Response to Request for Consultation on the National Broadband Network

Contents

Tables

Figures

Copyright Notice

This paper is copyrighted by the author with the following exclusion:

http://www.dbcde.gov.au/communications_for_business/funding_programs__and__support/national_broadband_network/consultation

"By lodging a submission interested parties will be consenting to the Commonwealth publishing the submission and making use of the submission as it may see fit in the development of the Request for Proposals documentation."

All other rights remain with the author.

Stephen Baxter
CTO PIPE Networks
steve@pipenetworks.com

Introduction

About PIPE

PIPE Networks (PIPE) is an Australian licensed carrier (#101) started in December 2001. PIPE owns and operates a network of peering and Internet interconnection points, the third largest metro dark fibre network in Australia as well as providing other services such as Ethernet and colocation.

PIPE recently announced its plans to build a submarine cable between Sydney and Guam with this project due for completion by July 2009.

PIPE is publicly listed on the ASX.

Approach to consultation

The guidance for this response was what was effectively contained at :

http://www.dbcde.gov.au/communications_for_business/funding_programs__and__support/national_broadband_network/consultation

Given that we had two weeks (over Easter) to respond to the no guidelines put out by the department, if there is to be a second iteration of this process before tender it would useful to know (without limitation) some or all of the following:

- what current models are being considered for FTTN from the technical and competition aspects;
- what timeframes for 98% coverage for any network build;
- scope of what the FTTN will cover – access tails to users, backhaul, creating a bundled offering etc;
- what price goals for end user services are being targeted; and
- what applications are expected to be used over the network;

This would greatly assist in drafting a response.

There is a lot more data that could be included if there was more time to collect it.

Data realities

Broadband is about the delivery of high speed data services to end users, these end users can then use the service provided to them to access online resources, interact with others, conduct commerce etc.

High speed does not always mean high throughput, the speed that an end user receives will depend greatly on:

- The lowest speed in the network between source and destination;

- The requirements of their application (how much data throughput does it require?); and

- Any external price/utilisation controls put on the user by the service provider for using a share of the network.

The last point goes toward bandwidth or data quotas. Quotas are used when the user has exceeded any imposed bandwidth limitations that result in the user connection being shaped (data throughput lowered by intelligent network devices to introduce delay, packet loss and other means to slow down the effective throughput).

In any consideration of high speed networks it is helpful to bear in mind just how much network capacity is required to deliver certain outcomes. The following table helps illustrate this, it shows how many gigabytes per month will be received over various speed connections (Y axis) versus the percentage of time that the connection is actually utilised at full speed (X axis):

Megabits per second	Percentage Utilisation 1%	2%	5%	10%	20%	30%	50%	60%	70%	80%	90%	100%
18.00	59	118	296	591	1,183	1,774	2,957	3,548	4,139	4,730	5,322	5,913
16.00	53	105	263	526	1,051	1,577	2,628	3,154	3,679	4,205	4,730	5,256
14.00	46	92	230	460	920	1,380	2,300	2,759	3,219	3,679	4,139	4,599
12.00	39	79	197	394	788	1,183	1,971	2,365	2,759	3,154	3,548	3,942
10.00	33	66	164	329	657	986	1,643	1,971	2,300	2,628	2,957	3,285
9M	30	59	148	296	591	887	1,478	1,774	2,070	2,365	2,661	2,957
7M	23	46	115	230	460	690	1,150	1,380	1,610	1,840	2,070	2,300
5M	16	33	82	164	329	493	821	986	1,150	1,314	1,478	1,643
4M	13	26	66	131	263	394	657	788	920	1,051	1,183	1,314
3M	10	20	49	99	197	296	493	591	690	788	887	986
1.5M	5	10	25	50	101	151	252	303	353	404	454	505
1M	3	7	17	34	67	101	168	202	235	269	303	336
768k	3	5	13	25	50	76	126	151	177	202	227	252
512k	2	3	8	17	34	50	84	101	118	135	151	168
256k	1	2	4	8	17	25	42	50	59	67	76	84
192k	1	1	3	6	13	19	32	38	44	50	57	63
128k	0	1	2	4	8	13	21	25	29	34	38	42
64k	0	0	1	2	4	6	11	13	15	17	19	21

Table 1. Data speeds and utilisation

This table shows that:

- a 64 kilobits per second line (ISDN speed) at 100% can transfer (in one direction) 21 gigabytes per month;
- a 1 megabit per second line at 50% can transfer (in one direction) 168 gigabytes per month;
- with a 12 megabit per second at 20% utilisation the user will receive (one way) 788 gigabytes – no current download plan in Australia can cater for such heavy network use.

The numbers in Table 1 are for one direction only (to the end user OR from the end user) – it needs to be noted that some ISPs meter both upload and download. However, in any FTTN where an asymmetric (more speed in one direction than the other) technology is most likely to be used the upstream (from the user) will be a lot smaller than the downstream (to the user). That upstream though can be in the megabits per second and the above table indicates the transfer capability of such capacity.

The next table shows how much data is received on a monthly basis by tuning into a certain number of hours per day of IPTV Standard Definition (SD) content and also IPTV High Definition (HD) content (details of numbers used for calculations can be found in Annex A):

Hours per day	SD TV - Data Required per Month (gigabytes)	HD TV - Data Required per Month (gigabytes)
1	20	108
2	41	216
3	61	324
4	81	432
5	101	540
6	122	648
7	142	756
8	162	864

Table 2. Data consumed using video

3 hours per day of HD content will consume almost 10% of a 12 megabit per second line or 324G per month (394G equates to 10% at 12 megabits per second); this is before Internet, VOIP or any other application. It also only assumes only one concurrent stream per connection – these speeds do offer the ability to deliver more than one SD stream and even one HD stream plus several SD streams – all of this would add to the throughput requirements.

Multicast v. Unicast

What is multicast

Multicast is ability to send one stream of traffic into a network and have the network intelligently deliver that stream of traffic to all parties that want it (have registered for it) without creating unnecessary duplications of that traffic through the network. IP multicast is what we refer to with respect to multicast, this requires use of IP address space 224.0.0.0 to 239.255.255.255.

Unicast is the sending of information from one party to another party across a network, whilst there are many paths for the traffic to travel through the network the data is only routed between the two parties.

Broadcast is where all parties receive traffic that is sent – regardless of whether they want it or not.

Need for IPTV and similar

Multicast is the most efficient way to distribute IPTV. If an SD TV stream requires 1.5 megabits per second and 1000 end users need to receive it then to do this via a unicast method will require the injection into a network of at least 1.5 gigabits per second – this will place a strain on the sending infrastructure, the backhaul between the ISP and the end users and, given the charging arrangements for some forms of backhaul, make the exercise commercially infeasible.

If this stream is to be sent to 1000 users on a multicast enabled network then only one copy of the stream is sent and the network ensures that it is only sent down paths to end users once. The sending infrastructure is not overloaded with 1000 different end points to transmit to, the network only sends one copy down each path to end users and as a result less of the network is used.

The above example is where there is only one IPTV channel and does not at all consider a multichannel environment. Given the consumer desire to be able to 'channel surf' (the rapid flicking around through an assortment of channels to find content of interest) having a serious lag between channel changes is not helpful. The process in multicast of having an end users software register for the channel does take a finite amount of time that will be noticed and found to be an inhibitor. To avoid this most large scale IPTV providers around the world have all channels on their service available at the access concentrator ensuring that the streams are ready for distribution to the end user at the press of a button. This places a large demand on the network (100 channels at 1.5 megabits per second for SD or 8 megabits per second for HD relates to a lot of traffic even when multicast).

When considering an FTTN operated by one party and the potential scope for backhaul charges to access end users on that network it is easy to see that very quickly the costs will add up to the point where it would be commercially infeasible to offer any form of IPTV.

