

# VPN Backup: Redundancy for WAN Circuits Using IPSec VPNs

By Kevin Holmes, CISSP, CCNP, CCDP

Many large organizations recognize the need for both router and circuit redundancy at the core of their network, but very few organizations have WAN redundancy at their small to medium sized sites. Even though these smaller remote sites may also need access to critical applications, the cost of providing redundancy via a second dedicated WAN circuit to smaller remote sites is rarely cost-justifiable. One of the frustrating ironies of being a Network Engineer may be that the very same executive or manager who denied your budget request to deploy WAN redundancy may some day scream "We've got to get this site back up ASAP! We're losing thousands of dollars per hour while it's down!" VPN backups over inexpensive broadband circuits may be the panacea that Network Managers are looking for.

How many different ways are there for your organization to lose WAN connectivity between remote sites, thereby losing access to critical applications? Let me count the ways.... Besides the obvious issues with the potential failure of networking hardware (router, CSU/DSU, etc.), the loss of a WAN circuit, or potential electrical power failures, there is also the possibility of human error. I have personally experienced outages due to indiscriminate electricians in the server room, telecommunications technicians that mistakenly thought their tests would be "non-intrusive", and overburdened network engineers that had too many telnet sessions going and issued a reload command to the wrong router.

## ISDN—It's Still Deployed Nowadays?

For the past several years, most networking professionals have been quick to proclaim ISDN technology to be dead, or perhaps archaic at the very least. However, ISDN circuits are still prevalent in many businesses, as it is a surprisingly common method of providing a backup path for business data when a circuit or router fails. Many of these same organizations have increased the size of their WAN circuits through the years as their traffic load has grown, but have not increased their corresponding ISDN backup bandwidth. A severely undersized circuit (regardless of the type of circuit technology) is virtually useless as a backup path during a WAN outage.

There are two primary problems with ISDN as a method of providing a failover path for remote site connectivity. Neither issue is related to the inherent functionality of the technology itself. ISDN works quite well for the right application.

The first issue is that not only do you pay monthly recurring costs to have an ISDN circuit available, but typically there are also per-minute charges for using ISDN when you need it. Much like a taxi, ISDN service is available when you need it, but quickly becomes quite expensive when "the meter begins to turn."

The second issue is that the amount of bandwidth ISDN offers is limited. ISDN BRI services only offer 128 Kbps of bandwidth per circuit. If you need more bandwidth than a single 128 Kbps circuit, yet less than a T1/E1 PRI circuit at approximately 1.5 or 2.0 Mbps, then you must "bond" multiple ISDN BRI circuits together using Multilink PPP protocol. Multiple ISDN BRI channels will also require additional networking hardware; either multiple ISDN interfaces on a router or an external ISDN Inverse Multiplexer. If your organization needs more bandwidth than can be delivered in a single T1 or E1 service, then your bandwidth issues become even more challenging and inflexible.

## Broadband Internet Access—The Last Nail In ISDN's Coffin?

In fairness to ISDN, broadband Internet services may not be able to match the minimal end-to-end delay qualities of an ISDN circuit, which is effectively a temporary, dedicated circuit. If an organization is transporting delay sensitive traffic such as voice or video, they should consider service providers that can offer Service Level Agreements (SLAs). SLAs can guarantee delay thresholds, availability levels, and minimum throughput.

Typically, the larger carriers are better poised to offer SLAs, as they can maintain control over quality of service if the organization's packets do not have leave their own BGP autonomous systems. However, this may require you to use one carrier for all of your Internet service, which is not always geographically feasible. In addition, the SLAs may increase your cost, minimizing or even eliminating any cost advantage an organization might derive from moving to broadband Internet service.

How does a business budget for a service that they hope they will never use, but anticipate that they may need to use? How frequently will your dedicated WAN circuits or your routers fail, thereby prompting usage of the ISDN circuit and incurring additional costs? And how long will your ISDN circuits be in use if you encounter such a failure? Consider also that within each circuit are multiple 64K "B Channels" (2 B channels for ISDN BRI circuits, 23 B channels for T1 ISDN PRI circuits, and 30 B channels for ISDN E1 circuits), and the telecom carriers charge per each B channel that is used. For example, in order to backup a heavily utilized E1 circuit, you may be charged for as many as 30 ongoing B channel calls.

