

DCMS AND HM TREASURY DIGITAL COMMUNICATIONS INFRASTRUCTURE STRATEGY

RESPONSE BY COLT TECHNOLOGY SERVICES

CONTENTS

Contents.....	1
1 About Colt.....	2
2 General Comments.....	3
3 Commentary on Scenarios	5
4 Regulation.....	6

1 ABOUT COLT

Colt is Europe's information delivery platform, enabling its customers to embrace the changing landscape of IT and communications so they can deliver, share, process, and store all of their vital business information.

Colt aims to inspire customers to think differently about the way they tackle their core business and technology issues. Colt's information delivery platform combines high performance, end-to-end control of integrated compute and network solutions with an agile and responsive approach that delivers an integrated experience across Europe and beyond.

Today, Colt runs a 23 country, 83,000 km network that includes metropolitan area networks in 42 major European cities, with direct fibre connections into over 20,000 buildings and 20 Colt data centres.

Colt enables its customers to deliver, share, process and store vital business information by bringing together 3 key elements:

- Pioneering European Ethernet and IP networks that seamlessly connect over 100 cities and achieves the industry's highest standards in performance, latency and security.
- Significant IT infrastructure and services across Europe, with 20 state-of-the-art data centres with tens of thousands of devices under management.
- Extensive expertise in creating integrated IT managed services, networking and communication solutions.

Colt is continuing to invest heavily in its ability to deliver integrated network and IT managed services. Colt is also helping to lead industry standards and certification for cloud services.

2 GENERAL COMMENTS

Colt Technology Services welcomes the opportunity to reply to the government's "Digital Communications Infrastructure Strategy" consultation. The essential premise of the consultation, that a strategic approach is needed to address the challenges posed by the unprecedented pace of technological change, is surely correct.

Colt agrees with the paper that market led solutions are preferable to government intervention. However, regulatory intervention is and will remain necessary to address entrenched market power. The type of intervention should be focused on the precise source of that market power, namely, the physical infrastructure, which is and will remain the largest single barrier to fibre network deployment. Moreover, Ofcom and the UK government need to refocus on the principle of "access at the deepest feasible level". Neither the services, nor the markets assumed to be addressed by those services, should be prescribed or circumscribed by regulatory intervention.

The focus of the DCIS review should be on allowing fibre connectivity to develop in the most efficient possible way, where possible by leveraging existing infrastructure. Today, the market penetration of high-quality optical fibre networks can be described as patchy at best. Many areas are well penetrated but often there is no coverage at all in locations within or adjacent to those areas, owing to the high cost of the civil works necessary to install the requisite network extensions.

As a leading international provider of business connectivity services, our comments are predominantly focused on the business sector. Business connectivity is relevant at two levels:

- As a service provided in its own right to business customers. Like residential customers, business customers are keen to ensure that the structure of the market is adequate to meet their communications needs in the years to come
- As a component in the value chain for services ultimately targeted at residential users.

It is important that any strategic decisions covering the electronic communications sector as a whole are cognisant of the ways in which business and residential supply and demand interact with one another. Business communications affect and are affected by developments in residential infrastructure, technology and usage models. Increased demand from residential customers and changing patterns of use create demands that must be met by the business sector. Equally, business and residential networks enjoy economies of scope: ducts and trenches deployed for one sector affect the costs of the other.

The degree of overlap should not, however, be overstated. There are substantial differences between business and residential needs that need to be borne in mind. For example, business customers are typically much more demanding with regard to service features such as symmetry, availability, technical performance, scalability, wide-area networking, SLAs and times to repair, etc.

A further factor that must not be ignored is the fact that, while fibre in residential networks is relatively young, in business networks it has been deployed for many years. This allows for greater scope for leveraging existing business communications infrastructure than is often realised.

