

Virtual Private Networks Coping with Complexity

Andreas Steffen

Security Group
Zürcher Hochschule Winterthur
CH-8401 Winterthur
andreas.steffen@zhwin.ch

Abstract: Large-scale deployment of virtual private networks with hundreds or thousands of clients means a constant battle with complexity that can only be won by setting up powerful authentication and authorization group policies. In this paper we are going to present some approaches for IP address, user, and access control management that have already been realized for the Linux FreeS/WAN IPsec stack or that are considered for implementation by the ZHW Security Group. First practical results from VPN production environments will be presented.

1. Introduction

The Security Group of the Zurich University of Applied Sciences in Winterthur, Switzerland (ZHW) is heavily involved in the development of the Linux “FreeS/WAN” IPsec stack (www.freeswan.org). We have contributed the X.509 certificate support to this popular OpenSource project and are currently moving towards the definition and implementation of IPsec security policies for large scale VPN deployment. It is our objective to make complex networks with hundreds or thousands of VPN clients manageable.

An especially demanding VPN application case is the so-called “road warrior” remote access scenario shown in Figure 1.




Figure 1: Road warrior remote access scenario

Published in “Security, E-Learning, E-Services – 17. DFN-Arbeitstagung über Kommunikationsnetze in Düsseldorf”, Jan von Knop, Wilhelm Haverkamp, Eike Jessen (Editors), GI-Edition - Lecture Notes in Informatics (LNI), P-44, Bonner Köllen Verlag 2003, pp. 289-302.

© Copyright 2003 Gesellschaft für Informatik e.V. (GI), Ahrstrasse 45, D-53175 Bonn.

In this particular constellation a mobile VPN user wants to access her company or campus network in a secure fashion from any arbitrary point in the global Internet. This means that the outer source address of the IPsec tunnel to be set up is assigned *dynamically* by the Internet service provider (ISP) at the local point of presence (POP). The same is true for many teleworkers who access the Internet from their home via an always-on ADSL or cable TV connection where often a daily change of the IP address is enforced by the network operator.

The road warrior case puts the VPN gateway securing the access to the company network into a difficult situation because it cannot identify the remote access clients on the basis of their IP source addresses. This precludes the use of pre-shared secrets as a means of authentication during *Main Mode* of the *Internet Key Exchange* (IKE) protocol [HC98], since the session key used to encrypt the identity in IKE message #5 as shown in Figure 2 depends also on the pre-shared secret. Thus we have a hen and egg problem: Without the a priori knowledge of the identity of a road warrior initiating a connection, the VPN gateway cannot select the correct pre-shared secret to decrypt IKE message #5 that in turn contains the identity information needed to identify the road warrior ...




Figure 2: Main Mode of the Internet Key Exchange (IKE) using pre-shared secrets

As a workaround *Aggressive Mode* is often used in low-end VPN solutions where the identity string ID_i of the road warrior is sent in unencrypted form. Unfortunately the $Hash_i$ field is also transmitted in the open, which creates a potential security hole by paving the way to an off-line dictionary attack on the pre-shared secret that was used to sign the hash.

Thus in order to avoid this potential weakness of *Aggressive Mode* and also to shield the identity of the remote access clients from prying eyes, IKE *Main Mode* using digital signatures and certificates as shown in Figure 3 should be used instead.




Figure 3: Main Mode of the Internet Key Exchange (IKE) using certificates

In this public key variant the symmetric session key encrypting the IKE exchange starting with message #5 depends solely on the Diffie-Hellman (DH) secret established by messages #3 and #4. Thus it becomes possible for the responder to extract the encrypted identity ID_i which in turn can be used to select the correct public key needed to verify the peer signature Sig_i . As a convenience most VPN implementations send along an optional X.509 certificate containing the required public key, so that it doesn't have to be fetched by other means, e.g. by a query to an LDAP server.

