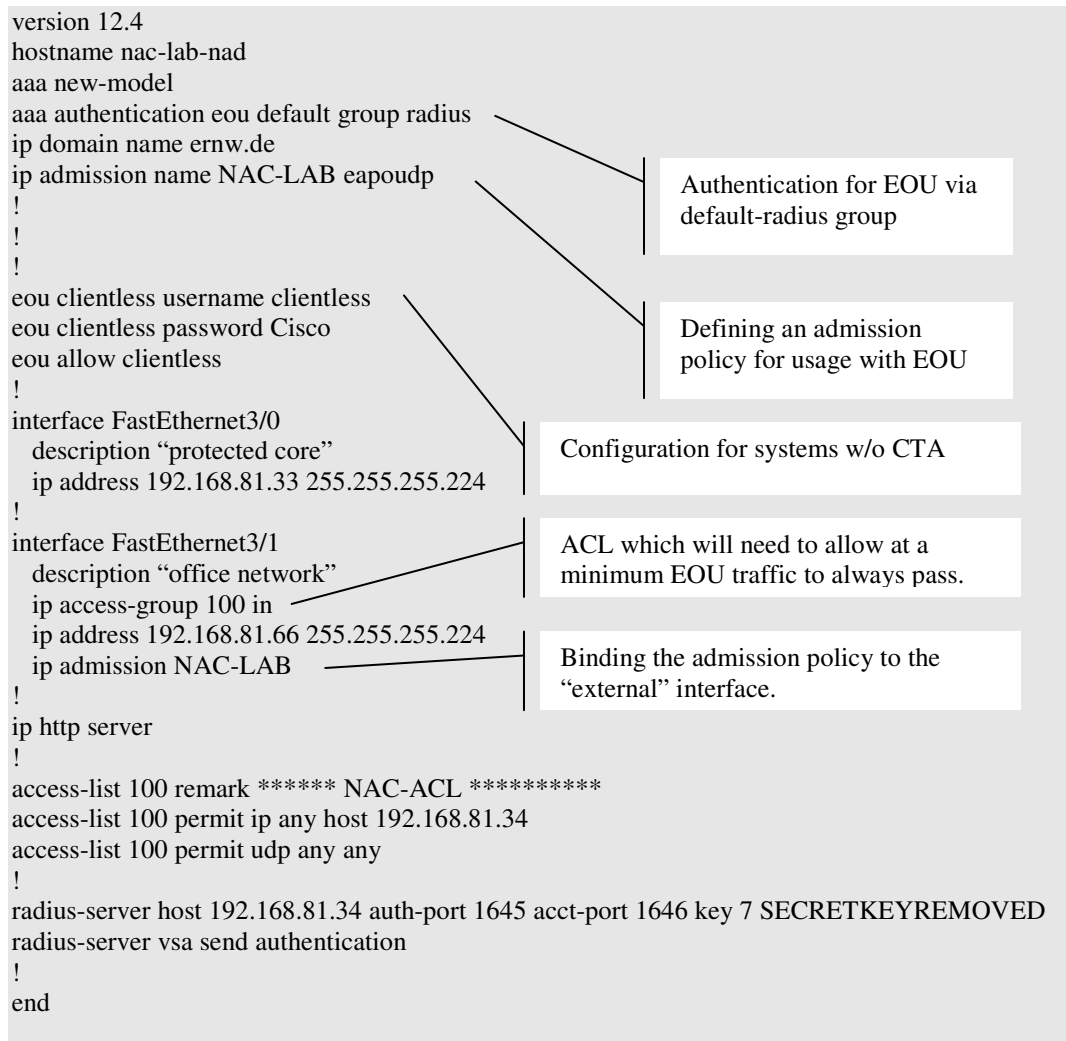


Figure 6 User notification setup in ACS and CTA client notification popup.

The configuration of the NAD (in our case a Cisco 3640) is quite simple. An “admission policy” is created and bound to the external interface. On the same interface an ACL is in place which allows



at least EoU traffic (UDP port 21862 to the ACS) and if access-restrictions for a client are configured as ACLs on the ACS, the ACL on the interface will be dynamically expanded with these downloadable ACLs. Additionally an optional configuration for clientless hosts can be configured. In our case clientless hosts are allowed and are authorized with username “clientless” and password “cisco” – this user (clientless) is configured on the ACS with a downloadable ACL for restricting access of these clients on the network.

When traffic passes the NAD from the office-network towards the core, the traffic is intercepted by the admission policy and the client is asked to provide the credentials. The NAD keeps track of all NAC-assessments and the current SPT of all clients. As can be seen in Figure 7 an agentless host (192.168.67.34) is treated according to the above configuration (authenticated as “clientless” with a password against the ACS user-database and hosts with CTA installed are validated via EAP. In the depicted case the CTA-host (192.168.67.24) has successfully passed all policy-checks and is assigned a SPT of “healthy”.

```

cisco - HyperTerminal
Datei Bearbeiten Ansicht Agrufen Übertragung ?
e=FastEthernet3/1
Mar  2 13:26:15.243: %EQU-6-AUTHTYPE: IP=192.168.67.24| AuthType=EAP
nad#
nad#
nad#show eou all
-----
Address          Interface      AuthType      Posture-Token  Age(min)
-----
192.168.67.34   FastEthernet3/1  CLIENTLESS    unknown        0
192.168.67.24   FastEthernet3/1  EAP           healthy        0
nad#
    
```

Figure 7 Output of „show eou all” on the NAD

Communication Flow for NAC-Layer3-IP

NAC-Layer3-IP uses PEAP over EAP over UDP as a transport-mechanism for the credentials. PEAP establishes a TLS-tunnel between the client and the ACS. Within this tunnel client can be optionally authenticated - which is not the case in Cisco NAC. Figure 8 shows a generic communication flow for NAC-Layer3-IP and Figure 11 shows the communication flow including the information that is carried within each step.

