Draft Minutes of Annual CCIRN Meeting
4-5 June 2005
PSNC, Poznan, Poland

Attendees

Asia-Pacific Delegation		
Jie An	CERNET	China
Shigeki Goto	APAN/Waseda University	Japan
George McLaughlin	APAN/AARNet	Australia
Francis Lee	SingAREN/Nanyang TU	Singapore
Xing Li (Co-Chair)	CERNET	China
European Delegation		
Claudio Allocchio	TERENA/GARR	Italy
Artur Binczewski	PSNC	Poland
Vasilis Maglaris	NTUA	Greece
Kevin Meynell (Scribe)	TERENA	-
Kees Neggers (Co-Chair)	SURFnet	The Netherlands
Dorte Olesen	TERENA/UNI·C	Denmark
Klaus Ullmann	DANTE/DFN	Germany
Karel Vietsch	TERENA	-
Latin American Delegation		
Florencio Utreras	CLARA	-
North American Delegation		
Heather Boyles	Internet2	United States
Warren Matthews	Georgia Tech	United States
Grant Miller	NCO IT R&D	United States
Don Riley	University of Maryland	United States
George Strawn (Co-Chair)	NSF	United States

1. Opening

Kees Neggers, as the co-chair from the continent hosting the meeting, opened the meeting. He re-iterated the proposed structure of the meeting: no long updates on general developments, but concentration on a small number of topics for discussion. Presentations or papers giving updates on continental developments could be made available via the CCIRN website.

2. Minutes of the last CCIRN meeting (Cairns, 3 July 2004)

The minutes were approved without change, with thanks to the author, Grant Miller.

3. Introductions by the Co-Chairs

George Strawn announced that Grant Miller would distribute three updates on paper in the meeting: a presentation by Joe Mambretti on Digital Media, a presentation by Kevin Thompson on NSF's IRNC program, and a spreadsheet on IPv6 deployment in the US. He went on to describe the reorganisations that had taken place in NSF.

As a result of the change in his own duties, George Strawn had decided to step down as CCIRN co-chair for North America. The new co-chair would be Heather Boyles. Her nomination had also been agreed by

CANARIE. Grant Miller was stepping down as CCIRN information co-ordinator for North America; Heather Boyles would look for a suitable successor.

As a consequence, Grant Miller, who had produced the minutes of the CCIRN meetings in the past years, would not be able to do so for this meeting and future ones. Kees Neggers suggested that for future meetings, the host organisation should provide the minute taker. Kevin Meynell volunteered to do this for the current meeting.

Xing Li announced that the next APAN Meeting would be held in Taipei on 23-27 August 2005. The APAN Secretariat was hosting and maintaining the CCIRN website. The APAN Secretariat had been moved from Korea to Thailand and appeared to be functioning well.

Florencio Utreras said that he was now working full-time for CLARA as their only employee. CLARA had been legally established in Uruguay with support from LACNIC (which was also established there), but the engineering, operations and database functions were contracted out to organisations in Brazil, Mexico and Chile respectively. The aim was to maintain as small an organisation as possible.

Kees Neggers recalled that there were two European research networking organisations, with DANTE responsible for pan-European production network services (e.g. GÉANT2), and TERENA responsible for network development activities. Recent activities had focused around the GN2 project, which involved DANTE, TERENA and most European NRENs, and encompassed a new pan-European backbone, a number of research activities, and so-called networking activities. One of those was the production of the TERENA Compendium (a survey of NRENs in Europe and the surrounding regions). The 2005 edition of the Compendium was currently being compiled and would be published in September/October, but some results were already available on the TERENA website (see http://www.terena.nl/compendium/).

4. Hybrid IP Optical Networking

4.1. The GÉANT2 Hybrid Network

Vasilis Maglaris gave a presentation about the GN2 project and their plans for a hybrid network to replace the existing GÉANT network. The GN2 project was part-funded by the European Commission and combined several activities necessary to provide research infrastructures and implement integrated services. The project partners were 30 NRENs plus DANTE and TERENA; GN2 had a projected budget of EUR 179 million over four years (of which the EC would provide just over 50%).

