Wireless Personal Communications (2006) 38: 125–141
DOI: 10.1007/s11277-006-9021-1

Virtual Private Ad Hoc Networking

JEROEN HOEBEKE, GERRY HOLDERBEKE, INGRID MOERMAN, BART DHOEDT and PIET DEMEESTER

Ghent University - IBBT - IMEC
Department of Information Technology (INTEC)
Gaston Crommenlaan 8, B-9050 Ghent, Belgium
E-mail: name.surname@intec.ugent.be

Abstract. The fact that a lot of applications require secure communication to take place only between a dynamic subset of distributed devices sharing a common context, is, from a network point of view, very challenging and demanding. Existing technologies such as VPN, P2P overlays or VLANs can only partially respond to these requirements. This observation is the key factor that has driven the proposal of the virtual private ad hoc network concept. Virtual private ad hoc networks (VPAN) are secure and self-organizing overlay networks on top of existing IP infrastructure that use ad hoc networking techniques to enable network connectivity. The underlying IP infrastructure can be the Internet, cellular networks, ad hoc networks, mesh networks . . . or combinations thereof. A virtual private ad hoc overlay network creates a transparent, shielded and trusted environment for the applications and services running on the participants' devices. The overlay uses internal addressing and ad hoc routing, thereby forming a virtual network on top of the physical infrastructure. In addition, the overlay must be self-organizing and self-maintaining upon member mobility or membership changes. This paper gives an overview of the potential applications, a high-level network architecture and the network challenges emerging from the novel concept of virtual private ad hoc networking.

Keywords: virtual private ad hoc networks, secure overlay, ad hoc networking, network virtualization, P2P

1. Introduction

The field of wireless and mobile communications has experienced an unprecedented growth during the last decade. Current 2G cellular systems have reached a high penetration rate, enabling world wide mobile connectivity. An increasing number of wireless LAN hot spots is emerging, allowing travellers with portable computers to surf the Internet from airports, railways, hotels or other public locations. Broadband Internet access is driving wireless LAN solutions in the home for sharing access between computers. In the meantime, 2G cellular networks are evolving to 3G, offering higher data rates, infotainment and location based or personalized services.

These technology trends illustrate how communication is evolving towards future large-scale and high-speed, all IP, 4th generation communication networks, enabling interconnectivity between a massive amount of devices and users anywhere, at anytime and from any device. On the other hand, these networks will overwhelm the user with available information, applications and services, a characteristic that is not always desired by the end user and that can result in potential security risks. Moreover, a lot of applications require secure communication to take place only between a dynamic subset of distributed devices sharing a common context (i.e. communication between your personal devices, communication with colleagues or friends . . .), a characteristic not reflected by our current and future communication networks.

Therefore, it is expected that apart from the evolution to integrate all devices into one large-scale IP network, the base network, an evolution towards network virtualization will take place, imposing a logical structure onto this base network [1]. These virtual networks will form a shielded and trusted environment for their participants, with its own internal routing, naming and addressing solutions, using the underlying base network as the enabler of connectivity and carrier of data. For the applications and services that use the virtual network, the underlying base network will be invisible. As a result, the base network will provide the end-to-end connectivity between all devices at anytime and at any place, but on top of it, massive amounts of virtual networks will be deployed, thereby logically structuring the network into small secure communities according to the needs of the end users.

In parallel, and part of the evolution towards 4th generation communication networks, alternative ways for mobile communication such as ad hoc networks and wireless mesh networks, in which mobile devices form self-creating, self-organizing and self-administering wireless networks, have received enormous interest form the research community the last decade [2]. Although these networks pave the way for many new and exciting applications, these types of networks have not yet witnessed mass market deployment and are not yet commercially viable despite the many research and development efforts. These efforts have led to many interesting protocols and network techniques, which can prove useful in other environments than wireless ad hoc and mesh networks. However, until now it has not been considered to deploy these ad hoc protocols or modifications thereof on top of existing network infrastructures.

In this paper, we will combine these two trends, namely evolution towards network virtualization and the migration of ad hoc networking techniques outside the realm of ad hoc or mesh networks, into the novel concept of virtual private ad hoc networks (VPANs). Section 2 further discusses the concept of VPANs and the requirements that need to be fulfilled in order to realize the concept and to exploit its full potential. In Section 3 we will investigate related work and point out to what extent it meets the VPAN requirements. The next two sections will present a high-level network and node architecture together with the potential applications of VPANs. Section 6 will discuss the research challenges faced with in order to realize the VPAN concept. Finally, the last section concludes the paper.

2. Concept and Requirements

In the next generation all IP communication networks it can be expected that dynamic subsets of distributed nodes will organize themselves into logical virtual networks, providing a secure and transparent overlay to their applications and services. This virtual network only uses the underlying IP network as a carrier of data and provides its own internal routing and addressing schemes. Applications and services can be given access rights to such an overlay, thereby operating within this secure and confined environment, with all security, networking and management details handled by the overlay.

