KNOSSOS

Technical Aspects of Unbundled Bitstream Service

20 May 2005

Index

Introduction

1 Attached as a schedule is a summary of my experience and that of our company, Knossos.

2 I have been asked by InternetNZ to comment on the draft UBS determination (the Draft Determination). The Regulated service is on the end-user side of the service to a service delivery point which is the trunk (that is, exchange) side of the ATM switch nearest to the DLSAM.

3 However, that regulated service does not stand in isolation. Further, to better understand the issues that impact the regulated service, it is desirable to review UBS as a whole, from the end-user to the access seeker. While the related regulated service is often called *"backhaul"*, it involves considerably more than mere backhaul in the sense of carrying traffic from Point A to Point B. The Access Seeker requires additional services, which either it or Telecom provides (currently under the commercial UBS services this is provided by Telecom). These components have an impact on the overall performance and metrics of the UBS service. I will deal with the position both generally and also specifically in relation to the regulated service which is the subject of the determination.

Unbundled Bitstream Service

4 It is understood that the proposed Unbundled Bitstream Service (UBS) will operate identically to the Commercial UBS currently used by some existing ISPs. It is therefore useful to describe the service and its technical elements.

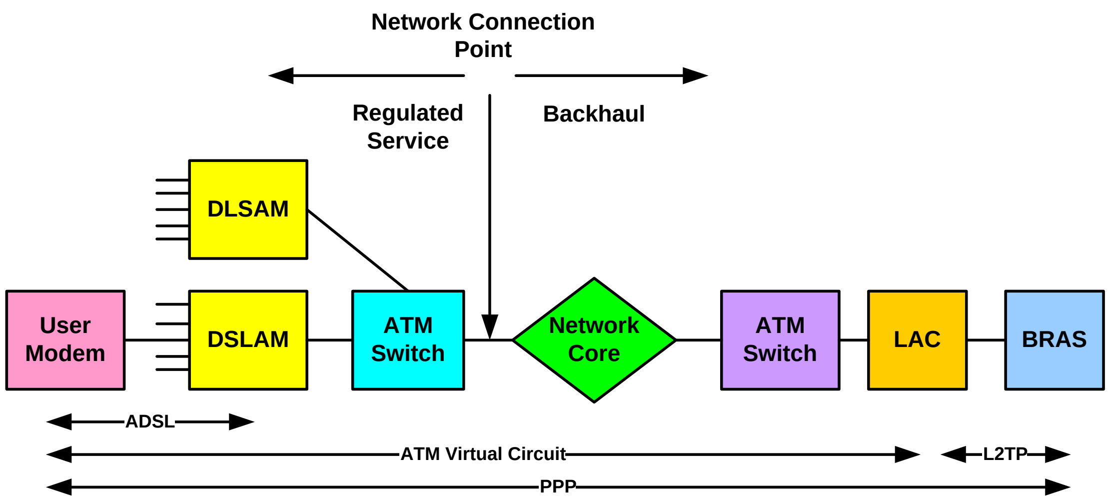

5 A user is attached to a DSLAM via copper, over which ADSL carries ATM cells. Over these cells, a virtual circuit connects a PPP (Point to Point Protocol) session from the user's ADSL modem through the DSLAM and across the network ATM core to the LAC (Layer 2 Access Concentrator). The PPP session is then carried via L2TP (Layer 2 Tunnelling Protocol) to the BRAS (Broadband Remote Access Server) operated by the service provider, or in the case of the regulated UBS, to the Access Seeker.

6 L2TP operates over IP, which in turn is delivered to the service provider/Access Seeker via an ATM virtual circuit. In Commercial UBS, access to the Bitstream service is restricted by dimensioning the ATM virtual circuit at a rate of 24 kbps per user. This happens between the LAC and the BRAS and this is the major restraint on

contention (see below for more details).

