Michael S Cox

Response to:

NBN Co consultation paper: proposed wholesale fibre bitstream products

This response follows the same layout, in terms of section numbers, as the original "NBN Co consultation paper: proposed wholesale fibre bitstream products" released December 2009. Within each section of this document are comments and responses corresponding to the original material and questions as documented in each specific section of the NBN Co consultation paper. The exception to this, is the addition of an appendix section in this response.

Overall the consultation paper reflected a very high quality of progress in the challenging area of defining and describing the intended wholesale bitstream products that will support a vibrant and successful competitive market. It is hoped that the responses included here will further help to optimise that competitive market, as well as ensure various additional social benefits and ensuring that NBNCo remain a viable ongoing entity within the market.

This response aims to represent the combined best interests of the Australian residential and business users (consumers) **first**; an open, competitive and efficient retail service provider marketplace (producers) **second**; and an economically viable NBNCo wholesale access provider (enabler) **third**. It is hoped that this hierarchy of consideration is aligned with NBNCo, its key stakeholder (the government) and the remaining stakeholders (end-users, carriers, businesses, vendors and suppliers).

Section 1 Introduction

Response to Background

Acknowledge the scope limits applied around the consultation paper. Some discussion on pricing models is included in the response, despite being described as out of scope for this particular engagement. Such discussion is necessary to highlight support for the proposed QoS models and some other associated areas.

Response to Summary of NBN Co's proposed wholesale fibre products

It is noted that NBNCo seeks to "*occupy as small a footprint as possible in the overall value chain, leaving retail service providers (RSPs) with significant ability to innovate …*". The challenge here of course, is knowing exactly where to balance Occam's[1] razor with Einstein's[2] razor. Based on comments made by Mike Quigley (during the January 2010 industry engagement sessions), regarding the existence of '*several organisations that are keen to offer independent layer-3 wholesale access services*', it appears that the overall choice to focus on layer-2 wholesale service products is most likely correct in the current Australian context.

A similar challenge facing NBNCo is deciding where and when "*contestable backhaul*" will exist and hence whether to offer the LEB product or the AEB product. This matter is not as clear-cut as the aforementioned. In particular, there are some concerns around statements such as "*only one PoI will be available for any FSA*". As such, several counter arguments will be offered that show that this approach may not necessarily support optimal market conditions nor behaviours. These comments can be found in the response to sections 6 and 7.

It is noted that NBNCo will leverage the capabilities of a GPON/Ethernet platform to provide products with support for security, QoS and multicast. This combination of technology and service offerings is well aligned to industry best practices and will serve the Australian market well.

It is noted that NBNCo intend to "*support access by multiple RSPs, a range of CPE and will include an interface for analogue telephony*". This combination of service capabilities is also well aligned with industry best practice. Again, the exact architecture for supporting multiple RSPs and the corresponding multitude of CPE will require some degree of balance and trade-off. A good litmus test for such architectures is ensuring that the simplest environments can always be solved using the simplest solutions whilst also ensuring that the most complex and demanding environments can still be accommodated in a reasonable fashion.

Response to NBN Co consultation and industry collaboration

Noted.

---

[1] Occam's razor can be paraphrased in this context as aiming to achieve some measure of perfection in a solution by removing complexity and simplifying to such a point that it precisely and only solves the problem at hand.

[2] Einstein's razor can be paraphrased in this context as aiming to achieve some measure of perfection in a solution by simplifying but not over-simplifying; noting that some problems are actually better addressed by a slightly more complex, generalised approach. The challenge is recognising when, where and why to stop simplifying.

Section 2 Building a fibre access network

It would be very useful for industry to be able to engage NBNCo in examining some of the trade-offs that are being considered prior to deploying such a large and extensive network. In order to achieve this, some of the underlying assumptions and definitions for the fibre plant model will need to be published. Agree that these should include at least:

- A definition of a service delivery or service termination point (rough correlate of a SIO in Telstra terms). Once a definition has been provided an estimate, per state/region of the number of service termination points (including residential, business and other (traffic lights etc)).

- An assumption of the relative deployment approaches such as percentage of aerial fibre, ducted fibre, micro-ducted fibre; the degree of opportunity for re-use of existing duct and aerial assets; and the estimated split of network infrastructure between the home and the fibre distribution hubs versus the fibre distribution hubs and the fibre access nodes.

- A more detailed estimate of the location of each PoI and hence the size of each FSA, the total number of FSAs and hence PoIs.

- A more detailed estimate on where and when internal/external ONTs may be used and the relative amounts of each.

- A more detailed description of the proposed fibre redundancy plan; i.e. 1+1 fibre overlay versus a n+m overbuild, the latter obviously lending itself to quick repair, as well as flexibility for delivery of on-demand point-to-point services.

- A more detailed description of the proposed roll-out plan FSA by FSA. It is noted from other engagements that early application releases are planned prior to general rollout, however more detail around the timing/size of these trials and the intended mass rollout plan would be useful for industry to assist in planning to support the rollout and to utilise services once the rollout is complete.

- A decision around the number and type of physical and logical port interfaces available on the customer premise ONT (including Ethernet ports, POTS ports and RF ports etc).

Section 3 NBN Co’s overall product objectives

Objectives used by NBNCo to guide development of wholesale product offering noted and agreed.

Section 4 Choice of layer in the vertical technology stack

Generally agree with the conclusion that a layer-2 product represents the best compromise of simplicity and functionality to ensure “*vibrant retail competition*” in the market. It is important to highlight the following consideration however: basic Layer-3 IPv4 and IPv6 connectivity services (packet forwarding exclusively) are increasingly being seen as commodity services with limited scope for differentiation beyond security assurance, QoS stream differentiation and efficient content caching/management. A key objective for the NBNCo network is to ensure “*maximum end-user choice in terms of both services and providers*”. There is a fine trade-off happening here that warrants very careful consideration and management.

This may be best explained via small digression to illustrate by example. Local community services are a good illustrative model for assessing basic reasonable and fair access in this larger context[3]. Local community information, health, council, library, volunteer and small businesses (independent newsagencies, independent video rental stores, community radio/video/education etc) should be able to efficiently provide a wide variety of “application services” directly[4] to their local population without having to deploy and manage significant network infrastructure assets, without having to virtually franchise through a larger centralised provider and without having to locate any of their operations outside of their local community. In order to achieve this, it is important to ensure that at least 2 (preferably 3 or more) truly open, layer-3 dedicated, wholesale service providers exist that can provide simple and cost-effective solutions to support this market in a local manner (non-centralised). If not, then it may be necessary in the future to reconsider the scope of wholesale service offerings in a similar manner to the considerations being given to implementation of backhaul for aggregated PoIs (i.e. based on “contestable access” assessments).

Although the same argument also holds for L2TP aggregators, it seems clear that sufficient wholesale competition probably already exists in the market and will continue to exist for these legacy services.

---

[3] It is recognised that the larger context will see the lion’s share of almost all services being offered primarily through a relatively small number of national retail service providers who themselves will offer compelling service bundles. This should not mean that the smaller and niche market opportunities are left unconsidered. In fact they should be promoted. The community and economy can gain significantly from non-mainstream ‘long-tail economics’ as evidenced by the success of community generated content in countries such as the United States where community access opportunities have existed for some time.

[4] Note - this does not at all imply that these local entities would ever be a direct customer of NBNCo, rather it is foreseen that they would directly interface with a (potentially local only, but not necessarily) layer-3 service aggregator who themselves would most likely be considered the RSP customer of NBNCo. The important consideration here is to ensure that neither the RSP customer, nor its subsidiary local community customers, should be forced to own or manage assets outside of the local community for purely local service delivery.