The nodes forming the overlay, called members, can be geographically distributed, belonging to networks with very different characteristics in terms of communication medium (e.g. wired versus wireless), mobility, bandwidth. . . Also, depending on the purpose of the overlay, its composition can change over time, with new members joining and existing members leaving. Consequently, these overlays are dynamic entities, both in composition and topology. In addition, the formation and maintenance of the overlay must take place with minimal or no user intervention. The above characteristics reveal many similarities with the characteristics encountered in ad hoc networks, encouraging the application of ad hoc network paradigms to the concept of network virtualization. Based on the above discussion, we define a VPAN as a secure and self-organizing virtual overlay network of distributed nodes, deploying ad hoc networking techniques to enable connectivity. In order to realize the challenging concept of VPANs, the following main requirements and characteristics need to be taken into account, some of which are closely related:

Membership configuration and management: A VPAN consists of a distributed subset of nodes, called members, of which the composition can change over time. Therefore, mechanisms to initialize new VPANs and to define, configure and manage its membership information are required.

Distributed operation: The members of a VPAN can be geographically dispersed, implying the need for distributed member discovery and VPAN formation and maintenance solutions.

Security: A VPAN needs to provide a secure communication environment. The security measures that need to be taken are located at different levels: access to the VPAN based on membership, application and service VPAN access rights and secure data transport. As such, security mechanisms for trust, encryption, authentication and authorization are indispensable.

Self-organization and mobility management: A VPAN should be self-organizing, meaning that the VPAN members need to be able to discover each other and form a secure overlay without user intervention. In addition, mobility management solutions are needed in order to maintain the VPAN in case of member mobility.

Application support: Users should be able to specify which applications, services, data... are reachable through or have access to a specific VPAN. As such the VPAN not only involves secure communication between a limited subset of nodes, but communication is also limited in terms of available applications, services and data.

Local private address space: Each VPAN will have its own local private address space, separated from the global IP address space used in the underlying base network. This address is VPAN based and independent of the number of interfaces or any public IP addresses assigned to these interfaces. Applications running within a VPAN use this private address independent of changes in the global address(es) of the node due to node mobility. As such, the VPAN can operate transparently and support session continuity without having to worry about the mobility management mechanisms (e.g. handover) in the underlying physical network. These mechanisms only ensure the end-to-end connectivity in the network on top of which the VPAN is established. In addition, as private addresses are used, the VPAN is not visible from the Internet directly, adding an additional level of security.

Ad hoc routing and tunnel management: Apart from its own addressing scheme, each VPAN also has its own internal routing mechanisms. As the composition and the topology of the VPAN can be dynamic, ad hoc routing techniques will be used for efficient internal routing. In addition, as a VPAN is an overlay network using private addresses, in many cases links between members are logical links, spanning multiple physical hops by the means of tunnel mechanisms. Consequently, VPAN forwarding should also encompass the notion of tunnels.

Scalability: Depending on the scenario the VPAN concept is used in, the number of members forming the VPAN can become quite large, making scalability a potential issue that has to be taken care of.

3. Related Work

During the past years several networking techniques and concepts have been proposed or deployed that capture some of the characteristics of the VPAN concept. In this section we will first give the main characteristics of related concepts and techniques, followed by an overview to what extent they meet the requirements imposed by the VPAN concept and where they fail to deal with all the VPAN challenges. This overview can be found in Table 1.

3.1. Virtual LANs

A virtual LAN, commonly known as a VLAN, is the logical grouping of a subset of devices belonging to an Ethernet system, possibly located on different segments [3]. These devices appear to be on a separate LAN, and thus broadcast domain, implemented on top of the physical network. VLANs are configured through software rather than hardware, which make them extremely flexible. All traffic is isolated within the VLAN and not visible for other VLANs defined on the same physical LAN. The membership of a VLAN can be defined in terms of switch port, MAC address or layer 3 information. Essentially, a VLAN is a logical layer 2 overlay and does not involve any IP addressing or routing.

3.2. Virtual Private Networks

A VPN is a private data network that makes use of the public telecommunications infrastructure, thereby maintaining privacy through the use of a tunnelling protocol and security procedures [4]. VPNs are mostly established between two sites or a, potentially mobile, client and a site, although more complex and dynamic setups involving multiple networks are possible [5]. VPNs can offer authentication of the tunnel endpoints, confidentiality and authenticity of data transferred between these endpoints [6]. Mostly, traffic is routed over the tunnels by manually updating the routing tables in the endpoints, although a routing daemon can be run in order to exchange tables.

3.3. P2P Application Level Overlays

In P2P application level overlays, applications running on distributed systems create logical links between each other using native Internet routing and standard IP addresses [7]. The result is a self-organizing semantic layer above the basic transport protocol level. P2P overlay networks exist in all flavours, offering a variety of features such efficient search operations, routing algorithms for optimizations, selection of nearby peers, anonymity, security . . . A detailed overview can be found in [8]. Contrary to the VPAN concept, application level overlays are triggered by and established between applications, whereas VPANs create communities accessible for a wide range of applications and services, offering much more flexibility in many potential scenarios.

3.4. Other

In [9] and [10], the concept of VIOLINs is proposed, isolated application-level virtual networks for virtual machine communications, that are created on top of an overlay infrastructure, which in turn is deployed on top of the Internet infrastructure. Each virtual machine has a complete protocol stack, with at the bottom a virtual interface, and can run any type of network application. In order to confine all communication within the VIOLIN, private addressing in tandem with UDP tunnelling is used: all traffic exchanged between two remote virtual interfaces is encapsulated in an application-level UDP tunnel. As the VIOLIN concept has been developed for use in grid computing, mechanisms not needed for VPANs, such as on-demand creation, deletion and migration of entities has been foreseen, together with topology adaptation based on the application requirements, computational and network resources. The VIOLIN is managed by its owner, having full administrator privileges.

