

Stefano Giacometti and Roberto Mameli

Tunnelling Effectiveness in the Access Environment

The current trend in the Internet access market is moving towards the SOHO (Small Office Home Office) target. Many analysts agree that the access market will continue to grow at a rate of approximately 30% to 40% annually through the year 2000. On the other hand, the availability of different access network infrastructures (often specialised for different kinds of media) leads to a full IP access solution able to overcome such a heterogeneity. This aspect is also justified by the ongoing convergence between telecom and datacom worlds, in which video, audio and data communications are delivered to the user in an integrated solution. In such a scenario the need for advanced access techniques arises.


This document proposes the use of tunnelling techniques in the context of a full IP access network. Tunnel mechanisms, which simulate point-to-point connections over connectionless networks, have been typically developed for applications based on virtual private networks (VPNs). IP tunnels solve traditional visibility problems that arise when we attempt to address a host within a private IP domain from a public one. Moreover the flexibility and the advanced security mechanisms offered to protect user data justify the introduction of the tunnel technology in the access domain.

Introduction

The telecommunication and networking worlds are going to be characterised by the 'convergence' process taking place during the last few years between telecommunications and datacommunications. Several technological and commercial factors justify this convergence. Historically telecommunications have been exploiting analogue connectivity, while data communications and networking have always been based on digital techniques. The development of powerful and cost-effective digital signal processors (DSPs) has

decreased the distance between these worlds: many analysts consider them as one of the main reasons that will lead from circuit switching to packet switching. This 'paradigm shift' will probably take place in the medium term and it will allow bandwidth efficiency to be increased up to a magnitude order. Internet protocol (IP) telephony is certainly one of the areas in which this phenomenon is rapidly growing. By combining the best features of real-time voice communications and data processing, new kinds of applications can be foreseen. Figure 1 demonstrates this

Figure 1—Network capacity demand over the years



phenomenon: it depicts the redistribution of the total network capacity demand among different services.

In such a scenario the access environment plays a relevant role. Currently, the Internet access method for residential and SOHO (small office home office) users is based on the switched telephone network and exploits the dial-up mechanism. Users connect to an Internet service provider (ISP) by means of an analogue modem or a digital interface for integrated services digital network (ISDN) lines; when a connection is established an IP address is negotiated and, only after this phase, IP packets can be exchanged by the point-to-point protocol (PPP). Note that before address negotiation the user PC does not have an IP address at all; that is, it does not belong to an IP domain, neither public nor private. However, following the previously outlined convergence process we can easily imagine a medium-term scenario in which the IP connectivity will start directly from the users' premises, to provide integrated access to a set of services, both involving real-time and non-real-time applications.

In the previously described full-IP-access scenario some problems arise, such as those typically related to visibility, security, transparency to applications and so on. The use of proper access methods can solve all of them. A possible solution is given by tunnelling: basically these techniques are used to transfer data from one network to another one by means of an internetwork infrastructure, eventually using different protocols. The data to be transferred are typically contained in a payload packet (or frame) that is encapsulated within a transport packet and sent through the carrier network. The additional header introduced by the encapsulation mechanism carries routing information to let the

Stefano Giacometti:

Ericsson Telecomunicazioni S.p.A.
Tel. +39-06-20410028
Fax +39-06-20410037
stefano.giacometti@ericsson.com

Roberto Mameli:

CoRiTeL
Tel. +39-06-20410038
Fax +39-06-20410037
mameli@coritel.it

payload packet traverse the transit internetwork. Once the payload packet has been encapsulated it can be routed between tunnel endpoints over the internetwork. The tunnel itself consists of the logical path through which the encapsulated packets travel. When the encapsulated packet reaches its destination on the internetwork, the frame is unencapsulated and forwarded to its final target. Note however that the intermediate internetwork can be any internetwork: in the context of virtual private networks, the latter usually coincides with the Internet, but this is not mandatory, since a different use of tunnelling could also imply a private transit internetwork.

