Introducing Virtual Private Overlay Network services in large scale Grid infrastructures

Francesco Palmieri
Federico II University, Napoli, Italy
Email: Francesco.Palmieri@unina.it

***Abstract*—The computational Grid concept is gaining great popularity as the best way to deliver access to a wide range of distributed computing and data resources. But, as Grids move from an experimental phase to real production and their deployment in the Internet significantly increases, controlling the security of a Grid application becomes imperative. The most significant Grid security issue is that the different sites composing the Grid will generally be managed by different organizations each with their own security mechanisms and policies. This makes any communication security arrangement on the entities participating to the Grid generally more difficult than if they were on the same LAN. In this paper, we propose a novel network resource abstraction for delivering dynamic on-demand Virtual Private Overlay connection services, into large-scale Grid environments. Such facility provides to Grid applications an illusion of dedicated layer-2 LAN connections that are fully comparable to a private network in performance, reliability, security and Quality of Service, but also provide topology and control plane virtualization to ensure better isolation also at the protocol and address space level. It may represent a technological breakthrough that can transform the overall connection paradigm in modern Grids, by reducing infrastructure costs, with the elimination of private circuits and long-distance direct connections, and increasing network coverage and flexibility by leveraging the Internet usage. As a proof of concept, the proposed facility has been implemented in a Grid Information Service prototype which was successfully tested on a small dedicated testbed infrastructure.**

***Index Terms*— Grid computing, security, MPLS control plane, Layer-2 VPN**

I. INTRODUCTION

In recent years, with the overwhelming success of the Internet, the landscape of computing and telecommunications is radically changing and the Grid technology is increasingly being looked upon as a natural application of the modern Internet for engaging in complex data processing tasks over resources which are distributed across the world. A typical Grid, consists of a large number of geographically distributed computing and storage resources (e.g., supercomputers, computer clusters, storage systems, data sources, instruments), usually spanning multiple administrative domains, interconnected through an high performance network, to be shared amongst its users as an aggregated, unified facility for supporting large-scale and data-intensive computing applications (e.g., molecular modeling for drug design, brain activity analysis, and high energy physics). Large computing endeavors (consisting of one or more "jobs") are then distributed over this network to these resources, and scheduled to fulfill requirements with the highest possible efficiency. A Grid offers a uniform and often transparent interface to its resources such that an unaware user can submit jobs to the Grid just as if he/she was handling a large virtual supercomputer. Recently, the Grid concept has been generalized to cover any virtual organization, defined as a dynamic collection of individuals and institutions which are required to share resources to achieve certain goals [1]. Thus the Grid will have applications in commerce and industry, supporting distributed collaborative design and engineering, or supporting distributed supply chains. Nevertheless, any distributed computing platform, including grids, needs to satisfy specific and often strict security and Quality of Service (QoS) demands. Without an adequate understanding of the security implications of a Grid, both the users and the system administrators who contribute with resources to a Grid can be subject to significant compromises. Thus the importance of data and application security issues assumes critical proportions as more and more industry and academic interests channelize their resources towards implementing such cross organizational computing infrastructures. However, existing approaches to security within distributed systems, usually based on access control policies enforced by firewalls or other kinds of packet filtering devices such as routers or layer-3 switches are stretched by the extreme conditions imposed by the modern Grids, and significant effort has been undertaken in, to provide support for secure use of resources without affecting the overall Grid functionality or computational efficiency. What clearly distinguishes grids from other platforms are its high dynamicity and complexity features, in terms of communication paradigms and protocols used, resulting in security requirements which cannot be addressed by existing access control technologies for distributed

Based on "Dynamic layer-2 VPN services for improving security in the Grid environment", by Francesco Palmieri which appeared in the Proceedings of the International Conference on Emerging Trends in Information and Communication Security (ETRICS 2006), Freiburg, Germany, June 2006, LNCS 3995/2006 pp. 45-59, Springer Verlag.

platforms. The elements of a grid are usually negotiated in a dynamic manner such that the trust relationship among these elements needs to be established during application execution time. There may not at all exist any direct security protocol among resources and processes which form this dynamic environment. Each resource belongs to a fixed administrative domain governed by its own security standards, policies and implementation within the domain. Grids that span several administrative sites and encourage the dynamic addition of resources are not likely to benefit from the security that static, centrally administered commercial firewalls or packet filtering routers provide. What is needed are some facilities that, while ensuring adequate end-to-end security features in terms of authentication, integrity and traffic isolation offer a totally dynamic and scalable "LAN extension" abstraction, so that as new resources are attached to the grid they can behave as belonging to the same LAN, without any apparent security concern. In this scenario, on-demand Layer-2 VPN technologies can be successfully applied to Grids, as they, offer all the above security features and can help to transparently bypass firewalls or any other filtering policy in order to prevent the performance and functionality penalties that may typically negatively affect high-end applications. Dynamic provisioning is needed in order to reduce management costs together with the number of Grid VPNs that the public networks have to support concurrently. The dynamicity relies on the availability of a suite of interfaces and protocols which perform discovery of available services, agreement negotiation and agreement establishment between initiators (the Grid user or proxy) and providers (e.g. Grid resource brokers). Of course, a control plane – capable of establishing, managing and tearing down services – is necessary for the actual provisioning of the service. Furthermore, an abstraction of such control plane services should be transparently made available to Grid applications in a way that is totally independent from the underlying network protocols and communication technologies. Here MultiProtocol Label Switching (MPLS), that has been now deployed to implement traffic engineering facilities on almost all the modern transport infrastructure making the Internet core, offers the essential features by providing the proper end-to-end label switched tunnels that will be useful to implement effective and flexible layer-2 VPN abstractions and services. Such VPNs, actually constitute the cornerstone for delivering innovative secure and reliable connection services to both large-scale Grid applications and data management/file transfer facilities which rely on Internet-based connectivity and storage access protocols not providing security and privacy. In fact, VPNs can support confidentiality and integrity by means of data isolation, i.e. by separating in intermediate forwarding devices the control and forwarding plane, the signaling and the routing information of each VPN. Layer-2 secure connections realized through VPN technology can also be used to dynamically cluster geographically dispersed resources belonging to the same Grid Virtual Organization. Clearly, in order to effectively use the above facilities in large-scale Grids (e.g., to be capable to address an increasing number of users ubiquitously), flexible, stable, scalable and QoS-aware layer-2 virtual private connection services are necessary. Accordingly, in this paper, we show how the security and privacy services offered by scalable on-demand layer-2 MPLS VPN services and the native MPLS traffic engineering facilities can be combined and successfully applied in large-scale Grid scenarios. Accordingly, we propose a novel network resource abstraction for implementing and managing on-demand Virtual Private Overlay Networks, providing an illusion of dedicated layer-2 connections that are fully comparable to a private network in performance, reliability, security and Quality of Service (QoS), but also provide topology and control plane virtualization to ensure better isolation also at the protocol and address space level. It has been implemented in a Grid Information Service prototype which was successfully tested on a dedicated testbed infrastructure. The paper is organized as follows: section 2 briefly sketches the basic background concepts behind the whole framework while section 3 and 4 respectively present the main security requirements and QoS needs in the Grid environment. The detailed components of the whole proposal are described in section 5. Finally, section 6 is dedicated to conclusions and final remarks.

