Place Oriented Virtual Private Networks

Stamatis Karnouskos, Ingo Busse, Stefan Covaci

German National Research Center for Information Technology
Research Institute for Open Communication Systems (GMD-FOKUS)
Kaiserin Augusta Allee 31, D-10589 Berlin, Germany
{Karnouskos|Busse|Covaci}@fokus.gmd.de
http://www.fokus.gmd.de

Abstract

Today's network infrastructures are inadequate to cover the needs of small groups or new requirements that pop-up daily from the ever-increasing demand for more sophisticated services. Place-Oriented VPNs are based on the agent technology and can be built on top of existing infrastructure by any user. We expect that in the future this flexible, dynamic and low cost approach that supports not only group working but also advanced services for commercial needs will be at great use.

1. Introduction

Internet is expanding rapidly. It can be characterized as a universal community. As it evolves it becomes a virtual society and therefore it will reflect in a great percentage the real world. In such a virtual world individuals as well as enterprises form groups with common characteristics. An infrastructure that supports group working is needed. Of course this is old news, as solutions already exist. Yet we think that these solutions are inadequate to cover the needs and the rapid evolution of requirements. Moreover they are designed for big enterprises and not for small groups of people.

Today expensive setup and maintenance of communication infrastructures has resulted in being used by large enterprises that can afford it. However these large enterprises cover less than 10% of potential customers. There are small groups and associations that simply cant afford to setup, configure and maintain their own VPN. Furthermore day-by-day advanced services are required e.g. by electronic commerce. Current infrastructure evolves slowly and can't cover to the full length the ever-increasing need for security, flexibility, effectiveness, interoperability, performance, low cost etc.

Extensive support for small groups of people or even individuals is the now days need. With the right tools those groups should be able to deploy their own VPN in minimal time. In parallel they should be able to take advantage of advanced services that the underlying infrastructure has to offer. The idea of VPNs based on Places we present here, deals successfully with these matters.

2. Virtual Private Networks (VPNs)

A VPN is a communications environment in which access is controlled to permit peer connections only within a defined community of interest and is constructed through some form of partitioning of a common underlying communications medium, where this underlying communications medium provides services to the network on a non-exclusive basis [1]. Such kind of networks are deployed within a public network and aim at providing a private working environment to its users while also taking advantage of the efficiencies of the underlying infrastructure.

There are various types of VPNs depending on the functions they perform as well as on the different methods for constructing them. The bottom idea is that VPNs provide an abstraction of a network infrastructure that is used by groups or enterprises to cover their needs.

2.1 Virtual Networks' Requirements

Virtual networks aim at supporting groups. As in real world groups, have several requirements. We identify here the most significant features that a VPN must support.

Admission Control: A policy specifies who can access the network and its services. By providing selectively the services to restricted members the need for customization of the underlying infrastructure is satisfied. An enforcement engine enforces the policy between the members whose privileges and obligations vary.

Information Path Selection: there is a need to control the path that information follows. This can be due to security, control, customization, effectiveness or other

reasons. The topology, as well as the paths within such a network need to be dynamically adjustable. This increases flexibility and security of the VPN. For instance, we might require that our flow follows a specific route via trusted hosts or dynamically adjust to available bandwidth and select always the highest bandwidth connections links.

Resource Management: Each VPN is assigned statically or dynamically some resources. These resources may be multiplexed by the underlying infrastructure or not. That doesn't affect VPNs which function on top of it. What we require is the ability to easily manage the resources provided to the VPN administrator and dynamically change them in order to achieve optimal usage of resources and optimal performance of our VPN. Such resources include communication bandwidth, CPU cycles, memory allocation, disk space allocation as well as other services provided by the underlying infrastructure.

Security/Privacy: Security and privacy are fundamental needs of an enterprise. We need authorization and identification schemes in order to screen and serve requests. The requirements may vary. Some tasks may require strong authentication of all involved parties while others may require anonymity.

Communication: Communication between members also in a non-standardized way is required. Users should be able to make wrappers and communicate even if using different end-point devices with different capabilities. Users also can exchange information and make decisions based on info received and on environmental state.

Effectiveness/Performance: The VPN must operate at highest performance and take advantage of the underlying network resources and services at maximum.