7 Access through the core network in the downstream direction is largely constrained by dimensioning the ATM virtual circuit from the BRAS. However it is possible to configure traffic shaping at the LAC or in the ATM network between the LAC and the DSLAMs, although for UBS and JetStream type services this is not currently, nor typically done. So, it would be possible to traffic shape at [?] (ie: within the regulated service) but that would be unusual. In the upstream direction, the traffic is constrained at the DSLAM to the upstream access rate, in the regulated UBS case, to 128 kbps.

Network Performance Measures

8 The following is a discussion of the Network Performance Measurement Parameters specified in Appendix A of the Draft Determination.

9 The metrics in Appendix A pivot around something called the “network connection point”. According to the Act and its schedule, this is the trunk side of the ATM switch nearest to the DSLAM. This should be made clear.

10 While the UBS service on the trunk side of that ATM switch is described by the Act as the backhaul service, this oversimplifies the position, as will be obvious from the description of the UBS above. Clearly, an Access Seeker wishing to use the UBS will need to use the core network up to the LAC and from the LAC to an interconnection point, which may be co-located with the LAC or subject to another interconnection agreement. Therefore, it is necessary to discuss the overall system

“95% of the time”

11 Under the Draft Determination, the performance measurement parameters, contention ratio, jitter, packet loss and latency are all defined as being required to be met “95% of the time”. This metric has two problems:

11.1 The metric is too low. “95%” seems like high availability to the lay person, but actually represents three minutes per hour, one hour and 20 minutes a day, one and a half days per month, or 18 days per year. These are rates that would be unacceptable to any reasonable service organisation.

11.2 Furthermore, having 72 minutes per day of leeway allows the network to operate outside the measurement parameters for the entire peak hour every day, without exceeding the overall network performance specification.

11.3 Secondly, the metric needs to be defined in terms of timeframes; is the metric met on an hourly, daily, monthly or yearly basis? How long can what is essentially a fault condition be allowed to last before it must be fixed?

12 The 95% limit is quite clearly at odds with the Service Availability figure of 99.3%, and should therefore be removed from these metrics. Failure to meet the metrics should constitute a fault, covered by the 99.3% Service Availability figure.

Contention Ratio

13 There are several problems with the contention ratio of 50:1 as defined in the Draft Determination.

14 Firstly, the contention ratio has a problem with definition. The Draft Determination requires "full speed" downstream ADSL, which may constrained by the ADSL line quality and distance, and by the Access Seeker's own traffic shaping, neither of which are under direct control of the Access Provider. This makes defining the peak bandwidth of a given ADSL line impossible without full knowledge of the actual line's and full knowledge of rate shaping within the Access Seeker's network

15 One could argue that the "peak bandwidth" is defined as the maximum reasonable throughput of the ADSL lines, i.e. 6-7 Mbps per line, regardless of quality or upstream traffic shaping.

16 Secondly, assuming that PIR is based on 1-2 Mbps per port, rather than the full capacity of the ADSL lines, experience suggests that a 50:1 contention ratio is too low. With Telecom's Commercial UBS product, interconnection is provisioned at 24 kbps per customer. A simple contention ratio for this product is difficult to calculate, as this same value is used for 256 kbps, 1 Mbps and 2 Mbps services. An even mixture of these speeds would give a contention ratio of 46:1, and a major user of Commercial UBS has provided figures that indicate a contention ratio of 37:1. ISPs using Commercial UBS (including the ISP that provided the 37:1 figure) report heavy congestion on interconnection circuits) ie: what's called *"backhaul"*). This suggests that 37:1 (and by extension the 50:1 specified by the Draft Determination) is too low.

17 The third issue with the contention ratio relates to infrastructure, interconnection and backhaul. Although the Draft Determination does not include backhaul, there is a risk that the contention ratio defined for the DSLAMs would set a precedent that could be carried through to a later backhaul definition. In any case, regulated UBS and Commercial UBS use essentially the same basic architecture for both interconnection and backhaul, and failure to meet adequate standards on interconnection circuits would render any other standards meaningless.

18 Similarly, infrastructure between the DSLAMs and the Network Connection Point must meet or exceed the standards specified in the Draft Determination; otherwise the standards are again meaningless

19 A major problem with a contention ratio, even if assigned on reasonable assumptions, is that it fails to take into account changing usage patterns. The overall trend in broadband is toward more traffic per user over time, representing a need to increase the contention ratio to match.

