AUSTRALIAN
COMMUNICATIONS
INDUSTRY FORUM

INDUSTRY DISCUSSION PAPER
QOS-BASED VOIP SERVICE INTERCONNECTIVITY

22 MARCH 2006

Industry Discussion Paper: QoS-Based VoIP Service Interconnectivity

ACIF MARCH 2006

Authors: Shara Evans and Richard Chirgwin, Market Clarity
(www.marketclarity.com.au)

Disclaimers

1. Notwithstanding anything contained in this Industry discussion paper, ACIF disclaims responsibility (including where ACIF or any of its officers, employees, agents or contractors has been negligent) for any direct or indirect loss, damage, claim, or liability any person may incur as a result of any:
   (i) reliance on or compliance with this Industry discussion paper;
   (ii) inaccuracy or inappropriateness of this Industry discussion paper; or
   (iii) inconsistency of this Industry discussion paper with any law.
2. The above disclaimers will not apply to the extent they are inconsistent with any relevant legislation.

Copyright

Request for Submissions

At the second Australian Communications Industry Forum (ACIF) VoIP Forum held in December 2005, Quality of Service (QoS) was identified as one of the major concerns relating to VoIP interconnectivity.

This paper examines the provisions of internet-based services that support Quality of Service (QoS) as a means of improving the VoIP user experience. It provides a mini-tutorial, including outlines of technical and business models relevant to the discussion of inter-domain peer-based QoS. An examination of whether an industry-wide approach to VoIP peering is also included. A number of international VoIP peering initiatives are discussed.

To further identify the steps that need to be taken by the industry to address VoIP QoS issues, ACIF seeks submissions in response to the questions and matters for comment raised in this paper. A glossary of acronyms is provided at the back of the paper to assist when reading the discussion paper. **A full list of the discussion paper questions is available in the Appendix of the document**.

About ACIF

ACIF is a member-funded organisation established in 1997 to lead industry involvement in defining the communications environment. ACIF provides a neutral forum in which all participants and end-users in the Australian communications industry can work together to foster an efficient, competitive environment.

Submissions should reach ACIF no later than **Friday 26 May 2006**.

Please send submissions (preferably in soft-copy form) to acif@acif.org.au or fax to 02 9964 6136 or post to:

Australian Communications Industry Forum
PO Box 444
Milsons Point NSW 1565

A copy of all submissions will be made publicly available on the ACIF website. Please note in your submission whether you wish for your submission to be treated as confidential.

TABLE OF CONTENTS

1 INTRODUCTION: INTER-DOMAIN PEER-BASED QUALITY OF SERVICE IN VOIP ENVIRONMENTS

This discussion paper is designed to examine the provision of Internet-based services that support Quality of Service (QoS) as a means of improving the VoIP user experience.

For private networks (LANs and WANs), Quality of Service support is considered necessary to maintain acceptable performance for real-time traffic such as VoIP, without subjecting that traffic to performance degradation caused by other corporate applications, particularly on limited private WAN links. Although the price of bandwidth has fallen considerably on a unit basis since the advent of competition, this has been matched by expanding demand for bandwidth, so corporate VoIP users find it still necessary to manage network QoS in such a way as to protect the integrity of their voice services.

Consumer VoIP services which traverse the public Internet are not supported by QoS technologies, for several reasons:

- Many ISPs do not use QoS-marking for traffic in their networks, preferring instead to rely on capacity management (purchasing more Internet capacity; expanding their Internal network capacity; and implementing technologies such as caching to avoid content unnecessarily traversing Internet links).
- While a growing number of providers use QoS-capable networks, QoS specifications are determined on an individual provider basis. To pass QoS tags between networks would therefore require a bilateral (or multi-lateral) arrangement in which providers agree not only to respect each others' QoS markings, but also on the specifications indicated by each others' markings. To avoid a raft of bilateral agreements, a multi-lateral agreement within an ACIF framework could simplify peering considerably (if the agreement was reached soon enough).
- If Inter-domain QoS also involves differentiated tariff mechanisms, it would also require a new layer of commercial arrangements between providers.

The scope of this document is restricted to the behaviour of conversations between Internet-connected VoIP clients. It does not consider the behaviour of VoIP-originated calls terminating to the PSTN.

While Australian Internet providers do not yet use inter-domain QoS, some traffic engineering can be accomplished by ISPs that participate in peering arrangements.

Peering allows ISPs to exchange traffic with each other, without having that traffic traverse upstream transit services. This reduces ISPs costs (since it limits the growth in transit service needs), and has the collateral benefit of reducing the number of route hops needed for sessions which traverse the peering link.

If ISPs begin to implement inter-domain QoS capabilities as part of their traffic management, the interaction between QoS regimes and ISP peering will need to be considered.

Issues that the industry should consider include:

(a) Is end-to-end QoS feasible or desirable for consumer Internet services?

(b) What are the technical prerequisites to achieving inter-domain QoS support on consumer Internet services?

(c) What level of co-operation between providers is necessary to achieve inter-domain QoS, and is such co-operation achievable? For example, are providers willing to give preferential treatment to traffic sourced from other providers' networks?

(d) Does a complete suite of industry technical standards exist to support inter-domain QoS? If not, is the local development of such standards an appropriate effort for the Australian industry?

(e) If service providers were to adopt QoS technologies and begin exchanging inter-domain QoS information, would this have an impact on the infrastructure already in place supporting inter-domain "data plane" peering? Are peering exchange providers willing or able to accept QoS requests from client networks, or would this negatively impact either the peering fabric or the services available to other customers of the QoS-based networks?

(f) Do emerging models for "signalling plane" peering in VoIP services have any role to play in supporting the development of inter-domain QoS?

(g) Is there a requirement for an industry-wide strategy to support the implementation of inter-domain QoS in Australia? Should this be limited to public Internet interconnectivity, or should it include carrier IP/MPLS networks as well? For example, would ACIF have a role in helping to coordinate the phased introduction of relevant industry standards and/or model terms and conditions as new technologies and/or services become available?

(h) To what extent should any peer-based inter-domain QoS initiatives support signalling plane activities? Should important network utility services such as DNS be considered as part of the suite of services requiring inter-domain QoS support?

(i) To what extent can both intra-domain and inter-domain QoS mechanisms improve the overall system reliability achieved by VoIP services?

(j) What alternatives exist to the creation of an inter-domain QoS fabric?

(k) Can ACIF play a role in facilitating the standardisation of contractual terms and conditions necessary for providers to standardise inter-domain QoS definitions, and to honour the QoS requirements of traffic originating outside their own networks?

(l) To what extent, if any, would the adoption of QoS regimes violate the more general Internet principle of "network neutrality"?

(m) How can QoS be implemented across today's Internet access networks? Is it worth considering consumer offerings that provide premium services, say with lower contention ratios, which are designed to support real-time applications? Is an industry-wide strategy required to address this specific area?

(n) The current discussion is based on Internet connectivity only. Is there also a need to discuss carrier-provided interconnect between PSTN-based and Internet-based VoIP services? To what extent can such a discussion progress at this time, given the lack of international standards for carrier grade VoIP interconnect from organisations such as the ITU and ETSI?

(o) Given the range of issues raised above, and elsewhere in this paper, what timeframes and priorities should be applied to each activity?

This document includes outlines of technical and business models relevant to the discussion of inter-domain peer-based QoS. These include:

- A background to VoIP technical models as relevant to inter-domain QoS and peering issues;
- A background discussion of Internet peering models;
- A background discussion of emerging VoIP peering models;
- Technical models for inter-domain QoS; and
- Business models for inter-domain QoS.

However, it should be emphasised that comprehensive technical tutorials are beyond the scope of this document.

In the course of researching this discussion paper, it has become clear to the authors that some movement towards resolving the interaction between VoIP services, QoS, and Internet peering is being undertaken overseas. Some of these are described in Section 8, *International Initiatives.*

These initiatives, however, are very much at a proof-of-concept stage, and are being undertaken as trials between a limited number of partners.

Early work is also underway at the IETF, which in November 2005 produced a new charter for a working group called "VOIPEER." In January 2006, the "SPEERMINT" (Session Peering for Multimedia Interconnect) working group replaced "VOIPEER." Its action items for 2006 are to produce informational RFCs covering the SPEERMINT routing architecture, and the message flows associated with its routing architecture. SPEERMINT expects to submit its first proposed standard in 2007.

It appears that Australia, through ACIF, is the only place where these issues are being approached from a whole-of-industry perspective.

2 VOIP BACKGROUND

2.1 VoIP: Discovery, Signalling, Management and Conversation

The behaviour of VoIP is fundamentally different to that of a PSTN telephone.

Although telecommunications networks are now digitised, the circuit model used in PSTN telecommunications presumes that a dedicated channel can be established between two endpoints; and that the endpoints make no contribution to network functions.

A simplified outline of the process of placing a PSTN call is shown in Figure 1, below.

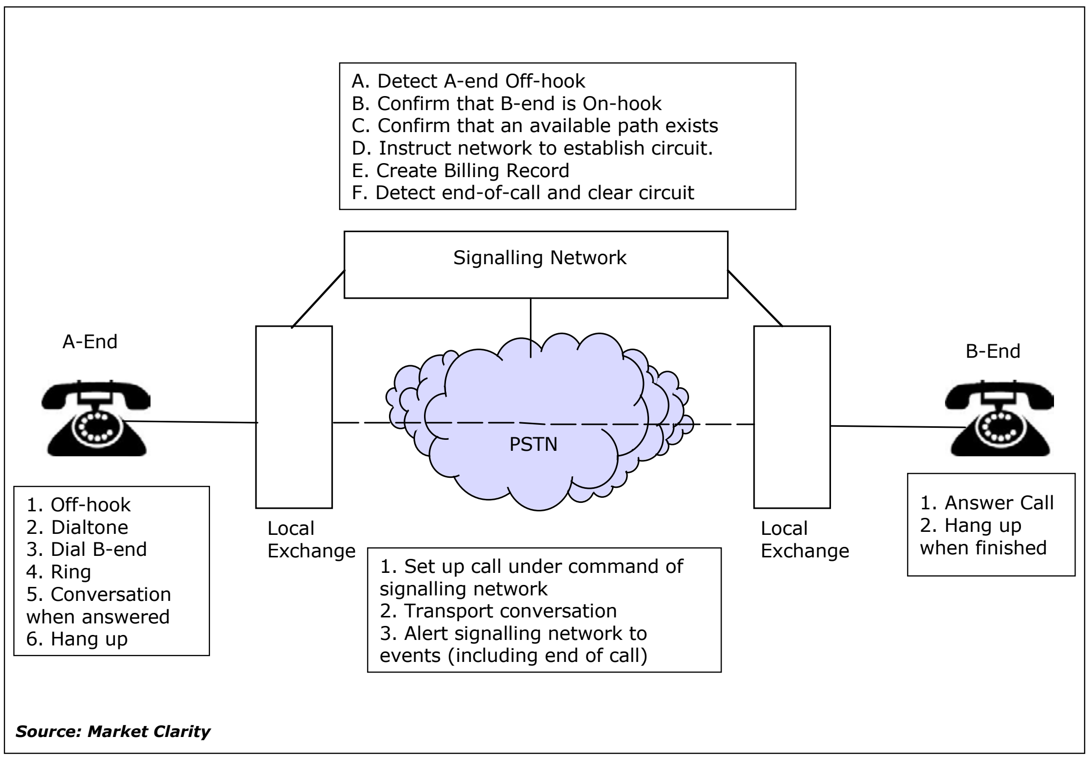

Figure 1
PSTN Signalling (Simplified)

This illustrates the importance of a part of the network overlooked by most users: the signalling system. The signalling system exists as an overlay network logically (if not physically) distinct from that part of the network which carries the conversations. This is described as out-of-band signalling.

If the call has to traverse different carriers' networks (for example, if an Optus customer calls a Telstra customer), the two carriers' signalling networks will negotiate the availability of a suitable path, and will collect billing information up to the two carriers' respective interconnect points for later settlement.

While some characteristics of the PSTN are made obsolete by the advent of IP-based networks, some of the functions outlined above must at least be replicated for a viable service. These include:

(a) **Discovery —** Traditional PSTN user discovery is rudimentary, being limited to identifying whether the called number exists (that is, whether the dialled number is associated with a physical port in the user's local exchange).

(b) **Signalling —** Whether on the PSTN or the IP network, signalling fulfils the crucial role of determining the availability of the called party, determining the availability of the path to the called party, and instructing the network to carry the call.

(c) **Management —** Both PSTN and IP networks need to collect at least a minimum of management data, such as the user database as well as network monitoring and usage data.

(d) **Conversation —** Both PSTN and IP networks need to successfully transport communications between the two parties, once a session is established.

To understand the role of peering in the VoIP market, it is necessary to understand the network "conversations," which take place before, during and after a VoIP call.

Aspects of these functions in PSTN and IP networks are compared in Table 1, below.

Table 1
Voice Session Functions in PSTN and IP Networks

Function	PSTN Networks	IP Networks
Discovery	• Existence of called party • State of called party (binary on/off hook)	• Rich discovery mechanisms possible, based on use of location servers • Users can self-manage their availability
Signalling	• Out-of-band signalling • Call success/failure dependent on end-to-end availability of network resources	• In-band signalling • Call success/failure dependent on instantaneous state of network resources between users
Management	• Out-of-band management • Management traffic traverses dedicated network	• In-band management • Carriage of management traffic dependent on network state
Conversation	• Symmetrical experience — both ends of conversation reasonably aware of the state of the connection	• Asymmetrical response — users may experience different service quality at each end of the conversation

2.1.1 VoIP Technical Models

Another important difference between PSTN call models and VoIP call models is that there currently exist at least three competing architectures for the transport of voice packets over the Internet protocol.

In the absence of any standard nomenclature for the different models, we refer to them as:

(a) The Standalone Peer-to-Peer Model — This is in relative disuse on the modern Internet, since it depends on users having fixed and easily discoverable IP addresses. It was, however, the model employed by early open source voice solutions such as Speak Freely. What distinguishes early VoIP products was that they required no VoIP-specific supporting infrastructure. By contrast, even when technically oriented users self-configure VoIP boxes, infrastructure such as STUN (Simple Traversal of UDP NAT) may be required.

(b) The Client-Server Model — Servers are deployed to perform address resolution and other signalling functions on behalf of the user agent. These servers may be located either in a VoIP service provider or within customer networks. In the VoIP service provider market, these servers may be accessed either over the public Internet or via private IP networks. Depending on implementation practices, some VoIP services will not function if the server becomes unavailable, even if the user agents are already communicating. In others, user agents can continue an established session in the absence of the server.

(c) The Node-based Peer-to-Peer Model — Skype is the best-known example of this mode of operation. Each client also performs a degree of processing on behalf of the network as a whole.

STUN **(Simple Traversal of UDP through NATs [Network Address Translation]) is a protocol for assisting devices behind a NAT firewall or router with their packet routing[1]. STUN is defined in IETF RFC 3489.**

As is discussed in Sections 4 and 5, these models have an impact on the effectiveness of QoS mechanisms. Each of these models is illustrated below. Please note that these are simplified diagrams in which not all network elements are shown.

