'Timely & Efficient NGA Investment'

Executive Summary

We are nowhere near the limits of the existing access network. We use less than 2% of the local loop capacity, but peak throughput is an issue because of bottlenecks and immature architecture on other parts of the network. We have a lot to do before we are ready for NGA.

There is an assumption that we need more local loop speed but I believe that this is flawed. What we need to do is give the market time to work out whether this is true or whether the speed gains can be achieved in other ways.

Distributed networking, local caching and storage and predictive off-peak downloads all based on Peer-to-Peer principles all allow the management of peak loads by deferring the transfer across the whole day.

P2P is not just about free-loading and copyright theft – that was just the sandbox in which the technology developed – the technology itself provides an intelligent routing layer onto a dumb network that gets data where it is needed *before* it is needed. Zero network latency…

On the copper itself, there are a number of projects that may deliver yet another step change in performance. ASSIA Inc, working with the University of Melbourne is one example – what they have found is that by minimising crosstalk you can deliver as much as 250 Mbps on existing copper loops. There are other, less credible, projects that promise even more.

The existing network has not yet been allowed to fully evolve so it would be premature to suggest that it needs to be replaced. More emphatically, it would be extremely dangerous to pre-empt what the most efficient solution was without giving some of the other options time to play themselves out. There is a False Sense of Certainty that FTTH is the answer.

New builds offer the opportunity to rollout fibre as cheaply as copper so in those cases, it would seem sensible to do this so that the fibre option can be fully evaluated. Elsewhere, replacement is an entirely different question as there is already a basic service with the potential to evolve.

What would seem clear is that evolution should happen from the core outwards, rather than from the access network inwards. Backhaul issues are being addressed, but there is a long way to go before the backhaul, exchange and cabinet facilities are at a level where an alternative operator has a real opportunity to innovate.

In general, improvements to copper performance come by shortening loops more than by using a better type of xDSL. The real issue in UK Broadband today is not Next Generation Access, it is the Digital Divide for Last Generation Access because of very long loops on existing infrastructure.

This is where policy and regulation should play. There should be universal service right to enough speed and capacity at a reasonable price. This might start at 512k with 2GB for £9.99 so that focus can be put on areas where this is not possible.

Solving such problems will undoubtedly involve shortening loops, perhaps by converting some cabinets into exchanges. This is an expensive task in rural areas so once policy has set the targets, regulation can work whether artificial assistance is required to make addressing this issue an economic proposition. It might also be a good sandbox for FTTC.

Where LLU has proved the market to be competitive, Policy and Regulation needs to back-off and let it evolve. The real job of Policy and Regulation is to focus on how to level the playing field for LGA where there is no competition.

Definitions and Key Data
I think it is important to detail what I understand by NGA as opposed to NGN before I pile to the questions themselves. I am using the following demarcations

- NGN represents the elements of the upstream network (towards the internet)
    - Around 130 core BT nodes and the fibre between them
    - OLO core networks (using wavelength or fibre)
    - **Does not include** backbones built on leased capacity
- Backhaul describes the portion of the network that connects the NGN to the DSLAM in either the exchange (LLU) or the street cabinet (SLU). It also includes any point to point long haul circuits used by this traffic using STM-x or Ethernet on a backbone network to a centralised core infrastructure.
- Access describes the connection from the exchange to the wall socket in the customer's premises