As regards business communications infrastructure, the key point to understand is that as supply and demand patterns change, markets are increasingly fragmenting. On the demand-side, more specific applications are creating demands for more specific technologies, creating widely varying demands

on local infrastructure. The ability of the supply-side to meet these changing demands varies enormously, even at the micro-geographic (or sub-postcode) level. No longer is it possible to rely on a “one-size-fits-all” solution to serve all customer demands. Today, a multitude of new technologies are on offer, each of which places differing demands on communications infrastructure. These include:

- The growth of cloud computing and remote data services, in turn driven by demand for myriad online services including multimedia content, health, financial services etc. Many aspects of our lives are coming online. Storage capacity and bandwidth need to become scalable, secure and protective of personal privacy.
- The increasing dispersal of business sites requiring integrated connectivity. This is seen in the growth of pan-European businesses with sites in multiple jurisdictions, requiring seamless connectivity. It is also seen in the growth in home and mobile working, requiring high-quality access to remote data and applications.
- The growth in demand for data transmission capacity, driven by the increase in access bandwidths delivered by FttX and 4G. This is particularly acute in mobile backhaul.
- The growth in machine-to-machine (M2M) communications. We include in this category scenarios where remote machines are monitored or controlled by central servers, and White Label network services, allowing service providers to provide apps or connected devices to users (for example health monitoring apps available on portable devices feeding into a health service).

These market trends illustrate what the business-to-business communications sector needs to deliver to the entire sector for the market to thrive.

The inadequacy of a “one-size-fits-all” solution serving every requirement of all sub-sectors of all markets should be apparent. One access regime, foisted on one infrastructure owner providing connectivity to all market participants, will not do the job. Indeed, there is not merely one market for electronic communications, even in the business sector. To the contrary, there are many markets, differing enormously in their requirements for quality of service (e.g. resilience, availability, latency, transparency, etc.).

To harness all the benefits on offer from technological progress, above all, the UK requires new investment. Not just any form of investment but rather, investment targeted at specific geographic and product market niches, focused on the individual needs of customers involved in driving the changes described above.

To meet the increasingly specific and demanding requirements of new technologies, the market requires a more differentiated suite of connectivity options than is presently available. Not all requirements are the same: a distributed computing system requires low latencies but not necessarily high bandwidths. A video conferencing system requires high bandwidths but not necessarily low latencies. An e-commerce platform requires high availability, but not necessarily high bandwidths or latencies. A mobile network reconfiguring its urban network from a few macro-cells to a larger number of pico-cells requires – above all – short delivery lead-times and the availability of connectivity at the granular, micro-geographic level. The fact that new technologies have widely differing performance requirements underlines the need for differentiation.

As the performance requirements of business-to-business connectivity solutions become more demanding, so does the need for competition and choice. Large national networks cannot fulfil all these requirements (either at retail or wholesale levels) and in the absence of genuine competition, have little incentive to try. While such players require incentives to serve the (capital intensive but relatively undifferentiated) requirements of residential users, this is cold comfort to business users, for whom residential connectivity options are inadequate.

For too long, policy towards business connectivity services has been hived-off into a category that regulators like to call “leased lines”. “Leased lines” is an old-fashioned term originating from an era when business connectivity options were much less diverse than they are today. Neither sellers nor buyers consider themselves as participants in a market for leased lines. Unfortunately, the term seems to have become embedded into the regulatory lexicon in a way that is entirely detached from how the market actually operates. This misunderstanding leads to errors of policy, particularly where the logic behind a policy decision is based (perhaps tacitly) on a belief that all “leased lines” are essentially the same.

3 COMMENTARY ON SCENARIOS

As a general remark, we believe scenarios need to examine in greater detail the “back-end” services that businesses require in order to deliver the “front-end” services required by consumers (and other businesses). Typically, these are defined in terms of technical performance requirements such as latency and jitter, availability, resilience, time to repair, installation time, performance monitoring, performance management, wide area networking, frame synchronisation and cloud infrastructure. The key issue here is the availability of a differentiated suite of products, at different price points such that more highly performing solutions are available at higher price points, and the more basic solutions are available at lower price points. And furthermore, it is important that the full suite of options is available in as wide a variety of geographic locations as is possible.

Scenario 1 describes a growth in IoT, while Scenario 2 describes a growth in machine-to-machine communications and home working. One thing these developments have in common is a requirement for the supply-side to be able to respond at the micro-location level. In other words: the 99% availability of (say) fibre connectivity at a postcode level is cold comfort if the requisite connectivity is not available in the precise location required. This gives rise to the possibility that geographic dimensions of the “digital divide” will become more granular than they are today. To ensure that this does not happen, it is important for the supply-side to be able to deliver a wide variety of connectivity methods in as wide a variety of locations as is possible.