[1] http://www.voip-info.org/wiki-STUN

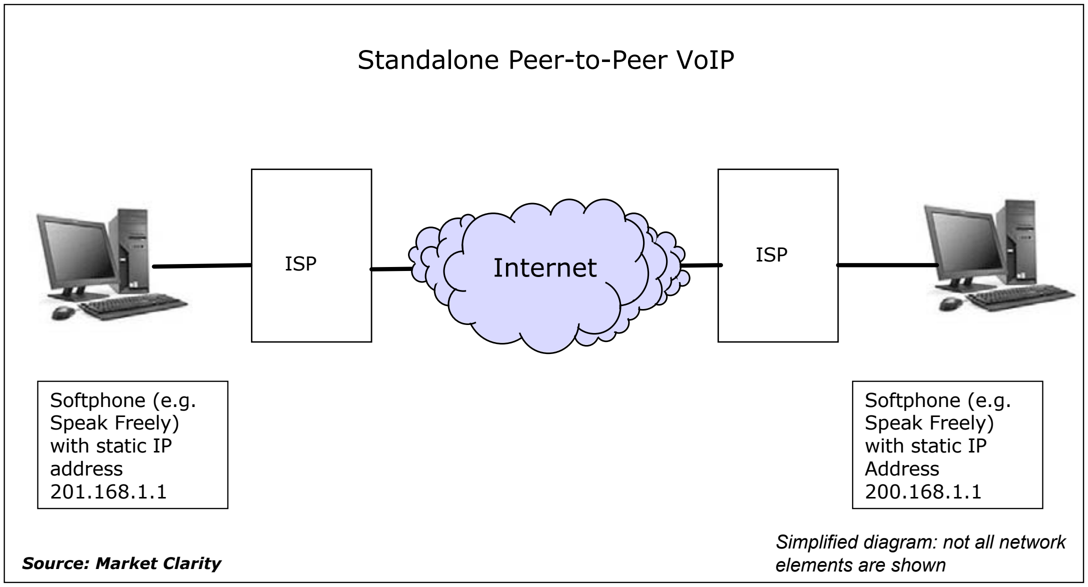

Figure 2
Standalone Peer-to-Peer VoIP

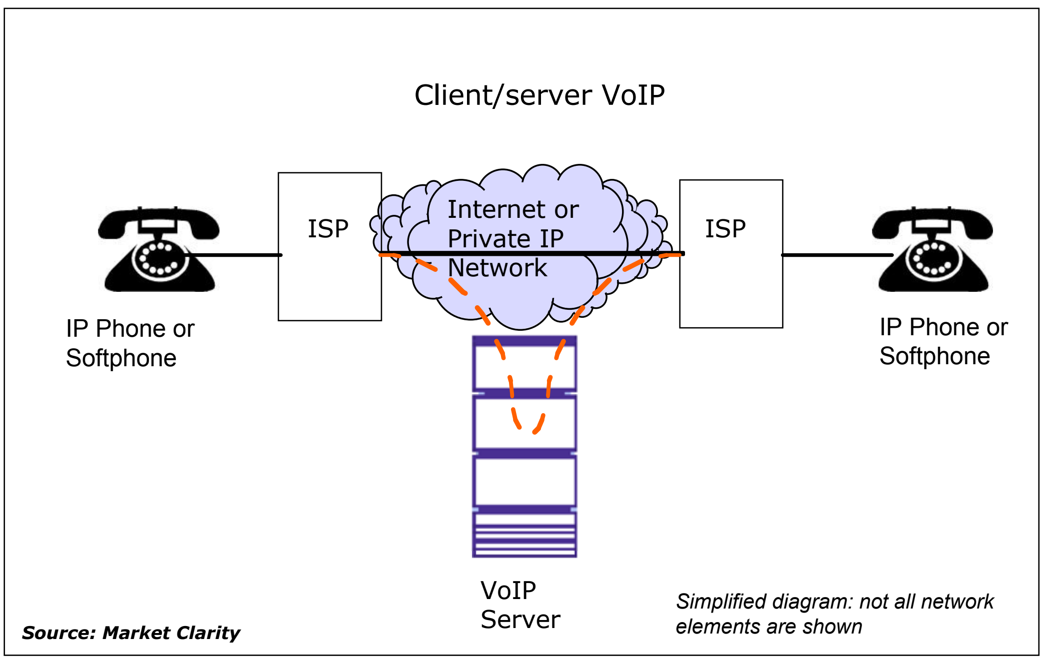

Figure 3
Client/Server VoIP

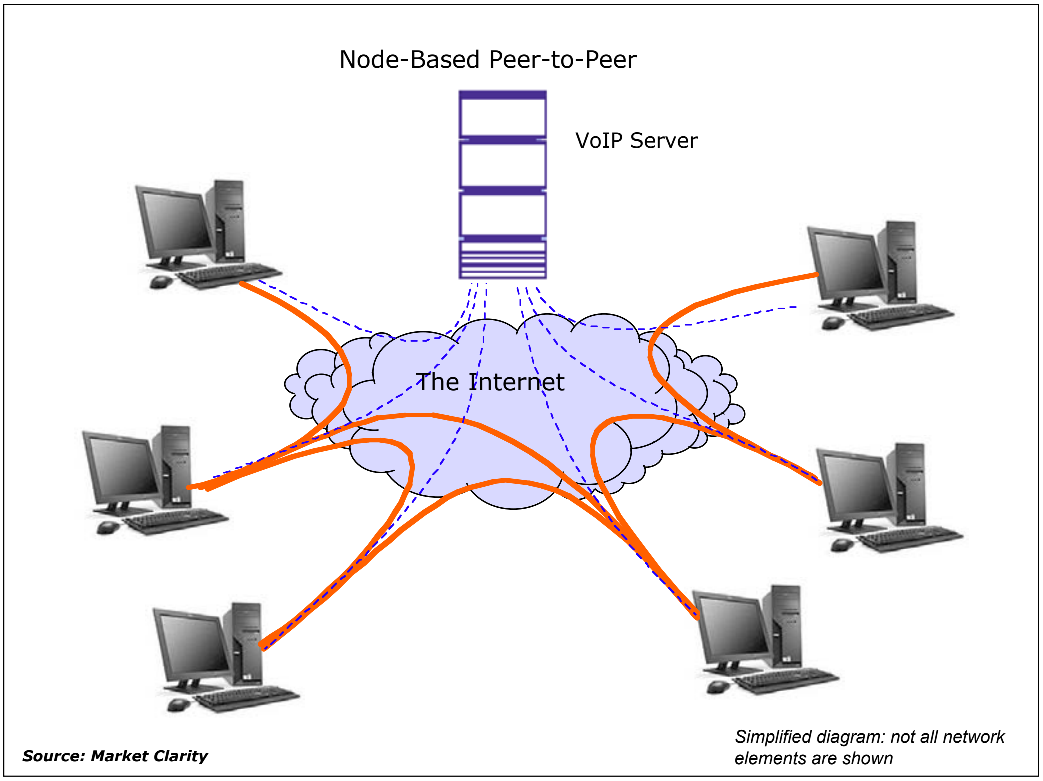

Figure 4
Node-based VoIP

In this model, clients combine the role of user agents and "nodes" which process traffic on behalf of other members of the network. While the Skype client can communicate past firewalls, to function as a node the client must have direct access to the Internet.

2.2 VoIP: Service Dependencies

It must be understood that the process of giving a VoIP client access to the network and the ability to make calls is far more complex than is typically illustrated in VoIP block diagrams.

Figure 5, below, illustrates in simplified form some of the components necessary to the successful operation of a VoIP client such as a broadband phone. Note that this illustration ignores some pragmatic aspects of VoIP deployment, such as SIP proxies or Session Border Controllers.

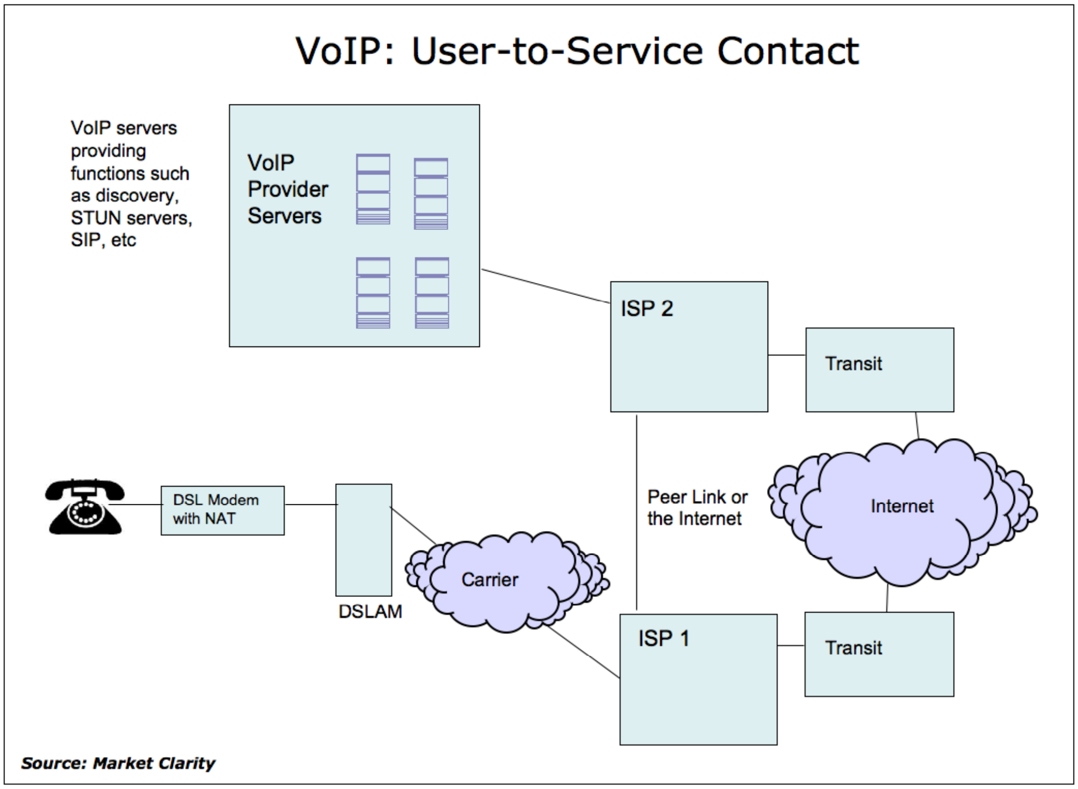

Figure 5
VoIP Service Enablement (Simplified)

The client at start-up has to:

(a) Obtain an IP address;

(b) Locate the VoIP server; and

(c) Register with the VoIP server.

When this three-step process is expanded with the various sessions required to make it happen, things become somewhat more complex, particularly when customer premises utilise firewalls. Presuming the customer is already logged into his/her ISP account, the conversations required, again in simplified form, are as follows:

(a) Obtain private IP address — The NAT firewall will assign an IP address (eg 192.168.1.10) to the broadband phone;

(b) Obtain the IP address of the STUN server — The client will initiate a DNS query to resolve the STUN server URL (for example, stun.voipprovider.com.au) with an IP address (299.299.1.52)[2];

---

2 A STUN Server (also just referred to as a server) is an entity that receives STUN requests, and sends STUN responses. STUN servers are generally attached to the public Internet. STUN sits along side a number of techniques to achieving NAT traversal; these include TURN (Traversal Using Relay NAT), ICE (Interactive Connectivity Establishment), UPnP (Universal Plug and Play) and Session Border controllers. Source: http://www.voip-info.org/wiki-STUN

(c) Register with the STUN server — the STUN server will associate the VoIP client with the public IP address of the DSL modem;

(d) Register with the VoIP server — The VoIP server will associate the user's public IP address (assigned by the user's ISP) with that user's VoIP account. To initiate the registration conversation, the client will issue a DNS query to resolve the service's URL (voiphost.voipprovider.com.au) with its IP address (299.299.1.1);

(e) The client will then send its login message to the VoIP server. The VoIP server can now associate the client's IP address (assigned by the client's ISP) with the client's VoIP user ID (sipuser@voiphost.voipserver.com.au), which will allow the client to send and receive VoIP packets (and therefore make and receive calls).

As with all Internet applications, VoIP is unable to function in the event that any of the components of the call (clients, servers or gateways) are unable to contact DNS servers.

2.3 VoIP QoS Overview

This Section presents a general overview of QoS and reliability considerations as they relate to VoIP.

2.3.1 **Network Performance**

In packet networks, QoS describes a network's ability to deliver traffic within parameters defined by the network manager as required for the acceptable performance of an application using that network. What constitutes acceptable performance depends on the latency/jitter sensitivity of different applications, as well as the application's (or user's) sensitivity to jitter and packet loss.

2.3.1.1 **Latency**

"Isochronicity" is the degree of real-time behaviour required by an application — in other words, whether the application remains functional and usable in the presence of end-to-end delays. Latency, the time taken for a packet to traverse between two points on a network, is used to measure that network's ability to support applications which require real-time or isochronous performance.

2.3.1.2 **Jitter**

Jitter describes variations in latency over time — a network able to deliver a packet within 30 milliseconds (ms) at one time may, due to congestion, suffer a delay of 250 ms for the next packet. This has minimal impact on traditional network applications such as e-mail or Web browsing.

However, for real-time applications such as a VoIP conversation, excessive jitter will disrupt users. It may also disrupt application behaviour: the ability of VoIP codecs to perform echo cancellation is partly dependent on the codec's measure of latency. Where latency varies, users may experience rising and falling levels of echo in addition to variable delay.

For echo cancellers to work well, latency must be constant (i.e. no jitter). However, jitter is an integral part of an IP network. Hence, it is essential that echo cancellers be located at a place that has no jitter source between the echo canceller and the telephone. For an IP phone, this means the echo canceller must be integral to the CPE, with responsibility only for removing echo in the local handset.

2.3.1.3 **Packet Loss**

Traditional Internet applications are designed to withstand packet loss. A traditional application such as e-mail can detect the loss of a packet and request a retransmission. Resilient applications are an important enabler of the Internet's ability to recover from congestion events: where traffic loads are excessive, network elements such as routers can discard packets and rely on the end applications to detect and correct the error.

Because a VoIP call is happening in near-real time, a dropped packet is noticeable, since the user will have noticed an interruption to the conversation before the network has had time to request a retransmission.

2.3.1.4 **Availability**

The availability of networks, and of different network elements, is also an important component of the QoS debate.

In Figure 5, above, which outlined the processes necessary to obtain access to a VoIP network, it can be seen that overall network availability depends on a number of service elements:

- **(a)** The customer premises: CPE used for Internet and VoIP connectivity, computers, IP phones, local area network (LAN), power supplies, and other peripherals.
- **(b)** The Internet service providing customer access;
- **(c)** The carrier network delivering customer traffic to the ISP;
- **(d)** The ISP's internal network and transit links;
- **(e)** The VoIP provider's network access (including the ISP service connecting its data centre); and
- **(f)** The VoIP provider's application servers.

Availability of a VoIP application will also depend on various "network utilities" without which Internet applications will fail. These include an ISP's access control (typically a Radius server) and DNS host.

While, the PSTN is well known for its design point of 99.999% availability, Internet access link availability is dependent on a range of different technologies, each with their own reliability characteristics. And, within the Internet cloud, individual ISPs may utilise their own design criteria, which may provide for very robust availability, or very poor availability[3].

[3] Availability is measured using the Mean Time Between Failure (MTBF) of all devices (and processes) in an end-to-end path.

VoIP services also introduce two further complexities: the customer premise environment, and the VoIP service provider's environment. At the customer premise, many discrete elements can impact service reliability. These include the CPE used for Internet connectivity, any VoIP specific CPE, computer equipment, local area networks — and the electrical power supply.

Within the VoIP provider environment, assuming that they are not also the Internet service provider, one needs to consider the following elements to determine service availability: the VoIP provider's network, the interconnectivity arrangements (if any) with other VoIP providers, VoIP application servers, and dedicated hardware such as session border controllers.

Where use of a service is sporadic, as with the telephone, user perception of downtime may be lower than that realised in practice — for example, scheduled downtime in an off-peak period will naturally impact fewer users than at times of high usage. When a service is used for multiple purposes, such as VoIP and Internet access, there is a greater likelihood that service outages will be noticed.

It may be worth examining how availability paradigms already developed for the PSTN may be applied to VoIP environments, in particular with reference to the categorisation of recovery mechanisms as outlined in Table 2, below.

Table 2
Voice Link Outage Classifications

Link Failure Duration	Classification	Technical Measures to Ameliorate the Impact in VoIP networks
Less than 150 ms	• Not considered unavailable, since users will not notice the break	• Link aggregation (IP/Ethernet and IP/MPLS-TE), backup LSPs, MPLS fast re-route
Less than 10 seconds	• Not considered unavailable if network recovers without losing call	• Rapid Spanning Tree Protocol (RSTP) for IP/Ethernet
Greater than 10 seconds	• Unavailable; will lose calls in progress • Successful redial indicates less severe problem • Inability to reconnect indicates service outage	• IP re-routing

2.3.2 The Impact of Network Behaviour on the User Experience

Table 3, below, outlines the sensitivity of different Internet applications to problems with network availability, latency, jitter, and packet loss.

Table 3
Application Sensitivity to Network Behaviour

Application	Availability Requirement	Latency Dependency	Jitter Dependency	Packet Loss Dependency
E-mail	• Moderate • Able to recover from interruptions to network connectivity	• Low • Timescale: Minutes • Application infrastructure resilient to very long delays	• None • E-mail not affected by jitter.	• None • Application protocols designed to recover from packet loss by requesting retransmission
Web browsing	• User-Dependent • Users will notice interruptions during activity • Some applications intolerant to interruptions	• Moderate • Timescale: Seconds • Users tolerant of moderate delays • Some applications may experience "time out" failures on excessive delays	• Low • Users unlikely to observe variable latency within wide limits	• Low • Application protocols designed to recover from packet loss • Users can reload affected Web pages
Domain Name System (DNS)	• High • All DNS-dependent applications fail if DNS server cannot be contacted	• Application-Dependent • Timescale: Seconds • While the DNS architecture is tolerant of delay, Real-time applications may experience "time out" failures on excessive DNS delays	• None	• Low • Applications designed to recover from packet loss by requesting retransmission

Application	**Availability Requirement**	**Latency Dependency**	**Jitter Dependency**	**Packet Loss Dependency**
Video Streaming	• User-Dependent • Users will notice interruptions during activity • Application-Dependent • Applications moderately tolerant to interruption	• Low • Timescale: seconds • Video streaming applications are designed to withstand interruptions • Users conditioned to download delays (for example during buffering)	• None	• Low • Applications designed to recover from packet loss by requesting retransmission
Instant Messaging (Text)	• User-Dependent • Users notice interruptions during activity • Application-Dependent • Some applications fail if unable to contact IM server	• Low • Timescale: seconds • Users cannot generally differentiate network behaviour from the behaviour of other chat participants	• None	• Low • Applications designed to recover from packet loss by requesting retransmission
VoIP	• User-Dependent • Users aware if network unavailable during activity • Application-Dependent • Application may fail if unable to contact VoIP server	• High • Timescale: 150 ms • Users sensitive to latency • Users and applications sensitive to jitter	• High • Users sensitive to variable delay • Applications such as codecs may be affected by excessive jitter	• High • Users will notice packet loss as loss of audio quality or interrupted conversation
Live Streaming (eg, IP-TV)	• User-Dependent • Users aware if network unavailable during activity	• High • Timescale: 1 second • Users sensitive to latency	• High • Users will experience jitter as loss of video quality	• High • Users will experience packet loss as loss of video quality

As the above discussions demonstrate, supporting VoIP with QoS mechanisms requires a consideration of a number of service characteristics.

