

Virtual Network Approach to Scalable IP Service Deployment and Efficient Resource Management

Yu Cheng, Ramy Farha, Ali Tizghadam, Myung Sup Kim, Massoud Hashemi, and Alberto Leon-Garcia, University of Toronto

James Won-Ki Hong, POSTECH

ABSTRACT

As the Internet evolves into a global all-service communication infrastructure, a key consideration is providing quality of service guarantees over IP with efficient resource utilization in a scalable, flexible, and automatic way. In this article we present a virtual network (VN) based architecture for scalable IP service deployment and efficient network resource management. Particularly considering a DiffServ/MPLS IP transport network supporting multiple VNs, we propose a dynamic approach for efficient bandwidth sharing among VNs. The bandwidth sharing is service-level-agreement-based; the spare capacity in underloaded VNs is adaptively and efficiently utilized, and SLA compliance for all the VNs involved is always guaranteed.

INTRODUCTION

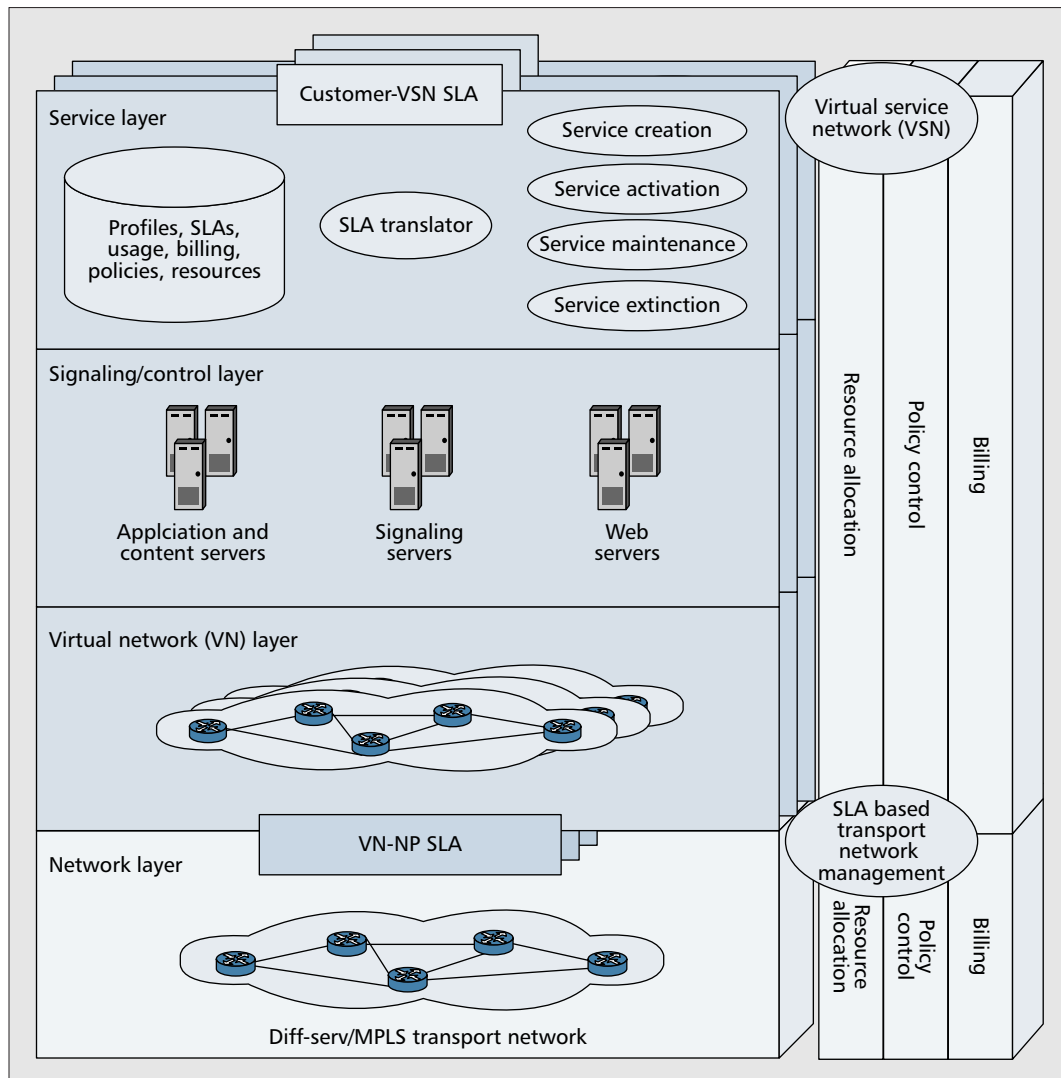
IP networks have been growing dramatically in size and functionality in the past decade. In addition to traditional best effort data services, quality of service (QoS) guaranteed telecommunication services have also started to be deployed over IP networks, such as voice over IP (VoIP). To reduce the time to market of new Internet services and lessen the operation/development/capital costs of service providers (SPs), it is necessary to develop efficient service and network management techniques by which the SPs can create and deploy QoS guaranteed Internet services in a scalable, flexible, and automatic way.

In this article we propose a virtual network (VN) [1] based service and network management (VSNM) architecture, for scalable IP service deployment and efficient resource management. A VN is a subset of physical network resources purchased by an SP from the IP transport network provider (NP) to build a logical service delivery infrastructure. For each VN, specific service management systems are installed to form a virtual service network

(VSN), which provisions certain QoS guaranteed Internet services to customers. In the VSNM architecture, good scalability can be achieved as the NP is freed from high-level application/service management and fine-grained flow-based QoS management, which will be addressed by specific VSNs, and only considers coarse-grained VN-based resource management. Customer or flow-level service management in a VSN is much easier to implement, considering that the scale of a VN is usually much smaller than that of the IP transport network. Each VSN can have its own service management system, and be provisioned and operated to support different service requirements.

