AARNet3 Design

During 2003 the RFP team worked on the design for the new national network. Unsurprising for a network now more than 6 years old they identified a number of critical issues that needed to be addressed in the design of any new replacement network. These issues included:

- Redundancy & Resilience
- Support for IPv4 and IPv6, unicast and multicast
- Traffic Accounting and Monitoring
- End to end performance measures
- Support QoS (diffserv)
- Support for large traffic flows

There are many single points of failure in the existing network and so to address some of them the team determined that the new network should be built around dual points of presence (POP) in the major cities. The POP sites within a city will be connected to each other and to other capital cities in a redundant fashion, creating a bracelet of rings across the country. These POP sites have been chosen so that at least one of the sites is in an inner suburb rather than have both in the CBD in the hope that this will assist clients in obtaining truly diverse fibre paths to each site and so gain additional resilience in their own connection to the network when they connect to each POP in the city.

While there has not been a large demand up until now within the wider client community for support of IPv6 the team saw this as an area of growth during the lifetime of the new network, especially given its profile within north Asia. Therefore IPv6 must be afforded the same priority within the new network as IPv4. A network that treated IPv6 as a second-class citizen was not going to be acceptable and so the type of traffic should not influence performance of the network. The team also saw growth in the use of multicast technologies, especially within the Access Grid community, and so the network had to support IPv4 multicast and the emerging IPv6 multicast standards.

While Netflow traffic analysis is largely seen as a billing mechanism in AARNet it is also used by the AARNet NOC for Denial of Service (DoS) mitigation and traffic analysis in general. Under any new network the team wanted to retain the ability to monitor traffic flows and so it needed to scale to the number and size of flows possible within a Gigabit network. It was clear that support for non sampled Netflow wasn't going to be possible in the class of equipment necessary to drive the new backbone and so rather than try to work within a sampled Netflow framework the decision was made to move the collection of data closer to the client. This largely drove the decision to provision an edge router to members, much like in the original AARNet model from 1990. This allows APL to continue to provide detailed Netflow based traffic analysis, providing not only a traffic breakdown for billing but also assisting in the identification of anomalies in the traffic flow that could indicate a DoS attack in progress.

An additional benefit of AARNet providing an edge device connecting directly to the client's network is that it provides APL with the opportunity to monitor end-to-end performance from the client's perspective.

Use of voice and video over AARNet is becoming more common and as these technologies are quite sensitive to changes to jitter APL not only wants to measure these effects but also try to ensure that they are minimised. This will result in APL deploying quite a sophisticated quality of service mechanism on the new network so that client's can signal to the network which traffic needs preferential treatment. This should ensure that the quality of voice and video calls doesn't degrade even in the face of a significant denial of service attack on the network.

Finally Internet2 has demonstrated that large bandwidth links, such as those to be deployed for AARNet3, will be under utilised unless the network is tuned to make better use of the links. One of the mechanisms that can be used is "jumbo frames". This allows the transmission of larger packets across the network and thus makes better use of the available capacity. This is only effective if it is available end to end and so clients will be encouraged to reconfigure their networks to use these "jumbo frames" where possible.

While the initial focus for building out the new AARNet3 network will be the inter-state trunk network APL will also be moving forward in utilising the Nextgen Network to supplement it with a new high speed regional network offering high speed connectivity to regional sites previously only experienced by metropolitan clients. The actual technology to be deployed on the regional component of AARNet3 will be determined in consultation with the clients in the regional areas to be serviced by the new network and will be the focus of a further RFQ in early 2004.

Coupled with a new national and regional network based on Nextgen Networks services APL will also be deploying multiple international connections to the United States, initially, to service our commodity and R&E connectivity requirements. The commodity services, to Palo Alto and Los Angeles, will allow AARNet to benefit from the more competitive market in the United States for Internet services. In each of these locations APL plans to take a commodity Internet transit service from two "Tier One" providers. Again this is in support of our diversity and redundancy requirements. Any single transit provider link should be able to accommodate the full AARNet traffic load but by utilising two providers at each POP site the network will be resilient to the loss of service from any one provider and/or POP site. APL also plans to "peer" with other providers, both content and service, at these locations to improve the performance, and thus user experience, in accessing services hosted by these providers.