Contribution Number

GIBSON: Global IP-Based Service-Oriented Network Architecture and IMS Case Study

September 8, 2006

Contributor: Ping Pan, Tom Nolle

Organization: Hammerhead Systems, CIMI Corp.

Author

Address

Address

City/State/Zip

Phone

Fax

E-mail: ppan@hammerheadsystems.com, tnolle@cimicorp.com

Working Group (if applicable)

Abstract

GIBSON (Global IP-Based Service-Oriented Network) is an architecture that is focused on providing data-plane service provisioning and management in the context of IPsphere. One of its key objectives is to interface with business routing at network edge/border to achieve end-to-end (or edge-to-edge) per-user-flow service guarantees.

In this document, we will provide an outline of the GIBSON architecture, and illustrate its operation in multi-provider IMS environment.

Declaration of IPR

IPsphere Forum disclaimer (copyright and IP).

Introduction and Motivation

Providing new services within the existing IP infrastructure will face significant challenges in creating a common service conception in a diversely owned and multi-technology network of the future:

1. Access/metro and backbone networks may belong to different carriers or business entities. For example, the access networks may belong to wireless service providers, while the backbone networks may belong to a national or global carrier that provide data transport services to multiple access or metro networks.
2. There will be technology differences among providers and also often between the metro/access and core networks of the same provider.
3. The combination of ownership diversity and technology diversity is reflected in a more complex set of management interfaces to control infrastructure, and this complexity will create problems if it is reflected upward to the service, business, and operations management layers.
4. There will be a considerable variability in the "value" of service relationships and thus in the per-service handling that can be justified. Some sessions will be handled individually (video, for example) and others will likely be handled in aggregated form (voice).
5. There may be regulatory issues such as intercept/surveillance that will have to be applied, or that will have to be routed around to avoid.
6. Since user traffic is transported as IP packets throughout the network, the backbone carriers may not have the ability or incentive to provide special treatment to *important* user flows. Subsequently, "hot-potato" type of routing polices are applied to inter-carrier traffic. The end users can only rely on application-level congestion and flow control, such as TCP, to regulate traffic. This practice will not likely to scale as end-user applications become more bandwidth-intensive and delay-sensitive.
7. There may be business issues associated with route selection that cannot be reflected in ordinary route processing using mechanisms like OSPF or IS-IS.
8. Services will have to be created at network edge and border. Service providers may offer new services, such as voice, video, security and VPN, from network edge. Given the competition from application service providers, the operation cost must be manageable.

One of the root causes of this problem is that the conception of the service must be created over all of the networks that contribute Elements, but that conception cannot be shared among the Elements because each will "see" only its own behavior. The current IPsphere process appears to assume that the service conception is held only in the Service Order Instance and that SMS Administration, SMS Parent, and SMS Child functions cooperate to make that service conception real in terms of network behaviors.

This paper presents a solution to all these problems created through the use of a Pseudowire facilitation layer shown in Figure 1, built above the current network protocols but mapping downward to all of the popular network technologies. A proposed expansion to the Pseudowire specification will enable this new layer to create all of the service relationships (p-p, m/p, etc.) and thus support the full range of business and residential services. Controlling this layer will require a single management interface per provider jurisdiction, and this will reduce operations cost and complexity.

We are recommending that the IPSF consider standardizing this new Pseudowire layer, which we call GIBSON, and to that end take the steps needed to establish its conformance with IPsphere goals and its value in IPsphere applications. Should the body find that standardizing beyond its scope; we ask that the body endorse GIBSON to a body of appropriate jurisdiction.

GIBSON Architecture and Operation Overview

GIBSON architecture provides a feasible solution for providers to transport service-guaranteed user traffic over multi-provider network. It interfaces with IPsphere SMS framework for service discovery and mapping, and leverages standard-based MPLS Pseudowire techniques to interface with network routers and switches. In essence, GIBSON architecture binds business services to network data-plane, and enables the providers to provision, manage and monitor customer traffic for new services, such as IMS and VoD (video-on-demand).

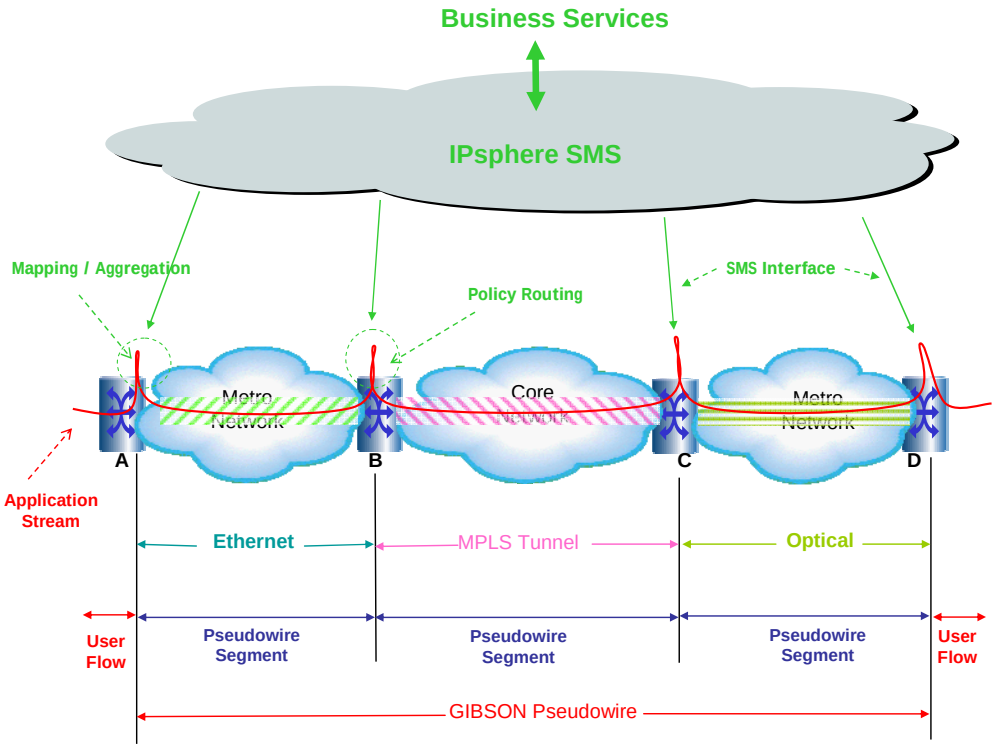

Figure 1: GIBSON: Pseudowire-based facilitation architecture

GIBSON (Global IP-Based Service-Oriented Network) architecture is to address the following:

1. Open interface for business service creation and provisioning
2. Operate in both intra-provider and inter-provider environment
3. Provide consistent edge-to-edge per-flow forwarding behavior
4. Flow type agnostics – capable of processing flows in any format
5. Support for "nesting" of pseudowires to facilitate traffic management and virtual service creation
6. Support for "virtual segments" that envelope multipoint service behavior created by capabilities like RFC 2547, to permit end-to-end multipoint delivery without N^2 *problems.*
7. Independent of underlying network transport tunneling mechanism
8. Applicable on all service devices, with less dependency on IP routing

For clarity, we outline the GIBSON operation in Figure-1: a provider network consists of two access, and one core networks, running Ethernet, optical (SONET/SDH) and IP/MPLS, respectively. And each network is administrated separately. The provider needs to support multiple services, including private leased line, video distribution and residential broadband data access.

In the example, the provider is to deliver high-quality, high-premium and long-duration video streams between end users. Through business arrangement and network planning, the provider decides to deliver user video traffic over border node A, B, C and D. Between each pair of border nodes, there may be multiple routers and switches operating with different types of control protocols. The network between A and B may be Ethernet network running IEEE 802.1ah Provider Backbone Bridging (PBB) and Provider Backbone Transport (PBT), the network between B and C may be running MPLS traffic engineering, and, finally, the network between C and D may operate over GMPLS.

The provider enables the IPsphere SMS framework to distribute policy information. At the connection setup time, the SMS will download the service-specific parameters, such as packet identification (i.e., RTP port numbers), a globally unique flow-id, and QoS information to the edge nodes A and D. Further all the nodes will get policy routing information from SMS. Each policy routing entry has the information such as: "for flow-id X, go to node Y".

Upon the reception of policy information, edge node A will trigger the establishment of a MPLS Pseudowire as defined in IETF PWE3 WG [PWE3]. The Pseudowire setup sequence will follow RFC4447, [Dry-Martini], [MHOP] and [PW-Switching]. Specifically, the Pseudowire type will be IP as defined in [L2VPN-ARP], and the flow-id and QoS information will be encoded in Pseudowire setup messages as defined in [MHOP].

When the Pseudowire setup messages arrived on node B and C, they will extract the flow-id information, search the policy routing database provided by the SMS, and route the Pseudowires accordingly. This will result in an edge-to-edge multiple-hop Pseudowire from A to D through B and C.

When video traffic arrives on the edge nodes A and D, they will map incoming packets that have the same RTP port numbers as provided from SMS to the established Pseudowire. At each hop, the traffic will receive the QoS guarantees as specified by the Pseudowire.

As shown in the example, GIBSON interfaces with SMS and leverages the existing Pseudowire technology to guarantee service-oriented traffic over the network. We refer the pseudowires that operate in the context of GIBSON as "GIBSON Pseudowires".

Background on MPLS Pseudowire

Before describing GIBSON in further detail, it may be important for us to evaluate the reasons for choosing Pseudowire as the key component for service facilitation.