20 Most importantly, the contention ratio has been defined[1] in relation to Telecom's JetStream suite of products. The ratio of 50:1 may work well for Telecom's JetStream customer base, using JetStream services, but may not work well for other providers. Take for example the two JetStream 2 Mbps plans offered by Telecom. They cap usage at 10 gigabytes per month. If that 10 GB was downloaded 7x24, the download speed would be approximately 30 kbps, or a contention ratio of 66:1. In reality, Internet usage is variable during the day, typically culminating in a peak hour during the evening at which time total usage exceeds average usage by a factor of two or more. This gives a contention ratio

[1] Workshop Transcript page 16 line 34.

of 33:1 for the 2Mbps service if all users of the service used the service to (but did not exceed) the usage cap.

21 The following graph shows usage over a typical day for a typical broadband ISP:

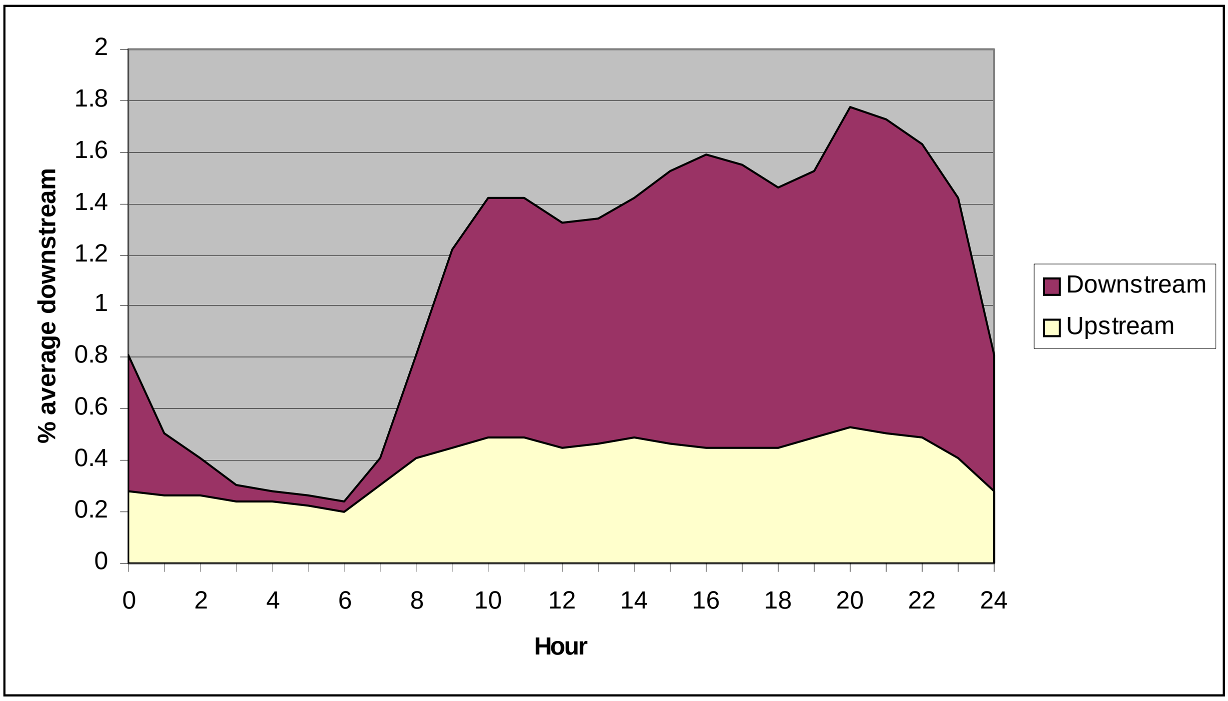

22 The following table lists the complete set of JetStream residential plans[2] with their fully utilised contention ratios:

Residential Plan	Go	Discover	Explorer	Adventure	Navigator
Downstream rate, kbps	256	1024	256	2048	2048
Allowance, GB	1	1	3	10	10
Price	39.95	44.95	49.95	69.95	69.95
Average download, 24/7 (kbps)	3.09	3.09	9.26	30.86	30.86
2:1 peak-time weighted download (kbps)	6.17	6.27	18.52	61.73	61.73
Contention ratio if fully utilised	42:1	166:1	14:1	33:1	33:1

23 With the exception of the Discover plan, all of the above plans, if fully utilised, would require a contention ratio of better than 33:1.