Response to Detailed assessment summary

- **In general**: Agree on evaluation points and assessments with exceptions as listed below. Also, it would be useful to have more detail about the rating scheme (weakest to strongest) applied here and in other sections in the future.
- **Support for differentiation and innovation**: Disagree. With the exception of managing allocated address space, a layer-3 IP product can provide the equivalent differentiation capabilities of a layer-2 one. All the product definitions of throughput, performance, availability, security etc apply pretty much equally to a layer-3 wholesale service as they do to a layer-2 wholesale service.
- **Healthy retail competition**: Agree. However, note that the argument applied for avoiding a layer-1 monopolist, equally applies for higher layer services. It is vital to ensure that a layer-3 service duopoly/monopoly does not eventuate, especially one where conflicting retail/wholesale interests persist.

Response to Questions

It is important to note that all connectivity infrastructure exists solely to allow users and systems to communicate and exchange information via applications. Many companies profit from this across all aspects and levels of the service supply chain. Almost all of the end-user value however, is derived from these applications--be they entertainment services, commerce services, socialisation services or support services. Today the majority of these applications operate on Windows or Unix based systems (maybe OSX, but that is neither here nor there for this discussion) and assume IP connectivity (which is the crucial point). Much of the vibrant market that will emerge to leverage Australia’s NBN will be based around novel application services being delivered ubiquitously to end-users (this also includes automated business to business or system to system applications).

Some of these high level application service providers will themselves be quite sophisticated and will offer multiple, often bundled, application services and may even have extensive network infrastructure of their own. Other application service providers may not be so sophisticated and they will seek to be provided with an open and efficient conduit into the market via a layer-3 wholesale service provider. Offerings that will be sought are likely to be a mix of hosted applications (applet hosting and cloud computing services for instance), hosted servers, hosted platforms and possibly extranet connectivity (non-hosted). It is **not** important that these services be provided by NBNCo itself, however it is **crucial** that they be provided by an independent and competitive layer-3 wholesale service provider who preferably do not themselves have a internal conflict of interest through offering similar, competing applications. In order to maximise end-user benefit it will be necessary to **ensure** such competition evolves in the market. NBNCo are not responsible for solving this issue, however it should not be permitted to fall by the wayside and potentially be lost.

Section 5 High level technology standards

Response to Ethernet

Agree, Ethernet represents world’s best practice for layer-2 service delivery.

Response to GPON

Agree, GPON represents the most cost effective and world’s best practice for large scale FTTH deployment so long as some provision for business, government and other point-to-point services is made (as described in the paper). Given that next generation GEPON will be compatible with existing and next generation GPON and that GPON is currently more advanced in terms of being a deployed and tested technology, then GPON clearly represents the best choice.

Response to Questions

The key issues around the specified preferred technologies are:

- definitions of Ethernet connection tagging standards (including the c-tag and s-tag application of 802.1ad), especially at the PoI;
- definitions of the QoS marking standards (including specification of the user priority bits in the 802.1Q tag);
    - Particular care is required with the QoS specification as all aggregate class of service behaviours must be fully described using only 3 bits, giving a total of only 8 differentiated characteristics. Included within these differentiated characteristics, must be support for CIR (committed) and PIR (excess/peak) markings, hence typically only 4 such classes of aggregate forwarding behaviour can be supported (more detail on this can be found in the response to section 8);
- specification of the mechanisms for supporting end-to-end OAM&P to RSPs across the GPON/Ethernet infrastructure;
- specification of the maximum permitted frame sizes across the GPON/Ethernet infrastructure - in particular support for at least mini-jumbo frames up to and including 2000[5] bytes in size (as opposed to the default 1522 which will be insufficient for some services). Full jumbo frames are addressed just below.

In response to the question on triggers for choosing point to point connectivity as opposed to standard PON delivery, the following at least are foreseen:

- Specialised security environments such as civil or national defence, some banking/utility/mining services and national security to name a few. Any environment where there may be intention to deploy quantum cryptography will require dedicated point-to-point fibre services.
- Specialised protocol environments such as large data centres, data silos and data warehouses which may utilise proprietary layer-2 networking protocols (fibre-channel etc) or have jumbo frame requirements that can be considered unreasonable in a mass-deployed Ethernet environment (i.e. typically greater than 4096/8192 bytes).

[5] IEEE 802.3as update specifies recommended support for frames up to and including 2000 bytes.

- Specialised bandwidth environments where the dedicated bandwidth between the connection points is much larger than the typical mass deployment bandwidths being offered and is required to be dedicated and available.
- Some quantity of dedicated fibre will also need to be reserved per FSA to support future lawful interception requirements.

Section 6 Location of Points of Interconnect for NBN Co wholesale fibre network

Wholly support the re-use of the Communications Alliance NBN reference architecture for industry engagement.

Response to Implication of layer choice for geographic scope/PoI location

The Ofcom diagram that indicates the geographical reach of various wholesale products is accurate for layer-1 and layer-2, however layer-3 “IPStream” products could also be easily made available at Regional, District and even Local Exchange locations, albeit at an increased cost (so in this respect it has inaccuracies).

With respect to the statement, “*it is intended that only one PoI will be available for each FSA”, which is primarily justified in the text based on “Simplicity/Cost management: The establishment of only one PoI for any given geography ensures simplicity in the network design (along with associated cost savings)*”, this may be materially detrimental to the overall product objectives of supporting “*differentiation and innovation by RSPs--by providing them with an experience as close as possible to owning their own network*” and promoting “*maximum end-user choice in terms of both services and providers*”. The primary case in point here is our local community consisting of both end-users and local application service providers (community health, small business, local council, public library etc). In such areas if only a single PoI is supported, in this case via the AEB product offering (with embedded backhaul component), then the local retailers and any RSP aggregator who would wish to function locally will be significantly disadvantaged by being forced to locate service infrastructure non-locally across to the aggregation PoI.

Having a single PoI also further complicates matters if in the future “contestable” backhaul and access becomes available, as a mass migration of services will have to be undertaken, possibly involving changes to existing commercial arrangements. In contrast, natural market dynamics and competition are enabled by offering LEB and AEB PoI facilities for any given FSA where initially “contestable” backhaul access did not exist. In such scenarios, a RSP would be free to choose between the NBNCo AEB offering; or the LEB offering coupled with either a competitors backhaul service, leveraging their own backhaul capability or even sans[6] backhaul for purely local service offerings. Such flexibility would solve future migration issues (by simply rendering them mute), would naturally ensure ongoing competitiveness of the AEB service offering by NBNCo (due to market competition opportunities) and ensures maximum market flexibility, especially for small local service providers and aggregators.

Lastly, the power utilities mostly desire connectivity between just the local transformer and the home, having an extensive communications network capability already in place to that level (the transformers). In this case, a local PoI will always be the preferred hand-off point despite the availability (or lack thereof) of contested backhaul. Whilst this situation could be handled as an exception, such an exception could be challenged by other players. More discussion regarding utilities and smart metering can be found in the responses to section 9.

[6] English preposition meaning "without" (someone asked).

Response to Questions

On the basis of the discussion above, it is recommended that NBNCo reconsider offering LEB PoI services for every FSA and where “contestable” backhaul does not exist, additionally also offering the AEB PoI service to help foster open access and RSP competition.

As far as the other questions are concerned, a preference exists for a threshold criteria of “at least three, or soon to be at least three” independent backhaul providers before a local FSA PoI is considered to be sufficiently “contested”. Game theory demonstrates that two players can easily form a legal, non-colluding, duopoly that clearly reduces and even eliminates the benefits gained in a real competitive market. With three or more players the situation rapidly becomes much harder to “game”.

