



aarnet

Australia's Academic and
Research Network

AARNet
GPO Box 1559
Canberra ACT 2601

www.aarnet.edu.au

Adelaide Office
Floor 7
10 Pulteney Street
The University of Adelaide SA 5005

Tel: (08) 8303 3936

Australian Communications Industry Forum
PO Box 444
Milsons Point NSW 1565
Email: acif@acif.org.au

Response to *QoS-based VoIP service interconnectivity*

Introduction

The AARNet network interconnects Australia's government research agencies and universities to each other and to the global Internet. The AARNet network has been running since 1989 and has supported quality of service for Voice over IP since 2000, complete with toll-bypass call hop-off and centralised number resolution. This long experience qualifies AARNet to comment upon ACIF's industry discussion paper *QoS-based VoIP service interconnectivity*.

We encourage ACIF's efforts in assisting the interconnection of IP telephony networks. The comments in our response are intended to assist ACIF's efforts.

It is not clear from the paper what process ACIF is undertaking with regard to VoIP service interconnection. We hope that ACIF's activities in this area lead to the issue of a set of guidelines for engineering the various qualities of Voice over IP interconnection; a set of small standards for the traffic reaching the various qualities of Voice over IP interconnection, including interconnections of PSTN-like quality¹; and the enhancement of the tutorial contained in ACIF's discussion paper.

General comments

Scope of paper

This paper is part of a process developing technical requirements for VoIP interconnection. However, the paper contains a great amount of tutorial discussion. We would have appreciated the tutorial discussion and the development of requirements for future advice, guidelines and standards being more clearly distinguished. For example, in naming MPLS and DiffServ as "true QoS" is the paper giving the author's opinion or it is recommending that an interconnection standard require MPLS and DiffServ?

The paper explores particular technologies and techniques for interconnection. AARNet would prefer that, where the paper is aiming to develop guidelines, the paper explore the requirements for interconnection, and leave the choice of technologies to meet those requirements to those seeking to implement an interconnection.

¹ Including Voice over IP connections to the Public Switched Telephone Network itself.

Delay, jitter, loss and dialling delay

A full exploration of the apportioning of delay, jitter, loss and dialling delay between interconnecting providers would be appreciated. This should be compatible with, and perhaps modelled upon, ITU-T's Q-unit scheme for wireline telephony.

The development of an ACIF guideline on these matters would allow providers to settle disputes over the proportion of the total acceptable delay, jitter, loss or dialling delay that another provider has consumed.

ACIF should consider whether the allowance for dialling delay allowance for an unsuccessful Voice over IP calls needs to be lower than the allowed wireline dialling delay of ten seconds. That would allow a failing Voice over IP call to be re-routed to the Public Switched Telephone Network. AARNet's experience is that a rapid indication of call failure is required for enterprise PABXs to successfully re-route calls.

Differing qualities of VoIP services and their requirements

The paper inconsistently discusses the resilience which Voice over IP service interconnection should support. Although AARNet3 is designed to achieve 99.999% availability we do not believe that this should be the norm for all ISPs and for all VoIP calls. Consumers would be better served by a range of voice offerings; ranging from "cheap and cheerful" through to "POTS replacement".

There is a question of how a consumer can correctly distinguish a "cheap and cheerful" from a "POTS replacement" service, which we discuss below. AARNet encourages ACIF to consider a branding for the Plain Old Telephone Service and POTS replacement services, so that consumers are aware of the clarity, reliability and resilience of their telephony service.

There are a range of VoIP services. AARNet finds this classification useful:

- *Cheap and cheerful.* These calls are originated from user-controlled terminals and use the best effort service. They do not use the ISP's VoIP control plane. Calls are not any more resilient than web traffic.
- *Goodish.* These calls are originated from user-controlled terminals. They use a VoIP differentiated service with policing. They use the ISP's control plane without admission control. In-progress calls are resilient to DoS attack and congestion. Placing new calls during a DoS attack can be problematic (as the control signalling uses best effort). The location finding equipment used by emergency services call centres fails to discover the location of a call.
- *POTS replacement.* These calls are originated from ISP-controlled terminals. They use a VoIP differentiated service with policing. Calls use the ISP's VoIP control plane with admission control. Calls are resilient to denial of service attack. The location-finding equipment used by emergency services call centres works to some extent, as a artificial Calling-Line Identifier representative of a site is associated with each terminal.