Within the IST-MAGNET project [11], in which the authors participate, Virtual Personal Overlay Networks are being developed. Such a network is a virtual overlay network that encompasses all of a person's devices independent of their physical location [12]. Networking solutions such as private addressing, dynamic tunnelling and agent techniques are being developed in order to realize this concept. Although the MAGNET concept shows some similarities with the VPAN concept described in this paper, its scope is much more limited as it is completely based on personal networking. As such, the VPAN concept opens up possibilities to a much broader range of applications and imposes additional challenges and research issues.

3.5. Conclusion

In Table 1, a summary is given of which techniques the discussed technologies offer for dealing with the requirements imposed by the VPAN concept and to what extent they meet the given VPAN requirements. This state-of-the-art overview reveals that, although existing solutions partially can meet the requirements of VPANs, none of them is capable to grasp all challenges imposed by this novel concept. In the following sections we will further discuss the potential VPAN architecture and the involved challenges and research issues.

4. High-Level Network and Node Architecture

4.1. High-Level Network Architecture

In this section, we will further explain the concept of VPANs, by describing its high-level network architecture. We will differentiate between three different types of VPANs as they have different implications on the network protocols and solutions for VPAN formation and maintenance.

4.1.1. *Localized VPAN*

In a localized VPAN, each member of the VPAN can reach every other member by only using VPAN members as relays (both wireless/wired links). As a consequence, the direct connectivity between all neighbouring members results in an overlay that encompasses all members. Normally, especially for this VPAN type established on top of an ad hoc network, no infrastructure is present to support VPAN formation, so distributed solutions are needed. Figure 1 shows an example of a localized VPAN, where all members are interconnected either wired or wireless without using any non-member nodes.

4.1.2. *Distributed Infrastructured VPAN*

In a distributed VPAN, connectivity between members has to use non-member nodes as relays in order to forward their traffic. In order to secure this traffic, dynamic tunnels between member

Table 1. Which techniques do existing technologies provide to deal with the VPAN requirements?

	VLAN	VPN	P2P Overlay	Virtual distributed environment (VIOLIN)
Protocol stack layer	Layer 2	Layer 3	Application layer	Application layer
Membership configuration and management	Membership based on switch port, MAC address or layer 3 information. Manual, semi-automated or fully automatic configuration.	Statically configured in tunnel endpoints. Central tunnel management in LeetNet.	Members are the peers, running P2P software.	Management is done by the owner of the VIOLIN, having administrator privileges.
Distributed operation	Limited to one Ethernet system.	Tunnel endpoints are distributed.	Distributed system, sometimes supported by centralized facilities (e.g. Napster).	Distributed overlay of virtual machines.
Security	Access based on VLAN membership. Traffic containment within VLAN. No other security mechanisms.	Authentication of tunnel endpoints, confidentiality and authenticity of data transferred between these endpoints.	In some case, features such as authentication, trust, anonymity, overlay access control… can be offered.	Depends on the network protocols used within the VIOLIN, as their packets are transferred over UDP tunnels.
Self-organization and mobility management	After VLAN configuration, the VLAN is automatically formed and maintained. Support of member mobility within the same Ethernet system.	No mobility management (through dynamic tunnelling). Some aspects of self-organization in dynamic VPNs.	Self-organization through direct or indirect discovery of other peers for overlay formation. Overlay topology control or routing optimizations can be used.	On-demand creation of virtual machines and the interconnecting virtual IP network. Dynamic topology adaptation possible.
Application support	Provides only layer 2 connectivity within a confined environment.	Provides only secure layer 3 connectivity between 2 endpoints or sites.	Application bound overlays (limited flexibility). The application enables the peers to offer services to and consume services from other peers.	Each virtual machine can run its own (distributed) network applications.
Local private address space (session continuity)	All members in a VLAN share a common – private or public – address space.	Addresses depend on schemes deployed in networks behind the endpoints.	Public IP addresses are used. Overlay is used for and to optimize application specific actions (e.g. search operations…)	Private IP addresses are used in order to confine all communications within the VIOLIN.

Continued on next page

Table 1. Continued

	VLAN	VPN	P2P Overlay	Virtual distributed environment (VIOLIN)
Dynamic internal (ad hoc) routing and tunnel management	Layer 2 switching, no routing involved.	Statically configured tunnels and routes in endpoints or routing daemon to exchange tables between endpoints in order to specify which traffic is sent over the tunnel.	Public IP addresses are used to obtain end-to-end P2P connectivity. No tunnelling needed.	VIOLIN has virtual routers for internal routing. Topology adaptation (adding/removing links and forwarding rules).