The main activity of GN2 was to deploy GÉANT2, which would combine IPv4/6 packet-switching with end-to-end circuit switching. This would provide a production IP service for regular users, with the possibility of Layer 1 and 2 services for high-end research users. A number of technical and policy issues still needed to be resolved, but for the first time, demanding research applications would have access to end-to-end bandwidth at up to 10 Gbps speeds.

GÉANT2 would largely be built from dark fibres, one of the benefits being that there would be only a small incremental cost for additional wavelengths. There would also be 2 x 10 Gbps fibres to North America to provide connectivity to Internet2, ESnet and CANARIE (including SINET); links to Latin America (ALICE), Mediterranean countries (EUMEDCONNECT) and East Asia (TEIN2); as well as to China and South Africa.

Amongst the issues to be considered for the future was whether the provisioning and ownership of cross-border fibre between neighbouring NRENs would necessitate the continuation of a monolithic interconnectivity model. In addition, the establishment of integrated collaborative environments such as VRVS and Access Grid was envisaged. Finally, the trend towards improved connectivity with other regions seemed certain to continue, along with steps to reduce the digital divide.

There was a reminder that the GÉANT2 Launch Event was being held on 14-15 June 2005 in Luxembourg City.

Kees Neggers added that several countries such as the Germany, Poland and the Netherlands were already committed to using hybrid networks, and many others were considering it, as more dark fibre became available. More information could be found in the Compendium.

Don Riley asked about the primary contact point for developing countries as NSF was looking to fund some connectivity in West Africa. Karel Vietsch replied that the remits of TERENA and DANTE did not extend to sub-Saharan Africa. RENATER had good contacts with some West African countries.

4.2. Hybrid IP/Optical Networking in the Asia-Pacific region

Xing Li gave a presentation about hybrid IP/optical networks in the Asia-Pacific region. The problems with BGP peering were that it became complicated when multiple parties were involved, and high-speed point-to-point links were really needed for long-haul services. The challenge was to define what was meant by optical networking as this included customer-owned dark fibre, leased lambdas, or point-to-point Layer 2 services. There were also different transport technologies such as SDH/SONET, Gigabit Ethernet and MPLS.

There were currently four APAN members with optical networks. AARNet had the right to use two fibre pairs on the NextGen network from Perth to Brisbane via Adelaide, Melbourne, Canberra and Sydney. They were currently tendering to utilise DWDM on one pair from Adelaide to Brisbane, as Adelaide to Perth was too expensive due to the need for solar regeneration. They will also link Sydney to Los Angeles through Hawaii via SXTransPORT. The AARNet3 STM-64c trunks will be rolled onto the new DWDM system, with most capacity for dedicated GE circuits. This would provide regional access to the network and private inter-campus links, with any surplus capacity to be utilised for user-controlled light paths.

CERNET had a number of dark fibres that utilised DWDM, but were upgrading to ELH that would support 40 x 10 Gbps. This provided coverage to much of the eastern part of China, as well as to a couple of the western provinces.

In Japan, T-LEX was fully operational and utilised four OC-192s, four OC-12s and 8 GEs. A DS-3 class extension to Bangkok was expected in October 2005, and long-haul OC-192 testing was underway from Tokyo to CERN and SURFnet. In addition, Super SINET was providing 10 Gbps IP connectivity, as well as site-to-site GE connections for advanced scientific research. MPLS VPNs were also available as an easier way of provisioning site-to-site networks.

In South Korea, KREONET2 (funded by the Ministry of Science and Technology) was a lambda-based network connecting 12 GigaPoPs at 5-10 Gbps, whilst SuperSIREN was an optical scientific network operating in Daejeon Science City. There were also connections with Seattle and Hong Kong as part of the GLORIAD network, and possibility of connections to Europe via TEIN2.

At the current time, the possibility of using hybrid IP and optical networks was really only at the experimental stage in the Asia-Pacific region. There was some interest in using them, but there were still many regulatory and economic constraints to overcome first. However, international collaborations were considered very important and should be encouraged in order to push national governments into action.

4.3. The HOPI testbed and plans for intercontinental collaboration

Heather Boyles gave a presentation on the HOPI testbed. Internet2, ESnet and National Lambda Rail were currently in the process of designing a new network architecture which aimed to utilise a combination of shared and dedicated paths. This design had been influenced by the experiences of the DRAGON (NSF-funded), CHEETAH (NSF and DoE funded) and OMNINET (NSF-funded) regional testbeds, as well as the UltraScienceNet (DoE-funded) and HOPI (Internet2-funded) national testbeds. The recommendations could be found in the 'Group A' report produced by NLR and Internet2 (see http://www.internet2.edu/resources/GroupAReport200511.pdf).