Access to multicast cloud

If a single operator does end up controlling all fixed line copper based broadband connections in Australia then we need to determine the best way that all competing parties can access the multicast portion/facilities on that network. From this point of view it is important that:

- multicast facilities exist on the network that are capable of being shared between all parties based on the proportion of network utilisation;
- standard based multicast be used; and
- the cost to access the multicast facilities are in line with the portion of network that gets used not, for example, the economic benefit delivered to the service provider because they can sell TV over the FTTN.

Quotas

Most ISPs today use a system of effective speed control by the imposition of a data transfer quota. This comes in various forms but usually consists of:

- an allocation for traffic transferred to, from or both directions on the user connection that may be segmented into peak and non peak usage; then
- either a charge for excess traffic or a rate limit applied in order to restrict the usage of the network by the user for the remainder of the billing period.

This system allows network operators to dimension their networks, control their costs and leads to a diversity of packages offered in the market with respect to price and download allocation. Without these rate limits networks would quickly fill with traffic and costs would be incurred where ISPs purchase usage based backhaul and are required to increase the capacity of their routing and switching equipment to deal with the influx of traffic.

Quota measuring or enforcement is a multicast environment is something that will need exploration depending on the outcome of this process.

The cost to deliver data to the user is covered later in this document.

Increase in inbound and outbound data

New applications have seen the profile of traffic usage change, where the demand for access has been for users to receive content, new ways to interact online as well as new ways to distribute content are seeing more demand placed on inbound and outbound channels.

Anecdotally ISP have reported shifts the pattern of traffic to where traffic from the customers is rising towards the volumes of traffic to the customers marking a shift in usage of end users.

New types of applications

Social Networking is the hype word given to new forms of Internet activities that allow users to interact with other users either by sharing content and/or experiences. With respect to network design social networking sites introduce a larger component of upload versus download than do traditional online sites – this is due to the interactive nature of these sites. These sites also tend to be built to be open and link to other user communities where possible allowing the sharing of content in a far easier manner.

The following highlights a few of the more popular social networking websites and tries to highlight their short term evolution and how that may impact throughput calculations for future Internet user behaviour. All of these applications tend to require high speed, but more importantly they require high throughput – from end to end.

YouTube (www.youtube.com) – a video sharing website that allows users to upload videos, create their own custom channels of content. Currently content is optimised for low speed content at 320x240 frame size and overall bandwidth lower than 256kbit/sec. The widespread use of cameras in phones and other devices and the increase in their resolution (5 mega pixel cameras very common) will lead to higher bit rate video on sites such as YouTube. Mix with this the ability to geo-reference by adding built in GPS systems for global location and instant upload of content from portable device to website and it is easy to see the demands from YouTube and similar sites increasing dramatically. This content cannot be multicast as it is an on demand stream that will become more personalised as time goes on with newer technologies such as being able to target commercials specific to the user receiving the content (this will also make it hard to cache).

Flickr/Image sharing websites (www.flickr.com) – this is very similar to YouTube in that users can upload content, the case of Flickr the content is typically still images. These can then be linked via user communities and other attributes as well as accessed by other applications (see Photosynth).

Amazon S3 (http://www.amazon.com/gp/browse.html?node=16427261) - online storage and on demand computing – the ability to serve web objects of up to 5 Gigabytes in size, paying only for what you use as an example. This is part of the

Amazon elastic computing initiative allowing for on demand almost everything with respect to computer resources.

Photosynth (http://labs.live.com/photosynth/) – this sites relies on many aspects of social networking to attempt to build models and collages of information uploaded to social networking websites. This allows models to be built using large collections of photos of a place or object to render the model. For a glimpse of what may be required for the future so far as image processing, network and server bandwidth is required a look at Photosynth will help highlight future network demands.

Joost (www.joost.com) – from the makers of Skype comes a peer to peer (not illegal downloads but using the peer to peer protocol) video service with real world content. Shows on this include the CSI franchise shows, Jericho, NCIS, Comedy Central, Reuters, CNN and many more. Currently shows only play on the PC but with time this is due to change. Bandwidth requirements (upload and download as it is peer to peer) for this applications are quite intense.

Video Conferencing – in general video conferencing can be used in a business to business mode, business to consumer (medical professional to patient consult for example; if the desire of the health sector to do it, ethical and privacy issues are overcome) and also consumer to consumer as a substitution for email, SMS and other person to person communication. The symmetrical nature of this service and the increase in video resolution will ensure that this type of service will place unique demands on any network.

Current Situation and bottlenecks

The current models for the delivery of high speed broadband for competitive carriers looks like:

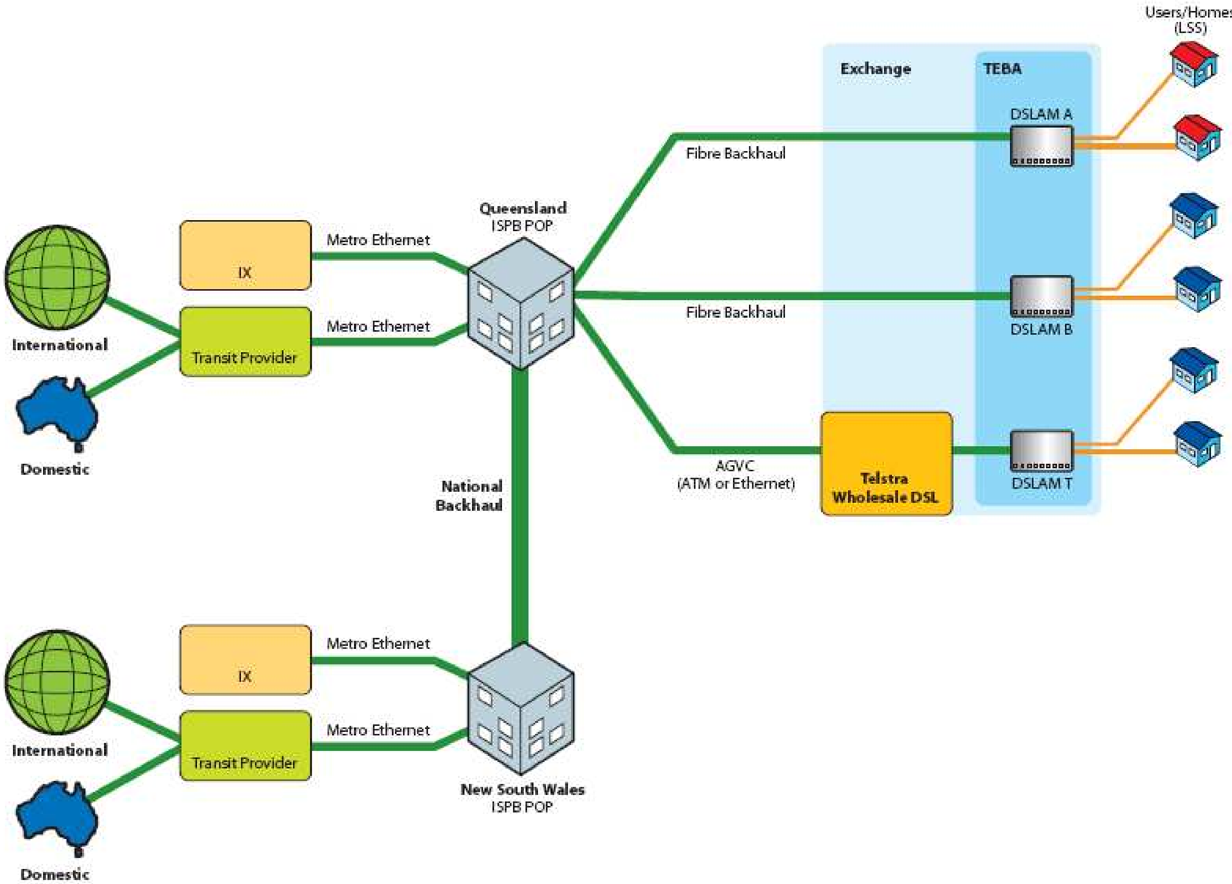

Figure 1. Typical Hybrid ADSL Service Provider

In Figure 1 (above) the ADSL service provider uses a combination of Telstra Wholesale DSL and their own DSLAMs. It is most cost effective to access the wholesale offering from a POP in each major capital city. In the diagram above the DSL access arrangement pictured from the Queensland POP is replicated also at other capital city POPs.