We present the notion of virtual network as an alternative to overlay networks (ONs) [2, 3, references therein]. Most overlay approaches build logical ONs on the best effort transport network and depend on application layer techniques to provide certain levels of QoS. However, without appropriate network layer QoS support, the network resources available to an ON are unpredictable, and different ONs impact each other due to uncontrolled network layer resource sharing; therefore, QoS guarantee is hard to achieve and a business service model difficult to develop. In the proposed VSNM model, we consider VNs/VSNs being constructed over a differentiated services (DiffServ)/multiprotocol label switching (MPLS) transport network, which is the promising IP network infrastructure for traffic engineering and scalable QoS management. The resources purchased by each VSN are guaranteed via a service level agreement (SLA) between the VSN and the NP. The business relationships between VSNs, the underlying transport network, and customers of the VSNs are clear: the VSNs purchase network resources from the NP to build logical service delivery infrastructure, and make revenue by selling QoS guaranteed Internet services to their customers.

In this article we give implementation details on VN resource management over a DiffServ/



■ **Figure 1.** Virtual-network-based service and network management architecture.

In VSNM, services are provisioned in a two-tier architecture: the NP provides transport services to VSNs via VN-NP SLAs; a VSN then provides upper layer application services to customers via customer-VSN SLAs, with the purchased VN as logical network infrastructure.

MPLS network. We leverage the TEQUILA architecture [4] to support various VN resource requirements via a uniform *per-VN per-service-class per-ingress/egress-pair* SLA format. Furthermore, we discuss the possibility to achieve a fully automatic network management by applying the *autonomic computing* [5] concept in the TEQUILA architecture. Autonomic management means that the network itself orchestrates the service deployment and resource management under the high-level policy guidance and interaction with the human network managers is limited to establishing the management policies according to the QoS and revenue objectives. We show that the TEQUILA architecture is in essence close to an autonomic control loop [5], albeit the achievement of a fully autonomic management system needs a concerted effort to update each function module to an autonomic element.

Another contribution of this article is that we propose a scheme for dynamic bandwidth sharing between VNs, with the objective of efficient resource utilization in the IP transport network. The novelty of the proposed resource sharing scheme is that the spare capacity from under-

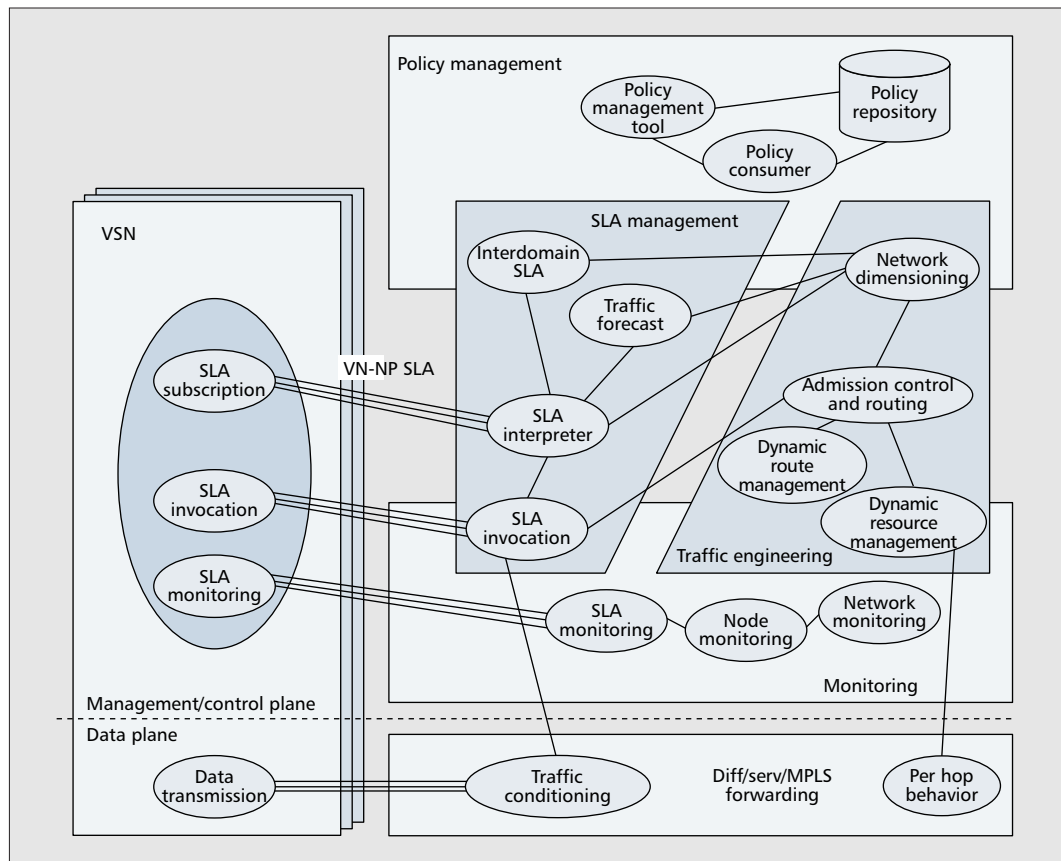
loaded VNs can be adaptively detected and automatically distributed to related overloaded VNs without human intervention, while SLA compliance of all VNs involved is guaranteed. The self-management nature of the proposed dynamic resource sharing technique can be integrated into the TEQUILA architecture as an autonomic element.

VN BASED MANAGEMENT ARCHITECTURE

The SLA-centric VSNM architecture is shown in Fig. 1. SLAs are bilateral contracts between SPs and customers that mainly define, among other things, the services provided, the QoS metrics associated with each service, liabilities on the part of the SP and the customer, and actions to be taken in specific circumstances. In VSNM, services are provisioned in a two-tier architecture: the NP provides transport services to VSNs via VN-NP SLAs;¹ a VSN then provides upper layer application services to customers via customer-VSN SLAs, with the purchased VN as the logical network infrastructure. For example, a

¹ The two parties involved in a VN-NP SLA are, in fact, the SP of the VSN and the NP of the transport network, where the VSN behaves as the customer and purchases network resources with certain QoS guarantees from the NP. We use the term VN-NP SLA to emphasize that the contracted resources are organized into a virtual network.