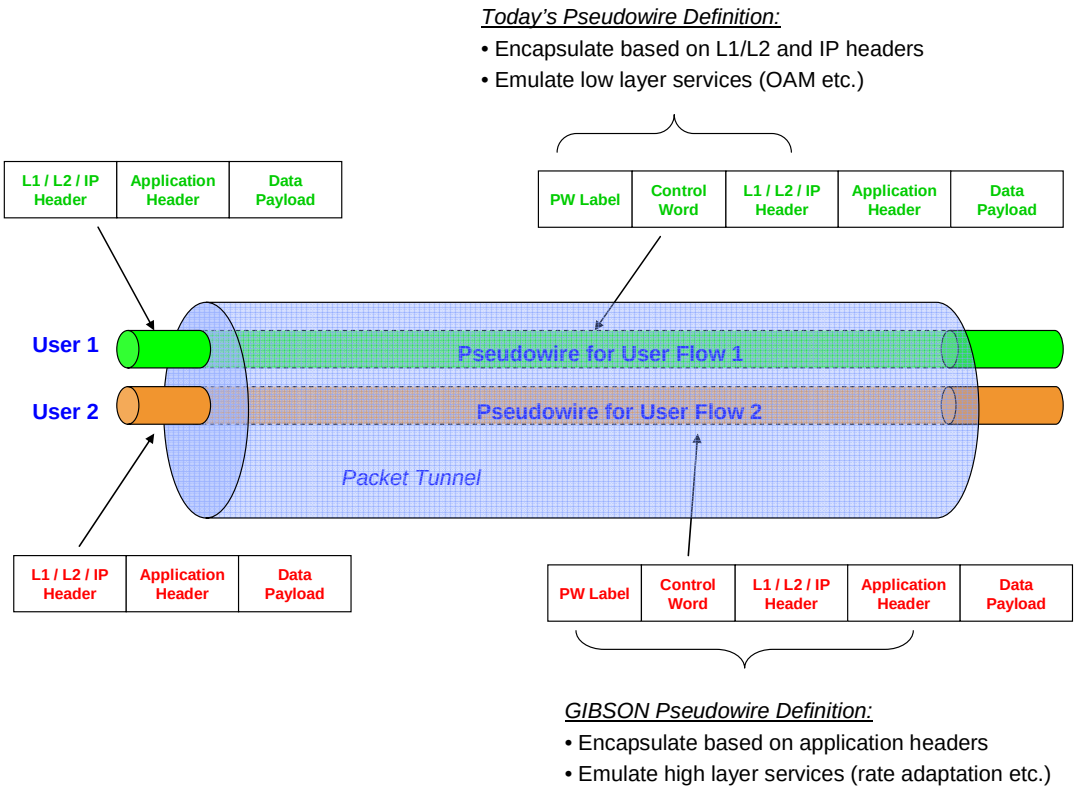

Figure 2: Pseudowire functionality in a nutshell

As shown in Figure 2, IETF-defined Pseudowire works as the following: to aggregate data flows over a shared physical or logical tunnel between network edges, the ingress edge node encapsulates a Pseudowire header to the packets. This header consists of a MPLS label, and a control word. At the egress edge node, it will stripe off the header, and forward the original packets. Each individual data flow has a unique MPLS label. The control word can be used for congestion control and OAM purposes. Per today's definition, a data flow could be a TDM trunk, ATM VPCI connection or an Ethernet VLAN flow. The edge nodes use IP routing and MPLS signaling protocols to setup Pseudowires over the network.

Followed by its initial deployment success, IETF has recently extended the functionality of Pseudowires. Some of new features include the ability to switch Pseudowires at network edge, provide QoS and OAM on per-Pseudowire basis, and support for protection and restoration.

It is important to notice the following Pseudowire characteristics:

- Transport agnostic: Pseudowires can transport data traffic over any physical or logical data tunnel which may be IP, MPLS, Ethernet, or even optical cross-connection.
- IP-friendly: Pseudowires are provisioned, controlled and operated via IP control plane
- Inter-network capable: PWE3 multi-hop [MHOP] and switching [SWITCH] techniques enable the providers to provision Pseudowires over multiple intra-domain or inter-domain networks.
- VPN capable: Pseudowire has been extended in IETF to create a nested topology for VPN applications, which include VPLS, VPWS and IPLS.
- SLA capable: Pseudowire technique can provide QoS [MHOP], protection and restoration [PROTECTION] and congestion control [CONGESTION] functionality at per-flow basis
- Flow type agnostic: Pseudowires can encapsulate any type of data flows. As defined today, pseudowires can encapsulate Layer-1 flows in SONET/SDH format (the technique is known as Circuit Emulation), Layer-2 flows such as ATM, Frame Relay, PPP and Ethernet, and IP. Note that all packets with in a Pseudowire always receive the same packet forwarding treatment throughout the network.

Base on the last point, we can easily add application-awareness dimension on Pseudowires. The deployment of new data services requires the service providers to control and management user traffic at per-application-stream granularity. A typical application stream may be encapsulated in RTP for session-based applications such as VoD and VoIP, or in MPEG for multimedia applications.

New applications bring an entirely different set of service requirements. For example, some applications can tolerate packet out-of-order delivery, some applications can tolerate packet drop, but not delay, and some applications require the user traffic to adapt to the change of available link bandwidth, but maintain constant-bit-rate at all time. Obviously, the traditional QoS mechanisms (IEEE 802.1p or DiffServ) are not adequate to handle such applications.

Today, most of the service providers must be able to support both legacy and new services. And the basic Pseudowire technology has been in deployment in their networks for years. Therefore, processing user traffic at Pseudowire-level will enable the providers to deliver a wide range of services over a common transport network, while be backward compatible with the existing routers and switches in the network.

As illustrated in the previous section, in the context of GIBSON and IPsphere, GIBSON Pseudowires embodies the following key attributes:

- GIBSON Pseudowires map and aggregate data flows from any layer (i.e., layer 1 to 7). The number of user flows to be aggregated into each pseudowire depends on services. For example, for high-bandwidth and long-duration VoD streams, the edge nodes may initiate and maintain one pseudowire per stream. For low-speed and short-lived VoIP sessions, the edge nodes may aggregate a large number of sessions into a single pseudowire. Once again, all packets within each pseudowire will receive the same SLA treatment throughout the network.
- At network border, GIBSON pseudowires are switched based on business-driven routing. As shown in Figure 1, nodes B and C situate at the border between metro and core networks, and are connected to edge nodes A and B over physical or logical tunnels. When they receive data packets from GIBSON pseudowires, the pseudowire traffic will only be forwarded to other tunnels based on the business bilateral/multilateral agreement.

- GIBSON pseudowires are provisioned as a result of SMS Administration, SMS Parent, and SMS Child communication. Gibson Pseudowires also recognize multipoint transport behavior and can exploit it at the service level to facilitate multipoint services

GIBSON Architecture in Detail

We refer GIBSON endpoints (or GIBSON-Enabled Nodes) as the devices that are responsible for processing user data flows by interfacing with IPsphere SMS child. Between the GIBSON endpoints, there could be one or multiple networks. The GIBSON endpoints always operate at network edge and/or border.

Figure-2 indicates the interfaces in GIBSON architecture. In GIBSON-enabled networks, all data flows, best-effort or otherwise, are always provisioned and established through the interface with SMS.