This paper is organised as follows. The next section provides a technical overview on tunnelling techniques. Then the use of tunnels in the access network is explained, together with the advantages and disadvantages of this approach. Finally, the main points of interest are summarised and conclusions drawn.

Technical Overview of Tunnelling

The idea of tunnelling is not recent, since some examples have already been used in the past (for example, SNA (System Network Architecture) tunnels or IPX (Internetwork Packet eXchange) tunnels for Novell NetWare, both over IP internetworks). But recently, especially due to the interest in emerging scenarios (for example, virtual private networks (VPNs) or mobile IP environments) some new tunnelling technologies have been introduced. These newer technologies include:

- point-to-point tunnelling protocol (PPTP),
- layer 2 tunnelling protocol (L2TP), and
- IP security (IPSec) tunnel mode.

VPNs and IP mobility are two emerging applications of tunnelling.

The basic idea behind VPNs is to use tunnelling in order to cross transparently a transport infrastructure (usually the Internet), making it possible for a remote user to reach its private network (for example, a corporate Intranet). In this way the remote user appears to belong to the private network domain and all the features concerning visibility, authentication and security, typically needed in a VPN scenario, are supported by the tunnelling technique. The mobile IP constitutes another application of tunnelling; in this case tunnelling is used between the foreign agent and the home agent. The former is the router in the visited network that manages visitor hosts, while the latter is the router in the home network that is aware of the current position of the mobile host. When the home agent receives data directed to the mobile host it forwards them through a tunnel to the foreign agent, which in turn delivers them to the target host. In this way the mobile host can keep its IP address (public and static at the same time) while travelling around the world.

Given its importance, it is worth providing a more detailed explanation of the tunnel mechanism from a technical point of view. Tunnels are based on the classical client-server paradigm: both the client and the server communicate by means of a tunnelling protocol, which can be either a layer 2 or a layer 3 protocol. Layer 2 protocols, such as PPTP and L2TP, encapsulate a layer 2 data unit (for example, a PPP frame) inside the transport packet of the carrier internetwork. In contrast, layer 3 protocols, like IPSec, encapsulate layer 3 data units (for example, IP packets) in an additional IP header before sending them across the IP internetwork.

A further distinction can be made between voluntary and compulsory tunnels. In the former case one of the tunnel endpoints coincides with the

end user, that is, the tunnel client is located in the user's equipment, and the virtual connection is set up on demand (see Figure 2). In the latter case, the tunnel endpoint is physically distinct from the remote user equipment; thus at first a connection to the tunnel client must be properly set up and only then data coming from the user can be tunnelled through the internetwork towards the tunnel server. In this case, the tunnel client typically allows different users to share the same tunnel; that is, there are no separate tunnels for each user as in the case of voluntary tunnels (see Figure 3).

In spite of the previously outlined differences, the basic mechanism is always the same. For example, when the tunnel client has to send a payload to the tunnel server it appends a tunnel data transfer protocol header to the payload. Then it sends the resulting encapsulated payload across the carrier network, which routes it to the tunnel server. The tunnel server accepts the packets, removes the tunnel data transfer protocol header, and forwards the payload to the target network. In this way the tunnel client becomes visible in the server domain; that is, it behaves as a node physically located within that environment. Note that this process is completely transparent to applications, which are not aware of the underlying tunnelling mechanism. Tunnelling also involves some other aspects: among them a relevant role is certainly played by security issues. These include not only aspects concerning user authentication (to avoid unauthorised access to a private Intranet), but also problems related to data encryption (in order to guarantee a high level of confidentiality to transactions).