All the VSN management functionalities are operated following high-level policies, and related information is read from or written into a knowledge database. The QoS satisfaction in VSNs is ultimately determined by whether the network resources contracted in the VN-NP SLAs can be guaranteed by the NP.



■ Figure 2. VTEQUILA architecture for a DiffServ/MPLS network supporting virtual networks.

VN can be purchased by a corporation to provision a QoS guaranteed virtual private network (VPN), or by a telecommunication service provider to provision VoIP or IPTV. Also note that VNs can be spawned from other VNs by suballocating resources, thus enabling scalability and software reuse [1].

In each VSN a layered management structure is used to provision services that usually have a four-phase life cycle: service creation, activation, maintenance, and extinction [6]. The top *service layer* directly communicates with the customers and expresses the service from a business perspective. For example, to create a new VoIP service, the SLA negotiated with a customer may be expressed as “voice quality at a given mean opinion score (MOS).” This business-level QoS will be automatically translated by an SLA translator into technical service parameters, such as packet loss rate, delay, and jitter requirements, and mapped to a DiffServ service class [7].² The VSN also needs to determine the scope of target customers and investigate the customers’ behavior, and automatically translates the information into traffic load estimation and statistical parameters characterizing the traffic profile.

The QoS specification and traffic load information will then be forwarded to the *signaling/control* and *VN layers* for resource planning. The signaling/control layer determines the amount and location of application, content, signaling, and Web servers to meet the service demand. These servers can be equipment inclusively

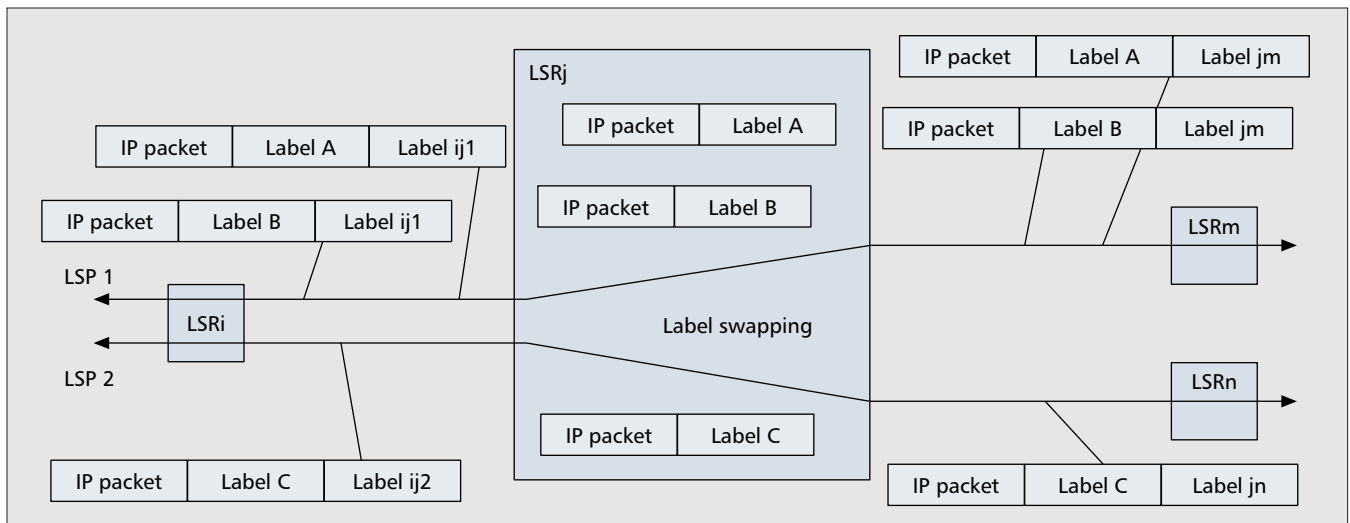
owned by the VSN or virtual resources purchased from some SPs. The VN layer determines the bandwidth/buffering/computing resources (in general, the network resources) that should be purchased from the NP. The resource provisioning of a VN should be carefully planned so that the QoS can be achieved in a cost-effective approach. After resource planning, the service instances can be activated via connection or flow admission control. For ongoing services, the traffic load and provisioned QoS will be monitored at all three layers, and the information will be used for billing, possible SLA renegotiation, and resource replanning. If some services are found unprofitable, they may be retired. All the VSN management functionalities are operated following high-level policies, and related information is read from or written into a knowledge database.

The QoS satisfaction in VSNs is ultimately determined by whether the network resources contracted in the VN-NP SLAs can be guaranteed by the NP. The remainder of this article focuses on VN-based resource management in the transport network.

VN IMPLEMENTATION OVER DIFFSERV/MPLS NETWORKS

In this section we present a VN implementation over the DiffServ/MPLS transport network by leveraging the TEQUILA architecture [4], which is a representative framework for resource management and traffic engineering in DiffServ/

² Hereafter, service class is sometimes briefly referred to as class for convenience.



■ **Figure 3.** Two-layer label stacking to support virtual networks.

MPLS networks. In TEQUILA, a DiffServ/MPLS network is operated in a *first plan, then take care* fashion, first through offline planning and dimensioning, and subsequently through dynamic operations and management functions. In Fig. 2 we replot the TEQUILA architecture in the VN context, called *virtual TEQUILA* (VTEQUILA) for convenience. In the following, we focus on explaining the VN extension parts to TEQUILA, revealing the autonomic concept implied in the architecture, and giving more implementation details.