Obviously the interconnection requirements of these services differ. It would have been helpful if the paper had specified which types of service were being included in which suggested guidelines.

VoIP services not provided by the ISP

The discussion paper assumes that a Voice over IP service must necessarily be offered completely by the ISP. Yet nothing prevents a customer from running their own cheap and cheerful Voice over IP. It would be in the interests of consumers if Voice over IP services which are fully supported by an ISP are clearly distinguished from those that are not.

This should not prevent the provider from purchasing some elements of their supported Voice over IP service, such as PSTN peering, number resolution and interception capability.

VoIP peering and the public switched telephone network

AARNet is concerned that the paper spends little time on PSTN-VoIP interconnection. There is real potential for PSTN-centric carriers to suppress the provision of clear, reliable and resilient Voice over IP services by offering sub-par QoS for PSTN-VoIP interconnection. We would encourage PSTN, GSM and 3G carriers to make their networks accessible from VoIP, as knowing the codec used by the called terminal allows the calling terminal to make a considered choice of originating codec.

We would encourage the development of a ACIF standard to prevent PSTN networks offering only sub-par QoS for PSTN interconnection.

Denial of service

The paper does not address the use of QoS in allowing a VoIP call to survive an Internet-wide denial of service attack (such as Code Red²). Denial of service attacks are a major cause of widespread link congestion and it is important that the level of denial of service traffic which leaks into the VoIP QoS class is minimised, else the VoIP traffic class will become unserviceable just as link congestion means that effective Quality of Service would be most useful.

For VoIP interconnection to AARNet³ we require that peers do sufficient admission control or policing so that the resilience of voice calls to an Internet-wide DoS is preserved. A H.323 proxy or a Session Border Controller can be used to enforce traffic admission based upon reservations made by the control plane.

The paper should outline what call admission control and what policing parameters adequately prevent a VoIP peering link from permitting a denial of service attack to spread across the connection. These parameters will then allow engineers to correctly position Session Border Controllers and other traffic admission equipment.

Consumer protection

Although beyond the scope of this discussion paper, AARNet encourages ACIF to consider consumer protection when providers describe their voice services.

ACIF is encouraged to develop a strategy and branding standards so that consumers are fully aware of the ability of the service to offer:

1. clear voice

² CERT advisory CA-2001-19, see www.cert.org/advisories/CA-2001-19.html. Also see CAIDA's analysis of the spread of the worm at www.caida.org/analysis/security/code-red.

2. location signalling to emergency services
3. robustness in the case of network congestion or failure, and reliability of infrastructure.
4. geographical coverage.

The branding should incorporate existing PSTN and mobile services, not just VoIP services. In the interests of consumers, where the provider does not make their own claims, ACIF should evaluate the service as best it can and issue its own assessment.

Consumers should be encouraged to have one service offering (2) and (3) before they consider installing a second service. For example, a GSM handset in an area with good coverage could meet requirements (2) and (3).

Perhaps ACIF should require providers offering service to households to ensure that the household has a service providing (2) and (3) before promoting or providing another telephony service.

Occupational health and safety regulators should be informed of the shortcomings of current IP telephony in placing calls to emergency services so that workplaces can be regulated accordingly.

ACIF should seek to reach an industry consensus on the basis for call billing terminology. For example, does a “call” include unsuccessful calls? From the provider's point of view it should: the majority of cost and resources is in the control plane setup of the call, and a unanswered call costs the provider roughly the same as an answered call. From the consumer's point of view it should not: no communication has been provided.

Similar confusion exists around other billing models: terms such as “call minutes” (to the second, or rounded to the minute), “bytes” (including link layer encapsulation or not) all have areas of vagueness. Differing definitions are problematic for consumers: at best the confusion makes prices impossible to compare; at worst the consumer may claim to have been misled.

AARNet has experienced occasional large variations when reconciling its records with bills from providers. We note that consumers are not in a position to make or keep extensive call and bandwidth records. As with the PSTN, the onus is on the provider to retain and present enough information to allow for a contestable bill.