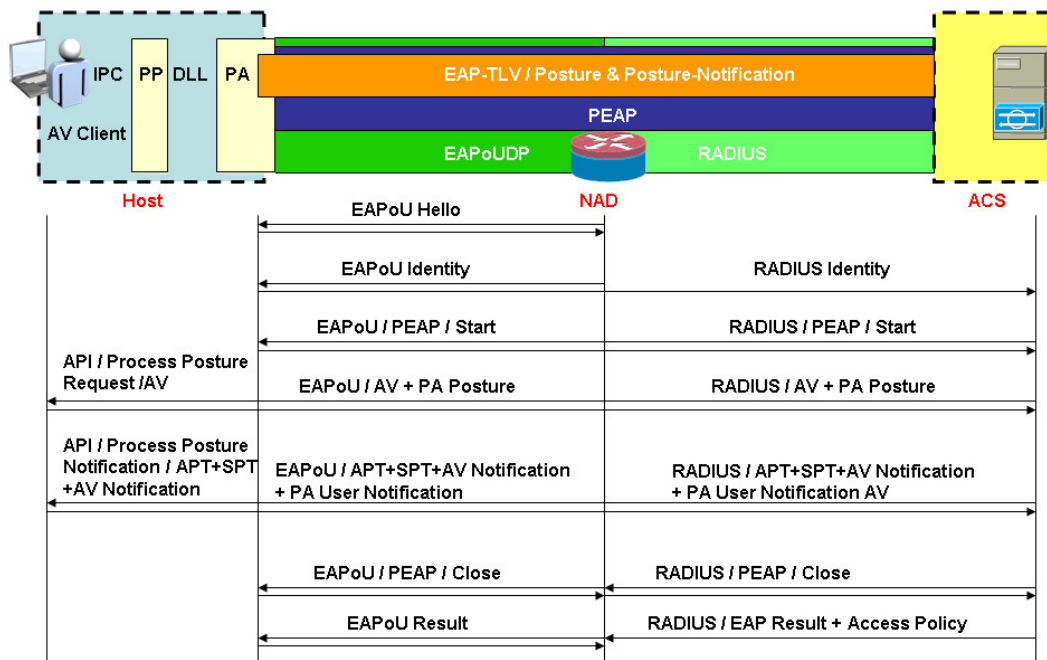


Figure 8 Detailed Communication Flow in NAC Layer-3 IP³

The Cisco NAC for Layer3-IP and Layer2-IP is built on PEAPv1 and EAP-TLV. The PEAP communication uses two distinguished phases. Phase 1 establishes a secure tunnel using EAP-TLS and authenticates the server using a server-certificate. The second phase includes an optional client-authentication (not implemented in NAC) and the exchange of arbitrary information – in the case of NAC this arbitrary information is composed of posture data and posture notifications using EAP-TLV as a means of encoding that information. Figure 10 shows the frame format for Cisco PEAP

³ Citing from the “Implementing Network Admission Control Phase One Configuration and Deployment” guide (available on <http://www.cisco.com>): „ Note that the router acts as a pass-through device at this point; it does not proxy any part of the PEAP session but merely reencapsulates the PEAP packets from UDP to RADIUS.”

(PEAPv1) and Figure 12 & Figure 13 show the format for EAP-TLV and EAP-TLV Vendor specific frames respectively.

In Cisco NAC PEAP is carried as the payload of EAP over UDP (EAPoU), leading to a generic frame-format shown in Figure 9.




Figure 9 Generic Frame Format

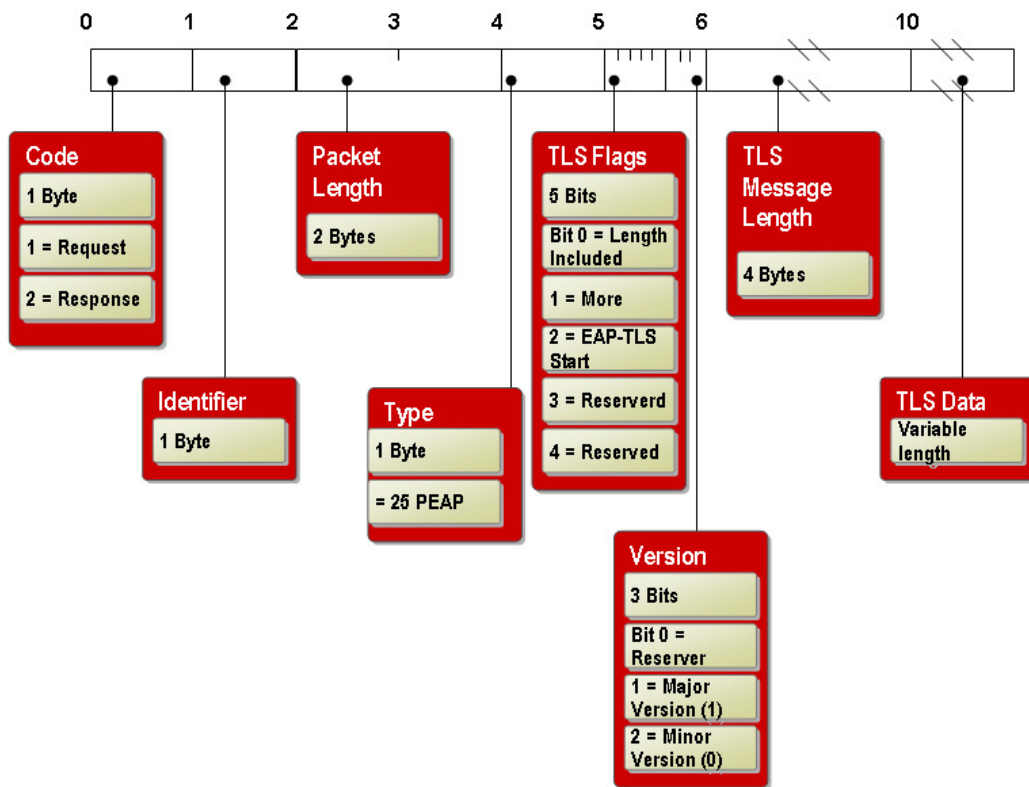


Figure 10 Cisco PEAP Frame Format

PEAP uses a multiple packets to establish the tunnel and ensure secure communications. A rough overview was given in Figure 8 – a more detailed communication flow for PEAPv1 tunnel establishment and decommissioning is given in Figure 11.



Figure 11 Detailed view of PEAPv1 packet flow




Figure 12 EAP-TLV Frame Format

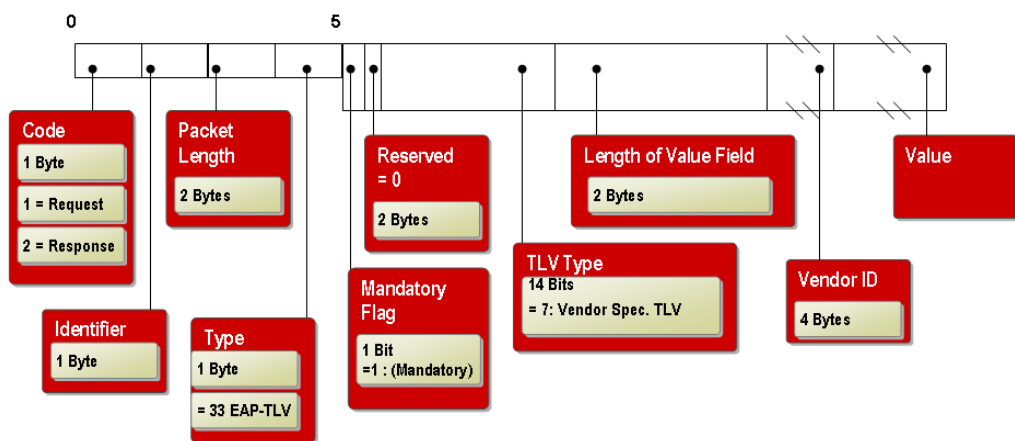


Figure 13 EAP-TLV Vendor Packet Format

Flawed by Design – Analysis of Cisco NAC

Looking at the three different variants of Cisco NAC it is seen, that they differ significantly in the security they provide for the network on at least two details:

	NAC-Layer 3 IP	NAC Layer 2 IP	NAC Layer 2 802.1x
Client Authentication	No intrinsic Client Authentication. In VPN scenarios there is a “VPN Authentication” which might be considered a “mitigating control”.	No intrinsic Client Authentication – and no means of “adding” such on top.	Client Authentication based on 802.1x
Restriction of access on local subnet.	It is not possible to restrict access to the local subnet via NAC.	It is not possible to restrict access to the local subnet via NAC.	Access to local subnet can be denied through “port shutdown” via NAC.