Unlike ISDN circuits, broadband Internet access is usually provided at a fixed monthly fee, regardless of usage. It is readily available and reasonably priced throughout much of the world, and is often available through multiple service providers in a given area. So, why not utilize these lower-cost broadband Internet services as your method of providing data backup through a VPN?

## Replacing ISDN With IPSec

For those who are unfamiliar with the term, a VPN is a Virtual Private Network. A VPN is "virtually private" because although the data is passed across a public network, it is first encoded, thereby obscuring sensitive data from the view of others. Even if a hacker is able to copy your transmissions, decoding your encoded data without being the intended recipient is roughly akin to trying to unscramble an egg and return the yolk and the white intact back into an unbroken shell.

The most common method for obscuring your data from the view of nefarious hackers and crackers prowling the Internet is through IPSec. IPSec is a suite of related protocols and algorithms that provide data encryption, authentication, and key management (IKE). IPSec can guaranty the confidentiality and integrity of your data, ensuring that it has not been revealed to any entity other than the intended sender/recipient pair, and that the data has not been altered in any way throughout the transmission. It can also provide proper authentication of the two entities that are communicating.

In addition to confidentiality, authentication and integrity, IPSec can also provide anti-replay protection. As an example of how this feature could be beneficial, anti-replay protection could preclude someone from capturing a data session containing a transfer of bank funds, and resubmitting it to a database until the account was depleted.

Two IPSec clients form what is typically referred to as a "tunnel". A tunnel is created by two endpoints that form a peer relationship even though they are not on neighboring networks. Tunnels are most often created either for security reasons (to obscure the data contained within the data packets being carried), or for protocol compatibility issues (e.g., you need to pass AppleTalk packets across an IP-only WAN).

IPSec can utilize various encryption methods. These various methods offer differing levels of encryption "strength" (prevention from being unencrypted by anyone other than the intended recipient). The stronger methods of encryption provide greater defense against attacks. However, these stronger encryption methods are usually more computationally complex, and require more processing "horsepower" from the VPN hardware devices.

If an organization can not afford the risk of exposing your data to a hacker, your competition, or the public at large, then it should choose the strongest encryption available. An example of strong encryption would be 3DES (typically pronounced triple dez) as opposed to DES (Data Encryption Standard). 3DES provides 168-bit encryption as opposed to only 56-bit encryption for DES.

Vendor hardware will typically offer many different options, and vendors may or may not charge more for utilizing 3DES software than for DES. In addition, you must consider that utilizing 3DES over DES will consume many more CPU processing cycles to encrypt and decrypt data packets. For this reason, the device serving as the IPSec tunnel endpoint can not handle as many 3DES encrypted packets as it could DES encrypted packets. The number of IPSec tunnels and the number of encrypted packets sent across these tunnels are critical factors in determining what type of VPN hardware will be utilized. Choosing stronger encryption requires more powerful VPN hardware, which is, of course, more expensive.

An even better option than 3DES encryption would be AES (Advanced Encryption Standard), assuming AES is fully supported by your hardware provider. AES (FIPS-197) is a standard of the NIST (National Institute of Standards and Technology), and can provide 128, 192 or 256 bit encryption. AES provides the unique benefit of being both faster and stronger, making it likely to be the prominent choice in the future as this newer technology becomes more widely supported.

In the NIST AES Question and Answer page ([http://www.nist.gov/public\\_affairs/releases/aesq&a.htm](http://www.nist.gov/public_affairs/releases/aesq&a.htm)), the NIST states the following in reference to the inherent defensive strength of AES:

Question: "What is the chance that someone could use the "DES Cracker"-like hardware to crack an AES key?"

Answer: "In the late 1990s, specialized "DES Cracker" machines were built that could recover a DES key after a few hours. In other words, by trying possible key values, the hardware could determine which key was used to encrypt a message. Assuming that one could build a machine that could recover a DES key in a second (i.e., try 255 keys per second), then it would take that machine approximately 149 thousand-billion (149 trillion) years to crack a 128-bit AES key. To put that into perspective, the universe is believed to be less than 20 billion years old."

That ought to fend off today's script-kiddies until they are dust and then some... and that was at 128 bit encryption, as opposed to the available 256 bit option. Beware though, because as computational horsepower increases throughout the coming years, this figure could potentially be reduced to less than a hundred trillion years...

Although there is typically not a differential in costs for the various encryption options, you need to be aware that the U.S Government restricts where you may utilize encryption processes above 56 bit, and you must be licensed accordingly to deploy stronger encryption software. Your hardware vendor should be able to provide information on the legality and availability of strong encryption software for your application and your intended country of use.