Interoperability: Users of one VPN should be able to interact (communicate, exchange info etc) with users of other VPN infrastructures using tools. Those basic tools have somehow to be based on standards and be universally understandable by the customer VPNs. Wrappers can be build on top of the basic standardized tools/services to provide more customized and advanced services.

New Service Deployment: It is desirable that VPN operators will be able to provide advanced services to their users by making partial or full use of the underlying infrastructure services or even by developing their own. This should be done in minimal time and cost.

Flexibility/Adaptivity: Everything starting from VPN topology to configuration of services has to be dynamically updateable. Adding a new link to a physical network might be a cost and time expensive action but adding a host or rerouting the info on a VPN has to be done at very short time and with minimal cost. The VPN has to react to environmental changes and adapt its behavior in order to avoid malfunctions and achieve flexibility and effectiveness.

Low Cost: Finally the above requirements have to be provided at a low cost since the target area is Small & Medium Enterprises (SMEs) as well as small groups or even individuals that cant afford expensive or complicated solutions.

The above requirements reflect the need for a network infrastructure that can be tailored easily to the needs of small groups e.g. SMEs or even individuals.

2.2 What today's Internet satisfies?

Considering Internet today we can clearly see that it doesn't satisfy the requirements set by the previous section. That comes without surprise, since Internet was designed with different goals in mind. There is no admission control, everyone can join. Connectivity control is also non-existent since any user can connect to or contact another user if his Internet address is known.

There is hardly any authority and its users are optionally requested to "behave" in a certain way. The last is the result of a democratic/anarchistic design philosophy based on which the Internet was created and currently operates. The route selection and topology is limited and usually only administrators can control it. Users can't specify special routes that their data should follow. The resource management in such a distributed environment is a hot research area but no concrete commercial solutions exist yet.

The security is not guaranteed but it is an optional add-on. Current efforts in this direction include the Security Architecture for the Internet Protocol (IPsec) [11] which is a standardized technology that supports authentication/encryption and tunneling. It enables interoperability among legacy VPNs (routers using IPsec) and is applicable to IPv4 and IPv6. Unfortunately this superstructure approach although it offers a good notion of security it can't achieve QoS goals in a best effort network as Internet and is functions in low level (packet level).

Connection control is nonexistent as everyone is connected almost to everything, Internet makes it easy among other things to impersonate someone, to hijack

confidential data and communications, to publish unauthorized info etc. Internet is not yet fully trusted for electronic commerce since it can't guarantee secure business transactions which require a flexible and sophisticated security scheme.

What is important from the research as well as the commercial point of view is that Internet connects millions of people around the globe. Internet offers the missing link between the companies and the customers and that is exactly its strongest advantage. Enterprises that want to access those potential customers have to do it via Internet.

3. The Agent Distributed Environment

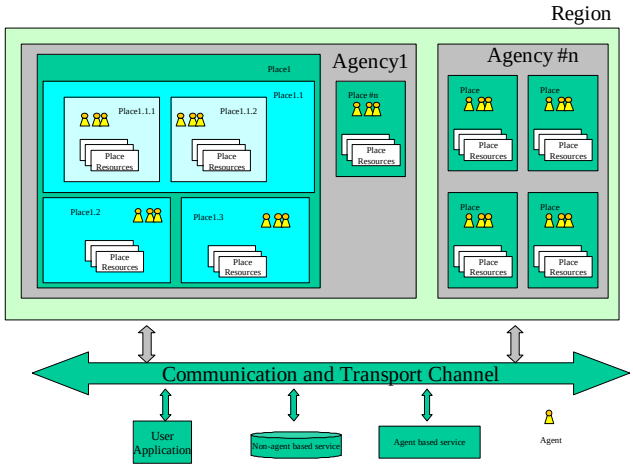

Figure 1. The distributed agent environment

An agency is the actual runtime environment where agents execute. Each agency runs to a host and has one or more Places. A Place provides a logical grouping of functionality within an agency. Agents [5] that reside within a place have something in common. This could mean created by the same user, or performing similar tasks, or even subjected to specific policy rules etc. An agency contains a number of places and each place can have several sub-places. The Region concept facilitates the management of the distributed components in the agent environment i.e. agencies, places and agents. Thus a region may reflect all agencies belonging to an organization or to a specific domain.

