

G'day,

A quick response the documents Tony just sent out:

I'm not exactly sure where Greg got the idea that you had to setup multiple static tunnels from CONUS institutions to some point in Denver. Yes, there are issues with key exchange and interoperability between different IPsec clients, but if you just stick with the a single client you're fine. For example, cisco's current client supports:

Windows 98, ME, NT 4.0, 2000, XP, Linux (Intel), Solaris (UltraSparc 32- and 64-bit), and Mac OS X 10.1 and 10.2 (Jaguar)

Above taken from:

<http://www.cisco.com/en/US/products/sw/secursw/ps2308/index.html>

You can read all about the cisco VPN client if you want to go there.

Once setup, you enter a username + password and off you go. You can use any application without alteration. Yes, I would expect RPSC to provide copies of the CISCO client and support said client. There are freeware clients that support IPsec vpn's but they are not guaranteed to work with the cisco products (for example <http://www.freeswan.org/>).

SSL based VPN's are attractive in that you have the potential to access specific services via a web interface. It's a bit of a downside in that if you have to use a custom application you end up having to install a bit of client software anyway. For instance, the product that Greg mentioned was Nortel Networks Alteon VPN hardware and to read all about that you can see:

<http://www.nortelnetworks.com/products/01/alteon/sslvpn/index.html>

But if you want to use all of your old applications seamlessly you have to install some client software anyway.

So, potentially you get the same thing from both solutions. You install some software on your machine and securely connect over the internet to pole. You can potentially get both services. For instance, cisco (can you tell I own stock in them?) is scheduled to come out with a module for their VPN concentrator 3000 (see: <http://www.cisco.com/en/US/products/hw/vpndevc/ps2284/index.html> for more details on this device) that will let it do both IPsec and SSL based VPN's.

The devil is then in questions of how easy is it to install clients for either solution, what's the price difference between the two options, and so on. But do we really have to tell Raytheon exactly how to implement the desired access to pole? Isn't that what the NSF pays them for? How about something like:

In order to have a successful science program at the South Pole, it is required that science groups be able securely access any of their computers from any node on the internet. Solutions to this problem exist. For instance, NASA GSFC uses Cisco's IPSec VPN technology to support roaming users. Many companies are using SSL based VPN's to provide the same service. Legacy applications must be supported and if any client software is required on the science users CONUS machine it is required to support platforms A, B, C, and D. RPSC will be expected to provide any necessary VPN clients to end users. RPSC is requested to investigate these and any other technologies, decide how they can provide the services required and report back to the SPUC for further review before any further action is taken.

I think this might be a little closer to what Bob was aiming for in his original document.

If we really have to tell them exactly how they should implement polar security, might it be worth while to spend some money testing? I'm not exceedingly familiar with Nortel's product line, but I know you can get a Cisco VPN box for a small network for ~ \$500 US or so. I could set one up somewhere and people could try using it and see what they think.

Good to here some discussion on the topic though. Keep in coming..

Cheers,

Matt Newcomb

Hi [Matt], Everyone,

Thank you, [Matt], for your comments. I'm off in a few hours to IETF (Internet Engineering Task Force), where as usual one of the topics will be, "why aren't people using IPSec?". So maybe a few additional remarks are appropriate.

And by the way, my wife and I between us own stock in both Cisco and Nortel, so my recommendations are unbiased :-)

[Matt] asked why I assumed that there would be permanent IPSec VPN tunnels set up to some central access point. The reason mentioned this is that this is how most real IPSec deployments work these days. There is much more traffic carried through permanent IPSec VPNs than through per-session IPSec.

If you have administrative control of the clients, so you can compel everyone to use the same client software, you have solved half the problem. The trouble is that the other half of the problem is very ugly.

It's not that IPSec VPNs are the trouble, it's the key exchange protocol needed to start the session. When IPSec VPNs were standardized, it was assumed that clients would be connected directly to the public internet, without firewalls.

Key exchange requires outbound and inbound connections on privileged ports, something that many firewalls block. This is similar to the problem of FTP through firewalls, except that there is no analog of “passive ftp” which avoids the problem. (It’s the inbound connection to the privileged port that’s the issue, in this case.)

If network address translation (NAT) is used, the outbound connection gets remapped from its privileged port to an unprivileged port and the key exchange fails. NAT is a real killer for many IPSec VPNs. You can do port redirection to avoid this——some home firewalls have this feature so you can use your IPSec VPN to access the office——but they have other limitations, for example only allowing one client at a time to use the VPN.

Because of these issues, many IPSec VPNs, even those that originally planned on using per-session authentication, ended up being set up as permanent tunnels. IPSec VPNs with per-session authentication are mostly used for people with no client side firewalling, (home offices or other public net access). In university and government agencies, which have already deployed firewalls and set policies for security, getting a per-session IPSec VPN working can be a real pain. (The pain is usually more political than technical.)