In Fig. 2 the centralized VTEQUILA system, implemented by a bandwidth broker (BB) according to the DiffServ terminology, manages the resource allocation to each VN. For simplicity, we mainly consider bandwidth as the resource to be managed. Note that the TEQUILA architecture can readily support VN without adding any new functional module. However, customers of the transport network are VSNs in VTEQUILA, and the VN based SLAs, instead of the class based SLAs³ in TEQUILA, are contracted at the customer-NP interface. Therefore, the functionality of the SLA interpreter should be enhanced to support the VN resource commitments. The VN-based resource allocation is not completely consistent with class-based resource allocation, as it is quite possible that a VN includes traffic for multiple service classes. The SLA interpreter should then be able to extract the per-class resource requirements from the VN-NP SLAs and forward such information to the traffic forecast (TF) module. We consider resource allocation with MPLS-based traffic engineering (TE) [4]. By knowing the network topology (memorized in the BB) and the long-term traffic load in each service class (from the TF), the network dimensioning will determine the label switched path (LSP) deployment over the network and calculate the bandwidth provisioning directives for each class at each link. The network dimensioning directives are forwarded to the admission control and routing module as “soft” resource partitions, leaving space for traffic fluctuations to be handled by dynamic route and resource management techniques.

It is noteworthy that the TEQUILA/VTEQUILA architecture is in essence close to a self-managed system, where human interventions are mainly in policy design and SLA negotiation. The modules in Fig. 2 can be conceptually mapped into the “monitor, analyze, plan, and execute” policy-based autonomic control loop [5]. The monitoring and planning (network dimensioning) parts are self-explanatory. The SLA interpreter and traffic forecast modules analyze the resource requirements from VNs. Online resource allocation is then executed by the admission control/routing module as well as dynamic route and resource management modules, following the network planning guidelines. The autonomic nature of the system is also reflected in that the dynamic management modules can trigger the network redimensioning and SLA renegotiation, if they cannot handle the network status deviation due to device failure or traffic load change. In the autonomic control loop, monitoring is essential for problem determination. Nevertheless, the current design only initiates the effort toward an autonomic network management system. A full autonomic implementation requires that each functional component within the management architecture is designed as an autonomic element and that they automatically interact with each other based on high-level policies [5]. In a later section we present a detailed design of dynamic resource management that can be integrated into the architecture as an autonomic element.

TWO-LAYER LABEL STACKING

To support VN-based resource allocation in a DiffServ/MPLS network, we propose a two-layer label stacking scheme as shown in Fig. 3 to achieve both DiffServ-aware TE and VN identification. With label stacking, a label switched path (LSP) between a pair of ingress/egress points can comprise a set of microflows from different VNs. In Fig. 3 LSP-1 is used to deliver traffic from both VN A and VN B. During forwarding, the outer label determines where to forward the packet and the DiffServ per hop

³ In the TEQUILA architecture, the term SLS is used instead of SLA to emphasize the technical service parameters contained in an SLA.

While the RMD-QOSM can well support admission control in a distributed way when SLA resource commitments remain static or are explicitly renegotiated, it can hardly support the inter-SLA resource sharing scheme discussed in this article.

behavior (PHB), and the label switched router (LSR) only applies label swapping to this outer label. The inner label is checked to measure VN-based bandwidth usage and QoS performance. Such information is used for VN-NP SLA monitoring.

UNIFORM SLA MANAGEMENT

In the VTEQUILA architecture, the VN-NP SLA resource commitment is translated by the SLA interpreter to a uniform per-VN per-class per-ingress/egress-pair SLA format for network dimensioning. Applying MPLS-based TE [4], we present a path-oriented dimensioning approach, which is also the basis for the dynamic resource sharing presented later. We assume that a set of parallel paths for each ingress/egress pair are computed offline and fixed as LSPs. All traffic traversing an ingress/egress pair is distributed among these LSPs for load balancing purposes. A *traffic trunk* is defined as a logical pipeline within an LSP, which is allocated a certain amount of capacity to serve the traffic associated with a certain SLA. Therefore, an LSP between an ingress/egress pair may carry multiple traffic trunks associated with different SLAs, and traffic belonging to different trunks can be discriminated by the label stacking scheme mentioned earlier. In this path-oriented approach all the per-VN, per-class, per-ingress/egress resource commitments are mapped to bandwidth allocation at each traffic trunk by network dimensioning.

The dimensioning problem is normally formulated as an optimization problem subject to the following constraints:

- The total bandwidth allocated to parallel traffic trunks associated with an SLA should meet the SLA resource commitment.
- The total bandwidth allocation at a link does not exceed the physical link capacity.

At each router, the total bandwidth allocation for a PHB is derived by summing the bandwidth allocation of all the trunks of the same class that traverse that router along a given output port. With feasible bandwidth allocation for each PHB, the specific scheduling algorithm can be designed and configured correspondingly to guarantee the resource allocation and packet-level QoS requirements.

In the VTEQUILA architecture, the network topology, trunk deployment, network dimensioning results, and real-time bandwidth usage are tracked in the BB and saved in a database. Such information is used to support explicit connection (or flow) admission control. With MPLS-based TE, admission control and routing are correlated and jointly controlled by the same module as shown in Fig. 2. Each time a new connection request arrives at a certain ingress router, it is forwarded to the BB where the admission control and routing module, by checking the stored status information, will select a traffic trunk according to the routing algorithm and make an admission decision according to the resource availability of the selected trunk. The decision will then be delivered back to the corresponding ingress router. If accepted, flow-related information is stored at ingress routers and bandwidth usage information is updated in

the BB. A detailed design of the data structures used in the BB and edge routers is presented in [8].

We suggest an *effective-bandwidth*-based admission control approach due to its capability for both packet-level and connection-level QoS control [9], and its convenience for resource management. Effective bandwidth is the minimum bandwidth that should be allocated to a connection or flow to achieve a satisfying application, or technically to guarantee the flow's packet-level QoS. Each VSN independently calculates an effective bandwidth for each associated service class⁴ to encapsulate the packet-level QoS, and then contract an equivalent flow-level capacity (i.e., the number of bandwidth guaranteed flows) in the VN-NP SLA to guarantee the connection-level QoS (e.g., the connection acceptance ratio). In the transport network, the BB ensures the VN resource allocation and determines on which traffic trunk the accepted flow should be placed. With effective bandwidth-based resource allocation, bandwidth usage and leftover capacity of each trunk (and therefore of each VN) can readily be obtained by tracking the arrival and completion of connections.