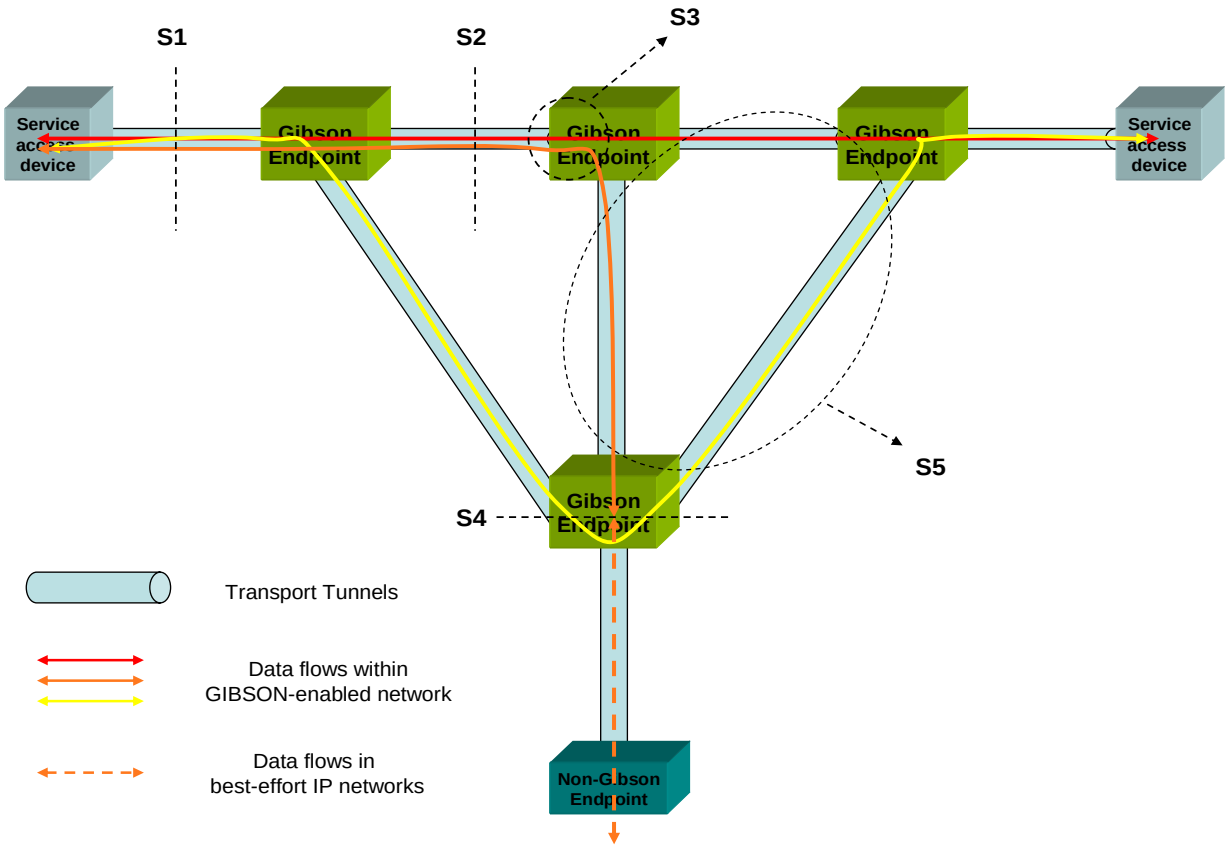

Figure 3: GIBSON architecture overview

S1 - Access Interface

This is where GIBSON Endpoints aggregate user flows from service access devices. The service access devices may include a wide range of equipment including IP DSLAM, PON, SONET/SDH MSPP and wireless base-stations. The service access devices do not need to be fully IP routing capable. Instead, they need to be efficient in supporting the services that they are designed to do.

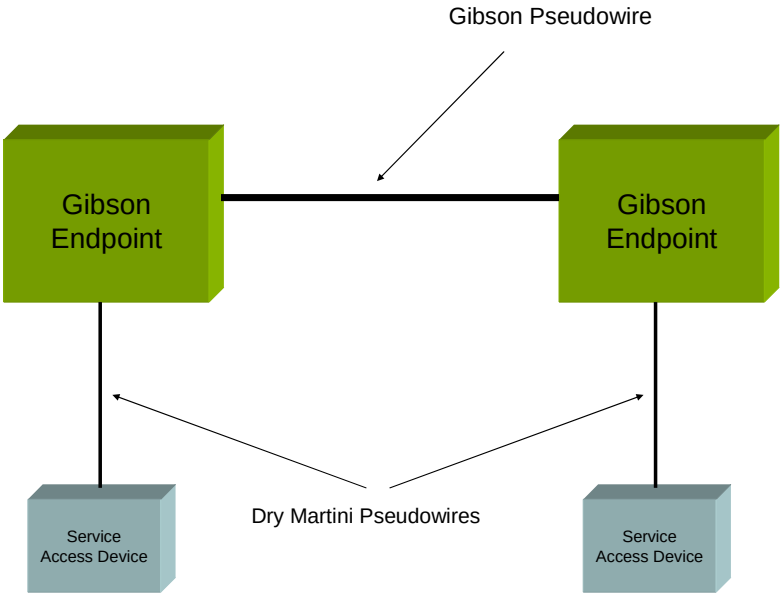

Figure 4: GIBSON S1 interface

Service access devices can deliver user flows in any native format (such as Ethernet VLAN and IP). Upon the reception of user flows from service access devices, the GIBSON endpoints will classify the incoming traffic based on the application criteria provided from SMS.

For example, two service access devices have established an IMS session via SIP. To provide service guarantees, the GIBSON Endpoints will read the IP source and destination address, and UDP source and destination port numbers for each incoming packet, and match them against the customer flow policies. If there is a match, the GIBSON Endpoint will encapsulate the packet with a Pseudowire header, and forward it over a pre-established Pseudowire. Each Pseudowire is provisioned with the SLA parameters provided by SMS.

Since the access network topology is relatively simple, the service access devices do not need to be routers. Therefore, there is little use in running complex IP routing and signaling protocols to setup pseudowires over MPLS networks, as defined by IETF [PWE3-CTRL]. The access devices may use a much simpler approach in Pseudowire setup as defined in [Dry-Martini].

S2 – Data Aggregation

Between two GIBSON endpoints, user data flows are encapsulated as Pseudowires and aggregated into transport tunnels according to the criteria provided from SMS.

Each GIBSON Pseudowire may consist of multiple user data flows. As an example, multiple VoIP calls (all in RTP) going between two GIBSON endpoints may share the same Pseudowire header. In this case, the GIBSON Pseudowire will be over-provisioned with a fixed bandwidth, which is computed based on call arrival and departure distribution at business level, and downloaded to the GIBSON endpoints via SMS. This would be a good technique to transport a large number of short-lived voice calls over packet networks with service guarantees. This can potentially reduce the number of control-plane messages during call setup.