Places are created dynamically (on demand) or statically (reside permanently on the node). They have their own resources assigned to them and managed by the place administrator. When a sub-place is created (e.g. place 1.1) then some resources are assigned to this place (from the original resources assigned to place 1). If another place is created within place 1.1 e.g. place 1.1.1 then part of the resources belonging to place 1.1 are assigned to this newly created place and so on. We have a nested form of places. This nested form extends not only to the resources but also to Policies etc. A security architecture has been proposed [8] for this reason. Agents reside and execute within these places but in order to do so they have to successfully pass authentication and authorization procedures. Each place has a policy of its own, managed by its administrator. Policies are also in a nested form. That means that for an agent to execute in place 1.1.1 it has first to pass successfully the policy of place 1.1 and place 1.

4. Agent-powered Node

Agent distributed systems are usually installed in a host computer. This host can also be a router/ switch (referred as node in this paper) which gives more advanced abilities to the Agency if it can control node's functions. Currently service providers don't have access to node's control environment, algorithms or states, which makes the deployment of new services impossible due to the close nature of the networks. However there are efforts [3] to open-up the current network infrastructure. Active network technology [2] is a new research effort in this direction which targets not only the openness of current infrastructure but also aims to move dynamic computation within the network and therefore making it more intelligent (and open up this intelligence) not just to its end-points but in also in the intermediate nodes. As hosts can adapt flexible and intelligent roles within such a programmable infrastructure, the support of more sophisticated applications that take optimal advance of the resources offered to them is boosted.

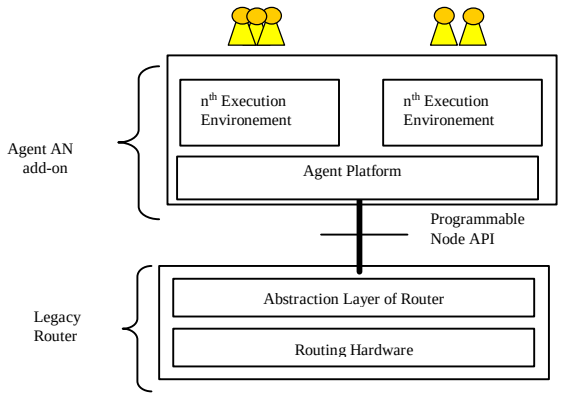

Figure 2. Agent-powered Node Architecture

In such an environment the user initiates agents that traverse the network and configure the nodes in order to achieve the best of the network's behavior during a particular task/application data flow.

The resources of the node can be accessed/controlled by visiting agents and according to the node's policy schemes. The node administrator is able to specify the policy of the node. According to this policy, agent's requests are denied or allowed. The authentication and authorization of the agent is based on the credentials it carries. Usually agent's contents are signed with the private key of its creator.

The agent enhanced node architecture can be seen in Figure 2.It consists of the following components:

An enhanced node (e.g. a programmable IP router). This enhanced node is accessed via an API for dynamic programming of its resources. It differs from the conventional nodes (e.g. a simple legacy router) in that its facilities are available for configuring/programming via a programmable node API to 3rd party entities. An entity can access and manage these facilities e.g. via a conventional SNMP, CMIP or even vendor-specific interfaces. Lately several router providers are developing such nodes that contain enhanced features such as QoS guarantees. Currently standardization efforts for open APIs are done via the IEEE P1520 working group [3] and Multi-service Switching Forum [4].

Agent add-on. This is a Mobile Agent Platform that provides the Execution Environments (EEs) and other services. Agents come to the EEs and execute. The router facilities are accessible from the agents via the agent platform. Agents not only have access to read the router's states or algorithms but also the ability to modify their behavior and therefore configure them according to their preference.

Agents that reside in EEs and via the facilities offered to them program the node. These agents can be mobile agents (e.g. the visiting agents) or even stationary intelligent agents that reside permanently on EEs offering services e.g. tuning of node's behavior according to their goals. Such stationary agents could be e.g. A resource management facility, a security facility etc

5. Place- Oriented VPNs

As discussed above each node hosts at least one agency. Based on these characteristic one could set-up and administrate a number of places in strategically located nodes, creating an art of virtual private network of places which we call PO-VPN (Place Oriented VPN).