The choice between DES, 3DES and AES essentially determines how data will be encrypted. IPSec uses one of two protocols for the authentication and encapsulation of what will be encrypted. The two protocols are AH (Authentication Header) and ESP (Encapsulating Security Payload). This creates another set of options to determine which of these four possible combinations you will need to use. However, the choices are easier than you might initially assume.

AH only provides authentication, but not confidentiality. It does not obscure your data from other users on the Internet. Consequently, we can obviously eliminate utilizing AH on the Internet altogether. By contrast, ESP always encapsulates and encrypts your data payload from the view of others. In addition, ESP can allow modifications to the IP header, thereby supporting NAT (Network Address Translation).

Both AH and ESP must be run in either transport mode or tunnel mode. AH or ESP will determine what will be done to the packet (authentication or both encryption and authentication), but transport mode or tunnel mode will determine how much of the packet will be authenticated or encrypted/authenticated. Transport mode will apply the specified protocol (AH or ESP) only to the payload of the IP packet. Tunnel mode will apply the protocol to both the payload and the original IP header before appending a new header to the encrypted packet.

## Potential Pitfalls

Although IPSec and NAT can be made to function in conjunction with each other, be forewarned they do not always do not always "play nicely together within the same sandbox". Consider that the purposes of these two protocols are diametrically opposed. NAT intentionally modifies your IP packets to make the packets appear as though they originated from a different source. In contrast, IPSec provides the capability to maintain the integrity of IP packets, insuring that they have not been modified along the path toward their destination.

Typically ESP in tunnel mode can survive the NAT processes of your firewall, but your particular implementation methods should be tested and verified before you set delivery time frames for a VPN backup deployment. The deployment of IPSec could conflict with your firewall's existing method of deploying NAT or NAPT (Network Address Port Translation, a.k.a. PAT). Another potential problem is that deploying IPSec-based VPN Backup could even violate your organization's security policies.

Where the Network Address Translation occurs relative to where the IPSec tunnel endpoints (peers) are located is of critical importance, and could easily circumvent any potential conflicts between IPSec and NAT. If you are utilizing VPN devices that translate addresses prior to IPSec encryption in the outbound direction, and decrypt packets prior to NAT translations on the inbound direction, NAT will not be an issue for site-to-site IPSec tunnels. This approach is common for site-to-site VPN tunnels, and because the NAT translation occurs before packets are forwarded into an IPSec tunnel, it eliminates any potential conflicts between NAT and IPSec (see figure 1).

IPSec is a complex subject with a wide variety of implementation options, and substantially exceeds the scope of this article. If you wish to familiarize yourself further with IPSec, a good choice would be to read the RFCs pertaining to IPSec and the various related cryptographic and security processes and procedures. Not only are these documents the definitive source of information for these international standards, but they are also priced appropriately for even the most frugal among us, as they are available for free. RFC 2401 (Security Architecture for the Internet Protocol) would be the critical document to begin with, as it deals with IPSec in general.

Other pertinent RFCs would depend upon your particular configuration options that you will utilize, such as whether or not to use SHA-1 (Secure Hash Algorithm) or SHA-2 as opposed to MD5 (Message Digest 5) as your hashing algorithm for HMAC message authentication. RFCs can be located at the following URL: <http://www.rfc-editor.org/cgi-bin/rfcsearch.pl>. For the very brave, you may wish to delve immediately into the manuals from your hardware vendor.

In addition to eliminating the need for ISDN circuits for backup use, local Internet access can also contribute to lowering your WAN circuit costs by diverting remote site Internet traffic off of the high cost long-distance WAN circuits and onto low cost local Internet circuits.

Quite often Internet traffic may constitute twenty-five percent or more of all of the data traffic being carried from remote sites back to core sites where Internet access is more likely to exist. For multi-national corporations, reducing the load on expensive international WAN circuit costs can provide tremendous savings where low-cost local Internet services are available at the remote sites. As an example, if you have an \$8,500 per month WAN circuit, and you are able to divert twenty-five percent of your traffic over a \$50 per month ADSL circuit, you may effectively save over \$2,000 per month.

If your security policies require that you maintain control over the content of Internet traffic being viewed by your user community, you can either install a proxy server at the remote site, or utilize content filtering control programs such as Websense or N2H2 (among others) on the local firewall.