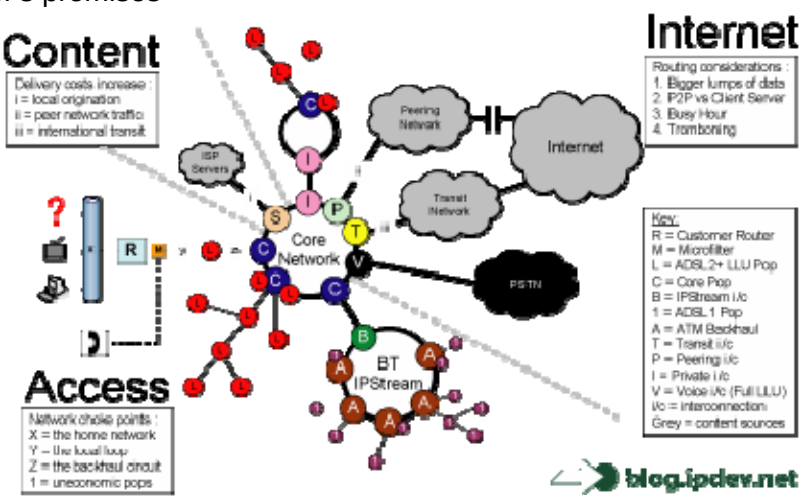

It is also worth splitting out the logical layers of the network in this discussion. I think you have to consider the Layer 1 (physical), Layer 2 (data), Layer 3 (IP) together and in isolation.

Backhaul today runs layer 2 tunnels to the core nodes before the traffic is routed in the NGN. The extent of Layer 3 deployment is a very important consideration in the effective use of existing resources.

The effective use of existing resources is the centre of my response. A 2 Mbps link is capable of delivering 642 GB per month of data and yet average UK internet usage is what? 5 GB – 10 GB? Less than 2%...

Question 1: When do you consider it would be timely and efficient for NGA investment to take place in the UK?

NGAs should be built wherever there is a new build project like Ebbsfleet. It is possible that all new estates can follow this model, although there needs to be consideration given to what constitutes a critical mass of homes in a project.

Replacing existing copper is altogether different, obviously in part because of the scale of the task. It will take a great deal of time to achieve - perhaps 10 years - so obviously delays getting started are a concern. But, there is another angle to consider first.

There are a small number of homes for whom even 1Mbps is a pipe dream today. There is an excellent response to the consultation by Stephanie Northen which says it far better than I can. The Digital Divide should not be thought of as an issue at the margins. Of course, Northern's case is extreme, but it should be noted that 20% can't get 1Mbps today and 40% can't get more than 4Mbps (source: BSG Pipe Dreams Report).

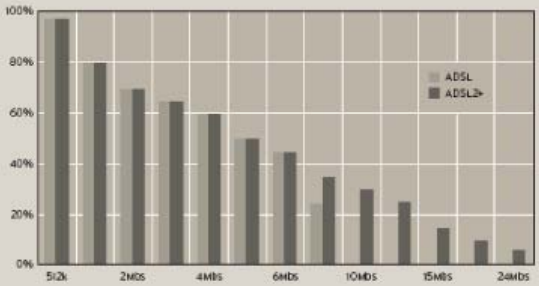

Figure 11: Peak rate access speeds of ADSL and ADSL 2+ available to UK households

This country need to address The Digital Divide first, before the need for speed at the top end. There needs to be a clear understanding of what basic services everyone has a right to be able to receive – either Ofcom or Government need to set benchmarks and work out policies to achieve these targets.

Conclusion 1.1

We need to level the playing field before we extend the divide further. Everyone should have access to 2 Mbps before anyone has access to 100 Mbps.

The Need for Speed (or Capacity?)

10 GB divided by 642 GB = 1.6%. We use less than 2% of the capacity of a 2 Mbps loop.

Of course capacity is not the same as speed and this is where the problem lies. There is an obsession with headline speed because it is one of the few “marketing features” in broadband today. We jumped to 512 kbps then 2 Mbps; we wanted 8 Mbps then 24 Mbps. Now we want 100 Mbps… But do we know why?

What is so big and so live and so important that you need 100 Mbps there and then? With 100 Mbps, you can send / receive 32 Terabytes a month. We certainly don’t need that yet.