The use of X.509 certificates [Ho99] usually necessitates the setup of a public key infrastructure (PKI) based on a certification authority (CA) that issues and eventually revokes user and/or host certificates. The CA can either be run in-house or optionally be outsourced to an official trust center. This additional overhead puts a considerable burden on the initial deployment of a VPN solution but the investment pays off quickly because certificate-based user management scales extremely well with an increasing number of VPN clients, as we will show in section 3. User certificates, either in themselves or in conjunction with X.509 attribute certificates [FH02], [GS02] also form the ideal basis for sophisticated *access control* schemes as detailed in section 4.

Since road warriors carry dynamic outer source IP addresses assigned to them by their ISPs, it is highly desirable that their inner source IP addresses belong to a special segment of the company or campus network's address range, thereby forming an *extruded net*. This can be achieved by assigning a *Virtual IP* to the remote VPN client either statically or dynamically as shown in Figure 1. The use of virtual IP addresses facilitates both the *firewalling* of incoming IP packets after the IPsec tunnels have been terminated by the VPN gateway as well as the *routing* of return packets from hosts in the company or campus subnets back to the road warriors. How virtual IPs can be distributed dynamically using the *DHCP-over-IPsec* protocol [Pa03] will be described in section 2.

Thus in the context of road warrior remote access via IPsec tunnel mode we have identified the following three main areas

- Virtual IP Address Management (section 2)
- User and Certificate Management (section 3)
- Access Control Management (section 4)

that we are going to treat in more detail in the ensuing sections of this publication.

2. Virtual IP Address Management

2.1 Legacy Concepts

The legacy of more than thirty years of modem-based dial-in history is weighing heavily on us! Because of the great success of the point-to-point (PPP) protocol and its auxiliary IP control protocol (IPCP) [Mc92] that allows the automatic assignment of a client IP address as well as the specification of DNS and WINS servers [Co95], these principles were readily inherited by the Layer 2 Tunneling Protocol (L2TP) [To99] which encapsulates PPP frames in UDP datagrams in order to tunnel them over the Internet - thus creating a virtual end-to-end “copper wire”. Because the IPCP functionality is not directly supported by the IKE protocol [HC98], a L2TP solution is often preferred in remote access scenarios. In order to make up for the lacking cryptographic security of layer 2 tunnels, L2TP must additionally be secured by IPsec [Pa01] as shown in the upper half of Figure 4. This is exactly the approach chosen by Microsoft for their VPN remote access solution in the Windows 2000/XP operating systems.




Figure 4: L2TP over IPsec Transport Mode vs. IPsec Tunnel Mode

If only IP packets are to be transported over a virtual private network connection then encapsulating them in PPP/L2TP frames embedded in UDP datagrams and secured by IPsec transport mode of course raises the question why straight IPsec tunnels are not used in the first place, as depicted in the lower half of Figure 4. Such a layer 3 setup becomes possible if the dynamic assignment of virtual IP addresses and DNS/WINS server information can be solved somehow.

A proprietary approach called “config mode” [DP01] that was initially proposed by Cisco and subsequently adopted by other VPN products, introduces vendor-specific configuration messages into the IKE protocol. This concept has some convincing advantages when user information (including the virtual IP address to be assigned) is centrally stored on an LDAP or RADIUS server. The VPN gateway can then directly retrieve the user information from the directory server and forward the information to the road warriors thanks to the in-band IKE communication channel. This argument in favor of “config mode” has led to the official inclusion¹ of a *Configuration Payload* in the IKEv2 protocol draft [Ka03] currently being specified by the IETF IPsec working group.

¹ following a fiery debate by mailing list members on the pros and cons of config mode in the light of the general consent that DHCP-over-IPsec was the proper way of dynamically assigning virtual IP addresses.

2.2 DHCP-over-IPsec

Traditionally one or several DHCP servers are responsible for the dynamic assignment of IP addresses plus auxiliary information to networked hosts. Such important aspects as the periodic renewal of the address leases, the efficient management of the available address pool and the proper reaction to timeouts must be handled by a DHCP server in a stable and reliable fashion. Road warriors accessing a company network over a VPN tunnel need the same kind of information for the configuration of their virtual IP interfaces. Thus it is just a logical consequence to rely on a DHCP server to provide these services whereas the VPN gateway will restrict itself to the transparent forwarding of DHCP information over the IPsec payload channel, only.