This is one reason why we believe that the development envisaged in Scenario 1 whereby “the access network will be delivered over a common physical layer, used by many operators with competing services housed in the same physical infrastructure”, (with the bundling of services provided within the aggregation layer of the network) should be regarded as an unfortunate outcome. In particular, the treatment of “the” access networks as though it were a single, homogeneous entity would discourage operators such as Colt, who are actively engaged in innovating different access network topologies and technologies, from bringing the fruits of their innovations to the market. Furthermore, it would reduce the degree of choice available to different customers (with different needs), and further cement geographical dimensions of the digital divide into (for example), regions with copious alternative infrastructure, which would yield a wide variety of high (and lower) performing solutions and those in which only the basic services are available.

In terms of the granularity of geographic variations in availability, the existence (to a greater or lesser extent) of a single supplier should be strongly discouraged. In such a world, the availability of connectivity (of the right type) in any given location would be dependent on where the location happens to be, relative to the topology of the network supplier. A network architecture optimised for one purpose (for example, residential connectivity) of necessity cannot be optimised for all purposes. It is therefore important that the supply side is able to respond with a variety of network topologies, architectures and technologies as required by the market, where possible by making efficient use of civil infrastructure already in place.

As a general point, we disagree that (as stated in the Scenario 2 commentary) “it will be left to the device and application layer to make the best use of the infrastructure available.” In our view, this greatly over-simplifies the extent to which applications have specific requirements in terms of the quality, architecture and technology inherent in the infrastructure. Related to this point: we believe that access technologies will remain siloed and that the de-siloing of access technology should not be an objective. This is because of the heavy inter-dependence between the access technology and the application.

In terms of scenario 3, we disagree that connectivity and coverage will be a hygiene factor any time in the near future. Different standards of connectivity are needed to meet different types of application. We would argue that the standard of connectivity available will become an increasingly important factor in determining business location. We do, however, agree with the statement that fibre access will be available at a more granular level than present, but that in order to facilitate this, it will be necessary to open the market to smaller, niche market providers, who are able to use locational advantages in expanding their network (such as proximity to existing infrastructure). This is one reason why it is important that the position of any given location (such as a building or street) relative to the dominant operator’s network hierarchy should not be the determinant of the services that are or are not available. The “virtualisation” of the access network is not a development that we envisage over the timescale considered. The differential performance of different network connectivity methods is – and will remain in our view – too important a factor for customers to be able to abstract from it.

We do – on the other hand and in line with scenario 3 – envisage a scenario in which the access network is more dynamic than today, but that is partly dependent on the policy framework allowing it to be so. In particular, it is important that the policy environment is more active in encouraging the right micro-local decisions to be made, which cumulatively add up to a more dynamic access environment across the UK as a whole.

4 REGULATION

As the consultation notes, Ofcom’s duties as set out in the Communications Act 2003 include:

- Encouraging investment and innovation
- Encouraging the availability and use of high-speed data services

The consultation also correctly notes that Ofcom’s professed policy since the strategic review has been to “encourage competition at the deepest level in the network where it is efficient and economically sustainable to do so.”

The problem, however, is that Ofcom has not followed this policy in the business connectivity market. This has had the effect of limiting investment and innovation in business infrastructure and thus denied the market the opportunity to invest efficiently, leveraging the interdependence between business and residential networks, thereby bringing the fruits of innovation to the widest possible community. In effect, we characterise Ofcom's policy in the business connectivity market as "shallow-remote", where a preferable approach from an infrastructure investment point of view, would be "deep-local".

"Shallow vs. Deep" is a product dimension that refers to the scope for individual service providers to provide their own differentiated products to meet the needs of specific classes of customer. In the UK, the vast majority of business connectivity is based on wholesale Ethernet services provided by Openreach. The main exception is Central London where substantial alternative infrastructure exists. In Central London, the competitive environment is "deeper" in that alternative infrastructure providers can offer a rich set of variants (at different price points) compared with the dominant provider's nation-wide wholesale Ethernet offer.

"Remote vs. Local" is a geographic dimension of competition. "Remote" competition occurs where a competitor accesses a wide geographic area by means of a small number of interconnect points. The competitor mainly uses a third party to handle the physical delivery of service into the customer's premise. Remote models of competition are not very scaleable: after gaining one customer in a given locality, the cost of serving a nearby second customer is approximately twice as much. The geographic dimension of competition correlates with the product dimension. Where the product dimension of competition is "deep", alternative operators compete by offering a direct (local) connection to their own fibre. By contrast, where it is shallow, there is relatively little benefit of geographic proximity to the customer (except possibly as a price arbitrage). This means that networks tend to be less well built out and therefore, remote from the customer.