Each transport tunnel may aggregate multiple GIBSON Pseudowires. The transport tunnel could be any type preferred by the service providers. It could be MPLS LSP established via RSVP-TE or LDP in core networks, and could be Ethernet PBT tunnels in metro networks. In addition, GIBSON supports the nesting of pseudowires so that pseudowires can be used to create transport tunnels that transit multiple lower-level tunnel technologies. See the section on “Virtual Segments” below.

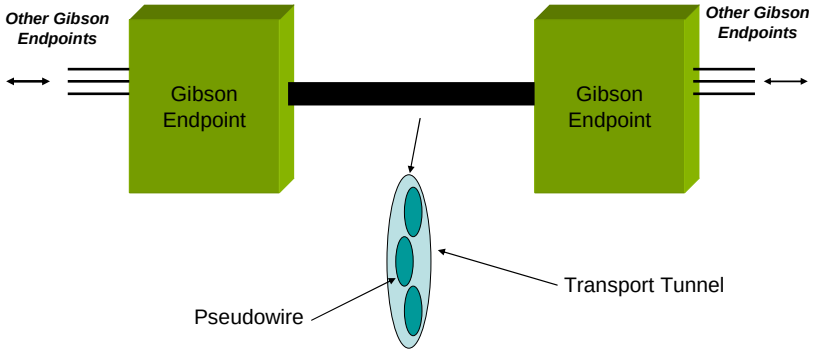

Figure 5: GIBSON S2 interface

The Pseudowire aggregation process may be bounded by the business policies distributed via IPsphere SMS. The policies may determine the allocation of pseudowires into transport tunnels, and the modification of transport tunnels to accommodate the Pseudowire network resource consumption.

S3 – Pseudowire Routing

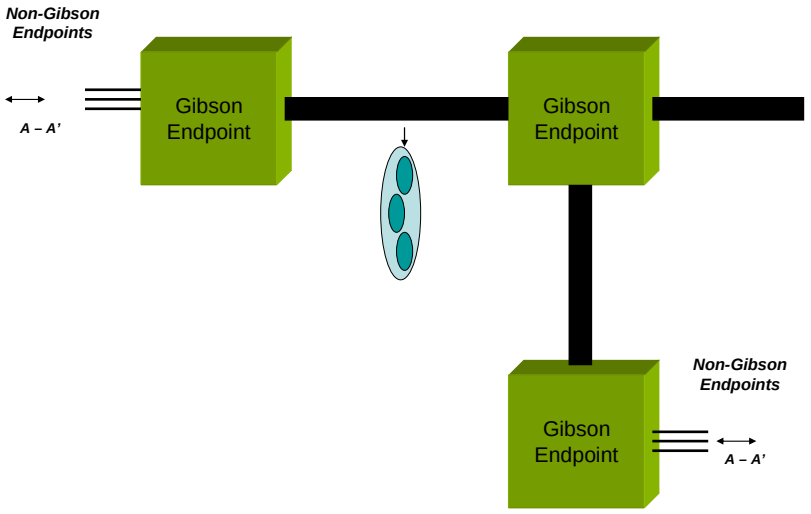

Figure 6: GIBSON S3 interface

Within the network, the intermediate GIBSON endpoints will terminate the transport tunnels, and switch the pseudowires into another set of tunnels toward their destination. At GIBSON Pseudowire setup time as defined in [MHOP], each Pseudowire request message will carry information such as QoS, protection and congestion control data and user flow information.

In some applications, such as IMS, the user flow and QoS information will be derived from the control protocols. The QoS can be piggybacked by the Pseudowire-initiating GIBSON Endpoints into the setup messages. The transit GIBSON Endpoints will use this information for routing. Another alternative is to store the QoS information on each provider's Router Servers.

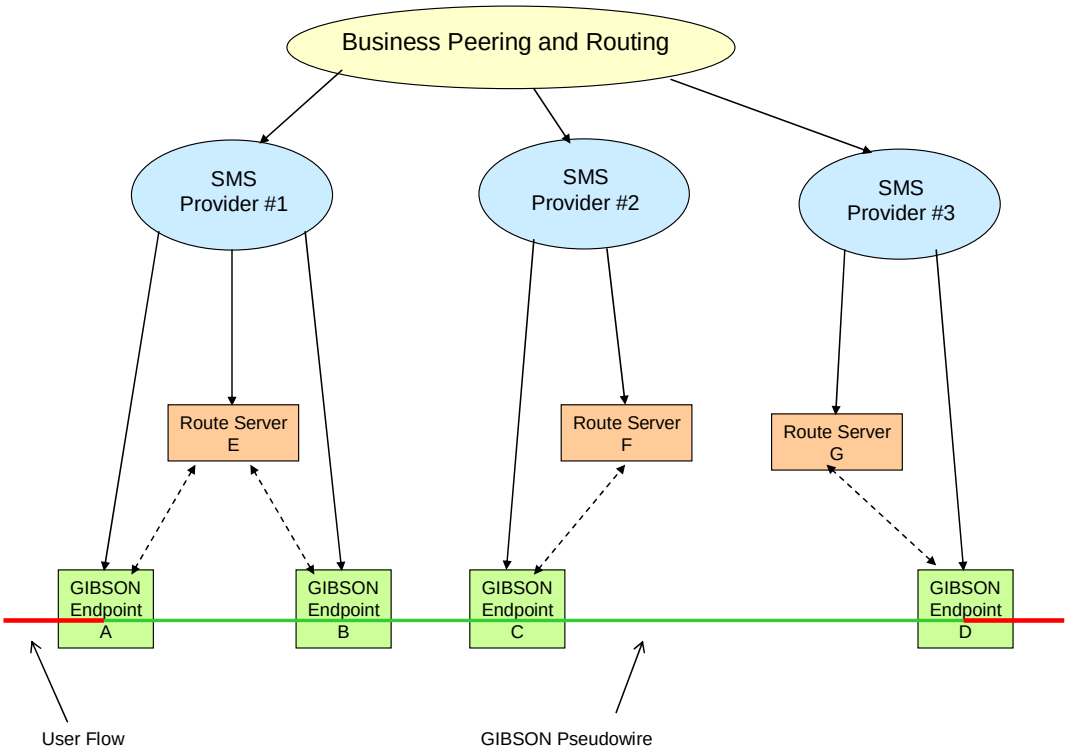

Figure 7: GIBSON Routing Illustration

Figure 7 illustrates the operation of GIBSON-based Pseudowire routing. There are two types of routing here: business-based, and topology-based. Business-based routing is conducted between service providers. As shown here, provider 1, 2 and 3 define the business routing constraints from their inter-provider bilateral or multilateral agreement. Through SMS, the business routing constraints will be downloaded to the Route Servers within each provider network.