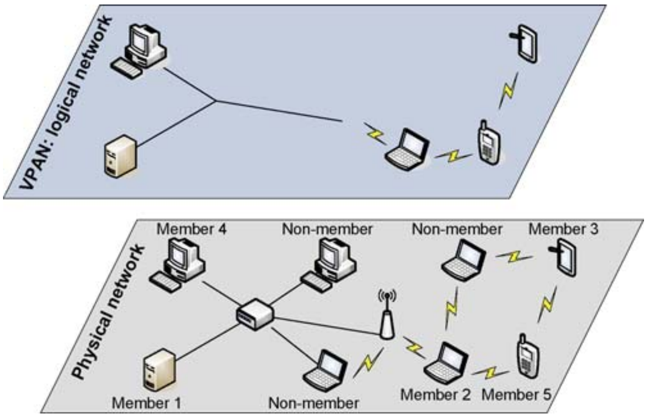

Figure 1. Localized VPAN.

nodes need to be established. In the infrastructured VPAN, the members are interconnected over the Internet, using nodes (mainly routers) in the infrastructure as relays. In this type of VPANs, infrastructure support, for instance from service providers or the naming system, can assist the VPAN membership management, member discovery, formation, routing and mobility management. In some networking aspects, the distributed infrastructured VPAN, when applied to personal nodes, resembles closely the approach adopted in the Personal Network concept [13]. Figure 2 gives an example of the distributed infrastructured VPAN.

4.1.3. *Distributed Ad Hoc VPAN*

Again, as it is distributed, connectivity between members has to use non-member nodes as relays in order to forward their traffic over end-to-end tunnels. However, in the ad hoc VPAN, illustrated in Figure 3, ad hoc nodes are used as relays. As no dedicated infrastructure can be assumed, VPAN membership management, member discovery, formation, routing and mobility management has to be done in a completely distributed manner. This will impose additional challenges to the network solutions, challenges highly similar to the ones encountered in ad hoc networks.

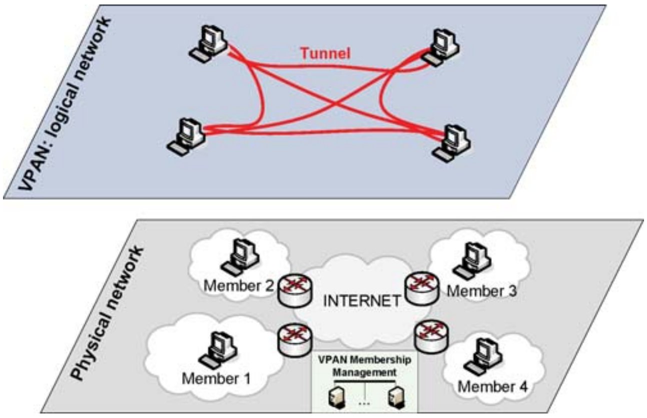

Figure 2. Distributed infrastructured VPAN.

4.1.4. *Hybrid VPANs*

In many potential scenarios, a combination of the above architectures will be encountered, resulting in hybrid VPANs and the required integration and harmonization of solutions.

4.2. High-Level Node Architecture

If we try to map the VPAN requirements onto the individual nodes, a conceptual high-level node architecture can be derived. Figure 4 presents the proposed high-level architecture of the protocol stack in a VPAN node. In traditional IP nodes, only one IP protocol stack responsible for routing, fragmentation. . . is present.

However, if we want the coexistence of multiple virtual overlay networks, each with their own private addressing scheme and routing protocol, multiple IP stacks, one for each VPAN, are required next to each other, where a specific VPAN stack is installed dynamically only when the specific VPAN is active.

Apart from traditional IP functionalities such as routing, decreasing time-to-live field . . . additional functionalities for VPAN support are needed, which can be divided into a data and

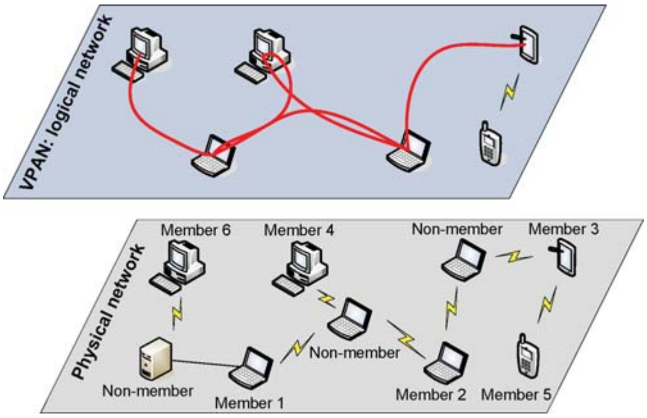

Figure 3. Distributed Ad Hoc VPAN.