Absolute Capacity

Headline Speed	GB per Month	Hours of 1080p	% of Avg Monthly Viewing	Hours of 720p	% of Avg Monthly Viewing	Hours of Standard	% of Avg Monthly Viewing
512 kbps	160	36	34%	57	53%	515	479%
1 Mbps	321	73	68%	113	106%	1,029	958%
2 Mbps	642	146	135%	227	211%	2,059	1,915%
4.6 Mbps	1,476	335	311%	522	485%	4,736	4,405%
8 Mbps	2,566	582	542%	908	844%	8,236	7,661%
20 Mbps	6,416	1,456	1,354%	2,269	2,111%	20,590	19,153%
50 Mbps	16,040	3,639	3,386%	5,672	5,277%	51,474	47,883%
100 Mbps	32,080	7,279	6,771%	11,345	10,553%	102,949	95,766%

Note: average monthly viewing is per person. Capacity is per connection (household)

File Size in GB	1 Hour
1080p	4.41
720p	2.83
Standard	0.31

Fill Factor

		Average Speed							
Average Usage GB per Month	Hours of 1080p	512 kbps	1 Mbps	2 Mbps	4.6 Mbps	8 Mbps	20 Mbps	50 Mbps	100 Mbps
2	0.5	1%	1%	0%	0%	0%	0%	0%	0%
5	1.1	3%	2%	1%	0%	0%	0%	0%	0%
10	2.3	6%	3%	2%	1%	0%	0%	0%	0%
20	4.5	12%	6%	3%	1%	1%	0%	0%	0%
50	11.3	31%	16%	8%	3%	2%	1%	0%	0%
100	22.7	62%	31%	16%	7%	4%	2%	1%	0%
200	45.4	125%	62%	31%	14%	8%	3%	1%	1%
500	113.4	312%	156%	78%	34%	19%	8%	3%	2%
1000	226.9	623%	312%	156%	68%	39%	16%	6%	3%

Note: average viewing is 100-120 hours per person per month (BARB). That is 440 to 530 GB per month at 1080p

Furthermore, local loop speed is only one of the factors in the throughput – which is surely what we are actually talking about. Speed is important but it degrades with distance as the chart from Akamai shows (source Telco 2.0). There is no point in having 100 Mbps access unless you can throughput at 100 Mbps, so you need to consider where content is physically distributed from and place that closer to the user.

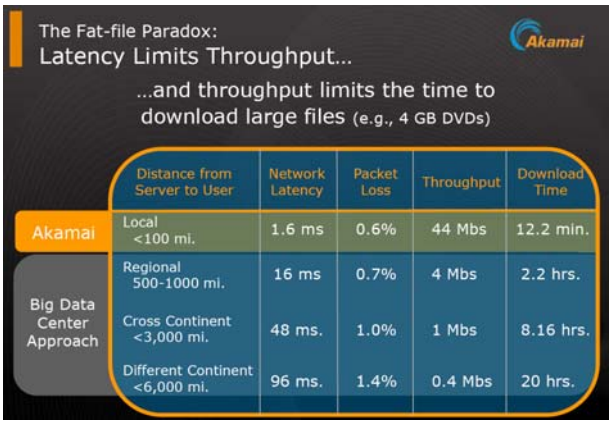

	Distance from Server to User	Network Latency	Packet Loss	Throughput	Download Time
Akamai	Local <100 mi.	1.6 ms	0.6%	44 Mbs	12.2 min.
Big Data Center Approach	Regional 500-1000 mi.	16 ms	0.7%	4 Mbs	2.2 hrs.
	Cross Continent <3,000 mi.	48 ms.	1.0%	1 Mbs	8.16 hrs.
	Different Continent <6,000 mi.	96 ms.	1.4%	0.4 Mbs	20 hrs.

If speed is important, we should first look to lower latency and ask what the options for doing that are. The falling cost of storage is highly relevant to the broadband market as that allows more content to be deployed deeper into the network, shortening routes and lowering latency.

Tromboning & Routing Inefficiencies

Bringing Level 3 into the backhaul network and even eventually into the NGA should be considered as one of the options because it would help minimise use of the backhaul and core networks and thus improve performance. The first time a connection is routed is when the Layer 2 data sees a Broadband Remote Access Server (BRAS) which are typically deep within the network. This long backhaul is the cause of the tromboning effect.

DSLAMs convert Layer 1 data to Layer 2 and a BRAS coverts Layer 2 to Layer 3. Until it reaches the BRAS and becomes Layer 3, the data is heading down a tunnel with no exits. This tromboning effect is an inefficient use of the Access and Backhaul networks.

