Journal of Telecommunications and the Digital Economy

Volume 7 Issue 3
September 2019

Publisher: Telecommunications Association Inc.

ISSN 2203-1693

JTDE Volume 7, Number 3, September 2019

Table of Contents

Editorial

Public Policy Discussion

Articles

History of Australian Telecommunications

The Future of Australia's NBN Continued

Editorial

Leith H. Campbell
Guest Editor

Abstract: TelSoc has held its first forum on the future of Australia's National Broadband Network (NBN). Three papers from that forum are published in this issue. TelSoc is planning a second forum, discussing the user potential of the NBN, in October 2019. The historical reprint in this issue is also NBN-related about online learning. The technical papers in this issue concern architectural issues in the Internet of Things and cybersecurity. The *Journal* welcomes further contributions on telecommunications and the digital economy.

Continuing Discussion on the Future of Australia's NBN

On 31 July 2019, TelSoc (the Telecommunications Association, publisher of this *Journal*) held a forum on the future ownership of Australia's National Broadband Network (NBN). Three papers from that forum are published in this issue. *The NBN Futures Forum: Discussing the future ownership of Australia's National Broadband Network* is a summary of the forum itself, including the four speakers' main points and the subsequent discussion. *NBN Futures: The Option of Merging NBN Co with InfraCo, as a Benefit to the Digital Economy* is an expansion of the speech given by Professor Peter Gerrand at the forum. *Getting the NBN Infrastructure We Need* is an account of Dr Jim Holmes' argument for keeping NBN Co in public ownership.

These papers are a first contribution to the debate on the future of the NBN and NBN Co, the builder and operator of the network. As such, they have been made open access so that they are freely available to all interested readers. We commend them to you.

The debate on the future of the NBN will continue during the final stages of the network rollout and subsequently as the Australian government considers how best to leverage its investment in the network. This will be a significant public-policy issue in the early 2020s. Public discussion meanwhile can help to clarify the issues around the future of the NBN and, hopefully, build consensus among all stakeholders.

http://doi.org/10.18080/jtde.v7n3.224

As a next step, TelSoc will be holding a second forum on **Realising the User Potential of the NBN** in Melbourne at lunchtime on Tuesday, 22 October 2019. Details can be found on the *telsoc.org* website. As in the first forum, there will be four speakers making brief opening remarks, followed by general discussion. We expect to publish outcomes from this forum in the next issue of the *Journal*.

The historical reprint in this issue continues the NBN theme by reproducing a paper from the *Telecommunications Journal of Australia* in 2013 entitled *E-learning: Supplementary or disruptive?* It considers the potentially disruptive effect of the NBN on online learning. Although it is only six years old, much has changed since then.

In This Issue

In this issue, in addition to the papers related to the NBN, there are several technical contributions.

S-MANAGE Protocol for Provisioning of IoT Applications on Demand proposes a layered architecture and a specific protocol for configuring and managing applications on IoT devices while hiding the complexity and diversity of those devices.

Sanctus: An Architecture for Trusted Products argues against the Balkanization of software and systems distribution and describes an architecture in which a trusted device or subsystem can be embedded in systems or products to provide cybersecurity and trust.

The *Journal*, Looking Forward

The *Journal* welcomes papers on telecommunications and the digital economy, including theory, public policy and case studies.

Technological change is a constant in telecommunications. The *Journal* is especially interested in papers on how new technologies – especially 5G – will affect Australian telecommunications consumers.

Regulation and competition are also continuing themes. We encourage papers on the topics of *International Telecommunications Legislation and Regulations* and *International Mobile Cellular Regulation and Competition* that reflect on where the global telecommunications market is now, how it got to where it is, and what is going to happen next.

Papers are invited on local and international topics in telecommunications and the digital economy, broadly conceived. The Editorial Advisory Board also values input from our readership, so please let us know what themes you would like to see in future issues.

All papers related to telecommunications and the digital economy are welcome and will be considered for publication after a double-blind peer-review process.

Leith H. Campbell

The NBN Futures Forum

Discussing the future ownership of Australia's National Broadband Network

Leith H. Campbell
Honorary Fellow, School of Engineering, University of Melbourne

Murray Milner
Principal Consultant, Milner Consulting Ltd

Abstract: On 31 July 2019, TelSoc held an NBN Futures Forum in Melbourne to outline possible future ownership options for Australia's National Broadband Network (NBN). Following an introduction to the objectives of the forum, four speakers outlined various options for future ownership of the NBN and identified pros and cons for these options. Then the floor was opened for questions from both the local and the wider virtual audience. The resulting conversation provided a useful insight into the range of social, economic, technical and policy issues that need to be considered in order to reach a balanced and properly informed view on the most appropriate future ownership model for the NBN.

Keywords: NBN, public policy

Introduction

On 31 July 2019, TelSoc held an NBN Futures Forum in Melbourne, Australia, to encourage a discussion of the options for future ownership of Australia's National Broadband Network (NBN) after the completion of the rollout to all premises in the country, due by 2021. The NBN is currently being built and managed by a government-owned entity, NBN Co. The legislation setting up NBN Co envisages the eventual sale of the company after completion of the initial rollout and several other steps, with the approval of the Parliament.

The initial rollout is now reaching its final stages. The Minister of Communications has announced (in the video described below) that 9.93 million premises are now "ready for service", out of an approximate total of 11 million. After an area is declared "ready to connect", there is an 18-month "migration window", during which time customers must transition to a

http://doi.org/10.18080/jtde.v7n3.202

service delivered on the NBN by their chosen retail service provider. Before a sale of NBN Co, legislation requires the following steps to be completed (Gregory, 2018): the Minister of Communications must declare that the NBN has been built and is fully operational; the Productivity Commission must hold an inquiry on matters relating to the NBN and a Parliamentary Joint Committee must consider its findings; the Minister of Finance must make a disallowable declaration that conditions are suitable to sell the NBN and the Parliament does not disallow this declaration. With these steps to be taken, it is then possible to contemplate a sale of NBN Co in 2022 or beyond.

With this timing in mind, TelSoc (the Telecommunications Association Inc., publisher of this *Journal*) aims to promote informed discussion among stakeholders, with a view to building consensus at least among expert opinion and, if possible, in the political sphere on the future ownership of the NBN in the years beyond 2022. The Forum held on 31 July 2019 will be the first in a coordinated series of such events over the next 18 months. The outcomes of all events and supporting documents will all be recorded in this *Journal*.

Options for NBN Ownership

If NBN Co is not to be retained in government ownership, there are several alternative ownership options (Gregory, 2018) depending on to whom NBN Co is sold and whether it is sold as a single entity or in several parts. In parallel with the development of the NBN, Telstra has announced (Irving, 2018) the creation of Telstra InfraCo, a standalone business unit providing telecommunications fixed-network infrastructure. This raises the possibility of merging, by one means or another, NBN Co and InfraCo into a single entity.

The four options for future ownership of NBN Co canvassed at the Forum were:

A. Merging NBN Co and [Telstra] InfraCo into a single wholesale network provider;
B. Retaining NBN Co in Government ownership;
C. Selling NBN Co as a single entity – or, perhaps, the option of selling the urban parts of NBN Co while retaining the less competitive regional parts in government hands;
D. Splitting NBN Co along access technology lines and selling each part.

The fourth option, of splitting NBN Co along technology lines, was first mooted by the Vertigan report (Vertigan, 2014) as a means of promoting infrastructure competition and encouraging private investment. Promoting competition and potential competitive threats to the NBN have been canvassed in previous articles in this *Journal* (McLaren, 2018; Pugh, 2019).

The four options for the future ownership of the NBN represent the most discussed alternatives and the Forum was organized around them. However, this list may not be exhaustive and other options will be explored as they become defined.

The NBN Futures Forum

This first forum was designed to promote discussion of the ownership options, rather than to reach any definitive conclusions. The main presenters were therefore limited to short speeches of 7-10 minutes each, permitting them to outline their option, and leaving time for questions and discussion by attendees, both those physically present and those viewing the forum online.

A video of the complete forum is available on the TelSoc website (TelSoc, 2019).

The Forum was opened by Mr John Burke, who chaired the event. He described the process of discussion and debate to be supported by the forum and urged participants to consider the future of the NBN over the next 10-15 years and what it should be, rather than be diverted by current issues or shortcomings. He suggested that an ideal outcome from the foreshadowed series of forums would be bipartisan support for a future purpose for the NBN and the structural settings that would support that purpose. He summarized statements supportive of the process from the Minister for Communications and the Shadow Minister.

The Minister for Communications, Cyber Safety and the Arts, the Honourable Paul Fletcher MP, had provided a short video introduction in support of the Forum. (Because of technical difficulties, the video was only fully played at the end of the Forum.) In it, he noted that 9.93 million premises were now ready for service and he emphasised the strong focus on delivery that had achieved this result and would continue. He supported the view that it was an appropriate time for discussion on how best to leverage the $51 billion investment in the NBN. He would, he said, be interested in the outcomes of the forums.

Option A. Integrating NBN and InfraCo

This option was outlined by Professor Peter Gerrand. He began by emphasising that, without national policy goals being set for the Digital Society and Digital Economy that the NBN is obliged to meet, its fate will continue to be left to the market. History, he believed, shows that the market has a habit of disappointing many end users. In particular, he suggested that innovative digital businesses need very high speed, symmetric Internet access at affordable tariffs in locations beyond the NBN's FTTP footprint in order to compete in the global economy.

He argued that combining NBN Co and InfraCo into a single "NetCo", could do more to support the digital economy in the long term than NBN Co alone. He described the several possibilities of government or private ownership of NetCo and noted that an important caveat would be that Telstra's ownership of InfraCo would first need to be reduced to a non-controlling level before its merger with NBN Co, if InfraCo were to purchase NBN Co. A major benefit would then be the ability of the merged NetCo to do what a Telstra-controlled InfraCo would not be expected to do: to support the entry of new competitors to Telstra in the 5G market.

Professor Gerrand has provided a more detailed account of this option, published elsewhere in this issue (Gerrand, 2019).

Option B. NBN remaining in Government ownership

Dr Jim Holmes spoke to this option. He considered that a preference for a particular ownership option should proceed from clearly articulated policy and strategy. He suggested that national inclusion and a universal service obligation were aspects that only a government could do. He described retention in government ownership for the time being as the "least-worst" option. He supported this argument through a comparison table (well received by attendees) showing the key attributes of each of the four options.

Dr Holmes has provided a more detailed description of his argument, published elsewhere in this issue (Holmes, 2019).

Option C. NBN sold as a (single?) entity

Mr Graeme Samuel AC, in speaking to this option, suggested that now was not the time to privatise NBN Co. There was, he said, negative speculation about the value of the NBN, for example relating to wholesale service costing and the potential challenge from 5G. In such an environment, there would be no premium in a sale price. It was necessary to reach a situation of "business as usual", a point with some market stability, including known competition. He felt that this point would be reached in 3-5 years.

The issue, he believed, was not *whether* to privatise but *how* to privatise NBN Co. Effective competition would be key. He suggested that the privatisation of Telstra had breached competition policy by not including either strong regulation or structural separation of Telstra. This historical example highlighted the market problem of vertical integration, something that must be avoided in the future ownership model for NBN Co. He noted in passing that adding InfraCo into the mix would be an inappropriate aggregation of resources and so should be avoided.

In acting on privatisation, the Government, he suggested, should not be seeking to maximize its financial return but, rather, it should act in the best interests of long-term public policy. The long-term interests of end users should be paramount. Any continuation of a universal service obligation should be made transparent. Good policy suggested that NBN Co should not be privatised as a single entity. For example, it may be necessary for some time to keep the regional NBN in government ownership.

Mr Samuel described himself as a strong advocate of the Vertigan principles (Vertigan, 2014). These, he said, represented sound policy, supporting competition and the long-term interests of end users. Competition, he noted, would bring consumer benefits, as well as promoting innovation and reducing the need for regulation.

In summary, he supported competition: infrastructure competition, to the maximum extent possible. This could best be brought about by creating competitive entities through disaggregation of the NBN business. It would also have the effect of reducing the burden of regulation.

Option D. NBN disaggregated by technology and sold

Mr Michael Cosgrave, Executive General Manager, Infrastructure Regulation Division, Australian Competition & Consumer Commission, spoke to this option. He noted that the recommendation to split NBN Co into technology-based businesses had been made by the Vertigan Committee (Vertigan, 2014) and had been supported by the ACCC. This had been the Vertigan inquiry's answer to what would be the most appropriate structure for delivery of future broadband.

He noted that Stephen Rue, NBN Co's CEO, had an opinion piece (Rue, 2019) in that day's *Australian Financial Review*, outlining the objectives for the NBN set by government. Mr Cosgrave summarised these objectives as three-fold: build the network; earn a financial return from the network; provide broadband availability across Australia. These would remain the objectives of the NBN, however it was delivered.

On the timing of privatisation, he suggested that disaggregation along technology lines had been deferred, not abandoned, by the Government. Privatisation is unlikely to occur in the immediate future, allowing time for debate and design of an appropriate competitive framework.

It had been Vertigan's conclusion (Vertigan, 2014) that infrastructure competition should be the basis for future wholesale broadband provision. The ACCC had looked at this issue since 2003 and most recently in 2018: it remained interested in models for infrastructure-based competition. The Vertigan proposal was not necessarily the only means of providing effective competition. For example, fibre technologies could be split from the rest.

The ACCC had also considered a geographic split of the NBN. For the less competitive areas – with access based on satellite or fixed wireless – a privatisation would have the benefit of making any subsidies (for example, for a universal service obligation) transparent.

An aggregation of NBN Co with Telstra InfraCo to create a NetCo would raise questions of competition and about likely future upgrades of current access technologies.

On the question of competition from 5G, Mr Cosgrave considered the jury was still out. He noted that Andy Penn, Telstra's CEO, at a National Press Club address that very day (Duke, 2019) would say that competition from 5G would be only at the margins.

Questions and discussion

Questions and discussion from those attending in person and online followed the introductory speeches.

Was the ACCC wrong to recommend 121 Points of Interconnection (PoIs) for the NBN?

> Mr Cosgrave emphasised that this was a government decision, albeit one based on advice from the ACCC. He described the ACCC as having made a balanced compromise between three competing options: 8-12 centralized PoIs; a hybrid model for transmission competition; and a Telstra-supported proposal for 600 PoIs. The compromise reached was 121 PoIs following detailed analysis of these options.

Given that fibre to the premises has much lower operating costs and higher reliability, leading to greater operating profits in areas served by FTTP, is a uniform national price only possible if NBN Co remains a single entity?

> Mr Samuel claimed to be unpersuaded by the need for a uniform national wholesale price. He suggested that improved prices would come from competition between wholesale providers and he would support any means of maximizing competition.

Is the continuing ownership by Telstra of pits and ducts a barrier to disaggregation of NBN Co or future wholesale competition?

> Mr Samuel considered that the mixed ownership of infrastructure could be a financial advantage, not the disadvantage generally assumed. He did not see that the Telstra ownership of pits and ducts would support the creation of NetCo, if the need for competition would later lead to disaggregation of the business.
>
> Professor Gerrand disagreed. He believed that the ownership by Telstra InfraCo of the pits and ducts supported the merger with NBN Co to create NetCo, which could be justified in terms of meeting national goals for digital equality and providing some competition to future "NBNs", such as 5G. He did not support infrastructure competition created artificially; instead, he believed that forms of competition should naturally occur, as, for example, from other technologies. He noted the example of a very fast train line, which may have no direct competition from other train lines but was considered beneficial because it supported national infrastructure goals. Mr Samuel thought the analogy with a train line was not valid, because it would always face multimodal competition, such as from buses and planes.

Are there national security implications in a sale or disaggregation of NBN Co, considering that now retail service providers are required to coordinate security issues with Government?

> Mr Samuel argued that security was sometimes used as a reason for not taking economically rational decisions. The industry collectively needed to focus on cyber risk, just as the banks did.

The Vertigan recommendation of horizontal disaggregation of NBN Co had been made in 2014, but since that time there have been new developments for the delivery of broadband, such as the introduction of 5G and new low Earth orbit satellites. Do these developments change the view of disaggregation recommended by Vertigan?

> Mr Cosgrave remarked that there would always be technological changes that could change the competitive positions. He noted that, in the NBN rollout, there had been some change in the mix of access technologies, with HFC being used for about 400,000 fewer premises than had been earlier envisaged. He questioned, however, how fundamentally the new developments affected the competitive landscape. He noted that 5G was already competing at the margins, but he remained cautious about competition from wireless technologies because of the enormous amounts of data currently being carried on fixed networks. He recognized that this was current thinking and may change over time.
>
> Dr Holmes thought the question raised a fundamental issue. He was concerned that a structural separation based on current technology and technology forecasts could lock the new entities so formed into the technologies associated with their initial assets more than would be desirable. Technology forecasts would always be changing. If a separation by technology was to be contemplated, it could only be planned at the end of the initial rollout.
>
> Mr Samuel claimed that the Vertigan recommendations were not dependent on fixed proportions of each technology. Instead, separation by technology was just a means for starting competition which would evolve over time as new technology innovations entered the market.

Dr Murray Milner from New Zealand, where a broadband policy has been seen to be successful, was invited to make some closing remarks. He suggested that there had been three critical factors in New Zealand's success: structural separation of Telecom NZ to create Chorus as a wholesale provider; a degree of geographic wholesale competition; and a separation of urban and rural rollouts. He noted that there were 10 Gbps services in operation today, much faster than any service currently delivered on the NBN. He remarked that it was very difficult to compete with fibre-to-the-premises solutions, unless some other service attribute is required, such as mobility.

Conclusion

This was the first of a planned series of forums on the topic of future ownership of the NBN. As such it was never intended to lead to definitive conclusions on the preferred ownership model. Instead, it was intended to start a conversation that would lead to better understanding of the options available and their pros and cons. This the forum did very well, with four different points of view presented followed by a robust discussion through audience participation. It was clear from the presentations and discussion that all the options considered would require more elaboration before it would be possible to make definitive judgments between them.

Two areas for further exploration that arose from the first forum were:

- What role does the NBN play in supporting a digital economy and digital society, including the long-term interests of end users and the role of competition?
- How will the changing technological landscape, including 5G (and perhaps 6G) and developments in fixed access, affect the value and competitive position of the NBN?

These and other topics will be the subject of future forums over the period 2019-2020.

References

Duke, J. (2019). 'Incredibly damaging for the NBN': Telstra boss warns 5G disruptors will take customers. *Sydney Morning Herald*, 31 July. Retrieved from https://www.smh.com.au/business/companies/incredibly-damaging-for-the-nbn-telstra-boss-warns-5g-disruptors-will-take-customers-20190731-p52cda.html

Gerrand, P. (2019). NBN Futures: The option of merging NBN Co with InfraCo, as a Benefit to the Digital Economy, *Journal of Telecommunications and the Digital Economy*, *7*(3), September. DOI: 10.18080/jtde.v7n3.198

Gregory, M. A. (2018). Australian Wholesale Telecommunications Reforms, *Journal of Telecommunications and the Digital Economy*, *6*(2), June, 1-34. DOI: 10.18080/jtde.v6n2.155

Holmes, J. R. (2019). Getting the NBN Infrastructure We Need, *Journal of Telecommunications and the Digital Economy*, *7*(3), September. DOI: 10.18080/jtde.v7n3.199

Irving, W. (2018). Establishing a standalone infrastructure business unit. *Telstra Exchange*, 20 June. Retrieved from https://exchange.telstra.com.au/establishing-standalone-infrastructure-business/

McLaren, G. (2018). What Now for Australia's NBN? *Journal of Telecommunications and the Digital Economy*, *6*(4), 31-62. DOI: 10.18080/jtde.v6n4.162

Pugh, N. (2019). The Wireless Threat to Fixed Broadband Services. *Journal of Telecommunications and the Digital Economy*, *7*(1), 7-19. DOI: 10.18080/jtde.v7n1.178

Rue, S. (2019). Don't lose sight of NBN Co's goals, *Australian Financial Review*, 31 July. Retrieved from https://www.afr.com/companies/telecommunications/don-t-lose-sight-of-nbn-co-s-goals-20190730-p52c6i

TelSoc. (2019). NBN Future Forum: Encouraging Debate on NBN Ownership Models. Retrieved from https://telsoc.org/event/national/2019-07-31/nbn_future_forum

Vertigan, M. (2014). *Independent cost-benefit analysis of broadband and review of regulation*. Department of Communications, 14 August. Retrieved from https://www.communications.gov.au/sites/g/files/net301/f/NBN-Market-and-Regulatory-Report.pdf

NBN Futures: The Option of Merging NBN Co with InfraCo, as a Benefit to the Digital Economy

Peter Gerrand

Honorary Professorial Fellow, The University of Melbourne

Abstract: No national goal has been set for how the National Broadband Network should provide competitive advantage for Australian small or medium enterprises (SMEs) participating in the global digital economy. This paper proposes robust national goals for how the NBN should serve both our digital society and our digital economy. From this perspective it considers the merits of merging NBN Co with InfraCo, and the pros and cons of public versus private ownership of the merged entity, "NetCo".

Keywords: Telecommunications policy, structural separation, National Broadband Network, Telstra, Australian telecommunications.

Introduction

The author was invited to discuss the merger of NBN Co with Telstra's InfraCo as one of four options proposed at TelSoc's NBN Futures Forum held in Melbourne on 31 July 2019.

A reason for this mission was evidently my authorship of the paper 'Revisiting the Structural Separation of Telstra' (Gerrand, 2004), which had some impact on policy discussions in the industry. In this paper I had advocated, ahead of the full privatization of Telstra in 2005, the structural separation of Telstra and the retention in government ownership of its fixed network wholesale business. The motivation was to create a self-funding entity which could roll out a high speed national broadband network without the need for additional government investment – as well as providing a level playing field for all retail service providers of fixed access broadband.

Graeme Samuel in his presentation at this same NBN Futures Forum suggested that "the privatisation of Telstra had breached competition policy by not including either strong regulation or structural separation of Telstra. This historical example highlighted the market problem of vertical integration, something that must be avoided in the future ownership model for NBN Co." (Campbell & Milner, 2019). I could not agree more.

http://doi.org/10.18080/jtde.v7n3.198

However, in 2006 the Australian Government chose to maximise the short-term financial returns from the full sale of an unseparated Telstra over alternative policy options[i], and gained a net $15.2b from the "T3" sale in that year (ANAO, 2008). The previous sales of 16% ('T1" in 1997) and 35% ("T2" in 1999) of the Government's shares in Telstra in 1998 achieved net returns of $14.0b (ANAO, 1998) and $15.9b (ANAO, 2000); thus, the total privatisation earned a net return $45.1b after deducting the costs of sales. Ironically, this figure exceeds the estimated cost to the government of funding the original NBN project, and comes close to the $51b expected to be spent on the project by its completion in 2020. But that is now water under the bridge.

The opportunity in 2006 to use Telstra's structurally separated, profitable wholesale fixed network business (now known as InfraCo) as the engine for rolling out a high-speed NBN was lost. Telstra's continuing dominance as a vertically integrated carrier not only prevented the entry of any significant infrastructure competition in the fixed broadband market, but kept most residential premises restricted to the entry-level ADSL technology, at a time when much higher speed access technologies were being introduced overseas.

When finally the Rudd government bit the bullet in 2009 and decided to fund a government-owned NBN, it planned to expend $41b on its national rollout, with future-proof FTTP to be connected to 93% of all premises. Due in part to major changes in the NBN's design in 2014 by Communications Minister Malcolm Turnbull under the Abbott government, the total government investment in the NBN has increased to $51b, with a legacy of the original FTTP now being available (but not necessarily connected) to only 21% of all premises (NBN Co, 2018).

Many would have considered that the InfraCo horse had bolted. However, in 2018 Telstra announced (Penn, 2018) that it planned to structurally separate most of its wholesale fixed network business as a stand-alone business unit, InfraCo, with the potential for divesting it. This raised afresh the policy option of merging InfraCo with the new broadband access network business, NBN Co, to create an end-to-end wholesale network business ("NetCo"), with possible added value to national infrastructure building in the future.

This is the background to the talk I gave on 31 July 2019. The following sections flesh out the presentation I made at that NBN Futures Forum. (A glossary of abbreviations is provided at the end of the article.)

Why Do We Need National Broadband Networks?

Before we consider any of the future ownership options for the NBN, we should ask the fundamental question: 'Why do we need the NBN?' Or, indeed, 'Why do we need NBNs?' in the plural: because in the ten-year time frame of this policy discussion, several other networks will emerge, in my view as early as 2020, that will satisfy most definitions of a 'national broadband network'.

My answer to that question is simple: we need NBNs to support national policy goals for both our Digital Economy and our Digital Society. I illustrate that point with Figure 1.

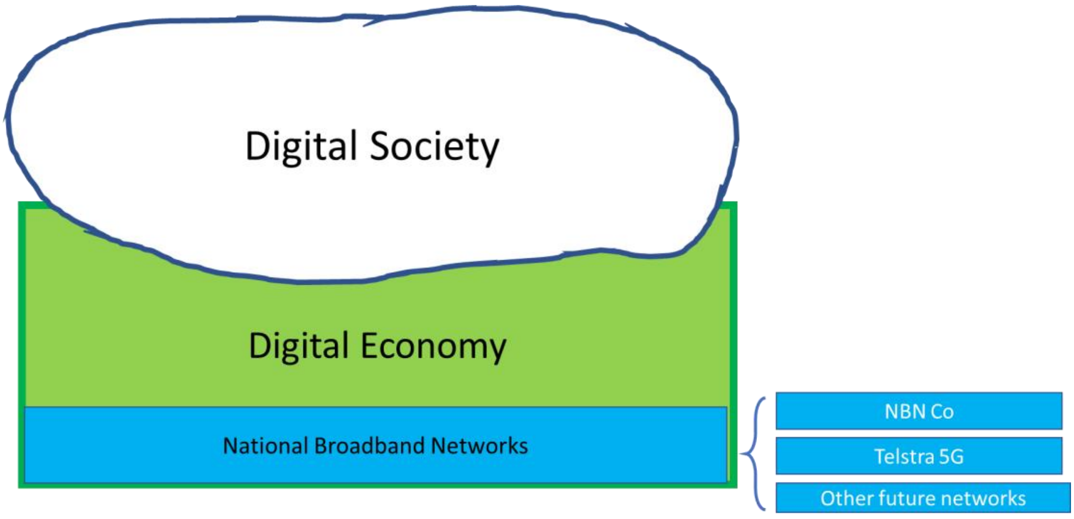

NBNs need to support the policy goals for both the Digital Economy and Digital Society

Figure 1. Why do we need the NBN?

A reasonable policy goal for our Digital Society is:

> All residents should have affordable broadband access to essential online services, irrespective of the location of their residence (to the maximum extent possible).

This goal is more focussed explicitly on consumer needs than the Universal Service Guarantee announced by the federal government on 5 December 2018 (Fifield, 2018), intended to replace the former Universal Service Obligation legislation after the NBN has been fully rolled out (Reichert, 2017). The Universal Service Guarantee will commit the government to providing affordable broadband (peak speeds unspecified) and telephone services in rural and remote areas. However, if essential e-health services for remote residences require, for example, 25/25 bandwidth in order to provide a high-quality video signal upstream for diagnostic purposes, the current rural NBN and hence the USG itself will not be fit for purpose.

A reasonable policy goal for our Digital Economy is:

> All Australian businesses should have access to broadband infrastructure at internationally competitive bandwidths and prices, irrespective of their location (to the maximum extent feasible).