A PO-VPN (Figure 3) is an alternative virtual private network. As in legacy VPNs it has its own resources and these are managed by its users in a policy related way. PO-VPNs offer numerous advantages over legacy VPNs and satisfy all the requirements listed in section 2.1. Current VPNs are not sufficiently flexible, have long set-up and deployment times, are not dynamic enough to accommodate rapid changes, don't respond to environmental changes, cant be controlled by the user, the software evolution is slow and creates problems if the VPN spawns different networks with different policies etc

PO-VPNs are built on top of existing physical networks and controlled by agents. They even can serve in the creation/deployment/management of legacy VPNs as we will explain later.

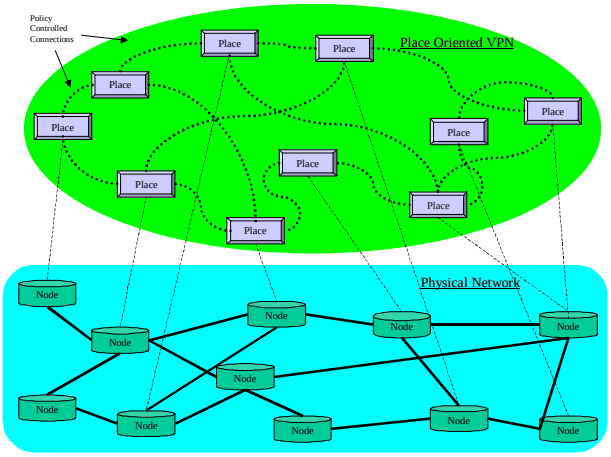

Figure 3. Place Oriented Virtual Private Network

5.1 Benefits and Satisfied Requirements

The above idea of place-oriented VPNs satisfies the requirements listed in section 2.1. Specifically we have:

5.1.1 Resource Management. Each place has its own resource manager or an interface to the agency's resource manager. Via that, the place administrator can manage the resources assigned to that specific place and cover the time changing needs of the inside executing agents. Visiting agents can also dynamically ask and get more resources if permitted by the policy. The resources assigned to a place are transparent (satisfies the virtual part of the VPN name) to the end-users of that place. These resources could be extended via CORBA interfaces and include resources also to other nodes. Then each place could have access to the resources of another place. This eases also the management issues as some resources e.g. disk space, could be managed as a total (the sum of disk space assigned to each place in the PO-VPN) and not individualistic per place.

5.1.2 Security, Privacy and Admission Control. PO-VPN is a virtual Private network, therefore security and privacy is a must in such an environment! Security cannot be an explicitly called service and should be always up-to-date (from the software point of view), flexible and dynamic. An extensive authentication/authorization scheme is supported for agents that want to execute in places of this alternative form of VPN.
The security architecture modeled in UML [14] in Figure 4 (extensive analysis in [8]) provides such a flexible scheme for agent systems. In short the components of the architecture are :

Policy Manager : It manages policies stored in the policy database. Each agent is signed by its user and has some rights concerning the accesses to the resources and services provided by a place. In the policy database those rights are specified either for specific agents or for groups of agents (e.g. agents belonging to a specific place) Both negative and positive policies enhance the user-friendliness and flexibility of the system

Credential Manager : The credential manager verifies the credentials of the agents and the code. A database serves for storing credentials (and avoiding time/computation expensive network connections to Certification Authorities). If a credential cant be verified locally (reasons: doesn't exist in the DB or the time of its validity has passed etc) then it is fetched via protocol e.g. LDAP [9] and it is verified. We use X.509 certificates although others could be used also e.g. PGP or SKIP, but these alternatives don't offer [10] any special advantages over X.509.

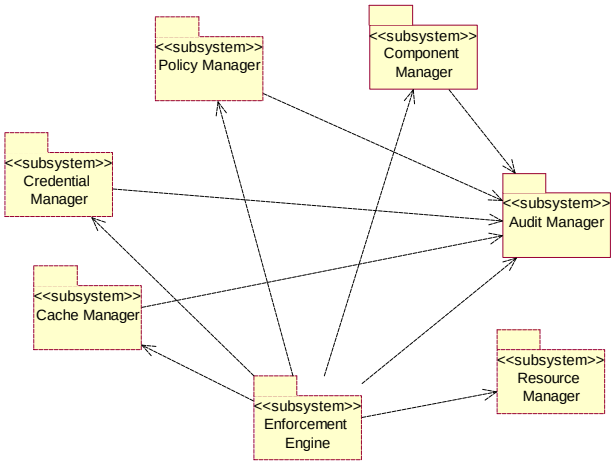

Figure 4. UML model of the security architecture

Component Manager : The component manager manages all requests concerning the preinstalled by the administrator components as well as the user components in the component database. The administrator can install components ,that are used by the above levels in order to construct services, and make it available to the users. Similarly if the policy allows it the users of a place can store in the node their code/components and use it (the agent doesn't have to implement it every time he visits a node), or even make it available to other users. Good examples are e.g. an encryption algorithm etc. The component DB is a repository of active code, algorithms and their implementations, protocols etc. By using this DB we guarantee that our offered services will use always the updated version of software, which in turn promotes security and lightweightness of agents.