In areas where there are not sufficient backhaul providers to be considered “contested”, clearly an opportunity exists for NBNCo to solve it, fairly simply. If insufficient players are willing to actually construct the required backhaul at the same time as NBNCo are cabling the FSA, NBNCo can simply proceed to construct its own backhaul and then offer both local and aggregated PoI facilities. In the event a player wishes to enter the competitive backhaul market, commercial arrangements can be made to transfer the backhaul assets from NBNCo to the new entrant conditionally upon continuance of supply of the AEB service to existing customers. If the entrant does not wish to do this, then they can still enter the market themselves, however it would not automatically become a “contested” PoI. This is obviously not a full scenario analysis, however it is a reasonable mechanism for consideration. It also highlights how offering concurrent local and non-local (district) PoI can actually simplify some matters.

Section 7 NBN Co wholesale fibre bitstream products definition

Response to The product offering

Noted and agreed.

Response to Local Ethernet Bitstream product

Noted and agreed. Product PoI components line up with a Logical interface description (vlan tags, priority tags etc) and a Physical interface description (Ethernet 1Gps optical). No mention is made of a corresponding logical and physical interface at the customer premise ONT. The assumption is that RSPs should assume either/both a physical Ethernet interface and/or a logical vlan tagged interface will be available. Would it be possible for NBNCo to publish an estimate of the number of customers expected to be covered by the LEB versus the AEB products at this point in time?

Response to Aggregated Ethernet Bitstream product

Noted and agreed. With regard to the text "*contended access to FSAs*", is there a plan to be able to offer larger physical interfaces such as 10G to ensure that the access would be available uncontended (if required and commercially viable)?

Response to Product conditions

Noted. See additionally the response in Section 6 regarding the intention to not offer the LEB product in conjunction with the AEB product. It is difficult to understand why this limitation has been applied to the product set based on the information provided so far. The text "*So, in a particular area where contestable backhaul is not available, only making the AEB product available will ensure that there is no competitive advantage or disadvantage to any wholesale customer irrespective of whether that wholesale customer self-provides backhaul or acquires the backhaul from a third party*" is not obvious.
For instance, assume that a FSA PoI is considered uncontested (else we have no issue as the LEB will be available and at least two (preference would be three) independent backhaul providers exist).
We have only three possible scenarios to consider.

1) No market player has backhaul to the local FSA PoI. All market players are naturally required to seek the AEB offering at the district PoI and hence no player will have an unfair advantage or disadvantage. Note that a local application service provider will be significantly disadvantaged here if forced to seek the AEB product non-locally due to the lack of a local LEB product. Hence offering the **LEB and AEB concurrently** can only **positively** benefit the marketplace. Offering only the AEB product would act to disadvantage local business service providers who could otherwise form a consortium aggregated RSP for local service delivery to their local community. If a player later enters the market by constructing backhaul then scenario 2 comes into effect.

2) One (maybe two) market players have backhaul (self owned or leased) to the local FSA PoI. If NBNCo were to force all players to take only the AEB product at the district PoI then--the player with the existing backhaul is immediately (unfairly)

disadvantaged as they cannot utilise their network asset and they cannot enter a competitive market for backhaul as NBNCo now have a monopoly on that market. The remaining players in the market are forced to accept the NBNCo AEB product with the bundled transit link charge until such time as NBNCo deem the local FSA PoI to be “contested”. As a result, they potentially miss out on competitive pricing of services in the meantime and are faced with a potentially disruptive migration exercise when the PoI finally becomes “contested”.
Potentially all players and ultimately the consumers are disadvantaged by this approach. If the **LEB and AEB were offered concurrently** then much of this argument would disappear and would only **positively** benefit the marketplace. The existing incumbent player could asses the economic viability of their backhaul asset and choose to either leverage it themselves exclusively (no disadvantage to them, no advantage to others), they could choose to abandon the asset and leverage the NBNCo AEB the same as all other players (no player advantaged or disadvantaged unfairly) or they could choose to partition some of the capacity and compete in the backhaul market against NBNCo (all players are advantaged by this).

3) Later in time, three or more players may have backhaul capability to the local FSA PoI. This case is almost identical to scenario 2 with the exception all players can make the value assessments of how they want to leverage their backhaul assets and that NBNCo may choose to deem the PoI as now sufficiently contested and withdraw the AEB product from the district PoI. Again, If the **LEB and AEB are offered concurrently** then it ensures a **positive** benefit for the marketplace, **simplifies** the product management and service delivery of NBNCo, ensures absolute **fairness** to all players, provides a **natural** market-led migration strategy for decommissioning the AEB product in the future if necessary and removes much of the **complexity** that exists around deeming an FSA PoI as ‘initially uncontested’ given the current implications of such a decision.

Likewise, mention is made that “*NBN Co considers that offering the LEB and AEB on a mutually exclusive basis is likely to best meet NBN Co’s objective of ensuring a level competitive playing field and is less likely to result in competitive distortions in the marketplace*”--this also is not self-evident and requires more information around the analysis for consideration.

Section 8 Important product elements

Supporting material for this section can be found in the published article, "Which Takes Precedence: Your New NGN or Your Current Business Model"[7] (Michael S Cox, IEC Annual Review of Broadband Communications Vol 3, 2007[8]). This concludes that in order to offer effective tiered-pricing and prevent marketplace arbitrage, it is **necessary** to enforce aggregate shared queues per forwarding class. It is very encouraging to see that NBNCo are considering a tiered QoS capability as this will significantly open up the market opportunities for service and price differentiation/competition. Why this is highly desirable will be the primary focus of this response supported by material in the appendix.

Response to Bandwidth Rates

Noted and agreed. Definitions of PIR and CIR can be further clarified as follows. The primary difference between PIR and CIR behaviour is the service level guarantee around individual packet discard based on conformance to the policy profile. However, before a more robust definition can be specified some additional parameters are required. A Committed Burst Size (CBS) is necessary and a Maximum Burst Size (MBS) is desirable. Now, measured over a period of time (NBNCo to define, in operational networks however over intervals of ms and seconds they can only be reasonably estimated, whilst over intervals of minutes and hours they can be assured) the service level agreement would state something along the lines of:

> Within a bitstream service, packets that are below the committed information rate will be delivered with 100% reliability. Packets that are above the committed information rate but are within the committed burst size will also be delivered with 100% reliability[9]. Combined the CIR and CBS represent the minimum guaranteed "good-put" that a service provider can expect assuming active connectivity (another measure). Packets that are within the peak information rate and within the maximum burst size will be delivered as best effort. For some classes this means that under periodic congestion they may not be reliably delivered. Packets in excess of the peak information rate and in excess of the maximum burst size will be discarded immediately. A more technical description of these and related concepts (network intrinsic performance measures) can be found in the appendix.

It is appropriate to note here that telecommunications transmission capacity is a common, finite and contended resource, which is thus subject to the "tragedy of the commons"[10]. It is **finite** as at any given point in time there is only so much total capacity available (it may grow over time subject to appropriate financing but is fixed at any given point in time), it is **common** as it must be shared amongst all users trying to utilise capacity on shared trunks and it is **contended** when user demands reach a peak and equal/exceed the maximum available capacity. Without additional constraints to limit use (fair use policies, download caps etc), users will be incentivised to abuse the resource to depletion (no spare capacity). Or users will be expected to pay for the total maximum cost instead of an averaged one

---

[7] Article extract available on request.

[8] All articles available for purchase via http://www.iec.org/pubs/pub.asp?pid=102&bsi=2&cat=cont

[9] The CBS in this context is roughly equivalent and analogous to ATM's CDVT (Asynchronous Transfer Mode, Cell Delay Variation Tolerance) setting and provides a soft-threshold for the committed traffic enforcer to allow for natural packet clumping and phasing effects.