As part of another initiative, ACIF may care to issue guidance as to what data is reasonable for providers to retain to allow a Voice over IP service with billing that consumers can contest.

Interconnection and peering

The discussion paper discusses peering in depth. But there is little difference in interconnection technology between running QoS across a peering fabric and running QoS to a transit provider. So it is unclear why the discussion is restricted to peering points, rather than interconnections in general.

AARNet participates in bilateral peering and multilateral peering. Our preference where

small numbers of peering partners is present is for bilateral peering. Then we can implement features such as jumbo frames, QoS and multicast without needing to wait for such features to become economic for a multilateral peering point to offer. Bilateral peering also gives the AARNet Network Operations Centre a direct contact with the other provider's Network Operations Centre, which significantly reduces operational risks.

Our preference where large numbers of peers are present is for a connection through dual ethernet fabrics, with control-plane neighbourings being formed with each participant.

AARNet notes that although approximately 40% of our off-network traffic is exchanged with peering points, the price of transit has been decreasing steadily to the point where it is becoming close to the costs of peering. This is not too surprising, as the price of both is dominated by international transmission costs. AARNet regards the choice between purchasing transit and participating in peering as a business decision and thus beyond the scope of future ACIF guidelines.

AARNet expects that peering of voice calls will be much more economic than peering of best effort traffic, since there is currently a substantial difference between the cost and price of telephony transit.

AARNet believes that there may be public good in promoting further peering within Australia, as at present Australia's Internet communications are very exposed to a failure of the Southern Cross Cable Network³. AARNet would encourage ACIF to explore to what extent the stability of Internet and telephony infrastructure within Australia would be increased if Australian Communications & Media Authority were to mandate domestic peering between larger providers.

Traffic engineering

The paper confuses quality of service and traffic engineering.

Traffic engineering is the ability to create a path through the network which is not the least-cost path. However, because Australia has an essentially linear network topology — Perth, Adelaide, Melbourne, Sydney, Brisbane — the need for a great deal of traffic engineering is substantially less than in other countries. Scenarios on mainland Australia where the least-cost path is not the path preferred for traffic engineering are rare.

MPLS supports Differentiated Services, and that is a good thing, but MPLS is not essential for the delivery of a robust voice service.

AARNet is greatly concerned that a guideline may require the use of MPLS between providers. This would imply the interconnection of providers' MPLS control planes. That would discourage interconnection, as no provider would be willing to take the risk that an error by an interconnected provider may effect their entire traffic engineered network.

Video conferencing

AARNet notes that voice telephony and point-to-point video conferencing use substantially

³ This undersea network connects Sydney to the west coast of the USA via New Zealand, Fiji and Hawai'i. See www.southerncrosscables.com. This modern cable system is capable of carrying the large amounts of data required for Internet connectivity. Although Australia seems to have a large number of cable systems, most are not capable of providing circuits greater than 155Mbps.

the same network infrastructure and encourages ACIF to consider them as one, rather than as distinct services.

Quality of service

AARNet encourages ACIF to suggest Differentiated Services Code Points which providers should use. The Differentiated Services architecture allows the DSCP to vary between providers. However, this does not work well with customer-provided terminals, such as VoIP handsets. Those terminals have no standardised and automatic means to discover the DSCP for the voice class, and the consumer experience would be much better if the default values in those terminals worked.

AARNet suggests using DSCP = EF for voice telephony traffic, DSCP = AF41 for video conferencing traffic and DSCP = CS3 for voice and video control plane traffic. It finds that these values are widely supported in current terminals.

To support older customer-owned terminals, at the customer edge we re-mark Precedence = 5 to DSCP = EF and re-mark Precedence = 4 to DSCP = AF41. We set unsupported Precedences and DSCPs to DSCP = 0.⁴ Note that the re-marking is designed so that the DSCP-marked traffic exiting AARNet is compatible with customer networks which use Precedence.⁵

AARNet encourages ACIF to suggest policing values for telephony traffic at inter-provider connections, perhaps in terms of Kbps of single leaky bucket shaping per Erlang. A common basis for calculating policing would simplify the debugging of end-to-end under-performance.