This is undoubtedly a contributor to the current state of affairs in which statistics on average levels of availability and coverage mask substantial geographic variations. It is a contributor to (for example) the poor availability of fibre connectivity in business parks and business districts that are adjacent to well-penetrated areas. It contrasts with a preferable state in which competitors have the opportunity to exploit gaps in BT's product set and/or geographic coverage, and develop them incrementally over time.

4.1 PASSIVE INFRASTRUCTURE ACCESS

The greatest single policy initiative that the UK could implement to improve the quality of communications infrastructure is to open up the passive infrastructure market in business connectivity. Using EU benchmarks of duct and dark fibre access, the cost of network construction via regulated passive access is as little as one tenth of the cost of network construction using new digging. Colt has evaluated generic business cases for investment in new fibre networks in most major UK cities and the implications are clear: with passive infrastructure access the business case for new investment works, whereas without it, it does not work. Colt is highly experienced in constructing new networks using regulated passive infrastructure access and has done so in several cities outside the UK. Benefits that liberating the market in passive infrastructure would bring include:

- Better quality infrastructure to currently under-served regions (for example, many business parks)

- A richer set of products available to local businesses. This would spread the availability of higher performing products to regions outside central London, and open up the possibility of new market niches serving customers who require better quality options than are available using residential product sets, but who do not require the premium priced classic business connectivity products.
- Better availability of infrastructure at the micro-local level, facilitating small-cell and IoT connectivity
- Better economics of incremental network expansion, as each instance of network expansion includes an option value for subsequent network expansion
- A richer set of options for customers providing residential services. Fibre networks originally constructed for business connectivity reasons could serve as a wholesale platform for the provision of local residential services (both access and backhaul)

A recent report by the Federation of Small Business (FSB)¹ shows that this constituency urgently needs better connectivity. The report singles out business parks as being particularly badly served by modern communications infrastructure. In our view, this is because business parks fall between two stools, typically lacking the volume to justify FTTC rollout, while lacking the budget to afford premium business connectivity products. This is precisely the type of situation where passive infrastructure access is able to provide answers. Business parks need an “in-between” product that is better quality than residential grade FTTC, while cheaper than premium products. In practice, this can be achieved by putting in place a 10 Gbit/s fibre ring with an individual connection to each building (or cluster of buildings). The connection would be contended at a ratio of (say) 5:1. This would offer FTTP directly to local small businesses and would offer a better quality connection (by far) than residential grade FTTC while being cheaper (by far) than a premium business product.

Colt does precisely this – providing FTTP to business parks – in continental Europe. It does so because it is able to use regulated duct access to implement the local fibre ring, and regulated dark fibre access to connect the fibre ring to its long distance network. This type of operational model is only possible where regulated passive infrastructure access is available.

Furthermore, we would argue that a dominant player such as BT does not have the right incentives to pursue this type of model of its own volition. The ability for alternative entrants to use raw network components to provide a platform for innovation and differentiated products targeted at specific market niches, is the type of disruptive competition that a dominant player is most keen to avoid. This is because their primary incentive is to maintain price discrimination by implementing a sharp demarcation between residential grade products and premium products. They do so by offering “one-size-fits-all” wholesale products (on which they hope their competitors will depend) that are targeted at either one market segment or another, with nothing in between.