[10] Summary of the "Tragedy of the Commons" via http://en.wikipedia.org/wiki/Tragedy_of_the_commons

(long term). This is an issue that needs to be managed and like most it involves establishing appropriate compromises. The good news, there exists an optimal compromise that works in both the consumers and providers benefit.

With respect to the dimensioning rules for calculating capacity required at the PoI, at the very least the interconnect capacity must match the CIR and PIR characteristics of the largest access service configured at the PoI and additionally should be such that for each priority class of service offered, the CIR should be configured to at least be the sum of all access service CIRs configured at the POI (per class). The PIR should be configured as either the sum of all the PIRs configured at the POI, the physical interconnect interface speed or a lesser speed chosen by the RSP (whichever is the smaller). Whether there should be a financial incentive to the RSP for choosing a lesser speed is the subject of a very complex analysis model. In general, the answer is yes, but only up-to a point (based on a statistical sum of each PIR per class). RSPs (and end-users) should be incentivised to adopt and leverage peak bandwidth capabilities. It is recognised that some scope for allowing RSPs to choose their own CIR contention ratios should be allowed at the PoI, however the financial incentives for doing so need to be carefully modelled and managed in order to avoid an arbitrage situation (such as always choosing a zero bps CIR).

Some further detail on dimensioning at a PoI: a given RSP PoI connection would at a high level be described as follows:

Access Services: 802.1Q vlan or 802.3 Ethernet port at customer ONT. CIR per class of service (CoS), CBS = 2-5 frames and PIR per class , MBS = 2-3 frames (default) defined per CoS.

Physical Interface: IEEE 802.3-2008 Optical Ethernet (SX or LX tbd), PIR = 1Gbps (default), Maximum Frame Size = 1530 bytes (default)

Logical Interfaces: 802.1Q/p tags for CoS Priority marking, 802.1ad c-tag (per end-user) and s-tag (per fibre access node).

- CBS = SUM (all Access_Service CBS),
- PBS = SUM (all Access_Service PBS).
- CIR = SUM (all Access_Service CIR) or RSP choose down to NBNCo_Defined_MinCIR_SUM,
- PIR = SUM (all Access_Service PIR), Physical_Interface_PIR or RSP choose down to NBNCo_Defined_MinPIR_SUM.

Conveniently NBNCo should also enjoy some additional statistical gain (Erlang gain) beyond these straight summed CIRs.

Finally, it would help the industry if NBNCo could place some figures around the amount of time that they typically expect the network to be congested. This is often achieved using percentile performance specifications, however a detailed description/report of seasonal, diurnal and societal traffic patterns would be very useful and would provide a certain level of service transparency. It is noted that NBNCo’s ability to do this is entirely dependant on their decision on whether to control pricing, cost and/or performance. Only two of these parameters can be fully controlled at any time, the third then becomes a complex function of market take-up and demand (which the best models today can only vaguely estimate). NBNCo will need to identify which two parameters it intends to manage and hence which parameter will subsequently be market determined. Please see appendix for further information on this.

Response to Traffic Management & Prioritisation

This section appears to be an interim specification of an otherwise advanced, yet still clearly developing, QoS architecture. Whilst it appears on the surface to contain all the key elements required for optimal market competition and fairness, the statements form an unclear and partially incomplete picture. For the sake of brevity, this response will attempt to describe an aligned approach for delivering multi-QoS along with supporting arguments in the appendix as to why it is desirable to construct it in this manner.

Considerations:

- QoS markings will utilise 802.1Q/p priority markings. These consist of 3 bits, providing a limited capability of describing, at most, 8 differential forwarding characteristics. As a result more abstract forwarding behaviours will need to be classified as opposed to individual application or service constructs.

- There is a need to support at least the following:
    - A Class of Service (CoS) suitable for carrying priority, **real-time** and **interactive** traffic such as voice, video conferencing services, online gaming etc;
    - a CoS for large scale **streaming** of video, streaming of audio, reliable bulk and **transactional** messaging/transfer, network storage etc;
    - a CoS for basic **economy** access services (non-premium internet and background services);
    - and most likely a CoS that sits somewhere in between the two best services and the cheapest service which would provide a **business** grade, **premium** multi-use capability.

Additional classes of service can be supported, however the abstract behaviour requirements of the service should be carefully considered for mapping into one of these four before assigning a new ID to them from the limited 802.1Q/p address space.

- There is a need to be able to differentiate between committed carriage and best-effort carriage - almost on a per-class basis (in-profile (within CIR/CBS) and out-of-profile (up to PIR/MBS)).

- Other specialised “Service Classes” can always exist outside of the aggregate CoS specification and should be identified (classified) and handled directly by the network infrastructure in an appropriate manner. Emergency, management and other forms of critical applications, true background services and active performance management services should be supported in this manner, independent of customer CoS streams.

-All bitstream services include a common specification of CIR/CBS and PIR/MBS, a common specification of availability and reliability, a common specification of minimum-latency (directly related to transmission distance).

- Individual CoS bitstreams are further differentiated from one another by specification of a combination of per class in-profile and out-of-profile packet loss probabilities, packet-latency-variation (above the common minimum-latency - also called jitter), pricing per-committed-megabit-per-second and/or network capacity dimensioning applied[11].

[11] Note, not all of these parameters can be absolutely specified and controlled as dynamic network utilisation and demand (mostly driven by end-users) will ultimately affect one or more of these characteristics. In other words, if you attempt to fix the performance characteristics and fix the selling prices, then the network dimensioning will have to be adjusted based on service take-up to appropriately match. Vice-versa if you attempt to fix the network dimensioning and fix the performance characteristics, the only way to ensure that the SLAs are maintained is to be able to dynamically vary the selling price in order to manipulate service demand per class. And likewise, if you attempt to fix the selling price and fix the network dimension (fix costs) then the service performance characteristics will dynamically respond according to market take-up and utilisation (this is probably the recommended/preferred approach for NBNCo and is similar to Paris Metro Pricing, Airline Seat Pricing and many other tiered pricing strategies).

Recommendations:

- Use the 8 available coding patterns to define 4 network CoS behaviours, each of which supports in-profile and out-of-profile packet marking and differentiation. These behaviours should be:

    i. **Expedited** - a service class where packets receive queue servicing priority above all other classes. In-profile packets are absolutely guaranteed timely delivery. Out-of-profile packets are carried as best-effort but still enjoy priority servicing. Network capacity is dimensioned with little to no overbooking to ensure packet loss is negligible and latency-variation is negligible.

    ii. **Assured** - a service class where packets receive a high level of delivery assurance due to the judicious deployment of network capacity and overflow buffering. In-profile packets are absolutely guaranteed delivery, Out-of-profile packets are carried as best-effort, Network buffering and capacity is dimensioned with little overbooking to ensure as reliable a delivery as possible however some accommodation for latency-variation by applications will be required as this is not an expedited service.

    iii. **Premium** - a service class where packets receive overall high performance characteristics by offering a lower level of latency-variation to the Assured class (by trading-off some burst discard-ability) without incurring the high costs associated with the absolute dimensioning of the Expedited class. In-profile packets are absolutely assured of delivery and out-of-profile packets are also assured a reasonably high probability of delivery (>98%) most of the time (unlike any of the other classes where out-of-profile is carried as best-effort)[12].

    iv. **Standard** - a service class where packets are able to utilise excess bandwidth when available but are forced to share contended bandwidth during peak periods. In-profile packets are absolutely assured delivery whilst out-of-profile packets are carried as best-effort with no assurances. Due to the shared contention nature of this class, latency-variation can oscillate throughout the day in response to aggregate utilisation. This is similar to the dimensioning and management of many Internet access services today (excluding the premium and cap-limited services).