It is worth observing that ITU-T recommends a jitter buffer in VoIP phones of 50 ms. However, many ISPs experience jitter higher than this threshold. As a result, some ISPs may have networks that are incompatible with ITU compliant CPE.

Should ACIF consider recommending jitter buffers greater than 50 ms for VoIP CPE so as to avoid future interoperability problems?

2.3.3 Other Factors Affecting the User Experience

While this paper is limited to an examination of network peering and inter-domain QoS, it should be kept in mind that the VoIP user experience is also subject to factors beyond the service provider community's control. Factors that may degrade the user experience even in the presence of the most favourable network conditions include:

(a) Low-quality or mis-configured client software;

(b) Low-quality or mis-configured IP phones or broadband phone adapters;

(c) Low-quality or mis-configured PC microphones, speakers, or sound cards;

(d) Unsuitable ambient environments (for example, where there is a high level of background noise or ambient echoes which cannot be corrected by software; or

(e) Excessive transcoding (for example, where a VoIP-originated call is transcoded at the PSTN gateway, and then transcoded again to terminate on a mobile phone).

3 INTERNET PEERING MODELS

This section is designed to equip readers unfamiliar with Internet peering concepts, with information so as to facilitate their understanding of the role of VoIP peering.

3.1 Defining Peering

In Internet terms, "peering" is an expression, which is narrow in definition but broad in application. It should be noted that not all peering models are applicable to VoIP services.

3.1.1 Communications without Servers

In one definition of "peer," any two Internet hosts that can exchange traffic with each other without requiring a controller can be considered peers.

Since there is no restriction on what kinds of hosts may initiate a communication session in which they exist in a peer relationship, the expression "peer" is applied to a wide variety of activities and applications.

3.1.2 Communications at a Single ISO Layer

A second usage of "peering" describes communications at the same ISO networking layer. It is in this sense that Skype can be considered as conducting "peer-to-peer" communications, since the application resident on one host communicates with the application resident on another host.

This second usage is also applied to the management of Internet traffic flows, which is discussed in this Section.

3.2 Exchanging Traffic as Peers

3.2.1 Internet Traffic Flows

Over time, the Internet developed into a layered approach to providing users with access to Internet hosts.

In the simplest model, a customer's request to a Web server may follow this path:

(a) Customer's HTTP request is sent to the customer's ISP (source ISP);

(b) If the requested Web server is not on the ISP's own network, the request is routed through the transit service used by that ISP for routing towards the destination;

(c) The request is routed to the ISP hosting the requested Web server (the destination ISP); and

(d) The destination ISP delivers the HTTP request to the requested Web server.

Transit services must be purchased by both the source ISP and the destination ISP. While a wide variety of tariff models are used in the sale of transit services, the most common model for transit purchases in Australia is for an ISP to acquire transit at facilities such as "Internet exchanges" with port speed used as the basis of its tariffs. Under this model, an ISP would purchase transit on a "per megabit per second, per month" basis.

ISPs therefore have to balance competing interests:

(a) The need to give customers access to requested network services and/or content at the best possible performance; and

(b) The ISP's need to minimise the cost of its transit services.

A number of techniques, such as caching, have been developed to help minimise an ISP's transit requirements, but these are outside the scope of this paper and are not relevant to VoIP services. Likewise, we do not consider a detailed discussion of transit tariff models in this document.

3.3 ISP Traffic Peering

"Peering" between ISPs describes the establishment of direct connection between ISPs' routers, so that those ISPs can directly exchange traffic.

For example, Fubar.com.au and foobar.net.au may find, upon analysing their traffic, that a large number of their users are sending e-mail to each other. The establishment of a peer connection between them would allow direct communication between their mail servers (Figure 6, below). Since the ISPs are reducing their cost of business, they each bear their share of the cost of establishing a peer link between their routers.

Establishing a peer connection in this way provides another benefit for ISPs and their customers: it reduces the number of router hops between hosts on participants' networks. Therefore if two ISPs are providing game hosting to their customers, players on the two networks will have a shorter path to the games[4] host than players on distant networks, with a consequent performance advantage.

In IETF terminology [4], this is referred to as Layer 3 peering, the "interconnection of two service providers for the purposes of exchanging IP packets which [are] destined for one (or both) of the peer's networks. Layer 3 peering is generally agnostic to the IP payload, and is frequently achieved using a routing protocol."

Figure 6, below, illustrates this model, referred to as bilateral peering.

[4] Terminology for Describing VoIP Peering and Interconnect, IETF, October 2005

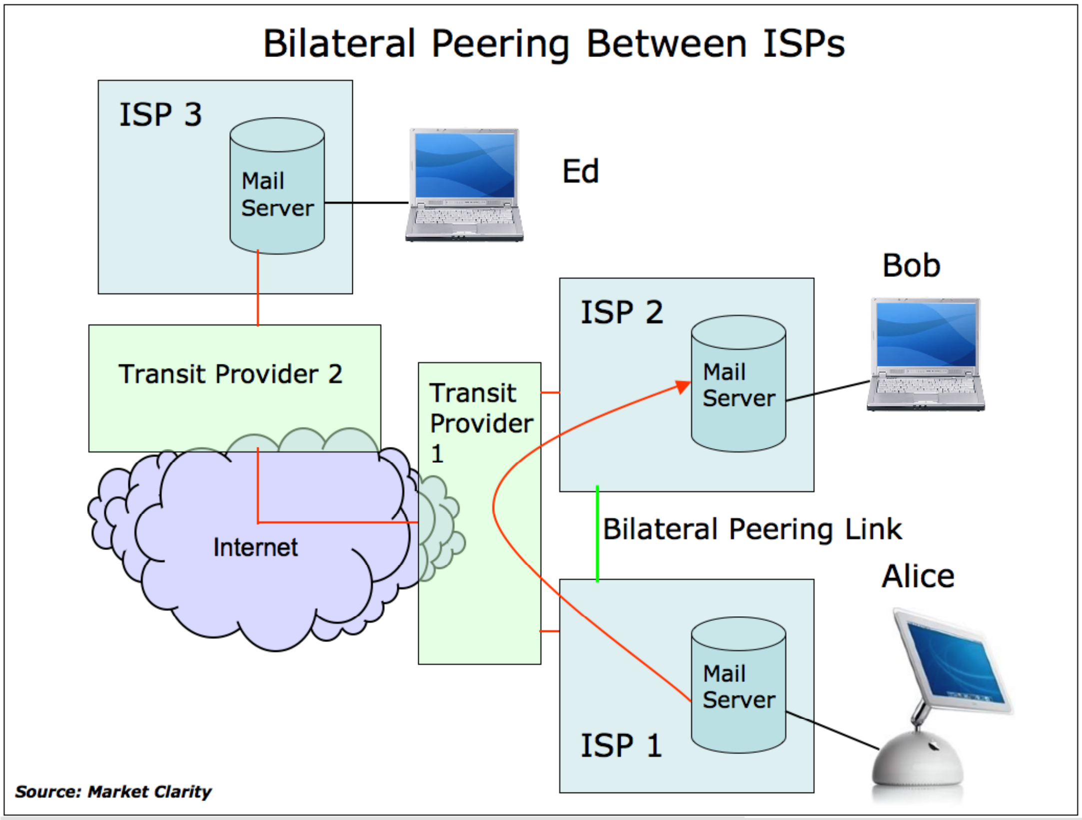

Figure 6
Bilateral Peering

If Alice sends an e-mail to Bob, ISP 1 and ISP 2, who have established a bilateral peering relationship, can route the message across that link and avoid having that message use their transit service. E-mails between Ed and either Alice or Bob have to use transit services to reach their destinations.

The existence of the bilateral peering link also reduces the route hops needed between Alice and Bob. Table 4, below, examines routes between hosts on the above network. Note that this Table does not take into account the large number of internal hops, which may arise for the following reasons:

(a) **Internal network segmentation —** ISPs and transit providers will create segments within their networks for reasons of security and traffic management;

(b) **Geographic diversity —** Additional router hops will be incurred where ISPs and transit providers need to transmit traffic between geographically dispersed locations.

Table 4
Peering and Route Length

Communication	Routes	Hops
IM between Alice and Bob (without peering)	• Customer router at ISP 1 • ISP 1 (internal hop) • ISP 1 to Transit Provider 1 • Transit Provider 1 (internal hop) • Transit Provider 1 to ISP 2 • ISP 2 (internal hop) • Customer router at ISP 2	• 7 Hops
IM between Alice and Bob (using bilateral peering link)	• Customer router at ISP 1 • ISP 1 (internal hop) • ISP 1 to ISP 2 • ISP 2 (internal hop) • Customer router at ISP 2	• 5 Hops
IM between Alice and Ed (without peering, across two transit provider networks)	• Customer router at ISP 1 • ISP 1 (internal hop) • ISP 1 to Transit Provider 1 • Transit Provider 1 (internal hop) • Transit Provider 1 to Transit Provider 2 • Transit Provider 2 (internal hop) • Transit Provider 2 to ISP 3 • ISP 3 (internal hop) • Customer router at ISP 3	• 9 Hops

Clearly, then, for applications sensitive to latency the creation of new adjacencies via peering arrangements benefits customers as well as ISPs.

3.3.1 **Multilateral Peering**

ISPs seeking to establish a peer relationship need a facility by which their router ports can be connected. This may take the form of a point-to-point data service, but this service can only interconnect two ISPs.

As the base of Internet customers expanded, ISPs identified the need for a model which allowed large numbers of peers to exchange traffic. Since Australian ISPs outside what is colloquially referred to as the "gang of four" (Telstra, Optus, AAPT and Verizon) were excluded from peering with those four providers, they could instead reduce their need for transit services by forming their own multilateral peering arrangements.

In a multilateral peering exchange, member ISPs share the costs of a data centre sufficient to house their switches and accept their incoming data links. Within the peering exchange, the ISPs establish routes between each other, so that traffic can pass between them on a peer basis without occupying transit links. Member ISPs are responsible for securing data services into the peering exchange.

This peering model is illustrated in Figure 7, below.

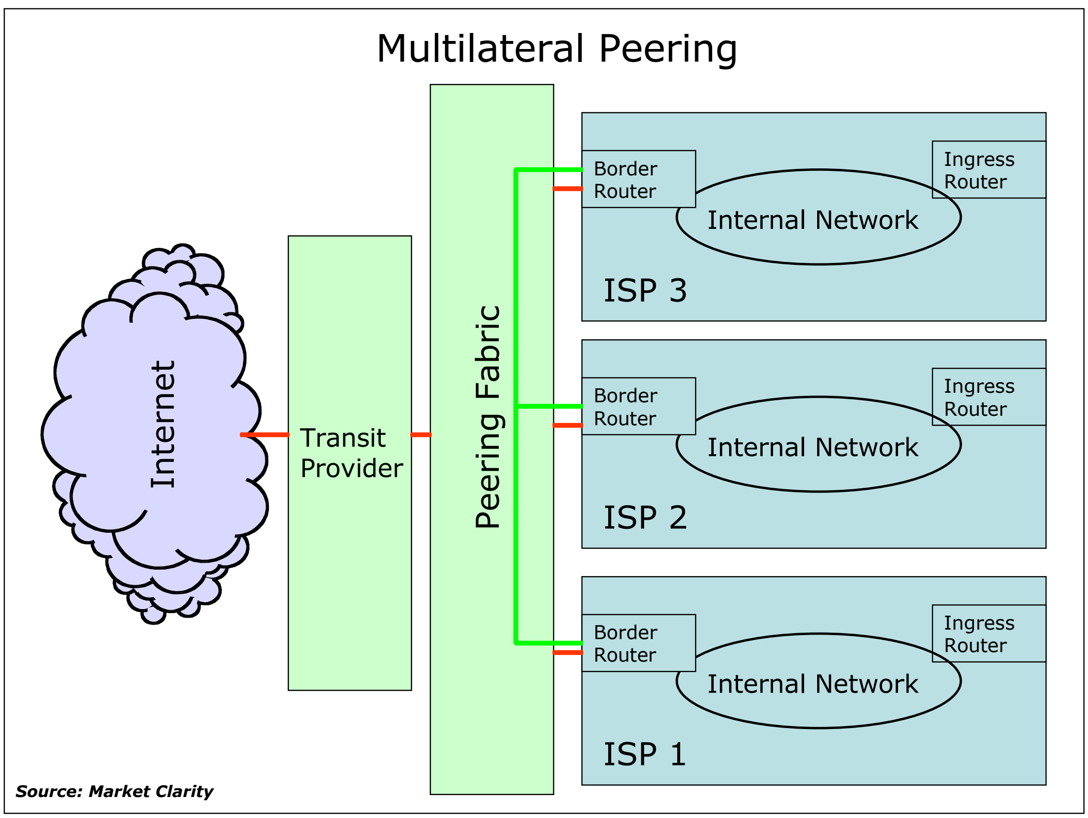

Figure 7
Multilateral Peering — Simplified Diagram

A sample list of peering activity in Australia is given in Table 5 below.

Table 5
Australian Peers

Organisation	URL	Comments
AdNAP	www.adnap.net.au	AdNAP is the Adelaide Network Access Point.
Ausix	www.ausix.net	The Ausix peering exchange operates in the GlobalCenter data centre in Melbourne.
Equinix	www.equinix.com	Commercial data centre operator supporting peering connections between customers
Pipe Networks	www.pipenetworks.com.au	Commercial carrier offering peering exchange services
SAIX	www.saia.asn.au	The South Australian Internet Association administers SAIX.
STIX	www.ix.singtel.com	STIX is a peering point operated by SingTel in Singapore.
VIX	www.vix.asn.au	ViX is the Victorian Internet Exchange.
WAIX	www.waia.asn.au/waix	The WA Internet Association operates WAIX.

This list should not be regarded as comprehensive. Australia's Tier-1 ISPs (AAPT, Optus, Telstra, and Verizon) peer with each other, and in addition, there are likely to exist bilateral and multilateral peering relationships not known to the authors.

4 EMERGING PEERING MODELS FOR VOIP

To understand the complexities of VoIP peering, it is necessary to examine the differences between data plane peering and control plane peering as they relate to VoIP.

4.1 Interconnection Without Peering

As VoIP grows in popularity, providers face a growing need to interconnect their users, because an inability to call users on other networks is a disincentive to VoIP adoption.

Where VoIP providers have established an in-dial capability, this interconnection can be achieved by having the call traverse the PSTN, as illustrated in Figure 8, below, in which the two discrete VoIP services are identified as "Green VoIP Provider" and "Blue VoIP Provider".

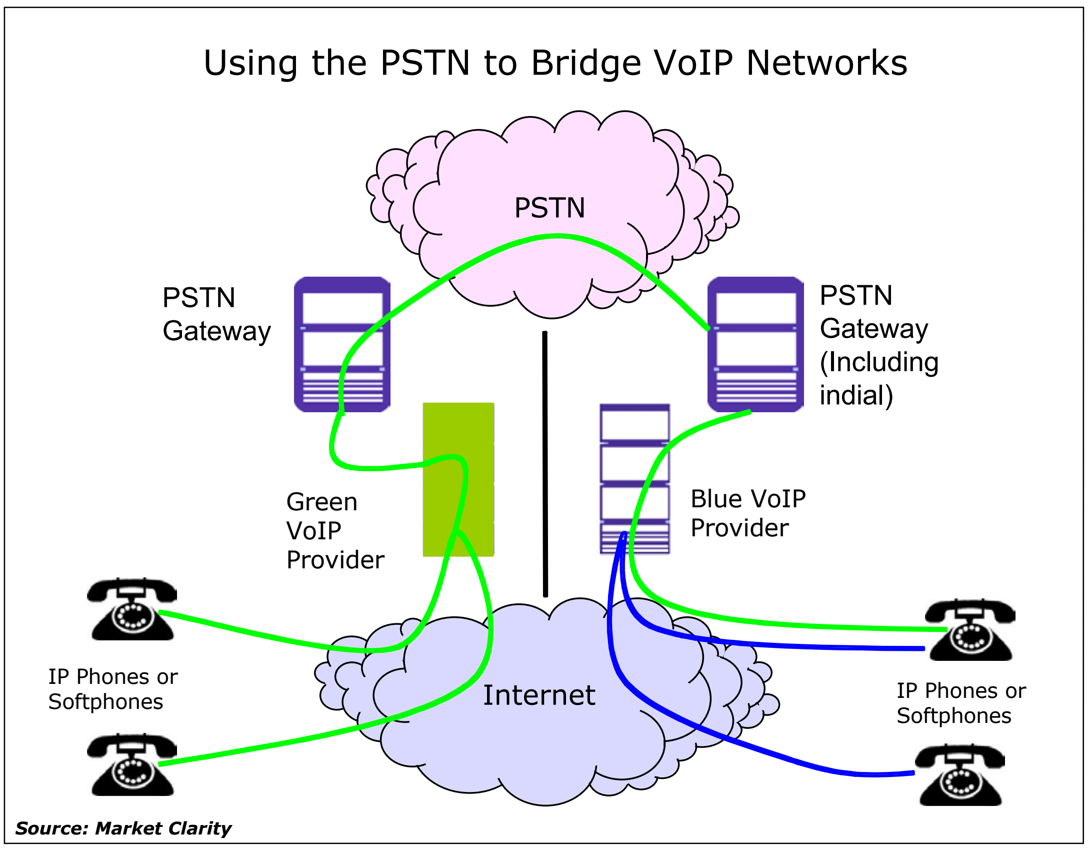

Figure 8
PSTN Delivery of VoIP-to-VoIP Calls

This approach is, however, not optimal for the following reasons:

(a) The call incurs PSTN termination charges, which erodes VoIP's competitive advantage; and

(b) The call may be transcoded two or more times in the VoIP-PSTN gateways (if different codecs are used), which erodes call quality.