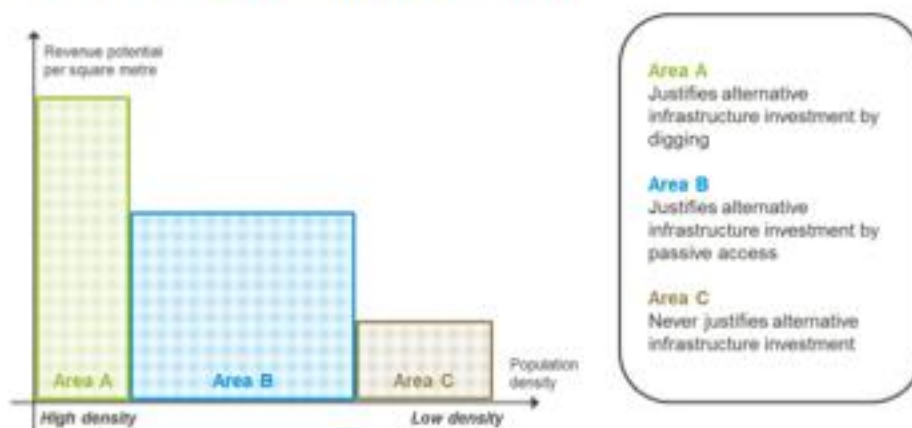
Considering its historic position as one of the leading innovators in communications policy and regulation, the UK has been remarkably cautious and slow to recognise the benefits of passive infrastructure access in the business connectivity market. This caution is difficult to understand, given the extensive experience available from other EU countries of this type of regulation. Colt understands Ofcom’s reasoning for this caution. Nevertheless, we believe it stems from an incorrect understanding of how the market operates and in particular, an incorrect belief that passive infrastructure access is only about “cheaper leased lines”, a view that itself is based on an incorrect

¹ <http://www.fsb.org.uk/policy/assets/FSB-The-Fourth-Utility.pdf>

belief that the business connectivity market is all (and only) about a single, homogeneous category of product that Ofcom calls “leased lines”. To the contrary, passive infrastructure access would unlock similar levels of innovation and investment – and more – in the business connectivity market as LLU has in residential markets.

Liberating the market for passive infrastructure would stimulate investment in “Area B” type regions as depicted in the diagram below.

Passive access and economy of density



Area A type regions – roughly corresponding to central London and other similar regions – justifies alternative infrastructure investment without further measures. Such regions already enjoy a rich variety of product options. Area B, however, which corresponds to the majority (by number) of business locations in the UK does not justify new network build under the current regulatory environment, but would do if a pro-competitive access regime were introduced.

Opening up the market for passive infrastructure access has been successfully carried out in several other EU countries. For Colt, it is now the default, business-as-usual method of deploying new network coverage and individual customer connections in France, Spain, Italy and Portugal. Implementation issues such as compatibility with existing price controls have successfully been addressed in other markets and there is no reason why this should not also be possible in the UK. Given that the UK is relatively late to the table in implementing this policy measure (as was the case with Local Loop Unbundling), this is a double-edged sword: although other markets have a jump-start on the UK in terms of reaping the benefits of passive infrastructure access, there is now a wealth of information on best-practice models and the “dos and don’ts” of implementation. For passive access to be fully effective, Colt’s experience is that it needs to incorporate at least the following characteristics:

- No limitation on downstream use. This is necessary for access seekers to avail themselves of the full economies of scope from network deployment. Business service providers such as Colt would be able to defray the cost of network installation by wholesaling fibre capacity to residential service providers, both for access and backhaul;
- Any-point-to-any-point connectivity. (For example, no restrictions such as “the ingress point must be outside the access provider’s exchange”). This is necessary to ensure that service providers with existing fibre assets can expand their networks incrementally, using existing

infrastructure, to adjacent areas. It is also necessary to ensure that the structure of alternative networks is not a mere copy of that already provided by the incumbent;

- No architectural limitation on access rights. In other words, this would prevent an access provider from denying access on the basis of whether the facility to which access is sought, is classified by the access provider as part of its access or backhaul network. Such classifications are arbitrary factors related to the access provider's own network architecture and should not be used to limit other network providers' ability to fill in coverage gaps, or coverage "not-spots".

If implemented in this way, passive infrastructure access would yield the following benefits:

- Stimulate the rollout of fibre-optic networks to business and residential markets
- Create a much more dynamic and responsive access market
- Reduce the number of fibre coverage "not-spots" within and adjacent to well penetrated areas
- Improve the extent of fibre availability to small and medium sized businesses, particularly in business parks
- Improve the availability of applications requiring access in specific locations such as mobile small-cells, IoT and M2M communications
- Allow greater product differentiation vis-à-vis BT's Ethernet leased line offer, including more "basic" services suitable for smaller businesses and higher performing services suitable for businesses with demanding requirements
- Allow greater penetration of innovative technologies and architectures in access networks
- Improve the availability of high-end connectivity methods in provincial cities
- Allow greater competition in SLAs, provisioning and upgrade performance
- Allow greater competition in pricing models
- Allow alternative infrastructure providers to internalise the option value of past investments, by reducing the cost of future investments