Figure 5: DHCP-over-IPsec scenario

Figure 5 shows how a dynamic IP address assignment scheme can be realized by using the freshly standardized DHCP-over-IPsec protocol [Pa03]:

- a) In a first phase an IKE Main Mode negotiation is used to create an ISAKMP security association (ISAKMP SA) by establishing a trust relationship between the road warrior and the VPN gateway through mutual authentication. This ISAKMP SA is then the basis for all subsequent IPsec SAs that will be negotiated by the two tunnel end points.

- b) Next an IKE Quick Mode negotiation sets up an IPsec SA with a subnet mask of 0.0.0.0/0 in order to be able to tunnel the subsequent DHCP DISCOVER broadcast message originating from the remote access client. Since such a global network mask might pose a potential security risk, this so-called DHCP SA is restricted to traffic between the udp/bootpc and udp/bootps ports on the client and server side, respectively. Because a company's DHCP server usually is not hosted on the same box as the VPN gateway, a DHCP relay is needed on the gateway in order to forward the DHCP DISCOVER message to a DHCP server located somewhere in the back of the secured intranet. As an additional security measure, the lifetime of the DHCP SA will be set to the absolute minimum time needed to handle the exchange of the initial DHCP DISCOVER broadcast and the returned DHCP REPLY message.
- c) As soon as the road warrior gets the inner IP address, a normal Quick Mode negotiation is started, connecting the inner virtual IP address of the VPN client via the IPsec tunnel with the desired company network[s]. Each time when the DHCP lease will be up for renewal, the directed DHCP REQUEST unicast message can be tunnelled to the VPN gateway using this normal payload IPsec SA, so that a separate DHCP SA does not have to be set up anymore.

2.3 Linux IPsec Support for DHCP-over-IPsec and Virtual IPs

The Linux FreeS/WAN IPsec implementation supports the DHCP-over-IPsec protocol on the server side by providing a special DHCP relay agent that is able to relay a virtual IP back to the road warrior that asked for it. The connection definition shown in Figure 6 can handle an arbitrary number of road warriors having distinct virtual IP addresses.

In FreeS/WAN notation the *right* and *left* sides are interchangeable but in our examples we assume *left* to designate the *local* side and *right* the *remote* side. Thus *right=%any* means that any outer IP source address will be accepted as long as the peer presents a valid and trusted X.509 certificate (*rightrsasigkey=%cert*). The virtual IPs must lie within the range defined by the *rightsubnetwithin* parameter. *left=%defaultroute* designates the current IP address of the outgoing gateway interface, *leftsubnet* specifies the protected company network and *leftcert* loads the gateway's own X.509 certificate.




Figure 6: Linux IPsec road warrior connection definition using virtual IPs

3. User and Certificate Management

3.1 User and Host Certificates

In large-scale VPN deployments, the only viable way of doing mutual peer authentication both in an efficient and secure way is to use X.509 certificates. As Figure 7 shows, each VPN end point must possess either a user certificate (Antje, Bodo) or a host certificate (Gateway) which it sends to the peer as part of the IKE Main Mode negotiation. Authentication is based on an RSA or DSA signature generated by encrypting a hash value with the private key of the VPN end point. The peer can then easily verify the signature by decrypting it with the public key contained in the certificate and then comparing hashes. For this authentication process to be secure, it is crucial that full trust into the peer certificate can be established. This can be done by including the root certificate of the CA that issued the user/host certificate on each VPN end point. Trust is thus transferred to the CA certificate. If multi-tier certification authorities are used then the whole trust chain must be available to each VPN client. The intermediate CA certificates can either be loaded statically or made available via IKE Main Mode.