The service provider then, at the POP, blends peered traffic as well as traffic from its other POPs and transit (global Internet connectivity) to provide a full service global solution to the end user.

The bottleneck in this situation is access to the AGVC for access to the wholesale ADSL users – this is controlled by the wholesale provider and typically has no controls with respect to protocol, speed or quality. In Figure 1 the service provider has mitigated this bottleneck by installing their own fibre based backhaul.

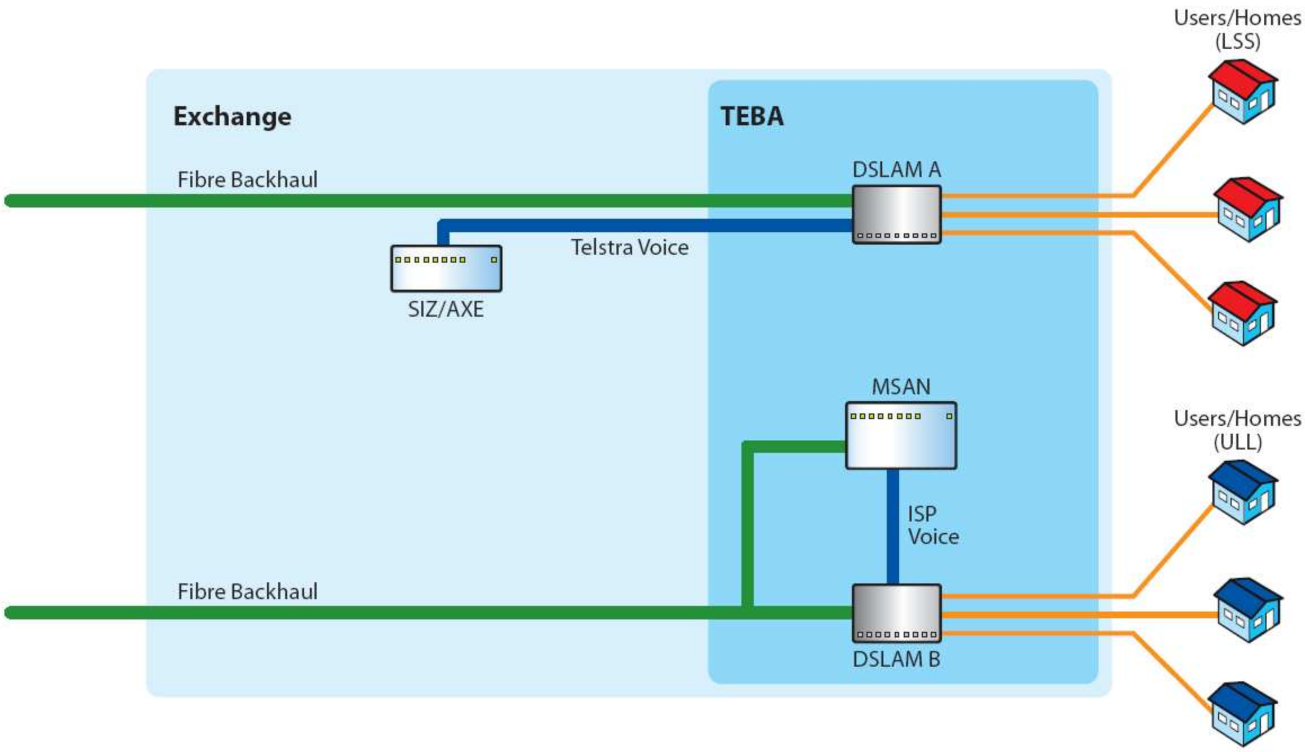

Figure 2. Inside TEBA

Figure 2 shows the general arrangement inside of the TEBA (Telstra Equipped Building – Telephone Exchange) facility. Using ULL competitive ADSL providers can also provide their own POTS voice service to the end user or the users can continue to acquire this from Telstra – in this case the service provider will use LSS. The provider uses their own DLSAM, which may or may not be voice enabled, it is most likely multicast enabled – enabling the sending of one to many data streams such as broadcast TV, scheduled video and audio content.

There is no ability under the wholesale ADSL model to provide a non Telstra POTS service other than through resale of Telstra dial tone. There is no ability at all to offer any multicast on the wholesale service. There are examples in the market now (TPG IPTV) of competitive carriers deploying this technology on their own networks.

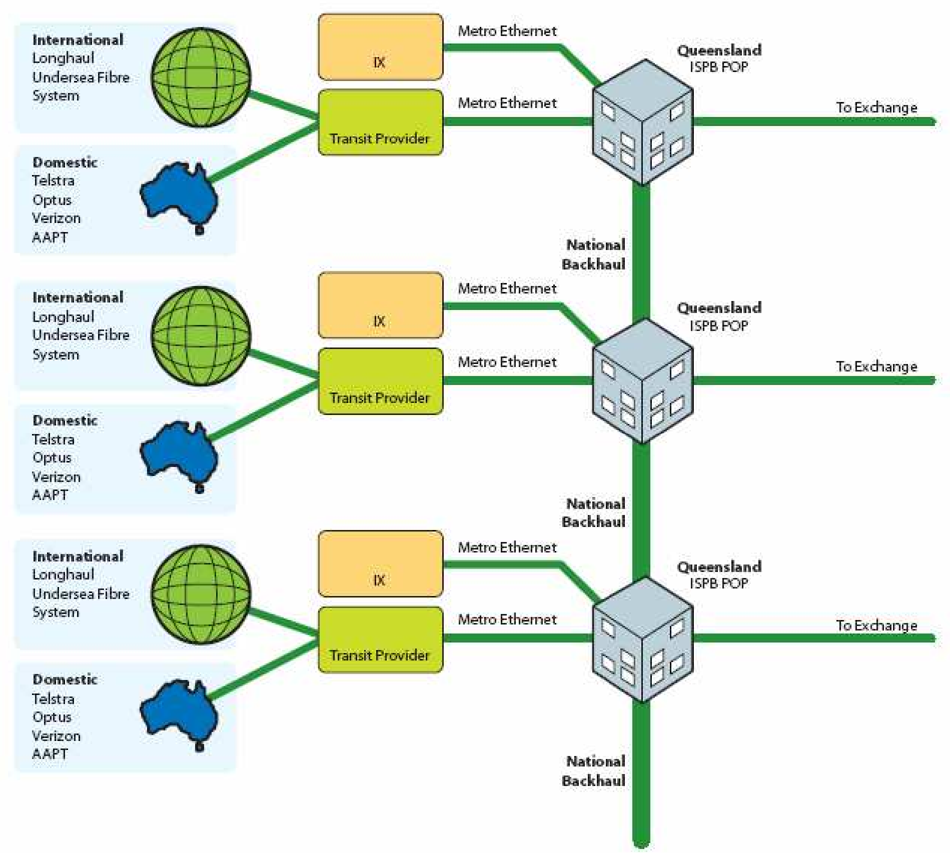

Figure 3. Onward Connectivity

In order to provide the full range of services to end users ISPs need to interconnect to other networks. Domestically these may be peering networks (such as PIPE Networks provides) or carriers that refuse to traditionally peer such as Telstra, Optus Verizon or AAPT/TNZ. Many of these carriers offer access to Australia based routes (even those from in the same city) at rates that are identical to traffic from overseas – this leads to little or no incentive to peer locally with these networks due to the inequity in the peering arrangements.