Table 4 Authentication & restrictions for Cisco NAC flavours.

A problem common to all NAC solutions is that the state of the Client is used to derive the access-level (one could also name this “trust level”) on the network. Why is that a problem? Because this information (state of client) has to be provided by the client – whom one doesn’t trust first place, or else one wouldn’t be thinking about using NAC to distinguish between “trustworthy” and “not trustworthy” clients. So the untrusted client is asked to provide credentials in order to judge its trustworthiness. The problem at hand reminds us of the Epimenides paradox – a logical problem named after the Cretan philosopher Epimenides of Kronos. It can be stated like that:

Epimenides was a Cretan who made one statement: "All Cretans are liars."

The next problem (which is a Cisco NAC design specific) is that in two out of three possible variants, there is no intrinsic client authentication whatsoever. We are talking about authorization without authentication. If you wear a police-shirt, NAC treats you as a police-officer without ever asking you for any proof – you look like a police officer, so you have to be one.

The design of Cisco NAC is either flawed or NAC was not developed as a security-mechanism but as something else. We can’t help but cite from the Cisco NAC website (<http://www.cisco.com/go/nac>). There is a section “NAC business benefits” and the first listed benefit is “dramatically improves security”. In my opinion it should read “NAC might dramatically improve your security”. So we postulate that the design is flawed. Feel free to prove us wrong.

The next obvious question is “is it possible to exploit this design-flaw?” and if the answer to that question is “yes” [which, it is] – than “under which circumstances and how can we exploit it and what could be achieved by that?” The “restrict access to local subnet” we quickly put away as “not interesting” – there is nothing to hack in that respect, it is simply a matter of “plug in your client” and be connected. The fact that the access-level is derived from client-provided information is the far more interesting aspect in NAC. By the way this is not specific to the Cisco solution but is common to most NAC approaches (Consentry being a notable exception to this). And the “inadequate authentication” or as we like to call it “authorization without authentication” is yet a different aspect and it stroked us as a very promising vulnerability. Thus the idea of “posture spoofing” was born.

We define “posture spoofing” as an attack where a legitimate or illegitimate client spoofs “NAC posture credentials” in order get access unrestricted network access. The remainder of the paper is mostly dedicated to “which steps are necessary to realize a ‘posture spoofing’ and under which circumstances are these attacks feasible”.

Attacker's Definition

Different attack-vectors for “posture spoofing” are available for different types of attackers. Therefore we will first define these attackers.

Insider: An insider is a legitimate user of a NAC-protected network. The client has a working installation of the CTA and valid user/machine-credentials for the network. Additionally the inside attacker has the certificate of the ACS installed in its certificate store and if 802.1x is being used with client-certificates, this attacker has a valid certificate. The insider simply wants to bypass restrictions placed on his machine (e.g. no “leet tools” allowed and NAC checks list of installed programs).

Outsider: An outsider is not a legitimate user of the NAC-protected network and wants to get unrestricted access to the network. The outsider has no valid user/machine-credentials and no working CTA installation.

Attacker's Location

The logical location of the attacker with respect to the network and the enforcing device does not play a major role as the EAPoU or EAPo802.1x traffic always needs to be enabled.

Without adumbration of results presented later, the attack-vectors for “posture spoofing” with respect to “Attacker type” and “NAC flavour” are presented in the following table:

	Insider	Outsider
NAC-L2-802.1x	DLL/Plug-In replacement Scripting Interface CTA replacement ⁴	None as to our current knowledge.
NAC-L2-IP	DLL/Plug-In replacement Scripting Interface CTA replacement	CTA replacement Scripting Interface
NACL-L3-IP	DLL/Plug-In replacement Scripting Interface CTA replacement	CTA replacement Scripting Interface

Table 5 Attack Vectors for posture spoofing

The crux with “posture spoofing” is that it has to be assumed that the attacker has no knowledge of the posture credentials which are required for “SPT healthy” in the actual setup. But, reading the official Cisco documentation it can be deduced, that the ACS actually requests credentials, thus promoting assisting an attacker with a classical “information disclosure” flaw which in itself seems harmless at first sight.

To sum it up, the analysis unveiled three design flaws:

1. Authorization without authentication.
2. Epimenides Paradox.
3. Information Disclosure through the ACS by providing the list of needed credentials.

⁴ In the case of 802.1x a self-coded alternative client would need to also need to be a 802.1x supplicant with EAP-FAST support. This seems a bit of an “overkill” for that attack vector but it should nevertheless be mentioned for completeness sake.

Approaching NAC@ACK

After deciding to head at a “Posture Spoofing” attack a systematic approach was developed. The approach encompassed an attack-tree (which is depicted in Figure 14) and a set of techniques. From the techniques and the attack-tree “hacking steps” were derived and every step was first analyzed for its’ feasibility.