- NBNCo should publish fixed network dimensioning and overbooking characteristics, along with fixed marginal cost based[13] pricing structures (as agreed/negotiated with the ACCC) per mbps (upstream and downstream independently) for the Expedited, Assured and Standard classes. Note the actual performance characteristics of out-of-profile packets within these classes cannot be specified as per earlier discussion. However, base latency and latency-variation can be bounded for each of these classes (based on chosen active components implementation) and in-profile packets are always absolutely assured of delivery (minimum good-put). It is highly desirable to exercise some degree of cross-subsidy between the Standard class and the other higher-performance classes as this will allow NBNCo to increase the availability and uptake of broadband services nationwide by pricing the Standard class potentially well-below marginal cost. This is the social-benefit of such a scheme. Everyone can afford to play.

- NBNCo can and should publish defined network performance characteristics for in-profile and out-of-profile packets as part of the Premium class, however if the ACCC wishes to fix the pricing then NBNCo will need to be aware that network dimensioning

[12] There are well established rules for dimensioning such networks (premium/business grade) which accommodate fractal traffic patterns. These typically assume a normal distribution of instantaneous traffic rates around measured average rates, however these assumptions are easily verified by periodic monitoring of active networks.

[13] This is merely a likely approach that may be undertaken to negotiating a satisfactory regulatory outcome for managing the natural monopoly of NBNCo. In general any of cost of service regulation, price cap regulation, performance based regulation, yardstick regulation or franchise regulation etc could be used by the ACCC.

costs can be neither controlled nor estimated, likewise if costs are to be controlled then the only way to ensure ongoing delivery performance is to allow NBNCo to control/establish service prices and hence manipulate demand. The recommendation here is to allow NBNCo to flexibly define both pricing and manage network dimensioning. Natural market dynamics will ensure that NBNCo are pricing competitively as well as non-destructively as other players in the business service market will exist (to compete on price and performance, albeit with less reach) and the fixed price nature of the other classes means they can always be utilised by RSPs if desired. If NBNCo were to price this service too aggressively their costs of maintaining the chosen price points and performance characteristics is likely to prove unsustainable. This ensures a self correcting market.

- Allow RSPs to choose the amount of CIR per class per ONT access service both upstream and downstream independently. NBNCo can then specify the appropriate CBS, PIR and MBS for each class. Each unit mbps for each class should have its own pricing structure and RSPs should then be able to freely choose which classes and associated access prices are suitable for their target application and end-user segments. RSPs can target single or multiple market segments at price points that are not just tolerable but more importantly are appealing. This allows the entire market to operate in an optimal manner. See appendix for more details.
- Finally, allow RSPs the ability to choose to contend their own CIR and PIR aggregates at each PoI by purchasing less capacity than the simple summations of all end-user access services would indicate. There should be some financial incentive available to the RSP for choosing to do this, however the actual incentive amounts will need to be carefully managed to ensure ongoing market stability, viability and fairness. This represents at least the concurrency or Erlang gain advantage for the RSPs.

Response to Security and VLANs

Noted and agreed.

This is an appropriate point to raise an opportunity that exists to ensure the security of the ONTs are delivered in a possibly optimal manner. From the “Realising our Broadband Future” engagement held in December 2009, it was clear that the utility industries (power, gas, water, sewer) were all exceptionally keen to leverage the NBN coverage and capabilities to enable them to deploy real smart grid technologies. They were also keen to not have to pay (or at least pay very little) for these connectivity services. This represents a fantastic opportunity for NBNCo. In exchange for providing a high quality yet low bitrate connection to every household in Australia, the utilities in turn (as a consortium) could provide NBNCo with a reliable trickle of power for the ONT (for of course, no cost). Not only would this provide every ONT with reliable power and every utility with reliable smart meter connectivity, it also provides the additional benefit of localised ‘earthing’. The obvious choice for enabling the provision of this power would be for the utilities smart meter to actually provide 802.3af PoE to the ONT via its Ethernet connection. This would be a highly desirable outcome with benefits that extend well beyond just the utilities and NBNCo (into the community at large). Battery backup strategies are discussed later in this section.

Response to Voice Option

Noted, will await the further engagement.

In the meantime, it would be desirable that the SIP based POTS service support the following service capabilities and characteristics:

- Ability for the end-user to have a reasonable choice of possible POTS providers as opposed to a forced acceptance of a default provider.
- Basic POTS will need to be offered at a very competitive wholesale price. World benchmark pricing is rapidly approaching AUD$10pcm.
- That the SIP client on the ONT being managed and configured by the chosen POTS provider in as direct a manner as is possible (preferably via a vlan tagged interface). The relay or proxy of SIP messages and data streams should be avoided if possible.
- Ensure that basic self provisioning and most importantly EMERGENCY services are available via the POTS port regardless of any other service activations or subscriptions (or lack thereof).

Response to Multicast

Noted, will await the further engagement.

Response to Optical Network Termination

Some comments regarding the listed options under consideration:

- Recommend 4 total Ethernet ports.
    - 1 connected to utilities smart meter which receives power over Ethernet
    - 2 ports to support dedicated RSP service delivery (untagged)
    - 1 port to support shared RSP service delivery (VLAN tagged)
- Support a single POTS interface. Additional and optional POTS services can be delivered via separate ATA.
- If PoE is to be supplied over ports 2-4 then power should be derived from external power supply (not fed from the power received over port 1), this power should also probably not be supplied from battery backup if utilised. Only emergency POTS needs be maintained during power outages. This ensures greatest survivability of service. See more on battery backup below.
- RF port--It's easy to do but it would maintain a legacy. Thankfully there are only two large RF infrastructure owners in Australia and both of them are vey likely to be NBNCo's first and largest RSPs. Maybe an economically sensible compromise can be negotiated? It will be interesting to see what the outcome is in later engagements.

Battery backup is somewhat more involved. It is highly desirable for NBNCo to be able to offer an equivalent 'life-line' style service capability to what the Cu based network can today. In order to do this, some form of continuous and reliable power backup is required. It is well understood that even the best power storage technologies available today have limited life spans and have significant environmental impacts. A reasonable compromise for managing this dilemma is:

1. Do not make battery backup a mandatory component of the service.
    a. Increasingly households are becoming less reliant on fixed line voice communications and increasingly have access to mobile handsets for basic voice communication. In the event of an emergency where local power is disrupted, there still exists a reasonable level of availability for the mobile network. This isn't where it ends though - there is more…

2. Do make battery backup an end-user selectable component of the service.
   a. End-users should be actively queried as to whether or not they require the service and appropriately briefed as to what actions they may need to take to ensure the service is delivered. These actions may involve a cost component. This will help to minimise the environmental impacts. This decision should not be forced at the time of installation and should be able to be changed easily at any time after.
3. If battery backup is required for ‘life line’ POTS then in order to minimise the service costs, the battery packs should be user-serviceable (minimise truck rolls).
   a. This means that the installation and removal of a battery pack should be straightforward, idiot-proof and preferably integrated into the design of the ONT (so that it becomes a simple matter of sliding it into place). Obviously this compartment in the ONT will need to be appropriately engineered to ensure ease of removal and insertion as well as offer appropriate environmental intrusion protection against water and dust. In the rare case that a user cannot service the battery themselves, a third party support service can be used. Such an approach may also offer a choice of battery packs and allows for improvements in battery technology over time.
   b. Battery presence and performance can be actively monitored by NBNCo and when batteries need to be replaced they can simply be mailed out along with instructions for removal and return of the old battery and installation of the new one. Additional instructions for calling a local service agent who can perform the removal and install for a fee can also be included.
4. If extensive/additional battery backup is required and/or protection of smart meter and other data services is also required then larger UPS style units can be deployed that interface via the external power port.
5. If security of batteries is a concern with externally mounted ONTs, then apart from the active monitoring of the battery state, batteries could be equipped with identification technology. Using physical security (lock and keys) is probably unnecessary and introduces other complications like key management.