II. BACKGROUND CONCEPTS

This section briefly introduces some of the basic concepts that will be useful to better explain the proposed Grid VPN paradigm, by presenting its architectural building blocks, ideology and the theory behind it.

A. Network Resource Virtualization

Resource Virtualization refers to the ability to abstract multiple instances of physical resources in an aggregate and uniform virtual resource view. The ability to virtualize the network resources is essential to provide true security, in terms of traffic and address space isolation, and QoS guarantees to Grid Applications. This is analogous to the manner in which the Grid Middleware provides a virtual view of the available computational and storage resources, such as CPU and memory, to each user. Emerging trends in network resource virtualization include simple Virtual Private Networks and Virtual Private Overlay Networks, the former being actually the more standardized of the two and the latter still being a subject of research studies. By pure definition, a virtual private network [2] is the interconnection of multiple sites through a set of circuit-switched paths, or virtual connections, thereby skirting security and performance issues of the very public Internet. Each virtual connection in the VPN provides an illusion of a dedicated data-path between two remote VPN endpoints. In reality each virtual connection traverses multiple physical links and switches that are shared with other traffic streams. VPN endpoints do not perform traffic forwarding or routing and act solely as traffic producers and consumers. The concept of VPN is not new, and legacy technologies such

as ISDN, Frame Relay or ATM have been used over the last decades as a basis for the implementation of this concept. Whatever the format or the technology behind it, a VPN provides a service functionally equivalent to a private network using resources of a public network such as the Internet. The VPN should be comparable to a private network in performance, reliability, management security and Quality of Service (QoS). Virtual Private Overlay Networks (VPON) provide the next level of resource abstraction where, in addition to the data-path virtualization, control plane is also virtualized. Figure 1 below shows a VPON in which several nodes in the network are interconnected by means of virtual connections. Each VPON has its own virtual topology, control plane (such as its routing protocols), and data plane (such as traffic processing and forwarding), thus the participating nodes have a finer level of control over their traffic streams instead of merely producing and consuming traffic, as in the case of VPNs. Each node in a VPON has a specific address and is connected by virtual connections to other participant nodes. Further, every node can participate in control plane activities such as intra-VPON routing protocols. In a Grid scenario, a VPN can be better conceived as a specific Grid service with its own control plane features (hence a VPON), in which customer connectivity amongst multiple sites is deployed on a shared infrastructure with the same access or security policies as a private independent network.

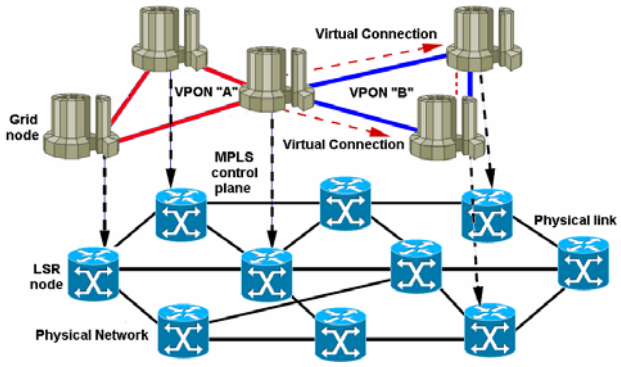

Figure 1. The Virtual Private Overlay Network scheme.

B. *The MPLS paradigm*

MPLS is a packet forwarding technique being standardized by IETF [3] that is actually considered the most promising technology for implementing the above network-based VPONs. MPLS uses labels to make forwarding decisions at the network node level, in contrast to the traditional destination-based hop-by-hop forwarding in IP networks. The key idea of MPLS is a strict separation between control and forwarding planes in the network functions as well as in the software and hardware architecture of the routers. In MPLS, the space of all possible forwarding options in a network domain is partitioned into "Forwarding Equivalence Classes" (FECs). For example, all the packets destined for a given egress may belong to the same FEC. The packets are labeled at the ingress depending on the FEC they belong to. Each of the intermediate nodes uses the label of incoming packet to determine its next hop, and also performs "label swapping," i.e., replaces the incoming label with the new outgoing label that identifies the respective FEC for the downstream node. Such a label-based forwarding technique reduces the processing overhead required for routing at the intermediate nodes, thereby improving their packet forwarding performance. Also, the label-merging procedure used by MPLS creates multipoint-to-point packet forwarding trees in contrast to a routing mesh in conventional network based on a similar paradigm such as ATM networks. This reduces considerably the size of forwarding table at the intermediate nodes, thereby improving their scalability. The MPLS encapsulation envelope is shown below.

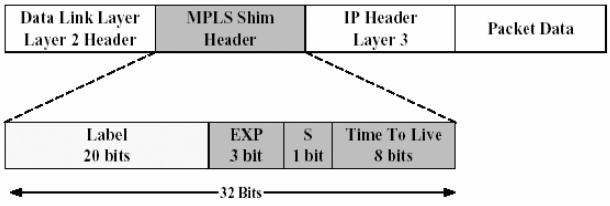

Figure 2. MPLS shim header.

While MPLS was originally conceived to improve the efficiency of packet forwarding in network equipments, it was soon realized that it could also provide other advanced features, such as Traffic Engineering and Virtual Private Networks capabilities. Both these facilities need predetermined paths to be established through the network to specific destinations. Once the paths, called label switched paths (LSPs), have been created, traffic is mapped onto them according to the dynamic needs of the traffic and their capabilities. The LSPs can thus be used to implement explicit virtual connections on the underlying transport network supporting precise QoS and traffic isolation constraints. MPLS introduces a circuit-switching paradigm on top of the basic packet-switching framework of the Internet. In fact it has been argued that the circuit-switching paradigm will become increasingly prevalent in the future, with core of the Internet being mainly circuit-switched and IP based packet-switching mainly thriving at the edges to provide best effort services. The fundamental advantage of circuit-switching paradigm is that it enables performance isolation between traffic streams that belong to different virtual connections - something that packet-switching by itself cannot guarantee. By performance isolation, we mean that the involved internet service providers (ISP) can prevent the performance of one virtual connection from being effected by a traffic stream belonging to another virtual connection. The ISP can provide a QoS guarantee to each virtual connection that is independent of other traffic streams sharing the physical network. Presence of QoS guarantees (such as average bandwidth, end-to-end delay, delay-jitter, loss rates or reliability) on virtual connections enables deployment of real-time services such as Voice over IP (VoIP), video conferencing, media streaming and real-time computing.