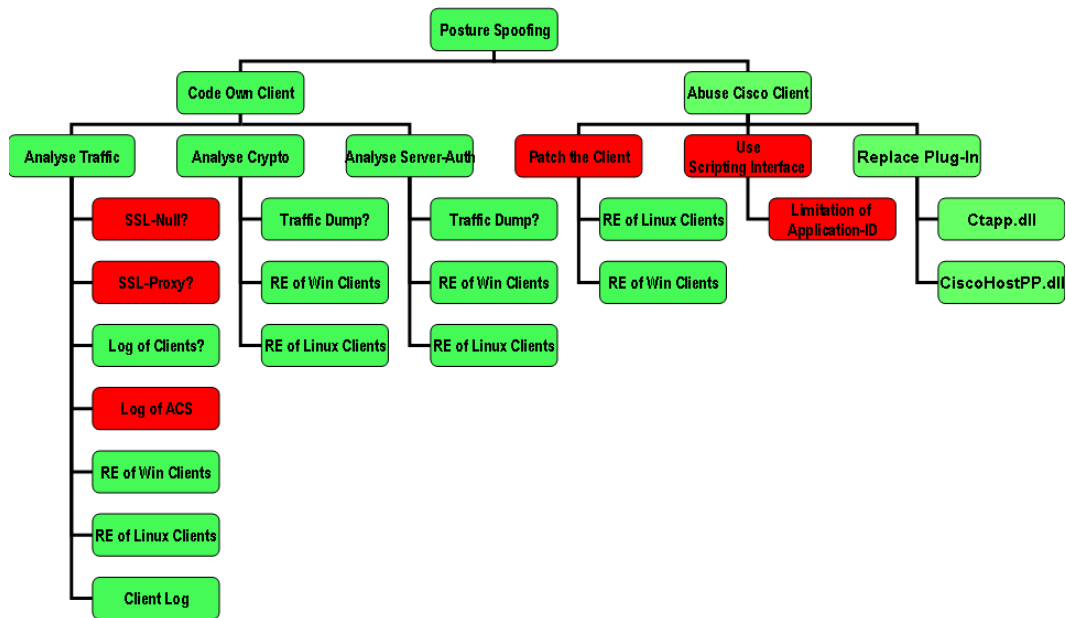


Figure 14 The Attack Tree. Green steps show “feasible” approaches whereas red steps have been shown to be not feasible or have no major benefit. The attack tree shown are not the first version but the latest version.

Brief primer on ‘hacking techniques’ for complex systems

The Cisco NAC framework is a proprietary closed-source complex system and hacking it requires a well defined set of tools and techniques. The process of hacking such complex systems is not comparable to the finding of Buffer Overflows the average reader might be more or less familiar with. FX did a very good job of describing a “complex system hack” with his “Analysing Complex Systems” presentation regarding Blackberry at the Blackhat-Europe 2006. A brief discussion of the techniques used shall be presented here before going into the results achieved. Some of these techniques yield results by themselves, others are either only useful in combination or serve as a staging process for a different technique.

Reverse Engineering: Reverse Engineering aims at uncovering the constructional elements of a product. In our case the product which needs reverse engineering is the binary version of CTA itself. The “tool-of choice” is the well-known disassembler IDAPro (thanks guys!)⁵.

Packet Sniffing: As we are examining a networked system it strikes at obvious, that “packet sniffing” is one technique which should be taken into account. On a second look one can see, that packet sniffing is of limited use only, as the communication is encrypted and the packet-format is undisclosed. Nevertheless, Wireshark/Ethereal is your friend.

Packet Diffing: Understanding encrypted packets in an undisclosed format is quite challenging. One way to tackle the problem is by ‘diffing’ packets. By “diffing” we mean extracting common and differing parts of two packets. These differences hint at packet-bytes used for common tasks as e.g. “packet-code” [request/response] or “communication identifier” [comparable to a session-id].

⁵ We will not go into the theory of RE – there are other people who are a lot better at that than we are and it is definitely out-of-scope for the paper. We assume you have a basic understanding of what RE is and how it works.

Multiple tools exist for diffing, even vim will serve as a simple binary diffing tool when combined with xxd.

Debugging / API-Monitoring: While reverse engineering or disassemblies provide the interested researcher with a snapshot of the program, it does not show the live process and the contents passed to the functions through the stack. That is where API-Monitoring and debugging have their say. Through attaching a debugger or api-monitor to the running process, it is possible to actually see the contents of the stack while the program is running. This is especially useful if the source code of the analysed program is not available. While we tested some different api-monitoring tool, we achieved the best results with “autodebug for windows”⁶.

Built-in capabilities: Many software-products come with built in logging and/or debugging capabilities of various degrees. If these are available, they usually prove to be a very good source of information and should be used extensively.

RTFM: Again this should be self-explanatory – reading all available documentation should be one of the first steps in any analysis. The original documentation may vary in quality but it usually contains many useful hints. And it is often worth reading more than once as some details in the documentation may be evaluated differently on the second read with more knowledge from other sources at hand.

Applying the techniques

The first big “want to have” was a clear-text version of the encrypted traffic. As the traffic is secured with SSL/TLS the first approaches were to either force SSL-Null or use an SSL-Proxy. The SSL-Null was discarded quite quickly as it would need tampering with the client and/or the ACS and at the beginning of the research we wanted to avoid that. SSL-Proxy was not feasible either, as the NAC implements “SSL over UDP” and we have no knowledge of a Proxy which is able to do that. In addition the Cisco CTA has the certificate of the ACS installed and it validates the certificate in the process of the PEAP negotiation, so tampering with the certificate (e.g. through an SSL-Proxy) would lead to a malfunction of the SSL negotiation and therefore NAC would not work at all. While reading the (quite comprehensive) Cisco documentation of the CTA we discovered the “ctalogd.exe” – the part of CTA which is responsible for logging and which is configured through the “ctalogd.ini”-file. Logging writes its entries into a plaintext file on the client and every message is assigned a log level. The ctalogd.ini allows for the filtering of the messages according to the log level and a log level of 15 means “log everything”. A test with log level 15 revealed that even decrypted packet-payloads are dumped – this was validated with the text from a user notification.

As to our knowledge the exact format of the packets is not documented anywhere (yet) [In February 2007 Cisco announced to publish the CTA as OpenSource but recalled that announcement later on⁷]– the best documentation for PEAPv1 is a Cisco Press book on wireless security⁸ and there is an ietf-draft for “EAP over UDP” (expired August 2002)⁹. First hand source are (of course) the packet-dumps themselves and so we developed a test suite with multiple different policies for packet-capturing. Figure 15 shows a sample packet capture.

⁶ Autodebug is available through <http://www.autodebug.com/>

⁷ <http://slashdot.org/article.pl?sid=07/02/08/2327220&from=rss>

⁸ Krishna Sankar et. al., „Cisco Wireless Lan Security” ISBN-13 978-1-58705-154-8, Nov. 2004, Cisco Press.

⁹ <http://folk.uio.no/paalee/publications/draft-engelstad-pana-eap-over-udp-00.txt>

The ctalogd.ini file with log level 15 enabled:

```
[main]
EnableLog=1
LogDir=c:\temp
[LogLevel]
PADaemon=15
NetTrans=15
PAPlugin=15
[...]
```

Logging enabled

Log level "15" – log everything possible.