4.2 Signal Plane Peering

For calls to traverse different VoIP services, providers need a means by which users are able to discover each other. When a user on Green VoIP Provider calls a user on Blue VoIP Provider, the following steps are necessary:

(a) Green VoIP Provider must recognise that while the called party is not present on its network, they are available at Blue VoIP Provider;

(b) Green VoIP Provider must request the location and availability of the user agent of the called party; and

(c) Blue VoIP Provider provides this location data and the call completes.

One means to do this is on a bilateral basis, in which providers establish a path between their servers allowing each other to seek IP address resolution for each other's networks.

This is illustrated on Figure 9, below.

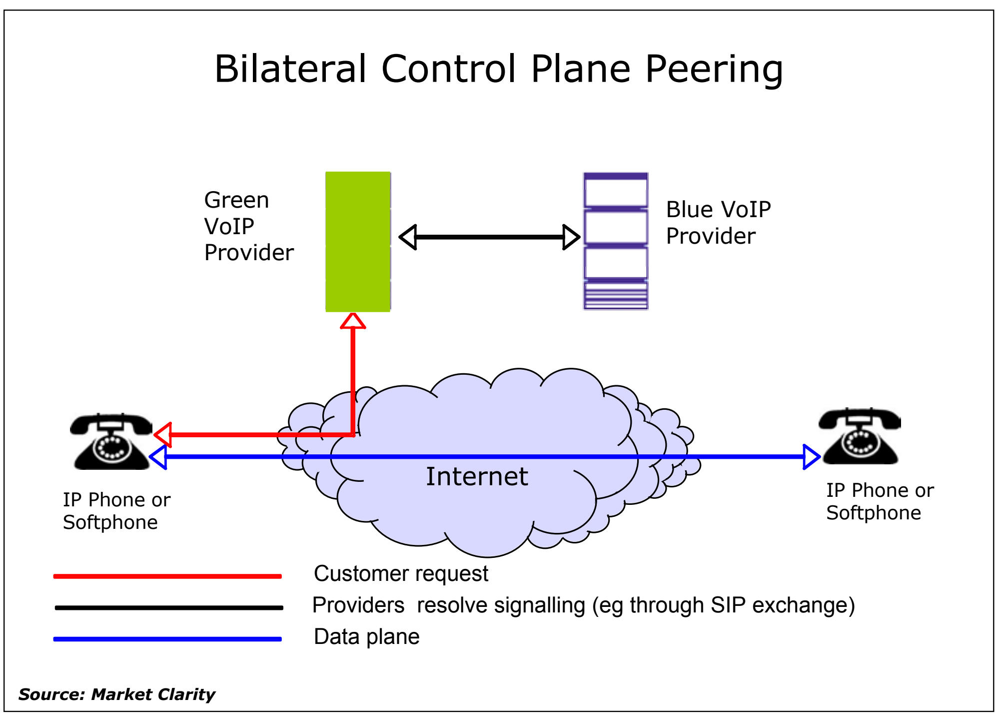

Figure 9
Bilateral Peering in VoIP

The IETF describes such VoIP peering as "Layer 5 peering[5]". The IETF definition is given as follows:

> "Layer 5 peering refers to interconnection of two service providers for the purposes of SIP signalling. Media streams associated with this signalling (if any) are not constrained to follow the same set of paths."

Note that in the Layer 5 peering case, there is no intervening network. That is, for purposes of this discussion, there is no such thing as a "Layer 5 Transit Network". According to the draft IETF definition, VoIP peering is defined to be a Layer 5 peering between two VoIP providers for purposes of routing real-time (or quasi-real time) call signalling between their respective customers.

This gives rise to issues, which may prevent this model from scaling upwards to operate on an industry-wide basis, specifically:

- **Trust** — While providers may trust each other on a case-by-case basis, for the VoIP industry to achieve any-to-any connection, this trust would have to expanded to an unrealistic level.
- **Administration** — An expanding network of one-to-one peers burdens providers with an expanding administration load.
- **Security** — Providers must manage responses to off-network requests such that they do not provide user information to untrusted parties (for example, to attackers spoofing the network identity of a VoIP provider). Also, this model allows end-users to see each other's public IP addresses (those assigned by the ISP's DNS), which reduces user security.

4.3 Data Plane Peering

VoIP media (the encapsulated voice packets) can be exchanged on a Layer 3 basis (that is, the peering exists only at the data plane to exchange IP packets terminating on the peer networks).

This is illustrated in Figure 10, below. Green VoIP Provider and Blue VoIP Provider have agreed to establish a connection between their services.

[5] Terminology for Describing VoIP Peering and Interconnect, IETF Draft, October 2005.

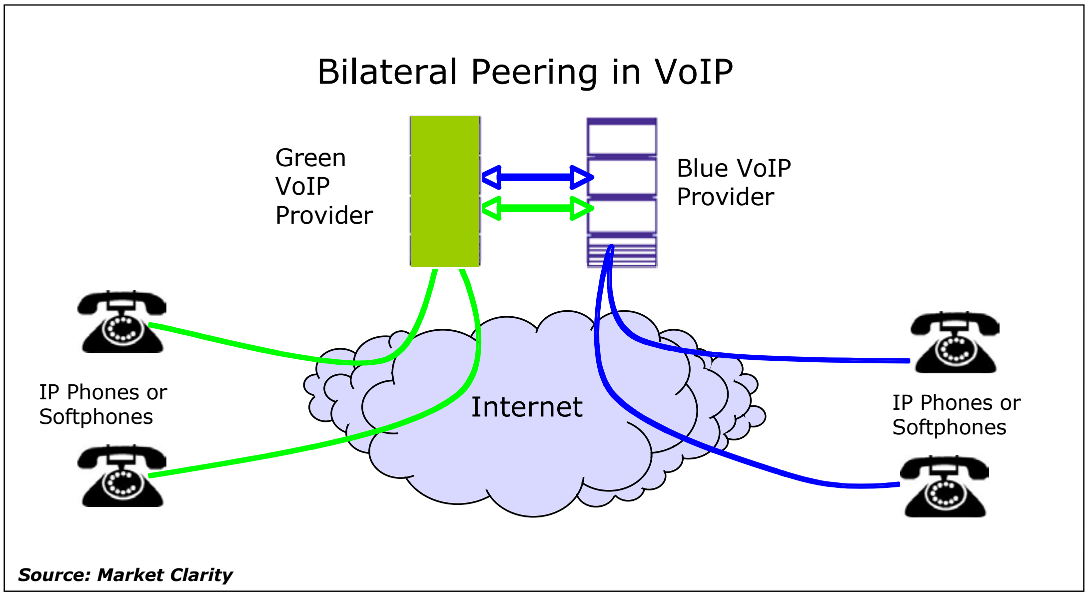

Figure 10
Bilateral Peering in VoIP

The exchange of traffic between VoIP providers partly replicates the bilateral peering model illustrated in Figure 9, above. However, instead of building a fabric that exchanges all IP traffic on a route basis, the VoIP providers are routing call traffic between subscribers.

Note that Figure 9 shows a logical connection rather than a physical connection. Data plane peering between VoIP services may involve traffic routed across the Internet and still facilitate users placing VoIP calls between networks.

As with bilateral IP peering, bilateral peering between VoIP providers does not scale well, requiring as it does a large number of individual contracts between providers. This would be overcome, of course, by employing the "peering exchange" model in the VoIP world.

Absent from this simplified model of data plane peering is the resolution of VoIP user addresses. This requires signalling plane transactions, which are discussed below.

4.4 Multilateral Signal Plane Peering

Another model for VoIP peering, employed by Pulver.com spin-off IPeerX (www.ipeerx.com), is outlined in simplified form in Figure 11 below.

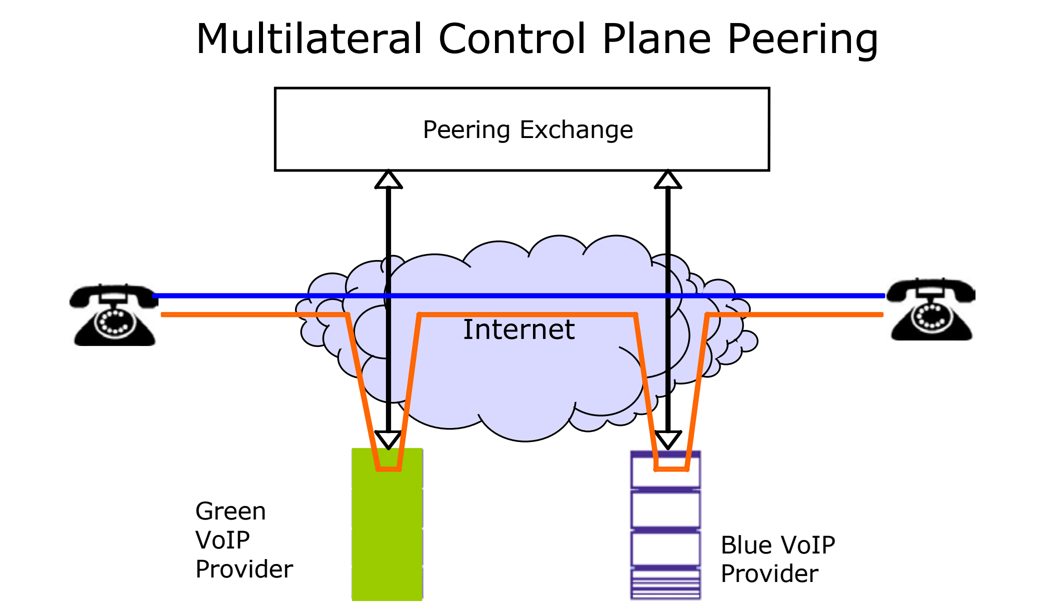

Source: Market Clarity

Figure 11
Multilateral Control Plane VoIP Peering

In this model, peering is not applied to the traffic between user agents (the "media path") but rather to the "signalling plane".

When a user on Green VoIP Provider attempts to contact a user on Blue VoIP Provider, the originating provider will fail to find the called number in its directory. Since it does not have access to Blue VoIP Provider's directory, Green VoIP Provider instead passes the request to the signalling exchange for resolution.

Multilateral peering may be accomplished in several ways, some of which are outlined below (a comprehensive list is beyond the scope of this paper):

(a) **Free Public Peering** — Individuals or organisations may host, on a co-operative basis, SIP or Asterisk servers through which VoIP networks supporting the appropriate protocols may resolve each other's user requests.

(b) **ENUM [IETF RFC 3761]** — Although in the early stages of development (for example, in Australia ENUM has not progressed beyond trials), ENUM allows cross-network address resolution on a peer basis.

(c) **Co-operative Multilateral Control Plane Peering** — As with ISPs, groups of VoIP providers may agree to jointly fund a system able to support control plane peering on a co-operative basis.

(d) **Commercial Multilateral Control Plane Peering** — The costs of hosting the peering server are recovered on a membership or transaction basis.

The purpose of the peering fabric is to handle user discovery and addressing requests from providers. The general model is as follows:

(a) Receiving a request for a user not known to its directory, Green VoIP provider sends a request across the Internet to the peering exchange.

(b) The peering exchange searches its database for Blue VoIP Provider.

(c) If Blue VoIP Provider is a member of the exchange, the IP address of its VoIP host is returned to Green VoIP Provider.

(d) Green VoIP Provider then requests a session with Blue VoIP Provider.

A further development of this model is that the peering exchange may also be able to provide route information allowing the VoIP session to be carried by a specialist VoIP carrier rather than the public Internet for part of its journey.

Peering may involve user agents being put directly in touch with each other; or it may involve the two providers acting as proxies for user agents. The user agents are obscured from each other for various reasons, including user privacy (and the concealment of end user IP addresses), the privacy of the providers' customer databases, and security of the providers' networks.

4.4.1 Session Border Controllers

Signalling and routing are not the only requirements for VoIP providers to exchange traffic. There already exists an array of different signalling protocols (SIP, H.323, Megaco and others) and different codecs.

This is increasingly resolved through the use of Session Border Controllers (SBCs). These perform proxy functions on behalf of clients, but add signalling protocol translation and transcoding (codec translation) to their functions.

SBCs also help protect user privacy by concealing end user IP addresses from other users.

4.5 Peering and Proprietary VoIP Services

While consumer VoIP services leverage public Internet infrastructure as the data transport, VoIP clients and servers are applications. As with all applications, developers have considerable scope to implement proprietary protocols and features governing interactions between clients and server, or interactions between clients.

The service's ability to communicate over the Internet is not impaired, since this depends on the developer correctly implementing the IP stack through which upper-layer functions access the network.

Proprietary application functions may also include proprietary signalling and/or codecs. An example is the popular Skype softphone.

Skype's proprietary signalling is of particular interest in this discussion, since as well as establishing connections to the service and placing PSTN-terminated calls using SkypeOut, its signalling protocol must be considered integral to Skype's distributed peer-to-peer processing model. The signalling is used to discover information about peers, which are available to act as "supernodes" in the Skype network, performing processing on behalf of other Skype nodes.

For the individual provider, however, proprietary signalling protocols may have undesirable side-effects:

- **Translation** — To exchange end-to-end Internet-based calls with other VoIP provider networks, resources must be devoted to developing gateway software to act as an interface between proprietary signalling protocols and open industry standards.
- **Feature Transparency** — End-to-end feature transparency becomes more difficult and more complex where those features are invoked by non-standard protocols.

As will be discussed in Section 5.4, below, Skype's architecture also has implications for peer-based exchange of QoS-categorised traffic.

4.6 VoIP Peering: Issues for Industry Consideration

In examining whether an industry-wide approach to VoIP peering is desirable, the industry may wish to include the consideration of the following questions:

(a) To what extent is industry-wide co-operation desirable in implementing inter-domain VoIP peering? Is co-operation achievable, or will it create a new layer of peering, which exists alongside of, or in competition with, other peering initiatives?

(b) What provider vulnerabilities may exist in the implementation of inter-domain VoIP peering, and what steps can providers take to manage any vulnerability that may arise?

(c) Is the adoption or endorsement of particular Internet standards to facilitate inter-domain VoIP peering desirable or feasible? For example, is there any benefit in promoting particular signalling standards (such as SIP), codecs, or in standardising approaches to common issues such as firewall traversal?

(d) What role may organisations such as ACIF have in promoting a co-operative approach to peering in the Australian VoIP industry?

(e) Would a national approach to inter-domain VoIP peering have any impact, desirable or otherwise, on Australian VoIP providers' ability to implement peering arrangements with international VoIP providers?

(f) To what extent would multi-lateral VoIP peering agreements place customer information at risk of exposure, and what industry agreements or standards are necessary to avoid any risk that may arise?

(g) If Australian VoIP providers choose to adopt an industry-wide approach to inter-domain VoIP peering, in what ways might this interact with the Quality of Service issues (discussed in Section 5 of this discussion paper)?

(h) In what way would any proposed new initiatives align or interact with existing activities such as the ENUM trial currently being administered by ACMA? What liaison should be conducted with existing activities?

5 QUALITY OF SERVICE

5.1 Introduction

Quality of Service (QoS) refers to the capability of a network to provide better service to selected network traffic, and is often referred to as traffic engineering. QoS mechanisms can be deployed across various technologies, and at various network layers.

In order to provide preferential service to a traffic type (such as VoIP), it must first be *identified*. Subsequent to this identification, the packet may or may not be *marked*. Where a packet is marked it can be sent along a pre-defined path, otherwise, identification must take place at each hop in the end-to-end path. Together these two tasks are referred to as *traffic classification*.

There are numerous **QoS mechanisms**, which can be employed at the IP Layer:

- **Admission Control**: Determines whether the flow can/should be allowed to enter the network;
- **Packet Classification**: Classifies the data based on admission control for desired treatment through the network, and may also include packet marking;
- **Bandwidth Management**: Determines if enough bandwidth is available, provisions devices for reserved bandwidth (if applicable), and enforces the traffic contract;
- **Queue Management**: Determines the behaviour of data within a queue; and
- **Queue Scheduling**: Determines the order and timing for different queues to send information onto the outbound link.

In combination with the above, the following mechanisms can be used to **manage bandwidth**:

- **Traffic Policing**: Measures the traffic to determine if it is out of profile, where packets that are determined to be out-of-profile can be dropped or marked differently;
- **Traffic Shaping**: Uses buffering to control traffic flows, therefore delaying some of the data, to ensure the traffic fits into the profile; and
- **Bandwidth Allocation**: Mechanisms such as RSVP, can be used to reserve bandwidth on a per-network, per-link or per-path basis.