Cache Manager : We use this component to cache security checks. In combination with the credential database, this component provides improvement into architecture's performance. The elements of cache DB are stored in a time-limited manner in order to avoid outdated info. Each time the policy changes, the security info accompanying the object that changed are deleted from the cache. Thus we can be sure that only up-to-date info is stored there and that we wont authorize a non-valid action.

Audit Manager : All actions are logged. 100% security cannot be guaranteed due to the multiple factors that interfere. By collecting the data generated by the network it is possible to analyze and trace back security breakouts or even prevent some intrusion attempts.

Enforcement Engine : The enforcement engine is used to enforce the policy based also on the results from the credential manager. It is the front-end interface of the architecture.

Resource Manager : As mentioned before this is the interface for configuring dynamically the resources assigned to a place. This is the job of the place administrator or of the authorized user.

The above approach guarantees a flexible and dynamic way to handle security decisions. Policies change at runtime adapting to intrusion attempts or change of interest, users can customize their environments and store code in the component DB with selective access rights etc. Furthermore the usage of SSL [13] provides a high level of security concerning external communications. Issues such as data integrity, secrecy,

authentication etc are covered by the usage of X.509 certificates.

Based on policy and the available tools (cryptographic algorithms in component DB) one can specify a wide range of access rights including place to place communication, change of node's behavior etc. The security in agent technology is a hot research area and new solutions pop up everyday. These future solutions could be easily imported in an open architecture as the one presented above.

5.1.3 Information path selection. The administrator can select the network topology he desires by selecting and owning places in strategically selected nodes. Therefore he can select via which network path his agents or the data flow passes. The end user also, can specify in his agent to execute only in specific places and exchange info with the environment under special conditions. So the user can have some control over the messages he sends or the data he exchanges via his agents.

5.1.4 Communication. Users communicate via the messages their agents exchange. Thus there can be a policy specific framework that defines with who a user is allowed to communicate/interact and which other users to ignore. Furthermore the agents are capable of updating the software on demand in the underlying infrastructure as well as implementing new protocols. So in a scenario two users would like to communicate via a specific protocol that is not available by the service provider. Then the agents implement that protocol and the users can communicate. Even in the occasion that two users use incompatible communication end-points agents could be used as wrappers and provide the missing connectivity link between these two users.

5.1.5 New Service deployment. As noted above agents can be used for the implementation of new protocols that don't exist or are not standardized. Furthermore agents can be used to deploy new services. By updating the underlying infrastructures components on demand and by sampling their services, agents are able to provide more sophisticated customized services available to other agents or a 3rd party application e.g. a multimedia filter. The service provider has some filters stored in the Component database. Agents can use these pre-installed filters with a specific flow. We suppose that the user wants to implement a new non-existent filter and even reuse some of the existing ones to provide a customized version that fits his needs. Then he can program an agent that would sample the abilities of the provider's filters and also add some customized extensions. Then the agent would be permanently positioned (stationary agent) within a place and filter incoming flow based on the new code. Furthermore clone agents are transmitted to other nodes (according to user's preferences) and filter the flow there as well. That results into maximum flexibility as the user is in control, no to mention the low network utilization and the customization of the flow within the network and at the end-points.

5.1.6 Effectiveness and performance. As PO-VPNs are user controlled/customized and dynamically adapt to the environmental changes, they provide flexibility not only to the end-users but also to the network itself. The network administrators can tune-up flows and nodes to operate in an optimal way. Examples are : compressing the flow between nodes in case of low capacity of the line, provide alternative routes in case of malfunction, update the security components of the node according to security bulletins etc. By providing an application/task specific optimization and adapting the network to these requirements we can obtain optimal use of the underlying infrastructure and boost performance.

