Integrity Check of Remote Computer Systems
Trusted Network Connect

Marian Jungbauer, Norbert Pohlmann

{jungbauer | pohlmann}@internet-sicherheit.de

Abstract

The economic dependence on the fast and inexpensive exchange of information that has arisen as a result of globalisation is leading to ever increasing levels of networking. The internet provides a communication infrastructure that is available worldwide. However, it does not offer any opportunity for trustworthy communication, as it is not possible to analyse the computer systems found in the network with respect to their system integrity and trustworthiness. The same also applies to intranets. Visitors and field workers who use their computer systems - for example notebooks - both outside and within the company network represent a threat to the company with these computer systems. Through the use of the computer systems outside the company network they are also working outside the protective measures and control area of the company's IT. Solution approaches such as Trusted Network Connect (TNC) provide methods for determining the integrity of end points which serve as a basis for trustworthy communication. The configurations of the end points can be measured on both the software and hardware level. Through the reconciliation of defined safety rules it is possible to realise policy-controlled access control.

1 Introduction

Only a few decades ago data and documents were exchanged both within and between firms either personally or by post. For the transport of sensitive data it was necessary to use trustworthy communication pathways, such as the company's own postal service.

Subsequently initial forms of electronic data transmission, e.g. transmission by fax or the first networks offered the opportunity of transmitting documents quickly from A to B. However, it was not possible to guarantee the confidential transmission of sensitive data.

Today, in the age of globalisation, the distances, quantity and importance of the information that is to be exchanged are increasing. Furthermore, there is increasing cost and time pressure on all operational processes. Classical, locally restricted networks (intranets) throughout the world are therefore becoming ever more intensively integrated into large, locally unrestricted company networks. Home and field workers require rapid and secure access to data in the company's network from any location. Transactions with other organisations, in particular in the B2B field, are increasingly made by electronic means.

Former static network infrastructures with clear system limits are therefore giving way more and more to heterogeneous and dynamic networks [Hart05].

1.1 The Problem

The flexibility and inexpensive use of the Internet is diminished by a lack of security which currently only permits restricted use by public authorities and industry in fields where security is critical.

Fieldworkers use their computer systems in many environments with various security requirements. Home workers use their PCs for private purposes. Employees take their notebooks home with them. These computer systems which are removed either temporarily or permanently from the protective measures of the company networks are exposed to considerably greater hazards. If they are compromised by malware, their reintegration into the company network (either directly or via the Internet) allows them to circumvent the company's own safety mechanisms and therefore jeopardise the company itself through the company network.

Today Virtual Private Networks (VPNs) are mainly used for connecting external computer systems to company networks. These offer both encoding and user authentication, but no analysis of system statuses and therefore of the trustworthiness of the accessing computer systems.

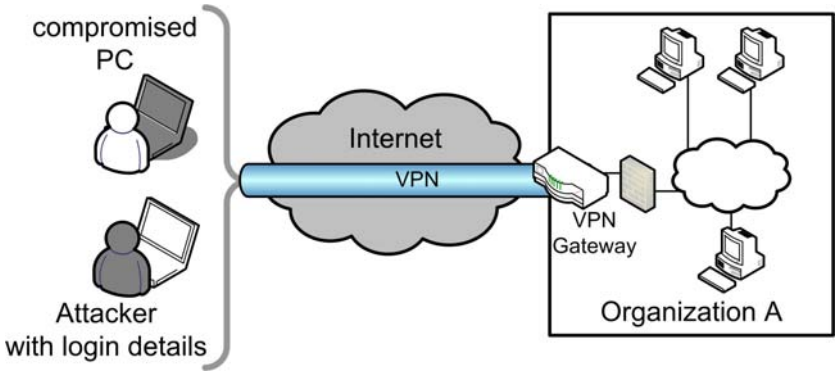

Fig. 1 Possible dangers by VPN

Figure 1 shows two possible dangers to a network from computer systems accessing via VPN:

- Due to the absence of an integrity analysis it is not possible to protect the computer network with its services against a computer system compromised by malware.
- It cannot be ascertained whether the computer system which is being communicated with is also the computer system that it claims to be. If an attacker obtains the access data of a VPN, he can use them for unauthorised access.

1.2 Requirements Placed on Today's Computer Networks

Current and future networks must be flexible and open in terms of their expansion. At the same time they should make trusted communication possible.

