



# Global Information Assurance Certification Paper

Copyright SANS Institute  
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

## Interested in learning more?

Check out the list of upcoming events offering  
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"  
at <http://www.giac.org/registration/gsec>

# Remote Access for ISP On-Call Engineers

A Case Study

© SANS Institute 2000 - 2002, Author retains full rights.

Bruce A. Kaalund  
GSEC –Practical Assignment  
Version 1.4 (amended April 8, 2002)

## Table of Contents

<b>SUMMARY .....</b>	<b>3</b>
<b>THE INITIAL SITUATION .....</b>	<b>3</b>
<b>CHALLENGES .....</b>	<b>4</b>
POLICY DEVELOPMENT .....	5
<i>User Specification</i> .....	6
<i>Equipment Specification</i> .....	6
VPN ACCESS .....	6
<i>Certificate Security</i> .....	7
PROCEDURE IMPLEMENTATION .....	8
<b>RESULTS .....</b>	<b>9</b>
PRIVACY ISSUE .....	9
<b>FUTURE ISSUES FOR IMPROVEMENT .....</b>	<b>9</b>
<b>REFERENCES .....</b>	<b>11</b>

© SANS Institute 2000 - 2002, Author retains full rights.

## Summary

Running an ISP is a full-time job. The users of the network expect it to be there every day, at any hour. Because of this, an ISP is expected to have its most senior technical engineers available at any time to solve a critical problem. The ideal situation would be for an ISP to staff three 8-hour shifts with everything from help desk telephone staff to highly talented engineers to deal with critical issues at any time. In reality, it is much easier to find network operations (NOC) staff to work the “swing” shift or the “graveyard” shift than it is to find an experienced engineer who will work those hours. So, to compensate, many ISPs have their engineers work the “regular” business day (8 AM to 5 PM), and put them on a rotating “on-call” schedule. In the event an issue occurs that the NOC staff cannot handle, the on-call engineer can be contacted to work on the issue.

Early on, we recognized a problem. We found that the majority of our engineers live anywhere from 30 minutes to over an hour’s drive from the office. This can be an issue in a critical failure, because it would impact the MTTR (mean time to repair) statistics of the ISP, leading to the perception of poor service. Since every employee of the ISP has broadband access to the network, the decision was made to provide access to the network devices remotely. This way, an engineer called at “oh dark thirty” to fix a device’s configuration would be able to connect their laptop to their broadband access. Using a VPN, the engineer would access the network device in question, and make the necessary change to the device.

## The Initial Situation

I am responsible for the security of a large broadband ISP. We have a large number of users to our system, spread out across the continental US, who depend upon the Internet access we provide to do everything from research homework assignments to conduct business transactions. Our users expect to have access to their e-mail or to view the latest trailer to the summer’s biggest blockbuster to purchase tickets to see that movie. This requires that the service be available at all times with the throughput the user expects.

With any system that is dependent upon electronic devices for its performance and capability, malfunctions can and do occur. When they occur, quick resolution is necessary to get users “surfing” again. It is inevitable that such a malfunction can occur outside of regular business hours, when our engineers are at home with their families, or at 2 AM, when they are asleep. Our NOC staff is well trained to handle simple issues (e. g., restarting a daemon) but in the event the issue requires a more talented individual, our engineers are on a rotating on-call schedule. The on-call engineer will be paged or called on their company-issued cell phone by the NOC, and will be expected to analyze and correct whatever issue they are presented with, in the ISP’s established MTTR standards.

This work required having the engineer drive into the office and access their PC to correct the issue. This is a problem for much of the engineering staff, as many live as much as an hour away from the office. Such a delay would have an impact on not only the MTTR, but would damage the user's perception of the network's reliability. In addition, since the engineers work a regular shift, it required a potentially tired engineer to come back to the office and work some more. Besides the morale issues this created, it was not healthy to have a tired engineer drive back to work, and risk an automobile accident due to falling asleep at the wheel.