5.1.7 Cost matters. Current VPN infrastructures may provide adequate solutions to big enterprises that can afford it. Unfortunately this leaves out SMEs as well as smaller groups of people that look for something with minimal cost. PO-VPNs provide this opportunity. As they can be set-up/deployed over the current physical networks and be customized dynamically by the user, they provide the long-wanted combination of effectiveness and cost. Multiple PO-VPNs can reside even in nested forms. Since the agreements are done by agents, long deployment times and re-configuration requests are handled automatically by the system without the need of human intervention.

5.1.8 Flexibility, Adaptivity. Flexibility is pushed to the maximum by the use of underlying offered services and the ability to dynamically make changes to the configuration of the PO-VPN and the services it offers. Stationary agents that reside on the place not only offer their services but also respond to the environment changes - which may be unpredicted - by reconfiguring or updating node's components. The node is not any more static but an responsive one as it can react and interact with the environment and the users in an automatic way which can also handle non-deterministic events and therefore adding intelligence, robustness and fault-tolerance to the whole infrastructure.

5.1.9 Interoperability. Current network infrastructures are heterogeneous both in hardware and in software matters. PO-VPNs based on agent technology and agent

themselves, are computer and transport independent (they depend only on the execution environment installed to the node) and therefore promote interoperability among systems and software. It is possible with agents to implement interactions with any legacy system and currently existing services and make it available to other heterogeneous agents e.g. via MASIF interface. Although the ground is pretty new, standardization efforts exist within many organizations e.g. Object Management Group [12], Foundation for Intelligent Physical Agents [7] and will be available in the near future.

5.2 Who can use PO-VPNs?

We try here to identify some groups that could benefit from PO-VPNs. Please note that only some areas have been targeted.

- Low-populated groups or even individuals that can't afford to build their own VPNs over a physical network. On the contrary they might find it easier to get a taste of an alternative style of VPNs that may suit their purposes and satisfy their needs. Furthermore even large organizations with a greater investment ability are interested in cost-effective solutions. Because PO-VPNs are easy to create, maintain and administrate they have numerous advantages over traditional VPNs.

- Companies that have a heterogeneous infrastructure and decide to co-operate for a project lifetime. Those companies already have their VPN and don't intend to change their infrastructure. By building PO-VPNs on top of their infrastructure they can use the advantages mentioned before (in 5.1) with their business partners in an immediate and cost-effective way.

- Groups that have requirements that change unpredictably by the time. This kind of groups are interested in a flexible management of VPN's allocated resources and also to immediate reaction to an unpredictably changing environment. Such groups can't use a static version of a VPN as they would have to lease a VPN that satisfies the highest overall amount of each of their requirements. But that results in a time, computation and cost expensive solution which at the end may not be adequate due to an unpredicted change (e.g. the bandwidth need required is higher than originally thought). In such cases a PO-VPN offers the flexibility to dynamically change the configuration of the network and its resources without human intervention.

- Potential users are also those groups whose lifetime is short and usually determined by some other external events. Such groups need to set-up and delete a VPN in minimal time. E.g. an auction is held in a specific date and time in a place and the interested agents would like to interact under high security conditions in the trusted environment of the provider. This flexibility to create and teardown such a virtual environment can be achieved by PO-VPNs but by no way in normal static legacy systems.

PO-VPNs are an alternative form of VPNs. Yet there is a connection between this form and the legacy VPNs. Since agents and their execution environments offer the advantages described above, it seems a pretty good idea to install such an environment within a router and control the functions of the router via agents. So places in a PO-VPN are actually turning to control/management places for legacy VPNs. Again all interactions with the system and the administrators are done in an automated manner via agents, once of course the human agreement exists. So agents again can add some more flexibility in the domain of negotiation, configuration and service management of the old-style VPNs. PO-VPNs are also a good candidate as an service/application for multiagent systems [6] area, as they are scalable, handle software evolution, promote open systems and are by their nature inhabited by multiple and diverse agents.

PO-VPNs cannot replace legacy VPN systems. They are not able to provide the full advantages of a legacy VPN. Though they can be installed on top of current infrastructures and provide task specific solutions for group working. Furthermore they can provide enhancements and rapid cost-effective solutions to specific problems and needs even per user. PO-VPNs should be seen as a value added service to the current static and low evolving infrastructures.