C. Layer 2 VPN services in the MPLS environment

Customers of the VPN services use shared facilities and equipment, which are managed, engineered and operated by a public network operator, either totally or partly. Traditionally, the most common way for cooperating organizations to build their own wide area networks was to set up a private communication infrastructure on top of a number of point-to-point or point-to-multipoint links (based on virtual circuits on a public switched communication service such as Frame-Relay or ATM) provided by a service provider. This model corresponds to what is usually known as "layer 2 VPN". Although layer 2 VPNs based on ATM or Frame Relay have been extensively deployed, several drawbacks related to this kind of VPN can be identified. First, the service provider VPN infrastructure is dependent on a single layer 2 technology (e.g., ATM, Frame Relay). In addition, the Internet infrastructure and the VPN infrastructure, even if they share the same physical network, need separate administration and maintenance. Finally, provisioning is difficult – for example, adding a site to an existing VPN is usually a complex task. Consequently, it can be easily seen that the above solution lacks of the sufficient scalability that is an essential prerequisite for all the new-generation services offered on the modern Internet. Furthermore, to offer the abilities required to establish a layer-2 virtual circuit between any two computers or clusters in a Grid environment, bandwidth allocation and management on the network must be dynamic. MPLS control plane protocols allow large-scale transport networks to be created and enable these networks to respond to on-demand requests for rate-guaranteed connectivity between multiple points in the network. These features make MPLS-based networks well suited to serve Grids supporting the realization of bandwidth guaranteed on-demand VPONs. The idea of transporting generic Layer 2 protocols over MPLS backbones has introduced the concept of the so-called "Layer 2 VPN" over MPLS. An MPLS-based layer 2 VPN allows the use of a single MPLS-based network infrastructure to offer a wide range of services, including IP traffic, layer 2 VPNs, layer 3 VPNs, MPLS traffic engineering and DiffServ-based QoS control. Easy migration from traditional layer 2 VPNs is a significant advantage of this model, as the two VPN types are indistinguishable from the customer's point of view. Here the VPN service is functionally equivalent to emulated leased lines and the service provider and the customer do not exchange layer 3 routing information. This model provides a clear separation between the customer's and provider's responsibilities. Basically, in an MPLS-based Layer 2 VPN the service provider uses an MPLS network to provide layer 2 services to the customer. The interior of an MPLS infrastructure on which VPN services are offered is made up of MPLS-aware provider (P) router devices forming the MPLS core that are not directly connect to any VPN-terminating router. Provider edge (PE) routers that surround the core devices enable the VPN functions of an MPLS network. MPLS core and PE routers work as label switch routers (LSR) that are devices capable of switching packets based on their MPLS-imposed labels. The VPN-terminating router is referred to as a customer edge router (CE) and thus a VPN consists of a group of CE routers connected to the MPLS backbone PE routers [4]. Only the PE routers are aware of the VPN. The CE routers are not aware of the underlying network. The CE routers perceive that they are connected via a private network. From the customer's point of view, a layer-2 MPLS VPN is exactly the same as a layer-2 VPN, with layer-2 circuits interconnecting the various sites. For example, a customer CE device may be configured with a Frame-Relay Data Link Connection Identifier (DLCI) on which to transmit to other CEs through the provider network, which appears as a traditional layer-2 cloud to the users. Within the service provider network, the layer-2 packets are transported in MPLS LSPs. The service provider does not participate in the customer's Layer-3 network routing. The establishment of emulated VCs, also called Virtual Leased Lines (VLL), or layer-2 point-to-point connectivity across an MPLS backbone is specified in the IETF drafts usually known as "drafts martini" [5] and [6]. These drafts define how MPLS can be used to support Layer 2 protocols such as Ethernet, Frame Relay or ATM. The first draft [5] concentrates on encapsulation methods, while the other [6] specifies signaling to set up point-to-point layer-2 circuits over an MPLS network. The following figure represents an example of an MPLS-based layer 2 VPN [7]. The connection between two customer's CE devices is composed of three segments: two CE-PE "attachment" VCs and one emulated VC in the core. The routing tables of the source CE router and the ingress and egress PE routers are indicated. Basically, the first CE router forwards the traffic to DLCIs 600 and 610 to sites B and C respectively, as in a normal Frame Relay network, whereas the ingress and egress PE routers perform the mapping between the DLCIs and the appropriate LSPs.

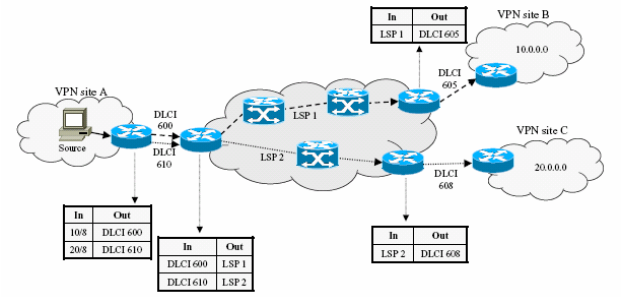

Figure 3. The MPLS Layer-2 VPN paradigm.

It should be noted that MPLS layer 2 VPNs make provisioning much easier in comparison to conventional layer 2 VPNs. In particular, adding a site to an existing VPN should simply require the configuration of the PE router connected to the new site, and not the reconfiguration of a high number of CEs. The IETF draft "An architecture for L2VPNs" [8], proposes a layer 2 VPN solution, which is based on the emulation of layer 2 circuits. In the service provider core, tunnels are

established using a proper tunneling technology (usually MPLS, but L2TP or IPSec should also be possible) to emulate layer 2 VCs. This draft can be seen as an evolution of a previous draft, called "MPLS-based Layer 2 VPNs" [9], now obsolete, which originally described how to build layer 2 CE-to-CE VPNs using MPLS in the provider core. The draft [8] is based on the "drafts martini" indicated above for encapsulation of data frames and for the signaling used to setup and maintain the emulated VCs. The need to specify an auto-discovery mechanism is indicated but no solution is proposed for the time being. Recently, the PPVPN IETF group has reutilized the VPLS concept (Virtual Private LAN Service, following a term originally defined in RFC2764 [10]) as a layer 2 service that emulates a LAN across a WAN [11][12]. The basic purpose of a VPLS is to offer layer-2 connectivity to multiple customer sites in a manner that is transparent to the CE devices. The service provider is responsible for transporting customer Layer 2 frames and switching them across the service provider network between customer sites. From the customer's point of view the service is equivalent to connecting the CE devices via a switch, i.e., all in the same broadcast domain/LAN segment.

III. Security requirements in the Grid environment

Available Research and development efforts within the Grid community have produced protocols, services, and tools that address the challenges arising when we seek to build scalable virtual organizations. What distinguishes a virtual organization is that it may gather individuals and/or institutions that have agreed to share resources and otherwise collaborate on an ad-hoc, dynamic basis, while they continue to belong to different real organizations, each governed by their own set of internal rules and policies. This poses a challenge when combined with the fact that an individual or institution may be a member of several virtual organizations simultaneously. From a security point of view, one is thus confronted with the need of policies and protection domains that may superpose, straddle, conflict and intersect one another in many different ways. Really, distributed computation and security, apparently seem to be two topics in direct conflict. Recently, with the aim of migrating grid applications from the localized LAN or MAN scenario to global Internet-based grid computing, the need for security guarantees is forcing also the Grid hosting organizations to implement stronger security hardening configurations on their firewalls or border routers, which prevent some of the communication functionality needed by many distributed applications. This can lead to debates over which functionality is more important. Within this context, we require harmonization and interoperability among protection domains while maintaining a clear separation of the security policies and mechanisms deployed by both virtual and real organizations.

