AJTDE Volume 3, Number 1, March 2015

Table of Contents

Thank you and looking forward

Mark A Gregory
RMIT University

Summary: On behalf of the Editorial Board, authors and readers of the Journal, the recently appointed Managing Editor Mark Gregory thanks the outgoing Managing Editor Peter Gerrand for his leadership, scholarly editing and hard work over the past 21 years.

Thank you

After 80 editions of the Journal since June 1994, when it was named the *Telecommunications Journal of Australia* and after November 2013 the *Australian Journal of Telecommunications and the Digital Economy*, Professor Peter Gerrand has stepped down from the role as Editor-in-Chief and Managing Editor.

On behalf of the Editorial Board, the many authors and the *Journal's* loyal readership, I would like to thank Peter for his excellent contribution over the past 21 years. Peter's hard work, determination, friendly and cheerful manner and willingness to publish points of view from both sides of a debate have ensured that the *Journal* has remained relevant and published articles of substance about the changing telecommunications landscape.

Peter's career in telecommunications spans many decades and his leadership as a senior telecommunications academic, first at RMIT University and later at the University of Melbourne, underpinned his broader involvement with the telecommunications industry.

Over the past 21 years the *Journal* has included articles on telecommunications-related public policy, technology, consumer issues, legislative and regulatory matters and what the future may hold for an ever-changing telecommunications industry. Peter's editorials have focused on the positive, and his inclusive management of the *Journal* has resulted in a superbly broad coverage of the key telecommunications issues. A highlight of Peter's contribution is the enduring high standard of the *Journal's* articles and the publication of views from the leaders in the regional telecommunications industry.

Looking forward

The Editorial Board will continue to publish a *Journal* that is focused on regional telecommunications, and aims to build upon past Editorial initiatives that have resulted in a steady stream of high quality papers being included in each issue.

http://doi.org/10.18080/ajtde.v3n1.1

There will be new initiatives commencing in 2015, including shifting the *Journal* to the *Open Journal Systems* publishing platform and building upon the "theme editor" approach to encourage people with specialised knowledge of aspects of the telecommunications industry to bring together papers for upcoming issues.

The *Open Journal Systems* publishing platform will automate the publishing process and permit the *Journal* to be broadly indexed and build an Impact Factor, which is an important requirement if the *Journal* wishes to continue to attract high quality Academic papers.

A medium term goal is to consolidate the *Journal's* past issues in a single publishing platform which will ensure that the *Journal* is available to the widest possible audience and is able to benefit from the considerable value of the papers published over past decades.

A key aspect of the *Journal's* success has been the quality of the papers published in each of the issues; and the initiative to introduce issue "themes", as well as including appropriate submitted papers that meet the Editorial guidelines, aims to provide readers with a selection of papers from which a broader understanding can be obtained.

The historical and book review papers that have become a feature over past years will continue to provide us with an entertaining look at past and current topics. In 2015 an Editorial goal is to introduce a public policy paper in each issue to provide a perspective on how the telecommunications industry should be shaped moving forward.

In the role of Managing Editor, my objective is to work with colleagues on the Editorial Board to publish the *Journal* whilst maintaining the high standards that have been a hallmark of the *Journal* over many decades. As the Managing Editor, I will work to ensure that the Telecommunications Association's inclusive, broadly-minded and politically neutral stance will continue to be the foundation of Editorial policy.

Mark A Gregory

http://doi.org/10.18080/ajtde.v3n1.1

The world's fastest wireless backhaul radio

A case study in industry-research collaboration

Dr Rowan Gilmore
EM Solutions Pty Ltd

Dr Xiaojing Huang
CSIRO Digital Productivity and Services (previously) and currently University of Technology Sydney

Dr Richard Harris
EMClarity Pty Ltd

Summary: Fibre is commonly perceived to be the dominant transport mechanism for transferring data from access points back to a central office, where it is aggregated onto the core network. However, high speed and long range wireless backhaul remains a cost-effective alternative to fibre networks. In some areas, wireless backhaul is dominant and becoming more and more attractive. However, commercially available wireless backhaul systems do not meet the requirements for both high speed and long range at the same time with sufficiently low latency for some applications. Traditional microwave systems can achieve long transmission range, but the data rates are then limited to a few hundred megabits per second. Multi-gigabit per second wireless communications can be achieved using millimetre-wave (mm-wave) frequency bands, especially in E-band, but the practical transmission range has then always been a major weakness.

In this article, the world's first 5Gbps radio solution – and the fastest commercial backhaul product – developed by EM Solutions Pty Ltd with the Commonwealth Scientific and Industrial Research Organisation (CSIRO) – is described. As well as achieving a state-of-the-art data rate, other key design features include maximal path length, minimal latency, and constant antenna pointing under wind and tower vibration.

Introduction

According to the Cisco visual networking index (VNI) global mobile data traffic forecast (Cisco), global mobile data traffic reached 1.5 Exabytes (1 Exabyte = 2^{60} bytes) per month at the end of 2013, and is expected to grow 10 times by 2018, a compound annual growth rate (CAGR) of 61 percent (Cisco 2014). The rapid increase in mobile data traffic shown in Figure 1 is mainly due to the widespread deployment of broadband wireless access services enabled by smart phones and tablets, as well as next generation networks, such as the worldwide interoperability for microwave access (WiMAX) (Fu et al. 2010: 50 - 58) and fourth generation long term evolution (LTE)– Advanced systems (Fodor et al. 2011: 84 - 91), which

http://doi.org/10.18080/ajtde.v3n1.2

can offer data rates from hundreds of Megabits per second (Mbps) to 1 Gigabit per second (Gbps). These high-speed access networks are attractive alternatives to wired access networks such as fibre optic networks, and provide broadband services in a cost-effective manner. In the meantime, network operators are also seeking solutions to support low latency applications, such as real-time machine-to-machine communications, interactive multiplayer gaming, high frequency trading, and cloud computing (Transmode).

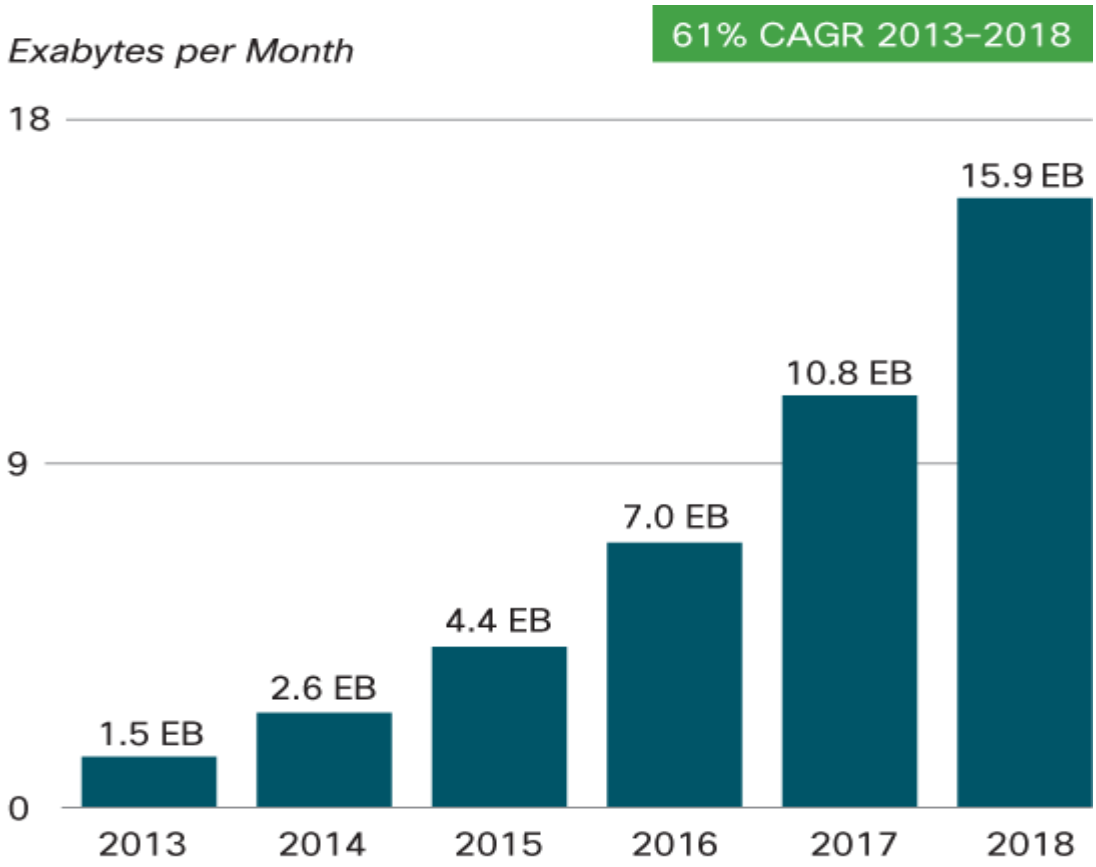

Figure 1: Cisco forecasts mobile data traffic of 15.9 Exabytes per month by 2018 (Cisco 2014)

Due to the ever-growing capacity required to support high speed broadband services, the backhaul network, which transfers data traffic from cell sites of a wireless access network to the core network or a switching centre, is under intense pressure.

There are a number of challenges to implement such backhaul. The first is how to achieve **higher data rate** or capacity up to multiple Gigabits per second. For example, if the capacity of a cell (or sector) in a broadband wireless access (BWA) base station is 1 Gbps, the backhaul capacity required by a three sector base station would be at least 3 Gbps. Sometimes the traffic from multiple base stations will be aggregated together across a backhaul link before reaching the core network. This will drive the backhaul capacity to a much higher rate, say 10-15 Gbps.

http://doi.org/10.18080/ajtde.v3n1.2

The second challenge is the **link distance** of the backhaul. To deliver BWA services to unserved areas, such as rural and regional areas that are often quite remote from the main telecommunication infrastructure, a long distance backhaul link is required. In countries with large geographical areas of low population density, such as Australia, long-distance high-data rate wireless backhaul links are essential to bring broadband services to the last 10-30% of the population. Of course, the data rate challenge is not as great in rural areas because of the sparsity of users, but over time this too may become an issue.

The third challenge is how to achieve **low latency** communications between end users across the backhaul networks. Although low latency has always been important for delivering high quality voice, video and data services in broadband networks, recent application requirements within many industry sectors, such as gaming and finance, have made low latency an even more important consideration in data transport. It becomes even more critical as the intelligence within the network is centralised to reduce costs, and functions previously delivered at the access node must now be performed within the core network.

Fibre is the obvious solution – and indeed the primary medium – to deliver leased synchronous digital services and Ethernet services. It is the first choice for high data rate backhaul at speeds from 155 Mbps up to 10 Gbps. However, due to its high installation expense, owning a fibre is a significantly capital-intensive (CAPEX) option. It is estimated that leased lines account for roughly 15 percent of typical network operating expenditure (OPEX). Wireless backhaul can be more cost-effective than leased T1/E1, DS3, or OC-3 lines. In addition to the economic benefits of ownership, wireless backhaul also allows service providers to retain end-to-end control of their data, gaining the security, stability and freedom associated with full control over their own network. For less populated rural areas, where the cost to lay fibre can be prohibitive, wireless backhaul may be the only viable solution. Finally, because radio propagates over the air faster than light travels through fibre, wireless backhaul can achieve lower latency than fibre.

The purpose of this article is to describe the cooperation between CSIRO and EM Solutions to develop the world's fastest microwave backhaul radio, over the longest link distances, with the lowest latency. This radio was developed for an application in the US financial markets, and achieved 5 Gbps backhaul speeds on links up to 25 km with latencies of the order of one to two microseconds per hop over the free space propagation delay. Such results are several times faster and longer than other commercially available microwave backhaul solutions, which (at the time of writing) typically have processing latencies measured in tens of microseconds and maximum throughputs of 1 or 2 Gbps, over link lengths approximately one-half of those achieved here for the same availability.

http://doi.org/10.18080/ajtde.v3n1.2

The Technology Challenges

As noted above, the key design challenges are to achieve a state-of-the-art data rate, while maximising path length, minimising latency, and maintaining antenna pointing.

Multi-gigabit mm-wave communications systems are commercially available. Typical mm-wave bands suitable for wireless backhaul application include the 60 GHz band and 70/80 GHz E-band. Commercial point-to-point links in the 60 GHz band with data rates of up to 1.25 Gbps are sold by several manufacturers. However, high propagation loss due to oxygen absorption in this band and regulatory requirements limit the communication range for outdoor applications to at most 0.5-0.8 km.

The recent availability of the E-band spectrum worldwide provides an opportunity for line of sight (LOS) links with longer range and higher data rates, ideally suited for fibre replacement and backhaul applications. The merits of E-band wireless communications include the vast, uncongested and inexpensive spectrum, where a total of 10 GHz of available RF bandwidth enables very high data rates beyond 10 Gbps and the use of small, highly directional antennas. In some countries, contiguous bandwidth of up to 4GHz is available at E-band. Current commercial solutions provide low output power at relatively low infrastructure cost. Current commercial suppliers include BridgeWave, LOEA, Proxim, E Band Communications, Elva, Siae and Huawei. These E-band links are ideally suited for short range (1-3 km for all practical purposes) fibre-quality wireless communications. At the time of writing, solutions up to 2 Gbps using higher order modulation schemes (64QAM) (Huawei 2012) are reportedly available, but these require strong signal to noise ratios (SNR). In more common use are the simpler but more robust modulation techniques, such as amplitude shift keying (ASK) or binary phase shift keying (BPSK) with spectral efficiencies below one bps/Hz, but which only achieve lower data rates. Link distances are rarely more than a few km, in the absence of rain.

Technologies enabling much higher data rates (up to 24 Gbps) were developed and reported several years ago by the CSIRO ICT Centre (Huang et al 2012a : 122-129) (Huang et al 2011). The key algorithms developed (Zhang et al 2012 : 589-599; Huang et al 2012b : 2113-2122) are applicable to systems where the radio channel bandwidth is greater than the Nyquist bandwidth of the associated analogue-to-digital converter (ADC) and digital-to-analogue converter (DAC). Such systems can be utilised for E-band full-duplex wireless links, and can achieve a spectral efficiency scalable from 2.4 to 4.8 bit/s/Hz using 8-phase shift keying (8PSK) to 64QAM modulation, enabling data rates from 12 to 24 Gbps. This has been proven by experimental results on a 6 Gbps prototype that achieved a spectral efficiency of 2.4 bps/Hz (Dyadyuk et al 2007: 2813-2821).

http://doi.org/10.18080/ajtde.v3n1.2

Since the first report, mixed signal processing technologies, higher speed ADC and DACs, and larger scale FPGA devices are now available, and the spectral efficiency of the E-band system above can be further improved. Combined with dual polarisation, a high speed E-band link of up to 50 Gbps data rate can be achieved in 5 GHz bandwidth. However, such data rates have not yet been demonstrated.

Because radio waves propagate through air faster than light travels through fibre, wireless links can achieve lower end-to-end latency. This makes E-band wireless links the medium of choice compared with fibre networks for applications where latency is a concern, such as high frequency trading. Even though the range per hop of the E-band link is still limited, the end-to-end latency of a multi-hop E-band system can be still lower than for fibre if the processing delay at each radio relay node remains small. In mission critical cases, where fibre links can be used as a backup for adverse weather conditions, the per-hop distance of an E-band link can be extended to over 20 km. This will further reduce the overall end-to-end latency and deployment cost for the multi-hop E-band link.

The classic way of increasing link range is to maximise the link budget. The link budget is essentially the signal transmit power increased by the sum of receiver and transmitter antenna gains, reduced by the signal to noise ratio required at the receiver to achieve the desired bit error rate at the modulation and bandwidth necessary for the required data throughput.

Increasing the transmit power is an obvious first choice. The most common commercial E-band power amplifiers achieve saturated output powers of 100 mW, but output powers of 1W are now becoming available. Changing to a 1W power amplifier can theoretically increase the link margin by up to 10dB. However, it is linear power, not saturated power, that is critical in radio communications. In addition, more complex modulation schemes (to achieve a higher bit rate within the same bandwidth) require greater linearity to preserve symbol amplitude than do simpler schemes such as BPSK or QPSK. This means the average transmit power is usually considerably backed-off from the headline saturated power.

The other large variable in the link budget is the antenna gain – doubling the diameter of the antenna increases the link budget by 6 dB per end, or 12 dB for each hop. The most common E-band antennas are 300mm in diameter; therefore using a 1200mm antenna achieves a significant improvement in link budget and consequently range. Unfortunately, a 1200mm E-band antenna will have a 3-dB beamwidth of just 0.25 degrees, a pencil-thin beam. This makes it impossible to manually align two ends of such a link several km apart, or to maintain such alignment when the antennas and their towers are subjected to wind or thermal stresses that arise during normal operation.

http://doi.org/10.18080/ajtde.v3n1.2

Commercialisation of a long-range, low latency, high speed E-band radio

With a lead customer in the financial industry requiring a high speed (upgradable to 10 Gbps) low latency radio that could complement a fibre network and achieve a per hop span of 25 km, EM Solutions and CSIRO together developed a radio solution that was able to meet the required specifications. These are summarised in Table 1 below.

Table 1 – Target backhaul E-band radio specifications

Frequency	E-band (70 to 86 GHz)
Throughput	5 Gbit/s full duplex
Data Interface	10 GbE Optical SFP IEEE 802.3ae compliant
Latency:	Significantly less than 1 - 2 microsec per link end to end (one hop), excluding propagation delay
Minimum operational system budget	Gross system budget exceeding 192 dB assuming a minimum receiver SNR requirement of 9.5dB
Overall BER	10^{-11} per hop
Pointing and Tracking	ISM K-band for tracking compliant with US FCC regulations, to maintain optimal lock up to a physical antenna angular deviation of ±7 degrees per end
Network Management System interface	SNMP based messaging interface via Ethernet port from each radio

CSIRO's focus was on the radio modem. A number of novel techniques were developed in order to achieve high power efficiency and low latency. One major architectural difference with modern lower frequency digital radios was the use of analogue in-phase and quadrature (I/Q) modulation at the IF stage. Using the entire channel bandwidth of over 4 GHz as a single channel presented a number of technical difficulties, particularly relating to maintaining quadrature and gain flatness across the band. This could be estimated and compensated. Other novel techniques included advanced channel estimation and equalisation, algorithmic-efficient transmitter and receiver filter design and implementation, and overall system optimisation to achieve both low latency and high performance at the same time. This required a tradeoff between the number of processing loops involved with channel estimation and the improvement in signal-to-noise ratio that was achieved, against the desired processing latency.

Compared with other commercially available high speed E-band radios, CSIRO's low latency E-band system achieved some distinctive advantages. First, it has low processing latency due

http://doi.org/10.18080/ajtde.v3n1.2

to the smart signal processing architecture and signalling protocol, even though complicated error correction coding, equalisation, and practical impairment compensation algorithms are employed. Second, it can achieve 10 Gbps data rate with relatively low level modulation (QPSK and 16 QAM) so that it offers higher power efficiency and hence longer distance than more complex modulation formats. Third, with its advanced channel estimation and equalisation techniques, it can cope with harsh channel conditions with tens of nanoseconds delay spread caused by analogue circuitry, cable reflection, and multipath propagation. Finally, it offers flexible switching between Ethernet traffic and ultra-low latency relay traffic and among multiple radios, so that a multi-hop E-band link can be configured with fibre backup, suitable for low latency application under all weather conditions.

EM Solutions' focus was on the commercialisation and packaging of the radio modem itself, and the design and manufacture of the antenna and feed, the automatic pointing system and the overall network element management system. Although the management system was straightforward, design of an antenna pair that would automatically maintain alignment along the boresight between both ends proved a challenge. Tower vibration of even a fraction of a degree, for instance due to wind or thermal variation, will twist the antennas to the extent that communications would be impossible between them. One solution would be to use a phased-array antenna with electronic beam steering. Conceptually simple, such a solution was considered too costly to develop from scratch. The approach we adopted was to mount a high-gain parabolic antenna on a two-axis gimbal system, with low friction contact-less motors mounted within the bearings to drive the antenna in either the azimuthal or elevation direction as required to maintain boresight. The resulting radio is shown in Figure 2.

http://doi.org/10.18080/ajtde.v3n1.2

Figure 2: The commercialised E10G radio, with the 1.2m antenna and gimbal system housed with a protective shroud. Portion of the automatic pointing control system and radio modem can be seen mounted at the rear of the antenna.

Determining the required direction to maintain alignment borrowed an idea from radar technology, a technique known as monopulse. This uses a broad beamwidth signal (a beacon) emanating from the remote end, and measures the phase difference across two slots either side of the central feed at the near end antenna. If the near end antenna is pointing directly towards the far end beacon, the phase difference will be zero. The phase difference will gradually increase as the antenna is pointed off-boresight. With two pairs of slots along a horizontal and vertical axis, a vector to indicate the centre of boresight can be generated from these phase measurements, and used to drive the motors to re-centre the antenna along boresight (where the phase error will be zero, and the received beacon and data signals a maximum).

Because for initial acquisition the beacon signal from the far end needs to be of broad beamwidth to fall within the capture angle of the near antenna, it cannot be the E-band data signal itself, since that is too narrow in beamwidth (0.25 degrees for a 1200mm antenna). In the EM Solutions' system, the acquisition beacon was generated at K-band by a second co-located antenna at each end, a horn with a broader beamwidth and low enough gain to meet the radiated power regulations of the band, so that it illuminated the receiver end regardless of its own vibration and in whatever direction the near end or far end antenna is pointing (within reason). A very narrowband tuned receiver is able to detect the transmitted signal since the frequency is known according to a defined frequency plan. A concept is shown in

http://doi.org/10.18080/ajtde.v3n1.2

Figure 3. Such a system proves remarkably effective at acquiring the remote end in a matter of seconds and maintaining lock even when the antennas at either end were violently shaken. This should be compared with fixed high gain antennas that can otherwise take hours to manually align using optical telescopes and maximum signal strength indicators, and that have no resilience to any small motion at either end.