It should also be noted that very few Australian government organisations peer their content, as a result service providers are forced to effectively pay for this content as if it was hosted overseas. The mere act of all Australian government organisations making their content available for peering would help make more content available for service providers to send to their customers that was not sourced from overseas or locally for the same price as overseas.

The lack of an equitable peering and interconnection framework in Australia means that we import far more data than we need to. It is cheaper to place content in the US so people do that despite that their audience may be based in

Australia. Under this model, where content is based overseas due to cheap hosting or content in Australia being given to Tier 1 providers solely, users pay the price in limited experience due to traffic shaping once quotas are reached or in excess traffic charges once limits are exceeded.

International connectivity is supplied by many players including the four large domestic networks of Telstra, Optus, Verizon and AAPT/TNZ or by going direct to the operators of submarine cable systems. In Australia these are Australian Japan Cable (owned in part by Telstra) and Southern Cross (owned in various parts by TNZ, Optus and Verizon) [See Note]. Newer entrants have entered the market in this respect with Telstra and PIPE Networks recently announcing cable ventures in the Pacific region – this should see positive pressure placed on pricing due to competitors entering the market, this has nothing to do with new capacity as the in place cables are at small fractions of total end of life capacity but everything to do new competitors.

Note: SEA-ME-WE 3 not considered due to location and capacity

Cost per dollar to access bandwidth used to deliver services to end users

The portions of network involved in getting data from source to destination are all important as they add into the cost of a broadband service - especially in a world where the cost to run the network is fed back to consumers in line with their usage (quotas and/or excess charges).

Generally the components in the network are:

1) **the last mile** – the portion between the end user device and the first electronic interface on the telecommunications network (the DSLAM, base station or even dialup modem on the ISP side – the access concentrator);

2) **backhaul** from the access concentrator to the service provider POP;

    a) In a wholesale model this may be in the form of Aggregated General Virtual Circuit (AGVC) where an access interface is purchased (either ATM or Gigabit Ethernet) then capacity is acquired over that interface by varying its throughput. This allows connectivity to all last mile connections you have with that provider over that interface. This is used typically to connect to many last mile end users that may be physically hooked up to many access concentrators into one network interface for delivery to the wholesale customer;

b) Where providers have installed their own DLSAM this could be purchased transmission from a provider in the form of dark fibre, SDH or Ethernet based point to point services and typically goes from exchange/access concentrator to POP;

3) **Domestic transmission**/intra Australia carriage – this is used to connect other POPs in other states or locations

4) **International IP** – sometimes called transit IP this is the process of interconnecting your network to the rest of the Internet;

5) **Domestic peering** – whereby you connect to other networks that make local interconnection available.

The following table has been constructed from feedback from customers of PIPE Networks and what we know of the charging for various elements of the network. It is broken down to a range for the price to access a megabit per second per month of capacity on the various legs of the network. Given the sensitive nature of the information and the sources I gathered it from I have presented it in an aggregated form here.

In the table below, all these costs are additive - you use some or all of all parts for your broadband service (a choice or combination of items 2a and 2b). Any FTTN or similar proposal that only attempts to deal with the access speeds without taking in account the impact of other costs and speeds in the network will not achieve anything new. Any proposal that removes the ability of providers to access unbundled portions threatens to remove the fantastic cost saving and network flexibility afforded by (for example) providers using their own backhaul.

Network Portion	**Cost per megabit/sec per month**
Last Mile (1)	Between $2 and $5
Backhaul (2a) ***(2a or 2b)***	Between $55 and $180
Backhaul (2b) ***(2a or 2b)***	Between $1.50 and $30
Domestic Transmission (3)	Between $20 and $1000
International IP (4)	Between $170 and $300
Domestic Peering (5)	Between $1 and $300

Any fibre to the node network seeks primarily to ‘fix’ the Last Mile ‘problem’ at the expense of losing the flexibility offered by unbundling where competitors can leverage their own backhaul and provide innovative products.

Walled gardens and time based quotas

Many service providers use the walled garden, a set of locally available content sources such as the ISP’s web server, FTP, mail server etc as cheaper sources of data that users can consume. This data is then metered at a different rate, often $0, than externally received content. Some ISPs include all of the content they can access by inexpensive peering in the walled garden.

Many ISPs take a time based approach to give customers a split quota, one quota available during the peak period for data access and the other for the traditionally quieter period. This allows users with bandwidth intensive applications to use their link after hours by scheduling downloads or other activity to take advantage of the timing break offered by their service provider.

Interconnection

Interconnection is the most important issue for any FTTN proposal. If competitive services other than plain resale of a bit stream service are to be provided then interconnection must be given highest priority.

Interconnection needs to occur for both data and voice (POTS delivered). A modern approach needs to be taken to both. For data - work undertaken by the DSL Forum (www.dslforum.org) needs to be used where applicable. For voice - SIP reference groups and interconnection standards and drafts should be used as the basis for determining a voice interconnection method.

Without technically and commercially feasible interconnection no innovation will occur. Interconnection needs to occur **as close to the network edge** as possible to allow providers to bring their own network elements instead of only having access to a bundled service from the network operator.

It can be seen worldwide that incumbent operators, if left to their own devices without competitive pressure are hesitant to introduce innovations. Under the current DSLAM based competition where interconnection occurs at the ULL/LSS pair the following has occurred as examples of innovation by non incumbent carriers:

- Quality of Service – service providers can control their own networks and configure their equipment to be able to sell a service that has built in Quality of Service. If an FTTN is bundled to the service provider (backhaul included for example) then you are forced to use the Quality of Service settings as dictated by the network owner – the same as everybody else – there is no opportunity for differentiation – except on price;

- ADSL2+ - this higher speed version of DSL was first introduced by competitive carriers;

- In Australia ADSL higher 1.5 megabits per second was introduced first by competitive carriers;

- Annex M – this variant of DSL allows for higher speed uploads under various conditions and was first introduced by competitive carriers;

- Naked DSL – the requirement to have a POTS line as a pre-requisite for a broadband connection;

- Multicast – Multicast is now being used by competitive carriers to provide IPTV; and

- Multi-line DSL/Mid-band Ethernet/SHDSL.bis – the ganging together of several ULL lines to create short haul services of large capacity essentially using DSL like technologies providing capacity of up to 45 megabits per second over copper.

Potential models for FTTN

Figures 4 through 8 at the end of this document deal with the various models that exist now and are potentially being looked at for FTTN.

Current competitive model

Figure 4 is the current model where access seekers (competitive carriers) install their own DSLAMs, acquire backhaul and roll their own product. They are free to set all of the parameters of their network and how it gets used.

Proposed G9 Model

Figure 5. Under the proposed G9 model access seekers can connect to a Local Access Point (LAP) in or near the current Telstra Exchange. They can then access the new network via a managed service from the LAP to the node then on the copper to the customer. This model allows for OK, close to edge, interconnection to the end user.

Access seekers are still free to setup their network, acquire their own backhaul and determine their own contention ratios, QoS, protocols etc from the LAP back into their network.

The G9 model did have a Transit Access Point (TAP) concept where the G9 group could provide backhaul to an access seeker but it was not mandatory. If new telecoms investment is to be encouraged this style of network element (TAP and included backhaul) should not be included as part of any new build as many carriers already perform this task now (at competitive exchanges). TAP style activities at the very least should not receive government funding.