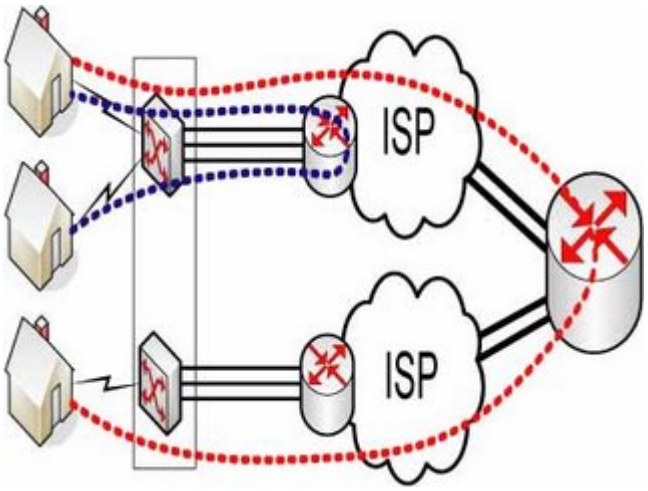

In the diagram above the blue line represents a session between two users on the same ISP. You can see the tromboning on the backhaul circuits. Similarly when the session is between users of different ISPs, there is even greater tromboning because the traffic has to go through a peering point.

Bringing routing deeper into the network makes it ever more like a grid - which might be what we need to get the best out of technologies like peer-to-peer. Shorter distances means lower latency and lower lower jitter, and a much greater scope to use the capacity on the fixed price, uncontended portion of the network at zero extra cost.

With DSL[1], the local loop is dedicated capacity between the exchange and the home. It costs a set rate that doesn't vary based on how much it is used. This is not the case at Layer 2, because more usage means more circuits and that means more cost.

Conclusion 1.2

The elimination of tromboning and the localisation of content reduces network load and latency. This leads to a more complex network, but one with better the performance.

[1] It is worth noting that cable and optical networks are typically configured differently and aggregate traffic in a ring rather than using a point to point topology like DSL from the exchange.

Alternatives have not been explored
We need to question how much bandwidth we need. Bandwidth is two things: speed and capacity (speed x time). While there is a speed problem constraining certain applications, there is not a capacity problem on the access network.

We use less than 2% of the DSL access network's capacity.

Furthermore, the cost of local storage is falling very quickly and predictive technology is improving that could enable a user to receive their files before they need them. This could enable much more efficient use of the aggregate capacity installed today by filling up off-peak troughs on the backhaul networks.

It has to be considered that the satellite and Freeview broadcast models will probably be preferable for mass market content for a good time to come, particularly when, as with the Sky+, local storage is incorporated at the edge to provide catch-up and on-demand services to broadcast. Broadcast will have a long-term advantage in live events like News and Sport.

The falling cost of storage means that the CDN can be extended into the home. Catch-up TV and other non-live media publications can use this to deliver their products while the network is quieter. In most cases, the time of delivery is flexible, and can be done in advance of publication by pulling down the files required ahead of time.

The internet should never need to achieve the encoding speed of a file – we don't need 8 Mbps internet to watch 8 Mbps HD. It can instead deliver any necessary files at any time of the day or night and it has the broadcast channels for live events.

Conclusion 1.3
We may not need more than 10Mbps (4 TBytes / month) for a considerable time and perhaps never for applications like HD TV. If its is watched live, it is surely most efficient to use broadcast, while if it is anything other than a live event, it can be sent in advance. In any case a file can be stored locally either off the broadcast or the internet feed.

Multicast
For whatever reason, Multicast is not enabled on BT's core network. This limits the use of a technology that offers an alternative to more capacity.

Why has it not been used? Is it because no attention has been paid to how multicast interconnection might work and how an aggregated network of active access components could deliver this, perhaps alongside a network of passive access components?

Multicast is another natural monopoly layer to the network like fibre (and other passive components). Unfortunately in a competitive market such as we have created, we lose the ability to cooperate for the greater good – multicast is one example of how this can be highly inefficient.