We encourage inter-provider policing to re-mark overflow packets into a best effort class or into a worse than best effort⁶ class. This allows a call to proceed in the face of some misconfiguration.

AARNet wishes to ensure that only traffic which has been admitted by the provider's control plane is exchanged at the inter-connection. For example, SIP calls could be regulated by a Session Border Controller, H.323 calls could be regulated by a H.323 Proxy.

Multicast traffic in the AARNet3 network can only use the best effort and Scavenger classes. Although we have experimented with video-conferencing over multicast we do not recommend that this traffic be admitted to the any QoS class which has preferential treatment. A similar precaution seems wise for voice traffic over multicast.

TCP seeks to estimate the amount of available bandwidth by measuring the round-trip time taken by data acknowledgements. This becomes less precise as the amount of non-congestion aware traffic in the network increases. Until this effect is better understood we suggest that a link carrying both TCP (best effort and Scavenger classes) and non-TCP traffic

4 It is unclear if this is the best alternative. It may be better to treat unsupported DSCPs as though they were DSCP = 0 but not to re-write the DSCP. This would allow unsupported QoS services to transit AARNet3, but allow these QoS services to be supported by downstream networks.

5 Customer networks attempting to interpret Precedence rather than DSCP will see voice traffic's DSCP = EF as Precedence = 5; video traffic's DSCP = AF41 as Precedence = 4; voice control traffic's DSCP = CS3 as Precedence = 3; Scavenger's traffic DSCP = CS1 as Precedence = 1. We are careful not to use other DSCPs which would confuse this backwardly-compatible interpretation of common Precedence values.

6 An example worse than best effort class is Internet2's Scavenger service (DSCP=CS1), which offers 1% of the link's capacity to the Scavenger class when the link is congested. See qbone.internet2.edu/qbss.

(voice and video classes) restrict the non-TCP traffic to 30% of the link bandwidth to avoid a dramatic fall in TCP goodput.

Alternatives to QoS

AARNet encourages ACIF to state performance requirements so that they are engineering goals rather than in terms of technology. This allows requirements to be met by the use of non-QoS mechanisms, such as a completely independent Voice over IP network.

QoS requires support from all devices which can deliver packets with an inconsistent delay or in an inconsistent order. As a result not all network devices require the use of QoS to deliver the engineering goals. An obvious example being an SDH bearer, which can neither delay nor reorder traffic.

Not all of the devices in a provider's network which do have jitter or packet re-ordering may support QoS. ACIF's guidelines should be stated in terms which allow workarounds of sufficient quality.

Billing

The most common billing regime for transit traffic from a US provider is a monthly bill related to the 95th percentile of the five minute traffic average, where the "traffic average" is the maximum of input bytes and output bytes sampled each five minutes.

A moment's reflection suggests that the small amount of bandwidth used by voice calls and the short duration of voice calls implies that a 95th percentile of all traffic model will not recover sufficient funds to cover the costs for the additional engineering and router features needed for a quality of service deployment.

Thus AARNet finds it unlikely that the basis for VoIP charging will be the same as the basis for best effort charging. Potential alternatives are to apply the current billing regime independently to each QoS class, with a differing tariff per class; or to use the number and duration of calls as recorded by a Session Border Controller; or a simple monthly subscription, which avoids the huge costs of a contestable per-call billing infrastructure.

Control plane robustness

Where a robust inter-connection is required ACIF should state some requirements for the control plane, such as no one single failure causing loss of control plane signalling where the forwarding plane is still operational.

Legal interception requirements

Where a call crosses an interconnection point where does the responsibility for having an interception capability fall? At the moment it falls upon all parties which are licensed carriers. What if neither participant is a carrier?

Is it reasonable to require providers to have a telephony interception requirement for "cheap and cheerful" traffic which may be on the provider's network without the knowledge or participation of the provider?

Network neutrality

AARNet does not have a well-developed view on the debate over network neutrality, as the debate presents a conundrum for voice services.

Firstly, network neutrality for voice could present a threat to a provider's carefully-designed QoS implementation. For example, if a provider is forced to accept and interpret DSCP-marked traffic from a neighbour as if the traffic were marked in its own network then then poor admission control or policing by the neighbour could undermine the QoS of the provider's network.