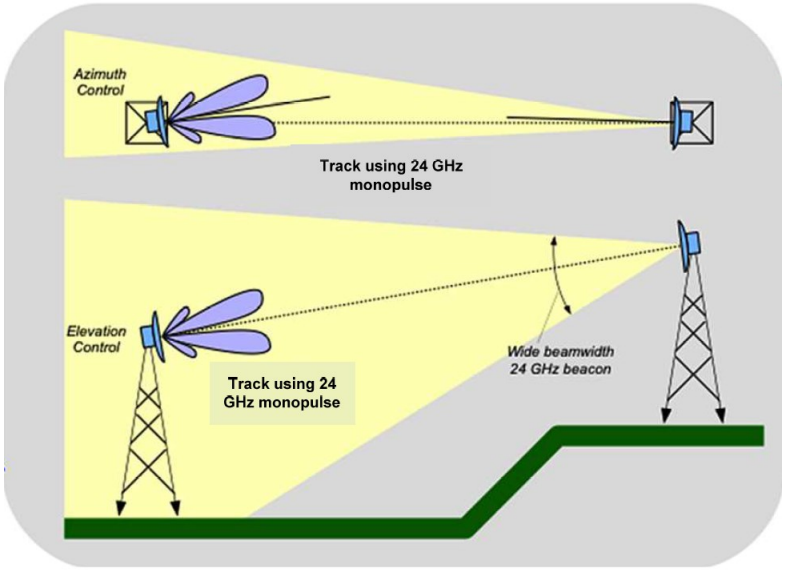

Figure 3: The concept of pointing using a K-band beacon transmitter at the remote end for monopulse detection at the near end, to steer the near end antenna back to boresight.

Field Results

The system was installed in a range crossing the upper reaches of the Brisbane River, between a University of Queensland site on the outskirts of Brisbane, at Pinjarra Hills, with line of sight visibility to a site 16km away owned by the Bureau of Meteorology at Springfield, near Ipswich. The system was transported using a special frame on a flat-bed truck. Lifting, placement, and tie down took no longer than 20 minutes. Lifting was simplified because the unit is neutrally balanced, and fitted with side hooks for steering lines. The antenna and its pointing system can be mounted to a standard 120 mm pole, using either a low weight normal-strength or high-strength mounting bracket. Figure 4 shows the unit being lifted and installed in place.

http://doi.org/10.18080/ajtde.v3n1.2

Figure 4: (a) One end of the E-band radio being lifted for mounting and (b)the radio installed at one end of the trial link.

Removal of the gimbal packing and fitting of the radome takes a further 20 minutes, and the electronics modules are then attached to the rear of the unit using hex keys within a similar period of time. Three cables are required, the first for the optical data payload, the second for optical control and monitoring, and the third to provide 48V DC power. Power up and initial scan takes approximately 2 minutes, and requires only coarse pointing towards the far end (which is barely visible). As the system automatically acquires the remote end, feedback on the pointing response, off angle error, bit error rate, and received signal strength is achieved through the user interface.

Upon start-up, each end searches for the wide angle beacon signal emanating from the other end. Each system then locks on to the beacon and moves to closed-loop pointing control. At this stage, depending on the motion at the near end, the system begins the search for a narrow beam, higher strength beacon signal that is transmitted using the main antenna reflector at the far end, rather than the low gain horn. This beacon signal is more directional than the first that was used for initial acquisition, and can assist general pointing performance due to its higher beacon signal to noise ratio, as well as in cases where multipath propagation might affect the broader beamwidth beacon. Gyroscopes that measure the acceleration along the two axes of the antenna are also used in a feedback loop to adjust the pointing direction.

Once the system has optimised its pointing, data communications at E-band can begin. Pointing accuracy was measured in real time to be within 50 millidegrees of true boresight. The QPSK constellation shown in Figure 5 was observed to be clear and achieved bit error rate less that 2E-13 for carrier to noise ratios better than 9.5 dB consistently over a path in excess of 15km.

http://doi.org/10.18080/ajtde.v3n1.2

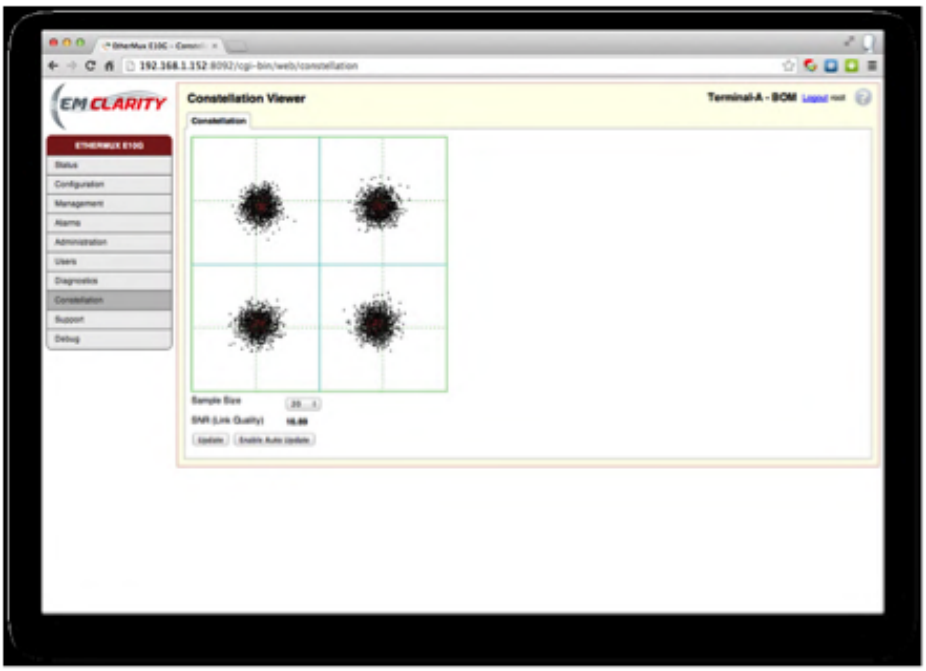

Figure 5: Constellation diagram of the received QPSK signal after downconversion from E-band to baseband.

The link was left running for a total of 30 days in various types of weather. Outages occurred during thunderstorms, since heavy rain is essentially impenetrable at E-band for distances exceeding 5 km (for this system; and 1km for other commercial systems). For example, on December 19, a typical summer day, the link margin was measured to be 17dB when the temperature was 27 degrees and humidity 43%. On December 12, when the humidity was greater than 90%, the link margin had dropped to 11.9dB over the same 16km path length (the e^{-11} BER threshold SNR of the system is 9 dB when running QPSK). In both cases, no bit errors were observed over a 24 hour period with data passing at a throughput of 5Gbps over the link. Figure 6 shows the recorded SNR and RSSI and other performance parameters as a function of time.

http://doi.org/10.18080/ajtde.v3n1.2

Figure 6: A performance log of the system during field trial. The upper traces show temperature of critical system components and the lower traces show SNR, Tx power and Rx power during a typical day of testing.

Latency could not be measured in the field since its measurement requires a loopback, and this would include a long and potentially variable propagation time over the transmission path. Instead, a known length of fibre cable was used for testing in the lab, and latency tests were conducted on the inter-link connection, with a latency measured to be well under two microseconds. The latency from the 10Gig Ethernet input port to the link side of the radio was also measured to be similar for small packets.

Comparison with Other Technology Options

There are a number of alternative ways to achieve further improvements in the data rate, such as employing higher order modulation, making use of multiple input multiple output (MIMO) antennas, improving link availability, and increasing transmit power.

The highest order modulation reported in a commercial microwave product is 1024QAM. Due to the DAC/ADC speed and resolution, as well as phase noise issues, such higher order modulation is possible only for narrow bandwidth (below 100 MHz) systems.

For wider bandwidth systems in mm-wave bands, the highest modulation so far reported is 64QAM. Employing high order modulation such as this incurs some system penalties, such as design and implementation cost, and reduced receiver sensitivity i.e. higher signal-to-

http://doi.org/10.18080/ajtde.v3n1.2

noise ratio (SNR) requirements, as well as reduced output power due to higher linearity requirements for the power amplifier. This reduction in both transmit power and receiver sensitivity will of course result in reduced link distance.

The use of line of sight (LOS) MIMO in microwave backhaul has been recently demonstrated by Ericsson. By combining a 2x2 LOS MIMO (2 receive and 2 transmit antennas relying on path diversity) with dual polarisation and 1024 QAM modulation, a spectral efficiency of 36 bps/Hz can be achieved, yielding 1 Gbps throughput in a 28 MHz channel (Hansryd et al. 2011). This requires two radios at each end. A similar announcement has been made by MIMOtech with its Starburst Janus, an ultra-high capacity packet radio for last mile backhaul, which utilizes a 4x4 LOS MIMO yielding a spectral efficiency of 25bps/Hz (*Microwave Journal* 2013)(Cellular 4G/LTE Channel/Industry News 2013). As a promising technology for future high speed wireless communications, LOS MIMO has also been proposed by the US Defense Advanced Research Projects Agency (DARPA) in its 100G program (DARPA-BAA-13-15) to develop a 100 Gbps RF backbone using E-band mm-wave frequency spectrum.

At E-band, any quoted operational distance is complicated by its relationship with link availability. Achieving high link availabilities, of the order of 99.99% or higher as required for critical telco links, normally entails very short hops, to overcome the effect of rain and humidity that may (rarely but occasionally) extend across the entire link and prove impenetrable to the signal. Thus hop lengths in the tropics must be shorter than those in the desert to achieve the same availabilities, even though the radio systems are identical. Use of automatic transmit power control to increase transmit power to the maximum level during a rain fade can allow the link to overcome the fading effects. Other techniques include adaptive coding and modulation (ACM) and adaptive rate (AR). Such schemes change either the modulation scheme or the channel bandwidth in response to worsening path loss. For mm-wave radios which only operate with low order modulations, AR may be a better solution. By taking advantage of the large available bandwidth, AR can keep the modulation constant but reduce the transmitted symbol rate and thus the data rate. This reduces the bandwidth, and hence the noise floor to improve the signal to noise ratio during rain fades.

Increasing transmit power is the most straightforward way to increase the link distance. However, this may create both regulatory and technical issues. Although lightly licensed, depending on the country of operation E-band transmission still requires satisfying Australian ACMA or United States FCC rules (ITU Region 3 or 2 respectively) with respect to radiated power. Technically, trying to increase power amplifier output power continues to be one solution for microwave systems. For mm-wave systems, using antenna arrays with beamforming will be another potential solution, where each individual antenna element can

http://doi.org/10.18080/ajtde.v3n1.2

have low power, but as a spatially combined system of many elements the total emitted power can be high. Challenges include MMIC integration, devising an efficient digital beamforming algorithm, and compensation for mutual coupling between elements. Adaptive beamforming (Guo et al. 2010; Huang et al. 2010 : 1770-1779) and electronic steering are promising future research directions to achieve even longer range inter-aircraft, aircraft-to-base station, and aircraft-to-vehicle communications.

Conclusion

With the advance of broadband wireless access and next generation mobile systems, backhaul infrastructure is being stressed by demand for higher data rates. As cost-effective alternatives to fibre, high speed and long distance wireless backhaul is becoming increasingly attractive. However, there are significant technical challenges such as achieving higher spectral efficiency and extended transmission range.

This article has described an E-band mm-wave backhaul radio modem coupled with a high power amplifier and an automatic pointing system to enable the use of 1.2m antennas at the end of each hop. This system achieved data rates of 5 Gbps over link distances of up to 25km, and maintained communications even under severe vibration of the tower due to wind effects. The measured latency was significantly less than two microseconds. This system resulted from a collaboration between CSIRO and EM Solutions, and resulted in the world's fastest commercially available radio, with the highest link budget for the data throughput, to achieve the longest link distances, and with the lowest latency.

References

Cisco. 2014. "Cisco Visual Networking Index: Globe Mobile Data Traffic Forecast Update, 2014-2019," white paper, available at: http://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/white_paper_c11-520862.html

Dyadyuk, V; Bunton, J; Pathikulangara, J; Kendall, R; Sevimli, O; Stokes, L; Abbott, D. 2007. "A Multi-Gigabit mm-Wave Communication System with Improved Spectral Efficiency". *IEEE Transactions on Microwave Theory and Techniques*, vol. 55, no. 12, pp. 2813-2821.

Fodor, G; Dahlman, E; Mildh, G; Parkvall, S; Reider, N; Miklós, G. 2011. "Evolution of LTE toward IMT-Advanced," *IEEE Communications Magazine*, vol. 49, no. 2, February 2011, pp. 84 – 91.

Fu, IK; Chen, YS; Cheng, P; Yuk, Y. 2010. "Multicarrier Technology for 4G WiMAX System," *IEEE Communications Magazine*, vol. 48, no. 8, August 2010, pp. 50 – 58.

http://doi.org/10.18080/ajtde.v3n1.2

Guo, Y. J; Bunton, J; Dyadyuk, V; Huang, X. 2010, "Multi-Stage Hybrid Adaptive Antennas," Australian Provisional Patent, AU2009900371 (2 February 2009), PCT published (20 August 2010) WO 2010/085854 A1.

Hansryd, J.; Edstam, J; "Microwave Capacity Evolution," *Ericsson Review*, June 2011, available: http://www.ericsson.com/news/110621_microwave_capacity_evolution_244188810_c

Huang, X; Guo, Y. J; Bunton, J. 2010. "A Hybrid Adaptive Antenna Array," *IEEE Transactions on Wireless Communications*, vol. 9, no. 5, May 2010, pp. 1770-1779.

Huang, X; Murray, B; Bunton, J; Guo, Y. J; Dyadyuk, V. 2011. "Wireless Data Communications". Australian Provisional Patent (AU2009903399P) (July 2009), PCT published (WO 2011/009157 A1

Huang, X; Guo, Y. J; Zhang, J; Dyadyuk, V. 2012a "A Multi-Gigabit Microwave Backhaul," *IEEE Communications Magazine*, vol. 50, no. 3, March 2012, pp. 122-129.

Huang, X; Guo, Y. J; Zhang, J. 2012b. "Sample Rate Conversion Using B-Spline Interpolation for OFDM Based Software Defined Radios," *IEEE Transactions on Communications*, vol. 60, no. 8, August 2012, pp. 2113-2122.

Huawei. 2012. "Huawei Debuts 2nd-Generation Ultra-Broadband E-Band Microwave," Press Release, 2 October 2012; available at : http://www.huawei.com/en/about-huawei/newsroom/press-release/hw-194598-e-band microwave.htm

Microwave Journal. 2013. Cellular 4G/LTE Channel/Industry News. 2013. "MIMOtech Launches Ultra High Capacity Radio for Last Mile LTE Backhaul," *Microwave Journal*, 20 May 2013, available at: http://www.microwavejournal.com/articles/19885-mimotech-launches-ultra-high-capacity-radio-for-last-mile-lte-backhaul

Transmode. nd. "Low Latency - How low can you go?" white paper, available at: http://www.transmode/en/ resources/whitepapers

Zhang, J; Huang, X; Cantoni, T; Guo, Y. J. 2012. "Sidelobe Suppression with Orthogonal Projection for Multicarrier Systems," *IEEE Transactions on Communications*, vol. 60, no. 2, pp. 589-599.

http://doi.org/10.18080/ajtde.v3n1.2

Internet Governance: Is it finally time to drop the training wheels?

Narelle Clark
ACCAN and Internet Society

Summary

In March 2014 the US Government announced its intent to transition away from the current system of oversight of core Internet functions, and move the obligations of the Internet Assigned Numbers Authority (IANA) over to the international multi-stakeholder community. The current contract is set to expire on 30 September 2015, and thenceforth a new globalised model has the opportunity to come into being.

This article describes the current Internet governance model, and the process towards a future mode of operation.

The Internet just works, right?

Over the last few decades we have come to accept that the Internet just works. Many also assume it is truly global and operating essentially without borders or centralised control. Those of us involved in the mechanics of making it work tend also to assume that it is controlled by the many bottom-up multi-stakeholder processes that routinely deliver Internet Protocol (IP) numbers[i], domain names and protocol identifiers when needed, through the many working groups that develop policy, protocols and processes pertaining to these key identifiers. While for the most part this is true, there is, however, one entity at the heart of this system: the Internet Assigned Numbers Authority (IANA) (iana.org), operated by the Internet Corporation for Assigned Names and Numbers (ICANN) (icann.org). That entity's operation is governed by a service contract with the United States government through its National Telecommunications and Information Administration (NTIA). Via this contract with ICANN, the NTIA has explicit sign-off over the entries to IANA within the name space, and a degree of contractual influence over the numbering system (NTIA 2012). This is an overarching, or stewardship, role that operates in addition to those accountability measures embedded in the policy process and systems that exist today to create, allocate and record IP numbers, protocol identifiers and domain names.

http://doi.org/10.18080/ajtde.v3n1.3

In March 2014, the NTIA announced "its intent to transition key Internet domain name functions to the global multi-stakeholder community" (NTIA 2014). This commitment is in line with others made to the Internet community since the time when the policy was originally formalised in 1998: "While international organisations may provide specific expertise or act as advisors to the new corporation, the U.S. continues to believe… that neither national governments acting as sovereigns nor intergovernmental organisations acting as representatives of governments should participate in management of Internet names and addresses." (NTIA 1998)

The key Internet identifiers have different and separate systems of creation and allocation at the day-to-day level. For IP numbers, the Regional Internet Registries (RIRs) (nro.org) allocate from the pool of IP numbers allocated to them by IANA to Internet Service Providers (ISPs) according to the rules established within both the RIRs themselves and at the overall level by the Internet Engineering Task Force (IETF) (ietf.org). The RIRs are member organisations, and each has its own policy working groups that determine the rules governing the allocation of IP numbers in an open and transparent manner.

Protocol identifiers are created by technical working groups of the IETF that are open to anyone to participate in, with the technical merit of each determined by the collective expertise of the working group. As each technical protocol is defined, the requisite changes to IANA are stated and tracked. Once a protocol is approved, the specific parameters are entered into the IANA database; with the rapid change in technology, this can amount to thousands of entries each year.

The naming system, however, is somewhat more complex, with two distinct sets of policy processes at work. In essence, there are two types of top level domains: generic top level domains (gTLDs), such as ***.com** and ***.org**, and country code domains (ccTLDs) such as ***.au** and ***.nz**. Apart from a number of historical top level domains, the gTLDs are generally subject to the policies established by the organisations operating the registries and conform to a registrar agreement in place with ICANN. For many ccTLDs the policies pertaining to the allocation of domain names are similarly developed by multi-stakeholder processes incorporating their own government's participation, and with varying degrees of government oversight. For some country codes, however, governments control the domain absolutely, and conversely there are some which do not have any control or influence at all. In all of these circumstances, changes to the entries in IANA both for the country codes and the generic top level domains have to be signed off by the NTIA before they can be implemented.

http://doi.org/10.18080/ajtde.v3n1.3

Back to the beginnings

So how does this situation exist where a single government holds such a degree of control over what is an international resource, or indeed of resources that are arguably those of other countries? With some consideration of history, this begins to make sense. In the earliest days of the Internet, network operators comprised a variety of researchers and technologists from universities and ISPs who simply got together and agreed the various technical matters that allowed the networks to function. Initially known as the Network Working Group (NWG), this group created detailed protocol specifications and conventions which were recorded in 'Requests for Comments' (RFC). These records were simultaneously meeting notes, process instructions and technical documents. In May 1972 Jon Postel, then at UCLA, wrote RFC 349, which stated:

> *I propose that there be a czar (me?) who hands out official socket numbers for use by standard protocols. This czar should also keep track of and publish a list of those socket numbers where host specific services can be obtained.* (Postel 1972)

The RFC then went on to propose a list of initial allocations. As the protocols and networks expanded, the records pertaining to core network functions expanded similarly, and the set of Internet Protocol (IP) numbers, names and protocol identifiers expanded well beyond the notebook Jon Postel could keep in his pocket. This list eventually became the basis of the Internet Assigned Numbers Authority, and was later defined as: "the technical team making and publishing assignments of Internet protocol technical parameters" (Carpenter et al 2000).

Along with a group of Internet pioneers, Jon Postel was originally a graduate student at UCLA. In this capacity he was funded – and the IANA work performed – as an unwritten component of various US Department of Defense research projects (ICANN SSAC 2014). As the work grew, and more people became involved, the NWG evolved to become the Internet Engineering Task Force (IETF). The IETF took responsibility for the ongoing creation of RFCs and, more specifically, the allocation and procedures for assignment of IP numbers, names and protocol identifiers. Postel remained as editor-in-chief and record keeper of RFC assignments, and contributed significantly to many RFCs.

IP addresses are allocated hierarchically, with the highest level allocations recorded by IANA. Below this level, the specific addresses were recorded separately from the key protocol identifiers. The Network Information Center (NIC) was established in 1970 as "an *ad hoc* thing, with no specific directives from ARPA" (Meyer 1970) to record IP address allocations, and these remained in a stand-alone track of RFCs documenting their assignment until 1990. The day-to-day assignment of Internet numbers was officially assumed by the Defense Data

http://doi.org/10.18080/ajtde.v3n1.3

Network - Network Information Center in 1987 under a US National Science Foundation contract.

The third major component of IANA's work emerged in the form of the domain name system, largely to support the delivery of email. In this component, the familiar 'name' forms of Internet addresses were encoded, and the rules around how individual computers could be addressed by name evolved. Once again, Jon Postel was at the heart of the foundational record keeping (Postel 1982).