Conclusion 1.4
Consideration need to be given to natural monopolies and whether it is necessarily a bad thing to **not** have competition. It may be that by trying to encourage competition (so that competition delivers investment), we may actually be killing any potential return on the value of the investment.

Question 2: Do you agree with the principles outlined for regulating next generation access?

This is a hard question to answer directly because the SRT principles are to a large extent "motherhood and apple pie". What is missing is the detail as to how these principles will be employed.

i) promote competition at the deepest levels of infrastructure where it will be effective and sustainable;

Um, yes ok. But what do you mean? What is this level of infrastructure, how do you define effective and how do you define sustainable? Effective and sustainable to who – eg. someone who already has X, Y and Z or a new entrant?

ii) focus regulation to deliver equality of access beyond those levels;

I think you need to specify the levels. I think there should be regulated equality of access to the following network elements: backhaul, local loop, space in exchange, space in cabinet and local hands because these are natural monopoly resources that any operator would need to build their own service.

Others might agree with the overall principle, but may disagree with some of these level: the devil is in the detail, and I think this consultation needs to base its principles at least one level of detail further down if it is not to gain a meaningless agreement.

iii) as soon as competitive conditions allow, withdraw from regulation at other levels

Ditto. What do you mean by "as soon as competitive conditions allow". I think you also need to describe the process for end-of-lifing a regulation, much as you would sunset a product. How is this requested, contested and implemented?

iv) promote a favourable climate…

Hard to say no to this one…

v) accommodate varying regulatory solutions for different products and where appropriate different geographies

This one on the other hand, is hard to agree with. There should be **consistent** regulatory solutions, ***not*** *varying* solutions. One of the biggest barriers to investment is uncertainty over alternative technologies and the treatment they get compared to the one that ABC co is interested in. Spectrum usage and future licensing of the analogue TV band threatens to offer an alternative to fixed NGAs depending on how/where it is implemented.

It also seems that licenses use can change, completely altering the basis on which investments either in this or another technology were based. The change in PCCW's Now license is just one example.

Every different technology should be treated equally. If there is an open access obligation on one, then there should be on all others or the situation is somewhat absurd. This may be uncomfortable to migrate to, but is surely a base principle that helps us remove uncertainty.

Such treatment should be technology agnostic, but the base principles in rural Wales (for all technology) need not be the same as the base principles (for all technology) in central London.

Question 3: How should Ofcom reflect risk in regulated access terms?

I think that the anchor product approach is correct, but I agree that definitions need to be tight. There are fundamental weaknesses in both other suggestions.

Price setting is too open to obfuscation – what is a cost, how is it attributed by product and how does a fixed price investment get reflected in a monthly recurring, usage based price? BT are the world-leaders in this… Cost-based pricing has turned into a game of confusion and hidden charges and that is bad for business everywhere.

Removing the burden to justify the price and going with a free market approach to monopoly pricing would remove much of the gaming that goes on, but is unlikely to lead to an attractive result. Price skimming is a clear risk of this approach.

One major charge I would like to see is the elimination of usage (or capacity) based charging on monopoly components because it is a distortion and one of the major weaknesses of the wholesale market. The backhaul cost per end-user increases to the wholesale customer where it remains flat for the wholesale monopoly provider because backhaul is a usage based price even though the build cost is largely fixed (the dig and the passive components).

At the same time, the ability of service providers to manage the volumes that drive these bills is limited to throttling. There are a number of alternative options, some of which require monopoly components like collocation space deeper in the network. More needs to be done to help ISPs manage the impact of usage on their costbase.

Conclusion 3.1
The cost of components required to build a competitive infrastructure should be fixed (per customer) or the monopoly carrier needs to provide tools that the service provider can use to manage usage – hosting space deeper in the network and traffic shaping ability – to help manage variable backhaul costs, on a similar anchor product basis.

Government Contracts
Major contracts like the award of local government network contracts can go a long way to provide the basis for a profitable investment in the network. This raises a multitude of issues that need to be considered.