Why should they be offered less?

The original NBN plan satisfied this objective for the 93% of premises intended to be served with FTTP, an access technology that can be upgraded cost-effectively to Gbps speeds as required. The current NBN, as we shall see in the next section, will probably only satisfy this policy objective in about one in four premises across Australia, in service areas that are largely located in the capital cities.

The government's most recent Statement of Expectations for the NBN, dated 24 August 2016, "expects the network will provide peak wholesale download data rates (and proportionate upload rates) of at least 25 megabits per second to all premises, and at least 50 megabits per second to 90 per cent of fixed line premises as soon as possible" (Department of Communications and the Arts, 2016). These speeds are fine for watching Netflix or SBS On Demand, but are ludicrously inadequate for innovative SMEs, and totally uncompetitive when compared with the 900/450 Mbps offerings now available, cheaply by Australian standards, to 70% of premises in New Zealand, as will be detailed in the next section.

Figure 1 also makes the point that we need to consider **the co-existence of other NBNs** in the near- to mid-term future, which could to some extent assist in satisfying national policy goals for both the digital society and (especially) the digital economy.

The first of these will be Telstra's 5G network, implemented through the T22 strategy, which Telstra first announced in June 2018. T22 implies an investment by Telstra of $3b to, *inter alia*, "lead the market and win in 5G". Telstra has already achieved its milestone of being 'network ready' for 5G in the first half of 2019 and is already marketing it. It plans to have "full rollout to capital cities, regional centres and other high demand areas by FY20", i.e. by June 2020 (Telstra, 2019).

Needless to say, Telstra aims to have "the largest, fastest, safest, smartest and most reliable next generation network" (Telstra, 2019). The sobriquet T22 suggests that their new 5G-based network and associated software applications will be fully implemented across the nation by June 2022.

Telstra's 'first mover advantage' in dominating the 5G market has been assisted considerably by the regulator (ACCC's) decision in May 2019 to block the proposed merger of TPG and Vodafone Australia, which left many industry observers surprised. Telstra has also been aided by the Australian government's decision in August 2018 to ban the use of the Huawei 5G network technology, which the TPG-Vodafone merger and other potential 5G market entrants were intending to use to provide significant cost advantage.

Telstra's CEO, Andy Penn, has recently made a point of saying that, if the NBN's wholesale pricing is not significantly reduced, "the $51 billion project has left itself at risk of losing customers to competitors using high-speed mobile technology" (Duke, 2019b).

In the ten-year timeframe of this forum's policy discussion, we can expect other network technologies with superior cost and/or performance to arise and provide further infrastructure competition to the current NBN.

However, it will be very difficult for new entrants to the 5G market to compete with the firmly entrenched and dominant Telstra T22 network. Many of the following next generation network technologies may also be first introduced by Telstra, once it has entrenched a dominant and profitable lead with 5G.

What We've Got: the Current NBN

At the time of this forum, on 31 July 2019, the most recent public information we have on the status of the NBN is regrettably more than a year old, as it appears in the company's FY 2018 Annual Report (NBN Co, 2018), issued in October last year.

Table 1 shows the progress by NBN Co in achieving "ready for service" status in designated service areas using seven alternative access technologies. Unfortunately for external analysts, the FTTN technology, widely found to be unfit for purpose, is mixed with the more versatile FTTB and FTTC access technologies.

Table 1. NBN statistics (Source: NBN Co (2018))

Technology	Premises "ready for service" (million)	Percentage of all premises
FTTP	1.7	21%
FTTN/FTTB/FTTC	4.0	49%
HFC	1.4	17%
Fixed Radio	0.6	7.4%
Satellite	0.4	4.9%
Total	**8.1**	

One can conclude from the current bandwidth limitations of NBN's FTTN, HFC, Fixed Radio and Satellite services, and the pricing of NBN Co's most popular product offerings, that the NBN is currently largely a network for a nation of Netflix watchers, using 25/5 or 25/10 peak speed products.

One cannot blame NBN Co's management for this: they have simply been responding to the market, in the absence of national policy goals[ii] that aim to provide competitive advantage to SMEs participating in the global Digital Economy. The broadband market is dominated by the millions of residential users largely needing downstream speeds for video entertainment. This swamps the needs of the thousands of innovative SMEs (in scattered locations across the country) needing symmetric speeds of at least 100 Mbps (preferably much higher) in order to send their high-density data files to customers, suppliers and collaborators.

By comparison, in New Zealand the majority of RSPs are offering a 900/450 Mbps Ultra Fibre service to SMEs, available to 70% of residential premises (i.e. across New Zealand's 'Ultra Fibre Broadband' FTTP footprint, serviced largely by Chorus). See for example Vocus's Fibre 900/450 product priced at NZ$137.42 per month, including GST, for unlimited data (Vocus, 2019).

If we set a national goal of providing competitive advantage to our SMEs operating in the global digital economy, one concludes that only the FTTP, FTTC and FTTB access technologies will provide the necessary symmetric and ultra-highspeed services they need. By this criterion, more than half of the current NBN is not fit for purpose.

To be fair to NBN Co, they are working on upgrade paths for their HFC and FTTN technologies, for when their RSP customers start demanding higher speeds. (The upgrade path for FTTN is basically FTTC or FTTB.) But upgrading the NBN to provide the broadband services at a par with those being offered in New Zealand, as just one competitive example, will require significant additional investment by the current or future owner of the NBN (Gregory, 2019).

What is the sale (or purchase) value of the NBN likely to be, on completion of its rollout? NBN Co's estimate for 2021, made in July 2018, is a net value of $10.4b (NBN Co, 2018). On the other hand, PWC made an estimate of $27.0b for the NBN's net value in 2024, in a report to Infrastructure Australia in February 2016 (Ramli, 2016). Just prior to publication of this paper, NBN Co released its 2020-23 Corporate Plan, forecasting an EBITDA of $3.2m in 2023, suggesting a net value of $19.2b (NBN Co, 2019), using the same x6 multiplier.

However, these estimates ignore the expected $49b debt to the Commonwealth, and the allowed $2b debt to the private sector, which would need to be brought to account following any sale of NBN Co (Department of Finance, n.d.). For example, the sale of NBN Co for $27b in 2024 could lead to a loss of $24b being added to the federal budget's bottom line in the financial year of the sale. Given the current government's aversion to incurring any budget deficit, one can understand their unwillingness to privatise the NBN within the term of the current parliament (Pearce, 2019a).

The three estimates above also ignore the cost of upgrading the NBN in order to remain competitive against future 5G offerings – and future low-earth-orbit multi-satellite offerings (Ritchie, 2019) – over the next five years. Mark Gregory has recently considered three alternative strategies for the upgrade. He concluded that the cheapest strategy, in which the retail customer would be expected to pay for the pit-to-premises lead-in, would cost NBN Co between $10b and $12.2b for a total upgrade of its existing FTTN, HFC and FTTC technologies to at least FPP (Fibre Premises Passed) (Gregory, 2019). Some fraction of this cost needs to be factored into any realistic evaluation of NBN Co's future liabilities.

Telstra's InfraCo

InfraCo is a business division within Telstra that consists of the following fixed network assets (Telstra, 2019; Penn, 2018):

- Telstra's national transit network plus international undersea cables *but not the fibre dedicated to supporting Telstra's mobile network.*
- Telstra's fixed access networks: all its ducts, pits and pipes; and the residual copper and HFC networks yet to be transferred to NBN Co.
- Telstra's more than 5,000 exchange buildings and data centres;
- *but not yet its own Operations Support System or its own billing system – these are located and managed in other divisions of Telstra.*

InfraCo has more than 200 wholesale customers, including NBN Co (which contributed about $1b of InfraCo's revenues in FY 2018), Telstra's external wholesale customers (largely 'corporate and government'), and its internal retail businesses (consumer, small business and enterprise) (Telstra, 2019; Penn, 2018).

At InfraCo's launch as a stand-alone business division in July 2018, Telstra's CEO, Andy Penn, valued its book value (i.e. replacement value) at $11b (Chirgwin, 2018). However, its EBITDA of $1.225b in FY19 (Telstra, 2019, p. 41), if annualised as $2.45b, suggests a conservative net value of $15b, possibly a lot more, provided it were not saddled with any of Telstra's corporate debt prior to sale.

In short, InfraCo itself is a large, profitable wholesale fixed network business – but one which Telstra regards as being surplus to its future requirements, and will consider "divesting", i.e. selling off, when opportune (Chirgwin, 2018). InfraCo's network is nationally strategic: it comprises the major part of Australia's transit broadband telecommunications network, connecting all of Australia's major cities by very high-bandwidth optical fibre links, and supporting most of the country's digital economy.

Why Merge NBN Co with InfraCo?

The first thing to note is that the network assets owned by these two businesses are quite complementary: InfraCo owns Australia's largest optical fibre *transit network*, and NBN Co owns Australia's largest set of broadband fixed *access networks*. InfraCo's ownership of the pits and pipes that NBN Co uses is a further area of exact complementarity. It is significant that NBN Co already depends upon InfraCo's transit network, and is a major customer.

http://doi.org/10.18080/jtde.v7n3.198

Secondly, the merged entity – let us call it "NetCo" – would have greater managerial alignment in meeting Australia's national infrastructure goals, most of which require advanced telecommunications technology. A concerted national effort to grow the Digital Economy would particularly benefit from the merger.

There are some significant synergies that would add $ value to the merged entity, NetCo:

- Reduced costs with a single Operations Support System and single wholesale billing system;
- No need for profit between the backhaul network (InfraCo) and the access network (NBN Co), thus ensuring that either output prices to customers can be reduced or profits to the owner increased – or a balance between these two aims, to make the merger acceptable to the key stakeholders.
- NetCo would have greater resources – financial, technical and managerial – to plan and design the technology upgrades essential for the NBN access networks to remain fit for purpose.

But, in addition, there are two important strategic benefits that the merger would deliver, from a national competition viewpoint.

Firstly, NetCo can provide a level playing field of wholesale backhaul services to the whole industry, including Telstra – which would not be achieved while InfraCo remains owned by Telstra.

Secondly, NetCo can do what the NBN alone cannot do, and what a Telstra-owned InfraCo would not do. It can offer cost-effective transit network services to new entrants into the 5G market, which is likely to become heavily dominated by Telstra, as discussed above. How NetCo can achieve this will be explained with the help of Figure 2 below.

5G networks will provide picocells with very high data transfers taking place in small radio footprints, often requiring the 5G radio stations to be attached to buildings or other existing structures, e.g. along streetscapes. They are expected to be well suited to managing communications for moving vehicles, such as driverless cars and other vehicles. NetCo can use its NBN-derived FTTP and FTTN service areas to provide fibre connections to 5G radio stations, as well as using its InfraCo-derived transit network as the most cost-effective backhaul in the country, to provide a 'single shop' wholesale fixed network service to new entrants into 5G and later networks. This should be much more economical for new 5G players than investing in their own transit networks.

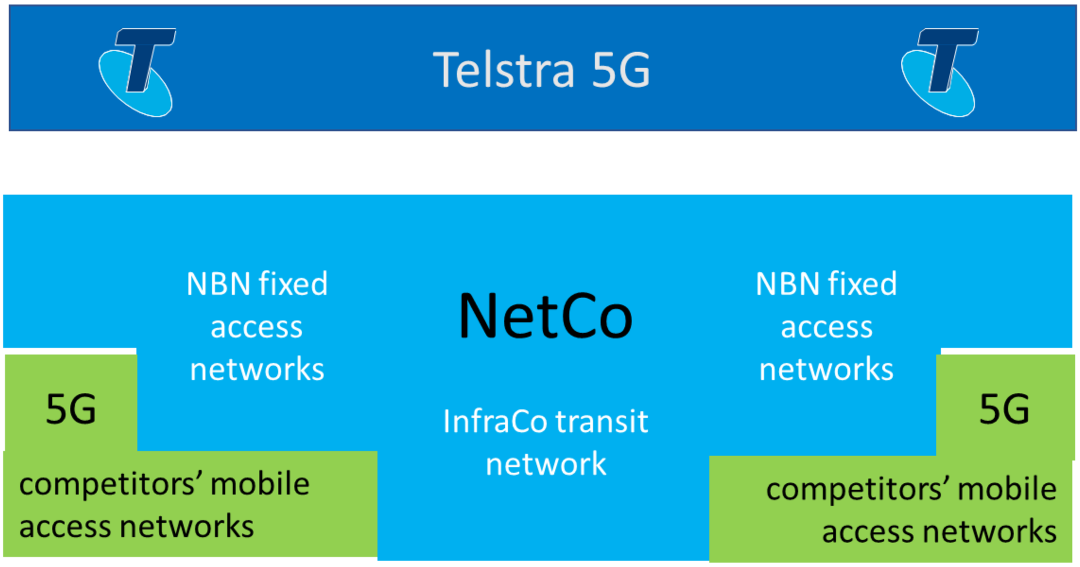
Telstra 5G
NetCo
NBN fixed access networks
NBN fixed access networks
InfraCo transit network
5G
5G
competitors' mobile access networks
competitors' mobile access networks

Figure 2. How NetCo can support new entrants to the 5G market

Telstra has chosen to publicly underplay the potential of 5G networks to substitute for NBN connections except in the context of warning of the likely consequences of NBN Co sustaining its current high wholesale prices (e.g. Kruger (2019)). However, a recent study of fixed-mobile substitution by Pugh (2019) with 1,446 participating households reveals that in 2017 "30% of existing fixed broadband households would consider switching to a wireless broadband service", in the context of existing retail prices for the NBN and for 3G/4G mobile services. The key reasons given by these respondents were "concerns over the NBN", "wireless faster than their fixed connection", "wireless is cheaper" and "portability" (Pugh, 2019). I take the view that 5G will be a considerable substitutional threat to the NBN in its poorer performing service areas – as well as providing opportunities to NBN Co as a 5G infrastructure provider.

Arguments Against Such a Merger

Infrastructure competition. The first counter argument arises from the traditional regulatory culture in this country, in which infrastructure competition is raised as being the highest policy goal[iii] – far higher than the goals of greatest interest to consumers or citizens, such as reduced retail prices, or more reliable performance, neither of which is necessarily guaranteed by infrastructure competition, but are simply expected to flow from it. Because NetCo, whether publicly or privately owned, will be a monopoly, this will be considered to be an extremely undesirable outcome. The traditionalists would rather InfraCo and NBN Co continued as separate entities, encouraged to invest competitively in each other's monopoly area.

But infrastructure competition does not necessarily lead to better outcomes for end users. Some of us remember the frantic competition between Optus and Telstra in the 1990s, to gain an early advantage in what was then called the cable TV market. The result was both carriers laying parallel HFC access networks in essentially the same streets to serve essentially the same set of two million premises, located in the most affluent suburbs of Australia's four largest cities. Any rational person would see that this duplication was a wasteful allocation of nationally valuable infrastructure, because another eight million homes missed out.

The obvious counterexamples to the belief in infrastructure competition as the essential solution for reducing prices occur in the electricity and energy distribution markets: it simply hasn't worked. If your prime aim is to control prices, then surely the best mechanism is price control. Ditto for control of service performance, in this case via the monitoring and regulation of Service Level Agreements.

Perceived lack of innovation. An argument is sometimes raised, usually by textbook economists, that monopolies lack the ability to innovate. This was clearly not true for the US private monopoly AT&T, whose Bell Labs produced a cornucopia of both fundamental discoveries and practical inventions from its founding in the 1920s onwards. Nor was it true in Australia for either the PMG Department or its successor, Telecom Australia, both public monopolies, which were sources of ongoing sequences of innovation in new network services and technologies, over a span of seventy years.

The obvious means for ensuring that NetCo, as either a public or private monopoly, would continue to innovate to remain fit for purpose, would be to ensure firstly that its management encouraged it, and secondly that its budget included the funding of a small R&D division to assess ongoing customer needs and evaluate new solutions for them. NBN Co has recently created an "insightLab" to find ways to improve its services (Crozier, 2018).

NetCo too large to be manageable? I was surprised that this objection was raised at the forum. The obvious response is to note that, if NetCo were created today, it would be significantly smaller in both size of workforce and range of business activities than Telstra has been during its most commercially successful years.

Anticompetitive structure. Communications Minister Paul Fetcher has emphatically (and understandably) ruled out the NBN being owned by any vertically integrated telco, such as Telstra: "That's baked into the legislation" (Duke, 2019a). But that does not rule out InfraCo purchasing NBN Co if Telstra's ownership of InfraCo were previously reduced, e.g. via an IPO, to a level at which Telstra had no effective control, e.g. below 10%. This is a scenario which Telstra's CEO, Andy Penn, seems to be positively contemplating in the next few years (Pearce, 2019b). The sale of a network that is no longer essential for Telstra's T22 strategy would provide a valuable financial contribution to its bottom line.

Ownership of NetCo: Public or Private?

If the merger of InfraCo with NBN Co is seen as being in the national interest, the question remains as to whether the merged entity should best be in public ownership, i.e. by the Australian government, or in private ownership. The first outcome would be achieved by NBN Co buying InfraCo; the second either through InfraCo buying NBN Co, or through an independent investor buying both and merging them.

As the pros and cons of the third possibility (the action of a third party) are virtually the same as for the second (InfraCo buying NBN Co), only the first two scenarios will now be examined.

Option A1: NBN Co, under government ownership, to buy InfraCo

Advantages. Firstly, the government's AAA investment rating gives it cheaper access to the ongoing finance necessary to keep the combined entity (NetCo) fit for purpose, with appropriate technology upgrades.

Secondly, NetCo under government ownership should be able to put the needs of users above the need to raise maximum dividends for its owners, enabling it to offer a cheaper range of wholesale products than if it were privatised.

Thirdly, the government as owner would have greater governance control over NetCo to comply with its national policies, either via legislation, regulation or simply Ministerial direction. The memory of the fully privatised Telstra's CEO, Sol Trujillo, locking horns with the federal governments in 2007-2009 over national broadband policy remains a good history lesson.

Disadvantages of government ownership. Firstly, there is the political risk of the investment returning a net loss to the federal budget if the government continues to allow infrastructure competitors to 'cherry pick' the most profitable parts of the market.

Secondly, we have seen in recent years considerable political aversion to any form of electoral risk; and we have observed a large amount of public criticism over the NBN's underperformance being directed towards the government rather than to NBN Co.

Thirdly, there is the tendency for Ministers to interfere in engineering and financial decisions for perceived electoral benefits, such as the mistakes in choosing FTTN over FTTP as the preferred access technology in 2009, and indeed the extra costs to NBN Co of managing the Multi-Technology Mix (Quigley, 2019).

Option B1: a publicly listed InfraCo to buy NBN Co

In the following, it is assumed that any ownership of InfraCo by Telstra, or indeed by any other vertically integrated carrier, has been reduced to a level providing no effective control of InfraCo, e.g. less than 10%. I do not suggest 0%, as in floating InfraCo it may be strategically important to new investors to see that InfraCo's major wholesale customer, Telstra, retains some financial motivation in continuing to use InfraCo's network services.

Advantages. Firstly, the synergies between InfraCo and NBN Co, described above, would provide a significant financial boost to the owner of the merged entity, NetCo.

Secondly, the merger would provide some insurance to InfraCo regarding the potential of a future privatised and unleashed NetCo to attack InfraCo's core backhaul business, e.g. by focussing only on the most profitable intercity routes.

Thirdly, the transfer of NBN Co's nation-building role to the private sector may minimise its electoral risk to the government. However, to meet national policy objectives (in support of the digital economy and digital society), NetCo will need to be subject to strong regulatory control of its pricing, its reach and its performance (in the form of monitored Service Level Agreements with its customers).

Fourthly, the sale would generate a windfall of $27b or more to the federal government. However, this outcome could be seen as a major disadvantage by the government of the day. Federal accounting practices will crystallise the government's investment loss in NBN Co (of currently $51b minus $27b = $24b) onto its bottom line, possibly dragging the federal budget into deficit in that year. This would be an unattractive outcome to any federal government that places the avoidance of budget deficits as a very high priority. Hence, the timing of the sale of NBN Co becomes politically quite crucial.

Disadvantages of a privatised NetCo. Firstly, there is the loss of long-term utility revenue to the government.

Secondly, a privately owned NetCo is likely to prioritise profits over service. (We have seen this with other recent privatisations.) However, there is a solution to this: creating regulation 'with teeth'.

Thirdly, a privately owned NetCo can be expected to use its financial muscle to lobby to change its charter, e.g. to increase its profits at the expense of universal reach or universal pricing. The solution to this is to incorporate the government's intended national goals for NetCo into strong legislation.

Lastly, major foreign ownership of NetCo would tend to accentuate the push for profits over other goals. The solution is to legislate strict ownership limits, as for Telstra Ltd.

Conclusions

It is vital to clarify our national policy goals for our digital society and our digital economy before deciding the future of the NBN. Otherwise, its future will largely be left to the market – and we have seen how well that has worked over the past twenty years! In particular, we note how "leaving it to the market", in the absence of national policy goals for how the NBN should support the digital economy, has created a national network which is largely optimised for passive Netflix watchers, given the government's weak statement of expectations for the NBN. Only the approximately 21% of premises who receive FTTP, as a legacy from the original NBN implementation, provide highly competitive advantage to SMEs operating in the global digital economy – and most of the FTTP service areas are located in the capital cities.

This paper suggests two worthy policy goals, one for the digital society and one for the digital economy:

1. All residents should have affordable broadband access to essential online services, irrespective of the location of their residence (to the maximum extent possible).
2. All Australian businesses should have access to broadband infrastructure at internationally competitive bandwidths and prices, irrespective of their location (to the maximum extent feasible).

If one supports those broad policy goals, one can assess and compare each "NBN Future" option in the light of how well it will achieve them.

In considering the option of merging Telstra's InfraCo business with NBN Co, the following observations have been made.

Firstly, it would not be acceptable (under either competition law or good policy) to allow InfraCo to buy NBN Co until InfraCo ceases to be controlled by Telstra or by any other vertically integrated carrier. But this independence can in fact be achieved by a process which Telstra's CEO seems to be favouring: InfraCo's divestment from Telstra, either by an IPO or by sale to an independent investor.

Secondly, NBN Co and InfraCo have entirely complementary networks. Their merger (as "NetCo") would serve the national interest better than NBN Co alone, creating an end-to-end fibre network that can support new entrants to the 5G (and later generation) mobile markets. Without this cost-effective assistance, new market entrants are likely to struggle to compete effectively with the market dominance which Telstra's 5G network is expected to achieve by 2022 at the latest.

Thirdly, NBN Co and InfraCo have synergies which can be crystallised via their merger to provide a balance between additional financial dividends to the new owner and reduced pricing to its wholesale customers. Or the owner can use the captured value of the merger to invest in the technology upgrades necessary for the network to become or remain fit for purpose, and hence more profitable.

Whether NetCo becomes a public or private monopoly, its pricing and performance will need to be strongly regulated. This is important because of NetCo's crucial role in supporting the national digital economy as a whole.

Lastly, the pros and cons of public versus private ownership have been considered. An economic rationalist (as distinct from a free-market religionist) would see that cost and governance advantages flow from public ownership. However, the aversion of governments these days to electoral risk and budget deficits may triumph over the need to achieve more practical policy goals for the country. The timing of the sale of NBN Co will therefore be quite crucial. The sale would seem to be an unattractive option until the fully rolled out NBN has time to pay off about half of the government's current $51b investment in it, in order that the net value of the transaction can have a negligible impact on the bottom line of that year's federal budget.

So, due to the perverse impact of federal accounting rules, it may be in the interests of both sides of politics for the NBN to remain in public ownership for an extended period of time.

Disclosure

The author has held shares in Telstra since "T1". This has not deterred him from authoring articles which favour the long-term interests of end users over the interests of shareholders.

References

ANAO [Australian National Audit Office]. (1998). Sale of One-third of Telstra, *Audit Report No. 10*, 19 October. Retrieved from https://www.anao.gov.au/sites/default/files/anao_report_1998-99_10.pdf

ANAO [Australian National Audit Office]. (2000). Second Tranche Sale of Telstra Shares, *Audit Report No. 20 2000-2001*, 30 November. Retrieved from https://www.anao.gov.au/sites/g/files/net616/f/anao_report_2000-2001_20.pdf

ANAO [Australian National Audit Office]. (2008). Third Tranche Sale of Telstra Shares, *Audit Report No. 43 2007-08*, 24 June. Retrieved from https://www.anao.gov.au/sites/default/files/ANAO_Report_2007-2008_43.pdf

Campbell, L.H. & Milner, M. (2019). The NBN Futures Forum. Discussing the Future Ownership of Australia's National Broadband Network, *Journal of*

Telecommunications and the Digital Economy, *7*(3), September. https://doi.org/10.18080/jtde.v7n3.202

Chirgwin, R. (2018). Telstra reveals radical restructure plan, *The Register*, 20 June. Retrieved from https://www.theregister.co.uk/2018/06/20/yet_another_restructure_for_struggling_telstra/

Crozier, R. (2018). NBN is building an “insight lab”, *itnews*, 6 April. Retrieved from https://www.itnews.com.au/news/nbn-co-is-building-an-insight-lab-488139

Department of Communications and the Arts. (2016). NBN Co Ltd Statement of Expectations, 24 August. Retrieved from https://www.communications.gov.au/publications/nbnstatementofexpectations

Department of Finance. (no date). NBN Co Ltd (nbn), *Commonwealth Government Business Enterprises*. Retrieved on 1 September 2019 from https://www.finance.gov.au/gbe-directors-guide/gbe/nbn/

Department of Industry, Innovation and Science. (2018). *Australia’s Tech Future: Delivering a strong, safe and inclusive digital economy*, 19 December. Retrieved from https://www.industry.gov.au/sites/default/files/2018-12/australias-tech-future.pdf

Duke, J. (2019a). Fletcher rules out NBN sale to Telstra, *Sydney Morning Herald*, 10 July. Retrieved from https://www.smh.com.au/business/companies/fletcher-rules-out-nbn-sale-to-telstra-20190709-p525j0.html

Duke, J. (2019b). ‘Incredibly damaging for the NBN’: Telstra boss warns 5G disruptors will take customers, *Sydney Morning Herald*, 31 July. Retrieved from https://www.smh.com.au/business/companies/incredibly-damaging-for-the-nbn-telstra-boss-warns-5g-disruptors-will-take-customers-20190731-p52cda.html

Fifield, M. (2018). Telecommunications Universal Service Guarantee, 5 December. Retrieved from https://www.minister.communications.gov.au/minister/mitch-fifield/news/telecommunications-universal-service-guarantee

Gerrand, P. (2004). Revisiting the Structural Separation of Telstra, *Telecommunication Journal of Australia*, *54*(3), 15-28. Republished online in Gerrand (2017).