Secondly and conversely, this argument allows a provider which controls a link layer (such as an ATM+ADSL network) to prevent other providers from offering a voice service as robust as the service which can be offered by the carrier which controls the link layer.

ACIF could effectively address the first point in the process of developing Voice over IP interconnection guidelines and standards, the second point seems to be a matter for the ACCC.

Comments on tutorial content

Chapter 2: VoIP background

Table 2 gives an ambitious time for "link aggregation, backup LSPs and MPLS fast reroute" at "less than 150ms". Given that the shortest undersea cable between the east coast of Australia and the west coast of USA has a round-trip time of 140ms, it seems unavoidable that users of all telephony and data services in Australia will occasionally face link failures of around 1s.

Table 3 has errors. E-mail has a low availability requirement; consider that e-mail was once transferred by FidoNet nodes which dialled each other nightly. Video streaming can cope with high delay and high jitter, since playback can be delayed to build a sizeable jitter buffer at the receiver. I wonder if "live streaming" and "video conferencing" have been confused.

ACIF recommending a particular jitter buffer size has little effect beyond setting the lowest tail bandwidth over which VoIP can have good quality.

This section lacks the most compelling reason for using QoS mechanisms: the defeat of congestion, from denial-of-service attacks and other sources.

Chapter 3: Internet peering models

See comments above, which are not repeated here.

"The most common model for transit services" is misleading. Transit is usually acquired complete with a tail circuit to the customer's site. Some customers build or lease tail circuits into an Internet exchange and take transit services there, but this is not as common in Australia as in USA. Some large customers or ISPs will lease an undersea circuit and purchase transit in USA. Billing models for transit traffic vary substantially: "per megabit per second per month" is not particularly common as it ignores the effect of convenient pipe sizes. An ISP might find it convenient to run gigabit ethernet across dark fibre to a transit provider, but they may find it substantially less convenient to be billed for 1000Mbps per

month. Thus most transit agreements measure the amount of traffic: megabytes per month is popular in Australia; 95th percentile of five minute average bandwidth is popular in USA.

Section 3.3.1 does not distinguish the flow of control traffic at a peering point. Some peering points interpose their BGP control plane, which requires an all or nothing approach to accepting routes⁷ and makes it difficult to enhance the control plane to include IPv6 and multicast traffic. Most peering points, such as PAIX and LINX, facilitate the flow of data but leave it to each peer to establish control plane neighbourings.

It is financially advantageous for a peer to offer as many routes to the peering point as possible, even those to which it has congested capacity. So the robustness claims for peering in the paper may not be as great as claimed.

Many transit providers offer service level agreements.

Outside of Australia, the price of transit has fallen remarkably. Even very large peering points overseas are finding it difficult to remain competitive with transit prices.

Chapter 4: Emerging peering models for VoIP

The description of Skype lacks one other important point: system administrators do not get to select the “supernode”. This has obvious risks, such as a supernode oversubscribing a mission-critical workstation.⁸

The other great advantage of using a Session Border Controller is that it gives a choke point for a legal telephony interception capability. For this reason all calls may be routed through a Session Border Controller, not just inter-provider calls.

Chapter 5: Quality of service

See comments above, which are not repeated here.

“QoS refers to the capability of a network to provide better service...” is incorrect. QoS mechanisms can also worsen performance, particularly to traffic that may be malicious.

The confusion of traffic engineering and quality of service is discussed above. The need for requirements to be met by non-QoS mechanisms is discussed above.

Section 5.2.1 should mention that ethernet does not support carrier loss signalling to all neighbours. Thus alternative mechanisms need to be used to detect a failed link. As part of setting a guideline for dialling time, ACIF will need to set a guideline for the time taken to alert neighbours to the failure of the inter-connecting fabric. Network engineers can then use metro ethernet OAM, Cisco UDLD, MPLS keepalive or a similar mechanism for detecting failure within the guideline.

⁷ Not entirely true. Some network layer peering points, such as Germany's DE-CIX, allow routes from a particular peer to be suppressed using a BGP community string. This feature is not yet widespread, as it relies upon a feature of the BGP software developed for DE-CIX.