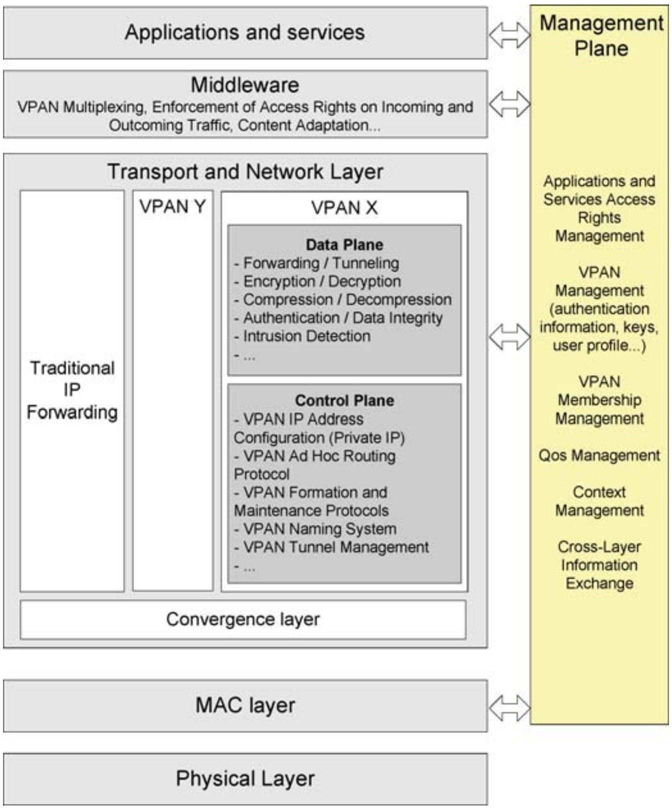
Applications and services
Middleware
VPAN Multiplexing, Enforcement of Access Rights on Incoming and Outcoming Traffic, Content Adaptation...
Transport and Network Layer
Traditional IP Forwarding
VPAN Y
VPAN X
Data Plane
- Forwarding / Tunneling
- Encryption / Decryption
- Compression / Decompression
- Authentication / Data Integrity
- Intrusion Detection
- ...
Control Plane
- VPAN IP Address Configuration (Private IP)
- VPAN Ad Hoc Routing Protocol
- VPAN Formation and Maintenance Protocols
- VPAN Naming System
- VPAN Tunnel Management
- ...
Convergence layer
MAC layer
Physical Layer
Management Plane
Applications and Services Access Rights Management
VPAN Management (authentication information, keys, user profile...)
VPAN Membership Management
Qos Management
Context Management
Cross-Layer Information Exchange

Figure 4. Conceptual high-level protocol stack architecture of a VPAN node.

control plane. In the data plane, the functionalities encompass, amongst others, forwarding, encryption, tunnelling ... In the control plane, we find all functionalities related to the intra-VPAN routing protocol, VPAN formation, tunnel management, private IP address configuration ... Below the multiple IP protocol stacks, a convergence layer is needed that is capable of demultiplexing incoming packets to the corresponding VPAN and forwarding outgoing packets to the corresponding interface. At the higher layer, a middleware system is needed. One of the main functionalities of this component is to act as a firewall for the resources and services at the higher layers that enforces access rights: only applications and services with the correct privileges are given access to a specific VPAN. As such, the middleware system in tandem with the VPAN protocol stack needs to create a confined and secure environment for applications and services. Also firewall functionality in terms of membership could be enforced here: packets coming from a specific VPAN are only forwarded to the higher layers if the originating member application has the required VPAN access rights.

In addition, the middleware system could also provide more complex functionalities such as content adaptation according to the context or capabilities and access rights of the receiver or generic support for distributed applications and services. Next to all layered functionalities, a management plane is needed that spans all layers. This plane needs to provide the user with functionalities for the management of the access rights of the applications and services, for the management of the VPAN and its members (e.g. storage of member profiles such as keys, authentication information ...), for the management and exchange of context and cross-layer information ...

Table 2. Potential VPAN scenarios

Application	VPAN
Emergency networks	Rescue people (police men, fire fighting) organized in teams
Military communications	Soldiers divided in separated military units, potentially hierarchically organized
Enterprise networking	Overlay network between collaborating people, within a department, between people at a construction site, monitoring networks, augmented reality support, virtual service providers . . .
Personal networking	Overlay network of all your personal devices, distributed virtual desktop
Education	Virtual classrooms, project collaborations
Entertainment	Multi-user games, closed P2P communities, augmented reality
Commercial and civilian environments	Cab network, touristic information, surveillance systems, building automation, e-health networks . . .

The above description is far from a complete solution, but only intends to present the generic high-level building blocks of a conceptual node architecture that are needed to fulfil the requirements imposed by VPAN.

5. Applications

As already stated, in many scenarios communication will take place between users or devices that share a common context and should not be visible or accessible for outsiders. In addition, accessibility will often be restricted to a limited number of applications, service and resources. These common characteristics are all captured by the proposed generic VPAN concept, which can act as an enabler for all these applications, without requiring individual, not reusable and often incompatible solutions at the application layer. Table 2 gives a list of some potential applications that can really benefit from the proposed VPAN concept. These applications can span all kinds of networks, including wireless and wired, infrastructured and infrastructureless communications and can all be mapped on one of the above high-level network architectures. In all these scenarios, the applications and users will benefit from the presence of the VPAN, as it transparently takes care of membership, connectivity, mobility, security, access rights enforcement . . . As can be seen, the potential applications span multiple domains, ranging from leisure to business, from limited to high security . . . Further, each user will simultaneously take part in multiple VPANs at the same time, making the coexistence of multiple VPANs next to each other a stringent requirement.

In the following section we will discuss the research challenges we are faced with for providing virtual private ad hoc networks that enable the applications presented.