More players, more formality

Thus we have the three essential parts of IANA: domain names, numbers, and protocol identifiers. The records were originally kept in an *ad hoc* fashion; however the increasing reliance on Internet systems, coupled with the commercialisation of Internet services that took place throughout the 1990s, meant there was a corresponding increase in the formalisation of relationships between the parties involved and the tasks undertaken. The first part of this formalisation took place within the IETF itself, and a number of RFCs were created that codified the relationship between IANA and the IETF, as well as the various rights and responsibilities of the entities involved.

New bodies emerged, such as the Internet Society (ISOC) (isoc.org) in 1991 to promote the open development, evolution, and use of the Internet for the benefit of all people throughout the world. ISOC then chartered the Internet Architecture Board (IAB) (iab.org) which took on architectural oversight of the IETF's work, and the Internet Engineering Steering Group (IESG) (iesg.org) which took on technical management of IETF activities and the Internet standards process.

The IESG administers the process according to the rules and procedures that have been ratified by the ISOC Board of Trustees (Bradner 1996). In 1998 the IESG changed the basic form for all Internet Drafts to ensure that they contain a mandatory section under the heading "IANA considerations". In this section any required changes to the registries operated by IANA are formally stated and captured as part of the mainstream standards-making process.

The numbers on record soon exceeded what could easily be managed by IANA, and the Regional Internet Registries were formed by the IETF in 1992 under the rationale and guidelines established in RFC 1366:

> *The major reason to distribute the registration function is that the Internet serves a more diverse global population than it did at its inception. This means that registries which are located in distinct geographic areas may be better able to serve the local community in terms of language and local customs.* (Gerich 1992)

http://doi.org/10.18080/ajtde.v3n1.3

Similarly, the number of names also exceeded that manageable in a simple list file, and by 1993 were largely being registered by the company Network Solutions. In 1995 Network Solutions commenced charging for the registration of domain names, and this rapidly altered the incentives around Internet governance. In 2000, Verisign acquired Network Solutions, which at the time operated several gTLDs under agreements with ICANN – as well as the overall root server containing the top level mappings between IP addresses and domain names.

In 1997, the IANA Functions were documented within the US Department of Energy's Tera-node Network Technology contract. These functions were specified to include:

1. Parameter assignment
2. Address management
3. Domain name system supervision

In February 2000, the NTIA entered into the first IANA Functions contract (NTIA 2000) with a purpose-built entity known as the Internet Corporation for Assigned Names and Numbers (ICANN), an organisation incorporated in 1998 as a Californian not-for-profit public benefit corporation. This original contract specified many basic corporate governance requirements, as well as technical and operational requirements for the conduct of IANA. The intention at this stage was that the NTIA's role would diminish over time, and once the organisation was fully established the NTIA would withdraw completely (Chehade 2015).

The functions that comprise the IANA have evolved over time. The current set of functions, defined in the latest version of the IANA Functions contract issued in July 2012 by NTIA and performed by ICANN, consist of:

1. DNS Root Zone Management
2. Internet Numbers Registry Management
3. Protocol Parameter Registry and *.ARPA TLD Management
4. Management of *.INT

Since 2000, a series of RFCs have been created that more explicitly set out the relationship and performance standards of ICANN and IANA. Today IANA also performs additional functions on behalf of the global Internet community, such as maintenance of the Time Zone Database, but these are independent of the IANA Functions contract.

The root zone continues to be operated by Verisign under contractual arrangements with both ICANN and NTIA.

The politics of oversight

Two essential features of the Internet are those of cooperation and agreement, and this is embodied by the nature of IANA. There is no legal compulsion for equipment vendors,

http://doi.org/10.18080/ajtde.v3n1.3

Internet Service Providers, or users broadly, to use IANA; however, without the consistency and coordination it offers, the Internet would clearly not have been the success it is today. Internet protocols rely on the *uniqueness* of the many identifiers at the core of operation. For example, where an IP address is used more than once, communications to the computer or device so designated cannot be reliably delivered. Where protocol identifiers are not unique, whole communications systems break down. The entire Internet is founded on these essential principles of collaboration.

Thus it was that the documentary and administrative structures performed for and on behalf of the IETF were not formally recognised in *contractual* language until the late 1990s, as the technical approaches and methods were always determined by agreement and recorded in technical specifications – the RFCs. As a result, the IANA functions can be viewed in two ways: as services to the IETF, and as activities performed under contract (ICANN SSAC 2014). Indeed in many legal jurisdictions the Memorandum of Understanding between the IETF and ICANN would be viewed as a form of contract; and certainly within the IETF community these services are well understood and the protocols for use well and truly the norm.

Correspondingly, the RIRs possessed similar structures developed and documented over time. While IANA retains ultimate responsibility for the entire address pool, RFC 2050 (Hubbard et al 1996) recognises that RIRs operate under the consensus of their respective regional Internet communities, using open policy development frameworks (APNIC n.d). Common to the RIRs and to the IETF is that policy is discussed openly and transparently, and that decisions are taken on mailing lists in order to ensure the widest possible participation and therefore highest technical rigour.

After the establishment of ICANN, names policy formulation also drew in larger groups of people and a range of working groups and bodies were created to recognise the different constituencies of use. ICANN is today a global multi-stakeholder forum comprising commercial entities, consumers, regulators and technologists. While the ICANN Board of Directors has the ultimate authority to approve or reject policy recommendations, three Supporting Organisations are responsible for developing and making policy recommendations to the Board and four Advisory Committees advise the Board. As of mid-2013, the Governmental Advisory Committee represented 125 nations (plus the African Union Commission, European Union and the Vatican). The Country Code Names Supporting Organisation represents more than 135 country code domains (ICANN Beginner's Guide).

http://doi.org/10.18080/ajtde.v3n1.3

External factors

Outside of the systems of Internet governance outlined above, increasing external pressures have mounted for further internationalisation of key Internet identifiers and the systems that surround them. Of particular note is that of the International Telecommunications Union (ITU), which has featured Internet matters on its agenda at its policy making forums – most notably the World Conference on International Telecommunications held in 2012 (Wentworth 2012). As a treaty organisation, and part of the United Nations, the ITU operates under a strict system of membership and accreditation, with governments holding all the voting controls. A number of proposals placed before the ITU over the last few years have tried to give the ITU rights to allocate and manage the IP address space as well as other policy matters relating to the Internet's function. To date, these have been steadfastly resisted by the Internet technical and policy community, largely on the basis that existing multi-stakeholder systems of Internet governance are inherently global in nature and have led to the open platform for permissionless innovation that we have today (Arrko 2013).

The US Congress also passed legislation in November 2014 "restricting the NTIA from using Federally-appropriated dollars to relinquish stewardship during fiscal year 2015 with respect to Internet domain name system functions" (Strickling 2015a). This means the NTIA may be prevented from accessing funds to terminate the IANA contract. With the contract due to lapse on 30 September 2015 this leaves the NTIA in a practically difficult position, but also a rather politically sensitive position, with parts of the US Republican Party taking a stance opposed to the transition.

Formulating Transition Proposals

In its official announcement of the intention to transition oversight of key Internet domain name functions to the global multi-stakeholder community, the NTIA established four principles for transition:

- Support and enhance the multi-stakeholder model;
- Maintain the security, stability, and resiliency of the Internet DNS;
- Meet the needs and expectation of the global customers and partners of the IANA services; and
- Maintain the openness of the Internet.

Furthermore, the NTIA stated that it "will not accept a proposal that replaces the NTIA role with a government-led or an inter-governmental organisation solution" (NTIA 2014).

In response to the NTIA announcement, ICANN convened a process which led to the formation of a representative organisation – known as the IANA Coordination Group (ICG) (ianacg.org) – to bring together a proposal for a future mode of operation and submit it to

http://doi.org/10.18080/ajtde.v3n1.3

the NTIA. The ICG commenced work in July 2014 and comprises 30 representatives drawn from across the Internet community, including business, consumer and government representatives alongside the various technical constituencies.

In September 2014, the ICG issued a request for proposals from the three operational communities comprising the direct customers of IANA, and directed this request to the IETF, the RIRs and ICANN. In true Internet fashion, each of these groups established open working groups using their existing processes to compile a response to the request, with both the IETF (Lear & Housley 2015) and the RIRs (Ng 2015) meeting the January 15, 2015 deadline. ICANN established a separate cross-community working group to bring a consensus proposal for the names function; however, that working group has predicted that it cannot produce a proposal until June 2015 (Dickinson 2015).

In parallel with the work to produce a proposal for a future mode of operation, ICANN continues its existing work to address internal accountability issues, as these measures address essential corporate governance issues that are required by the NTIA in order to undertake a transition. This work is also expected to be completed in June 2015. The NTIA will only consider a stewardship transition proposal alongside recommendations on how ICANN's accountability can be improved.

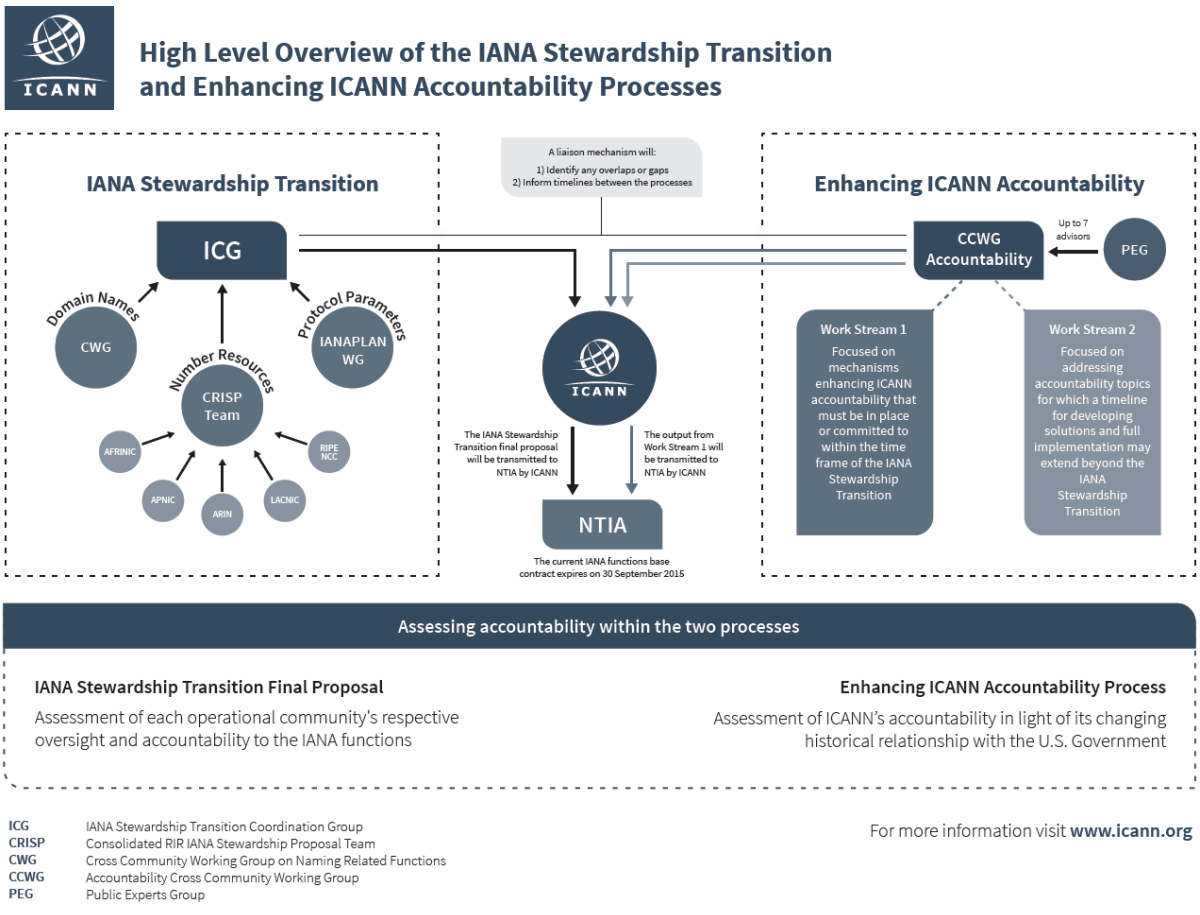

Figure 1 IANA Relationship Transition (ICANN 2015)

http://doi.org/10.18080/ajtde.v3n1.3

A significant amount of revenue is associated with the international domain name business. ICANN's own revenue for FY 2013 was well in excess of US $200m. Individual domain names can be sold for hundreds of thousands of dollars, and new top level domains introduced recently have sold for as high as US$6.7m as in the case of ***.tech**. As a result much pressure has been placed over ICANN's lifetime on its domain names policy decisions.

Hope for the future?

Both the IETF and RIR transition proposals are clear in expressing their satisfaction with ICANN as the IANA functions operator. The IETF proposal essentially describes its current mode of operation and oversight; the RIRs' was similar, but with two distinct differences – the RIR proposal calls for a new contract between ICANN and the RIRs; and for the movement of the IANA trademark and domain name iana.org to the IETF Trust, the legal entity holding intellectual property on behalf of the IETF under the auspices of the Internet Society.

In this sense, both the IETF and RIRs are clear that while there is the intention to strive to ensure the continuation of the existing systems of Internet governance, there is the case of last resort where the IANA functions could conceivably be removed from ICANN either separately or as a whole. It must be stressed at this point that this case for separation is one not intended to be entered into lightly, but only in the case of complete and systemic breakdown of the operations to the point that they cannot be remedied.

At the time of writing, a number of models were being discussed in the names community, with both internal and external options for structural separation of IANA and ICANN being mooted (Kuerbis 2015a; 2015b). These discussions are ongoing. Despite this, there does appear to be strong community support for a transition to occur within the near term (Strickling 2015b).

Should these proposals fail to converge on a single operating model before the 30 September 2015 deadline, the NTIA has the option to renew the contract for a two- or four-year period, or for some other specified period, such as a precise term within which to implement a transition. Given the increasing pressure for globalisation from other governments and a confident Internet policy community, it is likely that the US government will relinquish its singular role in the stewardship of key Internet functions in the near term. The potential for failure is also high, however, in that the US Congress may seek to politicise the transition as it moves into its next election cycle. Thus if the Internet community dithers – whether in the attempt to produce a perfect system of governance, or by failing to produce an effective model satisfying both the Internet community and the NTIA – then the opportunity may pass for at least another four years. Should the opportunity be missed then we can all expect

http://doi.org/10.18080/ajtde.v3n1.3

renewed and increasingly vigorous attempts at taking control of Internet naming, numbering and protocol identifiers by other players in the telecommunications landscape. This latter outcome can only lead to a reduction in trust and continuing pressure on the stability of the overall system.

Appendix

The Internet Ecosystem

The Internet is successful in large part due to its unique model: shared global ownership, open standards development, and freely accessible processes for technology and policy development.

The Internet's unprecedented success continues to thrive because the Internet model is open, transparent, and collaborative. The model relies on processes and products that are local, bottom-up and accessible to users around the world.

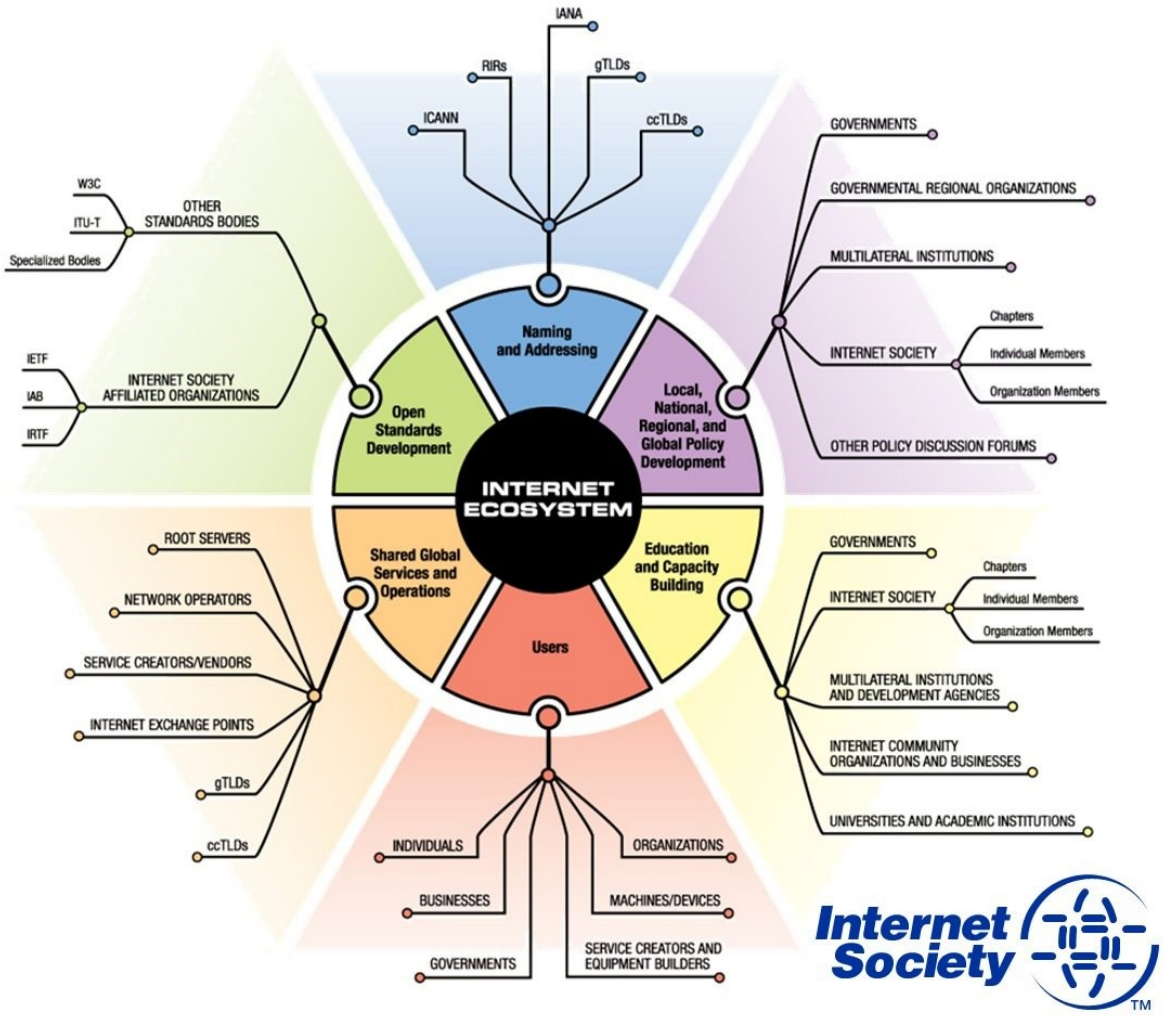

http://doi.org/10.18080/ajtde.v3n1.3

County-Code Top-Level Domains (ccTLDs) ccTLDs are operated according to local policies that are normally adapted to the country or territory involved. http://www.iana.org/domains/root/db/

Generic Top-Level Domains (gTLDs) gTLD registries operate sponsored and unsponsored generic Top-Level Domains according to ICANN policies. http://www.iana.org/domains/root/db/#

Governments Federal, State and local governments and their regulators have roles in setting policies on issues from Internet deployment to Internet usage.

Governmental Regional Organizations Governmental regional organizations include, but are not limited to, the African Union, the Asia-Pacific Economic Cooperation (APEC), the Asia-Pacific Telecommunity, the Caribbean Telecommunication Union (CTU), the Commonwealth of Nations, the European Union (EU), and the Inter-American Telecommunication Commission (CITEL). Governments sometimes like to coordinate policies related to the Internet for their regions.

Internet Architecture Board (IAB) The IAB is chartered as a committee of the Internet Engineering Task Force (IETF) and as an advisory body of the Internet Society (ISOC). Its responsibilities include architectural oversight of IETF activities, Internet Standards Process oversight and appeal, and the appointment of the RFC Editor. The IAB is also responsible for the management of the IETF protocol parameter registries. http://www.iab.org/

Internet Assigned Numbers Authority (IANA) IANA is responsible for the global coordination of the Domain Name System (DNS) Root, Internet Protocol (IP) addressing, and other Internet protocol resources. http://www.iana.org/

Internet Corporation for Assigned Names and Numbers (ICANN) ICANN is a not-for-profit public- benefit corporation that coordinates the system of unique names and numbers needed to keep the Internet secure, stable, and interoperable. It promotes competition and develops policy on the Internet's unique identifiers through its coordination role of the Internet's naming system. http://www.icann.org/

Internet Engineering Task Force (IETF) The IETF is a large, open, international community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet architecture and the smooth operation of the Internet. It is open to any interested individual. http://www.ietf.org/

Internet Community Organizations and Businesses Many Internet organizations and businesses encourage, train, and invest in Internet education and capacity building. Organizations include, but are not limited to, the RIRs, ICANN, regional and national network operators, and the Network Startup Resource Centre (NSRC), as well as vendors such as Afilias Limited, Alcatel-Lucent, Cisco, IBM, and Microsoft.

Internet Research Task Force (IRTF) The IRTF's mission is to promote research of importance to the evolution of the future Internet by creating focused, long-term, and small Research Groups working on topics related to Internet protocols, applications, architecture, and technology. http://www.irtf.org/

Internet Society (ISOC) ISOC promotes the evolution and growth of the global Internet. Through members, chapters, and partners, they are the hub of the largest international network of people and organizations that work with the Internet. http://www.internetsociety.org/

ISOC Chapters ISOC Chapters localize ISOC's core values and promote the Internet for their local communities. http://www.internetsociety.org/who-we-are/chapters

ISOC Individual Members ISOC Individual Members show commitment to ISOC's vision. http://www.internetsociety.org/who-we-are/our-members

ISOC Organization Members ISOC Organization Members support and contribute to ISOC and understand the need to take action collectively to ensure the Internet remains open, accessible, trusted, and secure. http://www.internetsociety.org/get-involved/join-community/organisations-and-corporations

International Telecommunication Union Telecommunication Standardisation Sector (ITU-T) The ITUT regularly convenes specialists drawn from industry, the public sector, and R&D entities worldwide to develop technical specifications that ensure that each piece of communications systems can interoperate seamlessly with the myriad elements that make up today's complex ICT networks and services. http://www.itu.int/ITU-T/

Internet Exchange Points (IXP) Regional and national IXPs provide physical infrastructure that allows network operators to exchange Internet traffic between their networks by means of mutual peering agreements.