One of the main disadvantages of tunnelling is certainly represented by the introduction of some amount of header overhead. This is mainly due to the encapsulation mechanism that can be better understood by looking

Figure 2—Voluntary tunnel





Figure 3—Compulsory tunnel



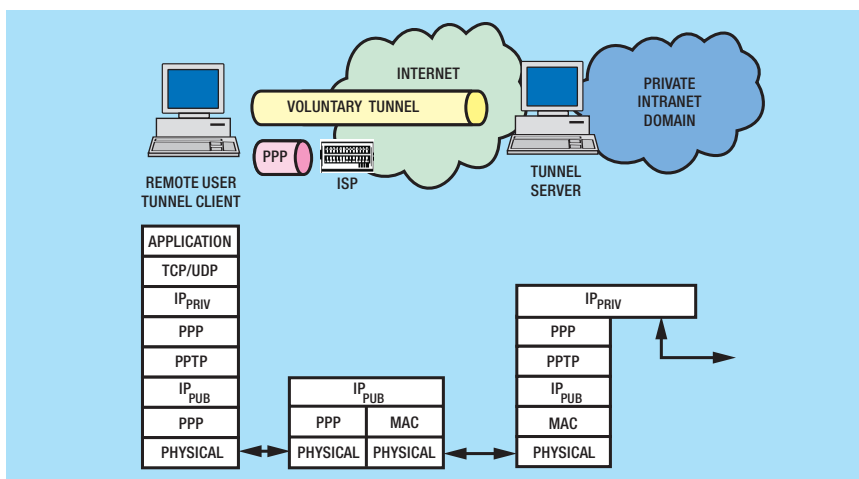


Figure 4—Access of a remote user to a private Intranet domain

at the protocol stack increase. In Figure 4, a typical scenario is represented, in which a remote user connects to its intranet by using a dial-up connection. At first the user sets up a PPP connection, by which it acquires a public IP address (that is, an IP address with a worldwide scope). After this it opens a voluntary tunnel to the tunnel server, by which it acquires a new IP address, whose scope is now limited to the private domain. From now on it is virtually located in the private domain and its packets are encapsulated (at the client side) and unencapsulated (at the server side) in order to be routed through the Internet without problems.

The tunnelling mechanism obviously involves several topics. The main ones are listed below along with a brief explanation. It is worth noticing that layer 2 protocols based on PPP usually inherit some of its mechanisms; for example, for authentication or address assignment.

- **Dynamic address assignment**
This is the first and foremost problem to be solved. In order to make the tunnel client visible in the server's domain, the tunnel server must assign an address of its domain to the client. Usually layer 2 tunnelling protocols support dynamic assignment of client addresses based on the network control protocol (NCP) negotiation mechanism of PPP. In contrast, layer 3 tunnelling schemes assume that an address has already been assigned prior to initiation of the tunnel. Schemes for assignment of addresses in the IPsec tunnel mode are currently under development and are not yet available.

- **User authentication** As previously stated this is one of the main security concerns. The approach followed by layer 2 tunnelling protocols usually relies on the user authentication schemes of PPP (such as password authentication protocol and challenge handshake authentication protocol), while layer 3 tunnelling schemes generally assume that the endpoints authenticate themselves reciprocally before tunnel establishment. IPsec represents an exception, since it provides authentication during tunnel set up (by means of ISAKMP—Internet security association and key management protocol)
- **Data encryption** This problem constitutes another relevant security topic. As before, layer 2 tunnelling protocols exploit PPP-based data encryption mechanisms. For example, Microsoft PPTP implementation uses Microsoft Point-to-Point Encryption (MPPE), based on the RSA/RC4 algorithm. In contrast, layer 3 tunnelling protocols provide themselves proper encryption mechanisms. For example, IPsec defines several optional data encryption methods, which are negotiated during the preliminary ISAKMP exchange. L2TP can address this problem in two different ways: the former relies on PPP mechanisms, while the latter uses IPsec encryption to protect the data stream from the client to the tunnel server.
- **Key management** This is strongly related to the previous

point. Layer 2 protocols generally use an initial key generated during user authentication and then refresh it periodically (an example of such a mechanism is given by the previously mentioned MPPE). Similarly IPsec explicitly negotiates a common key during the ISAKMP exchange and refreshes it periodically.