Response to Questions

A few remaining answers, needed to address questions not answered sufficiently by the above response:

Are there any other technical parameters that should be included?

Yes. Definitions and specifications for network reliability, network availability, measures of service connectivity and the intended model for providing service difficulties and faults support. Options include NBNCo providing initial point of contact for all end-user difficulties, escalating to the appropriate RSP after initial fault diagnosis or …

How to provide SPs with the ability to confirm connectivity and power?

Such interfaces should be implemented alongside the OSS ordering and provisioning interfaces. These interfaces should be fully automated wherever possible and would implement a number of service diagnostic and analysis capabilities. A customised feed of network alarms, correlated/filtered and delivered on a per SP basis, should also be generated by NBNCo wherever possible. Automated xml-rpc style (web services) interfaces should be defined and would permit an SP to query the status of any of their active service in the following manner:

Level 1 Diagnostic: Report NBNCo internal OSS view of the service status (active, inactive, fault …)

Level 2 Diagnostic: NBNCo OSS actively queries network state to report service status (active, inactive, fault …) (This would confirm power and possibly activity)
Level 3 Diagnostic: NBNCo OSS actively queries network service configuration and information - reports compliance of actual config against intended config. (This would confirm accuracy of service configuration)
Level 4 Diagnostic: NBNCo OSS actively tests basic connectivity and activity response of the service. This would absolutely confirm power and activity.
Level 5 Diagnostic: *This test is active and can affect applications operating across the service* NBNCo OSS actively evaluates the performance characteristics of the target service and reports.

Well, something along these lines will need to be defined. Such interfaces can be secure and efficient.

Should battery backup capabilities … be offered…?

In addition to the responses concerning battery backup already provided:

- The basic battery backup option for POTS protection should be paid for by the end-user, most likely via a small monthly charge through their chosen POTS Provider.
- An extended version of battery backup that would also protect data services should definitely be paid for by the end-user and would probably be easiest done through a third-party service agent or via NBNCo. As undesirable as that may sound, it is difficult to imagine alternative scenarios which may work fairly and effectively.
- The ability to opt-in or opt-out of either battery backup service should be completely flexible and available at any time. Australia’s residential population is too itinerate to have it being a once-off, at-install only decision.

Section 9 Conclusion and next steps

Overall response to Consultation program

By providing ubiquitous, high-speed broadband to the entire population, the NBN initiative places Australia at the forefront of communications capability worldwide. This will open up many new, certainly some unforseen, opportunities for improved communication, entertainment and collaboration between individuals, businesses and the government itself. Many of these applications will not only help to improve the quality of life for all Australians, it will also help the country to achieve ambitious goals regarding greenhouse gas emission reductions, improving health and educational services whilst reducing costs and creating significant new ICT industry opportunities. Being one of the first countries in the world to make such a network available to all of its citizens, we are also well positioned to then export the inevitable intellectual property we will learn from the construction, operations and subsequent innovative utilisation of our NBN.

It is greatly appreciated that NBNCo is allowing for open and constructive industry engagement as part of the early planning and designing phases. It would also be useful if NBNCo could continue this engagement as it develops the service pricing and bundling models. We want to get this right.

Response to Next Steps

Thankyou for permitting this opportunity to engage and look forward to being able to continue engagements with NBNCo as the network design and architecture evolves.

Appendix A Addendum Material

The contents of this appendix consist of several summary topics, none of which should be considered exhaustive, nor complete. They are included here in order to provide further detail and in some cases justify material included within the responses to each section.

Affects of Price and Performance versus Costs and Demand:

Performance measures fall into a number of broad categories. Some useful categories can be summarised as:

- User perceived performance. This is really a measure of how happy the user is with the service performance as a whole (sometimes referred to as quality of experience). It is based on a value assessment of the service that trades off the price paid for the service versus the utility derived from the service.
  User_Perceived_Performance = function (price, utility)
- Application assessed performance. This is a measure of how performant the application or service is assessed/deemed to be. It can be measured either internally by the service itself or via external approaches (MOS and r-scores for voice are examples of this) and is primarily affected by platform and system performance characteristics (Platform intrinsic measures). It is important to note that even though an application may be assessed as performing poorly (mobile telephony calls for instance), users may still perceive significant value (due to high utility of mobility).
  Application_Assessed_Performance = function (platforms and systems)
- Platform intrinsic performance. These are measures of internal, system specific, parameters that affect performance or are affected by load. Packet delay and packet loss across a network element are examples of intrinsic performance measures.
  Platform_Intrinsic_Performance = function (internal parameters)

Now it is quite clear that complex functions can be used to derive levels of platform, system and application service performance, however these do not translate easily into user-perceived performance measures. The best measure of user-perceived performance is via satisfaction surveys and service take-up monitoring. This usually needs to be done within individual market segments as: what one user group may perceive as representing an overall positive service experience, another could easily consider to be a negative service experience. Tiered pricing is one mechanism for addressing this imbalance of perception across a heterogeneous population.

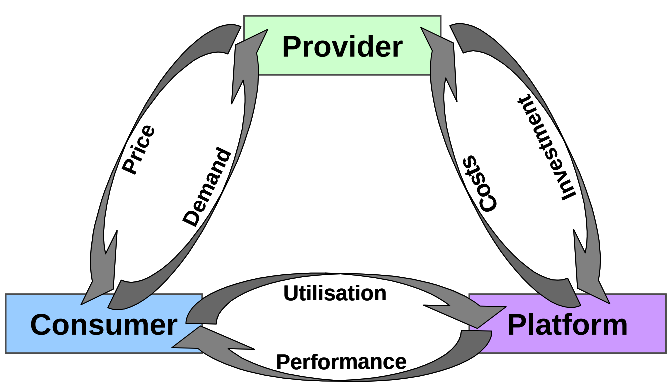

The figure above illustrates these relationships. From this, it becomes clear that prices and utility have the potential to drive demand for services (positive experience) or withdraw demand for services (negative experience). These behaviours in turn can vary the load presented to the platforms and systems involved in delivering the services. In response to this utilisation, the service provider can choose to increase or decrease capacity which in turn will affect the performance of the service. The service provider could alternatively choose to adjust the price of the service, either increasing it (and thus dropping demand) or decreasing them (thus increasing demand). It is clear that these feedbacks quickly become complex and recursive, each affecting the other as illustrated above. Eventually a marketplace should find states of equilibrium where the right sort of price, for the right sort of investment for the right sort of performance all balances out. In healthy markets, competition ensures this happens efficiently. In monopoly markets other unfriendly outcomes can occur unless actively prevented. I am confident that in NBNCo’s case, the outcome on this matter will be a very positive one, as it is being actively managed right from the outset and in a constructive fashion.

Although it is possible to constrain and manage some of these parameters, it is not possible to control them all, due to the interdependencies inherent within the feedback loops.

Network Intrinsic Performance Metrics:

A summary of some typical and example IETF derived network intrinsic performance measures.

