THE 'SUPERFIT' REGULATOR

MARTIN GEDDES describes how regulators can shape up to meet the realities of the digital age and give users broadband services that are fit for purpose

Telecoms policymakers everywhere face similar challenges. As networked applications become ever more critical to modern life, users are demanding that broadband services meet their growing needs. How can regulators develop the institutional and industry capability for managing increasingly complex digital supply chains?

I am based in the UK, so we can take the charter of Ofcom, the UK regulator, as an exemplar of the situation. Its remit for broadband and broadcast is encoded in the Communications Act 2003 (and elsewhere in government policies). It aims to:

- Develop UK digital infrastructure, making good experiences through investment and innovation
- Ensure retail services offer choice and transparency, by articulating fitness-for-purpose, thus allowing users to make rational trade-offs of cost and service level
- Ensure fair wholesale access for competitive markets, by instigating and refereeing BT's equivalence platform, with price regulation at boundaries
- Identify and manage economic bottlenecks as these shift, e.g. new markets for network capacity or content delivery vs old local loop monopoly
- Ensure efficient use of finite resources (spectrum, capital, labour), in particular to protect the weak (poor, disabled, remote, etc.) from exploitation.

There is much to celebrate in terms of the industry growing up to meet these essential needs. For instance in the UK, there has been a massive take-up of broadband, with widespread adoption of fibre to the curb (FTTC), and a very competitive mobile market. In many ways the UK is seen as a role model of competent and forward-looking regulation. All human endeavours are necessarily imperfect, and this is one of the better ones. But all regulators now face greater user demand from:

- Richer and more diverse applications, like IoT, telework, remote education, home healthcare, etc.
- Video moving to the internet, so you cannot any longer separate broadband from broadcast
- Increasing need for dependability, requiring a "safety case" as society comes to depend on applications working on a continued basis.

Then there are concurrent fundamental changes in the nature of the supply, which for the UK are:

- A growing requirement to move from FTTC to FTTP/H (premises, home) to improve reliability and capacity
- BT's wholesale arm facing economy of scale issues, especially as BT's retail arm has built its own infrastructure
- TETRA (terrestrial trunked radio) replacement with associated (and often unquantified) technical, commercial and political risks

● 5G is coming for capacity and capability, with challenging backhaul requirements for coverage
● Increasing demand for security and performance isolation between users and industrial internet uses (e.g. smart cities).

I see increasing commentator concern over the broadband experience, as the promised supply doesn't sufficiently satisfy future demand. Indeed, there is now a widespread user perception of broadband being the "unreliable utility", in stark contrast to its peers. You won't get many laughs at a stand-up gig talking about water or gas delivery, but broadband internet service is the target of jokes. My Australian friends tell me their unending NBN (National Broadband Network) comedy is increasingly tragic to live with.

The resulting stresses in the regulatory system exhibit themselves as symptoms like net neutrality, which is an expression of power battles between the edge and core over resource pricing and fairness. End users experience confusion over how to resolve service quality faults: is it the Wi-Fi, router setup, in-building wiring, protocol design, local access loop, ISP service, or the internet in general that's the problem?

I myself have faced this, struggling to "debug" my own poor home broadband experience, and I am supposed to be an industry expert! It has involved at least three items from the above list interacting to manufacture "unhappiness as a service".

As we can see, there are many facets to this challenge of digital experience quality regulation. The central problem is that there is no universally agreed framework to quantify the network, and to relate it to the user experience. The science of network performance and digital experience quality is immature.

For spectrum policy, the mathematics was all cracked in the 18th century, the science of electromagnetism in the 19th century, and the engineering (e.g. MIMO, multiple-input and multiple-output antennas) in the 20th century. Today in the 21st century we can focus on what policies to implement, not how to execute them.

Similarly, with computing, the mathematics was cracked in early 1930s, and the bulk of the science in the 1950s to 1980s. We don't worry about different regulation for Intel vs ARM processors, because we know they are fundamentally equivalent at some level. For data transmission, we have information theory from the 1940s: there is a strong theoretical basis with widespread buy-in to foundational concepts.

IT'S PACKET NETWORKING NOW

Packet networking is different. Yes, the statistical theory behind shared buffers was developed in the 17th to 20th centuries. But the theory and mathematics of complete distributed systems only became public knowledge in the 2000s. The formal definition of the resource "trading space" is very recent indeed. These conceptual underpinnings aren't yet in the textbooks or taught on university courses. I know, as I have been teaching them to the R&D labs of leading equipment vendors...

The obscurity and novelty of the core maths and science results in a diversity of metrics, measurement systems, and predictive models. They vary by network bearer technology, vendor and market. For regulators, this poses a danger of picking winners, and then regretting the choice. There is a lack of consensus in the regulatory community about both the problem and its solution.

Furthermore, there is an unclear mandate to solve this issue: regulation presumes the existence of the necessary fundamental concepts and tools. Go back to the start of the article: which of those mandates would legitimise Ofcom spending its limited resources with the Institution of Engineering and Technology to talk mathematical models of packet performance? Possibly none of them.

> **"There is now a widespread user perception of broadband being the 'unreliable utility'."**

As a result, regulators are forced to face a series of serious questions:
● How can regulators make progress at the practical level to measure and manage broadband QoE (quality of experience)?
● How can they ensure that digital infrastructure is fit for purpose?
● How can they develop the necessary human skills in digital experience quality?
● How can they redesign processes and policies to fit the maths and science?
● How can they acquire an adequate technical capability to get visibility through the end users' eyes?
● How can they anticipate and act with respect to managing quality at key regulatory, technical or economic boundaries?
● How can they engage with a systemic capability deficit that exceeds their own scope and remit?