24 It would appear that Telecom's customer base includes a large number of users who do not use their service to its intended capacity, thereby giving the 50:1 contention ratio quoted by Dr Milner.[3] An Access Seeker wishing to offer similar

[2] Extracted from: http://www.telecom.co.nz/chm/0,5123,203071-202469,00.html

[3] Workshop transcript page 17.

plans, but experiencing a smaller proportion of "under-active" users than Telecom, would not be able to offer the same level of service using the same contention ratio as Telecom. Indeed, the contention ratio specified by Telecom is a "nominal" figure, subject to variation with experience[4]; the ability to vary this figure according to specific loads is an advantage that an Access Seeker would not have if the metric is the sole determinant of contention ratio for the access seeker.

25 The opposite side of this argument is that, without some kind of limiting, an Access Seeker could require constant-bit-rate access to an ADSL client; clearly this is not what is sought from this Determination. What is required is that for equivalence of service to apply, an Access Seeker should be able to offer plans within the same range that is available to Telecom, and experience the same performance, regardless of differences in usage patterns between the Access Seeker's customers and Telecom's customers. This equivalence is not achieved by aligning the contention ratio of the Access Seeker to that experienced by Telecom.

26 The above discussion assumes that PIR is defined in terms of existing JetStream or Commercial UBS type plans; as pointed out in the first few paragraphs, this is difficult in an unrestricted environment. It may be more useful to define contention in terms of the capacity of the ADSL ports, irrespective of line quality or upstream shaping by the provider. Alternatively, it may be better to define contention in terms of absolute bits per second per user.

27 In terms of basing capacity on port speed, if the ports are capable of operating at a nominal 7 Mbps, a 50:1 contention ratio would indicate the nominal bandwidth seen at the Network Connection Point should be 140 kbps per port, which would be suitable for most retail applications. Similarly, a DSLAM capable of only 6 Mbps per port could be configured with a network interconnect speed of only 120 kbps per port.

28 In all cases, it should be possible for an Access Seeker to offer the equivalent of the JetStream plans described and experience full utilisation of these plans under typical peak/off-peak traffic patterns. This requires that there be at least 62 kbps per user of aggregate downstream bandwidth available

29 With 128 kbps upstream bandwidth, it is unlikely that upstream contention will be a limiting factor. Upstream trunk bandwidth should be allocated in the same dimensions as the corresponding downstream trunk bandwidth.

Packet Loss

30 The figure of 3% packet loss is simply inadequate. Any reputable Internet service provider would consider 3% packet loss on an otherwise unloaded circuit as a fault. And then the metric requires even that low standard only 95% of the time.

31 The following table illustrates the effect of packet loss on simple download speed, under test conditions:

Packet Loss	Local	National	International

[4] Workshop transcript page 42.

0.0 %	5.2 Mbps	4.8 Mbps	2.4 Mbps
0.5 %	4.3 Mbps	3.8 Mbps	1.5 Mbps
1.0 %	3.4 Mbps	2.9 Mbps	1.0 Mbps
2.0 %	2.5 Mbps	1.4 Mbps	0.37 Mbps
3.0 %	1.8 Mbps	1.1 Mbps	0.29 Mbps

32 "Local" represents traffic from a source within 20 ms of round-trip latency, e.g. a resource close to DSLAM with minimal additional network latency. "National" represents a source within 40 ms of round-trip latency, as would be typical for a resource accessed from Wellington by a user in Auckland. "International" represents a source 180 ms round-trip away, as is typical for a resource located in the continental US. Rate shaping for the above figures was 128 kbps up, 6 Mbps down.