Multilateral Institutions and Development Agencies Multilateral institutions include organizations that have multiple countries working in concert on Internet issues for policy development, education and capacity building. Organizations include, but are not limited to, the International Telecommunication Union (ITU), the ITU's Development Sector (ITU-D), the United Nations' UNESCO, and the World Intellectual Property Organization (WIPO).

Network Operators Network Operators include companies that provide access to the Internet. Regional Network Operator Groups (NOGs) provide collaboration and consultative opportunities for local operators and among NOGs globally.

Other Policy Discussion Forums Organizations include, but are not limited to, the Internet Governance Forum (IGF) and the Organisation for Economic Co-operation and Development (OECD), as well as national consultative forums, industry associations, and civil society organizations.

http://doi.org/10.18080/ajtde.v3n1.3

Regional Internet Registries (RIRs) RIRs oversee the allocation and registration of Internet number resources within a particular region of the world. Each RIR is a member of the Number Resource Organization (NRO). RIRs include AfriNIC, the Asia Pacific Network Information Centre (APNIC), the American Registry for Internet Numbers (ARIN), the Latin American and Caribbean Internet Addresses Registry (LACNIC) and the RIPE Network Coordination Centre. http://www.nro.net/

Root Servers DNS root name servers reliably publish the contents of one small file called a root zone file to the Internet. This file is at the apex of a hierarchical distributed database called the Domain Name System (DNS), which is used by almost all Internet applications to translate worldwide unique names like www.isoc.org into other identifiers; the web, e-mail, and other services use the DNS. http://www.root-servers.org/

Service Creators/Vendors Service Creators and Vendors provide software applications and experiences that utilize the Internet.

Specialized Standards Bodies Many organizations focus on specialized standards; some play key roles in the Internet. These organizations include, but are not limited to, the European Telecommunications Standards Institute (ETSI), the Identity Commons, the IEEE Standards Association, the ISO ANSI, the Liberty Alliance Project, Open Source Communities, and the Organization for the Advancement of Structured Information Standards (OASIS).

Universities and Academic Institutions Historically and continuing today, academic institutions play a critical role in educating students and business people. They also prototype and demonstrate hardware and software solutions that benefit the Internet.

Users People and organizations that use the Internet or provide services to others via the Internet.

World WideWeb Consortium (W3C) W3C is an international consortium where Member organizations, a full-time staff, and the public work together to develop Web standards. http://www.w3.org

The Internet Society is a non-profit organization founded in 1992 to provide leadership in Internet related standards, education, and policy. With offices in Washington, D.C., and Geneva, Switzerland, it is dedicated to ensuring the open development, evolution, and use of the Internet for the benefit of people throughout the world. More information is available at: http://InternetSociety.org.

1775 Wiehle Avenue, Suite 201, Reston, VA 20190-5108, USA Galerie Jean-Malbuisson 15, CH-1204 Genèva, Switzerland +1 703 439 2120 +41 22 807 1444

http://doi.org/10.18080/ajtde.v3n1.3

References

APNIC. n.d The RIR system, The Open Policy Framework. Retrieved from https://www.apnic.net/about-APNIC/organization/history-of-apnic/history-of-the-internet2/4

Arrko, J. 2013. Permissionless Innovation. Retrieved from http://www.ietf.org/blog/2013/05/permissionless-innovation/

Bradner, S. 1996. RFC 2026/BCP 9 The Internet Standards Process, Retrieved from http://tools.ietf.org/html/rfc2026.txt

Carpenter, B; Baker, F; Roberts, M. 2000. RFC 2860 Memorandum of Understanding Concerning the Technical Work of the Internet Assigned Numbers Authority, p2. Retrieved from http://tools.ietf.org/html/rfc2860

Chehade, F. 2015. ICANN Testimony before the U.S. Senate Committee on Commerce, Science, and Transportation Hearing: Preserving the Multistakeholder Model of Internet Governance Wednesday, February 25, 2015

Dickinson, S. 2015. Update on IANA Stewardship Discussion - 24 February 2015. Retrieved from https://www.icann.org/news/blog/update-on-iana-stewardship-discussion-24-february-2015

Gerich, E. 1992. RFC 1366 Guidelines for Management of IP Address Space, p2. Retrieved from http://tools.ietf.org/html/rfc1366.txt

Hubbard, K; Kosters, M; Conrad, D; Karrenberg, D; Postel, J. 1996. RFC 2050 Internet Registry IP Allocation Guidelines. Retrieved from http://tools.ietf.org/html/rfc2050.txt

ICANN. n.d. Beginner's Guide to Participating in ICANN. Retrieved from https://www.icann.org/en/about/learning/beginners-guides/participating-08nov13-en

ICANN. 2014. Security and Stability Advisory Committee (SSAC), 2014, SAC067 Overview and History of the IANA Functions. Retrieved from https://www.icann.org/en/system/files/files/sac-067-en.pdf

ICANN. 2015. IANA Stewardship Transition & Enhancing ICANN Accountability | Presentation, p2. Retrieved from http://singapore52.icann.org/en/schedule/sun-iana-stewardship-accountability/presentation-iana-transition-accountability-08feb15-en

http://doi.org/10.18080/ajtde.v3n1.3

Kuerbis, B. 2015a. The last third: Why the IANA transition for names is hard. Internet Governance Project iSchool @ Syracuse University Syracuse, NY USA 13244. Retrieved from http://www.internetgovernance.org/2015/02/10/the-last-third-why-the-iana-transition-for-names-is-hard/

Kuerbis, B. 2015b. The economic case for vertical separation of IANA. Internet Governance Project iSchool @ Syracuse University Syracuse, NY USA 13244. Retrieved from http://www.internetgovernance.org/2015/03/09/the-economic-case-for-vertical-separation-of-iana/

Lear, E; Housley, R 2015. Draft Response to the Internet Coordination Group Request for Proposals on the IANA protocol parameters registries Retrieved from http://datatracker..ietf.org/doc/draft-ietf-ianaplan-icg-response/

Meyer, E. 1970. Network Meeting Notes. Retrieved from http://www.rfc-editor.org/rfc/rfc82.txt.

Ng, C. 2015. CRISP submits final response to ICG. Retrieved from http://blog.apnic.net/2015/01/16/crisp-submits-final-response-to-icg/

NTIA. 1998. Statement of Policy on the Management of Internet Names and Addresses June 05, 1998, Docket Number: 980212036-8146-02. Retrieved from http://www.ntia.doc.gov/federal-register-notice/1998/statement-policy-management-internet-names-and-addresses

NTIA. 2000. NTIA PO: 40SBNT067020 IANA contract 2000. Retrieved from http://www.ntia.doc.gov/files/ntia/publications/ianacontract.pdf

NTIA. 2012, NTIA Contract: SA 1301-12-CN-0035. Retrieved from www.ntia.doc.gov/files/ntia/publications/sf_26_pg_1-2-final_award_and_sacs.pdf

NTIA. 2014. NTIA Announces Intent to Transition Key Internet Domain Name Functions. Retrieved from http://www.ntia.doc.gov/press-release/2014/ntia-announces-intent-transition-key-internet-domain-name-functions

Postel, J. 1972. RFC 349 Proposed Standard Socket Numbers, p1. Retrieved from http://tools.ietf.org/html/rfc349.txt

Postel, J. 1982. RFC 805 Computer Mail Meeting Notes. Retrieved from http://tools.ietf.org/html/rfc805.txt

http://doi.org/10.18080/ajtde.v3n1.3

Strickling, L. 2015a. Remarks by Assistant Secretary Strickling at the State of the Net Conference 1/27/2015 http://www.ntia.doc.gov/speechtestimony/2015/remarks-assistant-secretary-strickling-state-net-conference-1272015

Strickling, L. 2015b. Testimony of Assistant Secretary Strickling before the Senate Committee on Commerce, Science, and Transportation on "Preserving the Multistakeholder Model of Internet Governance". Retrieved from http://www.ntia.doc.gov/speechtestimony/2015/testimony-assistant-secretary-strickling-senate-committee-commerce-science-and-

Wentworth, S. 2012. Updates from WCIT. Retrieved from http://www.internetsociety.org/updates-wcit

Further information

IANA Coordination Group http://ianacg.org
IANA Protocol Registries https://www.iana.org/protocols
Internet Corporation for Assigned Names and Numbers http://icann.org
Internet Assigned Numbers Authority http://iana.org
Internet Engineering Task Force http://ietf.org
IETF Overview of RFC Document Series https://www.rfc-editor.org/RFCoverview.html
Number Resource Organization http://nro.net
Internet Society http://internetsociety.org
Internet Architecture Board http://iab.org/

---

[i] For ease of use, both IP addresses and autonomous system numbers are referred to herein as IP numbers.

http://doi.org/10.18080/ajtde.v3n1.3

Metadata Retention and the Internet

Geoff Huston
Asia Pacific Network Information Centre

Summary: The Metadata retention measures being considered in Australia make some sweeping assumptions about the semantics of IP addresses and their association with individual subscribers to the Internet. But are these assumptions warranted? The exhaustion of the free pool of IPv4 addresses has prompted a new generation of Internet services that treat IP addresses as ephemeral shared conversation tokens, and retaining address-use metadata in such an environment is an exercise in futility. The regulatory environment persists in treating the Internet in the same manner as the telephone network and as a network-centric service utility, while the revolutionary change that the Internet brought to the communications environment was to reverse the roles of network and attached device, and form a device-centric focus to communications. Unless our regulators can grasp the implications of this essential architectural change we will continue to see misplaced and ultimately futile regulatory measures imposed on the Internet, to the ultimate cost of the consumer.

> *A police officer on his beat late at night sees a drunken man intently searching the ground near a lamp post and asks him the goal of his quest. The inebriate replies that he is looking for his car keys, and the officer helps for a few minutes without success. He then asks whether the man is certain that he dropped the keys near the lamppost.*
>
> *"No," is the reply, "I lost the keys somewhere across the street." "Why look here?" asks the surprised and irritated officer. "The light is much better here," the intoxicated man responds with aplomb.*

I can't help but think that the situation in this rather old joke applies very precisely to the current Australian efforts to compel network operators, through some contemplated regulatory instrument, to record and retain network-collected data about their customers' online activities.

What I'd like to examine here is the emerging picture that while networks, and network operators, make convenient targets for such surveillance efforts, the reality of today's IP environments are far more complex. With various forms of overlays, tunnels, packet header transforms, proxies and application level middleware, the supposition that the source IP

http://doi.org/10.18080/ajtde.v3n1.4

address in an outbound IP packet is a rigid pointer to an individual customer is a triumph of wishful thinking. The information available in the packet headers of the transport IP packets bear less and less relation to persistent endpoint identities. Internet Access carriage network operators are increasingly ignorant about what their customers are doing. The result is that it is now quite common for Internet networks not to have the information that government agencies are seeking. Not only can moderately well-informed users hide their activities from their local network, but increasingly this has been taken out of the hands of users, as the applications we have on our smartphones, tablets and other devices are increasingly making use of the network in ways that are completely opaque to the network provider and, indeed, to the human user. Looking to the network and network operators for this stream of data about connected users and the IP addresses that they use is increasingly an exercise that appears to fall into the category of "security pantomime".

"Security pantomime" is a term I've seen applied in a number of security-related exercises where, like a pantomime, there is the superficial appearance of security, whereas in fact it's a parody of the real thing, and it's obviously ineffectual in achieving its supposed objectives.

The Australian Government's proposal to introduce legislation for metadata retention was reported in the *Daily Telegraph* on 5 August 2014, and was announced by the Prime Minister that same day in a press conference. The Minister for Communications, Malcolm Turnbull, reportedly did not know of the proposal until he read it in the newspaper, and is said to have asked in Cabinet whether his colleagues understood what metadata was.

In an interview two days later, the Attorney-General, Senator George Brandis, was unable to clearly explain the nature and scope of the data ISPs would be required to retain. He explained this proposal as retaining the domain names of sites visited by users, but not the contents of any session. But it was pointed out that the only way that an ISP would harvest domain names from network traffic would be to inspect the content traffic flows and pick out the domain names from the content stream. Time to call in the Minister of Communications, who explained that "metadata" meant that the ISPs would not be collecting domain names, but instead would be retaining a record of IP addresses used by clients. The attempts to explain these measures were no clearer at the end of the week than at the start. ISPs were going to be required to collect some of this "metadata" stuff, but no politician could give a clear and coherent view of precisely what this data actually was.

Time to call in more folk to try and explain, and David Irvine, the Director General of Australia's security organisation, ASIO, and Andrew Colvin, the Australian Federal Police's Deputy Commissioner for National Security, fronted the media after the government had bungled its attempt to explain the technical details behind its proposed legislation to force internet service providers to store non-content data for two years to aid law enforcement.

http://doi.org/10.18080/ajtde.v3n1.4

They confirmed they were after source IP addresses as opposed to destination IP addresses, namely the addresses associated with the web pages and services users are connecting to, alongside what was termed "non-content call data".

"We have been accessing that data for many years legally, and all that we are changing actually, or seeking to change, is that the data – which is held by the companies for a commercial purpose for billing or other reasons – be held in such a way that we can continue to have access to it in an environment where that access has begun to diminish a little bit," Irvine said. Both clarified that law enforcement and security intelligence agencies were only legally allowed to access source IP addresses under non-warranted metadata requests.

Let's take a step back and consider the question: Why are some Internet users apparently so concerned about measures to capture and retain network "metadata" on the Internet? It appears that they believe that the network is in a position to be privy to all our communications, both in terms of who we contact and what we do and say online, and were this information to be circulated in an uncontrolled manner, or even if there is a risk of uncontrolled disclosure, this level of data capture and retention comes uncomfortably close to sensitive aspects of the erosion of personal privacy. There are related concerns about agencies acting outside of their statutory powers and without due regard for legal process.

The starting position of this metadata meta-conversation is that it appears to be a commonly accepted truth that the network, and the network operator, is privy to the complete details of all our communications activities.

Why do we think that? And is it really the case?

It's often useful to compare the Internet to traditional telephony. Not only are folk more confident that they understand the basic concepts of telephony, but also it's often the case that when we talk about the Internet, we unconsciously borrow terms and concepts from telephony.

So I'll take a quick excursion into the operation of the traditional telephone network. Feel free to move on if you've known all this for the past forty years or so!

When you make a "call" in a traditional switched telephone network there is an initial exchange of data within what is known as the "control plane". The network takes the dialled number and maps that into a terminating location. A call request control message is sent to this location and the ensuing exchange of messages sets up a virtual circuit in the data plane of the network that links the two telephone endpoints. Termination of the call generates a further exchange of control messages to tear down this virtual circuit.

If we could look at a telephone network's control data relating to call establishment and

http://doi.org/10.18080/ajtde.v3n1.4

teardown we would have a form of “metadata” about telephony: a record of who is dialling whom and when; whether the call was answered; and, if so, how long the call lasted. But of course this control “metadata” is all about telephone numbers. If we combine this data with current telephone directory information then the combination of the two data sets can be transformed into a “metadata” log of who is talking to whom, when, and for how long. This per call control data does not track what is said, nor can it. It just tracks who is talking to whom. Even in mobile telephony networks, the network’s ability to track the location of handsets through triangulation of base station data can be used to provide the same metadata.

Telephone service operators have been collecting this per-call “metadata" for decades. Most telephone tariffs were based on a rental fee and a usage component, where the usage component was based on who you called, and how long each call lasted. Subscribers could get a copy of this call record data if they had a dispute over their bill, and more recently the telephone companies turned this call log information into a “feature” by generating itemised logs on your bill. Not only does it providing you with a comprehensive log of all the telephone calls you made in each billing period, it also illustrates that the telephone company is privy to all of your call behaviour on the telephone network. The public telephone directory maps my name to a telephone number, and the telephone company’s call records detail all of the calls I made or received. The two combined form this rich stream of metadata information about who is talking to whom, where and when. Who else is privy to this metadata? The directory is public. The telephone company retains the call information, and no doubt under the terms of an appropriate regulatory instrument certain third parties can also gain access to this information.

This is not a new development for telephony. It's been the case for decades. So if the collection and retention of per-call data in the telephone network has been a matter of deep and abiding concern on the part of the users of the telephone network, then they’ve not exactly been highly vocal on the topic. I suspect that we’ve all largely grown up with a telephone network that we know generates and retains this data, and we appear to accept that.

Now let’s switch context back to the Internet.

Given that the traditional telephone network of yesteryear generated a rich stream of metadata about each individual call that is made across the network, why can’t we just look for the same information in the Internet? Aren’t these both instances of large scale public communications systems with very similar properties? Indeed, many of the old telephone companies now also operate an Internet service, so they can’t be that different. My mobile phone is also a mobile Internet platform. It must be the same under the hood. So if we are on

http://doi.org/10.18080/ajtde.v3n1.4

the hunt for Internet metadata, why can't we use the same information, the same procedures, the same regulatory framework on the Internet that's we've used on the telephone network?

Or is this line of thought a case of looking for Internet metadata under a convenient lamppost, when in actual fact the network may not be privy to this form of metadata in the first case. End users can trivially hide their identifying details from the local access network, and sometimes they do so on their own and sometimes the network itself uses equipment that destroys any lasting semantic interpretation of an IP address.

There is a common architectural theme behind all of these questions. This common theme is that, architecturally, the Internet is not telephony. It's almost precisely the reverse. The Internet inverted the telephone model.

The telephone network consists of a "smart" network and dumb devices. Without the telephone network, a traditional telephone handset is a lump of useless plastic and metal. In contrast, the Internet is populated with computers and various incarnations of "smart devices". These connected devices make no critical demands of the network's functionality. The Internet network can be "dumb" and the connected devices will use this dumb network without any problem whatsoever. In the Internet model, these end devices generate datagrams and hand them into the network. The network can deliver these datagrams to their intended destination, or it may drop them. It can re-order them, and it can slice and dice them in order to squeeze through narrow network crevices. In this datagram delivery network every packet is an adventure. Normally such a level of unreliability and variability in the network service model would be inadequate to support a useful set of applications and services. But it's the responsibility of the network protocol software in the connected end devices to take these arriving datagrams and reassemble the original information stream. In the Internet it's the "handsets" and the applications that they run that form the true network, not the interior of the network's pipes and switches.

This reversal of roles between network and attached device has a profound consequence. Within the network there is no control data plane that establishes and tears down individual "calls". Within the network there is no "call" at all. When one computer establishes a connection to another computer, then within the framework of the original architecture of the Internet, this connection is a shared state between the two edge computers, and the network not only is not privy to this shared state, it has no need to be privy to this shared state. From this perspective networks do not collect call logs because there is no network-level control plane that causes the network to establish and tear down "calls" in the first place.

http://doi.org/10.18080/ajtde.v3n1.4

Oddly enough, this role reversal has not resulted in a simpler picture of networks. The result has been a more complex view of networks and even more complex application behaviours. This is now reaching the stage when we can question what exactly is the role of an IP address in today's Internet.

Historically, an IP address was a relatively stable endpoint identifier token: "Everybody who wants to converse with me pushes packets into the network that are addressed to my IP address." These days we are seeing networks whose use of an IP address is limited to a conversation endpoint token without any connotations of permanence or even uniqueness, and in response we are seeing applications whose use of an address is even more ephemeral, and IP addresses are redefined as an ephemeral transient endpoint token, whose lifetime does not even extend across the lifetime of the conversation, nor is the use of the IP address even purported to be unique at any given time.

If the hunt for Internet metadata is a hunt for the stable associations of end users with IP addresses, then, as we head down this path of redefining what is an IP address, wherever this metadata might be found, looking inside the network, as is proposed by this metadata retention proposal, would be a poor place to start.

How did we get to this rather curious situation? Let's look at a few technologies that illustrate some of the layers of complexity in current Internet networks.

Carrier Grade NATs

For many years the issue of forthcoming IP address exhaustion was been something that Internet Service Providers (ISPs) pushed over the wall to become a customer problem rather than addressing it as a network problem. ISPs assigned their customers a single public IP address, and left it to the customer to incorporate an edge network address translator (NAT) in their modem or similar to permit multiple devices in the home or office to share this single public address. But as the IP address shortfall pressures increased in intensity there were situations where even one unique IP address per connected customer would require more IP addresses than the ISP actually had. The response has been to deploy NATs in the interior of the network as carrier grade NATs (CGNs). This allows a number of customers (who themselves may have NATs as well) to share a smaller pool of unique public addresses.

Each time an interior device initiates an external connection attempt, the outbound packet is intercepted by the CGN, and an available external port number and IP address pair is pulled from the CGN's managed pool and assigned to this connection. In terms of IP address use, a single public IP address may be used by many users simultaneously. The efficiency of the CGN in terms of maximising the efficiency of use each public IP address is improved by increasing the pool of interior customers behind each CGN. In other words, there are

http://doi.org/10.18080/ajtde.v3n1.4

advantages to the network operator to configure the service network to use a small number of large CGNs, allowing the CGN operator to maximise the efficiency of this form of address sharing. The implication of this form of configuration is that an arbitrarily large number of end users would be using the same IP address within a given time window.