HOPI aims to build a complete hybrid network, as individual components were known to work, but have never been put together. The activity is looking at determining when a host should use a circuit-switched or packet-switched infrastructure, dynamic provisioning, duration of paths and scheduling, connectivity between different networks, and network operations and management issues.

Nodes had already been established at Los Angeles Equinix (CalTech and HENP), Washington (MAX/Dragon), and Chicago StarLight, and would shortly be installed at Seattle Pacific Northwest GigaPoP, and New York (NYSERnet). There would also be a circuit between New York and London to provide experimental capabilities between Internet2 and GÉANT. Other collaborations would include GLIF, ESNET, CANARIE and NII/SINET.

A call for proposals to provide the HOPI Testbed Support Centre had been issued, and four excellent responses had been received. The host organisation had not yet been selected, but the support centre would implement control-plane activities, co-ordinate application activities, and engineer and manage the facility.

The plan was for HOPI to offer 1 and 10 GE channels to end-sites, with either a dark fibre or provisioned service to the cross-connect switches. MPLS L2VPNs would also be offered across Abilene, which would allow connections from NLR and the regional optical networks. The testbed would be open to encourage experimentation and the involvement of the corporate community.

More information was available at http://hopi.internet2.edu.

4.4. Hybrid Networking in Latin America

Florencio Utreras gave a presentation on the status of hybrid networking in Latin America. Most operators in the region were still using PDH or SDH/SONET for their production networks, and it was only in Argentina, Brazil, Chile and Mexico where DWDM networks were available. Even then, no lambda services were currently available, and in any case, the regulatory regime in most countries did not allow anyone to operate their own network.

There were, however, three optical testbeds in Brazil (Project GIGA), Chile (G-REUNA) and Mexico (LILA Link). The GIGA network had connected 20 research and development centres in South-Eastern Brazil since May 2004, and offered IP/WDM with ethernet framing. G-REUNA was funded by the MECESup telecom operator and made 4 x 2.5 Gbps connections available between Santiago and Valparaiso for production IP services and optical experiments. Meanwhile, there was a CWDM link between CUDI (the Mexican NREN) in Tijuana, and CalREN (the Californian research and education network) in San Diego.

There were hopes that the current situation could be improved as there was an Inter-American Development Bank initiative to integrate telecommunications across South America. Optical networks could be piggy-backed onto new power lines, gas pipes and roads as these were constructed. In the meantime though, the provision of dark fibre between Argentina and Chile was currently being negotiated by CLARA.

4.5. The Global Lambda Integrated Facility

Kees Neggers gave a presentation about the Global Lambda Integrated Facility. It was no longer sufficient to connect researchers to the Internet as they now needed to be connected to each other. The idea of the GLIF was to build a new Grid computing paradigm where the central architectural element is optical networks that can support e-science applications.

The initial meeting to brainstorm the concept was held during TNC 2001 in Antalya, Turkey. This was followed up with invitational workshops in September 2001 in Amsterdam, in September 2002 in Amsterdam again, and then in August 2003 in Reykjavik. In Reykjavik it was agreed to form the GLIF, with the first meeting being held the following year in Nottingham, UK.

The GLIF is a worldwide collaborative initiative of NRENs and institutes that own or have access to lambdas, and is therefore positioned on the demand side of the community. It is an open initiative which has participants rather than members, with the aim of bringing together network resources to facilitate middleware and application development. Appropriate to its mission, GLIF has a lightweight governance structure with decisions collectively taken by the participants, and secretariat functions being provided by TERENA.

Four working groups had been formed to work on specific issues related to establishing integrated optical networks. The Governance Working Group would identify future goals in terms of lambdas, connections and applications support, and would decide what cross-domain policies needed to be put in place. The Research and Applications Working Group would identify the applications that could benefit from LambdaGrids and define the services required by user communities. The Technical Working Group would design and implement an international LambdaGrid infrastructure, identifying necessary equipment, connection requirements and engineering services. Finally, the Control Plan and Grid Integration Middleware would agree which interfaces and protocols would be used to connect the contributed lambda resources. Each of these working groups had their own mailing list for discussing their activities, and would all meet at the next annual GLIF workshop.