The security challenges faced in a Grid environment can be grouped into four categories: dynamicity, integration with existing systems and technologies, interoperability with different "hosting environments" (e.g., J2EE servers, .NET servers, Linux/Unix systems), and trust relationships among interacting hosting environments.

A. Dynamicity

One of the aims of a grid is to enable the sharing of vast amounts of distributed resources within large, dynamic and distributed communities of users, where the availability of resources, membership of communities (or virtual organizations) and access rights are continually changing and evolving. A grid is expected to provide an architecture that enables such a dynamic structure. These changing patterns of use add considerably to the already great challenge of allowing controlled access to remote resources owned and managed by third parties: issues of trust and liability become very important.

B. Integration

For both technical and pragmatic reasons, it is unreasonable to expect that a single security technology can be defined that will both address all Grid security challenges and be adopted in every hosting environment. Existing security infrastructures cannot be replaced overnight. Each domain typically has its own authorization infrastructure that is deployed, managed and supported. It will not typically be acceptable to replace any of these technologies in favor of a single model or mechanism. Thus, to be successful, a Grid security architecture needs to step up to the challenge of integrating with existing security architectures and models across platforms and hosting environments. This means that the architecture must be implementation agnostic, so that it can be instantiated in terms of any existing security mechanisms (e.g., Kerberos, PKI); extensible, so that it can incorporate new security services as they become available; and integrable with existing security services.

C. Interoperability

Services that traverse multiple domains and hosting environments need to be able to interact with each other, thus introducing the need for interoperability at multiple levels:

- At the protocol level, we require mechanisms that allow domains to exchange messages. This can be achieved, for example, via SOAP/HTTP.
- At the policy level, secure interoperability requires that each party be able to specify any policy it may wish in order to engage in a secure conversation - and that policies expressed by different parties can be made mutually comprehensible. Only then can the parties attempt to establish a secure communication channel and security context upon mutual authentication, trust relationship, and adherence to each other's policy.

At the identity level, we require mechanisms for identifying a user from one domain in another domain. This requirement goes beyond the need to define trust relationships and achieve federation between security mechanisms (e.g., from Kerberos tickets to X.509

certificates). Irrespective of the authentication and authorization model, which can be group-based, role-based or other attribute-based, many models rely on the notion of an identity for reasons including authorization and accountability. It would be nice if a given identity could be (pre)defined across all participating domains, but that is not realistic in practice. For any cross-domain invocation to succeed in a secure environment, mapping of identities and credentials must be made possible. This can be enforced at either end of a session through proxy servers or through trusted intermediaries acting as trust proxies.

D. The Trust Relationship

Grid service requests can span multiple security domains. Trust relationships among these domains play an important role in the outcome of such end-to-end traversals. A service needs to make its access requirements available to interested entities, so that they can request secure access to it. Trust between end points can be presumed, based on topological assumptions (e.g., in our case, a VPN), or explicit, specified as policies and enforced through exchange of some trust-forming credentials. In a Grid environment, presumed trust is rarely feasible due to the dynamic nature of the virtual organization relationships. Trust establishment may be a one-time activity per session or it may be evaluated dynamically on every request. The dynamic nature of the Grid in some cases can make it impossible to establish trust relationships among sites prior to application execution. Given that the participating domains may have different security technologies in their infrastructure (e.g., Kerberos, PKI) it then becomes necessary to realize the required trust relationships through some form of federation among the security mechanisms. The trust relationship problem is made more difficult in a Grid environment by the need to support the dynamic, user-controlled deployment and management of transient services. End users create such transient services to perform request-specific tasks, which may involve the execution of user code. For example, in a distributed data mining scenario, transient services may be created at various locations both to extract information from remote databases and to synthesize summary information.

IV. THE NEED FOR QOS

Together with the above security needs, the emerging data-intensive Grid applications usually have also strict requirements on the underlying networks in terms of throughput, latency and jitter. Many high-energy physics Grid applications desire high-speed networks capable of transferring bulk files in the order of terabytes at rates of 1Gbps or higher. Some Grid applications, such as those featuring interactive and high-resolution object rendering, desire not only high bandwidth but also low latency and low jitter. Consequently, best-effort IP networks such as the Internet cannot easily accommodate the Grid applications exemplified above at a reasonable cost and any kind of virtual private overlay network implemented as the abstraction of a mesh or multipoint interconnection of secure virtual circuits of on the Internet to interconnect Grid sites and applications must take into account their QoS requirements. Quality of service generally describes the assurance of a minimum available bandwidth and sufficiently low delay or packet loss for specific types of applications or traffic. Usually the delay and packet loss requirements, are combined together to describe a specific service class (premium, assured or best-effort service). The delay and/or packet loss bound guarantees requested in a specific service class could be defined either in a deterministic or statistical way. Deterministic QoS guarantees promise an absolute end-to-end bound for every packet carried by a virtual connection. On the other hand, with statistical guarantees, the end-to-end bound is accompanied with a small probability of violation. For applications that can tolerate occasional bound violations, statistical guarantees can help to reduce the resource requirement for each virtual connection. For instance a virtual connection carrying sensitive traffic between a couple of Grid applications might require average bandwidth of 10Mbps, and premium service class, defined by near zero packet loss, per-packet delay smaller than 50ms, and probability of violating the delay bound smaller than 10^{-3}. In both the legacy and the MPLS-based layer 2 VPN model, QoS guarantees are usually expressed in terms of maximum bandwidth guaranteed (Committed Information Rate) and available (Peak Information Rate) on a certain virtual circuit. The committed bandwidth guarantee is usually provided through the statistical nature of the Layer 2 service, but depends on the overbooking strategy of the network service provider (ISP). This means that the committed rate may not be actually guaranteed although the provider can provision a Minimum Information Rate across the Layer 2 infrastructure. QoS guarantees for virtual connections come at a price - the involved ISPs need to dedicate a portion of physical network resources (such as link capacity, buffer space and computation resources) for each virtual connection. From an ISP's point of view, the network resources need to be utilized in the most efficient manner possible. Thus the central problem faced by the involved ISPs becomes the assignment of available resources to each virtual connection so as to satisfy the following two objectives:

- The QoS guarantee for each virtual connection must be satisfied.
- The number of virtual connections admitted over the long term is maximized.

The first objective is to satisfy the performance requirements for each virtual connection. The second objective is to essentially maximize the overall network usage efficiency which in turn impacts the total revenue derived by the ISP.

In our work, we will use the most advanced MPLS-based network resource management and engineering techniques that can be employed to achieve these two (often conflicting) goals in the context of virtual private connections that require:

- a specific class of service treatment,

– long-term bandwidth guarantees
– protection against failure of any one link/node along the path of the virtual private network connection.

V. THE ARCHITECTURAL FRAMEWORK

In the global grid scenario the network and security services required by the geographically distributed applications may vary from basic end-to-end connectivity, like Internet access, to more complex isolation services, like QoS-guaranteed Virtual Private Networks. From the perspective of the networking community, the main challenge is to develop inter-domain protocols and facilities, so that security-effective and QoS-guaranteed virtual private end-to-end and multipoint connections can be set up and torn down across multiple carrier networks. On the other hand, the Grid community takes a top-down perspective and sees two major requirements lying ahead.