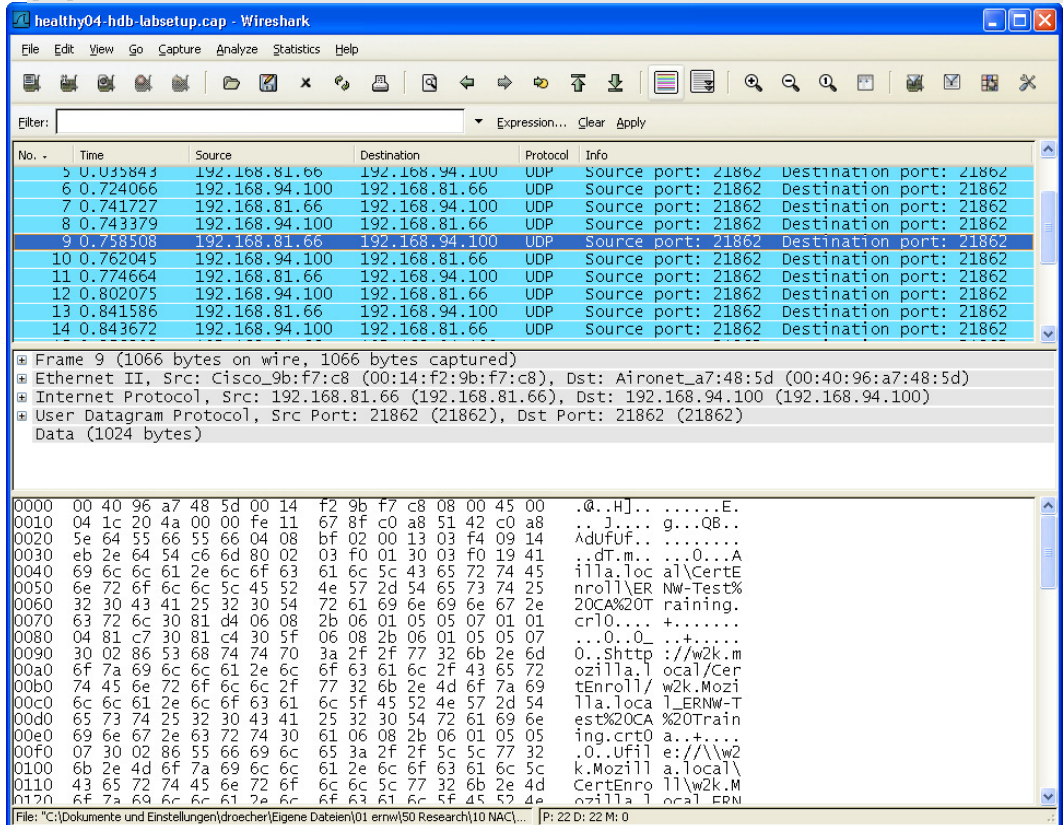


Figure 15 A sample packet capture. The payload shows the ACS certificate as part of the TLS establishment.

For further packet-analysis binary differences of the packet-payloads were taken. These “diffs” revealed bytes within the packets which are unchanged between “request” and “response” or between packets from the same source (e.g. all “client-packets” or all “ACS-packets”). A sample “diff” is shown in Figure 16.




Figure 16 Sample packet-diffing

```
Excerpt from a CTA logfile:  
65 16:23:13.343 04/26/2006 Sev=PktDump/13 CTAVSTLV/0x64300016  
Request message dump:  
000700D100000009800200C900000009000100100001000800000000000000000000010000200  
08000000000000009000100A9000700A14865727A6C696368656E20476C7565636B77756E73  
6368202D20496872205043206B6F6E6E7465206572666F6C6772656963682061757468656E746  
966697A696572742077657264656E20756E6420656E74737072696368742064657220536563757  
26974792050696F6C6963792E2049687265204E65747A7765726B7A7567616E67207769726420  
6E696368742065696E676573636872E46E6B7421800300020001  
66 16:23:13.359 04/26/2006 Sev=Info/4 PAPIPlugin/0x63200001  
Application Posture Result = Healthy  
67 16:23:13.359 04/26/2006 Sev=PktDump/13 PAPIPlugin/0x63200001  
Response message dump: 800300020001  
68 16:23:13.359 04/26/2006 Sev=Debug/7 CTA  
EapHandlePacket exit  
[...snipped...]  
70 16:23:13.359 04/26/2006 Sev=Info/4 PAPIPlugin/0x63200002  
System Posture Result = Healthy  
71 16:23:13.359 04/26/2006 Sev=Warning/2 PAPIPlugin/0xA3200012  
CTAPP received UserMsg Notification: Content = Herzlichen Glueckwunsch - Ihr PC konnte  
erfolgreich authentifiziert werden und entspricht der Security Piolicy. Ihre Netzwerkzugang wird  
nicht eingeschränkt!
```

User Notification:
"Herzlichen ..."
Convert to Hex:
%48%65%72%7a%6c%69
%63%68%65%6e%20

Reverse Engineering & Function Hooking Techniques

In order to get a better understanding of the inner workings of the CTA reversing techniques were applied. While the Windows version of CTA is compiled using some optimization methods, the Linux version is not optimized and so the focus of the reversing was first set to the Linux version. Figure 17 shows that the crypto within the CTA 2.0 for Linux is built around OpenSSL, meaning that it is possible to use existing crypto-libraries in coding an “alternative” NAC-client.

Besides revealing which (standard) crypto is being used, the reversing of the Linux client also sheds some light on the process-flow within the CTA itself and the core functions within the CTA. OpenSSL 0.9.7g is being used, which has known vulnerabilities¹⁰ - attacking the client through OpenSSL was not pursued up to now but may prove to be a successful attack vector, but we would not consider this hacking of the NAC framework.

Address	Length	Type	String
"" .rdata:1...	0000000E	C	FIPS routines
"" .rdata:1...	0000000E	C	OCSP routines
"" .rdata:1...	00000010	C	engine routines
"" .rdata:1...	00000015	C	DSD support routines
"" .rdata:1...	00000018	C	random number generator
"" .rdata:1...	00000010	C	PKCS12 routines
"" .rdata:1...	00000011	C	X509 V3 routines
"" .rdata:1...	0000000F	C	PKCS7 routines
"" .rdata:1...	0000000D	C	BIO routines
"" .rdata:1...	0000000D	C	SSL routines
"" .rdata:1...	00000018	C	elliptic curve routines
"" .rdata:1...	0000001A	C	common libcrypto routines
"" .rdata:1...	0000001C	C	configuration file routines
"" .rdata:1...	00000017	C	asn1 encoding routines
"" .rdata:1...	0000001A	C	x509 certificate routines
"" .rdata:1...	0000000D	C	dsa routines
"" .rdata:1...	0000000D	C	PEM routines
"" .rdata:1...	0000001B	C	object identifier routines
"" .rdata:1...	00000017	C	memory buffer routines
"" .rdata:1...	0000001A	C	digital envelope routines
"" .rdata:1...	00000018	C	Diffie-Hellman routines
"" .rdata:1...	0000000D	C	rsa routines
"" .rdata:1...	00000010	C	bignum routines
"" .rdata:1...	0000000F	C	system library
"" .rdata:1...	00000010	C	unknown library
"" .rdata:1...	00000013	C	.\crypto\ver\ver.c
"" .rdata:1...	00000014	C	int_err_get (err.c)
"" .rdata:1...	00000017	C	int_thread_get (err.c)
"" .rdata:1...	00000008	C	unknown
"" .rdata:1...	00000015	C	error:%08X:%s:%s:%s
"" .rdata:1...	0000000C	C	reason(%u)
"" .rdata:1...	0000000A	C	func(%u)
"" .rdata:1...	00000009	C	lib(%u)
"" .rdata:1...	0000001C	C	.\crypto\engine\mtb_digest.c
"" .rdata:1...	0000001B	C	.\crypto\engine\eng_init.c
"" .rdata:1...	00000029	C	Stack part of OpenSSL 0.9.7g 11 Apr 2005
"" .rdata:1...	00000017	C	.\crypto\stack\stack.c
"" .rdata:1...	00000019	C	.\crypto\buffer\buffer.c
"" .rdata:1...	00000027	C	RSA part of OpenSSL 0.9.7g 11 Apr 2005
"" .rdata:1...	00000017	C	.\crypto\rsa\rsa_lib.c