6. Challenges

The proposed VPAN concept involves a lot of requirements that need to be fulfilled in order to realize its potential. Addressing these requirements introduces numerous research challenges. The main challenges, together with some guidelines and thoughts on potential solutions are presented in this section.

6.1. VPAN Definition and Management

Before a VPAN can be formed and organize itself, the VPAN, its policies and its members, or its membership rules, first have to be defined. The creation of a new VPAN will be triggered by a service provider or, more commonly, by an individual person. In the former case, the service provider can define VPANs to which individuals can subscribe or be invited. In the latter case, when defining the VPAN, all information related to the VPAN definition and membership will be located in this individual member. However, in most application scenarios, the VPAN can be formed and operate in the presence of a subset of the members, in which the creator of the VPAN mostly does not have to be included. Therefore, the VPAN definition and membership information needs to be distributed over the other members that will be added to the VPAN or, in the case of an infrastructured VPAN, could be outsourced to a service provider. The above discussion makes clear that during the VPAN definition the following questions have to be answered, and their resulting policies have to be defined:

- Who is the creator of the VPAN?
- How and where is the VPAN definition and membership information stored and/or distributed?
- When can the VPAN be formed and become operational (if the creator is online, if all members are online, if a subset or quorum of the members is online)?
- How are new members added to the VPAN (membership policies)?
- Member authentication?

As soon as the VPAN has been defined, new members can be added to the VPAN. There are a number of potential solutions for membership management, of which the preferred one has to be defined in the VPAN membership policies during its creation. Nodes can be invited to join the VPAN by (one of) its members. If the existence of the VPAN can be consulted by non-VPAN nodes, nodes can request to join the VPAN by contacting the creator or one of the member nodes. Another solution could be to formally define membership rules and automating the processing of member addition by checking these rules against an (authenticated) user profile of the potential members. In all cases, it can further be defined if the addition of a new member needs to be confirmed by only one member, the creator of the VPAN, all existing members or a quorum of the members. Additionally, when new members are added to the VPAN, existing members need to be informed of, or at least able to consult, the new membership information. The choice of the solution adopted will highly depend on the type of VPAN (see Section 4.1) and the scenario used in. Apart from member addition, similar mechanisms need to be developed for member removal and banishment.

In the above discussion, we have considered membership as an abstract concept. However, in order to implement this concept, security issues such as member identification, mutual trust and authentication, key distribution... need to be considered, as they are indispensable for the secure formation and self-organization of the VPAN and inherently tied to the concept of communities. These issues are considered in the next subsection.

6.2. Security

As already stated, the membership concept inherently implies the notion of security. More precisely, the membership concept should allow the establishment of trust between the VPAN members in order to offer the following security services:

- Member identification and authentication: a member needs to be able to check whether another one belongs to the same VPAN and holds the identity it claims.
- Authentication of communication between VPAN members.
- Confidentiality of communication between VPAN members.

In the case of localized VPANs or distributed ad hoc VPANs, no infrastructure to support the establishment of trust relationships can be assumed and distributed solutions are needed, with minimal impact in terms of user effort. An overview of existing solutions and a novel solution, together with implementation hints, for distributed establishment of trust relationships can be found in [14]. In the case of a distributed infrastructured VPAN, solutions relying on a service provider or central entity could be used in order to support the membership management and trust establishment (e.g. using certificates).

Based on these trust relations, directly connected members can establish a secure link between each other. If not directly connected, they can establish a secure tunnel between them. Both these secure links and tunnels are prerequisites for the formation of the secure overlay.

6.3. VPAN Formation and Self-Organization

As soon as the VPAN has been defined, its policies have been created and trust relations have been established, VPAN formation and self-organization can take place. Again, different types of VPANs can result in different solutions, requiring harmonization when being deployed in hybrid VPANs.

In a localized VPAN, it is sufficient that every member node discovers it's neighbouring, i.e. directly connected, members in order to form the overlay. To this end, a neighbour discovery mechanism with mutual member authentication, based on the predefined trust relations, is needed. In a distributed infrastructured VPAN, mechanisms deployed in P2P overlay networks could be used for member discovery. Again, different options exist. First of all, a fixed anchor point that acts as a membership registration system (e.g. a management entity in the infrastructure, service provider support) could help the VPAN members to find out which other members are online. Also, P2P techniques using completely distributed discovery through the use of rendez-vous peers could be deployed as member discovery mechanism. In a distributed ad hoc VPAN, mechanisms deployed in ad hoc networks for path discovery could be deployed for member discovery. Based on membership knowledge, nodes can actively search for missing members. In addition, recovery mechanisms are needed in case of temporary connectivity loss to a previously found member due to mobility and could be enhanced based on prior knowledge of the location of that member.

These member discovery mechanisms form the basis of the formation of the VPAN overlay, as it is needed for secure link and tunnel establishment between member nodes. In the distributed VPAN types, the formation of the VPAN can be subject to different policies. First of all, the VPAN formation strategy can be classified as either always-on or on-demand, depending on the fact if connectivity between all members is maintained continuously or only when requested by the applications and services running on top of it. Further, different strategies regarding topology control could be deployed, ranging from minimal connectivity, for connecting all members or only the active applications and services, to a full mesh between all members.