A typical Router Server can be the one defined and studied in IETF PCE WG. It is to combine the business routing constraints and network internal topology and resource information, and compute the optimal path between network edges.

In Figure 7, there may be multiple routers and switches between GIBSON Endpoints. After mapping the data flow into Pseudowires, GIBSON Endpoint A will interface with Route Server E to find a suitable path to GIBSON Endpoint B. On transit GIBSON Endpoint C, it will again interface with Route Server F for the optimal path to reach GIBSON Endpoint D. This is how inter-domain GIBSON Pseudowires are established.

Another alternative for inter-domain routing is to simply use BGP (2547bis) to distribute GIBSON capability within the network. The GIBSON Endpoints can therefore learn about the overall network topology for Pseudowire forwarding. Note that this may be the preferred method for forwarding non-GIBSON-enabled pseudowires.

A collection of pseudowire/transport tunnels can be defined as a "virtual segment" in GIBSON, creating a cloud that can be addressed as an entity and used for routing services based on QoS or other properties. See the section on "Virtual Segments" below.

S4 – Pseudowire Termination

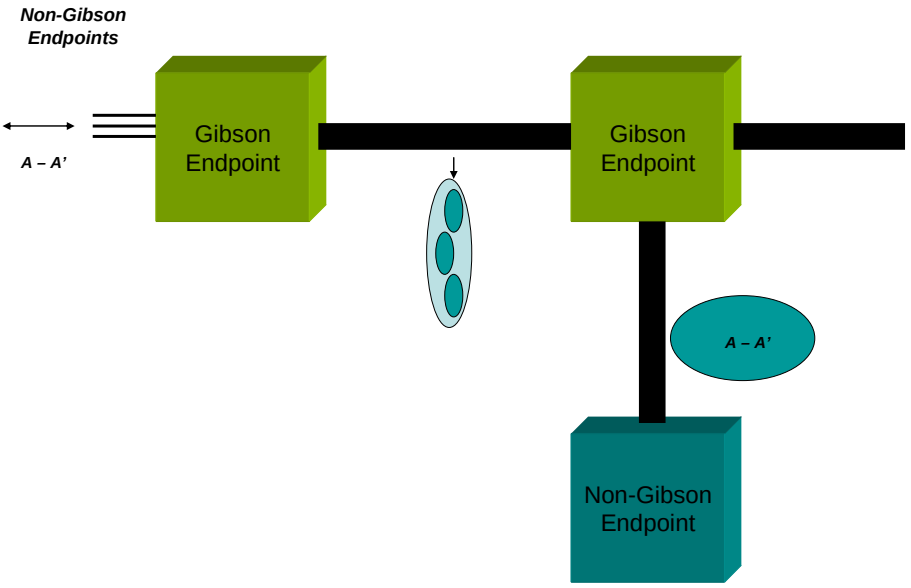

Figure 8: GIBSON S4 interface

When interfacing with Non-GIBSON-enabled nodes, the GIBSON endpoints will terminate the pseudowires and forward the original users traffic. If the network between GIBSON and Non-GIBSON Endpoints does not have per-flow service guarantees, the service providers will be notified about partial service guarantees on data flows. Flow information carried at Pseudowire setup time will permit the terminating process to map the flow correctly onto the egress network.

GIBSON "Virtual Segments"

A GIBSON pseudowire termination is a routable entity in GIBSON, so traffic can be routed onto a transport tunnel for delivery to the location where that tunnel terminates, using the nested pseudowire property of GIBSON that is analogous to nested LSPs or the ATM VCI/VPI hierarchy. We propose an additional enhancement to GIBSON to facilitate the definition of a set of GIBSON pseudowires as a virtual segment.

Virtual segment is a collection of GIBSON pseudowires that operate as a transport subnetwork. The virtual segment is bounded by GIBSON endpoints, and at these endpoints the services of the subnetwork are uniformly available. We propose to treat this virtual segment as a GIBSON segment, so that traffic can be routed onto it. When the GIBSON virtual segment is addressed as a hop in a GIBSON pseudowire, the routing logic will resolve the request to the nearest GIBSON endpoint on-ramp to the virtual segment.

GIBSON virtual segments can enclose any of the following:

1. A network of GIBSON pseudowires.

2. A network of MPLS LSPs or other tunnels.
3. A multipoint network based on 2547bis [RFC4364] or other multipoint technologies.

We offer more detail on the last option below.