From the outside, users positioned behind a CGN assume the IP address of the CGN itself. And because the same address may be used by multiple users' connections simultaneously, you need to use additional lookup keys to differentiate one user from the other. For low efficiency CGNs the source port and protocol field is sufficient. But if really high levels of address efficiency are required, CGN binding software may be configurable to operate in a "5-tuple binding" mode, where simultaneous connections to different external services can use identical source-side IP addresses and port numbers. In this situation the only way to disambiguate individual connections in the CGN log is to log both source and destination IP addresses, the source and destination port addresses, the transport protocol, and a consistent, accurate and finely granulated time stamp. But nobody collects that massive volume of information. In other words the keys to unlock the nature of the address transforms that are being performed by these CGN devices are not being collected and certainly not being stored.

What CGNs illustrate is that the Internet, unlike the telephone network, is highly heterogeneous. One class of endpoints have a stable well known IP address. This class of endpoints can initiate connections and receive connection requests. These endpoints are often not human users at all, but are more typically used by content servers. Web servers traditionally need a stable IP address, as do mail servers, cloud servers, and similar. It should be noted that even this is no longer always the case, and certain forms of cloud-based services have been observed to map a service name into different IP addresses depending on the assumed location of the client who is making the DNS query. Another class of endpoints live behind NATs. These endpoints no not have a permanent IP address and cannot receive incoming connection requests. When they initiate connections, the IP address they use to represent their own identity to the network depends on the local configuration of the network. If may be that the IP address is a stable address that is used across all connections from this user. With CGN on the path the picture changes radically, and the IP address that is used to identify the client is in fact just part of a vector that is used to identify the appropriate CGN-binding table entry. In this context the IP address does not relate to the user, and as the user makes further connections there is no assurance that the user will be assigned the same IP address. When CGNs are in play, knowledge of a single IP address may well be equivalent to knowledge of a single city, or a single region, rather than the desired knowledge of an individual.

http://doi.org/10.18080/ajtde.v3n1.4

Tunnelling

Tunnelling uses a different form of packet transformation, where, in its simplest form, a datagram becomes the payload of an enclosing datagram. This wrapping of one IP datagram in another is performed upon tunnel ingress, and the complementary unwrapping is performed at tunnel egress. When outside the tunnel the packet exposes its source and destination address to the network, but within the tunnel the source and destination addresses of the IP packet are the addresses of the tunnel ingress and egress points. The actual IP packet is carried as a payload of the tunnel packet. This hiding of the inner IP packet from the network can be further strengthened by encrypting the tunnel packet payload within the tunnel, using a cypher that is shared between the tunnel ingress and egress points. None of the activity within the tunnel is visible to the network that carries the tunnel traffic.

Tunnels can assume arbitrary forms. HTTPS proxy tunnels embed IP packets in the payload of a secure transport session that appears to be a secure web transaction. The rationale for this form of tunnelling is that in certain firewall configurations about the only packets that can be passed through the firewall without being hopelessly mangled are packets that are part of a secure web flow. Because the payloads of these packets are encrypted, placing the original IP packet into this stream as a payload is a useful technique. I've heard of IP buried in the DNS, and no doubt it's possible to embed IP into many application level protocols, but at some point the exercise becomes one of flag planting to prove that such convolutions are possible as distinct from providing a solution to a real world problem.

VPNs

Tunnelling is used in many forms of host based Virtual Private Network (VPN) services, where one end of the tunnel is the user's device, and all of the device's traffic is passed into the tunnel within the device. The device's interaction with the local network is limited to passing encrypted traffic to and from the tunnel egress point. The visible local IP address of the end device doesn't communicate with anyone other than the nominated tunnel egress. At the other end of a connection, the remote service receives a packet whose source address is associated with the tunnel egress point. In this form of tunnel configuration there is no point where the local IP address of the device and the remote address of the server are exposed together in the same packet header.

VPN + NATs

The VPN and NAT functionality can be combined by using a NAT at the common tunnel egress. The VPN client is provided with a private address by the VPN provider, and a secured

http://doi.org/10.18080/ajtde.v3n1.4

tunnel is set up between the client device and the VPN provider. When an outbound packet arrives at the tunnel egress point, a local NAT is used on the inner IP packet to transform the packet's source address to the local IP address of the NAT. Inbound packets arrive at the NAT, and a transformation is applied to the destination address fields, and the packet is then passed into the tunnel for transmission to the VPN client. Here again there is no point in the transmission path where there is an IP packet whose packet header contains the address of both the VPN client and the address of the remote service.

TOR

Into this combination of IP-in-IP tunnelling, payload encryption, VPNs and NATs we can add relaying, and the result is the TOR network. TOR is a form of relayed tunnelling where a packet is passed from relay to relay through tunnels, and at no point is there a clear text IP packet that contains a packet header with the actual IP addresses of the two parties who are communicating via TOR. TOR is an asymmetric protocol, in that both sides of a TOR conversation do not need to be TOR-aware, allowing a TOR client to connect to any conventional IP service. The service receives an IP packet whose source address is that of the TOR exit router, and its response is passed to this TOR point, which then applies a TOR tunnel wrapper and sends it back along a tunnel relay path to the original TOR client. One can go further with TOR and wrap it in WebSocket format and bounce the traffic though short-lived JavaScript proxies on browsers to further disguise the TOR traffic pattern so as to emulate conventional secure web transactions, for example, so that even traffic profiling scanners would be unable to distinguish a TOR tunnel from other unrelated traffic.

V6 Transition

Interestingly, the IPv6 transition has also become enmeshed into this story of increasing network complexity in their treatment of IP addresses in packet headers. As originally envisaged, network operators would progressively deploy IPv6 across their infrastructure in addition to IPv4, in a so-called dual stack configuration. End users would need to wait for their service provider before they could connect using IPv6. To some extent impatience took over, and it was not long before we saw various tunnel approaches used to connect "islands" of IPv6 across the oceans of an IPv4 substrate. Some of these used IP-in-IP (such as 6to4) while others used IP-in-UDP-in-IP (such as Teredo) in order to traverse NATs. These approaches have been taken up by network service providers who are attempting to address the twin issues of IPv4 address depletion and IPv6 deployment simultaneously.

Some approaches use a single protocol IPv4 network, and tunnel IPv6 using this infrastructure, such as 6RD, while other approaches use an IPv6 common substrate and

http://doi.org/10.18080/ajtde.v3n1.4

tunnel IPv4. The tunnelling can take the form of one of the various permutations of IP-in-IP, or, if the substrate is IPv6, the approach can use protocol translation and embed the original IPv4 addresses in the interface identifier part of the IPv6 packet headers. This translation can be mapped, so the transform is stateless, or it can be performed using a NAT-like function, using a stateful translation. There are now many approaches to IPv6 transition, and ISPs appear to be customising their choices from this selection rather than adopting a uniform approach.

So in a network that performs one of these transition mechanisms, what does an IP address signify? Is it an endpoint identity or an address of a translating unit? Is the address a synthesised compound of one or more IPv4 addresses embedded into the IPv6 address (such as Teredo or 6to4 and their form of embedding V4 addresses into a V6 address) or are the IPv4 addresses contained in the payload of the outer packet?

In this hybrid environment there are no clear and consistent outcomes. IP addresses can take on many forms and the address itself typically provides no clue as to its role. It may be a stable endpoint identity, or it may be an address which has significance only within a particular scope or context. It may not be a stable address, and may only have significance in conjunction with other fields in the IP packet header. It may be an ephemeral conversation token, or it may be just a one-off translation table index. It's only when you understand the context of the address, and understand the form and function of the units that have to manipulate the address that you can place any meaningful interpretation on what an IP address signifies.

V6

There is a point of view that many of these transitional complexities are due to the combination of the difficulties in this transition to IPv6 and the exhaustion of IPv4 addresses. The claim is that if we were operating an all-ipv6 network it would all be far simpler, we could dispose of these cumbersome and complex NATs, remove all this transitional complexity and IP addresses would once more become stable endpoint identifiers.

But that’s not going to happen.

IPv6 now has taken on so-called “privacy addresses” with a passion. Devices that use this function periodically generate a new 64-bit interface identifier based on some form of random number identification, and generate a new local synthetic address combining a common 64 bits of the network identifier part with a self-selected random lower 64 bits. In this world of privacy addressing IPv6 addresses have a dual personality. The upper 64 bits are a stable endpoint identifier that identifies a local network. But the lower 64 bits are again

http://doi.org/10.18080/ajtde.v3n1.4

ephemeral, and have no permanent existence. IPv6 hosts can use these privacy addresses for outbound connections, but cannot use them as stable connection points for inbound connection attempts.

In IPv4 the concept of private addresses and NATs has been around for a very long time. IPv6 has never been clear about its position on these concepts. Initially, IPv6 had no such concept. Local addresses were simply global addresses that were exclusively used in some local scope and were unreachable from elsewhere, akin to the original concept in IPv4. However, we subsequently saw the reservation of a very large block of IPv6 addresses to be used exclusively in local addressing scoped contexts, where the remainder of the network part of the address was locally defined. These Unique Local Address prefixes (ULAs) are now being used in the same manner as private addresses in the IPv4 realm. NATs also exist in IPv6, and they appear to come in a couple of flavours: NAT66 performs a NAT translation on the entire 128 bit address field, in a manner entirely analogous to a NAT in IPv4. NPTv6 is a variant of this function that performs an address transform only on the address prefix, leaving the remainder of the IPv6 address field unaltered.

Other forms of address transforms are encompassed by the SHIM6 work. This leverages the observation that in IPv6 it is possible for a device to have a number of IP addresses simultaneously. There are numerous motivations for this, including explicit address scoping as part of a site security framework, but the explicit area SHIM6 was addressing was the practice of edge networks "multi-homing" with multiple upstream transit service providers for greater resiliency. In IPv4 this is undertaken by using so-called provider independent space and asserting a unique entry in the global routing table. While this is possible in IPv6, there was the desire to be a bit kinder to the routing table, and see if there was a solution which used provider-based addressing from each of the upstream providers and allowed individual TCP sessions (indeed allowed any IP end-to-end context to be address agile without disturbing the upper level end-to-end transport sessions. One way to look at SHIM6 is that it embeds a NPTv6 function right into the host's IPv6 protocol stack, and via a dialogue with its counterpart at the other end of the connection, allows the local host to switch provider prefixes at will, and thereby move the transport session from one provider to another without disturbing the transport session. At a network level this would produce some peculiar outcomes. TCP sessions could simply appear within a network without any opening context of a SYN packet exchange, and equally they could disappear without any visible shutdown. Pretty clearly, a combination of SHIM6 and CGNs in IPv6 would be mutually incompatible. In the case of SHIM6, the IP address visible to the network is not necessarily usable without the earlier establishment of a SHIM6 context.

There are other aspects of IPv6 address management with a similar flavour. What they all

http://doi.org/10.18080/ajtde.v3n1.4

illustrate is that in IPv6 an IP address is not necessarily a stable end point identifier. Whatever simplicity is being regained in an all-IPv6 network, a consistent and simple semantic interpretation of all IPv6 addresses is not part of the package being offered.

MultiPath TCP

As a final illustration of the level of complexity that we see in the Internet today I'd like to highlight MultiPath TCP. In devices with two or more interfaces the protocol in the local device uses a local rule to determine which interface to use to send an outbound packet, and this choice is "sticky" at least for the level of granularity of a TCP connection. TCP cannot move from one connection to another, or even use multiple connections in parallel. Until MultiPath came along. With MultiPath TCP the local protocol stack is able to open up multiple TCP connections to the remote endpoint, and then distribute the application payload data stream across these multiple connections.

All this sounds a bit esoteric until you look closely at the device in your pocket. Often it has a WiFi interface and a cellular data interface. Normally there is a local configuration rule that says "when the WiFi interface is usable stop using cellular data". When cellular data was a highly priced limited service and WiFi was an abundant service with no marginal use tariff, this kind of local preference rule made a lot of sense. But increasingly we see "unlimited" market offerings for cellular data, and the speeds of cellular data are rising to a level that is comparable to many Wi-Fi services. If we had MultiPath TCP we could open up a connection on each interface and use both at once. The faster of the two would deliver more data, and we would optimise the speed of the overall transaction as a result. If the device in your pocket is an Apple device with a recent version of iOS, then chances are that when you use the Siri application the application will attempt to use MultiPath TCP if it can. All this may sound esoteric, but there are hundreds of millions of these particular devices out there, and as a result MultiPath TCP is an esoteric technology with a user base that numbers in the hundreds of millions! Any other technology would call itself "mainstream" based on those deployment metrics!

What does MultiPath TCP mean for the network? On any individual network, the network sees only fragments of the full data exchange. The complete content flow is being passed across multiple networks at once using multiple IP address pairs. It's possible that no single network carries the complete conversation flow, and even if you could look at the packet flows within each of these networks, the binding glue that identifies each sub-stream as part of a common TCP MultiPath stream is not necessarily visible to the network: that essential context information is held as state on the two end devices.

http://doi.org/10.18080/ajtde.v3n1.4

Where's Metadata?

There is no doubt that Internet service networks generate large volumes of data about the network, or "metadata". This data is used as aggregate data to operate the network, and used for network planning. It often drives the network's security subsystems. It's used for authentication and accounting of the network's customers. It's used to enforce various network policies and operational practices.

But does this torrent of metadata about how the network operates provide an insightful window on the identity and actions of its end customers? Do IP networks have the ability to generate the equivalent of the telephone call log?

Increasingly, the answer is "no", and the more we are provided with information on various surveillance efforts to use the network to extract such information, the response from end user systems and applications they use is even more in the direction where the users' actions are kept hidden from the network.

If you really want to understand what is happening in an end-to-end network, the best place to do so is at either end; and the best way to do so is within the application. But we have a rich regulatory history, gathered in the postal, telegraph and telephone industries of using the network as the point of regulatory control. We have various legislative frameworks that govern the way in which public telecommunications services can operate. We recycled this entire legacy framework when the Internet came along, and these days the locus of attention in terms of regulation remains the network. But the Internet does not use a network-centric architecture. As we've explored here, there are many ways in which users, and the applications that users run, can still function perfectly normally yet still hide their essential aspects of their communications from the network itself. But applications and servers are not conventionally subject to public regulation. What happens on a web site is the business of the web site's owner and operator and there are few, if any, regulatory constraints in the way in which the web site operates. Our response has been through generic measures in terms of consumer protection and privacy protective measures, but these measures are generic and often poorly applied in this context.

Looking within the network to try and piece together exactly what is happening end-to-end is a guessing game that is increasingly becoming a rather expensive and futile exercise. And forcing network operators to collect and retain their data for arbitrarily long period of time does not restore any form of rationality to the exercise. In the end all that is being collected here in the form of network metadata is just a truly massive pile of useless bits.

So when visible action is called for, and various agencies whose charter includes aspects of national security are called to stand up and be seen to be engaged with the national security

http://doi.org/10.18080/ajtde.v3n1.4

agenda, then it is perhaps no surprise that the network is the focus of their attention. Not because there is a rich vein of readily accessible information lurking there. Far from it. It's just that the network is a convenient, well lit and well established lamppost. But the information that they are searching for – that information truly is located elsewhere.

Glossary of terms

CGN Carrier-grade NAT, also known as large-scale NAT (LSN)

DNS The domain name system (DNS) is the way that Internet domain names are located and translated into Internet Protocol addresses. A domain name is a meaningful and easy-to-remember "handle" for an Internet address.

IP Internet Protocol

ISP Internet Service Provider

NAT Network Address Translation, an internet technology that translates the addresses of packets as they pass across the NAT. It enables a network to be locally numbered from addresses drawn from a private number space, but whenever such a privately addressed device communicates with the outside world the packets passed out to the public network will have the private addresses translated into a corresponding public address.

TCP Transmission Control Protocol. TCP is one of the main end-to-end transport protocols in TCP/IP networks. Whereas the IP protocol deals only with datagram packets, TCP enables two hosts to establish a connection and exchange streams of data. TCP guarantees delivery of data and also guarantees that packets will be delivered in the same order in which they were sent.

TOR TOR is free software for enabling anonymous communication. The name is an acronym derived from the original software project name: The Onion Router.

UDP User Datagram Protocol is a communications protocol that is similar in nature to the datagram protocol of IP itself. UDP does not offer reliable, sequenced flow control. UDP packets may be lost, reordered or even duplicated by the network. On the other hand UDP is a very lightweight protocol and is used extensively in the Internet for simple query/response services, such as the Domain Name System, or the Network Time Protocol.

ULA Unique Local Address – an IPv6 address defined in RFC 4193. It is the approximate IPv6 counterpart of the IPv4 private address. Unique local addresses are available for use in private networks, e.g. inside a single site or organisation or spanning a limited number of sites or organisations

http://doi.org/10.18080/ajtde.v3n1.4

Net Neutrality in the U.S. – 2015

Taking the Next Step

Robert Larribeau, Jr.
Retired Telecommunications Industry Analyst
San Francisco, California, USA

Summary: A U.S. Federal Court ruling in January 2014 overturned Net Neutrality rules issued in 2010 by the Federal Communications Commission (FCC), the body that regulates both the telecommunications and the cable industries in the U.S. This sparked significant support for establishing new rules to provide Net Neutrality and resulted in the submission of more than one million comments to the FCC, which broke all records. This led to the FCC adopting new Net Neutrality rules in February 2015. The FCC followed President Barack Obama's lead and classified the broadband operators as common carriers, which will require that they treat all of their customers and all content providers equally. As common carriers the broadband operators will not be able to favour one content provider over another or favour their own content services. It is very likely that these new rules will not settle the issue and will be challenged in Congress and in the courts. The Net Neutrality controversy will continue.

Introduction

I wrote an article called "A Fair Approach to Net Neutrality" that appeared in the July 2009 edition of this publication. It proposed an approach to Net Neutrality that recognised that traffic engineering is a necessary part of an Internet Network, which is required in order to optimise all of the different services carried on the Internet. It took the position that broadband operators should treat content providers on an equal basis and not use traffic engineering to favour one content provider over another or use traffic engineering to favour its own content services over those competing content providers. In this approach, for example, AT&T and Comcast could use traffic engineering to improve the performance of video streaming for all content providers but could not degrade the performance of the popular Netflix video streaming service in order to favour their own competing video services.

This approach to Net Neutrality has been gaining support. The rules that the FCC issued in 2010 are consistent with it. However, the rules that the FCC adopted in February 2015 will regulate broadband services as common carriers under Title II of the Communications Act of 1934 that has defined how the telephone services industry operated in the U.S. ever since.

http://doi.org/10.18080/ajtde.v3n1.5

Title II regulation was pushed by supporters of Net Neutrality because they believe that it would be accepted by the courts and would require the broadband operators to treat all of their customers and all the content providers equally and fairly. The broadband operators object, saying that regulation under Title II will make it much more difficult to introduce new, innovative services.

The fundamental problem with Net Neutrality is that it does not have a rigorous technical definition. Net Neutrality advocates talk about "fast lanes" and "treating all packets equally". It is not clear what this means. These concepts do not draw a clear line between acceptable traffic engineering practices and abuses that favour broadband operators' own content services or favour content from favoured providers.

Origins of Net Neutrality

In the 1990s the Internet became widely available and provided new information services such as the World Wide Web and new communication services such as emailing and instant messaging. It was revolutionary. Communication could be accomplished more quickly, more conveniently, and at little or no additional cost over the cost of a basic Internet connection. Emails could be delivered in seconds to any place in the world at no cost compared to being delivered in days at the cost of postal services. The World Wide Web provided instant access to information, provided a platform for social networking, created a soap box for ideas and opinions, and gave businesses an inexpensive way to communicate with their customers and suppliers.

The Internet also facilitated the formation of new online businesses and industries – Amazon for selling and distributing products, eBay for online auctions, and PayPal for online payments – to name just three of the most successful. These three companies became billion dollar businesses along with others including Google, Face Book, and SalesForce.com and made fundamental changes in how business is done in the U.S. and globally.

The benefits that came with the Internet were quickly and universally recognised. People liked the openness and the innovations that it brought. It was easy and inexpensive for users and businesses to connect to the Internet. In the 1990s it was also easy and inexpensive to set up an Internet Service Provider (ISP) business: all you needed was a couple of servers and a dialup communications multiplexor along with a connection to an Internet backbone provider, and many people did it.

During the 1990s most users accessed the Internet using dialup modems. These modem users relied on the ubiquitous switched voice telephone network to make a connection to their ISP. In the U.S. ISPs located their termination equipment within the users' local dialling areas to insure that there were no per-minute usage charges for Internet calls. This

http://doi.org/10.18080/ajtde.v3n1.5

approach allowed users to spend hours connected to the Internet for fixed monthly fees for their telephone and for their Internet service with no per-minute charges for communications.

The ease of setting up an ISP business, at least in a local area, helped to create a large number of ISP companies. The result was that Internet services were highly fragmented, and there were no companies that had significant market power. Even the large ISPs of that time including AOL, Prodigy, and CompuServe were not able to dominate Internet services either locally or nationally. This kept prices low and made it easy for users to change providers if service was unacceptable or if they found better services or prices with another ISP.

This started to change in the late 1990s with the introduction of ADSL and cable modem broadband services. These broadband technologies used the existing telephone local loop networks or cable hybrid fibre coaxial cable plants for their physical connections. Broadband required the addition of new electronic systems to provide the broadband connection to the user and the creation of a data network to aggregate the broadband traffic and interconnect with the global Internet. Broadband services no longer used the existing switched voice telephone network to connect to the Internet.

These ADSL and cable modem broadband networks concentrated the ISP business into the large telephone and cable television companies. The Federal Communications Commission (FCC), the communications regulatory body in the U.S., did issue regulations based on the Telecommunications Act of 1996 that provided a way for ISPs to set up competitive broadband businesses by giving them the ability to install their own broadband equipment in the telephone company's central offices. However, the cost for doing this was high enough that most of these new broadband competitors failed, and by the year 2000 it was clear that the large telephone companies and the cable modem companies would dominate the broadband market in the U.S. This gave these companies an unprecedented amount of power over the consumer Internet market. Different approaches to regulation governed the formation of successful competitive broadband operators in many countries in Europe and Asia and led to strong competition in broadband services there.