The 5th Annual GLIF Workshop would be held on 29-30 September 2005 in conjunction with iGRID 2005 in San Diego, USA.

More information was available at http://www.glif.is.

5. IPv6

5.1. IPv6 rollout and future planning in the Asia-Pacific region

Xing Li gave a presentation on the current status of IPv6 in the Asia-Pacific region. IPv6 allocations had been rising steadily over the past few years, and the largest consumers of IPv6 addresses were Japan, South Korea, Taiwan and China. Most sites currently appeared to operate on a dual-stack basis.

In Australia, both AARNet and Grangenet were already IPv6 native, and AARNet had peerings with Abilene, CA*net4 and TANET2 at Pacific Wave. They also had commercial peerings at the SIX, PAIX and LAIIX Internet Exchanges. Unfortunately, take-up at customer sites was slow, possibly due to firewall support issues. Many academic users wanted to use IPv6 from home, but were hampered by a lack of ISP support and no native peerings with AARNet. There was simply not much commercial interest in IPv6 at the present time, possibly due to the lack of any killer application.

China had a population of 1.2 billion which meant that if everyone wanted an IP address, it would require 72 Class A IPv4 addresses. If they each spent one hour per day online (at 28 kbps), this would generate 1.68 terabits per second, and if they each had a homepage of 1 MB, this would require 1200 terabytes of storage. It was therefore essential that China migrated to IPv6 if the Internet was to scale in the future, and CERNET2 had therefore been established as a native IPv6 network with dual-stack connections to other networks.

Hong Kong received its first IPv6 block in 2001, and now had four /32 blocks covering both commercial and academic networks. The Hong Kong Internet Exchange (HKIX) supported IPv6, and HARNET (the Hong Kong NREN) also had native IPv6 connections to its member institutions. The HARNET IPv6 testbed website could be found at http://www.i2.jucc.edu.hk/ipv6_testbed/ipv6.cgi.

IPv6 is still considered experimental in Japan, although most of the larger ISPs are interested in providing it. Many are running trials, and more than 50 ISPs are connected to the experimental NSPIXP6 Internet Exchange. WIDE were also looking at the different IPv6 implementations in order to analyse what effect they might have on the network if they were widely used, and there were VoIP and content delivery trials underway. JGN II was an experimental IPv6 backbone that also had connections to the US, and tunnelled connectivity to SINET.

In South Korea, the ANF IPv6 Task Force aimed to deploy a high-speed native IPv6 service (6GN) that would represent 10% of the country's total Internet traffic. The second phase had started in April 2005 and now comprised 11 members connected though 2 backbone networks (KOREN and KREONET2) and the 6NGIX Internet Exchange. A number of IPv6 applications were also being developed including multi-channel HDTV, videoconferencing, high-quality video streaming, and high-speed file transfer.

The Malaysian IPv6 Council was launched in December 2004 under the auspices of the Ministry of Energy, Water and Communications. This recommended that all ISPs must be IPv6 compliant by the end of 2005, all government agencies must be IPv6 compliant by the end of 2008, and the whole of Malaysia must use IPv6 by 2010. In conjunction with this, the National Advanced IPv6 Centre of Excellence was established in March 2005 to promote IPv6 research and development, study the impact of protocol transition, port existing applications to IPv6, and distribute and manage IPv6 addresses. MYREN (the Malaysian NREN) will enable dual-stack on its backbone by June 2005, but will also offer a separate 2 Mbps IPv6-only link to its 12 members. In addition, MY6 is an initiative of ISPs and telcos to establish a joint and open IPv6 network for conducting production level trials.

There were also plans by the Philippine, New Zealand and Pakistan NRENs to rollout IPv6 on their networks in the near future.

In conclusion, there was a great need for IPv6 due to the large population in the Asia-Pacific region. However, IPv6 was still considered experimental and few commercial ISPs were currently interested in using it. It was mostly the academic communities that were leading the efforts, although this may change if a good business model could be found.