There is an *upside diversity benefit* (Figure 6) that occurs if nodes can have a diverse fibre installed to them. Under this design the resilience of the Australian CAN will increase dramatically as only a small portion of the customer service would be non diverse (from node to user). This will increase the overall cost of the network and you may lose some (probably not all) upside benefit from the fact that you have distributed your exchange electronics evenly across Exchange Service Areas (ESA) by placing them onto footpaths etc. Only time will tell if this leads to more or less downtime versus the risk from cutting one of the main cables (copper trunk) as represented in Figure 4.

Interconnection at Super Node

The models represented in Figure 7 and Figure 8 are essentially a heavily bundled, hard to differentiate, easy to game and generally an undesirable model. Both of these models operate whereby all of the services provisioned to end users are taken back to a Super Node – most likely one or more of the large inner city exchanges. Access seekers are then forced to take service at that location and thus will inherit the network operators contention ratios, QoS settings, protocols, SLA etc. This means that every access seeker will offer the same service with competition based on price.

It needs to be noted that collection of users at a super node type facility will ensure that we remain with the current model as wholesale ADSL is now made available. This is identical in concept and higher level operation to what we already have.

The All Optical WDM model (Figure 7) is where all nodes are allocated a separate light frequency and the node’s traffic is multiplexed back to the supernode. In the diagram it is via the LAP but need not be, it could go straight back to the supernode giving NO opportunity for innovative carriers to access the service other than as a bundle.

Higher Protocol Aggregation model (Figure 8) is more like the current system of whole ADSL and in fact could be the current system. End user services get delivered to the access seeker where and when the network operator tells them, all contention rates, QoS, protocols etc are exactly the same across the network and the ability to innovate is lost.

Interconnection at Node

Due to the low numbers of customers connected at the node (sub 200 typically) interconnection at the node is feasible but challenging. To expect multiple providers to build a fibre path to the node for the opportunity to service less than 200 subscribers is hard to justify. There would also be numerous issues with being able to 'multi-tenant' the DSLAM at the node in order to partition it use between multiple access seekers.

Lack of options to un-bundle under certain models

The models in figures 4 through 8 are extremely hard to unbundle. Any attempt to bring all access seeker interconnection back to a super node, super exchange or large aggregation point makes this service extremely hard to unbundle. Those models that require interconnection of access seekers far away from the edge are dangerous for competition and innovation.

Range of Services

A variety of services are currently deployed over exchange based copper, these include alarm services over things like Permitted Attached Private Lines (PAPL), Megalink, frame relay, digital voice and many many others.

Will the new FTTN be able to support all of these services ? Will the equipment that the end users currently have working on these services be compatible ? If they are not then who is responsible ? In general who is responsible when a new network comes along and breaks a current commercial activity that was working perfectly before the switch ?

Any proposal that dictates a mandated shift over to a node based network will need to answer these questions. Vendors will say that all of these services will work but those of us with experience of vendors know that they often confuse what works now with :

- tomorrow;
- the next quarter; or
- the next software release.

or they end up providing a solution that supports all services but not all combinations at once. Removing choice from end users can lead to many unintended consequences.

Non Price terms

Churn

In a healthy competitive environment users are free to chose the provider that will best suit their needs. These decisions may be based on price or features (with a common bundled bit stream service features will not vary from provider to provider) but regardless users will churn.

Previous churn processes have been based on fear and the desire to make it as hard as possible for a user to move as the law will allow. Consumer sensible terms need to be determined that allow users not to be the victims of carriers seeking to inflict commercial damage on each other through burdensome churn process.

Provisioning

Like churn, the process that allows a user to quickly and painlessly connect to a provider (where no churn is involved) needs to be enshrined in the operation of the network. This will mostly deal with the network day 1 and any additional loops being added as suburbs change and grow.

Unbundled network elements

If there is any thought of a day 1 fully bundled service as the only option being offered to access seekers on a FTTN then prior to network switch on positive action is required to ensure there is a rock sold timetable for unbundling. This should include an independent party (regulator) having the power to set the timeframe if need be as well as prices.

Current investment in exchange based services

There has been over the last few years an explosion in competitive carriers installing their own equipment in Telstra exchanges, accessing LSS/ULL and providing innovative services to end users. All of this has been done in the face of extreme resistance and a bloated regulatory process regarding access, disputes and other issues. Despite this resistance though and because of the perseverance of the competitive industry many people in Australia can now access a broadband service other than the Telstra 1.5M/256k ADSL.

This has come about due to large investments by those competitive carriers. Any network that requires the removal of the competitive infrastructure will be accompanied by relevant calls for compensation as well as questions that may linger for future investors in Australian infrastructure – why invest here when all of the sudden the rules may change to strand your investment.

Dual use of CAN

The current investments can stay in place and not interfere with any use of the copper network for a FTTN deployment. In this way an FTTN can proceed without stranding current investment - despite calls to the contrary.

PIPE Networks commissioned a report on this matter in October 2007 to understand the technical issues with running parallel exchanged based ADSL and node based services, the report can be seen at :

http://www.pipenetworks.com/docs/media/ASX_07_10_31%20FttN%20Local%20Loop%20Impact%20Report.pdf

(http://tinyurl.com/2kbzvn)

In summary it was found :

"… there does not appear to be any technical or interference management basis that might justify removing the copper pairs between the exchange and a node once a FTTN network has been deployed."

Decisions need to be made on the future deployment of VDSL2 in order to ensure that spectral compatibility with exchange based services does occur. This is important as it is just as easy to make a decision with respect to the VDSL2 spectrum deployment plans that ensures poor compatibility with exchange based services for no other gain (other than to ensure exchange based services will effectively not work).

Overbuild protection

The recent G9 proposal included a request and suggested format for overbuild protection. Any overbuild protection needs to be carefully assessed as to what it is trying to achieve – if it is about the commercial protection of a potentially undesirable and expensive FTTN end user product (rumours of $59 and $85 for a basic phone service and 512k Internet package would fit the undesirable and expensive label) then perhaps we are building the wrong thing for end users.

If any arrangements for overbuild protection were to be put in place for commercial protection reasons it is recommended that they should also potentially consider for inclusion:

- protection from fixed wireless operators competing in the broadband space – remove the ability of mobile operators to compete with the monopoly FTTN;

- removal of all current HFC networks that may compete; and

- bans on satellite delivered broadband services in the FTTN service areas.

Issues that should be dealt with concurrently

Internet Interconnection

[Note: this is Internet Interconnection as distinct from interconnecting to the FTTN]

As stated earlier in this document there is a real cost associated in domestic peering, between $1 and $300 per megabit per second per month. The cost variation here is made up of:

- at the low end: private peering between parties;

- at the medium end: peering point services where many parties peer in a one to many relationship lowering overheads and increasing available content; and

- at the high end: this is paid peering from a Tier 1 – this is typically charged at the same rate as international traffic (transit) – so the price that a Tier 1 will charge for a unit of data from next door is the same as if it came from the West Coast USA !

This has lead to it being far cheaper for online businesses to host their content in the US. If you peer your traffic locally you receive the benefit of:

- local direct handoff via a path that is typically shorter than transiting a Tier 1 – this leads to better quality;

- vastly cheaper traffic where the cost to do so is usually fixed and more often than not a small internal cost transfer to build the relevant infrastructure to participate; and

- vastly improved quality – peered links typically have more head room and can therefore cater to bursty network activity better than transit links. Due to the costs of a transit link you do not buy these with lots of room for unexpected bursts of network activity – you buy just enough for now and many ISPs will even budget for congestion on these links for set durations each day – this is when the net goes slow. ISPs that make that decision do so in order to control costs and they add to the strata of providers out there giving choice to consumers.

