

July 2021
Geoff Huston

Another Volume in the Saga of the Decline and Fall of the Telco

The Swedish carrier group Telia has **recently announced** the sale of its international wholesale business to Polhelm Infra, an infrastructure investment manager jointly owned by a number of Swedish pension funds. Why would a telco operator sell off what was a core part of its operation to a pension fund?

The Internet was originally conceived as *a telephone network for computers*. (I should mention that this was not a concept that was unique to the Internet at the time. Most computer networks were designed using a similar transactional model.) By that I mean that, like the human users of a telephone network, computers could be both senders and receivers, conversations occurred between any two parties, and conversations had a distinct start and an equally clear end. The network was largely a passive fabric of connectivity, and a client of the network could only initiate conversations with other clients and could not converse directly with the network itself. This model was suited to the early mainframe model of the computer environment, where computers were shared by many users and used for many purposes at once.

However, the computer environment has changed radically over the past forty years, and with it the service requirements from a network have changed as well. Particularly so with relation to the Internet. The major shift has been the specialisation of the roles of client and server. Services are hosted within dedicated platforms that are constructed specifically for a server role. These days we use data centres, constructed as large-scale concentration points of storage and processing capability, as the platform for service delivery, while clients have disappeared into our pockets, or quietly embedded themselves into a myriad of *things* where they're all but invisible. The immediate corollary for the network's architecture was a transformation of the network into the same client/server distinction. Clients initiate network transactions with servers, and a subsequent bidirectional data stream ensues.

However, the client/server network had a somewhat unanticipated outcome. If servers never initiated a transaction with a client, then clients did not need to have a permanently available unique rendezvous address for other to use to initiate a network transaction. When a client was not talking with a server then it had no need of an IP endpoint address at all. There is an associated observation that a client could use many different client-side endpoint addresses to talk simultaneously to many different servers. The client's address only needed to be relatively unique to those used by other conversations, and a client had no requirement to present the same source address in all of its transactions with servers.

From the point of view of the client, an IP address is merely a conversation token, to be borrowed and used for a conversation with a server and then passed back into a shared pool. This form of sharing a pool of endpoint addresses between clients was implemented by Network Address Translators (NATs) and it is a foundational technology of today's Internet. Indeed, it is the reason why the exhaustion of the supply of unique IPv4 addresses has had no visible impact on the Internet, and the reason why IPv6 is still seen as an option, rather than an urgent necessity.

But even this client/server architecture has been a temporary model for the Internet. We have amassed abundant processing and storage capacity in the digital world. What is in short supply is time. The Earth has an unavoidable physical size, and there is an equally unavoidable delay in propagating a signal between two points. The further apart these two points, then the longer the propagation delay for a signal between those points.

However, we can use abundant processing and storage capacity to steer a path around this limitation. If we clone these service delivery points across the world and devise a way to direct each client to the closest instance of the service, then we can speed things up. Over the past decade or so that is exactly what we've been doing. These days Content Distribution Networks are the core of the Internet's operational model, and the aim is to position services in close physical proximity to every significant cluster of clients. Rather than having a network that brings diverse clients to a unique service delivery point, the network is now facilitating the selection of the best instance of a particular service delivery point from many potential alternatives.

These changes in the design of the Internet can be seen in many metrics. The speed and quality of service delivery have been improved by orders of magnitude in many for many users. The network is used more efficiently, and the costs associated with network transit have dropped dramatically.

Faster. Cheaper. Better. What's the problem?

Well, for most of us it's probably all good. There is however, one case where it's not quite so rosy, and that's the Internet transit service providers. Whether your assets are under the ocean or trenched across long overland distances it's somewhat challenging to change your asset base and reposition yourself to define a new engagement model with content-centric service network architectures. In some cases, we see these transit providers just holding on and riding through this change process and accept the inevitable consequences. Others are trying to pivot their business and get out while they can.

For that reason, I was drawn to this Telia announcement to divest itself of its wholesale international Internet transit business. One possible explanation here is that Telia is cutting and running before the residual value of its transit business and transmission assets plummet to the point where they are basically worthless. Of course, the reasons why a company divests itself of part of its activities is not always obvious, but assuming that companies act in terms of perceived self-interest and enhancing shareholder value it would appear that a number of Internet transit providers are seeing the transit market entering a phase of decline that could possibly be terminal and are looking for other investments that offer a better return.

It's not just the transit business, but the larger picture of cross-investment in foreign markets that could well be in decline. The period of the *techno-colonialisation* of the digital world, to put it somewhat cynically, where telcos whose home markets were saturated looked outward to less developed markets where they believed there was an opportunity for further growth and financial rewards, could well be drawing to a close. It's not just Telia selling off their international enterprises and investments, but they're joined by Norway's Telenor and Sweden's Tele2 are following a similar path of withdrawing from foreign markets. A summary of the changes over the past 5 years for these Nordic telcos is shown in Figure 1.

This shift is not isolated to the Nordic telcos of course. The wave of international investment in the telco sector started with the breakup of AT&T in the late 1980's and was followed by the progressive deregulation of many national telco sectors. The initial focus was the establishment of competitive phone services in national markets but was quickly followed by various Internet services, including wholesale and enterprise support services. Part of the rationale of this expansion into international markets was to service a domestic demand for international services in both the ISP and enterprise networks. However, the emergence of cloud-based services to the enterprise and retail sectors has undermined the ongoing value of this form of distributed service integration. A good recent example here is [Microsoft's Windows 365](#) service which takes the model of the now well-established cloud platform for applications and adds the **platform itself into the cloud**, making the desktop itself a cloud-based service.