5.2. IPv6 and Internet2 Update

Heather Boyles gave a presentation about the Internet2 IPv6 activities. Internet2 currently had 56 organisations and 26 connections using IPv6, as well as 34 peering partners: 3 federal, 1 corporate and 26 international.

The peering methods were via exchange points; direct peering to backbone routers; and peering via GigaPoPs, tunnels and BGP multi-hop. The current exchange points were MANLAN (I2/NYSERnet partnership), PacWave-Seattle, PacWave-LA, Starlight, NGIX East, NGIX West, and Palo Alto PAIX (from January 2005). The IPv6 and IPv6 multicast peering policy was open, with transit offered if desired.

The Abilene backbone had now been enabled for IPv6 unicast using BGP and IS-IS, with 6to4 tunnel relays at Indiana University and PSC. However, while the Abilene and GigaPoPs were IPv6-enabled, a 'last mile' problem existed. This may be because end-site hardware did not support IPv6, administrators were still not convinced of the value of IPv6, there were security concerns, many application protocols did not support IPv6, and there was still a lack of monitoring tools. Some of the security concerns were currently being addressed as the Abilene NOC now limited the prefixes being sent (as with IPv4), and provided some filtering for peer networks. A new IPv6 security mailing list had also been established.

The Abilene backbone was also enabled for IPv6 multicast, but was not currently part of the global M6Bone. NYSERnet, the New York research and education network was also configuring IPv6 multicast on their backbone.

Both IPv4 and IPv6 traffic was being measured on Abilene using BWCTL and OWAMP which showed that performance of the two protocols was virtually indistinguishable. IPv6 throughput lagged IPv4 by about 0.1-0.3%, which was almost negligible.

With respect to applications, North Carolina State University and Centaur Labs were streaming the WCPE and WZYC radio stations, whilst Georgia Tech was running IPv6 H.323 tests. CalREN was also adding IPv6 support to VRVS, and the Internet2 Detective application was being ported to IPv6.

A number of two-day IPv6 tutorials had been organised and the lecture notes and exercises were available at http://ipv6.internet2.edu/workshops/index.shtml.

5.3. The JET survey on IPv6 rollout

Grant Miller presented the results of the JET survey into expected organisational requirements. The Department of Defence was driving IPv6 rollout, but other governmental agencies were expected to start planning and implement testbeds. As a result, most networks had already implemented dual-stack, although there was still little IPv6 traffic on them. By way of anticipating future demand, a survey of different research areas had been undertaken to determine future requirements.

High Energy Physics had an average end-to-end throughput of 0.5 Gbps in 2003, but this was expected to rise to 100 Gbps in the next five years, and to 1000 Gbps within ten years. High bulk throughput was needed, although this was generally not time critical.

Climate Data Computation also had an average end-to-end throughput in 2003, but this was expected to rise to 160-200 Gbps in the next five years, and to 1000 Gbps within ten years. Again, this high bulk throughput was needed rather than time critical access.

Nano-science had not started in 2003, but expected to need 1 Gbps in the next five years, and 1000 Gbps with ten years. This involved remote control, so time-critical throughput (QoS) was required.

Fusion energy was using 60-500 Mbps in 2003, but was expected to need 190-500 Mbps in the next five years, and 1000 Gbps within ten years. Data tended to be bursty and time-critical.

Astrophysics typically used 1 TB per week end-to-end in 2003, but was looking to use multicast as much of its data was asymmetrical. Nevertheless, it expected to need 1000 Gbps within ten years.

Genomics typically used 1 TB per day in 2003, but expected hundreds more users over the next five years who would require high throughput and QoS.

Kees Neggers asked whether the US government had a mandatory IPv6 strategy. George Strawn replied this was not yet the case, but it was coming.

5.4. IPv6 rollout in GÉANT2 and Europe

Vasilis Maglaris gave a presentation about IPv6 rollout on GÉANT and elsewhere. The first pilot connections had been implemented by RedIRIS (Spain) and RENATER (France) in April 2003, followed by connections to Abilene, SINET and CANARIE, as well as the commercial networks of Telia and Global Crossing. A production unicast service was initiated on the GÉANT backbone from October 2003, followed by IPv6 multicast from January 2005.