Gerrand, P. (2017). Historical paper: The 2004 Proposal for the Structural Separation of Telstra, *Australian Journal of Telecommunications and the Digital Economy*, *5*(4), 70-86, December. https://doi.org/10.18080/jtde.v5n4.134

Gregory, M. A. (2019). How to Transition the National Broadband Network to Fibre To The Premises, *Journal of Telecommunications and the Digital Economy*, *7*(1), 57-67. https://doi.org/10.18080/jtde.v7n1.182

Kruger, C. (2019). Telstra CEO warns of NBN's 'unnatural distortion' of broadband market, *Sydney Morning Herald*, 30 July. Retrieved from https://www.smh.com.au/business/companies/telstra-ceo-warns-of-nbn-s-unnatural-distortion-of-broadband-market-20190730-p52c30.html?js-chunk-not-found-refresh=true

NBN Co. (2010). *NBN Co. Business Case Summary*, NBN Co Limited, 24 November. Retrieved from http://images.smh.com.au/file/2010/11/24/2061700/NBN%20Co%20%20Business%20Case%20Summary.pdf

NBN Co. (2018). *Annual Report 2018*. Retrieved from https://www.nbnco.com.au/content/dam/nbnco2/2018/documents/media-centre/nbn-co-annual-report-2018.pdf

NBN Co. (2019). *Corporate Plan 2020-23*. Released 29 August. Retrieved from https://www2.nbnco.com.au/corporate-information/about-nbn-co/corporate-plan

Pearce, R. (2019a). Comms minister says NBN privatisation 'some way away', *Computerworld*, 11 June. Retrieved from https://www.computerworld.com.au/article/662719/comms-minister-says-nbn-privatisation-some-way-away/

Pearce, R. (2019b). Path still open for InfraCo role in privatised NBN: Telstra CEO, *Computerworld*, 31 July. Retrieved from https://www.computerworld.com.au/article/664799/path-still-open-infraco-role-privatised-nbn-telstra-ceo/

Penn, A. (2018). T22 – our plan to lead, *Telstra Exchange*, 20 June. Retrieved from https://exchange.telstra.com.au/telstra2022-our-plan-to-lead/

Pugh, N. (2019). The Wireless Threat to Fixed Broadband Services. *Journal of Telecommunications and the Digital Economy*, *7*(1), 7-19. https://doi.org/10.18080/jtde.v7n1.178

Quigley, M. (2019). What happened to broadband in Australia?, *The Monthly*, March. Retrieved from https://www.themonthly.com.au/issue/2019/march/1551445200/michael-quigley/what-happened-broadband-australia

Ramli, D. (2016). NBN worth $27b despite $56b construction cost, says PWC, *Sydney Morning Herald*, 18 February. Retrieved from https://www.smh.com.au/business/nbn-worth-27-billion-despite-56-billion-construction-cost-says-pwc-20160217-gmwbd5.html

Reichert, C. (2017). USO to be axed in 2020 for Universal Service Guarantee, *ZDNet*, 20 December. at https://www.zdnet.com/article/uso-to-be-axed-in-2020-for-universal-service-guarantee/

Ritchie, G. (2019). Why Low-Earth Orbit Satellites are the New Space Race, *Bloomberg Businessweek*, 9 August. Retrieved from https://www.bloomberg.com/news/articles/2019-08-09/why-low-earth-orbit-satellites-are-the-new-space-race-quicktake

Telstra. (2019). *Financial results for the half-year ended 31 December 2018*. Retrieved from https://www.telstra.com.au/content/dam/tcom/about-us/investors/pdf%20F/140219-Financial-results-for-the-half-year-ended-31-December-2018.pdf

Vocus. (2019). Business Broadband Plans. Retrieved from https://www.vocus.co.nz/broadband-plans

Glossary

5G = 5th Generation mobile network

ACCC = Australian Competition and Consumer Commission

EBITDA = Earnings Before Interest, Tax, Depreciation and Amortization

FTTC = Fibre to the Curb

FTTN= Fibre to the Node

FTTP = Fibre to the Premises

FY18 = Financial Year 2018, i.e. July 2017 to June 2018

HFC = Hybrid Fibre Coaxial [cable]

IPO = Initial Public Offering, i.e. floating the business on the stock exchange

NBN = National Broadband Network

RSP = Retail Service Provider

SME = Small or Medium Enterprise

USG = Universal Service Guarantee

Endnotes

[i] The Telstra 3 sale objectives agreed by the then Minister for Finance and Administration in August 2005 were as follows (ANAO, 2008):

- achieve an appropriate financial return from the sale;
- promote orderly market trading of Telstra shares;
- secure a timely sale process, conducted to the highest standards of probity and accountability;
- support Australia's reputation as a sound international investment location;
- continue to build investor support for the Government's asset sale programme and broaden share ownership; and
- remove the Government's conflict of interest as owner and regulator of Telstra.

[ii] The federal government’s Strategy for Australia’s Tech Future, published in December 2018, aims to deliver “a strong, safe and inclusive economy, boosted by digital technology”. It admirably proposes that “Australians have access to world-class digital infrastructure in their personal and working lives” (Department of Industry, Science and Technology, 2018). Yet there is a major disconnect between the aim of “world-class digital infrastructure” and what the NBN is actually delivering for the majority of premises in Australia, as evidenced by Table 1 above. This new government strategy appears to have emerged too late to influence the design of the current NBN; and it is perhaps significant that it was not issued jointly by the Minister for Communications, Cyber Safety and the Arts.

[iii] See, for example, the opinions expressed by two of my fellow panellists at the 31 July 2019 NBN Futures Forum, cited in Campbell & Milner (2019) in this issue of the *Journal*.

Getting the NBN Infrastructure We Need

Dr Jim Holmes
Director, Incyte Consulting

Abstract: This article summarises the presentation given by the author at the TelSoc NBN Futures forum held in Melbourne on 31 July 2019. The author spoke in favour of retaining NBN Co in public ownership, at least for the medium term and until a long-term plan and evolution pathway has been established. Such a plan is needed to ensure that Australians have affordable access to world's best broadband service and that delivers social and economic inclusion. Dr Holmes argued that there are positive reasons for supporting public ownership, and that this was the least worst of the options available.

Keywords: NBN, public policy

Introduction

The discussion that we had under the aegis of TelSoc in Melbourne on 31 July 2019, and which we are now continuing through the *Journal*, is ostensibly about the structural and ownership options for the National Broadband Network and the future of NBN Co. But these are secondary questions that do not need to be considered until after we address the question of the sort of NBN infrastructure that we need as a society and as an economy.

Broader policy issues must determine industry structure and ownership, not the other way around. If we rush to judgment on ownership issues, without first settling the policy objectives for the sector or the context in which such decisions will play out over the long term, the result is almost certain to be simply an ideological preference for competition, private enterprise, socialism, or whatever. This would be a disservice to the current and future generations. Broad agreement on the policy objectives will give us the criteria for assessing proposals for industry structure and ownership, and therefore such policy issues need to be considered first.

Policy Objectives

Part of the conversation that TelSoc seeks to have is to create a workable consensus on the policy objectives that need to be achieved through the NBN.

http://doi.org/10.18080/jtde.v7n3.199

Generally, there is surprisingly little disagreement about most of the objectives – although some contention might occur about the relative priority of each objective.

The objectives that I would suggest as being relevant to the NBN are as follows, in no particular priority order:

- A key source of international competitive advantage;
- Affordable, reliable and accessible for all users;
- Permitting maximum competition for wholesale and other retail services and technologies;
- With a future development pathway that sustains the above;
- With a regulatory regime that delivers on all of the above.

The first thing to note with the above list is that the words used are general and will require further analysis as the objective is implemented. Admittedly, there is a risk that, at this level of generality, the objectives could become mantra without practical meaning. On the other hand, greater specificity risks losing focus altogether as we all rush down one rabbit hole after another.

The second thing to note is that the objectives are continuing. They are not going to be achieved once and for all. They need continuous effort.

The third thing to note is that the first two objectives are expressed as goals – albeit with moving goal posts. However, the last three are process objectives, which draw from other areas of policy preference in Australia. NBN policy is not something separate from the body of policies on which there is some measure of widespread public support in Australia, namely competition, economic progress and infrastructure development.

So, there is a public policy commitment to facilitating, where feasible and appropriate, the dynamic of competition as the engine to drive efficiency, innovation and value creation. It is not the only thing that matters, but it does matter. (The benefits of competition have been outlined previously by McLaren (2018).)

The fourth objective in the list above requires that our approach to the NBN recognises that we are involved for the longer term, and that structural and ownership proposals need to permit pathways to be developed, and to evolve, to address the changing social, market and technology environments that may emerge. We cannot lock in assumptions based on current markets, usage patterns and technologies.

The last objective listed recognises that NBN Co and its successor organisation(s) will have significant market power and be such fundamental infrastructure that it will require a carefully developed regulatory framework that, itself, must be regularly adjusted to keep pace with the changing environment.

Starting Points

Much NBN discussion seeks to review past decisions and to analyse the quality of past decision making in this area. This can be useful in a discussion about pathways for NBN future evolution, but all too often it serves to focus on the past and on blame. On many major issues across the whole spectrum of public policy – not just about the NBN – there are calls for a respectful national conversation. All too often these quickly deteriorate. Perhaps we can change that in the case of the NBN by resolutely focussing on where we actually are at present, and I would characterise this as:

- A collection of (mainly) fixed access networks that reflect a multi-technology mix and seek to extend the economic life of some older transmission assets (such as copper cabling);
- A delayed rollout that is now picking up and which should now be completed by 2022 (after the migration period);
- Public financing in the order of $A51B and a value (after estimated write downs) of around $A24B; and
- A Government commitment to privatisation at some point that would be as soon as practicable after the completion of the NBN and customer migration.

Further details about the current state of the NBN and future options can be found in Gregory (2018) and Gregory (2019).

The Government planned to privatise the NBN during the current parliamentary term, but the requirements for NBN completion and migration now make that extremely unlikely. The rollout will continue and be completed more or less on the new timetable. The extent of the final Government debt and the value that will need to be written down (or off) will be in the order of magnitude indicated above. These are large amounts, and a loss of this magnitude appears unavoidable as accounting begins to catch up with reality.

Public Ownership – More Than a Default Option

If nothing further happened once the current NBN building program and customer migration process are completed, then the NBN would remain in public ownership. In that sense public ownership is a default option. It is where we are at present, and – if nothing changes (legislation aside) – that is where we stay.

But public ownership for the NBN into the foreseeable medium-term future is more than a default option. During that period, it has positive features as outlined below, and is to be preferred over many of the other options now being put forward.

The key arguments for public ownership are:

- **Long term fundamental infrastructure**

The NBN is infrastructure that is fundamental. It is the basis on which Australia's digital economy and future online society will rest. Our national communications infrastructure will determine how we interrelate, learn, and participate socially and economically in the longer term. Its importance cannot be underestimated.

A long-term perspective is necessary to match the long term and fundamental nature of the infrastructure involved. We should not risk serious compromise by subjecting the NBN to the short-term profit-maximising imperatives of private ownership. At least we should not do that until some very important frameworks have been put in place and tested.

Taking this sort of infrastructure and the means for evolving it to meet emerging and developing needs is not something that can be specified in a sale process – at least not until a lot of other pre-conditions are also in place. Importantly, what serves the long-term interest cannot always be identified in a short-term process due to changing circumstances over time. Public ownership does not preclude efficient management – nor does it guarantee it. But it does facilitate a longer-term perspective for critical decision making, free of the short-term pressures for commercial levels of profit returns to shareholders.

Of course, private companies can manage their own long-term assets in the interests of their shareholders, but, where there is effectively a single service provider and the assets cannot be economically replicated, the overall public interest requires continuous government intervention, which in turn is restricted by the private property rights involved.

- **Avoiding privatisation of monopolies**

Picking up from the last point, there is a well understood recent tendency in Australia to privatise government enterprises that operate in competitive markets. Some might extend this to the privatising of government enterprises that should, or are planned to, operate in competitive markets. The restructuring and privatising of Qantas is an example of the more limited principle – the international airline industry was and remains competitive – and the privatisation has not since been called into question. Electricity generators and Telstra are in the category of enterprises that have been privatised for operation in competitive markets, and have proved to be much more troublesome.

In the case of Telstra, the enterprise was privatised in tranches, as a vertically integrated and very powerful incumbent operator. It was assumed that the legislative settings and the regulatory framework were adequate to ensure that the achievement of competition and other public policy goals would be facilitated. The worst assumption of all was that the Government and bureaucracy continued to treat Telstra as if it was an instrument for achieving national policies, much as it was prior to privatisation and to the market liberalisation of the early 1990s. They appeared to be surprised to find that Telstra had its own commercial agenda, even during the period when the government maintained a majority shareholding.

I am not arguing that Telstra should have remained in public ownership: only that an ideological commitment to complete privatisation and the failure to foresee some of the issues that have arisen since should not be repeated with NBN Co.

It is true that NBN Co is a purely wholesale operation and that the issues associated with Telstra being vertically integrated will not arise. But it will have substantial market power (even monopoly power outside inner city and some other urban areas) and the policy and regulatory framework that is intended to limit potential abuses of market power have not been developed for a privatised NBN Co, let alone road-tested.

Some of the proposals for restructuring and ownership of NBN are driven by a desire to encourage competition. The Vertigan Committee Report released in October 2014 (Vertigan, 2014) proposed disaggregation of NBN Co along technology lines to encourage inter-platform competition. Competition will be forthcoming from other sources in any case, with the development of mobile data technologies such as 5G[i], and based on the economic feasibility of network duplication in CBDs and some other urban settings. Future competition may be in some markets only, rather than across NBN’s complete range of services. However, competition will likely be as much about contestability rather than actual contests in the early post-rollout stages of the NBN. We can therefore reasonably plan around the NBN having effective market dominance for the medium-term and possibly longer.

- **Planning**

Following on from the last points, the imperative for national planning (and continuous plan review) cannot be overstated. This is much more than the outcomes from competitive markets, where planning is for competitive advantage and of limited transparency. The certainty generated from transparent national infrastructure planning is important not only for the NBN itself, and for its own financial and operating plans, but for all users.

If the NBN is to transform the economy and ensure that it becomes even more globally integrated and competitive, the industries that develop large-scale applications or are otherwise using the network need to understand how the NBN now and in future will affect their plans. Investment in nationwide applications and services is contingent on the certainty and transparency of NBN public planning. The same applies for major public sector programs based on e-Health, e-Education and e-Government.

I do not want to overstate the point, or to suggest that these developments that promote usage and economic and social inclusion will not happen. My argument is that delay, uncertainty and fragmentation have costs that can be significantly avoided through appropriate decisions for the NBN's ownership and industry structure.

- **Only Government can do certain things**

The provision of services required for social and economic inclusion, but which are not commercially feasible, can only be provided by Government. Government can employ private contractors to discharge universal service obligations, either directly or through programs run by Government-owned entities. However, in the case of the NBN, the planning and costs of programs for delivering above-cost essential services requires substantial planning, financing and forward procurement, especially satellite systems and capacity for the 7% of services that cannot be delivered by terrestrial technologies alone. These matters are best planned and executed by a publicly owned and publicly accountable enterprise, reporting to an equally publicly accountable Government and Parliament.

Public Ownership – Not Necessarily Forever

Public ownership of NBN need not be forever. Such a view would be as ideological as insisting that all enterprises in this industry (or any other) must be privately owned.

The development of a suitable policy and regulatory framework should also contemplate the pre-conditions that need to be satisfied before all or part of the NBN is privatised. That policy will not include an artificial timetable for ready disposability such as we have at present. This is not the time to argue what those pre-conditions might be, but they would be premised on the market for infrastructure broadband services having "settled down" and also that regulatory frameworks would have been stress-tested over several years.

Other Options

The table below was presented at the TelSoc event in Melbourne on 31 July 2019. It sets out some other options for structure and ownership that were raised by speakers at the event, and a preliminary assessment against the criteria discussed earlier in this paper.

Table 1. Some NBN Ownership and Structural Options Assessed and Compared

	InfraCo merger	Public ownership	Privatise	Technology disaggregation
1. International competitive advantage	Doubtful	Maybe	No, due to chronic short termism	No – competition seems the only consideration
2. Affordable, reliable and accessible to all	Unlikely	Yes	No	Maybe
3. Maximum competition for other wholesale operators	Limited at best	Yes, but subject to regulation	No	Yes, but subject to regulation
4. Future development pathway	Yes, but limited by past choices	Yes	Unlikely	Very unlikely
5. Effective regulation	Unclear – could be Telstra all over again	Yes, with public accountability	No – Telstra all over again	Unlikely

The first thing to notice is that the assessments are tentative. This is because the options are broadly defined and have yet to be detailed. The way in which each option is developed will undoubtedly seek to address the criteria in the left-hand column in greater detail. This applies to the public ownership option as well.

Apart from public ownership, I conceive of the other options, pending further clarification from protagonists, as follows:

- **InfraCo merger:** This option involves the merger of Telstra's backhaul network with NBN Co to form a comprehensive broadband network with both access and backhaul transmission capabilities. The option is open to public or private ownership, although only the latter would seem to be under contemplation, because the renationalisation of a significant part of Telstra is unlikely. The government has made it clear that Telstra is not a suitable buyer for NBN Co, because that would reinforce dominance in a vertically integrated organisation. This view aligns with Telstra's preference for selling InfraCo in any case. The main advantage claimed for an InfraCo merger is that the constituent parts are complementary and that there would be significant synergies that will improve the provision of a more comprehensive range of wholesale broadband services.
- **Privatisation:** This option is the one already included in legislation and the Government's announced policy. The claimed benefits are that it puts the enterprise on a commercial footing and that the imperatives of private ownership will ensure

that it operates efficiently and can access private capital for future investment in growth and development. Implicitly, the option relies on effective regulation to control a private monopoly or near-monopoly.

- **Disaggregation by technology:** This option picks up the recommendations of the Vertigan Committee (Vertigan, 2014), which favoured bringing forward platform competition by disaggregating NBN Co into separate enterprises defined by the assets and the related technologies that each would be allocated. Therefore, the new enterprises would be based on fibre and copper, HFC, fixed wireless and satellite, respectively. Presumably each of the new enterprises would seek to maximise the potential of their "starting technologies" but be free to build or lease infrastructure incorporating other technologies as required. The assumption remains that each of the disaggregated enterprises would be privatised at the earliest time, but that they need not be sold off at the same time. Conceivably, the satellite operation serving rural and regional areas could be sold last, or even not sold at all.

My preliminary assessments of each option against the criteria are:

- **InfraCo merger:** This option takes a large and troublesome enterprise and makes it larger. Whatever monopoly cultures are involved, they are likely to strengthen, possibly at the expense of achieving the full value of the synergies that are being claimed. The impact is likely to be negative on efficiency and, then, because costs will be higher than they otherwise might be, on delivering innovative and affordable services. This form of amalgamation is unlikely to facilitate competition in the sector – unless poor performance of the merged entity attracts competition. Lastly, whether our regulatory systems are able to effectively constrain and guide the behaviour of such a behemoth – especially in private hands – should not be assumed. This is time for regulatory scepticism.
- **Public ownership:** My preliminary assessment is already stated. However, we should recognise that the possibilities of under-performance, of regulatory non-compliance and of abuse of market power are equally present, whether an enterprise is in public or private ownership. We should maintain a sceptical mindset in both cases. However, the opportunities for public accountability are increased with public ownership, including reporting to the Parliament and Ministerial accountability.
- **Privatisation:** The short-term profit imperatives of private ownership structures and the need to deliver returns that are considered to be sufficiently commercial are the main concerns for this option, because of the impact on wholesale prices and ultimate affordability to broadband service users. If prices are constrained, profits are likely to be maintained by reducing service performance and quality below what they

might otherwise be. This will have serious economic knock-on effects. Private owners of public infrastructure require levels of certainty and this could constrain competition in the medium to longer term and require future government intervention or supplementary funding. Scepticism is our best guide here as well.

- **Disaggregation by technology:** This option really needs to be further spelled out because it is entirely unclear whether the complexity of implementation is worth the effort. There are other sources of platform competition as mentioned above; they do not have to come from within NBN Co. We should be sceptical about the incentives for the "baby NBN Cos" to move beyond the technology that defined their share of the starting assets, and whether, in private hands, there would not be substantial pressure to exploit to the maximum, rather than to the optimum, the assets that have been purchased. This option is unlikely to provide the sort of coherent national pathway for broadband service evolution and development that we need. It seems to be driven by a notion of platform competition and of competition generally, even though, in practice, it might well be a tortuous and ineffective way of delivering it. The assurances that private buyers will seek will, if delivered, lead to the same problems for disaggregated sales as for an aggregated (whole of NBN Co) sale.

Reality Test

The options that were discussed at the TelSoc event on 31 July 2019 and which are referred to here are in the nature of primary colours that are unlikely to result from the present exercise. All of the options, and more than have been discussed, will undoubtedly be qualified and modified by the interplay of ideas. They will take on shades and hues appropriate to ideas that need to be implemented in a complex economic, social and political environment. That process of refinement is no more than making them as fit for purpose as they can be, and is the reason that TelSoc has sought to be one platform in a useful public debate.

Preliminary Conclusion

My preliminary conclusion is that public ownership for at least the medium term is necessary to ensure that a long-term plan for long-term sustainability and public accountability is delivered. I reserve the right to modify that view as I hear new and relevant views out of the ensuing discussion.

References

Gregory, M. A. (2018). Australian Wholesale Telecommunications Reforms, *Journal of Telecommunications and the Digital Economy*, *6*(2), 1-34. https://doi.org/10.18080/jtde.v6n2.155

Gregory, M. A. (2019). How to Transition the National Broadband Network to Fibre To The Premises, *Journal of Telecommunications and the Digital Economy*, *7*(1), 57-67. https://doi.org/10.18080/jtde.v7n1.182

McLaren, G. (2018). What Now for Australia's NBN? *Journal of Telecommunications and the Digital Economy*, *6*(4), 31-62. https://doi.org/10.18080/jtde.v6n4.162

Vertigan, M. (2014). *Independent cost-benefit analysis of broadband and review of regulation: National Broadband Network Market and Regulatory Report*. Department of Communications, 14 August (Vertigan Committee Report to Minister for Communications). Retrieved from https://www.communications.gov.au/sites/g/files/net301/f/NBN-Market-and-Regulatory-Report.pdf

Endnote

[i] There is considerable diversity in the views of experts about the potential impact of 5G technologies and the development of mobile broadband generally on the NBN in future. The point being made is that the NBN is not immune from the effects of technology change and will need to adopt those technologies itself, or respond to them in appropriate ways.

S-MANAGE Protocol for Provisioning IoT Applications on Demand

Thi Minh Chau Nguyen
University of Technology Sydney, Faculty of Engineering & IT

Doan B. Hoang
University of Technology Sydney, Faculty of Engineering & IT

Abstract: Internet of Things (IoT)-based services have started making an impact in various domains, such as agriculture, smart farming, smart cities, personal health, and critical infrastructures. Sensor/IoT devices have become one of the indispensable elements in these IoT systems and services. However, their development is restricted by the rigidity of the current network infrastructure, which accommodates heterogeneous physical devices. Software-Defined Networking-Network Functions Virtualization (SDN-NFV) has emerged as a service-enabling solution, supporting network and network function programmability. Provisioning IoT applications on demand is a natural application of programmability. However, these technologies cannot be directly deployed in the sensing/monitoring domain due to the differences in the functionality of SDN network devices and sensor/IoT devices, as well as the limitation of resources in IoT devices. This paper proposes an S-MANAGE protocol that preserves the SDN-NFV paradigm but provides a practical solution in controlling and managing IoT resources for provisioning IoT applications on demand. S-MANAGE is proposed as a new southbound protocol between the software-defined IoT controller and its IoT elements. The paper presents the design of S-MANAGE and demonstrates its use in provisioning IoT services dynamically.

Keywords: Provisioning services on demand, Software-defined IoT model, Programming services, Network functions virtualization, Software-defined virtual sensor (SDVS)

Introduction

An Internet of Things (IoT) environment accommodates numerous IoT devices (we use 'IoT devices' to include networked sensors in this paper) with various sensing, computing, communicating, actuating capabilities, and resources. The number of IoT devices in the world is predicted to be about 6.58 billion by 2020 (Perera *et al.*, 2014). Deployment of an IoT application can become challenging owing to large geographical coverage of the application, limitation of resources of IoT devices, heterogeneity of the environment, and a huge number of these devices (Li, Xu & Zhao, 2018). IoT applications often overlay and share the deployments of IoT devices, and this presents difficulties and challenges in the interaction and sharing of information between the devices and the applications (Li, Xu & Zhao, 2018). Therefore, many efforts have been put into programming IoT devices to meet IoT application demands (Javed *et al.*, 2018).

Among solutions to the programmability of Wireless Sensor Networking/Internet of Things (WSN/IoT) systems, many proposals have taken advantage of the Software-Defined Networking (SDN) paradigm (Bera, Misra & Vasilakos, 2017). The SDN provides solutions to programmability, agility, flexibility and end-to-end connectivity challenges, which are associated with management of real-time traffic flows and dynamic traffic patterns (Deva Priya & Silas, 2019). The SDN approach addresses many existing problems concerning network management and provisioning resources required by network services. It can change the functionality of physical networks as well as devices in real time to meet requirements of IoT applications (Sood, Yu & Xiang, 2016). The SDN principle separates the network control plane from the data plane of networking devices and allows the provision of on-demand services through a programmable and logically centralised controller. Autonomous management of network devices is enabled under SDN. The SDN architecture comprises three main planes, which are application plane, control plane, and data plane. The control plane centrally controls and manages the behaviour of the whole network in the data plane via a Southbound Interface (SBI). To provide network services to the application plane, it uses a Northbound Interface (NBI) to expose the abstraction of the underlying network.

However, challenges remain when applying the SDN paradigm to the constrained WSN/IoT (Luo, Tan & Quek, 2012). As a fundamental element of the underlying resources that provide necessary data for IoT applications, an IoT system must of necessity not only control and manage the underlying resources but also orchestrate them to satisfy application demands. However, architectural solutions for provisioning various IoT applications are still immature. A majority of the proposed approaches are vertically integrated, so it is difficult for the infrastructure to handle various IoT application demands that require horizontal capabilities from other subsystems. While many attempts have been made to address IoT platform architectures and to provision IoT applications on demand, challenging issues remain: they include scalable and dynamic resource discovery and composition; context-awareness; integration of intelligence; interoperability; reliability; security and privacy; and system-wide scalability (Razzaque *et al.*, 2016). In this paper, we propose the S-MANAGE protocol as an enabler of the software-defined Internet of Things (SD-IoT) model, suggested in our previous work (Nguyen, Hoang & Dang, 2017).

The SD-IoT model adopts SDN and Network Functions Virtualization (NFV) principles and deploys these technologies to IoT devices to provide IoT applications on demand. However, it is not feasible to completely and directly apply the SDN technique to the resource-limited WSN/IoT environment, owing to its constraints (Kobo, Abu-Mahfouz & Hancke, 2017). In the model, the NFV technique is deployed to realise software-defined virtual sensors as a representation of the underlying IoT resources. This technology is thus applied readily to the WSN/IoT environment for creating a virtual representation of IoT devices that are utilised by multiple IoT applications simultaneously. This virtual representation offers a solution to enrich the features of limited IoT devices. By applying both SDN and NFV principles in the proposed SD-IoT model, diverse underlying sensor nodes can be programmed in accordance with IoT application requests.

In the SD-IoT model, the S-MANAGE protocol has been proposed as a communication bridge between the SD-IoT controller and the software-defined virtual sensor (SDVS) in each cluster. The controller sets up and configures the SDVS by using S-MANAGE, which is designed to deal with the constraints of IoT systems. It should be emphasised that the SDN OpenFlow protocol was designed to handle SDN network (routing) devices and it is not suitable for resource-constrained IoT devices, whose mission is different from that of SDN routers and switches. Furthermore, a separate protocol such as OF-CONFIG is often required to configure the network devices, and this introduces complexity to already constrained IoT devices. This paper investigates the design and the implementation of S-MANAGE. S-MANAGE is designed both to configure SDVSs and to control the behaviour of the underlying networked IoT resources. Major contributions of this paper are as follows:

1. It proposes a new southbound S-MANAGE protocol for programming the behaviour and configurational management of software-defined virtual sensors and their associated physical devices.
2. It proposes a programmability approach to provisioning IoT applications on demand.
3. It describes an implementation of the proposed protocol in the context of a software-defined Internet of Things system and provides implementation results.