33 Clearly, when performing TCP downloads (such as using the World Wide Web or FTP) from a distant site, the long latencies greatly exacerbate the effect of packet loss, due to the way that TCP congestion control is done.

34 Thus, packet loss has a dramatic effect on user experience, even for simple downloads. The effect becomes noticeable after even a small amount of loss, and drops transfer rates to 50% at 1% of loss. At 3%, transfer rates have dropped by 88% for international transfers (worse, e.g. for a resource in Europe), and is even down by 78% for national transfers.

35 Interactive services, such as remote login become noticeably "laggy" when there are even fairly small levels of packet loss, however "laggy" is more difficult to quantify and demonstrate than the download rates represented above. Other services, notably services relying on many short separate transactions, are also more sensitive to packet loss than one-way TCP traffic. Online games may also be significantly affected by packet loss, if packets updating a player's position or actions are lost.

36 Any packet loss standards must factor packet sizes. For example, a steady uncorrected error rate of one in every 10,000 bits transmitted would impact approximately 3% of 40 byte packets, but more than half of 1500 byte packets. This level of packet loss would make a connection completely unusable; and clearly must not me allowed to pass by measuring packet loss with small packets. One could (and indeed should) interpret the definition of "population of interest" as including large packets, however this should be explicit.

37 Given that noticeable degradation occurs at packet loss rates of less than 1%, a packet loss ratio of 1% or more of packets of any size for an extended period should be regarded as a fault. Even 1% should not be considered normal.

Latency

38 Large latencies (the time taken for a packet to get through a network link) affect interactive applications, such as online games. Downloads largely adapt to large latencies, although they do affect start-up times and recovery from packet loss.

39 Latency has two major sources: transmission and contention. Transmission latency represents the time taken to put bits on the wire, and is most affected by the raw data rate and interleaving. At 128 kbps, a 64 byte packet takes 4 ms to transmit. Interleaving adds a delay, (configured at 16 ms in current Telecom configurations[5]). So for example, four 64 byte packets (each taking 4 ms of line time) would be sent together, adding additional latency to the transmission latency if interleaving was not turned on.

40 A 1500 byte packet would take 94 ms to be sent at 128 kbps with no interleaving; interleaving would add another 16 ms.

41 Contention latencies occur when several ingress channels converge into one egress channel. When this happens, traffic is queued until the egress channel is free. If there is heavy contention for the egress channel, packets or cells from one ingress channel may be queued for a considerable period while packets from other ingress channels are sent.

42 Transmission latencies tend to be consistent, as the factors affecting them don't change much with conditions or usage. Contention latencies vary enormously with load; indeed most customer complaints about high latency stem from contention delays, either due to poor dimensioning or unusual traffic patterns.

43 Contention latencies for a single packet are not noticeably affected by the size of the packet, whereas transmission latencies are at least partly proportional to packet sizes. However, where packets are sent through congested ATM circuits (as they are between the user and the LAC), the packet will be divided into one or more 48 byte cell payloads. This causes the packet's latency through that link to be the sum of the latencies of the cells making up the packet.

44 The specification defines latency in one direction only: from the user into the network.[6] This is the direction most prone to both (a) contention latencies between the end-user and the network, and (b) lowest bandwidth. However, it does not in any way define the downstream value, which should meet the same standard.

45 The figure given for latency of 50 ms seems as a worst case, although one would expect a significantly lower average latency.

Jitter

46 The jitter figure of 500ms appears erroneous, considering the latency figure of 50ms; it would be impossible to have jitter of 500ms when the minimum possible packet latency is less than one ms and the greatest allowable is 50 milliseconds. Transmission latencies for larger packets of equal size should be roughly constant, unless core network trunks are heavily congested, so one would expect that a network with a short-packet latency of 50 ms would not experience more than 50ms jitter on larger packets.

47 A figure of 50 ms, consistent with a maximum small-packet latency of 50ms, would be appropriate.

---

[5] Workshop transcript page 68.