There were currently 25 NRENs natively connected via IPv6, with a further 4 connected via tunnels. There were also native connections to Abilene, CANARIE and ESnet; and tunnelled connections to SINET, APAN-KR (via RENATER), Telia and Global Crossing. The Abilene connection also provided onward transit to a number of other networks in the US, Asia-Pacific and South American regions. The GÉANT policy was to accept prefixes in the range /19 to /35 (with exceptions for root name servers), although 6Bone prefixes (3ffe::/16) were not accepted.

Native IPv6 multicast had been tested by the Multicast v6 Task Force and had been implemented as a pilot service on GÉANT since early-2005. 7 NRENs were connected natively, with a further 4 connected via tunnels. Connectively to the M6Bone was provided via RENATER.

Since the introduction of IPv6, traffic to and from GÉANT had increased from 13 terabits per month in May 2005, to 43 terabits per month in April 2005. The biggest users were NORDUnet (Scandinavia),

SWITCH (Switzerland), SURFnet (the Netherlands) and PSNC (Poland) respectively. At the present time, IPv6 accounted for 2% of the total traffic over GÉANT.

IPv6 was also being implemented on the networks associated with GÉANT. SEEREN had already implemented 6PE, Silk (Central Asia) had deployed IPv6 over DVB-S, and there were plans to run IPv6 on ALICE and EUMEDCONNECT.

There had been several EU-funded projects to promote, develop and test IPv6 technologies. 6NET (http://www.6net.org) involved a number of NRENs, academic institutes and commercial companies, and looked at all the aspects necessary for running a production IPv6 network. Euro6IX (http://www.euro6ix.net) involved a number of telecommunication operators and established a number of IPv6 Internet Exchanges. 6DISS (http://www.6diss.org) was a follow-up project that aimed to disseminate knowledge gained through these projects to the South-Eastern European, Mediterranean, Central Asia, Asia-Pacific, Sub-Saharan Africa, Southern Africa and Caribbean regions.

5.5. Deployment of IPv6 in Latin America

Florencio Utreras gave a presentation on the status of IPv6 deployment in Latin America. As of November 2004, there were 57 nodes in the context of the 6Bone project, which represented 3.83% of the total hosts worldwide. There were also 16 IPv6 networks, mostly in Mexico, but also in Argentina/ Chile, Brazil/Dominican Republic, and Uruguay.

RedCLARA (the South American backbone network) aimed to be dual-stack by October 2005. A /32 address block had been requested from LACNIC and deployment would start in July. CUDI (Mexico), REUNA (Chile), RNP (Brazil) and RETINA (Argentina) would start peering immediately.

CLARA had established an IPv6 Working Group which was producing technical recommendations on the use of IPv6. It would also coordinate the activities of the CLARA Engineering Group and CLARA NOC, and would liaise with other relevant IPv6 groups around the world.

In conjunction with this, the Latin American IPv6 Task Force was helping to promote IPv6 deployment in the region, and facilitating exchange of information and experiences. This was sponsored by LACNIC who were also offering free IPv6 address allocations for the first two years. Regular meetings had created a community of users, and the next meeting would be held on 29 June 2005 in Lima, Peru.

Heather Boyles asked whether anyone had undertaken a survey to determine which IPv6 components were missing or not of production quality. Kevin Meynell replied that 6NET had done this and would shortly be publishing a report (D2.5.3). He said he would notify the CCIRN list when this became available.

ACTION 20050604-1: Kevin Meynell to notify CCIRN mailing list when 6NET report on IPv6 deployment issues was published.

6. Intercontinental collaboration on PR and information dissemination

Karel Vietsch discussed improving collaboration on PR and information dissemination issues between the continental organisations. He said that European researchers and European research networkers were traditionally quite poor at public relations, and this had often resulted in the rather embarrassing situation where the European Commission, national governments and large parts of the European research networking community first heard about developments in Europe via North American sources.

TERENA had now established TF-PR, which was a task force for the PR staff of its member organisations. TF-PR aimed to improve communication between TERENA member organisations, as well as external bodies. In particular, it had established the PeaR electronic news agency where PR officers could upload news items of wider interest, and from which selected items could be displayed on the TERENA website (see http://www.terena.nl/news/pr/agency/).

Karel Vietsch asked how other organisations undertook their information dissemination, and whether it was possible to initiate a common exchange of news between continental organisations.