Consider using symmetrical circuits such as SDSL when VPN backup is the primary business driver for purchasing local Internet access. Although asymmetrical services may be suitable for standard Internet services where traffic flows are typically much heavier in the download direction, this type of circuit may not be as suitable for use as a backup circuit where traffic flows are likely to be more bidirectionally balanced. However, don't let the

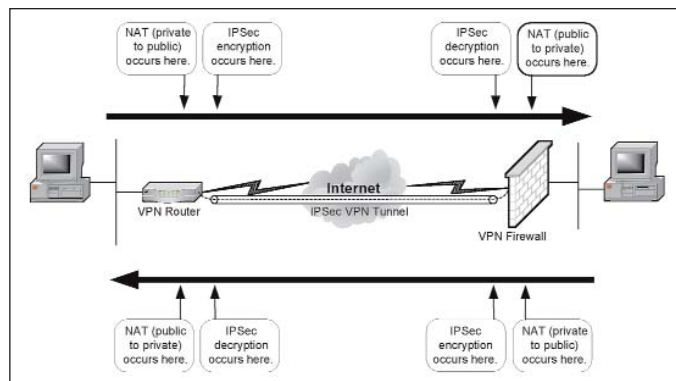


Figure 1: Site to site VPN Tunnels

name ADSL (Asymmetrical DSL) fool you. Carriers often provide ADSL in bi-directionally balanced increments such as 384K/384K, 784K/784K or 1.5Mbps/1.5Mbps. If your intent is to backup a T1 or E1 circuit, consider that T1 and E1 circuits are full duplex, and can carry approximately 1.5 Mbps (T1) or 2 Mbps (E1) in both the upload and the download direction at the same time. Heavily asymmetrical ADSL circuits will not be able to upload traffic in a manner consistent with the T1 or E1 circuit it is intended to back up.

## Routing Issues

Routing traffic correctly can become challenging when establishing a backup path for your network traffic. The goal is to make your traffic take the preferred route when it is available, and only take the backup path when you need traffic to failover to the only available path. Your routing metrics must be influenced to make the backup path less preferable to your primary circuit, although the two paths may appear to be the same number of router hops away. Tunneling may make the number of router hops much lower than the actual number of hops the VPN tunnel may take across the Internet. Dynamic routing protocols such as OSPF, IS-IS and Cisco's proprietary EIGRP have methods available for altering the cost or metric of one path over another.

Routing summaries can also provide a way of influencing routing paths. I once deployed a VPN backup for a new site located in Dallas. We utilized an Interworking service (Frame Relay to ATM) between Dallas and Houston as their primary (preferred) path, and an IPSec VPN backup between Dallas and Southern California. We made the preferred path look preferable by the use of multiple, longer-match route summaries as opposed to the VPN, which advertised only one individual, shorter-match route summary. The longest match rule of routing instructs a router to choose the route with the longest bit mask. For example, if a router knows of both a 22 bit mask and a 24 bit mask to a given site, it will choose the 24 bit mask because it is more specific.

However, I was initially quite puzzled to find that most of the traffic to the Dallas site was using the VPN backup path as opposed to the intended Interworking circuit path. However, this was not the case for all of the sites. The Houston site and all of the regional sites that fed into Houston would take the Interworking path as intended. However, all other global sites seemed to prefer to traverse the VPN backup path between Dallas and Southern California.

Upon further exploration, I discovered that although route summarization was the solution on one interface of the core router, it was causing a problem on another "upstream" interface. The Houston site advertised only one large routing summary to its neighboring routers,

and did not allow the more specific route summaries (longer route matches) over the preferred path to propagate outward. Because this very large route summary at Houston precluded the advertisement of the preferred Dallas path to other core sites, the VPN was being utilized because its advertised route was more specific than the large summary advertised at Houston.

The solution was to reduce the size of the summaries between core sites, but leaving out the summary range from the Dallas site. This allowed the preferred path via the Dallas Interworking circuit to be advertised throughout the network when it was available, and to disappear when there was a circuit or router failure.

The basic rule is to beware of any summaries between the router with the dedicated circuit and the router handling the VPN tunnels. Before you deploy a VPN Backup solution in a network utilizing route summarization, try to analyze the impact that these summaries may have at from the perspective of various sites. Once the solution is initially deployed, be sure and test via traceroute to ensure that your traffic is taking the intended path.