We would like to emphasize that the BB is just a *logical* central entity to implement network management in accordance with VTEQUILA. A physical central controller in a network is prone to become a congestion point in the network. Achieving the BB functions in a distributed way is preferred, although it is a problem yet to be solved. The resource management in DiffServ QoS model (RMD-QOSM) being developed in the Internet Engineering Task Force (IETF) Next Steps in Signaling (NSIS) working group proposes to add admission control to DiffServ networks using hop-by-hop resource reservation signaling, where edge routers keep fine-grained reservation states and interior routers use only one aggregated reservation state per class or no state at all. While the RMD-QOSM can well support admission control in a distributed way when SLA resource commitments remain static or are explicitly renegotiated, it can hardly support the inter-SLA resource sharing scheme discussed in this article.

So far, we have discussed VN based service and network management in the architecture level. In the following, we will focus on efficient bandwidth provisioning in both VNs and the transport network, which is the essential point concerning the SP or the NP for high revenue.

VIRTUAL NETWORK BANDWIDTH PROVISIONING

The key issue in VN bandwidth provisioning is to determine the appropriate amount of bandwidth to be purchased by the VSN from the NP in advance, so that the long-term, average net income of the VSN is maximized while the customer traffics are served with QoS guarantee. The optimal bandwidth provisioning needs to take various factors into account including the statistical characteristics of the traffic profile, the packet/connection level QoS requirements, the billing scheme, the investment paid to the NP,

⁴ We assume that traffic flows belonging to the same service class are allocated the same effective bandwidth.

and the penalty scheme for possible SLA/QoS violation. The bandwidth provisioning solution presented in [3] for a service overlay network is possible to be used for VN bandwidth provisioning. However, the provisioning technique can only be applied when all the traffic belongs to the same service class and a linear billing scheme is applied. In practice, it is quite possible that a VN may provide multiple classes of service to customers and different billing schemes may be applied to different services. Therefore, bandwidth provisioning of virtual networks is still an interesting and important open problem.

DYNAMIC BANDWIDTH SHARING

In the transport network, VN-based resource management can be considered as a way to partition resources. At the SLA level, the transport network resources are shared by a set of per-VN, per-class, per-ingress/egress-pair SLAs. Assuming a proper VN bandwidth provisioning approach available, the resource commitment in each SLA is to guarantee certain QoS metrics at an *engineered* traffic load (i.e., an estimation of the long-term [days or weeks] average traffic demand). The BB will find an optimal solution to accommodate all the SLA resource requirements by network dimensioning. However, due to the random nature of traffic, in operation the short-term (minutes or hours) traffic load may be higher or lower than the engineered load in an SLA, leading to *overloaded* or *underloaded* SLA states, respectively. With hard static resource partitioning, the overloaded SLAs will suffer degraded QoS while the spare resources in the underloaded SLAs are wasted. Therefore, in the VTEQUILA architecture, dynamic resource management is used to handle traffic load fluctuation for higher resource utilization and better QoS. Note that an SLA may alternate in overloaded and underloaded states around the long-term engineered load, where frequent network redimensioning is not an economic or even a capable approach to address efficient resource sharing.

A dynamic resource sharing technique based on virtual partitioning (VP) is proposed in [9] to exploit the spare capacity in underloaded SLAs. The cost of virtual partitioning is that the QoS of the underloaded SLAs cannot be guaranteed. QoS or SLA violation for underloaders is a serious problem that can encourage malicious overloading. Therefore, the authors in [9] propose to use a *penalty payment* from the service provider to the customer to compensate for possible QoS/SLA violations. The basic reason leading to SLA violation is that VP-based resource sharing is a static design, where a preconfigured VP scheme is used for SLA resource sharing in all possible load conditions. Moreover, in a VN environment the VSN and NP penalty schemes are correlated, as the VN-NP SLA violation will subsequently incur customer-VSN SLA violation. How to design the VSN and NP penalty schemes to guarantee customer satisfaction as well as high network revenue is a very complex problem. In the following we present an adaptively self-configured resource sharing technique called *bandwidth borrowing*, by which SLA compliance

and high resource utilization can be achieved simultaneously. For convenience, we consider a flow's packet-level QoS to be encapsulated by the notion of effective bandwidth, and refer to a bandwidth guaranteed flow as a *call*. An SLA handles QoS and resources at the call level.

SLA WITH CALL-LEVEL DIFFERENTIATION

To facilitate resource sharing, we propose that the definition of an SLA be extended with a statement of the QoS and resource commitment in an underloaded period and a *call-level differentiation* agreement as follows:

- A nominal capacity is allocated in the SLA according to the target call arrival rate to satisfy the target call blocking probability (CBP).

- During operation, according to the actual call arrival rate measured at the ingress router, two resource utilization states are associated with an SLA. The SLA is said to be in a *lendable state* if the actual call arrival rate is less than the engineered load, and the target CBP can be satisfied with a smaller serving capacity defined as the *QoS ensuring* (QoSE) *bandwidth*. Otherwise, the SLA is in the *unlendable state*.

- In the lendable state, the QoSE bandwidth of the SLA is guaranteed to its traffic flows to meet the QoS. The unused bandwidth within the nominal capacity can be exploited by all related SLAs.

- In the unlendable state, the nominal capacity of the SLA is guaranteed to its traffic flows. Furthermore, the SLA may accept overloaded traffic if borrowing bandwidth from lendable SLAs is possible. The traffic flows accepted with borrowed bandwidth are tagged as *out* profile calls, and the flows accepted with nominal capacity are considered as *in* profile calls.

- When an SLA returns to the unlendable state from the lendable state, the QoSE bandwidth is increased to the nominal capacity to reclaim resources of the SLA. Some out profile flows from borrower SLAs may be preempted during bandwidth claiming.