First, the dynamic nature of Grid computing calls for application-driven provisioning of secure end-to-end connections. Traditionally such connections are manually provisioned by the transport network administrator. However manual provisioning will not fit into the picture of Grid computing where connections between nodes and applications need to be set up and torn down on demand. The decisions as to when and where to set up these connections, how much bandwidth is needed and when to tear down each of them, are all parts of the Grid computing workflow. Consequently VPN connections should be viewed as dedicated physical wires or wire meshes that can be turned on and off by Grid applications. Such applications, however should not be aware of the underlying network layout or resource availability so that they can only transparently query the network control plane to drive the layer 2 topology to fit their needs.

Second, powerful and flexible interfaces for virtual circuit allocation and management, implementing VPON infrastructures, are needed between Grid applications and high speed transport networks. It is unlikely that a network carrier will dedicate all its optical network resources to one single Grid project. Instead, a network carrier shall divide its switching domains and high speed links into partitions, and each partition should be only visible to, and accessible, by the designated Grid project.

Accordingly, in our proposed architectural framework, an application program running on a computer should be able to dynamically request via a web-service interface, a layer-2 circuit to a distant computer and have this request filled cooperatively by the network devices on the end-to-end path between these computers. Control plane protocols define the procedures for the handling such on-demand calls, i.e., immediate requests for connectivity at a guaranteed rate. The adaptability/dynamicity feature of Grids makes support for immediate on-demand requests for bandwidth necessary in a suitable transport network, which may be a mesh of private or public shared networks, owned and managed by some cooperating service providers and/or enterprises. Anyway, the network must be a transparent cloud with respect to the Grid, so that all the necessary network operations have to be totally hidden to the customers.

A. Network Operations

We consider a network of label switching routers and communication links that may be under the administrative control of several cooperating ISP, realizing a common transport infrastructure. A subset of the routers are known to be ingress and egress points for the GRID network traffic and these are typically the customer edge devices directly attached to the ISP's point-of-presence locations, i.e. places where ISP's network interfaces with customer sites. First, the whole transport network involved in the implementation of the layer-2 VPN service must support MPLS to switch the traffic based in the MPLS labels. In most cases the customer service provider's sites will be located in different Autonomous Systems (ASes), different providers, so the VPN will transit through several domains (inter-domain MPLS VPN). There are no requirements for CE devices in order to map the logical connections to the remote sites - they have to be configured as if they were connected to a single bridged network or local area network. Also the Provider Routers, in the core do not have any information related to the VPN and only transfer the labeled packets from one PE to another in a transparent way. All the VPN intelligence is located in the PE. It is where the VPN connection originates and terminates, and where all the necessary tunnels are set up to connect to all the others PEs. As we already stated in section 2, there are several available strategies (expressed by different drafts) to implement layer-2 MPLS VPNs. The main difference between them is in the supported signaling protocol, that is vital to implement the label switched tunnels. The first one, (supported by Juniper Inc.) uses Border Gateway Protocol (BGP) while the other (supported by Cisco Systems Inc.) uses Label Distributed Protocol (LDP) for this purpose. Some of the benefits to use BGP as signaling protocol are that it allows for the auto-discovery of new sites, and is better supported at the inter-domain level. If we use BGP, when we add new sites we will only need to configure the PE connected to the new site. Moreover, BGP is a more scalable protocol, so we can use route reflectors or confederations to handle VPN deployment in complex inter-domain infrastructures. Anyway, they all, both solutions, have a common objective; to exchange VPN information generated inside an AS with the other remote ASes. The MAC addresses and connection ports of the users in the local sites will be known by the remote users. In our implementation we preferred the use of Multiprotocol Border Gateway Protocol-based (MP-BGP) signaling to distribute labeled VPN-IPv4 (Internet Protocol version 4) routes and VPN information between AS border or internal routers or router reflectors. We need to advertise the VPN information from one PE to the others, so we will configure one MP-BGP session from each PE to the rest of PEs. Note that some of these sessions will be external and others internal BGP sessions. Accordingly, we have to establish one internal MP-BGP session between the

loopback addresses of all the PE routes belonging to an AS and configure Label Switching Paths (LSPs) between them. At the network control plane level, for each virtual private network connection between two CE nodes a reserved LSP must be set up through the underlying network to carry a service guaranteed traffic stream from the ingress PE router to the egress PE router where the CE nodes are attached. By definition, LSP connections are unidirectional so that a different LSP is needed in both the directions. Each connection implemented through a single LSP is identified by a unique label. Such Virtual private network connections are long-lived connections possibly lasting several months at a stretch. Hence a single LSP is set up to carry customer's per-VPN aggregated traffic rather than a single small individual flow. We assume that all the small flows in the aggregate have similar QoS requirements and that the bandwidth of virtual connection is sufficient to accommodate the cumulative requirements of individual flows. For LSP management purpose, we need MPLS support and one signaling protocol, which can be LDP or RSVP. Clearly, we need a routing instance for each site we want to connect. To handle inter-domain connections, we could configure one normal BGP session between the AS border routers and extend the LSP from each PE to the others PEs through domains using LDP or RSVP, but there is another possible solution, this is, a new Network Layer Reachability Information (NLRI) family called labeled-unicast that results in labeled route exchanges between providers AS Border Routers (ASBRs) which establishes MPLS LSPs between the providers' PE routers. When the multi-point VPNs and the BGP sessions are established, the behavior of the final users will be as if they are in the same LAN and the transit networks from one user to others will be completely transparent. The whole architectural schema, from the network operations point of view is reported below.

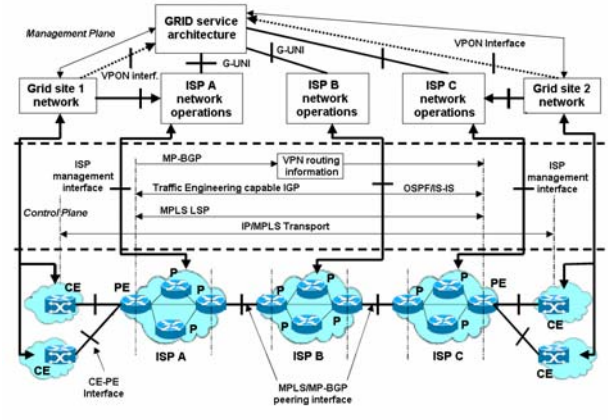

Figure 4. The networking architecture

B. Security Guarantees

VPN security in the MPLS environment is accomplished by using a combined data plane and control plane dual-layer protection approach for security.

The data plane protects against a packet from within a MPLS VPN from traveling outside of its VPN boundaries and from packets from outside a MPLS VPN traveling into the boundaries of a MPLS VPN. The label switching logic will ensure that all the MPLS routers will perform label-based selective packet forwarding on their VPN interfaces. This means that they will drop on the involved interfaces all the packets that do not belong to a specific MPLS VPN by examining the label of each packet. Thus all the information sent on a VPN connection will be strictly confined to the VPN participating nodes.