[6] That's implicit from the fact the measurement is to be done at the end-user's premises.

Upstream Throughput

48 The specification defines a maximum upstream throughput of 128 kbps, but does not specify a minimum. Given that the low upstream rate heavily constrains the use of the service, it is essential that the service perform at close to its maximum rate at all times. Therefore, strict limits need to be applied to the minimum performance of the upstream service. A short-term (5 minute average) minimum of 32 kbps and a longer-term (one hour average) minimum of 128 kbps, consistent with the downstream throughput specification, would be appropriate.

Downstream Throughput

49 One recognises that overall throughput is highly dependent on line quality, and may additionally be affected in short term bursts by radio frequency and electromagnetic interference. With ADSL, the downstream component operates at higher frequencies to the upstream component, and is therefore more prone to degradation. Such signal degradation may result in reduced throughput, and short-term bursts of interference may cause significant degradation.

50 The Draft Determination gives short and long-term minimums of 32 kbps over a 5 minutes average and 256 kbps over a 60 minute average minutes respectively. These would appear to be reasonable, if only one set of metrics is to be used covering worst-case conditions. One would expect better performance on circuits terminated close to (e.g. less than 2km) the DSLAM.

Equivalency of Service

51 It will be apparent from the above that any ISP (whether Telecom or another ISP) would find consistent failure to meet the above metrics as being unacceptable. This raises the question as to whether having these metrics in the first place provides a useful reference point. The draft determination highlights the importance of equivalency between services available to end users via Telecom and those available by the Access Seeker. They should be equivalent and the focus, I suggest, should be on this, rather than the specific metrics. The categories of metrics listed in the Appendix (such as latency, etc) can be measured and compared. If there are to be specific metrics (such as latency of 50ms), this can be regarded as a backstop measure, not that to which the Access Provider and the Access Seeker should ultimately aspire. To be added to those metrics should be performance requirements around OSS such as for provisioning, fault response, and so on.

Interleaving

52 Interleaving represents a trade-off between latency and reliability. Data transmitted over an ADSL line is divided into groups of small blocks, and error-correcting codes (ECC) are used to re-create any blocks within a group lost due to line noise. However, the re-creation process can only re-create a block if other blocks in the group are intact. If a noise burst destroys two blocks in a group, the whole group, and therefore the packet(s) contained in the group, is lost. Interleaving mitigates this effect by allowing groups to accumulate over a short period, and interleaving the blocks of each group. If a single noise burst destroys two adjacent blocks, each block belongs to a different group, and therefore both blocks can be re-created. The effect of interleaving is that there must be a wait-time for packets to accumulate, so they can be interleaved. This has an effect on latency, as enough data must be accumulated to fill all the groups specified by

the interleaving setting. (If data does not arrive during the wait time, empty blocks may be transmitted.)

53 ISPs using Commercial UBS with whom interleaving was discussed indicated no particular desire to turn interleaving off, preferring to make the trade-off in favour of reliability rather than reduced latency. This does not mean that the option to turn it off should not be available.

54 However, interleaving is not a simple on-off parameter. Interleaving is configured as a delay in which groups can accumulate prior to interleaving and transmission. Telecom configures an interleaving delay of 16 ms[7], but other values are possible. If interleaving is to be specified by the Access Seeker, the interleaving delay should also be able to be specified. There may be cases where equipment has only a fixed set of interleaving delays, in which case the best match should be selected.

Service Availability

55 The service availability figure of 99.3% represents a service outage of up to 2.5 days per annum. This percentage is not consistent with the 95% specified under Contention Ratio, Jitter, Packet Loss and Latency, and indeed all observers consulted would agree that of the two figures, 99.3% is the more reasonable for all parameters. However, even 99.3% falls short of most industry benchmarks.

56 It is not clear from the specification what "Service Availability" refers to. A customer of the Access Seeker sees only his or her own service, so a service that is not available to every user is not available. Thus, the Service Availability should be clarified to include each individual circuit used by the Access Seeker.