- **Data compression** Many tunnelling techniques also support data compression mechanisms. As before layer 2 tunnelling protocols are PPP-based; for example, the Microsoft implementations of both PPTP and L2TP use Microsoft Point-to-Point Compression (MPPC). The Internet Engineering Task Force (IETF) is investigating similar mechanisms (such as IP compression) for the Layer 3 tunnelling protocols.
- **Multi-protocol support** Another relevant topic is constituted by the possibility of supporting multiple payload protocols. This is a peculiar characteristic of layer 2 protocols, that are able to support various encapsulated protocols, such as IP, IPX, NetBEUI and so on. In contrast, layer 3 tunnelling protocols, such as IPsec, typically support only target networks that use IP.

Tunnelling in the Access Network

As previously outlined, the tunnel technique is normally used in order to support VPNs and IP mobility. However, the new emerging needs motivated by the convergence between telecommunications and datacommunications worlds justify its adoption in the access section of the network. The scenario we refer to is the one depicted in the Introduction. In such a situation there is an IP access provider (IAP) that allows full IP connectivity to residential and SOHO users. Note however that the access provider can be administratively distinct from the Internet service provider (ISP) and, moreover, it usually supports different services. In fact the former provides exclusively physical and IP connectivity, while the latter grants primarily worldwide visibility (that is, the possibility for a host located in a private domain to interact with other parties located anywhere in the world). Other types of service can also be provided by the ISP, such as

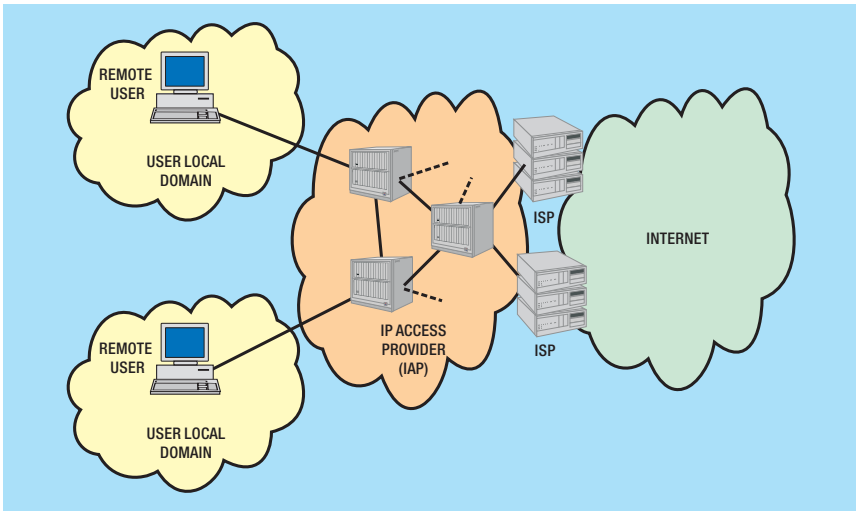


Figure 5—Full IP access scenario

mail, proxies or newsgroup services; the choice among different ISPs may also be supported by the IAP, making it possible for the user to select the one more suitable for his purposes (see Figure 5).

Consequently, the access environment can be administered by a private entity with only a small set of available public IP addresses. To allow the growth of the access network the administrator may take advantage of the statistical multiplexing of IP addresses allowed by the dynamic address negotiation of the tunnelling mechanism. This is one of the several advantages in the use of tunnelling in a full IP access scenario; the main ones are explained in more detail below.

- **Statistical multiplexing and flexibility in the address re-use**
The user (either a residential user or a small local area network (LAN), as in the case of SOHO environment) would be able to set up a tunnel on demand to access the Internet whenever necessary; for example, to make an IP telephone call. This is called *virtual dial up*, in contrast with the traditional dial up currently used. The main difference with the latter is, when not needed, the user is confined in its private domain; that is, with the user's private IP and with a limited scope. In this case, for example, the user would be able to make an IP telephone call with a scope limited to the access domain, thus without the need for a public IP address.
- **Transparency to applications**
The tunnel is completely transparent to applications, which do not

need to be modified to adapt to the tunnelling mechanism underlying. This happens because no address translation is performed and the visibility is obtained by packet encapsulation: some applications, for example, FTP data transfer, insert the IP address inside the packet payload, and a suitable address translation mechanism should be aware of it. In such cases we should have an address translation mechanism specific for each application, with all related problems of complexity and modularity. In contrast, the tunnelling mechanism encapsulates packets on the client side and extracts them at the server side in a completely transparent way.