TELENOR	TELIA	TELE2
Eurasian Exits	Eurasian Exits	Eurasian Exits
Russia (Veon, sold 2016-2019)	Turkey (Turkcell, sold 2017-2020)	Croatia (2020)
India (Uninor, sold 2018)	Russia (MegaFon, sold 2012-2017)	Kazakhstan (2019)
Hungary (sold 2018)	Afghanistan (Roshan, sold 2020)	Russia (2013)
Bulgaria (sold 2018)	Moldova (Moldcell, sold 2020)	Poland (2008)
Montenegro (sold 2018)	Kazakhstan (Kcell, sold 2018)	Hungary (2007)
Serbia (sold 2018)	Azerbaijan (Azercell, sold 2018)	Czech Republic (2006)
	Uzbekistan (Ucell, sold 2018)	
	Tajikistan (Tcell, sold 2017)	
	Nepal (Ncell, sold 2016)	
Acquisitions/Mergers	Acquisitions/Mergers	Acquisitions/Mergers
Finland (DNA, 2019)	Norway (TDC Norway, GET, 2018)	Sweden (Com Hem)
	Sweden (Telecom3, 2016)	
	Norway (Tele2, 2015)	
Remaining Markets	Remaining Markets	Remaining Markets
Norway	Sweden	Sweden
Sweden	Norway	Estonia
Finland	Finland	Lithuania
Denmark	Denmark	Latvia
Thailand (DTAC)	Estonia	Netherlands (T-Mobile, minority)
Bangladesh (GrameenPhone)	Lithuania	
Myanmar (exiting?)	Latvia (Tet, LMT)	
Malaysia (Digi, merging?)		

Figure 1 – Divestiture of International activities by Nordic Telcos - Source: Telegeography, <https://blog.telegeography.com/scandinavian-trio-show-strong-home-focus>

This change is one from content and service delivery as an overlay on a shared common network platform to content and services environments being tightly bound to a dedicated network infrastructure.

If we take some of the detail from [capacity reports](#) that Telegeography has made to industry forums, this trend is clearly visible. They've reported that in 2006 some 80% of international capacity was shared common IP transit capacity, while by 2018 this proportion was down to 40% of total international capacity. At the same time content providers' dedicated capacity inventory account for 54% of all international capacity by the end of 2018. This shift is most evident in the trans-Atlantic, trans-Pacific and intra-Asia markets. Further decline in the transit market and an associated shift to cloud and content distribution services is expected in the coming years.

It is also interesting to observe the decline in the role of North America as the major hub for interregional traffic. Some 97% of inter-regional capacity had a North American end point in 2004, and this has declined to some 80% of capacity using North America in 2018. In its place, Europe has grown as a hub, with inter-regional capacity that uses Europe servicing capacity to Asia, the Middle East and Africa. The overall picture of the western hemisphere being the hub for much of the international capacity inventory is unlikely to change dramatically in the near-term to mid-term future. While seafloor-based capacity in the Atlantic and the Pacific might involve longer paths than some Indian Ocean/Mediterranean routes, these extended paths represent a far lower political risk.

Perhaps the most significant aspect of this content and cloud capacity which impacts on the historical telco model is a *no sharing* regime. Cloud and content providers have expressed a strong preference for direct purchase of entire cable systems, or a dedicated fibre pair as a minimum level of capacity provisioning. This reduces unwelcome third-party dependencies and creates higher levels of assurance in meeting service and capacity evolution. While dedicated capacity measures were negligible in 2013, the intervening period has seen these cloud and content providers making subsequent investments in international cable systems that have resulted in wholly content-owned capacity representing 80% of the entire trans-Atlantic capacity, 50% of the trans-


Pacific capacity and 25% of the Latin American capacity by the end of 2018. Perhaps the critical issue here is that these cloud and content providers appear to eschew more intimate capacity sharing solutions, including wavelength-based segmentation. In trying to assess how far this dedicated capacity market segment will grow, the capacity predictions are somewhat clouded here. Part of the reason for this uncertainty is that this is no longer based on the old model of the Internet transit's *just in time* demand-pull model of content delivery to clients, but one that uses a CDN's cache pre-provisioning *just in case* push distribution of content. The larger the content volume being managed in this matter, and the greater the number of client access points (CDN data centres) the higher the demand for internal capacity to meet these content flooding demands.

But while these demands can be extrapolated to very high numbers, it should be stressed that the CDN industry is currently highly resourced, and while private capacity can be seen as an expensive luxury, it's a luxury that the content industry can currently afford. This edges out the more traditional capacity sharing broker, which is the telco wholesale carrier business model.

There was a time, and it was not so long ago, that the telcos dominated the communications sector. In many national environments they were the largest employer, the largest revenue earner, even the largest credit provider. Their position as a domestic monopoly created a secure revenue base that was immune from competitive pressures and the privileged position was all but unassailable. Deregulation opened up this sector to initial waves of like-for-like competition, but this has been quickly followed by a new wave of substitution competition where the rise of Internet-based services has shifted the nexus of value generation from carriage to content and service. The result was initially the commoditisation of the shared public carriage role, but the process of time and the successive improvements of the digital capacity of physical transmission systems has meant that dedicated transmission systems are now accessible to the large cloud and content providers and shared carriage systems have become a niche activity rather than the mainstream of the carriage sector.

With this trend we see the continual erosion of the residual value of shared common carriage systems. For the wholesale telco carriage provider this story of inexorable decline not exactly good news. No wonder some of them are bailing out!

I simply cannot let the title of this article pass without at least a passing mention to that epic work by the English historian Edward Gibbon. **The History of the Decline and Fall of the Roman Empire** is a six-volume epic, published between 1776 and 1789, still in print today. To quote from this work: “The story of its ruin is simple and obvious; and, instead of inquiring why the Roman empire was destroyed, we should rather be surprised that it had subsisted so long.”



Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.

Author

Geoff Huston B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region.

www.potaroo.net