Figure 17 CTA for Linux is built on OpenSSL libraries.

The two most promising attack-vectors are “plug-in replacement” and “alternative client”. Even though it would be very easy to implement, the “scripting interface” is deemed less feasible because of limitations placed on the “application id” and “vendor name” for posture-credentials generated through scripts. As mentioned earlier the vendor id for scripts is always “Cisco” and the application id has to be from the “private” range. As to now it is not clear if these restrictions can be circumvented in any way without replacing the scripting-interface-dll.

For the “alternative client” it is essential to understand how the payload is built and transported. The disassembly shows that the core function of the CTA for that purpose is “TlvHandlePacket” (Figure 18). A disassembly of the components of the CTA on its own reveals core functionality and

¹⁰ Chek BID 20246, 15071 and 20249

core functions but it can not trace an application while it is running. The possibility to trace an application at runtime and see calls to functions and contents of stacks and buffers is pivotal to understand what is actually going on within the application and for that reason “function hooking/api-monitoring” was used. The results from the disassembly were used to identify “interesting” functions for hooking into.

Figure 19 shows a call to “EapTlvHandlePacket” with posture data passed to that function. The buffer in contains a list of all installed Microsoft hotfixes concatenated with a pipe, the hostname of the machine (vm-xp1-nocta), the name of the logged-in user (droecher) and the operating system (Windows XP SP2). This result is somewhat surprising, as the policy defined on the ACS during the test-run had no rule regarding installed patches/hotfixes. A possible conclusion is, that the CTA always submits all available posture-credentials to the ACS and that the ACS then only evaluates the credentials which are configured in a policy. This seems to contradict the earlier stated design flaw “information disclosure”.