It should be noted that while traffic shaping is appropriate for traffic such as Web browsing or non-real-time media downloads, the use of buffers degrades performance for real-time traffic such as VoIP. Traffic policing is more appropriate for these traffic types.

Arguably, bandwidth over provisioning may also be included as a bandwidth management mechanism.

Once a packet has been classified, appropriate congestion management, queue management, link efficiency, and shaping/policing policies can be applied.

Traffic management can be accomplished at different network layers:

- **Layer 1** — Over-provisioning of bandwidth;
- **Layer 2** — Through the use of ATM service classes, Ethernet 802.1p prioritisation of frames, 802.1Q VLANs for the logical separation of traffic, or Metro Ethernet Forum bandwidth profile mechanisms; and
- **Layer 3** — MPLS, DiffServ and IP Precedence: DiffServ provides traffic prioritisation as well as traffic engineering, when used in conjunction with MPLS, whereas IP Precedence can be used to provide traffic prioritisation.

The primary goal of QoS is to provide priority to certain traffic flows including dedicated (or reserved) bandwidth, controlled jitter and latency, and improved packet loss characteristics. Another important characteristic of traffic engineering is ensuring that by providing priority for one or more traffic flows, that other non-prioritised flows are also serviced.

QoS can be described as the concept of applying and ensuring specific, quantifiable performance levels across a network, or networks.

Figure 12 below, provides a simplified view of key QoS techniques at various network layers.

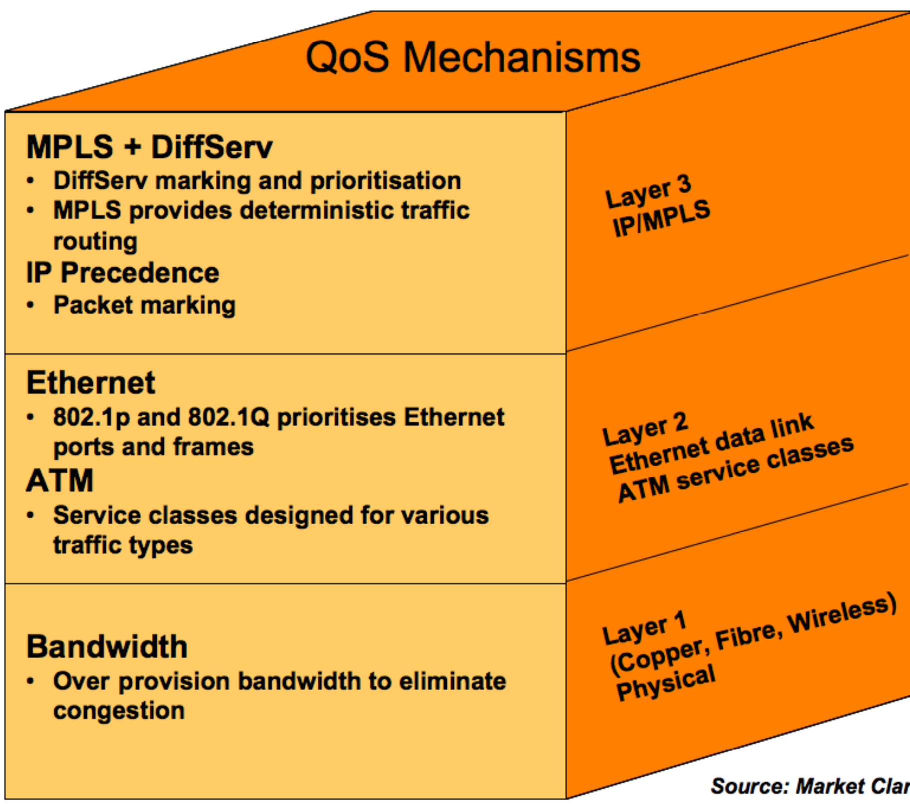

Source: Market Clarity

Figure 12
QoS Mechanisms

Within the VoIP interconnectivity framework, there are numerous Layer 3 traffic-engineering methodologies that must be considered:

- **IP Precedence Marking** — This is a Layer 3 (IP Layer) prioritisation scheme, which is defined in IETF RFC 1349. It uses three (3) bits from the Type of Service (TOS) field in an IPv4 header to indicate traffic priorities. A total of eight (8) priorities can be established. IP Precedence deals with traffic prioritisation, rather than bandwidth reservation or bounds on latency or jitter. As such, it does not meet all the goals of a traffic-engineered network.

- **Multi-Protocol Label Switching (MPLS)** — The use of MPLS labels allow routers to make forwarding decisions based on the contents of a simple label, rather than by performing complex route lookups based on destination IP address. Label Switched Paths (LSPs) are determined by an MPLS edge router, which analyses the contents of the IP header and selects an appropriate label with which to encapsulate the packet. This high level of control results in a network that provides a deterministic path between two points. It does not, however, provide a quality of service architecture.

  MPLS supports the concept of label stacking, which allows MPLS network operators to establish a network hierarchy, and to tunnel (encapsulate) customer or partner service provider information transparently across a network domain. Label stacks can be used to engineer arbitrarily complex processing of various traffic flows, such as VoIP, across a network domain, or indeed between multiple service provider domains (Inter-Domain).

  MPLS can force packets into specific paths. However, MPLS by itself does not provide for QoS.

- **Differentiated Services (DiffServ)** — DiffServ is a mechanism that provides a relative ordering of “aggregate” behaviour flows by making use of the IP header's TOS byte, renamed the DS byte. DiffServ includes traffic prioritisation as well as traffic engineering characteristics. In the DiffServ environment, standardised flows are mapped into service attribute categories called Per Hop Behaviours (PHBs) that can be engineered to support real-time traffic flows such as VoIP. PHB characteristics can be specified in quantitative terms such as throughput, delay, jitter or loss; or in terms of some relative priority of access to network resources. DiffServ is intended to supersede the existing definitions of the IPv4 TOS byte and the IPv6 Traffic Class byte.

  DiffServ provides scalable service discrimination without the need for per-flow state and signalling at every hop in the network. Policing is done at the edge of the network, and services can be defined on an end-to-end (Inter-Domain) basis or purely Intra-Domain (within a single network).

  DiffServ provides a QoS treatment to traffic aggregates, providing a scalable and operationally simple solution. However, because it does not influence a packet path, it cannot guarantee QoS.

- **MPLS and DiffServ** — By combining MPLS and DiffServ, network operators are provided with a scalable traffic classification schema and that provides traffic engineered (TE) path selection and bandwidth guarantees, enabling true QoS.

Traffic engineering allows a network operator to establish a deterministic path, and bypass the normal routed hop-by-hop paths in an IP network. In other words, it is used to steer traffic to parts of the network where capacity is available.

5.2 Inter-Domain QoS

The use of IP/MPLS Traffic Engineering (TE), allows for the creation of unidirectional paths that are independent of the Interior Gateway Protocol (IGP) shortest path computations. There are numerous benefits in using this mechanism for QoS support[6]:

- Eliminates Layer 2 overlays;
- Provides deterministic performance;
- Supports QoS guarantees;
- Intelligent use of network resources (capacity);
- The ability to transport non-IP traffic;
- Advanced protection schemes such as end-to-end path protection, diverse routing, and fast re-route; and
- Is broadly deployed across many vendor platforms.

As recently as February 2006, the MFA Forum (MPLS and Frame Relay Alliance Forum) and 15 of its vendor members, featuring more than 30 devices, demonstrated the interoperability, resiliency, multicast and Quality of Service (QOS) capabilities of various MPLS network services at the MPLS World Congress.

Goals of the test event included interoperability testing on end-to-end traffic engineering, protection and Layer 2 and Layer 3 multicast. Based on the need to support several services over a single converged backbone, the MFA revisited network protection interoperability and verified QoS enabled routing, a mandatory precondition for upcoming IETF RFCs on integrated MPLS quality and Differentiated Services (DiffServ-TE).

The MFA tests and other industry initiatives, are supported by improved tools and systems that are now available for the analysis of TE-related data including metrics, routing, traffic and path performance.

In the context of QoS-based VoIP interconnectivity, MPLS traffic engineered (TE) tunnels can be set up to carry traffic that has specific bandwidth requirements and specific attributes. For instance, two service providers could establish a TE tunnel to carry VoIP traffic between two VoIP media gateways or session border controllers, connected to two different IP/MPLS networks.

Alternatively, inter-domain QoS can be accomplished at Layer 2, either through the use of ATM service classes, or Ethernet 802.1p prioritisation of frames, 802.1Q VLANs for the logical separation of traffic, and Metro Ethernet Forum bandwidth profile mechanisms.

---

[6] Dr. Vishal Sharma, Metanoia Inc, Presentation to ACIF, "Inter-Domain TE and QoS: Some Key Aspects and Challenges" (February 2006)

5.2.1 Ethernet QoS

It should be remembered that ISPs may not desire the level of QoS granularity offered by MPLS. Their requirement may be for a simpler mechanism, where the only distinction is between best-effort traffic and a known level of QoS for VoIP conversations.

Ethernet-based QoS mechanisms can meet this requirement. ISPs could mark the p-bits in an Ethernet frame (IEEE 802.1p), or establish separate VLANs (IEEE 802.1Q and IEEE 802.1ad), for VoIP (and other real-time traffic). At an Ethernet point-of-interconnect (POI), member ISPs would provide prioritised routing based on p-bit markings or VLAN membership.

This approach does not exclude the use of MPLS, since the Ethernet port carrying QoS-marked traffic can be associated with an MPLS network if required.

5.2.2 Key Challenges for Inter-Domain QoS

Even with the availability of increasingly sophisticated TE tools, there remain a number of challenges that service providers face in implementing cross-domain traffic engineering. These include [6]:

- The lack of common industry definitions for various service classes;
- SLA monitoring and reporting mechanisms;
- Universal agreement on metrics, sampling rates and reporting characteristics;
- Scalable performance measures;
- Agreement on how to divide up the delay impairment budget across multiple network domains;
- Defining (and architecting) resilient end-to-end services;
- Ensuring security and confidentiality;
- Developing TE or QoS-based peering contracts for admission control; and
- Establishing charging or settlement mechanisms across providers.

Traffic engineering complexity increases with the number of provider domains, and where inter-domain TE is defined on a signalled basis, is likely to prove too complex for a first step at interconnectivity.

Rather than trying to solve all of these issues right now, current industry thinking is leaning towards manual first steps:

- Utilising specific markings for traffic with the same requirements (e.g. VoIP);
- Providing a defined network behaviour based on that marking; and
- Allowing all traffic with the same marking to compete for resources on the same basis.

While the concept of traffic engineering is complex in its own right, the interconnectivity of VoIP networks introduces an additional layer of complexity with the need to mediate between different VoIP protocols such as SIP, H.323. MCGP and IAX while accounting for transcoding issues caused by the use of different codecs.

It is also important that networks be designed to avoid unnecessary transcoding. Rather than performing transcoding as a matter of course, peer networks should favour end-points utilising a mutually acceptable choice of codecs wherever possible.

It is also possible that different "grades" of VoIP service may emerge, which may require different treatment from each other.

5.3 VoIP QoS

In Section 2.3, we provided an overview of VoIP QoS. In this section we look at VoIP QoS issues in additional detail.

The preservation of QoS within a domain is a manageable problem. VoIP, however, is an application subject to the effect of QoS in many domains. The domains included in a VoIP call that traverses the Internet and is passed between two different VoIP providers may include multiple Internet customer access networks, ISPs, transit providers, peer exchanges and PSTN gateways. The complexity of managing QoS on an end-to-end basis is illustrated in Figure 13, below.

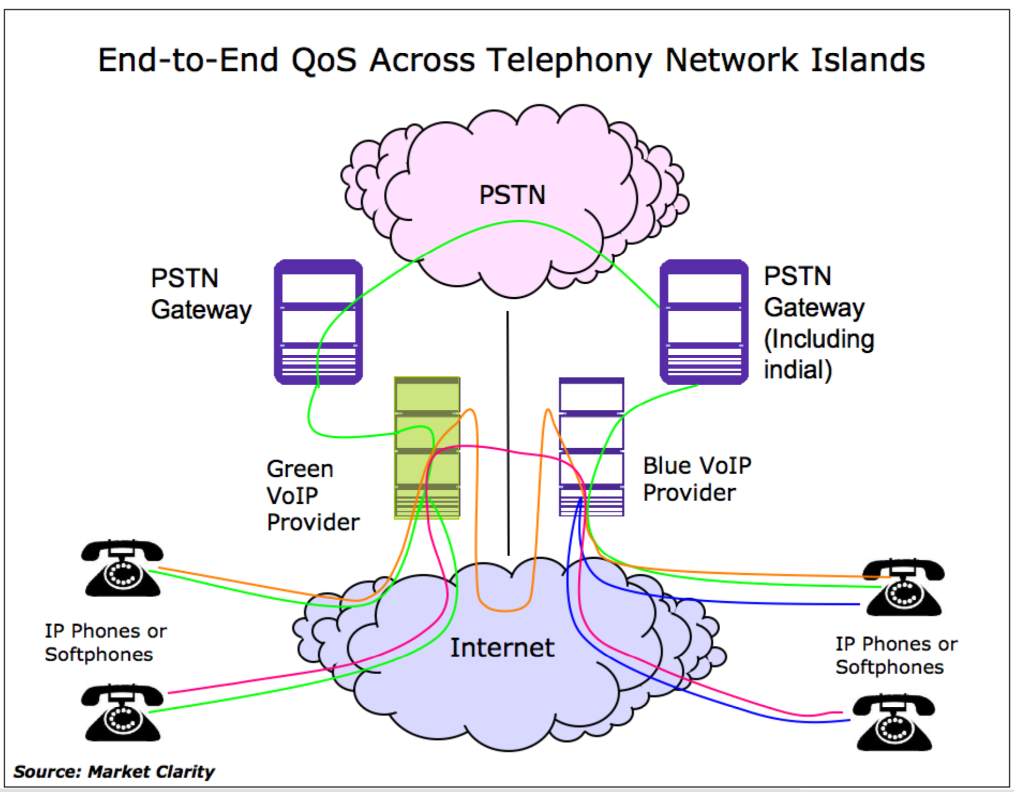

Figure 13
End-to-End QoS Across Telephony Network Islands

In Figure 13, we show a number of signalling and call path possibilities, to illustrate some of the complexities involved in VoIP QoS. For discussion purposes, we will assume that the Internet access links at each customer site provide suitable QoS characteristics.

- The blue line shows the path between two VoIP users within a single VoIP provider's network (Intra-Domain signalling and media). *With suitable access links, QoS can be engineered on an end-to-end basis.*
- The orange line shows the path between two VoIP users connected to different VoIP provider networks (Inter-Domain signalling and media). Call signalling between the two VoIP networks has been enabled, perhaps through a third party ENUM or SIP Proxy lookup service. Once the destination IP address is retrieved, the VoIP call media (the voice packets) are routed across the public Internet on a best effort basis. *Even with suitable access links, QoS cannot be engineered on an end-to-end basis.*
- The pink line shows the path between two VoIP users connected to different VoIP provider networks (Inter-Domain signalling and media) where the providers have an interconnect agreement (or peering relationship). Here, VoIP calls are routed between VoIP islands at a designated border, across the traffic-engineered networks of each VoIP provider. *With suitable access links, QoS can be engineered on an end-to-end basis.*
- The green line shows a more complex scenario: a three (3) party conference call. Two of the callers are subscribers to the Green Provider's VoIP service, and the third party is a Blue Provider subscriber. Because the two VoIP providers have not interconnected for either call signalling or call media, voice packets transiting between the Blue and Green VoIP Providers are translated (transcoded) between VoIP and TDM formats. *With suitable access links, suitable VoIP network design, and appropriate VoIP/PSTN gateways, end-to-end QoS can be achieved. However, each time voice is transcoded between formats there is a degradation of quality, as determined by the human ear. The codecs utilised by each end-user also impact the quality experience.*

As illustrated in these examples, a controlled end-to-end environment provides for the best human experience. This quality of service (as determined by the human listener) is simpler to guarantee when all elements in a call path are known either through a single provider's end-to-end management, or managed by a commercial agreement that specifies "QoS" (in some manner) within individual service provider VoIP and Internet domains. The specification of "QoS" may be dependent upon traffic markings, use of different ports for best effort and real-time traffic, a commercial understanding that a partner's network is of equivalent quality, or any other mechanism. In other words, it does not have to involve complex traffic engineering.

It does, however, have to provide a consistent service experience for the end user.

The goal of VoIP traffic engineering is to provide a high quality service, as perceived by service users. Ideally, the end user should have a PSTN-like experience every time they place or receive a VoIP call.

To provide a high quality service experience, QoS mechanisms recognise that different traffic types require different treatment within the network cloud (which may span multiple network islands).

It may also be desirable to allow for different "grades" of VoIP services, each with different network treatment.

And, in the real world, the QoS equation must also take into account the access link. Indeed, it is the access link that will prove most problematic in terms of guaranteeing service quality.

5.3.1 The Weakest Link

Perhaps the greatest challenge in creating an infrastructure able to preserve VoIP QoS on an end-to-end basis is the customer access network.