The remainder of the paper is organised as follows. Section II reviews related work. Section III describes the overall SD-IoT architecture. Section IV presents the design and the specification of the S-MANAGE protocol in terms of packet format, packet type, forwarding table, and configuring table. Section V describes an implementation scenario and evaluates its performance. Section VI concludes the paper.

Related Work

The most challenging task in applying the SDN paradigm to the WSNs/IoT environment is the design of the communication interface between the SDN-based controller and underlying IoT devices. However, existing proposals mainly suggest modification to the well-known OpenFlow SBI without an actual implementation: for example, Sensor OpenFlow (Luo, Tan & Quek, 2012), and SDWN (Costanzo *et al.*, 2012). The feasibility of application of SDN to WSNs has been demonstrated in the SDN-WISE (Galluccio *et al.*, 2015) proposal. The SDN-WISE SBI has been designed in accord with the OpenFlow protocol, thus enabling programmability of a WSN sink node's forwarding behaviour. The main components of the protocol are described and the design is evaluated via a real implementation. However, its main aim is to program a node's forwarding behaviour without concern for programming a sensor node's functionality. The details of these proposals are discussed in our previous work (Nguyen, Hoang & Chaczko, 2016).

Another work (Mahmud & Rahmani, 2011) deploys the OpenFlow technology in a WSN to enable a share of IoT resources for larger scale networks. They propose flow-sensors that communicate with an access point via the OpenFlow protocol. Their implementation results demonstrate how OpenFlow can be of benefit in controlling and monitoring sensor traffic flow, but there is no effort to make the OpenFlow protocol suitable for constrained sensor nodes.

Device and network management has been considered by the soft-WSN (Bera *et al.*, 2016) proposal. The proposed architecture is in accordance with the SDN paradigm and provides a network and device management approach for provisioning IoT application-aware services. However, the focus is on the design of the controller and the sensor node architecture. The controller is designed with two management policies, topology and device management. The communication between the two entities is based on traditional protocols, IEEE 802.15.4 and IEEE 802.11. There is no effort to solve the complexity of deployment of flow tables to the sensor nodes.

Software-Defined IoT Model

To reap the benefits of the SDN paradigm, the SD-IoT model is also structured in three layers – application, control, and data – as depicted in Figure 1.

The application layer is where developers can deploy their IoT applications without knowledge of the underlying IoT infrastructure.

The control layer accommodates the SD-IoT controller and its database. It is a bridge between the application layer and the data layer. It provides the application layer with a global view of the underlying resources as well as an efficient interface to control the underlying IoT resources. At the same time, it provides the underlying resources with an interface for updating their status and attributes, as well as their sensor services. With the knowledge of both requirements for and capabilities of the IoT resources, it can provide sensor services for IoT applications on demand.

The data layer hosts IoT devices or IoT infrastructure. Different from the SDN data layer, the SD-IoT data layer is designed with two sub-layers, called virtual and physical data layers. The virtual data layer is proposed as an interface between the SD-IoT controller and the physical IoT devices. The virtual data layer enables the controller to manage and control the underlying IoT resources in the physical data layer.

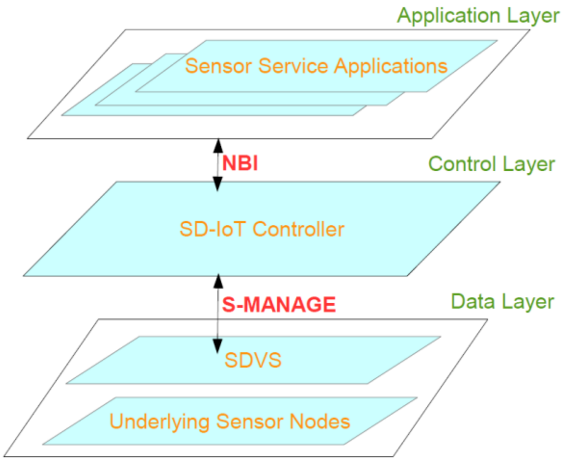

Figure 1. SD-IoT architecture

SD-IoT controller

The SD-IoT controller is responsible for i) processing application requests; ii) orchestrating resources; iii) updating knowledge of the underlying resources; and iv) controlling and managing the underlying resources according to IoT application demands.

To handle these responsibilities, the controller houses several core modules: application handler, resources manager and orchestrator, configuration manager, and a database for storing information concerning the underlying resources.

Specifically, the SD-IoT controller makes it possible for the application layer to specify its demands via an NBI. The controller interprets these requirements into the SD-IoT resource-specific language in order to orchestrate the resources to meet the IoT applications' requirements. To configure the underlying resources, the controller makes use of a southbound interface (SBI) defining resource-specific messages to configure the devices.

S-MANAGE protocol

The S-MANAGE protocol is proposed as an SBI between the SD-IoT controller and the virtual data layer. The protocol allows the controller to communicate with software-defined virtual sensors (SDVSs) in this layer. Moreover, via S-MANAGE, the controller can collect statistics regarding the physical IoT devices. It defines the message structure and message types exchanged between the controller and the virtual layer. These messages are for configuration management and control of behaviour of SDVSs.

Software-defined virtual sensor (SDVS)

Software-defined virtual sensors (SDVSs) are defined as representatives of physical/software sensor nodes or IoT devices located in the physical layer. It is configured with core features and attributes of a physical sensor node, and with a software-defined function (SDF). The core modules enable the SDVS to behave as the represented physical IoT device; the modules can interact with the represented devices via device-specific protocols. In addition, the SDF allows the controller to enhance the SDVS with processing, computing or forwarding functions. In particular, the forwarding and configuring functions are implemented in the SDVS as SDFs. The SDVS is connected to its underlying IoT device and acts on behalf of the represented device. In addition, it is installed with S-MANAGE protocol features, so it can communicate with the controller. In summary, an SDVS can be considered as a software layer that enriches its represented physical sensor/IoT device, allowing the SD-IoT controller to configure the device and program its behaviour. The design and implementation of SDVSs is the topic of another paper.

S-MANAGE Protocol

In traditional IP networks, routers are used to relay packets (or datagrams) according to the lookup table determined by a routing protocol. Packets are treated as independent elements not related to other packets that may belong to the same source-destination connection. Traffic flow is normally associated with packets belonging to an end-to-end TCP connection. In order to completely specify a flow at a router, a large number of identifiers are needed, including transport layer ID, network layer IP address, VLAN ID, MAC layer ID, and router port IP address. As a consequence, a flow in the OpenFlow protocol requires some 12 matching parameters to be identified. Clearly, this is not needed in sensor/IoT networks where the end devices are not routing devices in the traditional network. Many devices do not use TCP/IP; direct deployment of OpenFlow in WSN/IoT networks is not appropriate. Furthermore, the OpenFlow SDN network still requires OF-CONFIG or other protocols for device configuration. What we need is a streamlined protocol in WSN/IoT networks that can handle both configuration of the IoT devices and simple types of sensed data, but in the same spirit as flow in OpenFlow. S-MANAGE protocol is proposed to do just that. As the southbound protocol, the S-MANAGE protocol is proposed as a southbound interface between the SD-IoT controller and the virtual data layer. Via the SBI, the controller can both manage and configure the SDVS in this layer.

The S-MANAGE protocol is for managing and programming the SDVSs within the virtual data layer and indirectly via them to configure their represented physical devices. S-MANAGE makes it possible for the controller to program sensors or IoT devices, not only their forwarding behaviour but also other functionality.

The protocol is proposed according to the spirit of two protocols, OpenFlow (ONF, 2012a) and OF-CONFIG (ONF, 2012b). OpenFlow focuses on flow rules for setting, modification, and deletion, or adding rules for controlling forwarding behaviour of OpenFlow switches. Meanwhile, OF-CONFIG enables configuring an OpenFlow Switch itself as the setting of port number, IP address, or interfaces.

The protocol enables the management and configuration of representations of IoT devices to be based on two proposed instruction tables, called forwarding and configuring tables. The forwarding table instructs an SDVS on how to handle an arriving packet, while the configuring table guides the SDVS to configure its represented underlying nodes.

The protocol allows the controller to i) install instruction tables on the SDVS for configuration purposes, ii) get information concerning the SDVS's features, functions, and the status of its underlying sensors, and iii) collect statistics associated with the SDVS's operation, such as the number of processed packets or sensor services required by IoT applications.

In addition, via the protocol, the SDVS is able to i) update the controller with its status and attributes, and ii) ask for instructions on processing an incoming packet or configuring its underlying IoT devices.

S-MANAGE defines communication methods between the controller and an SDVS. It specifies exchanged message types between the two entities, the message format, the structure of instruction tables, and how the SDVS is programmed and should operate based on these tables' instructions. Details of the protocol design are described in the following sections.

S-MANAGE packet header

The S-MANAGE packet is comprised of a header and a payload. All S-MANAGE messages begin with an S-MANAGE header, as depicted in Figure 2. The header size is 10 bytes. It includes the following parts:

- Source Address (2 bytes) is an address of a source sending a packet.
- Destination Address (2 bytes) is a destination address of a packet.

- Next hop address (2 bytes) is an address of a hop in the list providing the path of a packet from the source to the destination.
- Type (1 byte) indicates a packet type.
- Packet length (1 byte) indicates the length of a packet including its header and payload.
- TTL (1 byte) is “time to live” of a packet. It is reduced by one at each hop.
- Message ID (1 byte) is an identifier of the packet type.

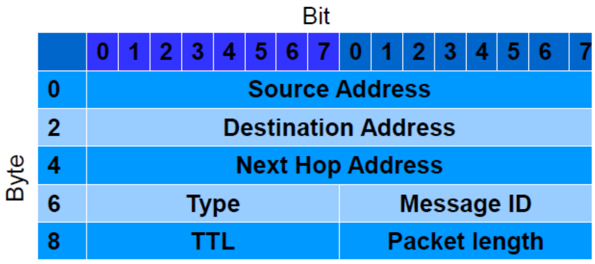

Figure 2. S-MANAGE packet header

Message types

The payload carries the content of a packet. Different types of packets carry different kinds of information that represent different purposes of a sender. Therefore, we define the following S-MANAGE message types to achieve the expected purposes.

The S-MANAGE message types are grouped into three categories, i) controller to SDVS, ii) asynchronous (SDVS to controller), and iii) symmetric (controller/SDVS to SDVS/controller). However, due to the constrained resources of the sensor nodes or IoT devices, the number of messages exchanged is minimised, and the messages are optimised.

Controller-to-SDVS message type

This message type is initiated by the SD-IoT controller and may or may not require a response. The messages for installation of forwarding and configuring instructions on the SDVS need no responses from the SDVS. This category includes messages such as SetForwardingInstruction, SetConfiguringInstruction, and ModifyConfiguration packets. However, if the controller demands an SDVS’s attributes or status, it needs the SDVS’s responses.

a) SetForwardingInstruction/SetConfiguringInstruction: Enable the controller to install a forwarding/configuring instruction on an SDVS’s forwarding/configuring table, and to respond to an SDVS’s requests for a forwarding/configuring instruction, respectively.
b) ModifyConfiguration: Modify a configuration instruction.
c) RequestFwdStats/RequestConfigStats: Get statistics of a forwarding or configuring instruction, respectively.

d) ResponseFwdStats/ResponseConfigStats: Sent from an SDVS to the controller whenever the SDVS receives a RequestFwdStats/RequestConfigStats message, respectively. These messages include information about the statistics of an instruction table or an instruction in the table.

e) RequestFeatures/ResponseFeatures: Get an SDVS's information about its sensor service list, the services' status, or the driver of the underlying IoT device.

Asynchronous message type (SDVS-to-Controller)

The message type is sent from an SDVS to the controller without any request from the controller. It enables the SDVS to ask for instruction on handling incoming packets, as well as to update the controller on changes in its underlying sensor nodes regarding their active/idle status or completion of a required task.

a) Report packet: Report the status and behaviour of an SDVS. Particularly, the controller will be updated by changes as follows.
    a. Update the controller on its features (ReportFeatures).
    b. Inform the controller about the removal of a configuration instruction from a configuring table (ReportConfigurationRemove).
    c. Notify the controller about a sensor node's battery level (ReportLowBatt).
    d. Notify the controller that a sensor is at its maximum level of handling requests, so it is unavailable to be assigned further tasks by the controller (ReportFullTask).
    e. Inform the controller about the completion of a required service by an SDVS (ReportCompletion).

b) Response message type: Send required services back to a required destination.

c) Request message type: An SDVS requests for an instruction for its operation. Particularly, if a SDVS cannot find an instruction for handling an incoming packet, it sends a Request packet to the controller, which uses its global knowledge of underlying network elements to respond to the request.

Symmetric message type (Controller/SDVS-to-SDVS/Controller)

This message type is initiated by the controller or an SDVS and sent periodically without solicitation from the other.

Hello message: This message is for an SDVS to notify its existence and for the controller to inform the SDVS that it has not received an update for the current period.

Forwarding table specifications

The forwarding table contains instruction entries as rows of the table. This table is composed of three main elements: matching window, action window, and statistic window (as presented in Figure 3). The matching window is matched against an incoming packet. If a match is found, a corresponding action in the action window is executed, then associated statistics are updated for the matching packet. Otherwise, the packet is forwarded to the controller. The controller figures out how to process the packet.

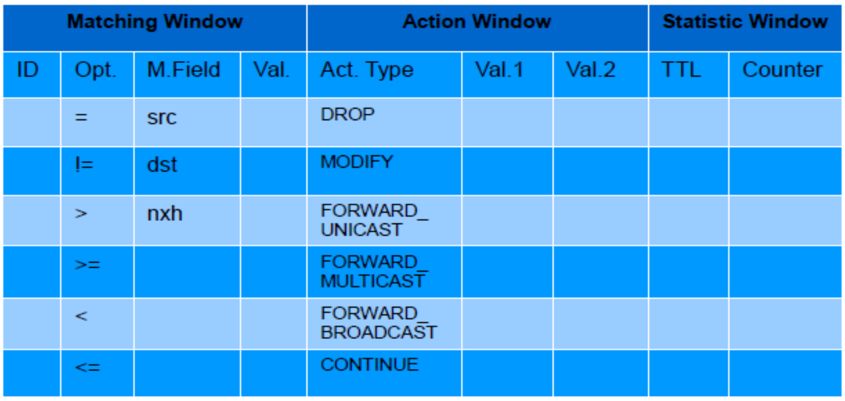

Matching Window				Action Window			Statistic Window	
ID	Opt.	M.Field	Val.	Act. Type	Val.1	Val.2	TTL	Counter
	=	src		DROP				
	!=	dst		MODIFY				
	>	nxh		FORWARD_UNICAST				
	>=			FORWARD_MULTICAST				
	<			FORWARD_BROADCAST				
	<=			CONTINUE				

Figure 3. Forwarding table structure

Matching window

It provides information for extracting needed values from an arriving packet header. The extracted values are matched against the specified values in the window to find a match for the incoming packet. The window is comprised of four parameters:

a) ID: Indicates an ID of a matching window of an instruction. It is used when an incoming packet needs to be matched with many matching fields since each forwarding entry allows matching of a field in a packet header. It enables multiple header fields of an incoming packet to be considered, while it does not require more memory for storing multiple matching windows for an instruction entry.

b) Matching Field: Indicates which part of packet header is compared to the specified value in the matching window, which means that not all packet header fields are necessarily matched against a forwarding entry.

c) Operator: Indicates a comparison method between the matching header field and the matching window Value. Operator values can be equal (=), different from (!=), higher than (>), higher than or equal to (>=), less than (<), less than or equal to (<=).

d) Value: Is compared to the extracted header field.

Action window

The window indicates a corresponding action for an instruction entry. The action window is composed of three parts: Action Type, value 1, and value 2. The value 1 and value 2 parts do not have a specific name, since they may represent values of different matching fields according to the action-type value.

a) Action Type: Indicates a type of action. Possible action types are FORWARD UNICAST, FORWARD MULTICAST, FORWARD BROADCAST, DROP, MODIFY, or CONTINUE.

b) Value 1: Different action types result in different meanings of Value 1. For example, the MODIFY action type requires a new value and the modified value. As for the CONTINUE action, the forwarding instruction ID needs to be specified, so the incoming packet needs to be matched against the instruction entry with the same ID. The FORWARD UNICAST, FOWARD MULTICAST, and FOWARD BROADCAST action types demand the unicast, multicast, and broadcast address, respectively.

c) Value 2: A replacement for the old value.

Statistic window

With a focus on efficiently programming of underlying IoT devices, their forwarding statistics would be necessary for an update of the network status. When a match is found, statistics related to the matched instruction are updated. The statistics are about Time To Live (TTL) and Counter.

a) TTL: Is a time to live of a forwarding instruction entry. It is decreased when the instruction table is updated. Its value depends on the required amount of time of an application request. It is gradually reduced to zero and is deleted when reaching zero.

b) Counter: Counts the number of packets matched against a forwarding entry.

Configuring table specifications

The configuring table provides an SDVS with instructions about configuration for its underlying IoT devices. Its structure is composed of two main windows: configuring and statistics (as presented in the Figure 4).

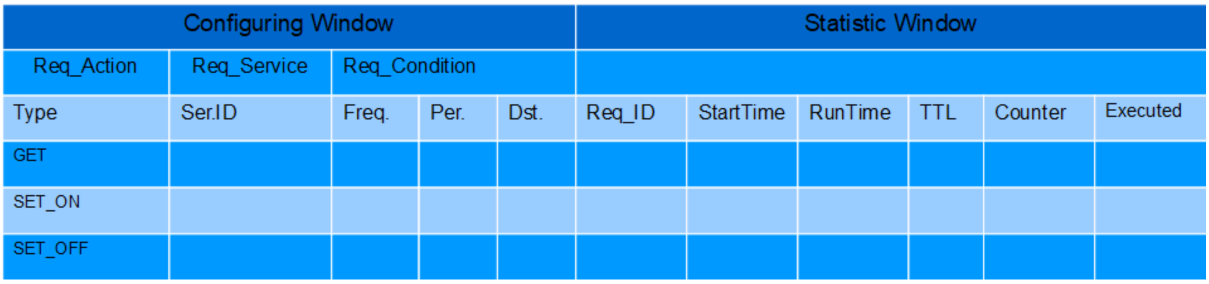

Configuring Window					Statistic Window					
Req_Action	Req_Service	Req_Condition								
Type	Ser.ID	Freq.	Per.	Dst.	Req_ID	StartTime	RunTime	TTL	Counter	Executed
GET										
SET_ON										
SET_OFF										

Figure 4. Configuring table structure

Configuring window

The configuring window includes three components: required services, required condition, and required action.

a) Required service: Indicates the required sensor service.

b) Required conditions: Indicates the conditions related to the required service.

a. Frequency: Specifies how often the required sensor service is achieved.
b. Period: Is an executing period of an instruction.
c. Destination Address: Specifies the destination of results returned by an SDVS. If there is no specified value, the destination is the controller.

c) Required action: Indicates an action type that is applied to the required service under the specified conditions.

Statistic window

The window shows the number of configuring instructions, and their operating time associated with an IoT application request. It includes information associated with an instruction, namely request ID, TTL, Counter, Operation time, and Executed status.

a) Request ID: Indicates which application request is associated with the configuring instruction. When the last configuring instruction of a ReqID is executed, the SDVS sends an acknowledgment to the controller about its completion of the required task.
b) TTL: is the existing time of a configuring instruction and is defined by application requirements. When it reaches zero, the related instruction is removed.
c) Counter: Shows the current number of requests for a sensor service and is used for updating a state of an SDVS. The state indicates a busy level of the SDVS. The higher the state number, the busier is the SDVS. The state is computed according to the total number of tasks that an SDVS performs and is updated in accordance with the counter statistics.
d) Operation time: Shows timing data related to an execution of an IoT application request. It provides information about the starting time and running time of an executed request. The information is essential for the orchestration function of the controller.
e) Executed: Specifies if an instruction has been executed or not. The executed status marked with “Y” means executed and with “N” means not executed.

Implementation of S-MANAGE in Provisioning IoT Sensor Services for IoT Applications

An implementation scenario

Our aim is provisioning IoT applications on demand by using the S-MANAGE protocol in the context of the SD-IoT model. Any request for IoT services is dynamically processed by the SD-IoT model. The system can orchestrate its underlying resources to handle multiple simultaneous sensor service demands (as shown in the Figure 5). According to its knowledge of the capability of the underlying resources, the system can i) obtain the availability of the resources and their current service-provisioning tasks; ii) provide appropriate responses to an application request, such as meet the request fully, or suggest an alternative that satisfies the request partially, or unable to provide the services because of insufficient resources; iii) handle simultaneous application requests and deal with conflicts among these requests; and iv) collect results corresponding to each application request.

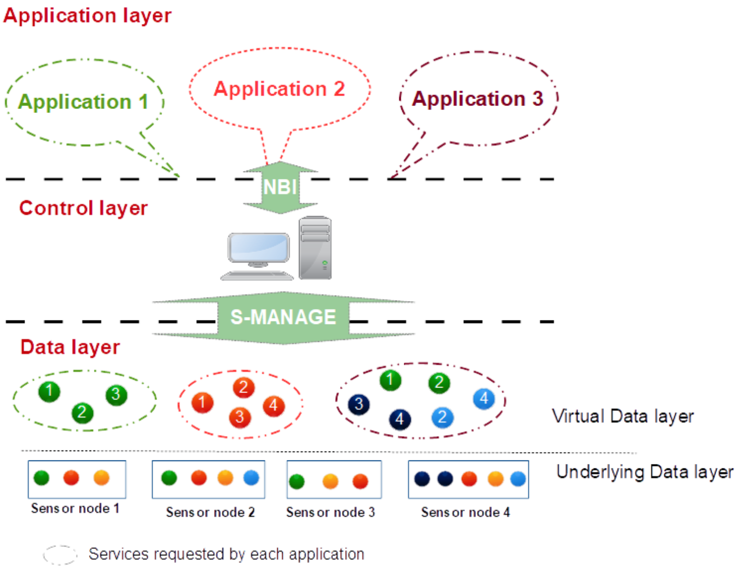

Figure 5. Implementation scenario

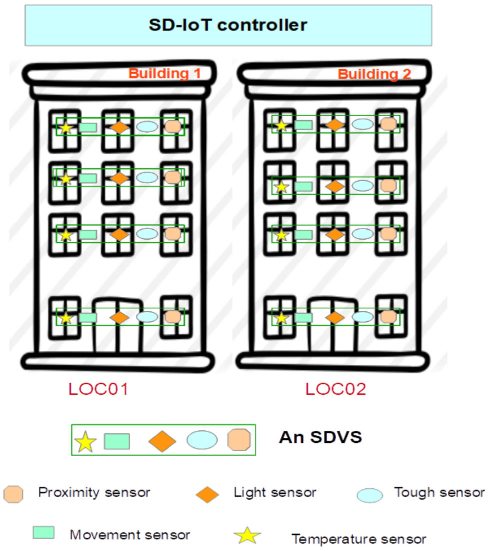

Figure 6. Implementation use case

For the sake of demonstration of a practical realization of the proposed protocol, we deploy the SD-IoT model that controls and manages two clusters of sensor nodes. The two resources are in two different locations. They can be orchestrated to provide sensor services for one or multiple IoT applications on demand. For the case study, the two clusters represent two buildings within a campus. Each building has four floors. Many types of sensors may be used on different floors, such as movement, temperature, proximity, touch, and light sensors (as presented in Figure 6).

A GUI interface is designed to enable users to indicate their sensor service types of interest and also to make specific demands for the required services. For example, they can indicate sensor services of interest, how long and how often they want to obtain the services. Moreover, they can choose the destination for the required services. An IoT application request is comprised of a set of these requirements.

Implementation setup

The implementation is developed from our preliminary implementation (Nguyen, Hoang & Dang, 2018). We also make use of Java dependencies support from the open source SDN-WISE platform (Milardo, 2017). To realise the operation of S-MANAGE, we build the SD-IoT model in which S-MANAGE provides a communication approach between the SD-IoT controller and its IoT elements. The SD-IoT model is a software platform written in Java and built in Netbeans 8.2. It is connected to a database built in MySQL. We have implemented three main software components of the proposed SD-IoT model: the SD-IoT controller, the S-MANAGE protocol, and the SDVS.

- The control module includes classes responsible for analysing application requests, orchestrating SDVS resources, generating instructions relating to the requests, networking and communicating with the SDVS.
- The Southbound interface module comprises classes for the construction of S-MANAGE messages, forwarding tables, and configuring tables of SDVSs.
- The virtual representation module is composed of classes for initiating instances of an SDVS.

We build a network where the controller communicates with its SDVSs. We establish a database in MySQL to store and update information regarding the SDVSs in the network, such as their sensor services, status, location, and attributes. The statistics from the forwarding and configuring tables are used to update the attributes, the status of the SDVSs and their underlying IoT devices. The database provides essential information for an operation of the controller’s core modules.

Implementation results

Implementation results demonstrate the expected features of the proposed S-MANAGE protocol in provisioning IoT applications on demand. S-MANAGE makes it possible for the controller to instruct IoT devices to achieve required services as well as forward results to required destinations. In addition, the protocol enables the controller to collect statistical information from the underlying IoT resources. Therefore, the controller can achieve the following results.

i) Programming its IoT resources via S-MANAGE according to an application request (Figure 7 and Figure 8).

ii) Responding dynamically to an application request about the service provisioning capability of the system according to its residual resources (as shown in Figure 9).

iii) Handling simultaneous application requests and conflicts over these requests (as demonstrated in Figure 10).

iv) Obtaining and displaying the status of multiple on-going application requests (as presented in Figure 11).