- **Availability and user friendliness**
Another advantage concerns the user, who would not be required to perform an expensive upgrade of its equipment, since newer operating systems include tunnelling support (for example, PPTP is

included in Microsoft Win 98/Win NT 4.0, and Windows 2000 will contain L2TP support). Moreover, user friendliness is to be considered. With these operating systems the tunnel set up is performed by the same operations needed for accessing an ISP by a modem, therefore with no complex procedures to learn.

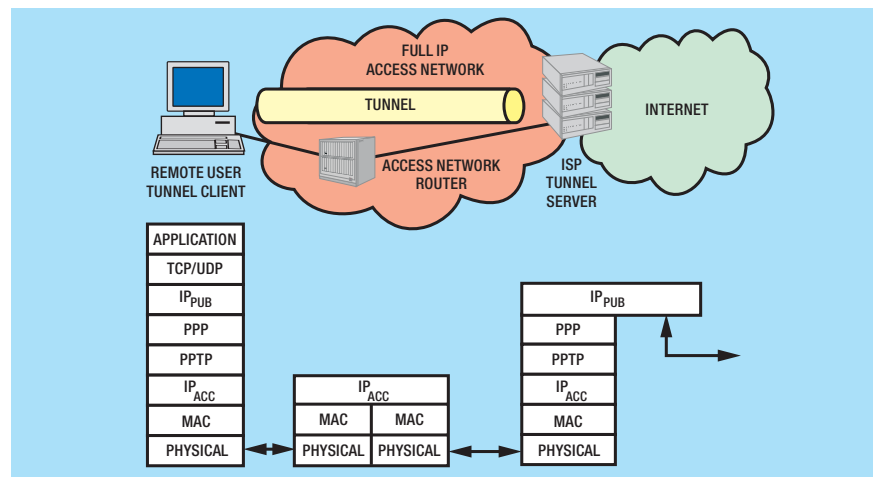
- **High security**
In the access environment some security concerns must be considered. First of all a reliable method for user authentication is required to avoid unauthorised access to service. Moreover, the need for confidential transactions may require data encryption. Both these aspects are covered by tunnel implementations; this is especially true if referred to IPSec, which is a protocol designed mainly for security purposes.

In the near term the previously mentioned advantages, along with the reliability and the widespread use of the tunnelling mechanism, justify its adoption in the context of a full IP access network. It represents a simple and effective solution in view of the rapid deployment of the scenario depicted in Figure 5.

There are obviously some open issues that should be addressed in order to achieve better performance. The main one concerns the overhead introduced by tunnelling; Figure 6 shows the situation obtained using PPTP as the tunnelling protocol in a full IP access scenario.

The situation is slightly different from the one depicted in Figure 4. Now the user does not have to set up a PPP connection to an ISP in order to acquire a public IP address. In fact, thanks to the full IP access

Figure 6 – Overhead introduced by a PPTP tunnel in the access



network, the user already belongs to an IP domain and therefore has an IP address (even if a private one); this suffices in the case of transactions within the access domain (for example, an IP call or a mail message to another user in the same domain). In contrast a tunnel must be properly set up to reach the public Internet (for example, for a data transfer from a content provider located outside the domain). Note that in the latter case the tunnel server is located in the ISP, which is therefore responsible for assigning a public IP address to the user, while routers in the access network are transparently crossed by the tunnel. The overhead obtained is due to the protocol stack increase caused by the encapsulation mechanism. It can be reduced by proper header compression techniques, such as those proposed for L2TP, or otherwise by using a layer 3 tunnelling protocol (for example, IPSec), which introduces fewer overheads. Along with this problem there are some other minor aspects to consider, that are outlined in the following paragraph along with a brief explanation.