Even today there are possibilities in existence for expanding networks in a flexible manner and equipping them with security services - for example via VPNs. However, there is a lack of security mechanisms which guarantee the trustworthiness of the computer systems employed by the user.

The aim of new security systems must be the verifiability of the trustworthiness of the participating computer systems and the creation of secure communication.

2 Trusted Network Communications

One speaks of trusted communication when the trustworthiness of all communication partners involved can be ascertained. Here it is important that the trustworthiness is considered separately for each communication direction. It is possible to conduct a confidential discussion with an individual whom one trusts without the other person having the same level of trust in oneself. In addition to the individuals concerned, the environment in which the communication takes place must also be trusted. If the communication takes place locally, the location must be classified as trustworthy, i.e. free of listening devices, for example. If the communication takes place over a certain distance, the transmission route - for example a postal service with all its employees and premises - must be trustworthy. Additionally, the trustworthiness depends on the requirements of the communication partner, who defines the requirements he places on confidential communication with the help of a policy. This policy might define which messenger service he considers to be trustworthy and how the delivery must be packaged in order to be able to detect compromised consignments.

With respect to communication via computer networks, all individuals involved in the communication must be reliably authenticated on the one hand, and on the other hand the integrity of all computer systems involved in the communication must be guaranteed.

Techniques already in existence, such as VPN, offer reliable verification of the identity of the individuals involved in the communication and the secure transmission of data via networks. However, there is a lack of appropriate security mechanisms which allow integrity verification of the endpoints used for the communication (computer systems).

The trustworthiness of a computer system depends on the overall status of all hardware and software components with their configurations. Measured integrity does not represent a standardised status of a computer system. Whether the trustworthiness of a computer system exists or not depends rather on the security guidelines (policies) of the communication partners. For example, an operator of a network can consider the trustworthiness of a computer system to be evidenced by the use of an up-to-date operating system, while another operator also demands up-to-date user software, for example the latest browser.

Only when all of the system components, i.e. hardware and software, defined in the policy of the network operator are in a desired and uncompromised state is the trustworthiness of the system considered to be secure. The problem here is that today any compromising - particularly of the operating systems - cannot be measured directly, but only indirectly through the existence of further software. This occurs, for example, through the use of an up-to-date virus scanner and an optimally adjusted personal firewall. Only if these programmes are installed, correctly configured and kept right up-to-date in terms of data can the probability of their being compromised be considered low.

This is the very point at which the new concepts for the establishment of integrity-tested network connections come in. Under the generic term "Network Access Control" (NAC) these concepts make it possible to verify the configuration of the end points when building up a network connection. Which configurations of the hardware and software of a computer system are permitted in a network is defined by the network operator by means of policies. For example, these policies stipulate the presence of a virus scanner with up-to-date virus signature, an installed and well-configured personal firewall and the latest patch level for the operating system and applications. Only if the policies are fulfilled is an enquiring end device granted access to the network and its services.

With respect to these new security concepts, one also speaks of a change to the protection strategy of networks and their services. As a result of the verification of the computer systems before network access, there is a changeover from the reaction to a threat to its prevention. While attempts are made today through the use of Intrusion Detection Systems (IDS) to detect compromised computer systems on the basis of abnormal measurement readings in network traffic - i.e. a computer system infected by malware must first behave "incorrectly" (reaction) before it can be discovered - preventive security concepts stop computer systems with a faulty or undesirable system configuration - which may therefore be compromised - from entering the network in the first place and using the services present there.

3 Trusted Network Connect

With the Trusted Network Connect (TNC) specification the Trusted Computing Group is developing its own NAC approach. The development is taking place through the Trusted Network Connect Subgroup [Trus06] with over 75 firms represented and is currently available (May 2007) in version 1.2 [Tru+06].

The aim is the development of an open, producer-independent specification for verifying endpoint integrity. This verification is fundamental to ascertaining the trustworthiness of a computer system.

TNC makes use of current security technologies for network access ("802.1x" and "VPN"), for the transport of messages ("EAP", "TLS" and "HTTPS") and for authentication ("Radius" and "Diameter") in order to enable easy integration into existing network infrastructures.

3.1 Phases

All functions provided by TNC can be classified into three phases:

The **Assessment Phase** comprises all actions from the attempt at establishing a connection with a TNC network to the decision on its integrity. The measured values of a computer system are compared by a server in the network on the basis of policies. This comparison makes it possible to take a decision on the integrity of the computer system.