57 Should the Service Availability figure represent just the core elements of the service, this should be treated as a carrier grade service, and the industry benchmark of 99.999% availability applied.

Summary

58 The critical parameters of Contention Ratio and Packet Loss are simply too low for a credible competitive service to be offered using UBS, as is easily demonstrated by simple tests and experience with Commercial UBS. In short, a service operating within these values could still be considered broken by any reasonable user or service provider.

59 It is tempting to request that the UBS be defined in terms of Telecom's own experience with retail services, such as JetStream, and in fact to a large degree this has been done in Telecom's earlier submissions to the Commission. However, this is a poor approach to defining a wholesale service, as the Telecom user base and service offerings may be a poor match for the services an Access Seeker may wish to offer, even in the same market space.

60 In particular, Telecom services tend to be a path of least resistance for many less sophisticated users; Telecom advertises heavily, and a large proportion of the JetStream user base had other non-Internet Telecom services prior to taking up JetStream. Thus it is likely that by comparison with other ISPs' customer bases,

---

[7] Workshop transcript page 68.

JetStream users on average represent a lower level of usage per user than those which have gone out of their way to use a service provider other than Telecom.

61 Thus, to assume that contention ratios and contingent parameters that work for experienced loads within the JetStream service should apply to other service providers is to say that all other service providers must have JetStream-style customers and offer JetStream-style services. Experience with Commercial UBS has already shown that the kinds of usage experienced by other ISPs' loads are severely constrained by these parameters.

62 Furthermore, by expressing quality-of-service parameters purely in terms of a retail service offered by Telecom means that other services can potentially offer cheaper, but poorer services that that offered by Telecom, but can not compete on quality. They can be worse, but not better.

Recommendations

63 Our recommendations regarding the Network Performance Measurement Parameters are as follows:

Contention Ratio

64 The contention ratio should be defined at 50:1 of the unrestricted ADSL port bandwidth irrespective of line quality or Access Seeker shaping. The available aggregate downstream trunk bandwidth at peak times must be not less than 62 kbps per user, and the available upstream trunk bandwidth per user should set to the same level as for the downstream.

Jitter

65 Jitter should be <50 ms; we believe the current figure of 500 ms to be erroneous.

Packet Loss

66 Packet loss should be less than 1%.

Latency

67 Latency should be < 50 ms one-way in either direction for 64 byte packets. For larger packets, the additional latency should scale with the ADSL bit rate.

Additional Limit (a): Upstream Throughput

68 Upstream bandwidth must operate at a one-hour average rate of not less than 128 kbps and a 5-minute average rate of not less than 32 kbps.

Additional Limits (b & c): Downstream Throughput

69 Consideration should be given to providing greater minimum throughput for service provided closer to the exchange.

Interleaving

70 In addition to being turned on or off, the interleaving delay should be a parameter than can be specified by the Access Seeker.

Service Availability

71 Service on each line should be available and operating within the Service Performance Measurement Parameters 99.3% of the time.

Stability of Unlimited Downstream Service

72 Question 5 of the Draft Determination asks whether there are risks of instability in provision of a bitstream access service with unlimited downstream speed and a 128 kbps upstream speed.

73 An unlimited downstream service is not strictly "unlimited"; it is limited by the actual capacity of the ADSL line, which is in turn dependant on line condition, distance etc. It is also limited by the ability of the applications to send data when the rate of acknowledgments is limited by the 128 kbps upstream speed. As such, it is no different to a limited ADSL line, except that the unlimited service can potentially deliver data at a higher rate, if the applications allow it.

74 It is always possible in the absence of proper network engineering to have instability caused by excessive traffic entering the system without sufficient rate limiting. However, this is true regardless of whether limited or unlimited service is offered.

Operational Support Systems

75 The Draft Determination makes no attempt to define technical details of the Operational Support Systems (OSS) required to support the UBS, beyond the requirement that:

> *"Telecom is accordingly required to provide a level of operational support to TelstraClear, whether manual or automated, such that there is no material difference in provisioning or fault repair in regard to the experience of retail customers whether retail services reliant on bitstream access are supplied to TelstraClear or Telecom." (Paragraph 276 of Draft Determination)*

OSS Experience with Commercial UBS

76 One large provider described the provisioning process of Commercial UBS. They noted:

76.1 The process was largely a manual one, involving filling out application forms provided as Microsoft Word documents, and emailing the completed document to Telecom for action.