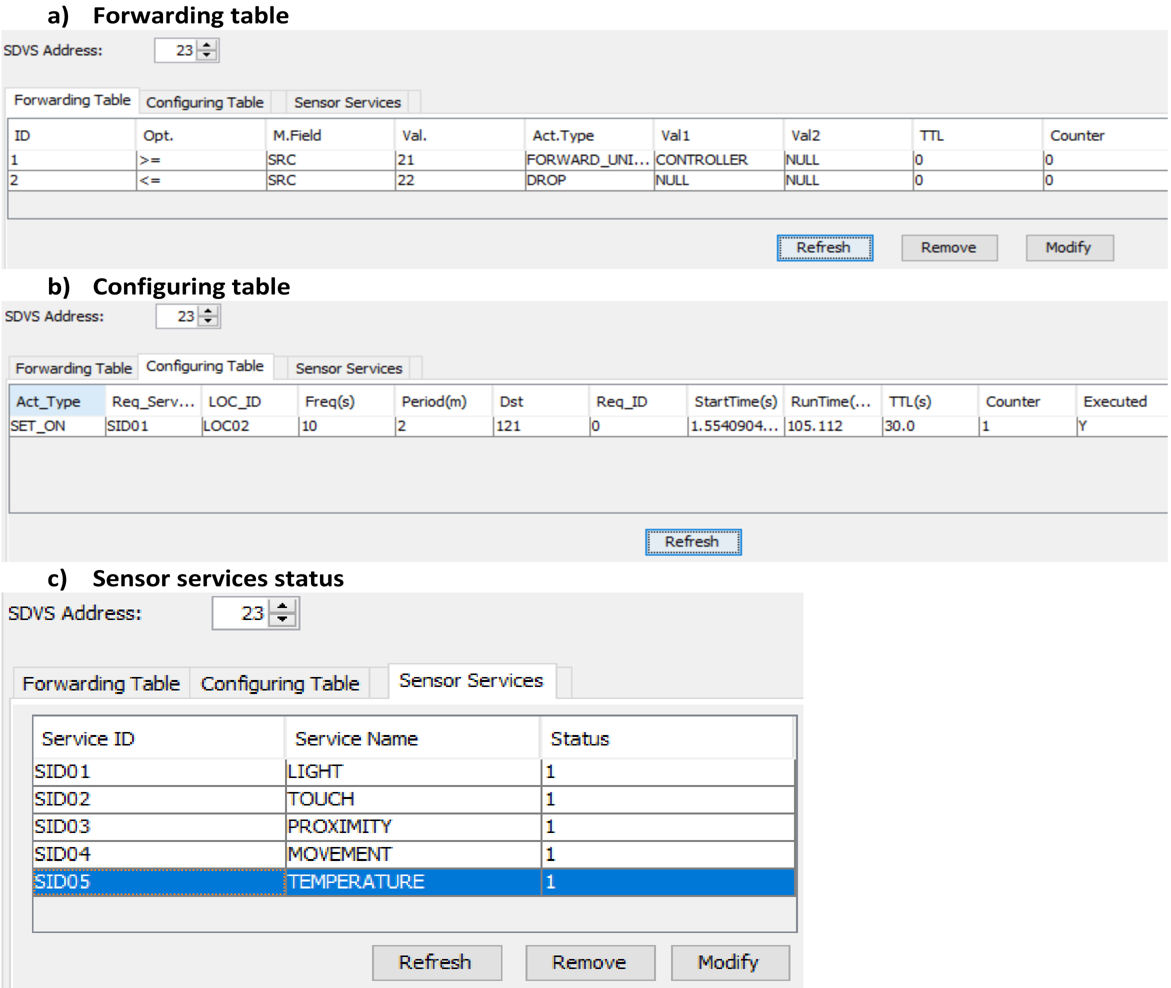

a) Forwarding table

SDVS Address: 23

ID	Opt.	M.Field	Val.	Act.Type	Val1	Val2	TTL	Counter
1	>=	SRC	21	FORWARD_UNI...	CONTROLLER	NULL	0	0
2	<=	SRC	22	DROP	NULL	NULL	0	0

b) Configuring table

SDVS Address: 23

Act_Type	Req_Serv...	LOC_ID	Freq(s)	Period(m)	Dst	Req_ID	StartTime(s)	RunTime(...	TTL(s)	Counter	Executed
SET_ON	SID01	LOC02	10	2	121	0	1.5540904...	105.112	30.0	1	Y

c) Sensor services status

SDVS Address: 23

Service ID	Service Name	Status
SID01	LIGHT	1
SID02	TOUCH	1
SID03	PROXIMITY	1
SID04	MOVEMENT	1
SID05	TEMPERATURE	1

Figure 7. Status of the SDVS before its configuration

The programmable function of S-MANAGE is demonstrated in Figure 7 and Figure 8. The two figures illustrate the status of an SDVS (SDVS03) before and after, respectively, it is programmed by the controller. In each figure, the status of the SDVS is presented, its forwarding instructions in (a), configuring instructions in (b), and sensor services status in (c). Differences between Figure 7 and Figure 8 are: i) both the forwarding and configuring tables of the SDVS03 are installed with one new instruction entry; and ii) changes in the status of the required service belonging to the SDVS. Via the installed configuring instruction, the SDVS can achieve the required services. Deploying the forwarding instruction, it knows how to forward results to the required destination. The result for the request is to change the status of the sensor service SID05 from 1 (ON) (as shown in Figure 7-c) to 0 (OFF) (as shown in Figure 8-c).

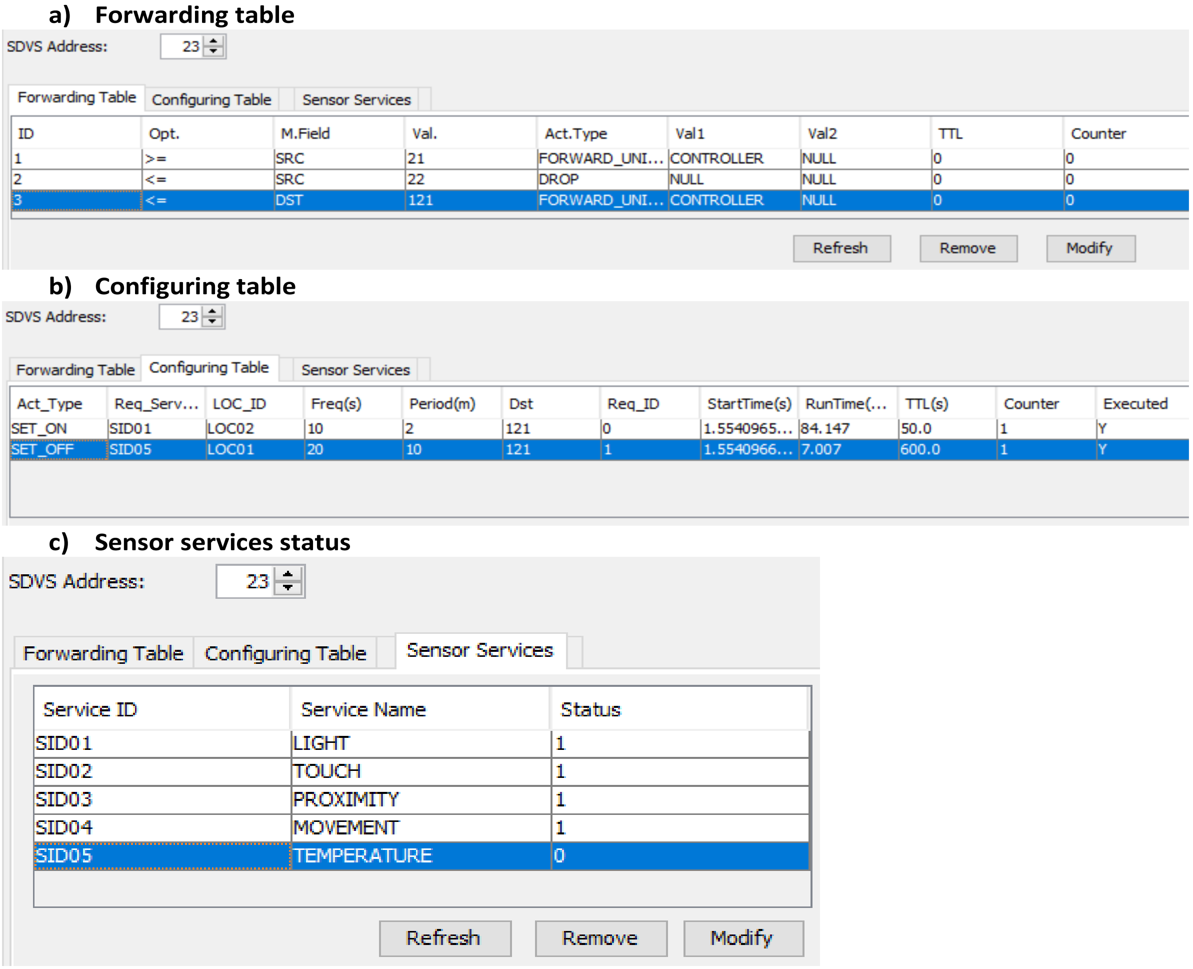

a) Forwarding table

SDVS Address: 23

ID	Opt.	M.Field	Val.	Act.Type	Val1	Val2	TTL	Counter
1	>=	SRC	21	FORWARD_UNI…	CONTROLLER	NULL	0	0
2	<=	SRC	22	DROP	NULL	NULL	0	0
3	<=	DST	121	FORWARD_UNI…	CONTROLLER	NULL	0	0

b) Configuring table

SDVS Address: 23

Act_Type	Req_Serv…	LOC_ID	Freq(s)	Period(m)	Dst	Req_ID	StartTime(s)	RunTime(…	TTL(s)	Counter	Executed
SET_ON	SID01	LOC02	10	2	121	0	1.5540965…	84.147	50.0	1	Y
SET_OFF	SID05	LOC01	20	10	121	1	1.5540966…	7.007	600.0	1	Y

c) Sensor services status

SDVS Address: 23

Service ID	Service Name	Status
SID01	LIGHT	1
SID02	TOUCH	1
SID03	PROXIMITY	1
SID04	MOVEMENT	1
SID05	TEMPERATURE	0

Figure 8. Status of the SDVS after its configuration

Moreover, thanks to the S-MANAGE protocol, the controller can muster the available IoT resources and orchestrate them to satisfy all the services whenever demanded. The S-MANAGE messages allow the controller to collect essential information about the updated status of the underlying IoT resources. If a request can be partially provisioned, the controller will also inform the application. Depending on the reply from the application, the controller performs its tasks based on the status table containing status of all SDVSs. The controller can program appropriate SDVSs to handle an incoming request according to its status (availability and capability). As shown in Figure 9, the controller provides appropriate responses to the application request in the case i) it can fully achieve all the required services (see Figure 9-a); ii) it partially achieves the required services and provides waiting time for obtaining the remaining required services (see Figure 9-b); or unable to provide the services because of insufficient resources (see Figure 9-c).

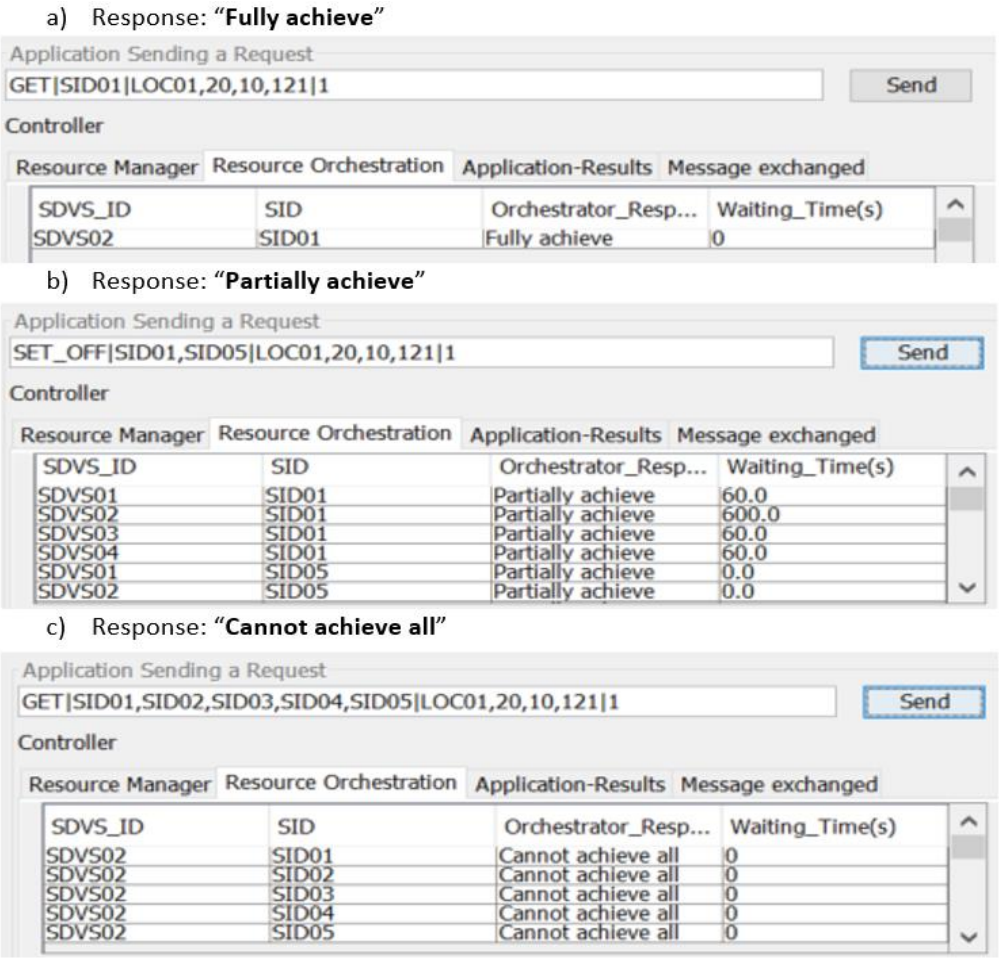

a) Response: "**Fully achieve**"

GET|SID01|LOC01,20,10,121|1

SDVS_ID	SID	Orchestrator_Resp...	Waiting_Time(s)
SDVS02	SID01	Fully achieve	0

b) Response: "**Partially achieve**"

SET_OFF|SID01,SID05|LOC01,20,10,121|1

SDVS_ID	SID	Orchestrator_Resp...	Waiting_Time(s)
SDVS01	SID01	Partially achieve	60.0
SDVS02	SID01	Partially achieve	600.0
SDVS03	SID01	Partially achieve	60.0
SDVS04	SID01	Partially achieve	60.0
SDVS01	SID05	Partially achieve	0.0
SDVS02	SID05	Partially achieve	0.0

c) Response: "**Cannot achieve all**"

GET|SID01,SID02,SID03,SID04,SID05|LOC01,20,10,121|1

SDVS_ID	SID	Orchestrator_Resp...	Waiting_Time(s)
SDVS02	SID01	Cannot achieve all	0
SDVS02	SID02	Cannot achieve all	0
SDVS02	SID03	Cannot achieve all	0
SDVS02	SID04	Cannot achieve all	0
SDVS02	SID05	Cannot achieve all	0

Figure 9. Dynamic response from the controller's resource orchestration to an application request

In addition, the system can handle multiple simultaneous application requests and resolve conflicts among these requests. As presented in Figure 10, in the control panel of the controller, the Resource Manager tab shows SDVSs' locations and their state. The Application-Results tab presents the current application requests and the status of the SD-IoT model's IoT application provision. Figure 10 illustrates three different states of the SDVS's functionality and corresponding tasks. In state 1, SDVS01 and SDVS02 are providing services SID01 and SID05 for the two application requests 1 and 2, respectively. Meanwhile, in state 2, SDVS02 receives another application request number 3 for the service SID05. The request cannot be processed immediately owing to the conflict between the two requests 2 and 3 for the same service. The request number 2 requires data from the sensor service, but the request number 3 requires deactivating the sensor service Therefore, SDVS02 delays the request number 3 until it completes the request number 1. In state 3, when the request number 1 is done, SDVS02 achieves the required service for the request number 3.

State 1: application requests and current-task status of involved SDVSs.

Current status of each SDVS

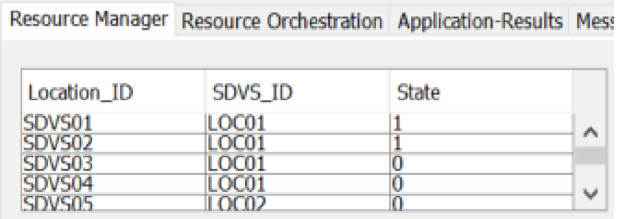

Location_ID	SDVS_ID	State
SDVS01	LOC01	1
SDVS02	LOC01	1
SDVS03	LOC01	0
SDVS04	LOC01	0
SDVS05	LOC02	0

Current application requests and related executed status

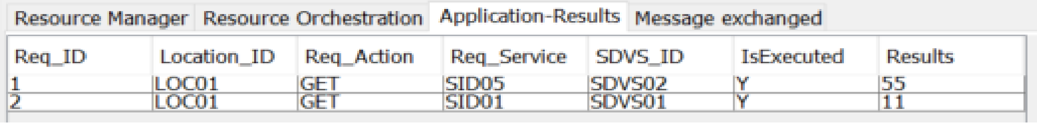

Req_ID	Location_ID	Req_Action	Req_Service	SDVS_ID	IsExecuted	Results
1	LOC01	GET	SID05	SDVS02	Y	55
2	LOC01	GET	SID01	SDVS01	Y	11

State 2: when there is an incoming request to turn off the required service SID05, all SDVSs in LOC01 have to be reconfigured. However, the SDVS01 is currently providing SID05 for another application: SDVS01 cannot start processing the incoming request for SID05.

Current status of each SDVS

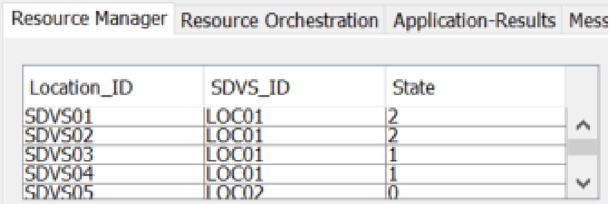

Location_ID	SDVS_ID	State
SDVS01	LOC01	2
SDVS02	LOC01	2
SDVS03	LOC01	1
SDVS04	LOC01	1
SDVS05	LOC02	0

Current application requests and related executed status

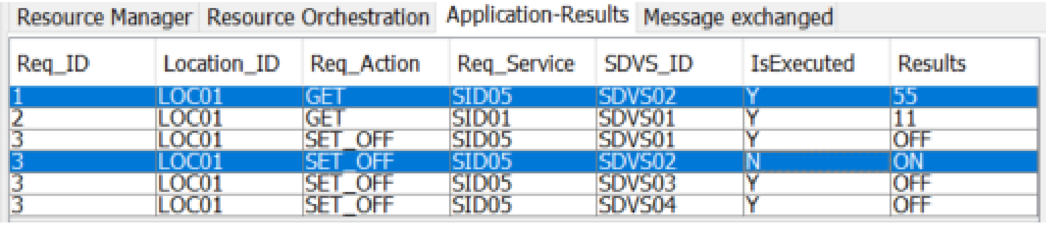

Req_ID	Location_ID	Req_Action	Req_Service	SDVS_ID	IsExecuted	Results
1	LOC01	GET	SID05	SDVS02	Y	55
2	LOC01	GET	SID01	SDVS01	Y	11
3	LOC01	SET_OFF	SID05	SDVS01	Y	OFF
3	LOC01	SET_OFF	SID05	SDVS02	N	ON
3	LOC01	SET_OFF	SID05	SDVS03	Y	OFF
3	LOC01	SET_OFF	SID05	SDVS04	Y	OFF

State 3: After releasing the task for request number 1, the SDVS02 processes the request number 3.

Current status of each SDVS

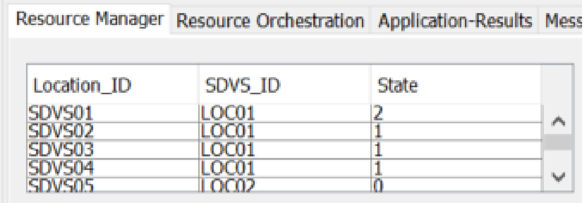

Location_ID	SDVS_ID	State
SDVS01	LOC01	2
SDVS02	LOC01	1
SDVS03	LOC01	1
SDVS04	LOC01	1
SDVS05	LOC02	0

Current application requests and related executed status

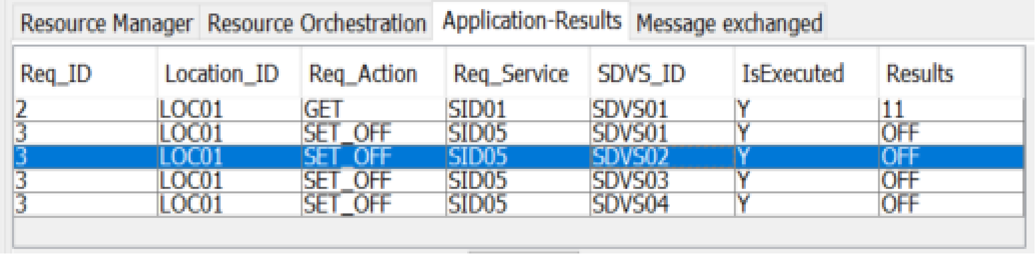

Req_ID	Location_ID	Req_Action	Req_Service	SDVS_ID	IsExecuted	Results
2	LOC01	GET	SID01	SDVS01	Y	11
3	LOC01	SET_OFF	SID05	SDVS01	Y	OFF
3	LOC01	SET_OFF	SID05	SDVS02	Y	OFF
3	LOC01	SET_OFF	SID05	SDVS03	Y	OFF
3	LOC01	SET_OFF	SID05	SDVS04	Y	OFF

Figure 10. Handling multiple application requests and solving conflicts among them

Figure 11 shows the status of all application requests and associated results. The Application-Results tab presents information about all application requests (represented by the Req_ID) and their execution status (see IsExecuted column: Y means Executed and N means Not-Executed). Moreover, the tab also displays required parameters regarding service type, location, related action, and associated results (see the Results column).

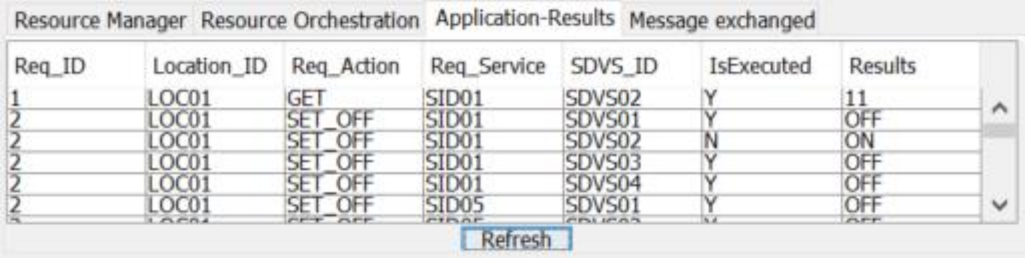

Req_ID	Location_ID	Req_Action	Req_Service	SDVS_ID	IsExecuted	Results
1	LOC01	GET	SID01	SDVS02	Y	11
2	LOC01	SET_OFF	SID01	SDVS01	Y	OFF
2	LOC01	SET_OFF	SID01	SDVS02	N	ON
2	LOC01	SET_OFF	SID01	SDVS03	Y	OFF
2	LOC01	SET_OFF	SID01	SDVS04	Y	OFF
2	LOC01	SET_OFF	SID05	SDVS01	Y	OFF

Figure 11. Status of on-going application requests and corresponding results

Conclusion

In this paper, we propose a design and implementation of the S-MANAGE protocol in order to address challenges of configuring and programming an IoT network and devices in provisioning IoT applications on demand. S-MANAGE is designed to configure functionalities of resource-constrained IoT devices through their virtual representations (SDVSs) and program their forwarding behaviours. Details of the design are provided. The implementation performance demonstrates the feasibility of the proposed protocol and its application. The proposal also enables further research and development on interoperability and orchestration of heterogeneous WSN/IoT devices for the provision of diverse IoT applications on demand.

References

Bera, S., Misra, S., Roy, S. K., & Obaidat, M. S. (2016). Soft-WSN: Software-Defined WSN Management System for IoT Applications. *IEEE Systems Journal, 12*(3), 2074-2081. doi:10.1109/JSYST.2016.2615761

Bera, S., Misra, S., & Vasilakos, A. V. (2017). Software-Defined Networking for Internet of Things: A Survey. *IEEE Internet of Things Journal, 4*(6), 1994-2008. doi:10.1109/JIOT.2017.2746186

Costanzo, S., Galluccio, L., Morabito, G., & Palazzo, S. (2012). Software Defined Wireless Networks: Unbridling SDNs. *2012 European Workshop on Software Defined Networking (EWSDN)*, October. doi:10.1109/EWSDN.2012.12

Deva Priya, I., & Silas, S. (2019). A Survey on Research Challenges and Applications in Empowering the SDN-Based Internet of Things. In: Peter, J., Alavi, A., & Javadi, B. (eds), *Advances in Big Data and Cloud Computing*, Advances in Intelligent Systems and Computing, vol. 750, Singapore: Springer. doi:10.1007/978-981-13-1882-5_39

Galluccio, L., Milardo, S., Morabito, G., & Palazzo, S. (2015). SDN-WISE: design, prototyping and experimentation of a stateful SDN solution for WIreless SEnsor networks. *2015 IEEE Conference on Computer Communications (INFOCOM)*, April-May. doi: 10.1109/INFOCOM.2015.7218418

Javed, F., Afzal, M. K., Sharif, M., & Kim, B. (2018). Internet of Things (IoT) Operating Systems Support, Networking Technologies, Applications, and Challenges: A Comparative Review. *IEEE Communications Surveys & Tutorials, 20*(3), 2062-2100. doi:10.1109/COMST.2018.2817685

Kobo, H. I., Abu-Mahfouz, A. M., & Hancke, G. P. (2017). A Survey on Software-Defined Wireless Sensor Networks: Challenges and Design Requirements. *IEEE Access, 5*, 1872-1899. doi:10.1109/ACCESS.2017.2666200

Li, S., Xu, L. D., & Zhao, S. (2018). 5G Internet of Things: A survey. *Journal of Industrial Information Integration, 10*, 1-9. doi:10.1016/j.jii.2018.01.005

Luo, T., Tan, H.-P., & Quek, T. Q. S. (2012). Sensor OpenFlow: Enabling Software-Defined Wireless Sensor Networks. *IEEE Communications Letters, 16*(11), 1896-1899. doi: 10.1109/LCOMM.2012.092812.121712

Mahmud, A., & Rahmani, R. (2011). Exploitation of OpenFlow in wireless sensor networks. *Proceedings of 2011 International Conference on Computer Science and Network Technology*, December. doi: 10.1109/ICCSNT.2011.6182029

Milardo, S. (2017). The stateful Software Defined Networking solution for the Internet of Things. Retrieved from https://github.com/sdnwiselab/sdn-wise-java

Nguyen, T. M. C., Hoang, D. B., & Chaczko, Z. (2016). Can SDN Technology Be Transported to Software-Defined WSN/IoT? *2016 IEEE International Conference on Internet of Things (iThings)* and *IEEE Green Computing and Communications (GreenCom)* and *IEEE Cyber, Physical and Social Computing (CPSCom)* and *IEEE Smart Data (SmartData)*, December. doi: 10.1109/iThings-GreenCom-CPSCom-SmartData.2016.63

Nguyen, T. M. C., Hoang, D. B., & Dang, T. D. (2017). Toward a programmable software-defined IoT architecture for sensor service provision on demand. *2017 27th International Telecommunication Networks and Applications Conference (ITNAC)*, Melbourne, November. doi: 10.1109/ATNAC.2017.8215419

Nguyen, T. M. C., Hoang, D. B., & Dang, T. D. (2018). A software-defined model for IoT clusters: Enabling applications on demand. *2018 International Conference on Information Networking (ICOIN)*, January. doi: 10.1109/ICOIN.2018.8343223

ONF [Open Networking Foundation]. (2012a). OpenFlow Switch Specification, Version 1.3.0 (Wire Protocol 0x04), ONF TS-006, 25 June. Retrieved from https://www.opennetworking.org/wp-content/uploads/2014/10/openflow-spec-v1.3.0.pdf

ONF [Open Networking Foundation]. (2012b). OpenFlow Management and Configuration Protocol (OF-Config 1.1), Version 1.1, ONF TS-005, 25 June. Retrieved from https://www.opennetworking.org/wp-content/uploads/2013/02/of-config-1.1.pdf

Perera, C., Liu, C. H., Jayawardena, S., & Min, C. (2014). A Survey on Internet of Things From Industrial Market Perspective. *IEEE Access, 2*, 1660-1679. doi:10.1109/ACCESS.2015.2389854

Razzaque, M. A., Milojevic-Jevric, M., Palade, A., & Clarke, S. (2016). Middleware for Internet of Things: A Survey. *IEEE Internet of Things Journal, 3*(1), 70-95. doi:10.1109/JIOT.2015.2498900

Sood, K., Yu, S., & Xiang, Y. (2016). Software-Defined Wireless Networking Opportunities and Challenges for Internet-of-Things: A Review. *IEEE Internet of Things Journal, 3*(4), 453-463. doi:10.1109/JIOT.2015.2480421

Sanctus: An Architecture for Trusted Products

Malcolm Shore
Canterbury University

Sherali Zeadally
University of Kentucky

Andy Clark
Royal Holloway College

Abstract: The last two decades have seen a fundamental shift in the manufacturing, sourcing and operation of technology, which has raised concerns in state security agencies about the cybersecurity risk to government and critical infrastructure. Sophisticated cyber attacks continue to be launched by state actors worldwide, while the engineering practices in common use have failed to deliver a commensurate improvement in technology cyber security. Cyber attacks continue to be successful against commercial networks, leading the US Government to encourage government agencies to look towards models such as zero-trust networking and tailored trustworthy spaces. There has been progress in product engineering, with formal methodologies such as Correctness by Construction (CbyC) successfully producing commercial products with increased trustworthiness. However, the adoption of these techniques has been limited, and governments are now increasingly resorting to an approach of technology Balkanization, where import and use of products and components may be restricted based on their country of origin. Even in the early stages of this strategy, the effect upon the economy is significantly adverse. We propose an alternative to technology Balkanization by combining trustworthy engineering approaches with the use of a national security component we call a sanctum which together can deliver sovereign trust.