Since the inception of the public Internet, service providers have sought to balance the cost of service provision with customer needs. In a dial-up ISP, over-subscription may cover:

- The number of incoming calls a dial-up ISP can accept simultaneously (the number of modems installed in its point-of-presence);
- The aggregate capacity provisioned in the ISP's internal network (which will be less than the aggregate modem capacity, since most users will have "idle periods" on their connections even when online); and
- The capacity of the ISP's Internet peer and transit connections.

Narrowband oversubscription is typically described only in terms of the modem:customer ratio, since the narrowband user is more likely to notice an engaged signal on a dial-up attempt than temporarily slow download performance.

Broadband uses a different access model. In some broadband architectures, customers have dedicated access to a broadband port with shared upstream bandwidth from the access port; other models share both the access network and the upstream bandwidth, as outlined in Table 6[7], below.

Table 6
Broadband Access Network Contention

Broadband Service Type	Customer Tail
DSL	Dedicated DSLAM port; shared capacity upstream from DSLAM.
HFC	Shared customer access network
Point-to-Point Wireless Broadband	Dedicated wireless transmitter/receiver pair to ISP point-of-presence; shared capacity upstream
Shared Wireless Broadband (eg: 802.16)	Shared customer access network
Broadband Over Powerlines	Shared customer access network

[7] Table 6 presents common deployment models for each technology. Individual service provider architectures may vary.

The economics of Internet services do not allow the elimination of access network contention.

However, there is a very great disparity between business DSL services (which are available as contention-free or 1:1 services at a price premium) and consumer services (which typically experience contention of greater than 40:1 and in extreme cases may exceed 100:1).

While VoIP sessions occupy only limited bandwidth, excessive contention in the access network can degrade VoIP performance in other ways as well. For example, a congested access network may cause traffic to be excessively buffered, resulting in latency outside acceptable levels for a voice service.

Where there is any possibility of congestion across a network, or across a series of network domains, end-to-end QoS mechanisms are required.

5.3.2 QoS-Based VoIP Peering

Today, most QoS implementations are confined to private networks, either providing QoS for environments such as corporate LANs and WANs, or in managed carrier networks to provide QoS-enabled private IP networks.

To provide QoS support for traffic on interconnected networks requires the ability for service providers to present QoS-marked (or otherwise characterised) traffic to each other at the network boundary, confident that the QoS-marked traffic will be handled in accordance with underlying service level guarantees.

There are several models in which inter-domain QoS may be implemented:

- **Project-based QoS Agreements** — Driven by the requirements of a customer project, two carriers may agree to apply similar QoS markings at an interconnect point. This approach is not scalable beyond a limited number of projects, since QoS needs to be managed and reported on a project-by-project basis.
- **Bilateral QoS Agreements** — Carriers may agree to interconnect QoS-enabled networks under standardised arrangements, to allow each other to 'productise' their respective QoS services while expanding the geographical reach of those services. The scalability of bilateral arrangements is limited by the need to manage relationships on a carrier-by-carrier basis.
- **Multilateral QoS Agreements** — This would mirror other multilateral models. At a traffic exchange point, such as a peering exchange, carriers may also agree to respect a given set of QoS markings.

Figure 14 illustrates a simplistic view of a multilateral QoS exchange point.

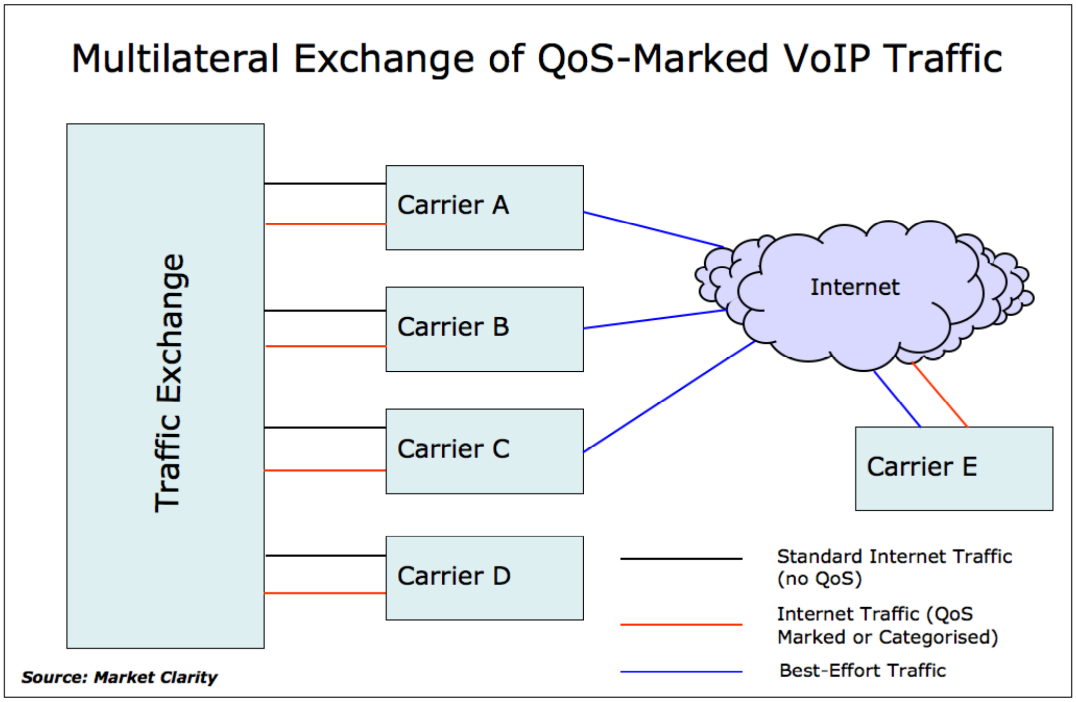

Figure 14
Multi-Lateral Exchange of QoS VoIP Marked Traffic

While the conceptual model shown in Figure 14 appears simple, it presents considerable complexities of implementation and management, including the need for common inter-domain QoS marking schemes (or a "trusted" QoS characterisation mechanism), managing the impact on routing infrastructure, and the preservation of QoS markings if traffic needs to traverse networks without QoS support.

We believe the multilateral model offers the best prospect for achieving the eventual goal of supporting high quality (as discernable to the human ear) VoIP services on an end-to-end basis, across service provider domains.

5.4 VoIP QoS: Issues for Industry Consideration

In examining whether an industry-wide approach to VoIP QoS is desirable, the industry may wish to include consideration of the following:

(a) **VoIP QoS Specifications** — The industry may wish to consider initiatives to foster the development of agreed QoS specifications for services intended to carry VoIP traffic.

(b) **QoS for Internet Access Networks** — The industry may wish to consider initiatives to foster the development of Internet access services designed to provide support for real-time traffic such as VoIP.

(c) **Co-ordination** — The industry may wish to consider how best to co-ordinate the development of inter-domain QoS mechanisms and VoIP peering mechanisms.

(d) **QoS Signalling** — As providers seek end-to-end Internet-based calls between their networks, the use of inter-provider signalling will expand. The industry may wish to consider whether it is desirable to incorporate QoS signalling mechanisms within VoIP call signalling; and if so, whether appropriate industry standards already exist or if a new development effort is required.

(e) **Cost Recovery** — If inter-domain QoS initiatives are instigated, the industry may wish to consider appropriate tariff models for these services.

(f) **QoS SLA Metrics** — The industry may wish to consider agreed mechanisms for measuring inter-domain QoS performance against SLAs, and processes for resolving failure to deliver to the SLA.

(g) **Consumer Information** — As differentiated consumer services emerge, the industry may wish to consider a consumer awareness strategy designed to inform users of the network QoS requirements necessary to deliver PSTN-like VoIP performance.

(h) **Proprietary VoIP Services** — The industry may wish to consider whether industry-wide initiatives are able to support VoIP services that use proprietary protocols.

(i) **Timing** — The industry may wish to consider the timing and relative priority given to QoS-related initiatives for VoIP services and Internet access networks.

6 INTER-DOMAIN QOS BUSINESS CASE

Cost recovery is a key question in any broad-based rollout of inter-domain QoS, since carriers and service providers alike will require justification for the investment both in network upgrades and ongoing system management.

Within a single provider's network, the business case for a QoS implementation is driven by:

(a) **Competition** — Providers can use QoS-enabled networks to create new customer services; or they may be driven to do so in response to other providers' QoS-based services.

(b) **Economics** — QoS-based networks can help manage the utilisation of providers' network infrastructure. Its proponents argue, therefore, that QoS capabilities can be justified on the basis of network cost of ownership.

These drivers, however, may not apply to the question of inter-domain QoS models. The inter-domain QoS business case has additional obstacles to overcome.

(a) **Competitive** — Inter-domain QoS involves co-operation between competitors. The expected benefits of an inter-domain QoS regime therefore must outweigh any perceived advantage a provider may accrue by maintaining its network as an "island" of QoS.

(b) **Technical** — As discussed in Section 5, above, QoS markings are only valuable on an inter-domain basis if they are meaningful to all parties to the QoS mechanism. The business case for inter-domain QoS must be sufficient to justify the overhead of creating a standardised set of QoS definitions.

(c) **Industry Fragmentation** — Another perceived risk would be that Internet providers taking part in an inter-domain QoS mechanism may face a temporary price disadvantage relative to providers which choose not to participate; or that any fragmentation may result in a lesser degree of interconnectedness.

(d) **External Criticisms** — Questions of QoS can easily be confused with the issue of network neutrality, and thus face criticism as creating "a two-tier" Internet. A multi-tiered Internet does not need to sacrifice principles of network neutrality, however, since within any tier, the key requirement is that carriers do not discriminate against other carriers.

6.1 Participants in Inter-Domain Peer-based QoS for VoIP

An end-to-end VoIP conversation would traverse a large number of elements owned by different parties. After the signalling path illustrated in Figure 9 (Section 4.2 above), the media path for VoIP customers on different networks may include:

(a) Retail ISPs associated with originating customer, VoIP providers, and destination customers;

(b) Carrier services (including DSLAMs) giving retail ISPs access to customer CPE;

(c) Wholesale ISP services;

(d) Transit services providing access to backbone Internet routes, and/or peer exchange ports between wholesale ISPs;

(e) Session border controllers at the edge of VoIP provider networks; and

(f) VoIP provider peer connections.

The interaction of all these elements needs to be considered in building the QoS business case, for the following reasons:

(a) To determine models which preserve QoS markings across unsupported network links;

(b) To broaden industry support for industry-wide QoS initiatives; and

(c) To broaden the base for cost recovery, should this be required.

6.2 Co-operation and Bilateralism

Bilateral models are attractive on a small scale, since they are likely to involve already trusted parties who are willing to create the necessary processes to implement inter-domain QoS.

Bilateral arrangements have other advantages as well: they could help build a trust base which could benefit the later development of multilateral peer-based inter-domain QoS mechanisms; and they would give participants experience with the associated business case which could inform multi-lateral efforts.

The risk is that if bilateralism proceeds too far ahead of any industry-wide initiatives, it could create an "installed base" of arrangements, which cannot later be aligned with the technologies, standards or processes used in multilateral initiatives.

Regardless of the timeframe needed to realise any multilateral arrangements in "live" services, the industry may wish to consider whether an early effort directed towards providing information on technologies and business processes will help to avoid the creation of "islands" of bilateral peer arrangements which cannot later be integrated in broader initiatives and services.

6.3 Encouraging Trust

Both bilateral and multilateral arrangements will require fostering of trust between a large number of organisations.

Initiatives that may help in fostering this trust could include:

(a) The development of model terms and conditions to place participants in multilateral efforts on a level playing field;

(b) The development of common tariff elements and SLA metrics; and/or

(c) The fostering of QoS and peering infrastructure which is managed independently of particular carriers or service providers, and which is accessible on a membership basis.

The industry may wish to suggest and consider initiatives that encourage the required trust between providers necessary for QoS-based peering.

6.4 Cost Recovery

Whether inter-domain QoS and VoIP peering models develop on a bilateral or multilateral basis, providers will need to recover various cost components.

Examples of cost recovery mechanisms may include:

(a) Membership based on the recovery of a service's foundation costs, ongoing operational costs, and the incremental cost of supporting new entrants to any proposed inter-domain QoS and VoIP peering facility;

(b) Differentiated charging based on the proportion of a service provider's port (or ports) which is set aside for high-QoS traffic;

(c) Differentiated charging based on separate ports which are set aside for high-QoS traffic;

(d) Transactional charges for control-plane VoIP peering services;

(e) Volume-based charging based on the percentage of traffic at a given QoS level; and/or

(f) Combinations of the above models.

The industry should consider whether any particular cost recovery mechanism is more likely to foster co-operation in industry-wide initiatives, and if so, how the favoured cost recovery mechanism may be implemented.

6.5 What other activities can be leveraged into QoS?

Much of the peering and QoS debate is also relevant to other emerging Internet applications, most particularly:

(a) **Multimedia personal communications** — Personal video services inherit VoIP's service requirements with the added requirement of higher bandwidth;

(b) **Video Services** — While best-effort Internet connections can achieve acceptable performance for medium-quality video downloads, they are inadequate for applications such as live HDTV streaming;

(c) **Gaming** — Players of online action games are keenly interested in high-performance Internet services, particularly with respect to low latency.

(d) **Support for deaf, hard of hearing, and speech-impaired individuals** — Work has commenced within the IETF on the technical considerations required for persons with speech or hearing impairments use of VoIP services. (See IETF RFCs 3351, 4103 and 4117.)

The industry may wish to consider:

(a) What other Internet applications or activities may benefit from inter-domain QoS and VoIP peering initiatives;

(b) How the requirements of other sectors may best be leveraged to support any proposed initiatives; and

(c) How to encourage QoS–enabled peering points to be "multi-service capable" from the outset.

7 COMMERCIAL ISSUES FOR INDUSTRY CONSIDERATION

This section identifies factors that may impact on the development of industry-wide strategies for inter-domain QoS and VoIP-based peering. These factors are described briefly and placed before the industry for consideration.

7.1 Prioritisation of Activity

The industry should agree on the relative priorities of any activities commenced to address VoIP QoS and peering issues.

7.2 Alternative Approaches

While this paper has concentrated on the potential to secure industry-wide co-operation to support VoIP peering and inter-domain QoS mechanisms, it must be recognised that other approaches may be adopted. These include:

- The creation of bilateral or multilateral arrangements under providers' own commercial decision-making;
- The pursuit of acceptable VoIP performance by other means (for example, provisioning more bandwidth); and
- The use of existing regulated numbering schemes in conjunction with public ENUM databases to support interconnectivity.

In examining whether industry-wide approaches are feasible or desirable, the industry may also wish to consider whether there are advantages to the alternatives listed above; or whether other viable alternative approaches exist.

7.3 QoS Islands

As providers increasingly deploy QoS-capable technologies within their own networks, there is a risk that they may use incompatible definitions of key network performance parameters such as availability, latency, jitter, packet loss, and restoration time. This may, in turn, inhibit the later adoption of inter-domain mechanisms that rely on consensus definitions for network behaviour.

The industry may wish to consider seeking submissions providing input to an agreed definition of QoS parameters relevant to VoIP services, so as to facilitate the early and broad adoption of these parameters alongside providers' existing or new QoS schemes.

7.4 Apathy

Any industry-wide initiative needs to attract support from a wide base of ISPs, carriers, VoIP providers, and other interested parties such as transit providers and peering exchanges.

If industry-wide inter-domain QoS and VoIP peering mechanisms are sought, the industry may wish to consider means to foster co-operation among the broader industry.

7.5 Security

Both inter-domain QoS and VoIP-based peering rely on the ability of trusted parties to submit requests of each other's network infrastructure, to have these requests honoured by the network of which the requests are made, and to have those requests traverse other networks intact.

For example, an ENUM server that responds to anonymous requests could be leveraged to create VoIP spam lists; or malicious users could leverage a network's QoS capabilities to gain inappropriate priority for their traffic.

The industry may wish to consider techniques and processes that foster inter-provider trust and maintain appropriate and efficient use of network resources.

7.6 Network Neutrality

Over the past few months, a debate has been raging over the concept of "network neutrality" and the possible creation of a "two-tiered Internet," in which Internet providers would start charging companies to access customers over their networks. Other industry commentators describe the issue as concerned with "traffic shaping" that would result in a two-tiered Internet, which is counter to the concept of network neutrality, where all data over the Internet is treated equally. One of the reasons for this concern is that an Internet provider might use traffic engineering to make their own VoIP service work well, while degrading the performance of other VoIP providers.

The adoption of inter-domain QoS mechanisms, particularly those associated with prioritising traffic on an application-specific basis, may give rise to concerns that the new services represent a threat to Internet network-neutrality.

The Wikipedia[8] describes network neutrality as: "a principle of Internet regulation with particular relevance to the regulation of broadband. It suggests that (1) to maximise human welfare, information networks ought be as neutral as possible as between various

[8] http://en.wikipedia.org/wiki/Network_neutrality

uses or applications, and (2) if necessary, government ought to intervene to promote or preserve the neutrality of the network."

However, the counter argument is that by providing traffic engineered services, applications such as VoIP that are timing sensitive, will perform better. Rather than detracting from current Internet performance, the addition of a new service class (or classes) is part of the natural progression of the Internet (the next generation Internet, or the second generation Internet), and that this is simply the introduction of new services, which can create new revenue opportunities.