- **One-way minimum transfer delay** (OWMD) is roughly the measure of the minimum elapsed time it takes for a packet to transit a given network path or subpath. OWMD can be derived from singleton metrics such as Type-P-One-Way-Delay-Minimum as defined in RFC2679 “A One-way Delay Metric for IPPM”, or from alternative approaches like calculating the minimum of a set of valid Delay metrics, such as defined in RFC3432 “Network Performance Measurement with periodic Streams”. This is a minimum estimate.
- **Round-trip transfer delay** (RTTD) is roughly the measure of the elapsed time it takes for a packet to transit a given network path or subpath from a point of origin to a far-end point of reflection whereby it is immediately returned back to the point of origin. RTTD can be derived from singleton metrics such as Type-P-Round-trip-Delay-Median and the Type-P-Round-trip-Delay-Percentile(50) as defined in RFC2681 “A Round-trip Delay Metric for IPPM” or as a linear combination of two simultaneous OWTD measures, one from the point of origin to the point of reflection and the other from the point of reflection to the same point of origin. This is a median estimate.
- **Packet Delay Variation** (PDV), often ambiguously referred to as “jitter”, is a measure of the mean variation in OWTD of a sequence or sample of packets. Unfortunately, there are several different methodologies from several different standards groups for calculating delay variation. For these purposes, PDV can to be derived from singleton metrics such as Type-P-One-way-ipdv-Percentile(50) as defined in RFC3393 “IP Packet Delay Variation” and using (P(min), P(i)) as the preferred selection function when calculating; or from alternatives such as calculating the median or mean of the absolute differences between each singleton Delay and the AveDelay, for a set of valid Delay metrics, as defined in RFC3432 “Network Performance Measurement with periodic Streams”. This is a median estimate.
- **Peak Packet Delay Variation** (PPDV) is a measure of the maximum variation in OWTD of a sequence or sample of packets. Unfortunately, there are several different and conflicting methodologies from several different standards groups for calculating delay variation metrics. For these purposes, PPDV can be derived from singleton metrics such as Type-P-One-way-peak-to-peak-ipdv as defined in RFC3393 “IP Packet Delay Variation” and instead of using the selection function (P(max),P(min)) it is recommended that (P(95th),P(min)) be used; or from alternatives such as calculating the range of the 95^{th} percentile of Delay and the minimum Delay as defined in RFC3432 “Network Performance Measurement with periodic Streams. This is a normalised maximum estimate.
- **One-way Packet Loss** (OWPL) is a measure of the relative number of packets that fail to be delivered correctly versus the total number of packets originally transmitted. There are several possible causes of packet loss, including but not limited to, packet

discard due to congestion management (active or passive), packet corruption due to transmission errors, packet loss due to insufficient processing resources and finally packet re-ordering due to network changes or operation. OWPL can to be derived from singleton metrics such as Type-P-One-way-Packet-Loss-Average as defined in RFC2680 “One Way Packet Loss Metric for IPPM” or from alternatives such as calculating the ratio of values with infinite or unspecified Delay versus the total number of values in a set of Delay metrics as defined in RFC3432 “Network Performance Measurement with periodic Streams”.

- **Connectivity** (CON) is an end-to-end measure of the ability for two hosts to communicate with each other across a cloud regardless of underlying performance. It is measured as the relative amount of time that two hosts are considered connected relative to the total duration of time elapsed, whilst connectivity was being tested. CON can be derived from singleton metrics such as Type-P-Instantaneous-Unidirectional-Connectivity, Type-P-Instantaneous-Bidirectional-Connectivity, Type-P-Interval-Unidirectional-Connectivity and Type-P-Interval-Bidirectional-Connectivity which can be used in combination to derive Type-P1-P2-Interval-Temporal-Connectivity as defined in RFC2678 “IPPM Metrics for Measuring Connectivity”.

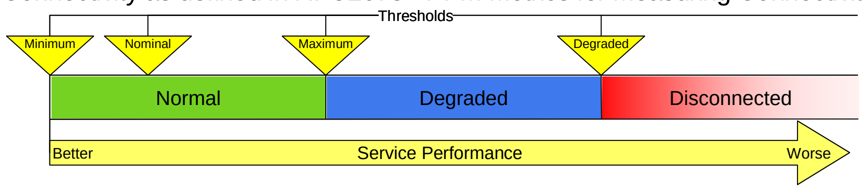

- **Throughput** (TPUT) is roughly a measure of a systems ability to transfer a continuous stream of data with varying temporal properties. It is usually described as a combination of both the sustained (minimum) and instantaneous (maximum) rate of transfer supported across a path or subpath, between two hosts without inducing excessive latency variation or loss. TPUT often provides two clearly differentiated latency and loss behaviours based on the different rates, the first is a guaranteed (committed) throughput specification and the second is a best efforts (burst) throughput specification. TPUT can to be specified and measured using various definitions within RFC3148 “A Framework for Defining Empirical Bulk Transfer Capacity Metrics”, RFC2697 “A Single Rate Three Color Marker”, RFC2698 “A Two Rate Three Color marker”, RFC2963 “A Rate Adaptive Shaper for Differentiated Services”, RFC2215 “General Characterisation Parameters for Integrated Service Network Elements”, RFC2216 “Network Element Service Specification Template”, RFC2211 “Specification of the Controlled-Load Network Element Service”, RFC2212 “Specification of Guaranteed Quality of Service” and numerous others which may include equivalent definitions. The preferred scheme would be to adopt the specification of the trTCM described as a 4-tuple consisting of the Committed Information Rate (CIR), the Committed Burst Size (CBS), the Peak Information Rate (PIR) and the Maximum Burst Size (MBS).
- **Packet Re-ordering** (PRO) is roughly a measure of the degree to which packets of a single microflow arrive in an order or sequence that differs from the sequence in which they were originally transmitted. Unfortunately the IPPM related specification for quantifying packet re-ordering is a draft work-in-progress and for this reason, there is no specification for limits on the PRO metric as it is currently undefined. Since packet re-ordering is generally considered an undesirable behaviour—designers and implementers should adopt approaches that minimise, or eliminate altogether, the possibility of packet re-ordering.
- **Availability**, or more accurately, **unavailability** could be based on ITU-T Recommendation G.826 “End-to-end error performance parameters and objectives for international, constant bit-rate digital paths and connections”. Availability (AVAIL) is thus defined as the availability of an item (Any part, device, subsystem, functional unit, equipment or system that can be individually considered) to be in a state to perform a required function at a given instant of time or at any instant of time within a given time

interval, assuming that the external resources, if required, are provided. The numerical value of availability is expressed as a percentage from 0 — 100% (representing a probability ranging from 0 — 1). Availability calculations take into account both the failure rates and the repair times of the system.

Diurnal, Societal and Seasonal traffic patterns:

Everything about society is fundamentally driven by patterns. These patterns fundamentally derive first, from the natural cycles of earths own rotation around its axis (called a diurnal cycle, it is the daily period of 24hrs) and it drives wake/sleep patterns; second, from the societal cycle of a working week (called a weekly cycle funnily enough and consisting of 7 days) which drives our work/leisure patterns; and third, from the natural cycle of the earths rotation around the sun (called a year and consisting of roughly 365.25 days) which drive our holiday breaks, mostly through the changes of the seasons. These patterns show up clearly and predictably in almost every human system - including telecommunications (unless the system is already heavily congested). These patterns are nearly 100% responsible for determining and causing the peak-period congestion we experience in so many areas. Only highly correlated, anomalous events, such as 9-11 (the terrorist disaster that occurred in the United States) can make an impact more significant than these traditional cycles. Examples of these cycles are illustrated below. This sort of demand and utilisation information is vital for capacity planning and useful for understanding performance variance. Note the illustration below has been detrended of any growth data, in reality these graphs exhibit the same patterns, just with a trended growth slope.