6.4. Addressing and Routing

Using the mechanisms described in the previous section, members are able to discover each other and to form an overlay that consists of either secure links or tunnels. Within the overlay, the members will use private IP addresses and will run an ad hoc routing protocol. As such, the overlay is used to transparently exchange VPAN data and control messages, using the public IP infrastructure as carrier.

As already stated, each VPAN will have its own local address space, from which each member is assigned one address, independent of its number of interfaces. All applications and services that communicate within the VPAN will use this address. This address space is confined within the VPAN and invisible to the outside world by tunnelling or link encryption. For automating the address assignment procedure, two options exist: stateful auto-configuration and stateless auto-configuration. For the former, mechanisms similar to the one deployed for the membership management could be used. The latter option requires the avoidance of duplicate addresses. Here, solutions using ad hoc duplicate address detection techniques or IPv6 stateless address auto-configuration could be deployed. Further, when running multiple VPANs within the same device, their address spaces need to be distinguishable, which could be achieved by assigning each VPAN a different address prefix.

Concerning the overlay ad hoc routing, different ad hoc routing techniques (reactive, proactive, hybrid, adaptive. . .) are possible and the choice of the most appropriate one will depend on multiple factors: application requirements and traffic, on-demand or always-on VPAN formation, context. . .

6.5. Member Mobility Management (VPAN Maintenance)

The members of a VPAN can be mobile, requiring adaptations to the VPAN overlay. For the localized VPAN efficient member (i.e. neighbour) discovery and (layer 2) link break detection mechanisms can improve VPAN maintenance. For the other VPAN types, member mobility resulting in a change of public IP address can cause the breakdown of tunnels established between members of the VPAN overlay, requiring dynamic tunnel reestablishment mechanisms and interaction with the membership management or member discovery framework. As private IP addresses are used within a VPAN, session continuity can be assured, provided the overlay and tunnel management mechanisms are efficient enough.

6.6. Application Middleware

As previously mentioned, the main functionality of this component is to act as a firewall for the resources and services at the higher layers: only applications and services with the correct privileges are given access to a specific VPAN. VPAN members should be able to specify to what extent their applications and services have access to the VPAN and to what extent other VPAN members have rights to access these applications and services. Further, applications and services need to be able to specify the VPAN they want to use. Therefore, a powerful interface between the VPAN protocol stack and the service and application layer is needed, in the form of a platform independent generic middleware system.

6.7. Others

The previous subsections presented the main challenges of the VPAN concept and some ideas to tackle them. Of course, this description is far from complete as many other challenges exist or will emerge: naming, QoS, context information to improve networking and management solutions, intrusion detection, dealing with multiple or even hierarchical VPANs, traffic optimization...

7. Conclusion

In the next-generation all-IP communication network, consisting of heterogeneous wired and wireless technologies, a lot of potential applications will require, or benefit from, the ability to securely communicate only between a dynamic subset of distributed devices. This characteristic of communication is often overlooked when considering next generation communication networks and existing technologies can only partially respond to the challenges involved. In this paper, we have presented Virtual Private Ad Hoc Networks, a concept that merges network virtualization and ad hoc networking techniques in order to create a transparent, trusted and self-organizing environment for applications and services. A VPAN can be defined as a secure and self-organizing virtual overlay network for applications and services on distributed nodes, deploying ad hoc networking techniques to enable connectivity. In this paper we have extensively discussed the main requirements, characteristics, high-level architecture and research challenges of this challenging concept. However, this paper leaves open as many questions as it addresses and hence should be seen as a proposal or guideline towards future research, but not as a definitive answer.

Acknowledgements

Jeroen Hoebeke is research assistant with the Fund for Scientific Research – Flanders (F.W.O.-V., Belgium). This research is partly funded by the Belgian Science Policy through the IAP V/11 contract, by The Institute for the Promotion of Innovation by Science and Technology in Flanders (IWT) through the contract No. 020152, by the Fund for Scientific Re-search – Flanders (F.W.O.-V., Belgium) and by the EC IST integrated project MAGNET (Contract 507102).

References

1. K. Birman, “The Next-Generation Internet: Unsafe at Any Speed?,” *Computer*, vol. 33 no. 8 , pp. 54–60, Aug. 2000.
2. J. Hoebeke, I. Moerman, B. Dhoedt, and P. Demeester, “An Overview of Mobile Ad Hoc Networks: Applications and Challenges,” *Journal of the Communications Network*, vol. 3, pp. 60–66, July–Sept. 2004.
3. J. Freeman and D. Passmore, The Virtual LAN Technology Report, Decisys, Inc., Sterling, VA, 1996.
4. VPN Consortium, VPN Technologies: Definitions and Requirements, http://www.vpnc.org/vpn-technologies.html, White Paper, July 2004.
5. LeetNet: The True Dynamic VPN, http://www.leetnet.org.
6. S. Khanvilkar and A. Khokhar, “Virtual Private Networks: An Overview with Performance Evaluation,” *IEEE Communication magazine*, vol. 42 no. 10, pp. 146–154, Oct. 2004.
7. D. Doval and D. O’Mahony, “Overlay Networks: A Scalable Alternative for P2P,” *IEEE Internet Computing*, vol. 7 no. 4, pp. 79–82, July/Aug. 2003.