SSL VPNs are attractive because web access is considered a “mission critical” function, so we can almost always get to TCP port 80 (http). There is no need to deploy software since all of the major web browsers support SSL (you’re using SSL when you go to a site whose URL begins with https:).

There is a legitimate concern that not every application is SSL aware and for some older applications might require a complicated workaround. This is likely to be true. On the other hand, the SSL VPN makes the simple stuff simple, so it is likely to give you most of what you want sooner, as opposed to IPSec VPNs, which will give you everything you want eventually. By simple stuff I mean basic telnet/ssh access and file transfer by ftp, SMB (Windows network disks) and NFS.

Let me close with a piece of advice: if Neoteris (the market leader in the SSL VPN space; Nortel is number 2 and Cisco a distant 3rd) goes public, buy. It’s an obvious acquisition candidate. Usual disclaimers apply.

As you can read, my experience in the private sector inclines me to the practical.

Best Wishes, Greg

Hi Matt and all,

I have come up with some additional information on VPNs that might help you decide what to do.

Most of the SSL VPNs being deployed are going to financial institutions and government. In particular, the Alteon 310 SSL box is meets FIPS (Federal Information Processing

Standards) for security, and can therefore be used to secure medical records under the Health Insurance Privacy Protection Act (HIPPA). Deutsche Bank has been rolling out an SSL VPN even though they have an already existing IPsec VPN system. They are using it to handle both web enabled and legacy applications.

Matt makes a good point about security, but the current SSL stacks no longer vulnerable to the well publicized man-in-the-middle timing attack. The main concern, both for SSL and IPsec VPNs, is client software containing backdoors and trojans. The only way to protect against this is to ensure that everyone gets their client software from a trusted source. There are some add-on products that do a sanity check of clients, looking for obvious compromises. Deutsche Bank uses one of these on their SSL VPN.

The particular attack that Matt mentioned in his note, webmitm, targets a flaw specific to Microsoft Internet Explorer. This web browser will accept a forged authentication. As I mentioned above, for either scheme security depends on 1) the host not being compromised and 2) the VPN client software not being compromised. This vulnerability falls into the second category. In reality, these kinds of problems are common to both types of systems. (Also bear in mind the remarks attributed to a staff member at NSA, who claimed that it was never necessary to break the cryptosystems to get access to computer data. There is always another, easier way.) Overall system security really relies much more on setting a reasonable set of policies that everyone can live with. This minimizes the temptation to create exceptions which introduce security holes.

Since I'm here at IETF, I'll make a remark on where IPsec VPNs are going, specifically will they get any easier to deploy? The answer is "sort of". The focus on IPsec VPNs is not on user provided VPNs (the kind you are considering, where the organization using the VPN also has administrative responsibility for it) but carrier provided VPNs. In a carrier provided VPN, you just buy VPN service from AT&T, MCI or whomever. The carrier is responsible for setting everything up. A standard is in the works to improve the manageability of key exchange (target at carriers), called Group Domain of Interpretation (GDOI). It is maybe a year away from completion. The user-provided IPsec VPNs are unlikely to change. They won't get easier to deploy, but they won't get harder.

If you're waiting for IPv6 to solve all of these problems, you're going to be waiting for a long time. The IPv6 standard may be done, but the standards which will guide the conversion process are still very contentious and incomplete.

As always, let me know if you have any additional questions or if I can provide other information.

Best Wishes, Greg Wright

Greetings from [Matt],

Hey, no problem. Again, I think the decision on implementation details is best left to the people actually doing the implementation (RPSC). We can say what functionality we

want, what platforms must be supported, and present a few example cases of organizations that provide the services we want securely. Greg have you got some good examples of research insitutions using SSL based VPN's? I gave Bob the example of NASA GSFC using per-session authentication for setting up an IPSec tunnel.

With exception to the web interface you can get complete network connectivity either way with some client software. It'd be nice to have for times when away from a properly configured machine, but is not in any way viable as a single solution for me. Eventually I've got to be able to use any application to connect to any science machine associated with me at pole.

What platforms should be on the list of required platforms for any client software? I'd add solaris, linux (intel), Mac OSx, Windows (all flavors). Are there any others that need to be supported for any possible client software?

Hey, if you are going to the IETF meetings go beat some sense into the IKE committee and have fun in Vienna. If you get some time, I haven't done much with an SSL based VPNs, but how vulnerable are they to man-in-the-middle attacks? I mean the SSL based webmail programs are rediculously easy to attack. For some impressively evil utilities check out the dsniff suite at: <http://www.monkey.org/~dugsong/dsniff/> In particular the webmitm utility.

Cheers,

Matt Newcomb