This consolidation of broadband services continued in the U.S. The FCC abandoned its attempt to encourage the formation of new competitive broadband companies based on sharing the facilities of the major telecom companies. The FCC came to believe that the competition between the telecom companies and the cable companies would generate enough competition to create a free market in broadband services and that sharing facilities was unnecessary. This left the ISP business in the hands of a small number of very large companies. As a result most of the small ISPs that were started in the 1990s went out of

http://doi.org/10.18080/ajtde.v3n1.5

business. The problem was that these large companies such as AT&T and Comcast were generally more interested in protecting the business that they already had rather than developing disruptive strategies that would bring higher speeds and lower prices as has occurred in Asia and Europe.

People in the U.S. already had a strong distrust for both the telecom and the cable companies. They just did not believe that these companies could be trusted to continue to develop a free and open Internet. This was confirmed in many people's minds in 2005 when AT&T's CEO Ed Whitacre set off a fire storm when he was quoted as saying that Google got a "free ride" on his network and that this unfairness could only be rectified by charging companies to ensure that their traffic reaches AT&T consumers quickly. People were up in arms over this statement and inundated the FCC with comments defending Net Neutrality to prevent the broadband operators for setting up so called "fast lanes" for those content providers who can pay. Whitacre's statement created a public relations problem for AT&T and the broadband industry in general. People believed that the Internet belonged to them and not to AT&T or Ed Whitacre and were afraid that AT&T, Comcast, and the other large broadband providers would take advantage of their position to provide undue influence over the Internet.

Even with all of this concern, few blatant violations of Net Neutrality principles have been identified in the U.S. Comcast was found to be blocking peer-to-peer file sharing of public files in 2007. Netflix complained that Comcast was discriminating against Netflix's video streaming service in favour of its own video services in 2012.

Technical Basis of Net Neutrality

The issues of Net Neutrality derive from the fundamental architecture and practices of the Internet. The Internet started as a government-funded project on a non-commercial basis. When it was formed there was no consideration of making money from the Internet. In fact the original Internet backbone funded by the U.S. National Science Foundation (NSF) was built as an academic network and was not supposed to carry commercial traffic.

A major result of this non-commercial approach was that there were no settlement charges. ISPs exchanged traffic with no charges to each other. There was an assumption that the traffic coming into a network was roughly equal to the traffic going out of it. Settlement charges would require a complex accounting infrastructure and generate little net revenue for most ISPs. The non-commercial and government-funded nature of the early Internet made settlement charging unnecessary.

http://doi.org/10.18080/ajtde.v3n1.5

On the other hand there was a long tradition of settlement charges in the switched telephone networks. Telephone companies charged on a call-by-call basis for all calls they carried that originated on the networks of other telephone companies. Per minute rates were very high through the 1980s and 1990s, so these settlement charges generated significant revenues for the telephone companies.

The lack of settlement charges in the Internet must have been difficult for the executives of telephone companies who were now in the broadband business to accept, since they were so accustomed to a business environment with settlement charges. This is why Ed Whitacre can say that Google pays nothing to AT&T for the traffic that it terminates on its broadband network, since there are no settlement charges between AT&T and Google. This was made worse from the perspective of the telephone and cable companies by the fact that the broadband operators were terminating much more traffic than they were generating. Their customers were receiving much more data than they were sending to the network.

This situation was exacerbated as video streaming services became more and more popular. This trend started with the success of YouTube and the increasing use of video in web sites. It has now moved to the streaming of TV shows and movies along with the move from Standard Definition (SD) to High Definition (HD) formats. Many people now have TVs, DVD/Blu-Ray players, or DVRs that support access to video streaming service such as Netflix, Vudu, and Hulu and spend many hours a week streaming video to their TVs over broadband networks.

All of the broadband operators, such as AT&T and Comcast, offer their own video services that compete with Netflix, Vudu, and Hulu. This puts AT&T and Comcast in direct competition with the video streaming companies and creates the potential for a significant Net Neutrality issue where AT&T or Comcast, for example, would throttle back streaming traffic from Netflix, Vudu, or Hulu to make their own service look better in comparison.

Net Neutrality is a bigger issue with the broadband service providers than with the Internet backbone providers. Content distribution networks significantly reduce the load of content distribution on the backbone. The content distribution networks collect Internet content in data centres in every major region and fulfil repeated requests for content locally without burdening the backbone. A piece of content is sent once across the backbone to the regional data centre and all further requests for that content are fulfilled locally.

The situation is different for the broadband service providers. They have to provide bandwidth for every time that someone requests a particular piece of content. This is particularly troublesome for video content, which requires high bandwidth, and is getting

http://doi.org/10.18080/ajtde.v3n1.5

worse with the broad adoption of HD TV and the impending adoption of 4K Ultra HD TV sets.

The concept of Net Neutrality lacks a rigorous technical definition. Supporters talk about Net Neutrality as "fast lanes" and are basically concerned about the fairness of broadband services. They do not want the broadband operators to favour one content provider over another, and they especially do not want the broadband operator to favour their own content services over other content service providers. So far traffic engineering and private peering are two technical tools commonly used in Internet architectures that are particularly troublesome for Net Neutrality supporters.

Net Neutrality and Traffic Engineering

The concept of Net Neutrality is bound up with traffic engineering. Packet network operators of all types now have powerful tools available that permit them to optimise the performance of their networks. They can apply deep packet inspection on their networks, which lets them look into every packet and make routing decisions based on the contents of the packet. They can route packets based on who sent them, who is receiving them, or the type of service that generates them. Traffic engineering tools are powerful and can be used for both good and bad ends.

Each Internet service has specific performance requirements associated with it. Internet browsing is bursty with long periods of inactivity while the user reads the page. It is also very resilient and can recover from packet loss or packet errors with little or no impact on the user's experience. Internet telephony requires relatively low bandwidth, but it needs packets to be delivered promptly, in order, and without loss. Packet delivery problems can cause clicks or pops in an Internet telephony call and can even make the call unusable in extreme cases. Video streaming requires high, constant bandwidth but is more tolerant of lost or error packets than Internet telephony.

Over-provisioning is the classic way to provide good quality of service over the Internet. Over-provisioning requires providing substantially more bandwidth than is required to support all of the services running. It is clear that over-provisioning will work, at least most of the time, but it is inefficient. The network operator must provide enough capacity so that the total amount of traffic is well below the capacity of the network.

Over-provisioning is certainly expensive; however, it cannot guarantee high quality services at all times. Traffic on the Internet has continued to grow at high rates with video streaming accounting for a higher and higher proportion of the total (CISCO 2015). It is difficult for

http://doi.org/10.18080/ajtde.v3n1.5

network operators to keep up with the growth and to maintain sufficient capacity margins to guarantee the correct transmission of all packets.

The traffic load on the Internet is highly dynamic. Traffic on the Internet can vary tremendously from day to day if not minute by minute. A viral YouTube video can drive traffic up, as can the introduction of a new iPhone by Apple or the release of a new line of lingerie by Victoria's Secret. Network operators have absolutely no control over these events, any of which can absorb all of their excess capacity in a matter of seconds.

One of the goals of traffic engineering is to let different services operate over the Internet and provide a good quality of service to all of them, or at least to the most important. In general it is more important to deliver Internet telephony packets correctly and in time than packets for web browsing.

Traffic engineering can prioritise traffic so the most highly affected services such as Internet telephony or video streaming get priority over lower priority services such as Internet browsing. Many people accept this kind of prioritisation, even though it does not treat all users equally and seems to violate the principles Net Neutrality.

The other side of the coin is that traffic engineering can be used by a network operator to put themselves in a favourable position relative to their competitors, or even to suppress entire classes of service. Netflix has expressed concerns that cable companies have degraded their service on their cable modem networks to make their own video services look better in comparison. There have also been complaints that network operators have suppressed peer-to-peer file transfer applications in order to protect the copyright of content. These kinds of actions have raised public ire and helped to generate strong public support for Net Neutrality.

Net Neutrality and Private Peering

Public peering points have been set up to allow Internet network operators to exchange traffic which each other. For example a public peering point may be used by AT&T and Verizon to exchange traffic with each other as well as all other Internet networks. Traffic that originates on AT&T's network that is destined for a user on Verizon's network may be exchanged at a public peering point and vice versa.

This is, in general, a good approach that provides connectivity between a broad set of networks. However, the major backbone providers long ago found that congestion at the public peering points could degrade the quality of service to other major backbone providers. It became common for the major backbone providers to set up private, direct private peering

http://doi.org/10.18080/ajtde.v3n1.5

connections with each other that bypass the public exchange points. This improved the quality of their services and relieved congestion at the public peering points.

What has happened more recently is that content providers have begun to generate significant amounts of traffic. Netflix by itself can generate more than one-third of all Internet traffic during busy periods. This seemed to be more traffic than the public peering points could handle, and seemed to cause a noticeable degradation in Netflix's video streaming service. In 2014 Netflix implemented private peering relationships directly with Comcast, Time Warner, AT&T, and Verizon – the largest broadband operators in the U.S. – and saw improvements in the quality of its video delivery as a result.

Net Neutrality advocates seized upon this as unfair because it allowed Netflix to purchase a "fast lane" that gave it preferred access to these broadband networks. Again, this is an issue of trust. The question is whether or not these broadband companies are limiting the availability of private peering connections to create an advantage for themselves or for preferred content providers. That is certainly a possibility, but no significant examples of such discriminatory private peering practices have been identified as yet.

However there is no transparency to these private peering arrangements. The networks with private peering arrangements do not disclose the terms of these agreements and in most cases do not disclose that these private peering arrangements have been put in place. This lack of transparency makes these private peering arrangements suspect in the minds of Net Neutrality advocates.

The Federal Communications Commission Takes Action

The Federal Communications Commission (FCC) regulates both the telecom and the cable companies in the U.S. The FCC has attempted to adopt policies that were consistent with the principles of Net Neutrality but has been blocked by the Federal Courts.

In 2010, the FCC approved an order (FCC 2010) that prevented network operators from blocking access to competitive services and web sites that included three Net Neutrality principles:

- **Transparency**: That all Internet Service Providers (ISPs) must transparently disclose to their subscribers and users all relevant information as to the policies that govern their network;
- **No Blocking**: That no legal content may be blocked; and
- **No Unreasonable Discrimination**: That ISPs may not act in a commercially unreasonable manner to harm the Internet, including favouring the traffic from an affiliated entity

http://doi.org/10.18080/ajtde.v3n1.5

This measure was denounced by Net Neutrality advocates as a capitulation to telecommunication companies that allowed them to discriminate on transmission speed, while pro-business advocates complained about any regulation of the Internet at all.

In early 2014 the Federal Courts ruled in a case brought by Verizon that the FCC has no authority to enforce two of these Net Neutrality rules since the broadband service providers have not been classified as Common Carriers under Title II of the Communications Act of 1934. Specifically, the court said that the "No Blocking" and the "No Unreasonable Discrimination" clauses were unconstitutional. The court did uphold the "Transparency" clause. This has thrown the FCC into a quandary and has reignited the public clamour for Net Neutrality. In response to this court ruling the FCC stated that it would propose new rules for Net Neutrality.

The Title II approach received strong support in November 2014 when President Barack Obama announced his support for it. He said that the FCC should create a new set of rules protecting Net Neutrality and ensuring that neither the cable company nor the phone company will be able to restrict what you can do or see online. He urged the FCC to reclassify the broadband operators such as Comcast and Verizon under Title II, giving the agency more power over how the companies operate.

The FCC received more than one million positive comments from the public for new rules that supported Net Neutrality. Supporters of Net Neutrality organised an “Internet Slowdown” on September 10, 2014 where participating web sites were purposely slowed down to demonstrate what would happen without Net Neutrality.

All of this resulted in the FCC issuing new rules in February 2015 (FCC 2015) that brought the fixed line broadband under regulation as common carriers under Title II of the Communications Act of 1934 using a “modernised, light-touch approach”.

The FCC first adopted three rules ban practices that it believes harm the Open Internet:

- **No Blocking:** broadband providers may not block access to legal content, applications, services, or non-harmful devices.
- **No Throttling:** broadband providers may not impair or degrade lawful Internet traffic on the basis of content, applications, services, or non-harmful devices.
- **No Paid Prioritisation:** broadband providers may not favour some lawful Internet traffic over other lawful traffic in exchange for consideration of any kind. This rule also bans ISPs from prioritising content and services of their affiliates.

http://doi.org/10.18080/ajtde.v3n1.5

The rules against blocking and throttling are to prohibit harmful practices that target specific applications or classes of applications. The ban on paid prioritisation is to ensure that there will be no "fast lanes"

The FCC established a standard to address any concerns from new services and practices. This standard is that ISPs cannot "unreasonably interfere with or unreasonably disadvantage" the ability of consumers to select, access, and use the lawful content, applications, services, or devices of their choosing; or of edge providers to make lawful content, applications, services, or devices available to consumers. The FCC will address questionable practices on a case-by-case basis based on this standard.

The existing transparency rule from its 2010 Net Neutrality rules, which was not struck down by the court remains in place. In addition the FCC will require that broadband providers disclose promotional rates, fees, surcharges, and data caps. Disclosures must also include packet loss as a measure of network performance, and provide notice of network management practices that can affect service. Small ISPs with 100,000 or fewer subscribers will temporarily be exempt from the transparency requirements for fixed and mobile providers.

The FCC also stated that other than paid prioritisation, an ISP may engage in reasonable network management. This recognises the need of broadband providers to manage the technical and engineering aspects of their networks.

- In assessing reasonable network management, the FCC will take account of the particular engineering attributes of the technology involved.
- The FCC stated that the network practice must be primarily used for and tailored to achieving a legitimate network management – and not business – purpose. For example, a provider can't cite reasonable network management to justify reneging on its promise to supply a customer with "unlimited" data.

The FCC will also make sure that a fixed line broadband provider's services that do not go over the public Internet, but are delivered over broadband access services, do not undermine the effectiveness of its Net Neutrality rules. Examples of these services include voice over IP services and IP-based pay TV services. The broadband operator's transparency disclosures will cover these offerings as well as their Internet offerings.

The FCC believes that for the first time it can address issues that may arise in the exchange of traffic between broadband providers and other networks and services. The FCC will be able

http://doi.org/10.18080/ajtde.v3n1.5

to hear complaints and take appropriate enforcement action if it determines the interconnection activities of ISPs are not just and reasonable.

The FCC did not put broadband wireless under Title II, so broadband wireless will not be subject to the same rules as fixed line broadband. This may change as broadband data grows in importance and becomes a target for Net Neutrality advocates.

Is Title II the Answer?

This is a controversial question that does not have a clear answer. Many people think that Title II will require the broadband operators to treat all of the content providers equally, and that Title II will prevent the network operators from prioritising traffic of one content provider over another.

Other people point out that there is nothing in Title II that would prevent broadband operators from using traffic management or offering high speed private peering connections that would improve the performance of a content provider's service across the operator's broadband network. If these people are right, Title II would not meet a fundamental requirement put forward by Net Neutrality advocates.

However, almost everybody agrees that Title II would require that the broadband operators to treat all content providers equally. It would prevent broadband operators from offering a preferential deal to one content provider without offering to other content providers on an equal basis. They would not be able to play favourites.

The Communications Act of 1934 and Title II resulted in the heavy regulation of voice telephony in the U.S. The FCC regulated it at the Federal level and each of the states regulated it at the state level. This made it difficult to introduce new services or new pricing models. The result discouraged innovation. The biggest change in public telephone service since the 1934 Communications Act was passed was probably the introduction of touch tone dialing in the 1960s, which slowly led to the introduction voice mail services and other voice response services over the next couple of decades. This is a very weak record of service innovation.

The FCC has chosen to apply only certain parts of Title II to the broadband operators. It seems possible that other parts of Title II could be added later based on pressure from Net Neutrality advocates and based on shifting opinions of the Commission itself as its membership changes over time. The American Cable Association has expressed concern that rate regulation could be added in the future (Light Reading 2015c).

http://doi.org/10.18080/ajtde.v3n1.5

A group of people who pioneered the implementation of Internet-based Voice over IP (VoIP) services objected to the FCC's Net Neutrality rules that brings broadband under Title II regulation. Title II regulation was a serious impediment to their services. They had to deal with regulators at both the national and the state level to address the technical differences between VoIP and PSTN services.

In the United States, the FCC required VoIP service providers to comply with requirements comparable to those for PSTN providers (Light Reading 2015a). VoIP operators in the US are required to support local number portability; make service accessible to people with disabilities; pay regulatory fees, universal service contributions, and other mandated payments; and enable law enforcement authorities to conduct surveillance, and provide a form of 911 emergency calling service.

The FCC's new rules state that 27 provisions of Title II and over 700 regulations adopted under Title II will not apply to broadband. Most importantly broadband operators will not be subject to utility-style rate regulation, including rate regulation, tariffs, and last-mile unbundling.

The FCC believes that its new rules are unlikely to have any negative financial effect on the broadband operators. Operators such as Sprint, Frontier, along with representatives of hundreds of smaller carriers that have already voluntarily adopted Title II regulation, have said that a light-touch, Title II classification of broadband will not depress investment.

The Next Steps for Net Neutrality

The problem with Net Neutrality arguments is that they address moral principles such as the "openness of the Internet". These arguments reflect the general lack of trust in the cable and telephone companies that operate the broadband networks and have created a major public relations problem for these companies. The arguments about Net Neutrality are not technical arguments about how to fairly apply traffic management or how to fairly set up private peering arrangements between content providers and broadband operators. Technical issues can be resolved. Arguments about moral principles or arguments that stem from distrust of the broadband operators are much more difficult to address.

It appears that the FCC's ruling is based on these moral arguments. This ruling will give the FCC the ability to closely monitor how the fixed line broadband operators run their networks and businesses based on the fear that they may take unfair advantage of their position rather than being based on past bad actions. There have been only a few violations of Net Neutrality that have been brought to the FCC. Net Neutrality is more of a public relations issue than a technical issue.

http://doi.org/10.18080/ajtde.v3n1.5

The FCC's new Net Neutrality rules will be enforced on a case by case basis (Light Reading 2015b). The rules do not clearly defined traffic management practices or business practices that violate Net Neutrality. The broadband operators may find only after completing an expensive development and marketing program that they are violating Net Neutrality rules and withdraw new services or modify their network operations.

Net Neutrality is a distraction from the most important broadband issues: the deployment of fibre and the creation of new business models to support advanced services. Fibre will provide broadband networks with the speeds that will needed for future services. The likely broad acceptance of 4K Ultra HD TVs among other new services will push many current broadband networks beyond their capacity. Fibre technology will solve this problem. Fibre deployments are likely to require new business models to support them. Raising the monthly price for a broadband is not likely to be the answer.

The broad availability of fibre access services will enable the deployment of new, high bandwidth services. It seems that the Internet is in the process of becoming the medium for all television services, especially for 4K Ultra HD and higher resolutions that current cable TV technology may not have enough bandwidth to support. Just as people are listening to all of their radio broadcasts over the Internet today, they are likely to watch all of their TV over the Internet.

Watching TV over the Internet benefits from new business models that could easily be construed to violate Net Neutrality. For example, a content provider may want to provide a portion of their subscription fees to the broadband operators to support the cost of deploying fibre networks. It seems fair that the content providers should support the investment required to deploy fibre networks, but it not clear that this will be allowed under these new rules.

For example, Netflix is being criticised for creating a new business model that would exempt its service from usage caps in Australia (Washington Post 2015). As part of its March 24, 2015 launch of the Netflix service in Australia, it was revealed that the broadband operator iiNet will exempt Netflix traffic from its customers' monthly bandwidth quotas (Gigaom 2015). This certainly makes Netflix more attractive for Australian customers since they will not have to track their usage and stop watching Netflix as they approach their cap. It is not clear that eliminating such arrangements in the best interest of the consumer. With the caps the consumer may have to upgrade to a more expensive broadband plan to get a higher cap. Consumers may well prefer that Netflix use some portion of their monthly fees to compensate the broadband operator for removing their caps for its content.

http://doi.org/10.18080/ajtde.v3n1.5

It seems possible that content distribution networks (CDNs) could be considered as a violation of Net Neutrality. A CDN distributes content from the source to regional data centres. The customer requesting the content will get it from a local data centre and will not have to go across the backbone the get the content from the source. This can significantly reduce the amount of traffic on the backbone; but it can also be interpreted as giving an advantage to content providers that can pay for CDN services and putting those who can't at a disadvantage. This would seem to violate Net Neutrality rules. However, CDNs are now an important resource for delivering content across the Internet, and challenging them could have significant deleterious effects.

It seems inevitable that the Net Neutrality controversy will continue, but it also seems that this does not have to be the case if the large broadband operators such as AT&T, Verizon, and Comcast would start treating Net Neutrality as a public relations issue rather than a regulatory issue. The first thing they should do is declare that Net Neutrality is a fundamental policy of their business. They should state that they will treat all of their customers and all of the content providers equally, and that they will apply network technologies only to provide the best experience for all of their customers and not to discriminate unfairly.

It may be too late to implement this approach in the U.S. Countries that are still developing Net Neutrality strategies should consider it. Managing Net Neutrality as a public relations problem is likely to be much more successful than slugging it out in the courts or with regulatory bodies. Declaring support for Net Neutrality would put the broadband operators on the right side of the moral argument and would defuse the public relations controversy. Maybe then we could all move on and focus on fibre deployment and creative new business models that will bring us more innovative new services.

References

Cisco. 2015. Cisco measures Internet traffic with its Visual Networking Index (VNI) and forecasts that IP video will be 79% of all IP traffic in 2018, up from 66% in 2013.