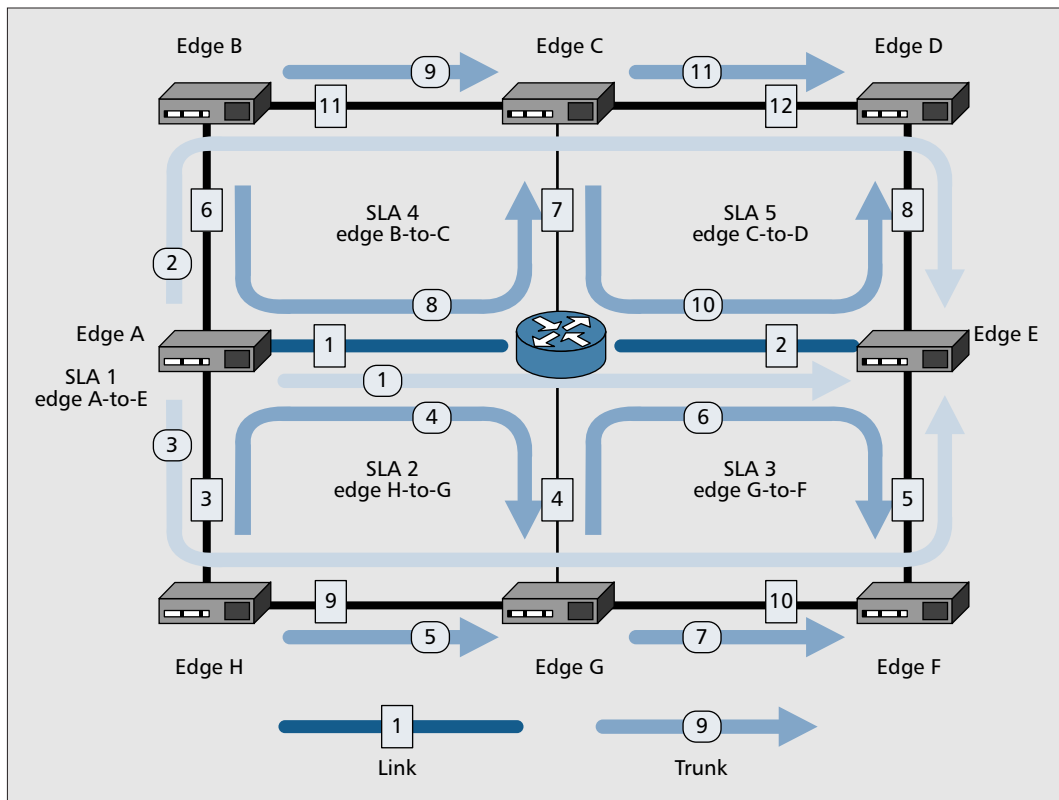
In the above SLA definition, the possible preemption of the out profile calls is considered as the QoS differentiation between the in profile and out profile traffic. The counterpart differentiation scheme at the packet level is the assured forwarding PHB. In addition to efficient resource utilization, call-level differentiation can bring a more customer-friendly service model. When a flow is to be served as an out profile call, a message can be sent to the customer (via the corresponding VSN) before the actual service regarding the flow admission status. The customer can then determine to continue or try at a later time, or send the most important information first.

CALL ADMISSION CONTROL

The proposed bandwidth borrowing scheme can be conveniently implemented using the traffic-trunk-based resource allocation discussed earlier. In the BB each traffic trunk has an information record including the nominal capacity, QoSE bandwidth (the SLA QoSE bandwidth is distributed among associated trunks), and current bandwidth usage. The traffic trunk information records are used by the BB to determine

When a flow is to be served as an out profile call, a message can be sent to the customer (via the corresponding VSN) before the actual service regarding the flow admission status. The customer can then determine to continue or try at a later time, or send the most important information first.

Although here we mainly consider the call-level QoS, the call-level differentiation and bandwidth borrowing concept can be applied to a very general scenario where a bandwidth requirement is able to be determined based on QoS, traffic load and management policies.



■ Figure 4. Network topology, SLAs, and trunk deployment for simulation.

whether a call can be accepted as in or out profile. Three basic principles guiding the admission control are:

- In an SLA, if current resource usage (SLA bandwidth usage can be obtained by summing up all related trunk usages) is less than QoSE bandwidth, a new call is treated as in profile, and its acceptance is guaranteed.
- Out profile calls can be accepted by fully exploiting unused bandwidth.
- If the BB finds there is no bandwidth available for an in profile call, some out profile calls should be preempted.

DYNAMIC SPARE BANDWIDTH DISTRIBUTION

In the bandwidth borrowing scheme, the spare bandwidth (nominal capacity minus QoSE bandwidth) is calculated at edge routers for related SLAs and then distributed to associated traffic trunks. The spare capacity at a certain link can be indirectly obtained by summing the spare capacity over all the trunks traversing the link. It is obvious that the distribution of spare capacity directly determines the resource utilization that can be achieved. A straightforward approach is for the spare capacity from an SLA to be evenly distributed to related traffic trunks. The even distribution may not be the best solution, because the bandwidth borrowing may not happen on all the routes, and the traffic loads and resource sharing levels on different routes, and therefore on different links, are different. Ideally, the QoSE bandwidth (correspondingly the spare bandwidth) should be distributed in such a way as to lead to maximum resource utilization.

It is very difficult, if not impossible, to derive a centralized optimal online distribution technique. Therefore, we propose an edge-based scheme for dynamic spare bandwidth distribution [8], which in some degree adds the distributional property to the BB.

Details of the bandwidth borrowing scheme are presented in [8, 10]. At this point we want to emphasize that the bandwidth borrowing scheme is designed with the objective of upgrading the dynamic resource management module in the VTEQUILA architecture to an autonomic element. In bandwidth borrowing, the bandwidth reservation for each SLA to guarantee QoS is adjusted adaptively, and the bandwidth sharing over the network is also dynamically adjusted according to network status and traffic load variations. Although here we mainly consider call-level QoS, call-level differentiation and bandwidth borrowing can be applied to a very general scenario where a bandwidth requirement is able to be determined based on QoS, traffic load, and management policies.