My peers in Engineering and Operations came to my security group to see if we could develop a solution to provide access to the network devices from a remote location, such as an engineer's home. This way, an engineer getting a call at "oh dark thirty" could simply go to their computer and access the device that needed to be tended to, and make the correction. This not only would help the MTTR and the reliability perception, it would prevent a possibly tired engineer from getting behind the wheel of their car, eliminating the potential issues that such a move may bring.

## Challenges

With the technology available, telecommuting has risen in popularity for many industries and governments. Although it may appear to be a trivial exercise to provide remote access to our engineers, it does come with some technical and managerial issues:

1. All of the devices that constitute the ISP network have private addresses (i. e., RFC1918 addresses, which are not routable over the Internet). All of our Internet users receive a public address. Our router access control lists strictly limit network device access to those computers who are on a specific subnet of the private network, denying access to all computers with a public address. So, an Internet user sitting at their PC in their home is prohibited from accessing any network device used to connect them through our network to the Internet.
2. Although we specified SSH access as the standard for logging in and working on our network devices, there are some devices that have poor or vulnerable implementations of SSH. Those devices are accessed via telnet, which transmits the passwords and session in the clear over the network. This makes the session prone to packet sniffing or session hijacking.<sup>1</sup>
3. A fallout of providing this service would be the "me too" syndrome. Once we open up remote access to the network devices, everybody would claim to need access to the devices from home. Many of those requiring access are those who would not be responsible for the device, they "just want to

---

<sup>1</sup> Garfinkel, page 495

look around”, or have the same toys everyone else has. This created an issue for my Security group due to lack of control of access. Lack of controlling access to a device violates the principle of least privilege, which “. . . dictates that processes and/or users should only be granted the minimum set of rights and privileges necessary to do their work.”<sup>2</sup> We do provide read only access to anyone who requires it. Unfortunately, some read-only access allows the user to be able to make limited changes. Such changes could have an effect on the performance of the device.

4. There was concern about allowing engineers to access the network using their home PCs. The primary issue concerned potential abuse. Should an engineer become disgruntled with the ISP, and decide to create mischief, giving them the ability to access the network devices from home would present an opportunity for them to do so without being in the office. Sure, we could detect the unauthorized intrusion, and if they are doing mischief, we would be able to document it. However, getting access to the home PC becomes an issue. Being personal property, we would have to convince a judge to issue a subpoena for the device. This would take some time, and the box could have been sanitized by the time we seized it, making forensics at best difficult.

In addition, there are the normal concerns about the insecure nature of home PCs.<sup>3</sup> These PCs are already connected to the Internet by the cable modem and are “. . . linked to the network all the time, yet they lack the protection and policy enforcement of (the) corporate firewall . . .”<sup>4</sup> As such, these PCs may “. . . bypass IT safeguards and provide a back door for threats to the network.”<sup>5</sup> The average user’s home PC may not have a firewall protecting it from outside network intrusion and the installation of Trojans. Many times, the latest anti-virus software updates haven’t been downloaded and kept up-to-date. Considering the fact that many home PCs are used for e-mail, this could place a vulnerable machine on to the network, without our being able to verify it meets our security requirements. In the event of an e-mail spread virus attack, we would not have access to assure that the PC is cleaned and protected.

We were able to deal with the above challenges by doing the following techniques described below.

### ***Policy Development***

The initial work that was done to implement a solution was the development of a policy for deciding who received access, and the procedures for implementing

---

<sup>2</sup> King, page 50

<sup>3</sup> Tweney, url: <http://www.business2.com/articles/web/0,1653,16680,FF.html>

<sup>4</sup> Yasin, url: <http://www.internetwk.com/story/INW20000503S0001>

<sup>5</sup> Donald, url: [http://rr.sans.org/encryption/remote\\_clients.php](http://rr.sans.org/encryption/remote_clients.php)

the policy. This required the development of an issue-specific policy, which “Address specific issues of concern to the organization”<sup>6</sup>.

## **User Specification**

Our policy was developed in conjunction with the heads of the operations and engineering groups of the ISP. It was these people who clearly specified who could access the network remotely; this helped to shape the policy and to garner leadership support for the policy. The decision was to only allow those personnel has access remotely who have on-call duties and would have responsibility for the configuration of a device. This eliminated those technicians whose job responsibilities included on-call changing out hardware, but did not involve software configuration. It also eliminated NOC personnel who had software configuration duties but did not have on-call duties. So, using the principle of least privilege, we wrote the policy to allow a limited set of engineers access the network devices from either home or on the road.