Conclusions

In this paper a novel use of tunnelling techniques is proposed, quite different from the classical applications in which tunnels are used to cross a public network infrastructure, either to reach a mobile host (for example, in the case of IP mobility) or a private domain (for example, in the case of VPNs). In more detail the idea is to use the tunnelling technique to transparently cross the private access environment (that, as previously mentioned, should be supposed full IP) in order to reach the public IP network (for example, Internet). The tunnelling mechanism well adapts to this purpose, since it allows advantages in terms of:

- statistical multiplexing and flexibility in the address re-use;
- transparency to applications;
- availability in standard operating systems and user friendliness; and
- high security.

In spite of all these features, there are some disadvantages to be kept in mind. Tunnelling introduces some amount of overhead, which limits the bandwidth for user data. In an access environment, where typically

resources are limited (especially a wireless one), this can represent a problem. Moreover additional time for encapsulation, unencapsulation and processing of tunnelled packets is required, possibly introducing unpredictable delays that could affect end-to-end delay (this is especially troublesome for real-time traffic). Finally, the implementation of QoS mechanisms (for example, following the differentiated or the integrated services paradigm) on tunnelled flows may not be simple. This can be easily understood by observing that both these QoS mechanisms exploit the knowledge of information contained in the internal IP header, which is hidden in the encapsulating packet.

None of these problems represents an insuperable obstacle; for example, the overhead due to the encapsulation mechanism can be reduced by proper header compression techniques or by using a layer 3 tunnelling protocol. Moreover delays introduced by the encapsulation-unencapsulation process are usually negligible if compared to the end-to-end delay experienced by the tunnelled flow. Finally the problem of supporting quality of service for tunnelled flows is currently under study by IETF: some solutions have already been proposed for the use of resource reservation protocol (RSVP) with IPSec and for a differentiated service extension for L2TP. All the reasons explained above justify the adoption of the tunnel mechanism in the access environment; tunnels provide a fine and easy way to deploy a solution to some of the intrinsic problems of a full IP access structure, such as those related to addressing and worldwide visibility.

References

- 1 CASEY, L. An extended IP VPN Architecture. draft-casey-vpn-extns-00.txt, Nov. 1998. <http://www.ietf.org/>
- 2 HAMZEH, K.; PALL, G. S.; VERTHEIN, W.; TAARUD, J.; LITTLE, W. A.; and ZORN, G. Point-to-Point Tunnelling Protocol (PPTP). draft-ietf-pppext-pptp-10.txt, April 1999. <http://www.ietf.org/>
- 3 TOWNSLEY, W. M.; VALENCIA, A.; RUBENS, A.; PALL, G. S.; ZORN, G.; and PALTER, B. Layer Two Tunneling Protocol (L2TP). draft-ietf-pppext-l2tp-15.txt, May 1999. <http://www.ietf.org/>
- 4 GUPTA, V. Secure, Remote Access over the Internet using IPSec. draft-gupta-ipsec-remote-access-01.txt. Nov. 1998. <http://www.ietf.org/>
- 5 Virtual Private Networking: an Overview. <http://msdn.microsoft.com/workshop/server/feature/vpnovw.asp?>
- 6 Remote Access: Understanding and Implementing Virtual Private Networking (VPN) Services. White paper, Bay Networks, <http://www.baynetworks.com/products/Papers/2746.html>

Biographies



Stefano Giacometti
Ericsson
Telecomunicazioni
S.p.A.

Stefano Giacometti received his degree in Electronic Engineering from University of Rome 'Tor Vergata' in 1997. He was then with CoRiTeL as a scholarship holder. His research interests focused on architectures and signalling protocols for broadband networks and integration of IP and ATM. He then joined CoRiTeL, being involved in ABR congestion control for ATM flow via satellite. He now he works for Ericsson Telecommunications, Research and Development department, being involved on high-bandwidth access and on real-time traffic on IP networks.



Roberto Mameli
CoRiTeL

Roberto Mameli received his degree in Telecommunication Engineering from the University of Rome 'La Sapienza' in 1997. Since September 1998 he has worked for CoRiTeL, where he takes part to the research activity of the SOFIA group (SOlution for Full IP Access). Within this project he is mainly interested in the investigation of problems related to quality of service and aspects concerning wireless access techniques and mobile IP issues.