If upon non-fulfilment of the policies the computer system is categorised as not having the required integrity, it moves to the **Isolation Phase**, in which the accessing computer system is isolated in a protected network area. Therefore any computer systems that are compromised with malware do not gain access to the network and the services offered there.

The **Remediation Phase** offers the isolated computer systems the opportunity to restore their integrity, for example through the installation of missing security software, and - after renewed verification - access to the network.

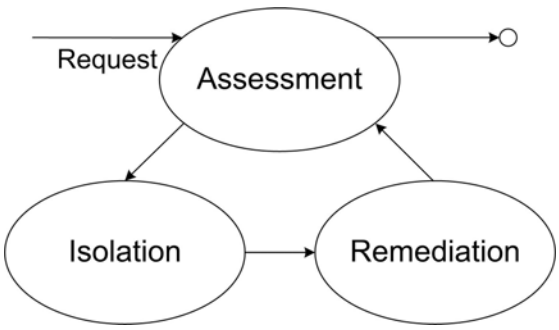

Fig. 2: Connection of the TNC phases

3.2 Structure

A fundamental distinction is made between three elements in the TNC specification.

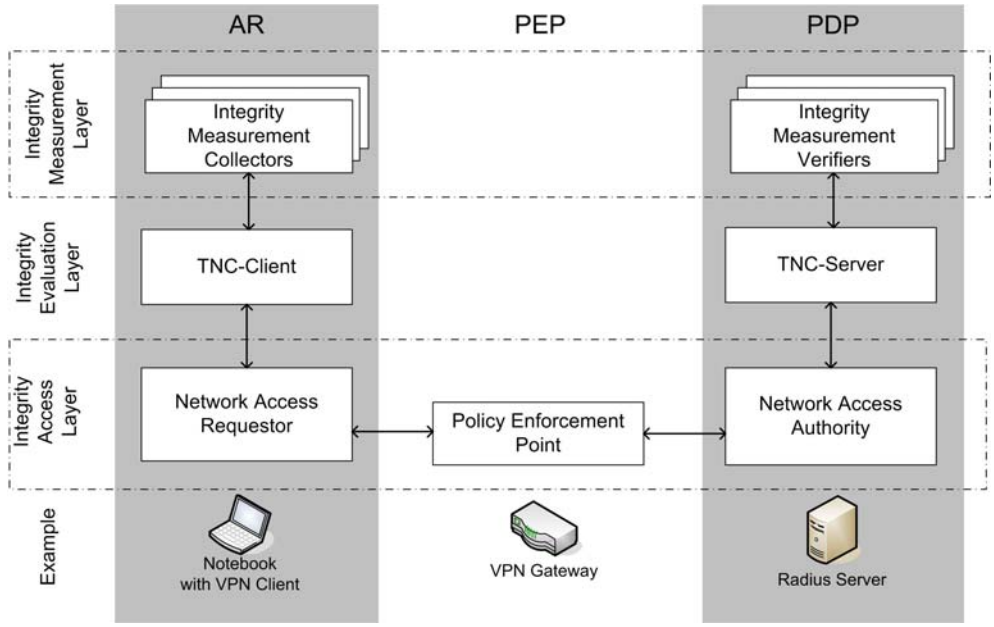

Fig. 3: Structure of TNC

The computer system by which a network connection to a TNC network is to be established is called the **Access Requestor** (AR). On the Access Requestor there are TNC components for a connection request, the transmission of measurement readings and the measurement itself.

Measurement of the individual components of the computer system is carried out by so-called "Integrity Measurement Collectors" (IMC). For each component to be measured there is a suitable IMC, for example one for virus scanners and one for the personal firewall. During the start-up of the system the IMCs are initialised by the TNC client on the accessing computer system in order to be able to collect measurement readings from the respective components during the establishment of a connection. The type of possible measurement readings is not initially restricted. For example, in the case of a virus scanner, information on the producer and the age of the virus signature may be important, whereas for a connected printer the version of the firmware and whether the printer has a fax function may be relevant. In order that an IMC can collect such detailed measurement readings it requires very precise knowledge of the hardware/software to be measured. Mostly such knowledge is only available to the producer of the hardware or software. The collaboration of the producer in the compilation or provision of an IMC is therefore essential.