Control plane security ensures that non-trusted peers can not inject routes into the MPLS VPN. This is accomplished by the use of the available MD5 authentication feature of BGP. Control plane security will also ensure that physical security of the routers is maintained to eliminate unauthorized access. A closed VPN is inherently secure since it has no connection to the Public Internet. If Internet access is needed one path can be setup to provide access. A single firewall can be placed on this path to provide a secure connection for the entire VPN. This is much easier to manage since policies will need to be maintained on only one firewall for the entire VPN. However security is strongly provided:

- At the edge of a provider network, ensuring packets received from a customer are placed on the correct VPN.
- At the backbone, VPN traffic is kept separate. Malicious spoofing (an attempt to gain access to a PE router) is nearly impossible because the packets received from customers are IP packets. These IP packets must be received on a particular interface or sub-interface to be uniquely identified with a VPN label.

From a security perspective, it is important to note that whereas MPLS layer 2 VPNs provide traffic isolation, similar to ATM or frame relay it does not include a mechanism to provide strict confidentiality through encryption. However, if the layer 2 separation provided by partitioned routers and reserved paths is not considered sufficient for the security requirements of the Grid users and strong encryption is required, strong cryptographic tunneling technologies such as IPSec and MPLS can be used together.

C. QoS and path protection support

The main requirement for the MPLS control plane to support the different service classes needed for implementing our Grid-empowered VPON Service is to ensure that packets get the appropriate QoS treatment by each LSR in the transport network. However, since the LSR by definition does not inspect the IP header, it is necessary to provide the required QoS information (i.e. the DSCP according to the DiffServ model) through the label header. Properly crafted LSPs between the PE nodes in the transport network, associated with guaranteed bandwidth and QoS properties will be used to support our VPON virtual connections. This technique inherits the basic QoS approach used in traditional ATM or Frame Relay VPNs (apart from the fact that the MPLS LSP

model is unidirectional, whereas in ATM or Frame Relay the connection is normally defined as bidirectional). One guaranteed bandwidth LSP can thus be established for each pair of CEs belonging to the same VPON, thus generating a mesh of LSPs, implementing the individual end-to-end CE-CE virtual connections. For this sake we first need to select a distinct primary route between the source and destination PE nodes that can transport the VPON's traffic during normal operations and also has sufficient resources to satisfy its QoS requirements. If protection is required we also need to select a backup route which can be kept ready to transport the VPON's traffic between the terminating nodes whenever any one link/node along the primary route fails. An important consideration is to select the primary and backup routes in such a manner that maximum number of virtual connection requests can be accommodated in the future. Thus route selection process should, as far as possible, avoid those physical links that are of critical importance to a large number of source-destination pairs. Since all MPLS packets in such LSPs always belong, from the QoS point-of view, to a single forwarding class, there is no need to indicate the forwarding class of each packet in a specific field of the MPLS header, because it can be directly derived from the label information. This approach to QoS support on MPLS is known as L-LSP (Label-Inferred LSP) based, to indicate that the service class information is inferred from the MPLS label. Stated in a more detailed way, each LSR that contributes to a hop into the L-LSPs, has packet scheduling logic that meets the QoS level defined by the class of service provisioned at each hop. Each LSR in the path examines the incoming label and determines the QoS treatment for the encapsulated packet. Establishing an L-LSP with bandwidth reservation means that QoS and bandwidth requirements for the LSP are signalled at the LSP establishment time. Such signalled bandwidth requirements may be used at establishment time by LSRs to perform admission control depending on the reserved resources provisioned. The above LSPs with reserved resources are established using either LDP or "extended" RSVP with a control-driven downstream-on-demand allocation approach, a scheme most commonly adopted in today MPLS networks because providing more network control (all LSRs belonging to the same LSP perform the label binding in an ordered manner) and better scalability in resource conservation. The LDP (or RSVP) module first checks the link admission control module of the outgoing interface to the next hop on the path to try reserving the required bandwidth. If successful, the remaining capacity of the link is diminished by the requested bandwidth and a Label Request message is sent to the next hop in the explicit route of that LSP, which also checks its link admission control to setup a reservation and so forth until the egress LSR of the explicit route is reached. The egress LSR then sends a Label Mapping message back to the originating LSR– following the reverse explicit route path – with the label information. If the LSP setup fails due to insufficient resources along the explicit path, an error message is sent back to the originating LSR, and the administrator would then try another path. Once the LSP is setup, the requested bandwidth would then be available end-to-end on the explicit route for the "sum" of all aggregate traffic in all the supported classes.

The great strength of L-LSP is its relationship to MPLS fast reroute high reliability restoration services. Packets arriving at the ingress LSR with premium service class QoS requirements, (that is an Expedited Forwarding DSCP value in a DiffServ environment) will be labeled for paths that are fast reroute capable (a backup path has been provisioned). Other non-premium packets (Best Effort) with the same destination can be sent into a shortcut tunnel that is not fast reroute capable. Such facility will be used to implement the required VPON protection classes in our Grid service paradigm.

D. Grid Service interface

Our VPON-secured Grid network will result in an overlay communication facility on top of the existing underlying lower layer networks, whose configuration, security policies and functional behaviors are assumed to be totally independent. The overlay Grid communication facilities must be managed by a standardized middleware stratum, offering well-defined secure service interfaces to the Grid applications. The core middleware technologies that have been widely deployed in the Grid community already include security solutions that support management of credentials and policies, together with resource management protocols and services that support secure remote access to computing and data resources, when computations span multiple institutions. We developed our interfaces basing on the above technologies to ensure that each on-demand access to the secure layer-2 communication Grid will be preceded by the necessary identification, authentication and authorization activities.

D.1 Communication and service reference model

Over the years, existing grid systems have stimulated a clear need for the existence of a well defined standard for possible protocols of secure communication between entities in a multi-enterprise grid system. Global Grid Forum (GGF) is the community involved actively in developing these standards and specifications for grid computing [13]. GGF has come up with a service oriented architecture which defines a set of basic capabilities and functionalities that address prime questions in grid systems that is known as Open Grid Services Architecture (OGSA). Industry efforts have rallied around Web services (WS) as an emerging architecture which has the ability to deliver integrated, interoperable solutions. A natural choice for implementing the VPON service interface on the Grid host sites is the Web Service Resource Framework (WSRF) [14] aiming at implementing some of the OGSA core services as Grid services, or better, web services enhanced for Grid applications. The implemented Web Service interfaces will be stateless and persistent, where data is not retained among invocations and services outlive their clients. They will also be compliant with the