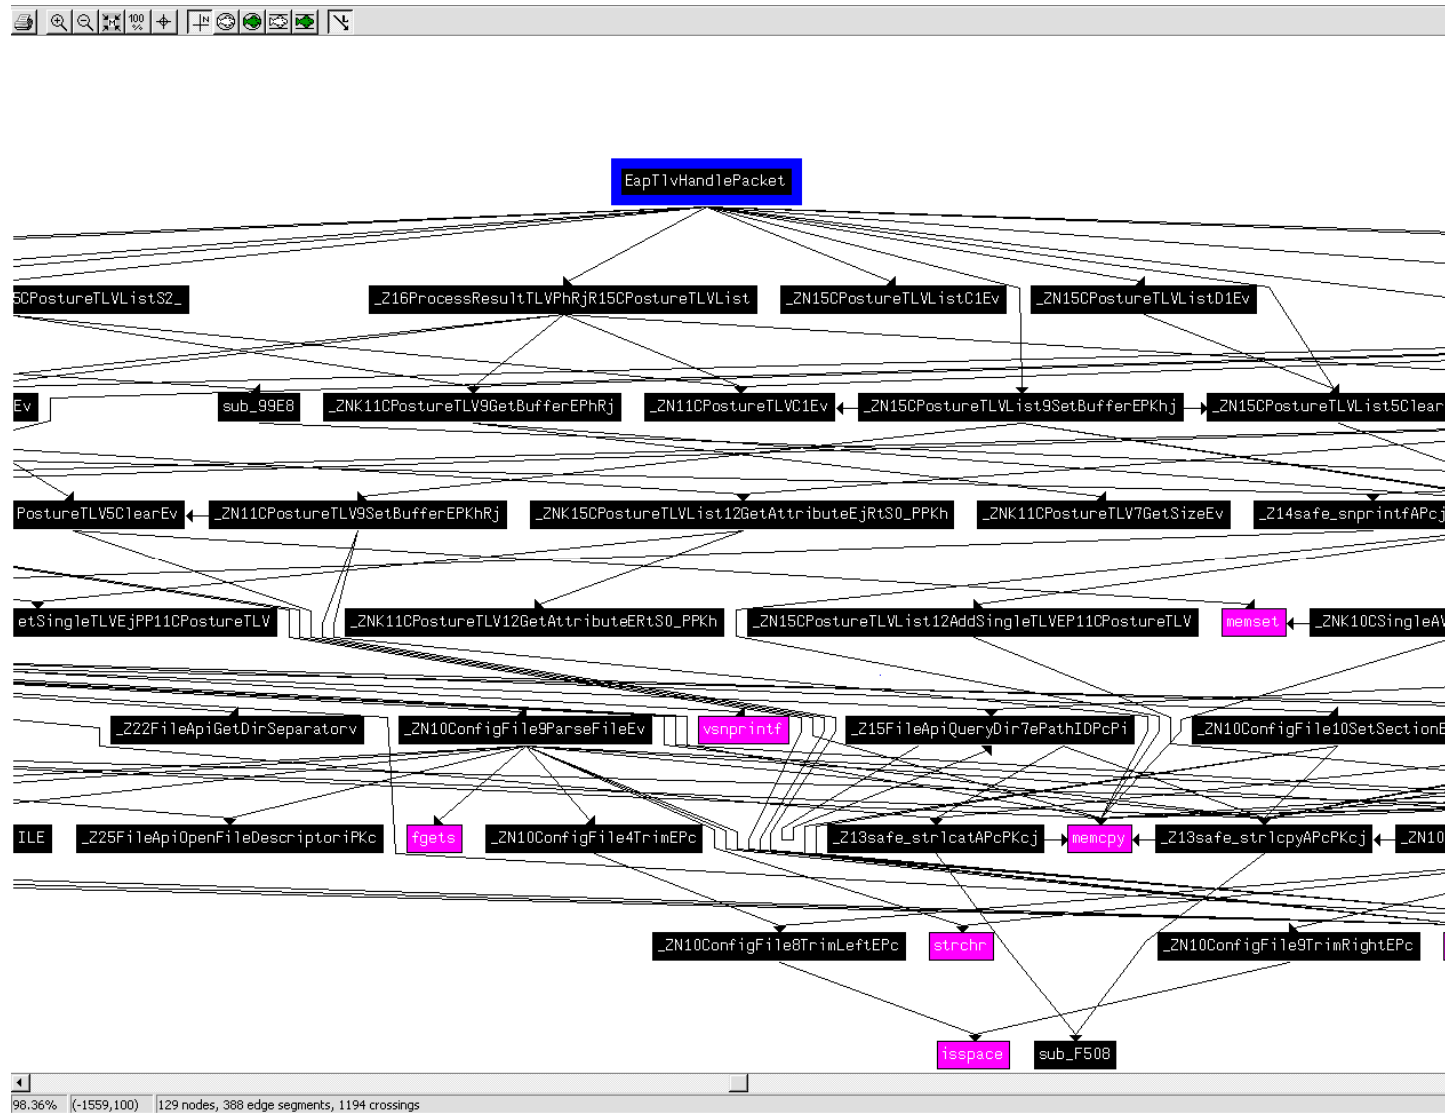


Figure 18 Core function „EapTlvHandlePacket“ from reversing of the CTA

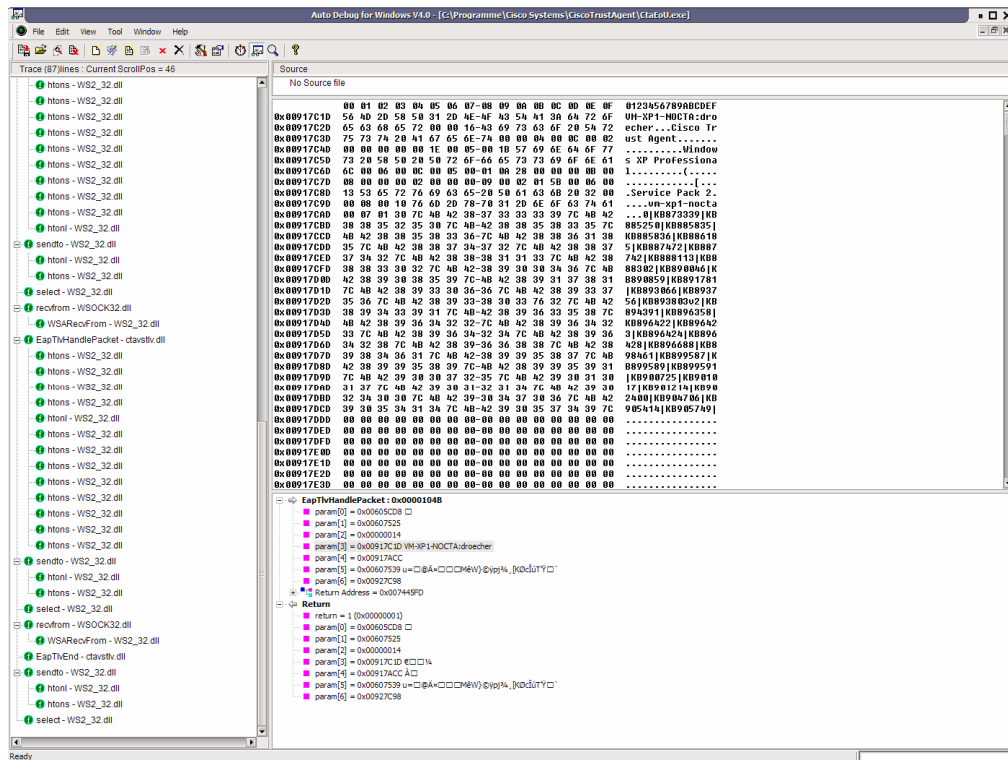


Figure 19 Function hooking reveals parameters passed to functions.

But even more promising than “alternative client” is the approach of “plug-in” replacement. If you recall the structure of the CTA (Figure 3) it becomes clear, that the (probably most commonly used) posture-data for “PA” and “Host” is collected using two corresponding plug-ins. The difference between a plug-in and a self-written script is that the script writes the posture-data to a file and invokes “ctasi.exe” whereas a plug-in hands the posture-data over to the CTA using exported DLL-calls. Additionally the “vendor id” and “application id” restrictions do not apply to plug-ins but only to scripts. So by analyzing different plug-ins for common exported functions it is possible to identify these functions (for an example of the disassembly see Figure 20) and thereby reconstruct the plug-in interface not publicly disclosed by Cisco. Three common functions have been identified in the plug-in feature of the CTA and the disassembly shows their functionality and structure (Figure 21). Disassembly of the public plug-in functions. Figure 21 shows parts of the disassemblies of the core functions of the plug-in interface.

The conclusion

Using the readily available documentation and the information provided above, the development of plug-ins for the CTA is not too difficult for a skilled programmer. And instead of just developing “a plug-in”, replacement-plug-ins for existing plug-ins can be developed (especially for the PA and Host plug-in – simply “replace” them with self-written plug-ins), therefore rendering Cisco NAC more or less useless in all “insider”-scenarios¹¹.

For “outsider-scenarios” the situation is somewhat more difficult. As the outsider is not in the possession of the ACS certificate and has no valid credentials, a simple “replace plug-in dlls” will not do. Even if the outsider would install the CTA and replace the plug-ins he would still be lacking the ACS certificate without which the CTA will fail to establish the PEAP session.

The outsider needs a complete “new” replacement-client. A key feature of this client (besides spoofing arbitrary posture credentials) would be to accept the ACS certificate as valid and use it for PEAP without it being installed in a local certificate-store beforehand. But again this “alternative client approach” is quite definitely possible without too much trouble as the modular CTA allows to “reuse” part of the program, as for example the PEAP-libraries provided by Cisco with the CTA

¹¹ This is a somewhat provocative statement. When talking about “mitigation in the next section, we will discuss how NAC can be useful and how it could be improved.

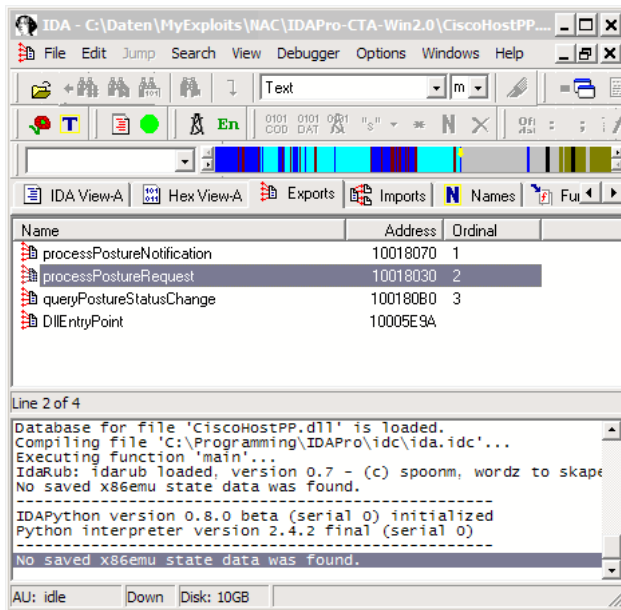


Figure 20 Sample disassembly of CiscoHostPP.dll plug-in showing the exported functions.