On the network side there are two TNC elements:

The **Policy Decision Point** (PDP) represents the counterpart to the Access Requestor (AR). This is a server whose task it is to collect the measured readings of an Access Requestor and formulate an access decision with the help of policies. This decision is then passed to the point executing the access.

The Network Access Authority (NAA) in the Policy Decision Point decides whether an AR should be given access or not. To this purpose the NAA asks the TNC server whether the integrity measurements of the AR conform to the security policy.

At the PDP so-called "Integrity Measurement Verifiers" (IMV) represent the counterpart to the IMCs of the AR. Here too there are several IMVs for the different security components. For each security component to be verified there must be a suitable IMV in addition to the IMC (see Figure 3). These compare the measured readings transmitted on the basis of the rules laid down in the policies and notify the TNC server in the PDP of the result. With the partial results this then takes an overall decision on the integrity of the computer system and informs the Policy Enforcement Point via the NAA of this decision.

The **Policy Enforcement Point** (PEP) is the TNC element at the entry point to the network. Its tasks are to receive and pass on connection requests and execute the actions decided by the PDP.

As the entry point the PEP represents the first connection point to the network that is addressed. Incoming connection requests from an AR are passed directly to the PDP. After a PDP has taken its decision on the integrity of the AR, it informs the PEP, which then has to allow or prevent access to the network with its services on the basis of this decision.

According to the TNC specification, a PEP can be an independent computer system or integrated into the PDP or other network equipment. It is therefore possible to integrate the PEP directly into a VPN gateway or, in order to leave existing network structures untouched, in front of or behind this gateway.

3.3 Fields of Application

TNC's specifications have been kept general in order that it can be used as flexibly as possible in many applications. The two most important fields of application are presented below in brief.

- **Protection of the Intranet against Attacks from Outside**

  Through the expansion of a network connection, e.g. VPNs, with TNC it is possible to verify the integrity of the computer systems before access to the network in order to prevent this in the case of non-fulfilment of the security policy. Additionally, it is possible to bind access data to certain computer systems in order to prevent access by third-party computer systems with stolen access data.

- **Direct Protection of the Intranet**

  Protection against attacks from within. With 802.1x a widespread method of authentication is supported. Through the equipping of all computer systems in the intranet with 802.1x supplicants that have been expanded with TNC functions it is possible to reliably analyse one's own computer systems and those of guests with respect to their integrity and therefore minimise possible dangers from compromised computer systems.

4 Alternative Approaches

Besides TNC there are further NAC approaches in existence. The most prominent representatives are Microsoft NAP and Cisco NAC, which are explained here in more detail. The alternative solutions, including those of Microsoft and Cisco, are - in contrast to the open TNC specification - proprietary and therefore not interoperable as a matter of principle (for more information see section 5).

4.1 Microsoft NAP

With "Microsoft Network Access Protection" (Microsoft NAP) Microsoft are developing their own NAC solution [Micr06a]. The intention is for Microsoft NAP to be available with the server version built up on Microsoft Vista. Client software is being developed for both Vista and Windows XP. Network access is controlled by means of existing technologies such as 802.1x and offers (amongst other things) support for VPNs. It is therefore largely hardware-independent. On the software side, software products from Microsoft are essential for use with the Network Policy Server (NPS) and the required clients.

4.2 Cisco NAC

Cisco Network Admission Control (Cisco NAC) is part of the "Self-Defending Network" strategy and is also one of the policy-based access control systems [Cisc04]. Here Cisco uses its own hardware entirely, i.e. special NAC-compatible hardware is required in the entire network, resulting in the enforced use of Cisco hardware. Cisco NAC is already available on the market.

4.3 Further Solutions

In addition to the three "major" solutions presented, there are many further approaches from firms such as Check Point, Juniper Networks, StillSecure, Symantec and Vernier Networks

5 Critical Consideration

Some aspects of the NAC concepts are discussed critically below. These include the trustworthiness of the measurement readings recorded, administration and - in particular - interoperability.

5.1 Trustworthiness of the Measurement Readings

The security of NAC solutions depends on the trustworthiness of the measurement readings. These must be correctly measured and transmitted without modification to the NAC server.