76.2 The forms were complex in format and including several macros, and which proved practically impossible to generate electronically, requiring a manual process to create and forward.

76.3 The documents required details that the provider considered irrelevant, as all these details would have already been available to Telecom, being associated with the telephone number provided. (Commercial UBS can only be provided on an existing telephone with an associated Telecom number.)

76.4 Any inconsistencies, including disagreement between details provided on the forms and details recorded against the existing telephone number, would cause the application to be rejected by Telecom. The rejection rate was described as “high”.

76.5 The existence of any outstanding service orders involving the telephone line or number would cause the application to be rejected by Telecom; the provider found this particularly onerous, as it meant that Commercial UBS and other telephone services could not be requested simultaneously.

76.6 Processing times had improved over the last few months. Average turnaround times of 30-40 days had been experienced early in the provision of Commercial UBS, but were now down to less than 20 days, but still in significantly excess of the 7-10 days promised.

76.7 There had been problems when transitioning from JetStream to Commercial UBS, where it appeared that there was a two-step process in disconnecting from JetStream and then connecting to UBS. The user could be without service between the two steps for some hours or even days. This appears to have been replaced with a single step process.

76.8 Fault reporting and investigation was described as inadequate, with very little reporting and often considerable difficulty convincing Telecom that there was in fact a fault.

OSS Recommendations

77 Electronic OSS should address most of the above issues. It should be possible for an Access Seeker to electronically enter orders into Telecom’s provisioning system, and be able to query the status of such orders.

78 There should be minimal impediments to such orders being processed; identification of the line should not require extensive consistency checks; a wholesale service should provide mechanisms to immediately confirm that an order will be carried out.

79 Such OSS should provide for new orders, Moves Adds & Changes and faults. All should require minimal information to be performed, e.g. a new connection should require no more than the telephone number to apply the new service to, and any technical information such as the interleave factor. There should be immediate confirmation that the order has been accepted, and status reporting as to the estimated time of completion and any issues or queries raised. All such OSS should be capable of being driven by automated systems as appropriate.

80 It is essential that OSS interfaces be simple, to allow rapid and inexpensive implementation by both Telecom and the Access Seeker.

Curriculum Vitae - Don Stokes

Don has been working in information technology since 1985 and in the Internet services industry since 1992, when he joined Victoria University of Wellington in the newly created position of Network Manager. As VUW was also the Wellington hub of the then embryonic Internet in NZ, he became intimately involved in the hub's operation, and managed the transition to a commercial Internet services provider. This operation was re-launched in 1995 as NetLink, and Don managed all technical operations until his departure in 1997.

From 1997 onward, Don operated a private consultancy. With John Rumsey, he formed Knossos Networks Ltd in 2001, to provide network consultancy, support services and provide Internet services to a wider customer base.

Some highlights of Don's career include:

- Developed carrier grade IP accounting system, 2004;
- Development of Domainz Ltd web and e-mail services, 2003;
- Implementation of DNS infrastructure for .NZ Registry Services for NZ TLD operations, 2002/2003;
- Convenor of InternetNZ committee investigating the "0867" dialup Internet interconnection arrangements, 2001;
- Member of InternetNZ Registry Review working group that recommended that NZ move to a shared registry implementation, 2000;
- Implementation of independent NZ DNS infrastructure for Domainz, 2000;
- Development of small-scale dial-up ISP, 1999;
- Redevelopment of ISP systems at Globe.Net, 1998/1999;
- Initial development of IP traffic accounting package for Paradise Net, 1998;
- Formation of NetLink and expansion of services from Internetworking Group leased lines operation to include dial-up, web and e-mail services etc, 1994/1995;
- Formation of VUW Internetworking Group, 1992.