Figure 21 Disassembly of the public plug-in functions

Thoughts on Mitigation

If you want to use Cisco NAC, go ahead – but first define what you want to achieve and then think about which NAC setup will fit your goals best. Then calculate the costs, think again, evaluate different approaches or different solutions (Cisco NAC is not the only solution on the market), think again, do a risk-analysis, calculate again, rethink and then you might decide, that you still want to use it. If that is true for you, then go ahead. But back to business – that the current version of Cisco NAC is flawed has been extensively shown. The question is – who can do what to make the solution “secure”? As we are not talking about a “coding error” leading to a buffer overflow but about a fundamental design-flaw, a simple “patch the client” will not work.

There are two parties who can improve the security of the solution: Cisco (of course) and customers using Cisco NAC.

What Cisco can do

Cisco could improve the design of the NAC framework in order to render attacks from insider and outsider useless. Here are some thoughts on what could be done:

Code Signing: Code Signing¹² the plug-ins and running only signed plug-ins would defeat plug-in replacement attacks thereby defeating insider attacks. We can not judge the effort needed to implement code signing but we would heartily welcome seeing signed code in any security product.

Mandatory Authentication: Strong mandatory authentication would stop outsider attacks against the NAC framework. Adding authentication (mandatory or, in a first step, optional) should be possible without too much of a change as PEAP is being used and PEAP has built-in authentication. Preferably we would like to see smart-card based authentication against a central user-directory – but that may be wishing for too much. The reasons for not having authentication in that framework can only be business-related – Cisco knows that implementing NAC is already a major effort and does not want to put additional stress on its clients by making authentication mandatory.

What customers can do

The security of Cisco NAC highly depends on the overall setup/installation of the client and on the NAC flavor chosen.

Strong Authentication: Whenever possible 802.1x-based NAC should be implemented in order to add strong authentication to the authorization process. If 802.1x is not feasible, other means of strong authentication should be implemented. In RAS-VPN scenarios for example, where NAC-Layer3-IP is the only NAC-flavor available, clients should be subjected to strong authentication on the VPN-device itself. The “strong authentication” mitigates threats posed by the “outside attacker”.

Least Privilege: All attack-vectors for “inside attackers” have a common characteristic. They need “tampering” with the CTA installation. In case of “plug-in”-replacement the authentic plug-ins are being replaced by self-written plug-ins. Two possible mitigations against this attack would be to (a) enforce strict access-rights on the plug-in files which prevent replacement (this being a preventive measure) and (b) use some sort of “file altering monitor” to detect changed plug-in files. The second listed mitigation is the less preferable one as it does not prevent but only detect. In case of “alternative client” neither “file monitoring” nor “file access restrictions” are possible mitigations. But an alternative client would need to open a listening UDP port 28162 – which could be restricted by a local firewall to the original CTA.

CSA instead of CTA: In addition to the CTA Cisco also offers a host based IDS in the name of “Cisco Security Agent” which also includes the CTA and has its own CTA plug in. The CSA monitors the integrity of the CTA and will prevent illegitimate changes to the CTA. This will mitigate threats posed by the “inside attacker”.

¹² A good introduction to code signing can be found here:
http://msdn.microsoft.com/workshop/security/authcode/intro_authenticode.asp

List of Figures

Figure 1 Components of a Cisco NAC Network.....	3
Figure 2 Sample communication-flows in Cisco NAC.....	5
Figure 3 Architectural Overview of the CTA.....	10
Figure 4 Sample NAC-L3 Setup.....	10
Figure 5 A sample policy which will result in an APT for PA of „healthy“ if the client is running Windows XP and has a hostname of „vm-xp-nocta“. The APT will be „quarantine“ if at least one of the tests fail.....	11
Figure 6 User notification setup in ACS and CTA client notification popup.	11
Figure 7 Output of „show eou all“ on the NAD	13
Figure 8 Detailed Communication Flow in NAC Layer-3 IP	13
Figure 9 Generic Frame Format.....	14
Figure 10 Cisco PEAP Frame Format.....	14
Figure 11 Detailed view of PEAPv1 packet flow	15
Figure 12 EAP-TLV Frame Format	15
Figure 13 EAP-TLV Vendor Packet Format.....	15
Figure 14 The Attack Tree. Green steps show “feasible” approaches whereas red steps have been shown to be not feasible or have no major benefit. The attack tree shown are not the first version but the latest version.....	18
Figure 15 A sample packet capture. The payload shows the ACS certificate as part of the TLS establishment.	20
Figure 16 Sample packet-diffing	20
Figure 17 CTA for Linux is built on OpenSSL libraries.	22
Figure 18 Core function „EapTlvHandlePacket“ from reversing of the CTA	24
Figure 19 Function hooking reveals parameters passed to functions.	25
Figure 20 Sample disassembly of CiscoHostPP.dll plug-in showing the exported functions.	26
Figure 21 Disassembly of the public plug-in functions.....	26
List of Tables	
Table 1 Feature Comparison.....	5
Table 2 Application Identifiers	6
Table 3 CTA included credentials	7
Table 4 Authentication & restrictions for Cisco NAC flavours.	16
Table 5 Attack Vectors for posture spoofing	17

References & Further Reading

General Cisco sites:

- <http://www.cisco.com/go/nac>
- <http://www.cisco.com/go/sdn>
- <http://www.cisco.com/go/cta>

Cisco CTA & NAC guides

- NAC deployment guide
http://www.cisco.com/application/pdf/en/us/guest/netsol/ns617/c649/cdccont_0900aec80417226.pdf
- CTA 2.1 Administrator Guide
http://www.cisco.com/application/pdf/en/us/guest/products/ps5923/c2001/ccmigration_09186a00807ec2f7.pdf

Books

- “Cisco Network Admission Control, Volume I: NAC Framework Architecture and Design”, ISBN-13: 978-1-58705-241-5;
- “Cisco Network Admission Control, Volume II: NAC Deployment and Troubleshooting”, ISBN-13: 978-1-58705-225-5
- “Cisco Wireless LAN Security”, ISBN-13: 978-1-58705-154-8

Acknowledgements

Without the support of some people all this research work would not have been possible – therefore we want to thank these persons:

- Enno Rey & Roland Fiege, ERNW GmbH, for providing us with lab-equipment and some spare time.
- Simon Rich, ERNW GmbH, for helping us with the coding of the client.
- Barbara Jantzen & Hannah Jantzen, for bearing with my absence.