Should these requirements be brought to the front of the rollout queue in order to provide a revenue base for the NGA network? Should these contracts default to the anchor tenant as part of a set of “soft guarantees” to the ROI? How can the effect of these indirect subsidies be measured and should they be controlled?

Conclusion 3.2
The award of government contracts that use a potential NGA network need to be on top of people’s minds when assessing the issue as there are a number of very significant implications of policy here.

New Builds
It is important to note different treatment of NGA versus LGA (Last Generation Access) applies only because the NGA replaces the LGA before the end of the LGA’s natural economic life. This is not the case with New Builds where there was never an LGA

Conclusion 3.3
Any preferential treatment of NGA investment should only apply where it replaces an LGA.

Question 4: Do you agree with the need for both passive and active access remedies to promote competition?

Passive access is defined in the document as direct access to physical elements like ducts, fibre or copper loops, while active access refers to products based on the physical elements (the example given is IP Stream, but also includes Ethernet handover products).

For me, this is a weakness of the framework. The active product described in the consultation document should be broken down further into:

1. the services which hand over at Layer 2 – eg. point to point data, backhaul
2. the services which hand over at Layer 3 – eg. IP Stream

Because passive capacity is physical and the Layer 2 product virtual, I don't believe they can be grouped together. But, I do not believe either that Layer 2 products can be analysed in combination with a Layer 3 option because there is the potential for genuine competition at Layer 3 active products for perhaps 75% of homes.

This is not the case with Layer 2 competition because there are very few operators with their own passive network backhaul to the exchange. The layer 2 service is a critical component in LLU, alongside a number of passive access products, but the passive capacity needs to be shared between service providers and that is where Layer 2 offers an intermediate options to layer 3 access.

Furthermore, layer 3 (routed) products can bring very different economics than layer 2 (point to point) services, especially in combination with deeper routing in the network and locally cached content described earlier in the document. It must be possible to provide a wholesale alternative offering that includes a fully managed UK IP network. This is certainly a different market from the provision of an Ethernet-based backhaul circuit.

Conclusion 4.1
There are three subsets of the market: passive (layer 1), data link (layer 2 and L2TP) and active (routed layer 3) components. LLU consists of passive and data link components while IP Stream today is an unrouted L2TP service so it would be a data link service. There is much scope for BT Wholesale to offer a routed Layer 3 service in the future as an alternative to IP Stream and this must be considered.

IP Stream Lockdown is a Major Market Distortion
For a period of time, IP Stream made it all happen. Now though, capacity based charging has flipped the business on its head – operators make more money the less a customer uses their services. The incentive to ISPs is to throttle growth in usage.

While cost per unit throughout the industry drops with increases in usage (because of scale gains), IP Stream has been locked in an uncompetitive position so as to enable the creation of a new generation of unbundler.

The time has come to rebase the prices on IP Stream to a competitive level because the creation of LLU competition in cities had hurt the price of IP Stream access in the villages. The creation of LLU has widened the digital divide, not just by improving service in the cities but by making the service worse for those who cannot get LLU.

Conclusion 4.2
IP Stream is vital element in a competitive market and services based on it need to be able to compete with LLU but this competition is distorted today. IP Stream should be allowed to resume competition, even if it at the expense of LLU.

IP Stream, Cross Subsidies and USO

The problem for IP Stream competing with LLU is that IP Stream needs to blend its costs between where it does and where it does not compete with LLU. A model whereby IP Stream is cheaper in locations where it has competition is a clear exploitation of the monopoly where it does not have such competition. There needs to be one price.

But this raises the key question: why should rural users be subsidised entirely by BT IP Stream users in the cities? It would seem that there should be a more equitable contribution to universal service.

Conclusion 4.3

All LLU circuits combined with IP Stream should contribute equally to deliver a basic universal service where there is no LLU and competition to make that happen.

Should BT Retail do LLU?

The perversion of the BT Retail / Wholesale split is that BT Retail is locked to Wholesale and the base IP Stream network. While it would clearly be cheaper to BT Retail in isolation to do what every other major player did and LLU, because it and Wholesale are part of the BT Group, it did not.