Keywords: Cybersecurity, Balkanization, trust

Introduction

The evolution of technology over the last two decades has been rapid, particularly in the telecommunications field. Simple internet services became the World Wide Web and have evolved into sophisticated cloud technologies; the local network, which allowed terminal access to connected servers, has evolved into the internet of everything (Chandhok, 2014). The bricks and mortar of society are continuing to give way through digital transformation into the smart cities and businesses of the future (Matt, Benlian & Hess, 2015; Chanias &

Hess, 2016). In pace with these changes, nation states and criminals have found ways to subvert technology for their own benefit at substantial cost to the rest of the world (Anderson *et al.*, 2013). Out of simple viruses and worms designed to be mere annoyances have emerged vast botnets of compromised zombie computers capable of launching devastating attacks on unsuspecting targets anywhere in the world. Simple malware attachments have evolved into sophisticated techniques, such as those used by the Platinum Group to attack computers even when they are powered off (Mimoso, 2017). One consequence of this evolution of malicious attacks is that trust in technology has plummeted. Concerns over the vulnerability of internet-connected systems continue to form the mainstream driver for cybersecurity, with increasingly sophisticated attack technologies being used by criminals and nation states (Alcaraz & Zeadally, 2015).

The threat to ICT systems has not gone unnoticed, and information security practices and standards have evolved over the years. The original United Kingdom's Department of Industry Code of Practice PD0003 was adopted as the British Standards Institute's BS7799 and, subsequently, by the International Standards Organization into what is now known as ISO27000: Code of Practice for Information Security Management System (ISO/IEC, n.d.). The US National Institute of Standards and Technology issued Special Publication 800-53: Security and Privacy Controls for Federal Systems and Organizations and more recently published the NIST Cybersecurity Framework, which adopts both ISO27000 and SP800-53 controls into a single framework for cybersecurity. Other recommendations such as the UK Cyber Essentials have been proposed, but did not achieve global recognition.

Recent attacks on the Ukrainian power grid, attributed by the Ukrainian authorities to Russia, have highlighted the continuing risks to critical infrastructure (Park, Summers & Walstrom, 2017). While much of the risk can be attributed to configuration weaknesses, technology flaws and inadequate operational management, nation-state subversion of the supply chain has been highlighted by, and become something of a hobby horse for, security agencies in certain countries. As national critical infrastructures become dependent upon more advanced computing and telecommunications technologies, some of which is sourced from countries considered to be potential adversaries, so distrust of technology becomes an increasingly significant issue and governments' responses become a major risk to the global economy (National Journal, 2018).

Various governments have attempted to develop a robust approach to technology trust. In the 1960s, the US Department of Defense introduced a set of trusted systems criteria in what was known as the Orange Book (US DoD, 1983), offering trust at levels from C2 through to A1. (C2 is the accepted entry level of assurance; A1 is the highest level applied to classified systems). The UK Government introduced an alternative scheme called the IT Security

Evaluation Criteria (ITSec), which decoupled security functionality from its level of assurance. While ITSec was a significant step forward in systems assurance, improvements were needed (Gehrke, Pfitzmann & Rannenberg, 1992). Eventually, in the late 1990s, the Orange Book and ITSec approaches merged into a single set of criteria recognized by the US, UK, Canada, Australia and New Zealand (a grouping referred to as the Five Eyes). This scheme, known as the Common Criteria (n.d.), is now recognized by 28 countries as the means of approving equipment for use by governments and national infrastructure.

The US has for many years controlled the export of advanced technologies (Clark, 2015), both military and those which may be sold for peaceful purposes but have the ability for dual use in military systems. These controls have had mixed success. Many countries in the 1990s agreed to limit the spread of a key technology – namely, encryption. Encryption systems were included as a category of strategic arms, with export controls being applied to the more powerful cryptographic products. These controls proved to be counter-productive, encouraging many countries to develop their own products in competition with products from, and outside the control of, the US. Furthermore, with the posting and exchange of high-grade cryptographic techniques and tools on the internet, any control over cryptography is now ineffective.

The US and other countries have increasingly focused on technology support for the fighter, and the US, in particular, has evolved its military strategy based on having information dominance (Miller, 2019). This requires that the US has the most advanced technologies and information-enabling and disabling systems in the world and is able to access more sophisticated microelectronics than its adversaries (Chappell, 2017). As China becomes increasingly capable in developing advanced technologies, and comes to the threshold of potential technology dominance, it becomes more difficult for the US to maintain information superiority. This is particularly serious for the United States as it increasingly sources military technology components offshore (NDIA, 2017). While the United States has traditionally had significant control over technology used globally, its dependence upon foreign components for military products puts the United States at increasing risk of technology blockade or subversion should military action involve countries of supply.

The Race to Balkanize

The need for a global framework for technology trust was addressed in the late 1990s with the establishment of the Common Criteria scheme and is based on a framework of increasingly trustworthy levels of technology evaluation. A similar approach was more recently introduced in the Cybersecurity Law issued in China in 2017 (Ning & Wu, 2017). However, this approach now appears to be inadequate for the US and Australia, both of which have suggested they cannot achieve sufficient assurance to enable Chinese vendors to participate in their 5G networks.

Telecommunications has become the first major sector of national infrastructure in which serious attention has been given to supply chain security risks. An example of this is the Australian implementation of telecommunications sector security reforms, in which carriers are required to notify the Government of any substantive network changes and use of certain vendors' technologies may be limited or banned. The US has also taken similar measures with respect to sourcing from China and Russia, such as banning products from being used by US Government agencies (Volz, 2017). Surprisingly, this new development has not so far extended to the use of non-Chinese technology manufactured in China that suffers from exactly the same risk.

The US has also expanded the scope of its technology export control justified through the sanctions process. In 2017, the US banned the export of technology components to ZTE, leaving them unable to continue operations. Only a late reversal of this ban after payment of fines enabled ZTE to survive.

Taken together, the emerging approach of banning use of certain technologies and blocking exports is potentially the start of a slide towards what can be characterized as "strategic technology Balkanization", in which the technology used in a country will be limited to that manufactured within its geopolitical bloc.

There is a downside to any country completely or partially blocking certain advanced technologies:

- Balkanizing technology in a global manufacturing environment means repatriating much of the nation's offshore technology manufacturing. The Organization for Economic Co-operation and Development has warned nations against the associated strategy of localization, noting that this would jeopardize the benefits individual users and businesses enjoy from integrating global communications and the digital economy (OECD, 2016). More concerning, as explained by Apple to the US Government, the US has neither the facilities nor the indigenous skills to do the manufacturing (Worstall, 2013).

- In the event a class of technology is blocked, there may need to be some alternative source. This is the case in the US for semiconductors that was highlighted in the National Defense Industry Association reports (NDIA, 2017). Semiconductor foundries have a limited life, after which new foundries have to be built to support the more advanced chipsets. The cost to the US Government to ensure its semiconductor industry remains active is substantial, with an advanced foundry costing around $10B-$15B.

- Not using the most advanced technology in its infrastructure may result in a nation being restricted to what over time will become a "second-world" legacy infrastructure. This is particularly concerning for those countries with strategies based on digital transformation for which legacy technology cannot deliver the required products and services. Without enhancing that technology, these countries will be unable to compete globally, resulting directly and indirectly in an adverse economic impact (Qiang, Rossotto & Kimura, 2009). Mühleisen argues that restricting the use of technologies or legislating against them is not beneficial (Mühleisen, 2018). He recommends the development of smart policies that can optimize the benefits of new technology.
- In terms of a military strategy, and as was the case with encryption controls, the use of Balkanization may have entirely the opposite result to what is intended. The victims – adversaries amongst them – may choose to redouble their efforts to develop their own advanced technologies and become self-sufficient, which would merely compound the problem.

Regardless, the US and Australia have adopted a technology Balkanization strategy and have banned all 5G mobile network technology that is coming out of China. In doing so, they have accepted the cost to their economies both directly, from more expensive networks, and indirectly, with the potential for trade repercussions (Letts, 2019). This delays their digital transformation and results in it being delivered with more expensive and less advanced network technology. In Australia, the Government's decision to ban 5G from Chinese manufacturer Huawei has led to Optus delaying its 5G roll out and TPG cancelling plans to build a 5G network. Technology innovation in China will continue to accelerate, leaving Balkanized nations even further behind the rest of the world in delivering digital transformation.

China for its part is also pursuing a form of Balkanization. It is pursuing two key initiatives: Made in China 2025 and Internet Plus. Action by the US in blocking component exports has underlined the need for China to even more aggressively pursue independence in technology, and increases the awareness in other countries of the downside risk of using US technology.

Trust – an Alternative to Balkanization

Trusted computing

One alternative to technology Balkanization is to develop an approach to technology that can be trusted regardless of its source. Trusted technology concepts were introduced in the Trusted Computing Base (TCB) books published by the US DoD in the 1960s, the most popular of which was the Orange Book (US DoD, 1983), which covered trusted operating systems. However, the adoption of TCB to ensure trustworthy computing in the military fell victim to the economic imperatives of commercial-off-the-shelf (COTS) solutions and the few trusted systems that were developed have long since disappeared.

Despite the demise of TCB, the concepts of technology trust have continued to evolve. A basic foundation for technology trust is to design components to be secure, and to have a verification process to confirm their implementation is true to design. Secure design has attracted a significant amount of research, while verification has been addressed in industry with schemes such as Common Criteria evaluation, and more effectively through initiatives such as the Huawei Deep Evaluation Cell in the UK (Katwala, 2019). These approaches, however, offer only assurance at a point in time and do not address the issue of in-service trustworthiness.

Existing literature on trust

Marsh formalizes trust as a computational concept (Marsh, 1994) and notes that the formalism would not only enable network nodes to reason with and about trust, but would also provide network managers with another way to assess their networks – a remarkable insight into what is now a critical problem. Marsh defines basic, general and situational trust and includes such concepts as blind trust, optimistic trust that will never decrease, pessimistic trust that will never increase, and distrust where past actions influence current trust. He argues that the concepts of blind trust and permanent distrust should be discarded, as they do not belong in a rational decision-making system. Situational trust reflects the idea that different agents may calculate trust for the same entity differently, depending on their situation – and this may change as the situation changes. Marsh introduces the idea of *utility*, where an agent seeks to maximize the utility of a node for economic benefit, a far-sighted view of how trust needs to be balanced with economic gain. An interesting view from Marsh is that trust is not transitive, contrary to the views of later researchers such as Grandison & Sloman (2000). Marsh notes the problems that can occur in real-world trust: trust is a subjective phenomenon and humans use trust in a fashion clouded by emotions, wants, needs, and so forth; that there is a need to assess the rationality of agents making trust decisions and there is no *a priori* reason to assume agents are always rational. However, he does provide rules that a rational trusting entity, human or automated, should follow and provides a formal trust model in terms of calculating situational trust and co-operation thresholds.

In their survey of trust in Internet applications, Grandison & Sloman (2000) explore the properties of trust relationships and note that trust is never absolute but operates within limits. They cover the issue of infrastructure trust, i.e. the trust in the workstation being used, the local network and the network servers, by referring to the Orange Book. Their conclusion in surveying trust is that trust is the belief that an entity will act dependably, securely and reliably within a specified context; that trust can change over time; and that trust management enables information to be collected in order to make trust decisions. These concepts are as relevant today as when they were published in 2000.

Saadi *et al.* (2011) propose a trust meta-model to enable heterogeneous trust management systems to interoperate using mediators, allowing the development of composite trust models. Their model relates to technical aspects of stakeholder trust within different system models, but the meta-model can be widely applied to the more generic issues of trust. Their model consists of three elements: 1) trust roles, abstract representations of stakeholder behaviour; 2) trust relations between stakeholders in the model; and 3) trust assessment to compute the trustworthiness of stakeholders. The model includes direct and indirect trust relationships, with indirect trust reflecting the transitive trust concept referred to by Grandison and Sloman but also including the concept of reputation-based trust.

Networks have been a specific focus area for trust modelling, and in particular mobile ad-hoc networks. Jaydep Sen (2010) has proposed a framework for distributed trust management in mobile ad-hoc networks that approaches the problem from the perspective of key distribution and misbehaviour detection. Nodes in the network are considered to be cooperative, malicious or selfish and detection of uncooperative behaviour can be calculated using node-based reputation scores. Sen notes that external methods of preventing attack cannot be used when a node may be compromised, as it may be operating within a security envelope provided by network encryption. Instead, an internal method is proposed in which every node in the network monitors the behaviour of its neighbours and, if any abnormal action is detected, it invokes an algorithm to determine whether the suspected node is indeed malicious. The framework is designed to handle a range of behaviours such as dropping adverse feedback, selective broadcast and packet dropping, and tampering.

A particular focus for network trust has been Byzantine attacks against packet forwarding – a specific case of the more generic issue of supply-chain malware insertion. A Byzantine attack is one in which one or more nodes in a network may exhibit malicious behaviour. Zouridaki *et al.* (2007) propose a Hermes scheme and propose improvements to its robustness to Byzantine attacks (Zouridaki, Mark & Hejmo, 2007). The Hermes network scheme combines first-hand information on the behaviour of neighbour nodes, and second-hand reputational information passed from other nodes. Their improvements include a punishment policy to discourage selfish behaviour. The trust measurement assesses the number of correctly forwarded packets relative to the number of incorrectly forwarded packets and has an associated confidence factor. Hermes includes the concept of opinion, to generalize the idea of trustworthiness to non-neighbouring nodes. Han, Ravindran & Jensen (2007) propose a gossip-based mechanism for information exchange that is robust against Byzantine attacks for denying and faking messages – sometimes called a black hole attack. Their research indicates that, with relatively few rounds of gossip, the mechanism is robust in the presence of Byzantine attacks. Goyal & Sharma (2014) provided a short survey of Byzantine attacks in mesh networks, reporting the classification of attack types and identifying some key papers. Another survey was carried out in 2015 by Sindhuja, Nasrinbanu & Elavarasi (2015) addressing malicious node detection in data fusion sensor networks. Geetha & Sreenath (2016) also survey Byzantine attacks on the routing protocols used in mobile ad-hoc networks. They identify a number of Byzantine attacks including black hole, sinkhole, wormhole, gray, flood rushing, selfish and overlay network attacks. They identify a range of mitigations, including trust-based, incentive-based, cryptography-based, and analytical approaches. Eschenauer, Gligor & Baras (2002) consider Byzantine attacks in the context of mobile ad-hoc networks which may not have a fixed infrastructure available, and propose the use of swarm intelligence for trust distribution. Zhang (2017) presents the Byzantine defense problem from a contemporary cloud-based carrier network perspective, with a set of adversaries that range from the curious to state actors. He presents a Cybersecurity 3.0 model, which has to deal with intrusions within the Observe-Orient-Decide-Act (OODA) approach, and with a resilient multi-tiered control hierarchy of protection, detection, response and recovery. This approach mitigates attacks launched from the outside inwards that may subvert a node.

A significant component of technology trust is the implementation of the technology. In their Manifesto for High Integrity Software, Croxford & Chapman (2005) report that the user community and the software industry have been driven to accept that software defects are inevitable. This is less so in industries operating safety critical systems. These range from aircraft fly-by-wire control systems and railway signalling systems, to software in medical devices and traffic lights. In developing safety critical software, there are three key questions: what are the hazards presented to safety by the software; what can software engineers do to reduce hazards to an acceptable level; and how can the developed system be safety certified? There have over the years been various standards for safety critical systems, including the general standard IEC61508: Functional Safety of Electrical/Electronic/Programmable Electronic Safety-related Systems (2010) and the now withdrawn IEC60880: Software for Computers in the Safety Systems of Nuclear Powers Stations (1986). The CANDU Computer Systems Engineering Centre of Excellence of the Atomic Energy of Canada has published standard CE-1001-STD: Standard for Developing Safety Critical Software (CANDU, 1999).

Teikari & Nevalainen (2014) provide a good summary of safety critical software standards, and specifically review a number of IEC standards. They identify a number of deficiencies, one of which is that security is inadequately covered, and there is limited reference to coding standards. However, IEC 62645 (2014) does address the requirements for security, specifically the prevention of, detection of and reaction to malicious cyber acts that could lead to an unsafe situation. These include malicious modifications affecting system integrity, malicious interference with information, data or resources, and malicious changes to hardware, firmware or software at the programmable logic controllers.

A significant approach to delivery of contemporary trustworthy technology has been published by MITRE in their Cyber Resiliency Engineering Framework (Bodeau *et al.*, 2012). The framework provides a defence-in-depth approach to the development of architectural resilience practices to address the cyber threat.

In the area of semiconductors, the DARPA SPADE programme demonstrates the ability to disaggregate trust and enable the co-existence of multi-source commercial semiconductor capabilities (Chappell, 2017). Specifically, SPADE is designed to address the risk of malicious insertion through using secure parts to monitor commercial components packaged together into a single ASIC. Other strategies include authentication at any stage in the supply chain, reverse engineering to verify the design, and disaggregation into functional subcomponents.

More informal concepts around zero trust have been published by industry. One such example is the Palo Alto zero-trust approach to network security (Palo Alto, 2014). Zero Trust in this context is a data-centric network design that puts micro-perimeters around specific data or assets to allow more-granular rules to be enforced. Zero Trust networks solve the "flat network" problem that helps attackers move undetected inside corporate networks so they can find and exfiltrate sensitive data, and is often implemented using network segmentation.

In 2011, the United States National Cybersecurity Centre issued a strategic plan to develop a trustworthy cyberspace (NSTC, 2011). The objective of this plan was to mitigate strategic cyberspace vulnerabilities and ensure that the United States gains the most it can from the evolving use of cyberspace. The plan revolved around deep research into the root causes of the cyberspace problem, to develop scientific foundations and maximize the research impact, to induce change and to accelerate the transition of research into practice. The plan included priority areas of designed-in security, tailored trustworthy spaces, making the system a moving target for cyberattack, and cyber economic incentives. Noting that the absence of mechanisms to establish trust has made cyberspace vulnerable to illicit exploitations, the research theme for Tailored Trustworthy Spaces (TTS) aims to provide flexible, adaptive, distributed trust environments that can support functional and policy requirements arising from a wide spectrum of activities in the face of an evolving range of threats.

Beyond trust

Beyond trust, other efforts have been focusing on making networks immune to cyber attack. Cyber immunity is a relatively new research field of advanced anomaly detection, which incorporates automated response to attacks. In his research, Wlodarczak (2017) describes a cyber immune system as a detection/response/recovery technology that is inspired by the human biological immune system (and aligns well with three of the five categories of controls in the NIST Cybersecurity Framework). Traditional firewall and intrusion detection systems using signature schemes often struggle to detect zero day attacks, but a cyber immune system is designed to look for symptoms of the attack and provide a defense mechanism, which can contain and eradicate any form of attack exhibiting these symptoms. The key to cyber immunity is good detection through predefined genetic rules (innate immunity) or through learning (adaptive immunity), having a low-to-zero level of false positives (known in the biological sense as autoimmunity), and an effective response for anything that is detected. It is likely that a healthy system would contain many different cyber immunity modules, each focused on a specific class of attacks. Wlodarczak suggests that this would be achieved using artificial intelligence and neural network techniques combined with machine learning.

Tailored trustworthy spaces

The NSTC Strategic Plan (NSTC, 2011) defines a tailored trustworthy space as a technology domain that “provides flexible, adaptive, distributed trust environments that can support functional and policy requirements arising from a wide spectrum of activities in the face of an evolving range of threats”. In other words, it is a cyberspace environment that provides a user with confidence in its security, using automated mechanisms to adjust the level of security based on the user's context to address an evolving range of threats. It can also be an isolated collection of devices, services, policies and data that interact securely, reliably, and with privacy. For the purposes of this paper, we take a product-centric viewpoint of tailored trustworthy spaces, such that a trustworthy space is the security domain within the product that can be trusted by users to maintain information security and operational integrity, using automated mechanisms to deliver cyber immunity services throughout the product.

The Sherwood Applied Business Security Architecture, or SABSA for short (Sherwood, Clark & Lynas, 2005), can be used to describe trustworthy technology spaces in terms of its broader concept of security domains. A SABSA domain is a complete description of a business space, in which there are people, processes and technology, all of which must be trustworthy according to the policy of the domain. SABSA describes the modelling of isolated and independent (interacting) domains and their associated inter-domain associations. This methodology allows for security attributes to be described within the domain and on the interacting links. The key elements of a security domain are that: its boundary is explicit; it has a common security policy; it has a security domain authority responsible for setting and ensuring the effectiveness of that policy; it interacts with other domains through a domain gateway; and the domain is responsible for enforcing its own security policy at the gateway. Domains can exist within a higher level domain known as the super domain, and are then subordinate domains, more commonly called subdomains. Subdomains inherit policy from their super domains, and may interpret it within the context of their own domain. The concept of SABSA domains can be applied to the enterprise as a whole, but also to describe security domains internal to an ICT system. Consequently, SABSA is a useful tool for describing tailored trustworthy spaces in the concept of products.

Early work on designing tailored trustworthy spaces that address in-service attacks was carried out for the US Department of Energy (Speicher, 2011). The focus was securing smart grid control systems with their underlying IP infrastructure using a combination of services which together form what is termed the *security fabric* of the TTS architecture. An important aspect of the security fabric framework is the use of secure silicon in addition to standard firmware-based management services, an approach now common in mobile devices with their embedded Trusted Execution Environment. Secure silicon provides for sensitive storage and processing as well as the trusted monitoring required in trustworthy spaces. The DoE design introduced the concept of a service-oriented architecture and a policy-driven managing device, which handled device communications with agents in other subordinate devices.

Main research contributions of this work

We summarise the main contributions of this work as follows:

- We introduce a trustworthy technology framework based on the Canadian CE-1001-STD standard for safety critical systems and the MITRE Cyber Resiliency Engineering Framework, which addresses the issues of national security at each stage from requirements capture to advanced operational capability.
- We introduce the concept of a component that we call a *sanctus,* a sovereign component which provides the trusted tailored technology space to use in product design.

- We demonstrate the application of our model to the design of more resilient smart grid systems.

Proposed Model of Sovereign Technology Trust

Starting with safety critical systems

Safety critical systems focus on the integrity of software and, while they do not have a focus on nation security, they do provide an insight into the engineering techniques that are required to deliver technology trust. The CANDU standard CE-1001-STD is an IEC-aligned practical example of the application of safety critical concepts and provides a model of engineering in six stages. This is shown in Table 1 with the associated tasks and outputs for each stage.

Table 1. CE-1001-STD Safety Critical Software Engineering

Stage	Activity	Documentation
Concept	Business Analysis	Design Input Document
Requirements Definition	Requirements Definition	Software Requirements Specification
	Requirements Review	Requirements Review Report
Design	Design	Software Design Description
	Design Review	Design Review Report
Code Implementation	Coding	Source Code
	Code Review	Code Review Report
Testing	Unit Testing	Unit Testing Procedures Unit Test Report
	Integration Testing	Integration Test procedures Integration Test Report
	Validation Testing	Validation Test procedures Validation Test Report
Verification	Hazards Analysis	Hazards Analysis Report
	Reliability Qualification	Reliability Qualification Report

Reviews are carried out through the Requirements, Design, and Implement stages. Testing validates engineering of the design functionality. Hazards Analysis, which gets its input from the Design Input Documentation, Software Requirements Specification, Software Design Description and Source Code, is intended to identify any input conditions or subsystem failures that could lead to the software shifting into an unsafe state.

Requirements Definition includes the identification of failure modes, and establishing requirements for fault tolerance and graceful degradation. The Design requires the identification of self-checks to enhance robustness to hardware failures or other system level hazards. Coding is required to defend against detectable run time errors such as buffer overflows.

Reliability Qualification involves defining and explicitly identifying the basis for a reliability hypothesis, and then defining tests that simulate the software's usage profile in order to provide evidence that the probability of failure of the software is sufficiently small for it to meet its reliability requirements.

With this six-stage process, the software used in nuclear energy generation is assured to be safe to an acceptable level.

Trustworthy technology framework

In developing our trustworthy technology framework, we have adopted the practical standard CE-1001-STD, aligned it with the MITRE Cyber Resiliency Engineering Framework, and further developed it with new architectural, design, development, and operational concepts to enable its use to address national security trustworthiness. This revised framework is shown in Table 2.

Table 2. Enhanced Framework for Technology Trustworthiness

Stage	Activity	Output/Outcome
Concept	Business Analysis	Design Input Document
Requirements Definition	Quality Driven Requirements Capture and Analysis	Software Requirements Specification
	Requirements Review	Requirements Review Report
Design	Zero Trust Design	Software Design Description
	Design Review	Design Review Report
Development	Secure Software Engineering	Source Code
	Code Review	Code Review Report
Testing	Unit Testing	Unit Testing Procedures (including security) Unit Test Report
	Integration Testing	Integration Test procedures (including security) Integration Test Report
	Validation Testing	Validation Test procedures (including security) Validation Test Report
Verification	Hazards Analysis	Hazards Analysis Report
	Reliability Qualification	Reliability Qualification Report
	Common Criteria Evaluation	CC Certificate
	Deep Security Evaluation	National Security Endorsement
Operate	Known attack detection	Real time blocking/alerting
	Resiliency	Real time response to ensure continuous operation
	Anomalous attack detection using national security algorithms	System heartbeat monitoring Real time blocking/alerting
Resilience	Byzantine attack detection	Isolation of malicious components
	Defence-in-Depth	Ensure no single control point of failure
	Cyber immunity	Auto response to, and recovery from, attack
	Survivability	Shutdown of non-essential functions

http://doi.org/10.18080/jtde.v7n3.197

The second stage of Requirements Definition becomes Quality Driven Requirements Capture and Analysis to ensure that the specification of the technology requirements is correct and provides a solid foundation for delivering a product fit-for-purpose. This is a critical step to minimize the amount of rework required for the product to achieve full user acceptance. The MITRE Cyber Resiliency Engineering Framework provides a set of practices, which can be represented as SABSA attributes, to ensure that we can deliver security requirements down through design and implementation. These are:

- Adaptive response, taking actions to respond to an attack based on its characteristics;
- Analytic monitoring, gathering and analyzing data continuously;
- Co-ordinated Defence, managing multiple distinct mechanisms to respond to attack;
- Deception, actions to confuse or misdirect an attacker;
- Diversity, using different technologies to limit the spread of an attack;
- Dynamic positioning, to dynamically relocate elements of the system;
- Dynamic representation, to support situational awareness;
- Non-Persistence, to defeat known-location attacks;
- Privilege-Restriction, to make it difficult for an attacker to gain escalated privileges;
- Realignment, to reduce the attack surface;
- Redundancy, to avoid single points of failure;
- Segmentation, to control access to sensitive resources;
- Substantiated integrity, to ensure that critical elements of the system have not been corrupted; and
- Unpredictability, to make gaining a foothold difficult for an attacker.