George McLaughlin said that AARNet employed mailing lists, a professional PR company, and had established contacts with specific publications.

Heather Boyles said that Internet2 generally used their mailing lists and website, but Bill St Arnaud still usually beat them to issuing news about their organisation. He had been running his well-known mailing list for many years and people tended to send information to him first.

Florencio Utreras said that CLARA had been using a news agency, but were trying to build up their website. The plan was to make information available in English, Spanish and Portugese.

Karel Vietsch added that CCIRN website needed to be updated and asked whether the APAN Secretariat could devote some time to this task. Regional representatives were asked to send updated links and information to the APAN Secretariat.

ACTION 20050604-2: Regional representatives to send updated links and information to the APAN Secretariat.

Karel Vietsch would informally discuss the organisation of PR and information dissemination with CLARA, APAN, Internet2 and CANARIE representatives, and prepare a further discussion about the topic in next year's CCIRN meeting.

7. The future of the CCIRN

A discussion paper about the future of the CCIRN had been circulated prior to the meeting (see TSec(05)051). The original activities of the CCIRN had largely been taken on by bi-lateral relationships of continental organisations, so the CCIRN had become more of an information sharing forum in recent years. Whilst this was still considered useful, the CCIRN meetings only had a small audience and the information sharing may be done more effectively at the regional conferences (e.g. TNC, APAN meetings). The members therefore needed to consider re-defining the role of the group, or whether to continue with it in the longer-term. It was proposed to position the CCIRN as the meeting forum of the continental networking organisations (TERENA, APAN, CLARA, Internet2 + CANARIE).

It was suggested that the CCIRN could also be used as a political lobbying body for the continental organisations, as the Internet Society now had wider goals beyond research networking. In conjunction with this, the CCIRN website could be developed into a portal for worldwide research networking, with links to the development and deployment activities that are ongoing in the continental organisations. This could be supplemented with information along the lines of the maps that Larry Landweber used to publish, showing the research network resources in each country.

Such activities would require some secretarial support, but such support functions could be provided by the staff of the continental organisations on a shared basis. The APAN Secretariat had already taken responsibility for managing the CCIRN website, so that was a step in the right direction.

Heather Boyles raised the question which role the CCIRN could play in addressing the problem of the digital divide world-wide. This was in particular a problem for those parts of the world (sub-Saharan Africa, Central Asia) that were outside the geographic scope of the existing continental research networking organisations.

Klaus Ullmann saw the CCIRN as a moderating committee for a number of activities in the areas of innovation, combating the digital divide etc. He and Dorte Olsen felt that a permanent subcommittee, such as the existing group of continental co-chairs and information co-ordinators, should take the responsibility for setting the agenda for topics to be discussed by the CCIRN.

With the points about the digital divide and the role of co-chairs and co-secretaries added, there was general support for the discussion note. It was agreed that the Terms of Reference of the CCIRN should be re-written on the basis of the discussion in this meeting.

ACTION 20050604-3: Karel Vietsch to draft revised Terms of Reference for the CCIRN, for adoption in the next CCIRN meeting.

8. Date and location of next meeting

As previously agreed, it was the turn of North America to host the next meeting, although some consideration should be given to adding Latin America to the rota. It was proposed that the next CCIRN meeting be held near the Internet2 Members Meeting on 24-26 April 2006 in Arlington, USA. Heather Boyles would make the meeting arrangements and announce the precise dates and location well in advance.

ACTION 20050604-4: Heather Boyles to announce precise dates and location of the next CCIRN meeting before January 2006.

9. Any Other Business and Close

Kees Neggers thanked George Strawn and Grant Miller for all the work they had done for the CCIRN in the past years. He thanked PSNC for their hospitality in hosting this meeting.

Summary of Actions

ACTION 20050604-1	Kevin Meynell to notify CCIRN mailing list when 6NET report on IPv6 deployment issues was published.	New
ACTION 20050604-2	Regional representatives to send updated links and information to the APAN Secretariat.	New
ACTION 20050604-3	Karel Vietsch to draft revised Terms of Reference for the CCIRN, for adoption in the next CCIRN meeting.	New
ACTION 20050604-4	Heather Boyles to announce precise dates and location of the next CCIRN meeting before January 2006	New