In many respects, Internet peering may already be undermining the concept of network neutrality, in that members of a peering point trade traffic in a more efficient and cost-effective manner. Yet, the Internet community has embraced peering because of its efficiencies and cost advantages. It can be argued that QoS-based VoIP interconnectivity is simply an extension of a widely used practice, and is inevitable.

It is arguable, however, that much of the debate rests on a misunderstanding of the requirements of network neutrality, since the Internet already has multiple tiers of access networks. A consumer may purchase a service that is underpinned by a heavily oversubscribed DSLAM, while a business customer can pay a premium for a "1:1" access service.

The emergence of QoS-managed services reflects a change in implementation rather than a chance in philosophy. Instead of building and provisioning physically discrete networks to enable quality differentiation, the technologies discussed in this document allow such services to be delivered over a single infrastructure.

An industry-wide agreement standardising the treatment of QoS would help protect the key aim of network neutrality, which is interconnectivity and fairness, since it would help ensure that carriers do not discriminate against traffic from other carriers.

In considering the implementation of inter-domain QoS for Internet services, the industry may wish to consider:

(a) Whether it is desirable to also put in place mechanisms or rules protecting the integrity of the public Internet.

(b) Whether it is desirable to consider the wider issues of a multi-tiered public Internet, where each "tier" may have specific network behaviours that are tied to application performance.

7.7 The Australian Regulatory Regime

Australia's *Telecommunications Act* includes the principle of any-to-any connectivity as a component of the long-term interest of end-users.

In considering the implementation of inter-domain QoS and VoIP-based peering, the industry may wish to consider:

(a) Whether the regulatory emphasis on any-to-any communications has any impact on the implementation processes adopted.

(b) The competitive implications of PSTN interconnect are overseen by the ACCC. In considering the implementation of inter-domain QoS and VoIP-based peering, should the industry structure its peering initiatives such as to avoid the need for later regulatory intervention?

7.8 Disability Support

Australia's regulatory environment also stipulates that telecommunications services should support users with disabilities such as hearing impairment. In the case of telephony services, this support includes the availability of amplified telephone handsets, and TTY (teletype) services.

Currently, VoIP providers are exempted from this requirement, but as these services are increasingly positioned as PSTN replacements, they will need to replicate the PSTN's support for TTY and other services which support the disabled.

Voice quality is important in this discussion, since a poor-quality conversation will have a greater impact on the conversations of the hearing-impaired.

The chief requirement for TTY support in VoIP conversations is that the text messages remain synchronised with the conversation (that is, that the text is carried in near real time). Current IETF standards work may be consulted as background to the progress of TTY support in Internet environments.

The industry may wish to consider the degree to which questions of disability support can be incorporated in work on other quality-related initiatives.

7.9 International Experience

In addition to the initiatives outlined in this document, a number of other international activities may exist not known to the authors. This may include national initiatives such as is taking place in the UK; or commercial activities similar to IPeerX and SIP-IX.

In addition, a large number of informal VoIP address resolution initiatives exist.

The industry may wish to consider conducting further research to discover the extent of national, commercial, and informal initiatives relevant to this project.

7.10 Impact on Existing Peer Relationships

Since existing bilateral and informal multilateral VoIP peering arrangements are known to exist, the industry should consider the way in which these arrangements may interact with any industry-wide initiatives.

A possible positive impact of this may be that existing informal peers may be able to offer technical and business case experience relevant to any initiatives undertaken on an industry-wide basis.

On the other hand, however, pre-existing groups may regard the emergence of industry-wide initiatives as threatening.

The industry may consider ways in which the aims of existing multilateral VoIP peering initiatives may be met within the framework of an industry-wide approach.

7.11 Branding

As VoIP-based peering and inter-domain QoS mechanisms develop, participants in these initiatives may find competitive advantage in promoting:

(a) Their ability to provide Internet-to-Internet VoIP calls between providers; and/or

(b) Their ability to provide high quality calls traversing the Internet on an end-to-end and cross-provider basis.

As these initiatives develop, the industry may wish to consider developing means by which participants may identify themselves or their relevant services.

8 INTERNATIONAL INITIATIVES

In this section we discuss three international VoIP peering initiatives:

1. IPeerX [USA, expanding to Europe and Asia Pacific];
2. The SIP-IX partnership of NeuStar, Equinix, Telehouse and the Amsterdam Internet Exchange (AMS-IX) [USA, Amsterdam and Tokyo]; and
3. UK-based initiatives by Ofcom and the Network Interoperability Consultative Committee (NICC).

The authors are aware of additional VoIP peering initiatives underway from CableLabs (USA), FiberNet Telecom Group (USA), Stealth Communications (USA and London), InfiniRoute (USA, London and Madrid), and Dutch cable operators (The Netherlands).

This is a new market, and new VoIP peering initiatives are emerging on a regular basis.

8.1 IPeerX

IPeerX, a spin-off from Pulver.com, offers control-plane peering on a transaction basis seeking to commercialise the experience of Pulver.com's Free World Dialup peering. The service currently claims a "federation" of more than 140 members.

8.1.1 About IPeerX

Peering is offered as part of a suite of services, including:

- PSTN Bypass
- E.164 Directory Listings
- Non-SIP and Non-Standard SIP Protocol Conversions
- Network Monitoring
- Rules Based Routing
- Clearing and Settlements

Peering is based on an IPeerX-operated database, which associates numbers with providers. The association between number and user remains proprietary to individual providers.

The peering environment supports ENUM, as well as offering a proxy-based service for providers not using ENUM.

The call process for providers using the proxy-based service is as follows:

- VoIP provider receives request for off-network number;
- VoIP provider sends SIP INVITE message to IPeerX proxy;
- Proxy queries database to see if call can be completed on an IPeerX member's network;

- If call can be completed on-network, IPeerX proxy sends SIP INVITE to terminating member;
- Terminating member completes call to terminating device.

Under ENUM, the originating provider instead requests an ENUM lookup from the IPeerX server.

Where calls cannot be completed (either under the proxy model or the ENUM service), the call reverts to the originating provider, which makes its own policy-based decision on call handling (for example, call failure, PSTN traversal, or reversion to another service).

Members' agreements are with IPeerX rather than with other members. This, the company says, eliminates the problem of managing authentication rules between a growing number of bilateral network partners.

IPeerX says its service is designed to facilitate fast and simple implementation for members. Upon joining the "federation", an IPeerX member is assigned a carrier code. The member then decides how best to configure its systems to reach IPeerX and is assigned a timeslot to test its configuration. IPeerX claims the process of joining its network can be accomplished in less than an hour.

IPeerX does not offer media-based services, nor does it operate data centre or rack space facilities.

8.1.2 **IPeerX and QoS**

By restricting its activities to the control plane, the IPeerX service enables what may be described as "added value" opportunities.

For example, while IPeerX does not route traffic, the results of an address lookup could in the future include route information, such as whether the destination provider is reachable through a higher-QoS service. This allows the originating member to choose between having the conversation routed through the public Internet or, if the VoIP provider is a customer of a carrier IP network such as Level 3 Communications, the call may transit that carrier link for part of its journey.

This would allow VoIP providers to choose managed networks to carry their traffic as near to its destination as possible, minimising the number of hops the call has to traverse the public Internet.

If such a model were successful, IPeerX contends it would provide at least an entry point to addressing QoS issues, since the managed IP carrier networks can support QoS.

However, the use of a private IP network in such a way does not necessarily impact, either positively or negatively, on the feasibility of, or development of, inter-domain QoS.

8.1.3 **Commercial Model**

IPeerX uses a transaction-based commercial model: member VoIP networks pay IPeerX for successful transactions.

A low transaction cost is designed to attract member networks and build a network effect: more member networks will result in more successful queries, which in turn should help attract still more member networks.

8.1.4 IPeerX — Other Commentary

ENUM — While IPeerX operates an ENUM server, the company expressed a concern that unrestricted use of ENUM could seriously damage the VoIP market. Within the IPeerX environment, ENUM transactions are only permitted between members.

IPeerX nominated "VoIP spam" as a serious risk, should ENUM servers be widely available on the Internet. Queries could be constructed which iteratively retrieve valid ENUM entries for use as the basis of automated VoIP calls over the Internet.

If VoIP spam were to become widespread, the company believes it would restrict the utility of VoIP to those users who have the technical knowledge and tools necessary to manage the security of their VoIP services.

Other Peering Providers — IPeerX also works with a number of VoIP media peering facilities, which allow IPeerX to provide the signalling.

8.2 The SIP-IX Partnership

In October 2005, Equinix a provider of network-neutral data centres and Internet exchange services (Internet peering services), and NeuStar a provider of "clearing house" and directory services entered into an agreement to jointly develop "the first in a new generation of services to enhance the interconnection of networks providing advanced services under Session Initiation Protocol (SIP)."

NeuStar, also offers a number portability service in North America, and was the operator of an ENUM trial (this was hosted at www.enum.org; the trial has now concluded).

8.2.1 About SIP-IX VoIP Peering

The Equinix/NeuStar project, SIP-IX, proposes both control plane and data plane peering, with NeuStar providing ENUM resolution (in the control plane) and Equinix providing IP data peering (the data plane) as well as hosting SIP servers to support customers who have not implemented ENUM.

NeuStar's involvement also helps peers maintain consistency of E.911 information on an inter-domain basis. Other participants in the project include Telehouse (a US-based provider of carrier-neutral collocation space and services) and the non-profit Amsterdam Internet Exchange (AMS-IX).

The service is being developed to support carrier-class VoIP peering to target carriers or service providers which operate both the VoIP infrastructure and the underlying network, rather than Internet-based VoIP services which do not operate significant network infrastructure.

Implementation of SIP-IX is taking place as a two-phase project:

(a) Phase One — The first phase is designed as an information-gathering exercise. It does not involve integration between Neustar and Equinix. Neustar is making SIP and ENUM repositories available over the Internet with no specific connection requirements. Its aim is to provide experience in provisioning and telephone number routing techniques.

(b) Phase Two — In the second phase of the project, SIP and ENUM repositories will be made available at each of the exchange points involved in the project (Equinix, AMS-IX, Telehouse, and the SIP-IX exchange in Tokyo). Carriers will be able to connect to the Neustar hosts either directly to Neustar or at the network excahnges. Those who take advantage of exchange access to the SIP-IX environment will also be able to use the peering fabric for the final media path of their calls.

8.2.2 **SIP-IX and QoS**

Equinix described the relationship between an initiative such as SIP-IX and the development of inter-domain QoS, as first depending on the degree to which service providers themselves implement QoS-based networks.

There is no uniform view regarding the value of inter-domain QoS, especially in the United States, where an excess of fibre build in the late 1990s and early in this decade is still suppressing the cost of bandwidth. As a result, the purchase of new network links is still perceived as simpler and less expensive than enabling and managing QoS schemes.

However, the use of capacity alone as a QoS mechanism ignores the importance of service reliability as a measure of service quality. Many best effort links are not provisioned across protected or self-healing capacity, relying instead on the ability of the network to route around an event such as a fibre cut, or an extremely high traffic event causing localised network congestion.

Where the supported services are not time-sensitive (for example, e-mail), new routes are likely to be discovered and used without serious user impact. However, the same event would terminate VoIP calls on the affected route.

Continuity of service is, therefore, an important consideration for a network owner seeking to position its VoIP service as a primary line replacement.

A degree of service quality management can also be achieved by using session border controllers in the following process:

(a) A Session Border Controller manages admission control from the Internet.

(b) SBC ports are numbered so as to use BGP netblocks to differentiate best-effort traffic from VoIP traffic.

This allows the controller to drop “best effort” packets while still recognising and processing VoIP packets: it acts as a form of application-specific routing at the peering point. It is, in essence, a “port-based” QoS mechanism in which a provider is associating a particular port in a co-location facility with traffic that requires higher QoS. This technique is used within SIP-IX to separate VoIP traffic from best-effort traffic.

While current industry standards support very fine-grained QoS schemes, Equinix favours a “minimalist” approach in which QoS definitions go no further than “gold, silver, and bronze” service classes. This greatly reduces the complexity involved in reaching cross-provider agreement about how to handle QoS markings which originate off-network.

Key inter-domain QoS challenges identified by Equinix include:

- The need for mechanisms which can preserve QoS markings across infrsastructure such as peering points;

- Persuading network owners (that is, ISPs and IP carriers) that drivers exist to justify offering VoIP traffic from competitors' networks priority over best-effort traffic from their own customers; and
- Establishing industry consensus as to whether there are particular parts of the network where QoS is best implemented — for example, whether it is more beneficial to focus on protecting the priority of VoIP traffic in the access network.

8.2.3 **SIP-IX Commercial Model**

At this early stage of development, the service is offered at no cost to participants. However, Equinix advised that the aim is to design a service that is aligned with the partners' existing business models.

Among the lead partners, this could mean that Neustar would charge a transaction fee for address resolution, while Equinix could base its business case on the sale of ports in its exchanges.

8.2.4 **Equinix — Other Commentary**

Equinix says it has designed the SIP-IX model to reduce the cost base of VoIP providers who use flat-rate billing models. In particular, the company believes the SIP-IX environment will become more attractive as providers' volumes rise.

The partners are also considering whether the opportunity exists to create a variety of "a la carte" services. For example, Neustar could carry out billing on behalf of service providers, or give customers access at the peering point to off-the-shelf E.911 and lawful intercept services. The peering point would also provide a suitable environment for authenticating providers to prevent malicious activities such as attackers spoofing the IP address of known VoIP providers.

Equinix has also observed the following partitions within the VoIP peering community: QoS-based VoIP providers are willing to peer with each other; but not with best-effort based VoIP providers. Within each of these service provider "camps" are further subdivisions, and providers tend to peer with other organisations that fit their profile. These partitions are illustrated in Table 7, below.

Table 7
VoIP Peering Partners

VoIP Network Classification	Easy Peering Mechanism	More Complex Peering Mechanism
QoS-based VoIP Network	Dedicated ports for VoIP traffic, where a variety of QoS/traffic management techniques may be used within each provider's network. QoS markings are preserved, but they don't matter in terms of the peering. The only IP traffic passed across the port is VoIP traffic. The receiving carrier then marks the service appropriately for their network.	Class based traffic filtering using packet marking and traffic engineering techniques. This is a more complex model, which would require commercial agreement on QoS markings and service class definitions.
	Cheap and Cheerful	**Pseudo-QoS**
Best Effort VoIP Network	Best effort network design principles.	Use of capacity over-provisioning to emulate QoS. This type of network will exhibit QoS based network characteristics, until such time as a route failure or Distributed Denial of Service (DDOS) attack occurs.

8.3 Activities in the UK

8.3.1 Ofcom Statements on Network Reliability

VoIP peering and QoS discussions are at a very early stage, with the UK regulator, Ofcom, releasing a consultation paper discussing VoIP regulation on 22 February, 2006 (*Regulation of VoIP Services — Statement and Further Consultation*[9]).

This consultative document is, in part, seeking to update an earlier consultative document issued in 2004, *New Voice Services: A consultation and interim guidance*[10].

In the latest consultation paper, Ofcom associates QoS primarily with questions of network reliability and integrity, particularly in association with access to emergency services, and the degree to which VoIP services need to replicate PSTN-like network availability.

9 http://www.ofcom.org.uk/consult/condocs/voipregulation/voipregulation.pdf

10 http://www.ofcom.org.uk/consult/condocs/new_voice/anew_voice/?a=87101

8.3.2 Network Interoperability Consultative Committee Efforts

The UK's NICC (Network Interoperability Consultative Committee) is working on issues of VoIP number interconnect as part of its more general brief to act as an enabler to network interoperability in the UK. NICC is a joint industry-regulatory initiative with representation from carriers, service providers, and Ofcom.

At an open forum held in November 2005, NICC made public the timeframes applicable to its current "NGN interoperability" work in a presentation entitled *The Role of NICC in the delivery of NGNs*[11].

The NICC project plan envisaged its current round of standards work reaching fruition by March 2006, with documentation posted on its Web site. Because the NICC only makes documents available to member organisations, it was not possible to assess the completeness of these interoperability specifications. However, it appears that these standards, once complete, will be returned to the industry, either for review or for implementation.

Inputs to the NICC project plan are listed in Table 8, below.

Table 8
NICC Project Plan Inputs

NICC Document Number	NICC Document Title
NICC Study No. 53	QoS in New Network Technologies
NICC Study No. 68	EGTPS
NICC Study No. 70	NGN UK Interconnect Signalling
NICC Study No. 71	IP Transport Interconnect Specification
NICC Study No. 72	IP Interconnect Management Requirements
NICC Study No. 73	IP Interconnect Security Requirements
NICC Study No. 74	IP Interconnect Transport Architecture
NICC Study No. 76	NGN Numbering, Naming and Addressing
NICC Study No. 77	NGN Test Specifications
NICC Study No. 78	Ring Back When Free

Ofcom perspective on NGNsand NICC[12], presented to the same forum by Ofcom, noted the incompleteness of various standards covering interconnection and signalling issues.

11 http://www.nicc.org.uk/nicc-public/Public/open_forums/NICC.pdf

12 http://www.nicc.org.uk/nicc-public/Public/open_forums/steve_unger.pdf

8.4 ETSI and ITU Standardisation Processes

Network convergence and the consequent issues of QoS and signalling are also the subject of work undertaken at the European Telecommunications Standards Institute (ETSI) and the International Telecommunications Union (ITU).