The lack of participation in this area by Tier 1 providers (at anything other than full transit pricing) has lead to this situation.

The sector of the Australian economy who are the biggest laggards when it comes to ensuring that their content is not given solely to a Tier 1 (for resale to the service provider industry marked up 20 times to a rate as if it came from USA) is Government and publically funded bodies. As a non exhaustive example:

www.australia.gov.au appears to only have connectivity via transit;
www.aarnet.edu.au appears to only have connectivity via transit;
www.csiro.au appears to only have connectivity via transit;
www.nsw.gov.au appears to only have connectivity via transit; and
www.ato.gov.au appears to only have connectivity via transit.

By merely stipulating that government funded entities should make their traffic available to local service providers and not solely a Tier 1 will go to start reducing the cost to access content. This will bring down the costs for end users to access government online services (a very worthy goal with or without a FTTN) and could help breathe life into the Australia content hosting industry. It will also save the government money.

Creating a more vibrant local content hosting industry in Australia should be seen as an easy to achieve national goal. If the cost to exchange traffic between ISPs and carriers is linked to the actual cost and not the price of international traffic then more content will be hosted in Australia.

This is a priority regardless of the price of transit, even with low price bandwidth coming online in Australia with the advent of new submarine cables, data in Australia will always be more expensive than in the USA (the main benchmark area for traffic pricing). With rising volumes service providers will always needs to be able to access low priced traffic for consumer needs.