PERFORMANCE EVALUATION

We present some simulation results from a small network, shown in Fig. 4, to demonstrate the performance of the bandwidth borrowing scheme. The units used for related measures are seconds for time, capacity unit (c-unit) for link/trunk/SLA capacity and efficient bandwidth usage, calls per second for call arrival rate, and c-unit per call for effective bandwidth. Five SLAs from different VSNs are supported by this network, and each SLA is served with parallel

Configurations	SLA-1	SLA-2	SLA-3	SLA-4	SLA-5	EBU
CP	0.3769	1.4758e-8	0.0100	0.0100	1.4758e-8	144.6668
VP	0.1718	0.0356	0.0688	0.0690	0.0367	159.4450
BR	0.1459	0.0028	0.0086	0.0087	0.0027	166.3356

Table 1. Call blocking probability and efficient bandwidth use in different configurations.

traffic trunks. Assume Poisson arrivals for each SLA and exponentially distributed call holding times with mean of 1. The effective bandwidth associated with each SLA is equal to 1. The target CBP for each SLA is 0.01. The engineered call arrival rates for SLA-1 to SLA-5 are (46.9, 29, 29, 29, 29), and the corresponding SLA capacity planning is (60, 40, 40, 40, 40) to achieve the target CBP for the engineered traffic load. The SLA capacity is evenly distributed to the related traffic trunks. The capacity is 60 for links 1 and 2, and 40 for the other links.

We consider a scenario where SLA-1 is overloaded and SLA-2 and SLA-5 are underloaded, and the online measured call arrival rates for SLA-1 to SLA-5 are (93.8, 14.4, 29, 29, 14.4). The QoSE bandwidths of SLA-2 and SLA-5 are then calculated equal to 23, and each has an initial even distribution of (11, 12) over its associated two trunks (fractional division of a capacity unit is assumed impossible). We compare the performance of the complete partitioning (CP), virtual partitioning (VP), and bandwidth borrowing (BR) schemes, and the sample path (of 8000s with 1,446,531 total call arrivals) is identically reproduced in all the simulations for fair comparison. Under each resource sharing scheme, the CBP for each SLA is measured, by which the total efficient bandwidth usage (EBU) over the network is calculated according to the approach given in [9]. All the results are presented in Table 1. Note that the results for the BR scheme are conservatively measured, where all the preempted calls are treated as blocked calls. In BR, SLA-1, SLA-3, and SLA-4 have 569,767, 941, and 1044 out profile admissions, respectively, of which 56,163, 21, and 21 calls have been preempted with the corresponding out profile preemption probability of 0.0986, 0.0223, and 0.0201.

With the CP scheme, traffic service in each SLA works independently. The CBP is directly obtained from the Erlang-B formula and serves as the performance benchmark. Obviously, CP leads to the worst QoS and resource utilization due to static resource allocation. Both the VP and BR schemes can significantly improve resource utilization. However, in the VP scheme, the overload SLA-1 leads to QoS violation in all the other SLAs, including both underloaded and normally loaded ones. With the BR scheme, resource utilization is even better than for VP, and SLA compliance is guaranteed. The good performance of the BR scheme stems from the more aggressive bandwidth usage by out profile calls, dynamic adjustment of spare bandwidth distribution, and protection of the in profile flows via the preemption scheme. We have com-

prehensively evaluated the performance of the BR scheme using large network topologies, heterogeneous effective bandwidths, and various load scenarios, and BR has robust good performance. We will present the complete results in a future paper due to the space limitations for this article.

CONCLUDING REMARKS

In this article we have presented a virtual network approach to IP service deployment and network resource management. The VN approach subdivides the IP transport network into multiple self-managed subsystems, virtual service networks, to achieve good scalability. Considering a DiffServ/MPLS transport network, we give a promising implementation to support VN-based resource management. Furthermore, we have discussed bandwidth provisioning and bandwidth sharing in the VN environment, where a novel dynamic resource sharing technique is proposed to achieve high resource utilization and SLA compliance simultaneously. In addition, we point out that service and network management are evolving toward a fully automatic or autonomic system. The self-management nature implied in the VN approach and the proposed dynamic resource sharing are consistent with the autonomic concept. For future work, we are developing techniques for fully autonomic SLA negotiation, service management, fault management, and security management in multidomain DiffServ/MPLS networks in order to achieve a full autonomic service and network management system. Efficient resource management where the VNs carry traffic that cannot be decomposed into calls (e.g., peer-to-peer traffic) is also under investigation.

REFERENCES

- [1] A. Leon-Garcia and L. G. Mason, "Virtual Network Resource Management for Next-Generation Networks," *IEEE Commun. Mag.*, vol. 41, no. 7, July 2003, pp. 102-09.
- [2] Z. Li and P. Mohapatra, "QRON: QoS-Aware Routing in Overlay Networks," *IEEE JSAC*, vol. 22, no. 1, Jan. 2004, pp. 29-40.
- [3] Z. Duan, Z.-L. Zhang, and Y. T. Hou, "Service Overlay Networks: SLAs, QoS, and Bandwidth Provisioning," *IEEE/ACM Trans. Net.*, vol. 11, no. 6, Dec. 2003, pp. 870-83.
- [4] P. Trimintzios et al., "A Management and Control Architecture for Providing IP Differentiated Services in MPLS-Based Networks," *IEEE Commun. Mag.*, vol. 39, no. 5, May 2001, pp. 80-88.
- [5] J. O. Kephart and D. M. Chess, "The Vision of Autonomic Computing," *IEEE Comp.*, vol. 36, no. 1, Jan. 2003, pp. 41-50.
- [6] TeleManagement Forum, <http://www.tmforum.org>
- [7] J. Babiarz, K. Chan, and F. Baker, "Configuration Guidelines for DiffServ Service Classes," Internet draft, draft-baker-diffserv-basic-classes-04, Oct. 2004.