If you use firewalls to create your IPSec VPN tunnels, you will most likely use static routing, as most firewalls do not support dynamic routing. Be aware that once a static route has been inserted, it may never disappear, even though the device it points to is no longer available and can not forward the traffic that is being directed to it. If one path is over a router and the backup path is over a firewall, the router may be able to advertise a preferable path dynamically over the router. This way, the static route pointing to the firewall will only be chosen when the dynamic route has disappeared due to a circuit or a router failure. Routers with firewall functionality are often ideally suited for VPN backups, as they can offer virtually instantaneous dynamic re-routing of data traffic upon the loss of the primary circuit.


## IPSec and Multicasts

This leads us to yet another "gotcha"... dynamic IGP routing protocols such as OSPF or EIGRP use multicast addresses as their destinations for router updates, and IPSec only carries unicast traffic. You can utilize a GRE tunnel (Generic Router Encapsulation, RFC 2784, <ftp://ftp.rfc-editor.org/in-notes/pdfrfc/rfc2784.txt.pdf>) to tunnel all of your traffic (including multicast-based traffic such as routing updates and video traffic) through the IPSec tunnel. Yes, it does sound convoluted to create a GRE tunnel and pass it through your IPSec tunnel, but it works, and it resolves the issue of IPSec not supporting multicasts and broadcasts.

Be aware that it does add even more overhead however, somewhat diminishing your throughput. It should also be noted that a GRE tunnel varies from an IPSec tunnel in that GRE does not automatically encrypt the traffic it carries. However, with this technique (GRE inside of IPSec), the traffic inside the GRE tunnel will ultimately be encrypted by IPSec.

You will need to review your VPN Backup design to make sure you eliminate any single points of failure. For example, it makes no sense to have both your primary WAN circuit and your VPN tunnel terminate on the same router, because you would obviously not have either path available if the router itself fails. Another benefit to using a router with firewall functionality is that you can utilize VRRP (Virtual Router Redundancy Protocol, <ftp://ftp.rfc-editor.org/in-notes/pdfrfc/rfc2338.txt.pdf>) or HSRP (Hot Standby Routing Protocol, <ftp://ftp.rfc-editor.org/in-notes/pdfrfc/rfc2281.txt.pdf>) for path redundancy on either or both sides of your link.

Without drilling into the differences between VRRP and HSRP, both protocols allow other devices to send traffic to a virtual gateway IP address that is shared by both the primary router (Master) and the sec-




**Figure 2: Virtual gateway address shared by VRRP/ASRP routers**

ondary router (Backup). When the primary router is functional, it will assume the responsibility for forwarding traffic. If the primary router fails for any reason, the secondary router will then assume the responsibility of forwarding packets. This capability is especially helpful for host devices that do not effectively support multiple gateway addresses. Figure 2 depicts the concept of a virtual gateway address shared by both VRRP/HSRP routers.

Test the VPN backup functionality before declaring it to be operational. Just because local Internet traffic is working does not mean that VPN backup functionality has been verified. These are two different functions with different rules and configuration requirements. Look for more than just whether or not your pings and traceroutes are taking the intended path for a given scenario. Test your applications. Due to the additional overhead consumed by IPSec or even worse, the combination of IPSec and GRE, data packets with the IP "Don't Fragment Bit" set may work over the dedicated circuit but fail over the VPN Backup.

## Conclusion

Using VPN backup represents a compromise for those that do not feel that the Internet is reliable enough to utilize as the primary WAN medium for site-to-site connectivity, but need to increase their network resiliency without obliterating their annual WAN budget. As converged data networks that support IP Telephony and multicast video become more common, the need for network resiliency will become critical. VPNs that carry delay-sensitive traffic such as IP voice or video should consider choosing IPSs that can provide stringent service level agreements (SLAs). With your traditional business applications, email, messaging and even voice (telephony) and video technology relying upon your data network for communications, few businesses will be able to afford to run without backup paths to their remote sites. VPNs with IPSec will meet this need. 

---

*Kevin Holmes has worked in information technology for the past nineteen years, focusing on networking and security. He received his Bachelor's degree from California State University at Long Beach in 1984. He is currently a Network Engineer with IBM Global Services, assigned to support Fluor Daniel, a Global 500 company. He has also worked for Netigy, EDS (Chevron account) and GTE (now Verizon). Kevin currently holds CISSP, CCNP, and CCDP certifications, and is a member of the Information Systems Security Association. He can be reached for comment at [kevin.holmes4@verizon.net](mailto:kevin.holmes4@verizon.net).*