5.3 Dynamic VPN provisioning Scenario

Here we will try to show how VPN provisioning can be done in a dynamic and flexible way, allowing the deployment of VPN (both the legacy ones and the alternative PO-VPNs) in minimal time. Agents here take the roles of the customer, the service provider and the network provider. The agents are signed with their users' private keys therefore authenticating that their users delegated them with the specific action. There are multiple network and service providers in the infrastructure on top of which we want to build our VPN. We assume that the customer agent (CA) interacts with the VPN service provider agent (SPA) and with the network provider agent (NPA). The CA generally has to

negotiate with all NPAs, but if a specific service provided by a service provider spawns multiple domains then the service provider makes the negotiations concerning the service with those domains. Therefore the CA finally can negotiate with the rest of NPAs that are not covered by the SPA.
The stages to follow are :

VPN Network Negotiation Stage :
In this stage the user has an idea of possible topologies of his future VPN. Now he or his agents contact various VPN network providers requesting the necessary resources or other special features concerning customer's requirements.

VPN Service Negotiation Stage : In this stage the user has multiple network topologies that suit him. Now he negotiates with a service provider in order to see which of the possible topologies support the services he desires. Have in mind that when a service spawns multiple domains then the negotiation for these domains is done by the service provider and not by the user itself. The service provider can also offer alternative solutions concerning the network topology etc. At the end of this stage the user has some network topologies that support the services he requires and a set of service level agreements (SLAs).

VPN Selection of the Final Network Topology : The CA either returns and reports the possible solutions to the initial goal of a VPN creation or if it is intelligent decides by itself for the best of the offered solutions. The final decision is made based on various facts e.g. connectivity bandwidth, error statistics, reputation of nodes contained in the topology, cost, security etc either by a human or an intelligent agent. Finally a network topology is selected and the final phase of agreement has to follow.

Deployment of the VPN and its services : The user requests that service and network providers set-up the services and the network connections. When everything is done the users of the VPN are informed that the VPN is ready for use. Management issues of the VPN are dealt by the responsible entity (e.g. VPN's administrator) for the whole life of the VPN. When the VPN is terminated then the service providers and the network providers will charge the responsible entity for the use of the services and the network. respectively. We have already said that a service might spawn various network nodes. In that case the network providers charge the service providers which in their turn add additional charges for the offered services and bill the VPN owner or the end user of the system.

The important thing is that via agents all can be configured automatically and PO-VPNs can be constructed or tear apart at minimal time and cost without human overhead. In the future a toolkit could automate further the whole process of provisioning, set-up and configuration of the VPN with minimal human intervention, probably only to the final stage or not at all. Further toolkits could be implemented and provide network applications with info on per node or even per place basis so that the applications can take optimal advantage of the environment (state, services, resources etc) and configure it to their task specific needs.

6. A Broker Application Scenario

With a broker scenario we will try to demonstrate the advantages the PO-VPNs offer. A broker office wants to connect people interested in buying or selling shares. Except from buyers and sellers, members of the community are stock exchange specialists (the advisors) that give their professional advice on investment issues. Other elements of the VPN except its members are the places where the auctions of shares take place, places where investors and/or customers could have private conversations etc.

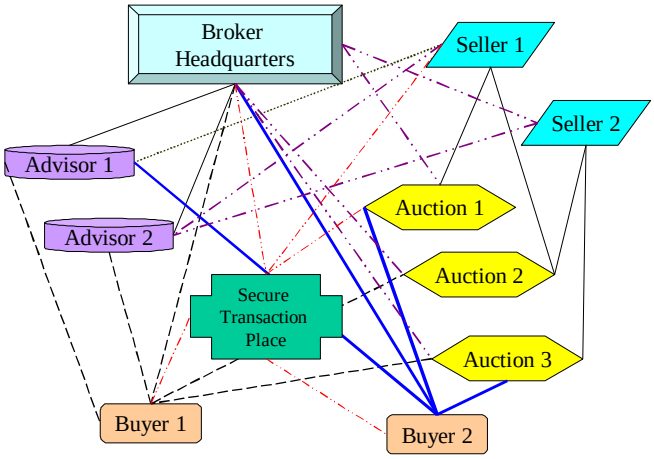

Figure 5. A Broker PO-VPN

With those requirements in mind the broker creates a PO-VPN connecting places of all entities mentioned before in various hosts (Figure 5) . These places could be assigned to buyers/sellers/advisors or to auctions e.g. a place for trading Internet shares, another for blue chip companies etc. The buyers and sellers communicate via the auction places set up in the PO-

VPN. Furthermore the broker provides one or more Secure Transaction places (STP). These places offer advanced security (e.g. verification of identities, check of certificate revocation lists, secure communications etc). So if a real transaction has to take place both sellers and buyers can do it in this STP, where the broker has the control and gives the security guarantees.