FCC. 2010. Preserving the Open Internet, FCC GN Docket No. 09-191. Available at: https://www.fcc.gov/rulemaking/09-191

FCC.2015. FCC Adopts Strong, Sustainable Rules to Protect the Open Internet. Available at: http://www.fcc.gov/document/fcc-adopts-strong-sustainable-rules-protect-open-internet

Gigaom. 2015. "Netflix won't count against iiNet broadband caps in Australia". 02 March, 2015. Available at: https://gigaom.com/2015/03/02/netflix-wont-count-against-iinet-broadband-caps-in-australia/

http://doi.org/10.18080/ajtde.v3n1.5

Light Reading. 2015a. “Internet Pioneers Decry Title II Rules” 02 March, 2015. Available at: http://www.lightreading.com/net-neutrality/internet-pioneers-decry-title-ii-rules/d/d-id/714129

Light Reading. 2015b. “Wheeler: We’ll Enforce Title II ‘Case-by-Case’ " 03 March, 2015 Available at: http://www.lightreading.com/net-neutrality/-wheeler-well-enforce-title-ii-case-by-case/d/d-id/714156

Light Reading. 2015c. “ACA Doesn’t Buy FCC’s Rate Reg Reassurances” 05 March, 2015 Available at: http://www.lightreading.com/net-neutrality/aca-doesnt-buy-fccs-rate-reg-reassurances/d/d-id/714237?

Washington Post. 2015. “Netflix tries to explain its apparent sudden flip-flop on net neutrality" Washington Post, 04 March, 2015. Available at: http://www.washingtonpost.com/blogs/the-switch/wp/2015/03/04/netflix-tries-to-explain-its-apparent-sudden-flip-flop-on-net-neutrality/

http://doi.org/10.18080/ajtde.v3n1.5

Data Privacy Law in the Asian region

Review of '*Asian Data Privacy Laws – Trade and Human Rights Perspectives*' by Graham Greenleaf

David Vaile
University of New South Wales

Summary: This article reviews *Asian Data Privacy Laws – Trade and Human Rights Perspectives'* by Graham Greenleaf, University of New South Wales

Review of 'Asian Data Privacy Laws – Trade and Human Rights Perspectives' by Graham Greenleaf, Oxford, October 2014, nearly 600 pages
Print ISBN-13: 9780199679669
DOI:10.1093/acprof:oso/9780199679669.001.0001
http://www.oxfordscholarship.com/view/10.1093/acprof:oso/9780199679669.001.0001/acprof-9780199679669

With the recent passing of Singapore's Lee Kwan Yew, there is renewed attention on the prospects for what some see as his brand of authoritarian corporatist development becoming the model for other Asian countries. But generalisations are often unhelpful. A new book from privacy scholar Professor Graham Greenleaf, *Asian Data Privacy Laws – Trade and Human Rights Perspectives*, offers many things for many readers, including a framework for robustly comparing the sort of human rights protection Singapore's privacy law offers with that in other countries.

The result of such a comparison is unexpectedly unflattering for Australians, for the lack of constitutional protections, limited and uncertain common law or equitable remedies for abuse of privacy, and a statutory data protection regime with promise but weakened by loopholes and relatively unused enforcement options, all sounds rather familiar. And some Australians could only dream of Singapore's reformed doctrine of privity of contract, a twist that could allow Singaporeans to enforce a contract made to benefit their privacy, even if they are not a party – over-zealous Internet data-mongers watch out!

Comparative studies of national data privacy laws, their underlying principles, and what constitutes effective administration of such laws, are still relatively uncommon except for the region around the European Union. Greenleaf's comprehensive and authoritative survey is the first work to examine data privacy laws across Asia in such detail, with an in-depth analysis of data privacy authorities and their powers in nine Asian countries and a lighter review of 20 more, as well as the international context in which these laws have developed,

common themes among them, regional trends and possible futures. Its keenest readers will be academics, regulators and policy-makers in the areas of data protection and privacy law, with legal practitioners in the Asian region and beyond also in line.

But anyone looking to see how the Asian century is manifesting in the area of privacy protection, including many in the technology and communications sector, will also find something of interest among the mass of detail on individual regimes or in the birds-eye view of the historical, cultural and legal context. Robert Gellman, former Chief Counsel to the Government Information Subcommittee in the US House of Representatives, observed that "as good as the national law chapters are, I found the introductory and concluding chapters even better."

Blair Stewart, New Zealand's Assistant Privacy Commissioner, noting Professor Greenleaf's deep roots in the region, recently spelled out why the book was produced: "up to now it has not been easy to get up-to-date and reliable information about the privacy laws in countries as disparate as Vietnam and Indonesia. How, for example, do we differentiate between the regulatory roles of South Korea's soup of agency acronyms like PIPC, MOSPA, KISA, PIDMC or KCC? How do we find out if there are controls on telemarketing in Singapore or on ID card numbers in Hong Kong or Japan (for your information, controls exist for all of them)? Clearly, a more authoritative source than Wikipedia is needed."

Part I Asia and international data privacy standards (Chapters 1 – 3) sets out the environment in which data privacy laws exist in Asia, the context and history, and the international standards that affect Asian privacy laws. While a dense and detailed exercise – US practitioner K Royal noted this section was 'not easy reading' – it provides the holistic frame that many will have been missing from a casual acquaintance with data protection in Asia, and the foundation on which the second part is built. Gellman notes the impressive breadth and depth of the book, and lauds the standards Greenleaf used to measure each country's national data protection law starting from this foundation.

Part II National data privacy laws in Asia (chapters 4–16) "reviews the most important national laws with a chapter on each country that describes in detail the country's legal framework, approach to privacy, relevant constitutional provisions, the substantive provisions of the privacy law, a description of the enforcement measures, and more." Greenleaf discusses and critiques the laws and their rationale for 28 Asian nations, plus a bit on North Korea (not surprisingly, the shortest and lowest point of the journey). There is more detail on more developed and central participants in the region, like Hong Kong and Malaysia, than on the more peripheral, such as Cambodia or Sri Lanka.

http://doi.org/10.18080/ajtde.v3n1.6

As well as analysis of all specialised data privacy laws in Asian countries, it covers constitutional and treaty protections, and protections in the general civil and criminal law – important in those countries still without specialised legislation.

Part III Regional comparisons, standards, and future developments (chapters 17–20) puts Asian privacy laws into a comparative scheme: sources of protection, scope, principles, liabilities and international dimensions; and then assesses privacy law enforcement. Royal notes, “this is the compelling read for privacy professionals who … take a risk-based approach to priorities” – the crunch issue in which commercial clients will hope readers take a deep interest! In addition, it analyses the international agreements and standards concerning data privacy that are relevant to Asia, including those of the European Union (EU), the Organisation for Economic Co-operation and Development (OECD), and the Asia-Pacific Economic Cooperation (APEC).

Finally, a survey of Asia’s prospects warrants cautious optimism about the future. Greenleaf writes, "In Asia, data privacy laws, or in some cases their enforcement, have not yet caught up with surveillance technologies and practices, but they are necessary, even though (as everywhere) they need to be supplemented with other modes of regulation. There are grounds for optimism, but not overconfidence, that in future they will restore a better balance between the human right of privacy and other interests."

This will no doubt become part of the reference library of a wide array of those engaged with human rights, trade negotiations and privacy law in the Asian region and beyond, assisting the reader to move beyond a passing awareness of a particular privacy law in an isolated example to a wide and deep apprehension of the shape and details of the laws of the region.

Contents

http://doi.org/10.18080/ajtde.v3n1.6

http://doi.org/10.18080/ajtde.v3n1.6

Public Telephone Cabinets In Australia

Simon Moorhead
Ericsson Australia & New Zealand

Summary: Two papers from the Telecommunications Journal of Australia in 1956 and 1960 respectively. The first provides an overview of public telephone cabinets in Australia and the second describes the state of the art, aluminium public telephone cabinet.

Introduction

It is hard to believe that public telephone cabinets have been around since the First World War in Australia. Before the ubiquitous mobile telephone, public telephone cabinets were situated in most popular city, metropolitan and country locations. They started as grand attachments to post offices and railway stations and transformed into practical enclosures sympathetic to climate, capital cost and maintenance.

Both papers are written by Mr H J Lewis, a Divisional Engineer attached the Telegraphs and Workshops Section of Central Administration of the Postmaster General's Department.

The first paper (TJA 1) provides a summary with photographs of the striking range of public telephone cabinets that were in operation in 1956 and summarises the desirable features of cabinets at that time.

The second paper (TJA 2) from 1960 describes the development of the aluminium public telephone cabinet which was state of the art and had many advantages over its predecessors.

Public telephones were deregulated in 1989 and are now supplied by Telstra and a number of other private companies.

Despite the penetration of mobile telephones, there is community concern over the reduction and location of public telephones in Australia, which form part of the Universal Service Obligations (USO). The social benefit of public telephones with respect to these USOs is monitored by the Australian Communications and Media Authority

References

TJA 1 - Lewis, H.J. 1956. 'Review of Public Telephone Cabinets in Australia', *Telecommunication Journal of Australia*, October 1956, Vol. 10, No. 5, pp. 155-157.

TJA 2 - Lewis, H.J. 1960. 'Review of Public Telephone Cabinets in Australia', *Telecommunication Journal of Australia*, October 1960, Vol. 12, No. 5, pp. 338-341.

http://doi.org/10.18080/ajtde.v3n1.7

REVIEW OF PUBLIC TELEPHONE CABINETS IN AUSTRALIA

*H. J. LEWIS**

Introduction

The climatic conditions in Australia vary greatly. The northern areas are sub-tropical with heavy rainfall in the coastal regions while a small highly populated area, consisting of southern Victoria and Tasmania, has a temperate climate with moderate rainfall. Generally speaking, provision has to be made for the warm conditions, as even in the coolest populated parts summer temperatures are likely to reach 100° F, while it is unusual for winter temperatures to fall much below 40° F.

Public telephone cabinets have been used extensively to relieve the shortage of private services as well as provide a public telephone service at key points throughout the cities and suburbs. They are generally good revenue-earners and so warrant attention to their design so that they provide an adequate service to the public. The cabinets are usually installed on the pavement or nature strip associated with the footway, close to the gutter or road kerb. They are also installed in banks of two or three outside railway stations or suburban post-offices where the calling rate is high.

There have been many cabinet designs built from several different materials during the last 60 years. The following sections of this review outline the chief features of the main types of cabinets still in use.

A.P.O. Former Standard Type Cabinet

This type of cabinet was introduced in 1933 and is constructed in several forms, namely—

(a) Glazed half length, 2 sides and door
(b) Glazed full length, 2 sides and door
(c) Glazed full length, all sides and door
or (d) Glazed half length, all sides and door.

Fig. 1.—Former A.P.O. Standard Cabinet, mounted on "cast in situ" Concrete Base.

The material used is chiefly wood, with galvanised-iron for roofing. The floor is linoleum-covered wood and the glazing is provided by ¼-inch drawn or plate glass. The cabinet is bolted down to a "cast in situ" concrete base. The general construction is shown in Fig. 1.

Experience has shown the following deficiencies in this design as generally used in Australia:—

(a) The wood joinery is a point of weakness and it is difficult to exclude the weather and maintain satisfactory jointing.
(b) The wood tends to rot around the lower portion of the cabinet due to ingress or absorption of moisture.
(c) The ventilation is inadequate.
(d) The cabinet requires frequent repainting (approximately every 3 years).

Several modifications have been made to improve ventilation, but they have not been adopted as standards pending a redesign of the cabinet to overcome the troubles listed under (a), (b) and (d).

Fig. 2—Steel Cabinet with Domed Asbestos-Cement Roof, and all but one row of Glazing fitted with Non-Actinic Glass.

B.P.O. Type Steel Cabinet

This type of cabinet is a close copy of the design used by the British Post Office and called Kiosk No. 3. It is made of sheet steel instead of the cast iron or concrete used by the B.P.O. Small quantities were made during the war period, but, due to failure of the surface by rust, and poor ventilation; production has been discontinued. The most notable point of design in this cabinet is the introduction of a cast asbestos-cement roof. This material has proved satisfactory and is being used on some of the recent designs. The general construction is shown in Fig. 2.

Fig. 3.—Precast "Pipe-Type" Concrete Cabinet with Wooden Door, concrete Floor forming mounting Base and Concrete Roof.

Concrete Types of Cabinet

Two main types of concrete cabinet have been made in Australia. The first of these is illustrated in Fig. 3, and was introduced in 1927. It consisted of a body made from a cylindrical precast pipe with two small half length windows and a wooden door. These cabinets had precast concrete roofs, some of which were flat with a small conical metal ventilator in the centre of the roof whilst others were conical and ornamented to represent small roofing tiles. The cabinets were fitted into a precast base which as well as providing the mounting, served as a floor. 6 x 1½ inch openings were made at floor level around the walls to act as drainage ports for the interior, and also provided ventilation, which is considered inadequate. The other type of concrete cabinet, shown in Fig. 4, was a flat-roofed, square shaped, fully glazed one which was cast in situ.

Both of these cabinets have been satisfactory from a maintenance viewpoint, although if damaged by a heavy blow such as given by a vehicle, they are

*Mr. H. J. Lewis is a Divisional Engineer attached to the Telegraphs & Workshops Section of the Central Administration.

Fig. 4—Concrete Cabinet cast in situ—Metal Door and Window.

more or less a total loss. They also have the disadvantage of requiring much "on-site" work during their installation, and are not readily shifted to another location.

Fig. 5.—"Head-Box" Type of Cabinet Installed in New South Wales.

Head-Box Types of Cabinets

As a measure of economy and for speedy erection during the recent war period, head-box types of cabinet were introduced. As illustrated in Fig. 5, these provided protection for the equipment and shelter for the user above waist level only. They are not greatly favoured by the public, but have proved useful in locations where pavement space was not available for a full sized cabinet. These cabinets were also used in Military Camps where a temporary public telephone service was provided. Despite the usage disadvantages of this type of cabinet, its maintenance costs are very low, its ventilation reasonably good, and it can provide very satisfactory services in special locations.

Special Tropical Types of Cabinet

As well as the louvre-glazed type of cabinet described in a following section of this review, two types of wooden panel-glazed cabinets have been produced in Queensland to suit the tropical climate. These cabinets are shown in Figs. 6 and 7. These cabinets have a somewhat similar appearance to the former standard cabinet. The roof ventilator has been deleted and additional ventilation introduced under the eaves. The lower glazing areas of the cabinet have been replaced with wooden louvres in one type, whilst in the later model a gap has been provided beneath the glazing areas to allow free ingress of air. The roof is galvanised iron and is unpainted. These cabinets suffer from the same faults of wood joinery exposed to the weather as the former standard type.

Fig. 6.—Tropical Type Cabinet Installed in Queensland.

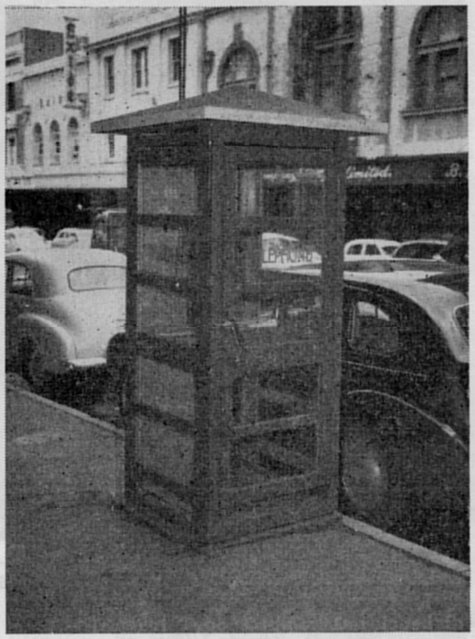

Fig. 7.—Tropical Type Cabinet Developed in Queensland in 1955.

Recent Types of Cabinet

Two recent design changes have chiefly been related to appearance and ventilation. Domed roofs similar to those used on the B.P.O. cabinets were introduced in Victoria on the former standard type cabinet. Initially these were made with a wooden framework covered with lead sheeting. The same design was repeated with cast aluminium and also cast asbestos cement. In the latter types the ventilation area was increased in the roof by adding small cast aluminium ventilators to the gables. At the same time the ceiling and cabinet body ventilating areas were increased, which greatly improved the ventilation in the cabinet. This type of roof is shown in Fig. 8.

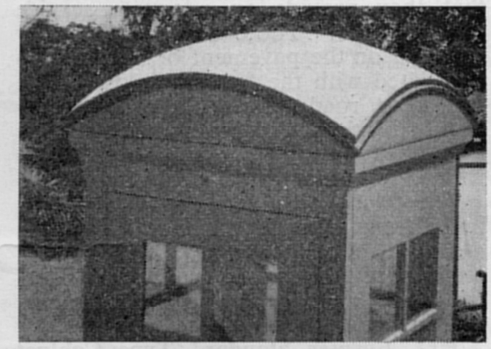

Fig. 8.—Domed Roof Used in Victoria on "Modified-Standard" Cabinets. Ventilators have been added to the Gables of recent models.

In 1951, following the manufacture of a steel cabinet of modernistic design, as shown in Fig. 9, in the Sydney Workshops, two new types of cabinet were

Fig. 9.—Experimental Louvre-Glazed Steel Cabinet Developed in New South Wales.

developed. The chief feature of these cabinets was the introduction of louvre glazing to provide excellent "cross" ventilation. Previous cabinets had been limited to a few square inches (except for the one described in the previous paragraph, which had 48 square inches) of "flue" ventilation whereas the louvres provided approximately 30 square inches on each side of the cabinet for "cross" ventilation as well as small "flue" ventilation in the roof to cool the roof cavity of the cabinet. Other features of this type of cabinet were an alternative roof

Fig. 10.—Recent Experimental Tropical Type Louvre-Glazed Wooden Cabinet.

with large eaves to provide more shade in tropical areas, improved door fittings, flush lighting, acoustic panelling around the upper portion of the walls and rubber flooring. This type is illustrated in Fig. 10.

Although this louvre-glazed cabinet appears to be satisfactory from a ventilation viewpoint, it still retains the same wood joinery principles which are a source of trouble in the former standard cabinet. It also requires frequent painting and is subject to deterioration due to wood rot around the lower portion of the cabinet.

Pending development of a cabinet of durable material, probably cast aluminium, the louvre-glazed cabinet with the domed roof shown in Fig. 11 has been adopted for present manufacture. Modifications have been incorporated in this cabinet which will limit some of the causes of deterioration previously experienced and enable all fittings to be made on a production basis. Previously the practice has been to purchase the basic cabinet shell and add the fittings at the time of installation. This is a costly practice which permits variations in methods of fitting the equipment and so spoils the appearance of the cabinet as well as developing unstandard arrangements. The chief modification introduced to the structure is the elimination of the wooden floor, one of the parts with the most wear and deterioration. In future this cabinet will be mounted directly on a concrete base with a durable surface, by four angle-iron pillars with a 1½ inch gap between the sides of the cabinet and the base. This method of mounting will isolate the wooden sides from the ground and reduce the failure on the lower rail due to capillary seepage of moisture, as well as provide a durable floor which can be readily cleaned.

Artificial lighting is provided by a flush panel fitting mounted in the ceiling, and the ceiling and upper portions of the wall are covered with an acoustic treatment which greatly reduces the "drumminess" of the cabinet and absorbs the external noise. The notice sign "Telephone" is provided on the top glass louvre by a sandwich joined with an epoxy resin adhesive. The lettering is a fused ceramic paint on the inner surface of the outer glass, with black characters on a yellow background. Provision is made for hidden light and communication cabling with two wooden corner strips rising up both the back interior corners of the cabinet. Cables can be brought in underground through these corner strips to the light and instrument, or where necessary, can also be brought in overhead, although the latter arrangement is not favoured.

Fig. 11.—Recently Adopted Louvre-Glazed Wooden Cabinet.

Summary of Desirable Features for Cabinets

The following points appear to be the desirable features for public telephone cabinets for the majority of the settled areas in Australia and for the greatest portion of year—

(a) Good ventilation is essential, with complete natural replacement of the air inside the cabinet at frequent intervals.
(b) Good natural lighting in the daytime and adequate artificial lighting after sunset which can be maintained by the street-lighting patrol.
(c) A complete view into the cabinet from all directions. This is necessary to limit vandalism and misuse of the cabinet.
(d) A floor surface which is durable and does not require repetitive protective treatment.
(e) An exterior surface which is durable and does not require repetitive protective treatment.
(f) Construction with materials which will withstand weather without deterioration, are readily available, repairable and can be fabricated without complex processes.
(g) Readily discernible by the public and yet aesthetically designed to harmonise with local buildings.
(h) Readily transported, erected or shifted.
(i) The fittings subject to damage such as windows, door fittings and lighting fixtures should be designed to facilitate quick replacement or repair in the field.
(j) Acoustic properties which limit the ingress of noise and the egress of conversation.
(k) Door to provide easy entry into the cabinet and easy access for cleaning.
(l) Shade, either by provision of eaves on the roof of the cabinet, or by placing the cabinet in a position where it is shaded by adjacent buildings or trees.
(m) Adequate designation signs either attached to the cabinets, or mounted on buildings or posts to direct callers to the public telephone.
(n) Provision for tidy and easy entry of lighting and transmission cables.
(o) Provision for interference-proof mounting of the equipment with fittings that are built into the cabinet.
(p) Easy access to the coin box for clearance purposes.
(q) Built-in notice frame and directory shelf.
(r) Interior of cabinet to be easily cleaned, finished in a hygienic manner and be not easily disfigured or damaged.

THE AUSTRALIAN ALUMINIUM PUBLIC TELEPHONE CABINET

*H. J. LEWIS**

INTRODUCTION

Public Telephone Cabinets are one of the three prominent structural facilities which the Postmaster-General's Department provides for public use. The other two are Post Offices and street Letter Receivers.

Not only is it essential that these facilities are efficient in their use and aesthetic in their appearance, but they must be durable and economical to provide and maintain.