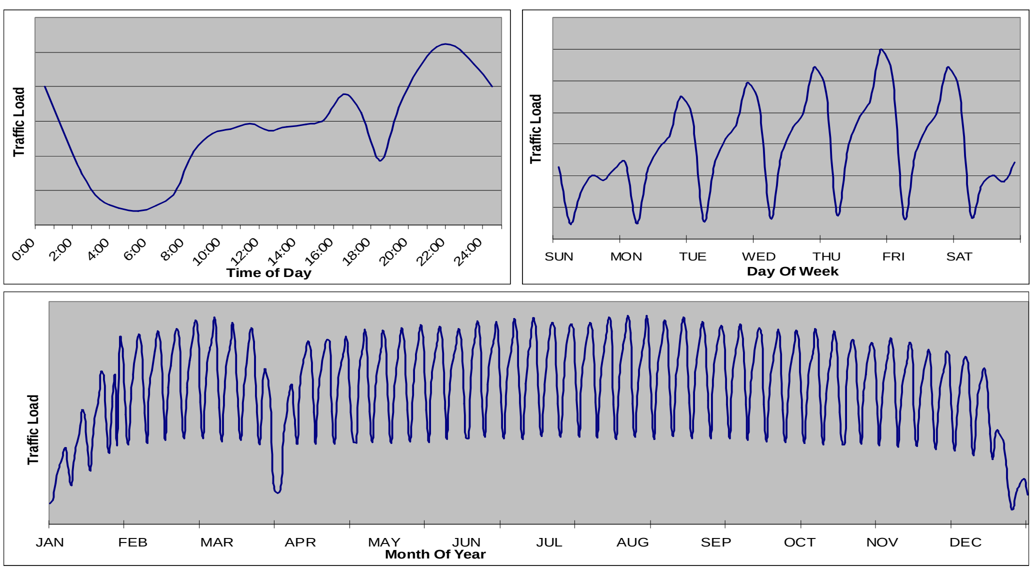

Although the busy hour load itself may be dominated by either the business market or the consumer market segments in any given service area, the actual cycle itself is usually determined by the average working day and school day cycles. This is due to the fact that both business and residential use of data services is directly correlated with each individual’s activities throughout the day and the general population’s activities are dominated by the majority cycle of individuals going to, and returning from, the workplace and associated morning, afternoon and evening. As a result, the busy hour is almost guaranteed to occur on a typical working day and conversely rarely occurs on weekends, holidays or other atypical workdays. This dependency on daily and weekly cycles, introduces some unique requirements with respect to managing network performance specifications or capacity planning estimations that are based on busy hour periods. Amongst these is the observation that in any given working week, one day typically exhibits a resource utilisation level that is significantly larger than the other days. Major holiday periods can also have significant, albeit temporary, effects on the busy hour load that occurs. These situations correspond to relevant major holiday periods (such as Easter, Thanksgiving in North America, New Years etcetera) and can even extend up to 56 days during extended holiday breaks (July Summer Holidays in the Northern Hemisphere and the combined Christmas plus New Years in the Southern Hemisphere as applicable). Even individual holidays can have a momentary, yet significant impact on the busy hour load, Mothers day for voice telephony is a classic example. Lastly, there also exist seasonal variations that can occur (Autumn, Winter, Spring, Summer), however these tend to align well with the extended holiday periods that are already associated with each hemisphere.

Market advantages derived from tiered pricing (abridged from article):

All common carriers and in fact all natural monopolies are faced with a common set of problems:

- Their infrastructure footprint often provides close to 100% coverage (as a result they enjoy economies of scale but also suffer a large, fixed base-cost overhead. In order to deal with this, they must achieve significant market penetration for their services).
- Regulatory constraints usually mean that their infrastructure must simultaneously meet the needs of all consumer levels from the lowest to highest performance demands and lowest to highest income brackets.

The only choice they have is to either utilise a fixed averaged cost structure or leverage a market based pricing strategy. If they choose the former then it is necessary for a competitive service aggregator to enter the market to offer the capability. This is not always desirable as it can add significant additional costs to the end-user services.

The latter is best illustrated using a common economic analysis. The primary purpose of market based pricing is to capture the market’s consumer surplus, although it can also include support for social welfare through increased output (this is the case with NBNCo). In a market with a single clearing price, some customers (usually the ones with low price elasticity) would be willing to pay more than the single market price, whilst others (the social welfare group) are not willing or possibly able to pay the single market price. Market pricing introduces the flexibility necessary to independently target each market segment at an appropriate price threshold. The diagram below illustrates the various effects of being able to tier pricing to various groups. In the single service price case, the addressable market for price (P) is represented by the corresponding demand for services (Q). The revenue opportunity is represented roughly by the area of the square. In the tiered service price case, individual prices P_1, P_2 and P_3 are offered (preferably in a manner that appeals to each market segment) and subsequently demand for services as represented by Q_1, Q_2 and Q_3 respectively, increases. More importantly, not only does the total addressable market increase, output increases and revenues increase. It is entirely possible that under the right set of circumstances and appropriate management, that the entire market will be able to afford some form of service (full social welfare support) without compromising the long-term viability of the producer.

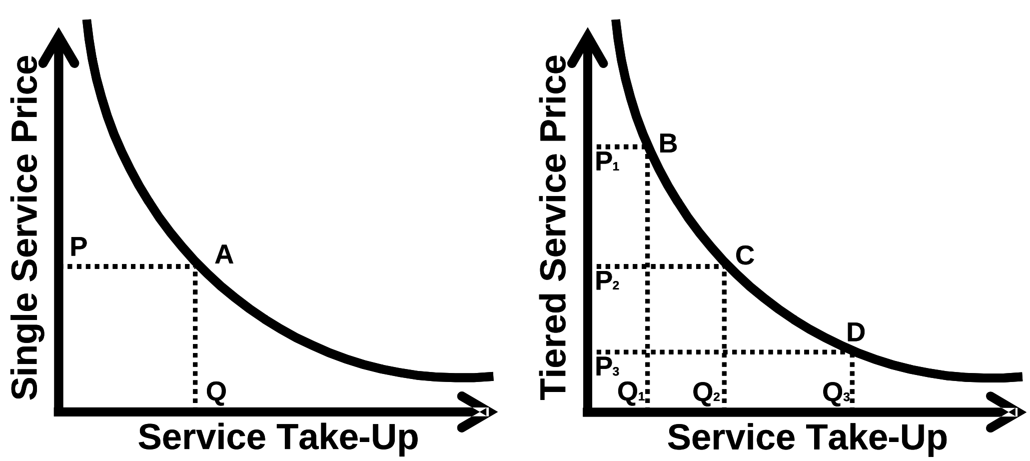

By choosing to leverage base infrastructure capabilities such as differentiated performance, a carrier can price differentiated services into the market. This allows the first level of consumers, the retail service providers and or aggregators to themselves select which entry points they desire and which market segments they are going to differentiate them into. Hence a rich portfolio of aggregators will exist that can leverage the differentiated capabilities and differentiated prices to develop market based value opportunities. This increases the uptake, increases margins, increases total market efficiency and reduces costs to consumers that have high price elasticity. This approach emulates to some degree, the traditional approach adopted by incumbent Telco’s who would normally have value-based price each legacy overlay network individually.

For the record, tiered pricing is very common across multiple industries, as the following list of examples illustrates:

- Volume based wholesale price discrimination where larger bulk quantity purchases are offered at a reduced per unit price (many wholesalers and component retailers). This could also include peak and off-peak pricing (utilities) etc.
- Premium product differentiation, whereby a similar product is marketed as premium and sold at a higher price (supermarkets and grocery stores). Other examples include the automobile industry (limited edition paint releases) and the airline industry (weekend away conditions, economy, business, first …) etc.
- Segmentation pricing whereby children, students and pensioners are offered reduced prices when compared with working adults (public transport, entertainment etc). Other examples include segmentation by gender (venue cover charges), occupation (education, volunteer, police force etc) and geography (local, interstate, international).

Without such tiered pricing strategies, many products and services may not be as affordable as they currently are for many people.