⁸ Consider the widespread use of the WinNY peer-to-peer file sharing program in Japan. A fault in this software has led to compromised machines at Maritime Self-Defence Force, ANA, JAL, Okayama Prefectural Police Force (this compromise revealing details of 1,500 investigations) and Mitsubishi Electric Plant Engineering (this compromise revealing details of nuclear power plants in Mihama, Tomari and Sendai).

Chapter 6: Business case

See comments above concerning separation of tutorial and guideline-making content, which are not repeated here.

Chapter 7: Commercial issues

See comments above concerning consumer protection, which are not repeated here.

Chapter 8: International activities

No response.

Responses to selected questions

Section 1

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

The desirability of end-to-end QoS depends upon the expected robustness of the telephony service. If this is a “cheap and cheerful” service to augment an existing telephony service then failure of the service has only monetary consequences. However, if the service replaces a PSTN service then a requirement that the service not fail due to, say, Internet-wide congestion from a worm is required. QoS is one way, but not the only way, of meeting this requirement.

There is a concern that the owner of a link layer, such as a ADSL access network, ATM transit network, or peering fabric, could restrict the deployment of PSTN replacement VoIP services by not offering access to QoS or alternative mechanisms.

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

The major requirements are: the traffic is marked with a reasonable DSCP; all parties admit or police traffic to prevent mis-use of the DSCP; all parties implement a per-hop behaviour that furthers the purpose of the QoS class; all parties have delay, jitter, loss that is suitable for the class of traffic; and control plane delays are within bounds acceptable for the service.

(c) ...are providers willing to give preferential treatment to traffic sourced from other provider's networks?

Providers are willing to preference traffic from other networks if they can be assured that the traffic is admitted, marked and policed correctly.

This is an issue where an intermediate provider does not provide the traffic class. For example, a user may mark a video conference with DSCP = AF41 but their provider may re-mark that to DSCP = 0. There is no way for another provider to recover that lost information.

There is an argument that providers should interpret non-supported DSCPs as best effort (DSCP = 0) traffic, but not re-mark the DSCP. This is a difficult thing to do when customers that do and do not support QoS marking are attached to the provider's network.

The admission and policing requirements lead to all but the most trusting of providers to

communicate via Session Border Controllers.

(e) ...or would this negatively impact either the peering fabric or the services available to other customers of the QoS-based networks?

If the peering point can classify connections into QoS-enabled and not QoS-enabled peers then the peering point can implement QoS support without effecting best effort peering.

The question for the peering point is “is activating advanced QoS features likely to lead to instability in the fabric's software”? And the potential for instability may encourage peering points to offer a distinct fabric for QoS-enabled peering, as they do now for multicast peering.

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

The signalling plane can determine if particular flows have valid admission to the network. Since the validity of flows is of concern at an interconnection then each provider may choose to perform admission control before handing the data off across the interconnection. For example, each provider could use a Session Border Controller.

The interconnection of the providers' control planes is problematic. No provider wishes for its network to be vulnerable to an error in a neighbouring network. For example, one of the design principles of the AARNet3 network is that no transit, peer or customer has access to the core routers' control plane or administration plane.

Control plane interconnection is only likely to occur in parts of the control plane which can be isolated from the control plane of the provider's network: for example, by offering other providers number lookup through a proxy server rather than the same infrastructure used by the provider.

(h) Should important network utilities such as DNS be considered as part of the suite of services requiring inter-domain QoS support?

No.

(n) ... Is there also a need to discuss carrier-provided interconnect between PSTN-based and Internet-based VoIP services? ...

Yes; to the extent that this is necessary to prevent no QoS or sub-par QoS being offered for VoIP interconnection to the PSTN.

Section 2.3.2

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

ACIF should recommend a jitter buffer size which allows international calls to be placed through at least three interconnection points.

ACIF should also guide providers as to acceptable delay, jitter, loss and dialling delay for the span of network which they cover.

Section 4.6

(b) What provider vulnerabilities may exist...

AARNet's experience is that low loss rates which are common in networks are particularly audible in VoIP calls.