Annex A. Specifications

SD Video Bandwidth – 1.5Mbits/sec
HD Video Bandwidth – 8.0Mbits/sec

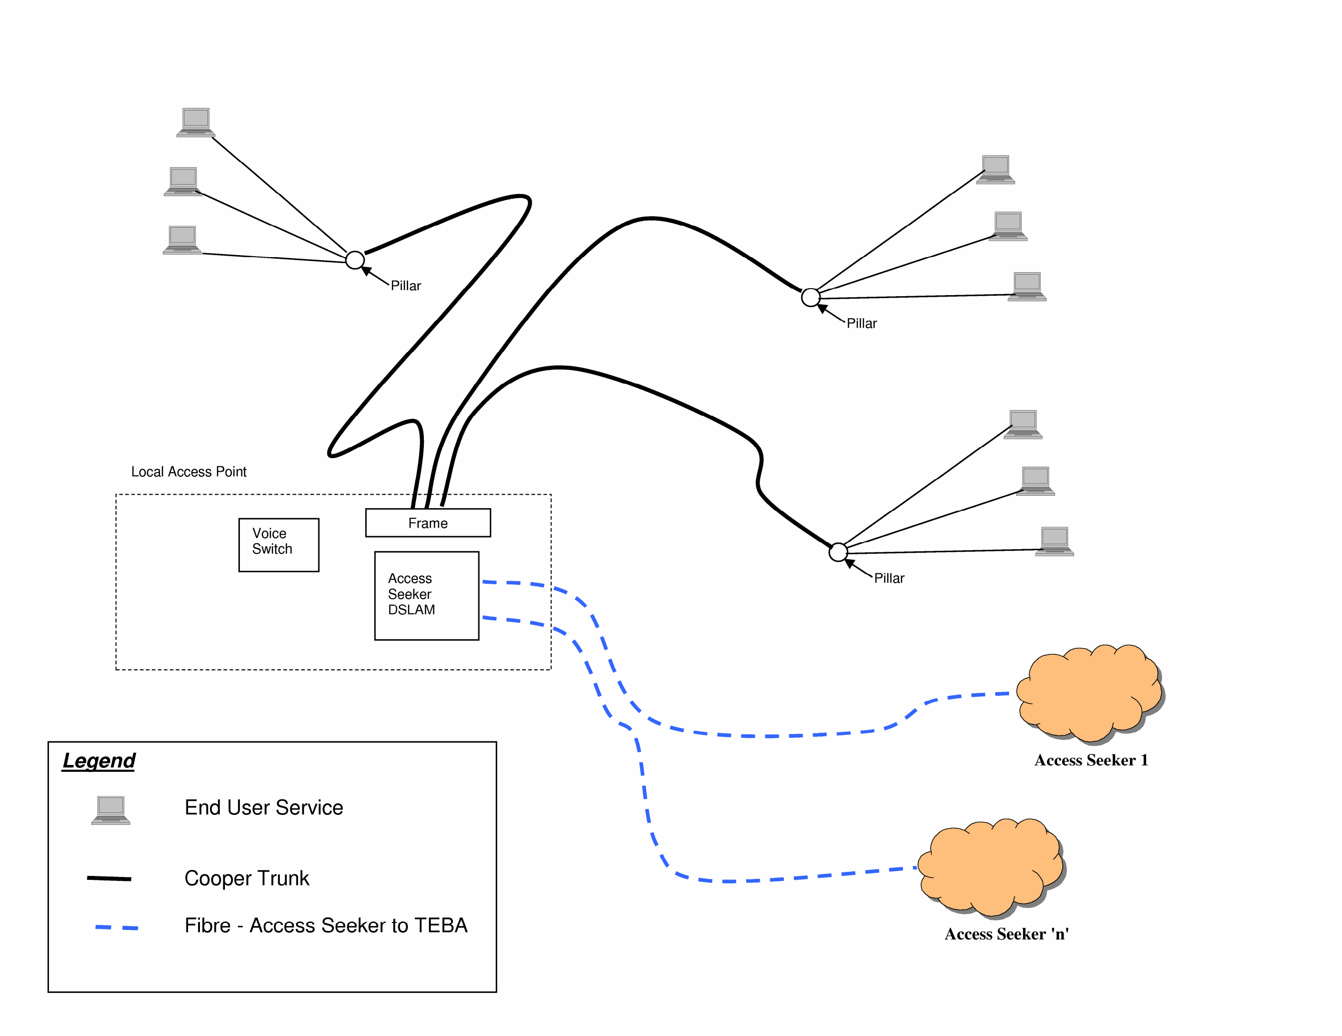

Figure 4. Current Competitive Model

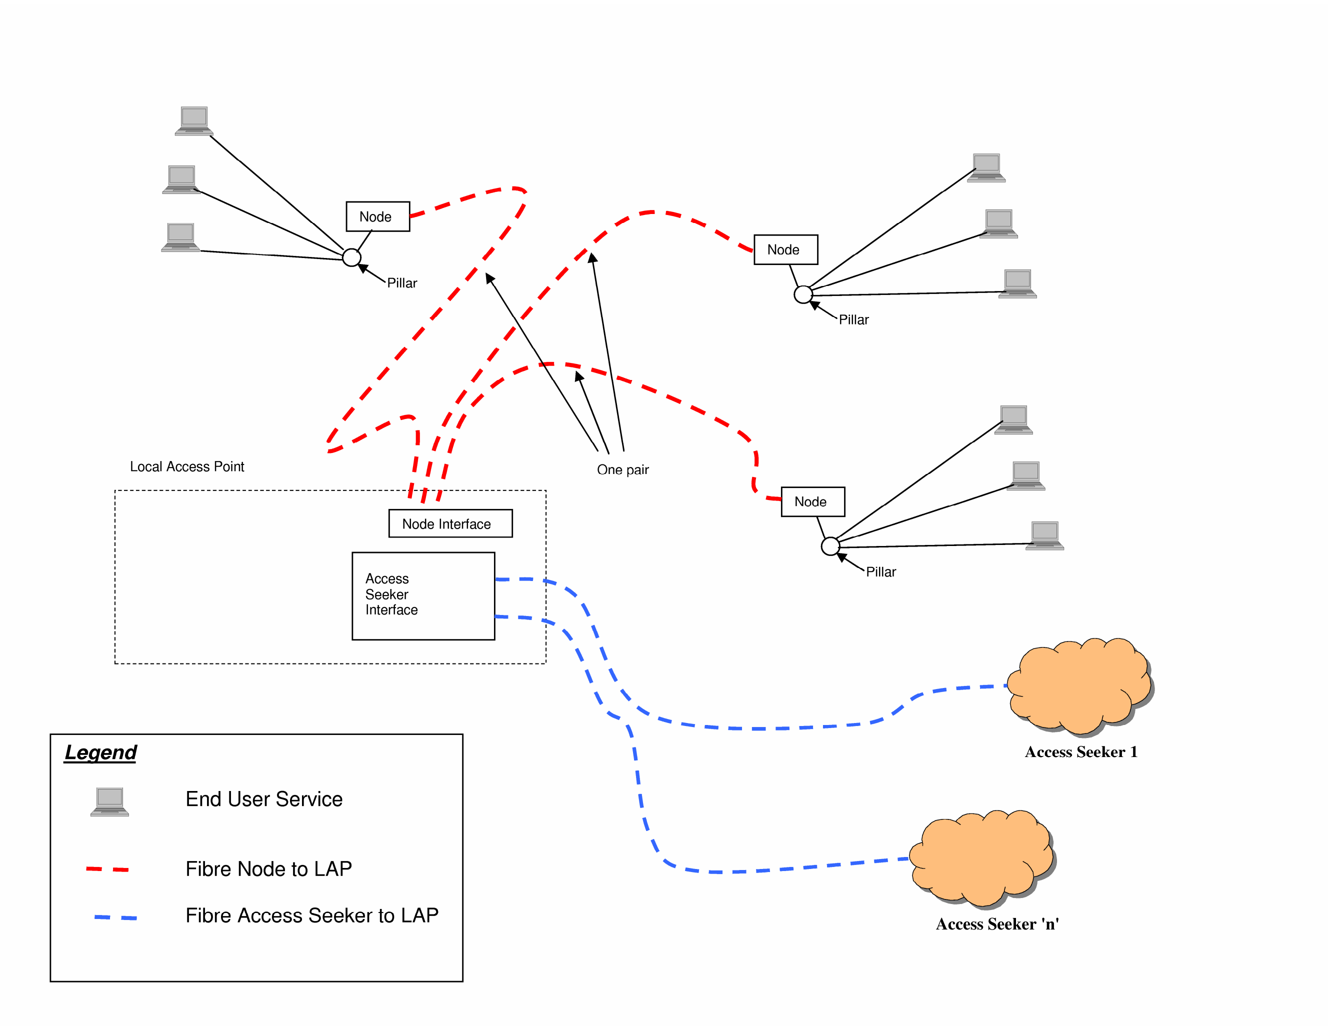

Figure 5. G9 Style FTTN