Because staffing is an issue in today's IT shops, a few of the engineers are actually contractors or consultants who are working an especially difficult problem. We were required by the head of engineering to allow these people to have remote access to the network, on a limited-time basis. We wrote them into the policy.

## **Equipment Specification**

Our policy also dictated that the remote access software must require the installation of a digital certificate, and could only be installed on a company-provided laptop PC. This makes the elimination of engineers using their home PC simple. Logically, a digital certificate can be used to authenticate a person to have access. But, in a practical sense, a digital certificate authenticates the *device* it is installed on to have access. By mandating the installation of the digital certificate on a company-provided laptop (or consultancy or contracting company supplied laptop) we effectively eliminate home PC's from being used. This has additional benefits. Now, should an intrusion be suspected, we would be able to seize the laptop when the engineer returns to the office, and perform the necessary forensics to gather the evidence. Now, this isn't foolproof, but it sure beats having to convince a judge to sign a subpoena! In addition, because we control the frequency of the anti-virus updates, we are assured that these devices are inoculated against the majority of viruses they could catch.

## **VPN Access**

Based on the policy, the next thing we did was to specify that remote access to the network would be through a Virtual Private Network (VPN) set up specifically for remote access users. For this, we used the Nokia CryptoCluster 2500, which comes with a built-in Certification Authority, used for “Verifying the identification

---

<sup>6</sup> Dulany, url: <http://rr.sans.org/policy/tech.php>

of communicating VPN gateways and mobile users . . .”<sup>7</sup> Using a VPN allows us to accomplish a couple of requirements. The first is to provide a secure channel across the Internet into the private network. A client installed on the user’s PC provides the encryption capability to form one end of the VPN tunnel, with a connection to our VPN gateway appliance located in our data center. Through the use of digital certificates, we were able to provide authentication of the user to the gateway, and is used in the Internet Key Exchange (IKE) negotiation. Second, the use of a VPN allows access to our private network from the public Internet. Along with the remote access client, each user will have the networks they are allowed to access on the private side of the VPN gateway specified prior to installation of the client.

This arrangement has the added advantage of allowing an engineer to access the network devices when they are outside of the home.<sup>8</sup> An engineer who is on the road in training or attending a meeting could be called upon to support the NOC. This would require the engineer to return to their hotel room, connect to the Internet (either using dial-up or the high-speed access the hotel may provide) and access the device using the VPN.

Using a remote access VPN helps to *reduce* the risk of packet sniffing or session hijacking of telnet sessions. Because the session data is riding an encrypted tunnel from the user’s PC to the VPN gateway, it becomes extremely difficult for someone with a sniffer to glean any valuable information from the collected packets. The risk is increased once it passes through the VPN gateway, as it traverses the private network in the clear. However, the private network is composed of circuits owned and operated by the ISP, and the devices on the network are hardened, and the network is audited and scanned on a regular basis, so this helps to reduce the risk of a rogue sniffer being placed on a device. However, the threat of internal personnel performing mischief remains a task we deal with on a regular basis; this is a subject for another time!

## Certificate Security

One question that arose during the design phase was the possibility of copying and/or removing the certificate and client from the laptop to another machine. We felt few of the engineers who received the access would try to subvert the policy by copying the client and the certificate off of their laptop and onto their home PC. My security engineers anticipated this effort, and tested the possibility before the rollout began.