This latter aspect has always presented a difficult engineering problem in regard to telephone cabinets, due to the number of variable influences such as climate, location, vandalism, and durability of materials.

In 1957, a committee was established to examine this problem and provide a cabinet suited to the Australian conditions. The original committee consisted of:—

Chairman: Mr. B. Edwards, Supervising Engineer, Telegraphs and Workshops.

Members: Mr. R. Lamb, Supervising Engineer, Melbourne Workshops; Mr. W. Murrell, Assistant Controller, Telecommunications Division; Mr. W. Waterworth, Senior Buildings Officer; Mr. A. McPherson, formerly Sectional Engineer, Telephone Equipment, now Superintending Engineer, Services, Victoria.

Mr. Edwards and Mr. Murrell have retired recently and Mr. McPherson transferred to other duties. These members have been replaced by:—

Mr. K. Smith, Sectional Engineer, Telephone Equipment.

Mr. L. Garrioch, Sectional Engineer, Workshops.

Mr. K. Richardson, Assistant Controller, Telecommunications Division.

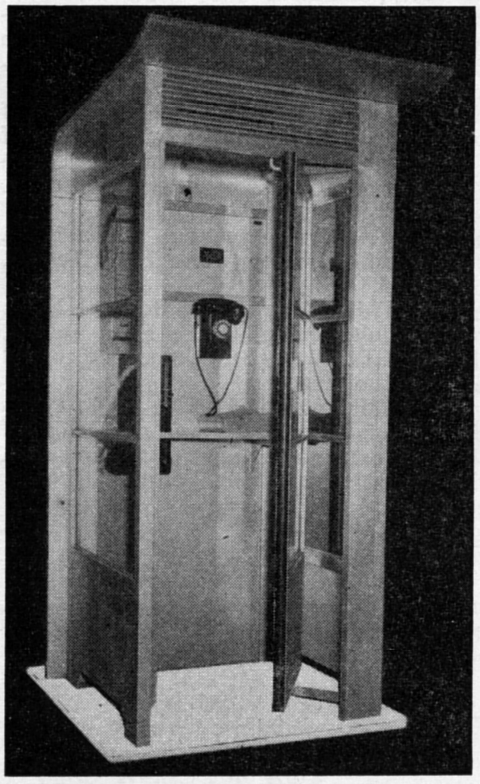

Fig. 1.—The Australian Aluminium P.T. Cabinet. This cabinet is at present installed at the corner of Spring and Collins Streets, Melbourne.

* See page 379.

PREVIOUS TYPES OF AUSTRALIAN CABINETS

Before describing the aluminium cabinet, a brief look at some of the previous designs and their weaknesses is of interest.

Early cabinets were usually built as attachments to post offices. They were made solidly with a small amount of glazing, poor lighting and ventilation, and, because of their bulk, were unsuitable for installation in residential streets. Fig. 2 shows a typical cabinet of this era.

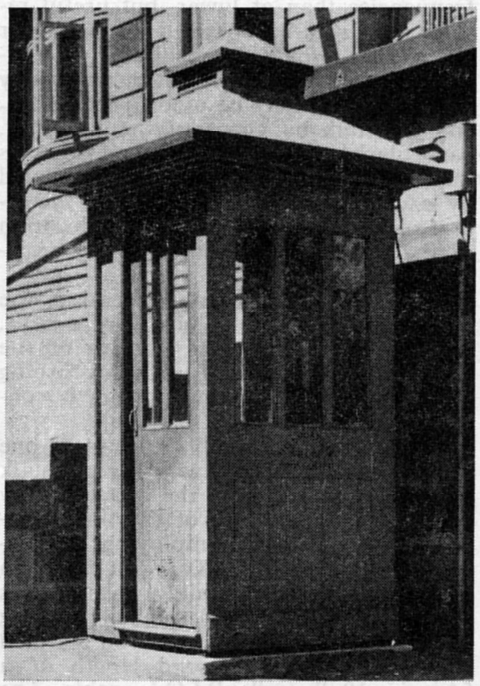

Fig. 2.—Early Wooden Half-Glazed Cabinet.

Subsequently, several designs of reinforced concrete were tried, some of which are still in existence (see Fig. 3). However, this material failed either due to corrosion of the reinforcing material, the weight of the cabinet causing subsidence of the footpath, the difficulty in repairing the concrete if damaged by a vehicle, or impossibility of moving it to a new site.

Pressed steel cabinets were also tried, but they cannot be considered successful due to the high cost of combating corrosion. In 1935, a wooden cabinet with fixed glazing in four styles was produced.

Style (a) Glazed all sides full length.
Style (b) Glazed three sides full length.
Style (c) Half-glazed all sides.
Style (d) Half-glazed three sides.

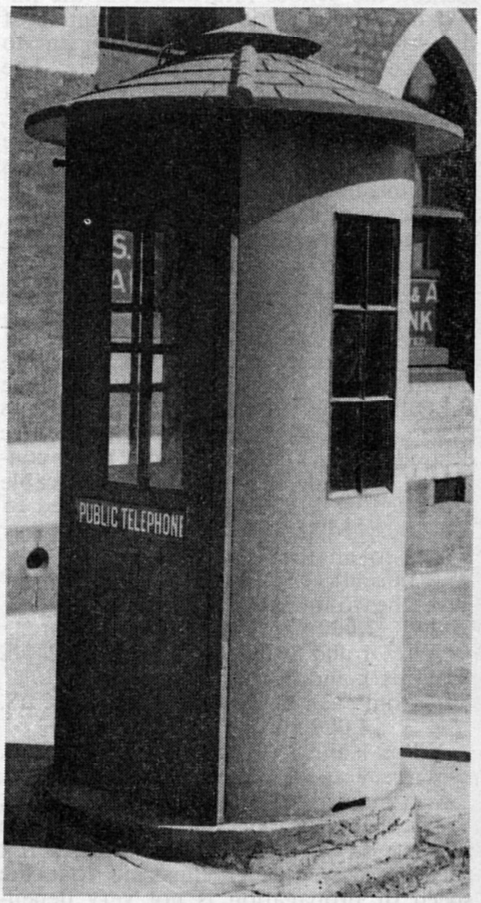

Fig. 3.—Reinforced Concrete Cabinet.

This type of cabinet originally had a hip-type roof with a lofted ventilator and wide eaves. Later models had a domed roof of either lead-covered timber, aluminium or asbestos-cement.

The cabinet was poorly ventilated, had a wooden floor which was subject to rot, particularly if the cabinet was mounted on wooden plinths and required frequent repainting to preserve it. Even so, the joinery in the lower portion of the sides and back failed due to capillary entry of ground moisture.

Many of the faults of the early wooden cabinets were reduced or eliminated in 1956, in the present wooden louvre-glazed cabinet. The wooden floor was omitted and the cabinet mounted on four metal angle brackets 1½″ above the concrete floor-base. Durable timber was specified, iron-work protected, and the lighting and acoustics improved with a flush-fitting and acoustic lining in the

ceiling. The louvre glazing provided most of the ventilation, and was assisted by ventilation through the ceiling and under the 1½" gap between sides and floor.

However, it is still necessary to repaint these cabinets at regular intervals to preserve their appearance and ensure long life for the joinery.

Fig. 4.—Wooden Cabinet with Fixed Glazing, Style (a).

DEVELOPMENT OF THE ALUMINIUM CABINET

In the light of this experience, the committee examined many overseas designs of cabinets, new materials (including plastics), and methods of manufacture which would allow "packaging" of a cabinet for long-distance transport.

At the outset, the requirements of a cabinet were determined by the committee, and are listed hereunder:—

Economics.

(i) Low maintenance costs.
(ii) Long site life.
(iii) Low installation, transport and storage costs.
(iv) Minimised first cost.
(v) Restriction of types, preferably to one.

Design.

(i) Easy to use.
(ii) Good ventilation.
(iii) Effective lighting, natural and artificial.
(iv) Readily discernible to the public.
(v) As good acoustic properties as possible.
(vi) Door to provide protection and privacy.

Technical.

(i) View into cabinet from at least three directions.
(ii) Design to suit shaded locations.
(iii) Provision for tidy and easy cable entries.
(iv) Protection of equipment.
(v) Easy access for coin clearance.
(vi) Built-in notice and directory shelf.
(vii) Cabinet to be easily cleaned.

A range of designs was prepared to enable the committee to consider this technical information in relation to both the desirable features, previously listed, and aesthetic, manufacturing, installation and usage viewpoints.

The previous practice of colouring departmental cabinets red was, by permission, departed from in the design of the new cabinet. This removed one of the main restrictions which had previously limited the improvement of cabinets. After survey of available materials aluminium was chosen as the material which fulfilled most of the requirements. This material blends well with present-day architecture, is readily procurable and workable, and offered a durable finish at reasonable cost.

Fig. 5.—Louvre-Glazed Wooden Cabinet.

The committee's first design of cabinet had a curved roof, solid back, fixed glass panelled sides and door, the latter folding into the cabinet, and a concrete base forming the floor.

Fig. 6 is an architectural sketch first depicting the new design.

Trial folding doors were fitted to wooden cabinets in service and their operation studied. Due to the strains that are placed on this type of door, the fault incidence was high and the welded joints tended to fail. It was observed that if a user had collapsed in the cabinet it was a near impossibility to gain access without dismantling or damaging the cabinet. Also, extra floor space is required for this type of door.

However, one feature which proved its worth was the partial opening of the door in its normal position. This feature has been retained on the present door but folding action has been abandoned.

A major contribution was made to the design of new cabinet by the development of a special door by Mr. L. C. Gemmell, then Sectional Draftsman of the Melbourne Workshops. The body of the door is a single panel which is pivoted on cantilever arms in such a way as to allow the door to swing inside and along the right hand wall of the cabinet. The door is stabilised and its movement controlled by an arm fixed to it and sliding in a track above the door frame. Fig. 7 shows the various positions of the door. A spring in the upper door pivot returns the door to its normal position. The door may be closed completely by a gentle push from the inside. No lubrication is required for the mechanism as Nylon is used for the bearings and slipper.

A spring at the hinge end of the guide track absorbs the opening jar of the door. The spring at the other end of the guide track normally holds the door partly open, but when the door is fully closed it acts to locate it in that position.

The aluminium-sheathed plywood used on the back and roof consists of waterproof plywood with a heat-bonded aluminium sheath. An aluminium capping, applied with aluminium screws and epoxy resin, seals the edge of the plywood. The surface of the aluminium on the inside of the cabinet is embossed with

Fig. 6.—Sketch of the First Design of Australian Aluminium P.T. Cabinet.

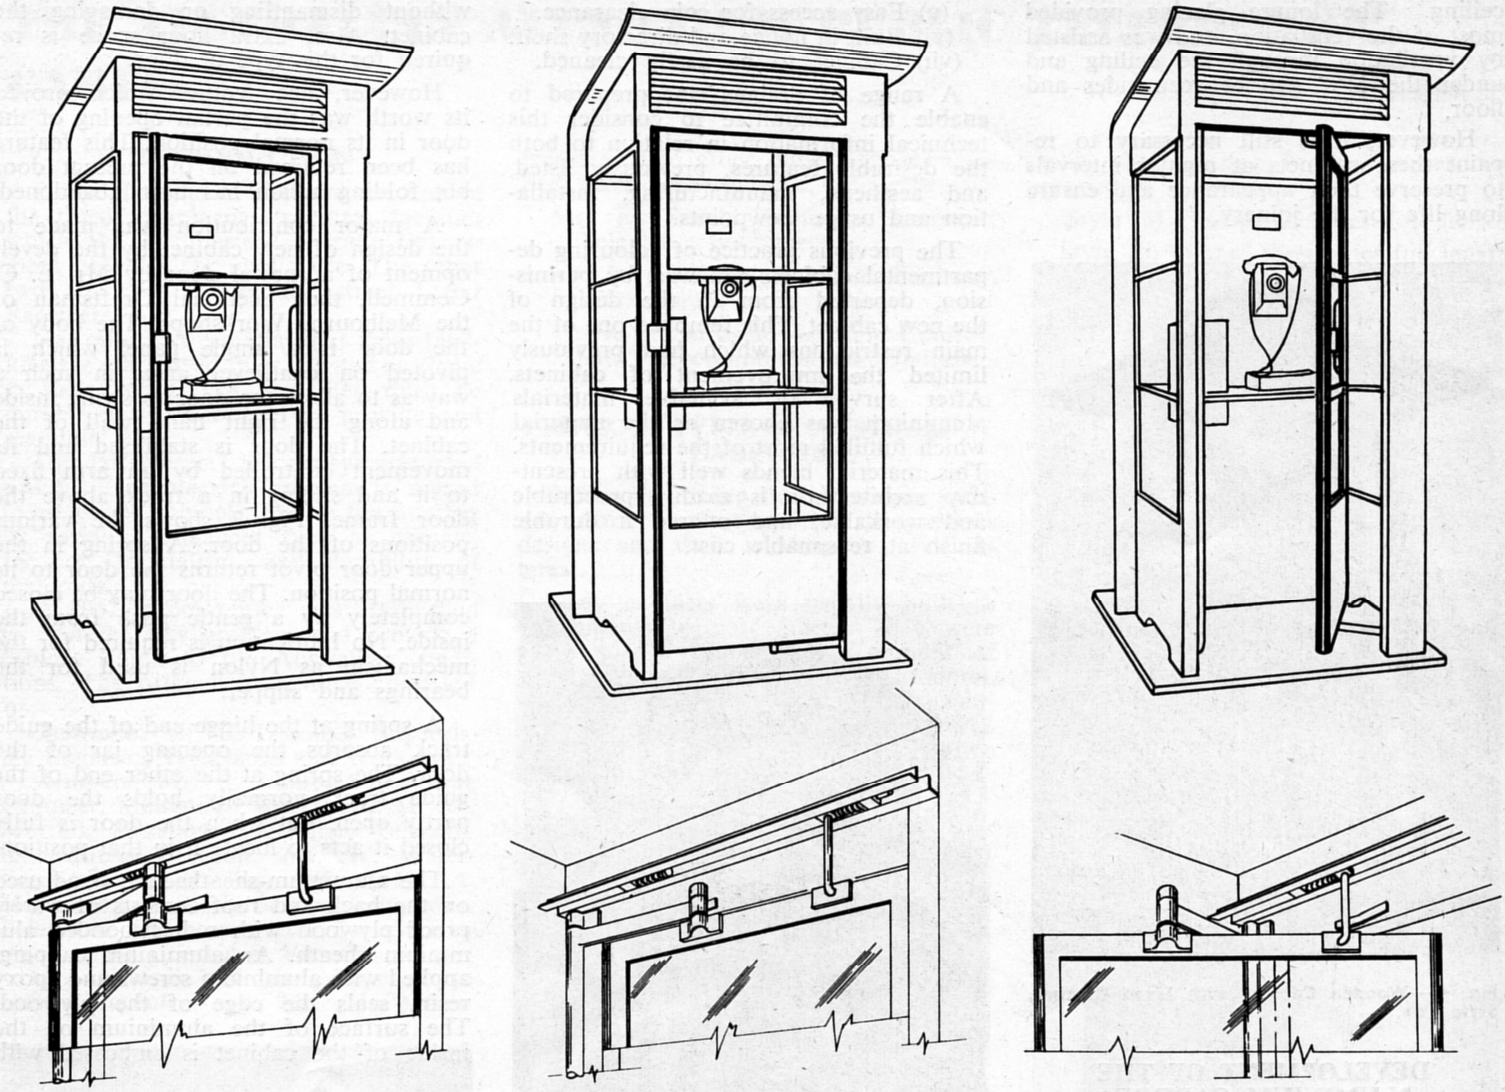

(A)—Normal Position. (B)—Fully Closed Position. (C)—Fully Open Position.

Fig. 7.—Various positions of the Gemmell Door.

a mesh pattern to restrict writing thereon, and the exterior is wire-brushed to give a satin finish.

The sides, floor, and door frame are constructed of special aluminium sections made of alloy AA 50S-T5, and ⅛″ sheet aluminium alloy AA 65S-T6. The first cabinet was made by welding these sections to form the desired framework. However, distortion was difficult to avoid and a fabricated assembly was designed, using the shapes of the sections to lock them together. The top and lower ⅛″ panels of the sides are respectively screwed and rivetted in position after the horizontal members have been fitted using self-tapping screws.

Glass in the sides and door is mounted and held in position by a "Neoprene" Strip designed for the purpose. This strip holds the ¼″ plate glass firmly but without pressure. The Neoprene strips and glass panels are held in position by aluminium strips which lock each other in turn, the final one being secured by one screw in the top strip of each panel.

The edges of the door are fitted with a rubber buffer strip. This eliminates any possibility of injury to fingers caught in between the door and the door frame. Also, it provides a weather seal when the door is closed.

The directory shelf is made of plywood with a mottled grey synthetic veneer covering on the face and edges. The shelf is mounted on right-angle brackets fixed to the back of the cabinet.

The coin attachment is mounted on a metal backplate previously screwed to the aluminium-sheathed plywood. One of the advantages of the latter material is that its wooden core provides a good medium for mounting the instrument, notice frame, ducting, etc.

Communication wiring is brought into the cabinet at ground level and runs in an aluminium duct up the right-hand back corner of the cabinet and across the back face to the wall-mounted telephone instrument and coin-attachment.

Lighting is provided, at present, by an incandescent lamp fitting centrally mounted on the ceiling. Supply wiring enters the cabinet at ground level passing up ducts in the left-hand corner of the cabinet, then across the back of the cabinet to a small switchboard, and then across to the light fitting. It is hoped to use fluorescent lighting when arrangements have been made regarding lamp replacement.

Ventilation is obtained in two ways. The door remains partly open when the cabinet is not in use and allows fresh air to continually enter the cabinet. When the door is closed, the 4″ gap between the sides and floor, and the louvre ventilator over the door, continue to provide ventilation to the user.

A fused ceramic symbol on the upper glass pane of each side provides an international notice for the public of the purpose of the structure.

Twelve of these cabinets have been made and are being field-tested throughout the Commonwealth. Six of the

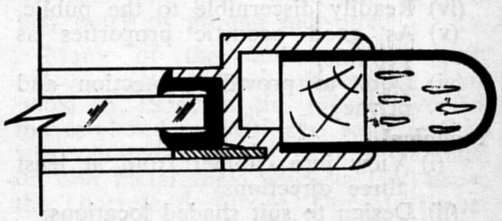

Fig. 8.—Section Through Door style showing "Neoprene" Glazing Strip Rubber and Buffer Strip.

twelve are finished with an anodised surface and the other six have a coating of Butyrate lacquer. Whilst the anodising process provides a very hard and durable surface, it is expensive and may not be justified, but the field tests will decide this matter.

PROPOSED MODIFIED DESIGN WITH FLAT ROOF

Examination of the costs of the present design of cabinet has led the committee to the development of a modified version with a flat roof without the overhanging back and roof. The louvre ventilator above the door has been omitted, and ventilation is provided by suspending the roof on angle brackets 1½" above the sides of the cabinet. Storm water is shed at the back of the cabinet.

Fig. 9 is a sketch of the proposed new model. The committee is continuing

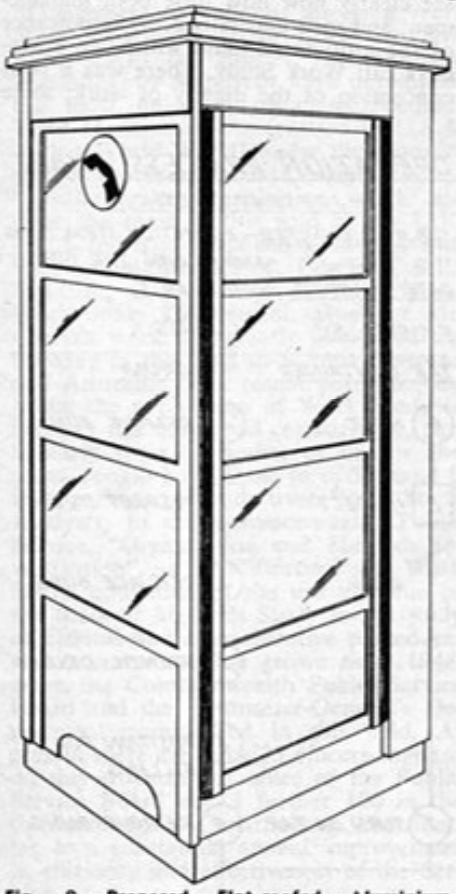

Fig. 9.—Proposed Flat-roofed Aluminium Cabinet.

development of this cabinet during the field trials of the other twelve. Reports and observations on the first cabinet have indicated that a big step has been taken towards the design of a cabinet which will meet the requirements of Australian conditions, be a serviceable cabinet, and have minimum maintenance requirements. No doubt further improvements will be made after the field trials, and another article will be published when design is finalised.

Acknowledgment

The author wishes to thank Mr. L. C. Gemmell for his assistance in preparing this article.

H. J. LEWIS, author of the article "The Australian Aluminium Public Telephone Cabinet" is a Divisional Engineer in the Workshops Section at Headquarters. He joined the Department as Radio Technician in 1939 and worked at the Melbourne Broadcast Studios and Lyndhurst Radio until qualifying as Engineer. Mr. Lewis then served in the Outer Metropolitan Lines Division during 1946-47, and in the Melbourne Workshops as Engineer and Divisional Engineer, Design and Plant, for the following four years. Subsequently he has been in the Workshops Section at Headquarters in various duties of Planning and Production work, and recently on Plant and Workshops Practices. Mr. Lewis had practical experience in the manufacture, installation, and maintenance of Telephone Cabinets, whilst at the Workshops and has since served on the Committee which designed the present louvre-glazed wooden cabinet. During the recent illness of Mr. Edwards he acted as Chairman of the Committee which designed the new aluminium cabinet.

H. J. LEWIS