So far, this is an internal BT issue, but there is an external effect because the prices that BT Wholesale customers pay are locked to those that BT Retail pays.

Clearly the choice for retail to unbundle is probably more costly to BT Group, but it highlights the artificial position of BT Wholesale and the prices that they offer. Charging excess profits to BT Retail is easy because they have no competition for that account and these prices can then also be applied to other monopoly markets by Wholesale.

Conclusion 4.4

The role of BT Wholesale and its supply relationship with BT Retail needs to be investigated and boundaries set. Does BT Retail have to treat BTW as its network supplier? Does BTW have to offer other ISPs an equivalent product to what they offer BT Retail? Everywhere or just where there is a monopoly?

Pricing Principles

This whole area comes down to a question of what pricing principles are adopted. These principles will determine the competitive state of the wholesale markets.

Conclusion 4.5

Passive elements and the Layer 2 components that are necessary in LLU should be available at fixed prices that are certain for the period during which the market is determined to exist. The same principles should be applied to a number of additional natural monopoly components like collocation space and local hands.

BT Wholesale should buy these components at the same rates as 3rd parties and should be free to set IP Stream prices as they see fit as long as a) the price paid by BT Retail is the same as prices charged to other ISPs, b) scale discounts are negligible, c) the price of IP Stream where there is LLU is the same as the price where there is not

In order to achieve c), there needs to be a Universal Service Charge applied to all broadband lines and distributed to subsidise the development of infrastructure and price equality in areas where LLU competition will not deliver this improvement.

Question 5: Do you consider there to be a role of direct regulatory or public policy intervention to create artificial incentives for earlier investment in NGA?

Emphatically not. The DSL evolution is only four years old and it cannot be possible to conclude that we need intervention or artificial incentives for earlier investment. Doing so would damage the value of investment in DSL and perhaps perversely even damage the rollout of NGA because of uncertainty about where, when and on what basis subsidy would be awarded.

Conclusion 5.1
At this stage in the market, regulatory and public policy needs to concentrate on the edges of the existing market – mitigating the Digital Divide – and not on widening it by making the rich richer.

What are the targets?
Ofcom needs to start addressing the issue or fair access to existing technologies for all before it starts to address the issue of access for a limited few to an NGA. In fact, there may even be a case to deliberately delay an NGA investment and to channel the funds to level the existing playing field before the “haves” again accelerate off into the distance.

There is a role here for either Ofcom or for Government to set out what this basic universal service requirement is. This should be done and those excluded dealt with before effort is put into NGA.

Conclusion 5.2
There should be a common baseline (speed, capacity, price) that everyone can expect: a universal access product available for an affordable price (TBD). Everyone should have access to a baseline.

NGA Policy
There are however public and regulatory policy positions which are unclear and are perhaps confusing the case for NGA. Unfortunately, this consultation is one of those initiatives clouding the horizon because it raises the possibility of public subsidy.

This consultation, having asked the question needs to put that question to rest in the final response – what are the direct and indirect subsidies that will and won’t be factored into the NGA investment case? How will technologies that compete with NGA, like Wireless and LGA also be regulated?

Conclusion 5.3
There should not be any artificial incentives in areas where LLU has proven that competition is viable. There needs to be a very clear statement on this subject, including where subsidies will be considered and where they will not and how indirect subsidies (as discussed earlier) will be applied.

Step by Step Upgrades

It's a cliché but we should be aware of trying to eat the elephant whole. NGA is simply too big a question to put into one bucket so the components need to be broken down and a passage plotted that takes us to where we want to be.

There is simply no point in upgrading access capabilities unless other elements of the network are in place to cope – you just shift the bottleneck and spend a lot of money for little return.

There are choke points in the infrastructure today, but these *are* being alleviated by investment. It was not long ago we were referring to a crisis in backhaul, but fibre is now being installed to connect even some very small exchanges (sub 1,000 homes).

Conclusion 5.3

Before we need access, we need a number of key components: we need data centres with power and we need these much deeper in the network. We need a commercial framework that takes care of the broadband incentive problem and we need stability.