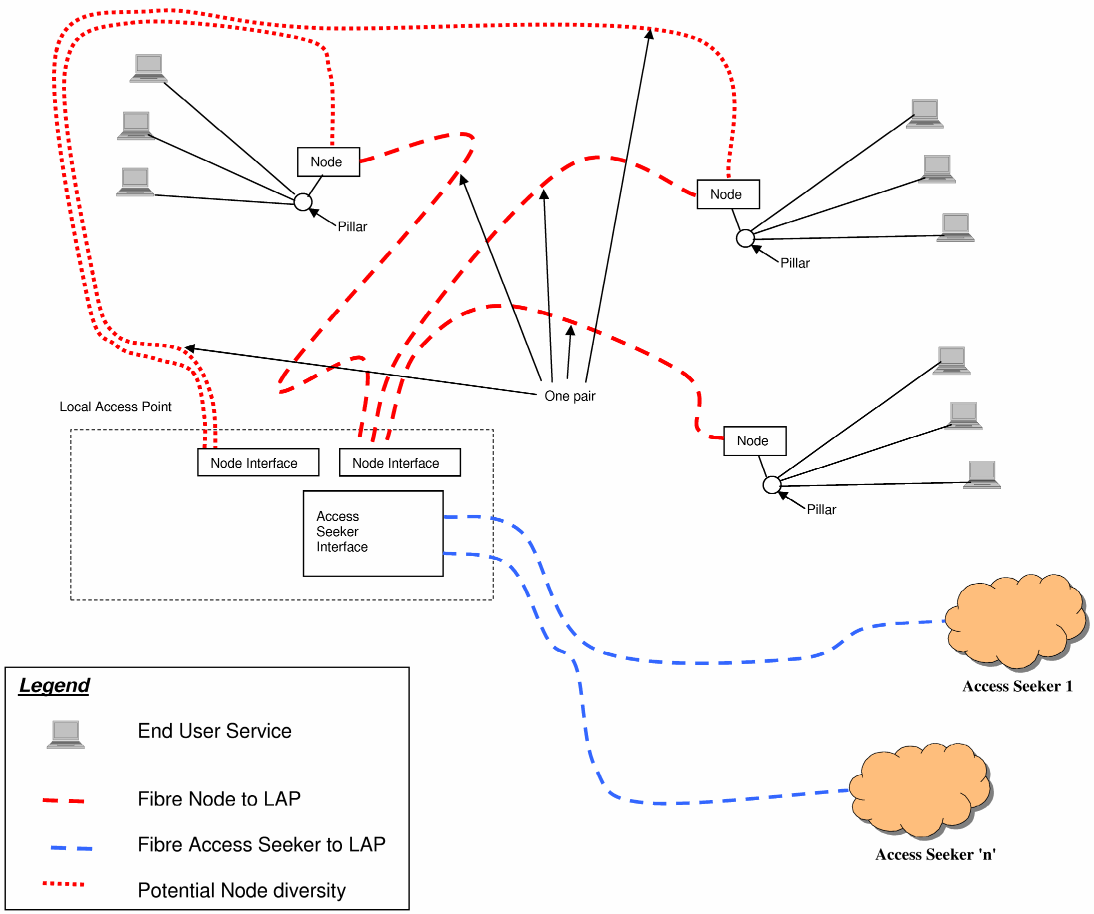

Figure 6. Upside diversity benefit

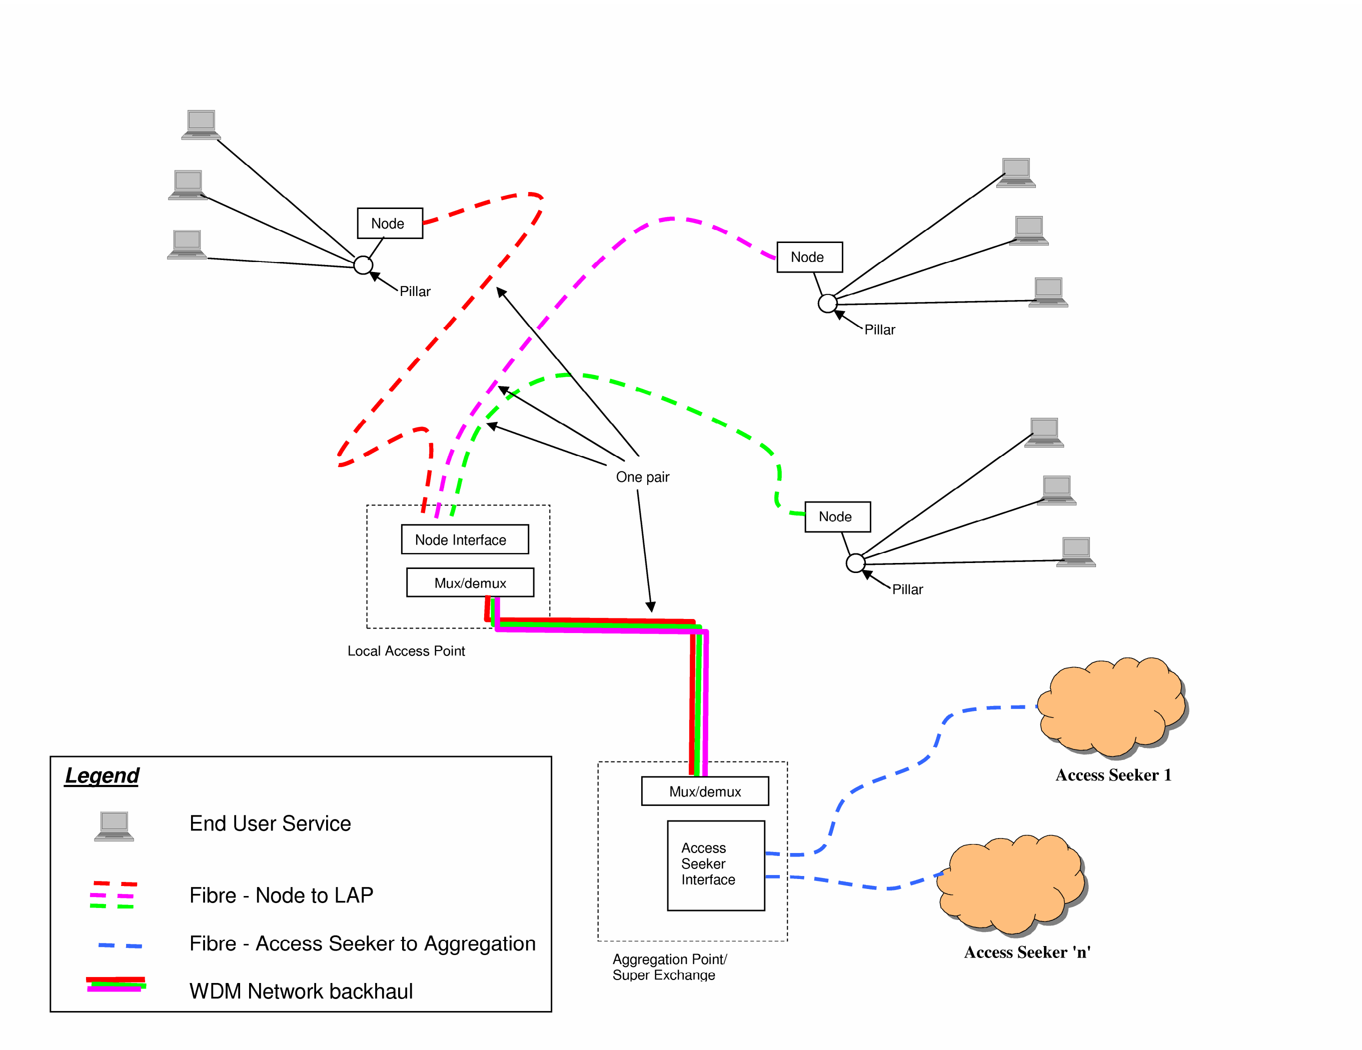

Figure 7. All Optical WDM model

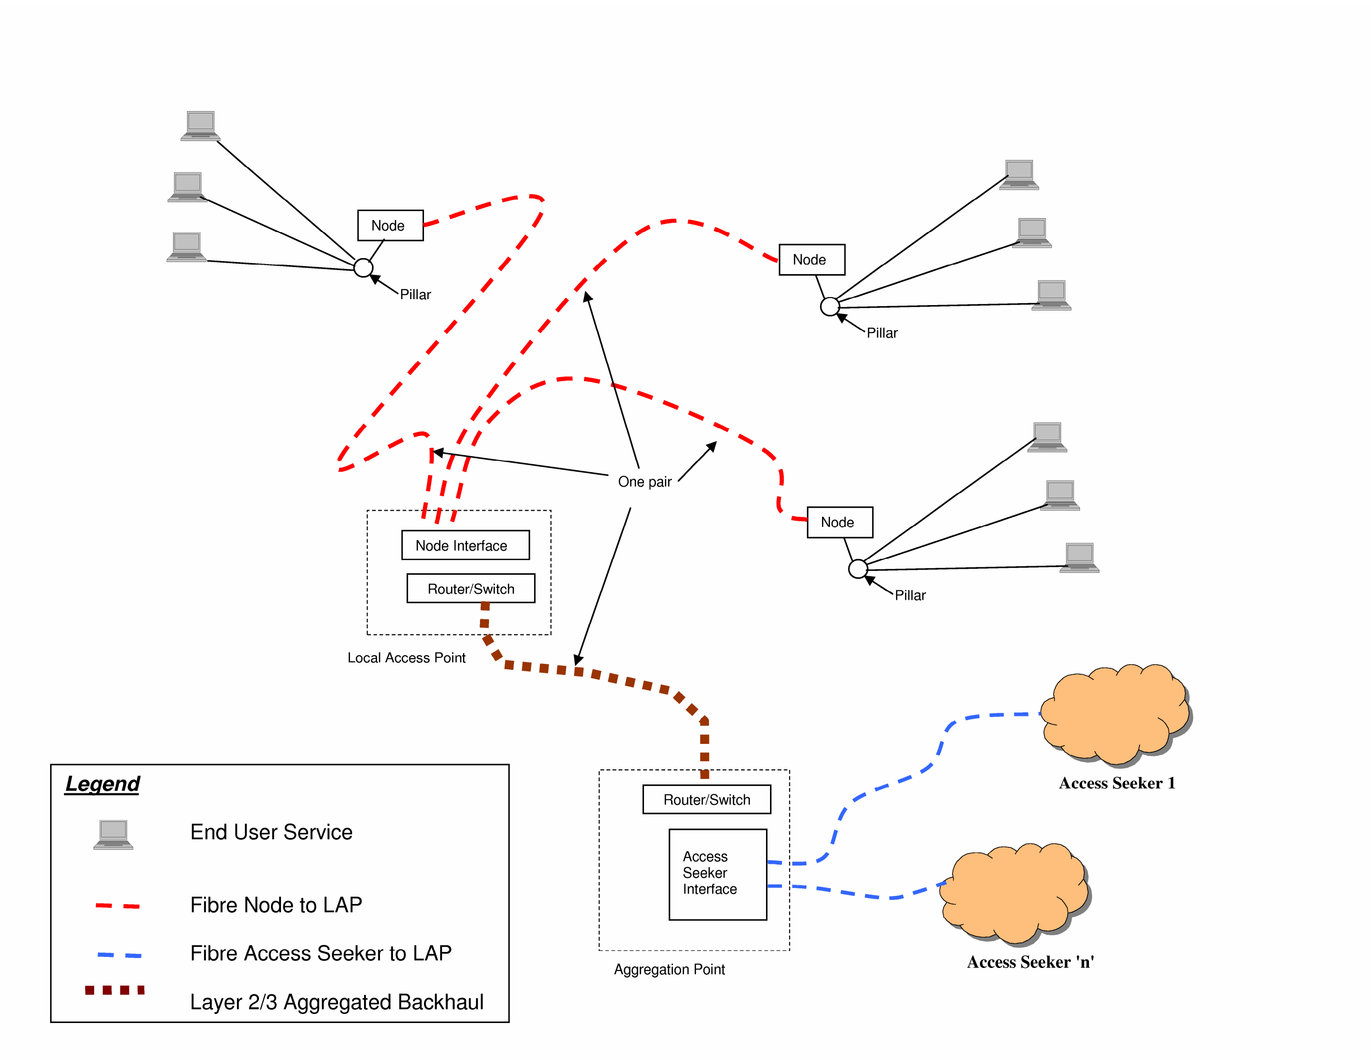

Figure 8. Higher Protocol Aggregation model