Figure 7: Authentication based on X.509 certificates

In the example of Figure 7 all end certificates have been issued by the Kool CA. Therefore the Kool CA certificate must be installed by each VPN end point in order to be able to establish the trust in the certificate received from the peer. Using the single connection definition from Figure 6, the gateway in Figure 7 will accept any road warrior that presents a valid user certificate issued by the Kool CA.

3.2 Certificate Revocation Lists

Putting trust into a CA certificate means that all end certificates issued by that CA are automatically trusted as well. Thus it is of utmost importance that an up-to-date certificate revocation list (CRL) [Ho99] is maintained by the CA which blacklists the serial numbers of all end certificates that have been revoked. How often an updated CRL is issued by the CA depends on the security policy that has been decided upon. Issuing intervals can range from a monthly, weekly, daily down to an hourly basis if unauthorized users or hosts must be locked out immediately. The VPN gateway and the VPN clients should periodically update their local copy of the CRL in step with the issuing intervals by downloading it from a HTTP and or LDAP server.

In order to give a hint to a VPN end point where a CRL can be downloaded from, one or several *crlDistributionPoints* [Ho99] can be embedded as a X.509v3 extension in each peer certificate. A *crlDistributionPoint* usually has the form of a Uniform Resource Indicator (URI) that can be used to automatically fetch a CRL from a HTTP or LDAP server.

Example of an HTTP URI in OpenSSL notation:

```
crlDistributionPoints=URI:http://www.kool.net/ca/cert.crl
```

Example of an LDAP URI in OpenSSL notation:

```
crlDistributionPoints=URI:ldap://ldap.kool.net/o=Kool AG,c=CH  
?certificateRevocationList?base  
?(objectClass=certificationAuthority)
```

Automatic fetching of CRLs based on *crlDistributionPoints* is supported by version 2.00 of Linux Free/SWAN. The necessary X.509 host and user certificates can be generated using the OpenSSL package by defining one or multiple *crlDistributionPoints* in the *openssl.cnf* configuration file.

3.3 Online Certificate Status Protocol

With an increasing number of users and frequent certificate revocations, CRLs can become quite bulky. Therefore a viable alternative to the download of huge revocation lists could be the use of the *Online Certificate Status Protocol* (OCSP) [My99]. With OCSP a VPN end point sends a request containing the serial number of the peer certificate to be verified to an OCSP server which returns a signed reply containing one of the indicators: *good*, *revoked* or *unknown*. The private key used to sign the response must belong either to the CA that issued the certificate in question, a Trusted Responder whose public key is trusted by the requester, or a CA Designated Responder (Authorized Responder) who holds a specially marked certificate issued directly by the CA, indicating that the responder may issue OCSP responses for that CA.

The Linux FreeS/WAN internet key exchange daemon currently does not support OCSP-based certificate verification whereas an OCSP server has been made available by the latest version 0.9.7 of the OpenSSL package.

4. Access Control Management

4.1 Joint Security Policies

Figure 8 shows a firewall / VPN gateway constellation that allows the easy implementation of a joint security policy encompassing both the termination of VPN tunnel connections and the selective access to different parts of the protected Intranet controlled by dedicated firewall rules. Using the user identity information available from a successful IKE negotiation a firewall could dynamically open selected parts of the network according to a predefined user profile. In order to implement such a joint scheme the firewall and the VPN software ideally should run on the same host computer but a variant where a separate VPN gateway is connected to the firewall via a dedicated network interface would also be feasible. The FreeS/WAN IPsec implementation is especially well suited because immediately after a successful set-up of a VPN connection an *updown* script is called that can insert any number of dynamic *iptables* firewall rules that will be enforced by the Linux 2.4 *netfilter* kernel module. Upon termination of a VPN tunnel the inserted rules will be automatically deleted and Intranet access is closed for that specific user.




Figure 8: Firewall / VPN gateway constellation allowing a joint security policy

In the following paragraphs we will discuss four different ways how a joint security policy could be implemented. Selective access can be based (in increasing order of complexity) on identity wildcards, intermediate certification authorities, attribute certificates, or Kerberos tickets.