S5 – Multipoint Transport

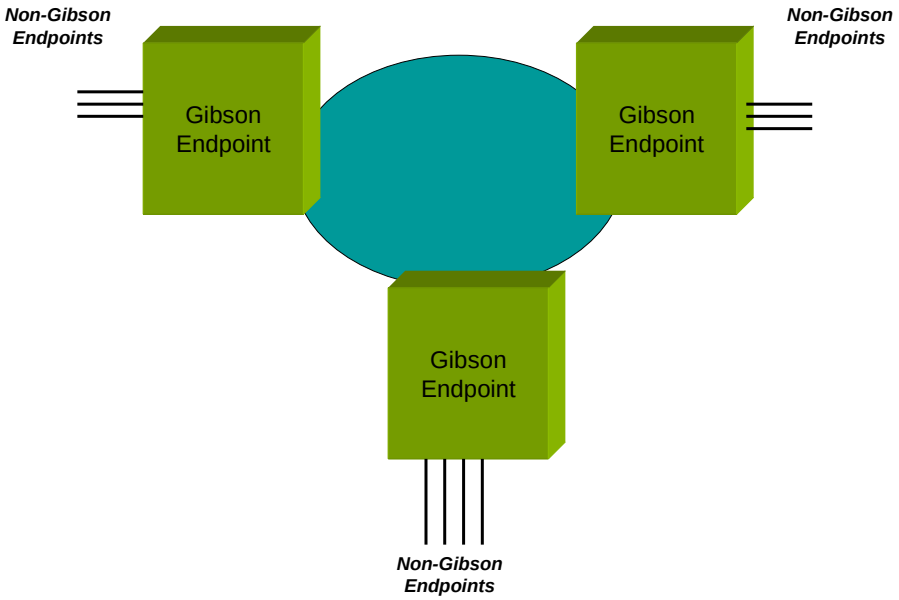

Figure 9: GIBSON S5 interface

GIBSON supports natural multipoint service behavior by enveloping the multipoint technology as a virtual segment. GIBSON pseudowires routed to the multipoint virtual segment will be forwarded to the nearest on-ramp (a GIBSON endpoint), where the outer GIBSON envelope will be stripped and the GIBSON endpoint destination address will be decoded from the header. The packet will then be forwarded to the GIBSON endpoint destination using the multipoint technology of the virtual segment. At that endpoint, a GIBSON header will be added and the pseudowire will continue as though the tunnel was continuous, but no pseudowire meshing will be required within the virtual segment.

The GIBSON endpoints can interface with IPsphere SMS to discover network topology. In turn, it can use the standard VPN mechanism, such as 2547bis option 2(b) [RFC4364], VPWS and IPLS to construct point-to-multipoint VPN networks. Note that the Pseudowire is not multipoint in itself, but it can be routed over a multipoint transport process.

A User Case for IMS

The GIBSON architecture can apply to a number of services. As shown in Figure-1, it can provide service guarantees for VoD applications in a heterogeneous networking environment. In this section, we will go over a more complex and dynamic application, IMS.

IMS users SIP for end-to-end session setup. Each session can be a VoIP call, a VoD session or a simple text messaging connection. The default data transport protocol is RTP. One of the key characters in SIP-based communication is that

control-plane is out-of-band. As a result, media traffic may take a path that is completely different from the one taken by SIP control messages. In today's network, the service providers cannot control the media traffic other than forwarding them as IP packets, and rely on DiffServ for QoS.

Since VoIP traffic does not require much bandwidth and well behaved (mostly CBR), service providers do not need to be overly concerned about end-to-end service guarantees. However, for high-volume broadband customers, video stream service guarantee may become important.

In this section, we will describe how GIBSON can provide service guarantees for SIP-based data traffic.

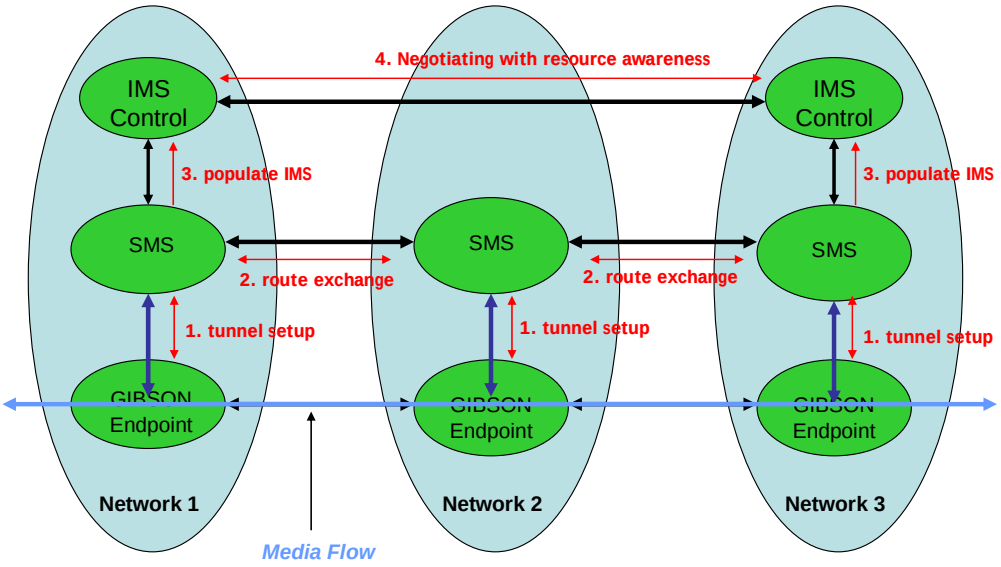

Figure 10: GIBSON for IMS (1)

1. Initially, GIBSON Endpoints in Network 1, 2 and 3 will setup transport tunnels among each others. The transport tunnels are either initiated by the SMS or by the network operators. In the latter case, the tunnel information will be reported to the SMS.
2. Through SMS, the networks will exchange the resource and topology information. The information will only be specific and relevant to the services that the network owners have agreed to support.
3. Access networks 1 and 3 have their own IMS (a.k.a. IMS domains). Through SMS, the IMS will populate the database (e.g. HSS) with relevant network resource information.
4. The end users in both networks will negotiate multimedia sessions via CSCF. During session negotiation, network resource availability will become one of the key parameters. If there is not enough resource between Network 1, 2 and 3 for a particular session, CSCF will reject the session.