My suggestion is that regulators should initiate a "superfit" transformation programme, akin to the "superfast" one of the past decade. The core purpose is to upgrade institutional capability in digital supply chain quality management. Note that I am not opposing or displacing the superfast model; it just has limits to its utility that we must transcend, since quantity is not synonymous with quality.

This is an issue of both organisational capability and national policy. It requires a coordinated response from all stakeholders in the ecosystem, including consumer groups, research institutes, and standards bodies, as well as the usual list of vendors, telecoms operators and experts in law and economics.

The timing is now right, as the maths, science and tools exist, albeit being immature. High-fidelity measurements can capture network performance in a user-centric way. It is possible to identify the dynamic performance effects due to statistical multiplexing, and separate them from the static architecture and configuration ones. There is also an inevitability to the transition, as networks become ever more dynamic – think SDN

(software defined networks), FTTH, 5G, distributed apps, smart antennas.

Not engaging with the "superfit" transformation challenge risks a legitimacy crisis. As vendors and service providers upgrade, you can foresee them outwitting the regulator — a bit like how high-frequency trading and derivatives have done in financial services. We may also see other regulators zoom ahead, exposing the laggard's credibility to act based on outdated models.

Alternatively, inaction may result in a severe user experience crisis as quantity fails to solve quality problems. Emerging markets may simply leapfrog established ones, as we have seen with cellular networks and services like mobile payments. For example, China "gets" infrastructure on an epic scale, and the ability of other countries to compete globally in digital services may prove illusory.

REGULATORY PRIORITY: EDUCATION

The regulatory priority is to educate oneself and key supply chain stakeholders: people first, then processes, and lastly clever new technology. This means regulators have to identify the internal and external game changers, and those who will help them be effective. These are the people who most need training in the core maths and engineering, and how to apply it.

This is a safe bet, as the science isn't going to change significantly: there's only one plausible answer. That said, regulators need independent advice and review of the core material, because we all require social proof as well as intellectual argument. This may involve formal contracted processes, as well as informal (e.g. from the Royal Society in the UK).

The subsequent step is to perform experimental measurements of infrastructure using these new techniques and their supporting tools. The experiments should be designed to meet regulatory objectives, which may substantially differ from commercial ones or those of the end user. These findings can then be promoted internally. This creates vital institutional understanding, and gathers feedback from the regulatory coal face.

Once this basic shared understanding is developed, the main planning task begins, to build a map of the upgrade journey. The needs a blueprint for organisational development, along the lines of the capability maturity models common for software engineering. The quality management systems theory is a well understood problem, and there are many other industries that have preexisting frameworks to draw on.

If we address the classic people, process and technology trio with a sensibly designed change programme, then we can progress from unmanaged chaos to a self-optimising system. The essential activity is to locate the present reality in that framework: we only change when we identify with what is, not what could be.

Nonetheless, we must also define the end state and ideal, so we know which way to head from wherever we are at. This will engage us in a paradigm change:

- From network-centric to user-centric
- From periods to instants
- From averages to distributions
- From separate silos to complete systems
- From state (quantity with quality) to stateflow (quantity of quality).

So where to begin? The kickoff action is to define a learning project to discover how to advance towards fully managed supply chain quality. This means finding a "corner case" which can be used to try the "upgraded" way of working. Good examples might be accessibility services for the disabled, or delivery of broadband to remote areas.

From this pioneer project, we can then construct the nucleus of the new organisational capability: the game changers need a "game changeable" context. This may require a temporary virtual organisation that draws together many different functions and competencies that cut across normal organisational boundaries. Inherent to the project is a process to disseminate the results: internally, to other regulators, and to the industry at large.

Once this internal baseline of engaged and educated staff exists, it is possible to roll out change in cycles of learning. An early priority is to engage with industry to define a broader framework for "superfit" services. The essential prerequisite is a language to describe the problem, and a means of defining what success might look like.

In the UK's case, this is also an opportunity to redefine the market interfaces, and so have significant impact on BT, Openreach and the regulated equivalence of inputs. What does it mean for a wholesale service to deliver the "same" performance to different users over varying geographies and bearer technologies? How can you be sure content providers are not being discriminated against?

This in turn opens up the opportunity for a broader industry upgrade, like those from dial-up to broadband, or analogue to digital, or fixed to mobile. The enlightened regulator can perform the necessary consultations to establish what is the right "superfit" vision, the best delivery approach, and how they can actually be delivered. Then the policy establishment can spearhead the change process.

The transformation timescale is 10+ years, but valuable results can be delivered relatively quickly. Each refactoring cycle can deliver tangible benefits tied to specific business processes. Examples might be fault isolation, retail service comparison, or service interoperability. This will create public awareness and buy-in to the "superfit" approach, resulting in a virtuous cycle of more transformation resources and ongoing service improvements.

THE 'SUPERECONOMIC' STEP

There is a necessary step for most that I term "supereconomic" on the way to "superfit" as in countries such as the UK users are still struggling to get basic broadband services. This means moving to more predictable relationships between the infrastructure and its economic utility – its value for money – shifting regulatory incentives from inputs to outcomes, where the emphasis is not so much on competition and certain technologies but on cooperation and decoupling services from network mechanisms.

Supereconomic is a transition zone from superfast to superfit. In mobile broadband, an example of superfast is 4G, supereconomic is 5G's "enhanced mobile broadband", and superfit is the desired (if overambitious) model for 5G's low-latency network "slices". Superfit then starts to deliver assured outcomes and is the regulatory end game.

MARTIN GEDDES *is a network performance scientist and a scholar of emerging telecoms technologies and business models. His company, Martin Geddes Consulting, is at martingeddes.com*