4.2 Identity Wildcards

According to the *Internet IP Security Domain of Interpretation for ISAKMP* [Pi98] and the proposed *IPsec PKI profile* [KR03] the following identity types can be used in an IKE Main Mode authentication based on X.509 certificates (see Figure 3)

- **ID_IPV4_ADDR / ID_IPV6_ADDR** (IPv4 or IPv6 address)
- **ID_FQDN** (Fully Qualified Domain Name)
- **ID_USER_FQDN** (User e-mail address)
- **ID_DER_ASN1_DN** (X.500 Distinguished Name)

For road warriors with dynamic network addresses it doesn't make much sense to use an IP address as an ID, so only the latter three identity types remain. Identities sent as part of IKE Main Mode messages #5 and #6 must be certified by corresponding entries the X.509 certificate since the identity must be bound to a public key that can be used by the peer to check the signature. An ID_DER_ASN1_DN must equal the subject distinguished name (DN) of the certificate whereas an ID_FQDN or an ID_USER_FQDN must be contained in the certificate as a *subjectAltName* X509v3 extension [Ho99].

```
conn research
    right=%any
    rightid="C=CH, O=Kool AG, OU=R&D, CN=*" /* ID_DER_ASN1_DN */
    leftsubnet=10.1.1.0/24

conn sales
    right=%any
    rightid="*@sales.kool.net                /* ID_USER_FQDN */
    leftsubnet=10.1.2.0/24

conn it-hosts
    right=%any
    rightid="*.it.kool.net                  /* ID_FQDN */
    leftsubnet=10.1.3.4/32
```

Figure 9: IPsec policies based on identity wildcards

Figure 9 shows how identity wildcards designated by the ‘*’ character can be employed to specify detailed access control policies:

- The first connection definition restricts access to the R&D subnet 10.1.1.0/24 to any user (CN=*) who belongs to the Research department (OU=R&D).
- The second connection is opened to members of the Sales department by using a wildcard in the e-mail address (*.sales.kool.net)
- The third definition enables access to the server 10.1.3.4 only for those machines that have a hostname belonging to the particular sub-domain assigned to the IT department (*.it.kool.net).

Linux FreeS/WAN supports wildcards in the relative distinguished name fields (C=, O=, OU=, CN=, etc.) of ID_DER_ASN1_DN identities, although with only minor additions to the source code, ID_FQDN and ID_USER_FQDN wildcards could be implemented as well.

This wildcard mechanism is not a proprietary invention but is actually mandated by section 5.5.2. “The Property MatchIdentityValue” of the *IPsec Configuration Policy Information Model* [JRV03] that is currently being drafted by the IETF IP Security Policy working group. A big advantage of the identity wildcard approach is the fact that a flat X.509 trust hierarchy can be used. As a serious drawback it should be mentioned that the use of wildcards requires careful planning of the field structure of the certificate distinguished names in the case of ID_DER_ASN1_DN identities or of the sub-domains if ID_FQDN or ID_USER_FQDN types are used. Once deployed, it will be difficult to introduce large changes in the access control scheme without replacing all issued certificates.

4.3 Intermediate Certification Authorities

As an alternative to identity wildcards, *Intermediate Certification Authorities* could be used to divide the users into groups with distinct access profiles. Figure 10 shows an example of how this could be done:

- The research network 10.1.1.0/24 is available exclusively to users whose certificates have been issued by the R&D CA, but cannot be accessed e.g. by members of the Sales department.
- The access to the sales network 10.1.2.0/24 is restricted to owners of certificates issued by the Sales CA.

This approach has the advantage that user management can be decentralized. The company Root CA issues an intermediate CA certificate to both the R&D and Sales departments. Each department can then flexibly issue an arbitrary number of end certificates on its own. Thus decision paths can be kept short and “red tape” can be minimized.