The good performance of the BR scheme stems from the more aggressive bandwidth usage by out profile calls, dynamic adjustment of spare bandwidth distribution, and protection of the in profile flows via the preemption scheme.

For future work, we are developing techniques for fully autonomous SLA negotiation, service management, fault management and security management in multi-domain DiffServ/MPLS networks in order to achieve a full autonomous service and network management system.

- [8] Y. Cheng, "Efficient Resource Allocation in Differentiated Services Networks," Ph.D. thesis, Univ. of Waterloo, 2003.
- [9] E. Bouillet, D. Mitra, and K. G. Ramakrishnan, "The Structure and Management of Service Level Agreements in Networks," *IEEE JSAC*, vol. 20, no. 4, May 2002, pp. 691–99.
- [10] Y. Cheng, W. Zhuang, and A. Leon-Garcia, "Call Level Service Differentiation for Efficient SLA Management," to appear, *IEEE GLOBECOM '05*.

BIOGRAPHIES

YU CHENG [S'01, M'04] (y.cheng@utoronto.ca) received B.E. and M.E. degrees from Tsinghua University, Beijing, China, in 1995 and 1998, respectively, and a Ph.D. degree from the University of Waterloo, Ontario, Canada, in 2003, all in electrical engineering. From September 2003 to August 2004 he was a postdoctoral fellow in the Department of Electrical and Computer Engineering at the University of Waterloo. Since September 2004 he has been a postdoctoral fellow in the Department of Electrical and Computer Engineering at the University of Toronto, Canada. His research interests include Internet QoS, traffic engineering, service and network management, and wireless/wireline interworking. He received a postdoctoral fellowship from the Natural Sciences and Engineering Research Council of Canada (NSERC) in 2004.

RAMY FARHA (rfarha@comm.utoronto.ca) is a Ph.D. candidate in the Department of Electrical and Computer Engineering, University of Toronto. He received his B.Eng. degree from the American University of Beirut in 2001, and his M.A.Sc. degree from the University of Toronto in 2003. He received the Natural Sciences and Engineering Research Council (NSERC) PGS award and the Ontario Graduate Scholarship (OGS) award, both in 2003. In addition, he holds the Distinguished Graduate Award from the American University of Beirut in 2001. His research interests include IP mobility, passive optical networks, peer-to-peer and overlay networks, and autonomous service management.

ALI TIZGHADAM (ali.tizghadam@utoronto.ca) is a Ph.D. student in the Department of Electrical and Computer Engineering, University of Toronto. He manages the Network Architecture and Service Creation Laboratory in the Communications group. He received his M.A.Sc. in 1994 from the University of Tehran in electrical engineering. He went into the industry for about 10 years where he gained an abundance of experience in telecommunications, especially in local exchange switches and access networks. He started his Ph.D. in 2003. His major interests are in autonomous core/metro network management, bio-inspired approaches to dynamic routing in telecommunication networks, and SIP as a promising control plane signaling protocol for next-generation networks.

MYUNG SUP KIM (myungsup.kim@utoronto.ca) received B.S., M.S., and Ph.D. degrees in computer science and engineering from Pohang University of Science and Technology

(POSTECH), Korea, in 1998, 2000, and 2004, respectively. Currently, he is a post-doctoral fellow in the Department of Electrical and Computer Engineering, University of Toronto. His research interests include service and network management, Internet traffic monitoring and analysis, application-level traffic analysis, and network-security attack detection and prevention. He is a member of KNOM.

Massoud Hashemi (hashemim@cc.iut.ac.ir) received B.S. and M.S. degrees from Isfahan University of Technology, Iran, in 1986 and 1987, respectively, and a Ph.D. from the University of Toronto in 1998. He was with the University of Toronto as a postdoctoral fellow from January 1998 to June 1999. He was a founding member of QoS Express in Toronto in 1999 and Accelight Networks in Ottawa in 2000. In both startup companies he had major contributions in the design and development of a carrier-grade multiservice terabit switch. He visited the University of Toronto from June 2004 to June 2005 as a research fellow. His research in this period included autonomous network management routing algorithms based on swarm intelligence, and overlay networks. He is currently an assistant professor with the Department of Electrical and Computer Engineering, Isfahan University of Technology, Iran. His current research includes autonomous traffic engineering, next-generation converged networks, and IDS systems.

ALBERTO LEON-GARCIA [F] (alberto.leongarcia@utoronto.ca) was co-founder and CTO of AcceLight Networks, Ottawa, Canada, which developed leading edge optical switching equipment. He holds the Jeffrey Skoll Chair in Computer Networks and Innovation, and is co-director of the Master of Engineering in Telecommunications program he founded in 1997. He is author of *Probability and Random Processes for Electrical Engineering* and coauthor with Indra Widjaja of *Communication Networks: Fundamental Concepts and Key Architecture*.

JAMES WON-KI HONG [SM] (jwkhong@postech.ac.kr) is an associate professor in the Department of Computer Science and Engineering, POSTECH. He received a Ph.D. degree from the University of Waterloo in 1991 and an M.S. degree from the University of Western Ontario in 1985. He has worked on various research projects in network and systems management, with a special interest in Web, Java, CORBA, and XML technologies. His research interests include network and systems management, distributed computing, and network monitoring and planning. He has served as Technical Chair (1998–2000), Vice Chair (2003–2005), a and Chair (2005–present) for IEEE ComSoc CNOM. He is also serving as Director of Online Content for the IEEE ComSoc (2004–2005). He is a NOMS/IM Steering Committee Member and a Standing Committee Member of APNOMS. He was technical co-chair of NOMS 200,0 and APNOMS '99. He was Finance Chair for IM 2005 and Finance Chair and Chair of the Local Planning Committee for NOMS 2004. He is an editorial advisory board member of JNSM and IJNM. He is also editor-in-chief of KNOM Review Journal. He is a member of KICS, KNOM, and KISS.