The best place to understand how the process worked, we need to understand how the certificate is generated, and how it is installed on the remote laptop. According to the Nokia Technical Note;

---

<sup>7</sup> Technical Note. url: <http://www.nokia.com/vpn/pdf/certificate.pdf>

<sup>8</sup> Torello. url: [http://rr.sans.org/encryption/remote\\_access.php](http://rr.sans.org/encryption/remote_access.php)



For mobile user configuration, the administrator on behalf of the user fills out the name and identification of the employee in the VPN Policy Manager software and clicks the Submit button. The Nokia VPN Gateway, acting as CA, generates a certificate for the remote client. The certificate is returned to the computer running the VPN Policy Manager software over an encrypted SSL link. CA data stored encrypted on the hard drive if the computer running the VPN Policy Manager.<sup>9</sup>

In addition, some key files were hidden on the hard drive after installation. Such a move would be unsuccessful, unless the engineer knew about and transferred all necessary files. Since the client was not downloadable off of the Internet, we had control over the distribution of both the client and the certificate. Both were kept in a safe controlled by Security. In the event an engineer came back and said their access didn't work, a test bed we set up in the lab was used to verify that the client installed worked on the laptop provided.

### ***Procedure Implementation***

The implementation of the policy is based on procedures developed by my staff to qualify engineers for access, and for the installation of the software and certificate. The procedure, which was approved by the heads of engineering and operations, reviewed by HR and Legal, and agreed to by my management peers, is as follows:

1. The engineer desiring access must complete a request for remote access form. This form asked for information considered important, such as home address, MAC address of the Ethernet port on the laptop, and network devices the engineer is authorized to access. If a contractor or consultant is requesting access, they must provide the name and address of the firm they work for. The engineer and their direct manager must sign the request. For consultants and contractors, the firm representative must also sign.
2. At the same time the remote access request is completed and turned in, a policy acknowledgement form is also required. This consists of the actual remote access policy, and a sign-off sheet that indicates the engineer has received, read, and is in full compliance with the policy. The engineer keeps the policy and returns the signed acknowledgement to my shop. If either the acknowledgement is not returned, or the request form is not completed in full, the request is denied, no exceptions.
3. Once the paperwork has been completed, and the request has been approved by Security, the engineer is instructed to hand carry their laptop to a security engineer for the installation of the software and certificate. The security engineer verifies that the laptop they are presented is the actual laptop listed on the request. Once the software and certificate have

---

<sup>9</sup> Technical Note. url: <http://www.nokia.com/vpn/pdf/certificate.pdf>

been installed, the engineer is instructed to access the network when they get home. The security engineer resolves any issues.

## Results

The implementation of the remote access solution was deemed a success by the leadership of the ISP, and the engineers who were approved to receive access. The defined policy was our strongest defense for denying access to others through the implementation of the principle of least access. This policy did create some friction, particularly with the field technicians whose responsibility consisted of hardware installation and replacement. They reasoned they needed to look at the device before heading out to make the repair. However the Operations process provided was detailed and thorough, so that the device was reviewed thoroughly by the NOC before a decision to dispatch a technician to make a repair was made. If a technician received a call to replace some hardware, all necessary investigation and verification was done before the tech received the call, and they were given specific instructions on what to replace. This eliminated their need to access the device remotely, it also reduced the number of times the on-call field technician was called to troubleshoot a problem.

## Privacy issue

An unexpected fallout of our procedure was the refusal of some engineers to supply their home address on the application, based on privacy concerns. We understood the concern, with the prevalence in identity theft.<sup>10</sup> They reasoned that this information could be found in the HR records, so we in Security could get this information there. However, Security does not have access to HR records. So we needed to have a record of where the engineer resided, as this was going to be the primary place they would use the access. As part of our security process, we locked all applications in a secure cabinet, and limited access to the cabinet to two people. Finally, we had their concerns over the situation and the process we developed reviewed by HR, Legal, and the Chief Privacy Officer of the ISP. All groups approved of our process and what we were doing. In the end, one engineer continued to refuse to give us his home address, and he never received the remote access client or certificate.

## Future Issues for Improvement

To this day, the remote access solution we implemented is working successfully. On-call engineers are able to make the necessary authorized configuration changes without having to drive in to the office at “oh dark thirty”. The process for obtaining remote access has become an accepted procedure for new engineers who management requires to be on call. HR includes Security in the resignation or dismissal of engineers, so we can promptly revoke their VPN certificate; the laptop is subsequently recovered and the hard drive re-imaged (client and certificate are erased in the process) before it is given to another person.