Another VPN application where CA-based security policies can be extremely useful are *Extranets* which give direct access to restricted areas of the company network to prime customers or important suppliers. Since these external users usually belong to different trust domains, foreign certificates must be accepted and trusted. This can be done by exchanging CA certificates with the partner firms. In order to keep the foreign users off sensitive parts of the corporate network, access can be restricted to company certificate holders, only.




Figure 10: IPsec policies based on intermediate certification authorities

Linux FreeS/WAN has built-in support of IPsec policies based on [intermediate] CAs by using the exact notation shown in Figure 10.

4.4 Attribute Certificates

X.509 attribute certificates (ACs) [FH02] introduce a clean separation of the tasks of user authentication and user authorization. ACs can contain an arbitrary number of *target*, *role* and *group* attributes and thus can make the most complex access control schemes feasible. Since attribute certificates usually possess a very short life time of 1..24 hours, they don't have to be revoked – they just expire!

The basic working principles of an attribute certificate scheme are shown in Figure 11.

- A Certification Authority (CA) issues long-lived user and host certificates that contain only fields that are rarely changed.
- A possibly decentralized Authorization Authority (AA) issues short-lived attribute certificates containing the current access control profile of each user.
- An AC is bound to a specific user or host certificate by including the serial number and issuing CA of the certificate holder.




Figure 11: Authorization scheme based on X.509 attribute certificates

Two ZHW students [GS02] have created a set of rudimentary tools that allow the generation of attribute certificates. Another group is currently developing a graphical management interface on top of the command line tools. A third group is integrating AC support into Linux FreeS/WAN. The tentative notation that will be used for specifying the connection attributes is shown in Figure 12. ACs will be fetched via LDAP queries.

```
conn research
  right=%any
  rightgroups=Research
  lefttarget=VPN
  leftsubnet=10.1.1.0/24

conn sales
  right=%any
  rightgroups=Sales
  lefttarget=VPN
  leftsubnet=10.1.2.0/24
```

Figure 12: IPsec policies based on attribute certificates

4.5 Kerberos Tickets

Kerberos V [KN93] is the default user authentication and access control scheme for the Windows 2000/XP operating systems. As a proprietary extension to the standard Kerberos ticket Microsoft has defined a so-called *Privilege Access Certificate* (PAC) which can be regarded as a kind of attribute certificate containing access control rights. Because of Microsoft's large market share, solutions based on Windows access control mechanisms cannot be ignored and are therefore currently being studied by the ZHW Security Group. A semester project conducted by two ZHW students [ST03] showed that Microsoft's use of Kerberos in setting up IPsec tunnels is highly proprietary, to say the least. Not only are special Vendor ID messages used, but private IKE payload types are also introduced. They are used to carry a Kerberos ticket from initiator to responder and a matching response back to the initiator. This is very strange in the light of the fact that ISAKMP [Ma98] lists a *Kerberos token* among the official certificate types that can be exchanged via the IKE protocol.

In our opinion, for the time being the use of Kerberos tickets for IPsec authentication will rather remain confined to Microsoft's Windows operating systems.

5. Practical Results

Although introduced only about a year ago, virtual IP addresses that are dynamically assigned via the DHCP-over-IPsec protocol are already in wide use. An experimental VPN remote access solution currently running at the Zurich University of Applied Sciences in Winterthur, Switzerland, supports IPsec policies based on identity wildcards in order to differentiate between staff and students. CA-based rules are also employed to give holders of foreign certificates limited access to the campus network. Full support of sophisticated IPsec policies using attribute certificates is currently being implemented and will be deployed and tested in the second half of this year. Unfortunately literature [COB03], [Th99], [JMT98] on the practical use of attribute certificates is rather sparse, so that we will tread on largely unknown territory.

6. Conclusions

We have shown that the complexity of setting up large virtual private networks can be coped with successfully if and only if the management of IP addresses, of users and certificates, and of the increasingly complex access control rights can be kept under tight control. The IP address management problem can be solved elegantly using the DHCP-over-IPsec protocol. If X.509 certificates are used for user and host authentication then selective access to network resources can be granted either on the basis of identity wildcards, intermediate certification authorities, or on the use of attribute certificates. Which one of these methods is going to prevail in the long run remains yet to be seen.