We use the SABSA security domain concepts to achieve design of trustworthy spaces. The principles of Quality Function Deployment (QFD) can be applied within SABSA, starting with requirements capture and continuing through the framework, to create products and solutions that faithfully deliver the full spectrum of stakeholder requirements. QFD is one of the recognized foundations of design for trustworthy software (Jayaswal & Patton, 2006). Applying the SABSA practices is the first step in preparing to deliver trustworthy software by design.

Designing against a concept of zero trust enables security to be maintained even when the product is deployed into an environment in which there is no trust. Zero trust design requires that data is independently protected, and micro-segmentation is used to protect sensitive subsystems. Going further, segmentation should ensure that any sensitive data storage and processing is carried out in a trustworthy technology space. This then minimizes the scope of proof or evaluation for security assurance.

An important enhancement to CE-1001-STD is to adopt a practice of Secure Software Engineering rather than a more generic approach to coding to minimize the opportunity for flaws to be introduced during the development phase. By minimizing defects during development of the software, a great deal of post-testing recoding can be avoided, increasing reliability and reducing implementation costs. Using a rigorous design and implementation approach, such as Correctness by Construction (Kourie & Watson, 2012), for critical functions provides high productivity and low defects; and the associated use of SPARK/Ada for developing the code enables formal verification of correctness. This is made practical by designing the security critical functions to be in a tailored trustworthy space, which can then be targeted for rigorous engineering.

Our enhanced framework introduces the concept of independent testing. Common Criteria Evaluation validates security claims using the global government-recognized evaluation scheme, and we suggest that this continues to be an adequate approach for many environments. Where it is not, then Deep Security Evaluation goes further by incorporating source code review checking for both implementation weaknesses and the existence of malicious code. Taken together, these test regimes contribute the independent verification necessary to validate vendor claims of trustworthy technology and to confirm product integrity.

It is not sufficient to take trust at the point of launch as assuring the product for whole of life, as products are updated and environments change over time. In particular, products operating in hostile environments, or which can be remotely reached from a hostile environment, are susceptible to in-use compromise. Our enhanced framework therefore includes operational monitoring to maintain a level of trust through the operational use of the product. This involves known attack detection, typically using some form of signature matching, and anomaly detection to detect unusual and suspicious behaviour. The more advanced anomaly detection systems are able to learn what normal network behaviour looks like in order to more effectively detect anomalies. These capabilities are available for deploying as network solutions, but the techniques can also be applied within the product internal design.

Our enhanced framework addresses the need for superior resilience through introducing the MITRE Defence in Depth model, and three advanced activities from the research domain that provide through-life capability to detect issues harmful to the system and recover from them.

- Byzantine attack monitoring will address the issue of latent and stealthy malware introduced during manufacturing or later in the supply chain. While there has been significant research in this area, it has yet to appear in commercial products. A form of Byzantine attack detection could be applied in a trustworthy space to identify anomalous behaviour outside that space.

- From the MITRE Cyber Resiliency Engineering Framework, we adopt the four top level goals of Anticipate, Withstand, Recover, and Evolve to enhance the SABSA Defence in Depth model, which is supported by controls that apply the MITRE resiliency engineering objectives of Understand, Prepare, Prevent, Continue, Constrain, Reconstitute, Transform and Re-Architect. This can be seen in MITRE's diagram of the goals (top) and objectives (bottom) as shown in Figure 1.

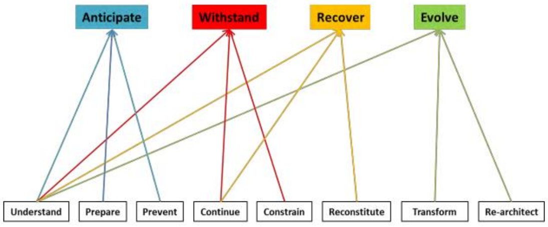

Figure 1. MITRE Cyber Resiliency Engineering Framework Goals and Objectives

- Signature-based quarantine techniques will evolve to more sophisticated advanced cyber immunity capability, which includes response mechanisms that affect not only shutdown of the attack but also "healing" of any damage done by the attack. Some progress in this area has occurred, with operating systems incorporating self-monitoring and service recovery – the first steps in the path to cyber immunity.

- Survivability, which requires the ability to fall back to a core set of critical activities in the event of overload or attack. This has not been common in industrial solutions, but can be seen in carrier mobile networks, where priority calls will be serviced even in congested networks by dropping non-priority calls. Applying the concepts of survivability to functions within a product or solution will improve the reliability of critical technology systems.

An important vector for critical infrastructure attack is the support and maintenance process, which is used to introduce changes to software in order to correct defects and provide new product features through code updates. Vendor product support is often provided remotely by a foreign national, and may require privileged access to the infrastructure. Attacks can occur by a malicious engineer uploading malware through legitimate access, or by a user-applied update being compromised. An example of the latter form of attack was the malicious code found in the CCleaner product (Collins & Hautala, 2017). The rigour applied during initial design and implementation needs to be applied for all subsequent changes to minimize the opportunity for this vector to be used, and operational monitoring will provide defence in depth.

A proposed model of sovereign technology trust

While Common Criteria evaluation was intended to deliver sufficient assurance to allow deployment of products into national infrastructures, the scheme has not been sufficient to satisfy sovereign security requirements. Despite Deep Security Evaluation facilities operating in UK, Germany and Canada, the US and Australian Governments (Varghese, 2019) continue to have insufficient trust in the world's leading technologies emerging from China.

A new model of technology trust is required to enable the deployment of globally sourced commercial products into national infrastructures. Based on our Cybersecurity Trust Framework, we propose an architectural approach to designing trustworthy technology that allows a sovereign trust module to be incorporated in a commercial product. This will then allow sovereign control of all security-related aspects of the product so that it can be deployed with assurance into the national infrastructure. We achieve this by isolating critical security information and functions to a specific domain or set of domains, which can be designed as the tailored trustworthy space within the product. The remainder of the product software can be untrusted. This is a logical extension of the design that we can see in contemporary mobile devices, where the Trusted Execution Environment allows critical security functions to be isolated in a trusted area inside the chip.

By designing the tailored trustworthy space as a discrete hardware component with secure, internationally standardized interfaces, a nation will be able to provide sovereign national security algorithms and secure key storage in a trustworthy module, which we call a *sanctus*. The sanctus could be a plug-in module of some form, either externally by the user or internally during national localisation of the product.

The sanctus as a component of the larger product would also need to be developed using our cybersecurity trust framework. It would need as a whole to be rigorously engineered in order to achieve a sovereign level of trust that can then be extended into any product in which it is used. As the secure heart of a product, it would need to have as small an attack surface as possible and ideally be formally proved.

In operational use, there would be three significant outcomes for the sanctus:

- **Secure the Information Flow**. A trustworthy product needs to be able to ensure the confidentiality and the integrity of information passing through it. The sanctus will need to have the ability to take control of interfaces so that any incoming information can be protected prior to being passed into the untrusted domain within the product. This is similar to the way in which a mobile trusted execution environment can take control of the keyboard interface for PIN entry.
- **Ensure Integrity**. Integrity is another key issue for effective security, and this means knowing exactly what software is running in the product. We can do this by demonstrating binary equivalence, meaning that the operational software matches a validated software release. By using a sanctus, securely loaded software signatures can be checked at start-up and during operation against the running code in a product. This ensures that the software has not been tampered with, and that only validated versions of software can run.
- **Monitor System Health**. For critical infrastructure, availability is often as important as confidentiality and integrity. The sanctus can be used to run real time checks of the operational state of the product and report back to a health monitoring system using a secure heartbeat mechanism. This will include the basic and advanced cybersecurity monitoring of information flows within the product and network flows touching the product and reporting any alerts via the heartbeat. These are likely to include nationally sensitive algorithms, and as such will be loaded into the sanctus rather than in the product itself. By doing this through the sanctus, an attacker cannot send spoofed reporting of health while the system is under attack or disabled.

Cybersecurity monitoring is complex. Contemporary cyber security products use one or all of signatures, learning schemes, and algorithms to deliver effective security. New cyber attacks are emerging all the time, and traditional anti-virus solutions require regular signature updates. Anomaly detection systems such as DarkTrace (2019) incorporate mechanisms to learn what normal network activity looks like, and to detect any deviations from the learned algorithms. Solutions such as Microsoft's CloudApp (2019) allows anomaly detection policies to be incorporated as algorithms. The sanctus will need to be able to support all these capabilities, as well as any advanced sovereign resilience features that have been developed.

Smart Metering Case Study

Smart grid security challenges

In this section, we apply our proposed architectural approach to the problem of securing distribution in a smart grid to demonstrate that it effectively addresses the known security challenges.

The major work on smart grid security emerged in 2010, in which smart grids are shown in a conceptual model covering generation, support systems, transmission, and distribution controlled via an operational centre, and being managed through a market approach in which service providers address the needs of customers. The business requirements for a power grid include reliability and resilience, self-healing against disruption events, and that it provides safe and efficient energy delivery (Amin, 2011). The customer impact of even a limited power loss event can be catastrophic. For example, in 2007, Mercury Energy deliberately disconnected power to a house in which a woman was on life support, resulting in her death (Henderson, 2007). There are also privacy concerns that must be addressed in smart grids – data sent regarding power usage can indicate the absence of a householder and be used to target houses for burglaries – and the data must be handled accordingly (Zeadally *et al.*, 2013).

In time of conflict or political tension, the smart grid could be an early target of critical infrastructure attack (Anderson & Fuloria, 2011). In past conflicts power and communications utilities have been targeted through air attacks or sabotage. The use of smart grids substantially reduces the cost and risk to the attacker by enabling the attack to be conducted through remote computer exploit. A significant amount of work has been carried out into identifying and mitigating smart grid threats (Otuoze, Mustafa & Larik, 2018); however, Alcaraz & Zeadally (2015) note that utilities typically have little experience of defending themselves against capable motivated cyber adversaries. The BlackEnergy cyber attack in 2015 (Lipovsky & Cherepanov, 2016) provided ample demonstration of the ability of threat actors to execute a denial of service and achieve social disruption.

Skopik *et al.* (2012) provide an insight into specific smart grid threats and vulnerabilities. They report that a smart grid system involves three tiers: the uplink from the smart meter; the backhaul to the application; and the smart grid application itself. Tier 1 attacks include local hardware and firmware manipulation and exploitation, potentially remotely, of design and implementation. Tier 2 attacks include network sniffing from the home or neighbourhood network, large scale meter takeover via malware spreading peer to peer, and backhaul concentrator node attacks. Tier 3 attacks are web attacks focusing on consumer and management services. The most likely attack from a national security perspective would come via a remote attack on the central operations management system, as the network should be fully protected against end-device penetration even should an attack on a home or neighbourhood network device be successful. However, meter misreporting of overloads could result in partial network shutdown. From a power operator perspective, the impact of these attacks falls into the category of power theft by manipulating recording or changing usage data, data theft, unit Denial of Service (DoS) by causing a meter to malfunction, or grid DoS by interfering with concentrators or mass meter compromise.

Alcaraz & Zeadally (2015) provide a further perspective on threats related to SCADA systems, including man-in-the-middle attacks, which can inject or modify control messages. These kinds of attacks can have a broad impact across the system, and are particularly significant for SCADA protocols, such as Modbus, that operate in plain text. The number of attacks on industrial control systems is increasing, and the key one for national security is denial of service. Their research indicates a number of technical approaches to achieving a Denial of Service, such as jamming of mobile components, flooding attacks, selective forwarding attacks, impersonation attacks, dropping and redirecting messages.

The attack vectors for each tier of the smart grid solution can be summarized as shown in Table 3, with the relevant adverse outcomes marked with an asterisk for the various types of attack.

Table 3. Network Attack Vectors

Tier	Attack	Power Theft	Data Theft	Unit DoS	Grid DoS
1. Water Meter	Hardware manipulation	*		*	
	Firmware manipulation	*	*	*	
	Impersonation/meter emulation	*			
	Exploitation of design weaknesses	*	*	*	
2. Utility	Network sniffing		*		
	Large scale meter takeover				*
	Message injection/modification	*	*		*
	Message redirect/drop	*			*
	Message Flooding				*
	Attacks on concentrator nodes		*		*
3. Web and Backend Applications	System penetration				*
	Theft of metering data	*			

To combat these threats, Otuoze, Mustafa & Larik (2018) have identified the key conceptual attributes relevant to ensure the smart grid is protected: authentic, available, reliable, confidential, integrity-assured, efficient, accessible, authenticated, robust, flexible, and resilient. To this we can add private, as identified by Zeadally *et al.* (2013), and the system itself needs to be trustworthy. It also needs to be as immune as possible to attacks, i.e. to have the capability to detect and trigger an immunity response. This provides the foundation for a quality driven design using a sanctus to enable zero trust throughout the smart grid.

Tier 1

Hardware manipulation of the meter is in general outside the scope of a sanctus and would need to be achieved using traditional tamper-proofing. However, designing the meter to incorporate a sanctus could address the remaining two methods of meter attack. Firmware manipulation could be detected using a sanctus through monitoring of the software integrity signature. Impersonation can be addressed as a Byzantine attack by using the sanctus to validate peer device signatures. Exploitation of design weakness would fall into either hardware or firmware exploitation.

Tier 2

The use of a negotiated encryption protocol using trusted encryption code in the sanctus would enable nationally-trusted encrypted communications to avoid data breach from network sniffing attacks. Large-scale meter takeover from viral malware moving laterally across the grid can be mitigated using monitoring modules in the sanctus. Message injection and dropping can be mitigated using Byzantine detection in the sanctus, and message modification through a man in the middle can be addressed by encrypting messages. Message flooding can be mitigated using an immunity mechanism in the sanctus that detects and reacts to a node flooding the grid by shutting it down. Concentrator nodes can be designed in a trustworthy manner in order to be resilient to attack.

Tier 3

Protecting the central smart-grid operations system and its back-end databases from attack is the most complex part of protecting the smart grid. The smart-grid software would be developed with security code and sensitive data in the sanctus, and using rigorous security engineering to ensure the code is trustworthy. Operation of the system would be designed to authenticate any critical command with a sanctus signature to ensure only sanctus sourced actions are taken. Operational monitoring would include the sanctus running the full range of configuration integrity checks across the server, automated account and access checks, and full anti-malware monitoring. Advanced resilience features might include a survival mode, which locks down the smart grid to a fixed safe state in the event of attack. Advanced security monitoring in the sanctus could prevent unauthorised access to the system to steal metering data. By having the security code and sensitive data protected in the sanctus, any compromise of the server would have limited impact.

Conclusion

In this paper we suggest that the move to technology Balkanization is a less than ideal solution to the problem of sovereignty in a global technology environment. We provide a survey of the literature on technology trust.

From this, we propose an alternative to technology Balkanization based on the Canadian CE-1001-STD safety-critical engineering process. This process is further developed to enable its application to be extended to address national security, using an architecture that incorporates a standardised sovereign component, called a sanctus, to ensure confidentiality, integrity and availability of sensitive data, algorithms, and processing. Using a smart grid case study, we have shown how this can address the known attacks in smart grids.

Further research is required to develop and trial a sanctus, and, in particular, to develop a robust interface standard that could be presented for international adoption.

References

Alcaraz, C., & Zeadally, S. (2015). Critical infrastructure protection: Requirements and challenges for the 21st century, *International Journal of Critical Infrastructure Protection*, *8*, 53-66, January. DOI: 10.1016/j.ijcip.2014.12.002

Amin, S. M. (2011). Smart Grid: Overview, Issues and Opportunities. Advances and Challenges in Sensing, Modeling, Simulation, Optimization and Control, *European Journal of Control*, *5*(6), 547-567. DOI: 10.3166/EJC.17.547–567.

Anderson, R., & Fuloria, S. (2011). *Smart meter security: a survey*. Available at: https://www.cl.cam.ac.uk/~rja14/Papers/JSAC-draft.pdf. Accessed 26 March 2019.

Anderson, R., Barton, C., Böhme, R., Clayton, R., van Eeten, M. J. G., Levi, M., Moore, T., & Savage, S. (2013). Measuring the Cost of Cybercrime. In: Böhme, R. (ed), *The Economics of Information Security and Privacy*, Berlin: Springer. Available at: https://www.researchgate.net/publication/263605690_Measuring_the_Cost_of_Cybercrime. Accessed 19 March 2019. DOI: 10.1007/978-3-642-39498-0_12

Bodeau, D. J., Graubart, R. D., Picciotto, J., & McQuaid, R. (2012). Cyber Resiliency Engineering Framework, *MITRE Technical Papers*. Available at: https://www.mitre.org/publications/technical-papers/cyber-resiliency-engineering-framework. Accessed 26 July 2019.

CANDU. (1999). Standard CE-1001-STD, Revision 2, Standard for Software Engineering of Safety Critical Software, CANDU Computer Systems Engineering Centre of Excellence, December. Available at: http://fm.csl.sri.com/VeriSure2015/talks/CE-1001-STD.pdf. Accessed 23 March 2019.

Chandhok, R. (2014). The Internet of Everything, *2014 IEEE Hot Chips 26 Symposium*, 1-29. DOI: 10.1109/HOTCHIPS.2014.7478826

Chanias, S., & Hess, T. (2016). Understanding Digital Transformation Strategy Formation: Insights from Europe's Automotive Industry, *20th Pacific Asia Conference on Information Systems*, Chiayi, Taiwan, June 2016.

Chappell, W. (2017). A Technology-Enabled New Trust Approach, DARPA presentation. Available at: https://www.darpa.mil/attachments/1TheCaseforSecureASICs_Slides.pdf. Accessed 26 March 2019.

Clark, D. (2015). U.S. Agencies Block Technology Exports for Supercomputer in China, *The Wall Street Journal*, 9 April. Available at: https://www.wsj.com/articles/u-s-agencies-block-technology-exports-for-supercomputer-in-china-1428561987. Accessed 23 March 2019.

Collins, K., & Hautala, L. (2017). Hackers hid malicious code in popular CCleaner software, *CNET*, 19 September. Available at: https://www.cnet.com/news/hackers-hid-malicious-code-in-popular-ccleaner-software/. Accessed 26 March 2019.

Common Criteria. (no date). Common Criteria Web Portal. Available at: https://www.commoncriteriaportal.org/. Accessed 27 August 2019.

Croxford, M., & Chapman, R. (2005). Correctness by Construction: A Manifesto for High-Integrity Software, *Crosstalk: Journal of Defense Software Engineering*, December. Available at: https://pdfs.semanticscholar.org/3516/e1cdf840fe0e697aa43dd6be9ab9de71120a.pdf. Accessed 27 August 2019.

DarkTrace. (2019). The Enterprise Immune System. Available at: https://www.darktrace.com/en/products/enterprise/. Accessed 27 August 2019

Eschenauer, L., Gligor, V., & Baras, J. (2002). On Trust Establishment in Mobile Ad-Hoc Networks, *Lecture Notes in Computer Science 2845*, September. DOI: 10.1007/978-3-540-39871-4_6

Geetha, A., & Sreenath, N. (2016). Byzantine Attacks and its Security measures in Mobile Adhoc Networks, *International Journal of Computing, Communications & Instrumentation Engineering*, *3*(1), 42-47. DOI: 10.15242/IJCCIE.AE0116013

Gehrke, M., Pfitzmann, A., & Rannenberg, K. (1992). Information Technology Security Evaluation Criteria (ITSEC) – a Contribution to Vulnerability?, *Proceedings of the 12th IFIP World Computer Congress on Education and Society* - Information Processing 92, *II*. Available at: https://pdfs.semanticscholar.org/facd/bd4b410670431e3f0ec2cf3dabcc7ef55545.pdf. Accessed 27 August 2019.

Goyal, S., & Sharma, V. (2014). Byzantine Attack on Wireless Mesh Networks: a Survey, *International Journal of Science, Engineering and Technology Research*, *3*(12), 3260-3264, December. Available at: http://ijsetr.org/wp-content/uploads/2014/12/IJSETR-VOL-3-ISSUE-12-3260-3264.pdf. Accessed 27 August 2019.

Grandison, T., & Sloman, M. (2000). A Survey of Trust in Internet Applications, *IEEE Communications Surveys & Tutorials*, *3*(4), 2-16, January. DOI: 10.1109/COMST.2000.5340804

Han, K., Ravindran, B., & Jensen, E. D. (2007). Byzantine-Tolerant Point-to-Point Information Propagation in Untrustworthy and Unreliable Networks, September 2007. Available at: https://www.researchgate.net/publication/220909574_Byzantine-Tolerant_Information_Propagation_in_Untrustworthy_and_Unreliable_Networks. Accessed 26 August 2019.

Henderson, H. (2007). Mercury introduces better systems after Muliaga death, *New Zealand Herald*, 2 July. Available at: https://www.nzherald.co.nz/nz/news/article.cfm?c_id=1&objectid=10447834. Accessed 27 August 2019.

IEC. (1986). IEC60880: Software for computers in the safety systems of nuclear power stations, International Electrotechnical Commission. Available at: https://webstore.iec.ch/publication/18251. Accessed 27 August 2019.

IEC. (2010). IEC61508: Functional safety of electrical/electronic/programmable electronic safety-related systems, International Electrotechnical Commission. Available at: https://www.iec.ch/functionalsafety/. Accessed 27 August 2019.

IEC. (2014). IEC62645: Nuclear power plants - Instrumentation and control systems - Requirements for security programmes for computer-based systems, International Electrotechnical Commission. Available at: https://webstore.iec.ch/publication/7311. Accessed 27 August 2019.

ISO/IEC [International Organization for Standardization/International Electrotechnical Commission]. (no date). ISO 27000 Series of Standards. Available at: https://www.itgovernance.co.uk/iso27000-family. Accessed 26 August 2019.

Jayaswal, B. K., & Patton, P. C. (2006). *Design for Trustworthy Software: Tools, Techniques, and Methodology of Developing Robust Software*. Upper Saddle River, NJ: Prentice Hall. ISBN 978-0131872509.

Katwala, A. (2019). Here's how GCHQ scours Huawei hardware for malicious code, *Wired*, 22 February. Available at: https://www.wired.co.uk/article/huawei-gchq-security-evaluation-uk. Accessed 26 August 2019.

Kourie, D. G., & Watson, B. W. (2012). *Correctness-By-Construction Approach to Programming*. Springer. ISBN 9783642279195.

Letts, S. (2019). China policy on Australian coal is 'as dark and impenetrable as night' and that's how it wants it, *ABC News*, 25 February. Available at: https://www.abc.net.au/news/2019-02-25/china-policy-on-australian-coal-dark-and-impenetrable/10843148. Accessed 23 March 2019.

Lipovsky, R., & Cherepanov, A. (2016). BlackEnergy trojan strikes again: Attacks Ukrainian electric power industry, *welivesecurity*. Available at: (https://www.welivesecurity.com/2016/01/04/blackenergy-trojan-strikes-again-attacks-ukrainian-electric-power-industry. Accessed 26 March 2019.

Marsh, S. P. (1994). *Formalising Trust as a Computational Concept*, PhD Thesis, University of Stirling. Available at: https://www.nr.no/~abie/Papers/TR133.pdf. Accessed 23 March 2019.

Matt, C., Benlian, A., & Hess, T. (2015). Digital Transformation Strategies, *Business & Information Systems Engineering*, *57*(5), 339-343, October. DOI: 10.1007/s12599-015-0401-5

Microsoft. (2019). Get instantaneous behavioral analytics and anomaly detection, 2 April. Available at: https://docs.microsoft.com/en-us/cloud-app-security/anomaly-detection-policy. Accessed 27 August 2019.

Miller, D. (2019). Information Dominance: The Philosophy, *GPF: Global Policy Forum*, 29 December. Available at: https://www.globalpolicy.org/component/content/article/154/26581.html. Accessed 23 March 2019.

Mimoso, M. (2017). Platinum APT First to Abuse Intel Chip Management Feature, *Threatpost News Wrap*, 9 June. Available at: https://threatpost.com/platinum-apt-first-to-abuse-intel-chip-management-feature/126166/. Accessed 23 March 2019.

Mühleisen, M. (2018). The Long and Short of The Digital Revolution, *Finance & Development*, *55*(2), 4-8, June.

National Journal. (2018). The Balkanization of Global Tech, *National Journal*, 30 April. Available at: https://www.nationaljournal.com/s/667253/balkanization-global-tech. Accessed 23 March 2019.

NDIA [National Defense Industrial Association]. (2017). Team 2 Summary: Trustable Access to Leading Edge Technology, *NDIA Trusted Microelectronics Joint Working Group*, July. Available at: https://www.ndia.org/divisions/%20working-groups/tmejwg/final-team-reports. Accessed 23 March 2019.

Ning, S., & Wu, H. (2017). China: Cybersecurity 2017. Available at: https://iclg.com/practice-areas/cybersecurity-laws-and-regulations/china. Accessed 27 August 2019.

NSTC [National Science and Technology Council]. (2011). Trustworthy Cyberspace: Strategic Plan for the Federal Cybersecurity Research and Development Program. Available at: https://www.nitrd.gov/pubs/Fed_Cybersecurity_RD_Strategic_Plan_2011.pdf. Accessed 27 August 2019.

OECD. (2016). The Economic Impact of Local Content Requirements. Available at: https://www.oecd.org/trade/topics/local-content-requirements/. Accessed 27 August 2019.

Otuoze, A., Mustafa, M. W., & Larik, R. M. (2018). Smart grids security challenges: Classification by sources of threats, *Journal of Electrical Systems and Information Technology*, *5*(3), 468–483, December. DOI: 10.1016/j.jesit.2018.01.001

Palo Alto. (2014). Getting Started With a Zero Trust Approach to Network Security, 25 March. Available at: https://www.bankinfosecurity.com/whitepapers/getting-started-zero-trust-approach-to-network-security-w-973. Accessed 27 August 2019.

Park, D., Summers, J., & Walstrom, M. (2017). Cyberattack on Critical Infrastructure: Russia and the Ukrainian Power Grid Attacks, *University of Washington, The Henry M. Jackson School of International Studies*, 11 October. Available at: https://jsis.washington.edu/news/cyberattack-critical-infrastructure-russia-ukrainian-power-grid-attacks/. Accessed 23 March 2019.

Qiang, C. Z.-W., Rossotto, C. M., & Kimura, K. (2009). Economic Impacts of Broadband, *World Bank Report*, Chapter 3. Available at: https://siteresources.worldbank.org/EXTIC4D/Resources/IC4D_Broadband_35_50.pdf. Accessed 23 March 2019.