The broker is also able to personalize his services based on the client's needs. E.g. if the client has a special interest on internet shares or IPOs (Initial Public Offerings) then broker's agents could inform the client of the new available offer each time and where the trading will take place, or if the price of a share has reached the limits set by the client, then the client could be warned in order to take some action. Other services offered could be e.g. a secure environment to buy and sell shares, high security places for exchanging info or even providing data to selected customers from the professional advisors. All these require privacy and controllable access which is handled by the security architecture.

In another scenario a new customer wants to join the broker's VPN and take advantage of the offered services. Then an extra place could be added to the PO-VPN (that of the customer's) where collaborative customer agents reside and process the info and services offered by the broker. Although the broker offers services to its customers he forbids direct connections between sellers and buyers (because that would cut off the provision he takes). Thus by place policy it is forbidden that selling and buying agents that belong to different customers are in the same place. If a trade has to be done then it is done only via broker's agents or in STPs who beyond controlling the transaction fulfill specific security requirements (e.g. based on the credentials and history of the person to which the agent belongs).

The whole environment is protected from external intrusions since agents and places are authenticated (both hold electronic certificates). By offering such services (personalization, security, up to date info etc) to its customers the broker is able to reach bigger market share. Traditional VPNs can't be considered as a practical alternative to such short-living groups or actions, nor to the dynamicity of the system (which may add/delete/modify places, policies and agents on demand) as they change infrequently and slowly.

7. Summary and Conclusions

The communication infrastructures that exist today are too expensive for wide exploitation by small groups that connect/cooperate over Internet. We have presented here the notion of PO-VPNs, an alternative form of VPNs that tackle successfully this kind of problem. After introducing and analyzing a number of requirements on group working, we presented at high level a feasible and cost-effective solution based on agent technology. Place Oriented Virtual Private Networks' aim is to provide an infrastructure for groups with special characteristics such as low-population, short time to live etc and in parallel promote sophisticated service deployment. PO-VPNs can be deployed on top of current networks (Internet, VPNs) and offer their customized services in a flexible, interoperable and cost effective way.

8. References

[1] "What is a VPN?", Paul Ferguson - Cisco Systems, Geoff Huston - Telstra Internet, April 1998.

[2] "Towards an Active Network Architecture", David L. Tennenhouse and David J. Wetherall. Keynote session of Multimedia Computing and Networking, San Jose, CA, January 1996.

[3] IEEE P1520 Project Web Site : http://www.ieee-pin.org/

[4] Multi-service Switching Forum web site http://www.msforum.org/

[5] Cetus Links on Mobile Agents : http://www.cetus-links.org/oo_mobile_agents.html

[6] "Multiagent Systems on the Net", Anupam Joshi and Munidar P. Singh, Communications of the ACM, March 1999/Vol 42, No 3, p. 39-40

[7] FIPA Web Site: http://www.fipa.org/

[8] "Agent Based Security for the Active Network Infrastructure", S. Karnouskos, I. Busse, S. Covaci, Active Networks - IWAN'99, Lecture Notes in Computer Science vol. 1653, Berlin, Germany, June-July 1999, ISBN 3-540-66238-3.

[9] Lightweight Directory Access Protocol (LDAP v3), RFC 2251. URL: http://info.internet.isi.edu/in-notes/rfc/files/rfc2251.txt

[10] "Overview of Certification Systems: X.509, CA, PGP and SKIP",E. Gerck, Meta-Certificate Group, Novware Softex/Unicamp Brazil.

[11] IETF RFC1825, "Security Architecture for the Internet protocol", R. Atkinson, August 1995.

[12] OMG Web Site : http://www.omg.org/

[13] IAIK-SSL Implementation. URL : http://jcewww.iaik.tu-graz.ac.at/IAIK_JCE/jce.htm

[14] Unified Modeling Language, Rational Software, URL : http://www.rational.com/uml/