8. K. Lua, J. Crowcroft, M. Pias, R. Sharma and S. Lim, “A Survey and Comparison of Peer-to-Peer Overlay Network Schemes,” *Communications Surveys & Tutorials, IEEE*, pp. 72–93, Second Quarter 2005.
9. P. Ruth, X. Jiang, D. Xu, and S. Goasguen, “Virtual Distributed Environments in a Shared Infrastructure,” *Computer*, vol. 38 no. 5, pp. 63–69, May 2005.
10. X. Jiang and D. Xu, “VIOLIN: Virtual Internetworking on OverLay INfrastructure,” Department of Computer Sciences Technical Report CSD TR 03-027, July 2003.
11. My Personal Adaptive Global Net, FP6-IST-IP-507102, http://www.ist-magnet.org.
12. I. Niemegeers and S. Heemstra de Groot, “From Personal Area Networks to Personal Networks: A User Oriented Approach,” *International Journal on Wireless Personal Communications*, vol. 22 no. 2, pp. 175–186, Aug. 2002.
13. L. Muñoz, L. Sanchez, J. Lanza, M. Alutoin. S. Lehtonen. D. Zeghlache, M. Girod Genet, W. Louati, J. Hoebeke, I. Moerman, G. Holderbeke, M. Ghader, and M. Jacobsson “A Proposal for Self-Organizing Networks,” *Wireless World Research Forum Meeting 15 (SIG 3) Paris*, White Paper, Dec. 2005.
14. N. Prigent, C. Bidan, J.-P. Andreaux, and O. Heen, “Secure long term communities in ad hoc networks,” *Proceedings of the 1st ACM workshop on Security of Ad Hoc and Sensor Networks (SASN 03)*, pp. 115–124, Aug. 2003.

Jeroen Hoebeke was born in Ghent, Belgium in 1979. In 2002 he received the Masters degree in engineering (Computer Science) from the University of Ghent. In August 2002, he joined the Broadband Communications Networks Group. His PhD research includes the development of adaptive routing protocol techniques for mobile ad hoc networks. His main research interests are in ad hoc wireless communications and, more generally, in broadband wireless communications. Within the European MAGNET project, he is actively involved in the development of a network architecture and demonstrator for Personal Networks, with a prime focus on routing and connectivity.

Gerry Holderbeke was born in Zottegem, Belgium in 1982. He graduated in Informatics at the University of Ghent in 2004. In August 2004 he joined the Broadband Communications Networks Group where he is currently working as a project developer. His research currently includes the development of an emulator for mobile ad hoc networks. His main research interests are in ad hoc networks and broadband wireless communications and involve routing, addressing and more generally, communication within mobile ad hoc networks and infrastructured networks. Within the European MAGNET project, he is actively involved in the development of a network architecture for Personal Networks, with a prime focus on the implementation of the routing architecture.

Ingrid Moerman was born in Gent, Belgium in 1965. She received the degree in Electrotechnical Engineering and the Ph.D degree from the Ghent University, Gent, Belgium in 1987 and 1992, respectively. Since 1987, she has been with the Interuniversity Micro-Electronics Centre (IMEC) at the Department of Information Technology (INTEC) of the Ghent University, where she conducted research in the field of optoelectronics. In 1997, she became a permanent member of the Research Staff at IMEC. Since 2000 she is part-time professor at the Ghent University. Since 2001 she has switched her research domain to broadband communication networks. She is currently involved in the research and education on broadband mobile & wireless communication networks and on multimedia over IP. The main research topics related to mobile & wireless communication networks are: wireless access to vehicles (high bandwidth & driving speed), adaptive QoS routing in wireless ad hoc networks, body area networks, protocol boosting on wireless links, design of fixed access/metro part, traffic engineering and QoS support in the wireless access network. Ingrid Moerman is author or co-author of more than 300 publications in the field of optoelectronics and communication networks.

Bart Dhoedt received a degree in Engineering from the Ghent University in 1990. In September 1990, he joined the Department of Information Technology of the Faculty of Applied Sciences, University of Ghent. His research, addressing the use of micro-optics to realize parallel free space optical interconnects, resulted in a PhD degree in 1995. After a 2 year post-doc in optoelectronics, he became professor at the Faculty of Applied Sciences, Department of Information Technology. Since then, he is responsible for several courses on algorithms, programming and software development. His research interests are software engineering and mobile & wireless communications. Bart Dhoedt is author or co-author of approximately 70 papers published in international journals or in the proceedings of international conferences. His current research addresses software technologies for communication networks, peer-to-peer networks, mobile networks and active networks.

Piet Demeester received the Masters degree in Electro-technical engineering and the Ph.D degree from the Ghent University, Gent, Belgium in 1984 and 1988, respectively. In 1992 he started a new research activity on broadband communication networks resulting in the IBCN-group (INTEC Broadband communications network research group). Since 1993 he became professor at the Ghent University where he is responsible for the research and education on communication networks. The research activities cover various communication networks (IP, ATM, SDH, WDM, access, active, mobile), including network planning, network and service management, telecom software, internetworking, network protocols for QoS support, etc. Piet Demeester is author of more than 300 publications in the area of network design, optimization and management. He is member of the editorial board of several international journals and has been member of several technical program committees (ECOC, OFC, DRCN, ICCCN, IZS, &).