With today's systems there is no possibility of guaranteeing correct measurement of the system status and its correct transmission. If the hardware or operating system of a computer system has been compromised, the measurement readings must also be considered as being no longer trustworthy, as they can be influenced by the malware at any time. However, as the measured values are to be used to discover the lack of integrity, the permanent risk of unnoticed falsification gives rise to long-term mistrust with respect to the measured values. This was demonstrated recently at the Black Hat Conference 2007 using Cisco NAC. By means of a modified Cisco Trust Agent (CTA) it was possible at all times - irrespective of the computer status - to gain access to a NAC-protected network [Heis07].

In order to get round this paradox, TNC offers a certain level of protection against manipulation of the hardware and the possibility of signing and therefore securing the transmission of the measured values through its optional and direct support for the Trusted Platform Module (TPM). However, without trustworthy determination of the measured values the level of security reached by the use of TPM is still limited.

Only with the introduction of suitable security platforms, such as the Turaya security platform of the EMSCB project [Emsc07] is it also possible to reliably determine the measurement readings for all security components.

However, this problem does not represent a specific problem of NAC approaches, but a general problem of today's computer systems that can be solved through the future use of security platforms built up on Trusted Computing.

5.2 Administration

Policy-based approaches involve an increased administration effort not only during the planning, but also during the operation of the network. This applies in particular to heterogeneous networks, as rules have to be defined for all end points located in the network with every imaginable configuration. Furthermore, there is also a danger that if the rules are too strict this will produce further side-effects. If two companies stipulate a virus scanner that is not a general product, i.e. producer-independent, but require special virus scanners from different producers, this may lead to incompatibilities on the notebook of an employee who works at both firms.

A further important point is the theme of patch management. It must be determined how the data of the policies can be kept up-to-date in an optimum manner. For example, it must be known at all times which version number the current virus signature of a virus scanner has, whether the personal firewall used is up-to-date (i.e. free from known gaps in security) and what the status of the patch database of the operating system and applications being used is. This information must be provided by the producers of the individual components and transmitted to the network operator. A network operator is therefore not only dependent on the producers - who have to equip their software with IMC and IMV functions - when building up the networks, but also during network operation. Here new forms of cooperation need to be found and contractual aspects of liability clarified.

It also has to be clarified how the patches that are to be incorporated can be obtained and installed. It would be possible to have a central patch server within the firm or to purchase them directly from the producers. On the one hand a central patch server minimises data traffic with the Internet. On the other hand it is difficult to keep all patches up-to-date. In practice a combination of both possibilities would therefore be ideal.

Further problems are imaginable in changing environments. If for reasons of compatibility (with other programs) an organisation stipulates an obsolescent software version and another organisation a later version, it is not possible to fulfil both policies. Constant up and downgrading of the versions is either not acceptable due to the effort involved, or simply not possible.

5.3 Interoperability and Standardisation

In spite of their similar structures and in most cases commonly used basic technologies, all of the NAC solutions on the market and under development are mostly proprietary and not compatible with one another. This is an obstacle to the increased use of such solutions, as the firms who decide on a solution that is on the market today have no guarantee that the selected solution will assert itself on the market. Through the selection of the solutions of producers that dominate the market there is therefore a particularly significant risk of the formation of de facto standards which will squeeze other competitors out of the market.

Furthermore, computer systems that are used in frequently changing environments, for example by fieldworkers, involve additional expenditure. These devices have to be prepared for all imaginable solutions.

Since the middle of 2006 several different attempts have been made to guarantee the interoperability of various solutions. These efforts are explained and discussed in brief below:

- In September 2006 Cisco and Microsoft announced their support for each other's solution [Micr06b]. This announcement does not mean that there is any direct adaptation of the architectures, but mainly instead support for both technologies on the client side. This means that the clients of both producers support networks with both Cisco NAC and Microsoft NAP. On the network side the network operators still have to decide on one of the two solutions.
- Various producers of NAC products (note: not Cisco NAC) design their products so that they are compatible with several solutions (example: Complete NAC from StillSecure). Here there is usually a concentration on the three "major" solutions of Cisco NAC, Microsoft NAP and TNC. This results in a competitive disadvantage for the solutions of other producers.
- Microsoft has made the Statement of Health (SoH) protocol available to the Trusted Computing Group. This specifies SoH as an additional interface (IF-TNCCS-SOH) between the TNC client and the TNC server [Trus07a] [Trus07b]. First of all this step offers the TNC-protected networks which support the new interface the advantage that as a client on the part of the AR the NAP clients supplied with the Windows Vista/Longhorn Server can be used. On the other hand it remains to be seen to what extent both interfaces can be administered jointly. If the administration effort should increase substantially, the two interfaces will probably not be able to coexist. Furthermore, there has been no statement so far concerning how new versions of the SoH protocol will be developed (i.e. whether together with other providers or by Microsoft alone) and whether these protocols will be published as an open specification.
- The only effort at standardisation in the NAC sector is currently being undertaken by the IETF with the "Network Endpoint Assessment Working Group". At the beginning of May 2007 the second version of the "Overview and Requirements" paper appeared [Ietf2007] for the "Network Endpoint Assessment" (NEA), in which amongst other items a reference model based on TNC, Cisco NAC and Microsoft NAP is described for concept formation. Members of this working group include Cisco and Symantec.