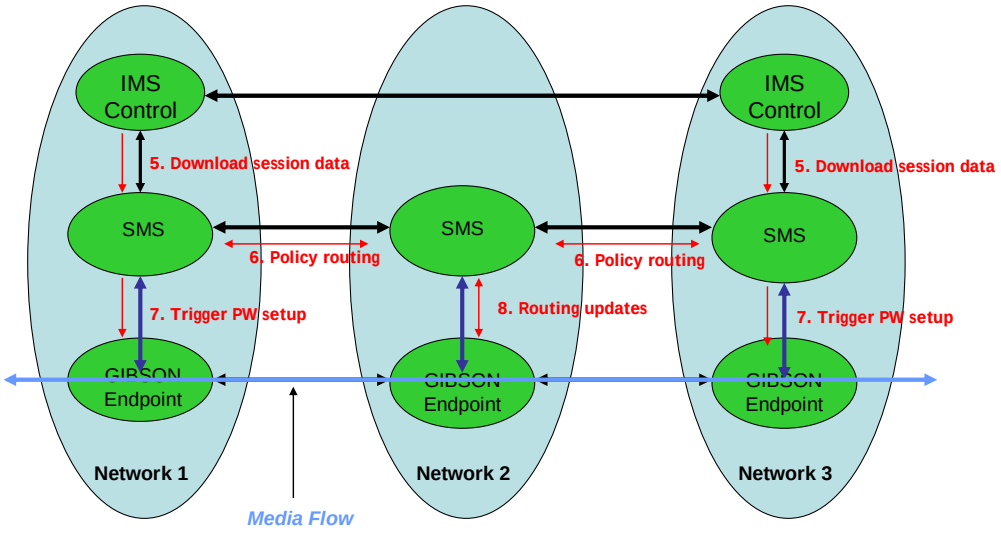

Figure 11: GIBSON for IMS (2)

5. When the sessions have been established, IMS will notify the SMS about the session data flow information (that is, RTP session).
6. The SMS exchanges the network resource information among each other.
7. The SMS downloads the data flow information to the GIBSON Endpoints in Network 1 and 3.
8. In transit network 2, the SMS will compute the routes from Network 1 and 3, and download the routing information to the GIBSON Endpoint. The route computation algorithm is governed by the service itself, which may override the classical SPF/BGP path computation. Further, the SMS is not necessary the right place for route computation. As we have described before, the SMS may interface with Route Servers for such task. The details of routing criteria will be studies in the future.

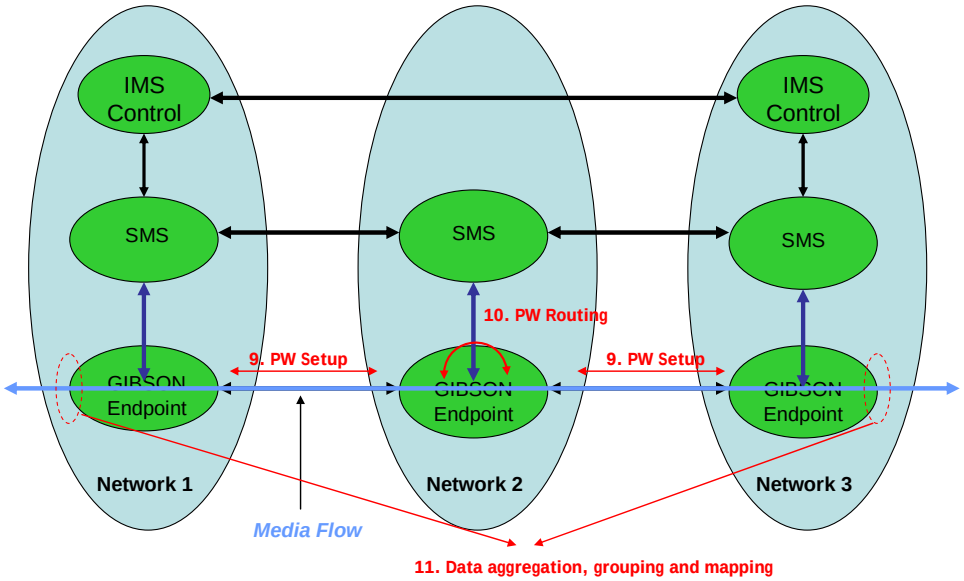

Figure 12: GIBSON for IMS (3)

9. The GIBSON Endpoints in Network 1 and 3 will setup pseudowires for the session media flow. The Pseudowire type depends on the application itself. In this case, the GIBSON Endpoint may choose IP as the Pseudowire type.
10. When Pseudowire setup messages arrive on Network 2, the GIBSON Endpoints would use pre-installed policy routes to direct the pseudowires toward Network 3.
11. When end-user data traffic in Network 1 and 3 arrive on the GIBSON Endpoints, they will be mapped to the corresponding Pseudowires. In this case, the GIBSON Endpoints will classify each packet base on its IP source and destination address, and UDP source and destination port numbers, all of which have been downloaded from the SMS in Step (7).

In this approach, two levels of admission control will take care. At control-plane, SMS provides network resource information to IMS, which may result in call admission control at SIP setup time. At data-plane, SMS provides per-session flow information to GIBSON Endpoints, which will aggregate each user flow in the form of Pseudowire to transport tunnels.

As a result, data flows will have QoS guarantees throughout the network.

Acknowledgement

Since the Ottawa IMS Workshop where we first presented the work, we have had discussion with several people, and received a detailed feedback from Mr. Christian Jacquenet. We thank them for their guidance.

References

[PWE3] http://www.ietf.org/html.charters/pwe3-charter.html

[PCE] http://www.ietf.org/html.charters/pce-charter.html

[L2VPN-ARP] IETF Draft, H. Shah, et al "ARP Mediation for IP Interworking of Layer 2 VPN"

[MHOP] IETF Draft, L. Martini, et al, "Dynamic Placement of Multi Segment Pseudo Wires", work in progress

[SWITCH] IETF Draft, L. Martini, et al, "Segmented Pseudo Wire", work in progress

[PROTECTION] IETF Draft, P. Pan, et al, "Pseudowire Protection", work in progress

[CONGESTION] IETF Draft, P. Pan, et al, "Pseudowire Congestion Control", work in progress

[Dry-Martini] IETF Draft, P. Pan, et al, "Dry-Martini: Supporting Pseudo-wires in Sub-IP Access Networks", work in progress

[PWE3-CTRL] IETF RFC 4447, L. Martini, et al, "Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)"

[RFC4364] IETF RFC 4364, E. Rosen, et al, "BGP/MPLS IP Virtual Private Networks (VPNs)"