GGF's OGSA specification [15] and, in addition, conform to widely used Web Services standards (WSDL, SOAP, and XML). It is reasonable to expect that in the future all Grid applications will be required to be OGSA-compliant [14]. OGSA defines Grid services as special Web services [17] that provide a set of well-defined interfaces that follow specific conventions [18], usually coordinated, with delegated authentication credentials, in a virtual organization. In other words OGSA enhances Web Services to accommodate requirements of the Grid. The fundamental concept behind OGSA is that it is a service-oriented Grid architecture powered by Grid services [16]. Despite the fact that OGSA represents a long-overdue effort to define a Grid architecture, it is a relatively new standard [16]. The Open Grid Service Infrastructure (OGSI) was the first set of formal and technical specifications of the concepts described in OGSA, but many problems were reported regarding these. In order to circumvent the discrepancies in the OGSI specifications a new standard is emerging, which is called Web Services Resource Framework (WSRF) [19]. WSRF represents a refactoring and evolution of OGSI that delivers essentially the same capabilities in a manner that is more in alignment with the Web Services community [17]. As such, it represents an important next step towards the larger goal of a comprehensive Open Grid Services Architecture that supports on-demand, utility computing, collaborative and other Grid scenarios within a Web services setting. The most valuable aspect of WSRF is that it effectively completes the convergence of the Web services and Grid computing communities. WSRF specifications build directly on core Web services standards, in particular WSDL, SOAP and XML, and exploit capabilities provided by WS-Addressing [20]. Since the proposed architecture is Web Services based it can be integrated with anything based on WSRF standard. In our proposal we explicitly refer to the Globus Toolkit [21] that implements a subset of OGSA services based on WSRF and to the Grid Security Infrastructure (GSI) services [22] providing inter-domain security protocols that bridge the gap between the different local security solutions at a Grid's constituent sites, to address the unique security requirements that arise in Grid environments.

D.2 Interface definitions

In the proposed architecture the transparent on-demand virtual layer-2 connection and security service provisioning is strictly related to basic connectivity services (like label switched path establishment) that should be hidden to the users. A fundamental construct underlying many of the required attributes of the Grid services architecture is that of service virtualization. It is virtualization of Grid services that underpins the ability to map common service semantic behavior seamlessly onto native platform facilities. For these reasons we proposed and developed a new service oriented abstraction that, based on the existing OGSA architecture and built on the Globus GSI toolkit, introduces a new secure connections layer, between the customers and the network infrastructure decoupling the connection service provisioning from the underlying network infrastructure implementation. On-demand allocation of VPON virtual connections requires on-line discovery of MPLS label-switched tunnel/path resource availability on the transport network to accommodate, if appropriate in terms of endpoints, bandwidth and QoS, new layer-2 associations on existing LSPs between the terminating network elements or create, if needed new ones. Grid middleware supports this by relying on information models responsible for capturing structures and relationships of the involved entities. To cope with the heterogeneity of the network infrastructure resources when making advanced reservations or engineering, we proposed a new technology-independent network resource abstraction: the Traffic Engineered Tunnel, modeling the available PE-to-PE LSPs on the underlying networks that can be used from the Grid for virtual connection transport. A centralized Tunnel Resource Broker keeps track of all the above available resources and interfaces with the MPLS network elements to cope with all the necessary network operations needed for handling the VPON connection facilities. For example, a dedicated bandwidth may be reserved between cooperating Grid applications connected in a layer-2 VPON so that based on network condition, Grid middleware can request, through the Tunnel Resource Broker, QoS or bandwidth constrained tunnels between relevant MPLS network elements. Once the service related tunnel resources are configured and provisioned, they have to be monitored from the performance and functionality point of views. Of course, this service too will be made available via the above resource broker. In detail, the proposed abstractions, supporting the VPON connectivity services concern:

- Connection Creation that allows a Layer-2 transparent connection with the specific attributes to be created between a pair of access points
- Connection deletion that allows an existing connection to be deleted
- Connection Status Enquiry that permits the status of certain connection parameters to be queried
- Connection Modification which allows parameters of an already established connection to be modified.

Each request to the Tunnel Resource Broker will be strongly authenticated against a Grid-wide PKI infrastructure through the GSI Generic Security Service (GSS) API [23] defining standard functions for verifying the identity of communicating parties, based on a Public Key Infrastructure where users authenticate to the grid using X.509 certificates. Thus the grid application or user must use its X.509 certificate provided in the GSI environment also to join to a layer-2 Grid association, identified by an existing VPON. The Grid Network Services interact with the Service Provider via the Grid User to Network Interface (GUNI) that implements the basic VPON functionalities and permits Grid applications to dynamic control and manage the underlying network resources according to the cooperation agreements stipulated between the Grid organization and the Service Providers owning the transport networks. Communication

between the GRID applications and the GUNI top level service interface will take place via SOAP/HTTP (eventually secured by SSL) using well-defined extended WSDL Grid Web service interface. Requests and responses conform to Web Services specifications, i.e., they are SOAP messages, carried in HTTP envelopes and transported over TCP/IP connections. The GRID Service Interface can announce its services by means of a Universal data base Description, Discovery and Integration (UDDI). About the specific GUNI implementation, the Extensible Markup Language (XML) appears to be the best candidate thanks to its representation format which can be useful to describe and transmit management information and Grid and network resources. Each network resource or node can be described by a set of XML interface elements. The overlay VPN topology can be represented by mutually referencing node interfaces through the attributes of the VPON termination elements. Note that every Interface can be characterized by the virtual link or LSP tunnel (identified by the addresses engaged) that in turn is characterized by a set of attributes (Service class, Bandwidth available, and Bandwidth utilized). The ability of the Service Interface to hide the complexity of the service provisioning permits to define simple XML-based messages capable of supporting high level services. In particular we want to describe the messages exchanged through GUNI related to the Grid layer-2 connection service:

- *Create_VPON (identifier)*: where an identifier uniquely associates to a new layer-2 VPON on a Grid. The details of VPON setup and configuration are totally hidden from the Grid applications and users.
- *Attach_VPON (source, existing_VPON, bandwidth, Qos)*: a Grid site joins a VPON by establishing a transparent secure layer-2 LAN-alike connection with the other nodes belonging to the secure Grid, with a guaranteed bandwidth and QoS service class such as Premium, Best Effort, etc., with its inherent path protection capabilities.
- *Leave_VPON (source, existing_VPON)*: a Grid site leaves an existing VPON.
- *Modify_VPON (source, existing_VPON, bandwidth, Qos)*: modifies the bandwidth and Qos parameters of an existing connection.
- *Query_VPON (source, existing_VPON)*: query the status of an existing VPON connection.

Every basic service function is in turn mapped to a set of UNI primitives for network resource setting. Commercial routers are not yet provided with standard UNI but, in general, are equipped with an application programming interface (API) based on XML that routers use to exchange information with the Tunnel Resource Broker. Using this interface it is possible to manage and monitor the available LSPs and relative traffic and performance parameters. In order to validate the service, a very simple prototype testing scenario was created, with three PCs running Linux, used as three grid nodes, operating in the Globus environment and interconnected across an MPLS transport network made with five M10 Juniper routers. The signaling interface between the Tunnel Resource Broker and the network elements has been implemented by using XMLscript language via TCP socket. The JUNOScript eXtensible Markup Language (XML) API [24] is used to exchange configurations and operational data between the Tunnel resource broker and the JUNOScript agent on the router in a tagged format. The client-server communication is session-based. Data retrieved from the router can be recast in different formats through the Extensible Stylesheet Language Transformations (XSLT). The prototype of the Tunnel Resource Broker has been implemented on another Linux server which is responsible for the creation, modification and deletion of dynamic LSPs needed by the VPON services and is the only device talking with the Juniper routers. Proper configuration is needed when MPLS paths are requested to enable MPLS-based VPNs and set up a symmetric path from the destination to the source domain. All the LSPs are configured on the LSP head-end router, which is the gateway of the Grid host joining to the VPON. In order to identify the device to configure, the broker uses an internal topology database from which network devices and routing information can be accessed. Configuration requires the definition of the LSP name (according to some naming conventions), of the associated Label-Inferred class of service and, possibly, of some additional terms such as the LSP bandwidth. The setup of intermediate routers is done automatically by a MPLS signaling protocol (RSVP-TE or CR-LDP) that is supported by all the intermediate domains toward the destination. The interface architectural model is sketched in fig. 5 below.