Bibliography

- [Co95] Cobb, S.: PPP Internet Protocol Control Protocol Extensions for Name Server Addresses. IETF RFC 1877, 1995.
- [COB03] Chadwick, D., Otenko, A., Ball, E.: Role-Based Access Control With X.509 Attribute Certificates. In: IEEE Internet Computing, Volume 7, Number 2, 2003; pp. 62-69.
- [DP01] Dukes, D., Pereira, R.: The ISAKMP Configuration Method. IETF Internet Draft <draft-dukes-ike-mode-cfg-02.txt>, 2001.
- [FH02] Farrell, S., Housley, R.: An Internet Attribute Certificate Profile for Authorization. IETF RFC 3281, 2002.
- [GS02] Gallizzi, U., Seiler A.: Attributszertifikate. Projektarbeit, Zürcher Hochschule Winterthur, 2002.
- [HC98] Harkins, D., Carrel, D.: The Internet Key Exchange (IKE). IETF RFC 2409, 1998.
- [Ho99] Housley, R., Ford, W., Polk, W., Solo, D.: Internet X.509 Public Key Infrastructure Certificate and CRL Profile (PKIX). IETF RFC 2459, 1999.
- [JMT98] Johnston, W., Mudumbai, S., Thompson, M.: Authorization and Attribute Certificates for Widely Distributed Access Control. In: Proc. IEEE 7th Int. Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises, 1998.
- [JRV03] Jason, J., Rafalow, L., Vyncke, E.: IPsec Configuration Policy Information Model. IETF Internet Draft <draft-ietf-ipsec-config-policy-model-07.txt>, 2003.
- [Ka03] Kaufman, C.: Internet Key Exchange (IKEv2) Protocol. IETF Internet Draft <draft-ietf-ipsec-ikev2-07.txt>, 2003.
- [KN93] Kohl, J., Neuman, C.: The Kerberos Network Authentication Services (V5). IETF RFC 1510, 1993.
- [KR03] Korver, B., Rescorla, E.: The Internet IP Security PKI Profile of ISAKMP and PKIX. IETF Internet Draft <draft-ietf-ipsec-pki-profile-02.txt>, 2003.
- [Ma98] Maughan, D., Schertler, M., Schneider, M., Turner, J.: Internet Security Association and Key Management Protocol (ISAKMP). IETF RFC 2408, 1998.
- [Mc92] McGregor, G.: The PPP Internet Protocol Control Protocol (IPCP). IETF RFC 1332, 1992.
- [My99] Myers, M., Ankney, R., Malpani, A., Galperin, S., Adams, C.: Online Certificate Status Protocol (OCSP). IETF RFC 2560, 1999.
- [Pa03] Patel, B., Aboba, B., Kelly, S., Gupta, V.: Dynamic Host Configuration Protocol (DHCPv4) Configuration of IPsec Tunnel Mode. IETF RFC 3456, 2003.
- [Pa01] Patel, B., Aboba, B., Dixon, W., Zorn, G., Booth, S.: Securing L2TP using IPsec. IETF RFC 3193, 2001.
- [Pi98] Piper, D.: The Internet IP Security Domain of Interpretation for ISAKMP. IETF RFC 2407, 1998.
- [ST03] Sanmiguel, J., Tümay, T.: IPsec Authentifizierung mit Kerberos unter Windows 2000. Projektarbeit, Zürcher Hochschule Winterthur, 2003.
- [Th99] Thompson, M., Johnston, W., Mudumbai, S., Hoo, G., Jackson, K., Essiari, A.: Certificate-based Access Control for Widely Distributed Resources. In: Proc. 8th Usenix Security Symposium, 1999.
- [To99] Townsley, W., Valencia, A., Rubens, A., Pall, G., Zorn, G., Palter, B.: Layer Two Tunneling Protocol "L2TP". IETF RFC 2661, 1999.