---

<sup>10</sup> U. S. Federal Trade Commission. url: <http://www.ftc.gov/bcp/online/pubs/credit/idtheft.pdf>

However, as with any process, there are opportunities for improvement:

1. **New Equipment.** In December of 2001, Nokia announced as a part of its partnership with Check Point Software, the CryptoCluster line of VPN appliances would be phased out<sup>11</sup>. This entailed the support for the current software, and the release and support of the promised next revision of software. But, they would soon stop producing the devices, and would eventually send the product line into the end-of-life cycle. This means we have to look at new means of providing this service once Nokia stops supporting the device. This review is expected to take a couple of years, so we can take advantage of improving technology, while using the presently working solution to improve the ISP's ROI on the equipment. Because we are using digital certificates in this solution, we may have to implement our own freestanding Certification Authority to continue providing remote access.
2. **Expanded User Base.** With the expansion of the products being offered by the ISP, we expect at some time in the future, employees who are presently serving as field technicians will handle the configuration of certain devices on the network. These employees will then be justified in requiring the remote access VPN solution. Such a change will present issues that need to be mitigated. How do we assure that a field tech, located across the continent, gets the client and certificate installed on the company laptop, and not on their home PC? Policy is fine, but there is no way to "trust but verify" if we just send a CD of the software to a tech and expect that it be installed on their company laptop. How do we determine if the tech really has a company-issued laptop, that the proper information was entered onto the application, and that the client and certificate are installed on that company-issued laptop? Who assumes the risk of our policy being violated by a rogue tech? These and other unforeseen issues have to be analyzed and mitigation strategies developed for them.
3. **Further Refinement of Read-Only Access.** Read-only access, in its larger sense, does not allow a technician or engineer to make configuration changes to a device. However, we have found on certain devices, read-only access provides the ability to make very limited changes to the device. Some of these changes may create a condition where a group of Internet users are denied access to their e-mail or the ability to surf. We need to do a detailed review of the capabilities of read-only access on each device, and determine the level of risk the ISP is exposed to. Based on that risk assessment, we may have to change the access levels of certain engineers and technicians, possibly denying access to specific staff.

---

<sup>11</sup> Greene, url: <http://www.nwfusion.com/newsletters/vpn/2001/01155369.html>

## References

Garfinkel, Simson, and Gene Spafford. Practical Unix and Internet Security, 2<sup>nd</sup> edition. O'Reilly and Associates, Inc. 1996. page 495

King, Christopher M, Curtis E. Dalton, and T. Ertem Osmanoglu. Security Architecture, Design, Deployment & Operations. Osborne/McGraw-Hill. 2001. page 50

Tweney, Dylan, "Are Home PCs a Backdoor Into Your Corporate Network?". Business 2.0. August 2, 2001. url: <http://www.business2.com/articles/web/0,1653,16680,FF.html>

Yasin, Rutrell. "Telecommuters On Security Alert". InternetWeek. May 3, 2000. url: <http://www.internetwk.com/story/INW20000503S0001>

Donald, G. Mac. "Secure Access of Network Resources by Remote Clients". SANS Information Security Reading Room. February 20, 2002. url: [http://rr.sans.org/encryption/remote\\_clients.php](http://rr.sans.org/encryption/remote_clients.php)

Dulany, Kevin M. "Security, It's Not Just Technical". SANS Information Security Reading Room. January 15, 2002. url: <http://rr.sans.org/policy/tech.php>

Technical Note. "The Nokia Built-in Certification Authority: An Immediate Means for Issuing Digital Certificates for Use in a Nokia VPN". Nokia Web Site. url: <http://www.nokia.com/vpn/pdf/certificate.pdf>

Torello, John. "Implementing Remote Access: Security, Usability and Management". SANS Information Security Reading Room. June 11, 2001. url: [http://rr.sans.org/encryption/remote\\_access.php](http://rr.sans.org/encryption/remote_access.php)

U. S. Federal Trade Commission. "ID Theft: When Bad Things Happen to Your Good Name". February 2002. url: <http://www.ftc.gov/bcp/online/pubs/credit/idtheft.pdf>

Greene, Tim. "A Happy Partnership". Network World VPNs Newsletter. December 17, 2001. url: <http://www.nwfusion.com/newsletters/vpn/2001/01155369.html>