The ETSI NGN Release 1 specifications were launched on 9 December 2005. To prepare the specification, ETSI progressively merged the following committees: Network Aspects (ETSI NA); Signalling, Protocols and Switching (SPS); Services and Protocols for Advanced Networks (SPAN); and Telecommunications and Internet Protocol Harmonisation Over Networks (TIPHON).

The first document is described as providing core specifications to bridge the fixed and mobile domains. Further work will address other aspects of network convergence.

The ITU had at the time of writing just completed a workshop in Las Vegas to report on the status of the NGN standardization process.

The ITU-T NGN group has so far produced 1,200 documents covering what the ITU describes as "fundamental framework areas for NGN", including:

- Services and capabilities
- Functional architecture and requirements
- QoS
- Control aspects
- Security issues
- Migration of current networks into NGN
- Future packet-based network requirements.

A full analysis of these initiatives is beyond the scope of this document. However, it appears from a perusal of the ITU-T NGN roadmap that the following timeframes are anticipated for work items relevant to this document:

- Performance Under Non-Heterogeneous Network Environment — Q1 2007
- Continuance of End-to-End QoS under different access network — H1 2007
- IP QoS Signalling Protocol — End 2006
- NTN Signalling Requirements — End 2006
- NGN Signalling Protocol — Q3 2007

Given the importance of international standards initiatives to the development of next generation networks, the Australian industry may wish to consider preparing a public analysis of the current status and roadmaps of relevant ETSI and ITU standards documents.

APPENDIX

This Appendix brings together the questions put to the industry throughout the rest of this discussion paper, with references to the section of text associated with each question.

Section 1: **Introduction**

Issues that the industry should consider include:

(a) Is end-to-end QoS feasible or desirable for consumer Internet services?

(b) What are the technical prerequisites to achieving inter-domain QoS support on consumer Internet services?

(c) What level of co-operation between providers is necessary to achieve inter-domain QoS, and is such co-operation achievable? For example, are providers willing to give preferential treatment to traffic sourced from other providers' networks?

(d) Does a complete suite of industry technical standards exist to support inter-domain QoS? If not, is the local development of such standards an appropriate effort for the Australian industry?

(e) If service providers were to adopt QoS technologies and begin exchanging inter-domain QoS information, would this have an impact on the infrastructure already in place supporting inter-domain "data plane" peering? Are peering exchange providers willing or able to accept QoS requests from client networks, or would this negatively impact either the peering fabric or the services available to other customers of the QoS-based networks?

(f) Do emerging models for "signalling plane" peering in VoIP services have any role to play in supporting the development of inter-domain QoS?

(g) Is there a requirement for an industry-wide strategy to support the implementation of inter-domain QoS in Australia? Should this be limited to public Internet interconnectivity, or should it include carrier IP/MPLS networks as well? For example, would ACIF have a role in helping to coordinate the phased introduction of relevant industry standards and/or model terms and conditions as new technologies and/or services become available?

(h) To what extent should any peer-based inter-domain QoS initiatives support signalling plane activities? Should important network utility services such as DNS be considered as part of the suite of services requiring inter-domain QoS support?

(i) To what extent can both intra-domain and inter-domain QoS mechanisms improve the overall system reliability achieved by VoIP services?

(j) What alternatives exist to the creation of an inter-domain QoS fabric?

(k) Can ACIF play a role in facilitating the standardisation of contractual terms and conditions necessary for providers to standardise inter-domain QoS definitions, and to honour the QoS requirements of traffic originating outside their own networks?

(l) To what extent, if any, would the adoption of QoS regimes violate the more general Internet principle of "network neutrality"?

(m) How can QoS be implemented across today's Internet access networks? Is it worth considering consumer offerings that provide premium services, say with lower contention ratios, which are designed to support real-time applications? Is an industry-wide strategy required to address this specific area?

(n) The current discussion is based on Internet connectivity only. Is there also a need to discuss carrier-provided interconnect between PSTN-based and Internet-based VoIP services? To what extent can such a discussion progress at this time, given the lack of international standards for carrier grade VoIP interconnect from organisations such as the ITU and ETSI?

(o) Given the range of issues raised above, and elsewhere in this paper, what timeframes and priorities should be applied to each activity?

Section 2.3.2: **The Impact of Network Behaviour on the User Experience**

Should ACIF consider recommending jitter buffers greater than 50 ms for VoIP CPE so as to avoid future interoperability problems?

Section 4.6: **VoIP Peering: Issues for Industry Consideration**

In examining whether an industry-wide approach to VoIP peering is desirable, the industry may wish to include the consideration of the following questions:

(a) To what extent is industry-wide co-operation desirable in implementing inter-domain VoIP peering? Is co-operation achievable, or will it create a new layer of peering, which exists alongside of, or in competition with, other peering initiatives?

(b) What provider vulnerabilities may exist in the implementation of inter-domain VoIP peering, and what steps can providers take to manage any vulnerability that may arise?

(c) Is the adoption or endorsement of particular Internet standards to facilitate inter-domain VoIP peering desirable or feasible? For example, is there any benefit in promoting particular signalling standards (such as SIP), codecs, or in standardising approaches to common issues such as firewall traversal?

(d) What role may organisations such as ACIF have in promoting a co-operative approach to peering in the Australian VoIP industry?

(e) Would a national approach to inter-domain VoIP peering have any impact, desirable or otherwise, on Australian VoIP providers' ability to implement peering arrangements with international VoIP providers?

(f) To what extent would multi-lateral VoIP peering agreements place customer information at risk of exposure, and what industry agreements or standards are necessary to avoid any risk that may arise?

(g) If Australian VoIP providers choose to adopt an industry-wide approach to inter-domain VoIP peering, in what ways might this interact with the Quality of Service issues (discussed in Section 5 of this discussion paper)?

(h) In what way would any proposed new initiatives align or interact with existing activities such as the ENUM trial currently being administered by ACMA? What liaison should be conducted with existing activities?

Section 5.4: **VoIP QoS: Issues for Industry Consideration**

In examining whether an industry-wide approach to VoIP QoS is desirable, the industry may wish to include consideration of the following:

(a) VoIP QoS Specifications — The industry may wish to consider initiatives to foster the development of agreed QoS specifications for services intended to carry VoIP traffic.

(b) QoS for Internet Access Networks — The industry may wish to consider initiatives to foster the development of Internet access services designed to provide support for real-time traffic such as VoIP.

(c) Co-ordination — The industry may wish to consider how best to co-ordinate the development of inter-domain QoS mechanisms and VoIP peering mechanisms.

(d) QoS Signalling — As providers seek end-to-end Internet-based calls between their networks, the use of inter-provider signalling will expand. The industry may wish to consider whether it is desirable to incorporate QoS signalling mechanisms within VoIP call signalling; and if so, whether appropriate industry standards already exist or if a new development effort is required.

(e) Cost Recovery — If inter-domain QoS initiatives are instigated, the industry may wish to consider appropriate tariff models for these services.

(f) QoS SLA Metrics — The industry may wish to consider agreed mechanisms for measuring inter-domain QoS performance against SLAs, and processes for resolving failure to deliver to the SLA.

(g) Consumer Information — As differentiated consumer services emerge, the industry may wish to consider a consumer awareness strategy designed to inform users of the network QoS requirements necessary to deliver PSTN-like VoIP performance.

(h) Proprietary VoIP Services — The industry may wish to consider whether industry-wide initiatives are able to support VoIP services that use proprietary protocols.

(i) Timing — The industry may wish to consider the timing and relative priority given to QoS-related initiatives for VoIP services and Internet access networks.

Section 6.2: **Co-operation and Bilateralism**

Regardless of the timeframe needed to realise any multilateral arrangements in "live" services, the industry may wish to consider whether an early effort directed towards providing information on technologies and business processes will help to avoid the creation of "islands" of bilateral peer arrangements which cannot later be integrated in broader initiatives and services.

Section 6.3: **Encouraging Trust**

The industry may wish to suggest and consider initiatives that encourage the required trust between providers necessary for QoS-based peering.

Section 6.4: **Cost Recovery**

The industry should consider whether any particular cost recovery mechanism is more likely to foster co-operation in industry-wide initiatives, and if so, how the favoured cost recovery mechanism may be implemented.

Section 6.5: **What other activities can be leveraged into QoS?**

The industry may wish to consider:

(a) What other Internet applications or activities may benefit from inter-domain QoS and VoIP peering initiatives;

(b) How the requirements of other sectors may best be leveraged to support any proposed initiatives; and

(c) How to encourage QoS–enabled peering points to be "multi-service capable" from the outset.

Section 7.1: **Prioritisation of Activity**

The industry should agree on the relative priorities of any activities commenced to address VoIP QoS and peering issues.

Section 7.2: **Alternative Approaches**

In examining whether industry-wide approaches are feasible or desirable, the industry may also wish to consider whether there are advantages to the alternatives listed above; or whether other viable alternative approaches exist.

Section 7.3: **QoS Islands**

The industry may wish to consider seeking submissions providing input to an agreed definition of QoS parameters relevant to VoIP services, so as to facilitate the early and broad adoption of these parameters alongside providers' existing or new QoS schemes.

Section 7.4: **Apathy**

If industry-wide inter-domain QoS and VoIP peering mechanisms are sought, the industry may wish to consider means to foster co-operation among the broader industry.

Section 7.5: **Security**

The industry may wish to consider techniques and processes that foster inter-provider trust and maintain appropriate and efficient use of network resources.

Section 7.6: **Network Neutrality**

In considering the implementation of inter-domain QoS for Internet services, the industry may wish to consider:

(a) Whether it is desirable to also put in place mechanisms or rules protecting the integrity of the public Internet.

(b) Whether it is desirable to consider the wider issues of a multi-tiered public Internet, where each "tier" may have specific network behaviours that are tied to application performance.

Section 7.7: **The Australian Regulatory Regime**

In considering the implementation of inter-domain QoS and VoIP-based peering, the industry may wish to consider:

(a) Whether the regulatory emphasis on any-to-any communications has any impact on the implementation processes adopted.

(b) The competitive implications of PSTN interconnect are overseen by the ACCC. In considering the implementation of inter-domain QoS and VoIP-based peering, should the industry structure its peering initiatives such as to avoid the need for later regulatory intervention?

Section 7.8: **Disability Support**

The industry may wish to consider the degree to which questions of disability support can be incorporated in work on other quality-related initiatives.

Section 7.9: **International Experience**

The industry may wish to consider conducting further research to discover the extent of national, commercial, and informal initiatives relevant to this project.

Section 7.10: **Impact on Existing Peer Relationships**

The industry may consider ways in which the aims of existing multilateral VoIP peering initiatives may be met within the framework of an industry-wide approach.

Section 7.11: **Branding**

As these initiatives develop, the industry may wish to consider developing means by which participants may identify themselves or their relevant services.

Section 8.4: **ETSI and ITU Standardisation Processes**

Given the importance of international standards initiatives to the development of next generation networks, the Australian industry may wish to consider preparing a public analysis of the current status and roadmaps of relevant ETSI and ITU standards documents.

GLOSSARY OF ACRONYMS

A complete technical glossary is beyond the scope of this document. Below is presented a glossary of acronyms for the convenience of readers.

ACCC	Australian Competition and Consumer Commission
ACIF	Australian Communications Industry Forum
ACMA	Australian Communications and Media Authority
ATM	Asynchronous Transfer Mode
BGP	Border Gateway Protocol
Codec	Coder/decoder
CoS	Class of Service
CPE	Customer Premises Equipment
DDOS	Distributed Denial of Service
DiffServ	Differential Services
DNS	Domain Name System
DSL	Digital Subscriber Line
DSLAM	DSL Access Multiplexer
ENUM	Electronic Numbers
ETSI	European Telecommunications Standards Institute
HDTV	High Definition Television
HFC	Hybrid Fibre-Coaxial
HTTP	Hypertext Transfer Protocol
IAX	Inter-Asterisk Exchange
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
IGP	Interior Gateway Protocol
IP	Internet Protocol
ISO	International Standards Organisation
ISP	Internet Service Provider
ITU	International Telecommunications Union
LAN	Local Area Network
LSP	Label Switched Path
Megaco	Media Gateway Control
MFA Forum	MPLS and Frame Relay Alliance Forum
MPLS	Multi-Protocol Label Switching
NAT	Network Address Translation
NGN	Next Generation Network
NICC	Network Interoperability Consultative Committee (UK)

Ofcom	Office of Communications (UK)
P2P	Peer-to-Peer
PHB	Per Hop Behaviour
PSTN	Public Switched Telephone Network
QoS	Quality of Service — used to describe the capacity of a network or service to function according to a given set of service parameters.
RFC	Request for Comment
RSVP	Resource Reservation Protocol
SBC	Session Border Controller
SIP	Session Initiation Protocol
SLA	Service Level Agreement
Speermint	Session Peering for Multimedia Interconnect
STUN	Simple Traversal of UDP NAT
TCP/IP	Transport Control Protocol / Internet Protocol
TDM	Time Division Multiplexing
TE	Traffic Engineered/Engineering
TOS	Type of Service
TTY	Teletype
UDP	User Datagram Protocol
URI	Universal Resource Identifier
URL	Universal Resource Locator
VLAN	Virtual Local Area Network
VoIP	Voice over Internet Protocol
WAN	Wide Area Network

REFERENCES

Industry Documents
CableLabs, VoIP Peering RFI (November 2005)
Alcatel Australia, Evan Stanbury, Contributions to the ACIF VoIP Working Group
Cisco, Internetworking Technologies Handbook (February 2003)
ETSI TISPAN home page: http://portal.etsi.org/portal_common/home.asp?tbkey1=TISPAN
IETF Draft (draft-meyer-voipeer-terminology-01), "Terminology for Describing VoIP Peering and Interconnect" (October 2005)
IETF Draft (sf-draft-twg-IP_PBX_SP_Interop-sibley-sipconnect-v3), "IP PBX / Service Provider Interoperability" (November 2005)
IETF, RFC 3270, "Multi-Protocol Label Switching (MPLS) Support of Differentiated Services" (May 2002)
IEFF, RFC 3031 "MPLS Architecture" (January 2001)
IETF, RFC 3032 "MPLS Label Stack Encoding" (January 2001)
IETF, RFC 2702, "MPLS Traffic Engineering" (September 1999)
IETF, RFC 3489, "STUN - Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs)" (March 2003)
IETF, RFC 4103, "RTP Payload for Text Conversation" (June 2005)
IETF, RFC 3351, "User Requirements for the Session Initiation Protocol (SIP) in Support of Deaf, Hard of Hearing and Speech-impaired Individuals" (August 2002)
ITU-T NGN home page: http://www.itu.int/ITU-T/ngn/index.phtml
ITU-T NGN Focus Group Proceedings (December 2005)
Metanoia Inc, Dr Vishal Sharma, Inter-Domain TE and QoS: Some Key Aspects and Challenges Presentation to ACIF (February 2006)
MFA Forum Press Release, "Fifteen Vendors to Demonstrate Interoperable Converged MPLS Services at MPLS World Congress" (7 February 2006)
MFA Forum White Paper, "Large Scale Multi-Vendor Layer 2 VPNs with MPLS, Paris 2005 Public Interoperability Event" (February 2005)
MFA Forum White Paper, "MPLS Inter-Carrier Interconnection (MPLS-ICI) Backgrounder" (December 2005)
Network Interoperability Consultative Committee (NICC), The Role of NICC in the delivery of NGNs (November 2005)

Interviews Conducted by Market Clarity for this Report
Steve Baxter, Technical Director, Pipe Networks
Dr. Vishal Sharma, President, Metanoia Inc (USA)

Lane Patterson, Director of Research and Development, Equinix (USA)
William Norton, Co-Founder and Chief Technical Liaison, Equinix (USA)
Kingsley Hill, CEO, IPeerX (USA)

Legislation
Telecommunications Act 1997
Telecommunications (Consumer Protection and Service Standards) Act 1999

The policy objective of the greatest practicable use of industry self-regulation without imposing undue financial and administrative burdens on industry is central to the regulatory scheme of the *Telecommunications Act 1997*.

ACIF was established to implement the policy of industry self-regulation. It is a company limited by guarantee and is a not-for-profit membership-based organisation. Its membership comprises carriers/carriage service providers, business and residential consumer groups, industry associations and individual companies.

ACIF's mission is to develop collaborative industry outcomes that foster the effective and safe operation of competitive networks, the provision of innovative services and the protection of consumer interests. In the development of Industry Codes and Technical Standards as part of its mission, ACIF's processes are based upon its principles of openness, transparency, consensus, representation and consultation. Procedures have been designed to ensure that all sectors of Australian society are reasonably able to influence the development of Standards and Codes. Representative participation in the work of developing a Code or Standard is encouraged from relevant and interested parties. All draft Codes and Standards are also released for public comment prior to publication to ensure outputs reflect the needs and concerns of all stakeholders.

Published by:
THE AUSTRALIAN
COMMUNICATIONS
INDUSTRY FORUM LTD

Level 9
32 Walker Street
North Sydney
NSW 2060 Australia

Correspondence:
PO Box 444
Milsons Point
NSW 1565

T 61 2 9959 9111
F 61 2 9954 6136
TTY 61 2 9923 1911
E acif@acif.org.au
www.acif.org.au
ABN 56 078 026 507