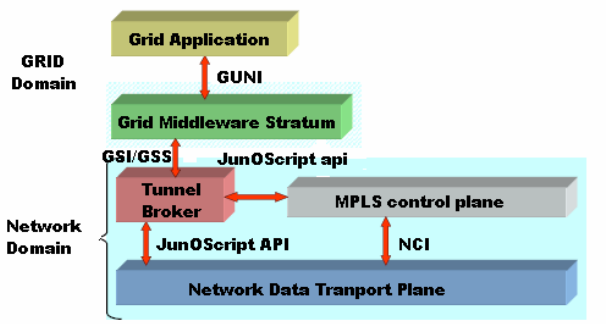

Figure 5. The interface model.

VI. CONCLUSIONS

In this paper, we presented a service-oriented framework that allows distributed Grid applications to transparently control their private and dedicated transport networks, and communicate as they were on the same local area network independently from the security policies and access control mechanisms implemented on the sites which they belong to. The Grid "virtual organization" paradigm can be achieved at layer 2 and thus extensions to existing services are provided to implement on-demand Virtual Private Overlay Network

services in Grids. The proposed framework is based on the MPLS VPN, which is the most flexible and scalable between the available technologies to implement dynamic on-demand tunnels through which the VPON services are implemented. The layer-2 network partitions and their interaction with the underlying network control plane have been abstracted using a secure web service interface. We were able to demonstrate that the VPON services for Grids proposed here are viable, by transparently and dynamically configuring on the underlying transport network some test Grid nodes in a VPON with different guaranteed bandwidth and packet forwarding behaviors.

REFERENCES

[1] I. Foster, C. Kesselman, S. Tuecke, "The Anatomy of the Grid: Enabling Scalable Virtual Organization", *International Journal Supercomputer Applications*, 2001.

[2] P. Ferguson, G, Huston, "What is VPN", *The Internet Protocol Journal* vol. 1 n. 1, 1998.

[3] E. Rosen, A. Viswanathan, R. Callon, "Multiprotocol Label Switching Architecture", *IETF RFC 3031*, 2001

[4] S. Previdi, "*Introduction to MPLS-BGP-VPN*", Proceedings of MPLS Forum, 2000.

[5] L. Martini, et al., "Encapsulation Methods for Transport of Layer 2 Frames Over IP and MPLS Networks", *IETF draft, draft-martini-l2circuit-encap-mpls-12.txt*, 2006.

[6] L. Martini, et al., "Transport of Layer 2 Frames Over MPLS", *IETF draft, draft-martini-l2circuit-trans-mpls-19.txt*, 2006.

[7] K. Kompella, et al., "*MPLS-based Layer 2 Virtual Private Networks*", http://www.juniper.net/techcenter/techpapers/200009.pdf, 2001.

[8] E. Rosen, et al., "An architecture for L2VPNs", *IETF draft, draft-ietfppvpn-l2vpn-00.txt*, 2001.

[9] K. Kompella, et al., "MPLS-based Layer 2 VPNs" , *IETF draft, draft-kompellampls-l2vpn-02.txt*, 2001.

[10] B. Gleeson, A. Lin, J. Heinanen, G. Armitage, A. Malis, "A Framework for IP Based Virtual Private Networks", *IETF RFC 2764*, 2000.

[11] L. Andersson, et al., "PPVPN L2 Framework", *IETF draft, draft-andersson-ppvpn-l2-framework-01.txt*, 2002.

[12] W. Agustyn, et al., "Requirements for Virtual Private LAN Services (VPLS)", *IETF draft, draft-augustyn-vpls-requirements-02.txt*, 2002.

[13] N. Nagaratnam, P. Janson, J. Dayka, A. Nadalin, F. Siebenlist, V. Welch, I. Foster, S. Tuecke, "*The Security Architecture for Open Grid Services*", http://www.cs.virginia.edu/~humphrey/ogsa-sec-wg/OGSA-SecArch-v1-07192002.pdf

[14] I. Foster, "*What is the Grid? A Three Point Checklist*", http://www-fp.mcs.anl.gov/~foster/Articles/WhatIsTheGrid.pdf, 2002

[15] I. Foster, et. al., "*Open Grid Services Architecture*", GGF draft-ggf-ogsa-spec-014, 2004.

[16] S. Tuecke, et al, "*Open Grid Services Infrastructure (OGSI)*", GGF Draft, GT3, Globus Toolkit 3, 2003

[17] S. Parastatidis, "*A Grid Application Framework based on Web Services Specifications and Practices*", Grid Application Framework White Paper, http://www.neresc.ac.uk/projects/gaf/, 2003

[18] L. Zhang, et al, "*Introduction of a Grid architecture and toolkit for building Grid solutions*", IBM developersWorks white paper, 2002

[19] I. Foster, et al, "*The WS-Resource Framework*", http://www.globus.org/wsrf/

[20] IBM, "*WS-Addressing*", IBM technical note, http://www-106.ibm.com/developerworks/library/specification/ws-add/

[21] Globus Project, "*The Globus Toolkit*", http://www-unix.globus.org/toolkit/

[22] V. Welch, F. Siebenlist, I. Foster, J. Bresnahan, K. Czajkowski, J. Gawor, C. Kessel-man, S. Meder, L. Pearlman, S. Tuecke, "*Security for Grid Services*", Twelfth International Symposium on High Performance Distributed Computing (HPDC-12), 2003.

[23] J. Linn, "Generic Security Service Application Program Interface, Version 2", *IETF RFC 2078*, 1997.

[24] Juniper Networks Inc., "*The JUNOScript API software*", (http://www.juniper.net/support/junoscript/.

Francesco Palmieri holds two Computer Science degrees from Salerno University, Italy. Since 1989, he worked for several international telecommunication companies on a variety of networking-related projects, concerned with nation-wide communication systems, network management, transport protocols, and IP networking. Since 1997 he leads the network management/operation center of the Federico II University, in Napoli, Italy. He has been closely involved with the development of the Internet in Italy in the last years, particularly within the academic and research sector, as a member of the Technical Scientific Committee and of the Computer Emergency Response Team of the Italian Academic and Research Network GARR. He is an active researcher in the fields of high performance/evolutionary networking and network security. He has published several papers in leading technical journals and conferences and given invited talks and keynote speeches. He can be reached at *Francesco.Palmieri@unina.it*.