AARNet is concerned that inadequate admission control and policing of the QoS voice class will allow a denial of service attack which denies voice service and denies best effort service.

Section 5.4

(a) VoIP QoS specifications.

AARNet encourages ACIF to use the IETF's Differentiated Services QoS architecture.

We encourage ACIF to offer gentle guidance to providers for DSCP values, as it is convenient for consumers if these are consistent and widely known. AARNet has revealed its DSCP values for this purpose and we are willing to share suggested policing and per-hop behaviours.

This gentle guidance will allow inexperienced providers to make good choices, whilst allowing those providers which are forced by circumstance to make other choices to do so.

(f) QoS SLA metrics

AARNet encourages ACIF to develop measurements of QoS service levels, both at the packet and at the perceptual level. We note that this is not as simple as it may first appear.

(g) Consumer information

AARNet encourages ACIF to develop a mechanism to allow consumers to readily distinguish a range of telephony from "cheap and cheerful" to "PSTN quality". We see no reason to limit this mechanism to Voice over IP telephony: it should be relevant to fixed line and mobile telephony.

Section 6.2

...avoid the creation of "islands" of bilateral peer arrangements which cannot later be integrated...

AARNet suggests that the Differentiated Services model is so dominant that only the most odd "islands" could not be integrated. The availability of Session Border Controllers which support both SIP and H.323 goes some way to addressing the most likely cause of uninteroperable islands.

Perhaps the concern with bilateral peerings is economic. Certainly participants in the early bilateral peering relationships for best effort traffic in Australia have proved to be extraordinarily reluctant to include other providers.

Section 6.5

What other Internet application or activities may benefit from inter-domain QoS and VoIP peering activities?

Video-conferencing uses much the same infrastructure as voice, but uses significantly more bandwidth. Where there is sufficient bandwidth voice interconnections can be extended to be voice and video-conferencing interconnections with little work.

AARNet strongly encourages providers to implement a "Scavenger" worst effort service. This provides a class for traffic which may be network misuse but which should be forwarded until congestion is encountered. For example, by placing ICMP Echo into this class the effectiveness of "ping flooding" is greatly reduced.

Section 7.2

...whether viable alternative approaches exist

It is for the inter-connecting providers to decide what is a viable alternative. For example, two providers may simply inter-connect their VoIP-only IP networks (such as two international calling card operators).

However ACIF's guidelines should set requirements so that a third provider would see delay, jitter, loss and dialling delay within usable bounds.

Section 7.3

...agreed definition of QoS parameters relevant to VoIP services...

As discussed above, AARNet would prefer gentle guidelines for the implementation of Differentiated Services (such as DSCPs and policing parameters) but firm guidelines on voice parameters (such as delay, jitter, loss and dialling delay) and a firm guideline that traffic in a class has been granted admission or strict policing to the class (as an extreme example, so that a general purpose PC can't simply send many Mbps of traffic in the voice class).

Section 7.8

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

Firstly, IP telephony providers should not be permitted to exclude the use of customer-provided terminal equipment. This allows people with disabilities to use terminals containing assistive technologies. Sufficient information should be made available so that customer-provided terminal equipment may be successfully configured.

Secondly, an estimate of clarity of an IP telephony service should be available to the consumer before purchase. Because of the high cost of assistive technologies it is unreasonable to expect a consumer to test the service and request a refund if it is unsuitable.

Thirdly, the industry should encourage augmentation of the TTY service by Internet technologies such as e-mail and instant messaging.

Fourthly, we would encourage the industry to squarely address disability support issues, rather than relying upon charities to do so. For example, few ISP help desks have TTY support.

Section 8.4

...a public analysis of the current status and roadmaps of relevant ETSI and ITU standards documents.

AARNet strongly encourages ACIF to adopt technologies and techniques that are widely available, widely used, and standardised. Where these objectives conflict we hope for a considered choice, not merely choosing a technology because it is the publication of an international standards organisation.

We assume the lack of mention of the IETF is an oversight, rather than expressing a preference for ITU technologies such as H.323 over IETF technologies such as SIP.

Glen Turner.

Network Engineer, AARNet.

Email: glen.turner@aarnet.edu.au

Tel: (08) 8303 3936.