Saadi, R., Rahaman, M. A., Issarny, V., & Toninelli, A. (2011). Composing Trust Models towards Interoperable Trust Management. In: Wakeman, I., Gudes, E., Jensen, C. D., & Crampton, J. (eds), *Trust management V*. IFIPTM 2011, *IFIP Advances in Information and Communication Technology*, *358*, 51–66. Berlin: Springer. DOI: /10.1007/978-3-642-22200-9_7

Sen, J. (2010). A Distributed Trust Management Framework for Detecting Malicious Packet Dropping Nodes in a Mobile Ad-hoc Network, *International Journal of Network Security & Its Applications*, *2*(4), 92-104, October. DOI: 10.5121/ijnsa.2010.2408

Sherwood, J., Clark, A., & Lynas, D. (2005). *Enterprise Security Architecture: A Business-Driven Approach*, CRC Press. ISBN:9781578203185

Sindhuja, K., Nasrinbanu, A., & Elavarasi, K. (2015). Survey on Malicious Node Detection and Reliable Data Fusion in MANET, *International Journal of Scientific Research Engineering & Technology*, *4*(3), 202-205, March.

Skopik, F., Ma, Z., Bleier, T., & Gruneis, H. (2012). A Survey on Threats and Vulnerabilities in Smart Metering Infrastructures, *International Journal of Smart Grid and Clean Energy*, *1*(1), 22-28.

Speicher, C. (2011). Security Fabric – Tailored Trustworthy Space, Part1: Flexibility Based on Policy Management, *Grid Interop Forum 2011*. Available at: https://www.gridwiseac.org/pdfs/forum_papers11/speicher_paper_part1_gi11.pdf. Accessed 26 March 2019.

Teikari, O., & Nevalainen, R. (2014). Comparison of software safety standards IEC 61508-3 and IEC 62138, *VTT Research Report* VTT-R-03820-14. Available at: https://www.vtt.fi/inf/julkaisut/muut/2014/VTT-R-03820-14.pdf. Accessed 23 March 2019.

US DoD. (1983). Department of Defence Trusted Computer System Evaluation Criteria. Available at: https://csrc.nist.gov/CSRC/media/Publications/white-paper/1985/12/26/dod-rainbow-series/final/documents/std001.txt. Accessed 26 August 2019.

Varghese, S. (2019). Huawei cyber testing centre rejection by Australia 'an old story', *ITWire*, 6 March. Available at: https://www.itwire.com/government-tech-policy/86272-huawei-cyber-testing-centre-offer-to-australia-an-old-story.html. Accessed 26 March 2019.

Volz, D. (2017). Trump signs into law U.S. government ban on Kaspersky Lab software, *Reuters*, 13 December. Available at: https://www.reuters.com/article/us-usa-cyber-kaspersky/trump-signs-into-law-u-s-government-ban-on-kaspersky-lab-software-idUSKBN1E62V4. Accessed 23 March 2019.

Wlodarczak, P. (2017). Cyber Immunity - A Bio-Inspired Cyber Defense System, *Lecture Notes in Computer Science*, *10209*, April. DOI: 10.1007/978-3-319-56154-7_19

Worstall, T. (2013). If Apple Brought iPhone Manufacturing To The US It Would Cost Them $4.2 billion, *Forbes*, 25 September. Available at: https://www.forbes.com/sites/timworstall/2013/09/25/if-apple-brought-iphone-manufacturing-to-the-us-it-would-cost-them-4-2-billion/#1fdce952115f. Accessed 23 March 2019.

Zeadally, S., Pathan, A-S. K., Alcaraz, C., & Badra, M. (2013). Towards Privacy Protection in Smart Grid, *Wireless Personal Communications*, *73*(1), 23-50, November.

Zhang, D. (2017). Intrusion Tolerance for CT Cloud Security, *RSA Conference*, Abu Dhabi.

Zouridaki, C., Mark, B. L., & Hejmo, M. (2007). Byzantine Robust Trust Establishment for Mobile Ad-hoc Networks, *Telecommunications Systems*, *35*(3-4), 189-206, August. DOI: 10.1007/s11235-007-9047-z

Zouridaki, C., Mark, B. L., Hejmo, M., & Thomas, R. K. (2007). HERMES: A Quantitative Trust Establishment Framework for Reliable Data Packet Delivery in MANETs, *Journal of Computer Security*, Special Issue on Security of Ad-Hoc and Sensor Networks, *15*(1), 3-38, January. DOI: 10.3233/JCS-2007-15102

E-Learning and the National Broadband Network

Simon Moorhead
Ericsson Australia and New Zealand

Abstract: A recent paper from February 2013 foreshadowing the dynamic changes in e-Learning from Australia's roll-out of the National Broadband Network.

Keywords: Telecommunications, History, National Broadband Network, Education

Introduction

This historic paper is only six years old and was selected to complement the other articles in this issue covering the future of the National Broadband Network (NBN). The NBN roll-out commenced around ten years ago and is well on the way to providing high-speed broadband access to most Australians across the continent. High-speed broadband provides many potential benefits, such as learning via electronic media (e-Learning), but also facilitates disruption, as we have seen with streaming services and social media, threatening some traditional information providers.

The paper (Barber, 2013) was written by James Barber of the University of New England and argues that the NBN will accelerate dramatic changes in education and teaching: in particular, the move away from bricks-and-mortar campuses towards global networks and the rise of mobile learning (m-Learning). When combined with Massive Open Online Courses, m-Learning will result in access to education becoming a universal human right.

Most readers would be aware of the proliferation of devices in the home which are enabled for e-Learning, such as smart TVs, smart mobile telephones and intelligent appliances. Combine this with the seemingly unlimited resources available to consumers on the internet, and you can appreciate the unprecedented learning opportunities, which the NBN will facilitate.

The historic paper concludes with the observations that "[u]ntil the invention of the Internet, universities did not have to be innovative because they have effectively had a monopoly" (p. 12.5) and "[l]et us hope that Australian universities embrace the opportunity that the NBN provides before it becomes a threat to them" (p. 12.6).

The disruption to traditional learning will continue as the technical capabilities of appliances expand and the NBN facilitates more diverse access to e-Learning resources.

http://doi.org/10.18080/jtde.v7n3.201

References

Barber, J. G. (2013). E-Learning: Supplementary or disruptive?, *Telecommunications Journal of Australia*. *63*(1), 12.1-12.6, February.

The Historic Paper

E-learning: Supplementary or disruptive?

James G. Barber
University of New England

The rollout of the National Broadband Network in Australia will accelerate dramatic changes in pedagogy and access that have been underway since the advent of the Internet. Among the most important of these are the move away from bricks-and-mortar campuses towards global learning networks that share expertise and resources, the blurring of the virtual and the material and the rise of mobile learning (m-learning). When combined with the proliferation of MOOCs (Massive Open Online Courses), m-learning will result in access to education becoming a universal human right.

Introduction

The advent of the Internet around two decades ago provided universities with three monumental opportunities. Any one of these should have been sufficient to alter the way teaching was conducted but in combination should have revolutionised university education.

- First, educators now had the ability to extend their reach globally.
- Second, anyone with an Internet connection suddenly had unprecedented access to information.
- And third, with the development of Web 2.0 technologies, digital communication began to provide a medium, and for many now the preferred medium, for interacting with friends and associates.

Despite this, it was only five years ago that I was lamenting how slow Australian universities had been to exploit educational technology in the service of learning and teaching. What I said at that time was:

> *"educational technology has yet to fulfill its promise to the extent witnessed in other sectors of the economy. The cottage craft of teaching and learning in universities has constrained IT's use. Indeed, in many universities, technology is often viewed as a barrier, even antithetical, to "genuine" education."*

How quickly things have moved since then. It seems that every day now, there is an announcement about some new online start-up, innovation or development in e-learning. Admittedly, the most exciting of these developments are occurring outside of Australia but since there are no geographic boundaries in cyberspace, Australian universities have no choice but to respond.

According to the National Center for Education Statistics (NCES), the number of U.S. students enrolled in at least one distance education course increased from 1 million to 12 million between 2002 and 2006, and the growth spurt shows no sign of abating. Research by the Sloan Consortium, for example, recently found that online college enrolments have continued to grow faster than the total population of college students. And respected market research firm, Ambient Insight, expects online enrolments in the U.S. to rise by a further 10 million in the next two years alone. By that time, the number of students taking all of their classes online will increase to 3.55 million while the number of students taking all of their courses in on-campus mode will actually decline.

It is difficult to obtain authoritative statistics on the rate of growth in Australia, but IBIS World recently estimated that the revenue from courses offered at least 80% online grew by 56% between 2008-9 and 2011-12 in Australia, while the number of online providers (i.e. those offering courses with more than 80% online content) grew by 40%, (from 764 to 1,082) over the same period.

More significant than the growth in uptake of online learning are changes that are now occurring in the nature and quality of e-learning itself. Until recently, the most common view of educational technology was that it was at best supplementary to conventional forms of instruction. The delivery of online courses mimicked the lecture-tutorial approach that has been the staple of university education for centuries. Lecture materials were provided, often in PDF form or, for the more advanced lecturers, as pod- or vodcasts. This material was then supplemented by online versions of tutorials, such as synchronous or asynchronous bulletin boards and chat rooms.

We are at last beginning to display more imagination in the application of educational technology as e-learning evolves into a genuine alternative to traditional teaching practices. As broadband further liberates us from the constraints of time and place, it will accelerate the movement of university education away from campuses, desktops and teacher-centred pedagogy towards learning networks, a merger between the virtual and the material, mobile learning, and a radically new, student-centred form of pedagogy.

Advances in Educational Technology

With the arrival in Australia of broadband, just about anything you could only do on a university campus will soon be available to you anywhere, anytime. Of course there will always be some learning activities, such as clinical placements and medical procedures that require students to be physically present somewhere at a specified time but the range of these activities will become narrower as the available bandwidth becomes broader. At this university, for example, certain aspects of emergency medicine are now taught by specialists from the medical school at the University of California (Irvine). Staff and students from UNE connect via telepresence technology to UCI's robotic simulation ward where they direct UCI's ward assistants to administer treatment in response to simulated medical emergencies. The responses of the robotic patient are observed by the students, vital signs are downloaded onto UNE laptops and displayed in real time on monitors adjacent to the telepresence screen. Meanwhile, staff and students at both universities work together to adjust interventions and debrief outcomes. Similarly, academic staff in chemistry and physics are designing practical classes whereby students can manipulate scientific equipment remotely and download their results onto PCs or i-Pads. With the aid of video cameras mounted around the laboratories, students watch the equipment come to life as they issue instructions from anywhere they can obtain a broadband Internet connection.

In place of the traditional campus, then, what we are witnessing is the emergence of nodes of activity or learning networks that are physically dispersed around the country and around the world. This movement into cyberspace will accelerate along with advances in hardware and software that are blending the virtual and the material. Virtual environments are being created that mimic the real world and provide us with a visceral sense of immersion. Some have even argued that the distinction between virtual and real will disappear altogether. This is because all surfaces, including the skin are potential interface points enabling users to issue and receive computer commands using their own body parts as touchpads.

We are also seeing rapid progress in augmented reality, where an overlay of data or interactivity is created on top of the real world around us. For example, the Google Android now allows you to point your phone at, say, a restaurant and bring up reviews and contact details of that restaurant, along with seating plans, similar restaurants in nearby locations,

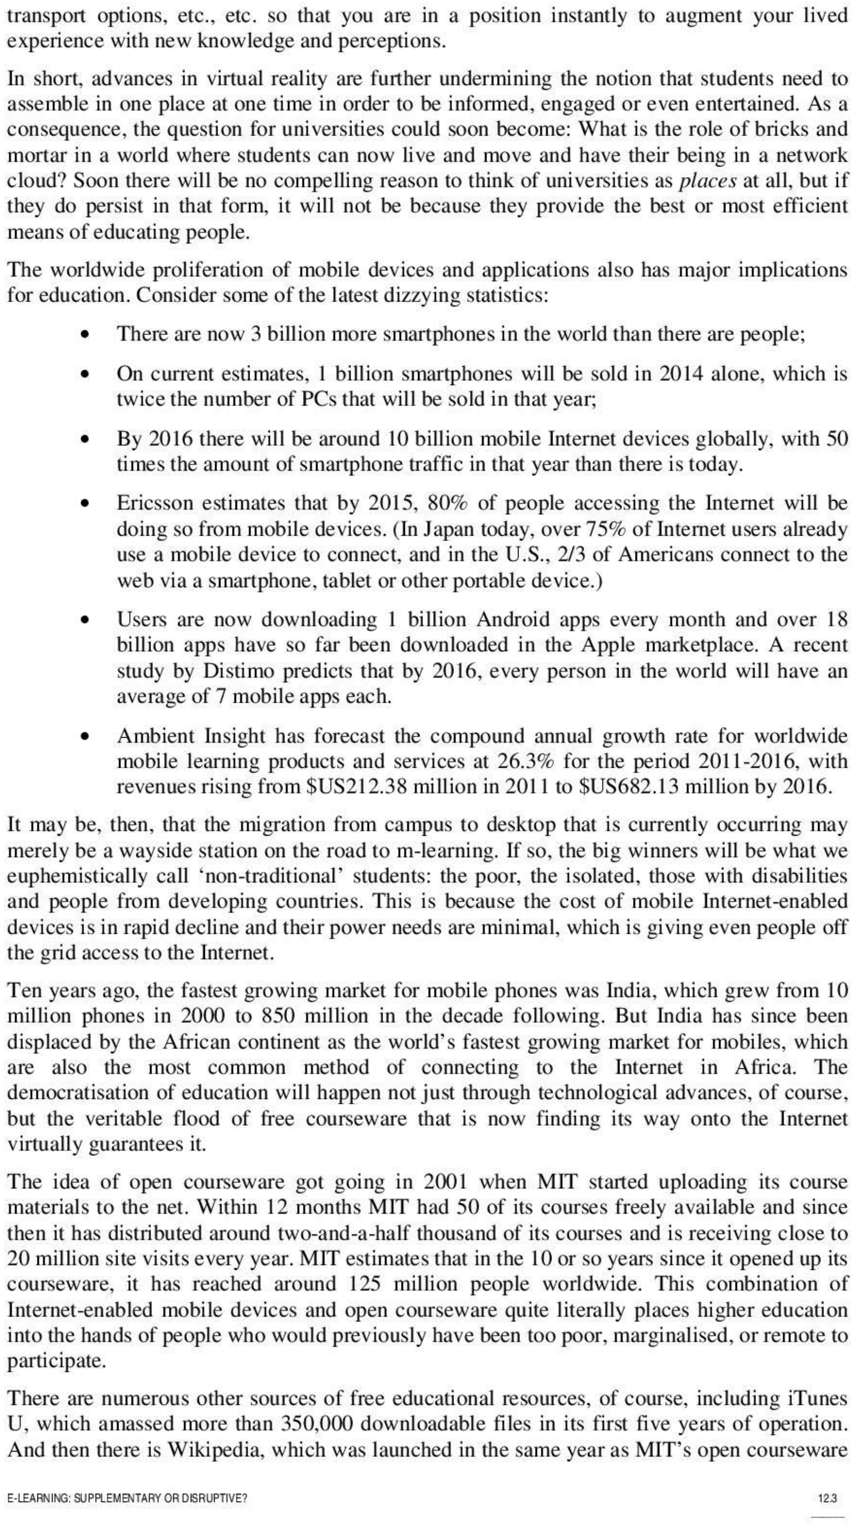

transport options, etc., etc. so that you are in a position instantly to augment your lived experience with new knowledge and perceptions.

In short, advances in virtual reality are further undermining the notion that students need to assemble in one place at one time in order to be informed, engaged or even entertained. As a consequence, the question for universities could soon become: What is the role of bricks and mortar in a world where students can now live and move and have their being in a network cloud? Soon there will be no compelling reason to think of universities as *places* at all, but if they do persist in that form, it will not be because they provide the best or most efficient means of educating people.

The worldwide proliferation of mobile devices and applications also has major implications for education. Consider some of the latest dizzying statistics:

- There are now 3 billion more smartphones in the world than there are people;
- On current estimates, 1 billion smartphones will be sold in 2014 alone, which is twice the number of PCs that will be sold in that year;
- By 2016 there will be around 10 billion mobile Internet devices globally, with 50 times the amount of smartphone traffic in that year than there is today.
- Ericsson estimates that by 2015, 80% of people accessing the Internet will be doing so from mobile devices. (In Japan today, over 75% of Internet users already use a mobile device to connect, and in the U.S., 2/3 of Americans connect to the web via a smartphone, tablet or other portable device.)
- Users are now downloading 1 billion Android apps every month and over 18 billion apps have so far been downloaded in the Apple marketplace. A recent study by Distimo predicts that by 2016, every person in the world will have an average of 7 mobile apps each.
- Ambient Insight has forecast the compound annual growth rate for worldwide mobile learning products and services at 26.3% for the period 2011-2016, with revenues rising from $US212.38 million in 2011 to $US682.13 million by 2016.

It may be, then, that the migration from campus to desktop that is currently occurring may merely be a wayside station on the road to m-learning. If so, the big winners will be what we euphemistically call ‘non-traditional’ students: the poor, the isolated, those with disabilities and people from developing countries. This is because the cost of mobile Internet-enabled devices is in rapid decline and their power needs are minimal, which is giving even people off the grid access to the Internet.

Ten years ago, the fastest growing market for mobile phones was India, which grew from 10 million phones in 2000 to 850 million in the decade following. But India has since been displaced by the African continent as the world’s fastest growing market for mobiles, which are also the most common method of connecting to the Internet in Africa. The democratisation of education will happen not just through technological advances, of course, but the veritable flood of free courseware that is now finding its way onto the Internet virtually guarantees it.

The idea of open courseware got going in 2001 when MIT started uploading its course materials to the net. Within 12 months MIT had 50 of its courses freely available and since then it has distributed around two-and-a-half thousand of its courses and is receiving close to 20 million site visits every year. MIT estimates that in the 10 or so years since it opened up its courseware, it has reached around 125 million people worldwide. This combination of Internet-enabled mobile devices and open courseware quite literally places higher education into the hands of people who would previously have been too poor, marginalised, or remote to participate.

There are numerous other sources of free educational resources, of course, including iTunes U, which amassed more than 350,000 downloadable files in its first five years of operation. And then there is Wikipedia, which was launched in the same year as MIT’s open courseware initiative and now contains more than 15 million articles (only 20% of which are in English), all of which are continually updated and corrected by subscribers themselves.

The enormous appeal of Wikipedia demonstrates another profound shift in the way that universities of the future will teach – the movement away from *acquisition* of knowledge as the fundamental purpose of education to incorporate its *creation* and *re*-creation by students themselves. This is unfamiliar territory for academics of my generation who were raised on the idea that only professional educators are qualified to teach. But Facebook, Wikipedia and blogging have radically undermined this assumption because all consist of information that is created by, not just communicated to, participants. The acts of teaching and learning are blurring as a consequence. Schooled on Google and Wikipedia, students today want to inquire, not rely on the professor. They want a conversation, not a lecture.

The most recent development in open courseware is of course the MOOC, which is an acronym for 'Massive Online Open Courses' in which huge numbers of students enrol in online courses, network with one another online and undertake online quizzes and self-directed learning. The term MOOC was first coined in 2008 but entered common parlance only towards the end of 2011 when Stanford University professor, Sebastian Thrun, offered to enrol students in his online robotics course free of charge and 160,000 people took up the offer. Buoyed by the success of his experiment, Thrun and his colleagues launched a free online university called Udacity in February of 2012 and within the first three months of operation had achieved over 100,000 enrolments.

Six days after Udacity, coursera.org was launched by a star-studded line-up of U.S. universities including the University of Pennsylvania; the University of Michigan; and Princeton and Stanford Universities. These universities offer their courseware free of charge online, and there is facility for students to interact with one another and take quizzes to monitor their progress. By the end of April, coursera had amassed a staggering 1,000,000 enrolments.

Not to be outdone, two weeks after coursera was launched, MIT and Harvard University joined forces to launch edX and on 1st August 2012, arguably the world's finest public university UC (Berkeley) threw in its lot with edX. Like other MOOC providers, edX also offers free online courseware to students around the world; its stated goal is to exceed one billion student enrolments in the next decade.

Effectiveness of e-Learning

Among the most common objections to e-learning continues to be that it is a very poor substitute for face-to-face teaching. Given the technological advances described earlier, however, this objection rather begs the question of whether the distinction between "face-to-face" and "online" has any real meaning in a post-NBN world. But even before the improvements that broadband will bring, there was solid evidence in support of e-learning methods.

Prior to the development of Web 2.0 technologies like MySpace and Facebook, there had been two major meta-analyses of the effectiveness of online education (Bernard et al 2004; Cavanaugh et al 2004). Meta-analysis is a technique for combining the statistical results of multiple research studies to obtain a composite estimate of the size of the effect. In this way, different studies using different online techniques and different measures of learning can be combined into a single study and a global measure of the effectiveness of online learning calculated.

The result of each experiment is first expressed as an *effect size*, which is the difference between the mean score for online learning and the mean score for face-to-face classes divided by the pooled standard deviation. Individual effect sizes are then combined into a single, overall index. Importantly, meta-analysis is only ever performed on studies that satisfy the most rigorous methodological standards, normally involving random allocation of students to classroom and online conditions.

Results of these pre-Web 2.0 meta-analyses suggested that there was no significant difference in learning outcome between distance education and face-to-face education. Findings of a more recent meta-analysis of job-related courses comparing Web-based and classroom-based learning (Sitzmann et al 2006) were even more positive. This study found online learning to be superior to classroom-based instruction in declarative knowledge (or knowing *that*) outcomes, with the two being equivalent in procedural knowledge (or knowing *how*) outcomes.

A more recent meta-analysis conducted by the United States Department of Education (Means et al 2010) combined effect sizes from 46 separate studies conducted since the advent of Web 2.0 technologies. The number of students in the various studies ranged from 16 to 1,857, with student ages ranging from an average of 13 to 44 years. Importantly also, the subject matter captured by the meta-analysis was very diverse, extending from medicine, to education, law enforcement, introductory science, problem-solving skills, computer science, optometry, veterinary science, and multi-media studies.

The overall finding of the meta-analysis was that classes with online learning (whether taught completely online or blended) actually produced better learning outcomes than classes with solely face-to-face instruction. Admittedly, the effect size was only moderate (+0.24) but it was sizeable enough to dispel the myth that face-to-face instruction is always the optimal form of teaching. This same meta-analysis compared three conditions:

(a) face-to-face only,

(b) online only, and

(c) blended or a combination of face-to-face with online learning.

Comparisons of these conditions revealed that the third actually produced the most powerful effects ($g+$ = +0.35 vs face-to-face; compared with $g+$ = +0.14 for online only), with the difference between online only and blended also being statistically significant. In short, then, as the technology has improved so has the relative advantage of online learning over face-to-face instruction and as broadband erodes the very distinction between online and face-to-face, the debate itself will be consigned to history.

Conclusion

Up to this point, I have written as if universities will continue to be the main providers of higher education and gatekeepers to the professions, but it would be a mistake for universities to take this privileged position for granted in the post-NBN world. In her confronting introduction to a U.S. Department of Education (2006) report on the future of higher education, Education Secretary Margaret Spellings wrote that:

> *"What we have learned over the last year makes clear that American higher education has become what, in the business world, would be called a mature enterprise: increasingly risk-averse, at times self-satisfied, and unduly expensive. It is an enterprise that has yet to address the fundamental issues of how academic programs and institutions must be transformed to serve the changing educational needs of a knowledge economy… .History is littered with examples of industries that, at their peril, failed to respond – or even to notice – changes in the world around them, from railroads to steel manufacturers. Without serious self-examination and reform, institutions of higher education risk falling into the same trap, seeing their market share substantially reduced and their services increasingly characterised by obsolescence"* (Spellings 2006, p.xii).

Until the invention of the Internet, universities did not have to be innovative because they have effectively had a monopoly. But competition from virtual institutions like ed-X, Udacity and coursera is changing all that. In his research on the implementation of new ideas, Harvard Professor Clayton Christensen found that disruptive innovations rarely come out of established enterprises. Even when a truly new way of doing things does occur to someone in a traditional organisation, established systems and standards take over and the idea is usually quashed. A new idea that is not dismissed entirely is almost inevitably modified to fit the way things are traditionally done, losing its innovation impact in the process. Let us hope Christensen is wrong. Universities must surely be among the very finest institutions that society has created. Let us hope that Australian universities embrace the opportunity that the NBN provides before it becomes a threat to them.

References

Bernard, Robert M; Abrami, Philip C; Lou, Yiping; Borokhovski, Evgueni; Wade, Anne; et al. 2004. How Does Distance Education Compare With Classroom Instruction? A Meta-Analysis of the Empirical Literature. *Review of Educational Research*, vol. 74 no. 3, 379-439. http://dx.doi.org/10.3102/00346543074003379

Cavanaugh, Cathy; Gillan, Kathy Jo; Kromrey, Jeff; Hess, Melinda; Blomeyer, Robert. 2004. The Effects of Distance Education on K-12 Student Outcomes: A Meta-Analysis

Christensen, Clayton M; Eyring, Henry J. *The innovative university: Changing the DNA of higher education from the inside out.* Jossey-Bass, 2011.

Means, Barbara; Toyama, Yukie; Murphy, Robert; Bakia, Marianne; Jones, Karla. 2010. Evaluation of Evidence-Based Practices in Online Learning: A Meta-Analysis and Review of Online Learning Studies. Technical Report. U.S. Department of Education, Washington, D.C..

Sitzmann, T, Kraiger, K, Stewart, D & Wisher, R. (2006) The comparative effectiveness of web-based and classroom instruction: A meta-analysis. *Personnel Psychology*, 59, 623-664. http://dx.doi.org/10.1111/j.1744-6570.2006.00049.x

Spellings, M. 2006. Commission on the Future of Higher Education. US Dept of Education, Washington DC.

Cite this article as: Barber, James, G. 2013. 'E-learning: Supplementary or disruptive?'. *Telecommunications Journal of Australia* 63 (1): 12.1-12.6. Available from: http://tja.org.au.

James G. Barber

University of New England

Professor Jim Barber is the University of New England's Vice-Chancellor and Chief Executive Officer.

Before taking up this position in February 2010, Professor Barber was Deputy-Vice Chancellor at the Royal Melbourne Institute of Technology (RMIT) University.

Professor Barber is a distinguished academic. After completing his PhD in experimental psychology, his research shifted into the applied fields of drug addiction and child welfare. His research record includes minimal interventions in the secondary prevention of drug addiction, and evidence-based social policy and child welfare. He is a winner of North America's Pro Humanitate Medal for his research in child welfare and a winner of the Vice Chancellor's Award for Excellence in Teaching from Flinders University.

Prior to moving to university senior executive positions in the higher education sector, his roles included that of Reader and then Professor of Social Work (La Trobe University and the University of Tasmania), Professor of Social Administration (Flinders University) and Dean at the University of Toronto.

Professor Barber's experience includes roles of Company Director on a number of national bodies, including Open Universities Australia (Australia's leading provider of fee-paying online degree programs), Jesuit Social Services Australia and Graduate Careers Australia.

Professor Barber has significant international education experience, most significantly taking on the additional role of interim President of RMIT International University of Vietnam. He has worked in regional universities and has a commitment to their important contribution in providing access to education, and also in driving economic prosperity and enhancing the morale, culture and identity of their regions.