6 Conclusion

Within the framework of increasingly stronger networking within and between firms over unsecured networks, an increase in the trustworthiness of network communication is essential. NAC solutions such as the TNC specification of the Trusted Computing Group provide the opportunity to analyse end points with respect to their integrity, and therefore contribute to an increase in trustworthiness. In contrast to other, proprietary solution approaches, TNC has a major advantage due to its openness. As a result of this openness, TNC is bound neither to the hardware nor to the software of specific producers. This increases its acceptance and adaptation by all producers of system components and network technology, which represents an important factor in its success.

It should be borne in mind however that in all approaches the trustworthiness of the components is not sufficiently guaranteed without the use of secure operating system structures and that therefore the level of trustworthiness that can currently be attained is limited.

As TNC does not require specific hardware such as TPMs or special operating system structures and also supports existing network infrastructures (or builds on them), it can already be readily integrated into existing networks.

References

[Cisc04] Cisco Systems GmbH, Die Evolution des Self-Defending Network, 2004 http://www.cisco.com/global/AT/pdfs/prospekte/Securtiy_CNAC_032004.pdf

[Emsc07] Das EMSCB-Projekt, www.emscb.de

[Heis07] News: Ciscos Netzwerkzugangskontrolle NAC ausgetrickst – März 2007 http://www.heise.de/newsticker/meldung/mail/87663

[Hart05] Michael Hartmann, Trusted Network Connect - Netzwerkhygiene auf hohem Niveau, 2005, Datenschutz und Datensicherheit (DuD)

[Ietf07] Network Endpoint Assessment (NEA): Overview and Requirements, Mai 2007 http://www.ietf.org/internet-drafts/draft-ietf-nea-requirements-02.txt

[JuPo06] M. Jungbauer, N. Pohlmann: „Vertrauenswürdige Netzwerkverbindungen mit Trusted Computing - Sicher vernetzt?“ IT-Sicherheit – Management und Praxis, DATAKONTEXT-Fachverlag, 6/2006

[Micr06a] Microsoft Corporation, Network Access Protection - Homepage 2006 http://www.microsoft.com/technet/network/nap/default.mspx

[Micr06b] Microsoft Corporation, NAP-Whitepaper, 2006 http://www.microsoft.com/technet/network/nap/naparch.mspx

[Micr06b] Microsoft Corporation, Cisco and Microsoft Unveil Joint Architecture for NAC-NAP Interoperability, 2006 http://www.microsoft.com/presspass/press/2006/sep06/09-06SecStandardNACNAPPR.mspx

[Trus06] Trusted Computing Group: Trusted Network connect Subgroup, 2006 https://www.trustedcomputinggroup.org/groups/network

[Tru+06] Trusted Computing Group, TCG Trusted Network Connect TNC Architecture for Interoperability, 2006 https://www.trustedcomputinggroup.org/specs/TNC/TNC_Architecture_v1_2_r4.pdf

[Trus07a] Microsoft and Trusted Computing Group Announce Interoperability, Mai 2007 https://www.trustedcomputinggroup.org/news/press/TNC_NAP_interop_release_final_may_18.pdf

[Trus07b] TCG TNC IF-TNCCS: Protocol Bindings for SoH, Mai 2007 https://www.trustedcomputinggroup.org/specs/TNC/IF-TNCCS-SOH_v1.0_r8.pdf

Index

Integrity Check

Isolation

Microsoft NAP – Microsoft Network Access Protection

NAC – Network Access Control

Policy

PDP – Policy Decision Point

PEP – Policy Enforcement Point

Remediation (Phase)

TNC – Trusted Network Connect

Trusted Network Connection

VPN – Virtual Private Network